User Guide for Cisco Security MARS Global Controller

Release 5.2.x
May 2007

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Customer Order Number:
Text Part Number: OL-13542-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

Modifying the equipment without Cisco's written authorization may result in the equipment no longer complying with FCC requirements for Class A or Class B digital devices. In that event, your right to use the equipment may be limited by FCC regulations, and you may be required to correct any interference to radio or television communications at your own expense.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

• Turn the television or radio antenna until the interference stops.

• Move the equipment to one side or the other of the television or radio.

• Move the equipment farther away from the television or radio.

• Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, *Packet*, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)

User Guide for Cisco Security MARS Global Controller
Copyright © 2006 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface xi

Introduction xi

Global Controller Overview xi

The Global Controller User Interface xi

About This Manual xii

Obtaining Documentation xii

Cisco.com xii

Product Documentation DVD xiii

Ordering Documentation xiii

Documentation Feedback xiii

Cisco Product Security Overview xiii

Reporting Security Problems in Cisco Products xiv

Product Alerts and Field Notices xiv

Obtaining Technical Assistance xv

Cisco Support Website xv

Submitting a Service Request xvi

Definitions of Service Request Severity xvi

Obtaining Additional Publications and Information xvi

CHAPTER 1 **Introduction** 1-1

Advantages 1-2

Basic Functions of the Global Controller 1-2

Incidents 1-2

Rules 1-3

Centralized Maintenance 1-3

Deployment 1-3

Incremental Deployment 1-3

Green-field, Multi-box Deployment 1-3

CHAPTER 2 **Configuring the Global Controller** 2-1

Adding Local Controllers 2-2

Topology Synchronization 2-3

Deleting Local Controllers 2-5
Importing the Security Certificates 2-6
Monitoring Local Controller Events from the Global Controller 2-10
Preparing to Add and Discover Devices 2-10
Adding Reporting Devices 2-11
Manual Configuration 2-11
Add a Device Manually 2-11
Configuring Supported Devices 2-12
L2 Discovery and Mitigation 2-12

CHAPTER 3 **Network Summary** 3-1
Global Controller Network Summary Page Concepts 3-1
Global Controller Technologies 3-1
Navigation within the MARS Appliance 3-1
Logging In 3-2
Basic Navigation 3-2
Help Page 3-4
Your Suggestions Welcomed 3-5
Summary Page 3-7
Dashboard 3-7
Recent Incidents 3-9
Sessions and Events 3-9
Data Reduction 3-10
Page Refresh 3-10
Diagrams 3-10
Manipulating the Diagrams 3-12
Display Devices in Topology 3-13
Network Status 3-13
Reading Charts 3-14
Hotspots 3-16
My Reports 3-16
To set up reports for viewing 3-16

CHAPTER 4 **Case Management** 4-1
Case Management Overview 4-1
Case Management Considerations for the Global Controller 4-3
Hide and Display the Case Bar 4-3
Create a New Case 4-4

Edit and Change the Current Case 4-5
Add Data to a Case 4-6
Generate and Email a Case Report 4-7

CHAPTER 5 **Incident Investigation and Mitigation** 5-1
Incidents Overview 5-1
The Incidents Page 5-2
Time ranges for Incidents 5-4
Incident Details Page 5-4
To Search for a Session ID or Incident ID 5-4
Incident Details Table 5-5
False Positive Confirmation 5-6
The False Positive Page 5-8
Virtual Private Network Considerations 5-8

CHAPTER 6 **Queries and Reports** 6-1
Queries 6-1
To Run a Quick Query 6-2
To Run a Free-form Query 6-2
To Run a Batch Query 6-3
To Stop a Batch Query 6-5
To Resubmit a Batch Query 6-5
To Delete a Batch Query 6-5
Selecting the Query Type 6-5
Result Format 6-6
Order/Rank By 6-8
Filter By Time 6-8
Use Only Firing Events 6-9
Maximum Number of Rows Returned 6-9
Selecting Query Criteria 6-9
To Select a Criterion 6-9
Query Criteria 6-11
Source IP 6-11
Destination IP 6-11
Service 6-12
Event Types 6-12
Device 6-12
Severity/Zone 6-12
Operation 6-13

Rule 6-13
Action 6-13
Saving the Query 6-13
Perform a Long-Duration Query Using a Report 6-14
View a Query Result in the Report Tab 6-16
Perform a Batch Query 6-16
Reports 6-19
Report Type Views: Total vs. Peak vs. Recent 6-20
Creating a Report 6-21
Create a New Report 6-22
Working With Existing Reports 6-23

CHAPTER 7 **Rules** 7-1
Rules Overview 7-1
Prioritizing and Identifying 7-1
Think Like a Black Hat 7-2
Planning an Attack 7-2
Back to Being the Admin 7-2
Types of Rules 7-3
Inspection Rules 7-3
Global User Inspection Rules 7-4
Drop Rules 7-4
Constructing a Rule 7-4
Working Examples 7-15
Example A: Excessive Denies to a Particular Port on the Same Host 7-15
Example B: Same Source Causing Excessive Denies on a Particular Port 7-15
Example C: Same Host, Same Destination, Same Port Denied 7-15
Working with System and User Inspection Rules 7-16
Change Rule Status—Active and Inactive 7-16
Duplicate a Rule 7-16
Edit a Rule 7-17
Add an Inspection Rule 7-18
Setting Alerts 7-20
Configure an Alert for an Existing Rule 7-20
Rule and Report Groups 7-21
Rule and Report Group Overview 7-22
Global Controller and Local Controller Restrictions for Rule and Report Groups 7-23
Add, Modify, and Delete a Rule Group 7-23
Add, Modify, and Delete a Report Group 7-26

Display Incidents Related to a Rule Group 7-28
Create Query Criteria with Report Groups 7-29
Using Rule Groups in Query Criteria 7-30

CHAPTER 8 **Sending Alerts and Incident Notifications** 8-1
Configure the E-mail Server Settings 8-4
Configure a Rule to Send an Alert Action 8-5
Create a New User—Role, Identity, Password, and Notification Information 8-10
Create a Custom User Group 8-12
Add a User to a Custom User Group 8-13

CHAPTER 9 **Management Tab Overview** 9-1
Activating 9-1
To activate a set of management additions or changes 9-1
Event Management 9-1
Search for an Event Description or CVE Names 9-2
To view a list of all currently supported CVEs 9-2
Event Groups 9-2
To filter by event groups or severity 9-2
Edit a Group of Events 9-2
Add a Group 9-3
IP Management 9-3
Search for an Address, Network, Variable, or Host 9-3
Filter by Groups 9-3
Edit a Group 9-4
Add a Group 9-4
Add a Network, IP Range, or Variable 9-4
Service Management 9-5
Search for a Service 9-5
Add a Group of Services 9-5
Edit a Group of Services 9-5
Add a Service 9-6
Edit a Service 9-6
Delete a Service 9-6
User Management 9-6
Add a New User 9-7
Add a Service Provider (Cell phone/Pager) 9-9
Search for a User 9-9

Edit or Remove a User 9-10
Create a User Group 9-10
Add or Remove a User from a User Group 9-10
Filter by Groups 9-11
Promoting Global User Roles on Local Controller 9-11

CHAPTER 10 **System Maintenance** 10-1

Setting Runtime Logging Levels 10-1
Viewing the MARS Backend Log Files 10-2
View the Backend Log 10-2
Viewing the Audit Trail 10-3
View an Audit Trail 10-3
Change the Default Password of the Administrator Account 10-3
Understanding Certificate and Fingerprint Validation and Management 10-4
Setting the Global Certificate and Fingerprint Response 10-5
Upgrading from an Expired Certificate or Fingerprint 10-6
Upgrade a Certificate or Fingerprint Interactively 10-6
Upgrade a Certificate Manually 10-6
Upgrade a Fingerprint Manually 10-6
Monitoring Certificate Status and Changes 10-7
Hardware Maintenance Tasks—MARS 110R, 110, 210, GC2R, and GC2 10-7
Field Replaceable Units 10-7
Removing and Replacing the Front Bezel 10-8
Removing the Chassis Cover 10-9
Replacing the RAID Battery Backup Unit 10-10
Procedure to Replace the Raid Battery Backup Unit 10-11
Hard Drive Troubleshooting and Replacement 10-14
Hard Drive Status LEDs 10-14
Partition Checking 10-14
Overview of RAID 10 Subsystem 10-15
Hotswapping Hard Drives 10-15
Viewing RAID Array Status with the raidstatus CLI Command 10-17
Hard Drive Slot Number Diagram 10-20
Procedure to Hotswap a Hard Drive 10-21
Hotswap CLI Example 10-22
Replacing a Hard Drive in the Hard Drive Carrier 10-24
Hot-swapping a Power Supply Unit 10-25
Diagnostic Beep Codes 10-26
Safety Information 10-27

Intended Application Uses 10-27
Equipment Handling Practices 10-27
Power and Electrical Warnings 10-27
Power Cord Warnings 10-28
System Access Warnings 10-28
Rack Mount Warnings 10-29
Electrostatic Discharge (ESD) 10-29
Battery Replacement 10-29
Cooling and Airflow 10-30
Laser Peripherals or Devices 10-30

APPENDIX A **Cisco Security MARS XML API Reference** A-1
XML Schema Overview A-1
XML Incident Notification Data File and Schema A-1
XML Incident Notification Data File Sample Output A-2
XML Incident Notification Schema A-4
Usage Guidelines and Conventions for XML Incident Notification A-4

APPENDIX B **Regular Expression Reference** B-1
PCRE Regular Expression Details B-1
Backslash B-2
Non-printing Characters B-3
Generic Character Types B-4
Unicode Character Properties B-5
Simple Assertions B-6
Circumflex and Dollar B-7
Full Stop (Period, Dot) B-8
Matching a Single Byte B-8
Square Brackets and Character Classes B-8
Posix Character Classes B-9
Vertical Bar B-10
Internal Option Setting B-10
Subpatterns B-11
Named Subpatterns B-12
Repetition B-12
Atomic Grouping and Possessive Quantifiers B-14
Back References B-15

Assertions B-16
Lookahead Assertions B-17
Lookbehind Assertions B-17
Using Multiple Assertions B-18
Conditional Subpatterns B-19
Comments B-20
Recursive Patterns B-20
Subpatterns as Subroutines B-21
Callouts B-22

APPENDIX C **Date/Time Format Specfication** C-1

APPENDIX D **System Rules and Reports** D-1
List of System Rules D-1
List of System Reports D-13

GLOSSARY

INDEX

Preface

Introduction

Global Controller helps network administrators, network operators, and security analysts be more productive by:

- reducing the amount of raw data you have to wade through
- enabling you to see your network security posture as it evolves
- identifying Hot Spots of malicious activity
- removing unwanted traffic from your network.

Global Controller Overview

The Global Controller is a Security Threat Mitigation (STM) system. It summarizes information about the health of your network as viewed through the reporting devices.

The Global Controller:

- collects all raw events,
- sessionizes them across different devices,
- fires default rules for incidents,
- determines false positives, and
- delivers consolidated information through diagrams, charts, queries, reports, and rules.

The Global Controller User Interface

The Global Controller system employs a Global Controller to monitor and manage Local Controllers and their monitored devices in the network. The Global Controller user interface uses a tabbed, hyperlinked, browser-based interface. If you have used the Web, you have used similar pages.

Note When using the Global Controller user interface, avoid using the browser's **Back** and **Forward** buttons. Using these buttons can lead to unpredictable behavior.

About This Manual

This manual describes the features and functionality of the Global Controller.

The layout of this manual is as follows:

- Chapter 1, "Introduction" — This chapter introduces the Global Controller and presents its basic features and deployment options.
- Chapter 2, "Configuring the Global Controller" — This chapter covers connecting to the Global Controller for the first time, and setting up and configuring your network security devices to connect to the Global Controller.

Part II: Monitoring Phase. This part concepts important to successfully using MARS to monitor your network. These concepts include defining inspection rules and investigating incidents.

- Chapter 3, "Network Summary" covers the Summary pages which includes the Dashboard, the Network Status, and the My Reports pages.
- Chapter 4, "Case Management" covers using cases to provide accountability and improve workflow.
- Chapter 5, "Incident Investigation and Mitigation" covers incidents and false positives and provides a starting point for configuring a Layer 2 path and mitigation to work with a MARS.
- Chapter 6, "Queries and Reports" covers working with scheduled and on-demand reports and queries. It also discussing using the real-time event viewer.
- Chapter 7, "Rules" covers defining and use inspection rules.
- Chapter 8, "Sending Alerts and Incident Notifications" explains how to configure the MARS to send an alert based on an inspection rule.
- Chapter 9, "Management Tab Overview" covers managing events, networks, variables, hosts, services, and MARS users.
- Chapter 10, "System Maintenance" covers some of the maintenance chores for the MARS.

Additionally, the following appendices are provided:

- Appendix A, "Cisco Security MARS XML API Reference" presents the XML schema used by MARS for XML-based notifications.
- Appendix B, "Regular Expression Reference" The syntax and semantics of the regular expressions supported by PCRE are described in this appendix.
- Appendix C, "Date/Time Format Specfication" The date/time field parsing is supported using the Unix strptime() standard C library function.
- Glossary — A glossary of terms as they relate to MARS.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. This section explains the product documentation resources that Cisco offers.

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

The Product Documentation DVD is a library of technical product documentation on a portable medium. The DVD enables you to access installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the HTML documentation and some of the PDF files found on the Cisco website at this URL:

http://www.cisco.com/univercd/home/home.htm

The Product Documentation DVD is created and released regularly. DVDs are available singly or by subscription. Registered Cisco.com users can order a Product Documentation DVD (product number DOC-DOCDVD= or DOC-DOCDVD=SUB) from Cisco Marketplace at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

Ordering Documentation

You must be a registered Cisco.com user to access Cisco Marketplace. Registered users may order Cisco documentation at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

If you do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Documentation Feedback

You can provide feedback about Cisco technical documentation on the Cisco Support site area by entering your comments in the feedback form available in every online document.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you will find information about how to do the following:

- Report security vulnerabilities in Cisco products
- Obtain assistance with security incidents that involve Cisco products
- Register to receive security information from Cisco

A current list of security advisories, security notices, and security responses for Cisco products is available at this URL:

http://www.cisco.com/go/psirt

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you have identified a vulnerability in a Cisco product, contact PSIRT:

- For emergencies only—security-alert@cisco.com

  An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- For nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Tip We encourage you to use Pretty Good Privacy (PGP) or a compatible product (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.*x* through 9.*x*.

Never use a revoked encryption key or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

If you do not have or use PGP, contact PSIRT to find other means of encrypting the data before sending any sensitive material.

Product Alerts and Field Notices

Modifications to or updates about Cisco products are announced in Cisco Product Alerts and Cisco Field Notices. You can receive these announcements by using the Product Alert Tool on Cisco.com. This tool enables you to create a profile and choose those products for which you want to receive information.

To access the Product Alert Tool, you must be a registered Cisco.com user. Registered users can access the tool at this URL:

http://tools.cisco.com/Support/PAT/do/ViewMyProfiles.do?local=en

To register as a Cisco.com user, go to this URL:

http://tools.cisco.com/RPF/register/register.do

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Support website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Support Website

The Cisco Support website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day at this URL:

http://www.cisco.com/en/US/support/index.html

Access to all tools on the Cisco Support website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Note Before you submit a request for service online or by phone, use the **Cisco Product Identification Tool** to locate your product serial number. You can access this tool from the Cisco Support website by clicking the **Get Tools & Resources** link, clicking the **All Tools (A-Z)** tab, and then choosing **Cisco Product Identification Tool** from the alphabetical list. This tool offers three search options: by product ID or model name; by tree view; or, for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Tip **Displaying and Searching on Cisco.com**

If you suspect that the browser is not refreshing a web page, force the browser to update the web page by holding down the Ctrl key while pressing **F5**.

To find technical information, narrow your search to look in technical documentation, not the entire Cisco.com website. After using the Search box on the Cisco.com home page, click the **Advanced Search** link next to the Search box on the resulting page and then click the **Technical Support & Documentation** radio button.

To provide feedback about the Cisco.com website or a particular technical document, click **Contacts & Feedback** at the top of any Cisco.com web page.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests, or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411
Australia: 1 800 805 227
EMEA: +32 2 704 55 55
USA: 1 800 553 2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—An existing network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of the network is impaired while most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Online Subscription Center is the website where you can sign up for a variety of Cisco e-mail newsletters and other communications. Create a profile and then select the subscriptions that you would like to receive. To visit the Cisco Online Subscription Center, go to this URL:

  http://www.cisco.com/offer/subscribe

- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco channel product offerings. To order and find out more about the *Cisco Product Quick Reference Guide*, go to this URL:

  http://www.cisco.com/go/guide

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- Cisco Press publishes a wide range of general networking, training, and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Internet Protocol Journal* is s a quarterly journal published by Cisco for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the *Internet Protocol Journal* at this URL:

  http://www.cisco.com/ipj

- Networking products offered by Cisco, as well as customer support services, can be obtained at this URL:

  http://www.cisco.com/en/US/products/index.html

- Networking Professionals Connection is an interactive website where networking professionals share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

  http://www.cisco.com/discuss/networking

- “What’s New in Cisco Documentation” is an online publication that provides information about the latest documentation releases for Cisco products. Updated monthly, this online publication is organized by product category to direct you quickly to the documentation for your products. You can view the latest release of “What’s New in Cisco Documentation” at this URL:

  http://www.cisco.com/univercd/cc/td/doc/abtunicd/136957.htm

- World-class networking training is available from Cisco. You can view current offerings at this URL:

  http://www.cisco.com/en/US/learning/index.html

CHAPTER 1

Introduction

The MARS Global Controller is a security threat mitigation (STM) appliance. Once you deploy multiple Local Controllers, you can deploy a Global Controller that summarizes the findings of two or more Local Controllers. In this way, the Global Controller enables you to scale your network monitoring without increasing the management burden. The Global Controller provides a single user interface for defining new device types, inspection rules, and queries, and it enables you to manage Local Controllers under its control. This management includes defining administrative accounts and performing remote, distributed upgrades of the Local Controllers. The Global Controller is available in two models—MARS GCm and MARS GC.

A Global Controller monitors two or more local zones. Each zone consists of a cluster of monitored devices and is managed by a Local Controller. The following diagram shows the relationship between the Global Controller and multiple local zones.

Figure 1-1 Relationship of Global Controller to Local Controller to Reporting/Mitigation Device

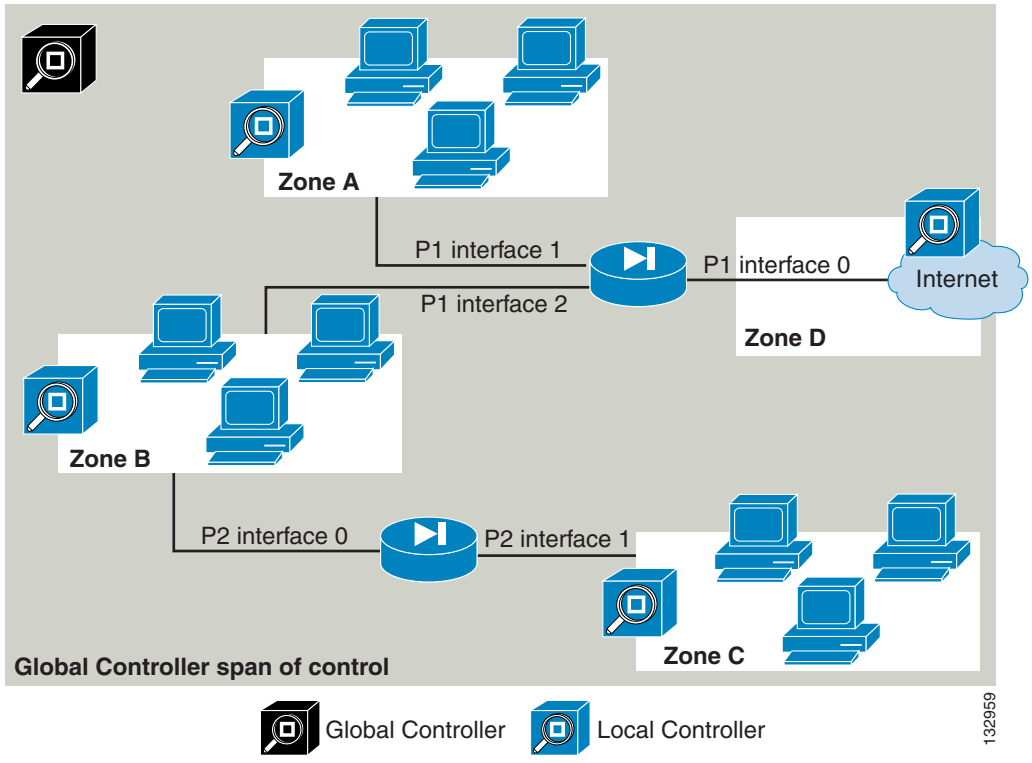

For more information about the architecture of a distributed MARS system, refer to the System Description, page 1-1 in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.

Advantages

The Global Controller/Local Controller architecture has the following advantages:

- It allows for centralized, distributed management of network topology.
- It lets remote sites view their own data while keeping data private between Global Controller and Local Controllers.
- You can view the entire network from the Global Controller.
- It enables linear scalability usnig a multi-layer hierarchy.
- You can use multiple Local Controllers to isolate departmental functions such as, host logging, NIDS, compliance, and for network profiling and anomaly detection.
- It preserves the WAN link by pushing up correlated information instead of raw data from monitoring device.

Basic Functions of the Global Controller

The Global Controller centrally manages a group of Local Controllers. Its user interface displays a listing of all the zones with their respective Local Controllers.

The Global Controller monitors and manages the network with a powerful suite of functions:

- Incidents
- Rules
- Queries and reports
- Centralized maintenance (for example, software upgrades of managed Local Controllers)

A Global Controller Admin user has the ability to create, edit or delete information on the Global Controller and its monitored Local Controllers. Information such as:

- Rules
- Reports and queries
- User, IP and service management
- Management grouping (for example, event and user groupings)

Incidents

The Global Controller can monitor any Local Controller at any time to receive data. It receives summarized information from all its Local Controllers and produces a merged summary of this data. The summary consists of global topologies and incidents reflecting network activities in each of its zones. The topologies and incidents can be drilled down to their subsets of paths and events at the zonal level.

The summaries provides an account of high-, medium-, and low-priority incidents. All network, port, protocol, applications, and events have to be global in scope to be on the Global Controller.

Rules

The Global Controller uses rules to monitor the zones that report to it. Rules that apply to multiple Local Controllers can be created on the Global Controller and pushed down to them from a central location. These rules trigger incidents that you can review at the global level.

Note Rules created on the Local Controller remain local. Incidents generated from these rules do not get push up to the Global Controller.

Centralized Maintenance

The Global Controller leaves most data archiving to the Local Controller. However, some basic archive/restore capability is provided at the global level.

The Global Controller centrally manages all upgrades to the Local Controllers. Global Controller manages Local Controller(s) that is running the same version of the software as it is.

Deployment

The Global Controller system's flexible architecture supports two types of deployment:

- Incremental Deployment, page 1-3
- Green-field, Multi-box Deployment, page 1-3

Incremental Deployment

In this scenario, a administrator deploys one or more Local Controller systems as standalone units. At a later date, the administrator decides to add a Global Controller to the scenario. The previously deployed Local Controllers must be upgraded to communicate with the new Global Controller.

To enable this communication, you must:

1. Create a zone for each Local Controller
2. Ensure that the reporting devices do not overlap among zones
3. Upgrade the Local Controller version to be the same as that of the Global Controller
4. Add the Local Controller as a monitored controller in the Global Controller.
5. The Global Controller is then configured to communicate with each Local Controller by exchanging security certificate information.

Once this communication is enabled, the Global Controller is able to receive information, such as incidents and rules, from the Local Controller.

Green-field, Multi-box Deployment

In this scenario, the network administrator decides from the very start to deploy two or more Local Controllers and a Global Controller to monitor them. In this case, the administrator must define the zones and their monitored devices ahead of time to complete a smooth installation.

CHAPTER 2

Configuring the Global Controller

Once you have performed the configuration tasks described in this chapter, a Global Controller administrator can create, edit, or delete user-defined settings and rules on the Global Controller and its monitored Local Controllers. These settings and rules include:

- Rules
- Reports and queries
- User, IP, and service management

To configure the Global Controller, you must perform several tasks before you can monitor the events and incidents reported by Local Controllers:

1. Configure the Global Controller to operate on your network. For more information on installing and configuring the Global Controller to connect to your network, see the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*.
2. Divide your network topology into locally controlled zones. For each zone identified, install and configure a Local Controller.
3. Add the reporting and mitigation devices in a zone to the Local Controller that monitors that zone. Also, configure the SNMP read-only community string settings for those devices to enable network discovery.
4. Add the zones to be monitored into Global Controller. Each zone is represented by a single Local Controller. By adding a Local Controller to the Global Controller, you are indicating that the Global Controller should monitor that local zone.

Note You can only add reporting devices to an active Local Controller.

5. Import the security certificate from each Local Controller into the Global Controller and vice versa. Sharing the security certificates among the appliances enables secure communications between a Local Controller and the Global Controller.
6. When a Global Controller and Local controller are separated by a firewall, open the following ports on both the inside and outside interfaces of the firewall to ensure proper operation of the Global Controller:

:

Port	Function
22	Secure Shell (SSH)

Port	Function
443	Hyper Text Transport Protocol with Secure Sockets Layer (HTTPS)
8444	Cisco Proprietary data synchronization with Local Controlle r

This chapter contains the following topics:

- Adding Local Controllers, page 2-2
- Importing the Security Certificates, page 2-6
- Monitoring Local Controller Events from the Global Controller, page 2-10
- Preparing to Add and Discover Devices, page 2-10
- Adding Reporting Devices, page 2-11
- Configuring Supported Devices, page 2-12
- L2 Discovery and Mitigation, page 2-12

Adding Local Controllers

Follow these steps to add a Local Controller to the Global Controller:

Step 1 Click **ADMIN > System Setup > Local Controller Management** to display the Zone Controller Information page, as shown in Figure 2-1**.**

Figure 2-1 Zone Controller Information Page

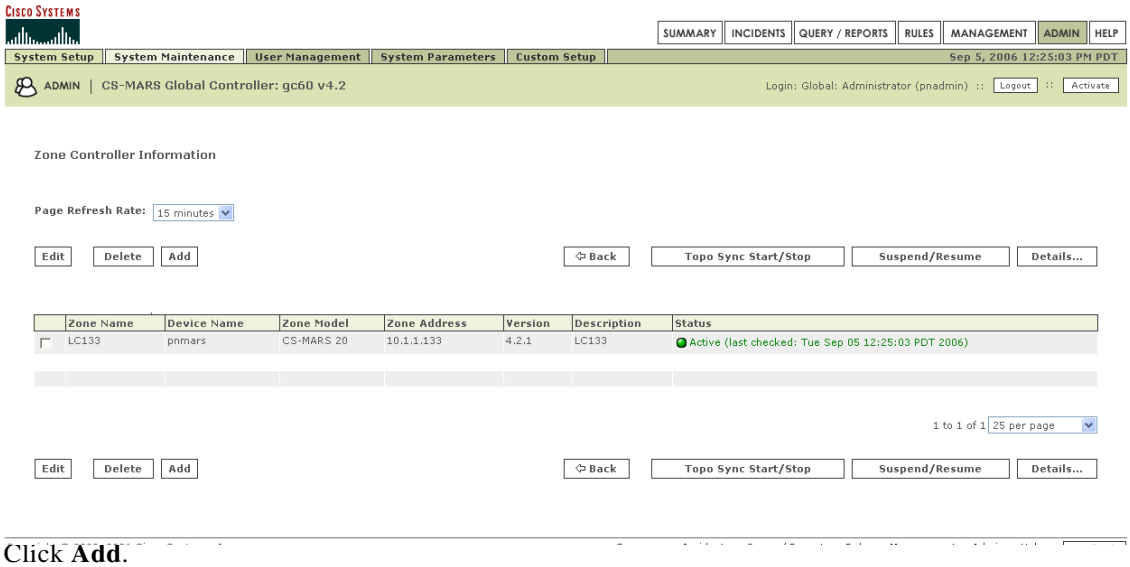

Step 2 Click **Add**.

A pop-up window appears in which you can add a Local Controller to the Global Controller.

Figure 2-2 Local Controller Information Page

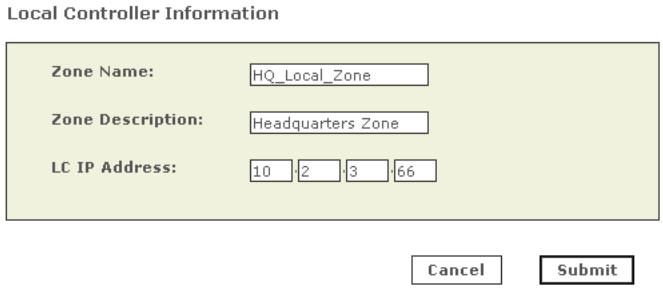

Step 3 Enter values for the following settings:

- **Zone Name**. Enter a name for this zone. This name is used to uniquely identify the networks within this zone relative to other zones. For example, many companies use the same private network addresses behind NATed gateways. The zone combined with the network address allows you to reuse the same network address on your private networks.
- **Zone Description**. Enter a description of the zone
- **LC IP Address**. Enter the IP address of the Local Controller that monitors this zone.

Step 4 Click **Submit** to save the values.

Before the Global Controller can communicate with the Local Controller, you must import the security certificate into the Global Controller. For more information, see Importing the Security Certificates, page 2-6.

Topology Synchronization

For the Global Controller to display a summarized and merged view of topology for its Local Controllers, topology data from all the Local Controllers must be pushed to the Global Controller. When you add a Local Controller to a Global Controller, the topology synchronization process begins and completes automatically.

When synchronized with Local Controllers, the Global Controller contains all the security and monitoring information of the Local Controllers (as displayed on **Admin > System Maintenance > Security and Monitor Devices)** and can display the combined topological maps of the Local Controllers with the following constraints:

- Devices common to Local Controllers are merged in the Global Controller topology. If you have a router listed on different Local Controllers, it only shows up once in topology graphs.
- Networks common to Local Controllers are not merged in the Global Controller topology, but are displayed as separate topologies even if they are the same network.

Topo Sync Start/Stop

When you change Local Controller topology or it otherwise becomes out-of-sync, you can re-synchronize the Local Controller and Global Controller by clicking **Topo Sync Start/Stop** on the Zone Controller Information Page. The **Status** field reports the current state of the synchronization process. Table 2-1 lists and describes all possible status messages.

An out-of-sync condition can occur when unexpected errors or events (device, software, network, etc.) disrupt communication between the Local and Global Controllers.

Suspend/Resume

The **Suspend/Resume** button toggles the communication link on and off between the Global Controller and the Local Controller. When suspended, the Local Controller cannot communicate with the Global Controller.

Note Incident, topology, and other information cannot be uploaded to the Global Controller when the Local Controller communication is suspended.

Table 2-1 Local Controller Status Messages on Zone Controller Page

Status Field Values	Description and Action
Active (last checked: (*Time_and_Date_last_checked*)	The Local Controller is online, connected, and synchronized with the Global Controller.
Suspended	Communications between the Local Controller and the Global Controller have been manually halted with the Suspend/Resume button. To re-establish communication, select the Local Controller and click **Suspend/Resume**.
Synchronizing (*progress*)	The Global Controller and Local Controller are comparing and updating their topology information tables.
Deleting in progress	The Global Controller is purging the selected Local Controller configuration and data from its database. If the Global and Local Controllers can communicate, the Local Controller is purging Global Controller configurations to change from monitor to standalone mode.
Not Responding (last checked: *Time_and_Date_last_checked*)	The Local Controller cannot be detected on the network. Check network status and connections.
Local Controller is online but is not responding (last checked: *Time_and_Date_last_checked*)	The Local Controller can be detected on the network, but does not respond. The problem or delay may clear, the status can return to Active.
Zone has standalone license	The Local Controller model indicated is not supported by the Global Controller.
Global controller license does not allow adding model PNMARS-100 for monitoring	The Local Controller model indicated is not supported by the Global Controller.
Global controller license does not allow adding model PNMARS-100X for monitoring	The Local Controller model indicated is not supported by the Global Controller.
Global controller license does not allow adding model PNMARS-200 for monitoring	The Local Controller model indicated is not supported by the Global Controller.
Zone version is different	The Global and Local Controllers are operating with different software versions. Update one or the other or both as appropriate.

Table 2-1 Local Controller Status Messages on Zone Controller Page

Status Field Values	Description and Action
Global license is Local Controller license	Enter the correct Global Controller license in the Global Controller at **Admin > System Maintenance > Set License Key**.
Global certificate not in LC or local certificate not on GC	Copy the Global Controller security certificate to the Local Controller, and the Local Controller security certificate to the Global Controller at **Admin > System Maintenance > Certificates**

Deleting Local Controllers

To delete a Local Controller from the Global Controller and return the Local Controller to Standalone mode, do the following steps:

Step 1 Click **ADMIN > System Setup > Local Controller Management,** to display the Zone Controller Information page, as shown in Figure 2-3.

Figure 2-3 Zone Controller Information Page

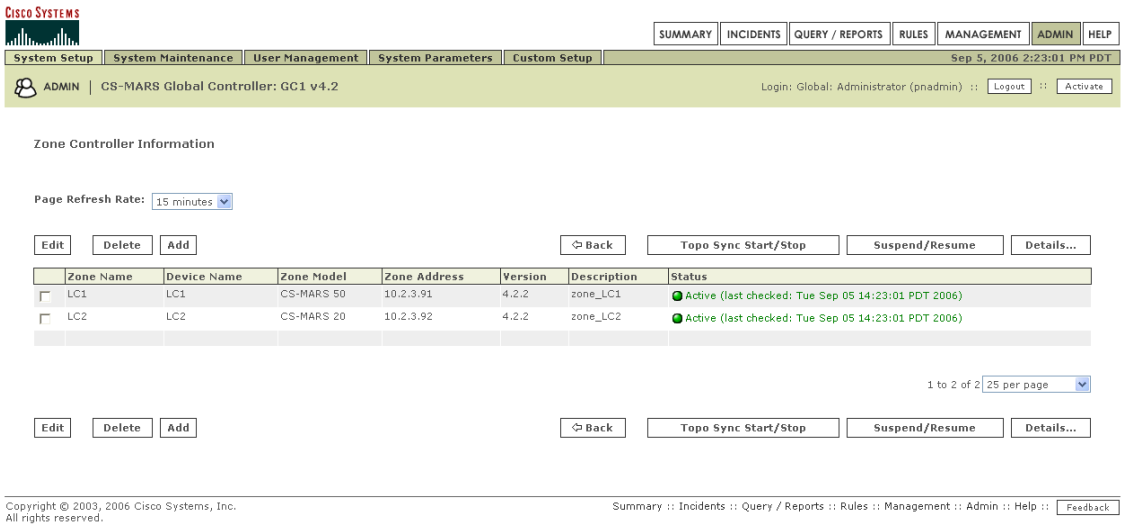

Step 2 Click the checkbox of the Local Controller to delete, and click **Delete**.

A Yes/No confirmation dialog box appears. Click **Yes** to remove configuration info and data from the Global and Local Controllers.

If the status of the Local Controller is **Not Responding**, a Continue/Cancel dialog box appears. Because the Global Controller cannot communicate with the Local Controller, clicking **Continue** removes only the Local Controller data from the Global Controller. To remove the Global Controller configuration information from the Local Controller, you must execute a **pnreset -s** CLI command on the Local Controller as explained in the following URL:

http://www.cisco.com/en/US/products/ps6241/products_installation_guide_chapter09186a008083b881.html#wp1239868

Note If you do not remove the Global Controller configuration from the Local Controller, errors may occur when the Local Controller attempts to contact the Global Controller. Moreover, the Local Controller cannot be added to a Global Controller until it is reset.

The duration of the deletion process varies with the amount of data to be deleted. A duration of many minutes is possible.

Importing the Security Certificates

Security certificates are used for secure communications between a web browser and the Global Controller, as well as between the Global Controller and any Local Controllers that are managed by the Global Controller. Every Global Controller comes with a default certificate which is unique to each Global Controller. However, users could choose to modify the default certificate using the `sslcert` CLI command. For more information on using the sslcert command, see sslcert, page A-52 in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System.*

Figure 2-4 Changing the Default Security Certificate

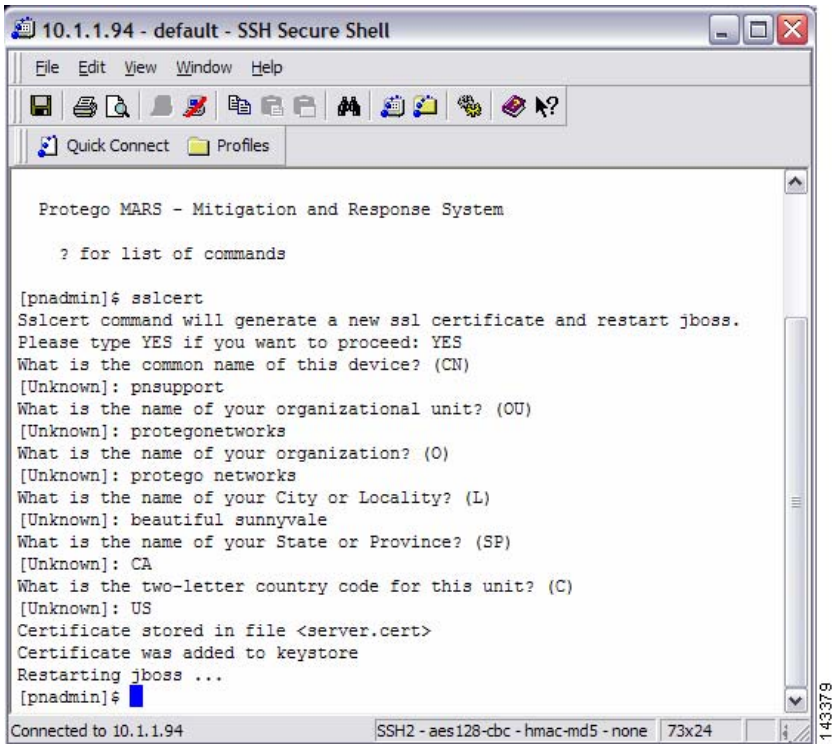

If you wish to install the certificate to an Internet Explorer browser, you must do it during the Global Controller login process.

When the Security Alert pop up appears, choose:

Step 1 **View Certificate**.

Step 2 **Install Certificate**. Then click **Next**.

Step 3 Select *Automatically Select the Certificate Based on the Type of Certificate*. Then click **Next**.

Step 4 Complete the Certificate Import process by clicking **Finish**.

Step 5 Select **Yes** to add the certificate to the Root Store.

Figure 2-5 Global Controller Login Security Alert

The security certificate is used for communication between a Global Controller and any Local Controllers that are managed by the Global Controller.

Although Global Controller and Local Controllers have default security certificates, the Global Controller certificate will need to be exported to all the Local Controllers manually. And all Local Controllers certificates will need to be exported to Global Controller.

To install a Global Controller security certificate on to Local Controllers, follow these steps:

Step 1 From the Global Controller, select **Admin > System Maintenance > Certificates**.

Step 2 Highlight the certificate, and press **Ctrl+C** to copy it.

Figure 2-6 Copy the Global Controller Security Certificate

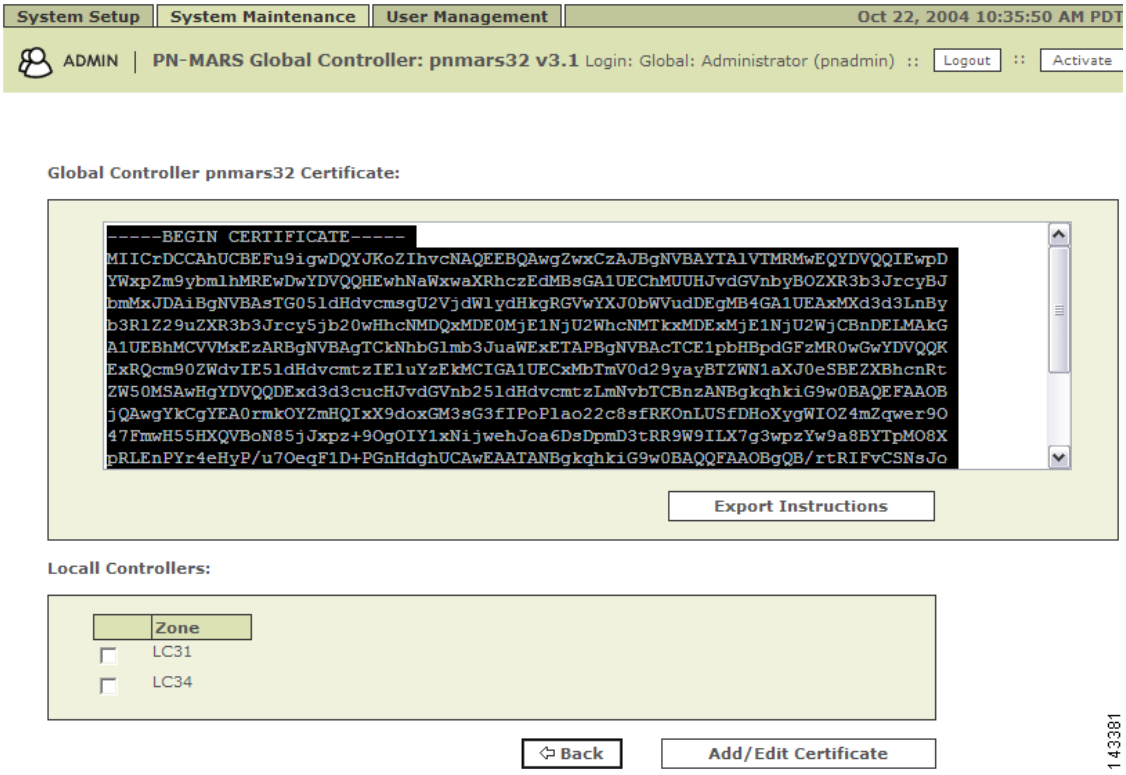

Step 3 Navigate to Local Controller **Admin > System Maintenance > Certificates**.

Step 4 Paste the Global Controller certificate into the **Global Controller Certificate** box.

Step 5 Repeat the process for all every Local Controller that the Global Controller is monitoring.

Figure 2-7 Apply the Global Controller Certificate to the Local Controller

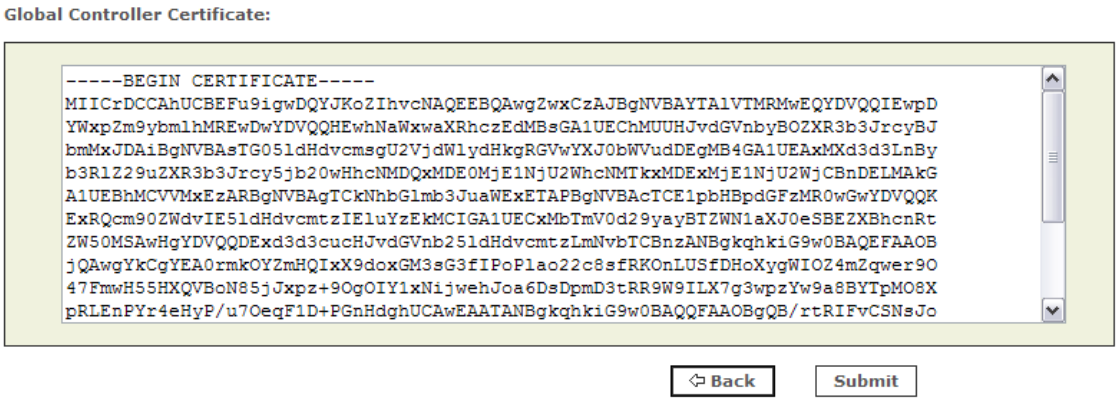

To install a Local Controller security certificate on to the Global Controller, follow these steps:

Step 1 From the Local Controller, select **Admin > System Maintenance > Certificates**.

Step 2 Highlight the certificate and press **Ctrl+C** to copy it.

Figure 2-8 Copy the Local Controller Security Certificate

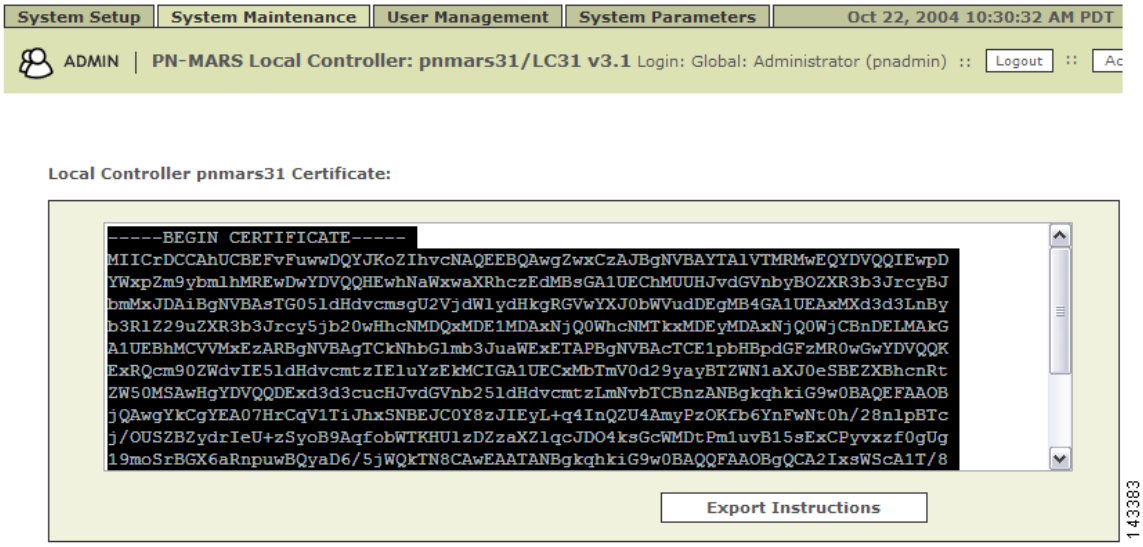

Step 3 From the Global Controller, select **Admin > System Maintenance > Certificates**.

Step 4 Select the specific zone from which this certificate was copied.

Figure 2-9 Select the Appropriate Local Controller

Step 5 Paste the Local Controller certificate to the **Global Controller Certificate** box.

Step 6 Repeat the process from all Local Controllers that are monitored by this Global Controller.

Figure 2-10 Apply the Local Controller Certificate to the Global Controller

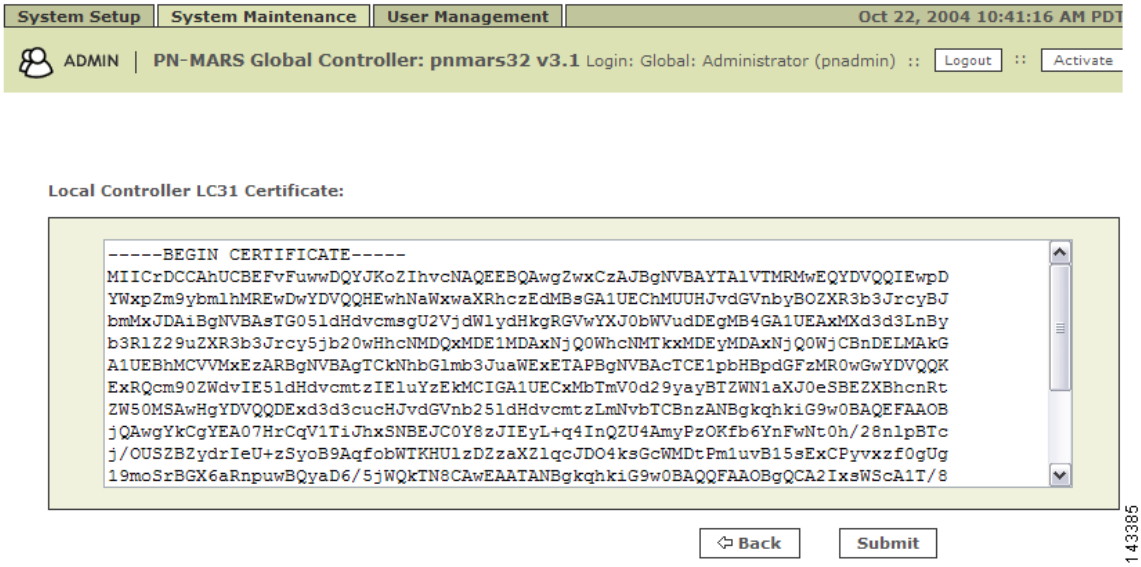

Monitoring Local Controller Events from the Global Controller

The various Local Controllers send summarized information to the Global Controller, which in turn compiles and collates it. There may be a reason you want to suspend, or temporarily hold back, information being sent from one of the Local Controllers. For example, if several of the Local Controller zones are compromised and sending many events at once, you may want to focus on isolating problems on one Local Controller at a time.

If you want to suspend the transmission of information from a Local Controller, follow these instructions:

Step 1 In the Zone Controller Information page, select the Local Controller you want to suspend.

Step 2 Click the **Suspend/Resume** button.

The Local Controller you selected disappears from the list of active Local Controllers, and its output is buffered until you select it and click **Suspend/Resume** again.

Follow the same procedure to resume output from the affected Local Controller.

Preparing to Add and Discover Devices

Before configuring the Global Controller to recognize reporting devices, be aware of the levels of operation supported by a Local Controller. To learn more about the levels of operation for the Local Controller s, see Levels of Operation, page 2-1 in the *User Guide for Cisco Security MARS Local Controller*

Adding Reporting Devices

After you have added the Global Controller's configuration information and rebooted it, you need to configure the third-party devices that report to the Global Controller. All of the event information that passes through these devices is distilled down and sessionized to the information that the Global Controller presents to you. The more information that you can provide for these devices, the clearer the picture you'll get when using the Global Controller.

Note For a list of devices supported by the Global Controller, see the Configuring Supported Devices, page 2-12.

Manual Configuration

In general, you have two choices for adding devices that you want to monitor into your Global Controller. You can create a seed file or you can add each device manually. Seed file support is limited to a few device types; see Configuring Supported Devices, page 2-12.

When manually configuring devices, select the devices that are most interesting to you. Once added, you can come back and edit them as necessary. Manual configuration is also useful when you add or change a single security device on your network. See Configuring Supported Devices, page 2-12 for more information about configuring individual devices.

Note Remember that you do not have to add all of the devices configuration information at once. You can start by adding the device's name and its access IP address. You can always return later, when the Global Controller starts to report to you, and provide more details.

Add a Device Manually

Step 1 Click **Admin > Security and Monitor Devices > Add**.

Figure 2-11 Selecting the Local Controller Zone

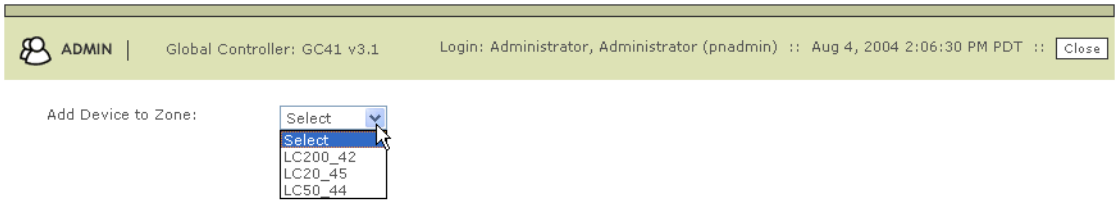

Step 2 Select the Local Controller **Zone** from the pull-down menu. This determines which Local Controller monitors the device. *You are then automatically logged into the Local Controller you have selected.* A pop-up window appears.

Figure 2-12 Entering the Device on the Local Controller

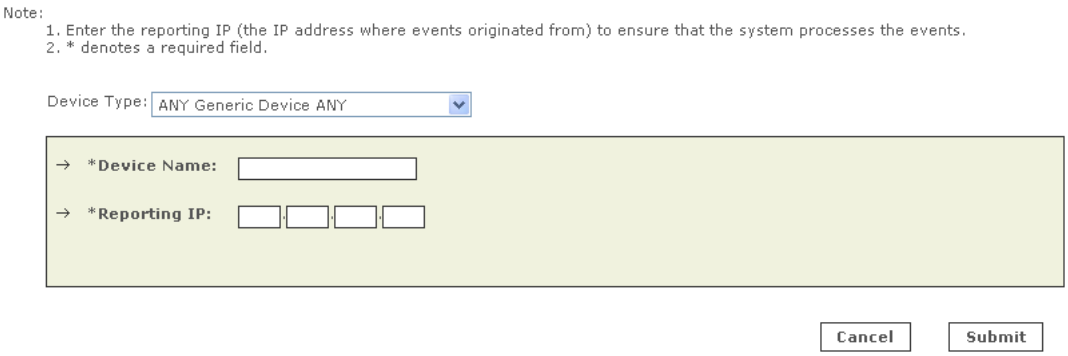

Step 3 Select the device from the pull-down menu.

Step 4 Enter the information needed to communicate with the device.

Step 5 Click the **Submit** button.

Newly added devices on the Local Controller are automatically discovered by the Global Controller.

For more information on installing individual devices, see Preparing to Add and Discover Devices, page 2-10.

Configuring Supported Devices

For most of the security and monitoring devices that you have report to Global Controller, set up and configuration is three-part. You need to:

- Open communication channels to the device.
- Add the appropriate communication information to the Global Controller.
- Make sure that firewalls and routers sitting between the Global Controller and the reporting device are configured to let event traffic pass.

For devices that use agents, modules, or sensors, you need to perform a couple of extra steps.

L2 Discovery and Mitigation

For information on L2 device discovery and mitigation, see Layer 2 Discovery and Mitigation, page 2-30 in the *User Guide for Cisco Security MARS Local Controller*.

CHAPTER 3

Network Summary

This chapter describes the web interface and the components of the Summary tab of the web interface.

Global Controller Network Summary Page Concepts

The Global Controller Summary page differs from the Local Controller summary page in the following ways:

- Devices common to Local Controllers are merged in the Global Controller topology. If you have a router listed on both Local Controllers LC1 and LC2, it only shows up once in topology graphs and on the Summary page.
- Networks common to Local Controllers are not merged in the Global Controller topology, but are displayed as separate topologies even if they are the same network.

Global Controller Technologies

The Global Controller is a complete threat mitigation Global Controller that combines network intelligence, ContextCorrelation™, SureVector™ analysis, and AutoMitigate™ capability in a high performance Global Controller indispensable to subvert real security incidents.

ContextCorrelation groups multiple events and network behavior across NAT boundaries in a session. System and user-defined correlation rules are then applied to multiple sessions to identify valid incidents – significantly reducing raw event data and prioritizing response.

SureVector analysis processes incidents to determine if threats are valid or have been countered by assessing the attack path components – end to end. The result eliminates false positives and resolved threats, and enables full path drill-down visualization and investigation.

AutoMitigate capability identifies available choke point devices along the attack vector and allows you to automate appropriate device commands that can mitigate the threat. The result responsively and accurately prevents or contains an attack by leveraging the infrastructure.

Navigation within the MARS Appliance

The MARS web interface runs within a single brower window. The MARS product functions are categorized with labeled tabs, each tab subdivided with subtabs.

Logging In

Step 1 To login to the Global Controller, enter its IP or DNS address into the browser address field. The login box appears.

Figure 3-1 Global Controller Login Box

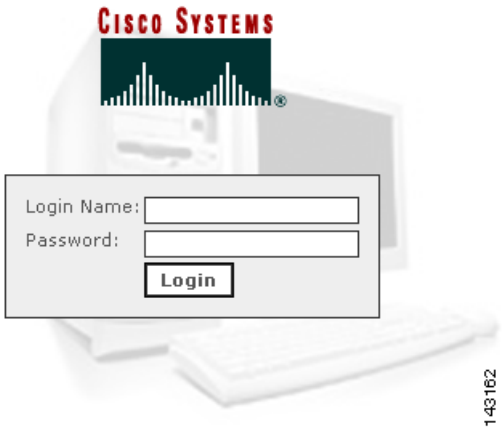

Step 2 Enter your login name and password. If you do not have a login name, contact your network administrator.

Step 3 Click **Login**.

The first page to appear after a login is the Summary tab Dashboard page. The duration of the delay in displaying information results from a combination of the following causes:

- How long the Global Controller has been powered up and connected to the network.
- Amount of traffic on your networks
- Reporting syslog levels of the reporting devices
- Size of the network
- The number and type of reporting devices

For most networks, the Summary page populates shortly after configuration. Some values are only relevant after an interval of time. For example, the values in the **24 Hour Events** and **24 Hour Incidents** tables.

Basic Navigation

The Global Controller uses a tab-based, hyperlinked user interface. When you mouse over an alphanumeric string or an icon that is a clickable hyper-link, the mouse cursor changes to a pointing finger cursor . Figure 3-2 shows some of the clickable objects on the Dashboard page.

Figure 3-2 Links, Icons, and Filters

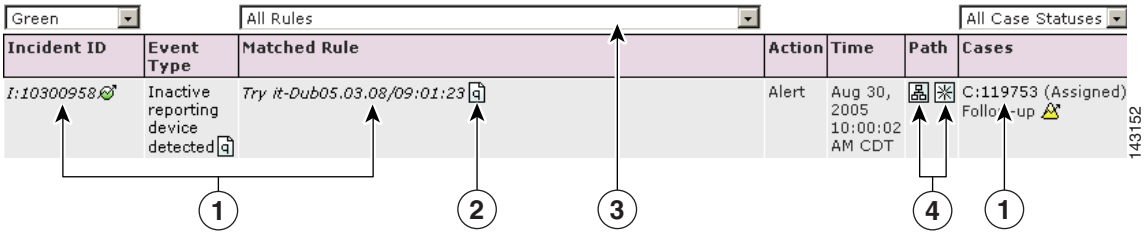

1	Link to the item's detail page or popup window.	**2**	Query icon links to query page. The corresponding query field is populated with the item.
3	Pulldown lists filter what is displayed.	**4**	Path icons launch Path or Incident Vector pop-up diagrams.

Click any of the seven tabs to navigate to the pages relevant to the tab's sub-tabs, as shown in Figure 3-3 though Figure 3-8.

Figure 3-3 Summary Tab

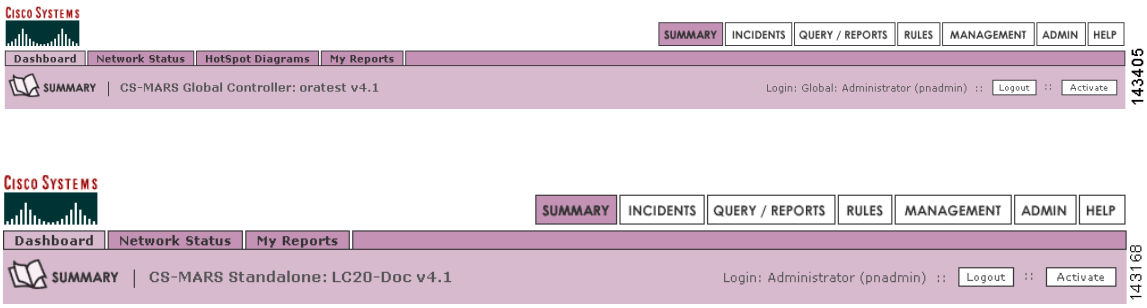

Figure 3-4 Incidents Tab

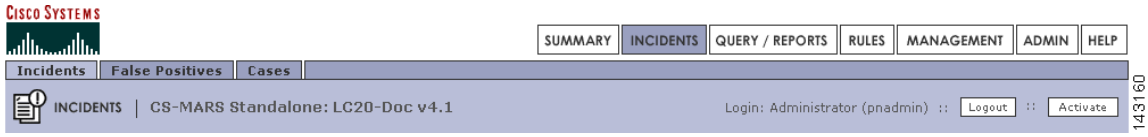

Figure 3-5 Query/Reports Tab

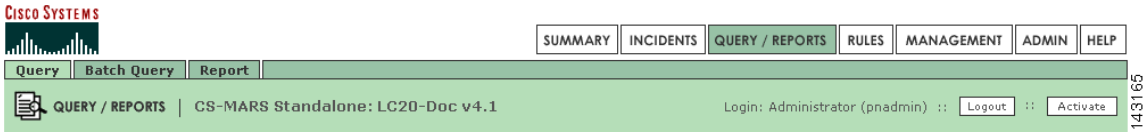

Figure 3-6 Rules Tab

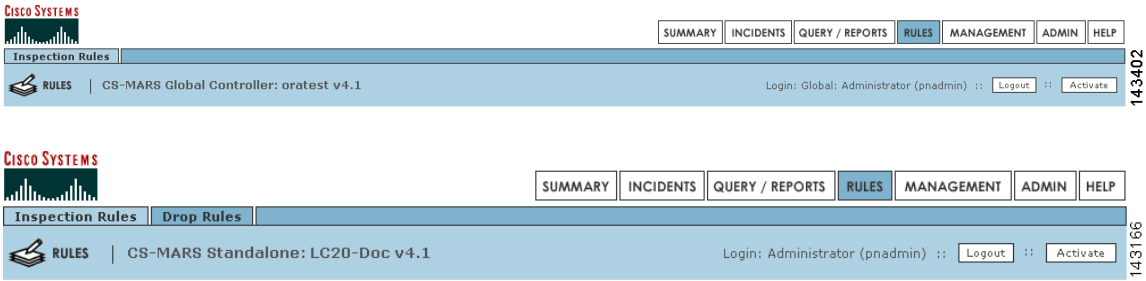

Figure 3-7 Management Tab

Figure 3-8 Administration Tab

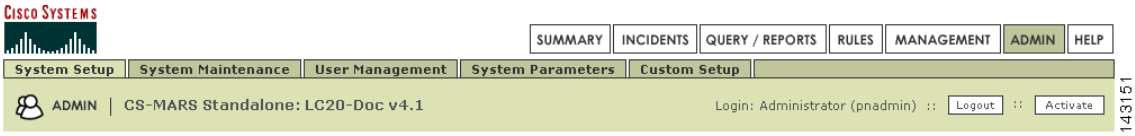

Figure 3-9 Help Tab

Help Page

The Help page, as shown in Figure 3-10, provides URLs to online documentation and a feedback form to submit constuctive comments to the MARS development engineering team.

Figure 3-10 Help Page

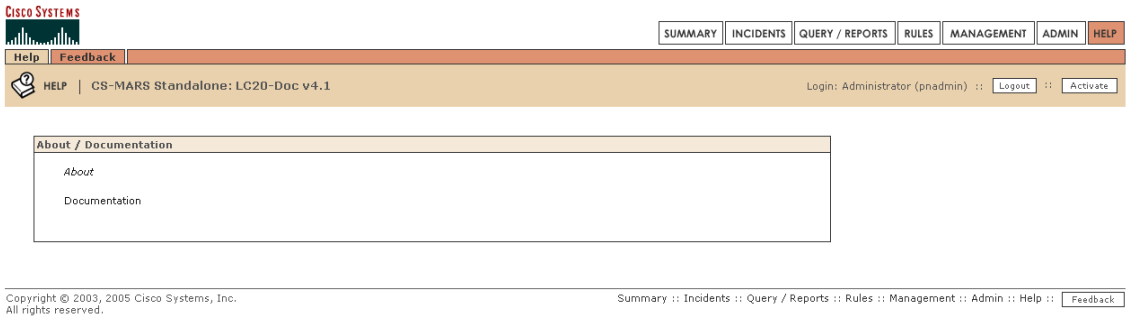

Click **About** to display the software version number running on the MARS.

Click **Documentation** to display URLs to MARS documentation on the Cisco Systems, Inc. website (http://www.cisco.com).

Your Suggestions Welcomed

The **Feedback** button appears at the bottom of most pages, a shown in Figure 3-10.

When you click the feedback button, or navigate to the Feedback page, the feedback dialog box appears, as shown in Figure 3-11.

Figure 3-11 Feedback Dialog Box

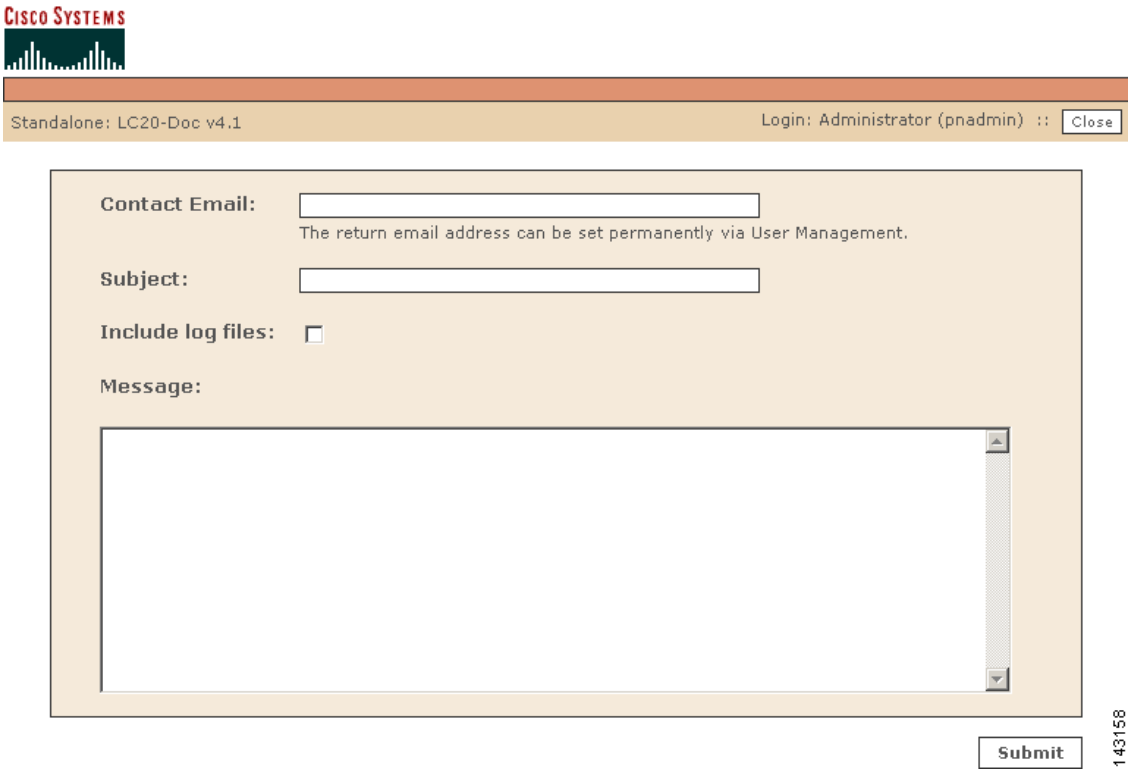

To send your comments to the MARS development engineering team, type in your email address and comments then click **Submit**. When you click the **Include log file** a MARS log file is sent with your message.

Summary Page

From the Summary pages, you can very quickly evaluate the state of the network. The Summary pages include the **Dashboard**, **Network Status**, **HotSpot Diagrams,** and **My Reports**, a shown in Figure 3-12.

Figure 3-12 Summary Tab

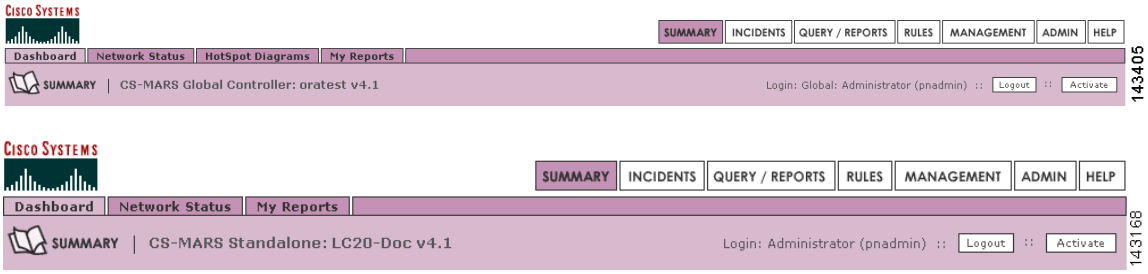

Dashboard

Note When you first view the Summary page after upgrading the Global Controller, expect a small delay while the Java Server pages recompile.

Figure 3-13 The Working Areas on the Dashboard

1	Subtabs	**5**	Tabs
2	Case Bar (Local Controller only)	**6**	Recent incidents information
3	Links to Cases assigned to you.	**7**	HotSpot and Attack diagrams
4	Charts		

Recent Incidents

The first feature to notice about the Dashboard are the recent incidents that have fired. The Global Controller comes with pre-defined rules, and these incidents are the result of those rules firing. These rules are generic, globally applicable, and should serve you well as a starting point once you begin to tune the Global Controller.

Figure 3-14 Drilling-down into Incidents

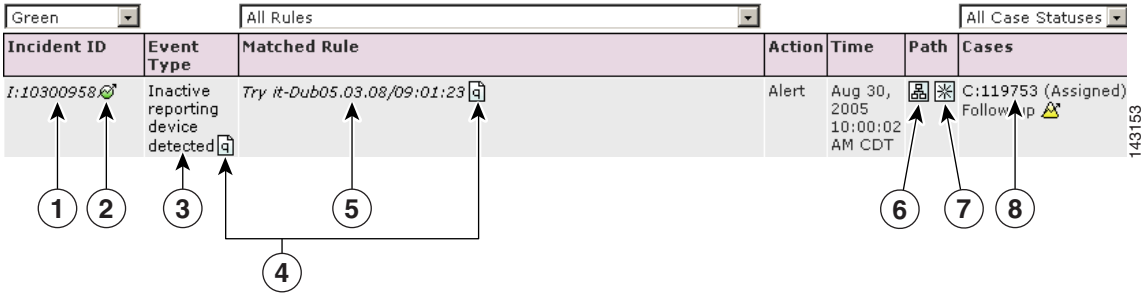

1	Link to the Incident sessions detail page	**5**	Link to the rule details page
2	Incident severity icons Red—Severe threat Yellow—Possible threat Green—Unlikely threat	**6**	Incident Path icon launches the topology diagram popup window
3	Link to the Event Type Details page	**7**	Incident Vector icon launches the incident attack vector diagram
4	Query icon links to Query page	**8**	Link to the View Case page

Sessions and Events

Within a given time window, a session is a collection of events that all share a common end-to-end:

- Source and destination address
- Source and destination port
- Protocol

Event sessionization aggregates event data making it easier to sort and examine. Event sessionization lets the system treat events as single units of information and helps you understand if an attack truly has materialized. It gives you the context of the attack by giving you all the events on that session.

Sessionization works across NAT (network address translation) boundaries – if a session traverses a device that does NAT on that session, the Global Controller is able to sessionize events even if they are reported by two devices on either side of that firewall.

Networks start to show immediate action in the events and sessions categories. Note that the 24 Hour Events table and the Events and Sessions chart are different ways of presenting the same information.

Data Reduction

Data Reduction is a representation of how much event data the Global Controller collapsed into sessions. For example a data reduction of 66% measures three events per session on the average – this number is dependent on many variables particular to your network.

Figure 3-15 Data Reduction

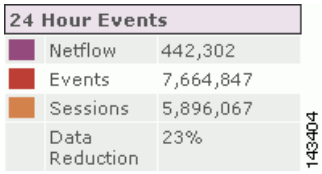

Page Refresh

The Page Refresh Rate polls the Global Controller according to the setting you assign. The default setting is fifteen minutes. The refresh setting remains the same until you log out. This setting only applies to the pages that have the Page Refresh pull-down.

Figure 3-16 Page Refresh

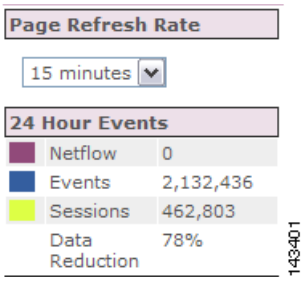

Note You can change the refresh rate with the dropdown list.

Diagrams

The Summary page has two diagrams: the Hot Spot Graph and the Attack Diagram. Global Controller uses the configuration and topology discovery information that were propogated up from the Local Controllers. The following table shows you the icons used in the diagrams.

You can start drilling-down into the diagrams by clicking any of the icons listed in Table 3-1 on page 3-11. You can start drilling-down attack paths in the Attack Diagram by clicking the Path icon . Drilling-down into these diagrams is one of the fastest ways to uncover real-time information about your network.

Figure 3-17 Clickable Hot Spots: Brown = Attackers & Red = Compromised

Note Clouds can represent collections of gateways in the Hotspot graph. A gateway cloud is a device that is unknown to the Global Controller. You can discover gateway clouds by clicking them if you have the SNMP information.

Table 3-1 Icons and States in Topology

	Healthy	Attacker	Compromised	Compromised and Attacking
Clouds		—	—	—
Firewall				
Reporting Host				
Host				
IDS				
Network				
Router				
Switch				
Global Controller (Global Controller or Local Controller)				

To see the diagrams, you need the Adobe SVG viewer plug-in. The Adobe SVG viewer plug-in should automatically install.

Note If you click **No** on the SVG auto-installer, the Global Controller does not prompt you to install it again. If you want to run the auto-installer, open the browser and click **Tools > Internet Options > General > Delete Cookies**.

Figure 3-18 The Hot Spot Graph and Attack Diagram

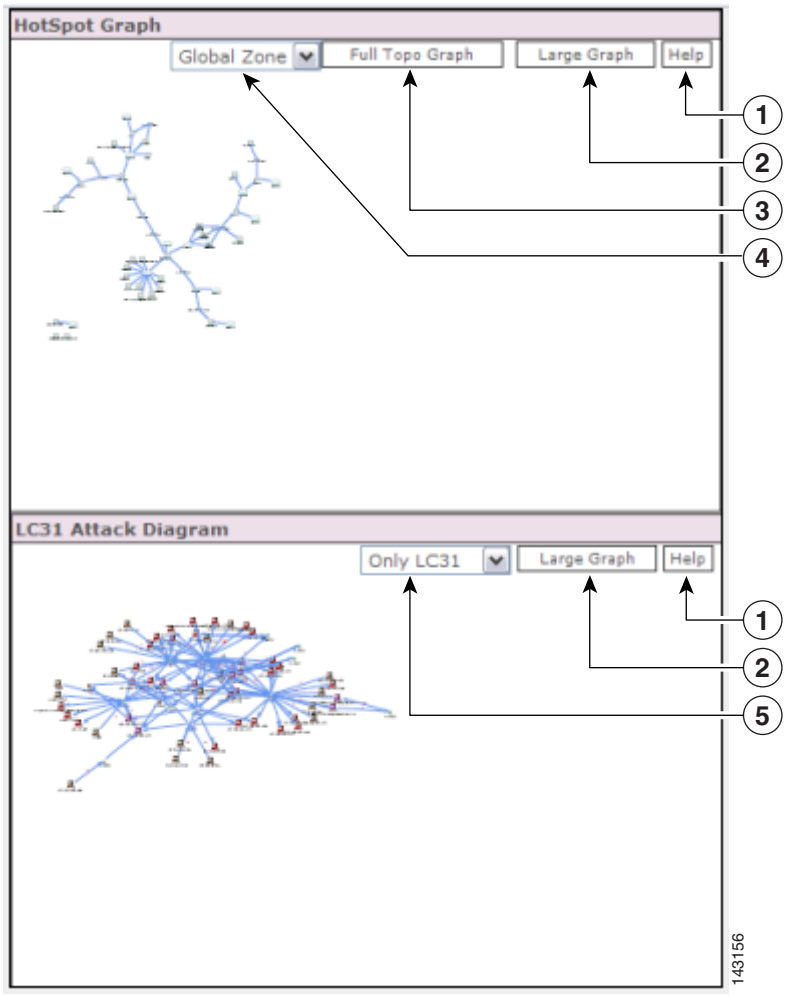

1	Displays SVG Help	**2**	Displays clouds for selected devices on a full page
3	Displays all devices on a full page	**4**	Selects zone to be displayed (Global Controller only)
5	Selects zone to be displayed (Global Controller only)		

Manipulating the Diagrams

- **Pull down** the menu labelled **Global Zone** to select an individual local zone.
- **Right-click** the diagram to zoom in and out, to reset the diagram to its original size, to set the diagram's viewing quality, to search, and to manipulate the SVG image.
- **Alt+click** to use the hand to move the image.
- **Ctrl+click** to use the magnifying glass to zoom in.
- **Ctrl+click and drag** to select an area.

- **Ctrl+shift+click** to use the magnifying glass to zoom out.

Note If the Global Controller discovers an unknown device, it displays that device using a unique name in the form of the string "eth" followed by a hyphen ("-"), followed by the IP address in 32 bit notation, such as "eth-168034561".

Display Devices in Topology

You can specify how to display a reporting device in the HotSpot Graph. By clicking the icon in the Device Display column, you can specify whether to display the device as an individual node on the graph or collapse it within a cloud. By having a device "hidden" in a cloud, you can cut down on the number of devices displayed in the graph, thus making it easier to read at a higher level.

A cloud identifies a collection of networks for which you do not want to define the complete physical topology. Much like when you draw a network diagram on a piece of paper, you can use a cloud to depict networks in which you have no direct interest, but which are needed to represent to complete the diagram. For example, you may want to display only gateway devices or mitigation devices, representing other reporting devices as part of a cloud.

To toggle the display status of a device, follow these steps:

Step 1 Click **Admin > Security and Monitor Devices**.

Step 2 Click the icon in the Device Display column of the device that you want to toggle.

Figure 3-19 The Device Display icons

The icon changes from a host icon to a host within a cloud or vice versa.

Step 3 Click **Activate**.

Network Status

The Network Status page is where you come to get the big picture. On the Network Status page, you can see the charts for:

- *Incidents*

Rated by severity.

- *Attacks: All - Top Rules Fired*

Rated by the highest number of incidents fired.

- *Activity: All - Top Event Types*

Rated by the highest numbers of events of that type.

- *Activity: All - Top Reporting Devices*

Rated by the total number of events reported by each security device.

- *Activity: All - Top Sources*

The top IP addresses that appear as session sources, ranked by session count.

- *Activity: All - Top Destinations*

The top IP addresses that appear as session destinations, ranked by session count.

For all of the charts on this page, you can set different time frames, the size of the chart, view the latest report, and so on, by clicking on the buttons in the chart's window.

Reading Charts

These are stacked charts. You can tell which severity of incident your network has most experienced for the day by looking for the dominant shade. In the figure below, low priority green incidents cover less area than high priority red incidents because they have occurred less often.

Figure 3-20 A Day's Events and Netflow with the Legend Displayed

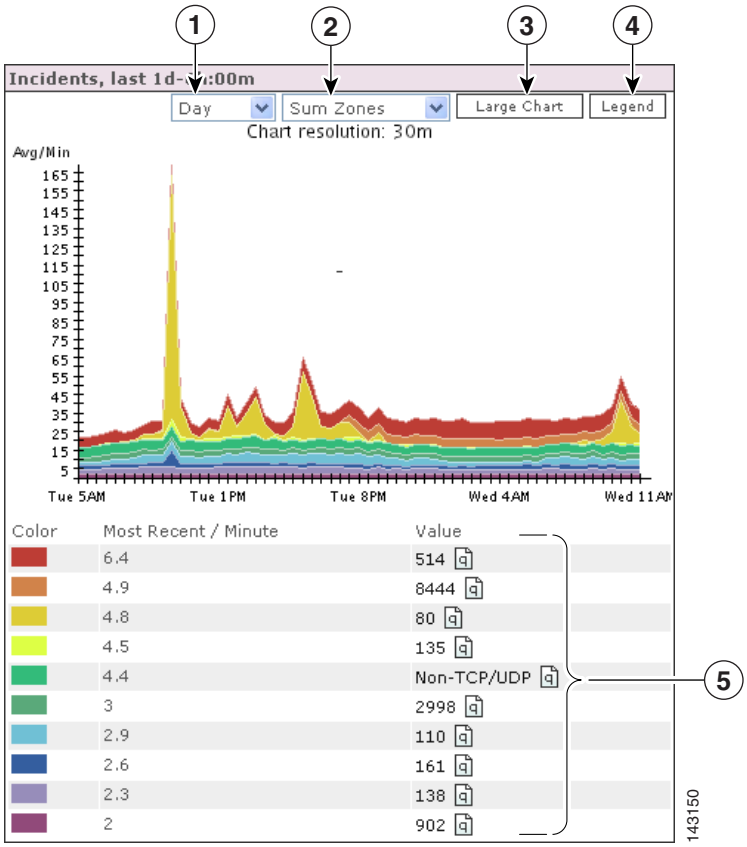

1	Displays values by hour, day, week, month, quarter (the last 3 months), or year.	**2**	Sets chart to represent the sum of all zones or each individual zone (Global Controller only).
3	Displays a larger version of the chart.	**4**	Displays the chart legend.
5	The chart legend		

To read the charts most efficiently, note that it is solely the thickness of a particular color that determines its value at that point – and that a spike (or drop) in any particular color could be caused by a spike (or drop) of a different color lower down in the stack.

A perfectly flat line indicates that Global Controller received no data during that time period.

Figure 3-21 A Flat Line in a Week's Top Rules Fired

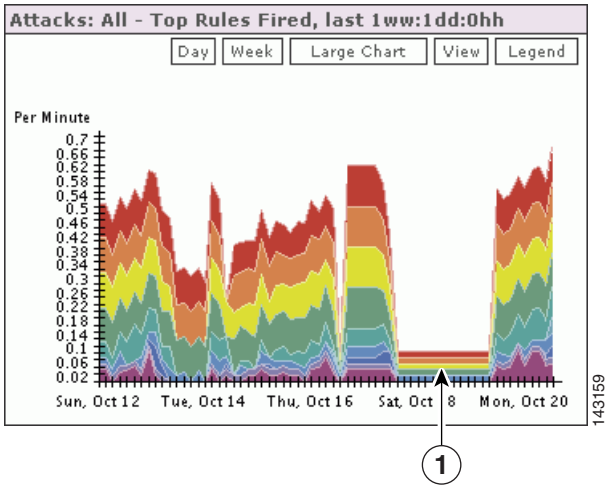

1	The flat line in the Top Rules Fired chart

In the following Incidents chart, you can see the top incidents for the week, starting eight days in the past.

Figure 3-22 Eight Days of Incidents

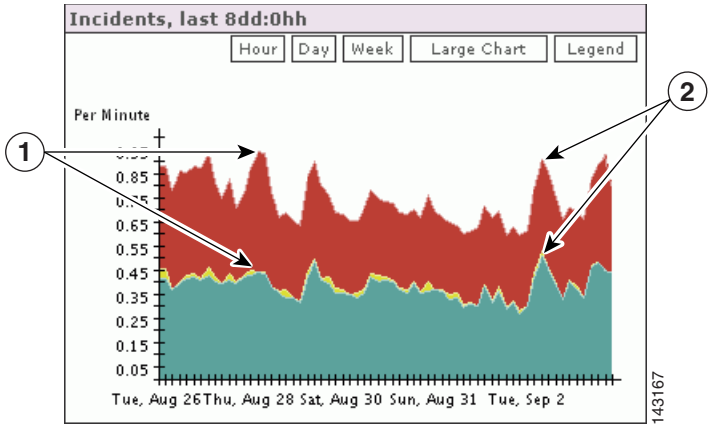

1	A more drastic spike in red is not offset by the green incident	2	Incident spikes are built upon each other

Hotspots

The Hotspots page contains topology graphs of the hotspots on each of the Local Controllers connected to your Global Controller. You can use the pull-down menu to select whether to view the hotspot for a single Local Controller or combined hotspots for all the Local Controllers connected to your Global Controller.

Clicking on the **Full Topo Graph** button displays a detailed graph of the topology; clicking the **Large Graph** button displays the attack on a full page. Clicking the Details button logs you into the Local Controller and displaysthe hotspot graph there.

My Reports

The My Reports page is where you can choose the reports that you want to view. As long as you are using the Global Controller with your log in name, the reports that you have selected appear here.

To set up reports for viewing

Step 1 Click the **Edit** button on the My Reports page.

Step 2 Select the radio button next to the report that you want to see as a chart.

Step 3 Click **Submit**.

Global Controller now displays the chart that you selected on the My Reports page.

Note Reports must be scheduled to run periodically, that is, every hour or every day. If you activate a report, allow for some time for the data to accumulate.

You can display any number of charts on the My Reports page, however expect slower loading times for large numbers of charts.

The reports that you can select from are pre-defined. When you create your own reports, you can select those to display. See Reports, page 6-19 for more information.

CHAPTER 4

Case Management

This chapter contains the following sections:

- Case Management Overview, page 4-1
- Hide and Display the Case Bar, page 4-3
- Create a New Case, page 4-4
- Edit and Change the Current Case, page 4-5
- Add Data to a Case, page 4-6
- Generate and Email a Case Report, page 4-7

Case Management Overview

The Case Management feature can capture, combine, and preserve user-selected MARS data within a specialized report called a case. The following data can be added to a case:

- Text annotations
- Incident ID page
- Incident device information (source IP address, destination IP address, reporting device)
- Session Information page
- Query Results page
- Build Report page
- Report Results page
- View Case page (the current case can reference another case)

Any user can create or alter any case. You can assign a case to a MARS user on the same machine, and can change the status of a case to assigned, resolved, or closed. The contents of a case are displayed by category on a single GUI page (View Case), and can be automatically assembled into a single HTML case document. You can email the Case Document to any MARS user account or user group.

Note When a case is closed, you can still email it, annotate it, add device information, and include a reference to another case.

Case information collected on incidents, sessions, queries, reports and mitigation logs are forensic evidence pertinent to the following:

- Audits (for example, regulatory compliance audits)
- Justifications for modifying ACLs or policy changes
- Notes for MARS false positive tuning
- Examples of allowed and prohibited behavior.

The case preserves and displays the selected data as it appeared when the data was added to the case, regardless of subsequent changes to the MARS state. For example, MARS data can be purged, topology can change from automatic discoveries or vulnerability scanning, and overall configuration can change when you edit rules or reports, but the data reported in the case remains the same as the time it was captured.

Note As of MARS software version 4.1.1 the Case Management feature replaces the incident escalation feature.

The Case Management homepage is the Cases subtab of the Incidents tab as shown in Figure 4-1.

Figure 4-1 Case Management Tab—Local Controller

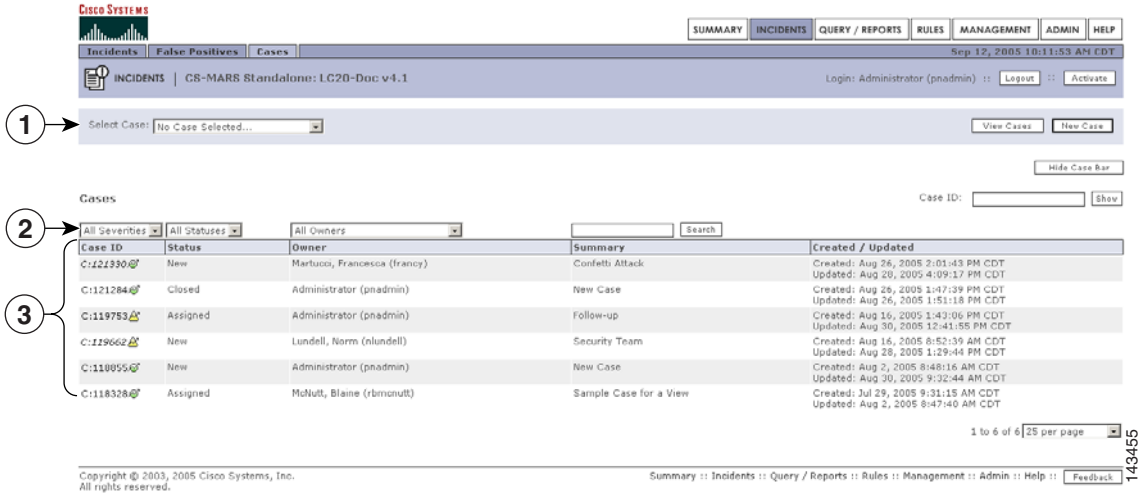

1	Case Bar	**2**	Dropdown Display Filters
3	Individual Cases		

All new, assigned, resolved and closed cases can be accessed from the Cases subtab.

To view the contents of a case, click the Case ID number of a case. The View Case page appears, as shown in Figure 4-2.

To generate an HTML document of the **View Case** page content that can be emailed, click **View Case Document** at the bottom of the **View Case** page. Graphs and charts plotted from reports are also captured in the Case Document.

Figure 4-2 The View Case Page—Local Controller

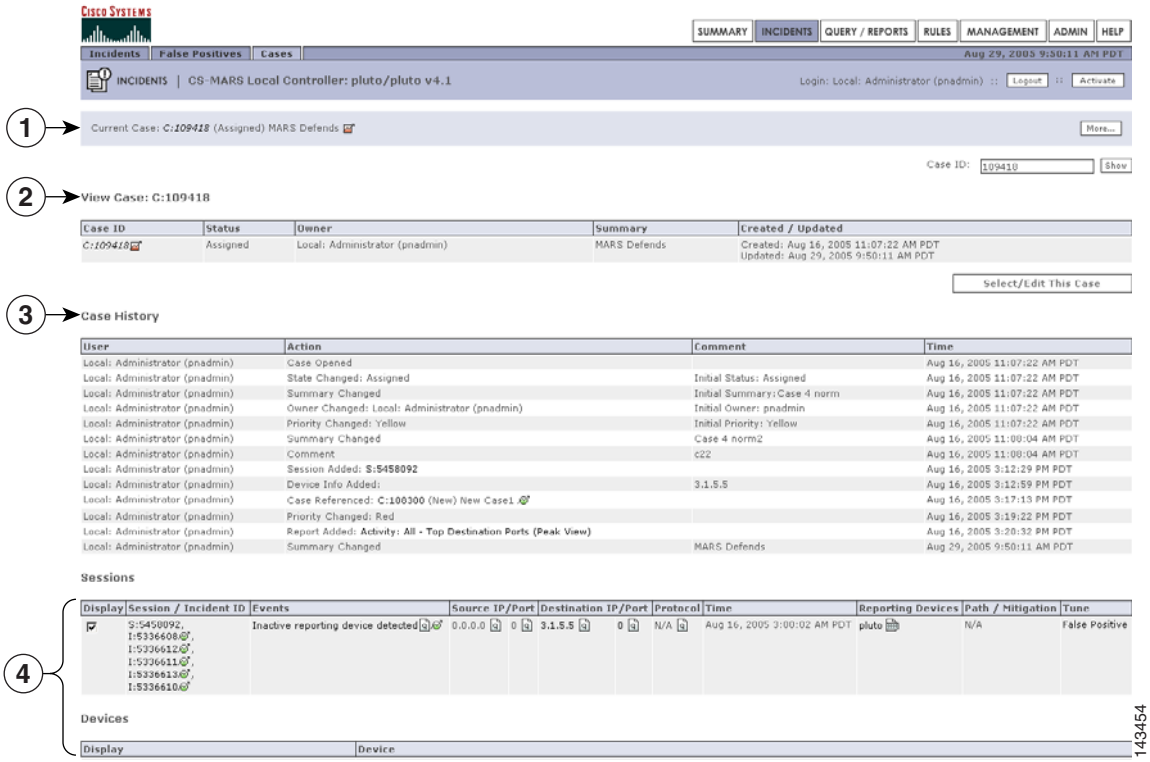

1	Case Bar—Identifies current case	**2**	View Case identifier—Shows the attributes of the case
3	Case History—Log of all changes made to the case	**4**	Summary of data added to the case

Case Management Considerations for the Global Controller

Case management on the Global Controller differs from the Local Controller implementation as follows:

- Cases are not created on a Global Controller.
  They can be viewed and modified.
- The Global Controller does not have a Case Bar.
  All Cases are selected from the Incident -> Cases page.
- The Cases page has an additional dropdown filter to display cases per Local Controller.

Hide and Display the Case Bar

The Case Bar displays by default. When displayed, the Case Bar appears at the top of each page. The Case Bar must be displayed to create or modify a case.

Hiding the Case Bar

To hide the Case Bar, perform the following steps:

Step 1 Navigate to the **Cases** subtab (**Incidents > Cases**), as shown in Figure 4-3.

Figure 4-3 Case Bar Displayed on the Incidents Page

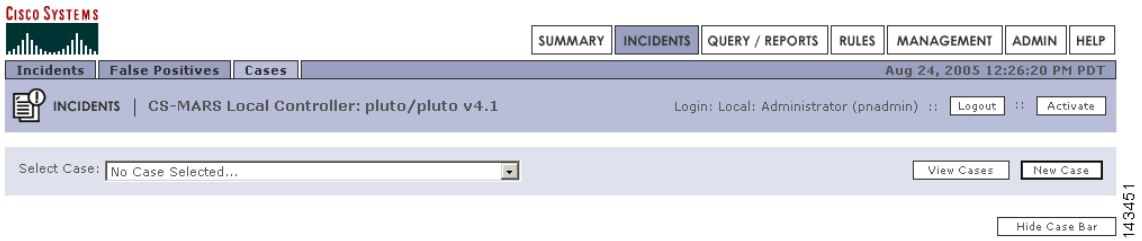

Step 2 Click **Hide Case Bar.**
The Case Bar no longer appears on all tabs, as shown in Figure 4-4**.**

Figure 4-4 Case Bar Hidden on the Incidents Page

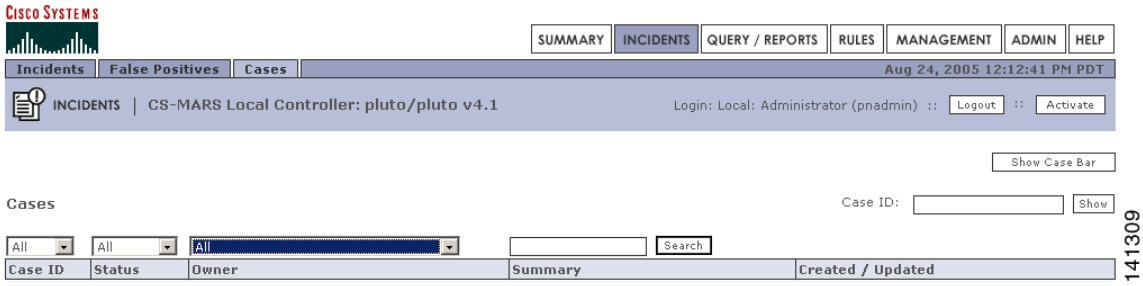

Displaying the Case Bar

To Display the Case Bar, follow these steps:

Step 1 Navigate to the **Cases** subtab (**Incidents > Cases**) as shown inFigure 4-4 .

Step 2 Click **Show Case Bar**
The Case Bar, as shown in Figure 4-3 now appears on all pages.

Create a New Case

To create a new case, perform the following procedure:

Step 1 Display the Case Bar as described in the section, Hide and Display the Case Bar.

Step 2 Click **New Case**.
The Add a New Case Dialog box appears, as shown in Figure 4-5.

Figure 4-5 Add a New Case Dialog Box

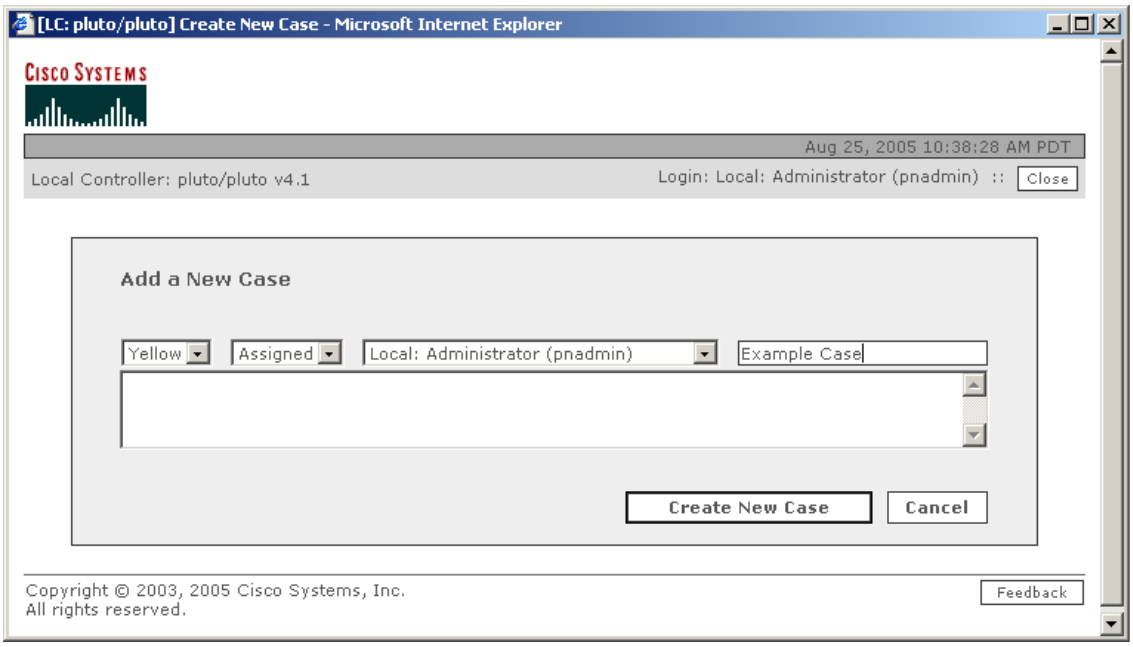

Step 3 Select a severity color, change the state from new to assigned if appropriate, select the owner, replace the default summary name (default is New Case).
Figure 4-5 shows a case with case summary of Example_Case, assigned to the administrator with a yellow priority color (default is Green).

Step 4 Type or paste any annotations into the text space.

Step 5 Click **Create New Case**.
The newly created case is numbered and becomes the current case displayed in the Case Bar as shown in Figure 4-6.

Figure 4-6 Case Bar Shows a Newly-Created Case as the Current Case

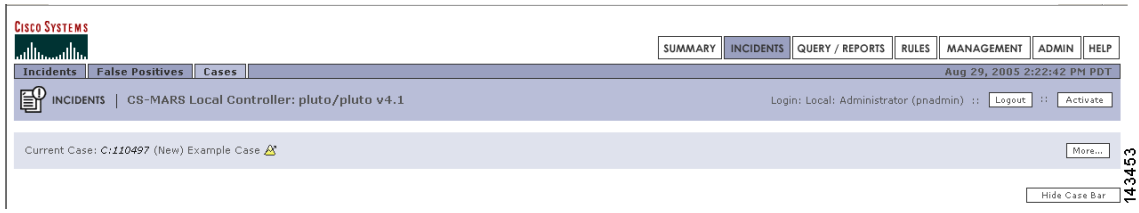

Proceed to the section Add Data to a Case for steps on how to combine various data into a single case.

Edit and Change the Current Case

Editing the Current Case

To edit the Current Case complete the following procedure:

Step 1 Display the Case Bar and click **More**.
The Case Bar Expands to expose the editing options, as shown in Figure 4-7.
See the section Hide and Display the Case Bar for procedures to display the case bar.

Figure 4-7 Expanded Case Bar

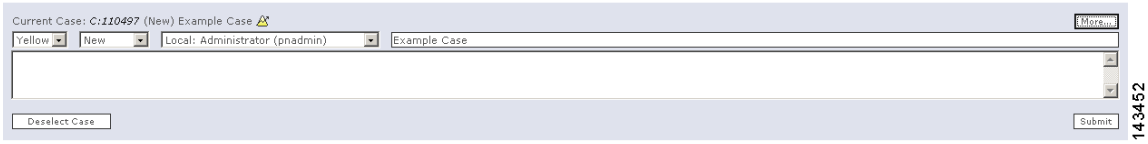

Step 2 Change the severity, status, owner, or summary of the case as required.

Step 3 Add an annotation in the text box as required.

Step 4 Click **Submit**

Deselecting the Current Case

To replace the Current Case case with another, complete the following procedure:

Step 1 Expand the Case Bar as explained in the previous procedure.

Step 2 Click **Deselect**.
The Case Bar drop-down list displays **No Case Selected. . .** as shown in Figure 4-4.

Step 3 To select a different Current Case, select a case from the Case Bar drop down list.

Add Data to a Case

To add data to a case, complete the following steps:

Step 1 Select the Current Case. See the section Edit and Change the Current Case for procedures on selecting the Current Case.

Step 2 Navigate to the page to be captured in the case. In the example, the Query page is selected.

Step 3 Click **Add this. . .** on the Case Bar.

Figure 4-8 Case Bar Add Button

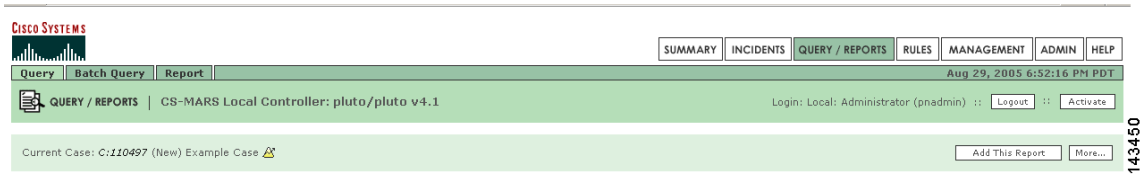

Step 4 To verify that the selected data was added to the case, click the case ID number in the Case Bar to display the View Case page.
In the example shown in Figure 4-8, the selected report should appear in the Reports section of the View Case page. A partial View Case page is shown in Figure 4-2.

Generate and Email a Case Report

You can generate a case report of the case data and email the report to any MARS user group or individual user account. The email event is logged in the case history listings on the View Case page.

To add a new user account or user group, see "Create a New User—Role, Identity, Password, and Notification Information" section on page 8-10.

Note Make sure that the MARS email server is configured. See "Configure the E-mail Server Settings" section on page 8-4 for further information.

To generate a case report and to email it, follow these steps:

Step 1 Select a case from the Cases page or from the Case Bar dropdown list.

Step 2 Click the Case ID number to navigate to the **View Case** page.

Step 3 Click the check box in an item's **Include** field to select or deselect that item for inclusion in the Case Document. By default, all items are selected.

Tip Click **Show Include** to show only those items selected for the Case Document. **Show Include** does not function for cases created in Cisco Security MARS version 4.1.1.

Step 4 Click **View Case Document** at the bottom of the **View Case** page.

MARS generates and displays the case report.

Step 5 Click **Email Case** at the bottom of the report page.

The Case Email dialog box appears, as shown in Figure 4-9.

Figure 4-9 Case Management Email Dialog Box

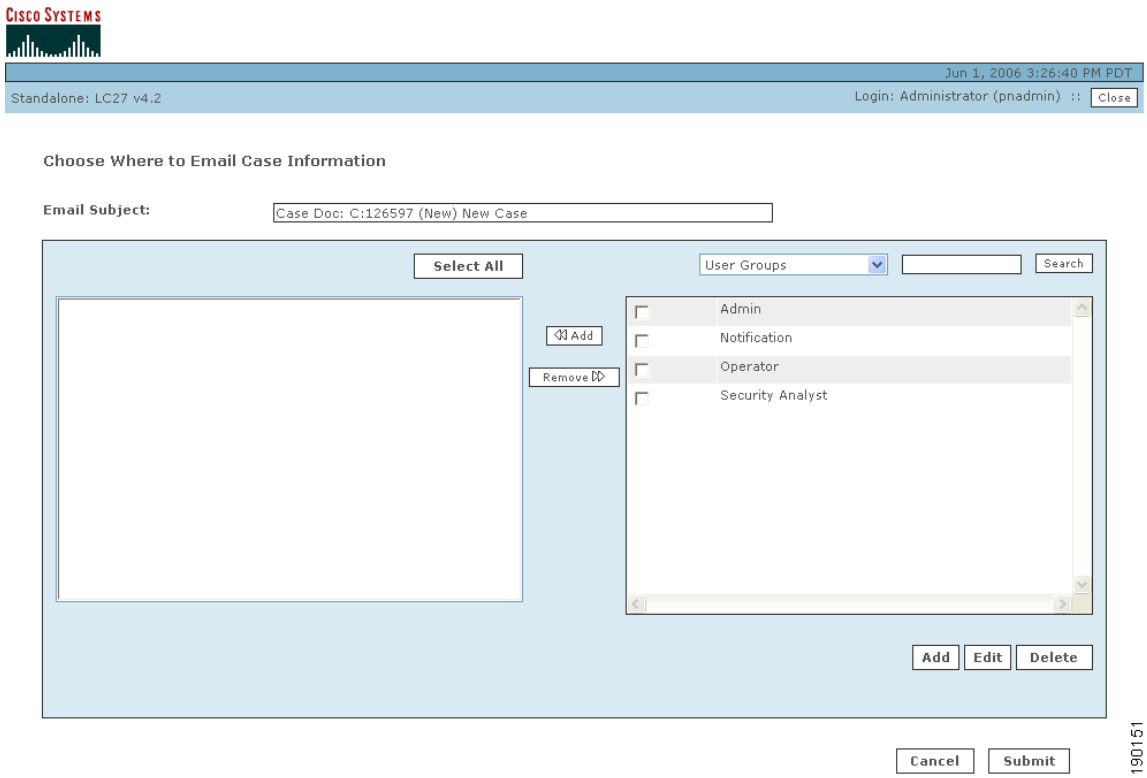

Step 6 Click the check box of the user groups or individual users you want to receive the Case Document, then click **<< Add**.

Tip Select **All Users** from the dropdown menu to display all individual user accounts.

The selected recipients appear in the left-hand area of the dialog box.

Step 7 Click **Submit** to send the Case Document to the recipients.

The email is sent and the case history is updated to show the email event as the lastest item of the case history.

CHAPTER 5

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 5-1 A DoS probe followed by a DoS attack

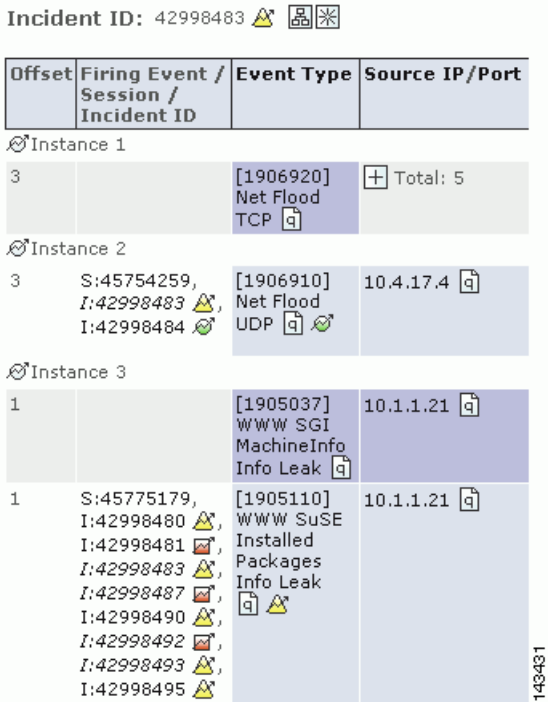

The Incidents Page

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 5-2 Global Controller Incidents Navigation Page

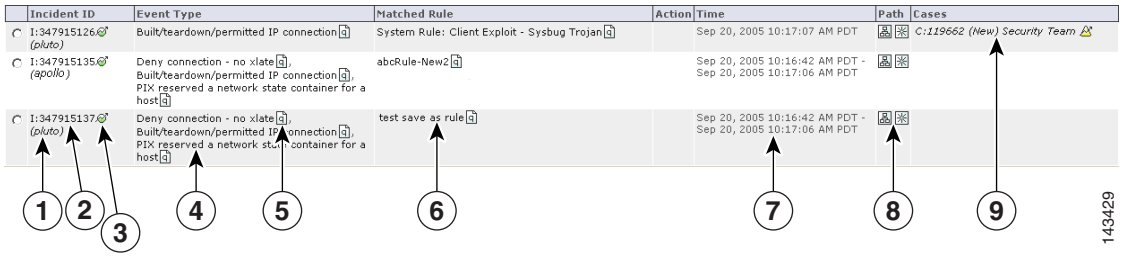

1	Name of the Local Controller reporting the incident, also links to the Local Controller Incident page.	**2**	Links to the Incident Detail page of the reporting Local Controller.
3	The incident severity indication icon	**4**	The events that compose the Incident. Links to the Event Type Details popup 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 5-4 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 5-3 The Incident Details Page

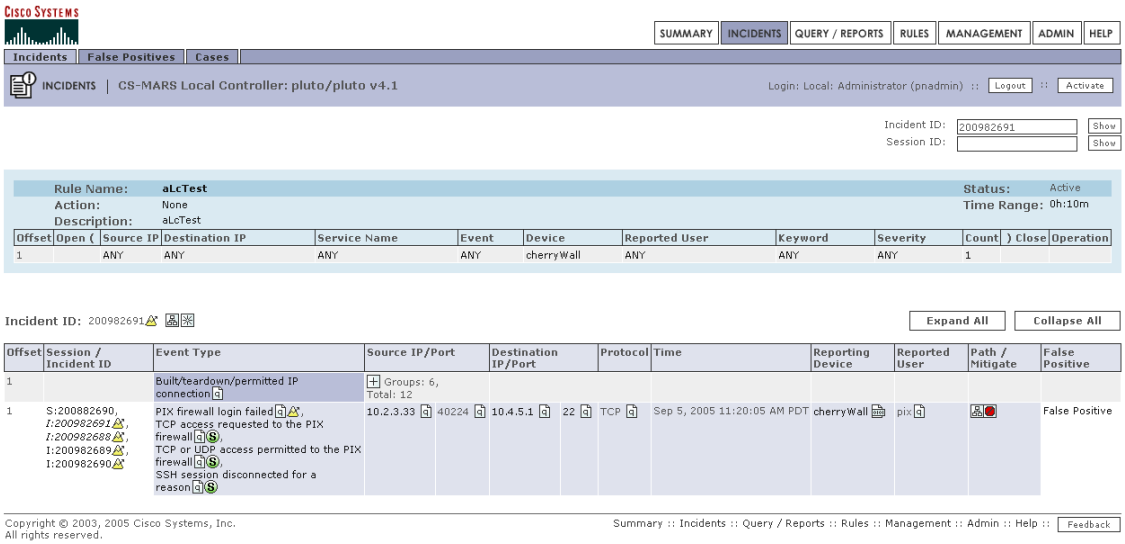

On the top of this page are the tools that let you search for Incident and Session ID and view the Matched Rule.

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 5-4 Expanding a Row in a Table'

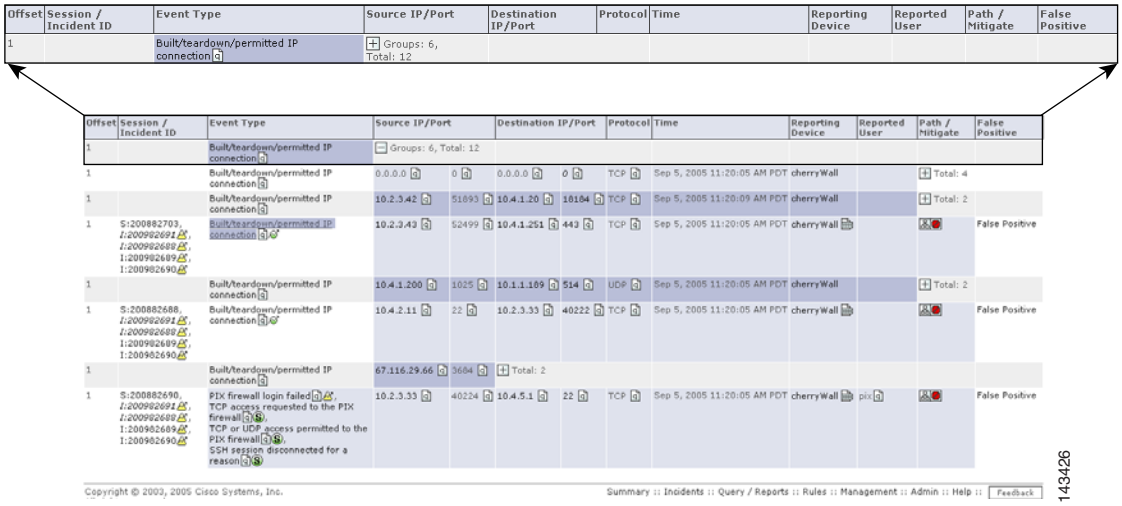

This high-density information table lets you drill deep into incidents. Click the Query icon anywhere on this page to query on a particular criteria. Click the Raw Events icon for raw events for a particular session. You can click the **Tune** link to tune incidents for False Positives, see The False Positive Page, page 5-8 or click the **Mitigate** link to mitigate an attack.

Figure 5-5 Incident Table

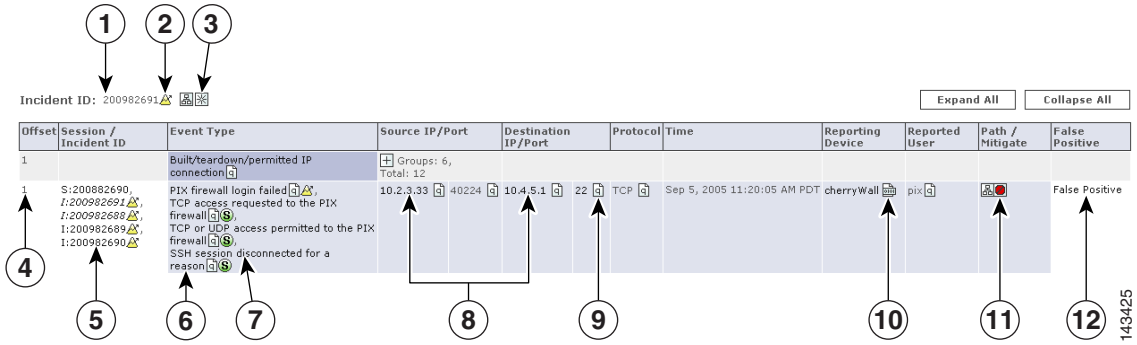

1	Incident ID	**2**	Severity icon
3	Path and Incident Vector icons. Launch popup windows to display Path and Incident Vector diagrams (L2 or L3 attack path information)	**4**	Offset number
5	Links to Session and Incident Detail pages of all incidents within the session	**6**	Links to the Event Type Details pages

7	Launchs 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 5-1.

Table 5-1 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 5-1, 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 the 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 5-6 False Positive Graph for a Local Controller

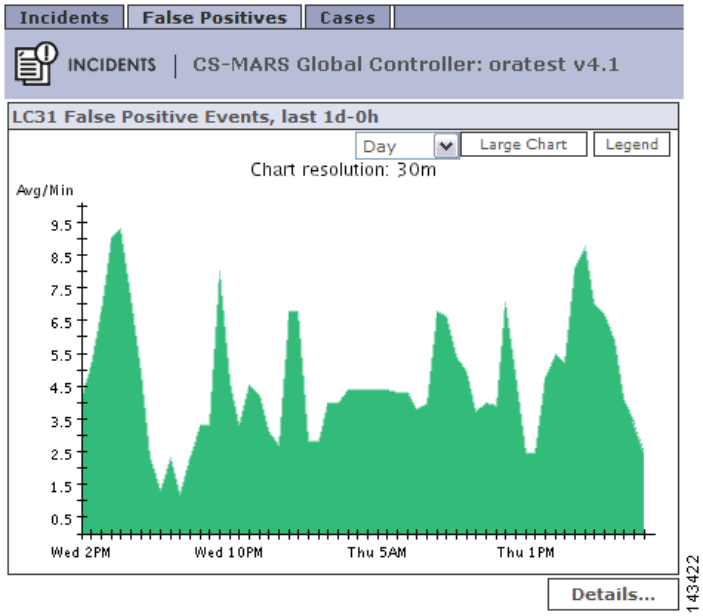

The False Positives page is where you can see groupings of False Positives for each Local Controller or for the sum of all the Local Controller zones. You can change the graph by selecting **Hour**, **Day**, **Week**, **Month**, **Quarter**, or **Year** from the first pull-down menu, and **Sum Zones**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

CHAPTER 6

Queries and Reports

This chapter discusses the following topics:

- Queries
- Perform a Long-Duration Query Using a Report
- Perform a Batch Query
- Reports

Queries

On the Query page, you can run reports as on-demand queries, or create your own query. Many links from other pages bring you to the query page, which then partially populate the query's criteria. Once you have submitted a query, you can save it as a report or a rule.

Queries performed at the Global Controller level are similar to those on an Local Controller, but also include the **Zone** parameter. You can run a query across one or more Local Controllers by specifying their zones. This enables a query at the Global Controller to select zone-specific objects.

When you submit a query from the Global Controller, it is sent out to the Local Controllers specified in the **Zone** parameter. The Local Controllers perform the actual query, send it back to the Global Controller, which then merges and presents the results at the global level.

Figure 6-1 The Global Controller Query Table

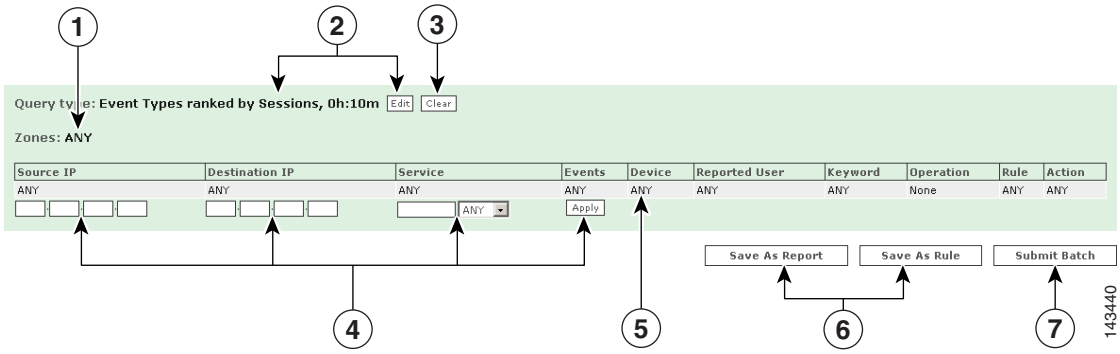

1	Click to select the Local Controller to query.	**2**	Click to set the query type and time range criteria.
3	Click **Clear** to return query values to default values.	**4**	Quick query fields permit entry of values without opening dialog box for the field.
5	Click on a field value to open the dialog box for that field.	**6**	Save the query as a report or as a rule.
7	Click **Submit Batch** to run the query**.**		

Except for the Zone parameter, running a query on the Global Controller is the same as running a query on a Local Controller.

To Run a Quick Query

Step 1 Enter a source IP, destination IP, or a service into the quick query field.

Step 2 Click the **Submit Inline** button to run the query.

Figure 6-2 Running a Quick Query

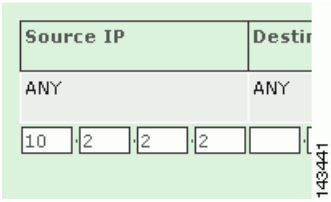

To Run a Free-form Query

Step 1 Enter a source IP, destination IP, or a service into the quick query field.

Figure 6-3 Running a free-form query

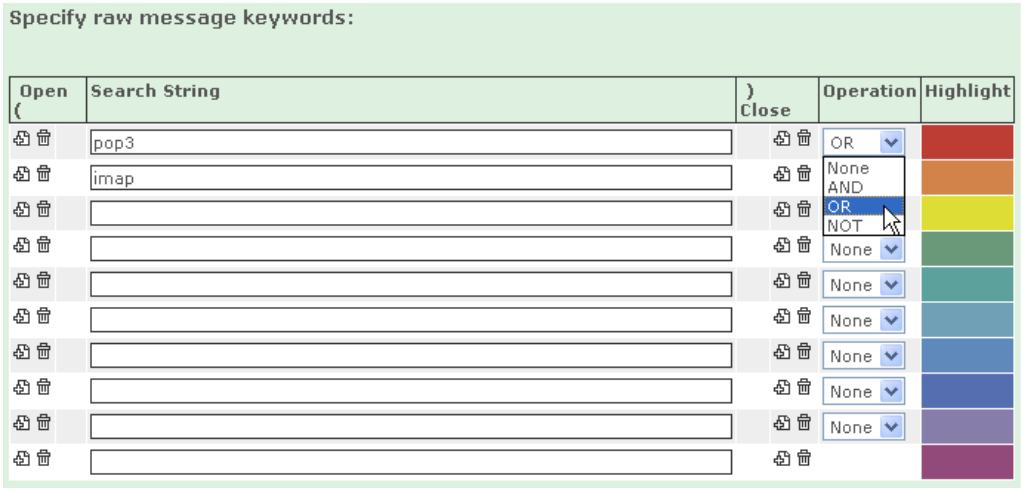

Step 2 Click the name of the query (**[None]** appears as the name if you have none saved) or Edit to enter the rest of the query. You can also click the parentheses icon ( )to add parentheses for nested queries or click the trash can icon ( ) to remove parentheses.

Step 3 Under Search String enter strings to query; under Operation, select the operation (**AND**, **OR**, **NOT**). For the final item in the list, select **None**.

Step 4 Click the **Apply** button.

Step 5 Click the **Submit** button to run the query.

Note The free-form query cannot be saved as a rule.

To Run a Batch Query

Step 1 Enter your data for either a simple or free-form query. If your query is expected to take a long time to run, instead of **Submit Inline**, you may given the option of having it run as a batch query.

Figure 6-4 Construct a Query to Run in Background (Batch Query)

Step 2 Click **Submit...** to make your selection.

Figure 6-5 Choosing the Query Submission Method

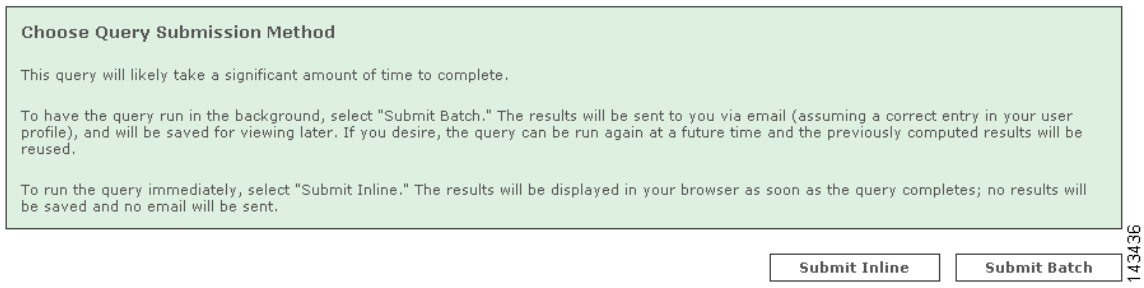

To submit as a standard inline query, click **Submit Inline**.To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the **Batch Query** tab.

If your query is very large, you may only be give the options of **Save as Rule**, **Save as Report**, or **Submit Batch**.

Figure 6-6 Change Query Criteria

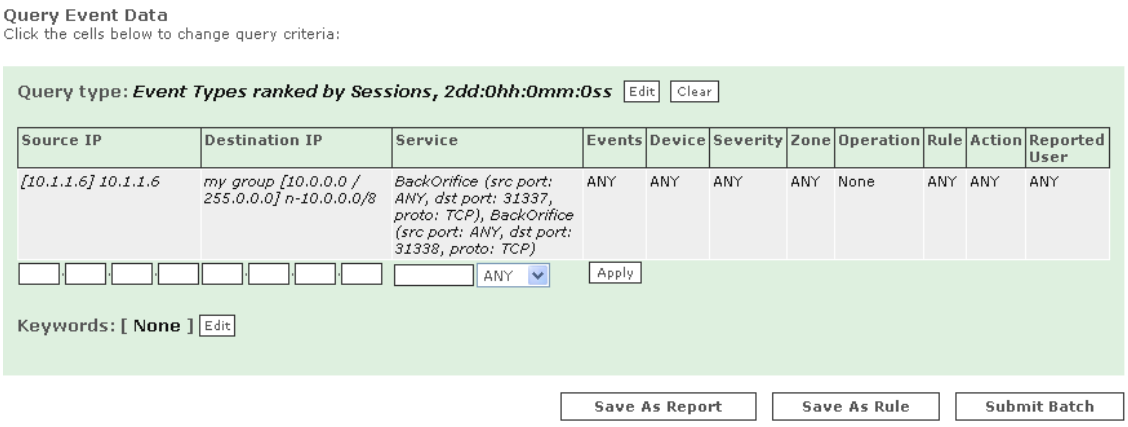

To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the **Batch Query** tab.

Figure 6-7 Select Batch Query

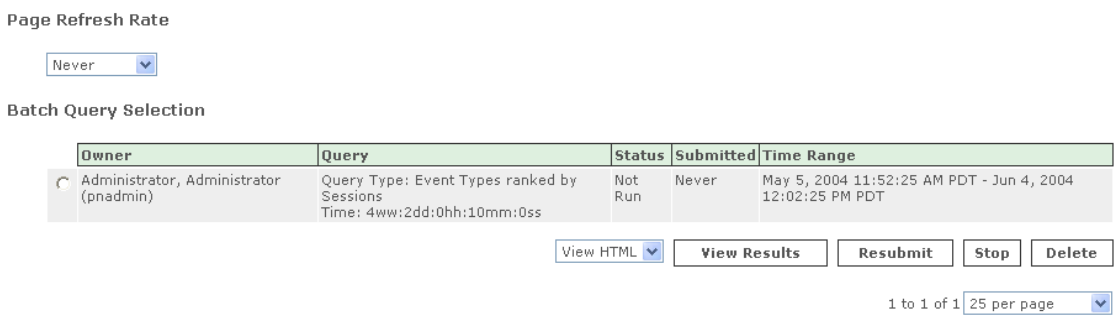

Step 3 To watch the status of the query in real-time, you can use the drop-down list to change the **Page Refresh Rate** from **Never** (the default) to 1 minute, 3 minutes, 5 minutes, 10 minutes, 15 minutes, or 30 minutes.

Step 4 To view the results of the batch query as it is running, click **View Results**. This can be done while the query is in progress.

If the email address in your user profile on the MARS is valid, the results of your batch query are emailed to you when the query has completed, and can also be viewed by clicking **QUERY / REPORTS > Batch Query > View Results**.

Note When you click **View Results** while the query is in progress, the results compiled up to that moment are recomputed. This can make the display take longer to appear than after the results are compiled.

To Stop a Batch Query

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Stop**. The **Status** of the query changes to **Finished**.

To Resubmit a Batch Query

You can resubmit a batch query if you want to restart it. A resubmitted batch query will use previously computed results, thus resulting in a faster query than one submitted for the first time.

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Resubmit**. The **Status** of the query changes to **In Progress**.

To Delete a Batch Query

Step 1 Click **QUERY/REPORTS**, then click the **Batch Query** tab.

Step 2 Click **Delete**.

Step 3 In the confirmation window, click **Delete** to confirm.

Note You can only see your own batch queries and their results. The batch queries of others and their results are not viewable by you, and your batch queries and their results are not viewable by others.

Selecting the Query Type

Figure 6-8 Clicking the Query Type or Edit link

You can select different query criteria by clicking the **Query Type** link or **Edit** button. This lets you determine a query's result format, rank, time, whether it only uses firing events, and the number of rows returned.

Figure 6-9 The Query Criteria: Result Page

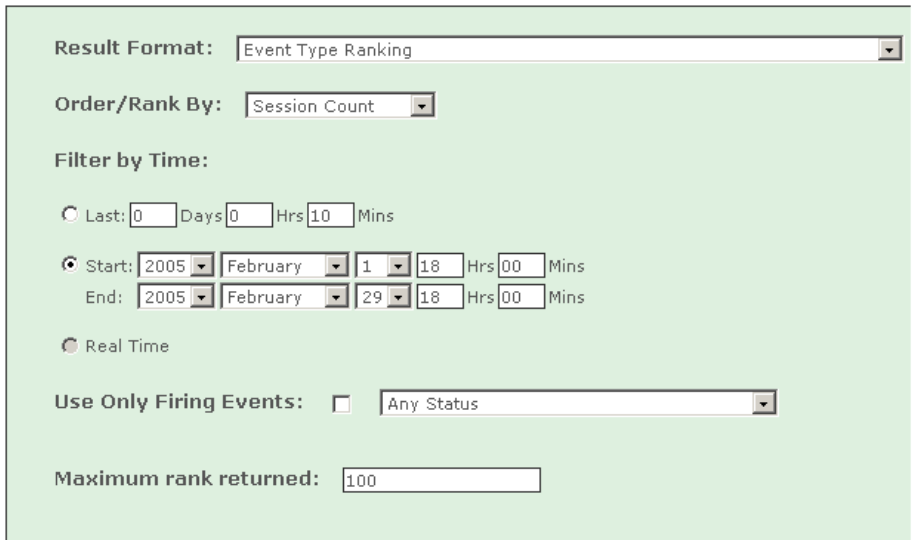

Result Format

- *Event Type Ranking*

Returns the most reported event types. Ranked by either: number of sessions containing at least one of the event type or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Event Type Group Ranking*

Returns either pre-defined or user defined grouped event types. Ranked by either: number of sessions containing at least one event type contained in the group or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Source IP Address Ranking*

Returns source IP addresses. Ranked by number of sessions with that source IP address or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Network Ranking*

Returns top networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Network Group Ranking*

Returns top network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Source Network Ranking*

Returns top source networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Source Network Group Ranking*

Returns top source network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination Network Ranking*

Returns top destination networks that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination Network Group Ranking*

Returns top destination network groups that exists in MARS. Ranked by either: number of sessions that contain events that meet the query criteria or by bytes transmitted in sessions that contain events that meet the query criteria. If a network is excluded, it is excluded from all results.

- *Destination IP Address Ranking*

Returns destination IP addresses. Ranked by either: number of sessions with that destination IP address or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Source Port Ranking*

Returns source ports. Ranked by either: number of sessions with that source port or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Destination Port Ranking*

Returns destination ports. Ranked by either: number of sessions with that destination port or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Protocol Ranking*

Returns most used protocols. Ranked by either: number of sessions with that protocol or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reporting Device Ranking*

Returns most active reporting devices. Ranked by either: number of sessions that contain events from the device or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reporting Device Type Ranking*

Returns most active reporting device types. Ranked by either: number of sessions that contain events from a device of that type or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Reported User Ranking*

Returns information about users from reporting devices such as: Windows clients, Solaris clients, etc. Ranked by either: number of sessions that contain events from a reported user or by bytes transmitted in sessions that contain events that meet the query criteria.

- *Matched Rule Ranking*

Returns top firing rules. Ranked by number of incidents.

- *Matched Incident Ranking*

Returns incidents. Ranked by either: number of sessions that contain events that meet the criteria that contributed to the incident or by bytes transmitted real time in sessions that contain events that meet the query criteria.

- *All Matching Sessions*

Returns all sessions that contain events that meet the criteria. Sessions that contain a common set of event types are grouped together. They are also sub-grouped by session source IP address and session destination IP address. Sessions in the same sub-group are ordered by time. Real Time results are available for this Result Type.

- *All Matching Events*

Returns events. Ranked by time with the most current first. Real Time results are available for this Result Type.

- *All Matching Event Raw Messages*

Returns the raw messages associated with events. Ranked by time with the most current first. Real Time results are available for this Result Type.

- *NAT Connection Report*

Returns NAT connections. Ranked by time with the most current first.

- *MAC Address Report*

Returns MAC addresses. Ranked by time with the most current first.

- *Unknown Event Report*

Returns events that are not fully processed by the MARS. In some cases, event information such as the five tuple (source IP, source port, destination IP, destination port, and protocol) might not be present, hence can not be queried in real time.

Order/Rank By

This selection determines the ranking or order of the query's results. These selections are determined by the kind of Result Format that you use when you run the query.

- *Session Count*

The number of sessions that contain events that meet the criteria that contributed to the incident.

- *Bytes Transmitted*

The number of bytes transmitted in sessions that contain events that meet the query criteria.

- *Time*

Most current results appear first.

- *Incident Count*

Largest number of incidents appear first.

Filter By Time

- *Last*

The present time minus the number of days, hours, and minutes entered.

- *Start/End*

Absolute literal time ranges defined by the date to the minute.

- *Real Time*

Streams rolling real-time results from recent past to current time. Result Formats that work in real time are: •All Matching Sessions, page 6-7, •All Matching Events, page 6-8, and •All Matching Event Raw Messages, page 6-8.

Real Time results appear in a normal browser window. Moving the scroll bar stops the "rolling" behavior. Clicking the Resume button on the bottom of the page allows the scrolling to resume.

Figure 6-10 Click the Resume Button to Start the Page Rolling

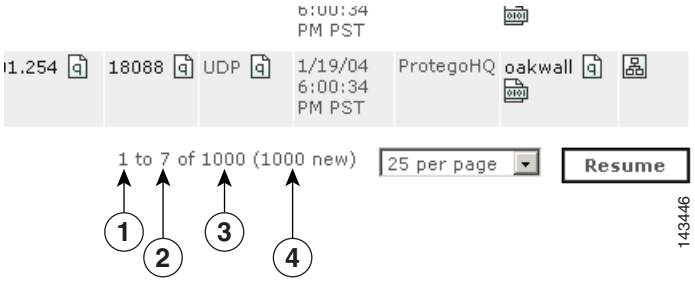

1	Top row visible	**2**	Bottom row visible
3	Total rows queried since start	**4**	Number of new queries pulled when this page last refreshed

Use Only Firing Events

Select this if you want only events that fired incidents to return information.

Maximum Number of Rows Returned

Select the number of rows that you want displayed.

Selecting Query Criteria

To Select a Criterion

Step 1 Select the criteria that you want to edit by clicking it.

Figure 6-11 Clicking any to narrow your criteria

Step 2 Move the items that you want to query from the right to the left of the filter by selecting the check box next to them, and clicking the Equal and Not Equal buttons.

Figure 6-12 Selecting Variables

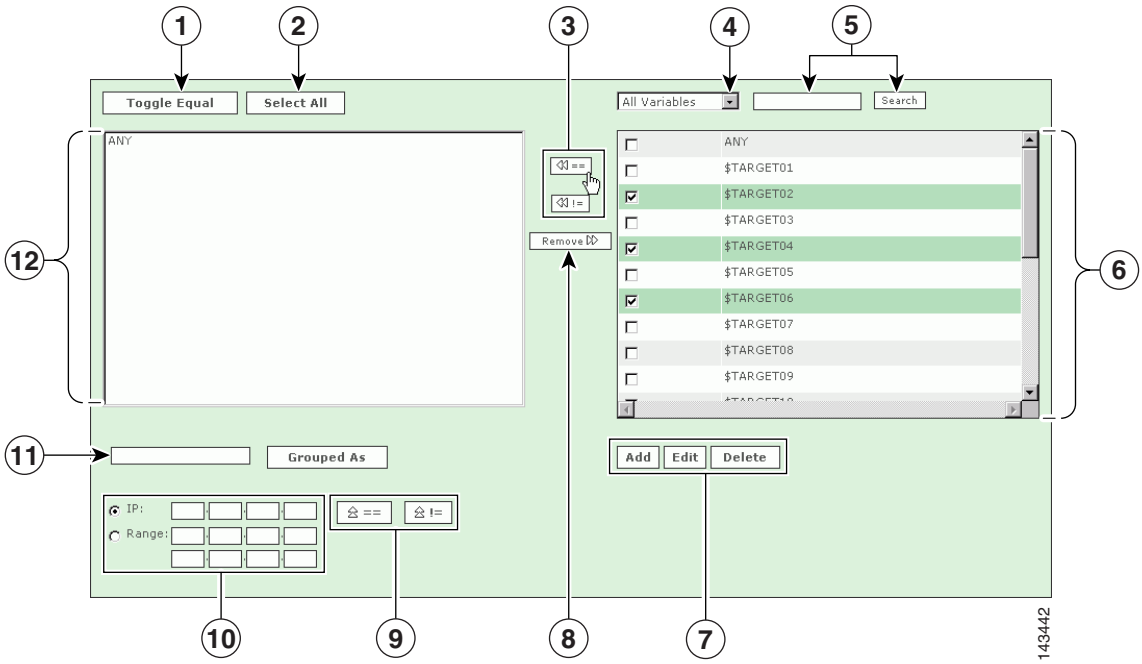

Step 3 You can select a variety of different variables, events, devices, addresses from the filter page. The following number correspond with the numbers in the preceding graphic:

1. Check the boxes next to the items in the **Sources Selected** field to select them, and click the **Toggle Equal** button to change them between equal and not equal.
2. Click the **Select All** button to select all items in the **Sources Selected** field. (Note: if you have items highlighted in the Sources Selected field, clicking **Select All** will de-select them.)
3. Use the **Equal** and **Not Equal** buttons to bring highlighted items from the **Sources Available** field into the **Sources Selected** field.
4. Filter sources from this drop-down list.
5. Enter search text, and click **Search** to move items that match the search criteria from the **Sources Available** field to the **Sources Selected** field.
6. To add a new item to the sources, click the **Add** button. To edit or delete an existing source, click the **Edit** or **Delete** button. See IP Management, page 9-3 for more information.
7. Click an item or items in the Sources Selected field, and use the **Remove** button.
8. To move IP values up into the Sources Selected field, click the **Equal** (Up) icon, or the **Not Equal** (Up) icon.
9. Check the radio button next to **IP** or **Range**, and enter an IP address or a range of IP addresses into their respective fields.
10. Select items in the Sources Selected field by clicking them. Enter a group name, and click the **Grouped As** button to group them.
11. Once you have chosen the query criteria that interests you, click **Apply** to return to the Query page.

Repeat this selection process for other query data.

Step 4 Click the **Submit** button to run the query.

Query Criteria

The following list describes the selections in the Query Event Data table.

Source IP

- *Pre NAT source addresses*

Specifies that the constraints entered are the session endpoints.

- *Post NAT source addresses*

Specifies that the constraints entered are the source as appearing at the destination.

- *ANY*

No constraint is placed on the source IP addresses.

- *Variables*

Signify any one IP address, only useful for queries in tandem with the same variable.

- *IP addresses*

IP addresses present on devices in the system or user entered dotted quads.

- *IP ranges*

The range of addresses between two dotted quads.

- *Networks*

Topologically valid networks.

- *Devices*

The hosts and reporting devices present in the system.

Destination IP

- *Post NAT destination addresses*

Specifies that the constraints entered are the session endpoints.

- *Pre NAT destination addresses*

Specifies that the constraints entered are the destination as appearing at the source.

- *ANY*

No constraint is placed on the source IP addresses.

- *Variables*

Any one IP address, only useful for queries in tandem with the same variable.

- *IP addresses*

IP addresses present on devices in the system or user entered dotted quads.

- *IP ranges*

The range of addresses between two dotted quads.

- *Networks*

Topologically valid networks.

- *Devices*

The hosts and reporting devices present in the system.

Service

- *ANY*

No constraint is placed on the source or destination ports or protocol.

- *Service variables*

Any one set of destination port and protocol, only useful for queries in tandem with the same variable.

- *Defined services*

Services on the database.

Event Types

- *ANY*

No constraint on the event type.

- *Event types*

Events that have been merged into types.

- *Event type groups*

Groups of event types.

Device

- *Devices*

The reporting devices present in the system. This restricts the query to a subset of the devices that report to the MARS.

Severity/Zone

- *ANY*

No constraint on the event type severity.

- *Green*

Low-severity events

- *Yellow*

Medium-severity events

- *Red*

High-severity events

- *Zone*

Events reported by devices in the indicated zone.

Operation

- *None*

Defines a single-line query.

- *AND*

Boolean “and” that defines a two or more line query.

- *OR*

Boolean “or” that defines a two or more line query.

- *FOLLOWED-BY*

Time conditional query (e.g.: Y must happen after X) that defines a two or more line query.

Rule

- *Empty field – Rules Chosen field*

When this field is empty, it acts like an ANY selection. No constraint is placed on the sub-set of events.

- *Rule*

Restricts the query to the sub-set of events that contributed to the incidents of the specified rules firing.

Action

- *Empty field – Empty Actions Chosen field*

When this field is empty, it acts like an ANY selection. No constraint is placed on the sub-set of events.

- *Actions*

Restricts the query to the sub-set of events that contributed to the incidents of rules that have the specified notifications as part of their actions. (See page ??? rules for more information?)

Saving the Query

You can save query criteria to re-use as reports or rules.

To save a query as a report

This takes the query that you are using and creates a report. For more information on creating reports, see Reports, page 6-19.

To save a query as a rule

This takes the query to the rules page, populating the rules with the selected query criteria. Likely, you must identify additional criteria to complete the rule. For more information on creating rules, see Rules, page 7-1.

Perform a Long-Duration Query Using a Report

This section explains how to create and view a long-duration query on the MARS. There are two ways to perform a long-duration query on the MARS:

1. **Modifying an existing report.**

Advantages:

- The report is compiled relatively quickly.
- You can compile data gathered over a longer time period

Disadvantage.

This type of query can only be used without any changes to query criteria other than time range, and can only be used with the following reports:

- Activity: All - Top Destination Ports
- Activity: All - Top Destinations
- Activity: All - Top Event Types
- Activity: All - Top Reporting Devices
- Activity: All - Top Sources
- Activity: Attacks Seen - Top Reporting Devices
- Activity: Denies - Top Destination Ports
- Activity: P2P Filesharing/Chat - Top Event Types
- Activity: Scans - Top Destination Ports
- Activity: Scans - Top Destinations
- Activity: Unknown Events - All Events
- Activity: Web Usage - Top Destinations by Sessions
- Activity: Web Usage - Top Sources
- Attacks: All - Top Rules Fired
- Attacks: All - Top Sources

2. **Performing a batch query.**

Advantages:

- You can modify any of the query criteria.
- Best suited for data that spans a short time period.

Disadvantages

- This type of query can be slow and may take a substantial amount of time to complete.
- Only Admin users can perform a batch query.

If you want to observe activity on your MARS over a long period, you can change the duration of time over an existing report that runs on a regular basis, such as hourly or daily, whether they are shipped with the MARS or created by you.

Note Trying to run a long-duration query using a report that only runs "on demand" has the same effect as running a query; it can take just as long because it has to compile data, whereas data from the regularly-run reports has been precompiled on an ongoing basis.

To query using a report, follow these steps:

Step 1 In the **QUERY / REPORTS** tab, click the **Reports** tab to obtain the Main Report window.

Figure 6-13 Main Report Window

Report Selection

	Name	Schedule	Format	Recipients	Query	Description	Status	Submitted	Time Range
○	Activity: All - NAT Connections	Run on demand only	Normal	None	Query Type: NAT connections ranked by Time Time: May 1, 2004 8:21:50 PM PDT - Jun 6, 2004 8:31:50 PM PDT	This report lists Network Address Translations performed on non-denied sessions as reported to MARS.	Finished: Jun 16, 2004 4:40:36 AM PDT	Jun 15, 2004 8:32:09 PM PDT	May 1, 2004 8:21:50 PM PDT - Jun 6, 2004 8:31:50 PM PDT
○	Activity: All - Top Destination Ports	Run on demand only	Trend	None	Query Type: Destination Ports ranked by Sessions Time: 1hh:0mm:0ss	This report ranks the UDP and TCP destination ports of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 10, 2004 4:17:02 PM PDT	Jun 10, 2004 4:16:58 PM PDT	Jun 10, 2004 3:16:58 PM PDT - Jun 10, 2004 4:16:58 PM PDT
◉	Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 17, 2004 2:15:52 PM PDT	Jun 17, 2004 2:15:52 PM PDT	Nov 30, 2003 1:05:52 PM PST - Jun 17, 2004 2:15:52 PM PDT

143798

Step 2 Navigate to and then click the radio button next to the regularly-scheduled report you want to modify (in this example, we use **Activity: All - Top Destinations**). Click the **Query** column to edit the report. The Build Report window appears.

Figure 6-14 Build Report window

Build Report
Click the cells below to define the report:

Name	Schedule	Format	Recipients	Query	Description	Status	Submitted	Time Range
Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 16, 2004 7:15:42 PM PDT	Jun 16, 2004 7:15:42 PM PDT	Nov 29, 2003 6:05:42 PM PST - Jun 16, 2004 7:15:42 PM PDT

Submit Previous Next

Time Range:

◉ Last: 200 Days 0 Hrs 10 Mins

○ Start: 2004 June 16 19 Hrs 42 Mins
End: 2004 June 16 19 Hrs 52 Mins

Submit Previous Next

143686

Step 3 In the lower portion of the Build Report window, change the **Time Range** the report (**Activity: All - Top Destinations**) covers to the duration you want it to cover.

Step 4 Click the **Submit** button to run the report and return to the Main Report window.

View a Query Result in the Report Tab

To view a query in the Report tab, follow these steps:

Figure 6-15 Main Report window (bottom)

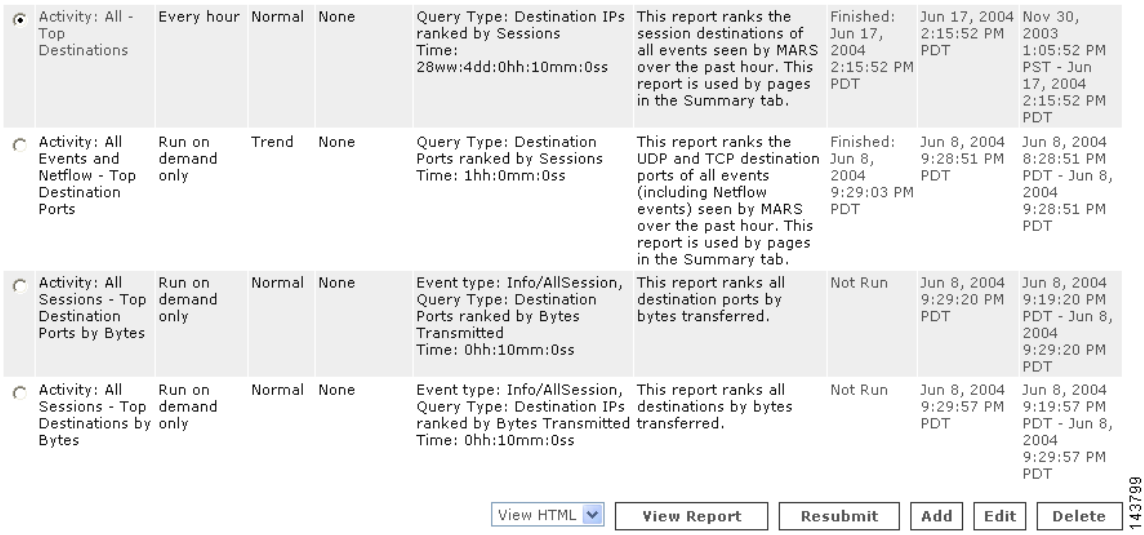

Activity: All - Top Destinations	Every hour	Normal	None	Query Type: Destination IPs ranked by Sessions Time: 28ww:4dd:0hh:10mm:0ss	This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 17, 2004 2:15:52 PM PDT	Jun 17, 2004 2:15:52 PM PDT	Nov 30, 2003 1:05:52 PM PST - Jun 17, 2004 2:15:52 PM PDT
Activity: All Events and Netflow - Top Destination Ports	Run on demand only	Trend	None	Query Type: Destination Ports ranked by Sessions Time: 1hh:0mm:0ss	This report ranks the UDP and TCP destination ports of all events (including Netflow events) seen by MARS over the past hour. This report is used by pages in the Summary tab.	Finished: Jun 8, 2004 9:29:03 PM PDT	Jun 8, 2004 9:28:51 PM PDT	Jun 8, 2004 8:28:51 PM PDT - Jun 8, 2004 9:28:51 PM PDT
Activity: All Sessions - Top Destination Ports by Bytes	Run on demand only	Normal	None	Event type: Info/AllSession, Query Type: Destination Ports ranked by Bytes Transmitted Time: 0hh:10mm:0ss	This report ranks all destination ports by bytes transferred.	Not Run	Jun 8, 2004 9:29:20 PM PDT	Jun 8, 2004 9:19:20 PM PDT - Jun 8, 2004 9:29:20 PM PDT
Activity: All Sessions - Top Destinations by Bytes	Run on demand only	Normal	None	Event type: Info/AllSession, Query Type: Destination IPs ranked by Bytes Transmitted Time: 0hh:10mm:0ss	This report ranks all destinations by bytes transferred.	Not Run	Jun 8, 2004 9:29:57 PM PDT	Jun 8, 2004 9:19:57 PM PDT - Jun 8, 2004 9:29:57 PM PDT

Step 1 At the bottom of the Main Report window, click the radio button next to the report (**Activity: All - Top Destinations**).

Step 2 From the drop-down list on the bottom of the Reports page, select either:

- **View HTML**: to view the report as an HTML file.
- **View CSV**: to view the report as a CSV (comma-separated values) file.

Step 3 Click the **View Report** button.

Note The **Status** column of the report lets you know whether the report has finished before viewing. You can view a partially-completed report, but it might not contain all the data you want to examine. You can also refresh the screen to update the **Status** column.

Perform a Batch Query

This type of long-duration query can take a long time to perform and is more suitable for a shorter duration of time.

Note Only Admin users can perform a batch query.

To perform a batch query, follow these steps:

Step 1 Click the **QUERY / REPORTS > Query** tab. The Query window appears.

Figure 6-16 Query window

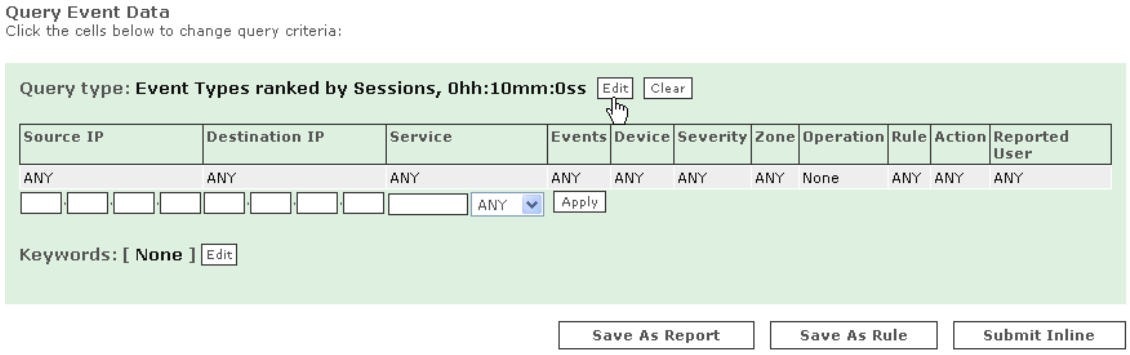

Step 2 In the Query window, click the **Edit** button to change the query criteria. The Query Event Data window appears.

Figure 6-17 Query Event Data window

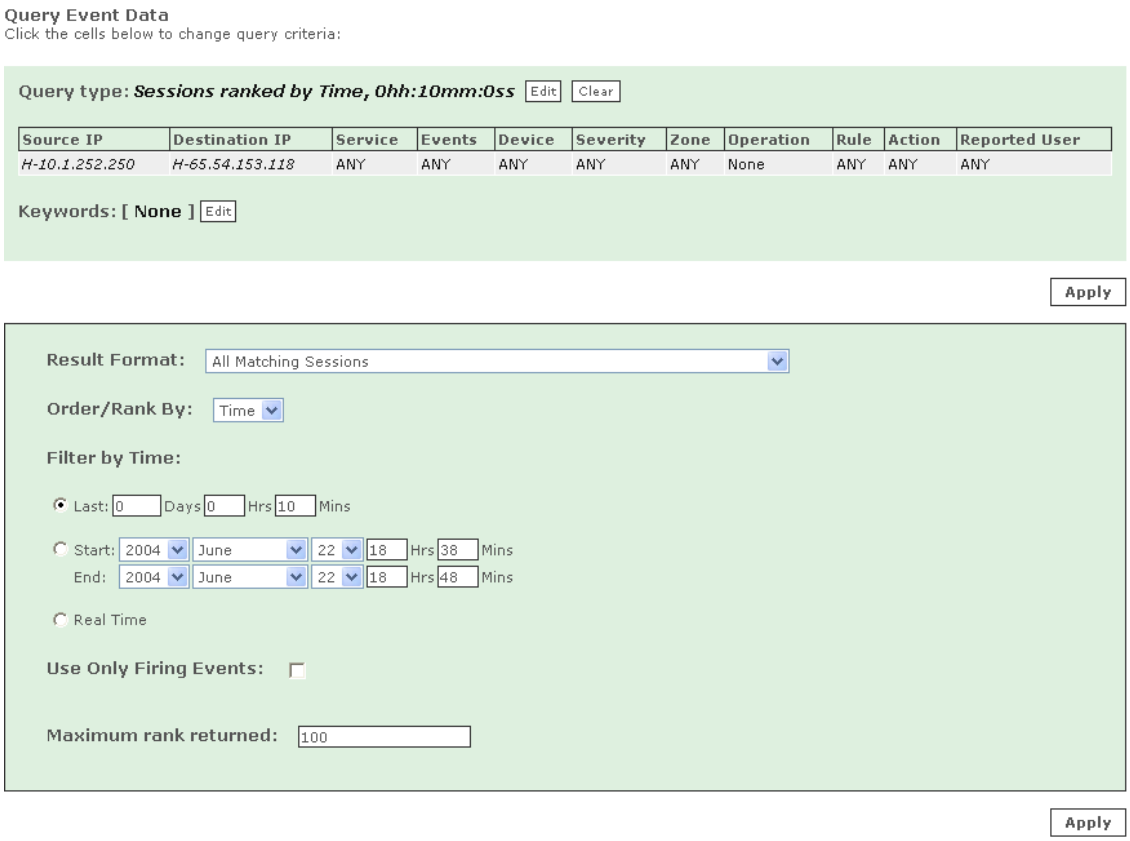

Step 3 In the Query Event Data window, you can change the query criteria. (For more information on query criteria, see Query Criteria, page 6-11). By clicking on various parameters you can change the nature of the query. In this case we are specifying a Source IP address of **10.1.1.6**, a Destination IP address range previously saved as **mygroup**, and setting the duration of the query to the past **2** days.Click either **Apply** button to apply your changes to the query. The Query Save/Submit window appears.

Figure 6-18 Query Save/Submit window

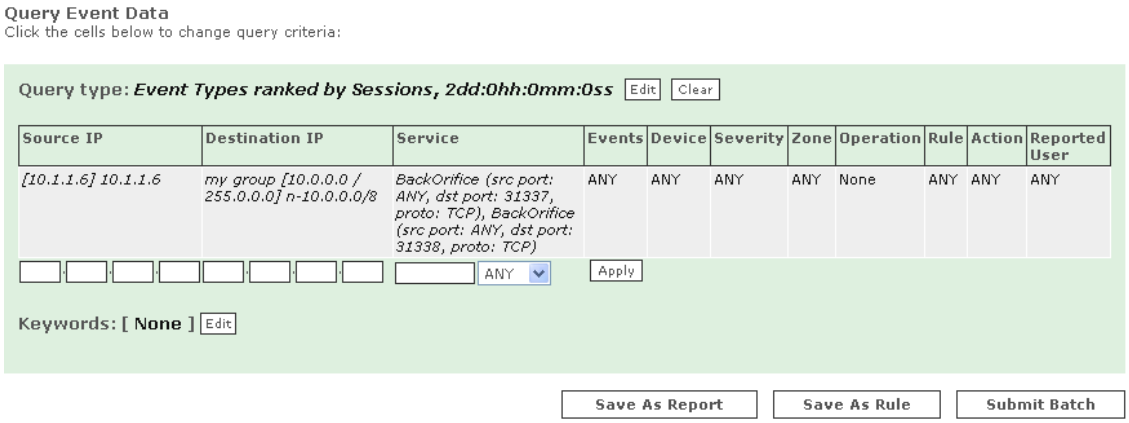

Step 4 The Query Save/Submit window asks you to choose from the options of **Save as Rule**, **Save as Report**, or **Submit Batch**. To submit your query as a batch query, click **Submit Batch**. Your query is submitted, and you are automatically taken to the Batch Query tab.

Figure 6-19 Batch Query tab

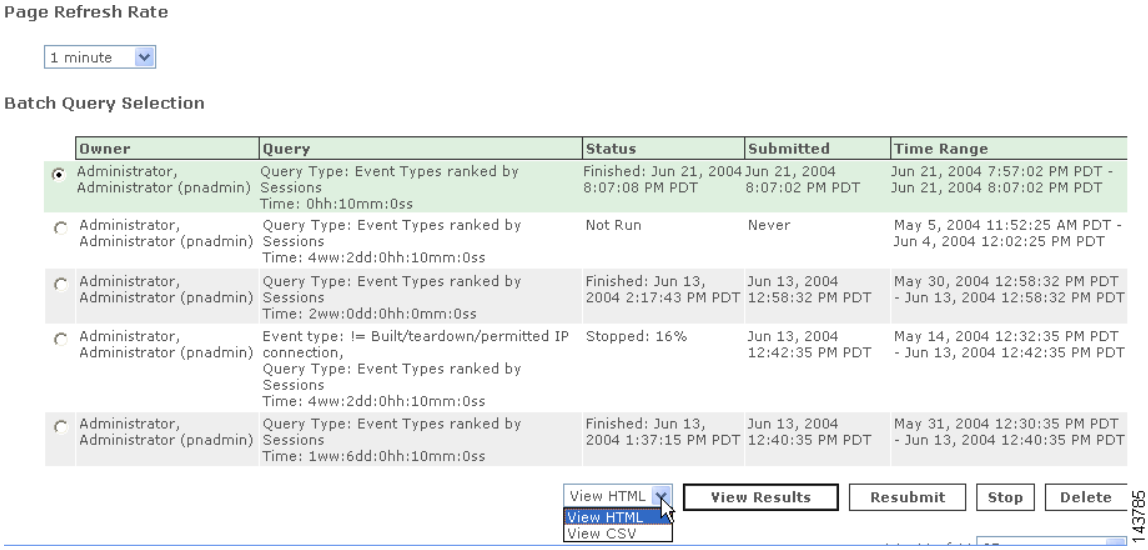

Step 5 To watch the status of the query in real-time, you can use the Batch Query tab drop-down list to change the **Page Refresh Rate** from **Never** (the default) to 1 minute, 3 minutes, 5 minutes, 10 minutes, 15 minutes, or 30 minutes.

Step 6 To view the results of the batch query as it is running, click the radio button next to your query (here it's highlighted in green) and click **View Results**. This can be done while the query is in progress.

If the email address in your user profile on the MARS is valid, the results of your batch query are emailed to you when the query has completed. You can also view the results of your batch query by clicking **QUERY / REPORTS > Batch Query > View Results**.

Note When you click **View Results** while the query is in progress, the results compiled up to that moment are recomputed. This can make the display take longer to appear than after the results are compiled.

Reports

Using the Reports page, you can build repeatable queries, edit and delete current reports, run reports, and view reports in either HTML or CSV (comma separated value) formats.

Reports performed at the Global Controller level are similar to those on an Local Controller, but also include the **Zone Collapsing** parameter. You can run a report across one or more Local Controllers by specifying their zones. This enables a report at the Global Controller to select zone-specific objects.

When you submit a report from the Global Controller, the report request is sent to the Local Controllers monitored by that Global Controller. Each Local Controller generates the report and sends summary data back to the Global Controller, which merges the results at the global level. The merged report is sent to any recipients, as defined by the report definition on the Global Controller.

When you view a report, you are viewing the last instance that ran. If you want to view an up-to-the-minute report, resubmit the report before viewing it.

Report results are purged from the database after a purge interval, as tabulated in Table 6-1.

Table 6-1 Maximum Database Retention Limits for Report Results

Cisco Security MARS Model	Maximum Number of Stored Reports[1]	Database Purge Interval[2]
CS-MARS-20-K9	1,000 ranking reports 5,000 event/session reports	3 months
CS-MARS-50-K9	1,000 ranking reports 5,000 event/session reports	3 months
CS-MARS-100-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-100E-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-200-K9	1,000 ranking reports 5,000 event/session reports	6 months
CS-MARS-GC-K9	1,000 ranking reports 5,000 event/session reports	12 months
CS-MARS-GCM-K9	1,000 ranking reports 5,000 event/session reports	12 months

1. Table values are for Cisco Security MARS Release 4.1.5. In Release 4.1.4 and prior, the maximum number of ranking reports is 100, maximum number of event/session reports is 1,000.
2. As of Cisco Security MARS Release 4.1.5. In Release 4.1.3, and 4.1.4, report results are retained for one year in the MARS database before they are automatically purged. In Releases prior to Release 4.1.3, report results are retained indefinately. The purge interval cannot be changed.

Report Type Views: Total vs. Peak vs. Recent

Where alerts provide up-to-the-minute views of high-priority incidents, reports aggregate sessions into different views. Reports correlate based on the three data points:

- Period of time
- Query criteria
- View type

The *period of time* defines boundaries around the analyzed session data based on when it was recorded. *Query criteria* restrict the set of sessions that will be aggregated to that which matches your criteria. Criteria can include source address, destination address, network service, event, reported user, and reporting device. The *view type* defines how to aggregate the matched data into a meaningful report view—one that matches the type of study in which you are interested.

Note In each view type, you can refine the report criteria to filter out expected activity—the data you know about. You can filter this activity by refining the query criteria. These criteria should be tuned to a specific network. Reports can be valuable in detecting behaviors beyond the normal traffic flows of your network. You can determine the expected activities using reports that are not filtered and vetting those results against normal network use.

MARS provides three view types, each of which restricts the matched sessions to a user-defined limit of *N*. The following view types exist:

- **Total View**. For each result type matching the query criteria, this view counts the occurrences of that result type that transpire during the specified time period. It presents the total count of the top *N* matched result types, ranked by number of sessions, as determined by which ones occurred most frequently over the period of time. You can use these reports to determine your network's condition relative to the studied sessions. For example, you can use this view to identify attacks that launched at frequent intervals. This view does not present spikes in network activity; it simply presents the top occurring result types.
- **Peak View**. Within MARS, all report result data is stored in 10-minute time slices. The Peak View studies each of the 10-minute time slices within the specified time period to which one contained the highest number of matched sessions for a specific result type. It also determines an additional nine peaks within the time period, where each peak identifies a unique result type relative to the other peaks.

Each peak value is charted relative to the other nine peaks. For each time slice containing a peak value, the Peak View lists the top *N* matched result types that occurred. It is possible to have multiple peaks within the same time slice, as it is the result type, not the time slice, that must be unique across peaks.

Note To be detected within this view, the result type must peak above normal traffic. Therefore, you must tune the query data to filter out expected traffic.

Unlike the Total View, the Peak View does not focus on the overall top occurring results, instead it identifies a high volume of traffic over a short time period. Its purpose is to detect temporary bursts of traffic on your network that overshadow normal traffic usage. These bursts identify possible issues, such as worm outbreaks.

- **Recent View**. This view is similar to Total View; however, it identifies the top *N* result types that occurred within the past hour. It then plots all occurrences of those result types over the selected time period.
- **CSV**. Generates the Total View but presents the report in the CSV format for processing by another tool or script. This option is intended for use with e-mail notifications where post-processing is required.

Creating a Report

You can create a report through the **Query** page, or you can create a report from scratch on the **Reports** page. These instructions detail creating a report from the **Reports** page, but are applicable to editing reports and to creating reports from the **Query** page.

Create a New Report

Step 1 On the Reports page, click the **Add** button.

Step 2 In the **Report Name** and **Report Description** fields, enter a report name and description. Click the **Next** button.

Step 3 Select the schedule parameters for the report.

Step 4 Select a format for the report's output. Under **View Type and Zone Collapsing**, select one of the following:

- **Total View/Sum Zones** - This view displays the summed total of the top N results over the specified time range.
- **Total View/List Zones** - This view displays the total, grouped by zone, of the top N results over the specified time range
- **Peak View/Sum Zones** - This view finds the top ten largest results in the time range, and displays the top ten results for the times when those peaks occurred.
- **Peak View/List Zones** - This view finds the top ten largest results in the time range, groups them by zone, and displays the top ten results for the times when those peaks occurred.
- **Recent View/Sum Zones** - This view finds the top N results from the past hour, and displays them versus their summed totals over the specified time range.
- **Recent View/List Zones** - This view finds the top N results from the past hour, groups them by zone, and displays them versus their summed totals over the specified time range.
- **CSV/Sum Zones** - This view displays the summed total of the top N results as a comma-separated values file. (See Report Type Views: Total vs. Peak vs. Recent, page 6-20).
- **CSV/List Zones** - This view displays the summed total of the top N results, grouped by zone, as a comma-separated values file. (See Report Type Views: Total vs. Peak vs. Recent, page 6-20).

Click **Next**.

Step 5 Select users in the Recipients Available field by expanding the user groups, clicking users or user groups, and clicking the **Add** button. See User Management, page 9-6 for more information.

Step 6 Repeat Step 5 for other users. Click **Next**.

Step 7 Build or modify the query. To edit the query time range, either click the Report type link or click the **Edit** button.

Note For more information on building queries, see Queries, page 6-1.

Step 8 Click **Apply** to save your changes; click **Next** when the query is complete.

Step 9 Click **Submit** to save your report.

Working With Existing Reports

To View a Report

Step 1 Click the radio button next to the report.

Step 2 From the drop-down list on the bottom of the page, select either:

- **View HTML**: to view the report as an HTML file.
- **View CSV**: to view the report as a CSV file.

Step 3 Click the **View Report** button.

Note If you chose to view the report as a CSV file, you need to save the file to your computer and open the CSV file in a third-party application.

To Run a Report

Step 1 Click the radio button next to the report.

Step 2 Click the **Run Now** button.

Note Due to caching issues, reports with a time range of less than one hour are not recommended.

See Table 6-1, "Maximum Database Retention Limits for Report Results" for information on how long report results are retained in the database per MARS model number.

To Delete a Report

Step 1 Click the radio button next to the report.

Step 2 Click the **Delete** button to delete the report.

Step 3 On the Delete Confirmation page, click **Delete**.

To Edit a Report

You can not edit system generated reports. Editing report criteria is meant for minor tweaking to previously generated report.

Step 1 Click the radio button next to the report.

Step 2 Click the **Edit** button to edit the report.

Step 3 Navigate using the **Previous** and **Next** buttons, or clicking on the report criteria.

Figure 6-20 Navigating to the Recipients column by clicking its criteria

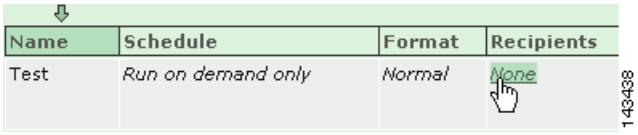

Step 4 Edit the report, and click the **Apply** button to apply changes to the report.

Step 5 Click the **Submit** button to finalize the report.

Note Changing the report's query criteria will not re-generate a new result. New edited criteria is based on the previously generated report. In some situation such as filtering out specific IP source, user should create a new report.

Note Email notification of a global generated report will be sent from the Global Controller and not the Local Controller.

CHAPTER 7

Rules

This chapter discusses MARS Inspection and Drop rules in the following sections:

- Rules Overview, page 7-1
- Constructing a Rule, page 7-4
- Working with System and User Inspection Rules, page 7-16
- Setting Alerts, page 7-20
- Rule and Report Groups, page 7-21

Rules Overview

An inspection rule is a real-time filter that detects interesting patterns of network activity. These patterns can signify attacks or false positives, and they inform you of network configuration errors and other anomalous network behavior. Rules on the Global Controller are propagated down to Local Controllers. When these rules are triggered, incidents are sent to the Global Controller.

An attack might be straightforward, or it could be a probe, an attack, and then a follow-up to the attack. Whatever the method of attack, attacks share common traits, and you can use rules to define these traits to identify and mitigate attacks.

Rules create incidents. Rules connect the information you receive from your networks' reporting devices, linking them together to form a chain of events that describes an unfolding intrusion. They classify incoming events as firing events by matching them against the rule criteria. They also determine when a false positive is either dropped completely or kept as information in the database.

A rule is either active or inactive. Active means the rule is operating and is being applied to incoming events. Inactive indicates that the rule is inoperative and not consuming CS-MARS resources. To view a list of all System Inspection rules, see Appendix D, "System Rules and Reports."

Note

Prioritizing and Identifying

Your first order of business is to prioritize your network's assets; in other words, figure out what is going to cost you the most money if it goes down. Next, identify your networks' most exploitable weaknesses. Choose which ones you are willing and able to close, and rank the remaining weaknesses by risk and exploitability.

Use this ranked list to guide your time and energy expenditures when customizing the CS-MARS rule set.

Think Like a Black Hat

Ignore for a moment the benign users who do legitimate business on your networks.

Get inside the mind of the black hat that wants to take your network down. The person who should concern you is the one with a plan.

Good plans have a sequence of steps, contingencies, and metrics to determine success or failure. The more fully you can anticipate these plans, the fewer attacks will be able to execute unhindered and unobserved. The black hat is looking for wide-open doors and easy access. Failing that, the black hat is going to look for specific and obvious exploitable weaknesses.

Planning an Attack

Start to detail your plan. You want to penetrate a network. You'd like to avoid detection and identification if possible. You want root access on a host.

How do you get root access? You do not have a preexisting account, and physical access isn't feasible. The first few options that come to mind are password guessing, password brute force, or exploiting a known weakness on the host.

You decide to exploit services running on the host, so you need to find out what it is running. To do this, you have a number of techniques: port scans, OS fingerprinting, banner probing, etc.

Once you've identified a vulnerable service or software, you can attack it with a catalogue of exploit software. Depending on what you find and your available exploits, there are a number of different effects, usually allowing you to execute arbitrary code.

You now own the host. What happens next is up to you. You have many options: you can install a root kit, you can crash the machine, etc. You have full access—you can do just about anything on to/from that host.

Back to Being the Admin

You must now express the plan in terms of information that is reported to you. This attack plan contains an attack with a follow up of some kind. You might write your plan like:

- probe
- attacker to target, buffer overflow
- attacker to target, root login (compromised host)

At this point, the black hat has compromised the host. What happens next is up to the attacker. This makes the next few steps especially hard to predict. They want to be able to manipulate the world, they want to make change. Your newly compromised host is the instrument for change. You can specify additional potential steps in the plan that make it even more urgent to take care of the situation immediately. Such as:

- target to FTP server, code download
- target to secondary target, buffer overflow

The attacker is now using your compromised host as a launching point for further attacks.

One you've mapped out the anticipated attack to watch for, you can define a monitoring plan. The following task flow outlines the tasks involved in implementing a monitoring plan:

Step 1 Ensure your reporting devices are providing all the data you need. This step involves ensuring that each device is generating logs about the events that you expect to occur as the result of the probes and attacks. Depending on the device type, this can involve several substeps, such as specify a logging level, enable logging for the specific event, and ensuring that the reporting device publishes events to the Local Controller appliance. It can also involve enabling administrative access to the reporting device from the Local Controller appliance.

Step 2 Configure CS-MARS to pull events from the reporting devices on your network. This step involves adding each reporting device to Local Controller. If the reporting device type is not directly supported, you must define a custom device type for the reporting device.

Step 3 Ensure that the event types that you need to study are accepted and processed by Local Controller. If they are not, you must define a custom log parser template for each event and a custom device template to which the custom log parser templates are associated. For device types supported by CS-MARS, this should not be necessary.

Note You cannot define a custom log parser template for a reporting device that is supported out of the box. In this case, to define log parser for an unsupported event type, you must still define a custom device type before you can define the log parser.

Step 4 Check to see if a system rule will capture the information that you want, otherwise write your own user inspection rule. Define user inspection rules that monitor for the event types and correlate those events into a structure that will help you identify the incident. You can also specify who should be notified and how if the rule fires.

Types of Rules

Inspection Rules

An inspection rule states the logic by which the CS-MARS tests whether or not a single network event or series of events is a noteworthy incident. An event or series of events with attributes that match the attributes specified in an inspection rule causes the rule to trigger (or "fire") to create an incident. Incidents may be attacks, network configuration errors, false positives, or just anomalous network activity. The over 100 inspection rules that ship with MARS are called System Inspection Rules. The number and structure of system rules are updated in signature upgrades and with more recent software releases. Both types of upgrades are performed from the Admin > System Maintenance > Upgrade page.

You can create custom inspection rules by editing or duplicating system inspection rules, by adding your own from the Inspection Rules page, or by using the Query interface. Customized inspection rules are called User Inspection Rules and are displayed on the Inspection Rules page.

Inspection rules can be created on both the Global Controller and the Local Controllers. Rules on the Global Controller are propagated down to Local Controllers. When these rules are triggered, incidents are sent to the Global Controller.

Global User Inspection Rules

Global Inspection Rules are inspection rules you create on a Global Controller then push to the Local Controller. From the Local Controller, you can edit only the Source IP Address, Destination IP Address, and Action fields of a Global Inspection Rule. To change the arguments of the other fields, you must edit the rule on the Global Controller. When you edit a global inspection rule on the Local Controller then edit it again on the Global Controller, the Global Controller version overwrites the Local Controller version. Global Inspection rule names are displayed with the prefix "Global Rule."

Rules on the Global Controller are propagated down to Local Controllers. When these rules are triggered, incidents are sent to the Global Controller.

Drop Rules

Drop rules allow false positive tuning on a MARS, and are defined only on the Local Controller Drop Rules page. They allow you to refine the inspected event stream by specifying events and streams to be ignored and whether those data should be stored in the database or discarded entirely. Drop rules are applied to events as they come in from a reporting device, after they have been parsed and before they have been sessionized. Events that match active drop rules are not used to construct incidents. Because the Global Controller does not receive events from reporting devices, rather it receives them from Local Controllers, you cannot define drop rules for the Global Controller.

Note For releases 4.2.3 and earlier of MARS, you cannot define drop rules for a NetFlow-based event. For these releases, tuning of NetFlow events must be performed on the reporting device.

Constructing a Rule

Each step of your plan corresponds to a line of a rule. Each line identifies a set of conditions. A rule can have a single line, two lines, or multiple lines. You link these lines together using the logical operators, "AND, OR, FOLLOWED-BY (in time)."

For more information on the conditions and operators found in a rule, see Table 7-1 on page 7-6.

The first step of the example plan, identified in Back to Being the Admin, page 7-2, involved probing the target host. You can express a probe by selecting the appropriate event type groups as the line's event type criteria. Also, you want to use dollar variables ($TARGET)[1] to constrain your host to ensure that

For more information on the conditions and operators found in a rule, see Table 7-1.

The first step of the example plan, identified in the section Back to Being the Admin, page 7-2, involved probing the target host. You can express a probe by selecting the appropriate event type groups as the line's event type criteria. Also, you want to use dollar variables ($TARGET)[2] to constrain your host to ensure that the probe and attacks that are reported have happened to the same host. Then you need to figure out the logical step for the next line. In this case, the probe could be optional depending on the time frame that the probe was sent and its subtlety.

1. A variable, such as ($TARGET), serves two purposes in the rule: 1.) It captures the number of times the same cell value is matched upon—the count for that cell, e.g., ten login failures from the same source address. 2.) It correlates the same value of a cell across rule lines, e.g., a probe from a source address AND an attack from that same source address.

2. A variable, such as ($TARGET), serves two purposes in the rule: 1.) It captures the number of times the same cell value is matched upon—the count for that cell, e.g., ten login failures from the same source address. 2.) It correlates the same value of a cell across rule lines, e.g., a probe from a source address AND an attack from that same source address.

Rule logic is simple. You have a row. Every row has cells. 'The logical expressions connecting different cells are “and,” while the expressions connecting items inside a cell are either “or” or “and not”, depending which clause is chosen—the equal to or not equal to.

By studying the system inspection rules, you can identify three commonly used rules: attempts, success likely, and failures, The most common rule structure is the basic three-line rule that identifies an attempted attack. It is expressed as:

```
(Probe AND
Attack) OR
Attack)
```

Note To clarify this pseudocode, keep in mind that uppercase AND, OR and FOLLOWED-BY identify a logical operator between two rule lines. Lowercase “and” identifies a logical operator between two cells. Lowercase “or” and “and not” identify a logical operator between two items within a cell.

Success likely rules extend the attempt rules by identifying suspicious activities originating from the attacked host. The general structure of these rules is:

```
((Probe AND
Attack) OR
Attack)) FOLLOWED BY
(Suspicious Activity[1]..Suspicious Activity[n])
```

Failures identify an event from a reporting device that the device classifies as a failure. Often, these rules simply match to known syslog or SNMP messages indicating some failure on the device. You can define alerts to keep you abreast of device failures. These rules follow one of two general structures: a one line failure—

```
Failure
```

—or multi-line failures separated by the *OR* operator—

`1..`*N* `Failure OR`

Failure

In the HTML interface, system rules are displayed in rows and columns. The row number is called the Offset. A rule can have more than one row (or offset), as shown in Figure 7-1.

Figure 7-1 Rule with Multiple Offsets

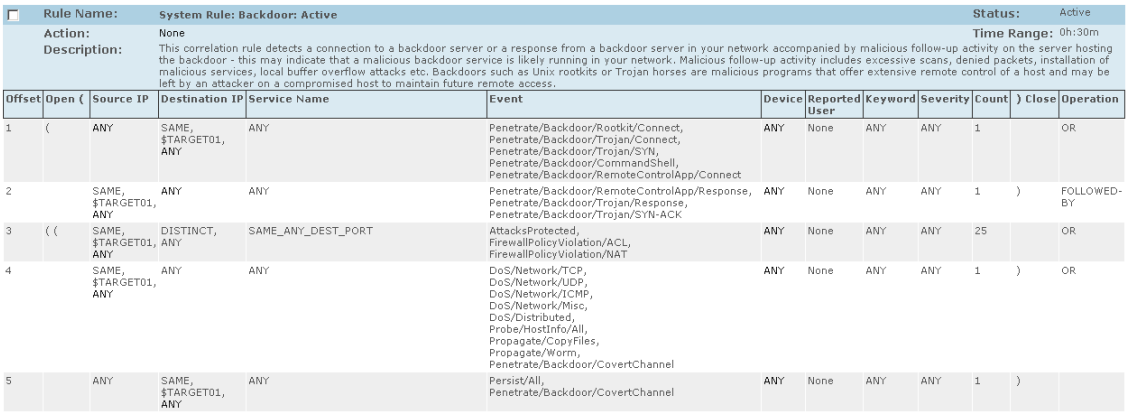

Rule Name: System Rule: Backdoor: Active — Status: Active
Action: None — Time Range: 0h:30m
Description: This correlation rule detects a connection to a backdoor server or a response from a backdoor server in your network accompanied by malicious follow-up activity on the server hosting the backdoor - this may indicate that a malicious backdoor service is likely running in your network. Malicious follow-up activity includes excessive scans, denied packets, installation of malicious services, local buffer overflow attacks etc. Backdoors such as Unix rootkits or Trojan horses are malicious programs that offer extensive remote control of a host and may be left by an attacker on a compromised host to maintain future remote access.

Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Reported User	Keyword	Severity	Count	) Close	Operation
1	(	ANY	SAME, $TARGET01, ANY	ANY	Penetrate/Backdoor/Rootkit/Connect, Penetrate/Backdoor/Trojan/Connect, Penetrate/Backdoor/Trojan/SYN, Penetrate/Backdoor/CommandShell, Penetrate/Backdoor/RemoteControlApp/Connect	ANY	None	ANY	ANY	1		OR
2		SAME, $TARGET01, ANY	ANY	ANY	Penetrate/Backdoor/RemoteControlApp/Response, Penetrate/Backdoor/Trojan/Response, Penetrate/Backdoor/Trojan/SYN-ACK	ANY	None	ANY	ANY	1	)	FOLLOWED-BY
3	((	SAME, $TARGET01, ANY	DISTINCT, ANY	SAME_ANY_DEST_PORT	AttacksProtected, FirewallPolicyViolation/ACL, FirewallPolicyViolation/NAT	ANY	None	ANY	ANY	25		OR
4		SAME, $TARGET01, ANY	ANY	ANY	DoS/Network/TCP, DoS/Network/UDP, DoS/Network/ICMP, DoS/Network/Misc, DoS/Distributed, Probe/HostInfo/All, Propagate/CopyFiles, Propagate/Worm, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)	OR
5		ANY	SAME, $TARGET01, ANY	ANY	Persist/All, Penetrate/Backdoor/CovertChannel	ANY	None	ANY	ANY	1	)	

143411

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Offset	The row number.	
Open (	Identifies the open of a clause. Clauses are used to compare one or more compound conditions in a rule.	Displays the open braces you create a clauses.
Source IP	IP address of the packet originator.	
	Variables	*ANY*—(Default). Signifies that the IP address for each count is any IP address. *SAME*—Signifies that the IP address for each count is the same IP address. This variable is local to its offset. *DISTINCT*— Signifies that the IP address for each count is a unique IP address. This variable is local to its offset. *$Target01 to $Target20*—The same variable in another field or offset signifies that the IP address for each count is the same IP address.
	Network Groups	*Defined network groups*—Topologically valid network groups as defined under Management > IP Management.
	Networks	Topologically valid network groups as defined under Management > IP Management.
	Devices	The hosts and reporting devices present in the system.
	IP addresses	IP addresses present on devices in the system or user entered dotted quads.
	IP ranges	The range of addresses between two dotted quads.

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Destination IP	IP address of the packet destination. Often referred to as the target.	
	Variables	*ANY*—(Default). Signifies that the IP address for each count is any IP address. *SAME*—Signifies that the IP address for each count is the same IP address. This variable is local to its offset. *DISTINCT*—Signifies that the IP address for each count is a unique IP address. This variable is local to its offset. *$Target01 to $Target20*—The same variable in another field or offset signifies that the IP address for each count is the same IP address.
	Network Groups—	*Defined network groups*—Topologically valid network groups as defined under Management > IP Management.
	Networks—	Topologically valid network groups as defined under Management > IP Management.
	Devices—The hosts and reporting devices present in the system.	The hosts and reporting devices present in the system.
	IP addresses—	IP addresses present on devices in the system or user entered dotted quads.
	IP ranges—The range of addresses between two dotted quads.	The range of addresses between two dotted quads.
Service Name	A TCP/IP-based network service, identified by protocol and port, defined within the packet.	

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
	Variables	**ANY**—(Default) No constraint is placed on the source or destination ports or protocol or port. SAME type variables signify that the specified destination port, source port and protocol are the same for each count. These variables are local to the offset. • SAME_ANY_DEST_PORT SAME_TCP_DEST_PORT SAME_UDP_DEST_PORT • SAME_ANY_SRC_PORT SAME_TCP_SRC_PORT SAME_UDP_SRC_PORT DISTINCT type variables signify that the specified destination port, source port and protocol are unique for each count. These variables are local to the offset. • DISTINCT_ANY_DEST_PORT DISTINCT_TCP_DEST_PORT DISTINCT_UDP_DEST_PORT Identical variables in different fields or offsets signify that the specified port and protocol for each count are identical to each other. • $ANY_BOTH_PORT5 • $ANY_DEST_PORT1 to ANY_DEST_PORT5 • $ANY_SRC_PORT1 • $TCP_BOTH_PORT1, $TCP_BOTH_PORT2 • $TCP_DEST_PORT1 to $TCP_DEST_PORT5 • $TCP_SRC_PORT1, $TCP_SRC_PORT2 • $UDP_BOTH_PORT1, $UDP_BOTH_PORT2 • $UDP_DEST_PORT1 to $UDP_DEST_PORT5 • $UDP_SRC_PORT1, $UDP_SRC_PORT2

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
	Defined services—One or more services defined under Management > Service Management.	
	Service groups—One or more service groups defined under Management > Service Management.	• Backdoor • Instant Messaging • Mail Retrieval • Online Game • P2P • Recent Backdoor • TCP-highport • UDP-highport • vulnerable-protocols
Event	Identifies one or more event types. An event type indicates some type of network activity or condition. Sometimes, events reported from different devices and different device types identify the same activity or condition, and therefore, they map to the same event type within MARS. Event types are sorted into event groups, such as "Probe/PortSweep/Stealth", to catch any of the network conditions identified by the group.	
	Variables—Signify any single event type defined under Management > Event Management, only useful for lines in tandem with the same variable.	• **ANY**—Any of the active event types can match this rule. • SAME • DISTINCT • $EVENT_TYPE01, $EVENT_TYPE10
	Event types—Events that have been merged into types.	• ANY • SAME • DISTINCT • All events
	Event type groups—Groups of event types.	• ANY • SAME • DISTINCT
	Red Severity Event Types—Displays all severe event types	
	Yellow Severity Event Types—Displays all yellow event types	

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
	Green Severity Event Types—Displays all green event types	
Device	The value of this condition can be one of the following:	
	Variables—Signify any single device defined under Admin > System Management > Security and Monitor Devices, only useful for lines in tandem with the same variable.	• **ANY**—(Default) Specifies that this rule is applied to events generated by any of the reporting devices defined in MARS. • SAME • DISTINCT • **Unknown Reporting Device**—Specifies that this rule is applied to events generated by any reporting device that is not defined in MARS. • $DEVICE01 to $DEVICE10
	• **Reporting Devices**—Identifies one or more hosts or reporting devices for which events are inspected. Valid values are one or more devices as defined under Admin > System Setup > Security and Monitor Devices.	
	Defined Device Types—	
Reported User	Identifies the active user on the host when this event was recorded. Not all events include this data. The value of this condition can be one of the following:	• **ANY**—No constraint is placed on the reported user. • **NONE**—(Default) Specifies that this condition should not be used to match this rule. • **Variables**—Signify any single user, only useful for lines in tandem with the same variable. • **Invalid User Name**—Specifies that this condition is met when the user name reported is invalid.

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Severity	The value of this condition can be one of the following:	• **ANY**—(Default) Specifies that this rule is applied to events of all severity levels. • **Green**—Restricts this rule to firing against low-severity events. • **Yellow**—Restricts this rule to firing against medium-severity events. • **Red**—Restricts this rule to firing against high-severity events.
Count	Identifies the number of items the event must occur before the condition is met. The value for this condition is a whole number ranging between 1 and 100. The default value is 1. **Note** Events of the same event type occurring in the same session in a three-second period increment the active count by one. This inherent threshold ensures that a event floods of the same type does not increase the active count arbitrarily and incorrectly fire the rule.	*Example usage*: When a backdoor rootkit install is detected, the count should be 1 as it is only going to be reported once and it is not something you expect to ever see on your network. However, if you are using deny messages to detect infected hosts, you may want the count value to be higher. For example, you may want to allow for several common mistakes, such as password failures, before firing a rule for the event. People accidentally mistype passwords, they don't accidentally install a rootkit.
Close	Identifies the close of a clause.	

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Operation	The value of this field can be one of the following:	• **None**—(Default) Defines a single-line rule or a simple condition. • **AND**—A boolean "and" used to construct a compound condition (two or more lines). This line and the next line must both be satisfied before the compound condition is met. • **OR**—A boolean "or" used to construct a compound condition (two or more lines). Either this line or the next line can be satisfied to meet the compound condition. • **FOLLOWED-BY**—Identifies a compound condition (two or more lines). specifically a sequential order of occurrence. Also referred to as a time conditional rule (e.g., Y must happen after X).The condition of this line must be met, and then the condition of the next line must be met before the compound condition is met.

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Time Range	Identifies the period of time over which the count value is augmented. For rules that have a Count value greater than one, the Time Range value determines how long the period should be before the count value is reset. For example, you can assume that if no more than three login attempts have occurred over a 10-minute period that counter can be reset.	Usage Guideline: The Time Range value combined with the Count value can affect the operation of your MARS. Each time an event is captured that satisfied a unique instance of an inspection rule, a monitoring session is constructed to track possible future occurrences until either the Count value is reached or the time period expires.

Table 7-1 Rule Fields and Arguments

Rule Field	Field Description and Arguments	Argument Descriptions
Action	Identifies the action that MARS will take when the rule is fired. Actions are user-defined alerts that include an action name and description, which also doubles at the message text provided in the alert. Each action can combine alert techniques, such as email and syslog. Each alert technique can have multiple values. For example, an action can generate two emails, a page, and a SNMP trap. Each rule can have multiple such actions. Alerts can be constructed using one or more of the following techniques: **Note** You will see the column Action/Operation. In this case, you can select either one of the following actions or one of the operators.	• **NONE**—(Default) This action states that no further action will be taken. When NONE value is selected, the firing of the rule causes an event record to be created and stored in MARS. Regardless of the selected action, this record is always created. • **Email**—Identifies the list of administrators to whom an alert should be sent. An e-mail address must be defined for the selected administrators. • **Syslog**—Identifies the list of hosts to whom an alert should be sent. You can select any number of devices to which you want a syslog message sent. • **Page**—Identifies the list of administrators to whom an alert should be sent. The message format is text. A pager number must be defined for the selected administrators. • **SNMP**—Lists the hosts to which a Simple Network Management Protocol (SNMP) alert can be sent. • **SMS**—List of users to receive notification by Short Message Service (SMS). The message can be up to 160 characters. An SMS number must be ten numbers and a domain name, for example, 1234567890@provider.com. • **Distributed Threat Mitigation (DTM)**— Lists the Cisco IOS Intrusion Prevention System (IPS) devices to which an IPS alert action can be sent (alarm, alarm and drop, or alarm and reset if it is a TCP session.) See the Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS, page 1 document for DTM configuration information.

Working Examples

The examples in this section demonstrate the use of variables, in particular, how to use variables to detect Deny patterns.

Note We recommend that you study the system inspection rules for more complex examples. To view a list of system rule names and descriptions, see Appendix D, "System Rules and Reports."

Note For a single offset rule, the variables SAME and SAME_ANY_DEST_PORT can be substituted in any of the examples for $TARGET01 and $ANY_DEST_PORT1, respectively. The "ANY" in $ANY_DEST_PORT1 means either UDP or TCP protocol.

Example A: Excessive Denies to a Particular Port on the Same Host

Figure 7-2 Rule for Excessive Denies to a Particular Port on the Same Host

Rule Name:	Example A							Status:	Active
Action:								Time Range:	0hh:0mm:10ss
Description:	Excessive denies to a particular port on the same host.								
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation
1		ANY	$TARGET01	$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	100	Training		

143408

In this example, the rule fires when 100 of the specified events occur from any source IP address to the same destination IP address, and the destination port numbers are identical.

Example B: Same Source Causing Excessive Denies on a Particular Port

Figure 7-3 Rule for Same Source Doing Excessive Denies on a Particular Port

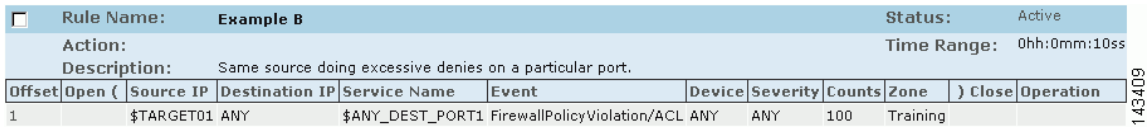

Rule Name:	Example B							Status:	Active
Action:								Time Range:	0hh:0mm:10ss
Description:	Same source doing excessive denies on a particular port.								
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation
1		$TARGET01	ANY	$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	100	Training		

143409

In this example, the rule fires when 100 of the specified events occur that have the source IP address, any Destination IP address, and identical destination port numbers.

Example C: Same Host, Same Destination, Same Port Denied

Figure 7-4 Rule for Same Host, Destination, Same Port Denied

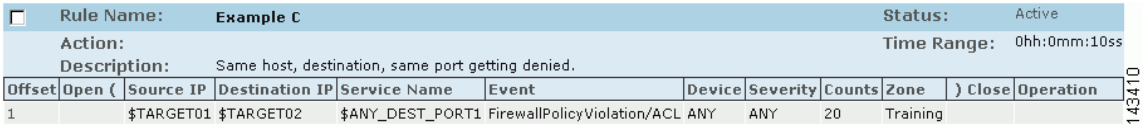

Rule Name:	Example C							Status:	Active
Action:								Time Range:	0hh:0mm:10ss
Description:	Same host, destination, same port getting denied.								
Offset	Open (	Source IP	Destination IP	Service Name	Event	Device	Severity	Counts	Zone	) Close	Operation
1		$TARGET01	$TARGET02	$ANY_DEST_PORT1	FirewallPolicyViolation/ACL	ANY	ANY	20	Training		

143410

In this example, the rule fires when 20 of the specified events occur that have the same source and destination addresses, and identical destination port numbers.

Working with System and User Inspection Rules

Navigate to the **Inspection Rules** page by clicking the **Rules** tab.

You can perform the following actions with Inspection Rules:

- Change the Source IP, Destination IP and Device fields of a System Inspection rule
- Duplicate any Inspection Rule then edit the fields to make a new User Inspection Rule
- Build a new User Inspection Rule with the Rule wizard
- Edit any field of a User Inspection Rule
- Make any rule active or inactive
- Edit, delete, or add, a Rule Group

Note When you add or edit a rule, you must click **Activate** to enable the changes.

Note Upgrade the MARS software regularly to obtain new and updated System Inspection rules. For more information, see the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*. To view a list of System Inspection rules, see Appendix D, "System Rules and Reports."

Change Rule Status—Active and Inactive

The CS-MARS correlation engine continuously tests only active rule criteria against incoming events to identify incidents. Inactive rules do not consume resources used for realtime operations.

To change the status of a rule, follow these steps:

Step 1 Navigate to the **Rules** > **Inspection Rules** page.

Step 2 Select the checkbox of the rule (or rules) to change.

Step 3 Click **Change Status**.
The selected rules are made inactive if active, and active if inactive and displayed on a different page.

Step 4 To display inactive rules, select **Inactive** from the View dropdown list. To display active rules, select **Active**.

Duplicate a Rule

Duplicating a rule creates a new rule that is a copy of an existing system or user inspection rule. You can edit all of the fields of a duplicate rule, but only the Source IP, Destination IP, and Device fields of a system inspection rule. The original rule is left unchanged after duplication.

Note You cannot delete a rule after it is created by **Duplicate** or **Add**.

To duplicate a rule, follow these steps:

Step 1 Select the checkbox of the rule to duplicate.

Step 2 Click **Duplicate**.
The name of duplicated rule is the name of the original rule extended with a timestamp of when the original was duplicated (for example, System Rule: Client Exploit - Sasser Worm Copied: 05.10.05/16:54:21). The name can be changed by editing the duplicate rule.

Edit a Rule

You can edit rules with inline editing, or with the rule wizard. To edit inline, you click the argument to edit. The rule wizard is invoked by selecting a rule to edit then clicking **Edit**. The rule wizard begins with the Rule Name field and progress through each subsequent field.

Note You only edit the Source IP, Destination IP, and Device fields of a system inspection rule. See Duplicate a Rule, page 7-16 for further information on modifying system inspection rules.

Edit a Rule with Inline Editing

You can perform inline editing to rules from the Incidents Detail page, or from the Inspections Rules page. To edit a rule with the Inline Editing, follow these steps:

Step 1 Click the Rule argument that you want to edit.
The edit page for the selected field appears.

Step 2 Change the argument, then click **Apply**.

Step 3 Repeat Step 1 as required.

Step 4 Add Open and Close parentheses as required then click **Submit**.
If no parentheses are required, just click **Submit**.

Step 5 Click **Activate** to include the rule in event correlation processing.

Edit a Rule with the Rule Wizard

The Rule Wizard can only be invoked from the Inspections Rule page.

To edit a rule with the Rule Wizard, follow these steps:

Step 1 Select the check box of the rule to edit.

Step 2 Click **Edit**.
The rule wizard page appears for the Rule Name field.

Step 3 Do one of the following actions:

- Change the argument of the field, then click **Apply**. Proceed to Step 6.
- Change the argument, then click **Next** to proceed to the next field.
- Click **Next** to proceed to the next field without changing the argument.
- Click **Previous** to go back to the previous field.
  Previous does not appear for the Rule Name page.

Step 4 Repeat • as required.

Step 5 Click **Apply** after making all edits.

Tip To skip to the end, click the Count argument, after which, only the **Action**, and **Time Range** fields must be reviewed.

Step 6 Add Open and Close parentheses as required then click **Submit**.
If no parentheses are required, just click **Submit**.

Step 7 Click **Activate** to include the rule in event correlation processing.

Note When you edit a rule on the Global Controller, the Local Controller receives only information pertinent to that Local Controller from the Global Controller. For example, if an edited Global Inspection rule is triggered only by a device that does not report to a specific Local Controller, the rule changes are not propagated to that Local Controller.

Add an Inspection Rule

Note Rules that you add are called User Inspection Rules.

Step 1 Navigate to the Inspection Rules page.

Step 2 Click **Add**.

Step 3 Enter a name and description for the rule, then click **Next**.

Step 4 Select Source IP address.

Figure 7-5 User Inspection Rule Wizard Form

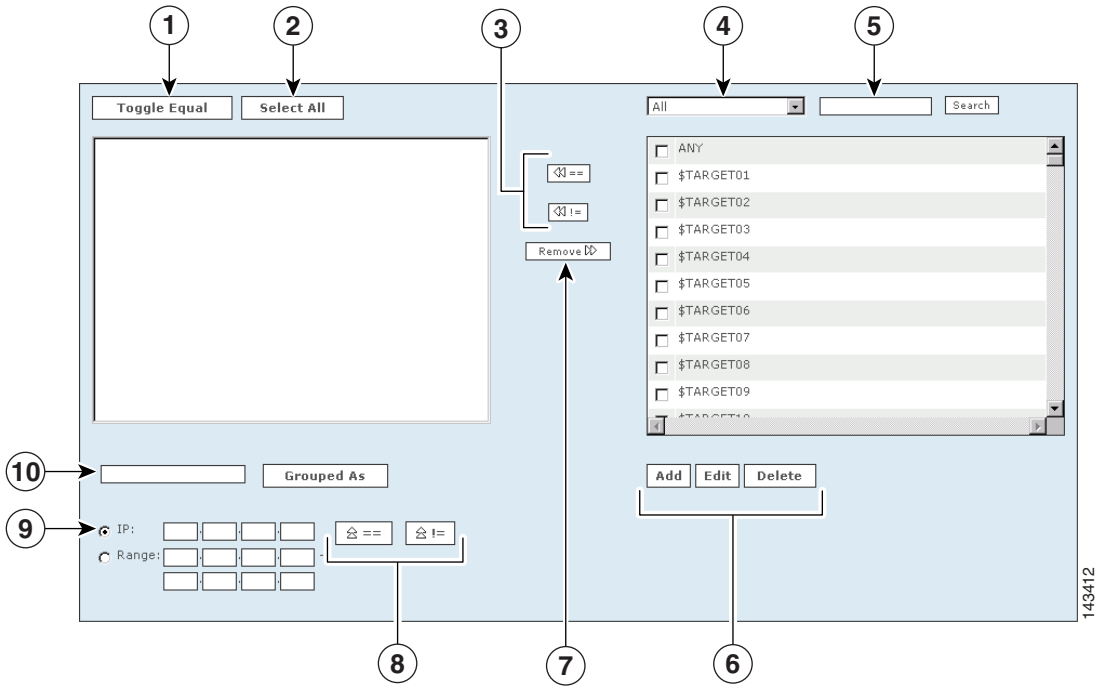

The following numbers correspond to the numbers shown in Figure 7-5.

1. Check the boxes next to the items in the **Sources Selected** field to select them, and click the **Toggle Equal** button to change them between equal and not equal.
2. Click the **Select All** button to select all items in the **Sources Selected** field. Items selected in the Sources Selected field are deselected when you click **Select All**.
3. Use the **Equal** and **Not Equal** buttons to bring highlighted items from the **Sources Available** field into the **Sources Selected** field.
4. Filter sources from this drop-down list.
5. Enter search text, and click **Enter** to move items that match the search criteria from the **Sources Available** field to the **Sources Selected** field.
6. To add a new item to the sources, click the **Add** button. To edit or delete an existing source, click the **Edit** or **Delete** button.
7. Click an item or items in the Sources Selected field, and use the **Remove** button.
8. To move IP values up into the Sources Selected field, click the **Equal** == up icon, or the **Not Equal** != up icon.
9. Check the radio button next to **IP** or **Range**, and enter an IP address or a range of IP addresses into their respective fields.
10. Select items in the Sources Selected field by clicking them. Enter a group name, and click the **Grouped As** button to group them.

Step 5 Follow the wizard, and select the values for the rule, clicking the **Next** button to progress to the next step.

Step 6 When you are asked, "Are you done defining the rule conditions," you can:

- Click the **Yes** button for a single line rule. Continue to add repetition requirements (counts), alert information, and valid time ranges for each line.

- Click the **No** button, to create a multi-line rule that uses an operator (OR, AND, or FOLLOWED BY). Return to Step 4 and continue to make your selections. Continue to add rule information, and click **Submit** when finished.
- Click the **Submit** button when finished.

Step 7 When the rule is complete, you need to activate it by clicking the **Activate** button.

Figure 7-6 Clicking the Activate button

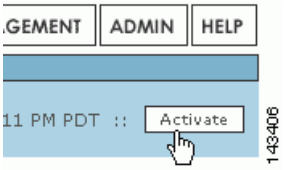

Note If you are creating or editing several rules, it is better for the system to click the **Activate** button for several changes rather than for each individual change.

Note For releases 4.2.3 and earlier of MARS, you cannot define drop rules for a NetFlow-based event. For these releases, tuning of NetFlow events must be performed on the reporting device.

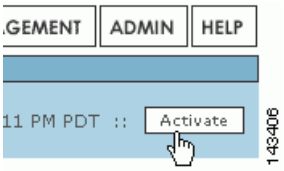

Setting Alerts

You have two options for learning about rules that have fired: you can log in and view the appropriate pages in the HTML interface or you can have MARS send alerts to external devices and users. Actions provide instructions to MARS on the second method.

Using Rules, you can alert a person if a rule has fired. The roles and groups you can choose are determined by the information you have entered in User Management. for more information on adding users into the Global Controller.

Configure an Alert for an Existing Rule

Step 1 Click on a rule argument.

Step 2 Click **Next** until the Action/Operation column is selected.

Step 3 Click the **Add** button to add users for an alert.

Step 4 Enter a **Name** and **Description** for the notification.

Step 5 Check the box next to the type of notification that you want to send. Your choices are:

- **Email** – select the roles or groups that you want to receive an email.
- **Syslog** – select the systems that you want to receive the syslogs.
- **Page** – select the roles or groups that you want to receive an electronic page on their pagers or cellular telephones.
- **SNMP** – select the systems that you want to receive the SNMP trap information.

Note For SNMP and Syslog, you need to configure the receiving systems for this feature to work.

Step 6 Click the **Change Recipient** button to add or edit recipients for alerts for that notification type (email, syslog, page, or SNMP).

Step 7 Check the box next to the role, group, or system that you want to receive alerts.

- Click the **Add** button to select recipients (to move them into the left field.)
- To remove recipients, click their names to highlight them (in the left field) and click the **Remove** button.

Step 8 Repeat steps 5 - 7 for all the alert selections that you want to include.

Step 9 Click the **Submit** button.

Step 10 Click the **Apply** button.

Note If a user adds an alert to a 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 and not the Global Controller

Rule and Report Groups

This section contains the following subsections:

- Rule and Report Group Overview, page 7-22
- Global Controller and Local Controller Restrictions for Rule and Report Groups, page 7-23
- Add, Modify, and Delete a Rule Group, page 7-23
- Add, Modify, and Delete a Report Group, page 7-26
- Display Incidents Related to a Rule Group, page 7-28
- Create Query Criteria with Report Groups, page 7-29
- Using Rule Groups in Query Criteria, page 7-30

Rule and Report Group Overview

Note To view a list of all System Inspection rules and reports, see Appendix D, "System Rules and Reports."

Rule and report groups help you manage rules and reports by speeding access to those rules and reports relevant to your task at hand. You can create groups, or use the groups provided with CS-MARS (System groups). Groups act as filters to limit the display of rules, reports, and incidents in the CS-MARS HTML interface. All groups can be modified or deleted.

CS-MARS provides over 100 system rules and 150 system reports. More can be added by creating custom rules and reports, and by performing periodic software updates. A rule or report group contains a subset of these rules or reports as members. Usually rules or reports within the same group have related functions (such as, reconnaissance activities, server attack, etc.). When you select a group from a dropdown filter, only those rules and reports that are members are displayed on the page. When you select a rule group on the Incidents page, only those incidents related to the rules of the selected group display. Report and rule groups can also be used when constructing queries.

For instance, there are at least 16 system rules that detect suspicious network access events and incidents, and 15 system reports to report this information. CS-MARS provides a system rule group and a system report group named "Access" that can filter the Inspection Rules, Incidents, and Report pages to display only those rules and reports related to monitoring access event (such as password attacks), thereby eliminating the need to search for the pertinent rules and reports within the complete rule and report pages or dropdown lists. CS-MARS provides system rule and report groups as listed in Table 7-2.

Table 7-2 Predefined Rule and Report Groups

System Report Groups	Corresponding System Rule Groups
System: Access	System: Access
System: All Events - Aggregate View	—
System: All Exploits - Aggregate View	—
System: COBIT DS3.3 - Monitoring and Reporting	—
System: COBIT DS5.10: Security Violations	—
System: COBIT DS5.19: Malicious software	—
System: COBIT DS5.20: Firewall control	—
System: COBIT DS5.2: Authentication and Access	—
System: COBIT DS5.4: User Account Changes	—
System: COBIT DS5.7: Security Surveillance	—
System: COBIT DS9.4: Configuration Control	—
System: COBIT DS9.5: Unauthorized Software	—
System: CS-MARS Distributed Threat Mitigation (Cisco DTM)	System: CS-MARS Distributed Threat Mitigation (Cisco DTM)
System: CS-MARS Incident Response	System: CS-MARS Incident Response
System: CS-MARS Issue	

Table 7-2 Predefined Rule and Report Groups (continued)

System Report Groups	Corresponding System Rule Groups
System: Client Exploits, Virus, Worm and Malware	System: Client Exploits, Virus, Worm and Malware
System: Configuration Changes	—
System: Configuration Issue	System: Configuration Issue
System: Database Server Activity	System: Database Server Activity
System: Host Activity	System: Host Activity
System: Network Attacks and DoS	System: Network Attacks and DoS
System: New Malware Outbreak (Cisco ICS)	System: New Malware Outbreak (Cisco ICS)
System: Operational Issue	System: Operational Issue
System: Reconnaissance	System: Reconnaissance
System: Resource Issue	System: Resource Issue
System: Resource Usage	—
System: Restricted Network Traffic	System: Restricted Network Traffic
System: SOX 302(a)(4)(A)	—
System: SOX 302(a)(4)(D)	—
System: Security Posture Compliance (Cisco NAC)	System: Security Posture Compliance (Cisco NAC)
System: Server Exploits	System: Server Exploits

Global Controller and Local Controller Restrictions for Rule and Report Groups

Global Controller and Local Controller rule and report groups have the following restrictions:

- Rule and report groups created on the Global Controller are pushed to all the Local Controllers.
- Rule groups created on a Local Controller are local to the Local Controller. They are not copied to the Global Controller or to other Local Controllers.
- Local Controller account holders can edit only the Source IP, Destination IP, and Device fields of a rule group created on a Global Controller.
- Local Controller account holders cannot edit Global Controller report groups.
- Local Controller account holders cannot delete Global Controller rule and report groups.

Note The procedures described in this section are valid for both the Local and Global Controllers, except that the Case Bar does not appear on the Global Controller HTML interface.

Add, Modify, and Delete a Rule Group

Adding a New Rule Group

To add a rule group follow these steps:

Step 1 Navigate to the Inspection Rules page, as shown in Figure 7-7.

Figure 7-7 Inspection Rules Page

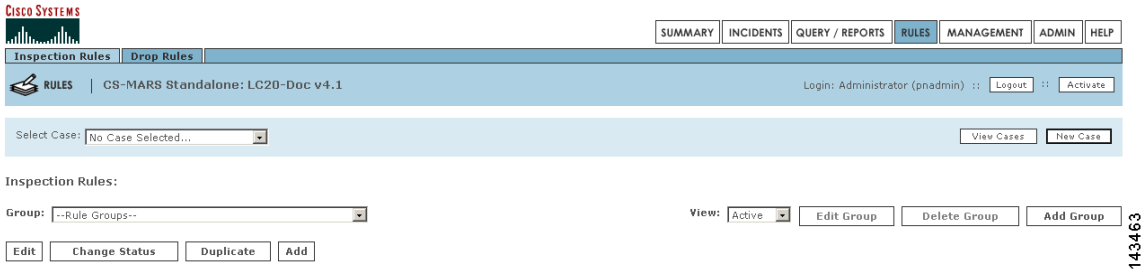

Step 2 Click **Add Group**.
The Add Group dialog box appears, as shown in Figure 7-8.

Figure 7-8 Add Group Dialog Box

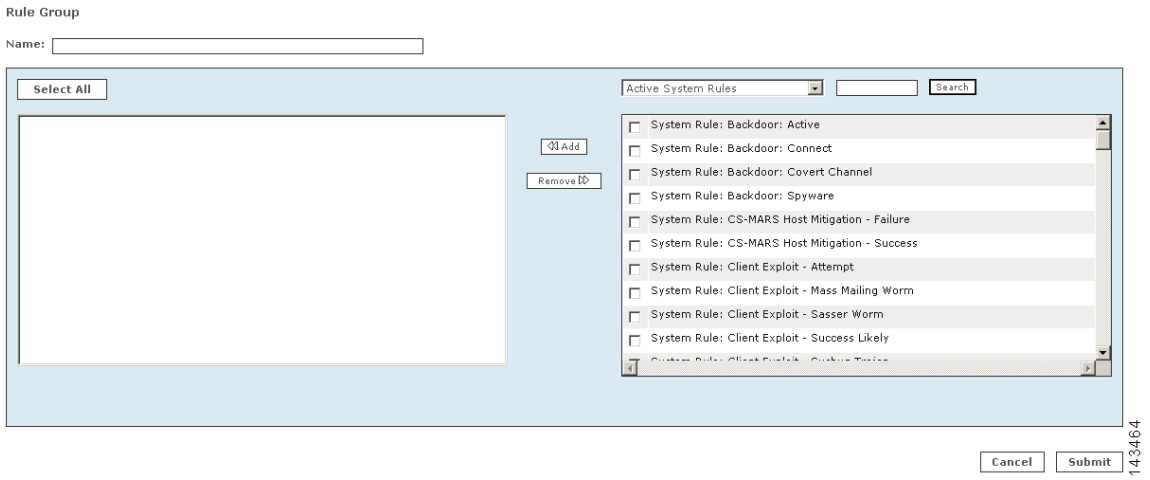

Step 3 Enter the new group name in the **Name** field.

Step 4 Click the checkboxes of the rules to be added to the new rule group.

Tip The dropdown list above the list of rules can limit the display of rules to active system rules, active user rules, or inactive rules. The search function displays only those rules that match a search string (for example, “New Malware Traffic Match.”). The asterisk wildcard character (*) is supported.

Step 5 Click **Add.**
The selected rules appear in the lefthand pane of the dialog box. To remove a rule from the group, highlight the item in the lefthand pane and click **Remove**.

Step 6 Click **Submit.**
The new rule group name appears in the **Group** dropdown filter on the Inspection Rules page, as shown in Figure 7-9. In this example, the new rule group name is “Example Group.” Because it is a user-created rule group, the rule group name appears without the prefix “System.” You can also click **Cancel** to return to the Inspection Rules page without creating a new rule group.

Figure 7-9 New Rule Group Appears on the Dropdown List of the Inspections Rules Page

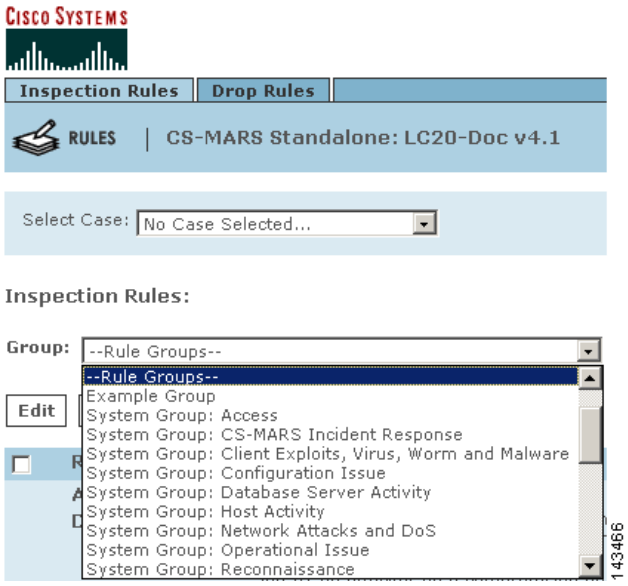

Modifying a Rule Group

To edit a rule group, follow these steps:

Step 1 Navigate to the Inspection Rules page, as shown in Figure 7-7.

Step 2 Select the rule group to edit in the **Group** pulldown filter.

Step 3 Click **Edit Group**.
The Add Group dialog box appears, as shown in Figure 7-8. The rule group name appears in the **Name** field, and the included rules appear as selected rules in the lefthand pane of the dialog box.

Step 4 To add additional rules, click the checkbox of all the rules to be added to the group, then click **Add**. To remove rules, highlight the items in the lefthand pane to remove, then click **Remove**.

Step 5 Click **Submit.**

Deleting a Rule Group

Step 1 Navigate to the Inspection Rules page, as shown in Figure 7-7.

Step 2 Select the rule group to delete in the **Group** pulldown filter.

Step 3 Click **Delete Group**.
The Delete Group dialog box appears listing the rules in the group to be deleted. You are prompted to confirm deletion.

Step 4 Click **Yes.**
The rule group no longer appears in the **Group** dropdown filters on the Incident and Inspection Rules pages.

Add, Modify, and Delete a Report Group

Adding a New Report Group

To add a report group follow these steps:

Step 1 Navigate to the Report page, as shown in Figure 7-10.

Figure 7-10 Reports Page

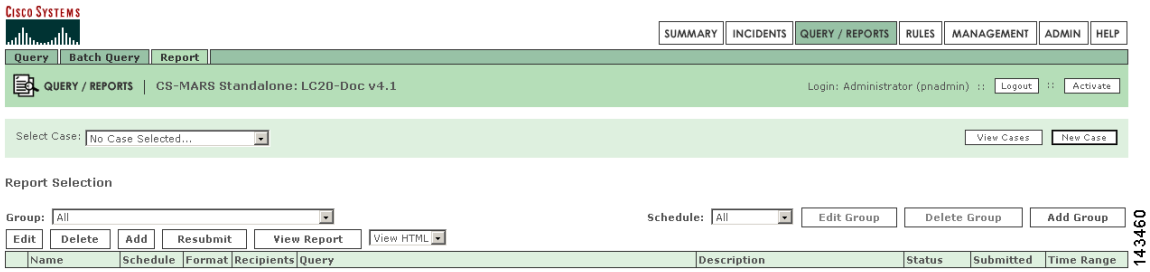

Step 2 Click **Add Group**.
The Add Group dialog box appears, as shown in Figure 7-11.

Figure 7-11 Add Report Group Dialog Box

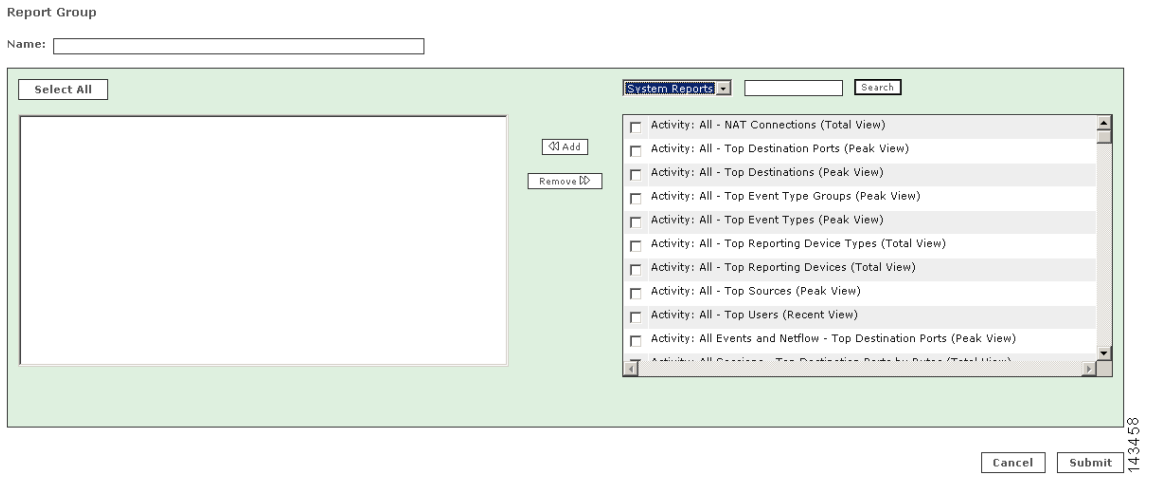

Step 3 Enter the new report group name in the **Name** field.

Step 4 Click the checkboxes of the reports to be added to the new report group.

Tip The dropdown filter above the list of reports can filter the display of reports to display system reports, user reports, or all reports. The search function displays only those reports that match a search string (for example, “Spy” for Spyware). The asterisk wildcard character (*****) is supported.

Step 5 Click **Add.**
The selected reports appear in the lefthand pane of the dialog box. To remove a report from the group, highlight the item in the lefthand pane and click **Remove**.

Step 6 Click **Submit.**
The new report group name appears in the **Group** dropdown list display filter on the Report page, as shown in Figure 7-12, and on the Query Page. Because it is a user-created report group, the report group name appears without the prefix "system." You can also click **Cancel** to return to the Report page without creating a new report group.

Figure 7-12 The New Report Group Appears on the Dropdown Filter of the Report Page

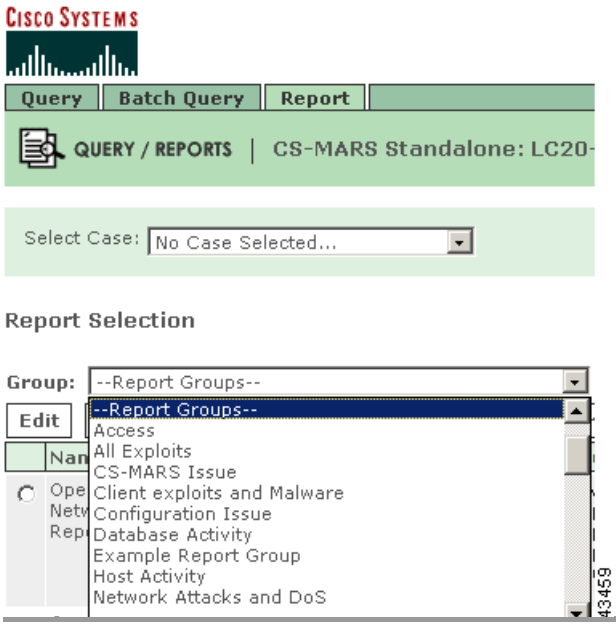

Modifying a Report Group

To edit a report group, follow these steps:

Step 1 Navigate to the Reports page, as shown in Figure 7-10.

Step 2 Select the report group to edit from the **Group** pull-down list.

Step 3 Click **Edit Group**.
The Add Report Group dialog box appears, as shown in Figure 7-11. The report group name appears in the **Name** field, and the reports that comprise the report group appear in the lefthand pane of the dialog box.

Step 4 To add additional reports, click the checkboxes of the reports to be added to the group, then click **Add**. To remove reports, highlight the items to remove in the lefthand pane, then click **Remove**.

Step 5 Click **Submit.**

Deleting a Report Group

Step 1 Navigate to the Reports page, as shown in Figure 7-10.

Step 2 Select the report group to delete in the **Group** pulldown filter.

Step 3 Click **Delete Group**.
The Delete Report Group dialog box appears listing the reports in the group to delete. You are prompted to verify deletion.

Step 4 Click **Yes.**
The report group no longer appears in the report group dropdown lists on the Report and Query pages.

Display Incidents Related to a Rule Group

To display incidents that occur from the firing of rules in a specific rule group, follow these steps:

Step 1 Navigate to the Incidents page.

Step 2 Select the rule group in the dropdown filter above the Matched Rules column, as shown in Figure 7-13. The Incidents page will display only those incidents that occurred from rules firing in the selected rule group.

Figure 7-13 Rule Group on Incidents Page

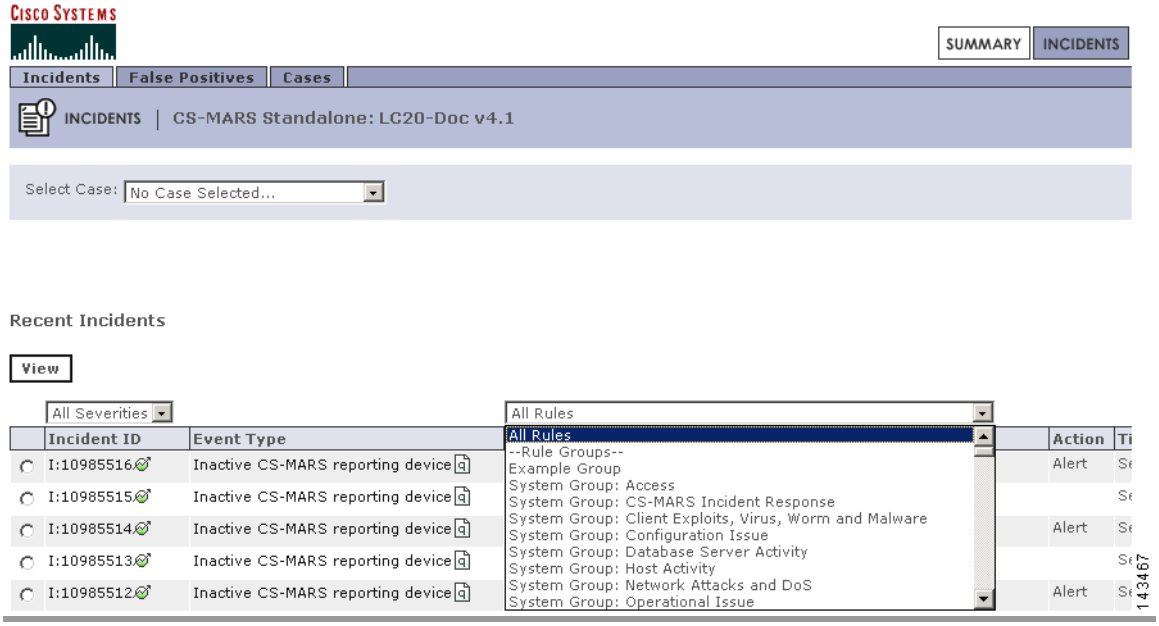

Create Query Criteria with Report Groups

To create queries from report groups, follow these steps:

Step 1 Navigate to the Query page.

Step 2 Select a report group in the **Load Report as On-Demand Query with Filter** dropdown filter, as shown in Figure 7-14.
Only the reports that comprise the report group can now display in the Select Report dropdown list, as shown in Figure 7-15.

Figure 7-14 Selecting A Report Group to Make a Query

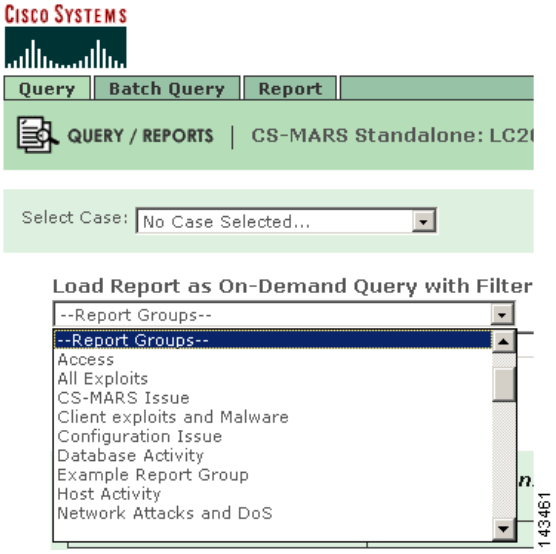

Figure 7-15 Selecting a Report Within the Report Group to Make a Query

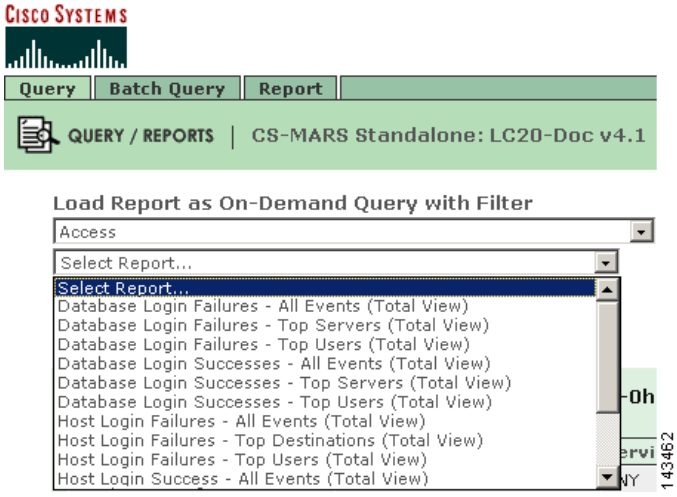

Step 3 Select the report in the secondary dropdown list.
The **Query** criteria are automatically populated per the selected report.

Using Rule Groups in Query Criteria

To populate the Rule field of the **Query Event Data** bar using rule groups, follow these steps:

Step 1 Navigate to the Query page.

Step 2 Click **Any** in the **Rules** field of the **Query Event Data** bar.
The Filter by Rule dialog box appears as shown in Figure 7-16.

Step 3 Select the rule group in the dropdown list above the list of rules, as shown in Figure 7-11.
The list of rules will display only those rules in the selected rule group.

Figure 7-16 Rule Group Used to Populate Rule Criterion in Query

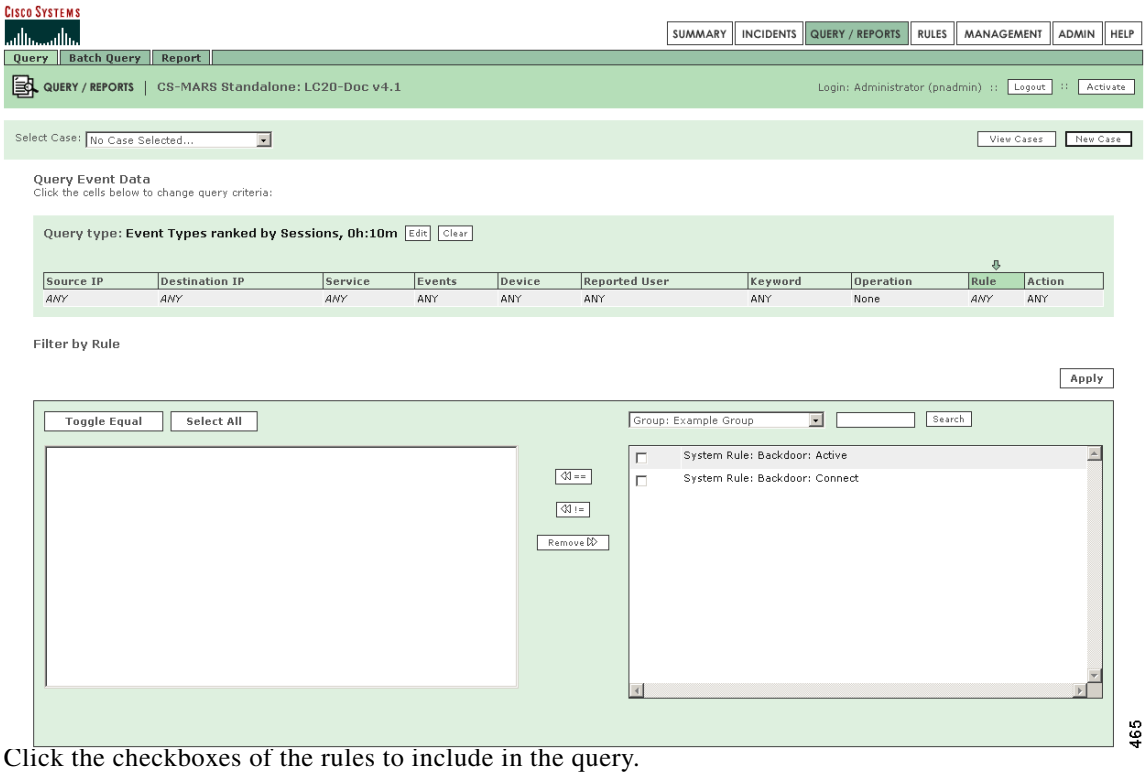

465

Step 4 Click the checkboxes of the rules to include in the query.

Step 5 Click **Add**. The selected items appear in the lefthand pane of the Query dialog box.
To remove rules, highlight the items to remove in the lefthand pane, then click **Remove**.

Step 6 Click **Apply**.
The selected rules appear in the **Rules** field of the **Query Event Data** bar.

CHAPTER 8

Sending Alerts and Incident Notifications

A Cisco Systems MARS alert action is a signal transmitted to people or devices as notification that a MARS rule has fired, and that an incident has been logged. Alert actions can only be configured through the Action parameter of a rule. An alert action determines which alert notification types are sent to which MARS user accounts or user groups. MARS can transmit alerts by the methods listed in Table 8-1.

Table 8-1 MARS Incident Notification Methods

Alert Notification Type	Description
Sent in Human-Readable Format	
• E-mail • XML Notification • Short Message Service (SMS) • Pager	E-mail, SMS, and pager alerts send the incident ID, matched rule name, severity, and incident time in email, SMS and pager formats respectively. You must login to the MARS to view all the incident details. XML notification sends an email notification of an incident with an attached XML data file (see Example 8-2). The XML data file contains the same incident details that can be viewed from the GUI, except for path and mitigation information. The XML data file can be sent as a plain-text file or as a compressed gzip file. The XML data filename is constructed with the incident ID number, for example `CS-MARS-Incident-13725095.xml`. You can parse and extract data from the XML file with a custom application. For example, you can integrate the XML data with trouble ticketing software. See Appendix A, "Cisco Security MARS XML API Reference," for further information on the MARS XML notification schema and usage guidelines. MARS SMS text message notifications can be up to 160 characters in length. Because the MARS SMS incident notification exceeds 160 characters, it is sent in three segments. Pager messages are sent through the MARS internal modem. MARS dials a carrier's IXO/TAP number and uses SNPP to transmit the alpha-numeric page. Pager notifications are still possible when the network is down. Pagers can often receive messages in places where mobile phones are inoperative or forbidden (for instance, hospitals).
Sent to a Device	
• SNMP trap • Syslog • Distributed Threat Mitigation	These alerts send the incident ID, matched rule severity, and incident time to devices or applications, all of which must be properly configured within the MARS device administration pages.

Table 8-2 provides links and description of related Alert Action configuration procedures. Although some of these procedures are documented elsewhere in this user guide, they are duplicated here for your convenience.

Table 8-2 Alert Notification Procedures

Alert Related Procedures	Description
Configure the E-mail Server Settings	To send Email, SMS, and XML notifications, MARS requires that you configure the E-mail Server settings.
Configure a Rule to Send an Alert Action	Complete this procedure to create or modify an alert action.
Create a New User—Role, Identity, Password, and Notification Information	Alert notifications can be sent only to user accounts configured on MARS. A new user account can be configured from the User Management tab, or when creating an alert action for a rule. This is where you enter the service provider phone numbers and email addresses for E-mail, SMS, Pager, and
Create a Custom User Group	Complete this procedure to create a MARS user group other than the default MARS user groups. Unlike default user groups, custom groups can be edited.
Add a User to a Custom User Group	Complete this procedure to include a newly created user account into a MARS user group.

Example 8-1 shows a typical email alert notification. Example 8-2 shows an XML notification with its attached XML data file. When compression is configured, the XML date file arrives as a GZIP compressed file.

Note Alert notifications cannot be customized.

Example 8-1 MARS Notification by Email

```
-----Original Message-----
From: notifier.Latest@serviceprovider.cisco.com [mailto:notifier.MyLatest@cisco.com]
Sent: Monday, May 15, 2006 8:48 AM
To: Naliza Mahda (Nalmah)
Subject: Incident Notification (green, Rule Name: System Rule: CS-MARS Database Partition
Usage)

The following incident occurred:

Start time:     Mon May 15 08:47:26 2006
End time:       Mon May 15 08:47:26 2006
Fired Rule Id:  134473
Fired Rule:     System Rule: CS-MARS Database Partition Usage
Incident Id:    597842933

For more details about this incident, please go to:

  https://MyLatest/Incidents/IncidentDetails.jsp?Incident_Id=597842933
  https://MyLatest.cisco.com/Incidents/IncidentDetails.jsp?Incident_Id=597842933
  https://10.2.3.7/Incidents/IncidentDetails.jsp?Incident_Id=597842933
  https://192.168.1.101/Incidents/IncidentDetails.jsp?Incident_Id=597842933
```

```
For all recent incidents, please go to:

  https://MyLatest/Incidents/
  https://MyLatest.cisco.com/Incidents/
  https://10.2.3.7/Incidents/
  https://192.168.1.101/Incidents/
```

Example 8-2 MARS XML Notification Email Attachment

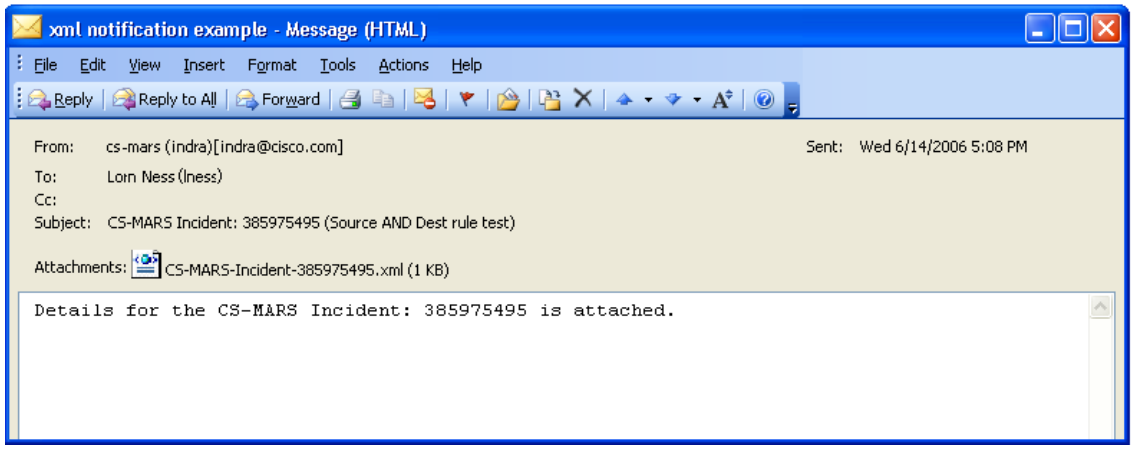

Configure the E-mail Server Settings

To send alert actions, MARS must be configured to communicate with an e-mail server. To configure the e-mail server settings, follow these steps:

Step 1 Click **Admin > Configuration Information**.

The Device Configuration window appears, as shown in Figure 8-1.

Figure 8-1 MARS Device Configuration Window

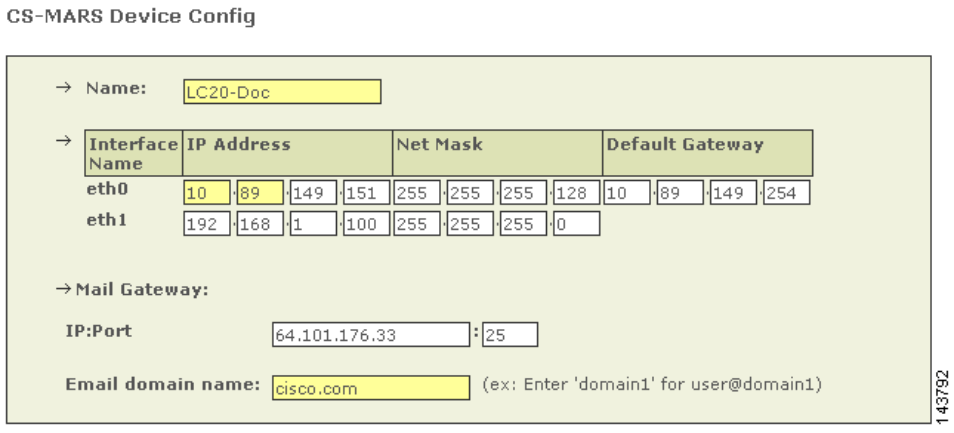

Step 2 In the **IP:Port** field of the **Mail Gateway** section, enter the IP address and **Email Domain Name** of your Mail Gateway server.

Step 3 Click the **Update** button at the bottom of the page to update the MARS configuration.

Configure a Rule to Send an Alert Action

To send alert notifications to individual users or groups of users, configure the Action parameters of a rule to create an alert action. This procedure configures alerts for pre-existing rules. When you create a rule, the Action parameters are configured after the count number parameter.

Note Drop rules do not have Action parameters and cannot trigger alerts.

To modify or create an alert for an existing rule, follow these steps:

Step 1 Click the **RULES** tab to navigate to the Inspection Rules page.

Step 2 Identify the Rule to configure, and click the value displayed in the **Action** field.

The Action Selection dialog box, as shown in Figure 8-2, appears below the rule description table. All previously defined alert actions are listed in the right-hand area of the Action dialog box. An alert action determines which alert notifications are sent to which users or user groups when the rule fires. You can edit or delete existing alert actions or create a new one.

Figure 8-2 Action Selection Dialog

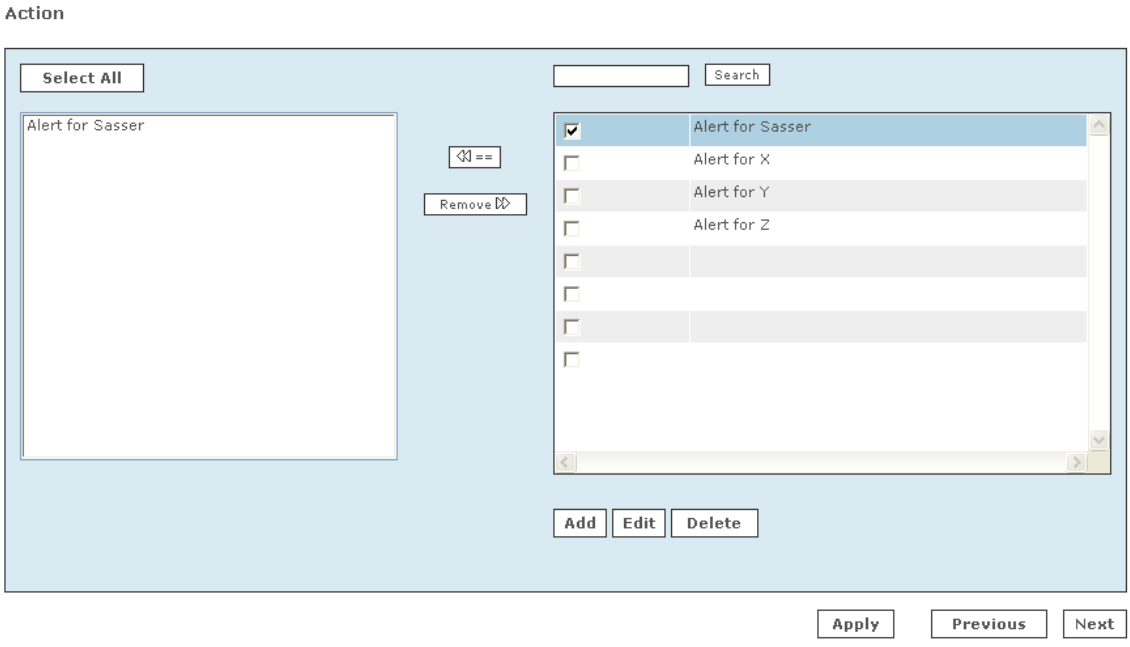

Step 3 Do one of the following five actions:

1. Remove an alert action currently applied to the rule.

  - In the left-hand area, pick the alert actions to remove with Ctrl+Click, then click **Remove >>**.

    The alert action is deleted from the left-hand area.
  - Proceed to Step 13 to complete the procedure.

- Apply an existing alert action to the rule.
  - In the right-hand area, click the check boxes of the alert actions you require, then click **<<==** .

    The alert action appears in the left-hand area.
  - Proceed to Step 13 to complete the procedure.
- Delete an existing alert action from MARS.
  - Click the check box of the alert action in the right-hand area, then click **Delete**.

    A delete verification window appears.
  - Click **Yes.**

    The alert action is deleted from the right-hand area.
  - Proceed to Step 13 to complete the procedure.
- Edit an existing alert action.
  - Click the check box of the alert action in the right-hand area, then click **Edit**.

    The Alert recipients page appears in a new window, as shown in Figure 8-3.
  - Proceed to Step 4 to complete the procedure.
- Create a new alert action.
  - Click **Add**.

    The Alert recipients page appears in an a new window, as shown in Figure 8-3.
  - Proceed to Step 4 to complete the procedure.

Figure 8-3 Alert Recipients Window

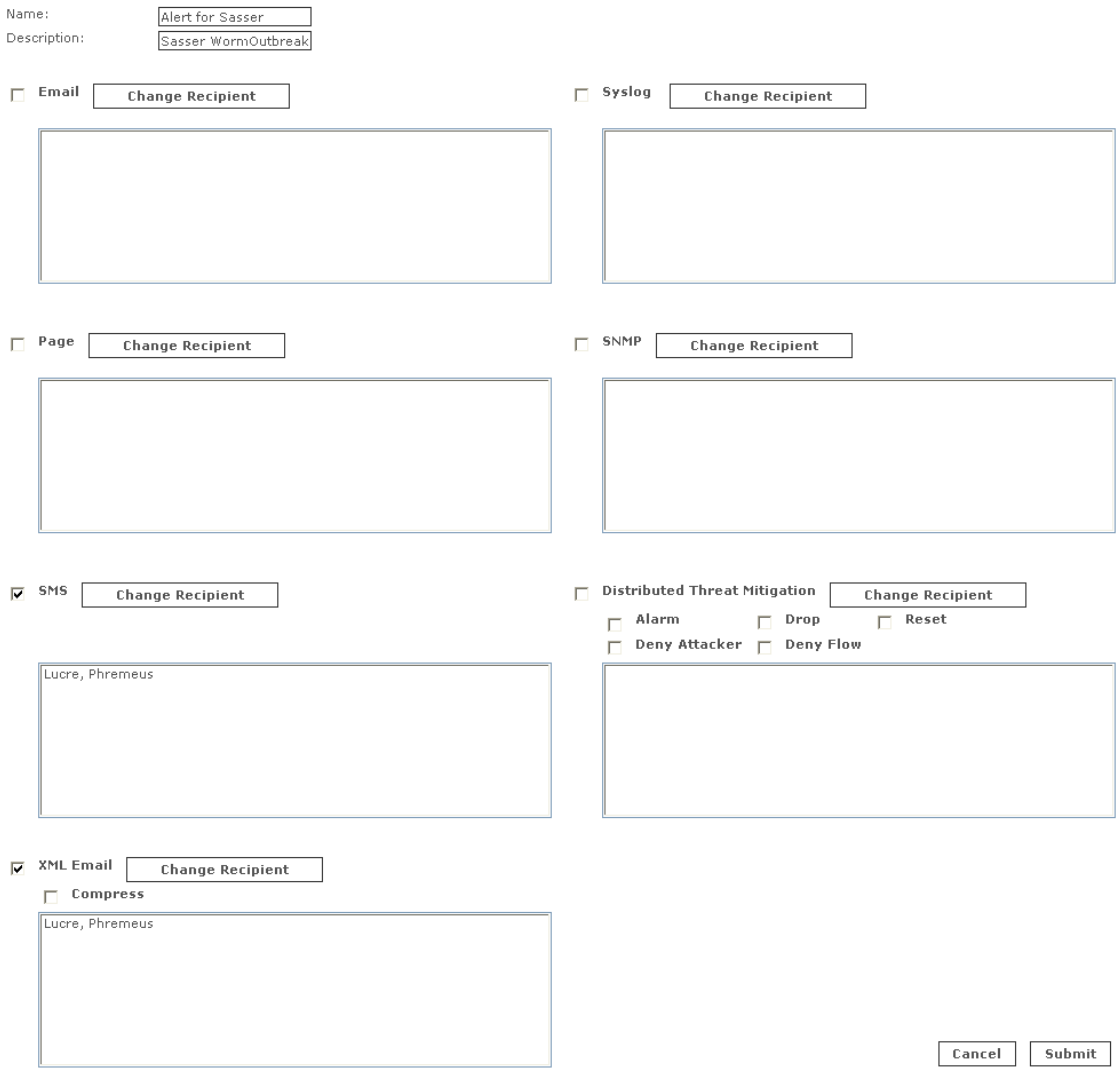

Step 4 For a new alert enter a name and description in the **Name** and **Description** fields. If editing an existing alert, you can modify the name or description.

Step 5 Click the check box of a notification type to select or deselect it.

Recipients for the notification types are as follows:

- **E-mail—**Users or user groups can receive an e-mail.
- **Page—**Users or user groups can receive an alpha-numeric electronic page on their pagers or pager-enabled mobile telephones.
- **SMS—**Users or groups can receive a text message on their SMS-enabled mobile telephones.
- **XML Email—**Users or groups can receive an email message with incident details appended in an XML data file. Click the **Compress** check box to send the XML data file as a compressed gzip file. For more information on this feature, see Appendix A, "Cisco Security MARS XML API Reference."
- **Syslog—**Specified devices can receive syslog messages.

- **SNMP—**Specified devices can receive SNMP trap information.
- **Distributed Threat Mitigation—**For more information on this feature, see Technology Preview: Configuring Distributed Threat Mitigation with Intrusion Prevention System in Cisco Security MARS, page 1.

Note For SNMP and Syslog, you must configure the receiving systems to receive notifications.

Step 6 Click the **Change Recipient** button to add or remove a recipient for a notification type.

For E-Mail, Page, SMS, and XML Email, the **Select** (recipient) dialog box appears, as shown in Figure 8-4.

Figure 8-4 Select Recipient Dialog Box

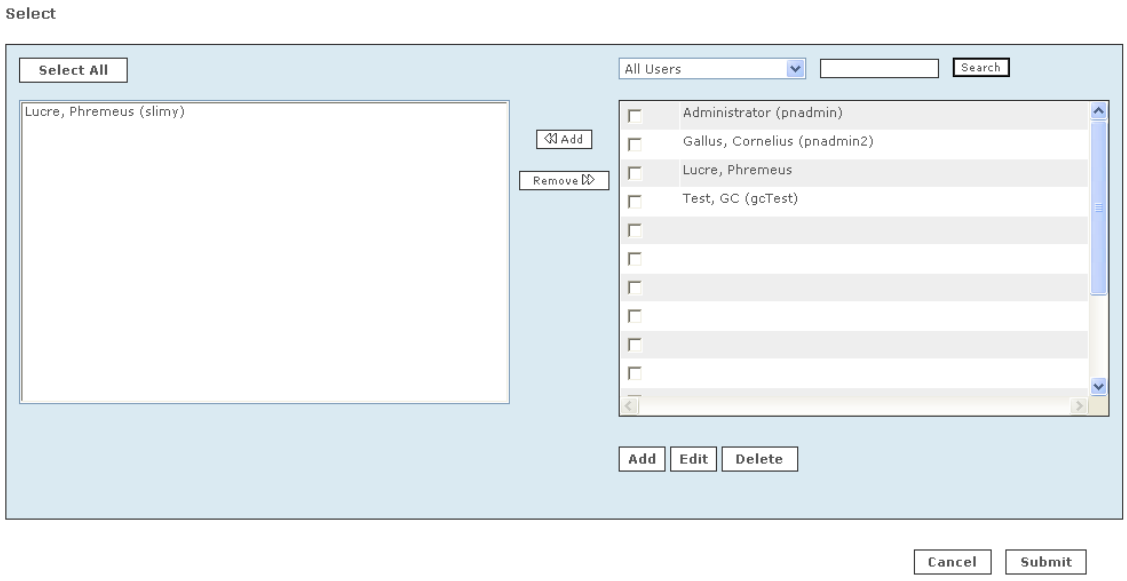

For Syslog and SNMP, the **Select** (device) dialog box appears, as shown in Figure 8-5.

Figure 8-5 Device Selection Page

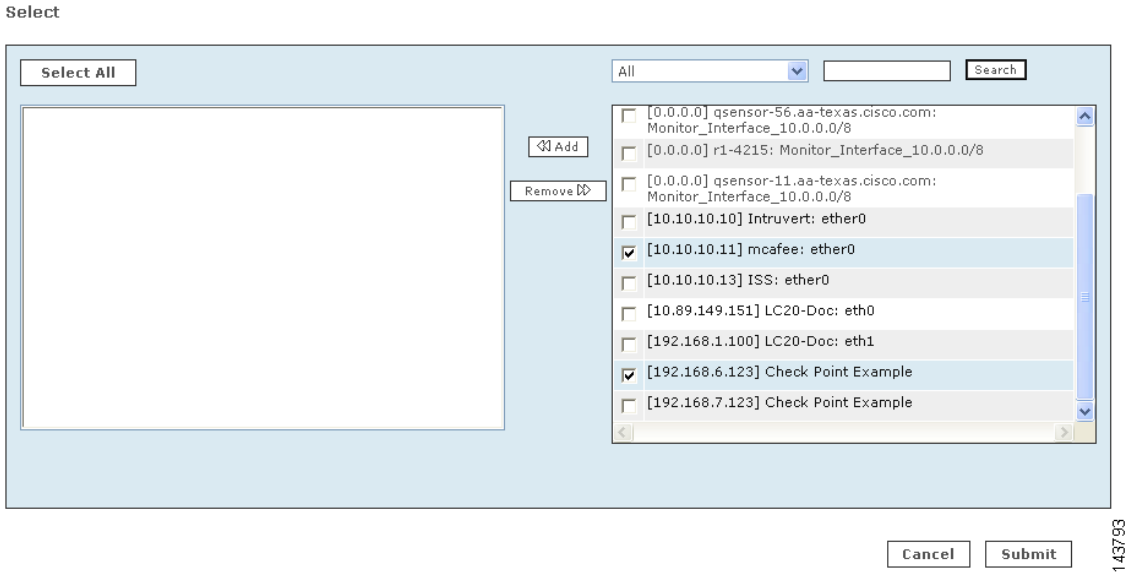

For Distributed Threat Management notification, the Select (IOS-IPS Devices) dialog appears (not shown).

Tip If you do not know the group to which a user or device belongs, select **All** from the dropdown list to view all users or devices.

Step 7 Click the check box next to the users or device you want to receive the notification, then click **<< Add**.

Your selections appear in the left-hand area. To remove items, Ctrl+click the items in the left-hand area, then click **Remove**. The items are then deleted from the left-hand area.

Step 8 If you are not adding a user, skip to Step 9. To add a new user, do the following substeps:

a. Click **Add.**

The User Configuration page appears in a separate window, as shown in Figure 8-6.

b. Enter the User Configuration information then click **Submit**.

You are returned to the Select Recipient Dialog Box.

For reference on user configuration fields, see the section, "Create a New User—Role, Identity, Password, and Notification Information"

c. Add the new user to the recipient list as described in Step 7.

Step 9 Click **Submit**.

You are returned to the Alert Recipients Window.

Step 10 Repeat Step 6 through Step 9 until you have assigned recipients to all the notification types you have selected.

Step 11 Click **Submit**.

You are returned to the Action Selection Dialog. Any newly-created or edited action alert appears in the right-hand area.

Step 12 Click the check boxes next to the action alerts to be sent when the rule fires. Click **<< Add**.

Your selections appear in the left-hand area.

Step 13 Click **Next**.

The Time Range dialog may or may not appear.

Step 14 Click **Next** if the Time Range dialog appears.

The Rule Summary table appears.

Step 15 Click **Submit** to save your changes to the rule.

Step 16 Verify that the alert actions you selected appear in the Action field of the rule description.

Note An inactive rule is made active by applying an alert action. To inactivate a rule, select the rule and click **Change Status**.

This ends the Configure a Rule to Send an Alert Action procedure.

Create a New User—Role, Identity, Password, and Notification Information

To create a new MARS user, complete the following steps:

New user accounts and user groups are created on the **Management > User Management** tab, or as a substep in creating an alert notification recipient (with the **Add** button on the Select [user] dialog).

Step 1 Navigate to the User Management page by either of the following methods:

- Click **Add** on the **Management > User Management** tab.
- Click **Add** on the Select (user) dialog box when creating an alert notification. See "Configure a Rule to Send an Alert Action" section on page 8-5.

The User Configuration page appears, as shown in Figure 8-6.

Figure 8-6 User Configuration Page

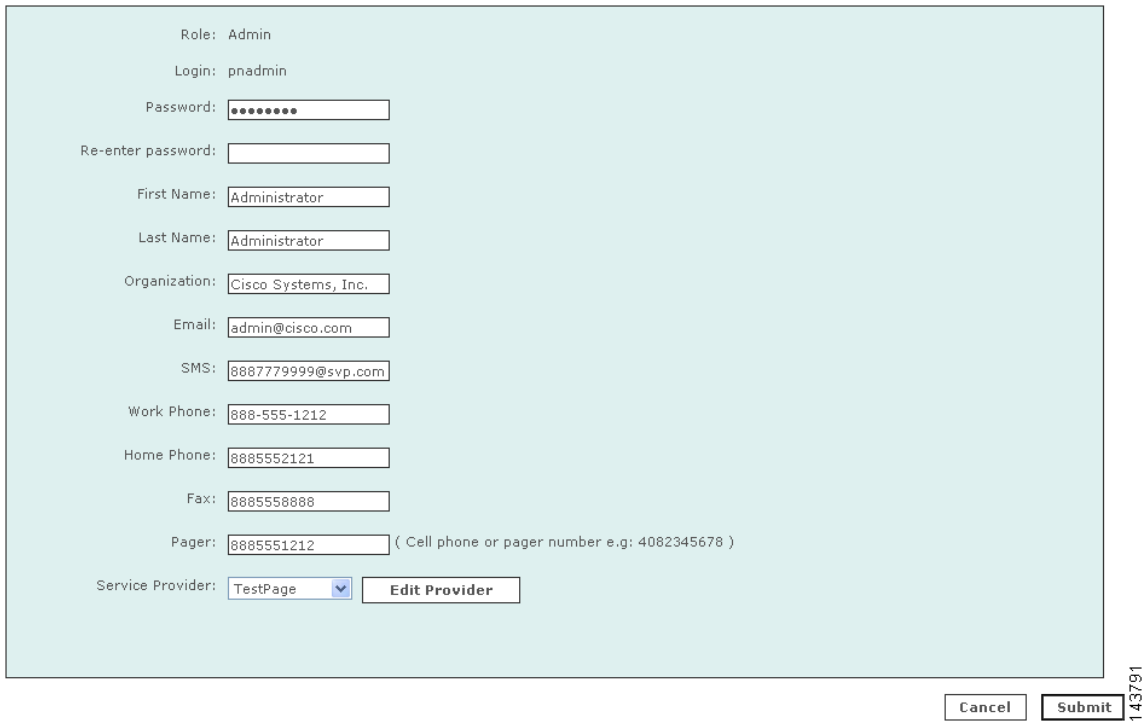

Step 2 From the **Role** field, select a **Role** for the user.

- **Admin**: has full use of the MARS.
- **Notification Only**: for a non-user of the MARS appliance, use this to send alerts to people who are not administrators, security analysts, or operators.
- **Operator**: has read-only privileges.
- **Security Analyst**: has full use of the MARS, except cannot access the Admin tab

Step 3 Create or change the user's password if necessary.

Step 4 Enter the user's credentials and personal information, which may include any of the following:

- First name
- Last name
- Organization name
- Email address
- Short Message Service (SMS) number—for example, 8885551212@servprov.com
- Work telephone number
- Home telephone number
- FAX number
- Pager number or ID— may also be a mobile telephone number, for example, 5552345678

Step 5 If you are not creating a notification by pager, go to Step 10.

Step 6 For notification by pager, you must specify a service provider (cell phone or pager company). From the Service Provider field, select **New Provider**.

This pull-down menu is populated as you add new providers.

Additional service provider information fields appear on the same page, as shown in Figure 8-7.

Figure 8-7 Service Provider Fields to Add or Change a Service Provider

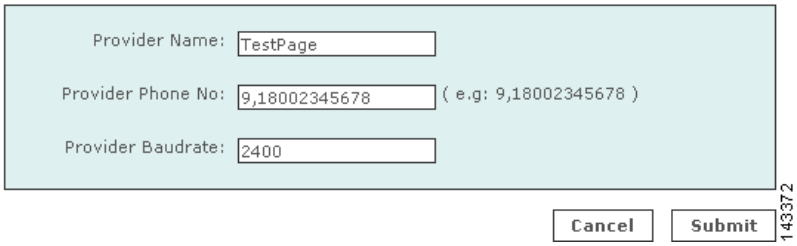

Step 7 In the **Provider Name** field, enter the name of the service provider.

Step 8 In the **Provider Phone No** field, enter the service provider's telephone number.

This is the number the service provider requires for accepting alpha-numeric messages using the IXO/TAP protocol. The format is like a regular phone number, such as: 18001234567. The format of 1-800-1234567 is also acceptable. If dialing "9" is required to access a number outside your private branch exchange, type a "9," before the full telephone number (for example, 9,1-800-1234567).

Step 9 In the **Provider Baudrate** field, enter the baud rate specified by the provider.

This is the baud rate the service provider requires for the specified phone number. Common values are 1200, 2400, 4800, and 9600.

Step 10 Click **Submit** to close the User Configuration page and return to the **User Management** tab.

This ends the Create a New User—Role, Identity, Password, and Notification Information procedure.

Create a Custom User Group

To create a custom user group in addition to the default groups created by MARS, complete the following procedure:

Step 1 Navigate to the **Management > User Management** tab.

Step 2 Click **Add Group**.

Step 3 In the **Name** field, enter a name for the group.

Step 4 To add users to the group, click the check box of users from the list on the right-hand area. Click **Add**.

The checked names appear in the left-hand side of the dialog box.

To remove users from the group, pick the users from the left-hand side with Ctrl+click. Click **Remove**.

The selected names appear in the right-hand side of the dialog box.

Step 5 Click **Submit**.

You are returned to the User Management tab.

This ends the Create a Custom User Group procedure.

Add a User to a Custom User Group

To include a user in a custom User Group, complete the following steps:

Note The user is automatically added to the User Group that corresponds to their role. Admin, Operator, Notification, and Security Analyst are system groups and cannot be edited.

Step 1 Navigate to the **Management > User Management** tab.

Step 2 Select the User Group to edit from the **Select Group** dropdown list.

The members of the group are displayed.

Step 3 Click **Edit Group**. The User Group dialog box appears.

Step 4 Check the users to add to the group from the list on the right hand side. Click **Add**. The checked names move to the left-hand area of the dialog box.

Step 5 Click **Submit**.

You are returned to the **User Management** tab.

This ends the Add a User to a Custom User Group procedure.

CHAPTER 9

Management Tab Overview

Revised: April 5, 2007, OL-13542-01

Use the management features in the Global Controller to assign: event, addressing, service, and user information. This information is used in rules, queries, and to determine false positives.

Activating

In general, you need to activate changes in the Management tabs if the changes are part of a rule.

To activate a set of management additions or changes

Step 1 When changes (or additions) are complete, activate them by clicking **Activate**.

Figure 9-1 Clicking the Activate Button

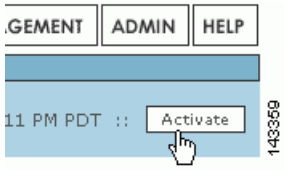

Event Management

To open the Event Management sub-tab, click the **Management > Event Management** tabs.

On the Event Management page, you can search and filter events and event groups, and work with groups of events.

Search for an Event Description or CVE Names

You can search for partial matches of event descriptions or Common Vulnerabilities and Exposures (CVE) names.

Step 1 Enter the text that you want to search for in the **Search** field.

Step 2 Click **Search**.

To view a list of all currently supported CVEs

Step 1 Enter CVE into the **Search** field.

Step 2 Click **Search**.

Event Groups

Using and creating event groups is one of the most powerful ways to leverage rules. You can take any of the events presented here, group them, and then use them with rules to concentrate your searches for attacks.

To filter by event groups or severity

From the appropriate list, select the group or severity.

Edit a Group of Events

Note You can not edit system-defined groups.

Step 1 Select the group in the **Select Group** list.

Step 2 Click **Edit Group**.

Step 3 Click each group in the Chosen and Available fields to highlight it. Click it again to de-highlight it.

Step 4 Click **Add** or **Remove** to move highlighted items as needed.

Step 5 Click **Submit**.

Add a Group

Step 1 Click **Add**.

Step 2 In the **Name** field, enter a name for the group.

Step 3 In the **Available** field, click each group that you want to add to highlight it. Click it again to de-highlight it.

Step 4 Click **Add**.

Step 5 Click **Submit**.

IP Management

The IP Management page, accessed by clicking **Management > IP Management**, enables the definition of network assets that you use as building blocks for inspection rules, drop rules, reports and queries, topology discovery schedules, and in defining reporting devices and mitigation devices. You can define assets as networks, IP ranges, or hosts. You can also defined named variables for use within inspection rules.

The vulnerability assessment information that you define for a host, specifically the operating system type and patch level and the known services that run on the host, assists MARS in determining false positives.

Tip You can filter the list of objects displayed by the View list box. This selection allows you to filter to hosts, networks, IP ranges, or variables.

Note A Global Controller pushes any global IP Management Groups to the active Local Controllers that it manages.

Search for an Address, Network, Variable, or Host

Step 1 Enter the text that you want to search for in the **Search** field.

Step 2 Click **Search**.

Filter by Groups

From the **Select Group** list, select the group.

Edit a Group

Step 1 Select **Management > IP Management**.

The IP Management page appears.

Step 2 Select the group in the **Select Group** list.

Step 3 Click **Edit Group**.

Step 4 Click each group in the **Chosen** and **Available** fields to highlight it. Click it again to de-highlight it.

Step 5 Click **Add** or **Remove** to move highlighted items as needed.

Step 6 Click **Submit**.

Add a Group

Step 1 Select **Management > IP Management**.

The IP Management page appears.

Step 2 Click **Add Group**.

Step 3 In the **Name** field, enter a name for the group.

Step 4 In the **Available** field, click a group to highlight it. To de-highlight an item, click it again.

Step 5 Click **Add** to move the selected Event Type Groups into the **Chosen** field.

Step 6 Click **Submit**.

Add a Network, IP Range, or Variable

Step 1 Select **Management > IP Management**.

The IP Management page appears.

Figure 9-2 Add a Network, IP Range, or Variable

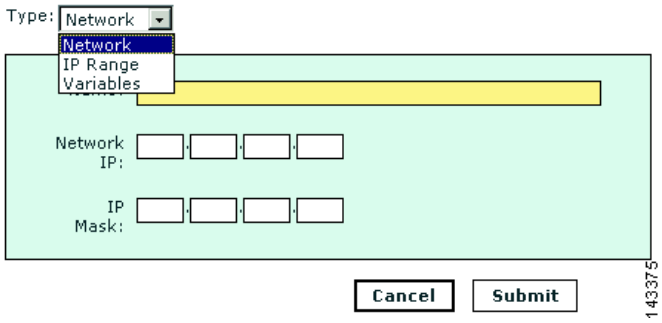

Step 2 Click **Add**.

Step 3 In the **Type** list select: network, IP range, or variable.

Step 4 For each type enter the appropriate information.

- Network: name, network IP, network mask
- IP range: name and range
- Variable: variable name

Step 5 Click **Submit**.

Service Management

To open the Service Management sub-tab, click the **Management > Service Management** tabs.

Service is a combination of source port, destination port and protocol. The Service Management page displays services and their descriptions, ports and protocols. On the Service Management page, you can work with the services on your networks.

Search for a Service

Step 1 Enter the text that you want to search for in the **Search** field.

Step 2 Click **Search**.

To filter by service groups

From the appropriate list, select the group.

Add a Group of Services

Step 1 Click **Add**.

Step 2 In the **Name** field, enter a name for the group.

Step 3 In the **Available** field, click items to select them, and click them again to de-select them.

Step 4 Click **Add**.

Step 5 Click **Submit**.

Edit a Group of Services

Note You can not edit system-defined groups.

Step 1 Select the group in the **Select Group** list.

Step 2 Click **Edit Group**.

Step 3 Click each group in the **Chosen** and **Available** fields to highlight it. Click it again to de-highlight it.

Step 4 Click **Add** or **Remove** to move the highlighted items as needed.

Step 5 Click **Submit**.

Add a Service

Step 1 Click **Add**.

Step 2 Enter the service's details.

Step 3 Click **Submit**.

Edit a Service

Step 1 Check the box next to the service.

Step 2 Click **Edit**.

Step 3 Make your changes, and click **Submit**.

Delete a Service

Step 1 Check the box next to the service.

Step 2 Click **Delete**.

Step 3 On the confirmation page, click **Yes**.

User Management

MARS supports local authentication of MARS users; user credentials are stored the MARS Appliance in SHA-1 cryptographic hash format. Each MARS Appliance only has one Administrative account, *pnadmin*. This account is the only account with privileges to access the command line interface via SSH or direct console connection.

The User Management page allows you to manage other users and administrators of the MARS system, including the roles and groups to which those users belong. On this page, you can define new user accounts, enabling access to specific features of the web interface. You can define user-specific notification settings for the user, such as a valid e-mail address or pager number. Some system-wide

settings, such as pager and cell phone service provider settings, are also accessible exclusively through this page. To access the User Management page, click either **Management > User Management** or **Admin > User Management**.

In MARS, four separate user roles exist that can be assigned to any user who needs to access the web interface:

- *Admin* has full read/write privileges. Users in this role can define new users with any desired role. Users in the role can change the password settings of the accounts in any user role.
- *Security Analyst* has full read privileges but is restricted to write for reports privileges. Users in this role can only define new users (and change passwords of users) with the Notifications Only role.
- *Operator* has read only privileges. Users in this role cannot define new users or change passwords, even of their own user account.
- *Notifications Only*. This user role has no permissions to access to the MARS web interface; use this role to identify users who will receive notifications, such as e-mail, SMS, or pager notifications.

No limit exists on the number of user accounts that can be defined in MARS.

While roles are system defined, you can define, edit, and delete user groups. For more information, see Create a User Group, page 9-10 and Add or Remove a User from a User Group, page 9-10.

Users created on the Global Controller are propagated down to the Local Controller with one notable exception: the user “pnadmin” is always local to the Global Controller or Local Controller on which it is first created.

When you create users with the same login name or the same first name/last name combination on both the Global Controller and a Local Controller, both appear in the list of users on the Local Controller: once as a local user, once as global.

Global users are maintained only on the Global Controller; local users are maintained only on individual Local Controllers. Users created on Local Controllers are not propagated up to the Global Controller. If you want a user of a Local Controller to have access to the Global Controller or any of its information, you must also create that user at the Global Controller level.

Good security practices suggest strong passwords for use with the MARS Appliances. When defining user names and password, keep the following guidelines in mind:

Login names and passwords:

- can be alphanumeric characters
- can contain special characters (!, @, #, etc.)
- *cannot* contain single or double quotes (‘or “)
- are case sensitive

Login names can have up to 20 characters. Passwords can have up to 64 characters.

Add a New User

Defining a new user involves specifying the user name, password, role, contact information, PGP key (Global Controller only), and notification information.

To add a new user, follow these steps:

Step 1 From the **Management > User Management** tab, click **Add**. The User Configuration page appears, as shown in Figure 9-3.

Figure 9-3 User Configuration Page

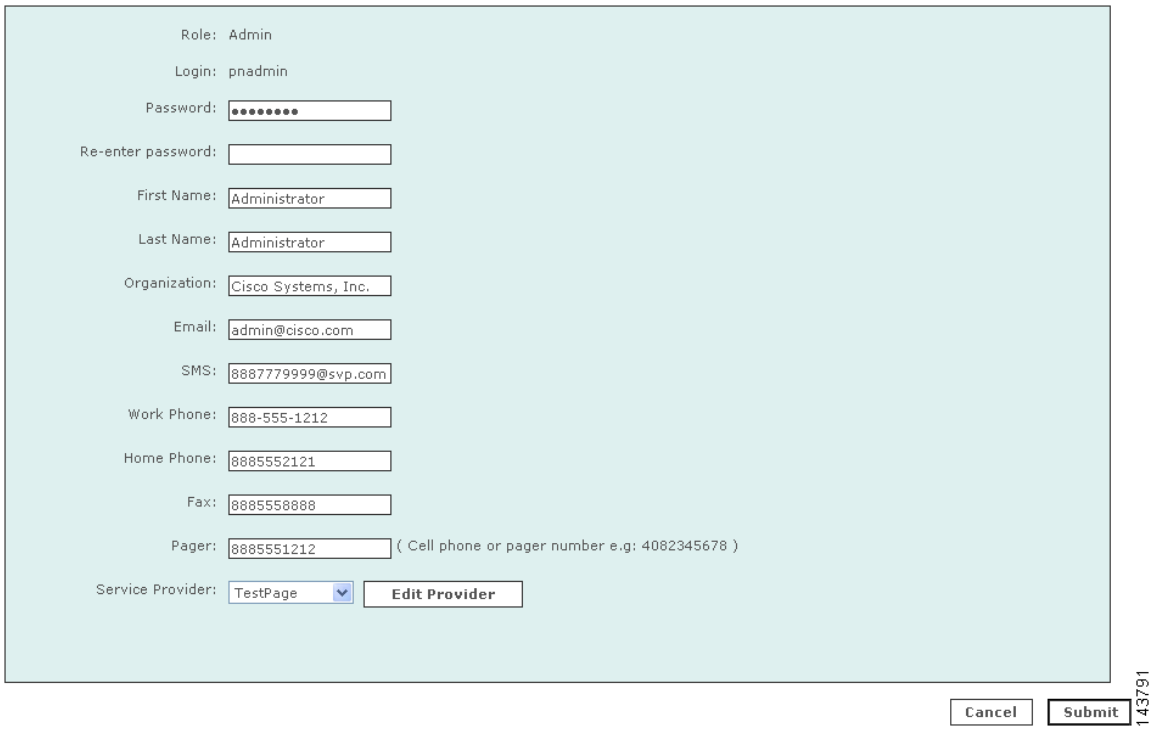

Step 2 From the **Role** field, select a **Role** for the user.

- **Admin**: has full use of Global Controller.
- **Notification Only**: for a non-user of the Global Controller appliance, use this to send alerts to people who are not admins, security analysts, or operators.
- **Operator**: has read-only privileges.
- **Security Analyst**: has full use of Global Controller, except cannot access the Admin tab

Step 3 Create or change the user's password if necessary.

Step 4 Enter the user's credentials and personal information.
The information can include the following:

- First name
- Last name
- Organization name
- Email address
- PGP Key
- Short Message Service (SMS) number—for example, 8885551212@servprov.com
- Work telephone number
- Home telephone number
- FAX number
- Pager number— may also be a mobile telephone number, for example, 5552345678

Step 5 If you are creating a notification by pager, go to the next section, "Add a Service Provider (Cell phone/Pager)", otherwise click **Submit** to complete the procedure for adding a user.

Add a Service Provider (Cell phone/Pager)

When configuring a notification by pager, add a service provider (cell phone or pager company) by completing the following procedure:

Step 1 From the **Service Provider** field, select **New Provider**. Additional fields appear, as shown in Figure 9-4.

The pull-down menu is populated as you add new service providers.

Figure 9-4 Select a New Provider and Provide Contact Details

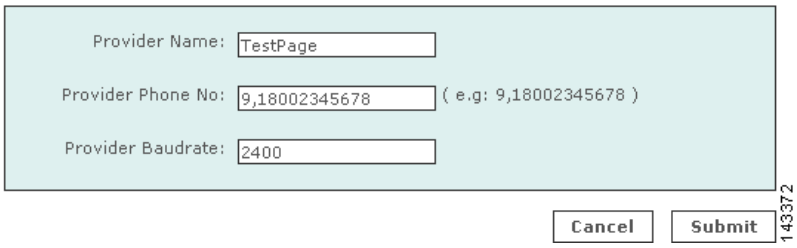

Step 2 In the **Provider Name** field, enter the name of the service provider.

Step 3 In the **Provider Phone No** field, enter the service provider's telephone number.

This is the number the service provider uses for accepting alpha-numeric messages using the IXO/TAP protocol. The format is like a regular phone number, such as: 18001234567. The format of 1-800-1234567 is also acceptable. If dialing "9" is required to access a number outside your private branch exchange, type a "9," before the full telephone number (for example, 9,1-800-1234567).

Step 4 In the **Provider Baudrate** field, enter the baud rate specified by the provider.

This is the baud rate the service provider requires for the specified phone number. Common values are 1200, 2400, 4800, and 9600.

Consult your service provider's website for more information on their baud rates.

Step 5 Click **Submit** to close the User Configuration page and return to the **User Management** tab.

Search for a User

Step 1 Enter the text that you want to search for in the **Search** field.

Step 2 Click **Search**.

Edit or Remove a User

Step 1 Form the **Management User tab**, check the box next to the user's name.

Step 2 Click **Delete** to delete the user.

Step 3 Click **Edit** to change the user's configuration information.
The User Configuration page appears.

Step 4 Edit the User Configuration page.

Step 5 Click **Submit**.

Create a User Group

Step 1 Click **Add Group**.

Step 2 In the **Name** field, enter a name for the group.

Step 3 To add to the group, check the users from the list on the right hand side. Click **Add**.
The checked names move to the lefthand side of the dialog box.

Step 4 To remove users from the group, select the users from the left hand side with Ctrl+click . Click **Remove**.
The selected names move to the righthand side of the dialog box.

Step 5 Click **Submit**.

Add or Remove a User from a User Group

To add or remove a user from a custom User Group, do the following steps:

Note Admin, Operator, Notification, and Security Analyst are system groups and cannot be edited. The user is automatically added to the User Group that corresponds to their role.

Step 1 Select the User Group from the **Select Group** field. The members of the group are displayed.

Step 2 Click **Edit Group**. The User Group dialog box appears.

Step 3 To add to the group, check the users from the list on the right hand side. Click **Add**.
The checked names move to the lefthand side of the dialog box.

Step 4 To remove users from the group, select the users from the left hand side with Ctrl+click . Click **Remove**.
The selected names move to the righthand side of the dialog box.

Step 5 Click **Submit**. You are returned to the **User Management** tab.

Filter by Groups

From the **Select Group** list, select the group. Only the members of the group are displayed.

Promoting Global User Roles on Local Controller

A global “Admin” user can log into the Local Controller and promote a global “System Analyst” or “Operator” user to a higher role. For example, a global “Operator” can be promoted to become an “Admin” or “System Analyst” on the Local Controller. However, his/her role as an “operator” on the Global Controller remains the same because the changes remain on the local controller and do not get pushed up to the Global Controller. Once these users get promoted to a higher role, they can’t be demoted afterward.

Global “Notification” users cannot be promoted given that these users have no login password information.

CHAPTER 10

System Maintenance

Revised: April 5, 2007, OL-13542-01

Much of the system maintenance information for the MARS Appliance is provided exclusively in the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System.*

The MARS Appliance requires little maintenance. To perform maintenance tasks, you can use the CLI or the web interface as needed. Some hardware maintenance tasks require physical access to the MARS Appliance.

This chapter contains the following sections:

- Setting Runtime Logging Levels, page 10-1
- Viewing the MARS Backend Log Files, page 10-2
- Viewing the Audit Trail, page 10-3
- Change the Default Password of the Administrator Account, page 10-3
- Understanding Certificate and Fingerprint Validation and Management, page 10-4
- Hardware Maintenance Tasks—MARS 110R, 110, 210, GC2R, and GC2, page 10-7

For information about upgrading, backing up, and restoring data on the MARS Appliance, see the following sections of the *Install and Setup Guide for Cisco Security Monitoring, Analysis, and Response System*:

- Performing Command Line Administration Tasks, page 6-1
- Checklist for Upgrading the Appliance Software, page 6-7
- Configuring and Performing Appliance Data Backups, page 6-24
- Recovery Management, page 6-34
- Configuring a Standby or Secondary MARS Appliance, page 6-39
- Guidelines for Restoring, page 6-40

Setting Runtime Logging Levels

To set the appliance's runtime logging levels, navigate to **Admin > System Maintenance > Set Runtime Logging Levels**. For typical use, it is best to leave this page set to its defaults.

When you have made your selections, click the **Change Logging Levels** button.

The following log levels are available:

- **Fatal**. Enables fatal logging messages. Fatal messages record very severe error events that will likely lead the application to abort.
- **Error**. Enables error and fatal logging messages. Error messages record error events that might still allow the application to continue running.
- **Warn**. Enables warning, error, and fatal logging messages. Warning messages record potentially harmful situations.
- **Info**. Enables informational, warning, error, and fatal logging messages. Informational messages highlight the progress of the application at coarse-grained level.
- **Debug**. Enables debug, informational, warning, error, and fatal logging messages. Debug messages record fine-grained informational events that are most useful to debug an application.
- **Trace**. Enables trace, debug, information, warning, error, and fatal logging messages. Trace messages record finer-grained informational events than debug messages.

Viewing the MARS Backend Log Files

To view the appliance's log files or to change their levels or source, navigate to **Admin > System Maintenance > View Log Files**.

Figure 10-1 Backend log viewing options

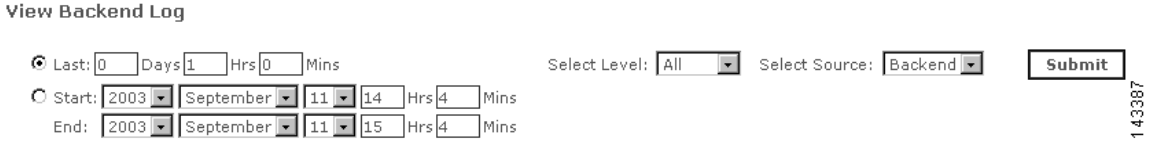

You can view the appliance's back-end logs either by selecting a number of days, hours, and minutes or you can view logs by selecting a start and ending date and time.

You can select the levels of logs that you want. Your choices are: All, Fatal, Error, Warn, Info, and Debug.

You can also choose the source of the files that you want to view. Select either Backend or GUI.

View the Backend Log

Step 1 Click the appropriate radio button:

- **Last**: The present time minus the number of days, hours, and minutes entered.
- **Start/End:** Absolute literal time ranges defined by the date to the minute.

Step 2 Select user, group, etc.

Step 3 Select the source.

Step 4 Click **Submit**.

Viewing the Audit Trail

You can track the activities of the appliance's users by analyzing the appliance's log files. To set the appliance's audit trail logs, navigate to **Admin > System Maintenance > View Audit Trail**. For typical use, it is best to leave this page set to its defaults.

You can view the user audit trails either by selecting a number of days, hours, and minutes, or you can view a specific interval by selecting a start and ending date and time.

View an Audit Trail

Step 1 Click the appropriate radio button:

- Last: DD-HH-MM
- Start/End: YY-MM-DD-HH-MM

Step 2 From the list, select the user or user group.

Step 3 Click **Submit**.

Change the Default Password of the Administrator Account

Good security practices require that you change the default password. We recommend using strong passwords for the MARS Appliance appliances.

Login names and passwords:

- can be alphanumeric characters
- are case sensitive
- can contain special characters (!, @, #, etc.)
- **cannot** contain single or double quotes (‘or “)

Login names can contain up to 20 characters. Passwords can contain up to 64 characters.

To change the default password and setup administrator notification, follow these steps:

Step 1 Click the **Management > User Management** tab.

Step 2 Check the box next to Administrator, and click **Edit**.

Step 3 Enter the new Administrator password and the Administrator e-mail address.

Step 4 Click **Submit**.

Understanding Certificate and Fingerprint Validation and Management

Many reporting devices use certificates or fingerprints to enable secure communications over SSL or SSH respectively. Beginning in 4.2.3, MARS performs a strict check of the certificate or fingerprint of the device or server to which it is attempting to connect.

Note Certificate validation does not follow the convention of presenting the client with a list of certificate authorities and using the selected one to validate individual certificates. Instead, the MARS Appliance compares the certificate presented by the reporting device with a previously stored instance of the certificate. If the two match, the presented certificate is considered valid. This approach allows MARS to validate certificates without knowledge of revocation lists and to operate in a network without an Internet connection.

Three options exist for specifying how MARS should respond during attempts to establish a secure connection. The three options are as follows:

- **Automatically always accept.** This option, which is compatible with previous releases, allows a MARS Appliance to connect to reporting devices regardless of how frequently the certificate or fingerprint changes because MARS automatically accepts and stores the replacement certificate or fingerprint for all devices. However, this option does not provide an opportunity to inspect and authorize the changes to the certificates or fingerprints. When a conflict is detected or when a new certificate or fingerprint is accepted, the event is logged to the internal log. The internal log entry includes the name of the process that detected the conflict and the IP address of the reporting device. The logs can be retrieved by queries and reports. See Monitoring Certificate Status and Changes, page 10-7 for more information on studying these events.
- **Accept first time and prompt on change (default).** This option accepts and stores a new certificate or fingerprint the first time MARS Appliance connects to a device. For subsequent connection attempts, the appliance checks the presented certificate or fingerprint against the stored value. If a conflict is detected, the session is refused unless the new certificate or fingerprint is manually accepted by the administrator. This option enables initial topology discovery to proceed without administrator intervention. Internal system logs of the initial acceptance, conflict detection, and acceptance of new change are created. The internal logs include the name of the process that detected the conflict, the IP address of the reporting device, and the username of the account used to accept the change.

  If, when a change is detected by a web interface process, the session times out before administrative intervention, the communication fails but no internal system log is generated to record the failure to accept the changed certificate or fingerprint. Also, if a back-end process initiates the request, such as auto discovery, then the session attempt always fails and no attempt to obtain administrative acceptance is initiated. In such cases, any data the MARS Appliance would normally ascertain from the device during such a session is not collected. This delay of data retrieval does not apply to syslogs forward to the MARS Appliance by the reporting device and it resumes once the new certificate is accept. The recommended method for manually kicking off the change detections is to use the Test Connectivity or Discover button on the reporting device.
- **Always prompt on new and changed.** This options requires an administrator to manually accept the certificate or fingerprint before MARS can establish the desired communications each time the certificate or fingerprint changes. During changes, the internal log includes the username of the account used to accept the change. If the communication times out before administrative intervention, the communication fails and an internal system log records the failure to accept the changed certificate or fingerprint.

The implication of each option varies based on which MARS service is attempting the connection, not in the enforcement of the option, but in the ability of the service to prompt for immediate administrative intervention. In other words, if the service is a GUI-based services, you will be prompted to accept the changed certificate or fingerprint. If the service is a backend service, the communications with the target device will fail and the event will be logged.

The following services and operations are affected by the global certificate/fingerprint response setting:

- Upgrade (SSL). When MARS uses the HTTPS option to download the upgrade package from the remote server specified on the Admin > System Maintenance > Upgrade page.
- Discovery operation. (SSH)
- Test Connectivity operation. (SSL)
- Cisco IDS, IPS, and IOS IPS router Event Processing (RDEP or SDEE over SSH)
- CSM Policy Query Integration (SSL)
- Qualys Report Discovery. (SSL)
- Graphgen process for mitigation operation (SSH and SSL)
- Device Monitor process for resource monitoring feature (SSH)
- DTM process (SSH)

Setting the Global Certificate and Fingerprint Response

The default response is to accept the certificate or fingerprint the first time MARS attempts to connect to the device, after which if a conflict is detected, then administrative intervention is required to update to the new certificate or fingerprint.

If this option is not the one that you wish to use, you can select from three options. The global setting for the conflict detection responses is located on the **Admin > System Parameters > SSL/SSH Settings** page.

To change the default certificate and fingerprint response, follow these steps:

Step 1 Log into the web interface using an account with Administrative privilege.

Step 2 Click the **Admin > System Parameters > SSL/SSH Settings**.

Step 3 Select one of the following options to define the global behavior that you require:

- Automatically always accept
- Accept first time and prompt when changed
- Always prompt on new and changed

For details on these options, see Understanding Certificate and Fingerprint Validation and Management, page 10-4.

Step 4 Click **Submit**.

Upgrading from an Expired Certificate or Fingerprint

If you have selected a global response option other than Automatically always accept (see Setting the Global Certificate and Fingerprint Response, page 10-5), you will at some time be required to update an expired certificate or fingerprint.

Two options exist for upgrading from an expired certificate or fingerprint. If you are logged in to the web interface when a GUI process detects a certificate or fingerprint conflict, you will be prompted to accept or reject the new value. Otherwise, if you are not logged in or a backend process detects the conflict, you must manually initiate a communication with the device. To determine the list of devices for which you must manually update the certificates or fingerprints, review the Activity: CS-MARS Detected Conflicting Certificates/Fingerprints report (see Monitoring Certificate Status and Changes, page 10-7).

The following procedures explain how to upgrade under the specific circumstances:

- Upgrade a Certificate or Fingerprint Interactively, page 10-6
- Upgrade a Certificate Manually, page 10-6
- Upgrade a Fingerprint Manually, page 10-6

Upgrade a Certificate or Fingerprint Interactively

An interactive upgrade refers to responding to a web interface prompt to update the certificate. This type of upgrade is available when you are logged into the GUI and a process, such as graphgen, prompts you to upgrade a certificate or fingerprint that conflicts with the previously accepted value. Click Yes to accept the new fingerprint or certificate.

Upgrade a Certificate Manually

A manual upgrade allows you to upgrade any certificate at any time due to any reason: session time out during interactive prompt, user error, detection of conflict by a backend process.

To manually upgrade to a new certificate, follow these steps:‘

Step 1 Log into the web interface using an account with Administrative privilege.

Step 2 Select the reporting device on the Admin > System Setup > Security and Monitor Devices page for which MARS has detected a certificate conflict. and click **Edit**.

Step 3 Click **Test Connectivity**.

The dialog box displays stating “Do you want to accept following certificate for the device named: <device_name>?.

Step 4 Verify the certificate value.

Step 5 If the value is correct. click **Yes**.

Upgrade a Fingerprint Manually

A manual upgrade allows you to upgrade any fingerprint at any time due to any reason: session time out during interactive prompt, user error, detection of conflict by a backend process.

To manually upgrade a fingerprint, follow these steps:

Step 1 Log into the web interface using an account with Administrative privilege.

Step 2 Select the reporting device on the Admin > System Setup > Security and Monitor Devices page for which MARS has detected a fingerprint conflict and click **Edit**.

Step 3 Click **Discover**.

The dialog box displays stating “Do you want to accept following fingerprint for the device named: <device_name>?.

Step 4 Verify the fingerprint value.

Step 5 If the value is correct, click **Yes**.

Monitoring Certificate Status and Changes

To support the certificate management features in MARS, the following system inspection rule exists:

- **System Rule: CS-MARS Failure Saving Certificates/Fingerprints.** This inspection rule indicates that MARS has failed to save a new or changed device SSL certificate or SSH key fingerprint based on either explicit user action or automatic accept as specified on the SSL/SSH Settings page.

In addition, the following reports appear under the System: CS-MARS Issue category.

- Activity: CS-MARS Accepted New Certificates/Fingerprints
- Activity: CS-MARS Accepted Conflicting Certificates/Fingerprints
- Activity: CS-MARS Detected Conflicting Certificates/Fingerprints
- Activity: CS-MARS Accepted New Certificates/Fingerprints'
- Activity: CS-MARS Failure Saving Certificates/Fingerprints
- Activity: CS-MARS Device Connectivity Errors

Hardware Maintenance Tasks—MARS 110R, 110, 210, GC2R, and GC2

- Field Replaceable Units, page 10-7
- Removing and Replacing the Front Bezel, page 10-8
- Removing the Chassis Cover, page 10-9
- Replacing the RAID Battery Backup Unit, page 10-10
- Hard Drive Troubleshooting and Replacement, page 10-14
- Hot-swapping a Power Supply Unit, page 10-25
- Diagnostic Beep Codes, page 10-26

Field Replaceable Units

Table 10-1 lists the field replaceable units (FRUs) supported for the MARS 110R, 110, 210, GC2R, and GC2 appliances.

Table 10-1 List of Field Replaceable Units for the Cisco Security MARS Appliances 5.X

FRU Description	FRU Part Number
SR2500 (Driskill 2) 750 Watt Power Supply Module	CS-MARS-D750-PS =
500 GB SATA-IO HDD (Seagate - Barracuda ES)	CS-MARS-S500-HD =
750 GB SATA-IO HDD (Seagate - Barracuda ES)	CS-MARS-S750-HD =
RAID Controller Back-Up Battery Unit	CS-MARS-X10-BB =
Rack-mount Kit	CS-MARS-X10-RAIL=

Removing and Replacing the Front Bezel

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

You must remove the front bezel to access the DVD ROM, hard drives, and control panel buttons. The bezel does not lock.

To remove the bezel, pull the bezel from the appliance, as shown in Figure 10-2.

To replace the bezel, line up the center notch on the bezel with the center guide on the rack handles, then push the bezel onto the front of the MARS Appliance until it clicks into place.

Figure 10-2 Removing the Front Bezel

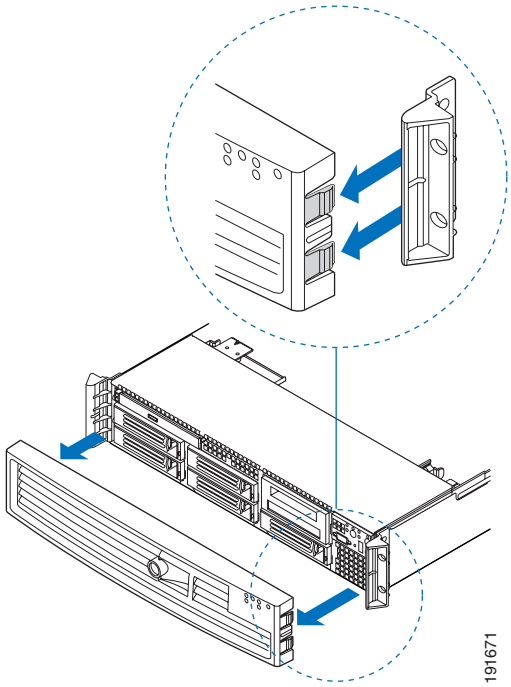

Removing the Chassis Cover

The MARS Appliance must be operated with the chassis cover in place to ensure proper cooling. Remove the top cover to add or replace components inside of the appliance. Before removing the chassis cover, power down the appliance and unplug all peripheral devices and the AC power cables.

Note A nonskid surface or a stop behind the MARS Appliance may be needed to prevent the MARS Appliance from sliding on your work surface.

Removing the Chassis Cover

Step 1 Observe all safety and ESD precautions. See "Safety Information" section on page 10-27."

Step 2 Turn off the appliance.

Step 3 Disconnect the AC power cords.

Step 4 Remove the safety screw if it is installed, as shown in callout A of Figure 3.

Step 5 While holding in the blue button at the top of the MARS Appliance (callout B), slide the top cover back until it stops, as shown in callout C of Figure 3.

Step 6 Insert your finger in the notch shown in callout D of Figure 3, then lift the cover upward to remove it.

Figure 3 Removing the MARS Appliance Cover

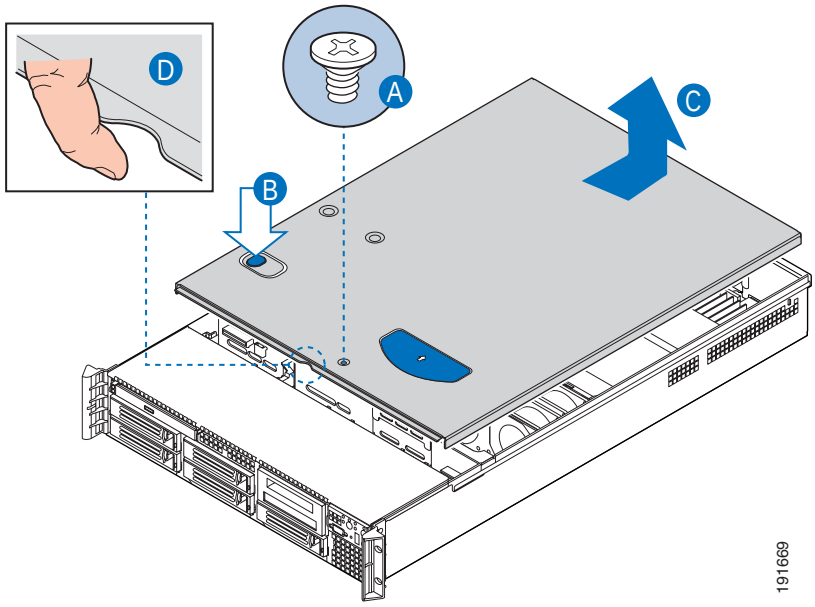

End of Procedure

Replacing the Chassis Cover

Step 1 Place the cover over the MARS Appliance so that the side edges of the cover sit just inside the MARS Appliance sidewalls.

Step 2 Slide the cover forward until it clicks into place.

Step 3 (Optional) Insert the safety screw at the center of the top cover if required.

Step 4 Reconnect the AC power cords.

End of Procedure

Replacing the RAID Battery Backup Unit

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

RAID Controller Back-Up Battery
Part number: CS-MARS-X10-BB=

The RAID Backup Battery Unit (RAID BBU) prevents RAID data loss by preserving data held in the RAID cache module during a power outage. The RAID BBU can provide up to 72 hours of battery power until the system power is restored.

The RAID BBU requires 24 hours to fully charge from when the appliance is first powered on, and is continually charged thereafter from the system power. The total charge capacity of the battery degrades over time. The **show healthinfo** CLI command reports the relative charge state of the RAID BBU.

There is a direct relationship between the relative charge and the battery backup time ($100\%_{charge} = 72_{hours}$). A 100 percent charge provides 72 hours RAID cache protection. Similarly, a 75 percent charge provides 54 hours of protection ($100\%_{charge} * .75 = 72_{hours} * .75$).

Make sure there is sufficient charge to provide RAID cache protection for the total probable hours the MARS Appliance could be without system power. For example, a 90.3% charge (65 hours) would allow 2 hours to manually restore system power if a total power outage occurred in an unattended facility between 17h00 Friday to 8h00 Monday (63 hours).

Example 10-1 displays BBU status information in an excerpt of the **show healthinfo** CLI command.

Example 10-1 RAID Battery Backup Unit show healthinfo Command Output

```
[pnadmin]$ show healthinfo
<snip>
BBU information :
Relative state of charge : 93 %
Full charge capacity : 920 mAh
Remain capacity : 858 mAh

<snip>
```

Summary of steps required to replace the RAID BBU:

1. Remove the chassis cover.
2. Remove the large air baffle.
3. Remove the RAID BBU.
4. Install the replacement RAID BBU.
5. Replace the large air baffle.
6. Replace the chassis cover.

Procedure to Replace the Raid Battery Backup Unit

Remove the Cover

Step 1 Observe all safety and ESD precautions. See "Safety Information" section on page 10-27.

Step 2 Power down the appliance and unplug all the AC power cables.

Step 3 Remove the chassis cover. For instructions, see the "Removing the Chassis Cover" section on page 10-9.

Remove the Large Air Baffle

Step 4 Write down how the cables are routed over and under the air baffle (if any). You will need to re-route these cables.

Step 5 Pull up on the air baffle to remove it, as shown in Figure 4. You may need to remove or hold cables out of the way.

Figure 4 Removing the Large Air Baffle

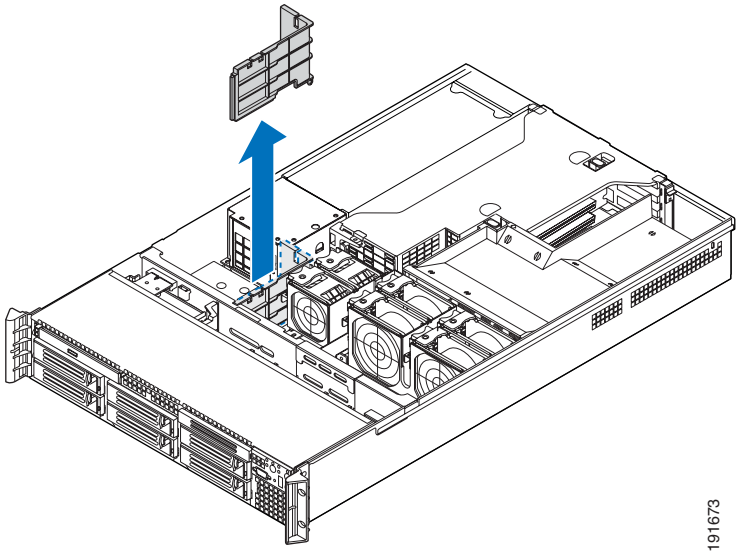

Remove the RAID BBU

Step 6 Disconnect the cable from the rear of the RAID battery backup unit and the mid-plane board as shown in callout A of Figure 5.

Step 7 Slide the RAID battery backup unit forward and lift it up from the appliance, as shown in callout B of Figure 5.

Figure 5 Removing the RAID Battery Backup Unit

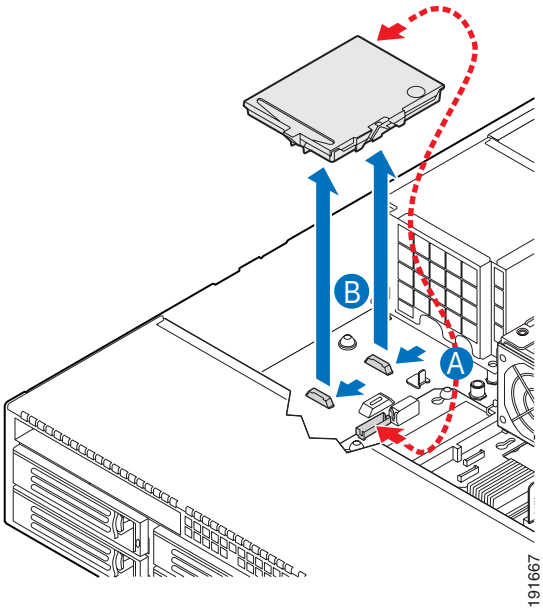

Install the Replacement RAID BBU

Step 8 Insert the RAID battery backup unit into the appliance and slide it back until it locks into place as shown in callout A of Figure 6.

Step 9 Attach the cable from the rear of the RAID battery backup unit to the mid-plane board as shown in callout B of Figure 6.

Figure 6 Installing the RAID Battery Backup Unit

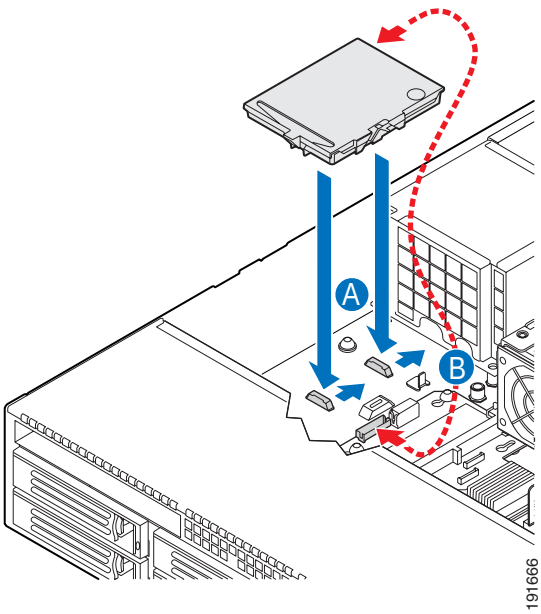

Replace the Large Air Baffle

Step 10 Lower the baffle into the appliance and snap it into the appliance board standoff.

Make sure to route the cables beneath the air baffle as were recorded in Step 4.

Replace the chassis cover.

Step 11 Replace the chassis cover.

Step 12 Reconnect the AC power cables to the power supplies.

End of Procedure

Hard Drive Troubleshooting and Replacement

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances and contains the following subsections:

- Hard Drive Status LEDs, page 10-14
- Partition Checking, page 10-14
- Overview of RAID 10 Subsystem, page 10-15
- Hotswapping Hard Drives, page 10-15
- Viewing RAID Array Status with the raidstatus CLI Command, page 10-17
- Procedure to Hotswap a Hard Drive, page 10-21
- Hotswap CLI Example, page 10-22
- Replacing a Hard Drive in the Hard Drive Carrier, page 10-24

Note Hard drives are also termed HDDs throughout this section.

Cisco Security MARS HDDs are Cisco field replaceable units (FRUs). The following table provides the correct FRU part numbers for your MARS appliance.

MARS Model	Hard Drive Descriptions and Part Numbers
CS-MARS-110R-K9 CS-MARS-110-K9	500 GB SATA-IO HDD Part number: **CS-MARS-S500-HD=**
CS-MARS-210-K9 CS-MARS-GC2R-K9 CS-MARS-GC2-K9	750 GB SATA-IO HDD Part number: **CS-MARS-S750-HD=**

Note Hard drives can consume up to 17 watts of power each. Drives are specified to run at a maximum ambient temperature of 45 °C.

Hard Drive Status LEDs

Each HDD has a status LED. A flickering green light indicates activity. The control panel has a status LED that flickers with any HDD activity.

Partition Checking

The appliance automatically runs checks on HDD partitions after the system has been re-booted 25–30 times, or if the appliance has not been re-booted in 180 days.

Overview of RAID 10 Subsystem

This section pertains to the following MARS Appliances equipped with a Serial ATA RAID controller card:

- CS-MARS-110R-K9
- CS-MARS-110-K9
- CS-MARS-210-K9
- CS-MARS-GC2R-K9
- CS-MARS-GC2-K9

The MARS RAID controller cards operate the hard drives in a RAID 10 configuration, also called RAID 1+0 because it combines the data handling techniques of RAID 1 and RAID 0. For additional information on RAID concepts and terminology, access the following URL: http://en.wikipedia.org/wiki/RAID

RAID 0 Data Striping

In a MARS RAID 10 configuration, half the total number of drives are arrayed as a single logical drive, wherein a data block is distributed across all of the physical drives in the logical drive using RAID 0 striping techniques. Data striping results in better performance for a data intensive application such as MARS, because hard drive random access times are minimized when data is read and written simultaneously from more than one physical hard drive.

RAID 1 Mirroring and Subunits

Half the number of drives in the MARS RAID 10 array mirror the RAID 0 virtual drive. Each physical drive in the RAID 0 array is mirrored by an identical physical drive using RAID 1 techniques. Data written to one of the drives within the RAID 0 array is simultaneously written to its dedicated RAID 1 partner, thereby providing fault tolerance through data redundancy. The RAID 1 hard drive pairs are listed in Table 10-4.

Rebuilding a Degraded Array

Either drive in a RAID 1 pair can serve in place of its partner should either drive become degraded (unavailable, physically inoperative, or data corrupted). A physical drive degraded but still physically operative can be rebuilt from the data of its undegraded partner and rejoin the array. An inoperative physical drive can be replaced with an operative one which is then rebuilt to join the array.

When any physical drive of the RAID 10 array is degraded, the entire array is considered degraded. While the array still functions, it is not working to its optimal throughput or redundancy capacity.

In a degraded RAID 10 array, data destined for a degraded physical drive is written to available space on the RAID 1 partner until the degraded drive can be rebuilt or replaced. Degraded drives are rebuilt in sequence, one rebuilding process must complete before the next process can begin. Between 200 and 300 minutes are required to rebuild a MARS subunit..

Hotswapping Hard Drives

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

An HDD can be hotswapped, that is, replaced without rebooting the MARS appliance. The hotswap actions can be summarized in the following five steps. The detailed procedure is in the section, Procedure to Hotswap a Hard Drive.

1. Establish a console connection to the MARS appliance.

2. Enter the **raidstatus** command to determine the status and the chassis HDD slot number of the HDD to hotswap.
3. Execute a **hotswap remove** *disk* command then remove the HDD.
4. Execute a **hotswap add** *disk* command then insert the replacement HDD.
5. Enter the **raidstatus** command to monitor the progress of the replacement HDD as it is rebuilt.

Use the **raidstatus** CLI command to view the status of the RAID array (virtual disk) and of the individual HDDs. Table 10-2 lists the status conditions that require an HDD to be hotswapped. These status conditions cause MARS to send an email alert to the adminstrator.

Caution Always use the **hotswap remove** *disk* CLI command before you remove a hard drive and **hotswap add** *disk* before you insert a hard drive. The *disk* argument is the hard drive slot number. Use the **hotswap list all** command to view hard drives that can be hotswapped.

The rebuilding process duration is between 200 and 300 minutes.

Note To match original performance, hotswapped HDDs should be the same make, model and size as the original HDDs.

Caution The RAID 10 array will not function if both both HDDs of any RAID 1 pair are removed or corrupted.

Table 10-2 HDD Actions for MARS 110R, 110, 210, GC2R, and GC2

Hard Drive Status	Possible Cause	Recommended Action
Failed	Unrecoverable error on previously operative HDD.	Hotswap with a new HDD.
Offline	The **hotswap remove** command was executed for this HDD.	Execute a **hotswap add** on the HDD if the HDD is otherwise good.
Unconfigured Good	An online HDD was removed and inserted without executing a **hotswap** command sequence.	Execute a **hotswap remove** and **hotswap add** on the HDD.
Unconfigured Bad	An online HDD was removed or inserted without executing a **hotswap** sequence and the HDD has a media error.	Hotswap with a new HDD.
N/A	The HDD slot is empty.	Insert a new HDD with the **hotswap add** command.

Failed Hard Drive Alert

MARS sends and email alert when the hard drive status changes from Online to Failed, Offline, Unconfigured Good, Unconfigured Bad, or N/A. Example 10-2 displays the contents of an e-mail alert sent to the administrator for a failed HDD. In the alert, the DISK number is the same as the chassis HDD slot or **raidstatus** PD number.

Example 10-2 MARS Hard Drive Replacement Alert

```
From: csmars-system.SJ-LC-86@cisco.com [mailto:csmars-system.SJ-LC-86@cisco.com]
Sent: Tuesday, March 20, 2007 12:22 PM
To: Liu Bang (liubang)
Subject: Hard disk failure (host: SJ-LC-86, disk No.: 4)
Importance: High

Hard disk failure: RAID error
-------------------------
 HOST   :   SJ-LC-86
 DISK   :   4
 STATUS :   FAILED
 MODEL  :   ST3750640NS
 SIZE   :   750GBs

Hard disk 4 on adapter a0 has failed. As a result, the disk array on adapter a0 is running
in degrade mode and is no longer fault tolerant. Please replace hard disk 4 as soon as
possible. Instructions for doing so can be found in the user's manual.
```

Viewing RAID Array Status with the raidstatus CLI Command

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

Example 10-3 displays the output of the **raidstatus** command executed on a Local Controller 210. In this example, HDDs p2 and p5 where hotswapped and are in the final stages of being rebuilt. Table 10-3 describes the **raidstatus** command output fields.

Example 10-3 MARS raidstatus CLI Command Output for MARS 110R, 110, 210, GC2R, and GC2

```
[pnmars]# raidstatus
Adapter Information:
------------------------------------------------------
Product Name     : Intel(R) RAID Controller SROMBSAS18E
Firmware Version :  1.02.00-0119
BIOS Version     : MT25

Adapter RaidType        Status     Stripe      Size             Cache
-------------------------------------------------------------------------------------------
a0      Raid-10         Degraded   64kB        2097151MB        Enabled


PD      Status  Size & Block                     Model                              Serial#
-------------------------------------------------------------------------------------------
p0      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD09ZNT
p1      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD07ZYK
p2      Rebuild 715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD091BZ
p3      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD09E3A
p4      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD0A03B
p5      Rebuild 715404MB [0x575466f0 Sectors]   ATA     ST3750640NS     E          3QD0A04G

Rebuild Progress on Device at Enclosure 20, Slot 2 Completed 71% in 279 Minutes.
Rebuild Progress on Device at Enclosure 20, Slot 5 Completed 60% in 259 Minutes.
===========================================================================================
```

Table 10-3 raidstatus CLI command for MARS 110R, 110, 210, GC2R, and GC2

Output Field	Description
RAID Controller Information Fields	
Product Name	RAID controller manufacturer and serial number
Firmware Version : 1.02.00-0119	Indicates version of the RAID controller firmware
BIOS Version : MT25	Indicates the RAID BIOS version. This is different from the system BIOS version.
RAID Array Information Fields (The RAID 10 Virtual Drive Information)	
Adapter	Identifier for the physical RAID controller.
RaidType	RAID Level of Array. MARS is always RAID 10.
Status	The current state of the RAID 10 virtual drive. • Optimal—All component HDDs are operating as configured. • Degraded—At least one of the component HDDs has failed or is offline. Troubleshooting is advised to prevent possible data loss. • Offline—The array is not available or is unusable.
Stripe	The MARS RAID 10 data stripe is always 64 KB.
Size	The available storage in megabytes of the RAID array.
Cache	The MARS RAID 10 array cache is always enabled.
Individual Hard Drive Information Fields	

Table 10-3 raidstatus CLI command for MARS 110R, 110, 210, GC2R, and GC2 (continued)

Output Field	Description
PD	p0–p5. The physical hard drive numbers.
Status	The current state of the physical HDD. • **Online**—The HDD is functioning normally within the RAID 10 array. • **Rebuild**—The HDD is being reimaged from its RAID 1 partner to restore full redundancy to a the virtual disk. The RAID 10 array efficiency is not yet optimal. • **Failed**—The HDD originally was Online, but now has an unrecoverable error. An email alert is sent to the administrator. • **Offline**—The HDD was removed by executing a **hotswap remove** command, but the HDD was not physically removed from the slot**.** An email alert is sent to the administrator. • **Unconfigured Good**—The HDD is usable, but the RAID information is out of sync with the RAID 1 partner. An email alert is sent to the administrator. • **Unconfigured Bad**— The firmware detected a media error on the hard drive. An online HDD was probably removed or inserted without executing a **hotswap** sequence and the HDD now has a media error. An alert is sent to the administrator. • **N/A**—There is no HDD in the slot. An email alert is sent to the administrator.
Size & Block	Size of the usable storage on the HDD
Model	The model number of the physical HDD
Serial#	The serial number of the physical HDD
Progress Messages	
`Rebuild Progress on Device at Enclosure 20, Slot 2 Completed 71% in 279 Minutes.`	Indicates the status, elapsed rebuilding time, and slot number of each physical drive being rebuilt.

Hard Drive Slot Number Diagram

Figure 10-7 shows the chassis HDD slot numbers of the MARS 110R, 110, 210, GC2R, and GC2. Table 10-4 shows how slot numbers correspond to PD numbers used in the **raidstatus** CLI.

Figure 10-7 HDD Slot Numbers—MARS 110R, 110, 210, GC2R, and GC2

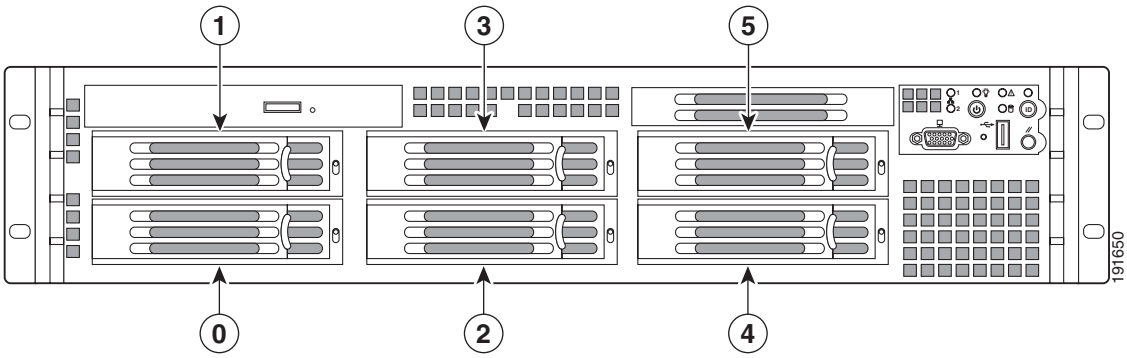

Table 10-4 Mapping HDD Slot Number to raidstatus CLI Command PD number—MARS 110R, 110, 210, GC2R, and GC2

MARS Appliance	Storage Capacity[1]	Chassis HDD Slot to PD Numbers	RAID 1 Pairs
MARS 110R, 110	1.5TB RAID 10 Hot-swappable Front accessible 6 X 500GB SATA-IO HDD 3.0Gbps, 7200 RPM, 16MB Buffer	Slot **0** is **p0** Slot **1** is **p1** Slot **2** is **p2** Slot **3** is **p3** Slot **4** is **p4** Slot **5** is **p5**	Slot **0** and Slot **1** Slot ***2*** and Slot **3** Slot **4** and Slot **5**
MARS 210, GC2R, GC2	2.0TB[2] RAID 10 Hot-swappable Front accessible 6 X 750GB SATA-IO HDD, 3.0Gbps, 7200 RPM, 16MB Buffer		

1. The stated storage capacity is the sum of the rated capacity of all the hard drives and does reflect bytes reserved for the RAID overhead on each drive.
2. Although there is a total of 4.5 TB storage, RAID 10 has a maximum size configuration of 2 TB Redundant, or 4 TB

Procedure to Hotswap a Hard Drive

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

The **hotswap** command requires the chassis slot number of the HDD. The **raidstatus** CLI command refers to physical drive (PD) numbers. The PD number is the chassis slot number. For example, p0 is slot 0. See Figure 10-7 to determine the physical location of the slot in the chassis.

To hotswap an HDD, complete the following steps:

Step 1 Remove the front bezel. See the "Removing and Replacing the Front Bezel" section on page 10-8.

Step 2 Establish a console connection with MARS.

Step 3 Identify the slot number of the HDD to replace with the **raidstatus** command.

Step 4 Enter **hotswap remove** *disk*.

A message informs you that it is safe to remove the HDD.

Note Make sure that you remove the correct physical HDD. If you remove the wrong one accidently then reinsert it, that HDD will register as Unconfigured Good.

Step 5 Pull out the black lever of the hard drive carrier and slide the carrier from the chassis as shown in callout B of Figure 8.

Figure 8 Removing a Hard Drive from the MARS Appliance Chassis

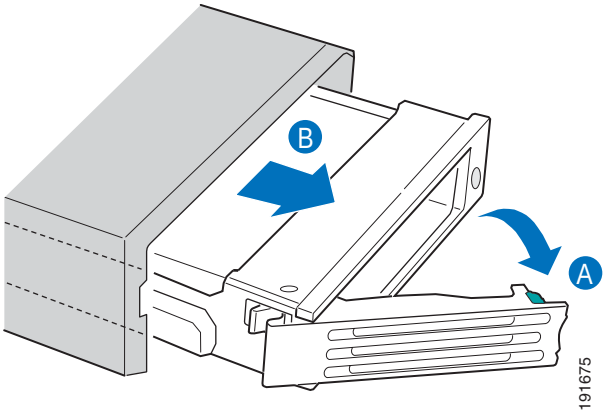

The **raidstatus** command should now report the slot status as N/A.

Step 6 At the CLI prompt, enter **hotswap add** *disk* . Be sure to use the same slot number as in Step 4.

Step 7 With the black lever in the fully open position, slide the replacement HDD and carrier into the chassis. The green latch at the front of the drive carrier must be to the right. Do not push on the black drive carrier lever until the lever begins to close by itself.

Step 8 When the black drive carrier lever begins to close by itself, push it closed to lock the drive assembly into place.

A console message informs you that the HDD (disk) is added successfully (to the logical array).

Step 9 Replace the front bezel. (See "Removing and Replacing the Front Bezel" section on page 10-8).

Step 10 From the CLI, enter **raidstatus** to verify that the HDD is being rebuilt into the RAID array.

The status message indicates the progress of the added HDD.

The rebuilding process can last from 200 to 300 minutes.

End of Procedure, Hotswapping Hard Drives.

Hotswap CLI Example

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

The following CLI output example hotswaps an HDD in slot 2 of a MARS 110.

Example 10-4 Hotswap Procedure, CLI Output Example

In the following example, a hard drive is hotswapped in slot 5 of a MARS 210. The hard drive status is verified with the **raidstatus** command:

```
[pnadmin]$ hotswap list all
Hardware RAID is found with 6 disks!
Disks available to be hotswapped:
        Disk 0 --- p0
        Disk 1 --- p1
        Disk 2 --- p2
        Disk 3 --- p3
        Disk 4 --- p4
        Disk 5 --- p5

[pnadmin]$ hotswap remove 5

Adapter: 0: EnclId-14 SlotId-5 state changed to OffLine.
Disk 5 can now be safely removed from the system.


[pnadmin]$ raidstatus
Adapter Information:
------------------------------------------------------
Product Name     : Intel(R) RAID Controller SROMBSAS18E
Firmware Version : 1.03.00-0211
BIOS Version     : MT30

Adapter RaidType        Status          Stripe  Size            Cache
-------------------------------------------------------------------------------
a0      Raid-10         Degraded        64kB    2097151MB       Enabled

PD      Status  Size & Block                    Model                    Serial#
-------------------------------------------------------------------------------
p0      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD09EEZ
p1      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD09CQT
p2      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD094KY
p3      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD08NZX
p4      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD09EWP
p5      Offline 715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E    3QD06AQ2


[pnadmin]$ hotswap add 5

Started rebuild progress on device(Encl-14 Slot-5)
Disk 5 has been successfully added to RAID
[pnadmin]$ raidstatus
Adapter Information:
------------------------------------------------------
```

```
Product Name     : Intel(R) RAID Controller SROMBSAS18E
Firmware Version :  1.03.00-0211
BIOS Version     : MT30

Adapter RaidType        Status          Stripe  Size            Cache
-------------------------------------------------------------------------------
a0      Raid-10         Degraded        64kB    2097151MB       Enabled

PD      Status  Size & Block                    Model                      Serial#
-------------------------------------------------------------------------------
p0      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD09EEZ
p1      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD09CQT
p2      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD094KY
p3      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD08NZX
p4      Online  715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD09EWP
p5      Rebuild 715404MB [0x575466f0 Sectors]   ATA     ST3750640NS E      3QD06AQ2

Rebuild Progress on Device at Enclosure 14, Slot 5 Completed 17% in 32 Minutes.
```

Replacing a Hard Drive in the Hard Drive Carrier

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

To match original performance, HDDs should be the same make, model and size as the original hard drives.

Step 1 Remove the four screws that attach the hard drive or empty retention device to the drive carrier, as shown in callout A of Figure 9.

Two screws are at each side of the retention device or the hard drive. Store the plastic retention device for future use.

Figure 9 Removing Hard Drive or Retention Device from Drive Carrier (Retention Device Shown Here)

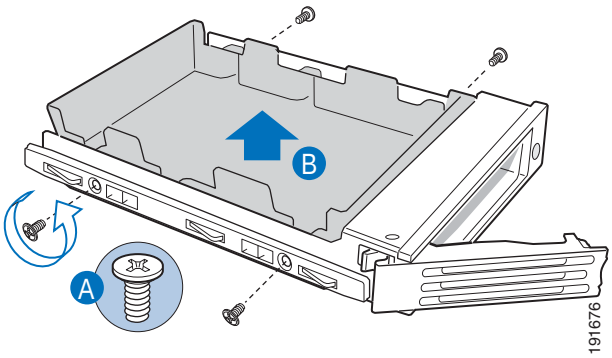

Step 2 Remove the hard drive from its wrapper and place it on an antistatic surface.

Step 3 With the hard drive circuit-side down, position the connector end of the drive so that it is facing the rear of the drive carrier, as shown in callout A of Figure 10.

Step 4 Align the holes in the drive to the holes in the drive carrier and attach it to the carrier with the screws that were attached to the plastic retention device, as shown in callout B of Figure 10.

Figure 10 Installing a Hard Drive into a Carrier

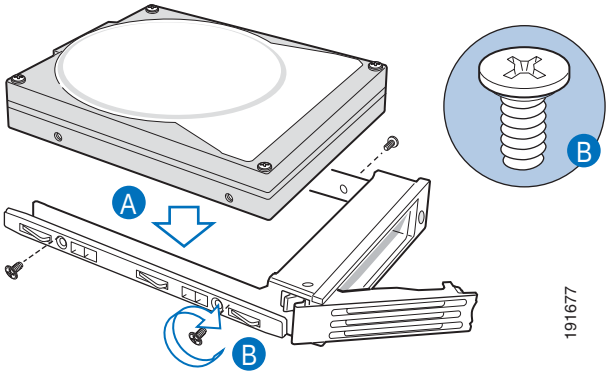

End of Procedure

Hot-swapping a Power Supply Unit

SR2500 (Driskill 2) 750 Watt Power Supply Module
Part number: CS-MARS-D750-PS =

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

Up to two power supply modules may be on a single AC line. The lower power supply (PS1) supplies most of the power requirements. The upper power supply (PS2) is the redundant power supply.

A power supply module can be replaced without powering down the system (hotswapped). Example 10-5 is an excerpt of the **show healthinfo** CLI command. The power supply unit should be evaluated for hotswapping if its status is "down."

Example 10-5 Power Supply Status in the show healthinfo CLI Command.

```
[pnadmin]$ show healthinfo
<SNIP>
Power Supply            Value    Status
----------------------------------------
PS1 AC Current  2.36 Amps        ok
PS2 AC Current  0.12 Amps        ok
PS1 +12V Current         21 Amps ok
PS2 +12V Current         0 Amps  ok
PS1 +12V Power  248 Watts        ok
PS2 +12V Power  0 Watts          ok
PS1 Status      0x01             ok
PS2 Status      0x09             ok

<SNIP>
```

To hotswap a power supply, do the following:

Step 1 Observe all safety and ESD precautions. See "Safety Information" section on page 10-27."

Step 2 Unplug the AC power cord of power supply to be replaced.

Step 3 Release the latch (callout A) and remove the power supply by pulling on the handle (callout B) as shown in Figure 11.

Figure 11 Removing Power Supply Module from the MARS Appliance

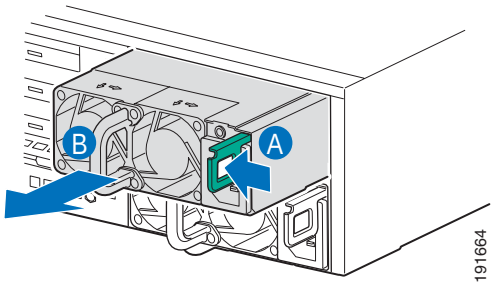

Step 4 Insert the replacement power supply module into the power supply cage until it clicks into place.

Step 5 Connect the AC power cord to the replacement power supply.

End of Procedure

Diagnostic Beep Codes

This section pertains only to the MARS 110R, 110, 210, GC2R, and GC2 appliances.

Table 10-5 lists Power-on Self Test (POST) error beep codes. Prior to system Video initialization, BIOS uses these beep codes to signal error conditions. The beep code is followed by a user visible code on the POST Progress LEDs (not shown). Beep codes are sounded each time the problem is discovered, such as on each power-up attempt, but they are not sounded continuously. The beep code sequence is read left to right. For example, 4–7 represents four beeps followed by seven beeps.

Table 10-5 POST Error Beep Codes

Number of Beeps	Error Message	Description
1, 2, or 3	Memory Error	Fatal memory error. Reseat the memory or replace the DIMMs with known good modules.
6	BIOS Error	The system has detected a corrupted BIOS in the flash part, and is rolling back to the last good BIOS.
4–7 or 9–11	System Error	Fatal error indicating a possible serious system problem.

Safety Information

These saftety instructions apply to all Cisco Security Monitoring, Analysis, and Response System models

To reduce the risk of bodily injury, electrical shock, fire, and equipment damage, read this section and observe all warnings and precautions before maintaining your Cisco Security MARS appliance.

Intended Application Uses

This product was evaluated as Information Technology Equipment (ITE), which may be installed in offices, schools, computer rooms, and similar commercial type locations. The suitability of this product for other product categories and environments (such as medical, industrial, residential, alarm systems, and test equipment), other than an ITE application, may require further evaluation.

Equipment Handling Practices

Reduce the risk of personal injury or equipment damage:

- Conform to local occupational health and safety requirements when moving and lifting equipment.
- Use mechanical assistance or other suitable assistance when moving and lifting equipment.
- To reduce the weight for easier handling, remove any easily detachable components.

Power and Electrical Warnings

Caution The power button, indicated by the stand-by power marking, DOES NOT completely turn off the system AC power, 5V standby power is active whenever the system is plugged in.

Warning **This unit might have more than one power supply connection. All connections must be removed to de-energize the unit.** Statement 1028

Caution Do not attempt to modify or use an AC power cord if it is not the exact type required. A separate AC cord is required for each system power supply.

Caution Some power supplies inCisco Security MARS appliances use Neutral Pole Fusing. To avoid risk of shock use caution when working with power supplies that use Neutral Pole Fusing.

Caution The power supply in this product contains no user-serviceable parts. Do not open the power supply. Hazardous voltage, current and energy levels are present inside the power supply. Return to manufacturer for servicing.

Caution When replacing a hot-plug power supply, unplug the power cord to the power supply being replaced before removing it from the server.

Caution To avoid risk of electric shock, turn off the server and disconnect the power cord, telecommunications systems, networks, and modems attached to the appliance before opening.

Power Cord Warnings

If an AC power cord was not provided with your product, purchase one that is approved for use in your country.

To avoid electrical shock or fire, check the power cords that will be used with the product as follows:

- Do not attempt to modify or use the AC power cord(s) if they are not the exact type required to fit into the grounded electrical outlets
- The power cord(s) must meet the following criteria:
- The power cord must have an electrical rating that is greater than that of the electrical current rating marked on the product.
- The power cord must have safety ground pin or contact that is suitable for the electrical outlet.
- The power supply cord(s) is/are the main disconnect device to AC power. The socket outlet(s) must be near the equipment and readily accessible for disconnection.
- The power supply cord(s) must be plugged into socket-outlet(s) that is /are provided with a suitable earth ground.

System Access Warnings

To avoid personal injury or property damage, the following safety instructions apply whenever accessing the inside of the product:

- Turn off all peripheral devices connected to this product.
- Turn off the system by pressing the power button to off.
- Disconnect the AC power by unplugging all AC power cords from the system or wall outlet.
- Disconnect all cables and telecommunication lines that are connected to the system.
- Retain all screws or other fasteners when removing access cover(s). Upon completion of accessing inside the product, refasten access cover with original screws or fasteners.
- Do not access the inside of the power supply. There are no serviceable parts in the power supply. Return to manufacturer for servicing.

- Power down the appliance and disconnect all power cords before adding or replacing any non hot-plug component.
- When replacing a hot-plug power supply, unplug the power cord to the power supply being replaced before removing the power supply from the appliance.

Caution If the appliance has been running, any installed processor(s) and heat sink(s) may be hot. Unless you are adding or removing a hot-plug component, allow the appliance to cool before opening the covers. To avoid the possibility of coming into contact with hot component(s) during a hot-plug installation, be careful when removing or installing the hot-plug component(s).

Caution To avoid injury do not contact moving fan blades. If your system is supplied with a guard over the fan, do not operate the appliance without the fan guard in place.

Rack Mount Warnings

The equipment rack must be anchored to an unmovable support to prevent it from tipping when a server or piece of equipment is extended from it. The equipment rack must be installed according to the rack manufacturer's instructions.

Install equipment in the rack from the bottom up, with the heaviest equipment at the bottom of the rack.

Extend only one piece of equipment from the rack at a time.

You are responsible for installing a main power disconnect for the entire rack unit. This main disconnect must be readily accessible, and it must be labeled as controlling power to the entire unit, not just to the appliance.

To avoid risk of potential electric shock, a proper safety ground must be implemented for the rack and each piece of equipment installed in it.

Electrostatic Discharge (ESD)

Caution ESD can damage disk drives, boards, and other parts. We recommend that you perform all procedures at an ESD workstation. If one is not available, provide some ESD protection by wearing an antistatic wrist strap attached to chassis ground -- any unpainted metal surface -- on your server when handling parts. Always handle boards carefully. They can be extremely sensitive to ESD. Hold boards only by their edges. After removing a board from its protective wrapper or from the server, place the board component side up on a grounded, static free surface. Use a conductive foam pad if available but not the board wrapper. Do not slide board over any surface.

Battery Replacement

Caution Do not attempt to recharge a battery. Do not attempt to disassemble, puncture, or otherwise damage a battery.

Warning **There is the danger of explosion if the battery is replaced incorrectly. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.** Statement 33

Cooling and Airflow

Carefully route cables as directed to minimize airflow blockage and cooling problems. For proper cooling and airflow, operate the system only with the chassis covers installed. Operating the system without the covers in place can damage system parts.

To install the covers:

- Check first to make sure you have not left loose tools or parts inside the system.
- Check that cables, add-in boards, and other components are properly installed.
- Attach the covers to the chassis according to the product instructions.

Laser Peripherals or Devices

To avoid risk of radiation exposure and/or personal injury:

- Do not open the enclosure of any laser peripheral or device
- Laser peripherals or devices have are not user serviceable
- Return to manufacturer for servicing

APPENDIX A

Cisco Security MARS XML API Reference

This appendix provides resources for creating XML applications that integrate Cisco Security MARS XML data into third-party applications.

XML Schema Overview

The XML schema are written in conformance with the standard World Wide Web Consortium (W3C) XML schema language. A schema by definition, describes all data and data structures required to create your application. Many XML development environments provide enough capability to view the schema in a way that you can identify all components, their relationships, constraints, attributes, annotations, and usage guidelines at a glance. Some applications generate hyperlinked reference documentation. By providing sufficient documentation and annotation tags within the schemas, Cisco supports such documentation generating applications.

Table A-1 lists resources for XML development.

Table A-1 XML Resources

Resource Description	URL
W3C XML Schema standards forum with resource links	http://www.w3.org/XML/Schema
General XML description with resource links	http://en.wikipedia.org/wiki/Xml
Online XML Tutorials	http://www.w3schools.com/xml/default.asp

XML Incident Notification Data File and Schema

XML incident notification sends an email notification of an incident with an attached XML data file. The XML data file contains all incident details that can be viewed on the GUI except for Path/Mitigation data. The XML data file can be sent as a plain-text file or as a compressed gzip file. The filename is constructed with the incident ID number, for example `CS-MARS-Incident-13725095.xml`. The compressed version of the same data file would be `CS-MARS-Incident-13725095.xml.gz`

An XML application can be written to parse and extract data from the XML incident notification data file for integration into third-party software, such as a trouble ticketing system, or helpdesk software.

Table A-2 lists the documentation for the Cisco Security MARS XML incident notification feature.

Table A-2 Related XML Incident Notification Documents

Resource Description	Resource Location
Configuring XML incident notification on MARS	Chapter 8, "Sending Alerts and Incident Notifications"
A ZIP file containing the XML incident notification schema	http://www.cisco.com/application/x-zip-compressed/en/us/guest/products/ps6840/c1225/ccmigration_09186a00806ba94b.zip
A hyper-linked component reference, generated from the XML incident notification schema	http://www.cisco.com/application/x-zip-compressed/en/us/guest/products/ps6840/c1225/ccmigration_09186a00806c27e3.zip
Sample XML incident notification data generated by MARS	Appendix A, ""Example A-1

XML Incident Notification Data File Sample Output

Example A-1 is XML incident notification data generated by the events that trigger the rule "CS-MARS Database Partition Usage."

Example A-1 XML Incident Notification Data File Contents

```
<?xml version="1.0" encoding="UTF-8"?>
<CSMARS-NOTIFICATION>
  <Header>
    <Version>1.0</Version>
    <GenTimeStamp>May 15, 2006 8:48:02 AM PDT</GenTimeStamp>
    <CSMARSHostIpAddr_eth0>10.2.3.7</CSMARSHostIpAddr_eth0>
    <CSMARSHostIpAddr_eth1>192.168.1.101</CSMARSHostIpAddr_eth1>
    <CSMARSHostName>MyLatest</CSMARSHostName>
    <CSMARSZoneName />
    <CSMARSVersion>4.2.1</CSMARSVersion>
  </Header>
  <Data>
    <Incident id="597842933">
      <StartTime>May 15, 2006 8:47:26 AM PDT</StartTime>
      <EndTime>May 15, 2006 8:47:26 AM PDT</EndTime>
      <Severity>LOW</Severity>
      <Session id="597744001">
        <Instance>0</Instance>
        <RuleOffset>1</RuleOffset>
        <SessionEndPoints>
          <Source ipaddress="0.0.0.0" />
          <Destination ipaddress="10.2.3.7" />
          <SourcePort>0</SourcePort>
          <DestinationPort>0</DestinationPort>
          <Protocol>-1</Protocol>
        </SessionEndPoints>
        <Event id="597744001">
          <EventType id="125755" />
          <TimeStamp>May 15, 2006 8:47:26 AM PDT</TimeStamp>
          <ReportingDevice id="50" />
          <RawMessage>Mon May 15 08:47:26 PDT 2006 &lt;13&gt;%MARS-3-100026 CS-MARS
MyLatest : Current database partition pn_event_session_8 utililization has reached 75%;
next database partition pn_event_session_9 containing data between Thu Apr 20 11:59:13 PDT
2006 and Fri Apr 21 11:32:17 PDT 2006 will be purged approximately at Mon May 15 11:56:02
PDT 2006.</RawMessage>
```

```
            <FalsePositiveType>NOT_AVAILABLE</FalsePositiveType>
            <EventEndPoints>
              <Source ipaddress="0.0.0.0" />
              <Destination ipaddress="10.2.3.7" />
              <SourcePort>0</SourcePort>
              <DestinationPort>0</DestinationPort>
              <Protocol>-1</Protocol>
            </EventEndPoints>
            <NATtedEndPoints>
              <Source ipaddress="0.0.0.0" />
              <Destination ipaddress="10.2.3.7" />
              <SourcePort>0</SourcePort>
              <DestinationPort>0</DestinationPort>
              <Protocol>-1</Protocol>
            </NATtedEndPoints>
            <FiringEventFlag>true</FiringEventFlag>
          </Event>
        </Session>
        <Rule id="134473">
          <Name>System Rule: CS-MARS Database Partition Usage</Name>
          <Description>This rule indicates that the current CS-MARS database partition
filled up to 75% of its capacity and the next database partition will be purged soon to
create space for new events. The estimated purge times are in the event message. This is
normal CS-MARS activity and will result in old events and incidents to purged from CS-MARS
database. Users are urged to archive CS-MARS data to prevent permanent data
loss.</Description>
        </Rule>
        <NetworkAddressObj id="0">
          <IPAddress>0.0.0.0</IPAddress>
          <MAC />
          <DNSName />
          <DynamicInfo>
            <HostName />
            <MACAddress />
            <AAAUser />
            <EnforcementDeviceAndPort />
            <ReportingDevice />
            <StartTime>Dec 31, 1969 4:00:00 PM PST</StartTime>
            <EndTime>Dec 31, 1969 4:00:00 PM PST</EndTime>
            <UpdateTime>Dec 31, 1969 4:00:00 PM PST</UpdateTime>
          </DynamicInfo>
        </NetworkAddressObj>
        <NetworkAddressObj id="167904007">
          <IPAddress>10.2.3.7</IPAddress>
          <MAC>
            <MACAddress>00:30:48:83:25:d9</MACAddress>
            <LastUpdateTime>May 15, 2006 6:59:09 AM PDT</LastUpdateTime>
          </MAC>
          <DNSName>MyLatest</DNSName>
          <Device id="50" />
          <DynamicInfo>
            <HostName />
            <MACAddress />
            <AAAUser />
            <EnforcementDeviceAndPort />
            <ReportingDevice />
            <StartTime>Dec 31, 1969 4:00:00 PM PST</StartTime>
            <EndTime>Dec 31, 1969 4:00:00 PM PST</EndTime>
            <UpdateTime>Dec 31, 1969 4:00:00 PM PST</UpdateTime>
          </DynamicInfo>
        </NetworkAddressObj>
        <EventTypeObj id="125755">
          <Name>1000029</Name>
```

```
        <Description>CS-MARS DB partition filling up causing the next partition to be
purged soon</Description>
        <Severity>LOW</Severity>
        <CVE />
      </EventTypeObj>
      <DeviceObj id="50">
        <Name>MyLatest</Name>
        <NetBiosName />
        <DefaultGateway>10.2.3.1</DefaultGateway>
        <OperatingSystem id="0" />
        <InterfaceAddressObj id="117924">
          <Name>eth0</Name>
          <IPAddress>10.2.3.7</IPAddress>
          <MAC>
            <MACAddress>00:30:48:83:25:d9</MACAddress>
            <LastUpdateTime>May 15, 2006 6:59:09 AM PDT</LastUpdateTime>
          </MAC>
        </InterfaceAddressObj>
        <InterfaceAddressObj id="123040">
          <Name>eth1</Name>
          <IPAddress>192.168.1.101</IPAddress>
          <MAC />
        </InterfaceAddressObj>
      </DeviceObj>
    </Incident>
  </Data>
</CSMARS-NOTIFICATION>
```

XML Incident Notification Schema

The XML incident notification schema document (csmars-incident-notification-v1_0.xsd) can be downloaded from the the following URL:

http://www.cisco.com/application/x-zip-compressed/en/us/guest/products/ps6840/c1225/ccmigration_09186a00806ba94b.zip

Usage Guidelines and Conventions for XML Incident Notification

All XML incident notification elements are defined in the XML incident notification schema. A WinZip archive containing a component reference document generated from the schema is available for your convenience at the following URL:

http://www.cisco.com/application/x-zip-compressed/en/us/guest/products/ps6840/c1225/ccmigration_09186a00806c27e3.zip

You can generate a similar document with the application of your choice, or view components, their relationships, constraints, attributes, annotations, and usage guidelines within your XML development environment.

MARS uses a best effort approach to create XML incident notification data. If an error occurs during data compilation, MARS does not stop the process, but sends the data, even if it is partial. Validating the data file against the schema would result in errors for these cases.

The following conventions are observed for XML incident notification data:

- Character encoding is Unicode Transformation Format 8 (UTF-8)
- The reported time zone would be the time zone of the local controller reporting the incident

- Raw messages from reporting devices are XML-escaped in the data file. Your XML parser should be able to unescape XML data.
- If there is no value for an element available from MARS, the element is included in the data file as an empty node. For instance, a DNS name may not be available for a device.
- All date formats are **Mmm dd, yyyy hh:mm:ss AM TZD**
  - **Mmm** is the month (Jan, Feb, Mar. . . Dec)
  - **dd** is the day (1–9, 10–31)
  - **yyyy** is the year (0000–9999)
  - **hh**:**mm**:**ss** is hours, minutes, seconds

    hh are 1–9, 10–12

    mm are 00–60

    ss are 00–60
  - **AM** or **PM**
  - **TZD** is time zone designator (PDT, PST, MDT, MST, etc.)

APPENDIX B

Regular Expression Reference

- PCRE Regular Expression Details, page B-1
- Backslash, page B-2
- Circumflex and Dollar, page B-7
- Full Stop (Period, Dot), page B-8
- Matching a Single Byte, page B-8
- Square Brackets and Character Classes, page B-8
- Posix Character Classes, page B-9
- Vertical Bar, page B-10
- Internal Option Setting, page B-10
- Subpatterns, page B-11
- Named Subpatterns, page B-12
- Repetition, page B-12
- Atomic Grouping and Possessive Quantifiers, page B-14
- Back References, page B-15
- Assertions, page B-16
- Conditional Subpatterns, page B-19
- Comments, page B-20
- Recursive Patterns, page B-20
- Subpatterns as Subroutines, page B-21
- Callouts, page B-22

PCRE Regular Expression Details

The syntax and semantics of the regular expressions supported by PCRE are described below. Regular expressions are also described in the Perl documentation and in a number of books, some of which have copious examples. Jeffrey Friedl's "Mastering Regular Expressions", published by O'Reilly, covers regular expressions in great detail. This description of PCRE's regular expressions is intended as reference material.

The original operation of PCRE was on strings of one-byte characters. However, there is now also support for UTF-8 character strings. To use this, you must build PCRE to include UTF-8 support, and then call pcre_compile() with the PCRE_UTF8 option. How this affects pattern matching is mentioned in several places below. There is also a summary of UTF-8 features in the section on UTF-8 support in the main PCRE page.

A regular expression is a pattern that is matched against a subject string from left to right. Most characters stand for themselves in a pattern, and match the corresponding characters in the subject. As a trivial example, the pattern

```
The quick brown fox
```

matches a portion of a subject string that is identical to itself. The power of regular expressions comes from the ability to include alternatives and repetitions in the pattern. These are encoded in the pattern by the use of *metacharacters*, which do not stand for themselves but instead are interpreted in some special way.

There are two different sets of metacharacters: those that are recognized anywhere in the pattern except within square brackets, and those that are recognized in square brackets. Outside square brackets, the metacharacters are as follows:

```
\      general escape character with several uses
  ^      assert start of string (or line, in multiline mode)
  $      assert end of string (or line, in multiline mode)
  .      match any character except newline (by default)
  [      start character class definition
  |      start of alternative branch
  (      start subpattern
  )      end subpattern
  ?      extends the meaning of (
         also 0 or 1 quantifier
         also quantifier minimizer
  *      0 or more quantifier
  +      1 or more quantifier
         also "possessive quantifier"
  {      start min/max quantifier
```

Part of a pattern that is in square brackets is called a "character class". In a character class the only metacharacters are:

```
\      general escape character
  ^      negate the class, but only if the first character
  -      indicates character range
  [      POSIX character class (only if followed by POSIX syntax)
  ]      terminates the character class
```

The following sections describe the use of each of the metacharacters.

Backslash

The backslash character has several uses. Firstly, if it is followed by a non-alphanumeric character, it takes away any special meaning that character may have. This use of backslash as an escape character applies both inside and outside character classes.

For example, if you want to match a * character, you write * in the pattern. This escaping action applies whether or not the following character would otherwise be interpreted as a metacharacter, so it is always safe to precede a non-alphanumeric with backslash to specify that it stands for itself. In particular, if you want to match a backslash, you write \\.

If a pattern is compiled with the PCRE_EXTENDED option, whitespace in the pattern (other than in a character class) and characters between a # outside a character class and the next newline character are ignored. An escaping backslash can be used to include a whitespace or # character as part of the pattern.

If you want to remove the special meaning from a sequence of characters, you can do so by putting them between \Q and \E. This is different from Perl in that $ and @ are handled as literals in \Q...\E sequences in PCRE, whereas in Perl, $ and @ cause variable interpolation. Note the following examples:

```
Pattern            PCRE matches   Perl matches

\Qabc$xyz\E        abc$xyz        abc followed by the contents of $xyz
\Qabc\$xyz\E       abc\$xyz       abc\$xyz
\Qabc\E\$\Qxyz\E   abc$xyz        abc$xyz
```

The \Q...\E sequence is recognized both inside and outside character classes.

Non-printing Characters

A second use of backslash provides a way of encoding non-printing characters in patterns in a visible manner. There is no restriction on the appearance of non-printing characters, apart from the binary zero that terminates a pattern, but when a pattern is being prepared by text editing, it is usually easier to use one of the following escape sequences than the binary character it represents:

```
\a        alarm, that is, the BEL character (hex 07)
\cx       "control-x", where x is any character
\e        escape (hex 1B)
\f        formfeed (hex 0C)
\n        newline (hex 0A)
\r        carriage return (hex 0D)
\t        tab (hex 09)
\ddd      character with octal code ddd, or backreference
\xhh      character with hex code hh
\x{hhh..} character with hex code hhh... (UTF-8 mode only)
```

The precise effect of \cx is as follows: if x is a lower case letter, it is converted to upper case. Then bit 6 of the character (hex 40) is inverted. Thus \cz becomes hex 1A, but \c{ becomes hex 3B, while \c; becomes hex 7B.

After \x, from zero to two hexadecimal digits are read (letters can be in upper or lower case). In UTF-8 mode, any number of hexadecimal digits may appear between \x{ and }, but the value of the character code must be less than 2**31 (that is, the maximum hexadecimal value is 7FFFFFFF). If characters other than hexadecimal digits appear between \x{ and }, or if there is no terminating }, this form of escape is not recognized. Instead, the initial \x will be interpreted as a basic hexadecimal escape, with no following digits, giving a character whose value is zero.

Characters whose value is less than 256 can be defined by either of the two syntaxes for \x when PCRE is in UTF-8 mode. There is no difference in the way they are handled. For example, \xdc is exactly the same as \x{dc}.

After \0 up to two further octal digits are read. In both cases, if there are fewer than two digits, just those that are present are used. Thus the sequence \0\x\07 specifies two binary zeros followed by a BEL character (code value 7). Make sure you supply two digits after the initial zero if the pattern character that follows is itself an octal digit.

The handling of a backslash followed by a digit other than 0 is complicated. Outside a character class, PCRE reads it and any following digits as a decimal number. If the number is less than 10, or if there have been at least that many previous capturing left parentheses in the expression, the entire sequence is taken as a back reference. A description of how this works is given later, following the discussion of parenthesized subpatterns.

Inside a character class, or if the decimal number is greater than 9 and there have not been that many capturing subpatterns, PCRE re-reads up to three octal digits following the backslash, and generates a single byte from the least significant 8 bits of the value. Any subsequent digits stand for themselves. For example:

```
  \040   is another way of writing a space
  \40    is the same, provided there are fewer than 40 previous capturing subpatterns
  \7     is always a back reference
  \11    might be a back reference, or another way of writing a tab
  \011   is always a tab
  \0113  is a tab followed by the character "3"
  \113   might be a back reference, otherwise the character with octal code 113
  \377   might be a back reference, otherwise the byte consisting entirely of 1 bits
  \81    is either a back reference, or a binary zero followed by the two characters
"8" and "1"
```

Note that octal values of 100 or greater must not be introduced by a leading zero, because no more than three octal digits are ever read.

All the sequences that define a single byte value or a single UTF-8 character (in UTF-8 mode) can be used both inside and outside character classes. In addition, inside a character class, the sequence \b is interpreted as the backspace character (hex 08), and the sequence \X is interpreted as the character "X". Outside a character class, these sequences have different meanings (see Unicode Character Properties, page B-5).

Generic Character Types

The third use of backslash is for specifying generic character types. The following are always recognized:

```
\d     any decimal digit
\D     any character that is not a decimal digit
\s     any whitespace character
\S     any character that is not a whitespace character
\w     any "word" character
\W     any "non-word" character
```

Each pair of escape sequences partitions the complete set of characters into two disjoint sets. Any given character matches one, and only one, of each pair.

These character type sequences can appear both inside and outside character classes. They each match one character of the appropriate type. If the current matching point is at the end of the subject string, all of them fail, since there is no character to match.

For compatibility with Perl, \s does not match the VT character (code 11). This makes it different from the the POSIX "space" class. The \s characters are HT (9), LF (10), FF (12), CR (13), and space (32).

A "word" character is an underscore or any character less than 256 that is a letter or digit. The definition of letters and digits is controlled by PCRE's low-valued character tables, and may vary if locale-specific matching is taking place (see "Locale support" in the **pcreapi** page). For example, in the "fr_FR" (French) locale, some character codes greater than 128 are used for accented letters, and these are matched by \w.

In UTF-8 mode, characters with values greater than 128 never match \d, \s, or \w, and always match \D, \S, and \W. This is true even when Unicode character property support is available.

Unicode Character Properties

When PCRE is built with Unicode character property support, three additional escape sequences to match generic character types are available when UTF-8 mode is selected. They are:

```
\p{xx}   a character with the xx property
\P{xx}   a character without the xx property
\X       an extended Unicode sequence
```

The property names represented by *xx* above are limited to the Unicode general category properties. Each character has exactly one such property, specified by a two-letter abbreviation. For compatibility with Perl, negation can be specified by including a circumflex between the opening brace and the property name. For example, \p{^Lu} is the same as \P{Lu}.

If only one letter is specified with \p or \P, it includes all the properties that start with that letter. In this case, in the absence of negation, the curly brackets in the escape sequence are optional; these two examples have the same effect:

```
\p{L}
\pL
```

The following property codes are supported:

```
C     Other
Cc    Control
Cf    Format
Cn    Unassigned
Co    Private use
Cs    Surrogate

L     Letter
Ll    Lower case letter
Lm    Modifier letter
Lo    Other letter
Lt    Title case letter
Lu    Upper case letter

M     Mark
Mc    Spacing mark
Me    Enclosing mark
Mn    Non-spacing mark

N     Number
Nd    Decimal number
Nl    Letter number
No    Other number

P     Punctuation
Pc    Connector punctuation
```

```
Pd    Dash punctuation
Pe    Close punctuation
Pf    Final punctuation
Pi    Initial punctuation
Po    Other punctuation
Ps    Open punctuation

S     Symbol
Sc    Currency symbol
Sk    Modifier symbol
Sm    Mathematical symbol
So    Other symbol

Z     Separator
Zl    Line separator
Zp    Paragraph separator
Zs    Space separator
```

Extended properties such as "Greek" or "InMusicalSymbols" are not supported by PCRE.

Specifying caseless matching does not affect these escape sequences. For example, \p{Lu} always matches only upper case letters.

The \X escape matches any number of Unicode characters that form an extended Unicode sequence. \X is equivalent to

```
(?>\PM\pM*)
```

That is, it matches a character without the "mark" property, followed by zero or more characters with the "mark" property, and treats the sequence as an atomic group (see below). Characters with the "mark" property are typically accents that affect the preceding character.

Matching characters by Unicode property is not fast, because PCRE has to search a structure that contains data for over fifteen thousand characters. That is why the traditional escape sequences such as \d and \w do not use Unicode properties in PCRE.

Simple Assertions

The fourth use of backslash is for certain simple assertions. An assertion specifies a condition that has to be met at a particular point in a match, without consuming any characters from the subject string. The use of subpatterns for more complicated assertions is described below. The backslashed assertions are:

```
\b     matches at a word boundary
\B     matches when not at a word boundary
\A     matches at start of subject
\Z     matches at end of subject or before newline at end
\z     matches at end of subject
\G     matches at first matching position in subject
```

These assertions may not appear in character classes (but note that \b has a different meaning, namely the backspace character, inside a character class).

A word boundary is a position in the subject string where the current character and the previous character do not both match \w or \W (i.e. one matches \w and the other matches \W), or the start or end of the string if the first or last character matches \w, respectively.

The \A, \Z, and \z assertions differ from the traditional circumflex and dollar (described in the next section) in that they only ever match at the very start and end of the subject string, whatever options are set. Thus, they are independent of multiline mode. These three assertions are not affected by the

PCRE_NOTBOL or PCRE_NOTEOL options, which affect only the behaviour of the circumflex and dollar metacharacters. However, if the *startoffset* argument of **pcre_exec()** is non-zero, indicating that matching is to start at a point other than the beginning of the subject, \A can never match. The difference between \Z and \z is that \Z matches before a newline that is the last character of the string as well as at the end of the string, whereas \z matches only at the end.

The \G assertion is true only when the current matching position is at the start point of the match, as specified by the *startoffset* argument of **pcre_exec()**. It differs from \A when the value of *startoffset* is non-zero. By calling **pcre_exec()** multiple times with appropriate arguments, you can mimic Perl's /g option, and it is in this kind of implementation where \G can be useful.

Note, however, that PCRE's interpretation of \G, as the start of the current match, is subtly different from Perl's, which defines it as the end of the previous match. In Perl, these can be different when the previously matched string was empty. Because PCRE does just one match at a time, it cannot reproduce this behaviour.

If all the alternatives of a pattern begin with \G, the expression is anchored to the starting match position, and the "anchored" flag is set in the compiled regular expression.

Circumflex and Dollar

Outside a character class, in the default matching mode, the circumflex character is an assertion that is true only if the current matching point is at the start of the subject string. If the startoffset argument of pcre_exec() is non-zero, circumflex can never match if the PCRE_MULTILINE option is unset. Inside a character class, circumflex has an entirely different meaning (see Square Brackets and Character Classes, page B-8 and Posix Character Classes, page B-9).

Circumflex need not be the first character of the pattern if a number of alternatives are involved, but it should be the first thing in each alternative in which it appears if the pattern is ever to match that branch. If all possible alternatives start with a circumflex, that is, if the pattern is constrained to match only at the start of the subject, it is said to be an "anchored" pattern. (There are also other constructs that can cause a pattern to be anchored.)

A dollar character is an assertion that is true only if the current matching point is at the end of the subject string, or immediately before a newline character that is the last character in the string (by default). Dollar need not be the last character of the pattern if a number of alternatives are involved, but it should be the last item in any branch in which it appears. Dollar has no special meaning in a character class.

The meaning of dollar can be changed so that it matches only at the very end of the string, by setting the PCRE_DOLLAR_ENDONLY option at compile time. This does not affect the \Z assertion.

The meanings of the circumflex and dollar characters are changed if the PCRE_MULTILINE option is set. When this is the case, they match immediately after and immediately before an internal newline character, respectively, in addition to matching at the start and end of the subject string. For example, the pattern /^abc$/ matches the subject string "def\nabc" (where \n represents a newline character) in multiline mode, but not otherwise. Consequently, patterns that are anchored in single line mode because all branches start with ^ are not anchored in multiline mode, and a match for circumflex is possible when the *startoffset* argument of **pcre_exec()** is non-zero. The PCRE_DOLLAR_ENDONLY option is ignored if PCRE_MULTILINE is set.

Note that the sequences \A, \Z, and \z can be used to match the start and end of the subject in both modes, and if all branches of a pattern start with \A it is always anchored, whether PCRE_MULTILINE is set or not.

Full Stop (Period, Dot)

Outside a character class, a dot in the pattern matches any one character in the subject, including a non-printing character, but not (by default) newline. In UTF-8 mode, a dot matches any UTF-8 character, which might be more than one byte long, except (by default) newline. If the PCRE_DOTALL option is set, dots match newlines as well. The handling of dot is entirely independent of the handling of circumflex and dollar, the only relationship being that they both involve newline characters. Dot has no special meaning in a character class.

Matching a Single Byte

Outside a character class, the escape sequence \C matches any one byte, both in and out of UTF-8 mode. Unlike a dot, it can match a newline. The feature is provided in Perl in order to match individual bytes in UTF-8 mode. Because it breaks up UTF-8 characters into individual bytes, what remains in the string may be a malformed UTF-8 string. For this reason, the \C escape sequence is best avoided.

PCRE does not allow \C to appear in lookbehind assertions (described below), because in UTF-8 mode this would make it impossible to calculate the length of the lookbehind.

Square Brackets and Character Classes

An opening square bracket introduces a character class, terminated by a closing square bracket. A closing square bracket on its own is not special. If a closing square bracket is required as a member of the class, it should be the first data character in the class (after an initial circumflex, if present) or escaped with a backslash.

A character class matches a single character in the subject. In UTF-8 mode, the character may occupy more than one byte. A matched character must be in the set of characters defined by the class, unless the first character in the class definition is a circumflex, in which case the subject character must not be in the set defined by the class. If a circumflex is actually required as a member of the class, ensure it is not the first character, or escape it with a backslash.

For example, the character class [aeiou] matches any lower case vowel, while [^aeiou] matches any character that is not a lower case vowel. Note that a circumflex is just a convenient notation for specifying the characters that are in the class by enumerating those that are not. A class that starts with a circumflex is not an assertion: it still consumes a character from the subject string, and therefore it fails if the current pointer is at the end of the string.

In UTF-8 mode, characters with values greater than 255 can be included in a class as a literal string of bytes, or by using the \x{ escaping mechanism.

When caseless matching is set, any letters in a class represent both their upper case and lower case versions, so for example, a caseless [aeiou] matches "A" as well as "a", and a caseless [^aeiou] does not match "A", whereas a caseful version would. When running in UTF-8 mode, PCRE supports the concept of case for characters with values greater than 128 only when it is compiled with Unicode property support.

The newline character is never treated in any special way in character classes, whatever the setting of the PCRE_DOTALL or PCRE_MULTILINE options is. A class such as [^a] will always match a newline.

The minus (hyphen) character can be used to specify a range of characters in a character class. For example, [d-m] matches any letter between d and m, inclusive. If a minus character is required in a class, it must be escaped with a backslash or appear in a position where it cannot be interpreted as indicating a range, typically as the first or last character in the class.

It is not possible to have the literal character "]" as the end character of a range. A pattern such as [W-]46] is interpreted as a class of two characters ("W" and "-") followed by a literal string "46]", so it would match "W46]" or "-46]". However, if the "]" is escaped with a backslash it is interpreted as the end of range, so [W-\]46] is interpreted as a class containing a range followed by two other characters. The octal or hexadecimal representation of "]" can also be used to end a range.

Ranges operate in the collating sequence of character values. They can also be used for characters specified numerically, for example [\000-\037]. In UTF-8 mode, ranges can include characters whose values are greater than 255, for example [\x{100}-\x{2ff}].

If a range that includes letters is used when caseless matching is set, it matches the letters in either case. For example, [W-c] is equivalent to [][\\^_`wxyzabc], matched caselessly, and in non-UTF-8 mode, if character tables for the "fr_FR" locale are in use, [\xc8-\xcb] matches accented E characters in both cases. In UTF-8 mode, PCRE supports the concept of case for characters with values greater than 128 only when it is compiled with Unicode property support.

The character types \d, \D, \p, \P, \s, \S, \w, and \W may also appear in a character class, and add the characters that they match to the class. For example, [\dABCDEF] matches any hexadecimal digit. A circumflex can conveniently be used with the upper case character types to specify a more restricted set of characters than the matching lower case type. For example, the class [^\W_] matches any letter or digit, but not underscore.

The only metacharacters that are recognized in character classes are backslash, hyphen (only where it can be interpreted as specifying a range), circumflex (only at the start), opening square bracket (only when it can be interpreted as introducing a POSIX class name - see the next section), and the terminating closing square bracket. However, escaping other non-alphanumeric characters does no harm.

Posix Character Classes

Perl supports the POSIX notation for character classes. This uses names enclosed by [: and :] within the enclosing square brackets. PCRE also supports this notation. For example,

```
[01[:alpha:]%]
```

matches "0", "1", any alphabetic character, or "%". The supported class names are

```
alnum    letters and digits
alpha    letters
ascii    character codes 0 - 127
blank    space or tab only
cntrl    control characters
digit    decimal digits (same as \d)
graph    printing characters, excluding space
lower    lower case letters
print    printing characters, including space
punct    printing characters, excluding letters and digits
space    white space (not quite the same as \s)
upper    upper case letters
word     "word" characters (same as \w)
xdigit   hexadecimal digits
```

The "space" characters are HT (9), LF (10), VT (11), FF (12), CR (13), and space (32). Notice that this list includes the VT character (code 11). This makes "space" different to \s, which does not include VT (for Perl compatibility).

The name "word" is a Perl extension, and "blank" is a GNU extension from Perl 5.8. Another Perl extension is negation, which is indicated by a ^ character after the colon. For example,

```
[12[:^digit:]]
```

matches "1", "2", or any non-digit. PCRE (and Perl) also recognize the POSIX syntax [.ch.] and [=ch=] where "ch" is a "collating element", but these are not supported, and an error is given if they are encountered.

In UTF-8 mode, characters with values greater than 128 do not match any of the POSIX character classes.

Vertical Bar

Vertical bar characters are used to separate alternative patterns. For example, the pattern

```
gilbert|sullivan
```

matches either "gilbert" or "sullivan". Any number of alternatives may appear, and an empty alternative is permitted (matching the empty string). The matching process tries each alternative in turn, from left to right, and the first one that succeeds is used. If the alternatives are within a subpattern (defined below), "succeeds" means matching the rest of the main pattern as well as the alternative in the subpattern.

Internal Option Setting

The settings of the PCRE_CASELESS, PCRE_MULTILINE, PCRE_DOTALL, and PCRE_EXTENDED options can be changed from within the pattern by a sequence of Perl option letters enclosed between "(?" and ")". The option letters are

```
i  for PCRE_CASELESS
m  for PCRE_MULTILINE
s  for PCRE_DOTALL
x  for PCRE_EXTENDED
```

For example, (?im) sets caseless, multiline matching. It is also possible to unset these options by preceding the letter with a hyphen, and a combined setting and unsetting such as (?im-sx), which sets PCRE_CASELESS and PCRE_MULTILINE while unsetting PCRE_DOTALL and PCRE_EXTENDED, is also permitted. If a letter appears both before and after the hyphen, the option is unset.

When an option change occurs at top level (that is, not inside subpattern parentheses), the change applies to the remainder of the pattern that follows. If the change is placed right at the start of a pattern, PCRE extracts it into the global options (and it will therefore show up in data extracted by the **pcre_fullinfo()** function).

An option change within a subpattern affects only that part of the current pattern that follows it, so

```
(a(?i)b)c
```

matches abc and aBc and no other strings (assuming PCRE_CASELESS is not used). By this means, options can be made to have different settings in different parts of the pattern. Any changes made in one alternative do carry on into subsequent branches within the same subpattern. For example,

```
(a(?i)b|c)
```

matches "ab", "aB", "c", and "C", even though when matching "C" the first branch is abandoned before the option setting. This is because the effects of option settings happen at compile time. There would be some very weird behaviour otherwise.

The PCRE-specific options PCRE_UNGREEDY and PCRE_EXTRA can be changed in the same way as the Perl-compatible options by using the characters U and X respectively. The (?X) flag setting is special in that it must always occur earlier in the pattern than any of the additional features it turns on, even when it is at top level. It is best to put it at the start.

Subpatterns

Subpatterns are delimited by parentheses (round brackets), which can be nested. Turning part of a pattern into a subpattern does two things:

Step 1 It localizes a set of alternatives. For example, the pattern :

```
cat(aract|erpillar|)
```

matches one of the words "cat", "cataract", or "caterpillar". Without the parentheses, it would match "cataract", "erpillar" or the empty string.

Step 2 It sets up the subpattern as a capturing subpattern. This means that, when the whole pattern matches, that portion of the subject string that matched the subpattern is passed back to the caller via the *ovector* argument of **pcre_exec()**. Opening parentheses are counted from left to right (starting from 1) to obtain numbers for the capturing subpatterns.

For example, if the string "the red king" is matched against the pattern

```
the ((red|white) (king|queen))
```

the captured substrings are "red king", "red", and "king", and are numbered 1, 2, and 3, respectively.

The fact that plain parentheses fulfil two functions is not always helpful. There are often times when a grouping subpattern is required without a capturing requirement. If an opening parenthesis is followed by a question mark and a colon, the subpattern does not do any capturing, and is not counted when computing the number of any subsequent capturing subpatterns. For example, if the string "the white queen" is matched against the pattern

```
the ((?:red|white) (king|queen))
```

the captured substrings are "white queen" and "queen", and are numbered 1 and 2. The maximum number of capturing subpatterns is 65535, and the maximum depth of nesting of all subpatterns, both capturing and non-capturing, is 200.

As a convenient shorthand, if any option settings are required at the start of a non-capturing subpattern, the option letters may appear between the "?" and the ":". Thus the two patterns

```
(?i:saturday|sunday)
(?:(?i)saturday|sunday)
```

match exactly the same set of strings. Because alternative branches are tried from left to right, and options are not reset until the end of the subpattern is reached, an option setting in one branch does affect subsequent branches, so the above patterns match "SUNDAY" as well as "Saturday".

Named Subpatterns

Identifying capturing parentheses by number is simple, but it can be very hard to keep track of the numbers in complicated regular expressions. Furthermore, if an expression is modified, the numbers may change. To help with this difficulty, PCRE supports the naming of subpatterns, something that Perl does not provide. The Python syntax (?P<name>...) is used. Names consist of alphanumeric characters and underscores, and must be unique within a pattern.

Named capturing parentheses are still allocated numbers as well as names. The PCRE API provides function calls for extracting the name-to-number translation table from a compiled pattern. There is also a convenience function for extracting a captured substring by name. For further details see the pcreapi documentation.

Repetition

Repetition is specified by quantifiers, which can follow any of the following items:

```
a literal data character
the . metacharacter
the \C escape sequence
the \X escape sequence (in UTF-8 mode with Unicode properties)
an escape such as \d that matches a single character
a character class
a back reference (see next section)
a parenthesized subpattern (unless it is an assertion)
```

The general repetition quantifier specifies a minimum and maximum number of permitted matches, by giving the two numbers in curly brackets (braces), separated by a comma. The numbers must be less than 65536, and the first must be less than or equal to the second. For example:

```
z{2,4}
```

matches "zz", "zzz", or "zzzz". A closing brace on its own is not a special character. If the second number is omitted, but the comma is present, there is no upper limit; if the second number and the comma are both omitted, the quantifier specifies an exact number of required matches. Thus

```
[aeiou]{3,}
```

matches at least 3 successive vowels, but may match many more, while

```
\d{8}
```

matches exactly 8 digits. An opening curly bracket that appears in a position where a quantifier is not allowed, or one that does not match the syntax of a quantifier, is taken as a literal character. For example, {,6} is not a quantifier, but a literal string of four characters.

In UTF-8 mode, quantifiers apply to UTF-8 characters rather than to individual bytes. Thus, for example, \x{100}{2} matches two UTF-8 characters, each of which is represented by a two-byte sequence. Similarly, when Unicode property support is available, \X{3} matches three Unicode extended sequences, each of which may be several bytes long (and they may be of different lengths).

The quantifier {0} is permitted, causing the expression to behave as if the previous item and the quantifier were not present.

For convenience (and historical compatibility) the three most common quantifiers have single-character abbreviations:

```
*    is equivalent to {0,}
+    is equivalent to {1,}
?    is equivalent to {0,1}
```

It is possible to construct infinite loops by following a subpattern that can match no characters with a quantifier that has no upper limit, for example:

```
(a?)*
```

Earlier versions of Perl and PCRE used to give an error at compile time for such patterns. However, because there are cases where this can be useful, such patterns are now accepted, but if any repetition of the subpattern does in fact match no characters, the loop is forcibly broken.

By default, the quantifiers are "greedy", that is, they match as much as possible (up to the maximum number of permitted times), without causing the rest of the pattern to fail. The classic example of where this gives problems is in trying to match comments in C programs. These appear between /* and */ and within the comment, individual * and / characters may appear. An attempt to match C comments by applying the pattern

```
/\*.*\*/
```

to the string

```
/* first comment */  not comment  /* second comment */
```

fails, because it matches the entire string owing to the greediness of the .* item.

However, if a quantifier is followed by a question mark, it ceases to be greedy, and instead matches the minimum number of times possible, so the pattern

```
/\*.*?\*/
```

does the right thing with the C comments. The meaning of the various quantifiers is not otherwise changed, just the preferred number of matches. Do not confuse this use of question mark with its use as a quantifier in its own right. Because it has two uses, it can sometimes appear doubled, as in

```
\d??\d
```

which matches one digit by preference, but can match two if that is the only way the rest of the pattern matches.

If the PCRE_UNGREEDY option is set (an option which is not available in Perl), the quantifiers are not greedy by default, but individual ones can be made greedy by following them with a question mark. In other words, it inverts the default behaviour.

When a parenthesized subpattern is quantified with a minimum repeat count that is greater than 1 or with a limited maximum, more memory is required for the compiled pattern, in proportion to the size of the minimum or maximum.

If a pattern starts with .* or .{0,} and the PCRE_DOTALL option (equivalent to Perl's /s) is set, thus allowing the . to match newlines, the pattern is implicitly anchored, because whatever follows will be tried against every character position in the subject string, so there is no point in retrying the overall match at any position after the first. PCRE normally treats such a pattern as though it were preceded by \A.

In cases where it is known that the subject string contains no newlines, it is worth setting PCRE_DOTALL in order to obtain this optimization, or alternatively using ^ to indicate anchoring explicitly.

However, there is one situation where the optimization cannot be used. When .* is inside capturing parentheses that are the subject of a backreference elsewhere in the pattern, a match at the start may fail, and a later one succeed. Consider, for example:

```
(.*)abc\1
```

If the subject is "xyz123abc123" the match point is the fourth character. For this reason, such a pattern is not implicitly anchored.

When a capturing subpattern is repeated, the value captured is the substring that matched the final iteration. For example, after

```
(tweedle[dume]{3}\s*)+
```

has matched "tweedledum tweedledee" the value of the captured substring is "tweedledee". However, if there are nested capturing subpatterns, the corresponding captured values may have been set in previous iterations. For example, after

```
/(a|(b))+/
```

matches "aba" the value of the second captured substring is "b".

Atomic Grouping and Possessive Quantifiers

With both maximizing and minimizing repetition, failure of what follows normally causes the repeated item to be re-evaluated to see if a different number of repeats allows the rest of the pattern to match. Sometimes it is useful to prevent this, either to change the nature of the match, or to cause it fail earlier than it otherwise might, when the author of the pattern knows there is no point in carrying on.

Consider, for example, the pattern \d+foo when applied to the subject line

```
123456bar
```

After matching all 6 digits and then failing to match "foo", the normal action of the matcher is to try again with only 5 digits matching the \d+ item, and then with 4, and so on, before ultimately failing. "Atomic grouping" (a term taken from Jeffrey Friedl's book) provides the means for specifying that once a subpattern has matched, it is not to be re-evaluated in this way.

If we use atomic grouping for the previous example, the matcher would give up immediately on failing to match "foo" the first time. The notation is a kind of special parenthesis, starting with (?> as in this example:

```
(?>\d+)foo
```

This kind of parenthesis "locks up" the part of the pattern it contains once it has matched, and a failure further into the pattern is prevented from backtracking into it. Backtracking past it to previous items, however, works as normal.

An alternative description is that a subpattern of this type matches the string of characters that an identical standalone pattern would match, if anchored at the current point in the subject string.

Atomic grouping subpatterns are not capturing subpatterns. Simple cases such as the above example can be thought of as a maximizing repeat that must swallow everything it can. So, while both \d+ and \d+? are prepared to adjust the number of digits they match in order to make the rest of the pattern match, (?>\d+) can only match an entire sequence of digits.

Atomic groups in general can of course contain arbitrarily complicated subpatterns, and can be nested. However, when the subpattern for an atomic group is just a single repeated item, as in the example above, a simpler notation, called a "possessive quantifier" can be used. This consists of an additional + character following a quantifier. Using this notation, the previous example can be rewritten as

```
\d++foo
```

Possessive quantifiers are always greedy; the setting of the PCRE_UNGREEDY option is ignored. They are a convenient notation for the simpler forms of atomic group. However, there is no difference in the meaning or processing of a possessive quantifier and the equivalent atomic group.

The possessive quantifier syntax is an extension to the Perl syntax. It originates in Sun's Java package.

When a pattern contains an unlimited repeat inside a subpattern that can itself be repeated an unlimited number of times, the use of an atomic group is the only way to avoid some failing matches taking a very long time indeed. The pattern

```
(\D+|<\d+>)*[!?]
```

matches an unlimited number of substrings that either consist of non-digits, or digits enclosed in <>, followed by either ! or ?. When it matches, it runs quickly. However, if it is applied to

```
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
```

it takes a long time before reporting failure. This is because the string can be divided between the internal \D+ repeat and the external * repeat in a large number of ways, and all have to be tried. (The example uses [!?] rather than a single character at the end, because both PCRE and Perl have an optimization that allows for fast failure when a single character is used. They remember the last single character that is required for a match, and fail early if it is not present in the string.) If the pattern is changed so that it uses an atomic group, like this:

```
((?>\D+)|<\d+>)*[!?]
```

sequences of non-digits cannot be broken, and failure happens quickly.

Back References

Outside a character class, a backslash followed by a digit greater than 0 (and possibly further digits) is a back reference to a capturing subpattern earlier (that is, to its left) in the pattern, provided there have been that many previous capturing left parentheses.

However, if the decimal number following the backslash is less than 10, it is always taken as a back reference, and causes an error only if there are not that many capturing left parentheses in the entire pattern. In other words, the parentheses that are referenced need not be to the left of the reference for numbers less than 10. See Non-printing Characters, page B-3 for further details of the handling of digits following a backslash.

A back reference matches whatever actually matched the capturing subpattern in the current subject string, rather than anything matching the subpattern itself (see Subpatterns as Subroutines, page B-21 for a way of doing that). So the pattern

```
(sens|respons)e and \1ibility
```

matches "sense and sensibility" and "response and responsibility", but not "sense and responsibility". If caseful matching is in force at the time of the back reference, the case of letters is relevant. For example,

```
((?i)rah)\s+\1
```

matches "rah rah" and "RAH RAH", but not "RAH rah", even though the original capturing subpattern is matched caselessly.

Back references to named subpatterns use the Python syntax (?P=name). We could rewrite the above example as follows:

```
(?<p1>(?i)rah)\s+(?P=p1)
```

There may be more than one back reference to the same subpattern. If a subpattern has not actually been used in a particular match, any back references to it always fail. For example, the pattern

```
(a|(bc))\2
```

always fails if it starts to match "a" rather than "bc". Because there may be many capturing parentheses in a pattern, all digits following the backslash are taken as part of a potential back reference number. If the pattern continues with a digit character, some delimiter must be used to terminate the back reference. If the PCRE_EXTENDED option is set, this can be whitespace. Otherwise an empty comment (see Comments, page B-20) can be used.

A back reference that occurs inside the parentheses to which it refers fails when the subpattern is first used, so, for example, (a\1) never matches. However, such references can be useful inside repeated subpatterns. For example, the pattern

```
(a|b\1)+
```

matches any number of "a"s and also "aba", "ababbaa" etc. At each iteration of the subpattern, the back reference matches the character string corresponding to the previous iteration. In order for this to work, the pattern must be such that the first iteration does not need to match the back reference. This can be done using alternation, as in the example above, or by a quantifier with a minimum of zero.

Assertions

An assertion is a test on the characters following or preceding the current matching point that does not actually consume any characters. The simple assertions coded as \b, \B, \A, \G, \Z, \z, ^ and $ are described above.

More complicated assertions are coded as subpatterns. There are two kinds: those that look ahead of the current position in the subject string, and those that look behind it. An assertion subpattern is matched in the normal way, except that it does not cause the current matching position to be changed.

Assertion subpatterns are not capturing subpatterns, and may not be repeated, because it makes no sense to assert the same thing several times. If any kind of assertion contains capturing subpatterns within it, these are counted for the purposes of numbering the capturing subpatterns in the whole pattern. However, substring capturing is carried out only for positive assertions, because it does not make sense for negative assertions.

Lookahead Assertions

Lookahead assertions start with (?= for positive assertions and (?! for negative assertions. For example,

```
\w+(?=;)
```

matches a word followed by a semicolon, but does not include the semicolon in the match, and

```
foo(?!bar)
```

matches any occurrence of "foo" that is not followed by "bar". Note that the apparently similar pattern

```
(?!foo)bar
```

does not find an occurrence of "bar" that is preceded by something other than "foo"; it finds any occurrence of "bar" whatsoever, because the assertion (?!foo) is always true when the next three characters are "bar". A lookbehind assertion is needed to achieve the other effect.

If you want to force a matching failure at some point in a pattern, the most convenient way to do it is with (?!) because an empty string always matches, so an assertion that requires there not to be an empty string must always fail.

Lookbehind Assertions

Lookbehind assertions start with (?<= for positive assertions and (?<! for negative assertions. For example,

```
(?<!foo)bar
```

does find an occurrence of "bar" that is not preceded by "foo". The contents of a lookbehind assertion are restricted such that all the strings it matches must have a fixed length. However, if there are several alternatives, they do not all have to have the same fixed length. Thus

```
(?<=bullock|donkey)
```

is permitted, but

```
(?<!dogs?|cats?)
```

causes an error at compile time. Branches that match different length strings are permitted only at the top level of a lookbehind assertion. This is an extension compared with Perl (at least for 5.8), which requires all branches to match the same length of string. An assertion such as

```
(?<=ab(c|de))
```

is not permitted, because its single top-level branch can match two different lengths, but it is acceptable if rewritten to use two top-level branches:

```
(?<=abc|abde)
```

The implementation of lookbehind assertions is, for each alternative, to temporarily move the current position back by the fixed width and then try to match. If there are insufficient characters before the current position, the match is deemed to fail.

PCRE does not allow the \C escape (which matches a single byte in UTF-8 mode) to appear in lookbehind assertions, because it makes it impossible to calculate the length of the lookbehind. The \X escape, which can match different numbers of bytes, is also not permitted.

Atomic groups can be used in conjunction with lookbehind assertions to specify efficient matching at the end of the subject string. Consider a simple pattern such as

```
abcd$
```

when applied to a long string that does not match. Because matching proceeds from left to right, PCRE will look for each "a" in the subject and then see if what follows matches the rest of the pattern. If the pattern is specified as

```
^.*abcd$
```

the initial .* matches the entire string at first, but when this fails (because there is no following "a"), it backtracks to match all but the last character, then all but the last two characters, and so on. Once again the search for "a" covers the entire string, from right to left, so we are no better off. However, if the pattern is written as

```
^(?>.*)(?<=abcd)
```

or, equivalently, using the possessive quantifier syntax,

```
^.*+(?<=abcd)
```

there can be no backtracking for the .* item; it can match only the entire string. The subsequent lookbehind assertion does a single test on the last four characters. If it fails, the match fails immediately. For long strings, this approach makes a significant difference to the processing time.

Using Multiple Assertions

Several assertions (of any sort) may occur in succession. For example,

```
(?<=\d{3})(?<!999)foo
```

matches "foo" preceded by three digits that are not "999". Notice that each of the assertions is applied independently at the same point in the subject string. First there is a check that the previous three characters are all digits, and then there is a check that the same three characters are not "999". This pattern does *not* match "foo" preceded by six characters, the first of which are digits and the last three of which are not "999". For example, it doesn't match "123abcfoo". A pattern to do that is

```
(?<=\d{3}...)(?<!999)foo
```

This time the first assertion looks at the preceding six characters, checking that the first three are digits, and then the second assertion checks that the preceding three characters are not "999".

Assertions can be nested in any combination. For example,

```
(?<=(?<!foo)bar)baz
```

matches an occurrence of "baz" that is preceded by "bar" which in turn is not preceded by "foo", while

```
(?<=\d{3}(?!999)...)foo
```

is another pattern that matches "foo" preceded by three digits and any three characters that are not "999".

Conditional Subpatterns

It is possible to cause the matching process to obey a subpattern conditionally or to choose between two alternative subpatterns, depending on the result of an assertion, or whether a previous capturing subpattern matched or not. The two possible forms of conditional subpattern are

```
(?(condition)yes-pattern)
(?(condition)yes-pattern|no-pattern)
```

If the condition is satisfied, the yes-pattern is used; otherwise the no-pattern (if present) is used. If there are more than two alternatives in the subpattern, a compile-time error occurs.

There are three kinds of condition. If the text between the parentheses consists of a sequence of digits, the condition is satisfied if the capturing subpattern of that number has previously matched. The number must be greater than zero. Consider the following pattern, which contains non-significant white space to make it more readable (assume the PCRE_EXTENDED option) and to divide it into three parts for ease of discussion:

```
( \( )?    [^()]+    (?(1) \) )
```

The first part matches an optional opening parenthesis, and if that character is present, sets it as the first captured substring. The second part matches one or more characters that are not parentheses. The third part is a conditional subpattern that tests whether the first set of parentheses matched or not. If they did, that is, if subject started with an opening parenthesis, the condition is true, and so the yes-pattern is executed and a closing parenthesis is required. Otherwise, since no-pattern is not present, the subpattern matches nothing. In other words, this pattern matches a sequence of non-parentheses, optionally enclosed in parentheses.

If the condition is the string (R), it is satisfied if a recursive call to the pattern or subpattern has been made. At "top level", the condition is false. This is a PCRE extension. Recursive patterns are described in the next section.

If the condition is not a sequence of digits or (R), it must be an assertion. This may be a positive or negative lookahead or lookbehind assertion. Consider this pattern, again containing non-significant white space, and with the two alternatives on the second line:

```
(?(?=[^a-z]*[a-z])
\d{2}-[a-z]{3}-\d{2}  |  \d{2}-\d{2}-\d{2} )
```

The condition is a positive lookahead assertion that matches an optional sequence of non-letters followed by a letter. In other words, it tests for the presence of at least one letter in the subject. If a letter is found, the subject is matched against the first alternative; otherwise it is matched against the second. This pattern matches strings in one of the two forms dd-aaa-dd or dd-dd-dd, where aaa are letters and dd are digits.

Comments

The sequence (?# marks the start of a comment that continues up to the next closing parenthesis. Nested parentheses are not permitted. The characters that make up a comment play no part in the pattern matching at all.

If the PCRE_EXTENDED option is set, an unescaped # character outside a character class introduces a comment that continues up to the next newline character in the pattern.

Recursive Patterns

Consider the problem of matching a string in parentheses, allowing for unlimited nested parentheses. Without the use of recursion, the best that can be done is to use a pattern that matches up to some fixed depth of nesting. It is not possible to handle an arbitrary nesting depth. Perl provides a facility that allows regular expressions to recurse (amongst other things). It does this by interpolating Perl code in the expression at run time, and the code can refer to the expression itself. A Perl pattern to solve the parentheses problem can be created like this:

```
$re = qr{\( (?: (?>[^()]+) | (?p{$re}) )* \)}x;
```

The (?p{...}) item interpolates Perl code at run time, and in this case refers recursively to the pattern in which it appears. Obviously, PCRE cannot support the interpolation of Perl code. Instead, it supports some special syntax for recursion of the entire pattern, and also for individual subpattern recursion.

The special item that consists of (? followed by a number greater than zero and a closing parenthesis is a recursive call of the subpattern of the given number, provided that it occurs inside that subpattern. (If not, it is a "subroutine" call, which is described in the next section.) The special item (?R) is a recursive call of the entire regular expression.

For example, this PCRE pattern solves the nested parentheses problem (assume the PCRE_EXTENDED option is set so that white space is ignored):

```
\( ( (?>[^()]+) | (?R) )* \)
```

First it matches an opening parenthesis. Then it matches any number of substrings which can either be a sequence of non-parentheses, or a recursive match of the pattern itself (that is a correctly parenthesized substring). Finally there is a closing parenthesis.

If this were part of a larger pattern, you would not want to recurse the entire pattern, so instead you could use this:

```
( \( ( (?>[^()]+) | (?1) )* \) )
```

We have put the pattern into parentheses, and caused the recursion to refer to them instead of the whole pattern. In a larger pattern, keeping track of parenthesis numbers can be tricky. It may be more convenient to use named parentheses instead. For this, PCRE uses (?P>name), which is an extension to the Python syntax that PCRE uses for named parentheses (Perl does not provide named parentheses). We could rewrite the above example as follows:

```
(?P<pn> \( ( (?>[^()]+) | (?P>pn) )* \) )
```

This particular example pattern contains nested unlimited repeats, and so the use of atomic grouping for matching strings of non-parentheses is important when applying the pattern to strings that do not match. For example, when this pattern is applied to

```
(aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa()
```

it yields "no match" quickly. However, if atomic grouping is not used, the match runs for a very long time indeed because there are so many different ways the + and * repeats can carve up the subject, and all have to be tested before failure can be reported.

At the end of a match, the values set for any capturing subpatterns are those from the outermost level of the recursion at which the subpattern value is set. If you want to obtain intermediate values, a callout function can be used (see Subpatterns as Subroutines, page B-21 and the **pcrecallout** documentation). If the pattern above is matched against

```
(ab(cd)ef)
```

the value for the capturing parentheses is "ef", which is the last value taken on at the top level. If additional parentheses are added, giving

```
\( ( ( (?>[^()]+) | (?R) )* ) \)
   ^                        ^
   ^                        ^
```

the string they capture is "ab(cd)ef", the cont

ents of the top level parentheses. If there are more than 15 capturing parentheses in a pattern, PCRE has to obtain extra memory to store data during a recursion, which it does by using **pcre_malloc**, freeing it via **pcre_free** afterwards. If no memory can be obtained, the match fails with the PCRE_ERROR_NOMEMORY error.

Do not confuse the (?R) item with the condition (R), which tests for recursion. Consider this pattern, which matches text in angle brackets, allowing for arbitrary nesting. Only digits are allowed in nested brackets (that is, when recursing), whereas any characters are permitted at the outer level.

```
< (?: (?(R) \d++  | [^<>]*+) | (?R)) * >
```

In this pattern, (?(R) is the start of a conditional subpattern, with two different alternatives for the recursive and non-recursive cases. The (?R) item is the actual recursive call.

Subpatterns as Subroutines

If the syntax for a recursive subpattern reference (either by number or by name) is used outside the parentheses to which it refers, it operates like a subroutine in a programming language. An earlier example pointed out that the pattern

```
(sens|respons)e and \1ibility
```

matches "sense and sensibility" and "response and responsibility", but not "sense and responsibility". If instead the pattern

```
(sens|respons)e and (?1)ibility
```

is used, it does match "sense and responsibility" as well as the other two strings. Such references must, however, follow the subpattern to which they refer.

Callouts

Perl has a feature whereby using the sequence (?{...}) causes arbitrary Perl code to be obeyed in the middle of matching a regular expression. This makes it possible, amongst other things, to extract different substrings that match the same pair of parentheses when there is a repetition.

PCRE provides a similar feature, but of course it cannot obey arbitrary Perl code. The feature is called "callout". The caller of PCRE provides an external function by putting its entry point in the global variable *pcre_callout*. By default, this variable contains NULL, which disables all calling out.

Within a regular expression, (?C) indicates the points at which the external function is to be called. If you want to identify different callout points, you can put a number less than 256 after the letter C. The default value is zero. For example, this pattern has two callout points:

```
(?C1)\dabc(?C2)def
```

If the PCRE_AUTO_CALLOUT flag is passed to **pcre_compile()**, callouts are automatically installed before each item in the pattern. They are all numbered 255.

During matching, when PCRE reaches a callout point (and *pcre_callout* is set), the external function is called. It is provided with the number of the callout, the position in the pattern, and, optionally, one item of data originally supplied by the caller of pcre_exec(). The callout function may cause matching to proceed, to backtrack, or to fail altogether. A complete description of the interface to the callout function is given in the **pcrecallout** documentation.

Last updated: 09 September 2004
Copyright © 1997-2004 University of Cambridge.

APPENDIX C

Date/Time Format Specfication

The date/time field parsing is supported using the Unix `strptime()` standard C library function.

The **strptime()** function is the converse function to **strftime()** and converts the character string pointed to by *s* to values which are stored in the *tm* structure pointed to by *tm*, using the format specified by *format*. Here ***format*** is a character string that consists of field descriptors and text characters, reminiscent of scanf(3). Each field descriptor consists of a % character followed by another character that specifies the replacement for the field descriptor. All other characters in the *format* string must have a matching character in the input string, except for whitespace, which matches zero or more whitespace characters in the input string.

The **strptime()** function processes the input string from left to right. Each of the three possible input elements (whitespace, literal, or format) are handled one after the other. If the input cannot be matched to the format string the function stops. The remainder of the format and input strings are not processed.

The supported input field descriptors are listed below. In case a text string (such as a weekday or month name) is to be matched, the comparison is case insensitive. In case a number is to be matched, leading zeros are permitted but not required.

%%

The **%** character.

%a or **%A**

The weekday name according to the current locale, in abbreviated form or the full name.

%b or **%B** or **%h**

The month name according to the current locale, in abbreviated form or the full name.

%c

The date and time representation for the current locale.

%C

The century number (0-99).

%d or **%e**

The day of month (1-31).

%D

Equivalent to %m/%d/%y. (This is the American style date, very confusing to non-Americans, especially since %d/%m/%y is widely used in Europe. The ISO 8601 standard format is %Y-%m-%d.)

%H

The hour (0-23).

%I

The hour on a 12-hour clock (1-12).

%j

The day number in the year (1-366).

%m

The month number (1-12).

%M

The minute (0-59).

%n or **%t**

Arbitrary whitespace.

%p

The locale's equivalent of AM or PM. (Note: there may be none.)

%r

The 12-hour clock time (using the locale's AM or PM). In the POSIX locale equivalent to %I:%M:%S %p. If *t_fmt_ampm* is empty in the LC_TIME part of the current locale then the behaviour is undefined.

%R

Equivalent to %H:%M.

%S

The second (0-60; 60 may occur for leap seconds; earlier also 61 was allowed).

%T

Equivalent to %H:%M:%S.

%U

The week number with Sunday the first day of the week (0-53). The first Sunday of January is the first day of week 1.

%w

The weekday number (0-6) with Sunday = 0.

%W

The week number with Monday the first day of the week (0-53). The first Monday of January is the first day of week 1.

%x

The date, using the locale's date format.

%X

The time, using the locale's time format.

%y

The year within century (0-99). When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969-1999); values in the range 00-68 refer to years in the twenty-first century (2000-2068).

%Y

The year, including century (for example, 1991).

Some field descriptors can be modified by the E or O modifier characters to indicate that an alternative format or specification should be used. If the alternative format or specification does not exist in the current locale, the unmodified field descriptor is used.

The E modifier specifies that the input string may contain alternative locale-dependent versions of the date and time representation:

%Ec

The locale's alternative date and time representation.

%EC

The name of the base year (period) in the locale's alternative representation.

%Ex

The locale's alternative date representation.

%EX

The locale's alternative time representation.

%Ey

The offset from %EC (year only) in the locale's alternative representation.

%EY

The full alternative year representation.

The O modifier specifies that the numerical input may be in an alternative locale-dependent format:

%Od or **%Oe**

The day of the month using the locale's alternative numeric symbols; leading zeros are permitted but not required.

%OH

The hour (24-hour clock) using the locale's alternative numeric symbols.

%OI

The hour (12-hour clock) using the locale's alternative numeric symbols.

%Om

The month using the locale's alternative numeric symbols.

%OM

The minutes using the locale's alternative numeric symbols.

%OS

The seconds using the locale's alternative numeric symbols.

%OU

The week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.

%Ow

The number of the weekday (Sunday=0) using the locale's alternative numeric symbols.

%OW

The week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.

%Oy

The year (offset from %C) using the locale's alternative numeric symbols.

%F

Equivalent to **%Y-%m-%d**, the ISO 8601 date format.

%g

The year corresponding to the ISO week number, but without the century (0-99).

%G

The year corresponding to the ISO week number. (For example, 1991.)

%u

The day of the week as a decimal number (1-7, where Monday = 1).

%V

The ISO 8601:1988 week number as a decimal number (1-53). If the week (starting on Monday) containing 1 January has four or more days in the new year, then it is considered week 1. Otherwise, it is the last week of the previous year, and the next week is week 1.

%z

An RFC-822/ISO 8601 standard time zone specification.

%Z

The timezone name.

Similarly, because of GNU extensions to *strftime*, %k is accepted as a synonym for %H, and %l should be accepted as a synonym for %I, and %P is accepted as a synonym for %p. Finally

%s

The number of seconds since the epoch, i.e., since 1970-01-01 00:00:00 UTC. Leap seconds are not counted unless leap second support is available.

APPENDIX D

System Rules and Reports

This appendix presents the list of system rules and reports and provides a brief description of their intended use.

This chapter contains the following topics:

- List of System Rules, page D-1
- List of System Reports, page D-13

List of System Rules

This topic defines the complete list of system rules issued with this release.

- **System Rule: Backdoor: Active.**

  This correlation rule detects a connection to a backdoor server or a response from a backdoor server in your network accompanied by malicious follow-up activity on the server hosting the backdoor - this may indicate that a malicious backdoor service is likely running in your network. Malicious follow-up activity includes excessive scans, denied packets, installation of malicious services, local buffer overflow attacks etc. Backdoors such as Unix rootkits or Trojan horses are malicious programs that offer extensive remote control of a host and may be left by an attacker on a compromised host to maintain future remote access.

- **System Rule: Backdoor: Connect.**

  This correlation rule detects a connection to a backdoor server or a response from a backdoor server in your network - there may or may not be any follow-up activity on the destination host. Backdoors (e.g. Rootkits, Trojan Horse programs) and command shells provide extensive remote control of a host and may be left by an attacker on a compromised host to maintain future remote access.

- **System Rule: Backdoor: Covert Channel.**

  This correlation rule detects communication over covert channels - this means DMZ services such as HTTP, DNS, ICMP, FTP, SMTP etc. are being misused to tunnel inappropriate traffic via those ports. DMZ services are chosen since firewalls permit them but may not perform deep protocol inspection. Either the source or the destination in this event may be compromised.

- **System Rule: Backdoor: Spyware.**

  This rule detects spyware e.g. Gator, Bonzi etc. installed on hosts or requests to hosts with spyware installed. Spyware are malicious applications that can be installed on a computer without the knowledge of the user, e.g. when one visits a web site or clicks on an advertising link or installs file

sharing freeware such as KaZaA, iMesh, and AudioGalaxy. Once installed, the spyware automatically runs each time the host PC is started and records URLs visited, the username, password, and credit card information used, and then sends this information to the spyware writers.

- **System Rule: Client Exploit - Attempt.**

  This rule detects a client workstation exploit - this means a workstation is either downloading executable content via Web or email or sending web requests that contain scripts or is the target of an (client side) exploit via protocols such as IRC, DHCP, DNS, P2P Worms.

- **System Rule: Client Exploit - Mass Mailing Worm.**

  This signature detects excessive amount of e-mail (at least 20/min) from a single host. To sharpen this rule for non-mail server hosts, create a group of mail server hosts and then create an exception by excluding these hosts in the source of this rule.

- **System Rule: Client Exploit - Sasser Worm.**

  This correlation rule detects a successful infection spread of the Sasser worm - an attack on port 445 followed by the any of the following (a)command shell connection to the victim on port 9996, (b) an FTP connection back to the victim on port 5554, (c) excessive scans on port 445 from the victim. This indicates that both the source and the destinations are likely infected with the Sasser worm. This worm exploits the Microsoft Windows vulnerability as described in Microsoft Security Bulletin MS04-011

- **System Rule: Client Exploit - Success Likely.**

  This correlation rule detects a client workstation exploit followed by the client performing anomalous activities. Client exploits include download of dynamically executable content via Web or email, web requests containing scripts, client side exploits via protocols such as IRC, DHCP, DNS, P2P Worms. Client anomalous activities include the client originating excessive denies and scans, attempting to connect to backdoors, propagating worms over the network. The presence of such activities may indicate that the client exploit is successful.

- **System Rule: Client Exploit - Sysbug Trojan.**

  This correlation rule detects a Sysbug Trojan exploit on a client workstation - the workstation downloaded executable content via email and the code executed and likely opened up Sysbug Trojan service on port 5555 to which other machines attempted to connect. Here, the source represents the client workstation and the destination represents the systems to which a connection is made after the trojan is installed.

- **System Rule: Configuration Issue: Firewall.**

  This rule detects configuration errors reported by a firewall - this may cause certain traffic to be dropped by the firewall.

- **System Rule: Configuration Issue: Server.**

  This rule detects configuration errors reported by a server - this may cause certain services to be not available at the server.

- **System Rule: Connectivity Issue: IOS IPS DTM.**

  This rule detects connectivity issues between CS-MARS and IOS - CS-MARS may not be able to dynamically turn on ACTIVE signatures on IOS.

- **System Rule: CS-MARS Database Partition Usage.**

This rule indicates that the current CS-MARS database partition filled up to 75% of its capacity and the next database partition will be purged soon to create space for new events. The estimated purge times are in the event message. This is normal CS-MARS activity and will result in old events and incidents to purged from CS-MARS database. Users are urged to archive CS-MARS data to prevent permanent data loss.

- **System Rule: CS-MARS Host Mitigation - Failure.**

  This rule triggers when CS-MARS is unable to successfully mitigate a host after having tried a few times.

- **System Rule: CS-MARS Host Mitigation - Success.**

  This rule triggers when CS-MARS is able to successfully mitigate a host.

- **System Rule: Database Privileged Command - Failures.**

  This correlation rule detects multiple failed attempts from the same database user to execute privileged database commands.

- **System Rule: DoS: Network - Attempt.**

  This rule detects network level denial of service (DoS) attacks along with relevant reconnaissance activity that may have preceded the attacks. Such attacks can create a dramatic increase in overall network traffic.

- **System Rule: DoS: Network - Success Likely.**

  This correlation rule detects the simultaneous occurrence of network level denial of service (DoS) attacks along with related events such as traffic anomaly (e.g. ICMP echo request/reply or TCP SYN/FIN anomaly), network devices reporting high utilization, excessive scans or denies in the network etc. This may indicate that the network is under denial of service attack.

- **System Rule: DoS: Network Device - Attempt.**

  This correlation rule detects attacks on network devices (such as switches, routers, firewalls) along with relevant reconnaissance activity that may have preceded these attacks. Such attacks if successful, can crash the network devices and create a denial of service for the network segment containing these devices.

- **System Rule: DoS: Network Device - Success Likely.**

  This correlation rule detects attacks on network devices (such as switches, routers, firewalls) along with (a) local high usage conditions reported by the device and (b) relevant reconnaissance activity that may have preceded these attacks.

- **System Rule: Inactive CS-MARS Reporting Device.**

  This rule detects reporting devices that have not reported an event in the last hour. For chatty devices such as firewalls and IDS, this may indicate connectivity issues or an issue with the device themselves. This rule should be scoped down to only include chatty network infrastructure devices.

- **System Rule: Local Attack - Attempt.**

  This correlation rule detects attacks on hosts by logged on users. Such attacks include local buffer overflow attacks, sym link attacks etc.

- **System Rule: Local Attack - Success Likely.**

  This correlation rule detects attacks on hosts by locally logged on users followed by the server performing anomalous activities - such activities include excessive denies and scans, connection to backdoors, attempts to propagate worms etc. The presence of such activities may indicate that the host is compromised.

- **System Rule: Misc. Attacks: Access Web Customer Data.**

This correlation rule detects malicious attempts to access customer data stored by web applications, preceded by reconnaissance attempts to that host, if any. Customer data typically contains sensitive information such as purchasing history, credit card numbers etc.

- **System Rule: Misc. Attacks: Application Admin Escalation.**

  This correlation rule detects attempts by a non-administrative user to perform administrative functions for Web applications by bypassing the required authentication. Several web applications have vulnerabilities that may allow an attacker to do so. These attempts may be preceded by reconnaissance attempts to that host.

- **System Rule: Misc. Attacks: ARP Poisoning.**

  This correlation rule detects ARP Poisoning attacks preceded by reconnaissance attempts to that host, if any.

- **System Rule: Misc. Attacks: Evasion.**

  This correlation rule detects generic attempts by an attacker to bypass network IDS systems. The attempts may be preceded by reconnaissance attempts to that host.

- **System Rule: Misc. Attacks: Identity Spoofing.**

  This correlation rule detects attempts to used spoofed source IP addresses.

- **System Rule: Misc. Attacks: Replay.**

  This correlation rule detects replay attacks on a host, preceded by reconnaissance attempts to that host, if any. Successful replay attacks may allow the attacker to gain access by bypassing authentication.

- **System Rule: Misc. Attacks: Session Hijacking.**

  This correlation rule detects attempts to hijack a TCP connection to that host, preceded by reconnaissance attempts to that host, if any.

- **System Rule: Misc. Attacks: TCP/IP Protocol Anomaly.**

  This correlation rule detects events that indicate errors in standard TCP/IP headers - these may be caused by broken protocol implementations on the source host or may be malicious attempts by the source host to test the robustness of protocol implementations on the destination host.

- **System Rule: Modify Host: Database Object - Failures.**

  This correlation rule detects multiple failed attempts from the same database user to modify database objects such as tables, indices etc.

- **System Rule: Modify Host: Database User/Group - Failures.**

  This correlation rule detects multiple failed attempts from the same database user to modify database user groups

- **System Rule: Modify Host: Files.**

  This rule detects attempts to modify files on a host.

- **System Rule: Modify Host: Logs.**

  This rule detects attempts to modify log files on a host.

- **System Rule: Modify Host: Registry.**

  This rule detects attempts to modify windows registry entries on a host.

- **System Rule: Modify Host: Security.**

  This rule detects attempts to modify the security settings on a host.

- **System Rule: Modify Host: Service.**

This rule detects attempts to modify the settings of services on a host.

- **System Rule: Modify Host: User Group.**

  This rule detects attempts to modify the user group definitions on a host.

- **System Rule: Modify Network Config.**

  This rule detects attempts to modify the configurations on a network device such as routers, switches, firewalls etc.

- **System Rule: Modify Server: SCADA Modbus.**

  This rule detects attempts to modify the counters and diagnostics on a Modbus Servers. Modbus protocol is the defacto standard in industrial control communications and is the protocol of choice in a Supervisory Control and Data Acquisition (SCADA) communications network, where the Programmable logic controllers (PLCs) act as Modbus servers.

- **System Rule: Network Activity: Chat/IM - Active.**

  This rule detects person-to-person Chat or Instant Messenger protocol activity.

- **System Rule: Network Activity: Chat/IM - File Transfer.**

  This rule detects file transfers via person-to-person Chat or Instant Messengers along with increase in network traffic if any. File transfer is not a normal use of Chat/IM and is suspicious. In addition, files shared with other IM users could contain viruses or other backdoor programs.

- **System Rule: Network Activity: Excessive Denies - Host Compromise Likely.**

  This correlation rule detects a large frequency (excess of 10/sec) of denies from a particular host to a particular destination port. This is a typical behavior of a compromised host looking to exploit hosts with a specififc vulnerability.

- **System Rule: Network Activity: Excessive IRC.**

  This correlation rule detects excessive Internet relay Chat (IRC) connections from the same source - this indicates that a Remote Admin Trojan (RAT) is likely running on the source and is likely compromised.

- **System Rule: Network Activity: P2P File Sharing - Active.**

  This rule detects person-to-person file sharing activity via applications such as KaZaa, Napster, EDonkey, Gnutella, Bearshare etc.

- **System Rule: Network Activity: P2P File Sharing - File Transfer.**

  This rule detects a file transfer via a person-to-person file sharing application such as KaZaa, Napster, EDonkey, Gnutella, Bearshare etc. along with increase in network traffic if any. The programs may consume significant amount of network bandwidth and furthermore, inappropriate materials possibly containing viruses and backdoors may be distributed.

- **System Rule: Network Activity: Recreational.**

  This rule detects recreational activities such as games, visiting adult web sites etc.

- **System Rule: Network Activity: Uncommon Traffic.**

  This rule detects traffic that are not common in modern networks, for example (a) uncommon ICMP types - ICMP Router advertisement, ICMP Timestamp request/reply etc., (b) packets with uncommon TCP/IP options such source routing, timestamp etc, (c) standard protocols such as SMTP, HTTP, POP3 running on non-standard ports, (d) uncommon protocols such as FSP.

- **System Rule: Network Activity: Windows Popup Spam.**

  This correlation detects excessive traffic (likely pop up spam) from the same source to the Windows Messenger service.

- **System Rule: Network Errors - Likely Routing Related.**

  This rule detects a large frequency of denied packets or ICMP destination unreachable events between the same source, destination pair - this may indicate a network routing error and may be caused by periodic retransmission attempts by TCP or the application itself (e.g. DNS).

- **System Rule: New Malware Discovered.**

  This rule detects that Cisco Incident Control Server (ICS) has received information about a new virus/worm/malware outbreak. ICS is going to deploy ACLs or signatures to routers and IPS devices

- **System Rule: New Malware Prevention Deployed.**

  This rule detects that Cisco Incident Control Server (ICS) has successfully deployed ACLs or signatures to routers and IPS devices in an attempt to prevent a newly discovered virus/worm/malware outbreak.

- **System Rule: New Malware Prevention Deployment Failed.**

  This rule detects that Cisco Incident Control Server (ICS) has failed to deploy ACLs or signatures to routers and IPS devices for preventing a new virus/worm/malware outbreak.

- **System Rule: New Malware Traffic Match.**

  This correlated rule detects a traffic pattern that (a) matches a worm pattern: same source to many distinct destinations and (b) matches the ACLs and signatures deployed by Cisco Incident Control Server (ICS) in response to a newly discovered virus/worm/malware outbreak.

- **System Rule: Operational Issue: Firewall.**

  This rule detects operational errors (e.g. bad network connectivity, failover errors, internal software/hardware errors) reported by a firewall - this may indicate that the firewall is not functioning properly.

- **System Rule: Operational Issue: IDS.**

  This rule detects operational errors reported by a intrusion detection system (IDS) - this may indicate that the device is not functioning properly.

- **System Rule: Operational Issue: Router / Switch.**

  This rule detects operational errors reported by non-security network devices such as routers and switches.

- **System Rule: Operational Issue: Server.**

  This rule detects operational errors reported by a host or by applications on a host - this may indicate that either the host or the specific application on the host is not functioning properly.

- **System Rule: Password Attack: Database - Attempt.**

  This correlation rule detects a password guessing attack to a database server, preceded by reconnaissance attacks to the host, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Database - Success Likely.**

  This correlation rule detects a password guessing attack on a database server followed by a successful logon. The attack may be preceded by reconnaissance attacks to the host. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Disabled Accounts.**

  This rule detects repeated failed password attempts on locked, expired or disabled accounts on a host

- **System Rule: Password Attack: FTP Server - Attempt.**

This correlation rule detects a password guessing attack to an FTP server, preceded by reconnaissance attacks to the host, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: FTP Server - Success Likely.**

  This correlation rule detects a password guessing attack on a FTP server followed by a successful logon. The attack may be preceded by reconnaissance attacks to the host. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Mail Server - Attempt.**

  This correlation rule detects a password guessing attack on a mail server (SMTP, POP, IMAP), preceded by reconnaissance attacks to the host, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Mail Server - Success Likely.**

  This correlation rule detects a password guessing attack on a mail server (SMTP, POP, IMAP) followed by a successful logon. The password attack may be preceded by reconnaissance attacks to the host. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Misc. Application - Attempt.**

  This correlation rule detects attempts to retrieve application passwords or multiple login failures while authenticating to a particular application. These attempts can be optionally preceded by reconnaissance attempts. Authentication failures may sometimes be caused by a user forgetting the password. The applications covered by this rule exclude common ones such as Mail, FTP, SSH, Telnet, SNMP, Network/File/Print share, for which there are special rules.

- **System Rule: Password Attack: Network Share - Attempt.**

  This correlation rule detects a password guessing attack on a network share, preceded by reconnaissance attacks, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Network Share - Success Likely.**

  This correlation rule detects a password guessing attack on a network share, followed by a successful logon. The password attack may be preceded by reconnaissance attacks to the host. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Remote VPN Access - Attempt.**

  This correlation rule detects a password guessing attack while authenticating to a remote access service (e.g. Windows L2TP, PPTP based RAS, IPSec etc.), preceded by reconnaissance attacks, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Remote VPN Access - Success Likely.**

  This correlation rule detects a password guessing attack while authenticating to a remote access service (e.g. Windows L2TP, PPTP based RAS, IPSec etc.), followed by a successful logon. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: SNMP - Attempt.**

This correlation rule detects attempts to retrieve SNMP community strings or access SNMP information by guessing SNMP community strings. Many SNMP installations have easily guessable passwords by default. The password attack may be preceded by reconnaissance attacks to the host.

- **System Rule: Password Attack: SNMP - Success Likely.**

  This correlation rule detects a likely successful SNMP community string guessing attack - such an attack consists of a community string guessing attempt followed by a SNMP modification at the target host. The attack may be preceded by reconnaissance attacks to the host.

- **System Rule: Password Attack: System - Attempt.**

  This correlation rule detects attempts a to retrieve system passwords or multiple login failures while authenticating to a particular system/domain via telnet, SSH or local console/terminal logon. These attempts can be optionally preceded by reconnaissance attempts. Authentication failures may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: System - Success Likely.**

  This correlation rule detects a successful password attack to gain system level access to a host or to a windows domain- such an attack consists of a successful login occurring after attempts to retrieve passwords or guess passwords while authenticating to that host. The password attack may be preceded by reconnaissance attacks to the host. Authentication failures may sometimes be caused by a user forgetting the password.

- **System Rule: Password Attack: Web Server - Attempt.**

  This correlation rule detects a password guessing attack to a Web server, preceded by reconnaissance attacks to the host, if any. A password guessing attack consists of multiple login failures and may sometimes be caused by a user forgetting the password.

- **System Rule: Password Scan: Disabled Accounts: Distinct Hosts.**

  This rule detects repeated failed password attempts on locked, expired or disabled accounts on distinct hosts.

- **System Rule: Password Scan: Disabled Accounts: Same Host.**

  This rule detects repeated failed password attempts on distinct locked, expired or disabled accounts on a host.

- **System Rule: Password Scan: Distinct Hosts.**

  This rule detects repeated failed password attempts on distinct hosts.

- **System Rule: Password Scan: Same Host.**

  This rule detects repeated failed password attempts on multiple distinct accounts on the same host.

- **System Rule: Resource Issue: CS-MARS.**

  This rule detects resource issues with the CS-MARS device, e.g. dropped events or netflow, etc.

- **System Rule: Resource Issue: Host.**

  This rule detects resource issues at a host, e.g. event log being full, disk near capacity, too many logged in users etc.

- **System Rule: Resource Issue: IOS IPS DTM.**

  This rule detects that a Cisco IOS router has too little memory for running the required set of ACTIVE IPS signatures. CS-MARS was not successful in downloading the complete ACTIVE signature set.

- **System Rule: Resource Issue: Network Device.**

This rule detects resource issues at a network device, e.g. router, switch, firewall or IDS. Such issues include high CPU usage, a firewall reaching session limit, insufficient memory etc.

- **System Rule: Scans: SCADA Modbus.**

  This correlation rule detects scans targeted at Modbus servers. Modbus protocol is the defacto standard in industrial control communications and is the protocol of choice in a Supervisory Control and Data Acquisition (SCADA) communications network, where the Programmable logic controllers (PLCs) act as Modbus servers.

- **System Rule: Scans: Stealth.**

  This rule detects highly suspicious scans that are performed by sending malformed TCP/IP packets with an intent to discover host and application characteristics such as OS name, OS version etc. A vulnerability assessment tool such as Nmap can generate such scans. The source of the scans, if from inside the trusted network, must be investigated to see if it is from an authorized source. A MARS appliance may be performing such a test as part of false positive analysis.

- **System Rule: Scans: Targeted.**

  This rule detects scans that are either (a) targeted at a host to identify its operating environment, such as users on a host, DNS version, RPC services open etc. or (b) targeted at a well-known service to determine the set of host that offer that service.

- **System Rule: Security Posture: Audit Server Issue - Network wide.**

  This rule detects excessive number of logs indicating network wide audit server issues - the indications can come from many hosts staying in TRANSITION posture state for too long or many AAA server reporing Audit Server communication problems. These events may indicate that the audit server is having difficulty in auditing and updating the end host security posture status from TRANSITION state. A host enters the TRANSITION state when it is not running the Cisco Trust Agent (CTA) software and requires an out-of-band audit by an audit server to move it out of TRANSITION state to any one of HEALTHY, INFECTED, QUARANTINE, CHECKUP or UNKNOWN states. A host in a TRANSITION state is likely to have limited or no network access.

- **System Rule: Security Posture: Audit Server Issue - Single Host.**

  This rule detects excessive number of logs indicating audit server issues for a single host - the indications can come from the host staying in TRANSITION posture state for too long or AAA server reporing Audit Server communication problems for the same host. These events may indicate that the audit server is having difficulty in auditing and updating the end host security posture status from TRANSITION state. A host enters the TRANSITION state when it is not running the Cisco Trust Agent (CTA) software and requires an out-of-band audit by an audit server to move it out of TRANSITION state to any one of HEALTHY, INFECTED, QUARANTINE, CHECKUP or UNKNOWN states. A host in a TRANSITION state is likely to have limited or no network access.

- **System Rule: Security Posture: Excessive NAC Status Query Failures - Network wide.**

  This rule detects excessive network-wide NAC status query failures reported by distinct end host, Network Access Device (NAD) combinations. A Status query failure indicates a change in posture detected by the Cisco Trust Agent (CTA) after the initial authorization. Excessive status query failures may indicate a sign of end point instability caused by the user enabling or disabling agents. Excessive status query failures reported by distinct NAD and end host combinations may indicate a critical software problem..

- **System Rule: Security Posture: Excessive NAC Status Query Failures - Single Host.**

  This rule detects excessive NAC status query failures from the same end host. A Status query failure indicates a change in posture detected by the Cisco Trust Agent (CTA) after the initial authorization. Excessive status query failures may indicate a sign of end point instability caused by the user enabling or disabling agents. The end host may be compromised; at least this behavior is suspicious.

- **System Rule: Security Posture: Excessive NAC Status Query Failures - Single NAD.**

  This rule detects excessive NAC status query failures from distinct hosts to the same Network Access Device (NAD). A Status query failure indicates a change in posture detected by the Cisco Trust Agent (CTA) after the initial authorization. Excessive status query failures may indicate a sign of end point instability caused by the user enabling or disabling agents. Excessive status query failures from distinct hosts reported by the same NAD may indicate a problem at the NAD.

- **System Rule: Security Posture: Infected - Network wide.**

  This rule detects that many distinct hosts are reporting INFECTED security posture status for an excessive period of time. This implies that a significant number of hosts are having trouble getting cleaned.

- **System Rule: Security Posture: Infected - Single Host.**

  This rule detects that a particular host is reporting INFECTED security posture status for an excessive period of time. This implies that the host is having trouble getting cleaned.

- **System Rule: Security Posture: Quarantined - Network wide.**

  This rule detects that many distinct hosts are reporting QUARANTINED security posture status for an excessive period of time. This implies that a significant number of hosts are having trouble getting DAT file updates.

- **System Rule: Security Posture: Quarantined - Single Host.**

  This rule detects that a particular host is reporting QUARANTINE security posture status for an excessive period of time. This implies that the host is having trouble getting DAT file updates.

- **System Rule: Server Attack: Database - Attempt.**

  This correlation rule detects attacks on a database server, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, denial of service attempts, SQL Injection and other remote command execution attempts using database server privileges.

- **System Rule: Server Attack: Database - Success Likely.**

  This correlation rule detects specific attacks on a database server followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to a database server include buffer overflows, denial of service attempts, SQL Injection and other remote command execution attempts using database server privileges.

- **System Rule: Server Attack: DNS - Attempt.**

  This correlation rule detects specific attacks on a DNS host, preceded by reconnaissance attempts targeted to that host, if any. Attacks on a DNS host includes buffer overflow attempts, denial of service attempts.

- **System Rule: Server Attack: DNS - Success Likely.**

  This correlation rule detects likely successful attacks on a DNS host - an attack is successful if it is followed by suspicious activity on the targeted DNS server. Suspicious activity includes the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host.

- **System Rule: Server Attack: FTP - Attempt.**

  This correlation rule detects attacks on a FTP server, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, remote command execution attempts using FTP server privileges, denial of service attempts.

- **System Rule: Server Attack: FTP - Success Likely.**

This correlation rule detects specific attacks on a FTP server followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to a FTP server include buffer overflows, remote command execution attempts using FTP server privileges, denial of service attempts.

- **System Rule: Server Attack: Login - Attempt.**

  This correlation rule detects attacks on login services on a host, preceded by reconnaissance attempts targeted to that host, if any. Login services include Telnet, SSH, R-protocols such as Rsh, Rlogin, Rexec etc. The attacks include buffer overflows, privilege escalation attempts to become root, denial of service attempts etc.

- **System Rule: Server Attack: Login - Success Likely.**

  This correlation rule detects specific attacks on login services on a host (e.g. Telnet, SSH, R-protocols such as Rsh, Rlogin, Rexec etc.) followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to a login server include buffer overflows, remote command execution attempts using the server privileges, denial of service attempts.

- **System Rule: Server Attack: Mail - Attempt.**

  This correlation rule detects attacks on mail services (SMTP, POP, IMAP) on a host, preceded by reconnaissance attempts targeted to that host, if any. The attacks to mail services include buffer overflows, remote command execution attempts, privilege escalation attempts to become root, denial of service attempts etc.

- **System Rule: Server Attack: Mail - Success Likely.**

  This correlation rule detects specific attacks on mail services (SMTP, POP, IMAP) on a host followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to a mail server include buffer overflows, remote command execution attempts using server privileges, denial of service attempts.

- **System Rule: Server Attack: Misc. - Attempt.**

  This correlation rule detects attacks on miscellaneous services (i.e. other than DNS, FTP, HTTP, Mail, FTP, RPC, Telnet, SSH, R-protocols) on a host, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, remote command execution attempts, privilege escalation attempts to become root, denial of service attempts etc.

- **System Rule: Server Attack: Misc. - Success Likely.**

  This correlation rule detects specific attacks on miscellaneous services (i.e. other than DNS, FTP, HTTP, Mail, FTP, RPC, Telnet, SSH, R-protocols) on a host followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks include buffer overflows, remote command execution attempts using server privileges, denial of service attempts etc.

- **System Rule: Server Attack: RPC - Attempt.**

  This correlation rule detects attacks on RPC services on a host, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, remote command execution attempts, privilege escalation attempts to become root, denial of service attempts etc.

- **System Rule: Server Attack: RPC - Success Likely.**

This correlation rule detects specific attacks on RPC services on a host followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to RPC services include buffer overflows, remote command execution attempts using system privileges, denial of service attempts.

- **System Rule: Server Attack: SCADA Modbus - Attempt.**

  This correlation rule detects attacks on Modbus servers, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, denial of service attempts etc. Modbus protocol is the defacto standard in industrial control communications and is the protocol of choice in a Supervisory Control and Data Acquisition (SCADA) communications network, where the Programmable logic controllers (PLCs) act as Modbus servers.

- **System Rule: Server Attack: Sniffer - Attempt.**

  This correlation rule detects denial of service attacks on a host in promiscuous host (e.g. a network IDS host).

- **System Rule: Server Attack: Sniffer - Success Likely.**

  This correlation rule detects denial of service attacks on a host in promiscuous host (e.g. a network IDS host) followed by the destination host reporting functionally anomalous behavior.

- **System Rule: Server Attack: SNMP - Attempt.**

  This correlation rule detects attacks on SNMP implementation on a host, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, privilege escalation attempts to become root, etc.

- **System Rule: Server Attack: SNMP - Success Likely.**

  This correlation rule detects specific attacks on SNMP implementation on a host followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks to RPC services include buffer overflows, remote command execution attempts using system privileges, denial of service attempts.

- **System Rule: Server Attack: Web - Attempt.**

  This correlation rule detects attacks on a web server, preceded by reconnaissance attempts targeted to that host, if any. The attacks include buffer overflows, remote command execution attempts, denial of service attempts etc.

- **System Rule: Server Attack: Web - Success Likely.**

  This correlation rule detects specific attacks on a web server followed by suspicious activity on the targeted host. Suspicious activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc. The attack may be preceded by reconnaissance attempts targeted to that host. The attacks include buffer overflows, remote command execution attempts, denial of service attempts etc.

- **System Rule: State Change: Host.**

  This correlation rule detects significant host status change events such as system failing, rebooting, interface cards coming up and down, audit log filling up or getting deleted etc...

- **System Rule: State Change: Network Device.**

  This correlation rule detects significant network status state change events such as system failing, failover occuring, interface cards coming up and down etc.

- **System Rule: State Change: SCADA Modbus.**

This rule detects Modbus servers restarting. Modbus protocol is the defacto standard in industrial control communications and is the protocol of choice in a Supervisory Control and Data Acquisition (SCADA) communications network, where the Programmable logic controllers (PLCs) act as Modbus servers.

- **System Rule: Sudden Traffic Increase To Port.**

  This rule detects scans statistically significant increase in traffic to a particular port.

- **System Rule: Virus Found - Cleaned.**

  This rule indicates that virus scanning software detected and was able to clean a virus.

- **System Rule: Virus Found - Not Cleaned.**

  This rule indicates that virus scanning software detected but was unable to clean a virus.

- **System Rule: Vulnerable Host Found.**

  This rule detects a vulnerable host in the network - such hosts typically run old vulnerable protocols (e.g. SSH version 1, Rexec) or authenticate using plaintext passwords.

- **System Rule: Worm Propagation - Attempt.**

  This correlation rule detects worm propagation via means such as SMTP, TFTP, and network shares.

- **System Rule: Worm Propagation - Success Likely.**

  This correlation rule detects worm propagation via means such as SMTP, TFTP, and network shares accompanied by suspicious follow-up activity at the target destination host. Suspicious follow-up activity may include the host scanning the network, creating excessive firewall deny traffic, a backdoor opening up at the server etc.

List of System Reports

This topic defines the complete list of system reports issued with this release.

- **[MARS Internal: Netflow: Top Destination Ports].**

  This report ranks the destination ports in events seen by MARS. This is for internal use only.

- **[MARS Internal: Netflow: Top Destination Ports].**

  [MARS Internal: Netflow: Top Destination Ports]

- **[MARS Internal: Netflow: Top Hosts/Destination Ports Byte Count].**

  This report ranks the destination ports in events seen by MARS. This is for internal use only.

- **[MARS Internal: Netflow: Top Hosts/Destination Ports Byte Count].**

  [MARS Internal: Netflow: Top Hosts/Destination Ports Byte Count]

- **[MARS Internal: Netflow: Top Hosts/Destination Ports Flow Count].**

  This report ranks the destination ports in events seen by MARS. This is for internal use only.

- **[MARS Internal: Netflow: Top Hosts/Destination Ports Flow Count].**

  [MARS Internal: Netflow: Top Hosts/Destination Ports Flow Count]

- **Activity: AAA Based Access - All Events.**

  This report details AAA based access (e.g. to the network or to specific devices).

- **Activity: AAA Based Access - All Events.**

  Activity: AAA Based Access - All Events

- **Activity: AAA Based Access Failure - All Events.**

  This report details all failed AAA (e.g. RADIUS, TACACS) based access attempts. Typically mechanisms such as 802.1x, network device access, Cisco NAC use AAA servers for access control.

- **Activity: AAA Based Access Failure - All Events.**

  Activity: AAA Based Access Failure - All Events

- **Activity: AAA Failed Auth - All Events.**

  This report displays event details on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credentials.

- **Activity: AAA Failed Auth - All Events.**

  This report displays event details on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credentials.

- **Activity: AAA Failed Auth - Top NADs.**

  This report ranks the Network Access Devices (NADs) based on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credentials.

- **Activity: AAA Failed Auth - Top NADs.**

  This report ranks the Network Access Devices (NADs) based on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credenti

- **Activity: AAA Failed Auth - Top Users.**

  This report ranks the users based on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credentials.

- **Activity: AAA Failed Auth - Top Users.**

  This report ranks the users based on failed AAA authentications. This report covers the following cases: regular AAA auth, 802.1x auth, L2 IP and L3 IP auth, L2 802.1x auth. An authentication may fail because of policy misconfiguration on the AAA server or wrong user credentials.

- **Activity: Accounts Locked - All Events.**

  This report details events that indicate locked computer accounts because of excessive login failures.

- **Activity: Accounts Locked - All Events.**

  Activity: Accounts Locked - All Events

- **Activity: Accounts Locked - Top Hosts.**

  This report ranks the hosts by the accounts locked.

- **Activity: Accounts Locked - Top Hosts.**

  Activity: Accounts Locked - Top Hosts

- **Activity: All - NAT Connections.**

  This report lists Network Address Translations performed on non-denied sessions as reported to MARS.

- **Activity: All - NAT Connections.**

  This report lists Network Address Translations performed on non-denied sessions as reported to MARS.

- **Activity: All - Top Destination Ports.**

  This report ranks the UDP and TCP destination ports of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.

- **Activity: All - Top Destination Ports.**

  Activity: All - Top Destination Ports

- **Activity: All - Top Destinations.**

  This report ranks the session destinations of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.

- **Activity: All - Top Destinations.**

  Activity: All - Top Destinations

- **Activity: All - Top Event Type Groups.**

  This report ranks event type groups by reported events that belong to each group. The event type groups give a general feeling about the type of network activity reported to MARS.

- **Activity: All - Top Event Type Groups.**

  Activity: All - Top Event Type Groups

- **Activity: All - Top Event Types.**

  This report ranks the event types of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.

- **Activity: All - Top Event Types.**

  Activity: All - Top Event Types

- **Activity: All - Top Reporting Device Types.**

  This report ranks security device types by the number events reported by devices of each particular type.

- **Activity: All - Top Reporting Device Types.**

  Activity: All - Top Reporting Device Types

- **Activity: All - Top Reporting Devices.**

  This report ranks security devices by the total number of events reported by each device. This report is used by pages in the Summary tab.

- **Activity: All - Top Reporting Devices.**

  Activity: All - Top Reporting Devices

- **Activity: All - Top Sources.**

  This report ranks the session sources of all events seen by MARS over the past hour. This report is used by pages in the Summary tab.

- **Activity: All - Top Sources.**

  Activity: All - Top Sources

- **Activity: All - Top Users.**

This report tracks the most frequent logins and other user activity by showing the most active user names.

- **Activity: All - Top Users.**

  Activity: All - Top Users

- **Activity: All Events and Netflow - Top Destination Ports.**

  This report ranks the UDP and TCP destination ports of all events (including Netflow events) seen by MARS over the past hour. This report is used by pages in the Summary tab.

- **Activity: All Events and Netflow - Top Destination Ports.**

  Activity: All Events and Netflow - Top Destination Ports

- **Activity: All Sessions - Top Destination Ports by Bytes.**

  This report ranks all destination ports by bytes transferred.

- **Activity: All Sessions - Top Destination Ports by Bytes.**

  Activity: All Sessions - Top Destination Ports by Bytes

- **Activity: All Sessions - Top Destinations by Bytes.**

  This report ranks all destinations by bytes transferred.

- **Activity: All Sessions - Top Destinations by Bytes.**

  Activity: All Sessions - Top Destinations by Bytes

- **Activity: Attacks Prevented - Top Reporting Devices.**

  This report ranks security devices by the number of attacks prevented.

- **Activity: Attacks Prevented - Top Reporting Devices.**

  Activity: Attacks Prevented - Top Reporting Devices

- **Activity: Attacks Seen - Top Event Types.**

  This report ranks the top attack event types.

- **Activity: Attacks Seen - Top Event Types.**

  Activity: Attacks Seen - Top Event Types

- **Activity: Attacks Seen - Top Reporting Devices.**

  This report ranks security devices by the number of attack events logged. The security devices can be firewalls, NIDS and HIDS.

- **Activity: Attacks Seen - Top Reporting Devices.**

  Activity: Attacks Seen - Top Reporting Devices

- **Activity: Backdoor - Top Destinations.**

  This report ranks the hosts that respond to backdoor connection attempts.

- **Activity: Backdoor - Top Destinations.**

  Activity: Backdoor - Top Destinations

- **Activity: Backdoor - Top Event Types.**

  This report ranks the events that detect some form of backdoor activity. A backdoor may be created by an attacker on a compromised host. A backdoor event can be either an attempt to connect to a backdoor or a response from a server running a backdoor.

- **Activity: Backdoor - Top Event Types.**

Activity: Backdoor - Top Event Types

- **Activity: Backdoor - Top Hosts.**

  This report ranks the hosts that respond to backdoor connection attempts. This means that the hosts are likely infected and running backdoors.

- **Activity: Backdoor - Top Hosts.**

  Activity: Backdoor - Top Hosts

- **Activity: CS-MARS Host Mitigation - Failure - All Events.**

  This report lists failed CS-MARS mitigation attempts - these can result from improper network connectivity or device access credentials.

- **Activity: CS-MARS Host Mitigation - Failure - All Events.**

  This report lists failed CS-MARS mitigation attempts - these can result from improper network connectivity or device access credentials.

- **Activity: CS-MARS Host Mitigation - Success - All Events.**

  This report lists successful mitigations from CS-MARS.

- **Activity: CS-MARS Host Mitigation - Success - All Events.**

  This report lists successful mitigations from CS-MARS.

- **Activity: Database Login Failures - All Events.**

  This report lists the event details for all database login failure events.

- **Activity: Database Login Failures - All Events.**

  Activity: Database Login Failures - All Events

- **Activity: Database Login Failures - Top Servers.**

  This report ranks the database servers by the number of login failures.

- **Activity: Database Login Failures - Top Servers.**

  Activity: Database Login Failures - Top Servers

- **Activity: Database Login Failures - Top Users.**

  This report ranks the users by the number of login failures.

- **Activity: Database Login Failures - Top Users.**

  Activity: Database Login Failures - Top Users

- **Activity: Database Login Successes - All Events.**

  This report lists event details for all successful database login events.

- **Activity: Database Login Successes - All Events.**

  Activity: Database Login Successes - All Events

- **Activity: Database Login Successes - Top Servers.**

  This report ranks the database server hosts by the number of successful logins.

- **Activity: Database Login Successes - Top Servers.**

  Activity: Database Login Successes - Top Servers

- **Activity: Database Login Successes - Top Users.**

  This report ranks the database users by the number of successful logins.

- **Activity: Database Login Successes - Top Users.**

  Activity: Database Login Successes - Top Users
- **Activity: Database Object Modification Failures - All Events.**

  This report lists the event details for all failed database object modification attempts.
- **Activity: Database Object Modification Failures - All Events.**

  Activity: Database Object Modification Failures - All Events
- **Activity: Database Object Modification Failures - Top Users.**

  This report ranks the users by the number of failed database object modification attempts.
- **Activity: Database Object Modification Failures - Top Users.**

  Activity: Database Object Modification Failures - Top Users
- **Activity: Database Object Modification Successes - All Events.**

  This report lists the event details for all successful database object modification attempts.
- **Activity: Database Object Modification Successes - All Events.**

  Activity: Database Object Modification Successes - All Events
- **Activity: Database Object Modification Successes - Top Users.**

  This report ranks the number of users by the number of successful database object modifications.
- **Activity: Database Object Modification Successes - Top Users.**

  Activity: Database Object Modification Successes - Top Users
- **Activity: Database Privileged Command Failures - All Events.**

  This report lists event details for all privileged database command execution failures.
- **Activity: Database Privileged Command Failures - All Events.**

  Activity: Database Privileged Command Failures - All Events
- **Activity: Database Privileged Command Failures - Top Users.**

  This report ranks the users by failed privileged database command execution attempts.
- **Activity: Database Privileged Command Failures - Top Users.**

  Activity: Database Privileged Command Failures - Top Users
- **Activity: Database Privileged Command Successes - All Events.**

  This report lists the event details for all successful privileged database commands executed.
- **Activity: Database Privileged Command Successes - All Events.**

  Activity: Database Privileged Command Successes - All Events
- **Activity: Database Privileged Command Successes - Top Users.**

  This report ranks the users by successful privileged database commands executed.
- **Activity: Database Privileged Command Successes - Top Users.**

  Activity: Database Privileged Command Successes - Top Users
- **Activity: Database Regular Command Failures - All Events.**

  This report lists the event details for all failed non-privileged database command execution attempts.
- **Activity: Database Regular Command Failures - All Events.**

Activity: Database Regular Command Failures - All Events

- **Activity: Database Regular Command Failures - Top Users.**

  This report ranks the users by the number of non-privileged database command execution attempts.
- **Activity: Database Regular Command Failures - Top Users.**

  Activity: Database Regular Command Failures - Top Users
- **Activity: Database Regular Command Successes - All Events.**

  This report lists the event details for all successful non-privileged database command executions.
- **Activity: Database Regular Command Successes - All Events.**

  Activity: Database Regular Command Successes - All Events
- **Activity: Database Regular Command Successes - Top Users.**

  This report ranks the users by successful non-privileged database command executions.
- **Activity: Database Regular Command Successes - Top Users.**

  Activity: Database Regular Command Successes - Top Users
- **Activity: Database User/Group Change Failures - All Events.**

  This report lists the event details for all failed database user/group modification attempts.
- **Activity: Database User/Group Change Failures - All Events.**

  Activity: Database User/Group Change Failures - All Events
- **Activity: Database User/Group Change Failures - Top Users.**

  This report ranks the users by the number of failed database user/group modification attempts.
- **Activity: Database User/Group Change Failures - Top Users.**

  Activity: Database User/Group Change Failures - Top Users
- **Activity: Database User/Group Change Successes - All Events.**

  This report lists the event details for all successful database user/group modifications.
- **Activity: Database User/Group Change Successes - All Events.**

  Activity: Database User/Group Change Successes - All Events
- **Activity: Database User/Group Change Successes - Top Users.**

  This report ranks the users by the successful database user/group modifications performed.
- **Activity: Database User/Group Change Successes - Top Users.**

  Activity: Database User/Group Change Successes - Top Users
- **Activity: Denies - Top Destination Ports.**

  This report ranks the destination ports to which attacks have been targetted but denied.
- **Activity: Denies - Top Destination Ports.**

  Activity: Denies - Top Destination Ports
- **Activity: Denies - Top Destinations.**

  This report ranks the destination hosts to which attacks have been targeted but denied.
- **Activity: Denies - Top Destinations.**

  Activity: Denies - Top Destinations
- **Activity: Denies - Top Sources.**

This report ranks attack sources by the number of denied connection attempts.

- **Activity: Denies - Top Sources.**

  Activity: Denies - Top Sources

- **Activity: Host Admin Login Success - All Events.**

  This report details successful administrative login events to hosts.

- **Activity: Host Admin Login Success - All Events.**

  This report details successful administrative login events to hosts.

- **Activity: Host Login Failures - All Events.**

  This report records all host login failure details.

- **Activity: Host Login Failures - All Events.**

  Activity: Host Login Failures - All Events

- **Activity: Host Login Failures - Top Destinations.**

  This report ranks hosts by the number of logon failures recorded.

- **Activity: Host Login Failures - Top Destinations.**

  Activity: Host Login Failures - Top Destinations

- **Activity: Host Login Failures - Top Users.**

  This report ranks host users by failed login attempts.

- **Activity: Host Login Failures - Top Users.**

  Activity: Host Login Failures - Top Users

- **Activity: Host Login Success - All Events.**

  This report details all host login success event details

- **Activity: Host Login Success - All Events.**

  Activity: Host Login Success - All Events

- **Activity: Host Login Success - Top Host.**

  This report ranks hosts by successful logins.

- **Activity: Host Login Success - Top Host.**

  Activity: Host Login Success - Top Host

- **Activity: Host Object Access - All Events.**

  This report records all Microsoft Windows Object Access events from Windows Event Logs.

- **Activity: Host Object Access - All Events.**

  Activity: Host Object Access - All Events

- **Activity: Host Privilege Escalation - All Events.**

  This report provides details for events that represent an user attempting to increase access rights on a particular host. Such attempts can happen remotely or from the local console and can be reported by Network or Host IDS devices or the hosts themselves

- **Activity: Host Privilege Escalation - All Events.**

  This report provides details for events that represent an user attempting to increase access rights on a particular host. Such attempts can happen remotely or from the local console and can be reported by Network or Host IDS devices or the hosts themselves

- **Activity: Host Privilege Escalation - Top Hosts.**

  This report records ranks the hosts by access privilege escalation attempts attempted against them. Such attempts can happen remotely or from the local console and can be reported by Network or Host IDS devices or the hosts themselves

- **Activity: Host Privilege Escalation - Top Hosts.**

  Activity: Host Privilege Escalation - Top Hosts

- **Activity: Host Privileged Access - All Events.**

  This report records all Microsoft Windows Host Privileged Access events from Windows Event Logs.

- **Activity: Host Privileged Access - All Events.**

  Activity: Host Privileged Access - All Events

- **Activity: Host Process Tracking - All Events.**

  This report records all Microsoft Windows Detailed Process Tracking events from Windows Event Logs.

- **Activity: Host Process Tracking - All Events.**

  Activity: Host Process Tracking - All Events

- **Activity: Host Registry Changes - All Events.**

  This report records the events signalling Microsoft Windows registry changes.

- **Activity: Host Registry Changes - All Events.**

  Activity: Host Registry Changes - All Events

- **Activity: Host Registry Changes - Top Host.**

  This report ranks hosts by the number of Microsoft Windows registry changes reported.

- **Activity: Host Registry Changes - Top Host.**

  Activity: Host Registry Changes - Top Host

- **Activity: Host Security Policy Changes - All Events.**

  This report lists all policy changes on a host affecting host security. These events are typically reported by Host IDS and host agents.

- **Activity: Host Security Policy Changes - All Events.**

  This report lists all policy changes on a host affecting host security. These events are typically reported by Host IDS and host agents.

- **Activity: Host Security Policy Changes - Top Host.**

  This report ranks hosts by the number of security policy changes on that host.

- **Activity: Host Security Policy Changes - Top Host.**

  Activity: Host Security Policy Changes - Top Host

- **Activity: Host System Events - All Events.**

  This report records the Microsoft Windows system events, e.g. startup, shutdown, LSA registration, audit event discards, etc.

- **Activity: Host System Events - All Events.**

  Activity: Host System Events - All Events

- **Activity: Host User/Group Management - All Events.**

This report recordss user group management events reported by hosts.

- **Activity: Host User/Group Management - All Events.**

  Activity: Host User/Group Management - All Events

- **Activity: Host User/Group Management - Top hosts.**

  This report ranks hosts by user group management events reported.

- **Activity: Host User/Group Management - Top hosts.**

  Activity: Host User/Group Management - Top hosts

- **Activity: IDS Evasion - Top Event Types.**

  This report ranks the events that detect an attempt by an attacker to evade detection by Network IDS systems. This may be web-based obfuscation attacks, fragmentation attacks or TCP/IP based attacks.

- **Activity: IDS Evasion - Top Event Types.**

  Activity: IDS Evasion - Top Event Types

- **Activity: Inactive Reporting Device - Top Devices.**

  This report lists devices that are configured to be reporting to CS-MARS bt haven't reported any event in the last hour.

- **Activity: Inactive Reporting Device - Top Devices.**

  This report lists devices that are configured to be reporting to CS-MARS bt haven't reported any event in the last hour.

- **Activity: IOS IPS DTM Successful Signature Tuning - All Events.**

  This report lists all successful IOS IPS signature download activities - both adition and deletion. CS-MARS Distributed Threat Mitigation (DTM) turns on ACTIVE IPS signatures on IOS routers.

- **Activity: IOS IPS DTM Successful Signature Tuning - All Events.**

  This report lists all successful IOS IPS signature download activities - both adition and deletion. CS-MARS Distributed Threat Mitigation (DTM) turns on ACTIVE IPS signatures on IOS routers.

- **Activity: IRC - All Events.**

  This report lists all IRC activities. Typically, worms deposit executables on infected hosts that initiate IRC connections.

- **Activity: IRC - All Events.**

  This report lists all IRC activities. Typically, worms deposit executables on infected hosts that initiate IRC connections.

- **Activity: Network Usage - Top Destination Ports By Bytes.**

  This report ranks the top destination ports by bytes sent and transmitted.

- **Activity: Network Usage - Top Destination Ports By Bytes.**

  This report ranks the top destination ports by bytes sent and transmitted.

- **Activity: Network Usage - Top Destination Ports.**

  This report ranks destination ports by number of network sessions. This report requires that the syslog level of routers or firewalls be set to high to be able to capture session events. This report provides a general usage pattern of the network.

- **Activity: Network Usage - Top Destination Ports.**

Activity: Network Usage - Top Destination Ports

- **Activity: New Malware Discovered - All Events.**

  This report lists all the new virus/worm/malware outbreaks discovered by Cisco Incident Control Server.

- **Activity: New Malware Discovered - All Events.**

  This report lists all the new virus/worm/malware outbreaks discovered by Cisco Incident Control Server.

- **Activity: New Malware Prevention Deployment Failure - All Events.**

  This report lists all devices to which ACL and signature deployment attempts by a Cisco Incident Control Server, in response to a new virus/worm/malware outbreak, failed.

- **Activity: New Malware Prevention Deployment Failure - All Events.**

  This report lists all devices to which ACL and signature deployment attempts by a Cisco Incident Control Server, in response to a new virus/worm/malware outbreak, failed.

- **Activity: New Malware Prevention Deployment Success - All Events.**

  This report lists all destinations (Cisco IOS IPS devices and IPS appliances) to which Cisco Incident Control Server has deployed new ACLs and signatures in respond to a new virus/worm/malware outbreak.

- **Activity: New Malware Prevention Deployment Success - All Events.**

  This report lists all destinations (Cisco IOS IPS devices and IPS appliances) to which Cisco Incident Control Server has deployed new ACLs and signatures in respond to a new virus/worm/malware outbreak.

- **Activity: New Malware Traffic Match - All Events.**

  This report details the traffic sources and the enforcing devices that match the ACLs and signatures deployed by the Cisco Incident Control Server in response to a newly discovered malware outbreak.

- **Activity: New Malware Traffic Match - All Events.**

  This report details the traffic sources and the enforcing devices that match the ACLs and signatures deployed by the Cisco Incident Control Server in response to a newly discovered malware outbreak.

- **Activity: New Malware Traffic Match - Top Sources.**

  This report lists the top sources that match the ACLs or signatures dynamically deployed by Cisco Incident Control Server in response to a new virus/worm/malware outbreak. This indicates that these sources are likely infected.

- **Activity: New Malware Traffic Match - Top Sources.**

  This report lists the top sources that match the ACLs or signatures dynamically deployed by Cisco Incident Control Server in response to a new virus/worm/malware outbreak. This indicates that these sources are likely infected.

- **Activity: P2P Filesharing/Chat - All Events.**

  This event details all P2P File sharing or Chat event details.

- **Activity: P2P Filesharing/Chat - All Events.**

  This event details all P2P File sharing or Chat event details.

- **Activity: P2P Filesharing/Chat - Top Event Types.**

This event ranks events detecting person-to-person file sharing protocol and chat protocol activity. File sharing protocols such as KaZaa, Napster, EDonkey and chat protocols such as IRC, Hotline and instant messaging protocols may not be suitable in business environments.

- **Activity: P2P Filesharing/Chat - Top Event Types.**

  Activity: P2P Filesharing/Chat - Top Event Types

- **Activity: P2P Filesharing/Chat - Top Hosts.**

  This report ranks hosts involved in P2P Filesharing and chat protocol activity. Such protocols may not be suitable in business environments.

- **Activity: P2P Filesharing/Chat - Top Hosts.**

  Activity: P2P Filesharing/Chat - Top Hosts

- **Activity: Recreational - All Events.**

  This event details all users involved in recreational activities such as games, specific web sites such as gambling etc.

- **Activity: Recreational - All Events.**

  This event details all users involved in recreational activities such as games, specific web sites such as gambling etc.

- **Activity: Recreational - Top Sources.**

  This report ranks the source addesses involved in recreational activities such as games, adult web sites, stock sites etc.

- **Activity: Recreational - Top Sources.**

  Activity: Recreational - Top Sources

- **Activity: Remote Access Login - All Events.**

  This report details of remote access login events (IPSec, SSLVPN, PPP, L2TP etc)

- **Activity: Remote Access Login - All Events.**

  This report details of remote access login events (IPSec, SSLVPN, PPP, L2TP etc)

- **Activity: Remote Access Login - Top User.**

  This report ranks users by remote access logins (PPP, L2TP, PPTP, IPSec).

- **Activity: Remote Access Login - Top User.**

  This report ranks users by remote access logins (PPP, L2TP, PPTP, IPSec).

- **Activity: Remote Access Login Failures - All Events.**

  This event details all failed remote access login event details.

- **Activity: Remote Access Login Failures - All Events.**

  This event details all failed remote access login event details.

- **Activity: Scans - Top Destination Ports.**

  This report ranks destination ports by the total number of events detecting scanning activity for that port. Scans involve activities such as searching for alive hosts, open services on such hosts and detecting host configuration and application settings.

- **Activity: Scans - Top Destination Ports.**

  Activity: Scans - Top Destination Ports

- **Activity: Scans - Top Destinations.**

This report ranks hosts by the total number of events detecting scanning activity directed to that host. Scans involve activities such as searching for alive hosts, open services on such hosts and detecting host configuration and application settings.

- **Activity: Scans - Top Destinations.**

  Activity: Scans - Top Destinations

- **Activity: Scans - Top Sources.**

  This report ranks an attack sources by the total number of events detecting scanning activity for certain services. Scans involve activities such as searching for alive hosts, open services on such hosts and detecting host configuration and application settings.

- **Activity: Scans - Top Sources.**

  Activity: Scans - Top Sources

- **Activity: Security Posture: Healthy - Top Users.**

  This report lists the users in a HEALTHY Security Posture State. A Healthy security posture implies that the posture of the host is up to date, policy compliant and does not need attention.

- **Activity: Security Posture: Healthy - Top Users.**

  Activity: Security Posture: Healthy - Top Users

- **Activity: Security Posture: NAC - Top NADs and Tokens.**

  This report displays the Network Access Devices (NADs) handling Network Admission Control transcations along with the tokens assigned by each of them.

- **Activity: Security Posture: NAC - Top NADs and Tokens.**

  This report displays the Network Access Devices (NADs) handling Network Admission Control transcations along with the tokens assigned by each of them.

- **Activity: Security Posture: NAC - Top NADs.**

  This report ranks the network access devices (NADs) handling Network Admission Control transcations.

- **Activity: Security Posture: NAC - Top NADs.**

  This report ranks the network access devices (NADs) handling Network Admission Control transcations.

- **Activity: Security Posture: NAC - Top Tokens.**

  This report shows the network wide distribution of NAC tokens. The possible token values are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC - Top Tokens.**

  This report shows the network wide distribution of NAC tokens. The possible token values are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC Agentless - Top Hosts.**

  This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC Agentless - Top Hosts.**

This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, I

- **Activity: Security Posture: NAC Agentless - Top NADs.**

  This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC Agentless - Top NADs.**

  This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, I

- **Activity: Security Posture: NAC Agentless - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC Agentless - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that do not have Cisco Trust Agent (CTA) software. In this case, the posture validation is done either locally by the Network Access Device or via the Audit Server. The possible NAC tokens values in this report are HEALTHY, CHECKUP, I

- **Activity: Security Posture: NAC Audit Server Issues - All Events.**

  This report ranks the end hosts for which the AAA server is having an issue with obtaining the right security posture token from the audit server. These hoend sts do not have the Cisco Trust Agent (CTA) running and they depend on an Audit Server for obtaining the proper Security Posture Token.

- **Activity: Security Posture: NAC Audit Server Issues - All Events.**

  This report ranks the end hosts for which the AAA server is having an issue with obtaining the right security posture token from the audit server. These hoend sts do not have the Cisco Trust Agent (CTA) running and they depend on an Audit Server for obtaining the proper Security Posture Token.

- **Activity: Security Posture: NAC End Host Details - All Events.**

  This report details all the NAC related messages from the Network Access Devices (NAD) and AAA servers. Choose a source IP address or user to see the messages for one end host.

- **Activity: Security Posture: NAC End Host Details - All Events.**

  This report details all the NAC related messages from the Network Access Devices (NAD) and AAA servers. Choose a source IP address or user to see the messages for one end host.

- **Activity: Security Posture: NAC Infected/Quarantine - All Events.**

  This report reports the event details for the hosts that are in an INFECTED or QUARANTINE state. The QUARANTINE hosts must do Anti-virus DAT file updates before network access and the INFECTED hosts must be cleaned before network access.

- **Activity: Security Posture: NAC Infected/Quarantine - All Events.**

This report reports the event details for the hosts that are in an INFECTED or QUARANTINE state. The QUARANTINE hosts must do Anti-virus DAT file updates before network access and the INFECTED hosts must be cleaned before network access.

- **Activity: Security Posture: NAC Infected/Quarantine - Top Hosts.**

  This report details the hosts that are in an INFECTED or QUARANTINE state. The QUARANTINE hosts must do Anti-virus DAT file updates before network access and the INFECTED hosts must be cleaned before network access.

- **Activity: Security Posture: NAC Infected/Quarantine - Top Hosts.**

  This report details the hosts that are in an INFECTED or QUARANTINE state. The QUARANTINE hosts must do Anti-virus DAT file updates before network access and the INFECTED hosts must be cleaned before network access.

- **Activity: Security Posture: NAC L2 802.1x - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that use Layer 2 IEEE 802.1x method to validate their posture. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC L2 802.1x - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that use Layer 2 IEEE 802.1x method to validate their posture. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC L2IP - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that use Layer 2 IP method to validate their posture. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC L2IP - Top Tokens.**

  This report captures the distribution of NAC tokens for end hosts that use Layer 2 IP method to validate their posture. The possible NAC tokens values in this report are HEALTHY, CHECKUP, INFECTED, QUARANTINE, UNKNOWN. The TRANSITION token is excluded since it is an intermediate state.

- **Activity: Security Posture: NAC Static Auth - Top Hosts.**

  This report captures the hosts that are configured as static exceptions on the Network Access Device (NAD). For these hosts, the NAD directly permits network access without consulting the posture validation server.

- **Activity: Security Posture: NAC Static Auth - Top Hosts.**

  This report captures the hosts that are configured as static exceptions on the Network Access Device (NAD). For these hosts, the NAD directly permits network access without consulting the posture validation server.

- **Activity: Security Posture: NAC Static Auth - Top NADs.**

  This report captures the Network Access Device (NAD) that are permitting end hosts into the network as static exceptions. For these end hosts, the NAD directly permits network access without consulting the posture validation server.

- **Activity: Security Posture: NAC Static Auth - Top NADs.**

This report captures the Network Access Device (NAD) that are permitting end hosts into the network as static exceptions. For these end hosts, the NAD directly permits network access without consulting the posture validation server.

- **Activity: Security Posture: NAC Status Query Failure - Top Hosts.**

  This report details the top hosts that failed the status queries from the Network Access Devices (NAD). Such failures occur after initial authorization whenever there is a change in posture detected by the Cisco Trust Agent (CTA) on the end host. Such failures may be caused by user frequently enabling or disabling CTA agents.

- **Activity: Security Posture: NAC Status Query Failure - Top Hosts.**

  This report details the top hosts that failed the status queries from the Network Access Devices (NAD). Such failures occur after initial authorization whenever there is a change in posture detected by the Cisco Trust Agent (CTA) on the end host. Such failures may be caused by user frequently enabli

- **Activity: Security Posture: Not Healthy - All Events.**

  This report lists the detailed events for users whose security posture is not up to date, ie. in either a CHECKUP, QUARANTINE or INFECTED state. The software on these hosts need to be upgraded. The CHECKUP hosts may need DAT file updates, the QUARANTINE hosts must do DAT file updates before network access and the INFECTED hosts must be remediated before network access.

- **Activity: Security Posture: Not Healthy - All Events.**

  This report lists the detailed events for users whose security posture is not up to date, ie. in either a CHECKUP, QUARANTINE or INFECTED state. The software on these hosts need to be upgraded. The CHECKUP hosts may need DAT file updates, the QUARANTINE hosts must do DAT file updates before network

- **Activity: Spyware - All Events.**

  This event details all spyware events.

- **Activity: Spyware - All Events.**

  This event details all spyware events.

- **Activity: Spyware - Top Hosts.**

  This report ranks the hosts running spyware applications. Spywares are malicious applications that installs and runs on hosts, collect the username, passwords, and credit card information and send this information to the spyware writers.

- **Activity: Spyware - Top Hosts.**

  Activity: Spyware - Top Hosts

- **Activity: Stealth Scans - Top Sources.**

  This report ranks attackers by the amount of stealth scanning activity. Such activities include sending crafted packets to detect host operating systems and other vulnerabilities. Vulnerability scanners may generate such events.

- **Activity: Stealth Scans - Top Sources.**

  Activity: Stealth Scans - Top Sources

- **Activity: Sudden Traffic Increase To Port - All Destinations.**

  This report lists hosts that exhibit anomalous behavior by suddenly receiving statistically significant volume on a TCP/UDP port or ICMP traffic.

- **Activity: Sudden Traffic Increase To Port - All Destinations.**

This report lists hosts that exhibit anomalous behavior by suddenly receiving statistically significant volume on a TCP/UDP port or ICMP traffic.

- **Activity: Sudden Traffic Increase To Port - All Sources.**

  This report lists hosts that exhibit anomalous behavior by suddenly sending statistically significant volume on a TCP/UDP port or ICMP traffic.

- **Activity: Sudden Traffic Increase To Port - All Sources.**

  This report lists hosts that exhibit anomalous behavior by suddenly sending statistically significant volume on a TCP/UDP port or ICMP traffic.

- **Activity: Uncommon or Anomalous Traffic - All Events.**

  This report details uncommon or anomalous traffic such as unused TCP options, uncommon ICMP traffic, non-standard traffic on standard port, tunneled traffic etc.

- **Activity: Uncommon or Anomalous Traffic - All Events.**

  This report details uncommon or anomalous traffic such as unused TCP options, uncommon ICMP traffic, non-standard traffic on standard port, tunneled traffic etc.

- **Activity: Unknown Events - All Events.**

  This report tracks the events that are unknown to MARS.

- **Activity: Unknown Events - All Events.**

  Activity: Unknown Events - All Events

- **Activity: Virus/Worms - Top Event Types.**

  This report ranks the events that detect virus or worm activity in the network.

- **Activity: Virus/Worms - Top Event Types.**

  Activity: Virus/Worms - Top Event Types

- **Activity: Virus/Worms - Top Infected Hosts.**

  This report ranks hosts that are propagating virus and worms via SMTP, POP, IMAP, network shares etc.

- **Activity: Virus/Worms - Top Infected Hosts.**

  Activity: Virus/Worms - Top Infected Hosts

- **Activity: Virus: Detected - Top Users.**

  This report ranks users/workstations by viruses detected.

- **Activity: Virus: Detected - Top Users.**

  Activity: Virus: Detected - Top Users

- **Activity: Virus: Infections - Top Users.**

  This report ranks users/workstations by viruses detected and not cleaned.

- **Activity: Virus: Infections - Top Users.**

  Activity: Virus: Infections - Top Users

- **Activity: Vulnerable Host Found via VA Scanner.**

  This report lists vulnerable hosts and associated vulnerabilities found by importing information from Vulnerability Analysis (VA) scanners.

- **Activity: Vulnerable Host Found via VA Scanner.**

This report lists vulnerable hosts and associated vulnerabilities found by importing information from Vulnerability Analysis (VA) scanners.

- **Activity: Vulnerable Host Found.**

  This host lists all vulnerable hosts found by IDS or VA scanners

- **Activity: Vulnerable Host Found.**

  This host lists all vulnerable hosts found by IDS or VA scanners

- **Activity: Web Usage - Top Destinations by Bytes.**

  This report ranks the web servers by bytes transferred.

- **Activity: Web Usage - Top Destinations by Bytes.**

  Activity: Web Usage - Top Destinations by Bytes

- **Activity: Web Usage - Top Destinations by Sessions.**

  This report ranks the top web destinations by session count.

- **Activity: Web Usage - Top Destinations by Sessions.**

  Activity: Web Usage - Top Destinations by Sessions

- **Activity: Web Usage - Top Sources.**

  This signature ranks source addresses based on web use.

- **Activity: Web Usage - Top Sources.**

  Activity: Web Usage - Top Sources

- **Attacks: All - All Events.**

  This event details details (event type, destination, source) for all attack events.

- **Attacks: All - All Events.**

  This event details details (event type, destination, source) for all attack events.

- **Attacks: All - Top Destinations.**

  This report ranks hosts by the number of attacks targetted at each host.

- **Attacks: All - Top Destinations.**

  Attacks: All - Top Destinations

- **Attacks: All - Top Event Type Groups.**

  This report ranks event type groups that appear in fired correlation rules. The event type groups give a general feeling about the network activity classified as part of an attack by MARS.

- **Attacks: All - Top Event Type Groups.**

  Attacks: All - Top Event Type Groups

- **Attacks: All - Top Rules Fired.**

  This report ranks rules fired over the past hour by number of incidents. This provides a general feeling about the attack activity in the network. This report is used by pages in the Summary tab.

- **Attacks: All - Top Rules Fired.**

  Attacks: All - Top Rules Fired

- **Attacks: All - Top Sources.**

  This report ranks the sources of attack events seen by MARS over the past hour.

- **Attacks: All - Top Sources.**

  Attacks: All - Top Sources

- **Attacks: Client Exploits - Top Sources.**

  This report ranks hosts by the number of exploits originating from each host.

- **Attacks: Client Exploits - Top Sources.**

  Attacks: Client Exploits - Top Sources

- **Attacks: Database Server - Top Event Types.**

  This report ranks attacks on database servers such as MS SQL Server, Oracle and Sybase.

- **Attacks: Database Server - Top Event Types.**

  Attacks: Database Server - Top Event Types

- **Attacks: FTP Server - Top Event Types.**

  This report ranks attacks on FTP servers.

- **Attacks: FTP Server - Top Event Types.**

  Attacks: FTP Server - Top Event Types

- **Attacks: Identity Spoofing - Top Event Types.**

  This report ranks events that represent attempts by an attacker to spoof his/her identity over the past hour.

- **Attacks: Identity Spoofing - Top Event Types.**

  Attacks: Identity Spoofing - Top Event Types

- **Attacks: Login Services - Top Event Types.**

  This report ranks attacks on servers providing login services and remote shells. Examples include Telnet, SSH and Berkeley r-protocols.

- **Attacks: Login Services - Top Event Types.**

  Attacks: Login Services - Top Event Types

- **Attacks: Mail Server - Top Event Types.**

  This report ranks attacks on Mail servers (SMTP, POP, IMAP).

- **Attacks: Mail Server - Top Event Types.**

  Attacks: Mail Server - Top Event Types

- **Attacks: Network DoS - Top Event Types.**

  This report ranks attacks that represent network wide denial of service attempts. Such attacks may include crashing or rebooting an inline network device such as router, firewall or switch or increasing network load by creating TCP, UDP or ICMP traffic.

- **Attacks: Network DoS - Top Event Types.**

  Attacks: Network DoS - Top Event Types

- **Attacks: Password - All Events.**

  This report details all password attack events.

- **Attacks: Password - All Events.**

  This report details all password attack events.

- **Attacks: Password - Top Destinations.**

This report ranks hosts by the number of password attacks attempted on them. Passwords attacks include attempts to (a) capture passwords, either remotely or locally and (b) guess passwords. Password guessing attempts are recorded as authentication failures by IDS and hosts.

- **Attacks: Password - Top Destinations.**

  Attacks: Password - Top Destinations

- **Attacks: Password - Top Event Types.**

  This report ranks password retrieving and guessing attacks. The password can be system passwords or application passwords.

- **Attacks: Password - Top Event Types.**

  Attacks: Password - Top Event Types

- **Attacks: Password: Locked Accounts - All Events.**

  This report details password attacks on locked/disabled/expired accounts.

- **Attacks: Password: Locked Accounts - All Events.**

  Attacks: Password: Locked Accounts - All Events

- **Attacks: Password: Restricted Times - All Events.**

  This report details all events that indicate login failures at restricted times - the hosts are specifically configured to disallow access at these hours.

- **Attacks: Password: Restricted Times - All Events.**

  Attacks: Password: Restricted Times - All Events

- **Attacks: RPC Services - Top Event Types.**

  This report ranks attacks on RPC based applications.

- **Attacks: RPC Services - Top Event Types.**

  Attacks: RPC Services - Top Event Types

- **Attacks: SANS Top 20 - Top Event Types.**

  This report ranks the attacks that have been included in SANS Top 20 list.

- **Attacks: SANS Top 20 - Top Event Types.**

  Attacks: SANS Top 20 - Top Event Types

- **Attacks: SNMP - Top Event Types.**

  This report ranks SNMP based attacks over the past hour.

- **Attacks: SNMP - Top Event Types.**

  Attacks: SNMP - Top Event Types

- **Attacks: Uncommon or Anomalous Traffic - Top Event Types.**

  This report ranks the events that represent uncommon or anomalous traffic. Uncommon traffic involves ICMP types and TCP/IP options not in common usage or standard traffic on non-standard ports. Anomalous traffic includes traffic that violate IETF or other well known protocol specifications.

- **Attacks: Uncommon or Anomalous Traffic - Top Event Types.**

  Attacks: Uncommon or Anomalous Traffic - Top Event Types

- **Attacks: Virus/Worms - Top Sources.**

  This report ranks addresses that are the source of virus/worm propagation attempts.

- **Attacks: Virus/Worms - Top Sources.**

  Attacks: Virus/Worms - Top Sources

- **Attacks: Web Server/App - Top Event Types.**

  This report ranks attacks on web servers or applications built on top of web servers over the past hour.

- **Attacks: Web Server/App - Top Event Types.**

  Attacks: Web Server/App - Top Event Types

- **Configuration Changes: Network - All Events.**

  This event details all the configuration changes in network devices.

- **Configuration Changes: Network - All Events.**

  This event details all the configuration changes in network devices.

- **Configuration Changes: Network - Top Event Types.**

  This report summarizes configuration changes to network devices such as firewalls, routers and switches over the past hour.

- **Configuration Changes: Network - Top Event Types.**

  Configuration Changes: Network - Top Event Types

- **Configuration Changes: Server - All Events.**

  This event details all configuration changes on hosts (reported by OS or Host IDS agents)

- **Configuration Changes: Server - All Events.**

  This event details all configuration changes on hosts (reported by OS or Host IDS agents)

- **Configuration Changes: Server - Top Event Types.**

  This report summarizes configuration changes to servers over the past hour.

- **Configuration Changes: Server - Top Event Types.**

  Configuration Changes: Server - Top Event Types

- **Configuration Changes: Server - Top Reporting Devices.**

  This report summarizes the configuration changes per server over the past hour.

- **Configuration Changes: Server - Top Reporting Devices.**

  Configuration Changes: Server - Top Reporting Devices

- **Configuration Issues: Network - All Events.**

  This report lists details for events that indicate configuration error on network devices.

- **Configuration Issues: Network - All Events.**

  This report lists details for events that indicate configuration error on network devices.

- **Configuration Issues: Network - Top Reporting Devices.**

  This report summarizes the events that may indicate certain configuration related problems in network devices such as firewalls, routers and switches.

- **Configuration Issues: Network - Top Reporting Devices.**

  Configuration Issues: Network - Top Reporting Devices

- **Configuration Issues: Server - All Events.**

This report lists details for all events that indicate configuration errors on hosts or host applications.

- **Configuration Issues: Server - All Events.**

  This report lists details for all events that indicate configuration errors on hosts or host applications.

- **Configuration Issues: Server - Top Reporting Devices.**

  This report summarizes the events that may indicate certain configuration related problems in servers. These are likely to be Host IDS events.

- **Configuration Issues: Server - Top Reporting Devices.**

  Configuration Issues: Server - Top Reporting Devices

- **Connectivity Issue: IOS IPS DTM - All Events.**

  This report lists connectivity issues between CS-MARS and IOS IPS devices. Connectivity issues may prevent CS-MARS from turning on ACTIVE signatures on IOS IPS.

- **Connectivity Issue: IOS IPS DTM - All Events.**

  This report lists connectivity issues between CS-MARS and IOS IPS devices. Connectivity issues may prevent CS-MARS from turning on ACTIVE signatures on IOS IPS.

- **Detailed NAC Report.**

  Detailed NAC Report

- **Detailed NAC Report.**

  Detailed NAC Report

- **Operational Issues: Network - All Events.**

  This report lists details about all operational issues on network devices.

- **Operational Issues: Network - All Events.**

  This report lists details about all operational issues on network devices.

- **Operational Issues: Network - Top Reporting Devices.**

  This report summarizes the events that may indicate operational issues with network devices such as routers, firewalls and Network IDS systems.

- **Operational Issues: Network - Top Reporting Devices.**

  Operational Issues: Network - Top Reporting Devices

- **Operational Issues: Server - All Events.**

  This report lists details about events that indicate operational errors on hosts or host applications.

- **Operational Issues: Server - All Events.**

  This report lists details about events that indicate operational errors on hosts or host applications.

- **Operational Issues: Server - Top Reporting Devices.**

  This report summarizes the events that may indicate operational issues with servers.

- **Operational Issues: Server - Top Reporting Devices.**

  Operational Issues: Server - Top Reporting Devices

- **Resource Issues: CS-MARS - All Events.**

  This report lists event details for all events related to resource issues with the CS-MARS device, e.g. dropped events or netflow, etc.

- **Resource Issues: CS-MARS - All Events.**

This report lists event details for all events related to resource issues with the CS-MARS device, e.g. dropped events or netflow, etc.

- **Resource Issues: IOS IPS DTM - All Events.**

  This report lists event details that indicate certin IOS IPS routers running low on memory for CS-MARS Distributed Threat Mitigation (DTM). Because of low memory, CS-MARS may not be able to download and activate the complete set of ACTIVE IPS signatures to those IOS IPS devices.

- **Resource Issues: IOS IPS DTM - All Events.**

  This report lists event details that indicate certin IOS IPS routers running low on memory for CS-MARS Distributed Threat Mitigation (DTM). Because of low memory, CS-MARS may not be able to download and activate the complete set of ACTIVE IPS signatures to those IOS IPS devices.

- **Resource Issues: IOS IPS DTM - Top Devices.**

  This report lists IOS IPS routers that are running low on memory for CS-MARS Distributed Threat Mitigation (DTM). Because of low memory, CS-MARS may not be able to download and activate the complete set of ACTIVE IPS signatures to IOS IPS devices.

- **Resource Issues: IOS IPS DTM - Top Devices.**

  This report lists IOS IPS routers that are running low on memory for CS-MARS Distributed Threat Mitigation (DTM). Because of low memory, CS-MARS may not be able to download and activate the complete set of ACTIVE IPS signatures to IOS IPS devices.

- **Resource Issues: Network - All Events.**

  This report lists event details for all events related to resource issues on network devices such as IDS, routers, firewalls etc.

- **Resource Issues: Network - All Events.**

  This report lists event details for all events related to resource issues on network devices such as IDS, routers, firewalls etc.

- **Resource Issues: Network - Top Reporting Devices.**

  This report summarizes the events that represent resource issues with network devices such as firewalls, routers and switches.

- **Resource Issues: Network - Top Reporting Devices.**

  Resource Issues: Network - Top Reporting Devices

- **Resource Issues: Server - All Events.**

  This report lists event details for all resource issues on hosts. These are reported by Host IDS or Operating System logs.

- **Resource Issues: Server - All Events.**

  This report lists event details for all resource issues on hosts. These are reported by Host IDS or Operating System logs.

- **Resource Issues: Server - Top Reporting Devices.**

  This report summarizes the events that represent resource issues with servers. These are likely to be Host IDS events.

- **Resource Issues: Server - Top Reporting Devices.**

  Resource Issues: Server - Top Reporting Devices

- **Resource Utilization: Bandwidth: Inbound - Top Interfaces.**

This report ranks the inbound bandwidth utilization of the interfaces on the devices managed by PN-MARS.

- **Resource Utilization: Bandwidth: Inbound - Top Interfaces.**

  Resource Utilization: Bandwidth: Inbound - Top Interfaces

- **Resource Utilization: Bandwidth: Outbound - Top Interfaces.**

  This report ranks the outbound bandwidth utilization of interfaces on devices managed by Pn-MARS.

- **Resource Utilization: Bandwidth: Outbound - Top Interfaces.**

  Resource Utilization: Bandwidth: Outbound - Top Interfaces

- **Resource Utilization: Concurrent Connections - Top Devices.**

  This report ranks the number of concurrent connections established through the devices managed by PN-MARS.

- **Resource Utilization: Concurrent Connections - Top Devices.**

  Resource Utilization: Concurrent Connections - Top Devices

- **Resource Utilization: CPU - Top Devices.**

  This report ranks the CPU utilization of the devices managed by PN-MARS.

- **Resource Utilization: CPU - Top Devices.**

  Resource Utilization: CPU - Top Devices

- **Resource Utilization: CS-MARS - All Events.**

  This report lists event details for all events related to CS-MARS resource utilization, e.g. database partitions, etc.

- **Resource Utilization: CS-MARS - All Events.**

  This report lists event details for all events related to CS-MARS resource utilization, e.g. database partitions, etc.

- **Resource Utilization: Errors: Inbound - Top Interfaces.**

  This report ranks by error rate on the inbound interfaces of the devices managed by PN-MARS.

- **Resource Utilization: Errors: Inbound - Top Interfaces.**

  Resource Utilization: Errors: Inbound - Top Interfaces

- **Resource Utilization: Errors: Outbound - Top Interfaces.**

  This report ranks by error rate on the outbound interfaces of the devices managed by PN-MARS.

- **Resource Utilization: Errors: Outbound - Top Interfaces.**

  Resource Utilization: Errors: Outbound - Top Interfaces

- **Resource Utilization: Memory - Top Devices.**

  This report ranks the memory utilization of the devices managed by PN-MARS.

- **Resource Utilization: Memory - Top Devices.**

  Resource Utilization: Memory - Top Devices

GLOSSARY

#

5-tuple (Quintuple) The five pieces of data found within all IP-based network packets: source IP address, source port, destination IP address, destination port, and protocol. You can define inspection rules, queries, and reports using the data found in the 5-tuple.

A

(\

Access IP Address This is the IP address that MARS uses to connect to the device and to get its configuration information. MARS needs this address for NAT-related session correlation, attack path calculation, and mitigation enter access information.

Activate Making changes or edits known to the MARS after submitting changes.

D

Devices The hosts and reporting devices present in the system.

Discovery The act of identifying, either automatically or manually, devices in networks.

Dynamic Vulnerability Scanning The MARS STM probes selected networks, and their components, for vulnerabilities.

E

Event A security event reported to the MARS STM appliance. Events have: types, sources, destinations, reporting devices, etc.

Event Types Groups of similar security events. An event type is the normalized signature from a reporting device.

F

False Positive An event that resembles a valid security threat, but is not.

Firing Events An event that contributed to a rule firing.

I

Incident Incidents are collections of events and sessions that meet the criteria for a rule, having helped to cause it to fire.

Incident Instances An instance of an incident.

M

MI B management information base

mitigate To stop a detected attack or anomaly. The method of mitigation varies based on network composition and configuration.

O

Offset The offset of a firing event is the line number of the rule criteria that this firing event matches.

P

Pre NAT Source Address Session endpoints.

Post NAT Source Address The source as appearing at the destination.

Post NAT Destination Address Session endpoints.

Pre NAT Destination Address The destination as appearing at the source.

Q

Query A user-defined request to the database for information.

R

Report A user-defined request to the database on an automatic or on-demand basis.

Reporting Device A discovered device that reports information – usually in the form of logs – to a MARS STM appliance.

Reporting IP Address This is the IP address as it appears to MARS. This address is where the logs (syslog, SNMP traps, LEA) come from.

Rule The sub-set of events that contributed to the incidents of the specified rules firing.

S

Service A protocol and range of IP addresses.

Session A session is a collection of events that all share a common source and destination, which were reported within a given time window. For example, usually the events in a session map well to the events generated between the opening and closing of a TCP/IP connection.

Sessionize Combining event data from multiple reporting devices to reconstruct the occurrence of a session. Sessionizing takes two forms: reconstructing a session-oriented protocol, such as TCP, where the initial handshake and the session tear down and reconstructing a sessionless protocol, such as UDP, where the initial start and session end times are defined more based on first and last packets tracked within a restricted time period. In other words, packets that fall outside of the time period are considered part of different sessions.

T

True Positive A valid security threat.

U

Unreported device A device from which the MARS Appliance receives events, such as syslog messages, SNMP notifications, or NetFlow events, but the device is not defined in the appliance. Without a definition, MARS is unable to correlate events correctly as it needs to know which message format to use in parsing.

T

True Positive A valid security threat.

INDEX

A

Action 5-3
Activate button 7-17, 7-18, 7-20, 9-1
adding
  cell phone number 8-11, 9-9
  devices 2-11
    manually 2-11
  event groups 9-3
  inspection rules 7-18
  pager number 8-11, 9-9
  service 9-6
  user 8-10, 9-7
  user group 9-10
adding IP groups 9-4
adding service provider 8-11, 9-9
admin roles, see user management 9-7
Adobe SVG 3-11
alert
  action 7-14
    Distributed Threat Management 7-14
    Email 7-14
    NONE 7-14
    Page 7-14
    SMS 7-14
    SNMP 7-14
    Syslog 7-14
  hard drive 10-16
alerts 8-1
all matching event raw messages 6-8
all matching events 6-8
all matching sessions 6-7
attack diagram 3-10
attack paths
  L2 5-5
  L3 5-5
audit trail 10-3

B

beep code 10-26
bytes transmitted 6-8

C

cell phone paging 8-11, 9-9
certificate
  monitor status 10-6
  upgrading from expired or fingerprint 10-6
changing
  inspection rule status 7-16
Collapse All 5-5
Common Vulneratbilities and Exposures 9-2
creating
  report 6-22
CVE 9-2

D

data reduction 3-10
default certificate response
  change 10-5
default fingerprint response
  change 10-5
default password
  change 10-3

deleting service 9-6
destination IP address ranking 6-7
destination network group ranking 6-7
destination network ranking 6-7
destination ranking 6-7
diagnostics
  beep codes 10-26
diagrams
  attack 3-10
display format
  query 6-6

E

editing
  inspection rules 7-17
  IP groups 9-4
  service 9-6
  user 9-10
event groups 9-3
event management 9-1
  editing 9-2
Event Type 5-3
event type group ranking 6-6
event type ranking 6-6
Expand All 5-5
expired certificate 10-6

F

false positives
  tuning 5-5
fingerprint validation 10-4

G

Global Controller xi
  adding Local Controllers to 2-2
  and Local Controllers 2-10, 3-1, 5-1, 6-1, 7-1, 7-3, 9-7
  Network Summary page 3-1
  queries 6-1
  rules 7-1, 7-3
  user interface xi
  user management 9-7
Global Controller
  overview 1-1

H

hard drive
  failure alert 10-16
  hotswap procedure for MARS 110R, 110, 210, GC2R, and GC2 10-21
  raidstatus command 10-15
  replacing in carrier 10-24
  slot number diagram, MARS 110R, 110, 210, GC2R, and GC2 10-20
hardware maintenance
  MARS 110, 110R, 210, GC2R, GC2 10-14
Hot Spot Graph 3-10
hotswap
  hard drives 10-15
  power supply 10-25
  procedure for MARS 110R, 110, 210, GC2R, and GC2 10-21

I

incident count 6-8
Incident Details page 5-4
Incident ID 5-3
Incident Path 5-3
incidents 3-9
  action 5-3
  event type 5-3
  incident ID 5-3
  incident path 5-3

- incident vector 5-3
- instances 5-6
- matched rule 5-3
- severity 5-3
- time 5-3
- time ranges 5-4

incidents table
- navigation 5-3

incident table 5-5
Incident Vector 5-3
inspection rule
- activate and inactive 7-16

inspection rules
- adding 7-18
- editing 7-17

inspection rule status
- changing 7-16

instances
- incidents 5-6

IP groups
- adding 9-4
- editing 9-4

IP management 9-3
- adding
  - IP range 9-4
  - network 9-4
  - variable 9-4

L

L2 attack path 5-5
L3 attack path 5-5
Local Controller 2-10, 3-1, 5-1, 6-1, 7-1, 7-3, 9-7
log files 10-2

M

MAC address report 6-8
management
- events 9-1
- IP 9-3
- service 9-5
- user 9-6

MARS
- audit trail 10-3
- log files 10-2

matched incident ranking 6-7
Matched Rule 5-3
matched rule ranking 6-7
mitigate 5-5

N

NAT connection report 6-8
network group ranking 6-6
network ranking 6-6
Network Status tab
- Incidents 3-13
- Top Destinations 3-14
- Top Event Types 3-13
- Top Sources 3-14

O

Order/Rank By 6-8
order by 6-8
- bytes transmitted 6-8
- incident count 6-8
- session count 6-8
- time 6-8

P

pager 8-11, 9-9
password
- change default 10-3

post NAT destination addresses 6-11
post NAT source addresses 6-11
pre NAT destination addresses 6-11
pre NAT source addresses 6-11
protocol ranking 6-7

Q

queries
action
ANY 6-13
actions 6-13
destination IP 6-11
ANY 6-11
devices 6-12
IP addresses 6-11
IP ranges 6-11
networks 6-11
post NAT destination addresses 6-11
pre NAT destination addresses 6-11
devices 6-12
display format
all matching event raw messages 6-8
all matching events 6-8
all matching sessions 6-7
destination IP address ranking 6-7
destination ranking 6-7
event type group ranking 6-6
MAC address report 6-8
matched incident ranking 6-7
matched rule ranking 6-7
NAT connection report 6-8
protocol ranking 6-7
reporting device ranking 6-7
reporting device type ranking 6-7
source IP address ranking 6-6
source port ranking 6-7
unknown event report 6-8
use only firing events 6-9
event type grouping 6-12
event types 6-12
ANY 6-12
operation
AND 6-13, 7-12
FOLLOWED-BY 6-13, 7-12
none 6-13, 7-12
OR 6-13, 7-12
result format
destination network group ranking 6-7
destination network ranking 6-7
event type ranking 6-6
network group ranking 6-6
network ranking 6-6
reported user ranking 6-7
source network group ranking 6-6
source network ranking 6-6
rule 6-13
ANY 6-13
save as
reports 6-13
rules 6-13
service
ANY 6-12
defined services 6-12
service variables 6-12
severity
ANY 6-12
green 6-12
red 6-12
yellow 6-12
source IP
ANY 6-11
devices 6-11
IP addresses 6-11
IP ranges 6-11
networks 6-11
post NAT source addresses 6-11
pre NAT source addresses 6-11

variables 6-11
time range
last 6-8
start and end times 6-8
zone 6-12
query
display format 6-6
Query page 6-1

R

rank by 6-8
bytes transmitted 6-8
incident count 6-8
session count 6-8
time 6-8
removing
user 9-10
report
adding 6-22
delete 6-23
edit 6-23
new 6-22
reported user ranking 6-7
reporting device ranking 6-7
reporting device type ranking 6-7
reports
viewing 6-16, 6-23
reports, view type, CSV 6-21
reports, view type, recent 6-21
reports,view type, total 6-21
report views, CSV 6-21
report views, peak, reports, view type, peak 6-21
report views, recent 6-21
report views, total 6-21
rules
destination IP
ANY 7-7
devices 7-7
DISTINCT 7-7
IP addresses 7-7
IP ranges 7-7
Network Groups 7-7
networks 7-7
SAME 7-7
variables 7-7
device 7-10
ANY 7-10
Unknown Reporting Device 7-10
variables 7-10
event type grouping 7-9
event types 7-9
ANY 7-9
variables 7-9
reported user
ANY 7-10
Invalid User Name 7-10
NONE 7-10
variables 7-10
service
ANY 7-8
defined groups 7-9
defined services 7-9
service variables 7-8
severity
ANY 7-11
green 7-11
red 7-11
yellow 7-11
source IP
devices 7-6
IP addresses 7-6
IP ranges 7-6
Network Groups 7-6
networks 7-6
variables 7-6
runtime logging 10-1

S

see CVE 9-2
service
  adding 9-6
  deleting 9-6
  editing 9-6
  editing groups 9-5
service group
  adding 9-5
service management 9-5
service provider
  adding 8-11, 9-9
services
  adding group 9-5
session count 6-8
setting
  runtime logging levels 10-1
Severity icons 5-3
Short Message Service
  See SMS. 7-14
Simple Network Management Protocol
  See SNMP. 7-14
source IP address ranking 6-6
source network group ranking 6-6
source network ranking 6-6
source port ranking 6-7
SSH
  fingerprint validation 10-4
SSL
  certificate validation 10-4
stacked charts 3-14

T

table
  incidents 5-5
Time 5-3
time ranges
  incidents 5-4
Topology
  toggle device display 3-13
tuning
  false positives 5-5

U

unknown event report 6-8
use only firing events 6-9
user
  adding 8-10, 9-7
  editing 9-10
  removing 9-10
user group
  adding 9-10
user management 9-6
  roles defined 9-7

V

validation
  fingerprint 10-4
variables 6-11, 7-6, 7-7