

Administration Guide

NetIQ[®] Sentinel[™]

April 2014



Legal Notice

NetIQ Sentinel is protected by United States Patent No(s): 05829001.

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, NETIQ CORPORATION PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

For purposes of clarity, any module, adapter or other similar material ("Module") is licensed under the terms and conditions of the End User License Agreement for the applicable version of the NetIQ product or software to which it relates or interoperates with, and by accessing, copying or using a Module you agree to be bound by such terms. If you do not agree to the terms of the End User License Agreement you are not authorized to use, access or copy a Module and you must destroy all copies of the Module and contact NetIQ for further instructions.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of NetIQ Corporation, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of NetIQ Corporation. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. NetIQ Corporation may make improvements in or changes to the software described in this document at any time.

U.S. Government Restricted Rights: If the software and documentation are being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), in accordance with 48 C.F.R. 227.7202-4 (for Department of Defense (DOD) acquisitions) and 48 C.F.R. 2.101 and 12.212 (for non-DOD acquisitions), the government's rights in the software and documentation, including its rights to use, modify, reproduce, release, perform, display or disclose the software or documentation, will be subject in all respects to the commercial license rights and restrictions provided in the license agreement.

© 2014 NetIQ Corporation and its affiliates. All Rights Reserved.

For information about NetIQ trademarks, see <http://www.netiq.com/company/legal/>.

Contents

About this Book and the Library	11
About NetIQ Corporation	13
Part I Getting Started	15
1 Understanding Sentinel Applications	17
1.1 Accessing the Sentinel Web Console	17
1.2 Accessing the Sentinel Control Center	18
2 Adding a License Key	19
2.1 Adding a License Key By Using the Web Interface	19
2.2 Adding a License Key through the Command Line	19
3 Security Considerations	21
3.1 Basic Security Considerations	21
3.1.1 Traditional Installation	21
3.1.2 Appliance Installation	22
3.2 Securing Sentinel Data	22
3.3 Best Practices	22
3.3.1 Changing Passwords	22
3.3.2 Changing Collector Manager or Correlation Engine User Credentials	23
3.3.3 Enforcing Password Policies for Users	24
3.3.4 Securing Communication with Collector Managers and Event Sources	24
3.3.5 Securing Communication with Secondary Storage	24
3.3.6 Auditing Sentinel	24
3.3.7 Determining if Data was Tampered	24
3.3.8 Using CA Signed Certificates	28
3.4 Network Communication Options	30
3.4.1 Communication between Sentinel, Collector Manager, and Correlation Engine	31
3.4.2 Communication between Sentinel and NetFlow Collector Manager	32
3.4.3 Communication between Sentinel and the Sentinel Control Center and Solution Designer Client Applications	32
3.4.4 Communication between the Server and the Database	32
3.4.5 Communication with Web Browsers	32
3.4.6 Communication between the Database and Other Clients	33
3.5 Sensitive Data Locations	33
3.6 Implementing Intruder Detection and Lockout Mechanisms	34
3.7 Applying Updates for Security Vulnerabilities in Embedded Third-Party Products	35
Part II Configuring Roles and Users	37
4 Configuring Roles and Users	39
4.1 Overview	39
4.2 Creating Roles	40
4.2.1 Creating a Role	41

4.3	Configuring Password Complexity	42
4.4	Creating Users	43
5	Configuring LDAP Authentication	45
5.1	Overview	45
5.2	Prerequisites	46
5.2.1	Exporting the LDAP Server CA Certificate	46
5.2.2	Enabling Anonymous Search in the LDAP Directory	46
5.3	Setting Up LDAP Authentication	46
5.4	Logging in by Using LDAP User Credentials	49
5.5	Configuring Multiple LDAP Servers for Failover	49
	Part III Collecting and Routing Event Data	51
6	Configuring Data Storage	53
6.1	Types of Data	53
6.1.1	Raw Data	53
6.1.2	Event Data	58
6.2	Configuring Secondary Storage Locations	60
6.2.1	Supported Storage Options	60
6.2.2	Types of Secondary Storage	61
6.2.3	Configuring Secondary Storage	61
6.2.4	Changing the Secondary Storage Location	63
6.3	Configuring Disk Space Usage	64
6.4	Verifying and Downloading Raw Data Files	64
6.5	Configuring Data Retention Policies	65
6.5.1	Raw Data Retention Policy	66
6.5.2	Event Data Retention Policies	66
6.5.3	Rules for Applying a Retention Policy	67
6.6	Configuring Data Synchronization	67
6.6.1	Overview	67
6.6.2	Creating a Data Synchronization Policy	70
6.6.3	Managing Data Synchronization	72
6.7	Viewing Primary and Secondary Storage Capacity	72
6.8	Using Sequential-Access Storage for Long Term Data Storage	73
6.8.1	Determining What Data You Need to Copy to Tape	73
6.8.2	Backing Up Data	73
6.8.3	Configuring Storage Utilization	74
6.8.4	Configuring Data Retention	74
6.8.5	Copying Data to Tape	74
6.8.6	Restoring Data	75
6.9	Data Deletion Policy	76
7	Configuring Agentless Data Collection	79
7.1	Before You Begin	79
7.2	Configuring Data Collection for Syslog Event Sources	80
7.2.1	Parsing Logic for Syslog Messages	80
7.2.2	Configuring Syslog Servers	80
7.2.3	Configuring Client Authentication for the SSL Syslog Server	81
7.3	Configuring Data Collection for the Novell Audit Server	83
7.3.1	Specifying the Audit Server Settings	83
7.3.2	Setting the Audit Server Options	84
7.4	Configuring Data Collection for Other Event Sources	87

7.4.1	Accessing Event Source Management	87
7.4.2	Viewing Data in Event Source Management	88
7.4.3	Searching for Event Sources	93
7.4.4	Installing Plug-Ins	94
7.4.5	Updating a Connector or a Collector Plug-In	95
7.4.6	Adding Components to Sentinel	95
7.4.7	Connecting to Event Sources	97
7.4.8	Exporting Configurations	100
7.4.9	Importing Configurations	101
7.4.10	Debugging	101
7.4.11	Troubleshooting	105
7.5	Managing Event Sources	105
7.5.1	Viewing the Event Sources Page	105
7.5.2	Changing the Data Logging Status of Event Sources	108
8	Configuring Agent-Based Data Collection	109
9	Managing Event Sources	111
9.1	Viewing the Event Sources Page	111
9.1.1	Viewing Event Sources	112
9.1.2	Configuring Event Sources	113
9.1.3	Viewing Collector Managers	113
9.1.4	Viewing Event Source Servers	114
9.1.5	Viewing Collector Plug-Ins	114
9.2	Filtering Event Sources	115
9.2.1	Filtering Event Sources by Name	115
9.2.2	Filtering Event Sources by Health Status	116
9.2.3	Filtering by Event Sources Event Search Results	116
9.2.4	Filtering Event Sources by Collector Managers	116
9.2.5	Filtering Event Sources by Event Source Servers	116
9.2.6	Filtering Event Sources by Collector Plug-Ins	117
9.2.7	Changing the Data Logging Status of Event Sources	117
9.2.8	Changing the Associated Collector Plug-In for Event Sources	117
10	Configuring Event Routing Rules	119
10.1	Creating an Event Routing Rule	119
10.2	Ordering Event Routing Rules	120
10.3	Activating or Deactivating an Event Routing Rule	121
11	Mapping Events	123
11.1	Overview	123
11.1.1	Maps	123
11.1.2	Mapping Events	124
11.2	Default Maps	126
11.3	Accessing Map Definitions	127
11.4	Adding Map Definitions	127
11.5	Adding a Number Range Map Definition	128
11.6	Updating Map Data	131
11.7	Using Maps for Event Configuration	132
11.8	Renaming Event Fields	133
12	Linking to Additional Sentinel Systems	135
12.1	Benefits	135

12.2	Supported Platforms	135
12.3	Prerequisite	135
12.4	Configuring Sentinel Link	136
Part IV Integrating with External Systems		137
13 Configuring Actions		139
13.1	Overview	139
13.2	Understanding the Action Manager Interface	140
13.3	Managing Actions	141
13.3.1	Adding an Action	142
13.3.2	Debugging Actions	143
13.4	Managing Action Plug-Ins	144
13.4.1	Understanding the Action Plug-In Manager Interface	144
13.4.2	Importing an Action Plug-In	145
14 Configuring Integrators		147
14.1	Overview	147
14.2	Managing Integrators	148
14.2.1	Configuring the Default Integrators	148
14.2.2	Adding an Integrator	151
14.2.3	Viewing Integrator Health Details	151
14.3	Managing Integrator Plug-Ins	153
14.3.1	Importing an Integrator Plug-In	153
15 Integrating Identity Information		155
15.1	Overview	155
15.2	Integration with Identity Management Systems	156
15.3	Leveraging Identity Information	158
16 Managing Data Feeds		159
16.1	Using Data Feeds for Botnet Detection	159
17 Detecting Vulnerabilities and Exploits		161
17.1	Understanding Advisor	161
17.2	Understanding Exploit Detection	162
17.2.1	How Exploit Detection Works	162
17.2.2	Generating the Exploit Detection File	164
17.2.3	Viewing the Events	164
17.3	Introduction to the Advisor Interface	164
17.4	Processing the Advisor Feed	166
17.4.1	Processing the Feed Files Manually	166
17.4.2	Processing the Feed Files Automatically	167
17.5	Configuring the Advisor Products for Exploit Detection	167
17.6	Downloading the Advisor Feed	167
17.6.1	Configuring the Sentinel Server for Automated Downloads	168
17.6.2	Creating a Download Configuration	168
17.6.3	Downloading the Feed Instantly	169
17.6.4	Audit Events for the Download Manager	170
17.6.5	Downloading the Advisor Feed Manually	170
17.7	Viewing the Advisor Status	170

17.8	Viewing the Advisor Data	172
17.8.1	Active Views	172
17.8.2	Search results in the Web Interface.	172
17.9	Resetting the Advisor Password	172
17.10	Deleting the Advisor Data	173
Part V Monitoring Your Network		175
18	Searching and Reporting Events in a Distributed Environment	177
18.1	Overview	177
18.2	Configuring Servers for Distributed Searching and Reporting	179
18.2.1	Enabling Distributed Search	179
18.2.2	Using the Administrator Credentials to Add a Search Target Server	179
18.2.3	Using the Opt-in Password to Add a Search Target Server	180
18.3	Searching for Events	182
18.4	Managing the Distributed Search Results	182
18.5	Viewing the Search Activities.	183
18.6	Running Reports	184
18.7	Managing the Distributed Setup Configuration	184
18.7.1	Editing the Search Target Server Details	184
18.8	Troubleshooting	185
18.8.1	Permission Denied.	185
18.8.2	Connection Down	185
18.8.3	Unable to View Raw Data	186
18.8.4	Problems Adding Search Target	186
18.8.5	Certain Events Are Only Visible from the Local System	186
18.8.6	Cannot Run Reports on the Target Servers	186
18.8.7	Different Users Get Different Results	186
18.8.8	Cannot Set the Admin Role as the Search Proxy Role	186
18.8.9	Error Logs	186
19	Visualizing Network Traffic	189
19.1	Understanding Network Flow Data Collection	189
19.1.1	NetFlow Collector Manager	189
19.1.2	Supported NetFlow Fields	190
19.2	Collecting Network Flow Data for Specific Tenants	191
19.3	Visualizing and Analyzing Network Flow Data	191
19.4	Viewing and Managing the Network Flow Data Collection.	192
19.5	Customizing the NetFlow Configuration.	192
19.5.1	Tuning the Network Flow Data Collection	192
19.5.2	Changing the User Name and Password in the NetFlow Collector Manager	193
Part VI Managing Solution Packs		195
20	Using Solution Packs	197
20.1	Overview	197
20.2	Solution Pack Components	198
20.3	Using the Import Plugin Wizard to Import a Solution Pack.	199
20.4	Using the Solution Manager.	200
20.4.1	Launching the Solution Manager.	200
20.4.2	Solution Manager Interface	200
20.5	Installing and Managing Solution Packs	202

20.5.1	Viewing the Contents of a Solution Pack	203
20.5.2	Installing Content from Solution Packs	203
20.5.3	Configuring Controls	204
20.5.4	Implementing a Control	206
20.5.5	Testing a Control	206
20.5.6	Uninstalling a Control	207
20.5.7	Viewing Solution Pack Status	207
20.5.8	Deleting a Solution Pack	209
20.6	Installing an Edited Solution Pack	209
20.7	Solution Designer	210
21	Creating Solution Packs	211
21.1	Accessing the Solution Designer	211
21.2	Creating a Solution Pack	211
21.3	Adding Content to a Solution Pack	212
21.3.1	Sentinel Content	212
21.3.2	Using Placeholders	214
21.3.3	File Attachments	214
21.4	Initializing Dynamic Lists Through Solution Pack	214
21.5	Documenting a Solution Pack	215
21.5.1	Description	215
21.5.2	Implementation Steps	215
21.5.3	Testing Steps	215
21.6	Synchronizing Content	215
21.7	Handling Inter-control Dependency	216
21.8	Managing a Solution Pack	216
21.8.1	Adding a Node to a Control	216
21.8.2	Moving Nodes	217
Part VII	Managing Your Sentinel Environment	219
22	Managing Active Searches and Reports	221
23	Viewing Events Per Second Statistics	223
23.1	Viewing a Graphical Representation of the Events Per Second Value	223
23.2	Viewing the Events Per Second Value of Event Source Servers	224
24	Configuring Sentinel for High Availability	225
25	Backing Up and Restoring Data	227
25.1	Parameters for the Backup and Restore Utility Script	227
25.2	Running the Backup and Restore Utility Script	229
26	Customizing Sentinel Settings	233
26.1	Optimizing the Operating System	233
26.2	Configuring the Resources for Event Partition Compression	234
26.3	Setting the Grace Period to Close Event Data Partitions	234
26.4	Compressing the Storage Index on Primary Partition	234
26.5	Configuring Memory for the Sentinel Server	235
26.6	Setting the Raw Data Limit	235

A	Command Line Utilities	239
A.1	Managing the Sentinel Services	239
A.2	Sentinel Scripts	240
A.3	Running the Report Development Utility	241
A.4	Getting the .jar Version Information	242
A.5	Changing the Hostname of a Sentinel Server	242
A.6	Importing or Exporting Event Association Data	242
A.7	Managing the Internal Database	244
A.7.1	Commands	244
A.7.2	Options	244
A.8	Cleaning Up the Internal Database	245
A.8.1	Prerequisites	245
A.8.2	Using the clean_db.sh Script	245
A.9	Managing the Sentinel Server	246
A.9.1	Commands	246
A.9.2	Options	246

About this Book and the Library

The *Administration Guide* provides the administration information and tasks required to manage a Sentinel deployment.

Intended Audience

This guide is intended for Sentinel administrators and consultants.

Other Information in the Library

The library provides the following information resources:

Installation and Configuration Guide

The Installation and Configuration Guide provides an introduction to NetIQ Sentinel and explains how to install and configure Sentinel.

User Guide

Provides conceptual information about Sentinel. This book also provides an overview of the user interfaces and step-by-step guidance for many tasks.

About NetIQ Corporation

We are a global, enterprise software company, with a focus on the three persistent challenges in your environment: Change, complexity and risk—and how we can help you control them.

Our Viewpoint

Adapting to change and managing complexity and risk are nothing new

In fact, of all the challenges you face, these are perhaps the most prominent variables that deny you the control you need to securely measure, monitor, and manage your physical, virtual, and cloud computing environments.

Enabling critical business services, better and faster

We believe that providing as much control as possible to IT organizations is the only way to enable timelier and cost effective delivery of services. Persistent pressures like change and complexity will only continue to increase as organizations continue to change and the technologies needed to manage them become inherently more complex.

Our Philosophy

Selling intelligent solutions, not just software

In order to provide reliable control, we first make sure we understand the real-world scenarios in which IT organizations like yours operate — day in and day out. That's the only way we can develop practical, intelligent IT solutions that successfully yield proven, measurable results. And that's so much more rewarding than simply selling software.

Driving your success is our passion

We place your success at the heart of how we do business. From product inception to deployment, we understand that you need IT solutions that work well and integrate seamlessly with your existing investments; you need ongoing support and training post-deployment; and you need someone that is truly easy to work with — for a change. Ultimately, when you succeed, we all succeed.

Our Solutions

- ♦ Identity & Access Governance
- ♦ Access Management
- ♦ Security Management
- ♦ Systems & Application Management
- ♦ Workload Management
- ♦ Service Management

Contacting Sales Support

For questions about products, pricing, and capabilities, contact your local partner. If you cannot contact your partner, contact our Sales Support team.

Worldwide:	www.netiq.com/about_netiq/officelocations.asp
United States and Canada:	1-888-323-6768
Email:	info@netiq.com
Web Site:	www.netiq.com

Contacting Technical Support

For specific product issues, contact our Technical Support team.

Worldwide:	www.netiq.com/support/contactinfo.asp
North and South America:	1-713-418-5555
Europe, Middle East, and Africa:	+353 (0) 91-782 677
Email:	support@netiq.com
Web Site:	www.netiq.com/support

Contacting Documentation Support

Our goal is to provide documentation that meets your needs. If you have suggestions for improvements, click **Add Comment** at the bottom of any page in the HTML versions of the documentation posted at www.netiq.com/documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Contacting the Online User Community

Qmunity, the NetIQ online community, is a collaborative network connecting you to your peers and NetIQ experts. By providing more immediate information, useful links to helpful resources, and access to NetIQ experts, Qmunity helps ensure you are mastering the knowledge you need to realize the full potential of IT investments upon which you rely. For more information, visit <http://community.netiq.com>.

Getting Started

Sentinel is a Security Information and Event Management (SIEM) system that receives information from many sources throughout an enterprise, standardizes it, prioritizes it and presents it to you to make threat, risk and policy related decisions. For detailed information about Sentinel and its components, see [“Understanding Sentinel”](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

This section provides information about the following:

- ♦ [Chapter 1, “Understanding Sentinel Applications,”](#) on page 17
- ♦ [Chapter 2, “Adding a License Key,”](#) on page 19
- ♦ [Chapter 3, “Security Considerations,”](#) on page 21

1 Understanding Sentinel Applications

Sentinel presents events in both Sentinel Web interface and also in the Sentinel Control Center.

Sentinel Web interface helps you to view existing events and interact with those events. Some of the actions that you can perform include the following:

- ♦ Configure data collection for event sources such as syslog, audit, and so on.
- ♦ Analyze trends in data.
- ♦ Correlate event data.
- ♦ Visualize and analyze network flow data.
- ♦ Generate reports.

Sentinel Control Center (SCC) is used primarily for administration and configuration, but it also helps you to view events in near real time. Some of the actions that you can perform include:

- ♦ Configure data collection for event sources such as Windows, database, Check Point server, and so on.
- ♦ View events in real time.
- ♦ Add contextual information to events.
- ♦ Configure Actions and Integrators.
- ♦ Manage Solution Packs

1.1 Accessing the Sentinel Web Console

- 1 Launch a supported Web browser.
- 2 Specify the URL of the Sentinel Web interface:

`https://<IP_Address/DNS_Sentinel_server>:8443`

IP_Address/DNS_Sentinel_server is the IP address or DNS name of the Sentinel server and 8443 is the default port for the Sentinel server.

- 3 Log in as a user with permissions to access the desired feature.

or

If it is your first time logging in, log in with the administrator name and password specified during the installation.

By default, the user name is admin. The password is specified during the installation of Sentinel. For more information, see “[Performing Interactive Installation](#)” in the *[NetIQ Sentinel Installation and Configuration Guide](#)*.

1.2 Accessing the Sentinel Control Center

The Sentinel Control Center is accessed through the Sentinel Web console.

- 1 Log in to the Sentinel Web console:

`https://<IP_Address/DNS_Sentinel_server:8443>`

IP_Address/DNS_Sentinel_server is the IP address or DNS name of the Sentinel server and 8443 is the default port for the Sentinel server.

- 2 In the toolbar, click **Applications**.
- 3 Click **Launch Control Center**.

NOTE: NetIQ Corporation leverages our partner company Novell for certificate signing services under a trust model. Although the certificate says Novell, it is trusted through NetIQ Corporation.

- 4 Click **Yes** to accept the security certificate.
- 5 Specify a user name and password of a user that has rights to access the SCC, then click **Login**.
- 6 Click **Accept** or **Accept Permanently** to accept the security certificate.

The Sentinel Control Center launches in a new window.

2 Adding a License Key

You can add a license key when installing Sentinel. This section provides information about adding the license key after the Sentinel installation.

If you are using the temporary license key, you must add the enterprise license key before the temporary key expires to avoid any interruption in the Sentinel functionality. For information about how to purchase the license, see the [Sentinel Product Web site](#).

You can add a license key either by using the Web interface or through the command line.

- [Section 2.1, “Adding a License Key By Using the Web Interface,” on page 19](#)
- [Section 2.2, “Adding a License Key through the Command Line,” on page 19](#)

2.1 Adding a License Key By Using the Web Interface

- 1 Log in to the Sentinel Web interface as an administrator.
- 2 Click **About > Licenses**.
- 3 In the Licenses section, click **Add License**.
- 4 Specify the license key in the **Key** field. After you specify the license, the following information is displayed in the Preview section:
 - Features:** The features that are available with the license.
 - Hostname:** This field is for internal Novell use only.
 - Serial:** This field is for internal Novell use only.
 - EPS:** Event rate built into the license key. Beyond this rate, Sentinel generates warnings but will continue to collect data.
 - Expires:** Expiry date of the license. You must specify a valid license key before the expiry date to prevent an interruption in functionality.
- 5 Click **Save**.

2.2 Adding a License Key through the Command Line

If you are using the Sentinel traditional installation, you can add the license through the command line by using the `softwarekey.sh` script.

- 1 Log in to the Sentinel server as root.
- 2 Change to the `/opt/novell/sentinel/bin` directory.
- 3 Enter the following command to change to the novell user:

```
su novell
```
- 4 Specify the following command to run the `softwarekey.sh` script.

```
./softwarekey.sh
```

- 5** Enter 1 to insert the license key.
- 6** Specify the license key, then press Enter.

3 Security Considerations

This section provides information on how to securely maintain your Sentinel server.

- ♦ [Section 3.1, “Basic Security Considerations,” on page 21](#)
- ♦ [Section 3.2, “Securing Sentinel Data,” on page 22](#)
- ♦ [Section 3.3, “Best Practices,” on page 22](#)
- ♦ [Section 3.4, “Network Communication Options,” on page 30](#)
- ♦ [Section 3.5, “Sensitive Data Locations,” on page 33](#)
- ♦ [Section 3.6, “Implementing Intruder Detection and Lockout Mechanisms,” on page 34](#)
- ♦ [Section 3.7, “Applying Updates for Security Vulnerabilities in Embedded Third-Party Products,” on page 35](#)

3.1 Basic Security Considerations

Sentinel has undergone security hardening before being released. This section describes some of the hardening mechanisms used in Sentinel.

- ♦ [Section 3.1.1, “Traditional Installation,” on page 21](#)
- ♦ [Section 3.1.2, “Appliance Installation,” on page 22](#)

3.1.1 Traditional Installation

- ♦ All unnecessary ports are turned off.
- ♦ Whenever possible, a service port listens only for local connections and does not allow remote connections.
- ♦ Files are installed with least privileges so that the least number of users can read the files.
- ♦ Default passwords are not used.
- ♦ Reports against the database are run as a user that only has SELECT permissions on the database.
- ♦ All Web interfaces require HTTPS.
- ♦ All communication over the network uses SSL by default and is configured to require authentication.
- ♦ User account passwords are encrypted by default when they are stored on the file system or in the database.

3.1.2 Appliance Installation

In addition to the points mentioned in [“Traditional Installation” on page 21](#), the appliance has undergone the following additional hardening:

- ♦ Only the minimally required packages are installed.
- ♦ The firewall is enabled by default and all unnecessary ports are closed in the firewall configuration.
- ♦ Sentinel is automatically configured to monitor the local operating systems syslog messages for audit purposes.

3.2 Securing Sentinel Data

Because of the highly sensitive nature of data in Sentinel, you must keep the computer physically secure and in a secure area of the network. To collect data from event sources outside the secure network, use a remote Collector Manager. For more information about remote Collector Managers, see [“Installing Additional Collector Managers and Correlation Engines”](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

Sentinel is compatible with disk encryption technologies. These technologies provide a higher level of data privacy when they are used on file systems where Sentinel stores its data. However, software-based encryption technologies, such as dm-crypt, have a significant CPU overhead, and they can dramatically reduce the performance of Sentinel by 50% or more. However, hardware-based encryption technologies have a much lower impact on the performance of the rest of the system and are available from leading hard drive manufacturers.

3.3 Best Practices

Use the following best practices to secure your Sentinel server:

- ♦ [Section 3.3.1, “Changing Passwords,” on page 22](#)
- ♦ [Section 3.3.2, “Changing Collector Manager or Correlation Engine User Credentials,” on page 23](#)
- ♦ [Section 3.3.3, “Enforcing Password Policies for Users,” on page 24](#)
- ♦ [Section 3.3.4, “Securing Communication with Collector Managers and Event Sources,” on page 24](#)
- ♦ [Section 3.3.5, “Securing Communication with Secondary Storage,” on page 24](#)
- ♦ [Section 3.3.6, “Auditing Sentinel,” on page 24](#)
- ♦ [Section 3.3.7, “Determining if Data was Tampered,” on page 24](#)
- ♦ [Section 3.3.8, “Using CA Signed Certificates,” on page 28](#)

3.3.1 Changing Passwords

To increase security, you can change the passwords of the system users created during the installation of Sentinel. There are two different types of users:

- ♦ [“Administration Users” on page 23](#)
- ♦ [“Operating System Users” on page 23](#)
- ♦ [“Application and Database Users” on page 23](#)

Administration Users

The `admin` user is the administrator user for Sentinel applications. You can use the administrator credentials to log in to the Sentinel Web interface. The password is set during the installation process.

Operating System Users

The Sentinel server installation creates a `novell` system user and a `novell` group that owns the installed files within the `install_directory`. The user's home directory is set to `/home/novell`. The `novell` user does not have a password and cannot log in to the operating system unless you assign a password after installation.

Application and Database Users

Sentinel application users are native database users and their passwords are protected by the native database platform, unless you have configured LDAP authentication. These users have only read access to certain tables in the database so that they can execute queries against the database. Users authenticated by LDAP do not have read access on the database.

dbauser: The `dbauser` is created as a superuser who can manage the database and is typically the user who can log in to pgAdmin for troubleshooting purposes. The password for the `dbauser` is the same as the `admin` user password specified during installation. The password must meet PostgreSQL database password standards.

appuser: The `appuser` is used to connect to the database for regular operations that do not require a superuser. The password for the `appuser` is the same as the password for the `admin` user specified during installation.

To modify the password for `admin`, `dbauser`, or `appuser`, use the `configure.sh` script. For more information, see [Appendix A, "Command Line Utilities," on page 239](#). When you change the password by using the script, Sentinel updates the password in all relevant places without any manual intervention. However, if you try to change the password by any other method, the password does not get updated in all of the relevant files and some parts of Sentinel might stop working.

NOTE: There is also a PostgreSQL database user that owns the entire database, including system database tables. By default, the PostgreSQL database user is set to `NOLOGIN`, so that no one can log in as the PostgreSQL user.

3.3.2 Changing Collector Manager or Correlation Engine User Credentials

The ActiveMQ user credentials for the Collector Manager and the Correlation Engine are stored in the `/<install_dir>/etc/opt/novell/sentinel/config/activemqusers.properties` file located in the Sentinel server. You can add additional ActiveMQ users or change the password for the existing users by modifying the `activemqusers.properties` file. For more information, see ["Adding a Custom ActiveMQ User for the Collector Manager or Correlation Engine"](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

If you modify the password for an existing ActiveMQ user or create a new user by modifying the `activemqusers.properties` file in the Sentinel server, you must copy the same changes to all of the remote Collector Managers and Correlation Engines and then restart the services. Otherwise, the Collector Manager and the Correlation Engine would not be able to connect to the Sentinel server.

3.3.3 Enforcing Password Policies for Users

To achieve robust password policy enforcement in Sentinel, use Sentinel's built-in password complexity policy enforcement capability. For more information about configuring password complexity, see [Section 4.3, “Configuring Password Complexity,” on page 42](#).

You can also use an LDAP directory to authenticate Web application users. To enable this option by using the Sentinel Web interface, see [Chapter 5, “Configuring LDAP Authentication,” on page 45](#). This option has no effect on accounts used by back-end services, which continue to authenticate through PostgreSQL.

3.3.4 Securing Communication with Collector Managers and Event Sources

You can configure Sentinel to securely collect data from various event sources. However, secured data collection is determined by the specific protocols supported by the event source. For example, the Check Point LEA, Syslog, and Audit Connectors can be configured to encrypt their communication with event sources.

For more information about the possible security features that you can enable, see the [Sentinel Plugins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

3.3.5 Securing Communication with Secondary Storage

You must consider the security implications before deciding the type of secondary storage location to use. If you are using CIFS or NFS servers as secondary storage locations to store the Sentinel event data and raw data, remember that these protocols do not offer data encryption. An alternative is to use direct attached storage (Primary, formerly known as local or SAN), which does not have the same security vulnerabilities. If you choose to use CIFS or NFS, it is important to configure the CIFS or NFS server to maximize the security of your data.

For more information about configuring the secondary storage location server settings, see [Section 6.2, “Configuring Secondary Storage Locations,” on page 60](#).

3.3.6 Auditing Sentinel

Sentinel generates audit events for many actions performed manually and also for actions performed internally for system activities. Sentinel tags these events with the `Sentinel` tag. To include these events in a report, perform a search by using the `rv145:Sentinel` query and select **include system events**. However, you must have the necessary permissions to view system events. For more information, see [Chapter 4, “Configuring Roles and Users,” on page 39](#).

Sentinel provides reports that are preconfigured to include only the events tagged with the `Sentinel` tag.

A well-audited Sentinel system not only audits events occurring within Sentinel, but also the infrastructure on which Sentinel is running. You can set up data collection from the computers and the devices that make up the Sentinel infrastructure and tag them with the `Sentinel` tag to enable a complete auditing of the systems that can affect the behavior of Sentinel. For appliance installations, Sentinel is automatically configured to monitor the local operating system's syslog messages for audit purposes.

3.3.7 Determining if Data was Tampered

There are two approaches to verify that the event was not tampered while in storage.

Event Data Approach

The event data approach involves proving that a particular event of interest has not been tampered with. At a high level, this involves verifying that the partition that the event is stored in has not been tampered. Since Sentinel computes integrity hashes at the partition level and not the per-event level, the integrity check must be done at the partition level.

You can verify the integrity of event data by checking if the data in the secondary storage location has been tampered. Immediately after an event data partition is copied from primary storage to secondary storage, a hash is computed on the copy of the partition in the secondary storage. You can verify the integrity of event data using the hash.

The hash is computed as follows:

1. In the event partition, the data in the following files are concatenated in the following order:
 - a. `index.sqfs`
 - b. All the files in the `events.evt` directory, in alphabetical order.
2. The concatenation of the files is hashed using the SHA-256 algorithm.
3. The hash is base64 encoded. The base64 encoded hash value is stored in the row associated with the event partition under the HASH column of the IXLOG_PART table. The event partition directory name is stored under the NAME column of the IXLOG_PART table.

After the event partition is copied to the secondary storage, the hash value populates in the HASH column of the IXLOG_PART table. You can determine if the integrity of the event partition has been compromised by recomputing the hash of the event partition and comparing it with the HASH value in the IXLOG_PART table.

NOTE: This approach depends on the concept that the hash is stored separately and securely from the event data. The hash is stored in the authenticated Sentinel database whereas the event data is stored on the file system (not within the database). You can further protect the hashes by taking regular backups of the Sentinel database and storing the backups in an even more secure location. You can retrieve the hashes later to check the event data integrity.

To verify whether the event data was tampered:

- 1 Determine the partition the event is stored in:
 - 1a Export the event of interest to a CSV file and including the RetentionPolicyID (rv171) field in the export.
 - 1b Note down the value in the RetentionPolicyID (rv171) field. It is a unique ID of the retention policy under which the event is stored.
 - 1c Find the exact partition by executing database queries by running the following command as the novell user on the Sentinel server:

```
db.sh sql SIEM dbauser "select name, part_date, hash, state, part_id,
ret_pol_id from IXLOG_PART where ret_pol_id='<retention_policy_ID>'"
```

where `<retention_policy_ID>` is the value in the RetentionPolicyID (rv171) field determined in [Step 1b](#).
 - 1d Determine the exact partition by comparing the value in the EventTime field of the event with the partition dates. The partition dates are in UTC. If you are viewing the EventTime in your local time, you need to convert the time to UTC to find the right partition date.
- 2 Find the hash stored in the database for the partition by running the following command:

```
db.sh sql SIEM dbauser "select name, hash, state from IXLOG_PART where
name='<partition_name>'"
```

where *<partition_name>* is the name of the partition

The hash value is in base64 format. You need to convert the hash value to hexadecimal to compare it with the value returned by sha256sum.

- 3 Convert the hash value found into hexadecimal.

```
echo "<hash>" | base64 -d - | hexdump -e '/1 "%02x" '
```

For example,

```
echo "lUrp+hejhDbyb59ZRpoQ88vpA8eiZfI2ySVCibMoDXo=" | base64 -d - |
hexdump -e '/1 "%02x" '
```

- 4 Calculate the hash of an event partition by executing the following command in the partition directory.

```
cat index.sqfs events.evt/* | sha256sum
```

- 5 Compare the hash values calculated in [Step 2](#) and [Step 3](#). If they match, the event partition has not been tampered with. If they do not match, the integrity of the file has been compromised.

Raw Data Approach

The raw data approach involves finding the raw data associated with the normalized event and proving the raw data has not been tampered.

Sentinel stores the raw data files in one of the following locations:

- ♦ Primary storage location: *<Sentinel data directory>/rawdata/online*
- ♦ Secondary storage location: *<Sentinel secondary storage directory>/rawdata_archive*

If your secondary storage is NFS or CIFS, the NFS/CIFS share is automatically mounted to the */var/opt/novell/sentinel/data/archive_remote* directory on the Sentinel server. If the secondary storage is SAN, the NFS/CIFS share is mounted to the configured directory.

Each raw data file is a .gz compressed file. For more information about raw data files, see [Section 6.1.1, "Raw Data," on page 53](#).

To identify the raw data associated with the normalized event:

- 1 Perform the steps described in section, [Event Data Approach](#). These steps are important because the data in the event is required to find the associated raw data record.
- 2 In the event, find the value in the following fields:
 - ♦ **RawDataRecordId:** The ID of the raw data record that was normalized to create the event is stored in this field. For example, B926DF62-462C-1031-8FE2-000C29E90B7D.
 - ♦ **EventSourceID:** The ID of the event source the data came from. For example, 9DA14E20-4595-1031-BE22-000C29E90B7D. In some cases, the display name is shown for the EventSourceID, such as: sles11sp2:Syslog:Map Output (universal).
 - ♦ **SentinelProcessTime:** the time when Sentinel processed the data. This information is useful as an approximation of which raw data log file the data is stored in.
- 3 In the Sentinel Web interface, click **Storage > Raw Data**.
- 4 Identify the event source in the list that exactly matches the EventSourceID.
- 5 Using the SentinelProcessTime, find the raw data files that have a date that is approximately around this time.

- 6 Download the raw data files that might have the raw data record.
- 7 Open the file and search for the RawDataRecordId.
- 8 After identifying the right raw data file, verify the integrity of the file by clicking **Verify Integrity**.

To determine if the deleted raw data files were tampered:

- ♦ Verify the sequence number of JSON records. All JSON records have the same ChainID with a monotonically increasing ChainSequence number starting with zero. There are no gaps or missing numbers in the ChainID sequence. If a new ChainID is present, its ChainSequence begins with zero. If there are gaps in the sequence of numbers, the records were either tampered or were manually deleted.
- ♦ Verify the RawDataHash against the RawData. To do this, convert the RawData value to a sequence of bytes in UTF-8 format. Calculate a 256 SHA digest against those bytes. Convert the digest to a HEX string, and compare the string with the value in RawDataHash. If they are not identical, either the RawData or the RawDataHash file was tampered.

If, for example, you want to compute the hash of a file on the file system on Linux, specify a command similar to the following:

```
sha256sum F6673C60-573A-102D-ADE0-003048306A7C/2010-06/15-1600.gz
```

For example, if you want to query the database for the hash of a file, you can specify a command similar to the following:

```
db.sh sql SIEM dbauser "select FILE_HASH from RAW_DATA_FILES_INFO where
FILE_NAME='/F6673C60-573A-102D-ADE0-003048306A7C/2010-06/15-1600.gz';"
```

However, there is a possibility that a person tampered the files in such a way that the tampering cannot be detected, because the person also recomputed the sequence number or the RawDataHash. To determine if the raw data files were tampered, you can also use the hash key values of each raw data file stored in the database. The Sentinel server calculates a hash key value for every raw data file and stores it in the RAW_DATA_FILES_INFO in the database.

The table has the following columns:

- ♦ **FILE_NAME:** This column contains the relative file name in the following format:
<Event Source UUID>/<Date>/<RawDataFile>
- ♦ **STATE:** This column indicates if the raw data file is in the primary storage location or the secondary storage location. If the value is ARCHIVED, the raw data file is in the secondary storage location. If the value is ONLINE or COMPRESSED, the raw data file is in the primary storage location. If the value is DELETED, it indicates that the file is deleted from the disk and was not present either in primary or in secondary storage location.
- ♦ **FILE_HASH:** The hash value is computed when the files are closed for writing. Therefore, only files in the COMPRESSED or ARCHIVED state have a hash value. The FILE_HASH column contains a SHA256 hash key value computed over the contents of the file. The file is treated as a stream of binary bytes to compute the hash. The hash is stored as a HEX string (lowercase).

To determine if a file is tampered, compute the SHA-256 hash, convert it to a HEX string (lowercase), then compare this computed value with the hash value stored in the RAW_DATA_FILES_INFO. If the values are different, it indicates that either the file or the database has been tampered.

To determine if the files were deleted in an unauthorized way, you can scan the records in the RAW_DATA_FILES_INFO and look for files whose STATE value is ARCHIVED, ONLINE, or COMPRESSED. You can ignore those marked DELETED. If the STATE value is ARCHIVED, the raw data file should be in the secondary storage location. If the STATE value is ONLINE or COMPRESSED, the raw data file should be in the primary storage location or the secondary storage location.

3.3.8 Using CA Signed Certificates

Sentinel uses several digital, public-key certificates as part of establishing secure TLS/SSL communications. During the initial configuration of Sentinel, these certificates are self-signed. In some circumstances, it might be necessary to obtain certificates digitally signed by a certificate authority (CA).

You can replace the self-signed certificate with a certificate signed by a well-known CA, such as VeriSign, Thawte, or Entrust. You can also replace the self-signed certificate with a certificate digitally signed by a less common CA, such as a CA within your company or organization.

NOTE: There are many well-known CAs and identifying which CAs are most commonly used varies with country.

This section provides information about various certificates used in Sentinel, instructions about configuring the TLS/SSL certificates to get them digitally signed by a CA, and then importing the digitally signed certificates into Sentinel:

- ♦ [“Types of Certificates” on page 28](#)
- ♦ [“Configuring the TLS/SSL Certificates” on page 29](#)
- ♦ [“Using a Signed Certificate on Distributed Components” on page 30](#)

Types of Certificates

- ♦ [“Web Server Certificate” on page 28](#)
- ♦ [“Java Messaging Service Certificates” on page 28](#)
- ♦ [“SSL Proxy Server Certificate” on page 28](#)

Web Server Certificate

The Web server certificate is used for the following purposes:

- ♦ With Web browsers to connect to the Sentinel Web user interface.
- ♦ Establish trust relationships for the REST API calls between Sentinel instances. For example, it is used when configuring Distributed Search.

If the Web server certificate is not signed by a well-known CA and you connect to the Sentinel Web user interface, Sentinel displays the `Connection is Untrusted` message.

Java Messaging Service Certificates

The Java Messaging Service (JMS) certificates include the following:

- ♦ Broker Certificate
- ♦ Client Certificate

The JMS certificates are used to establish secure communications between various components of Sentinel, including the Sentinel server and remote Collector Managers.

SSL Proxy Server Certificate

The Client Proxy Server certificate is used to establish secure communication between the Sentinel server and client applications, such as the Sentinel Control Center or the Solution Designer. This certificate is not used with the Sentinel Web interface.

Configuring the TLS/SSL Certificates

Sentinel provides the `ssl_certs` command line tool that helps with the certificate signing process. This tool is available at:

```
<installation_root>/opt/novell/sentinel/setup/ssl_certs
```

You can run the tool in an interactive mode or with the command line specifications and options. To see the command line options, change to the directory that contains `ssl_certs`, then run the `-help` command.

Configuring the TLS/SSL certificates involves the following steps:

- ♦ [“Generating a Certificate Signing Request” on page 29](#)
- ♦ [“Getting the Certificate Signing Requests Signed by the CA” on page 29](#)
- ♦ [“Importing the Digitally Signed Certificates into Sentinel” on page 29](#)

Generating a Certificate Signing Request

To obtain a digitally signed certificate, you must first generate a certificate signing request, which is presented to the CA. To generate one or more certificate signing requests, perform the following steps on the Sentinel server:

- 1 Log in as the `novell` user, or switch to the `novell` user.
- 2 Change to the `setup` directory:

```
cd $APP_HOME/setup
```
- 3 Run the `./ssl_certs` script.
- 4 Enter 1 to generate certificate signing requests.
- 5 Specify the service for which you want to obtain the signed certificates.
- 6 Specify a filename where the certificate signing request must be saved.
The default filename is based on the internal name of the certificate entry.
- 7 Select another service if necessary, or enter 5 to exit from the service options.
- 8 Enter 4 to exit from the TLS/SSL certificate configuration.

The certificate signing requests are now saved in the specified files.

Getting the Certificate Signing Requests Signed by the CA

- 1 Submit the certificate signing requests to the CA for signature.
- 2 Obtain the signed certificate files from the CA.

The details of how this is done depend on the CA. For more information, consult your CA.

Importing the Digitally Signed Certificates into Sentinel

Copy the files that contains the digital certificates signed by the CA to the Sentinel server. If the files are signed by an enterprise or organizational CA rather than a well-known CA, you must copy the CA's self-signed root certificate to the Sentinel server.

To import the certificate files to the Sentinel server:

- 1 Log in as the `novell` user, or switch to `novell` user.
- 2 Change to the `setup` directory:

```
cd $APP_HOME/setup
```

- 3 Run the `.ssl/certs` command.
- 4 (Conditional) For certificates that are signed by the enterprise or organizational CAs, enter 2, then specify the name of the file that contains the CA root certificates.
- 5 (Conditional) For certificates that are signed by a well-known CA such as Verisign or Entrust, enter 3.
- 6 Select the service for which you obtained the signed certificates.
- 7 Specify the name of the file that contains the CA's signed digital certificate.
- 8 Select another service if necessary, or enter 5 to select Done and exit from the service option.
- 9 Enter 4 to exit from the TLS/SSL certificate configuration.
- 10 Restart Sentinel.

Using a Signed Certificate on Distributed Components

There can be conflicts with the certificates on the distributed components: remote Collector Manager, remote NetFlow Collector Manager, and remote Correlation Engine, depending on the order you installed these components and generated certificates. A common cause of conflicts occurs if you installed the components in the following order:

1. Installed the Sentinel server.
2. Installed a Remote Collector Manager, NetFlow Collector Manager, or a Correlation Engine.
3. You obtained and imported a digitally signed certificate into the Sentinel server first.

If you do things in this order, the certificates on the distributed components do not match the certificates on the Sentinel server. This causes errors when trying to access the Sentinel Web interface and other Sentinel services.

To fix this issue you must import the digitally signed certificate into the distributed components.

3.4 Network Communication Options

Various components of Sentinel communicate across the network, and there are different types of communication protocols used throughout the system. All of these communication mechanisms affect the security of your system.

- ♦ [Section 3.4.1, "Communication between Sentinel, Collector Manager, and Correlation Engine," on page 31](#)
- ♦ [Section 3.4.2, "Communication between Sentinel and NetFlow Collector Manager," on page 32](#)
- ♦ [Section 3.4.3, "Communication between Sentinel and the Sentinel Control Center and Solution Designer Client Applications," on page 32](#)
- ♦ [Section 3.4.4, "Communication between the Server and the Database," on page 32](#)
- ♦ [Section 3.4.5, "Communication with Web Browsers," on page 32](#)
- ♦ [Section 3.4.6, "Communication between the Database and Other Clients," on page 33](#)

3.4.1 Communication between Sentinel, Collector Manager, and Correlation Engine

Sentinel processes include the Sentinel server, Correlation Engine, and Collector Manager.

The communication between these server processes is by default over SSL through the ActiveMQ message bus. The processes use the following configuration information in the `${esecurity.config.home}/config/configuration.xml`:

```
<jms brokerURL="failover://(ssl://
${activemq.ip.server}:${activemq.port.userapps}?wireFormat.maxInactivityDuration=3
0000)?randomize=false" interceptors="compression"
keystore="${esecurity.config.home}/config/.activemqclientkeystore.jks"
keystorePassword="${sentinel.keystore.password}" password-
file="${esecurity.config.home}/config/activemqusers.properties"
username="${activemq.client.username}"/>
```

The jms strategy shown in this XML snippet defines how the Sentinel process connects to the server. This snippet defines the client-side settings of the connection.

Table 3-1 XML Entries in the `configuration.xml` File

XML Entry	Description
ssl://	Indicates that SSL is used for secure connection. You should not modify this value.
\${activemq.ip.server}	The hostname or IP address where the Java message service (JMS) server is running.
\${activemq.port.userapps}	The port that the JMS server is listening on. The default value is 61616.
?wireFormat.maxInactivityDuration=30000)?randomize=false"	This is where ActiveMQ configuration parameters are passed to the transport mechanism. These entries should be modified only if you are an ActiveMQ expert.
interceptors="compression"	Enables compression over the connection. You should not modify this value.
keystore="\${esecurity.config.home}/config/.activemqclientkeystore.jks"	The path to the Java keystore that is used to check if the server is trusted.
keystorePassword="\${sentinel.keystore.password}"	The password to the Java keystore file.
password-file="\${esecurity.config.home}/config/activemqusers.properties"	The location of the file containing the password to present to ActiveMQ for authenticating the connection.
username="\${activemq.client.username}"	The user name to present to ActiveMQ for authenticating the connection. This corresponds to a ActiveMQ user name in the password-file.

The server-side settings are defined in the `/etc/opt/novell/sentinel/config/activemq.xml` file. For instructions about how to edit the `activemq.xml` file, see the [ActiveMQ Web site \(http://activemq.apache.org/\)](http://activemq.apache.org/). However, NetIQ does not support modifying of the server-side settings.

3.4.2 Communication between Sentinel and NetFlow Collector Manager

The NetFlow Collector Manager sends NetFlow data to the Sentinel server over HTTPS connections.

This connection utilizes a signed certificate from the Sentinel server to ensure that the NetFlow Collector Manager communicates with the Sentinel server specified during NetFlow Collector Manager installation. During the NetFlow Collector Manager installation, you are prompted to accept the certificate from the Sentinel server. Upon accepting the certificate, the Sentinel server stores the certificate in the `/etc/opt/novell/sentinel/config/.webserverkeystore.jks` file. When the NetFlow Collector Manager initiates a communication with the Sentinel server, the server validates the connection to ensure that the certificate matches.

In addition to the certificate check, the NetFlow Collector Manager also authenticates a SAML token with a valid Sentinel user name and password, which has the appropriate permission to send network flow data to Sentinel.

3.4.3 Communication between Sentinel and the Sentinel Control Center and Solution Designer Client Applications

The Sentinel Control Center and Solution Designer client applications use SSL communication through the SSL proxy server by default.

The client applications use SSL by reading the following information in `/etc/opt/novell/sentinel/config/configuration.xml`:

```
<strategy active="yes" id="proxied_client"
location="com.esecurity.common.communication.strategy.proxystrategy.ProxiedClients
strategyFactory">
  <transport type="ssl">
    <ssl host="10.0.0.1" port="10013" keystore="./novell/sentinel/
.proxyClientKeystore" />
  </transport>
</strategy>
```

3.4.4 Communication between the Server and the Database

The protocol used for communication between the server and the database is defined by a JDBC driver.

Sentinel uses the PostgreSQL driver (`postgresql-version.jdbc3.jar`) to connect to the PostgreSQL database, which is a Java (Type IV) implementation. This driver supports encryption for data communication. To download the driver, refer to the [PostgreSQL Download Page \(http://jdbc.postgresql.org/download.html\)](http://jdbc.postgresql.org/download.html). To configure the encryption, refer to [PostgreSQL Encryption Options \(http://www.postgresql.org/docs/8.1/static/encryption-options.html\)](http://www.postgresql.org/docs/8.1/static/encryption-options.html).

NOTE: Turning on encryption has a negative impact on the performance of the system. Therefore, this security concern needs to be weighed against your performance needs. The database communication is not encrypted by default for this reason. Lack of encryption is not a major concern because communication with the database occurs over the localhost network interface.

3.4.5 Communication with Web Browsers

The Web server is by default configured to communicate via HTTPS. For more information, see the [Jetty documentation \(http://wiki.eclipse.org/Jetty/Howto/Configure_SSL\)](http://wiki.eclipse.org/Jetty/Howto/Configure_SSL).

3.4.6 Communication between the Database and Other Clients

You can configure the PostgreSQL SIEM database to allow connections from any client computer that uses pgAdmin or another third-party application.

The PostgreSQL database is compiled with the `--with-openssl` flag. You can configure it to use encrypted communication, although that is not the default setting. Typically all database communication in Sentinel is performed locally and not over the network.

To allow pgAdmin to connect from any client computer, add the following line in the `/var/opt/novell/sentinel/3rdparty/postgresql/data/pg_hba.conf` file:

```
host all all 0.0.0.0/0 md5
```

If you want to limit the client connections that are allowed to run and connect to the database through pgAdmin, specify the IP address of the host in the above line.

The following line in the `pg_hba.conf` file is an indicator to PostgreSQL to accept connections from the local computer so that pgAdmin is allowed to run only on the server.

```
host all all 127.0.0.1/32 md5
```

To allow connections from other client computers, you can add additional host entries in the `pg_hba.conf` file.

To provide maximum security, by default, PostgreSQL only allows connections from the local computer.

3.5 Sensitive Data Locations

For certain components, passwords must be stored so that they are available to the components when the system needs to connect to a resource such as a database or an event source. In this case, the password is first encrypted to avoid unauthorized access to the clear-text password.

Even if the password is encrypted, you must ensure that the access to the stored password data is protected to avoid password exposure. For example, you can set permissions to ensure that files with sensitive data are not readable by other users.

Database credentials are stored in the `/etc/opt/novell/sentinel/config/object-component.ConnectionManager.properties` file.

```
username=appuser
database=SIEM
password=<password>
```

The following database tables store passwords (/certificate) in encrypted format. You must limit access to these tables.

- ♦ **EVT_SRC**: column: `evt_src_config` column data
- ♦ **evt_src_collector**: column: `evt_src_collector_props`
- ♦ **evt_src_grp**: column: `evt_src_default_config`
- ♦ **md_config**: column: data
- ♦ **integrator_config**: column: `integrator_properties`
- ♦ **md_view_config**: column: `view_data`
- ♦ **esec_content**: column: `content_context`, `content_hash`
- ♦ **esec_content_grp_content**: column: `content_hash`

- ♦ **sentinel_plugin:** column: content_pkg, file_hash

Sentinel stores both configuration data and event data in the following locations:

Table 3-2 Locations for Configuration Data and Event Data

Components	Location for Configuration Data	Location for Event Data
Sentinel server	<p>The database tables and file system at <code>/etc/opt/novell/sentinel/config</code>.</p> <p>This configuration information includes the encrypted database, event source, integrators, and passwords.</p>	<p>The database (for example, <code>CORRELATED_EVENTS</code> and <code>EVT_RPT_*</code> tables) and the file system at <code>/var/opt/novell/sentinel/data/eventdata</code>, <code>/var/opt/novell/sentinel/data/rawdata</code>, <code>/var/opt/novell/sentinel/data/server.cache</code>, <code>/var/opt/novell/sentinel/data/map_data</code>, and <code>/var/opt/novell/sentinel/3rdparty/mongodb</code>.</p>
Collector Manager	<p>The file system at <code>/etc/opt/novell/sentinel/config</code>. The most sensitive configuration information is the client key pair used to connect to the message bus.</p>	<p>Event data might be cached on the file system during error conditions such as the message bus being down or event overflow. This event data is stored in the <code>/var/opt/novell/sentinel/data/collector_mgr.cache</code> directory.</p>

3.6 Implementing Intruder Detection and Lockout Mechanisms

Sentinel supports intruder detection and lockout to prevent potential brute-force attacks. Sentinel provides several configurable parameters that help you implement intruder detection and lockout mechanisms.

- ♦ **failedAuthDelay:** Specifies the duration in milliseconds that a subsequent authentication request must wait after a failed authentication for a specific user. The default value is 2000 (2 seconds). If the value is 0, the delay is disabled. This wait period is calculated for each user. If an authentication request for User A fails, it does not cause a delay for an authentication request for User B.
- ♦ **intruderDetectInterval:** Specifies the time period in milliseconds in which consecutive failed authentication requests for a user must occur for Sentinel to identify the failures as a possible intruder detection. For example, if the value is 300000 (5 minutes) and four failed authentication requests happen within 4 minutes, but the 5th consecutive request happens 5:01 (minutes:seconds) later than the 1st failed request, Sentinel does not consider the requests suspicious. If the value is 360000 (6 minutes) and the same sequence of failed requests happen, Sentinel considers the requests to be suspicious. The default value for this parameter is 300000 (5 minutes).
- ♦ **intruderDetectMaxFailedAttempts:** Specifies the number of consecutive, failed authentication requests that must occur for Sentinel to consider a user name during the `intruderDetectInterval` for the requests as suspicious. If the value is 0 then intruder detection and lockout is disabled. The default value for this parameter is 5.
- ♦ **intruderDetectLockPeriod:** Specifies the duration that a Sentinel user account remains locked when the user account is automatically locked in response to a suspicious series of failed authentication requests. If the value is 0, automatically locked accounts are not automatically unlocked. They must be unlocked manually by an administrator. The default value for this parameter is 900000 (15 minutes).

- ♦ **intruderDetectAdminAutoLock:** Specifies whether or not the Sentinel admin account is subject to automatic locking in response to a series of failed authentication requests. The default is `false` since a denial-of-service attack exists in which an attacker can continually lock the built-in admin account, unless there is a separate administrator account.

The values listed above are defined in the `AuthenticationService` component of the `/etc/opt/novell/sentinel/config/server.xml` file. After you make any manual modifications to the values, you must place this component in the component properties override file `/etc/opt/novell/sentinel/config/obj-component.AuthenticationService.properties` to ensure that the modified settings are not lost during an upgrade.

3.7 Applying Updates for Security Vulnerabilities in Embedded Third-Party Products

Sentinel contains embedded third-party products such as JRE, Jetty, PostgreSQL, and ActiveMQ. Sentinel includes patches to address the security vulnerabilities (CVE) for these products when updates for Sentinel are released.

However, each of these products has its own release cycle, which means that there might be CVEs that are discovered before a Sentinel update is released. You need to separately review the CVEs for each embedded third-party product, and decide whether to apply these updates to your Sentinel system outside of the Sentinel updates.

If you decide to apply patches to address these CVEs outside of a Sentinel update, use the instructions in TID 7008820 in the [NetIQ Support Knowledge Base](#).

II Configuring Roles and Users

This section provides information about configuring roles and users that can use Sentinel.

- ♦ [Chapter 4, “Configuring Roles and Users,” on page 39](#)
- ♦ [Chapter 5, “Configuring LDAP Authentication,” on page 45](#)

4 Configuring Roles and Users

In Sentinel, you can add, edit, and delete roles. You can also grant different permissions at the role level, and edit the details of user and role profiles.

- ♦ [Section 4.1, “Overview,” on page 39](#)
- ♦ [Section 4.2, “Creating Roles,” on page 40](#)
- ♦ [Section 4.3, “Configuring Password Complexity,” on page 42](#)
- ♦ [Section 4.4, “Creating Users,” on page 43](#)

4.1 Overview

You can create different user roles and assign them different permissions. Role assignment helps you control users access to functionality, data access based on fields in the incoming events, or both. Each role can contain any number of users. Users belonging to the same role inherit the permissions of the role they belong to. You can set multiple permissions for a role.

Sentinel has the following roles by default:

Administrator: A user in this role has administrative rights in the Sentinel system. You cannot delete users in this role. Administrative rights include the ability to perform user administration, data collection, data storage, search operations, rules, report, dashboard, and license management.

You cannot modify or delete the administrator role.

Database Administrator: A user in this role has access to events coming from database event sources. The Collector parsing the data from the event source determines the type of the event source (database). A user in this role can view data that matches filter `rv32: "DB"` and search remote targets.

Incident Administrator: A user in this role can manage incidents in the system and control incidents being handled by other users.

NetFlow Provider: A user in this role can send NetFlow data to Sentinel and manage NetFlow Collector Managers. A user in this role can also view and analyze the NetFlow data.

Network Administrator: A user in this role can administer network infrastructure devices, such as routers, switches, and VPNs. This role has access to events coming from devices in the category `NETD` or `VPN` (as determined by the Collector parsing the data) or from event sources with the `Network` tag. Set the `Network` tag on network infrastructure event sources to allow users in this role to view the events. A user in this role can view data that matches filter `rv32: "NETD" OR rv32: "VPN" OR rv145: "Network"`, and can search remote targets.

Network Security Administrator: A user in this role can administer network security infrastructure devices, such as firewalls, IDSs, and Web proxies. This role has access to events coming from devices in the category `AV`, `FW`, or `IDS` (as determined by the Collector parsing the data) or from event sources with the `NetworkSecurity` tag. Set the `NetworkSecurity` tag on network infrastructure

event sources to allow users in this role to view the events. A user in this role can view data that matches filter `rv32:"AV" OR rv32:"FW" OR rv32:"IDS" OR rv145:"NetworkSecurity"`, and can search remote targets.

PCI Compliance Auditor: A user in this role has access to view events that are tagged with at least one of the regulation tags such as PCI, SOX, HIPAA, NERC, FISMA, GLBA, NISPOM, JSOX, and ISO/IEC_27002:2005, and can view system events, view the Sentinel configuration data, and search remote targets.

Report Administrator: A user in this role can run reports, view, rename and delete report results, add and delete report templates and report results, run reports on configuration database, export all reports, and save search results as a report. A Report Administrator can also tag report templates and report results. The Report Administrator can search report templates and report results based on these tags.

Search Proxy User: This is a system role for proxy users. This role is critical to setting up another Sentinel system to access your local Sentinel system using the Distributed Search feature.

Security Policy Administrator: A user in this role can implement the security policies within the system for users to access anomaly detection, correlation, incident remediation, and iTRAC workflows.

System Event Monitor: A user in this role can monitor the Sentinel system for errors or outages. This role has access only to events coming from Sentinel systems. A user in this role can also access data coming from event sources that Sentinel is dependent on. For example, you can tag operating systems on which Sentinel and the Collector Managers are running with a Sentinel event source tag so that the users in this role can monitor problems in the operating systems. A user in this role can view data that matches filter `rv145:"Sentinel"`, view system events, and search remote targets.

Unix Administrator: A user in this role has access to events from operating system event sources that are not Windows computers. The type of the event source is determined by verifying the Collector parsing data and also by verifying if a Windows tag is present. A user in this role can view data that matches filter `(rv32:"OS" NOT ((("Microsoft?Active?Directory*" NOT msg:"Microsoft?Active?Directory*") OR ("Microsoft?Windows*" NOT msg:"Microsoft?Windows*")))) NOT rv145:"Windows"` and search remote targets.

User: A user in this role can manage dashboards, run reports, view and rename reports, and delete report results.

Windows Administrator: A user in this role can administer Windows computers. This role has access to data generated by Windows event sources. The type of the event source is determined by verifying the Collector parsing the data. If data from a Windows event source is not being processed by the Active Directory or the Windows Collector, add the Windows tag to event sources to indicate that Windows data is being collected from the event source. This enables the Windows administrator to access the data. A user in this role can view data that matches filter `(rv32:"OS" AND ((("Microsoft?Active?Directory*" NOT msg:"Microsoft?Active?Directory*") OR ("Microsoft?Windows*" NOT msg:"Microsoft?Windows*")))) OR rv145:"Windows"` and search remote targets.

4.2 Creating Roles

Roles allow you define what a user can manage and what data they can view. Permissions are granted to the role, and then the user is assigned to the role.

4.2.1 Creating a Role

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Users** in the toolbar.
- 3 Click **Create** in the **Roles** section to create a new role.
- 4 Use the following information to create the role:

Role name: Specify a unique name for the role. A role name should not exceed 40 characters.

Description: Specify a description of the role.

Users with this role can: Select the permissions that a role grants to users assigned to the role.

- ♦ **View all data:** Select this option to allow users to view all the data in the Sentinel system. If you select this option, you must select one or more of the following permissions:
 - ♦ **Manage Correlation Engine/Rules:** Allows users to manage Correlation rules and all data associated with these rules. The Correlation feature is displayed in the Web interface only if this permission is selected.
 - ♦ **Manage Reports:** Allows users to view and manage the data in reports.
 - ♦ **Manage and View Security Intelligence Dashboards:** Allows user to view, create, and manage the Security Intelligence dashboards and the data displayed in the dashboards. The Security Intelligence option is displayed in the Web interface only if this permission is selected.
 - ♦ **View Security Intelligence Dashboards:** Allows user to view the Security Intelligence dashboards and the data displayed in the dashboards. The Security Intelligence option is displayed in the Web interface only if this permission is selected.
- ♦ **View the following data:** Select this option to allow users to view only selected data in the Sentinel system.
 - ♦ **Only events matching the filter:** Allows users to view only the events returned by the specified search query. For example, if you set the filter value to `sev:5`, users with this permission can view only events of severity five in a search.

For more information about using filters, see [“Configuring Filters”](#) in the *NetIQ Sentinel User Guide*.

Select one or more of the following permissions to use when viewing the filtered data:
 - ♦ **View NetFlow data:** Allows users to view and analyze the network flow data.
 - ♦ **Search Remote Targets:** When this permission is set on a role, all members of that role can perform searches on event sources that are in a distributed location.

For more information on distributed searching and reporting, see [Chapter 18, “Searching and Reporting Events in a Distributed Environment,”](#) on page 177.
 - ♦ **View asset data:** Allows users to view asset data.
 - ♦ **View asset vulnerability data:** Allows users to view vulnerability data.
 - ♦ **View data in the embedded database:** Allows users to view the data in the embedded database.
 - ♦ **View people browser:** Allows users to view the data in the Identity Browser.
 - ♦ **View system events:** Allows users to view the Sentinel system events.
- ♦ **Incidents:** Select one of the followings permissions that enable users to manage incidents:
 - ♦ **View incidents assigned to user:** Allows a user to view any incident that is assigned to them.

- ♦ **View or create incidents and add events to incidents:** Allows users to create incidents and add events to the incidents.
- ♦ **Create, modify and execute actions on assigned incidents:** Allows users to create, modify, and execute actions on incidents that are assigned to them.
- ♦ **Manage all aspects of incidents: create, modify and delete:** Allows users to manage all incidents.
- ♦ **Miscellaneous:** Assign miscellaneous permissions as necessary:
 - ♦ **Create and use Active Views:** When this permission is set on a role, all members of this role can access and use **Active Views** in the Sentinel Control Center. For more information about Active Views, see “[Viewing Events](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ **Manage Tags:** When this permission is set on a role, all members of this role can create, delete, and modify tags, and associate tags to different event sources. For more information about tags, see “[Configuring Tags](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ **Proxy for Authorized Search Initiators:** When this permission is set on a role, the members of this role can accept searches from remote targets. For more information, see [Chapter 18, “Searching and Reporting Events in a Distributed Environment,” on page 177](#).
 - ♦ **Send NetFlow data:** Allows users to send network flow data from the NetFlow Collector Manager to the Sentinel server.
 - ♦ **Share search filters:** When this permission is set on a role, all members of this role can share search filters that they have created. For more information about sharing filters, see “[Configuring Filters](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ **Solution Designer access:** When this permission is set on a role, all members of this role can access Solution Designer. For more information, see [Section 20.7, “Solution Designer,” on page 210](#).
 - ♦ **View and execute event actions:** When this permission is set on a role, all members of this role can view events and execute actions on the selected events. For more information, see “[Manually Performing Actions on Events](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ **View detailed internal system state data:** When this permission is set on a role, all members of this role can view detailed internal system state data by using a JMX client.

5 Click **Save**.

To create users for this role, see [Section 4.4, “Creating Users,” on page 43](#).

4.3 Configuring Password Complexity

A complex password improves security by preventing password guessing attacks. Sentinel provides a set of password validation rules that help you maintain a complex password for all local user passwords. You can select the desired validation rules as applicable for your environment.

You can configure the password validation rules in the `/etc/opt/novell/sentinel/config/passwordrules.properties` file. The validation rules apply only to the local user passwords and not LDAP user passwords. For existing users, validation rules apply only after the users update their password.

By default, all the validation rules are disabled and commented with `#`. To enable validation rules, uncomment the rules, specify the values for the rules, and save the file.

The following table describes the password complexity validation rules:

Table 4-1 Password Complexity Rules

Validation Rule	Description
MINIMUM_PASSWORD_LENGTH	Specifies the minimum number of characters required in a password.
MAXIMUM_PASSWORD_LENGTH	Specifies the maximum number of characters allowed in a password.
UNIQUE_CHARACTER_LENGTH	Specifies the minimum number of unique characters required in a password. For example, if the UNIQUE_CHARACTER_LENGTH value is 6 and a user specifies the password as "aaaabbccc", the Sentinel does not validate the password because it contains only 3 unique characters a, b, and c.
LOWER_CASE_CHARACTERS_COUNT	Specifies the minimum number of lowercase characters required in a password.
UPPER_CASE_CHARACTERS_COUNT	Specifies the minimum number of uppercase characters required in a password.
ALPHABET_CHARACTERS_COUNT	Specifies the minimum number of alphabetic characters required in a password.
NUMERIC_CHARACTERS_COUNT	Specifies the minimum number of numeric characters required in a password.
NON_ALPHA_NUMERIC_CHARACTERS_COUNT	Specifies the minimum number of non-alphanumeric or special characters required in a password. The rule considers only the following non-alphanumeric characters: ` ~ ! @ # \$ % ^ & * () - _ = + [{] } \ ; : ' " < , > . / ?
RESTRICTED_WORDS_IN_PASSWORD	Specifies the words that are not allowed in a password. The restricted words are case-insensitive. You can specify multiple words separated by a comma. For example, RESTRICTED_WORDS_IN_PASSWORD= admin,password,test

4.4 Creating Users

Adding a user in the Sentinel system creates an application user who can then log in to Sentinel. You also assign roles when you create the user.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Users**.
- 3 Click **Create** in the **Users** section.
- 4 Specify the name and e-mail address of the user.

The fields with an asterisk (*) are mandatory, and the user name must be unique.

A user name cannot exceed 30 characters, and you can use extended characters when you create it.

- 5 Select a role for the user.

6 Select the authentication type:

Local: Select this option for the server to authenticate the user login against the internal database. By default, the **Local** option is selected.

Directory: The **Directory** option is enabled only if you have configured the Sentinel server for LDAP authentication. Select this option for the server to authenticate the user login against an LDAP directory.

7 (Conditional) If you specified Local for the authentication type in [Step 6](#), specify any user name in the Username field and continue with [Step 9](#).

8 (Conditional) If you specified Directory for the authentication type in [Step 6](#), specify the user name according to the settings you used when you configured LDAP, then continue with [Step 11](#).

- ♦ **If you selected Yes for Anonymous Search:** The user name must be the same as the LDAP directory user name.
- ♦ **If you selected No for Anonymous Search and did not specify the domain name:** The user name does not need to be the same as the LDAP directory user name.

You must also specify the **LDAP User DN**. If Base DN was set, the Base DN is appended to the relative user DN to construct the absolute user DN.

For example, if the Base DN was set to `o=netiq` and the absolute user DN is `cn=sentinel_ldap_user,o=netiq` only the relative user DN for example, `cn=sentinel_ldap_user` can be specified.

When some reserved special characters are used as literals in an **LDAP User DN**, they must be escaped with a backslash (`\`). The following characters must be escaped:

- ♦ A space or '#' character occurring at the beginning of the string
- ♦ A space character occurring at the end of the string
- ♦ Any one of the characters, +, ", \, <, > or ;

For more information, see [LDAPv3 Distinguished Names \(http://www.ietf.org/rfc/rfc2253.txt\)](http://www.ietf.org/rfc/rfc2253.txt).

For example, if the **LDAP User DN** contains a ',' (comma) as a literal, specify the LDAP User DN as follows:

```
CN=Test\,User,CN=Users,DC=netiq,DC=com
```

eDirectory or Active Directory might require additional characters to be escaped. Refer the eDirectory or Active Directory documentation for any additional characters to be escaped.

- ♦ **If you selected No for Anonymous Search and specified the domain name:** The user name must be the same as the LDAP directory user name.

9 Specify a password in the **Password** field.

NOTE: For local user password, ensure that the password adheres to the password complexity validation rules. For more information, see [“Configuring Password Complexity” on page 42](#).

10 Re-enter the password in the **Verify** field.

11 The **Title**, **Office #**, **Ext**, **Mobile #**, and **Fax**. fields are optional. The phone number fields allow any format. Make sure you enter a valid phone number so that the user can be contacted directly.

12 Click **Save**.

5 Configuring LDAP Authentication

Sentinel supports LDAP authentication in addition to database authentication. You can configure a Sentinel server for LDAP authentication to enable users to log in to Sentinel with their LDAP directory credentials.

NOTE: Sentinel LDAP authentication has been tested with Novell eDirectory and Microsoft Active Directory. Other LDAP compliant directories might be used, but have not been tested. If an issue is encountered with a directory that has not been tested, support will be provided to the extent that the issue can be reproduced on one of the tested directories.

- ♦ [Section 5.1, “Overview,” on page 45](#)
- ♦ [Section 5.2, “Prerequisites,” on page 46](#)
- ♦ [Section 5.3, “Setting Up LDAP Authentication,” on page 46](#)
- ♦ [Section 5.4, “Logging in by Using LDAP User Credentials,” on page 49](#)
- ♦ [Section 5.5, “Configuring Multiple LDAP Servers for Failover,” on page 49](#)

5.1 Overview

LDAP authentication can be performed either using an SSL connection or an unencrypted connection to the LDAP server.

You can configure the Sentinel server for LDAP authentication either with or without using anonymous searches on the LDAP directory.

NOTE: If anonymous search is disabled on the LDAP directory, you must not configure the Sentinel server to use anonymous search.

- ♦ **Anonymous:** When you create Sentinel LDAP user accounts, the directory user name must be specified and the user distinguished name (DN) does not need to be specified.

When the LDAP user logs in to Sentinel, the Sentinel server performs an anonymous search on the LDAP directory based on the specified user name, finds the corresponding DN, then authenticates the user login against the LDAP directory by using the DN.

- ♦ **Non Anonymous:** When you create Sentinel LDAP user accounts, the user DN must be specified along with the user name.

When the LDAP user logs in to Sentinel, the Sentinel server authenticates the user login against the LDAP directory by using the specified user DN and does not perform any anonymous search on the LDAP directory.

There is an additional approach applicable only for Active Directory. For more information, see [“Domain Name:” on page 47](#).

5.2 Prerequisites

- [Section 5.2.1, “Exporting the LDAP Server CA Certificate,” on page 46](#)
- [Section 5.2.2, “Enabling Anonymous Search in the LDAP Directory,” on page 46](#)

5.2.1 Exporting the LDAP Server CA Certificate

If you want to connect to the LDAP server by using an SSL connection and the LDAP server certificate is not signed by a well-known certificate authority (CA), you must export the LDAP server CA certificate to a Base64-encoded file.

- ♦ **eDirectory:** See “Exporting an Organizational CA’s Self-Signed Certificate” (<http://www.novell.com/documentation/edir88/edir88/?page=/documentation/edir88/edir88/data/a7elxuq.html>).

To export an eDirectory CA certificate in iManager, the Novell Certificate Server plug-ins for iManager must be installed.

- ♦ **Active Directory:** See “How to enable LDAP over SSL with a third-party certification authority” (<http://support.microsoft.com/kb/321051>).

5.2.2 Enabling Anonymous Search in the LDAP Directory

To perform LDAP authentication using anonymous search, you must enable anonymous search in the LDAP directory. By default, anonymous search is enabled in eDirectory and is disabled in Active Directory.

- ♦ **eDirectory:** See `ldapBindRestrictions` in section [Attributes on the LDAP Server Object](#) (<http://www.novell.com/documentation/edir88/edir88/data/agq8auc.html>).
- ♦ **Active Directory:** Enabling anonymous binds for Active Directory requires two steps. These steps are the same for both Windows 2003 and Windows 2008 Active Directory.

- ♦ **Enable Anonymous LDAP Operations:** By default, anonymous LDAP operations are disabled in Active Directory. You must enable anonymous LDAP operations in Active Directory by setting the `dsHeuristics` attribute to an appropriate value.

For more information, see [Anonymous LDAP operations in Windows 2003 AD](#) (http://www.petri.co.il/anonymous_ldap_operations_in_windows_2003_ad.htm).

- ♦ **Assign Permissions to the ANONYMOUS LOGON User:** The Read and List Contents permissions must be assigned to the ANONYMOUS LOGON user.

For more information, see [Granting anonymous read access](#) (http://www.petri.co.il/anonymous_ldap_operations_in_windows_2003_ad.htm).

5.3 Setting Up LDAP Authentication

- 1 Log in to Sentinel Web interface as a user in the administrator role.
- 2 Click **Users** in the toolbar.
- 3 On the Users page, click the **LDAP Settings** tab.
- 4 Specify the following to configure LDAP authentication:

Host: Specify the hostname or the IP address of the LDAP server.

This is a required field if you select the SSL option.

SSL: Select this option if you want to connect to the LDAP server by using a Secure Socket Layer (SSL) connection.

Port: Specify the port number for the LDAP connection. The default SSL port number is 636 and the default non-SSL port number is 389.

Certificate File Path: Specify the path of the CA certificate file for the LDAP server.

This field should be used only if you selected the SSL option and if the LDAP server certificate is not signed by well-known CA and is not trusted by default.

Anonymous Search: Select **Yes** to perform anonymous searches or select **No** if you do not want to perform anonymous searches on the LDAP directory.

Base DN: Specify the root container to search for users, such as o=netiq for eDirectory or cn=users,dc=example,dc=co for Active Directory.

- ♦ **If Anonymous Search is Yes:** Specify the root container in the LDAP directory to search for users.

This is optional for eDirectory, and mandatory for Active Directory. For eDirectory, if the Base DN is not specified, the entire directory is searched to locate the users.

- ♦ **If Anonymous Search is No:** Specify the root container in the LDAP directory that contains the users.

This is mandatory if you are using Active Directory and if you set a domain name. For all other cases, this is optional.

Search Attribute: Specify the LDAP attribute holding the user login name. This is used to search for users.

For example:

- ♦ eDirectory:

uid

- ♦ Active Directory:

sAMAccountName

This field is available only if you selected **Yes** for Anonymous Search.

Domain Name: Specify the name of the Active Directory domain.

This is an additional approach applicable only for Active Directory for performing LDAP authentication without using anonymous search.

When you specify the Domain Name, username@domainname (userPrincipalName) is used to authenticate the user before searching for the LDAP user object.

For example, test.example.com

This field is applicable only for Active Directory and is available only if you selected **No** for Anonymous Search.

NOTE: If **Base DN** is set and **Domain Name** is not set, the **Base DN** is appended to the relative user DN to construct the absolute user DN.

For example, if the Base DN is set to o=netiq and the absolute user DN is cn=sentinel_ldap_user,o=netiq when the LDAP user account is created, only the relative user DN of cn=sentinel_ldap_user can be specified.

5 Click **Test Connection** to test whether the LDAP connection is successful.

5a Specify the test credentials to connect to the LDAP server:

If Anonymous Search is Yes: Specify the user name and password.

If you selected No for Anonymous Search and did not specify the Domain Name:

Specify the user DN and password. The user DN can be relative to the Base DN.

The **User DN** is based on the RFC 2253 standard. According to RFC 2253, when some reserved special characters are used as literals in a **User DN**, they must be escaped with a backslash (\). The following characters must be escaped:

- ♦ A space or # character occurring at the beginning of the string
- ♦ A space character occurring at the end of the string
- ♦ One of the characters , + , " , \ , < , > or ;

For more information, see [RFC 2253 \(http://www.ietf.org/rfc/rfc2253.txt\)](http://www.ietf.org/rfc/rfc2253.txt).

For example, if the **User DN** contains a comma (,) as a literal, specify the **User DN** as follows:

```
CN=Test\,User,CN=Users,DC=netiq,DC=com
```

eDirectory or Active Directory might require additional characters to be escaped. Refer the eDirectory or Active Directory documentation for any additional characters to be escaped.

If you selected No for Anonymous Search and specified the Domain Name: Specify the user name and password.

5b Click **Test** to test the LDAP connection.

A message is displayed that indicates whether the connection is successful.

If there is an error, review the configuration details you provided and test the connection again. You can determine the cause of the failure by examining the `/var/opt/novell/sentinel/log/server0.0.log` file. You must ensure that the test connection is successful before saving the LDAP settings.

6 Click **Save** to save the LDAP settings.

On successful configuration:

- ♦ The `LdapLogin` section of the `/etc/opt/novell/sentinel/config/auth.login` file is updated. For example:

```
LdapLogin {
    com.sun.security.auth.module.LdapLoginModule required
    java.naming.ldap.factory.socket="com.esecurity.common.communication.ProxyL
dapSSLSocketFactory"
    userProvider="ldap://10.0.0.1:636/o=netiq"
    userFilter="( & (uid={USERNAME}) (objectclass=user) )"
    useSSL=true;
};
```

- ♦ The LDAP server CA certificate, if provided, is added to a keystore named `/etc/opt/novell/sentinel/config/.ldapkeystore.jks`.

After saving the LDAP settings successfully, you can create LDAP user accounts to enable users to log in to Sentinel by using their LDAP directory credentials.

NOTE: You can also configure the Sentinel server for LDAP authentication by running the `ldap_auth_config.sh` script in the `/opt/novell/sentinel/setup` directory.

The script also supports command line options. To view the command line options, run the script as follows:

```
/opt/novell/sentinel/setup/ldap_auth_config.sh --help
```

5.4 Logging in by Using LDAP User Credentials

After you successfully configure the Sentinel server for LDAP authentication, you can create Sentinel LDAP user accounts. For more information on creating LDAP user accounts, see [Section 4.4, “Creating Users,” on page 43](#).

After you create the LDAP user account, you can log in to the Sentinel by using your LDAP user name and password.

5.5 Configuring Multiple LDAP Servers for Failover

To configure one or more LDAP servers as failover servers for LDAP authentication:

- 1 Log in to the Sentinel server as root user.
- 2 Switch to the novell user:

```
su - novell
```

- 3 Change to the `/etc/opt/novell/sentinel/config` directory:

```
cd /etc/opt/novell/sentinel/config/
```

- 4 Open the `auth.login` file for editing:

```
vi auth.login
```

- 5 Update the `userProvider` in the `LdapLogin` section to specify multiple LDAP URLs. Separate each URL by a blank space.

For example:

```
userProvider="ldap://primary_server_IP:port/BaseDN ldap://  
failover_server_IP:port/BaseDN"
```

For Active Directory, ensure that the BaseDN in the LDAP URL is not blank.

For more information on specifying multiple LDAP URLs, see the description of the `userProvider` option in “Class LdapLogin Module” (<http://java.sun.com/javase/6/docs/jre/api/security/jaas/spec/com/sun/security/auth/module/LdapLoginModule.html>).

- 6 Save the changes.

If you are using an SSL connection to the LDAP server and if the LDAP server certificate is not signed by a well-known CA, you must perform the following additional steps:

- 1 Export the certificate of each failover LDAP server and copy the certificate file to the `/etc/opt/novell/sentinel/config` directory on the Sentinel server.

For more information, see [Section 5.2.1, “Exporting the LDAP Server CA Certificate,” on page 46](#).

- 2 Ensure that you set the necessary ownership and permissions of the certificate file for each LDAP server.

```
chown novell:novell /etc/opt/novell/sentinel/config/<cert-file>
```

```
chmod 600 /etc/opt/novell/sentinel/config/<cert-file>
```

- 3 Add each LDAP server certificate to the keystore named `.ldapkeystore.jks`.

```
/opt/novell/sentinel/jre/bin/keytool -importcert -noprompt -trustcacerts -file  
<certificate-file> -alias <alias_name> -keystore /etc/opt/novell/sentinel/  
config/.ldapkeystore.jks -storepass password
```

Replace *<certificate-file>* is the LDAP certificate filename and *<alias_name>* with the alias name for the certificate to be added.

IMPORTANT: Ensure that you specify the alias. If no alias is specified, the keytool takes *mykey* as the alias by default. When you import multiple certificates into the keystore without specifying an alias, the keytool reports an error that the alias already exists.

In some environments, the Sentinel server might not connect to the failover LDAP server if the Sentinel server times out before it finds that the primary LDAP server is down. In such cases, perform the following additional steps to ensure that the Sentinel server connects to the failover LDAP server without timing out:

- 1 Open the `sysctl.conf` file for editing:

```
vi /etc/sysctl.conf
```

- 2 Ensure that the `net.ipv4.tcp_syn_retries` value is set to 3. If the entry does not exist, add the entry. Save the file:

```
net.ipv4.tcp_syn_retries = 3
```

- 3 Execute the following commands for the changes to take effect:

```
/sbin/sysctl -p
```

```
/sbin/sysctl -w net.ipv4.route.flush=1
```

- 4 Open the `server.conf` file for editing:


```
vi /etc/opt/novell/sentinel/config/server.conf
```

- 5 Set the Sentinel server time out value to 60 seconds by appending a new parameter in the *Java Additional Parameters* section as follows:

```
wrapper.java.additional.53=-Desecurity.remote.timeout=60
```

- 6 Restart the Sentinel server:

```
/etc/init.d/sentinel restart
```



Collecting and Routing Event Data

This section provides information about collecting and routing events.

- ♦ [Chapter 6, “Configuring Data Storage,” on page 53](#)
- ♦ [Chapter 7, “Configuring Agentless Data Collection,” on page 79](#)
- ♦ [Chapter 8, “Configuring Agent-Based Data Collection,” on page 109](#)
- ♦ [Chapter 9, “Managing Event Sources,” on page 111](#)
- ♦ [Chapter 10, “Configuring Event Routing Rules,” on page 119](#)
- ♦ [Chapter 11, “Mapping Events,” on page 123](#)
- ♦ [Chapter 12, “Linking to Additional Sentinel Systems,” on page 135](#)

6 Configuring Data Storage

Sentinel stores raw data and compressed event data on the primary location. You can configure Sentinel to store the data in a secondary location for long-term storage.

The data files are deleted from the primary and secondary storage locations on a configured schedule. Raw data retention is governed by a single raw data retention policy. Data retention is governed by a set of event data retention policies. All of these policies are configured by the Sentinel administrator.

- ♦ [Section 6.1, “Types of Data,” on page 53](#)
- ♦ [Section 6.2, “Configuring Secondary Storage Locations,” on page 60](#)
- ♦ [Section 6.3, “Configuring Disk Space Usage,” on page 64](#)
- ♦ [Section 6.4, “Verifying and Downloading Raw Data Files,” on page 64](#)
- ♦ [Section 6.5, “Configuring Data Retention Policies,” on page 65](#)
- ♦ [Section 6.6, “Configuring Data Synchronization,” on page 67](#)
- ♦ [Section 6.7, “Viewing Primary and Secondary Storage Capacity,” on page 72](#)
- ♦ [Section 6.8, “Using Sequential-Access Storage for Long Term Data Storage,” on page 73](#)
- ♦ [Section 6.9, “Data Deletion Policy,” on page 76](#)

6.1 Types of Data

Sentinel receives two separate but similar data streams from the Collector Managers: the event data and the raw data. The data is moved from the primary, compressed, file-based storage to a user-configured, compressed secondary storage location on a regular basis.

- ♦ [Section 6.1.1, “Raw Data,” on page 53](#)
- ♦ [Section 6.1.2, “Event Data,” on page 58](#)

6.1.1 Raw Data

The raw data files are unprocessed events that are received by the Connector and sent directly to the Sentinel message bus. This data is written to the Sentinel server. When the event is sent to the message bus, the following additional information is also sent without altering the original event:

- ♦ SHA-256 hash of the event
- ♦ Chaining indicator, which is reset to 0 whenever the Sentinel event source is restarted
- ♦ Raw Data ID (in s_RV25)
- ♦ Event source, Connector, Collector, and Collector Manager node IDs

Sentinel receives all raw data without being filtered. Because the raw data is not searched or used to generate reports, the data is not indexed.

- ♦ [“Raw Data Storage” on page 54](#)
- ♦ [“Raw Data Representation” on page 55](#)
- ♦ [“Disabling Raw Data Collection” on page 58](#)

Raw Data Storage

Sentinel compresses the raw data and stores it in partitions that are based on the time and the event source. New raw data files are created every hour.

Sentinel stores the raw data files in one of the following locations:

- ♦ Primary storage location: `<Sentinel data directory>/rawdata/online`
- ♦ Secondary storage location: `<Sentinel archive directory>/rawdata_archive`

The compressed raw data files are moved from the primary storage to the secondary storage location.

The following table describes the directory structure of the raw data in the primary storage under the installation directory:

Table 6-1 Raw Data Directory Structure

Directory Structure	Description
<code>/data</code>	The primary directory for all data storage.
<code>/data/rawdata</code>	The subdirectory where all raw data is stored.
<code>/data/rawdata/online</code>	The directory where all the raw data in the primary storage is stored.
<code>/data/rawdata/ online/EventSource UUID</code>	<p>There is one subdirectory for each event source under the <code>online</code> subdirectory. That subdirectory contains all raw data received from that event source.</p> <p>The subdirectory name is the universally unique identifier (UUID) of the event source (for example, E20D0840-1E0A-102C-9F30-000C2949BA91).</p>
<code>/data/rawdata/ online/EventSource UUID/Month</code>	<p>Data in the event source subdirectory is partitioned by month. Each month has its own subdirectory.</p> <p>The subdirectory name is in the yyyy-mm format. For example, 2009-05 indicates May 2009.</p>

Directory Structure	Description
/data/rawdata/ online/EventSource UUID/Month/1 Hour Data Files	<p>Each file in the Month directory contains data received during a specific one-hour period. Most data in the file has a time stamp that is within the one-hour period.</p> <p>The name of the file indicates the day of the month and the one-hour period that is represented.</p> <p>The filename format is dd-hhmm.extension.</p> <p><i>dd</i> is the day of the month.</p> <p><i>hh</i> is the hour of the day.</p> <p><i>mm</i> is the minute of the hour.</p> <p>The extension is either .gz. or .open.</p> <p>NOTE: Raw data files are compressed and have the extension .gz. However, when the raw data file is being written into, the raw data file appears with the extension .open.</p> <p>For example:</p> <p>A filename of 08-0000.gz indicates that the file contains compressed data received on the 8th day of the month between 12.00 a.m. and 01.00 a.m.</p> <p>A filename of 08-1300.open indicates that the file contains data received on the 8th day of the month between 01.00 p.m. and 02.00 p.m.</p>

If the raw data files are stored in the primary storage location, the full path name of the file is in the following format:

<Sentinel data directory>/rawdata/online/<event source UUID>/<Date>/<RawDataFile>

For example:

/var/opt/novell/sentinel/data/rawdata/online/A75CF6A0-4948-102D-A615-000C29A9C3DB/2010-05/24-0600.gz

In this example, */var/opt/novell/sentinel/data* is the data directory for Sentinel.

If the raw data files are stored in the secondary storage location, the full path name would be as follows:

<Sentinel archive directory>/rawdata_archive/<event source UUID>/<Date>/<RawDataFile>

For example:

/sentinel_archive_data/rawdata_archive/A75CF6A0-4948-102D-A615-000C29A9C3DB/2010-05/24-0600.gz

In this example, */sentinel_archive_data* is the secondary storage directory configured by the user.

Raw Data Representation

Each raw data event is represented as a single line in a raw data file. Each line is a JSON object with the following format:

```

{
  "EventDate": "<date>",
  "EventRecordID": " <event record uuid>",
  "RawData": "<raw data>",
  "RawDataHash": "<SHA256 hash of raw data, in hex format>",
  "EventSourceManagerID": "<uuid of event source manager>",
  "CollectorID": "<uuid of collector>",
  "EventSourceID": "<uuid of event source>",
  "ChainID": "<chain ID>",
  "ChainSequence": "<Sequence number>"
}

```

The following table describes each of the fields in the raw data event:

Table 6-2 Raw Data Representation

Field Name	Description
EventDate	<p>The date and time when Sentinel received this event and not the date and time when the event occurred.</p> <p>Example: "05/24/2010 06:15:06.676"</p>
EventRecordID	<p>The unique ID identifying the raw data record.</p> <p>Example: "595829C0-1C8F-102C-A922-000C2949BA91"</p> <p>If an event was generated as a result of parsing a raw data record, this ID is set in the event RecordID field. Because of filtering, not all raw data records result in an event.</p>
RawData	The original raw data received by the event source.
RawDataHash	<p>The SHA-256 hash of the RawData value represented as a HEX string. The hash is calculated by converting the RawData value to a UTF-8 string and then performing the hash over that string.</p> <p>To detect tampering, each raw data event is stored with a SHA-256 hash value.</p> <p>Example: cc661009e2f3dc565c0c7fe25b705219004dcd8132c0b0a7e987bfdcb55e49cf</p>
EventSourceID	<p>The UUID of the event source from which the raw data originated.</p> <p>Example: A2A0C600-1C6C-102C-A781-000C2949BA91</p>
EventSourceGroupID	<p>The UUID of the event source group (Connector) to which the event source was connected when the raw data was received.</p> <p>Example: A2A0C600-1C6C-102C-A77A-000C2949BA91</p> <p>Different raw events from the same event source can have different event source group IDs, because event sources can be moved from one Connector to another.</p>
CollectorID	<p>The UUID of the Collector that the Connector and event source were connected to when the raw data was received.</p> <p>Different raw events from the same event source can have different Collector IDs, because event sources and event source groups can be moved from one Collector to another.</p> <p>Example: A2A0C600-1C6C-102C-A779-000C2949BA91</p>

Field Name	Description
EventSourceManagerID	<p>The UUID of the Event Source Manager (Collector Manager) object where this raw data was received.</p> <p>Example: C76D2820-C395-1029-BB86-001321B5C0B3</p>
ChainID	<p>A random number that identifies a raw data chain. Whenever an event source is stopped and restarted between generation of raw data events, a new ChainID number is generated.</p> <p>To detect tampering, each raw data event is stored with a ChainID and a ChainSequence number.</p> <p>Example: 1241630654754</p>
ChainSequence	<p>A sequence number within a particular raw data chain.</p> <p>The raw data events in a given raw data chain must have an uninterrupted sequence of numbers starting with 0. In addition, all raw data events in a given raw data chain must appear sequentially in the files, with no other chains intermixed. If a raw data chain can span files, the sequence should continue uninterrupted into the file that represents every hour during which raw data was received.</p> <p>Example: 4</p> <p>If no raw data is received for the one-hour period, the file records only from the next arrival of raw data. Nonetheless, the raw data chain sequence should continue uninterrupted until a new raw data chain begins. A new raw data chain is signaled by a changed ChainID value, and a ChainSequence value of zero (0).</p>

The following examples show three raw data records:

```
{
  "EventDate": "05\24\2010 06:15:06.676",
  "EventRecordID": "A75CF6A0-4948-102D-A61C-000C29A9C3DB",
  "RawData": "Sep 22 10:22:00 testhost Message #100",
  "RawDataHash": "7003c0e0be4ddf43a3b49026a37483f59c7f839950f581ec9fde5dea43da90f5",
  "EventSourceManagerID": "C76D2820-C395-1029-BB86-001321B5C0B3",
  "CollectorID": "A75CF6A0-4948-102D-A613-000C29A9C3DB",
  "EventSourceGroupID": "A75CF6A0-4948-102D-A614-000C29A9C3DB",
  "EventSourceID": "A75CF6A0-4948-102D-A615-000C29A9C3DB",
  "ChainID": "1274696106664",
  "ChainSequence": "0"
}
{
  "EventDate": "05\24\2010 06:15:07.358",
  "EventRecordID": "A75CF6A0-4948-102D-A624-000C29A9C3DB",
  "RawData": "Sep 22 10:22:00 testhost Message #99",
  "RawDataHash": "f5681ba965144d2d22b13188767d94540b5fe57904afcee5821854bde2afca72",
  "EventSourceManagerID": "C76D2820-C395-1029-BB86-001321B5C0B3",
  "CollectorID": "A75CF6A0-4948-102D-A613-000C29A9C3DB",
  "EventSourceGroupID": "A75CF6A0-4948-102D-A614-000C29A9C3DB",
  "EventSourceID": "A75CF6A0-4948-102D-A615-000C29A9C3DB",
  "ChainID": "1274696106664",
  "ChainSequence": "0"
}
```

```

    "ChainSequence": "1"
  }
  {
    "EventDate": "05/24/2010 06:15:07.988",
    "EventRecordID": "A75CF6A0-4948-102D-A62A-000C29A9C3DB",
    "RawData": "Sep 22 10:22:00 testhost Message #98",
    "RawDataHash": "98435b5dba95633699b88d07782109876e8ceb4169d567602f2c92657118645d",
    "EventSourceManagerID": "C76D2820-C395-1029-BB86-001321B5C0B3",
    "CollectorID": "A75CF6A0-4948-102D-A613-000C29A9C3DB",
    "EventSourceGroupID": "A75CF6A0-4948-102D-A614-000C29A9C3DB",
    "EventSourceID": "A75CF6A0-4948-102D-A615-000C29A9C3DB",
    "ChainID": "1274696106664",
    "ChainSequence": "2"
  }
}

```

Disabling Raw Data Collection

By default, raw data collection is enabled for the Collector Manager on the Sentinel server. Collecting raw data can impact the performance of the server or the remote Collector Manager. Perform the following procedure on any Collection Manager where you want to disable raw data collection:

- 1 Open the `/etc/opt/novell/sentinel/config/event-router.properties` file in a text editor.
This is the default location of the file.
- 2 Change `esecurity.router.event.rawdata.send=true` to `esecurity.router.event.rawdata.send=false`.
- 3 Save the file, then restart the Collector Manager.

6.1.2 Event Data

Event data is created as a result of a Collector parsing and normalizing raw data. For more information about event processing and parsing, see [Chapter 7, “Configuring Agentless Data Collection,” on page 79](#).

Filtering rules are available on the event source, Connector, and Collector, which can selectively prevent the Collector from parsing raw data. Filtering rules avoid the overhead of parsing and normalizing data you do not need for further processing or analysis, and free up hardware resources for more important tasks. These rules do not affect the storage of the raw data. However, event data can be dropped after it is created by the parsing and normalization logic of the Collector by configuring an event routing rule to selectively drop the event data. This is useful when it is more convenient to define the rule on normalized data rather than non-normalized (raw) data. For more information, see [Chapter 10, “Configuring Event Routing Rules,” on page 119](#).

Sentinel closes the event data partitions after one day, and no more events are written to the closed partitions. Even though the duration for event data partitions is one day, a grace period of 10 minutes is given to accommodate events arriving late. You can change the grace period as necessary. For more information, see [Section 26.3, “Setting the Grace Period to Close Event Data Partitions,” on page 234](#).

By default, after the partitions are closed, Sentinel copies a compressed copy of the partition to secondary storage, but also retains the uncompressed copy on primary storage as a fast-access cache for searching. When the primary storage reaches its maximum disk usage, Sentinel deletes the copy in the primary storage and the copy in the secondary storage remains online for searching.

NOTE: However, if disk space in primary storage is at a premium, you can compress these partitions as soon as they are closed to save the disk space on the primary storage. This requires additional I/O to compress and store on the primary partition, which means that the supported EPS rate will be significantly lower. Also, searches on these partitions will be slower. Therefore, this option is only

suitable for lower EPS rates and if you want to get the most out of primary storage space. For information about compressing the storage index on primary partitions, see [Section 26.4, “Compressing the Storage Index on Primary Partition,”](#) on page 234.

The partitions are laid out as follows:

Primary storage: Open partitions + Most recent N days of closed partitions

- ♦ **Open partitions:** The partitions that new data is being written to.
- ♦ **Most recent N days of closed partitions:** As many of the most recently written closed partitions that can fit in primary storage.

Secondary storage: Most recent N days of closed partitions + The rest of the online data

- ♦ **The rest of the online data:** The older closed partitions that primary storage no longer has room to hold. This data is online (searchable), just like the data in primary storage.

Sentinel stores the primary storage partitions in the `/var/opt/novell/sentinel/data/eventdata` directory, which is on the local file system. Sentinel creates partitions based on the dates and retention policies.

A central partition index is maintained in the database that keeps track of all the existing partitions and their location.

The following table describes the directory structure under the installation directory where event data is stored:

Table 6-3 *Event Data Directory Structure*

Directory Structure	Description
<code>/data</code>	The primary directory for all data storage.
<code>/data/eventdata</code>	The subdirectory where all event data is stored.
<code>/data/eventdata/YYYYMMDD_<classid></code>	<p>A partition consists of the events for a single day (midnight-midnight UTC) within a given data retention class and is held within a subdirectory named <code>YYYYMMDD_<class-id></code>.</p> <p>YYYYMMDD: is the UTC date stamp.</p> <p><class_id>: is a UUID identifier associated with the data retention class.</p>
<code>/data/eventdata/YYYYMMDD_<class_id>/events.evt</code>	<code>events.evt</code> contains the binary event data for the partition. The format of the binary event data is stored as a Reliable Persistent Random Access Compressed Stream.
<code>/data/eventdata/YYYYMMDD_<class_id>/index</code>	The index directory contains the Lucene index for the partition.

6.2 Configuring Secondary Storage Locations

All closed event data files are copied from the primary storage location to the secondary storage location. The original files are retained on primary storage to facilitate faster searches. However, if the primary storage disk space usage nears a user-defined threshold, duplicate data files on the primary storage area are deleted from the primary storage and remain only on secondary storage.

- ♦ [Section 6.2.1, “Supported Storage Options,” on page 60](#)
- ♦ [Section 6.2.2, “Types of Secondary Storage,” on page 61](#)
- ♦ [Section 6.2.3, “Configuring Secondary Storage,” on page 61](#)
- ♦ [Section 6.2.4, “Changing the Secondary Storage Location,” on page 63](#)

6.2.1 Supported Storage Options

Sentinel supports the following types of storage options:

- ♦ **Primary Storage or SAN:** The primary storage or Storage Area Network (SAN) option includes storage that is attached directly to the Sentinel computer. This option provides the best combination of performance, security, and reliability.
- ♦ **CIFS:** The Common Internet File System (CIFS) is a native Windows protocol. It is also known as the Server Message Block (SMB) protocol in later implementations. The latest implementation from Microsoft is referred to as SMB 2.
- ♦ **NFS:** The NFS protocol requires significant configuration to optimize performance and security, and it is recommended only if you already have a well-established NFS infrastructure in your environment.

If the secondary storage is an NFS server, additional configuration is necessary to ensure that the Sentinel server has the necessary permissions. For more information, see [“Exporting the Secondary Storage Volume” on page 62](#).

WARNING: Only one Sentinel server should be configured to use a particular secondary storage directory (remote share). Configuring the same secondary storage location across multiple Sentinel servers might cause system failure.

The primary storage must use a different partition than the partition that is used for the secondary storage.

- ♦ The system monitors the disk usage of both primary storage and secondary storage, freeing space on primary storage when it fills up. If both storage locations share the same underlying file system partition, the way in which the partition usage changes as a result of deleting data confuses the system and could result in undesirable behavior.
- ♦ The event data is first copied to secondary storage rather than moved, because there is an assumption that these are two different disk partitions. If they are in same disk partition instead of being on the different disk partition, the storage usage monitoring is confused by how the usage is changing and could result in undesirable behavior.

6.2.2 Types of Secondary Storage

You can enable and configure secondary storage for raw data and event data stored on the Sentinel server.

- ♦ [“Raw Data Storage” on page 61](#)
- ♦ [“Event Data Storage” on page 61](#)

Raw Data Storage

Raw data files are compressed and have the .gz extension. When the data is currently being written into, the raw data file appears with the .open extension.

If secondary storage is configured and enabled, Sentinel copies the compressed raw data files to the configured secondary storage location every 15 minutes.

Event Data Storage

If secondary storage is enabled, Sentinel moves the closed files to secondary storage every midnight UTC and also whenever the server starts. These files are compressed in the primary storage location, but the file indexes are compressed before moving to the secondary storage. If the secondary storage location is not configured or if there is any problem while moving the closed files, Sentinel attempts to move the files to secondary storage every 60 seconds until it succeeds.

6.2.3 Configuring Secondary Storage

The NFS, CIFS/SMB, and SAN must be configured so that Sentinel has read and write permissions.

For CIFS/SMB and NFS, if multiple Sentinel instances are moving the closed partitions to the same secondary storage location, ensure that each Sentinel instance has its own unique directory on that secondary storage location.

- ♦ [“Configuring a SAN/Local Directory as a Secondary Storage Location” on page 61](#)
- ♦ [“Configuring an CIFS/SMB Server as a Secondary Storage Location” on page 62](#)
- ♦ [“Configuring an NFS Server as a Secondary Storage Location” on page 62](#)

Configuring a SAN/Local Directory as a Secondary Storage Location

Configuring a SAN/Local directory as a secondary storage location is the preferable configuration for best performance, security, and reliability.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Configuration**.
- 3 From the Data Storage Location section, select **SAN (locally mounted)** as the secondary storage location.
- 4 In the **Location** field, specify the local directory path or the location on which the storage area network (SAN) is mounted.
The SAN partition must be manually mounted before the location is specified.
- 5 Click **Test** to check if the write permissions for the specified location are available.
- 6 Click **Save** to configure the specified secondary storage location.

Configuring an CIFS/SMB Server as a Secondary Storage Location

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Configuration**.
- 3 In the Data Storage Location section, select **CIFS**.
- 4 Specify the following information:
 - Server:** Specify the IP address or hostname of the computer where the CIFS server, also known as the SMB server, is configured.
 - Share:** Specify the share name of the SMB or CIFS server. The mounted shares are unmounted when the server stops and are mounted again when the server starts. If the configured share unmounts, the Sentinel server detects this and mounts it again.
 - Username:** Specify the user name (if one is assigned) to access the share.
 - Password:** Specify the password (if one is assigned) to access the share.
 - Mount Options:** Specifies the options that are used while mounting the secondary storage location of the SMB or the CIFS server.

You can specify new mount options. For more information about the available NFS mount options, see the [mount.cifs \(8\) - Linux man page \(http://linux.die.net/man/8/mount.cifs\)](http://linux.die.net/man/8/mount.cifs).

The default mount options are `file_mode=0660,dir_mode=0770`.
- 5 (Optional) Click **Restore Defaults** to restore the default mount options.
- 6 Click **Test** to mount the SMB or CIFS server and to check the write permissions on the server.
- 7 Click **Save** to configure the specified secondary storage location.

Configuring an NFS Server as a Secondary Storage Location

The NFS protocol requires significant configuration to improve performance and security. NFS server is recommended only when you already have a well-established NFS infrastructure in your environment.

- ♦ “[Exporting the Secondary Storage Volume](#)” on page 62
- ♦ “[Configuring NFS as a Secondary Storage Location](#)” on page 63

Exporting the Secondary Storage Volume

The NFS server needs to export (share) the secondary storage volume to the Sentinel server so that the secondary storage is readable by the `root` user on the Sentinel server. The settings described in this section indicate one of the methods to achieve this readability:

- ♦ The NFS server must have a user and a group with a UID and a GID that correspond to the `novell` user and group on the Sentinel server.

In the following examples, the user on the NFS server is `novell` with `UID=5555` and the group is `novell` with `GID=5555`. The Sentinel server has the hostname `log-manager-server`, which can be resolved by the NFS server.

- ♦ The `novell` user and group must own the secondary storage destination directory on the NFS server. In the following examples, the secondary storage destination is `/archive`.
- ♦ Map the `root` user on the Sentinel server to the `novell` user and group on the NFS server.

- ♦ **Linux:** Add the following line to the `/etc/exports` file:

```
/archive sentinel=server(rw,root_squash,anonuid=5555,anongid=5555)
```

- ♦ **Solaris:** Add the following line to the `/etc/dfs/dfstab` file:

```
/usr/bin/share -F nfs -o sec=sys,rw=sentinel-server,anon=5555 -d "/archive"
/archive
```

- ♦ **HP-UX:** Add the following line to the `/etc/exports` file:

```
archive -access=sentinel-server,anon=5555
```

You can speed up the archiving process by adding the `async` option. However, this might increase the risk of lost or corrupt secondary storage if the NFS server crashes. The following examples demonstrate using the `async` option:

- ♦ **Linux:** Add the following line to the `/etc/exports` file:

```
/archive sentinel=server(rw,root_squash,anonuid=5555,anongid=5555),async
```

- ♦ **HP-UX:** Add the following line to the `/etc/exports` file:

```
archive -access=sentinel-server,anon=5555,async
```

For information about security recommendations for NFS, see [Chapter 3, “Security Considerations,” on page 21](#).

Configuring NFS as a Secondary Storage Location

Configure the secondary storage as follows:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.

- 2 Click **Storage > Configuration**.

- 3 In the Data Storage Location section, select the **NFS** option.

- 4 Specify the following information:

Server: Specify the IP address or hostname of the computer where the NFS server is configured.

Share: Specify the share name of the NFS server.

The mounted shares are unmounted when the server stops and are mounted again when the server starts. If the configured share unmounts, the Sentinel server detects this and mounts it again.

Mount Options: Specifies the options that are used while mounting the secondary storage location of the NFS server.

You can also specify new mount options. For more information about the available NFS mount options, see [NFS \(5\) Linux Programmer’s Manual \(http://unixhelp.ed.ac.uk/CGI/man-cgi?nfs+5\)](http://unixhelp.ed.ac.uk/CGI/man-cgi?nfs+5).

The default mount options are `soft,proto=tcp,retrans=1,timeo=60`.

- 5 (Optional) Click **Restore Defaults** to restore the default mount options.

- 6 Click **Test** to verify the configuration of the NFS server and to check the write permissions on the server.

This procedure tests a subset of all the settings that are necessary for the NFS server and client.

- 7 Click **Save** to configure the specified secondary storage location.

6.2.4 Changing the Secondary Storage Location

- 1 Log in to the Sentinel Web interface as a user in the administrator role.

- 2 Click **Storage > Configuration**.

- 3 In the Data Storage Location section, select **Change Location**. The **Change Location** option is displayed only if the secondary storage location is configured.

4 Click **Change Location**.

5 Select the option to disable data collection.

You can select this option to avoid filling the primary storage before Sentinel moves the data to the new location. If this option is not selected and if the primary storage is filled before the new data storage location is configured, Sentinel deletes the oldest data to make space for the incoming data.

6 Configure the new data storage location.

For more information about configuring the NIFS or SMB/CIFS or primary/SAN secondary storage locations, see [“Configuring Secondary Storage” on page 61](#).

7 Click **Save** to save the changes and configure the new secondary storage location.

8 Manually copy the files from the old secondary storage location to the new secondary storage location.

9 After copying the files, select **Copy Done** to start data storage at the new location.

6.3 Configuring Disk Space Usage

If secondary storage is enabled, Sentinel copies the data to the secondary storage location after two days, and a local copy remains until space is available. Sentinel moves the raw data to the secondary storage location after approximately one hour.

1 Log in to the Sentinel Web interface as a user in the administrator role.

2 Click **Storage > Configuration**.

In the **Disk Space Usage** section, the **Primary storage size** field displays the total storage size currently used by Sentinel.

3 Specify the primary storage utilization value:

- ♦ Specify a value to start deleting the data storage from primary storage when the specified value is reached.
- ♦ Specify a value to stop deleting the data storage from primary storage when the specified value is reached.

These are the settings at which Sentinel starts (and stops) deleting the duplicate data files that are in primary storage. Sentinel keeps the copies in the primary storage, until this disk usage threshold is reached, to speed up searches and reports. When the threshold is reached and files are deleted, the secondary storage becomes the sole location for the data.

The **Secondary storage size** field specifies the value of the secondary storage space.

4 Specify the maximum archive size to be used as part of the total available archive size.

6.4 Verifying and Downloading Raw Data Files

The raw data files for each event source are compressed and moved to secondary storage every hour and the file hash is computed for secondary storage files. The file hash is used to check the integrity of the files in the secondary storage.

1 Log in to the Sentinel Web interface as a user in the administrator role.

2 Click **Storage > Raw Data**.

3 In the **Event source hierarchy** field, select the desired Collector and Connector combination from the drop-down list.

- 4 In the **Event Source** field, select the event source from the drop-down list.
The **Event Source** field displays the list of associated event sources (hostnames or IP addresses) after the **Event source hierarchy** field is populated.
- 5 In the table, click **Select All** to select all the files in the table.
or
Select each file separately.
The table displays the list of primary and secondary storage raw data files for the selected event source. The **Verify Integrity** and **Download** options are enabled only when you select a file from the table.
- 6 Click **Verify Integrity** to verify the integrity of the selected files in the secondary storage by comparing the hash values for the selected files in the secondary storage.
Sentinel computes the hash and updates the database for the files in the secondary storage, but not for the local raw data files. Because the raw data files are updated until they are moved to secondary storage, the hash value cannot be computed or updated for these files. It is not possible to check the integrity of the local raw data files.
- 7 Select the raw data file, click **Download** to download the selected secondary storage and local raw data files.
The selected files are downloaded in the form of a ZIP file that contains a .csv (comma separated values) file. If the secondary storage files are selected, the ZIP file also contains a hash file corresponding to each of the secondary storage files downloaded.
Sentinel uses the SHA-256 algorithm to generate the file hash. The generated hash is Base64 encoded.
- 8 Select **Save File** and click **OK**.

6.5 Configuring Data Retention Policies

The data retention policies control when data should be deleted from the system. A retention policy contains a filter that is used to identify the events for which the retention policy applies and the minimum and maximum number of days these events should be kept in the system.

You can configure one or more data retention policies to control the duration for which specific types of events are retained in Sentinel. Except for the Raw Data Retention policy, all of the configured policies apply to the event data.

The configured retention policies are displayed in the data retention policy table. By default, the data retention policy table is refreshed every 30 seconds to reflect the changes made by multiple administrators.

- ♦ [Section 6.5.1, “Raw Data Retention Policy,” on page 66](#)
- ♦ [Section 6.5.2, “Event Data Retention Policies,” on page 66](#)
- ♦ [Section 6.5.3, “Rules for Applying a Retention Policy,” on page 67](#)

6.5.1 Raw Data Retention Policy

The raw data retention policy controls the duration for which the raw data is kept in the system before it is deleted. The raw data retention policy cannot be deleted or disabled. However, you can modify the **Keep at most** and **Keep at Least** values, which determine the maximum and minimum number of days to keep the raw data file.

The process to delete raw data files runs every time the server is started, every hour because that is when the raw data files are closed, and whenever the **Keep at most** value is changed. All the files exceeding the retention time are removed permanently from the primary and secondary storage locations.

6.5.2 Event Data Retention Policies

The event data retention policies control the duration for which different types of event data are kept in the system before being deleted.

To create a data retention policy:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Configuration**.
- 3 In the Data Retention section, click **Create**.
- 4 Use the following information to create the data retention policy.

Policy name: Specify a name for the retention policy.

The policy name must be unique and must contain alphanumeric characters.

Filter: Click the **Filter** icon to select a saved or shared filter.

or

Specify a filter value. The filter value uses the same syntax as searches.

For example, assume that the **Filter** field contains a filter such as *sev:[3 TO 5] AND (evt:"SyslogNICListener")*. This filter value matches all the events with a severity of 3, 4 or 5 and event name SyslogNICListener.

For more information, see "[Configuring Filters](#)" in the *NetIQ Sentinel User Guide*.

For example, use *sev:[0 TO 1]* to define a retention policy that applies to all events with a severity of 0 or 1.

Keep at least: Specify the minimum number of days to retain the events in the system. The value must be a valid positive integer.

Keep at most: (Optional) Specify the maximum number of days for which the events should be retained in the system.

The value must be a valid positive integer and must be greater than or equal to the **Keep at least** value. If no value is specified, the system retains the events in the system until the space is available.

- 5 Click **Save**. The newly created policy is displayed in the data retention table.

The table also contains the following additional columns:

Size: Displays the amount of space used to store the events for each retention policy.

Events: Displays the number of events for the selected retention policy.

The policies are sorted in alphabetical order by policy name. The default retention policy is always shown as the last policy in the list.

For more information, see [Section 6.9, “Data Deletion Policy,” on page 76](#).

6.5.3 Rules for Applying a Retention Policy

An event could match the filter criteria of multiple data retention policies.

To determine which data retention policy will apply to an event and, therefore, how long an event will be retained before deleting it from the primary and secondary data stores, apply the following rules:

1. If an event meets the criteria of only one data retention policy filter, that data retention policy is applied to the event.
2. If an event does not meet the criteria for any of the data retention policies, the default data retention policy is applied to that event.
3. If an event meets the criteria for more than one of the data retention policies, the following guidelines are used to determine which data retention policy should be applied:
 - ♦ If the maximum retention period of a policy is shorter than the others, that policy is applied. (If the maximum retention period is not specified for a policy, the policy is considered to have a long maximum retention period.)
 - ♦ If multiple matching policies have the same shortest maximum retention period, the policy with the longest minimum retention period is applied.
 - ♦ If multiple matching policies have the same shortest maximum retention period and the same longest minimum retention period, the system arbitrarily applies one of the policies.

6.6 Configuring Data Synchronization

Sentinel provides the ability to synchronize data to an external database, so that you can use third-party or custom reporting systems to search the data in the external database with more advanced tools than what are provided in Sentinel.

- ♦ [Section 6.6.1, “Overview,” on page 67](#)
- ♦ [Section 6.6.2, “Creating a Data Synchronization Policy,” on page 70](#)
- ♦ [Section 6.6.3, “Managing Data Synchronization,” on page 72](#)

6.6.1 Overview

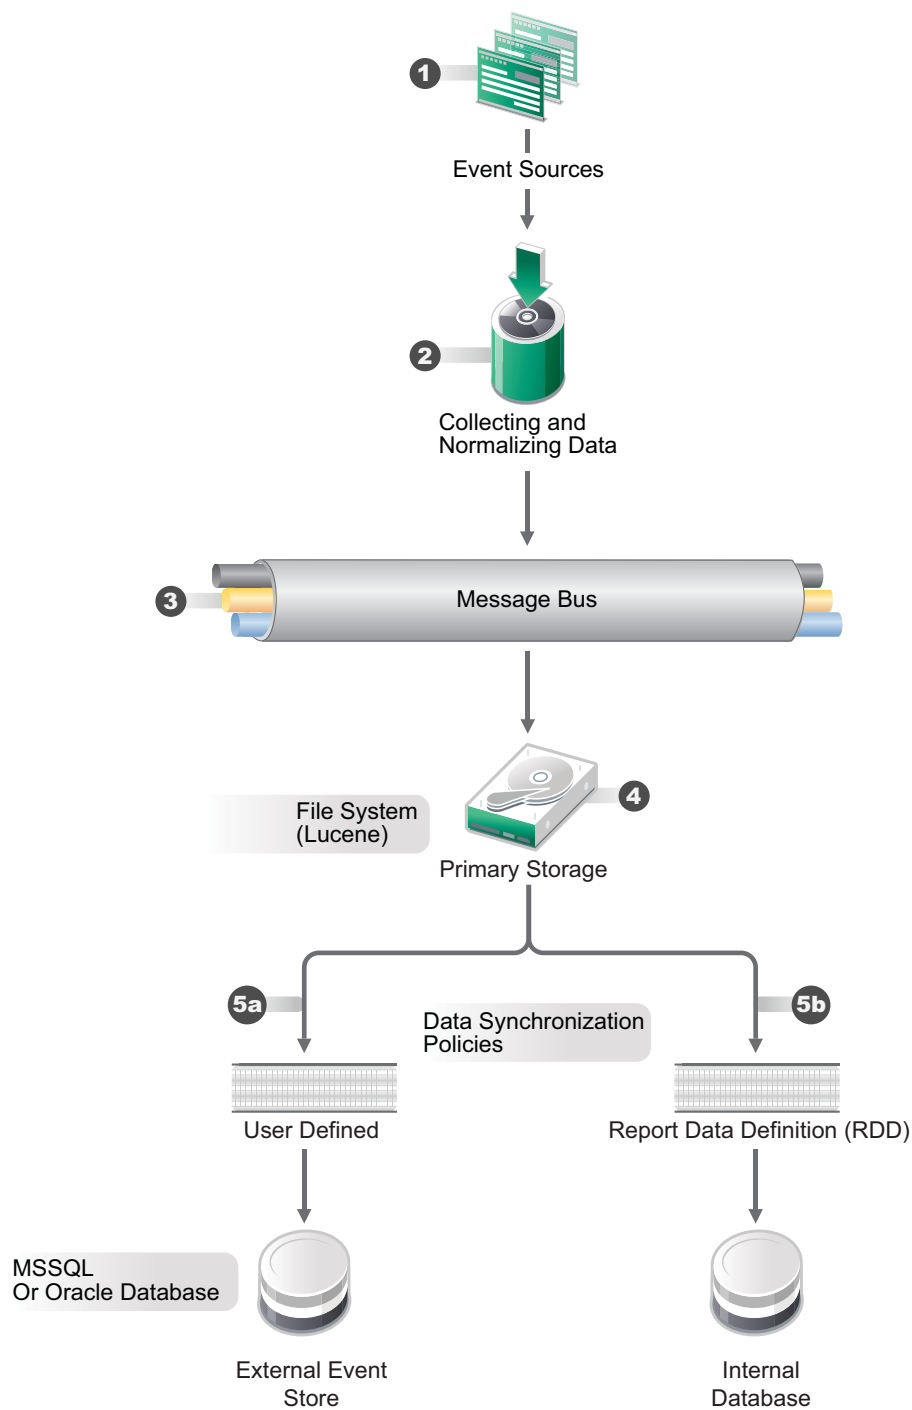
Sentinel can store the data in an external database by synchronizing a subset of the data that Sentinel gathers.

Sentinel uses the following process to synchronize data:

1. Sentinel gathers the events from the Event Sources through the Connectors.
2. Sentinel uses Collectors to normalize the event data.
3. The normalized event data is then sent to the Sentinel message bus.
4. The event data is then stored and indexed in the file system in the primary storage.

5. The data synchronization policies allow events in the primary storage to be copied and stored in PostgreSQL and external SQL databases.
 - a. User-defined data synchronization policies synchronize the filtered event data to an external SQL database. The external SQL database can be Microsoft SQL 2008, Microsoft SQL 2012, Oracle 11G, DB2, or Sybase.
 - b. Report Data Definitions (RDD) generate system data synchronization policies that are used to copy event data into tables in the internal PostgreSQL database. These data synchronization policies cannot be edited or deleted. Reports that rely on an RDD will search internal database tables for events instead of the primary storage. These kinds of reports search internal tables instead of the event store because they utilize more complex SQL SELECT statements that need to join event data to the data in other tables in the internal database.

Figure 6-1 Data Synchronization



Sentinel allows you to partition tables if they are in the internal PostgreSQL database. When you choose to partition a table in the internal PostgreSQL database, a new table partition is created for each days worth of data.

Partitions are only used with RDD data sync policies. Partitioning has advantages and disadvantages:

Advantages

- ♦ If a retention period is in force, old data can be deleted quickly. When data has aged, it is much quicker to drop a partition than it is to delete individual table records.
- ♦ Reports that query on the event time field might be quicker, because it is only necessary to search the partitions that have the specified event times.

Disadvantages

- ♦ Reports that do not query on event time might be slower where there are multiple of partitions, because every partition must be searched.
- ♦ Each partition causes one or more schema items to be created and managed by the database system. If there is no retention period, the number of partitions just keeps growing.

6.6.2 Creating a Data Synchronization Policy

When Sentinel syncs data to an external database, it is not as fast when it writes to the file system. Therefore, you need to ensure that you use filters on the data sync policies to synchronize only the most important data. Consider the following factors based on your business needs for data sync policies:

- ♦ The CPU, RAM, and disk capacity of the Sentinel system
- ♦ Number of EPS scaled per system
- ♦ Number of searches and reports running on a Sentinel system
- ♦ Filters added to the data sync policies

To create a data synchronization policy:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Data Synchronization**.
- 3 Click **Create** to create a new data synchronization policy.
- 4 Use the following information to create the data synchronization policy:

Filter query: Select a saved filter to use in the data synchronization policy.

This filter determines which events are stored in the external database. For more information, see “[Configuring Filters](#)” in the *NetIQ Sentinel User Guide*.

Policy name: Specify a name for the data synchronization policy.

Retention period: Specify how many days to retain the events in the external database.

Start data synchronization time: Specify when to start synchronizing events to the external database.

Batch size: Specify how many events are sent to the external database at once.

Sleep period: Specify the length of time that the data synchronization process sleeps before checking to see if there are more events to process.

Schedule: Select when the data is synchronized to the external database.

- ♦ **All the time:** Synchronizes events to the external database constantly.
- ♦ **Custom:** Allows you to configure specific time periods to perform data synchronization so that it does not occur when the system is busy.

If you select **Custom**, specify the following information to set the custom synchronization time:

- ♦ **Day of the Week:** Select the day of the week, or select **Everyday**.
- ♦ **Start time:** Specify the time to start the synchronization process. You can enter 24:00 hour time and it is converted to 12:00 hour time.
- ♦ **Duration:** Specify the synchronization period in minutes.

If you do not see the data in the database tables immediately, you need to wait for the next synchronization cycle.

- 5 Use the following information to define the connection to the external database:

Database type: Select the type of external database.

Host name: Specify the host name of the server where the external database is installed.

Port: Specify the port used to connect to the external database.

User name: Specify the name of the user that authenticates to the external database.

Password: Specify the password of the database user.

Database: Specify a unique name for the external database.

Field Mapping: Allows you to map fields in the event to fields in the external database.

- ♦ [“Creating a Table for Event Data Synchronization” on page 71](#)
- ♦ [“Using an Existing Table for Event Data Synchronization” on page 72](#)

- 6 Click **Save** to create the data synchronization policy.

Creating a Table for Event Data Synchronization

- 1 Complete [Step 1](#) through [Step 5](#) in [Section 6.6.2, “Creating a Data Synchronization Policy,” on page 70](#).

- 2 Click **Field Mapping**.

- 3 Select **Create table**.

- 4 Use the following information to create the table:

Table name: Specify a name for the table.

Table Space (Optional): Specify a tablespace for the table.

Index Space (Optional): Specify a tablespace for the index.

Summarize Events: Select this option if you want a summary of events during a specific period.

Summary Period (Minutes): If you selected **Summarize Events**, you must specify the amount of time in minutes to summarize events.

- 5 Map the fields in the table to the desired fields.

- 6 Click **Create Table**.

- 7 Click **Save**.

Using an Existing Table for Event Data Synchronization

NOTE: Starting from Sentinel 7.0.1 version, the InitiatorServiceName (sp) and TargetServiceName (dp) fields size is increased from 32 to 256 characters to accommodate more characters in these fields. If you have created a Data Sync policy in Sentinel 7.0 that synchronizes either or both of the event fields, you need to modify the target column size in the external database table to reflect the increased size of the fields.

- 1 Complete [Step 1](#) through [Step 5](#) in [Section 6.6.2, “Creating a Data Synchronization Policy,”](#) on [page 70](#).
- 2 Click **Field Mapping**.
- 3 Select **Select existing table**.
- 4 Browse to select an existing table you want to use, then click **OK**.
- 5 (Optional) Select the **Summarize Events** option if you want a summary of events during a specific period.
- 6 (Optional) If you selected **Summarize Events**, specify the amount of time in minutes to summarize events.
- 7 Change the field mappings for the desired fields.
- 8 Click **Save**.

6.6.3 Managing Data Synchronization

You can edit, delete, and view the status of each data synchronization policy you create on the Data Synchronization page. If your policy is a custom synchronization policy and you perform a resynchronization, the data synchronizes during the next synchronization cycle.

6.7 Viewing Primary and Secondary Storage Capacity

The Health page displays primary and secondary data capacity. For more information about configuring secondary storage, see [Section 6.2, “Configuring Secondary Storage Locations,”](#) on [page 60](#).

To view the primary storage and secondary storage capacity:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Health**.

The Health page of Sentinel also displays the current storage capacity and also forecasts the storage capacity for both primary and secondary storage.

WARNING: When the system is running out of primary disk storage space, a warning message is displayed and a system audit event is logged. To avoid data loss, you must increase the primary storage space.

6.8 Using Sequential-Access Storage for Long Term Data Storage

Sentinel requires data to be on a storage system that supports random access, such as data on your typical hard drive. It does not support directly interfacing with the data stored on tape.

You can search the raw data by using tools such as `egrep` or a text editor, but this search might not be sufficient for your requirements. The search mechanism provided by Sentinel on event data is more powerful than these tools.

The high-level approach to configure Sentinel is to retain data for a longer duration so you can perform searches and run reports on the data you regularly need to access, and to copy the data to tape before Sentinel deletes it. To search or run reports on data that was copied to tape, but deleted from Sentinel, copy the data from the tape back to Sentinel.

- ♦ [Section 6.8.1, “Determining What Data You Need to Copy to Tape,” on page 73](#)
- ♦ [Section 6.8.2, “Backing Up Data,” on page 73](#)
- ♦ [Section 6.8.3, “Configuring Storage Utilization,” on page 74](#)
- ♦ [Section 6.8.4, “Configuring Data Retention,” on page 74](#)
- ♦ [Section 6.8.5, “Copying Data to Tape,” on page 74](#)
- ♦ [Section 6.8.6, “Restoring Data,” on page 75](#)

6.8.1 Determining What Data You Need to Copy to Tape

There are two types of data in Sentinel:

- ♦ **Raw Data:** For more information on raw data, see [Section 6.1.1, “Raw Data,” on page 53](#).
- ♦ **Event Data:** For more information on event data, see [Section 6.1.2, “Event Data,” on page 58](#).

If you want to perform searches or reports on the data, copy both the raw data and the event data to tape so that you can copy both sets of data back into Sentinel when the data is needed. If you want to store data only to comply with legal requirements, copy only the raw data to the tape.

6.8.2 Backing Up Data

Events should be moved to secondary storage regularly. The following types of data can be backed up in Sentinel:

Configuration Data: This option includes non-event or raw data backup. It is faster because it contains small amount of data, including all the installation directories except the data directory.

Data: This option backs up all the data in the primary storage and secondary storage directories. This option takes a longer time to finish.

Secondary storage directories can be located on a remote computer.

Best practices for data backup include the following:

- ♦ Periodically export all the Event Source Management configurations and save them. When the environment is relatively stable, you can generate a full Event Source Management export including the entire tree of the Event Source Management components. This action captures the plug-ins and the configuration of each node. You must back up the resulting `.zip` file and move it to secondary storage.

If changes such as updating plug-ins or adding nodes are made to Event Source Management later, you must export the configuration and save it again.

- ♦ Back up the entire installation directory so there is no risk of manual mistakes and the process is quicker.

6.8.3 Configuring Storage Utilization

You should configure primary and secondary storage space to store data before the data is deleted from the Sentinel server. While configuring the storage space, ensure that your storage system is not 100% utilized to avoid undesirable behaviors such as data corruption. Additionally, you should also have additional space in your secondary storage to copy data from tape back into Sentinel. You can do this by decreasing the archive utilization setting.

6.8.4 Configuring Data Retention

You can configure the duration for the data to remain on the disk before it is deleted. If your hard drive storage space is not sufficient to store data long enough to meet your legal requirements, you can use tape storage to store data beyond the specified duration.

You must configure data retention policies so that the data that you want to search and report is retained within the Sentinel server until you no longer need it. Additionally, a data retention policy should ensure that Sentinel is not prematurely deleting the data because of storage utilization limits. If the storage utilization limit is exceeded and you notice that the data is being prematurely deleted, change the data retention policy to expand the data storage space.

6.8.5 Copying Data to Tape

You can set up a process to copy raw data and event data to tape, depending on the data that you need.

The following sections describe how each type of data is stored in Sentinel so that you can set up copy operations to copy the data out of Sentinel onto tape.

- ♦ [“Copying Raw Data to Tape” on page 74](#)
- ♦ [“Copying Event Data to Tape” on page 75](#)

Copying Raw Data to Tape

Raw data partitions are individual files. They are created every hour. Raw data files are compressed and have the .gz extension.

The directory hierarchy in which the raw data files are placed is organized by the event source and the date of the raw data. You can use this hierarchy to periodically copy a batch of raw data files to tape. For more information on raw data directory hierarchy, see [Table 6-1, “Raw Data Directory Structure,” on page 54](#).

You cannot copy files that are in the process of being compressed. You must wait until the raw data files are compressed and moved to secondary storage before copying them to tape. The presence of a .log file with the same name as the zip file indicates that the file is still in the process of being compressed. You must also ensure that the raw data files are copied to the tape before the interval configured in the Raw Data Retention policy expires so that the data is not lost.

Copying Event Data to Tape

Event data partitions are created every 24 hours. Event data is stored in the `data/eventdata` directory with subdirectory names prefixed with the year, month, and day when the partition was created (yyyymmdd). For example, the path to a complete event data partition, relative to the installation directory, is `data/eventdata/20090101_408E7E50-C02E-4325-B7C5-2B9FE4853476`. You can use this hierarchy to know when a partition is closed. Subdirectories whose date is at least 48 hours old should be in the closed state.

For more information about the event data directory hierarchy, see [Table 6-3, “Event Data Directory Structure,” on page 59](#).

You should wait until event data partitions have been copied to secondary storage before copying them to tape. Before you copy, ensure that the directory is not currently being copied from primary storage. To do this, see if there is a primary storage directory partition of the same name. If the corresponding primary storage directory partition is not present, the secondary storage directory partition is not being copied. If the corresponding primary storage directory partition is still present, sure that all of the files in the primary storage directory partition are also in the secondary storage directory partition and that they are all of the same size. If they are all present and of the same size, it is highly likely that they are not currently being copied.

6.8.6 Restoring Data

The event data restoration feature enables you to restore old, lost, or deleted event data. You can also restore the data from other systems. You can select and restore the event partitions in the Sentinel Web interface. You can also control when these restored event partitions expire.

NOTE: The Data Restoration feature is a licensed feature. This feature is not available with the free or trial licenses. For more information, see [“Understanding License Information”](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

- ♦ [“Enabling Event Data for Restoration” on page 75](#)
- ♦ [“Viewing Event Data Available for Restoration” on page 75](#)
- ♦ [“Restoring Event Data” on page 76](#)
- ♦ [“Configuring Restored Event Data to Expire” on page 76](#)

Enabling Event Data for Restoration

To enable event data for restoration, you must copy the event directories that you want to restore to one of the following locations:

- ♦ The primary storage data directory on the Sentinel server. For example, `/var/opt/novell/sentinel/data/eventdata`
- ♦ The `eventdata_archive` directory of the secondary storage directory on the Sentinel server.

Viewing Event Data Available for Restoration

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Storage > Configuration**.
The Data Restoration section does not initially display any data.
- 3 Click **Find Data** to search and display all event data partitions available for restoration.

The Data Restoration table chronologically lists all the event data that can be restored. The table displays the date of the event data, the name of event directory, and the location. The **Location** column indicates whether the event directory was found in the primary storage directory of Sentinel or in the configured secondary storage directory.

- 4 Continue with [“Restoring Event Data” on page 76](#) to restore the event data.

Restoring Event Data

- 1 Select the check box in the **Restore** column next to the partition that you want to restore.
The **Restore Data** button is enabled when the Data Restoration section is populated with the restorable data.
- 2 Click **Restore Data** to restore the selected partitions.
The selected events are moved to the **Restored Data** section. It might take approximately 30 seconds for the **Restored Data** section to reflect the restored event partitions.
- 3 (Optional) Click **Refresh** to search for more restorable data.
- 4 To configure the restored event data to expire according to data retention policy, continue with [“Configuring Restored Event Data to Expire” on page 76](#).

NOTE: You cannot restore the secondary storage data if the novell user ID (UID) and the group ID (GID) are not the same on both the source (server that has the secondary storage data) and destination (server where the secondary storage data is being restored). In such a scenario, you need to unsquash and squash the squash file system.

Configuring Restored Event Data to Expire

The restored partitions do not expire by default, according to any data retention policy checks. To enable the restored partitions to return to the normal state and also to allow them to expire according to the data retention policy, select **Set to Expire** for data that you want to expire according to the data retention policy, then click **Apply**.

The restored partitions that are set to expire are removed from the Restored Data table and returned to normal processing.

It might take about 30 seconds for the Restored Data table to reflect the changes.

6.9 Data Deletion Policy

Sentinel deletes the data types in their listed order until the required space is available.

1. All partitions (both primary storage and secondary storage) are deleted as soon as the **keep at most** time limit of their retention policy is reached.
2. Partitions that are successfully moved to secondary storage. The oldest is moved first until none are left or until the desired amount of space is available.
3. Partitions that are not yet moved to secondary storage, but have reached their retention policy's **keep at most** time limit. The partitions that are close to the **keep at most** limit are deleted first, until none are left or until the desired amount of space is available.

If at least half of the desired space is not yet free, partitions are deleted early, on the assumption that incoming data is more important than old data.

4. Partitions that are not yet moved to secondary storage that have reached their policy's **keep at most** time limit. The partitions that are close to the **keep at most** limit are moved first, until none are left or at least half of the desired amount of space is available, but the current UTC day partitions are not deleted.

7 Configuring Agentless Data Collection

Sentinel can collect data from a wide range of event sources, such as intrusion detection systems, firewalls, operating systems, routers, databases, switches, mainframes, antivirus applications, and Novell applications. A modular architecture divides the task of protocol-level connections (Connectors) and the parsing logic (Collectors) for specific event sources.

Sentinel supports a wide variety of Connectors and also includes a variety of Collectors. The configuration required to integrate a new event source with Sentinel varies, depending on the type of event source and the communication method selected.

You should review the Collector and Connector documentation for any new event source integration to ensure that all available features are enabled.

The configuration required to integrate a new event source with Sentinel varies depending on the type of event source and the communication method selected.

- ♦ [Section 7.1, “Before You Begin,” on page 79](#)
- ♦ [Section 7.2, “Configuring Data Collection for Syslog Event Sources,” on page 80](#)
- ♦ [Section 7.3, “Configuring Data Collection for the Novell Audit Server,” on page 83](#)
- ♦ [Section 7.4, “Configuring Data Collection for Other Event Sources,” on page 87](#)
- ♦ [Section 7.5, “Managing Event Sources,” on page 105](#)

7.1 Before You Begin

- ♦ To download the new and updated Collector and Connector plug-ins, and for documentation about Collectors and Connectors, see the [Sentinel Plug-ins Web site](#).
- ♦ For more information about editing Collectors that are already included in Sentinel, see the [Sentinel Plug-ins SDK Web site](#).

NOTE: Sentinel 7.0 and later do not support legacy Collectors. When you run a legacy Collector, an error is displayed in the **Status** field of the Collector and the Collector does not start. You must use an equivalent JavaScript Collector, then move the legacy Collector event source nodes to the newly added JavaScript Collector.

7.2 Configuring Data Collection for Syslog Event Sources

Sentinel is preconfigured to accept Syslog data from Syslog event sources that send data over TCP (port 1468), UDP (port 1514), or SSL (port 1443). You can also configure Sentinel to listen on additional ports. To collect Syslog data, configure your Syslog event sources to send data to one of these ports.

The following sections describe how you can configure the event sources to send data to Sentinel and how you can configure new Syslog ports to receive data:

- ♦ [Section 7.2.1, “Parsing Logic for Syslog Messages,” on page 80](#)
- ♦ [Section 7.2.2, “Configuring Syslog Servers,” on page 80](#)
- ♦ [Section 7.2.3, “Configuring Client Authentication for the SSL Syslog Server,” on page 81](#)

7.2.1 Parsing Logic for Syslog Messages

Sentinel can receive, store, and search against events from any Syslog source. If it recognizes the data source, Sentinel automatically chooses the most suitable Collector to parse the data, parses the data into events, and stores the data in the configured secondary storage location.

You can filter the collected data to drop any unwanted events. Messages from recognized data sources are parsed into fields such as `target IP address` and `source username`. Messages from unrecognized data sources are placed in a single field for storage, searching, and reporting.

The Generic Event Collector collects and processes data from unrecognized event sources that have suitable Connectors. If the data was generated by a supported event source, the Generic Event Collector analyzes the received data and attempts to parse the information. If the Generic Event Collector does not understand the message, it does minimal parsing and places the text in the **Message** (Msg) field.

7.2.2 Configuring Syslog Servers

When you point your Syslog event sources to Sentinel, Sentinel automatically creates an event source entry to track data received from the event source. An entry is created for each unique IP address or hostname that appears in the header portion of the Syslog messages. This entry enables you to identify the computers generating the Syslog messages, regardless of whether they are being aggregated by a Syslog relay or not. It also enables you to manage how the data is processed.

To add or remove Syslog servers, use the Event Source Management interface. For more information, see [“Configuring Data Collection for Other Event Sources” on page 87](#).

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection > Event Source Servers**.
- 3 In the **Syslog Server** section, specify the SSL, TCP, and UDP port numbers for the Syslog servers.
The default ports for SSL, TCP, and UDP are 1443, 1468, and 1514 respectively.
- 4 To start or stop the data collection for each of the Syslog servers, select the On or Off options next to them.
- 5 To change the port values, specify a valid port value.
The following table shows the status messages you see after entering the valid or non-valid port values.

Status Icon	Message
Green Check Mark	If the specified port is valid and is not in use, a port is valid and open message is displayed.
Red Cross	If the specified port is not valid (non-numeric or not between 1 to 65535), a port is not valid message is displayed.
Red Cross	If the specified port is valid but it is already in use, or if the Syslog server does not have permission to use it, a port is valid but not open message is displayed.

- 6 Set the appropriate client authentication and server key pairs settings for the SSL Syslog server.
For more information about setting the client authentication, see [“Configuring Client Authentication for the SSL Syslog Server” on page 81](#).
The SSL Syslog server is automatically restarted if any changes are made here.
- 7 (Optional) Click **Reset** to restore the previous settings.
- 8 Click **Save** to save the new settings.
The **Save** button is disabled until you specify a valid port for all of the servers.

7.2.3 Configuring Client Authentication for the SSL Syslog Server

The client authentication settings determine how strictly the SSL Syslog server verifies the identity of Syslog event sources that are attempting to send their data. You should use a strict client authentication policy that is applicable in your environment to prevent rogue Syslog event sources from sending undesired data into Sentinel.

Open: No authentication is required. Sentinel does not request, require, or validate a certificate from the event source.

Loose: A valid X.509 certificate is required from the event source, but the certificate is not validated. It does not need to be signed by a certificate authority.

Strict: A valid X.509 certificate is required from the event source, and it must be signed by a trusted certificate authority. If the event source does not present a valid certificate, Sentinel does not accept its event data.

- ♦ [“Creating a Trust Store” on page 81](#)
- ♦ [“Importing a Certificate into the Trust Store” on page 82](#)
- ♦ [“Server Key Pair” on page 82](#)

Creating a Trust Store

For strict authentication, you must have a trust store that contains the public certificate of the certificate authority (CA) that signed the event source certificate. After you have a DER or PEM certificate, you can create the trust store by using the TruststoreCreator utility that comes with Sentinel.

- 1 Log in to the Sentinel server as novell user.
- 2 Go to the `/var/opt/novell/sentinel/data/updates/done` folder.
- 3 Use the following command to extract the `syslog_connector.zip` file:

```
unzip syslog_connector.zip
```

- 4 Copy the `TruststoreCreator.sh` or `TruststoreCreator.bat` file to the machine that has the certificates.

or

Copy the certificates to the computer with the `TruststoreCreator` utility.

- 5 Run the `TruststoreCreator.sh` utility as follows:

```
TruststoreCreator.sh -keystore /tmp/my.keystore -password password1 -certs /  
tmp/cert1.pem,/tmp/cert2.pem
```

In this example, the `TruststoreCreator` utility creates a keystore file called `my.keystore` that contains two certificates (`cert1.pem` and `cert2.pem`). It is protected by the password `password1`. The keystore file must be imported into the trust store.

Importing a Certificate into the Trust Store

For strict authentication, the administrator can import a certificate by using the **Import** button. This helps ensure that only authorized event sources are sending data to Sentinel. The trust store must include public certificate of the certificate authority (CA) that signed the event source certificate.

The following procedure must be run on the computer that has the trust store on it. You can open a Web browser on the computer with the trust store or move the trust store to any computer with a Web browser.

NOTE: If the CA is signed by another CA, you must import the chain of CA certificates until the root CA.

To import a trust store:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection > Event Source Servers**.
- 3 In the Syslog Server section, select the **Strict** option under **Client authentication**.
- 4 Click **Browse** and browse to the trust store file (for example, `my.keystore`)
- 5 Specify the password for the truststore file.
- 6 Click **Import**.
- 7 (Optional) Click **Details** to see more information about the trust store.
- 8 (Optional) Click **Reset** to restore the previous settings.
- 9 Click **Save**.

After the trust store is successfully imported, you can click **Details** to see the certificates included in the trust store.

Server Key Pair

Sentinel is installed with a built-in certificate, which is used to authenticate the Sentinel server to the event sources. This certificate can be overridden with a certificate signed by a public certificate authority (CA).

To replace the built-in certificate:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection > Event Source Servers**.

- 3 In the Syslog Server section, under **Server key pairs**, select **Custom**.
- 4 Click **Browse** and browse to the trust store file.
- 5 Specify the password for the trust store file.
- 6 Click **Import**.
- 7 (Optional) If there is more than one public-private key pair associated with the file, select the desired key pair, then click **OK**.
- 8 (Optional) Click **Details** to see more information about the server key pair.
- 9 (Optional) Click **Reset** to restore the previous settings.
- 10 Click **Save**.

7.3 Configuring Data Collection for the Novell Audit Server

Many Novell products send data to Sentinel by using a Platform Agent. Data is received by an Event Source Server called the Audit Server, which is packaged with Sentinel. The Audit Server is similar in many ways to a Syslog server.

The following sections describe the procedure to configure the Audit Server port to receive data and also to set the Audit Server options:

- [Section 7.3.1, “Specifying the Audit Server Settings,” on page 83](#)
- [Section 7.3.2, “Setting the Audit Server Options,” on page 84](#)

7.3.1 Specifying the Audit Server Settings

To specify the data collection settings for the Audit Server:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection > Event Source Servers**.
- 3 In the **Audit Server** section, select the **On** or **Off** options to start or stop the data collection for the Audit Server.
- 4 In the Audit Server section, specify the port on which the Sentinel server listens to messages from the event sources.

For more information about setting the port, see [“Port Configuration and Port Forwarding for the Audit Server” on page 84](#).

- 5 Set the appropriate client authentication and server key pairs settings.

For more information about client authentication, see [“Client Authentication for the Audit Server” on page 85](#).

The Audit server and all related Audit Connectors are automatically restarted if any changes are made here.

- 6 Select the Sentinel server behavior when the number of events received exceeds the buffer capacity.

WARNING: If you select **Drop oldest messages**, there is no supported method to recover dropped messages,

- 7 Select **Idle Connection** to disconnect event sources that have not sent data for a certain period of time.

The event source connections are automatically re-created when they start sending data again.

8 Specify the number of minutes before an idle connection is disconnected.

9 (Optional) Select **Event Signatures** to receive a signature with the event.

To receive a signature, the Platform Agent on the event source must be configured properly. For information on validating event signatures, see the Signing Events section in the [Audit Platform Agent Guide](#).

10 (Optional) Click **Reset** to restore the previous settings.

11 Click **Save** to save the new settings.

The **Save** button is disabled until a valid port is specified for the server.

These settings might affect data collection for several servers (for example, multiple eDirectory instances). However, they do not start or stop services on the event source computers. Changes on this page take effect immediately.

To view the health of the Audit Server and its event sources, see [Section 7.5, “Managing Event Sources,”](#) on page 105.

7.3.2 Setting the Audit Server Options

Administrators can change the settings for how Sentinel listens for data from the event source applications, set the port on which Sentinel listens, and select the type of authentication between the event source and Sentinel.

- ♦ [“Port Configuration and Port Forwarding for the Audit Server”](#) on page 84
- ♦ [“Client Authentication for the Audit Server”](#) on page 85

Port Configuration and Port Forwarding for the Audit Server

By default, Sentinel listens on port 1289 for messages from the server. When the port is changed, the system checks whether the specified port is valid and open.

Binding to ports lower than 1024 requires `root` privileges, so you should use a port higher than 1024. You can change the source devices to send data to a higher port or use port forwarding on the Sentinel server.

To change the Platform Agent to send data to a different port:

- 1 Log in to the event source computer.
- 2 Open the `logevent` file for editing. The file location depends on the operating system:
 - ♦ Linux: `/etc/logevent.conf`
 - ♦ Windows: `C:\WINDOWS\logevent.cfg`
 - ♦ NetWare: `SYS:\etc\logevent.cfg`
 - ♦ Solaris: `/etc/logevent.conf`
- 3 Set the `LogEnginePort` parameter to the desired port.
- 4 Save the file.
- 5 Restart the Platform Agent.

The method varies by operating system and application. Restart the computer or refer to the application-specific documentation on the [Novell Documentation Web site \(http://www.novell.com/documentation\)](http://www.novell.com/documentation) for more instructions.

To configure port forwarding on the Sentinel server:

- 1 Log in to the Sentinel server operating system as root (or su to root).
- 2 Open the `/etc/init.d/boot.local` file for editing.
- 3 Add the following command at the end of the bootstrap process:

```
iptables -A PREROUTING -t nat -p protocol --dport incoming port -j DNAT --to-destination IP:rerouted port
```

Replace *protocol* with `tcp` or `udp`, *incoming port* with the port where the messages are arriving, and *IP:rerouted port* with the IP address of the local computer and an available port above 1024.

- 4 Save the changes.
- 5 Reboot the server. If you cannot reboot immediately, run the `iptables` command in [Step 3](#) from a command line.

Client Authentication for the Audit Server

The event sources send their data over an SSL connection, and the Client authentication setting for the Sentinel server determines what kind of authentication is performed for the certificates from the Audit Server on the event sources.

Open: No authentication is required. Sentinel does not request, require, or validate a certificate from the Event Source.

Loose: A valid X.509 certificate is required from the Event Source, but the certificate is not validated. It does not need to be signed by a certificate authority.

Strict: A valid X.509 certificate is required from the Event Source, and it must be signed by a trusted certificate authority. If the Event Source does not present a valid certificate, Sentinel does not accept its event data.

- ♦ [“Creating a Trust store” on page 85](#)
- ♦ [“Importing a Certificate into the Trust Store” on page 86](#)
- ♦ [“Server Key Pair” on page 86](#)

Creating a Trust store

For strict authentication, you must have a trust store that contains the public certificate of the certificate authority (CA) that signed the event source certificate. After you have a DER or PEM certificate, you can create the trust store by using the `CreateTrust store` utility that comes with Sentinel.

- 1 Log in to the Sentinel server as `novell`.
- 2 Go to `/var/opt/novell/sentinel/data/updates/done`.
- 3 Unzip the `audit_connector.zip` file:

```
unzip audit_connector.zip
```
- 4 Copy `TruststoreCreator.sh` or `TruststoreCreator.bat` to the computer with the certificates or
Copy the certificates to the computer with the `TruststoreCreator` utility
- 5 Run the `TruststoreCreator.sh` utility:

```
TruststoreCreator.sh -keystore /tmp/my.keystore -password password1 -certs /  
tmp/cert1.pem,/tmp/cert2.pem
```

In this example, the TruststoreCreator utility creates a keystore file called `my.keystore` that contains two certificates (`cert1.pem` and `cert2.pem`). It is protected by the password `password1`.

Importing a Certificate into the Trust Store

For strict authentication, the administrator can import a certificate. This helps ensure that only authorized event sources are sending data to Sentinel. The trust store must include public certificate of the certificate authority (CA) that signed the event source certificate.

The following procedure must be run on the machine that has the trust store on it. You can open a Web browser on the machine with the trust store or move the trust store to any machine with a Web browser.

NOTE: If the CA is signed by another CA, then you must import the chain of CA certificates until the root CA.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection > Event Source Servers**.
- 3 In the Audit Server section, select the **Strict** option under **Client authentication**.
- 4 Click **Browse** and browse to the trust store file (for example, `my.keystore`)
- 5 Specify the password for the trust store file.
- 6 Click **Import**.
- 7 (Optional) Click **Details** to see more information about the trust store.
- 8 (Optional) Click **Reset** to restore the previous settings.
- 9 Click **Save**.

After the trust store is imported successfully, you can click **Details** to see the certificates included in the trust store.

Server Key Pair

Sentinel is installed with a built-in certificate, which is used to authenticate the Sentinel server to the event sources. This certificate can be overridden with a certificate signed by a public certificate authority (CA).

To replace the built-in certificate:

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection > Event Source Servers**.
- 3 In the Audit Server section, under **Server key pairs**, select **Custom**.
- 4 Click **Browse** and browse to the trust store file.
- 5 Specify the password for the trust store file.
- 6 Click **Import**.
- 7 (Optional) If there is more than one public-private key pair in the file, select the desired key pair and click **OK**.
- 8 (Optional) Click **Details** to see more information about the server key pair.
- 9 (Optional) Click **Reset** to restore the previous settings.
- 10 Click **Save**.

7.4 Configuring Data Collection for Other Event Sources

The Event Source Management (Live View) interface provides a set of tools to manage and monitor connections between Sentinel and the event sources that provide data to Sentinel. The graphical interface shows the current event sources and the software components that are processing data from that event source. Each component can be easily deployed to integrate the devices in the enterprise, and it can be monitored in real time within the Event Source Management interface. Some Connectors and Collectors must be configured in Event Source Management, such as the WMS connector for Windows, Database connectors, and SDEE connectors for Cisco devices.

NOTE: If you are using openSUSE11.1, update your JRE to the latest version. Then use the Java Web Start (javaws) launcher command to launch the Event Source Management.

You can perform the following tasks through the Event Source Management window:

- ♦ Add or modify connections to event sources by using Configuration wizards.
- ♦ View the real-time status of the connections to event sources.
- ♦ Import or export configuration of event sources to or from the Live View.
- ♦ View and configure Connectors and Collectors that are installed.
- ♦ Import or export Connectors and Collectors from or to a centralized repository.
- ♦ Monitor data flowing through the Collectors and Connectors.
- ♦ Design, configure, and create the components of the event source hierarchy, and execute required actions by using these components.
- ♦ [Section 7.4.1, “Accessing Event Source Management,” on page 87](#)
- ♦ [Section 7.4.2, “Viewing Data in Event Source Management,” on page 88](#)
- ♦ [Section 7.4.3, “Searching for Event Sources,” on page 93](#)
- ♦ [Section 7.4.4, “Installing Plug-Ins,” on page 94](#)
- ♦ [Section 7.4.5, “Updating a Connector or a Collector Plug-In,” on page 95](#)
- ♦ [Section 7.4.6, “Adding Components to Sentinel,” on page 95](#)
- ♦ [Section 7.4.7, “Connecting to Event Sources,” on page 97](#)
- ♦ [Section 7.4.8, “Exporting Configurations,” on page 100](#)
- ♦ [Section 7.4.9, “Importing Configurations,” on page 101](#)
- ♦ [Section 7.4.10, “Debugging,” on page 101](#)
- ♦ [Section 7.4.11, “Troubleshooting,” on page 105](#)

7.4.1 Accessing Event Source Management

- 1 Log in to the Sentinel Web interface as a user in the administrator role.

`https://<IP_Address/DNS_Sentinel_server>:8443`

IP_Address/DNS_Sentinel_server is the IP address or DNS name of the Sentinel server and 8443 is the default port for the Sentinel server.

- 2 In the toolbar, click **Applications > Launch Control Center**.

or

In the toolbar, click **Collection > Advanced > Launch Control Center**.

- 3 Log in to the Sentinel Control Center.
- 4 In toolbar, click **Event Source Management** > **Live View**.

7.4.2 Viewing Data in Event Source Management

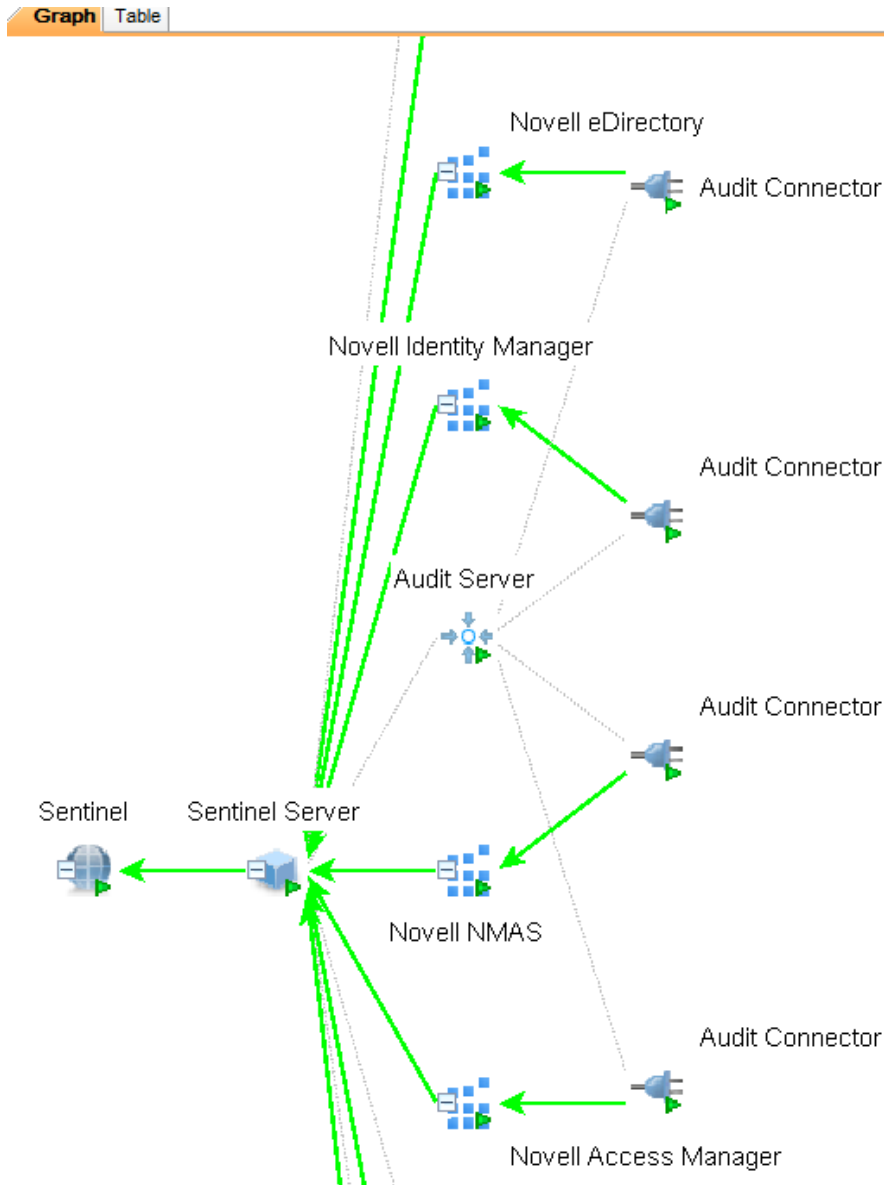
In the Event Source Management, you can view configuration data in different views:

- ♦ [“Graphical ESM View” on page 88](#)
- ♦ [“Table ESM View” on page 90](#)
- ♦ [“How the Components Are Displayed” on page 91](#)
- ♦ [“Component Status Indicators” on page 91](#)
- ♦ [“Right-Click Menus” on page 92](#)

Graphical ESM View

The graphical view is the default view in Event Source Management. In the graphical view, you can view the status of a Collector and access the configuration settings of Collectors and Collector objects as a graph of connected nodes.

Figure 7-1 Graphical View



By default, the Health Monitor Display frame displays in the graphical view. The data can be displayed in several different layouts. The default layout in graph is the Hierarchic Left to Right layout. You can change between these layouts by selecting the layout format from the drop-down list in the toolbar.

TIP: Click in the graphical ESM view and use the "+" or "-" sign to zoom in or zoom out, or use the mouse wheel to zoom in and zoom out.

In the graphical view, the lines connecting the components are color-coded to indicate data flow.

Green Line: Indicates that data is flowing between the components.

Grey Line: Indicates that the connection is not live and there is no data flow.

Blue dashed Line: Indicates the logical relation of Event Source Servers to their associated Collector Managers and event sources.

To improve the manageability and performance of the graphical display, Sentinel automatically collapses any node with 20 or more immediate child node. This is especially useful for Connectors such as Syslog or Novell Audit that have the ability to automatically configure a large number of event sources.

In collapsed state, a node displays the number of immediate child nodes, such as WMI Connector (3) [Collector name (Number of immediate children)]. The Children panel of a collapsed node shows the immediate children of that node, each of which can be managed in the same way as nodes in the tabular ESM view.


NOTE: Event Source Server node do not have the plus or minus sign after their names even if they contain child nodes.

The parent node can take several minutes to expand if it has a large enough number of child nodes to potentially cause the user interface to become unresponsive.

Table ESM View

The components visible in the graphical view of ESM can also be viewed in a table. You can view the status of a Collector in the table and access the configuration settings of Collectors and Collector related objects.

Figure 7-2 Event Source Management Table View



The screenshot shows the 'Table' tab selected in the Event Source Management interface. The table has three columns: Name, Configured Status, and Actual Status. The rows list various event sources, including Sentinel, Sentinel Server, Syslog Server TCP, Syslog Server UDP, Generic Event, Audit Server, Novell Access Manager, Cisco Firewall 6.1r2, and Syslog Connector. The 'Syslog Server TCP' row is highlighted in blue.

Name	Configured Status	Actual Status
Sentinel	On	On
Sentinel Server	On	On
Syslog Server TCP	On	On
Syslog Server UDP	On	On
Generic Event	On	On
Audit Server	On	On
Novell Access Manager	On	On
Cisco Firewall 6.1r2	On	On
Syslog Connector	On	On

The columns in the ESM table view are as follows:

Configured Status: The On state the object is configured to be in. This is the state that is stored in the database, which does not necessarily match the actual On state of the object. For example, the two states do not match if a parent object is turned off or if there is an error.

Actual Status: The On state of the object as reported by the actual running Collector Manager.

Connection Info (populated for Event Sources only): A text description of the event source connection.







Error: A text description of an error that occurred in the running object.

TIP: Use the **Table** and **Graph** tabs to change to tabular and graphical views.

How the Components Are Displayed

ESM displays the information on the Collectors and other components in a hierarchy specific to ESM.








Table 7-1 *Components of the ESM Hierarchy*

	Sentinel	<p>The single Sentinel icon represents the main Sentinel Server that manages all events collected by the Sentinel system.</p> <p>The Sentinel object is installed automatically through the Sentinel installer.</p>
	Collector Manager or the Sentinel Server	<p>The Collector Manager display name in the ESM is Sentinel Server.</p> <p>Each icon represents another instance of a Collector Manager process. Multiple Collector Manager processes can be installed throughout the enterprise. As each Collector Manager process connects to Sentinel, the objects are automatically created in ESM.</p>
	Collector	<p>Collectors instantiate the parsing logic for data from a particular event source. Each Collector icon in ESM refers to a deployed Collector script as well as the runtime configuration of a set of parameters for that Collector.</p>
	Connector	<p>Connectors are used to provide the protocol-level communication with an event source, using industry standards like Syslog, JDBC, and so forth. Each instance of a Connector icon in ESM represents the Connector code as well as the runtime configuration of that code.</p>
	Event Source Server	<p>An Event Source Server (ESS) is considered part of a Connector, and is used when the data connection with an event source is inbound rather than outbound. The ESS represents the daemon or server that listens for these inbound connections. The ESS caches the received data, and one or more Connectors connect to the ESS to retrieve a set of data for processing. The Connector requests only the data from its configured event source (defined in the metadata for the event source) that matches additional filters.</p>
	Event Source	<p>The event source represents the actual source of data for Sentinel. Unlike other components, this is not a plug-in, but is a container for metadata, including runtime configuration about the event source. In some cases a single event source can represent many real sources of event data, such as if multiple devices are writing to a single file.</p>

Component Status Indicators

Indicators are used to represent various states as follows:

Table 7-2 *Component Status Indicators*

	Stopped	Indicates that the component is stopped.
	Running	Indicates that the component is running.
	Warning	Indicates that a warning is associated with the component. At this time, this warning indicator is primarily used to show when the configured state and actual state of a component differ. (That is, a component is configured to be running, but the actual state of the component is stopped.)
	Error	Indicates that an error is associated with the component. See the individual component's status display for details about the error.
	Reporter Time is Skewed	Indicates when the time of a component differs from the main server's time. (The difference is greater than a predefined time threshold.)
	Debug	Indicates that the component is in Debug mode. Only a Collector can be in Debug mode.
	Unknown	Indicates that the status of the object in the ESM panel is not yet known.

Right-Click Menus

The Health Monitor Display View provides a set of right-click menus that helps you execute a set of actions, as described below:

The right-click actions available depend on the object you click.

Status Details: View all information known about the status of the selected object.

Start: Run an object.

The selected object starts only after the parent nodes starts and is running.

Stop: Stops the running object.

Edit: Modifies the editable information (Filter information, Object name and so on).

Debug: Debugs the Collector. You must stop the running Collector before you debug it.

Move: Moves the selected object from its current parent object to another parent object. You can move objects from a Live View to the scratch pad and vice versa.

Clone: Creates a new object that has its configuration information prepopulated with the settings of the currently selected object. This allows you to quickly create a large number of similar event sources without retyping the same information over and over again. You can clone objects from Live View to the scratch pad and vice versa. Cloning an object copies all the settings except the Run status. New objects created using the Clone command are always in the Stopped state after creation.

Remove: Deletes a selected object from the system.

Contract: Collapses the child nodes into this node. This option is only available on parent nodes that are currently expanded.

Expand: Expands the child nodes of this node. This option is only available on parent nodes that are currently collapsed.

Add Collector: Opens an Add Collector Wizard that guides you through the process of adding a Collector to the selected Collector Manager.

Add Connector: Opens an Add Connector Wizard that guides you through the process of adding a Connector to the selected Collector.

Add Event Source: Opens an Add Event Source Wizard that guides you through the process of adding an event source to the selected Connector.

Open Raw Data Tap: Lets you view the live stream of raw data from an event source or flowing through the selected object.

Open Active View: Opens an Active View window that displays only events that have been generated by data from or flowing through the selected object.

Zoom: In the graphical view, zoom in on the selected object.

Show in Tabular/Graphical View: Lets you switch between the graphical view and the tabular view and automatically selects the object that is selected in the current view. When switching to the graphical view, it also zooms in on the selected object.

Raw Data Filter: Filters the raw data flowing through the selected node. The raw data filter is available on Collectors, Connectors, and event sources. If a filter specifies to drop data, the data to be dropped is not passed to the parent node and is not converted into events.

Import Configuration: Imports the configuration of ESM objects.

Export Configuration: Exports the configuration of ESM objects

Add Event Source Server: Adds an Event Source Server to the selected Collector Manager

Add Collector Manager: In Scratch pad mode, you can add a Collector Manager to the scratch pad by using this option. In the Live view, Collector Manager objects are created automatically as each Collector Manager connects to the Sentinel system.

When you select multiple objects in the ESM panel and right-click, the following options are available:

Start: Starts all objects

Stop: Stops all objects

Remove selected objects: Removes the selected objects along with their children.

7.4.3 Searching for Event Sources

You can use the Attribute Filter panel to search for event sources.

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

- 2 In the Attribute Filter panel, use the following information to display objects you want:

Search: Specify the name of the objects you want displayed.

Limit to: Select the types of objects to display.

Status: Select the status of the objects to display.

As you define each filter, the display is automatically updated.

7.4.4 Installing Plug-Ins

Although some Sentinel components are preinstalled with the Sentinel system, you should also check the [Sentinel Plug-ins Web site](#) to download the latest versions on the plug-ins.

- ♦ [“Installing a Connector Plug-In” on page 94](#)
- ♦ [“Installing a Collector Plug-In” on page 94](#)

Installing a Connector Plug-In

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).
- 2 In the toolbar, click **Tools > Import plug-in**.
- 3 Select **Import Collector Script or Connector plug-in package file (.zip, .clz, .cnz)**.
- 4 Click **Next**.
- 5 Click **Browse**.
- 6 Browse to and select the Connector plug-in package file, then click **Open**.
- 7 Click **Next**.
- 8 (Conditional) If the Connector already exists in the plug-in repository, select to replace the existing plug-in with the new plug-in by clicking **Next**.
- 9 (Conditional) In the plug-in details window, select **Deployed Plug-ins** to deploy the plug-in from this window.

For more information, see [“Access Event Source Management.” on page 97](#).

- 10 Click **Finish**.

When you add a plug-in to Sentinel, it is placed in the plug-in repository, which enables Sentinel components on other machines to start using the plug-in without adding the plug-in separately.

Installing a Collector Plug-In

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).
- 2 In the toolbar, click **Tools > Import plug-in**.
- 3 Select **Import Collector Script or Connector plug-in package file (.zip, .clz, .cnz)**.
or
Select **Import Collector Script from directory**.
- 4 Click **Next**.
- 5 Click **Browse**.
- 6 Browse to and select the Collector script from a file or directory, then click **Open**.
- 7 Click **Next** to display the plug-in details window.
- 8 Select **Deploy Plug-in** to deploy the plug-in from this window.

For more information on the deployment procedure, see [“Access Event Source Management.” on page 97](#).

- 9 Click **Finish**.

7.4.5 Updating a Connector or a Collector Plug-In

If a new version of a Connector or Collector is released, you can update the Sentinel system and any deployed instances of the Connector or Collector.

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).
- 2 In the toolbar, click **Tools > Import plug-in**.
- 3 Select **Import Collector Script or Connector plug-in package file (.zip, .clz, .cnz)**.
or
Select **Import Collector Script from directory**.
- 4 Click **Next**.
- 5 Click **Browse**.
- 6 Browse to and select the Connector or Collector plug-in package file, then click **Open**.
- 7 Click **Next**.
- 8 Click **Next** to update an already-imported Connector or Collector.
- 9 In the plug-in details window, select **Update Deployed Plug-ins** to update any currently deployed plug-ins that use this Connector or Collector.
- 10 Click **Finish**.

When you add a plug-in to Sentinel, it is placed in the plug-in repository, which enables Sentinel components on other machines to start using the plug-in without adding the plug-in separately.

7.4.6 Adding Components to Sentinel

After the plug-ins are installed in the Event Source Management, you must add the different components to your Sentinel solution.

- ♦ [“Adding a Collector” on page 95](#)
- ♦ [“Adding a Connector” on page 96](#)
- ♦ [“Adding an Event Source” on page 96](#)
- ♦ [“Adding an Event Source Server” on page 96](#)

Adding a Collector

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).
- 2 In the main ESM display, locate the Collector Manager where the Collector will be associated.
- 3 Right-click the Collector Manager, then select **Add Collector**.
- 4 Follow the prompts in the Add Collector Wizard.
These prompts are unique for each Collector. For details, see the specific Collector documentation at the [Sentinel Plug-ins Web page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).
- 5 Click **Finish**.

The Collector Script enables the ESM panel to prompt you for parameter values as well as enable Event Source Management to automatically select supported connection methods that work well with the Collector script.

Adding a Connector

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

- 2 In the main ESM display, locate the Collector where the new Connector will be associated
- 3 Right-click the Collector, then select **Add Connector**.
- 4 Follow the prompts in the Add Connector Wizard.

These prompts are unique for each Connector. For details, see the specific Connector documentation at the [Sentinel Plug-ins Web page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- 5 Click **Finish**.

Adding an Event Source

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

- 2 In the main ESM display, locate the Connector where the new event source will be associated.
- 3 Right-click the Connector, then select **Add Event Source**.

These prompts are unique for each event source that is associated with the Connector. For details, see the specific Connector documentation at the [Sentinel Plug-ins Web page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- 4 Follow the prompts in the Add Event Source Wizard.
- 5 Click **Finish**.

Adding an Event Source Server

Certain event source Connectors (such as the Syslog Connector) require a process to collect data from the actual data source. These processes are called Event Source Servers. They collect data from the data source and then serve it to the event source Connector. Event Source Servers must be added and associated to any event source Connectors that require a server.

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

- 2 Right-click the Collector Manager, then select **Add Event Source Server**.
- 3 Select a Connector that supports your device, then click **Next**.

If you do not have any connectors in the list that supports your device, click **Install More Connectors**. For more information on installing a Connector plug-in, see [“Installing a Collector Plug-In” on page 94](#).

- 4 Configure the various parameters for the server that is associated with the selected Connector. For example, Syslog Connector, NAudit Connector, and so on.

These parameters are unique for each Connector. For details, see the specific Connector documentation at the [Sentinel Plug-ins Web page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- 5 Click **Next**.
- 6 Specify a name for the Event Source Server.
- 7 (Optional) If you want this server to run, select **Run**.
- 8 Click **Finish**.

In the Health Monitor Display frame, the Event Source Server is displayed with a dashed line showing which the Collector Manager it is associated with.

This Add Event Source Server Wizard can also be initiated from within the Add Connector Wizard if a compatible Event Source Server has not yet been added.

7.4.7 Connecting to Event Sources

There are many different ways to add an event source. The following procedures walk you through the process.

- ♦ “Prerequisites” on page 97
- ♦ “Connecting to the Event Source” on page 97
- ♦ “Creating a New Collector and Connector” on page 99
- ♦ “Using an Existing Collector” on page 99
- ♦ “Using an Existing Connector” on page 100

Prerequisites

Make sure you have the following prerequisites:

Collector Script: Collector scripts can be downloaded from the [Sentinel Plug-ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html) or built with the Collector Builder.

Connector: Connectors can be downloaded from the [Sentinel Plug-ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html). There are also some Connectors included in the installed Sentinel system, but there might be more recent versions on the Web site.

Documentation: Check the documentation for each Connector and Collector, because they have different configuration steps for the event source. The documentation is located on the [Sentinel Plug-ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html). Make sure you download the documentation when you download the Connector and Collector.

Event Source Configuration: You must have configuration information for the event source.

Connecting to the Event Source

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).
- 2 In the toolbar, click **Tools > Connect to Event Source**.
The event source types for the compatible Collector parsing scripts are listed here.
- 3 Select the desired Event Source.
You can click **Add More** to import an event source not listed.
- 4 After the event source is selected, click **Next**.
- 5 Select a Collector script from the list.

You can click **Install More Scripts** to install additional Collector scripts that support your Event Source.

For more information on installing a Collector script, see [“Installing a Collector Plug-In” on page 94](#).

6 Click **Next**.

7 Select a connection method from the list.

There are many different types of Connectors. Depending on the type of Connector you select, there are additional configuration screens.

You can click **Install More Connectors** to install additional Connectors.

For more information, see [“Installing a Connector Plug-In” on page 94](#) to install connectors.

8 Click **Next**.

9 Use the following information to select how to manage the event source connection, then proceed to [Step 10](#).

- ♦ [“Creating a New Collector and Connector” on page 99](#)
- ♦ [“Using an Existing Collector” on page 99](#)
- ♦ [“Using an Existing Connector” on page 100](#)

Based on the existing Collectors and Connectors in your system that are compatible with your new event source, one or more of these options might be unavailable.

10 Use the following information to configure the event source:

Name: Specify a unique name for the event source.

Run: Select **Run** if you want the event source to run automatically.

Details: Allows you to see the details of the plug-in.

Alert if no data received in specified time period: Select this option to receive notifications if no data is received during the specified time period.

Limit Data Rate: Use this option to limit the maximum number of records the Connector receives per second.

Trust Event Source Time: (Optional) Select this option to set the event time to the time the event occurred, rather than the time Sentinel received the data.

You can also set this option while configuring an event source. If the **Trust Event Source Time** option is selected, all data flowing through the Collector has the event time set to the time the event occurred, even if the event sources do not have this option selected.

Set Filters: Allows you to set filters on the data in the event source.

11 Click **Next**.

12 Click **Test Connection** to test the event source.

12a Click the **Data** tab to view the data in the event source.

It takes a few seconds for the raw data to be displayed in the **Data** tab.

12b Specify the maximum number of rows to control the number of raw data records obtained at one time.

12c Click the **Error** tab to view if there are any errors in the configuration of the event source.

12d Click **Stop** to stop the test.

13 Click **Finish**.

The Collector parsing script is executed on the same system as the Collector Manager that you select here.

Creating a New Collector and Connector

Use the following information to create a new Collector and Connector to manage the event source connection. This procedure is a continuation of [Step 9 on page 98](#).

- 1 Select **Create a new Collector and Connector**, then click **Next**.

- 2 Select the Collector Manager you want to use, then click **Next**.

- 3 Change any of the Collector properties, then click **Next**.

- 4 Use the following information to configure the Collector:

Name: Specify a unique name for the Collector.

Run: Select **Run** if you want to run the Collector automatically.

Details: Allows you to view the details of the plug-in.

Alert if no data received in specified time period: Select this option to receive notifications if no data is received during the specified time period.

Lime Data Rate: Use this option to limit the maximum number of records the Collector receives per second.

Trust Event Source Time: (Optional) Select this option to set the event time to the time the event occurred, rather than the time Sentinel received the data.

You can also set this option while configuring an event source. If the **Trust Event Source Time** option is selected, all data flowing through the Collector has the event time set to the time the event occurred, even if the event sources do not have this option selected.

Set Filters: Allows you to set filters on the data in the Collector.

- 5 Click **Next**.

There is a different configuration page displayed depending on the type of Connector you selected in [Step 7 on page 98](#). For the Connector-specific documentation, see the [Sentinel Plugins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- 6 Use the following information to configure the Connector:

Name: Specify a unique name for the Connector.

Run: Select **Run** if you want to run the Connector automatically.

Details: Allows you to view the details of the plug-in.

Alert if no data received in specified time period: Select this option to receive notifications if no data is received during the specified time period.

Limit Data Rate: You can limit the maximum number of records the Connector receives per second.

Set Filters: Allows you to set filters on the data in the Connector.

Copy Raw Data to a file: Select this option, then specify a location where you want to copy the raw data coming from the event source.

- 7 Click **Next**, then continue with [Step 10 on page 98](#).

Using an Existing Collector

If you are using an existing Collector, but want to create a new Connector to manage the Event Source connection, use the following information to complete the procedure from [Step 9 on page 98](#).

- 1 Select **Use an Existing Collector**, then click **Next**.

- 2 Select the Collector you want to use, then click **Next**.

There is a different configuration page displayed depending on the type of Connector you selected in [Step 7 on page 98](#). For the Connector-specific documentation, see the [Sentinel Plug-ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- 3 Use the following information to configure the Connector:

Name: Specify a unique name for the Connector.

Run: Select **Run** if you want to run the Connector automatically.

Details: Allows you to view the details of the plug-in.

Alert if no data received in specified time period: Select this option to receive notifications if no data is received during the specified time period.

Limit Data Rate: Use this option to limit the maximum number of records the Connector receives per second.

Set Filters: Allows you to set filters on the data in the Connector.

Copy Raw Data to a file: Select this option, then specify a location where you want copy the raw data coming from the event source.

- 4 Click **Next**, then continue with [Step 10 on page 98](#).

Using an Existing Connector

If you are using an existing Connector, but want to create a new Collector to manage the event source connection, use the following information to continue the procedure from [Step 9 on page 98](#).

- 1 Select **Use an Existing Connector**, then click **Next**.
- 2 Select the Collector Manager you want to use, then click **Next**.
- 3 Change any of the Collector properties, then click **Next**.
- 4 Use the following information to configure the Collector:

Name: Specify a unique name for the Collector.

Run: Select **Run** if you want to run the Collector automatically.

Details: Allows you to view the details of the plug-in.

Alert if no data received in specified time period: Select this option to receive notifications if no data is received during the specified time period.

Limit Data Rate: Use this option to limit the maximum number of records the Collector receives per second.

Trust Event Source Time: (Optional) Select this option to set the event time to the time the event occurred, rather than the time Sentinel received the data.

You can also set this option while configuring an event source. If the **Trust Event Source Time** option is selected, all data flowing through the Collector has the event time set to the time the event occurred, even if the event sources do not have this option selected.

Set Filters: Allows you to set filters on the data in the Collector.

- 5 Click **Next**, then continue with [Step 10 on page 98](#).

7.4.8 Exporting Configurations

Event Source Management allows you to export the configuration of Event Source Management objects along with the associated Collector scripts and the Connector plug-ins.

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87.](#)

2 Click File > Export Configuration.

or

Right-click an object, then click **Export Configuration**.

3 Select which nodes you want to export, then click Next.

4 Select the Collector scripts to export, then click Next.

5 Select the Connector plug-ins to export, then click Next.

6 Click Browse, then browse to a location to save the export.

7 Specify a file name, then click Save.

The export information is saved as a .zip file.

8 Click Next.

9 Review the items to be exported, then click Finish.

7.4.9 Importing Configurations

Event Source Management allows you to import the configuration files that you export. The configuration files contain configuration information for Event Source Management objects along with the associated Collector scripts and Connector plug-ins.

1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87.](#)

2 Click File > Import Configuration.

3 Click Browse and browse to and select the configuration file, then click Open.

The configuration files are .zip files.

4 Click Next.

5 Select the nodes to import, then click Next.

6 Select the Collector scripts to import, then click Next.

7 Select the Connector plug-ins to import, then click Next.

8 Review the items to import, then click Finish.

7.4.10 Debugging

Sentinel's Collectors are designed to be easily customizable and to be created by customers and partners. The debugging interface analyzes the Collector code running in place on the Collector Manager.

For more information on customizing or creating new Collectors, obtain the Developer's Kit for Sentinel at the [Sentinel SDK Web site \(http://www.novell.com/developer/develop_to_sentinel.html\)](http://www.novell.com/developer/develop_to_sentinel.html).

- ♦ [“Collector Workspace and Collector Directory” on page 102](#)
- ♦ [“Debugging JavaScript Collectors” on page 102](#)
- ♦ [“Generating a Flat File using the Raw Data Tap” on page 104](#)

Collector Workspace and Collector Directory

Collectors are simple text scripts that are run by a Collector Manager. The handling of these scripts is a bit complex:

1. The code for all Collectors is stored in a plug-in repository on the central Sentinel server when the Collectors are imported.

Location: *sentinel/data/plugin_repository* on the Sentinel server.
2. The runtime configuration for the Collector (when it is configured to run on a particular Collector Manager) is stored separately in the Sentinel database.
3. When a Collector is actually started in the Collector Manager, the Collector plug-in is deployed to the Collector Manager, the runtime configuration is applied, and the code is started. Any pre-existing instance of the Collector code on that Collector Manager is overwritten.

Location: *sentinel/data/collector_mgr.cache/collector_instances* on each Collector Manager.

4. In order to edit a Collector, you need to use the ESM Debugger **Download** button, which copies the Collector to the local Collector workspace on the client machine (the machine where you are running SCC). Edits are made against that local copy and then uploaded back into the central plug-in repository.

Location: *sentinel/data/collector_workspace* on the client application machine.

Debugging JavaScript Collectors

The debugger for JavaScript Collectors can be used to debug any JavaScript Collector.

- 1 Access Event Source Management.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

- 2 Select a Collector to debug in the Live View.

- 3 Select the debug mode:

Live Mode: Requires that the Collector Manager is currently running. For more information, see [“Live Mode” on page 103](#).

Stand-alone Mode: Allows you to run the Collector in debug mode without a Collector Manager running. For more information, see [“Stand-alone Mode” on page 103](#).

- 4 Right-click the Collector and select **Stop**, then click **Debug**.

The following describe how to use the JavaScript debug window:

Debug: Launches the JavaScript file in this window.

Upload/Download: Upload or download a JavaScript file here. You can download an existing JavaScript file, edit it, and upload it again to continue debugging.

Context: Displays the variable that the debugger is pointing to and its value.

Expression: Watch the values of a selected parameter here.

You can use the following options when debugging a Collector:

Run	Starts debugging.
Pause	Pauses debugging.
Step Into	Steps to the next line in the script.
Step Over	Steps over a function.
Step Out	Steps out of a function.
Stop	Stops debugging.

- ♦ [“Hot Keys” on page 103](#)
- ♦ [“Live Mode” on page 103](#)
- ♦ [“Stand-alone Mode” on page 103](#)

Hot Keys

When the source code window has the focus in the debugger, you can use the following hot keys:

- ♦ Ctrl+F to find a string in the source code
- ♦ Ctrl+G to go to a line number
- ♦ Ctrl+M to find the parenthesis or brace that matches the selected parenthesis or brace

You can also open a script file, set break-point, step through the script code, and watch variables and methods values at each step.

Live Mode

- ♦ Live debug mode requires that the Collector Manager associated with the Collector is running.
- ♦ In Live debug mode, Input to the script comes from actual event sources connected to the Collector. To get data from a specific event source, you must right-click and start the desired event source via the Event Source Management display. Starting or stopping event sources can be done any time during the debug session.

If no event source is started during the debug session, no data is available in the buffer for the Collector and you see the Collector script’s readData method blocking.

- ♦ In Live debug mode, output from the script is via live Sentinel Events. The Events can be viewed on the Active Views display.

When you are in Live debug mode, the script engine is executed on the local computer rather than the actual computer that the associated Collector Manager is running on. The Connectors and event sources still runs on the same box as the Collector Manager. When you are running debug mode, data is automatically routed from the event sources to the script engine running in debug on the local box.

Stand-alone Mode

- ♦ Stand-alone debug mode allows you to debug a Collector even if the associated Collector Manager is not running.

- For stand-alone mode, input to the script comes from an input file rather than a live event source. Specify the path to a raw data file that is used as input. For Collectors that use a DB Connector, the input file is a text file with log data in nvp format and for a Collector that uses the File Connector, the input file is a text file with log data in CSV format.
- For stand-alone mode, Output from the script is to an output file rather than to live events. You must specify the path to the output file that the script uses for output. If you specify an output file that does not exist, the system creates the file for you.

To debug in Stand-alone mode:

- 1 In Event Source Management, right-click the Collector to debug.
- 2 Select **Stop**.
- 3 Select **Debug**.
- 4 Select **Stand-alone Mode**, then specify a path for the input and output files.
If you specify an output file that does not exist, the system creates the file for you.
- 5 Click **OK** to display the Debug Collector window.
- 6 In the Debug Collector window, click **Run**.
In the Source text area, the source code of the Collector appears and stops at the first line of the text script.
- 7 Click the left side bar to toggle a breakpoint in the script code.
- 8 Click **Step Into** to go to the next breakpoint.
- 9 Click **Pause** to pause debugging whenever you want.
- 10 After debugging is complete, click **Stop** to stop debugging.
- 11 Click the **Upload/Download** tab in the debugger window.
- 12 Click **Download**, then specify a location to download the script file.
- 13 Open the file with any JavaScript editor, then make your edits.
- 14 Save the file, then click **Upload**.
- 15 Debug the uploaded script to have a Collector Script ready to use.

Generating a Flat File using the Raw Data Tap

Occasionally when debugging, it might be helpful to view Connector output data. In addition to the **Raw Data Tap** right-click option for nodes in the Sentinel Control Center, Sentinel also includes an option to save the raw data from a Connector to a file for further analysis.

To save raw data from a deployed Connector to a file:

- 1 Access Event Source Management.
For more information, see [Section 7.4.1, "Accessing Event Source Management," on page 87](#).
- 2 Right-click the Connector node, then click **Edit**.
- 3 Click the **Configure Connector** tab.
- 4 Select the **Copy Raw Data** to a file.
- 5 Specify (or browse to and select) a path on the Collector Manager machine where the raw data is saved.

IMPORTANT: The account running the Sentinel service on the Collector Manager machine must have permissions to write to the file location.

7.4.11 Troubleshooting

If the help does not launch, there is a cache file on the local machine that is running the Event Source Management that must be deleted.

- 1 Exit Event Source Management and the Sentinel Control Center.
- 2 On the local machine running Event Source Management, search for the `.novell` directory.
- 3 Delete the `sentinel` subdirectory in the `.novell` directory.
- 4 Launch Event Source Management, then click **Help**.

For more information, see [Section 7.4.1, “Accessing Event Source Management,” on page 87](#).

7.5 Managing Event Sources

The Event Sources interface displays the health of the event source and the volume of data being received from it in events per second. The event sources page lists all the event sources, such as Syslog, Audit, File, and Database, that are configured in the Event Source Management interface.

You can refine the displayed event sources by selecting Collector Managers, Event Source Servers, and Collector plug-ins. You can also specify a filter on the event source name and select particular event source health states you want to view. All of these selections and filters are stored on a per-user basis, so that each time you log into the Sentinel server you can view event sources that match your last selections. You can also perform filtering based on tags. For more information, see “[Configuring Tags](#)” in the *NetIQ Sentinel User Guide*.

- ♦ [Section 7.5.1, “Viewing the Event Sources Page,” on page 105](#)
- ♦ [Section 7.5.2, “Changing the Data Logging Status of Event Sources,” on page 108](#)

7.5.1 Viewing the Event Sources Page

The Event Sources page consists of different sections that allow you to perform different functions.

Collector Managers: Lists all the Collector Managers associated with the Sentinel system. It also displays the state and details about the Collector Managers.

Event Source Servers: Lists all the Event Source Servers associated with the Sentinel system. It also displays the state of the Event Source Servers.

Collector Plug-ins: Lists all the Collector plug-ins associated with the Sentinel system. You can also view the details about the installed plug-ins.

The Event Sources section in the right pane lists the event sources based on the options selected from the left pane.

NOTE: The Event Sources page shows event sources that were already configured or automatically detected. To manually configure additional event sources, use the Event Source Management user interface described in “[Configuring Data Collection for Other Event Sources](#)” on page 87.

- ♦ [“Viewing Event Sources” on page 106](#)
- ♦ [“Viewing Collector Managers” on page 107](#)
- ♦ [“Viewing Event Source Servers” on page 107](#)
- ♦ [“Viewing Collector Plug-Ins” on page 108](#)

Viewing Event Sources

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Event Sources page is displayed.

Each column in the Event Source section has different information:

Health Icon: The colored icon indicates the event source health.

- ♦ **Green:** Indicates that the event source is healthy and Sentinel has received data from it.
- ♦ **Red:** Indicates that the Sentinel server is reporting an error about connecting to or receiving data from this event source.
- ♦ **Gray:** Indicates that the event source is turned off. Sentinel is not processing any data from it.
- ♦ **Orange:** Indicates that the event source is running with some warnings.

You can sort the event sources based on their health status.

Name: Displays the name given to the Event Source by the system (if it was auto-created) or by a user. For Syslog Event Sources, if the Event Source was auto-created by the system, the name is a combination of the hostname/IP address and the Collector connection mode the event source is using.

You can rename any Event Source at any time through the Event Source Management interface.

You can sort the Event Sources in alphabetical order based on their names.

Collector Plug-in: Displays the name of the Collector plug-in that the event source is connected to.

This is the name of the Collector plug-in, not the name of the Collector instance. You can sort the event sources based on Collector plug-in name.

Drop: Indicates whether data from the associated event source should be dropped.

- ♦ **YES:** If **Drop Data** is set to YES, all data received from the event source is dropped. This means that the data is not saved and events are not generated.
- ♦ **NO:** If **Drop Data** is set to NO, all data from the event source is saved and events are generated. When it is set to NO, data is always saved, regardless of whether a filter is set on the event source using the Event Source Management user interface. However, if a filter is set, events might not be generated if the filter causes the data to be ignored.

You can sort the event sources based on the drop data status.

Create Date: Specifies the date and time when the event source was created. You can sort the event sources based on when they were created.

EPS: Displays the events per second value received from the event source. You can sort the event sources based on their events per second value.

If you see a value of less than one (<1) in this column, it indicates that the EPS rate is greater than zero, but less than one.

- 3 To select or deselect an event source, select the check box next to the event source.
To select all the available event sources, select the check box at the top of the column.
- 4 To sort the event sources by **Health**, **Name**, **Collector Plug-in**, **Drop Data**, **Create Date**, and **EPS** values, click the column header. The selected column header is displayed in bold.

When you first click a column header, the event sources are arranged in ascending order. A blue down-arrow is displayed to indicate that the sort order is ascending. When you click the column header for the second time, the sort order is changed to descending, and a blue up-arrow is displayed to indicate that the sort order is descending.

- 5 To view additional information about an event source, click the **Name** or **EPS** value of an event source. A dialog box displays the additional information.

Viewing Collector Managers

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Collector Manager section is displayed in the Event Sources page.

Health: Indicates the health of the Collector Managers. You can sort the Collector Managers based on their health status.

Name: Displays the names of the Collector Managers. You can sort the Collector Managers in alphabetical order based on their names.

EPS: Displays the events per second value received from the event sources. You can sort the Collector Managers based on the events per second value.

- 3 To select or deselect a Collector Manager, select the check box next to the Collector Manager.

To select all the available Collector Managers, select the check box located at the top of the column.

The right pane displays the list of event sources connected to the selected Collector Managers.

If none of the Collector Managers are selected, the event sources table displays all the configured event sources.

- 4 To sort the Collector Managers by **Health**, **Name**, and **EPS** values, click the column header. The selected column header displays in bold text.
- 5 To get additional information about the Collector Managers, click the **Name** or **EPS** value column. A dialog box displays the additional information.

Viewing Event Source Servers

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Event Source Servers section is displayed.

Health: Indicates the health of the Event Source Server. You can sort the Event Source Servers based on their health status.

Name: Displays the names of the Event Source Server used to parse the data from the event sources (for example, Syslog Server SSL). You can sort the event source server in alphabetical order based on their names.

EPS: Displays the events per second value received from the event sources. You can sort the event source servers based on the events per second value.

- 3 To sort the Event Source Servers by **Health**, **Name**, and **EPS** values, click the column header. The selected column header displays in bold text.
- 4 To view additional details, click the **Name** or **EPS** value column. A dialog box displays the additional information.

Viewing Collector Plug-Ins

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

Health: Indicates the aggregate health of all event sources that are connected to the Collector plug-in.

With the exception of the green icon (healthy state), the icon does not necessarily mean that all event sources connected to the Collector plug-in are in the state indicated by the icon.

The red icon (error state) indicates that one or more event sources connected to the Collector plug-in are in an error state. To get a detailed information, click the **Name** or **EPS** column value to view help information.

Name: Displays the names of the Collector plug-in used to parse the data from the event sources (for example, Cisco Firewall 6.1r1).

This is the name of the Collector plug-in, not the name of the Collector instance. You can sort the event sources based on Collector plug-in name.

EPS: Displays the events per second value received from the event sources. You can sort the Collector based on the events per second value.

- 3 To select or deselect the Collector plug-ins, select the check box next to the Collector plug-in.
To select all the available Collector plug-ins, select the check box at the top of the column.
- 4 To sort the Collector plug-ins by **Name** or **EPS** values, click the appropriate column header. The selected column header displays in bold text.

The **Collector Instances** field displays the number of instances of the Collector plug-in. Clicking the **Collector Instances** field displays a **Collectors** window with a list of Collector instances associated with the Collector plug-in:
- 5 Click the **Collector Plug-in** column to display a dialog box with additional information about the Collector plug-in.

7.5.2 Changing the Data Logging Status of Event Sources

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.
- 3 To change the data logging status for one or more event sources, select the event sources from the list.
- 4 Click the **Configure** button in the table, then select either **Drop Data** or **Allow Data**,

Drop Data: If **Drop Data** is selected, the selected event sources drop all the events received. Messages are not sent to the Collectors the selected event sources are connected to.

Allow Data: If **Allow Data** is selected, the selected event sources forward events received to the Collectors they are connected to.

If you select a large number of event sources to change, it might take some time to complete. The Event Sources list does not show the Drop state (**YES** or **NO**) until after the changes are complete and the display is refreshed from the database.

8 Configuring Agent-Based Data Collection

Agent Manager provides host-based data collection that complements agentless data collection by allowing you to:

- ♦ Access logs not available from the network
- ♦ Operate in tightly-controlled network environments
- ♦ Improve security posture by limiting attack surface on critical servers
- ♦ Provide enhanced reliability of data collection during times of network interruption

Agent Manager allows you to deploy agents, manage agent configuration, and act as a collection point for events flowing into Sentinel. For more information about Agent Manager, see the Agent Manager documentation.

9 Managing Event Sources

The Event Sources interface displays the health of the event source and the volume of data being received from it in events per second. The event sources page lists all the event sources, such as Syslog, Audit, File, and Database, which are configured in the Event Source Management interface.

You can refine the displayed event sources by selecting Collector Managers, Event Source Servers, and Collector plug-ins. You can also specify a filter on the event source name and select particular event source health states you want to view. All of these selections and filters are stored on a per-user basis, so that each time you log into the Sentinel server you can view event sources that match your last selections. You can also perform filtering based on tags. For more information, see [“Configuring Tags”](#) in the *NetIQ Sentinel User Guide*.

- ♦ [Section 9.1, “Viewing the Event Sources Page,”](#) on page 111
- ♦ [Section 9.2, “Filtering Event Sources,”](#) on page 115

9.1 Viewing the Event Sources Page

The Event Sources page consists of different sections that allow you to perform different functions.

Collector Managers: Lists all the Collector Managers associated with the Sentinel system. It also displays the state and details about the Collector Managers.

Event Source Servers: Lists all the Event Source Servers associated with the Sentinel system. It also displays the state of the Event Source Servers.

Collector Plug-ins: Lists all the Collector plug-ins associated with the Sentinel system. You can also view the details about the installed plug-ins.

The Event Sources section in the right pane lists the event sources based on the options selected from the left pane.

NOTE: The Event Sources page shows event sources that were already configured or automatically detected. To manually configure additional event sources, use the Event Source Management user interface described in [“Configuring Data Collection for Other Event Sources”](#) on page 87.

- ♦ [Section 9.1.1, “Viewing Event Sources,”](#) on page 112
- ♦ [Section 9.1.2, “Configuring Event Sources,”](#) on page 113
- ♦ [Section 9.1.3, “Viewing Collector Managers,”](#) on page 113
- ♦ [Section 9.1.4, “Viewing Event Source Servers,”](#) on page 114
- ♦ [Section 9.1.5, “Viewing Collector Plug-Ins,”](#) on page 114

9.1.1 Viewing Event Sources

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Event Sources page is displayed.

Each column in the Event Source section has different information:

Health Icon: The colored icon indicates the event source health.

- ♦ **Green:** Indicates that the event source is healthy and Sentinel has received data from it.
- ♦ **Red:** Indicates that the Sentinel server is reporting an error about connecting to or receiving data from this event source.
- ♦ **Gray:** Indicates that the event source is turned off. Sentinel is not processing any data from it.
- ♦ **Orange:** Indicates that the event source is running with some warnings.

You can sort the event sources based on their health status.

Name: Displays the name given to the Event Source by the system (if it was auto-created) or by a user. For Syslog Event Sources, if the Event Source was auto-created by the system, the name is a combination of the hostname/IP address and the Collector connection mode the event source is using.

You can rename any Event Source at any time through the Event Source Management interface.

You can sort the Event Sources in alphabetical order based on their names.

Collector Plug-in: Displays the name of the Collector plug-in that the event source is connected to.

This is the name of the Collector plug-in, not the name of the Collector instance. You can sort the event sources based on Collector plug-in name.

Store raw data: Indicates whether Sentinel stores raw data from the associated event source.

- ♦ **Yes:** Sentinel stores all data received from the event source regardless of filter set on the event source.
- ♦ **No:** Sentinel does not store the data received from the event source and does not generate events.

You can sort the event sources based on the **Store raw data** status.

Parse: Indicates whether Sentinel parses the data received from the event source.

- ♦ **All:** Sentinel parses all the data received from the event source.
- ♦ **Filtered:** Sentinel parses only the filtered data received from the event source.
- ♦ **None:** Sentinel does not parse the data received from the event source.

NOTE: If the **Store raw data** option is set to No, Sentinel does not parse the data.

Create Date: Specifies the date and time when the event source was created. You can sort the event sources based on when they were created.

EPS: Displays the events per second value received from the event source. You can sort the event sources based on their events per second value.

If you see a value of less than one (<1) in this column, it indicates that the EPS rate is greater than zero, but less than one.

- 3 To select or deselect an event source, select the check box next to the event source.
To select all the available event sources, select the check box at the top of the column.

- 4 To sort the event sources by **Health**, **Name**, **Collector Plug-in**, **Drop Data**, **Create Date**, and **EPS** values, click the column header. The selected column header is displayed in bold.

When you first click a column header, the event sources are arranged in ascending order. A blue down-arrow is displayed to indicate that the sort order is ascending. When you click the column header for the second time, the sort order is changed to descending, and a blue up-arrow is displayed to indicate that the sort order is descending.

- 5 To view additional information about an event source, click the **Name** or **EPS** value of an event source. A dialog box displays the additional information.

9.1.2 Configuring Event Sources

- 1 Log in to the Sentinel Web Console as an Administrator, Select **Collection > Event Sources**.
- 2 In the **Event Sources** tab, select one or more event source.

NOTE: If you select multiple event sources, the settings apply to all the selected event sources.

- 3 Click the **Settings** icon.
 - ♦ **Start:** Sentinel starts collecting raw data received from the event sources. Sentinel starts only the event sources that are in the stopped state. If the event sources are already in the start state, they remain unchanged.
 - ♦ **Stop:** Sentinel stops collecting raw data received from the event sources. Sentinel stops only the event sources that are in the start state. If the event sources already in the stopped state, they remain unchanged.
 - ♦ **Delete:** Deletes the selected event sources.
 - ♦ **Collector Plug-in:** Select the Collector plug-in to connect to the event sources.
 - ♦ **Tags:** Select the tags to set on the event sources.
 - ♦ **Configure:** Select Configure to set the following options in the **Configure Event Sources** window:
 - ♦ **No data alert:** Click **Edit** to configure notifications if Sentinel does not receive data from the event sources.
 - ♦ **Alert if no data received in specified time period:** Select this option to receive notifications if no data is received during the specified time period.
 - ♦ **Send repeated alerts every time period:** Select this option to receive repeated notifications at every time interval (specified in the **time period**).
 - ♦ **Time zone:** Select the time zone for the event source.
 - ♦ **Trust event source time** (Optional) Select this option to set the event time to the time the event occurred, rather than the time Sentinel received the data.
 - ♦ **Store raw data:** Select Yes to store the raw data from the event sources. If you select No, Sentinel does not store the raw data.

Click **Save** to apply the selected settings.

9.1.3 Viewing Collector Managers

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Collector Manager section is displayed in the Event Sources page.

Health: Indicates the health of the Collector Managers. You can sort the Collector Managers based on their health status.

Name: Displays the names of the Collector Managers. You can sort the Collector Managers in alphabetical order based on their names.

EPS: Displays the events per second value received from the event sources. You can sort the Collector Managers based on the events per second value.

- 3 To select or deselect a Collector Manager, select the check box next to the Collector Manager.

To select all the available Collector Managers, select the check box located at the top of the column.

The right pane displays the list of event sources connected to the selected Collector Managers.

If none of the Collector Managers are selected, the event sources table displays all the configured event sources.

- 4 To sort the Collector Managers by **Health**, **Name**, and **EPS** values, click the column header. The selected column header displays in bold text.
- 5 To get additional information about the Collector Managers, click the **Name** or **EPS** value column. A dialog box displays the additional information.

9.1.4 Viewing Event Source Servers

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

The Event Source Servers section is displayed.

Health: Indicates the health of the Event Source Server. You can sort the Event Source Servers based on their health status.

Name: Displays the names of the Event Source Server used to parse the data from the event sources (for example, Syslog Server SSL). You can sort the event source server in alphabetical order based on their names.

EPS: Displays the events per second value received from the event sources. You can sort the event source servers based on the events per second value.

- 3 To sort the Event Source Servers by **Health**, **Name**, and **EPS** values, click the column header. The selected column header displays in bold text.
- 4 To view additional details, click the **Name** or **EPS** value column. A dialog box displays the additional information.

9.1.5 Viewing Collector Plug-Ins

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.

Health: Indicates the aggregate health of all event sources that are connected to the Collector plug-in.

With the exception of the green icon (healthy state), the icon does not necessarily mean that all event sources connected to the Collector plug-in are in the state indicated by the icon.

The red icon (error state) indicates that one or more event sources connected to the Collector plug-in are in an error state. To get a detailed information, click the **Name** or **EPS** column value to view help information.

Name: Displays the names of the Collector plug-in used to parse the data from the event sources (for example, Cisco Firewall 6.1r1).

This is the name of the Collector plug-in, not the name of the Collector instance. You can sort the event sources based on Collector plug-in name.

EPS: Displays the events per second value received from the event sources. You can sort the Collector based on the events per second value.

- 3 To select or deselect the Collector plug-ins, select the check box next to the Collector plug-in.

To select all the available Collector plug-ins, select the check box at the top of the column.

- 4 To sort the Collector plug-ins by **Name** or **EPS** values, click the appropriate column header. The selected column header displays in bold text.

The **Collector Instances** field displays the number of instances of the Collector plug-in. Clicking the **Collector Instances** field displays a **Collectors** window with a list of Collector instances associated with the Collector plug-in:

- 5 Click the **Collector Plug-in** column to display a dialog box with additional information about the Collector plug-in.

9.2 Filtering Event Sources

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.
- 3 Select the desired criteria to filter event sources.

You can use one or more of the following options to filter the event sources:

- ♦ [Section 9.2.1, “Filtering Event Sources by Name,” on page 115](#)
- ♦ [Section 9.2.2, “Filtering Event Sources by Health Status,” on page 116](#)
- ♦ [Section 9.2.3, “Filtering by Event Sources Event Search Results,” on page 116](#)
- ♦ [Section 9.2.4, “Filtering Event Sources by Collector Managers,” on page 116](#)
- ♦ [Section 9.2.5, “Filtering Event Sources by Event Source Servers,” on page 116](#)
- ♦ [Section 9.2.6, “Filtering Event Sources by Collector Plug-Ins,” on page 117](#)
- ♦ [Section 9.2.7, “Changing the Data Logging Status of Event Sources,” on page 117](#)
- ♦ [Section 9.2.8, “Changing the Associated Collector Plug-In for Event Sources,” on page 117](#)

9.2.1 Filtering Event Sources by Name

To filter the event sources by name, type a name value in the filter text box, then click **Filter**.

Matching is case insensitive. The name value can contain wildcard characters. Use * to match zero or more characters and use ? to match one character. If no wildcard characters are specified in the name value, it is assumed that the name value is intended to mean contains <name value>, or *<name value>.*.

For example, an event source value of abc is interpreted as *abc*. Some examples of common filter types are:

- ♦ If the event source name starts with abc, enter the filter value as abc*.
- ♦ If the event source name ends with abc, enter the filter value as *abc.
- ♦ If the event source name contains abc, enter the filter value as abc or *abc*.

The Event Source table displays the list of event sources whose names match the value entered in the filter input box.

9.2.2 Filtering Event Sources by Health Status

To view the event sources based on the health status, select the **Healthy**, **Warning**, **Error**, or **Offline** check boxes.

The Event Source table displays the list of event sources with the selected health states.

If none of the health states are selected, health state filtering is not performed. It is essentially equivalent to selecting all four health states.

In the Event Source section, click the **Next**, **Previous**, **First**, and **Last** arrow links to scroll through all the event sources. The event source section displays 30 Event Sources per page.

9.2.3 Filtering by Event Sources Event Search Results

To view the event search result for an event source, select the event source from the list and click the **Search** link.

A search is performed using the universally unique identifier (UUID) of the event source (for example, `rv24 : "2CBFB8A0-F24B-102C-A498-000C"`).

If multiple event sources are selected for search, the `rv24 : <UUID>` expressions are combined with the OR operator in the search filter expression.

9.2.4 Filtering Event Sources by Collector Managers

To display the event sources connected to particular Collector Managers, select one or more Collector Managers from the Collector Managers section.

If none of the Collector Managers are selected, event source filtering is not performed based on the Collector Managers. This is not the same as selecting all Collector Managers, because it also includes event sources that are not connected to any Collector Manager.

To select or deselect Event Source Servers, select the check boxes next to the Event Source Servers.

9.2.5 Filtering Event Sources by Event Source Servers

To display only event sources connected to particular Event Source Servers, select one or more Event Source servers from the Event Source Servers section.

If none of the Event Source Servers are selected, event source filtering is not performed based on the Event Source servers. This is not the same as selecting all Event Source Servers, because it also includes event sources that are not connected to any Event Source Server.

To select or deselect Event Source Servers, select the check boxes next to the Event Source Servers.

9.2.6 Filtering Event Sources by Collector Plug-Ins

To display only those event sources connected to particular Collector plug-ins, select one or more Collector plug-ins from the **Collectors Plug-ins** section.

If none of the Collector plug-ins are selected, event source filtering is not performed based on the Collector plug-in. It is essentially equivalent to selecting all of the Collector plug-ins.

9.2.7 Changing the Data Logging Status of Event Sources

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.
- 3 To change the data logging status for one or more event sources, select the event sources from the list.
- 4 Click the **Configure** button in the table, then select either **Drop Data** or **Allow Data**.
Drop Data: If **Drop Data** is selected, the selected event sources drop all the events received. Messages are not sent to the Collectors the selected event sources are connected to.
Allow Data: If **Allow Data** is selected, the selected event sources forward events received to the Collectors they are connected to.

If you select a large number of event sources to change, it might take some time to complete. The Event Sources list does not show the Drop state (**YES** or **NO**) until after the changes are complete and the display is refreshed from the database.

9.2.8 Changing the Associated Collector Plug-In for Event Sources

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection** in the toolbar, then click the **Event Sources** tab.
- 3 Select the event sources from the list, then click the **Configure** button in the toolbar.
- 4 Select the **Collector Plug-in** option.
The Set Collector Plug-in window is displayed with the **Collector Plug-in Name** and **Supported Devices** information.
- 5 Select a new Collector plug-in, then click **Set**.

The event sources are connected to the selected Collector plug-in.

If you select a large number of event sources to change, it might take some time to complete. The Event Sources list does not show the new Collector plug-in until after the changes are complete and the display is refreshed from the database.

10 Configuring Event Routing Rules

You can configure event routing rules to evaluate and filter all incoming events and deliver selected events to designated output actions. For example, each severity 5 event can be logged to a file.

You can configure event routing rules to filter events based on one or more of the searchable fields. You can associate each event routing rule with one or more of the configured actions. You can also assign tags to group the events logically.

Sentinel evaluates the event routing rules on a first-match basis in top-down order and applies the first matched event routing rule to events that match the filter criteria.

- ♦ [Section 10.1, “Creating an Event Routing Rule,” on page 119](#)
- ♦ [Section 10.2, “Ordering Event Routing Rules,” on page 120](#)
- ♦ [Section 10.3, “Activating or Deactivating an Event Routing Rule,” on page 121](#)

10.1 Creating an Event Routing Rule

You can create a filter-based event routing rule and then assign one or more configured actions that are executed to handle or output the events that meet the event routing rule criteria.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Routing** in the toolbar.
- 3 Click **Create**, then use the following information to create a new event routing rule:

Name: Specify a unique name for the event routing rule.

Filter: Select a saved filter to use in creating event routing rule. This filter determines which events are stored in the event store. For more information, see [“Configuring Filters”](#) in the *NetIQ Sentinel User Guide*.

Select tag: (Optional) Select a tag for tagging the filter. The tag makes the filter more specific. For more information, see [“Configuring Tags”](#) in the *NetIQ Sentinel User Guide*.

Route to the following services: Select where the information is routed. The options are:

- ♦ **All:** Routes the event to all services including Correlation, Security Intelligence, and Anomaly Detection.
- ♦ **Event store only:** Routes the event to the event store only.
- ♦ **None (drop):** Drops or ignores the events.

Perform the following actions: (Optional) Select an action to be performed on every event that meets the filter criteria. The following default actions are available for event routing rules:

- ♦ **Log to File:** For configuration information, see [“Configuring the File Integrator” on page 148](#).
- ♦ **Log to Syslog:** For configuration information, see [“Configuring the Syslog Integrator” on page 151](#)



- ♦ **Send Events via Sentinel Link:** For configuration information, see [“Configuring the Sentinel Link Integrator” on page 148.](#)
- ♦ **Send SNMP Trap:** For configuration information, see [“Configuring the SNMP Integrator” on page 150.](#)

NOTE: When you associate an action with routing rules, ensure that you write rules that match a small percentage of events, if the rule triggers a Javascript action. If the rules trigger actions frequently, the system might backlog the actions framework. This can slow down the EPS and might affect the performance of the Sentinel system. If the rule triggers non-Javascript actions like Sentinel Link, then there is no limitation.

For the actions to work, you must have configured the Integrator associated with each action for your environment.

The actions listed here are different than the actions displayed in the **Event Actions** tab in the Sentinel Web interface, and are distinguished by the `<EventRouting>` attribute in the `package.xml` file created by the developer.

Adding or Removing Actions: You can add more than one action to perform on the events that meet the filter criteria:

- ♦ Click  to select additional actions to be performed.
- ♦ Click  to remove the selected action for this event routing rule.

- 4 Click **Save** to save the event routing rule.

The newly created event routing rule appears at the end of the rules list under the **Event Routing Rules** tab. By default, this new event routing rule is active.

10.2 Ordering Event Routing Rules

When there is more than one event routing rule, the event routing rules can be reordered by dragging them to a new location. Events are evaluated by event routing rules in the specified order until a match is made, so you should order the event routing rules accordingly. More narrowly defined event routing rules and more important event routing rules should be placed at the beginning of the list.

The first routing rule that matches the event based on the filter is processed. For example, if an event passes the filter for two routing rules, only the first rule is applied. The default routing rule cannot be reordered. It always appears at the end.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Routing** in the toolbar.
The **Event Routing Rules** tab is displayed.
Existing event routing rules appear on the page.
- 3 Mouse over the icon to the left of the event routing rule numbering to enable drag-and-drop.
The cursor changes.
- 4 Drag the event routing rule to the correct place in the ordered list.
When the event routing rules are ordered, a success message is displayed.

10.3 Activating or Deactivating an Event Routing Rule

New event routing rules are activated by default. If you deactivate an event routing rule, incoming events are no longer evaluated according to that event routing rule. If there are already events in the queue for one or more actions, it might take some time to clear the queue after the event routing rule is deactivated. If the **On** check box next to the event routing rule is selected, the event routing rule is activated. If the **On** check box is not selected, the event routing rule is deactivated.

- 1 Log in to the Sentinel Web interface as a user in the administrator role.

- 2 Click **Routing** in the toolbar.

The **Event Routing Rules** tab is displayed.

Existing event routing rules appear on the page.

- 3 To activate the event routing rule, select the check box next to each event routing rule in the **Enabled** column.

If the event routing rule is activated, a success message is displayed.

- 4 To deactivate the event routing rule, select the check box next to each event routing rule in the **Enabled** column.

When the event routing rule is deactivated, a success message is displayed.

11 Mapping Events

Sentinel provides the ability to use mapping to inject additional information into events. This increases Sentinel's ability to analyze events, execute correlation rules, or provide detailed reports.

- ♦ [Section 11.1, "Overview," on page 123](#)
- ♦ [Section 11.2, "Default Maps," on page 126](#)
- ♦ [Section 11.3, "Accessing Map Definitions," on page 127](#)
- ♦ [Section 11.4, "Adding Map Definitions," on page 127](#)
- ♦ [Section 11.5, "Adding a Number Range Map Definition," on page 128](#)
- ♦ [Section 11.6, "Updating Map Data," on page 131](#)
- ♦ [Section 11.7, "Using Maps for Event Configuration," on page 132](#)
- ♦ [Section 11.8, "Renaming Event Fields," on page 133](#)

11.1 Overview

There are two major components of mapping:

- ♦ [Section 11.1.1, "Maps," on page 123](#)
- ♦ [Section 11.1.2, "Mapping Events," on page 124](#)

11.1.1 Maps

A map is a collection of keys and associated values defined in a simple text CSV (comma-separated value) file format. You can enrich your event data by using maps to add additional information such as host and identity details to the incoming events from your source devices. This additional information can be used for advanced correlation and reporting. The system supports several built-in maps as well as custom user-defined maps. For more information, see [Section 11.7, "Using Maps for Event Configuration," on page 132](#).

Maps that are defined in Sentinel are stored in two ways:

- ♦ Built-in maps are stored in the database, updated via APIs in Collector code, and automatically exported the Mapping service.
- ♦ Custom maps are stored as CSV files and can be updated on the file system or via the Map Data Configuration UI, then loaded by the Mapping service.

In both cases, the CSV files are kept on the central Sentinel server but changes to the maps are distributed to each Collector Manager and applied locally. This distributed processing ensures that mapping activity does not overload the main server.

11.1.2 Mapping Events

Event Mapping is a mechanism that allows you to add data to an event by using data already in the event to reference and pull in data from an outside source.

Because virtually any data set can be made into a map, Event Mapping is useful for incorporating data from elsewhere in your organization into the event stream. Some opportunities that Event Mapping provides are:

- ♦ Regulatory compliance monitoring
- ♦ Policy compliance
- ♦ Response prioritization
- ♦ Enabling security data to be analyzed related to business operations
- ♦ Enhancing accountability

When an Event Mapping is defined, it is applied system-wide to all events from all Collectors. Sentinel automatically distributes map data to all processes that perform event mappings and also keeps the map data in these processes up-to-date.

Event Mapping has four main parts:

- ♦ **Controller:** Stores all map information
- ♦ **Distributor:** Automatically redistributes modified maps to those processes that register for the map
- ♦ **Monitor:** A monitor to detect changes in map source data
- ♦ **Generator:** Generates maps from source data

One application of Event Mapping is Sentinel's Asset Data functionality. Asset information is collected and stored in the Sentinel database asset schema and is represented by a Physical Asset entry. Soft assets, such as services and applications, are represented by an entry that is linked to a Physical Asset. The primary automated update mechanism for asset data is through an asset Collector reading data from a scanner such as Nmap. The asset Collector automates the retrieval of asset information by reading asset data from the scanner and populating the asset schema tables with this data. For Event Mapping, asset information is mapped from the destination IP and source IP.

There are two types of data sources:

- ♦ **External:** A Collector populates that value in the event ID.
- ♦ **Referenced from Map:** Data is retrieved from a map to populate the event ID.

Figure 11-1 Data Sources

Map Key Field	Event Tag
PhysicalAssetName	SourceIP

In the above illustration, the SourceAssetName tag is populated from the map called Asset (which has asset.csv as its map data source file). The specific value for SourceAssetName is taken from the AssetName column from the Asset map. The PhysicalAssetName column is set as the key. When the InitIP tag of the event matches one of the source IP values in the PhysicalAssetName column of the map, the row with the matching key is used to intersect the AssetName Column. For instance, in the following example the IP address corresponds to AssetName Finance35.

Table 11-1 Populating an Event ID with Data from a Map

PhysicalAssetName	CustomerID	MacAddress	AssetName
10.0.0.1			Marketing0
10.0.0.2			Marketing02
10.0.0.3			ProgramMgmt03
10.0.0.4			Finance34
10.0.0.5			Finance35

You can have more than one column set as a key if you do not want the map to be a Range Map (Range Maps can only have one key column, with that column type set to NumberRange). For instance, with column type set to String, the AttackId tag has the DeviceName (name of the security device) and DeviceAttackName columns set as keys and uses the NormalizedAttackID column in the AttackNormalization map for its value. In a row where the DeviceName event ID matches the data in Device map column and the DeviceAttackName matches the data in the AttackSignature map column, the value for AttackId is the value in the NormalizedAttackID column. The configuration for the Event Mapping just described is:

Figure 11-2 Event Mapping Configuration

Table 11-2 Device and Attack Signature Corresponds to Asset Name

Device	Attack Signature	Normalized Attack ID	
Secure	BackDoorProbe (TCP 1234)	3	Trojan: Backdoor.Subeven
Secure	BackDoorProbe (TCP 1999)	3	Trojan: Backdoor.Subeven
Dragon	RWALLD: SYSLOG-FORMAT	4	Sun Microsystems Solaris rwall Elevated

Device	Attack Signature	Normalized Attack ID	
Snort	RPC TCP rwall request	4	Sun Microsystems Solaris rwall Elevated
Snort	RPC UDP rwall request	4	Sun Microsystems Solaris rwall Elevated
Snort	WEB-IIS foxweb.dll access	12	Microsoft Exchange Server Arbitrary Code
RealSecure	SMTP_Exchange_Verb_DoS	12	Microsoft Exchange Server Arbitrary Code

11.2 Default Maps

Maps defined in this tool work together with the **Referenced from Map** data source setting for individual fields. The following built-in maps are available:

- ♦ **Identity:** Contains information about identities and the accounts associated with them. Data is added to the Identity map through the [Collector API \(http://www.novell.com/developer/develop_to_sentinel.html\)](http://www.novell.com/developer/develop_to_sentinel.html) and Identity Tracking Module for Sentinel. Data is then extracted to the `identityAccountMap.csv` file.
 - ♦ **Keys:** User name and domain, TenantName (mapped to both Initiator and Target users).
 - ♦ **Data added to event:** User identity (an internal GUID), full name, department, workforce ID, and e-mail address.
- ♦ **Asset:** Contains information about hosts in the environment. Data is added to the Asset map using a [Collector API \(http://www.novell.com/developer/develop_to_sentinel.html\)](http://www.novell.com/developer/develop_to_sentinel.html) and Collectors such as the Generic Asset Collector, stored in the database, and then extracted to the `asset.csv` file.
 - ♦ **Keys:** IP address and TenantName (mapped to Initiator, Target, Observer, and Reporter hosts).
 - ♦ **Data added to event:** Host identity (an internal GUID), function, class, department, and criticality.
- ♦ **Country:** Contains information about which physical location hosts reside in, including country, latitude, and longitude. Data is downloaded from a commercial IP location database and added to the `IpToCountry.csv` file using the Generic IP Geolocation Collector.
 - ♦ **Keys:** IP address and TenantName (mapped to Initiator, Target, and Observer hosts).
 - ♦ **Data added to event:** Country, latitude, and longitude.
- ♦ **Exploit:** Contains information about vulnerabilities present in and attacks on enterprise hosts. Information is added to the Exploit map by the Advisor service, stored in the database, and extracted to the `exploitDetection.csv` file.
 - ♦ **Keys:** Target IP address, IDSAttackName, IDSName, and TenantName.
 - ♦ **Data added to event:** Vulnerability flag.
- ♦ **CustomerHierarchy:** Contains a hierarchical list of tenants that are generating event data. This can be used by security providers that collect data for multiple third parties or departments to provide a hierarchic namespace for users and hosts. Data is added to the `customerhierachy.csv` file manually.
 - ♦ **Keys:** TenantName.
 - ♦ **Data added to event:** TenantHierarchy fields.

11.3 Accessing Map Definitions

- 1 Access the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.

- 2 Click the **Configuration** tab.
- 3 In the menu bar, click **Configuration > Map Data Configuration**.

The main Mapping GUI displays a list of the maps that have been defined for the system. Default Sentinel maps cannot be edited or deleted.

11.4 Adding Map Definitions

- 1 Access the map definitions.

For more information, see [Section 11.3, “Accessing Map Definitions,”](#) on page 127.

- 2 Click **Add**, then use the following information to create the map definition:

Name: Specify the name of the map definition.

File Name: Select whether the file is local or remote, then browse to and select your map definition.

- ♦ **Local File:** Allows you to browse for your file on your local file system (on the machine where Sentinel Control Center was launched).
- ♦ **Remote File:** Allows you to select from existing map source data files on the Sentinel server. The remote file points to `/var/opt/novell/sentinel/data/map_data`.

Map Definition: Use the following information to define the map. As you configure each setting and filter, the data preview is automatically updated to allow you to preview your data and ensure your data is being parsed as expected.

- ♦ **Delimiter:** Character used to separate the data into rows in the map data source file. Usually a comma, but other delimiters such as pipe, tab, and semicolon are supported. You can also specify other delimiters in the **Other** field.
- ♦ **Start at row:** Some input map files contain header rows that do not contain real data. This option specifies the number of rows to skip from the top of the map data source file.
- ♦ **Column names:** Specify the column names. These names are used later to configure which columns are matched against event fields and which columns are injected into event data.
- ♦ **Key columns:** A key is a unique identifier for the row of data in the map data. If more than one column is selected as a key, then all event fields must match each of the corresponding selected key columns.

When a column is set as a key, it does not appear in the **Column** drop-down field.

- ♦ **Column types:** Specify the data type for the column. The currently supported data types are:
 - ♦ **String:** A string is treated as a simple sequence of characters, and can include any characters except the specified delimiter. Use strings for any data including numeric fields.
 - ♦ **Number Range:** A number range (NumberRange) is a range of numbers. For example, 10 through 200 is represented as 10-200. To use the range map functionality, a map definition must have only one key column and the key column must be of type NumberRange. If there are any other key columns, or if the key column is of a different type, the mapping service does not consider the map to be a range map. For more information, see [Section 11.5, “Adding a Number Range Map Definition,”](#) on page 128.

- ♦ **Active columns:** When a column is marked as active, the data in the column is distributed to processes using maps. All key columns must be active. Only active columns (but not key columns) can be selected as the Map Column under the **Event Configuration** tab.
 - ♦ **Column filtering:** A row can be explicitly included or excluded based on matching criteria for a particular column. This can be used to exclude rows from the map source data that are not needed or that interfere with your mapping.
- 3 After you finish configuring all parameters and filters for the definition, click **Finish**.
 - 4 (Conditional) If you selected **Local File** above, you are prompted to upload your file. Specify a file name, then click **OK**.

You can have more than one column set as a key if you do not want the map to be a Range Map (Range Maps can only have one key column, with that column type set to NumberRange). For instance, with column type set to String, the AttackId tag has the DeviceName (name of the security device) and DeviceAttackName columns set as keys and uses the NormalizedAttackID column in the AttackNormalization map for its value. In a row where the DeviceName event ID matches the data in Device map column and the DeviceAttackName matches the data in the AttackSignature map column, the value for AttackId is the value in the NormalizedAttackID column.

11.5 Adding a Number Range Map Definition

To use the range map functionality, a map definition can have one key column of type NumberRange and zero or more other key columns of type String. For example, you can create a map that lists the allowed maintenance time for each individual server in an enterprise, for which you need to match both the hostname and the time range. If there are any other key columns, or if the key column is of a different type, the mapping service does not consider the map to be a range map.

To create a range map, select a single column to be the key of the map and select NumberRange as the type of the column. The format of the data in a column of type NumberRange must be “m-n”, where m is the minimum number in the range and n is the maximum number in the range (for example, 10-200). The maximum number in the range is not included in the range (m, n). This means a range of 10-200 only uses numbers equal to 10 to 199. An example set of data has the first column defined as the key column:

```
1-2, AA
2-4, AA
4-12, BB
10-20, BB
30-31, BB
100-200, AA
110-120, CC
```

When the source CSV file is loaded into the system, any common or overlapping number ranges are collapsed into a single entry as follows:

Table 11-3 Transformation

FROM	TO
1-2, AA	1-4, AA
2-4, AA	4-20, BB
4-12, BB	30-31, BB
10-20, BB	100-110, AA
30-31, BB	110-120, CC
100-200, AA	120-200, AA
110-120, CC	

Figure 11-3 Number Range Map Definition

The first 500 rows are shown

	Column 1	Column 2
Name:	Range	Value
Type:	NumberRange	String
Key:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Active:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Row 0	1-4	AA
Row 1	4-20	BB
Row 2	30-31	BB
Row 3	100-110	AA
Row 4	110-120	CC
Row 5	120-200	AA

An example event configuration on the above map might look like this:

Figure 11-4 Event Configuration

CustomerVar82	Data Source <input type="radio"/> External <input checked="" type="radio"/> Referenced from Map Map Name: <input type="text" value="Maps/RangeMap"/> Map Column: <input type="text" value="Value"/> Key Configuration: <table border="1"> <tr> <th>Map Key Field</th> <th>Event Tag</th> </tr> <tr> <td>Range</td> <td>CustomerVar97</td> </tr> </table>	Map Key Field	Event Tag	Range	CustomerVar97
Map Key Field		Event Tag			
Range		CustomerVar97			
CustomerVar83					
CustomerVar84					
CustomerVar85					
CustomerVar86					
CustomerVar87					
CustomerVar88					
CustomerVar89					
SARBOX					
HIPAA					
GLBA					
FISMA					

In the scenario above, CustomerVar97 must contain a numeric value. That value is compared against each NumberRange defined in the RangeMap until a match is found. The corresponding row from the map is returned and used to set CustomerVar89, as follows:

CustomerVar97 = 1; CustomerVar89 will be set to AA
CustomerVar97 = 4; CustomerVar89 will be set to BB
CustomerVar97 = 300; CustomerVar89 will not be set

For Sentinel event fields that are defined as having an IP address or Date datatype, Sentinel internally converts those fields to an integer representation of the values of that field.

Fields that are defined as IP address fields are:

- ♦ SourceIP (sip)
- ♦ TargetIP (dip)
- ♦ SourceTranslateIP (sxiip)
- ♦ TargetTranslateIP (dxiip)
- ♦ ObserverIP (obsip)

Fields that are defined as Data fields are:

- ♦ EventTime (dt)
- ♦ ObserverEventTime (det)
- ♦ SentinelProcessTime (spt)
- ♦ BeginTime (bngt)
- ♦ EndTime (endt)
- ♦ CustomerVar11 to CustomerVar20 (cv11 to cv20)
- ♦ ReservedVar11 to ReservedVar20 (rv11 to rv20)

The conversions for IP address are simple octet mappings to decimal integers, which you can calculate using any server online calculator. The following example shows a numerical range equivalent to an IP range of 10.0.0.0 to 10.0.2.255.

167772160-167772415,AAA
167772416-167772671,BBB
167772672-167772927,CCC

Using the same setup as the previous example, if:

- ♦ In Key Configuration, the Event ID used for comparison is set to TargetIP for the range Map Key field.
- ♦ The Map column returned from the map to set CustomerVar89 is defined as a value, as displayed in the second column below.

If an event contains a target IP of 10.0.1.14 (equivalent to numerical value of 167772430), the output for column CustomerVar89 within the event is BBB.

Dates are represented as an integer number of seconds since midnight January 1, 1970. Data and time ranges can be used in maps in a similar fashion as the IP address sample above.

Sentinel supports the following number ranges:

- ♦ Range from negative number to negative number (for example, -234-34)
- ♦ Range from negative number to positive number (for example, -234-34)
- ♦ Range from positive number to positive number (for example, 234-236)
- ♦ Single number range (negative) (for example, -234). In this case, the min and the max will both be -234.
- ♦ Single number range (positive) (for example, 234). In this case, the min and the max will both be 234.

- ♦ Range from negative number to max number (for example, -234-). In this case, the min will be -234 and the max will be $(2^{63} - 1)$.
- ♦ Range from positive number to max number (for example, 234-). In this case, the min will be 234 and the max will be $(2^{63} - 1)$.

NOTE: In all cases, the min must be less than or equal to the max (for example, “-234- -235” is NOT valid).

11.6 Updating Map Data

Updating allows you to replace the map source data file of a map on the server with another file. Your new map source data file must have the same delimiter, number of columns, and overall structure as the existing map data source file in order for the map to function properly after the update. The new map source data file should differ from the existing file only by the values that appear in the columns. If the new map source data file has a different structure than the existing file, use the Edit feature to update the map definition.

Map updates can be performed on demand from the Sentinel Control Center. To set up an automated process to update map data, you can run an equivalent process from the command line using `map_updater.sh`.

There are two map locations: the location referenced by the Event Map Configuration (which is a user-defined location) and the location where Sentinel stores its internal representation of the map (`/var/opt/novell/sentinel/data/map_data`). The internal representation of the map should never be manually updated.

To update the map data from the Sentinel Control Center:

- 1 If you haven't already done so, create a CSV file containing the new map source data.
This file can be generated (for example, from a data dump script), created manually, or be an edited version of the existing map data source file. If necessary, you can obtain the existing map data source file from `/var/opt/novell/sentinel/data/map_data`.
- 2 Access a map definition.
For more information, see [Section 11.3, “Accessing Map Definitions,” on page 127](#).
- 3 Expand the folder of interest and select the mapping, then click **Update**.
- 4 Select the new map data source file by clicking **Browse** and selecting the file with the new map data.
After you select the file, the data from the new map data source file displays under the **New** tab. The map data you are replacing is under the **Current** tab.
- 5 Deselect or leave the default setting for **Backup Existing Data On Server**.
Enabling this option puts a backup of the existing map data source file in the `/var/opt/novell/sentinel/data/map_data` folder. The prefix of the name of the backup map data source file is the name of the existing map data source file. The end of the filename includes a set of random numbers followed by the `.bak` suffix. For example: `vuln_attacks10197.bak`.
- 6 Click **OK**.
The data from the new map data source file is uploaded to the server, replacing the contents of the existing map data source file. After the source data is completely uploaded, the map data is regenerated and distributed to map clients such as, Collector Manager.

11.7 Using Maps for Event Configuration

After you have created a map, you must decide where the mapped data is injected into the event. You must configure each Event to use mapping:

- 1 Access the **Configuration** tab in the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).

- 2 Click **Event Configuration** in the navigation pane.

or

In the toolbar, click **Configuration > Event Configuration**.

- 3 Select an entry from the Event Columns, then use the following information to configure the event using a map:

Name: Displays the name of the event ID selected in the **Event Columns** field.

Referenced from Map: Click **Referenced from Map** to configure the event ID to be populated with data from a map.

The default option of **External** keeps the value the Collector put in the event ID (if any).

Map Column: Click the **Map Name** drop-down list, then select one of the available maps.

The maps listed are the default maps or a map you have created.

Map Column: Click the **Map Column** field drop-down list, then select a **Map Column** name.

Depending on your Map Name choice in the previous step, these values vary.

- ♦ **All other choices:** Names of active columns within the map definition that are not set as a key (for example, `CustomerId` column in `Asset` or `NormalizedAttackId` column in `AttackNormalization`)
- ♦ **_EXIST_:** This is a special Map Column that exists in every map. If this Map Column is selected, a 1 is placed in the event ID if the key is in the map data. If the key is not in the map data, a 0 is placed in the event ID.

Key Configuration: For each row in the table, select the event ID in the **Event ID** column that is matched against the map key column specified in the corresponding **Map Key Field** column. The rows in the Key Configuration table depend on the Map Name selected.

A key is a unique identifier for the row of data in the map data.

- 4 Click **Apply**.

Clicking **Apply** saves the changes you made for the currently selected event column in a temporary buffer. If you don't click **Apply**, the changes you made to the previously selected event column are lost when you select a different event column. Changes aren't be saved to the server until you click **Save**.

- 5 If you want to edit the event mapping of another event column, repeat the steps above. Remember to click **Apply** after editing the event mapping of each event column.

- 6 Click **Save**.

Clicking **Save** saves your changes to the server. The save function saves all changes stored in the temporary buffer when you clicked **Apply**.

11.8 Renaming Event Fields

The Event Configuration window allows you to rename the event field. When you do this, the event names are changed in the Sentinel Control Center and the Sentinel Web interface.

Renaming event fields does not change the event ID in Collector scripts or in internal Sentinel representations of the event. For example, if you rename the event name `BeginTime` to `StartTime`, the ID will still remain as `bgnt`. Any references to this ID in Correlation or Filters still work, even if they were originally written using `bgnt`.

To rename an event field:

- 1 Access the Configuration tab in the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).
- 2 Click **Event Configuration** in the navigation pane.
or
In the menu bar, click **Configuration > Event Configuration**.
- 3 Select an entry in the **Event Fields** column.
- 4 Specify a new value for the event name in the **Name** field.
- 5 Click **Apply**, then click **Save**.
- 6 Close and launch Sentinel Control Center for the changes to be visible.

12 Linking to Additional Sentinel Systems

Sentinel Link is a mechanism that provides the ability to hierarchically link multiple Sentinel systems, including Sentinel Log Manager, Sentinel, and Sentinel Rapid Deployment. You can hierarchically link two or more Sentinel systems to forward filtered events from one Sentinel system to another for further evaluation.

- ♦ [Section 12.1, “Benefits,” on page 135](#)
- ♦ [Section 12.2, “Supported Platforms,” on page 135](#)
- ♦ [Section 12.3, “Prerequisite,” on page 135](#)
- ♦ [Section 12.4, “Configuring Sentinel Link,” on page 136](#)

12.1 Benefits

- ♦ Multiple Sentinel servers, local or distributed, can be linked in a hierarchical manner. In this setup, Sentinel servers can manage a large volume of data, retaining raw data and event data locally, while forwarding important events to a central Sentinel server for consolidation.
- ♦ One or more Sentinel servers can forward important data to either a Sentinel server, Sentinel Log Manager server, or a Sentinel Rapid Deployment server. These systems provide real-time visualization of data, advanced correlation and actions, workflow management, and integration with identity management systems.
- ♦ Multiple Sentinel, Sentinel Log Manager, or Sentinel Rapid Deployment servers can be hierarchically linked to monitor the consolidated event information.
- ♦ One or more Sentinel, Sentinel Log Manager, or Sentinel Rapid Deployment servers can forward important events to a Sentinel server for event consolidation.

12.2 Supported Platforms

- ♦ Sentinel 7.0
- ♦ Sentinel 6.1 Rapid Deployment Hotfix 2 or later
- ♦ Sentinel 6.1 Service Pack 1 Hotfix 2 or later
- ♦ Sentinel Log Manager 1.2 or later

12.3 Prerequisite

Before you forward events from the sender machine, ensure that the Sentinel Link server is up and running on the receiver machine.

12.4 Configuring Sentinel Link

In a Sentinel Link Solution setup, the Sentinel system that forwards the events is called the sender and the Sentinel system that receives the events is called the receiver. You can simultaneously link multiple Sentinel systems to a single receiver system.

To configure a Sentinel link, you need to configure at least two systems: the sender machine and the receiver machine. For further details on configuring Sentinel Link, see the [Sentinel Link Overview Guide \(http://www.novell.com/documentation/sentinel70/sentinel_link_overview/data/bookinfo.html\)](http://www.novell.com/documentation/sentinel70/sentinel_link_overview/data/bookinfo.html).

IV Integrating with External Systems

This section provides information about integrating Sentinel with external systems.

- ♦ [Chapter 13, “Configuring Actions,” on page 139](#)
- ♦ [Chapter 14, “Configuring Integrators,” on page 147](#)
- ♦ [Chapter 15, “Integrating Identity Information,” on page 155](#)
- ♦ [Chapter 16, “Managing Data Feeds,” on page 159](#)
- ♦ [Chapter 17, “Detecting Vulnerabilities and Exploits,” on page 161](#)

13 Configuring Actions

An Action is a configured instance of an Action plug-in. Actions are used to execute some type of action in Sentinel, either manually or automatically. Actions can be triggered by routing rules, by manually executing an event or incident operation, and by correlation rules.

- ♦ [Section 13.1, “Overview,” on page 139](#)
- ♦ [Section 13.2, “Understanding the Action Manager Interface,” on page 140](#)
- ♦ [Section 13.3, “Managing Actions,” on page 141](#)
- ♦ [Section 13.4, “Managing Action Plug-Ins,” on page 144](#)

13.1 Overview

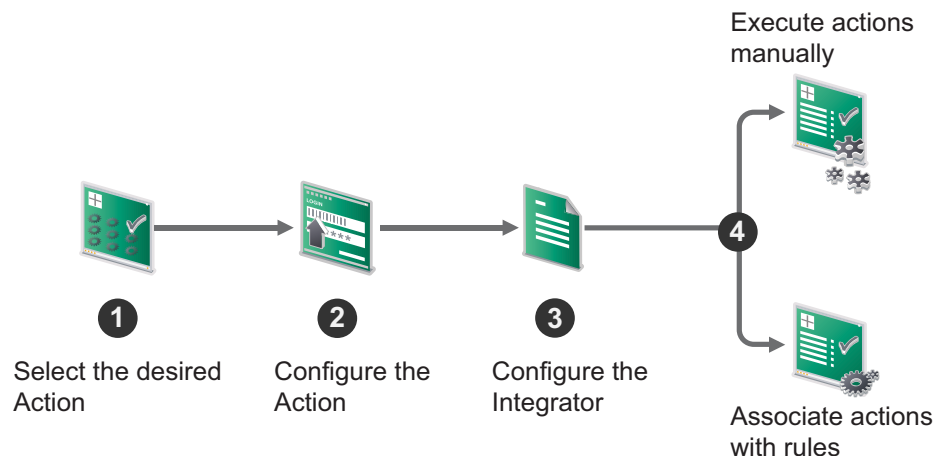
Sentinel provides a list of preconfigured Actions that should be useful in most standard situations. You can use the default Actions and reconfigure them as necessary, or you can add new Actions.

NOTE: Only users in the administrator role can configure and manage Actions.

An Action can be executed on its own, or it can make use of an Integrator instance configured from an Integrator plug-in. Integrators provide the ability to connect to an external system, such as an LDAP, SMTP, or SOAP server, to execute an action.

The general process for using an Action to perform remediation is shown in the following figure:

Figure 13-1 *Actions Workflow*



1. Determine the best type of Action plug-in that should be used to perform your desired action.
2. Configure the appropriate Action plug-in with appropriate parameter settings for your environment.

For more information, see [Section 13.3.1, “Adding an Action,” on page 142](#).

3. If the Action requires an Integrator, configure the appropriate Integrator.


To determine the required Integrators for an Action, see the documentation that is available with the Action on the [Sentinel Plug-ins Web site](#). Alternatively, you can view a specific Action's documentation by clicking the **Help** button while configuring that Action in the Action Manager.

4. For information on configuring the Integrator, see [Section 14.2, “Managing Integrators,” on page 148](#).
5. Execute actions manually or associate actions to rules for the action to fire automatically when the rule fires:
 - ♦ For information on executing an action in an Incident, see “[Executing Incident Actions](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ For information on executing an action on events that meet the event routing rule criteria, see [Section 10.1, “Creating an Event Routing Rule,” on page 119](#).
 - ♦ For information on executing an action on events in Active Views, see “[Executing Actions on Events](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ For information on executing actions on events in Search results, see “[Executing Actions](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ For information on associating an action to a Correlation rule, see “[Associating Actions to a Rule](#)” in the *NetIQ Sentinel User Guide*.

13.2 Understanding the Action Manager Interface

The Action Manager allows you to configure actions that can be executed in various contexts throughout the Sentinel system. You can also use the Action Manager to manage Actions and Action plug-ins

Figure 13-2 Action Manager



The Action Manager window lists the preconfigured actions. You can configure these Actions further as necessary.

- ♦ **Actions:** To configure an Action, select the desired action, then click **View/Edit**.
- ♦ **Add:** Allows you to add an action. For more information, see [Section 13.3.1, “Adding an Action,” on page 142](#).
- ♦ **Manage Plug-ins:** Allows you to import new and updated action plug-ins, and manage the action plug-ins. For more information, see [Section 13.4, “Managing Action Plug-Ins,” on page 144](#).

13.3 Managing Actions

An Action is a configured instance of an Action plug-in. There can be one or more instances of an Action plug-in with different parameters or settings. A few Actions are available by default. You can also add additional actions as required.



- ♦ [Section 13.3.1, “Adding an Action,” on page 142](#)
- ♦ [Section 13.3.2, “Debugging Actions,” on page 143](#)

13.3.1 Adding an Action

- 1 Launch the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).

- 2 Launch the Action Manager:

- ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.
- ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.

- 3 Click  **Add**.

- 4 To create an Action, select an existing Action plug-in from the available action types in the **Action** drop-down. Alternatively, you can import another plug-in by clicking the **Add Action Plug-in** button.

The parameters for the selected plug-in are displayed. For Actions provided by NetIQ, more information about configuration and the available parameters are available in the help file for the Action.

- 5 (Conditional) If the selected Action plug-in requires an Integrator, the **Add Integrator** button is displayed to allow you to add the Integrator for this action. Click **Add Integrator** and select the appropriate Integrator for the action.

- 6 Specify the attribute values for the type of action selected.

You can reference an event field value in an attribute value as `$event_field_name`. For example, in the send Email action, if you want to reference the event field tenant name in the subject attribute, you can specify the parameter `$TenantName` in the subject attribute. Similarly, you can reference the correlation rule metadata by specifying the parameters `%RuleName%`, `%RuleDescription%`, `%RuleId%`, or `%RuleLg%` in actions or in custom correlation events.

- 7 Click **Save**.

- 8 Execute actions manually or associate actions to Correlation rules for the action to fire automatically when the rule fires:

- ♦ For information on executing an action in an Incident, see “[Executing Incident Actions](#)” in the *NetIQ Sentinel User Guide*.
- ♦ For information on executing an action on events that meet the event routing rule criteria, see [Section 10.1, “Creating an Event Routing Rule,” on page 119](#).
- ♦ For information on associating an action to a Correlation rule, see “[Associating Actions to a Rule](#)” in the *NetIQ Sentinel User Guide*.
- ♦ To execute actions on events in Active Views or Search results, you must first add the action in the Event Actions Configuration. For more information, see “[Assigning Actions to Events](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ For information on executing an action on events in Active Views, see “[Executing Actions on Events](#)” in the *NetIQ Sentinel User Guide*.
 - ♦ For information on executing actions on events in Search results, see “[Executing Actions](#)” in the *NetIQ Sentinel User Guide*.

Each individual Action plug-in defines where it can be used and what data it requires as input. Every Action plug-in has certain performance characteristics relating to how quickly it can execute, reset, and be ready for the next event. When an Action instance is created, it inherits the characteristics of the selected Action plug-in. For better performance, not all Actions are available for all the different

Action modes in Sentinel. For example, Actions based on the Send E-mail Action plug-in do not appear in Event Routing rules because you might not want to receive messages with a large event stream every time the rule fires.

For information on where an Action plug-in can be used, refer to the Action Modes section in the specific Action plug-in document.

13.3.2 Debugging Actions

You can debug the Action files from the Sentinel Control Center by using the Action debugging option. The debugger is a local debugger that executes scripts on which the Sentinel Control Center is running. The debugger instantiates a debug session from the Sentinel server machine.



Only actions that are executed in an Incident can be debugged. Therefore, a prerequisite to debug an action is to execute that action in an Incident. For more information, see “[Executing Incident Actions](#)” in the *NetIQ Sentinel User Guide*.

The Action debugger has the following controls:

Table 13-1 Debugger Controls

Action	Description
Run	Runs the script until the next breakpoint is encountered.
Step In	Steps into a function, one line at a time.
Pause	Pauses the running script.
Stop	Stops the script.
Step Over	Steps over a function to the next line in the script.
Step Out	Steps out of the function to the next line in the script.

To debug an action:

- 1 Execute an action in an Incident.
For more information, see “[Executing Incident Actions](#)” in the *NetIQ Sentinel User Guide*.
- 2 In the Sentinel Control Center toolbar, click the **Action Debugging**  icon.
- 3 Click  to start the debugging process. The debugger panel displays the source code and positions the cursor on the first line of the script.
You can debug the script as many times as needed. To debug the script by using a different incident, close the Debug JavaScript Correlation Action window and repeat the debugging process.

13.4 Managing Action Plug-Ins

The Sentinel system includes several Action plug-ins by default, but you can download updates, documentation, and additional Actions from the [Sentinel Plug-ins Web site](#). For more information on specific Actions, see the documentation that is available with the Action. Alternatively, you can view a specific Action's documentation by clicking the **Help** button while configuring that Action in the Action Manager.


Although common Action plug-ins can be obtained from NetIQ, you can create and manage your own Action plug-ins. New Action plug-ins can be developed by using the [Sentinel Plug-in SDK](http://www.novell.com/developer/develop_to_sentinel.html) (http://www.novell.com/developer/develop_to_sentinel.html).

- ♦ [Section 13.4.1, “Understanding the Action Plug-In Manager Interface,” on page 144](#)
- ♦ [Section 13.4.2, “Importing an Action Plug-In,” on page 145](#)



13.4.1 Understanding the Action Plug-In Manager Interface

The Action Plug-in Manager window lists the available Action plug-ins and allows you to import and manage action plug-ins.




Figure 13-3 Action Plug-In Manager



The Action Plug-in Manager window includes the following options:

- ♦ **Add** : Allows you to import or update an Action plug-in.
- ♦ **Delete** : Allows you to delete an Action plug-in.

NOTE: You can delete an Action plug-in only if there are no Action instances configured to use it. For example, if an Action is referenced by an Integrator, the Action cannot be deleted.

- ♦ **Refresh** : Refreshes the list of plug-ins.
- ♦ **View Details** : Displays the plug-in details such as the version number, release date, and a brief description of the Action.
- ♦ **Add Auxiliary file** : Allows you to add an auxiliary file to the Action plug-in. Auxiliary files are occasionally required if NetIQ cannot ship a specific file within the plug-in for licensing reasons, or if users want to customize or extend the Action in some way.




13.4.2 Importing an Action Plug-In

You can download and import an action plug-in package file or you can import an action plug-in file that is in a directory.




When you import a file from a directory, it is important to define the required objects correctly so that the Actions are available for the appropriate Sentinel features.

- ♦ “Downloading and Importing an Action Plug-In” on page 145
- ♦ “Importing an Action Plug-In File from a Directory” on page 145

Downloading and Importing an Action Plug-In

- 1 Download the appropriate Integrator plug-in from the [Sentinel Plug-Ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).
- 2 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 3 Launch the Action Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.
- 4 Click **Manage Plug-ins**, then click the **Import**  icon.
- 5 To import an Action plug-in package file, select **Import an Action plug-in file (.zip, acz)**, then click **Next**.
- 6 Click **Browse** to select the location of the package file.
The file must be in .zip or acz format.
- 7 .(Conditional) If you have selected an Action file that already exists, the Replace Existing Plug-in window is displayed. Click **Next** to replace the existing plug-ins.
- 8 Review the plug-in details, then click **Finish**.

Importing an Action Plug-In File from a Directory

- 1 Download the appropriate Integrator plug-in from the [Sentinel Plug-Ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).
- 2 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 3 Launch the Action Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Action Manager** or click the  icon in the toolbar.
- 4 Click **Manage Plug-ins**, then click the **Import**  icon.
- 5 Select **Import an Action plug-in from directory**.
- 6 Click **Browse** to select the location of the directory, then click **Next**.

- 7 Review the details and make changes as necessary.

If you have selected an Action file that already exists, the Replace Existing Plug-in window is displayed. Click **Next** to replace the existing plug-ins.

- 8 Select the objects you want to associate with the action, then click **Next**.

This affects where the Action is available in the Sentinel interface.

- 9 (Optional) Customize the plug-in by adding or editing the parameters.

For example, if you want to include the attachment option for the Send E-mail action, you can add a parameter named "Attachment" and configure it appropriately. Also, you must modify the code in the action script to handle the new parameter.

- 10 Click **Add** to add parameters that can be set when an Action is configured.

This option should be used for any file that expects to receive information as parameters.

- 10a Specify the parameter name.

The name used here should be identical to the one used in the JavaScript API method `scriptEnv.getParameter("<name>")` in the script that is being imported.

- 10b Select the parameter type from the **Type** drop-down list.

- ♦ **String:** Passes the string value to the parameter.
- ♦ **Boolean:** The parameter can take a True or False value.
- ♦ **Integrator:** Select an Integrator for the action.
- ♦ **Event Tag:** Select an Event Tag for the parameters.
- ♦ **Severity:** Select a severity for the parameters.

NOTE: The **Options** area is available only for the String type parameters.

- 10c (Optional) Specify a description.

- 10d Click **Next**.

- 11 Review the plug-in details, then click **Finish**.

14 Configuring Integrators

- ♦ [Section 14.1, “Overview,” on page 147](#)
- ♦ [Section 14.2, “Managing Integrators,” on page 148](#)
- ♦ [Section 14.3, “Managing Integrator Plug-Ins,” on page 153](#)

14.1 Overview

Integrators are plug-ins that can be used in Sentinel to extend the features and functionality of Sentinel remediation actions. Integrators allow Sentinel to connect to external systems, such as an LDAP server, SMTP server, or SOAP server. Actions use Integrators to interact with other systems. For example, you can set an attribute in Novell eDirectory, an LDAP server, to enable or disable a user, edit details, and so on. You can also start an Identity Manager workflow, such as a provisioning request, by using SOAP calls.

The general process for using an Integrator to perform remediation actions involves the following steps:

1. Determine the best type of Integrator to access the external system with which you want to interact.
2. Import and configure the appropriate Integrator to connect to the external system. Or, you can use the Integrators that are configured by default.
3. Configure the appropriate Action. Or, you can use the Actions that are configured by default. For more information, see [Chapter 13, “Configuring Actions,” on page 139](#).
4. Perform additional configuration, if desired, to associate the action with a deployed correlation rule or an event menu action.
5. The action is executed:
 - ♦ When an associated correlation rule fires.
 - ♦ When you execute actions on the selected events in the Sentinel Web interface > **Event Actions** tab.
 - ♦ When you execute actions on the selected Incident in the Sentinel Control Center > **Incidents > Actions > Execute Incident Action**.
 - ♦ When a you select an action for the selected event in the Sentinel Control Center > **Active View > right-click menu**.

For more information on specific Integrators, see the documentation that is available with the Integrators. Alternatively, you can view a specific Integrator’s documentation by clicking the **Help** button in the Integrator Manager after configuring that Integrator.

NOTE: Only users in the administrator role can configure and manage Integrators.

14.2 Managing Integrators

An Integrator is a configured instance of an Integrator plug-in. There can be one or more instances of an Integrator plug-in with different parameters or settings. A few Integrators are available by default. You can also add additional Integrators as required.

- ♦ [Section 14.2.1, “Configuring the Default Integrators,” on page 148](#)
- ♦ [Section 14.2.2, “Adding an Integrator,” on page 151](#)
- ♦ [Section 14.2.3, “Viewing Integrator Health Details,” on page 151](#)



14.2.1 Configuring the Default Integrators

The default Integrators installed with Sentinel are not configured. To use these default Integrators with the actions in your Sentinel system, you must configure the Integrators for your environment.

- ♦ [“Configuring the File Integrator” on page 148](#)
- ♦ [“Configuring the Sentinel Link Integrator” on page 148](#)
- ♦ [“Configuring the Sentinel Mail Integrator” on page 149](#)
- ♦ [“Configuring the SNMP Integrator” on page 150](#)
- ♦ [“Configuring the Syslog Integrator” on page 151](#)

Configuring the File Integrator

The File Integrator is used with the **Log to File** action.



- 1 Access the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).
- 2 Launch the Integrator Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 3 Select **File** in the Integrators column, then click the **File Configuration** tab.
- 4 Change the default filename and location where the events are stored.
By default, the location is `../data/log_to_file_events.txt`.
- 5 Click **Save**.

Configuring the Sentinel Link Integrator

The Sentinel Link Integrator is used to connect multiple Sentinel systems. For more information, see [Chapter 12, “Linking to Additional Sentinel Systems,” on page 135](#).

- 1 Access the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).

2 Launch the Integrator Manager:

- ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.

3 Select **Sentinel Link** in the Integrators column, then click the **Sentinel Link Connector** tab.

4 Use the following information to configure the Sentinel Link Integrator:

Host Name: Specify the hostname or IP address of the destination Sentinel system where a Sentinel Link Connector is configured.

Port Number: Specify the port number for the destination Sentinel system. The default port is 1290.

Encrypted (HTTPS)/Not Encrypted (HTTP): Select whether the connection to the destination Sentinel system is encrypted or not encrypted. If you selected Encrypted, there are additional fields to configure:

♦ **Server Validation Mode:** Select one of the following:

- ♦ **None:** The Integrator does not validate the receiver's certificate.
- ♦ **Strict:** The Integrator always verifies the receiver's certificate when connecting to the receiver. If this option is selected, the Integrator immediately attempts to retrieve the receiver's certificate over the network and validate that it is issued by an authorized CA.

If the certificate is not validated, it is still presented to the user to accept or reject. The certificate is considered to be valid if the user accepts it. When a validated certificate is acquired, it is stored in the Integrator's configuration. From now on, the Integrator allows communication only with a receiver that provides that certificate during the initial connection setup.

♦ **Integrator Key Pair:** Select one of the following:

- ♦ **None:** The receiver system does not validate the sender certificates. Select this option if the receiver's client authentication type is configured to **Open**.
- ♦ **Custom:** The receiver system validates the sender certificates. Select this option if the receiver's client authentication type is configured to **Strict**. If the receiver system performs a strict validation, it imports a trust store, which contains all the sender certificates that it trusts.

After selecting this option, click the **Import Key Pair** button to import a key pair. The key pair you import must match one of the certificates that is included in the trust store that is imported by the receiver system.

5 Click **Save**.



Configuring the Sentinel Mail Integrator

All Sentinel events that meet the filter criteria for which the **Send an E-mail** action is defined are sent to the associated SMTP relay and e-mail addresses by the Sentinel Mail Integrator.

To configure the Sentinel Mail Integrator:

1 Access the Sentinel Control Center.



For more information, see [Section 1.2, "Accessing the Sentinel Control Center,"](#) on page 18.

- 2 Launch the Integrator Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 3 Select **Sentinel Mail** in the Integrators column, then click the **Connection** tab.
- 4 Use the following information to configure the Sentinel Mail Integrator:
 - Host:** Specify the hostname or IP address of an available SMTP server.
 - Port:** Specify the port number of an available SMTP server. The default port is 25.
 - From (default):** Specify an address to using or sending the e-mail messages are sent. The default value is siem@yourcompany.com.
 - User:** If the SMTP server requires authentication, specify a user name.
 - Password:** Specify the password for authentication to the SMTP server.
- 5 Click **Save**.

Configuring the SNMP Integrator

All Sentinel events that meet the filter criteria for which the Send SNMP Traps action is defined are sent to the specified SNMP addresses.



To configure the SNMP Integrator:

- 1 Access the Sentinel Control Center.
For more information, see [Section 1.2, "Accessing the Sentinel Control Center," on page 18](#).
- 2 Launch the Integrator Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 3 Select **snmp** in the Integrators column, then click the **Server Configuration** tab.
- 4 Use the following information to configure the SNMP server:
 - Host:** Specify the IP address or hostname of the SNMP server you want to send the trap to.
 - Port:** Specify the port number for the SNMP server. The default port is 162.
 - Community String (Password):** Specify the community string (password) to access the SNMP management system. If no community string is specified, the Integrator sets the default value to public.
 - OID:** Specify the desired ASNL object ID you want to associate with this message. If no object ID is specified, the Novell Audit internal OID is used (2.16.840.1.113719.1.347.3.1).
- 5 Click **Save**.

Configuring the Syslog Integrator

All Sentinel events that meet the filter criteria for the **Send to Syslog** action are sent to the specified syslog server.



To configure the Syslog Integrator:

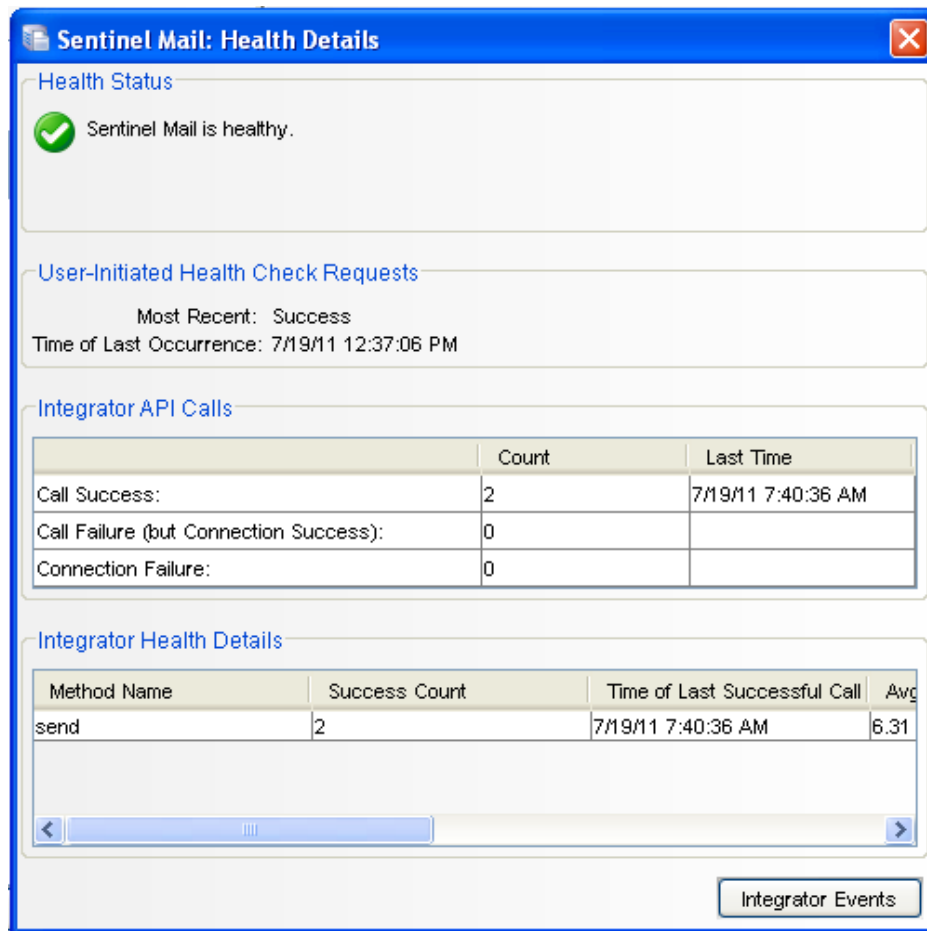
- 1 Access the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 2 Launch the Integrator Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 3 Select **Syslog** in the Integrators column, then click the **Server Configuration** tab.
- 4 Use the following information to configure the syslog Integrator.
 - Host:** Specify the host name or IP address of the syslog server.
 - Protocol:** Select the protocol used to connect to the syslog server.
 - Port:** Specify the port number used to connect to the syslog server.
 - Default Facility:** Select the default facility for the syslog server from the drop-down list. The facility allows you to classify the syslog messages.
 - Stream Encoding:** Select the encoding standard for the syslog Integrator.
- 5 Click **Save**.

14.2.2 Adding an Integrator

The specific steps to configure an Integrator depend on the type of Integrator. The steps are described in detail in documents that come with the Integrators. Documentation for installed plug-ins can be viewed by selecting an Integrator in the Integrator Manager and clicking **Help**. Or, you can refer to the document that comes with the Integrator plug-in.

14.2.3 Viewing Integrator Health Details

- 1 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 2 Launch the Integrator Manager:
 - ♦ (Conditional) If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ (Conditional) If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 3 In the Integrator Manager window, select an Integrator from the left pane.
- 4 Click **see details**.



The Health Details window displays the Refresh Health State, time of last occurrence, its method calls and the related events of the selected Integrator configuration.

- ♦ **Integrator API Calls:** Indicates the status of the connection and the method calls used from the API of the selected integrator. For more information on the JavaScript plug-in, see [Chapter 13, "Configuring Actions,"](#) on page 139.
 - ♦ **Call Success Count:** Displays the number of times the connection was established successfully and the methods were called successfully from the API. Time of Last Occurrence displays the time when the connection and the method call were successful.
 - ♦ **Call Failure (but Connection Success) Count:** Displays the number of times the connection was established successfully but the method call failed. Time of Last Occurrence displays the last time when the connection was successful and the method call failed.
 - ♦ **Connection Failure Count:** Displays the number of times the connection failed. Time of Last Occurrence displays the last time when the connection and method call failed.

NOTE: The most recent success or failure time is shown in the overall health status for the configured Integrator.

- ♦ **Integrator Health Details:** Provides information about the success of the API methods called in the JavaScript action files associated with the Integrator. It provides information specific to the methods called.
 - ♦ **Method Name:** Name of the API method used in the JavaScript.

Success Count: Number of times the API method executed successfully.

Time of Last Successful Call: The time at which the method was last successfully executed.

Average Successful Run Time: Average time to make a successful method call.

Error Count: Number of times the API method failed.

Time of Last Error Call: The time at which the method call failed.

Average Error Run Time: Average time to make a failed method call.




NOTE: The most recent success or failure time is shown in the overall health status of the method.

14.3 Managing Integrator Plug-Ins

Sentinel includes several Integrators plug-ins by default, but you can download updates and additional Integrators from the [Sentinel Plug-ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

- ♦ [Section 14.3.1, “Importing an Integrator Plug-In,” on page 153](#)

14.3.1 Importing an Integrator Plug-In

- 1 Download the appropriate Integrator plug-in from the [Sentinel Plug-Ins Web site \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).
- 2 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).
- 3 Launch the Integrator Manager:
 - ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
 - ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Integrator Manager** or click the  icon in the toolbar.
- 4 Click **Manage Plug-Ins**.
- 5 In the Integrator Plug-in Manager window, click the  icon.
- 6 In the Plug-in Import Type window, select **Import an Integrator plug-in file (.zip)**, then click **Next**.
- 7 In the Choose Plug-in Package File window, click **Browse** to locate the Integrator zip file you want to import to the plug-in repository, then click **Open**.
- 8 (Conditional) If you have selected an Integrator file that already exists, the Replace Existing Plug-in window is displayed. Click **Next** to replace the existing plug-ins.
- 9 Click **Next**.
- 10 In the Plug-in Details window, select **Launch Integrator Configuration Wizard** to deploy the plug-in after importing it.
- 11 Click **Finish**.

After importing the Integrator plug-in, you must configure the Integrator plug-in for your environment. For more information, see [Section 14.2.1, “Configuring the Default Integrators,” on page 148](#).

15 Integrating Identity Information

This chapter provides information about integrating Sentinel with identity management systems.

- ♦ [Section 15.1, “Overview,” on page 155](#)
- ♦ [Section 15.2, “Integration with Identity Management Systems,” on page 156](#)
- ♦ [Section 15.3, “Leveraging Identity Information,” on page 158](#)

15.1 Overview

Users in an IT environment have multiple accounts and sometimes multiple account identifiers per user. Normally, if a user has multiple account identifiers, Sentinel treats each identifier as a unique account.

This means a user could log in to Active Directory and an LDAP directory and Sentinel tracks these events, but Sentinel does not realize these events are related to the same user account.

In order to overcome this issues, Sentinel provides an integration framework to identity management systems to track the identities of for each user account and what events those identities have performed.

This integration provides functionality on several levels:

- ♦ The Identity Browser provides the ability to look up the following information about a user:
 - ♦ Contact information
 - ♦ Accounts associated with that user
 - ♦ Most recent authentication events
 - ♦ Most recent access events
 - ♦ Most recent permissions changes
- ♦ The Identity Browser lets you do a lookup from events
- ♦ Reports and Correlation rules provide an integrated view of a user's true identity, even across multiple systems on which the user has separate accounts. For example, accounts like COMPANY\testuser; > cn=testuser,ou=engineering,o=company, and TUser@company.com can be mapped to the actual person who owns the accounts.

By displaying information about the people initiating a given action or people affected by an action, incident response times are improved and behavior-based analysis is enabled.

Sentinel provides an optional integration with Novell Identity Manager. The figures and descriptions in this section are based on Identity Manager.

Sentinel synchronizes identity information with major identity management systems and stores local copies of key information about each identity. The following table summarizes the commonly used information provided:

Name	Description
AccountGUID	Auto-generated internal ID
Name	User name that references the account, generally provided by the user to log in.
ID	The numeric or other identifier that represents the account in the event source. This ID is used for resolution when the user name is not available.
Authority	The realm within which this account is unique. Collectors calculate the realm based on event information.
Status	The status of the account
IdentityGUID	A reference to the identity that owns this account

The identities stored by Sentinel are then linked with accounts created on endpoint systems by the identity management system. This helps Sentinel associate the correct identity information with the native events from those endpoint system. Some identity information is injected directly into the inbound event by using the mapping service. The remaining identity information, such as photograph and contact information, is accessible through the Identity Browser.

The identity information injected into the event can be used for correlation and for performing actions on the identities that are associated with detected activity. For example, Sentinel is able to see multiple failed logins from a given person and not just an account. A detected violation could trigger disabling activities for all accounts associated with an identity.

15.2 Integration with Identity Management Systems

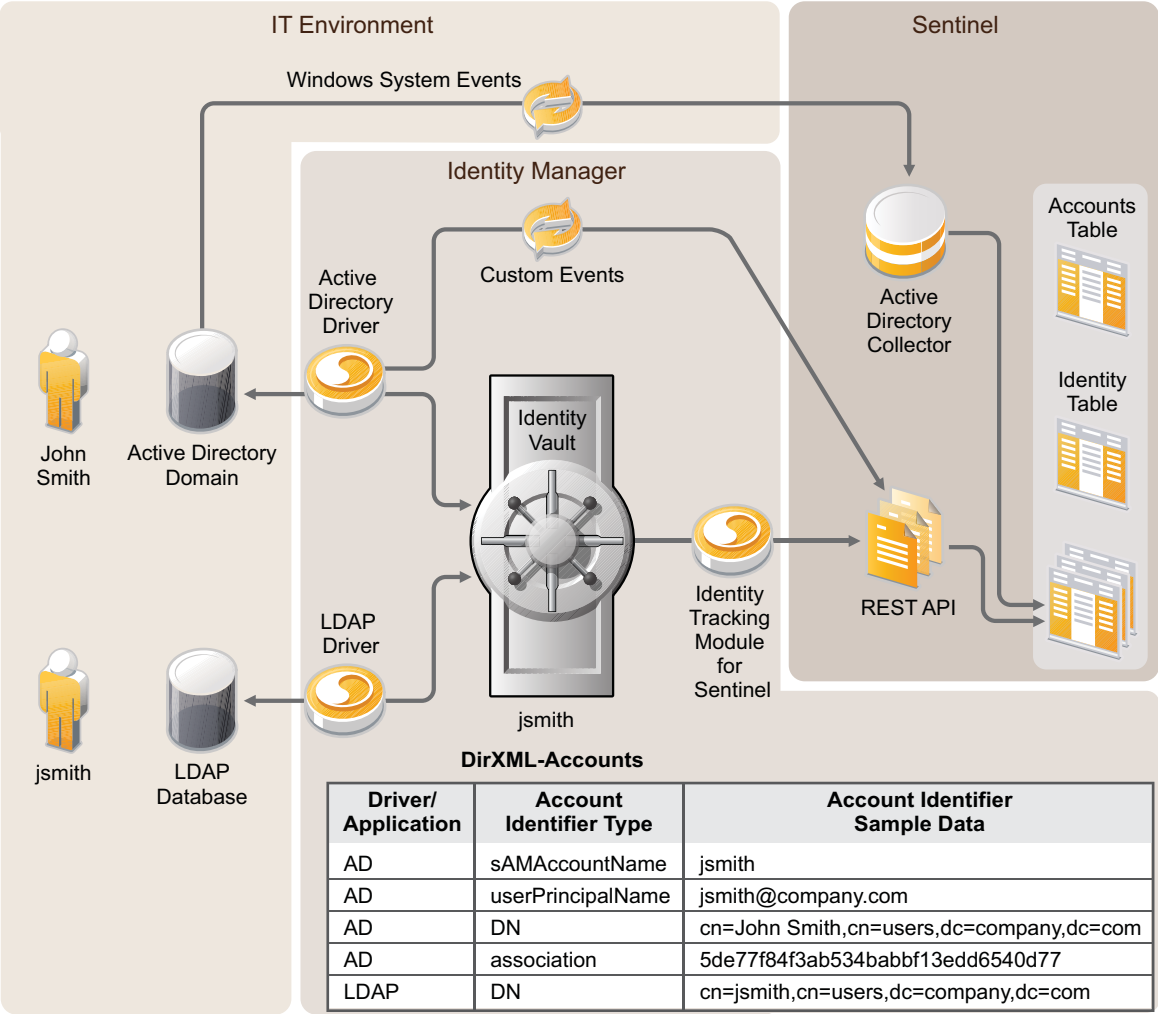
NetIQ provides a solution that integrates Sentinel with Novell Identity Manager. For other identity systems, similar integration can be achieved by writing an identity synchronization collector that uses the Identity API. For more information, see the [Novell Plug-in SDK Web site for Sentinel](#).

The examples and illustrations in this section use Identity Manager to explain how integration works among identity management systems.

Sentinel integration with Identity Manager is provided using the Identity Manager driver called the Identity Tracking Module for Sentinel. Standard event collection must also be set up with the systems to which Identity Manager has provisioned accounts using the Collectors available at the [Sentinel Plug-ins Web site](#).

After Sentinel and Identity Manager are installed, the Identity Tracking Module for Sentinel sends identity and account information from the Identity Vault to the Sentinel REST API, which populates the Sentinel database.

Figure 15-1 Sentinel Integration with IDM



The time required to initially populate the Sentinel database depends on the amount of data in the Identity Vault; identity information including photographs requires significantly more time to load.

The Identity Tracking Module for Sentinel also keeps the identity information synchronized as information is updated in the Identity Vault during normal Identity Manager operations.

After the identity and account information are loaded, a map named `IdentityAccount` is automatically generated in `/var/opt/novell/sentinel/data/map_data`. The map contains the following information:

- Account Name
- Authority
- Customer Name
- Identity GUID
- Full Name
- Department
- Job Title

- ♦ Manager GUID
- ♦ Account Status

IMPORTANT: An identity can have multiple accounts but one account cannot be assigned to multiple identities.

The identity map is automatically applied to all events from Collectors to look for an identical match between the information in the event and key fields in the map. The table below shows the fields that are populated if all of the map key fields and event data exactly match. These mappings are automatically configured and are not editable.

Label	Populated by which Column from IdentityAccount Map
InitUserDepartment	Department
InitUserFullName	Full Name
InitUserIdentity	Identity GUID
TargetUserDepartment	Department
TargetUserFullName	Full Name
TargetUserIdentity	Identity GUID

NOTE: To find a match, the event fields and map key fields must match exactly. This might require modifications to existing Collectors to enable them to parse or concatenate data to make these fields match the data from the Identity Vault.

After these fields are added to the event by the mapping service, they are used by Correlation rules, remediation actions, and reports in the Identity Tracking Solution Pack. In addition to using the content included in the Solution Pack, users can also perform the following actions:

- ♦ Create Correlation rules based on identity in addition to account name. This allows you to look for similar events from a single user, which provides a more comprehensive view than looking at events from a single account
- ♦ Create reports that show identity, including all accounts associated with a user
- ♦ Use the Identity Browser to get more information about users and their activities

15.3 Leveraging Identity Information

For information about searching and viewing user profiles of the identities synchronized from the identity management system, see [“Searching and Viewing User Identities”](#) in the *NetIQ Sentinel User Guide*.

16 Managing Data Feeds

Data feeds allow Sentinel to maintain a current set of data from Web sites that are regularly updated by external sources. For example, a typical feed might provide a list of known compromised hosts, and if any connections are made to those hosts, the source of the connection becomes a suspect.

Solution Packs in Sentinel provide Data Feed plug-ins that allow Sentinel to consume external intelligence data, which provides contextual information. For example, data feed plug-ins for Palevo can download lists of known botnet IP addresses and domains. You can leverage these lists in correlation rules to detect communications to known botnets in your network.

Sentinel includes Solution Packs by default. New data feed plug-ins are frequently bundled in the latest Solution Packs, which are uploaded to the [Sentinel Plug-ins Web site](#). To get the latest data feed plug-ins, download the relevant Solution Packs from the [Sentinel Plug-ins Web site](#).

You can configure the data feeds plug-ins to either automatically download data feeds at specified intervals or download feeds manually at the desired time. You can download the data feeds in the Sentinel Web interface, click **Collection > Plug-ins**.

16.1 Using Data Feeds for Botnet Detection

Solution Packs such as the Threat Intelligence Solution Pack provide data feeds plug-ins that include Palevo tracker, SpyEye tracker, and Zeus tracker. These data feeds plug-ins help you detect connections to botnet servers. Data feed plug-ins define the source of the feed and how to process the data in that feed.

Along with the data feed plug-ins, Solution Packs also install useful correlation rules, dynamic lists, and filters. You must deploy the correlation rules to detect botnet connections to your network environment. For information about deploying correlation rules, see “[Deploying Rules in the Correlation Engine](#)” in the *NetIQ Sentinel User Guide*.

To detect botnet connections, you must first enable the data feed plug-ins to download botnet feeds. The botnet feeds Web sites are updated on a regular basis as new IP addresses and domains are reported as botnets. For information about downloading the feeds, see the Solution Pack documentation.

17 Detecting Vulnerabilities and Exploits

Sentinel Advisor, powered by Security Nexus, is an optional data subscription service that provides device-level correlation between real-time events, from intrusion detection and prevention systems, and from enterprise vulnerability scan results. Advisor acts as an early warning service and detects attacks against vulnerable systems by providing normalized attack information. It also provides the associated remediation information.

The Advisor subscription is optional. However, it is necessary if you want to use the Sentinel Exploit Detection or the Advisor Reporting features.

- ♦ [Section 17.1, “Understanding Advisor,” on page 161](#)
- ♦ [Section 17.2, “Understanding Exploit Detection,” on page 162](#)
- ♦ [Section 17.3, “Introduction to the Advisor Interface,” on page 164](#)
- ♦ [Section 17.4, “Processing the Advisor Feed,” on page 166](#)
- ♦ [Section 17.5, “Configuring the Advisor Products for Exploit Detection,” on page 167](#)
- ♦ [Section 17.6, “Downloading the Advisor Feed,” on page 167](#)
- ♦ [Section 17.7, “Viewing the Advisor Status,” on page 170](#)
- ♦ [Section 17.8, “Viewing the Advisor Data,” on page 172](#)
- ♦ [Section 17.9, “Resetting the Advisor Password,” on page 172](#)
- ♦ [Section 17.10, “Deleting the Advisor Data,” on page 173](#)

17.1 Understanding Advisor

The Advisor service and its corresponding Exploit Detection feature depend on the mappings between the attacks against enterprise assets and the known vulnerabilities of those assets. The Advisor and the Exploit Detection features require the following data to work with the Advisor products:

- ♦ **Vulnerability scan data:** The vulnerability scanners check enterprise assets for known vulnerabilities. The scanned data can then be loaded into the Sentinel database to serve as referential information, by using the Collectors that support Advisor.
- ♦ **Advisor mapping data:** The Advisor data contains information about known threats, including attacks and vulnerabilities. The Advisor service gathers information from various vulnerability and intrusion detection vendors, and creates mappings between abstract vulnerabilities and attacks.

Security Nexus provides the Advisor feed data that contains information about known security vulnerabilities and threats, and also provides normalization of intrusion detection signatures and vulnerability scans. The Advisor data feed is updated on a regular basis as new attacks and vulnerabilities are reported. The updates are available at the [Novell download Web site \(https://secure-www.novell.com/sentinel/download/advisor/\)](https://secure-www.novell.com/sentinel/download/advisor/).

NOTE: The initial Advisor data feed is installed by default on the Sentinel server at `/var/opt/novell/sentinel/data/updates/advisor`. However, you must purchase an additional license from Novell to download the updated Advisor feed.

- ♦ **Real-time attack data:** Intrusion detection systems report real-time attacks against enterprise assets. However, this data does not indicate the impact of the attacks.

The real-time attacks that are generated as events are loaded into the Sentinel database by using the intrusion detection systems or vulnerability type Collectors.

17.2 Understanding Exploit Detection

- ♦ [Section 17.2.1, “How Exploit Detection Works,” on page 162](#)
- ♦ [Section 17.2.2, “Generating the Exploit Detection File,” on page 164](#)
- ♦ [Section 17.2.3, “Viewing the Events,” on page 164](#)

17.2.1 How Exploit Detection Works

Exploit detection instantly sends notification when an attack is attempting to exploit a vulnerable system. The Exploit Detection feature depends on the following:

- ♦ Both vulnerability scanners and the intrusion detection systems must report vulnerabilities and attacks against the same set of systems. In Sentinel, systems are identified by their IP addresses and their tenant name. The tenant name is a namespace identifier that prevents overlapping IP ranges from matching incorrectly. The tenant name can be the name of the customer, division, department, and so forth that owns this event data.
- ♦ The vulnerability scanner and intrusion detection system products must be supported by the Advisor service. This data uses specific product identifiers to ensure proper matching.
- ♦ The specific reported attacks and vulnerabilities must be known to the Advisor service and Exploit Detection.

All Collectors shipped by NetIQ meet these requirements, as long as they are declared as being supported by Advisor. To write your own vulnerability or intrusion detection Collector, or to modify one of the shipped Collectors, refer to the [Sentinel Plug-in SDK \(http://developer.novell.com/wiki/index.php?title=Develop_to_Sentinel\)](http://developer.novell.com/wiki/index.php?title=Develop_to_Sentinel) for specific information about which event and vulnerability fields must be filled in to support this service.

The following table lists the supported products with their associated device type (IDS for intrusion detection system, VULN for vulnerability scanners, and FW for firewall).

Table 17-1 Supported Products and the Associated Device Types

Supported Products	Device Type	RV31 Value
Cisco Secure IDS	IDS	Secure
Enterasys Dragon Host Sensor	IDS	Dragon
Enterasys Dragon Network Sensor	IDS	Dragon
Intrusion.com (SecureNet_Provider)	IDS	SecureNet_Provider

Supported Products	Device Type	RV31 Value
ISS BlackICE PC Protection	IDS	XForce
ISS RealSecure Desktop	IDS	XForce
ISS RealSecure Network	IDS	XForce
ISS RealSecure Server	IDS	XForce
ISS RealSecure Guard	IDS	XForce
Sourcefire Snort/Phalanx	IDS	Snort
Symantec Network Security 4.0 (ManHunt)	IDS	ManHunt
Symantec Intruder Alert	IDS	Intruder
McAfee IntruShield	IDS	IntruShield
TippingPoint	IPS	TippingPoint
eEYE Retina	VULN	Retina
Foundstone Foundscan	VULN	Foundstone
ISS Database Scanner	VULN	XForce
ISS Internet Scanner	VULN	XForce
ISS System Scanner	VULN	XForce
ISS Wireless Scanner	VULN	XForce
Nessus	VULN	Nessus
nCircle IP360	VULN	nCircle IP360
Qualys QualysGuard	VULN	QualysGuard
Cisco IOS Firewall	FW	Secure

To enable exploit detection, the Sentinel Collectors must populate several variables as expected. Collectors built by NetIQ populate these variables by default. The Collectors are available for download on the [Sentinel Plug-ins Web site](#). For information on configuring the Collectors, Connectors, and Event Sources, see [Section 7.4, “Configuring Data Collection for Other Event Sources,”](#) on page 87.

- ♦ In intrusion detection systems and vulnerability Collectors, the RV31 (IDSName) variable in the event must be set to the value in the RV31 column in [Table 17-1](#). This string is case sensitive.
- ♦ In the intrusion detection systems Collector, the DIP (TargetIP) must be populated with the IP address of the machine that is being attacked.
- ♦ In the intrusion detection systems Collector, RT1 (IDSAttackName) must be set to the attack name or attack code for that intrusion detection system.
- ♦ In the intrusion detection systems and vulnerability Collectors, the RV39 (TenantName) value must be populated. For a standard corporation, the value can be anything. For a Managed Security Service Provider (MSSP), the tenant name should be set for the individual customer. For either type of company, the value in the intrusion detection systems Collector must exactly match with the value in the vulnerability Collector.

These values are used by the Mapping Service to populate the VULN field in the event. This value is used to evaluate the incoming events to determine whether a vulnerability is exploited or not. When the vulnerability field (VULN) equals 1, the asset or destination device is exploited. If the vulnerability field equals 0, the asset or destination device is not exploited.

17.2.2 Generating the Exploit Detection File

When you run the intrusion detection system or vulnerability type Collectors, events from all the selected products are scanned for possible attacks and vulnerabilities, and the product name and the tenant name are mapped to the Advisor product name and the tenant name. If the events match successfully, the exploit information (IP address, device name, attack name, and tenant name) is updated in the `exploitdetection.csv` file in the `/var/opt/novell/sentinel/data/map_data` directory.

The initial mapping time might take up to 30 minutes. However, you can modify the time by changing the value of the `minregenerateinterval` property in the `ExploitDetectDataGenerator` component of the `server.xml` file. The time is given in milliseconds. For example, you can change the time from 1800000 (30 minutes) to 180000 (3 minutes).

NOTE: You must restart the Sentinel service for the changes to take effect.

17.2.3 Viewing the Events

To view events that indicate a possible exploitation, search for events whose Vulnerability value is 1.

The value in the Vulnerability field conveys the following:

- ♦ **1:** The asset or destination device is possibly exploited.
- ♦ **0:** The asset or destination device is not exploited.

NOTE: If the Vulnerability field is blank, the `exploitdetection.csv` file is not generated.

For more information on viewing events, see [Section 17.8, “Viewing the Advisor Data,”](#) on page 172.

17.3 Introduction to the Advisor Interface


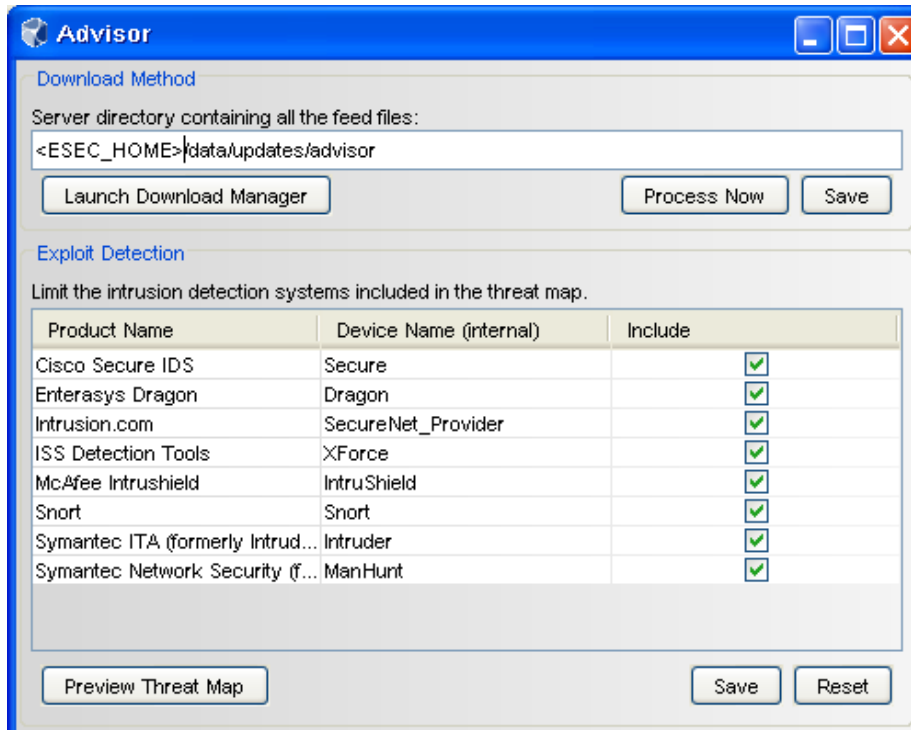
The Advisor feature is available in the Sentinel Control Center. You can access Advisor either by clicking the **Advisor** tab or by clicking the Advisor icon  in the toolbar. The Advisor configuration interface is available only for users in the administrator role.

Figure 17-1 Advisor Window



- ♦ **Download Method:** Enables you to process the Advisor feed files manually and launch the Download Manager feature to configure the Sentinel server for automated processing of the feed files. For more information on processing the Advisor feed, see [Section 17.4, “Processing the Advisor Feed,”](#) on page 166.
- ♦ **Exploit Detection:** Lists the vulnerable products that are included in the feed files, and enables you to configure the products for exploit detection. For more information, see [Section 17.5, “Configuring the Advisor Products for Exploit Detection,”](#) on page 167.

NOTE: The Exploit Detection section initially displays a blank list unless you process the initial Advisor feed that was loaded during Sentinel installation. For more information, see [Section 17.4, “Processing the Advisor Feed,”](#) on page 166.

- ♦ **Preview Threat Map:** The Preview Threat Map window lists the top 5000 entries of the `exploitdetection.csv` file. This list displays the attacks that attempt to exploit your machine. To view the threat map: click **Preview Threat Map**.

Figure 17-2 Viewing the Threat Map

TargetIP	IDSName	IDSAAttackName	TenantName
10.22.81.145	XForce	17889	unknown
10.22.81.145	XForce	2358	unknown
10.22.81.145	XForce	3992	unknown
10.22.81.145	TippingPoint	2176:SMB:Null S...	unknown
10.22.81.145	TippingPoint	2176:SMB:Null S...	unknown
10.22.81.145	TippingPoint	2176:SMB:Null S...	unknown
10.22.81.145	SecureNet_Provi...	FTP Client[STAF ...	unknown
10.22.81.145	SecureNet_Provi...	FTP Client[STAF ...	unknown
10.22.81.145	Snort	FTP EXPLOIT ST...	unknown
10.22.81.145	Snort	FTP EXPLOIT ST...	unknown
10.22.81.145	ManHunt	HTTP MS IE Malf...	unknown
10.22.81.145	IntruShield	NETBIOS-NS:Se...	unknown
10.22.81.145	IntruShield	NETBIOS-SS:PI...	unknown

Ready Refresh Refreshed At: 1/20/12 11:26:23 AM

NOTE: This list is blank unless the `exploitdetection.csv` file has been generated.

17.4 Processing the Advisor Feed

You can process the Advisor feed files manually or you can configure the Sentinel server to automatically process the feed files at scheduled time intervals.

- ♦ [Section 17.4.1, “Processing the Feed Files Manually,” on page 166](#)
- ♦ [Section 17.4.2, “Processing the Feed Files Automatically,” on page 167](#)

17.4.1 Processing the Feed Files Manually

- 1 Launch the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).

- 2 Click **Advisor**, then click **Advisor Configuration** in the **Navigator** panel.

- 3 In the Advisor window, select the directory where you downloaded the latest Advisor feed files.

The initial Advisor feed is loaded at `/var/opt/novell/sentinel/data/updates/advisor`.

- 4 Click **Process Now** to process and load the feed files into the Sentinel database.

After the feed files are processed, the products included in the feed files are displayed in the Exploit Detection section.

- 5 (Optional) Click **Save** to save the location of the Advisor directory.

17.4.2 Processing the Feed Files Automatically

You can use the Download Manager to configure the Sentinel server to automatically process the feed files after they are downloaded. For more information on the Download Manager, see [Section 17.6.1, “Configuring the Sentinel Server for Automated Downloads,”](#) on page 168.

17.5 Configuring the Advisor Products for Exploit Detection

The Exploit Detection section of the Advisor window lists the names of the Advisor products and the device names that are included in the feed files.

- 1 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 2 Click **Advisor**, then click **Advisor Configuration** in the **Navigator** panel.
- 3 In the Exploit Detection section, select the products that need to be included for Exploit Detection by selecting the corresponding check boxes.
- 4 (Conditional) To remove any product from the list, deselect the corresponding check box.
- 5 Click **Save** to save the changes made to the Advisor products list.
After the product list is saved, the `exploitdetection.csv` file is updated. For more information on exploit detection, see [“Generating the Exploit Detection File”](#) on page 164.
- 6 (Optional) Click **Reset** to undo the changes made to the Exploit Detection products list.

For more information on exploit detection, see [Section 17.2, “Understanding Exploit Detection,”](#) on page 162.

17.6 Downloading the Advisor Feed

To access the updated Advisor feed, you must download the feed and process it into the Sentinel database. You can download the Advisor feed manually or you can configure the Sentinel server for automated downloads at fixed intervals.

NOTE: To download Advisor updates, you must purchase the Advisor Data Subscription and obtain the credentials.

- ♦ [Section 17.6.1, “Configuring the Sentinel Server for Automated Downloads,”](#) on page 168
- ♦ [Section 17.6.2, “Creating a Download Configuration,”](#) on page 168
- ♦ [Section 17.6.3, “Downloading the Feed Instantly,”](#) on page 169
- ♦ [Section 17.6.4, “Audit Events for the Download Manager,”](#) on page 170
- ♦ [Section 17.6.5, “Downloading the Advisor Feed Manually,”](#) on page 170

17.6.1 Configuring the Sentinel Server for Automated Downloads

You can use the Download Manager to configure the Sentinel server for automated Advisor feed downloads. The Download Manager enables you to configure the Sentinel server for automated downloading and processing of the Advisor feed files at fixed intervals. The Download Manager notifies the Sentinel processes to process the downloaded feed whenever a feed is downloaded.

For example, you can configure the Sentinel server to download the Advisor feed at fixed intervals. After the feed is downloaded, the Download Manager notifies the Advisor processes to process the downloaded feed and load it into the Sentinel database.



The Download Manager window lists the download configurations and displays the status of the previous download for each download configuration. By default, a download configuration is available that is partially configured for Advisor. You can use the same configuration to download an Advisor feed, by specifying the login credentials and other details, and activating the download configuration.

17.6.2 Creating a Download Configuration

- 1 Launch the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,” on page 18](#).

- 2 Launch the Download Manager:

- ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Download Manager** or click the  icon on the toolbar.
- ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Download Manager** or click the  icon on the toolbar.

- 3 Click **Add** to configure the download feed.

- 4 Specify the following information:

- ♦ **Repository Name:** Name of the repository from which you are downloading the data. The repository name must be unique and meaningful to the feed that you are downloading. For example, to download Advisor data, specify the repository name as NetIQ Advisor.
- ♦ **URL:** URL where the download feed is located. For example, to download Advisor data, specify the Advisor URL (<https://secure-www.novell.com/sentinel/download/advisor/feed/>).

NOTE: The Advisor feed can also be manually downloaded from the Novell Download Web site (<https://secure-www.novell.com/sentinel/download/advisor/>). However, the manual download URL does not work in Download Manager.

- ♦ **Authentication:** Specify the authentication details.
 - ♦ **Anonymous:** Select the check box to download the information as an anonymous user. If **Anonymous** is selected, the **Username** and **Password** fields are disabled. You can only access the URLs that do not require authentication.
 - ♦ **Username and Password:** Specify valid credentials to log in to the URL of the repository.
- ♦ **Network:** If you do not have direct access to the server where the download feed is located, you can download the advisor feed through a proxy server.
 - ♦ **Use a proxy server:** Select the check box if you are using a proxy server to download the feed.

The proxy server through which you are connecting should have access to the server where the download feed is located.

- ♦ **Address:** Specify the proxy server name or server IP address.
- ♦ **Port:** Specify the port number through which you connect to the proxy server.
- ♦ **Username and Password:** Specify valid credentials to log in to the proxy server.
- ♦ **Download Directory:** Specify the location and name of the directory where you want to save the feed. Ensure that you specify the absolute path.

The directory is created on the Sentinel server at the specified path while downloading the feed.

If the specified directory already has existing files that are not in sync with the Sentinel server files, these files are deleted when the new download starts.

- ♦ **Enable:** Select this check box to activate the configuration to download the feed.
- ♦ **Schedule Interval:** Specify the time interval for the download by selecting one of the following options from the drop-down list:
 - ♦ **6 Hours:** Schedules the download for every six hours.
 - ♦ **12 Hours:** Schedules the download for every 12 hours.
 - ♦ **Daily:** Schedules the download every day.
 - ♦ **Weekly:** Schedules the download once a week.
 - ♦ **Monthly:** Schedules the download once a month.
- ♦ **Feed Type:** Select Advisor from the drop-down list to notify the Advisor processes on the Sentinel server that the Advisor feed is downloaded.

5 (Optional) Click **Validate** to validate the URL and the login credentials.

The URL and its credentials are validated and a confirmation message is displayed. If the validation fails, you must provide a valid URL and login credentials.

6 Click **Save** to save the configuration settings.

The specified directory path is validated and the download configuration that you created is displayed in the Download Manager window. If the directory path is invalid, you are prompted to specify a valid directory path.



17.6.3 Downloading the Feed Instantly

You can immediately download the feed for a selected download configuration, regardless of the time interval scheduled for the selected download configuration.

1 Launch the Sentinel Control Center.

For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.

2 Launch the Download Manager:

- ♦ If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Download Manager** or click the  icon on the toolbar.
- ♦ If the **Configuration** menu is enabled, click the **Configuration** menu > **Download Manager** or click the  icon on the toolbar.

3 Click **Download Now**.

The Advisor data feed is downloaded if you have specified the appropriate URL.

The **Download Now** button is disabled when the download is in progress for the selected configuration.

17.6.4 Audit Events for the Download Manager

The Download Manager generates an audit event whenever you perform any of the following actions:

- ♦ Create a download configuration
- ♦ Edit a download configuration
- ♦ Delete a download configuration
- ♦ Download the feed

NOTE: An audit event is generated regardless of whether the download status indicates success or failure.

You can create appropriate correlation rules to receive notifications of these audit events through e-mail. For more information on creating correlation rules, see “[Correlating Event Data](#)” in the *NetIQ Sentinel User Guide*.

17.6.5 Downloading the Advisor Feed Manually

- 1 Log in to the Novell download Web site (<https://secure-www.novell.com/sentinel/download/advisor/>) (<https://secure-www.novell.com/sentinel/download/advisor/>) by using your Novell eLogin user name and password.

The Novell eLogin username and password must be associated with the Advisor license.

- 2 Download all the .zip and .md5 files.
- 3 Copy the downloaded feed files to the Sentinel server.

To process the downloaded feed, you must provide the location where you have saved the feed in the Advisor window. The default location is `/var/opt/novell/sentinel/data/updates/advisor`.

17.7 Viewing the Advisor Status

The Advisor Status window lists the products NetIQ supports for Advisor and also displays the status of the last five feed files that have been processed or are being processed.

The Advisor Status window displays the Advisor information only if the feed files are processed and loaded into the database.

To view the Advisor status:

- 1 Launch the Sentinel Control Center.
For more information, see [Section 1.2, “Accessing the Sentinel Control Center,”](#) on page 18.
- 2 Click **Advisor** or click **Advisor Status** in the **Navigator** panel.

Figure 17-3 Advisor Status

The screenshot shows a window titled "Advisor Status" with two main sections: "Supported Products" and "Status".

Supported Products Table:

Product Name	Product Type	Number of Signatu...	Last Update
Bugtraq	Knowledge Base	35077	Fri Jan 22 03:10:22 I...
Cisco Secure IDS	IDS	4414	Wed Jan 20 10:34:2...
CVE	Knowledge Base	42188	Wed Jan 20 10:34:2...
eEye Retina	Vulnerability	9578	Sat Jan 23 15:10:19 ...
Enterasys Dragon	IDS	8024	Wed Jan 20 10:34:0...
Intrusion.com	IDS	4545	Fri Jan 15 03:51:02 I...
ISS Detection Tools	IDS	3142	Fri Jan 15 03:51:03 I...

Status Table:

Feed File Name	Process Start	Process End
advnxsfeed.12.zip	Mon Jan 25 15:10:19 IST 2010	Mon Jan 25 15:11:09 IST 2010
advnxsfeed.11.zip	Sun Jan 24 15:10:18 IST 2010	Sun Jan 24 15:10:50 IST 2010
advnxsfeed.10.zip	Sat Jan 23 15:10:18 IST 2010	Sat Jan 23 15:10:33 IST 2010
advnxsfeed.9.zip	Fri Jan 22 15:10:18 IST 2010	Fri Jan 22 15:10:32 IST 2010
advnxsfeed.8.zip	Fri Jan 22 03:10:33 IST 2010	Fri Jan 22 03:10:40 IST 2010

At the bottom, there is a "Ready" status indicator, a "Refresh" button, and a timestamp "Refreshed At: 1/25/10 3:10:49 PM".

Table 17-2 Advisor Status

Fields	Description
Product Name	Name of the product supported by NetIQ for Advisor. For example: Cisco Secure IDS and Enterasys Dragon Host Sensor.
Product Type	Shows whether the product type is a Vulnerability, Intrusion Detection System (IDS), or Firewall.
Number of Signatures	Shows the number of signatures for the product by Nexus.
Last Update	Time stamp indicating when the product was last updated.
Feed File Name	Shows the name of the feed files that have been processed and are currently being processed.
Process Start	Time stamp indicating when processing the feed file started.
Process End	Time stamp indicating when processing the feed file finished. The process end time is blank if there is an error that halts processing or if processing is still in progress.

17.8 Viewing the Advisor Data

Make sure you have completed the following prerequisites to view the Advisor data:

- ♦ The Advisor feed must be up-to-date, processed, and loaded into the Sentinel database.
- ♦ The selected event is from a product supported by Advisor and has the Vulnerability field value set to 1.

You can view the Advisor data from the following options:

- ♦ [Section 17.8.1, “Active Views,” on page 172](#)
- ♦ [Section 17.8.2, “Search results in the Web Interface,” on page 172](#)

17.8.1 Active Views

- 1 In the Sentinel Control Center, click **Active Views**.
- 2 In the real-time events table, right-click an event that has the Vulnerability field value set to 1.
- 3 Click **Analyze > Advisor Data**.

17.8.2 Search results in the Web Interface


- 1 In the Sentinel Web interface, click **Filters > Exploit Detected Events** or specify `vul : 1` in the **Search** field, then click **Search**.

All events that are likely to have exploited a known vulnerability are displayed.

You can also view the report for these events by selecting the desired events, then clicking **Event Operations > View Advisor report**.

17.9 Resetting the Advisor Password

If you have configured automated downloads for Advisor and if your Novell eLogin password changes, you need to also change your Advisor password. You can change your Advisor eLogin password by using the Download Manager feature.

- 1 Open the Download Manager by doing one of the following:
 - ♦ Click **Configuration > Download Manager**.
 - ♦ Click the Download Manager  icon on the toolbar.
 - ♦ In the Advisor window, click **Launch Download Manager**.
- 2 Select the download configuration for which you want to change the password, then click **Edit**.
The Edit window is displayed.
- 3 Specify the new password in the **Password** field.
- 4 (Optional) Click **Validate** to validate the URL and the login credentials.
The URL and its credentials are validated and a confirmation message is displayed. If the validation fails, you must provide a valid URL and the login credentials.
- 5 Click **Save** to save the configuration settings.

17.10 Deleting the Advisor Data

You can delete the Advisor data from the Sentinel database by running the `clean_db` script. For more information on running the `clean_db` script, see [Appendix A, “Command Line Utilities,” on page 239](#).

V Monitoring Your Network

You can configure Sentinel to analyze your events, detect patterns, set up baselines, configure workflows to act on the events, manage and remediate threats, visualize network activities, and more. For more information about using these features, see the [NetIQ Sentinel User Guide](#).

This section provides information about the following:

- ♦ [Chapter 18, “Searching and Reporting Events in a Distributed Environment,”](#) on page 177
- ♦ [Chapter 19, “Visualizing Network Traffic,”](#) on page 189

18 Searching and Reporting Events in a Distributed Environment

The Sentinel Distributed Search feature enables you to search for events and run reports not only on your local Sentinel server, but also on other Sentinel and Sentinel Log Manager servers distributed across the globe.

- ♦ [Section 18.1, “Overview,” on page 177](#)
- ♦ [Section 18.2, “Configuring Servers for Distributed Searching and Reporting,” on page 179](#)
- ♦ [Section 18.3, “Searching for Events,” on page 182](#)
- ♦ [Section 18.4, “Managing the Distributed Search Results,” on page 182](#)
- ♦ [Section 18.5, “Viewing the Search Activities,” on page 183](#)
- ♦ [Section 18.6, “Running Reports,” on page 184](#)
- ♦ [Section 18.7, “Managing the Distributed Setup Configuration,” on page 184](#)
- ♦ [Section 18.8, “Troubleshooting,” on page 185](#)

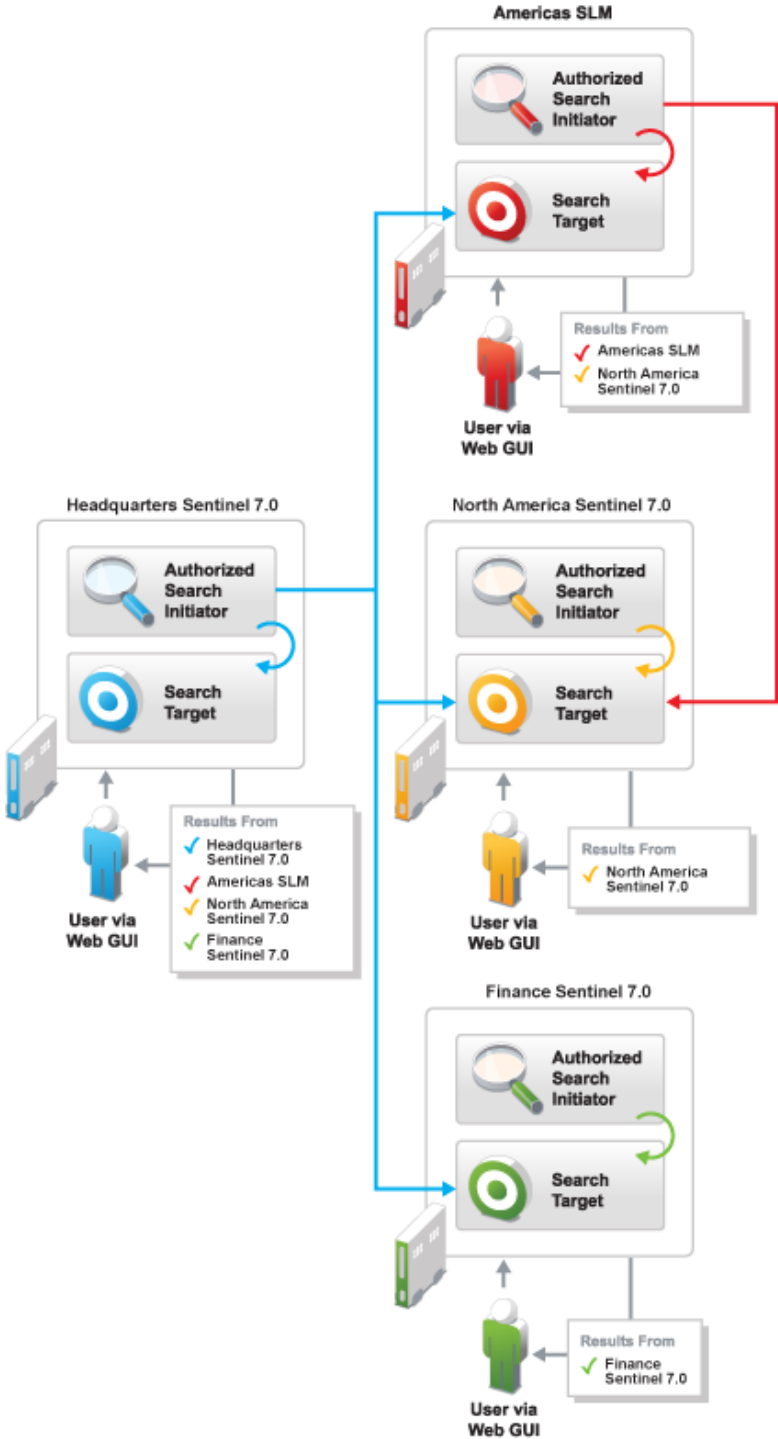
18.1 Overview

The Distributed Search feature facilitates searching events and reporting event data from local and remote Sentinel servers. When you are working with multiple servers, you can perform a search or run a report on just one server and have it automatically run a search or report across the selected remote servers. The server on which the search is initiated is referred to as the search initiator, and the remote servers are referred to as the search targets or target servers.

When you run a search or report on the search initiator, search queries are sent to each selected target server. The target server authenticates the search initiator server, using a password that is exchanged during configuration. Event data is returned to the search initiator, where it is merged, sorted, and consolidated for presentation. Individual search results display the target servers from which they were received. The search status for each server is available for viewing and troubleshooting.

[Figure 18-1](#) shows how you can set up the Sentinel servers across the globe for distributed searching and reporting.

Figure 18-1 Distributed Search



18.2 Configuring Servers for Distributed Searching and Reporting

To configure a search initiator for distributed search, you must first enable distributed searching on the search initiator server.

After you enable distributed search, you need to add target servers to the search initiator server. If you know the administrator user name and password for the target server, you can add the target server directly from the search initiator.

If you do not know the administrator user name and password for a target server, you can set up the search initiator with an opt-in password that allows administrators of target servers to add their target servers to the search initiator. When you do this, administrators of target servers do not need to share their user names and passwords with you. You must share the opt-in password with the target server administrator.

- ♦ [Section 18.2.1, “Enabling Distributed Search,” on page 179](#)
- ♦ [Section 18.2.2, “Using the Administrator Credentials to Add a Search Target Server,” on page 179](#)
- ♦ [Section 18.2.3, “Using the Opt-in Password to Add a Search Target Server,” on page 180](#)

18.2.1 Enabling Distributed Search

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Search Setup** in the toolbar, then click **Search Targets**.
- 3 Select **This local server and other search targets**.
- 4 Do one of the following to add target servers to your search initiator:
 - ♦ If you are the administrator of the search initiator and you know the administrator user name and password on the target server, continue with [Section 18.2.2, “Using the Administrator Credentials to Add a Search Target Server,” on page 179](#).
 - ♦ If you are the administrator of the search initiator and you do not know the administrator user name and password on the target server, you can set an opt-in password to allow administrators of target servers to add their target servers to the search initiator. Continue with [Section 18.2.3, “Using the Opt-in Password to Add a Search Target Server,” on page 180](#).

18.2.2 Using the Administrator Credentials to Add a Search Target Server

If you are the administrator of the search initiator and you know the administrator user name and password on the target server, you can add the target server while you are logged in to your search initiator server.

IMPORTANT: You should ensure that the target server that you add is able to communicate with the search initiator. The target server should be able to communicate through TCP/IP. The IP address or host name of the target server must be accessible through firewalls, NATs, etc. You can use the ping

command to ensure that there is communication from both ways. If there is a communication failure between the servers, an error is displayed in the extended status page. For more information, see [Section 18.4, “Managing the Distributed Search Results,” on page 182](#).

- 1 If you are continuing from [Section 18.2.1, “Enabling Distributed Search,” on page 179](#), skip to [Step 5](#); otherwise, continue to [Step 2](#).

- 2 Log in to the Sentinel Web interface as a user in the administrator role.

- 3 Click **Search Setup** in the toolbar, then click **Search Targets**.

- 4 Select **This local server and other search targets**.

- 5 Click the **Create Search Target** link.

- 6 Specify the following information:

IP Address/DNS Name: IP address or the DNS name of the target server.

Port: Port number of the target server. The default port number is 8443. The target server and search initiator do not need to be on the same port.

User Name: User name to log in to the target server. This must be a user with administrator privileges.

Password: Password associated with the user name.

- 7 Click **Login**, then click **Accept** after verifying that the certificate information is correct.

- 8 Use the following information to configure the target server:

The Create Search Target page displays a lists of the various proxy roles on the target server.

Name: Specify a descriptive name that you want to give to the search target.

This helps you to easily identify the target server by a name instead of by its IP address or DNS name.

Search Proxy Role: Select a search proxy role that you want to assign to the search initiator.

When the search initiator makes search requests to the target server, the proxy role's security filter is used when performing the search. Only events that pass the proxy role's security filter are returned to the search initiator server.

Only roles that have the `Proxy for Authorized Search Initiators` permission are listed. This permission is required for the target server to accept and process incoming search requests from the search initiator server.

- 9 Click **OK**.

The server information is listed in the **Search Targets** list.

You can now search events or view event reports from the target server. For more information, see [Section 18.3, “Searching for Events,” on page 182](#) and [Section 18.6, “Running Reports,” on page 184](#).

18.2.3 Using the Opt-in Password to Add a Search Target Server

In organizations where administrative control of Sentinel servers is decentralized, it might violate the security policy to share administrator passwords. However, Sentinel allows you to share a limited-purpose opt-in password to add target servers, which is more secure than requiring full administrator credentials. If you are not the administrator of the target server, you can set an opt-in password in the search initiator server, then provide the opt-in password to the target server administrators to allow them to opt in to the search initiator server.

When a target server opts in to the search initiator, a message is sent to the search initiator server requesting that it be added to the list of target servers maintained by the search initiator server. The request authorizes the search initiator to access event data on the target server. The search initiator

requires an opt-in password to verify that the opt-in request has originated from a valid target server. During the opt-in process, the search initiator and the target server exchange the appropriate password, which allows the target server to authenticate the search requests from the search initiator.

This procedure is similar to adding a target server, but it is done from the target server instead of the search initiator server.

- ♦ [“Setting the Opt-In Password” on page 181](#)
- ♦ [“Authorizing a Search Initiator Server” on page 181](#)

Setting the Opt-In Password

- 1 Log in to the search initiator server as a user in the administrator role.
- 2 Click **Search Setup** in the toolbar.
The Search configuration page that is displayed has two tabs: **Search Targets** and **Authorized Search Initiators**. Ensure that the **Search Targets** tab is selected.
- 3 Select **This local server and other search targets**.
- 4 Click **Set Opt-in Password**.
- 5 Specify the opt-in password, then click **Set Password**.
- 6 Continue with [“Authorizing a Search Initiator Server” on page 181](#) to add the target server to the search initiator.

Authorizing a Search Initiator Server

- 1 Log in to the target server as a user in the administrator role.
- 2 Click **Search Setup** in the toolbar.
The Search configuration page that is displayed has two tabs: **Search Targets** and **Authorized Search Initiators**. The **Search Targets** tab is selected by default.
- 3 Select the **Authorized Searcher Initiators** tab.
- 4 Check the **Allow these authorized search initiators to search this server as a search target** box.
- 5 Click the **Create** link.
The Create Authorized Search Initiator page is displayed.
- 6 Specify the following information:
IP Address/DNS Name: The IP address or the DNS name of the search initiator.
Port: Port number of the search initiator. This is the port number on which the search initiator listens for incoming opt-in requests. The default port number is 8443.
Opt-in Password: The opt-in password that you configured on the search initiator. You must obtain this password from the administrator of the search initiator.
- 7 Click **OK**.
The Confirm Certificate page is displayed.
- 8 Verify the certificate information, then click **Accept**.
The Create Authorized Search Initiator page is displayed that lists the various proxy roles on the search target servers.
- 9 In the **Name** field, specify a descriptive name that you want to give to the search initiator server.
This helps you to easily identify the search initiator server by a name instead of by its IP address or DNS name.

- 10 Select a proxy role that you want to assign to the search initiator.

When the search initiator makes search requests to the target server, the proxy role's security filter is used when performing the search. Only events that pass the proxy role's security filter are returned to the search initiator.

Only roles in the target server that have the `Proxy for Authorized Search Initiators` permission are listed. This permission is required for the target server to accept and process incoming search requests from the search initiator.

- 11 Click **OK**.

The search initiator is added to Authorized Search Initiators list and is enabled by default.

The target server gets is also added under the Search Targets list in the search initiator server. Alternatively, you can click the **Refresh** link to see the target server in the Search Targets list.

18.3 Searching for Events

Searching for events in a distributed environment is similar to how you perform a search on your local server, except that you perform the search on the selected target servers and can also include your local server.

- 1 Log in to the search initiator server as a user with Search Remote Targets permission.
- 2 Click **New Search**.
- 3 Click the **Targets** link under the **Search** field.

A page is displayed that lists the target servers that you have added, including the local server. The target servers that are disabled are also displayed, but they are dimmed.

- 4 Select the check boxes next to the target servers on which you want to perform a search, then click **OK**.
- 5 Specify the search criteria in the search field, then click **Search**.

If you do not specify any search criteria, the search initiator server runs a default search for all events with severity 3 to 5.

The Search Results page displays the events from the selected target servers and the local server (if selected). The search results are filtered through the combination of the security filter and permissions of the logged-in user and the security filter and permissions of the search proxy role on the target servers. For information on the distributed search results, see [Section 18.4, "Managing the Distributed Search Results," on page 182](#).

18.4 Managing the Distributed Search Results

The Search Results page displays the events from the selected target servers and the local server, based on the search criteria you specified. Each event displays the target server information from which the event is being retrieved.

You can expand the event results to see the details by clicking the **All details++** link.

For non-internal events, the **get raw data** link is displayed. You can view the raw data only if your role's security filter is set to view all event data.

NOTE: For search results that come from the target servers, the role that is used to retrieve raw data is not the role of the logged-in user that is performing the search on the search initiator server, but the role that is assigned to the search initiator server on the target server.

You can view the status of the search in the extended status page while a search is in progress as well as when the search has finished. To access the extended status page, click the **Displaying N of M events from X targets** link that appears at the top of the refinement panel.

The extended status page displays the following information:

- ♦ **Target Name:** The descriptive name of the target server. If you did not specify a descriptive name for the target server, this field displays the IP address or DNS name of the target server.
- ♦ **Events Available:** Indicates the number of events that have actually been retrieved from the target server. The value is displayed as **N of M events available**, where N is the number of events that have been retrieved so far and M is the total number of events on the target server that match the search criteria.
- ♦ **Retrieval Rate (EPS):** An approximate rate of how fast the events were retrieved from a specific target server.
- ♦ **Status:** Displays the error messages, if any (generally in red). In addition to error messages, this field also displays the status of the search.
 - ♦ **Running:** Indicates that the search is still running on the search target server.
 - ♦ **Buffering events for display:** Indicates that the search is finished on the target server, but the search initiator server is retrieving events from the target server and buffering them for display.
 - ♦ **Paused buffering events for display:** Indicates that the search is finished on the target server, and the search initiator has paused while retrieving events from the search target. The search initiator reads ahead a few pages from the last page that you scrolled down to. When it has buffered enough pages ahead, it pauses so that events are not buffered unnecessarily.
 - ♦ **Searching, paused buffering events for display:** This is similar to pausing and buffering events for display, except that the search is not yet complete on the target server.
 - ♦ **Done buffering:** Indicates that the search is complete on the target server, and all of the results are retrieved by the search initiator and queued for display.

You can further refine the distributed search results and perform various actions based on your requirements. For more information, see “[Searching Events](#)” in the *NetIQ Sentinel User Guide*.

18.5 Viewing the Search Activities

You can view the search activities that are being done on the target server by the search initiator server. You can see what queries are being done and how frequently they are being done. Based on the search activity, you might want to assign a more/less restrictive proxy role or even disable access to the target server.

You can also refine the search activity query. For example, you can change the date range to see what queries have been performed today or yesterday or in the last hour, or you can drill down to see the queries that were made by particular users on the search initiator.

- 1 Log in to the target server as a user in the administrator role.
- 2 Click **Search Setup** in the toolbar.

The Search configuration page that is displayed has two tabs: **Search Targets** and **Authorized Search Initiators**. The **Search Targets** tab is selected by default.

- 3 Select the **Authorized Searcher Initiators** tab.

A list of search initiator servers is displayed under the Authorized Search Initiators list.

- 4 Click the **Search Activities** link for the search initiator server for which you want to view the search activities.

The search activities page is displayed that lists the audit events that are retrieved from all of the distributed search requests the target server has received from that particular search initiator.

18.6 Running Reports

Running reports in a distributed environment is similar to running reports on your local server, except that you select the target servers from which you want view the reports while specifying the report parameters.

- 1 Log in to the search initiator as a user with Search Remote Targets permission.

- 2 From the Reports section, select the report you want to run, then click **Run**.

The Run Report page is displayed.

- 3 Click the link for **Targets**.

A page is displayed that lists the target servers that you have added, including the local server. The target servers that are disabled are also displayed, but they are dimmed.

- 4 Select the target servers from which you want to view the reports, then click **OK**.

- 5 Specify the other parameters for the report.

For more information on these parameters, see “[Reporting](#)” in the *NetIQ Sentinel User Guide*.

- 6 Click **Run**.

A report results entry is created and listed under the selected report. For more information on managing reports, see “[Reporting](#)” in the *NetIQ Sentinel User Guide*.

18.7 Managing the Distributed Setup Configuration

- ♦ [Section 18.7.1, “Editing the Search Target Server Details,”](#) on page 184

18.7.1 Editing the Search Target Server Details

- 1 Log in to the search initiator as a user in the administrator role.

- 2 Click **Search Setup** in the toolbar.

- 3 Select the **Search Targets** tab.

A list of search target servers is displayed under the Search Targets list.

- 4 Click the **Edit** link for the search target server for which you want to modify the details., then edit the information.

Do not change the IP address unless the IP address of the search target server has actually changed. The **Edit** page is not intended to allow you to change this target server entry to be a different search target server. To add a new target server, see [Section 18.2.2, “Using the Administrator Credentials to Add a Search Target Server,”](#) on page 179. If you accidentally change the IP address or DNS name to a different server, change it back to the correct server, or authentication of search requests fails.

- 5 (Optional) To change the proxy role on the target server as necessary:
 - 5a Click **View/Change**.
 - 5b Log in to the target server.
 - 5c Select a proxy role, then click **OK**.
- 6 Click **Save**.

18.8 Troubleshooting

You can perform some basic troubleshooting to ensure that you have successfully configured the search initiator for distributed search. This section lists the most common issues and the probable causes for these issues.

- ♦ [Section 18.8.1, “Permission Denied,” on page 185](#)
- ♦ [Section 18.8.2, “Connection Down,” on page 185](#)
- ♦ [Section 18.8.3, “Unable to View Raw Data,” on page 186](#)
- ♦ [Section 18.8.4, “Problems Adding Search Target,” on page 186](#)
- ♦ [Section 18.8.5, “Certain Events Are Only Visible from the Local System,” on page 186](#)
- ♦ [Section 18.8.6, “Cannot Run Reports on the Target Servers,” on page 186](#)
- ♦ [Section 18.8.7, “Different Users Get Different Results,” on page 186](#)
- ♦ [Section 18.8.8, “Cannot Set the Admin Role as the Search Proxy Role,” on page 186](#)
- ♦ [Section 18.8.9, “Error Logs,” on page 186](#)

18.8.1 Permission Denied

After doing a distributed search, check the extended status page to view the search status. If the search is not successful, check the following possible causes:

- ♦ The target server administrator might have disabled distributed searching on the target server. To enable distributed search on the target server, see [Step 4 in “Authorizing a Search Initiator Server” on page 181](#).
- ♦ The target server administrator might have disabled the search initiator server for distributed searching. Ensure that the search initiator server is enabled in the target server. For more information, see [“Authorizing a Search Initiator Server” on page 181](#).
- ♦ The role that you used for connecting might not have the `Search Remote Targets` permission.

18.8.2 Connection Down

- ♦ Network issues in your organization.
- ♦ Sentinel servers or Sentinel services might be down.
- ♦ Connection time-out.
- ♦ The IP address or the port number of the target server has changed, but the search initiator configuration might not be updated.

18.8.3 Unable to View Raw Data

The Proxy group that is assigned to the search initiator might not have the `view all events` permission to view the raw data.

18.8.4 Problems Adding Search Target

The search initiator server and target server might not be communicating with each other. Ensure that the firewall and NAT are set up properly to allow communication in both directions. Ping both ways to test.

18.8.5 Certain Events Are Only Visible from the Local System

You might not be able to view the events from the target servers for one of the following reasons:

- ♦ The trial license might be expired. You must purchase an enterprise license to reactivate this feature to view the events from the target servers.
- ♦ The user who has logged in to the search initiator has one set of permissions on the local data such as view all data, view system events, security filter settings, and so on. The search proxy group has another set of permissions, possibly more restrictive. Therefore, certain types of data, such as raw data, system events, and PCI events, might be returned only from the local system and not the target server.

18.8.6 Cannot Run Reports on the Target Servers

The trial license might be expired. You must purchase an enterprise license to reactivate this feature to run the reports from the target servers.

18.8.7 Different Users Get Different Results

Different users might have different security filters or other permissions and therefore get different results from a distributed search.

18.8.8 Cannot Set the Admin Role as the Search Proxy Role

This is by design, for security reasons. Because the data viewing rights for the admin are unrestricted, it is not desirable to allow the admin role to be the search proxy role.

18.8.9 Error Logs

You can also determine the cause of a search failure by examining the log file on the search initiator server. The default location for the log file is `/var/opt/novell/sentinel/log`. For example, you might see one of the following messages:

```
Invalid console host name 10.0.0.1
```

```
Error sending target request to console host 10.0.0.1
```

```
Error getting certificate for console host 10.0.0.1
```

```
Authentication credentials in request to opt-in to console 10.0.0.2 were rejected
```

```
Request to opt-in to console 10.0.0.2 was not authorized
```

Error sending target request to console host 10.0.0.1

19 Visualizing Network Traffic

To perform a complete investigation and analysis of a security event, you might want to monitor network activities in detail. Sentinel helps you monitor your enterprise network by collecting, visualizing, and analyzing network flow data.

Network flow data helps you with the following types of analysis:

- ♦ Monitor network activities in near real time and those that occurred at the time of a security event for a given IP address.
- ♦ Analyze the change in network activity before and after a security event.
- ♦ Determine the impact of a security event on the resources of an affected system. For example, whether the network traffic into or out of a host changed after the security event.
- ♦ Track network propagation behavior for attacks such as viruses, bots, and DDOS.
- ♦ Remediate issues and verify the solution by network flow inspection. For example, you can verify whether you need to create a firewall rule to prevent such security issues.

This chapter provides information about the following:

- ♦ [Section 19.1, “Understanding Network Flow Data Collection,” on page 189](#)
- ♦ [Section 19.2, “Collecting Network Flow Data for Specific Tenants,” on page 191](#)
- ♦ [Section 19.3, “Visualizing and Analyzing Network Flow Data,” on page 191](#)
- ♦ [Section 19.4, “Viewing and Managing the Network Flow Data Collection,” on page 192](#)
- ♦ [Section 19.5, “Customizing the NetFlow Configuration,” on page 192](#)

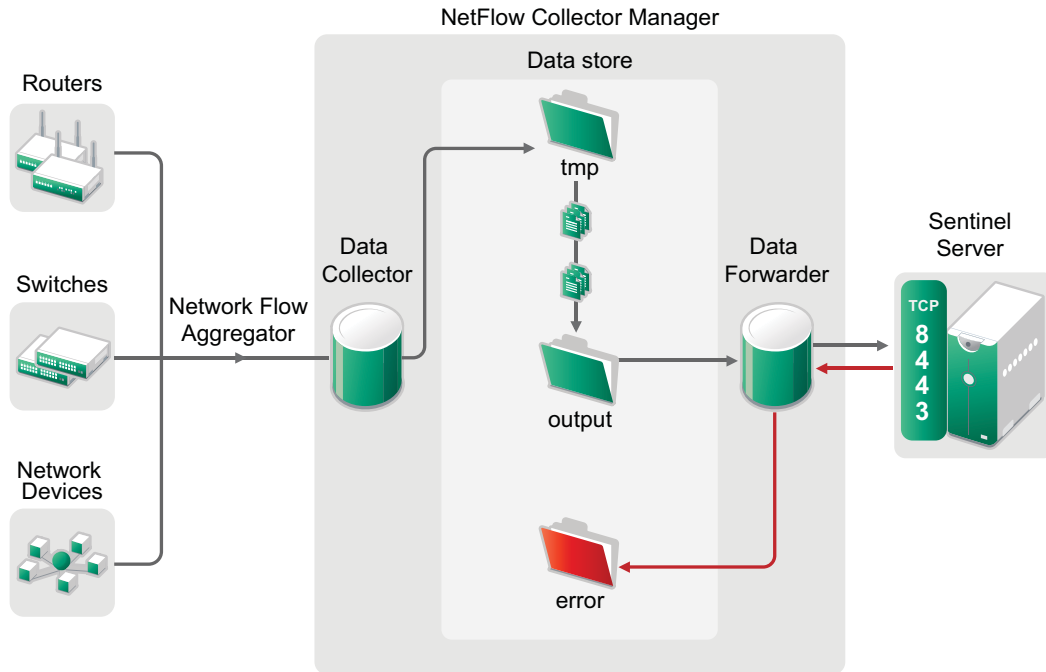
19.1 Understanding Network Flow Data Collection

Sentinel supports NetFlow versions v5, v9, and v10 (IPFIX). The NetFlow Collector Manager collects network flow data from supported routers, switches, and other network devices.

19.1.1 NetFlow Collector Manager

The NetFlow Collector Manager collects and processes network flow data. The following figure illustrates how NetFlow Collector Manager works:

Figure 19-1 NetFlow Collector Manager



The NetFlow Collector Manager includes the following components:

- ♦ **Network flow aggregator:** Collects and aggregates the raw network flow data from the connected devices.
- ♦ **Data collector:** Requests the network flow data from the network flow aggregator store every 10 minutes (configurable) and sends the data to the data store.
- ♦ **Data store:** Stores batches of network flow data ready to be sent to Sentinel.
- ♦ **Data forwarder:** Forwards any network flow data waiting in the data store by using the HTTPS protocol.
 - ♦ If the Sentinel server is not reachable, the data remains in the output store until it reaches the maximum queue size. The data forwarder continues to send the buffered data when the Sentinel server is available.
 - ♦ The number of files that the output store can store depends on the queue size. You can configure the queue size in the `/etc/opt/novell/sentinel/config/netflow-collector-configuration.properties` file. For more information, see [“Customizing the NetFlow Configuration” on page 192](#).
 - ♦ If an error occurs during the data transfer to Sentinel, the data forwarder will move the associated batch files to the error store. You can send this data again by manually copying the data from the error store to the output store.

To install the NetFlow Collector Manager, you must have administrator privileges or the Send NetFlow data permission. For more information, see [“NetFlow Collector Manager Installation”](#) in the *“NetIQ Sentinel Installation and Configuration Guide.”*

19.1.2 Supported NetFlow Fields

The NetFlow Collector Manager collects only specific fields from the network flow data. The following table provides information about the specific fields that the NetFlow Collector Manager collects:

Table 19-1 Supported NetFlow Fields

Field	Description
Start time	Start time of network flow from an initiator to the target point.
End time	The time at which network flow reached the target point.
SIP	The IP address of the source from which network flow started.
SP	The port number of the initiator.
DIP	The IP address of the target that received network flow.
DP	The port number of the target.
Protocol	The IP protocol used for the communication between the initiator and the target IP.
Number of flows	The number of flows that happened between the initiator and the target.
Number of packets	The number of packets transferred between the initiator and the target.
Number of bytes	The number of bytes transferred between the initiator and the target.

19.2 Collecting Network Flow Data for Specific Tenants

If you are in a multi-tenant environment and need to segregate data from different divisions or customers, you must install individual NetFlow Collector Managers for each tenant to monitor the network on each tenant. After you install the NetFlow Collector Manager, you must assign the tenant name appropriately so that network flow data for specific tenants are stored in separate network flow tables for each tenant.

To collect network flow data for specific tenants:

- 1 Log in to the Sentinel Web console as an administrator.
- 2 Click **Collection** > **NetFlow**.
- 3 Click **Edit** for the NetFlow Collector Manager for which you want to assign tenant.
- 4 Specify the tenant name in the **Tenant** field.

NOTE: The tenant name should correspond to the TenantName (rv39) event field. For information about TenantName (rv39) field, see **Tips** in the Sentinel Web console.

- 5 Click **Save**.

19.3 Visualizing and Analyzing Network Flow Data

Sentinel provides a graphical representation of the statistical network flow data that helps you identify and analyze suspicious activities in your network. You can view the network flow data in near real time for a specific department, for a specific customer, a specific event, an IP address, the entire network, or a time range.

NOTE: To view network flow data, you must either be an administrator or have the View NetFlow data permission.

For more information about viewing and analyzing network flow data, see the “[Visualizing and Analyzing Network Flow Data](#)” in the *NetIQ Sentinel User Guide*.

19.4 Viewing and Managing the Network Flow Data Collection

You can view the list of NetFlow Collector Managers sending network flow data to the Sentinel server. You can also view details, such as the time stamp that Sentinel last received data, number of network flow records accepted and rejected by the Sentinel server, and so on.

You can manage the NetFlow Collector Managers by performing actions, such as renaming the NetFlow Collector Manager, assigning the tenant name, blocking the data from a specific NetFlow Collector Manager, and deleting the NetFlow Collector Manager.

To view and manage network flow data collection:

- 1 Log in to the Sentinel server as an administrator.
- 2 Click **Collection > NetFlow**.
- 3 Perform the necessary action.

NOTE: Deleting the NetFlow Collector Manager in the Collections Web interface only disconnects the connection between the specific NetFlow Collector Manager from the Sentinel server. To uninstall the NetFlow Collector Manager, see [“Uninstalling the NetFlow Collector Manager”](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

19.5 Customizing the NetFlow Configuration

This section provides information about customizing your NetFlow configuration.

- ♦ [Section 19.5.1, “Tuning the Network Flow Data Collection,”](#) on page 192
- ♦ [Section 19.5.2, “Changing the User Name and Password in the NetFlow Collector Manager,”](#) on page 193

19.5.1 Tuning the Network Flow Data Collection

You can configure the following NetFlow data collection parameters in the `/etc/opt/novell/sentinel/config/netflow-collector-configuration.properties` file:

Table 19-2 NetFlow Data Collection Configurable Parameters

Parameter	Default Value	Description
netflow.collection.interval	10	Specifies the data collection interval in minutes.
netflow.forwarder.batchsize	6	Number of aggregated records (in thousands) per batch, sent from the NetFlow Collector Manager to the Sentinel server.
rest.endpoint.host	<IP_address>	The IP address or hostname of the Sentinel server that receives network flow data. If you are specifying the hostname, you should specify the fully qualified domain name (FQDN).
netflow.collector.id	<Collector_Manager_ID>	The UUID of the NetFlow Collector Manager. Sentinel identifies the NetFlow Collector Manager by this ID. If you have assigned the NetFlow Collector Manager to a specific tenant, Sentinel uses the UUID to map data from the NetFlow Collector Manager to a specific tenant. Therefore, you should not modify the UUID.

Parameter	Default Value	Description
netflow.collection.pool_size	10	Number of simultaneous collection processes that the NetFlow Collector Manager runs. The maximum value is 20.
netflow.reader.nfdump.daemon.port	3578	The port used by the NetFlow Collector Manager to collect raw network flow data from network devices.
rest.endpoint.port	8443	The port used by the NetFlow Collector Manager to connect to the Sentinel server to send data.
netflow.rawdata.lifetime	1	<p>The retention period (in weeks) for the raw network flow data to be stored in the NetFlow Collector Manager.</p> <p>The minimum value is 10M (10 minutes). You can also set the value in M - Minutes, H - Hours, d - Days, and w - Weeks.</p> <p>NOTE: M, H, d, and w are case sensitive.</p>
netflow.error.queueSize	5760	The number of output files in queue to be sent to Sentinel. The maximum value is 5760. When the queue size reaches the maximum value, the data collector deletes the oldest output file.

To configure network flow data collection:

- 1 Log in to NetFlow Collector Manager computer.
- 2 Open the `/etc/opt/novell/sentinel/config/netflow-collector-configuration.properties` file in an editor.
- 3 Change the parameters as necessary. For information about the parameters, see [Table 19-2, “NetFlow Data Collection Configurable Parameters,”](#) on page 192.

NOTE: For security reasons, the NetFlow Collector Manager stores the encrypted password in the `netflow-collector-configuration.properties` file and you cannot change the password in this file. To change the user name and password, you must run the `./configure.sh` script. For more information, see [“Changing the User Name and Password in the NetFlow Collector Manager”](#) on page 193.

- 4 Save the changes.
- 5 Restart the NetFlow Collector Manager.

19.5.2 Changing the User Name and Password in the NetFlow Collector Manager

To change the NetFlow Collector Manager user name and password:

- 1 Log in to the NetFlow Collector Manager computer.
- 2 Change to the following directory:
`/opt/novell/sentinel/bin`
- 3 Run the NetFlow configuration script as follows:
`/opt/novell/sentinel/setup/configure.sh`
- 4 Change the user name and password.
- 5 Restart the NetFlow Collector Manager.

VI Managing Solution Packs

This section provides information about using the built-in solution packs and creating your own solution packs.

- ♦ [Chapter 20, “Using Solution Packs,” on page 197](#)
- ♦ [Chapter 21, “Creating Solution Packs,” on page 211](#)

20 Using Solution Packs

Solution Packs allow NetIQ partners and customers to create and easily manage solutions to specific business problems.

- ♦ [Section 20.1, “Overview,” on page 197](#)
- ♦ [Section 20.2, “Solution Pack Components,” on page 198](#)
- ♦ [Section 20.3, “Using the Import Plugin Wizard to Import a Solution Pack,” on page 199](#)
- ♦ [Section 20.4, “Using the Solution Manager,” on page 200](#)
- ♦ [Section 20.5, “Installing and Managing Solution Packs,” on page 202](#)
- ♦ [Section 20.6, “Installing an Edited Solution Pack,” on page 209](#)
- ♦ [Section 20.7, “Solution Designer,” on page 210](#)

20.1 Overview

Solution Packs provide a framework within which sets of content can be packaged into controls, each of which is designed to enforce a specific business or technical policy. The control can use any of the detection, filtering, alerting, and response features of Sentinel, as well as provide documentation on control status and enforcement. By managing the set of content as a unit within the control, the Solution Pack solves dependency problems and simplifies implementation.

Controls within a Solution Pack can include the following types of content:

- ♦ Correlation rule deployments, including deployment status and associated Correlation rules, Correlation actions, JavaScript plug-ins and integrators, and Dynamic Lists
- ♦ Reports
- ♦ Filters
- ♦ Searches
- ♦ iTRAC workflows, including associated roles
- ♦ Event enrichment, including map definitions and event meta tag configuration
- ♦ Other associated files added when the Solution Pack is created, such as documentation, example report PDFs, or sample map files.

For example, a Solution Pack can package content related to governance and regulatory compliance into a comprehensible and easily enforceable framework that is easy to deploy. NetIQ and its partners offer extended Solution Packs for such regulations or other customer needs.

Solution Packs are created with the Solution Designer application. Using this tool, a user creates the Solution Pack, associated controls and documentation, and then associates Sentinel content with each control. The entire package is then exported as a ZIP file. For more information on creating a Solution Pack and adding content to it, see [Section 20.7, “Solution Designer,” on page 210](#).

The ZIP file containing the Solution Pack is first imported into an existing Sentinel system by using the Solution Packs Manager in the Sentinel Control Center. You then use the Solution Manager to install the imported Solution Pack.

The Solution Manager also displays the implementation and testing steps in the Solution Pack and tracks the status of each control. At any time, users can generate a detailed document with the implementation status for each control. You can also use the Solution Manager to install the predefined Sentinel Core Solution Pack.




NOTE: Only users in the administrator role can access the Solution Packs Manager.

20.2 Solution Pack Components

Solution Packs consist of categories, controls, content, and content groups, all in a hierarchy of related components.





The below describes each level in a Solution Pack hierarchy.











Table 20-1 *Solution Pack Hierarchy Levels*

	Solution Pack	A Solution Pack is the root node in the content hierarchy. Each Solution Pack can contain one or multiple Category nodes.
	Category	A category is a conceptual classification. Each Category can contain one or multiple controls.
	Control	A control is another level of classification, which often corresponds to a particular control defined by a set of regulations. Each control can contain one or multiple content groups.
N/A	Content Group	A content group is a set of related content. There are several types of content groups, such as reports, Correlation rules, and event configurations, each with its own icon.

The below describes the types of content groups and what the content that they contain:

Table 20-2 *14-2: Content Groups*

	Event Configuration	A content group that contains a map definition and the configuration of one or more related Sentinel meta tags. This icon is also used for the meta tag configuration definition.
	Map	A map definition instance.
	Workflow	An iTRAC Workflow template and any associated roles. This icon is also used for the iTRAC workflow template itself.
	Role	A role used in a workflow.

	Correlation Rule	A Correlation rule, the namespace in which it is stored, and any associated Correlation actions or Dynamic Lists. This icon is also used for the Correlation rule definition.
	Namespace	A namespace instance in which the Correlation rule is stored
	JavaScript Action Plugin	A JavaScript action plug-in.
	JavaScript Action	A configured JavaScript action instance.
	Integrator Plugin	An Integrator plug-in.
	Integrator	A configured Integrator instance.
	Action	An action configuration for a Correlation action.
	Correlation Rule Deployment	The Correlation rule deployment.
	Report	A report. This icon is also used for the .rpt report file.
	Dynamic List	A Dynamic List.

20.3 Using the Import Plugin Wizard to Import a Solution Pack

Solution Packs are available from several sources. They can be downloaded from the [Sentinel Plugins download Web page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html). They can be provided by a NetIQ partner, or they can be created from content in your own Sentinel system.

The first step in using a Solution Pack is to import the .zip file into the system through the Import Plugin Wizard. When a Solution Pack is imported, the .zip file is copied to the server. The actual contents of the Solution Pack are not available in the target Sentinel system until the controls are installed through the Solution Manager. For more information, see [Section 20.5.2, “Installing Content from Solution Packs,” on page 203](#).

If you import an updated version of a Solution Pack, you are prompted to replace the existing plug-in.

- 1 Launch the Sentinel Control Center as a user who has the administrator role.
- 2 Launch the Solution Packs Manager.
 - ♦ (Conditional) If the Configuration menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Solution Packs** or click the icon in the toolbar.
 - ♦ (Conditional) If the **Configuration** menu is enabled, click the **Configuration** > **Solution Packs** or click the icon in the toolbar.

The Solution Packs Manager window is displayed.

- 3 Click the **Import** icon in the Solution Packs Manager window to open the Import Plugin Wizard.
- 4 Select **Import Solution package plugin file (.zip)**, then click **Next** to display the Choose Plugin Package File window.

- 5 Use the **Browse** button to locate the Solution Pack to import to the plug-in repository. Select a .zip file and click **Open**.
- 6 Click **Next**.

If you selected a solution pack that already exists, the Replace Existing Plugin window displays.
- 7 If you want to replace the existing plug-in, click **Next**.

The Plug-in Details window displays the details of the plug-in to be imported.
- 8 Select the **Launch Solution Manager** check box if you want to install the plug-in after importing the Solution Pack.

If you select the **Launch Solution Manager** check box, the Solution Manager displays the Solution Pack that is imported.
- 9 Click **Finish**.

20.4 Using the Solution Manager

You can use the Solution Manager for several different tasks:

- ♦ Installing a control and the child content for the control into the Sentinel system. When the content is initially installed, its status is Not Implemented.
- ♦ Implementing a control by configuring event source systems and Sentinel to use the content associated with the control. Solution Packs include detailed documentation describing implementation steps.
- ♦ Testing a control to verify the content associated with the control. Solution Packs include detailed documentation describing testing steps.
- ♦ [Section 20.4.1, “Launching the Solution Manager,” on page 200](#)
- ♦ [Section 20.4.2, “Solution Manager Interface,” on page 200](#)

20.4.1 Launching the Solution Manager

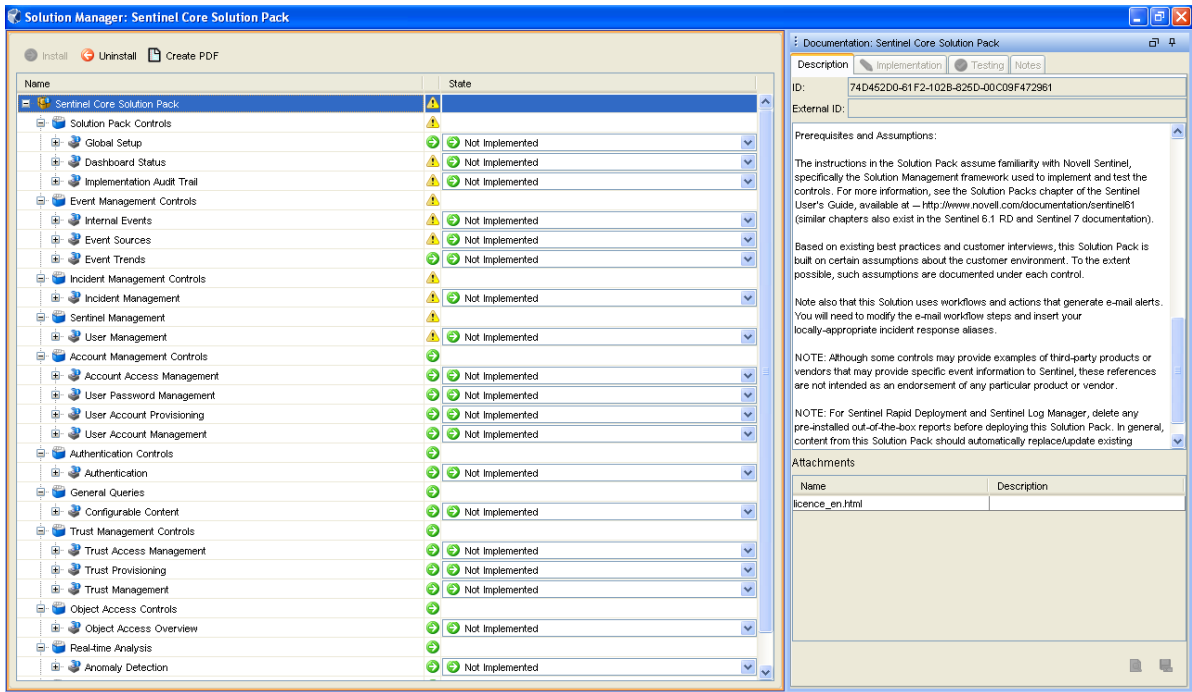
You can launch the Solution Manager by:

- ♦ Selecting **Launch Solution Manager** in the Import Plugin Wizard.
- ♦ Double-clicking the imported Solution Pack in the Solution Packs Manager.
- ♦ In the Solution Packs Manager window, select a Solution Pack and then click the **Open with Solution Manager** icon.

20.4.2 Solution Manager Interface

The Solution Manager window is divided into two panels that contain information about the extracted Solution Pack, along with comparison information, status information, and documentation:

Figure 20-1 Solution Manager



- ◆ “Content Panel” on page 201
- ◆ “Content Comparison” on page 202
- ◆ “Out Of Sync Status” on page 202
- ◆ “Documentation Panel” on page 202

Content Panel

The Content panel provides information about the extracted Solution Pack files. This panel displays a hierarchical view of the category, control, content group, and various types of content. All parent nodes reflect the overall state of the controls they contain. This means that parent nodes have an inherited status based on their child content.

The Content panel has the following columns:



- ◆ **Name:** Displays the name of the node.
- ◆ **Installed:** Indicates whether the content is installed in the target Sentinel system. If not, this column is blank.
- ◆ **State:** This column is available for the control node. This column contain a list with the following values:
 - ◆ **Not Implemented:** The default state when the control is first imported.
 - ◆ **Implemented:** Indicates that you have fully implemented the content by using the associated documentation.
 - ◆ **Tested:** Indicates that you have fully tested the content for this control by using the associated documentation.

NOTE: Because of the regulatory significance of implementing controls, status changes for each control are tracked for auditing purposes.

Content Comparison

When the Solution Pack is opened, the Solution Manager compares the contents of the Solution Pack to other Solution Pack content from different Solution Packs or previous versions of the same Solution Pack.

Table 20-3 *Content Status*

	Installed	Indicates that the content is already installed in the target Sentinel system. The version is the same in the opened Solution Pack and the previously installed Solution Pack.
	Out of Sync	Indicates that a different version of the content is already installed in the target Sentinel system. A difference in name, definition, or description could trigger an Out of Sync status.

Out Of Sync Status

The Out of Sync icon indicates that content in the newly opened Solution Pack differs from a version that was previously installed by another Solution Pack (either a different Solution Pack or a previous version of the same Solution Pack). The name, definition, or description of the content might be different.

NOTE: The Solution Manager only compares content from different Solution Packs (or different versions of the same Solution Pack) for installed content. It does not compare content that has not yet been installed. It also does not compare Solution Pack content to content in the target system, so manual changes to content in the Sentinel Control Manager are not reflected in Solution Manager.

Documentation Panel

The Documentation panel displays descriptive information provided to describe the Solution Pack when it was created in the Solution Designer. For more information on the Solution Designer, see [Section 20.7, “Solution Designer,” on page 210](#).

The following informational tabs, populated and edited in the Solution Designer, are available:

- ♦ **Description:** A description of the selected node. You can view attachments and their descriptions in the **Description** tab.

You can add text to the **External ID** field to refer to specific regulations or corporate IDs.
- ♦ **Implementation:** Instructions for implementing the selected control.
- ♦ **Testing:** Instructions for testing the selected control.
- ♦ **Notes:** Use this tab for any notes related to the control, including user comments on the testing or implementation process.

20.5 Installing and Managing Solution Packs

- ♦ [Section 20.5.1, “Viewing the Contents of a Solution Pack,” on page 203](#)
- ♦ [Section 20.5.2, “Installing Content from Solution Packs,” on page 203](#)
- ♦ [Section 20.5.3, “Configuring Controls,” on page 204](#)

- ♦ [Section 20.5.4, “Implementing a Control,” on page 206](#)
- ♦ [Section 20.5.5, “Testing a Control,” on page 206](#)
- ♦ [Section 20.5.6, “Uninstalling a Control,” on page 207](#)
- ♦ [Section 20.5.7, “Viewing Solution Pack Status,” on page 207](#)
- ♦ [Section 20.5.8, “Deleting a Solution Pack,” on page 209](#)

20.5.1 Viewing the Contents of a Solution Pack

To use the Solution Manager and view the contents of a Solution Pack, a user must belong to the administrator role.

- 1 Click the **Configuration** menu and select **Solution Packs** to display the Solution Packs Manager window:
- 2 Double-click a Solution Pack in the Solution Packs window to display the Solution Manager.

20.5.2 Installing Content from Solution Packs

To use the content of a Solution Pack in the Sentinel Control Center, you must install the Solution Pack or selected controls in a Sentinel System (also known as the “target” Sentinel system).

When you install either a Solution Pack or an individual control, all of the child nodes are installed. Only fully defined controls can be installed. For controls that contain placeholders, the Install option is disabled.

- 1 In the Solution Packs Manager window, double-click a Solution Pack to open the Solution Manager.

Alternatively, you can select a Solution Pack that you want to install and click **Open with Solution Manager** icon.

For more information, see [Section 20.4.1, “Launching the Solution Manager,” on page 200](#).

- 2 Select a Solution Pack or a control that you want to install. Click **Install**.

Alternatively, right-click a Solution Pack or control and select **Install**.

The Install Control Wizard opens. If you select a Solution Pack, all the controls in that Solution Pack display. If you select an individual control, then that control is displayed in the Install Control Wizard window.

- 3 Click **Next** to display the Install Content window.

If Correlation rules and actions are included in the Solution Pack, you need to proceed through several additional screens until you reach the Install Content window. For more information, see [“Correlation Rules and Actions” on page 204](#).

- 4 Click **Install**.

- 5 After the installation is complete, click **Finish**.

If the installation fails for any content item in the control, the Solution Manager rolls back all the content in that control to Uninstalled. In this situation, create a new content or Solution Pack using the Solution Designer, import it and then try installing it

- 6 Configure the control, then implement it according to the instructions in the documentation for the control.
 - ♦ For information on configuring the controls or content you just installed, continue with [Section 20.5.3, “Configuring Controls,” on page 204.](#)
 - ♦ For information on implementing the control in order to use the content, continue with [Section 20.5.4, “Implementing a Control,” on page 206.](#)

Correlation Rules and Actions

Correlation rules are deployed to a specific Correlation Engine. During the control installation, a screen shows the Correlation Engines in the target Sentinel system and the rules that are already running on those engines. Based on the number and complexity of the rules running on the engines, you can decide where you will deploy the Correlation rule.

Correlation rules deploy in an Enabled or Disabled state, depending on their status in the source Sentinel system when the Solution Pack was created.

If an Execute Script Correlation action is associated with the Correlation rule, the Solution Manager attempts to install the associated JavaScript code on all Correlation Engines. If a Correlation Engine is offline, you can still deploy the rule. However, the status of the rule will be disabled in the Sentinel Web interface after importing.

If an Execute Command Correlation action is associated with the Correlation rule, the Solution Manager installs the command and its arguments, but the script or utility must be manually configured on the Correlation Engine. This might require installing the utility, configuring permissions, or manually copying a script file to the proper directory on the Correlation Engine.

In a default installation, the proper directory for the script file is `/opt/novell/sentinel/bin/actions`.

If a JavaScript Action is associated with the Correlation rule, the Solution Manager installs the Action configuration, the Action plug-in, and the associated integrator configuration and Integrator plug-in (if it is needed).

20.5.3 Configuring Controls

- ♦ [“Duplicate Content within a Solution Pack” on page 204](#)
- ♦ [“Content with the Same Name in the Target Sentinel System” on page 205](#)
- ♦ [“Resolving Out of Sync Content” on page 205](#)
- ♦ [“Copying a Map File” on page 205](#)

Duplicate Content within a Solution Pack

If two separate controls contain identical content and one control is installed successfully, the status of the duplicate content in the other control is changed to Installed. The remaining child nodes in the second control stay uninstalled.

Each content item is only installed once. If the same content item (for example, an iTRAC workflow or a Correlation rule) is included in more than one control, it is only installed once. Therefore, if you install one of those controls, the content displays with an installed status in the other control. In this scenario, the Solution Manager might show that the content for the second control is only partially installed.

Content with the Same Name in the Target Sentinel System

If the Solution Manager detects content with the same name but a different unique identifier in the target Sentinel system, the Solution Manager installs the content with a unique ID appended to the name. For example, the rule from the Solution Pack might be named Unauthorized Firewall Change (1). The existing rule in the Sentinel system is unchanged.

NOTE: To prevent confusion for end users, NetIQ recommends that one of these rules be renamed.

Resolving Out of Sync Content

Out of Sync status indicates that a different version of the content in the Solution Pack has been installed in the Sentinel target system by another Solution Pack or a previous version of the same Solution Pack. The Solution Manager only compares content from different Solution Packs (or different versions of the same Solution Pack) for installed content. Before you implement a control, you need to resolve this out of sync status.

- 1 Open the Solution Pack in the Solution Manager for which you want to resolve the out of sync content.
- 2 Select the out of sync content (not the control or category) in the Solution Manager.
- 3 Select the content, right-click, then select **Out of synchronization content detail**.
A message displays information about which Solution Pack is the source of the out of sync content.
- 4 Compare the content in the two Solution Packs to determine which version you want to keep.
- 5 Uninstall the control from the Solution Pack that you do not want to keep. For more information, see [Section 20.5.6, “Uninstalling a Control,” on page 207](#).
Resolve the out of sync issue before installing the new Solution Pack.
- 6 Reinstall the control with the content you want to keep.
- 7 Implement and test as necessary.

Copying a Map File

If the Solution Pack that you install contains a Map control, you need to copy the associated .csv file to the system where you are importing and installing the Solution Pack. This file is used by the mapping service for event enrichment. Data from this .csv file is used to populate the tag when specified conditions are met for all incoming events.

When you create a Solution Pack using a Map control, the map definition file (.csv) that is used to create the Map is not bundled in the Solution Pack. Therefore, when you install this Solution Pack on any other Sentinel server, you do not get the expected behavior. This is because this Map looks for the required information in the map definition file (.csv) in the var/opt/novell/sentinel/data/map_data folder of the Sentinel Server to populate the tag. If the correct .csv file is not there, the Map control does not work properly.

You must copy the map definition file that you used to create the Map to the var/opt/novell/sentinel/data/map_data folder whenever you install any Solution Pack that has a Map control.

20.5.4 Implementing a Control

The steps on how to implement a control is added when the Solution Pack is created in the Solution Designer. The steps might include instructions for the following types of implementation actions:

- ♦ Scheduling automatic report execution.
- ♦ Enabling auditing on source devices.
- ♦ Copying an attached script for an Execute Command Correlation Action to the appropriate location on the correlation engines.

You only need to follow the instructions that are in the Implementation tab for the control.

To implement a control:

- 1 Open a Solution Pack in Solution Manager.
For more information, see [Section 20.4.1, “Launching the Solution Manager,” on page 200](#).
- 2 Select a control.
- 3 Click the **Implementation** tab in the Documentation panel.
- 4 Follow all of the instructions in the **Implementation** tab.
- 5 Add notes to the **Notes** tab of the Documentation panel as necessary to document progress or necessary changes from the recommended implementation steps.
- 6 When the implementation is complete, select the control and change the status drop-down to **Implemented**.
- 7 An audit event is generated and sent to the Sentinel Control Center.

Because of potential legal and regulatory implications, the status for a control should only be changed after all of the implementation steps have been successfully completed.

20.5.5 Testing a Control

After a control is implemented, the content should be tested to verify that it is working as expected. Testing might require steps such as running a report or generating a failed login. The testing instructions for each control are added when the Solution Pack is created in the Solution Designer.

To test a control:

- 1 Open a Solution Pack in Solution Manager.
For more information, see [Section 20.4.1, “Launching the Solution Manager,” on page 200](#).
- 2 Select a control.
- 3 Click the **Testing** tab in the Documentation panel.
- 4 Follow all of the instructions in the **Testing** tab.
- 5 Add notes to the **Notes** tab of the Documentation panel as necessary to document progress or necessary changes from the recommended testing steps.
- 6 When the testing is complete, select the control and change the status drop-down to **Tested**.
- 7 An audit event is generated and sent to the Sentinel Control Center.

Because of potential legal and regulatory implications, the status for a control should only be changed after all of the testing steps have been successfully completed.

20.5.6 Uninstalling a Control

Controls are often used to meet legal or regulatory requirements. After they are implemented and tested, controls can be uninstalled after careful consideration.

When a control is uninstalled, the status for the control reverts to Not Implemented and child content is deleted from the Sentinel system. There are a few exceptions and special cases:

- ♦ Dependencies are checked to ensure that no content that is still in use is deleted. Some examples of this include a Dynamic List that is used by a Correlation rule created in the target Sentinel system, a report that is used in a control that is still installed, an iTRAC workflow template that is used in a Solution Pack that is still installed, or a folder that still contains other content.
- ♦ Reports (.rpt files) copied to a local system cannot be removed if the uninstall is performed from a Sentinel Control Center on a different machine.
- ♦ JavaScript files associated with Execute Script Correlation actions remain on the Correlation Engines.
- ♦ Maps (.csv files) and the data they contain are not deleted.
- ♦ Roles associated with workflows are not deleted.
- ♦ iTRAC workflow processes that are already in progress continue until completion even if the iTRAC workflow is uninstalled.

To uninstall a control:

- 1 Open a Solution Pack in Solution Manager.

For more information, see [Section 20.4.1, “Launching the Solution Manager,” on page 200](#).

- 2 Right-click the control you want to uninstall and select **Uninstall**.

Alternatively, click the **Uninstall** icon. The Controls To Uninstall window appears.

- 3 Click **Next**.

If the control you are uninstalling includes one or more reports, you are prompted whether to uninstall the reports from the server. Ideally, this information was recorded on the **Notes** tab when the reports were installed.

- 4 Click **Next** to display the Uninstall Content window.

- 5 Click **Uninstall**. The selected contents are uninstalled.

Local reports cannot be uninstalled from a different Sentinel Control Center machine than the one that was used for the installation, or if the files were copied to a new location after installation. If the Solution Manager cannot find the .rpt files in the expected location, a message is logged in the Sentinel Control Center log file.

- 6 Click **Finish**.

20.5.7 Viewing Solution Pack Status

There are several sources of information about the status of a Solution Pack.

- ♦ [“Viewing the Status in the Solution Manager” on page 208](#)
- ♦ [“Generating Status Documentation” on page 208](#)
- ♦ [“Audit Events in the Sentinel Control Center” on page 208](#)

Viewing the Status in the Solution Manager

You can view the status of Solution Pack contents in the Solution Manager. For more information, see [Section 20.4.1, “Launching the Solution Manager,” on page 200](#).

- ♦ **None/Blank:** No status indicator for a control indicates that the associated content has not been installed yet.
- ♦ **Not Implemented:** When none or some of the contents of a control are installed, the control is in the Not Implemented state. If the same content is installed by another control, a control might be Not Implemented even if some of its child content is Installed.
- ♦ **Implemented:** Indicates that a user has completed all of the implementation steps and manually set the control status to Implemented. For more information, see [Section 20.5.4, “Implementing a Control,” on page 206](#).
- ♦ **Tested:** Indicates that a user has completed all of the testing steps and manually set the control status to Tested. For more information, see [Section 20.5.5, “Testing a Control,” on page 206](#).
- ♦ **Out of Sync:** Indicates that a different version of the content in the Solution Pack has been installed in the Sentinel target system by another Solution Pack or a previous version of the same Solution Pack. For more information, see [“Out Of Sync Status” on page 202](#)

Generating Status Documentation

The information about the Solution Pack can be exported in PDF format. The report contains details about every node in the Solution Pack, including the category, control, and content group.

To generate Solution Pack status documentation:

- 1 Open a Solution Pack in the Solution Manager for which you want to generate a status report.
- 2 Click **Create PDF**. The Report Options window displays.
 - ♦ **Show status information:** Select this option to show deployment status for each control (Not Installed, Not Implemented, Implemented, or Tested) and whether it's Out of Sync.
 - ♦ **Include content nodes:** Select this option to include information about the child content for each control in the documentation.
- 3 Select Show status and Show individual content if desired.
- 4 To view the documentation, click **Preview**.
(Conditional) If this is the first time a PDF has been opened from the Sentinel Control Center, you might need to locate Acrobat Reader.
- 5 (Conditional) Click **Browse**, and locate the Acrobat Reader, then click **OK**.
The report is opened in the PDF format.
- 6 To save the PDF, click **Browse**. Navigate the location where you want to save the PDF and specify a filename.
- 7 Click **Save**.

Audit Events in the Sentinel Control Center

All major actions related to Solution Packs and controls are audited by the Sentinel system, with information about which user performed the action. The following events are visible in the Sentinel Control Center and are stored in the Sentinel database:

- ♦ Solution Pack is imported.
- ♦ Control is installed.

- ♦ Control status is changed to Implemented.
- ♦ Control status is changed to Tested.
- ♦ Control status is changed to Not Implemented.
- ♦ Control is uninstalled.
- ♦ Notes are modified for a control.
- ♦ Solution Pack is deleted.

20.5.8 Deleting a Solution Pack

Solution Packs are often used to meet legal or regulatory requirements. After they are implemented and tested, Solution Packs can be deleted after careful consideration. All deletions are audited by the Sentinel system and sent to both the Sentinel Control Center and the Sentinel database.

You cannot delete a Solution Pack without uninstalling the controls first. For more information, see [Section 20.5.6, “Uninstalling a Control,” on page 207](#). If you do not uninstall the controls and try to delete a solution pack, you are notified that content is still deployed.

- 1 Open the Solution Packs Manager.
- 2 Select the Solution Pack that you want to remove, then click **Remove**.
You are prompted to delete the Solution Pack.
- 3 Click **Yes** to delete.

20.6 Installing an Edited Solution Pack

When a Solution Pack is modified and saved through the **Save** or **Save As** options in the Solution Designer, it is considered to be a new version of the original Solution Pack. When the new version is imported, it replaces any older versions of the original Solution Pack. There is no immediate impact on any installed content in the target Sentinel system.

After the new Solution Pack is installed, its behavior varies, depending on the status of the original Solution Pack's content.

- ♦ If the content from the original Solution Pack was not installed yet, the content is simply replaced. When a user installs content, the new content is installed to the target Sentinel system.
- ♦ If the content from the original Solution Pack was installed (Not Implemented), Implemented, or Tested, the original content is compared to the new content.
- ♦ If the content version is the same, the original content is still valid and no action is necessary.
- ♦ If the content version is different, the content status is set to Out of Sync. You must decide how to resolve the synchronization issue. For more information, see [“Out Of Sync Status” on page 202](#).
- ♦ If the content did not exist in the original Solution Pack, it is displayed in Solution Manager as not installed. You can install, implement, and test the new content.
- ♦ If the content existed in the original Solution Pack but has been deleted from the modified Solution Pack, it does not appear in the Solution Manager.

NOTE: The Solution Manager only handles differences in the contents of Solution Packs. It does not recognize manual content changes that are performed after content is installed.

20.7 Solution Designer

You can use the Solution Designer to package and export different contents, such as a Correlation rule with associated actions and Dynamic Lists. These items can be selected and packaged with their configuration to a .zip file. You can then view or select the content of the file in the Solution Manager. For more information, see [Chapter 21, “Creating Solution Packs,” on page 211](#).

21 Creating Solution Packs

You can use the Solution Designer to package and export different contents, such as a Correlation rule with associated actions and dynamic lists. The content can be selected and packaged with its configuration in a ZIP file. You can then view or select the content of the ZIP file by using the Solution Manager. For more information on the Solution Manager, see [Chapter 20, “Using Solution Packs,” on page 197](#).

To use the Solution Designer, you must have the correct permission. All roles contain the permission for the Solution Designer except for the PCI Compliance Audit role and the Search Proxy User role. For more information, see [Chapter 4, “Configuring Roles and Users,” on page 39](#).

- ♦ [Section 21.1, “Accessing the Solution Designer,” on page 211](#)
- ♦ [Section 21.2, “Creating a Solution Pack,” on page 211](#)
- ♦ [Section 21.3, “Adding Content to a Solution Pack,” on page 212](#)
- ♦ [Section 21.4, “Initializing Dynamic Lists Through Solution Pack,” on page 214](#)
- ♦ [Section 21.5, “Documenting a Solution Pack,” on page 215](#)
- ♦ [Section 21.6, “Synchronizing Content,” on page 215](#)
- ♦ [Section 21.7, “Handling Inter-control Dependency,” on page 216](#)
- ♦ [Section 21.8, “Managing a Solution Pack,” on page 216](#)

21.1 Accessing the Solution Designer

- 1 Log in to the Sentinel Web interface as a user with permissions to access the Solution Designer.
- 2 In the toolbar, click **Applications**.
- 3 Click **Launch Designer**.
- 4 Click **Yes** to accept the security certificate.

NOTE: NetIQ Corporation leverages our partner company Novell for certificate signing services under a trust model. Although the certificate says Novell, it is trusted through NetIQ Corporation.

- 5 Specify a user name and password of a user with permission to access the Solution Designer.
- 6 Click **Login**.
- 7 Click **Accept** or **Accept Permanently** to accept the security certificate.

21.2 Creating a Solution Pack

You can use the Solution Designer to create a Solution Pack with existing content objects (for example, Actions, Event Actions, Filters, Searches, Correlation Rules, Dynamic Lists, or iTRAC workflow templates) from Sentinel. The Solution Designer analyzes the dependencies for a content

object and include all necessary components in the Solution Pack. For example, a Correlation Rule deployment includes a Correlation Rule definition, one or more actions, and the ability to create an incident using a workflow. The Solution Designer includes the Correlation Rule, the associated correlation actions, the iTRAC template, and the roles associated with the iTRAC template in the Solution Pack.

IMPORTANT: To add a content object to a Solution Pack, it must already exist in Sentinel. Content objects cannot be created in the Solution Designer.

To create a new Solution Pack:

- 1 Access the Solution Designer.
For more information, see [Section 21.1, “Accessing the Solution Designer,”](#) on page 211.
- 2 Click **File > New**.
An empty Solution Pack is displayed in the Solution Pack panel.
- 3 Add Categories, Controls, Content Groups, and content placeholders.
For detailed instructions, see [Section 21.3, “Adding Content to a Solution Pack,”](#) on page 212.
- 4 Add file attachments to the hierarchy nodes as desired.
For detailed instructions, see [“File Attachments”](#) on page 214.
- 5 Click **File > Save**.
- 6 Browse to and select a location to save the Solution Pack, then specify a name for the Solution Pack.
- 7 Click **Save** to save the Solution Pack.
The Solution Pack is saved in a .zip format.

Although you can save a Solution Pack with empty placeholders, you cannot install controls in the Solution Manager unless all placeholders have been filled with content.

21.3 Adding Content to a Solution Pack

A vital part of creating a Solution Pack is adding content to the controls. Each control can have one or more types of content associated with it.

- ♦ [Section 21.3.1, “Sentinel Content,”](#) on page 212
- ♦ [Section 21.3.2, “Using Placeholders,”](#) on page 214
- ♦ [Section 21.3.3, “File Attachments,”](#) on page 214

21.3.1 Sentinel Content

The same general procedure is used to add all types of Sentinel content to a Solution Pack. The Sentinel content palette includes the following:

- ♦ Actions
- ♦ Correlation Rule deployments, including their deployment status (enabled or disabled) and associated Correlation rules, Correlation Actions, and Dynamic Lists
- ♦ Event Actions
- ♦ Reports
- ♦ Filters

- ♦ Searches
- ♦ iTRAC workflows, including associated roles
- ♦ Event enrichment, including map definitions and event metatag configuration
- ♦ Other associated files added when the Solution Pack is created, such as documentation, example report PDFs, or sample map files.

Adding Sentinel Content to a Control

To add Sentinel content to a control:

- 1 Access the Solutions Designer.
For more information, see [Section 1.1, “Accessing the Sentinel Web Console,” on page 17](#).
- 2 Open or create a Solution Pack.
- 3 Click the appropriate panel to display the available content:

- ♦ Actions
- ♦ Correlation
- ♦ Event Actions
- ♦ Event Enrichment
- ♦ Filters
- ♦ iTRAC
- ♦ Jasper Reports
- ♦ Searches

- 4 Drag the item and drop it into the control.

If you try to drag and drop pre-existing content in the Solution Designer, the existing content is highlighted. After you drop the content, a message prompt indicates that similar content exists.

Setting Content Properties

You can set properties to a content to indicate it is designed for specific Sentinel platforms. Content that is designed in newer versions of Sentinel might not be supported in older versions because of changes in the Sentinel schema. If you try to install a Control on an unsupported Sentinel platform, the installation does not proceed and shows an “Out of date” error.

To set the properties:

- 1 Right-click a content, then select **Properties**.
- 2 (Conditional) For Correlation rules, select **Automatically deploy during installation** to deploy Correlation rules automatically during the solution pack installation.
- 3 Select **Minimum Required Versions**, and then specify the Sentinel versions.
- 4 Click **Apply**.

21.3.2 Using Placeholders

If the user is not ready to associate content with a control, an empty placeholder can be used instead.

- 1 Click the **Correlation**, **Event Actions**, **Actions**, **Filters**, **Event Enrichment**, **iTRAC**, or **Jasper Report** button in the Content Palette to open the panel for the type of placeholder you want to add.
- 2 Drag and drop the placeholder to the appropriate control in the Solution Pack panel.
- 3 Rename the placeholder, if desired.

To replace a placeholder with content:

- 1 Click the **Correlation**, **Event Actions**, **Filters**, **Event Enrichment**, **iTRAC**, or **Jasper Report** button in the Content Palette to open the panel for the type of placeholder you want to add.
- 2 Drag and drop the appropriate Content Group from the Content Palette to the placeholder in the Solution Pack panel or select the appropriate Content Group, then click **Add Selected Content**.

You can set properties for placeholders to indicate whether a placeholder is designed for specific Sentinel platforms. Placeholders that are designed in newer versions of Sentinel might not be supported in older versions because of changes in the Sentinel schema. If you try to install a placeholder on an unsupported Sentinel platform, the install does not proceed and shows an “Out of date” error.

To set the properties:

- 1 Right-click the placeholder, then select **Properties**.
- 2 Select **Minimum Required Versions**, then specify the Sentinel versions.
- 3 Click **Apply**.

21.3.3 File Attachments

You can attach a file or files to any node in the hierarchy. The content in the attachment is included in the Solution Pack. These files can include anything useful for a user who must deploy the Solution Pack, such as a PDF view of a report, sample map data for event enrichment, or a script for an Execute Command Correlation Action. These files can be added, deleted, viewed, renamed, or saved to the local machine.

21.4 Initializing Dynamic Lists Through Solution Pack

The Correlation rules in solution packs require some data in the dynamic lists for it to work properly. The solution pack framework includes the ability to automatically populate the dynamic lists with data when you install a solution pack.

To populate a dynamic list when you install a solution pack:

- 1 Create a text file with the values that you want to add to the dynamic list. Add each different value on a separate line.
- 2 In the Solution Designer, expand the Correlation content, and then select the dynamic list.
- 3 Click **Add a new attachment** in the Attachment panel, and attach the file that you created in [Step 1](#).

All the values in the Dynamic list are persistent. For more information, see “[Configuring Dynamic Lists](#)” in the *NetIQ Sentinel User Guide*.

21.5 Documenting a Solution Pack

The Solution Designer provides three different categories of documentation to help you create the documentation for the Solution Pack you are creating.

- ♦ [Section 21.5.1, “Description,” on page 215](#)
- ♦ [Section 21.5.2, “Implementation Steps,” on page 215](#)
- ♦ [Section 21.5.3, “Testing Steps,” on page 215](#)

21.5.1 Description

Allows you to provide a detailed description about the Solution Pack for your users.

21.5.2 Implementation Steps

Lets you add the steps required to implement the content in the target Sentinel system to the **Implementation** tab of the Documentation panel. The steps might include instructions for the following types of implementation actions:

- ♦ Populating a .csv file that is used by the mapping service for event enrichment.
- ♦ Scheduling automatic report execution
- ♦ Enabling auditing on source devices.
- ♦ Copying an attached script for an Execute Command Correlation Action to the appropriate location on the correlation engines.

After the content implementation, the content should be tested to verify that it is working as expected.

21.5.3 Testing Steps

Lets you add the steps required to test the content in the target Sentinel system to the **Testing** tab of the Documentation panel. The steps can include instructions for the following types of testing activities:

- ♦ Running a report and verifying that data is returned.
- ♦ Generating a failed login in a critical server and verifying that a correlated event is created and assigned to an iTRAC workflow.

21.6 Synchronizing Content

If you modify the content in the source system, the content in the source system and the content in the original Solution Pack can be out of synchronization. To synchronize the content, do one of the following:

- ♦ For content with no dependencies, drag and drop the content from the Content Palette onto the control.

The modified content is immediately updated. For example, a report has no dependencies.

- ♦ For content with dependencies, the dependencies are checked and updates are made when you click the **Synchronize All Content** icon or when you save the Solution Pack. However, you need to ensure that the system that you are connected to has the latest content.
- ♦ To synchronize specific content based on any content group, right-click the content or a content group and click **Synchronize this content**. Using this menu ensures that only the content and the contents within that group are synchronized.

When an action uses the Send Email action, this action always appears as Out of Synchronization. This is expected and does not cause an error.

21.7 Handling Inter-control Dependency

You can specify any control as a required control in the Solution Designer. This ensures that the control marked as required is also installed when a user chooses to install any other control first. For example, you can mark the global setup control as a required control, which is then installed when the user installs any other control from a solution pack.

You can also specify if you want to overwrite an existing control during installation. For example, if you include a newer version of a White Label Template and want to ensure that this newer version is automatically installed with a new install of solution pack, you can enable the overwrite properties.

To mark a control as required:

- 1 In the Solution Designer, select the control that you want to mark as required.
- 2 Right-click the control and select **Properties**.
- 3 (Conditional) Select **Required** if you want to ensure that this control is also installed while installing any specific control from a solution pack.
- 4 (Conditional) Select **Enable Overwrite** if you want to automatically install this control with a new install of solution pack.
- 5 Click **Apply**.

21.8 Managing a Solution Pack

All content in a Solution Pack is hierarchically organized into categories, controls, and content groups.

- ♦ [Section 21.8.1, “Adding a Node to a Control,” on page 216](#)
- ♦ [Section 21.8.2, “Moving Nodes,” on page 217](#)

21.8.1 Adding a Node to a Control

- 1 Select a node in the Solution Pack panel.
- 2 Right-click the node, then select **Create**
or
Click **Create** in the Solution Pack panel heading.

21.8.2 Moving Nodes

Category, control, and content group nodes can be created in any order and then reordered or moved to a different parent in the hierarchy.

To move a node to another branch in the hierarchy, drag and drop a node to its new parent node. A control can be moved to a new category. A content group can be moved to a new control.

To reorder a node, drag and drop it on top of the node it should appear after in the Solution Pack.

VII Managing Your Sentinel Environment

This section provides information about managing your Sentinel server.

- ♦ [Chapter 22, “Managing Active Searches and Reports,” on page 221](#)
- ♦ [Chapter 23, “Viewing Events Per Second Statistics,” on page 223](#)
- ♦ [Chapter 24, “Configuring Sentinel for High Availability,” on page 225](#)
- ♦ [Chapter 25, “Backing Up and Restoring Data,” on page 227](#)
- ♦ [Chapter 26, “Customizing Sentinel Settings,” on page 233](#)

22 Managing Active Searches and Reports

Sentinel provides an option to monitor and manage active searches and reports on the Sentinel server for the purpose of resource management. You can view all the searches and reports currently active on the Sentinel server, determine which long-running searches or reports are no longer needed, and stop them as necessary.

Sentinel helps you monitor search and report activities and determine whether a search or a report is not retrieving events as expected or whether a search or a report is retrieving more than the expected events, which might indicate that the search or the report needs to be tuned. It also helps you determine if too many searches or reports are running, and helps identify long-running searches and reports that might slow down the system. Searches and reports that consume a lot of memory are a potential liability to a healthy system and should be carefully reviewed to ensure that the search query is specified properly. You can also stop the searches and reports that are no longer needed and thereby free up system resources.

To manage active searches:

- 1 Log in to the Sentinel Web interface as an administrator.

`https://<IP_Address/DNS_Sentinel_server:8443>`

The `IP_Address/DNS_Sentinel_server` is the IP address or DNS name of the Sentinel server and `8443` is the default port for the Sentinel server.

- 2 Click **Storage** in the toolbar, then click **Search Jobs**.

The Search Jobs lists all the active event search jobs running in the system, including searches that are initiated when users perform activities, such as:

- ♦ Run a search in the Search interface.
- ♦ View events that fire a correlation rule.
- ♦ View events processed when testing a correlation rule.
- ♦ Generate a report or drill down into report results.
- ♦ Select filters to view events that match the filter criteria.
- ♦ Select tags to view the events that are tagged with the specified criteria.
- ♦ View events from the dashboard, anomaly, continuation breakdown, and so forth.
- ♦ Analyze, investigate, or view triggered events in Active Views.

The Search Jobs page refreshes every 30 seconds.

You can view the following search details. Mouse over each field for information on what the field indicates:

- ♦ **Duration:** The time spent to search events in the event store.
- ♦ **Status:** Whether a search job is pending, running, finished, finished with errors, or canceled.
- ♦ **Owner:** The user who initiated the search. For search jobs initiated by the system, the owner is indicated as "System."

- ♦ **Type:** Indicates the following:
 - ♦ **System:** Search jobs that are run for maintenance purposes. For example, to clean up invalid references to events from the database.
 - ♦ **User:** Search jobs started by users either through the Search interface or through the REST API.
 - ♦ **Reports:** Search jobs started by users, but used for getting event results for reports.
 - ♦ **Data sync:** Search jobs started to support the Data Synchronization feature.
 - ♦ **Distributed:** Search jobs initiated by a remote server (distributed search.)
 - ♦ **Start:** The time the search started searching for events.
 - ♦ **Accessed:** The time elapsed since the search was initiated.
 - ♦ **More:** Provides detailed information such as the IP address of the machine that initiated the search, events processed, search criteria, and so forth.
- 3** (Conditional) To stop any active search jobs, select the search jobs you want to stop, then click **Stop selected**.

To manage active reports:

- 1** Log in to the Sentinel Web console as an administrator.

`https://<IP_Address/DNS_Sentinel_server:8443>`

The *IP_Address/DNS_Sentinel_server* is the IP address or DNS name of the Sentinel server and *8443* is the default port for the Sentinel server.

- 2** Click **Storage** in the toolbar, then click **Report Jobs**.
- 3** (Conditional) To stop any report jobs, select the report jobs you want to stop, then click **Stop selected**.
- 4** (Conditional) To delete any report jobs, select the report jobs you want to delete, then click **Delete selected**


23 Viewing Events Per Second Statistics

- ♦ [Section 23.1, “Viewing a Graphical Representation of the Events Per Second Value,” on page 223](#)
- ♦ [Section 23.2, “Viewing the Events Per Second Value of Event Source Servers,” on page 224](#)

23.1 Viewing a Graphical Representation of the Events Per Second Value

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection** in the toolbar.
- 3 In the **Overview** section, view the events per second (eps) value of the incoming events in the last one minute.

The graph shows the 90-day statistics of all the events coming to the Sentinel server. The graph also includes an EPS indicator that enables you to determine whether the current EPS rate is exceeding the licensed EPS rate or is close to the licensed EPS rate.



23.2 Viewing the Events Per Second Value of Event Source Servers

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Collection** in the toolbar, then click the **Event Sources** tab.

The **EPS** column of the **Event Source Servers** section specifies the events per second value received from all the Event Source Servers.

24

Configuring Sentinel for High Availability

You can install Sentinel in an Active-Passive High Availability mode, which allows Sentinel to fail over to a redundant cluster node in case of hardware or software failure. NetIQ Consulting and NetIQ partners can help you implement Sentinel high availability and disaster recovery. For more information about configuring Sentinel for High Availability, see “[Configuring Sentinel for High Availability](#)” in the *NetIQ Sentinel Installation and Configuration Guide*.

25 Backing Up and Restoring Data

The Sentinel backup and restore utility is a script that performs a backup of the system data and also lets you restore the data at any given point in time. However, this utility cannot be used to back up a Collector Manager or the operating system components and configuration.

NOTE: You can restore the data only on the same version of Sentinel in which the data was backed up because there might be changes between Sentinel versions, which might make the data incompatible.

You can back up the following data:

- ♦ **Configuration data:** Data stored in the `config`, `data`, `3rdparty/postgresql`, and `3rdparty/jetty` directories, and the data in the Sentinel database. This data includes configuration files, property files, and keystore files. The Sentinel database contains various configuration information related to users, plug-ins, Collectors, Connectors, and filters.

NOTE: The configuration data can be critical and you should always include the configuration data in the backup.

- ♦ **Event data:** Dynamic event data and raw event data stored in the `data/eventdata` and `/var/opt/novell/sentinel/data/rawdata` directories. The event data also includes event associations stored in the `/var/opt/novell/sentinel/data/eventdata/exported_associations` directory. The event associations data includes correlated event association data and the incident event association data.
- ♦ **Secondary storage data:** The closed event data files that have been moved to the secondary storage.
- ♦ **Runtime data:** Dynamic file-based queues used by plug-ins, Sentinel Link, and other Sentinel components. This includes the data in the `data/plugindata` and the `/var/opt/novell/sentinel/data/sentinel_link.queues` directories.
- ♦ **Security Intelligence data:** The Security Intelligence data stored in the MongoDB database.
- ♦ **Sentinel logs:** Log files generated by Sentinel and stored in the `/var/opt/novell/sentinel/log` directory.
- ♦ [Section 25.1, “Parameters for the Backup and Restore Utility Script,” on page 227](#)
- ♦ [Section 25.2, “Running the Backup and Restore Utility Script,” on page 229](#)

25.1 Parameters for the Backup and Restore Utility Script

The following lists the various command line parameters that you can use with the `backup_util.sh` script:

Table 25-1 Backup and Restore Script Parameters

Parameters	Description
-m backup	Backs up of the specified data.
-m restore	<p>Restores the specified data. The restore mode of the script is interactive and allows you to specify the data to be restored from the backup file.</p> <p>The restore parameter can be used in the following scenarios:</p> <ul style="list-style-type: none">♦ System Failure: In the event of a system failure, you must first reinstall Sentinel and then use the <code>backup_util.sh</code> script with the restore parameter to restore the most recent data that you backed up.♦ Data Loss: In the event of data loss, use the <code>backup_util.sh</code> script with the restore parameter to restore the most recent data that you had backed up. <p>You must restart the Sentinel server after you restore any data because the script might make several modifications to the database.</p>
-m info	Displays information for the specified backup file.
-m simple_event_backup	Backs up events located in a specified directory.
-m simple_event_restore	Restores events into a specified directory.
-c	Backs up the configuration data.
-e	Backs up the event data. All event partitions are backed up except the current online partition. If the backup is being performed with the Sentinel server shut down, the current online partition is also included in the backup.
-dN	<p>Backs up the event data for the specified number of days. The -dN option backs up the primary storage event data stored for the last N days. Based on the current data retention policy settings, many days of events might be stored on the system. Backing up all of the event data might not always be necessary and might not be desirable. This option allows you to specify how many days to include when backing up the event data. For example, -d7 includes only the event data from the last week in the backup. -d0 just includes the data for the current day. -d1 includes the data from the current day and previous day. -d2 includes the data from the current day and two days ago.</p> <p>Online backups (that is, backups performed while the system is running) only back up the closed event partitions, which means partitions one day old or older. For online backups, a value of -d1 is the appropriate specification for the number of days.</p>
-u	<p>Specifies the user name to use when backing up the event associations data. If the user name is not specified, "admin" is used as the default value.</p> <p>This parameter is required only when backing up the event associations data.</p>
-p	<p>Specifies the user password when backing up the event associations data.</p> <p>This parameter is required only when backing up the event associations data.</p>
-x	<p>Specifies a file name that contains the user password when backing up the event associations. This is an alternative to the -p option.</p> <p>This parameter is required only when backing up the event associations data.</p>
-f	Enables you to specify the location and name of the backup file.

Parameters	Description
-l	Includes the log files in the backup. By default, the log files are not backed up unless you specify this option.
-r	Includes the runtime data in the backup. Runtime data can only be backed up if the Sentinel server is shut down, because the data is dynamic. This means that this parameter can only be used in combination with the -s option (described below). If -s is not specified, this parameter is ignored.
-b	<p>Backs up the NetFlow data collections and the baseline Security Intelligence, and not the entire MongoDB database. The following baseline data is backed up:</p> <ul style="list-style-type: none"> ◆ configs ◆ anomalydefs ◆ baselines ◆ baselines.ID.URN ◆ paths.UUID.URN ◆ anomalydeployment
-i	Backs up the entire MongoDB database including the Security Intelligence database collections and NetFlow data collections in the backup.
-s	Shuts down the Sentinel server before performing the backup. Shutting down the server is necessary to back up certain dynamic data such as the Runtime data and the current primary storage partitions. By default, the server is not shut down before performing the backup. If this option is used, the server restarts automatically after the backup is complete.
-w	Backs up the raw event data.
-z	Only available with the simple_event_backup and simple_event_restore options. Specifies the location of the event data directory, such as where the event data is collected during a simple_event_backup and where the event data is placed during a simple_event_restore.

25.2 Running the Backup and Restore Utility Script

- 1 Open a console, and navigate to the `/opt/novell/sentinel/bin` directory as the novell user.
- 2 Enter `backup_util.sh`, along with the necessary parameters for the data that you want to back up or restore.

NOTE: You can restore the data only on the same version of Sentinel in which the data was backed up because there might be changes between Sentinel versions, which might make the data incompatible.

For more information on the different parameters, see [Table 25-1](#). The following lists examples of how to specify the parameters:

Syntax	Action
<pre> backup_util.sh -m backup -c -e -i -l -r -w -s -u admin -x <mypassword.txt> -f / var/opt/novell/ sentinel/data/ <my_full_backup>.tar .gz </pre>	Shuts down the Sentinel server and performs a full system backup.
<pre> backup_util.sh -m backup -c -e -i -l -w -u admin -x <mypassword.txt> -f / var/opt/novell/ sentinel/data/ <my_weekly_backup>.t ar.gz </pre>	Performs an online backup without shutting down the server. This backup includes everything except online event data and dynamic runtime data.
<pre> backup_util.sh -m backup -b -c -e -d7 - u admin -x <mypassword.txt> -f / var/opt/novell/ sentinel/data/ <my_weekly_backup>.t ar.gz </pre>	Performs an online backup with event data just from the last week. This backup includes configuration data, the baseline Security Intelligence collections, and the event data for the last 7 days. Event data older than 7 days is not backed up because that data can be extracted selectively, if necessary, from an older backup.
<pre> backup_util.sh -m backup -c -f /var/ opt/novell/sentinel/ data/ config_backup.tar.gz </pre>	Performs a local backup of the configuration data. This is a minimal backup of the system without any event data.
<pre> backup_util.sh -m backup -e -f /var/ opt/novell/sentinel/ data/ events_backup.tar.gz </pre>	Performs a local backup of the event data. This is a minimal backup of the primary storage event data.
<pre> backup_util.sh -m backup -e -d5 -f / var/opt/novell/ sentinel/data/ events_5days_backup. tar.gz </pre>	Performs a local backup of the event data from the last 5 days. This is a minimal backup of the primary storage event data from the last five days.
<pre> backup_util.sh -m info -f /var/opt/ novell/sentinel/ data/ config_backup.tar.gz </pre>	Displays the backup information for the specified backup file.
<pre> backup_util.sh -m simple_event_backup -e -z /opt/archives/ archive_dir -f /opt/ archives/ archive_backup.tar.g z </pre>	<p>Performs a backup of event data on the machine where the secondary storage directory is located.</p> <p>If the /opt/archives/archive_dir is not located in the server, you might need to copy the backup_util.sh script to the machine where the secondary storage is located and then run the simple_event_backup command from that machine.</p> <p>Alternatively, you can also use any 3rd party backup tool to backup the event directories on secondary storage.</p>

Syntax	Action
<pre> backup_util.sh -m restore -f /var/opt/ novell/sentinel/ data/ config_backup.tar.gz </pre>	Restores the data from the specified filename.
<pre> backup_util.sh -m simple_event_restore -z /opt/archives/ archivedir -f /opt/ archives/ archive_backup.tar.g z </pre>	Performs a restore of backed-up secondary storage data.

- 3 (Conditional) If you have restored any data, restart the server because the script might make several modifications to the database.
- 4 Use the Data Restoration feature to restore the extracted partitions. For more information, see [Section 6.8.6, “Restoring Data,” on page 75](#).

26 Customizing Sentinel Settings

This section provides information about configuring the default Sentinel settings to optimize the performance.

- ♦ [Section 26.1, “Optimizing the Operating System,” on page 233](#)
- ♦ [Section 26.2, “Configuring the Resources for Event Partition Compression,” on page 234](#)
- ♦ [Section 26.3, “Setting the Grace Period to Close Event Data Partitions,” on page 234](#)
- ♦ [Section 26.4, “Compressing the Storage Index on Primary Partition,” on page 234](#)
- ♦ [Section 26.5, “Configuring Memory for the Sentinel Server,” on page 235](#)
- ♦ [Section 26.6, “Setting the Raw Data Limit,” on page 235](#)

26.1 Optimizing the Operating System

Sentinel sets a default limit of 65536 open files for the Sentinel processes. This is sufficient for most installations. However, there might be some scenarios that may cause the Sentinel server to need more open files, such as running a large number of searches or reports concurrently, or running searches and reports that span unusually large time ranges. In such cases you might consider increasing the default open files limit to allow more than 65536 open files.

NOTE: For optimal performance, NetIQ Corporation recommends you not to decrease the default limit.

Perform the following steps to change the open file limits:

- 1 Log in to the Sentinel server as the root user.
- 2 Open the `/etc/security/limits.conf` file.
- 3 Change the value in the following lines:

```
novell soft nfile 65536
novell hard nfile 65536
```

NOTE: Setting the soft limits is optional, however, setting the hard limits is mandatory.

- 4 Save the changes.
- 5 Restart the Sentinel server.

26.2 Configuring the Resources for Event Partition Compression

You can specify the number of processors for mksquashfs to use when compressing the index on the event data. This capability enables the mksquashfs utility to make more use of additional CPUs that may be available on some systems.

- 1 Log in to the system as the root user.
- 2 Open the `/etc/opt/novell/sentinel/config/configuration.properties` file.
- 3 Add or edit the `mksquashfs.numprocessors` property and specify the desired value. This value specifies the number of processors you want to allow mksquashfs to use.

```
mksquashfs.numprocessors=4
```

- 4 Save the changes.
- 5 Restart the Sentinel server.

26.3 Setting the Grace Period to Close Event Data Partitions

The event data partitions are closed after one day, and no more events are written to them. Even though the duration for event data partitions is one day, a grace period of 10 minutes is given to accommodate events arriving late.

You can change the default value as necessary. To customize the grace time period to close a partition, perform the following steps:

- 1 Log in to the system as the root user.
- 2 Open the `/etc/opt/novell/sentinel/config/configuration.properties` file.
- 3 Edit the `sentinel.events.online.opengraceminutes` property to specify the desired value.
By default, Sentinel does not set any value to this property. If no value is set, Sentinel considers 10 minutes as the grace time period to close a partition.
- 4 Restart the Sentinel server.

26.4 Compressing the Storage Index on Primary Partition

You can compress the partitions in primary storage as soon as they are closed to save disk space. This requires additional I/O to compress and store on the primary partition, which means that the supported EPS rate will be significantly lower. Also, searches on these partitions will be slower. Therefore, this option is only suitable for lower EPS rates and if you want to get the most out of primary storage space.

To compress the storage index on the primary partition:

- 1 Log in to the system as the root user.
- 2 Open the `/etc/opt/novell/sentinel/config/configuration.properties` file.
- 3 Set the `localstorage.indexes.compress` property to `true`.
- 4 Restart the Sentinel server.

26.5 Configuring Memory for the Sentinel Server

By default, the Sentinel server consumes 75% of the computer memory. The `setmemory.sh` script in Sentinel dynamically calculates and allocates the default memory settings for the Sentinel server based on the amount of physical memory. The `setmemory.sh` script is located in the `/etc/opt/novell/sentinel/bin` directory.

To make changes to the default memory settings, you must create a `setmemory.properties` file. The default location of this file is `/etc/opt/novell/sentinel/config/setmemory.properties`.

You can set the following configuration parameters in the `setmemory.properties` file:

- ♦ **JAVA_MEM_SERVER:** The maximum heap memory (Xmx in MB) allocated to the Java process. The maximum setting for this parameter is 16384 (or 16 GB). Removing or increasing this limit can have a detrimental effect on system performance, including a reduced event rate (EPS).
- ♦ **JAVA_MEM_REPORT_PROCESS:** The maximum memory allocated to the report processes. The maximum setting for this parameter is 16384 (or 16 GB). Removing or increasing this limit can have a detrimental effect on system performance, including a reduced event rate (EPS).
- ♦ **JAVA_MEM_PERMGEN:** The maximum permanent generation memory (in MB) allocated to the process. By default, this value is 10% of the remaining allocated memory for the Sentinel server.
- ♦ **JAVA_MEM_BROKER:** The maximum amount of memory allocated for the message bus broker. This affects how many connections the message bus broker can accept. By default, this value is 10% of the remaining allocated memory for the Sentinel server.
- ♦ **BROKER_MAX_CON:** The maximum number of connections the message bus broker can accept. By default, this value is 10% of the remaining allocated memory for the Sentinel server.
- ♦ **CORRELATION_INPUT_BUFFER_MAX_SIZE:** The memory allocated to hold the Correlation events. By default, this value is 10% of the remaining allocated memory for the Sentinel server. On the remote Correlation Engine, this value is 50% of the allocated memory for the Sentinel server.

When the server starts, these memory settings in the `setmemory.properties` file override the default settings.

26.6 Setting the Raw Data Limit

Sentinel now sets a limit of 100,000 when copying raw data to prevent raw data from unnecessarily consuming the disk space and causing the system to become unstable. When the limit reaches 100,000, Sentinel automatically deselects the **Copy Raw Data to a file** option in the Connector configuration window and stops copying the raw data. If you want to collect the raw data again, edit the Connector and select the **Copy Raw Data to a file** option.

You can also change the default value as necessary.

- 1 Log in to the system as novell user.
- 2 Open the `/etc/opt/novell/sentinel/config/server.conf` file.
- 3 Edit the `wrapper.java.additional.60=-DMAX_DUMP_SIZE=10000` property to the desired value.
- 4 Restart the Sentinel server.

VIII Appendix

This appendix provides information about managing some of the Sentinel activities by using the command line.

- ♦ [Appendix A, “Command Line Utilities,” on page 239](#)

A Command Line Utilities

The command line utilities included with Sentinel are useful for managing and configuring many lower level functions of the system.

- ♦ [Section A.1, “Managing the Sentinel Services,” on page 239](#)
- ♦ [Section A.2, “Sentinel Scripts,” on page 240](#)
- ♦ [Section A.3, “Running the Report Development Utility,” on page 241](#)
- ♦ [Section A.4, “Getting the .jar Version Information,” on page 242](#)
- ♦ [Section A.5, “Changing the Hostname of a Sentinel Server,” on page 242](#)
- ♦ [Section A.6, “Importing or Exporting Event Association Data,” on page 242](#)
- ♦ [Section A.7, “Managing the Internal Database,” on page 244](#)
- ♦ [Section A.8, “Cleaning Up the Internal Database,” on page 245](#)
- ♦ [Section A.9, “Managing the Sentinel Server,” on page 246](#)

A.1 Managing the Sentinel Services

There is a command line utility included with Sentinel is useful for managing and configuring many lower level functions of the system. The utility is located in `/usr/sbin/rcsentinel`.

The utility has the following options to manage the Sentinel services:

rcsentinel start: Starts the Sentinel service.

rcsentinel stop: Stops the Sentinel service.

rcsentinel status: Displays the status of the Sentinel service.

rcsentinel restart: Restarts the Sentinel service.

rcsentinel try-restart: Restarts the Sentinel service if the Sentinel service is running.

rcsentinel force-reload: Forces the Sentinel service to reload the Sentinel configuration.

rcsentinel startdb: Starts the PostgreSQL database.

rcsentinel stopdb: Stops the PostgreSQL database.

rcsentinel force_stopdb: Forces the PostgreSQL database to stop.

rcsentinel startSldb: Starts the Security Intelligence database.

rcsentinel stopSldb: Stops the Security Intelligence database.

rcsentinel version: Displays the version of the Sentinel service.

rcsentinel -p, --priority=<integer>: Specifies the process priority.

rcsentinel -h, --help: Displays all options for the rcsentinel utility.

rcsentinel -l, --log-file=FILE: Sends log messages to a file.

rcsentinel -q, --quiet: Displays fewer messages.

rcsentinel -v, --verbose: Displays more messages.

A.2 Sentinel Scripts

Sentinel provides operational scripts that are appropriate for use during normal operations. These scripts are located in the following directories:

- ♦ /opt/novell/sentinel/bin
- ♦ /opt/novell/sentinel/setup

For most scripts that require arguments, running the scripts without arguments provides details about how to use the arguments.

- ♦ [Table A-1, “Operational Scripts in /opt/novell/sentinel/bin,” on page 240](#)
- ♦ [Table A-2, “Operational Scripts in /opt/novell/sentinel/setup,” on page 241](#)

Table A-1 Operational Scripts in /opt/novell/sentinel/bin

Script File	Description
backup_util.sh	Use this script to back up and restore Sentinel event data and configuration data. For more information, see Chapter 25, “Backing Up and Restoring Data,” on page 227 .
db.sh	Allows you to manage the PostgreSQL database without Sentinel running. For more information, see Section A.7, “Managing the Internal Database,” on page 244 .
clean_db.sh	Allows you to clean up data from the PostgreSQL database without Sentinel running. For more information, see Section A.8, “Cleaning Up the Internal Database,” on page 245 .
softwarekey.sh	Use this script to add and view a license key through the command line.
event_assoc_data.sh	Use this script to import or export the event association data. For more information, see Section A.6, “Importing or Exporting Event Association Data,” on page 242 .
server.sh	Use this script to manually manage the Sentinel server with this script. For more information, see Section A.9, “Managing the Sentinel Server,” on page 246 .
report_dev_setup.sh	Use this utility to set up the report development environment. For more information, see Section A.3, “Running the Report Development Utility,” on page 241 .
updateServerLocale.sh	<p>This utility provides an option to change the language of Sentinel server process. The Sentinel server messages that are displayed on the user interface appear in the language selected by this script.</p> <p>If the appliance language is changed through WebYast, you can use this script to change the language of the Sentinel process in the server.</p>

Script File	Description
versionreader.sh	Displays the version of the .jar files for Sentinel. For more information, see Section A.4, “Getting the .jar Version Information,” on page 242.

Table A-2 Operational Scripts in /opt/novell/sentinel/setup

Script File	Description
configure.sh	This utility runs on a remote Correlation Engine or Collector Manager if your change the hostname of the Sentinel server. For more information, see Section A.5, “Changing the Hostname of a Sentinel Server,” on page 242.
ldap_auth_config.sh	This utility helps you configure Sentinel to receive LDAP authentications. For more information, see Chapter 5, “Configuring LDAP Authentication,” on page 45.
ssl_certs	This utility helps with the certificate signing process to configure an SSL connection to the Sentinel server. For more information, see Section 3.3.8, “Using CA Signed Certificates,” on page 28.

A.3 Running the Report Development Utility

You can use the /opt/novell/sentinel/bin/report_dev_setup.sh utility to set up the report development environment. This utility does the following:

- Opens the PostgreSQL database port so that other systems can connect to the database.
- Updates the firewall to allow connection on the PostgreSQL database port.
- Modifies the database configuration files (postgresql.conf and pg_hba) so that other applications can connect to the database. The database configuration files are located at /var/opt/novell/sentinel/3rdparty/postgresql/data.
- Changes the rptuser password, if necessary, and saves it in an encoded format in the obj-component.JasperReportingComponent.properties file. This password can also be changed in the database.
- Collects the required Sentinel jar files, xml files, and the keystore file for report development and creates a tar file sentineljarsforireport.tar file in the /opt/novell/sentinel/bin directory.

To run this utility:

- 1 Log in as novell user.
- 2 Change to the following directory:

```
cd /opt/novell/sentinel/bin/
```

- 3 Run the following command:

```
./report_dev_setup.sh
```

A warning message is displayed, indicating that the Sentinel server will be restarted after the script is executed.

- 4 To continue running the script, press 1.
- 5 Specify the root password when prompted.

The script opens the database port, updates the firewall configuration files, and modifies the configuration files and database files.

- 6 When you are prompted to change the `rptuser` password, continue without changing the password.

or

Specify a password for `rptuser` and reconfirm the password.

The `rptuser` password is randomly generated during the installation of Sentinel. It is a recommended practice to change it here.

The Sentinel server restarts.

For information or help on commands, use the following command:

```
./report_dev_setup.sh -h
```

A.4 Getting the .jar Version Information

You can gather version information for Sentinel `.jar` files for troubleshooting purposes:

- 1 Log in to the Sentinel server as a user in the administrator role.
- 2 Go to the `/opt/novell/sentinel/bin` directory.
- 3 At the command line, specify the `./versionreader.sh <path/jar file name>`.

Running the script without any arguments gives the version of the installed Sentinel server. For more information on the arguments that can be used, use the `--help` command.

A.5 Changing the Hostname of a Sentinel Server

If the hostname of the Sentinel server changes, you must update any remote Correlation Engines or Collector Managers to point to the updated hostname. To do this, you run the `configure.sh` script on each remote machine.

The `configure.sh` script is located in `/opt/novell/sentinel/setup`. Run this script on each remote machine to update the hostname and IP address of the Sentinel server in the remote machine's configuration files.

The `configure.sh` file also allows you to change the password of the `appuser`. The `appuser` is an internal Sentinel identity that is used to establish a connection and interact with the PostgreSQL database.

A.6 Importing or Exporting Event Association Data

Sentinel provides the `event_assoc_data.sh` script that allows you to export event association data from the database to the file system, as well as import previously exported event association data from the file system back into the database. The script is `event_assoc_data.sh` located in the `/opt/novell/sentinel/bin` directory.

There are two types of event association data:

Incident events: There is a record in the database for every event that is associated with an incident, including what partition the event came from. When a partition is deleted, all incident events records for the partition are exported to a file on the file system, and the records are then deleted from the database. The file name is `incidents_events.json`.

Correlated events: There is a record in the database for every trigger event that is associated with a correlated event. The record also indicates what partition the correlated event belongs to. When a partition is deleted, all correlated event records for the partition are exported to a file on the file system, and the records are then deleted from the database. The file name is `correlated_events.json`.

When you export event association data, it is saved to the files in the following default directory structure `/var/opt/novell/sentinel/data/eventdata/exported_associations/<partition name>/*.json`

When a partition is restored from backup, the system automatically attempts to import the event association records for the partition. The `.json` file must be restored to the correct directory structure when the event association records are restored. If these files are not restored, the event association records are not imported, but the partitions are restored without this information. The event association records for the partitions are not available.

You can use the following options with for the `event_assoc_data.sh` file:

-i, --import: Imports event association data. This option works only on partitions that are currently in the restored state, but have not yet imported the event association data.

-x, --export: Exports event association data. This option works only on partitions that are currently in the deleted state, but have not exported their event association data.

-d, --days=<integer>: Specify the last number of days of the partitions.

-s, --startdate=<date>: Specify a start date and end date to select partitions in the specified date range.

-e, --enddate=<date>: Specify an end date and start date to select partitions in the specified date range.

--date=<date>: The utility selects the partitions with the specified date. You can use this option multiple times to select multiple dates.

-u, --user=<user name>: Specify the name of the user with administrative privileges to the Sentinel server.

-p, --password=<user password>: Specify the password of the administrative user.

--host=<host name>: Specify the host name or IP address of the Sentinel server.

--port=<port>: Specify the port number for communication to the Sentinel server. If this option isn't specified, the default port of 8443 (HTTPS) or 8000 (HTTP) is used.

--https: If this option is used, the utility communicates over HTTPS.

--http: If this option is used, the utility communicates over HTTP.

-h, --help: Displays the help options.

-l, --log file=FILE: Logs messages from the utility to the file name specified in the parameter.

--no-banner: Suppresses banner messages.

-q, --quiet: Displays fewer messages.

-v, --verbose: Displays more messages.

A.7 Managing the Internal Database

Sentinel provides a `db.sh` script that allows you to manage the internal database. You can use this script if you need to start the database without starting Sentinel so you can perform maintenance tasks. You can also use this script to run SQL commands against the internal database.

The script `db.sh` is located in the `/opt/novell/sentinel/bin` directory. The script has commands and options. You must be logged in as the user that installed Sentinel for the script to work. The command must come first, followed by the option.

For example: `./db.sh status --log-file=sentinel_status.txt`

This command writes the status of the internal database to the log file named `sentinel_status.txt`.

- ♦ [Section A.7.1, “Commands,” on page 244](#)
- ♦ [Section A.7.2, “Options,” on page 244](#)

A.7.1 Commands

You can use the following commands with the `db.sh` script.

start: Starts the internal database without starting the Sentinel server.

stop: Stops the internal database.

force_stop: Forces the database to stop when the Sentinel service is still running.

status: Displays the status of the internal database.

sql <db name> <user name> <sql statement>: Allows you to send SQL commands to the internal database.

restart: Restarts the internal database.

force_restart: Forces the database to restart when the Sentinel service is still running.

try-restart: Tries to restart the internal database.

reload: Reloads the internal database.

force-reload: Forces a reload of the database.

A.7.2 Options

You can use the following options with the `db.sh` script. The options must start with a `-` or `--` to be executed.

-w, --wait=<seconds>: Allows you to specify the amount of time to wait for the database to start or stop.

-h, --help: Displays help information for the script.

-l, --log-file=FILE: Logs messages to the specified file name.

--no banner: Suppresses banner messages.

-q, --quiet: Displays fewer messages.

-v, verbose: Displays more messages.

A.8 Cleaning Up the Internal Database

Sentinel provides a `clean_db.sh` script that allows you to clean up redundant data from the Sentinel database. You can delete data such as incidents, identities, assets, Advisor data, and vulnerabilities individually. You can run this script even without Sentinel running. For example, an improperly configured correlation rule might create hundreds of unwanted incidents in the database. Or, the identity information might encounter an error when someone attempts to delete the `IdentityAccountMap.csv` file. In such a situation, you can use this script to remove the unusable identity information.

The script `clean_db.sh` is located in the `/opt/novell/sentinel/bin` directory.

WARNING: Because this script is designed to delete information from your database, it should be used carefully and only after understanding the implications.

- ♦ [Section A.8.1, “Prerequisites,” on page 245](#)
- ♦ [Section A.8.2, “Using the `clean_db.sh` Script,” on page 245](#)

A.8.1 Prerequisites

- ♦ Ensure that you have permission to run the script. Only the user who installed Sentinel has permission to run this script.
- ♦ Ensure that the database is started and is running.

A.8.2 Using the `clean_db.sh` Script

- 1 In the Terminal mode, log in to Sentinel by using the credentials that were used to install Sentinel.

This script cannot be run by the `root` user.

- 2 Go to `<install_directory>/bin`, then specify `clean_db.sh` to run the script.

The following menu is displayed:

```
Which objects would you like to cleanup?
(1) Incidents
(2) Identities
(3) Assets
(4) Advisor
(5) Vulnerabilities
(6) Incidents and Identities
(7) All
(q) Quit without action
```

- 3 At the prompt, indicate which objects you want to remove from the database.
- 4 Specify the following information to connect to the PostgreSQL database:

```
Database server hostname (Press ENTER for default localhost)=>
Database name (Press ENTER for default SIEM) =>
Database username (press ENTER for default dbauser) =>
```

The database connection is verified before proceeding to the next step. If the connection was not successful, the script exits.

- 5 (Conditional) If you select 1 to delete Incidents data, several options are displayed. Select one of the options and specify the required information:
 - ♦ **Delete Incidents By Query:** Specify a custom `SELECT` query. For example:

```
select inc_id from incidents where inc_id=500
```

Ensure that SELECT statement does not include quotation marks.

- ♦ **Delete Incidents By Id:** Specify the ID of the Incident that you want to delete. For example:

```
101
```

- ♦ **Quit without action:** Specify q to exit from the script.

- 6 You are prompted to confirm data cleanup. Specify `start` to start the data cleanup or specify `abort` to quit without performing the data cleanup.

The results of the data cleanup are written to the log file. You should review the log file for any errors and retry.

If Identities data is being cleaned up, the script cleans up the Identities information from the database tables, and deletes the Identity Account Map file (`identityAccountMap.csv`).

NOTE: If you have a distributed Sentinel install, you might need to manually connect to the main Sentinel server to delete the `identityAccountMap.csv` file.

A.9 Managing the Sentinel Server

You can use the `server.sh` script to manually manage the Sentinel server. This script is located in the default directory of `/opt/novell/sentinel/bin`. The script has commands and options. The command must come first, followed by the option.

- ♦ [Section A.9.1, “Commands,” on page 246](#)
- ♦ [Section A.9.2, “Options,” on page 246](#)

A.9.1 Commands

You can use the following commands with the `server.sh` script:

start: Starts the Sentinel server.

stop: Stops the Sentinel server.

status: Displays the status of the Sentinel server.

restart: Restarts the Sentinel server.

try-restart: Tries to restart the Sentinel server.

force-reload: Forces a reload of the Sentinel server.

startdb: Starts the internal Sentinel database.

stopdb: Stops the internal Sentinel database.

force_stopdb: Forces the internal database to stop.

version: Displays the version of the Sentinel server.

A.9.2 Options

You can use the following options with the `server.sh` script:

-p, --priority=<integer>: Specifies the process priority for the Sentinel server.

-h, --help: Displays the help options.

-l, --log-file=FILE: Logs messages to a file you specify.

--no-banner: Suppresses banner messages.

-q, --quiet: Displays fewer messages.

-v, --verbose: Displays more messages.

