



NetIQ[®] Sentinel[™]

User Guide

February 2015

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.

© 2015 NetIQ Corporation. 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
1 Introduction to the Sentinel Interface	15
1.1 Sentinel Web Interface	15
1.2 Sentinel Control Center	15
1.3 Solution Designer	15
2 Viewing Events in Real-Time	17
2.1 Viewing Events in Real-Time Views	17
2.1.1 Creating an Event View	17
2.1.2 Viewing Events	18
2.2 Viewing Events in Active Views	18
2.2.1 Accessing the Active Views Tab	19
2.2.2 Reconfiguring Total Display Time	20
2.2.3 Viewing Real-Time Events	20
2.2.4 Managing Events	21
3 Searching Events	33
3.1 Running an Event Search	33
3.2 Viewing Search Results	35
3.2.1 Summary View	36
3.2.2 Detailed View	36
3.3 Refining Search Results	38
3.4 Saving a Search Query	39
3.4.1 Saving a Search Query as a Search Template	39
3.4.2 Saving a Search Query as a Filter	40
3.4.3 Saving a Search Query as a Report Template	40
3.4.4 Saving a Search Query as a Routing Rule	43
3.4.5 Saving a Search Query as a Retention Policy	43
3.4.6 Creating a Dashboard	44
3.5 Performing Event Operations	44
3.5.1 Executing Actions	44
3.5.2 Exporting the Search Results to a File	45
3.5.3 Adding Events to an Incident	45
3.5.4 Creating an Incident	46
3.5.5 Adding Events to a Correlation Rule	46
3.5.6 Creating a Correlation Rule by Using Events	46
3.5.7 Viewing Identity Details of Events	47
3.5.8 Viewing Advisor Report	47
3.5.9 Viewing Asset Data	47
3.5.10 Viewing Vulnerabilities	48
4 Configuring Filters	49
4.1 Overview	49
4.2 Introducing the Filters Interface	49
4.2.1 Filters Panel	50

4.2.2	Filter Criteria	50
4.3	Creating a Filter	52
4.3.1	Creating a Filter by Using the Edit Criteria Dialog	52
4.3.2	Creating a Filter by Using a Search Query	53
4.4	Sample Filters	54
4.4.1	View Events of Severity 3 to 5 from a System in China	54
4.4.2	Determine if User "Bob Smith" Tried to Log In after His Account was Disabled	54
4.4.3	View Events from Two Subnets and Share the Filter with Network Administrators	55
4.4.4	Find all Events that Include the Words "database" and "service," and exclude "test"	55
4.5	Viewing Events by Using Filters	56
4.6	Managing Filters	56
4.6.1	Editing a Filter	56
4.6.2	Deleting a Filter	56

5 Correlating Event Data 57

5.1	Overview	57
5.1.1	How Correlation Works	58
5.1.2	Correlation Rules	58
5.1.3	Correlation Engine	62
5.2	Accessing the Correlation User Interface	62
5.3	Understanding the Correlation Interface	63
5.3.1	Correlation Panel	63
5.3.2	Correlation Rule Builder	64
5.4	Creating Correlation Rules	68
5.4.1	Understanding the Correlated Event	69
5.4.2	Creating a Simple Rule	70
5.4.3	Creating a Sequence Rule	71
5.4.4	Creating a Composite Rule	72
5.4.5	Creating a Free-Form Rule	74
5.4.6	Creating Correlation Rules From Search Results	74
5.5	Associating Actions to a Rule	75
5.6	Testing a Correlation Rule	75
5.7	Sample Correlation Rules	76
5.7.1	Detecting Critical Events from an Intrusion Detection System	76
5.7.2	Detecting a Spreading Attack	77
5.7.3	Detecting an Attack that Came from Outside the Firewall	77
5.8	Deploying Rules in the Correlation Engine	78
5.9	Viewing Correlated Events	78
5.10	Customizing Correlated Event	79
5.11	Managing Correlation Rules	79
5.11.1	Viewing the Rule Dashboard	80
5.11.2	Editing a Rule	81
5.11.3	Deleting a Rule	81
5.12	Managing the Correlation Engine	81
5.12.1	Using the Correlation Engine Dashboard	81
5.12.2	Stopping or Starting a Correlation Engine	83
5.12.3	Renaming a Correlation Engine	84

6 Visualizing and Analyzing Alerts 85

6.1	Viewing and Triaging Alerts in Alert Views	85
6.1.1	Creating an Alert View	85
6.1.2	Viewing Alerts	86
6.1.3	Creating Incidents from Alerts	88
6.2	Analyzing Alert Dashboards	88
6.2.1	Creating the Alert Dashboard	89

6.2.2	Analyzing Alerts	89
6.2.3	Customizing the Alert Dashboard	90
6.3	Troubleshooting	90
6.3.1	Unable to View Alerts in the Dashboard and Alert Views	91
7	Analyzing Trends in Data	93
7.1	Overview	93
7.1.1	Terminology	93
7.1.2	How Security Intelligence Works	95
7.1.3	Permissions for Security Intelligence	96
7.2	Creating a Dashboard	96
7.2.1	Creating a Dashboard by Using a Filter.	97
7.3	Understanding the Dashboard Interface	97
7.4	Creating Baselines	98
7.5	Configuring Anomaly Detection	99
7.5.1	Creating an Anomaly Definition	99
7.5.2	Deploying an Anomaly Definition.	100
7.5.3	Undeploying an Anomaly Definition.	100
7.5.4	Managing Anomalies	101
7.6	Viewing Anomaly Events	101
7.7	Managing Dashboards	103
7.7.1	Viewing a Dashboard.	103
7.7.2	Renaming a Dashboard.	104
7.7.3	Deleting a Dashboard	104
7.8	Troubleshooting	104
7.8.1	The Create Button Is Not Displayed	104
7.8.2	The Main Graph and the Time Slider Are Not Synchronized.	104
7.8.3	Both Names for a Renamed Anomaly Are Displayed in the Filter.	104
7.8.4	Dashboard Date Range Not Updated to in Real Time.	105
8	Visualizing and Analyzing Network Flow Data	107
9	Configuring Dynamic Lists	109
9.1	Creating a Dynamic List	109
9.1.1	Using the Sentinel Control Center to Create a Dynamic List.	109
9.1.2	Using the Correlation Rule Builder to Create a Dynamic List	110
9.2	Managing Dynamic Lists	111
9.2.1	Editing a Dynamic List	111
9.2.2	Deleting a Dynamic List	111
9.2.3	Removing Dynamic List Elements	112
10	Leveraging Identity Information	113
10.1	Overview	113
10.2	Searching and Viewing User Identities.	113
10.2.1	Accessing the Identity Browser	113
10.2.2	Performing a Search	114
10.2.3	Searching.	114
10.2.4	Viewing Profile Details	115
10.2.5	Viewing Activity	116
11	Manually Performing Actions on Events	117
11.1	Accessing Event Actions	117
11.2	Prerequisites for Assigning Actions to Events	117

11.3	Assigning Actions to Events	117
11.4	Configuring Event Actions	118
11.4.1	Creating a New Event Action	118
11.4.2	Cloning an Event Action	119
11.4.3	Moving an Event Action	119
11.4.4	Deleting an Event Action	119

12 Configuring Tags 121

12.1	Overview	121
12.2	The Tags Interface	121
12.3	Creating a Tag	122
12.4	Managing Tags	122
12.4.1	Sorting Tags	122
12.4.2	Adding and Removing Tags from Favorites	123
12.4.3	Viewing and Modifying Tags	123
12.5	Performing Text Searches for Tags	123
12.6	Deleting Tags	123
12.7	Associating Tags with Objects	124
12.7.1	Associating Tags with Event Routing Rules	124
12.7.2	Associating Tags with Event Sources	124
12.7.3	Associating Tags with Collector Managers	124
12.7.4	Associating Tags with Event Sources Servers	124
12.7.5	Associating Tags with Collector Plug-Ins	125
12.7.6	Associating Tags with Report Results and Report Definitions	125
12.8	Viewing Tagged Events	125

13 Reporting 127

13.1	Importing Report Definitions	127
13.2	Creating Reports	128
13.3	Scheduling a Report	131
13.4	Grouping Reports Based on Category	132
13.5	Viewing Events	132
13.5.1	Viewing Events Based on Report Criteria	132
13.5.2	Viewing Events Based on Report Result Criteria	132
13.5.3	Finding Reports Based on Words	133
13.5.4	Finding Reports Based on Tags	133
13.6	Renaming a Report Result	133
13.7	Marking Report Results as Read or Unread	133
13.8	Managing Favorite Reports	134
13.8.1	Adding Reports as Favorites	134
13.8.2	Removing Favorite Reports	134
13.9	Associating Tags with Report Results and Report Definitions	134
13.10	Exporting Report Definitions and Report Results	135
13.10.1	Exporting a Single Report Definition	135
13.10.2	Exporting Multiple Report Definitions	135
13.10.3	Exporting All Report Definitions	135
13.10.4	Exporting a Report Result	136
13.11	Deleting Reports	136
13.11.1	Deleting a Report Definition	136
13.11.2	Deleting Multiple Report Definitions	136
13.11.3	Deleting a Report Result	137
13.11.4	Deleting Multiple Report Results	137
13.12	White Label Template Report	137

14 Viewing Compliance to Configuration Policies	139
14.1 Viewing Secure Configuration Manager Events and Compliance Details	139
15 Viewing Change Guardian Events	141
16 Configuring Incidents	143
16.1 Accessing Incidents	143
16.2 Creating Incidents	143
16.3 Managing Incidents	144
16.3.1 Viewing an Incident	144
16.3.2 Attaching Workflows to Incidents	145
16.3.3 Adding Attachments to Incidents	145
16.3.4 Adding Notes to Incidents	145
16.3.5 Executing Incident Actions	145
16.3.6 E-mailing an Incident	146
16.4 Adding an Incident View	146
17 Configuring iTRAC Workflows	147
17.1 Overview	147
17.2 Accessing the iTRAC Administration Tools	148
17.3 Using the Template Manager	149
17.3.1 Default Templates	149
17.4 Template Builder Interface	150
17.5 Creating a Template	152
17.6 Managing Templates	152
17.6.1 Viewing or Editing a Template	152
17.6.2 Copying a Template	153
17.6.3 Deleting a Template	153
17.7 Steps	153
17.7.1 Start Step	153
17.7.2 Manual Steps	154
17.7.3 Decision Steps	155
17.7.4 Mail Steps	155
17.7.5 Command Steps	155
17.7.6 Activity Steps	156
17.7.7 End Step	157
17.8 Adding Steps to a Workflow	157
17.8.1 Adding a Step from the Step Palette	157
17.8.2 Adding a Step in the Process Builder	157
17.8.3 Adding an Activity Step	158
17.8.4 Adding an End Step	158
17.9 Managing Steps	158
17.9.1 Copying a Step	158
17.9.2 Modifying a Step	159
17.9.3 Editing a Manual Step	159
17.9.4 Editing a Decision Step	159
17.9.5 Editing a Mail Step	160
17.9.6 Editing a Command Step	160
17.9.7 Deleting a Step	160
17.10 Transitions	161
17.10.1 Unconditional Transitions	161
17.10.2 Conditional Transitions	162
17.10.3 Creating an Expression	163
17.10.4 Else Transitions	164

17.10.5	Timeout Transitions	164
17.10.6	Alert Transitions	165
17.10.7	Error Transition	166
17.10.8	Managing Transitions	167
17.11	Activities	167
17.11.1	Incident Command Activity	168
17.11.2	Incident Internal Activity	168
17.11.3	Incident Composite Activity	168
17.12	Creating iTRAC Activities	169
17.13	Managing Activities	170
17.13.1	Editing an Activity	170
17.13.2	Exporting an Activity	170
17.13.3	Importing an Activity	170
17.14	Managing iTRAC Roles	171
17.14.1	Adding a Role	171
17.14.2	Deleting a Role	171
17.14.3	Viewing the Role Details	171
17.15	Process Management	172
17.15.1	Instantiating a Process	172
17.15.2	Automatic Step Execution	172
17.15.3	Manual Step Execution	173
17.15.4	Display Status	173
17.15.5	Displaying the Status of a Process	173
17.15.6	Changing Views in Process Manager	174
17.15.7	Starting or Terminating a Process	174

18 Managing Work Items 177

18.1	Overview	177
18.2	Understanding the Work Item Summary Interface	177
18.3	Viewing a Work Item	178
18.4	Processing a Work Item	179
18.5	Managing Work Items Of Other Users	179

A Search Query Syntax 181

A.1	Basic Search Query	181
A.1.1	Case Insensitivity	182
A.1.2	Special Characters	182
A.1.3	Operators	182
A.1.4	The Default Search Field	183
A.1.5	Tokenized Fields	184
A.1.6	Non-Tokenized Fields	186
A.2	Wildcards in Search Queries	186
A.2.1	Wildcards in Tokenized Fields	186
A.2.2	Quoted Wildcards	187
A.2.3	Leading Wildcards	187
A.3	The notnull Query	188
A.4	Tags in Search Queries	188
A.5	Regular Expression Queries	189
A.6	Range Queries	189
A.7	IP Addresses Query	190
A.7.1	CIDR Notation	190
A.7.2	Wildcards in IP Addresses	190

B	Correlation Rule Expression Syntax	193
B.1	Event Fields	193
B.2	Event Operations	194
B.2.1	Filter Operation	194
B.2.2	Trigger Operation	197
B.2.3	Window Operation	198
B.2.4	Gate Operation	200
B.2.5	Sequence Operation	200
B.2.6	Distinct Operation	201
B.3	Operators	201
B.3.1	Flow Operator	201
B.3.2	Union Operator	202
B.3.3	Intersection Operator	202
B.4	Order of Operators	202

About this Book and the Library

The *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.

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.

Administration Guide

Provides the administration information and tasks required to manage a Sentinel deployment.

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>.

1 Introduction to the Sentinel Interface

Sentinel is a Security Information and Event Management (SIEM) solution that receives information from many sources throughout an enterprise, standardizes it, prioritizes it, and presents it to you to make threat, risk, and policy decisions.

There are different tools to help you take advantage of all of the features Sentinel has to offer: You must have necessary permissions to access these tools.

- ♦ [Section 1.1, “Sentinel Web Interface,” on page 15](#)
- ♦ [Section 1.2, “Sentinel Control Center,” on page 15](#)
- ♦ [Section 1.3, “Solution Designer,” on page 15](#)

1.1 Sentinel Web Interface

The Sentinel Web interface is the main user interface for viewing and interacting with Sentinel data.

For more information about the user interface and its options, see [“Understanding Sentinel Applications”](#) in the *NetIQ Sentinel Administration Guide*.

1.2 Sentinel Control Center

Sentinel presents the collected data in the Sentinel Web interface as well as the Sentinel Control Center (SCC). For more information on the Sentinel Control Center, see [“Understanding Sentinel Applications”](#) in the *NetIQ Sentinel Administration Guide*.

1.3 Solution Designer

You can use the Solution Designer to package and export different contents, such as filters, reports, searches, and Correlation rules with associated actions and dynamic lists. For more information on Solution Designer, see [“Solution Designer”](#) in the *NetIQ Sentinel Administration Guide*.

2 Viewing Events in Real-Time

An event represents a normalized log record reported to Sentinel from a third-party security device, network or application device, or from an internal Sentinel source. There are several types of events:

- ◆ External events (events received from a security device) such as:
 - ◆ An attack detected by an intrusion detection system (IDS)
 - ◆ A successful login reported by an operating system
 - ◆ A customer-defined situation such as a user accessing a file
- ◆ Internal events (events generated by Sentinel), including:
 - ◆ A correlation rule being disabled
 - ◆ The database filling up

Sentinel displays events in near real-time in a graphical and tabular representation. You can view and analyze events in Real-time Views in the Sentinel Web interface and also in Active Views in the Sentinel Control Center. Viewing events in Real-time Views enables you to view events without having to log in to the Sentinel Control Center. To perform any event operations, you can launch the **Search** interface. However, in Active Views, you can view the events and also perform event operations in the Active Views table.

- ◆ [Section 2.1, “Viewing Events in Real-Time Views,” on page 17](#)
- ◆ [Section 2.2, “Viewing Events in Active Views,” on page 18](#)

2.1 Viewing Events in Real-Time Views

To view events in Real-time Views, you must first create an event view.

The administrator can see all events in the system. The events other users see may be limited by a security or tenant filter. For information about permissions to view events, see “[Creating Roles](#)” in the *NetIQ Sentinel Administration Guide*.

2.1.1 Creating an Event View

To view and analyze events, you must first create an event view in the **Real-time Views** panel.

To create an event view:

- 1 Log in to the Sentinel Web console.
- 2 Click **Real-time Views > Events > Create**.
- 3 Specify the following information:
 - ◆ **Name:** Specify a unique name for the event view.
 - ◆ **Sharing:** Select either of the following options:
 - ◆ **Public:** Allow everyone to view the event view. In the **Public** mode, other users can only view the events but cannot modify the event view. You are still the owner of the event view.

- ♦ **Private:** Only you are able to view the event view.
 - ♦ **Criteria:** Specify the criteria to view specific events.
 - ♦ **Event Attribute:** Select the attribute based on which you want to categorize the event data.
 - ♦ **Tenant:** If you are in a multi-tenant environment, select a tenant name for which you want to view events. The default tenant allows you to view events from all tenants. If you select a tenant, only users of that tenant can view the events in this event view.
This option is available only if you are an administrator in a multi-tenant environment.
 - ♦ **Chart Type:** Specify the chart type in which you want to view the event data.
 - ♦ **Y Axis:** Select either of the following options:
 - ♦ **Event Count:** Displays a graph with number of events for the specified time range.
 - ♦ **Event Count per Second:** Displays a graph with event rate for the specified time range.
 - ♦ **Time range:** Select the time range for which you want to view the event data.
 - ♦ **Display Interval:** Select the time interval between two data points.
- 4 Click **Save** to save the event view configuration.

2.1.2 Viewing Events

To view an event view:

- 1 Click **Real-time Views > Events**.
- 2 Select the event view and click **Open the event view**.

Sentinel provides a graphical representation of events for the specified criteria. The chart automatically refreshes after the interval specified in **Display Interval**.

As you are viewing event data, you can perform the following actions in the chart:

- ♦ Mouse over the data points in the chart to view the number of incoming events or events per second for a specific timestamp.
- ♦ Click any category in the legend to filter the view by the legend items.
- ♦ Click and drag the mouse to zoom the view for a specific time range.

The event view enables you to view only the summarized event data. To view the event details or perform any event operations, you can do either of the following:

- ♦ Click a specific area in the chart to open the Search interface with the list of events represented in that area.
- ♦ Click **Search Events**.

For information about viewing event details and performing event operations, see [Section 3.2, “Viewing Search Results,” on page 35](#) and [Section 3.5, “Performing Event Operations,” on page 44](#).

2.2 Viewing Events in Active Views

Sentinel displays events in near real-time in Active Views tab in the Sentinel Control Center. As you view events, you can also perform relevant operations on the events, including:

- ♦ Create incidents or add the events to an existing incident
- ♦ For correlated events, you can view the events that triggered the correlated event

- ◆ Execute actions on the events
- ◆ View the Advisor and asset data

For more information, see [Section 2.2.4, “Managing Events,” on page 21](#).

NOTE: You must have the Create and use Active Views permission to view and manage events in Active Views.

2.2.1 Accessing the Active Views Tab

When you launch the Sentinel Control Center, the Active Views tab is the first tab that is displayed. If the **Active Views** tab is not displayed, you can access it by:

- 1 Log in to the Sentinel Web interface, then click **Applications** in the toolbar.
- 2 Click **Launch Control Center**.
- 3 Log in to the Sentinel Control Center as a user with Create and use Active Views permission.
The **Active Views** tab is displayed.

Active Views provides two types of views, which display the events in tables and graphs.

Table format displays the variables of the events as columns in a table. You can sort the information in the grid by clicking the column name.





Graphical format displays events as graphs. You can change the chart types by right-clicking anywhere in the chart and by selecting the desired chart type. You can also view the events that match the filter criteria by right-clicking anywhere in the chart and by selecting the **Drill-down** option from the menu. The events are displayed in the Sentinel Web interface.

The event table and the snapshot are the two types of Active Views.

Near Real Time Event Table: Displays the events in graphs with the following features:

- ◆ Holds up to 750 events per 30-second period. If there are more than 750 events, the events are prioritized to display correlated events first, then events that are sent to the GUI by using routing rules, then all remaining events.
- ◆ By default, the client maintains a 24-hour period of cached events. You can configure the time. For more information, see [Section 2.2.2, “Reconfiguring Total Display Time,” on page 20](#).
- ◆ By default, the smallest possible display interval of an active view is 30 seconds. This is represented by a gray line in the event table.

Figure 2-1 Gray Line- Indicating the Smallest Possible Display Interval

Severity	EventTime	XDASOutcomeName	XDASTaxonomyName
	1/27/12 5:21:33 PM	XDAS_OUT_SUCCESS	XDAS_AE_QUERY_DATA_ITEM_CONTENTS
	1/27/12 5:21:33 PM	XDAS_OUT_SUCCESS	XDAS_AE_QUERY_DATA_ITEM_CONTENTS
	1/27/12 5:21:24 PM	XDAS_OUT_SUCCESS	XDAS_AE_CREATE_SESSION
	1/27/12 5:21:24 PM	XDAS_OUT_SUCCESS	XDAS_AE_AUTHENTICATE_ACCOUNT

If there are more than 750 events per 30-second time period, a red separation line indicates that there are more events than are displayed.

Figure 2-2 Red Line- Indicating More Events to Display

Severity	EventTime	XDASOutcomeName ▾	Tags
0	1/27/12 3:29:49 PM	XDAS_OUT_SUCCESS	Sentinel
1	1/27/12 3:29:49 PM	XDAS_OUT_SUCCESS	Sentinel
1	1/27/12 3:29:49 PM	XDAS_OUT_SUCCESS	Sentinel

- ◆ When you save user preferences, the system continues to collect data for four days. For instance, if you save your preferences, then log out and log back in the following day, your Active View displays data as if you never logged off.
- ◆ If an Active View is created and not saved, it continues to collect data for an hour. Within that hour, if an identical Active View is created, the Active View displays data for the last hour.

Snapshot: Time-stamped view of a **Real Event View** table.

Several features make an Active View unique.

- ◆ The filter assigned to an Active View
- ◆ The z-axis attribute
- ◆ The security filter assigned to a user

You can change event names to user-friendly names and the new names are populated throughout the system. For more information, see “[Renaming Event Fields](#)” in the *NetIQ Sentinel Administration Guide*.

2.2.2 Reconfiguring Total Display Time

Active View Properties allows you to configure the cached time in each client. The default cache time value in an Active View is 24 hours.

To reconfigure the total display time:

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,”](#) on page 19.
- 2 In the menu, click **Active Views > Properties**.
- 3 Select the maximum display time, then click **OK**.
The new values do not take effect until you restart the Sentinel Control Center.

2.2.3 Viewing Real-Time Events

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,”](#) on page 19.
- 2 In the menu, click **Active Views > Create Active View**.
- 3 Use the following information to define the display properties:
 - Event Attribute (Z Axis):** In drop-down list, select the desired attribute.
 - Filter:** Specify the filter, or click the **Browse** button to select a defined filter.
 - Display Events:** Select where to display event.
- 4 To customize additional settings, click **Next**, then continue with [Step 5](#).
or

To use the default settings, click **Finish**. The following default values are selected:

- ◆ A display interval and refresh rate of 30 seconds
- ◆ Total display time of 15 minutes
- ◆ Y-axis as the event count
- ◆ A chart type of Stacked Bar 2D

5 Define the statistical parameters:

Display Interval: Select the display interval. This is the interval to display events.

Refresh Rate: Select the refresh rate. This is the rate at which Active Views refresh.

Total Display Time: Specify the amount of time to display the chart.

Y-axis: Set the Y axis to be either the total event count or event count per second.

6 Click **Next**.

7 Select your chart type from the drop-down list:

- ◆ Stacked Bar 2D
- ◆ Bar 3D
- ◆ Line
- ◆ Ribbon

8 Click **Finish**.

The five buttons to the left of the chart perform the following functions:

Lock Chart or Unlock Chart: Used when performing a drill-down, zoom in, zoom out, zoom to selection, and when saving a chart as an `.html` file. When you click the **Lock** button, you have additional options:

- ◆ **Zoom In:** Zooms in without changing any of the time settings of the chart.
- ◆ **Zoom Out:** Zooms out without changing any of the time settings of the chart.
- ◆ **Zoom to Selection:** Zooms in on a selection of time intervals of events.
- ◆ **Snapshot Active View:** Saves an `.html` file with the chart as an image and events in a tabular format.

Increase Display Interval: Increases the display time interval for incoming events.

Decrease Display Interval: Decreases the display time interval for incoming events.

Increase Display Time: Increases the time interval along the x-axis.

Decrease Display Time: Decreases the time interval along the x-axis.

2.2.4 Managing Events

In the table view, you can manage one or more events. Select the event or events you want to manage, then right-click and select the option.

- ◆ [“Showing and Hiding Event Details” on page 22](#)
- ◆ [“Sending Mail Messages about Events and Incidents” on page 22](#)
- ◆ [“Creating Incidents” on page 23](#)
- ◆ [“Adding Events to an Incident” on page 23](#)
- ◆ [“Viewing Events That Trigger Correlated Events” on page 23](#)

- ◆ “Executing Actions on Events” on page 24
- ◆ “Investigating an Event or Events” on page 24
- ◆ “Accessing the Active Browser” on page 26
- ◆ “Viewing Advisor Data” on page 27
- ◆ “Viewing Asset Data” on page 27
- ◆ “Viewing Vulnerabilities” on page 28
- ◆ “Viewing User Information” on page 30
- ◆ “Viewing the Targets” on page 30
- ◆ “Managing Columns” on page 30
- ◆ “Taking a Snapshot of a Navigator Window” on page 31

Showing and Hiding Event Details

To show details in a Real Time Event Table of the Navigator or Snapshot, double-click or right-click an event, then click **Show Details**.

The details display in the left panel of the Real Time Event Table.

To close the Event Details window in a Real Time Event Table of the Navigator or Snapshot, right-click an event, then click **Show Details**.

Sending Mail Messages about Events and Incidents

To send mail messages from within the Sentinel Control Center, you must have an SMTP integrator that is configured with connection information and with the `SentinelDefaultEMailServer` set to `True`. For more information on configuring the SMTP Integrator, see “[Configuring Integrators](#)” in the *NetIQ Sentinel Administration Guide*.

To send an event message by e-mail:

- 1 In a Real Time Event Table, right-click an event or a group of events, then select **Send Email**.
The e-mail is sent to the specified recipients.

To e-mail an incident:

- 1 After you save your incident, click the **Incidents** tab, **Incidents > Incidents View**.
- 2 Click **All Incidents** in the **Switch View** drop-down list located at the bottom right corner.
- 3 Double-click an incident.
- 4 Click the **Email Incident** button.
- 5 Provide the following information:
 - ◆ Email Address
 - ◆ Email Subject
 - ◆ Email Message
- 6 Click **OK**.

The e-mail messages have `.html` attachments that address incident details, events, assets, vulnerabilities, advisor information, attachment information, incident notes, and incident history.

Creating Incidents

Creating an incident is useful in grouping a set of events together to indicate a pattern of interest, such as an attack.

You must have the View or create incidents and add events to incidents permission to create incidents.

If events are not initially displayed in a newly created incident, it is probably because of a lag in the time between display in the Real Time Events window and insertion into the event store. It takes a few minutes for the original events to be inserted into the event store and displayed in the incident.

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19.](#)
- 2 In a Real Time Event Table, select an event or a group of events and right-click.
- 3 Select **Create Incident**.
- 4 Follow the steps in [Section 16.2, “Creating Incidents,” on page 143.](#)

Adding Events to an Incident

You must have the View or create incidents and add events to incidents permission to add events to Incidents.

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19.](#)
- 2 In a Real Time Event Table, select an event or a group of events and right-click.
- 3 Select **Add To Incident**.
- 4 Click **Browse**, then click **Search** to list the available incidents.
You can define your criteria to better search for a particular incident or incidents in the Select Incident window.
- 5 Select an incident, then click **Add**.
- 6 Click **OK**.

If events are not initially displayed in a newly created incident, it is probably because of a lag in the time between display in the Real Time Events window and insertion into the event store. It takes a few minutes for the original events to be inserted into the event store and display in the incident.

Viewing Events That Trigger Correlated Events

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19.](#)
- 2 In the Real Time Event Table, right-click a correlated event, then select **View Trigger Events**.
The **View Trigger Events** option is enabled only for Correlated events. [“Viewing the Targets” on page 30.](#) The events that generated the correlation event are displayed in the Sentinel Web interface.

Executing Actions on Events

You can execute the following actions on events:

- ◆ Log to File
- ◆ Log to Syslog
- ◆ Send Email
- ◆ Send SNMP trap
- ◆ Send via Sentinel Link

1 Access the **Actives Views** tab.

For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19](#).

2 In a Real Time Event Table, select an event or a group of events, right-click, then select the action you want to execute.

Investigating an Event or Events

You can use the Active Views tab to investigate events in two ways:

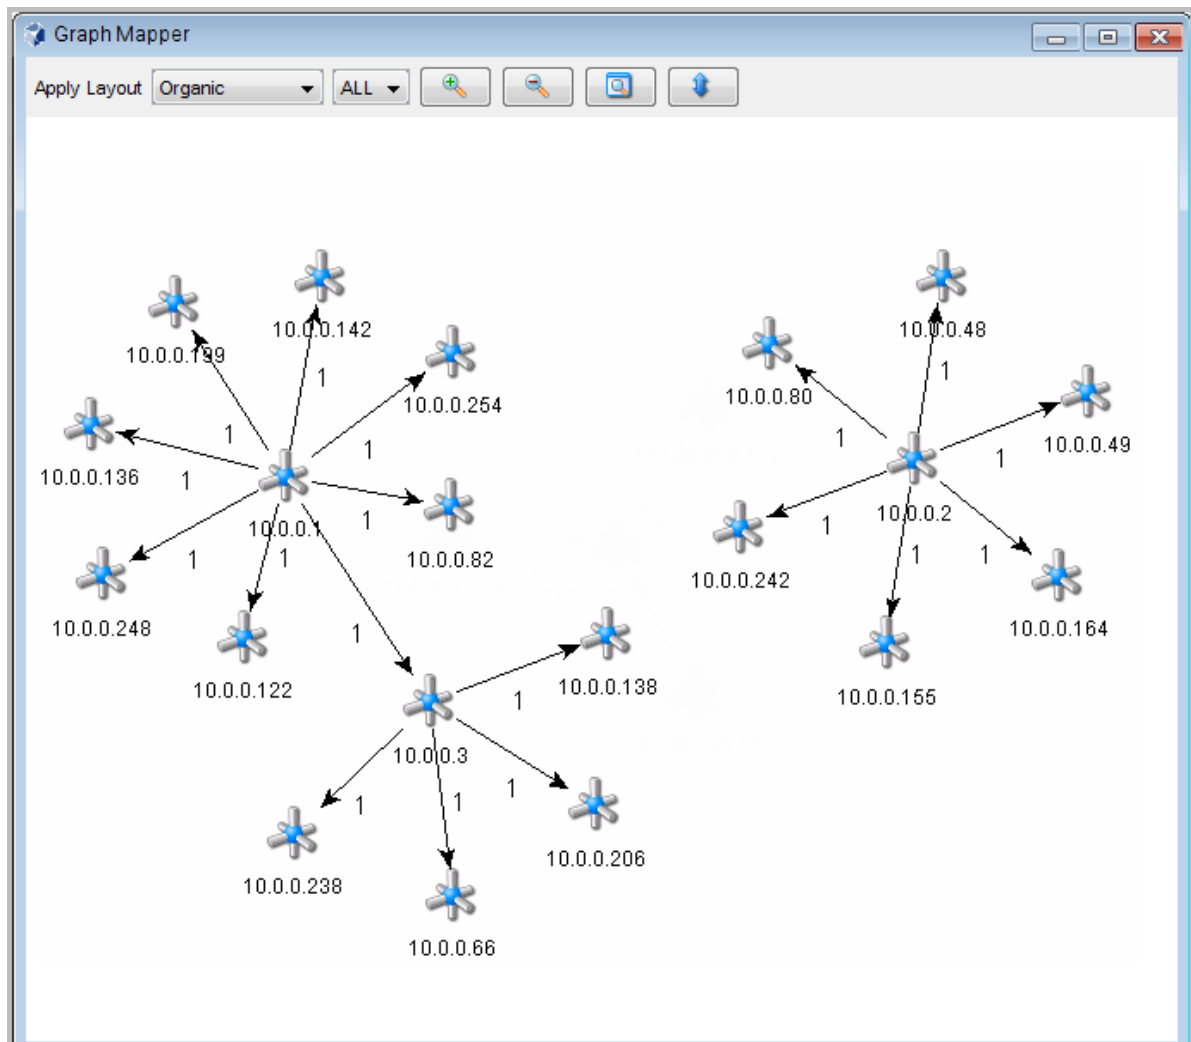
- ◆ You can perform an event query for the last hour on a single event for:
 - ◆ Other events with the same target IP address
 - ◆ Other events with the same source (initiator) IP address
 - ◆ Other targets with the same event name

NOTE: You cannot perform a query on a null (empty) field.

- ◆ You can graphically display the mappings between any two fields in the selected events. This is particularly useful to view the relationship between the initiators (IP, port, event, sensor type, Collector) and the targets (IP, port, event, sensor type, Collector name) of the selected events, but any fields can be used

The following illustration shows initiator IP addresses mapped to target IP addresses.

Figure 2-3 Graph Mapper



- ◆ [“Using the Investigate Option for an Event Query” on page 25](#)
- ◆ [“Using the Investigate Option with the Graph Mapper” on page 26](#)

Using the Investigate Option for an Event Query

You can perform an event query within the last hour for events similar to the selected event.

- 1 In a Navigator or Snapshot window, right-click an event, then click **Investigate** and select one of the options below:

Show More Events to this target: Events with the same destination IP address

Show More Events from this source: Events with the same initiator IP address

What are the target objects of this event?: Events with the same event name as the selected event.

The chosen event information is displayed in the Sentinel Web interface.

Using the Investigate Option with the Graph Mapper

To create a graph map:

- 1 In a Real Time Event Table, right-click one or more events, then select **Investigate > Show Graph**.
- 2 Select the desired information in the **From** and **To** fields, then click **Finish**.
- 3 In the Graph Mapper window, select how the information is displayed:
 - ◆ Circular
 - ◆ Hierarchical
 - ◆ Organic
 - ◆ Orthogonal

Accessing the Active Browser

The Active Browser provides the ability to browse through a selected set of data to look for patterns and perform investigation. You can view the selected events in the Active Views in the Active Browser. You can also perform all the right-click activities that are available in Active Views in the Active Browser.

- ◆ [“Viewing Events in the Active Browser” on page 26](#)
- ◆ [“Searching in the Active Browser” on page 26](#)
- ◆ [“Adding Attributes in the Active Browser” on page 26](#)

Viewing Events in the Active Browser

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19](#).
- 2 Select the events you want to view in the Active Browser.
- 3 Right-click the events, then select **View in Active Browser**.
The selected events display in the Active Browser window.

Searching in the Active Browser

- 1 Access the Active Views tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19](#).
- 2 Specify the value or text you want to search for in the **Search** field, then click the **Search** button.

You can move between the various searches by using the **Forward** and **Backward** buttons above the search field.

Adding Attributes in the Active Browser

- 1 Access the Active Views tab.
For more information, see [Section 2.2.1, “Accessing the Active Views Tab,” on page 19](#).
- 2 Click the **Add an attribute for categorization** button.
- 3 Select an attribute from the list.
- 4 Click **OK**.

Viewing Advisor Data

Advisor provides a cross-reference between real-time intrusion detection system (IDS) attack signatures and Advisor's knowledge base of vulnerabilities. The supported intrusion detection systems are listed in “[Detecting Vulnerabilities and Exploits](#)” in the *NetIQ Sentinel Administration Guide*.

To view Advisor data:

- 1 In a Real Time Event Table, right-click one or more events, then click **Analyze > Advisor Data**.
If the **IDS AttackName** field is properly populated, a report similar to the one below is displayed in the Sentinel Web interface.

Advisor Report

Table of Contents:

- Summary
- Details
- Selected Events

Summary

Attack	Reference
TippingPoint,2176: SMB: Null Session SetUp	Advisor: 299 CVE IDs: 1999-0519, 1999-0520 BugTraq IDs: No info found on this attack in the Advisor database

Details

[top](#)

(id 299)

Microsoft Windows NetBIOS Shares Access Control Weakness

Description:
Microsoft Windows NetBIOS Shares Access Control

Scenario:

Impact:
Loss Of Confidentiality, Loss Of Integrity

Attack category:
Misconfiguration

Solution:

Patches:

4
Urgency

7
Severity

Viewing Asset Data

This feature allows you to view and save your view as an HTML file of your Asset Report. To view the Asset data, you must run an asset management tool such as NMAP and bring the results into Sentinel by using an Asset Collector. The available data for viewing are:

- ◆ Hardware
 - ◆ MAC Address
 - ◆ Name
 - ◆ Type
 - ◆ Vendor
 - ◆ Product
 - ◆ Version
 - ◆ Value
 - ◆ Criticality

- ◆ Network
 - ◆ IP Address
 - ◆ Hostname
- ◆ Software
 - ◆ Name
 - ◆ Type
 - ◆ Vendor
 - ◆ Product
 - ◆ Version
- ◆ Contacts
 - ◆ Order
 - ◆ Name
 - ◆ Role
 - ◆ Email
 - ◆ Phone Number
- ◆ Location
 - ◆ Location
 - ◆ Address

If both the Source IP and Destination IP are populated in an event, the asset data is displayed for both. If either of them is populated, the respective asset data is displayed.

To view asset data:

- 1 In a Real Time Event Table, right-click one or more events, then click **Analyze > Asset Data**.
The Asset data is displayed in the Sentinel Web interface.

Viewing Vulnerabilities

Vulnerability Visualization provides a textual or graphical representation of the vulnerabilities of selected destination systems. Vulnerabilities for the selected destination IPs can be seen for the current time or for the time of the selected events.

Vulnerability Visualization requires that a vulnerability Collector is running and is adding vulnerability scan information to the Sentinel database. The [Sentinel Plug-in Web site](#) provides Collectors for several industry-standard vulnerability scanners, and additional vulnerability Collectors can be written by using the Collector Builder.

NOTE: Vulnerability Collectors are distinct from Event Collectors and use different commands.

There are several Vulnerability Visualization views:

- ◆ HTML
- ◆ Graphical
 - ◆ Circular
 - ◆ Organic

- ◆ Hierarchical
- ◆ Orthogonal

The HTML view is a report view that lists relevant fields, depending on which vulnerability scanner you have:

- ◆ IP
- ◆ Host
- ◆ Vulnerability
- ◆ Port/protocol

The graphical display is a rendering of vulnerabilities that links them to an event through common ports. There are four available views:

- ◆ Organic view
- ◆ Hierarchical view
- ◆ Circular view
- ◆ Orthogonal view

The graphical display has four panels:

- ◆ Graph panel
- ◆ Tree panel
- ◆ Control panel
- ◆ Details/events panel

The graph panel display associates vulnerabilities to a port/protocol combination of a resource (IP address). For example, if a resource has five unique port/protocol combinations that are vulnerable, there are five nodes attached to that resource. The resources are grouped together under the scanner that scanned the resources and reported the vulnerabilities. If two different scanners are used, such as ISS and Nessus, there are two independent scanner nodes that have vulnerabilities associated with them.

NOTE: Event mapping takes place only between the selected events and the vulnerability data returned.

The tree panel organizes data in same hierarchy as the graph. The tree panel also allows users to hide or show nodes at any level in the hierarchy.

The control panel exposes all the functionality available in the display. This includes:

- ◆ Four different algorithms to display
- ◆ The ability to show all or selected nodes that have events mapped to them
- ◆ Zooming in and out of selected areas of the graph

There are two tabs. When you are in the **Details** tab, click on a node displays node details. When you are in the **Events** tab, clicking an event associated with a node displays a table in a Real Time window.

To run a vulnerability visualization:

- 1 In a Real Time Event Table, right-click one or more events, then click **Analyze** and select one of the following options:

Current Vulnerability: Queries the database for vulnerabilities that are active (effective) at the current date and time.

Event Time Vulnerability: Queries the database for vulnerabilities that were active (effective) at the date and time of the selected event.

Viewing User Information

NetIQ provides optional integration with Novell Identity Manager. With this integration, user identity information is added to incoming events when the account name matches one from Novell Identity Manager. When the `InitUserIdentity` or `TargetUserIdentity` column is populated in an event, a right-click menu is enabled to open the user's page in the Identity Browser.

When you right-click an event and select **Show Identity Details**, you can choose to view the identity of the Initiator user, the Target user, or both. The Identity Browser opens and shows identifying information about the user (or users) from the identity management system, all the accounts to which the user is provisioned, and the recent activity by that user. For more information on the Identity Browser, see [Chapter 10, "Leveraging Identity Information,"](#) on page 113.

Viewing the Targets

You can view the targets of the events in the Real Time Event Table.

- 1 In an existing Real Time Event Table, right-click an event, then click **Target**.
- 2 Select one of the following options:
 - ◆ **ping**
 - ◆ **Name server lookup**
 - ◆ **Trace Route**
 - ◆ **Whois?**

The selected information is displayed.

Managing Columns

You can manage the columns displayed in the **Active Views** tab.

- 1 Access the **Active Views** tab.
For more information, see [Section 2.2.1, "Accessing the Active Views Tab,"](#) on page 19.
- 2 Click **Active Views > Event Real Time > Manage Columns**.
- 3 Use the **Add** and **Remove** buttons to move column titles between the **Available columns** list and the **Show these columns in this order** list.
- 4 Use the **Insert** button to insert an available column item into a specific location.
- 5 Use the up-arrow and down-arrow to arrange the order of the columns as you want them to display in the Real Time Event Table.
The top to bottom order of column titles in the Manage Column dialog box determines the left to right order of the columns in the Real Time Event Table.
- 6 To save your changes, click **OK**.
- 7 (Optional) If you want your columns to display the next time you open the Sentinel Control Center, click **File > Save Preferences**.

Taking a Snapshot of a Navigator Window

You must have user permission to take a snapshot.

A snapshot is useful to study events of interest because the Navigator refreshes automatically and the alert or alerts of interest scroll off the screen. Also, within a snapshot, you can sort by column.

- 1 Access the **Active Views** tab.

For more information, see [Section 2.2.1, “Accessing the Active Views Tab,”](#) on page 19.

- 2 In the menu, click **Active Views > Event Real Time > Snapshot**.

A Snapshot window opens and is added to the Snap Shots folder list under Active Views in the Navigator. The graphical display is not part of the snapshot.

To sort columns in a snapshot, click any column header once to sort by ascending value and twice to sort by descending value.

3 Searching Events

Sentinel provides an option to perform a search on events. You can search the local data in the primary storage (`/data` directory) or the stored data in compressed format at the configured secondary storage location. With the necessary configuration, you can also search system events generated by Sentinel, and view the raw data for each event. By default, events are returned in a reverse chronological order. This sort order relates to how the events are stored in the file system partitions.

You can also search Sentinel servers that are distributed across different geographic locations. For more information, see [“Configuring Data Federation”](#) in the *NetIQ Sentinel Administration Guide*.

- ◆ [Section 3.1, “Running an Event Search,”](#) on page 33
- ◆ [Section 3.2, “Viewing Search Results,”](#) on page 35
- ◆ [Section 3.3, “Refining Search Results,”](#) on page 38
- ◆ [Section 3.4, “Saving a Search Query,”](#) on page 39
- ◆ [Section 3.5, “Performing Event Operations,”](#) on page 44

3.1 Running an Event Search

By default, the search results include all events generated by the Sentinel system operations. These events are tagged with the `sentinel` tag. If no query is specified and you click **Search** for the first time after the Sentinel installation, the default search returns all events with severity 0 to 5. Otherwise, the Search feature reuses the last specified search query.

To search for a value in a specific field, use the ID of the event name, a colon, and the value. For example, to search for an authentication attempt to Sentinel by user2, use the following text in the search field:

```
evt>LoginUser AND sun:user2
```

An advanced search can narrow the search for a value to a specific event field. The advanced search criteria are based on the event IDs for each event field and the search logic for the index. Advanced searches can include the product name, severity, source IP, and the event type. For example:

- ◆ `pn:NMAS AND sev:5`

This searches for events with the product name NMAS and severity five.

- ◆ `sip:10.0.0.01 AND evt:"Set Password"`

This searches for the initiator IP address 10.0.0.1 and a “Set Password” event.

Multiple advanced search criteria can be combined by using various operators. The advanced search criteria syntax is modeled on the search criteria for the Apache Lucene open source package. For more information on building search criteria, see [Appendix A, “Search Query Syntax,”](#) on page 181.

Performing a Search

To perform a search:

- 1 Log in to the Sentinel Web interface:

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

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

- 2 In the **Reports and Searches** panel, click **New search**.

- 3 You can perform a search by using any of the following:

- ♦ **Search criteria:** Specify the search criteria in the **Search** field.
For information on creating search criteria, see [Appendix A, “Search Query Syntax,” on page 181](#).
- ♦ **Add Criteria:** Click **Add Criteria** and select from the criteria listed, click **Add**, and then click **Search**. You can select criteria from the list of criteria or filter the criteria based on recent criteria, tags, or filters.
 - ♦ **Show recent criteria:** Select a search criterion from the recent search history. The search history displays a maximum of 15 search expressions. Select the criteria, click **Show recent criteria**, and then click **Add**.
 - ♦ **Show Tags:** You can search events that have a particular tag. Click **Show Tags**, that lists the tags in the system. Select the tags, and then click **Add**.
 - ♦ **Show Filters:** You can reuse existing filters to perform a new search. Click **Show Filters** that lists the existing filters. Select the filter on which you want to perform the search, and then click **Add**.

You can combine multiple criteria, tags, or filters by using the **And** or **Or** condition.

- 4 (Optional) Select a time period for the search.

- ♦ The default is **Last 1 hour**.
- ♦ **Custom** allows you to select a start date and time and an end date and time for the query. The start date should be earlier than the end date, and the time is based on the machine's local time.
- ♦ **Whenever** searches all available data, without any time constraints.

- 5 (Optional) If you have administrator privileges, you can select other Sentinel servers for the search.

If you have data federation configured, you can perform a search on other Sentinel servers. For more information, see [“Configuring Data Federation”](#) in the *NetIQ Sentinel Administration Guide*.

- 6 Click **Search**.

The search results are displayed. For information on the search results, see [Section 3.2, “Viewing Search Results,” on page 35](#).

- 7 (Optional) Modify the search criteria by clicking **Edit Criteria**.

- 8 (Optional) Modify the search results by selecting the desired event fields in the search results

To add an AND or Or condition to the existing criteria, left-click the event field, select the required fields, and then specify the desired condition.

- 9 Click **Search**.

- 10 (Conditional) To save the search query, see [Section 3.4, “Saving a Search Query,” on page 39](#).





3.2 Viewing Search Results

Searches return a set of events. When results are sorted by relevance, only the top 50,000 events can be viewed. When results are sorted by time, all the events in the system are displayed.

Occasionally, the search engine might index events faster than they are inserted into the data directory. If you run a search that returns events that were not added in the data directory, you get a message indicating that some events match the search query, but they are not found in the `data` directory. If you run the search again later, the events are added to the `data` directory and the search is shown as successful.




NOTE: If time is not synchronized across your server, client, and event sources, you might get unexpected results from your search. This is especially a problem if searches are performed on time durations such as Custom, **Last 1 hour**, and **Last 24 hours** where display results are based on the time zone of the machine on which the search is performed.







The information in each event is grouped into the following categories:

Category	Icon	Description
General	No icon	Generic information about the event, such as severity, date, time, product name, and taxonomy.
Initiator		The source that caused the event to occur. The source can be a device, network port, etc.
Target		The object that is affected by the event. The object can be a file, database table, directory object, etc.
Observer		The service that observed the event activity.
Reporter		The service that reported the event activity.
Tags	No icon	Tags that the events are being tagged with.
Customer value	No icon	Fields set by the customer.
Retention period	No icon	Retention period of the event.






The initiator, target, and observer can be hosts, services, and accounts. In some cases, the initiator, target, and observer can be all the same, such as a user modifying this or her own account. In other cases, the initiator, target, and observer can be different, such as an intrusion detection system detecting a network attack. If an event field has no data, it is not displayed in the results.

Event fields are grouped according to the following categories:

Group	Icon	Description
Host		The initiator or target host information. For example, initiator host IP, target hostname, or target host ID.
User		The initiator or target user information. For example, the initiator username, initiator user department, target user ID, or target username.
Service		The initiator or target service information. For example, the target service name, target service component, or initiator service name.

Group	Icon	Description
Domain		Domain information of both the host and user. For example, the target host domain and initiator username.
IPCountry		The country information of the initiator and target trust. For example, the target host country.
Target trust		The target trust and target domain information of the event that was affected. The name can be a group, role, profile, etc.
Target data		The target data name and data container information. The data name is the name of the data object, such as a database table, directory object, or file that was affected by the event. The data container is the full path for data object.
Tenant name		The name of the tenant that owns the event data, applied to all the events in the inbound stream from a given Collector. The tenant name can be the name of the customer, division, department, etc.
Vulnerability		A flag that indicates whether Exploit Detection has matched this attack against known vulnerabilities in the target.

Each event type is represented by a specific icon. The following table lists the icons that represent the various types of events:

Icon	Type of Event
	Audit event
	Performance event
	Anomaly event
	Correlation event
	Unparsed event

You can view the search results in the summary view and in the detailed view. When you mouse over an event field, the information about the field is displayed.

- ◆ [Section 3.2.1, “Summary View,” on page 36](#)
- ◆ [Section 3.2.2, “Detailed View,” on page 36](#)

3.2.1 Summary View

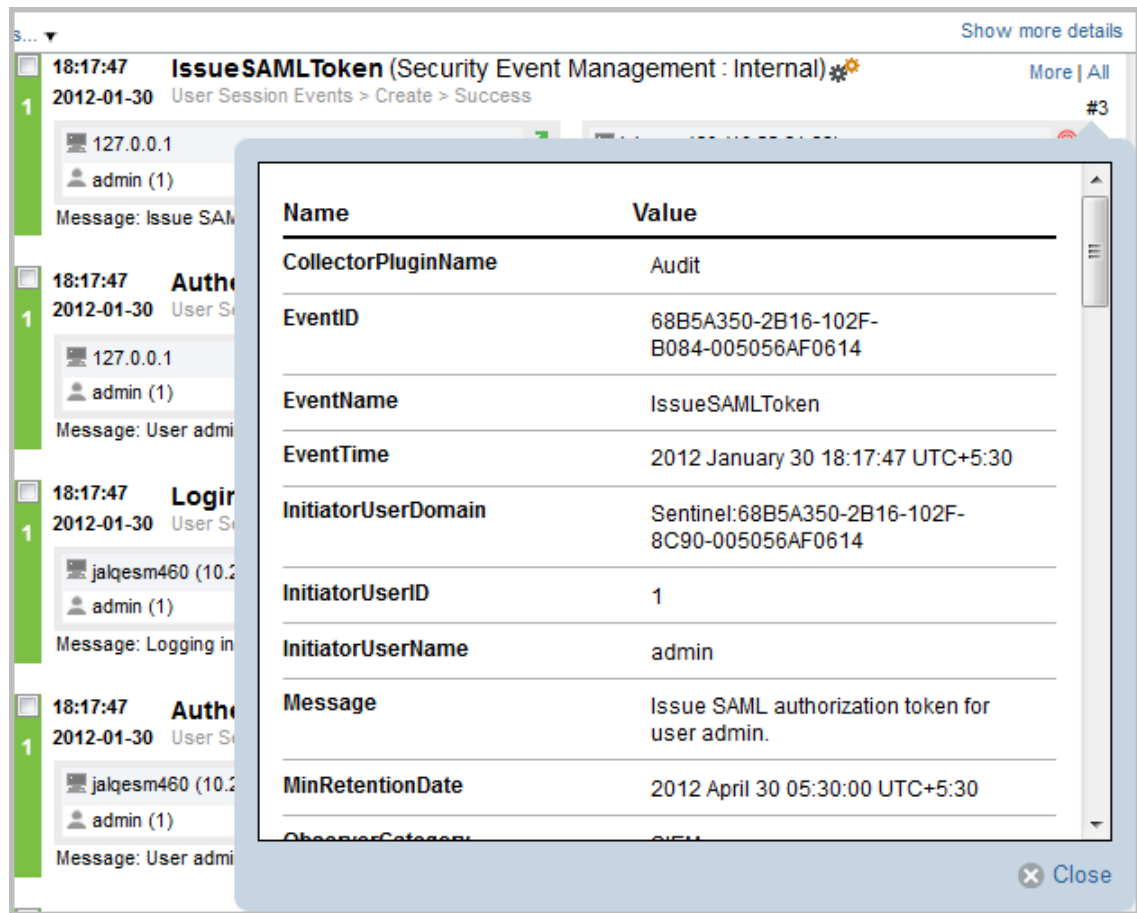
The Summary view of the search results displays the basic information about the event. The basic information includes severity, date, time, product name, taxonomy, and observer category for the event.

3.2.2 Detailed View

- 1 To view the report details, click the **More** link at the top right corner of the search results.

This displays details such as host/user domain information, IPCountry information, extended target fields like TargetTrust and TargetData, Observer and Reporter fields, customer set variables, default data retention duration information for any individual event, and the tags set for the event.

- 2 To view all the details of an event, click the **All** link.



- 3 To view details about all events, click the **Show more details** link at the top of the search results page.

You can expand or collapse the details for all events on a page by using the **Show more details** or **Show less details** link.

- 4 (Optional) Click the **get raw data** link to open a new **Raw Data** tab with event source hierarchy and event source fields populated, based on the information received from the event.

The **get raw data** link is available only for users in the administrator role.

If the search result is a system or an internal event, the **get raw data** link does not appear.

To verify and download the raw data files, see [“Verifying and Downloading Raw Data Files”](#) in the *NetIQ Sentinel Administration Guide*.

3.3 Refining Search Results

The search refinement panel can be used to narrow the search results by selecting one or more values for an event field. You can refine the results for one or more event fields.

The set of event fields that is displayed in the search refinement panel is configurable on a per-user basis.

For performance considerations, the maximum sample size used to calculate the event field value statistics is 50,000 events. The actual sample size is displayed in the field count label as `Field counts based on the first <sample-size> events where <sample-size>` is replaced by the actual sampling size.

To refine search results:

- 1 Log in to the Sentinel Web interface.

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

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

- 2 In the **Reports and Searches** panel, click **New Search**.
- 3 Specify the search criteria, then click **Search**.

For more information on how to run an event search, see [“Running an Event Search” on page 33](#).

- 4 Click **fields** in the REFINE section. The Select Event Fields window is displayed.
- 5 To refine the search, select the event fields from the available fields, then click **Save**.

The selected event fields are displayed in the **REFINE** panel.

A count at the right side of each event field displays the number of unique values that exist for that event field in the data directory. The calculation is based on the first 50,000 events found.

The event field selection is on a per-user basis. Each user can have a different set of selected event fields.

- 6 Click each event field to view the unique values for that event field.

For example, if the search results contain events that had severities 1, 2, 5, and 4, the event field is displayed as **Severity (4)**.

The top 10 unique values are initially displayed in the order of most frequent to least frequent.

The value next to the check box represents the unique value for that event field and the value at the far right represents the number of times the value appears in the search result.

If there are multiple unique values occurring the same number of times in a search, the values are sorted by the most recent occurrence of the value.

For example, if events of severity 1 and 4 occurred 34 times in the search results, and an event of severity 4 was logged most recently, the unique value 4 appears at the top of the list.

To display the unique values in the order of least frequent to most frequent, click **reverse**.

When there are more than 10 unique values, you can view and filter either the top 10 or the bottom 10 unique values. You cannot refine your search on both the conditions at the same time.

In the following scenarios, the number of events returned from a refined search is greater than the number of values listed for an event field:

- ◆ If the refinement performs a new search with additional terms intersected with the initial search string, such as by using an AND operator, the new search is run against all events in the system, including the result set from the initial search. If new events that came into the system match the refined search, they are shown in the resulting set and the event count is greater than the field value count.
- ◆ If there are more than 50,000 events, the event field statistics are calculated only on the first 50,000 events.

There could be an event field value that occurs 50 times in the first 50,000 events, but it could occur 1,000 times in all other stored events. In this scenario, the displayed value count is 50, but when the search is refined with this value it returns 1,000 events.

7 Click OK.

Selected event field values are listed under the event field in the **REFINE** panel.

The right panel displays the refined search results, which contain only the selected values.

8 Repeat Step 4 through Step 7 to further refine the search.

9 (Optional) Click **clear to clear the selected unique event field values from the **REFINE** panel and to return to the original search results.**

10 (Optional) Click **add to search to add the refined search values to the current search tab and to recalculate the search statistics.**

If you have already added the event field value to the current search tab, clicking **clear** does not return to the previous search results.

3.4 Saving a Search Query

You can save a search query, then repeat it as desired. To save a search query, you must first perform a search. When you are satisfied with the search results, you save the search query.

NOTE: You must have the necessary permission to access the specific options. For example, only users in the Report Administrator role can save the search query as a report template.

- ◆ [Section 3.4.1, “Saving a Search Query as a Search Template,” on page 39](#)
- ◆ [Section 3.4.2, “Saving a Search Query as a Filter,” on page 40](#)
- ◆ [Section 3.4.3, “Saving a Search Query as a Report Template,” on page 40](#)
- ◆ [Section 3.4.4, “Saving a Search Query as a Routing Rule,” on page 43](#)
- ◆ [Section 3.4.5, “Saving a Search Query as a Retention Policy,” on page 43](#)
- ◆ [Section 3.4.6, “Creating a Dashboard,” on page 44](#)

3.4.1 Saving a Search Query as a Search Template

1 Perform and refine a search until you are satisfied with the search results.

For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).

2 Click **Save as, and then click **Save search**.**

3 Specify a unique name for the search and provide an optional description.

- 4 Specify the following information in the **Default Parameters** section:

Data sources: Displays the number of servers that Sentinel will search for events. This option is useful if data federation is enabled. To select the data sources you want to search, click **selected data sources**, then select the data sources.

Email to: To e-mail the report template to others, specify the e-mail address. To send the report template to more than one person, specify multiple e-mail addresses separated by a comma.

Result limit: Specify the number of results to be stored in the search template. By default, 1000 results are stored in a report template.

- 5 Click **Save**.

3.4.2 Saving a Search Query as a Filter

- 1 Perform a search, and refine the search results as desired.

For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).

- 2 When you are satisfied with the search results, click **Save as**, then click **Save search as filter**.

- 3 Specify a unique name for the filter and an optional description.

- 4 In the drop-down list, select one of the following options to specify the access for this filter:

- ♦ **Private:** Allows you to make this filter private. Other users cannot view or access this filter.
- ♦ **Public:** Allows you to share this filter with all users.
- ♦ **Users in same role:** Allows you to share this filter with users who have the same role as yours.
- ♦ **Users in selected roles:** Allows you to share this filter with users in specific roles. If you select this option, a blank field is displayed where you can specify the roles. As you type the role name, a list of roles is displayed.

Select one or more roles.

NOTE: This option is available only for users in the administrator role.

- 5 Click **Save**.

The saved filter is listed in the Filters panel. For more information on filters, see [Chapter 4, “Configuring Filters,” on page 49](#).

3.4.3 Saving a Search Query as a Report Template

You can save the search query as a search report.

NOTE: You must have the Manage Reports permission to save the search query as a report template.

- 1 Perform a search, and refine the search results as desired.

For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).

- 2 When you are satisfied with the search results, click **Save as**, then click **Save search as report**.

- 3 Specify the following parameters:

Parameter	Description
Report name	Specify a unique name for the report. The name should not exceed 200 characters.
Based on	Select the base report from which you want to create the report. You can view a sample report by clicking the View Sample button.
Description	The description is automatically displayed based on the report that is selected and you can edit the description.
Criteria	Criteria is automatically populated based on the report selected and is not editable.
Additional criteria	Specify additional search criteria to the existing criteria. To build a new criteria on your own, click Edit Criteria . To build a new criteria from available system objects containing criteria, click Add Criteria . The criteria that you add here is appended to the existing criteria.
Data sources	Select the source machines on which the reports can be run by clicking the selected data sources link. You can select data sources only if your Sentinel is configured for data federation. For more information, see “ Configuring Data Federation ” in the <i>NetIQ Sentinel Administration Guide</i> .
Additional Criteria	Specify additional criteria to refine the results. The criteria that you specify here can be edited while scheduling the report. If you specify Criteria name , the name is displayed at the end of the report results. NOTE: This parameter is not available for all reports.
Time Zone	Specify the time zone with which you want to populate the report. When you schedule the report, the time zone that you specify here is displayed in the report data. For example, if the Time Zone is set to US/Pacific-New time, the report data displays the selected time zone. By default, it displays the time zone that is set in the client system. NOTE: This parameter is not available for all reports.

Parameter	Description
Date Range	<p>If the report includes time period parameters, choose the date range. All time periods are based on the local time for the browser. The From Date and the To Date automatically change to reflect the option you selected.</p> <ul style="list-style-type: none"> ◆ Current Day: Shows events from midnight of the current day until 11:59:00 PM of the current day. If the current time is 8:00:00 AM, the report shows 8 hours of data. ◆ Previous Day: Shows events from midnight yesterday until 11:59:00 PM yesterday. ◆ Week To Date: Shows events from midnight Sunday of the current week until the end of the selected day. ◆ Previous Week: Shows events for the last seven days. ◆ Month to Date: Shows events from midnight the first day of the current month until the end of the selected day. ◆ Previous Month: Shows events for a month, from midnight of the first day of the previous month until 11:59:00 PM. of the last day of the previous month. ◆ Custom Date Range: Shows events for a period whose start and end date are chosen. If you select Custom Date Range, set the start date (From Date) and the end date (To Date) for the report.
Group By	<p>Group the events according to specific event field by selecting the event field from the Group by drop-down list.</p> <p>NOTE: This parameter is not available for all reports.</p>
Language	<p>Choose the language in which the report labels and descriptions should be displayed. The possible values are English, French, German, Italian, Japanese, Traditional Chinese, Simplified Chinese, Spanish, or Portuguese.</p> <p>The default value is the language with which the current user logged in, if that language is supported by the report. If the report does not support the language, the report's default language (typically English) is used.</p> <p>The data in the report is displayed in the language that was originally used by the event source.</p>
Email to	<p>Specify an e-mail address in the Email to field. If you want to mail the report to more than one user, separate the e-mail addresses with a comma.</p>
Result limit	<p>Specify the number of results to be displayed or stored when you run or schedule the report. By default, 1000 results are stored.</p> <p>If you specify a value in Group By field, the result limit is based on grouping.</p>

4 Click **Save** to save the search as report definition.

You can see the saved report definition in the **Reports and Searches** panel in the Sentinel Web interface. To view the reports, see [“Viewing Events” on page 132](#).

3.4.4 Saving a Search Query as a Routing Rule

You must be in the administrator role to save the search query as a routing rule.

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,”](#) on page 33 and [Section 3.3, “Refining Search Results,”](#) on page 38.
- 2 When you are satisfied with the search results, click **Save as**, then click **Save search as routing rule**.
- 3 Specify a name for the rule.
- 4 (Conditional) To associate one or more tags to the events, click **Select tag**, select the desired tags, then click **Set**.
- 5 Select where you want to route the events to:
 - ♦ **All:** Events are routed to all Sentinel services, including Correlation and Security Intelligence.
 - ♦ **Event store only:** Events are sent directly to the event store, and are not displayed in Active Views and the search results page.
 - ♦ **None (drop):** Events are dropped or ignored, and are not sent to any Sentinel service.
- 6 Select one or more actions to be performed on each event that meets the search criteria. Click the plus and minus icons to add and remove actions.
- 7 Click **Save**.

3.4.5 Saving a Search Query as a Retention Policy

You must be in the administrator role to save the search query as a retention policy.

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,”](#) on page 33 and [Section 3.3, “Refining Search Results,”](#) on page 38.
- 2 When you are satisfied with the search results, click **Save as**, then click **Save search as retention policy**.
- 3 Specify a name for the retention policy.
- 4 In the **Keep at least** field, specify the minimum number of days to retain the events in the system. The value must be a valid positive integer.
- 5 (Optional) In the **Keep at most** field, 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 in primary storage.
- 6 Click **Save**.
The newly created policy is displayed in the data retention table. For more information on retention policies, see [“Configuring Data Retention Policies”](#) in the *NetIQ Sentinel Administration Guide*.

3.4.6 Creating a Dashboard

You must have the Manage and View Security Intelligence Dashboards permission to create a dashboard.

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,”](#) on page 33 and [Section 3.3, “Refining Search Results,”](#) on page 38.
- 2 When you are satisfied with the search results, click **Save as**, then click **Save search as dashboard**.
- 3 Specify the following information to create the dashboard:
 - ◆ **Name:** Specify a unique name for the dashboard.
 - ◆ **Classifier:** Select the classifier that determines the categories displayed in the dashboard. Click the **Info** link for information on each category.
 - ◆ **Data Retention Period:** Select how long the data for the dashboard is retained.
- 4 Click **Create dashboard** to create the dashboard.

The dashboard is displayed in a new browser tab. A new dashboard is empty because it has not had time to collect any data. For more information on dashboards, see [Chapter 7, “Analyzing Trends in Data,”](#) on page 93.

3.5 Performing Event Operations

You can use the events in the search results to perform various tasks as you view the search results.

- ◆ [Section 3.5.1, “Executing Actions,”](#) on page 44
- ◆ [Section 3.5.2, “Exporting the Search Results to a File,”](#) on page 45
- ◆ [Section 3.5.3, “Adding Events to an Incident,”](#) on page 45
- ◆ [Section 3.5.4, “Creating an Incident,”](#) on page 46
- ◆ [Section 3.5.5, “Adding Events to a Correlation Rule,”](#) on page 46
- ◆ [Section 3.5.6, “Creating a Correlation Rule by Using Events,”](#) on page 46
- ◆ [Section 3.5.7, “Viewing Identity Details of Events,”](#) on page 47
- ◆ [Section 3.5.8, “Viewing Advisor Report,”](#) on page 47
- ◆ [Section 3.5.9, “Viewing Asset Data,”](#) on page 47
- ◆ [Section 3.5.10, “Viewing Vulnerabilities,”](#) on page 48

3.5.1 Executing Actions

Only users in the following roles can execute actions on events:

- ◆ Administrator
- ◆ Incident Administrator
- ◆ Security Policy Administrator
- ◆ User

To execute actions on events:

- 1 Perform a search, and refine the search results as desired.

For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).

- 2 In the search results, select the events on which you want to execute actions.
- 3 Click **Event operations > Show action panel**.
- 4 In the **Event Actions** panel > **Actions** field, select the desired actions, then click **Execute**.
The results of the actions are displayed in the **Results** field. For more information on executing actions, see [Chapter 11, “Manually Performing Actions on Events,” on page 117](#).

3.5.2 Exporting the Search Results to a File

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events you want to export to a file.
- 3 Click **Event operations > Export to file**.
- 4 Specify the following information:
File Name: Specify a name for the file to which you want to export the search results.
Event Limit: Specify the maximum number of events to be saved. The event limit must be less than the number of events you selected and the maximum event limit is 200000.
All the search results are written into a `.csv` file. These files are then compressed into a `.zip` file for downloading.
- 5 (Optional) You can remove the event fields that you do not want to export to the file. Click **Choose Fields**, then clear the selections for the fields that you do not want to export to the file.
By default, the null fields are excluded and not exported to file.
- 6 Click **Export** to export the search result to a file.
A download file dialog box is displayed with an option to open or save the `.zip` file.
- 7 Select the desired option, then click **OK**.

3.5.3 Adding Events to an Incident

You must have the View or Create Incidents and Add Events to Incidents permission to add events to incidents.

For more information on Incidents, see [Chapter 16, “Configuring Incidents,” on page 143](#).

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events you want to add to an incident.
- 3 Click **Event Operations > Add to incident**.

NOTE: Ensure that incidents are available. If there are no incidents available, then you need to create one. For more information on creating incidents see [Section 3.5.4, “Creating an Incident,” on page 46](#).

- 4 Click **Search** to view all the available incidents.

- 5 (Optional) To view incidents based on categories, select a category from the **GroupBy** drop-down list.
- 6 Select the incident to which you want to add events.
- 7 Click **OK**.

3.5.4 Creating an Incident

You can create an incident from a group of events representing something of interest. For example, group together similar events or group together a set of different events that indicate a pattern of interest such as an attack.

You must have the View or Create Incidents and Add Events to Incidents permission to create incidents.

For more information on Incidents, see [Chapter 16, “Configuring Incidents,” on page 143](#).

To create an incident from events:

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events you want to add to an incident.
- 3 Click **Event operations > Create incident**.
- 4 Use the following information to create the incident:
 - Title:** Specify a title for the incident.
 - Description:** Specify a description of the incident.
 - Severity:** Select the severity of the incident from the drop-down list.
 - Priority:** Select the priority of the incident from the drop-down list.
 - Category:** Select the category of the incident from the drop-down list.
 - Responsible:** Select the user that is responsible to investigate and close the incident.
 - iTRAC:** Select an iTrac workflow to use to manage the incident.
- 5 Click **OK** to create the incident.

3.5.5 Adding Events to a Correlation Rule

You must have the Manage Correlation Engine and Rules permission to create a Correlation rule. For more information on creating a Correlation rule by using events, see [Section 5.4.6, “Creating Correlation Rules From Search Results,” on page 74](#).

3.5.6 Creating a Correlation Rule by Using Events

You must have the Manage Correlation Engine and Rules permission to create a Correlation rule. For more information on creating a Correlation rule by using events, see [Section 5.4.6, “Creating Correlation Rules From Search Results,” on page 74](#).

3.5.7 Viewing Identity Details of Events

If Sentinel is integrated with Identity Management systems, you can view the user identity details of events. You must have the View People Browser permission to view the Identity details.

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events for which you want to view the identity details.
- 3 Click **Event operations > Show identity details**.
- 4 Select whether you want to view the identity of the Initiator user, the Target user, or both.

For more information on identity details, see [Chapter 10, “Leveraging Identity Information,” on page 113](#).

3.5.8 Viewing Advisor Report

The following are the prerequisites to view the Advisor data:

- ♦ The Advisor feed must be up-to-date, processed, and loaded into the Sentinel database.
- ♦ The selected event must be from a product supported by Advisor and it must have the Vulnerability field value set to 1.

To view the Advisor data:

- 1 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.
- 2 In the search results, select the events for which you want to view the Advisor data.
- 3 Click **Event operations > View Advisor report**.
The Advisor report is displayed in a new tab.
For more information on Advisor, see [“Detecting Vulnerabilities and Exploits” in the *NetIQ Sentinel Administration Guide*](#).

3.5.9 Viewing Asset Data

You must have the View Asset Data permission to view the asset data of the selected events. You can view the asset information related to a machine or device from which you are receiving events. To view the asset data, you must run the asset management Collector and ensure that the asset data is being added to the Sentinel database.

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events for which you want to view the asset data.
- 3 Click **Event operations > View assets**.
The asset data is displayed in a new tab.
For more information on asset data, see [“Viewing Asset Data” on page 27](#).

3.5.10 Viewing Vulnerabilities

You must have the View asset vulnerability data permission to view the Vulnerability data. You can view the vulnerabilities of the selected destination systems. To view the Vulnerability data, you must run the Vulnerability Collector and ensure that the Vulnerability scan information is being added to the Sentinel database.

Vulnerabilities can be seen for the current time or for the event time.

- ♦ **View Vulnerabilities at current time:** This report queries the database for vulnerabilities that are active (effective) at the current date and time, and displays the relevant information.
- ♦ **View Vulnerabilities at time of event:** This report queries the database for vulnerabilities that were active (effective) at the date and time of the selected event, and displays the relevant events.

To view the Vulnerability report:

- 1 Perform a search, and refine the search results as desired.
For more information, see [Section 3.1, “Running an Event Search,” on page 33](#) and [Section 3.3, “Refining Search Results,” on page 38](#).
- 2 In the search results, select the events for which you want to view the Vulnerability data.
- 3 (Conditional) To view vulnerabilities at the current time, click **Event operations > View Vulnerabilities at current time**.
- 4 (Conditional) To view vulnerabilities at the time of the event, click **Event operations > View Vulnerabilities at time of event**.

For more information on the vulnerability data, see [“Viewing Vulnerabilities” on page 28](#).

4 Configuring Filters

- ◆ [Section 4.1, “Overview,” on page 49](#)
- ◆ [Section 4.2, “Introducing the Filters Interface,” on page 49](#)
- ◆ [Section 4.3, “Creating a Filter,” on page 52](#)
- ◆ [Section 4.4, “Sample Filters,” on page 54](#)
- ◆ [Section 4.5, “Viewing Events by Using Filters,” on page 56](#)
- ◆ [Section 4.6, “Managing Filters,” on page 56](#)

4.1 Overview

The Filters feature in Sentinel allows you to customize the event search and prevent data overload. This feature provides a Add or Edit criteria dialog that helps you build search queries ranging from simple to complex. You can save a search query as a filter and reuse it as required, so you can perform a search by selecting the filter rather than specifying the query manually every time.

You can reuse filters while using or configuring Sentinel features, such as:

- ◆ [Configuring Data Synchronization](#). For more information, see [“Configuring Data Synchronization”](#) in the *NetIQ Sentinel Administration Guide*.
- ◆ [Configuring a Data Retention policy](#). For more information, see [“Configuring Data Retention Policies”](#) in the *NetIQ Sentinel Administration Guide*.
- ◆ [Configuring the data visibility settings for a role](#). For more information, see [“Creating a Role”](#) in the *NetIQ Sentinel Administration Guide*.
- ◆ [Creating dashboards](#). For more information, see [Section 7.2, “Creating a Dashboard,” on page 96](#)
- ◆ [Configuring event routing rules](#). For more information, see [“Configuring Event Routing Rules”](#) in the *NetIQ Sentinel Administration Guide*.
- ◆ [Viewing real-time events in Active Views](#). For more information, see [Chapter 2, “Viewing Events in Real-Time,” on page 17](#).

Sentinel provides a list of filters by default. You can also create your own filters.

4.2 Introducing the Filters Interface

- ◆ [Section 4.2.1, “Filters Panel,” on page 50](#)
- ◆ [Section 4.2.2, “Filter Criteria,” on page 50](#)

4.2.1 Filters Panel

To access the Filters panel, click **Filters** in the navigation panel on the left of the Sentinel Web interface. The Filters panel lists the default filters, filters you create, and the filters that other users have shared with you.

The Filters panel includes the following:

- ♦ **Find Filters:** Allows you to search the specified filter. Specify the filter name, description of the filter, or keywords to search for a filter.
- ♦ **Create:** Launches the filter similar to search that allows you to specify the filter criteria.
- ♦ **My Filters:** Lists the default filters and the filters you created.
- ♦ **Shared Filters:** Lists the filters that other users have shared with you.

To view events based on filters, select the desired filter. The associated events are displayed in the search results panel.

4.2.2 Filter Criteria

You can specify the filter criteria by using either the Add Criteria or Edit Criteria dialog.

Add Criteria

Creating a filter is similar to performing a search. For more information, see [Section 3.1, “Running an Event Search,” on page 33](#).

The Add Criteria provides the predefined criteria list from which you can select the required criteria. You can filter the criteria based on recent criteria, tags, or filters.

- ♦ **Show only recent criteria:** Select a search criterion from the recent search history. The search history displays a maximum of 15 search expressions. Select the criteria, click **Show only recent criteria**, and then click **Add**.
- ♦ **Show only tags:** You can search events that have a particular tag. Click **Show only tags** to list the tags in the system. Select the tags, and then click **Add**.
- ♦ **Show only filters:** You can reuse existing filters to perform a new search. Click **Show only filters** to list the existing filters. Select the filter on which you want to perform the search, and then click **Add**.

You can combine multiple criteria, tags, or filters by using the **And** or **Or** condition. After adding the criteria, you can test the filter by clicking **Test Filter**.

Edit Criteria

The Edit Criteria provides a list of parameters required to build search criteria ranging from simple to complex. You can either select the parameters, or you can manually specify the search criteria.

For information on building search queries, see [Appendix A, “Search Query Syntax,” on page 181](#).

The Edit Criteria dialog box includes the following elements:

Table 4-1 Edit Criteria Dialog Box Elements

Element	Description
Criteria	<p>If you select Structured, this field displays the criteria formed by the parameters you select. You cannot modify or specify the filter criteria.</p> <p>If you select Free-form, you can manually specify the filter criteria.</p>
Structured	Allows you to select the various parameters to build the filter criteria.
Free-form	<p>Allows you to manually specify the filter criteria rather than selecting from the available parameters.</p> <p>The search criteria is based on the standard Lucene syntax with some Sentinel extensions. For information on creating a filter criteria (search query), see Appendix A, "Search Query Syntax," on page 181.</p> <p>If this option is selected, the following elements are not displayed:</p> <ul style="list-style-type: none">◆ Event fields◆ Criteria fields◆ Field details
Exclude system events	Select this option to exclude Sentinel internal events such as audit events and performance events from the search results.
Event fields	<p>Displays a categorized list of possible event fields you can add to the filter criteria. You can expand each category to display the set of fields in that category. If you know the name of the field you want, specify the name in the Search field. The event category list will adjust to present only matching fields.</p> <p>For more information on event fields, click Tips located at the top right of the Sentinel Web interface.</p>
Criteria fields	<p>Lists a set of overlay criteria that you can use on top of per-field searches. The following fields are displayed by default:</p> <ul style="list-style-type: none">◆ All data: Performs a search across all event fields. For more information, see Section A.1.4, "The Default Search Field," on page 183 in Appendix A, "Search Query Syntax," on page 181.◆ Tags: Events can be tagged in various ways to help identify relationships between events. Queries that include a "Tags" search will look at the event tags (rv145) for matches.◆ Taxonomy: Events are also classified using a number of taxonomic categories for the action, outcome, and so on. Queries that include a "Taxonomy" search will search for specific classes of events. For more information on taxonomy, see Sentinel Taxonomy.

Element	Description
Field details	<p>The fields in this section vary depending on the event or criteria fields you select. For example:</p> <ul style="list-style-type: none"> ◆ For tokenized fields, you can specify the words that you want to include or exclude in the filter criteria. For information on the tokenized and non-tokenized fields, click Tips located at the top right of the Sentinel Web interface. ◆ For non-tokenized fields, you can specify a value or a range of values. ◆ For taxonomy fields, specific taxonomy options are displayed. ◆ For date attributes, a date-time calendar is displayed as you type the date. You can select a date. ◆ For fields that contain internal Sentinel UUIDs, such as the CollectorID field, the corresponding Sentinel object names are displayed and can be selected.
Condition: AND OR	Allows you to specify the AND or OR condition between the criteria fields. These options are available when you add additional event criteria to the criteria fields.
Cancel	Allows you to cancel the filter creation process.
Search	Runs a search to test the filter before saving it.

4.3 Creating a Filter

Filter expressions are simple math expressions and simple evaluations. Filters work on selection sets by matching events against the specified criteria. If the match is TRUE, the event is displayed in Active Views or search results, or passed to other functions. If the match is FALSE, the event is blocked.

For example, consider a search query that is written as follows:

```
(sip:"10.0.0.1")
```

Events whose source IP address is 10.0.0.1 are included in the filter.

You must use the event field ID to represent an event name. Click the **Tips** link on the top right of the Sentinel Web interface for a list of event field names and their IDs.

For more information building search queries, see [Appendix A, “Search Query Syntax,” on page 181](#).

While creating a filter, you can specify whether you want to share a filter with other users. You must have the Share Search Filters permission to share filters with everyone or with users in the same role as yours. If you are a user in the administrator role, you can share filters with users in a different role.

You can create filters either by using the Add or Edit Criteria dialog or by using the **Save as** icon in the Search panel.


- ◆ [Section 4.3.1, “Creating a Filter by Using the Edit Criteria Dialog,” on page 52](#)
- ◆ [Section 4.3.2, “Creating a Filter by Using a Search Query,” on page 53](#)

4.3.1 Creating a Filter by Using the Edit Criteria Dialog

- 1 Log in to 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.

- 2 In the navigation panel, click **Filters > Create**.
 - 3 Select one of the following methods to create a search criteria:
 - ♦ To build the search criteria by selecting parameters, make sure that **Structured** is selected, select the parameters, then continue with [Step 4](#).
For information on these parameters, see [Table 4-1, “Edit Criteria Dialog Box Elements,” on page 51](#).
 - ♦ To manually specify the search query rather than selecting the listed parameters, select **Free-form**. In the **Criteria** field, specify the search query, then continue with [Step 4](#).
For information on creating a search query, see [Appendix A, “Search Query Syntax,” on page 181](#).
 - 4 (Conditional) If you do not want to include Sentinel internal events in the search, select **Exclude system events**.
 - 5 Click **Search** to search events according to the specified filter criteria.
By default, the search is performed on events that were generated within the last 1 hour.
 - 6 Review the search results to verify that the filter is retrieving the expected events.
 - 7 (Optional) You can modify the search query by selecting one or more event field values from the search results, or you can click **Edit search filter**, then make necessary changes.
 - 8 When you are satisfied with the search results, click , then click **Save as new filter**.
 - 9 Specify a name for the filter and an optional description.
 - 10 In the drop-down list, select one of the following options to specify the access for this filter:
 - ♦ **Private**: Allows you to make this filter private. Other users cannot view or access this filter.
 - ♦ **Public**: Allows you to share this filter with all users.
 - ♦ **Users in same role**: Allows you to share this filter with users who have the same role as yours.
 - ♦ **Users in selected roles**: Allows you to share this filter with users in specific roles. If you select this option, a blank field is displayed where you can specify the roles. As you type the role name, a list of roles is displayed.
Select one or more roles.
-
- NOTE:** This option is available only for users in the administrator role.
-

- 11 Click **Save**.

4.3.2 Creating a Filter by Using a Search Query

You can save a search query as a filter and use this filter to perform searches when required rather than specifying the search query again. For more information on creating a filter by using a search query, see [Section 3.4.2, “Saving a Search Query as a Filter,” on page 40](#).

4.4 Sample Filters

This section lists a few examples on how you can create filters.

- ◆ [Section 4.4.1, “View Events of Severity 3 to 5 from a System in China,” on page 54](#)
- ◆ [Section 4.4.2, “Determine if User “Bob Smith” Tried to Log In after His Account was Disabled,” on page 54](#)
- ◆ [Section 4.4.3, “View Events from Two Subnets and Share the Filter with Network Administrators,” on page 55](#)
- ◆ [Section 4.4.4, “Find all Events that Include the Words “database” and “service,” and exclude “test,” on page 55](#)

4.4.1 View Events of Severity 3 to 5 from a System in China

- ◆ In the **Edit Criteria > Event fields**, select **SourceHostCountry**.

For more information on the Filter Criteria, see [Section 4.2.2, “Filter Criteria,” on page 50](#).

- ◆ The name should match any string that contains the name “China.” For example, “ChinaBeijing.” Specify `china*` in the **Value** field.
- ◆ The severity of the events must be 3 to 5:
 - ◆ In **Event fields**, select **Severity**.
 - ◆ In the **Values that range from** field, specify 3 TO 5.

NOTE: If you are familiar with the search query syntax, you can directly specify the query in the **Criteria** field as follows:

```
(rv29:china*) AND (sev:[3 TO 5])
```

For more information on the search query syntax, see [Appendix A, “Search Query Syntax,” on page 181](#).

Click **Search** to view events that match the specified criteria.

4.4.2 Determine if User “Bob Smith” Tried to Log In after His Account was Disabled

- ◆ In the **Edit Criteria > Event fields**, select the following:
 - ◆ **InitiatorUserName**
 - ◆ **TargetUserName**
 - ◆ **EffectiveUserName**

For more information, see [Section 4.2.2, “Filter Criteria,” on page 50](#).

- ◆ Select the **OR** condition.
- ◆ Specify "Bob Smith" in the **Value** field.
- ◆ To determine if the user has logged in, or tried to log in, select **Taxonomy** in **Criteria fields**.

NOTE: You can also select the appropriate event fields if you are familiar with the values to be specified for the event fields. Taxonomy is a classification of events where events of similar type are grouped together. It helps you search events based on the taxonomy classification rather than you specifying the specific event names and their values.

- ◆ In the **Field details**, select the following:
 - ◆ From the **Class** drop-down list, select **User Session Events**.
 - ◆ From the **Identifier** drop-down list, select **Create**.
 - ◆ For **Outcome**, select **Success**, then select **Failure**.


NOTE: If you are familiar with the search query syntax, you can directly specify the query in the **Criteria** field as follows:

```
(xdasclass:2 AND xdasid:0 AND (xdasoutcome:0 OR xdasoutcome:1)) AND (iufname:"Bob Smith")
```

For more information on taxonomy, see [Sentinel Taxonomy](#).

Click **Search** to view the events that match the specified criteria.

4.4.3 View Events from Two Subnets and Share the Filter with Network Administrators

- ◆ Select subnets:
 - ◆ In **Edit criteria > Event fields**, select **SourceIP**.
 - ◆ In **Field details > Value**, specify the subnet, for example, 172.17.0.0/16.
 - ◆ Repeat the above two steps to specify another subnet.
- ◆ The events must be from either of the subnets. Therefore, select **OR** as the condition.
- ◆ Click **Search** to view events that match the specified criteria.
- ◆ The filter must be shared with network administrators:
 - ◆ In the search results panel, click , then click **Save as new filter**.
 - ◆ Specify an intuitive name and an optional description.
 - ◆ From the drop-down list, select **Share with roles**, then select **Network Administrator**.
- ◆ Click **Save**.

4.4.4 Find all Events that Include the Words “database” and “service,” and exclude “test”

- ◆ In **Edit Criteria > Criteria fields**, select **All data**.
- ◆ You want to find events that include words “database” and “service,” and exclude “test.” Therefore, in **Field details**, specify the following:
 - ◆ In the **All of these words** field, specify `database service`.
 - ◆ In the **Exclude these words** field, specify `test`.


NOTE: If you are familiar with the search query syntax, you can directly specify the query in the **Criteria** field as follows:

```
_data:(database AND service) NOT _data:test
```

The `_data` field allows you to search for words that might appear in any event field. For more information, see [“The Default Search Field”](#) in [Appendix A, “Search Query Syntax,”](#) on page 181.

Click **Search** to view the events that match the specified criteria.

4.5 Viewing Events by Using Filters


You can use filters to view events either by selecting the desired filter in the **Filters** panel or by using the **Filter**  icon in the search results panel. For more information, see [Chapter 3, “Searching Events,”](#) on page 33.

4.6 Managing Filters

You can edit and delete only the filters that you created. The default filters and the filters that other users have shared with you cannot be edited or deleted.

- ◆ [Section 4.6.1, “Editing a Filter,”](#) on page 56
- ◆ [Section 4.6.2, “Deleting a Filter,”](#) on page 56

4.6.1 Editing a Filter

- 1 In the Sentinel Web interface, select **Filters**, then select the filter that you want to edit.
The filter criteria are displayed in the search results panel.
- 2 (Conditional) To edit the filter name, description, or the filter type, click **Edit**.
- 3 Make the necessary changes, then click **OK**.
- 4 (Conditional) To edit the filter criteria, click **Edit search filter**.
For more information on the criteria you can select, see [Section 4.2.2, “Filter Criteria,”](#) on page 50.
- 5 Make the necessary changes, then click **Search**.
- 6 Click , then select **Save <filter_name> filter**.

4.6.2 Deleting a Filter

- 1 In the Sentinel Web interface, select **Filters > My filters**, then select the filter that you want to delete.
- 2 Click .
- 3 Click **Delete** to confirm deletion.

5 Correlating Event Data

A single event viewed in the system might not necessarily draw your attention. But when you correlate a set of similar or comparable events in a given period, you might identify a potential problem. Sentinel helps you correlate events by using the rules you create and deploy in the Correlation Engine, so you can take appropriate action to mitigate any problems.

- ◆ [Section 5.1, “Overview,” on page 57](#)
- ◆ [Section 5.2, “Accessing the Correlation User Interface,” on page 62](#)
- ◆ [Section 5.3, “Understanding the Correlation Interface,” on page 63](#)
- ◆ [Section 5.4, “Creating Correlation Rules,” on page 68](#)
- ◆ [Section 5.5, “Associating Actions to a Rule,” on page 75](#)
- ◆ [Section 5.6, “Testing a Correlation Rule,” on page 75](#)
- ◆ [Section 5.7, “Sample Correlation Rules,” on page 76](#)
- ◆ [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#)
- ◆ [Section 5.9, “Viewing Correlated Events,” on page 78](#)
- ◆ [Section 5.10, “Customizing Correlated Event,” on page 79](#)
- ◆ [Section 5.11, “Managing Correlation Rules,” on page 79](#)
- ◆ [Section 5.12, “Managing the Correlation Engine,” on page 81](#)

5.1 Overview


Correlation adds intelligence to security event management by automating analysis of the incoming event stream to find patterns of interest. Correlation allows you to define rules that identify critical threats and complex attack patterns so that you can prioritize events and initiate effective incident management and response.

- ◆ [Section 5.1.1, “How Correlation Works,” on page 58](#)
- ◆ [Section 5.1.2, “Correlation Rules,” on page 58](#)
- ◆ [Section 5.1.3, “Correlation Engine,” on page 62](#)

5.1.1 How Correlation Works

The following illustration shows how Correlation works:

Figure 5-1 Correlation Workflow



1. A user creates a Correlation rule.
2. The user associates one or more actions to the Correlation rule.
3. The user deploys the rule in the Correlation Engine.
4. The Correlation Engine processes events from the real-time event stream to determine whether they should trigger any of the active rules to fire the associated actions.
5. If events match the rule criteria, correlation events are generated and associated actions are executed.

Sentinel's correlation is near real-time and depends on the time stamp of the individual events. When an event arrives at the Correlation Engine, the engine reorders the events in a buffer based on the event time stamp (dt) field so that the events are evaluated in time order. This is done partly to evaluate sequence rules in which the rule only fires if events occur in a specific order.

The buffer is 30 seconds long, so if the event time stamp (dt) is more than 30 seconds older than the Collector Manager time stamp, the event is not evaluated. To minimize false time differences, you must use an NTP (Network Time Protocol) server to synchronize the time settings on the relevant machines. For more information, see "Configuring Time" in the *NetIQ Sentinel Installation and Configuration Guide*.

5.1.2 Correlation Rules

Correlation rules define a pattern of events that should trigger a rule. You can create rules that range from simple to extremely complex. For example:

- ♦ High severity event from a finance server
- ♦ High severity event from any server brought online in the past 10 days
- ♦ Five failed logins in 2 minutes

- ◆ Five failed logins to the same server from the same username in 2 minutes
- ◆ Intrusion detection event targeting a server, followed by an attempted login to root originating from the same server within 60 seconds

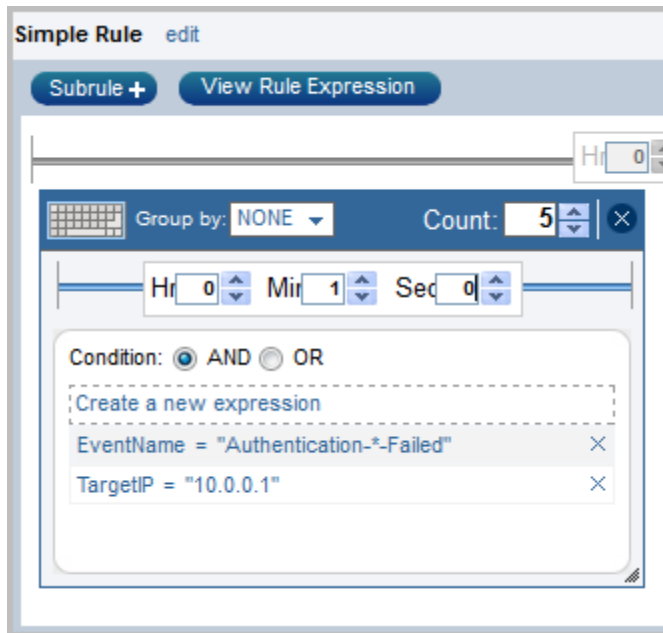
A rule can have one or more subrules:

- ◆ “Simple Rule” on page 59
- ◆ “Sequence Rule” on page 59
- ◆ “Composite Rule” on page 60
- ◆ “Free-form Rule” on page 61

Simple Rule

A simple rule has just one subrule. You can specify additional criteria if you want the rule to fire when all or any of the specified criteria are met. You can also specify the number of times the event should occur for the rule to fire. For example, to monitor a situation with five failed logins within a minute on a finance server, you can define a simple rule, as shown in the following figure:

Figure 5-2 Simple Rule

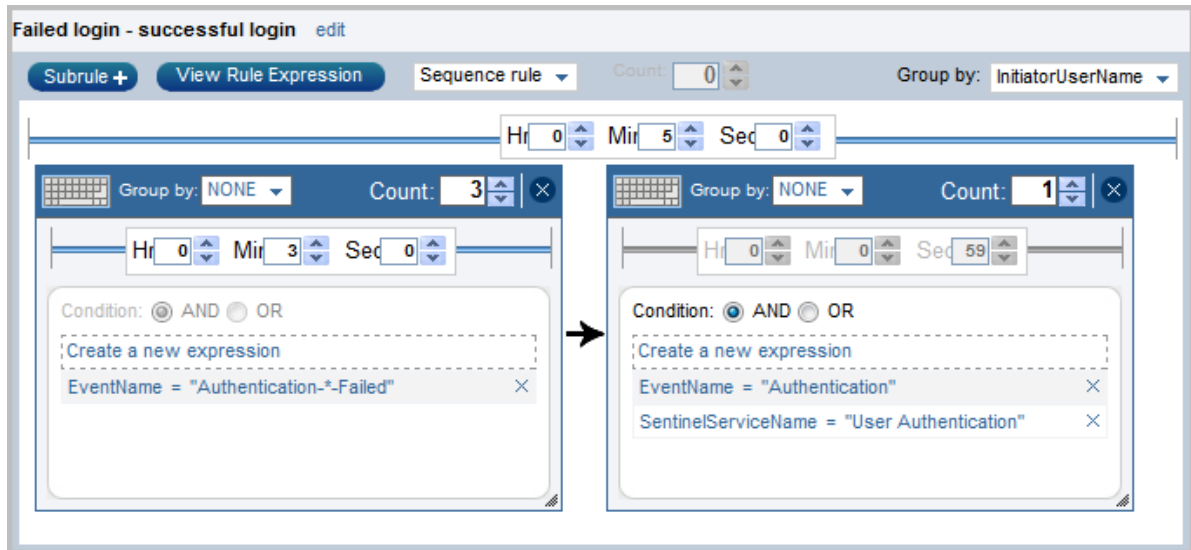


For information on creating a simple rule, see [Section 5.4.2, “Creating a Simple Rule,”](#) on page 70.

Sequence Rule

A sequence rule has two or more subrules that fire in a sequence. You can use a sequence rule when you want the rule to fire if its subrules meet the specified criteria in the specified sequence within the defined time frame. For example, to monitor a situation where there has been a successful login after three failed logins by the same user within five minutes, you can define the rule as shown in the following figure:

Figure 5-3 Sequence Rule



A sequence rule can also fire even when a single event can satisfy multiple or all of the subrules within the time frame and the times at which they were satisfied are in ascending order. For information about creating a sequence rule, see [Section 5.4.3, “Creating a Sequence Rule,”](#) on page 71.

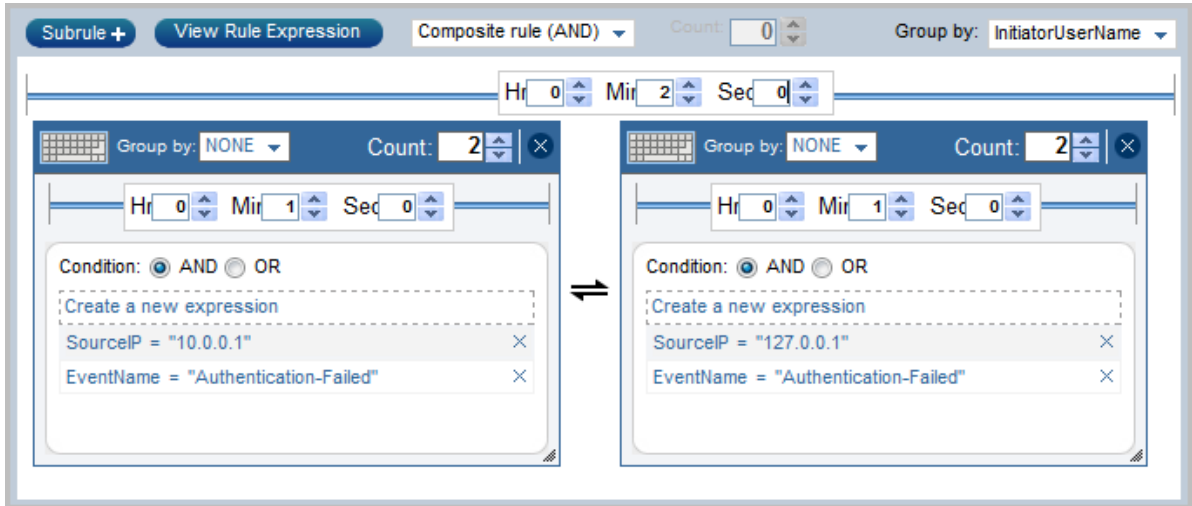
Composite Rule

A composite rule has two or more subrules that fire according to the criteria you define. There are two types of composite rules:

- ◆ **Composite (AND):** Indicates that all subrules must fire.
- ◆ **Composite (OR):** Indicates that a specified number of subrules must fire.

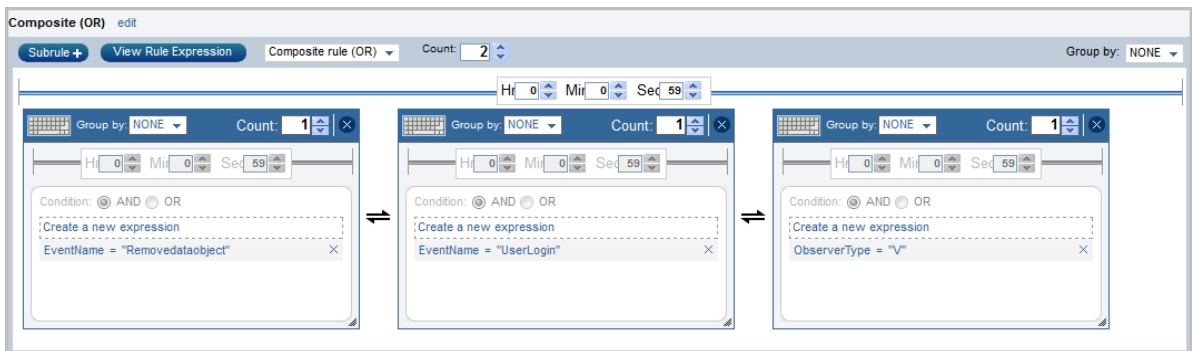
For example, to monitor a situation where there have been failed logins on a finance server and a database server within two minutes, you can create a composite (AND) rule, as shown in the following figure:

Figure 5-4 Composite (AND) Rule



Similarly, for example, if you have three or more subrules and you want the rule to fire if a maximum of two subrules meet the specified criteria, you can create a composite (OR) rule, as shown in the following figure:

Figure 5-5 Composite (OR) Rule



A composite rule can also fire even when a single event can satisfy multiple or all of the subrules. For information on creating a composite rule, see [Section 5.4.4, “Creating a Composite Rule,” on page 72](#).

Free-form Rule

If you are familiar with the rule expression syntax, you can create Correlation rules by manually specifying the rule expression. You can use free-form rules to create complex rules by using additional operators such as Window, Intersection, and Union.

For information on the rule expression syntax, see [Appendix B, “Correlation Rule Expression Syntax,” on page 193](#).

For more information on creating a free-form rule, see [Section 5.4.5, “Creating a Free-Form Rule,” on page 74](#).

5.1.3 Correlation Engine

To monitor events according to the Correlation rules, you must deploy the rules in the Correlation Engine. When an event occurs that satisfies the rule criteria, the Correlation Engine generates a correlation event describing the pattern.

NOTE: Events that are sent directly to the event store or dropped by event routing rules are not processed by the Correlation Engine.

The Sentinel Correlation Engine provides specific advantages over database-centric Correlation Engines.

- ♦ By relying on in-memory processing rather than database inserts and reads, the Correlation Engine performs during high steady-state volumes as well as during event spikes when under attack, which is the time when correlation performance is most critical.
- ♦ The correlation volume does not slow down other system components, so the user interface remains responsive, especially with high event volumes.
- ♦ The Correlation Engine can add events to incidents after an incident has been created.
- ♦ You can deploy multiple Correlation Engines, each on its own server, without the need to replicate configurations or add databases. The Correlation Engine is built with a pluggable framework that allows the addition of new Correlation Engines. Independent scaling of components provides cost-effective scalability and performance.

NOTE: You cannot install more than one Correlation Engine on a single system. You can install additional Correlation Engines on remote systems, and then connect them to the Sentinel server.

For more information about installing the Correlation Engine, see [“Installing Collector Managers and Correlation Engines”](#) in the *NetIQ Sentinel Installation and Configuration Guide*.

5.2 Accessing the Correlation User Interface

NOTE: You must have the Manage Correlation Engine and Rules permission to access the Correlation interface.

- 1 Log in to the Sentinel Web interface as a user with the Manage Correlation Engine and Rules permission.

`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 navigation panel, click **Correlation**.

The Correlation panel is displayed. For more information, see [Section 5.3.1, “Correlation Panel,” on page 63](#).

You can also access the Correlation user interface from the search results panel. For information on creating Correlation rules from the search results panel, see [Section 5.4.6, “Creating Correlation Rules From Search Results,” on page 74](#).

5.3 Understanding the Correlation Interface

The Correlation interface includes the following:

- ◆ [Section 5.3.1, “Correlation Panel,” on page 63](#)
- ◆ [Section 5.3.2, “Correlation Rule Builder,” on page 64](#)

5.3.1 Correlation Panel

The Correlation panel lists the rules and the Correlation Engines installed on your system.

The Correlation panel includes the following options:

- ◆ [“Search Field” on page 63](#)
- ◆ [“Create Link” on page 63](#)
- ◆ [“More Link” on page 63](#)
- ◆ [“Rules” on page 63](#)
- ◆ [“Engines” on page 64](#)

Search Field

The **Search**  field allows you to search for the specified rule in your system.

Create Link



The **Create** link launches the Correlation Rule Builder that helps you to create correlation rules.

More Link




The **More** link allows you to select multiple rules to delete them.

Rules

The Rules section lists all the available rules in the system. The icon next to the rule indicates the status of the rule:




- ◆ **Enabled** : The rule is deployed in the Correlation Engine and is enabled to process events.
- ◆ **Disabled** : The rule is deployed in the Correlation Engine, but the rule is disabled and is not processing events.
- ◆ **No icon**: The rule is not deployed in the Correlation Engine.

To view the details of any rule, click the rule. When you select or click any rule, the following icons are displayed:

- ◆ **Search** : Searches for the events that meet the rule criteria. The events are displayed in the search results panel.
- ◆ **Edit** : Allows you to edit the rule.
- ◆ **Delete** : Allows you to delete the rule.

Engines

The Engines section lists the Correlation Engines installed in the system. The icon next to the Correlation Engine indicates the status of the Correlation Engine.

- ♦ **Start** : The Correlation Engine is started and is processing the deployed correlation rules.
- ♦ **Stop** : The Correlation Engine is stopped. When the Correlation Engine is stopped, all in-memory data is preserved and no new correlation events are generated.
- ♦ **Offline** : The remote Correlation Engine is off.

Click any engine to view the details, such as the rules deployed in this engine and information about the engine. For more information, see [Section 5.12.1, “Using the Correlation Engine Dashboard,” on page 81](#).

5.3.2 Correlation Rule Builder

The Correlation Rule Builder helps you to create correlation rules and includes the following:

- ♦ [“Command Buttons” on page 64](#)
- ♦ [“Rule Builder Elements” on page 64](#)
- ♦ [“Subrule Window” on page 65](#)
- ♦ [“Expression Builder” on page 67](#)
- ♦ [“Actions Panel” on page 68](#)

Command Buttons

The following command buttons are available:

- ♦ **Subrule**: Adds a subrule window in the rule builder. You can add additional subrules to create a sequence or composite rule.
- ♦ **View Rule Expression/Hide Rule Expression**: Displays or hides the expression of the rule. This is a toggle button.
- ♦ **Save Rule**: Saves the rule in the Sentinel database.
- ♦ **Save As**: Allows you to save the rule with another name.
- ♦ **Test Rule**: Tests the rule against the events in your system. For more information, see [Section 5.6, “Testing a Correlation Rule,” on page 75](#).

Rule Builder Elements

Table 5-1 Common Rule Builder Elements



Element	Description	Action
edit	Allows you to edit the rule name	Click the edit link.


Element	Description	Action
Rule type	<p>Lists the types of rules:</p> <ul style="list-style-type: none"> ◆ Sequence ◆ Composite (AND) ◆ Composite (OR) <p>This list is displayed only if there is more than one subrule.</p>	Select an appropriate rule type for the rule you want to create.
Count	<p>This field is enabled only for composite (OR) rules and if there are more than 2 subrules.</p> <p>It indicates the maximum number of subrules that should meet the specified criteria for the rule to fire. For example, if you have 5 subrules and you specify the Count as 3, the rule fires if one, two, or three subrules meet the specified criteria.</p>	<p>Specify a number or use the up/down arrow keys to select a number.</p> <p>NOTE: The value should always be less than the number of subrules in the rule.</p>
Group by	<p>Lists the attributes you can use to group the correlation events.</p> <p>The Group by list is enabled if there are two or more subrules.</p>	Select one or more attributes. For example, to group events by username, select <code>initusername</code> .
Time frame Hr: Min: Sec	Indicates the time within which the specified criteria in the subrules should be satisfied for the rule to fire.	Specify the time in hours, minutes, or seconds.

Subrule Window

The subrule window allows you to specify the expressions (criteria) for the rule and lists the various expressions that you have created for a subrule.

Table 5-2 Subrule Window Elements


Element	Description	Action
Toggle icons 	Toggles between a structured rule and a free-form rule.	Click the icon to toggle to the free-form or structured view.
Group by	<p>Lists the attributes you can use to group the correlation events.</p> <p>The Group by list is enabled only if the Count is greater than 1.</p>	Select one or more attributes. For example, to group events by username, select initusername .
Count	Indicates the number of times the expressions must meet the specified criteria for the subrule/rule to fire.	Specify a number or use the up/down arrow keys to select the number.
Close icon 	Closes the subrule window.	Click the icon to close the subrule window.

Element	Description	Action
Time frame Hr: Min: Sec	Indicates the time within which the specified criteria in the subrule should be satisfied for the rule to fire.	Specify the time in hours, minutes, or seconds. For example, if you want the rule to fire within 2 minutes, specify 2 in the Min field.
Condition: AND OR	Determines whether the subrule should fire when all or any of the conditions in the expression are met, according to the selection. These option buttons are enabled only if there are two or more expressions in a subrule.	Select one of the conditions.
Create a new expression	Allows you to create a new expression. Displays the Expression Builder. For more information, see “Expression Builder” on page 67 .	Click the link to create a new expression.
Delete expression 	Deletes the expression.	Click the icon to delete the expression.

Expression Builder

The Expression Builder allows you to select various parameters required to create an expression for the rule. The various parameters include Attributes, Operator, and Value. These parameters are interdependent, and changing one of them affects the validity of others.

Figure 5-6 Expression Builder



Attribute: Displays a categorized list of possible event fields that can be used to create a Correlation rule. Each category can be expanded to display the set of fields in that category. If you know the name of the field you want, specify the name in the **Search** field. The event category list adjusts to present only matching fields.

For information on the various event fields, click **Tips** located at the top right corner.

Operator: Lists the various operators. The list varies depending on the selected attribute type. For example:

- ♦ For all attributes, the =, <, >, !=, <=, >=, inlist, isnull, not inlist, and not isnull operators are available.
- ♦ For string attributes, the match regex operator is available.
- ♦ For IP attributes, the match subnet operator is available.
- ♦ For tag attributes, the contains operator is available.

For more information on using the operators, see “[Filter Operation](#)” in [Appendix B, “Correlation Rule Expression Syntax,”](#) on page 193.

Value: This field varies, depending on the attribute and operator. For example:


- ◆ For the isnull and not isnull operators, no value can be chosen.
- ◆ For the inlist and not inlist operators, the available dynamic lists are displayed. You can also create a new dynamic list if necessary. For more information on dynamic lists, see [Section 9.1, “Creating a Dynamic List,” on page 109](#).
- ◆ For the Severity attribute, the severity list is displayed.
- ◆ For date attributes, a date-time calendar is displayed.
- ◆ For xdas taxonomy attributes, the taxonomy builder is displayed.
- ◆ For numeric attributes, only numbers are accepted.
- ◆ For the Sensor type attribute, the list of sensor types is displayed.
- ◆ For the Tags (rv145) attribute, the list of available tags is displayed.
- ◆ For Collector fields, a list of Collectors is displayed.

You can also select one or more event attributes as the value by using the **Show Attributes** option.


Actions Panel

The Actions panel lists the actions associated with the rule, allows you to associate actions to the rule, and allows you to define when the action should execute.

The Actions panel includes the following icons:

- ◆ **Action execution criteria** : Allows you to specify when the rule should initiate the action. When you click this icon, the Action Execution Criteria dialog box is displayed:

You can select one of the following options:

- ◆ **Perform actions every time the rule fires:** The action executes each time the rule criteria are met.
- ◆ **Perform actions at most every:** The action executes at most every specified time interval. By default, this option is selected and the time interval is set to 1 hour. This is to ensure that the rule does not fire continually and overutilize resources.
- ◆ **Add action** : Allows you to associate actions to the rule. When you click this icon, a list of actions is displayed. Select one or more actions that you want to associate to the rule.

For more information on associating actions to a rule, see [Section 5.5, “Associating Actions to a Rule,” on page 75](#).

5.4 Creating Correlation Rules

The procedure to create various types of Correlation rules is the same for all rule types, except for a few steps that are specific to each rule type. Events are evaluated by rules in the specified order until a match is made, so you should order subrules accordingly. More narrowly defined subrules and more important subrules should be placed at the beginning of the list. When creating correlation rules, you can also define what you want to see in the correlated event.

- ◆ [Section 5.4.1, “Understanding the Correlated Event,” on page 69](#)
- ◆ [Section 5.4.2, “Creating a Simple Rule,” on page 70](#)
- ◆ [Section 5.4.3, “Creating a Sequence Rule,” on page 71](#)
- ◆ [Section 5.4.4, “Creating a Composite Rule,” on page 72](#)

- ◆ [Section 5.4.5, “Creating a Free-Form Rule,” on page 74](#)
- ◆ [Section 5.4.6, “Creating Correlation Rules From Search Results,” on page 74](#)

5.4.1 Understanding the Correlated Event

By default, correlated events display the event name and message as the rule name and description respectively.

The correlated event field values depend on the number of events that triggered the correlated event and whether any of the event fields are grouped.

- ◆ In case of a single event triggering the correlated event, all the event field values are copied to the correlated event. If one or more event fields are grouped, only the grouped event fields are copied to the correlated event.
- ◆ In case of multiple events triggering a correlated event, the grouped event field values are copied to the correlated event. If none of the event fields are grouped, no values are copied to the correlated event.

The following table describes the default event fields in a correlated event. You can also customize the event field values to suit your requirements. For more information, see [Section 5.10, “Customizing Correlated Event,” on page 79](#).

Table 5-3 Default Correlated Event Fields

Correlation Event Field	ID	Sample Value	Description
EventName	evt	LoginUser	The name of the correlation rule.
EventTime	dt	2014-02-10T05:21:29.047Z	The time when the last trigger event was fired.
Message	msg	Rule triggered for every successful login	The description in the correlation rule.
ObserverCategory	rv32	SIEM	For a correlated event, this event field is always set to SIEM.
ObserverServiceComponent	rv150	SessionServices	This value is same as that of the last trigger event.
ObserverTZ	estz	Asia/Kolkata	The time zone in which the correlation engine is located.
ObserverType	st	C	For a correlated event, the event field is always set to C.
SentinelProcessingComponent	rt2	LoginUser	The correlation rule name.
SentinelProcessingComponent ID	rv123	CC72FBA4-711D-1031-8046-005056A56C5B	This is the ID of the correlation rule. The correlation rule ID remains the same even though the correlation rule name changes.
SentinelServiceComponentName	sres	LoginUser	It is the name of the correlation rule.

Correlation Event Field	ID	Sample Value	Description
SentinelServiceName	res	Correlation	For a correlated event, this event field is always set to Correlation.
Severity	sev	4	For a correlated event, this event field is always set to 4.
XDASClass	xdasclass	2	This value is same as that of the last trigger event.
XDASDetail	xdasdetail	0	This value is same as that of the last trigger event.
XDASIdentifier	xdasid	0	This value is same as that of the last trigger event.
XDASOutcome	xdasoutcome	0	This value is same as that of the last trigger event.
XDASOutcomeName	xdasoutcomename	XDAS_OUT_SUCCESS	This value is same as that of the last trigger event.
XDASProvider	xdasprov	0	This value is same as that of the last trigger event.
XDASRegistry	xdasreg	0	This value is same as that of the last trigger event.
XDASTaxonomyName	xdastaxname	XDAS_AE_CREATE_SESSION	This value is same as that of the last trigger event.

For more information on correlated event fields, click **Tips** in the Sentinel Web Console. For more information on the event taxonomy and event fields, see [Sentinel Taxonomy](#).

NOTE: By default, Sentinel correlates the correlated events received from remote Sentinel servers. If you do not want the correlation rules to consider remote correlated events, set the following property in the `/etc/opt/novell/sentinel/config/server.xml` file to `false` and restart the Sentinel server:

```
<property name="correlateRemoteCorrelationEvents">false</property>
```

5.4.2 Creating a Simple Rule

A simple rule has just one subrule. You can specify additional criteria if you want the rule to fire when all or any of the specified criteria are met. You can also specify the number of times the event should occur for the rule to fire.

- 1 Launch the Correlation Rule Builder.

For more information, see [Section 5.2, "Accessing the Correlation User Interface,"](#) on page 62.


- 2 Click **Create**.

- 3 In the Subrule window, click **Create a new expression**.

The Expression Builder is displayed. For more information, see ["Expression Builder" on page 67](#).

- 4 Select the criteria for the rule, then click **OK**.


The specified criteria are displayed in the subrule window.

- 5 (Conditional) Specify additional expressions as necessary:
 - 5a Repeat [Step 3](#) and [Step 4](#).
 - 5b Select either of the following conditions:
 - ♦ **AND**: Use this condition if you want the subrule to fire when the conditions in all of the expressions are met.
 - ♦ **OR**: Use this condition if you want the subrule to fire when the condition in either of the expressions is met.
 - 5c (Conditional) You can group events based on the distinct values of event fields or group events by same values of event fields. Select the **Group by** drop-down list, drag and drop the desired event fields in the **Group By Fields** or **Distinct Fields** list depending on how you want to group the events.
 - 5d In the **Count** field, specify the number of times the expressions must meet the specified for the rule to fire. If the count is greater than 1, the **Hr**, **Min**, and **Sec** fields are enabled.
 - 5e Specify the time frame within which the subrule should fire.
 - 5f (Conditional) If the count is greater than one and if there are any grouped event fields, by default only the grouped event field values are copied to the correlated event. If you want to copy all the event field values from the last event that triggered the correlated event, deselect **Copy only group by fields from the trigger events**.
- 6 (Optional) To associate one or more actions to the rule, click  in the Actions panel.
For more information on associating actions, see [Section 5.5, “Associating Actions to a Rule,”](#) on [page 75](#).
- 7 (Optional) To test whether the rule works as expected, click **Test Rule**.
For more information on testing the rule, see [Section 5.6, “Testing a Correlation Rule,”](#) on [page 75](#).
- 8 Click **Save As**.
- 9 Specify a name for the rule and an optional description, then click **OK**.
- 10 Deploy the rule in the Correlation Engine so that events can be processed according to the rule.
For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,”](#) on [page 78](#).

5.4.3 Creating a Sequence Rule

A sequence rule has two or more subrules that fire in sequence. You can use a sequence rule when you want the rule to fire if its subrules meet the specified criteria in the specified sequence within the defined time frame.


- 1 Launch the Correlation Rule Builder.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on [page 62](#).
- 2 Click **Create**.
- 3 In the Subrule window, click **Create a new expression**.
The Expression Builder is displayed. For more information, see [“Expression Builder”](#) on [page 67](#).
- 4 Select the criteria for the rule, then click **OK**.
The specified criteria are displayed in the subrule window.

- 5 (Conditional) Specify additional expressions as necessary:
 - 5a Select either of the following conditions:
 - ♦ **AND:** Use this condition if you want the subrule to fire when the conditions in all of the expressions are met.
 - ♦ **OR:** Use this condition if you want the subrule to fire when the condition in either of the expressions is met.
 - 5b (Conditional) You can group events based on the distinct values of event fields or group events by same values of event fields. Select the **Group by** drop-down list, drag and drop the desired event fields in the **Group By Fields** or **Distinct Fields** list depending on how you want to group the events.
 - 5c In the **Count** field, specify the number of times the expressions must meet the specified criteria for the rule to fire. If the Count is greater than 1, the **Hr**, **Min**, and **Sec** fields are enabled.
 - 5d Specify the time frame within which the subrule should fire.
 - 5e (Conditional) If the count is greater than one and if there are any grouped event fields, by default only the grouped event field values are copied to the correlated event. If you want to copy all the event field values from the last event that triggered the correlated event, deselect **Copy only group by fields from the trigger events**.
- 6 To add additional subrules, click **Subrule**, then repeat [Step 3](#) through [Step 5](#) to specify the subrule criteria.
- 7 In the Rule Type drop-down list, select **Sequence rule**.
- 8 Specify the time frame within which the rule should fire.
- 9 (Optional) To associate one or more actions to the rule, click  in the Actions panel.
For more information on associating actions, see [Section 5.5, "Associating Actions to a Rule," on page 75](#).
- 10 (Optional) To test whether the rule works as expected, click **Test Rule**.
For more information on testing the rule, see [Section 5.6, "Testing a Correlation Rule," on page 75](#).
- 11 Click **Save As**.
- 12 Specify a name for the rule and an optional description, then click **Save**.
- 13 Deploy the rule in the Correlation Engine so that events can be processed according to the rule.
For more information, see [Section 5.8, "Deploying Rules in the Correlation Engine," on page 78](#).

5.4.4 Creating a Composite Rule




A composite rule has two or more subrules that fire according to the criteria you define.

- 1 Launch the Correlation Rule Builder.
For more information, see [Section 5.2, "Accessing the Correlation User Interface," on page 62](#).
- 2 Click **Create**.
- 3 In the Subrule window, click **Create a new expression**.
The Expression Builder is displayed. For more information, see ["Expression Builder" on page 67](#).
- 4 Select the criteria for the rule, then click **OK**.
The specified criteria are displayed in the subrule window.


- 5 (Conditional) Specify additional expressions as necessary:
 - 5a Select either of the following conditions:
 - ♦ **AND:** Use this condition if you want the subrule to fire when the conditions in all of the expressions are met.
 - ♦ **OR:** Use this condition if you want the subrule to fire when the condition in either of the expressions is met.
 - 5b (Conditional) You can group events based on the distinct values of event fields or group events by same values of event fields. Select the **Group by** drop-down list, drag and drop the desired event fields in the **Group By Fields** or **Distinct Fields** list depending on how you want to group the events.
 - 5c In the **Count** field, specify the number of times the expressions must meet the specified criteria for the rule to fire. If the Count is greater than 1, the **Hr**, **Min**, and **Sec** fields are enabled.
 - 5d Specify the time frame within which the subrule should fire.
 - 5e (Conditional) If the count is greater than one and if there are any grouped event fields, by default only the grouped event field values are copied to the correlated event. If you want to copy all the event field values from the last event that triggered the correlated event, deselect **Copy only group by fields from the trigger events**.
- 6 Complete [Step 1](#) through [Step 5](#) in [Section 5.4.2, "Creating a Simple Rule,"](#) on page 70.
- 7 To add additional subrules, click **Subrule**, then repeat [Step 3](#) through [Step 5](#) to specify the subrule criteria.
- 8 In the Rule Type drop-down list, select **Composite rule**.
- 9 Select one of the following:
 - ♦ **Composite Rule (AND):** The rule fires if all the subrules meet the specified criteria within the defined time frame.
 - ♦ **Composite Rule (OR):** The rule fires if any of the subrules meets the specified criteria within the defined time frame.
- 10 (Conditional) If you selected Composite Rule (OR), use the **Count** field to specify the number of subrules that should meet the specified criteria.
 The value in the **Count** field must be less than the number of subrules. For example, if there are 5 subrules and you specify the count as 3, the rule fires if 3 or more subrules meet the specified criteria.
- 11 Specify the time frame within which the rule should fire.
- 12 (Optional) To associate one or more actions to the rule, in the Actions panel, click .
 For more information on associating actions, see [Section 5.5, "Associating Actions to a Rule,"](#) on page 75.
- 13 (Optional) To test whether the rule works as expected, click **Test Rule**.
 For more information on testing the rule, see [Section 5.6, "Testing a Correlation Rule,"](#) on page 75.
- 14 Click **Save As**.
- 15 Specify an intuitive name for the rule and an optional description, then click **Save**.
- 16 Deploy the rule in the Correlation Engine so that events can be processed according to the rule.
 For more information, see [Section 5.8, "Deploying Rules in the Correlation Engine,"](#) on page 78.

5.4.5 Creating a Free-Form Rule

If you are familiar with the rule expression syntax, you can create correlation rules by manually specifying the rule expression. You can use free-form rules to create complex rules by using additional operators such as Window, Intersection, and Union.

- 1 Launch the Correlation Rule Builder.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.
- 2 Click **Create**.
- 3 In the subrule window, click  to switch to the free-form view.
- 4 Specify the criteria for the rule.
As you type the rule expression, the Free-form editor validates the rule expression syntax and indicates errors if the syntax is wrong.
For more information on the rule expression syntax, see [Appendix B, “Correlation Rule Expression Syntax,”](#) on page 193.
- 5 (Optional) Click  to view the rule in a structured format.
Free-form expressions that include the Window operator or a combination of AND and OR operators are not supported in the structured view.
- 6 (Optional) To associate one or more actions to the rule, in the Actions panel, click .
For more information on associating actions, see [Section 5.5, “Associating Actions to a Rule,”](#) on page 75.
- 7 (Optional) To test whether the rule works as expected, click **Test Rule**.
For more information on testing the rule, see [Section 5.6, “Testing a Correlation Rule,”](#) on page 75.
- 8 Click **Save As**.
- 9 Specify an intuitive name for the rule and an optional description, then click **Save**.
- 10 Deploy the rule in the Correlation Engine so that events can be processed according to the rule.
For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,”](#) on page 78.

5.4.6 Creating Correlation Rules From Search Results

- 1 In the search results panel, select the events from which you want to create a Correlation rule.
- 2 In the **Events Operations** drop-down list, select one of the following:
 - ♦ **Add to correlation rule:** Adds the selected events to an existing rule.
 - ♦ **Create correlation rule:** Creates a new rule with the selected events.
- 3 (Conditional) If you selected **create correlation rule**, the Correlation Rule Builder is displayed. The events that you selected to build the rule are displayed below the rule builder. Skip to [Step 5](#).
- 4 (Conditional) If you selected **add to correlation rule**, the Add events to an existing rule window is displayed that lists the rules in the system.
Select a rule, then click **OK**.
The Correlation Rule Builder is displayed. The events that you selected to build the rule are displayed below the rule builder.
- 5 From the event list, drag the attributes that you want to add to the rule to the Subrule window.
- 6 (Optional) To associate one or more actions to the rule, in the Actions panel, click .

For more information on associating actions, see [Section 5.5, “Associating Actions to a Rule,” on page 75](#).

- 7 (Optional) To test whether the rule works as expected, click **Test Rule**.

For more information on testing the rule, see [Section 5.6, “Testing a Correlation Rule,” on page 75](#).

- 8 Click **Save As**.

- 9 Specify an intuitive name for the rule and an optional description, then click **Save**.

- 10 Deploy the rule in the Correlation Engine so that events can be processed according to the rule.
For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#).

5.5 Associating Actions to a Rule


You can configure one or more actions to a rule. The associated actions are executed when the rule fires.

- 1 Launch the Correlation interface.

For more information, see [Section 5.2, “Accessing the Correlation User Interface,” on page 62](#).


- 2 In the Correlation panel, click any rule to which you want to associate actions, then click .

The Correlation Rule Builder is displayed.

- 3 In the Actions panel, click  to associate one or more actions to the rule.

The list of actions is displayed.

- 4 Select the actions that you want to associate with the rule, then click **OK**.

- 5 Click  to define when the action should execute.

- 6 Select one of the following:

- ♦ **Perform actions every time the rule fires:** Sentinel creates a correlated event and executes the associated action each time the rule fires.
- ♦ **Perform actions at most every:** Sentinel creates a correlated event and executes the action at most every specified time interval. By default, this option is selected and the time interval is set to 1 hour. This is to ensure that rule does not fire continually and overutilize resources. You can also define the maximum number of trigger events to be associated with the correlated event. For more information, see [“Configuring the Number of Trigger Events to be Associated with a Correlated Event”](#) in the *NetIQ Sentinel Administration Guide*.

- 7 Click **OK**.

- 8 Click **Save Rule**.



NOTE: If you modified a deployed rule, you must redeploy the rule in the Correlation Engine for the changes to take effect. For information on deploying a rule, see [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#).

5.6 Testing a Correlation Rule

You can determine whether the rule is working as expected by testing a rule on the events that are already in the system before deploying it to monitor real-time events.

- 1 Launch the Correlation interface.

For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.

- 2 In the Correlation panel, click any rule that you want to test, then click .
- 3 Click **Test Rule**.
- 4 Specify the time frame during which you want to test the rule.
- 5 (Optional) Click  to filter events that the rule should process.
- 6 Click **Test Rule**.

The test takes some time, depending on the specified criteria. After the test is complete, the test results are displayed.

The test results display the rule details:

- ♦ **Status:** Indicates whether the test is running, stopped, or completed. The test stops when the rule has fired at least 20 times during the test process. This ensures that the rule is working as expected and saves time when there are many events.
- ♦ **Started at:** The date/time when the rule started to fire.
- ♦ **Finished at:** The date/time when the test stopped.

The indicators (dots) indicate when the rule fired. The white dot indicates a single correlation event. The black dot indicates multiple correlation events generated within a short period of time. Click the indicator to see the event details.

- 7 Click the **Close** icon to close the test results.

5.7 Sample Correlation Rules



This section provides a few examples on how you can create correlation rules. For more examples, see [Appendix B, “Correlation Rule Expression Syntax,”](#) on page 193.

- ♦ [Section 5.7.1, “Detecting Critical Events from an Intrusion Detection System,”](#) on page 76
- ♦ [Section 5.7.2, “Detecting a Spreading Attack,”](#) on page 77
- ♦ [Section 5.7.3, “Detecting an Attack that Came from Outside the Firewall,”](#) on page 77

5.7.1 Detecting Critical Events from an Intrusion Detection System

This example identifies critical events from an intrusion detection system and sends an e-mail to the Administrator.

- ♦ Launch the Correlation Rule Builder. In the **Correlation** panel, click **Create**.
- ♦ In the Subrule window, click **Create a new expression**.
- ♦ Specify that the events must be from an intrusion detection system (IDS):
 - ♦ In the Expression Builder > **Event Fields**, select **ObserverCategory**.
 - ♦ Ensure that the “=” operator is selected.
 - ♦ In the Value field, specify `IDS`, then click **OK**.
- ♦ Identify critical events:
 - ♦ Add another expression. In the Subrule window, click **Create a new expression**.
 - ♦ In the Expression Builder > **Event Fields**, select **Severity**.


- ♦ Select **>=** as the operator.
- ♦ In the **Value** field, select **4**, then click **OK**.
- ♦ If events are found, send an e-mail to the administrator:
 - ♦ In the Actions panel, click  to associate the action with the rule.
 - ♦ Select **Send E-mail**.
 - ♦ Click  to update the action execution criteria.
 - ♦ Select **Perform actions everytime the rule fires**, then click **OK**.
- ♦ Click **Save Rule**.
- ♦ Deploy the rule in the Correlation Engine.

For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#).
- ♦ Search events that match the rule criteria.

For more information, see [Section 5.9, “Viewing Correlated Events,” on page 78](#).

5.7.2 Detecting a Spreading Attack

This example creates a Correlation rule that indicates whether the source of an attack was previously the destination of an attack (within 15 minutes.) Because this involves comparing a current event set with a past event set, it uses the window operation.

- ♦ In the Subrule window, click  to switch to the free-form editor.
- ♦ Specify the expression as follows:



```
filter(e.TaxonomyLevel1="Attack") flow window(w.dip=e.sip,
filter(e.rv51="Attack"), 15m)
```
- ♦ Click **Save Rule**.
- ♦ Deploy the rule in the Correlation Engine.

For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#).
- ♦ Search events that match the rule criteria.

For more information, see [Section 5.9, “Viewing Correlated Events,” on page 78](#).

5.7.3 Detecting an Attack that Came from Outside the Firewall

This example creates a Correlation rule that checks whether an intrusion detection system attack event seen inside your network came through your firewall in the last 10 seconds.

- ♦ In the Subrule window, click  to switch to the free-form editor.
- ♦ Specify the expression as follows:


```
filter(e.TaxonomyLevel1="Attack") flow window(w.dip=e.sip,
filter(e.rv32="FW"), 10)
```
- ♦ Click **Save Rule**.
- ♦ Deploy the rule in the Correlation Engine.

For more information, see [Section 5.8, “Deploying Rules in the Correlation Engine,” on page 78](#).
- ♦ Search events that match the rule criteria.

For more information, see [Section 5.9, “Viewing Correlated Events,” on page 78](#).

5.8 Deploying Rules in the Correlation Engine

You can deploy the Correlation rules either from the rule dashboard or from the Correlation Engine dashboard.

To deploy rules from the Correlation Engine dashboard:



- 1 Launch the Correlation interface.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,” on page 62.](#)
- 2 In the Correlation panel, select the rule that you want to deploy and click the **Open Overview** icon.
- 3 In the Deploy/Undeploy section, select the engine to which you want to deploy the rule, then click **Deploy**.


To deploy rules from the Correlation Rule dashboard:

- 1 Launch the Correlation interface.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,” on page 62.](#)
- 2 In the Correlation panel, click the engine to which you want to deploy rules.
The Correlation Engine dashboard is displayed.
- 3 In the Available rules section, select the rule or rules that you want to deploy, then click **Deploy**.

5.9 Viewing Correlated Events

Correlated events contain detailed information about the trigger events. To view correlated events, perform the following:

- 1 Launch the Correlation interface.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,” on page 62.](#)
- 2 In the Correlation panel, select any rule, then click .
- 3 (Optional) Click  to see the correlated event fields and their values. For more information, see [Table 5-3 on page 69.](#)

The events that match the rule criteria are displayed in the search results panel. The correlated events are displayed with the  icon.

You can use the event field IDs to create search queries to find specific correlated events. For example, if you want to search for the correlated events that were generated because of the correlation rule `LoginUser`, specify the following query in the **Search** field:

```
st:C AND rt2:LoginUser
```

For more information about searching for events, see [Section 3.1, “Running an Event Search,” on page 33.](#)

- 4 (Optional) Click **View triggers** to view the events that generated the correlated event.

5.10 Customizing Correlated Event

Correlated events include specific event fields with some default values. For more information, see [Table 5-3 on page 69](#) in [Section 5.4.1, “Understanding the Correlated Event,” on page 69](#).

You can customize the default correlated event field values. For example, when a specific event pattern is detected and you want the rule to create correlated events with a high severity value, you can customize the Severity field of the correlated event. Similarly, if you want to allow a specific tenant to view correlated events specific to their tenant, you can set the TenantName field to the tenant associated with the rule.

You can customize the correlated event when performing either of the following:

- ◆ **Creating the correlation rule:**
 - ◆ You can use the **Customize correlated event** option that allows you to customize almost all the correlated event fields except for the Sentinel internal fields.
 - ◆ Any customization in the correlated event field values is applicable only to the specific correlation rule.
- ◆ **Configuring the Generate Custom Correlation Event action:**
 - ◆ When specifying the attribute values you can either reference to a correlated event field value or the last event that triggered the correlated event. However, you can customize only certain event fields.
 - ◆ You can associate the customizations made in the action to multiple correlation rules.

For more information, see [“Managing Actions”](#) in the *NetIQ Sentinel Administration Guide*.

To customize correlated event field values:

- 1 In the **Correlation** panel, select the correlation rule, then click the Edit icon.
- 2 In the correlation rule builder, click **Customize correlated event**.
- 3 Drag the event fields you want to edit from the **Available event fields** section to the **Selected event fields** section.
- 4 Click the **Edit** icon to specify the value for the event field.

By default, the value is same as the event field value of the last event that triggered the correlated event. You can also reference other event field values in the **Value** field by enclosing the event field in `$`.

For example, for a login failure event, in the Message field you may want to refer to the InitiatorUserName, you can specify the Value as `Login failure by sun`.

- 5 Click **OK**.
- 6 After customizing the correlated event fields, click **Set** to save the correlated event.
- 7 Click **Save Rule**.

The customized correlated event values are applied only for correlated events generated after the customization.

5.11 Managing Correlation Rules

- ◆ [Section 5.11.1, “Viewing the Rule Dashboard,” on page 80](#)
- ◆ [Section 5.11.2, “Editing a Rule,” on page 81](#)
- ◆ [Section 5.11.3, “Deleting a Rule,” on page 81](#)

5.11.1 Viewing the Rule Dashboard

The Rule dashboard displays overall information of the rule. The Rule dashboard helps you to deploy or undeploy a rule on a Correlation Engine, and also helps you manage the rule status. After a rule is deployed, you can monitor the health of the rule, activities such as the events processed, and the memory usage by the rule.

To view the Rule dashboard, select the desired rule in the Correlation panel.

After a rule is deployed, the dashboard displays the following information:


- ◆ **Rule health statistics:** Indicates the overall performance of the rule and enables you to monitor the activities of the rule.
 - ◆ **Activity statistics:** Indicates the activities of the rule since it was deployed in the Correlation Engine:
 - ◆ **Fire count:** The number of times the rule fired. You can use this information to discover a rule that fires more than expected and that might need to be tuned, or to discover a rule that does not fire as often as you would expect. In either case, this tab guides you to the rules that are the most and least active. The search icon allows you to view the events generated since the rule was deployed or enabled.
 - ◆ **Fire rate:** The number of times the rule has been fired relative to the events processed by the rule. This statistic is similar to fire count in that it gives an indication of how active a rule is. However, instead of giving a raw count, the fire rate gives a percentage that is relative to the number of events a rule has processed.
 - ◆ **EPS utilization:** The processing time this rule consumes relative to the capacity of the engine. This statistic provides an estimate of the amount of engine capacity a given rule is currently consuming. Rules that are more complex, have time-consuming actions, or fire frequently consume more capacity. You can use this statistic to determine whether the rule needs to be tuned or perhaps moved to another Correlation Engine for scalability reasons.
 - ◆ **Events processed:** The number of events processed by the rule since the rule was deployed.
 - ◆ **Total processing time:** Total time spent by the Correlation Engine processing the rule since it was deployed or enabled.
 - ◆ **Memory statistics:** Indicates the memory consumed by the rule:
 - ◆ **Estimated memory utilization:** Gives a snapshot of roughly how much memory a rule is consuming. Rules consume memory when they have discriminators specified for fields with multiple values (through the **Group by** list), and when rules hold events in memory for operations like the advanced “window” operation. Rules that consume a lot of memory are a potential liability to a healthy system and should be carefully reviewed to ensure they are properly written or possibly moved to another Correlation Engine for scalability reasons.
 - ◆ **Events in memory:** Number of events held in memory by the rule.
 - ◆ **Cardinality:** Number of strings and related structures held in memory by the rule.
- ◆ **Deploy/Undeploy:** Lists the available Correlation Engines. You can select an engine, then click **Deploy** or **Undeploy** to add or remove the rule in the Correlation Engine.
- ◆ **Associated actions:** Lists the actions associated with the rule.
- ◆ **Status:** Indicates the current status of the rule. You can also use this option to enable or disable the rule.

5.11.2 Editing a Rule

1 Launch the Correlation interface.

For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.

2 Do either of the following:

- ◆ In the Correlation panel, select the rule for which you want to modify the details, then click .
- ◆ In the Rule dashboard window, select the rule for which you want to modify the details, then click **Edit rule**.

The rule details are displayed in the Correlation Rule Builder.

3 Make the necessary changes, then click **Save Rule**.

NOTE: If you modify a deployed rule, you must redeploy the rule in the Correlation Engine for the changes to take effect. For information on deploying a rule, see [Section 5.8, “Deploying Rules in the Correlation Engine,”](#) on page 78.

5.11.3 Deleting a Rule

You can delete rules that are not deployed in the Correlation Engine. To delete a deployed rule, you must first undeploy the rule from the Correlation Engine.

1 Launch the Correlation interface.

For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.

2 Select the rule that you want to delete, then click .

3 Click **OK** to confirm deletion.

5.12 Managing the Correlation Engine

- ◆ [Section 5.12.1, “Using the Correlation Engine Dashboard,”](#) on page 81
- ◆ [Section 5.12.2, “Stopping or Starting a Correlation Engine,”](#) on page 83
- ◆ [Section 5.12.3, “Renaming a Correlation Engine,”](#) on page 84

5.12.1 Using the Correlation Engine Dashboard

The Correlation Engine dashboard provides an overall picture of the health of the engine and the various rules deployed to it. The dashboard provides information on the activity of an engine and provides insights about which rules are behaving as expected and which rules might need additional tuning.

To view the Correlation Engine dashboard, select the desired Correlation Engine in the Correlation panel.

The dashboard displays a simple view of all the rules deployed to that engine and a list of rules available to deploy. You can also see some general information about the engine, such as its current state (running, stopped, or offline) and how long the engine has been in that state. There is also a summary that shows how many events the engine has processed, and an indication of current utilization of the engine (EPS utilization). In general, you can think of this number as something

analogous to CPU utilization, but this metric indicates how much of the capacity of the engine is currently utilized. As the number of rules and the complexity of rules increases along with the current EPS (events per second) rate of the system, you can expect this number to grow larger.

The **Deployed rules comparison** tab provides a more granular view of engine activity. The various tabs represent various statistics, and you can select a given tab to sort the rules (ascending or descending) by these statistics to get a clear picture of how various rules are behaving and consuming resources.

The Correlation Engine dashboard also allows you to manage the engine and the rules in your system:

- ◆ [“Managing the Correlation Engine” on page 82](#)
- ◆ [“Managing Deployed Rules” on page 82](#)
- ◆ [“Comparing Deployed Rules” on page 82](#)
- ◆ [“Viewing Engine Details” on page 83](#)
- ◆ [“Viewing Available Rules” on page 83](#)

Managing the Correlation Engine

The following options are available to manage the Correlation Engine:

- ◆ **Stop:** Stops an active Correlation Engine. When the engine stopped, it does not monitor the events against the deployed rules.
- ◆ **Undeploy all:** Undeploys all the deployed rules from the engine.
- ◆ **Rename:** Allows you to rename the engine.

Managing Deployed Rules

This section lists the number of rules and the rules deployed in the engine.

The following options are displayed when you mouse over a rule:

- ◆ **View:** Opens the Rule dashboard to provides overall information on the rule. For more information, see [Section 5.11.1, “Viewing the Rule Dashboard,” on page 80](#).
- ◆ **Disable:** Allows you to disable the rule. When a rule is disabled, it does not process the events.
- ◆ **Undeploy:** Undeploys the rule from the Correlation Engine.

Comparing Deployed Rules

This section helps you compare the rules based on parameters such as fire count, EPS capacity, and memory utilization. You can sort the rules as desired by using the up-arrow and down-arrow icons.

- ◆ **Fire count:** The number of times the rule has fired since it was deployed or enabled. You can use this information to discover a rule that fires more than expected and that might need to be tuned, or to discover a rule that does not fire as often as you would expect. In either case, this tab guides you to the rules that are the most and least active.
- ◆ **Last fired:** The last time the rule fired since it was deployed. This statistic is useful for determining the rules that are currently active and inactive in the system. A rule might have fired frequently, but it has not fired recently. Or, a rule might fire infrequently, but it fired recently. This tab gives you a real-time picture of what is active in the engine at any given time.

- ♦ **Fire rate:** The number of times the rules have fired relative to the events processed by the engine. This statistic is similar to fire count in that it gives an indication of how active a rule is. However, instead of giving a raw count, the fire rate gives a percentage that is relative to the number of events a rule has processed. This normalizes the metric, and rules that were recently deployed can be compared with rules that were deployed at an earlier time.
- ♦ **EPS utilization:** The events processing time the rule consumes relative to the capacity of the engine. This statistic provides an estimate of the amount of engine capacity a given correlation rule is currently consuming. Rules that are more complex, have time-consuming actions, or fire frequently consume more capacity. You can use this statistic to identify rules that need to be tuned or perhaps moved to another correlation engine for scalability reasons.
- ♦ **Memory utilization:** The estimated memory utilization of the rule. In addition to EPS utilization, which provides a good picture of how much time a rule consumes relative to the total available processing time of the engine, the memory utilization gives a snapshot of roughly how much memory a rule is consuming. Rules consume memory when they have discriminators specified for fields with multiple values (through the **Group by** list), and when rules hold events in memory for operations like the advanced “window” operation. Rules that consume a lot of memory are a potential liability to a healthy system and should be carefully reviewed to ensure they are properly written or possibly moved to another engine for scalability reasons.

Viewing Engine Details

This section lists the Correlation Engine details and so you can monitor the performance of the Correlation Engine.

- ♦ **Engine ID:** The Correlation Engine ID.
- ♦ **Engine Name:** The name of the Correlation Engine.
- ♦ **Host IP:** The IP address of the host machine where the Correlation Engine is installed.
- ♦ **Hostname:** The hostname of the machine where the Correlation Engine is installed
- ♦ **State:** Whether the status of the Correlation Engine is running or stopped.
- ♦ **Events Processed:** The number of events processed by the deployed rules since they were deployed.
- ♦ **Last changed state:** The time the Correlation Engine status was last changed.
- ♦ **EPS utilization:** The processing time the Correlation Engine consumes relative to the capacity of the engine.

Viewing Available Rules

This section lists the available rules in the system that are not deployed in the Correlation Engine. It also allows you to select rules and deploy them in the Correlation Engine.

5.12.2 Stopping or Starting a Correlation Engine

The Correlation Engine is in the Start mode by default and keeps processing events for the rules deployed in the engine. You can determine when the Correlation Engine should process the events and start or stop the Correlation Engine accordingly.

- 1 Launch the Correlation interface.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.
- 2 In the **Correlation > Engines** section, select the Correlation Engine that you want to stop.
The Correlation Engine dashboard is displayed.

- 3 (Conditional) To stop the engine, click **Stop**.
The engine is stopped and does not process any events.
- 4 (Conditional) To start the engine, click **Start**.

5.12.3 Renaming a Correlation Engine

- 1 Launch the Correlation interface.
For more information, see [Section 5.2, “Accessing the Correlation User Interface,”](#) on page 62.
- 2 In the **Correlation > Engines** section, select the Correlation Engine that you want to rename.
The Correlation Engine dashboard is displayed.
- 3 Click **Rename**, which is above the Rules Deployed section.
- 4 Modify the name of the engine, then click anywhere outside the field.

6 Visualizing and Analyzing Alerts

Alerts notify you of what is most important for you to look at. Alerts can relate to threats to IT resources or performance thresholds such as system memory full or IT resources not responding. Correlation rules define the patterns that you are alerted to. Sentinel automatically associates the relevant events and identities with the alert to help you determine the root cause of a potential threat.

Sentinel provides a graphical and tabular representation of alerts. Visualizing alerts helps you identify and analyze potential threats against your IT resources. You can view and analyze alerts in the following panels in the Sentinel Web console:

- ♦ **Alert Views:** The real-time alert views show you the alerts that are most important to look at and enable you to view and manage alert details. Charts provide a summary of alerts and the table provides a prioritized list of all the alerts. Alert view also enables you to perform alert triage operations such as changing states of an alert, assigning alerts to users or roles, adding information to the knowledge base, and so on. You can further drill-down into each alert to view the alert details such as trigger events, user identities involved, alert history, and so on. For more information about alert view, see [Section 6.1, “Viewing and Triaging Alerts in Alert Views,” on page 85](#).
- ♦ **Dashboards:** Alert dashboards enable you to perform powerful exploration and analysis of alerts. For example, you can find out the average time taken by owners to close alerts, the correlation rule generating the maximum number of alerts, geographical locations from where the alerts with high severity are detected, and so on. The alert dashboard enables you to create customized charts and tables for analysis. You can filter and refine the data further as you select certain areas in the charts and use the query and filter options. For more information about alert dashboard, see [Section 6.2, “Analyzing Alert Dashboards,” on page 88](#).

This chapter provides information about the following:

- ♦ [Section 6.1, “Viewing and Triaging Alerts in Alert Views,” on page 85](#)
- ♦ [Section 6.2, “Analyzing Alert Dashboards,” on page 88](#)
- ♦ [Section 6.3, “Troubleshooting,” on page 90](#)

6.1 Viewing and Triaging Alerts in Alert Views

To view and analyze alerts, you must first create the alert view. To configure the alert view, you must either be an administrator or have the Manage Alerts permission. For more information, see [“Visualizing and Analyzing Alerts” in *NetIQ Sentinel Administration Guide*](#).

- ♦ [Section 6.1.1, “Creating an Alert View,” on page 85](#)
- ♦ [Section 6.1.2, “Viewing Alerts,” on page 86](#)
- ♦ [Section 6.1.3, “Creating Incidents from Alerts,” on page 88](#)

6.1.1 Creating an Alert View

To view and analyze alerts, you must first configure the alert view.

To create an alert view:

- 1 Log in to the Sentinel Web console.
- 2 Click **Real-time Views > Alert Views > the Create** icon.
- 3 Specify the following information:
 - ◆ **Name:** Specify a name for the alert view.
 - ◆ **Sharing:** Select either of the following options:
 - ◆ **Public:** Allow everyone to view the alert view. In the public mode, you are the owner of the alert view and other users cannot edit it.
 - ◆ **Private:** Only you will be able to view the alert view.
 - ◆ **Data sources:** Add other data sources from which you want to view alerts. For information about data sources, see “[Configuring Data Federation](#)” in the *NetIQ Sentinel Administration Guide*.
 - ◆ **Criteria:** Specify the criteria to filter the alerts.
 - ◆ **Time range:** Specify the time range for which you want to view alerts.
 - ◆ **Use alert period:** Select Created or Modified to view the alerts that were created or modified in the specified time range.
 - ◆ **Tenant:** If you are in a multi-tenant environment, select the department or the tenant name for which you want to view alerts.

NOTE: This option is displayed only if you are an administrator in a multi-tenant environment. For information about multitenancy, see “[Configuring Sentinel for Multitenancy](#)” in the *NetIQ Sentinel Administration Guide*.

- 4 Click **Save** to save the alert view.

6.1.2 Viewing Alerts

Sentinel provides a graphical and tabular representation of alerts that match the specified alert criteria. The charts represent the alerts overview information classified by Priority, State, and Severity. The alert view table displays only distinct alerts. Duplicate alerts are rolled up to a single distinct alert. For more information about rolling up of duplicate alerts, see “[Configuring Alert Creation](#)” in the *NetIQ Sentinel Administration Guide*. The alert view table provides information about an alert such as, severity, priority, owner, state, occurrences, and so on. Alerts from Sentinel servers in distributed location are distinguished by the Remote icon (🌐) next to the name of the alert. You can view the IP address of the remote Sentinel server by moving the mouse over the name of the alert.

NOTE: The alerts displayed in the alert view are filtered depending on the manage alert permissions and the tenancy of the user role. For more information about permission to manage alerts, see “[Configuring Roles and Users](#)” in the *NetIQ Sentinel Administration Guide*.

To create customized charts and tables for further analysis of the alert data, see [Section 6.2, “Analyzing Alert Dashboards,”](#) on page 88.

To view alert views:

- 1 In the Sentinel Web interface, click **Real-time Views > Alert Views**.
- 2 Select the desired alert view and click the **Open alerts view** icon.

As you visualize and monitor alerts, you can perform the following activities in the alert view:

- ◆ Mouse over the charts to determine the number of alerts based on alert states, priority, and severity.
- ◆ Sort alerts based on one or more columns in the table. Press Shift+click to select multiple columns to sort. By default, the alert view table displays alerts based on the time when the alerts were triggered. Therefore, the latest alerts are listed on the top in the table.
- ◆ Assign alerts to a user or a role, including yourself or your role.
- ◆ Modify the alert state to indicate the progress on the alert investigation.
- ◆ Add comments to the alert to indicate the changes you made to the alert, which helps you to keep an up-to-date record of the alert investigation. For example, you can add comments when you change the state of a specific alert or when you have gathered more information about the alert. Providing specific comments allows you to accumulate knowledge about a particular instance of the alert and track how a particular condition was addressed. Comments are important in tracking the alert, particularly if the process of resolving the alert spans several users or roles.
- ◆ View events that triggered the alert and drill-down further to the extent of viewing the user identities that triggered the event by clicking the **View details** icon in the alert view table.

The Alert Details page displays a detailed information about an alert including the following:

- ◆ **Source:** Displays the correlation rule that generated the alert. You can also annotate the correlation rule by adding information to the knowledge base so that future alerts generated by this correlation rule include the associated historical information.
- ◆ **Knowledge Base:** Knowledge base is a repository that contains information about the conditions that resulted in the alert. It can also include information about resolution of a particular alert, which can help others resolve similar alerts in the future. Over time, you can collect a valuable knowledge base about the alert specific to a tenant or an enterprise.

For example, an employee has recently joined the organization and does have the access permissions to a secured server. But this employee might not have been added yet to the authorized users list. Therefore, an alert is generated every time the employee tries to access the server. In such a case, you can add a note in the alert knowledge base to indicate that the “employee is approved to access the server, but is not yet listed in the authorized users list. This alert can be ignored and set to low priority.”

NOTE: To view or edit the knowledge base, you must be an administrator or have the **View Knowledge Base** or **Edit Knowledge Base** permissions.

- ◆ **Alert Fields:** Displays the alert fields that provide the following information:
 - ◆ who and what caused the alert
 - ◆ the assets affected
 - ◆ the taxonomic categories of the action that caused the alert, the outcome, and so on. For more information on taxonomy, see [Sentinel Taxonomy](#).

For more information about alert fields, click **Tips** on the top-right corner of the Sentinel Web interface.

- ◆ **Trigger Events:** Displays the events that triggered the correlated event associated with the alert. You can determine the conditions that triggered the alert by examining the trigger events. By default, the Alert Details page displays 10000 trigger events per alert. You can also define this number as necessary. For more information, see “[Configuring the Number of Trigger Events to be Displayed in the Alert View](#)” in the *NetIQ Sentinel Administration Guide*.

- ♦ **Show history:** Displays the changes made to the alert, which helps you track any actions taken on the alert.
- ♦ **Identities:** Displays the list of users involved in the alert. This information helps you to investigate about the users involved in the alert and monitor their activities.

6.1.3 Creating Incidents from Alerts

If the potential threat that caused the alert is critical such that it needs immediate action or the threat indicates any security compromise, escalate the threat by creating an incident. You can associate the trigger events that caused the alert to an incident.

NOTE: You must be an administrator or have the **View or create incidents and add events to incidents** permission to create incidents.

To create incidents from alerts:

- 1 In the **Alert Details** page > **Source** section, click the **Open correlation overview** icon.
- 2 In the **Rule Health Statistics** section, click the **Search events** icon.
The search result displays all the correlated events generated by the correlation rule associated to the alert.
- 3 In the search results, click **View triggers** for a specific correlated event.
The search result displays all the events that triggered the correlated event.
- 4 In the search results, select the events you want to associate to the incident. Click **Event Operations > Create Incidents or Add to incident**. For more information about creating an incident from events, see [Section 3.5.4, “Creating an Incident,” on page 46](#). For information about adding events to an existing incident, see [Section 3.5.3, “Adding Events to an Incident,” on page 45](#).

6.2 Analyzing Alert Dashboards

The alert dashboard provides a high-level overview of the alerts in your organization. It enables you to analyze and study common patterns in alerts, such as types of alerts, geographical locations from where the alerts originated, oldest open alerts, and alerts that took longest time to close.

For example, you are a Security Operations Center manager in a multi-tenant environment, and you want to analyze and investigate alerts in detail and also understand how your team is handling the alerts. You can perform the following analysis in the alert dashboard:

- ♦ **Investigate Alerts:** You can view the alerts generated over time, number of open alerts versus closed alerts, top correlation rules generating the most number of alerts, oldest open alerts, any spikes in alerts at a specific time range, and so on.
- ♦ **Monitor team performance:**
 - ♦ The type of alerts the team has been working on
 - ♦ How the alert load is distributed among top owners
 - ♦ Time taken to close alerts of specific priorities
 - ♦ Find the team member owning the most number of alerts
 - ♦ Team members that took longest to investigate alerts

- ♦ **Monitor performance against tenant service-level agreement (SLA):** You can view alerts from various tenants, analyze the most number of alerts from a specific tenant, time taken to investigate or close alerts for a specific tenant compared to other tenants, and so on.

The Alert dashboard provides a customizable and an easy-to-configure interface that helps you to view and investigate alerts in detail.

To create or view alerts in the dashboard, you must either be an administrator or have the permission to manage alerts. Depending on the alert permissions and the tenant you belong to, Sentinel displays the relevant alerts in the dashboard.

6.2.1 Creating the Alert Dashboard

To create the alert dashboard:

- 1 In the Sentinel Web console, click **Dashboards** > **Alert** > click the **Create alert dashboard** icon. The Alert dashboard view opens in a new window. For information about analyzing alerts, see [Section 6.2.2, “Analyzing Alerts,” on page 89](#).
- 2 (Optional) Customize the default dashboard to suit your requirements. For more information about customizing the alert dashboard, see [Section 6.2.3, “Customizing the Alert Dashboard,” on page 90](#).
- 3 Click the **Save** icon to save the customized dashboard.

6.2.2 Analyzing Alerts

The Alert dashboard provides some pre-configured panels that provide information about alerts such as the following:

- ♦ **Overview:** Displays a time series chart that shows alerts generated in Sentinel over time. You can inspect the time series charts for any spikes, which can indicate increase in attacks in your organization. You can drag and select the time period when the spike occurred to zoom into the alerts. As you select the specific time range, Sentinel filters the dashboard for alerts in the selected time range. Also, you can find out the geographical locations from where the alerts originated.

To view geographical locations from where the alerts originated, ensure that the `IpToCountry.csv` file is populated by using the Generic IP Geolocation Collector. For more information, see the Generic IP Geolocation Collector on the [Sentinel Plug-ins Website \(https://www.netiq.com/support/sentinel/plugins/\)](https://www.netiq.com/support/sentinel/plugins/).

- ♦ **Alert Load:** Provides information about the alerts at a granular level such as the following:
 - ♦ Topmost alerts in your enterprise
 - ♦ Alert distribution among top alert owners
 - ♦ Total number of alerts in individual alert states
 - ♦ Number of alerts received from each tenant
 - ♦ Total number of alerts based on priority
- ♦ **Performance rows:** Provides statistical information about how efficiently alerts are investigated and closed based on priority, correlation rule, alert owners, and tenants.
- ♦ **Details:** Provides detailed alerts information such as the oldest open alerts, number of times the duplicate alerts were rolled up, and all alert fields.

The alert dashboard displays all alerts in your local Sentinel server. To view alerts from other Sentinel servers, you need to view the alerts in the Alert Views. For more information, see [Section 6.1, “Viewing and Triaging Alerts in Alert Views,” on page 85](#). The alert dashboard displays only distinct alerts. Duplicate alerts are rolled up to a single distinct alert. For more information about the rolling up of duplicate alerts, see “[Configuring Alert Creation](#)” in the *NetIQ Sentinel Administration Guide*.

To view the alert dashboard:

- 1 In the Sentinel Web Console, click **Dashboards > Alert**.
- 2 Select the desired alert dashboard and click the **Open alert dashboard** icon.
- 3 As you visualize and monitor alerts, you can perform the following activities in the alert dashboard:
 - ♦ Mouse over specific areas in the charts to view more information.
 - ♦ Select desired areas in the chart to filter the alert data. As you select a specific area in the chart, Sentinel filters the alerts in rest of the charts and tables in the dashboard. Click **Filtering** to remove the applied filters and go back to the unfiltered view.
- 4 (Optional) You can customize the default view and save the dashboard. For information about customizing the dashboard, see “[Customizing the Alert Dashboard](#)” on page 90.
- 5 (Conditional) To perform various operations on alerts such as closing an alert, assigning alerts to a user, and so on, see [Section 6.1.2, “Viewing Alerts,” on page 86](#).

6.2.3 Customizing the Alert Dashboard

Sentinel leverages Kibana, a browser-based analytics and search dashboard, that helps you to visualize and analyze data. The dashboard is divided into rows and panels. You can create rows and add panels as required to display various charts. You can drag and drop the rows or panels to arrange them on the dashboard. You can also configure the size, style, and type of panels for a row. For more information, see [rows and panels](#) in Kibana documentation. As you visualize and monitor alerts, you can customize the alert dashboard as follows:

- ♦ Apply multiple filters to refine the alert data using the Query and Filter options. For more information about queries and filters, see [queries and filters](#) in Kibana documentation.
- ♦ Customize the rows and panels according to your requirements, for example, drag and drop the panels to arrange them on the dashboard. Click **Configure Dashboard > Rows** tab to arrange, create, or remove rows in the dashboard. For example, if you have to view the tenant specific information frequently, move the tenant performance row to the top or if you do not want to view the detailed tables, you can remove the Details row.
- ♦ Click the **configure** option in the panels to customize the size and style of panels.
- ♦ Add markers in the time series panel to display additional information such as owners, severity, and so on.

6.3 Troubleshooting

This section lists issues, which may occur when viewing alerts, along with the solution to resolve the issues.

6.3.1 Unable to View Alerts in the Dashboard and Alert Views

The alert dashboard and the charts in the alert view do not refresh or display new alerts. However, the table in the alert view displays the newly generated alerts. This issue could happen because of a corrupt alert index. For more information, see [“Unable to View Alerts in the Dashboard and Alert Views”](#) in the *NetIQ Sentinel Administration Guide*.

7 Analyzing Trends in Data

The following sections describe how the Sentinel Security Intelligence feature analyzes trends in data and how to use the Security Intelligence feature.

- ◆ [Section 7.1, “Overview,” on page 93](#)
- ◆ [Section 7.2, “Creating a Dashboard,” on page 96](#)
- ◆ [Section 7.3, “Understanding the Dashboard Interface,” on page 97](#)
- ◆ [Section 7.4, “Creating Baselines,” on page 98](#)
- ◆ [Section 7.5, “Configuring Anomaly Detection,” on page 99](#)
- ◆ [Section 7.6, “Viewing Anomaly Events,” on page 101](#)
- ◆ [Section 7.7, “Managing Dashboards,” on page 103](#)
- ◆ [Section 7.8, “Troubleshooting,” on page 104](#)

7.1 Overview

The Correlation capability in Sentinel provides security analysis with the ability to look for “known knowns” occurring within the enterprise in real time. Sentinel provides another way to analyze data by looking for the “known unknowns” in the events generated within the enterprise.

This second type of analysis is called Security Intelligence in Sentinel. It allows you to perform anomaly-based analysis so you can find deviations from the normal trends of your enterprise.

A key aspect in looking for deviations depends upon the nature of input data. Input data is typically a collection of various instances, records, or points. Each data instance itself consists of attributes that define the type of input. Examples of different types of input data within the SIEM domain are:

- ◆ Spatial (location of various entities)
- ◆ Graphical (relation between the various entities)
- ◆ Sequences (time series)

The Security Intelligence feature in Sentinel focuses on statistical analysis of time series data to enable analysts to identify and analyze anomalies either by an automated statistical engine or by visual representation of the statistical data for manual interpretation. The following concepts collectively define the Security Intelligence feature in Sentinel:

- ◆ [Section 7.1.1, “Terminology,” on page 93](#)
- ◆ [Section 7.1.2, “How Security Intelligence Works,” on page 95](#)
- ◆ [Section 7.1.3, “Permissions for Security Intelligence,” on page 96](#)

7.1.1 Terminology

The following terminology is used with Security Intelligence:

Security Intelligence: The feature that allows you to create dashboards with baselines that define what is normal in your network. An analysis of the real-time data lets you see if there are any anomalies, and an e-mail is sent to notify someone that anomalies are occurring.

Dashboard: Displays data matching a particular filter, categorizes data using a particular classifier, and looks for particular anomalies as specified in the dashboard's set of anomaly rule definitions. This is not just a front-end Web interface, but also a back-end engine because it is always looking for anomalies, even if the Web interface is not open.

Classifier: Determines the categories displayed in the dashboard.

Filter: Determines the scope of events displayed in the dashboard.

MongoDB: A database that stores the Security Intelligence data.

Category: A class or division of events with shared characteristics as defined by the classifier.

Anomaly: Something that deviates from what is standard, normal, or expected when compared to the user's selected baseline.

Anomaly Definition: A set of principles, configurable by a user, that describes the threshold and circumstances as defined in the anomaly rules.

Chart: The graph in the dashboard that depicts the statistical data and baseline.


Time Slider: The Web interface component that allows a user to quickly visualize and navigate the time range of data.

Breakdown: The Web interface component that displays the top ten values of a select number of event fields to show a user details on the most prevalent event values in the data they are viewing.

7.1.2 How Security Intelligence Works

The following diagram depicts how the Security Intelligence feature works.

Figure 7-1 How Security Intelligence Works



1. A user creates a dashboard in the Sentinel Web interface.
2. The dashboard consists of a filter and classifier that create a segment of data from the normalized events in the message bus.
3. The segment of data is stored in the MongoDB database for up to 28 days.
4. The user then defines a baseline of normal activity for the environment.
5. The user creates anomaly definitions to look for real-time events that are occurring outside of the defined baseline.
6. The anomaly engine analyzes the events in the MongoDB database.

7. This information is displayed in the dashboard.
8. If an anomaly is detected, an anomaly event is generated (e-mail can be sent with the information about the anomaly).

7.1.3 Permissions for Security Intelligence

The Security Intelligence option is displayed in the Web interface if the user has one of the following permissions:

- ♦ Manage and View Security Intelligence Dashboards
- ♦ View Security Intelligence Dashboards

A user is assigned these permissions while creating a role. For more information, see “[Creating Roles](#)” in the *NetIQ Sentinel Administration Guide*.

Using the Security Intelligence option, you can view, create, and manage dashboards.

7.2 Creating a Dashboard

- 1 Log in to the Sentinel Web interface as a user with the Manage and View Security Intelligence Dashboards permission.

- 2 On the left side of the page, click **Security Intelligence > Dashboards > Create**.

The Create Dashboard page opens in a new tab.

- 3 Use the following information to create the dashboard:

Name: Specify a unique name for the dashboard.

Classifier: Select the classifier that determines the categories displayed in the dashboard. The options are:

- ♦ **Taxonomy Outcome**
- ♦ **Device Activity**
- ♦ **Tag Activity**
- ♦ **Http**
- ♦ **Exploit**

Filter: Specify a filter to determine the scope of events displayed in the dashboard, or select a predefined filter. By default it displays the **(sev:[0 TO 5])** filter.

To search for an event field, specify the short name of the field, a colon, and the value. For example, `notnull: xdstaxname` displays all events. For more information, see [Chapter 4, “Configuring Filters,”](#) on page 49.

Data retention period: Select how long the data for the dashboards should be retained.

By default the Security Intelligence MongoDB database retains the data for 4 weeks.

Load historical data: Select this option to view historical data, which provides more context when analyzing the data.

- 4 Click **Create dashboard**.

The newly created empty dashboard is displayed because it has not had time to collect any data.

After few minutes, you can see the event data in the dashboard.

7.2.1 Creating a Dashboard by Using a Filter

You can use a filter search query as a dashboard filter and create a dashboard. For more information on creating a dashboard by using a filter, see [Section 3.4.6, “Creating a Dashboard,”](#) on page 44.

7.3 Understanding the Dashboard Interface

The dashboard displays the analysis of the data.

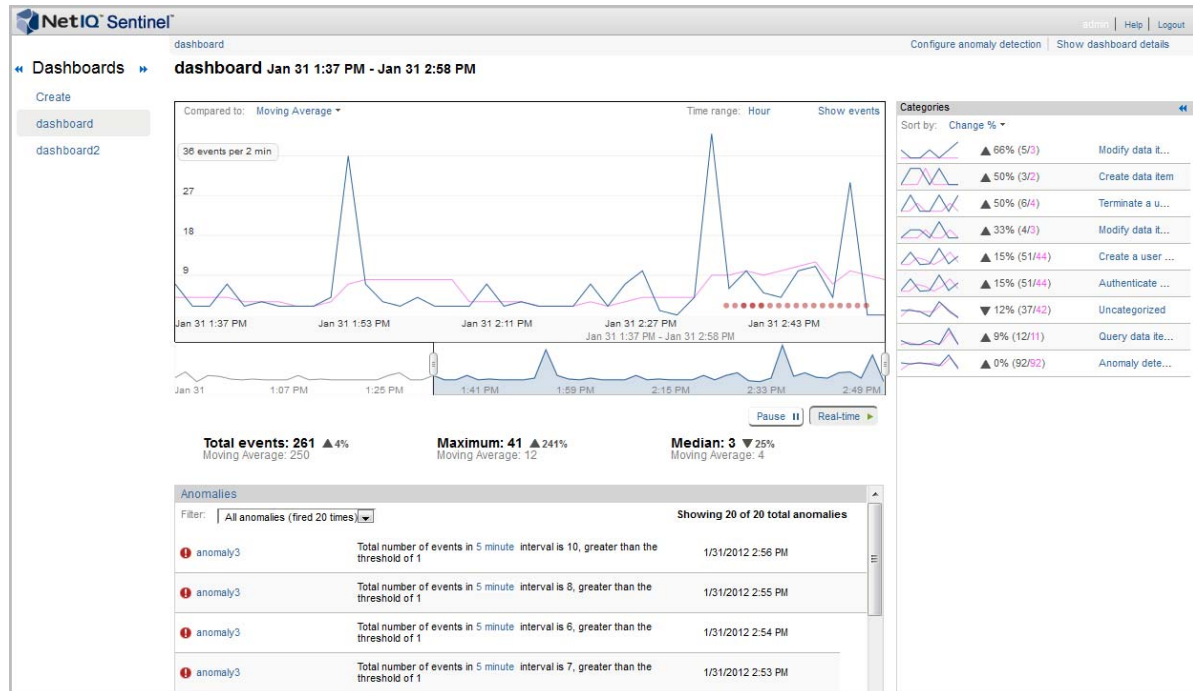


Chart: The graph displays the events, the anomalies, and the baseline.

Compared to: Displays the types of baselines for comparison with the flowing data. Baseline is the referenced line that is displayed in blue color, whereas the actual events that are flowing in to the system is displayed in brown. You can compare the flowing data to the following:

- ◆ **Moving Average:** Data that is flowing in to the system is compared to the average of the data.
- ◆ **Previous Day:** If you have one day data stored in Security Intelligence database (MongoDB), then you can compare the flowing data with the previous day data.
- ◆ **Previous Week:** If you have one week of data stored in Security Intelligence database (MongoDB), then you can compare the flowing data with the previous week data.
- ◆ **Create Baseline:** Allows you to create a custom baseline. You must have at least a week's worth of data before you can create a baseline. For more information on creating a custom baseline, see [Section 7.4, “Creating Baselines,”](#) on page 98.

Time range: Displays the time range between which we can see the data. When you create a dashboard, the Time range shows **Hour**. It then displays **Day** and **Week** as you have one day or one week data stored in Security Intelligence database (MongoDB). The Hour option does not appear, if the Security Intelligence database (MongoDB) have 28 days of data.

Show Events: Displays the list of events for the selected time range in the Sentinel Web interface. The total events in the Sentinel Web interface matches with the total events in the dashboard. However, the total events count in the Sentinel Web interface does not match with the total events in the dashboard in the following cases:

- ◆ If the leftmost and rightmost time point on the dashboard are included. To match the events in the dashboard and Sentinel Web interface, select the time point between the leftmost and rightmost time point on the dashboard.
- ◆ By default, Sentinel includes internal audit events in the dashboard search results. To exclude the internal events, create a dashboard by using the filter `(sev:[0 TO 5]) -st:A -st:I -st:Y`. For more information, see [Section 7.2.1, “Creating a Dashboard by Using a Filter,” on page 97](#).

Time Slider: The time slider allows you to change the amount of information displayed in the dashboard. It allows you to zoom in or zoom out for a specific time period. As you move the time slider, the graph changes accordingly.

Time Slider Data Summaries: Below the graph, a summary of the time slider data is displayed. The data that is

Anomalies: Displays the anomalies that have occurred during the lifetime of the dashboard. To view the details of the anomaly, you can click on anomaly name. This displays the anomaly detail page.

Categories: The **Categories** panel on the right of the dashboard displays the categories of the current time range at the current level of the dashboard. It provides the ability to drill down and find more information about the categories of events. This section displays lines identifying changes from the baseline indicators of the categories. You can sort the category list by percent change, reference count or current count.

Clicking a specific category in the list on the right displays the data for just that category. It changes the main graph to show the events in that category list. The totals in the main section changes to reflect the current category. It also displays the following sections in the bottom of the main panel.

- ◆ **Category anomalies:** Displays anomalies happened in the current time window for the selected category.
- ◆ **Category breakdown:** Displays the attributes of the selected category. Only top 10 values of the selected category are displayed in the UI. You can click any value under the **Top 10 Values** list to view the list of events in the Sentinel Web interface.

The **Categories** list on the right changes to **Subcategories** and displays attributes of the selected category. You can sort them as per your requirement.

7.4 Creating Baselines

You can create a baseline to use for anomaly detection. You must have at least a week’s worth of data before you can create a baseline.

- 1 Click the desired dashboard under the dashboard heading.
- 2 Click the option displayed in the **Compare to** field.
or
Click **Show dashboard details**.
- 3 Click **Create baseline**.

- 4 Read the confirmation message, then click **Create**.
- 5 After the baseline is created if you want to update the baseline with more amount of data then you can click **Regenerate baseline**.

After the baseline is created, you can use it to create an anomaly definition as described in [Section 7.5, “Configuring Anomaly Detection,” on page 99](#).

7.5 Configuring Anomaly Detection

After you create a baseline, you can configure anomalies to use with the information gathered in the dashboard. This allows you to receive alerts when events occur outside of the baseline.

- ♦ [Section 7.5.1, “Creating an Anomaly Definition,” on page 99](#)
- ♦ [Section 7.5.2, “Deploying an Anomaly Definition,” on page 100](#)
- ♦ [Section 7.5.3, “Undeploying an Anomaly Definition,” on page 100](#)
- ♦ [Section 7.5.4, “Managing Anomalies,” on page 101](#)

7.5.1 Creating an Anomaly Definition

- 1 Log in to the Sentinel Web interface as a user with the Manage Dashboard permission.
- 2 Click the desired dashboard under the dashboard heading, then click **Configure anomaly detection**.
- 3 Click **Create anomaly definition**.

The Anomaly detection definition details screen is displayed.

- 4 Use the following information to create the anomaly definition:

Anomaly name: Specify a unique name for the anomaly.

Comparison type: Select and define the anomaly type. The options are:

- ♦ **Threshold:** When the number of a specific type of events exceeds a specified limit, Sentinel triggers an anomaly event. For example, if you set the threshold for login failures to five and if more than five failed logins occur, Sentinel triggers an anomaly event.
- ♦ **Moving Average:** Moving averages are calculated over a specific period of time. All averages in that period are recalculated to remove noise and deviations which results in the moving average. Sentinel triggers an anomaly event if the moving average deviates from the normal averages. For example, in the holiday seasons, the internet traffic for e-commerce websites might spike which might result in abnormal average compared to the rest of the year. Sentinel triggers an anomaly event indicating the deviation in the moving average.
- ♦ **Ratio:** Provides a comparison between the number of different types of events. If the ratio of a specific event type compared to the other type exceeds beyond a specified limit, Sentinel triggers an anomaly event. For example, if a significant number of events were reported for viruses as compared to network attacks.
- ♦ **Historical:** Provides a comparison of the number of current events with the events received in the past. For Example, if the historical data reports the number of invalid logins per day in the range of 100-150 and if the current number of invalid logins is 1000, Sentinel triggers an anomaly event.
- ♦ **Baseline:** Provides a comparison to an established baseline. A baseline is usually the accepted or agreed upon values of event data. You must have a custom baseline to use this option. For more information, [Section 7.4, “Creating Baselines,” on page 98](#). If there is a

deviation from the baseline, Sentinel triggers an anomaly event. For example, if the event stream baseline is 1000 per second and if the event stream rate increases or decreases, Sentinel triggers an anomaly event.

As per your requirement, you can select the `Comparison` type and specify the anomaly definition.

Anomaly description: Specify a description for the anomaly. The description is displayed in the anomaly event.

Anomaly state: Define the state of the anomaly by selecting any one of the following:

- ♦ **Always active:** You can use this option to keep the anomaly definition active always and trigger when the specified anomaly definition is met.
- ♦ **Only active for selected days and times:** You can use this option to define the specific times for anomaly definition to trigger. When you select this option, it displays a default time grid. You can change the time grid and specify different time periods for the same anomaly definition by holding the Ctrl key.

NOTE: The timing that is displayed in the time grid is the local time.

Notification information: Select the information to define the notification information.

- ♦ **Severity:** Select the severity of the notification. The options are 0 to 5.
- ♦ **After this anomaly definition fires:** Specify the notification time gap to send e-mail or events after an anomaly triggers.

Optionally send notification via e-mail after the anomaly triggers: Fill in the following fields to send an e-mail when the anomaly triggers.

- ♦ **E-mail address:** Specify the e-mail addresses of the people who should receive notification when the anomaly occurs. Separate multiple e-mail addresses with commas.
- ♦ **Subject:** Specify a subject for the e-mail.
- ♦ **Message:** Specify a message for the e-mail to explain the anomaly that occurred.

5 Click **Save**.

6 Continue with [“Deploying an Anomaly Definition”](#) on page 100.

7.5.2 Deploying an Anomaly Definition

After the anomaly definition is created, it must be deployed to be applied to the dashboard.

- 1 In the Sentinel Web interface, click **Security Intelligence > Dashboard**, then select the dashboard where you created the anomaly definition.
- 2 Click **Configure anomaly detection**.
The Anomaly detect screen is displayed.
- 3 Mouse over the anomaly definition you want to deploy, then click **Deploy**.
You receive a message that the anomaly definition was deployed.

7.5.3 Undeploying an Anomaly Definition

To undeploy the anomaly definition:

- 1 In the Sentinel Web interface, click **Security Intelligence > Dashboard**, then select the dashboard where you created the anomaly definition.
- 2 Click **Configure anomaly detection**.

- 3 Mouse over the anomaly definition you want to undeploy, then click **Undeploy**.
- 4 Click **Undeploy** again to verify that you want to perform this action.
You receive a message that the anomaly definition was undeployed.

7.5.4 Managing Anomalies

You can perform the following management tasks on the anomalies:

- ♦ [“Editing an Anomaly” on page 101](#)
- ♦ [“Deleting an Anomaly” on page 101](#)

Editing an Anomaly

- 1 In the Sentinel Web interface, click **Security Intelligence > Dashboard**, then select the dashboard where you created the anomaly definition.
- 2 Click **Configure anomaly detection**.
- 3 Mouse over the anomaly you want to edit, then click **Edit**.
- 4 Make any desired changes to the anomaly definition, then click **Save**.


Deleting an Anomaly

- 1 In the Sentinel Web interface, click **Security Intelligence > Dashboard**, then select the dashboard where you created the anomaly definition.
- 2 Click **Configure anomaly detection**.
- 3 Mouse over the anomaly you want to delete, then click **Delete**.
- 4 Click **Delete** again to verify that you want to perform this action.

7.6 Viewing Anomaly Events

When an anomaly is detected, Sentinel generates an anomaly event. Anomaly event fields contain detailed information about the anomaly.

To view the anomaly events:

- 1 In the Sentinel Web Console, in the left pane, expand **Filters > My filters**, click **Anomaly Events**, click  .
- 2 To view the event field values for an anomaly event, in the search results, click **All** next to the anomaly event.

The following table describes the various event fields in an anomaly event:

Anomaly Event Field	ID	Sample Value	Description
BeginTime	bgnt	2014-01-06T07:13:00.000Z	The start of the time range when the anomaly was detected.

Anomaly Event Field	ID	Sample Value	Description
EndTime	endt	2014-01-06T07:17:00.000Z	The end of the time range when the anomaly was detected.
EventName	evt	FailedLogins:AbnormalFailedLogins	The name of the anomaly definition.
EventTime	dt	2014-01-06T07:18:54.285Z	The time when the anomaly event was generated.
Message	msg	abnormal failed login activity	The description in the anomaly definition.
ObserverCategory	rv32	SIEM	For an anomaly event, this event field is always set to SIEM.
ObserverServiceComponent	rv150	/Create a user session/ Failure	The classifier path which contains the categories displayed in the dashboard.
ObserverTZ	estz	Asia/Kolkata	The time zone in which the anomaly engine is located.
ObserverType	st	Y	For an anomaly event, the event field is always set to Y.
SentinelProcessingComponent	rt2	AbnormalFailedLogins	The anomaly definition name.
SentinelProcessingComponentID	rv123	2F38BBCA-1A39-42A9-9873-D2C4CE732B0D	This is the UUID of the dashboard which is associated with the anomaly definition. The UUID remains the same even though the dashboard name changes.
SentinelServiceComponentID	rv124	B7E6B2A7-CDB1-40A8-AA33-8AE99284DE6B	This is the ID of the anomaly definition. The ID remains the same even though the anomaly definition name changes.
SentinelServiceComponentName	sres	FailedLogins	This is the dashboard name associated with the anomaly definition.
SentinelServiceName	res	SecurityIntelligence	For an anomaly event, this event field is always set to SecurityIntelligence.
Severity	sev	5	The severity in the anomaly definition.
XDASClass	xdasclas s	11	For an anomaly event, this event field is always set to 11.
XDASDetail	xdasdeta il	12	For an anomaly event, this event field is always set to 12.

Anomaly Event Field	ID	Sample Value	Description
XDASIdentifier	xdasid	13	For an anomaly event, this event field is always set to 13.
XDASOutcome	xdasoutcome	1	For an anomaly event, this event field is always set to 1.
XDASOutcomeName	xdasoutcomeName	XDAS_OUT_THRESHOLD_EXCEEDED	For an anomaly event, this event field is always set to XDAS_OUT_THRESHOLD_EXCEEDED.
XDASProvider	xdasprov	0	For an anomaly event, this event field is always set to 0.
XDASRegistry	xdasreg	0	For an anomaly event, this event field is always set to 0.
XDATaxonomyName	xdastaxname	XDAS_AE_ANOMALY	For an anomaly event, this event field is always set to XDAS_AE_ANOMALY.

For more information on anomaly event fields, click **Tips** in the Sentinel Web Console. For more information on the event taxonomy and event fields, see [Sentinel Taxonomy](#).

You can use the event field IDs to create search queries to find specific anomaly events. For example, if you want to search for the anomaly events that were generated because of the anomaly definition `AbnormalFailedLogins`, specify the following query in the **Search** field:

```
st:y AND rt2:AbnormalFailedLogins
```

For more information about searching for events, see [Section 3.1, “Running an Event Search,” on page 33](#).

7.7 Managing Dashboards

You cannot change the filter or classifier for an existing dashboard, because this changes all of the data. If you want to change the filter or classifier, you must create a new dashboard.

However, you can perform the following management tasks on dashboards:

- ◆ [Section 7.7.1, “Viewing a Dashboard,” on page 103](#)
- ◆ [Section 7.7.2, “Renaming a Dashboard,” on page 104](#)
- ◆ [Section 7.7.3, “Deleting a Dashboard,” on page 104](#)

7.7.1 Viewing a Dashboard

- 1 In the Sentinel Web interface, click **Security Intelligence > Dashboard**, then select the dashboard that you want to view.
- 2 Click **Show dashboard details** to see the following information about the dashboard:
 - ◆ The classifiers used to create the dashboard.
 - ◆ The filter used to create the dashboard.

- ♦ When the dashboard was created.
- ♦ The amount of time data is retained for the dashboard.

7.7.2 Renaming a Dashboard

- 1 In the Sentinel Web interface, click **Security Intelligence** > **Dashboard**, then select the dashboard that you want to rename.
- 2 Click **Show dashboard details** in the toolbar.
- 3 Click **Rename** in the toolbar, then rename the dashboard.
- 4 Click **Save** to save the change.

7.7.3 Deleting a Dashboard

- 1 In the Sentinel Web interface, click **Security Intelligence** > **Dashboard**, then select the dashboard that you want to delete.
- 2 Click **Show dashboard details** in the toolbar.
- 3 Click **Delete**.
- 4 Click **Delete** again to verify that you want to perform this action.

7.8 Troubleshooting

- ♦ [Section 7.8.1, “The Create Button Is Not Displayed,” on page 104](#)
- ♦ [Section 7.8.2, “The Main Graph and the Time Slider Are Not Synchronized,” on page 104](#)
- ♦ [Section 7.8.3, “Both Names for a Renamed Anomaly Are Displayed in the Filter,” on page 104](#)
- ♦ [Section 7.8.4, “Dashboard Date Range Not Updated to in Real Time,” on page 105](#)

7.8.1 The Create Button Is Not Displayed

If you access the Security Intelligence feature and the **Create** button is not displayed, the MongoDB database is not running. The solution is to restart the Sentinel system.

7.8.2 The Main Graph and the Time Slider Are Not Synchronized

The main graph and the time slider can display different data because of the potential differences in granularity.

7.8.3 Both Names for a Renamed Anomaly Are Displayed in the Filter

If an anomaly definition display name changes but the anomaly was fired in both the old name and the new name within the selected time range, the anomaly shows two filters: the old anomaly display name with its firing count and new anomaly display name with its firing count.

If you filter on the new anomaly, only the new anomaly name is shown in the list, but the Show x of y anomaly message shows x as the total count of the anomaly fired under the new name and y as the total count of the anomaly fired in both the new and old display names.

The x will never equal y if there are events fired in both names within selected date range.

For example, assume that there is an anomaly definition of DemoDef that was renamed to DemoDef-nameChanged within the selected time range and it has fired under the old name 60 times and under the new name 180 times. In the filter drop-down list, both anomalies are displayed, showing DemoDef (fired 60 times) and DemoDef-nameChanged (fired 180 times) for a total of 240. If you filter on DemoDef, the filter message displays, Showing 60 of 240 anomalies. If you filter on DemoDef-nameChanged, the filter message displays, Showing 180 of 240 anomalies.

7.8.4 Dashboard Date Range Not Updated to in Real Time

When the dashboard is bigger, the time granularity is bigger, so it page refreshes slowly. In another words, the date range for the bigger display takes longer to show a difference. Even though the auto refresh timer is the same, it might need to fire two times before you see any change in data.

8 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 real-time for a specific department, for a specific customer, a specific event, an IP address, or a time range.

To view and analyze the network flow data, you must first configure Sentinel for network flow data collection. To configure network flow data collection, you must either be an administrator or have the Send NetFlow data permission. For more information about configuring network flow data collection, see “[Visualizing Network Traffic](#)” in the *NetIQ Sentinel Administration Guide*.

NOTE: To view the network flow data, you must have the View NetFlow data permission.

To view the NetFlow data:

- 1 Log in to the Sentinel Web interface as a user with the View NetFlow data permission.
- 2 Perform either of the following:
 - ♦ **To view real-time network flow data for a specific department or a tenant:**
 1. In the navigation panel, click **Real-time Monitoring > NetFlow > Create**.
 2. Specify the following information:
 - Tenant:** If you are in a multi-tenant environment, select the department or the tenant name for which you want to view network flow data. Otherwise, select **Default**.
 - Address:** Specify the IP address for which you want to view the network flow data. By default, Sentinel provides the network flow data for the entire network.
 - Time Range:** Specify the time range for which you want to view the network flow data.
 3. Click **Monitor NetFlow**.
 - ♦ **To view network flow data for a specific event:**
 1. Perform a search to view the desired events.
 2. In the search results, click the NetFlow icon for the Source IP address or the Destination IP address of the event.

Sentinel provides a graphical representation the network flow data, which automatically refreshes every 10 minutes. You can hover the mouse over the graphs to monitor the incoming and outgoing number of bytes, packets, and flows for the specified tenant or IP address. Sentinel also helps you analyze the network traffic by providing a summary of the top 10 hosts and top 10 ports sending data to the specified tenant or IP address. Similarly, you can view the top 10 hosts and top 10 ports to which the specified tenant or IP address is sending data.

If you want to analyze the data in detail for a specific time, you can take a snapshot of the graph and analyze the data at the desired time. You can also compare the changes in the network traffic before and after the security event.

9 Configuring Dynamic Lists

Dynamic Lists are distributed list structures that can be used to store string elements, such as IP addresses, server names, or usernames. The lists are then used within a Correlation rule for a quick lookup to see whether an incoming event includes an element from the Dynamic List. Some examples of Dynamic Lists include:

- ♦ Terminated user lists
- ♦ Suspicious user watchlist
- ♦ Privileged user watchlist
- ♦ Authorized ports and services list
- ♦ Authorized server list

NOTE: You must have the Manage Correlation Engine and Rules permission to create and manage Dynamic Lists.

9.1 Creating a Dynamic List

A Dynamic List can be built using the text values for any event ID. Elements can be added to the list manually or automatically whenever a Correlation rule fires.

Regardless of how the values were added, an element can be of the following types:

- ♦ **Persistent:** The element is active until it is manually removed or until the maximum list size is reached.
- ♦ **Transient:** The element is active only for a specified time after being added to the list.

Dynamic Lists can be created either in the Sentinel Control Center or in the Correlation Rule Builder:

- ♦ [Section 9.1.1, “Using the Sentinel Control Center to Create a Dynamic List,” on page 109](#)
- ♦ [Section 9.1.2, “Using the Correlation Rule Builder to Create a Dynamic List,” on page 110](#)

9.1.1 Using the Sentinel Control Center to Create a Dynamic List

1 Launch the Sentinel Control Center.

1a Log in to 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.

1b In the tool bar, click **Applications**.

1c Click **Launch Control Center**.

1d Click **Yes** to accept the security certificate.

- 1e Specify a username and password of a user that has rights to access the SCC, then click **Login**.
- 1f Click **Accept** or **Accept Permanently** to accept the security certificate and display the SCC.
- 2 Launch the Dynamic Lists window:
 - ♦ (Conditional) If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
 - ♦ (Conditional) If the **Configuration** menu is enabled, click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
- 3 Click **Add**.
- 4 Specify a name for the Dynamic List.

The name must start with a character. The name can contain only letters, digits, or underscores. The name cannot be changed after you create the Dynamic List. Therefore, specify a descriptive name.
- 5 To add elements, click **Add**.
- 6 Specify a name for the list element.
- 7 To keep the element active until it is manually removed or until the maximum list size is reached, select **Make Persistent**, then click **OK**.

or

To keep the element active only for a specific time, use the **Transient elements life span** fields to specify how long the element remains active.

The time period can range from minutes to 90 days.
- 8 Specify the maximum number of elements you want in the Dynamic List.

The maximum list size can be 100,000.
- 9 Click **OK**.

9.1.2 Using the Correlation Rule Builder to Create a Dynamic List

You can create Dynamic Lists while creating a Correlation rule. This option is provided in the Correlation Rule Builder to help you complete the rule creation process without switching to the Sentinel Control Center, and also if you want to just create an empty Dynamic List.

- 1 Log in to 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.
- 2 Select **Correlation** from the navigation panel.
- 3 In the Subrule window, click **Create a new expression**.
- 4 In the **Expression Builder**, select an appropriate event field from **Attributes**.
- 5 In the **Operator** list, select **inlist** or **not inlist**.
- 6 In the **Value** section, click **Create**.
- 7 Specify the following information for the list:
 - ♦ **List name:** A descriptive name for the Dynamic List. The name must start with a character. The name can contain only letters, digits, or underscores.

- ♦ **Transient elements life span:** The time for the element to remain active. The time can range from 1 hour to 90 days.
- ♦ **Maximum number of elements:** The maximum number of elements the list should include.



8 Click **OK**.

The Dynamic List is created. However, you must launch the Sentinel Control Center to add elements to the list. You can complete the Correlation rule creation process, then add elements to the list. For more information on adding elements, see [Section 9.1.1, “Using the Sentinel Control Center to Create a Dynamic List,” on page 109](#).

9.2 Managing Dynamic Lists

- ♦ [Section 9.2.1, “Editing a Dynamic List,” on page 111](#)
- ♦ [Section 9.2.2, “Deleting a Dynamic List,” on page 111](#)
- ♦ [Section 9.2.3, “Removing Dynamic List Elements,” on page 112](#)

9.2.1 Editing a Dynamic List



- 1 Launch the Sentinel Control Center. For more information, see [Step 1 in Section 9.1.1, “Using the Sentinel Control Center to Create a Dynamic List,” on page 109](#).
- 2 Launch the Dynamic Lists window:
 - ♦ (Conditional) If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
 - ♦ (Conditional) If the **Configuration** menu is enabled, click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
- 3 In the Dynamic Lists window, select the dynamic list you want to edit, then click **View/Edit**.
- 4 Make the necessary changes, then click **OK**.

The name cannot be changed.

NOTE: If you make changes to the transient life span on an active Dynamic List, the changes do not apply to elements already in the list. Elements that are already in the Dynamic List retain their original life span.

9.2.2 Deleting a Dynamic List

You can delete only the dynamic lists that are not being used by Correlation rules or actions:

- 1 Launch the Sentinel Control Center. For more information, see [Step 1 in Section 9.1.1, “Using the Sentinel Control Center to Create a Dynamic List,” on page 109](#).
- 2 Launch the Dynamic Lists window:
 - ♦ (Conditional) If the **Configuration** menu is not enabled, click the **Configuration** tab, then click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
 - ♦ (Conditional) If the **Configuration** menu is enabled, click the **Configuration** menu > **Dynamic Lists** or click the  icon in the toolbar.
- 3 In the Dynamic Lists window, select the dynamic list you want to delete, then click **Delete**.
- 4 Click **Yes** to confirm deletion.

9.2.3 Removing Dynamic List Elements

There are several ways an element can be removed from a Dynamic List:

- ♦ The element can be removed manually:
In the Dynamic List Properties window, select the element you want to delete, then click **Delete**.
- ♦ The element can be removed by a Correlation rule action.
- ♦ The element's life span can expire.
- ♦ If the maximum number of elements for a Dynamic List is reached, elements are removed from the list to keep the list at or below the maximum list size. The transient elements are removed from oldest to newest before any persistent elements are removed.

10 Leveraging Identity Information

This section provides information about integrating Sentinel with identity management systems.

- ♦ [Section 10.1, “Overview,” on page 113](#)
- ♦ [Section 10.2, “Searching and Viewing User Identities,” on page 113](#)

10.1 Overview

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.

NOTE: Only administrators can integrate Sentinel with identity management systems. For more information, see [“Integrating Identity Information”](#) in the *NetIQ Sentinel Administration Guide*.

10.2 Searching and Viewing User Identities

The Identity Browser in Sentinel allows you to search and view user profiles of the identities in the Sentinel database that have been synchronized from the identity management system. In addition to information from the identity management system, the Identity Browser also shows recent user activity that has been collected through the Sentinel Collectors.

10.2.1 Accessing the Identity Browser

- 1 Log in to the Sentinel Web interface.
- 2 Click **People** on the left side.
The Identity Browser is displayed.

10.2.2 Performing a Search

The Identity Browser allows you to search for people to view what they have been doing. You can use the search box or click the arrow next to the search box for more options. As you start typing the information in the search field, the data is automatically displayed.

- 1 Access the Identity Browser.

For more information, see [Section 10.2.1, “Accessing the Identity Browser,” on page 113](#).

- 2 Search for a user by typing in the search box. The search box is dynamic
or

Click the arrow next to the search box to display more search fields.

You can type letters to view all identities whose first or last name starts with the letters. For example, if you type the names Abraham, Abdullah, and so on are matched. For more information, see [Section 10.2.3, “Searching,” on page 114](#).

- 3 If you used the search box, skip to [Step 5](#). Otherwise, specify the search value for the users you are searching for.

For more information, about the search fields, see [“Using the Search Fields” on page 115](#).

- 4 Click **Search** to perform the search, then click **Close**.

- 5 Click on the user name to view the information about the user.

- 6 Proceed to [Section 10.2.4, “Viewing Profile Details,” on page 115](#) to view the details about the user.

10.2.3 Searching

You can search for users by using the search box or by using the search fields.

- ♦ [“Using the Search Box” on page 114](#)
- ♦ [“Using the Search Fields” on page 115](#)

Using the Search Box

The search box automatically uses the following logic to interpret the text you enter:

- ♦ All letters and no spaces searches for the given name or surname.
- ♦ All letters and a space between letter groups searches for the given name and surname. The surname match is a starts-with, unless there is a trailing space.
- ♦ All letters with a comma in the middle is a match of the surname and given name. The given name match is starts-with unless there is a trailing space.
- ♦ Anything with a @ in it is a starts-with match for e-mail address.
- ♦ All digits, or letters and digits but no telephone punctuation characters is a starts-with match for workforce ID.
- ♦ Digits in addition to a leading +, and spaces, hyphens, periods, or parentheses is a starts-with match for a telephone number.
- ♦ Alphanumeric, or all numeric with no spaces, or all numbers with spaces is a starts-with match for the workforce ID.

Using the Search Fields


You can search for many values, including custom values, in the search fields. The following is a list of the fields you can search:

- ◆ Given Name
- ◆ Surname
- ◆ Telephone
- ◆ E-mail
- ◆ Position
- ◆ Department
- ◆ Office Location Code
- ◆ Workforce ID
- ◆ Vault Name
- ◆ Customer ID
- ◆ DN
- ◆ Custom Value Name
- ◆ Custom Value

10.2.4 Viewing Profile Details

After you have performed the search, (see [Section 10.2.2, “Performing a Search,”](#) on page 114), the user name, photo, position, department, e-mail, and telephone number are displayed.

Figure 10-1 User Information Displayed



Click **Profile** to see detailed information about the user and all of the accounts that belong to this user.

You can use the clipboard functionality to copy the data of the user's profile and account information. Click the clipboard icon to the left of the user's photo and their information is now in the clipboard. You can paste this information into a text editor.

10.2.5 Viewing Activity

You can view the recent activity of the user through the Identity Browser.

- ◆ Authentication information
- ◆ Access events
- ◆ Permission changes

Select one or more of these options, then click **Show Recent Activity**. The activity is displayed in the search panel of the Sentinel Web interface.

11 Manually Performing Actions on Events

Administrators can configure event routing rules to automatically perform specific actions on events, as described in “[Configuring Event Routing Rules](#)” in the *NetIQ Sentinel Administration Guide*. However, Sentinel allows users to manually perform actions on events returned in searches.

This allows users to perform the desired actions as they are viewing events.

- ◆ [Section 11.1, “Accessing Event Actions,” on page 117](#)
- ◆ [Section 11.2, “Prerequisites for Assigning Actions to Events,” on page 117](#)
- ◆ [Section 11.3, “Assigning Actions to Events,” on page 117](#)
- ◆ [Section 11.4, “Configuring Event Actions,” on page 118](#)

11.1 Accessing Event Actions

To access the event actions:

- 1 Log in to the Sentinel Web interface as a user who is a member of one of the following roles:
 - ◆ Administrator
 - ◆ Incident Administrator
 - ◆ Security Policy Administrator
 - ◆ User
- 2 Click **Event Actions** on the left.

If **Event Actions** is not displayed, you do not have the proper rights to access this feature.

11.2 Prerequisites for Assigning Actions to Events

In order to manually assign actions to select events, you must perform the following tasks:

- Configure the Integrators for the Actions:** There are default actions available to perform. However, the integrators for these actions must be configured before the actions can be executed. For more information, see “[Configuring the Default Integrators](#)” in the *NetIQ Sentinel Administration Guide*.
- Perform a Search:** You must have the results of a search available in order to perform actions on selected events. For more information about searching, see [Chapter 3, “Searching Events,” on page 33](#).

11.3 Assigning Actions to Events

- 1 Access **Event Actions** in the Sentinel Web interface.
For more information, see [Section 11.1, “Accessing Event Actions,” on page 117](#).
- 2 Perform a search.
If you have not performed a search, you cannot execute an action.

- 3 Select the events in the search to perform actions on.
- 4 In the **Actions** field, select the desired action from the drop-down box.
You can add or remove items from this list. For more information, see [Section 11.4, “Configuring Event Actions,” on page 118](#).
- 5 Click **Execute**.
- 6 In the **Results** field, view the results of the action.

11.4 Configuring Event Actions

Administrators can control what actions can manually performed on events.

- ♦ [Section 11.4.1, “Creating a New Event Action,” on page 118](#)
- ♦ [Section 11.4.2, “Cloning an Event Action,” on page 119](#)
- ♦ [Section 11.4.3, “Moving an Event Action,” on page 119](#)
- ♦ [Section 11.4.4, “Deleting an Event Action,” on page 119](#)

11.4.1 Creating a New Event Action

If you have imported new action plug-ins into Sentinel and you want to create an event action for the action plug-in:

- 1 Access the Sentinel Control Center.

- 1a Log in to 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.

- 1b In the tool bar, click **Applications**.
 - 1c Click **Launch Control Center**.
 - 1d Click **Yes** to accept the security certificate.
 - 1e Specify a username and password of a user that has rights to access the SCC, then click **Login**.
 - 1f Click **Accept** or **Accept Permanently** to accept the security certificate and display the SCC.
- 2 Click the **Configuration** tab.
 - 3 From the menu, click **Configuration > Event Actions Configuration**.
 - 4 Click **Add** to add a new action.
 - 5 Use the following information to create the new event action:
Name: Specify a unique name for the event action.
Description: Specify a description for the new event action.
Action: Select the desired action from the drop-down list.
 - 6 Click **Add Action**.
 - 7 Create a new action by following the instructions in [“Adding an Action”](#) in the *NetIQ Sentinel Administration Guide*.

11.4.2 Cloning an Event Action

You can clone an existing event action and give it another name.

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 11.4.1, "Creating a New Event Action," on page 118.](#)
- 2 Click the **Configuration** tab.
- 3 On the menu, click **Configuration** > Event Actions Configuration.
- 4 Select an event action, then click **Clone**.
- 5 Change the name to a unique name, then click **OK**.

11.4.3 Moving an Event Action

You can display the event actions in any order. The order in the Event Actions Configuration page is the order that is shown in the Sentinel Web interface.

To move event actions:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 11.4.1, "Creating a New Event Action," on page 118.](#)
- 2 Click the **Configuration** tab.
- 3 On the menu, click **Configuration** > **Event Actions Configuration**.
- 4 Select the event action you want to move, then click **Up** or **Down** until the event action is in the correct location.

11.4.4 Deleting an Event Action

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 11.4.1, "Creating a New Event Action," on page 118.](#)
- 2 Click the **Configuration** tab.
- 3 From the menu, click **Configuration** > **Event Actions Configuration**.
- 4 Select the event action you want to delete, then click **Delete**.
- 5 Click **Yes** to confirm the deletion.

12 Configuring Tags

Tags are user-defined values that can be used to logically group data collection objects such as event sources, event source servers, Collector Managers, Collector plug-ins, event routing rules, report templates, and report results. Tags help you to filter object lists for the data collection objects and also to augment incoming data. You can search for events, report templates, and report definitions that are tagged with a particular tag.

NOTE: Only users in the Manage Tags role can create and manage tags.

- ◆ [Section 12.1, “Overview,” on page 121](#)
- ◆ [Section 12.2, “The Tags Interface,” on page 121](#)
- ◆ [Section 12.3, “Creating a Tag,” on page 122](#)
- ◆ [Section 12.4, “Managing Tags,” on page 122](#)
- ◆ [Section 12.5, “Performing Text Searches for Tags,” on page 123](#)
- ◆ [Section 12.6, “Deleting Tags,” on page 123](#)
- ◆ [Section 12.7, “Associating Tags with Objects,” on page 124](#)
- ◆ [Section 12.8, “Viewing Tagged Events,” on page 125](#)

12.1 Overview

You can associate objects with more than one tag. You can, for example, create tags related to regulations (PCI) or compromised systems or network infrastructure such as routers, switches, and firewalls. Some organizations need to define data retention or data viewing policies based on the geographic location, so tags can be used to tag event sources based on different locations.



When ESM objects such as event sources, event servers, Collector Managers, or Collector plug-ins are tagged, all the events from those ESM objects are tagged with that value. The tag value is placed in a reserved variable, `rv145`. However, events generated before tagging the ESM objects are not tagged. Sentinel does not perform retroactive tagging of data that is already stored because it is not an accepted practice to modify events that are already stored.

You must have the appropriate permission to view events that are tagged with specific tags. For example, only users in the PCI Compliance Auditor role can view events that are tagged with at least one of the regulation-related tags such as PCI, SOX, HIPAA, NERC_CIP, FISMA, GLBA, NISPOM, JSOX, and ISO/IEC_27002:2005.

12.2 The Tags Interface

The Tags interface lists the tags available in the system and allows you to manage the tags. You can perform text-refined searches to find the tags that you are looking for. The interface also provides options such as maintaining a list of favorite tags and searching tagged events.

As you mouse over a tag, you can see the icons available to manage the tag. The number next to each tag indicates the number of objects associated with the tag. For more information on creating new tags, see [Section 12.3, “Creating a Tag,” on page 122](#).


The **Tag**  icon is available in various parts of the Sentinel interface, which allows you to quickly add tags to the desired data collection objects such as event sources, event source servers, Collector Managers, Collector plug-ins, report templates, and report results. When you click the **Tag**  icon, the Tags dialog box is displayed that allows you to select tags and to create new tags.

12.3 Creating a Tag

- 1 Log in to the Sentinel Web interface as a user in the Manage Tags 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 Select **Tags** in the navigation panel on the left or click the **Tag**  icon in the appropriate data object interface to which you want to associate tags.
- 3 Click **Create**.
- 4 Specify a name for the tag.

Tags have the following naming conventions, and a warning message is displayed if the name you specify does not comply with the following conventions:

- ♦ Tag names should not be more than 20 characters.
- ♦ There should not be any white space as part of the tag name.
- ♦ A tag name is not case-sensitive. You cannot create two tags with identical names except for capitalization. For example, you cannot have the tag names IDM and idm, because both are perceived as the same name.

- 5 Specify an optional description for the tag.

If the tag name is available, a message is displayed.

If a tag with the same name already exists, a message is displayed indicating the name is not unique. You must specify a different name for the tag.

- 6 Click **Save**.

12.4 Managing Tags

- ♦ [Section 12.4.1, "Sorting Tags," on page 122](#)
- ♦ [Section 12.4.2, "Adding and Removing Tags from Favorites," on page 123](#)
- ♦ [Section 12.4.3, "Viewing and Modifying Tags," on page 123](#)

12.4.1 Sorting Tags


You can sort tags either based on their names or based on the number of objects associated with the tags.

- 1 Log in to the Sentinel Web interface as a user in the Manage Tags role.
- 2 Select **Tags** in the navigation panel, then click **More**.
- 3 (Conditional) To sort the tags in the alphabetical order, select **Sort by Name**.
- 4 (Conditional) To sort the tags based on the number of objects associated with them, select **Sort by Count**.

The Tags are sorted according to the selection.


12.4.2 Adding and Removing Tags from Favorites

You can add your frequently used tags to the Favorites section so that it is easier to locate them and associate them with objects. When a tag is added to the Favorites section, it is removed from the Other section.

- 1 Log in to the Sentinel Web interface as a user in the Manage Tags role.
- 2 Select **Tags** in the navigation panel on the left.
- 3 To add or remove a tag from Favorites, select the tag, then click the **Favorites**  icon.

12.4.3 Viewing and Modifying Tags

You can modify only the description of a tag. The tag name cannot be modified because it might be used to tag events and other data collection objects, and it is not an accepted practice to modify events that are already stored. Therefore, to modify the name of a tag, you must create a new tag.


- 1 Log in to the Sentinel Web interface as a user in the Manage Tags role.
- 2 Select **Tags** in the navigation panel on the left.
- 3 Select the tag that you want to edit, and click the **Edit**  icon.
- 4 Modify the description as necessary, then click **Save**.

12.5 Performing Text Searches for Tags

This option is useful when you want to look for a particular tag.

- 1 Log in to the Sentinel Web interface as a user in the Manage Tags role.
- 2 Select **Tags** in the navigation panel on the left.
- 3 To search for a particular tag, specify the name or description of the tag or a keyword. To search for multiple tags, specify the tag names separated by the space character.
The tag that matches the keyword is displayed.

12.6 Deleting Tags

- 1 Log in to the Sentinel Web interface as a user in the Manage Tags role.
- 2 Select **Tags** in the navigation panel on the left.
- 3 Select the tag that you want to delete, then click the **Delete**  icon.
The Sentinel tag is a system tag that tags all Sentinel internal events, and cannot be deleted.
- 4 Click **Delete** to confirm deletion.

12.7 Associating Tags with Objects


You can associate tags with event sources, event source servers, Collector Managers, Collector Plug-ins, event routing rules, and reports and report templates. You can add more than one tag to a data collection object. However, the `rv145` field, which stores the tag value, can hold a maximum of 256 characters. Therefore, the maximum number of tags that you can associate with an object depends on the length of the tag name.

- ◆ [Section 12.7.1, “Associating Tags with Event Routing Rules,” on page 124](#)
- ◆ [Section 12.7.2, “Associating Tags with Event Sources,” on page 124](#)
- ◆ [Section 12.7.3, “Associating Tags with Collector Managers,” on page 124](#)
- ◆ [Section 12.7.4, “Associating Tags with Event Sources Servers,” on page 124](#)
- ◆ [Section 12.7.5, “Associating Tags with Collector Plug-Ins,” on page 125](#)
- ◆ [Section 12.7.6, “Associating Tags with Report Results and Report Definitions,” on page 125](#)


12.7.1 Associating Tags with Event Routing Rules

- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Click **Routing** in the toolbar, then click **Create**.
- 3 Specify a name and filter criteria for the rule.
- 4 Click **Select tag**, then select the tags that you want to associate with the rule.
- 5 Click **Set**.

12.7.2 Associating Tags with Event Sources


- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection > Event Sources**.
- 3 Select the event sources that you want to associate with the tag.
- 4 Select the **Configure**  icon, then select **Tags**.
- 5 Select the tags you want to associate, then click **Set**.

12.7.3 Associating Tags with Collector Managers


- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection > Event Sources**.
- 3 From the **Collector Managers** section, select one or more of the Collector Managers that you want to associate with the tag.
- 4 Select the **Configure**  icon, then select **Tags**.
- 5 Select the tags you want to associate, then click **Set**.

12.7.4 Associating Tags with Event Sources Servers

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

- 3 From the **Event Source Servers** section, select one or more event source servers that you want to associate with the tag.
- 4 Select the **Configure**  icon, then select **Tags**.
- 5 Select the tags you want to associate, then click **Set**.

12.7.5 Associating Tags with Collector Plug-Ins


- 1 Log in to the Sentinel Web interface as a user in the administrator role.
- 2 Select **Collection > Event Sources**.
- 3 From the Collector Plugins section, select one or more of the Collector plug-ins that you want to associate with the tag.
- 4 Select the **Configure**  icon, then select **Tags**.
- 5 Select the tags you want to associate, then click **Set**.

12.7.6 Associating Tags with Report Results and Report Definitions

NOTE: When a tag is set on a report definition, the report results under the report definition inherit the tag by default. Inherited tags for a report result appear disabled in the Tag selector dialog box.

- 1 Log in to the Sentinel Web interface.
- 2 Select **Reports** in the navigation panel on the left.
- 3 Select the report result or the report definition that you want to associate with a tag.
- 4 Do one of the following:
 - ♦ Select **Tags** from the **more** drop-down list.
 - ♦ Click **Edit** at the bottom left pane.
- 5 Select one or more tags that you want to associate with selected reports.
- 6 Click **Set**.

12.8 Viewing Tagged Events

- 1 Log in to the Sentinel Web interface.
- 2 Do either of the following:
 - ♦ From the Tags panel, select the tag for which you want to view events, then select **Search**.
 - ♦ In the **Search** field, click the **Tag**  icon, select the desired tags, then click **OK**. Click **Search**.
 - ♦ In the **Search** field, specify `rv145:<tagname>` or `@<tagname>` as the search criteria, then click **Search**.

13 Reporting

Sentinel includes a variety of Solution Packs that contain out-of-the-box reports, some of which are device-specific. Some of the reports allow users to specify the columns Sentinel displays in the results. These reports use a Lucene-based query language.

When the Sentinel page is loaded for the first time, Sentinel loads the report definitions in the system and displays them in the **Reports and Searches** panel. You can also import the pre-packaged reports using Solution Packs. For more information, see [Section 13.1, “Importing Report Definitions,” on page 127](#).

You can create a report from an existing report and search for events based on the report definitions. You can also schedule, export, and email the reports.

You can schedule reports on Sentinel servers that are distributed across different geographic locations. For more information, see “[Configuring Data Federation](#)” in the *NetIQ Sentinel Administration Guide*.

All report results depend on the data viewing permissions associated with the user's role. For more information, see “[Configuring Roles and Users](#)” in the *NetIQ Sentinel Administration Guide*.

- ◆ [Section 13.1, “Importing Report Definitions,” on page 127](#)
- ◆ [Section 13.2, “Creating Reports,” on page 128](#)
- ◆ [Section 13.3, “Scheduling a Report,” on page 131](#)
- ◆ [Section 13.4, “Grouping Reports Based on Category,” on page 132](#)
- ◆ [Section 13.5, “Viewing Events,” on page 132](#)
- ◆ [Section 13.6, “Renaming a Report Result,” on page 133](#)
- ◆ [Section 13.7, “Marking Report Results as Read or Unread,” on page 133](#)
- ◆ [Section 13.8, “Managing Favorite Reports,” on page 134](#)
- ◆ [Section 13.9, “Associating Tags with Report Results and Report Definitions,” on page 134](#)
- ◆ [Section 13.10, “Exporting Report Definitions and Report Results,” on page 135](#)
- ◆ [Section 13.11, “Deleting Reports,” on page 136](#)
- ◆ [Section 13.12, “White Label Template Report,” on page 137](#)

13.1 Importing Report Definitions

Reports in Sentinel are designed as plug-ins (.zip or .rpz files that include the report definition in addition to the metadata and resources used by the report). User in the *Manage Reports* role can import new or updated reports into Sentinel.

The primary sources for new or updated reports are:

- ◆ **Solution Packs:** Solution Packs provide a framework where sets of content are packaged into controls, each of which is designed to enforce a specific business or technical policy. They are created in Sentinel Solution Designer and contain different types of plug-ins, including Sentinel reports. For more information on Solution Packs, see the [Sentinel Plug-ins Web Site](#).

- ♦ **JasperForge iReport:** You can modify or write reports by using JasperForge iReport, which is a graphical report designer for JasperReports. iReport is an open source report development tool that is available for download from JasperForge.org (as of the time of this publication).

New or modified reports can include additional fields that are not presented in the Sentinel interface. They must adhere to the file and format requirements of the report plug-ins. For more information about database fields and file and format requirements for report plug-ins, see the [Sentinel SDK Web site](#).

Importing Reports

You can either import a single or multiple report definitions.

- 1 Log in to the Sentinel Web interface.
- 2 In the **Reports and Searches** panel, click **Import reports or searches**.
- 3 Browse and select the report definition file that you want to import from your local computer, and then click **Import**. You can import only .zip, .rpz, .spz, and .rdz files.
- 4 Click **OK**.

Sentinel imports the new report definitions from the selected file. You can schedule a report immediately, if necessary.

Sentinel verifies the unique ID of the report to determine whether an older or identical version of the report already exists in the report repository. If it does, Sentinel displays the details of both the reports so that the user can decide whether to cancel the action or replace the existing report with the current report.

13.2 Creating Reports

You can create a new report definition based on the report definitions included in existing Sentinel reports or from those reports you imported to the system.

NOTE: You cannot create a report from a user created report.

You can create a report by using the desired parameters such as a From and a To date, add additional criteria to the existing report, and save the report definitions with a unique name.

Use the following procedure to create a report:

- 1 Log in to the Sentinel Web interface as a user with the Manage Reports permission.
- 2 In the **Reports and Searches** panel, select the report definition from which you want to create a new report.
- 3 Click **Create report**.
- 4 To create a report, specify the following parameters:

Parameter	Description
Report name	Specify a unique name for the report. The name should not exceed 200 characters.
Based on	Select the base report from which you want to create the report. You can view a sample report by clicking the View Sample button.
Description	Sentinel automatically displays the description based on the report you selected. You can edit the default description.
Criteria	Sentinel automatically populates the criteria based on the report you selected. This criteria is not editable.
Additional criteria	Specify additional search criteria to the existing criteria. To define additional criteria click Edit Criteria . To define criteria from available system objects containing criteria, click Add Criteria . The criteria that you define here is appended to the existing criteria.
Data sources	Select the source computers on which you want to run the reports by clicking the selected data sources link. You can select the data sources only if Sentinel is configured for data federation. For more information, see “ Configuring Data Federation ”in the <i>NetIQ Sentinel Administration Guide</i> .
Additional Criteria	Specify additional criteria to refine the results. The criteria that you specify here can be edited while scheduling the report. If you specify Criteria name , Sentinel displays the name at the end of the report results. NOTE: This parameter is not available for all reports.
Time Zone	Specify the time zone with which you want to populate the report. When you schedule the report, Sentinel displays this time zone in the report data. For example, if the time zone is set to US/Pacific-New time, the report data displays the selected time zone. By default, it displays the time zone that is set in the client system. NOTE: This parameter is not available for all reports.

Parameter	Description
Date Range	<p>If the report includes time period parameters, choose the date range. All time periods are based on the local time for the browser. The From Date and the To Date automatically change to reflect the option you selected.</p> <ul style="list-style-type: none"> ◆ Current Day: Shows events from midnight of the current day until 11:59:00 PM of the current day. If the current time is 8:00:00 AM, the report shows 8 hours of data. ◆ Previous Day: Shows events from midnight yesterday until 11:59:00 PM yesterday. ◆ Week To Date: Shows events from midnight Sunday of the current week until the end of the selected day. ◆ Previous Week: Shows events for the last seven days. ◆ Month to Date: Shows events from midnight the first day of the current month until the end of the selected day. ◆ Previous Month: Shows events for a month, from midnight of the first day of the previous month until 11:59:00 PM of the last day of the previous month. ◆ Custom Date Range: Shows events for a period whose start and end date are chosen. If you select Custom Date Range, set the start date (From Date) and the end date (To Date) for the report.
Group By	<p>Group the events according to specific event field by selecting the event field from the Group by drop-down list.</p> <p>NOTE: This parameter is not available for all reports.</p>
Language	<p>Select the language in which you want the report labels and descriptions displayed. The possible values are English, French, German, Italian, Japanese, Traditional Chinese, Simplified Chinese, Spanish, or Portuguese.</p> <p>The default value is the language with which the current user logged in, if that language is supported by the report. If the report does not support the language, the report's default language (typically English) is used.</p> <p>Sentinel displays the data in the report in the language that was originally used by the event source.</p>
Email to	<p>Specify an e-mail address in the Email to field. If you want to email the report to more than one user, separate the e-mail addresses with a comma.</p>
Result limit	<p>Specify the number of results you want Sentinel to display or store when you schedule the report. By default, 1000 results are stored.</p> <p>If you specify a value in Group By field, the result limit is based on grouping.</p>

5 Click **Create**.

13.3 Scheduling a Report

You have to schedule a report to see the report results. You can schedule a report to run immediately, once, or at a specified time. All Sentinel reports come with a sample report. When you schedule a report, the report runs at the scheduled time and the report results are saved in a PDF format. When you schedule to run a report immediately, it searches the events related to the report's query from midnight till the time that you ran the report. Sentinel displays the results in the Sentinel console and also stores the results as a PDF.

NOTE: The report results in the PDF are different than the results in reports run immediately. The report results in the PDF are for the time range that you specified while scheduling a report. When you run a report immediately, the report includes events from midnight to the time you ran the report.

You can schedule a report using the desired parameters, such as a From and a To date, and save the report results with a name of your choice. After the report runs, you can view it in the Sentinel console or in PDF. For more information on viewing the reports results, see [Section 13.5, "Viewing Events,"](#) on page 132.

Because the reports run asynchronously, you can simultaneously perform other tasks in the application while you run reports. If the Sentinel server was restarted while a report was processing, you can either cancel or restart the report. If you restart the report, it runs with the same parameters that were used the first time. If the report was scheduled with a relative time setting, such as **Week to Date**, the time period for rerunning the report is based on the current date and time and not the date and time when the report was initially scheduled.

Use the following procedure to schedule a report:

- 1 Log in to the Sentinel Web interface as a user with the Manage Reports permission.
- 2 In the **Reports and Searches** panel, select the report you want to schedule, then click **Schedule**.
- 3 You can schedule the report to run immediately or schedule it to run later, either once or on a recurring basis. For scheduled reports, choose a frequency and specify a time (**Start Time**) for the report to run except if you select **Now**. The report runs based on the time settings of the Sentinel server.
 - ♦ **Now:** This is the default. It runs the report immediately.
 - ♦ **Once:** Runs the report once at the specified date and time.
 - ♦ **Daily:** Runs the report once a day at the specified time.
 - ♦ **Weekly:** Runs the report once a week on the same day at the specified time.
 - ♦ **Monthly:** Runs the report on the same day of the month every month, starting at the specified date and time. For example, if the start date and time is May 26, 2010 4:00:00 p.m., the report runs on the 26th day of the month at 4:00:00 p.m. every month.
- 4 Specify a unique name to identify the report results. By default, the name of the report is `Report 1`. Next time you schedule a report for the same report definition, the name of the report by default displays `Report2`.
- 5 (Conditional) If Sentinel is configured for data federation, click the **selected data sources** link in the **Data sources** section to select the data sources on which the reports can be run. For more information on data federation, see "[Configuring Data Federation](#)" in the *NetIQ Sentinel Administration Guide*.
- 6 Displays the additional criteria that you specified in "[Additional Criteria](#)" on page 41. You can combine more than one criteria with `And` or `Or` operator.

- 7 To schedule a report, specify other parameters that are similar to **Create report**. For more information, see [Section 13.2, “Creating Reports,” on page 128](#).
- 8 (Conditional) For user-defined reports, you can specify the number of results stored in the **MaxResults** field. By default, Sentinel stores 1000 results. The number you specify here overrides the number you specified in the **Result Limit** parameter on the Create Report screen.
- 9 Click **Schedule**.

13.4 Grouping Reports Based on Category

You can group the reports based on the category. For example, ISO, Sentinel Core, and so on. To categorize the reports, click **More options**, select **Group by**, and then click **Category**. The report definitions are grouped based on different categories segregated with the title for each category. The title displays the number of reports that are available under each category.

13.5 Viewing Events

- ◆ [Section 13.5.1, “Viewing Events Based on Report Criteria,” on page 132](#)
- ◆ [Section 13.5.2, “Viewing Events Based on Report Result Criteria,” on page 132](#)
- ◆ [Section 13.5.3, “Finding Reports Based on Words,” on page 133](#)
- ◆ [Section 13.5.4, “Finding Reports Based on Tags,” on page 133](#)

13.5.1 Viewing Events Based on Report Criteria

You can view events for the existing report definitions that are loaded in Sentinel or for the report that you create. To view the events, select a report definition and click **Search events**. Specify the required parameters and click **Search**. The search results are displayed. For more information, see [Section 3.2, “Viewing Search Results,” on page 35](#).

When you schedule a report to run immediately, the results display in the Sentinel interface. You can interact with the results as you would do for any search, such as refining the search or performing an action on the results. For more information, see [Section 3.3, “Refining Search Results,” on page 38](#).

13.5.2 Viewing Events Based on Report Result Criteria

The reports are loaded and displayed in the **Reports and Searches** panel on the left of the Sentinel Web console. Click **More options > Show all reports** to view all reports, click **More options > Show only scheduled reports** to view only the scheduled reports, or click **More options > Show only unread reports** to view the reports that are not read.

The report results for each user varies depending on the data security settings configured for the role of that user. All the report results are ordered by the creation time. If there is more than one report, the **Show more** link is displayed that displays other report results.

In the **Reports and Searches** panel, the report definitions show the number of unread reports with a blue dot next to them. If you have grouped the report definitions by category, the category title shows the number of unread reports under that category.

A blue dot next to the report result indicates that the report result is unread. For more information, see [“Marking Report Results as Read or Unread” on page 133](#).

Viewing the Report Results in PDF

The scheduled report results are saved as PDF. To view the report results, select the report result in the **Reports** panel, and then click **View PDF**. Select the PDF reader and click **OK**. Report results are organized from the newest to the oldest. The bottom of the PDF displays the parameter values used to run the report.

The **Reports and Searches** panel on the left side of the page displays a status pane at the bottom left corner of the page. The status pane displays the parameter information or the schedule information associated with the selected report. It also displays the error messages, if the report execution resulted in an error.

To view the sample report to find out how the completed report looks, click the sample report, and then click **View report PDF**.

13.5.3 Finding Reports Based on Words

You can find a report based on words in the report definition. In the **Reports and Searches** panel, specify the words for which you want to search in the **Find** field. The search is based on the individual words. For example, if you specify Audit Trail, Sentinel displays all reports that have Audit or Trail in the report definition title.

13.5.4 Finding Reports Based on Tags

If you have tagged report definitions, you can find reports based on tags by clicking **Tags** next to the **Find** field. Select the tags on which you want to search the reports, and then click **Find**. The **Reports and Searches** panel displays only those reports that are associated with the selected tag. For associating a report with a tag, see [Section 13.9, “Associating Tags with Report Results and Report Definitions,” on page 134](#).

13.6 Renaming a Report Result

- 1 In the **Reports and Searches** panel, select a report result.
- 2 Do one of the following:
 - ◆ Click **More options** and select **Rename**.
 - ◆ Double-click the report name in the status pane.
- 3 Specify a name in the bottom left status pane.
- 4 Click **Rename**.

The selected report result is renamed under the report definition.

13.7 Marking Report Results as Read or Unread

When a report result is created under a report definition, the report result is in unread state. An unread report result appears with a blue dot next to the report result in the **Reports and Searches** panel. When you view a report result, the blue dot is removed to indicate that the report has been read. You can also manually mark a report result as read or unread without viewing it by selecting the report and click **Mark read** or **Mark unread** respectively, from the options displayed.

In the **Reports and Searches** panel, each of the report definitions shows the number of unread reports next to it.

NOTE: The reports marked as read or unread are on a per-user basis. Each user can have a different set of read or unread reports.

13.8 Managing Favorite Reports

You can set the most used reports and searches as favorites. You can also create folders to store your favorite reports and searches, which helps you to locate and manage them easily.

- ♦ [Section 13.8.1, “Adding Reports as Favorites,” on page 134](#)
- ♦ [Section 13.8.2, “Removing Favorite Reports,” on page 134](#)

13.8.1 Adding Reports as Favorites

You can mark individual report definitions as Favorites so that they are easier to find.

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, select the report definition that you want to mark as favorite, and then click `Add to Favorites`.

The report definitions that are marked as favorites are displayed first than the other report definitions that are not marked as favorites. If you have grouped the report definitions by category, the report definitions that are marked as favorites are displayed at the top of each category.

13.8.2 Removing Favorite Reports

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, select the report definition that you want to remove from the Favorites, and then click `Remove from Favorites`.

13.9 Associating Tags with Report Results and Report Definitions

You can associate a tag on report definition or report results. For more information tags, see [Chapter 12, “Configuring Tags,” on page 121](#)

NOTE: When a tag is set on a report definition, the report results under the report definition inherit the tag by default. Inherited tags for a report result appear disabled in the Tag selector dialog box.

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, select the report result or the report definition that you want to associate with a tag and click **Tag**.
- 3 Select one or more tags that you want to associate with selected report definitions or report results.
- 4 Click **Set**.

NOTE: You can associate same tag for multiple reports simultaneously. Click **More** and then select **Select multiple reports and searches**, select the report definitions, and then click **Tag**. Select one or more tags that you want to associate, and then click **Set**.

13.10 Exporting Report Definitions and Report Results

You can export report definitions from one Sentinel instance and export it into another instance of Sentinel. You can either export single report definition or all report definitions. You can save individual report results as .zip file and send to another user.

NOTE: When you export report definitions, only the report definitions are exported. None of the associated report results are exported.

- ◆ [Section 13.10.1, “Exporting a Single Report Definition,” on page 135](#)
- ◆ [Section 13.10.2, “Exporting Multiple Report Definitions,” on page 135](#)
- ◆ [Section 13.10.3, “Exporting All Report Definitions,” on page 135](#)
- ◆ [Section 13.10.4, “Exporting a Report Result,” on page 136](#)

13.10.1 Exporting a Single Report Definition

- 1 Log in to the Sentinel Web console.
- 2 Select a report definition in the **Reports and Searches** panel, and click **Export**.
The Opening <Selected Report Name>.zip dialog box is displayed with the option to save the file on your local machine.
- 3 Save the file to the location you prefer.

13.10.2 Exporting Multiple Report Definitions

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, click **More options**, then select **Select multiple reports or searches** option.
- 3 Select the report definitions that you want to export.
- 4 Click **Export**.
The Opening reportexport.zip dialog box is displayed with the option to save the file.
- 5 Save the file to a location you prefer.

13.10.3 Exporting All Report Definitions

You can use the **Export All** option to export all reports as a .zip file.

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, click **Export All**.
The Opening reportexport.zip dialog box is displayed with the option to save the file.
- 3 Save the file to a location you prefer.

13.10.4 Exporting a Report Result

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, select the report results that you want to export, and then click **View PDF**.
A list of available report results appear.
A dialog box appears with an option to export and save the file as *<Report Def name>_<Report result name>.pdf*.
- 3 Click **Save File** option, and then click **OK** to save the file.

13.11 Deleting Reports

You can delete a report definition or a report result. If a report definition is deleted, all associated report results are also deleted.

- [Section 13.11.1, “Deleting a Report Definition,” on page 136](#)
- [Section 13.11.2, “Deleting Multiple Report Definitions,” on page 136](#)
- [Section 13.11.3, “Deleting a Report Result,” on page 137](#)
- [Section 13.11.4, “Deleting Multiple Report Results,” on page 137](#)

13.11.1 Deleting a Report Definition

- 1 Log in to the Sentinel Web console.
- 2 Select a report definition in the **Reports and Searches** panel, and then click the **Delete** icon.
- 3 Click **Delete** to confirm deletion.

13.11.2 Deleting Multiple Report Definitions

You can select multiple report definitions and delete all of them.

- 1 Log in to the Sentinel Web console.
- 2 In the **Reports and Searches** panel, click **More options**, then select **Select multiple reports or searches**.
- 3 A check box is displayed next to each report definition in the **Reports and Searches** panel. Select the check boxes to select the report definitions that you want to delete.
If no report definitions are selected, the **Tag**, **Export**, and **Delete** links are disabled.
The **Delete(x)** icon in the **Reports and Searches** panel, shows the number of selected report definitions, where **(x)** is the number of selected report results.
- 4 Click **Delete(x)**.
- 5 Click **Delete** again to confirm deletion.

13.11.3 Deleting a Report Result

- 1 Log in to the Sentinel Web console.
- 2 Select a report result under a report definition in the **Reports and Searches** panel, and then click **Delete**.
- 3 Click **Delete** again to confirm deletion.

13.11.4 Deleting Multiple Report Results

You can select multiple report results and delete all of them.

- 1 Log in to the Sentinel Web console.
- 2 Click **More**, then select **Select Multiple Results**.
- 3 A check box is displayed next to each report result in the **Reports and Searches** panel. Select the check boxes to select the report results.
If no report results are selected, the **Tag**, **Export**, and **Delete** links are disabled.
- 4 Click **Delete(x)**.
- 5 Click **Delete** again to confirm deletion.

13.12 White Label Template Report

Sentinel delivers an out-of-the-box White Label report template, which is available in the Reports and Searches panel. This panel displays the reports and searches alphabetically. If you have grouped the reports by category, this template is available under the Sentinel Core category.

Sentinel uses this template to present the report results. If you delete this template or if you do any updates to this template, you need to manually import this White Label template in to the Sentinel UI. You can also customize this template to include your own header, footer, and logo as per your organization's needs.

To customize the White Label Template:

- 1 In the **Reports and Searches** panel, select the White Label Template report definition, and then click **Export**.
- 2 Save the file to your local computer.
- 3 Create a new directory.
- 4 Extract the file contents to the new directory using any zip extraction tool.
- 5 In the new directory, open the **Resources** directory. In this directory, you can modify the following files:
 - ♦ **Header/Footer.jrxml**: Contains the Jasper report layout descriptions. You can modify the layout of fields, text, or images in the header and footer, but you must ensure that the overall size of the header and footer does not change. You can manually edit the XML file or use iReport to modify them.
 - ♦ **Header/Footer*.properties**: Contains the strings from the layout file localized into various languages. You can modify the strings that appear in the header or footer by editing this file. Ensure that the new strings do not exceed the space allocated to them. For information on editing `.properties` file, see [Oracle Java documentation](#).

- ♦ **Logo.jpg**: Contains the logo that appears in the White Label footer. You can replace this file with another image. Ensure that the size of the new image is exactly the same size of the existing image.
- 6 After you are done with the changes, use a zip tool to re-zip the modified report template.
 - 7 In the **Reports and Searches** panel, click **Import reports or searches**, browse to this zip file, and then click **Import**.

NOTE: If the directory structure is different than the original zip file, then the import process displays an error.

- 8 Schedule a report definition and view the report results to ensure the changes you made are applied correctly.

14 Viewing Compliance to Configuration Policies

Sentinel is a compliance monitoring system that helps you verify whether your enterprise is compliant with internal policies, information security standards such as PCI DSS and ISO 27000 series, and government regulations such as Sarbanes-Oxley, HIPAA, GLBA, and FISMA.

Sentinel extends its compliance monitoring capability by integrating seamlessly with your existing security management solutions, such as NetIQ Secure Configuration Manager (SCM). Integration with Secure Configuration Manager helps you to assess system configurations against regulatory requirements, security best practices, and corporate IT policies to demonstrate compliance and manage information security risk. This integration helps you to view the security and audit information from both Sentinel and Secure Configuration Manager in a single interface.

Sentinel is auto-configured to receive events from Secure Configuration Manager. You must configure Sentinel to receive compliance details associated with the Secure Configuration Manager events. For information about configuring Sentinel to receive compliance details from Secure Configuration Manager, see “[Viewing Compliance to Configuration Policies](#)” in the *NetIQ Sentinel Administration Guide*.

14.1 Viewing Secure Configuration Manager Events and Compliance Details

To view the compliance details, search for Secure Configuration Manager events by using the appropriate query, for example, `(sev:[0 TO 5]) AND pn:(Secure OR Configuration OR Manager)`, then click the **View compliance details** icon associated with the event. For more information about searching events in Sentinel, see “[Searching Events](#)” on page 33.

Compliance details provide information about compliance to configuration policies of various assets in your IT environment based on rules configured in the Secure Configuration Manager. Compliance details also provide information about the risk to the organization because of the non-compliance of the assets and the validation results of security checks in the policies.

The compliance details can be based on assets, policy, or both assets and policy. The compliance details provides information about the following:

- ♦ **Introduction:** This section provides information about the assets, user accounts, policy, and the timestamp when the compliance was checked.
- ♦ **Risk Analysis:** This section displays charts representing the risks involved based on the compliance status of the assets and the importance of the asset in the organization.
- ♦ **Security Check Details:** This section displays the validation results for security checks in the policies.

For more information about understanding compliance details, see [Understanding Report Results](#) in the *NetIQ Secure Configuration manager User Guide*.

15 Viewing Change Guardian Events

If the Change Guardian server is configured to send its events to Sentinel, you can view Change Guardian events in Sentinel. To verify whether Change Guardian is configured, contact your system administrator.

To view the Change Guardian events and the event details, run a search by using the appropriate query, for example, `pn:Change Guardian`, then click **Search**. For more information about searching events in Sentinel, see [“Searching Events” on page 33](#).

Change Guardian events indicates activities such as:

- ◆ Changes or access to critical files, platforms and systems
- ◆ Activities of privileged users

Each event provides detailed information about the file that was changed, the actual changes made to the file, and so on. For more information about the event details, see [Understanding Event Information](#) in the *Change Guardian User Guide*.

To view the event details, click the **Change Guardian** icon associated with the event.

16 Configuring Incidents

In Sentinel, a set of related events (for example, a possible attack) can be grouped together form an incident. An incident in open state alerts you to investigate, resolve, and close the incident. For example, the resolution to an attack might be to close a port, block a source IP, or rebuild a machine.

Incidents are created automatically as a result of a correlation rule being triggered, or they are created manually by a security analyst monitoring incoming data or querying past data.

- ♦ [Section 16.1, “Accessing Incidents,” on page 143](#)
- ♦ [Section 16.2, “Creating Incidents,” on page 143](#)
- ♦ [Section 16.3, “Managing Incidents,” on page 144](#)
- ♦ [Section 16.4, “Adding an Incident View,” on page 146](#)


16.1 Accessing Incidents

You access the incidents through the Sentinel Control Center. You need to have appropriate permissions to access this tab. Only an Administrator has controls to enable or disable access to the features of incidents for a user.

- 1 Log in to the Sentinel Web interface as a user with permissions to access incidents.
- 2 Click **Applications** in the toolbar.
- 3 Click **Launch Sentinel Control Center**.
- 4 Log in to the Sentinel Control Center as a user with permissions to access incidents.
- 5 Click **Incidents**.

The Incidents are displayed.

16.2 Creating Incidents

- 1 Access the **Incidents** tab in the Sentinel Control Center.
For more information, see [Section 16.1, “Accessing Incidents,” on page 143](#).
- 2 In the menu, click **Incidents > Create Incident**.
or
Click the **Create Incident**  button in the toolbar.
- 3 Use the following information to create the incident:
 - Title:** Specify the title of the incident.
 - State:** Select the state of the incident from the drop-down list.
 - Severity:** Select the severity of the incident from the drop-down list.
 - Priority:** Select the priority of the incident from the drop-down list.
 - Category:** Select the category of the incident from the drop-down listor

Create your own category by clicking the button next to the **Category** field, then click **Add**. You must specify a name and a description of the new category.

Responsible: Select the user that is responsible to investigate and close the incident.

Description: Specify a description of the incident.

Resolution: Specify the steps required to resolve the incident.

4 Click **Create**.

The Incident ID is automatically generated after you click **Create**.

You can also create incidents from the Sentinel Web interface. For more information, see [Section 3.5.4, “Creating an Incident,” on page 46](#).

After the incident is created, proceed to [Section 16.3, “Managing Incidents,” on page 144](#) to manage the incident.

16.3 Managing Incidents

- ◆ [Section 16.3.1, “Viewing an Incident,” on page 144](#)
- ◆ [Section 16.3.2, “Attaching Workflows to Incidents,” on page 145](#)
- ◆ [Section 16.3.3, “Adding Attachments to Incidents,” on page 145](#)
- ◆ [Section 16.3.4, “Adding Notes to Incidents,” on page 145](#)
- ◆ [Section 16.3.5, “Executing Incident Actions,” on page 145](#)
- ◆ [Section 16.3.6, “E-mailing an Incident,” on page 146](#)

16.3.1 Viewing an Incident

1 Click **Incidents** in the Sentinel Control Center.

For more information, see [Section 16.1, “Accessing Incidents,” on page 143](#).

2 From the menu, click **Incidents > Display Incident View Manager**

or

Click the **Display Incident View Manager**  button in the toolbar.

3 Select the desired Incident in the Incidents View window.

When you view an incident, you see the tabs listed below where you can perform Incident related activities. As you investigate and remediate an Incident, additional information can be added to these tabs.

Events: Lists events attached to this Incident. For more information, see [Section 3.5.4, “Creating an Incident,” on page 46](#).

Assets: Lists assets affected by the events of this Incident.

Vulnerability: Lists asset vulnerabilities.

Advisor: Displays Asset attack and alert information.

iTRAC: Allows you to add a workflow to Incident.

History: Lists the activities performed on the current Incident.

Attachments: Allows you to add an attachment to the Incident created in the system.

Notes: Allows you to add notes to the Incident.

16.3.2 Attaching Workflows to Incidents

- 1 In the Incidents View window, select the desired Incident.
- 2 Click the **iTRAC** tab.
- 3 Select a workflow from the iTRAC process drop-down list.

For more information about workflows, see [Chapter 17, “Configuring iTRAC Workflows,”](#) on page 147.

- 4 Click **Save**.

You can attach only one workflow to an Incident.

16.3.3 Adding Attachments to Incidents

- 1 In the Incidents View window, select the desired Incident.
- 2 Click the **Attachments** tab, then click **Add**.
- 3 Click **Browse**, then navigate to the attachment and select it.
- 4 Specify the required information, or accept the default entries.
- 5 Click **OK**, then click **Save**.

You can right-click the attachment to view it or save it to your local hard drive.

16.3.4 Adding Notes to Incidents

- 1 In the Incidents View window, select the desired Incident.
- 2 Click the **Notes** tab, then click **Add**.
- 3 Specify your notes, then click **OK**.
- 4 Click **Save** to update the Incident.

To edit or delete the note, select a note in the **Notes** tab of the Incident window, right-click the note, then select **edit** or **delete**.

16.3.5 Executing Incident Actions

Any configured Javascript action or iTRAC activity can be executed on an Incident.

- 1 In the Incidents View window, select the desired Incident.
- 2 In the menu, click **Action > Execute Incident Action**.
or
Click the **Execute Incident Action** button.
- 3 Select an Action or click the **Add Action** button to create a new one.
- 4 Click **Execute**.


If the action is a Javascript Action, a window opens to show the progress of the action.

- 5 To add the command output to the Incident, click the **Attach to Incident** button.

The action output is saved and can be viewed from the **Attachments** tab of the Incident.

16.3.6 E-mailing an Incident

To e-mail an Incident using the preinstalled E-mail Incident action, you must have an SMTP Integrator configured with valid connection information and with the property SentinelDefaultEMailServer set to "true". For more information, see the SMTP Integrator documentation available at the [Sentinel Plug-in Web site](#).

- 1 In the Incidents View window, select the desired Incident.
- 2 Click the **Email Incident**  icon.
- 3 Specify the required information.
- 4 Select which HTML attachments should be included in the mail message: the events included in the incident, assets, vulnerabilities, Advisor attacks, incident history, attachments, and notes.
- 5 Click **OK**.

16.4 Adding an Incident View

- 1 Click the **Incidents** tab in the Sentinel Control Center.
For more information, see [Section 16.1, "Accessing Incidents," on page 143](#).
- 2 Click the **Manage Views** drop-down, then select **Add View**.
- 3 Specify a name in the **Option Name** field. Click each button (listed below) to specify the options.
 - ♦ **Fields:** The variables of the events attached to Incidents are displayed as fields. By default, all the fields are arranged as columns in the Incidents View. You can add or remove columns, and arrange the order of the columns by using the up and down arrows.
 - ♦ **Group By:** Allows you to set rules to group Incidents.
 - ♦ **Sort:** Allows you to set rules to sort the Incidents.
 - ♦ **Filter:** Allows you to set filters. Only Incidents that match the filter are displayed in the Incidents View.
 - ♦ **Leaf Attribute:** Allows you to select attributes from the list, which is displayed as the first column in the Incidents View.
- 4 Click **Save**.

17 Configuring iTRAC Workflows

This chapter provides information about using iTRAC workflows to automate and track incidents.

- ♦ [Section 17.1, “Overview,” on page 147](#)
- ♦ [Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148](#)
- ♦ [Section 17.3, “Using the Template Manager,” on page 149](#)
- ♦ [Section 17.4, “Template Builder Interface,” on page 150](#)
- ♦ [Section 17.5, “Creating a Template,” on page 152](#)
- ♦ [Section 17.6, “Managing Templates,” on page 152](#)
- ♦ [Section 17.7, “Steps,” on page 153](#)
- ♦ [Section 17.8, “Adding Steps to a Workflow,” on page 157](#)
- ♦ [Section 17.9, “Managing Steps,” on page 158](#)
- ♦ [Section 17.10, “Transitions,” on page 161](#)
- ♦ [Section 17.11, “Activities,” on page 167](#)
- ♦ [Section 17.12, “Creating iTRAC Activities,” on page 169](#)
- ♦ [Section 17.13, “Managing Activities,” on page 170](#)
- ♦ [Section 17.14, “Managing iTRAC Roles,” on page 171](#)
- ♦ [Section 17.15, “Process Management,” on page 172](#)

17.1 Overview

iTRAC workflows are designed to provide a simple, flexible solution for automating and tracking an enterprise’s incident response processes. iTRAC leverages Sentinel’s internal incident system to track security or system problems from identification (through correlation rules or manual identification) through resolution.

Workflows can be built using manual and automated steps. Advanced features such as branching, time-based escalation, and local variables are supported. Integration with external scripts and plugins allows for flexible interaction with third-party systems. Comprehensive reporting allows administrators to understand and fine-tune the incident response processes.

NOTE: Access to manage iTRAC templates, activities, and processes can be enabled on a user-by-user basis by any user with the ability to change user permissions.

The iTRAC system uses three Sentinel objects that can be defined outside the iTRAC framework:

- ♦ **Incident:** Incidents within Sentinel are groups of events that represent an actionable security incident, plus associated state and meta-information.

Incidents are created manually or through Correlation rules. They can be associated with a workflow process. They can be viewed on the **Incidents** tab.

- ♦ **Activity:** An activity is a predefined automatic unit of work, with defined inputs, command-driven activity, and outputs (for example, automatically attaching asset data to the incident or sending an e-mail).

Activities can be included in a workflow template and executed during workflow processes, or they can be executed within an incident.

- ♦ **Role:** Sentinel users can be assigned to one or more roles. Manual steps in the workflow processes can be assigned to a role. For more information, see [Section 17.14, “Managing iTRAC Roles,”](#) on page 171.

iTRAC workflows have four major components that are unique to iTRAC:

- ♦ **Step:** A step is an individual unit of work within a workflow; including manual steps, decision steps, command steps, mail steps, and activity-based steps. Each step displays as an icon within a given workflow template.
- ♦ **Transition:** A transition defines how the workflow moves from one state (activity) to another. A transition is determined by an analyst action, by the value of a variable, or by the amount of time elapsed.
- ♦ **Templates:** A template is a design for a workflow that controls the flow of execution of a process in iTRAC. The template consists of a network of manual and automated steps that combine activities and criteria for transition between the steps.

Workflow templates define how an incident is responded to after a process based on that template is instantiated. A template can be associated with many incidents.

- ♦ **Processes:** A process is a specific instance of a workflow template that is actively being tracked by the workflow system. It includes all the relevant information for the instance, including the current step in the workflow, the associated incident, the results of steps, attachments, and notes. Each workflow process is associated to one incident.

NOTE: On a system with 16GB of RAM and 8 core CPU, you can run a maximum of 1000 processes on a single Sentinel server instance.

17.2 Accessing the iTRAC Administration Tools

There are multiple tools that allow you to create iTRAC workflows in Sentinel. To access these tools:

- 1 Access the Sentinel Control Center.

- 1a Log in to 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.

- 1b In the toolbar, click **Applications**.

- 1c Click **Launch Control Center**.

- 1d Click **Yes** to accept the security certificate.

- 1e Specify a username and password of a user that has rights to access the SCC, then click **Login**.

- 1f Click **Accept** or **Accept Permanently** to accept the security certificate and display the SCC.

- 2 Click the **iTRAC** tab, then click iTRAC in the toolbar.

All of the different administrative tools are listed here.

Process Manager: Manages the instantiated workflow processes.

Activity Manager: Defines the activities used in the iTRAC workflows.

Template Manager: Defines the templates used in the iTRAC workflows.

iTRAC Role Manager: Assigns roles that are used by the iTRAC workflows to assign work items to groups of users.


17.3 Using the Template Manager

The Template Manager can be used to create, view, modify, copy, or delete a template. Within the Template Manager you can add, delete, copy, view, and edit templates. Templates can be sorted into folders for easy management

In the Template Manager, you can:

- ♦ Create new workflow templates.
- ♦ Edit or copy existing templates.
- ♦ Define workflow steps:
 - ♦ Mark steps as Manual or Automated
 - ♦ Include a description of a step or include instructions for iTRAC users
- ♦ Define transitions between steps:
 - ♦ Transition type
 - ♦ Escalation procedures
 - ♦ Timeout and alert attributes

Figure 17-1 iTRAC Workflow




17.3.1 Default Templates

iTRAC is shipped with several templates to use as examples. The process and activity attributes for these templates are set to pre-defined values. Users can modify them to suit their requirements. The default templates are:

- ♦ AlertTimeoutExample
- ♦ TwoStepSimpleExample
- ♦ ConditionalTransitionExample
- ♦ CommandExample

17.4 Template Builder Interface

Figure 17-2 Template Builder Interface



You see the following panes in the Template Builder window:








- ◆ **Process Tree:** This pane displays the steps, transitions and variables added to the template. Users can add steps or variables, and edit or remove steps, variables and transitions.
 - To perform an action on a step, variable, or transition:
 - ◆ Expand the relevant group in the tree.
 - ◆ Select and right-click an existing attribute.
 - ◆ Select the action you want to perform.
- ◆ **Process:** This is the main GUI for viewing and creating a Workflow template. For more information on creating a Workflow template, see "[Section 17.5, "Creating a Template," on page 152](#)".
- ◆ **Quick Edit:** Select a step or transition to see its properties. This pane allows you to edit process attributes.
 - To edit the details of steps by using Quick Edit:
 1. Click the Process Attribute value in the Quick Edit Pane.
 - The attribute values are highlighted, indicating Edit Mode.
 2. Modify the value and click anywhere outside the Quick Edit frame to save the new value.
- ◆ **Messages:** This pane displays messages if steps or transitions are incomplete. You must resolve any issues listed here before saving the template.
- ◆ **Overview:** This pane displays an overview of the entire template.
- ◆ **Step Palette:** There are four types of steps in the Step Palette. You can drag and drop the steps into the Process pane.
 - ◆ Decision step
 - ◆ Mail step

- ◆ Manual step
- ◆ Command step
- ◆ **Activities:** The activities added in the Activity Manager are shown in this pane and can be added to a workflow template. The user can also add, edit and remove activities. For more information, see [Section 17.13, “Managing Activities,” on page 170](#).

IMPORTANT: Use caution when editing or deleting an Activity that is already in use.

The following icons are used in the Template Builder to represent the steps:

Table 17-1 Template Builder Icons

Icon	Description
	Start Step: All workflow templates have a start step.
	Decision Step: This step provides different execution paths depending on the value of a variable defined in a previous step.
	Mail Step: This step sends a pre-written e-mail.
	Manual Step: This step indicates that manual work must be performed, often outside the Sentinel system (For example, telephoning the owner of the affected system or analyzing the results of a scan).
	Activity Step: This step is a predefined set of activities.
	Command Step: This step executes a command or script on the iTRAC workflow server, which is usually installed in the same place as the Data Access Service (DAS). The output of the command can be stored in a string variable and used as input to a decision step.
	End Step: This step signifies the completion of a workflow process.

17.5 Creating a Template

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Click Add to display the iTRAC Template Builder window.
- 5 In the Process Details window, provide a name and description (optional) of the template, then click **OK**.
- 6 Drag and drop a step from the Step Palette or an activity from the Activities pane into the Process window.
- 7 Add as many steps and activities as needed to create the template.
- 8 Right-click the step where you need to add transition, then click **Add Transition**.
The transition is added after the selected step.

NOTE: Any step (except for the end step) can have one or more exit transition lines. A decision step must have at least two exit lines.

- 9 Right-click each final step in the template, then click **Add End Transition**.
- 10 Look at the message pane at the bottom of the iTRAC Template Builder to find any messages with warnings or errors about incomplete steps, then fix any problems you find.
- 11 When the template is complete, click **File > Save**.
or
Click the **Save** button to save the template.

17.6 Managing Templates

After creating a template, you can modify, copy, or delete it.

- ♦ [Section 17.6.1, “Viewing or Editing a Template,” on page 152](#)
- ♦ [Section 17.6.2, “Copying a Template,” on page 153](#)
- ♦ [Section 17.6.3, “Deleting a Template,” on page 153](#)

17.6.1 Viewing or Editing a Template

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 In the toolbar, click **iTRAC > Template Manager**.
- 3 Select a template, then click View/Edit to display the Template Builder.

17.6.2 Copying a Template

One way to create a new workflow template is to copy one of the default templates and modify it.

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select a template, then click Copy to display the Template Builder with the copied template
- 5 Specify a new name, then edit the template.

17.6.3 Deleting a Template

If you delete a template, any instantiated workflow processes that are based on that template finish normally.

To delete a template:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select a template, then click **Delete**.
- 5 Click **Yes** to confirm you want to delete the template.

17.7 Steps

Steps are the basic components of a template. Every template must have a start step and an end step. The start step exists by default. You can also add the following types of steps to a template:

- ♦ [Section 17.7.1, “Start Step,” on page 153](#)
- ♦ [Section 17.7.2, “Manual Steps,” on page 154](#)
- ♦ [Section 17.7.3, “Decision Steps,” on page 155](#)
- ♦ [Section 17.7.4, “Mail Steps,” on page 155](#)
- ♦ [Section 17.7.5, “Command Steps,” on page 155](#)
- ♦ [Section 17.7.6, “Activity Steps,” on page 156](#)
- ♦ [Section 17.7.7, “End Step,” on page 157](#)

17.7.1 Start Step

Every workflow template must have one start step. The transition from a start step is always unconditional.

17.7.2 Manual Steps

This type of step indicates that manual work must be performed. Every manual step in a template must be assigned to a role. The users in that role are notified through a worklist item when an instantiated workflow process reaches the manual step. When a user accepts the worklist item, it is removed from the queues of the other users in that role. For more information about worklists and stepping through a workflow process, see [Section 18.2, “Understanding the Work Item Summary Interface,” on page 177](#).

The description of the step should indicate what work needs to be performed. The user is expected to perform that work and then acknowledge completion.

A manual step includes the following attributes:

- ◆ Name of step
- ◆ Role
- ◆ Variables
 - ◆ Delete
 - ◆ Add
- ◆ Description

Variables

The user can also be asked to set one or more variables to appropriate values. Four variable types can be assigned to manual steps: Integer, Boolean, String, and Float. The variable can be set to an explicit default value during the step definition, or the user can set the value at run-time as part of the workflow process. The value can be optional or required.

The value of the variable can be used as part of a conditional transition to determine the path the workflow follows. It can also be used later as part of a conditional transition from a decision step to determine the workflow path.

NOTE: If the value is to be used later as part of a decision step, it should be marked Required.

For example, an integer variable can be set by the user to hold the event rate. Output transitions from the manual step can be defined so that if the event rate is greater than 500, one path is followed, and another path is followed if the event rate is less than 500.

To create a variable:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148](#).
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Click the **Add** button in upper left corner to open a new template.
or
Select an existing template, then click **View/Edit**.
- 5 Right-click **Variables** in the Process Tree, then select the type of variable to add.
- 6 Use the following information to define the selected variable:
Name: Specify a name for the variable.

Variable Type: Select the variable type. The options are:

- ◆ INTEGER
- ◆ BOOLEAN
- ◆ FLOAT
- ◆ STRING

Default Value: Specify a default value for the selected variable. The Boolean variable has only the options of True or False.

Description: (Optional) Specify a description for the variable.

7 Click **OK** to save the variable.

From a manual step, you can set conditional, unconditional, timeout, or alert transitions.

17.7.3 Decision Steps

This type of step selects different exit transitions, depending on the values of variables defined in prior steps. See [Section 17.7.2, “Manual Steps,” on page 154](#) for the available variable types. The decision step itself is very simple; you can edit only the step name and description. The workflow path is determined by the transitions.

From a decision step, you can set conditional and else transitions. Every decision step must have an else transition and at least one conditional transition. The else transition leads to a workflow path that is followed if none of the criteria for the conditional transitions are met.

17.7.4 Mail Steps

This step sends a pre-written e-mail. A mail step includes the following attributes:

- ◆ Name of step
- ◆ To addressee
- ◆ From addressee
- ◆ Subject of e-mail
- ◆ Body of e-mail

From a mail step, you can set a conditional, unconditional, timeout, alert, or error transition. An error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

17.7.5 Command Steps

A command step executes an operating system command or script (shell, batch, Perl and so on). The name of the command can be explicitly provided or set as a string variable, and parameters can be passed in the same manner. Output from the command can also be placed back into a string variable.

A command step includes the following attributes:

- ◆ Name of step
- ◆ Description

- ◆ Command (Can be explicit or variable-driven)
- ◆ Arguments (Can be explicit or variable-driven)
- ◆ Output Variable

NOTE: The command (a script file that refers to the command) must be stored in the `/opt/novell/sentinel/bin/actions` directory on the iTRAC workflow server. Symbolic links are not supported

Variables

The command output can also be used to set a variable to the appropriate values. Command steps must use String variable types.

The value of the variable can be used as part of a conditional transition to determine the path the workflow follows. It can also be used later as part of a decision step to determine the workflow path.

For example, a command step can return a value of 0 for failure and 1 for success. This output can be assigned to a variable, and then a conditional transition or a decision step can use this value to determine which workflow path to take.

The command and its arguments can each be specified explicitly by the person designing the workflow or can be set as a string variable. If either the command or the argument is set as a String variable, there must be a previous step in the template where the variable is set to a String value.

From a command step, you can set conditional, unconditional, timeout, alert, or error transitions. An error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

17.7.6 Activity Steps

An activity step is a type of automated step that can be used in a workflow template. Activity steps are created in the Activity Manager and can consist of internal Sentinel operations or external scripted operations. After activity steps are created, the user can select from a library of these activities and include them into a workflow. For more information on creating each type of predefined activity, see [Section 17.12, "Creating iTRAC Activities," on page 169](#).

An activity step includes the following attributes:

- ◆ Name
- ◆ Description
- ◆ Activity Assignment

From an activity step, you can set conditional, unconditional, timeout, alert, or error transitions. An error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

17.7.7 End Step

Every workflow template must have an end step to complete every branch of the workflow path.

17.8 Adding Steps to a Workflow

Steps can be added to a workflow by using the Step Palette or by right-clicking in the Process Builder. When you add steps to a workflow, a yellow entry field indicates an invalid entry.

- ◆ [Section 17.8.1, “Adding a Step from the Step Palette,” on page 157](#)
- ◆ [Section 17.8.2, “Adding a Step in the Process Builder,” on page 157](#)
- ◆ [Section 17.8.3, “Adding an Activity Step,” on page 158](#)
- ◆ [Section 17.8.4, “Adding an End Step,” on page 158](#)

17.8.1 Adding a Step from the Step Palette

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Click the **Add** button in upper left corner to open a new template.
or
Select an existing template, then click **View/Edit**.
- 5 Drag and drop a step from the Step Palette.
- 6 Right-click the step, then select **Edit Step**.
- 7 Edit the details of the step, then click **Save**.

17.8.2 Adding a Step in the Process Builder

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click an existing step in the Process Builder, then click **Insert New**.
- 6 Select the type of step you want to add:
 - ◆ Manual
 - ◆ Command Step
 - ◆ Mail Step
 - ◆ Decision Step
- 7 Edit the details of the step, then click **Save**.

17.8.3 Adding an Activity Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Drag and drop an activity from the Activity Pane to the Process Builder.

17.8.4 Adding an End Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click the last step with no transition, then select **Add End Transition**.

17.9 Managing Steps

Steps can be copied, edited, or deleted.

- ♦ [Section 17.9.1, “Copying a Step,” on page 158](#)
- ♦ [Section 17.9.2, “Modifying a Step,” on page 159](#)
- ♦ [Section 17.9.3, “Editing a Manual Step,” on page 159](#)
- ♦ [Section 17.9.4, “Editing a Decision Step,” on page 159](#)
- ♦ [Section 17.9.5, “Editing a Mail Step,” on page 160](#)
- ♦ [Section 17.9.6, “Editing a Command Step,” on page 160](#)
- ♦ [Section 17.9.7, “Deleting a Step,” on page 160](#)

17.9.1 Copying a Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, click **View/Edit**.
- 5 Right-click an existing step, then select **Copy Step**.
The step window opens in edit mode with all the attributes of the selected step.
- 6 Specify a name for the new step.
- 7 Edit step attributes as necessary, then click **OK**.

17.9.2 Modifying a Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click an existing step, then select **Edit Step**.
- 6 Edit the step attributes, then click **OK**.

17.9.3 Editing a Manual Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template that contains a manual step, then click **View/Edit**.
- 5 Right-click the manual step, then select **Edit Step**.
- 6 Change the name for the step.
- 7 Attach a role to this step by selecting a role from the drop-down list.
For more information on roles, see [Section 17.14, “Managing iTRAC Roles,” on page 171.](#)
- 8 Click **Associate** to associate a variable, then select the variable from the list or create new variables to be associated.
- 9 Set a default value for the variable as desired.
- 10 Check **Read-Only**, if this variable is to be forced to the default value.
- 11 Click the **Description** tab, then provide a description for this step.
- 12 Click **Preview** to preview the step you created.
- 13 Click **OK** to save the step.

17.9.4 Editing a Decision Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template that contains a decision step, then click **View/Edit**.
- 5 Right-click a decision step, then select **Edit Step**.
- 6 Change the name of the step.
- 7 Click the **Description** tab to provide a description for this step.
- 8 Click **OK** to save the step.

17.9.5 Editing a Mail Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2](#), “[Accessing the iTRAC Administration Tools](#),” on page 148.
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template that contains a mail step, then click **View/Edit**.
- 5 Right-click a mail step, then select **Edit Step**.
- 6 Change the name of the step.
- 7 Specify the **To** and **From** e-mail addresses, then specify a subject for the e-mail in the **General** tab.
- 8 Click the **Body** tab, then compose the e-mail message.
- 9 Click **OK** to save the step.

17.9.6 Editing a Command Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2](#), “[Accessing the iTRAC Administration Tools](#),” on page 148.
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template that contains a command step, then click **View/Edit**.
- 5 Right-click a command step, then select **Edit Step**.
- 6 Change the name for this step.
- 7 Specify the path and name of the command or script to execute relative to the `/opt/novell/sentinel/bin/actions` directory.
- 8 (Conditional) Select **Use Variables**, if you want to run a command or script referenced in a variable that is populated during the workflow process.
- 9 Specify any command line arguments to pass to the command or script.
- 10 (Conditional) Select **Use Variables**, if you want to run a command or script referenced in a variable that is populated during the workflow process.
- 11 Specify a variable to hold output from the command or script.
Any standard output is placed into these variables.
- 12 Click the **Description** tab to provide a description for this step.
- 13 Click **OK** to save the step.

17.9.7 Deleting a Step

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2](#), “[Accessing the iTRAC Administration Tools](#),” on page 148.
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click an existing step, then select **Delete Step**.
- 6 In the Alert Message window, select **Yes** to confirm deletion.

17.10 Transitions

Transitions are used to connect steps. There are several types of transitions:

- ◆ Unconditional
- ◆ Conditional
- ◆ Timeout
- ◆ Alert
- ◆ Else
- ◆ Error

A transition can have the following attributes:

- ◆ Name
- ◆ Description
- ◆ Destination: The step where the transition links
- ◆ Expression
- ◆ Timeout Values

Different steps have different properties and therefore they are associated with different transition types.

Table 17-2 Steps and Valid Transitions

Step Type	Valid Transitions
◆ Decision	◆ Conditional ◆ Else
◆ Manual	◆ Unconditional ◆ Timeout ◆ Alert
◆ Command	◆ Unconditional
◆ Mail	◆ Timeout
◆ Activity	◆ Alert ◆ Error

17.10.1 Unconditional Transitions

An unconditional transition must always be used from a start step. Manual, command, activity, and mail steps can also have unconditional transitions. The only parameter for an unconditional transition is the next step.

The transition is carried out when the current step is completed (unless a timeout transition is configured and the timeout period elapses).

To add an unconditional transition:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click an existing step, then select **Add Transition**.
- 6 Use the following information to create the unconditional transition:
 - Name:** Specify a name for the transition that is displayed in the Process Builder.
 - Type:** Select **Unconditional** for the transition type.
 - Destination:** Select the next step in the workflow for the unconditional transition.
 - Description:** Specify a description for the transition.
- 7 Click **OK** to save the transition.

17.10.2 Conditional Transitions

Select an exit path based on an expression using iTRAC variables set in a manual or command step.

You can add conditional transitions only from a decision step to any other step.

When you create a conditional transition, the conditional expressions can be based on comparing a variable that is populated during the workflow process to a specific value or to another variable populated during the workflow process. Multiple conditional expressions can be combined or nested using the AND and OR operators.

To add a conditional transition:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 Select an existing template, then click **View/Edit**.
- 4 Right-click an existing decision step, then select **Add Transition**.
- 5 Use the following information to create the conditional transition:
 - Name:** Specify a name for the conditional transition.
 - Type:** Select **Conditional** for the transition type.
 - Destination:** Select the next step in the workflow for the conditional transition.
 - Expression:** Create an expression for the conditional transition. See [Section 17.10.3, “Creating an Expression,” on page 163](#) for instructions.
 - Description:** Specify a description for the conditional transition.
- 6 Click **OK** to save the conditional transition.

17.10.3 Creating an Expression

Each conditional transition contains an expression that defines the condition. Use the following procedure to create a transition:

- 1 Verify you have completed [Step 2](#) through [Step 4](#).
- 2 Click **Set** to create the expression.
- 3 Click **EXP** to add the first expression.

The evaluation expression is an expression that evaluates to True or False during the workflow process.

- 4 Use the following information to define the expression:

Relations: Select how the expression compares the conditional transition:

- ◆ **Variables and Variable:** Compares a variable to another variable.
- ◆ **Variables and Values:** Compares a variable to a constant value.


Attribute: Select a variable from the drop-down list or create a new one if desired.

Condition: Select a condition from the drop-down list.

The condition list varies depending on the type of attribute variable chosen.

- ◆ **Boolean:** The options are:
 - ◆ equals
 - ◆ not equals
- ◆ **Float:** The options are:
 - ◆ is exactly
 - ◆ is not
 - ◆ is <
 - ◆ is <=
 - ◆ is >
 - ◆ is >=
- ◆ **Integer:** The options are:
 - ◆ is exactly
 - ◆ is not
 - ◆ is <
 - ◆ is <=
 - ◆ is >
 - ◆ is >=
- ◆ **String:** The options are:
 - ◆ startsWith
 - ◆ endsWith
 - ◆ equals
 - ◆ equalsIgnoreCase
 - ◆ matches
 - ◆ is empty
 - ◆ is not empty

Value: Either select an existing value or define a new value.

- 5 Click **OK**.
- 6 If a second expression is desired, select the root folder .
- 7 Repeat [Step 3](#) through [Step 5](#) as needed.
- 8 To nest expressions or to use the OR operator, click the appropriate operator button, then drag and drop expressions onto that operator.
By default, all expressions at the root level are separated by AND operators.
- 9 When the expression is complete, click **OK**.
You can edit or delete an existing expression by using the **Edit** and **Delete** buttons in the Expression window.
- 10 Continue with [Step 5 on page 162](#) to finish creating the conditional transition.

17.10.4 Else Transitions

An else transition leads to a path that is taken from a decision step when the criteria for the conditional transitions are not met. This transition only applies to decision steps, and every decision step must have an else transition. The workflow path with the else transition is only followed if none of the criteria for the conditional transitions is met.

You can add else transitions only from a decision step to any other step.

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, "Accessing the iTRAC Administration Tools," on page 148](#).
- 2 Click the **iTRAC** tab.
- 3 Select an existing template, then click **View/Edit**.
- 4 Right-click an existing Decision step, then select **Add Transition**.
- 5 Use the following information to create the else transition:
 - Name:** Specify a name for the else transition.
 - Type:** Select **Else** for the transition type.
 - Destination:** Select the next step in the workflow for the else transition.
 - Description:** Specify a description for the else transition.
- 6 Click **OK** to save the else transition.

17.10.5 Timeout Transitions

A timeout transition leads to a path that is taken when a user-specified amount of time (minutes, hours, or days) elapses after a base time, which is either `step_activated_time` or `step_accepted_time`. `Step_activated_time` is the time that iTRAC activates this step within the workflow process. `Step_accepted_time` is the time when a user accepts, or takes ownership, of the worklist item for this step. If the timeout period passes without the step being completed, control moves to the next step.

Timeout transitions can be set for a manual step or a command step. `Step_accepted_time` is only relevant for manual steps and should not be selected for a command step.

This transition is represented by a red line.

To add a timeout transition:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 Select an existing template, then click **View/Edit**.
- 4 Right-click an existing manual or command step, then select **Add Transition**.
- 5 Use the following information to create the timeout transition:

Name: Specify a name for the timeout transition.

Type: Select **Timeout** for the transition type.

Destination: Select the next step in the workflow for the timeout transition.

Timeout Details: Click **Set** to specify the timeout details:

- ♦ **Time:** Specify the timeout value.
- ♦ **Unit:** Select the unit of the timeout value. The options are:
 - ♦ Minutes
 - ♦ Hours
 - ♦ Days
- ♦ **Base Time:** Select the base time to use with the timeout transition.
 - ♦ **Step Accepted Time:** The time that iTRAC activates this step within the workflow process
 - ♦ **Step Activated Time:** The time when a user accepts, or takes ownership, of the worklist item for this step.

Description: Specify a description for the timeout transition.

- 6 Click **OK** to save the timeout transition.

17.10.6 Alert Transitions

An alert transition leads to a path that is taken when a user-specified amount of time (minutes, hours, or days) elapses after `step_activated_time` or `step_accepted_time`. At this point, the workflow process is usually escalated to a user who can intervene and take action.

`Step_activated_time` is the time that iTRAC activates this step within the workflow process.

`Step_accepted_time` is the time when a user accepts (or takes ownership) of the worklist item for this step.

If the alert time period passes without the step being completed, the workflow process branches into two active paths. The original step remains active for user intervention. The alert path is also initiated. For example, the alert path might escalate the workflow process to the attention of a supervisor, although the main path is still open and the original owner still has the option to complete the worklist item. Another example is that if a command is taking too long to run, you might want to alert an analyst to investigate the delay or possibly run the command manually.

Alert transitions can be set for a manual step or a command step. `Step_accepted_time` is only relevant for manual steps and should not be selected for a command step.

This transition is represented by a yellow line.

To add an alert transition:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, "Accessing the iTRAC Administration Tools," on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 Select an existing template, then click **View/Edit**.
- 4 Right-click an existing manual or command step, then select **Add Transition**.
- 5 Use the following information to create the alert transition:
 - Name:** Specify a name for the alert transition.
 - Type:** Select **Alert** for the transition type.
 - Destination:** Select the next step in the workflow for the alert transition.
 - Alert Details:** Click **Set** to specify the alert details:
 - ♦ **Time:** Specify the alert value.
 - ♦ **Unit:** Select the unit of the alert value. The options are:
 - ♦ Minutes
 - ♦ Hours
 - ♦ Days
 - ♦ **Base Time:** Select the base time to use with the alert transition.
 - ♦ **Step Accepted Time:** The time that iTRAC activates this step within the workflow process
 - ♦ **Step Activated Time:** The time when a user accepts, or takes ownership, of the worklist item for this step.
 - Description:** Specify a description for the alert transition.
- 6 Click **OK** to save the alert transition.

17.10.7 Error Transition

An error transition leads to a path that is taken if an automated step cannot successfully finish. Error transitions can be used for command, mail, and activity steps (for example, if a command step fails to execute).

Error transitions should typically lead to some kind of notification. For example, an error transition might lead to a manual step in which the user is instructed to manually run a process that previously failed.

The error transition is taken only if the iTRAC call to the command, mail, or activity step fails. If there is an internal error with the command script or the mail server fails, this does not satisfy the conditions for an error transition.

To add an error transition:

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, "Accessing the iTRAC Administration Tools," on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 Select an existing template, then click **View/Edit**.
- 4 Right-click an existing command, mail, or activity step, then select **Add Transition**.
- 5 Use the following information to create the error transition:

Name: Specify a name for the error transition.

Type: Select **Error** for the transition type.

Destination: Select the next step in the workflow for the error transition.

Description: Specify a description for the error transition.

6 Click **OK** to save the error transition.

17.10.8 Managing Transitions

After creating a transition, you can edit or delete the transition.

- ♦ [“Editing a Transition” on page 167](#)
- ♦ [“Deleting a Transition” on page 167](#)

Editing a Transition

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Double-click an existing transition line.
- 6 Edit the transition as needed.
- 7 Click **OK** to save the changes.

Deleting a Transition

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Template Manager**.
- 4 Select an existing template, then click **View/Edit**.
- 5 Right-click an existing step, then select **Remove Transition**.
- 6 Click **Yes** to confirm the deletions of the transition.

17.11 Activities

An activity is very similar to a command step, except that activities are reusable and cannot use input or output variables. The activities pane shows a library of user-defined, reusable activities that can reduce the amount of configuration necessary when building templates.

Activities are exported or imported as `.xml` files. These files can be exported or imported from one system to another.

iTRAC activities can be used in iTRAC templates to define a workflow step, or they can be manually executed from within an incident. Sentinel provides three types of actions that can be used to build activities:

- ♦ [Section 17.11.1, “Incident Command Activity,” on page 168](#)
- ♦ [Section 17.11.2, “Incident Internal Activity,” on page 168](#)
- ♦ [Section 17.11.3, “Incident Composite Activity,” on page 168](#)

17.11.1 Incident Command Activity

An incident command activity enables you to launch a specific command with or without arguments. The following fields from the incident associated with the workflow process can be used as input to the command:

- ♦ DIP [Target IP]
- ♦ DIP: Port
- ♦ RT1 (IDSAttackName)
- ♦ SIP [Initiator IP]
- ♦ SIP: Port
- ♦ Text (incident information in name value pair format)

The command (a script file that refers to the command) must be stored in the `/opt/novell/sentinel/bin/actions` directory on the Sentinel server.

17.11.2 Incident Internal Activity

An incident internal activity enables you to email and attach information from the Sentinel database to the incident associated with the workflow process. Each of these options has a prerequisite:

- ♦ **Vulnerability for the Initiator IP address (SIP) or the Target IP address (DIP):** Requires that you run a vulnerability scanner and bring the results of the scan into Sentinel by using a Vulnerability (or information) Collector
- ♦ **Advisor attack-related data:** Requires the purchase and installation of the optional Advisor data subscription service.
- ♦ **Asset data:** Requires that you run an asset management tool such as NMAP and bring the results into Sentinel by using an Asset Collector.

To send mail messages from within the Sentinel Control Center, you must have an SMTP integrator that is configured with connection information and with the `SentinelDefaultEMailServer` property set to `True`.

17.11.3 Incident Composite Activity

An incident composite activity enables you to combine one or more existing commands and internal activities.

17.12 Creating iTRAC Activities

1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, "Accessing the iTRAC Administration Tools,"](#) on page 148.

2 Click the **iTRAC** tab.

3 In the toolbar, click **iTRAC > Activity Manager**.

4 Select an existing activity, then click **View/Edit**.

or

Click the **Add** button to create a new activity.

5 Use the following information to define the activity type:

Type: Select the activity type you want to create:

- ◆ Incident Command Activity
- ◆ Incident Composite Activity
- ◆ Incident Internal Activity

Usage: The **Usage** field is populated with the type of activity you selected.

Name: Specify a name for the activity

Description: Specify a description for the activity.

6 Configure the necessary settings for the type of activity you chose:

Incident Command Activity: Use the following information to create the incident command activity:

- ◆ **Command:** Specify the command the activity performs.
- ◆ **Arguments:** Select the type of arguments for this command. The options are:
 - ◆ None
 - ◆ Incident Output
 - ◆ Custom

Incident Composite Activity: Select one or more of the activities to create a composite activity.

The options are:

- ◆ CollectionActivity
- ◆ EraditionaActivity
- ◆ ContainmentActivity

Incident Internal Activity: Configure the incident internal activity to e-mail the output to a specific address and attach the output to the incident associated with the workflow process.

- ◆ **Mail and Attach:** Select whether you want to e-mail the vulnerability and Advisor data.
- ◆ **Mail Details:** Specify the details of the e-mail that will be sent. You must define the **To**, **From**, and **Subject** fields.

7 View and confirm the details you chose in the Summary page, then click **Finish**.

17.13 Managing Activities

After creating an activity, you can modify, import or export it.

- ♦ [Section 17.13.1, “Editing an Activity,” on page 170](#)
- ♦ [Section 17.13.2, “Exporting an Activity,” on page 170](#)
- ♦ [Section 17.13.3, “Importing an Activity,” on page 170](#)

17.13.1 Editing an Activity

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Activity Manager**.
- 4 Select activity you want to edit, then click **View/Edit**.
- 5 Edit information in **General**, **Attachments** and **Mail** tabs.
- 6 Click **OK** to save the changes.

17.13.2 Exporting an Activity

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Activity Manager**.
- 4 Click **Import/Export**.
- 5 Use the following information to export the activity:
 - Action:** Select **Export Activity**.
 - File Name:** Click **Explore**, then browse to and select the file you want to export this information to.
 - File Path:** This is automatically populated when you select the export file.
- 6 Click **Next**.
- 7 Select one or more activities to be exported.
- 8 Click **Next**, then click **Finish**.

17.13.3 Importing an Activity

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148.](#)
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Activity Manager**.
- 4 Click **Import/Export**.
- 5 Use the following information to import the `.xml` file that contains the activity:
 - Action:** Select **Import Activity**.
 - File Name:** Click **Explore**, then browse to and select the file you want to import.

File Path: This is automatically populated when you select the import file.

6 Click **Next**.

The list of activities is displayed.

7 Click **Next**, then click **Finish**.

17.14 Managing iTRAC Roles

Sentinel users can be assigned to one or more roles. Manual steps in the workflow processes can be assigned to a role.

- ♦ [Section 17.14.1, “Adding a Role,” on page 171](#)
- ♦ [Section 17.14.2, “Deleting a Role,” on page 171](#)
- ♦ [Section 17.14.3, “Viewing the Role Details,” on page 171](#)

17.14.1 Adding a Role

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148](#).
- 2 Click the **iTRAC** tab.
- 3 In the menu bar, click **iTRAC > iTRAC Role Manager**.
- 4 Click **Add Role**.
- 5 Specify a unique name for the role.
- 6 Click **Add** to specify which users must be members of this role.
- 7 Click **OK**.

17.14.2 Deleting a Role

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148](#).
- 2 Click the **iTRAC** tab.
- 3 In the menu bar, click **iTRAC > iTRAC Role Manager**.
- 4 Right-click the role you want to delete, click **Delete Role**.
- 5 Click **Yes** to confirm deletion.

17.14.3 Viewing the Role Details

- 1 Access the Sentinel Control Center. For more information, see [Step 1 in Section 17.2, “Accessing the iTRAC Administration Tools,” on page 148](#).
- 2 Click the **iTRAC** tab.
- 3 In the menu bar, click **iTRAC > iTRAC Role Manager**.
- 4 Right-click the role for which you want to view details, click **iTRAC Role Details**.

17.15 Process Management

Process management allows you to view the incident's progress in the workflow or terminate a workflow process.

Process execution is the time period during which the process is operational, with process instances being created and managed.

When an iTRAC process is executed or instantiated in the iTRAC server, a process instance is created, managed, and eventually terminated by the iTRAC server in accordance with the process definition. As the process progresses towards completion or termination, it executes various activities defined in the workflow template based on the criteria for the transitions between them. The iTRAC workflow server processes manual and automatic steps differently.

An iTRAC process must be created with a single associated incident; there is therefore a one-to-one match between iTRAC processes and incidents. Not all incidents are attached to a process. Only one incident can be associated to an iTRAC process instance.

- ◆ [Section 17.15.1, "Instantiating a Process," on page 172](#)
- ◆ [Section 17.15.2, "Automatic Step Execution," on page 172](#)
- ◆ [Section 17.15.3, "Manual Step Execution," on page 173](#)
- ◆ [Section 17.15.4, "Display Status," on page 173](#)
- ◆ [Section 17.15.5, "Displaying the Status of a Process," on page 173](#)
- ◆ [Section 17.15.6, "Changing Views in Process Manager," on page 174](#)
- ◆ [Section 17.15.7, "Starting or Terminating a Process," on page 174](#)

17.15.1 Instantiating a Process

An iTRAC process can be instantiated in the iTRAC server by associating an incident to an iTRAC process by any of the following three methods

- ◆ Associating an iTRAC process to the incident at the time of incident creation
- ◆ Associating an iTRAC process to incident after an incident has been created
- ◆ Associating an iTRAC process to an incident through correlation

For more information on associating a process to an incident, see [Chapter 16, "Configuring Incidents," on page 143](#).

17.15.2 Automatic Step Execution

When the process instance executes an automatic activity step, command step, or mail step, it executes the associated activity or command defined in the template and stores the result in process variables. It then transitions to the next step in the iTRAC template.

For example, an activity might be defined to ping a server. When this activity is executed in a workflow process, the activity runs and attaches the results to the associated incident.

17.15.3 Manual Step Execution

On encountering a manual step, the iTRAC server sends out notifications in the form of work items to the assigned resource. If the step was assigned to a role, a work item is sent to all users within the role. The iTRAC server then waits for the user to complete the work item before proceeding to the next activity.




For more information, see [Section 18.2, “Understanding the Work Item Summary Interface,”](#) on [page 177](#).

17.15.4 Display Status

The Display Status option monitors the progress of a process. As the process instance progresses from one activity, the user can track the progress visually by clicking the Refresh button. The process monitor also provides an audit trail of all the actions performed by the iTRAC server when executing the process.


Activities that are running, completed, or terminated are represented by the following icons.

Table 17-3 Display Status Icons

Icon	Description
	Indicates that the activity is running.
	Indicates that the activity is completed.
	Indicates that the activity is terminated.

17.15.5 Displaying the Status of a Process

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2, “Accessing the iTRAC Administration Tools,”](#) on [page 148](#).
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Display Process Manager**.
- 4 Click the down-arrow on the **Switch Views** button to select a view or create a new view.
- 5 In the Process Manager window, right-click a process, then select **Actions > Display Status**.
The current step is highlighted in red.



6 To close, click the X in the upper right corner.

17.15.6 Changing Views in Process Manager

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2](#), “[Accessing the iTRAC Administration Tools](#),” on page 148.
- 2 Click the **iTRAC** tab.
- 3 In the toolbar, click **iTRAC > Display Process Manager**.
- 4 Click the **Manage View** drop-down list, then select **Edit Current View**.
- 5 Change the following options to change your current view:
 - ◆ Fields
 - ◆ Group by
 - ◆ Sort
 - ◆ Filter
 - ◆ Leaf Attribute
- 6 Click **Apply**, then click **Save**.

17.15.7 Starting or Terminating a Process

- 1 Access the Sentinel Control Center. For more information, see [Step 1](#) in [Section 17.2](#), “[Accessing the iTRAC Administration Tools](#),” on page 148.
- 2 Click the **iTRAC** tab.

- 3 In the toolbar, click **iTRAC > Display Process Manager**.
- 4 In the Process View Manager window, right-click a process, then select **Actions > Start Process** or **Terminate Process**.

18 Managing Work Items

- ◆ [Section 18.1, “Overview,” on page 177](#)
- ◆ [Section 18.2, “Understanding the Work Item Summary Interface,” on page 177](#)
- ◆ [Section 18.3, “Viewing a Work Item,” on page 178](#)
- ◆ [Section 18.4, “Processing a Work Item,” on page 179](#)
- ◆ [Section 18.5, “Managing Work Items Of Other Users,” on page 179](#)

18.1 Overview

A work item is a workflow task assigned to a particular user or role in the iTRAC application. The individual activities to be performed to complete an iTRAC process are listed as work items in Work Item Summary in the Sentinel Control Center. For more information on iTRAC processes, see [Chapter 17, “Configuring iTRAC Workflows,” on page 147](#). You can access the work items from any tab in the Sentinel Control Center.

NOTE: To have access to a work item, you must assign it to you or acquire the work item management permission. If you have the work item management permission, you can manage work items of other users.

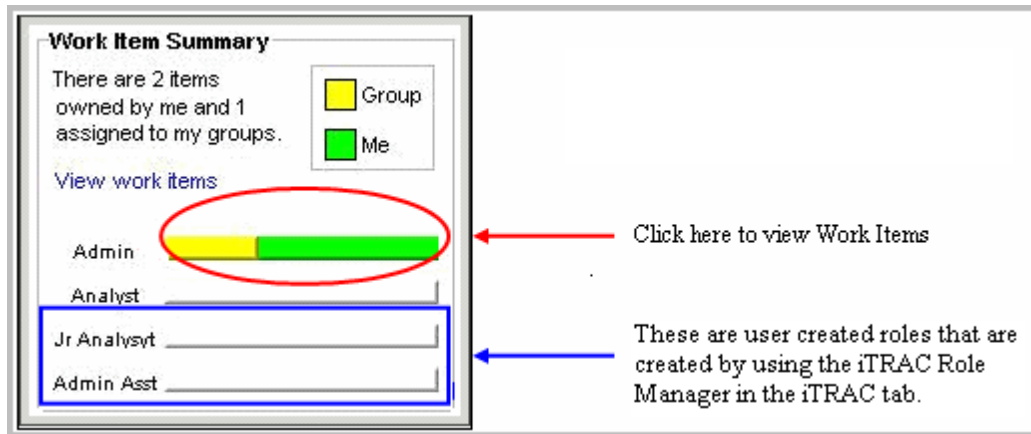
18.2 Understanding the Work Item Summary Interface

The Work Item Summary lists the work items allocated to a user as an individual and as a member of a group. It can be used as an incident workflow to-do list for a user who is a part of the incident response process. In the Work Item Summary, you can access the work items and perform different tasks:

- ◆ View the details of a work item
- ◆ Process the work item to complete the task


In the Work Item Summary, work items are grouped by current user and by other users with similar roles. The following example is for a user who is a member of the Admin, Analyst, Jr Analyst, and Admin Asst groups.

Figure 18-1 Work Item Summary Example 1



The following example is for a user who is a member of the Analyst group who has a process assigned to his role (group).

Figure 18-2 Work Item Summary Example 2



18.3 Viewing a Work Item

- 1 In the Work Item Summary, click the yellow or green bar.

A work item list for the group or the current user displays and shows the name and ID of the incident, the workflow process name, and the step name and description

- 2 Double-click any work item and click **View Details**.

The Work Item Details window appears and shows the process details, including any detailed instructions included by the iTRAC workflow developer and any variables that need to be set in the step.

- 3 Click **Process Overview** to view an overview of the entire iTRAC process.

- 4 Click **Incident** to view the details of the associated incident.

- 5 To take responsibility for this work item, click **Acquire**. Otherwise, click **Cancel**.

NOTE: Any changes to the incident from this page must be saved. There is a **Save** button on the toolbar and a **Save** button at the bottom of the page.

The information on the **Process Details** and **Process Overview** tabs is defined by the iTRAC workflow designer. For more information on creating workflow templates, see [Chapter 17, "Configuring iTRAC Workflows,"](#) on page 147.

18.4 Processing a Work Item

A work item can be accessed from any part of the main tabbed Sentinel Control Center interface.

- ♦ You can process a work item in a group even if you have logged in as a different user. However, you cannot acquire a step if you have logged in as a different user.
- ♦ The work item remains with the user of a group who has acquired it.
- ♦ Consecutive steps are dependent. If two consecutive steps are assigned to the same role, the user who acquires the first step is also assigned the second step.
- ♦ Non-consecutive steps are independent. For example, if a workflow proceeds from steps that are assigned to the Tier 1 Analyst group to the Tier 2 Analyst group and then back to the Tier 1 Analyst group, the third step is available to the entire Tier 1 Analyst group. It is not assigned to the individual user who handled the first step.

To process a work item, you must accept and complete the work item:

- 1 In the Work Item Summary, click the yellow or green bar. A work item list for the group or the current user displays.
- 2 To assign an iTRAC process to you, select the process and click **Acquire**.
The Work Item Summary changes from yellow to green.

NOTE: When you acquire (accept) a work item, it is removed from the queues of all other users in the same role. The work item can be returned to the group by clicking **Release**.

- 3 Click **View Details**.

The current step within a work item is highlighted in red.

- 4 To take action on the step, click the **Process Details** tab.

Depending on the type of variable (Integer, String, Boolean and Float) in a manual step, you can click the down-arrow and select a value. If needed, you can add comments or add an attachment.

In all other cases, the steps are automatic.

- 5 Click **Complete** to complete the process.

Completing the work item signals the completion of the task to the iTRAC server. The updateable variables from the work item are processed by the server to move to the next step, which depends on how the workflow is defined. The work item is removed from the user's worklist and appears in the worklist of the individual or role associated with the next step in the process.

18.5 Managing Work Items Of Other Users

The Administration function allows an administrative user to release a work item from a specific user back to everyone in a role. This is beneficial if a work item is already in process but the assigned user cannot complete the work.

- 1 Log in into Sentinel as a user with iTRAC – Manage Work Items Of Other Users user rights.
- 2 In the Summary pane, click **View work items**.
- 3 In the Work Items window, set the following:
 - ♦ **User:** Name of the user that has acquired the process
 - ♦ **Group:** Name of the group that the user belongs to. In this example, the user belongs to the Analyst group.

- ♦ **Owner:** Select **<All>** (all processes acquired or not), **me** (acquired processes), or **Group** (un-acquired processes).
- ♦ **Process:** Name of the process.

In this example, all processes are acquired by jr1, who belongs to the Analyst group.

- 4 To release the work item, select the work item and click Release. Release changes to Acquire (not available).

In this example, only a member of the Analyst group can acquire this work item.

A Search Query Syntax

Sentinel uses the Lucene query language for searching events. This section provides an overview of how to use the Lucene query language to perform searches in Sentinel. For more advanced features, see [Apache Lucene - Query Parser Syntax](#).

For information on the event fields in Sentinel, click **Tips** on the top right corner in the Sentinel Web interface. A table is displayed that lists the event names and their IDs.

- ◆ [Section A.1, “Basic Search Query,” on page 181](#)
- ◆ [Section A.2, “Wildcards in Search Queries,” on page 186](#)
- ◆ [Section A.3, “The notnull Query,” on page 188](#)
- ◆ [Section A.4, “Tags in Search Queries,” on page 188](#)
- ◆ [Section A.5, “Regular Expression Queries,” on page 189](#)
- ◆ [Section A.6, “Range Queries,” on page 189](#)
- ◆ [Section A.7, “IP Addresses Query,” on page 190](#)

A.1 Basic Search Query

A basic query is a search for a value on a field. The syntax is as follows:

```
msg:<value>
```

The field name (msg) is separated from the value by a colon.

For example, to search for a phrase that includes the word “authentication,” you can specify the search query as follows:

```
msg:authentication
```

Or, to search for events of severity 5, you can specify the search query as follows:

```
sev:5
```

If the value has spaces or other delimiters in it, you should use quotation marks. For example:

```
msg:"value with spaces"
```

Sentinel classifies event fields as either tokenized fields or non-tokenized fields. A tokenized field is indexed and is searched differently than a non-tokenized field.

- ◆ [Section A.1.1, “Case Insensitivity,” on page 182](#)
- ◆ [Section A.1.2, “Special Characters,” on page 182](#)
- ◆ [Section A.1.3, “Operators,” on page 182](#)
- ◆ [Section A.1.4, “The Default Search Field,” on page 183](#)
- ◆ [Section A.1.5, “Tokenized Fields,” on page 184](#)
- ◆ [Section A.1.6, “Non-Tokenized Fields,” on page 186](#)

A.1.1 Case Insensitivity

Indexing and searching in Sentinel is not case-sensitive. For example, the following queries are all equivalent:

```
msg:AdMin
msg:admin
msg:ADMIN
```

A.1.2 Special Characters

If you include special characters as part of a search, the special characters must be escaped. These characters are as follows:

```
+ - && | | ! ( ) { } [ ] ^ " ~ * ? : \ /
```

Use “\” before the character you want to escape. For example, to search for ISO/IEC_27002:2005 in the rv145 (Tag) field, use the following query:

```
rv145:ISO\IEC_27002\:2005
```

You can also use quotation marks around the query:

```
rv145:"ISO/IEC_27002:2005"
```

If the value contains quotation marks, you must escape it by using the “\” character instead of quotation marks. For example, to search for “system “mail” service” in the `initiatorservicename` field, you must specify the query as follows:

```
sp:"system \"mail\" service"
```

For more information on quoting wildcard characters, see [Section A.2.2, “Quoted Wildcards,” on page 187](#).

A.1.3 Operators

Lucene supports AND, OR, and NOT Boolean operators, which allow words to be combined. Boolean operators must be always capitalized.

- ◆ [“OR Operator” on page 182](#)
- ◆ [“AND Operator” on page 183](#)
- ◆ [“NOT Operator” on page 183](#)
- ◆ [“Operator Precedence” on page 183](#)

OR Operator

The OR operator is the default conjunction operator. If there is no Boolean operator between two clauses, the OR operator is used. The OR operator links two clauses and finds a matching event if either of the clauses is satisfied. The symbol `||` can be used in place of the word OR. For example, consider the following query:

```
sun:admin OR dun:admin
```

This query finds events whose initiator username or target username is “admin.” The following query produces the same result because OR is used by default:

```
sun:admin dun:admin
```

AND Operator

The AND operator links two clauses and finds a matching event only if both clauses are satisfied. The symbol && can be used in place of the word AND. For example, consider the following query:

```
sun:admin AND dun:tester
```

This query finds events whose initiator username is admin and the target username is tester.

NOT Operator

The NOT operator excludes events that match the clause after the NOT. The symbol ! can be used in place of the word NOT. For example, consider the following query:

```
sev:[0 TO 5] NOT st:I NOT st:A NOT st:P
```

This query matches all events whose severity is between 0 and 5, but excludes those whose sensor type is I (internal), A (audit), or P (performance); that is, it excludes Sentinel internal events.

The NOT operator cannot be used by itself because it is a way to exclude events from a set that has been found by other search terms. For example, consider the following query:

```
NOT st:I NOT st:A NOT st:P
```

This query might seem like it should return all events where the sensor type is not I, A, or P. However, it is an invalid query because a query cannot begin with the NOT operator.

Operator Precedence

Parentheses can be used in the usual way to change operator precedence. They can be nested to any depth, as shown in the following examples:

```
(sun:admin OR dun:admin) AND (sip:10.0.0.1 OR sip:10.0.0.2)
((sun:admin OR dun:admin) AND (sip:10.0.0.1 OR sip:10.0.0.2)) OR (msg:user AND
evt:authentication)
```

A.1.4 The Default Search Field

Lucene uses a default search field, which is the field that is searched if no field is specified. In Sentinel, `_data` is the default search field. This field is a concatenation of all non-empty fields in the event. It is indexed and searched as a tokenized field. The result is that you can search for words that might appear in any event field.

For example, suppose you have two non-tokenized fields in an event, `sun` (initiatorusername) and `dun` (targetusername). The `sun` field has the following value:

```
report-administrator
```

The `dun` field has the following value:

```
system-tester
```

The `_data` field contains the concatenation of these fields separated by a single space character:

```
report-administrator system-tester
```

Because the `_data` field is a tokenized field, the words “report,” “administrator,” “system,” and “tester” are indexed and searchable. The following queries would find this event:

```
report
_data:report
report-administrator
_data:report-administrator
report tester
```

In addition, the following queries also find this event:

```
sun:report-administrator
dun:system-tester
```

A.1.5 Tokenized Fields

Fields that are classified as tokenized fields are parsed into individual words for indexing. Therefore, a search occurs only on words within the field value. Characters that are considered to be word delimiters are not searchable, nor are words that are considered to be stop words. Lucene removes extremely common words to save disk space and speed up searching. These words are ignored in search filters. Currently, the following stop words are removed:

- ◆ a
- ◆ an
- ◆ and
- ◆ are
- ◆ as
- ◆ at
- ◆ be
- ◆ but
- ◆ by
- ◆ for
- ◆ if
- ◆ in
- ◆ into
- ◆ is
- ◆ it
- ◆ no
- ◆ not
- ◆ of
- ◆ on
- ◆ or
- ◆ such
- ◆ that
- ◆ the
- ◆ their
- ◆ then

- ♦ there
- ♦ these
- ♦ they
- ♦ this
- ♦ to
- ♦ was
- ♦ will
- ♦ with

When it does a search, Lucene examines all of the words in a field and tries to match words in the search value. For example, suppose that you specify a search for messages containing the following value:

```
msg:"user-authentication failed on the server"
```

The words that are parsed within this value are “user,” “authentication,” “failed,” and “server.” These are the only search words that would match this value. “On” and “the” are omitted because they are stop words.

The value has the hyphen character (-) between some words. Hyphens are treated as word delimiters, so Lucene does not search for hyphens. Consider, the following query:

```
msg:"user-authentication"
```

The results might not be exactly what you expect. The query search value matches the value, but not because it is matching the hyphen. It matches because Lucene first parses the words in the search value and identifies the words “user” and “authentication.” Lucene then matches those words against values that have the words “user” and “authentication” with no intervening words in between. This query would also match the following value, even though there is no hyphen between “user” and “authentication”:

```
user authentication has failed on the server
```

Consider the following query:

```
msg:"failed on server"
```

This query has the stop word, “on,” which is ignored. However, the stop word does affect the relative positioning that is expected to be between words when evaluating a value to see if it matches. The “failed on server” search matches any phrase where the words “failed” and “server” are separated by exactly one word. It does not matter what the word is because the separating word is a stop word and is ignored. Thus, the above query would match all of the following:

```
failed on server
```

```
failed-on server
```

```
failed a server
```

```
failed-a-server
```

Proximity indicators created by using the ~ character followed by a value, make this more complicated. The query dictates an expected distance between words. In the “failed on server” query, the expected distance between “failed” and “server” is one word. The proximity indicator specifies how much variance there can be from the expected distance. For example, consider the following query, where a proximity indicator of one (~1) is specified:

```
msg:"failed on server"~1
```

This query indicates that the distance between “failed” and “server” could be plus or minus one from the expected distance, which is one because of the stop word “on.” Thus, the distance could be 1, 1-1 (0), or 1+1 (2). Thus, all of the following would match:

```
failed on server
```

```
failed on the server
```

```
failed finance server
```

As of Lucene version 3.1, word parsing is done according to word break rules outlined in the Unicode Text Segmentation algorithm. For more information, see [Unicode Text Segmentation](#).

For information on tokenized fields in Sentinel, in the Sentinel Web interface click **Tips** on the top right corner of the Sentinel Web interface. A table is displayed that lists all the event fields and whether an event field is searchable or not.

A.1.6 Non-Tokenized Fields

Fields that are classified as non-tokenized fields are parsed fully for indexing. Thus, a search occurs on full field values. For example, to search events whose initiatoruserfullname (iufname) field has the value “Bob White”, you must specify the query as follows:

```
iufname:"Bob White"
```

A.2 Wildcards in Search Queries

Sentinel supports wildcards in search values but not in regular expressions:

- ◆ The asterisk (*) matches zero or more characters.
- ◆ The questions mark (?) matches any one character.

For example:

- ◆ **adm*test:** Matches admtest, ADMTEST, admintest, adMINtEst (note the lack of case sensitivity).
- ◆ **adm?test:** Matches adm1test and AdMatest. Does not match admtest or ADMINTEST because it must have exactly one character between "adm" and "test."
- ◆ [Section A.2.1, “Wildcards in Tokenized Fields,” on page 186](#)
- ◆ [Section A.2.2, “Quoted Wildcards,” on page 187](#)
- ◆ [Section A.2.3, “Leading Wildcards,” on page 187](#)

A.2.1 Wildcards in Tokenized Fields

Wildcards are applied differently to tokenized fields and non-tokenized fields. Wildcards for tokenized fields match only words that were parsed from the value and not the entire value. For example, if you specify the search query `msg:authentication*failed` to search for the message `The user authentication has failed on the server`, it does not return the events with this message. This is because “*” does not match anything between “authentication” and “failed.” However, it matches any words that begin with “authentication” and end with “failed.” For example, it returns results if any of the following words are used: “authenticationhasfailed,” “authenticationuserfailed,” and “authenticationserverfailed.” For tokenized fields, all matching that uses wildcard searches is done on the words within the value and not on the full value.

A.2.2 Quoted Wildcards

- ◆ [“Tokenized Fields” on page 187](#)
- ◆ [“Non-Tokenized Fields” on page 187](#)

Tokenized Fields

When wildcards are quoted, they are not treated as wildcards, but as word delimiters. For example, consider the following query:

```
msg:"user* fail*"
```

The search value `"user* fail*"` is parsed into two words, “user” and “fail.” The semantic is “find any event where the `msg` field contains “user” AND “fail” words in that order, and there are no intervening words between them.” Thus, it does not match the following value:

```
The user authentication has failed on the server.
```

This is because the wildcard is not treated as a wildcard but as a word delimiter.

Non-Tokenized Fields

When wildcards are quoted, they are treated as literal characters to search. For example, if the query is: `sun:"adm*,"` it returns the following values:

```
adm*
```

```
ADM* (case-insensitive)
```

The query does not return the following values:

```
admin
```

```
ADMIN
```

A.2.3 Leading Wildcards

Leading wildcards are not valid in searches because Lucene does not allow the `*` or `?` characters to be the first character of a search value. For example, the following queries are invalid:

- ◆ **sun:*adm*** The semantic is “find any event whose initiator username value contains the letters a, d, and m in sequence.”
- ◆ **sun:*tester** The semantic is “find any event whose initiatorusername value ends with “tester.”
- ◆ **sun:*** The semantic is “find any event whose initiator username field is non-empty.”

Because this is an important type of query, Sentinel provides an alternative way to accomplish this. For more information, see [Section A.3, “The notnull Query,” on page 188](#).

A.3 The notnull Query

You might need to find events where some field is present, or non-empty. For example, to find all events that have a value in the sun field, you can specify the query as `sun:*`

The query does not return the expected results because Lucene does not support wildcards to be the first character of a search value. However, Sentinel provides an alternate solution. For every event, Sentinel creates a special field called `notnull`. The `notnull` field is a list of all fields in the event that are not null (not empty). For example, if there is an event that has values in the `evt`, `msg`, `sun`, and `xdasid` fields, the `notnull` field contains the following value:

```
evt msg sun xdasid
```

The `notnull` field is a tokenized field, so the following kinds of queries are possible:

- ◆ **notnull:sun** Finds all events whose sun field has a value.
- ◆ **notnull:xdas*** Finds all events where any field beginning with the name "xdas" has a value.

When a `notnull` field is added in Lucene, creating, indexing, and storing this field adds a cost to processing each event as CPU needs to create and index the field and it also requires additional storage space. If you want to disable storing the list of non-empty fields in the `notnull` field, set the following property in the `/etc/opt/novell/sentinel/config/configuration.properties` file:

```
indexedlog.storenotnull=false
```

Save the file and restart the Sentinel server. All events received after this property was set do not have a `notnullfield` associated.

NOTE: If you disable the `notnull` field, do not use the `notnull` field in search filters, rule filters, or policy filters because the results might be incorrect and unpredictable.

A.4 Tags in Search Queries

The Tag field (`rv145`) is a tokenized field that has special parsing rules for words. The parsing rules enable you to search on tags that include non-alphanumeric characters. However, the only word delimiters are white space characters such as the blank and the tab. This is because tags do not include white space in their names. For example, the following queries find the event if the event is tagged with the `ISO/IEC_27002:2005` tag and the `NIST_800-53` tag:

```
rv145:"ISO/IEC_27002:2005"
```

```
rv145:"iso/iec_27002:2005"
```

```
rv145:"ISO/IEC_27002*"
```

```
rv145:nist_*
```

The slash (/), hyphen (-), and colon (:) characters are significant in the search value because, unlike other tokenized fields, the parsing rules for `rv145` do not treat them as a word delimiter. Also, the search is not case sensitive.

The following queries would not find the event:

```
rv145:"ISO IEC_27002 2005"
```

```
rv145:"iso *"
```

A.5 Regular Expression Queries

Regular expression queries allow you to search events that match a pattern. For example, to search for an initiator user name that ends with the character "a", you can specify the search query as follows:

```
sun:/. *a/
```

If you need to include special characters in your query, you must escape special characters by preceding them with the backslash (\) character. For example, to search for an initiator user name that ends with the character "\$", you can specify the search query as follows:

```
sun:/. *\$/
```

For more information about using special characters, see [Section A.1.2, "Special Characters," on page 182](#).

NOTE: Regular expression queries utilize significantly more system resources than other kinds of queries because they are unable to leverage the more efficient data structures available in the index. Executing regular expression queries take longer than other kinds of queries and potentially pull system resources from other components of the system. Therefore, use regular expression queries carefully and narrow the breadth of the search as much as possible by using time range and non-regular expression criteria terms.

A.6 Range Queries

Range queries allow you to find events where a field value is between a lower bound and an upper bound. Range queries can be inclusive or exclusive of the upper and lower bounds. Whether a particular value falls in the specified range is based on lexicographic character sorting. Inclusive ranges are denoted by square brackets []. Exclusive ranges are denoted by curly brackets {}.

For example, consider the following query:

```
sun:[admin TO tester]
```

This query finds events whose sun field has values between admin and tester, inclusive. Note that "TO" is capitalized.

However, if you change the query as follows:

```
sun:{admin TO tester}
```

The query now finds all events whose sun field is between admin and tester, not including admin and tester.

Some event fields such as sev and xdasid are numeric. In Sentinel, range queries on numeric fields are based on numeric sorting and not on lexicographic character sorting. For example, consider the following query:

```
xdasid:[1 TO 7]
```

This query returns events whose xdasid value is 1, 2, 3, 4, 5, 6, or 7. If the range evaluation was based on lexicographic sorting, it would incorrectly match 10, 101, 100001, 200, and so on.

A.7 IP Addresses Query

There are several extensions that Sentinel has implemented for searching on IP addresses. Specifically, there are a number of convenient ways to specify IP address ranges. These are explained in the following sections:

- ◆ [Section A.7.1, “CIDR Notation,” on page 190](#)
- ◆ [Section A.7.2, “Wildcards in IP Addresses,” on page 190](#)

A.7.1 CIDR Notation

Sentinel supports the Classless Inter-Domain Routing (CIDR) notation as a search value for IP address fields such as sip (initiator IP) and dip (target IP) for specifying an IP address range. The notation uses a combination of an IP address and a mask, as follows:

```
"xxx.xxx.xxx.xxx/n"
```

In this notation, n is the number of high order bits in the value to match. For example, consider the following query:

```
sip:"10.0.0.0/24"
```

This query returns events whose sip field is an IPv4 address ranging from 10.0.0.0 to 10.0.0.255.

The same notation works for IPv6 addresses. For example, consider the following query:

```
sip:"2001:DB8::/48"
```

This query returns events whose sip field is an IPv6 address ranging from 2001:DB8:: to 2001:DB8:0:FFFF:FFFF:FFFF:FFFF:FFFF.

A.7.2 Wildcards in IP Addresses

You can use only the asterisk character (*) in the IP address search values to specify ranges of IP addresses. You cannot use the question mark (?) character.

In IPv4 addresses, an asterisk (*) can be used at any of the positions in the quad format. In IPv6 addresses, an asterisk (*) can be used between colons to specify a 16-bit segment. For example, all of the following queries are valid on the sip field:

```
sip:10.*.80.16
```

```
sip:10.02.*.*
```

```
sip:10.*.80.*
```

```
sip:CAFE*::FEED
```

```
sip:CAFE*:FADE*::FEED
```

If an asterisk (*) is used in one of the quad positions in an IPv4 address or between colons in an IPv6 address, it cannot be combined with other digits. For example, all of the following queries are invalid:

```
sip:10.*7.80.16
```

```
sip:10.10*.80.16
```

```
sip:CAFE:FA*::FEED
```

```
sip:CAFE:*DE::FEED
```

Because the question mark (?) is not allowed, the following queries are invalid:

`sip:10.10?.80.16`

`sip:10.?.80.16`

`sip:CAFE:FA??:FEED`

`sip:CAFE:??DE:FEED`

B Correlation Rule Expression Syntax

Correlation rules are written to match specific events or sequences of events by using field references, comparison and match operators on the field contents, and operations on sets of events.

The Correlation Engine loads the rule definition and uses the rules to evaluate, filter, and store events in memory that meet the criteria specified by the rule. Depending on the rule definition, a correlation rule might fire according to several different criteria:

- ♦ The value of one field or multiple fields
- ♦ The comparison of an incoming event to past events
- ♦ The number of occurrences of similar events within a defined time period
- ♦ One or more subrules firing
- ♦ One or more subrules firing in a particular order

This section provides a basic overview of how to build Correlation rules and the various parameters required to build a rule.

- ♦ [Section B.1, “Event Fields,” on page 193](#)
- ♦ [Section B.2, “Event Operations,” on page 194](#)
- ♦ [Section B.3, “Operators,” on page 201](#)
- ♦ [Section B.4, “Order of Operators,” on page 202](#)

B.1 Event Fields

All operations function on event fields, which can be referred to by their names or by their IDs within the rule expression. For a full list of event field names and their IDs, in the Sentinel Web interface, click **Tips** on the top right corner of the Sentinel Web interface.

The event field name or its ID must also be combined with a prefix to designate whether the event field is part of the current event (*e*) or a past event that is stored in memory. For the Window operator, the stored events are prefixed with (*w*).

Examples:

```
e.dip (Destination IP for the current event)
w.dip (Destination IP for any stored event)
```

IMPORTANT: If you rename an event field by using the Event Configuration utility in the Sentinel Control Center, use the new name when writing rules. In all cases, rules are stored internally with the fixed IDs and the names are translated dynamically when viewed.

B.2 Event Operations

Event operations evaluate, compare, and count events. Each operation works on a set of events, receiving a set of events as input and returning a set of events as output. The current event processed by a rule often has a special meaning within the language semantics. The current event is always part of the set of events output by the operation unless the set is empty. If an input set of an operation is empty, then the operation is not evaluated.

The individual correlation operations evaluate sets of events and, if matches are produced, generates sets of corresponding events. You can chain together multiple correlation operations by using the flow operator. The most common use of a chain construct is to begin a chain with a filter to select a subset of events. The subset is then grouped by operations, such as Trigger, or by cross-event comparisons, such as Window. For more information on the flow operator, see [Section B.3.1, “Flow Operator,” on page 201](#).

```
<operation 1> [flow <operation N> ...]
```

<operation 1> is a fully-specified operation.

In the above example, the output from one operation is treated as input to the next operation, much like the UNIX pipe operator. The last operation in a chain's output causes the whole rule to fire, generating a correlated event.

- ◆ [Section B.2.1, “Filter Operation,” on page 194](#)
- ◆ [Section B.2.2, “Trigger Operation,” on page 197](#)
- ◆ [Section B.2.3, “Window Operation,” on page 198](#)
- ◆ [Section B.2.4, “Gate Operation,” on page 200](#)
- ◆ [Section B.2.5, “Sequence Operation,” on page 200](#)
- ◆ [Section B.2.6, “Distinct Operation,” on page 201](#)

B.2.1 Filter Operation

A filter consists of evaluation expressions that evaluate the current event from the real-time event stream. It compares the event field values with user-specified values by using a wide set of comparison and match operators. If a match is found, the output set is the matched event, and then the filter resets.

The syntax for the filter operation is:

```
<filter expression> ::= "filter("<evaluation expression 1> [NOT|AND|OR <evaluation expression 2>] [...] [NOT|AND|OR <evaluation expression n>]")"
```

<evaluation expression 1..N> is an expression constructed of one event field reference (name or ID) and a comparison or match operator.

```
<evaluation expression> ::= "e." <event_name | event_ID> <comparison operator | match operator> <user_specified_value>
```

Multiple evaluation expressions can be combined by using the standard Boolean operators (AND, OR, and NOT) and grouping parentheses.

For example, the following rule detects whether the current event has a severity of 4 and the event name contains either “FW” or “Comm.”

```
filter(e.sev = 4 and (e.evt match regex ("FW") or e.evt match regex ("Comm")))
```

The filter operation supports the following operators:

- ♦ [“Boolean Operators” on page 195](#)
- ♦ [“Standard Arithmetic Operators” on page 195](#)
- ♦ [“Match Regex Operator” on page 195](#)
- ♦ [“Match Subnet Operator” on page 196](#)
- ♦ [“Inlist and Not Inlist Operators” on page 196](#)
- ♦ [“IsNull Operator” on page 196](#)

Boolean Operators

Filter expressions can be combined by using the Boolean operators AND, OR, and NOT. The filter Boolean operator precedence (from the highest to the lowest precedence) is described in the following table:

Table B-1 Boolean Operators

Operator	Meaning	Operator Type	Associativity
NOT	logical NOT	unary	None
AND	logical AND	binary	left to right
OR	logical OR	binary	left to right

Standard Arithmetic Operators

Standard arithmetic operators can be used to build a condition that compares the value of an event ID and a user-specified value, which is either a numeric value or a string field. The standard arithmetic operators in Sentinel are =, <, >, !=, <=, and >=.

Examples:

```
filter(e.sev > 3)
filter(e.BeginTime < 1179217665)
filter(e.iufname != "Administrator")
```

Match Regex Operator

The match regex operator is used to match an event field value by using standard regular expressions. This operator is used only for string type tags.

Examples:

Match the exact phrase (case sensitive):

```
filter(e.evt match regex ("LoginUser"))
```

Match the exact phrase (not case sensitive):

```
filter(e.evt match regex ("(?i)^TeSt EvenT$"))
```

The value includes the phrase (not case sensitive):

```
filter(e.evt match regex ("(?i)EsT EVen"))
```

For more information on the match regex operator, see [Regular Expressions](#).

Match Subnet Operator

The match subnet operator is used to match event field IP addresses to a subnet in standard CIDR notation. This operator is used only for IP address fields.

Example:

```
filter(e.dip match subnet (10.0.0.1/22))
filter(e.sip = 10.0.0.1 or e.sip=10.0.0.2) and (e.dpint=80)
```

For more information on CIDR notation, see [CIDR](#).

Inlist and Not Inlist Operators

The inlist operator is used to perform a lookup on an existing Dynamic List of string values. This operator returns TRUE if the event field value is present in the list. The Inlist and Not Inlist operators are case sensitive. For more information on Dynamic Lists, see [Chapter 9, “Configuring Dynamic Lists,” on page 109](#).

For example, the following expression is used to evaluate whether the source IP address of the current event is present in a Dynamic List named MailServerList. If the source IP address is present in this list, the expression evaluates to TRUE.

```
filter(e.sip inlist MailServerList)
```

As another example, this filter expression combines the NOT operator and the inlist operator. This expression evaluates to TRUE if the source IP address is not present in the Dynamic List named MailServerList.

```
filter(not (e.sip inlist MailServerList))
```

The following expression is used to evaluate whether the event name of the current event equals “File Access” and the InitiatorUserName is not present in a Dynamic List named AuthorizedUsers. If both conditions are true for the current event, the expression evaluates to TRUE.

```
filter(e.evt="File Access" and not(e.sun inlist AuthorizedUsers))
```

IsNull Operator

The isnull operator returns TRUE if the event field value is equal to null. For example:

```
Filter(isnull(e.sip))
```

Output Sets

- ◆ The output of this expression is either the empty set (if the evaluation expression evaluates to FALSE) or a set containing the current event (if the evaluation expression evaluates to TRUE).
- ◆ If the filter is the last or only operation of a correlation rule, the output set of the filter is used to construct a correlated event. The input event that matched the filter is associated with the correlated event.
- ◆ If the filter is not the last operation of a correlation rule (that is, if the filter is followed by a flow operator), the output set of a filter is used as the input set to other operations through the flow operator.

Additional Information

The filter operator can be used to compare event field values with other event field values within the same event. For example:

```
filter(e.sip=e.dip)
```

B.2.2 Trigger Operation

The trigger operation counts a number of events for a specified duration. The trigger command defines a threshold count of events within a time condition and applies an optional discriminator that splits events into unique buckets on which to apply the threshold count and time window conditions. Trigger itself does not apply any filtering. Every event that enters the trigger function is put in a bucket and is counted. If you want to use trigger on certain type of events, you must prefilter the events and then flow that output to trigger.

Syntax

```
<trigger expression> ::= "trigger("<threshold count>","<evaluation period>","discriminator("<list of event fields>")")"
```

<threshold count> is an integer value specifying the number of events in a single bucket that are necessary for the rule to fire.

<evaluation period> indicates the duration for which old events are kept in the bucket and counted towards the threshold. The duration can be seconds (s), minutes (m), hours (h), and days (d.).

discriminator (<list of event fields>) specifies the set of fields to use to segregate each event into its unique bucket. All fields are logically combined to define the bucket.

For example, the following rule detects if 5 events with the same source IP address happen within 10 seconds.

```
trigger(5,10,discriminator(e.sip))
```

Note that the discriminator is used to split the input stream of events into distinct buckets where all events in a bucket have identical data for the specified fields (SourceIP in the example above.) Multiple fields can be used to create buckets where all events in a single bucket have, for example, the same SourceIP and TargetUserName. However, you cannot count across multiple buckets. For example, you cannot use the Trigger operation to detect conditions where a source system contacts N distinct target systems.

Output Sets

- ♦ If the specified count is reached within the specified duration, a set of events containing all of the events maintained by the trigger is output.
- ♦ When receiving a new input set of events, a trigger first discards the outdated events (events that have been maintained for more than the duration) and then inserts the current event. If the number of resulting events is greater than or equal to the specified count, the trigger outputs a set containing all of the events.
- ♦ If a trigger is the last operation (or the only operation) of a correlation rule, then the output set of the trigger is used to construct a correlated event. The raw events associated with the correlated events are the trigger operation output set of events with the current event listed first.

B.2.3 Window Operation

The Window operation compares the current event to a set of past events that are stored in a window that is defined as part of the Window operation. Window allows you to temporarily store events of interest and then use the data within those stored events to filter the incoming raw events.

For example, suppose you want to detect if someone attempts to log in to an account several times and fails, but then guesses that password and succeeds in logging in. You would collect the failed login events and store them in the window, then compare new successful login events against those stored events and match the username. If a match is found, then the user first failed to log in and then did so successfully. Any event field can be matched against any other event field for more sophisticated correlations.

How the Window Operation Works

The incoming raw stream of events is split before it arrives at the window operation:

- ♦ **Storage filter (w.event tag):** One stream of raw events is matched against the storage filter and, if any events match, they are placed in the window. Note that the current event is not part of this stream (that is, this stream is delayed by one event.) Events in the window are referred to with the "w." prefix, which can be interpreted as "the set of events in the window."
- ♦ **Evaluation filter (e.event tag):** The other stream of raw events is filtered and then passed, one at a time, into the Window operation. Each event (the current event) is matched against the set of events stored in the window by using the defined comparison expression. The current event passing into the Window operation is referred to with the "e." prefix, which can be interpreted as "the current event."

You can define two filters for the Window operation to constrain resource usage. You define the first filter within the Window operation itself to set the storage filter. You define the second filter as a prefix to the Window operation using a standard in-line filter that sets the evaluation filter. These filters are optional, but recommended.

To ensure that the current event does not match itself in cases where the comparison operator is matching on the same field and a single event would match both the evaluation and storage filters, the set of events in the window does not include the current event.

NOTE: It is critical to design your storage and evaluation filters carefully to minimize resource usage by the Window operation. The window stores the event UUID and any fields necessary for comparison for all events that match the storage filter, so constraining the set of matched events is important to reduce memory use.

Syntax

The syntax is as follows:

```
<window expression> ::= "window(" <comparison expression> "," <storage filter> ","  
<storage_time_period> ")"
```

<comparison expression> matches a current event field against a field in the set of past events stored in the window. The standard comparison operators are supported.

```
<comparison expression> ::= "w." <event_ID | event_name> <comparison operator> "e."  
<event_ID | event_name>
```

Multiple comparison expressions can be combined by using the standard Boolean AND, OR, and NOT operators. Note that it is also possible to compare stored event fields against literals, in which case the Window operation ignores the contents of the current event and always copies it to the output as long as an event that matches the literal is in the window.

`<storage filter>` is an embedded filter expression that defines the storage filter. All events that match this filter are stored in the window for later comparison until they expire.

```
<storage filter> ::= <filter_rule>
```

`<filter_rule>` is a simple expression using the filter operator. For more information, see [Section B.2.1, "Filter Operation," on page 194](#).

`<storage_time_period>` is the total time for which a single event is stored in the window. The storage time period can be seconds (s), minutes (m), hours (h), and days (d.)

```
<storage_time_period> ::= <1 or more digits> "s" | "m" | "h" | "d"
```

Examples

The following rule detects whether the current event has a source IP address that matches the source IP address of an event that happened within the past 60 seconds, with the past events limited to those whose source IP address is within the specified subnet. This Window would typically be preceded by an evaluation filter to restrict the set of events evaluated by the Window.

```
window(w.sip = e.sip, filter(e.sip match subnet (10.0.0.10/22), 60)
```

As another example, the following rule is a domino type of rule. An attacker exploits a vulnerable system and uses it as an attack platform. This Window would typically be preceded by an evaluation filter to restrict the set of events evaluated by the Window.

```
filter(e.XDASTaxonomyName = "XDAS_AE_IDS_PROBE" OR e.XDASTaxonomyName =  
"XDAS_AE_IDS_PENETRATE") flow window((e.sip = w.dip AND e.dp = w.dp AND e.evt =  
w.evt), filter(e.XDASTaxonomyName = "XDAS_AE_IDS_PROBE" OR e.XDASTaxonomyName =  
"XDAS_AE_IDS_PENETRATE"), 1h)
```

The following rule identifies a potential security breach after a denial of service attack. The rule fires if the destination of a denial of service attack has a service stopped within 60 seconds of the attack.

```
filter(e.rv51="Service" and e.rv52="Stop") flow window (e.sip = w.dip,  
filter(e.XDASTaxonomyName = "XDAS_AE_DOS"), 60)
```

Output Sets

- ◆ If the Window operation matches a current event against the window based on the comparison expression, the output set is the incoming event plus all matching past events
- ◆ If no events in the window match the current event, the output set is empty.
- ◆ If a window is the last or only operation of a correlation rule, the output set of the window is used to construct a correlated event. The raw events associated with the correlated events are the window operation output set of events with the current event first.

Additional Information

- ◆ All window simple evaluation expressions must include an event ID in the form w.[event_ID].
- ◆ Every event coming in to the Correlation Engine that passes the storage filter is put into the window of past events except for the most recent, or current, event.

- ♦ If no storage filter expression is defined, all events coming into the Correlation Engine are stored by the window. With extremely high event rates or long duration storage time period, this might require a large amount of memory.
- ♦ To minimize memory usage, only the relevant parts of the past events, not all event ID values, are maintained in memory.

B.2.4 Gate Operation

The Gate operation is used to combine multiple subrules together to detect conditions where several distinct activities happen within a given time period. The order in which each activity occurred is not considered.

The gate operation is made up of one or more nested subrules and can be configured to fire if some, any, or all of the subrules fire within a specified time. The subrules can be a simple rule or another composite rule. For more information on composite rules, see [“Composite Rule” on page 60](#).

Syntax

```
<gate expression> ::= "gate("<subrule 1>","<subrule 2>","<subrule
n>","<mode>","<evaluation period>","discriminator ("<list of event fields>"))"
```

<subrule 1...N> rules are the rule definitions for 1 to n subrules. Each subrule is an independent, valid, correlation rule.

<mode> can be one of the subrules that must be triggered for the Gate operation to trigger.

<evaluation period> indicates the time period over which the specified number of subrules must fire in order for the whole Gate rule to fire.

discriminator (*<list of fields>*) acts similar to the discriminator in the Trigger operation. The output from each subrule is placed into unique buckets based on the data in the fields specified in this discriminator. The *<mode>* count and *<evaluation period>* is then applied to each bucket separately.

For example, the following rule is a typical perimeter security IDS inside/outside rule:

```
filter(e.sev > 3) flow gate(filter(e.sn = "in"), filter(e.sn = "out"), all, 60s,
discriminator(e.dip, e.evt))
```

B.2.5 Sequence Operation

Sequence rules are similar to gate rules, except that all subrules must fire in sequence for the overall Sequence operation to fire.

The subrules can be a simple rule or another composite rule.

Syntax

Syntax:

```
<sequence expression> ::= "sequence("<subrule 1>","<subrule 2>","<subrule
n>","<evaluation period>","discriminator("<list of event fields>"))"
```

<subrule 1...N> are the rule definitions for 1 to n subrules. Each subrule is an independent, valid, correlation rule.

<evaluation period> is a time period expressed in seconds (s), minutes (m), or hours (h).

discriminator (<list of fields>) acts similar to the discriminator in the Trigger operation; the output from each subrule is placed into unique buckets based on the data in the fields specified in this discriminator. The <evaluation period> is then applied to each bucket separately.

For example, this rule detects three failed logins by a particular user in 10 minutes followed by a successful login by the same user.

```
sequence (filter(e.evt="failed logins") flow trigger(3, 600, discriminator(e.sun, e.dip)), filter(e.evt="goodlogin"), 600, discriminator(e.sun, e.dip))
```

B.2.6 Distinct Operation

The Distinct operation considers only the unique values in events. You can use this operation in scenarios where you need to differentiate scans versus flood type attacks. In flood type attacks the overall quantity is important, with scans the overall unique information disseminated is important.

The following are some examples where you can use the Distinct operation:

You want a rule to trigger if five events have unique source IP addresses but the same destination IP address within a 10 second period. You can create a free-form rule as follows:

```
trigger(5,10, discriminator(e.dip, distinct e.sip))
```

You want a rule to trigger if three severity 5 events occur from distinct combination of initiator IP addresses and initiator users within a 60 second period. You can create a free-form rule as follows:

```
filter(e.sev = 5) flow trigger(3, 60, discriminator(distinct e.sip, distinct e.sun))
```

You want a rule to trigger if three severity 5 events occur from the same initiator IP address but distinct initiator users within a 60 second period. You can create a free-form rule as follows:

```
filter(e.sev = 5) flow trigger(3, 60, discriminator(e.sip, distinct e.sun))
```

You want a rule to trigger if three severity 5 events occur from the same initiator user but distinct initiator IP addresses within a 60 second period. You can create a free-form rule as follows:

```
filter(e.sev = 5) flow trigger(3, 60, discriminator(distinct e.sip, e.sun))
```

B.3 Operators

Operators are used to transition between operations or expressions. The following fundamental operators are used between operations:

- ♦ [Section B.3.1, "Flow Operator," on page 201](#)
- ♦ [Section B.3.2, "Union Operator," on page 202](#)
- ♦ [Section B.3.3, "Intersection Operator," on page 202](#)

B.3.1 Flow Operator

The output set of events of the left side operation is the input set of events for the right side operation. Flow is typically used to transition from one correlation operation to the next.

For example:

```
filter(e.sev = 5) flow trigger(3, 60)
```

The output of the filter operation is the input of the trigger operation. The trigger only counts 3 events with severity equal to 5.

B.3.2 Union Operator

The union operator is the union of the left side operation output set and the right side operation output set. The resulting output set contains events from either the left side operation output set or the right side operation output set, without duplicates.

For example:

```
filter(e.sev = 5) union filter(e.sip = 10.0.0.1)
```

is equivalent to

```
filter(e.sev = 5 or e.sip = 10.0.0.1)
```

B.3.3 Intersection Operator

The intersection operator is the intersection of the left side operation output set and the right side operation output set. The resulting output set contains events that are common to both the left side operation output set and the right side operation output set without duplicates.

For example:

```
filter(e.sev = 5) intersection filter(e.sip = 10.0.0.1)
```

is equivalent to

```
filter(e.sev = 5 and e.sip = 10.0.0.1)
```

B.4 Order of Operators

The operator precedence (from the highest to the lowest) is as follows:

Table B-2 Operator Precedence

Operator	Meaning	Operator Type	Associativity
flow	The output set becomes the input set	binary	left to right
intersection	Set intersection (remove duplicates)	binary	left to right
union	Set union (remove duplicates)	binary	left to right