



Sentinel™

Installation and Configuration Guide

March 2020

Legal Notice

© Copyright 2001-2020 Micro Focus or one of its affiliates.

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

For additional information, such as certification-related notices and trademarks, see <http://www.microfocus.com/about/legal/>.

Contents

About this Book and the Library	11
Part I Understanding Sentinel	13
1 What is Sentinel?	15
Challenges of Securing an IT Environment	15
The Solution That Sentinel Provides	16
2 How Sentinel Works	19
Event Sources	21
Sentinel Event	21
Mapping Service	22
Streaming Maps	23
Collector Manager	23
Collectors	23
Connectors	23
ArcSight SmartConnectors	24
Agent Manager	24
Sentinel Data Routing and Data Storage	24
Event Visualizations	25
Correlation	25
Security Intelligence	25
Incident Remediation	25
iTrac Workflows	26
Actions and Integrators	26
Searching	26
Reports	26
Identity Tracking	27
Event Analysis	27
Part II Planning Your Sentinel Installation	29
3 Implementation Checklist	31
4 Understanding License Information	33
Sentinel Licenses	34
Evaluation License	34
Free License	35
Enterprise Licenses	35

5 Meeting System Requirements	37
Connector and Collector System Requirements	37
Virtual Environment	37
6 Deployment Considerations	39
Data Storage Considerations	39
Planning for Traditional Storage	40
Planning for Scalable Storage	43
Sentinel Directory Structure	46
Advantages of Distributed Deployments	46
Advantages of Additional Collector Managers	47
Advantages of Additional Correlation Engines	47
All-In-One Deployment	48
One-Tier Distributed Deployment	48
One-Tier Distributed Deployment with High Availability	49
Two-Tier and Three-Tier Distributed Deployment	50
Three-Tier Deployment with Scalable Storage	51
7 Deployment Considerations for FIPS 140-2 Mode	55
FIPS Implementation in Sentinel	55
RHEL NSS Packages	55
SLES NSS Packages	56
FIPS-Enabled Components in Sentinel	56
Data Connections Affected by FIPS Mode	57
Implementation Checklist	57
Deployment Scenarios	58
Scenario 1: Data Collection in Full FIPS 140-2 Mode	58
Scenario 2: Data Collection in Partial FIPS 140-2 Mode	59
8 Ports Used	61
Sentinel Server Ports	61
Local Ports	61
Network Ports	61
Sentinel Server Appliance Specific Ports	63
Collector Manager Ports	63
Network Ports	63
Collector Manager Appliance Specific Ports	64
Correlation Engine Ports	65
Network Ports	65
Correlation Engine Appliance Specific Ports	65
Scalable Storage Ports	65
9 Installation Options	67
Traditional Installation	67
Appliance Installation	68

Part III Installing Sentinel	69
10 Installation Overview	71
11 Installation Checklist	73
12 Installing and Configuring Elasticsearch	75
Prerequisites	75
Installing and Configuring Elasticsearch	75
Securing Data in Elasticsearch	78
Installing the Elasticsearch Security Plug-In	79
Providing Secure Access to Additional Elasticsearch Clients	80
Updating the Elasticsearch Plug-In Configuration	81
Performance Tuning for Elasticsearch	81
Redeploying Elasticsearch Security Plug-In	82
13 Installing and Setting Up Scalable Storage	85
Installing and Configuring CDH	86
Prerequisites	86
Installing and Configuring CDH	87
Enabling Scalable Storage	87
14 Traditional Installation	89
Performing Interactive Installation	89
Sentinel Server Standard Installation	89
Sentinel Server Custom Installation	90
Collector Manager and Correlation Engine Installation	93
Performing a Silent Installation	96
Installing Sentinel as a Non-root User	97
15 Appliance Installation	101
Prerequisites	101
Installing the Sentinel ISO Appliance	102
Installing Sentinel	102
Installing Collector Managers and Correlation Engines	103
Installing the Sentinel OVF Appliance	104
Installing Sentinel	104
Installing Collector Managers and Correlation Engines	105
Post-Installation Configuration for the Appliance	106
Registering for Updates	106
Creating Partitions for Traditional Storage	107
Configuring Scalable Storage	108
Configuring the Appliance with SMT	108
16 Installing Additional Collectors and Connectors	111
Installing a Collector	111
Installing a Connector	111

17 Verifying the Installation	113
Part IV Configuring Sentinel	115
18 Configuring Time	117
Understanding Time in Sentinel	117
Configuring Time in Sentinel	119
Configuring Delay Time Limit for Events	119
Handling Time Zones	119
19 Securing Data in Elasticsearch	121
20 Enabling Event Visualization	123
Prerequisite	123
Enabling Event Visualization.	123
21 Modifying the Configuration after Installation	125
22 Configuring Out-of-the-Box Plug-Ins	127
Viewing the Preinstalled Plug-Ins	127
Configuring Data Collection	127
Configuring Solution Packs.	127
Configuring Actions and Integrators	128
23 Enabling FIPS 140-2 Mode in an Existing Sentinel Installation	129
Enabling Sentinel Server to Run in FIPS 140-2 Mode	129
Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines	130
24 Operating Sentinel in FIPS 140-2 Mode	131
Configuring Distributed Search in FIPS 140-2 Mode	131
Configuring LDAP Authentication in FIPS 140-2 Mode	132
Updating Server Certificates in Remote Collector Managers and Correlation Engines	133
Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode	133
Agent Manager Connector	134
Database (JDBC) Connector	135
Sentinel Link Connector	135
Syslog Connector	136
Windows Event (WMI) Connector	137
Sentinel Link Integrator	138
LDAP Integrator	138
SMTP Integrator	139
Syslog Integrator	139
Using Non-FIPS Enabled Connectors with Sentinel in FIPS 140-2 Mode	140
Importing Certificates into FIPS Keystore Database	140
Reverting Sentinel to Non-FIPS Mode	141
Reverting Sentinel Server to Non-FIPS mode	141

Reverting Remote Collector Managers or Remote Correlation Engines to Non-FIPS mode	141
25 Adding a Consent Banner	143
26 Limiting the Number of Concurrent Active Sessions	145
27 Ending Inactive Sessions	147
28 Configuring IP Flow Data Collection	149
Part V Upgrading Sentinel	151
29 Implementation Checklist	153
30 Prerequisites	155
Saving the Custom Configuration Information	155
Saving the server.conf File Settings	155
Saving the jetty-ssl File Settings	155
Extending the Retention Period for Event Associations Data	155
Pre-Upgrade Configuration for SSDM	156
Change Guardian Integration	156
31 Upgrading Sentinel Traditional Installation	157
Upgrading Sentinel	157
Upgrading Sentinel as a Non-root User	159
Upgrading the Collector Manager or the Correlation Engine	160
Upgrading the Operating System	161
32 Upgrading the Sentinel Appliance	163
Prerequisites for Upgrading the Appliance	163
Upgrading the Operating System to SLES 12 SP3	163
Migrating Data from MongoDB to PostgreSQL	165
Upgrading the Appliance	166
Upgrading through the Appliance Update Channel	166
Upgrading through SMT	168
Upgrading through Sentinel Appliance Manager	169
Applying Operating System Patches	170
33 Troubleshooting	171
Cleaning Up Data From PostgreSQL When Migration Fails	171
Unable to Run the Migration Script	172
Cannot Connect to Servers or Other Components through Appliance	172
Error When Upgrading the Appliance	173

34 Post-Upgrade Configurations	175
Removing Data from MongoDB	175
Synchronizing the postgresql.conf file	175
Securing Data in Elasticsearch	176
Configuring Event Visualizations	176
Configuring IP Flow Data Collection	176
Configuring SmartConnectors that collect IP Flow data	177
Uninstalling Existing NetFlow Collector Managers	177
Post-Upgrade Configuration for Sentinel Scalable Data Manager	177
Install Elasticsearch Security Plug-In	178
Updating Spark Applications on YARN	178
Enabling Sentinel Features	179
Updating Dashboards and Visualizations in Sentinel Scalable Data Manager	179
Adding the JDBC DB2 Driver	180
Configuring Data Federation Properties in Sentinel Appliance	180
Registering Sentinel Appliance for Updates	181
Updating External Databases for Data Synchronization	181
Updating Permissions for Users Who Send Data from Other Integrated Products to Sentinel	181
35 Upgrading Sentinel Plug-Ins	183
Part VI Migrating Data from Traditional Storage	185
36 Migrating Data to Scalable Storage	187
Data You Can Migrate	188
Migrating Configuration Data	189
Backing Up Data on the Source Server	189
Restoring Data on the Target Server	190
Migrating Event Data and Raw Data	190
Migrating Alerts	190
Updating Sentinel Clients	191
Importing ESM Configuration	191
37 Migrating Data to Elasticsearch	193
38 Migrating Data	195
Part VII Deploying Sentinel for High Availability	197
39 Concepts	199
External Systems	199
Shared Storage	199
Service Monitoring	200
Fencing	200

40 System Requirements	201
41 Installation and Configuration	203
Initial Setup	204
Shared Storage Setup	205
Configuring iSCSI Targets	206
Configuring iSCSI Initiators	207
Sentinel Installation	209
First Node Installation	209
Subsequent Node Installation	211
Cluster Installation	212
Cluster Configuration	213
Resource Configuration	216
Secondary Storage Configuration	217
42 Configuring Sentinel HA as SSDM	219
43 Upgrading Sentinel in High Availability	221
Prerequisites	221
Upgrading a Traditional Sentinel HA Installation	221
Upgrading Sentinel HA	222
Upgrading the Operating System	224
Upgrading a Sentinel HA Appliance Installation	227
Upgrading through the Zypper patch	228
Upgrading through the Sentinel Appliance Management Console	229
44 Backup and Recovery	233
Backup	233
Recovery	233
Transient Failure	233
Node Corruption	234
Cluster Data Configuration	234
Part VIII Appendices	235
A Troubleshooting	237
Failed Installation Because of an Incorrect Network Configuration	237
The UUID Is Not Created for Imaged Collector Managers or Correlation Engine	238
Sentinel Main Interface is Blank in Internet Explorer After Logging in	238
Sentinel Does Not Launch in Internet Explorer 11 in Windows Server 2012 R2	238
Sentinel Cannot Run Local Reports with Default EPS License	239
Synchronization Needs to be Started Manually in Sentinel High Availability After You Convert the Active Node to FIPS 140-2 Mode	239
Sentinel Main Interface Displays Blank Page After Converting to Sentinel Scalable Data Manager	239
The Event fields Panel is Missing in the Schedule Page When Editing Some Saved Searches	240

Sentinel Does Not Return Any Correlated Events When You Search for Events for the Deployed Rule	
with the Default Fire Count Search	240
Security Intelligence Dashboard Displays Invalid Baseline Duration When Regenerating a Baseline	240
Sentinel Server Shuts Down When Running a Search If There Are Large Number of Events in a Single Partition	240
Error While Using the report_dev_setup.sh Script to Configure Sentinel Ports for Firewall Exceptions on Upgraded Sentinel Appliance Installations.	241

B Uninstalling 243

Uninstallation Checklist	243
Uninstalling Sentinel	243
Uninstalling the Sentinel Server	243
Uninstalling the Collector Manager and Correlation Engine	244
Post-Uninstallation Tasks	244

About this Book and the Library

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

Intended Audience

This guide is intended for Sentinel administrators and consultants.

Other Information in the Library

The library provides the following information resources:

Administration Guide

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

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.

Understanding Sentinel

This section provides detailed information about Sentinel and how it provides an event management solution for your organization.

- ◆ [Chapter 1, “What is Sentinel?”](#) on page 15
- ◆ [Chapter 2, “How Sentinel Works,”](#) on page 19

1 What is Sentinel?

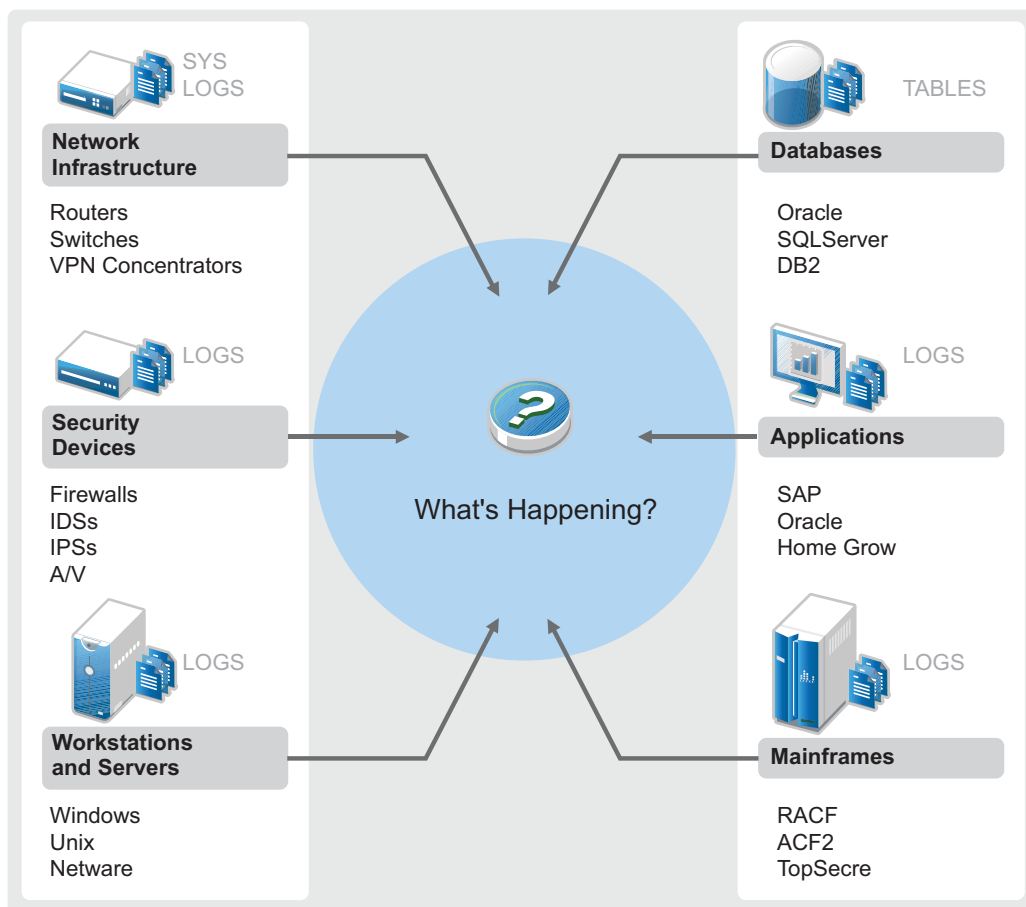
Sentinel is a security information and event management (SIEM) solution and also a compliance monitoring solution. Sentinel automatically monitors the most complex IT environments and provides the security required to protect your IT environment.

- ♦ [“Challenges of Securing an IT Environment”](#) on page 15
- ♦ [“The Solution That Sentinel Provides”](#) on page 16

Challenges of Securing an IT Environment

Securing your IT environment is a challenge because of the complexity of the environment. Typically, there are many applications, databases, mainframes, workstations, and servers in your IT environment, and all these entities generate logs of events. You might also have security devices and network infrastructure devices that generate logs of events in your IT environment.

Figure 1-1 What Happens in Your Environment



Challenges arise because of the following facts:

- ◆ There are many devices in your IT environment.
- ◆ The logs are in different formats.
- ◆ The logs are stored in different locations.
- ◆ The volume of information captured in the log files is large.
- ◆ It is impossible to determine event triggers without manually analyzing the log files.

To make the information in the logs useful, you must be able to perform the following:

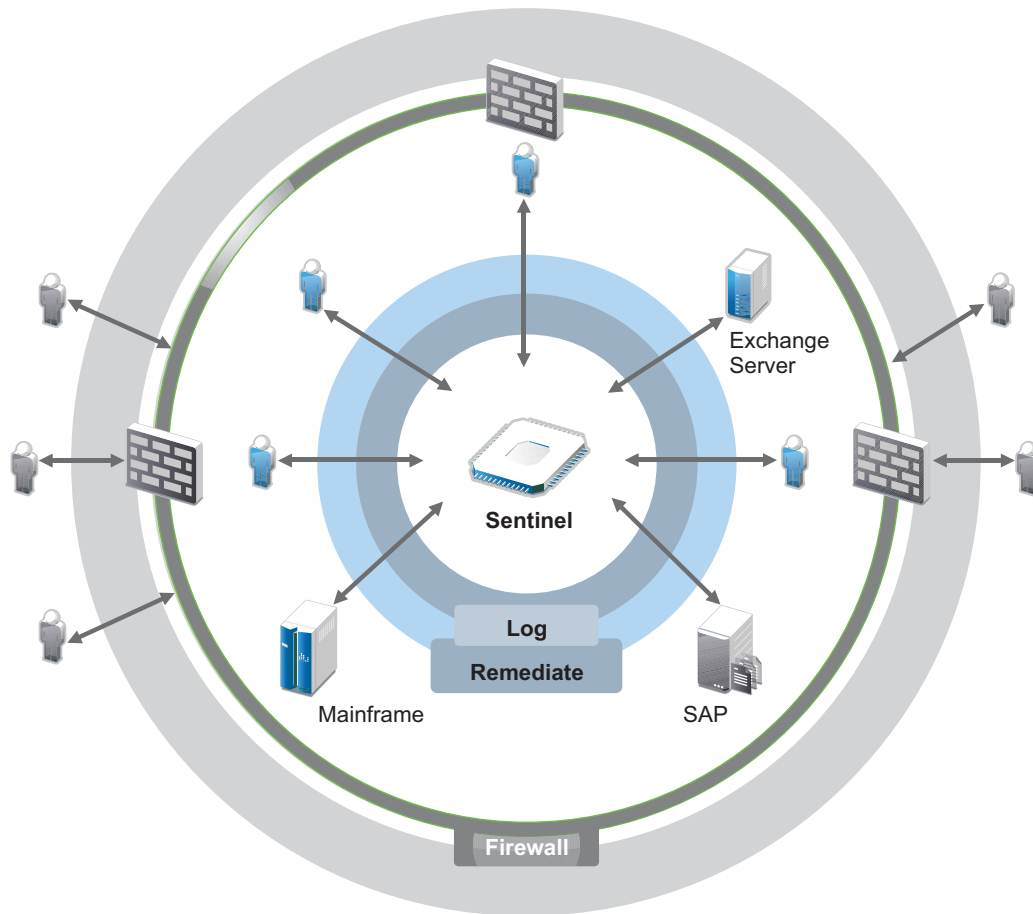
- ◆ Collect the data.
- ◆ Consolidate the data.
- ◆ Normalize disparate data into events that you can easily compare.
- ◆ Map events to standard regulations.
- ◆ Analyze the data.
- ◆ Compare events across multiple systems to determine if there are security issues.
- ◆ Send notifications when the data does not comply with the norms.
- ◆ Take action on notifications to comply with business policies.
- ◆ Generate reports to prove compliance.

After you understand the challenges of securing your IT environment, you need to determine how to secure the enterprise for and from the users without impacting the user experience. Sentinel provides the solution.

The Solution That Sentinel Provides

Sentinel acts as the central nervous system to the enterprise security. It gathers data from across your entire infrastructure—applications, databases, servers, storage, and security devices. It analyzes and correlates the data, and makes the data actionable, either automatically or manually.

Figure 1-2 The Solution That Sentinel Provides



With Sentinel, you know what is happening in your IT environment at any given point, and you have the ability to connect the actions taken on resources to the people taking those actions. This allows you to determine user behavior and effectively monitor activities to prevent malicious activities.

Sentinel achieves this by:

- ◆ Providing a single solution to address IT controls across multiple security standards.
- ◆ Address the gap between what should happen and what is actually happening in your IT environment.
- ◆ Helping you to be compliant to security standards.
- ◆ Providing out-of-the-box compliance monitoring and reporting programs.

Sentinel automates log collection, analysis, and reporting processes to ensure that IT controls are effective in supporting threat detection and audit requirements. Sentinel provides automated monitoring of security events, compliance events, and IT controls. It allows you to take immediate action if there is a security breach or non-compliant event occurring. Sentinel also allows you to gather summary information about your environment, which you can share with your key stakeholders.

2 How Sentinel Works

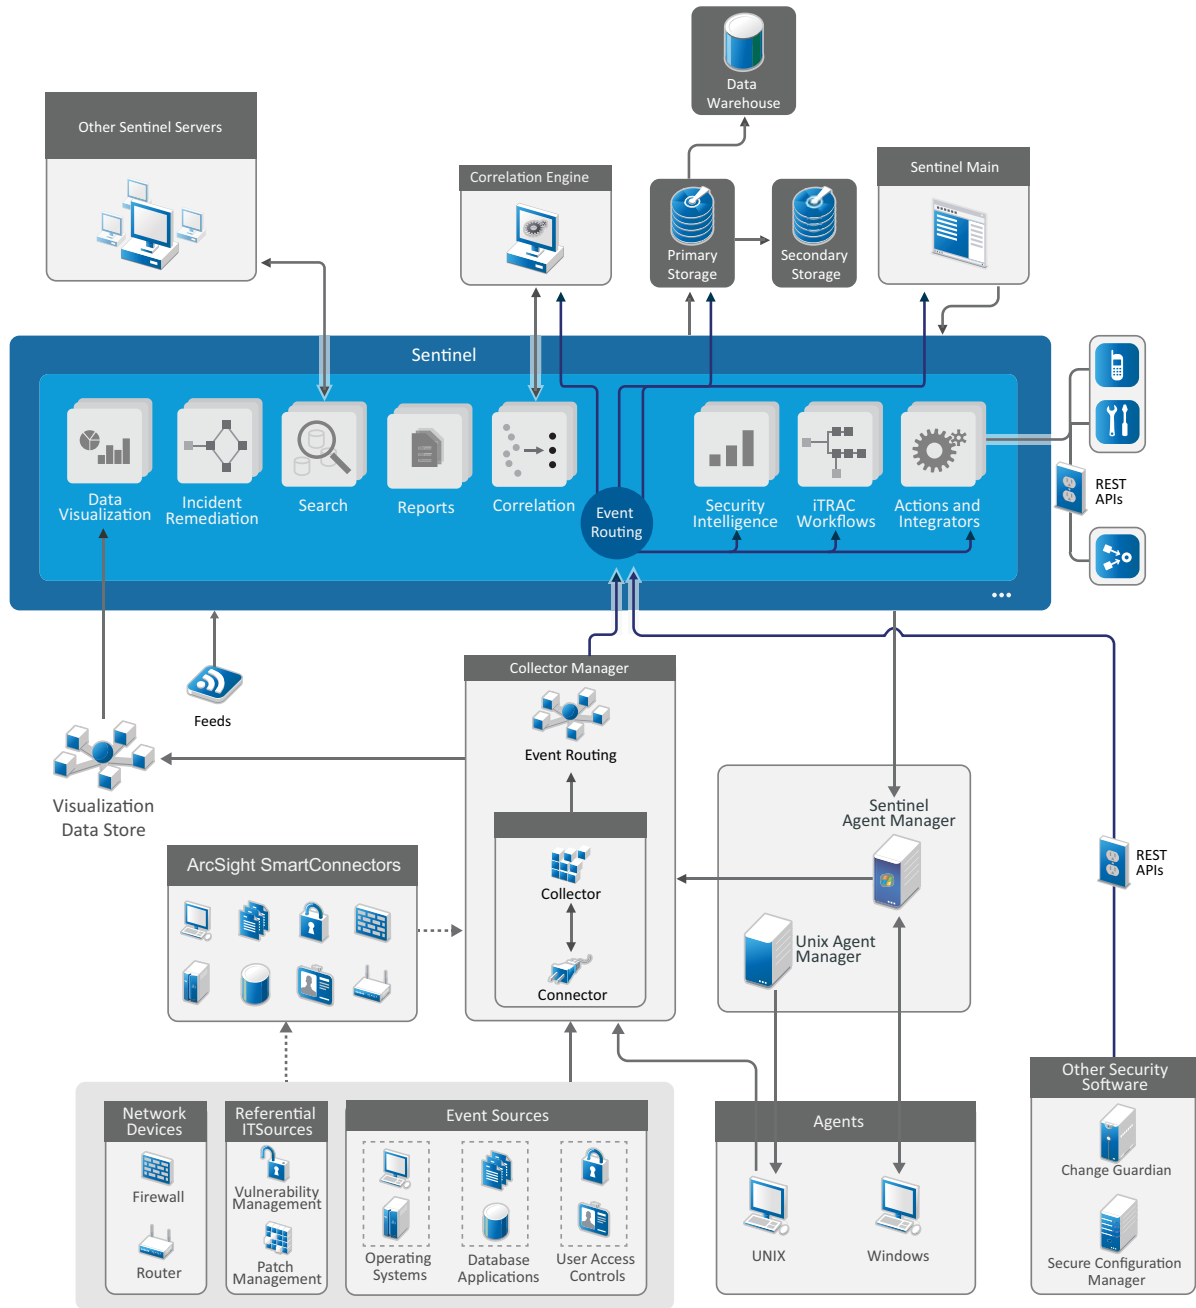
Sentinel continuously manages security information and events across your IT environment to provide a complete monitoring solution.

Sentinel does the following:

- ◆ Gathers logs, events, and security information from the various sources in your IT environment.
- ◆ Normalizes the collected logs, events, and security information into a standard Sentinel format.
- ◆ Stores events in a file-based data storage or Hadoop-based scalable storage with flexible, customizable data retention policies.
- ◆ Collects IP Flow data and helps you monitor network activities in detail.
- ◆ Provides the ability to hierarchically link multiple Sentinel systems, including Sentinel Log Manager.
- ◆ Allows you to search for events on your local Sentinel server, and also on other Sentinel servers distributed across the globe.
- ◆ Performs a statistical analysis that allows you to define a baseline and then compares it to what is occurring, to determine if there are unseen problems.
- ◆ Correlates a set of similar or comparable events in a specific duration to determine a pattern.
- ◆ Organizes events into incidents for efficient response management and tracking.
- ◆ Provides reports based on real time and historical events.

The following figure illustrates how Sentinel works with traditional storage as the data storage option:

Figure 2-1 Sentinel Architecture



The following sections describe Sentinel components in detail:

- ◆ “Event Sources” on page 21
- ◆ “Sentinel Event” on page 21

- ◆ “Collector Manager” on page 23
- ◆ “ArcSight SmartConnectors” on page 24
- ◆ “Agent Manager” on page 24
- ◆ “Sentinel Data Routing and Data Storage” on page 24
- ◆ “Event Visualizations” on page 25
- ◆ “Correlation” on page 25
- ◆ “Security Intelligence” on page 25
- ◆ “Incident Remediation” on page 25
- ◆ “iTrac Workflows” on page 26
- ◆ “Actions and Integrators” on page 26
- ◆ “Searching” on page 26
- ◆ “Reports” on page 26
- ◆ “Identity Tracking” on page 27
- ◆ “Event Analysis” on page 27

Event Sources

Sentinel gathers security information and events from various sources in your IT environment. These sources are called event sources. Typically, the following are the event sources on your network:

Security Perimeter: Security devices including hardware and software used to create a security perimeter for your environment, such as firewalls, intrusion detective systems (IDS), and virtual private networks (VPN).

Operating Systems: Various operating systems running in the network.

Referential IT Sources: The software used to maintain and track assets, patches, configuration, and vulnerability.

Applications: Various applications installed in the network.

User Access Control: Applications or devices that allow users access to company resources.

For more information about collecting events from event sources, see “[Collecting and Routing Event Data](#)” in the *Sentinel Administration Guide*.

Sentinel Event

Sentinel receives information from devices, normalizes this information into a structure called an event, categorizes the event, and then sends the event for processing.

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 adds category information (taxonomy) to events, to make it easier to compare events across systems that report events differently. Events are processed by the real time display, correlation engine, dashboards, and the back end server.

An event comprises more than 200 fields; event fields are of different types and of different purposes. There are some predefined fields such as severity, criticality, destination IP address, and destination port.

There are two sets of configurable fields:

- ◆ Reserved fields: For Sentinel internal use to allow extension of functionality in future.
- ◆ Customer fields: For customer use to allow customization.

The source for a field can either be external or referential:

- ◆ The value of an external field is set explicitly by the device or the corresponding Collector. For example, a field can be defined to be the building code for the building containing the asset mentioned as the destination IP address of an event.
- ◆ The value of a referential field is computed as a function of one or more other fields using the mapping service. For example, a field can be computed by the mapping service using a customer defined map, using the destination IP address from the event.
- ◆ [“Mapping Service” on page 22](#)
- ◆ [“Streaming Maps” on page 23](#)

Mapping Service

The Mapping Service propagates business relevance data throughout the system. This data can enrich events with referential information.

You can enrich your event data by using maps to add additional information, such as host and identity information, to the incoming events from your source devices. Sentinel can use this additional information for advanced correlation and reporting. Sentinel supports several built-in maps, and also customized user-defined maps.

Maps that are defined in Sentinel are stored in two ways:

- ◆ Built-in maps are stored in the database, updated internally, and automatically exported to the Mapping service.
- ◆ Custom maps are stored as CSV files and can be updated on the file system or by using the Map Data Configuration User Interface, then loaded by the Mapping service.

In both cases, the CSV files are kept on the central Sentinel server but changes to the maps are distributed to each Collector Manager and applied locally. This distributed processing ensures that mapping activity does not overload the main server.

Streaming Maps

The Map Service employs a dynamic update model and streams the maps from one point to another, avoiding the accumulation of large static maps in the dynamic memory. This is relevant in a mission-critical, real-time system such as Sentinel where a steady, predictive, and agile movement of data, independent of any transient load on the system is required.

Collector Manager

Collector Manager manages data collection, monitors system status messages, and performs event filtering. The main functions of Collector Manager include the following:

- ◆ Collecting data through the use of Connectors.
- ◆ Parsing and normalizing data through the use of Collectors.

Collectors

Collectors collect the information from the Connectors and normalize it. They perform the following functions:

- ◆ Receiving raw data from the Connectors.
- ◆ Parsing and normalizing the data:
 - ◆ Translating event-source specific data into Sentinel specific data.
 - ◆ Enriching events by changing the information in the events in a format Sentinel can read.
 - ◆ Event-source specific filtering of events.
- ◆ Adding business relevance to events through the mapping service:
 - ◆ Mapping events to Identities.
 - ◆ Mapping events to Assets.
- ◆ Routing events.
- ◆ Passing the normalized, parsed, and formatted data to the Collector Manager.
- ◆ Sending health message to the Sentinel server.

For more information about Collectors, see the [Sentinel Plug-ins website](#).

Connectors

Connectors provide connections from the event sources to the Sentinel system.

Connectors provide the following functionalities:

- ◆ Transportation of raw event data from the events sources to the Collector.
- ◆ Connection-specific filtering.
- ◆ Connection error handling.

ArcSight SmartConnectors

Sentinel leverages ArcSight SmartConnector to collect events from various types of event sources not directly supported by Sentinel. SmartConnectors collect events from supported devices, normalizes events into the Common Event Format (CEF), and forwards them to Sentinel through the Syslog Connector. The Connector then forwards the events to Universal Common Event Format Collector for parsing.

For more information about configuring Sentinel with SmartConnectors, see the Universal Common Event Format Collector documentation on the [Sentinel Plug-ins Website](#).

Agent Manager

Agent Manager provides host-based data collection that complements agentless data collection, by allowing you to perform the following tasks:

- ♦ Access logs that are not available through the network.
- ♦ Operate in tightly-controlled network environments.
- ♦ Improve security posture by limiting attack surface on critical servers.
- ♦ Provide enhanced reliability of data collection during times of network interruption.

Agent Manager allows you to deploy agents, manage agent configuration, and also acts as a collection point for events flowing into Sentinel. For more information about Agent Manager, see the [Agent Manager documentation](#).

Sentinel Data Routing and Data Storage

Sentinel provides multiple options for routing, storing, and extracting the collected data. By default, Sentinel receives the parsed event data and the raw data from the Collector Managers. Sentinel stores the raw data to provide a secure evidence chain and routes the parsed event data according to the rules you define. You can filter the parsed event data, send it to storage or to real-time analytics, and route it to external systems. Sentinel further matches all the event data that is sent to storage to user-defined retention policies. The retention policies control when event data should be deleted from the system.

Depending on the events per second (EPS) rate and your deployment requirements, you can choose to use the traditional, file-based data storage or the Hadoop-based scalable storage as the data storage option. For more information, see [“Data Storage Considerations” on page 39](#).

Event Visualizations

Sentinel provides event visualizations that present data in charts, tables, and maps. These visualizations make it easier to visualize and analyze large volumes of events including IP Flow events. You can also create your own visualizations and dashboards.

Event visualizations are available by default in Sentinel with scalable storage. In a traditional storage set up, event visualizations are available only if you enabled visualization data store (Elasticsearch) to store and index data. For more information about enabling Elasticsearch, see [“Configuring the Visualization Data Store”](#) on page 42.

Correlation

A single event might seem trivial, but in combination with other events, it might warn you of a potential problem. Sentinel helps you correlate such events by using the rules you create and deploy in the Correlation Engine, and take appropriate action to mitigate any problems.

Correlation adds intelligence to security event management by automating the 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. For more information about correlation, see [“Correlating Event Data”](#) in the *Sentinel User Guide*.

To monitor events according to the correlation rules, you must deploy the rules in the Correlation Engine. When an event occurs that matches the rule criteria, the Correlation Engine generates a correlation event describing the pattern. For more information, see [“Correlating Event Data”](#) in the *Sentinel User Guide*.

Security Intelligence

The correlation capability of Sentinel provides you the ability to look for known patterns of activity, which you can analyze for security, compliance, or any other reason. The Security Intelligence capability looks for activity that is out of the ordinary, which might be malicious, but does not match any known pattern.

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. For more information, see [“Analyzing Trends in Data”](#) in the *Sentinel User Guide*.

Incident Remediation

Sentinel provides an automated incident response management system that enables you to document and formalize the process of tracking, escalating, and responding to incidents and policy violations. It also provides two-way integration with trouble-ticketing systems. Sentinel enables you to react promptly and resolve incidents efficiently. For more information, see [“Configuring Incidents”](#) in the *Sentinel User Guide*.

iTrac Workflows

iTRAC workflows 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 starting identification (through correlation rules or manual identification) through resolution.

You can build workflows using manual and automated steps. iTrac workflows support advanced features such as branching, time-based escalation, and local variables. Integration with external scripts and plug-ins allows flexible interaction with third-party systems. Comprehensive reporting allows administrators to understand and fine-tune the incident response processes. For more information, see ["Configuring iTRAC Workflows"](#) in the *Sentinel User Guide*.

Actions and Integrators

Actions, either manually or automatically, execute some type of action, such as sending an email. You can trigger Actions by routing rules, by manually executing an event or incident operation, and by correlation rules. Sentinel provides a list of preconfigured Actions. You can use the default Actions and reconfigure them as necessary, or you can add new Actions. For more information, see ["Configuring Integrators"](#) in the *Sentinel Administration Guide*.

An Action can execute on its own, or it can make use of an Integrator instance configured from an Integrator plug-in. Integrator plug-ins extend the features and functionality of Sentinel remediation actions. Integrators provide the ability to connect to an external system, such as an LDAP, SMTP, or SOAP server, to execute an action. For more information, see ["Configuring Integrators"](#) in the *Sentinel Administration Guide*.

Searching

Sentinel provides an option to perform a search on events. With the necessary configuration, you can also search system events generated by Sentinel, and view the raw data for each event. For more information, see ["Searching Events"](#) in the *Sentinel User Guide*.

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

Reports

Sentinel provides you the ability to run reports on the gathered data. Sentinel is packaged with a variety of customizable reports. Some reports are configurable, which allow you to specify the columns to be displayed in the results.

You can run, schedule, and e-mail reports in the PDF format. You can also run any report as a search and then work with the results as you can do with a search, such as refining the search or performing an action on the results. You can also run reports on Sentinel servers distributed across different geographic locations. For more information, see ["Reporting"](#) in the *Sentinel User Guide*.

Identity Tracking

Sentinel provides an integration framework to identity management systems, to track the identities of each user account and the events those identities perform. Sentinel provides user information such as contact information, user accounts, recent authentication events, recent access events, permission changes, and so on. By displaying information about the users initiating a specific action or the users affected by an action, Sentinel improves incident response time and enables behavior-based analysis. For more information, see “[Leveraging Identity Information](#)” in the *Sentinel User Guide*.

Event Analysis

Sentinel provides a powerful set of tools to help you find and analyze critical event data easily. Sentinel optimizes the system for maximum efficiency in any type of analysis, and provides methods to transition from one type of analysis to another easily, for seamless transitions.

Investigating events in Sentinel often starts with the near real-time Event Views. Although more advanced tools are available, Event Views display filtered event streams along with summary charts that you can use for simple, quick analysis of event trends and event data, and identification of specific events. Over time, you can build up tuned filters for specific classes of data, such as output from correlation. You can use Event Views as a dashboard, which shows an overall operational and security posture.

You can then use the interactive search to perform detailed analysis of events. This allows you to quickly and easily search for and find data related to a specific query, such as activity by a specific user or on a specific system. By clicking on the event data or using the left-hand refinement pane, you can zero in on specific events of interest quickly.

When analyzing hundreds of events, the reporting capabilities of Sentinel provide custom control over event layout and can display large volumes of data. Sentinel makes this transition easier, by allowing you to transfer the interactive searches built up in the Search interface into a reporting template. This instantly creates a report that displays the same data but in a format better suited for a larger number of events.

Sentinel includes many reporting templates for this purpose. There are two types of reporting templates:

- ◆ Templates that are fine-tuned to display particular types of information, such as authentication data or user creation.
- ◆ General purpose templates that allow you to customize groups and columns on the report interactively.

Over time, you will develop commonly-used filters and reports that make your workflows easier. Sentinel supports storing this information and distributing it with people in your organization. For more information, see the *Sentinel User Guide*.



Planning Your Sentinel Installation

The following chapters guide you through planning your Sentinel installation. If you want to install a configuration that is not identified in the chapters that follow, or if you have any questions, contact [Technical Support](#).

- ◆ Chapter 3, “Implementation Checklist,” on page 31
- ◆ Chapter 4, “Understanding License Information,” on page 33
- ◆ Chapter 5, “Meeting System Requirements,” on page 37
- ◆ Chapter 6, “Deployment Considerations,” on page 39
- ◆ Chapter 7, “Deployment Considerations for FIPS 140-2 Mode,” on page 55
- ◆ Chapter 8, “Ports Used,” on page 61
- ◆ Chapter 9, “Installation Options,” on page 67

3 Implementation Checklist

Use the following checklist to plan, install, and configure Sentinel.

If you are upgrading from a previous version of Sentinel, do not use this checklist. For information about upgrading, see [Part V, “Upgrading Sentinel,”](#) on page 151.

<input type="checkbox"/> Tasks	See
<input type="checkbox"/> Review the product architecture information to learn about Sentinel components.	Part I, “Understanding Sentinel,” on page 13.
<input type="checkbox"/> Review the Sentinel licensing information to determine whether you need to use the evaluation license or the enterprise license of Sentinel.	Chapter 4, “Understanding License Information,” on page 33.
<input type="checkbox"/> Assess your environment to determine the hardware configuration. Ensure that the computers on which you install Sentinel and its components meet the specified requirements.	Chapter 5, “Meeting System Requirements,” on page 37.
<input type="checkbox"/> Determine the type of deployment suitable for your environment based on the events per second (EPS). Determine the number of Collector Managers and Correlation Engines you need to install to improve performance and load balancing.	Chapter 6, “Deployment Considerations,” on page 39.
<input type="checkbox"/> Review the latest Sentinel release notes to understand the new functionality and the known issues.	Sentinel Release Notes
<input type="checkbox"/> Install Sentinel.	Part III, “Installing Sentinel,” on page 69.
<input type="checkbox"/> Configure Sentinel.	Part IV, “Configuring Sentinel,” on page 115.
<input type="checkbox"/> Sentinel includes out-of-the-box correlation rules. Some correlation rules are configured by default, to execute an action that sends an email when the rule fires, such as the Notify Security Admin action. Therefore, you must configure the mail server settings in the Sentinel server by configuring the SMTP Integrator and the Send Email action.	SMTP Integrator and Send Email action documentation on the Sentinel Plug-ins website .
<input type="checkbox"/> Install additional Collectors and Connectors as needed in your environment.	Chapter 16, “Installing Additional Collectors and Connectors,” on page 111.
<input type="checkbox"/> Install additional Collector Managers and Correlation Engines as needed in your environment.	Part III, “Installing Sentinel,” on page 69.

4 Understanding License Information

Sentinel comprises a broad spectrum of functionality, which caters to various needs of its many customers. You can choose a licensing model that fulfills your needs.

The Sentinel platform provides the following two licensing models:

- ♦ **Sentinel Enterprise:** A full-featured solution that enables all the core, real-time visual analytics functions and many additional features. Sentinel Enterprise focuses on SIEM use cases such as real-time threat detection, alerting, and remediation.
- ♦ **Sentinel for Log Management:** A solution for log management use cases such as the ability to collect, store, search, and report on data.

Sentinel for Log Management represents a substantial upgrade from the functionality provided in Sentinel Log Manager 1.2.2, and in some cases, significant parts of the architecture have changed. To plan your upgrade to Sentinel for Log Management, see the [Sentinel FAQ page](#).

Depending on the solution(s) and add-ons you purchase, you can buy the appropriate license keys and entitlements to enable the right functionality within Sentinel. Though the license keys and entitlements govern basic access to product features and downloads, you should refer to your purchase agreement and the End-User License Agreement for additional terms and conditions.

The following table outlines the specific services and features available on each of the solutions:

Table 4-1 Sentinel Services and Features

Services and Features	Sentinel Enterprise	Sentinel for Log Management
Core Functionality	Yes	Yes
<ul style="list-style-type: none">♦ Event collection, parsing, normalization, and taxonomic classification♦ Non-event data collection (asset data, vulnerability data, and user identity data)♦ In-line contextual mapping♦ Event storage with retention policies and non-repudiation♦ Event routing to traditional storage (internal and external)♦ Event searches and visualization♦ IP Flow collection, storage, and visualization♦ Reporting♦ Federal Information Processing Standard Publication 140-2 (FIPS 140-2) enablement♦ Manually-triggered actions♦ Manual incident creation and management		

Services and Features	Sentinel Enterprise	Sentinel for Log Management
Sentinel Link	Yes	Yes
Data Synchronization	Yes	Yes
Event data restoration from archive	Yes	Yes
Data Federation (distributed search)	Yes	Yes
Scalable Storage	Yes	Yes
Correlation	Yes	No
<ul style="list-style-type: none"> ◆ Real-time event pattern correlation ◆ Actions triggered by correlation rules ◆ Alerts triage ◆ Alert visualization 		
Security Intelligence	Yes	No
<ul style="list-style-type: none"> ◆ Anomaly rules ◆ Real-time statistical analysis 		

Sentinel Licenses

This section provides information about the types of Sentinel licenses.

- ◆ [“Evaluation License” on page 34](#)
- ◆ [“Free License” on page 35](#)
- ◆ [“Enterprise Licenses” on page 35](#)

Evaluation License

The default evaluation license allows you to use all the features of Sentinel Enterprise for a specific evaluation period with unlimited EPS subject to the capacity of your hardware. For information about the features available in Sentinel Enterprise, see [Table 4-1, “Sentinel Services and Features,” on page 33](#).

The expiration date of the system is based on the oldest data in the system. If you restore old events to your system, Sentinel updates the expiration date accordingly.

After the evaluation license expires, Sentinel runs with a basic, free license that enables a limited set of features and a limited event rate of 25 EPS. This is applicable only if Sentinel is configured with traditional storage.

In scalable storage deployments, Sentinel will no longer store events and raw data when the evaluation license expires.

After you upgrade to an enterprise license, Sentinel restores all functionality. To prevent any interruption in functionality, you must upgrade the system with an enterprise license before the evaluation license expires.

Free License

The free license allows you to use a limited set of features with a limited event rate of 25 EPS. The free license is applicable only for Sentinel with traditional storage.

The free license allows you to collect and store events. When the EPS rate goes above 25, Sentinel stores the events received, but does not display the details of those events in the search results or reports. Sentinel tags these events with the `OverEPSLimit` tag.

The free license does not provide real-time features. You can restore all the functionality by upgrading the license to an enterprise license.

NOTE: Technical support and product updates are not available for the free version of Sentinel.

Enterprise Licenses

When you purchase Sentinel, you receive a license key through the customer portal. Depending on the license you purchase, your license key enables features, data collection rates, and event sources. There might be additional license terms that are not enforced by the license key, therefore read your license agreement carefully.

To make changes to your licensing, contact your account manager.

You can add the enterprise license key either during the installation or any time thereafter. To add the license key, see [“Adding a License Key”](#) in the *Sentinel Administration Guide*.

5 Meeting System Requirements

A Sentinel implementation can vary based on the needs of your IT environment, so you should contact [Consulting Services](#) or any of the Sentinel partners prior to finalizing the Sentinel architecture for your environment.

NOTE: Before you install Sentinel, ensure that your environment is secure and up-to-date with the latest security updates.

For information about the recommended hardware, supported operating systems, appliance platforms, and browsers, see the [Sentinel Technical Information website](#).

- ♦ [“Connector and Collector System Requirements” on page 37](#)
- ♦ [“Virtual Environment” on page 37](#)

Connector and Collector System Requirements

Each Connector and Collector has its own set of system requirements and supported platforms. See the Connector and Collector documentation on the [Sentinel Plug-ins website](#).

Virtual Environment

Sentinel is supported on VMware ESX servers. When you set up a virtual environment, the virtual machines must have two or more CPUs. To achieve performance results that are same as the physical machine testing results on ESX or in any other virtual environment, the virtual environment should provide the same memory, CPUs, disk space, and I/O as the physical machine recommendations.

For information about physical machine recommendations, see [Sentinel System Requirements](#).

6 Deployment Considerations

Sentinel has a scalable architecture that can grow to handle the load you need to place on it. This chapter provides an overview of the most important considerations to make when scaling a Sentinel deployment. A [Technical Support](#) or a [Partner Services](#) professional can work with you to design the Sentinel system that is suitable for your IT environment.

- ♦ [“Data Storage Considerations” on page 39](#)
- ♦ [“Advantages of Distributed Deployments” on page 46](#)
- ♦ [“All-In-One Deployment” on page 48](#)
- ♦ [“One-Tier Distributed Deployment” on page 48](#)
- ♦ [“One-Tier Distributed Deployment with High Availability” on page 49](#)
- ♦ [“Two-Tier and Three-Tier Distributed Deployment” on page 50](#)
- ♦ [“Three-Tier Deployment with Scalable Storage” on page 51](#)

Data Storage Considerations

Depending on the EPS rate, you can choose to use traditional storage or scalable storage to store and index your Sentinel data. Your Sentinel deployment depends on the data storage option you choose to use.

Table 6-1 Comparison between Traditional Storage and Scalable Storage

Traditional Storage	Scalable Storage
By default, data is stored in file-based traditional storage and indexing is done locally on the Sentinel server.	Data is stored in Hadoop-based scalable storage and uses scalable distributed indexing mechanism to index data.
In addition to file-based data storage, you can also choose to store and index events in the Visualization Data Store to leverage data visualization capabilities. For more information, see “Configuring the Visualization Data Store” on page 42.	
Seamlessly scales up to approximately 20000 EPS. Beyond that you must add additional Sentinel servers to scale up to much higher EPS.	Seamlessly scales out to a very large EPS, for example, 1 million events per second.
Data collection is load-balanced across several Sentinel servers. Therefore, data is spread across different Sentinel servers and should be managed individually.	Data collection is managed by a single Sentinel server. Therefore, data management and resource management is centralized on a single Sentinel server.
Data is labeled tenant-wise but not segregated tenant-wise on disk.	Data is labeled and segregated on disk tenant-wise.
Data replication and availability must be done either manually or by using expensive storage mechanisms such as SAN disk.	Data replication and availability is cost-effective since Hadoop runs on commodity hardware.

- ◆ [“Planning for Traditional Storage”](#) on page 40
- ◆ [“Planning for Scalable Storage”](#) on page 43
- ◆ [“Sentinel Directory Structure”](#) on page 46

Planning for Traditional Storage

File-based data storage has a three-tier structure:

Online Storage	Primary storage, formerly known as local storage.	Optimized for quick writes and fast retrieval. Stores the most recently collected event data and the most frequently searched event data.
	Secondary storage, formerly known as network storage. (optional)	Optimized to reduce space usage on optionally less expensive storage while still supporting fast retrieval. Sentinel automatically migrates data partitions to the secondary storage.
NOTE: Using the secondary storage is optional. Data retention policies, searches, and reports operate on event data partitions regardless of whether they are residing on primary or secondary storage, or both.		
Offline Storage	Archival storage	When the partitions are closed, you can back up the partition to any file storage service, such as Amazon Glacier. You can temporarily re-import the partitions for use in long-term forensic analysis whenever necessary.

You can also configure Sentinel to extract the event data and event data summaries to an external database by using data synchronization policies. For more information, see [“Configuring Data Synchronization”](#) in the *Sentinel Administration Guide*.

When you install Sentinel, you must mount the disk partition for primary storage in the location where Sentinel will be installed, by default the `/var/opt/novell` directory.

The entire directory structure under the `/var/opt/novell/sentinel` directory must reside on a single disk partition to ensure correct disk usage calculations. Else, the automatic data management capabilities might delete event data prematurely. For more information about the Sentinel directory structure, see [“Sentinel Directory Structure”](#) on page 46.

As a best practice, ensure that this data directory is located on a separate disk partition than the executables, configuration, and operating system files. The benefits of storing variable data separately include easier backup of sets of files, simpler recovery in case of corruption, and provides additional robustness if a disk partition fills up. It also improves the overall performance of systems where smaller file systems are more efficient. For more information, see [Disk Partitioning](#).

NOTE: There is a limitation in ext3 file systems for file storage, which prevents a directory from having more than 32000 files or subdirectories. You can use XFS file system if you are going to have a large number of retention policies or if you are going to retain the data for longer periods of time, such as an year.

- ♦ [“Using Partitions in Traditional Installations”](#) on page 41
- ♦ [“Using Partitions in Appliance Installations”](#) on page 42
- ♦ [“Best Practices for Partition Layout”](#) on page 42
- ♦ [“Configuring the Visualization Data Store”](#) on page 42

Using Partitions in Traditional Installations

On traditional installations, you can modify the disk partition layout of the operating system before installing Sentinel. The administrator should create and mount the desired partitions to the appropriate directories, based on the directory structure described in [“Sentinel Directory Structure”](#) on page 46. When you run the installer, Sentinel is installed into the pre-created directories resulting in an installation that spans multiple partitions.

NOTE:

- ♦ You can use the `--location` option while running the installer to specify a different top-level location than the default directories to store the file. The value that you pass to the `--location` option is prepended to the directory paths. For example, if you specify `--location=/foo`, the data directory will be `/foo/var/opt/novell/sentinel/data` and the config directory will be `/foo/etc/opt/novell/sentinel/config`.
 - ♦ You must not use filesystem links (for example, soft links) for the `--location` option.
-

Using Partitions in Appliance Installations

If you are using the DVD ISO appliance format, you can configure the partitioning of the appliance filesystem during installation by following the instructions in the YaST screens. For example, you can create a separate partition for the `/var/opt/novell/sentinel` mount point to place all data on a separate partition. However, for other appliance formats, you can configure the partitioning only after installation. You can add partitions and move a directory to the new partition by using the SuSE YaST system configuration tool. For information about creating partitions after the installation, see [“Creating Partitions for Traditional Storage” on page 107](#).

Best Practices for Partition Layout

Many organizations have their own documented best-practice partition layout schemes for any installed system. The following partition proposal is intended to guide organizations without any defined policy, and considers Sentinel specific use of the filesystem. Generally, Sentinel adheres to the [Filesystem Hierarchy Standard](#) where practicable.

Partition	Mount point	Size	Notes
Root	/	100GB	Contains operating system files and Sentinel binaries/configuration.
Boot	/boot	150MB	Boot partition
Primary storage	<code>/var/opt/novell/sentinel</code>	Calculate using the System Sizing Information .	This area will contain the primary Sentinel collected data, and other variable data such as log files. This partition can be shared with other systems.
Secondary storage	Location based on the type of storage, NFS, CIFS, or SAN.	Calculate using the System Sizing Information .	This is the secondary storage area, which can be mounted locally as shown or remotely.
Archival storage	Remote system	Calculate using the System Sizing Information .	This storage is for archived data.

Configuring the Visualization Data Store

Sentinel provides event visualizations that present data in charts, tables, and maps. These visualizations make it easier to visualize and analyze large volumes of events. You can also create your own visualizations and dashboards.

Sentinel leverages Kibana, a browser-based analytics and search dashboard, that helps you to search and visualize events. Kibana accesses data from visualization data store (Elasticsearch) to present events in dashboards. By default, Sentinel includes an Elasticsearch node that stores and indexes only alerts. You must enable event visualization to store and index events in Elasticsearch.

When you enable Elasticsearch to store and index data, Sentinel indexes only some specific event fields required for visualizations and stores the indexed fields in Elasticsearch. Sentinel creates a dedicated index for each day and uses the UTC timezone (midnight-midnight) to calculate the index date. The index name is in the `security.events.normalized_yyyyMMdd` format. For example, the index `security.events.normalized_20160101` contains all events that with an event time of January 01, 2016.

Configuring the visualization data store involves the following:

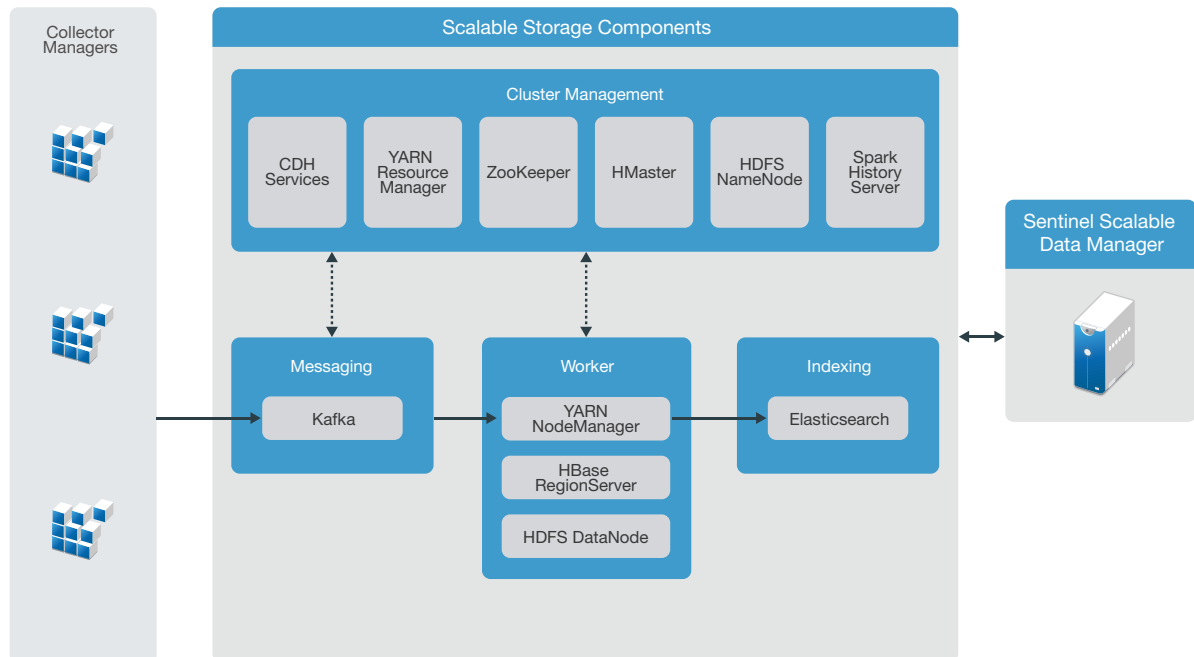
- ❑ **Installing Elasticsearch nodes in a cluster mode:** By default, Sentinel includes an Elasticsearch node. For optimal performance and stability of the Sentinel server, it is mandatory that you install additional Elasticsearch nodes in a cluster mode. For more information, see [Chapter 12, “Installing and Configuring Elasticsearch,” on page 75](#).
- ❑ **Enable event visualization:** Event visualization is disabled by default. To enable event visualization, see [Chapter 20, “Enabling Event Visualization,” on page 123](#).
- ❑ **Performance tuning:** Sentinel automatically configures certain Elasticsearch settings for optimal performance. You can customize these settings as needed. For example, you can modify the event fields you want Elasticsearch to index. For more information, see [“Performance Tuning for Elasticsearch” on page 81](#).

Planning for Scalable Storage

Sentinel uses Cloudera’s Distribution Including Apache Hadoop (CDH) framework to store and manage large data. For indexing events, Sentinel uses a scalable, distributed indexing engine called Elasticsearch from Elastic.

The following illustration explains the various components used in scalable storage:

Figure 6-1 Scalable Storage Architecture



- ◆ **Messaging:** Sentinel uses Apache Kafka as the scalable messaging system that receives normalized events and raw data from Collector Managers. Collector Managers send raw data and event data to Kafka clusters.

By default, Sentinel creates the following Kafka topics:

- ◆ **security.events.normalized:** Stores all the processed and normalized event data including system generated events and internal events.
- ◆ **security.events.raw:** Stores all the raw data from the event sources.

Event and raw data follow the Apache Avro schema. For more information, see [Apache Avro documentation](#). The schema files are available in the `/etc/opt/novell/sentinel/scalablestore` directory.

- ◆ **Worker:** This node hosts real-time processing and storage jobs. Apache Spark does large-scale data processing in real time such as segregating events based on tenant IDs, requesting large volume of data and storing data to system of record (SOR), and scalable indexing.

Apache HBase is a distributed and scalable Hadoop-based data store. It is used as an SOR for normalized events and raw data, segregated by tenant IDs.

Based on the tenant ID, Sentinel creates a separate namespace for each tenant. For example, the namespace for the default tenant is 1. Under each namespace, Sentinel creates the following tables and stores data based on the event time.

- ◆ **<tenant_ID>:security.events.normalized:** Stores all the processed and normalized event data including system generated events and internal events.
- ◆ **<tenant_ID>:security.events.raw:** Stores all the raw data from the event sources.
- ◆ **Cluster Management:** This node hosts all the masters and cluster management services. Apache ZooKeeper acts as a centralized service for maintaining configuration information, naming services, providing distributed synchronization, and providing group services.

- ♦ **Indexing:** Sentinel uses Elasticsearch as the scalable and distributed indexing engine for indexing events. You can access data from Elasticsearch for searching and visualizing events.

Sentinel creates a dedicated index for each day and uses the UTC timezone (midnight-midnight) to calculate the index date. The index name is in the `security.events.normalized_YYYYMMdd` format. For example, the index `security.events.normalized_20160101` contains all events that with an event time of January 01, 2016. For optimal performance, Sentinel indexes only some specific event fields. You can modify the event fields you want Elasticsearch to index. For more information, see [“Performance Tuning for Elasticsearch” on page 81](#).

Scalable Storage Configuration

When you enable scalable storage, the Sentinel server user interface is trimmed down to just cater to some of the Sentinel features such as data collection, correlation, event routing, search and visualize events, and perform certain administrative activities. This trimmed down version of Sentinel is referred to as Sentinel Scalable Data Manager (SSDM). For other Sentinel capabilities such as Security Intelligence, conventional searching, and reporting, you must install separate instances of Sentinel with traditional storage and route the specific event data from SSDM to Sentinel by using Sentinel Link.

The following list provides information about the services and features not available in SSDM:

- ♦ Reports
- ♦ Security Intelligence
- ♦ Performing event operations during search
- ♦ Testing correlation rules
- ♦ Incident creation and management
- ♦ Manually performing Actions on events
- ♦ Data Synchronization
- ♦ iTRAC Workflows
- ♦ Forensic analysis on the events that trigger the correlated event
- ♦ Viewing event attachments for Secure Configuration Manager and Change Guardian events

Enabling scalable storage is a one-time configuration, which cannot be reverted. If you want to disable scalable storage and switch to traditional storage, you must re-install Sentinel.

The following checklist provides a high-level information about the tasks you need to perform to configure scalable storage:

Table 6-2 Scalable Storage Configuration Checklist

Tasks	See
<input type="checkbox"/> Review the deployment information to understand how you need to deploy Sentinel with scalable storage.	“Three-Tier Deployment with Scalable Storage” on page 51
<input type="checkbox"/> Review the prerequisites and complete all the required tasks.	Chapter 13, “Installing and Setting Up Scalable Storage,” on page 85.
<input type="checkbox"/> Enable scalable storage. You can enable scalable storage either during installation or post-installation. In upgrade installations, you can enable scalable storage only after you upgrade Sentinel.	To enable scalable storage during installation, perform a custom installation of Sentinel. See “Sentinel Server Custom Installation” on page 90. To enable scalable storage post-installation or post-upgrade, see Enabling Scalable Storage Post-Installation in the <i>Sentinel Administration Guide</i> .
<input type="checkbox"/> Configure CDH components and Elasticsearch with Sentinel.	Configuring Scalable Storage in the <i>Sentinel Administration Guide</i> .

Sentinel Directory Structure

By default, the Sentinel directories are in the following locations:

- ◆ The data files are in `/var/opt/novell/sentinel/data` and `/var/opt/novell/sentinel/3rdparty` directories.
- ◆ Executables and libraries are stored in the `/opt/novell/sentinel` directory.
- ◆ Log files are in the `/var/opt/novell/sentinel/log` directory.
- ◆ Temporary files are in the `/var/opt/novell/sentinel/tmp` directory.
- ◆ Configuration files are in the `/etc/opt/novell/sentinel` directory.
- ◆ The process ID (PID) file is in the `/home/novell/sentinel/server.pid` directory.

Using the PID, administrators can identify the parent process of Sentinel server and monitor or terminate the process.

Advantages of Distributed Deployments

By default, the Sentinel server includes the following components:

- ◆ **Collector Manager:** Collector Manager provides a flexible data collection point for Sentinel.
- ◆ **Correlation Engine:** Correlation Engine processes events from the real-time event stream to determine whether they should trigger any of the correlation rules.
- ◆ **Elasticsearch:** An optional data storage component to store and index data. By default, Sentinel includes an Elasticsearch node. If you expect large EPS, more than 2500, you must deploy additional Elasticsearch nodes in a cluster.

IMPORTANT: In production environments, you should set up a distributed deployment because it isolates data collection components on a separate computer, which is important for handling spikes and other anomalies with maximum system stability.

This section describes the advantages of distributed deployments.

- ◆ [“Advantages of Additional Collector Managers” on page 47](#)
- ◆ [“Advantages of Additional Correlation Engines” on page 47](#)

Advantages of Additional Collector Managers

Sentinel server includes a Collector Manager by default. However, for production environments, distributed Collector Managers provide much better isolation when large volumes of data is received. In this situation, a distributed Collector Manager may become overloaded but the Sentinel server will remain responsive to user requests.

Installing more than one Collector Manager in a distributed network provides the following advantages:

- ◆ **Improved system performance:** Additional Collector Managers can parse and process event data in a distributed environment, which increases the system performance.
- ◆ **Additional data security and decreased network bandwidth requirements:** If the Collector Managers are co-located with event sources, then filtering, encryption, and data compression can be performed at the source.
- ◆ **File caching:** Additional Collector Managers can cache large amounts of data while the server is temporarily busy archiving events or processing a spike in events. This feature is an advantage for protocols such as syslog, which do not natively support event caching.

You can install additional Collector Managers at suitable locations in your network. These remote Collector Managers run Connectors and Collectors, and forward the collected data to the Sentinel server for storage and processing. For information about installing additional Collector Managers, see [Part III, “Installing Sentinel,” on page 69](#).

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

Advantages of Additional Correlation Engines

You can deploy multiple Correlation Engines, each on its own server, without the need to replicate configurations or add databases. In environments with large numbers of correlation rules or extremely high event rates, it is advantageous to install more than one Correlation Engine and redeploy some rules to the new Correlation Engine. Multiple Correlation Engines provide the ability to scale as the Sentinel system incorporates additional data sources, or as event rates increase. For information about installing additional Correlation Engines, see [Part III, “Installing Sentinel,” on page 69](#).

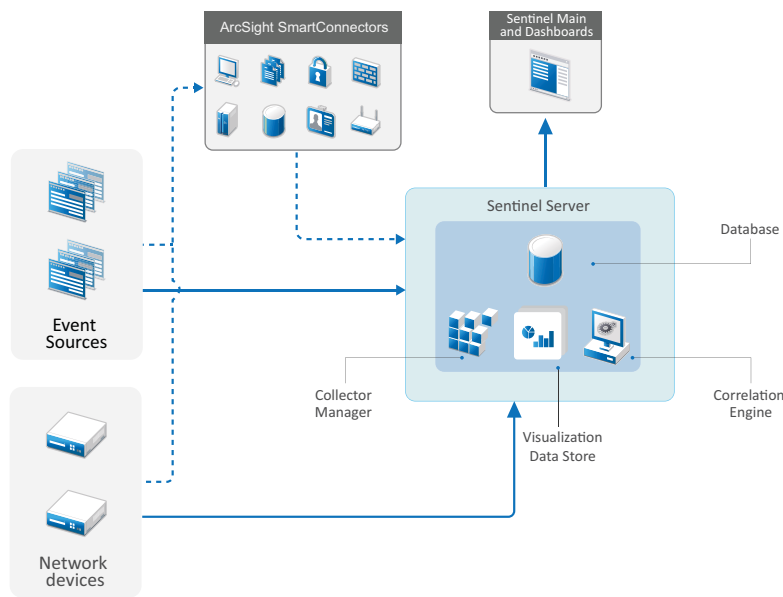
NOTE: You cannot install more than one Correlation Engine on a single system. You can install additional Correlation Engines on remote systems, and then connect them to the Sentinel server.

All-In-One Deployment

The most basic deployment option is an all-in-one system that contains all of the Sentinel components on a single computer. All-in-one deployment is suitable only if there is a small load on the system and you do not need to monitor Windows machines. In many environments, unpredictable and fluctuating loads, and resource conflicts between components can cause performance issues.

IMPORTANT: For production environments, you should set up a distributed deployment because it isolates data collection components on a separate computer, which is important for handling spikes and other anomalies with maximum system stability.

Figure 6-2 All-In-One Deployment

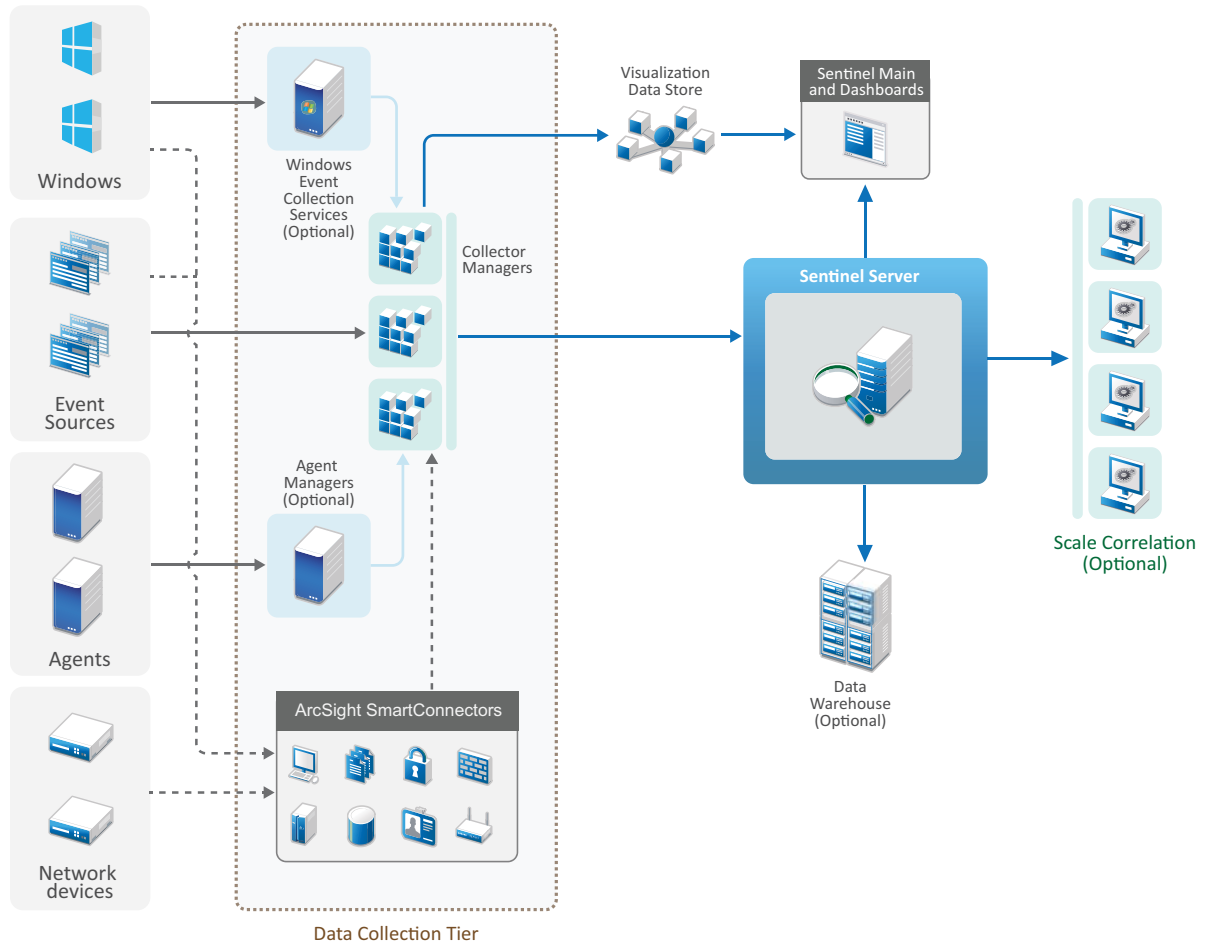


One-Tier Distributed Deployment

The one-tier deployment adds the ability to monitor Windows computers and to handle a larger load than the all-in-one deployment. You can scale out data collection and correlation by adding Collector Manager and Correlation Engine computers that offload processing from the central Sentinel server. In addition to handling the load of events and correlation rules, remote Collector Managers and Correlation Engines also free up resources on the central Sentinel server to service other requests such as event storage and searches. As the load gets higher on the system, the central Sentinel server will eventually become a bottleneck and you need a deployment with more tiers to scale out further.

Optionally, you can configure Sentinel to copy event data to a data warehouse, which can be useful to offload custom reporting, analytics, and other processing to another system.

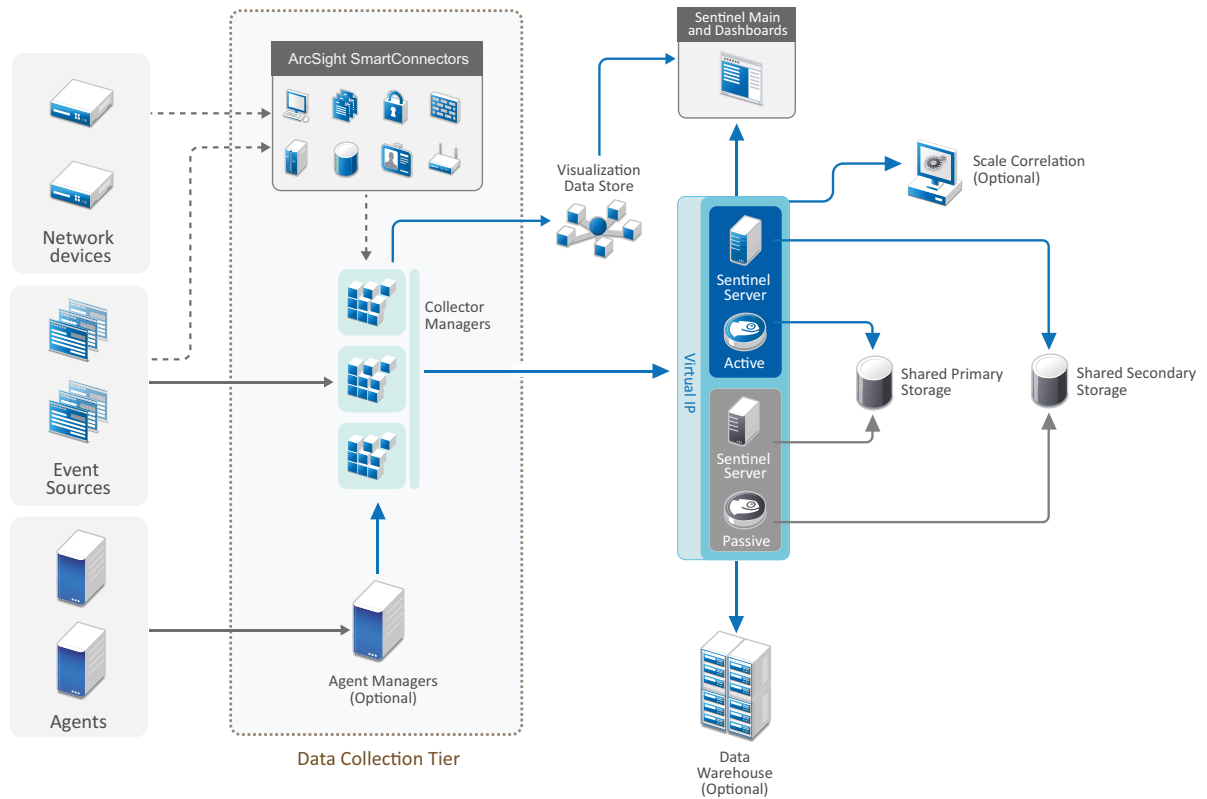
Figure 6-3 One-Tier Distributed Deployment



One-Tier Distributed Deployment with High Availability

The one-tier distributed deployment shows how it can be turned into a highly available system with fail-over redundancy. For more information about deploying Sentinel in High Availability, see [Part VII, “Deploying Sentinel for High Availability,”](#) on page 197.

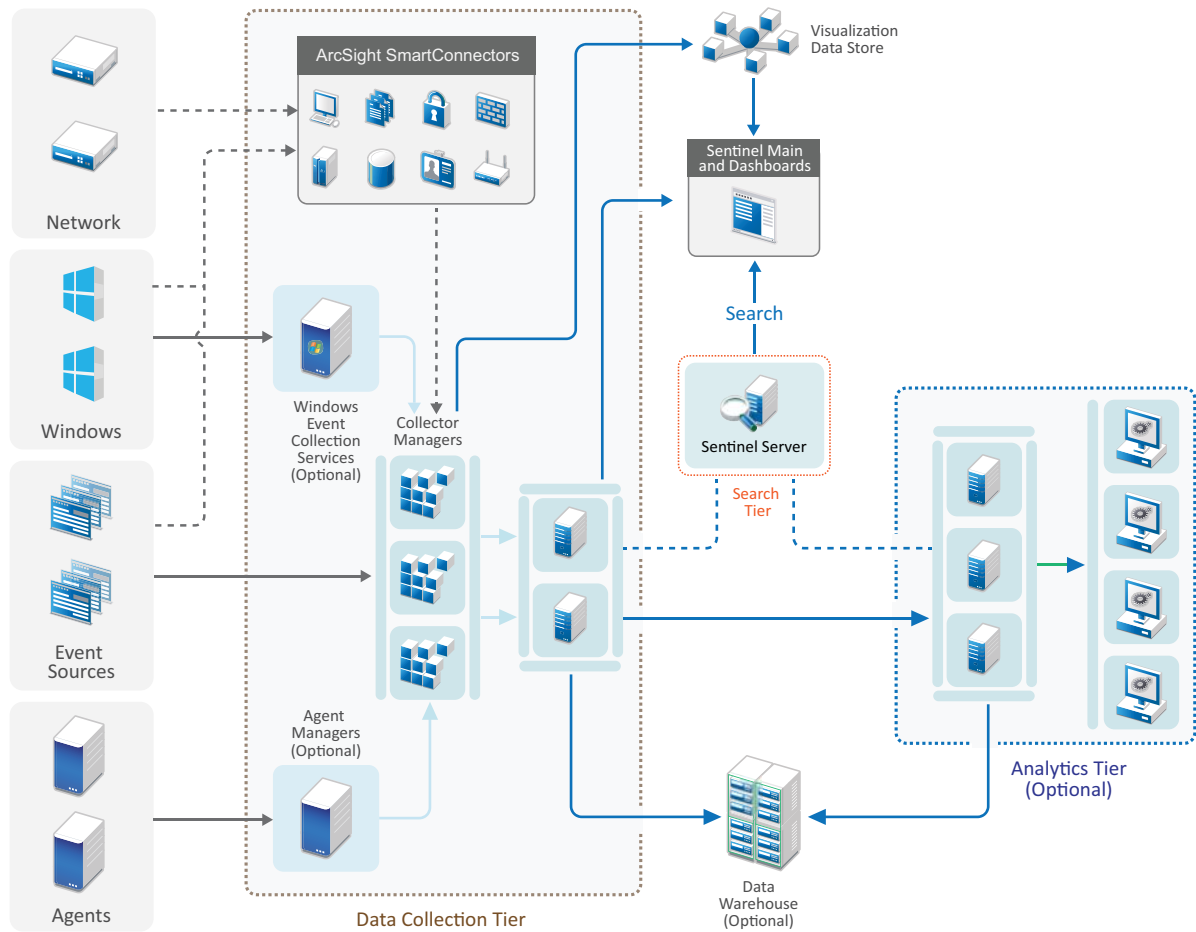
Figure 6-4 One-Tier Distributed Deployment with High Availability



Two-Tier and Three-Tier Distributed Deployment

These deployments enable you to surpass the load handling capabilities of a single central Sentinel server and share the processing load across multiple Sentinel instances by leveraging Sentinel Link and Sentinel Data Federation features. The data collection is load-balanced across several Sentinel servers, each having several Collector Managers, as shown in the Data Collection Tier. If you want to perform event correlation or security intelligence, you can optionally forward data up to the Analytics Tier using Sentinel Link. The Search Tier provides a convenient single access point for searching across all systems in all other tiers by using Sentinel Data Federation. As the search request is federated across several instances of Sentinel, this deployment also has search load-balancing properties useful in scaling to handle a heavy search load.

Figure 6-5 Two-Tier and Three-Tier Distributed Deployment



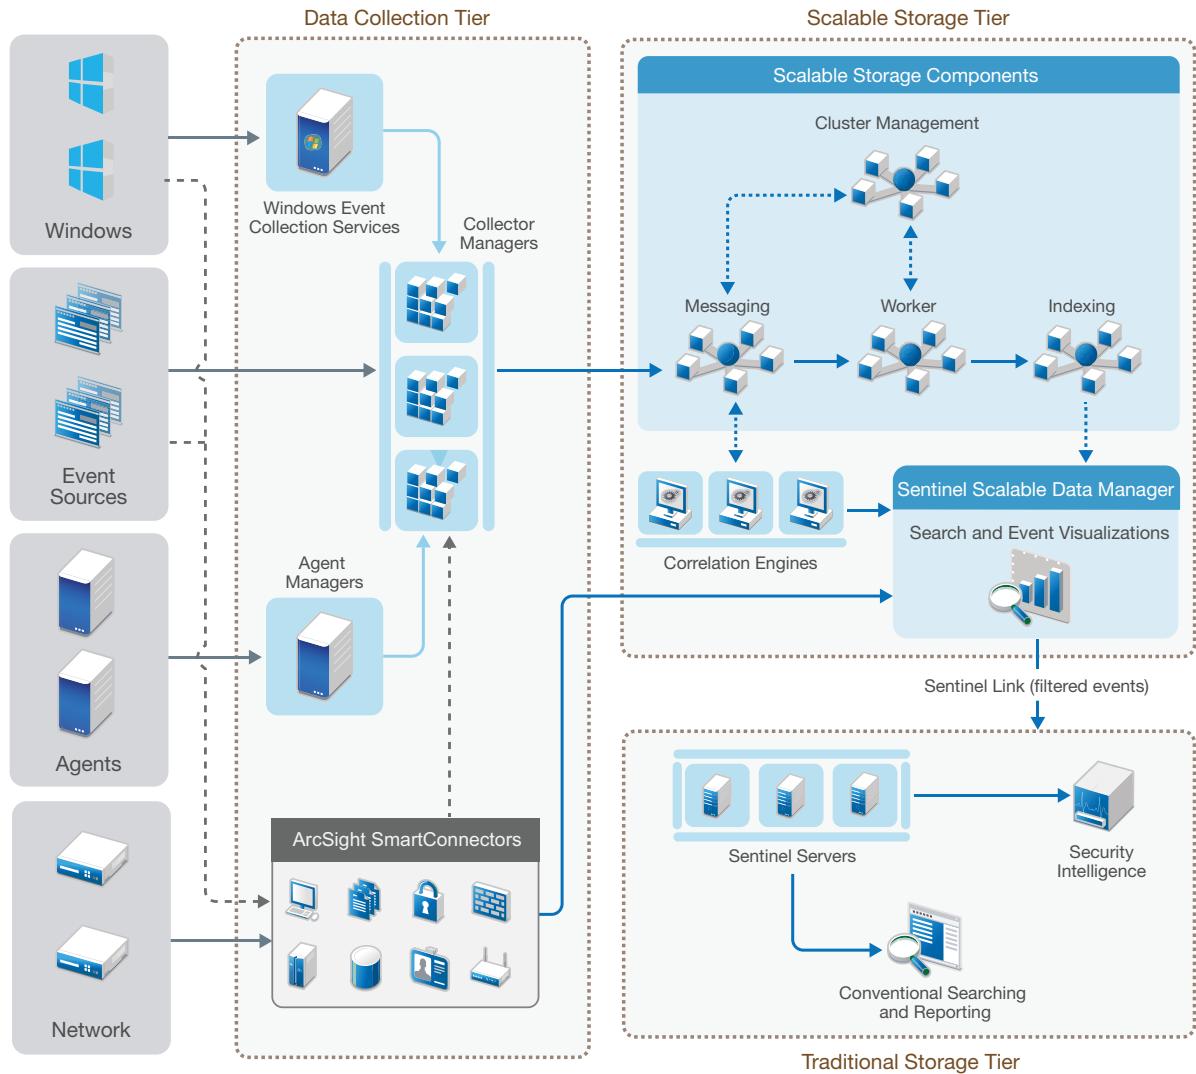
Three-Tier Deployment with Scalable Storage

For large data storage and data processing needs where you do not want to distribute events across multiple Sentinel servers and duplicate configuration settings across multiple instances, you can set up a three-tier distributed deployment with scalable storage. This deployment enables you to store and manage large data by using a single Sentinel server with scalable storage versus using multiple Sentinel servers.

You can set up a new Sentinel server with scalable storage or upgrade your existing Sentinel server to enable scalable storage.

Depending on the Sentinel capabilities you want to use, you can determine how you want to set up your Sentinel deployment.

Figure 6-6 Three-Tier Deployment for Scalable Storage



This deployment includes the following tiers:

- Data Collection Tier:** For collecting events from a wide range of event sources. Optionally, if you want to retain your existing data collection setup with traditional storage Sentinel and still leverage the scalable storage capabilities, you can forward the desired events directly from traditional storage to scalable storage by using the `data_uploader.sh` script. For more information, see [Chapter 36, “Migrating Data to Scalable Storage,”](#) on page 187.
- Scalable Storage Tier:** For storing, indexing, and analyzing large data. The SSDM server in this tier enables you to manage data collection and correlation, and provides other SSDM capabilities. To use Sentinel capabilities not available in SSDM, you can set up the Traditional Storage tier. You can also forward the collected data to any other SIEM systems or enable other business intelligence tools to query the data or perform analytics directly on your Hadoop distribution using the widely supported Hadoop, Kafka, Spark, and Elasticsearch APIs.

- ♦ **Traditional Storage Tier:** For Sentinel capabilities such as Security Intelligence, conventional searching, and reporting, you must install separate instances of Sentinel with traditional storage. You can configure event routing rules to forward the desired events from SSDM to Sentinel by using Sentinel Link.

You can perform searching and reporting using any of the Sentinel servers in the Traditional Storage Tier. Optionally, you can set up a separate Search Tier that provides a convenient single access point for searching and reporting across all Sentinel servers in the Traditional Storage Tier. For searching events in the scalable storage, use the search option in SSDM.

For more information about installing and setting up scalable storage, see [Chapter 13, “Installing and Setting Up Scalable Storage,”](#) on page 85.

7 Deployment Considerations for FIPS 140-2 Mode

You can optionally configure Sentinel to use Mozilla Network Security Services (NSS), which is a FIPS 140-2 validated cryptographic provider, for its internal encryption and other functions. The purpose of doing so is to ensure that Sentinel is 'FIPS 140-2 Inside' and is compliant with United States federal purchasing policies and standards.

Enabling Sentinel FIPS 140-2 mode causes communication between the Sentinel Server, Sentinel remote Collector Managers, Sentinel remote Correlation Engines, the Sentinel Main interface, and the Sentinel Control Center to use FIPS 140-2 validated cryptography.

IMPORTANT: FIPS mode is supported only for Sentinel. Sentinel is not supported if the operating system is in FIPS mode.

- ◆ “FIPS Implementation in Sentinel” on page 55
- ◆ “FIPS-Enabled Components in Sentinel” on page 56
- ◆ “Data Connections Affected by FIPS Mode” on page 57
- ◆ “Implementation Checklist” on page 57
- ◆ “Deployment Scenarios” on page 58

FIPS Implementation in Sentinel

Sentinel uses the Mozilla NSS libraries that are provided by the operating system. Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) have different set of NSS packages.

The NSS cryptographic module provided by RHEL 6.3 and later is FIPS 140-2 validated. The NSS cryptographic module included in SLES 11 are not yet officially FIPS 140-2 validated, but work is in progress to get the SUSE module FIPS 140-2 validated. Once the validation is available, no necessary changes to Sentinel are anticipated to provide 'FIPS 140-2 Inside' on the SUSE platform.

For more information about RHEL FIPS 140-2 certification, see <https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/2711> and <https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/1837>.

RHEL NSS Packages

Sentinel requires the following 64-bit NSS packages to support FIPS 140-2 mode:

- ◆ nspr-*
- ◆ nss-sysinit-*
- ◆ nss-util-*
- ◆ nss-softokn-freebl-*

- ♦ nss-softokn-*
- ♦ nss-*
- ♦ nss-tools-*

If any of these packages are not installed, you must install them before enabling FIPS 140-2 mode in Sentinel.

SLES NSS Packages

Sentinel requires the following 64-bit NSS packages to support FIPS 140-2 mode:

- ♦ libfreebl3-*
- ♦ mozilla-nspr-*
- ♦ mozilla-nss-*
- ♦ mozilla-nss-tools-*

If any of these packages are not installed, you must install them before enabling FIPS 140-2 mode in Sentinel.

FIPS-Enabled Components in Sentinel

The following Sentinel components provide FIPS 140-2 support:

- ♦ All Sentinel platform components are updated to support FIPS 140-2 mode.
- ♦ The following Sentinel plug-ins that support cryptography are updated to support FIPS 140-2 mode:
 - ♦ Agent Manager Connector 2011.1r1 and later
 - ♦ Database (JDBC) Connector 2011.1r2 and later
 - ♦ File Connector 2011.1r1 and later (only if the file event source type is local or NFS)
 - ♦ LDAP Integrator 2011.1r1 and later
 - ♦ Sentinel Link Connector 2011.1r3 and later
 - ♦ Sentinel Link Integrator 2011.1r2 and later
 - ♦ SMTP Integrator 2011.1r1 and later
 - ♦ Syslog Connector 2011.1r2 and later
 - ♦ Windows Event (WMI) Connector 2011.1r2 and later
 - ♦ Check Point (LEA) Connector 2011.1r2 and later
 - ♦ Syslog Integrator 2011.1r1 and later

For more information about configuring these Sentinel plug-ins to run in FIPS 140-2 mode, see [“Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode”](#) on page 133.

The following Sentinel Connectors that support optional cryptography are not yet updated to support FIPS 140-2 mode at the time of release of this document. However, you can continue to collect events using these Connectors. For information about using these Connectors with Sentinel in FIPS 140-2 mode, see [“Using Non-FIPS Enabled Connectors with Sentinel in FIPS 140-2 Mode” on page 140](#).

- ◆ Cisco SDEE Connector 2011.1r1
- ◆ File Connector 2011.1r1 - The CIFS and SCP functionalities involve cryptography and will not work in FIPS 140-2 mode.
- ◆ Audit Connector 2011.1r1
- ◆ SNMP Connector 2011.1r1

The following Sentinel Integrators that support SSL are not updated to support FIPS 140-2 mode at the time of release of this document. However, you can continue to use unencrypted connections when these Integrators are used with Sentinel in FIPS 140-2 mode.

- ◆ Remedy Integrator 2011.1r1 or later
- ◆ SOAP Integrator 2011.1r1 or later

Any other Sentinel plug-ins that are not listed above do not use cryptography and are not affected by enabling FIPS 140-2 mode in Sentinel. You do not need to perform any additional steps to use them with Sentinel in FIPS 140-2 mode.

For more information about the Sentinel plug-ins, see [Sentinel Plug-ins website](#). If you want to request that any of the plug-ins that has not yet been updated be made available with FIPS support, please submit a request using [Bugzilla](#).

Data Connections Affected by FIPS Mode

If Sentinel is in FIPS 140-2 mode, you cannot make encrypted connections to Microsoft SQL Server. This consideration affects the following types of Sentinel operations:

- ◆ Data synchronization policies to SQL Server
- ◆ Sentinel server communicating with the Agent Manager database
- ◆ Database Connector collecting data from SQL Server

Implementation Checklist

The following table provides an overview of the tasks required to configure Sentinel for operation in FIPS 140-2 mode.

Tasks	For more information, see...
Plan the deployment.	“Deployment Scenarios” on page 58 .

Tasks	For more information, see...
<p>Determine whether you need to enable FIPS 140-2 mode during the Sentinel installation or you want to enable it in future.</p> <p>To enable Sentinel in FIPS 140-2 mode during the installation, you need to select the Custom or Silent installation method during the installation process.</p>	<p>“Sentinel Server Custom Installation” on page 90.</p> <p>“Performing a Silent Installation” on page 96</p> <p>Chapter 23, “Enabling FIPS 140-2 Mode in an Existing Sentinel Installation,” on page 129</p>
<p>Configure Sentinel Plug-ins to run in FIPS 140-2 Mode.</p>	<p>“Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode” on page 133.</p>
<p>Import certificates into the Sentinel FIPS Keystore.</p>	<p>“Importing Certificates into FIPS Keystore Database” on page 140</p>

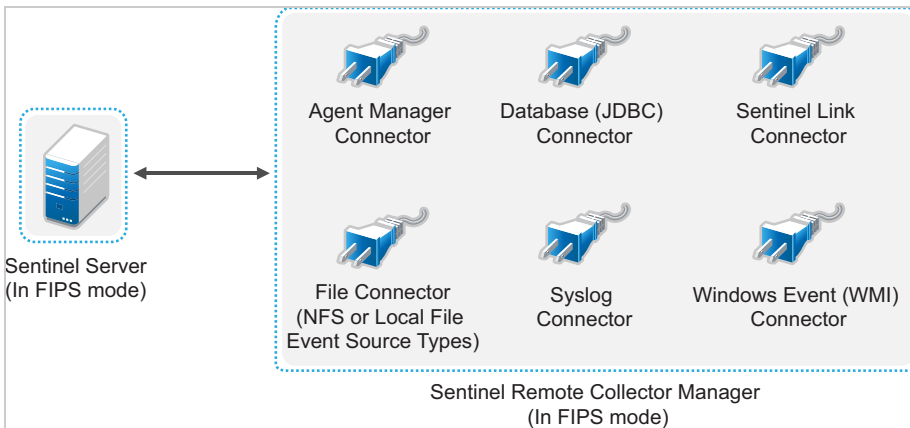
NOTE: Back up your Sentinel systems before beginning the conversion to FIPS mode. If the server must be reverted to non-FIPS mode at a later time, the only supported method for doing so involves restoring from a backup. For more information about reverting to non-FIPS mode, see [“Reverting Sentinel to Non-FIPS Mode” on page 141.](#)

Deployment Scenarios

This section provides information about the deployment scenarios for Sentinel in FIPS 140-2 mode.

Scenario 1: Data Collection in Full FIPS 140-2 Mode

In this scenario, data collection is done only through the Connectors that support FIPS 140-2 mode. We assume that this environment involves a Sentinel server and data is collected through a remote Collector Manager. You may have one or more remote Collector Managers.



You must perform the following procedure only if your environment involves data collection from event sources using Connectors that support FIPS 140-2 mode.

- 1 You must have a Sentinel server in FIPS 140-2 mode.

NOTE: If your Sentinel server (freshly installed or upgraded) is in non-FIPS mode, you must enable FIPS on Sentinel server. For more information, see [“Enabling Sentinel Server to Run in FIPS 140-2 Mode”](#) on page 129.

2 You must have a Sentinel remote Collector Manager running in FIPS 140-2 mode.

NOTE: If your remote Collector Manager (freshly installed or upgraded) is running in non-FIPS mode, you must enable FIPS on the remote Collector Manager. For more information, see [“Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines”](#) on page 130.


3 Ensure that FIPS server and remote Collector Managers communicate with each other.

4 Convert remote Correlation Engines if any to run in FIPS mode. For more information, see [“Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines”](#) on page 130.

5 Configure Sentinel plug-ins to run in FIPS 140-2 mode. For more information, see [“Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode”](#) on page 133.

Scenario 2: Data Collection in Partial FIPS 140-2 Mode

In this scenario, data collection is done using Connectors that support FIPS 140-2 mode and Connectors that do not support FIPS 140-2 mode. We assume data is collected through a remote Collector Manager. You may have one or more remote Collector Managers.



To handle data collection using Connectors that support and those that do not support the FIPS 140-2 mode, you should have two remote Collector Managers - one running in FIPS 140-2 mode for FIPS supported Connectors, and another running in non-FIPS (normal) mode for Connectors that do not support the FIPS 140-2 mode.

You must perform the following procedure if your environment involves data collection from event sources using Connectors that support FIPS 140-2 mode and Connectors that do not support FIPS 140-2 mode.

- 1 You must have a Sentinel server in FIPS 140-2 mode.

NOTE: If your Sentinel server (freshly installed or upgraded) is in non-FIPS mode, you must enable FIPS on Sentinel server. For more information, see [“Enabling Sentinel Server to Run in FIPS 140-2 Mode” on page 129.](#)

- 2 Ensure that one remote Collector Manager is running in FIPS 140-2 mode, and another remote Collector Manager continues to run in non-FIPS mode.
 - 2a If you do not have a FIPS 140-2 mode enabled remote Collector Manager, you must enable FIPS mode on the remote Collector Manager. For more information, see [“Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines” on page 130.](#)
 - 2b Update the server certificate on the non-FIPS remote Collector Manager. For more information, see [“Updating Server Certificates in Remote Collector Managers and Correlation Engines” on page 133.](#)
- 3 Ensure that the two remote Collector Managers communicate with FIPS 140-2 enabled Sentinel server.
- 4 Configure the Remote Correlation Engines if any to run in FIPS 140-2 mode. For more information, see [“Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines” on page 130.](#)
- 5 Configure the Sentinel plug-ins to run in FIPS 140-2 mode. For more information, see [“Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode” on page 133.](#)
 - 5a Deploy Connectors that support FIPS 140-2 mode in the remote Collector Manager running in FIPS mode.
 - 5b Deploy the Connectors that do not support FIPS 140-2 mode in the non-FIPS remote Collector Manager.

8 Ports Used

Sentinel uses various ports for external communication with other components. For the appliance installation, the ports are opened on the firewall by default. However, for the traditional installation, you must configure the operating system on which you are installing Sentinel to open the ports on the firewall.

- ♦ [“Sentinel Server Ports” on page 61](#)
- ♦ [“Collector Manager Ports” on page 63](#)
- ♦ [“Correlation Engine Ports” on page 65](#)
- ♦ [“Scalable Storage Ports” on page 65](#)

Sentinel Server Ports

The Sentinel server uses the following ports for internal and external communication.

Local Ports

Sentinel uses the following ports for internal communication with database and other internal processes:

Ports	Description
TCP 27017	Used for the Security Intelligence configuration database.
TCP 28017	Used for the web console for Security Intelligence database.
TCP 32000	Used for internal communication between the wrapper process and the server process.
TCP 9200	Used for communication with alert indexing service using REST.
TCP 9300	Used for communication with alert indexing service using its native protocol.

Network Ports

For Sentinel to work correctly, ensure that the following ports are open on the firewall:

Ports	Direction	Required/Optional	Description
TCP 5432	Inbound	Optional. By default, this port listens only on loopback interface.	Used for the PostgreSQL database. You do not need to open this port by default. However, you must open this port when you develop reports by using the Sentinel SDK. For more information, see the Sentinel Plug-in SDK .

Ports	Direction	Required/ Optional	Description
TCP 1099 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).
TCP 1289	Inbound	Optional	Used for Audit connections.
UDP 1514	Inbound	Optional	Used for syslog messages.
TCP 8443	Inbound	Required	Used for HTTPS communication.
TCP 1443	Inbound	Optional	Used for SSL encrypted syslog messages.
TCP 61616	Inbound	Optional	Used for incoming connections from Collector Managers and Correlation Engines.
TCP 10013	Inbound	Required	Used by the Sentinel Control Center and Solution Designer.
TCP 1468	Inbound	Optional	Used for syslog messages.
TCP 10014	Inbound	Optional	Used by the remote Collector Managers to connect to the server through the SSL proxy. However, this is uncommon. By default, remote Collector Managers use the SSL port 61616 to connect to the server.
TCP 8443	Outbound	Optional	If data federation is used, the port initiates a connection to other Sentinel systems to perform distributed search.
TCP 389 or 636	Outbound	Optional	If LDAP authentication is used, the port initiates a connection to the LDAP server.
TCP/UDP 111 and TCP/UDP 2049	Outbound	Optional	If secondary storage is configured to use NFS.
TCP 137, 138, 139, 445	Outbound	Optional	If secondary storage is configured to use CIFS.
TCP JDBC (database dependent)	Outbound	Optional	If data synchronization is used, the port initiates a connection to the target database using JDBC. The port that is used is dependent on the target database.
TCP 25	Outbound	Optional	Initiates a connection to the email server.
TCP 1290	Outbound	Optional	When Sentinel forwards events to another Sentinel system, this port initiates a Sentinel Link connection to that system.
UDP 162	Outbound	Optional	When Sentinel forwards events to the system receiving SNMP traps, the port sends a packet to the receiver.
UDP 514 or TCP 1468	Outbound	Optional	This port is used when Sentinel forwards events to the system receiving Syslog messages. If the port is UDP, it sends a packet to the receiver. If the port is TCP, it initiates a connection to the receiver.
TCP 7443	Inbound	Optional	This port allows a Sentinel system to receive events from other SIEM software such as Change Guardian and Secure Configuration Manager.

Sentinel Server Appliance Specific Ports

In addition to the above ports, the following ports are open for appliance.

Ports	Direction	Required/Optional	Description
TCP 22	Inbound	Required	Used for secure shell access to the Sentinel appliance.
TCP 4984	Inbound	Required	Also used by the Sentinel appliance for the update service.
TCP 289	Inbound	Optional	Forwarded to 1289 for Audit connections.
TCP 443	Inbound	Optional	Forwarded to 8443 for HTTPS communication.
UDP 514	Inbound	Optional	Forwarded to 1514 for syslog messages.
TCP 1290	Inbound	Optional	Sentinel Link port that is allowed to connect through the SuSE Firewall.
UDP and TCP 40000 - 41000	Inbound	Optional	Ports that can be used when configuring data collection servers, such as syslog. Sentinel does not listen on these ports by default.
TCP 443 or 80	Outbound	Required	Initiates a connect to the appliance software update repository on the Internet or a Subscription Management Tool service in your network.
TCP 80	Outbound	Optional	Initiates a connection to the Subscription Management Tool.
TCP 7630	Inbound	Required	Used by the High Availability Web Konsole (Hawk).
TCP 9443	Inbound	Required	Used by the Sentinel Appliance Management Console.
TCP 1098 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).
TCP 7443	Inbound	Required	Used by the HTTP Server Connector.

Collector Manager Ports

The Collector Manager uses the following ports to communicate with other components.

Network Ports

For Sentinel Collector Manager to work properly, ensure that the following ports are open on the firewall:

Ports	Direction	Required/ Optional	Description
TCP 1289	Inbound	Optional	Used for Audit connections.
UDP 1514	Inbound	Optional	Used for syslog messages.
TCP 1443	Inbound	Optional	Used for SSL encrypted syslog messages.
TCP 1468	Inbound	Optional	Used for syslog messages.
TCP 1099 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).
TCP 61616	Outbound	Required	Initiates a connection to the Sentinel server.
TCP 8443	Outbound	Required	Initiates a connection to the Sentinel web server port. Leave this port open only during installation and configuration of Collector Manager.
TCP 7443	Inbound	Required	Used by the HTTP Server Connector.

Collector Manager Appliance Specific Ports

In addition to the above ports, the following ports are open for the Sentinel Collector Manager appliance.

Ports	Direction	Required/ Optional	Description
TCP 22	Inbound	Required	Used for secure shell access to the Sentinel appliance.
TCP 4984	Inbound	Required	Also used by the Sentinel appliance for the update service.
TCP 289	Inbound	Optional	Forwarded to 1289 for Audit connections.
UDP 514	Inbound	Optional	Forwarded to 1514 for syslog messages.
TCP 1290	Inbound	Optional	This is the Sentinel Link port that is allowed to connect through the SuSE Firewall.
UDP and TCP 40000 - 41000	Inbound	Optional	Used when configuring data collection servers, such as syslog. Sentinel does not listen on these ports by default.
TCP 443	Outbound	Required	Initiates a connection to the appliance software update repository on the Internet or a Subscription Management Tool service in your network.
TCP 80	Outbound	Optional	Initiates a connection to the Subscription Management Tool.
TCP 9443	Inbound	Required	Used by the Sentinel Appliance Management Console.
TCP 1098 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).
TCP 7443	Inbound	Required	Used by the HTTP Server Connector.

Correlation Engine Ports

The Correlation Engine uses the following ports to communicate with other components.

Network Ports

For the Sentinel Correlation Engine to work correctly, ensure that the following ports are open on the firewall:

Ports	Direction	Required/Optional	Description
TCP 1099 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).
TCP 61616	Outbound	Required	Initiates a connection to the Sentinel server.
TCP 8443	Outbound	Required	Initiates a connection to the Sentinel web server port. Leave this port open only during installation and configuration of Correlation Engine.

Correlation Engine Appliance Specific Ports

In addition to the above ports, the following ports are open on Sentinel Correlation Engine appliance.

Ports	Direction	Required/Optional	Description
TCP 22	Inbound	Required	Used for secure shell access to the Sentinel appliance.
TCP 4984	Inbound	Required	Also used by the Sentinel appliance for the update service.
TCP 443	Outbound	Required	Initiates a connection to the appliance software update repository on the Internet or a Subscription Management Tool service in your network.
TCP 80	Outbound	Optional	Initiates a connection to the Subscription Management Tool.
TCP 9443	Inbound	Required	Used by the Sentinel Appliance Management Console.
TCP 1098 and 2000	Inbound	Required	Used together by monitoring tools to connect to Sentinel server process using Java Management Extensions (JMX).

Scalable Storage Ports

For SSDM to communicate successfully with CDH and Elasticsearch, ensure that the ports you specify during scalable storage configuration are open on the firewall in addition to the ports required by Cloudera and the ports listed in the [Sentinel Server Ports](#) section.

9 Installation Options

You can perform a traditional installation of Sentinel or install the appliance. This chapter provides information about the two installation options.

Traditional Installation

The traditional installation installs Sentinel on an existing operating system, by using the application installer. You can install Sentinel in the following ways:

- ♦ **Interactive:** The installation proceeds with user inputs. During installation, you can record the installation options (user inputs or default values) to a file, which you can use later for silent installation. You can either perform a standard installation or a custom installation.

Standard Installation	Custom Installation
Uses the default values for the configuration. User input is required only for the password.	Prompts you to specify the values for the configuration setup. You can either select the default values or specify the necessary values.
Installs with default evaluation key.	Allows you to install with the default evaluation license key or with a valid license key.
Allows you to specify the admin password and uses the admin password as the default password for both dbauser and appuser.	Allows you to specify the admin password. For dbauser and appuser, you can either specify new password or use admin password.
Installs the default ports for all the components.	Allows you to specify ports for different components.
Installs Sentinel in non-FIPS mode.	Allows you to install Sentinel in FIPS 140-2 mode.
Uses traditional storage to store raw data and events.	Allows you to use scalable storage to store raw data and events.
Authenticates users with the internal database.	Provides the option set up LDAP authentication for Sentinel in addition to the database authentication. When you configure Sentinel for LDAP authentication, users can log in to the server by using their Novell eDirectory or Microsoft Active Directory credentials.

For more information about interactive installation, see [“Performing Interactive Installation” on page 89](#).

- ♦ **Silent:** If you want to install multiple Sentinel servers, Collector Managers, or Correlation Engines in your deployment, you can record the installation options during the standard or custom installation in a configuration file and then use the file to run an silent installation. For more information on silent installation, see [“Performing a Silent Installation” on page 96](#).

Appliance Installation

The appliance installation installs both the SLES operating system and Sentinel.

The Sentinel appliance is available in the following formats:

- ◆ An OVF appliance image
- ◆ An ISO appliance image

For more information about appliance installation, see [Chapter 15, “Appliance Installation,”](#) on [page 101](#).



Installing Sentinel

This section provides information about installing Sentinel and additional components.

- ♦ [Chapter 10, “Installation Overview,” on page 71](#)
- ♦ [Chapter 11, “Installation Checklist,” on page 73](#)
- ♦ [Chapter 12, “Installing and Configuring Elasticsearch,” on page 75](#)
- ♦ [Chapter 13, “Installing and Setting Up Scalable Storage,” on page 85](#)
- ♦ [Chapter 14, “Traditional Installation,” on page 89](#)
- ♦ [Chapter 15, “Appliance Installation,” on page 101](#)
- ♦ [Chapter 16, “Installing Additional Collectors and Connectors,” on page 111](#)
- ♦ [Chapter 17, “Verifying the Installation,” on page 113](#)

10 Installation Overview

The default Sentinel installation installs the following components in the Sentinel server:

- ♦ **Sentinel server and Web server processes:** The Sentinel server process processes requests from other components of Sentinel and enables seamless functionality of the system. The Sentinel server process handles requests, such as filtering data, processing search queries, and managing administrative tasks that include user authentication and authorization.

The Sentinel Web server allows secure connection to the Sentinel Main interface.

- ♦ **PostgreSQL database:** Sentinel has a built-in database that stores Sentinel configuration information, asset and vulnerability data, identity information, incident and workflow status, Security Intelligence, alerts data, and so on.
- ♦ **Elasticsearch:** Indexes events and alerts for searching and visualization. An optional data storage component to store and index data. By default, Sentinel includes an Elasticsearch node. If you expect large EPS, more than 2500, you must deploy additional Elasticsearch nodes in a cluster
- ♦ **Collector Manager:** Collector Manager provides a flexible data collection point for Sentinel. The Sentinel installer installs a Collector Manager by default during installation.
- ♦ **Correlation Engine:** Correlation Engine processes events from the real-time event stream to determine whether they should trigger any of the correlation rules.
- ♦ **Sentinel plug-ins:** Sentinel supports a variety of plug-ins to expand and enhance system functionality. Some of these plug-ins are preinstalled. You can download additional plug-ins and updates from the [Sentinel Plug-ins website](#). Sentinel plug-ins include the following:
 - ♦ Collectors
 - ♦ Connectors
 - ♦ Correlation rules and actions
 - ♦ Reports
 - ♦ iTRAC workflows
 - ♦ Solution packs

11 Installation Checklist

Ensure that you have completed the following tasks before you start the installation:

- Verify that your hardware and software meet the system requirements listed in [Chapter 5, “Meeting System Requirements,”](#) on page 37.
- If there was a previous installation of Sentinel, ensure that there are no files or system settings remaining from a previous installation. For more information, see [Appendix B, “Uninstalling,”](#) on page 243.
- If you plan to install the licensed version, obtain your license key from the [Customer Care Center](#).
- Ensure that the ports listed in [Chapter 8, “Ports Used,”](#) on page 61 are opened in the firewall.
- For the Sentinel installer to work properly, the system must be able to return the hostname or a valid IP address. To do this, add the hostname to the `/etc/hosts` file to the line containing the IP address, then enter `hostname -f` to make sure that the hostname is displayed properly.
- Synchronize time by using the Network Time Protocol (NTP).
- If you plan to deploy Sentinel with scalable storage configuration, ensure that you have installed CDH and Elasticsearch. For more information about deploying Sentinel with scalable storage, see [“Installing and Setting Up Scalable Storage”](#) on page 85.
- On RHEL systems:** For optimal performance, the memory settings must be set appropriately for the PostgreSQL database. The SHMMAX parameter must be greater than or equal to 1073741824.

To set the appropriate value, append the following information in the `/etc/sysctl.conf` file:

```
# for Sentinel Postgresql
kernel.shmmax=1073741824
```

- For traditional installations:**
 - ◆ The operating system for the Sentinel server must include at least the Base Server components of the SLES server or the RHEL server. Therefore, ensure that the following packages are installed before you install Sentinel:
 - ◆ `bc`
 - ◆ `bash`
 - ◆ `coreutils`
 - ◆ `gettext`
 - ◆ `glibc`
 - ◆ `grep`
 - ◆ `libgcc`
 - ◆ `libstdc`
 - ◆ `lsof`
 - ◆ `openssl`

- ♦ sed
- ♦ insserv
- ♦ net-tools
- ♦ zlib (up to SLES 12.x and RHEL 7.x, 8.x)
- ♦ python-libs (up to SLES 12.x and RHEL 7.x)
- ♦ netstat (up to SLES 12.x and RHEL 7.x) or ss (for SLES 15 and later)
- ♦ pam-modules (available only when you install Legacy-Module on SLES 15.x)

❑ **For Sentinel with traditional storage:**

To view event visualizations, set the virtual memory by adding the property `vm.max_map_count=262144` in the `/etc/sysctl.conf` file.

12 Installing and Configuring Elasticsearch

For scalable and distributed indexing of events you must install Elasticsearch in cluster mode. The Elasticsearch cluster you install for Sentinel must be used to index only Sentinel data.

- ♦ “Prerequisites” on page 75
- ♦ “Installing and Configuring Elasticsearch” on page 75
- ♦ “Securing Data in Elasticsearch” on page 78
- ♦ “Performance Tuning for Elasticsearch” on page 81
- ♦ “Redeploying Elasticsearch Security Plug-In” on page 82

Prerequisites

Complete the following prerequisite before you install the external Elasticsearch nodes:

- ♦ If you have installed the external nodes Elasticsearch 5.6.13 with Sentinel 8.3 or prior version, uninstall the Elasticsearch, and then install the Elasticsearch 7.7.0 version. For installation, see [Installing and Configuring Elasticsearch](#).
- ♦ Based on your EPS rate, deploy Elasticsearch in a cluster mode with the number of nodes and the number of replicas as recommended in [Sentinel System Requirements](#).
- ♦ Set the file descriptors by adding the following properties in the `/etc/security/limits.conf` file:

```
elasticsearch hard nofile 65536
elasticsearch soft nofile 65536
elasticsearch soft as unlimited
```

NOTE: After you complete the above prerequisites, run the `sysctl -p` command to reload the changes made to the files.

Installing and Configuring Elasticsearch

You must install Elasticsearch and the required plug-ins on each external node of the Elasticsearch cluster.

To install and configure Elasticsearch:

- 1 Install the JDK version supported by Elasticsearch.
- 2 Download the certified version of Elasticsearch RPM. For information about the certified version of Elasticsearch and the download URL, see the [Sentinel System Requirements](#) page.
- 3 Install Elasticsearch:

```
rpm -ivh elasticsearch-<version>.rpm
```

- 4 Complete the tasks as mentioned on-screen in the RPM post-installation instructions.
- 5 Ensure that the Elasticsearch user has access to Java.
- 6 Configure the `/etc/elasticsearch/elasticsearch.yml` file by updating or adding the following information:

Property and Value	Notes
discovery.seed_hosts: [<IP of the master eligible elasticsearch node in the cluster>,<IP of the master eligible elasticsearch node in the cluster>,<IP of the master eligible elasticsearch node in the cluster>, and so on]	
cluster.name: <Elasticsearch_cluster_name>	The cluster name that you specify must be same for all the nodes.
node.name: <node_name>	The node name must be unique for each node.
network.host: _<networkInterface>:ipv4_	
thread_pool.write.queue_size: 300	
thread_pool.search.queue_size: 10000	Once the search queue size reaches its limit, Elasticsearch discards any pending search requests in queue. You can increase the search queue size based on the below calculation: threadpool.search.queue_size = Average number of widget queries per user for a dashboard x number of shards (per day index) x number of days (search duration)
index.codec: best_compression	
path.data: ["/<es1>", "/<es2>"]	Spread data across multiple independent disks or locations to reduce the disk I/O latency. Configure multiple paths for storing Elasticsearch data. For example /es1, /es2, and so on. For best performance and manageability, mount each path to a separate physical disk (JBOD).

- 7 Update the default Elasticsearch heap size in the `<Sentinel_installation_path>/etc/elasticsearch/jvm.options` file.
The heap size must be 50% of the server memory. For example, on a 24 GB Elasticsearch node, allocate 12 GB as the heap size for optimal performance.
- 8 Repeat all of the above steps on each external Elasticsearch nodes of the Elasticsearch cluster.

- 9 In the Sentinel server Elasticsearch node, configure the `<Sentinel_installation_path>/opt/novell/sentinel/3drparty/elasticsearch/config/elasticsearch.yml` as follows:
- 9a Ensure that the values of `cluster.name` and `discovery.seed_hosts` in the `elasticsearch.yml` file are same as the `elasticsearch.yml` file in external Elasticsearch node.
 - 9b (Conditional) If you have not enabled the **Event Visualization** component, then specify the localhost IP address followed by the IP address of the local Elasticsearch node in the `network.host` property as follows:


```
network.host:
  ["127.0.0.1", "<IP_address_of_the_Elasticsearch_node_in_Sentinel>"]
```
- 10 (Conditional) For Sentinel with traditional storage, add the external Elasticsearch nodes IP addresses to the `ServerList` property in the `<Sentinel_installation_path>/etc/opt/novell/sentinel/config/elasticsearch-index.properties` file.
For example: `ServerList=<Elasticsearch_IP1>:<Port>,<Elasticsearch_IP2>:<Port>`
- 11 Restart Sentinel:
- ```
rcsentinel restart
```
- 12 Restart each external Elasticsearch node:
- ```
/etc/init.d/elasticsearch restart
```
- 13 Verify that the ES cluster is formed:
- ```
curl -X GET http://<Elasticsearch_IP_of_the_Sentinel_server>:<Port>/_cat/nodes?v
```
- 14 Ensure that all the existing alert data and event data (if available), it must be moved to the external Elasticsearch nodes.
- 15 For optimal performance and stability of the Sentinel server, configure the Elasticsearch node in the Sentinel server as a dedicated `master-eligible` node so that all the event visualization data is indexed in external Elasticsearch nodes:
- 15a Stop the internal node (Sentinel server)
 

```
rcsentinel stopES
```
  - 15b Set the following in the internal node:
 

```
node.master: true
node.data: false
node.ingest: false
```
  - 15c Run `elasticsearch-node repurpose` to clean all the shards
 

```
<Sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/bin/elasticsearch-node -v repurpose
```
  - 15d Start the internal Elasticsearch node
 

```
rcsentinel startES
```
  - 15e Restart each external Elasticsearch node:
 

```
/etc/init.d/elasticsearch restart
```
- 16 Proceed with [“Securing Data in Elasticsearch” on page 78.](#)

---

## IMPORTANT

Whenever an external Elasticsearch node goes down, the ES cluster restarts automatically, due to which, there might be a temporary issue in launching dashboards through Kibana and alert search.

When the Sentinel server is restarted ensure that you restart the external Elasticsearch nodes as well.

---

## Securing Data in Elasticsearch

Elasticsearch cluster nodes can be accessed by various clients such as the following:

- ◆ Sentinel: to fetch and present event data in the Event Visualization dashboard.
- ◆ Spark jobs running in the YARN NodeManager nodes: to perform bulk indexing of the events received from Kafka. (for SSDM)
- ◆ Collector Manager: to perform bulk indexing of events in Sentinel with traditional storage.
- ◆ Other external clients: to perform custom operations such as custom analytics.

Sentinel provides a security plug-in for Elasticsearch named **elasticsearch-security-plugin** that authenticates and authorizes access to Elasticsearch.

The plug-in uses either a SAML token or a whitelist for validation depending on how the clients connect:

- ◆ When a client sends a SAML token along with the request, the plug-in authenticates the token against the Sentinel authentication server. Upon successful authentication, the plug-in allows access only to the filtered events that the client is authorized for.

For example, the Event Visualization dashboard (client) displays only those events from Elasticsearch that a user's role is authorized to view.

For information about roles and permissions, see [“Creating a Role”](#) in the *Sentinel Administration Guide*.

- ◆ When a client cannot send a SAML token, the plug-in checks its 'whitelist of legitimate clients. Upon successful validation, the plug-in allows access to all events without filtering.
- ◆ When a client does not send a valid SAML token or is not allowed by the whitelist, the plug-in considers it as an illegitimate client and denies access to the client.

This section provides information about installing and configuring the Elasticsearch Security plug-in:

- ◆ [“Installing the Elasticsearch Security Plug-In”](#) on page 79
- ◆ [“Providing Secure Access to Additional Elasticsearch Clients”](#) on page 80
- ◆ [“Updating the Elasticsearch Plug-In Configuration”](#) on page 81

# Installing the Elasticsearch Security Plug-In

You must install the Elasticsearch security plug-in in each node of the Elasticsearch cluster and also in the Elasticsearch node included in Sentinel.

## To install the elasticsearch-security-plugin-in on the Elasticsearch node included in Sentinel:

- 1 Log in to the Sentinel main or SSDM server.
- 2 Set the path for the JAVA\_HOME environment variable as follows:

```
export JAVA_HOME=/<Sentinel_installation_path>/opt/novell/sentinel/jdk/
```

- 3 Install the plug-in:

**For Linux, log in as the user that Elasticsearch is running as and run the following command:**

```
<sentinel_installation_path>/opt/novell/sentinel/3rdparty/
elasticsearch/bin/elasticsearch-plugin install file://localhost/
<Sentinel_installation_path>/etc/opt/novell/sentinel/scalablestore/
elasticsearch-security-plugin*.zip --verbose
```

When prompted to continue with installation, enter *y*.

- 4 (Conditional) If Elasticsearch is not listening on the default HTTP port (9200), you must update the Elasticsearch port number in each entry of the *<Sentinel\_installation\_path>/opt/novell/sentinel/3rdparty/elasticsearch/plugins/elasticsearch-security-plugin/elasticsearch-ip-whitelist.txt* file.

For more information, see [“Providing Access to Elasticsearch Clients by Using Whitelist” on page 80](#).

- 5 Restart the indexing services in Sentinel using the command:

```
rcsentinel stopES
rcsentinel startES
```

## To install the elasticsearch-security-plugin-in on external Elasticsearch nodes:

Perform the following steps on each node in the Elasticsearch cluster:

- 1 Log in to the Sentinel main or SSDM server.
- 2 Copy the *<Sentinel\_installation\_path>/etc/opt/novell/sentinel/scalablestore/elasticsearch-security-plugin\*.zip* file to a temporary location on each node in the Elasticsearch cluster.

- 3 Install the plug-in:

**For Linux, log in as the user that Elasticsearch is running as and run the following command:**

```
<elasticsearch_install_directory>/bin/elasticsearch-plugin install
file://localhost/<full path of elasticsearch-security-plugin*.zip file>
--verbose
```

When prompted to continue with installation, enter *y*.

- 4 (Conditional) If Elasticsearch is not listening on the default HTTP port (9200), you must update the Elasticsearch port number in each entry of the *<elasticsearch\_install\_directory>/plugins/elasticsearch-security-plugin/elasticsearch-ip-whitelist.txt* file.

For more information, see [“Providing Access to Elasticsearch Clients by Using Whitelist” on page 80](#).

- 5 In the `elasticsearch.yml` file, add the property (`xpack.security.enabled: false`)
- 6 Restart Elasticsearch.

## Providing Secure Access to Additional Elasticsearch Clients

By default, trusted clients, such as SSDM server (for the Event Visualization Dashboard) and YARN NodeManagers, Sentinel server (for the Event Visualization Dashboard) and RCM have access to Elasticsearch. If you want to use additional Elasticsearch clients, you must provide secure access to those additional clients either by using SAML token or whitelist.

### Providing Access to Elasticsearch REST Clients by Using SAML Token

If you are using a REST client to access Elasticsearch, you can include a SAML token in the request header as follows:

- 1 Obtain a SAML token from the Sentinel authentication server. For more information, see the REST API documentation available in Sentinel.

Click [Help](#) > [APIs](#) > [Tutorial](#) > [API Security](#) > [Obtaining a SAML Token \(Logon\)](#).

- 2 Use the SAML token in the subsequent REST requests: include the SAML token in the Authorization header of each request made by the REST client. Specify the header name as `Authorization` and the header value as the `<SAML token>` obtained in Step 1.

### Providing Access to Elasticsearch Clients by Using Whitelist

By default, Sentinel auto-populates a whitelist with the IP addresses of the trusted Elasticsearch clients, such as the SSDM server (for the Event Visualization Dashboard) and YARN NodeManagers, Sentinel server (for the Event Visualization Dashboard) and RCM. The Elasticsearch security plug-in grants access to Elasticsearch for all the clients listed in its whitelist.

To provide access to additional clients that do not send a valid Sentinel token, you must add the IP address of the client and the HTTP port number of the Elasticsearch server to the whitelist in the `IP address:port` format. You must ensure that the external clients you add in the whitelist are legitimate and trustworthy to prevent any unauthorized access.

#### To update the whitelist:

- 1 Log in to the Sentinel server or Elasticsearch node as the user which Elasticsearch is running as.
- 2 Add the entry `<Elasticsearch_Client_IP>:<Target_Elasticsearch_HTTP_Port>` in the file:
  - ♦ `<Sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/plugins/elasticsearch-security-plugin//elasticsearch-ip-whitelist.txt` for Elasticsearch node included in Sentinel.
  - ♦ `<elasticsearch_install_directory>/plugins/elasticsearch-security-plugin/elasticsearch-ip-whitelist.txt` for external Elasticsearch nodes.

If there are multiple entries, add each entry in a new line and save the file.

- 3 Repeat the above steps in each node of the Elasticsearch cluster.



## Updating the Elasticsearch Plug-In Configuration

In cases where you modify the scalable storage components' IP address/hostname and port number or the Elasticsearch version and port number, you must update the Elasticsearch plug-in configuration files accordingly.

**Perform the following steps on each node of the Elasticsearch cluster:**

- 1 Log in to the Elasticsearch node as the user which Elasticsearch is running as.
- 2 (Conditional) If you modified YARN NodeManager IP addresses, SSDM or Sentinel server IP address, RCM IP addresses, or the Elasticsearch port number, update the whitelist accordingly to ensure that the Elasticsearch security plug-in grants access to the Elasticsearch clients.  
  
If you are configuring SSDM or Sentinel in HA mode, add entries for the physical IP address of each active node and passive node of the HA cluster.  
  
If you modify the physical IP address of any node of the HA cluster or add a new node to the HA cluster, update the whitelist with the physical IP addresses of the modified or the newly added nodes.  
  
For more information, see [“Providing Access to Elasticsearch Clients by Using Whitelist” on page 80](#).
- 3 (Conditional) If you modified the SSDM IP address, Sentinel server IP address or web server port number, update the `authServer.host` and `authServer.port` properties in the following files and restart Elasticsearch:
  - ♦ `<Sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/plugins/elasticsearch-security-plugin/plugin-configuration.properties` for Elasticsearch node included in Sentinel.
  - ♦ `<elasticsearch_install_directory>/plugins/elasticsearch-security-plugin/plugin-configuration.properties` for external Elasticsearch nodes.  
If you are configuring SSDM or Sentinel in HA mode, set the `authServer.host` property to the virtual IP address of the HA cluster.  
  
If you modify the virtual IP address of the HA cluster, update the `authServer.host` property to the modified virtual IP address.
- 4 (Conditional) If you upgraded Elasticsearch to a newer version, update the `elasticsearch.version` property in the following files and restart Elasticsearch:
  - ♦ `<Sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/plugins/elasticsearch-security-plugin/plugin-descriptor.properties` for Elasticsearch node included in Sentinel.
  - ♦ `<elasticsearch_install_directory>/plugins/elasticsearch-security-plugin/plugin-descriptor.properties` for external Elasticsearch nodes.

## Performance Tuning for Elasticsearch

Sentinel automatically configures the Elasticsearch settings described in the table below. You can customize the Elasticsearch settings as needed.

To customize the default settings:

**For traditional storage:** Open the `<Sentinel_installation_path>/etc/opt/novell/sentinel/config/elasticsearch-index.properties` file and update the properties listed in the table as required.

**For scalable storage:** In the SSDM home page, click **Storage > Scalable Storage > Advanced Properties > Elasticsearch**.

**Table 12-1** *Elasticsearch Properties*

| Property                                     | Default Value                                                                                                                                                                                                                                                                                                                        | Notes                                                                                                                                                                                                                            |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| elasticsearch.events.lucenefilter (Optional) |                                                                                                                                                                                                                                                                                                                                      | Specify a filter to send only specific events to Elasticsearch for indexing.<br><br>For example: If you specify the value as <code>sev:[3-5]</code> , events with severity value only between 3 and 5 are sent to Elasticsearch. |
| index.fields                                 | id,dt,rv171,msg,ei,evt,xdatastaxname,xdasoutcomename,sev,vul,rv32,rv39,rv159,dhn,dip,rv98,dp,fn,rv199,dun,tufname,rv84,rv158,shn,sip,rv76,sun,iufname,sp,iudep,rv198,rv62,st,tid,srcgeo,destgeo,obsgeo,rv145,estz,estzmonth,estzdiy,estzdim,estzdiw,estzhour,estzmin,rv24,tudep,pn,xdasclass,xdasid,xdasreg,xdasprov,iuident,tuident | Indicates the event fields that you want Elasticsearch to index.                                                                                                                                                                 |
| es.num.shards                                | 5                                                                                                                                                                                                                                                                                                                                    | Indicates the number of primary shards per index.<br><br>You can increase this default value when the shard size goes beyond 50 GB.                                                                                              |
| es.num.replicas                              | 1                                                                                                                                                                                                                                                                                                                                    | Indicates the number of replica shards that each primary shard should have.<br><br>A minimum of 2 node cluster is recommended considering failover and high availability.                                                        |

## Redeploying Elasticsearch Security Plug-In

You must redeploy; that is, uninstall and reinstall the Elasticsearch security plug-in in the Elasticsearch node included in Sentinel and external Elasticsearch nodes in the following scenarios:

- ◆ Adding or modifying remote Collector Manager IP addresses.
- ◆ Uninstalling remote Collector Managers.
- ◆ Enabling Scalable Storage post-installation.

To redeploy Elasticsearch security plug-in:

**1** Log in to the Sentinel server or Elasticsearch node as the user which Elasticsearch is running as.

**2** Uninstall the plug-in using the following command:

- ◆ For Elasticsearch included in Sentinel: `<Sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/bin/elasticsearch-plugin remove elasticsearch-security-plugin`
- ◆ For external Elasticsearch: `<elasticsearch_install_directory>/bin/elasticsearch-plugin remove elasticsearch-security-plugin`

**3** Reinstall the plug-in:

- ◆ For Elasticsearch included in Sentinel: `<sentinel_installation_path>/opt/novell/sentinel/3rdparty/elasticsearch/bin/elasticsearch-plugin install file://localhost/<Sentinel_installation_path>/etc/opt/novell/sentinel/scalablestore/elasticsearch-security-plugin*.zip --verbose`
- ◆ For external Elasticsearch: `<elasticsearch_install_directory>/bin/elasticsearch-plugin install file://localhost/<full path of elasticsearch-security-plugin*.zip file copied from the Sentinel machine> --verbose`

**4** Restart Elasticsearch using the following command:

- ◆ For the Elasticsearch node included in Sentinel:

```
rcsentinel stopES
rcsentinel startES
```

- ◆ For external Elasticsearch nodes:

```
sudo systemctl restart elasticsearch.service
```



# 13 Installing and Setting Up Scalable Storage

Complete the prerequisites listed in the following table to set up scalable storage as the data storage option for Sentinel:

**Table 13-1** Prerequisites to Enable Scalable Storage

| ❑ Tasks                                                                                                                                                                                                                                                  | See                                                                    |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| ❑ Determine the number of Hadoop distribution cluster and Elasticsearch cluster nodes you need to configure based on the EPS rate and number of replicas needed.                                                                                         | Sentinel System Requirements                                           |
| Determine the certified version of CDH and Elasticsearch.                                                                                                                                                                                                |                                                                        |
| ❑ CDH, Elasticsearch, and Sentinel have their own platform support matrix. Review the platform support matrix for each of these products and determine the platform you want to use.                                                                     | CDH support matrix in Cloudera documentation.                          |
| For Elasticsearch, RPM install is recommended because the RPM includes the init script. This will install Elasticsearch as a service and enables it to automatically stop and start during reboot and upgrades, and does not overwrite the config files. | Elasticsearch support matrix in Elasticsearch documentation.           |
| Elasticsearch RPM installation is not supported on SLES 11. Therefore, determine a suitable platform for Elasticsearch.                                                                                                                                  | <a href="#">Sentinel Support Matrix</a> .                              |
| ❑ Install and configure CDH in cluster mode.                                                                                                                                                                                                             | <a href="#">“Installing and Configuring CDH” on page 86.</a>           |
| ❑ Install and configure Elasticsearch in cluster mode.                                                                                                                                                                                                   | <a href="#">“Installing and Configuring Elasticsearch” on page 75.</a> |
| ❑ Enable scalable storage in Sentinel.                                                                                                                                                                                                                   | <a href="#">“Enabling Scalable Storage” on page 87</a>                 |

# Installing and Configuring CDH

This section provides information about the specific settings required for Sentinel when installing and configuring CDH. For detailed information about CDH installation and configuration, you must refer to the certified version of Cloudera documentation.

Sentinel works with Cloudera Express, the free edition of CDH. Sentinel also works with Cloudera Enterprise, which requires the purchase of a license from Cloudera and includes numerous capabilities not available in the Cloudera Express edition. If you choose to begin with Cloudera Express and later discover you need the capabilities available with Cloudera Enterprise, you can upgrade the cluster after purchasing the license from Cloudera.

- ◆ [“Prerequisites” on page 86](#)
- ◆ [“Installing and Configuring CDH” on page 87](#)

## Prerequisites

Before you install CDH, you must set up the hosts as per the following prerequisites:

- ◆ Complete the prerequisites mentioned in the [Cloudera documentation](#).
- ◆ Use ext4 or XFS file system for better performance.
- ◆ CDH needs a few operating system packages that do not get installed by default. Therefore, you must mount the respective operating system DVD. The Cloudera installation instructions guide you about the packages to install.
- ◆ For SLES operating systems, CDH requires the `python-psycopg2` package. Install the `python-psycopg2` package. For more information, see [openSUSE documentation](#).
- ◆ If you are using virtual machines, reserve the disk space required on the file system when you create virtual machines nodes. For example, in VMware, you can use thick provisioning.
- ◆ Ensure that Sentinel and CDH cluster nodes are in the same timezone.
- ◆ Set swappiness of all the hosts to 1 in the `/etc/sysctl.conf` file by adding the following entry:

```
vm.swappiness=1
```

To apply this setting immediately, run the following command:

```
sysctl -p
```

- ◆ The JDK version in CDH must be at least the same JDK version used in Sentinel. If the JDK version available in CDH is lower than the Sentinel JDK, you must follow the instructions to install the JDK manually versus installing the JDK available in the CDH repository.

Install JDK by using the archive binary file (`.tar.gz`) because the JDK RPM installation causes issues when using the `manage_spark_jobs.sh` script to submit Spark jobs on YARN.

To determine the JDK version used in Sentinel, see the [Sentinel Release Notes](#).

## Installing and Configuring CDH

Install the certified version of CDH. For information about the certified version of CDH, see the [Technical Information for Sentinel](#) page. Refer to the certified version of [Cloudera documentation](#) for installation instructions.

Perform the following while you install CDH:

- ♦ (Conditional) If the installation fails during embedded PostgreSQL database installation, perform the following steps:

```
mkdir -p /var/run/postgresql
```

```
sudo chown cloudera-scm:cloudera-scm /var/run/postgresql
```

- ♦ When choosing the software installation type in the **Select Repository** window, ensure that **Use Parcels** is selected and select Kafka in **Additional Parcels**.
- ♦ When you add services, ensure that you enable the following services:
  - ♦ Cloudera Manager
  - ♦ ZooKeeper
  - ♦ HDFS
  - ♦ HBase
  - ♦ YARN
  - ♦ Spark
  - ♦ Kafka

---

**NOTE:** The Spark history server and HDFS NameNode must be installed on the same node for system reliability. For information about the scalable storage architecture, see [“Planning for Scalable Storage” on page 43](#).

---

When enabling the above services, configure high availability for the following:

- ♦ HBase HMaster
- ♦ HDFS NameNode
- ♦ YARN ResourceManager
- ♦ (Conditional) If the installer does not deploy the client configuration due to missing Java path, open a new browser session and manually update the Java path as follows:

Click **Hosts > All Hosts > Configuration** and specify the correct path in the **Java Home Directory** field.

## Enabling Scalable Storage

You can enable scalable storage either during installation or post-installation of Sentinel. When you enable scalable storage during installation, Sentinel configures CDH components with default values. Some of these configurations are permanent and cannot be changed. For example, the default number of partitions for Kafka topics is 9 and this value cannot be changed.

If you want to change the default values, you must enable scalable storage after you install Sentinel and then set the configurations for CDH components as desired.

For traditional installations, you can enable scalable storage either during Sentinel installation or after Sentinel installation. For appliance installations, you can enable scalable storage only after installation.

In upgrade installations, you can enable scalable storage only after you upgrade Sentinel.

Before you proceed with enabling scalable storage, keep the list of IP addresses or hostnames and port numbers of Kafka, HDFS NameNode, YARN NodeManager, ZooKeeper, and Elasticsearch nodes handy. You need this information when you enable scalable storage.

To enable scalable storage during Sentinel installation, see [“Sentinel Server Custom Installation” on page 90](#).

To enable scalable storage after Sentinel installation or upgrade, see [“Enabling Scalable Storage Post-Installation”](#) in the *Sentinel Administration Guide*.



# 14 Traditional Installation

This chapter provides information about the various ways to install Sentinel.

- “Performing Interactive Installation” on page 89
- “Performing a Silent Installation” on page 96
- “Installing Sentinel as a Non-root User” on page 97

## Performing Interactive Installation

This section provides information about standard and custom installation.

- “Sentinel Server Standard Installation” on page 89
- “Sentinel Server Custom Installation” on page 90
- “Collector Manager and Correlation Engine Installation” on page 93

## Sentinel Server Standard Installation

Use the following steps to perform a standard installation:

- 1 Download the Sentinel installation file from the [Downloads website](#):
- 2 Specify at the command line the following command to extract the installation file.

```
tar zxvf <install_filename>
```

Replace *<install\_filename>* with the actual name of the install file.

- 3 Change to the directory where you extracted the installer:

```
cd <directory_name>
```

- 4 Specify the following command to install Sentinel:

```
./install-sentinel
```

or

If you want to install Sentinel on more than one system, you can record your installation options in a file. You can use this file for an unattended Sentinel installation on other systems. To record your installation options, specify the following command:

```
./install-sentinel -r <response_filename>
```

- 5 Specify the number for the language you want to use for the installation, then press Enter.  
The end user license agreement is displayed in the selected language.
- 6 Press the Spacebar to read through the license agreement.
- 7 Enter *yes* or *y* to accept the license and continue with the installation.

The installation might take a few seconds to load the installation packages and prompt for the configuration type.

- 8 When prompted, specify `1` to proceed with the standard configuration.

Installation proceeds with the default evaluation license key included with the installer. At any time during or after the evaluation period, you can replace the evaluation license with a license key you have purchased.

- 9 Specify the password for the administrator user `admin`.

- 10 Confirm the password again.

This password is used by `admin`, `dbauser`, and `appuser`.

The Sentinel installation finishes and the server starts. It might take few minutes for all services to start after installation because the system performs a one-time initialization. Wait until the installation finishes before you log in to the server.

To access the Sentinel Main interface, specify the following URL in your web browser:

```
https://IP_AddressOrDNS_Sentinel_server:8443/sentinel/views/main.html
```

Where `IP_AddressOrDNS_Sentinel_server` is the IP address or DNS name of the Sentinel server and `8443` is the default port for the Sentinel server.

## Sentinel Server Custom Installation

If you are installing Sentinel with a custom configuration, you can customize your Sentinel installation by specifying your license key, setting a different password, specifying different ports, and so on.

- 1 If you want to enable scalable storage, complete the prerequisites specified in [Chapter 13, “Installing and Setting Up Scalable Storage,”](#) on page 85.
- 2 Download the Sentinel installation file from the [Downloads website](#):
- 3 Specify at the command line the following command to extract the installation file.

```
tar zxvf <install_filename>
```

Replace `<install_filename>` with the actual name of the install file.

- 4 Specify the following command in the root of the extracted directory to install Sentinel:

```
./install-sentinel
```

or

If you want to use this custom configuration to install Sentinel on more than one system, you can record your installation options in a file. You can use this file for an unattended Sentinel installation on other systems. To record your installation options, specify the following command:

```
./install-sentinel -r <response_filename>
```

- 5 Specify the number for the language you want to use for the installation, then press Enter.  
The end user license agreement is displayed in the selected language.
- 6 Press the Spacebar to read through the license agreement.
- 7 Enter `yes` or `y` to accept the license agreement and continue with the installation.

The installation might take a few seconds to load the installation packages and prompt for the configuration type.

**8** Specify `2` to perform a custom configuration of Sentinel.

**9** Enter `1` to use the default evaluation license key

or

Enter `2` to enter a purchased license key for Sentinel.

**10** Specify the password for the administrator user `admin` and confirm the password again.

**11** Specify the password for the database user `dbauser` and confirm the password again.

The `dbauser` account is the identity used by Sentinel to interact with the database. The password you enter here can be used to perform database maintenance tasks, including resetting the admin password if the admin password is forgotten or lost.

**12** Specify the password for the application user `appuser` and confirm the password again.

**13** Change the port assignments for the Sentinel services by entering the desired number, then specifying the new port number.

**14** After you have changed the ports, specify `7` for done.

**15** Enter `1` to authenticate users using only the internal database.

or

If you have configured an LDAP directory in your domain, enter `2` to authenticate users by using LDAP directory authentication.

The default value is `1`.

**16** *If you want to enable Sentinel in FIPS 140-2 mode*, enter `y`.

**16a** Specify a strong password for the keystore database and confirm the password again.

---

**NOTE:** The password must be at least seven characters long. The password must contain at least three of the following character classes: Digits, ASCII lowercase letters, ASCII uppercase letters, ASCII non-alphanumeric characters, and non-ASCII characters.

If an ASCII uppercase letter is the first character or a digit is the last character, they are not counted.

---

**16b** If you want to insert external certificates into the keystore database to establish trust, press `y` and specify the path for the certificate file. Otherwise, press `n`

**16c** Complete the FIPS 140-2 mode configuration by following the tasks mentioned in [Chapter 24, “Operating Sentinel in FIPS 140-2 Mode,”](#) on page 131.

**17** *If you want to enable scalable storage*, enter `yes` or `y` to enable scalable storage.

---

**IMPORTANT:** Once you enable scalable storage, you cannot revert the configuration unless you re-install Sentinel.

---

**17a** Specify the IP addresses or hostnames and port numbers of the scalable storage components.

**17b** (Conditional) If you want to exit scalable storage configuration and proceed with Sentinel installation, enter `no` or `n`.

**17c** After the Sentinel installation is done, complete the scalable storage configuration mentioned in the section [“Post-Installation Configuration for Scalable Storage”](#) on page 92.

The Sentinel installation finishes and the server starts. It might take few minutes for all services to start after installation because the system performs a one-time initialization. Wait until the installation finishes before you log in to the server.

To access the Sentinel Main interface, specify the following URL in your web browser:

```
https://IP_AddressOrDNS_Sentinel_server:8443/sentinel/views/main.html
```

Where `<IP_AddressOrDNS_Sentinel_server>` is the IP address or DNS name of the Sentinel server and `8443` is the default port for the Sentinel server.

## Post-Installation Configuration for Scalable Storage

- 1 Log in to the SSDM server.
- 2 Clear your browser cache to view the Sentinel version you installed.
- 3 To view events and alerts, add the Elasticsearch node included in SSDM to the Elasticsearch cluster you have setup for scalable storage:

In the local Elasticsearch node, open `<Sentinel_installation_path>/opt/novell/sentinel/3drparty/elasticsearch/config/elasticsearch.yml` file and add the following information:

- ◆ `cluster.name: <Elasticsearch_cluster_name>`
- ◆ `node.name: <node_name>`
- ◆ `discovery.zen.ping.unicast.hosts:["<FQDN of elasticsearch node1>", "<FQDN of elasticsearch node2>", and so on"]`

In all the external Elasticsearch nodes, open `/<Sentinel_installation_path>/etc/elasticsearch/elasticsearch.yml` and update

```
discovery.zen.ping.unicast.hosts:["<FQDN of elasticsearch node1>", "<FQDN of elasticsearch node2>", and so on"]
```

---

**NOTE:** Ensure that the values of the parameters in the local `elasticsearch.yml` file and the `elasticsearch.yml` file in external Elasticsearch nodes are same except `network.host` and `node.name` as these values are unique to the node.

---

- 4 Restart the indexing services using the command:

```
rcsentinel stopES
rcsentinel startES
```

- 5 Complete the scalable storage configuration as mentioned in the following sections:
  - ◆ [“Securing Data in Elasticsearch” on page 78](#)
  - ◆ [Performance Tuning Guidelines](#) in the *Sentinel Administration Guide*
  - ◆ [Processing Data](#) in the *Sentinel Administration Guide*

# Collector Manager and Correlation Engine Installation

By default, Sentinel installs a Collector Manager and a Correlation Engine. For production environments, set up a distributed deployment because it isolates data collection components on a separate machine, which is important for handling spikes and other anomalies with maximum system stability. For information about the advantages of installing additional components, see [“Advantages of Distributed Deployments” on page 46](#).

You can install more than one Collector Manager or Correlation Engine.

---

**IMPORTANT:** You must install the additional Collector Manager or the Correlation Engine on separate systems. The Collector Manager or the Correlation Engine must not be on the same system where the Sentinel server is installed.

---

You can record the installation parameters during interactive installation and then use the recorded files for an unattended installation on other systems. You can specify the following files to record the installation:

- ◆ *<Response\_file>*: Records the installation parameters you specify during installation.
- ◆ *<Configuration\_file>*: Specify this file only if you have multiple Sentinel servers. You can use this file to connect Collector Manager and Correlation Engine to a different Sentinel server than the one recorded in the response file. During interactive installation, it creates placeholders for Sentinel server details. You can later update this file with the relevant Sentinel server details and use it along with the response file during unattended installation.

---

**NOTE:** This option is available only in Sentinel 8.2 SP3 or later.

---

**Installation Checklist:** Ensure that you have completed the following tasks before starting the installation.

- ◆ Make sure that your hardware and software meet the minimum requirements. For more information, see [Chapter 5, “Meeting System Requirements,” on page 37](#).
- ◆ Synchronize time by using the Network Time Protocol (NTP).
- ◆ A Collector Manager requires network connectivity to the message bus port (61616) on the Sentinel server. Before you start installing the Collector Manager, make sure that all firewall and network settings are allowed to communicate over this port.

## To install the Collector manager and the Correlation engine:

- 1 Launch the Sentinel Main interface by specifying the following URL in your web browser:

```
https://IP_AddressOrDNS_Sentinel_server:8443/sentinel/views/main.html
```

Where *<IP\_AddressOrDNS\_Sentinel\_server>* is the IP address or DNS name of the Sentinel server and *8443* is the default port for the Sentinel server.

Log in with the username and password specified during the installation of the Sentinel server.

- 2 In the toolbar, click **Downloads**.
- 3 Click **Download Installer** under the required installation.
- 4 Click **Save File** to save the installer to the desired location.

- 5 Specify the following command to extract the installation file.

```
tar zxvf <install_filename>
```

Replace *<install\_filename>* with the actual name of the install file.

- 6 Change to the directory where you extracted the installer.
- 7 (Conditional) To install without recording the installation, specify the following command:

- ♦ **For Collector Manager:**

```
./install-cm
```

- ♦ **For Correlation Engine:**

```
./install-ce
```

- 8 (Conditional) To install and record the installation, do one of the following:

- ♦ (Conditional) If you have a single Sentinel server, specify the following command:

- ♦ **For Collector Manager:**

```
./install-cm -r <response_filename>
```

- ♦ **For Correlation Engine:**

```
./install-ce -r <response_filename>
```

- ♦ (Conditional) If you have multiple Sentinel servers, specify the following command:

- ♦ **For Collector Manager:**

```
./install-cm -r <response_filename> -c <configuration_filename>
```

- ♦ **For Correlation Engine:**

```
./install-ce -r <response_filename> -c <configuration_filename>
```

For more information about using the response file or the configuration file, see [“Performing a Silent Installation” on page 96](#).

- 9 Specify the number for the language you want to use for the installation.  
The end user license agreement is displayed in the selected language.
- 10 Press the Spacebar to read through the license agreement.
- 11 Enter *yes* or *y* to accept the license agreement and continue with the installation.  
The installation might take a few seconds to load the installation packages and prompt for the configuration type.
- 12 When prompted, specify the appropriate option to proceed with the Standard or Custom configuration.
- 13 Enter the default Communication Server Hostname or IP Address of the machine on which Sentinel is installed.
- 14 (Conditional) If you chose Custom configuration, specify the following:
  - 14a Sentinel server communication channel port number.
  - 14b Sentinel Web server port number.

- 15 When prompted to accept the certificate, run the following command in the Sentinel server to verify the certificate:

For FIPS mode:

```
<Sentinel_installation_path>/opt/novell/sentinel/jdk/jre/bin/keytool -
list -keystore
<Sentinel_installation_path>/etc/opt/novell/sentinel/config/
.activemqkeystore.jks
```

For Non-FIPS mode:

```
<Sentinel_installation_path>/opt/novell/sentinel/jdk/jre/bin/keytool -
list -keystore
<Sentinel_installation_path>/etc/opt/novell/sentinel/config/
nonfips_backup/.activemqkeystore.jks
```

Compare the certificate output with the Sentinel server certificate displayed in [Step 13](#).

---

**NOTE:** If the certificate does not match, the installation stops. Run the installation setup again and check the certificates.

---

- 16 Accept the certificate if the certificate output matches the Sentinel server certificate.
- 17 Specify credentials of any user in Administrator role. Enter the user name and the password.
- 18 (Conditional) If you chose Custom configuration, enter *yes* or *y* to enable FIPS 140-2 mode in Sentinel and continue with the FIPS configuration.
- 19 (Conditional) If your environment uses multi-factor or strong authentication, you must provide the Sentinel client id and Sentinel client secret. For more information about authentication methods, see “[Authentication Methods](#)” in the *Sentinel Administrator Guide*.

To retrieve the Sentinel client ID and Sentinel client secret, go to the following URL:

```
https://Hostname:port/SentinelAuthServices/oauth/clients
```

Where:

- ◆ *Hostname* is the host name of the Sentinel server.
- ◆ *Port* is the port Sentinel uses (typically 8443).

The specified URL uses your current Sentinel session to retrieve the Sentinel client ID and Sentinel client secret.

- 20 (Conditional) If you have enabled Event Visualization, you must add the Collector Manager to the Elasticsearch whitelist. For more information, see “[Providing Access to Elasticsearch Clients by Using Whitelist](#)” on page 80.
- 21 Continue with the installation as prompted until the installation is complete.

# Performing a Silent Installation

The silent or unattended installation is useful if you need to install more than one Sentinel server, Collector Manager, or Correlation Engine in your deployment. You can record the installation parameters during the interactive installation and then run the recorded files on other systems.

- ◆ Ensure that you have recorded the installation parameters to a file. For information about creating the response file, see:
  - ◆ “Sentinel Server Standard Installation” on page 89
  - ◆ “Sentinel Server Custom Installation” on page 90
  - ◆ “Collector Manager and Correlation Engine Installation” on page 93.

---

**NOTE:** For Collector Manager and Correlation Engine, use the configuration file to connect Collector Manager and Correlation Engine to a different Sentinel server than the one recorded in the response file. Update this file with the relevant Sentinel server details and use it along with the response file during unattended installation.

---

**To enable FIPS 140-2 mode, ensure that the response file includes the following parameters:**

- ◆ ENABLE\_FIPS\_MODE
- ◆ NSS\_DB\_PASSWORD

**To perform a silent installation:**

- 1 Download the installation files from the [Downloads website](#).
- 2 Log in as `root` to the server where you want to install Sentinel or Collector manager or Correlation engine.
- 3 Specify the following command to extract the install files from the tar file:

```
tar -zxvf <install_filename>
```

Replace `<install_filename>` with the actual name of the install file.

- 4 (Conditional) To install Sentinel server in silent mode, specify the following command:

```
./install-sentinel -u <response_filename>
```

The installation proceeds with the values stored in the response file.

If you installed a Sentinel server, it might take few minutes for all services to start after installation, because the system performs a one-time initialization. Wait until the installation finishes before you log in to the server.

- 5 (Conditional) To install Collector Managers, specify the following command:

- ◆ To use the response file:

```
./install-cm -u <response_filename>
```

- ◆ To use the response file and the configuration file:

```
./install-cm -u <response_filename> -i <configuration_filename>
```

- 6 (Conditional) To install Correlation Engines, specify the following command:

- ◆ To use the response file:



```
./install-ce -u <response_file>
```

- ◆ To use the response file and the configuration file:

```
./install-ce -u <response_filename> -i <configuration_filename>
```

- 7 (Conditional) If you chose to enable FIPS 140-2 mode for the Sentinel server**, complete the FIPS 140-2 mode configuration by following the tasks mentioned in [Chapter 24, “Operating Sentinel in FIPS 140-2 Mode,”](#) on page 131.

## Installing Sentinel as a Non-root User

If your organizational policy does not allow you to run the full installation of Sentinel as `root`, you can install Sentinel as a `non-root` user; that is, as the `novell` user. In this installation, a few steps are performed as a `root` user, then you proceed to install Sentinel as the `novell` user created by the `root` user. Finally, the `root` user completes the installation.

When installing Sentinel as a `non-root` user, you should install Sentinel as the `novell` user. Non-root installations other than the `novell` user is not supported, although the installation proceeds successfully.

---

**NOTE:** When installing Sentinel in an already existing, non-default directory, ensure that the `novell` user has ownership permissions to the directory. Run the following command to assign ownership permissions:

```
chown novell:novell <non-default installation directory>
```

---

- 1 Download the installation files from the [Downloads website](#).
- 2 Specify the following command at the command line to extract the install files from the tar file:

```
tar -zxvf <install_filename>
```

Replace `<install_filename>` with the actual name of the install file.

- 3 Log in as `root` to the server where you want to install Sentinel as `root`.
- 4 Specify the following command:

```
./bin/root_install_prepare
```

A list of commands to be executed with root privileges is displayed. If you want the non-root user to install Sentinel in non-default location, specify the `--location` option along with the command. For example:

```
./bin/root_install_prepare --location=/foo
```

The value that you pass to the `--location` option `foo` is prepended to the directory paths.

This also creates a `novell` group and a `novell` user, if they do not already exist.

- 5 Accept the command list.  
The displayed commands are executed.
- 6 Specify the following command to change to the newly created non-root user; that is, `novell`:

```
su novell
```

7 (Conditional) To do an interactive installation:

7a Specify the appropriate command depending on the component you are installing:

---

| Component          | Command                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Sentinel server    | <b>Default location:</b> <code>./install-sentinel</code><br><b>Non-default location:</b> <code>./install-sentinel --location=/foo</code> |
| Collector Manager  | <b>Default location:</b> <code>./install-cm</code><br><b>Non-default location:</b> <code>./install-cm --location=/foo</code>             |
| Correlation Engine | <b>Default location:</b> <code>./install-ce</code><br><b>Non-default location:</b> <code>./install-cm --location=/foo</code>             |

---

7b Continue with [Step 10](#).

8 (Conditional) To perform silent installation of Sentinel Server, ensure that you have recorded the installation parameters to a file. For information on creating the response file, see [“Sentinel Server Standard Installation” on page 89](#) or [“Sentinel Server Custom Installation” on page 90](#).

8a To install, specify the following command:

**Default location:** `./install-sentinel -u <response_filename>`

**Non-default location:** `./install-sentinel --location=/foo -u <response_filename>`

8b Continue with [Step 13](#).

9 (Conditional) To perform silent installation of Collector Manager or Correlation Engine, ensure that you have recorded the installation parameters to a file.

---

**NOTE:** Use the configuration file to connect Collector Manager and Correlation Engine to a different Sentinel server than the one recorded in the response file. Update this file with the relevant Sentinel server details and use it along with the response file during unattended installation.

---

For information on creating the response file or the configuration file, see [“Collector Manager and Correlation Engine Installation” on page 93](#).

9a Specify the appropriate command depending on the component you are installing:

| Component          | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Collector Manager  | <ul style="list-style-type: none"> <li>◆ To use the response file: <ul style="list-style-type: none"> <li>◆ <b>Default location:</b> <code>./install-cm -u &lt;response_filename&gt;</code></li> <li>◆ <b>Non-default location:</b> <code>./install-cm --location=/foo -u &lt;response_filename&gt;</code></li> </ul> </li> <li>◆ To use the response file and the configuration file: <ul style="list-style-type: none"> <li>◆ <b>Default location:</b> <pre>./install-cm -u &lt;response_filename&gt; -i &lt;configuration_filename&gt;</pre> </li> <li>◆ <b>Non-default location:</b> <pre>./install-cm --location=/foo -u &lt;response_filename&gt; -i &lt;configuration_filename&gt;</pre> </li> </ul> </li> </ul> <p>The installation proceeds with the Sentinel server values from the configuration file and the other installation parameter values stored in the response file.</p> |
| Correlation Engine | <ul style="list-style-type: none"> <li>◆ To use the response file: <ul style="list-style-type: none"> <li>◆ <b>Default location:</b> <code>./install-ce -u &lt;response_filename&gt;</code></li> <li>◆ <b>Non-default location:</b> <code>./install-ce --location=/foo -u &lt;response_filename&gt;</code></li> </ul> </li> <li>◆ To use the response file and the configuration file: <ul style="list-style-type: none"> <li>◆ <b>Default location:</b> <pre>./install-ce -u &lt;response_filename&gt; -i &lt;configuration_filename&gt;</pre> </li> <li>◆ <b>Non-default location:</b> <pre>./install-ce --location=/foo -u &lt;response_filename&gt; -i &lt;configuration_filename&gt;</pre> </li> </ul> </li> </ul> <p>The installation proceeds with the Sentinel server values from the configuration file and the other installation parameter values stored in the response file.</p> |

**9b** Continue with [Step 13](#).

**10** Specify the number for the language you want to use for the installation.

The end user license agreement is displayed in the selected language.

**11** Read the end user license and enter `yes` or `y` to accept the license and continue with the installation.

The installation starts installing all RPM packages. This installation might take a few seconds to complete.

**12** You are prompted to specify the mode of installation.

- ◆ If you select to proceed with the standard configuration, continue with [Step 8](#) through [Step 10](#) in “Sentinel Server Standard Installation” on page 89.
- ◆ If you select to proceed with the custom configuration, continue with [Step 8](#) through [Step 15](#) in “Sentinel Server Custom Installation” on page 90.

**13** Log in as a `root` user and specify the following command to finish installation:

```
./bin/root_install_finish
```

The Sentinel installation finishes and the server starts. It might take few minutes for all services to start after installation because the system performs a one-time initialization. Wait until the installation finishes before you log in to the server.

To access the Sentinel Main interface, specify the following URL in your web browser:

```
https://IP_AddressOrDNS_Sentinel_server:8443/sentinel/views/main.html
```

Where `IP_AddressOrDNS_Sentinel_server` is the IP address or DNS name of the Sentinel server and `8443` is the default port for the Sentinel server.

# 15 Appliance Installation

The Sentinel appliance is a ready-to-run software appliance based on the Micro Focus common appliance framework. The appliance combines a hardened SLES operating system and the Sentinel software integrated update service to provide an easy and seamless user experience that allows you to leverage existing investments. The Sentinel appliance provides a Web-based user interface to configure and monitor the appliance.

Depending on the Sentinel version, the appliance installer installs the certified SLES operating system:

- ◆ For 8.2, the appliance installer installs SLES 12 SP3.
- ◆ For 8.2 SP2, the appliance installer installs SLES 12 SP4.

The Sentinel appliance image is packaged in both ISO and OVF formats that can be deployed to the virtual environments. For information about supported virtualization platforms, see [Sentinel System Requirements](#).

- ◆ [“Prerequisites” on page 101](#)
- ◆ [“Installing the Sentinel ISO Appliance” on page 102](#)
- ◆ [“Installing the Sentinel OVF Appliance” on page 104](#)
- ◆ [“Post-Installation Configuration for the Appliance” on page 106](#)

## Prerequisites

Ensure that the environment where you are going to install Sentinel as ISO appliance meets the following prerequisites:

- ◆ Before you install the Sentinel appliance, review new functionality and known issues in the certified SLES [Release Notes](#).
- ◆ (Conditional) If you are installing Sentinel ISO appliance on bare metal hardware, download the appliance ISO disk image from the support site, and make a DVD.
- ◆ Ensure that the minimum hard disk space is 50 GB for the installer to make the automatic partition proposal.
- ◆ Ensure that your system has a minimum memory of 4 GB for the installation to complete. If the memory is less than 4 GB, the installation fails. If the memory is more than 4 GB but less than the recommended size 24 GB, the installation displays a message that you have less memory than is recommended.

# Installing the Sentinel ISO Appliance

This section provides information about installing Sentinel, Collector Managers, and Correlation Engines using the ISO appliance image. This image format allows you to generate a full disk image format that can be deployed directly to hardware, either physical (bare metal) or virtual (uninstalled virtual machine in a hypervisor) by using a bootable ISO DVD image.

- ◆ “Installing Sentinel” on page 102
- ◆ “Installing Collector Managers and Correlation Engines” on page 103

## Installing Sentinel

**To install the Sentinel ISO appliance:**

- 1 Download the ISO virtual appliance image from the [Download Website](#).
- 2 (Conditional) If you are using a hypervisor:  
Set up the virtual machine using the ISO virtual appliance image and power it on.  
or  
Burn the ISO image into a DVD, set up the virtual machine using the DVD, and then power it on.
- 3 (Conditional) If you are installing the Sentinel appliance on bare metal hardware:
  - 3a Boot the physical machine from the DVD drive with the DVD.
  - 3b Follow the installation wizard on-screen instructions.
  - 3c Select **Install sentinel server <version>**.
- 4 Select the language of your choice.
- 5 Select the keyboard layout.
- 6 Click **Next**.
- 7 Read and accept the SUSE Enterprise Server Software License Agreement. Click **Next**
- 8 Read and accept the Sentinel Server Appliance License Agreement. Click **Next**
- 9 Set the Sentinel appliance passwords, NTP configuration, and the time zone.  
Set `vaadmin` user credentials for logging on to Sentinel Appliance Management Console.

---

**NOTE:** After installation, you can change the NTP configuration and time zone in the following ways:

- ◆ Go to the command prompt and enter `yast->Network Services->NTP Configuration`
- ◆ Go to Sentinel Appliance Management Console and click **Time**.

If the time appears out of sync immediately after the install, run the following command to restart NTP:

```
rcntp restart
```

---

- 10 On the Sentinel Server Appliance Network Settings page, specify the hostname and domain name. Select either **Static IP Address** or **DHCP IP Address**.
- 11 Click **Next**.

- 12 (Conditional) If you have selected **Static IP Address** in Step 10, specify the network connection settings.
- 13 Click **Next**.
- 14 Set the password for Sentinel user `admin`, then click **Next**.  
Appliance is installed.
- 15 Make a note of the appliance IP address that is shown in the console.
- 16 Log in as `root` user at the console to log in to the appliance.  
Enter the username as `root` and enter the password you set in [Step 9](#).
- 17 Proceed with “[Post-Installation Configuration for the Appliance](#)” on page 106.

## Installing Collector Managers and Correlation Engines

The procedure to install a Collector Manager or a Correlation Engine is similar to the procedure to installing Sentinel except that you need to download the appropriate ISO appliance file from the [Download website](#).

- 1 Complete Step 1 through Step 13 in “[Installing Sentinel](#)” on page 102.

The installation checks for the available memory and disk space. If the available memory is less than 1 GB, the installation will not let you proceed and the **Next** button is greyed out.

- 2 Specify the following configuration for the Collector Manager or the Correlation Engine:

- ◆ **Sentinel Server Hostname or IP Address:** Specify the host name or IP address of the Sentinel server that the Collector Manager or Correlation Engine should connect to.
- ◆ **Sentinel Communication Channel Port:** Specify the Sentinel server communication channel port number. The default port number is 61616.
- ◆ **Sentinel Web Server Port:** Specify the Sentinel web server port. The default port is 8443.
- ◆ **User name with Administrator role:** Specify username of any user in Administrator role.
- ◆ **Password for user with Administrator role:** Specify the password for the user name you have specified in the above field.

- 3 (Conditional) If your environment uses multi-factor or strong authentication, you must provide the Sentinel client id and Sentinel client secret. For more information about authentication methods, see “[Authentication Methods](#)” in the *Sentinel Administrator Guide*.

To retrieve the Sentinel client ID and Sentinel client secret, go to the following URL:

```
https://Hostname:port/SentinelAuthServices/oauth/clients
```

Where:

- ◆ *Hostname* is the host name of the Sentinel server.
- ◆ *Port* is the port Sentinel uses (typically 8443).

The specified URL uses your current Sentinel session to retrieve the Sentinel client ID and Sentinel client secret.

- 4 Click **Next**.
- 5 Accept the certificate when prompted.
- 6 Make a note of the appliance IP address that is shown in the console.

The console displays a message that this appliance is the Sentinel Collector Manager or Correlation Engine depending on what you chose to install, along with the IP address. The Console also displays the Sentinel server user interface IP address.

- 7 Complete [Step 16](#) through [Step 17](#) in “[Installing Sentinel](#)” on [page 102](#).

## Installing the Sentinel OVF Appliance

This section provides information about installing Sentinel, Collector Manager, and Correlation Engine as an OVF appliance image.

The OVF format is a standard virtual machine format that is supported by most hypervisors, either directly or by a simple conversion. Sentinel supports OVF appliance with two certified hypervisors, but you can also use it with other hypervisors.

- ♦ “[Installing Sentinel](#)” on [page 104](#)
- ♦ “[Installing Collector Managers and Correlation Engines](#)” on [page 105](#)

### Installing Sentinel

To install the Sentinel OVF appliance:

- 1 Download the OVF virtual appliance image from the [Download Website](#).
- 2 In your hypervisor's management console, import the OVF image file as a new virtual machine. Allow the hypervisor to convert the OVF image into the native format if you are prompted to do so.
- 3 Review the virtual hardware resources allocated to your new virtual machine to ensure that they meet the Sentinel requirements.
- 4 Power on the virtual machine.
- 5 Select the language of your choice.
- 6 Select the keyboard layout.
- 7 Click **Next**.
- 8 Read and accept the SUSE Enterprise Server Software License Agreement. Click **Next**.
- 9 Read and accept the Sentinel Server Appliance License Agreement. Click **Next**.
- 10 Set the Sentinel appliance passwords, NTP configuration and the time zone.  
Set `vaadmin` user credentials for logging on to Sentinel Appliance Management Console.

---

**NOTE:** After installation, you can change the NTP configuration and time zone in the following ways:

- ♦ Go to the command prompt and enter `yast->Network Services->NTP Configuration`
- ♦ Go to Sentinel Appliance Management Console and click **Time**.

If the time appears out of sync immediately after the install, run the following command to restart NTP:



```
rcntp restart
```

---

- 11 On the Sentinel Server Appliance Network Settings page, specify the hostname and domain name. Select either **Static IP Address** or **DHCP IP Address**.
- 12 Click **Next**.
- 13 (Conditional), If you have selected **Static IP Address** in Step 11, specify the network connection settings.
- 14 Click **Next**.
- 15 Set the Sentinel admin password, then click **Next**.  
It might take a few minutes for all services to start after installation because the system performs a one-time initialization. Wait until the installation finishes before you log in to the server.
- 16 Make a note of the appliance IP address that is shown in the console. Use the same IP address to access the Sentinel Main interface.

## Installing Collector Managers and Correlation Engines

To install a Collector Manager or a Correlation Engine on a VMware ESX server as an OVF appliance image:

- 1 Complete Step 1 through Step 14 in [“Installing Sentinel” on page 104](#).  
The installation checks for the available memory and disk space. If the available memory is less than 1 GB, the installation will not let you proceed and the **Next** button is greyed out.
- 2 Specify the host name/IP address of the Sentinel server that the Collector Manager should connect to.
- 3 Specify the Communication Server port number. The default port is 61616.
- 4 Specify credentials of any user in Administrator role. Enter the user name and the password.
- 5 (Conditional) If your environment uses multi-factor or strong authentication, you must provide the Sentinel client id and Sentinel client secret. For more information about authentication methods, see [“Authentication Methods”](#) in the *Sentinel Administrator Guide*.

To retrieve the Sentinel client ID and Sentinel client secret, go to the following URL:

```
https://Hostname:port/SentinelAuthServices/oauth/clients
```

Where:

- ♦ *Hostname* is the host name of the Sentinel server.
- ♦ *Port* is the port Sentinel uses (typically 8443).

The specified URL uses your current Sentinel session to retrieve the Sentinel client ID and Sentinel client secret.

- 6 Click **Next**.
- 7 Accept the certificate.

- 8 Click **Next** to complete the installation.

When the installation is complete, the installer displays a message indicating that this appliance is the Sentinel Collector Manager or the Sentinel Correlation Engine depending on what you chose to install, along with the IP address. It also displays the Sentinel server user interface IP address.

## Post-Installation Configuration for the Appliance

After you install Sentinel, you need to perform additional configuration for the appliance to work properly.

- ♦ [“Registering for Updates” on page 106](#)
- ♦ [“Creating Partitions for Traditional Storage” on page 107](#)
- ♦ [“Configuring Scalable Storage” on page 108](#)
- ♦ [“Configuring the Appliance with SMT” on page 108](#)

### Registering for Updates

You must register the Sentinel appliance with the appliance update channel to receive Sentinel and latest operating system updates. To register the appliance, you must first obtain your appliance registration code or the appliance activation key from the [Customer Care Center](#).

Based on the operating system you install, you can register for updates in the following ways:

- ♦ If you are using SLES 12 SP3 or later, you can register using the Sentinel appliance management console.
- ♦ If you are using SLES 12 SP3 or later, you can register using commands.
- ♦ [“Register Using Sentinel Appliance Management Console” on page 106](#)
- ♦ [“Register Using Commands” on page 107](#)

### Register Using Sentinel Appliance Management Console

To register using the Sentinel appliance management console:

- 1 Launch Sentinel appliance by doing either of the following:
  - ♦ Log in to Sentinel. Click **Sentinel Main** > **Appliance**.
  - ♦ Specify the following URL in your web browser: `https://<IP_address>:9443`.
- 2 Log in as either `vaadmin` or `root` user.
- 3 Click **Online Update** > **Register Now**.
- 4 In the **Email** field, specify the email ID to which you want to receive updates.
- 5 In the **Activation Key** field, enter the registration code.
- 6 Click **Register** to complete the registration.

## Register Using Commands

To register using commands:

- 1 Log in to the Sentinel server as the `root` user.
- 2 Specify the following commands:
  - ♦ To register server, specify: `suse_register -a regcode-sentinel="<registration_code>" -a email="<email_ID>"`
  - ♦ To register Collector Manager, specify: `suse_register -a regcode-sentinel-collector="<registration_code>" -a email="<email_ID>"`
  - ♦ To register Correlation Engine, specify: `suse_register -a regcode-sentinel-correlation ="<registration_code>" -a email="<email_ID>"`
  - ♦ To register Sentinel in high availability, specify: `suse_register -a regcode-sentinel-ha ="<registration_code>" -a email="<email_ID>"`

For the email parameter, specify the email ID to which you want to receive updates.

## Creating Partitions for Traditional Storage

The information in this section is applicable only if you want to use traditional storage as the data storage option.

As a best practice, ensure that you create separate partitions to store Sentinel data on a different partition than the executables, configuration, and operating system files. The benefits of storing variable data separately include easier backup of sets of files, simpler recovery in case of corruption, and provides additional robustness if a disk partition fills up. For information about planning your partitions, see [“Planning for Traditional Storage” on page 40](#). You can add partitions in the appliance and move a directory to the new partition by using the YaST tool.

Use the following procedure to create a new partition and move the data files from its directory to the newly created partition:

- 1 Log in to Sentinel as `root`.
- 2 Run the following command to stop the Sentinel on the appliance:

```
/etc/init.d/sentinel stop
```
- 3 Specify the following command to change to `novell` user:

```
su - novell
```
- 4 Move the contents of the directory at `/var/opt/novell/sentinel` to a temporary location.
- 5 Change to `root` user.
- 6 Enter the following command to access the YaST2 Control Center:

```
yast
```
- 7 Select **System > Partitioner**.
- 8 Read the warning and select **Yes** to add the new unused partition.  
For information about creating partitions, see [Using the YaST Partitioner](#) in the *SLES 11 documentation*.
- 9 Mount the new partition at `/var/opt/novell/sentinel`.

10 Specify the following command to change to `novell` user:

```
su - novell
```

11 Move the contents of the data directory from the temporary location (where it was saved in [Step 4](#)) back to `/var/opt/novell/sentinel` in the new partition.

12 Run the following command to restart the Sentinel appliance:

```
/etc/init.d/sentinel start
```

## Configuring Scalable Storage

To enable and configure scalable storage as the data storage option, see [“Configuring Scalable Storage”](#) in the *Sentinel Administration Guide*.

## Configuring the Appliance with SMT

In secured environments where the appliance must run without direct Internet access, you can configure the appliance with the Subscription Management Tool (SMT), which enables you to upgrade the appliance to the latest versions of Sentinel as they are released. SMT is a package proxy system that is integrated with Customer Center and provides key Customer Center capabilities.

- ♦ [“Prerequisites” on page 108](#)
- ♦ [“Configuring the Appliance” on page 109](#)
- ♦ [“Upgrading the Appliance” on page 109](#)

## Prerequisites

Before you configure the appliance with SMT, ensure that you meet the following prerequisites:

- ♦ Get the Customer Center credentials to get Sentinel updates. For more information about getting the credentials, contact [Technical Support](#).
- ♦ Ensure that SLES 11 SP3 is installed with the following packages on the computer where you want to install SMT:
  - ♦ `htmlDoc`
  - ♦ `perl-DBIx-Transaction`
  - ♦ `perl-File-Basename-Object`
  - ♦ `perl-DBIx-Migration-Director`
  - ♦ `perl-MIME-Lite`
  - ♦ `perl-Text-ASCIITable`
  - ♦ `yum-metadata-parser`
  - ♦ `createrepo`
  - ♦ `perl-DBI`
  - ♦ `apache2-prefork`
  - ♦ `libapr1`
  - ♦ `perl-Data-ShowTable`

- ◆ perl-Net-Daemon
- ◆ perl-Tie-IxHash
- ◆ ftik
- ◆ libapr-util1
- ◆ perl-PIRPC
- ◆ apache2-mod\_perl
- ◆ apache2-utils
- ◆ apache2
- ◆ perl-DBD-mysql
- ◆ Install SMT and configure the SMT server. For more information, see the following sections in the [SMT documentation](#):
  - ◆ SMT Installation
  - ◆ SMT Server Configuration
  - ◆ Mirroring Installation and Update Repositories with SMT
- ◆ Install the `wget` utility on the appliance computer.

## Configuring the Appliance

Perform the following steps to configure the appliance with SMT:

- 1 Enable the appliance repositories by running the following commands in the SMT server:
 

```
smt-repos -e Sentinel-Server-8-OS-Updates sle-12-x86_64
smt-repos -e Sentinel-Server-8-Prod-Updates sle-12-x86_64
smt-repos -e Sentinel-Collector-Manager-8-OS-Updates sle-12-x86_64
smt-repos -e Sentinel-Collector-Manager-8-Prod-Updates sle-12-x86_64
smt-repos -e Sentinel-Correlation-Engine-8-OS-Updates sle-12-x86_64
smt-repos -e Sentinel-Correlation-Engine-8-Prod-Updates sle-12-x86_64
```
- 2 Configure the appliance with SMT by performing the steps in the “[Configuring Clients to Use SMT](#)” section of the [SMT documentation](#).

## Upgrading the Appliance

For information about upgrading the appliance, see “[Upgrading the Sentinel Appliance](#)” on [page 163](#).



# 16 Installing Additional Collectors and Connectors

By default, all released Collectors and Connectors are installed when you install Sentinel. If you want to install a new Collector or Connector released after the Sentinel release, use the information in the following sections.

- ♦ “Installing a Collector” on page 111
- ♦ “Installing a Connector” on page 111

## Installing a Collector

Use the following steps to install a Collector:

- 1 Download the desired Collector from the [Sentinel Plug-ins website](#).
- 2 From **Sentinel Main**, click the **admin** drop-down, then click **Applications**.
- 3 Click **Launch Control Center** to launch the Sentinel Control Center.
- 4 In the toolbar, click **Event Source Management > Live View**, then click **Tools > Import plugin**.
- 5 Browse to and select the Collector file you downloaded in [Step 1](#), then click **Next**.
- 6 Follow the remaining prompts, then click **Finish**.

To configure the Collector, see the documentation for the specific Collector on the [Sentinel Plug-ins website](#).

## Installing a Connector

Use the following steps to install a Connector:

- 1 Download the desired Connector from the [Sentinel Plug-ins website](#).
- 2 From **Sentinel Main**, click the **admin** drop-down, then click **Applications**.
- 3 Click **Launch Control Center** to launch the Sentinel Control Center.
- 4 In the toolbar, select **Event Source Management > Live View**, then click **Tools > Import plugin**.
- 5 Browse to and select the Connector file you downloaded in [Step 1](#), then click **Next**.
- 6 Follow the remaining prompts, then click **Finish**.

To configure the Connector, see the documentation for the specific Connector on the [Sentinel Plug-ins website](#).





# 17 Verifying the Installation

You can determine whether the installation is successful by performing either of the following:

- ◆ Verify the Sentinel version:

```
/etc/init.d/sentinel version
```

- ◆ Verify whether the Sentinel services are up and running and functioning in FIPS or Non-FIPS mode:

```
/etc/init.d/sentinel status
```

- ◆ Verify whether the Web services are up and running:

```
netstat -an |grep 'LISTEN' |grep <HTTPS_port_number>
```

The default port number is 8443.

- ◆ Launch Sentinel:

1. Launch a supported web browser.
2. Specify the URL of Sentinel:

```
https://IP_AddressOrDNS_Sentinel_server:8443
```

Where *IP\_AddressOrDNS\_Sentinel\_server* is the IP address or DNS name of the Sentinel server and *8443* is the default port for the Sentinel server.

3. Log in with the administrator name and password specified during the installation. The default username is admin.



# IV Configuring Sentinel

This section provides information about configuring Sentinel and the out-of-the-box plug-ins.

- ♦ [Chapter 18, “Configuring Time,” on page 117](#)
- ♦ [Chapter 19, “Securing Data in Elasticsearch,” on page 121](#)
- ♦ [Chapter 20, “Enabling Event Visualization,” on page 123](#)
- ♦ [Chapter 21, “Modifying the Configuration after Installation,” on page 125](#)
- ♦ [Chapter 22, “Configuring Out-of-the-Box Plug-Ins,” on page 127](#)
- ♦ [Chapter 23, “Enabling FIPS 140-2 Mode in an Existing Sentinel Installation,” on page 129](#)
- ♦ [Chapter 24, “Operating Sentinel in FIPS 140-2 Mode,” on page 131](#)
- ♦ [Chapter 25, “Adding a Consent Banner,” on page 143](#)
- ♦ [Chapter 26, “Limiting the Number of Concurrent Active Sessions,” on page 145](#)
- ♦ [Chapter 27, “Ending Inactive Sessions,” on page 147](#)
- ♦ [Chapter 28, “Configuring IP Flow Data Collection,” on page 149](#)



# 18 Configuring Time

The time of an event is very critical to its processing in Sentinel. It is important for reporting and auditing purposes as well as for real-time processing. This section provides information about understanding time in Sentinel, how to configure time, and handling time zones.

- ♦ [“Understanding Time in Sentinel” on page 117](#)
- ♦ [“Configuring Time in Sentinel” on page 119](#)
- ♦ [“Configuring Delay Time Limit for Events” on page 119](#)
- ♦ [“Handling Time Zones” on page 119](#)

## Understanding Time in Sentinel

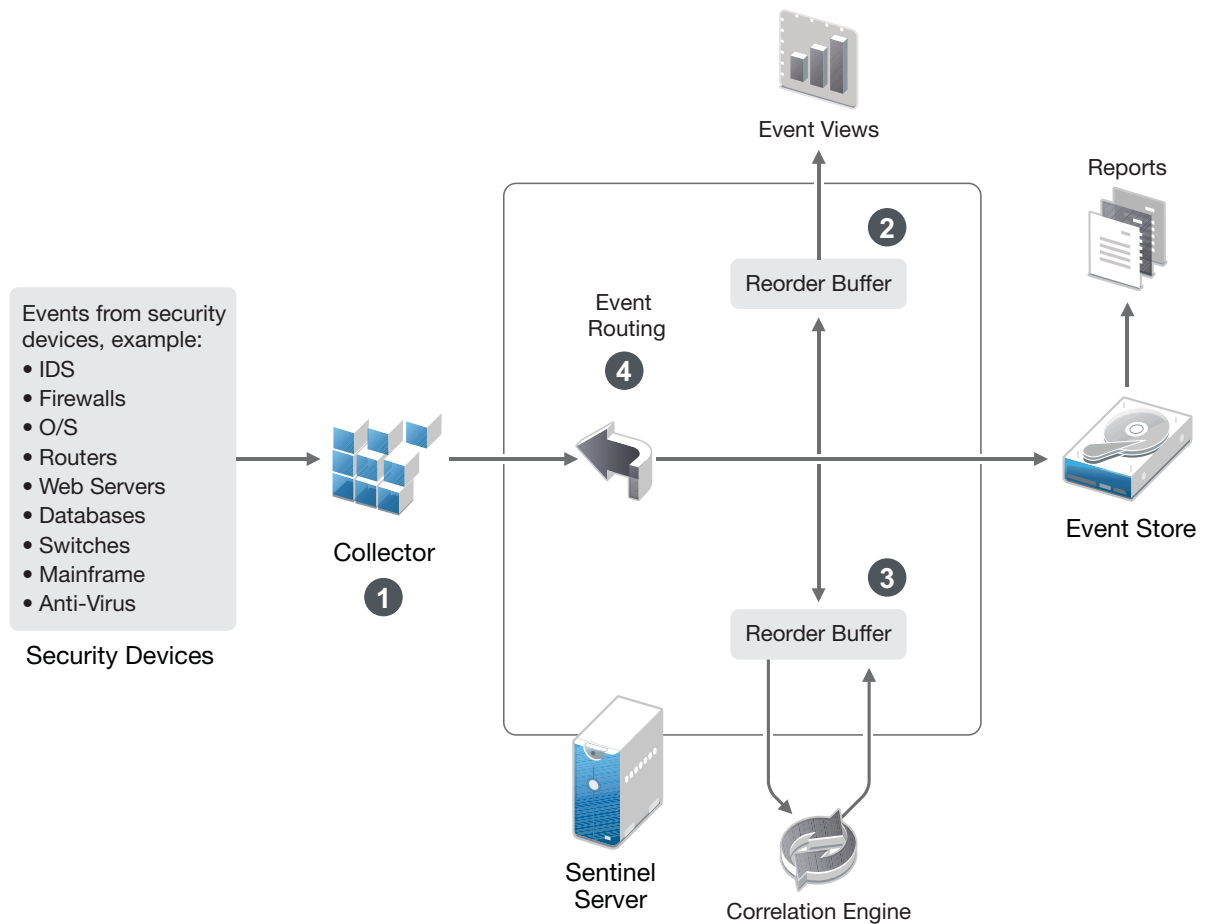
Sentinel is a distributed system that is made up of several processes distributed through out your network. In addition, there can be some delay introduced by the event source. To accommodate this, the Sentinel processes reorder events into a time-ordered stream before processing.

Every event has three time fields:

- ♦ **Event Time:** This is the event time used by all analytical engines, searches, reports, and so on.
- ♦ **Sentinel Process Time:** The time Sentinel collected the data from the device, which is taken from the Collector Manager system time.
- ♦ **Observer Event Time:** The time stamp the device put in the data. The data might not always contain a reliable time stamp and can be quite different than the Sentinel Process Time. For example, when the device delivers data in batches.

The following illustration explains how Sentinel does this in a traditional storage setup:

Figure 18-1 Sentinel Time



1. By default, the Event Time is set to the Sentinel Process Time. The ideal, however, is for the Event Time to match the Observer Event Time, if it is available and trustworthy. It is best to configure data collection to **Trust Event Source Time** if the device time is available, accurate, and properly parsed by the Collector. The Collector sets the Event Time to match the Observer Event Time.
2. The events that have an Event Time within a 5 minute range from the server time (in the past or future) are processed normally by Event Views. Events that have an Event Time more than 5 minutes in the future do not show in the Event Views, but are inserted into the event store. Events that have an Event Time more than 5 minutes in the future and less than 24 hours in the past still are shown in the charts, but are not shown in the event data for that chart. A drill-down operation is necessary to retrieve those events from the event store.
3. Events are sorted into 30-second intervals so that the Correlation Engine can process them in chronological order. If the Event Time is more than 30 seconds older than the server time, the Correlation Engine does not process the events.
4. If the Event Time is older than 5 minutes relative to the Collector Manager system time, Sentinel directly routes events to the event store, bypassing real-time systems such as Correlation Engine and Security Intelligence.

## Configuring Time in Sentinel

The Correlation Engine processes time-ordered streams of events and detects patterns within events as well as temporal patterns in the stream. However, sometimes the device generating the event might not include the time in its log messages.

To configure time to work correctly with Sentinel, you have two options:

- ◆ Configure NTP on the Collector Manager and deselect **Trust Event Source Time** on the event source in the Event Source Manager. Sentinel uses the Collector Manager as the time source for the events.
- ◆ Select **Trust Event Source Time** on the event source in Event Source Manager. Sentinel uses the time from the log message as the correct time.

To change this setting on the event source:

- 1 Log in to Event Source Management.  
For more information, see “[Accessing Event Source Management](#)” in the *Sentinel Administration Guide*.
- 2 Right-click the event source you want to change the time setting for, then select **Edit**.
- 3 Select or deselect the **Trust Event Source** option on the bottom of the **General** tab.
- 4 Click **OK** to save the change.

## Configuring Delay Time Limit for Events

When Sentinel receives events from event sources, there may be a delay between the time the event was generated and the time Sentinel processes it. Sentinel stores the events with large delays in separate partitions. If many events are delayed over a long period of time, it may be an indicator of an incorrectly configured event source. This might also decrease the Sentinel performance as it attempts to handle the delayed events. Since the delayed events may be the result of a misconfiguration and, therefore, may not be desirable to store, Sentinel allows you to configure the acceptable delay limit for the incoming events. The event router drops the events that exceed the delay limit. Specify the delay limit in the following property in the `configuration.properties` file:

```
esecurity.router.event.delayacceptthreshold = <time in milliseconds>
```

You can also have a listing periodically logged to the Sentinel server log file showing the event sources from which events are received that are delayed beyond a specified threshold. To log this information, specify the threshold in the following property in the `configuration.properties` file:

```
sentinel.indexedlog.eventdelay.reportthreshold= <time in milliseconds>
```

## Handling Time Zones

Handling time zones can become very complex in a distributed environment. For example, you might have an event source in one time zone, the Collector Manager in another, the back-end Sentinel server in another, and the client viewing the data in yet another. When you add concerns such as

daylight saving time and the many event sources that don't report what time zone they are set to (such as all syslog sources), there are many possible problems that need to be handled. Sentinel is flexible so that you can properly represent the time when events actually occur, and compare those events to other events from other sources in the same or different time zones.

In general, there are three different scenarios for how event sources report time stamps:

- ♦ The event source reports the time in UTC. For example, all standard Windows Event Log events are always reported in UTC.
- ♦ The event source reports in local time, but always includes the time zone in the time stamp. For example, any event source that follows RFC3339 in structuring time stamps include the time zone as an offset; other sources report long time zone IDs such as Americas/New York, or short time zone IDs such as EST, which can present problems because of conflicts and inadequate resolutions.
- ♦ The event source reports local time, but does not indicate the time zone. Unfortunately, the extremely common syslog format follows this model.

For the first scenario, you can always calculate the absolute UTC time that an event occurred (assuming that a time sync protocol is in use), so you can easily compare the timing of that event to any other event source in the world. However, you cannot automatically determine what the local time was when the event occurred. For this reason, Sentinel allows customers to manually set the time zone of an event source by editing the Event Source node in the Event Source Manager and specifying the appropriate time zone. This information does not affect the calculation of DeviceEventTime or EventTime, but is placed into the ObserverTZ field, and is used to calculate the various ObserverTZ fields, such as ObserverTZHour. These fields are always expressed in local time.

In the second scenario, if the long-form time zone IDs or offsets are used, you can convert to UTC to get the absolute canonical UTC time (stored in DeviceEventTime), but you can also calculate the local time ObserverTZ fields. If a short-form time zone ID is used, there is some potential for conflicts.

The third scenario requires the administrator to manually set the event source time zone for all affected sources so that Sentinel can properly calculate the UTC time. If the time zone is not properly specified by editing the Event Source node in the Event Source Manager, then the DeviceEventTime (and probably the EventTime) can be incorrect; also, the ObserverTZ and associated fields might be incorrect.

In general, the Collector for a given type of event source (such as Microsoft Windows) knows how an event source presents time stamps, and adjusts accordingly. It is always good policy to manually set the time zone for all Event Source nodes in the Event Source Manager, unless you know that the event source reports in local time and always includes the time zone in the time stamp

Processing the event source presentation of the time stamp happens in the Collector and on the Collector Manager. The DeviceEventTime and the EventTime are stored as UTC, and the ObserverTZ fields are stored as strings set to local time for the event source. This information is sent from the Collector Manager to the Sentinel server and stored in the event store. The time zone that the Collector Manager and the Sentinel server are in should not affect this process or the stored data. However, when a client views the event in a web browser, the UTC EventTime is converted to the local time according to the web browser, so all events are presented to clients in the local time zone. If the users want to see the local time of the source, they can examine the ObserverTZ fields for details.



# 19 Securing Data in Elasticsearch

Sentinel leverages Kibana, a browser-based analytics and search dashboard, which helps you to visualize events and alerts in dashboards. Sentinel stores and indexes alerts in Elasticsearch. You can configure Sentinel to also store and index events in Elasticsearch to leverage event visualization capabilities. Sentinel dashboards access data from Elasticsearch to present events and alerts in dashboards. To ensure that the dashboards display only the data that a user's role is authorized to view and to prevent unauthorized data access in Elasticsearch, you must install the Elasticsearch Security plug-in. For more information, see [“Securing Data in Elasticsearch” on page 78](#).



# 20 Enabling Event Visualization

In a scalable storage setup, event visualizations are available by default. In a traditional storage set up, event visualizations are available only if you enabled the visualization data store (Elasticsearch) to store and index data.

- ♦ [“Prerequisite” on page 123](#)
- ♦ [“Enabling Event Visualization” on page 123](#)

## Prerequisite

For scalable and distributed indexing of events in production environments, you must set up additional Elasticsearch nodes in a cluster mode. To install and configure Elasticsearch in a cluster mode, see [“Installing and Configuring Elasticsearch” on page 75](#).

## Enabling Event Visualization

To enable event visualization:

- 1 Log in to the Sentinel server as the `novell` user.
- 2 Open the `/etc/opt/novell/sentinel/config/configuration.properties` file.
- 3 (Conditional) If you are using Sentinel in High Availability (HA) mode, ensure that the `sentinel.ha.cluster` property is set to `true` for all the nodes on the cluster.
- 4 Set the `eventvisualization.traditionalstorage.enabled` to `true`.
- 5 Refresh the user interface after few minutes to view event visualizations.

You should now see all the dashboards enabled in the **My Sentinel** user interface. Launch any dashboard, the Threat Hunting dashboard for example, and click **Search**. The dashboard displays all the events generated in the last 1 hour.

- 6 (Optional) Event visualization dashboards display only the events processed after you enabled event visualization. To view existing events present in file-based storage, you must migrate data from file-based storage to Elasticsearch. For more information, see [Chapter 37, “Migrating Data to Elasticsearch,” on page 193](#).

---

**NOTE:** Enabling or disabling event visualization generates an exception, as it restarts Sentinel indexing services. This exception is expected and you can ignore this exception.

---



# 21 Modifying the Configuration after Installation

After installing Sentinel, if you want to enter the valid license key, change the password, or modify any of the assigned ports, you can run the `configure.sh` script to modify them. The script is available in the `/opt/novell/sentinel/setup` folder.

- 1 Shut down Sentinel using the following command:

```
rcsentinel stop
```

- 2 Specify the following command at the command line to run the `configure.sh` script:

```
./configure.sh
```

- 3 Specify 1 to perform a standard configuration or specify 2 to perform a custom configuration of Sentinel.
- 4 Press the Spacebar to read through the license agreement.
- 5 Enter `yes` or `y` to accept the license agreement and continue with the installation.

The installation might take a few seconds to load the installation packages.

- 6 Enter 1 to use the default evaluation license key

or

Enter 2 to enter a purchased license key for Sentinel.

- 7 Decide whether you want to keep the existing password for the `admin` administrator user.

- ♦ If you want to keep the existing password, enter 1, then continue with [Step 8](#).
- ♦ If you want to change the existing password, enter 2, specify the new password, confirm the password, then continue with [Step 8](#).

The `admin` user is the identity used to perform administration tasks through the Sentinel Main interface, including the creation of other user accounts.

- 8 Decide whether you want to keep the existing password for the `dbauser` database user.

- ♦ If you want to keep the existing password, enter 1, then continue with [Step 9](#).
- ♦ If you want to change the existing password, enter 2, specify the new password, confirm the password, then continue with [Step 9](#).

The `dbauser` account is the identity that Sentinel uses to interact with the database. The password you enter here can be used to perform database maintenance tasks, including resetting the admin password if the admin password is forgotten or lost.

- 9 Decide whether you want to keep the existing password for the `appuser` application user.

- ♦ If you want to keep the existing password, enter 1, then continue with [Step 10](#).

- ♦ If you want to change the existing password, enter 2, specify the new password, confirm the password, then continue with [Step 10](#).

The `appuser` account is an internal identity, which Sentinel java process uses to establish connection and interact with the database. The password you enter here is used to perform database tasks.

**10** Change the port assignments for the Sentinel services by entering the desired number, then specifying the new port number.

**11** After you have changed the ports, specify 7 for done.

**12** Enter 1 to authenticate users using only the internal database.

or

If you have configured an LDAP directory in your domain, enter 2 to authenticate users by using LDAP directory authentication.

The default value is 1.

# 22 Configuring Out-of-the-Box Plug-Ins

Sentinel is preinstalled with the default Sentinel plug-ins available at the time of the Sentinel release.

This chapter provides information about how to configure the out-of-the-box plug-ins.

- ♦ [“Viewing the Preinstalled Plug-Ins” on page 127](#)
- ♦ [“Configuring Data Collection” on page 127](#)
- ♦ [“Configuring Solution Packs” on page 127](#)
- ♦ [“Configuring Actions and Integrators” on page 128](#)

## Viewing the Preinstalled Plug-Ins

You can see the list of plug-ins preinstalled in Sentinel. You can also see the plug-ins versions and other metadata, which helps you determine whether you have the latest version of a plug-in.

**To view the plug-ins installed in your Sentinel server:**

- 1 Log in as an administrator to the Sentinel Main interface at `https://<IP address>:8443`, where 8443 is the default port for the Sentinel server.
- 2 Click **Plug-ins > Catalog**.

## Configuring Data Collection

For information about configuring Sentinel for data collection, see [“Collecting and Routing Event Data”](#) in the *Sentinel Administration Guide*.

## Configuring Solution Packs

Sentinel ships with a wide variety of useful out-of-the-box content that you can use immediately to meet many of your analysis needs. Much of this content comes from the pre-installed Sentinel Core Solution Pack and Solution Pack for ISO 27000 Series. For more information, see [“Using Solution Packs”](#) in the *Sentinel Administration Guide*.

Solution Packs allow categorization and grouping of content into controls or policy sets that are treated as a unit. The controls in the Solution Packs are pre-installed to provide you with this out-of-the-box content, but you have to formally implement or test those controls by using the Sentinel Main interface.

If a certain amount of rigor is desired to help show that your Sentinel implementation is working as designed, you may use the formal attestation process built into the Solution Packs. This attestation process implements and tests the Solution Pack controls just as you would implement and test

controls from any other Solution Pack. As part of this process, the implementer and tester will attest that they have completed their work; these attestations will then become part of an audit trail that can be examined to demonstrate that any particular control was properly deployed.

You can do the attestation process by using the Solution Manager. For more information on implementing and testing the controls, see [“Installing and Managing Solution Packs”](#) in the *Sentinel Administration Guide*.

## Configuring Actions and Integrators

For information about configuring the out-of-the-box plug-ins, see the specific plug-in documentation available on the [Sentinel Plug-ins website](#).



# 23 Enabling FIPS 140-2 Mode in an Existing Sentinel Installation

This chapter provides information about enabling FIPS 140-2 mode in an existing installation of Sentinel.

---

**NOTE:** These instructions assume that Sentinel is installed at the `/opt/novell/sentinel` directory. The commands must be executed as the `novell` user.

---

- ♦ [“Enabling Sentinel Server to Run in FIPS 140-2 Mode” on page 129](#)
- ♦ [“Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines” on page 130](#)

## Enabling Sentinel Server to Run in FIPS 140-2 Mode

To enable the Sentinel Server to run in FIPS 140-2 mode:

- 1 Log in to the Sentinel server.
- 2 Switch to `novell` user (`su novell`).
- 3 Browse to the Sentinel bin directory.
- 4 Run the `convert_to_fips.sh` script and follow the on-screen instructions.
- 5 (Conditional) If your environment uses multi-factor or strong authentication, you must run the `create_mfa_fips_keys.sh` script and follow the on-screen instructions.

---

**NOTE:** While the script is running, it requires the password for the nss database.

---

- 6 (Conditional) If your environment uses multi-factor or strong authentication, you must provide the Sentinel client id and Sentinel client secret. For more information about authentication methods, see [“Authentication Methods”](#) in the *Sentinel Administrator Guide*.

To retrieve the Sentinel client ID and Sentinel client secret, go to the following URL:

`https://Hostname:port/SentinelAuthServices/oauth/clients`

Where:

- ♦ *Hostname* is the host name of the Sentinel server.
- ♦ *Port* is the port Sentinel uses (typically 8443).

The specified URL uses your current Sentinel session to retrieve the Sentinel client ID and Sentinel client secret.

- 7 Restart the Sentinel server.
- 8 Complete the FIPS 140-2 mode configuration by following the tasks mentioned in [Chapter 24, “Operating Sentinel in FIPS 140-2 Mode,” on page 131](#).

# Enabling FIPS 140-2 Mode on Remote Collector Managers and Correlation Engines

You must enable FIPS 140-2 mode on the remote Collector Manager and Correlation Engine if you want to use FIPS-approved communications with the Sentinel server running in FIPS 140-2 mode.

## To enable a remote Collector Manager or Correlation Engine to run in FIPS 140-2 mode:

- 1 Login to the remote Collector Manager or Correlation Engine system.
- 2 Switch to `novell` user (`su novell`).
- 3 Browse to the `bin` directory. The default location is `/opt/novell/sentinel/bin`.
- 4 Run the `convert_to_fips.sh` script and follow the on-screen instructions.
- 5 Restart the Collector Manager or Correlation Engine.
- 6 Complete the FIPS 140-2 mode configuration by following the tasks mentioned in [Chapter 24, "Operating Sentinel in FIPS 140-2 Mode,"](#) on page 131.

# 24 Operating Sentinel in FIPS 140-2 Mode

This chapter provides information about configuring and operating Sentinel in FIPS 140-2 mode.

- ♦ [“Configuring Distributed Search in FIPS 140-2 Mode” on page 131](#)
- ♦ [“Configuring LDAP Authentication in FIPS 140-2 Mode” on page 132](#)
- ♦ [“Updating Server Certificates in Remote Collector Managers and Correlation Engines” on page 133](#)
- ♦ [“Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode” on page 133](#)
- ♦ [“Importing Certificates into FIPS Keystore Database” on page 140](#)
- ♦ [“Reverting Sentinel to Non-FIPS Mode” on page 141](#)

## Configuring Distributed Search in FIPS 140-2 Mode

This section provides information about configuring distributed search in FIPS 140-2 mode.

### Scenario 1: Both the source and the target Sentinel servers are in FIPS 140-2 mode

To allow distributed searches across multiple Sentinel servers running in FIPS 140-2 mode, you need to add the certificates used for secure communication to the FIPS keystore.

- 1 Log in to the distributed search source computer.
- 2 Browse to the certificate directory:

```
cd <sentinel_install_directory>/config
```

- 3 Copy the source certificate (`sentinel.cer`) to a temporary location on the target computer.
- 4 Import the source certificate into the target Sentinel FIPS keystore.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

- 5 Log in to the distributed search target computer.
- 6 Browse to the certificate directory:

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

- 7 Copy the target certificate (`sentinel.cer`) to a temporary location on the source computer.
- 8 Import the target system certificate into the source Sentinel FIPS keystore.
- 9 Restart the Sentinel services on both the source and target computer.

### Scenario 2: The source Sentinel server is in non-FIPS mode and the target Sentinel server is in FIPS 140-2 mode

You must convert the Web server keystore on the source computer to the certificate format and then export the certificate to the target computer.

- 1 Log in to the distributed search source computer.
- 2 Create the Web server keystore in certificate (.cer) format:

```
<sentinel_install_directory>/jdk/jre/bin/keytool -export -alias
webserver -keystore <sentinel_install_directory>/config/
.webserverkeystore.jks -storepass password -file <certificate_name.cer>
```

- 3 Copy the distributed search source certificate (Sentinel.cer) to a temporary location on the distributed search target computer.
- 4 Log in to the distributed search target computer.
- 5 Import the source certificate into the target Sentinel FIPS keystore.  
For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140.](#)
- 6 Restart Sentinel services on the target computer.

### Scenario 3: The source Sentinel server is in FIPS mode and the target Sentinel server is in non-FIPS mode

- 1 Log in to the distributed search target computer.
- 2 Create the Web server keystore in certificate (.cer) format:

```
<sentinel_install_directory>/jdk/jre/bin/keytool -export -alias
webserver -keystore <sentinel_install_directory>/config/
.webserverkeystore.jks -storepass password -file <certificate_name.cer>
```

- 3 Copy the certificate to a temporary location on the distributed search source computer.
- 4 Import the target certificate into the source Sentinel FIPS keystore.  
For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140.](#)
- 5 Restart the Sentinel services on the source computer.

## Configuring LDAP Authentication in FIPS 140-2 Mode

To configure LDAP authentication for Sentinel servers running in FIPS 140-2 mode:

- 1 Get the LDAP server certificate from the LDAP administrator, or you can use a command. For example,

```
openssl s_client -connect <LDAP server IP>:636
```

and then copy the text returned (between but not including the BEGIN and END lines) into a file.

- 2 Import the LDAP server certificate into the Sentinel FIPS keystore.  
For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140.](#)
- 3 Navigate to the **Sentinel Main** interface as a user in the administrator role and proceed with configuring LDAP authentication.

For more information, see [“LDAP Authentication Against a Single LDAP Server Or Domain”](#) in the *Sentinel Administration Guide*.

---

**NOTE:** You can also configure LDAP authentication for a Sentinel server running in FIPS 140-2 mode by running the `ldap_auth_config.sh` script in the `/opt/novell/sentinel/setup` directory.

---

## Updating Server Certificates in Remote Collector Managers and Correlation Engines

To configure existing remote Collector Managers and Correlation Engines to communicate with a Sentinel server running in FIPS 140-2 Mode, you can either convert the remote system in FIPS 140-2 mode or you can update the Sentinel server certificate to the remote system and leave the Collector Manager or Correlation Engine in non-FIPS mode. Remote Collector Managers in FIPS mode may not work with event sources that do not support FIPS or that require one of the Sentinel Connectors that are not yet FIPS-enabled.

If you do not plan to enable FIPS 140-2 mode on the remote Collector Manager or Correlation Engine, you must copy the latest Sentinel server certificate to the remote system, so that the Collector Manager or Correlation Engine, can communicate with the Sentinel server.

To update the Sentinel server certificate in the remote Collector Manager or Correlation Engine:

- 1 Log in to the remote Collector Manager or Correlation Engine computer.
- 2 Switch to `novell` user (`su novell`).
- 3 Browse to the `bin` directory. The default location is `/opt/novell/sentinel/bin`.
- 4 Run the `updateServerCert.sh` script and follow the on-screen instructions.

## Configuring Sentinel Plug-Ins to Run in FIPS 140-2 Mode

This section provides information about configuring various Sentinel plug-ins to run in FIPS 140-2 mode.

---

**NOTE:** These instructions are provided assuming that you have installed Sentinel at the `/opt/novell/sentinel` directory. Run all the commands as `novell` user.

---

- ♦ [“Agent Manager Connector” on page 134](#)
- ♦ [“Database \(JDBC\) Connector” on page 135](#)
- ♦ [“Sentinel Link Connector” on page 135](#)
- ♦ [“Syslog Connector” on page 136](#)
- ♦ [“Windows Event \(WMI\) Connector” on page 137](#)
- ♦ [“Sentinel Link Integrator” on page 138](#)
- ♦ [“LDAP Integrator” on page 138](#)
- ♦ [“SMTP Integrator” on page 139](#)

- ♦ [“Syslog Integrator” on page 139](#)
- ♦ [“Using Non-FIPS Enabled Connectors with Sentinel in FIPS 140-2 Mode” on page 140](#)

## Agent Manager Connector

Follow the below procedure only if you have selected the **Encrypted (HTTPS)** option when configuring the networking settings of the Agent Manager Event Source Server.

### To configure the Agent Manager Connector to run in FIPS 140-2 mode:

- 1 Add or edit the Agent Manager Event Source Server. Proceed through the configuration screens until the Security window is displayed. For more information, see the *Agent Manager Connector Guide*.
- 2 Select one of the options from the *Client Authentication Type* field. The client authentication type determines how strictly the SSL Agent Manager Event Source Server verifies the identity of Agent Manager Event Sources that are attempting to send data.

- ♦ **Open:** Allows all the SSL connections coming from the Agent Manager agents. Does not perform any client certificate validation or authentication.
- ♦ **Strict:** Validates the certificate to be a valid X.509 certificate and also checks that the client certificate is trusted by the Event Source Server. New sources will need to be explicitly added to Sentinel (this prevents rogue sources from sending unauthorized data).

For the **Strict** option, you must import the certificate of each new Agent Manager client into the Sentinel FIPS keystore. When Sentinel is running in FIPS 140-2 mode, you cannot import the client certificate using the Event Source Management (ESM) interface.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

---

**NOTE:** In FIPS 140-2 mode, the Agent Manager Event Source Server uses the Sentinel server key pair; importing the server key pair is not required.

---

- 3 If server authentication is enabled in the agents, the agents must additionally be configured to trust the Sentinel server or the remote Collector Manager certificate depending on where the Connector is deployed.

**Sentinel server certificate location:** `/etc/opt/novell/sentinel/config/sentinel.cer`

**Remote Collector Manager certificate location:** `/etc/opt/novell/sentinel/config/rcm.cer`

---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), the Agent Manager agent must trust the appropriate certificate file.

---

## Database (JDBC) Connector

Follow the below procedure only if you have selected the *SSL* option when configuring the database connection.

### To configure the Database Connector to run in FIPS 140-2 mode:

- 1 Before configuring the Connector, download the certificate from the Database server and save it as `database.cert` file into the `/etc/opt/novell/sentinel/config` directory of the Sentinel server.

For more information, refer to the respective database documentation.

- 2 Import the certificate into the Sentinel FIPS keystore.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

- 3 Proceed with configuring the Connector.

## Sentinel Link Connector

Follow the below procedure only if you have selected **Encrypted (HTTPS)** option when configuring the networking settings of the Sentinel Link Event Source Server.

### To configure the Sentinel Link Connector to run in FIPS 140-2 mode:

- 1 Add or edit the Sentinel Link Event Source Server. Proceed through the configuration screens until the Security window is displayed. For more information, see the *Sentinel Link Connector Guide*.
- 2 Select one of the options from the *Client Authentication Type* field. The client authentication type determines how strictly the SSL Sentinel Link Event Source Server verifies the identity of Sentinel Link Event Sources (Sentinel Link Integrators) that are attempting to send data.
  - ♦ **Open:** Allows all the SSL connections coming from the clients (Sentinel Link Integrators). Does not perform any Integrator certificate validation or authentication.
  - ♦ **Strict:** Validates the Integrator certificate to be a valid X.509 certificate and also checks that the Integrator certificate is trusted by the Event Source Server. For more information, refer to the respective database documentation.

For the **Strict** option:

- ♦ If the Sentinel Link Integrator is in FIPS 140-2 mode, you must copy the `/etc/opt/novell/sentinel/config/sentinel.cer` file from the sender Sentinel machine to the receiver Sentinel machine. Import this certificate into the receiver Sentinel FIPS keystore.

---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), you must import the appropriate custom certificate file.

---

- ♦ If Sentinel Link Integrator is in non-FIPS mode, you must import the custom Integrator certificate into the receiver Sentinel FIPS keystore.

---

**NOTE:** If the sender is Sentinel Log Manager (in non-FIPS mode) and the receiver is Sentinel in FIPS 140-2 mode, the server certificate to be imported on the sender is the `/etc/opt/novell/sentinel/config/sentinel.cer` file from the receiver Sentinel machine.

---

When Sentinel is running in FIPS 140-2 mode, you cannot import the client certificate using the Event Source Management (ESM) interface. For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

---

**NOTE:** In FIPS 140-2 mode, the Sentinel Link Event Source server uses the Sentinel server key pair. Importing the server key pair is not required.

---

## Syslog Connector

Follow the below procedure only if you have selected the **SSL** protocol when configuring the network settings of the Syslog Event Source Server.

### To configure the Syslog Connector to run in FIPS 140-2 mode:

- 1 Add or edit the Syslog Event Source Server. Proceed through the configuration screens until the Networking window is displayed. For more information, see the *Syslog Connector Guide*.
- 2 Click **Settings**.
- 3 Select one of the options from the *Client Authentication Type* field. The client authentication type determines how strictly the SSL Syslog Event Source Server verifies the identity of Syslog Event Sources that are attempting to send data.
  - ♦ **Open:** Allows all the SSL connections coming from the clients (event sources). Does not perform any client certificate validation or authentication.
  - ♦ **Strict:** Validates the certificate to be a valid X.509 certificate and also checks that the client certificate is trusted by the Event Source Server. New sources will have to be explicitly added to Sentinel (this prevents rogue sources from sending data to Sentinel).

For the **Strict** option, you must import the certificate of the syslog client into the Sentinel FIPS keystore.

When Sentinel is running in FIPS 140-2 mode, you cannot import the client certificate using the Event Source Management (ESM) interface.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

---

**NOTE:** In FIPS 140-2 mode, the Syslog Event Source Server uses the Sentinel server key pair. Importing the server key pair is not required.

---

- 4 If server authentication is enabled in the syslog client, the client must trust the Sentinel server certificate or the remote Collector Manager certificate depending on where the Connector is deployed.

**The Sentinel server certificate file** is in the `/etc/opt/novell/sentinel/config/sentinel.cer` location.

**The Remote Collector Manger certificate file** is in `/etc/opt/novell/sentinel/config/rcm.cer` location.



---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), the client must trust the appropriate certificate file.

---

## Windows Event (WMI) Connector

To configure the Windows Event (WMI) Connector to run in FIPS 140-2 mode:

- 1 Add or edit the Windows Event Connector. Proceed through the configuration screens until the Security window is displayed. For more information, see the *Windows Event (WMI) Connector Guide*.
- 2 Click **Settings**.
- 3 Select one of the options from the *Client Authentication Type* field. The client authentication type determines how strictly the Windows Event Connector verifies the identity of the client Windows Event Collection Services (WECS) that are attempting to send data.

- ♦ **Open:** Allows all the SSL connections coming from the client WECS. Does not perform any client certificate validation or authentication.
- ♦ **Strict:** Validates the certificate to be a valid X.509 certificate and also checks that the client WECS certificate is signed by a CA. New sources will need to be explicitly added (this prevents rogue sources from sending data to Sentinel).

For the **Strict** option, you must import the certificate of the client WECS into the Sentinel FIPS keystore. When Sentinel is running in FIPS 140-2 mode, you cannot import the client certificate using the Event Source Management (ESM) interface.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

---

**NOTE:** In FIPS 140-2 mode, the Windows Event Source Server uses the Sentinel server key pair. Importing the server key pair is not required.

---

- 4 If server authentication is enabled in the Windows client, the client must trust the Sentinel server certificate or the remote Collector Manager certificate depending on where the Connector is deployed.

**The Sentinel server certificate file** is in the `/etc/opt/novell/sentinel/config/sentinel.cer` location.

**The remote Collector Manager certificate file** is in the `/etc/opt/novell/sentinel/config/rcm.cer` location.

---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), the client must trust the appropriate certificate file.

---

- 5 If you want to automatically synchronize the event sources or populate the list of event sources using an Active Directory connection, you must import the Active Directory server certificate into the Sentinel FIPS keystore.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

## Sentinel Link Integrator

Follow the below procedure only if you have selected the **Encrypted (HTTPS)** option when configuring the network settings of the Sentinel Link Integrator.

### To configure the Sentinel Link Integrator to run in FIPS 140-2 mode:

- 1 When Sentinel Link Integrator is in FIPS 140-2 mode, server authentication is mandatory?. Before configuring the Integrator instance, import the Sentinel Link Server certificate into the Sentinel FIPS keystore:

- ◆ **If Sentinel Link Connector is in FIPS 140-2 mode:**

If the Connector is deployed in the Sentinel server, you must copy the `/etc/opt/novell/sentinel/config/sentinel.cer` file from the receiver Sentinel machine to the sender Sentinel machine.

If the Connector is deployed in a remote Collector Manager, you must copy the `/etc/opt/novell/sentinel/config/rcm.cer` file from the receiver remote Collector Manager machine to the receiver Sentinel machine.

Import this certificate into the sender Sentinel FIPS keystore.

---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), you must import the appropriate custom certificate file.

---

- ◆ **If Sentinel Link Connector is in non-FIPS mode:**

Import the custom Sentinel Link Server certificate into the sender Sentinel FIPS keystore.

---

**NOTE:** When the Sentinel Link integrator is in FIPS 140-2 mode and the Sentinel Link Connector is in non-FIPS mode, use the custom server key pair on the connector. Do not use the internal server key pair.

---

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

- 2 Proceed with configuring the Integrator instance.

---

**NOTE:** In FIPS 140-2 mode, the Sentinel Link Integrator uses the Sentinel server key pair. Importing the Integrator key pair is not required.

---

## LDAP Integrator

### To configure the LDAP Integrator to run in FIPS 140-2 mode:

- 1 Before configuring the Integrator instance, download the certificate from the LDAP server and save it as `ldap.cert` file into the `/etc/opt/novell/sentinel/config` directory of the Sentinel server.

For example, use

```
openssl s_client -connect <LDAP server IP>:636
```

and then copy the text returned (between but not including the BEGIN and END lines) into a file.

- 2 Import the certificate into the Sentinel FIPS keystore.

For more information about importing the certificate, see [“Importing Certificates into FIPS Keystore Database” on page 140](#).

- 3 Proceed with configuring the Integrator instance.

## SMTP Integrator

The SMTP Integrator supports FIPS 140-2 from version 2011.1r2 and later. No configuration changes are required.

## Syslog Integrator

Perform the following procedure only if you have selected the Encrypted (SSL) option when configuring the network settings of the Syslog Integrator.

### To configure the Syslog Integrator to run in FIPS 140-2 mode:

- 1 When Syslog Integrator is in FIPS 140-2 mode, server authentication is mandatory. Before configuring the Integrator instance, import the Syslog Server certificate into the Sentinel FIPS keystore:
  - ♦ **If the Syslog Connector is in FIPS 140-2 mode:** If the Connector is deployed in the Sentinel server, you must copy the `/etc/opt/novell/sentinel/config/sentinel.cer` file from the receiver Sentinel server to the sender Sentinel server.  
  
If the Connector is deployed in a remote Collector Manager, you must copy the `/etc/opt/novell/sentinel/config/rcm.cer` file from the receiver remote Collector Manager computer to the receiver Sentinel computer.  
  
Import this certificate into the sender Sentinel FIPS keystore.  

---

**NOTE:** When using custom certificates that are digitally signed by a certificate authority (CA), you must import the appropriate custom certificate file.

---
  - ♦ **If Syslog Connector is in non-FIPS mode:** You must import the custom Syslog Server certificate into the sender Sentinel FIPS keystore.  

---

**NOTE:** When the Syslog Integrator is in FIPS 140-2 mode and the Syslog Connector is in non-FIPS mode, use the custom server key pair on the connector. Do not use the internal server key pair.

---

### To import certificates to the FIPS Keystore Database:

1. Copy the certificate file to any temporary location on the Sentinel server or remote Collector Manager.
  2. Go to the `/opt/novell/sentinel/bin` directory.
  3. Run the following command to import the certificate into the FIPS keystore database, and then follow the on-screen instructions:

```
./convert_to_fips.sh -i <certificate file path>
```
  4. Enter `yes` or `y` when prompted to restart the Sentinel server or remote Collector Manager.
- 2 Proceed with configuring the Integrator instance.

---

**NOTE:** In FIPS 140-2 mode, the Syslog Integrator uses the Sentinel server key pair. You do not need to import the Integrator key pair.

---

## Using Non-FIPS Enabled Connectors with Sentinel in FIPS 140-2 Mode

This section provides information about how to use non-FIPS enabled Connectors with a Sentinel server in FIPS 140-2 mode. We recommend this approach if you have sources that do not support FIPS or if you want to collect events from the non-FIPS Connectors in your environment.

### To use non-FIPS connectors with Sentinel in FIPS 140-2 mode:

- 1 Install a Collector Manager in non-FIPS mode to connect to the Sentinel server in FIPS 140-2 mode.

For more information, see [Part III, “Installing Sentinel,” on page 69](#).

- 2 Deploy the non-FIPS Connectors specifically to the non-FIPS remote Collector Manager.

---

**NOTE:** There are some known issues when non-FIPS Connectors such as Audit Connector and File Connector are deployed on a non-FIPS remote Collector Manager connected to a Sentinel server in FIPS 140-2 mode. For more information about these known issues, see the [Sentinel Release Notes](#).

---

## Importing Certificates into FIPS Keystore Database

You must insert certificates into the Sentinel FIPS keystore database to establish secure (SSL) communications from the components that own those certificates to Sentinel. You cannot upload certificates by using the Sentinel user interface when FIPS 140-2 mode is enabled. You must manually import the certificates into the FIPS keystore database.

For event sources that are using Connectors deployed to a remote Collector Manager, you must import the certificates to the FIPS keystore database of the remote Collector Manager rather than the central Sentinel server.

### To import certificates to the FIPS Keystore Database:

- 1 Copy the certificate file to any temporary location on the Sentinel server or remote Collector Manager.
- 2 Browse to the Sentinel bin directory. The default location is `/opt/novell/sentinel/bin`.
- 3 Run the following command to import the certificate into the FIPS keystore database, and then follow the on-screen instructions:

```
./convert_to_fips.sh -i <certificate file path>
```

- 4 Enter `yes` or `y` when prompted to restart the Sentinel server or remote Collector Manager.

# Reverting Sentinel to Non-FIPS Mode

This section provides information about how to revert Sentinel and its components to non-FIPS mode.

- ♦ [“Reverting Sentinel Server to Non-FIPS mode” on page 141](#)
- ♦ [“Reverting Remote Collector Managers or Remote Correlation Engines to Non-FIPS mode” on page 141](#)

## Reverting Sentinel Server to Non-FIPS mode

You can revert a Sentinel server running in FIPS 140-2 mode to non-FIPS mode only if you have taken a backup of your Sentinel server before converting it to run in FIPS 140-2 mode.

---

**NOTE:** When you revert a Sentinel server to non-FIPS mode, you will lose the events, incident data, and configuration changes made to your Sentinel server after converting to run FIPS 140-2 mode. The sentinel system will be restored back to the last restoration point of non-FIPS mode. You should take a backup of the current system before reverting to non-FIPS mode for future use.

---

### To revert your Sentinel server to non-FIPS mode:

- 1 Log in to the Sentinel server as the `root` user.
- 2 Switch to the `novell` user.
- 3 Browse to the Sentinel bin directory. The default location is `/opt/novell/sentinel/bin`.
- 4 Run the following command to revert your Sentinel server to non-FIPS mode, and follow the on-screen instructions:

```
./backup_util.sh -f <backup_file_name.tar.gz> -m 'restore'
```

For example, if `non-fips2013012419111359034887.tar.gz` is the backup file, run the following command:

```
./backup_util.sh -f non-fips2013012419111359034887.tar.gz -m 'restore'
```

- 5 Restart the Sentinel server.

## Reverting Remote Collector Managers or Remote Correlation Engines to Non-FIPS mode

You can revert remote Collector Managers or remote Correlation Engines to non-FIPS mode.

### To revert a remote Collector Managers or a remote Correlation Engine to non-FIPS mode:

- 1 Login to the remote Collector Manager or remote Correlation Engine system.
- 2 Switch to `novell` user (`su novell`).
- 3 Browse to the bin directory. The default location is `/opt/novell/sentinel/bin`.
- 4 Run the `revert_to_nonfips.sh` script and follow the on-screen instructions.
- 5 Restart the remote Collector Manager or remote Correlation Engine.



# 25 Adding a Consent Banner

Sentinel allows you to display a consent banner before login. You can specify the content of the banner as per your requirements. After you add the consent banner, you must accept the terms in the consent banner every time you log in to Sentinel.

**To add a consent banner:**

- 1 Log in to the Sentinel server as the `novell` user.
- 2 Browse to `/<Sentinel_installation_path>/var/opt/novell/sentinel/3rdparty/jetty/webapps/ROOT/siemdownloads`.
- 3 Add a text file with the name, `USER_AGREEMENT.txt`.
- 4 Enter the user agreement text.
- 5 Save the file.
- 6 Launch the Sentinel to view the consent banner.

Sentinel now displays the consent banner on the login screen.

---

**NOTE:** You must manually back up the `USER_AGREEMENT.txt` file before you upgrade Sentinel.

---





# 26 Limiting the Number of Concurrent Active Sessions

In Sentinel 8.2 SP3 or later, you can limit the number of concurrent active sessions you want to allow per user, tenant, or both. Limiting the number of sessions prevents attackers from launching sessions beyond the allowed limit, in case of an attack.

If you limit sessions per user and tenant, users will not be able to launch sessions after the total number of sessions launched by multiple users, reaches the limit allowed for the tenant.

Sentinel does not limit concurrent sessions, by default. You must configure this limit manually.

---

**NOTE:** This feature is available only in non MFA mode.

---

## To limit the number of concurrent active sessions:

- 1 Log in to the Sentinel server.
- 2 Open the `<sentinel_installation_path>/etc/opt/novell/sentinel/config/configuration.properties` file.
- 3 (Conditional) To configure the limit per tenant, set the `concurrent.overall.sessions` property to a required value.
- 4 (Conditional) To configure the limit per user, set the `concurrent.per.user.sessions` property to a required value.
- 5 Save the file.
- 6 Restart the Sentinel server.



# 27 Ending Inactive Sessions

In Sentinel 8.2 SP3 or later, you can configure Sentinel to end a session if there is no user activity for a specified duration. Sentinel displays a warning a minute before the specified duration ends. If a user keeps a session inactive till this duration, Sentinel logs the user out of the session.

Sentinel does not track user inactivity, by default. You must manually configure Sentinel to end inactive sessions.

## To set the inactivity timeout period:

- 1 Log in to the Sentinel server.
- 2 Open the `<sentinel_installation_path>/etc/opt/novell/sentinel/config/ui-configuration.properties` file.
- 3 Set the desired value for the `user.inactivity.time` property in milliseconds.
- 4 Refresh the browser through which you have logged into Sentinel.



# 28 Configuring IP Flow Data Collection

Sentinel leverages ArcSight SmartConnectors that help you monitor your enterprise network by collecting IP Flow data. SmartConnectors collect IP Flow data as events, which allow you to:

- ◆ Use existing Collector Managers to collect IP Flow data.
- ◆ Leverage IP Flow data in several areas of Sentinel such as visualizations, event routing, data federation, reports, and correlation.
- ◆ Apply data retention policies to IP Flow data, which allows you to store this data for the desired duration.

To configure IP Flow data collection, you must install and configure the ArcSight SmartConnector. While configuring, ensure that you configure the relevant SmartConnectors that collect IP Flow data.

For information about configuring SmartConnectors, see the Generic Universal CEF Collector documentation on the [Sentinel Plug-ins Website](#).



# V Upgrading Sentinel

This section provides information about upgrading Sentinel and other components.

- ♦ [Chapter 29, “Implementation Checklist,” on page 153](#)
- ♦ [Chapter 30, “Prerequisites,” on page 155](#)
- ♦ [Chapter 31, “Upgrading Sentinel Traditional Installation,” on page 157](#)
- ♦ [Chapter 32, “Upgrading the Sentinel Appliance,” on page 163](#)
- ♦ [Chapter 33, “Troubleshooting,” on page 171](#)
- ♦ [Chapter 34, “Post-Upgrade Configurations,” on page 175](#)
- ♦ [Chapter 35, “Upgrading Sentinel Plug-Ins,” on page 183](#)





# 29 Implementation Checklist

Before you upgrade Sentinel, review the following checklist to ensure a successful upgrade:

**Table 29-1** *Implementation Checklist*

| <input type="checkbox"/> | Tasks                                                                                                       | See                                                      |
|--------------------------|-------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| <input type="checkbox"/> | Ensure that the computers on which you install Sentinel and its components meet the specified requirements. | <a href="#">Sentinel System Requirements</a>             |
| <input type="checkbox"/> | Review the supported operating system release notes to understand the known issues.                         | <a href="#">SUSE Release Notes</a>                       |
| <input type="checkbox"/> | Review the Sentinel release notes to see new functionality and understand the known issues.                 | <a href="#">Sentinel Release Notes</a>                   |
| <input type="checkbox"/> | Complete the tasks mentioned in Prerequisites.                                                              | <a href="#">Chapter 30, "Prerequisites," on page 155</a> |



# 30 Prerequisites

- ♦ [“Saving the Custom Configuration Information” on page 155](#)
- ♦ [“Extending the Retention Period for Event Associations Data” on page 155](#)
- ♦ [“Pre-Upgrade Configuration for SSDM” on page 156](#)
- ♦ [“Change Guardian Integration” on page 156](#)

## Saving the Custom Configuration Information

### Saving the server.conf File Settings

If you have set any custom configuration parameter values in the `server.conf` file, save those values in separate files before the upgrade.

To save your custom configuration information:

- 1 Log in to the Sentinel server as the `novell` user and go to the `/etc/opt/novell/sentinel/config/` directory.
- 2 Create a configuration file named `server-custom.conf` and add your custom configuration parameters in this file.

Sentinel applies the saved custom configuration in these configuration files during the upgrade.

### Saving the jetty-ssl File Settings

Sentinel 8.1 includes an updated version of Jetty. The updated version of Jetty includes changes to its file structure.

If you have modified the `/etc/opt/novell/sentinel/3rdparty/jetty/jetty-ssl.xml` file in previous versions of Sentinel, such as excluding any ciphers, save those changes in a separate file before the Sentinel upgrade.

Once the Sentinel upgrade is complete, copy those changes to the `/etc/opt/novell/sentinel/3rdparty/jetty/jetty-ssl-context.xml` file and restart Sentinel.

## Extending the Retention Period for Event Associations Data

Starting from Sentinel 7.4.4, the default retention period for event associations data is 14 days. If you are upgrading from a Sentinel version prior to 7.4.4, the retention period you had set for event associations data will be overridden to 14 days after the upgrade. To avoid this, you can set the retention period to a desired value by adding a property in the `configuration.properties` file. For more information, see [“Configuring the Retention Period for the Event Associations Data”](#) in the *Sentinel Administration Guide*.

## Pre-Upgrade Configuration for SSDM

The upgrade process will update files related to Spark applications. To use the updated files, you must restart the Spark job and reset all Spark checkpoints on Kafka topics. To prevent data loss due to resetting the Kafka topic checkpoint, you must pause forwarding data from Collector Managers to Kafka before you upgrade SSDM. While data forwarding is paused, the data will be stored on the Collector Manager until data forwarding is resumed. Once the Spark application is done processing data that was forwarded to Kafka before forwarding was paused, the checkpoint can safely be reset without data loss.

### To pause event forwarding from Collector Manager to Kafka:

- 1 In Sentinel Main, click **Storage > Scalable Storage > Advanced Configuration > Kafka**.
- 2 Add the following property and set it to true:  
`pause.events.tokafka`
- 3 Click **Save**.

## Change Guardian Integration

Sentinel is compatible with Change Guardian 4.2 and later. To receive events from Change Guardian, you must first upgrade the Change Guardian server, Agents, and Policy editor to version 4.2 or later to ensure Sentinel continues to receive events from Change Guardian, post-upgrade.

# 31 Upgrading Sentinel Traditional Installation

The procedures in this chapter guide you through upgrading the Sentinel.

You can upgrade from Sentinel 8.2 or later.

The upgrade process does the following:

- ♦ Migrates Security Intelligence data and alerts data from MongoDB to PostgreSQL.  
Sentinel now stores Security Intelligence data, alerts data, and so on in PostgreSQL instead of MongoDB. The upgrade process will first migrate this data to PostgreSQL and if successful, will automatically proceed with the upgrade. If the data migration is unsuccessful, you cannot upgrade Sentinel.
- ♦ Generates a cleanup script that you can use to remove data and MongoDB related RPMs from MongoDB.

The data stored in MongoDB is retained as a backup.

- ♦ [“Upgrading Sentinel” on page 157](#)
- ♦ [“Upgrading Sentinel as a Non-root User” on page 159](#)
- ♦ [“Upgrading the Collector Manager or the Correlation Engine” on page 160](#)
- ♦ [“Upgrading the Operating System” on page 161](#)

## Upgrading Sentinel

Use the following steps to upgrade the Sentinel server:

**To upgrade the Sentinel server:**

- 1 Back up your configuration, then create an ESM export.  
For more information about backing up data, see [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide*.
- 2 (Conditional) If you have customized the configuration settings in the `server.xml`, `collector_mgr.xml`, or `correlation_engine.xml` files, ensure that you have created appropriate properties files named with the obj-component id to ensure that the customizations are retained after the upgrade. For more information, see [“Maintaining Custom Settings in XML Files”](#) in the *Sentinel Administration Guide*.
- 3 Download the latest installer from the [Download Website](#).
- 4 Log in as `root` to the server where you want to upgrade Sentinel.
- 5 Specify the following command to extract the install files from the tar file:  

```
tar xfz <install_filename>
```

  
Replace `<install_filename>` with the actual name of the install file.
- 6 Change to the directory where the install file was extracted.

7 Specify the following command to upgrade Sentinel:

```
./install-sentinel
```

8 To proceed with a language of your choice, select the number next to the language.

The end user license agreement is displayed in the selected language.

9 Read the end user license, enter `yes` or `y` to accept the license, then continue with the installation.

10 Select the required migration option. It migrates Security Intelligence data and Alerts data from MongoDB to PostgreSQL.

If you select the **Only upgrade without migrating data** option, the Sentinel server should be up and running.

---

**WARNING:** Ensure that you select the appropriate option because you cannot repeat this procedure after the upgrade is successful.

---

If your data is migrated successfully, the data that was stored in MongoDB is retained as a backup, and then the Sentinel upgrade process will automatically proceed.

The upgrade might take several minutes to complete.

11 (Conditional) If the data migration is not successful:

11a Clean up the partially migrated data. For more information, see [“Cleaning Up Data From PostgreSQL When Migration Fails”](#) on page 171.

11b Repeat from [Step 7](#) to [Step 10](#) above until you upgrade Sentinel.

12 Clear your web browser cache to view the latest Sentinel version.

13 (Conditional) If the `delete_old_cluster.sh` file is located in the `bin` folder (`/opt/novell/sentinel/3rdparty/postgresql/bin`), which means the PostgreSQL database has been upgraded to a major version (for example, 8.0 to 9.0). Clear the old PostgreSQL files from the PostgreSQL database. The folder path might be different in case of custom path installations.

To clear the old PostgreSQL files:

13a Switch to the `novell` user.

```
su novell
```

13b Browse to the `bin` folder:

```
cd /opt/novell/sentinel/3rdparty/postgresql/bin
```

13c Delete all the old PostgreSQL files by using the following command:

```
./delete_old_cluster.sh
```

14 Log in to Sentinel and verify the migrated data such as alerts, Security Intelligence data, and so on.

15 The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB”](#) on page 175.

16 To upgrade Collector Manager systems and Correlation Engine systems, see [“Upgrading the Collector Manager or the Correlation Engine”](#) on page 160.

# Upgrading Sentinel as a Non-root User

If your organizational policy does not allow you to run the full upgrade of Sentinel as `root`, you can upgrade Sentinel as another user. In this upgrade, a few steps are performed as a `root` user, then you proceed to upgrade Sentinel as another user created by the `root` user.

- 1 Back up your configuration, then create an ESM export.

For more information on backing up data, see “[Backing Up and Restoring Data](#)” in the *Sentinel Administration Guide*.

- 2 (Conditional) If you have customized the configuration settings in the `server.xml`, `collector_mgr.xml`, or `correlation_engine.xml` files, ensure that you have created appropriate properties files named with the obj-component id to ensure that the customizations are retained after the upgrade. For more information, see “[Backing Up and Restoring Data](#)” in the *Sentinel Administration Guide*.

- 3 Download the installation files from the [Downloads Website](#).

- 4 Specify the following command at the command line to extract the install files from the tar file:

```
tar -zxvf <install_filename>
```

Replace `<install_filename>` with the actual name of the install file.

- 5 Log in as `root` to the server where you want to upgrade Sentinel.

- 6 Extract the `squashfs` RPM from the Sentinel install files.

- 7 Install the `squashfs` on the Sentinel server.

```
rpm -Uvh <install_filename>
```

- 8 Specify the following command to change to the newly created non-root `novell` user: `novell`:

```
su novell
```

- 9 (Conditional) To do an interactive upgrade:

- 9a Specify the following command:

```
./install-sentinel
```

To upgrade Sentinel in a non-default location, specify the `--location` option along with the command. For example:

```
./install-sentinel --location=/foo
```

- 9b Continue with [Step 11](#).

- 10 (Conditional) To do a silent upgrade, specify the following command:

```
./install-sentinel -u <response_file>
```

The installation proceeds with the values stored in the response file. The Sentinel upgrade is complete.

- 11 Specify the number for the language you want to use for the upgrade.

The end user license agreement is displayed in the selected language.

- 12 Read the end user license and enter `yes` or `y` to accept the license and continue with the upgrade.

- 13 Select the migration option. It migrates Security Intelligence data and alerts data from MongoDB to PostgreSQL.

---

**WARNING:** Ensure that you select the appropriate option because you cannot repeat this procedure after the upgrade is successful.

---

If your data is migrated successfully, the data that was stored in MongoDB is retained as a backup, and then the Sentinel upgrade process will automatically proceed.

The upgrade might take several minutes to complete.

- 14 (Conditional) If the data migration is not successful:
  - 14a Clean up the migrated data. For more information, see [“Cleaning Up Data From PostgreSQL When Migration Fails”](#) on page 171.
  - 14b Repeat from [Step 7](#) to [Step 13](#) above until you upgrade Sentinel.
- 15 Clear your web browser cache to view the latest Sentinel version.
- 16 (Conditional) If the `delete_old_cluster.sh` file is located in the bin folder (`/opt/novell/sentinel/3rdparty/postgresql/bin`), which means the PostgreSQL database has been upgraded to a major version (for example, 8.0 to 9.0). Clear the old PostgreSQL files from the PostgreSQL database. The folder path might be different in case of custom path installations.

To clear the old PostgreSQL files:

  - 16a Switch to the novell user.

```
su novell
```
  - 16b Browse to the bin folder:

```
cd /opt/novell/sentinel/3rdparty/postgresql/bin
```
  - 16c Delete all the old PostgreSQL files by using the following command:

```
./delete_old_cluster.sh
```
- 17 Log in to Sentinel and verify the migrated data such as alerts, Security Intelligence data, and so on.
- 18 The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB”](#) on page 175.

## Upgrading the Collector Manager or the Correlation Engine

Use the following steps to upgrade the Collector Manager or the Correlation Engine:

- 1 Back up your configuration and create an ESM export.

For more information, see [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide*.
- 2 Navigate to the **Sentinel Main** interface as a user in the administrator role.
- 3 Select **Downloads**.
- 4 Click **Download Installer** in the Collector Manager Installer section.
- 5 Save the installer file on the respective Collector Manager or Correlation Engine server.



6 Copy the file to a temporary location.

7 Extract the contents of the file.

8 Run the following script:

**For Collector Manager:**

```
./install-cm
```

**For Correlation Engine:**

```
./install-ce
```

9 Follow the on-screen instructions to complete the installation.

10 (Conditional) For custom installations, run the following command to synchronize configurations between Sentinel server, Collector Manager, and Correlation Engine:

```
/opt/novell/sentinel/setup/configure.sh
```

## Upgrading the Operating System

This version of Sentinel includes a set of commands to use during the operating system upgrade procedure. These commands ensure Sentinel works correctly after you upgrade the operating system. Before you upgrade Sentinel, ensure that you refer system requirements for compatibility. For information, see [Sentinel System Requirements](#).

---

### IMPORTANT

- ◆ If you want to upgrade from SLES 11 SP4, you must first upgrade to SLES 12 SP3 or SLES12 SP4 and then upgrade to SLES 15 or SLES 15 SP1.
- ◆ If you are on RHEL version 7.7, do not upgrade to version 8.x because this is not supported by Red Hat. For more information, see the *Red Hat documentation*.

---

Use the following steps to upgrade your operating system:

1 On the Sentinel server where you want to upgrade your operating system, log in as one of the following:

- ◆ root user
- ◆ Non-root user

2 Open a command prompt and change to the directory where the Sentinel install file was extracted.

3 Stop the Sentinel services:

```
rcsentinel stop
```

4 (Conditional) If Sentinel was in FIPS mode before the operating system upgrade, NSS database files must be manually upgraded by running the following command:

```
certutil -K -d sql:/etc/opt/novell/sentinel/3rdparty/nss -X
```

Follow the on-screen instructions to upgrade the NSS database.

Give full permissions to novell user for the following files:

```
cert9.db
key4.db
pkcs11.txt
```

**5** Upgrade your operating system.

**6** (Conditional) When upgrading to SLES 15 or SLES 15 SP1, the following warning is displayed:

```
warning: Unsupported version of key: V3
```

You can either ignore the warning or perform a workaround to prevent the warning from being displayed. For more information about the workaround, see the [SLES documentation](#).

**7** (Conditional) If you are upgrading from RHEL version 7.6 to version 8.x, you may see a message that certain novell RPMs will be removed. If you see this message, you must

**7a** Upgrade to Sentinel 8.3

**7b** Back up Sentinel data.

**7c** Upgrade the RHEL operating system.

**7d** Restore Sentinel data.

**8** (Conditional) If you use Mozilla Network Security Services (NSS) 3.29, two dependent RPM files `libfreebl3-hmac` and `libsoftokn3-hmac` are not installed.

Manually install the following RPM files: `libfreebl3-hmac` and `libsoftokn3-hmac`.

**9** (Conditional) If you are upgrading from SLES12SP4 to SLES15SP1 in FIPS mode, you must first upgrade the SLES operating system, apply the latest operating system patches, and then start Sentinel.

**10** (Conditional) For RHEL 7.x, run the following command to check whether there are any errors in the RPM database:

```
rpm -qa --dbpath <install_location>/rpm | grep novell
```

```
Example: # rpm -qa --dbpath /custom/rpm | grep novell
```

**10a** If there are any errors, run the following command to fix the errors:

```
rpm --rebuilddb --dbpath <install_location>/rpm
```

```
Example: # rpm --rebuilddb --dbpath /custom/rpm
```

**10b** Run the command mentioned in Step 7 to ensure that there are no errors.

**11** Repeat this procedure on the following:

- ♦ Collector Managers
- ♦ Correlation Engines

**12** Restart the Sentinel service:

```
rcsentinel restart
```

This step is not applicable for Sentinel HA.

# 32 Upgrading the Sentinel Appliance

The procedures in this chapter guide you through upgrading the Sentinel appliance.

Sentinel now stores Security Intelligence data, alerts data, and so on in PostgreSQL instead of MongoDB. You can upgrade the appliance only after you migrate data from MongoDB to PostgreSQL successfully.

The data stored in MongoDB is retained as a backup and you can delete it after upgrading Sentinel.

- ♦ [“Prerequisites for Upgrading the Appliance” on page 163](#)
- ♦ [“Upgrading the Appliance” on page 166](#)
- ♦ [“Applying Operating System Patches” on page 170](#)

## Prerequisites for Upgrading the Appliance

Before you upgrade, complete the following prerequisites:

1. You must have Sentinel 8.2 or later installed.
2. You must have SLES 12 SP3 or SLES 12 SP4 installed.
  - a. (Conditional) If you have SLES 11 SP4, you must upgrade to SLES 12 SP3. For more information about upgrading the SLES operating system, see [“Upgrading the Operating System to SLES 12 SP3” on page 163](#).
  - b. (Conditional) If you have upgraded to SLES 12 SP3 earlier, you must download the latest version of the post upgrade utility and run it. This is important to ensure that the upgrade works fine.  
  
Download and run the latest utility from [Sentinel Appliance 8.2 SLES11SP4 to SLES12SP3 Migration Tools B](#) in the [Micro Focus Patch Finder](#) website.
3. Migrate Security Intelligence data, alerts data, and so on from MongoDB to PostgreSQL. You can perform this only after completing the previous prerequisites. For more information about migrating data, [“Migrating Data from MongoDB to PostgreSQL” on page 165](#).

You must run the migration script even when you do not have data to migrate because the migration script generates a cleanup script. You can use the cleanup script to remove MongoDB data, which will be redundant after you upgrade Sentinel.

## Upgrading the Operating System to SLES 12 SP3

You must upgrade the operating system because:

- ♦ Sentinel is now available only on the SLES 12 channel. Therefore, to continue receiving Sentinel and operating system updates, you must first upgrade the operating system to SLES 12 SP3 before you upgrade Sentinel.
- ♦ You can leverage Sentinel Appliance Manager capabilities. Sentinel Appliance Manager provides a simple Web-based user interface that helps you to configure and manage the appliance.

## To upgrade the operating system and configure the appliance:

- 1 Stop the Sentinel services:

```
rcsentinel stop
```

- 2 (Conditional) If Sentinel was in FIPS mode before the operating system upgrade, NSS database files must be manually upgraded by running the following command:

```
certutil -K -d sql:/etc/opt/novell/sentinel/3rdparty/nss -X
```

Follow the on-screen instructions to upgrade the NSS database.

Give full permissions to `novell` user for the following files:

```
cert9.db
key4.db
pkcs11.txt
```

- 3 (Conditional) If you are using Mozilla Network Security Services (NSS) 3.29, two dependent RPM files `libfreebl3-hmac` and `libsoftokn3-hmac` are not installed.

Manually install the following RPM files: `libfreebl3-hmac` and `libsoftokn3-hmac`.

- 4 Download the SLES 12 SP3 installer and the post-upgrade utility from the [Micro Focus Patch Finder](#) website. For Sentinel HA, download the SLES 12 SP3 HA file as well.
- 5 Follow the installation prompts to upgrade the operating system. For Sentinel HA, when prompted to install additional add-on products, select the location where you have downloaded the SLES 12 SP3 HA file and proceed with the upgrade.

For more information about upgrading to SLES 12 SP3, see the [SLES documentation](#).

---

**IMPORTANT:** You will be prompted to register for SLES 12 SP3 during the upgrade. However, skip the registration. Registering for updates in this screen will only register for SLES 12 SP3 updates from the SUSE Customer Channel, which is not supported. Also, you will not receive Sentinel updates. Therefore, register for updates only after completing Step 9 to receive both Sentinel and SLES 12 SP3 updates from the Sentinel appliance update channel.

---

- 6 During the upgrade process, SLES renames the `/etc/sysctl.conf` file to `/etc/sysctl.conf.rpmsave` as a back up and creates a new `/etc/sysctl.conf` file. After you upgrade, copy the contents of the `/etc/sysctl.conf.rpmsave` file to the `/etc/sysctl.conf` file. Open the `sysctl.conf` file and search for `# Added by sentinel vm.max_map_count`. Move this setting to the next line as follows:

Change

```
net.core.wmem_max = 67108864# Added by sentinel vm.max_map_count :
65530
vm.max_map_count = 262144
```

to

```
net.core.wmem_max = 67108864
Added by sentinel vm.max_map_count : 65530
vm.max_map_count = 262144
```

- 7 (Conditional) For Sentinel HA, complete the steps mentioned in the following sections:
  - ◆ [“Configuring iSCSI Targets” on page 225](#)

- ♦ [“Configuring iSCSI Initiators” on page 226](#)
  - ♦ [“Configuring the HA Cluster” on page 226](#)
- 8 To configure the appliance, run the post-upgrade utility from the command prompt:

8a Untar the file:

```
tar -xvf <post_upgrade_utility_installer_filename>.tar.gz
```

8b Change to the directory where you extracted the utility:

```
cd <post_upgrade_utility_installer_filename>
```

8c To configure the appliance, run the following script:

```
./appliance_SLESISO_post_upgrade.sh
```

---

**NOTE:** Do not run this script remotely since this script involves network reconfiguration.

---

8d Follow the on-screen instructions to complete the configuration.

This script reconfigures the installed packages and configures packages for managing appliance.

- 9 Using your existing registration code, register for updates again to receive Sentinel and latest operating system updates. For more information, see [“Registering for Updates” on page 106](#).

## Migrating Data from MongoDB to PostgreSQL

You must migrate Security Intelligence data, alerts data, and so on from MongoDB to PostgreSQL by running the migration script.

The migration script does the following:

- ♦ Migrates Security Intelligence data and alerts data to PostgreSQL.
- ♦ Generates a cleanup script that you can use to remove data and MongoDB related RPMs from MongoDB.

---

**WARNING:** After you migrate data, you must upgrade Sentinel before you start or restart Sentinel. This is to ensure that there is no loss of data coming to Sentinel.

---

**To migrate data:**

- 1 Download `Mongo_To_PostgreSQL_Migration_Utility_8.3.0.0-5575.tar.gz` from the [Download Website](#) website.
- 2 Untar the files.
- 3 Log in to the appliance console as a `novell` user.

---

**IMPORTANT:** Run the migration script from terminal of the machine. Do not use an emulation terminal software like PuTTY or MobaXterm.

---

- 4 Run the following script: `mongo_to_pgsql_migration.sh`.
- 5 Select the migration option as per your requirements.

---

**WARNING:** Ensure that you select the appropriate option because you cannot repeat this procedure after the migration is successful.

---

If your data is migrated successfully, a confirmation message will be displayed on the screen. You can now upgrade the appliance.

- 6 (Conditional) If the data migration is not successful:
  - 6a Clean up the migrated data. For more information, see [“Cleaning Up Data From PostgreSQL When Migration Fails”](#) on page 171.
  - 6b Repeat this procedure to migrate data.
- 7 (Conditional) If you see the following error when running the migration script, complete the tasks mentioned in [“Unable to Run the Migration Script”](#) on page 172:

```
8101server:/opt # su novell
novell@8101server:/opt>
novell@8101server:/opt> ./mongo_to_pgsql_migration.sh
./mongo_to_pgsql_migration.sh: line 25: /bin/setenv.sh: No such file or
directory
Cannot execute ./mongo_to_pgsql_migration.sh as novell
novell@8101server:/opt>
novell@8101server:/opt> exit
exit
8101server:/opt #
8101server:/opt # ./mongo_to_pgsql_migration.sh
./mongo_to_pgsql_migration.sh: line 25: /bin/setenv.sh: No such file or
directory
Cannot execute ./mongo_to_pgsql_migration.sh as root
```

## Upgrading the Appliance

You can upgrade both Sentinel and the SLES operating system through the Sentinel Appliance Manager, Appliance Update Channel, or Subscription Management Tool (SMT). You must first complete the prerequisites listed in [“Prerequisites for Upgrading the Appliance”](#) on page 163 and then upgrade the appliance.

- ♦ [“Upgrading through the Appliance Update Channel”](#) on page 166
- ♦ [“Upgrading through SMT”](#) on page 168
- ♦ [“Upgrading through Sentinel Appliance Manager”](#) on page 169

## Upgrading through the Appliance Update Channel

You can upgrade Sentinel by using Zypper. Zypper is a command line package manager that allows you to perform an interactive upgrade of appliance. In instances where user interaction is required to complete the upgrade, such as an end user license agreement update, you must upgrade the Sentinel appliance using Zypper.

**To upgrade the appliance through the Appliance update channel:**

- 1 Back up your configuration, then create an ESM export.

For more information, see “[Backing Up and Restoring Data](#)” in the *Sentinel Administration Guide*.

- 2 (Conditional) If you have customized the configuration settings in the `server.xml`, `collector_mgr.xml`, or `correlation_engine.xml` files, ensure that you have created appropriate properties files named with the obj-component id to ensure that the customizations are retained after the upgrade. For more information, see “[Maintaining Custom Settings in XML Files](#)” in the *Sentinel Administration Guide*.
- 3 Log in to the appliance console as the `root` user.
- 4 Run the following command:

```
/usr/bin/zypper patch
```
- 5 (Conditional) If the installer displays a message that you must resolve dependency for the OpenSSH package, enter the appropriate option to downgrade the OpenSSH package.
- 6 (Conditional) If the installer displays a message that indicates change in the `ncgOverlay` architecture, enter the appropriate option to accept the architecture change.
- 7 (Conditional) If the installer displays a message that you must resolve dependency for some appliance packages, enter the appropriate option to uninstall the dependent packages.
- 8 Enter `Y` to proceed.
- 9 Enter `yes` to accept the license agreement.
- 10 Open the `/etc/sysctl.conf` file and search for `# Added by sentinel vm.max_map_count`. Move this setting to the next line as follows:  
Change

```
net.core.wmem_max = 67108864# Added by sentinel vm.max_map_count : 65530
vm.max_map_count = 262144
```

to

```
net.core.wmem_max = 67108864
Added by sentinel vm.max_map_count : 65530
vm.max_map_count = 262144
```
- 11 Restart the Sentinel appliance.
- 12 (Conditional) If Sentinel is installed on a custom port or if the Collector Manager or the Correlation Engine is in FIPS mode, run the following command:

```
/opt/novell/sentinel/setup/configure.sh
```
- 13 Clear your web browser cache to view the latest Sentinel version.
- 14 (Conditional) If the PostgreSQL database has been upgraded to a major version (for example, 8.0 to 9.0 or 9.0 to 9.1), clear the old PostgreSQL files from the PostgreSQL database. For information about whether the PostgreSQL database was upgraded, see the Sentinel Release Notes.
  - 14a Switch to `novell` user.

```
su novell
```
  - 14b Browse to the `bin` folder:

```
cd /opt/novell/sentinel/3rdparty/postgresql/bin
```

**14c** Delete all the old PostgreSQL files by using the following command:

```
./delete_old_cluster.sh
```

- 15** (Conditional) To upgrade the Collector Manager or the Correlation Engine, follow [Step 3](#) through [Step 12](#).
- 16** (Conditional) If you are running Sentinel in an HA environment, repeat these steps on all nodes in the cluster.
- 17** Restart Sentinel.
- 18** Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence data, and so on.
- 19** The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).

## Upgrading through SMT

In secured environments where the appliance must run without direct internet access, you can configure the appliance with Subscription Management Tool (SMT) that allows you upgrade the appliance to the latest available versions.

### To upgrade the appliance through SMT:

- 1** Ensure that the appliance is configured with SMT.  
For more information, see [“Configuring the Appliance with SMT” on page 108](#).
- 2** Back up your configuration, then create an ESM export.  
For more information, see [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide*.
- 3** (Conditional) If you have customized the configuration settings in the `server.xml`, `collector_mgr.xml`, or `correlation_engine.xml` files, ensure that you have created appropriate properties files named with the obj-component id to ensure that the customizations are retained after the upgrade. For more information, see [“Maintaining Custom Settings in XML Files”](#) in the *Sentinel Administration Guide*.
- 4** Log in to the appliance console as the `root` user.
- 5** Refresh the repository for upgrade:

```
zypper ref -s
```
- 6** Check whether the appliance is enabled for upgrade:

```
zypper lr
```
- 7** (Optional) Check the available updates for the appliance:

```
zypper lu
```
- 8** (Optional) Check the packages that include the available updates for the appliance:

```
zypper lp -r SMT-http_<smt_server_fqdn>:<package_name>
```



9 Update the appliance:

```
zypper up -t patch -r SMT-http_<smt_server_fqdn>:<package_name>
```

10 Open the `/etc/sysctl.conf` file and search for `# Added by sentinel vm.max_map_count`. Move this setting to the next line as follows:

Change

```
net.core.wmem_max = 67108864# Added by sentinel vm.max_map_count : 65530
```

```
vm.max_map_count = 262144
```

to

```
net.core.wmem_max = 67108864
```

```
Added by sentinel vm.max_map_count : 65530
```

```
vm.max_map_count = 262144
```

11 Restart the appliance.

```
rcsentinel restart
```

12 (Conditional) If Sentinel is installed on a custom port or if the Collector Manager or the Correlation Engine is in FIPS mode, run the following command:

```
/opt/novell/sentinel/setup/configure.sh
```

13 (Conditional) To upgrade the Collector Manager or the Correlation Engine, follow [Step 4](#) through [Step 12](#).

14 (Conditional) If you are running Sentinel in an HA environment, repeat these steps on all nodes in the cluster.

15 Restart Sentinel.

16 Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence data, and so on.

17 The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up disk space, you can delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).

## Upgrading through Sentinel Appliance Manager

### To upgrade the appliance through Sentinel Appliance Manager:

- 1 Launch the Sentinel appliance by doing either of the following:
  - ◆ Log in to Sentinel. Click **Sentinel Main** > **Appliance**.
  - ◆ Specify the following URL in your web browser: `https://<IP_address>:9443`.
- 2 Log in either as a `vaadmin` or a `root` user.
- 3 (Conditional) Register for updates if you have not done it earlier. For more information, see [“Registering for Updates” on page 106](#).
- 4 Click **Online Update**.
- 5 To install the displayed updates, click **Update Now** > **OK**.
- 6 To apply the installed updates, click **Reboot**.

- 7 Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence data, and so on.
- 8 The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, you can delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).

## Applying Operating System Patches

To apply operating system patches:

- 1 Launch the Sentinel appliance by doing either of the following:
  - ♦ Log in to Sentinel. Click **Sentinel Main > Appliance**.
  - ♦ Specify the following URL in your web browser: `https://<IP_address>:9443`.
- 2 Log in either as a vaadmin or a root user.
- 3 Click **Online Update**.
  - 3a (Conditional) Register for updates if you have not done it earlier. For more information, see [“Registering for Updates” on page 106](#).
  - 3b To install the displayed updates for the operating system, click **Update Now > OK**.
- 4 To apply the installed updates, click **Reboot**.

# 33 Troubleshooting

- ♦ [“Cleaning Up Data From PostgreSQL When Migration Fails” on page 171](#)
- ♦ [“Unable to Run the Migration Script” on page 172](#)
- ♦ [“Cannot Connect to Servers or Other Components through Appliance” on page 172](#)
- ♦ [“Error When Upgrading the Appliance” on page 173](#)

## Cleaning Up Data From PostgreSQL When Migration Fails

If the migration fails, you must delete the data that has been partially moved to the PostgreSQL database and then run the migration script again.

---

**WARNING:** Do not perform this procedure if migration is successful. This script deletes all the migrated data.

---

### To clean up the partially migrated data:

- 1 Ensure that the PostgreSQL database is up and running.
- 2 Log in to the Sentinel server as the `novell` user.
- 3 Go to the location where you have extracted the Sentinel installer or migration utility.
- 4 Run the `./db_migration_failure_cleanup.sh` script to delete the partially migrated data.
- 5 Run the `rm db_migration_failure_cleanup.sh` command to delete the `db_migration_failure_cleanup.sh` file.

To continue with the traditional upgrade, see [Chapter 31, “Upgrading Sentinel Traditional Installation,” on page 157](#).

To continue with the appliance upgrade, migrate data from MongoDB to PostgreSQL. For information, see [“Migrating Data from MongoDB to PostgreSQL” on page 165](#).

# Unable to Run the Migration Script

You may see the following error when running the migration script to move data to PostgreSQL:

```
8101server:/opt # su novell
novell@8101server:/opt>
novell@8101server:/opt> ./mongo_to_pgsql_migration.sh
./mongo_to_pgsql_migration.sh: line 25: /bin/setenv.sh: No such file or
directory
Cannot execute ./mongo_to_pgsql_migration.sh as novell
novell@8101server:/opt>
novell@8101server:/opt> exit
exit
8101server:/opt #
8101server:/opt # ./mongo_to_pgsql_migration.sh
./mongo_to_pgsql_migration.sh: line 25: /bin/setenv.sh: No such file or
directory
Cannot execute ./mongo_to_pgsql_migration.sh as root
```

This error may occur if you had upgraded appliance to Sentinel 8.2 from a prior version because the `bashrc` may have been modified during a prior upgrade.

To prevent this error, you must update the `bashrc` file.

## To update the `bashrc` file:

- 1 Open the `bashrc` file:

```
/home/novell/.bashrc
```

- 2 (Conditional) If the file does not include the following properties, add them:

```
APP_HOME="/opt/novell/sentinel"
export PATH="$APP_HOME/bin:$APP_HOME/bin/actions:$PATH"
```

- 3 Run the migration script again. For more information, see [“Migrating Data from MongoDB to PostgreSQL” on page 165](#).

# Cannot Connect to Servers or Other Components through Appliance

Prior installations of Sentinel may include the IP address as 127.0.0.2 in the `/etc/hosts` file if you had chosen [Assign hostname to loopback address](#) option during installation. This may cause communication issues with other servers or components. You must edit the file and remove this IP address.

## To remove the IP address:

- 1 Open the `/etc/hosts` file.
- 2 Comment the entry for IP address 127.0.0.2.
- 3 Save the file.

# Error When Upgrading the Appliance

When you are upgrading the appliance and you had upgraded to 8.2 or later from a prior version, you may see the following error:

```
(104/134) Installing: kernel-default-4.12.14-95.45.1.x86_64
.....
.....[error]
Installation of kernel-default-4.12.14-95.45.1.x86_64 failed:
Error: Subprocess failed. Error: RPM failed: installing package kernel-
default-4.12.14-95.45.1.x86_64 needs 4MB on the /boot filesystem
```

This is a known issue with the SUSE operating system and not with Sentinel. Therefore, to resolve this issue, follow the workaround provided in the [SUSE documentation](#).



# 34 Post-Upgrade Configurations

This chapter includes the post upgrade configurations.

- ♦ “Removing Data from MongoDB” on page 175
- ♦ “Synchronizing the postgresql.conf file” on page 175
- ♦ “Securing Data in Elasticsearch” on page 176
- ♦ “Configuring Event Visualizations” on page 176
- ♦ “Configuring IP Flow Data Collection” on page 176
- ♦ “Post-Upgrade Configuration for Sentinel Scalable Data Manager” on page 177
- ♦ “Adding the JDBC DB2 Driver” on page 180
- ♦ “Configuring Data Federation Properties in Sentinel Appliance” on page 180
- ♦ “Registering Sentinel Appliance for Updates” on page 181
- ♦ “Updating External Databases for Data Synchronization” on page 181
- ♦ “Updating Permissions for Users Who Send Data from Other Integrated Products to Sentinel” on page 181

## Removing Data from MongoDB

After you upgrade Sentinel you will not need the data stored in MongoDB. You can delete this data to free up the disk space.

**To clean up storage space:**

- 1 Log in to the Sentinel server as `root` user.
- 2 Go to `<sentinel_installation_path>/opt/novell/sentinel/bin`.
- 3 Run the following script:  

```
./mongoDB_cleanup.sh
```

## Synchronizing the postgresql.conf file

During upgrade the `postgresql.conf` file from a prior version is renamed to `postgresql.conf_old`. A new `postgresql.conf` file is created for 8.3. The new `postgresql.conf` file contains configurations to improve the performance of the Security Intelligence dashboard. Therefore, you must retain this file. In case, your customizations are not included in the new file, edit the new `postgres.sql` file. Both the files are at the following location: `/var/opt/novell/sentinel/3rdparty/postgresql/data/`

## Securing Data in Elasticsearch

Sentinel leverages Kibana, a browser-based analytics and search dashboard, which helps you to visualize events and alerts in dashboards. Sentinel stores and indexes alerts in Elasticsearch. You can configure Sentinel to also store and index events in Elasticsearch to leverage event visualization capabilities. Sentinel dashboards access data from Elasticsearch to present events and alerts in dashboards. To ensure that the dashboards display only the data that a user's role is authorized to view and to prevent unauthorized data access in Elasticsearch, you must install the Elasticsearch Security plug-in. For more information, see [“Securing Data in Elasticsearch” on page 78](#).

## Configuring Event Visualizations

Sentinel provides event visualizations that present data in charts, tables, and maps. These visualizations make it easier to visualize and analyze large volumes of data such as events, IP Flow events, and alerts. You can also create your own visualizations and dashboards.

Sentinel leverages Kibana, a browser-based analytics and search dashboard, that helps you to search and visualize events. Kibana accesses data from visualization data store (Elasticsearch) to present events in dashboards. By default, Sentinel includes an Elasticsearch node. You must enable event visualization to store and index events in Elasticsearch. For more information, see [“Configuring the Visualization Data Store” on page 42](#).

---

**NOTE:** Some of the Sentinel dashboards that leverage Kibana do not load after you upgrade to Sentinel 8.2. This issue occurs because Elasticsearch and Kibana versions have been upgraded in Sentinel 8.2, and the existing Kibana index file is not compatible with the upgraded versions of Elasticsearch and Kibana. To fix this issue, you must manually delete the existing Kibana index file and recreate a new Kibana index file. For more information, see the [Knowledge Base Article 7022736](#).

---

## Configuring IP Flow Data Collection

Sentinel leverages ArcSight SmartConnectors that help you monitor your enterprise network by collecting IP Flow data. SmartConnectors collect IP Flow data as events and are therefore considered for EPS count. This allows you to:

- ◆ Use existing Collector Managers to collect IP Flow data.
- ◆ Leverage IP Flow data in several areas of Sentinel such as visualizations, event routing, data federation, reports, and correlation.
- ◆ Apply data retention policies to IP Flow data, which allows you to store this data for the desired duration.

The IP Flow functionality is now enabled by default. You must install and configure the ArcSight SmartConnector to collect IP Flow data,.



Sentinel no longer includes NetFlow capabilities including NetFlow views. With SmartConnectors collecting IP Flow data as events you can use existing Collector Managers to collect NetFlow data. So, you will no longer need NetFlow Collector Managers to collect NetFlow data. Therefore, you can uninstall any existing NetFlow Collector Managers.

- ♦ [“Configuring SmartConnectors that collect IP Flow data” on page 177](#)
- ♦ [“Uninstalling Existing NetFlow Collector Managers” on page 177](#)

## Configuring SmartConnectors that collect IP Flow data

Install and configure the ArcSight SmartConnector. While configuring, ensure that you configure the relevant SmartConnectors that collect IP Flow data.

For information about configuring SmartConnectors, see the Generic Universal CEF Collector documentation on the [Sentinel Plug-ins Website](#).

## Uninstalling Existing NetFlow Collector Managers

**To uninstall existing NetFlow Collector Managers:**

- 1 Log in to the NetFlow Collector Manager computer with the same user permission that you used to install the NetFlow Collector Manager.

- 2 Change to the following directory:

```
/opt/novell/sentinel/setup
```

- 3 Run the following command:

```
./uninstall-sentinel
```

- 4 Enter `y` to uninstall the Collector Manager.

The script first stops the service and then uninstalls the Collector Manager completely.

## Post-Upgrade Configuration for Sentinel Scalable Data Manager

- ♦ [“Install Elasticsearch Security Plug-In” on page 178](#)
- ♦ [“Updating Spark Applications on YARN” on page 178](#)
- ♦ [“Enabling Sentinel Features” on page 179](#)
- ♦ [“Updating Dashboards and Visualizations in Sentinel Scalable Data Manager” on page 179](#)

## Install Elasticsearch Security Plug-In

In addition to external Elasticsearch nodes, Sentinel now includes an local Elasticsearch node by default for data visualization. Hence, you must install an Elasticsearch plug-in for the local Elasticsearch. For more information, see [“Installing the Elasticsearch Security Plug-In” on page 79](#).

As the Elasticsearch and Kibana used in Sentinel are upgraded, you must redeploy all the Elasticsearch security plug-ins in the existing Elasticsearch nodes. For more information about redeploying Elasticsearch security plug-in, see [“Redeploying Elasticsearch Security Plug-In” on page 82](#).

## Updating Spark Applications on YARN

During the Sentinel upgrade, some of the Spark application files are also updated. You must re-submit the Spark applications with these updated files by performing the following steps:

- 1 Log in to the SSDM server as the `novell` user and copy the files to the Spark history server where HDFS NameNode is installed:

```
cd /etc/opt/novell/sentinel/scalablestore
scp SparkApp-*.jar avroevent-*.avsc avrorawdata-*.avsc spark.properties
log4j.properties manage_spark_jobs.sh
root@<hdfs_node>:<destination_directory>
```

where `<destination_directory>` is any directory where you want to place the copied files. Also, ensure that the `hdfs` user has full permissions to this directory.

- 2 Log in to the `<hdfs_node>` server as the root user and change the ownership of the copied files to `hdfs` user:

```
cd <destination_directory>
chown hdfs SparkApp-*.jar avroevent-*.avsc avrorawdata-*.avsc
spark.properties log4j.properties manage_spark_jobs.sh
```

Assign executable permission to the `manage_spark_jobs.sh` script.

- 3 Ensure that the Spark jobs have completed processing all the data:

Go to YARN ResourceManager Web user interface and view each Sentinel Spark application. The Spark Streaming application data will show the input rate drop to zero when all data has been processed from Kafka.

- 4 Run the following command to stop data processing:

```
./manage_spark_jobs.sh stop
```

- 5 Clear the data processing checkpoint:

```
sudo -u hdfs hadoop fs -rm -R -skipTrash /spark/checkpoint
```

where `/spark/checkpoint` is the checkpoint directory.

- 6 Run the following script to re-submit the Spark jobs:

```
./manage_spark_jobs.sh start
```

The above command takes a while to complete the submit process.

- 7 (Optional) Run the following command to verify the status of the submitted Spark jobs:

```
./manage_spark_jobs.sh status
```

- 8 Resume event forwarding to Kafka for Spark to start processing events:
  - 8a In Sentinel Main, click **Storage > Scalable Storage > Advanced Configuration > Kafka**.
  - 8b Set the following property to false:

```
pause.events.tokafka
```
  - 8c Click **Save**.

## Enabling Sentinel Features

When you upgrade from SSDM 8.0.x.x, some of the Sentinel features added in Sentinel 8.1 and later are not available by default. You must manually enable those features in the `/etc/opt/novell/sentinel/config/ui-configuration.properties` file.

- 1 Log in to the Sentinel server as `novell` user.
- 2 Open the `/etc/opt/novell/sentinel/config/ui-configuration.properties` file.
- 3 Change the following properties to false:

```
alerts.hideUI
solutionDesigner.launcher.hideUI
correlation.hideUI
scc.configurations.solutionPacks.hideUI
people.hideUI
permission.knowledgeBase.hideUI
scc.menuBarItem.toolsMenu.hideUI
scc.toolBarItem.peopleBrowser.hideUI
integration.hideUI
```

- 4 Refresh the Sentinel browser.

## Updating Dashboards and Visualizations in Sentinel Scalable Data Manager

You must update dashboards and visualizations after upgrading SSDM, so that the enhancements included in the latest version for dashboards and visualizations are applied.

When you upgrade SSDM, dashboards and visualizations are not updated by default. However, you can update them manually after the upgrade. You can update dashboards and visualizations by deleting the existing dashboards and visualizations and running the `load_kibana_data.sh` script, which installs latest dashboards and visualizations.

---

**IMPORTANT:** The customizations you might have done in dashboards and visualizations will be lost when you update dashboards and visualizations.

---

To update dashboards and visualizations:

- 1 Log in to the SSDM web interface and go to Event Visualization.
- 2 In Event Visualization, go to **Settings > Objects > Dashboards**.
- 3 Select the dashboards you want to update, and click **Delete**.
- 4 Click **Visualizations**. Select the visualizations you want to update, and click **Delete**.

- 5 Log out of the SSDM web interface.
- 6 Log in to the SSDM server as the `novell` user.
- 7 Go to the `/opt/novell/sentinel/bin` directory.
- 8 Run the `load_kibana_data.sh` using the following command:  

```
./load_kibana_data.sh http://<ip address>:<port>> <alerts/events/misc>
```

 For example:  

```
./load_kibana_data.sh http://127.0.0.1:9200 alerts
```

```
./load_kibana_data.sh http://127.0.0.1:9200 events
```
- 9 Log in to the SSDM web interface and go to Event Visualization to view the updated dashboards and visualizations.

## Adding the JDBC DB2 Driver

After upgrading Sentinel, add the correct JDBC Driver and configure it for data collection and data synchronization, by performing the following steps:

- 1 Copy the correct version of the IBM DB2 JDBC driver (`db2jcc-*.jar`) for your version of the DB2 database in the `/opt/novell/sentinel/lib` folder.
- 2 Ensure that you set the necessary ownership and permissions for the driver file.
- 3 Configure this driver for data collection. For more information, see the [Database Connector documentation](#).

## Configuring Data Federation Properties in Sentinel Appliance

Perform the following procedure after upgrading Sentinel appliance, so that data federation does not display any errors in the environment where two or more NICs are configured:

- 1 In the authorized requestor server, add the following property in the `/etc/opt/novell/sentinel/config/configuration.properties` file as follows:  

```
sentinel.distsearch.console.ip=<one of the authorized requestor's IP addresses>
```
- 2 In the data source server, add the following property in the `/etc/opt/novell/sentinel/config/configuration.properties` file as follows:  

```
sentinel.distsearch.target.ip=<one of the data source's IP addresses>
```
- 3 Restart Sentinel:  

```
rcsentinel restart
```
- 4 Log in to the authorized requestor server and click Integration. If the data source you want to add is already present, delete it and add it again using one of the IP addresses you specified in Step 2.

Similarly, add authorized requestors using the IP addresses you specified in Step 1.

## Registering Sentinel Appliance for Updates

If you have upgraded the operating system, you must re-register the Sentinel appliance to receive Sentinel and latest operating system updates. You can use your existing registration key to re-register for updates. To register the appliance, see [“Registering for Updates” on page 106](#).

## Updating External Databases for Data Synchronization

Starting from Sentinel 8.x, the size of the `Message (msg)` event field has been increased from 4000 to 8000 characters to accommodate more information in the field.

If you have created a data synchronization policy in previous versions of Sentinel that synchronizes the `Message (msg)` event field to an external database, you must increase the size of the appropriate mapped column in the external database accordingly.

---

**NOTE:** The above step is applicable only if you are upgrading previous versions of Sentinel to 8.x.

---

## Updating Permissions for Users Who Send Data from Other Integrated Products to Sentinel

Sentinel 8.2 SP1 and later provide a new permission, `Send events and attachments`, that enables you to allow only designated users to send events and attachments from Change Guardian or Secure Configuration Manager to Sentinel. When you upgrade to Sentinel 8.2 SP1 and later, Sentinel automatically assigns this permission to users in the Administrator role. For non-administrator users who send events or attachments to Sentinel, you must manually assign this permission. Unless you assign this permission, Sentinel will no longer receive events or attachments from Change Guardian or Secure Configure Manager.

Updating this permission is applicable only if Sentinel is integrated with Change Guardian or Secure Configuration Manager. For more information, see [Creating Roles](#) in the [Sentinel Administration Guide](#).



# 35 Upgrading Sentinel Plug-Ins

The upgrade installations of Sentinel do not upgrade the plug-ins unless a particular plug-in is not compatible with the latest version of Sentinel.

New and updated Sentinel plug-ins, including Solution Packs, are frequently uploaded to the [Sentinel Plug-ins website](#). To get the latest bug fixes, documentation updates, and enhancements for a plug-in, download and install the latest version of the plug-in. For information about installing a plug-in, see the specific plug-in documentation.





# VI Migrating Data from Traditional Storage

Migrating data from Sentinel with traditional storage allows you to leverage your existing Sentinel data and the time you have invested in it. To migrate data from Sentinel with traditional storage, the Sentinel version on both the source and the target Sentinel servers must be same. For example, if you want to migrate data from Sentinel 8.1 (source) to Sentinel 8.2 (target), you must first upgrade Sentinel 8.1 to Sentinel 8.2 and then start with the data migration process.

This section provides information about migrating existing data to the desired data store component.

- ♦ [Chapter 36, “Migrating Data to Scalable Storage,” on page 187](#)
- ♦ [Chapter 37, “Migrating Data to Elasticsearch,” on page 193](#)
- ♦ [Chapter 38, “Migrating Data,” on page 195](#)



# 36

## Migrating Data to Scalable Storage

You may have a single Sentinel server or multiple Sentinel servers with traditional storage. The data migration process you need to follow depends on how you want to set up and maintain your Sentinel deployment.

**Table 36-1** *Data Migration Process for Your Sentinel Deployment*

| Sentinel Deployment                                                                                                                                                              | Migration Process                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| You have a single Sentinel server and you plan to upgrade your existing Sentinel server to scalable storage.                                                                     | Migrate the event data and raw data from traditional storage to scalable storage once you upgrade your Sentinel server and enable scalable storage.<br><br>For more information, see <a href="#">Chapter 38, “Migrating Data,”</a> on page 195.                                                   |
| You have a single Sentinel server with traditional storage and you want to set up another Sentinel server for scalable storage so that you can use all the features in Sentinel. | Use the Backup and Restore utility to migrate data from Sentinel with traditional storage to Sentinel with scalable storage.<br><br>For information about using the backup and restore utility, see <a href="#">“Backing Up and Restoring Data”</a> in the <i>Sentinel Administration Guide</i> . |

---

| Sentinel Deployment                                                                                                                                                                                                                                              | Migration Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>You have a multi-tier setup where you have multiple Sentinel servers and you plan to set up a new Sentinel server or use one of the existing servers for scalable storage. You need to migrate configuration data in addition to event data and raw data.</p> | <p>In a multi-tier setup, you can identify one of the traditional Sentinel servers that has most of your data and then use the backup and restore utility to migrate the data.</p> <p>If you need to back up data from the rest of your Sentinel servers, you must migrate the configuration data, event data, and raw data from those servers using a different approach described later in this section. You must also manually re-create some of the configurations.</p> <p>You cannot use the backup and restore utility to migrate your data from multiple servers because the utility overrides existing data when you restore. For example, if you have already restored data from Server A and then you try to restore data from Server B, this utility overrides the data already restored from Server A.</p> <p>Therefore, to understand the data migrate process involved, follow the instructions provided in the following sections in the same order:</p> <ul style="list-style-type: none"><li>◆ <a href="#">Data You Can Migrate</a></li><li>◆ <a href="#">Migrating Configuration Data</a></li><li>◆ <a href="#">Migrating Data</a></li><li>◆ <a href="#">Migrating Alerts</a></li><li>◆ <a href="#">Updating Sentinel Clients</a></li><li>◆ <a href="#">Importing ESM Configuration</a></li></ul> |

---

## Data You Can Migrate

You can migrate event data, raw data, and some of the configuration data. You must manually re-create the rest of the configuration, which cannot be migrated.

**Table 36-2** Configurations You Can Migrate and Configurations You Need to Re-Create

| Configurations You Can Migrate                                                                                                                                                                                                        | Configurations You Need to Re-Create                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>◆ Correlation rules</li> <li>◆ Actions</li> <li>◆ Maps</li> <li>◆ Filters</li> <li>◆ Threat feeds</li> <li>◆ ESM configuration</li> <li>◆ Alerts except for the Knowledge Base data</li> </ul> | <ul style="list-style-type: none"> <li>◆ Tenants, Roles, Users, and LDAP configuration</li> <li>◆ Event and alert routing rules</li> <li>◆ Data and alert retention policies</li> <li>◆ Dashboards</li> <li>◆ Real-time views</li> <li>◆ Identity information</li> <li>◆ Feeds configuration</li> <li>◆ Action and Integrator plug-in configuration</li> <li>◆ Security configuration</li> </ul> |

## Migrating Configuration Data

Before you migrate event data, you must first migrate the configuration data to the Sentinel target server. You can back up some of the configuration by using Solution Designer and the Export and Import options in Event Source Management (ESM). You must manually re-create the rest of the configuration data, which cannot be backed up or exported.

- ◆ [“Backing Up Data on the Source Server” on page 189](#)
- ◆ [“Restoring Data on the Target Server” on page 190](#)

## Backing Up Data on the Source Server

You must back up the necessary data by using various options in Sentinel.

- ◆ [“Using Solution Packs” on page 189](#)
- ◆ [“Using the Export Configuration Option in ESM” on page 190](#)

## Using Solution Packs

Back up the following configuration on the source server by using Solution Designer:

**Table 36-3** Configuration Data

| Data                                       | Notes                                                                                                                                                              |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> Correlation rules | Create separate controls for each Correlation Engine so that you can migrate the rules separately to specific Correlation Engines.                                 |
| <input type="checkbox"/> Actions           | You can back up only JavaScript actions and not legacy actions such as dynamic list and create incident.                                                           |
| <input type="checkbox"/> Event Enrichment  | Sentinel also backs up the associated maps to the event fields. Therefore, you do not need to re-create the associated maps after restoring event enrichment data. |

| Data                             | Notes                                                                                             |
|----------------------------------|---------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> Filters | Backs up all custom filters.                                                                      |
| <input type="checkbox"/> Feeds   | The Solution Pack backs up only the Feed plug-ins but does not back up the plug-in configuration. |

For information about backing up data in Solution Designer, see [“Creating Solution Packs”](#) in the *Sentinel Administration Guide*.

## Using the Export Configuration Option in ESM

Back up your data collection configuration by using the Export configuration option in ESM. For more information, see [“Exporting Configurations”](#) in the *Sentinel Administration Guide*.

## Restoring Data on the Target Server

- ♦ [“Installing Configuration Data from Solution Pack”](#) on page 190
- ♦ [“Manually Re-Creating Configuration”](#) on page 190

## Installing Configuration Data from Solution Pack

Import the configuration data you backed up on the source server by using Solution Designer. For more information, see [“Installing Content from Solution Packs”](#) in the *Sentinel Administration Guide*.

Rename any duplicate names of objects such as Filters, Actions, and Correlation Rules. By default, all Filters are Public when you import them on the target server. Re-assign the permission for each filter manually.

## Manually Re-Creating Configuration

Apart from the configuration data you imported from the Solution Pack, you must manually re-create all other configurations. For more information about the configurations you need to re-create manually, see [Table 36-2, “Configurations You Can Migrate and Configurations You Need to Re-Create,”](#) on page 189.

## Migrating Event Data and Raw Data

To migrate event data and raw data, see [Migrating Data](#).

## Migrating Alerts

You can use the backup and restore utility to migrate Alerts from the source server to target server. For alerts, this utility restores the events that triggered the alert. However, it does not restore the associated correlation rule and knowledge base information.

Use the following commands to back up and restore Alerts:

For backing up:  
`./backup_util.sh -A`

For restore:  
`./backup_util.sh -m restore -f <backup_file_path>`

For Alerts, you have an option to either override or to append to existing data. Choose the desired option.

Although the above command backs up and restores the Security Intelligence data, you cannot use that data because Security Intelligence is not available in SSDM.

For detailed information about using the backup and restore utility, see [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide*.

## Updating Sentinel Clients

You must update any existing Collector Managers and Correlation Engines configurations such that they start communicating with the target Sentinel server. For more information, see [“Updating Sentinel Clients”](#) in the *Sentinel Administration Guide*.

---

**NOTE:** Although you have already migrated event data from the source server, you must run the data migrate script again to migrate any event data that might have arrived during or after this data migration process. For more information, see [Chapter 38, “Migrating Data,” on page 195](#).

---

## Importing ESM Configuration

Import the data collection configuration you used on the source server by using the Import Configuration option in the ESM user interface. For more information, see [“Importing Configurations”](#) in the *Sentinel Administration Guide*.





# 37 Migrating Data to Elasticsearch

Sentinel stores data in file-based traditional storage and indexes data locally on the Sentinel server by default. When you enable event visualization, Sentinel stores and indexes data in Elasticsearch in addition to file-based traditional storage. The dashboards display only the events processed after you enabled event visualization. To view existing events present in file-based storage, you must migrate data from file-based storage to Elasticsearch. To migrate data to Elasticsearch, see [Chapter 38, “Migrating Data,” on page 195](#).



# 38 Migrating Data

You can use the `data_uploader.sh` script to migrate data to one of the following data storage components:

- ♦ **Kafka:** You can migrate both event and raw data to Kafka. Run the script individually for event data and raw data. The script migrates the data to the Kafka topics.

You can specify customizations such as compressing data during the migration, sending data in batches, and so on. To specify these customizations, create a properties file and add the required properties in key-value format. For example, you can add properties as follows:

```
compression.type=lz4
```

```
batch.size=20000
```

For information about Kafka properties, see [Kafka documentation](#). Set the properties and their values at your discretion because the script does not validate these properties.

---

**NOTE:** Ensure that the Sentinel server is able to resolve all Kafka broker hostnames to valid IP addresses for the entire Kafka cluster. If DNS is not setup to enable this, add the Kafka broker hostnames to the `/etc/hosts` file of the Sentinel server.

---

- ♦ **Elasticsearch:** You can migrate only event data to Elasticsearch. Before you migrate the data, ensure that you have enabled event visualization. For more information, see [“Enabling Event Visualization” on page 123](#).

The script transfers data for the date range (from and to) you specify. When you run the script, it displays the mandatory and optional parameters you should specify to initiate the data migration and also the information about the relevant properties to use for the desired data storage component.

The script must be run as `novell` user. Therefore, ensure that the data directories and any files you specify have appropriate permissions for `novell` user. By default, the script migrates data from primary storage. If you want to migrate data from secondary storage, specify the appropriate path for secondary storage when running the script.

### To migrate data:

- 1 Log in to the Sentinel server as the `novell` user.
- 2 Run the following script:

```
/opt/novell/sentinel/bin/data_uploader.sh
```
- 3 Follow the on-screen instructions and run the script again with the required parameters.

The migrated data will have the retention period as set in the target server.

Once the data migration is done, the script records the status such as partitions migrated successfully, partitions failed to migrate, number of events migrated, and so on. For partitions with previous day and current day's date, the data transfer status will show IN\_PROGRESS considering events that may come in late.

Run the script again in scenarios where the data migration did not complete successfully or where the data migration status for partitions still indicate IN\_PROGRESS. When you re-run the script, it first checks the status file to understand the partitions that were already migrated and then continues to migrate only the remaining ones. The script maintains the logs in the `/var/opt/novell/sentinel/log/data_uploader.log` directory for troubleshooting purposes.

# VII Deploying Sentinel for High Availability

This section provides information about how to install Sentinel in an active-passive high availability mode, which allows Sentinel to fail over to a redundant cluster node in case of hardware or software failure. For more information on implementing high availability and disaster recovery in your Sentinel environment, contact [Technical Support](#).

---

**NOTE:** High availability (HA) configuration is supported only on the Sentinel server. However, Collector Managers and Correlation Engines can still communicate with the Sentinel HA server.

---

- ◆ [Chapter 39, “Concepts,” on page 199](#)
- ◆ [Chapter 40, “System Requirements,” on page 201](#)
- ◆ [Chapter 41, “Installation and Configuration,” on page 203](#)
- ◆ [Chapter 42, “Configuring Sentinel HA as SSDM,” on page 219](#)
- ◆ [Chapter 43, “Upgrading Sentinel in High Availability,” on page 221](#)
- ◆ [Chapter 44, “Backup and Recovery,” on page 233](#)



# 39 Concepts

High availability refers to a design methodology that is intended to keep a system available for use as much as is practicable. The intent is to minimize the causes of downtime such as system failures and maintenance, and to minimize the time it will take to detect and recover from downtime events that do occur. In practice, automated means of detecting and recovering from downtime events quickly become necessary as higher levels of availability must be achieved.

For more information about high availability, see the [SUSE High Availability Guide](#).

- ♦ “External Systems” on page 199
- ♦ “Shared Storage” on page 199
- ♦ “Service Monitoring” on page 200
- ♦ “Fencing” on page 200

## External Systems

Sentinel is a complex multi-tier application that depends upon and provides a wide variety of services. Additionally, it integrates with multiple external third-party systems for data collection, data sharing, and incident remediation. Most HA solutions allow implementors to declare dependencies between the services that should be highly available, but this only applies to services running on the cluster itself. Systems external to Sentinel such as event sources must be configured separately to be as available as required by the organization, and must also be configured to properly handle situations where Sentinel is unavailable for some period of time such as a failover event. If access rights are tightly restricted, for example if authenticated sessions are used to send and/or receive data between the third-party system and Sentinel, the third-party system must be configured to accept sessions from or initiate sessions to any cluster node (Sentinel should be configured with a virtual IP address for this purpose).

## Shared Storage

All HA clusters require some form of shared storage so that application data can be quickly moved from one cluster node to another, in case of a failure of the originating node. The storage itself should be highly available; this is usually achieved by using Storage Area Network (SAN) technology connected to the cluster nodes using a Fibre Channel network. Other systems use Network Attached Storage (NAS), iSCSI, or other technologies that allow for remote mounting of shared storage. The fundamental requirement of the shared storage is that the cluster can cleanly move the storage from a failed cluster node to a new cluster node.

There are two basic approaches that Sentinel can use for the shared storage. The first locates all components (application binaries, configuration, and event data) on the shared storage. On failover, the storage is unmounted from the primary node and moved to the backup node; which loads the entire application and configuration from the shared storage. The second approach stores the event data on shared storage, but the application binaries and configuration reside on each cluster node. On failover, only the event data is moved to the backup node.

Each approach has benefits and disadvantages, but the second approach allows the Sentinel installation to use standard FHS-compliant install paths, allows for verification of the RPM packaging, and also allows for warm patching and reconfiguration to minimize downtime.

This solution will guide you through an example of the process of installing to a cluster that uses iSCSI shared storage and locates the application binaries/configuration on each cluster node.

## Service Monitoring

A key component of any highly available environment is a reliable, consistent way to monitor the resource(s) that should be highly available, along with any resource(s) that they depend on. The SLE HAE uses a component called a Resource Agent to perform this monitoring - the Resource Agent's job is to provide the status for each resource, plus (when asked) to start or stop that resource.

Resource Agents must provide a reliable status for monitored resources in order to prevent unnecessary downtime. False positives (when a resource is deemed to have failed, but would in fact recover on its own) can cause service migration (and related downtime) when it is not actually necessary, and false negatives (when the Resource Agent reports that a resource is functioning when in fact it is not operating properly) can prevent proper use of the service. On the other hand, external monitoring of a service can be quite difficult - a web service port might respond to a simple ping, for example, but may not provide correct data when a real query is issued. In many cases, self-test functionality must be built into the service itself to provide a truly accurate measurement.

This solution provides a basic OCF Resource Agent for Sentinel that can monitor for major hardware, operating system, or Sentinel system failure. At this time the external monitoring capabilities for Sentinel are based on IP port probes, and there is some potential for false positive and false negative readings. We plan to improve both Sentinel and the Resource Agent over time to improve the accuracy of this component.

## Fencing

Within an HA cluster, critical services are constantly monitored and restarted automatically on other nodes in the case of failure. This automation can introduce problems, however, if some communications problem occurs with the primary node; although the service running on that node appears to be down, it in fact continues to run and write data to the shared storage. In this case, starting a new set of services on a backup node could easily cause data corruption.

Clusters use a variety of techniques collectively called fencing to prevent this from happening, including Split Brain Detection (SBD) and Shoot The Other Node In The Head (STONITH). The primary goal is to prevent data corruption on the shared storage.



# 40 System Requirements

When allocating cluster resources to support a high availability (HA) installation, consider the following requirements:

- ❑ (Conditional) For HA appliance installations, ensure that the Sentinel HA appliance with a valid license is available. The Sentinel HA appliance is an ISO appliance that includes the following packages:
  - ◆ Operating system: SLES 12 SP4
  - ◆ SLES High Availability Extension (SLES HAE) package
  - ◆ Sentinel software (including HA rpm)
- ❑ (Conditional) For traditional HA installations, ensure that the following are available:
  - ◆ Operating system: SLES 12 SP3 or later
  - ◆ SLES HAE ISO image with valid licenses
  - ◆ Sentinel installer (TAR file)
- ❑ (Conditional) If you are using the SLES operating system with kernel version 3.0.101 or later, you must manually load the watchdog driver on the computer. To find the appropriate watchdog driver for your computer hardware, contact your hardware vendor. To load the watchdog driver, perform the following:
  1. In the command prompt, run the following command to load the watchdog driver in the current session:

```
/sbin/modprobe -v --ignore-install <watchdog driver name>
```
  2. In the `/etc/init.d/boot.local` file, add the following line to ensure that the computer automatically loads the watchdog driver at every boot time:

```
/sbin/modprobe -v --ignore-install <watchdog driver name>
```
- ❑ Ensure that each cluster node that hosts the Sentinel services meet the requirements specified in [Chapter 5, “Meeting System Requirements,” on page 37](#).
- ❑ Ensure that sufficient shared storage is available for the Sentinel data and application.
- ❑ Ensure that you use a virtual IP address for the services that can be migrated from node to node on failover.
- ❑ Ensure that your shared storage device meets the performance and size characteristics requirements specified in [Chapter 5, “Meeting System Requirements,” on page 37](#). Use a standard SLES virtual machine configured with iSCSI Targets as shared storage.

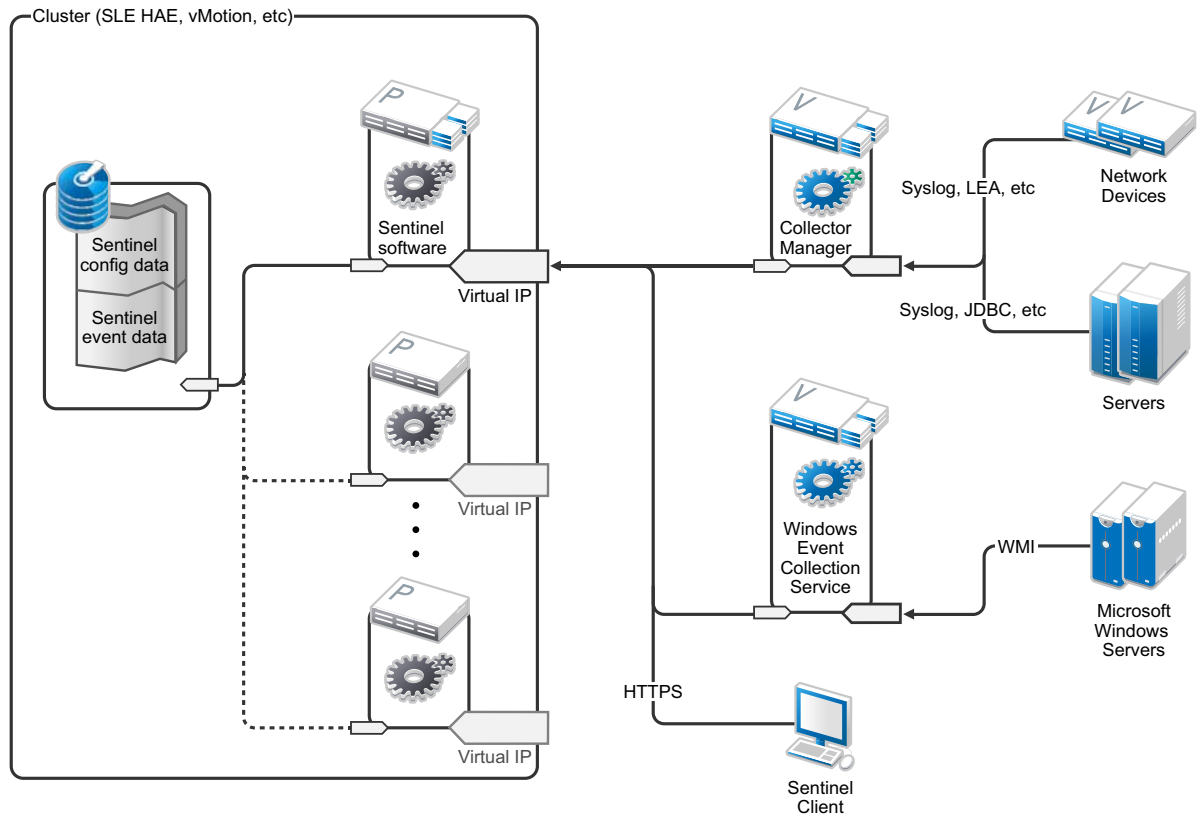
For iSCSI, you should use the largest Message Transfer Unit (MTU) supported by your hardware. Larger MTUs benefits the storage performance. Sentinel might experience issues if latency and bandwidth to storage is slower than recommended.
- ❑ Ensure that you have a minimum of two cluster nodes that meet the resource requirements for running Sentinel in the customer environment. Two SLES virtual machines is recommended.

- ❑ Ensure you create a method for the cluster nodes to communicate with the shared storage, such as FibreChannel for a SAN. Use a dedicated IP address to connect to the iSCSI Target.
- ❑ Ensure you have a virtual IP address that can be migrated from one node to another node in a cluster to serve as the external IP address for Sentinel.
- ❑ Ensure you have at least one IP address per cluster node for internal cluster communications. You can use a simple unicast IP address, but multicast is preferred for production environments.

# 41 Installation and Configuration

This chapter provides the steps for installing and configuring Sentinel in a high availability (HA) environment.

The following diagram represents an active-passive HA architecture.



- ◆ [“Initial Setup” on page 204](#)
- ◆ [“Shared Storage Setup” on page 205](#)
- ◆ [“Sentinel Installation” on page 209](#)
- ◆ [“Cluster Installation” on page 212](#)
- ◆ [“Cluster Configuration” on page 213](#)
- ◆ [“Resource Configuration” on page 216](#)
- ◆ [“Secondary Storage Configuration” on page 217](#)

# Initial Setup

Configure the computer hardware, network hardware, storage hardware, operating systems, user accounts, and other basic system resources per the documented requirements for Sentinel and local customer requirements. Test the systems to ensure proper function and stability.

Use the following checklist to guide you through initial setup and configuration.

|                          | Checklist Items                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | The CPU, RAM, and disk space characteristics for each cluster node must meet the system requirements defined in <a href="#">Chapter 5, “Meeting System Requirements,” on page 37</a> based on the expected event rate.                                                                                                                                                                                                                                                    |
| <input type="checkbox"/> | The disk space and I/O characteristics for the storage nodes must meet the system requirements defined in <a href="#">Chapter 5, “Meeting System Requirements,” on page 37</a> based on the expected event rate and data retention policies for primary and secondary storage.                                                                                                                                                                                            |
| <input type="checkbox"/> | If you want to configure the operating system firewalls to restrict access to Sentinel and the cluster, refer to <a href="#">Chapter 8, “Ports Used,” on page 61</a> for details of which ports must be available depending on your local configuration and the sources that will be sending event data.                                                                                                                                                                  |
| <input type="checkbox"/> | Ensure that all cluster nodes are time-synchronized. You can use NTP or a similar technology for this purpose.                                                                                                                                                                                                                                                                                                                                                            |
| <input type="checkbox"/> | <ul style="list-style-type: none"><li>◆ The cluster requires reliable host name resolution. Enter all internal cluster host names into the <code>/etc/hosts</code> file to ensure cluster continuity in case of DNS failure.</li><li>◆ Ensure that you do not assign a host name to a loopback IP address.</li><li>◆ When configuring host name and domain name while installing the operating system, deselect <a href="#">Assign Hostname to Loopback IP</a>.</li></ul> |

You can use the following configuration:

- ◆ (Conditional) For traditional HA installations:
  - ◆ Two cluster node VMs running SLES 11 SP4 or SLES 12 SP1 or later.
  - ◆ (Conditional) You can install X Windows if you require GUI configuration. Set the boot scripts to start without X (runlevel 3), so you can start them only when needed.
- ◆ (Conditional) For HA appliance installations: Two HA ISO appliance based cluster node virtual machines. For information about installing the HA ISO appliance, see [“Installing Sentinel” on page 102](#).
- ◆ The nodes will have one NIC for external access and one for iSCSI communications.
- ◆ Configure the external NICs with IP addresses that allow for remote access through SSH or similar. For this example, we will use 172.16.0.1 (node01) and 172.16.0.2 (node02).
- ◆ Each node should have sufficient disk for the operating system, Sentinel binaries and configuration data, cluster software, temp space, and so forth. See the SLES and SLES HAE system requirements, and Sentinel application requirements.
- ◆ One virtual machine running SLES 11 SP4 or SLES 12 SP1 or later that is configured with iSCSI Targets for shared storage
  - ◆ (Conditional) You can install X Windows if you require GUI configuration. Set the boot scripts to start without X (runlevel 3), so you can start them only when needed.

- ◆ The system will have two NICs: one for external access and one for iSCSI communications.
- ◆ Configure the external NIC with an IP address that allows for remote access using SSH or similar. For example, 172.16.0.3 (storage03).
- ◆ The system should have sufficient space for the operating system, temp space, a large volume for shared storage to hold Sentinel data, and a small amount of space for an SBD partition. See the SLES system requirements, and Sentinel event data storage requirements.

---

**NOTE:** In a production cluster, you can use internal, non-routable IP addresses on separate NICs (possibly a couple, for redundancy) for internal cluster communications.

---

## Shared Storage Setup

Set up your shared storage and ensure that you can mount it on each cluster node. If you are using FibreChannel and a SAN, you might need to provide physical connections as well as additional configuration. Sentinel uses this shared storage to store the databases and event data. Ensure that the shared storage is appropriately sized accordingly based on the expected event rate and data retention policies.

Consider the following example of a shared storage setup:

A typical implementation might use a fast SAN attached using FibreChannel to all the cluster nodes, with a large RAID array to store the local event data. A separate NAS or iSCSI node might be used for the slower secondary storage. As long as the cluster node can mount the primary storage as a normal block device, it can be used by the solution. The secondary storage can also be mounted as a block device, or could be an NFS or CIFS volume.

---

**NOTE:** Configure your shared storage and test mounting it on each cluster node. However, the cluster configuration will handle the actual mount of the storage.

---

Perform the following procedure to create iSCSI Targets hosted by a SLES virtual machine:

- 1 Connect to `storage03`, the virtual machine you created during [Initial Setup](#), and start a console session.
- 2 Run the following command to create a blank file of any desired size for Sentinel primary storage:

```
dd if=/dev/zero of=/localdata count=<file size> bs=<bit size>
```

For example, run the following command to create a 20GB file filled with zeros copied from the `/dev/zero` pseudo-device:

```
dd if=/dev/zero of=/localdata count=20480000 bs=1024
```

- 3 Repeat steps 1 and 2 to create a file for secondary storage in the same way.

For example, run the following command for the secondary storage:

```
dd if=/dev/zero of=/networkdata count=20480000 bs=1024
```

---

**NOTE:** For this example you created two files of the same size and performance characteristics to represent the two disks. For a production deployment, you can create the primary storage on a fast SAN and the secondary storage on a slower iSCSI, NFS, or CIFS volume.

---

Perform the steps provided in the following sections to configure iSCSI target and initiator devices:

- ♦ “[Configuring iSCSI Targets](#)” on page 206
- ♦ “[Configuring iSCSI Initiators](#)” on page 207

## Configuring iSCSI Targets

Perform the following procedure to configure `localdata` and `networkdata` files as iSCSI Targets.

For more information about configuring iSCSI targets, see [Creating iSCSI Targets with YaST](#) in the SUSE documentation.

- 1 Run YaST from the command line (or use the Graphical User Interface, if preferred): `/sbin/yast`
- 2 Select **Network Devices** > **Network Settings**.
- 3 Ensure that the **Overview** tab is selected.
- 4 Select the secondary NIC from the displayed list, then tab forward to Edit and press `Enter`.
- 5 On the **Address** tab, assign a static IP address of 10.0.0.3. This will be the internal iSCSI communications IP address.
- 6 Click **Next**, then click **OK**.
- 7 (Conditional) On the main screen:
  - ♦ If you are using SLES 12 SP1 and later, select **Network Services** > **iSCSI LIO Target**.

---

**NOTE:** If you do not find this option, go to **Software** > **Software Management** > **iSCSI LIO Server** and install the iSCSI LIO package.

---

- 8 (Conditional) If prompted, install the required software:
  - ♦ For SLES 12 SP1 and later: `iscsiliotarget RPM`
- 9 (Conditional) If you are using SLES 12, perform the following steps on all the nodes in the cluster:
  - 9a Run the following command to open the file which contains the iSCSI initiator name:

```
cat /etc/iscsi/initiatorname.iscsi
```
  - 9b Note the initiator name which will be used for configuring iSCSI initiators:  
For example:

```
InitiatorName=iqn.1996-04.de.suse:01:441d6988994
```

These initiator names will be used while configuring iSCSI Target Client Setup.
- 10 Click **Service**, select the **When Booting** option to ensure that the service starts when the operating system boots.
- 11 Select the **Global** tab, deselect **No Authentication** to enable authentication, and then specify the necessary credentials for incoming and outgoing authentication.

The **No Authentication** option is enabled by default. However, you should enable authentication to ensure that the configuration is secure.

---

**NOTE:** We recommend that you use the different password for the iSCSI target and initiator.

---

- 12 Click **Targets** and then click **Add** to add a new target.

The iSCSI Target will auto-generate an ID and then present an empty list of LUNs (drives) that are available.

- 13 Click **Add** to add a new LUN.

- 14 Leave the LUN number as 0, then browse in the **Path** dialog (under Type=fileio) and select the `/localdata` file that you created. If you have a dedicated disk for storage, specify a block device, such as `/dev/sdc`.

- 15 Repeat steps 13 and 14, and add LUN 1 and select `/networkdata` this time.

- 16 Leave the other options at their defaults, and click **Next**.

- 17 (Conditional) If you are using SLES 12, click **Add**. When prompted for Client Name, specify the initiator name you have copied in Step 9. Repeat this step to add all the client names, by specifying the initiator names.

The list of client names will be displayed in the Client List.

You do not have to add the client initiator name for SLES 15 and later.

- 18 (Conditional) If you have enabled authentication in Step 11, provide authentication credentials.

Select a client, select **Edit Auth > Incoming Authentication**, and specify the user name and password. Repeat this for all the clients.

- 19 Click **Next** again to select the default authentication options, then **Finish** to exit the configuration. Accept if prompted to restart iSCSI.

- 20 Exit YaST.

---

**NOTE:** This procedure exposes two iSCSI Targets on the server at IP address 10.0.0.3. At each cluster node, ensure that it can mount the local data shared storage device.

---

## Configuring iSCSI Initiators

Perform the following procedure to format the iSCSI initiator devices.

For more information about configuring iSCSI initiators, see [Configuring the iSCSI Initiator](#) in the SUSE documentation.

- 1 Connect to one of the cluster nodes (node01) and start YaST.
- 2 Select **Network Devices > Network Settings**.
- 3 Ensure that the **Overview** tab is selected.
- 4 Select the secondary NIC from the displayed list, then tab forward to Edit and press Enter.
- 5 Click **Address**, assign a static IP address of 10.0.0.1. This will be the internal iSCSI communications IP address.
- 6 Select **Next**, then click **OK**.
- 7 Click **Network Services > iSCSI Initiator**.

- 8 If prompted, install the required software (`iscsiclient` RPM).
- 9 Click **Service**, select **When Booting** to ensure the iSCSI service is started on boot.
- 10 Click **Discovered Targets**, and select **Discovery**.
- 11 Specify the iSCSI Target IP address (10.0.0.3).  
(Conditional) If you have enabled authentication in Step 11 in [“Configuring iSCSI Targets” on page 206](#), deselect **No Authentication**. In the **Outgoing Authentication** field, enter the user name and the password you configured during iSCSI target configuration.  
Click **Next**.
- 12 Select the discovered iSCSI Target with the IP address 10.0.0.3 and then select **Log In**.
- 13 Perform the following steps:
  - 13a Switch to Automatic in the **Startup** drop-down menu.
  - 13b (Conditional) If you have enabled authentication, deselect **No Authentication**.  
The user name and the password you have specified in Step 11 should be displayed in the **Outgoing Authentication** section. If these credentials are not displayed, enter the credentials in this section.
  - 13c Click **Next**.
- 14 Switch to the **Connected Targets** tab to ensure that we are connected to the target.
- 15 Exit the configuration. This should have mounted the iSCSI Targets as block devices on the cluster node.
- 16 In the YaST main menu, select **System > Partitioner**.
- 17 In the System View, you should see new hard disks of the following types (such as `/dev/sdb` and `/dev/sdc`) in the list:
  - ◆ In SLES 11 SP4: IET-VIRTUAL-DISK
  - ◆ In SLES 12 SP1 or later: LIO-ORG-FILEIOTab over to the first one in the list (which should be the primary storage), select that disk, then press Enter.
- 18 Select **Add** to add a new partition to the empty disk. Format the disk as a primary partition, but do not mount it. Ensure that the **Do not mount partition** option is selected.
- 19 Select **Next**, and then **Finish** after reviewing the changes that will be made.  
The formatted disk (such as `/dev/sdb1`) should be ready now. It is referred to as `/dev/<SHARED1>` in the following steps of this procedure.
- 20 Go to **Partitioner** again and repeat the partitioning/formatting process (steps 16-19) for `/dev/sdc` or whichever block device corresponds to the secondary storage. This should result in a `/dev/sdc1` partition or similar formatted disk (referred to as `/dev/<NETWORK1>` below).
- 21 Exit YaST.
- 22 (Conditional) If you are performing a traditional HA installation, create a mount point and test mounting the local partition as follows (the exact device name might depend on the specific implementation):

```
mkdir /var/opt/novell
mount /dev/<SHARED1> /var/opt/novell
```



You should be able to create files on the new partition and see the files wherever the partition is mounted.

- 23** (Conditional) If you are performing a traditional HA installation, to unmount:

```
umount /var/opt/novell
```

- 24** (Conditional) For HA appliance installations, repeat steps 1-15 to ensure that each cluster node can mount the local shared storage. Replace the node IP address in step 5 with a different IP address for each cluster node.
- 25** (Conditional) For traditional HA installations, repeat steps 1-15, 22, and 23 to ensure that each cluster node can mount the local shared storage. Replace the node IP address in step 6 with a different IP address for each cluster node.

## Sentinel Installation

There are two options to install Sentinel: install every part of Sentinel onto the shared storage using the `--location` option to redirect the Sentinel installation to the location where you have mounted the shared storage or install only the variable application data on the shared storage.

Install Sentinel to each cluster node that can host it. After you install Sentinel the first time, you must perform a complete installation including the application binaries, configuration, and all the data stores. For subsequent installations on the other cluster nodes, you will only install the application. The Sentinel data will be available once you have mounted the shared storage.

### First Node Installation

- ♦ [“Traditional HA Installation” on page 209](#)
- ♦ [“Sentinel HA Appliance Installation” on page 210](#)

### Traditional HA Installation

- 1 Connect to one of the cluster nodes (node01) and open a console window.
- 2 Download the Sentinel installer (a tar.gz file) and store it in `/tmp` on the cluster node.
- 3 Perform the following steps to start the installation:
  - 3a Execute the following commands:

```
mount /dev/<SHARED1> /var/opt/novell
cd /tmp
tar -xvzf sentinel_server*.tar.gz
cd sentinel_server*
./install-sentinel --record-unattended=/tmp/install.props
```
  - 3b Specify 2 to select Custom Configuration when prompted to select the configuration method.
- 4 Run through the installation, configuring the product as appropriate.
- 5 Start Sentinel and test the basic functions. You can use the standard external cluster node IP address to access the product.

- 6 Shut down Sentinel and dismount the shared storage using the following commands:

```
rscsentinel stop
umount /var/opt/novell
```

This step removes the autostart scripts so that the cluster can manage the product.

```
cd /
insserv -r sentinel
```

## Sentinel HA Appliance Installation

The Sentinel HA appliance includes the Sentinel software that is already installed and configured. To configure the Sentinel software for HA, perform the following steps:

- 1 Connect to one of the cluster nodes (node01) and open a console window.
- 2 Navigate to the following directory:

```
cd /opt/novell/sentinel/setup
```

- 3 Record the configuration:

- 3a Execute the following command:

```
./configure.sh --record-unattended=/tmp/install.props --no-start
```

This step records the configuration in the file `install.props`, which is required to configure the cluster resources using the `install-resources.sh` script.

- 3b Specify 2 to select Custom Configuration when prompted to select the configuration method.

- 3c When prompted for password, specify 2 to enter a new password.

If you specify 1, the `install.props` file does not store the password.

- 4 Shut down Sentinel using the following command:

```
rscsentinel stop
```

This step removes the autostart scripts so that the cluster can manage the product.

```
insserv -r sentinel
```

- 5 Move the Sentinel data folder to the shared storage using the following commands. This movement allows the nodes to utilize the Sentinel data folder through shared storage.

```
mkdir -p /tmp/new
mount /dev/<SHARED1> /tmp/new
mv /var/opt/novell/sentinel/* /tmp/new
umount /tmp/new/
```

- 6 Verify the movement of the Sentinel data folder to the shared storage using the following commands:

```
mount /dev/<SHARED1> /var/opt/novell/sentinel
umount /var/opt/novell/sentinel
```

## Configuring the Appliance

Perform the following steps to configure the appliance with SMT:

- 1 Enable the appliance repositories by running the following commands in the SMT server:

```
smt-repos -e Sentinel-Server-HA-8-OS-Updates sle-12-x86_64
smt-repos -e Sentinel-Server-HA-8-Prod-Updates sle-12-x86_64
```

- 2 Configure the appliance with SMT by performing the steps in the [“Configuring Clients to Use SMT”](#) section of the [SMT documentation](#).

## Subsequent Node Installation

- ♦ [“Traditional HA Installation”](#) on page 211
- ♦ [“Sentinel HA Appliance Installation”](#) on page 212

Repeat the installation on other nodes:

The initial Sentinel installer creates a user account for use by the product, which uses the next available user ID at the time of the install. Subsequent installs in unattended mode will attempt to use the same user ID for account creation, but the possibility for conflicts (if the cluster nodes are not identical at the time of the install) does exist. It is highly recommended that you do one of the following:

- ♦ Synchronize the user account database across cluster nodes (manually through LDAP or similar), making sure that the sync happens before subsequent installs. In this case the installer will detect the presence of the user account and use the existing one.
- ♦ Watch the output of the subsequent unattended installs - a warning will be issued if the user account could not be created with the same user ID.

## Traditional HA Installation

- 1 Connect to each additional cluster node (node02) and open a console window.
- 2 Execute the following commands:

```
cd /tmp
scp root@node01:/tmp/sentinel_server*.tar.gz .
scp root@node01:/tmp/install.props .
tar -xvzf sentinel_server*.tar.gz
cd sentinel_server*
./install-sentinel --no-start --cluster-node --unattended=/tmp/
install.props
insserv -r sentinel
```

## Sentinel HA Appliance Installation

- 1 Connect to each additional cluster node (node02) and open a console window.
- 2 Execute the following command:

```
insserv -r sentinel
```

- 3 Stop Sentinel services.

```
rcsentinel stop
```

- 4 Remove Sentinel directory.

```
rm -rf /var/opt/novell/sentinel/*
```

At the end of this process, Sentinel should be installed on all nodes, but it will likely not work correctly on any but the first node until various keys are synchronized, which will happen when we configure the cluster resources.

## Cluster Installation

You must install the cluster software only for traditional high availability (HA) installations. The Sentinel HA appliance includes the cluster software and does not require manual installation.

**Use the following procedure to set up the SLES High Availability Extension with a Sentinel-specific Resource Agents overlay:**

- 1 Install the cluster software on each node.
- 2 Register each cluster node with the cluster manager.
- 3 Verify that each cluster node appears in the cluster management console.

---

**NOTE:** The OCF Resource Agent for Sentinel is a simple shell script that runs a variety of checks to verify if Sentinel is functional. If you do not use the OCF Resource Agent to monitor Sentinel, you must develop a similar monitoring solution for the local cluster environment. To develop your own, review the existing Resource Agent, stored in the `Sentinelha.rpm` file in the Sentinel download package.

---

- 4 Install the core SLE HAE software according to the [SLE HAE Documentation](#). For information about installing SLES add-ons, see the [Deployment Guide](#).
- 5 Repeat step 4 on all cluster nodes. The add-on will install the core cluster management and communications software, as well as many Resource Agents that are used to monitor cluster resources.
- 6 Install an additional RPM to provide the additional Sentinel-specific cluster Resource Agents. The HA RPM can be found in the `novell-Sentinelha-<Sentinel_version>*.rpm` file, stored in the default Sentinel download, which you unpacked to install the product.
- 7 On each cluster node, copy the `novell-Sentinelha-<Sentinel_version>*.rpm` file into the `/tmp` directory, then run the following commands:

```
cd /tmp
```

```
rpm -i novell-Sentinelha-<Sentinel_version>*.rpm
```

# Cluster Configuration

You must configure the cluster software to register each cluster node as a member of the cluster. As part of this configuration, you can also set up fencing and Shoot The Other Node In The Head (STONITH) resources to ensure cluster consistency.

---

**IMPORTANT:** The procedures in this section use `rcopenais` and `openais` commands, which work only with SLES 11 SP4. For SLES 12 SP2 and later, use the `systemctl pacemaker.service` command.

For example, for the `/etc/rc.d/openais start` command, use the `systemctl start pacemaker.service` command.

---

## Use the following procedure for cluster configuration:

For this solution you must use private IP addresses for internal cluster communications, and use unicast to minimize the need to request a multicast address from a network administrator. You must also use an iSCSI Target configured on the same SLES virtual machine that hosts the shared storage to serve as a Split Brain Detection (SBD) device for fencing purposes.

### SBD Setup

- 1 Connect to `storage03` and start a console session. Run the following command to create a blank file of any desired size:

```
dd if=/dev/zero of=/sbd count=<file size> bs=<bit size>
```

For example, run the following command to create a 1MB file filled with zeros copied from the `/dev/zero` pseudo-device:

```
dd if=/dev/zero of=/sbd count=1024 bs=1024
```

- 2 Run YaST from the command line or the Graphical User Interface: `/sbin/yast`
- 3 Select **Network Services** > **iSCSI Target**.
- 4 Click **Targets** and select the existing target.
- 5 Select **Edit**. The UI will present a list of LUNs (drives) that are available.
- 6 Select **Add** to add a new LUN.
- 7 Leave the LUN number as 2. Browse in the **Path** dialog and select the `/sbd` file that you created.
- 8 Leave the other options at their defaults, then select **OK** then **Next**, then click **Next** again to select the default authentication options.
- 9 Click **Finish** to exit the configuration. Restart the services if needed. Exit YaST.

---

**NOTE:** The following steps require that each cluster node be able to resolve the hostname of all other cluster nodes (the file sync service `csync2` will fail if this is not the case). If DNS is not set up or available, add entries for each host to the `/etc/hosts` file that list each IP address and its hostname (as reported by the `hostname` command). Also, ensure that you do not assign a hostname to a loopback IP address.

---

Perform the following steps to expose an iSCSI Target for the SBD device on the server at IP address 10.0.0.3 (`storage03`).

## Node Configuration

Connect to a cluster node (node01) and open a console:

- 1 Run YaST.
- 2 Open **Network Services** > **iSCSI Initiator**.
- 3 Select **Connected Targets**, then the iSCSI Target you configured above.
- 4 Select the **Log Out** option and log out of the Target.
- 5 Switch to the **Discovered Targets** tab, select the **Target**, and log back in to refresh the list of devices (leave the **automatic** startup option and deselect **No Authentication**).
- 6 Select **OK** to exit the iSCSI Initiator tool.
- 7 Open **System** > **Partitioner** and identify the SBD device as the 1MB IET-VIRTUAL-DISK. It will be listed as `/dev/sdd` or similar - note which one.
- 8 Exit YaST.
- 9 Execute the command `ls -l /dev/disk/by-id/` and note the device ID that is linked to the device name you located above.
- 10 (Conditional) Execute one of the following commands:
  - ♦ If you are using SLES 11 SP4:  
`sleha-init`
  - ♦ If you are using SLES 12 SP1 or later:  
`ha-cluster-init`
- 11 When prompted for the network address to bind to, specify the external NIC IP address (172.16.0.1).
- 12 Accept the default multicast address and port. We will override this later.
- 13 Enter `y` to enable SBD, then specify `/dev/disk/by-id/<device id>`, where `<device id>` is the ID you located above (you can use `Tab` to auto-complete the path).
- 14 (Conditional) Enter `N` when prompted with the following:  

```
Do you wish to configure an administration IP? [y/N]
```

To configure an administration IP address, provide the virtual IP address during “[Resource Configuration](#)” on page 216
- 15 Complete the wizard and make sure no errors are reported.
- 16 Start YaST.
- 17 Select **High Availability** > **Cluster** (or just **Cluster** on some systems).
- 18 In the box at left, ensure **Communication Channels** is selected.
- 19 Tab over to the top line of the configuration, and change the **udp** selection to **udpu** (this disables multicast and selects unicast).
- 20 Select to **Add a Member Address** and specify this node (172.16.0.1), then repeat and add the other cluster node(s): 172.16.0.2.
- 21 (Conditional) If you have not enabled authentication, then select **Security** from the left panel, clear **Enable Security Auth**.
- 22 Select **Finish** to complete the configuration.

**23** Exit YaST.

**24** Run the command `/etc/rc.d/openais restart` to restart the cluster services with the new sync protocol.

Connect to each additional cluster node (node02) and open a console:

**1** Run YaST.

**2** Open **Network Services > iSCSI Initiator**.

**3** Select **Connected Targets**, then the iSCSI Target you configured above.

**4** Select the **Log Out** option and log out of the Target.

**5** Switch to the **Discovered Targets** tab, select the **Target**, and log back in to refresh the list of devices (leave the **automatic** startup option and deselect **No Authentication**).

**6** Select **OK** to exit the iSCSI Initiator tool.

**7** (Conditional) Execute one of the following commands:

- ♦ If you are using SLES 11 SP4:

```
sleha-join
```

- ♦ If you are using SLES 12 SP1 or later:

```
ha-cluster-join
```

**8** Enter the IP address of the first cluster node.

(Conditional) If the cluster does not start correctly, perform the following steps:

**1** Run the command `crm status` to check if the nodes are joined. If the nodes are not joined, restart all the nodes in the cluster.

**2** Manually copy the `/etc/corosync/corosync.conf` file from node01 to node02, or run the `csync2 -x -v` on node01, or manually set the cluster up on node02 through YaST.

**3** (Conditional) If the `csync2 -x -v` command you run in Step 1 fails to synchronize all the files, perform the following procedure:

**3a** Clear the `csync2` database in the `/var/lib/csync2` directory on all the nodes.

**3b** On all the nodes, update the `csync2` database to match the filesystem without marking anything as needing to be synchronized to other servers:

```
csync2 -cIr /
```

**3c** On the active node, perform the following:

**3c1** Find all the differences between active and passive nodes, and mark those differences for synchronization:

```
csync2 -TUXI
```

**3c2** Reset the database to force the active node to override any conflicts:

```
csync2 -fr /
```

**3c3** Start synchronization to all the other nodes:

```
csync2 -xr /
```

**3d** On all the nodes, verify that all the files are synchronized:

```
csync2 -T
```

This command will list only the files that are not synchronized.

4 Run the following command on node02:

**For SLES 11 SP4:**

```
/etc/rc.d/openais start
```

**For SLES 12 SP1 and later:**

```
systemctl start pacemaker.service
```

(Conditional) If the `xinetd` service does not properly add the new `csync2` service, the script will not function properly. The `xinetd` service is required so that the other node can sync the cluster configuration files down to this node. If you see errors like `csync2 run failed`, you may have this problem.

To resolve this issue, execute the `kill -HUP `cat /var/run/xinetd.init.pid`` command and then re-run the `sleha-join` script.

5 Run `crm_mon` on each cluster node to verify that the cluster is running properly. You can also use 'hawk', the web console, to verify the cluster. The default login name is `ishacluster` and the password is `linux`.

(Conditional) Depending on your environment, perform the following tasks to modify additional parameters:

1 To ensure that in a single node failure in your two-node cluster does not unexpectedly stop the entire cluster, set the global cluster option `no-quorum-policy` to `ignore`:

```
crm configure property no-quorum-policy=ignore
```

---

**NOTE:** If your cluster contains more than two nodes, do not set this option.

---

## Resource Configuration

Resource Agents are provided by default with SLE HAE. If you do not want to use SLE HAE, you need to monitor these additional resources using an alternate technology:

- ♦ A filesystem resource corresponding to the shared storage that the software uses.
- ♦ An IP address resource corresponding to the virtual IP address by which the services will be accessed.
- ♦ The PostgreSQL database software that stores configuration and event metadata.

**Use the following procedure for resource configuration:**

The `crm` script helps you for cluster configuration. The script pulls relevant configuration variables from the unattended setup file generated as part of the Sentinel installation. If you did not generate the setup file, or you wish to change the configuration of the resources, you can use the following procedure to edit the script accordingly.

1 Connect to the original node on which you installed Sentinel.

---

**NOTE:** This must be the node on which you ran the full Sentinel install.

---

2 Edit the script so that it appears as follows, where `<SHARED1>` is the shared volume you created previously:

```
mount /dev/<SHARED1> /var/opt/novell
```



```
cd /usr/lib/ocf/resource.d/novell
```

```
./install-resources.sh
```

- 3 (Conditional) You might have issues with the new resources coming up in the cluster. If you experience this issue, run the following command on node02:

**For SLES 11 SP4:**

```
/etc/rc.d/openais start
```

**For SLES 12 SP1:**

```
systemctl start pacemaker.service
```

- 4 The `install-resources.sh` script will prompt you for a couple values, namely the virtual IP address that you would like people to use to access Sentinel and the device name of the shared storage, and then will auto-create the required cluster resources. Note that the script requires the shared volume to already be mounted, and also requires the unattended installation file which was created during Sentinel install to be present (`/tmp/install.props`). You do not need to run this script on any but the first installed node; all relevant config files will be automatically synced to the other nodes.
- 5 If your environment varies from this recommended solution, you can edit the `resources.cli` file (in the same directory) and modify the primitives definitions from there. For example, the recommended solution uses a simple Filesystem resource; you may wish to use a more cluster-aware cLVM resource.
- 6 After running the shell script, you can issue a `crm status` command and the output should look like this:

```
crm status
```

---

```
Last updated: Thu Jul 26 16:34:34 2012
Last change: Thu Jul 26 16:28:52 2012 by hacluster via crmd on node01
Stack: openais
Current DC: node01 - partition with quorum
Version: 1.1.6-b988976485d15cb702c9307df55512d323831a5e
2 Nodes configured, 2 expected votes
5 Resources configured.
```

---

```
Online: [node01, node02]
stonith-sbd (stonith:external/sbd): Started node01
Resource Group: sentinelgrp
 sentinelip (ocf::heartbeat:IPaddr2): Started node01
 sentinelfs (ocf::heartbeat:Filesystem): Started node01
 sentineldb (ocf::novell:pgsql): Started node01
 sentinelserver (ocf::novell:sentinel): Started node01
```

- 7 At this point the relevant Sentinel resources should be configured in the cluster. You can examine how they are configured and grouped in the cluster management tool, for example by running `crm status`.

## Secondary Storage Configuration

Perform the following steps to configure the secondary storage so that Sentinel can migrate event partitions to less-expensive storage:

---

**NOTE:** This process is optional and the secondary storage does not need to be high-availability in the same way that you configured the rest of the system. You can use any directory, mounted from a SAN or not, NFS, or CIFS volume.

---

- 1 In the Sentinel Main interface, in the top menu bar, click **Storage**.
- 2 Select **Configuration**.
- 3 Select one of the radio buttons under Secondary storage not configured

Use a simple iSCSI Target as a network shared storage location, in much the same configuration as the primary storage. In your production environment, your storage technologies might differ.

Use the following procedure to configure the secondary storage for use by Sentinel:

---

**NOTE:** For iSCSI Target, the target will be mounted as a directory for use as secondary storage. You must configure the mount as a filesystem resource similar to the way the primary storage filesystem is configured. This was not automatically set up as part of the resource installation script since there are other possible variations.

---

- 1 Review the steps above to determine which partition was created for use as secondary storage (`/dev/<NETWORK1>`, or something like `/dev/sdc1`). If necessary create an empty directory on which the partition can be mounted (such as `/var/opt/netdata`).

- 2 Set up the network filesystem as a cluster resource: use the Sentinel Main interface or run the command:

```
crm configure primitive sentinelnetfs ocf:heartbeat:Filesystem params
device="/dev/<NETWORK1>" directory="<PATH>" fstype="ext3" op monitor
interval=60s
```

where `/dev/<NETWORK1>` is the partition that was created in the Shared Storage Setup section above, and `<PATH>` is any local directory on which it can be mounted.

- 3 Add the new resource to the group of managed resources:

```
crm resource stop sentinelgrp
crm configure delete sentinelgrp
crm configure group sentinelgrp sentinelip sentinelfs sentinelnetfs
sentinelldb sentinelserver
crm resource start sentinelgrp
```

- 4 You can connect to the node currently hosting the resources (use `crm status` or `Hawk`) and make sure that the secondary storage is properly mounted (use the `mount` command).

- 5 Log in to the Sentinel Main interface.

- 6 Select **Storage**, then select **Configuration**, then select the **SAN (locally mounted)** under Secondary storage not configured.

- 7 Type in the path where the secondary storage is mounted, for example `/var/opt/netdata`.

Use simple versions of the required resources, such as the simple Filesystem Resource Agent. You can choose to use more sophisticated cluster resources like cLVM (a logical-volume version of the filesystem) if required.

# 42 Configuring Sentinel HA as SSDM

This chapter provides information about configuring a Sentinel HA setup as SSDM. These instructions are applicable for both traditional and appliance installations.

To configure Sentinel HA setup as SSDM:

- 1 Install and configure scalable storage for Sentinel. For more information, see [Chapter 13, “Installing and Setting Up Scalable Storage,”](#) on page 85.
- 2 Enable scalable storage on the active node. For more information, see “[Enabling Scalable Storage Post-Installation](#)” in the *Sentinel Administration Guide*.

- 3 Run the following command on the active node:

```
csync2 -x -v
```

This synchronizes the SSDM configuration to all the passive nodes.

- 4 (Conditional) If the `csync2 -x -v` command you run in Step 3 fails to synchronize all the files, perform the following steps:

- 4a Clear the `csync2` database (in the `/var/lib/csync2` directory) on all the nodes.

- 4b Run the following command on all servers to update the `csync2` database to match the filesystem, but without marking anything as needing to be synchronized to other servers:

```
csync2 -cIr /
```

- 4c Run the following command to find all the differences between authoritative server and remote servers, and mark for synchronization:

```
csync2 -TUXI
```

- 4d Run the following command to reset the database to force the current server to be winner on any conflicts:

```
csync2 -fr /
```

- 4e Run the following command to start a synchronization to all the other servers:

```
csync2 -xr /
```

- 4f Run the following command to verify that all the files are synchronized:

```
csync2 -T
```

This command will not list any files if the synchronization is successful.



# 43 Upgrading Sentinel in High Availability

When you upgrade Sentinel in an HA environment, you should first upgrade the passive nodes in the cluster, then upgrade the active cluster node.

- ♦ “Prerequisites” on page 221
- ♦ “Upgrading a Traditional Sentinel HA Installation” on page 221
- ♦ “Upgrading a Sentinel HA Appliance Installation” on page 227

## Prerequisites

- ♦ Download the latest installer from the [Downloads website](#).
- ♦ If you are using the SLES operating system with kernel version 3.0.101 or later, you must manually load the watchdog driver in the computer. To find the appropriate watchdog driver for your computer hardware, contact your hardware vendor. To load the watchdog driver, perform the following:

1. In the command prompt, run the following command to load the watchdog driver in the current session:

```
/sbin/modprobe -v --ignore-install <watchdog driver name>
```

2. Add the following line to the `/etc/init.d/boot.local` file to ensure that the computer automatically loads the watchdog driver at every boot time:

```
/sbin/modprobe -v --ignore-install <watchdog driver name>
```

## Upgrading a Traditional Sentinel HA Installation

This section provides information about upgrading a traditional Sentinel installation, and also about upgrading the operating system in a traditional Sentinel installation.

Sentinel now uses PostgreSQL instead of MongoDB to store Security Intelligence data and alerts data.

On the active node, the upgrade process does the following:

- ♦ Migrates Security Intelligence data, alerts data, and so on from MongoDB to PostgreSQL.

Sentinel now stores Security Intelligence data and alerts data in PostgreSQL instead of MongoDB. The upgrade process will first migrate this data to PostgreSQL and if successful, it will automatically proceed with the upgrade. If the data migration is unsuccessful, you cannot upgrade Sentinel.

- ♦ Generates a cleanup script that you can use to remove data and MongoDB related RPMs from MongoDB.

The data stored in MongoDB is retained as a backup and you can delete it after upgrading Sentinel.

- ♦ [“Upgrading Sentinel HA” on page 222](#)
- ♦ [“Upgrading the Operating System” on page 224](#)

## Upgrading Sentinel HA

- 1 Enable the maintenance mode on the cluster:

```
crm configure property maintenance-mode=true
```

Maintenance mode helps you to avoid any disturbance to the running cluster resources while you update Sentinel. You can run this command from any cluster node.

- 2 Verify whether the maintenance mode is active:

```
crm status
```

The cluster resources should appear in the unmanaged state.

- 3 Upgrade the passive cluster node:

- 3a Stop the cluster stack:

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain accessible and avoids fencing of nodes.

- 3b Log in as `root` to the server where you want to upgrade Sentinel.

- 3c Extract the install files from the tar file:

```
tar xfz <install_filename>
```

- 3d Run the following command in the directory where you extracted the install files:

```
./install-sentinel --cluster-node
```

- 3e After the upgrade is complete, restart the cluster stack:

```
rcpacemaker start
```

Repeat [Step 3](#) for all passive cluster nodes.

- 3f Remove the autostart scripts so that the cluster can manage the product.

```
cd /
```

```
insserv -r sentinel
```

- 4 Upgrade the active cluster node:

- 4a Back up your configuration, then create an ESM export.

For more information about backing up data, see [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide*.

- 4b Stop the cluster stack:

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain accessible and avoids fencing of nodes.

- 4c Log in as `root` to the server where you want to upgrade Sentinel.

- 4d Run the following command to extract the install files from the tar file:

```
tar xfz <install_filename>
```

**4e** Run the following command in the directory where you extracted the install files:

```
./install-sentinel
```

**4f** Select the desired migration option.

---

**WARNING:** Ensure that you select the appropriate option because you cannot repeat this procedure after the upgrade is successful.

---

If your data is migrated successfully, the upgrade process will automatically proceed with the upgrade.

The upgrade process retains the data that was stored in MongoDB as a backup.

**4g** (Conditional) If the data migration is not successful:

**4g1** Clean up the unsuccessful migrated data. For more information, see [“Cleaning Up Data From PostgreSQL When Migration Fails” on page 171](#)

**4g2** Repeat this procedure to upgrade Sentinel.

**4h** After the upgrade is complete, start the cluster stack:

```
rcpacemaker start
```

**4i** Remove the autostart scripts so that the cluster can manage the product.

```
cd /
```

```
insserv -r sentinel
```

**4j** Run the following command to synchronize any changes in the configuration files:

```
csync2 -x -v
```

**5** Disable the maintenance mode on the cluster:

```
crm configure property maintenance-mode=false
```

You can run this command from any cluster node.

**6** Verify whether the maintenance mode is inactive:

```
crm status
```

The cluster resources should appear in the Started state.

**7** (Optional) Verify whether the Sentinel upgrade is successful:

```
rcsentinel version
```

**8** Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence data, and so on.

**9** The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).

## Upgrading the Operating System

This section provides information about how to upgrade the operating system to a major version, such as upgrading from SLES 11 to SLES 12, in a Sentinel HA cluster. When you upgrade the operating system, you must perform few configuration tasks, to ensure that Sentinel HA works correctly after you upgrade the operating system.

Perform the steps described in the following sections:

- ♦ [“Upgrading the Operating System” on page 224](#)
- ♦ [“Configuring iSCSI Targets” on page 225](#)
- ♦ [“Configuring iSCSI Initiators” on page 226](#)
- ♦ [“Configuring the HA Cluster” on page 226](#)

## Upgrading the Operating System

To upgrade the operating system:

- 1 Log in as `root` user to any node in the Sentinel HA cluster.
- 2 Run the following command to enable the maintenance mode on the cluster:  

```
crm configure property maintenance-mode=true
```

The maintenance mode helps you to avoid any disturbance to the running cluster resources while you upgrade the operating system.
- 3 Run the following command to verify whether the maintenance mode is active:  

```
crm status
```

The cluster resources should appear in the unmanaged state.
- 4 Ensure that you have upgraded Sentinel to version 8.2 or later on all the cluster nodes.
- 5 Ensure that all the nodes in the cluster are registered with SLES and SLESHA.
- 6 Perform the following steps to upgrade the operating system on the passive cluster node:
  - 6a Run the following command to stop the cluster stack:  

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain inaccessible and avoids fencing of nodes.
  - 6b Upgrade the operating system. For more information, see [Upgrading the Operating System](#).
- 7 Repeat Step 6 on all the passive nodes to upgrade the operating system.
- 8 Repeat Step 6 on the active node to upgrade the operating system on it.
- 9 Repeat Step 6b to upgrade the operating system on shared storage.
- 10 Ensure that the operating system on all the nodes in the cluster is upgraded to SLES12 SP3.



## Configuring iSCSI Targets

To configure iSCSI targets:

- 1 On the shared storage, check if the iSCSI LIO package is installed. If it is not installed yet, go to YaST2 Software Management and install the iSCSI LIO package (`iscciliotarget` RPM).

- 2 Perform the following steps on all the nodes in the cluster:

- 2a Run the following command to open the file which contains the iSCSI initiator name:

```
cat /etc/iscsi/initiatorname.iscsi
```

- 2b Note the initiator name which will be used for configuring iSCSI initiators:

For example:

```
InitiatorName=iqn.1996-04.de.suse:01:441d6988994
```

These initiator names will be used while configuring iSCSI Target Client Setup.

- 3 Click **Service** and select the **When Booting** option to ensure that the service starts when the operating system boots.

- 4 Select the **Global** tab, deselect **No Authentication** to enable authentication, and then specify the user name and the password for incoming and outgoing authentication.

The **No Authentication** option is enabled by default. However, you should enable authentication to ensure that the configuration is secure.

- 5 Click **Targets**, and click **Add** to add a new target.

- 6 Click **Add** to add a new LUN.

- 7 Leave the LUN number as 0, browse in the **Path** dialog (under **Type=fileio**) and select the `/localdata` file that you created. If you have a dedicated disk for storage, specify a block device, such as `/dev/sdc`.

- 8 Repeat steps 6 and 7, and add LUN 1 and select `/networkdata` this time.

- 9 Repeat steps 6 and 7, and add LUN 2 and select `/sbd` this time.

- 10 Leave the other options at their default values. Click **Next**.

- 11 Click **Add**. When prompted for Client Name, specify the initiator name you have copied in Step 2. Repeat this step to add all the client names, by specifying the initiator names.

The list of client names will be displayed in the Client List.

- 12 (Conditional) If you have enabled authentication in Step 4, provide the authentication credentials you specified in Step 4.

Select a client, select **Edit Auth > Incoming Authentication**, and specify the user name and password. Repeat this for all the clients.

- 13 Click **Next** to select the default authentication options, and then click **Finish** to exit the configuration. Restart iSCSI if prompted.

- 14 Exit YaST.

## Configuring iSCSI Initiators

To configure iSCSI initiators:

- 1 Connect to one of the cluster nodes (node01) and start YaST.
- 2 Click **Network Services** > **iSCSI Initiator**.
- 3 If prompted, install the required software (`iscsiclient` RPM).
- 4 Click **Service**, and select **When Booting** to ensure that the iSCSI service is started on boot.
- 5 Click **Discovered Targets**.

---

**NOTE:** If any previously existing iSCSI targets are displayed, delete those targets.

---

Select **Discovery** to add a new iSCSI target.

- 6 Specify the iSCSI Target IP address (10.0.0.3).  
(Conditional) If you have enabled authentication in Step 4 in [“Configuring iSCSI Targets” on page 225](#), deselect **No Authentication**. In the **Outgoing Authentication** section, enter the authentication credentials you specified while configuring iSCSI targets.  
Click **Next**.
- 7 Select the discovered iSCSI Target with the IP address 10.0.0.3 and select **Log In**.
- 8 Perform the following steps:
  - 8a Switch to Automatic in the **Startup** drop-down menu.
  - 8b (Conditional) If you have enabled authentication, deselect **No Authentication**.  
The user name and the password you have specified should be displayed in the **Outgoing Authentication** section. If these credentials are not displayed, enter the credentials in this section.
  - 8c Click **Next**.
- 9 Switch to the **Connected Targets** tab to ensure that you are connected to the target.
- 10 Exit the configuration. This should have mounted the iSCSI Targets as block devices on the cluster node.
- 11 In the YaST main menu, select **System** > **Partitioner**.
- 12 In the System View, you should see new hard disks of the LIO-ORG-FILEIO type (such as `/dev/sdb` and `/dev/sdc`) in the list, along with already formatted disks (such as `/dev/sdb1` or `/dev/<SHARED1`).
- 13 Repeat steps 1 through 12 on all the nodes.

## Configuring the HA Cluster

To configure the HA cluster:

- 1 Start YaST2 and go to **High Availability** > **Cluster**.
- 2 If prompted, install the HA package and resolve the dependencies.  
After the HA package installation, Cluster—Communication Channels is displayed.
- 3 Ensure that `Unicast` is selected as the Transport option.

- 4 Select **Add a Member Address** and specify the node IP address, and then repeat this action to add all the other cluster node IP addresses.
- 5 Ensure that **Auto Generate Node ID** is selected.
- 6 Ensure that the HAWK service is enabled on all the nodes. If it is not enabled, run the following command to enable it:
 

```
service hawk start
```
- 7 Run the following command:
 

```
ls -l /dev/disk/by-id/
```

The SBD partition ID is displayed. For example, `scsi-1LIO-ORG_FILEIO:33caaa5a-a0bc-4d90-b21b-2ef33030cc53`.

Copy the ID.
- 8 Open the `sbdd` file (`/etc/sysconfig/sbdd`), and change the ID of `SBD_DEVICE` with the ID you have copied in step 7.
- 9 Run the following commands to restart the pacemaker service:
 

```
rcpacemaker restart
```
- 10 Run the following commands to remove the autostart scripts, so that the cluster can manage the product.
 

```
cd /
insserv -r sentinel
```
- 11 Repeat steps 1 through 10 on all the cluster nodes.
- 12 Run the following command to synchronize any changes in the configuration files:
 

```
csync2 -x -v
```
- 13 Run the following command to disable the maintenance mode on the cluster:
 

```
crm configure property maintenance-mode=false
```

You can run this command from any cluster node.
- 14 Run the following command to verify whether the maintenance mode is inactive:
 

```
crm status
```

The cluster resources should appear in the Started state.

## Upgrading a Sentinel HA Appliance Installation

You can upgrade to Sentinel from Sentinel 8.2 or later. You can upgrade both Sentinel and the SLES Operating System through the Sentinel Appliance Manager or Zypper (Appliance Update Channel).

Sentinel now uses PostgreSQL instead of MongoDB to store Security Intelligence data and alerts data. Before you upgrade the appliance on the active node, you must first migrate your data from MongoDB to PostgreSQL. You will be able to upgrade the appliance only if you have successfully migrated your data to PostgreSQL.

- ♦ [“Upgrading through the Zypper patch” on page 228](#)
- ♦ [“Upgrading through the Sentinel Appliance Management Console” on page 229](#)

## Upgrading through the Zypper patch

You must register all the appliance nodes through Sentinel Appliance Manager before the upgrade. For more information, see [“Registering for Updates” on page 106](#). If you do not register the appliance, Sentinel displays a yellow warning.

- 1 Enable the maintenance mode on the cluster.

```
crm configure property maintenance-mode=true
```

Maintenance mode helps you to avoid any disturbance to the running cluster resources while you update the Sentinel software. You can run this command from any cluster node.

- 2 Verify whether the maintenance mode is active.

```
crm status
```

The cluster resources should appear in the unmanaged state.

- 3 Upgrade the passive cluster node:

- 3a Stop the cluster stack.

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain inaccessible and avoids fencing of nodes.

- 3b Complete prerequisites 1 and 2 listed in [“Prerequisites for Upgrading the Appliance” on page 163](#)

- 3c Download the updates for Sentinel.

```
zypper -v patch
```

- 3d (Conditional) If the installer displays a message that you must resolve dependency for the OpenSSH package, enter the appropriate option to downgrade the OpenSSH package.

- 3e (Conditional) If the installer displays a message that indicates change in the ncgOverlay architecture, enter the appropriate option to accept the architecture change.

- 3f (Conditional) If the installer displays a message that you must resolve dependency for some appliance packages, enter the appropriate option to uninstall the dependent packages.

- 3g After the upgrade is complete, start the cluster stack.

```
rcpacemaker start
```

- 4 Repeat Step 3 for all the passive cluster nodes.

- 5 Upgrade the active cluster node:

- 5a Back up your configuration, then create an ESM export.

For more information on backing up data, see [“Backing Up and Restoring Data”](#) in the [Sentinel Administration Guide](#).

- 5b Stop the cluster stack.

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain inaccessible and avoids fencing of nodes.

- 5c Complete the prerequisites listed in [“Prerequisites for Upgrading the Appliance” on page 163](#).

**5d** Download the updates for Sentinel.

To upgrade Sentinel:

```
zypper -v patch
```

**5e** (Conditional) If the installer displays a message that you must resolve dependency for the OpenSSH package, enter the appropriate option to downgrade the OpenSSH package.

**5f** (Conditional) If the installer displays a message that indicates change in the ncgOverlay architecture, enter the appropriate option to accept the architecture change.

**5g** (Conditional) If the installer displays a message that you must resolve dependency for some appliance packages, enter the appropriate option to uninstall the dependent packages.

**5h** After the upgrade is complete, start the cluster stack.

```
rcpacemaker start
```

**5i** Run the following command to synchronize any changes in the configuration files:

```
csync2 -x -v
```

**6** Disable the maintenance mode on the cluster.

```
crm configure property maintenance-mode=false
```

You can run this command from any cluster node.

**7** Verify whether the maintenance mode is inactive.

```
crm status
```

The cluster resources should appear in the Started state.

**8** (Optional) Verify whether the upgrade is successful:

```
rcsentinel version
```

**9** Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence dashboards, and so on.

**10** The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).

## Upgrading through the Sentinel Appliance Management Console

**To upgrade through the Sentinel Appliance Management Console:**

**1** (Conditional) If you have already installed Sentinel 8.2 and `sentinel_sles_iso_os_post_upgrade-release-73.tar.gz`, you will need to reinstall a few RPMs that were uninstalled by the `sentinel_sles_iso_os_post_upgrade-release-73.tar.gz` installer. For more information about installing these RPMs, see [Knowledge Base Article 7023543](#).

**2** Run the following command on the active node or a passive node in the cluster, to enable maintenance mode:

```
crm configure property maintenance-mode=true
```

Maintenance mode helps you to avoid any disturbance to the running cluster resources while you update Sentinel.

**3** Run the following command to verify whether the maintenance mode is active:

```
crm status
```

The cluster resources should be displayed in the unmanaged state.

**4 Upgrade all the passive cluster nodes first:**

**4a** Run the following command to stop the cluster stack:

```
rcpacemaker stop
```

Stopping the cluster stack ensures that the cluster resources remain inaccessible and avoids fencing of nodes.

**4b** Run the following command to verify whether the 9443 port is listening on the active node to access appliance:

```
netstat -na | grep 9443
```

**4c** (Conditional) Run the following command if the 9443 port is not listening:

```
systemctl restart vabase vabase-jetty vabase-datamodel
```

**4d** Complete prerequisites 1 and 2 listed in [“Prerequisites for Upgrading the Appliance” on page 163](#)

**4e** Launch the appliance by doing either of the following:

- ♦ Log in to Sentinel. Click **Sentinel Main > Appliance**.
- ♦ Specify the following URL in your web browser: `https://<IP_address>:9443`.

**4f** (Conditional) If you are unable to launch Sentinel Appliance Management Console:

**4f1** Go to `/var/opt/novell` in the active node and copy the following files to `/var/opt/novell/` in each passive node:

- ♦ `datamodel-service`
- ♦ `ganglia`
- ♦ `jetty`
- ♦ `python`
- ♦ `va`

**4f2** In each passive node, set the file permission as `vabase-jetty` for the files in the `jetty` folder:

1. Go to `/var/opt/novell/jetty`.
2. Run the following command:

```
chown -R vabase-jetty:vabase-jetty *
```

**4f3** Run the following command to restart the `vabase` services:

```
systemctl start vabase-jetty vabase-datamodel vabase
```

**4f4** Run the following command to verify that port 9443 is listening in all the available nodes:

```
netstat -na | grep 9443
```

**4g** Log in as `vaadmin`.

- 4h** Click **Online Update**.
  - 4h1** (Conditional) Register for updates if you have not done it earlier. For more information, see [“Registering for Updates” on page 106](#).
  - 4h2** To install the displayed updates for Sentinel and the operating system, click **Update Now > OK**.
  - 4h3** To apply the installed updates, click **Reboot**.
  - 4h4** After reboot, check the version on the top right corner of the screen to verify if the upgrade is successful.
- 4i** After the upgrade is complete, restart the cluster stack.

```
rcpacemaker start
```
- 5** Upgrade the active cluster node.
  - 5a** Complete the prerequisites listed in [“Prerequisites for Upgrading the Appliance” on page 163](#).
  - 5b** Repeat Step 3 for the active cluster node.
- 6** Run the following command on the active node or a passive node in the cluster to disable maintenance mode:

```
crm configure property maintenance-mode=false
```
- 7** Run the following command on the active node or a passive node in the cluster to verify if maintenance mode is not active:

```
crm status
```
- 8** Log in to Sentinel and verify if you are able to see the migrated data such as alerts, Security Intelligence dashboards, and so on.
- 9** The data in MongoDB is now redundant because Sentinel 8.3 and later will store data only in PostgreSQL. To clear up the disk space, delete this data. For more information, see [“Removing Data from MongoDB” on page 175](#).





# 44 Backup and Recovery

The highly available failover cluster described in this document provides a level of redundancy so that if the service fails on one node in the cluster, it will automatically failover and recover on another node in the cluster. When an event like this happens, it's important to bring the node that failed back into an operational state so that the redundancy in the system can be restored and protect in the case of another failure. This section talks about restoring the failed node under a variety of failure conditions.

- ♦ [“Backup” on page 233](#)
- ♦ [“Recovery” on page 233](#)

## Backup

While a highly available failover cluster like the one described in this document provides a layer of redundancy, it is still important to regularly take a traditional backup of the configuration and data, which would not be easy to recover from if lost or corrupted. The section [“Backing Up and Restoring Data”](#) in the *Sentinel Administration Guide* describes how to use Sentinel's built-in tools for creating a backup. These tools should be used on the active node in the cluster because the passive node in the cluster will not have the required access to the shared storage device. Other commercially available backup tools could be used instead and may have different requirements on which node they can be used.

## Recovery

- ♦ [“Transient Failure” on page 233](#)
- ♦ [“Node Corruption” on page 234](#)
- ♦ [“Cluster Data Configuration” on page 234](#)

## Transient Failure

If the failure was a temporary failure and there is no apparent corruption to the application and operating system software and configuration, then simply clearing the temporary failure, for example rebooting the node, will restore the node to an operational state. The cluster management user interface can be used to fail back the running service back to the original cluster node, if desired.

## Node Corruption

If the failure caused a corruption in the application or operating system software or configuration that is present on the node's storage system, then the corrupted software will need to be reinstalled. Repeating the steps for adding a node to the cluster described earlier in this document will restore the node to an operational state. The cluster management user interface can be used to fail back the running service back to the original cluster node, if desired.

## Cluster Data Configuration

If data corruption occurs on the shared storage device in a way that the shared storage device can't recover from, this would result in the corruption affecting the entire cluster in a way that cannot be automatically recovered from using the highly available failover cluster described in this document. The section “[Backing Up and Restoring Data](#)” in the *Sentinel Administration Guide* describes how to use Sentinel's built-in tools for restoring from a backup. These tools should be used on the active node in the cluster because the passive node in the cluster will not have the required access to the shared storage device. Other commercially available backup and restore tools could be used instead and may have different requirements on which node they can be used.

# VIII Appendices

- ◆ [Appendix A, “Troubleshooting,” on page 237](#)
- ◆ [Appendix B, “Uninstalling,” on page 243](#)



# A Troubleshooting

This section contains some of the issues that might occur during installation, along with the actions to work around the issues.

- ♦ [“Failed Installation Because of an Incorrect Network Configuration” on page 237](#)
- ♦ [“The UUID Is Not Created for Imaged Collector Managers or Correlation Engine” on page 238](#)
- ♦ [“Sentinel Main Interface is Blank in Internet Explorer After Logging in” on page 238](#)
- ♦ [“Sentinel Does Not Launch in Internet Explorer 11 in Windows Server 2012 R2” on page 238](#)
- ♦ [“Sentinel Cannot Run Local Reports with Default EPS License” on page 239](#)
- ♦ [“Synchronization Needs to be Started Manually in Sentinel High Availability After You Convert the Active Node to FIPS 140-2 Mode” on page 239](#)
- ♦ [“Sentinel Main Interface Displays Blank Page After Converting to Sentinel Scalable Data Manager” on page 239](#)
- ♦ [“The Event fields Panel is Missing in the Schedule Page When Editing Some Saved Searches” on page 240](#)
- ♦ [“Sentinel Does Not Return Any Correlated Events When You Search for Events for the Deployed Rule with the Default Fire Count Search” on page 240](#)
- ♦ [“Security Intelligence Dashboard Displays Invalid Baseline Duration When Regenerating a Baseline” on page 240](#)
- ♦ [“Sentinel Server Shuts Down When Running a Search If There Are Large Number of Events in a Single Partition” on page 240](#)
- ♦ [“Error While Using the report\\_dev\\_setup.sh Script to Configure Sentinel Ports for Firewall Exceptions on Upgraded Sentinel Appliance Installations” on page 241](#)

## Failed Installation Because of an Incorrect Network Configuration

During the first boot, if the installer finds that the network settings are incorrect, an error message is displayed. If the network is unavailable, installing Sentinel on the appliance fails.

To resolve this issue, properly configure the network settings. To verify the configuration, use the `ifconfig` command to return the valid IP address, and use the `hostname -f` command to return the valid hostname.

# The UUID Is Not Created for Imaged Collector Managers or Correlation Engine

If you image a Collector Manager server (for example, by using ZENworks Imaging) and restore the images on different machines, Sentinel does not uniquely identify the new instances of the Collector Manager. This happens because of duplicate UUIDs.

You must generate a new UUID by performing the following steps on the newly installed Collector Manager systems:

- 1 Delete the `host.id` or `sentinel.id` file that is located in the `/var/opt/novell/sentinel/data` folder.
- 2 Restart the Collector Manager.

The Collector Manager automatically generates the UUID.

# Sentinel Main Interface is Blank in Internet Explorer After Logging in

If the Internet Security Level is set to High, a blank page appears after logging in to Sentinel and the file download pop-up might be blocked by the browser. To work around this issue, you need to first set the security level to Medium-high and then change to Custom level as follows:

1. Navigate to **Tools > Internet Options > Security** and set the security level to **Medium-high**.
2. Make sure that the **Tools > Compatibility View** option is not selected.
3. Navigate to **Tools > Internet Options > Security tab > Custom Level**, then scroll down to the **Downloads** section and select **Enable** under the **Automatic prompting for file downloads** option.

# Sentinel Does Not Launch in Internet Explorer 11 in Windows Server 2012 R2

When you use Windows Server 2012 R2, Sentinel does not launch in Internet Explorer 11 due to the default security configurations of Internet Explorer 11. You must manually add Sentinel to the list of trusted sites before launching Sentinel.

## To Add Sentinel to the List of Trusted Sites

- 1 Open Internet Explorer 11.
- 2 Click **Settings** icon > **Internet Options > Security tab > Trusted Sites > Sites**
- 3 Add Sentinel host to the list of trusted sites.

## Sentinel Cannot Run Local Reports with Default EPS License

If your environment has the default 25 EPS license and you run a report, the report fails with the following error: `License for Distributed Search feature is expired`

To run reports in the same JVM as Sentinel, complete the following steps:

- 1 Log in to the Sentinel server and open the `/etc/opt/novell/sentinel/config/object-component.JasperReportingComponent.properties` file.
- 2 Locate the `reporting.process.oktorunstandalone` property.
- 3 (Conditional) If the property is not in the file, add it.
- 4 Set the property to `false`. For example:  

```
reporting.process.oktorunstandalone=false
```
- 5 Restart Sentinel.

## Synchronization Needs to be Started Manually in Sentinel High Availability After You Convert the Active Node to FIPS 140-2 Mode

**Issue:** When you convert the active node to FIPS 140-2 mode in Sentinel HA, the synchronization to convert all the passive nodes to FIPS 140-2 mode is not performed completely. You must start the synchronization manually.

**Workaround:** Manually synchronize all passive nodes to FIPS 140-2 mode as follows:

- 1 Log in as the root user on the active node.
- 2 Open the `/etc/csync2/csync2.cfg` file.
- 3 Change the following line:  

```
include /etc/opt/novell/sentinel/3rdparty/nss/*;
```

  
to  

```
include /etc/opt/novell/sentinel/3rdparty/nss;
```
- 4 Save the `csync2.cfg` file.
- 5 Start the synchronization manually by running the following command:  

```
csync2 -x -v
```

## Sentinel Main Interface Displays Blank Page After Converting to Sentinel Scalable Data Manager

**Issue:** After you enable SSDM, when you log in to the Sentinel Main interface, the browser displays a blank page.

**Workaround:** Close your browser and log in to the Sentinel Main interface again. This issue only happens once, the first time you log in to the Sentinel Main interface after you enable SSDM.

## The Event fields Panel is Missing in the Schedule Page When Editing Some Saved Searches

**Issue:** When editing a saved search upgraded from Sentinel 7.2 to a later version, the **Event fields** panel, used to specify output fields in the search report CSV, is missing in the schedule page.

**Workaround:** After upgrading Sentinel, recreate and reschedule the search to view the **Event fields** panel in the schedule page.

## Sentinel Does Not Return Any Correlated Events When You Search for Events for the Deployed Rule with the Default Fire Count Search

**Issue:** Sentinel does not return any correlated events when you search for all correlated events that were generated after the rule was deployed or enabled, by clicking the icon next to **Fire count** in the **Activity statistics** panel in the Correlation Summary page for the rule.

**Workaround:** Change the value in the **From** field in the Event Search page to a time earlier than the populated time in the field and click **Search** again.

## Security Intelligence Dashboard Displays Invalid Baseline Duration When Regenerating a Baseline

**Issue:** During Security Intelligence baseline regeneration, the start and finish dates for the baseline are incorrect and display 1/1/1970.

**Workaround:** The correct dates are updated after the baseline regeneration is complete.

## Sentinel Server Shuts Down When Running a Search If There Are Large Number of Events in a Single Partition

**Issue:** Sentinel server shuts down when you run a search if there are a large number of events indexed in a single partition.

**Workaround:** Create retention policies in such a way that there are at least two partitions open in a day. Having more than one partition open helps reduce the number of events indexed in partitions.

You can create retention policies that filter events based on the `estzhour` field, which tracks the hour of the day. Therefore, you can create one retention policy with `estzhour:[0 TO 11]` as the filter and another retention policy with `estzhour:[12 TO 23]` as the filter.

For more information, see “[Configuring Data Retention Policies](#)” in the *Sentinel Administration Guide*.



# Error While Using the `report_dev_setup.sh` Script to Configure Sentinel Ports for Firewall Exceptions on Upgraded Sentinel Appliance Installations

**Issue:** Sentinel displays an error when you use the `report_dev_setup.sh` script to configure Sentinel ports for firewall exceptions.

**Workaround:** Configure Sentinel ports for firewall exceptions through the following steps:

1 Open the `/etc/sysconfig/SuSEfirewall2` file.

2 Change the following line:

```
FW_SERVICES_EXT_TCP=" 443 8443 4984 22 61616 10013 289 1289 1468 1443
40000:41000 1290 1099 2000 1024 1590"
```

to

```
FW_SERVICES_EXT_TCP=" 443 8443 4984 22 61616 10013 289 1289 1468 1443
40000:41000 1290 1099 2000 1024 1590 5432"
```

3 Restart Sentinel.



# B Uninstalling

This appendix provides information about uninstalling Sentinel and post-uninstallation tasks.

- ♦ [“Uninstallation Checklist” on page 243](#)
- ♦ [“Uninstalling Sentinel” on page 243](#)
- ♦ [“Post-Uninstallation Tasks” on page 244](#)

## Uninstallation Checklist

Use the following checklist to uninstall Sentinel:

- Uninstall the Sentinel server.
- Uninstall the Collector Manager and Correlation Engine, if any.
- Perform post-uninstallation tasks to complete the Sentinel uninstallation.

## Uninstalling Sentinel

An uninstall script is available to help you remove a Sentinel installation. Before performing a new installation, you should perform all of the following steps to ensure there are no files or system settings remaining from a previous installation.

---

**WARNING:** These instructions involve modifying operating system settings and files. If you are not familiar with modifying these system settings and files, please contact your system administrator.

---

## Uninstalling the Sentinel Server

Use the following steps to uninstall the Sentinel server:

- 1 Log in to the Sentinel server as `root`.

---

**NOTE:** You cannot uninstall Sentinel server as a non-root user, if the installation is performed as a `root` user. However, a non-root user can uninstall the Sentinel server if the installation was performed by non-root user.

---

- 2 Access the following directory:

```
<sentinel_installation_path>/opt/novell/sentinel/setup/
```

- 3 Run the following command:

```
./uninstall-sentinel
```

- 4 When prompted to reconfirm that you want to proceed with the uninstall, press `y`.  
The script first stops the service and then removes it completely.

## Uninstalling the Collector Manager and Correlation Engine

Use the following steps to uninstall the Collector Manager and Correlation Engine:

- 1 Log in as `root` to the Collector Manager and Correlation Engine computer.

---

**NOTE:** You can not uninstall the remote Collector Manager or remote Correlation Engine as a non-root user, if the installation was performed as a `root` user. However, a non-root user can uninstall, if the installation was done as a non-root user.

---

- 2 Go to the following location:

```
/opt/novell/sentinel/setup
```

- 3 Run the following command:

```
./uninstall-sentinel
```

The script displays a warning that the Collector Manager or Correlation Engine, and all associated data will be completely removed.

- 4 Enter `y` to remove the Collector Manager or Correlation Engine.

The script first stops the service and then removes it completely. However, the Collector Manager and Correlation Engine icon is still displayed in inactive state in the Sentinel Main interface.

- 5 (Conditional) If you have enabled Event Visualization, you must redeploy the Elasticsearch security plug-in. For more information, see [“Redeploying Elasticsearch Security Plug-In” on page 82](#).

- 6 Perform the following additional steps to manually delete the Collector Manager and Correlation Engine in the Sentinel Main interface:

**Collector Manager:**

1. Access **Event Source Management > Live View**.
2. Right-click the Collector Manager you want to delete, then click **Delete**.

**Correlation Engine:**

1. Navigate to the **Sentinel Main** interface as an administrator.
2. Expand **Correlation**, then select the Correlation Engine that you want to delete.
3. Click the **Delete** button (garbage can icon).

## Post-Uninstallation Tasks

Uninstalling the Sentinel server does not remove the Sentinel Administrator User from the operating system. You must manually remove that user.

After you uninstall Sentinel, certain systems settings remain. These settings should be removed before performing a “clean” installation of Sentinel, particularly if the Sentinel uninstallation encountered errors.

To manually clean up the Sentinel system settings:

- 1 Log in as `root`.
- 2 Ensure that all Sentinel processes are stopped.
- 3 Remove the contents of `/opt/novell/sentinel` or wherever the Sentinel software was installed.
- 4 Make sure no one is logged in as the Sentinel Administrator operating system user (novell by default), then remove the user, the home directory, and the group.  

```
userdel -r novell
groupdel novell
```
- 5 Restart the operating system.

