



# CHAPTER 23

## Monitoring and Troubleshooting

---

The Operations tab on the Cisco Identity Services Engine (ISE) home page, also known as the dashboard, provides integrated monitoring, reporting, alerting, and troubleshooting, all from one centralized location.

This chapter explains monitoring and troubleshooting functions and tasks, and covers the following topics:

- [Understanding Monitoring and Troubleshooting, page 23-1](#)
- [Configuring Devices for Monitoring, page 23-3](#)
- [Cisco ISE Dashboard Monitoring, page 23-3](#)
- [Monitoring the Network, page 23-10](#)
- [Troubleshooting the Network, page 23-29](#)
- [Obtaining Additional Troubleshooting Information, page 23-40](#)
- [Monitoring Administration, page 23-45](#)



### Note

For a list of inherent known issues and workarounds associated with monitoring and troubleshooting, see the [Release Notes for the Cisco Identity Services Engine, Release 1.1](#).

---

## Understanding Monitoring and Troubleshooting

Monitoring and troubleshooting is a comprehensive identity solution for all Cisco ISE run-time services, using the following components:

- **Monitoring**—Provides a real-time presentation of meaningful data representing the state of access activities on a network. This insight allows you to easily interpret and effect operational conditions.
- **Troubleshooting**—Provides contextual guidance for resolving access issues on networks. You can then address user concerns and provide resolution in a timely manner.
- **Reporting**—Provides a catalog of standard reports that you can use to analyze trends and monitor system performance and network activities. You can customize reports in various ways, and save your changes for future use.

The Cisco ISE dashboard provides visibility into configured policies, authentication and authorization activities, profiled endpoints, postured sessions, and guest activities. Likewise, monitoring and troubleshooting capabilities include the following:

- A real-time summary of system activity and individual services, as well as a comprehensive at-a-glance view of network activity.
- A web-based user interface that simplifies generating and accessing predefined and custom reports.
- Various alert capabilities, including rules and triggers on authentication activity, that allows for early detection of operation or trends.

The data that is gathered by monitoring functionality is accessible from the central administration console, known as the Cisco ISE dashboard. When you log into the administration console, the real-time data, as shown in [Figure 23-1](#).

The dashboard shows the activity on the Network Privilege Framework (NPF), and provides drill-down capabilities for information on the various components. For information on the dashlets and metric meters that comprise the dashboard, see [Cisco ISE Dashboard Monitoring, page 23-3](#).

The NPF is composed of the following three tiers:

**Table 23-1 NPF Tiers**

Tier	Specifications
1	Access control based on identity using 802.1x, MAC authentication bypass (MAB), the Cisco ISE Profiler service
2	Access control based on identity using 802.1x, MAB, Profiler, guest provisioning of the Network Admission Control (NAC) manager, central web authentication
3	Access control based on identity and posture using 802.1x, MAB, Profiler, guest provisioning of the NAC manager, central web authentication

NPF authentication and authorization generates a flow of events. The events from the different sources are then collected by Cisco ISE monitoring and troubleshooting tools and summarized. You can view the authentication and authorization results on the dashboard, or choose to run any number of reports. For more information, see [Chapter 24, “Reporting.”](#)

**The NPF authentication and authorization event flow uses the following process:**

- 
- Step 1** NAD performs an authorization or flex authorization.
- Step 2** An unknown, agentless identity is profiled with web authorization.
- Step 3** RADIUS server authenticates and authorizes the identity.
- Step 4** Authorization is provisioned for the identity at the port.
- Step 5** Unauthorized endpoint traffic is dropped.
- 

## User Roles and Permissions

Monitoring and troubleshooting capabilities are associated with default user roles. The tasks you are allowed to perform are directly related to your assigned user role. For more information on the user roles and their assigned permissions, see [Understanding the Impact of Roles and Admin Groups, page 2-19](#).

## Monitoring and Troubleshooting Database

The Cisco ISE monitoring service collects and stores data in a specialized Monitoring database. The rate and amount of data utilized to monitor network functions may require a node dedicated solely to monitoring. If your Cisco ISE network collects logging data at a high rate from Policy Service ISE nodes or network devices, a Cisco ISE node dedicated to monitoring is recommended.

To manage the information stored in the Monitoring database, administrators are required to perform full and incremental backups of the database. This includes purging unwanted data, and then restoring the database. For more information, see [Monitoring Administration, page 23-45](#).

## Configuring Devices for Monitoring

The Monitoring ISE node receives and uses data from devices on the network to populate the dashboard display. To enable communication between the Monitoring ISE node and the network devices, switches and network access devices (NADs) must be configured properly.

For information on how to configure these devices, see the following:

- [Set the Logging Source-Interface for ISE Monitoring, page C-9](#)
- [Configure NADs for ISE Monitoring, page C-10](#)

## Cisco ISE Dashboard Monitoring

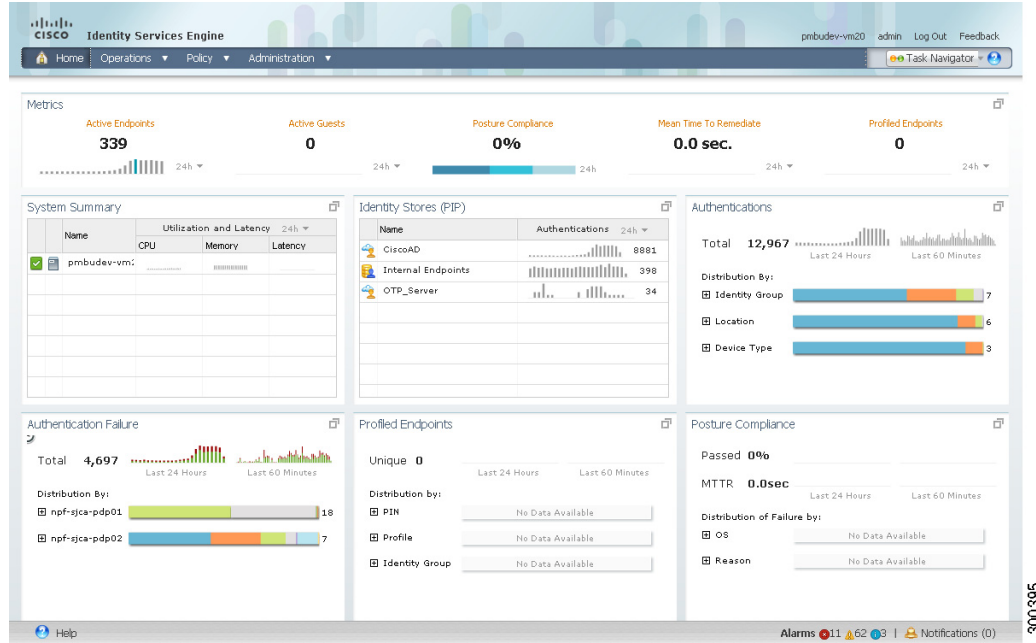
The Cisco ISE dashboard (Home) is the landing page that appears after you log into the Cisco ISE administration console. The dashboard is a centralized management console consisting of metric meters along the top of the window, with dashlets below. This section describes the features functions that comprise the dashboard, as they are represented in the following the graphical user interface elements:

- [Dashlets, page 23-4](#)
- [Metric Meters, page 23-9](#)

Dashboard real-time data provides an at-a-glance status of the devices and users that are accessing your network, as well as a system health overview.

**Note**

You must have Adobe Flash Player installed on the Administration ISE node to be able to view the dashlets and metric meters on the dashboard.

**Figure 23-1** The Cisco ISE Dashboard

The Alarms icon at the bottom right of the Cisco ISE window provides instant access to alarm summaries. Mouse over the Alarms icon to display a pop-up page with a list of recent alarms. You can run filters on the list to view only the alarms of a specific nature. Or, you can drill down for detailed information on individual alarms.

Default alarms include ISE AAA health, ISE process status, ISE system health, and ISE system diagnostics.

#### For more information:

For information on how to interpret and use the data that is shown on the Cisco ISE dashboard, see the following sections:

- [Simplifying Complex Data, page 2-7](#)
- [Managing Alarms, page 23-11](#)
- [Drilling Down for Details, page 2-15.](#)

## Dashlets

Dashlets are individual panels on the dashboard, dashlets summarize important statistics about the devices and users accessing the network. They also provide information about the overall health and security of the network. Each dashlet contains an independent function, and can display the statistical data that is related to its function in various ways. This section explains the purpose and functions of the standard dashlets.



#### Note

You can click a sparkline in a dashlet to generate a report showing relevant logs. Sparklines are a method of visualizing data with vertical lines that depict trends over time. Taller bars mean there was a higher load at a particular time.

Hovering a cursor over the elements of a dashlet brings up a tooltip with detailed information. Tooltip values for a sparkline reflect the specified time interval.

For example, a sparkline with the 24 hour time interval 14 March 3:00 AM, means the sparkline value is calculated based on logs from 3:00 AM to 4:00 AM on that date. Likewise, a sparkline for the 60 minute interval 14 March 3:01:00 AM, means the sparkline value is calculated based on logs from 3:01:00 to 3:02:00 on that date.


## System Summary

The System Summary dashlet focuses on the health of the distributed identity services system deployment. This dashlet provides data for all the nodes on your network, providing an at-a-glance view of node performance, such as CPU, memory, and latency utilization. Sparklines represent a percentage of CPU usage over a specified time increment. For more information, see [Sparklines, page 2-14](#).

The color of the system status icon indicates the health of your system:

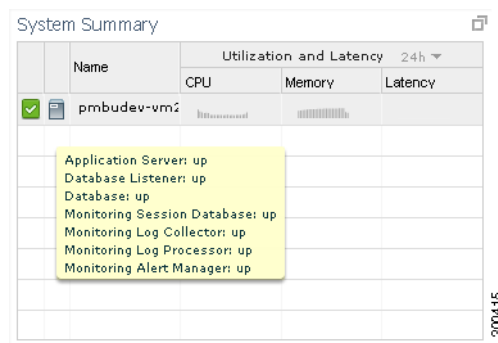
- Healthy = Green
- Warning = Yellow
- Critical = Red
- No information = Gray

**Figure 23-2 System Summary Dashlet**



When you hover the cursor over the health icon, a quick view dialog appears showing detailed information on system health, as shown in [Figure 23-3](#).


**Figure 23-3 System Summary Quick View Display**



## Identity Stores

The Identity Stores dashlet for policy information points (PIP) focuses on the Microsoft Active Directory infrastructure, providing data on the number of authentications for users and devices, as well as the health of the servers. Internal user attributes and the credential information that was most used to authenticate users and hosts for a given time range is also shown.

**Figure 23-4 Identity Stores Dashlet**

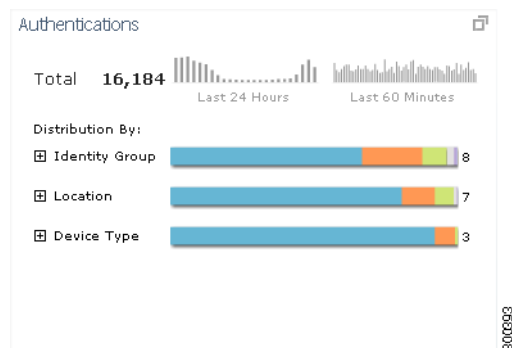


## Authentications

The Authentications dashlet shows passed and failed network authentications, providing data on the user or type of device, location, and the identity group to which the user or device belongs. The sparklines along the top of the dashlet represent distribution over the last 24 hours and the last 60 minutes.

When you hover your cursor over a stack bar or sparkline, a tooltip provides detailed information. [Figure 23-5](#) shows data for all authentication attempts that are made on the network, both passed and failed.

**Figure 23-5 Authentications Dashlet**



## Authentication Failure


The Authentication Failure dashlet focuses on authentication failures, providing information on the nature of the failures. Total counts are shown across the top, while below is a breakdown of statistics by individual node and individual errors.

When you hover your cursor over a stack bar or sparkline, a tooltip provides detailed information. Sparklines use color to convey passed or failed authentication status at a glance. Green represents passed authentications, and red represents failed authentications.

You can quickly assess the nature of failures that occur on your network with the following information:

- Total count of authentication failures in the last 24 hours
- Authentication trend (60 minutes to 24 hours), marking failures with a different color
- Distribution across all ISE nodes
- Distribution of reasons for failure
- Failure reason trend per Policy Service
- Visual health cues: green = pass, yellow = warning, red = failure

**Figure 23-6 Authentication Failure Dashlet**



## Profiled Endpoints

The Profiled Endpoint dashlet focuses on the endpoints on the network that have matched profiles, providing profile data for each endpoint. For example, the statistics allow you to determine the type of device, its location, and its IP address. The sparklines along the top of the dashlet represent endpoint activity over the last 24 hours and last 60 minutes.

You can expand the following data categories for more information:


- PIN—Place in network
- Profile—Profiler policy
- Identity Group—Includes both user and endpoint identity groups, as applicable



### Note

The Profiled Endpoint dashlet represents the total number of endpoints that have been profiled on the network for the last 24 hours, including those that are unknown. It is not a representation of how many endpoints are currently active on the network. Sparkline metrics at the top of the dashlet show time specific values for the last 24 hours and 60 minutes.

For information on Profiled Endpoints dashlet, see the [“Profiled Endpoints Dashlet” section on page 18-6](#).

**Figure 23-7** *Profiled Endpoints Dashlet*

## Posture Compliance

The Posture Compliance dashlet focuses on the health of the network, providing information on the users who are accessing the network and whether they meet posture compliance. Data is shown on the devices that are currently connected to the network. The stack bars show noncompliance statistics that are arranged according to operating system and other criteria. Sparklines represent the percentage of compliant versus noncompliant posture attempts.

- **Passed**—Overall average percentage (%) of compliant posture attempts for the last 24 hours and 60 minutes.

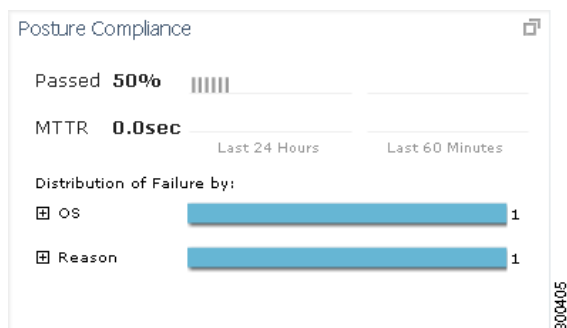


### Note

When you hover a cursor over a sparkline, the tooltip shows the average percentage of compliant posture attempts for a specific time period.

- **MTTR**—Mean Time To Remediate (MTTR). The time difference between an endpoint moving from a non-compliant to a compliant state is used to determine the mean time to remediate (MTTR). The endpoint MAC address is used as the key to calculate the MTTR.
- **OS**—Operating system
- **Reason**—Reason for compliance or noncompliance

For information on Posture Compliance dashlet, see the [“Posture Compliance Dashlet”](#) section on [page 20-8](#).

**Figure 23-8** *Posture Compliance Dashlet*



## Metric Meters

Metric meters are graphs that appear along the top section of the dashboard. Their data is refreshed every minute to provide real-time at-a-glance information.




### Note

You can click the main number display in a metric meter to display relevant detailed report data.

### Active Endpoints

The Active Endpoints metric meter shows data representing the endpoints connected to the network. The change indicator shows the difference in the number of active endpoints between refreshes.


**Figure 23-9 Active Endpoints Metric Meter**



### Active Guests

The Active Guests metric meter shows data representing the current active guests on the network. The change indicator shows the difference in count between the current refresh and the last refresh.

**Figure 23-10 Active Guests Metric Meter**



### Posture Compliance

The Posture Compliance metric meter shows the (average) percentage of hosts that are connected to the system that were compliant with posture rules over the last 24 hours. The black line superimposed on the color-coded bar changes dynamically to show compliance. The color-coded bar beneath remains static, showing a progression from lowest to highest compliancy.


**Figure 23-11 Posture Compliance Metric Meter**



## Mean Time to Remediate

The Mean Time to Remediate metric meter shows the average time that it takes for hosts that are connected to the network to move from a noncompliant state to a compliant state.


**Figure 23-12** Mean Time to Remediate Metric Meter



## Profiled Endpoints

The Profile Endpoints metric meter shows data representing the total number endpoints that have been profiled on the network for the last 24 hours, including those that are unknown.

**Figure 23-13** Profiled Endpoints Metric Meter



# Monitoring the Network

This section discusses the ways in which you can monitor your Cisco ISE network, and covers the following topics:

- [Monitoring Network Process Status, page 23-10](#)
- [Managing Alarms, page 23-11](#)
- [Available Alarm Rules, page 23-18](#)
- [Monitoring Live Authentications, page 23-25](#)
- [Monitoring Data Collections, page 23-28](#)

## Monitoring Network Process Status

You can view process status for the network from the Cisco ISE dashboard using the System Summary dashlet. For example, when processes like the application server or database fail, an alarm is generated and you can view the results using the System Summary dashlet.

To view process status, complete the following steps:

- 
- Step 1** Select **System Summary** dashlet. A detailed real-time report appears.
- Step 2** Review the following information for the processes that are running on the network:
- Name of the process
  - CPU and memory utilization
  - Time since process started running
- 

**For more information:**

See [Appendix A, “User Interface Reference.”](#)

#### Troubleshooting Topics

- [Cisco ISE Monitoring Dashlets Not Visible with Internet Explorer 8, page D-11](#)

## Managing Alarms

This section introduces Cisco ISE alarms, schedules, and rules which you can configure to effectively monitor your network. You can view them and specify alarms to notify you when critical system conditions occur. Notifications automatically appear in the Operations > Alarms > Inbox, but you can also receive notification of events through email and syslog messages.

This section covers the following topics:

- [Understanding Alarms, page 23-11](#)
- [Viewing, Editing, and Resolving Alarms, page 23-13](#)
- [Viewing and Filtering Alarm Schedules, page 23-14](#)
- [Creating, Editing, and Deleting Alarm Schedules, page 23-15](#)
- [Creating, Assigning, Disabling, and Deleting Alarm Rules, page 23-16](#)

## Understanding Alarms

This section covers the basics of alarms and notifications, and covers alarm categories, schedules and rules (or thresholds), alarm notifications, alarm syslog targets, license enforcement alarms, and RADIUS authentication alerts.

There are two basic categories of alarms: alarm rules and system alarms. See [Available Alarm Rules, page 23-18](#), for descriptions of the standard Cisco ISE alarm rules that you can customize for your network.

Default alarms include ISE AAA health, ISE process status, ISE system health, and ISE system diagnostics.

## Alarm Rules

Alarm rules notify you of specified events in log data that is collected from ISE nodes. For example, you can configure alarm rules to notify you about system health, process status, and authentication activity or inactivity.

You define conditions, or rules, on data sets, the time period for applying the alarm rule, the severity of the alarm, and how the notifications should be sent. When alarm rule conditions are met, an alarm is triggered. There are many alarm rule categories that allow you to monitor various types of system behavior.

## System Alarms

System alarms notify you of critical conditions that are encountered on the network. They also provide informational status of system activities, such as data purge events. You cannot create or delete system alarms, because they are predefined. However, you can configure how you want to be notified when they occur, or disable them entirely. When you enable system alarms, they are sent to the alarms inbox.

System alarms do not have an associated schedule and are sent immediately after an event occurs. You can only enable or disable system alarms as an entire group, not on an individual basis. For a list of the various types of system alarms and instructions on how to set them, see [Configuring System Alarm Settings](#), page 23-54.

## Schedules and Alarm Rules

A schedule consists of one or more continuous or noncontinuous periods you define when you create a alarm rule. For example, you can create a schedule that is active from 8:00 a.m. (0800) to 5:00 p.m. (1700) Monday through Friday. When you assign this schedule to an alarm rule, the rule is evaluated and the alarm is generated only during the specified active period.

Alarm rules are evaluated periodically, with the cycle frequency depending on the number of enabled rules. For example, if there are 1–20 enabled alarm rules, the evaluation cycle might occur every two (2) minutes. For 21–50 enabled rules, the evaluation cycle might occur every three (3) minutes, and 51–100 enabled rules every five (5) minutes.

**Note**

---

There is a current limitation that restricts the number of rules to a maximum of 100.

---

When an evaluation cycle begins, each enabled alarm rule is evaluated. If the schedule allows the rule to be executed, the conditions are also evaluated. An alarm is triggered when the conditions of a specified rule are met.

## Alarm Notifications

Alarm notifications are generated based on alarm rule conditions, and are evaluated over a specified time period, or schedule. An alarm notification is sent whenever a rule condition is reached or a system alarm is generated.

Alarm notifications are contained in the following locations:

- Alarm inbox—Contains the information that is on the alarm details page. The alarm details usually include one or more links to relevant reports to help you investigate the event that triggered the alarm. You can add comments, and change the status to indicate that it has been acknowledged or closed.

The alarm inbox can contain up to 5000 alarms, the most recent alarms appearing at the top. Alarms that have been acknowledged or closed are removed from the list.

- Email notification—Contains the information that is on the alarm details page. You can configure a list of recipients, and you can indicate whether you wish to receive notifications in plaintext or HTML format.
- Syslog message—Sent to the Linux or Microsoft Windows machines that you have configured as alarm syslog targets. You can configure up to two alarm syslog targets.
- Alarm summary—Shows a listing of the most recent alarms in a pop-up window when you mouse over the Alarms icon in the right corner of the Global Toolbar at the bottom of the Cisco ISE window. Click an alarm link to view details of the alarm.

For more information, see [Specifying Email Settings, page 23-53](#) and [Configuring System Alarm Settings, page 23-54](#).

## Alarm Syslog Targets

Alarm syslog targets are the destinations to which syslog messages are sent. Alarm notifications are sent in the form of syslog messages. You must have a configured syslog server on your network to receive syslog messages. For more information, see [Configuring Alarm Syslog Targets, page 23-54](#).

## License Enforcement Alarms

License enforcement alarms count concurrent endpoints or users and verify that number against the total amount that is allowed for a particular license. When the count exceeds the amount that is allowed by a license, a syslog is sent indicating that the license count has been exceeded.

## Viewing, Editing, and Resolving Alarms

You can view alarms that met configured alarm rules in the alarms inbox or in the Global Toolbar slide-up window.

The alarm inbox displays a list of recent alarms, which you can select from to view the alarm details. After viewing information for an alarm, you can edit its status, assign the alarm to an administrator, and add notes to track the event.

The Global Toolbar shows the current number of alarms, and the slide-up window displays a read-only list of alarms.



### Note

Move your cursor over any field on the page to view context-sensitive help for the feature.

## Viewing Alarm Summaries

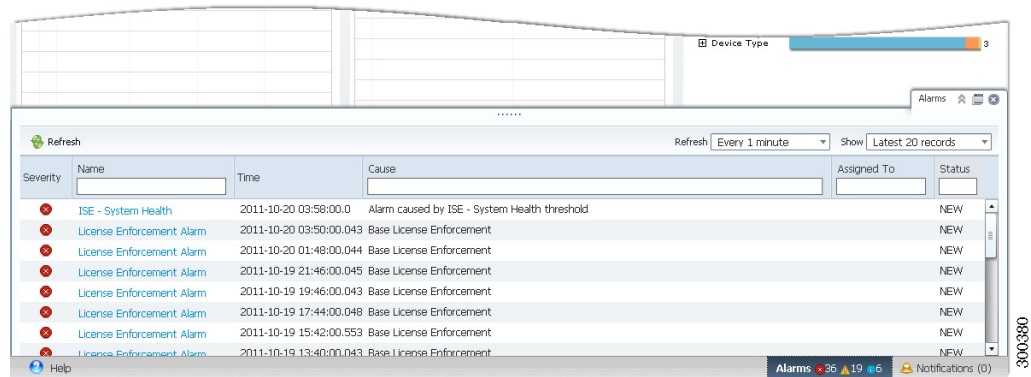
You can view a list of recent alarms from the alarm summary window that you access from the global toolbar. The global toolbar is always available at the bottom of the Cisco ISE window.

**To view a list of alarms, complete the following steps:**

- Step 1** Go to the toolbar at the bottom of the Cisco ISE window and move your mouse over the Alarms icon in the far-right corner. A slide-up window appears, showing a list of recent alarms.
- Step 2** (Optional) Choose the **Refresh Rate** or **Show** options to modify the window display.

- Step 3** (Optional) Choose the **Name**, **Cause**, **Assigned To**, or **Status** option. Enter the required information, and then click the arrow that appears in the right-hand corner of the field.
- Step 4** Click the alarm link to view a detailed description of the event that prompted the alarm. A new page appears.

**Figure 23-14 Alarm Summary Window**



## Using the Alarm Box to View, Edit, and Resolve Alarms

The following task shows you how to use the alarm inbox to view and edit alarms.

**To view and edit an alarm in the alarm inbox, complete the following steps:**

- Step 1** Choose **Operations > Alarms > Inbox**. The Alarms Inbox page appears, with a list of the recent alarms.
- Step 2** To view and edit an alarm, check the check box to the left of the alarm Name, and click **Edit**.
- Step 3** To change the status of the alarm, click the **Status** tab and do the following:
- Choose the appropriate option from the **Status** pull-down menu: **New**, **Acknowledged**, or **Closed**.
  - Assign the alarm to an administrator by entering a name or email address in the **Assigned** field.
  - Add any comments in the Notes field, and click **Submit**.
- You are returned to the alarms inbox.
- Step 4** To resolve an alarm, go back to the inbox, select the check box next to the alarm, and do one of the following:
- To close an alarm, click **Close**, enter Closing Notes in the dialog that appears, and click **Close** again.
  - To delete an alarm, click **Delete**, and verify the action by clicking **Yes** in the dialog that appears.

## Viewing and Filtering Alarm Schedules

You can view a list of all available alarm schedules, and then narrow the results by filtering for specified criteria.

**Note**

Move your cursor over any field on the page to view context-sensitive help for the feature.

**To view and filter alarm schedules, complete the following steps:**

- 
- Step 1** Select **Operations > Alarms > Schedules**. A list of alarm schedules appears.
- Step 2** To search for a specific type of alarm, enter the search criteria in the **Filter** field and click **Go**. The results are displayed.
- Step 3** To return to the complete list of alarms, click **Clear**.
- 

## Creating, Editing, and Deleting Alarm Schedules

You can create alarm schedules to specify when alarm rules are run, and then edit and delete schedules as necessary. Alarm schedules can run at different times of the day throughout a seven-day (week) period. The default alarm schedule is nonstop, monitoring events 24 hours a day, 7 days a week.

**Note**

Move your cursor over any field on the page to view context-sensitive help for the feature.

### Creating an Alarm Schedule

The following task shows you how to create and save alarm schedules.

**To create an alarm schedule, complete the following steps:**

- 
- Step 1** Choose **Operations > Alarms > Schedules**.
- Step 2** Click **Create**.
- Step 3** In the appropriate fields, enter a unique name and a meaningful description to describe the schedule.
- Step 4** Define the days and times for the schedule in one of the following ways:
- Click individual squares to select or deselect the hours and days of the alarm schedule. Squares fill with color when they are selected, and they are blank when they are deselected.  
Click **Clear All** or **Undo All** to clear the schedule and start again.
  - Click **Select All** to create a nonstop alarm schedule that runs 24 hours a day, 7 days a week.  
Use **Clear All** or **Undo All** to clear the schedule and start again.
- Step 5** Click **Submit** to save the schedule, or click **Cancel** to exit without creating a schedule.
- If you submitted the schedule, it appears in the list of schedules.
-

## Editing or Deleting an Alarm Schedule

The following task shows you how to edit and delete an alarm schedule.

**To edit or delete an alarm schedule, complete the following steps:**

- 
- Step 1** Select **Operations > Alarms > Schedules**. A list of schedules appears.
- Step 2** Check the check box to the left of a schedule name, and do one of the following:
- To remove a selected alarm from the list, click **Delete**, and then click **Yes** to confirm the action.
  - To modify a selected alarm, click **Edit**, and then do one of the following:
    - Select and deselect squares to modify the days and times. Squares fill with color when they are selected, and are blank when deselected.
    - Click **Clear All** or **Undo All** to clear the schedule and start again, defining a new schedule.
    - Click **Select All** to create a nonstop alarm schedule that runs 24 hours a day, 7 days a week.
- Step 3** Click **Submit** to save your changes, or **Cancel** to exit without saving the changes.
- 

## Creating, Assigning, Disabling, and Deleting Alarm Rules

You define alarm rule conditions (also known as rules) on data sets, the time period for (applying) the alarm rule, the severity of the alarm, and how the notifications should be sent. Due to the time element, an alarm rule must be linked to an alarm schedule.

This section shows you how to create an alarm rule and assign it to a schedule. It then shows you how to delete an alarm rule.

### Prerequisite

You should have created an alarm schedule, as described in [Creating, Editing, and Deleting Alarm Schedules, page 23-15](#).

### Creating and Assigning an Alarm Rule

One of the requirements for creating an alarm rule is that you assign it to a schedule. The following task shows you how to create an alarm rule, and then assign it to a schedule.

The following default alarm rules are shown in the user interface:

- ISE - AAA Health
- ISE - Process Status
- ISE - System Errors
- ISE - System Health

You can create these alarm rules using the following procedure:

- Passed Authentication
- Failed Authentication
- Authentication Inactivity
- Authenticated But No Accounting Start



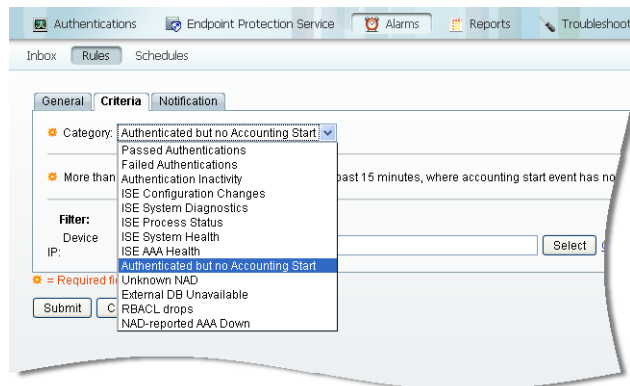
- Unknown NAD
- External DB Unavailable
- RBACL Drops
- NAD-Reported AA Down

**Note**

Move your cursor over any field on the page to view context-sensitive help for that feature.

To create an alarm rule and assign it to a schedule, complete the following steps:

- Step 1** Choose **Operations > Alarms > Rules** and do one of the following:
- To create a copy of an existing alarm rule select the name of the rule, or the check box next to the name, and click **Duplicate**.
  - To create a new rule, click **Create** and proceed with the rest of the steps in this task.
- Step 2** On the General tab, enter a name and description for the alarm rule, and select a schedule from the drop-down list.
- Step 3** Click the **Criteria** tab and do the following:
- Select a rule category from the drop-down list.
  - Specify the required details for the category.
  - (Optional) Specify any other criteria, as desired.



- Step 4** Click the **Notifications** tab and choose a severity level from the drop-down list. Then, specify Email Notification and Syslog Notification, as desired.
- Step 5** Click **Submit** to create the rule, or click **Cancel** to quit without creating the rule.

**For more information:**

See [Available Alarm Rules, page 23-18](#), for descriptions of the standard Cisco ISE alarm rules that you can customize for your network.

## Disabling or Deleting an Alarm Rule

You can disable an alarm rule, which turns it off without removing it. Or you can delete the alarm rule entirely.

**To disable or delete an alarm rule, complete the following steps:**

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Operations &gt; Alarms &gt; Rules</b> .  |
| <b>Step 2</b> | Select the check box next to the alarm rule you want to turn off or remove.  |
| <b>Step 3</b> | To turn off the alarm rule, click <b>Disable</b> .<br><br>To turn back on a disabled alarm rule, select the check box next to the rule and click <b>Enable</b> . |
| <b>Step 4</b> | To permanently remove the selected alarm rule, click <b>Delete</b> . Then click <b>Yes</b> in the dialog prompt to finalize the action.                          |
- 

**For more information:**

See [Available Alarm Rules, page 23-18](#), for descriptions of the standard Cisco ISE alarm rules that you can customize for your network.

## Available Alarm Rules

Cisco ISE provides the following standard categories for alarm rules. You can use the following alarm rules in their default form, or customize them to meet your needs:

- [Passed Authentication, page 23-19](#)
- [Failed Authentication, page 23-19](#)
- [Authentication Inactivity, page 23-20](#)
- [ISE Configuration Changes, page 23-20](#)
- [ISE System Diagnostics, page 23-21](#)
- [ISE Process Status, page 23-21](#)
- [ISE Health System, page 23-21](#)
- [ISE AAA Health, page 23-22](#)
- [Authenticated But No Accounting Start, page 23-22](#)
- [Unknown NAD, page 23-22](#)
- [External DB Unavailable, page 23-23](#)
- [RBACL Drops, page 23-24](#)
- [NAD-Reported AA Down, page 23-24](#)

## Passed Authentication

When Passed Authentication rules are evaluated, passed authentications (such as RADIUS) that occurred during a specified time interval (up to the previous 24 hours) are examined. These authentication records are grouped by a common attribute, such as instance, user, identity group, and so on. The number of records within each of these groups is computed. If the count for any of these groups exceeds the specified rule, an alarm is triggered.

For example, a rule that is configured for passed authentications greater than 1000 in the past 20 minutes for an instance is evaluated. The following table shows the three instances that passed authentications. An alarm was triggered, because at least one instance passed more than 1000 authentications in the past 20 minutes.

Cisco ISE Instance	Passed Authentication Count
New York Cisco ISE	1543
Chicago Cisco ISE	879
Los Angeles Cisco ISE	2096

For example, if you set up another rule for passed authentication less than 3 in the last 20 minutes for a user, the alarm will be generated if the passed authentication is less than 3, provided there was at least one authentication attempt. Zero is not considered as a value for alarm generation.

**Note**

You can specify one or more filters to limit the passed authentications that are considered for rule evaluation. Each filter is associated with a particular attribute in the authentication records, and only the records with a filter value that matches the specified value are counted. If you specify multiple filters, only the records that match all the filter conditions are counted. You can modify the fields in the **Criteria** tab to create a rule with the passed authentication criteria.

**For more information:**

See [Passed Authentications, page A-6 of Appendix A, “User Interface Reference.”](#)

## Failed Authentication

When the Failed Authentication rule is evaluated, failed authentications (such as RADIUS) that occurred during a specified time interval (up to the previous 24 hours) are examined. These authentication records are grouped by a common attribute, such as Cisco ISE instance, user, identity group, and so on. The number of records within each of these groups is computed. If the count that is computed for any of these groups exceeds the specified rule, an alarm is triggered.

For example, the rule reflected in the table is configured with failed authentications greater than 10 in 2 hours for Device IP. If failed authentications have occurred for four IP addresses in the past two hours, such as shown in the following table, an alarm is triggered. At least one Device IP has greater than 10 failed authentications in the past 2 hours.

Device IP	Failed Authentication Count
a.b.c.d	13
e.f.g.h	8

Device IP	Failed Authentication Count
i.j.k.l	1
m.n.o.p	1

**Note**

You can also modify the fields in the Criteria tab to create a rule with the failed authentication criteria.

You can specify one or more filters to limit the failed authentications that are considered for rule evaluation. Each filter is associated with a particular attribute in the authentication records, and only those records whose filter value matches the value that you specify are counted. If you specify multiple filters, only the records that match all the filter conditions are counted.

**For more information:**

See [Failed Authentications](#), page A-8 of [Appendix A, “User Interface Reference.”](#)

## Authentication Inactivity

When the Authentication Inactivity rule is evaluated, it examines authentications (such as RADIUS) that occurred during a specified time interval, up to the previous 31 days. If no authentications have occurred, an alarm is triggered. You can specify filters to generate an alarm if no authentications are seen for a particular instance or device IP address during the time interval.

If the specified time interval for authentication inactivity is less than the time taken to complete an aggregation job, then the alarm is suppressed.

**Note**

You can modify the fields in the Criteria tab to define rule criteria based on authentications that are inactive.

**For more information:**

See [Authentication Inactivity](#), page A-9 of [Appendix A, “User Interface Reference.”](#)

## ISE Configuration Changes

The ISE Configuration Changes alarm is generated when configuration changes, such as adding, updating, or deleting a user or policy, and the like, are made to the server. Cisco ISE then examines the configuration changes made during the interval between the previous and current alarm evaluation cycles. If one or more changes were made, an alarm is triggered. For example, a new user is added, an existing user is updated, and another user is deleted, causing the alarm to be triggered. Installing new software can also trigger a configuration change alarm.

You can specify one or more filters to limit which configuration changes are considered for rule evaluation. Each filter is associated with a particular attribute in the records, and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**For more information:**

See [ISE Configuration Changes](#), page A-9 of [Appendix A, “User Interface Reference.”](#)

## ISE System Diagnostics

When the ISE System Diagnostics rule is evaluated, the system diagnostic records that were generated during the specified interval are examined. If one or more diagnostics were generated at or above the specified security level, an alarm is triggered.

**Note**

Cisco ISE system diagnostics are generated for internal operational diagnostic data, depending on the specified severity level.

You can specify one or more filters to limit which system diagnostic records are considered for rule evaluation. Each filter is associated with a particular attribute in the records and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**For more information:**

See [ISE System Diagnostics, page A-10 of Appendix A, “User Interface Reference.”](#)

## ISE Process Status

When the ISE Process Status rule is evaluated and one or more failures are detected, an alarm is triggered. You can limit the check to particular processes, a particular Cisco ISE instance, or both.

For example, when processes like the application server or database fail, an alarm is generated and you can view the results using the System Summary dashlet.

**Note**

You can modify the fields in the Criteria tab to define rule criteria based on Cisco ISE process status.

**For more information:**

See [ISE Process Status, page A-10 of Appendix A, “User Interface Reference.”](#)

## ISE Health System

When the ISE Health System rule is evaluated, system health parameters are examined as a result of values exceeding the rule for a specified time interval (up to the previous 60 minutes). These health parameters include percentage of CPU utilization, percentage of memory consumption, and so on. If any parameters exceed the rule, an alarm is triggered. By default, the rule applies to all Cisco ISE instances. However, you can choose to limit the check to just a single Cisco ISE instance.

**Note**

You can modify the fields in the Criteria tab to define rule criteria for Cisco ISE system health.

**For more information:**

See [ISE System Health, page A-11 of Appendix A, “User Interface Reference.”](#)

## ISE AAA Health

When the ISE AAA Health rule is evaluated, ISE health parameters that exceeded the rule for the specified time interval (up to the previous 60 minutes) are examined. Cisco ISE monitors the following parameters:

- RADIUS throughput
- RADIUS latency

If any of the parameters exceed the rule, an alarm is triggered. By default, the rule applies to all monitored Cisco ISE instances. However, you can choose to limit the check to just a single Cisco ISE instance.

**Note**

---

You can modify the fields in the Criteria tab as needed.

---

**For more information:**

See [ISE AAA Health](#), page A-11 of [Appendix A, “User Interface Reference.”](#)

## Authenticated But No Accounting Start

When the Authenticated But No Accounting Start rule is evaluated, it determines whether a specified number of authenticated sessions have occurred in the past 15 minutes, where an accounting start event has not been received for a device IP.

These events are grouped by device IP address. If the occurrences for a device IP exceeds the specified of the rule, an alarm is triggered. You can set a filter to limit the evaluation to a single device IP.

**Note**

---

You can modify the fields in the Criteria tab to define rule criteria for authenticated sessions for a device IP.

---

**For more information:**

See [Authenticated But No Accounting Start](#), page A-12 of [Appendix A, “User Interface Reference.”](#)

## Unknown NAD

When the Unknown NAD rule is evaluated, the RADIUS failed authentications that occurred during the specified time interval (up to the previous 24 hours) are examined. The failed authentications with the failure reason “unknown NAD” are identified. The unknown NAD authentication records are grouped by a common attribute, such as Cisco ISE instance, user, and so on. A count of the records within each of the groups is computed, and if the records for any group exceed the specified rule, an alarm is triggered.

Take the following rule for example: Unknown NAD count greater than 5 in the past 1 hour for a Device IP

In our example, after one hour, the failed authentications with an “unknown NAD” failure reason occur for two different device IP addresses. An alarm is triggered as a result, because at least one device IP address has a count greater than 5. The following table shows the data for this example.

Device IP	Count of Unknown NAD Authentication Records
a.b.c.d	6
e.f.g.h	1

You can specify one or more filters to limit failed authentications that are considered for rule evaluation. Each filter is associated with an attribute in the records, and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**Note**

You can modify the fields in the Criteria tab to define rule criteria based on authentications that have failed because of an unknown NAD.

**For more information:**

See [Unknown NAD](#), page A-12 of [Appendix A, “User Interface Reference.”](#)

## External DB Unavailable

When the External DB Unavailable rule is evaluated, RADIUS failed authentications that occur during a specified time interval (up to the previous 24 hours) are examined. The failed authentications with the “external DB unavailable” failure reason are then determined. Authentication records with this failure reason are grouped by a common attribute, such as Cisco ISE instance, user, and so on. A count of the records within each of these groups is computed. If the count of records for any group exceeds the rule, an alarm is triggered.

Take the following rule for example: External DB Unavailable count greater than 5 in the past 1 hour for a Device IP

In our example, after one hour, the failed authentications with an “external DB unavailable” failure reason occur for two different device IP addresses. An alarm is triggered, because at least one device IP address has a count greater than 5. The following table shows the data for this example.

Device IP	Count of External DB Unavailable Authentication Records
a.b.c.d	6
e.f.g.h	1

You can specify one or more filters to limit the failed authentications considered for rule evaluation. Each filter is associated with an attribute in the records, and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**Note**

You can modify the fields in the Criteria tab to define rule criteria based on an external database to which Cisco ISE is unable to connect.

**For more information:**

See [External DB Unavailable](#), page A-13 of [Appendix A, “User Interface Reference.”](#)

## RBACL Drops

When the RBACL Drops rule is evaluated, Security Group Access RBACL drops that occurred during a set time interval (up to the previous 24 hours) are examined. The RBACL drop records are grouped by a particular common attribute, such as NAD, SGT, and so on. The number of records for group is computed. If the count for any group exceeds the rule, an alarm is triggered.

Take the following rule for example: RBACL drops greater than 10 in the past 4 hours by an SGT

In our example, RBACL drops occur for two different source group tags in a four-hour period. An alarm is triggered, because at least one SGT has a count greater than 10. The following table shows the data for this example.

SGT	Count of RBACL Drops
1	17
3	14

You can specify one or more filters to limit the RBACL drop records that are considered for rule evaluation. Each filter is associated with a particular attribute in the RBACL drop records, and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**Note**

You can modify the fields in the Criteria tab to define the RBACL Drops rule.

**For more information:**

See [RBACL Drops, page A-13](#) of [Appendix A, “User Interface Reference.”](#)

## NAD-Reported AA Down

For the NAD-Reported AAA rule, NAD-reported AAA down events occurring during a specified interval (up to the previous 24 hours) are examined. The AAA down records are then grouped by a particular common attribute, such as device IP address or device group, and a count of records within each group is made. If the count for any group exceeds the specified rule, an alarm is triggered.

Take, for example, the following rule configuration: AAA down count greater than 10 in the past 4 hours by a Device IP

In our example, in the past 4 hours, NAD-reported AAA down events occurred for 3 different device IP addresses, triggering an alarm because at least one device IP address has a count greater than 10. The following table shows the data for this example.

Device IP	Count of NAD-Reported AAA Down Events
a.b.c.d	15
e.f.g.h	3
i.j.k.l	9



You can specify one or more filters to limit the AAA down records, that are considered for rule evaluation. Each filter is associated with a particular attribute in the AAA down records and only those records that match the filter condition are counted. If you specify multiple filter values, only the records that match all the filter conditions are counted.

**Note**

You can modify the fields in the Criteria tab to define rule criteria based on the AAA downtime that a network access device reports.

**For more information:**

See [NAD-Reported AAA Downtime](#), page A-14 of [Appendix A, “User Interface Reference.”](#)

## Monitoring Live Authentications

You can monitor recent RADIUS authentications as they happen from the Live Authentications page. The page displays the top 10 RADIUS authentications in the last 24 hours. This section explains the functions of the Live Authentications page.

The Live Authentications page provides a tabular account of recent RADIUS authentications, in the order in which they happen.

**Note**

The Last update shown at the bottom of the Live Authentications page shows the current server date, time, and timezone.

**Figure 23-15** *Live Authentications Page*

Time	Status	Username	Endpoint ID	IP Address	Network Device	Device Port	Authorization Profiles	Identity Group	Posture
Oct 31,11 06:28:56.092 PM	✗	bulmark	90:00:4E:44:D1:19		WNBU-WLC1				
Oct 31,11 06:28:55.290 PM	✓	achakreb	DC:2B:61:53:2C:D4		WNBU-WLC1		PermitAccess	Profiled	NotApp
Oct 31,11 06:28:47.609 PM	✗	bulmark	00:1D:E0:CA:EC:3F		WNBU-WLC1				
Oct 31,11 06:28:42.763 PM	✗	rasinha	24:AB:81:94:C8:84		WNBU-WLC3				
Oct 31,11 06:28:37.830 PM	✓	narakris	C8:BC:C8:E6:60:83		WNBU-WLC1		PermitAccess	Profiled	NotApp
Oct 31,11 06:28:23.327 PM	✗	mandunn	F8:1E:DF:DE:12:40		WNBU-WLC1				
Oct 31,11 06:28:19.418 PM	✓	achakreb	DC:2B:61:53:2C:D4		WNBU-WLC1		PermitAccess	Profiled	NotApp
Oct 31,11 06:28:17.087 PM	✓	pchebrol	00:21:6A:94:78:8A		WNBU-WLC1		PermitAccess	Unknown	NotApp
Oct 31,11 06:28:16.835 PM	✓	wwinslow	28:6A:BA:33:29:37		WNBU-WLC3		PermitAccess	Profiled	NotApp
Oct 31,11 06:28:14.297 PM	✓	#ACSACL#-IP-PRE-PC			bxb22-11-alph...				
Oct 31,11 06:28:14.210 PM	✓	00:30:48:90:F4:58	00:30:48:90:F4:58	10.86.120.21	GigabitEthernet1/0/1<	CWA-Redirect	Unknown	Pending	

The Live Authentication data categories that are shown by default include the following:

- **Time**—Shows the time that the log was received by the collection agent. This column is required and cannot be deselected.
- **Status**—Shows if the authentication was successful or a failure. This column is required and cannot be deselected.
- **Details**—Brings up a report when you click the magnifying glass icon, allowing you to drill down and view more-detailed information on the selected authentication scenario. This column is required and cannot be deselected.
- **Username**—Shows the username that is associated with the authentication.
- **Endpoint ID**—Shows the unique identifier for an endpoint, usually a MAC or IP address.
- **IP Address**—Shows the IP address of the endpoint device.
- **Network Device**—Shows the IP address of the network access device.
- **Device Port**—Shows the port number at which the endpoint is connected.
- **Authorization Profiles**—Shows an authorization profile that was used for authentication.
- **Identity Group**—Shows the identity group that is assigned to the user or endpoint, for which the log was generated.
- **Posture Status**—Shows the status of posture validation and details on the authentication.
- **Event**—Shows the event status.
- **Failure Reason**—Shows a detailed reason for failure, if the authentication failed.


Optionally, you can choose to show the following categories:

- **Auth Method**—Shows the authentication method that is used by the RADIUS protocol, such as Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAPv2), IEE 802.1x or dot1x, and the like.
- **Authentication Protocol**—Shows the authentication protocol used, such as Protected Extensible Authentication Protocol (PEAP), Extensible Authentication Protocol (EAP), and the like.
- **Security Group**—Shows the group that is identified by the authentication log.
- **Server**—Indicates the Policy Service from which the log was generated.
- **Session ID**—Shows the session ID.

You can choose to view all of the columns, or to display only selected data columns. After selecting the columns that you want to appear, you can save your selections.

**To access and modify the Live Authentications display, complete the following steps:**

- 
- Step 1** Choose **Operations > Authentications**. The Live Authentications page appears.
  - Step 2** To change the data refresh rate, select a time interval from the drop-down menu.
  - Step 3** To manually update the data, click the **Refresh** icon on the Live Authentications menu bar.
  - Step 4** To change the number of records that appear, choose one of the following from the **Show** drop-down menu: Latest 20 Records, Latest 50 Records, Latest 100 Records.

- Step 5** To specify a time interval, choose one of the following from the **within** drop-down menu:
- Last 24 hours (the default)
  - Last 12 hours
  - Last 6 hours
  - Last 3 hours
  - Last 60 minutes
  - Last 30 minutes
  - Last 10 minutes
  - Last 5 minutes
  - Last 60 seconds
- Step 6** To change the columns that are shown, click **Add or Remove Columns**, and from the drop-down menu, do any of the following:
- Deselect a check box to remove the column from the display. The check mark disappears.
-  **Note** The Time, Status, and Details columns are essential and cannot be deselected.
- Select an empty check box to add the column to the display.
  - Select **Restore to Default** to reset the display to the default set of columns.
  - Select **Show All Columns** to automatically display all columns. The changes appear automatically.
- Step 7** Click **Save** at the bottom of the drop-down menu to save your modifications, or click **Cancel** to discard your changes.
- 

#### Troubleshooting Topics

- [RADIUS Accounting Packets \(Attributes\) Not Coming from Switch, page D-5](#)
- [RADIUS Server Error Message Entries Appearing in Cisco ISE, page D-14](#)
- [RADIUS Server Connectivity Issues \(No Error Message Entries Appearing in Cisco ISE\), page D-15](#)

## Monitoring Guest Activity

A guest is a type of user that has limited permissions, such as restricted network access and time duration. For example, a guest might not have access to the company's internal network, and the account expires after eight hours.

You can monitor guests that are currently on the network through the authentications that are generated by these accounts. One way to do this would be to set alarm rules for all users of type guest, and then monitor the live authentications.

To monitor guest activity, complete the following steps:

- 
- Step 1** Create an alarm, as described in [Creating, Editing, and Deleting Alarm Schedules](#), page 23-15.
- Step 2** Specify a rule for [Passed Authentication](#), page 23-19, [Failed Authentication](#), page 23-19, or [Authentication Inactivity](#), page 23-20 for all users of type guest, as described in [Creating and Assigning an Alarm Rule](#), page 23-16.
- Step 3** Calculate guest user activity by [Monitoring Live Authentications](#), page 23-25.
- 

#### Troubleshooting Topics

- [RADIUS Accounting Packets \(Attributes\) Not Coming from Switch](#), page D-5
- [RADIUS Server Error Message Entries Appearing in Cisco ISE](#), page D-14
- [RADIUS Server Connectivity Issues \(No Error Message Entries Appearing in Cisco ISE\)](#), page D-15

## Monitoring Data Collections

Monitoring functionality collects log and configuration data from nodes on your Cisco ISE network, stores the data in the Monitoring database, and processes it to generate reports and alarms. You can view the details of the logs that are collected from any of the servers in your deployment.

To monitor data collections for system performance and health, complete the following steps:

- 
- Step 1** Follow the procedure for [Creating, Editing, and Deleting Alarm Schedules](#), page 23-15.
- Step 2** Follow the procedure for [Creating, Assigning, Disabling, and Deleting Alarm Rules](#), page 23-16 using any combination of the following alarm rules:
- [ISE System Diagnostics](#), page 23-21
  - [ISE Process Status](#), page 23-21
  - [ISE Health System](#), page 23-21
  - [ISE AAA Health](#), page 23-22
- Step 3** Follow the procedure for [Specifying Email Settings](#), page 23-53.
- Step 4** Follow the procedure for [Configuring Alarm Syslog Targets](#), page 23-54.
- Step 5** Follow the procedure for [Viewing Log Collections](#), page 23-53.
- 

#### For more information:

See the [Alarms](#), page A-3 of [Appendix A, "User Interface Reference."](#)

#### Troubleshooting Topics

- [RADIUS Accounting Packets \(Attributes\) Not Coming from Switch](#), page D-5
- [RADIUS Server Error Message Entries Appearing in Cisco ISE](#), page D-14
- [RADIUS Server Connectivity Issues \(No Error Message Entries Appearing in Cisco ISE\)](#), page D-15

# Troubleshooting the Network

This section covers the following topics:

- [Viewing and Editing Failure Reasons, page 23-29](#)
- [Troubleshooting Network Access, page 23-29](#)
- [Performing Connectivity Tests, page 23-30](#)
- [Using Diagnostic Troubleshooting Tools, page 23-31](#)

## Viewing and Editing Failure Reasons

The Failure Reason Editor allows you to view and edit the description of a failure reason, as well as providing instructions on how to resolve the problem.

To view and edit failure reasons, complete the following steps:

- 
- Step 1** Choose **Administration > System > Settings**.
- Step 2** In the navigation panel on the left, expand **Monitoring** and select **Failure Reason Editor**. A list of failure reasons appears in the right panel.
- Step 3** To view a failure reason, do one of the following:
- Select a radio button or name link for a failure reason from the list.
  - Enter a text string in the Filter text box, click **Filter**, and select a failure from the results.
- Step 4** To edit a failure reason, do the following:
- a. Click the radio button to the left of the name. The button turns green when selected.
  - b. Click **Edit**.
  - c. In the appropriate field, enter or modify a description, then enter or modify resolution steps.
  - d. Click **Submit** to save your changes, or click **Cancel** to quit without saving any changes.
- 

**For more information:**

See [Troubleshoot, page A-40](#) of [Appendix A, “User Interface Reference.”](#)

## Troubleshooting Network Access

You can troubleshoot network access for a specific user, device, or search criteria based on attributes that are related to the authentication requests. You do this by running an Authentication Failure Code Lookup report.



**Note**

If the MAC address value that you provide is not in the prescribed format, it is assumed to be a username, and a user authentication summary report is run for the chosen time range and protocol.

To troubleshoot network access based on authentication requests, complete the following steps:

- 
- Step 1** Choose **Operations > Reports > Catalog**.
- Step 2** In the Reports list panel on the left, choose **Failure Reason**.
- Step 3** In the panel on the right, click the radio button next to **Authentication Failure Code Lookup**.
- Step 4** Follow the instructions described in [Running, Viewing, and Navigating Reports, page 24-3](#), and consider the following:
- If you provide the Username or MAC Address value in the format aa-bb-cc-dd-ee-ff, the report is run for this MAC address.
  - If you provide the Username or MAC Address value in any other format, the value is considered a username, and the report is run for that user.
  - If you leave the Username or MAC Address field empty, a report using the default parameters is run for the chosen protocol and time range (similar to running a RADIUS authentication report in the catalog pages).
  - If you provide a valid MAC address value for the Username or MAC Address field and choose the Summary View option, an endpoint summary report is run. Irrespective of the protocol that you choose, an endpoint summary report is always run for the RADIUS protocol.
- Step 5** Review the report data to troubleshoot your network access problem.
- 

**For more information:**

See [Troubleshooting RADIUS Authentications, page 23-31](#).

## Performing Connectivity Tests

Failed authentications can be caused by connection problems. Troubleshooting tools functionality allows you to perform connectivity tests to check for connectivity issues. You can enter the hostname or the IP address of the network device with which you are trying to connect and execute the following commands from the web interface: **ping**, **traceroute**, and **nslookup**. The output is displayed in the dashboard window.

To perform connectivity tests, complete the following steps:

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, select **General Tools > Connectivity Tests**.
- Step 3** In the appropriate field, enter the hostname or IP address for a connection that you want to test.
- Step 4** Do any of the following:
- Click **Ping** to view the packets sent and received, packet loss (if any), and the time it takes for the test to complete.
  - Click **Traceroute** to view the intermediary IP addresses (hops) between the Cisco ISE node and the specified hostname or IP address, and the time it takes for each hop to complete.

- Click **Nslookup** to view the server and IP address of your tested domain name server hostname or IP address.
- 

**For more information:**

See [Policy](#), page A-54 of [Appendix A](#), “User Interface Reference.”

## Using Diagnostic Troubleshooting Tools

Diagnostic Tools help you to diagnose and troubleshoot problems on an Cisco ISE network. Detailed instructions on how to resolve the problem is also provided. You can use these tools to evaluate the configuration of any network device on your network, including Security Group Access devices, and troubleshoot passed and failed authentications.

This section covers the following diagnostic procedures:

- [Troubleshooting RADIUS Authentications](#), page 23-31
- [Executing a Network Device Command](#), page 23-33
- [Evaluating a Network Device Configuration](#), page 23-33
- [Troubleshooting Posture Data](#), page 23-34
- [Troubleshooting with TCP Dump](#), page 23-35
- [Comparing SGACL Policies](#), page 23-37
- [Comparing SXP-IP Mappings](#), page 23-38
- [Comparing IP-SGT Pairs](#), page 23-38
- [Comparing SGT Devices](#), page 23-39

## Troubleshooting RADIUS Authentications

Use the RADIUS Authentication diagnostic tool to troubleshoot RADIUS authentications issues.

**To search and select a RADIUS authentication for troubleshooting, complete the following steps:**

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel at the left, select then select **General Tools > RADIUS Authentication Troubleshooting**.
- Step 3** Specify the following information:
- Username—Enter the username of the user whose authentication you want to troubleshoot, or click **Select** to choose the username from a list.
  - MAC address—Enter the MAC address of the device that you want to troubleshoot, or click **Select** to choose the MAC address from a list.
  - Audit Session ID—Enter the audit session ID that you want to troubleshoot.
  - NAS IP—Enter the NAS IP address, or click **Select** to choose the NAS IP address from a list.
  - NAS Port—Enter the NAS port number, or click **Select** to choose a NAS port number from a list.

- **Authentication Status**—Choose the status of your RADIUS authentication from the Authentication Status drop-down list. The available options are as follows:
  - Pass or Fail
  - Pass
  - Fail
- **Time Range**—Select a time range from the drop-down list.



**Note** If you selected a Custom time range, specify the Start Date-Time and End Date-Time.

- **Failure Reason**—View and edit the description of a failure reason.
- **Fetch Number of Records**—Choose the number of records that you want to fetch from the drop-down list: 10, 20, 50, 100, 200, or 500.

**Step 4** Click **Search** to display the RADIUS authentications that match your search criteria.

The Search Result table is populated with the results of your search. The following fields appear in the table: Time, Status, Username, MAC Address, Audit Session ID, Network Device IP, Failure Reason, and Allowed Protocol.

**Step 5** Select a RADIUS authentication record from the table and click **Troubleshoot**.

The Expert Troubleshooter begins to troubleshoot your RADIUS authentication. You are prompted for additional input, if required.

**Step 6** Click **User Input Required**, modify the fields as needed, and then click **Submit**.

The Progress Details page appears, providing a summary. You may be prompted for additional input, if required. If additional input is required, click **User Input Required** and enter the necessary information.

**Step 7** Click **Done**.

The Progress Details page refreshes periodically, displaying tasks that are performed as troubleshooting progresses.

**Step 8** After the troubleshooting is complete, click **Show Results Summary**.

**Step 9** Click **Done** to return to view a diagnosis, steps to resolve the problem, and troubleshooting summary.

#### For more information:

See [RADIUS Authentication Troubleshooting—Progress Details](#), page A-42 of Appendix A, “User Interface Reference.”

#### Troubleshooting Topics

- [RADIUS Accounting Packets \(Attributes\) Not Coming from Switch](#), page D-5
- [RADIUS Server Error Message Entries Appearing in Cisco ISE](#), page D-14
- [RADIUS Server Connectivity Issues \(No Error Message Entries Appearing in Cisco ISE\)](#), page D-15



## Executing a Network Device Command

The Execute Network Device Command diagnostic tool allows you to run the **show** command on any network device from the centralized Cisco ISE dashboard. The results are exactly what you would see on a console, and can be used to identify problems in the configuration of the device.

To run the **show** command on any network device, complete the following steps:

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
  - Step 2** In the panel on the left, choose **General Tools > Execute Network Device Command**.
  - Step 3** Enter the following information in the appropriate fields:
    - Network Device IP—The IP address of the network device
    - Command—A **show** command, such as **show run** or **show vlan**
  - Step 4** Click **Run** to execute the command on the specified network device. The Progress Details page appears, prompting you for additional input.
  - Step 5** Click **User Input Required**, and modify the fields as necessary.
  - Step 6** Click **Submit** to run the command on the network device, and view the output.
- 

**For more information:**

See [Progress Details, page A-44](#) of [Appendix A, “User Interface Reference.”](#)

## Evaluating a Network Device Configuration

You can use this diagnostic tool to evaluate the configuration of a network device and identify any configuration problems. The Expert Troubleshooter compares the configuration of the device with the standard configuration.

To evaluate the configuration of a network device, complete the following steps:

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
  - Step 2** In the panel on the left, choose **General Tools > Evaluate Configuration Validator**.
  - Step 3** Enter the Network Device IP address of the device whose configuration you want to evaluate, and specify other fields as necessary.
  - Step 4** Select configuration options to compare against the recommended template. A green check mark means the option is selected. Click the option again to deselect. Choose from the following:
    - Web Authentication—Check this check box to compare the web authentication configuration.
    - Profiler Configuration—Check this check box to compare the Profiler configuration.
    - CTS—Check this check box if you want to compare Security Group Access configuration.
    - 802.1X—Check this check box if you want to compare the 802.1X configuration, and choose one of the following options:
      - Open Mode
      - Low Impact Mode (Open Mode + ACL)

- High Security Mode (Closed Mode)

- Step 5** Click **Run**. The Progress Details page appears, prompting you for additional input.
- Step 6** Click **User Input Required**, and modify the fields as necessary.
- A new window appears, prompting you to select the interfaces for the configuration analysis.
- Step 7** Check the check boxes next to the interfaces that you want to analyze, and click **Submit**. The Progress Details page appears.
- Step 8** Click **Show Results Summary**.
- 

**For more information:**

See [Progress Details](#), page A-44 of [Appendix A](#), “User Interface Reference.”

## Troubleshooting Posture Data

The Posture Troubleshooting tool helps you find the cause of a posture check failure to identify the following:

- Which endpoints were successful in posture and which were not.
- If an endpoint failed in posture, what steps failed in the posture process.
- Which mandatory and optional checks passed and failed.

You determine this information by filtering requests based on parameters, such as username, MAC address, posture status, and so on.

**To troubleshoot posture incidents, complete the following steps:**

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, choose **General Tools > Posture Troubleshooting**.
- Step 3** Specify the following parameters:
- Username—Enter the username to filter on.
  - MAC Address—Enter the MAC address to filter on, using format: xx-xx-xx-xx-xx-xx
  - Posture Status—Select one of the following authentication status filters:
    - Any
    - Compliant
    - Noncompliant
    - Unknown
  - Failure Reason—Enter the failure reason, or click **Select** to choose a failure reason from a list.
  - Time Range—Select a time range filter from the drop-down list.



**Note** If you selected a Custom time range, specify the Start Date-Time and End Date-Time.

---

- Fetch Number of Records—Select the number of records you want displayed at one time from the drop-down list: 10, 20, 50, 100, 200, or 500.

**Step 4** Click **Search**.

The search results appear in the window, displaying time, status, username, MAC address, and failure reason for each event.

**Step 5** To find an explanation and determine a resolution for an event, select the event in the list and click **Troubleshoot**.**For more information:**

See [Egress SGACL Policy, page A-48](#) of [Appendix A, “User Interface Reference.”](#)

## Troubleshooting with TCP Dump

The tcpdump utility monitors the contents of packets on a network interface that match a given boolean expression. You can use the tcpdump utility to troubleshoot problems on your network. Cisco ISE troubleshooting diagnostic tools provide an intuitive user interface for this utility.

This section shows you how to use the TCP Dump feature directly from the Cisco ISE dashboard, and covers the following topics:

- [Monitoring and Saving Packets, page 23-35](#)
- [Saving a Dump File, page 23-36](#)

**Warning**

**Starting a TCP Dump automatically deletes a previous dump file. To save a previous dump file, perform the [Saving a Dump File, page 23-36](#) before you begin a new TCP Dump session.**

### Monitoring and Saving Packets

This procedure shows you how to configure TCP Dump options and then collect data from the network traffic to help you troubleshooting a network issue.

**To monitor packets on the network, complete the following steps:**

**Step 1** Select **Operations > Troubleshoot > Diagnostic Tools**.**Step 2** In the left-hand navigation pane, select **General Tools > TCP Dump**.**Step 3** Select a **Network Interface** to monitor from the drop-down menu.

This is the interface upon which the network traffic is monitored, or sniffed.

**Step 4** Set Promiscuous Mode to **On** or **Off** by clicking the radio button. The default is On.

Promiscuous mode is the default packet sniffing mode. It is recommended that you leave it set to On. In this mode the network interface is passing all traffic to the system's CPU.

**Step 5** In the Filter field, enter a boolean expression on which to filter.

Standard tcpdump filter expressions are supported, such as the following:

host 10.0.2.1 and port 1812

**Step 6** Click **Start** to begin monitoring the network.



**Note**

An In Progress status appears when you start the utility. You can navigate to another page in the user interface and later return. The In Progress status displays how many bytes generated so far, and is updated every 30 seconds until the process ends or you manually stop the process.

The date, time, format, and size of the file are shown at the bottom of the pane.

**Step 7**

Click **Stop** when you have collected a sufficient amount of data, or wait for the process to conclude automatically after accumulating the maximum number of packets (500,000).



**Note**

You must have Adobe Flash Player installed on the Administration ISE node to be able to view the tcpdump.

**Next Step**

- [Saving a Dump File, page 23-36](#)

**Troubleshooting Topics**

- [Policy Service ISE Node Not Passing Traffic, page D-6](#)

## Saving a Dump File

This procedure shows you how to save a dump file that you can use for troubleshooting purposes.

**Prerequisite**

You should have successfully completed [Monitoring and Saving Packets, page 23-35](#).

**To download a previous dump file, complete the following steps:**

- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the left-hand navigation pane, select **General Tools > TCP Dump**.
- Step 3** Select a Format from the drop-down menu. Human Readable is the default.
- Step 4** Click **Download**, navigate to the desired location and then click **Save**.
- Step 5** To get rid of the previous dump file without saving it first, click **Delete**.

**Figure 23-16 TCP Dump**
**Note**

You can also access tcpdump through the Cisco ISE command line interface (CLI). For more information, see the [Cisco Identity Services Engine CLI Reference Guide, Release 1.1](#).

## Comparing SGACL Policies

For devices that are enabled with the Security Group Access solution, an SGACL is assigned for every source and destination SGT pair based on the egress policy matrix that is configured in Cisco ISE. The egress policy diagnostic tool uses the following process for its comparison:

1. Connects to the device with IP address that you provided, and obtains the access control lists (ACLs) for each source and destination SGT pair.
2. Checks the egress policy that is configured in Cisco ISE and obtains the ACLs for each source and destination SGT pair.
3. Compares the SGACL policy that is obtained from the network device with the SGACL policy that is obtained from Cisco ISE.
4. Displays the source and destination SGT pair if there is a mismatch. Also, displays the matching entries as additional information.

**To compare SGACL policies using the Egress (SGACL) Policy tool, complete the following steps:**

- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, select **Security Group Access Tools > Egress (SGACL) Policy**.
- Step 3** Enter the **Network Device IP** address of the Security Group Access device whose SGACL policy you want to compare.
- Step 4** Click **Run**. The Progress Details page appears, prompting you for additional input.
- Step 5** Click **User Input Required** and modify the fields as necessary.

- Step 6** Click **Submit**. The Progress Details page appears with a brief summary of the results.
- Step 7** Click **Show Results Summary** to view the diagnosis and suggested resolution steps.

**For more information:**

See [Egress SGACL Policy, page A-48](#) of [Appendix A, “User Interface Reference.”](#)

## Comparing SXP-IP Mappings

Security Group Access devices communicate with their peers and learn their SGT values. The Security Exchange Protocol (SXP)-IP Mappings diagnostic tool connects to the device whose IP address you provide and lists the IP addresses of the peer devices and SGT values. You must select one or more of the device peers. This tool connects to each of the peers that you select, and it obtains their SGT values to verify that these values are the same as the values that it learned earlier.

**To compare SXP-IP mappings between a device and its peers, complete the following steps:**

- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, choose **Security Group Access Tools > SXP-IP Mappings**.
- Step 3** Enter the network device IP address of the network device, and click **Select**.
- Step 4** Click **Run**, and then click **User Input Required** and modify the necessary fields.  
The Expert Troubleshooter retrieves Security Group Access SXP connections from the network device and again prompts you to select the peer SXP devices.
- Step 5** Click **User Input Required**, and enter the necessary information.
- Step 6** Check the check box of the peer SXP devices for which you want to compare SXP mappings, and enter the common connection parameters.
- Step 7** Click **Submit**. The Progress Details page appears with a brief summary of the results.
- Step 8** Click **Show Results Summary** to view the diagnosis and resolution steps. The Results Summary page appears.

**For more information:**

See [SXP-IP Mappings, page A-49](#) of [Appendix A, “User Interface Reference.”](#)

## Comparing IP-SGT Pairs

For devices that are enabled with the Security Group Access solution, each user is assigned an SGT value through RADIUS authentication. The IP User SGT diagnostic tool connects to the network device (whose IP address you provide) and performs the following tasks:

1. Obtains a list of all IP-SGT assignments on the network device.
2. Checks the RADIUS authentication and accounting records for each IP-SGT pair to find out the IP-SGT-User value that assigned most recently.
3. Displays the IP-SGT pairs in a tabular format, and identifies whether the SGT values that were most recently assigned and those that are on the device are the same or different.

To compare IP-SGT values on a device with the most assigned SGT, complete the following steps:

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, select **Security Group Access Tools > IP User SGT**.
- Step 3** Specify the following:
- **Network Device IP**—Enter the IP address of the network device.
  - **Username**—Enter the username of the user whose records you want to troubleshoot.
  - **User IP Address**—Enter the IP address of the user whose records you want to troubleshoot.
  - **SGT**—Enter the user SGT value.
- Step 4** Click **Run**. You are prompted for additional input.
- Step 5** Click **User Input Required**, and modify the fields as necessary, and then click **Submit**.
- Step 6** Click **Show Results Summary** to view the diagnosis and resolution steps.
- 

**For more information:**

See [IP User SGT](#), page A-51 of [Appendix A, “User Interface Reference.”](#)

## Comparing SGT Devices

For devices that are enabled with the Security Group Access solution, each network device is assigned an SGT value through RADIUS authentication. The Device SGT diagnostic tool connects to the network device (whose IP address you provide) and performs the following tasks:

1. Obtains the network device SGT value.
2. Checks the RADIUS authentication records to determine the SGT value that was assigned most recently.
3. Displays the Device-SGT pairs in a tabular format, and identifies whether the SGT values are the same or different.

To compare the device SGT with the recently assigned SGT value, complete the following steps:

- 
- Step 1** Choose **Operations > Troubleshoot > Diagnostic Tools**.
- Step 2** In the panel on the left, choose **Security Group Access Tools > Device SGT**.

**Step 3** Specify the following:

- Network Device IPs—Enter the network device IP addresses (whose device SGT you want to compare with a Cisco ISE-assigned device SGT) separated by commas.
- Use Common Connection Parameters—Select this check box to use the following common connection parameters for comparison:
  - Username—Enter the username of the network device.
  - Password—Enter the password.
  - Protocol—Choose the protocol from the Protocol drop-down list box. Valid options are: Telnet and SSHv2. Telnet is the default option. If you choose SSHv2, SSH connections must be enabled on the network device.
  - Port—Enter the port number. The default port number for Telnet is 23 and SSH is 22.
- Enable Password—Enter the enable password if it is different from your login password.
- Same as login password—Select this check box if your enable password is the same as your login password.

**Step 4** Click **Run**.**Step 5** Click **Show Results Summary** to view the results of the device SGT comparison.

The Results Summary page appears with the diagnosis, resolution, and troubleshooting summary.

---

**For more information:**

See [Device SGT](#), page A-53 of [Appendix A, “User Interface Reference.”](#)

## Obtaining Additional Troubleshooting Information

Cisco ISE allows you to download support and troubleshooting information from the administrative user interface. You can use the support bundle to prepare diagnostic information for the Cisco Technical Assistance Center (TAC) to troubleshoot problems with Cisco ISE.

**Note**

The support bundles and debug logs provide advanced troubleshooting information for Cisco TAC and are difficult to interpret. You can use the various reports and troubleshooting tools that Cisco ISE provides to diagnose and troubleshoot issues that you are facing in your network. See [“Troubleshooting the Network”](#) section on page 23-29 for more information.

---

This section contains the following topics:

- [Downloading Support Bundles](#), page 23-40
- [Downloading Debug Logs](#), page 23-43

## Downloading Support Bundles

You can download the support bundle to your local computer as a simple tar.gz file. The support bundle will be named with the date and time stamps in the format `ise-support-bundle-mm-dd-yyyy-hh-mm.tar.gz`. The browser prompts you to save the support bundle to an appropriate location.



You can configure the logs that you want to be part of your support bundle. For example, you can configure logs from a particular service to be part of your debug logs. See the [“Understanding Debug Log Configuration” section on page 14-8](#) for more information.

The logs that you can download are categorized as follows:

- Full configuration database—The Cisco ISE configuration database is downloaded in a human readable XML format. When you are trying to troubleshoot issues, you can import this database configuration in another Cisco ISE node to recreate the scenario.
- Debug logs—Captures bootstrap, application configuration, run time, deployment, monitoring and reporting, and public key infrastructure (PKI) information.

Debug logs provide troubleshooting information for specific ISE components. See the [“Downloading Debug Logs” section on page 23-43](#) for more information. To enable debug logs, see [Chapter 14, “Logging”](#). If you do not enable the debug logs, all the informational messages (INFO) will be included in the support bundle.

- Local logs—Contains syslog messages from the various processes that run on Cisco ISE.
- Core files—Contains critical information that would help identify the cause of a crash. These logs are created when the application crashes and includes heap dumps.
- Monitoring and reporting logs—Contains information about the alerts and reports.
- System logs—Contains Cisco Application Deployment Engine (ADE)-related information.

You can download these logs from the Cisco ISE CLI by using the **backup-logs** command. For more information, see the [Cisco Identity Services Engine CLI Reference Guide, Release 1.1](#).

**Note**

For Inline Posture nodes, you cannot download the support bundle from the Cisco ISE user interface. You must use the **backup-logs** command from the Cisco ISE CLI to download logs for Inline Posture nodes.

If you choose to download these logs from the administrative user interface, you can do the following:

- Download only a subset of logs based on the log type such as debug logs or system logs.
- Download only the latest “n” number of files for the selected log type. This option allows you to control the size of the support bundle and the time taken for download.

Monitoring logs provide information about the monitoring, reporting, and troubleshooting features.

**Prerequisite:**

To perform the operations that are described in the following procedure, you must have Super Admin or System Admin privileges. See [Cisco ISE Admin Group Roles and Responsibilities](#) for more information on the various administrative roles and the privileges that are associated with each of them.

**To download support bundles, complete the following steps:**

- 
- Step 1** Choose **Operations > Troubleshoot > Download Logs**.
- Step 2** From the Appliance node list navigation pane on the left, choose the node from which you want to download the support bundles.
- The Support Bundle page appears, as shown in [Figure 23-17](#) appears. Your support bundle is populated with the parameters that you choose. For specific instructions on Debug Logs, see [“Downloading Debug Logs” section on page 23-43](#).

**Figure 23-17** Download Logs Parameters

**Step 3** Check the check boxes next to the logs that you want to download, and then specify one of the following, as appropriate:

- **All** to include all the selected log files
- **Include most recent** and enter the number of files to include
- **Include files from last** and enter the number of days

If you include all the logs, your support bundle will be excessively large and the download will take a lot of time. To optimize the download process, choose to download only the most recent *n* number of files.

**Step 4** Enter the encryption key for the support bundle, and then re-enter the encryption key.

**Step 5** Click **Create Support Bundle**.

**Step 6** Click **Download** to download the newly created support bundle.

The support bundle is a zip file that is downloaded to the client system that is running your application browser. You must extract the contents of the zip file and untar the tar.gz file to view the logs.

#### Next Step:

See [“Downloading Debug Logs” procedure on page 23-43](#) for information on how to obtain debug logs for specific components.

## Downloading Debug Logs

Debug logs provide troubleshooting information for various Cisco ISE components. While reporting problems, you might be asked to enable these debug logs on ISE and send these logs for diagnosis and resolution of your problems.

Obtaining debug logs is a two-step process:

1. Configure the components for which you want to obtain the debug logs on the Debug Log Configuration page. To configure debug logs for various components, see [“Understanding Debug Log Configuration” section on page 14-8](#) and [“Configuring Debug Log Level” section on page 14-9](#).  
[Table 23-2](#) provides a list of components and the corresponding debug logs that it generates.
2. Download the debug logs.

**Table 23-2**      **Debug Log Configuration: Components and the Corresponding Debug Logs**

Component	Debug Log
runtime-AAA	<i>prrt.log</i>
runtime-config	<i>prrt.log</i>
runtime-logging	<i>prrt.log</i>
NotificationTracker	<i>ise-tracking.log</i>
ReplicationTracker	<i>ise-tracking.log</i>
CacheTracker	<i>ise-tracking.log</i>
pep-auth-manager-test	<i>ise-psc.log</i>
net-securent	<i>ise-psc.log</i>
posture	<i>ise-psc.log</i>
provisioning	<i>ise-psc.log</i>
swiss	<i>ise-psc.log</i>
client	<i>ise-psc.log</i>
prrt-JNI	<i>ise-prrt.log</i>
profiler	<i>profiler.log</i>
cisco-mnt	<i>ise-psc.log</i>
guest	<i>ise-psc.log</i>
guestportal	<i>ise-psc.log</i>
sponsorportal	<i>ise-psc.log</i>
guestauth	<i>ise-psc.log</i>
epm-pap	<i>ise-psc.log</i>
epm-pdp	<i>ise-psc.log</i>
epm-pip	<i>ise-psc.log</i>
epm-pap-api.services	<i>ise-psc.log</i>
org-apache	<i>ise-psc.log</i>
org-apache-digester	<i>ise-psc.log</i>
org-displaytag	<i>ise-psc.log</i>

**Table 23-2**      **Debug Log Configuration: Components and the Corresponding Debug Logs**

Component	Debug Log
org-apache-cxf	<i>ise-psc.log</i>
identity-store-AD	<i>ise-psc.log</i>
mnt-collector	<i>mnt-collector.log</i>
mnt-alert	<i>mnt-alert.log</i>

**Prerequisite:**

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations that are described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See [Cisco ISE Admin Group Roles and Responsibilities](#) for more information on the various administrative roles and the privileges that are associated with each of them.

**To download debug logs, complete the following steps:**

- 
- Step 1** Choose **Operations > Troubleshoot > Download Logs**.
- Step 2** From the Appliance node list navigation pane on the left, choose the node from which you want to download the debug logs.
- The Support Bundle and Debug Logs page appears.
- Step 3** Click the **Debug Logs** tab.
- A list of debug log types and debug logs is displayed. This list is based on your debug log configuration. See “[Understanding Debug Log Configuration](#)” section on page 14-8 for more information.
- Step 4** Click the log file that you want to download and save it to the system that is running your client browser.
- You can repeat this process to download other log files, as needed. The following are additional debug logs that you can download from the Debug Logs page:
- isebootstrap.log—Provides bootstrapping log messages
  - monit.log—Provides watchdog messages
  - pki.log—Provides the third-party crypto library logs
  - iseLocalStore.log—Provides logs about the local store files
  - ad\_agent.log—Provides Microsoft Active Directory third-party library logs
  - catalina.log—Provides third-party logs
-

# Monitoring Administration

The rate and amount of data that is utilized by Monitoring functions requires a separate database on a dedicated node that is used for these purposes.

Like Policy Service, Monitoring has a dedicated database that requires administrators to perform maintenance tasks, such as the topics covered in this section:

- [Backing Up and Restoring the Monitoring Database, page 23-45](#)
- [Viewing Log Collections, page 23-53](#)
- [Specifying Email Settings, page 23-53](#)
- [Configuring System Alarm Settings, page 23-54](#)
- [Configuring Alarm Syslog Targets, page 23-54](#)

## Backing Up and Restoring the Monitoring Database

Monitoring functionality handles large volumes of data. Over time, the performance and efficiency of the node depends on how well you manage that data. To increase efficiency, it is recommended that you back up the data and transfer it to a remote repository on a regular basis. You can automate this task by scheduling automatic backups.



### Note

If you register a secondary Monitoring ISE node, it is recommended that you first back up the primary Monitoring ISE node and then restore the data to the new secondary Monitoring ISE node. This ensures that the history of the primary Monitoring ISE node is in sync with the new secondary node as new changes are replicated. For more information, see [Performing On-Demand Backups, page 23-50](#) and [Restoring the Monitoring Database, page 23-51](#).

Due to the size of the Monitoring database, the backup process can take a while to complete. To save time, you can perform incremental backups, after first completing an initial full database backup. A recommended step, purging unwanted data during the backup process permanently deletes data from the database, and can be configured as an automatic process.

You cannot back up or restore (on-demand and scheduled) the monitoring database for 1.5 hours before or after the purge process. The purge process begins at 4 a.m. (0400), and you cannot back up or restore the monitoring database between 2:30 a.m. (0230) and 5:30 a.m. (0530).



### Warning

**For scheduled backup and purge to work properly for a redundant Monitoring ISE node pair, you must create and specify the same repository, or repositories, for both the primary and secondary nodes. The repository is not automatically synced between the primary and nodes. For more information, see [Configuring Repositories, page 15-3](#).**

This section shows you how to effectively manage the Monitoring database and optimize disk space through the following tasks:

- [Configuring Data Purging, page 23-46](#)
- [Scheduling Full and Incremental Backups, page 23-48](#)
- [Performing On-Demand Backups, page 23-50](#)
- [Restoring the Monitoring Database, page 23-51](#)

**Note**

Every administrator account is assigned one or more administrative roles. Depending upon the roles that are assigned to your account, you may not be able to see or perform the options or perform the procedures that are described in this section. For more information, see [Understanding the Impact of Roles and Admin Groups](#), page 2-19.

## Configuring Data Purging

The purging process allows you to manage the size of the Monitoring database by configuring the following options:

- **Percentage of Disk Space**—Specifies a usage threshold for the Monitoring database as a percentage (%) of total used disk space. The default for the user-configurable option is 80 percent. The maximum value allowed is 100 percent.

When a purge operation triggers, if the actual used database disk space is greater than the configured threshold, the purge operation removes all data from the Monitoring database tables prior to the data retention window (as specified in the Maximum Stored data period field described below).

- **Maximum Stored Data Period**—Specifies the number of months to retain data during a purge. The default is three (3) months. This value is utilized when the disk space usage threshold for purging (Percentage of Disk Space) is met.

**Note**

For this option, each month consists of 30 days. The default of three months equals 90 days.

- **Data Repository**—Specifies the repository in which to backup data prior to purging. You select the repository from the drop-down menu. If a repository is not specified, the data is purged without prior backup. For information on how to specify a repository, see [Configuring Repositories](#), page 15-3.

### Conditions and Rules for Monitoring database Purging

- The purge process executes once every 24 hours at 4 AM.

Purging is always based on the database consumed disk space percentage. Only when the used database space is equal to or exceeds the user specified allowed percentage (by default 80%, which is user configurable), does the purging process begin purging the tables. Otherwise, the purging process is skipped.

- If the Monitoring database disk usage is greater than 95 percent of the threshold setting, an information (INFO) alarm is generated indicating that the database size is too large.
- If the Monitoring database disk usage is greater than 100 percent or above the threshold setting, a backup runs. Monitoring data that is older than the data retention window setting (the default is three months, or 90 days, as each month consists of 30 days) is removed from the database. An information (INFO) alarm is generated after the purge completes.

A purge process runs, creating a status history report that you can view by going to **Operations > Reports > System > Data Management > Monitoring Node > Purging History**. An information (INFO) alarm is generated when the purge completes.

**Note**

If you have not specified a repository, the data is not backed up.

- If the Monitoring database disk usage is greater than 125 percent of the threshold setting, a backup is not performed. Data that is older than the data retention window setting is automatically removed from the database.

A purge process runs, creating a status history report that you can view by going to **Operations > Reports > System > Data Management > Monitoring Node > Purging History**. An information (INFO) alarm is generated when the purge completes.

- You must configure repositories for backup and data purging on both the primary and secondary Monitoring ISE nodes, using the same repositories for each. This is important for the backup and purging features to work properly. Purging takes place on both the primary and secondary nodes of a redundant pair, and the repository is not automatically synced between the nodes.

For example, if the primary node uses two repositories for backup and purging, you must specify the same repositories for the secondary node. For more information, see [Configuring Repositories, page 15-3](#) and [Backing Up and Restoring the Monitoring Database, page 23-45](#).

- If the Cisco ISE node has Administration and Monitor personas (standalone or distributed deployment), a scheduled backup and restore pertains to both Administration and Monitoring data.
- In a distributed environment with a dedicated Monitor ISE node, a scheduled backup includes both Monitor and Administration content. However, since the Administration ISE node is remote on the network, the Administration data that is backed up from the Monitor ISE node might be out of date.

For this reason, it is recommended that you sync the dedicated Monitor ISE node with the Administration ISE node, after the Monitor ISE node restore is complete.

- An on-demand backup only backs up monitoring data.
- You cannot run an on-demand or scheduled backup for 1.5 hours before or after the purge process.

## Purging Unwanted Data

Purging is based on the percentage of consumed disk space for the database. When the consumed disk space for the database is equal to or exceeds the threshold (default 80 percent), the purging process starts. Purging always checks the Monitoring database disk space limit before proceeding.

The maximum stored data period is based on 30-day months, not calendar months. For example, if the server date is April 16, 2011 and the maximum stored data period is set to 1 month, a purge triggered on April 16, 2011 retains data from March 17, 2011 through April 15, 2011.

The purging process triggers once a day at 4:00 AM (a non-configurable default). If disk space usage is met or over the specified limit, the purge executes and runs in the background. If the limit has not been reached, purging is skipped.




### Warning

**For scheduled backup and purge to work properly on the nodes of a Monitoring redundant pair, you must configure the same repository, or repositories, on both the primary and secondary nodes using the CLI. The repositories are not automatically synced between the two nodes.**

### Prerequisite

Configure a data repository where data is backed up prior to purging. You can configure a data repository for a Monitoring ISE node using the **repository** command in the system command line interface (CLI). For more information on CLI commands, see the [Cisco Identity Services Engine CLI Reference Guide, Release 1.1](#).

To configure data purging, complete the following steps:

- 
- Step 1** Choose **Administration > System > Maintenance**.
- Step 2** In the left-hand Maintenance panel, choose **Data Management**, expand the **Monitoring Node**, and select **Data Purging**.
- Step 3** In the Data Purging panel, do the following:
- Enter a numerical percentage value for allowed disk space usage. This threshold triggers a purge when disk space usage meets or exceeds [Conditions and Rules for Monitoring database Purging, page 23-46](#).
  - Choose a data repository from the drop-down list. If no repository is specified, a backup does not occur.
  - Choose the maximum stored data period (in months) from the drop-down list. The default is three months.
-  **Note** For this option, each month consists of 30 days. The default of three months equals 90 days.
- 
- Step 4** Click **Submit**.
- Step 5** Verify the success of the data purge by viewing the Purging History report. For more information, see [System Reports, page 24-11](#).
- 

### Next Steps

Proceed with one of the of the following tasks:

- [Scheduling Full and Incremental Backups, page 23-48](#)
- [Performing On-Demand Backups, page 23-50](#)

## Scheduling Full and Incremental Backups

You can schedule full-backups to run automatically at a specified day and time. You need to perform a full database backup before you begin scheduling incremental backups. Incremental backups backup only the data that has changed since the last backup, allowing you to save time and disk space.



### Note

You cannot schedule a full or incremental backup for 1.5 hours before or after the purge process (between 2:30 a.m. (0230) and 5.30 a.m. (0530)). Also, full and incremental backups must be scheduled two hours apart.

### Prerequisite

Before you begin either procedure you should have successfully set purging options, as described in [Configuring Data Purging, page 23-46](#).

## Scheduling Full Backups

By default, scheduled monthly backups occur on last day of month, scheduled weekly backups occur last day of week, and scheduled daily backups occur at the time specified.



To configure a full database backup, complete the following steps:

- 
- Step 1** Choose **Administration > System > Maintenance**.
- Step 2** In the left-hand Maintenance panel, choose **Data Management**, expand the **Monitoring Node**, and select **Scheduled Backup**.
- Step 3** Enter an **Encryption Key**. This key is used to encrypt and decrypt the backup file.
- Step 4** Make sure that the **Incremental Backup** radio button is set to **On**. If it is set to **Off**, the full database backup will not start.
- Step 5** Specify **Configure Full Monitoring Database Backup** options as follows:
- Select a data repository from the drop-down list.  
For information on how to specify a repository, see [Configuring Repositories, page 15-3](#).
  - Schedule the time that the backup will be performed by selecting hours, minutes, and AM or PM from the drop-down lists.
  - Select the frequency of the backup from the drop-down list. Determine if it will be daily, weekly, or monthly.
- Step 6** Click **Submit**.
- Step 7** Verify the success of the backup by viewing the Backup History report. For more information, see [System Reports, page 24-11](#).
- Step 8** If the backup fails, check the following:
- Make sure that no other job or backup is running in parallel.
  - Check the available disk space for the configured repository.
    - If the database disk usage is greater than 120 GB, but less than 150 GB (125 percent of the total database size of 120 GB), monitoring functions may wait until another purge is performed before continuing with the backup.
    - If the database disk usage is greater than 150 GB, a purge occurs whether or not a backup has occurred, to reduce the database disk usage is below 120 GB.
  - Verify whether the repository is configured.
- 

#### Next Step

- [Restoring the Monitoring Database, page 23-51](#)

## Scheduling Incremental Backups

Incremental backups save time and disk space, and allow you to configure the frequency and time backups occur. Incremental backups store data updates in a separate location, so it is important that you perform an initial full backup before starting incremental backups.



#### Note

Perform a full database backup before scheduling incremental backups. If you disable the incremental backup feature, run a full backup before returning to incremental backups. This precaution will ensure that all your data is complete and current.

### Prerequisites

You should have successfully run a full backup of the Monitoring database, before you attempt to perform an incremental backup. For more information, see [Scheduling Full Backups, page 23-48](#) or [Performing On-Demand Backups, page 23-50](#).

To schedule incremental backups, complete the following steps:

- 
- Step 1** Choose **Administration > System > Maintenance**.
  - Step 2** In the left-hand Maintenance panel, choose **Data Management**, expand the **Monitoring Node**, and select **Scheduled Backup**.
  - Step 3** Enter an **Encryption Key**. This key is used to encrypt and decrypt the backup file.
  - Step 4** Make sure that the **Incremental Backup** radio button is set to **On**. If it is set to **Off**, the incremental backup will not start.
  - Step 5** Specify **Configure Incremental Monitoring Database Backup** options as follows:
    - a. Select a data repository from the drop-down list.  
For information on how to specify a repository, see [Configuring Repositories, page 15-3](#).
    - b. Schedule the time that the backup will be performed by selecting hours, minutes, and AM or PM from the drop-down lists.
    - c. Select the frequency of the backup from the drop-down list. Determine if it will be daily, weekly, or monthly.  
Scheduled monthly backups occur on last day of month; scheduled weekly backups occur last day of week; and scheduled daily backups occur at the time specified.
  - Step 6** Click **Submit**.
  - Step 7** Verify the success of the backups by viewing the Backup History report. For more information, see [System Reports, page 24-11](#).
- 

### Next Steps

Restore data from an incremental backup, start with the initial full backup and continue through the latest incremental backup. For more information on restoring data, see [Restoring the Monitoring Database, page 23-51](#).

## Performing On-Demand Backups

You can perform an immediate full backup of the Monitoring database at any time, as long as no other backup is already in progress. If another backup process is running, you must wait for it to complete before you can start an on-demand backup.



#### Note

An on-demand backup only backs up monitoring data.



#### Note

You cannot perform an on-demand backup for 1.5 hours before or after the purge process (between 2:30 a.m. (0230) and 5:30 a.m. (0530)).

**Prerequisite**

You should have configured data purging, as described in [Purging Unwanted Data, page 23-47](#).

**To generate a full backup immediately, complete the following steps:**

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Administration &gt; System &gt; Maintenance</b> .  |
| <b>Step 2</b> | In the left-hand Maintenance panel, choose <b>Data Management</b> , expand the <b>Monitoring Node</b> , and select <b>Full Backup On Demand</b> .  |
| <b>Step 3</b> | Select a data repository from the drop-down list.<br><br>If no repository is specified, the data will be purged and no backup occurs. For information on how to specify a repository, see <a href="#">Configuring Repositories, page 15-3</a> .<br><br>Enter an <b>Encryption Key</b> . This key is used to encrypt and decrypt the backup file. |
| <b>Step 4</b> | Click <b>Backup Now</b> .  |
| <b>Step 5</b> | Verify the success of the backup by viewing the Backup History report. For more information, see <a href="#">System Reports, page 24-11</a> .  |
- 

**Next Step**

- [Restoring the Monitoring Database, page 23-51](#)

## Restoring the Monitoring Database

You can restore data from an incremental or full backup using the Data Restore feature. If you choose to restore incremental backup data, the full data backup is restored first, followed by all subsequent incremental backups in sequential order.

The process for restoring the Monitoring database is different depending on the type of deployment. The following sections explain how to restore the Monitoring database in a standalone deployment and distributed deployments.

**Standalone Deployment Restore**

In a standalone deployment where Administration and Monitoring personas are both running on the Cisco ISE node, restoring a Monitoring database backup also restores the Administration database. For more information, see [Restoring a Monitor backup in a Standalone Environment, page 23-52](#).

**Distributed Deployment Restore**

There are two possible scenarios for restoring a Monitoring backup:

- Restoring a Monitoring backup to a Cisco ISE node with Administration and Monitoring personas.
- Restoring a Monitoring backup to a Cisco ISE node with only a Monitoring persona.

For more information, see [Restoring a Monitor Backup in a Distributed Environment, page 23-52](#).

**Warning**

**If you attempt to restore data to a node other than the one from which the data was taken, you must configure the logging target settings to point to the new node. This ensures that the monitoring syslog are sent to the correct node. For more information, see [Configuring Alarm Syslog Targets, page 8-22](#).**

## Restoring a Monitor backup in a Standalone Environment

Use the following procedure to restore the Monitoring database to a standalone node.

### Prerequisites

You should have successfully performed the following procedures:

- [Configuring Data Purging, page 23-46](#)
- [Scheduling Full and Incremental Backups, page 23-48](#) or [Performing On-Demand Backups, page 23-50](#).

**To restore incremental and full backup data, complete the following steps:**

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Administration &gt; System &gt; Maintenance</b> .  |
| <b>Step 2</b> | In the left-hand Maintenance panel, choose <b>Data Management</b> , expand the <b>Monitoring Node</b> , and select <b>Data Restore</b> .   |
| <b>Step 3</b> | Select the name of an incremental or full backup from the list.<br><br>If an incremental backup file is selected, all previous incremental backups are shown, along with the initial full backup.<br><br>Enter the <b>Encryption Key</b> used during the backup. |
| <b>Step 4</b> | Click <b>Restore</b> .   |
- 

## Restoring a Monitor Backup in a Distributed Environment

Use the procedures outlined in this section to restore a Monitor backup in a distributed environment.

### Prerequisites

You should have successfully performed the following procedures:

- [Configuring Data Purging, page 23-46](#)
- [Scheduling Full and Incremental Backups, page 23-48](#) or [Performing On-Demand Backups, page 23-50](#).

**To restore a Monitor backup to a Cisco ISE node with Administration and Monitor personas:**

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Prepare to promote another Cisco ISE node as the primary Administration ISE node, by syncing the node with the existing primary node you want to backup. For more information, see <a href="#">Synchronizing Primary and Secondary Nodes in a Distributed Environment, page 15-12</a> .<br><br>This ensures that the configuration of the Cisco ISE node you are going to promote is up to date |
| <b>Step 2</b> | Promote the newly synced Administration ISE node to primary status. For more information, see <a href="#">Configuring a Primary Administration Cisco ISE Node, page 9-11</a> .  |
| <b>Step 3</b> | Prepare to deregister the node to be backed up by assigning the Monitor persona to another node in the deployment. For more information, see <a href="#">Changing Node Personas and Services, page 9-23</a> .   |



---

**Note** A deployment must have at least one functioning Monitor ISE node.

---

- Step 4** Deregister the node to be backed up. For more information, see [Removing a Node from Deployment, page 9-26](#).
- Step 5** Restore the Monitor backup to the newly deregistered node, as described in [Restoring a Monitor backup in a Standalone Environment, page 23-52](#).
- Step 6** Register the newly restored node with the current Administration ISE node. For more information, see [Registering and Configuring a Secondary Node, page 9-13](#).
- Step 7** Promote the newly restored and registered node as the primary Administration ISE node. For more information, see [Configuring a Primary Administration Cisco ISE Node, page 9-11](#).

**To restore a Monitor backup to a Cisco ISE node with only a Monitor persona:**

- Step 1** Prepare to deregister the node to be restored by assigning the Monitor persona to another node in the deployment. For more information, see [Changing Node Personas and Services, page 9-23](#).



**Note** A deployment must have at least one functioning Monitor ISE node.

- Step 2** Deregister the node to be restored. For more information, see [Removing a Node from Deployment, page 9-26](#).



**Note** Wait until the deregistration is complete before proceeding with the restore. The node must be in a standalone state before you can continue with the restore.

- Step 3** Restore the Monitoring backup to the newly deregistered node, as described in [Restoring a Monitor backup in a Standalone Environment, page 23-52](#).
- Step 4** Register the newly restored node with the current Administration ISE node. For more information, see [Registering and Configuring a Secondary Node, page 9-13](#).
- Step 5** Promote the newly restored and registered node as the primary Administration ISE node. For more information, see [Configuring a Primary Administration Cisco ISE Node, page 9-11](#).

## Viewing Log Collections

Monitoring functions collects log and configuration data, stores the data, and then processes the collected data to generate reports and alarms. You can view the details of the logs that are collected from any of the servers in your deployment. For more information, see [Chapter 14, “Logging.”](#)

## Specifying Email Settings

For use with monitoring log messages, you can specify the email server email address and the name that is displayed for this address. For more information, see [Configuring Email Settings, page 8-20](#).



**Note** Depending upon the roles that are assigned to your account, you may or may not be able to perform the operations or see the options that are described in the following procedure. For more information, see [Understanding the Impact of Roles and Admin Groups, page 2-19](#).

## Configuring System Alarm Settings

System alarms notify you of critical conditions that are encountered. System alarms are standard and cannot be created or deleted. You can enable and disable system alarms, and you can configure how you receive notification. You can choose to send alarm notifications through email and as syslog messages.

For instructions on how to set system alarms, see [Configuring System Alarm Settings, page 8-21](#).

**Note**

To send syslog messages successfully, you must configure alarm syslog targets, which are syslog message destinations. See [Configuring Alarm Syslog Targets, page 8-22](#).

**For more information:**

See [System Alarm Settings, page A-60](#) of Appendix A, “User Interface Reference.”

## Configuring Alarm Syslog Targets

If you configure monitoring functions to send system alarm notifications as syslog messages, you need a syslog target to receive the notification. Alarm syslog targets are the destinations where alarm syslog messages are sent.

You must also have a system that is configured as a syslog server to be able to receive syslog messages. You can create, edit, and delete alarm syslog targets. For more information, see [Configuring Alarm Syslog Targets, page 8-22](#).

**Warning**

**Cisco ISE monitoring requires that the logging source-interface configuration use the network access server (NAS) IP address. For information on how to configure a switch for Cisco ISE monitoring, see [Set the Logging Source-Interface for ISE Monitoring, page C-9](#).**

**For more information:**

See [Alarm Syslog Targets, page A-59](#) of Appendix A, “User Interface Reference.”