



Zenoss Resource Manager Administration Guide

Release 5.1.x

Zenoss, Inc.

www.zenoss.com

Zenoss Resource Manager Administration Guide

Copyright © 2016 Zenoss, Inc. All rights reserved.

Zenoss and the Zenoss logo are trademarks or registered trademarks of Zenoss, Inc., in the United States and other countries. All other trademarks, logos, and service marks are the property of Zenoss or other third parties. Use of these marks is prohibited without the express written consent of Zenoss, Inc., or the third-party owner.

Amazon Web Services, AWS, and EC2 are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

Flash is a registered trademark of Adobe Systems Incorporated.

Oracle, the Oracle logo, Java, and MySQL are registered trademarks of the Oracle Corporation and/or its affiliates.

Linux is a registered trademark of Linus Torvalds.

RabbitMQ is a trademark of VMware, Inc.

SNMP Informant is a trademark of Garth K. Williams (Informant Systems, Inc.).

Sybase is a registered trademark of Sybase, Inc.

Tomcat is a trademark of the Apache Software Foundation.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All other companies and products mentioned are trademarks and property of their respective owners.

Part Number: 1012.16.291

Zenoss, Inc.
11305 Four Points Drive
Bldg 1 - Suite 300
Austin, Texas 78726

Contents

| | |
|---|-----------|
| About this guide..... | 7 |
| Chapter 1: Using Resource Manager..... | 8 |
| Initial Login..... | 8 |
| Interface and Navigation..... | 10 |
| Customizing the Dashboard..... | 14 |
| Search..... | 17 |
| Navigating the Event Console..... | 18 |
| Running a command from the browser interface..... | 22 |
| Running a Command from the CLI..... | 22 |
| Working with Triggers and Notifications..... | 23 |
| Advanced User Interface Configuration..... | 37 |
| LDAP Authentication..... | 37 |
| Verifying Connectivity and Credentials Outside of Resource Manager..... | 41 |
| Configuring Local Authentication as a Fallback..... | 41 |
| Chapter 2: Visualizing your Environment..... | 43 |
| Dynamic Service View..... | 43 |
| Datacenter View..... | 47 |
| Chapter 3: Adding, Discovering and Modeling Devices..... | 51 |
| Add a Single Device..... | 51 |
| Add Multiple Devices..... | 52 |
| Discovering Devices..... | 52 |
| Modeling Devices..... | 55 |
| About Modeler Plugins..... | 58 |
| Debugging the Modeling Process..... | 59 |
| Chapter 4: Working with Devices..... | 60 |
| Viewing the Device List..... | 60 |
| Working with Devices..... | 61 |
| Managing Devices and Device Attributes..... | 73 |
| Chapter 5: Configuration Properties..... | 77 |
| Configuration Properties Inheritance and Override..... | 77 |
| Configuration Property Types..... | 80 |
| Device Configuration Properties..... | 80 |
| Event Configuration Properties..... | 89 |
| Network Configuration Properties..... | 90 |
| Chapter 6: Monitoring Templates..... | 92 |
| Creating Templates..... | 92 |
| Renaming Templates..... | 92 |
| Template Binding..... | 92 |

| | |
|--|------------|
| Example: Defining Templates in the Device Hierarchy..... | 94 |
| Example: Applying Templates to Multiple Areas in the Device Hierarchy..... | 94 |
| Chapter 7: Basic Monitoring..... | 95 |
| Availability Monitoring..... | 95 |
| Monitoring Using ZenCommand..... | 104 |
| SNMP Monitoring..... | 106 |
| Monitoring Devices Remotely Through SSH..... | 107 |
| Chapter 8: Performance Monitoring..... | 110 |
| About Monitoring Templates..... | 111 |
| Template Binding..... | 111 |
| Data Sources..... | 112 |
| Data Points..... | 113 |
| Data Point Aliases..... | 114 |
| Thresholds..... | 116 |
| Performance Graphs..... | 124 |
| Performance Data Retention..... | 126 |
| Chapter 9: Event Management..... | 128 |
| Basic Event Fields..... | 128 |
| Other Fields..... | 130 |
| Details..... | 131 |
| De-Duplication..... | 131 |
| Auto-Clear Correlation..... | 132 |
| Event Consoles..... | 133 |
| Event Sources..... | 139 |
| Creating Events Manually..... | 139 |
| Event Classes..... | 141 |
| Mapping and Transformation..... | 142 |
| Event Life Cycle..... | 145 |
| Capturing Email Messages as Events..... | 147 |
| SNMP Traps and Event Transforms..... | 148 |
| Chapter 10: Production States and Maintenance Windows..... | 153 |
| Production States..... | 153 |
| Maintenance Windows..... | 154 |
| Chapter 11: Organizers and Path Navigation..... | 157 |
| Classes..... | 157 |
| Groups..... | 159 |
| Systems..... | 160 |
| Locations..... | 160 |
| Component Groups..... | 163 |
| Inheritance..... | 164 |
| Chapter 12: Distributed Monitoring..... | 166 |
| About Collectors..... | 166 |
| About Hubs..... | 166 |

| | |
|---|------------|
| About Resource Pools..... | 167 |
| Navigating Collectors and Hubs..... | 167 |
| Collector Data Storage..... | 169 |
| Deleting Collectors..... | 169 |
| Adding Devices to Collectors..... | 169 |
| Managing Collector Daemons..... | 169 |
| Chapter 13: User Commands..... | 171 |
| Defining Global User Commands..... | 171 |
| Running Global User Commands..... | 172 |
| Defining User Commands for a Single Device..... | 172 |
| Running User Commands for a Single Device..... | 173 |
| Defining User Commands for All Devices in an Organizer..... | 173 |
| Running User Commands for Devices in an Organizer..... | 174 |
| User Command Example: Echo Command..... | 174 |
| Chapter 14: Managing Users..... | 175 |
| Creating User Accounts..... | 175 |
| Editing User Accounts..... | 175 |
| User Groups..... | 178 |
| Roles..... | 179 |
| Device Access Control Lists..... | 179 |
| Chapter 15: Reporting..... | 183 |
| Cisco UCS Reports..... | 183 |
| Device Reports..... | 184 |
| Enterprise Reports..... | 186 |
| Event Reports..... | 195 |
| Performance Reports..... | 196 |
| Chapter 16: ZenPacks..... | 199 |
| Displaying the list of installed ZenPacks..... | 199 |
| ZenPack information resources..... | 200 |
| Preparing to install or upgrade a ZenPack..... | 200 |
| Installing or upgrading a ZenPack..... | 201 |
| Removing a ZenPack..... | 202 |
| Creating a new ZenPack..... | 204 |
| Chapter 17: General Administration and Settings..... | 205 |
| Recovering from a POSKey error during login..... | 205 |
| Events Settings..... | 205 |
| Rebuilding the Events Index..... | 207 |
| Audit Logging..... | 207 |
| Setting Portlet Permissions..... | 212 |
| Support Bundles..... | 213 |
| Backup and Restore..... | 215 |
| Snapshot and Rollback..... | 219 |
| Adding storage for application data or snapshots..... | 220 |
| Working with the Job Manager..... | 221 |
| Versions and Update Checks..... | 223 |

| | |
|---|------------|
| Appendix A: SNMP Device Preparation..... | 224 |
| Net-SNMP..... | 224 |
| SNMP v3 Support..... | 224 |
| Community Information..... | 225 |
| System Contact Information..... | 226 |
| Extra Information..... | 226 |
| | |
| Appendix B: Syslog Device Preparation..... | 227 |
| Forwarding Syslog Messages from UNIX/Linux Devices..... | 227 |
| Forwarding Syslog Messages from a Cisco IOS Router..... | 227 |
| Forwarding Syslog Messages from a Cisco CatOS Switch..... | 228 |
| Forwarding Syslog Messages using Syslog-ng..... | 228 |
| | |
| Appendix C: TALES Expressions..... | 230 |
| Examples..... | 230 |
| TALES Device Attributes..... | 231 |
| TALES Event Attributes..... | 232 |
| | |
| Appendix D: Monitoring Large File Systems..... | 234 |
| Procedure..... | 234 |
| | |
| Appendix E: Managing Multi-Realm Networks..... | 235 |
| Example System..... | 235 |
| System Setup..... | 236 |
| Notes..... | 237 |
| | |
| Glossary..... | 238 |

About this guide

Zenoss Resource Manager Administration Guide provides an overview of Zenoss Resource Manager (Resource Manager) architecture and features, as well as procedures and examples to help use and configure the system.

Related publications

| Title | Description |
|---|--|
| <i>Zenoss Resource Manager Administration Guide</i> | Provides an overview of Resource Manager architecture and features, as well as procedures and examples to help use the system. |
| <i>Zenoss Resource Manager Configuration Guide</i> | Provides required and optional configuration procedures for Resource Manager, to prepare your deployment for monitoring in your environment. |
| <i>Zenoss Resource Manager Installation Guide</i> | Provides detailed information and procedures for creating deployments of Control Center and Resource Manager. |
| <i>Zenoss Resource Manager Planning Guide</i> | Provides both general and specific information for preparing to deploy Resource Manager. |
| <i>Zenoss Resource Manager Release Notes</i> | Describes known issues, fixed issues, and late-breaking information not already provided in the published documentation set. |
| <i>Zenoss Resource Manager Upgrade Guide</i> | Provides detailed information and procedures for upgrading deployments of Resource Manager. |

Additional information and comments

If you have technical questions about this product that are not answered in this guide, please visit the [Zenoss Support](#) site or contact Zenoss Support.

Zenoss welcomes your comments and suggestions regarding our documentation. To share your comments, please send an email to docs@zenoss.com. In the email, include the document title and part number. The part number appears at the end of the list of trademarks, at the front of this guide.

1

Using Resource Manager

The browser interface of Resource Manager provides a variety of means for navigating and managing your environment.

Initial Login

The first time you log in to Resource Manager, you will immediately be taken to a startup wizard where you will perform the following tasks:

- Set your admin password
- Set your personal login
- Discover devices (optional)
- Add Infrastructure (optional)

- 1 Launch your Resource Manager application the first time by clicking on the Virtual Host Name in Control Center. You will be presented with the following page showing you the initial steps to follow:

Figure 1: Installation Wizard

Zenoss Installation Wizard

This wizard will guide you through the initial setup of Zenoss. Click **Get Started** to begin.

The screenshot displays the Zenoss Installation Wizard interface. It features three main steps in light blue boxes:

- Step 1: Setup Users** (with a group of people icon): Set the admin password and create your user account.
- Step 2: Network Discovery** (with a computer icon): Discover devices to monitor.
- Step 3: Add Infrastructure** (with a server rack icon): Manually add the devices in your infrastructure.

A grey button labeled "Get Started »" is positioned at the bottom right of the wizard area.

- 2 Click **Get Started** to begin the wizard. The Setup Users page appears.

Figure 2: Setup Users

Step 1: Setup Users

Set admin password

The admin account has extended privileges, similar to Linux's root or Windows' Administrator. Its use should be limited to administrative tasks.

Password Must:

- Contain 8 or more characters
- Contain at least one number
- Contain at least one upper and lower case character

Admin password:

Confirm password:

Create your account

Enter information for your personal user account. You'll use this to perform most tasks.

Username:

Password:

Retype password:

Your email:

« Previous
Next »

- 3 Set your admin password and create your personal user account. Click **Next**. The Network Discovery page appears.

Figure 3: Network Discovery

Step 2: Network Discovery

Devices found via Discovery will be placed in the /Discovered Device Class

Networks/Range

Enter one or more networks (such as 10.0.0.0/24) or IP ranges (such as 10.0.0.1-50):

Discover

SNMP

Community Strings:

SSH Authentication

Username:

Password:

Windows Authentication

Administrator Username:

Password:

Discoveries

| Status | Networks | Credentials | Duration | Job Log | Remove |
|---|----------|-------------|----------|---------|--------|
| <i>Add network discoveries using the above form</i> | | | | | |

« Previous
Next »

- 4 You can discover devices by setting networks or IP ranges on this page. If you are not ready to discover devices, you can skip this page and add devices later. For more information about the fields on this page see [Adding, Discovering and Modeling Devices](#) on page 51. When you are ready to continue, click **Next**. The Add Infrastructure page appears.

Figure 4: Add Infrastructure

Step 3: Add Infrastructure

Category

- CiscoUCS
- ControlCenter
- HTTP
- KVM
- Network
- Ping
- Power
- Printer
- Server
- Storage

Type

/CiscoUCS

Connection Information

Enter multiple similar devices, separated by a comma, using either hostname or IP Address:

Username: admin

Password: *****

Port: 443

Use SSL?:

Add

Devices

| Status | Host | Credentials | Type | Duration | Job Log | Remove | Retry |
|---|------|-------------|------|----------|---------|--------|-------|
| Add infrastructure using the above form | | | | | | | |

Previous


Finish

- You can add devices from different categories on this page. For each device, select the category, type and enter the connection information. If you are not ready to add devices, you can skip this page. When you are ready to continue, click **Done**. You will be taken to the Dashboard view of Resource Manager.

Note When you launch Resource Manager in the future, you will go directly to the login screen.

Interface and Navigation

After you install the system and navigate to the interface from your Web browser, the Dashboard appears. The Dashboard provides at-a-glance information about the status of your IT infrastructure. It is the primary window into devices and events that the system enables you to monitor.

Figure 5: Dashboard

The Dashboard can be customized to display a variety of information including:

- System information resources and other Web pages, such as internal portal pages
- Important error-level device events
- Geographical high-level view
- "Troubled" devices
- Devices in a certain production state You can also create additional dashboards that can be cloned from existing dashboards and can also be restricted to certain users or user groups.

Key Dashboard and interface areas include:

- Navigation menu
- User information area
- Portlets
- System Network Map

Navigation

The Navigation menu lets you access major system features. In addition to the Dashboard, the menu is divided among several functional areas:

- **EVENTS**- Guides you to the event management area, where you can monitor event status, triggers, and event transforms. You also can track changes made to events.
- **INFRASTRUCTURE**- Offers access to network infrastructure, including, devices, networks, processes, and services.
- **REPORTS**- Allows you access to pre-defined and configurable reports.
- **ADVANCED**- Provides access to monitoring templates, collectors, MIBs, system settings, and tuning.

User Information Area

Figure 6: User Information Area

The User information area offers information and selections:

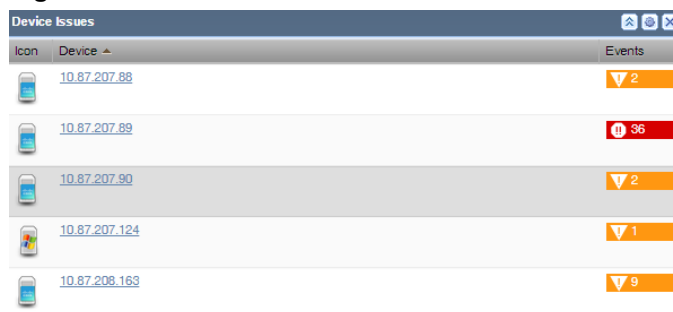
- **Search**- Search area to find information within the application. Click the down arrow in the search box to manage your saved searches.
- **Login ID**- The ID of the user currently logged in appears at the far left of this area. Click the ID to edit user settings, such as authentication information, roles, and groups. (You also can access user settings from the **ADVANCED > Settings > Users** page.)
- **Sign Out**- Click to log out of the system.
- **Help** icon - Click to access product documentation.

Portlets

The main content of the Dashboard contains portlets, which provide information about the system and your infrastructure. Portlets can display:

- **Device Issues**- Displays a list of devices, associated with color-coded events of critical, error, or warning severity levels. Click a device name to view details, or click an event to go to the event console for the device.

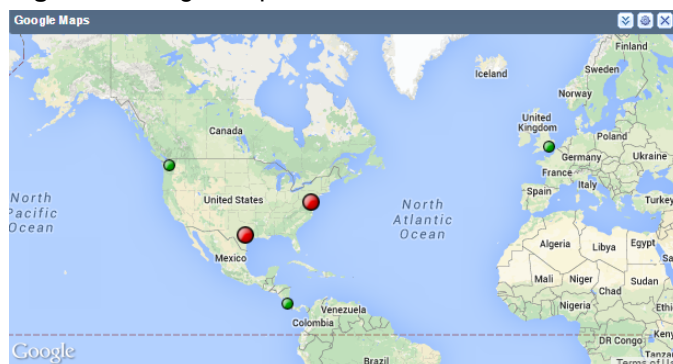
Figure 7: Device Issues Portlet



| Icon | Device | Events |
|------|-------------------------------|--------|
| | 10.87.207.88 | 2 |
| | 10.87.207.89 | 36 |
| | 10.87.207.90 | 2 |
| | 10.87.207.124 | 1 |
| | 10.87.208.163 | 9 |

- **Google Maps** (device locations) - Shows configured locations and configured network connections.

Figure 8: Google Maps Portlet



- **Daemon Processes Down**- Contains system self-monitoring information.
- **Production States**- Shows devices assigned to a particular production state. If needed, you can define multiple production states to display.
- **Site Window**- Initially provides links to resources such as product guides, forums, and training events.

You can customize the Site Window portlet to display content from any URL; however Zenoss recommends that you keep a portlet with the default URL so that you can stay up-to-date with Zenoss training and product updates. You can have multiple Site Window portlets defined on your Dashboard.

- **Device Chart**- Allows the display of a graph of multiple data points for a selected device class.
- **Event View**- Displays the list of events similar to the view on the Event console.


- **HTML Portlet**- Displays HTML content. You must use HTML markup in this portlet. If you want to populate a portlet with content from a given URL, use the Site Window portlet instead.
- **Network Map**- Displays a network map for a defined network being monitored. You can define the refresh interval and level of depth of your map.
- **Open Events Chart**- Displays a chart of the number of open events grouped by severity. You can define the event class to use as well as the number of past days for which to show events.
- **Watch List**- Allows the display of high-level status device classes, groups, systems, event classes, and locations that you select.

Network Map

The Network Map represents your network's Layer 3 topology. From the map, you can quickly determine the status of each device by its highlighted color.

To access the network map, select Infrastructure, and then select Network Map.

Figure 9: Network Map



Choosing the Network to Display

The network displayed is configured for each user. From user preferences, modify Network Map Start Object to select a network, and then click **Save**.

Viewing Device and Network Details

Double-click a device or network icon in the map to focus on it. Focusing on a node:

- Centers it on the map
- Shows links from the node, based on the number of hops selected

Alternatively, you can type the name or IP address of a device or network in the Selected Device or Network field, and then click **Refresh** to focus on that node.

To see detailed information about a device or network, select it in the map, and then click **Go to Status Page**.

Note When you select a node, the network map displays only the links that are currently loaded into the map. It does not download and display new link data.

Loading Link Data

To load link data for a node:

- 1 Double-click the node on the map to focus on it, or enter the device name or IP address in the Selected Device or Network field.
- 2 Select the number of hops to download and display by sliding the counter.
- 3 Click **Refresh**.

Filtering by Device Type

You can filter the devices that appear on the network map. To do this, select a filter from the Device Class Filter list of options. For example, to show only Linux devices on the map, select /Server/Linux from the list of options, then click **Refresh**.

Adjusting Viewable Hops

You can adjust the number of hops that appear on the network map. Use the Number of Hops slider, which adjusts the number of hops from 1 to 4.

Adjusting the Network Map

Use the Repulsion slider to expand or contract the icons that appear on the map. Move the slider right to expand the icons, or left to contract them.


Select the **Fit to Window** option to bring all displayed icons into the viewable area.

Customizing the Dashboard

You can customize the dashboard by:

- Creating multiple dashboards
- Selecting the portlets you want to view
- Arranging portlets
- Defining who can view the dashboard
- Changing the Dashboard column layout

The following screenshot of a sample dashboard shows the Add icon menu activated:



Adding a New Dashboard

A default admin dashboard is created when you launch the system. This dashboard can be customized by administrators, but you may want to create other dashboards that display distinctive information or are targeted to a specific type of user (including just yourself). You can create as many additional dashboards as you choose. You can customize them by selecting who can view the dashboard as well as selecting and customizing portlets to display the most important information. Users that are not administrators will initially see a read-only version of the default admin dashboard and can then create additional dashboards for their use.

Note You cannot delete the default admin dashboard.

To create an additional dashboard:

- 1 From the **Add** icon on the Dashboard controls, select **New Dashboard**. The Add a New Dashboard dialog appears.

- 2 Enter a Dashboard Name. When this dashboard name is displayed in the Dashboards drop-down list, the user name who created it will be appended in parentheses as part of the name. This gives everyone who can see the dashboard an indication of who created it.

- 3 Select who can view this dashboard. If you want a User Group to view this dashboard, the group must already be created in the system and the user creating this new dashboard must be a member of the group. You cannot add a dashboard and assign it to a group you are not part of.
- 4 Select the number of columns to display in the dashboard. The default is 3.
- 5 If you want to clone the new dashboard from the previously viewed dashboard, select the check box. Otherwise, you will begin with a completely blank dashboard.
- 6 Click **Create**.

Adding Portlets

You can customize your dashboard by adding portlets that display information you are interested in. Your dashboard can display more than one of the same portlet type. For example, you could have several Device Chart portlets with each one showing a different device class.

To add a portlet to the Dashboard:

- 1 Click the **Add** icon in the top right of the Dashboard main area and select **Add portlet**.
- 2 Select a portlet from the drop-down list.

The portlet appears at the top right of the Dashboard main area.

Arranging Portlets

To arrange portlets, click the portlet header and drag the portlet to any location on the Dashboard. Other portlets rearrange depending on the location you drop it.

Editing the Dashboard Settings

You can customize your dashboard to display a different number of columns or limit access to the dashboard.

To edit the dashboard settings:


- 1 Click the Action icon in the upper-right side of the Dashboard. The Edit Dashboard window appears.

- 2 Change the value of the users who can view this dashboard, if needed. If you want a User Group to be able to view this dashboard, the group must already be defined in the system and the user editing the dashboard must be a member of the user group in order to see that group in the drop-down list.
- 3 Change the number of columns to use in this dashboard if needed.
- 4 Select the **Lock from updates?** check box to prevent editing of the dashboard.
- 5 Click **Save**.

Working with Portlets

There are several options to control the portlet display.

Figure 10: Dashboard Portlet Controls



- Click the **Collapse** icon to collapse the portlet so that only the title is displayed on the dashboard.
- Click the **Pop Out** icon to show the portlet in a full screen view. Click **Close** to return to the dashboard view.
- Click the **Edit** icon to edit the portlet settings. You can edit the title of the portlet, its height and how often it refreshes. Some portlets may have additional configuration options. A preview of the portlet is provided on the right side of the dialog box. Click **Save** to update the portlet configuration.
- Click the **Close** icon to remove the portlet from the dashboard.

In tabular portlets, you can control the display by sorting columns as well as adding and hiding columns.


- To sort based on a column, hover over the column header and click the arrow to display the sort and display options.
- To add or hide columns, hover over the **Columns** entry and check or clear the boxes of the columns to add or hide.

| Status | Severity | Resource | Component | Event Class | Summary | First Seen | Last Seen | Count |
|--------|----------|------------------|---------------|-------------|-------------------|-------------------------|-------------------------|-------|
| ... | ... | | | | | | | |
| | Warning | qa-centos-7-s... | | | SNMP agent dow... | 2016-09-27 05:08:22 ... | 2016-09-27 10:48:50 ... | 69 |
| | Warning | qa-centos-5-s... | | | SNMP agent dow... | 2016-09-27 05:08:59 ... | 2016-09-27 10:49:44 ... | 69 |
| | Warning | Cisco 10.171... | | | | 2016-09-27 05:09:18 ... | 2016-09-27 10:47:02 ... | 69 |
| | Warning | Cisco 10.171... | si0 | /Perf/In | | 2016-09-27 05:09:19 ... | 2016-09-27 10:47:02 ... | 69 |
| | Error | nexus-5k.zeno... | VSAN0080... | /Status | Status | 2016-09-27 05:09:49 ... | 2016-09-27 10:48:56 ... | 69 |
| | Error | nexus-5k.zeno... | VSAN0001... | /Status | Severity | 2016-09-27 05:09:50 ... | 2016-09-27 10:48:57 ... | 69 |
| | Error | N6001-2 | Fex-108 A/... | /Status | Resource | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-108 P... | /Status | Component | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-109 A/... | /Status | Event Class | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-115 A/... | /Status | Summary | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-131 P... | /Status | First Seen | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-114 P... | /Status | Last Seen | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | Error | N6001-2 | Fex-110 A/... | /Status | Count | 2016-09-27 05:12:24 ... | 2016-09-27 10:47:36 ... | 68 |
| | | | | | Production State | | | |

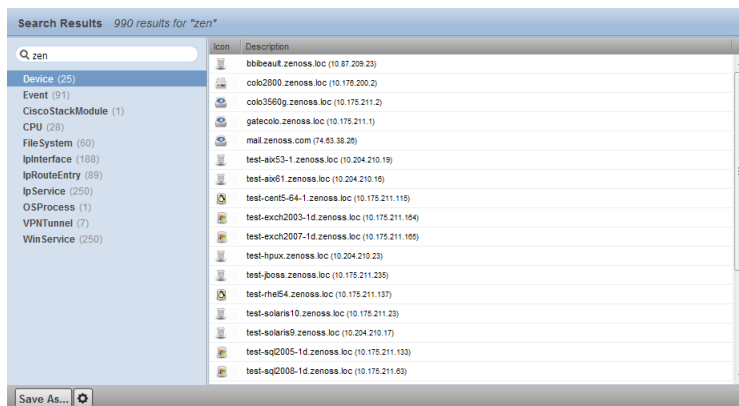
Search

The Resource Manager search facility supports locating devices and other system objects, as well as events and services.

In the Resource Manager interface, the search feature is part of the user information area. Enter part or all of a name in the search box at the top right of the interface. The system displays matches, categorized by type.

Figure 11: Search Results

To view all search results, click the indicator at the top of the list.

Figure 12: All Search Results

From here, you can display search results by category. Click in the left panel to filter search results by a selection.

You can save the search to access later.

- 1 Click **Save As** (at the bottom left of the Search Results page). The Save Search As dialog appears.
- 2 Enter a name for the search, and then click **Submit**.


You can access saved searches from:

- Action menu located at the bottom of the Search Results page.
- Search box located at the top of the interface. Click the arrow, and then select **Manage Saved Searches**.

Navigating the Event Console

The event console is the system's central nervous system, enabling you to view and manage events. It displays the repository of all events that are detected by the system.

To access the event console, click **Events** in the Navigation menu. The Event Console appears.

Figure 13: Event Console

Sorting and Filtering Events

You can sort and filter events that appear in the event console to customize your view.

You can sort events by any column that appears in the event console. To sort events, click a column header. Clicking the header toggles between ascending and descending sort order.

Filter options appear below each column header. A match value can be any full string or a subset of a string, optionally with the wildcard (*) contained in the values in that column. You can also use " | | " (OR), or " ! ! " (NOT) expressions to further target your filters. For example, typing ! windows in the Event Class filter will return all the non-Windows device events.

Figure 14: Event Console Filter Options

| Status | Severity | Resource | Component | Event Class | Summary | First Seen | Last Seen | Count |
|-------------------------------------|--------------|----------|------------------------------|--|---------|------------------------|------------------------|-------|
| <input checked="" type="checkbox"/> | New | ics1 | Fan_3142 | Fan 2 In Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |
| <input checked="" type="checkbox"/> | Acknowledged | ics1 | extpol_reg_clients_client... | UCS Domain ucs1-faba is registered with UCS Central without a valid license. | | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| <input type="checkbox"/> | Suppressed | ics1 | Fan_Module_314 | Fan 2 In Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | | 2015-07-29 08:45:41 am | 2015-07-29 08:45:41 am | 1 |
| <input type="checkbox"/> | Closed | ics1 | fabric_san_A_phys-foes... | FCoE uplink port 116 is down | | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| <input type="checkbox"/> | Cleared | ics1 | Fan_3142 | Fan 2 In Fan Module 1-4 under chassis 3 operability: inoperable | | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |
| <input type="checkbox"/> | Aged | ics1 | Fabric_FC_SAN_PortChan... | san port-channel 5 on fabric interconnect A oper state: failed, reason: No operatio... | | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| | | ics1 | Fabric_FC_SAN_PortChan... | san Member 127 of Port-Channel 5 on fabric interconnect A is down, membership... | | 2015-07-29 08:46:12 am | 2015-07-29 10:29:40 am | 4 |

You can filter the events that appear in the list in several ways, depending on the field type:

- **Resource** - Enter a match value to limit the list.
- **Component** - Enter a match value to limit the list.
- **Event Class** - Enter a match value to limit the list.
- **Summary** - Enter a match value to limit the list.
- **First Seen** - Enter a value or use a date selection tool to limit the list.
- **Last Seen** - Enter a value or use a date selection tool to limit the list.
- **Count** - Enter a value to filter the list, as follows:
 - N - Displays events with a count equal to N .
 - $:N$ - Displays events with a count less than or equal to N .
 - $M:N$ - Displays events with a count between M and N (inclusive).
 - $M:-$ - Displays events with a count greater than or equal to M .

To clear filters, select **Configure > Clear filters**.

You also can re-arrange the display order of columns in the event console. Click-and-drag column headers to change their display.

Creating an Actionable View

For users that are not Administrators, there is an option that will filter the list of events to show only those that are not read-only for the user's permission level, and enable the action buttons above the event table header.

To turn on the actionable view, click **Configure** and select the **Only show actionable events** check box. The view is changed to show only events that can have an action performed on them based on the user's permission level. For more information, see [Managing Events](#) on page 21.

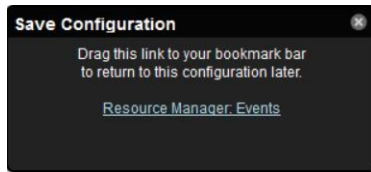
Saving a Custom View

This procedure describes how to save a custom event console view by bookmarking it for quick access later.

- 1 Select **Configure > Save this configuration**.

- The interface displays a dialog containing a link to the current view.
- 2 Select the link, and then drag it to the bookmarks area of your browser window.

Figure 15: Saving a Custom View (Bookmark)




The browser adds a link titled "Resource Manager: Events" to the bookmarks list. You may want to re-title the bookmark, particularly if you choose to save more than one event console view.

Refreshing the View

You can refresh the list of events manually or specify that they refresh automatically. To manually refresh the view, click **Refresh**. You can manually refresh at any time, even if you have an automatic refresh increment specified.

To configure automatic refresh, select one of the time increments from the Refresh list. By default, automatic refresh is enabled and set to refresh each minute.

Figure 16: Automatic Refresh Selections



Viewing Event Details

You can view details for any event in the system. To view details, double-click an event row.

Note Do not double-click on or near the device (resource) name, component, or event class in the row. Doing this displays details about that entity, rather than information about the event.

The Event Detail area appears.

Figure 17: Event Detail



To see more information about the event, click **Event Details**.

You can use the Log field (located at the bottom of the area) to add specific information about the event. Enter details, and then click **Add**.

Selecting Events

To select one or more events in the list, you can:


- Click a row to select a single event.
- Ctrl-click rows to select multiple events, or Shift-click to select a range of events.
- Click **Select > All** to select all events.

Managing Events

You can manage events from the event console. After making a selection by clicking on the row of the event, you can:

- Acknowledge the event
- Close the event
- Reclassify the event, associating it with a specific event class
- Return the event to New status (revoke its Acknowledged status)
- Reopen the event
- Add a note to the log

You also can add an event from the event console. This feature is useful for testing a specific condition by simulating an event.

Figure 18: Event Management Options

Running a command from the browser interface

Resource Manager allows commands to be run through the browser interface. You can run commands on a single device or on a group of devices.

The system includes several built-in commands, such as `ping` and `traceroute`.

To run commands from the browser interface:

- 1 Navigate to the **INFRASTRUCTURE** tab.
- 2 In the **Devices** list, select one or more devices.
To select a device, click anywhere in the row—except on the link.
- 3 Click **Commands** and select a command from the list.

Figure 19: Command Output

You can resize the command output window. You also can stop automatic scrolling by de-selecting the Autoscroll option at the bottom right corner of the output window.

Running a Command from the CLI

There are many commands in Resource Manager that need to be run from the command-line interface (CLI). Since Resource Manager uses containers to deploy the application, you must first attach to the container through the Control Center CLI before using Resource Manager commands that you see in this guide.

To use the Control Center CLI, you need a Linux shell account on hosts in the Control Center pool and the account needs to be a member of the `docker` group. Control Center uses the `serviced` application service orchestrator to manage the container-based system. This technology is based on Docker.

The general syntax of a `serviced` command is:

```
serviced [global options] command [command options] [arguments...]
```

For more details on the `serviced` command and its syntax, see the Control Center documentation by clicking the Help icon in that interface.

Some sample commands:

- `serviced service list` - lists all the running services (and the value of its `SERVICEID`) on the host (to avoid an error when invoking the same command on a host that is different than the master host, you can add the `--endpoint` flag).
- `serviced service stop zenhub` - stops the zenhub service
- `serviced service attach $SERVICEID bash` - attaches to a service identified by its `SERVICEID`.
- `serviced service attach zope/0 tail -F /opt/zenoss/log/event.log - tail` (and follow) the event log of the first running Zope service.
- `serviced service status zope` - show status of Zopes.

The following shows the commands to execute to make a change to the Zope container, commit the change, and restart all the Zenoss services:

```
serviced service shell -i -s mychange zope bash
[Make your changes, then exit]
serviced snapshot commit mychange
serviced service start Zenoss.resmgr
```

Working with Triggers and Notifications

You can create *notifications* to send email or pages, create SNMP traps, or execute arbitrary commands in response to an event. Notifications also can be used to notify other management systems, and to execute arbitrary commands to drive other types of integration. How and when a notification is sent is determined by a *trigger*, which specifies a rule comprising a series of one or more conditions.

To set up a notification, you must:

- Create a trigger, selecting the rules that define it
- Create a notification, selecting one or more triggers that cause it to run
- Choose appropriate options and subscribers, depending on the notification type

Read the following sections to learn about:

- Setting up triggers and trigger permissions
- Setting system SMTP settings for notifications
- Setting up notifications and notification permissions

Working with Triggers

Setting up a trigger involves:

- Creating the trigger and the rules that define it
- Setting trigger permissions

Creating a Trigger

This procedure describes how to create a trigger.

- 1 Select **EVENTS > Triggers**.

The Triggers page displays all defined triggers.

- 2 Click the **Add** icon.

- 3 In the **Add Trigger** dialog, enter a name for the trigger, and then click **Submit**.

Only uppercase letters, lowercase letters, digits, and underscores are allowed in trigger names.

The trigger is enabled and added to the list.

To complete the trigger, proceed to the next task.

Editing a Trigger

This procedure describes how to edit a trigger.

- 1 Select **EVENTS > Triggers**.

The Triggers page displays all defined triggers.

- 2 Open the **Edit Trigger** dialog of the trigger to edit.

In the **Triggers** table, double-click a trigger, or select a trigger and then click the **Action** icon.

Figure 20: Edit Trigger dialog

- 3 Define rules for the trigger, and then click **SUBMIT**.

Trigger rules combine Boolean logic with event values to decide whether to send notifications. A trigger can use any or all of its rules to make a decision, and each rule can have subordinate rules or branches.

Note Device production states can change during maintenance windows. If you want the same trigger to apply during maintenance windows, be sure to edit your trigger to account for all production states that apply to your trigger.

Setting Global Trigger Permissions

You can set global permissions for viewing, editing, and managing triggers. Global permissions are given to any user with "manage" permission, which includes:


- Admin, Manager, and ZenManager roles
- Trigger owner

Edit global permissions from the Users tab on the Edit Trigger dialog.

Global options are:

- **Everyone can view** - Provides global view permission.
- **Everyone can edit content** - Provides global update permission.
- **Everyone can manage users** - Provides global manage permission.

Figure 21: Edit Trigger - Users Tab



Setting Individual Trigger Permissions

You can grant permissions to individual users. For each user added, you can select:

- **Write** - Select this option to grant the user permission to update the trigger
- **Manage** - Select this option to grant the user permission to manage the trigger.

To set an individual's trigger permissions:

- 1 Select a user from the drop-down list in the Users section of the Edit Trigger dialog.
- 2 Click **Add**. The user is added.
- 3 Assign permissions by selecting the appropriate check box(es).
- 4 Optionally, add additional user trigger permissions by repeating this procedure.
- 5 When you are finished, click **Submit**.

To remove an individual's trigger permissions:

- 1 Select the row of the user's permissions.
- 2 Click the **Remove** icon.
- 3 Optionally, remove other user trigger permissions by repeating this procedure.
- 4 When you are finished, click **Submit**.

Working with Notifications

Setting up a notification involves:

- Creating the notification

- Defining notification content (for email- or page-type notifications)
- Defining the SNMP trap host (for SNMP trap-type notifications)
- Defining commands to run (for command-type notifications)
- Setting notification permissions
- Setting up notification schedules

Creating a notification

This procedure describes how to edit a notification.

- 1 Select **EVENTS > Triggers**.
- 2 In the left panel, select **Notifications**.

Figure 22: Notifications

| Event Console | | Event Archive | | Event Classes | | Triggers | |
|---------------|------------|---------------|---------|---------------|---------|------------------------|-----------|
| Triggers | | Notifications | | | | Notification Schedules | |
| Notifications | | | | | | | |
| Enabled | ID | Trigger | Action | Subscribers | Enabled | ID | Start |
| Yes | rancid-run | rancid-run | command | 0 | No | test | 1/14/2015 |

The **Notifications** area includes a table of notifications and a table of notification schedules.

- 3 Click the **Add** icon.
- 4 In the **Add Notification** dialog, provide a name and specify an action.
 - a Enter the **Id** of the notification.
Spaces are not allowed in notification Ids.
 - b Associate an action with the notification.
For more information, see [Notification actions](#) on page 26.
 - c Click **SUBMIT**.

To complete the notification, proceed to the next task.

Notification actions

| Action | Description |
|------------|---|
| Command | Invoke a shell command when events occur. Common uses of this action include: <ul style="list-style-type: none"> ▪ <i>Auto-remediation of events.</i> You can use <code>ssh</code> or <code>winexe</code> to restart services on Linux and Windows devices. ▪ <i>Integration with external systems.</i> For example, opening tickets in an incident management system. ▪ <i>Extending alerting mechanisms.</i> Resource Manager supports email and pagers as alerting mechanisms "out of the box" through normal alerting rules. |
| Email | Sends an HTML or text email message to authorized subscribers when an event matches a trigger rule. |
| Page | Pages authorized subscribers when an event matches a trigger rule. |
| Syslog | Sends a message to the syslog. |
| SNMP Trap | Sends an SNMP trap when an event matches a trigger rule. |
| WinCommand | Sends or clears a Windows CMD command. |

Editing a notification

This procedure describes how to edit a notification.

- 1 Select **EVENTS > Triggers**.
- 2 In the left panel, select **Notifications**.

Figure 23: Notifications

| Notifications | | | | | Notification Schedules | | |
|---------------|------------|------------|---------|-------------|------------------------|------|-----------|
| Enabled | ID | Trigger | Action | Subscribers | Enabled | ID | Start |
| Yes | rancid-run | rancid-run | command | 0 | No | test | 1/14/2015 |

The **Notifications** area includes a table of notifications and a table of notification schedules.

- 3 Open the **Edit Notification** dialog of the notification to edit.
In the **Notifications** table, double-click a notification, or select a notification and then click the **Action** icon.

Figure 24: Edit Notification

- 4 Define the settings for the notification.
For more information, see [Notification settings](#) on page 27.

Notification settings

| Setting | Description |
|---------------------------------|--|
| Enabled | Check the checkbox to enable the notification. |
| Send Clear | Send a notification when the problem is resolved by a clearing event. |
| Send only on Initial Occurrence | Send a notification only when the first triggering event occurs. |
| Delay (seconds) | <p>The minimum number of seconds to wait before performing a notification. A delay prevents notifications of transient problems and multiple notifications of the same problem.</p> <p>For example, if five events that match the trigger occur in 45 seconds, a delay of 60 seconds ensures that only one notification is sent. Also, if a triggering event repeats 15 seconds after the initial event, followed by a clearing event at 45 seconds, a 60-second delay ensures that no notifications are sent.</p> |

| Setting | Description |
|------------------|--|
| Repeat (seconds) | The interval between repetitions of the notification, in seconds. The notifications repeat until the triggering event is resolved. |

Defining Notification Content

To define notification content, click the **Content** tab of the notification.

For email-type notifications, you can use the default configuration for the following fields, or customize them to your needs:

- **Body Content Type** - Select HTML or text.
- **Message (Subject) Format** - Sent as the subject of the notification.
- **Body Format** - Sent in the notification.
- **Clear Message (Subject) Format** - Sent when a notification clears.
- **Body Format** - Sent when a notification clears.
- **From Address for Emails** - Sent as email address of sender
- **Various SMTP settings** - Used to define SMTP host, port, username, and password. To set these system-wide, go to the **ADVANCED > Settings** page.

Figure 25: Define Notification Content (Email)

Edit Notification - ExampleServiceEmailNotification (email)

Notification **Content** Subscribers

Body Content Type:

Message (Subject) Format:

Body Format:

```

Device: ${evt/device}
Component: ${evt/component}
Severity: ${evt/severity}
Time: ${evt/lastSeen}
Message:
${evt/message}
<a tal:attributes="href urls/eventUrl">Event Detail</a>
<a tal:attributes="href urls/ackUrl">Acknowledge</a>
<a tal:attributes="href urls/closeUrl">Close</a>
<a tal:attributes="href urls/eventsUrl">Device Events</a>
The following events are causes of this event:

<tal:block tal:repeat="item esa/causes">
Impact Chain: ${item/impactChain}
Device: ${item/evt/device}
Component: ${item/evt/component}
Severity: ${item/evt/severity}
Time: ${item/evt/lastSeen}
Message:
${item/evt/message}
<a tal:attributes="href item/urls/eventUrl">Event Detail</a>
<a tal:attributes="href item/urls/ackUrl">Acknowledge</a>
<a tal:attributes="href item/urls/closeUrl">Close</a>
<a tal:attributes="href item/urls/eventsUrl">Device
Events</a>
</tal:block>

```

Clear Message (Subject) Format:

Body Format:

```

Event: '${evt/summary}'
Cleared by: '${evt/id}'
At: ${evt/stateChange}
Device: ${evt/device}
Component: ${evt/component}
Severity: ${evt/severity}
Message:
${evt/message}
<a href="${urls/reopenUrl}">Reopen</a>

```


For page-type notifications, you can use the default configuration for the following fields, or customize them to your needs:

- **Message (Subject) Format** - Sent as the subject of the notification.
- **Clear Message (Subject) Format** - Sent when a notification clears.

Figure 26: Edit Notification Content (Page)

Notification Content Variables

Within the body of your email, page, and command notifications, you can specify information about the current event, in the form:

```
'${objectname/objectattribute}'
```

Note Do not escape event command messages and event summaries. For example, write this command as: `${evt/summary}` (rather than `echo '${evt/summary}'`).

Object names may be `evt`, `evtSummary`, or `urls`; or for clearing event context, `clearEvt` and `clearEventSummary`. For each object name, the following lists show valid attributes (for example, `'${evt/DevicePriority}'`):

Table 1: evt/ and clearEvt/

| Value | Description |
|----------------|--|
| DevicePriority | value of the priority of the device |
| agent | Typically the name of the daemon that generated the event. For example, an SNMP threshold event has <i>zenperfsnmp</i> as its agent. |
| clearid | id of the event this clear event will clear |
| component | component this event is related to |
| count | how many times this event occurred |
| created | when the event was created |

| Value | Description |
|----------------|--|
| dedupid | dynamically generated fingerprint that allows the system to perform de-duplication on repeating events that share similar characteristics |
| device | device this event is related to |
| eventClass | class of this event |
| eventClassKey | Free-form text field that is used as the first step in mapping an unknown event into an event class. |
| eventGroup | Free-form text field that can be used to group similar types of events. This is primarily an extension point for customization. Currently not used in a standard system. |
| eventKey | Free-form text field that allows another specificity key to be used to drive the de-duplication and auto-clearing correlation process. |
| eventState | state of the event |
| evvid | unique id for the event |
| facility | the syslog facility |
| firstTime | UTC Time. First time that the event occurred. |
| ipAddress | IP address |
| lastTime | UTC time. Most recent time that the event occurred. |
| manager | value of manager |
| message | a message communicated by the event |
| nvid | windows event id |
| ownerid | ownerid |
| priority | syslog priority |
| prodState | production state of the device |
| severity | the severity of the event expressed as a number (0, 1, 2, 3, 4, or 5) |
| severityString | the severity of the event expressed as a string (Clear, Debug, Info, Warning, Error, or Critical) |
| stateChange | last time that the event status changed |
| status | the status of the event |
| summary | a short message summarizing the event |

Note Some of the values in the following table are direct duplicates of fields on `evt`. For example, `uuid` -> `evt.evid`.

Table 2: eventSummary/ and clearEventSummary/

| Value | Description |
|------------|-------------|
| uuid | evt.evid |
| occurrence | evt.count |

| Value | Description |
|--------------------------|---|
| status | evt.eventState |
| first_seen_time | evt.firstTime |
| status_change_time | evt.stateChange |
| last_seen_time | evt.lastTime |
| count | evt.count |
| current_user_uuid | UUID of the user who acknowledged this event |
| current_user_name | name of the user who acknowledged this event. |
| cleared_by_event_uuid | the UUID of the event that cleared this event (for events with status == CLEARED). |
| notes | event notes |
| audit_log | event audit log |
| update_time | last time a modification was made to the event |
| created_time | evt.lastTime |
| fingerprint | evt.dedupid |
| event_class | evt.eventClass |
| event_class_key | evt.eventClassKey |
| event_class_mapping_uuid | If this event was matched by one of the configured event class mappings, it contains the UUID of that mapping rule. |
| actor | event actor |
| summary | evt.summary |
| message | evt.message |
| severity | evt.severity |
| event_key | evt.eventKey |
| event_group | evt.eventGroup |
| agent | evt.agent |
| syslog_priority | evt.priority |
| syslog_facility | evt.facility |
| nt_event_code | evt.ntevid |
| monitor | evt.monitor |
| tags | event tags |

Table 3: urls/

| Value | Description |
|----------|---------------------------------|
| ackUrl | URL for acknowledging the event |
| closeUrl | URL for closing the event |

| Value | Description |
|-----------|---|
| reopenUrl | URL for reopening the event |
| eventUrl | URL for viewing the event |
| eventsUrl | URL for viewing events for the relevant device, or all events |

ZenPacks also can define additional notification actions, and can extend the context available to notifications to add additional objects or attributes.

Defining the SNMP Trap Host

For SNMP trap-type notifications, enter information or make selections on the Content tab of the notification:

- **SNMP Trap Destination-** Specify the host name or IP address where the trap should be sent.
- **SNMP Community-** Specify the SNMP community. By default, this is public.
- **SNMP Version-** Select v2c (default) or v3.
- **SNMP Port-** Specify the SNMP port. Typically, this is 162.

SNMP traps sent as a result of this notification are defined in the ZENOSS-MIB file. You can find this MIB file on any Resource Manager server at `$ZENHOME/share/mibs/site/ZENOSS-MIB.txt`.

Note In order to execute a command using `$ZENHOME (/opt/zenoss` for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for `serviced` commands.

Figure 27: Edit Notification Content (SNMP Trap)

The screenshot shows a window titled "Edit Notification - test_trap (trap)" with three tabs: "Notification", "Content", and "Subscribers". The "Content" tab is active. It contains four configuration fields:

- SNMP Trap Destination:** A text input field containing "traphost".
- SNMP Community:** A text input field containing "public".
- SNMP Version:** A dropdown menu with "v2c" selected.
- SNMP Port (usually 162):** A spinner control with "162" displayed.

At the bottom of the dialog are two buttons: "SUBMIT" and "CANCEL".

Defining Commands to Run

For Command-type notifications, you must specify the command to run when configured triggers are matched. Do this on the Content tab of the notification. Configure these fields:

- **Command Timeout** - By default, 60 seconds.
- **Command** - Command to run when a trigger is matched.
- **Clear Command** - Optional command to run when the triggering event clears.
- **Environment variables** -

Figure 28: Edit Notification Content (Command)

Global Notification Permissions

By establishing permissions, you can control which users have the ability to view, manage, and update notifications. Permissions are granted based on the user's assigned role. The following table lists account roles and their associated notification permissions:

| Role | Permissions |
|--|--|
| Admin, Manager, ZenManager | Users assigned the Admin, Manager, or ZenManager roles can view, update, and manage any notification. |
| Notification owner | When a user creates a notification, he is designated the owner of that notification. During the life of the notification, the owner can view, update, and manage it. |
| All other users (including those assigned ZenUser and ZenOperator roles) | Must be specifically granted permissions through the interface to view, edit, or manage notifications. |

You can set global permissions for viewing, updating and managing a notification. Global permissions are given to any user with "manage" permission, which includes:

- Admin, Manager, and ZenManager roles
- Notification owner

Edit global permissions from the Subscribers tab on the Edit Notification Subscription panel.

Global options are:

- **Everyone can view** - Provides global view permission.
- **Everyone can edit content** - Provides global update permission.
- **Everyone can manage subscriptions** - Provides global manage permission.

Permission checks occur when the data is sent to the browser and when any action occurs. To determine where a user can make modifications to a particular tab, permission checks are performed on global roles, then managerial roles, and then individual roles. Any role that provides the required permission will allow that permission's associated behavior.

Figure 29: Edit Notification

| Type | Subscribers | Write | Manage |
|------|---------------|--------------------------|--------------------------|
| user | adm in (User) | <input type="checkbox"/> | <input type="checkbox"/> |

Setting Individual Notification Permissions

You can grant permissions to individual users or groups. For each user or group added, you can select:

- **Write** - Select this option to grant the user or group permission to update the notification.
- **Manage** - Select this option to grant the user or group permission to manage the notification.

You can manually enter in the name of a user or group, or select one from the list of options.

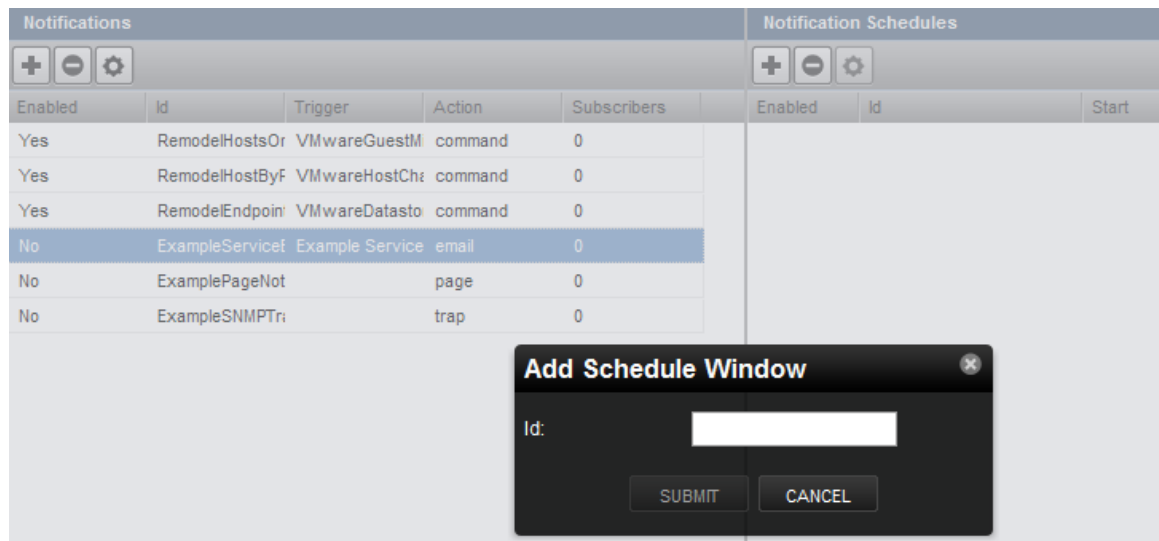
Setting Up Notification Schedules

You can establish one or more notification schedules for each defined notification.

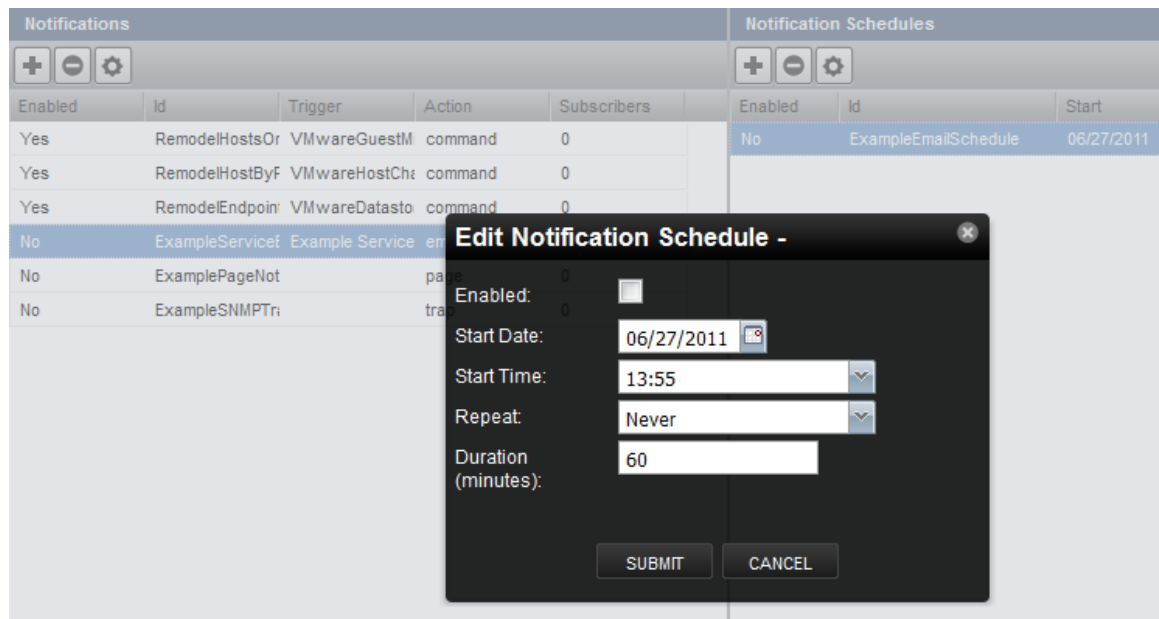
Note If a notification has one or more notification schedules, then notifications are sent only if a schedule is active when a notification arrives.

To set up a schedule:

- 1 In the **Notifications** area, select a notification.
- 2 In the **Notification Schedules** area, click the **Add** button.

Figure 30: Add Schedule Window dialog

- 3 In the **ID** field, enter an identifier for the schedule, and then click **Submit**.
- 4 In the **Notification Schedules** area, double-click the newly-added schedule.

Figure 31: Edit Notification Schedule

The fields of the **Edit Notification Schedule** dialog are described in the next topic.

- 5 To save the new schedule, click **Submit**.

Notification schedule fields

| Field | Description |
|-------------------|--|
| Enabled | Check the box to enable the schedule. |
| Start Date | The calendar date of the start of the schedule. |
| Start Time | The time of day of the start of the schedule. |
| Repeat | The schedule frequency, one of the following values: |

| Field | Description |
|---------------------------|---|
| | <ul style="list-style-type: none"> ▪ Never ▪ Daily ▪ Every Weekday ▪ Weekly ▪ Monthly ▪ First Sunday of the Month |
| Duration (Minutes) | The period of time during which the notification window is active. |

Advanced User Interface Configuration

To access advanced user interface configuration options, select **Advanced > Settings**, and then select **User Interface** in the left panel. This area lets you configure options such as how data loads, how much data is loaded, and filter and search options.

Select options or enter information:

- **Enable Infinite Grids for Events**- Disable this option to turn off infinite grids. By default, this feature is enabled.
- **Enable Live Filters**- Disable this option to turn off the live filters feature. If you disable this feature, you will need to press enter on all search and filter inputs.
- **Enable Incremental Tree Loading on the Infrastructure Page**- Enable this option to load the infrastructure tree one node at a time. If disabled (the default), then the infrastructure tree is loaded all at once. You might enable this option if you have a complex hierarchy of organizers and device classes and want to improve your UI load response time.
- **Show Tree Event Severity Icons**- Disabling this option may speed up page loading.
- **Enable Tree Filters**- If disabled, then tree filters (the text input area that allows you to filter the information displayed) are hidden on all pages. This option is enabled by default.
- **Show Page Statistics**- Enable this option to show debugging information. By default, this option is not enabled.
- **Device Grid Buffer Size**- Specify the number of device data rows to fetch from the server for each buffer request. The default buffer size is 100 rows.
- **Component Grid Buffer Size**- Specify the number of component data rows to fetch from the server for each buffer request. The default buffer size is 50 rows.
- **Event Console Buffer Size**- Specify the number of event rows to fetch from the server for each buffer request. The default buffer size is 200 rows.
- **Device Move Job Threshold**- Specify the limit at which devices are moved immediately. If the number of devices to be moved exceeds this threshold, then the move occurs in a job; otherwise, they are moved immediately. The default value is 5 devices.
- **Job Notification Refresh Interval**- Specify the refresh interval for the job notification dialog. The default time is 10 seconds.
- **Job Grid Buffer Size**- Specify the number of job data rows to fetch from the server for each buffer request. The default buffer size is 100 rows.

When complete, click **Save** to save your changes.

LDAP Authentication

You can use your existing LDAP authentication infrastructure, such as Active Directory or OpenLDAP, to enable single sign-on to the Resource Manager interface. With this capability, you can use the user management tools with which you are familiar to enable your Windows users to use their Windows credentials to authenticate to the

Resource Manager interface. This saves you from having to manually create user accounts and separately maintain passwords.

For those LDAP properties that are mapped, changes you make in LDAP are updated in Resource Manager. (You must clear your browser cache, or log out and back in, for properties changes to propagate immediately.)

LDAP Configuration Information

Before configuring LDAP authentication, you should gather the following information from your LDAP or Active Directory administrator:

- Host name or IP address of an Active Directory global catalog server (for Active Directory authentication)
- Host name or IP address of an LDAP server (for other LDAP server authentication)
- User's base distinguished name (DN)
- Manager DN
- Manager password
- Groups base DN
- Optionally, list of Active Directory groups to map to Resource Manager roles

Configuring LDAP Authentication

You can configure LDAP authentication at initial setup, or from the Settings area of the interface:

- While in the setup wizard, at Step 2: Specify or Discover Devices to Monitor, click **LDAP Setup** (located at the bottom right of the wizard panel).
- From the interface, select **ADVANCED > LDAP** and click the **Add** icon.

The first panel (Add LDAP Servers) of the LDAP Configuration wizard appears.

Figure 32: LDAP Configuration Wizard (Add LDAP Servers)

New LDAP Configuration

1. Add LDAP Servers

Type: Active Directory

Host: Port: SSL? Skip cert verification?

Manager Credentials

Server Type: Active Directory Other LDAP

Manager DN:
Example: cn=admin,cn=users,dc=example,dc=com

Manager Password:

- 1 Enter information and make selections:
 - **Host**- Enter the host name or IP address of an Active Directory global catalog server (for Active Directory authentication) or the host name or IP address of an LDAP server (for Other LDAP server types).
 - **Port**- Optionally, change the server port number. By default, the port number is 389.
 - **SSL**- Select if using SSL. When you select this option, the default port number adjusts to 636.
 - **Skip cert verification?**- If you are using a self-signed certificate, select this check box to skip its verification. Requires OpenLDAP 2.4 or higher.
- 2 Optionally, click **Add Server** to add another LDAP server. To remove a server from the list, click **Remove**.
- 3 Enter information and make selections in the Manager Credentials area:
 - **Server Type**- Select a server type (Active Directory or Other LDAP).
 - **Manager DN**- Enter the distinguished name of a user in the domain administrators group. An example that follows the user's base DN is:


```
cn=admin, cn=users, dc=example, dc=com
```
 - **Manager Password**- Enter the password for the Manager DN.
- 4 Optionally, click **Validate** to ensure your setup is valid.
- 5 Click **Next**. The second panel (Configure LDAP Plugin) of the LDAP Configuration wizard appears.

Figure 33: LDAP Configuration Wizard (Configure LDAP Plugin)

New LDAP Configuration

2. Configure LDAP Plugin

Server Type: Active Directory

LDAP Configuration ID: win2008-ad.example.com

Login Name Attribute: sAMAccountName

Users Base DN: dc=Users, dc=example, dc=com
Example: dc=Users, dc=example, dc=com

Groups Base DN: dc=Groups, dc=example, dc=com
Example: dc=Groups, dc=example, dc=com

User Filter: (cn=Organization.*)
Example: (cn=Organization.*)

Group Filter: (cn=IT Admins)
Example: (cn=IT Admins)

Default User Roles: ZenUser

PREVIOUS NEXT CANCEL

- 6 Enter information and make selections:
 - **Login Name Attribute**- Select the LDAP record attribute used as the user name.

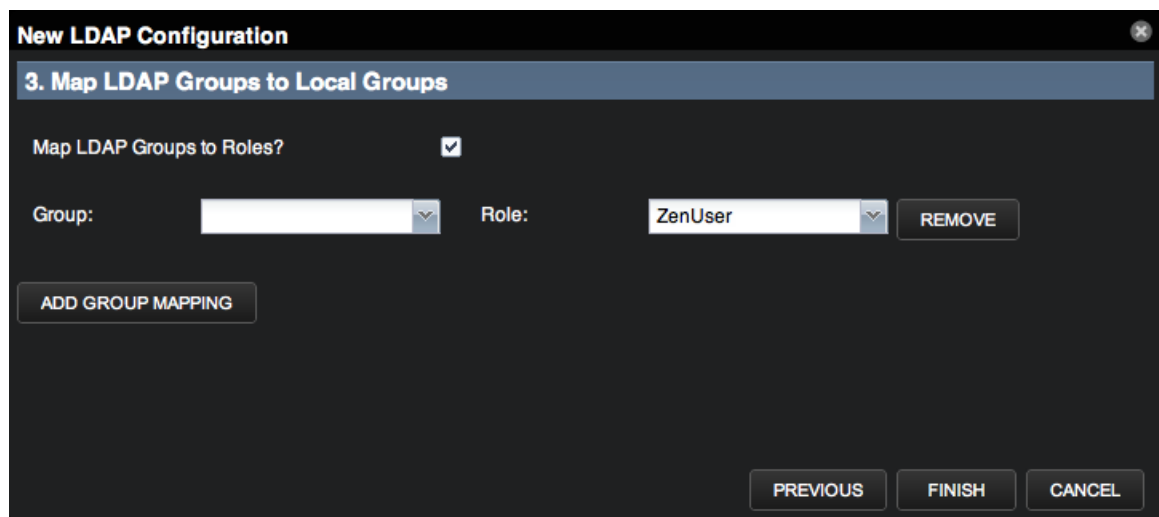
Note You can edit the list of selections by adding attributes on the Mappings page of the LDAP configuration area (**ADVANCED > LDAP**).

- **Users Base DN**- Enter the user's base distinguished name. For example, if your domain is ad.example.com, then your user's base DN might be:

```
dc=Users,dc=example,dc=com
```

- **Groups Base DN**- Enter the DN for the branch of your LDAP database that contains group records. These group records are of the LDAP class "groupOfUniqueNames," and the entry CN attribute constitutes the group name.
 - **User Filter**- Specify a free-form LDAP filter expression to be added to the default user search filter. The default user search filter and this additional search filter are combined as an AND expression. Records must satisfy both filters to be found using the various user searches. Any value specified in this field must follow correct LDAP search filter syntax.
 - **Group Filter**- Specify a free-form LDAP filter expression to be added to the default group search filter. The default group search filter and this additional search filter are combined as an AND expression. Records must satisfy both filters to be found using the various group searches. Any value specified in this field must follow correct LDAP search filter syntax.
 - **Default User Roles**- Specify one or more roles (by multi-selecting from the drop-down list) to be given to all users authenticated from your LDAP tree. Zope expects all users - anonymous as well as authenticated - to have the role Anonymous.
- 7 Click **Next**. The third panel (Map LDAP Groups to Local Groups) of the LDAP Configuration wizard appears.

Figure 34: LDAP Configuration Wizard (Map LDAP Groups to Local Groups)



- 8 Enter information and make selections:
- **Map LDAP Groups to Roles?**- Select this option if you want to control user roles within the Resource Manager Web interface by using Active Directory groups, instead of controlling the roles directly from within the system.

Note If you choose to use this option, then you should add the following groups to LDAP:

- Resource Manager Managers
 - Resource Manager Users
-
- **Group**- Select the LDAP group to map to a Resource Manager role.
 - **Role**- Select the Resource Manager role to map the LDAP group.
- 9 Optionally, click **Add Group Mapping** to map another group. To remove a mapped group, click **Remove**.
- 10 Click **Finish** to complete LDAP configuration.

After setup, you can edit your LDAP configuration settings from the Settings, Configuration Options, and Mappings tabs.

The Search tab allows you to locate user records on your LDAP server. Select from the list of search parameters, and optionally enter a search term, and then click **Search**. Search results return on the lower portion of the page.

Figure 35: LDAP Configuration - Search

| dn | cn | memberOf | sn | mail | givenName | sAMAccountName |
|------------------------|--------|------------------------|------|--------------------------|-----------|----------------|
| CN=User1,CN=Users,D... | User1 | CN=Group1,CN=Users,... | User | user1@zenoss-testing... | Number1 | User1 |
| CN=User2,CN=Users,D... | User2 | CN=Group1,CN=Users,... | User | user2@zenoss-testing... | Number2 | User2 |
| CN=User3,CN=Users,D... | User3 | CN=Group1,CN=Users,... | User | user3@zenoss.zenoss... | Number3 | User3 |
| CN=User4,CN=Users,D... | User4 | CN=Group1,CN=Users,... | User | user4@zenoss-testing... | Number4 | User4 |
| CN=User5,CN=Users,D... | User5 | CN=Group1,CN=Users,... | User | user5@zenoss-testing... | Number5 | User5 |
| CN=User6,CN=Users,D... | User6 | CN=Group1,CN=Users,... | User | user6@zenoss-testing... | Number6 | User6 |
| CN=User7,CN=Users,D... | User7 | CN=Group1,CN=Users,... | User | user7@zenoss-testing... | Number7 | User7 |
| CN=User8,CN=Users,D... | User8 | CN=Group1,CN=Users,... | User | user8@zenoss-testing... | Number8 | User8 |
| CN=User9,CN=Users,D... | User9 | CN=Group1,CN=Users,... | User | user9@zenoss-testing... | Number9 | User9 |
| CN=User10,CN=Users,... | User10 | CN=Group1,CN=Users,... | User | user10@zenoss-testing... | Number10 | User10 |
| CN=User11,CN=Users,... | User11 | CN=Group2,CN=Users,... | User | user11@zenoss-testing... | Number11 | User11 |
| CN=User12,CN=Users,... | User12 | CN=Group2,CN=Users,... | User | user12@zenoss-testing... | Number12 | User12 |
| CN=User13,CN=Users,... | User13 | CN=Group2,CN=Users,... | User | user13@zenoss-testing... | Number13 | User13 |
| CN=User14,CN=Users,... | User14 | CN=Group2,CN=Users,... | User | user14@zenoss-testing... | Number14 | User14 |

Verifying Connectivity and Credentials Outside of Resource Manager

You can verify that your credential information is valid from the Resource Manager server by using the `ldapsearch` command. To install this command, use the following for RPM-based systems:

```
# yum -y install openldap-clients
```

As the `zenoss` user on the Resource Manager server:

```
ldapsearch -LLL -x -b 'BaseDN' -D 'Bind DN' -W -H ldap://LDAP_server-name
\sAMAccountName=* member
```

Configuring Local Authentication as a Fallback

You can use local authentication as a fallback in the event that the LDAP server is unavailable. The local authentication plugin is called `userManager`.

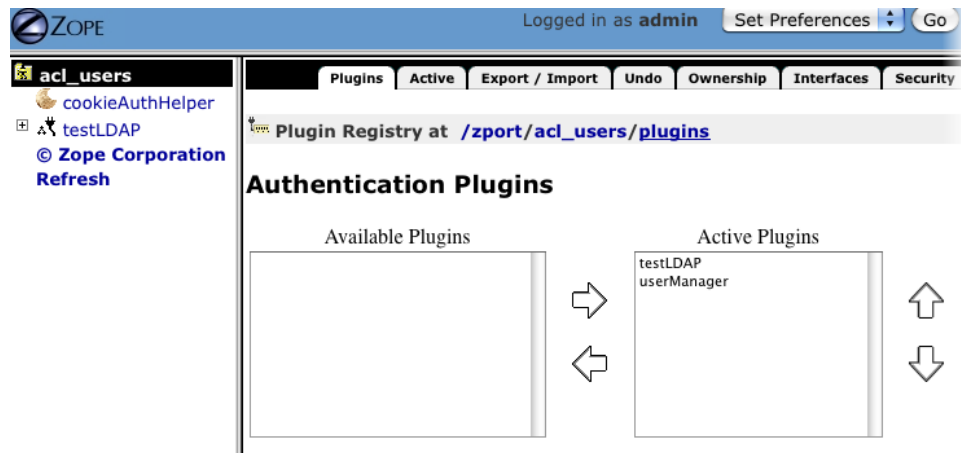
- 1 Verify that the `userManager` plugin is available.
 - a In a web browser, navigate to the Zope Management Interface.

Replace *Zenoss-Host* with the hostname or IP address of your Resource Manager server:

```
https://Zenoss-Host/zport/acl_users/manage
```

- b In the **Name** column, click **plugins**.
- c Click **Authentication Plugins**.

Figure 36: Authentication Plugins



- d Verify that your LDAP plugin is first in the list of active plugins, and that the userManager plugin is second.
- 2 Create one or more user accounts.

The accounts created in this step enable access when the LDAP server is unavailable. If you use the same account name in this step as the LDAP account name, the user need only remember the "fallback" password.

- a In Resource Manager, navigate to **ADVANCED > Users > Add > New User**.
- b Create one or more user accounts.

Note This account must be created before the account with the same name is authenticated by the LDAP server. The passwords defined when creating accounts in Resource Manager are only valid when the LDAP server is unavailable.


Visualizing your Environment

2

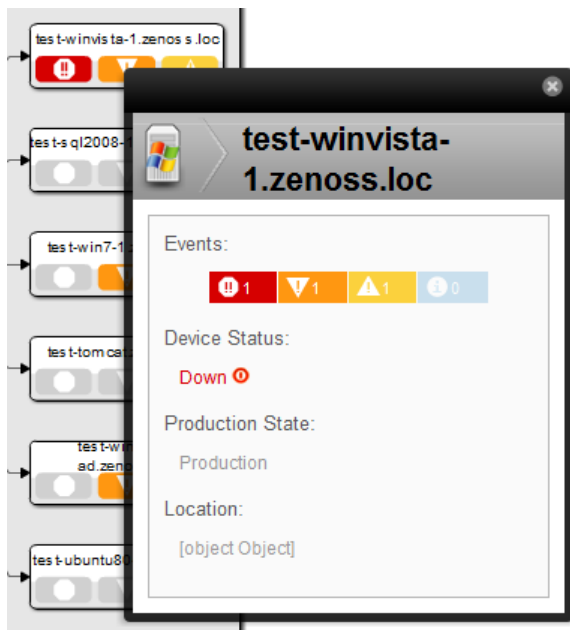
Dynamic Service View

Resource Manager provides a dynamic visualization of system objects and their relationships to other objects.

You can access the dynamic view from groups, systems, and locations. Depending on the object type, different relationships are illustrated. Each dynamic view shows related objects in a graph. Each object in that graph displays its associated event information.


Figure 37: Dynamic Service View: Locations Graph

When you click an object in the graph, the "inspector" panel appears. This panel provides detailed information about the object and links directly to it. Information that appears in the inspector depends on the object type selected.

Figure 38: Dynamic Service View: Inspector Panel

View controls appear to the right of the graph. These allow you to adjust your view:

- **Overview** - Toggles display on and off of the graph overview illustration.
- **Magnifier** - Toggles on and off the magnifier, which allows you to magnify selected portions of the graph.
- **Zoom In** - Zooms in on the graph.
- **Zoom Out** - Zooms out on the graph.
- **Fit View** - Fits the graph to the browser page.
- **Refresh** - Refreshes the graph.

Figure 39: Dynamic View: View Controls

Dynamic View of Organizers

The dynamic view of organizers shows objects that can impact the status of the organizer, such as other organizers and devices. This view also shows relationships between devices and a virtual infrastructure, such as VMware or Cisco UCS objects monitored by the system, as well as storage information.

To access the dynamic view for an organizer (such as a group, system, or location):

- 1 From Infrastructure > Devices, select the organizer in the devices hierarchy.

- 2 Click **Details**.
- 3 Select **Dynamic Service View**.

Dynamic View of Devices

The dynamic view of devices shows the relationship between a device and monitored components.

To access the dynamic view for a device:

- 1 From **INFRASTRUCTURE > Devices**, click a device in the device list. The device overview page appears.
- 2 Select **Dynamic View** in the left panel.

Dynamic View of Cisco UCS Devices

On Cisco UCS devices, the dynamic view shows the components and relationships that make up a Cisco UCS domain. The following list describes the components that are listed in the Dynamic View for various types of UCS devices:

- UCS Classic
 - UCS Domains
 - UCS Fabric Interconnects
 - UCS Fabric Ports
 - UCS IO Cards
 - UCS Fabric Extenders
 - UCS Chassis
 - UCS Blades
 - UCS Racks
 - UCS Service Profiles (only those bound to servers)
- UCS Mini
 - UCS Domains
 - UCS FI-IO Modules
 - UCS Chassis
 - UCS Blades
 - UCS Service Profiles (only those bound to servers)

Dynamic View of VMware Hosts

On VMware Hosts (ESX servers), the dynamic view shows the relative VMware elements that are connected to the host, such as:

- VMs that currently are running on the Host
- Data stores that are mounted by the Host
- Clusters to which the Host belongs

Dynamic View of Storage Devices

Storage devices such as NetApp Filers have two dynamic views:


- **Physical Storage View** - Shows the device's storage enclosures and associated hard disks.
- **Logical Storage View** - Shows the logical storage arrangement that the storage device presents, such as file systems and raid groups.

Datacenter View

Resource Manager provides a visual representation of devices (such as servers or blades) and device containers (such as racks or chassis).

With this feature, you can create a custom view that represents a physical space (such as a data center) by customizing the view background. You can then overlay this view with active representations of your devices and device containers.

Figure 40: Custom View



For each device or device container, the system can generate a rack view, which diagrams the physical location of devices in a chassis or rack. Each represented device provides at-a-glance information about its status.

Figure 41: Rack View



Configuring

Before a device or sub-location can appear in Datacenter View:

- At least one organizer must be configured
- At least one device or sub-organizer must be included in a location

To see the auto-generated rack view, you must set a rack slot value for the device.

Working with the List View

The List View provides a view of your devices (or, if configured, the Rack View). Follow these steps to access the List View:

- 1 From the interface, select **Infrastructure**.
- 2 In the Devices hierarchy, select a location, group, or system.
- 3 Click **Details**.
- 4 Select **Datacenter View**.

Note After you create a Custom View, that view appears by default.

Working with the Custom View

The Custom View lets you create a visual representation of your physical space (such as a data center).

To access the Custom View, from the Diagram selection, click Custom View.

You can edit the Custom View to:

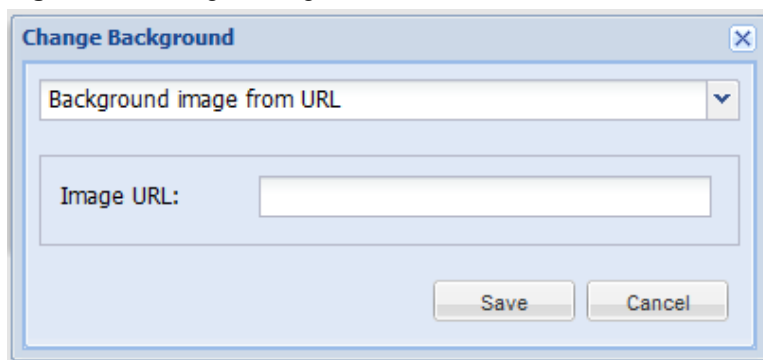
- Add or change a background image
- Move or resize device images
- Remove the view

Adding a Background Image to the Custom View

Follow these steps to create a custom view and add a background image to the view:

- 1 From the Datacenter View page (accessed from the Diagram selection), click **Custom View**.
- 2 Click **Edit** to enable edit mode.
The Edit button highlights to indicate that it is active, and Options selections become available.
- 3 Select **Options > Change Background**.
- 4 In the **Change Background** dialog, select **Background Image from URL**.
- 5 In the **Image URL** field, enter an image location, and then click **Save**.
Any image format and size supported by your browser can be used.

Figure 42: Change Background



Removing the Custom View Background Image

To remove the current background image from the Custom View:

- 1 From the Custom View area, click **Edit**.
- 2 Select **Options > Change Background**.

- 3 In the Change Background dialog, select No background image from the list of options.
- 4 Click **Save**.
The image no longer appears in the view.

Working with Devices in the Custom View

Devices in the custom view can be moved and resized. To work with devices in this view, click **Edit**. You can then drag devices to a specific location in the view, and resize them to accurately represent your physical space.

You also can view device details from this view. Click the device to go to its Status page.

Note To access device status, you cannot be in edit mode.

Removing the Custom View

Removing the custom view removes the view and custom background image, if any. To remove a custom view:

- 1 From the Datacenter View page (accessed from the Diagram selection), click **Custom View**.
- 2 Click Edit to enable edit mode.
- 3 Select **Options > Remove Custom View**.
The custom view no longer appears by default. If you select **Custom View**, devices still appear in the view; however, they are reset to default positions and sizes.

Activating the Auto-Generated Rack View

First, ensure that the device is included in a location. Then follow these steps to make devices visible in Datacenter View.

- 1 Select a device, from the list of Devices, click **Details**, and then click **Edit**.
- 2 Enter values for Rack Slot.
For more information about specifying values, see the next topic.
- 3 Click **Save**.

The device appears in Datacenter View.

- In the List View, it appears as part of a rack illustration. The rack illustration becomes the default image in the List View.
- In the Custom View, it appears as a single device image.

Note You can customize the device image by modifying the zIcon configuration property in the device class.

Rack view specification syntax


The specification of a rack view features two required and two optional key-value pairs, which are described in the following table. The value for each key is a positive integer.

| Key | Type | Description |
|-----|----------|--|
| ru | Required | The lowest rack unit used by the device. |
| rh | Required | The height of the device, in rack units. |
| st | Optional | The rack slot. |
| sc | Optional | The slot capacity; only for chassis devices. |

The syntax of a rack view specification is a comma-separated list. The following example specifies a 1u device in the second rack unit for a device named `beta.zenoss.loc`:

```
ru=2, rh=1
```

Figure 43: Example rack view specification



In the preceding example, no rack slot value is included, because there is only one device.

3

Adding, Discovering and Modeling Devices

Modeling is the process by which the system:

- Populates the device database
- Collects information about the devices in the system (such as operating system type or file system capacity)

The system models devices when they are added to the database, either manually or through the discovery process.

Add a Single Device

Follow these steps to add a single device:

- 1 From the Navigation menu, select **Infrastructure**. The Devices page appears.
- 2 Select **Add a Single Device** from the Add Devices icon. The Add a Single Device dialog appears.

Figure 44: Add a Single Device

- 3 Enter information or make selections to add the device:
 - **Hostname or IP address** - Enter the network (DNS) name or IP address of the device.

- **Device Class** - Select a device class to which this device will belong. For example, if the new device is a Linux server, then select `/Server/Linux`.
 - **Collector** - By default, this is `localhost`. Select a collector for the device.
 - **Model** - By default, this option is selected. If you do not want the device to be modeled when it is added, de-select this option.
- 4 Optional: To display additional fields, click **More**. On the expanded page, you can:
- Enter device-specific details.
 - Edit SNMP settings.
 - Set hardware and operating system information.
 - Add device comments.

Note **Hostname or IP address, Device Class, and Model** are the only required selections when adding a device. Information you add might conflict with information the system discovers about the device. Therefore, you can continue (click **Add**) without additional information or selections.

An exception is if you are adding a Cisco router in a device class other than `/Network`. In this case, before adding the device, set the value of the `zIfDescription` configuration property to `True`. Changing the value gives you additional information about Cisco routers. By default, this option is set to `True` for the `/Network` class.

- 5 Click **Add**.
- 6 Optional: To view the Add Device job in progress, click **View Job Log** in the notification that appears when you add the device.
When the job completes, the device is added in the selected device class.

Add Multiple Devices


Follow these steps to manually add multiple devices:

- 1 From the Navigation menu, select **Infrastructure**. The Devices page appears.
- 2 From the Add Devices icon, select **Add Multiple Devices**. The Add Infrastructure page appears.
- 3 Select the category, type, and connection information for each device you want to add.
- 4 Click **Add** to add the device to the system. You will see the status at the bottom of the page. Continue to add more devices until you are done.
- 5 When you have completed adding your devices, click **Done**.

Discovering Devices

You can provide network or IP address range information so that the system can discover your devices.

Follow these steps to discover devices:

- 1 From the Navigation menu, select **Infrastructure**. The Devices page appears.
- 2 Click the Add Devices icon  and select **Discover Networks** from the drop-down list. The **Network Discovery** page appears.

Network Discovery

Devices found via Discovery will be placed in the /Discovered Device Class

| Networks/Range | SNMP | SSH Authentication | Windows Authentication |
|--|--|--|--|
| Enter one or more networks (such as 10.0.0.0/24) or IP ranges (such as 10.0.0.1-50): <input type="text"/> | Community Strings: <input type="text" value="public"/> <input type="text" value="private"/> | Username: <input type="text"/> Password: <input type="text"/> | Administrator Username: <input type="text"/> Password: <input type="text"/> |
| <input type="button" value="Discover"/> | | | |

- For each network or IP range in which you want the system to discover devices, enter an address or range. For example, you might enter a network address in CIDR notation 192.0.2.0/24 or a range of IP addresses 192.0.2.1-50

Note Trying to add a /16 or /8 network can take a very long time, and may have unintended consequences.

- For each network or IP range, specify the Windows, SSH, or SNMP credentials you want Resource Manager to use on the devices it discovers. You can enter only one of each. Resource Manager will attempt to use the same credentials on each device it discovers within the networks or IP ranges specified, but will not try to automatically classify the devices.
- Click **Discover**. The discovery process iterates through every IP address in the networks and IP ranges you specify, adding each device that responds to a ping request. Further, the process adds information to any device that responds to an SNMP, WinRM, or SSH request.

Note Resource Manager uses Advanced Encryption Standard (AES) with a 256-bit key size to encrypt all passwords, and stores them in the Zope object database.

The system places discovered routers in the device path /Network/Router. Devices are placed in the /Discovered device class.

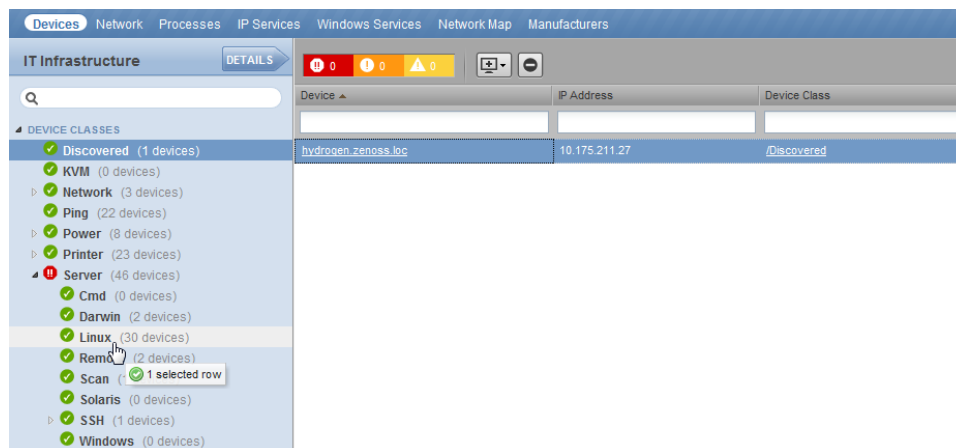
Classifying Discovered Devices

Once discovery is complete, you must move discovered devices (placed, by default, in the /Discovered class) to an appropriate device class in the hierarchy. Moving devices to their correct hierarchy location makes it possible for monitoring to begin.

Servers are organized by operating system. If the system discovers Windows devices, for example, you might choose to relocate them to /Server/Windows. Similarly, you might choose to classify discovered Linux devices in /Server/Linux (if you want to monitor and model using SNMP), or /Server/SSH/Linux (if you want to monitor and model using SSH).

To classify discovered devices:

- Select one or more discovered devices (highlight one or more rows) in the device list.
- Drag the selected devices to the new device class in the tree view.

Figure 45: Classifying Discovered Devices

The Move Devices dialog appears.

- 3 Click **OK**. The list of devices refreshes, and the devices now appear in the newly selected class.

Updating Device Authentication Details

For each device added to the database and set to its proper device class, Resource Manager may require additional or different authentication information before it can gather device information and monitor the device.

For example, for a device in the `/Server/Windows` class, you must supply your Windows user name and password before the system can monitor the device. To do this:

- 1 Click a device name in the devices list. The Device summary page appears.
- 2 Select **Configuration Properties** from the left panel.
- 3 Double-click the `zWinRMUser` configuration property to display the Edit Config Property dialog.
- 4 Enter your Windows user name in the Value field, and then click **Submit**.
- 5 Double-click the `zWinRMPasswd` configuration property to display the Edit Config Property dialog.
- 6 Enter your Windows password in the Value field, and then click **Submit**.

Similarly, for a device in the `/Server/SSH/Linux` class, you must supply your SSH user name and password. Set these values in the device's `zCommandUsername` and `zCommandPassword` configuration properties.

Note After making changes, you should remodel the device to ensure the authentication changes are valid.

Note Resource Manager uses Advanced Encryption Standard (AES) with a 256-bit key size to encrypt all passwords, and stores them in the Zope object database.

Adding or Editing Information on a Device Record

You may want to add or edit details about a device.

To add or edit information:

- 1 Click a device name in the devices list. The Device overview page appears.
- 2 You can select values to change, or click the "edit" link adjacent to a label to edit that value. Enter or change information in one or more areas, and then click **Save** to save your changes.

Modeling Devices

To model devices, the system can use:

- SSH
- WinRM
- SNMP (legacy option)

Note SSH and WinRM are the recommended options.

The modeling method you select depends on your environment, and on the types of devices you want to model and monitor.

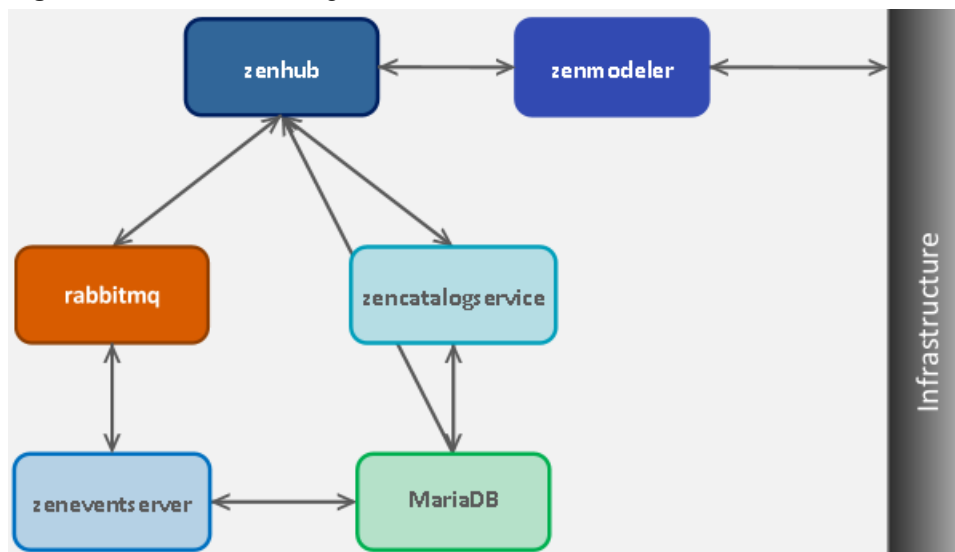
By default the system remodels each known device every 720 minutes (12 hours).

Note You can change the frequency with which devices are remodeled. Edit the value of the Modeler Cycle Interval in the collector's configuration.

For larger deployments, modeling frequency may impact performance. In such environments, you should stop the `zenmodeler` daemon and run the modeling process once daily from a cron job.

The following figure shows how the `zenmodeler` daemon fits into the modeling devices portion of the Resource Manager architecture. Note that MariaDB has replaced ZenDS (MySQL) in version 5 and stores the object database (ZODB).

Figure 46: Device Modeling Architecture



Configuring Windows Devices to Provide Data Through SNMP

By default, Windows may not have SNMP installed. To install SNMP on your particular version of Windows, please refer to the Microsoft documentation.

After setting up and configuring the SNMP service, you must set the `zSnmCommunity` string in Resource Manager to match, to obtain SNMP data.

If you want processor and memory monitoring, install *SNMP-Informant* on the device.

To collect Windows event logs or log files from a Windows box using syslog, install the *Syslog Agent* Windows add-on.

Configuring Linux Devices to Provide Data Through SNMP


To configure a Linux machine for monitoring, it must have SNMP installed. A good Linux SNMP application is net-snmp. Download, install, and configure net-snmp to then use SNMP to monitor Linux devices.

Modeling Devices Using SSH/COMMAND

You can gather additional information by running commands on the remote device and interpreting the results. This provides a more scalable and flexible way to gather information that may not be available through any other means.

Each built-in modeling command plugin is differentiated by the platform on which it runs. To determine the platform for the device you want to model, run the `uname` command in a shell on the device.

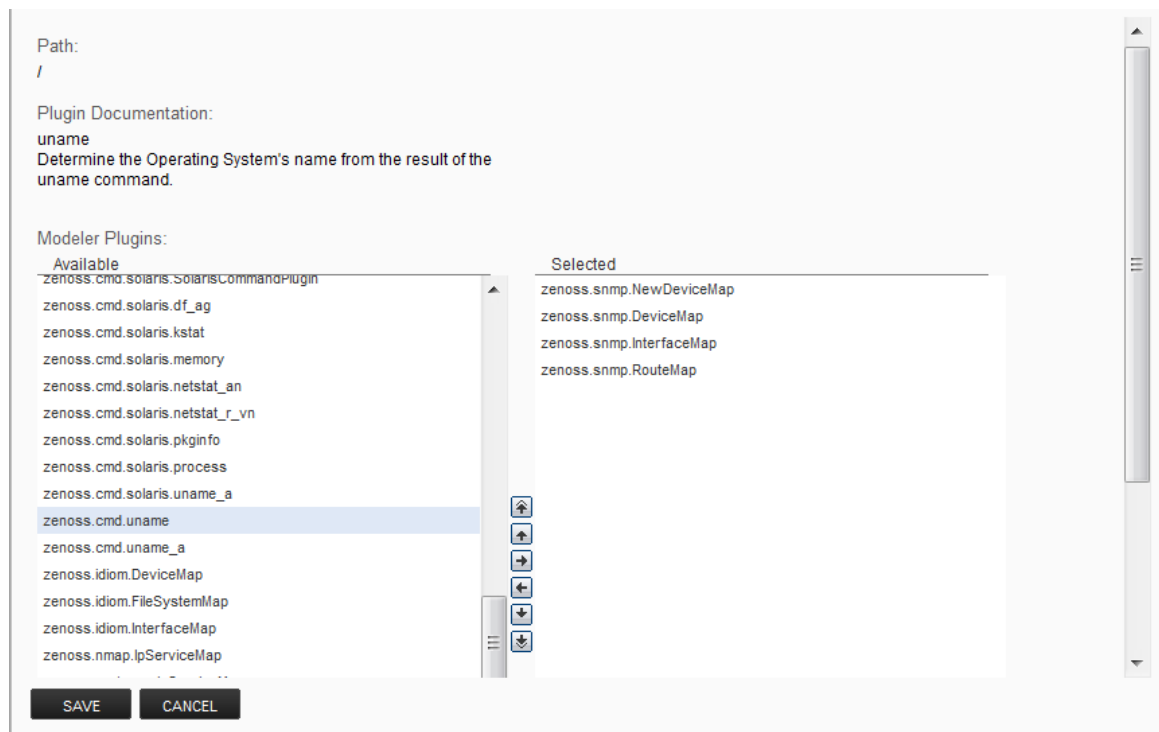
To model a device using command plugins, first add the device by using the protocol "none," and then choose the plugins you want to apply:

- 1 From the Navigation menu, select **Infrastructure**.
- 2 Click the Add Devices  icon and select **Add a Single Device** from the drop-down list. The Add a Single Device window appears.
- 3 Enter values for Name or IP and Device Class.
- 4 Clear the Model Device option.
- 5 Click **Add**.
- 6 After adding the device, select the device name in the devices list. The Device Overview page appears.
- 7 In the left panel, select **Configuration Properties**.
- 8 If necessary, set the values of the `zCommandUsername` and `zCommandPassword` configuration properties to the user name and password of the device (or set up authentication by using RSA/DSA keys.)

Note If using RSA keys for a device or device class, change the value of the `zKeyPath` configuration property to:

```
~/.ssh/id_rsa
```

- 9 In the left panel, select **Modeler Plugins**. The list of plugins appears. The left column displays available plugins; the right column shows those currently selected.
- 10 Select `zenoss.cmd.uname` from the Available list, and then use the right arrow control to move it to the Selected list on the right. Use the controls to place it at the top of the list.

Figure 47: Add Plugin

- 11 Use the left arrow control to move the other Selected plugins from the Selected list to the Available list.
- 12 Click **Save**.
- 13 Model the device by clicking the **Model Device** button.

Using Device Class to Monitor Devices Using SSH

The `/Server/Command` device class is an example configuration for modeling and monitoring devices using SSH. The `zCollectorPlugins` have been modified (see the section titled "Modeling Using SSH/Command"), and the device, file system, and Ethernet interface templates used to gather data over SSH have been created. You can use this device class as a reference for your own configuration; or, if you have a device that needs to be modeled or monitored via SSH/Command, you can place it in this device class to use the pre-configured templates and configuration properties. You also must set the `zCommandUsername` and `zCommandPassword` configuration properties to the appropriate SSH login information for each device.

Modeling Devices Using Port Scan

You can model IP services by doing a port scan, using the *Nmap Security Scanner*. You must provide the full path to your system's `nmap` command.

To determine where `nmap` is installed, at the command line, enter:

```
which nmap
```

If your system returns a result similar to:

```
/usr/bin/which: no nmap in (/opt/zenoss/bin:/usr/kerberos/bin:/usr/local/bin:/bin:/usr/bin:/opt/zenoss/bin)
```

then nmap is not installed. Install it, and then try again.

After locating the nmap command (including the directory beginning with /), enter the following as the zenoss user on the Resource Manager server:

```
cd $ZENHOME/libexec ln -s
    Full_Path_to_nmap
```

Note In order to execute a command using \$ZENHOME (/opt/zenoss for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for serviced commands.

To model a device using a port scan:

- 1 Select the device in the device list.
- 2 In the left panel, select **Modeler Plugins**.
- 3 Select the zenoss.nmap.ipServiceMap plugin in the list of Available plugins, and then use the right arrow control to move it to the list of Selected plugins.
- 4 Click **Save**.
- 5 Remodel the device by clicking the **Model Device** button.

Using the /Server/Scan Device Class to Monitor with Port Scan

The /Server/Scan device class is an example configuration for modeling devices by using a port scan. You can use this device class as a reference for your own configuration; or, if you have a device that will use only a port scan, you can place it under this device class and remodel the device.

About Modeler Plugins

Resource Manager uses plug-in maps to map real world information into the standard model. Input to the plug-ins can come from SNMP, SSH or Telnet. Selection of plug-ins to run against a device is done by matching the plug-in name against the zCollectorPlugins configuration property.

- **DeviceMap**– Collects basic information about a device, such as its OS type and hardware model.
- **InterfaceMap**– Collects the list of network interfaces on a device.
- **RouteMap**– Collects the network routing table from the device.
- **IpServicesMap**– Collects the IP services running on the device.
- **FileSystemMap**– Collects the list of file systems on a device.

Viewing and Editing Modeler Plugins for a Device

Plugins are controlled by regular expressions that match their names. To view a list of plugins for any device:

- 1 Click the device name in the devices list.
The Device summary page appears.
- 2 In the left panel, select **Modeler Plugins**.
The Modeler Plugins page appears.

Adding Plugins

To add a plugin to a device:

- 1 Use the right arrow control to move one or more plugins from the Available list (on the left) to the Selected list (on the right).
- 2 Click **Save**.

Reordering Plugins

Plugins run in the order in which they are listed. To re-order plugins, use the up and down arrow controls, and then click **Save**.

Deleting Plugins from a Device

To delete a plugin from a device, use the left arrow control to move the plugin from the Selected list to the Available list.

Debugging the Modeling Process

You can run the modeler from the command line against a single device. This feature is useful when debugging issues with a plugin.

By passing the `--collect` command to the modeler, you can control which modeler plugins are used. For example, the following command runs only the interface plugin against the `build.zenoss.loc` device:

- 1 Log in to the Control Center host as a user with `serviced` CLI privileges.
- 2 Attach to the `zenmodeler` service.

```
serviced service attach zenmodeler
```

- 3 Change to the `zenoss` user.

```
su - zenoss
```

- 4 Run the `zenmodeler` command.

```
$ zenmodeler run -v10 --collect=IpInterface -d build.zenoss.loc
```

If the command returns any stack traces, forward these details to Zenoss Support for assistance:

- Command you ran
- Stack trace or stack traces returned
- Version of your Resource Manager instance
- OS version and patch level for the remote device

Working with Devices

This chapter provides information and procedures for managing devices in the system.

Viewing the Device List

The device list shows all devices in the system. From this view, you can search for devices and perform a range of management tasks on all devices.

To access the device list, select Infrastructure from the Navigation menu.

Figure 48: Device List

| Device | IP Address | Device Class | Production State | Events |
|---|----------------|----------------------|------------------|--------|
| 32-bit Machines | | /VMware/Esxwin2_Colo | Production | |
| 64-bit Machines | | /VMware/Esxwin2_Colo | Production | |
| auslinbot.zenoss.loc | 10.87.209.6 | /Server/Linux | Production | |
| Dev_Resource_Pool | | /VMware/Esxwin2_Colo | Production | |
| doc-dev.zenoss.loc | 10.175.211.217 | /Server/Linux | Production | ▲ 1 |
| ec2-50-16-148-220.compute-1.amazonaws.com | 50.16.148.220 | /Server/Linux | Production | ● 1 |
| ec2-50-19-130-2.compute-1.amazonaws.com | 50.19.130.2 | /Server/Linux | Production | ● 1 |
| ec2-72-44-49-24.compute-1.amazonaws.com | 72.44.49.24 | /Server/Linux | Production | ▼ 1 |
| EC2Manager | | /AWS/EC2 | Production | |
| esx1.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |
| esx10.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |
| esx10.storage1 | | /VMware/Esxwin2_Colo | Production | |
| esx12.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |
| esx12.storage1 | | /VMware/Esxwin2_Colo | Production | |
| esx13.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |
| esx13.storage1 | | /VMware/Esxwin2_Colo | Production | |
| esx14.Storage1 | | /VMware/Esxwin2_Colo | Production | |
| esx14.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |
| esx15.zenoss.loc | | /VMware/Esxwin2_Colo | Production | |

Devices Hierarchy

Devices are organized in the tree view by:

- Devices
- Groups
- Systems

- Locations
- Component Groups

Click the indicator next to each category name to expand it and see included devices.

Managing Multiple Devices from the Device List

You can perform some management tasks for more than one device at a time. You can:

- Move devices to a different class
- Assign devices to groups, systems, and locations
- Remove devices
- Perform actions such as assigning priority, production state, or collector
- Lock devices

Working with Devices

To view details for a single device, click its name in the device list. The device overview page appears.

Figure 49: Device Overview

The screenshot displays the Cisco UCS Manager interface for a specific device. At the top, a navigation bar includes 'Devices', 'Networks', 'Processes', 'IP Services', 'Windows Services', 'Network Map', and 'Manufacturers'. Below this, a status bar shows the device IP '10.87.208.163', its name '/CiscoUCS/UCS-Manager', and various status indicators: a red '2' (errors), a yellow '61' (warnings), a green '13' (info), and a blue '49' (events). The device status is 'Up' (green), and its production state is 'Production' and priority is 'Normal'.

The main content area is divided into several sections:

- Overview:** A sidebar menu with options like Overview, Events, Components, Graphs, Component Graphs, Modeler Plugins, Software, Custom Properties, Configuration Properties, Device Administration, Dynamic View, Dependency View, and Monitoring Templates (Device (/CiscoUCS)).
- Device Information:** Fields for Device ID (10.87.208.163), Connection Information (admin 443 True), Uptime (104d:02h:35m:19s), First Seen (2016-01-14 10:40:59 am), Last Change (2016-01-14 11:12:20 am), Model Time (2016-01-14 10:42:21 am), Locking (Unlocked), and a 'Connect to this device' link.
- Configuration Fields:**
 - Device Name: 10.87.208.163
 - Production State: Production
 - Priority: Normal
 - Tag: (empty)
 - Serial Number: (empty)
 - Rack Slot: 0
- Collector and Manufacturer:** Collector: localhost; Hardware Manufacturer: Cisco; Hardware Model: UCS; UCSM Version: UCSM 2.2(5c).
- Systems, Groups, Location, Links, and Comments:** Systems: None; Groups: None; Location: None; Links: (empty); Comments: (empty text area).

At the bottom, there are buttons for '+', a gear icon, 'Commands', 'Model Device', and a '0 Jobs' indicator.

Event status is shown in the "event rainbow" at the top of the page. Other key information that appears at the top of the device overview page includes:

- Device name
- IP address used to communicate with the device
- Device status (shows the current results of a ping test)
- Production state (Pre-Production, Production, Test, Maintenance, or Decommissioned)
- Priority

When you open the page, device overview information displays. This view provides classification and status information. From here, you can edit device information (indicated by text fields or edit links). Editable fields include:

- Device Name
- Production State
- Priority
- Tag
- Serial Number
- Rack Slot
- Collector
- Hardware and software manufacturer and model
- Systems
- Groups
- Location

The Links area displays links between the device and other external systems.

The left panel of the device overview page allows you to access other device management views, such as:

- Events
- Components
- Graphs (Performance)
- Component Graphs
- Modeler Plugins
- Software
- Custom Properties
- Configuration Properties
- Device Administration
- Dynamic View
- Dependency View
- Monitoring Templates

Information that appears here varies depending on device type.

Events

Detailed information about events, scoped to the device, appears in the Events view. From here, you can:

- Sort event and event archive information by a range of categories
- Classify and acknowledge events
- Filter events by severity, state, or by one of several categories


Components

The Components view provides information about the different types of device components, including:

- IPService
- WinService
- IpRouteEntry
- IpInterface
- CPU
- FileSystem

To access components information, select Components in the left panel, and then select a component type. The components available will vary based on the type of device.

Figure 50: Device (Components)



The status of each device component type, as shown by the color of its indicator, is determined by the collective status of the monitored components of the same type. For example, if the IPService status is green, then all monitored IPServices on the device are functioning normally. If there is an event related to a monitored IPService, then the highest severity event associated with that component is displayed.

Note If there is an event unrelated to a known component, then the system places it in the component type Other.

From this view, you can:

- Lock components
- Turn on or off component monitoring
- Delete components

Disabling component monitoring

There may be occasions when you want to stop monitoring certain components of your monitored resources.

To disable monitoring on one or more components:

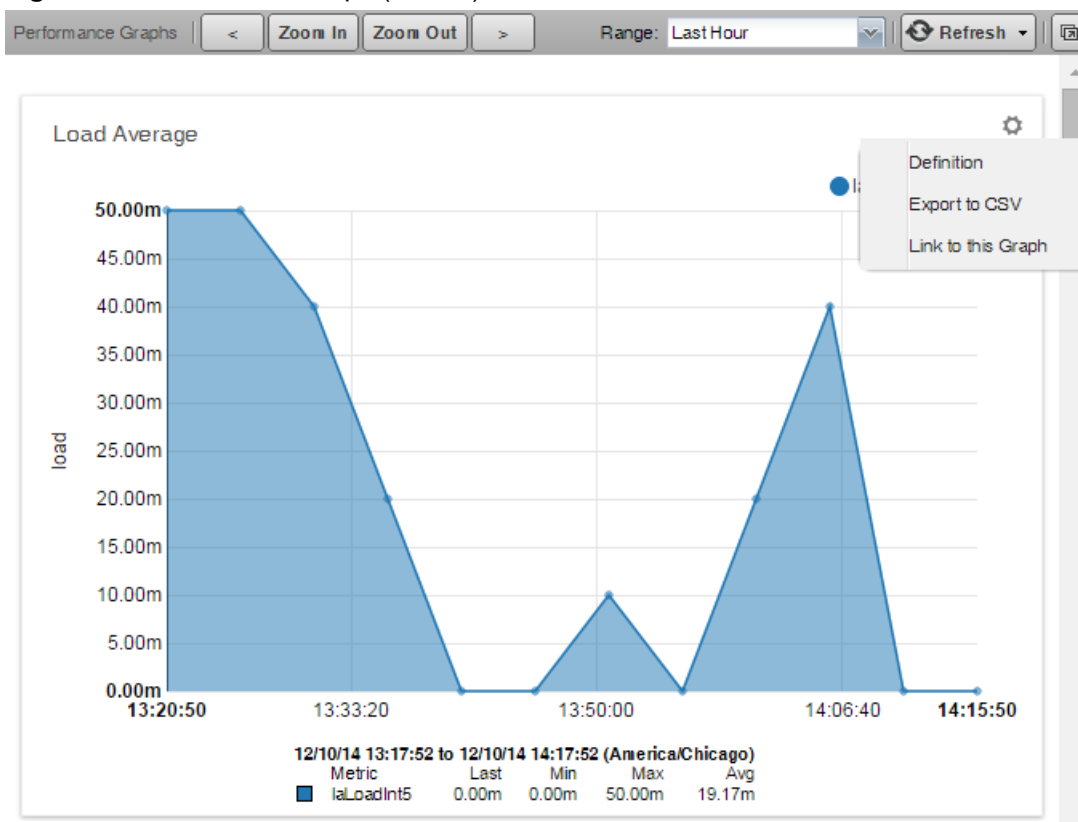
- 1 On the Device overview page, select the component group.
- 2 In the Component list, select the components for which you want to disable monitoring.
- 3 Click **Action > Monitoring**.
- 4 Click **NO** to disable monitoring.

You may want to clean up the Event log of any events that were created by these components prior to the disabling of monitoring, see [Closing Events](#) on page 138.

Graphs (Performance)

The Graphs view shows performance graphs defined for the device. To access graphs, select **Graphs** in the left panel.

Figure 51: Performance Graph (Device)



Note You can use the arrow key and magnifying glass controls on the sides of each graph to change the graph view, scrolling through or zooming in or out of a graph.

You can control these performance graph options:

- **Range** - Select the span of time displayed in the graph. You can select:
 - Last Hour
 - Yesterday
 - Last Week

- Last 30 days
- Last Year
- Custom
- **Refresh** - Modify the refresh value (by default, 30 minutes) by clicking the drop-down list. Setting the refresh rate to manual requires you to click the Refresh button each time you want an updated graph.
- **Full Screen** - Click the Full Screen icon to display only the graphs in your browser.

Each graph can be further controlled by clicking the Action icon located in the upper-right corner of the graph. You can perform the following actions:

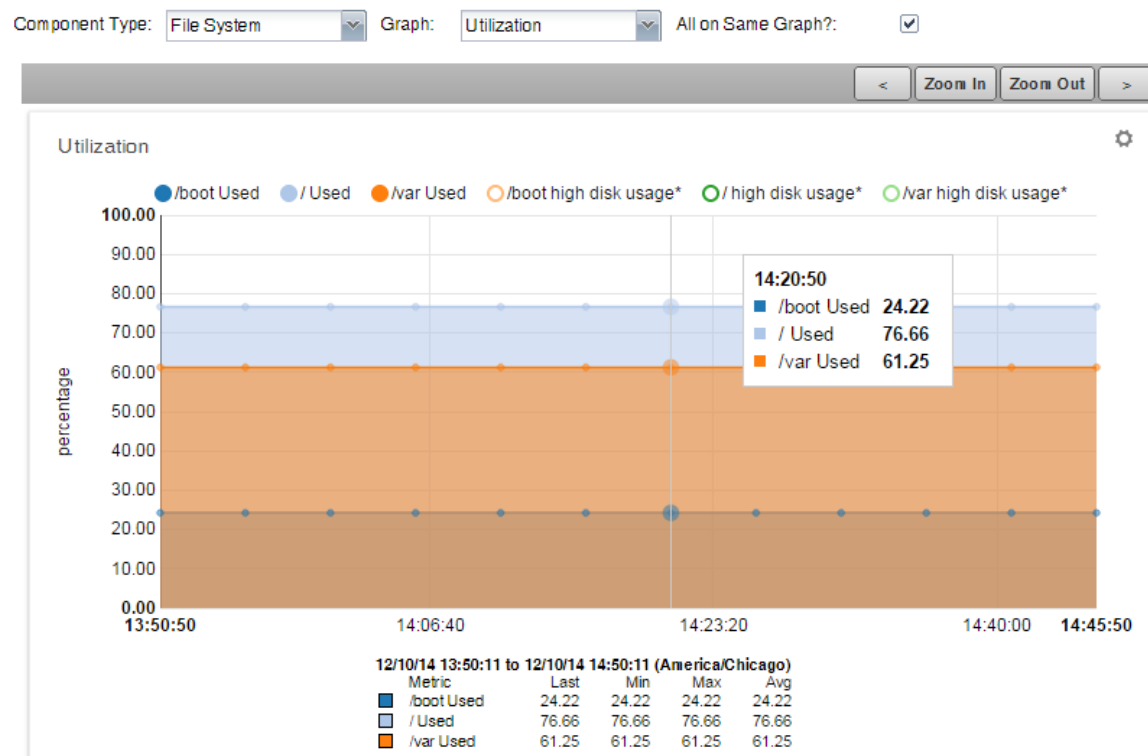
- **Definition** - View the JSON definition.
- **Export to CSV** - Export the datapoints as a .csv file for use in a spreadsheet.
- **Link to this Graph** - Generate a link to this graph that can be saved to your browser bookmarks or use the URL to directly point to the graph in another Web page or dashboard. For example, you may want to show this graph in the Zenoss Dashboard. You can create a new Site Window portlet and insert the URL to this graph.

For more information about performance monitoring and performance graphs, see [Performance Monitoring](#) on page 110.

Component Graphs

The Component Graphs view shows component graphs defined for the device. To access these graphs, select **Component Graphs** in the left panel. The following figure shows a file system utilization graph that has three metrics all displayed on the same graph.

Figure 52: Component Graph (Device)



Note You can use the arrow key and magnifying glass controls on the sides of each graph to change the graph view, scrolling through or zooming in or out of a graph.

You can control these component graph options:

- **Component Type** - Drop-down of available components based on the type of device being monitored.
- **Graph** - Drop-down of available graphs based on the selected component type.
- **All on Same Graph?** - Select this check box to display all the metrics on one graph. Clear this check box to have a separate graph for each metric.

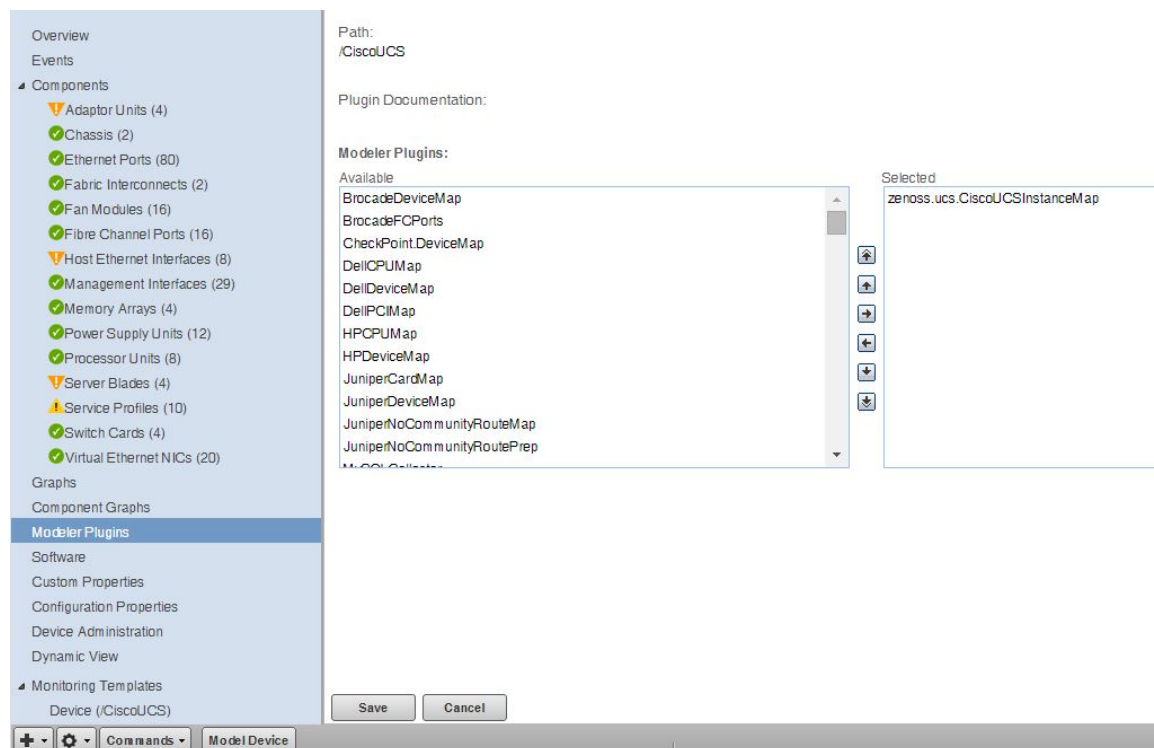
Each graph can be further controlled by clicking the Action icon located in the upper-right corner of the graph. You can perform the following actions:

- **Definition** - View the JSON definition.
- **Export to CSV** - Export the datapoints as a .csv file for use in a spreadsheet.
- **Link to this Graph** - Generate a link to this graph that can be saved to your browser bookmarks or use the URL to directly point to the graph in another Web page or dashboard. For example, you may want to show this graph in the Zenoss Dashboard. You can create a new Site Window portlet and insert the URL to this graph.

Modeler Plugins

Use the Modeler Plugins view to manage plugins that are run against a device. To access plugins, select Modeler Plugins in the left panel.

Figure 53: Device (Modeler Plugins)



Software

The Software view lists software installed on the device. The details provided in this area depend on the method used to model the device.

Listed software links into the system's inventory of software in your IT infrastructure. You can view this inventory from the Manufacturers link on the sub-navigation menu.

To access software information, select Software in the tree view.

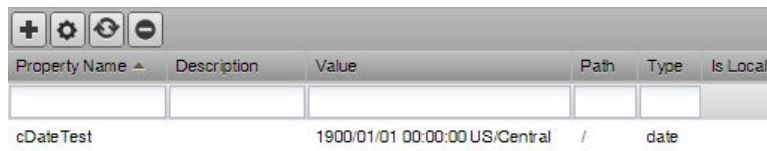
Custom Properties

Use the Custom Properties view to edit the values of custom properties associated with a device.

To display the Custom Properties view, click on a device name in the device list, and then select **Custom Properties** in the left panel. You can perform the following actions on custom properties:

- Add
- Edit
- Refresh
- Delete

Figure 54: Device (Custom Properties)



| Property Name | Description | Value | Path | Type | Is Local |
|---------------|-------------|--------------------------------|------|------|----------|
| cDateTest | | 1900/01/01 00:00:00 US/Central | / | date | |

Note The Custom Properties view allows you to edit the value of a custom property on an individual device, but not to define new custom properties for device classes. For more information about custom properties, refer to knowledge base article ["How To Add Custom Properties"](#) on the [Zenoss Support](#) site.

Configuration Properties

From the Configuration Properties view, you can configure certain configuration properties for a device. Further, you can delete local properties from a device.

To access configuration properties, select Configuration Properties in the left panel.

Figure 55: Device (Configuration Properties)

| Is Local | Category | Name | Value | Path |
|----------|------------------|-----------------------------|--------------------|------|
| | Cisco | zCiscoACEUseSSL | true | / |
| | Cisco | zCiscoRemodelEventClassKeys | | / |
| | Cisco UCS | zCiscoUCSCIMCEventsInterval | 60 | / |
| | Cisco UCS | zCiscoUCSCIMOPerInterval | 300 | / |
| | Cisco UCS | zCiscoUCSManagerPort | 443 | / |
| | Cisco UCS | zCiscoUCSManagerUseSSL | true | / |
| | Cisco UCS | zCiscoUCSManagerUser | admin | / |
| | Modeler Controls | zCollectorClientTimeout | 180 | / |
| | Modeler Controls | zCollectorDecoding | utf-8 | / |
| | Misc | zCollectorLogChanges | false | / |
| | zencmmand | zCommandCommandTimeout | 15 | / |
| | zencmmand | zCommandExistenceTest | test -f %s | / |
| | zencmmand | zCommandLoginTimeout | 10 | / |
| | zencmmand | zCommandLoginTries | 1 | / |
| | zencmmand | zCommandPassword | | / |
| | zencmmand | zCommandPath | \$ZENHOME/libexec | / |
| | zencmmand | zCommandPort | 22 | / |
| | zencmmand | zCommandProtocol | ssh | / |
| | zencmmand | zCommandSearchPath | | / |
| | zencmmand | zCommandUsername | | / |
| | Control Center | zControlCenterHost | \$(here/managerip) | / |
| | Control Center | zControlCenterModelCycle | 3600 | / |
| | Control Center | zControlCenterPassword | | / |
| | Control Center | zControlCenterPerfData | onn | / |

For detailed information about working with configuration properties, see [Configuration Properties](#) on page 77.

Device Administration

Use the Device Administration view to:

- Add, delete, and run custom user commands
- Manage maintenance windows
- Determine who holds administration capabilities for the device, and their roles

To access administration options, select **Device Administration** in the left panel.

Figure 56: Device Administration

The screenshot displays the Device Administration interface. On the left is a navigation menu with options like Devices, Events, and Device Administration. The main area is divided into three sections:

- Maintenance Windows:** A table with columns: Enabled, Name, Start, Duration, Repeat, and State. It lists three windows: '1st of Month' (Monthly), 'Every Thursday Night' (Weekly), and 'One Time Testing' (Never).
- User Commands:** A table with columns: Name and Command. It lists commands like 'DNS forward', 'DNS reverse', 'ping', 'snmpwalk', and 'traceroute' with their respective command strings.
- Administrators:** A table with columns: Name, Role, Email, and Pager. It lists one administrator: 'calchrist' with the role 'Manager'.

Each section includes pagination controls and a 'DISPLAYING' status indicator.

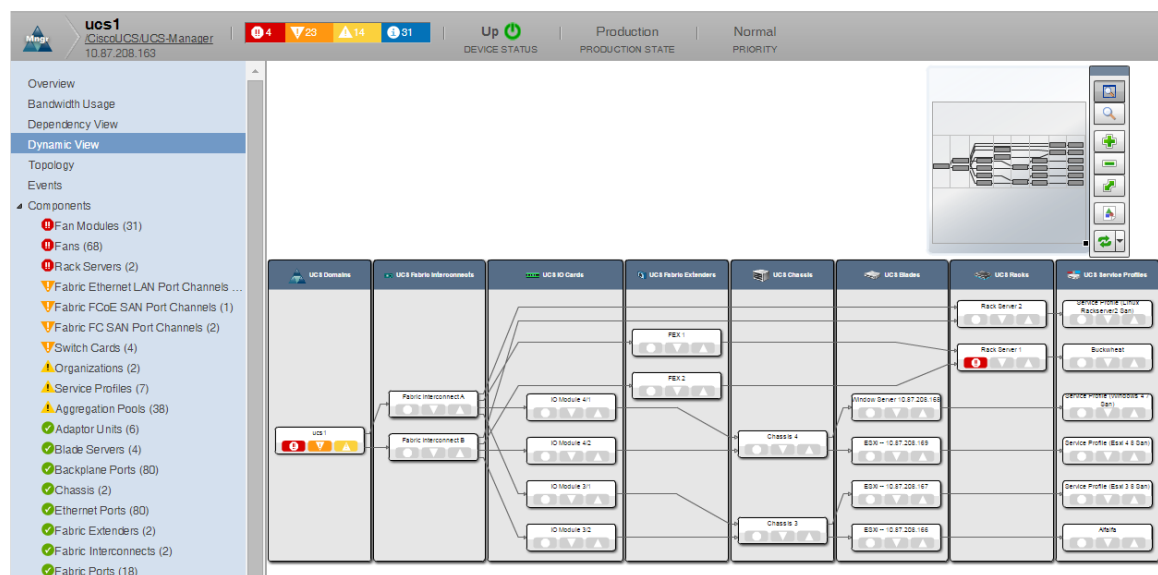
See these topics for more information about device administration tasks:

- [Maintenance Windows](#) on page 154
- [Defining User Commands for a Single Device](#) on page 172
- [Defining User Commands for All Devices in an Organizer](#) on page 173
- [Adding Administrators](#) on page 177

Dynamic View

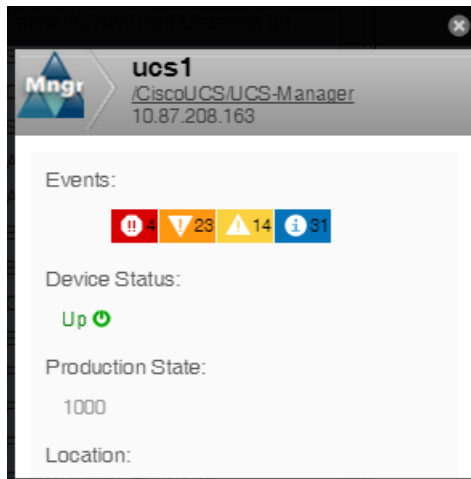
Resource Manager provides a dynamic visualization of system objects and their relationships to other objects.

You can access a dynamic view from a device overview, a group, a system, or a location. Depending on the object type, different relationships are illustrated. Each dynamic view shows related objects in a graph. Each object in that graph displays its associated event information.



When you click an object in the graph, the "inspector" panel appears. This panel provides detailed information about the object and links directly to it. Information that appears in the inspector depends on the object type selected.

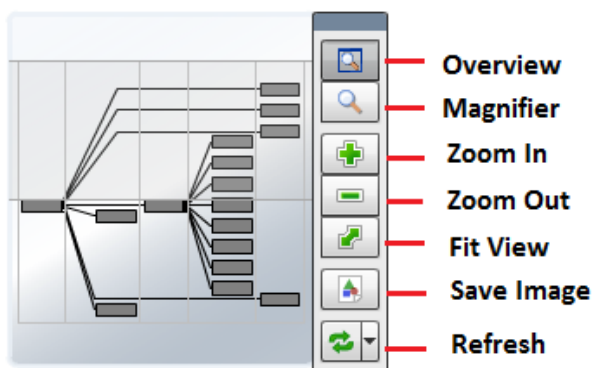
Figure 57: Dynamic View: Inspector Panel



View controls appear to the right of the graph. These allow you to adjust your view:

- **Overview** - Toggles display on and off of the graph overview illustration.
- **Magnifier** - Toggles on and off the magnifier, which allows you to magnify selected portions of the graph.
- **Zoom In** - Zooms in on the graph.
- **Zoom Out** - Zooms out on the graph.
- **Fit View** - Fits the graph to the browser page.
- **Save Image** - Saves the dynamic view as a .png image.
- **Refresh** - Refreshes the graph.

Figure 58: Dynamic View: View Controls



Dynamic View of Devices

The dynamic view of devices shows the relationship between a device and monitored components.

To access the dynamic view for a device:

- 1 From **INFRASTRUCTURE > Devices**, click a device in the device list. The device overview page appears.
- 2 Select **Dynamic View** in the left panel.

Dependency View

The dependency view of a device shows the resources that are dependent on the selected device as well as those resources that the device is dependent on. In this view, you can see the resource, associated component, and any events related to the component.

There are several places where you can see a dependency view:

- Device Overview page
 - Navigate to the device overview page and select **Dependency View**
- Device Component page
 - Navigate to the device overview page and select the component you are interested in. In the **Display** drop-down list, select **Dependencies**.
- Group Details page (including Groups, Systems, and Locations)
 - Navigate to the Group, System or Location name and click **Details**. Then, select **Dependency View**

Regardless of how you navigate to the Dependency View, the functionality remains the same. The following shows a sample Dependency View.

Figure 59: Dependency View

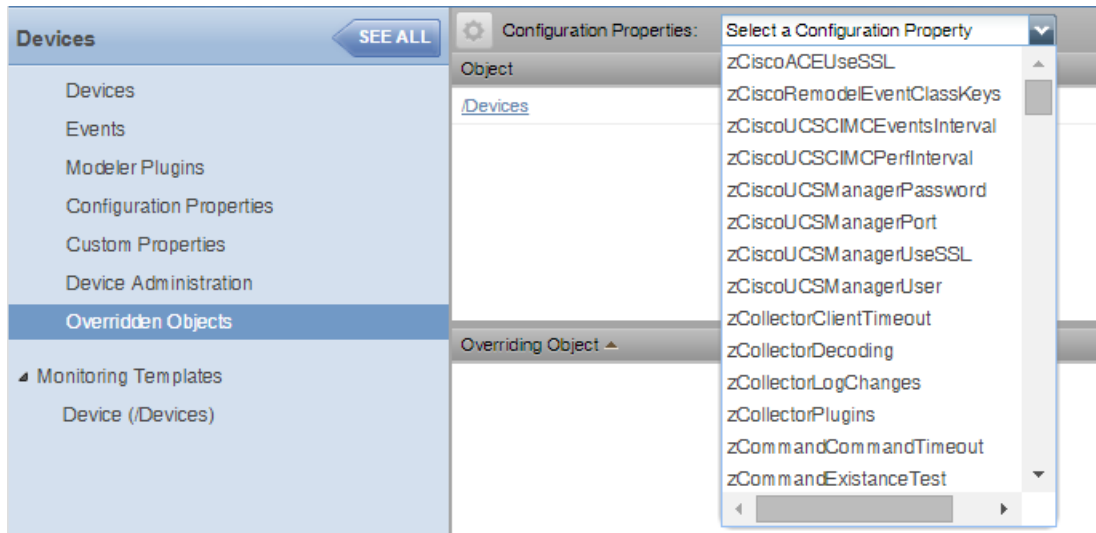
- 1 Click **Dependents** to see the resources that are dependent on the selected device or component. Click **Dependencies** to see the resources that this device or component is dependent on. For example, if you click the Dependency View of a Hyper-V server, you will see that VMs, datastores, and other resources are dependent on the Hyper-V server, while the server itself is dependent on the Host CPUs, HDDs, and Network Adapters (seen after clicking **Dependencies**).
- 2 Click on a Component to expand its dependents or dependencies.
- 3 Click **Groups** to select the resources that you want to display in the view.
- 4 Export the data displayed by clicking the **Export** button. A .csv file is exported with the data as it is presented in the view.
- 5 Click a resource name to be taken to its overview page which displays performance graphs and other details.

All selections made with respect to the display of the view will be saved so that when a particular user selects another dependency view, Resource Manager renders the dependency view using the same parameters. This is only applicable for the same user. There is no global setting available for the dependency view.

Overriden Objects

Use the Overriden Objects view to see the objects that have overrides on their configuration properties. This view is available when looking at details of all devices.

To display the Overriden Objects view, navigate to the **Infrastructure > Devices** pages and click **Details**. Then, click **Overriden Objects** from the left-column menu.



Select a configuration property from the drop-down list to view the overridden objects for that property. Double-click the row of the overriding object to open an edit dialog box.

Note Do not click the link of the overriding object. You will be taken to that object's page in the infrastructure view. Instead, double-click the clear area of the row of the overriding object to view the Edit Configuration Property dialog.

Monitoring Templates

Monitoring templates determine how the system collects performance data for devices and device components.

To access monitoring templates, expand Monitoring Templates in the left panel, and then select Device. The page shows all of the monitoring templates that are bound by name to this device.

Figure 60: Device (Monitoring Templates)

| Data Sources | | | | Thresholds | |
|------------------------------|---------|------|--------------|------------|--|
| Data Points by Data Source ▲ | | | | Name | |
| Source | Enabled | Type | | | |
| laLoadInt1 | true | SNMP | high load | | |
| laLoadInt15 | true | SNMP | low CPU idle | | |
| laLoadInt5 | true | SNMP | | | |
| memAvailReal | true | SNMP | | | |
| memAvailSwap | true | SNMP | | | |
| memBuffer | true | SNMP | | | |
| memCached | true | SNMP | | | |
| ssCpuIdle | true | SNMP | | | |
| ssCpuRawWait | true | SNMP | | | |
| ssCpuSystem | true | SNMP | | | |
| ssCpuUser | true | SNMP | | | |
| ssIORawReceived | true | SNMP | | | |
| ssIORawSent | true | SNMP | | | |
| sysUpTime | true | SNMP | | | |

| Graph Definitions | |
|--------------------|--|
| Name | |
| Load Average | |
| CPU Utilization | |
| Memory Utilization | |
| IO | |

For detailed information about monitoring templates, see [Performance Monitoring](#) on page 110.

Managing Devices and Device Attributes

Read the information and procedures in this section to learn about specific device management tasks, including:

- Clearing heartbeat events
- Pushing configuration changes to the system
- Locking device configuration
- Renaming devices
- Remodeling devices
- Setting the device manage IP address

Clearing Heartbeat Events

If you have devices configured to send a recurring event that is mapped to a heartbeat class, you can clear stale heartbeat events.

To clear the heartbeat events associated with a device:

- 1 Navigate to **ADVANCED > Settings**.
- 2 In the left panel, select **EVENTS**.
- 3 At the bottom of the Event Configuration page, click the **Clear** button in the Clear Event Heartbeats section. The system displays a brief message banner.

Pushing Configuration Changes to Resource Manager

When you make a configuration change, it is automatically propagated to all the remote collectors. If you think that your change has not been propagated correctly, then you can manually force a configuration "push" to the collectors.

To push configuration changes:

- 1 Navigate to the device in the device list.
- 2 At the bottom of the device overview page, select **Push Changes** from the Action menu. The Push Changes dialog appears.
- 3 Click **OK**. The system pushes changes to the collectors and displays a confirmation message of the action.

Locking Device Configuration


You can lock a device's configuration to prevent changes from being overwritten when remodeling the device. Two levels of locking are available. You can lock the configuration from deletion and updates, or solely from deletion.

Note Device locking prevents changes and deletion due to remodeling. It does not prevent manual changes and deletion.

To edit lock selections for a device configuration:

- 1 Navigate to the device in the device list.
- 2 At the bottom of the device overview page, select **Locking** from the Action menu. The Lock Device dialog appears.

Figure 61: Lock Device Dialog



- 3 Select the type of lock you want to implement or remove.
- 4 To send events when actions are blocked by a lock action, select the "Send an event..." option. The lock or unlock action is implemented on the device, and the system displays a confirmation message of the action.

Renaming a Device

Because the system uses the manage IP to monitor a device, the device name may be different than its fully qualified domain name (FQDN). The device name must always be unique in the system.

To rename a device:

- 1 Navigate to the device in the device list. Click the device name.
- 2 On the device overview page, edit the Device Name field with the new device name.
- 3 Click **Save**. The system renames the device and displays a confirmation message of the action.

Reidentifying a Device

You may want to change the Device ID in the system. This is different than changing the Device Name. If you change the ID, you will lose all graph data for this device.

To reidentify a device:

- 1 Navigate to the device in the device list. Click the device name to open the Device Overview page.
- 2 At the bottom of the Device Overview page, select **Reidentify Device** from the Action menu.
- 3 Enter a new ID for the device.
- 4 Click **Submit**.

Remodeling a Device

Remodeling forces the system to re-collect all configuration information associated with a device. Normally, the system models devices every 720 minutes; however, if you want to remodel a device immediately, follow these steps:


- 1 Navigate to the device in the device list and click on the Device name.
- 2 At the bottom of the Device Overview page, click the **Model Device** button. The system remodels the device. A dialog appears that shows progress of the action.

Resetting the Device Manage IP Address

You might want to reset the manage IP address if the IP address of a device has changed and you want to maintain the historical data at the original IP address. To reset the manage IP address of a device:

- 1 Navigate to the device in the device list.
- 2 At the bottom of the device overview page, select Reset/Change IP Address from the Action menu. The Reset IP dialog appears.

Figure 62: Reset IP Dialog



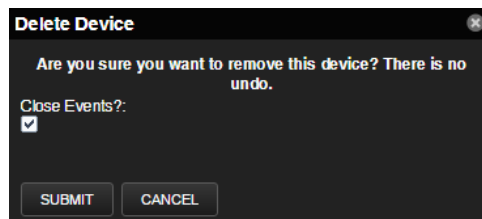
- 3 Enter the new IP address for the device, or leave the field blank to allow the IP address to be set by DNS.
- 4 Click **Save**. The IP address for the device is reset.

Deleting a Device

To delete a device from the system:

- 1 Navigate to the **INFRASTRUCTURE** page.
- 2 Select the device you want to remove from the system by clicking on its row. You can select multiple devices by Ctrl-clicking or Shift-clicking the devices. Be sure to click on the row in an area that is not defined by a link. The Delete Device dialog appears.

Figure 63: Delete Device



- 3 Optional: Change the selection to close current events for the device. By default, event data is removed.
- 4 Click **Submit**. The system removes the device(s) and associated data (if selected), and displays a confirmation message of the action.

Dumping and Loading Devices

The system allows you to export a list of your devices to a text file to import into another system instance. You can do this by using the `zenbatchdump` command on the CLI.

You cannot use `zenbatchdump` on devices whose classes were defined in a ZenPack.

- 1 Log in to the Control Center host as a user with `serviced` CLI privileges.

- 2 Attach to the zenhub service.

```
serviced service attach zenhub
```

- 3 Change to the zenoss user.

```
su - zenoss
```

- 4 Export the device list to a text file.

```
zenbatchdump > mydevicelist.txt
```

This command writes the names of your devices (including their device classes, groups, and systems) to a file named `mydevicelist.txt`.

To load these devices into another instance (or reload them into the same instance), while attached to the zenhub service of the system instance where you want the devices to be discovered, run the command:

```
zenbatchload mydevicelist.txt
```

The system attempts to discover each of the devices in the text file.

For additional ways to add and discover devices, see [Adding, Discovering and Modeling Devices](#) on page 51.

5

Configuration Properties

Configuration properties are individual values that you can set up on major system entities.

- **Devices** *Device configuration properties* control the way that devices are monitored.
- **Events** *Event configuration properties* control the rules that process events as they are received by the system.
- **Networks** *Network configuration properties* control options that are used when you perform network discovery.

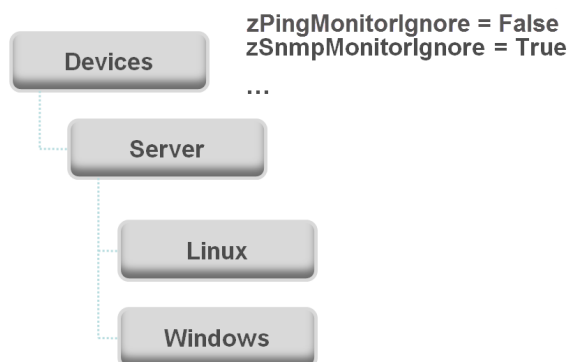
Configuration properties and values can be added to the ZenPacks you create, allowing you to customize the system when you add ZenPacks.

Configuration Properties Inheritance and Override

The following diagram illustrates a portion of the standard device class hierarchy. (A *device class* is a special type of organizer used to manage how the system models and monitors devices.)

At the root of the device hierarchy is the **Devices** object. All device class configuration properties are defined here. Their values are the default values for the entire hierarchy.

Figure 64: Device Class Hierarchy



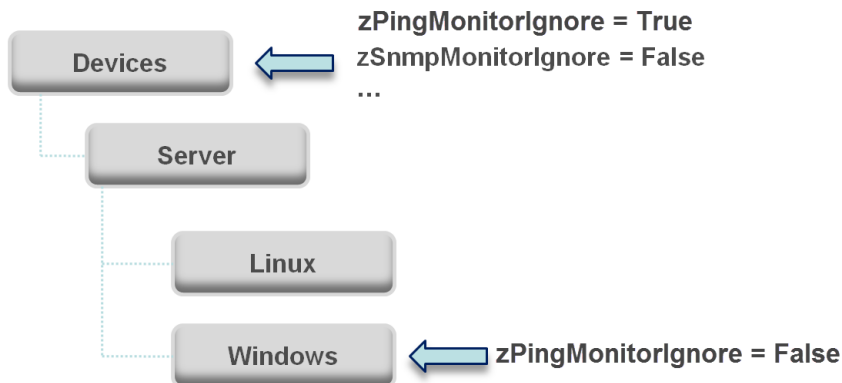
The illustration shows two defined configuration properties:

- **zPingMonitorIgnore**- Turns off all daemons that use ping. By default, its value at the root of the hierarchy is False.
- **zSNMPMonitorIgnore** - Turns off all daemons that use SNMP. By default, its value at the root of the hierarchy is True.

Through *inheritance*, properties that are defined at the root of the hierarchy apply to all objects beneath that node. So, at the `/Devices/Server/Linux` level of the device class hierarchy, the value of these two properties is the same as at `/Devices`, even though the property is not set explicitly at `/Devices/Server/Linux`. Inheritance simplifies system configuration because default values that are set at the root level apply to all devices regardless of their device class.

To further customize the system, you can change a specific configuration property further down the hierarchy without having to change the definitions of other configuration properties. As shown in the following illustration, the value of `zPingMonitorIgnore` is changed so that ping monitoring is performed at the `/Devices/Server/Windows` level.


Figure 65: Device Class Hierarchy - Locally Defined Value (Override)



This locally defined value for `zPingMonitorIgnore` *overrides* the value that is set at the root of the hierarchy. No other properties at this level are affected by this local change; they continue to inherit the value that is set at the root.

Configuration properties allow you to configure the system at a very granular level, down to a particular device. For example, in the following illustration, the device named `dev.zenoss.com` has the value of SNMP community set to private. This value overrides the root value of public.

Figure 66: Device Class Hierarchy - Value Set on Device



If you change the SNMP Community value of `dev.zenoss.com` to public, it matches the value that is set at the root, but is still explicitly defined. Only if you remove the locally defined property does it again inherit the value of the property that is set at the root.

Viewing Properties from the Browser Interface

This section further illustrates the characteristics of configuration properties from the browser interface perspective. The following figure shows device configuration properties that are defined at the root level.

To view configuration properties:

- 1 From the Navigation bar, select **Infrastructure**.

The Devices page appears.

- 2 Click **Details**.
- 3 Select **Configuration Properties**.

Figure 67: Defined Device Configuration Properties - Root Level

| Configuration Properties ⚙️ 🔄 Delete Local Copy | | | | |
|---|------------------|---------------------------|--|------|
| Is Local | Category | Name ▲ | Value | Path |
| Yes | CiscoUCS | zCiscoUCSManagerPass word | | / |
| Yes | CiscoUCS | zCiscoUCSManagerPort | 80 | / |
| Yes | CiscoUCS | zCiscoUCSManagerUseSSL | false | / |
| Yes | CiscoUCS | zCiscoUCSManagerUser | | / |
| Yes | Modeler Controls | zCollectorClientTimeout | 180 | / |
| Yes | Modeler Controls | zCollectorDecoding | latin-1 | / |
| Yes | Misc | zCollectorLogChanges | false | / |
| Yes | zencommand | zCommandCommandTimeout | 15 | / |
| Yes | zencommand | zCommandCycleTime | 60 | / |
| Yes | zencommand | zCommandExistanceTest | test -f %s | / |
| Yes | zencommand | zCommandLoginTimeout | 10 | / |
| Yes | zencommand | zCommandLoginTries | 1 | / |
| Yes | zencommand | zCommandPassword | | / |
| Yes | zencommand | zCommandPath | \$ZENHOME/libexec | / |
| Yes | zencommand | zCommandPort | 22 | / |
| Yes | zencommand | zCommandProtocol | ssh | / |
| Yes | zencommand | zCommandSearchPath | | / |
| Yes | zencommand | zCommandUsername | | / |
| Yes | Misc | zDeviceTemplates | Device | / |
| Yes | Misc | zEC2Secret | | / |
| Yes | Modeler Controls | zFileSystemMapIgnoreNames | | / |
| Yes | Modeler Controls | zFileSystemMapIgnoreTypes | other ram virtualMemory removableDisk floppyDisk compactDisk ramDisk f | / |
| Yes | Misc | zFileSystemSizeOffset | 1 | / |

DISPLAYING 1 - 22 OF 127 ROWS

As shown in the figure above, the `zCollectorClientTimeout` configuration property has a default value of 180. In the following figure, the value is set to 170 at the `/Server/Linux` device class, which overrides the default value at this node of the hierarchy.

Figure 68: zCollectorClientTimeout Configuration Property - Local Value Set

| Configuration Properties | | | | |
|--------------------------|------------------|--------------------------|------------|--|
| Is Local | Category | Name | Value | Path |
| | CiscoUCS | zCiscoUCSManagerPassword | | / |
| | CiscoUCS | zCiscoUCSManagerPort | 80 | / |
| | CiscoUCS | zCiscoUCSManagerUseSSL | false | / |
| | CiscoUCS | zCiscoUCSManagerUser | | / |
| Yes | Modeler Controls | zCollectorClientTimeout | | /Server/Linux/devices/t-cent5a-64.zeno |
| | Modeler Controls | zCollectorDecoding | | / |
| | Misc | zCollectorLogChanges | | / |
| | zencommand | zCommandCommandTimeout | 15 | / |
| | zencommand | zCommandCycleTime | 60 | / |
| | zencommand | zCommandExistenceTest | test -f %s | / |
| | zencommand | zCommandLoginTimeout | 10 | / |
| | zencommand | zCommandLoginTries | 1 | / |

DISPLAYING 1 - 13 OF 127 ROWS

To remove the override and once again inherit the value from the root of the hierarchy:

- 1 Select the property in the list.
- 2 Click **Delete Local Copy**.

The Delete Local Property dialog box requests confirmation of the action.

- 3 Click **OK**.

Configuration Property Types

Configuration properties can be one of the following types:

- **String**- Text value that can be ASCII or Latin-1 encoded
- **Integer**- Whole number
- **Float**- Number that can have a decimal value
- **Boolean**- True or False
- **Lines**- List of values separated by a return. The system stores these as an array.
- **Password**- Character-masked password value

Device Configuration Properties

To view and edit device configuration properties at the root level:

- 1 From the Navigation bar, select **Infrastructure**.
- 2 In the tree view, click **Details**.
- 3 Select **Configuration Properties**.

To view and edit device configuration properties that are set at a specific device class, navigate to that class, click **Details**, and then select **Configuration Properties**.

You can also view and edit device configuration at the individual device level. Select the device from the device list, and then select **Configuration Properties** from the left panel.

The following table lists device configuration properties.

Table 4: Device Configuration Properties

| Property Name | Property Type | Description |
|------------------------------|---------------|--|
| zCiscoACEUseSSL | boolean | Whether to use SSL when connecting to ACE XML API. |
| zCiscoNXAPIInterval | int | Value in seconds of the rediscovery interval using the Cisco NX-API protocol. Default value is 300. |
| zCiscoRemodelEventClassKeys | lines | Allows you to modify the list of SNMP traps that will cause Zenoss to schedule an immediate remodeling of the device from which the trap was sent. |
| zCiscoUCSCIMCEventsInterval | int | Event collection interval in seconds. Default is 60. |
| zCiscoUCSCIMCPerfInterval | int | Metric collection interval in seconds. Default is 300. |
| zCiscoUCSManagerPassword | password | UCS Manager username password |
| zCiscoUCSManagerPerfInterval | int | Seconds between UCS Manager statistics collections. Default value is 300. |
| zCiscoUCSManagerPort | int | Port used to connect to the UCS Manager or CIMC XML APIs. Default is 443 and typically should not be changed. |
| zCiscoUCSManagerUseSSL | boolean | Whether to use SSL when connecting to the UCS Manager or CIMC XML APIs. Default is true and typically should not be changed. |
| zCiscoUCSManagerUser | string | UCS Manager username |
| zCollectorClientTimeout | int | Allows you to set the timeout time of the collector client in seconds |
| zCollectorDecoding | string | Converts incoming characters to Unicode. |
| zCollectorLogChanges | boolean | Indicates whether to log changes. |
| zCommandCommandTimeout | float | Specifies the time to wait for a command to complete. |
| zCommandExistenceTest | string | DEPRECATED - No longer used. |
| zCommandLoginTimeout | float | Specifies the time to wait for a login prompt. |
| zCommandLoginTries | int | Sets the number of times to attempt login. |
| zCommandPassword | password | Specifies the password to use when performing command logins and SSH. |
| zCommandPath | string | Sets the default path where ZenCommand plug-ins are installed on the local Resource Manager box (or on a remote box where SSH is used to run the command). |
| zCommandPort | int | Specifies the port to connect to when performing command collection. |
| zCommandProtocol | string | Establishes the protocol to use when performing command collection. Possible values are SSH and telnet. |
| zCommandSearchPath | lines | Sets the path to search for any commands. |
| zCommandUsername | string | Specifies the user name to use when performing command collection and SSH. |

| Property Name | Property Type | Description |
|---------------------------------|---------------|---|
| zControlCenterHost | string | Control Center Host Name. Defaults to device name. |
| zControlCenterModelCycle | int | Control Center Modeling interval. Defaults to 3600s. |
| zControlCenterPassword | password | Control Center username password |
| zControlCenterPerfCycle | int | Control Center Performance Collection interval. Defaults to 300s. |
| zControlCenterPort | int | Port for Control Center. Defaults to HTTPS TCP/443 |
| zControlCenterUser | string | Control Center username |
| zDBInstances | *** | ***instancecredentials. This setting is only relevant when the <code>zenoss.winrm.WinMSSQL</code> modeler plugin is enabled. Multiple instances can be specified to monitor multiple SQL Server instances per server. The default instance is <code>MSSQLSERVER</code> . Fill in the user and password to use SQL authentication. Leave the user and password blank to use Windows authentication. |
| zDeviceTemplates | lines | Sets the templates associated with this device. Linked by name. |
| zEnablePassword | boolean | True or False to specify use of password for Cisco routers. |
| zFileSystemMapIgnoreNames | string | Sets a regular expression of file system names to ignore. |
| zFileSystemMapIgnoreTypes | lines | Do not use. |
| zFileSystemSizeOffset | int | SNMP typically reports the total space available to privileged users. Resource Manager (like the <code>df</code> command) reports capacity based on the space available to non-privileged users. The value of <code>zFileSystemSizeOffset</code> should be the fraction of the total space that is available to non-privileged users. The default reserved value is 5% of total space, so <code>zFileSystemSizeOffset</code> is preset to <code>.95</code> . If the reserved portion is different than 5%, then adjust the value of <code>zFileSystemSizeOffset</code> accordingly. The fraction should be set according to the value $(Used + Avail) / Size$ when the <code>df -PkH</code> command is run at the command line. |
| zHardDiskMapMatch | string | Regular expression that uses the disk ID in the <code>diskstats</code> output to filter disk activity statistics for inclusion in performance monitoring. |
| zIcon | lines | Specifies the icon to represent the device wherever device icon is shown, such as on the network map and device status page. |
| zIdiomPassword | password | IDIOM API password for Cisco IDS/IPS devices. |
| zIdiomUsername | string | IDIOM API username for Cisco IDS/IPS devices. |
| zIfDescription | boolean | Shows the interface description field in the interface list. |
| zInterfaceMapIgnoreDescriptions | string | Filters out interfaces based on description. |

| Property Name | Property Type | Description |
|---------------------------------|---------------|---|
| zInterfaceMapIgnoreNames | string | Filters out interfaces that should not be discovered. If you want to use an expression to define this property, note that only Python regular expressions are valid. |
| zInterfaceMapIgnoreTypes | string | Filters out interface maps that should not be discovered. |
| zIpServiceMapMaxPort | int | Specifies the highest port to scan. The default is 1024. |
| zJBossJmxManagementAuthenticate | boolean | DEPRECATED - No longer used. |
| zJBossJmxManagementPassword | password | JMX password |
| zJBossJmxManagementPort | int | The port number used to gather JMX information |
| zJBossJmxManagementUsername | string | JMX username for authentication. |
| zJmxAuthenticate | boolean | True or False to enable/disable authentication. |
| zJmxManagementPort | int | Port that enables JMX management |
| zJmxPassword | password | JMX username password |
| zJmxUsername | string | JMX username |
| zKeyPath | string | Sets the path to the SSH key for device access. |
| zLDAPBaseDN | string | The Base Distinguished Name for your LDAP server. Typically this is the organization's domain name (for example, dc=foobar,dc=com) |
| zLDAPBindDN | string | The Distinguished Name to use for binding to the LDAP server, if authentication is required. |
| zLDAPBindPassword | string | The password to use for binding to the LDAP server, if authentication is required. |
| zLDMsAutodiscover | boolean | Specify true or false for autodiscovery of LDMs. |
| zLTMVirtualServerIgnoreNames | string | Regular expression that can be used to prevent matching LTM Virtual Servers from being modeled. |
| zLinks | string | Specifies a place to enter any links associated with the device. |
| zLocalInterfaceNames | string | Regular expression that uses interface name to determine whether the IP addresses on an interface should be incorporated into the network map. For instance, a loopback interface "lo" might be excluded. |
| zLocalIpAddresses | string | Specifies IP addresses that should be excluded from the network map (for example. 127.x addresses). If you have addresses that you reuse for connections between clustered machines they might be added here as well. |
| zMaxOIDPerRequest | int | Sets the maximum number of OIDs to be sent by the SNMP collection daemons when querying information. Some devices have small buffers for handling this information so the number should be lowered. |
| zMySQLConnectionString | string | Enables setting the path, type and credentials for the MySQL connections. For example: Path: /server/ |

| Property Name | Property Type | Description |
|------------------------------|---------------|--|
| | | <i>Linux:/devices/localhost.localdomain, Type: multinecredentials, MYSQL connection credentials: User Password Port</i> |
| zMySQLPassword | password | MySQL user password |
| zMySQLPort | string | MySQL connection port |
| zMySQLTimeout | int | MySQL timeout. Default value is 30s. |
| zMySQLUsername | string | MySQL username |
| zNetAppSSL | boolean | Boolean true or false to enable SSL use with the NetApp Monitor. |
| zNmapPortscanOptions | string | Options used on nmap when scanning ports. Used in IpServiceMap. |
| zPingMonitorIgnore | boolean | Whether or not to ping the device. |
| zProdStateThreshold | int | Production state threshold at which Resource Manager will begin to monitor a device. |
| zPropertyMonitorInterval | int | Polling interval of the configured property data sources. System-wide setting. Default is 300s. |
| zPythonClass | string | DO NOT USE |
| zRancidGroup | string | RANCID group attribute. Controls what <code>router.db</code> file the device is written to. Can be set at the device class or device level. Default is <code>router</code> on the <code>/Network/Router/Cisco</code> class. |
| zRancidRoot | string | File system directory where RANCID is installed. It may be NFS mounted from the RANCID server. Default is <code>/opt/rancid</code> |
| zRancidType | string | RANCID type attribute. Controls what device type is written to the <code>router.db</code> file. Can be set at the device class or device level. Default is <code>cisco</code> on the <code>/Network/Router/Cisco</code> class. |
| zRancidUrl | string | Base URL |
| zRouteMapCollectOnlyIndirect | boolean | Only collect routes that are indirectly connected to the device. |
| zRouteMapCollectOnlyLocal | boolean | Only collect local routes. (These usually are manually configured rather than learned through a routing protocol.) |
| zRouteMapMaxRoutes | int | Sets maximum number of routes to collect. Default value is 500. |
| zSnmpAuthPassword | password | The shared private key used for authentication. Must be at least 8 characters long. |
| zSnmpAuthType | string | Use "MD5" or "SHA" signatures to authenticate SNMP requests |
| zSnmpCollectionInterval | int | Defines, in seconds, how often the system collects performance information for each device. |

| Property Name | Property Type | Description |
|----------------------------|---------------|--|
| zSnmCommunities | lines | Array of SNMP community strings that ZenModeler uses when collecting SNMP information. When you set this property, communities are tried in order; the first in the list that is successful is used as zSnmCommunity. If none is successful, then the current value of zSnmCommunity is used. The default value for the entire system is "public." |
| zSnmCommunity | string | Community string to be used when collecting SNMP information. If it is different than what is found by ZenModeler, it will be set on the modeled device. |
| zSnmContext | string | Configures zSNMP context to map logical network entity such as a topology or protocol instance. |
| zSnmDiscoveryPorts | int | List of UDP ports to try when performing SNMP discovery. Defaults to 161 if not set. |
| zSnmEngineId | string | SNMPv3 engine ID for the device. Will be discovered when SNMPv3 is used. |
| zSnmMonitorIgnore | boolean | Whether or not to ignore monitoring SNMP on a device. |
| zSnmPort | int | Port that the SNMP agent listens on. |
| zSnmPrivPassword | password | The shared private key used for encrypting SNMP requests. Must be at least 8 characters long. |
| zSnmPrivType | string | "DES" or "AES" cryptographic algorithms. |
| zSnmSecurityName | string | The Security Name (user) to use when making SNMPv3 requests. |
| zSnmTimeout | float | Timeout time in seconds for an SNMP request |
| zSnmTries | int | Amount of tries to collect SNMP data |
| zSnmVer | string | SNMP version used. Valid values are v2c, v3 |
| zSshConcurrentSessions | int | Maximum number of sessions supported by the remote device's MAX_SESSIONS parameter. Common values for AIX are 2 or 10. |
| zStatusConnectTimeout | float | The amount of time that the zenstatus daemon should wait before marking an IP service down. |
| zSugarCRMBase | - | DEPRECATED - No longer used. |
| zSugarCRMPassword | password | Password for the zSugarCRMUsername user. |
| zSugarCRMTestAccount | - | DEPRECATED - No longer used. |
| zSugarCRMUsername | string | Username allowed to log in to the Sugar CRM server. |
| zSysedgeDiskMapIgnoreNames | string | Regular expression used by zenoss.snmp.SysedgeDiskMap modeler plugin. Disks with matching names will not be modeled. |
| zTelnetEnable | boolean | When logging into a Cisco device issue the enable command to enable access during command collection. |

| Property Name | Property Type | Description |
|----------------------------------|---------------|---|
| zTelnetEnableRegex | string | Regular expression to match the enable prompt. |
| zTelnetLoginRegex | string | Regular expression to match the login prompt. |
| zTelnetPasswordRegex | string | Regular expression to match the password prompt. |
| zTelnetPromptTimeout | float | Time to wait for the telnet prompt to return. |
| zTelnetSuccessRegexList | lines | List of regular expressions to match the command prompt. |
| zTelnetTermLength | boolean | On a Cisco device, set term length to Zero. |
| zTomcatJ2EEApplicationName | string | Used to construct MBean names for a specific application deployed on Tomcat, typically used for JSP and Servlet statistics. |
| zTomcatJ2EEServerName | string | Used to construct MBean names for a specific application deployed on Tomcat, typically used for JSP and Servlet statistics. |
| zTomcatJmxManagementAuthenticate | - | DEPRECATED - No longer used. |
| zTomcatJmxManagementPassword | password | JMX password. |
| zTomcatJmxManagementPort | int | The port number used to gather JMX information. |
| zTomcatJmxManagementUsername | string | JMX username for authentication. |
| zTomcatListenHost | string | The hostname on which Tomcat is listening for web requests. This is used to construct MBean names |
| zTomcatListenPort | string | The Tomcat connector, which is a port and protocol (http, jk...) that Tomcat is listening on. This is used to construct MBean names that monitor bytes, error and requests on that connector. |
| zTomcatServletName | string | Specific Servlet name to monitor. |
| zTomcatServletUri | string | URI of Servlet to monitor. |
| zTomcatWebAppUri | string | URI path for a Tomcat web application. Used to construct MBean names. |
| zVCloudPassword | password | Password for the zVCloudUsername user. |
| zVCloudPort | int | Value of the cell port number |
| zVCloudUsername | string | Username for the cell in the form username@organization. For example, if you want to define the cell administrator, enter administrator@system. |
| zVSphereEndpointHost | string | vSphere host name |
| zVSphereEndpointPassword | password | vSphere username password |
| zVSphereEndpointUseSsl | boolean | vSphere boolean true or false for SSL use |
| zVSphereEndpointUser | string | vSphere username |
| zVSphereModelCache | lines | vSphere model cache |

| Property Name | Property Type | Description |
|------------------------------------|---------------|--|
| zVSphereModelIgnore | lines | vSphere model ignore |
| zVSpherePerfDelayCollectionMinutes | int | Value of how long to lag performance data collection. Default value is 0. |
| zVSpherePerfMaxAgeMinutes | int | Default value is 28 minutes. |
| zVSpherePerfParallelQueries | int | Default value is 6. |
| zVSpherePerfQueryChunkSize | int | Value of how many performance requests to make at a time. Default value is 250. |
| zVSpherePerfQueryRaw20 | boolean | Default value is <code>true</code> . |
| zVSpherePerfQueryTimeout | int | Default value is 200. |
| zVSpherePerfQueryVcChunkSize | int | Default value is 64. |
| zVSpherePerfQueryVcRaw20 | boolean | Default value is <code>false</code> . |
| zVSpherePerfRecoveryMinutes | int | Default value is 240 minutes. |
| zVSpherePerfWindowSize | int | DEPRECATED - No longer used. |
| zWBEMPassword | password | WBEM password |
| zWBEMPort | int | Value of the WBEM port number. Default value is 5989. |
| zWBEMUseSSL | boolean | True or false value to use SSL. Default value is <code>true</code> . |
| zWBEMUsername | string | WBEM username |
| zWebLogicJmxManagementAuthenticate | - | DEPRECATED - No longer used. |
| zWebLogicJmxManagementPassword | password | JMX password |
| zWebLogicJmxManagementPort | int | The port number used to gather JMX information |
| zWebLogicJmxManagementUsername | string | JMX username for authentication |
| zWebTxAgent | string | Default value is <code>zenWebTx/1.0</code> . |
| zWebTxPassword | password | WebTx password |
| zWebTxRealm | string | WebTx realm |
| zWebTxUser | string | WebTx user |
| zWebsphereAuthRealm | string | Used for HTTP basic authentication. This field is not required, and is empty by default. |
| zWebsphereNode | string | Used by the provided template to build the queries for the data to collect. You must supply a value for this field. For example: <code>server.A</code> |
| zWebspherePassword | password | Used for HTTP basic authentication. This field is not required, and is empty by default. |
| zWebsphereServer | string | Used by the provided template to build the xpath queries for the data to collect. You must supply a value for this field. For example: <code>server.AB</code> , There is no default value. |

| Property Name | Property Type | Description |
|---------------------|---------------|---|
| zWebsphereURLPath | string | Path to the PMI servlet on a WebSphere instance. The default value is the default path on a WebSphere installation: wasPerTool/servlet/perfservlet |
| zWebsphereUser | string | Used for HTTP basic authentication. This field is not required, and is empty by default. |
| zWinKDC | string | IP address or Fully Qualified Domain Name of a valid Windows domain controller. Must be set if domain authentication is used. |
| zWinKeyTabFilePath | string | This property is currently used and reserved for future use when keytab files are supported. |
| zWinPerfmonInterval | int | Interval, in seconds, at which Windows Perfmon datapoints will be collected. Default value is 300. It is possible to override the collection interval for individual counters. |
| zWinRMPassword | password | Password for the user defined by zWinRMUser |
| zWinRMPort | int | The port on which the Windows server is listening for WinRM or WS-Management connections. Default value is 5985. It is uncommon for this to be configured as anything else. |
| zWinRMServerName | string | <p>This property should only be used in conjunction with domain authentication when the DNS PTR record for a monitored server's managed IP address does not resolve to the name by which the server is known in Active Directory. For example, if myserver1 is known as myserver1.ad.example.com by Active Directory and is being managed by IP address 192.51.100.21, but 192.51.100.21 resolves to www.example.com, you will have to set zWinRMServerName to myserver1.ad.example.com for domain authentication to work.</p> <p>If many Windows servers in your environment don't have DNS PTR records that match Active Directory, it is recommended that you set the name of the Zenoss device to be the fully-qualified Active Directory name and set zWinRMServerName to <code>\${here/titleOrId}</code> at the /Server/Microsoft/Windows device class. This avoids the necessity of setting zWinRMServerName on every device.</p> <p>It is recommended to leave zWinRMServerName blank if local authentication is used, or DNS PTR records match Active Directory. This allows Zenoss to not rely on DNS resolution while monitoring, and avoids the overhead of configuring zWinRMServerName.</p> |

| Property Name | Property Type | Description |
|------------------|---------------|---|
| zWinRMUser | string | The syntax used for zWinRMUser controls whether Zenoss will attempt Windows local authentication or domain (kerberos) authentication. If the value of zWinRMUser is <i>username</i> , local Windows authentication will be used. If zWinRMUser is <i>username@example.com</i> , domain authentication will be used. The zWinKDC and potentially the zWinRMServerName properties become important. |
| zWinScheme | string | This must be set to either http or https. Default value is http. |
| zWinTrustedKDC | string | Windows Trusted KDC |
| zWinTrustedRealm | string | Windows Trusted Realm |

Event Configuration Properties

To view and edit event configuration properties at the root level:

- 1 From the Navigation menu, select **Events**, and then select **Event Classes**.
- 2 From the drop-down list, select **Configuration Properties**.

To view and override event configuration properties for a specific event class:

- 1 Navigate to that class, and then select **Configuration Properties**.
- 2 Double-click the configuration property that you want to change. The **Edit Config Property** dialog box appears.
- 3 Make your changes and click **Submit**. The presence of **Yes** in the **Is Local** column indicates an overriding value. If you want to later remove the override, select the configuration property and click **Delete Local Copy**.

The following table lists event configuration properties.

Table 5: Event Configuration Properties

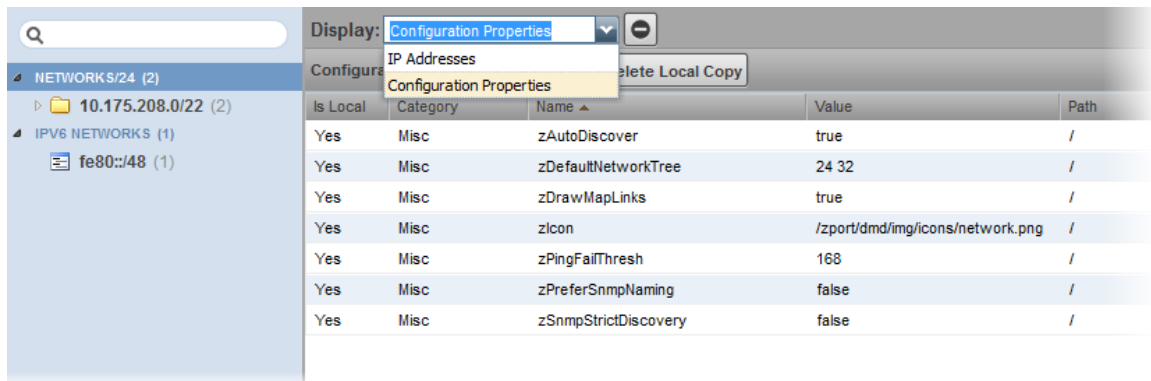
| Property Name | Property Type | Description |
|--------------------------|---------------|--|
| zEventAction | string | Specifies the database table in which an event will be stored. Possible values are: status, history and drop. Default is status, meaning the event will be an “active” event. History sends the event directly to the history table. Drop tells the system to discard the event. |
| zEventClearClasses | lines | Lists classes that a clear event should clear (in addition to its own class). |
| zEventSeverity | int | Overrides the severity value of events from this class. Possible values are 0-Critical, 1-Error, 2-Warning, 3-Info, 4-Debug, 5-Clear, and -1-Default. |
| zFlappingIntervalSeconds | int | Defines the time interval to check for event flapping (changing severity level repeatedly). Default value is 3600 seconds. |

| Property Name | Property Type | Description |
|--------------------|---------------|---|
| zFlappingSeverity | int | Drop-down list to set the severity to check for event flapping. If the severity level on an event changes from this value a certain number of times (zFlappingThreshold) within a certain time range (zFlappingIntervalSeconds) then an event flapping event is generated. Possible values include: 5-Critical, 4-Error, 3-Warning, 2-Info, 1-Debug, and 0-Clear. |
| zFlappingThreshold | int | Number of times an event severity must flap within an interval. One of the parameters to define in order to generate event flapping events. |

Network Configuration Properties

To view and edit network configuration properties, from the Navigation menu, select **Infrastructure > Networks**. The Networks page appears, listing networks by IP address.

Figure 69: Networks



You can view and change network configuration property values and inheritance selections. From the **Display** menu, select **Configuration Properties**.

The following table lists network configuration properties.

Table 6: Network Configuration Properties

| Property Name | Property Type | Description |
|---------------------|---------------|---|
| zAutoDiscover | boolean | Specifies whether zendisc should perform auto-discovery on this network. (When performing network discovery, this property specifies whether the system should discover devices and subnetworks on the network.) |
| zDefaultNetworkTree | lines | A network subnet is automatically created for each modeled device, based on that device's subnet mask setting. To create higher-level subnets automatically from the discovery and modeling processes, add the specific subnet mask breakpoints. For example: 8, 16. If you then model a device with, for example, an IP address of 192.0.2.0, and a subnet mask of 255.255.255.0 (corresponding to a /24 subnet), device |

| Property Name | Property Type | Description |
|----------------------|---------------|---|
| | | discovery will create a 192.0.0.0/8 network containing 192.0.2.0/16, containing 192.0.2.0/24, containing your device. |
| zDrawMapLinks | boolean | Calculating network links "on the fly" is resource-intensive. If you have a large number of devices that have been assigned locations, then drawing those map links may take a long time. You can use this property to prevent the system from drawing links for specific networks (for example, a local network comprising many devices that you know does not span multiple locations). |
| zIcon | string | Use to specify device icons that appear on the device status page, Dashboard, and network map. |
| zPingFailThresh | int | Specifies the number of pings sent without being returned before zendisc removes the device. |
| zPreferSnmpNaming | boolean | Specifies that when network discovery occurs, it uses the device name comes from SNMP rather than reverse DNS. |
| zSnmpStrictDiscovery | boolean | Specifies that if SNMP does not exist on the device during network discovery, ignore the device. |

6

Monitoring Templates

The system stores performance configuration data in *templates*. Templates contain other objects that define where and how to obtain performance data, thresholds for that data, and data graphs.

You can define a template anywhere in the device class hierarchy, or on an individual device.

Templates are divided among three types:

- Device
- Component
- Interface

Creating Templates

You can create a template by overriding an existing template. To override a template:

- 1 Navigate to the template you want to copy.
- 2 From the Action menu, select **Copy/Override Template**. The Copy/Override dialog box appears.
- 3 Select the bound template to override, and then click **Submit**. The copied template appears in the list of templates as locally defined.

Renaming Templates

To rename an existing template:

- 1 Select **Advanced > Monitoring Templates**.
- 2 Expand the organizer containing the template to be renamed, and then the class containing the template.
- 3 From the Action menu, select **View and Edit Details**. The Edit Template Details dialog appears.
- 4 Enter a new name in the Name field.
- 5 Click **Submit**.

Template Binding

The determination of which templates apply to what objects is called *binding*. Templates are bound in different ways, depending on the objects to which they are bound.

Device Templates

Device templates are applied to devices, one to each device. The system employs a single rule to bind device templates to devices: the value of the zDeviceTemplates property. For most device classes, this is "Device."

Common device templates are:

- Device
- MySQL
- Apache
- Active Directory
- MExchangeIS
- MSSQLServer
- IIS


For the Server/Linux/MySQL device class, the zDeviceTemplates property might contain, for example, "Device" and "MySQL." The system would collect CPU and memory information by using the Device template, and MySQL-specific metrics by using the MySQL template.

Binding Templates

To bind a device template to a device class or device:

- 1 From the devices list, select a device class or device.
- 2 On the Overview page, select **Bind Templates** from the Action menu. The Bind Templates dialog appears.

Figure 70: Bind Templates



- 3 Move templates between the Available and Selected lists using the arrows.
- 4 Click **Save**.

Resetting Bindings

Resetting template bindings removes all locally bound templates and uses the default template values. To reset bindings for a selected device or device class:

- 1 Select **Reset Bindings** from the Action menu.

The Reset Template Bindings dialog appears.

- 2 Click **Reset Bindings** to confirm the action.

Component Templates

Component templates are named exactly according to the name of the underlying class that represents a component. For example, the FileSystem template is applied to file systems. Component templates can be applied multiple times to each device, depending on how many of the device's components match the template. Configuration properties do not control the application of component templates.

Note Component templates should not be manually bound.

Common component templates are:

- FileSystem, HardDisk, IPService, OSProcess, WinService
- Fan, PowerSupply, TemperatureSensor
- LTMVirtualServer, VPNTunnel

Interface Templates

Interface templates are applied to network interfaces by using a special type of binding. Instead of using the name of the underlying class, the system looks for a template with the same name as the interface type. You can find this type in the details information for any network interface.

If Resource Manager cannot locate a template that matches the interface type, then it uses the ethernetCsmacd template.

Example: Defining Templates in the Device Hierarchy

You add a new device at /Devices/Server/Linux named Example1Server. You have not edited the value of its zDeviceTemplates property, so it inherits the value of "Device" from the root device class (/Devices). Resource Manager looks to see if there is a template named Device defined on Example1Server itself. There is not, so it checks /Devices/Server/Linux. There is a template named Device defined for that device class, so that template is used for Example1Server. (There also is a template named Device defined at the root level (/Devices), but the system does not use this one because the template at /Devices/Server/Linux overrides it.)

Example: Applying Templates to Multiple Areas in the Device Hierarchy

You want to perform specific monitoring of servers running a certain Web application, but those servers are spread across several different device classes. You create a template at /Devices called WebApplication with the appropriate data sources, thresholds and graphs. You then append the name "WebApplication" to the zDeviceTemplates configuration property for the devices classes, the individual devices running this Web application, or both.

7

Basic Monitoring

Availability Monitoring

The availability monitoring system provides active testing of the IT infrastructure, including:

- Devices
- Network
- Processes
- Services

Availability monitoring is facilitated by:

- **ZenPing**- The system's Layer-3 aware, topology-monitoring daemon. ZenPing performs high-performance, asynchronous testing of ICMP status. The most important element of this daemon is that Resource Manager has built a complete model of your routing system. If there are gaps in the routing model, the power of ZenPing's topology monitoring will not be available. If there are gaps, this issue can be seen in the `zenping.log` file.
- **ZenStatus**- Performs active TCP connection testing of remote daemons.

Zenmodeler discovers the routes to each device in the network. The system tries not to use Internet routing tables, relying instead on Zenmodeler to discover the relationships and create its own network map.

If any known route is broken, then only one ping event is generated by the outage. Any additional outages will only flag that device and the next time a ping sweep occurs the errors beyond the known router will not occur.

This monitoring model breaks down if the routers do not share their routing tables and interface information.

Zenping Correlation

By default, zenping attempts to correlate all zenping instances on the same collector host. For example, if there are two collectors deployed on Host A, they would participate in the same correlation activity. If zenping instances are on different physical hosts, they must share the same redis instance. In this case, the setting `--redis-url` which controls what redis instance zenping is talking to needs to be set to point to the same redis instance.

There are other correlation settings to use on the command line that you need to be aware of:

- `--disable-correlator`: This will set the ping correlation value to `False` and will turn off correlation.
- `--correlation-backend`: Default value is `distributed`, which is the default behavior when using the redis instance. If you set this to `simple`, no communication with redis is available and correlation will only be performed within the same zenping instance.

- `--connected-ip-suppress`: Defaults to disabled. This is used when zenping does not or cannot monitor all the interfaces of routers along a traceroute path. In this case, zenping can use the model information to infer if an IP along the traceroute path belongs to a device (like a router). It will use this information to determine if the root cause for a ping down is the "connected" device.
- `--delay-count`: Defaults to 0. This means that as soon as a ping down is detected an event is sent. When using distributed ping correlation, zenping instances are not in sync. `--delay-count` can be set to 1 or higher in order to allow zenping instances to share the ping state before sending events. It can also be used to avoid sending events after N ping downs. For example, you may only want to send an event after 3 ping downs.

Controlling Ping Cycle Time

Follow these steps to modify the ping cycle time:

- 1 Open Control Center and click your instance of Resource Manager to open the application overview page.
- 2 In the Services section, click **zenping** to open the zenping Overview page.
- 3 In the Configuration Files section, click **Edit** next to the `/opt/zenoss/etc/zenping.conf` file. The Edit Configuration window appears.
- 4 Uncomment the `zenhubpinginterval` line and edit the default value of 30 to your desired ping cycle time.
- 5 Click **Save**.

Using the Predefined /Ping Device Class

The `/Ping` device class is a configuration for devices that you want to monitor only for availability. The system does not gather performance data for devices placed in this class. You can use the `/Ping` device class as a reference for your own configuration; or, if you have a device that you want to monitor solely for availability, you can place it under this class.

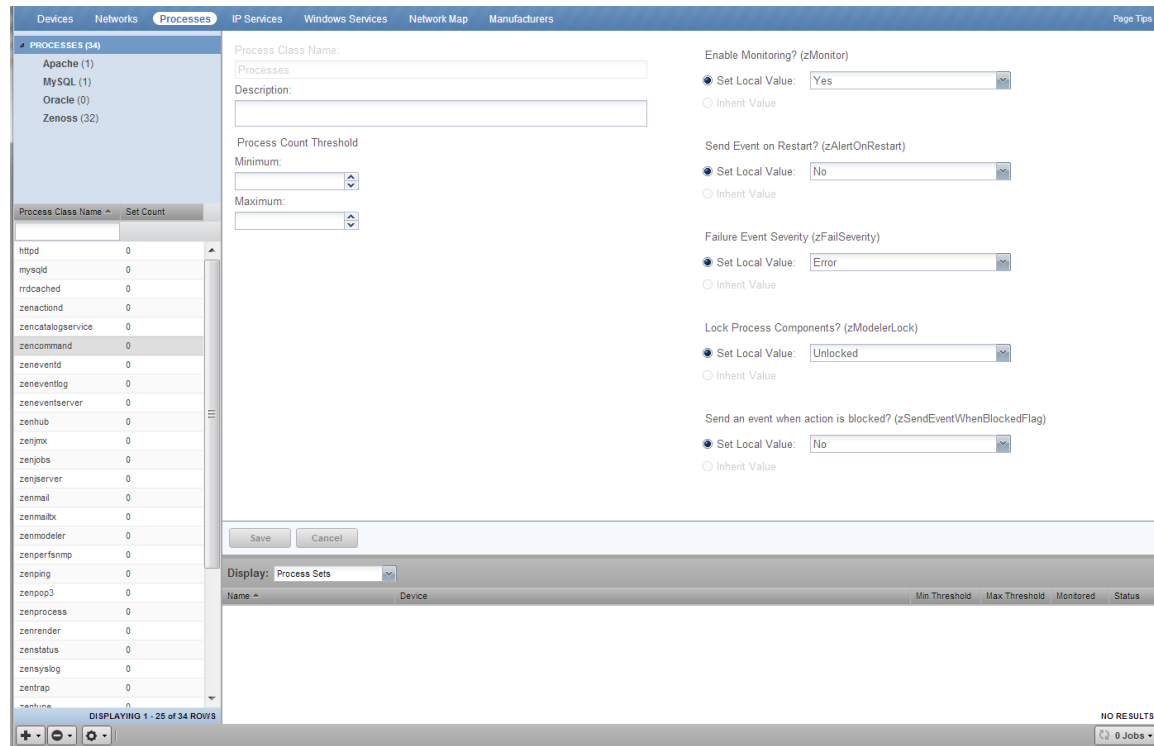
Monitoring Processes

Resource Manager provides process availability monitoring for hosts that support SNMP or SSH access. Process monitoring features include:

- Process classes, defined by Python regular expressions. Classes may generate one or more process sets, each containing one or more process instances.
- Process sets may include process instances running on multiple hosts. This captures related or redundant processes, enabling a more wholistic view of data center services.
- Process set names, to replace the often-cryptic names of process instances with descriptive labels.
- Process set locking, to maintain continuity of data collection if the members of a given process set are not running during modeling.
- A testing dialog, to discover and refine the sets a class generates.

Use the Processes page (**Infrastructure > Processes**) to create and manage process classes and process sets.

Figure 71: Processes page



The tree view shows process class organizers (at the top) and the list of process classes in each organizer (the rest of the view). You may filter the list with the active search field, at the top of the list.

Example: Creating a process class

This section provides an example of using process availability monitoring to create a new process class, for database processes. The database runs on a Linux host, and the following output is a partial list of the results of the `ps axho args` command on the database host. Process monitoring uses the output of that command (or its equivalent) as input for regular expression matching.

```
ora_pmon_orcl ora_psp0_orcl ora_vktm_orcl ora_gen0_orcl ora_diag_orcl
ora_dbrm_orcl ora_dia0_orcl ora_mman_orcl ora_dbw0_orcl ora_lgwr_orcl
ora_ckpt_orcl ora_smon_orcl ora_reco_orcl ora_mmon_orcl ora_mmln_orcl
ora_d000_orcl ora_s000_orcl ora_s001_orcl ora_s002_orcl ora_s003_orcl
ora_s004_orcl ora_s005_orcl ora_s006_orcl ora_s007_orcl ora_s008_orcl
ora_s009_orcl ora_p000_orcl ora_p001_orcl ora_p002_orcl ora_p003_orcl
ora_p004_orcl ora_qmnc_orcl ora_n000_orcl ora_l000_orcl ora_l001_orcl
ora_l002_orcl ora_l003_orcl
```

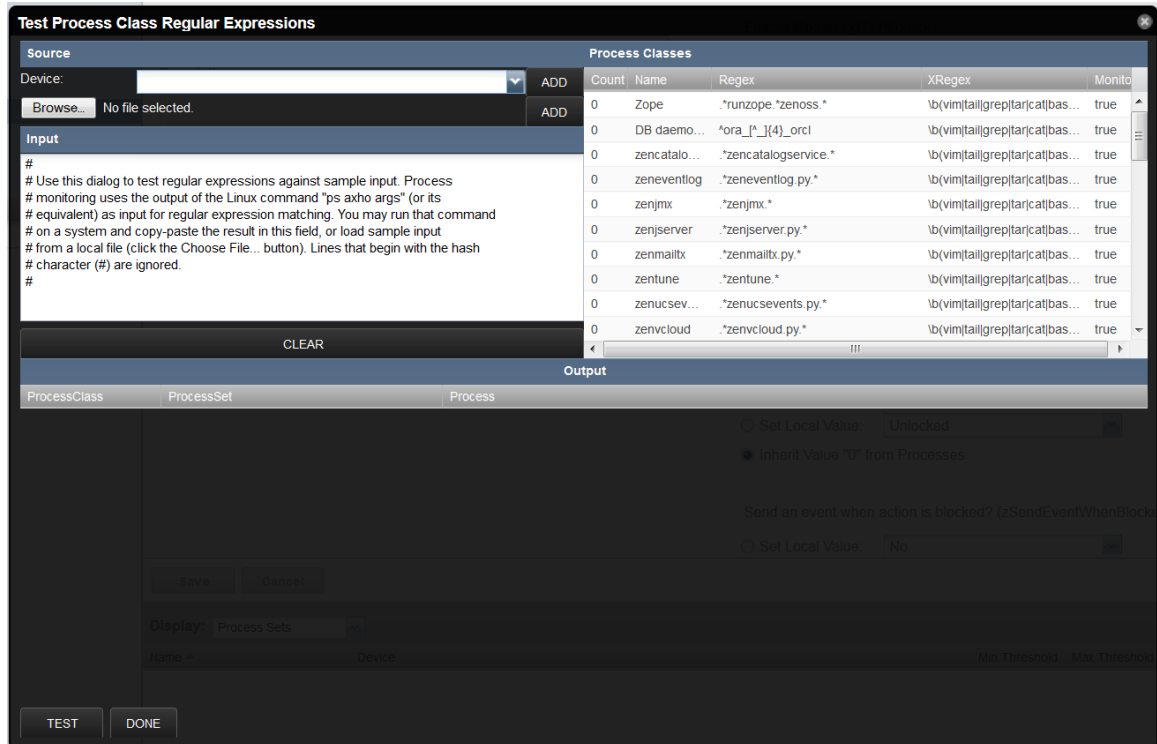
The following subsections provide procedures for creating a process class that captures a selection of the preceding process instances in process sets.

Test existing process classes

The existing process classes may already capture the process sets you want. Follow these steps to test the existing process classes.

- 1 Log in to the Resource Manager user interface.
- 2 Navigate to **Infrastructure > Processes**.
- 3 In the lower-left corner of the tree view, click the Action menu, and select **Test All Process Classes Regular Expressions**.

Figure 72: Test all process classes dialog



The top-right portion of the dialog displays all of the process classes, in the order in which they are evaluated.

- 4 Select the process names in the previous section, and copy them into a paste buffer.
- 5 In the **Test Process Class Regular Expressions** dialog, select all of the existing text in the **Input** area, and then paste the process names from the buffer. Alternatively, you may paste the process names into an empty file, save the file on the system from which your browser is launched, and then use the **Choose File** button to insert the process names.
- 6 At the bottom-left corner of the dialog, click **Test**.
- 7 If any process sets are created, they are displayed in the list area, above the **Test** button.

Create an organizer

Complete the previous section, and then follow these steps to create a process class organizer.

- 1 Log in to the Resource Manager user interface.
- 2 Navigate to **Infrastructure > Processes**.
- 3 In the lower-left corner of the tree view, click the **Add** menu, and select **Add Process Class Organizer**.
- 4 In the **Add Process Class Organizer** dialog, enter `Database`, and then click **Submit**.

Create a process class

Complete the previous section, and then follow these steps to create a process class.

- 1 At the top of the tree view, double-click **Processes**, the root organizer to open it.
- 2 Select **Database**.
- 3 In the lower-left corner of the tree view, click the **Add** menu, and select **Add Process Class**.
- 4 In the **Add Process Class** dialog, enter `DB daemons`, and then click **Submit**.
- 5 At the top of the tree view, double-click **Processes** to open it, and then select **Database**.

Define the regular expression series of a process class

Complete the previous section, and then follow these steps to create the series of regular expressions that define a process class, and generate process sets.

- 1 In the list area of the tree view, select **DB daemons**.

Figure 73: Process class definition page

Process Class Name:

Description:

Matching Rules (performance metrics will start over if changed)

Include processes like:

Exclude processes like:

Replace command line text:

With:

Process Count Threshold

Minimum:

Maximum:

Enable Monitoring? (zMonitor)

Set Local Value:

Inherit Value "Yes" from Processes

Send Event on Restart? (zAlertOnRestart)

Set Local Value:

Inherit Value "No" from Processes

Failure Event Severity (zFailSeverity)

Set Local Value:

Inherit Value "Error" from Processes

Lock Process Components? (zModelerLock)

Set Local Value:

Inherit Value "0" from Processes

Send an event when action is blocked? (zSendEventWhenBlockedFlag)

Set Local Value:

Inherit Value "No" from Processes

- 2 In the **Description** field, enter Database daemons.
- 3 In the **Include processes like** field, replace DB daemons with a *Python regular expression* that selects the database processes. For example, `ora_[^_]{4}_orcl`. The example regular expression selects all of the processes in the sample process instance list.
- 4 In the **Exclude processes like** field, enter a regular expression to remove processes from the results of the preceding regular expression. The default entry excludes common user commands. The default entry does not exclude any of the processes in the sample process instance list.
- 5 The next two fields, **Replace command line text** and **With**, work together to simplify the names of process sets.
 - In the **Replace command line text** field, enter a regular expression containing one or more pattern groups.
 - In the **With** field, enter replacement text, along with the sequence number of one or more of the pattern groups defined in the previous field.

For example, to create four pattern groups for database processes, enter the following regular expression in the **Replace command line text** field:

```
^(ora_)([a-z])(.{4})(orcl)
```

To use two of the pattern groups in the replacement text, enter the following text in the **With** field:

```
DB [\4] daemons starting with [\2]
```

Each unique replacement generated by the combination of the text plus the inserted pattern sequences becomes a process set. In the case of this example, pattern group 4 does not vary, but pattern group 2 does. So the

number of process sets generated by this class will equal the number of unique alphabetic characters found in the first position after the first underscore.

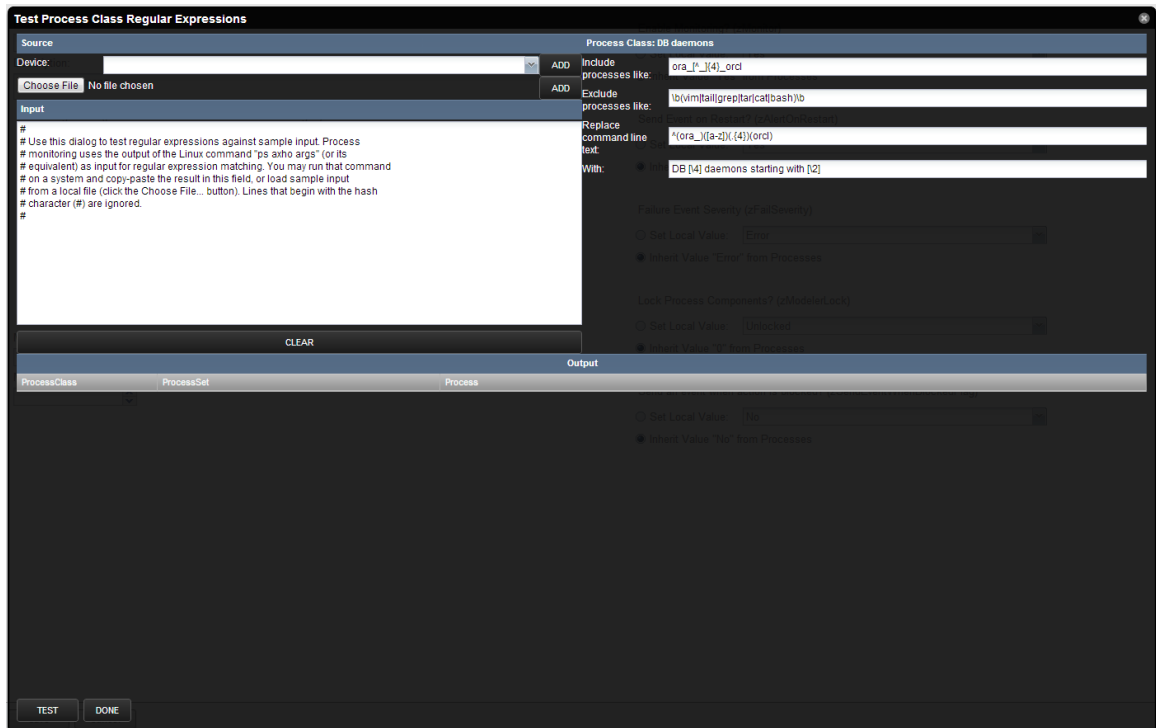
- 6 Click **Save**.

Test a process class

Complete the previous section, and then follow these steps to test a single process class.

- 1 In the lower-left corner of the tree view, click the Action menu, and select **Test Process Class Regular Expressions**.

Figure 74: Test process class dialog



The top-right portion of the dialog displays the regular expression series that defines this process class.

- 2 Select the process names in the previous section, and copy them into a paste buffer.
- 3 In the **Input** area, select all of the existing text, and then paste the process names from the buffer.
- 4 At the bottom-left corner of the dialog, click **Test**.
- 5 The **Output** area displays each individual match, along with the count of unique process sets. You may refine the regular expressions and retest as often as you like.
- 6 Click **Done**. Changes made to regular expressions in this dialog are copied to the process class definition page. However, the changes are not saved until you click the **Save** button on that page.

Test and review the process class sequence

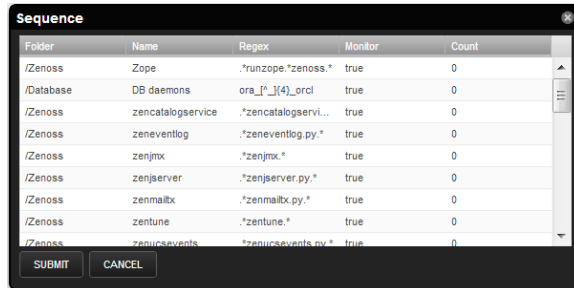
The order in which process classes are evaluated is significant. During modeling, each time a process matches a class, the matching process is put into a process set, and then removed from the list of processes that are passed on to the next class in the sequence. New process classes are inserted into the process class sequence automatically, and may not be in the appropriate position in the sequence.

Complete the previous section, and then follow these steps to test and review the process class sequence.

- 1 In the lower-left corner of the tree view, click the Action menu, and select **Test All Process Classes Regular Expressions**.
- 2 Select the process names in the previous section, and copy them into a paste buffer.

- 3 In the **Input** area, select all of the existing text, and then paste the process names from the buffer.
- 4 At the bottom-left corner of the dialog, click **Test**.
- 5 The number of processes matched and process sets created in this test should be identical to the results of testing the process class alone. If they are not, follow these steps to adjust the process class sequence.
 - a In the **Test Process Class Regular Expressions** dialog, click **Done**.
 - b From the Action menu, select **Change Sequence**.

Figure 75: Sequence dialog



- c Scroll through the list of process classes, and then select the class to move.
 - d Drag the class to an earlier (higher) location in the sequence.
 - e Click **Submit**.
- 6 Re-open the **Test Process Class Regular Expressions** dialog, and re-test the sequence.

Test the process class on a host

Process sets are created during modeling. To test a process class, choose a device host that is configured for SNMP or SSH access, and model it manually.

Note For more information about device support for process monitoring, refer to the release notes.

Complete the previous section, and then follow these steps to test the process class on a host.

- 1 Navigate to **Infrastructure > Devices**.
- 2 Select a host that is configured for SNMP or SSH access, and is running process that match the new class. For example, the list of processes used in this section is collected from a VirtualBox virtual appliance downloaded from the [Oracle Technology Network](#).
- 3 Open the host's **Overview** page. From the Action menu, select **Model Device**. When modeling completes, the **OS Processes** section of the tree view is updated to include the new process sets.
- 4 Navigate to **Infrastructure > Processes**.
- 5 In the tree view, select the **DB daemons** class.

Figure 76: Process class page with process sets

Process Class Name:

Description:

Matching Rules (performance metrics will start over if changed)

Include processes like:

Exclude processes like:

Replace command line text:

With:

Process Count Threshold

Minimum:

Maximum:

Enable Monitoring? (zMonitor)

Set Local Value:

Inherit Value "Yes" from Processes

Send Event on Restart? (zAlertOnRestart)

Set Local Value:

Inherit Value "No" from Processes

Failure Event Severity (zFailSeverity)

Set Local Value:

Inherit Value "Error" from Processes

Lock Process Components? (zModelerLock)

Set Local Value:

Inherit Value "0" from Processes

Send an event when action is blocked? (zSendEventWhenBlockedFlag)

Set Local Value:

Inherit Value "No" from Processes

Display:

| Name | Device | Min Threshold | Max Threshold | Monitored | Status |
|------------------------------------|---------------|---------------|---------------|-----------|--------|
| DB [orc] daemons starting with [c] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [d] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [g] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [i] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [l] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [m] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [n] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [p] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [q] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [r] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [s] | 10.87.110.103 | | | true | Up |
| DB [orc] daemons starting with [v] | 10.87.110.103 | | | true | Up |

DISPLAYING 1 - 13 of 13 ROWS

The process sets found on the host are displayed in the list at the bottom of the page.

Process class options

The process class page includes the options described in the following sections.

Process Count Threshold

You may set minimum and maximum values for the number of process instances included in a process set. The threshold values apply to all of the process sets generated by a class. The minimum and maximum values are inclusive. That is, if the minimum is 3 and the maximum is 5, then 3, 4, and 5 are all valid process instance counts.

You may define a threshold as an exclusive range. If the minimum is 5 and the maximum is 3, then 4 is an invalid process instance count.

Monitoring Options

- **Enable Monitoring (zMonitor)**

To disable monitoring for all process sets generated by this class, set the local value to No.

- **Send Event on Restart (zAlertOnRestart)**

To send an event when monitoring restarts, set the local value to Yes.

- **Failure Event Severity (zFailSeverity)**

To specify a non-default event severity for the failure of process sets generated by this class, set a local value.

Process Set Locking

Process sets are generated at modeling time. Since modeling recurs regularly, a given modeling run may result in a missing process set, due to a transient absence of one or more process instances. To prevent this from happening, set the local value of the **Lock Process Components? (zModelerLock)** field to one of the following options.

- **Lock from Deletes**

Prevent deletion of process sets generated by this process class if modeling returns empty sets.

- **Lock from Updates**

Prevent updates to process sets generated by this process class if modeling returns new process sets.

The final option, **Send an event when action is blocked? (zSendEventWhenBlockedFlag)**, is used only when a process set is locked. If you lock the process sets of a class, you may set the local value of this field to Yes, and an event will be created when a process set would have been either deleted or updated during modeling.

Monitoring IP Services

The IP Services page (**Infrastructure > IP Services**) lets you manage and monitor IP services that are running on your network.

Figure 77: IP Services

The screenshot shows the IP Services configuration page. On the left, a tree view displays the hierarchy: IPSERVICE (3756) > Privileged (684) > Registered (3069). Below this is a table listing various services with their names, ports, and counts. The details area on the right allows for configuration of the selected service (IpService). It includes fields for Name and Description, and radio button options for monitoring and failure severity. At the bottom, there are Save and Cancel buttons, a Display dropdown menu, and a table for monitoring service instances.

| Name | Port | Count |
|---------------|------|-------|
| 1ci-smcs | 3091 | 0 |
| 3com-amp3 | 629 | 0 |
| 3com-net-mgmt | 2391 | 0 |
| 3Com-nsd | 1742 | 0 |
| 3com-tsmux | 106 | 0 |
| 3com-webview | 2339 | 0 |
| 3com-faxrpc | 3446 | 0 |
| 3d-nfsd | 2323 | 0 |
| 3ds-lm | 1538 | 0 |
| 3i-l1 | 1511 | 0 |
| 3m-image-lm | 1550 | 0 |

The tree view lists all monitored IP services. Filter this list by using the active search area at the top of the view.

The details area shows:

- Service class description
- TCP port
- Associated service keys

To add or change details for a service class, enter or change information, and then click **Save**.

The lower section of the page lists currently running services in this class (by device), and shows their monitoring status. You can also display Configuration Properties by selecting that from the drop-down list.

Enabling IP Service Monitoring

You can choose to monitor:

- Individual services
- Service classes

When monitoring a service class, you can choose not to monitor one or more individual services in the class. For example, the SMTP service class is monitored by default, but may not be a critical service on some devices. In this case, you can disable its monitoring on those devices.

Note If a service is configured to listen only on local host (127.0.0.1), then it is not monitored by default.

Note When adding a new IP service that uses a port value higher than 1024, you need to increase the value of `zIpServiceMapMaxPort` to a number higher than the port you are monitoring.

To enable monitoring for a service class or service:

- 1 In the tree view, select the service class or service to monitor.
- 2 Make one or more selections:
 - **Enable Monitoring (zMonitor)** - By default, Inherit Value is selected for all services below the IPService node. When selected, the service class or service will inherit monitoring choices from its parent. If you want to individually enable monitoring choices, select the Set Local Value option, and then select a value.
 - **Failure Event Severity (zFailSeverity)** - By default, Inherit Value is selected for all services below the IPService node. When selected, the service class or service will inherit severity level choices from its parent. If you want to individually select severity levels, select the Set Local Value option, and then select a value.
- 3 Click **Save** to save your choices.

Using the Predefined /Server/Scan Device Class

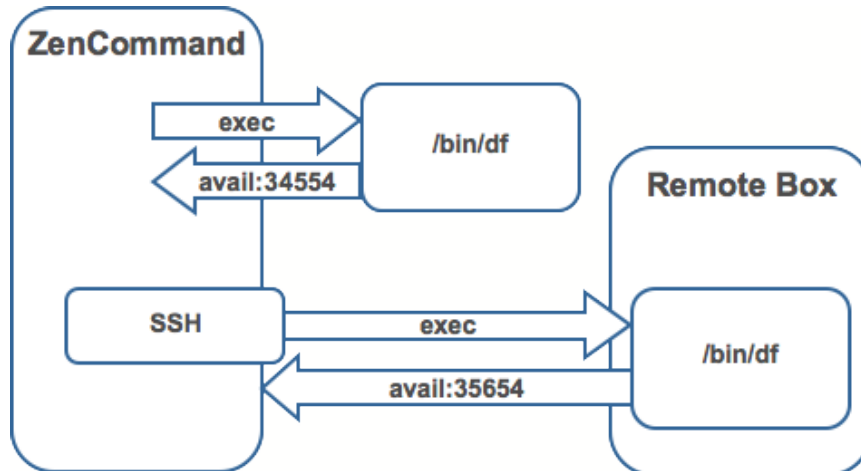
The predefined /Server/Scan device class is an example configuration for monitoring TCP services on devices using a port scan. If you have a device that you want to monitor for service availability alone, you can place it under this device class. The system will not collect performance data for devices in this class.

Monitoring Windows Services

The Windows Services page (**Infrastructure > Windows Services**) has been deprecated. Use the WinService monitoring template instead.

Monitoring Using ZenCommand

Resource Manager has the ability to run Nagios® and Cacti plug-ins through the ZenCommand process. ZenCommand can run plugins locally and remotely by using a native SSH transport. When run, the system tracks the return code of each plug-in and then creates events with plug-in output. Additionally, it can track performance information from a plug-in.

Figure 78: Running ZenCommand

Plugin Format for ZenCommands

Nagios® plugins are configured by using a command template.

A template named “Device” will bind to all devices below the template definition. Within each template is a list of commands that will run. The commands can be any program that follows the Nagios® plug-in standard. Inputs are command line arguments; output is the first line of stdout, plus a return code.

Note Resource Manager return codes differ from Nagios return codes, as follows:

| Value | Resource Manager | Nagios |
|-------|------------------|----------|
| 0 | Clear | OK |
| 1 | Data Source | WARNING |
| 2 | Data Source+1 | CRITICAL |
| 3 | Data Source | UNKNOWN |

For complete information about Nagios plugin guidelines, browse to this location:

<http://nagiosplug.sourceforge.net/developer-guidelines.html>

A Nagios® command has several fields:

- name – Specifies the name of the command object.
- enabled – Indicates whether this command should be used on a given device.
- component – Specifies the component name to use when zencommand sends events to the system.
- event class – Specifies the event class to use when sending events to the system.
- severity – Sets the default severity to use when sending events to the system.
- cycle time – Sets the frequency a command should be run (in seconds).
- command template – Specifies the command to run.

The command template string is built by using Zope TALEX expressions. Several variables are passed when evaluating the template. They are:

- zCommandPath – Path to the zencommand plug-ins on a given box it comes from the configuration property zCommandPath. zCommandPath is automatically added to a command if a path is absent from the beginning of the command.

- `devname` – Device name of the device against which the command is being evaluated.
- `dev` – Device object against which the command is being evaluated.
- `here` – Context of evaluation. For a device, this is equivalent to `dev` for a component (such as a file system or interface). This is the component object.
- `compname` – If this command evaluates against a component, specifies its name as a string.
- `now` – Current time.

Template values are accessed like shell variables. They are the same as the expression syntax used in the appendix titled TALEX Expressions (in this guide).

Testing ZenCommands

You can test ZenCommand data sources by using the `zetestcommand` shell script.

- 1 Log in to the Control Center host as a user with `serviced` CLI privileges.
- 2 Attach to the `zenhub` service.

```
serviced service attach zenhub
```

- 3 Change to the `zenoss` user.

```
su - zenoss
```

- 4 Run the `zetestcommand` script.

```
zetestcommand -d DeviceID --datasource=DataSourceName
```

where `DeviceID` is the ID of the device on which you want to run the command, and `DataSourceName` is the name of a data source on a template associated with the device.

The `zetestcommand` script prints the results of the command to standard output.

SNMP Monitoring

OID represent the data points where the data for the graphs comes from. Sometimes the reason that a graph is not appearing is because the OID for the particular graph is not valid for the device. You can test this validity using the command line to see if you can return a value. To test the validity of an OID data point giving performance data:

- 1 Log in to the Control Center host as a user with `serviced` CLI privileges.
- 2 Attach to the `Zenoss.resmgr` service.

```
serviced service attach Zenoss.resmgr
```

- 3 Change to the `zenoss` user.

```
su - zenoss
```

- 4 Run the `snmp get` command for one of the OIDs

In this case, use the command:

```
$ snmpget -v 2c -cpublic build .1.3.6.1.4.1.2021.4.14.0
```

If the OID is valid it will return a value.

Here are some basic SNMP commands to gather certain information.

- a** Walk a basic system MIB.

```
snmpwalk -v 2c -cpublic <device name> system
```

- b** Walk an interface description

```
snmpwalk -v 2c -cpublic <device name> ifDescr
```

- c** Get a single value.

```
snmpget -v 2c -cpublic <device name> ifDescr.2
```

- d** Detailed description of an OID value.

```
snmptranslate -Td RFC1213-MIB::ifDescr
```

- e** Convert a name to a raw OID.

```
snmptranslate -On RFC1213-MIB::ifDescr
```

- f** Convert a raw OID to a short name

```
snmptranslate -OS .1.3.6.1.2.1.2.2.1.2
```

Monitoring Devices Remotely Through SSH

You can monitor devices remotely through SSH. Follow the steps in the following sections to set up remote monitoring.

Changing Resource Manager to Monitor Devices Remotely Using SSH

You must edit system properties for the group where you want to collect remote information using SSH.

- 1 Navigate to the device class path you want to monitor remotely. You can apply this monitoring for a device or a device class path.
- 2 Change the configuration properties value for the group. After selecting the device class, click **Details**, and then select **Configuration Properties**. The Configuration Properties page appears.

Figure 79: Device Class Configuration Properties

| Is Local | Category | Name | Value | Path |
|----------|------------------|-----------------------------|--------------------|------|
| | Cisco | zCiscoACEUseSSL | true | / |
| | Cisco | zCiscoRemodelEventClassKeys | | / |
| | Cisco UCS | zCiscoUCSCIMCEventsInterval | 60 | / |
| | Cisco UCS | zCiscoUCSCIMCPerInterval | 300 | / |
| | Cisco UCS | zCiscoUCSMManagerPassword | ***** | / |
| | Cisco UCS | zCiscoUCSMManagerPort | 443 | / |
| | Cisco UCS | zCiscoUCSMManagerUseSSL | true | / |
| | Cisco UCS | zCiscoUCSMManagerUser | admin | / |
| | Modeler Controls | zCollectorClientTimeout | 180 | / |
| | Modeler Controls | zCollectorDecoding | utf-8 | / |
| | Misc | zCollectorLogChanges | false | / |
| | zencommand | zCommandCommandTimeout | 15 | / |
| | zencommand | zCommandExistenceTest | test -f %s | / |
| | zencommand | zCommandLoginTimeout | 10 | / |
| | zencommand | zCommandLoginTries | 1 | / |
| | zencommand | zCommandPassword | | / |
| | zencommand | zCommandPath | \$ZENHOME/libexec | / |
| | zencommand | zCommandPort | 22 | / |
| | zencommand | zCommandProtocol | ssh | / |
| | zencommand | zCommandSearchPath | | / |
| | zencommand | zCommandUsername | | / |
| | Control Center | zControlCenterHost | \$(here/managerip) | / |

You must make changes to the following configuration properties:

- zCollectorPlugins
- zCommandPassword
- zCommandPath
- zCommandUsername
- zSnmppMonitorIgnore

The following table lists sample values set up for remote devices. These have a pre-shared key (with no password) set up from the collector to the remote boxes. (It also can use password authorization if the password is entered into zCommandPassword.)

| Configuration Properties | Value |
|--------------------------|---|
| zCollectorPlugins | snmp portscan |
| zCommandPassword | The SSH password for the remote machine. |
| zCommandPath | The path to zenplugin.py |
| zCommandUsername | The SSH user name for the remote machine. |
| zSnmppMonitorIgnore | True |

Two passes are required for full modeling. The first pass obtains the platform type (so that the system knows which plugins to run). The second pass provides detailed data on interfaces and file systems.

- 1 Log in to the Control Center host as a user with serviced CLI privileges.
- 2 Attach to the zenmodeler service.

```
serviced service attach zenmodeler
```

- 3 Change to the zenoss user.

```
su - zenoss
```

- 4 Run the zenmodeler command.

```
$ zenmodeler run -d DeviceName
```

where *DeviceName* is the fully qualified device name.

- 5 Run the zenmodeler command a second time to use the plugins the command gathered on the first pass.

Using the Predefined /Server/Command Device Class

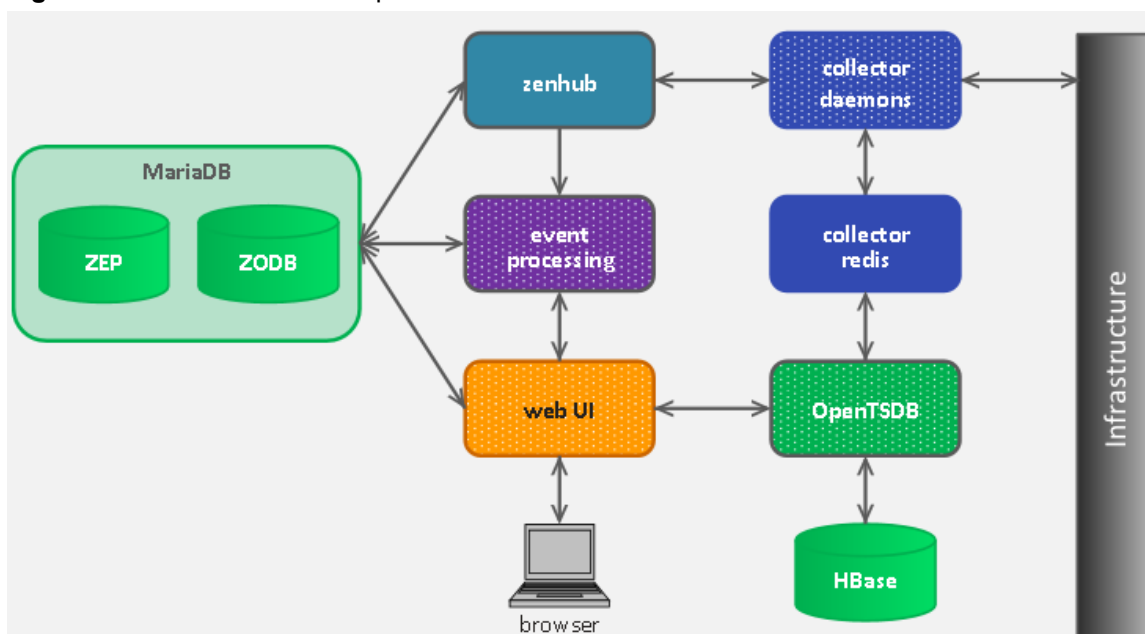
The /Server/Command device class is an example configuration for modeling and monitoring devices using SSH. The configuration properties have been modified (as described in the previous sections), and device, file system, and Ethernet interface templates that gather data over SSH have been created.

You can use this device class as a reference for your own configuration; or, if you have a device that needs to be modeled or monitored via SSH/Command, you can place it under this device class to use the pre-configured templates and configuration properties. You must set the zCommandUsername and zCommandPassword properties to the appropriate SSH login information for each device.

Performance Monitoring

Resource Manager stores device and daemon performance metrics directly into a time series database (OpenTSDB) that is run on top of an HBase instance. Writing directly into OpenTSDB has eliminated the need for RRD files to be stored on the collectors. The following image shows how the collector daemons fit into the data collection portion of the Resource Manager architecture.

Figure 80: Data Collection Simplified Architecture



Resource Manager uses several methods to monitor performance metrics of devices and device components. These are:

- **ZenPerfSNMP**- Collects data through SNMP from any device correctly configured for SNMP monitoring.
- **ZenWinPerf**- ZenPack that allows performance monitoring of Windows servers.
- **ZenCommand**- Logs in to devices (by using telnet or ssh) and runs scripts to collect performance data.
- **Other ZenPacks**- Collect additional performance data. Examples include the ZenJMX ZenPack, which collects data from enterprise Java applications, and the HttpMonitor ZenPack, which checks the availability and responsiveness of Web pages.

Regardless of the monitoring method used, the system stores performance monitoring configuration information in *monitoring templates*.

About Monitoring Templates

Monitoring templates determine how the system collects performance data for devices and device components. You can define monitoring templates for device classes and individual devices.

Templates comprise three types of objects:

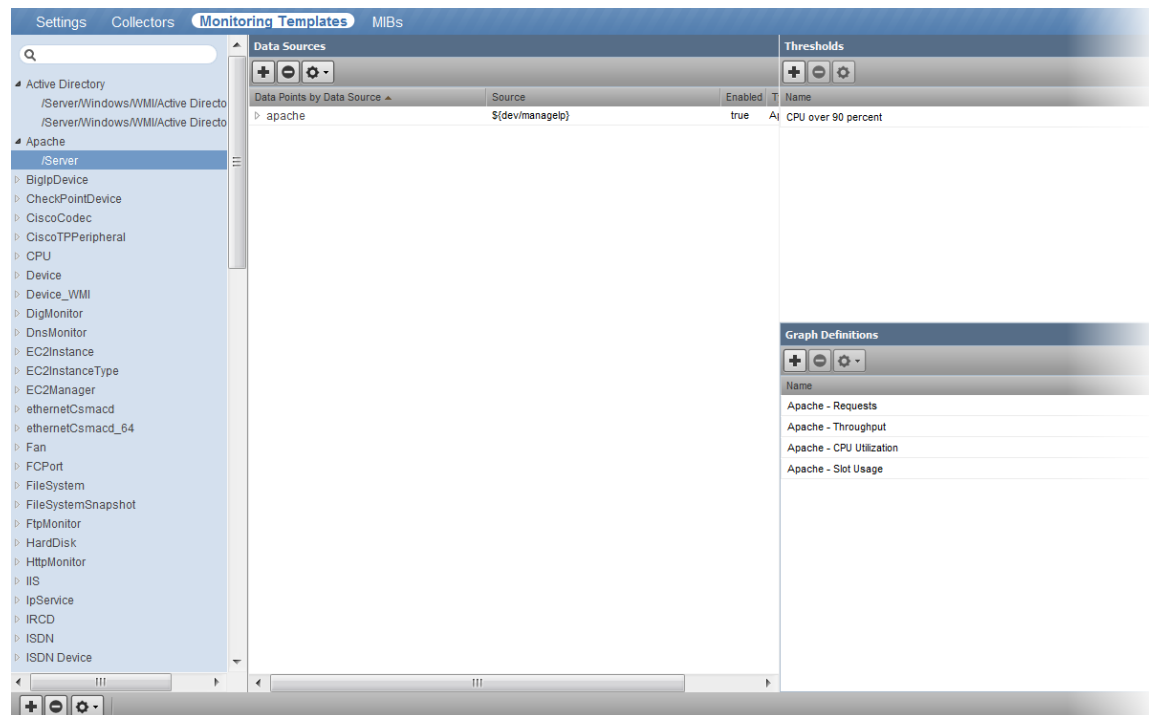
- **Data Sources**- Specify the exact data points to collect, and the method to use to collect them.
- **Thresholds**- Define expected bounds for collected data, and specify events to be created if the data does not match those bounds.
- **Graph Definitions**- Describe how to graph the collected data on the device or device components.

Before the system can collect performance data for a device or component, it must determine which monitoring templates apply. This process is called *template binding*.

Viewing Monitoring Templates

To view monitoring templates, select Advanced from the Navigation menu, and then select Monitoring Templates.

Figure 81: Monitoring Template for Load Average Graph



Template Binding

Before the system can collect performance data for a device or component, it must determine which templates apply. This process is called template binding.

First, the system determines the list of template names that apply to a device or component. For device components, this usually is the meta type of the component (for example, FileSystem, CPU, or HardDisk). For devices, this list is defined by the `zDeviceTemplates` configuration property.

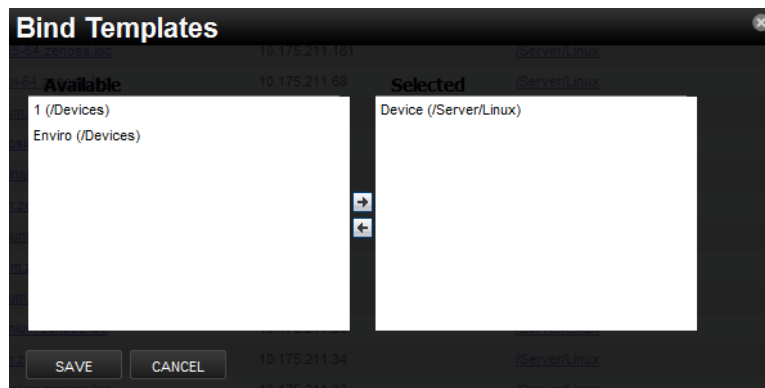
After defining the list, the system locates templates that match the names on the list. For each name, it searches the device and then searches the device class hierarchy. Resource Manager uses the lowest template in the hierarchy that it can locate with the correct name, ignoring others of the same name that might exist further up the device class hierarchy.

Binding Templates

To edit the templates bound to a device:

- 1 From the Navigation menu, select **Infrastructure**. The device list appears.
- 2 Select a device in the device list.
- 3 Select **Bind Templates** from the Action menu. The Bind Templates dialog appears.

Figure 82: Bind Templates



- 4 Select a template from the Available list and move it to the Selected list to bind it to the selected device.
- 5 Click **Save**.

Data Sources

Data sources specify which data points to collect and how to collect them. Each monitoring template comprises one or more data sources. The system provides two built-in data source types: SNMP and COMMAND. (Other data source types are provided through ZenPacks.)

- **SNMP**- Define data to be collected via SNMP by the ZenPerfSNMP daemon. They contain one additional field to specify which SNMP OID to collect. (Many OIDs must end in .0 to work correctly.) Because SNMP data sources specify only one performance metric, they contain a single data point.
- **Command**- Specify data to be collected by a shell command that is executed on the Resource Manager server or on a monitored device. The ZenCommand daemon processes COMMAND data sources. A COMMAND data source may return one or more performance metrics, and usually has one data point for each metric.

Shell commands used with COMMAND data sources must return data that conforms to the Nagios plug-in output specification. For more information, see the section titled Monitoring Using ZenCommand.

Adding a Data Source

To add a data source to a monitoring template:

- 1 Select **Advanced** from the Navigation menu, and then select **Monitoring Templates**.
- 2 In the tree view, select the monitoring template to which you want to add a data source.
- 3 In the Data Sources area, click the **Add** button. The Add Data Source dialog appears.
- 4 Enter a name for the data source and select the type, and then click **Submit**. The data source is added to the list in the Data Sources area.

- 5 Double-click the data source in the list. The Edit Data Source dialog appears.
- 6 Enter or select values to define the data source.

Data Points

Data sources can return data for one or more performance metrics. Each metric retrieved by a data source is represented by a data point.

You can define data points to data sources with all source types except SNMP and VMware. Because these data source types each rely on a single data point for performance metrics, additional data point definition is not needed.

To add a data point to a data source:

- 1 From the Navigation menu, select **Advanced > Monitoring Templates**.
- 2 In the Data Sources area, highlight the row containing a data source.
- 3 Select **Add Data Point** from the Action menu. The Add Data Point dialog appears.
- 4 Enter a name for the data point, and then click **Submit**.

Note For COMMAND data points, the name should be the same as that used by the shell command when returning data.

- 5 Double-click the newly added data point to edit it. Enter information or make selections to define the data point:
 - **Name**- Displays the name you entered in the Add a New DataPoint dialog.
 - **RRD Type**- Specify the RRD data source type to use for storing data for this data point. Available options are:
 - **COUNTER**- Saves the rate of change of the value over a step period. This assumes that the value is always increasing (the difference between the current and the previous value is greater than 0). Traffic counters on a router are an ideal candidate for using COUNTER.
 - **GAUGE**- Does not save the rate of change, but saves the actual value. There are no divisions or calculations. To see memory consumption in a server, for example, you might want to select this value.

Note Rather than COUNTER, you may want to define a data point using DERIVED and with a minimum of zero. This creates the same conditions as COUNTER, with one exception. Because COUNTER is a "smart" data type, it can wrap the data when a maximum number of values is reached in the system. An issue can occur when there is a loss of reporting and the system (when looking at COUNTER values) thinks it should wrap the data. This creates an artificial spike in the system and creates statistical anomalies.

- **DERIVE**- Same as COUNTER, but additionally allows negative values. If you want to see the rate of change in free disk space on your server, for example, then you might want to select this value.
- **ABSOLUTE**- Saves the rate of change, but assumes that the previous value is set to 0. The difference between the current and the previous value is always equal to the current value. Thus, ABSOLUTE stores the current value, divided by the step interval.
- **Create Command**- Enter an RRD expression used to create the database for this data point. If you do not enter a value, then the system uses a default applicable to most situations.

For details about the rrdcreate command, go to:

<http://oss.oetiker.ch/rrdtool/doc/rrdcreate.en.html>

- **RRD Minimum**- Enter a value. Any value received that is less than this number is ignored.
 - **RRD Maximum**- Enter a value. Any value received that is greater than this number is ignored.
- 6 Click **Save** to save the defined data point.


Data Point Aliases

Performance reports pull information from various data points that represent a metric. The report itself knows which data points it requires, and which modifications are needed, if any, to put the data in its proper units and format.

The addition of a data point requires changing the report.

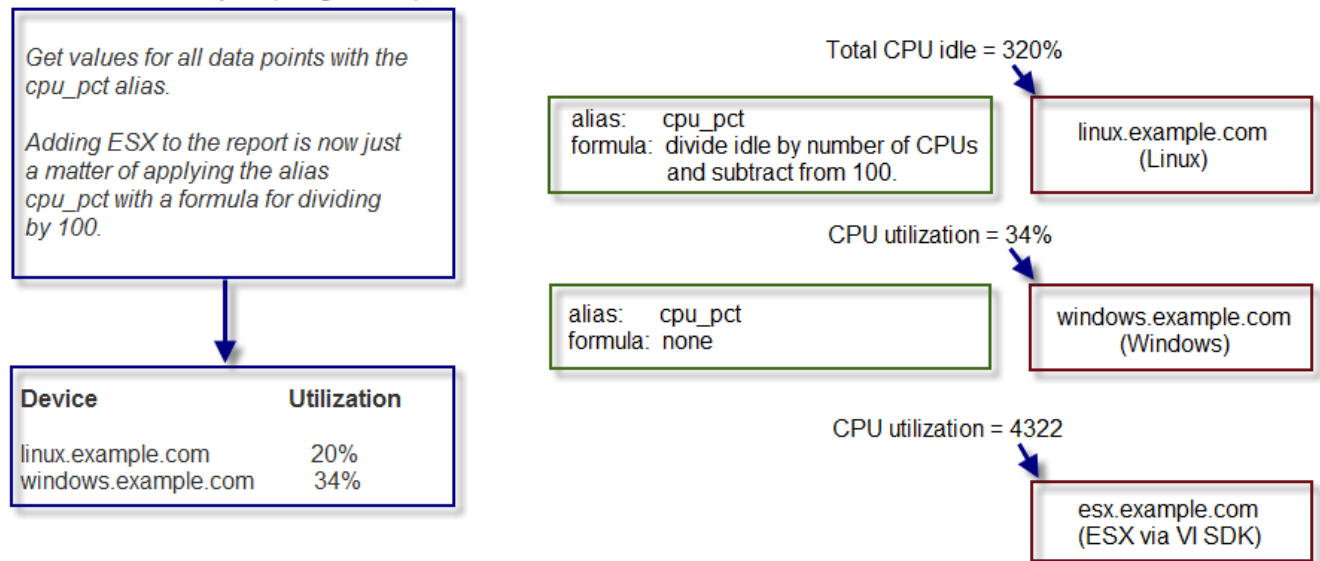
Figure 83: CPU Utilization Report

CPU Utilization Report



To allow for more flexibility in changes, some reports use *data point aliases*. Data point aliases group data points so they can be more easily used for reporting. In addition, if the data points return data in different units, then the plugin can normalize that data into a common unit.

An alias-based report looks up the data points that share a common alias string, and then uses them. This approach allows you to add data points without changing the report.

Figure 84: Alias-Based CPU Utilization Report**CPU Utilization Report** (using aliases)

In the simplest cases, data from the target data points are returned in units expected by a report. For cases in which data are not returned in the same units, an alias can use an associated formula at the data point. For example, if a data point returns data in kilobytes, but the report expects data in bytes, then the formula multiplies the value by 1024.

Alias Formula Evaluation

The system evaluates the alias formula in three passes.

Reverse Polish Notation

When complete, the alias formula must resolve to a Reverse Polish Notation (RPN) formula. For the simple conversion of kilobytes into bytes, the formula is:

```
1024, *
```

For more information on RPN formulas, browse to this site:

http://oss.oetiker.ch/rrdtool/doc/rrdgraph_rpn.en.html

Using TALES Expressions in Alias Formulas

For cases in which contextual information is needed, the alias formula can contain a TALES expression that has access to the device as context (labeled as "here"). The result of the TALES evaluation should be an RRD formula.

For example, if the desired value is the data point value divided by total memory, the formula is:

```
${here/hw/totalMemory}, /
```

For more information about TALES, see *TALES Expressions* on page 230.

Using Python in Alias Formulas

You also can embed full Python code in an alias formula for execution. The code must construct a string that results in a valid RRD formula. To signal the system to evaluate the formula correctly, it must begin with:

```
__EVAL:
```

Using the same example as in the previous section (division by total memory), the formula is:

```
__EVAL:here.hw.totalMemory + ",/"
```


Adding a Data Point Alias

To add an alias to a data point:

- 1 Navigate to a data source on a monitoring template.
- 2 Double-click a data point in the list to edit it. The Edit Data Point dialog appears.
- 3 Enter the alias name and the formula.

Note If the data point returns values in the desired units, then leave the Formula value blank.

Figure 85: Add Data Point Alias



- 4 Click **Save**.

Reports That Use Aliases

For information about reports that use aliases, refer to the chapter titled "Reporting."

The following table shows performance reports that use aliases, and the aliases used. To add data points to a report, add the alias, and then ensure the values return in the expected units.

CPU Utilization

| Alias | Expected Units |
|-----------------|----------------|
| loadAverage5min | Processes |
| cpu_pct | Percent |

Thresholds

Thresholds define expected bounds for data points. When the value returned by a data point violates a threshold, the system creates an event. There are several threshold types available:

- MinMax
- ValueChange
- CiscoStatus

- PredictiveThreshold

There are many thresholds already defined in the system. You can see all the defined thresholds in the Defined Thresholds report which can be accessed on the **REPORTS > Enterprise Reports** menu.

The following sections describe each type of threshold and how to create a new one and edit an existing one.

Additional threshold types may be provided through installed ZenPacks.

MinMax Threshold

MinMax thresholds inspect incoming data to determine whether it exceeds a given maximum or falls below a given minimum. You can use a MinMax threshold to check for these scenarios:

- *The current value is less than a minimum value.* To do this, you should set only a minimum value for the threshold. Any value less than this number results in creation of a threshold event.
- *The current value is greater than a maximum value.* To do this, you should set only a maximum value for the threshold. Any value greater than this number results in creation of a threshold event.
- *The current value is not a single, pre-defined number.* To do this, you should set the minimum and maximum values for the threshold to the same value. This will be the only "good" number. If the returned value is not this number, then a threshold event is created.
- *The current value falls outside a pre-defined range.* To do this, you should set the minimum value to the lowest value within the good range, and the maximum value to the highest value within the good range. If the returned value is less than the minimum, or greater than the maximum, then a threshold event is created.
- *The current value falls within a pre-defined range.* To do this, you should set the minimum value to the highest value within the bad range, and the maximum value to the lowest value within the bad range. If the returned value is greater than the maximum, and less than the minimum, then a threshold event is created.

ValueChange Threshold

ValueChange thresholds inspect incoming data to determine whether a status change has occurred and if so issues an event based on the defined severity.

CiscoStatus Threshold

The CiscoStatus threshold is a special threshold that uses preconfigured maps of numeric values returned by SNMP datasources to Resource Manager event severities. The following is an example OID and its mapped values. For a complete list of OIDs and values that are supported, visit http://wiki.zenoss.org/CiscoStatus_Threshold.

CISCO-ENTITY-FRU-CONTROL-MIB::cefcModuleOperStatus

- (1) unknown - Critical
- (2) OK - Clear
- (3) disabled - Clear
- (4) OK (diag failed) - Warning
- (5) boot - Warning
- (6) self-test - Warning
- (7) failed - Critical
- (8) missing - Critical
- (9) mismatch w/parent - Critical
- (10) mismatch w/config - Critical
- (11) diag-failed - Critical
- (12) dormant - Critical
- (13) out of service (admin) - Info

- (14) out of service (environ) - Critical
- (15) powered down - Critical
- (16) powered up - Critical
- (17) power denied - Critical
- (18) power cycle - Warning
- (19) OK (power warning) - Warning
- (20) OK (power critical) - Critical
- (21) sync in progress - Clear
- (22) upgrading - Critical
- (23) OK (auth failed) - Critical

PredictiveThreshold

The predictive threshold allows you to use data from the past to project a value in the future and send an event if the future projected threshold is breached during a pre-defined time period. Resource Manager uses a linear projection algorithm for all projections. For example, you could predict total raw storage capacity of a disk or the bandwidth utilization percentage. Several predictive thresholds are already defined in the system. You can see all the defined thresholds in the Defined Thresholds report which can be accessed on the **REPORTS > Enterprise Reports** menu.

Predictive thresholds are also used for creating a trendline on a graph. You can add the threshold to the management of graph points so that a trendline will be added to the graph.

Adding Thresholds

Follow these steps to define a threshold for a data point:

- 1 From the Navigation menu, select **ADVANCED > Monitoring Templates**.
- 2 Click on the template that contains the data point you want to use in your threshold. The Data Sources window is populated with folders containing all the data points being collected.
- 3 Select the data point by opening the appropriate folder in the data source and clicking on the data point row.
- 4 In the Thresholds area, click the **Add** icon. The Add Threshold dialog appears.
- 5 Enter a name and select the threshold type, then click **Add**. The threshold name and type is displayed in the Thresholds window.

Editing MinMax thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a MinMax threshold:

- 1 Double-click the threshold in the list. The Edit Threshold dialog appears.

Figure 86: Edit Threshold

Edit Threshold

Name: High Utilization Rx

Description:

Type: MinMaxThreshold

Explanation:

Resolution:

DataPoints:

| | | |
|-------------------------|---|-----------------|
| fcErrStats_crcRx | | fcStats_bytesRx |
| fcErrStats_discardRx | | |
| fcErrStats_discardTx | | |
| fcErrStats_linkFailures | ↑ | |
| fcErrStats_rx | ↑ | |
| fcErrStats_signalLosses | → | |
| fcErrStats_syncLosses | ← | |
| fcErrStats_tooLongRx | ↓ | |
| fcErrStats_tooShortRx | ↓ | |
| fcErrStats_tx | ↓ | |
| fcStats_bytesTx | | |
| fcStats_packetsRx | | |
| fcStats_packetsTx | | |

Severity: Warning

Enabled

Minimum Value:

Maximum Value: (here..linkSpeed or 1e10) / 8 * .9

Event Class: /Perf/Interface

Escalate Count: 0

SAVE CANCEL

2 Enter or select values to define the threshold:

- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
- **Description**- Description of the threshold that you entered on the Add a New Threshold dialog. The description is included in each event that is created from this threshold.
- **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
- **Explanation**- Information field where a user can enter information about what the event means. This field is included in each event that is created from this threshold
- **Resolution**- Information field where a user can enter information about what to do to resolve the event. This field is included in each event that is created from this threshold.
- **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column.
- **Severity**- Select the severity level of the first event triggered when this threshold is breached.

- **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
- **Minimum Value**- If this field contains a value, then each time one of the select data points falls below this value an event is triggered. This field may contain a number or a Python expression. When using a Python expression, the variable *here* references the device or component for which data is being collected. For example, a 90% threshold might be specified as:

```
(here.linkSpeed or 1e10) /8 * .9
```

The division by 8 is needed because interface speed frequently is reported in bits/second, where the performance data is bytes/second.

- **Maximum Value**- If this field contains a value, then each time one of the selected data points goes above this value an event is triggered. This field may contain a number or a Python expression.
 - **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
 - **Escalate Count**- Enter the number of consecutive times this threshold can be broken before the event severity is escalated by one step. A value of zero (0) indicates that the severity will not escalate.
- 3 Click **Save** to confirm the edits.

Editing ValueChange thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a ValueChange threshold:

- 1 Double-click the threshold in the list. The Edit Threshold dialog appears.

Figure 87: Edit Threshold

- 2 Enter or select values to define the threshold:

- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
- **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.

- **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column. When the data point changes status, an event will be triggered.
 - **Severity**- Select the severity level of the first event triggered when this threshold is breached.
 - **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
 - **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
- 3 Click **Save** to confirm the edits.

Editing CiscoStatus thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a ValueChange threshold:

- 1 Double-click the threshold in the list. The Edit Threshold dialog appears.

Figure 88: Edit Threshold

Edit Threshold

Name:

Type:

DataPoints:

- cvnVethIfAdditionalState_cvnVethIfAdditional
- ifHCInBroadcastPkts_ifHCInBroadcastPkts
- ifHCInMulticastPkts_ifHCInMulticastPkts
- ifHCInOctets_ifHCInOctets
- ifHCInUcastPkts_ifHCInUcastPkts
- ifHCOutBroadcastPkts_ifHCOutBroadcastPkt
- ifHCOutMulticastPkts_ifHCOutMulticastPkts
- ifHCOutOctets_ifHCOutOctets
- ifHCOutUcastPkts_ifHCOutUcastPkts
- ifInDiscards_ifInDiscards
- ifInErrors_ifInErrors
- ifOutDiscards_ifOutDiscards
- ifOutErrors_ifOutErrors

ifOperStatus_ifOperStatus

Severity:

Event Class Key:

Enabled

Event Class:

- 2 Enter or select values to define the threshold:
 - **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
 - **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
 - **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column. When the data point changes status, an event will be triggered.

- **Severity**- Select the severity level of the first event triggered when this threshold is breached.
 - **Event Class Key**- The event class key from the data point you selected will be shown here.
 - **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
 - **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
- 3 Click **Save** to confirm the edits.

Editing Predictive Thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

- 1 Double-click the threshold in the list. The Edit Threshold dialog appears.

Figure 89: Edit Threshold

Edit Threshold

Name: Projected High Pct Rx

Description:

DataPoint: bandwidthUtilizationPctRx_bandwidthUtilizationPctf

Event Class: /Capacity

Type: PredictiveThreshold

Enabled

Severity: Info

Aggregate Function: max

Projection Algorithm

Algorithm: linear

Projection Algorithm Parameters:

Alerting

Amount of Data Used in Projection: 10 days

Send an Event if the Threshold is Breached in the Next: 10 days

Minimum Value:

Maximum Value: 90

SAVE CANCEL

- 2 Enter or select values to define the threshold:
- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
 - **Description**- Description of the threshold that you entered on the Add a New Threshold dialog. The description is included in each event that is created from this threshold.
 - **Data Point**- Select the data point to which this threshold will apply.
 - **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
 - **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
 - **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
 - **Severity**- Select the severity level of the first event triggered when this threshold is breached.
 - **Aggregate Function**- The type of function to use when analyzing the past data. For example, do you look at the peak values in the past or average values. The default value is max.
 - **Algorithm**- Algorithm to use for the projection. Value is linear.

- **Projection Algorithm Parameters**- Some algorithms need additional parameters. Enter them as required.
- **Amount of Data Used in Projection**- Set the amount of historical data to be used when calculating the projection.
- **Send an Event if the Threshold is Breached in the Next**- Set the timeframe in days, weeks, or months. If the threshold is breached within this time period in the future, an event will be sent.
- **Minimum Value**- If this field contains a value, then each time one of the select data points falls below this value an event is triggered. This field may contain a number or a Python expression. An expression is required when using gauge or calculated values. When using an expression, the variable *here* references the device or component for which data is being collected. For example, a 90% threshold for interface speed might be specified as:

```
(here.linkSpeed or 1e10) / 8 * .9
```

The division by 8 is needed because interface speed frequently is reported in bits/second, where the performance data is bytes/second.

- **Maximum Value**- If this field contains a value, then each time one of the selected data points goes above this value an event is triggered. This field may contain a number or a Python expression. An expression is required for gauge or calculated values.
- 3 Click **Save** to confirm the edits.


Adding a Trendline to a Graph

Trendlines show you the projected utilization of a device or component over a specified range of time. The following procedure is tailored to be performed on a CiscoUCS device. In order to add a trendline to a graph, the threshold that you want to use must already be created in the system.

To add a trendline to an existing graph:

- 1 Navigate to **Infrastructure > Devices**, then select a **CiscoUCS** device.
The selected device's Overview page is displayed.
- 2 From the Devices panel, select **Components > Chassis**, and then select a chassis from the Chassis list.
- 3 Change the **Display** drop-down menu to **Templates**.
- 4 To the right of **Graph Definitions**, click **+** to create a new graph.
- 5 In the **Add Graph Definition** dialog, enter a name for the graph.
Alternatively, click the name of an existing graph to add the trendline to that graph.
- 6 Select the new graph, then from the gear icon, select **Manage Graph Points**.
The Manage Graph Points dialog is displayed.
- 7 Click **+ > Threshold**.

Figure 90: Manage Graph Points



- 8 Select the name of the new predictive threshold and click **Submit**.
- 9 In the **Manage Graph Points** dialog, click **Save**.
- 10 To view the trendline, change the **Display** drop-down menu to **Graphs**.
- 11 Scroll down to the new (or existing graph) to view the newly added trendline.

Note By default only the last 10 days of data is used in the trendline calculation. There may not be enough data to see the trendline appear on the graph. However, you can view the projected exhaustion dates in the lower right corner.

Performance Graphs

You can include any of the data points or thresholds from a monitoring template in a *performance graph*.

Graphs are no longer static images, but are rendered using the NVD3.js JavaScript library.

To define a graph:

- 1 Select Advanced from the Navigation menu, and then select Monitoring Templates.
- 2 In the Graph Definitions area, click the **Add** icon.

The Add Graph Definition dialog appears.

- 3 Enter a name for the graph, and then click **Submit**.
- 4 Double-click the graph in the list to edit it. Enter information or select values to define the graph:
 - **Name**- Optionally edit the name of the graph you entered in the Add a New Graph dialog. This name appears as the title of the graph.
 - **Height**- Enter the height of the graph, in pixels.
 - **Width**- Enter the width of the graph, in pixels.
 - **Units**- Enter a label for the graph's vertical axis.
 - **Logarithmic Scale**- Select True to specify that the scale of the vertical axis is logarithmic. Select False (the default) to set the scale to linear. You might want to set the value to True, for example, if the data being graphed grows exponentially. Only positive data can be graphed logarithmically.
 - **Base 1024**- Select True if the data you are graphing is measured in multiples of 1024. By default, this value is False.
 - **Min Y**- Enter the bottom value for the graph's vertical axis.
 - **Max Y**- Enter the top value for the graph's vertical axis.
 - **Has Summary**- Select True to display a summary of the data's current, average, and maximum values at the bottom of the graph.

Figure 91: Graph Definition

View and Edit Graph Definition

Name:

Height:

Width:

Units:

Logarithmic Scale:

Base 1024:

Min Y:

Max Y:

Has Summary:

- 5 Click **Submit** to save the graph.

Graph Points

Graph points represent each data point or threshold that is part of a graph. You can add any number of graph points to a graph definition by adding data points or thresholds.

From the Graph Definitions area of the Monitoring Templates page:

- 1 From the Action menu, select **Manage Graph Points**. The Manage Graph Points dialog appears.
- 2 From the Add menu, add a data point, threshold, or custom graph point.
- 3 Select values, then click **Submit**.

Note Thresholds are always drawn before other graph points.

Re-sequencing Graph Points

To re-sequence graph points, drag a graph point row in the Manage Graph Points dialog. (Click-and-drag from an "empty" part of the row.)

DataPoint Graph Points

DataPoint graph points draw the value of data points from the template on a graph.

Adding DataPoint Graph Points

To define a DataPoint graph point:

- 1 From the **Add** menu on the Manage Graph Points dialog, select **Data Point**. The Add Data Point dialog appears.
- 2 Select one or more data points defined in this template. On data point graph point is created for each data point you select from the list.
- 3 Optional: Select the **Include Related Thresholds** option. If selected, then any graph points are created for any thresholds that have been applied to the selected data points as well.

4 Click **Submit**.

Editing DataPoint Graph Points

Double-click the name of the graph point to go to its edit page. Enter information or select values to edit the graph point:

- **Name**- This is the name that appears on the Graph Definition page. By default, it appears in the graph legend.
- **Line Type**- Select Line to graph the data as a line. Select Area to fill the area between the line and the horizontal axis with the line color. Select None to use this data point for custom RRD commands and do not want it to be explicitly drawn.
- **Line Width**- Enter the pixel width of the line.
- **Stacked**- If selected, then the line or area is drawn above the previously drawn data. At any point in time on the graph, the value plotted for this data is the sum of the previously drawn data and the value of this data point now. You might set this value, for example, to assess total packets if measuring packets in and packets out.
- **Format**- Specify the RRD format to use when displaying values in the graph summary. For more information on formatting strings, go to:

http://oss.oetiker.ch/rrdtool/doc/rrdgraph_graph.en.html

- **RPN**- Optionally enter an RPN expression that alters the value of the data being graphed for the data point. For example, if the data is stored as bits, but you want to graph it as bytes, enter an RPN value of "8,/" to divide by 8. For more information about RPN notation, go to:

<http://oss.oetiker.ch/rrdtool/tut/rpntutorial.en.html>

- **Limit**- Optionally specify a maximum value for the data being graphed.
- **Consolidation**- Specify the RRD function used to graph the data point's data to the size of the graph. Most of the time, the default value of AVERAGE is appropriate.
- **Color**- Optionally specify a color for the line or area. Enter a six-digit hexadecimal color value with an optional two-digit hex value to specify an alpha channel. An alpha channel value is only used if 'stacked' is True.
- **Legend**- Name to use for the data in the graph legend. By default, this is a TALEX expression that specifies the graph point name. The variables available in this TALEX expression are here (the device or component being graphed) and graphPoint (the graph point itself).
- **Available RRD Variables**- Lists the RRD variables defined in this graph definition. These values can be used in the RPN field.

Editing Threshold Graph Points

Threshold graph points graph the value of thresholds from the template.

To edit a threshold graph point, double-click it in the list:

You can edit values for Name, Color, and Legend for a threshold graph point. Refer to the definitions in the section titled Editing DataPoint Graph Points for more information.

Performance Data Retention

Resource Manager stores all performance data (a.k.a., metrics) in HBase using OpenTSDB. The default retention policy saves performance data for 90 days. To change the default, the time to live (TTL) must be adjusted on the OpenTSDB column families in HBase.

Note TTL is defined in seconds.

Once a TTL value is changed, the data retention will adjust on the next major HBase compaction, which by default is once per day.

Changing the Performance Data Retention Time

To change the performance data retention time from the default value of 90 days:

- 1 Log in to the Control Center master host as a user with `sudo` and `docker` privileges.
- 2 Stop the `opentsdb` writer service.

```
serviced service stop opentsdb/writer
```

- 3 Execute the following command to list all the services and their `SERVICEID` values. Take note of the `SERVICEID` for the `opentsdb` reader service. It will be used as an argument in the following step.

```
serviced service list
```

- 4 Execute the following command where `$id` is the `opentsdb` reader `SERVICEID` and `$ttl` is your desired TTL value, in seconds.

```
serviced service shell $id /opt/opentsdb/set-opentsdb-table-ttl.sh  
$ttl
```

- 5 Start the `opentsdb` writer service.

```
serviced service start opentsdb/writer
```

Event Management

Events, and the graphs generated from performance monitoring, are the primary operational tools for understanding the state of your environment.

Basic Event Fields

To enter the event management system, an event must contain values for the device, severity, and summary fields. If an event is missing any of these fields, then Resource Manager rejects it.

Basic event fields are:

- device
- ipAddress
- eventState
- severity
- summary
- message
- evid

device and ipAddress Fields

The device field is a free-form text field that allows up to 255 characters. Resource Manager accepts any value for this field. If the device field contains an IP address, then the system queries for devices with a matching address. If it finds a match, it changes the device field to the found device identifier.

The ipAddress field is a free-form text field. This field is not required. If the system cannot successfully locate a device based on the event's device field content, it attempts to find the device based the event ipAddress field content, if present.

Resource Manager automatically adds information to incoming events that match a device. Fields added are:

- **prodState** - Specifies the device's current production state.
- **Location** - Specifies the location (if any) to which the device is assigned.
- **DeviceClass** - Classifies the device.
- **DeviceGroups** - Specifies the groups (if any) to which the device is assigned.
- **Systems** - Systems (if any) to which the device is assigned.
- **DevicePriority** - Priority assigned to the device.

For more information about these fields, refer to the chapters titled "Production States and Maintenance Windows" and "Organizers and Path Navigation."

eventState Field

The eventState field defines the current state of an event. This field is often updated after an event has been created. Values for this numeric field are 0-6, defined as follows:

| Number | Name | Description |
|--------|--------------|--|
| 0 | New | State given to an event when it is initially created in the system. |
| 1 | Acknowledged | State given to an event when a user has acknowledged the event. |
| 2 | Suppressed | State given to an event that has been suppressed via an event transform. |
| 3 | Closed | State given to an event that was closed as the result of a user action. |
| 4 | Cleared | State given to an event that was cleared by a corresponding clear event. |
| 5 | Dropped | State given to an event that was dropped via an event transform. These events are never persisted by the system. |
| 6 | Aged | State given to an event that was automatically closed by the system according to the severity and last seen time of the event. |

severity Field

The severity field defines the severity of the event. Values for this numeric field are 0-5, defined as follows:

| Number | Name | Color |
|--------|----------|--------|
| 0 | Clear | Green |
| 1 | Debug | Grey |
| 2 | Info | Blue |
| 3 | Warning | Yellow |
| 4 | Error | Orange |
| 5 | Critical | Red |

summary and message Fields

The summary and message fields are free-form text fields. The summary field allows up to 255 characters. The message field allows up to 4096 characters. These fields usually contain similar data.

The system handles these fields differently, depending on whether one or both are present on an incoming event:

- If only summary is present, then the system copies its contents into message and truncates summary contents to 128 characters.

- If only message is present, then the system copies its contents into summary and truncates summary contents to 128 characters.
- If summary and message are both present, then the system truncates summary contents to 128 characters.

As a result, data loss is possible only if the message or summary content exceeds 65535 characters, or if both fields are present and the summary content exceeds 128 characters.

To ensure that enough detail can be contained within the 128-character summary field limit, avoid reproducing information in the summary that exists on other fields (such as device, component, or severity).

Other Fields

Events include numerous other standard fields. Some control how an event is mapped and correlated; others provide information about the event.

The following table lists additional event fields.

| Field | Description |
|---------------|--|
| dedupid | Dynamically generated fingerprint that allows the system to perform de-duplication on repeating events that share similar characteristics. |
| component | Free-form text field (maximum 255 characters) that allows additional context to be given to events (for example, the interface name for an interface threshold event). |
| eventClass | Name of the event class into which this event has been created or mapped. |
| eventKey | Free-form text field (maximum 128 characters) that allows another specificity key to be used to drive the de-duplication and auto-clearing correlation process. |
| eventClassKey | Free-form text field (maximum 128 characters) that is used as the first step in mapping an unknown event into an event class. |
| eventGroup | Free-form text field (maximum 64 characters) that can be used to group similar types of events. This is primarily an extension point for customization. Currently not used in a standard system. |
| stateChange | Last time that any information about the event changed. |
| firstTime | First time that the event occurred. |
| lastTime | Most recent time that the event occurred. |
| count | Number of occurrences of the event between the firstTime and lastTime. |
| prodState | Production state of the device, updated when an event occurs. This value is not changed when a device's production state is changed; it always reflects the state when the event was received by the system. |
| agent | Typically the name of the daemon that generated the event. For example, an SNMP threshold event will have zenperfsnmp as its agent. |
| DeviceClass | Device class of the device that the event is related to. |
| Location | Location of the device that the event is related to. |
| Systems | Pipe-delimited list of systems that the device is contained within. |
| DeviceGroups | Pipe-delimited list of systems that the device is contained within. |
| facility | Only present on events coming from syslog. The syslog facility. |
| priority | Only present on events coming from syslog. The syslog priority. |

| Field | Description |
|-------------------|--|
| ntheid | Only present on events coming from Windows event log. The NT Event ID. |
| ownerid | Name of the user who acknowledged this event. |
| clearid | Only present on events in the archive that were auto-cleared. The evid of the event that cleared this one. |
| DevicePriority | Priority of the device that the event is related to. |
| eventClassMapping | If this event was matched by one of the configured event class mappings, contains the name of that mapping rule. |
| monitor | In a distributed setup, contains the name of the collector from which the event originated. |

Details

In addition to the standard fields, the system also allows events to add an arbitrary number of additional name/value pairs to events to give them more context.

De-Duplication

Resource Manager uses an event "de-duplication" feature, based on the concept of an event's fingerprint. Within the system, this fingerprint is the "dedupid." All of the standard events that the system creates as a result of its polling activities are de-duplicated, with no setup required. However, you can apply de-duplicating to events that arrive from other sources, such as syslog, SNMP traps, or a Windows event log.

The most important de-duplication concept is the *fingerprint*. An event's fingerprint (or dedupid) is composed of a pipe-delimited string that contains these event fields:

- device
- component (can be blank)
- eventClass
- eventKey (can be blank)
- severity
- summary (omitted from the dedupid if eventKey is non-blank)

When the component and eventKey fields are blank, a dedupid appears similar to:

```
www.example.com || /Status/Web || 4 | WebTx check failed
```

When the component and eventKey fields are present, a dedupid appears similar to:

```
router1.example.com | FastEthernet0/1 | /Perf/Interface | threshName
```

When a new event is received by the system, the dedupid is constructed. If it matches the dedupid for any active event, the existing event is updated with properties of the new event occurrence and the event's count is incremented by one, and the lastTime field is updated to be the created time of the new event occurrence. If it does not match the dedupid of any active events, then it is inserted into the active event table with a count of 1, and the firstTime and lastTime fields are set to the created time of the new event.

The following illustration depicts a de-duplication scenario in which an identical event occurs three times, followed by one that is different in a single aspect of the dedupid fingerprint.

Figure 92: Event De-Duplication

| | Event | dedupid | count |
|-------------------|--|--------------------------------|-------|
| Warning Threshold | device: router1 eventClass: /Perf/CPU eventKey: cpuThresh severity: 3 | router1//Perf/CPUcpuThreshold3 | 1 |
| Warning Threshold | | router1//Perf/CPUcpuThreshold3 | 2 |
| Warning Threshold | | router1//Perf/CPUcpuThreshold3 | 3 |
| Error Threshold | device: router1 eventClass: /Perf/CPU eventKey: cpuThresh severity: 4 | router1//Perf/CPUcpuThreshold4 | 1 |

If you want to change the way de-duplication behaves, you can use an event transform to alter one of the fields used to build the dedupid. You also can use a transform to directly modify the dedupid field, for more powerful cross-device event de-duplication.

Auto-Clear Correlation

The auto-clearing feature is similar to the de-duplication feature. It also is based on the event's fingerprint. The difference is which event fields make up the fingerprint, and what happens when a new event matches an existing event's fingerprint.

All of the standard events created as a result of polling activities do auto-clearing by themselves. As with de-duplication, you would invoke auto-clearing manually only to handle events that come from other sources, such as syslog, a Windows event log, or SNMP traps.

If a component has been identified for the event, then the auto-clear fingerprint consists of these fields:

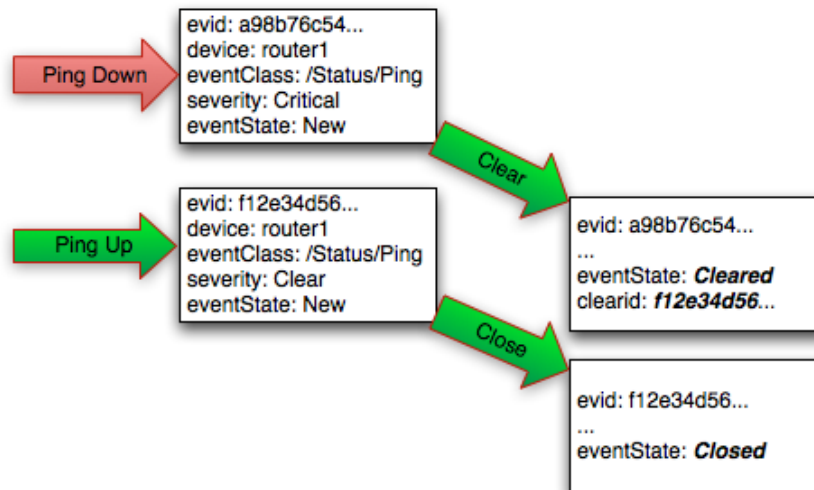
- If component UUID exists:
 - component UUID
 - eventClass (including zEventClearClasses from event class configuration properties)
 - eventKey (can be blank)
- If component UUID does not exist:
 - device
 - component (can be blank)
 - eventKey (can be blank)
 - eventClass (including zEventClearClasses from event class configuration properties)

When a new event comes into the system with a special 0 (Clear) severity, Resource Manager checks all active events to see if they match the auto-clear fingerprint of the new event. All active events that match the auto-clear fingerprint are updated with a Cleared state, and the clearid field is set to the UUID of the clear event. After a configurable period of time, all events in a closed state (Closed, Cleared, and Aged) are moved from the active events table to the event archive.

If an event is cleared by the clear event, it is also inserted into the active events table with a status of Closed; otherwise, it is dropped. This is done to prevent extraneous clear messages from filling your events database.

The following illustration depicts a standard ping down event and its associated clear event.

Figure 93: Event Auto-Clear



If you need to manually invoke the auto-clearing correlation system, you can use an event transform to make sure that the clear event has the 0 (Clear) severity set. You also need to ensure that the device, component, and eventClass fields match the events you intend to clear.

Note Avoid making clear events too generic; otherwise, you may inadvertently clear a wider range of events than you intend.

Event Consoles

Resource Manager features multiple event consoles that allow you to view and manage events. Each console shows different events subsets, depending on your current context.

Event consoles are:

- **Master**- To access this console, click Events on the Navigation menu. You can view all events from this console.
- **Contextual**- Contextual event consoles are found throughout the system. Each time you see an Events selection for a device, device organizer, component, or event class, you can view event information that has been automatically filtered to show events specific to the current context.

Master Event Console

The master event console is the system's central nervous system, enabling you to view and manage events. It displays the repository of all events that have been collected by the system.

Figure 94: Event Console

| Status | Severity | Resource | Component | Event Class | Summary | First Seen | Last Seen | Count |
|--------|----------|------------------------------|-----------|-------------|--|------------------------|------------------------|-------|
| ⚠ | ucs1 | Fan 3/4/2 | | | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |
| ⚠ | ucs1 | extpol_reg_clients_client... | | | UCS Domain ucs1-faba is registered with UCS Central without a valid license. | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| ⚠ | ucs1 | Fan Module 3/4 | | | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | 2015-07-29 08:45:41 am | 2015-07-29 08:45:41 am | 1 |
| ⚠ | ucs1 | fabric_san_A_phys-fcoes... | | | FCoE uplink port 1/8 is down | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| ⚠ | ucs1 | Fan 3/4/2 | | | Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |

Customizing the Event Console

You can add or delete data columns to customize your event console view. To configure the columns to be displayed on the Event Console:

- 1 Navigate to the Event Console (**Events > Event Console**).
- 2 Click the **Configure** button and select **Adjust columns** from the drop-down list. The Column Configuration window appears.

Figure 95: Event Console Column Configuration

Column Configuration

| Available | | Selected |
|---------------------|---|-------------|
| Agent | | Status |
| Cleared by Event ID | | Severity |
| Collector | | Resource |
| Device Class | | Component |
| Device Priority | | Event Class |
| Event Class Key | | Summary |
| Event Class Mapping | | First Seen |
| Event Group | | Last Seen |
| Event ID | | Count |
| Event Key | ↑ | |
| Fingerprint | ↑ | |
| Groups | → | |
| IP Address | → | |
| Location | ← | |
| Message | ← | |
| NT Event Code | ↓ | |
| Owner | ↓ | |
| Production State | | |
| State Change | | |
| Syslog Facility | | |
| Syslog Priority | | |
| Systems | | |

- 3 Move the names of the columns from the Available to the Selected column to display them. You can order the Selected column by using the arrow keys. The order of the selected column names determines the left-to-right display on the Event Console.
- 4 Click **Submit**.

Selecting Events

To select one or more events in the event console, you can:

- Click a row to select a single event
- Ctrl-Click rows to select multiple events, or Shift-Click to select a range of events

Sorting and Filtering Events

You can sort and filter events by any column that appears in the master event console.

To sort events, click a column header. Clicking the header toggles between ascending and descending sort order. Alternatively, hover over a column header to display its control, and then select Sort Ascending or Sort Descending.

Filter options appear below each column header.

Figure 96: Event Console Filter Options

| Status | Severity | Resource | Component | Event Class | Summary | First Seen | Last Seen | Count |
|-------------------------------------|--------------|----------|------------------------------|--|--|------------------------|------------------------|-------|
| <input checked="" type="checkbox"/> | New | ps31 | Fan 3/42 | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |
| <input checked="" type="checkbox"/> | Acknowledged | ps31 | extpol_reg_clients_client... | UCS Domain ucs1-faba is registered with UCS Central without a valid license. | UCS Domain ucs1-faba is registered with UCS Central without a valid license. | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| <input type="checkbox"/> | Suppressed | ps31 | Fan Module 3/4 | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable | 2015-07-29 08:44:33 am | 2015-07-29 08:45:41 am | 1 |
| <input type="checkbox"/> | Closed | ps31 | fabric_san_A_phys-foes... | FCoE uplink port 1/8 is down | FCoE uplink port 1/8 is down | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| <input type="checkbox"/> | Cleared | ps31 | Fan 3/42 | Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable | Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 3 |
| <input type="checkbox"/> | Aged | ps31 | Fabric_FC_SAN_Port-Chas... | san port-channel 5 on fabric interconnect A oper state: failed, reason: No operatio... | san port-channel 5 on fabric interconnect A oper state: failed, reason: No operatio... | 2015-07-29 08:44:33 am | 2015-07-29 10:29:40 am | 4 |
| | | ps31 | Fabric_FC_SAN_Port-Chas... | san Member 1/27 of Port-Channel 5 on fabric interconnect A is down, membership:... | san Member 1/27 of Port-Channel 5 on fabric interconnect A is down, membership:... | 2015-07-29 08:46:12 am | 2015-07-29 10:29:40 am | 4 |

You can filter the events that appear in the list in several ways, depending on the field type. Date fields (such as First Seen and Last Seen) allow you to enter a value or use a date selection tool to limit the list. For other fields, such as Device, Component, and Event Class, enter a match value to limit the list.

The Count field allows you to filter the list when compared to a value. To search on count:

- *N*- Displays events with a count equal to *N*.
- *:N*- displays events with a count less than or equal to *N*.
- *M:N*- Displays events with a count between *M* and *N*(inclusive).
- *M:-* Displays events with a count greater than or equal to *M*.

To clear filters, select **Configure > Clear filters**.

Working with Live Search

By default, the system uses a "live search" feature to help you locate information. From the event console, you can search for information by:

- **Device** (name) - Device name searches:
 - Are case-insensitive.
 - Are tokenized on whitespace (meaning that any searches that span whitespace and do not start with a complete token will return no results).
 - If quoted, return only exact matches.
- **Component** - Component searches:
 - Are case-insensitive.
 - Are tokenized on whitespace (meaning that any searches that span whitespace and do not start with a complete token will return no results).
 - If quoted, return only exact matches.
- **Summary**- Summary searches:

- Are case-insensitive.
- Are tokenized on whitespace (meaning that any searches that span whitespace and do not start with a complete token will return no results).
- **Event class** - Event class searches:
 - Are case-insensitive.
 - Are tokenized on / (slash). If the search begins with a slash, and ends with a slash or asterisk, then event classes are searched by using a "starts with" approach. If a search starts with a slash and ends with any other character, then event classes are searched by using an exact match for the event class. If a search does not begin with a slash, then event classes are searched by using a sub-string match on each event class.
- **IP Address** - IP address searches (for IPv4 and IPv6 values):
 - Are tokenized by . (period) and : (colon). For example, the following searches would return a result of 129.168.1.100:
 - 168
 - 168.1
 - 129.16*
 - *29
- **First Seen, Last Seen, State Change** - This field is not tokenized; date searches are converted to numeric representations, and then ranges using these representations are created. Search values are inclusive. Searches on date fields will search from the value entered. Any results that match the value or any value in the future are returned. The following searches would return the First Seen time of 2014-05-04 15:52:52:
 - First Seen: 2014-05-01 00:00:00
 - First Seen: 2014-05-04 15:52:52

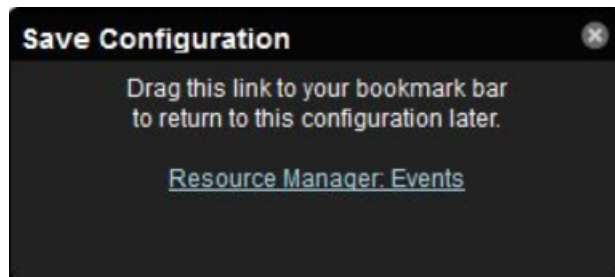
With live search enabled (the default behavior), the system filters available information immediately. It presents increasingly refined information with each character you type in the search window. When disabled, search responds only after you enter one or more characters and then press Enter.

Saving an Event Console View

You can save your event console view by bookmarking it for quick access later. Perform the following:

- 1 Select **Configure > Save this configuration**. A dialog containing a link to the current view appears.
- 2 Click and drag the link to the bookmarks link on your browser's menu bar. A link titled "Event Console" appears in your bookmarks list.

Figure 97: Saving a Custom View (Bookmark)




Note You may want to re-title the bookmark, particularly if you choose to save more than one event console view.

Refreshing the View

You can refresh the list of events manually or specify that they refresh automatically. To manually refresh the view, click **Refresh**. You can manually refresh at any time, even if you have an automatic refresh interval specified.

To set up automatic refresh, select one of the time increments from the Refresh list.

Figure 98: Automatic Refresh Selections



Viewing Event Details

You can view details for any event in the system. To view details, double-click an event row.

The Event Details area appears.

Figure 99: Event Details



To see more information about the event, click the Event Management, Device State, Event Data, or Event Details link. To display the event information in a new window, click the icon located at the top right.

You can use the Log area to add specific information about the event. Enter details, and then click **Add**.

Acknowledging Events

You may want to mark an event as "acknowledged" to indicate, for example, that you have taken action to remedy a problem. To mark events as acknowledged:

- 1 Select one or more events in the event console view.

- 2 Click the **Acknowledge Events** icon. A check mark appears for each acknowledged event.

Returning Events to New Status

You may want to return a previously acknowledged event to "new" status (revoke its "acknowledged" status). Perform the following:

- 1 Select one or more events in the event console view.
- 2 Click the **Unacknowledge Events** icon. A check mark no longer appears in the event row, and the event is returned to "new" status.

Classifying Events

Classifying events lets you associate events shown as /Unknown with a specific event class. To classify an unknown event, an event class key must be specified for the event.

To classify events:

- 1 Select one or more /Unknown events in the event console view.
- 2 Click the **Reclassify an Event** icon. The Classify Events dialog appears.
- 3 Select an event class from the list of options, and then click **Submit**.

Note You can also classify events from the event archive.

Closing Events

When you no longer want to actively monitor an event (after you acknowledge it, for example), you can specify to close the event and move it to the event archive according to a configured event archive interval. To do this:

- 1 Select one or more events in the event console view.
- 2 Click the **Close Events** icon. The selected events are closed and moved to the archive at the specified interval.

To view events in the event archive, select **EVENTS > Event Archive**.

Note Users with no assigned role can view all events in the archive.

- 3 Click the **Refresh** icon to update the event list.
The closed events are removed from the display in the event console view.

Reopening Events

You can reopen events in the active event console that are in the Closed, Cleared, or Aged state.

To reopen events:

- 1 Select one or more Closed, Cleared, or Aged events.
- 2 Click the **Reopen Events** icon. The selected events are returned to active status.

Note You cannot re-open a closed event if another active event with the same fingerprint exists. Before you can re-open the closed event, you must close the new event.

Exporting Event Data

You can export data from the event console to a comma-separated value (.csv) or XML file. You can select individual events (to export only those events), or make no selections (to export all events that match the current filter criteria).

To export events:

- 1 Select one or more events.
- 2 Select **Export > CSV** or **Export > XML**. By default, the exported file is named `events.Extension`.

Creating Events

To create events from the event console, click the **Add an Event** icon.

For more information about manual event creation, see the section titled "Creating Events Manually."

Event Sources

Events come into the system in two ways. *Generated events* are created as a result of active polling. *Captured events* are transmitted by external actions into the system.

Generated Events

These standard daemons are responsible for generating events. They automatically perform appropriate de-duplication and auto-clearing.

- **zenping** - Ping up/down events
- **zenstatus** - TCP port up/down events
- **zenperfsnmp** - SNMP agent up/down events, threshold events
- **zencommand** - Generic status events, threshold events
- **zenprocess** - Process up/down events, threshold events
- **zenwin** - Windows service up/down events

Captured Events

Captured events are those events that the system does not specifically know will occur in advance. De-duplication is performed on these events, but in some cases may need to be tuned. By default, no auto-clearing is done on captured events. Event transforms must be used to create the auto-clear correlations.

These standard daemons are responsible for collecting captured events:

- **zensyslog**- Events created from syslog messages.
- **zentrap**- Events created from SNMP traps and informs.

Any ZenPacks you install may optionally include their own daemons.

Creating Events Manually

You can manually create events. While this is not something you would do as part of normal system operation, it can be helpful when you are attempting to test mappings and transforms you have created.

Creating Events through the User Interface

To create events manually through the user interface:

- 1 Navigate to Events, then click the **Add an Event** icon. The Create Event dialog appears.

Figure 100: Create Event Dialog

- 2 Complete the basic event fields. Event class mappings are applied only for events that do not already have an event class.

Creating Events from the Command Line

To send events from the command line, use the `zensendevent` command.

- 1 Log in to the Control Center host as a user with `serviced` CLI privileges.
- 2 Attach to the `zenhub` service.

```
serviced service attach zenhub
```

- 3 Change to the `zenoss` user.

```
su - zenoss
```

- 4 Run the `zensendevent` command.

```
zensendevent Options "Event_Summary_Text"
```

Common options include:

- `-d DEVICE, --device=DEVICE`
- `-i IPADDRESS, --ipAddress=IPADDRESS`
- `-y EVENTKEY, --eventkey=EVENTKEY`
- `-p COMPONENT, --component=COMPONENT`
- `-k EVENTCLASSKEY, --eventclasskey=EVENTCLASSKEY`
- `-o OTHER, --other=OTHER`
- `-s SEVERITY, --severity=SEVERITY`
- `-c EVENTCLASS, --eventclass=EVENTCLASS`
- `-f INPUT_FILE, --file=INPUT_FILE`

Example

The following example shows how to use the `zensendevent` script to simulate a ping down event:

```
zensendevent -d router1.example.com -s Critical -c /Status/Ping "Router down"
```

Event Classes

Event classes are a simple organizational structure for the different types of events that the system generates and receives. This organization is useful for driving alerting and reporting. You can, for example, create an alerting rule that sends you an email or pages you when the availability of a Web site or page is affected by filtering on the `/Status/Web` event class.

Following is a subset of the default event classes. You can create additional event classes as needed.

- `/Status` - Used for events affecting availability.
 - `/Status/Ping` - Ping up/down events
 - `/Status/Snmp` - SNMP up/down events
 - `/Status/Web` - Web site or page up/down events
- `/Perf` - Used for performance threshold events.
 - `/Perf/CPU` - CPU utilization events
 - `/Perf/Memory` - Memory utilization or paging events
 - `/Perf/Interface` - Network interface utilization events
 - `/Perf/Filesystem` - File system usage events
- `/App` - Application-related events.
- `/Change` - Events created when the system finds changes in your environment.

Event Class Configuration Properties

Just as device classes and devices have configuration properties, so do event classes and event class mappings. Configuration properties are applied hierarchically, with the most specific property being applied.

The following configuration properties are available on event classes and class mappings.

- **zEventAction**- How and where affected events are stored when they occur.
 - **status**- Active events table with original event state
 - **history**- Active events table with closed event state
 - **drop**- Events are not stored
- **zEventClearClasses**- Optional list of event class names whose active events will be cleared by clear events occurring in this class.
- **zEventSeverity**- The severity of affected events is changed to this value unless the Default value (-1) is used.
- **zFlappingIntervalSeconds**- Defines the time interval to check for event flapping (changing severity level repeatedly). Default value is 3600 seconds.
- **zFlappingSeverity**- Define the severity to check for event flapping. If the severity level on an event changes from this value a certain number of times (`zFlappingThreshold`) within a certain time range (`zFlappingIntervalSeconds`) then an event flapping event is generated. Possible values include: 5-Critical, 4-Error, 3-Warning, 2-Info, 1-Debug, and 0-Clear.
- **zFlappingThreshold**- Number of times an event severity must flap within an interval. One of the parameters to define in order to generate event flapping events.

A good example of how the system uses the event class configuration properties is found in the `/Status` event class. Within the `/Status` event class' configuration properties, `zEventAction` is set to "history" and `zEventSeverity` is set to "Default". This means that events sent with this event class are sent into the active events table with an initial state of closed, and the event severity unchanged.

For more information about event manipulation techniques, see the section titled "Mapping and Transformation."

Mapping and Transformation

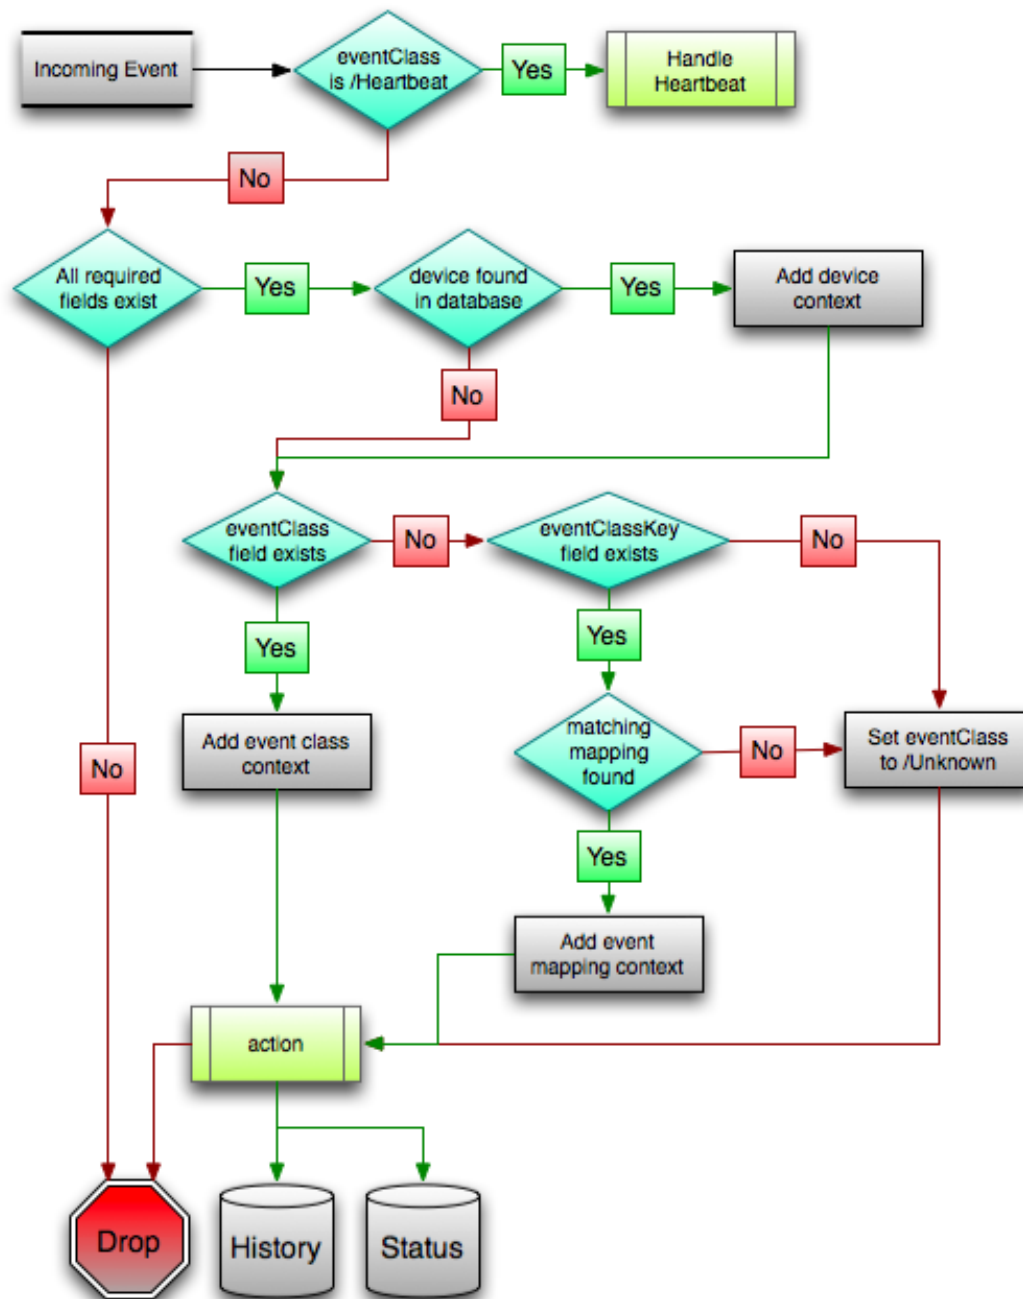
The event mapping and transformation system allows you to perform a wide range of operations, from altering the severity of certain events to altering nearly every field on an event, based on complex rules.

You cannot alter the following fields through event transformation. (This is because they are set after transformation has been performed.)

- evid
- firstTime
- lastTime
- count

The following illustration shows the path followed by an incoming event in the event mapping system.

Figure 101: Event Processing



The mapping and transformation process begins with the "eventClass field exists" decision. This also is one of the more important differentiators in how you must handle a particular type of event.

Event Class Mappings

To view event class mappings, select **EVENTS > Event Classes**, and then select **Mapping Instances** in the drop-down list. This allows you to see all event class mappings in a single location. The ID column shows the mapping's event class.

You can create event class mappings directly from the event classes, but this requires that you know the event class key. A simpler way to create event class mappings is through the event console:

- 1 Select an event that you want to match in the event console.
- 2 Click the **Reclassify an Event** icon. The Classify Events dialog appears.
- 3 Select the event class to which you want to map the event, and then click **Submit**. This creates the event class mapping with the correct event class key, and example text against which you can develop your regular expression.

When editing an event class mapping, you can control which events it will match, as well as other properties:

- **Matching tab**
 - **Event Class Key**- Must match the incoming event's Event Class Key field for this mapping to be considered as a match for events.
 - **Rule**- Provides a programmatic secondary match requirement. It takes a Python expression. If the expression evaluates to True for an event, this mapping is applied.
 - **Regex**- The regular expression match is used only in cases where the rule property is blank. It takes a Perl Compatible Regular Expression (PCRE). If the regex matches an event's message field, then this mapping is applied.
 - **Explanation**- Free-form text field that can be used to add an explanation field to any event that matches this mapping.
 - **Resolution**- Free-form text field that can be used to add a resolution field to any event that matches this mapping.
- **Transforms tab**- Takes Python code that will be executed on the event only if it matches this mapping. For more details on transforms, see the section titled "Event Class Transform."
- **Configuration Properties tab**- Listing of Configuration Properties defined for this event class.
- **Sequence tab**- Sequence number of this mapping. This number determines the order in which mappings with the same event class key are evaluated.

Mappings have the same configuration properties as event classes. Any configuration property set locally on a mapping will override the same property set on the event class. This works in the same hierarchical, most specific match, concept that device class and device configuration properties work.

When a captured event occurs, it will not have a pre-defined event class. For this type of event, you must create an event class mapping if you want to affect the event. If a captured event occurs and none of the event class mappings in the system match it, its event class will be set to /Unknown, and it will retain all of the default properties with which it began.

The next step of evaluation for events without an event class is to check the Event Class Key field. This controls which event class mapping the event will match. If the event has a blank event class key, or its event class key does not match any event class mappings in the system, the special "defaultmapping" event class key is searched for instead. This provides for a way to map events even if they have a blank or unpredictable event class key.

Event Class Mapping Sequence

The sequence area of an event class mapping (select Sequence in the left panel) allows you to provide more than one mapping for the same event class key. In this case, the sequence is evaluated in ascending order until a full (rule or regex) match is found.

For example, suppose a router is sending in unclassified events that need to be mapped to two event classes:

- /Events/Router/fanDown
- /Events/Router/fanUnknown

The event class key for both has been sent to "router", but one has a message of "Fan Down" and the other has no message at all. The mapping on /Events/Router/fanDown has an event class key of "router" and a regex of "Fan Down." The mapping on /Events/Router/fanUnknown has only an event class key of "router" and (in this

example) no regex. Because the fanUnknown mapping matches the fanDown events, the evaluation of fanDown needs to occur first.

You can modify the evaluation of mappings with the same event class key in the Sequence area of any of those event class mappings. In the previous example, you could go to either mapping, select Sequence, and both mappings would be displayed. You can set one to 0, and the other to 1. (You can enter other values, but they will be changed to the shortest list of integers, starting with 0.) Setting fanDown to 0 and fanUnknown to 1 will ensure that the events will be mapped properly.

Event Class Transform

When a generated event occurs, it has an event class assigned to it. This causes the event class mapping step to be skipped. The only way to affect the fields of one of these events is through the event class' configuration properties and transform.

To access the transform for an event class:

- 1 Navigate to the event class from **Events > Event Classes**.
- 2 From the drop-down list, select **Transforms**.
- 3 Enter information into the dialog (as Python code), and then click the **Save** button in the upper-right corner. As you develop your transform, you can revert back to the last saved state by clicking the **Revert this Transform** button.

The objects available in this Python context are `evt` (the event); and, if the event matches a device that exists in the system database, a `device` object.

The following example shows how you can validate that a device object exists before using it to drop events from a particular location.

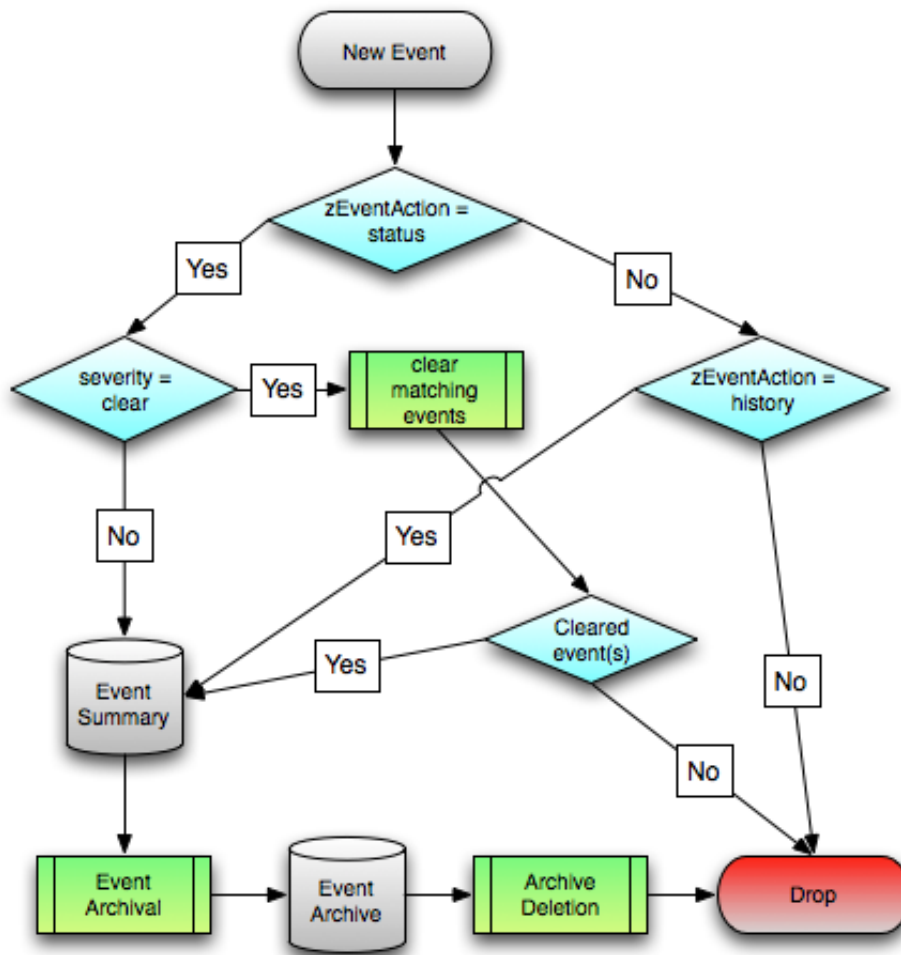
```
if device and "Hawaii" in device.getLocationName(): evt._action = "drop"
```

Event Life Cycle

In addition to manual methods for getting events into the status table or event archive, there are automated processes that move events from status into the archive. The *event life cycle* is defined as all of the ways that events can be added to, moved in, and deleted from the database.

The following illustration depicts the event life cycle.

Figure 102: Event Life Cycle



Automatic Event Aging

From the Event Configuration page (**ADVANCED > Settings > Events**), you can set up automatic aging of events. Aging of events will automatically update active events that match the severity and aging threshold to a status of Aged. After the configured event archive interval, all Closed, Aged, and Cleared events are moved to the event archive.

Properties that control this behavior are:

- **Don't Age This Severity and Above** - Options are Age All Events, Critical, Error, Warning, Info, Debug, and Clear. By default, this value is set to Error, meaning that all events with a status of Error or Critical are not aged.
- **Event Aging Threshold (minutes)** - Set the time value, in minutes, that an event must reach before it is aged. By default, this is 240 minutes.
- **Event Aging Interval (milliseconds)** - The interval when events are scanned to perform autoaging. By default, this is 60000 milliseconds (60 sec).
- **Event Aging Limit** - The maximum number of events to age in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- **Event Archive Threshold (minutes)** - Specify the number of minutes since a closed event was last seen before it is moved to the event archive. The minimum value is 1; the maximum value is 43200.
- **Event Archive Interval (milliseconds)** - The interval when events are scanned for moving to the archive. By default, this is 60000 milliseconds (60 sec).

- **Event Archive Limit** - The maximum number of events to archive in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- **Delete Archived Events Older Than (days)** - The number of days that events in the event archive are saved. By default, they are kept in the archive for 90 days. The minimum value is 1 and the maximum value is determined by the range of event archive partitions. With the default configuration, the maximum value is 1000 days.
- **Default Syslog Priority** - Specify the default severity level assigned to an event coming from zensyslog if no priority can be determined from the event.
- **Default Availability Report (days)** - Enter the number of days to include in the automatically generated Availability Report. This report shows a graphical summary of availability and status.
- **Max Event Size in Bytes** - The maximum size of an event that will be processed in bytes. Events that are too large will be logged and dropped. Events that will become too big will have their details overwritten with new details. By default, this is 32768 bytes.
- **Summary Index Interval (milliseconds)** - The default indexing interval of the event summary in milliseconds. By default, this is 1000 milliseconds (1 sec).
- **Archive Index Interval (milliseconds)** - The default indexing interval of the event archive in milliseconds. By default, this is 30000 milliseconds (30 sec).
- **Index Limit** - The number of events to index in each index interval. By default, this is 1000 events.
- **Event Time Purge Interval (days)** - The number of days that event occurrence time are kept. By default, they are kept for 7 days. The minimum value is 1 and the maximum value is determined by the range of event time partitions. With the default configuration, the maximum value is 7 days.
- **Enable Event Flapping Detection** - Select this check box if you wish to enable event flapping detection. If an event is created and then cleared *flapping_threshold* times in *event_flapping_interval* time then an event of event flapping event class is created.
- **Event Flapping Event Class** - The event class under which generated flapping events belong.
- **Clear Event Heartbeats** - Click **Clear** to clear the event heartbeats.

Automatic Archived Event Cleanup

You can set up automatic purging of events from the event archive from the Event Configuration page (**ADVANCED > Settings > Events**). When events are purged, they can be recovered only from backups.

The property that controls this behavior is Delete Archived Events Older Than (days). Acceptable values are between 1 and 1000 (days).

Capturing Email Messages as Events

ZenMail and ZenPop allow you to capture email messages as events. This capability can be useful for situations in which embedded systems (such as WAPs, NAS devices, or RAID controllers) rely on email notification for events.

ZenMail

ZenMail serves as an SMTP server that you can bind to a specific TCP port. You can then configure your embedded system to send mail to the Resource Manager server explicitly by using the server's IP address as the relay.

ZenMail supports these configuration directives:

- `${ZENHOME}/bin/zenmail` (no arguments) - Default operation. Binds to port 25 on all ports and listens for email messages to arrive. Ignores the TO field in the email and uses the FROM address as the device IP address.
- `${ZENHOME}/bin/zenmail --listenPort` - Bind to the port provided. Useful in situations in which an SMTP server is already running on the Resource Manager server and you do not want to interfere with the existing mail delivery system. Semantics are the same as the no argument version (FROM address is used as the device IP).

Note In order to execute a command using `$ZENHOME (/opt/zenoss` for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for `serviced` commands.

ZenPop3

ZenPop3 allows you to retrieve event email from a POP server. ZenPop3 supports these configuration directives:

- `--usessl`- Issue the STARTTLS command to the POP server and attempt to transfer email messages using SSL encryption. This is required if retrieving mail from Google.
- `--nodelete`- Do not issue the DELE command after retrieving all messages. Typically this is used during initial testing so that you do not have to resend test messages to the POP account. Some email systems (such as Google) do not actually delete messages when the DELE command is issued.
- `--pophost`- The hostname or IP address of the POP server from which to retrieve messages.
- `--popport`- The TCP port the POP server listens on. Defaults to 110. Used in situations where the POP provider listens on another port (for example, Google on port 995).
- `--popuser`- The user name that contains email messages to retrieve.
- `--poppass`- The password to use for the user name provided.
- `--cycletime`- The time to sleep between polls. After all email is retrieved, ZenPop3 sleeps for this amount of time before waking up and attempting to pull new email.

Translating Message Elements to the Event

Resource Manager translates various message elements to the event, as follows:

- **FROM Field**- If the FROM field is an IP address, then the system associates the event with the device with the same IP address. If the FROM field is a fully qualified domain name, then the system resolves it to an IP address, and then performs the device association using the resolved IP address. The resolution of hostname uses "A" records rather than "MX" records.
- **TO Field**- The system ignores the TO field in the email message. ZenMail accepts email to any user and domain name combination. ZenPop also drops the TO field, and uses only the FROM field.
- **SUBJECT Field**- ZenMail and ZenPop use the SUBJECT as the event summary.
- **Message Body**- ZenMail and ZenPop use the first mime attachment as the event details. The system ignores secondary message bodies (typically HTML-encoded versions of the message). It also ignores attachments (such as files).

SNMP Traps and Event Transforms

An SNMP trap is a message that is initiated by a network element and sent to the network management system. Often, traps indicate a failure of some sort, such as a router message indicating a power supply failure, or a printer message indicating an "out-of-ink" condition.

If an SNMP trap enters the system, and Resource Manager cannot identify the event (the event is classified as "/Unknown"), then you can classify the event so that the system handles it consistently.

Classifying SNMP Traps

To classify an SNMP trap event:

- 1 From the Event Console, select the unknown event or events.
- 2 Click the **Reclassify an event** icon. The Classify Events dialog appears.
- 3 Select /App, and then click **Submit**.

To edit this classification:

- a From the Navigation area, select **Events > Event Classes**.
- b Ensure Mapping Instances is displayed.
- c Select the event map you created.
- d In the left panel, select **Edit** from the **Action** icon.

The edit page appears. This page contains rules used to map the event to the /App category. This rule, since it matches the trap by a specific OID, is all that is needed.

In the Transform area, you can enter code to modify the summary. For example, if you want to set the summary string to "Spam Filter Detects Virus," then you can enter:

```
evt.summary = "Spam Filter Detects Virus"
```

A trap has a header with some standard information, followed by a sequence of attribute/values.

You have indicated you want the value for the OID ".1.3.6.1.4.1.9789.1500.2.5" as the summary. If you had the MIB loaded, you could do this:

```
evt.summary = evt.spamFilterDetectsVirus
```

However, the OID and the data is still in there. Instead, use the slightly more cryptic:

```
evt.summary = getattr(evt, ".1.3.6.1.4.9789.1500.2.5", "Unexpected missing OID")
```

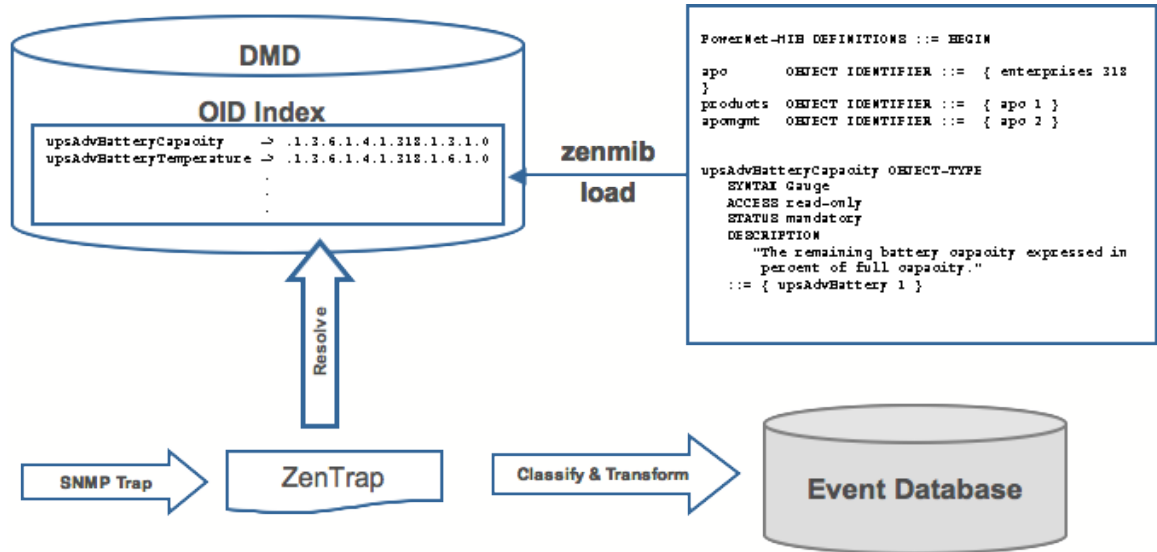
The "device" object for the event has been made available, as well:

```
evt.summary = getattr(evt, ".1.3.6.1.4.9789.1500.2.5", "Unexpected missing OID") \ + " from device " + device.getId()
```

Resource Manager uses MIBs to translate SNMP traps that contain raw OID values. Loading a MIB into the system allows it to translate numeric OIDs such as .1.3.6.1.2.1.1.6 into descriptive phrases like "sysLocation". It also makes it easier to manipulate the events in an event mapping.

Note After loading MIBs into the system, be sure to restart the `zentrap` service so that it can retrieve the new MIB information.

Figure 103: SNMP TRAP Transform



Following is a small demonstration MIB.

```

NOTIFICATION-TEST-MIB DEFINITIONS ::= BEGIN
IMPORTS
ucdavis FROM UCD-SNMP-MIB
NOTIFICATION-TYPE FROM SNMPv2-SMI
;
demonotifs OBJECT IDENTIFIER
::= { ucdavis 991 }
demo-notif NOTIFICATION-TYPE
OBJECTS { sysLocation }
STATUS current
DESCRIPTION "Just a test notification"
::= { demonotifs 17 }
END

```

Example: Sending Test Traps

To send an SNMP trap:

- 1 From the command line, enter the following command:

```

$ snmptrap -v 2c -c public localhost '
1.3.6.1.4.1.2021.991.1.3.6.1.2.1.1.6 s \ "Device in Austin"

```

- 2 Save this demonstration MIB into a file.
- 3 Send the trap.
- 4 Open the Event Console and find the trap you sent.
- 5 Send this event to the event archive.
- 6 Load some MIBs into the system so that this OID is translated into a better format:
 - a Copy the demonstration MIB into `$ZENHOME/share/mibs/site`.

Note In order to execute a command using `$ZENHOME` (`/opt/zenoss` for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for `serviced` commands.

- b** Run zenmib to load it:

```
$ zenmib run -v 10 DEBUG:zen.zenmib:TRAP-TEST-MIB.mib
INFO:zen.zenmib:Unable to find a file \ providing the MIB UCD-SNMP-
MIB ...
```

- c** The MIB loaded, but is missing some other definitions. Copy them:

```
$ cp /usr/share/snmp/mibs/SNMPv2-MIB.txt $ZENHOME/share/mibs/site \
$ cp /usr/share/snmp/mibs/UCD-SNMP-MIB.txt $ZENHOME/share/mibs/site
```

- d** Run zenmib again and load the definitions into the system:

```
$ zenmib run -v 10
```

- e** Restart the zentrap daemon to retrieve the new MIB information:

```
$ zentrap restart
```

- f** Send the trap a second time:

```
$ snmptrap -v 2c -c public localhost '
1.3.6.1.4.1.2021.13.991 .1.3.6.1.2.1.1.6s \ "Device in Austin"
```

- g** Check the event. Make sure the count is 1. If the count is 2, send the event to the event archive and send the trap again. Look at the Details tab. Now you should see something like this:

```
sysLocation Device in Austin
```

You should also see that the event summary changes from:

```
snmp trap 1.3.6.1.4.1.2021.13.991 from localhost
```

to:

```
snmp trap ucdExperimental from localhost
```

Transforming Events with Event Mappings

To modify events as they arrive, create an event map through the user interface:

- 1 Create an event class.
- 2 Go to the event console and create an event mapping in this class from the existing event.
- 3 Edit the map.
- 4 In the Transform area, update the event with detail data. The entry field allows you to insert Python scripts. The event is provided as "evt" and the device as "dev." In this case, extract the sysLocation event detail and make it the summary with:

```
evt.summary = evt.sysLocation
```

- 5 Save the event mapping.

If you move the event to the event archive and resend the trap, the summary for the trap should now read the device name in the location you assigned.

If you encounter problems with the transform, check the `zentrap.log` file for errors that occurred.

Event Transforms Based on Event Class

When an event arrives in the system, you can change values (such as severity). For example, you can make the summary more informative, or change severity according to text within the summary.

Each event class allows for a short Python script to be executed when an event arrives.

Example

A user may want full file system threshold events on `/data` to be critical. Add the following Python script in the Threshold Transform of `/Events/Perf/Filesystem`:

```
if evt.component == '/data' and evt.severity != 0: evt.severity = 5
```

Like event mappings for event class keys, "evt" and "dev" objects are available in the script of the transform.

SNMP Trap filtering at the collector level

zentrap supports SNMP trap filtering. For details on the filter definition, see the `/opt/zenoss/etc/zentrap.filter.conf` file in the `zentrap` service definition. For details on how to enable this feature see the following KB article [How to setup a filter for SNMP Traps at the collector level \(zentrap\) in version 5.1.1](#).

Production States and Maintenance Windows

10

Production state determines the level of monitoring and alerting applied to an individual device. Typically, alerting rules specify that the system will monitor and create events for devices that are in the "Production" production state.

Maintenance windows are planned time periods used to temporarily modify alerting rules so that event-generated alerts are temporarily halted during the window.

Production States

Production state determines whether a device is monitored, and can be used to control several elements of the event system, such as whether an event will produce a remote alert (email or page).

Choose a production state for a device based on whether you want:

- The device to be monitored
- The device to appear on the dashboard
- Alerting to occur

The following table lists production states and their characteristics.

| Production State | Devices Monitored? | Appear on Dashboard? |
|------------------|--------------------|----------------------|
| Production | yes | yes |
| Pre-Production | yes | no |
| Test | yes | no |
| Maintenance | yes | may appear |
| Decommissioned | no | no |

When you add a device to the system, its default state is Production. You may want to add triggers and notifications to alert you to various conditions that occur in the system, such as production state changes or a severity level being reached. For example, you can set up a trigger when a device is in either a production or a maintenance state and has a severity of Error or higher. You can then notify users when this trigger condition is met. For more information, see [Working with Triggers](#) on page 23.

Setting the Production State for Devices

To set the production state for a device:

- 1 Click a device name in the list of devices. The Device Overview page appears.

- 2 Select a production state from the list of options, and then click **Save**.
- 3 Optional: To set the production state for a group of devices, perform the following:
 - a Select a category of devices (by class, group, system, or location) from the hierarchy.
 - b From the list of options in the Production State column, select a production state. The newly selected state is applied to all devices that appear in the list.

Figure 104: Select Production State (Multiple Devices)

| Device | IP Address | Device Class | Production State | Events |
|-------------------------|---------------|-----------------------|------------------|--------|
| ucs-central13 | 10.87.208.139 | /CiscoUCS/UCS-Central | Production | 1 |
| ucs-central132 | 10.87.209.144 | /CiscoUCS/UCS-Central | Pre-Production | 10 |
| aus-ucs11 | 10.87.208.11 | /CiscoUCS/UCS-Manager | Test | 9 |
| aus-ucs14 | 10.87.208.14 | /CiscoUCS/UCS-Manager | Maintenance | 16 |
| aus-ucs20 | 10.87.208.20 | /CiscoUCS/UCS-Manager | Decommissioned | 1 |
| ucs-mini3 | 10.87.209.202 | /CiscoUCS/UCS-Manager | Production | 3 |
| ucs1 | 10.87.208.163 | /CiscoUCS/UCS-Manager | Production | 2 |
| ucspm-stable_zenoss_loc | 10.87.208.152 | /ControlCenter | Production | |

Maintenance Windows

Maintenance windows allow scheduled production state changes of a device or all devices in a system, group, or location. You might want to set up a maintenance window, for example, to change a device's production state while you perform configuration changes or reboot a device.

Note In lieu of setting up a maintenance window, you can change the production state for a device manually at the time you want to make changes.

When the maintenance window starts, the production state of the device is set to the value of Start Production State (for example, Maintenance). When the maintenance window closes, the production state of the device reverts to the value of Stop Production State (the state the device was in prior to Maintenance).

Maintenance windows do not prevent notifications from being triggered on the device. If you want to define the notifications you receive during the maintenance window, you will need to set up an appropriate trigger for the device production state that you set during your maintenance window. For more information, see [Working with Triggers](#) on page 23.

Maintenance Window Events

When a maintenance window starts, an event is created with the following information:

- depuid - zenactions | Resource | MaintenanceWindowName | TargetOrganizerOrDevice
- prodState - StartProductionState
- severity - Info
- summary/message - Maintenance window starting MaintenanceWindowName for TargetOrganizerOrDevice
- eventClass - /Status/Update
- eventClassKey - mw_change
- maintenance_devices - TargetOrganizerOrDevice
- maintenance_window - MaintenanceWindowName

When a maintenance window stops, an event is created with the following information:

- severity - Clear
- summary/message - Maintenance window stopping *MaintenanceWindowName* for *TargetOrganizerOrDevice*
- prodState - -99 (meaning "unknown.")

Maintenance window events auto-clear, meaning that stop events clear start events.

Creating and Using Maintenance Windows

You can create a maintenance window for an individual device or group of devices (all devices, a device class, group, system, or location) in the devices hierarchy.

Create a Maintenance Window for a Single Device

Use this procedure to create a maintenance window for a device.

- 1 Log in to the Resource Manager browser interface, and then navigate to **INFRASTRUCTURE > Devices**.
- 2 In the content area, click the name of the device.
- 3 In the sidebar, click **Device Administration**.
- 4 In the Maintenance Window toolbar, click **Add**.
- 5 In the **Add New Maintenance Window** dialog, specify the attributes of the new maintenance window.
The attributes allow you to specify whether or not to enable the window, the time of day and date when the maintenance window starts, the duration of the window, and whether and how the window repeats.
The **Window Production State** field allows you to categorize the production state of the device during the maintenance window. When the window ends, the device is returned to the production state it was in when it entered the maintenance window.
- 6 At the bottom of the **Add New Maintenance Window** dialog, click **SUBMIT**.

Create a Maintenance Window for a Group of Devices


Use this procedure to create a maintenance window for a group of devices.

- 1 Log in to the Resource Manager browser interface, and then navigate to **INFRASTRUCTURE > Devices**.
- 2 In the content area, select a group of devices, and then click **DETAILS**, located at the top of the sidebar.
- 3 In the sidebar, click **Device Administration**.
- 4 In the Maintenance Window toolbar, click **Add**.
- 5 In the **Add New Maintenance Window** dialog, specify the attributes of the new maintenance window.
The attributes allow you to specify whether or not to enable the window, the time of day and date when the maintenance window starts, the duration of the window, and whether and how the window repeats.
The **Window Production State** field allows you to categorize the production state of the device during the maintenance window. When the window ends, the device is returned to the production state it was in when it entered the maintenance window.
- 6 At the bottom of the **Add New Maintenance Window** dialog, click **SUBMIT**.

Managing Maintenance Windows

Once you have created maintenance windows for your devices or groups of devices, you can quickly manage these instances on the Maintenance Windows screen.

- 1 Navigate to the Maintenance Window screen. This is the same place where you initially created the maintenance window (Device Administration link on Device Overview page). On this screen you can perform any of the following by clicking the appropriate icon:



- Add a new maintenance window
 - Delete the selected maintenance window
 - Edit the selected maintenance window (can also double-click a maintenance window row)
 - Toggle the selected maintenance window from enabled to disabled and vice-versa. The Enabled column will switch values.
- 2 Ensure that your changes are reflected in the Maintenance Window screen.

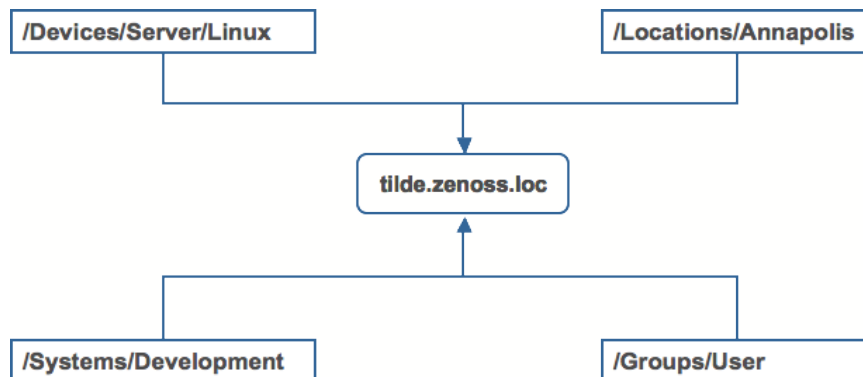
Organizers and Path Navigation

You can group system objects, including devices, sub-systems, components, configuration properties, and templates. A device, for example, can belong to multiple classifications, including:

- Device classes
- Groups
- Systems
- Locations

In the following illustration, the device `tilde.zenoss.loc` belongs to five different classifications. Configuration properties and monitoring settings for each of these groups are applied to this device.

Figure 105: Device Groupings



Classes

The most important organizers are *classes*, which comprise:

- Device classes
- Event classes
- Service classes
- Product classes

Templates and configuration properties can be inherited based on class. These attributes can be overwritten further down the class hierarchy, all the way down to the individual component level. The class hierarchy includes all defined and standard classes and sub-classes.


The following procedures are illustrated using device classes and sub-classes, but the same concepts apply to event classes, service classes, and product classes. When you add a device to the system, you should (after providing the network name or IP address) specify its device class. Templates and configuration properties can be set at any level in the device class hierarchy.

Viewing Device Classes

To view device classes and the devices they contain, select **INFRASTRUCTURE** from the Navigation menu.

The device list appears. The top of the devices hierarchy lists device classes. Click a class name to view devices in that class, or expand it to show sub-classes.

Figure 106: Devices Hierarchy



An indicator appears next to each listed class to show whether there are events associated with any devices in that class.

Adding Classes

To create a device class:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.
- 2 Select the parent location in the device class hierarchy where you want to add a new child class.
- 3 Click the **Add** icon. The Add Device Class dialog appears.

- 4 Enter a name and description for the new device class, and then click **Submit**. The new device class appears in the hierarchy under the parent device class. You can now move devices into this class. Select one or more devices in the device list, and then drag them to the new class.

Moving a Class

To move a class in the hierarchy:

- 1 Select the class in the hierarchy.
- 2 Drag the class to its new location. The Move Organizer confirmation dialog appears.
- 3 Click **OK** to confirm the action. The class appears at its new location in the hierarchy.

Setting Configuration Properties at the Class Level

To set configuration properties at the device class level:

- 1 Select a device class in the devices hierarchy.
- 2 Click **Details > Configuration Properties**. The Configuration Properties page for the selected device class appears.

Figure 107: Device Class Configuration Properties

| Is Local | Category | Name | Value | Path |
|----------|------------------|-----------------------------|--------------------|------|
| | Cisco | zCiscoACEUseSSL | true | / |
| | Cisco | zCiscoRemodeEventClassKeys | | / |
| | Cisco UCS | zCiscoUCSCIMCEventsInterval | 60 | / |
| | Cisco UCS | zCiscoUCSCIMCPerInterval | 300 | / |
| | Cisco UCS | zCiscoUCSMManagerPassword | ***** | / |
| | Cisco UCS | zCiscoUCSMManagerPort | 443 | / |
| | Cisco UCS | zCiscoUCSMManagerUseSSL | true | / |
| | Cisco UCS | zCiscoUCSMManagerUser | admin | / |
| | Modeler Controls | zCollectorClientTimeout | 180 | / |
| | Modeler Controls | zCollectorDecoding | utf-8 | / |
| | Misc | zCollectorLogChanges | false | / |
| | zencommand | zCommandComimandTimeout | 15 | / |
| | zencommand | zCommandExistanceTest | test -f %s | / |
| | zencommand | zCommandLoginTimeout | 10 | / |
| | zencommand | zCommandLoginTries | 1 | / |
| | zencommand | zCommandPassword | | / |
| | zencommand | zCommandPath | \$ZENHOME/libexec | / |
| | zencommand | zCommandPort | 22 | / |
| | zencommand | zCommandProtocol | ssh | / |
| | zencommand | zCommandSearchPath | | / |
| | zencommand | zCommandUsername | | / |
| | Control Center | zControlCenterHost | \$(here/manager/p) | / |

- 3 Edit configuration properties definitions as desired, and then click **Save**. Definitions are applied to all devices currently in, and added to, this class (unless overridden at a lower level in the hierarchy).

Groups

Groups are functional divisions that allow you to assign attributes to multiple objects with similar functions. Groups can be used, for example, to arrange objects along departmental lines. Groups do not appear on the Dashboard.

Adding Groups

To add a group:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.
- 2 Select the parent location in the groups hierarchy where you want to create a child group.

- 3 Click the **Add** icon. The Add Group dialog appears.
- 4 Enter a name and description for the group, and then click **Submit**. The group appears in the hierarchy. You can now move devices to this group. Select one or more devices in the device list, and then drag them to the new group.

Moving a Group

To move a group:

- 1 Select the group in the hierarchy.
- 2 Drag the group to its new location. The Move Organizer confirmation dialog appears.
- 3 Click **OK** to confirm the action. The group appears at its new location in the hierarchy.

Systems

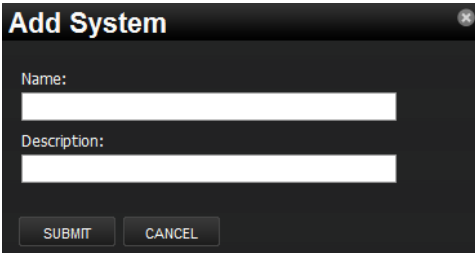
Systems are intended to follow virtual setups, such as those in a network setup or systems grouped by functionality.

Adding Systems

To add a system or sub-system:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.
- 2 Select the parent location in the systems hierarchy where you want to create a new child system.
- 3 Click the **Add** icon. The Add System dialog appears.

Figure 108: Add System



- 4 Enter a name and description for the system, and then click **Submit**. The system appears in the hierarchy. You can now move devices to this system. Select one or more devices in the device list, and then drag them to the new system.

Moving a System

To move a system:

- 1 Select the system in the hierarchy.
- 2 Drag the system to its new location. The Move Organizer confirmation dialog appears.
- 3 Click **OK** to confirm the action. The system appears at its new location in the hierarchy.

Locations

Locations are logical groupings for physical systems. They can be as general as city and state, or as specific as rack or closet. Locations do not appear on the Dashboard.

Adding Locations

To add a location:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.

- 2 Select the parent level in the Locations hierarchy where you want to create a child location.
- 3 Click the **Add** icon. The Add Location dialog appears.
- 4 Enter a name and description for the location, and then click **Submit**. The location appears in the hierarchy. You can now move devices to this location. Select one or more devices in the device list, and then drag them to the new location.

Moving a Location

To move a location:

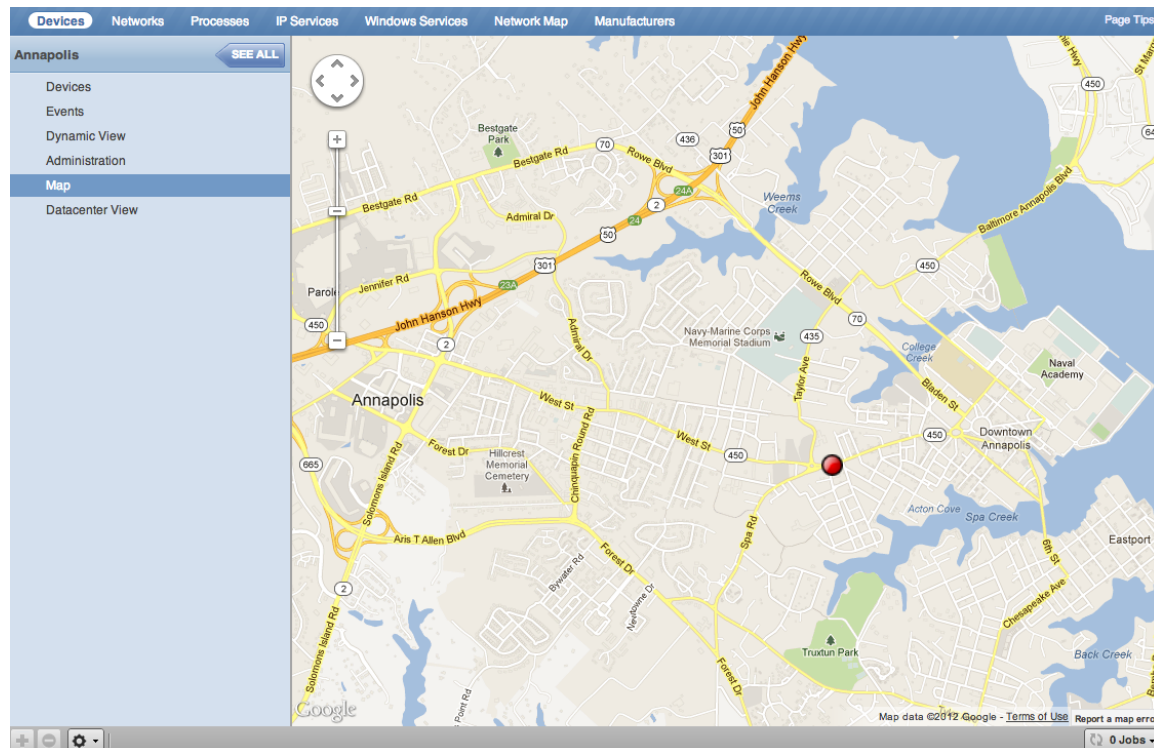
- 1 Select the location in the hierarchy.
- 2 Drag the location to its new place. The Move Organizer confirmation dialog appears.
- 3 Click **OK** to confirm the action.

Integration with Google Maps

The system can map locations by using a Google Maps mashup feature that sets the location's Address property to a valid Google Maps address. The selected location appears on the map as a dot. The color of the dot represents the highest severity of any event on any device in that location.

Network connections that span locations are represented on the map by lines. Each line color matches the status of the connection it represents.

Figure 109: Network Location - Google Map



To view the Google Map for a network location:

- 1 From the Navigation menu, select Infrastructure.
 - The device list appears.
- 2 In the hierarchy, select a location.
- 3 Click **Details**.

4 Select **Map**.

The network map appears.

Note You also can view the network location map from the Google Maps portlet on the Dashboard.

Enabling Google Maps on Your System

Before you can use the Google Maps feature, you must register your Resource Manager server key (UUID) with Zenoss. Typically, this is done when you license the product.

You can verify whether your server key is registered. From the Resource Manager interface, select **Advanced > Licensing**. If the Licensing Report indicates that no license is associated with your server key, contact Zenoss Support at the [Zenoss Support](#) site.

Be sure to provide your Resource Manager server key; this is required to generate the Google Maps API key.

Setting an Address for a Location

Follow these steps to set a location address:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.
- 2 Select a location in the hierarchy.
- 3 Select **Edit** from the **Action** icon. The Edit Organizer dialog appears.
- 4 In the Description field, enter a description for the location.
- 5 In the Address field, enter a complete address that can be resolved by [Google Maps](#). The address must include a valid zip code.
- 6 Click **Submit**. The selected address for the location is created. You must add at least one device to the location for the location "dot" to appear on the map.

Clearing the Google Maps Cache

Sometimes there are issues with drawing the maps and seeing the network status of locations or connections. Clearing the Geocode cache will solve these problems. To clear the Geocode cache:

- 1 Select **Infrastructure** from the Navigation menu. The device list appears.
- 2 Select a location in the hierarchy.
- 3 Select **Clear Geocode Cache** from the **Action** icon. A confirmation dialog appears.
- 4 Click **OK**.

Network Links

If two devices in the same network are in different map-able locations, a line is on the map representing a network connection between the two. If there are multiple separate network connections between the same two locations, only one line is drawn. The color of the line represents the highest severity of any events affecting the connection. These are determined by:

- A ping down event on the device at either end of the connection; or
- Any event on the interface at either end of the connection.

Drawing Map Links (zDrawMapLinks Configuration Property)

Calculating network links on the fly is an time-intensive procedure. If you have a large number of devices that have been assigned locations, drawing those links on the map may take a long time.

To save time, you can tell the system not to attempt to draw links for specific networks. You might want to do this, for example, for a local network comprising many devices that you know does not span multiple locations.

To edit the value for this property:

- 1 Select **Infrastructure > Networks** from the Navigation menu. The Networks page appears.
- 2 Select a network or sub-network for which you want to disable map links.
- 3 Display configuration properties for the network.
- 4 Double-click the `zDrawMapLinks` configuration property in the list. The Edit Config Property dialog appears.
- 5 De-select the value (uncheck the box), and then click **Submit**.

Note This setting will be inherited by networks or sub-networks below this selection in the hierarchy. If you have few networks for which links would be drawn, you might want to disable map links on `/Networks`, enabling it only on a network where you know a location-spanning WAN connection exists.

Google Maps Example

This example will show you how to:

- Create and display Google map links of devices
 - Send a test event to see how map links are affected by system changes
- 1 Disable map links. (Refer to the procedure in Drawing Map Links for instructions.)
 - 2 Create two locations: "New York" and "Los Angeles." (Refer to the procedure in Adding Locations for instructions.)
 - 3 Enter the following values in the Address field of the Add Location dialog: "New York, NY" and "Los Angeles, CA" respectively.
 - 4 Set the location of a device to New York. Locate another device on the same network and set its location to Los Angeles.
 - 5 Select Locations in the hierarchy, click **Details**, and then select Map.

New York and Los Angeles are represented by dots on the map; however, no link is drawn between these locations.

- 6 Select Networks and re-enable map linking.
- 7 Select Infrastructure, then select Locations in the hierarchy.
- 8 Click **Details**, and then select Map.

A green line is now drawn between New York and Los Angeles.

- 9 Send an event with a severity of Critical to the device in New York. (For information about creating events, refer to the chapter titled Event Management.) Do not specify a component.
- 10 Return to the Locations map.

The dot representing New York is now red, but the link between New York and Los Angeles remains green.

- 11 Navigate to the New York device and determine the ID of the component that is connected to the network shared with the Los Angeles device.
- 12 Send another test event, this time specifying that component.
- 13 Return to the Locations map.

The link between New York and Los Angeles is now red.

Component Groups

Use Component Groups to view or manage component resources as a logical group. When you add components to a component group, you can view all of the events for the group in a single location, display component graphs with all of the components on the same graph, and use the Dynamic and Dependencies views to see the group's dependents and dependencies. You can also monitor or lock individual components in the group or the entire group itself.

Creating and Viewing Component Group Information

To create and view information about component groups:

- 1 Navigate to **INFRASTRUCTURE > Devices**.
- 2 Select **COMPONENT GROUPS**, then click the **Add** icon at the bottom of the page.
The **Add Component Group** dialog is displayed.
- 3 Enter a suitable **Name**, and an optional **Description**, then click **SUBMIT**.
The new component group appears under **COMPONENT GROUPS**.
- 4 Under **COMPONENT GROUPS** at the top of the page, click the **Add** icon to open the **Add to Component Group** dialog.
- 5 In the dialog's search field, enter a component type, such as Blade Servers for example.
A list of components is displayed in the **Search Results** table.
- 6 Select one or more components, then click **ADD**.
- 7 Continue to search for and add components or click **CLOSE**.
The components are added to the group. You can now work with the new component group. For example:
- 8 To view Component Graphs:
 - a Highlight the new group and click **Details** at the top of the page.
 - b Click **Component Graphs**, then check **All on Same Graph**.

Note When using the **All on Same Graph** functionality, ensure that no more than 10 items are being displayed on the same graph for best usability.

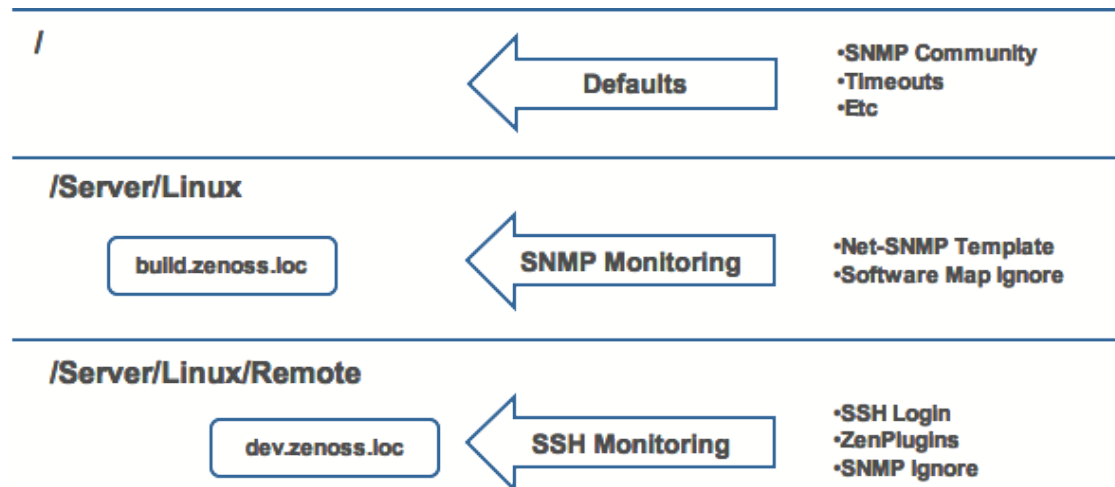
- 9 To view events for the component group, click **EVENTS**.
- 10 To disable monitoring on one or more components:
 - a Select the components you want to disable.
 - b Click **Action > Monitoring**.
 - c Click **NO** to disable monitoring.

Inheritance

Inheritance is defined by how many attributes are applied to a device at different levels in the device hierarchy.

The following diagram shows an example of how and where configuration properties can be set throughout the device class tree.

Figure 110: Device Class Tree and Inheritance



In this example, you can see that the default properties can be set at the highest level (/). However, as you travel further down the hierarchy, you see that you can override any of the configuration properties that are set at the root level.

The next two lines show how the device tree further defines properties for Linux servers. For example, to set up and use SNMP monitoring for all Linux servers (inclusive of) build.zenoss.loc, you could change these properties at the /Server/Linux level.

Further, if you wanted to change how you collect information for remote Linux servers, you could create a sub-group in /Server/Linux called /Server/Linux/Remote, setting these servers to use SSH monitoring and changing the associated properties for that sub-group.

All of these configuration properties and groupings co-exist, with any changes made lower in the hierarchy taking priority.

Distributed Monitoring

Resource Manager supports additional performance collection and event monitoring daemons, on the Resource Manager master host, and on other hosts.

This feature allows you to:

- Distribute processor, disk, and network load across multiple servers.
- Collect performance and event data from remote networks that cannot be reached by the Resource Manager master host.
- Configure more than one set of monitoring settings, such as different cycle times, for the `zenperfsnmp` daemon.

At installation, Resource Manager is configured with one hub and one collector (each named `localhost`) on the master host. The initial hub and collector may not be removed.

About Collectors

A *collector* is a set of collection daemons, on the Resource Manager host or another host, that shares a common configuration. That configuration contains values, such as:

- Number of seconds between SNMP collections cycles
- Default discovery networks
- Maximum number of `zenprocess` parallel jobs

Each collector has its own copy of each of the Resource Manager collection daemons. For example, Resource Manager initially contains collection daemons with names like `zenperfsnmp`, `zenprocess`, and `zenping`.

About Hubs

Resource Manager supports multiple hubs. A *hub* represents an instance of the `zenhub` daemon, through which all collector daemons communicate with the object and event databases. Hubs are now associated with a resource pool which helps with better collector distribution.

All collectors must belong to exactly one hub; however, a hub can have many collectors associated with it. All hubs (and indirectly all collectors) refer to the same object and event databases. Typically, only very large systems with more than five collectors (or more than 1,500 devices) benefit from multiple hubs.

The `ZooKeeper` service on the HBase cluster now maintains configuration data and provide distributed synchronization and group services.

About Resource Pools

A resource pool is a collection of compute, network, and storage resources that allows for better collector distribution. A default single resource pool is defined at startup. You can create additional resource pools by identifying them with a unique pool ID. You can then add hosts and assign them to the desired pool as well as adding collectors directly into a resource pool for better distribution.

There is no longer any need to update hubs or collectors with software updates. As soon as a service is restarted it will use the latest committed changes from the resource pool.

Navigating Collectors and Hubs

To view and manage collectors and hubs:

- 1 Log in as the Resource Manager user.
- 2 From the navigation menu, select **ADVANCED > Control Center**. The Control Center All Services page appears.

Figure 111: Control Center - All Services

The screenshot shows the 'All Services' page in the Control Center. At the top, there are navigation tabs: Settings, Move Devices, Control Center (selected), Monitoring Templates, and Jobs. Below the tabs is a toolbar with icons for adding, deleting, and refreshing, along with buttons for Start, Stop, and Restart. The main content is a table with the following columns: Name, Type, Uptime, AutoStart, Restart, State, and Host.

| Name | Type | Uptime | AutoStart | Restart | State | Host |
|-------------------|-----------|------------------------|-------------------------------------|---------|-------|-------------------------|
| CentralQuery | daemon | 6 days, 9:06:34.813614 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| HMaster | daemon | 6 days, 9:06:26.678772 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| Imp4MariaDB | daemon | | <input type="checkbox"/> | | Down | |
| Imp4OpenTSDB | daemon | | <input type="checkbox"/> | | Down | |
| localhost | hub | | | | ... | |
| localhost | collector | | | | ... | |
| collectorredis | daemon | 6 days, 9:06:15.053303 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| MetricShipper | daemon | 6 days, 9:06:33.888581 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zencommand | daemon | 6 days, 9:06:17.002472 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenmail | daemon | | <input type="checkbox"/> | | Down | |
| zenmodeler | daemon | 6 days, 9:06:15.942412 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenperfsnmp | daemon | 6 days, 9:06:15.928226 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenping | daemon | 6 days, 9:06:16.972205 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenpropertymon... | daemon | 6 days, 9:06:17.017761 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenpython | daemon | 6 days, 9:06:16.987216 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenucsevents | daemon | 6 days, 9:06:16.956418 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenvsphere | daemon | 5 days, 0:15:49.037176 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zminion | daemon | 6 days, 9:06:15.901294 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| zenhub | daemon | 6 days, 9:06:15.875667 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| mysqliadb-events | daemon | 6 days, 9:06:24.785442 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |
| mysqliadb-model | daemon | 6 days, 9:06:24.798263 | <input checked="" type="checkbox"/> | | Up | ucspm-stable.zenoss.loc |

The page lists existing hubs and collectors in hierarchical form. Hubs are listed at the top level; collectors are nested below the hub to which they belong.

From this page, you can:

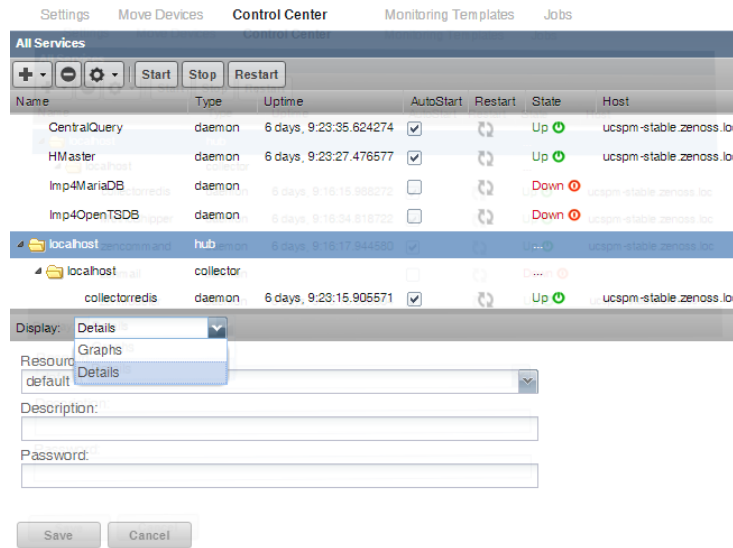
- Add a hub or a collector
- Delete a hub (which also deletes its associated collectors) or a collector

Note You cannot delete the default hub and collector (localhost)

- View and edit hub settings
- Configure associated monitoring and performance templates

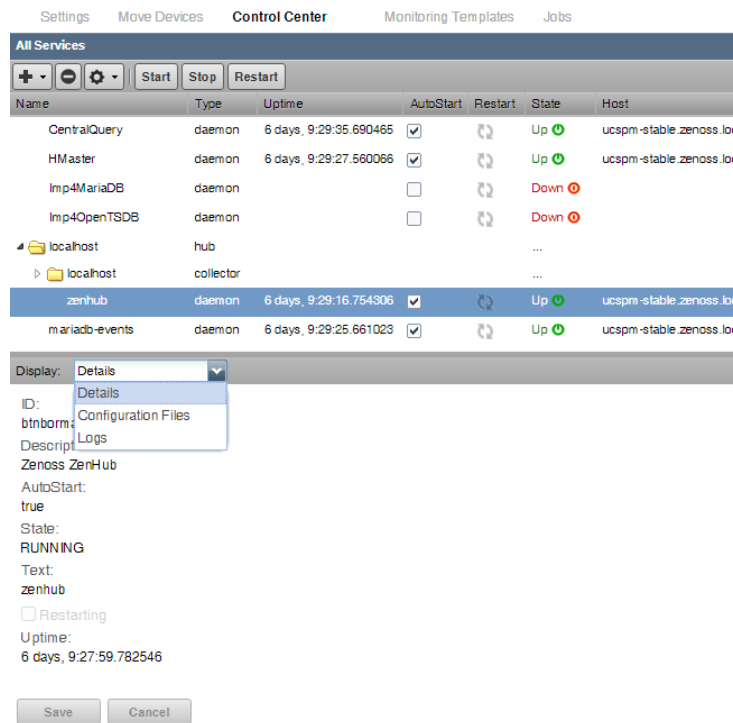
Select a hub to display details and graphs. The Resource Pool ID for the hub is displayed and can be changed if needed. You can add a description and password if needed.

Figure 112: Control Center - Hub Details



Select the zenhub daemon to view the details about the daemon, its logs, and to view and edit its configuration. Use the buttons on the top of the window to start, stop, or restart it.

Figure 113: Control Center - Daemon Details



Collector Data Storage

Resource Manager no longer uses RRD files on the collectors for data storage. We have created a centralized storage framework for this data which uses a Redis key-value store on the collector and then ships that data to an OpenTSDB (time series database) instance that runs on Hadoop and HBase.

Deleting Collectors

When you delete a collector, its devices are left without an assigned collector. Zenoss recommends that you reassign assigned devices prior to deleting a collector.

To delete a collector, click the name of the hub where the collector exists from the main collectors page. The Hub overview page appears. From the list of Collectors, select the collector you want to delete. From the Action menu, select Delete Collector.

When you delete collectors using this Resource Manager instance, they are not removed or "uninstalled" in any way from the collector device. They continue to exist on the device until manually removed through the file system.

Adding Devices to Collectors

Adding devices to collectors occurs when you add the device to Resource Manager.

- 1 Log in to the Resource Manager user interface as a user with ZenManager or Manager privileges.
- 2 Navigate to **Infrastructure > Devices**.
- 3 From the **Add Device** icon, select **Add a Single Device**.
- 4 In the **Add a Single Device** dialog, use the **Collector** field to select a local or remote collector for the device.
- 5 Populate the remaining fields as required, and then click **Add**.


The device appears in the **Devices** list, located at the bottom of the Collector Display Devices page.

Moving Devices Between Collectors

To move devices from one collector to another.

- 1 From the Navigation menu, select **ADVANCED > Move Devices**. The Move Devices screen appears.

Figure 114: Move Devices



- 2 Select the source collector(s) you want to move, then select a target collector to move the devices to from the drop-down list.
- 3 Click **Submit** to confirm. Resource Manager moves the devices to the newly selected collector.

Managing Collector Daemons

Collector daemons appear under the collector name on the Control Center All Services page and can be started, stopped and restarted from there. You can click on the name of a daemon to display details, configuration files, and logs.

Note Only one `zentrap` instance can be run on a server, as it must bind to the SNMP trap port (162).


User Commands

13

User commands allow you to execute arbitrary shell commands from the Resource Manager master server. A user command is executed on the server rather than the remote device unless the command explicitly uses SSH to connect to the remote device.

You can define and run user commands on a device or organizer (device class, system, group, location, or component group). You also can define commands globally. The User Commands menu bar shows the various functions that can be used in the User Commands screen. See the following sections for detailed instructions on adding and running user commands on specific devices or groups of devices.

Figure 115: User Commands Menu Bar



Defining Global User Commands

Global commands appear in the Commands list of options located at the top of the Devices page.

To define global user commands:


- 1 Select **Advanced > Settings** from the Navigation menu.
- 2 In the left panel, select **Commands**.
- 3 In the Define Commands area, select **Add User Command** from the Action menu. The Add User Command dialog appears.
- 4 Enter a name for the command, and then click **OK**. Only letters, numbers, and underscores are allowed in command names. Spaces are not allowed. The Define Commands page appears.
- 5 In the Description field, enter a description of what the command will do.
- 6 In the Command section, enter the TALEX expression-based command you want to run on the device.
- 7 Enter your system account password for confirmation, and then click **Save**.
The command is saved and added to the commands menu.

Note Global commands also can be edited from a specific device. Changes to a global command from a device are not limited to that device.

Running Global User Commands

To run a global user command, select one or more devices in the devices list, and then select a command from the Commands list of options.

Figure 116: Global User Commands



Defining User Commands for a Single Device

To define a user command for a device:

- 1 Select **INFRASTRUCTURE** from the Navigation menu. The Devices page appears.
- 2 Click a device name from the list to open the Device Overview page.
- 3 In the left panel, select **Device Administration**.
- 4 In the User Commands area, click the **Add a User Command** icon. The Add New User Command dialog appears.

Figure 117: Add New User Command

- 5 Enter the following information about the user command:
 - **Name** - Name of the user command.
 - **Description** - Description of what the command will do.
 - **Command** - TALES expression-based command you want to run.
 - **Confirm Password** - Enter your system account password for confirmation.

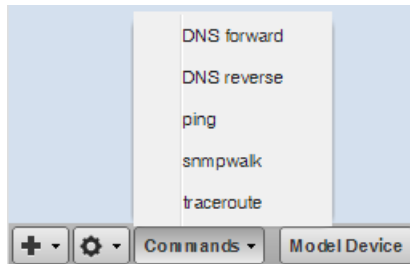
- 6 Click **Submit**. The command is saved and added to the user commands menu.
- 7 Optional: Test the command by selecting the command from the list and clicking the **Run** icon.

Running User Commands for a Single Device

To run a command defined for a single device:

- 1 Navigate to the **Infrastructure > Devices** page.
- 2 Click the device name in the device list to open the Device Overview page.
- 3 Select the command from the Commands list of options located at the bottom of the page.

Figure 118: Global User Commands Menu



Defining User Commands for All Devices in an Organizer

To define a user command for all devices in an organizer:

- 1 On the **INFRASTRUCTURE** page, select a device organizer in the devices hierarchy (e.g., /Server/Linux).
- 2 Click **Details**.
- 3 In the left panel, select **Device Administration**.
- 4 In the User Commands area, click the **Add a User Command** icon. The Add New User Command dialog appears.

Figure 119: Add New User Command

- 5 Enter the following information about the user command:
 - **Name** - Name of the user command.
 - **Description** - Description of what the command will do.
 - **Command** - TALES expression-based command you want to run.

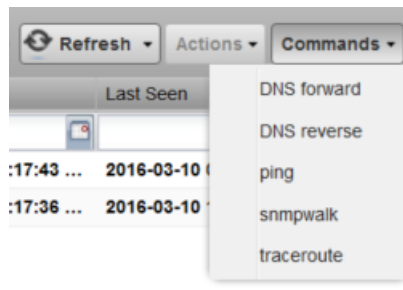
- **Confirm Password** - Enter your system account password for confirmation.
- 6 Click **Submit**. The command is saved and added to the user commands menu.
 - 7 Optional: Test the command by selecting the command from the list and clicking the **Run** icon.

Running User Commands for Devices in an Organizer

To run a command defined for devices in an organizer:

- 1 On the Infrastructure page, select a device organizer.
- 2 Select one or more devices in the filtered view.
- 3 Select the command from the Commands list located at the top of the page.

Figure 120: Global User Commands



User Command Example: Echo Command

This example shows how to create an echo user command. You can see the use of TALES expressions in the definition of this command.

- 1 Add a command called "echoDevice"
- 2 In the command definition, echo the name and IP address of the device:

```
echo name = ${here/id} ip = ${here/manageIp}
```

In a TALES expression, `here` is the object against which the expression is executed. Some TALES expressions in the system have other variables (such as `evt` for event, and `dev` or `device` for the device). See the Appendix titled TALES Expressions for more information about TALES expressions syntax.

- 3 Select a device and then run the command.
- 4 Edit the command to add more information:

```
echo name = ${here/id} ip = ${here/manageIp} hw = ${here/
getHWProductName}
```

- 5 Run the command against a group of devices and view the command outputs.

Managing Users

Every user within Resource Manager has a unique user ID, which allows an administrator to assign group permissions and alerting rules that are unique to each user. Unique IDs also help ensure secure access to the system.

To create and manage user accounts, you must be logged in to the system `admin` account, or as a user with extended privileges.

Creating User Accounts

To create a user account:

- 1 From the Navigation menu, select **ADVANCED**. The Settings page appears.
- 2 In the left panel, select **Users**. The users and groups administration page appears.
- 3 From the Action icon, select **Add New User**. The Add User dialog appears.
- 4 In the Username field, enter a unique name for the account.
- 5 In the Email field, enter the user account email address. Any alerts that you set up for this user will be send to this address.
- 6 Click **OK**. The user appears in the User List.

After creating the account, edit the account to provide a password and additional user details.

Editing User Accounts

To access and edit user account information:

- 1 In the Users list, click the name of the user you want to edit. The edit user page appears. The following example shows the admin user.

Figure 121: Edit User

2 Make changes to one or more settings:

- **Reset Password** - Facilitates user self-service by allowing a user to reset his or her own password. Click to reset and email the new password to the email address associated with the user's account.
- **User Preferences** - Resets all preferences such as grid columns and filters to their default values.
- **Roles** - Assign one or more roles (user privileges) to the user. To edit or assign roles, you must be a system Admin or be assigned the Manager role. For more information about user roles, and for a list of available roles and the privileges they provide, see [Roles](#) on page 179.
- **Groups** - Specify one or more groups to which this user belongs.
- **Email** - Enter the user's email address. To verify that the address is valid, click the test link.
- **Pager** - Enter the user's pager number.
- **Default Page Size** - Controls how many entries (by default) appear in tables. Enter a value for the default page size. The default value is 40.
- **Default Admin Role** - Select the default role that this user will have for administered objects associated with him or her.
- **Network Map Start Object** - Specify the default view for this user in the network map.
- **Time Zone** - Specify the time zone to be displayed on all charts and graphs within the product.
- **Set New Password / Confirm New Password** - Enter a new password for the user and confirm the entry.

3 Enter your password, and then click **Save** to confirm and save the changes for the user.

Associating Objects with Specific Users


You can associate any object in the system with a particular user, for monitoring or reporting purposes. Once associated with a user, you can then assign the user a specific role that applies to his privileges with respect to that object.

For more information about object-specific roles, see [Roles](#) on page 179.

To create an object association:


- 1 From **ADVANCED > Settings**, select **Users** in the left panel.
- 2 Select a user.
- 3 From the Edit page, select **Administered Objects** in the left panel. The list of administered objects appears.

Figure 122: Administered Objects - Add Object



- 4 Select an object type from the Administered Objects Action menu. You can add:
 - Device
 - Device class
 - System
 - Group
 - Location
- 5 Specify the component you want to add as an administered object, and then click **OK**. The object appears in the Administered Devices list for the user.

Figure 123: Administered Objects - Objects Added



- 6 Optional: Change the role that is associated for this user on this object.

Note The default role assigned to the user for an administered object is specified by the Default Admin Role field on the Edit page.

- 7 Click **Save** to save changes.

Adding Administrators

You also can associate an object with a user by adding an administrator to the object. Perform the following:

- 1 Navigate to the object you want to add to the user's list of administered objects.
- 2 Select **Device Administration**.

Figure 124: Administered Objects - Add Administrator

The screenshot displays the 'Administered Objects' page in the Zenoss Resource Manager. The left sidebar shows the navigation menu with 'Device Administration' selected. The main content area is divided into three sections:

- Maintenance Windows:** A table with columns: Enabled, Name, Start, Duration, Repeat, State. It contains three rows:

| Enabled | Name | Start | Duration | Repeat | State |
|---------|----------------------|---------------------|--------------|-----------------------|-------------|
| No | 1st of Month | 2014/12/01 02:00:00 | 01:00:00 hrs | Monthly: day of month | Maintenance |
| Yes | Every Thursday Night | 2014/12/11 22:00:00 | 30:00 mins | Weekly | Maintenance |
| Yes | One Time Testing | 2014/12/20 08:00:00 | 04:00:00 hrs | Never | Test |
- User Commands:** A table with columns: Name, Command. It contains five rows:

| Name | Command |
|-------------|--|
| DNS forward | host \${device/id} |
| DNS reverse | host \${device/manage/ip} |
| ping | \${device/pingCommand} -c2 \${device/mana... |
| snmpwalk | snmpwalk -\${device/zSnmpVer} -c\${device/z... |
| traceroute | \${device/tracerouteCommand} -q 1 -w 2 \${d... |
- Administrators:** A table with columns: Name, Role, Email, Pager. It contains one row:

| Name | Role | Email | Pager |
|------------|---------|---------|-------|
| cgilchrist | Manager | cgil... | |

- 3 Click the **Add Administrator** icon in the Administrators area. The Add Administrator dialog appears.
- 4 Select an administrator from the list, and then click **OK**. The administrator appears in the object's Administrators list. The object is added to the user's Administered Objects list.

User Groups

Resource Manager allows you to create user groups. By grouping users, you can aggregate rules and apply them across multiple user accounts.

Viewing User Groups

To view user groups, select **ADVANCED > Settings**, and then select **Users** from the left panel.


The groups area shows each user group and the users assigned to that group.

Creating User Groups

You can create user groups to aggregate rules and apply them across multiple user accounts.

To create a user group:

- 1 Navigate to **ADVANCED > Settings**.
- 2 In the left panel, select **Users**. The Users page appears.
- 3 From the Groups area Action menu, select **Add New Group**. The Add Group dialog appears.
- 4 In the Group field, enter a name for this user group, and then click **OK**. The group name appears in the Groups list.
- 5 Click the name of the group you created. The Users in Group page appears.
- 6 From the Action menu, select **Add User**. The Add User to Group dialog appears.

Figure 125: Add User to Group

- 7 From the User list of selections, select one or more users you want to add to the group, and then click **OK**. The user or users you select appear in the list of users for this group.

You also can choose administered objects and alerting rules for this user group. These alerting rules will apply to all users in the group. The user's original alerting rules and objects will also apply.

Roles

A role is a group of permissions that you can assign to users or groups.

The following table lists available roles.

| Role | Permissions |
|-------------|--|
| ZenUser | Provides global read-only access to system objects. |
| ZenManager | Provides global read-write access to system objects. |
| Manager | Provides global read-write access to system objects. Additionally provides read-write access to the Zope object database. |
| ZenOperator | Provides event management. Combine the ZenOperator role with the ZenUser role to allow users read-only access to the system, but also allow them to acknowledge and close events, move events to history, and add log messages to events. You can associate the ZenOperator role with an individual device, a device class, or a group of devices. |

Device Access Control Lists

About Device Access Control Lists (ACL)

Resource Manager supports fine-grained security controls. For example, this control can be used to give limited access to certain departments within a large organization or limit a customer to see only his own data. A user with limited access to objects also has a more limited view of features within the system. As an example, most global views, such as the network map, event console, and all types of class management, are not available. The device list is available, as are the device organizers: systems, groups, and locations. A limited set of reports can also be accessed.

Permissions and Roles

Actions in the system are assigned permissions. For instance to access the device edit screen you must have the “Change Device” permission. Permissions are not assigned directly to a user; instead, permissions are granted to roles, which are then assigned to a user. A common example is the ZenUser role. Its primary permission is “View,” which grants read-only access to all objects. ZenManagers have additional permissions such as “Change Device,” which grants them access to the device edit screen. When you assign a role to a user using the Roles field (on the Edit page), it is global.

Administered Objects

Device ACLs provide limited control to various objects within the system. Administered objects are the same as the device organizers: Groups, Systems, and Locations and Devices. If access is granted to any device organizer, it flows down to all devices within that organizer. To assign access to objects for a restricted user, you must have the Manager or ZenManager roles. The system grants access to objects is granted using the user's or user group's administered objects. To limit access, you must not assign a “global” role to the user or group.

Users and Groups

Users and user groups work exactly as they would normally. See the section in the User Management section of this guide dealing with users and groups.

Assigning Administered Object Access

For each user or group there is an Administered Objects selection, which lets you add items for each type of administered object. After adding an object you can assign it a role. Roles can be different for each object, so a user or group might have ZenUser on a particular device but ZenManager on a location organizer. If multiple roles are granted to a device through direct assignment and organizer assignment the resulting permissions will be additive. In the example above, if the device was within the organizer the user would inherit the ZenManager role on the device.

Portlet Access Control

Within Resource Manager, portlet access can be controlled. This is important for device ACLs.

Example: Restricted User with ZenUser Role

To create a restricted user with a ZenUser role:

- 1 As admin or any user account with Manager or ZenManager role, create a user named `actest`. Set a password for the user.
- 2 Make sure that no role is assigned to the user.
- 3 Edit the user's administered objects.
- 4 Add an existing device to the user. The device's role will default to ZenUser.
- 5 Log out of your browser, or open a second browser and then log in as `actest`.
- 6 Select **INFRASTRUCTURE**. You should see only the device you assigned to `actest`.
- 7 Navigate to the device and notice that the edit capabilities are not available. This is because you are in read-only mode for this device.

Example: Restricted User with ZenOperator Role

The ZenUser role from the previous section allows read-only access to devices. By adding the ZenOperator role to specific devices, device classes, or groups of devices, a user will be able to acknowledge and close events, move events to history, and add log messages to events.

To add the ZenOperator role to specific devices, device classes, or groups of devices:

- 1 Select the user name whose role must be changed on certain devices.
- 2 In the left-hand pane, click Administered Objects.
- 3 Click the Action icon and choose the device, device class, or other device organizer to which you want to grant the ZenOperator role.
- 4 Select the ZenOperator role from the drop-down menu for the newly selected device, device class, or device organizer.

The user now has the ZenUser role for all devices in this instance, with the exception of the device(s) selected above which function under the ZenOperator role.

Example: Restricted User with ZenManager Role

To restrict a user to a specific role:

- 1 Change the acltest user's role to "ZenManager" on the device. (You must do this as a user with ZenManager global rights.)
- 2 Go back to the acltest user's administered objects and set the role on the device to ZenManager.
- 3 As acltest, navigate back to the device. You now have access to edit the device.

Example: Adding Device Organizers

To add a device organizer:

- 1 Go to Groups and create a group called "RestrictGroup."
- 2 Go to the acltest user's administered objects and add the group to the user.
- 3 Logged in as acltest, notice that groups can be added to a user.
- 4 Place a device within this group and as acltest you should not only see the device within the group but also in the device list.

Restricted User Organizer Management

- 1 Give the acltest user ZenManager on your restricted group.
- 2 As acltest, you can now add sub-organizers under the restricted group.

Viewing Events

A user in restricted mode does not have access to the global event console. The available events for the user can be seen under his organizers.

Detailed Restricted Screen Functionality

Dashboard

By default, the dashboard is configured with three portlets:

- Object Watch List
- Device Issues
- Production State

These have content that will be restricted to objects for a given user.

Device List

The device list is automatically filtered to devices of a restricted user scoped to accessible devices. There are no menu items available.

Device Organizers

Device organizers control groups of devices for a restricted user. Every device added to the group will be accessible to the user. Permissions will be inherited down multiple tiers of a device organizer.

Reporting

Reports are limited to device reports and performance reports.

Reporting

15

Resource Manager provides many useful summaries of monitored resources at the **REPORTS** tab of the Resource Manager web interface.

Note If you experience any stairstepping in your graphs, you may want to change the reporting collection interval in Resource Manager. For example, setting the reporting collection interval to 60 minutes tells Resource Manager to update the API-driven reporting data at that interval which is different than the native collection interval.

To view a report, select the report's name in the left column.

Cisco UCS Reports

For each of the following Cisco UCS reports, you can click **Export all** to generate a `.csv` file showing the data in spreadsheet form.

Free Memory Slots

This report lists the number of free memory slots in each component grouped by domain. There are no filtering capability in this report.

Report Contents

| Column | Content |
|-------------------|---|
| Name | The domain name. |
| Component | The component being reported on. |
| Free Slots | The number of free memory slots on the corresponding component. At the bottom of this column, a list of total free memory slots is displayed. |

Click **Export all** to export the report as a comma-separated value (`.csv`) file.

Hardware Inventory

This reports lists the inventory of the UCS devices being monitored.

Report Contents

| Column | Content |
|--------------|--|
| UCS Device | The name of the UCS device. Clicking the link takes you to the device overview page. |
| Component | Lists components and related sub-components of the device. Clicking a link takes you to the appropriate device component page. |
| Manufacturer | The manufacturer of the component. |
| Model | The model number of the component. |
| Serial # | The serial number of the component. |
| Description | Detailed information about the component. |

Click **Export all** to export the report as a comma-separated value (.csv) file.

Device Reports

All Devices

A summary of each device Resource Manager is monitoring.

| Column | Content |
|---------|--|
| Name | The name of the device. |
| Class | The Resource Manager device class associated with the device. |
| Product | The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object. |
| State | The device's production state. Valid states include <code>Production</code> , <code>Pre-Production</code> , <code>Test</code> , and <code>Maintenance</code> . |
| Ping | The result of the most recent ping of the device. |
| SNMP | The result of the most recent attempt to gather data through the device's SNMP agent. |

All Monitored Components

A summary of each component Resource Manager is monitoring.

| Column | Content |
|-------------|--|
| Device | The name of the device which contains the component, with a link to its overview page. |
| Component | The name of the component, with a link to its overview page. |
| Type | The class associated with the component. |
| Description | A description of the component; typically, the component's name. |
| Status | The state of the component as of the most recent attempt to gather monitoring data. |

Device Changes

A summary of the devices in which changes were detected during the most recent collection of model data.

| Column | Content |
|-------------------|---|
| Name | The name of the changed device, with a link to its overview page. |
| Class | The Resource Manager device class associated with the device. |
| First Seen | The timestamp of the initial collection of modeling data for the device. |
| Collection | The timestamp of the collection in which a change was detected before the most recent collection. |
| Change | The timestamp of the most recent collection in which a change was detected. |

MAC Addresses (MAC Address Inventory)

A list of the unique device name, interface ID, and MAC address combinations in the Resource Manager database.

| Column | Content |
|---------------------|--|
| Device | The name of a device, with a link to its overview page. |
| Interface ID | The ID of a network interface, with a link to its overview page. |
| MAC address | A MAC address. |

Model Collection Age

A summary of devices that were not available for modeling data collection during the most recent 48 hour period.

| Column | Content |
|-------------------|--|
| Name | The name of the changed device, with a link to its overview page. |
| Class | The Resource Manager device class associated with the device. |
| First Seen | The timestamp of the initial collection of modeling data for the device. |
| Collection | TBD |
| Change | The timestamp of the most recent collection attempt. |

New Devices

The list of devices that were discovered and added to Resource Manager recently.

| Column | Content |
|-------------------|--|
| Name | The name of the changed device, with a link to its overview page. |
| Class | The Resource Manager device class associated with the device. |
| First Seen | The timestamp of the initial collection of modeling data for the device. |
| Collection | TBD |
| Change | TBD |

Ping Status Issues

A list of the devices which were down during the most recent collection of monitoring data.

| Column | Content |
|--------------|---|
| Name | The name of the device. |
| Class | The Resource Manager device class associated with the device. |

| Column | Content |
|----------------|--|
| Product | The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object. |
| State | The device's production state. Valid states include <code>Production</code> , <code>Pre-Production</code> , <code>Test</code> , and <code>Maintenance</code> . |
| Ping | The result of the most recent ping of the device. |
| SNMP | The result of the most recent attempt to gather data through the device's SNMP agent. |

SNMP Status Issues

A list of the devices for which no SNMP agent responded during the most recent collection of monitoring data.

| Column | Content |
|----------------|--|
| Name | The name of the device. |
| Class | The Resource Manager device class associated with the device. |
| Product | The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object. |
| State | The device's production state. Valid states include <code>Production</code> , <code>Pre-Production</code> , <code>Test</code> , and <code>Maintenance</code> . |
| Ping | The result of the most recent ping of the device. |
| SNMP | The result of the most recent attempt to gather data through the device's SNMP agent. |

Software Inventory

A list of the software installed in the devices and components which Resource Manager monitors.

| Column | Content |
|---------------------|---|
| Manufacturer | The name of the company that makes the software product. |
| Product | The name of the software product. |
| Count | The total number of devices or components on which the software is installed. |

Enterprise Reports

For each of the following Enterprise reports, you can click **Export all** to generate a `.csv` file showing the data in spreadsheet form.

Cisco Inventory

This report lists all the Cisco devices being monitored.

Report Filtering

| | |
|---------------|---|
| Device Class: | <input type="text" value="/Network/Cisco"/> |
| Group: | <input type="text" value="All"/> |
| | <input type="button" value="Update"/> |

Device Class

The device class to use for filtering. The default is `/Network/Cisco`.

Group

The specific group to consider when running the report. The default is `All`. The group could be an internal department, a location, a customer, etc.

To generate or refresh the report, click **Update**.

Report Contents

| Column | Content |
|-------------------|---|
| Name | The name of the Cisco device. Clicking the link takes you to the device overview page. |
| IP Address | Lists the IP address of the Cisco device. |
| Model | The model of the device. Clicking the link takes you to the Manufacturers Overview page for that product. |
| Serial # | The serial number of the Cisco device. |
| Type | The type of Cisco product, e.g., Device. |

Customized Performance Templates

| Customized Performance Templates | |
|--|---------------------------------|
| Target | Template |
| www.zenoss.com | Ping Local Copy |
| github.com | Device |
| apache1.zenoss.lbc | HttpMonitor |
| localhost.localdomain | HttpMonitor |
| oracle1.zenoss.lbc | Device |
| export all | |

Data Sources in Use

This reports lists the data sources defined in the system.

Report Contents

| Column | Content |
|-------------------------|--|
| Device Class | The device class of the data source in use. |
| Template | The template associated with the data source in use. |
| Data Source Name | The name of the data source. |
| Data Source Type | The type of connection for the data source. |
| ZenPack | The ZenPack related to the data source. |

Datapoints Per Collector

This report shows the number of devices and data points per collector, which is useful for gauging how much monitoring load is on each collector.

| Devices Per Collector | | | | |
|----------------------------|---------------------------|---------------------------|--------------|-----------------|
| Hub | Collector | Hostname | Device Count | Datapoint Count |
| localhost | localhost | localhost | 563 | 45764 |
| export all | | | | |

Report Contents

| Column | Content |
|------------------------|--|
| Hub | Name of the hub the collector is on |
| Collector | Name of the collector the devices are on |
| Hostname | Hostname location |
| Device Count | Total number of devices on the collector |
| Datapoint Count | Total number of data points being generated on the collector |

Defined Thresholds

The Defined Thresholds report provides details about all thresholds defined in the system. The report links to the target of each threshold. The target can be a device class, individual device, or individual component.

This report is useful for administering the system. You can use it to quickly identify which threshold events can occur within the system, and the severity of those events.

Report Contents

| Column | Content |
|------------------|--|
| Target | The device class of the defined threshold. |
| Template | The template associated with the defined threshold. |
| Threshold | The name of the defined threshold. Clicking the link takes you to the Performance Template where the threshold is defined. |
| Severity | The severity level assigned to the alert when the threshold is reached. |
| Enabled | The enabled status of the threshold. |

Event Time to Resolution

This report shows, for each user, the total time taken to acknowledge or clear events. Results are organized by user. It is helpful for tracking response time SLAs in a NOC-type environment.

Report Contents

| Column | Content |
|-------------|---|
| User | The user responsible for taking action on an event. |

| Column | Content |
|-----------------|--|
| Severity | The severity of the event. |
| Time to Ack | The amount of time to acknowledge the event. |
| Time to Archive | The amount of time to clear and archive the event. |

Interface Utilization

The Interface Utilization report shows the average, maximum, and minimum input and output traffic rate. The report is generated in a tabular form with the calculations shown by each interface.

Report Filtering

Root Organizer: Device Filter:
 Start Date: End Date:
 Trendline Type:
 Show All Interfaces

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click **select**. The default range is the week ending with the current date.

Trendline Type

The calculation used to determine the trendline. The only available selection is **Linear**.

To generate or refresh the report, click **Update**. If you want to show all interfaces on the report, check the **Show All Interfaces** box before clicking **Update**.

Report Contents

| Column | Content |
|-------------|---|
| Device | The name of the device that contains the interface, with a link to its overview page. |
| Interface | The name of the interface, with a link to its overview page. |
| Description | A description of the interface. |
| Speed | The maximum transfer rate the interface supports, per second. |
| In Avg | The average input traffic through the interface, per second. |
| Out Avg | The average output traffic through the interface, per second. |
| In Max | The maximum input traffic through the interface, per second. |

| Column | Content |
|----------------------------------|---|
| Out Max | The maximum output traffic through the interface, per second. |
| In Min | The minimum input traffic through the interface, per second. |
| Out Min | The minimum output traffic through the interface, per second. |
| Forecasted In Exhaustion | Projected date of exhaustion for input traffic. |
| Forecasted Out Exhaustion | Projected date of exhaustion for output traffic. |

Note In order for the In Max and Out Max values to be calculated, you need to have at least 8 hours worth of data, otherwise a N/A value will be returned.

The following shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the pop-up list. In this example, there are 5 pages of results from device 10.171.100.107 followed by one page of results from device 10.171.100.109, etc.

| | | | | | | |
|----------------|------------------|--------|------|------|------|------|
| 10.171.100.107 | Ethernet108_1_46 | 23.3Kb | 1.4b | 1.4b | 3.6b | 3.4b |
| 10.171.100.107 | Ethernet108_1_47 | 22.5Kb | 1.4b | 1.4b | 3.4b | 3.2b |
| 10.171.100.107 | 10.171.100.107 | 22.8Kb | 1.4b | 1.3b | 3.0b | 3.4b |
| 10.171.100.107 | 10.171.100.107 | 23.8Kb | 1.4b | 1.3b | 3.0b | 2.9b |
| 10.171.100.107 | 10.171.100.107 | 22.8Kb | 1.4b | 1.4b | 3.4b | 3.3b |
| 10.171.100.107 | 10.171.100.107 | 22.3Kb | 1.3b | 1.4b | 3.2b | 3.3b |
| 10.171.100.107 | 10.171.100.107 | 23.7Kb | 1.4b | 1.3b | 3.0b | 3.0b |
| 10.171.100.107 | 10.171.100.109 | 22.7Kb | 1.4b | 1.3b | 3.4b | 3.0b |
| 10.171.100.107 | 10.171.100.14 | 23.8Kb | 1.4b | 1.4b | 3.2b | 3.2b |
| 10.171.100.107 | 10.171.100.14 | 22.9Kb | 1.3b | 1.4b | 3.0b | 3.4b |
| 10.171.100.107 | 10.171.100.58 | 23.3Kb | 1.4b | 1.4b | 3.3b | 3.3b |
| 10.171.100.107 | perf2-switch | 23.6Kb | 1.4b | 1.3b | 3.1b | 2.8b |

1 of 1017 << 10.171.100.107 >> show all export all Page Size 100 ok

Click **Export all** to export the report as a comma-separated value (.csv) file.

Interface Volume

The Interface Volume report shows the total input and output traffic for the report period, along with a calculation of the input and output traffic per day. The report is generated in a tabular form with the calculations shown by each interface.

Report Parameters

Device Class: Device Filter:

Start Date: End Date:

Trendline Type:

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date**End Date**

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click **select**. The default range is the week ending with the current date.

Trendline Type

The calculation used to determine the trendline. The only available selection is **Linear**.

To generate or refresh the report, click **Update**.

Report Contents

| Column | Content |
|----------------------------------|---|
| Device | The name of the device that contains the interface, with a link to its overview page. |
| Interface | The name of the interface, with a link to its overview page. |
| Description | A description of the interface. |
| Speed | The maximum transfer rate the interface supports, per second. |
| In Vol | The input traffic through the interface for the time period of the report. |
| In Vol/day | The input traffic through the interface, per day. |
| Out Vol | The output traffic through the interface for the time period of the report. |
| Out Vol/day | The output traffic through the interface, per day. |
| Total Vol | Total volume of traffic through the interface for the time period of the report. |
| Forecasted In Exhaustion | Projected date of exhaustion for input traffic. |
| Forecasted Out Exhaustion | Projected date of exhaustion for output traffic. |

The following shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the pop-up list. In this example, there are 5 pages of results from device 10.171.100.107 followed by one page of results from device 10.171.100.109, etc.

| | | | | | | |
|----------------|------------------|--------|------|------|------|------|
| 10.171.100.107 | Ethernet108_1_46 | 23.3Kb | 1.4b | 1.4b | 3.6b | 3.4b |
| 10.171.100.107 | Ethernet108_1_47 | 22.5Kb | 1.4b | 1.4b | 3.4b | 3.2b |
| 10.171.100.107 | Ethernet108_1_48 | 22.8Kb | 1.4b | 1.3b | 3.0b | 3.4b |
| 10.171.100.107 | 10.171.100.107 | 23.8Kb | 1.4b | 1.3b | 3.0b | 2.9b |
| 10.171.100.107 | 10.171.100.107 | 22.8Kb | 1.4b | 1.4b | 3.4b | 3.3b |
| 10.171.100.107 | 10.171.100.107 | 22.3Kb | 1.3b | 1.4b | 3.2b | 3.3b |
| 10.171.100.107 | 10.171.100.107 | 23.7Kb | 1.4b | 1.3b | 3.0b | 3.0b |
| 10.171.100.107 | 10.171.100.109 | 22.7Kb | 1.4b | 1.3b | 3.4b | 3.0b |
| 10.171.100.107 | 10.171.100.14 | 23.8Kb | 1.4b | 1.4b | 3.3b | 3.2b |
| 10.171.100.107 | 10.171.100.14 | 22.9Kb | 1.3b | 1.4b | 3.0b | 3.4b |
| 10.171.100.107 | 10.171.100.14 | 23.3Kb | 1.4b | 1.4b | 3.3b | 3.3b |
| 10.171.100.107 | 10.171.100.68 | 23.6Kb | 1.4b | 1.3b | 3.1b | 2.8b |
| 10.171.100.107 | perf2-switch | 23.6Kb | 1.4b | 1.3b | 3.1b | 2.8b |

1 of 1017 | 10.171.100.107 | show all | export all | Page Size | 100 | ok

Click **Export all** to export the report as a comma-separated value (.csv) file.

Maintenance Windows

This report shows all defined windows that are active during a selected time period.

To change the reporting time period, enter Start and End dates (or click **Select** to select dates from a calendar). Click **Update** to refresh the report.

Report Contents

| Column | Content |
|-------------------|--|
| Name | Name of the maintenance window. |
| Target | |
| Start Date | Date the maintenance window becomes active. |
| Duration | Length of time for the maintenance window to be in effect. |
| Repeat | Repeat interval of the maintenance window. |

Network Topology

Shows the layout of the network, according to the routes that Resource Manager understands, starting from the collector and ending at the remote devices associated with the collector.

The report does not return data if the host on which the Resource Manager collector is running does not have a device created in the device management database (DMD) object which stores the basic model of the network in the Zope database (ZODB). Create a device representing the collector in the DMD, and then run report again.

An invalid route entry (for example, 'Missing link here' value in the Route column) indicates that Resource Manager cannot determine how to route from one device to another. Correct this by adding a network interface to the model (no new hardware required) and then adding a new route entry from the last device in the route to the device (the IP address shown at the far right of the table).

Report Filtering

Report Settings

Collectors: localhost ▾

Show valid routes?:

Show invalid routes?:

[Update](#)

Collectors

Select the collector to base the report on. Connections from this collector to associated devices will be shown in the report.

Show valid routes?

Select this check box to show valid connections between the selected collector and remote devices.

Show invalid routes?

Select this check box to show invalid connections between the selected collector and remote devices. An invalid route indicates that a route from the collector to the device could not be determined.

To generate or refresh the report, click **Update**.

Report Contents

| Column | Content |
|------------------|---|
| Collector | Name of the collector. |
| Route | Route from the collector to the device. |

| Column | Content |
|--------------------------|--|
| Device IP Address | IP address of the device. |
| Device Name | Name of the device. |
| Repeat | Repeat interval of the maintenance window. |

Notifications and Triggers by Recipient

This report lists all of the notifications that the system has sent.

| Column | Content |
|---------------------|--|
| Name | The recipient name. |
| Address | The address associated with the recipient. |
| Notification | The content of the notification. |
| Triggered by | The trigger that initiated the notification. |

Organizer Availability

This report provides the availability percentage of all network organizers in the system. It can be filtered by organizer, event class, component, event severity, and date.

You can report on the availability of device classes, locations, systems, or groups within a defined time frame. This report offers two reporting modes:

- **Averaged** - Defines the organizer as available for the average availability time for all devices contained in it.
- **Coalesced** - Defines the organizer as available only if all devices are available during a certain time period.

Report Filtering

Organizer Availability Filtering

| | | |
|--|--|---|
| Root Organizer: <input type="text" value="/Devices"/> | Component: <input type="text"/> | Event Class: <ul style="list-style-type: none"> /Status/Nagios /Status/Ntp /Status/OSProcess /Status/Perf /Status/Ping |
| Start: <input type="text" value="05/16/2014"/> <input type="button" value="select"/> | Severity: <input type="text" value="Critical"/> | |
| End: <input type="text" value="05/23/2014"/> <input type="button" value="select"/> | Summation: <input type="radio"/> Averaged <input checked="" type="radio"/> Coalesced | |
| <input type="button" value="Update"/> | | |

Root Organizer

The device class to use for filtering. The default is /Devices.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click **select**. The default range is the week ending with the current date.

Component

The specific group to consider when running the report. The default is All. The group could be an internal department, a location, a customer, etc.

Severity

The severity level to filter by. The default is `Critical`. If another level is wanted, select it from the drop-down list.

Summation

Select between `Averaged` and `Coalesced` depending on how you want to define the organizer as available. The default is `Coalesced`.

Event Class

The event class to use for filtering. The default is `/Status/Ping`.

To generate or refresh the report, click **Update**.

Note If you export the report, be sure to format the percentage columns to show percentages instead of decimal values.

Organizer Graphs

This reports shows graphical data about a given organizer. The information displayed varies depending on the type of organizer. For example, selecting `/Devices/Network/Cisco` as an organizer displays graphs on CPU and Memory Utilization, Throughput, Errors, etc.

Report Filtering

Organizer Graphs Filtering

Organizer: Filter: Start:

Organizer

The class to use for filtering. The default is `/Groups`.

Filter

Additional text filter to refine results. The default is `. *` which returns all graphs available for the selected organizer.

Start Date**End Date**

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click **select**. The default range is the week ending with the current date.

To generate or refresh the report, click **Update**.

User Event Activity

A list of users or groups and the number of events each has acknowledged and archived.

Report Filtering

Start: End: Group By:

The following fields define the report.

Start Date**End Date**

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click **select**. The default range is the week ending with the current date.

Group By

The type of report to create, **User** or **Group**.

To generate the report, click **Update**.

Report Contents

| Column | Content |
|-----------------------|--|
| User or Group | The column label and content depend on the report selection. User The name of a user, with a link to the user's USER SETTINGS page. Group The name of a group, with a link to the group's Users in Group page. |
| # Acknowledged | The total number of events acknowledged by the user or group during the reporting period. |
| # Archived | The total number of events archived (cleared) by the user or group during the reporting period. |

Users Group Membership (User to Group Mapping)

A list of Resource Manager users and the groups to which they belong.

| Column | Content |
|---------------|--|
| User | The name of a user, with a link to the user's USER SETTINGS page. |
| Groups | The list of groups to which the user belongs. Each name includes a link to the group's Users in Group page. |

Event Reports**All EventClasses (All Event Classes)**

A list of each item in the event hierarchy in Resource Manager. Each item (class) includes the total number of subclasses, instances, and events associated with the class.

| Column | Content |
|------------------|---|
| Name | The event class name. |
| Sublasses | The total number of subclasses associated with the event class. |
| Instances | The total number of instances of the class and its subclasses. |
| Events | The total number of events associated with the class. |

All EventMappings (All Event Mappings)

A list of each item in the event mapping hierarchy in Resource Manager. Each item (event mapping) includes its key and example text, and a count of the events associated with the event mapping.

| Column | Content |
|----------------------|---|
| Name | The name of the event mapping, which includes its location in the event class hierarchy, and its key. |
| EventClassKey | The unique identifier of the event mapping. |
| Evaluation | A portion of the example associated with the event mapping. |
| Events | The total number of events associated with the event mapping. |

All Heartbeats

A list of all Resource Manager daemons, showing the number of seconds elapsed since each daemon sent a heartbeat event.

| Column | Content |
|------------------|--|
| Device | The device on which the daemon is running. |
| Component | The name of the daemon. |
| Seconds | The number of seconds elapsed since the daemon sent a heartbeat event. |

Performance Reports

Availability Report

Shows the percentage of time that a device or component is considered available. You can filter this report on device, component, event class, or severity. You can also limit the time frame for the availability.

The value the percent available is calculated by first summing the duration of all events of a particular class with a production state of `Production` and with a severity greater than or equal to a specified severity. This sum is then divided by the total duration of the time range, and then subtracted from 1 and multiplied by 100 to get the percent available, as in the following equation:

$$1 - ((\text{Total event down time}) / (\text{total duration})) * 100$$

Note Events whose `firsttime` and `lasttime` fields are the same are not used in the calculation. These could represent an event that occurs and is subsequently cleared by the next event, or an event that has happened only once in the specific date range.

CPU Utilization

Shows monitored interfaces, devices, load averages, and % utility. You can customize start and end dates, and summary type (average or maximum).

This report uses data point aliases. (For more information about data point aliases, see [Data Point Aliases](#) on page 114.) To add data points to a report, add the alias, and then ensure the values return in the expected units.

| Alias | Expected Units |
|-----------------|----------------|
| loadAverage5min | Processes |
| cpu_pct | Percent |

Filesystem Util Report

Shows total bytes, used bytes, free bytes, and percentage utilization for each device. You can customize start and end dates, and summary type (average or maximum).

This report uses data point aliases. (For more information about data point aliases, see [Data Point Aliases](#) on page 114.) To add datapoints to a report, add the alias, and then ensure the values return in the expected units.

| Alias | Expected Units |
|----------------------------|----------------|
| usedFilesystemSpace__bytes | bytes |

Interface Utilization

Shows the traffic through all network interfaces monitored by Resource Manager.

Columns included in the report:

- **Device** - Interface's device
- **Interface** - Interface
- **Speed** - Interface's rated bandwidth, in bits per second
- **Input** - Average traffic going out of the interface, in bits per second
- **Output** - Average traffic coming in to the interface, in bits per second
- **Total** - Total average traffic across the interface, in bits per second
- **% Util** - Average fraction of the interface's bandwidth consumed

This report uses data point aliases. (For more information about data point aliases, see [Data Point Aliases](#) on page 114.) To add data points to a report, add the alias, and then ensure the values return in the expected units.

| Alias | Expected Units |
|---------------------|----------------|
| inputOctets__bytes | bytes/sec |
| outputOctets__bytes | bytes/sec |

Memory Utilization

Provides system-wide information about the memory usage for devices in Resource Manager.

Columns included in the report:

- **Device** - Name of the device
- **Total** - Total memory
- **Available** - Available memory
- **Cache Memory** - Cache memory
- **Buffered Memory** - Buffered memory
- **% Util** - Percentage of memory utilized

The report uses data point aliases. (For more information about data point aliases, see [Data Point Aliases](#) on page 114.) To add data points to the report, add the alias, and then ensure the values return in the expected units.

| Alias | Expected Units |
|------------------------|----------------|
| memoryAvailable__bytes | bytes |
| memoryBuffered__bytes | bytes |
| memoryCached__bytes | bytes |

Threshold Summary

Provides information about the devices that are approaching or exceeding their thresholds.

Columns included in the report:

- **Device** - Name of the device
- **Component** - Component affected
- **Event Class** - Event class. You can filter which event classes are displayed from the Event Class drop-down list.
- **Count** - Number of occurrences
- **Duration** - Length of time
- **%** - Percentage

ZenPacks

16

ZenPacks extend and modify the system to add new functionality. This can be as simple as adding new device classes or monitoring templates, or as complex as extending the data model and providing new collection daemons.

You can use ZenPacks to add:

- Monitoring templates
- Data sources
- Graphs
- Event classes
- User commands
- Reports
- Model extensions
- Product definitions

Simple ZenPacks can be created completely within the user interface. More complex ZenPacks require development of scripts or daemons, using Python or another programming language.

ZenPacks can be distributed for installation on other Resource Manager systems.

Displaying the list of installed ZenPacks

Displaying installed ZenPacks in the browser interface

Follow these steps to display the ZenPacks that are installed on Resource Manager, in the Resource Manager browser interface.

- 1 In the Resource Manager browser interface, select the **ADVANCED** tab.

Note The action menu in the left column does not include an option to create a ZenPack. For more information, see [Creating a new ZenPack](#) on page 204.

- 2 In the left column, select **ZenPacks**.

The following figure shows an example list of ZenPacks.

Displaying installed ZenPacks in the command-line interface

To perform this procedure, you need a user account with `serviced` CLI privileges on the Control Center master host.

Follow these steps to display the ZenPacks that are installed on Resource Manager, in the command-line interface.

- 1 Log in to the Control Center master host as a user with `serviced` CLI privileges.
- 2 Display the list of installed ZenPacks.

```
serviced service run zope zenpack list
```

ZenPack information resources

Zenoss provides numerous ZenPacks that add and extend system functionality. The Resource Manager web interface includes a link (the question mark icon) to the documentation of the ZenPacks included in your installation of Resource Manager, and *ZenPack Catalog* provides detailed descriptions of all of the ZenPacks developed by Zenoss.

You may also create your own ZenPacks, or download and install ZenPacks developed by others. The following list identifies ZenPack resources.

- [ZenPack Discussion Forum](#)
- [ZenPack Development Forum](#)
- [Public Zenoss repositories on GitHub](#)

Preparing to install or upgrade a ZenPack

Perform this procedure to minimize the amount of time Resource Manager is unavailable during a ZenPack install or upgrade.

- 1 Log in to your workstation, and start a web browser.
- 2 Download the ZenPack to install or upgrade from the [Zenoss Support](#) site.
Contact your Zenoss representative for login credentials.
- 3 Copy the ZenPack egg file to a local directory on the Control Center master host.

- a Create a directory for the ZenPack egg file.

The directory must be local (not mounted) and must be readable, writable, and executable by all users.

The following commands create a directory in `/tmp`:

```
mkdir /tmp/zenpack && chmod 777 /tmp/zenpack
```

- b Use the file transfer command or utility you prefer to copy the file.
- c Change the file permissions.

The ZenPack egg file must have the same permissions as its parent directory.

```
chmod 777 /tmp/zenpack/*.egg
```

- 4 Optional: Install ZenPack dependencies, if necessary.

A ZenPack may require packages or other software not included in the ZenPack egg file. Perform the following substeps to ensure the dependencies are available.

- a Log in to the Control Center master host as a user with Control Center CLI privileges.
- b Start an interactive shell in a Zope service container.

In the following command, the `-s` flag saves and tags the changes you make. Replace *MyTag* with a short name that describes the dependencies you are installing.

```
serviced service shell -i -s MyTag zope bash
```

The `serviced` daemon starts a Bash shell, and logs you in as the `root` user.

- c** Install required dependencies.

For example, to install the `terminus` font for X Windows in Ubuntu Linux, enter the following command:

```
apt-get install xfonts-terminus
```

You may enter any number of commands to install the required dependencies.

- d** Return to the Control Center master host shell session.

```
exit
```

- e** Create a snapshot and commit your changes.

```
serviced snapshot commit MyTag
```

- f** Restart the Resource Manager application.

```
serviced service restart Zenoss.resmgr
```

The restart command briefly pauses and then restarts Resource Manager and all of its child services. Performing the restart ensures all of the services have access to the ZenPack's dependencies.

Installing or upgrading a ZenPack

The ZenPack installation and upgrade procedure requires stopping and starting Resource Manager.

ZenPacks can introduce fundamental changes to the ways in which Resource Manager operates, so the safest way to install a ZenPack is to suspend operations temporarily.

Note Do not use this procedure to install or upgrade the Zenoss Service Impact ZenPacks, ZenPacks.zenoss.ImpactServer or ZenPacks.zenoss.Impact. For more information, refer to *Zenoss Service Impact Installation Guide for Resource Manager 5.x* or *Zenoss Service Impact Installation Guide for Resource Manager 4.2*.

Before performing this procedure, complete all of the steps in [Preparing to install or upgrade a ZenPack](#) on page 200.

- 1 Log in to the Control Center browser interface.
- 2 In the **Applications** table, identify the name of the Resource Manager instance to modify.
- 3 Scroll down to the **Services** table, and then collapse the subservices of the **Infrastructure** and **Zenoss** services.
- 4 Stop the **Zenoss** service, and then verify its subservices are stopped.
 - a In the **Actions** column of the **Services** table, click **Stop**.
 - b In the **Service** column, click **Zenoss**, and then scroll down to verify all subservices are stopped.
- 5 Create a snapshot.
 - a Log in to the Control Center master host as a user with `serviced` CLI privileges.

- b Create a snapshot.

```
serviced service snapshot Zenoss.resmgr
```

On completion, the `serviced` command returns the ID of the new snapshot.

- 6 Restart required services.
 - a Click the **Zenoss** service to expand its subservices.
 - b Restart the following subservices.
 - Impact (if Zenoss Service Impact is installed)
 - `zeneventserver`
 - `Zope`

Note In the **Services** table, the Failing icon (a red circle with an exclamation point) in the **Status** column represents the cumulative result of one or more customized health checks. To view the status of individual health checks, move the pointer over the icon, which displays a pop-up.

- 7 Change directory to the directory in which the ZenPack egg file is located.
For example:

```
cd /tmp/zenpack
```

- 8 Install the ZenPack.

```
serviced service run zope zenpack-manager install ZenPack-File.egg
```

The daemons a ZenPack provides (if any) are packaged in Docker containers, and installed as child services of the current instance of Resource Manager.

- 9 In the Control Center browser interface, stop the `zeneventserver` and `zope` services, then start the **Zenoss** service that you stopped previously.

Note You must perform this step exactly in order so that the underlying `zeneventserver` and `zope` services get restarted.

Removing a ZenPack

The ZenPack removal procedure requires stopping and starting Resource Manager.

Note Removing a ZenPack can have unexpected consequences.

- Removing a ZenPack removes all objects provided by the ZenPack, as well as all objects that depend on code provided by the ZenPack.
- Removing a newer version of a ZenPack to install an older version fails if the newer version includes migration code.
- Removing a ZenPack that installs a device class removes the device class, any contained device classes, and all devices in that class.

Often, the safest choice is not to remove a ZenPack.

Note Do not use this procedure to remove the Zenoss Service Impact ZenPacks, `ZenPacks.zenoss.ImpactServer` and `ZenPacks.zenoss.Impact`. For more information, refer to *Zenoss Service Impact Installation Guide for Resource Manager 5.x* or *Zenoss Service Impact Installation Guide for Resource Manager 4.2*.

To avoid negative consequences, Zenoss recommends the following practices.

- Review the documentation of the ZenPack to remove for information about classes and daemons (services) associated with it.
 - Delete data sources provided by the ZenPack to remove.
- 1 Log in to the Control Center browser interface.
 - 2 In the **Applications** table, identify the name of the Resource Manager instance to modify.
 - 3 Scroll down to the **Services** table, and then collapse the subservices of the **Infrastructure** and **Zenoss** services.
 - 4 Stop the **Zenoss** service, and then verify its subservices are stopped.
 - a In the **Actions** column of the **Services** table, click **Stop**.
 - b In the **Service** column, click **Zenoss**, and then scroll down to verify all subservices are stopped.
 - 5 Create a snapshot.
 - a Log in to the Control Center master host as a user with `serviced` CLI privileges.
 - b Create a snapshot.

```
serviced service snapshot Zenoss.resmgr
```

On completion, the `serviced` command returns the ID of the new snapshot.

- 6 Restart required services.
 - a Click the **Zenoss** service to expand its subservices.
 - b Restart the following subservices.
 - Impact (if Zenoss Service Impact is installed)
 - `zeneventserver`
 - `Zope`

Note In the **Services** table, the Failing icon (a red circle with an exclamation point) in the **Status** column represents the cumulative result of one or more customized health checks. To view the status of individual health checks, move the pointer over the icon, which displays a pop-up.

- 7 Obtain the exact name of the ZenPack to remove.

```
serviced service run zope zenpack list
```

The first item of each line of output is the full name of an installed ZenPack.

- 8 Remove the ZenPack.
Replace *ZenPack-Name* with the full name of the ZenPack to remove:

```
serviced service run zope zenpack-manager uninstall ZenPack-Name
```

- 9 Optional: Remove services associated with the removed ZenPack, if necessary.

The daemons a ZenPack provides (if any) are packaged in Docker containers, and installed as child services of the current instance of Resource Manager.

Note Some ZenPacks provide services upon which other ZenPacks rely. Make sure the service you remove is not needed by another ZenPack.

The following example removes the `zenwebtx` service, which is provided by the `ZenPacks.zenoss.ZenWebTx` ZenPack.

```
serviced service rm zenwebtx
```

- 10 In the Control Center browser interface, stop the `zeneventserver` and `zope` services, then start the **Zenoss** service that you stopped previously.

Note You must perform this step exactly in order so that the underlying `zeneventserver` and `zope` services get restarted.

Creating a new ZenPack

To perform this procedure, you need a user account with `serviced` CLI privileges on the Control Center master host.

This procedure demonstrates how to create a ZenPack for customized monitoring templates, customized event classes and mappings, device MIBs, and similar items which require no customized Python code. For more advanced ZenPack development, refer to the [zenpacklib site](#).

- 1 Log in to the Control Center master host as a user with `serviced` CLI privileges.
- 2 Create a ZenPack.
Replace `ZenPacks.myOrg.myPackName` with the name of the ZenPack to create:

```
serviced service run zope zenpack create \  
ZenPacks.myOrg.myPackName
```

- 3 Restart the Zope service.

```
serviced service restart zope
```

General Administration and Settings

17

Recovering from a POSKey error during login

To perform this procedure, you need an account on the Control Center master host that has permission to use the Control Center command-line interface.

Resource Manager automatically recovers from most causes of POSKey errors that prevent logging into the Resource Manager interface. If a POSKey error does occur, perform this procedure.

- 1 Log in to an account on the Control Center master host that has permission to use the Control Center command-line interface.
- 2 Stop the Zope and authentication services.

```
serviced service stop zope && serviced service stop zauth
```

- 3 Monitor the services until they are stopped.

```
watch serviced service status zope
watch serviced service status zauth
```

- 4 Rebuild the session database.

```
serviced service run mariadb rebuild_zodb_session
```

- 5 Restart the authentication and Zope services.

```
serviced service start zauth && serviced service start zope
```

Events Settings

You can adjust events settings for:

- Events database connection
- Event maintenance

Changing Events Database Connection Information

To edit events database connection settings, make changes in the `zeneventserver.conf` file. You can edit the file directly, or run a configuration script.

Configurable database connection settings are:

- **JDBC Hostname (zep.jdbc.hostname)** - Specify the IP address of the host.
- **JDBC Port (zep.jdbc.port)** - Specify the port to use when accessing the events database.
- **JDBC Database Name (zep.jdbc.dbname)** - Specify the database name.
- **JDBC Username (zep.jdbc.username)** - Specify the user name for the database.
- **JDBC Password (zep.jdbc.password)** - Specify the password for the database.

To edit these values, run the `zeneventserver` configuration script, as follows:

```
zeneventserver-config -u zep.jdbc.Name=Value
```

Where *Name* is the partial setting name and *Value* is the value you want to specify for the setting.

Changing Events Maintenance Settings

To edit maintenance settings, make changes to one or more fields on the Event Configuration page (**ADVANCED > Settings > Events**):

- **Don't Age This Severity and Above** - Options are Age All Events, Critical, Error, Warning, Info, Debug, and Clear. By default, this value is set to Error, meaning that all events with a status of Error or Critical are not aged.
- **Event Aging Threshold (minutes)** - Set the time value, in minutes, that an event must reach before it is aged. By default, this is 240 minutes.
- **Event Aging Interval (milliseconds)** - The interval when events are scanned to perform autoaging. By default, this is 60000 milliseconds (60 sec).
- **Event Aging Limit** - The maximum number of events to age in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- **Event Archive Threshold (minutes)** - Specify the number of minutes since a closed event was last seen before it is moved to the event archive. The minimum value is 1; the maximum value is 43200.
- **Event Archive Interval (milliseconds)** - The interval when events are scanned for moving to the archive. By default, this is 60000 milliseconds (60 sec).
- **Event Archive Limit** - The maximum number of events to archive in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- **Delete Archived Events Older Than (days)** - The number of days that events in the event archive are saved. By default, they are kept in the archive for 90 days. The minimum value is 1 and the maximum value is determined by the range of event archive partitions. With the default configuration, the maximum value is 1000 days.
- **Default Syslog Priority** - Specify the default severity level assigned to an event coming from zensyslog if no priority can be determined from the event.
- **Default Availability Report (days)** - Enter the number of days to include in the automatically generated Availability Report. This report shows a graphical summary of availability and status.
- **Max Event Size in Bytes** - The maximum size of an event that will be processed in bytes. Events that are too large will be logged and dropped. Events that will become too big will have their details overwritten with new details. By default, this is 32768 bytes.
- **Summary Index Interval (milliseconds)** - The default indexing interval of the event summary in milliseconds. By default, this is 1000 milliseconds (1 sec).
- **Archive Index Interval (milliseconds)** - The default indexing interval of the event archive in milliseconds. By default, this is 30000 milliseconds (30 sec).
- **Index Limit** - The number of events to index in each index interval. By default, this is 1000 events.
- **Event Time Purge Interval (days)** - The number of days that event occurrence time are kept. By default, they are kept for 7 days. The minimum value is 1 and the maximum value is determined by the range of event time partitions. With the default configuration, the maximum value is 7 days.

- **Enable Event Flapping Detection** - Select this check box if you wish to enable event flapping detection. If an event is created and then cleared *flapping_threshold* times in *event_flapping_interval* time then an event of event flapping event class is created.
- **Event Flapping Event Class** - The event class under which generated flapping events belong.
- **Clear Event Heartbeats** - Click **Clear** to clear the event heartbeats.

Rebuilding the Events Index

If you encounter inconsistent search results, you can rebuild the events index.

- 1 Log in to an account on the Control Center master host that has permission to use the Control Center command-line interface.
- 2 Stop zeneventserver:

```
serviced service stop zeneventserver
```

- 3 Attach to the Resource Manager application.

```
serviced service attach Zenoss.resmgr
```

- 4 Delete the index:

```
rm -rf $ZENHOME/var/zeneventserver/index
```

Note In order to execute a command using `$ZENHOME` (`/opt/zenoss` for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for `serviced` commands.

- 5 Exit the container session.

```
exit
```

- 6 Restart zeneventserver:

```
serviced service start zeneventserver
```

Depending on the number of events in the database, it may take a significant amount of time for indexing to complete. Until every event is indexed, the number of events shown in the event console may be inconsistent.

Audit Logging

The audit log tracks user actions in syslog or log files. The system maintains logged information in a format optimized for searching and reporting.

Logged information can appear in several locations:

- Log file
- Rotating log files (limited by time or size)
- syslog

By default, the `$ZENHOME/log/audit.log` file stores the latest 10MB of data, with three rolling backups.

Note In order to execute a command using \$ZENHOME (/opt/zenoss for the zenoss user), you must be attached to the container holding the Resource Manager application. See the Control Center documentation for serviced commands.

Configuring the Audit Logs

Settings in the \$ZENHOME/etc/audit_log.conf configuration file determine the location and content of logged information output.

The audit_log.conf and audit_log.conf.example files are created at installation (if they do not exist).

An entry in the audit log indicates that a user attempted an action, but does not always indicate whether that action was successful. For example, a log entry stating that a user added a device simply indicates that the user created a job to add the device; however, the job could still fail when it runs at a later time.

As shown in the following sample, the configuration file contains examples and instructions for each of the output methods.

```
## Audit Log configuration file
## ## Initially this outputs up to 10 megs to ZENHOME/log/audit.log with
## 3 backups.
##
## To output to the syslog or somewhere else:
## - Uncomment the desired handlers and formatters, or create your own.
## - Update the "keys" lists under [handlers] and [formatters].
## - Update the "handlers" list under [logger_audit].
## - Restart Zope with "zenwebserver restart".
##
## To change the log severity level:
## - Update "level" under [logger_audit]
##
## This file has all the features of the Python logging file format:
## http://docs.python.org/library/logging.config.html#configuration-file-
## format

[loggers]
## DO NOT CHANGE
keys=audit

##
##
## List all output handlers here. (part 1 of 3)
## This should match part 3 below.
##
## Example: keys=syslog,file,rotatingfile,timedrotatingfile,console
##
##
[handlers]
keys=rotatingfile

##
##
## List all string formatters here. (part 2 of 3)
##
## Example: keys=syslog,file,console
##
##
[formatters]
```



```

keys=file

[logger_audit]
## DO NOT CHANGE
qualname=zen.audit
propagate=0

##
##
## This is the severity level of all audit messages.
## (DEBUG, INFO, WARNING, ERROR, CRITICAL)
##
## You can override the level of individual handlers below,
## or keep them as NOTSET to use this default level.
##
##
level=INFO

##
##
## List all output handlers here. (part 3 of 3)
## This should match part 1 above, except "handlers=" not "keys=".
##
## Example: handlers=syslog,file,rotatingfile,timedrotatingfile,console
##
##
handlers=rotatingfile
##### Output Handlers

## SysLog
##
## See http://docs.python.org/library/logging.handlers.html#sysloghandler
##
## Here are typical configurations:
##
## Linux: args=('/dev/log', handlers.SysLogHandler.LOG_USER)
## OS/X : args=('/var/run/syslog', handlers.SysLogHandler.LOG_USER)
## UDP : args=('localhost', handlers.SYSLOG_UDP_PORT),
##          handlers.SysLogHandler.LOG_USER)
##
##
##[handler_syslog]
##class=handlers.SysLogHandler
##level=NOTSET
##formatter=syslog
##args=()

## File
##
## See http://docs.python.org/library/logging.handlers.html#filehandler
##
## To store in ZENHOME/log:
class=Products.ZenUtils.configlog.ZenFileHandler
## To store elsewhere: class=FileHandler
##
## Format and example:
## args=(filename, mode, encoding, delay)
## args=('audit.log', 'a', None, True)
##
##
##[handler_file]

```

```

##class=Products.ZenUtils.configlog.ZenFileHandler
##level=NOTSET
##formatter=file
##args=('audit.log', 'a', None, True)

## RotatingFile
##
## See http://docs.python.org/library/
logging.handlers.html#rotatingfilehandler
##
## To store in ZENHOME/log:
class=Products.ZenUtils.configlog.ZenRotatingFileHandler
## To store elsewhere: class=handlers.RotatingFileHandler
##
## Format:
## args=(filename, mode, maxBytes, backupCount, encoding, delay)
##
## Example of one 10-meg file in ZENHOME/log/
## args=('audit.log', 'a', 10000000, 0, None, True)
##
## Example of ten 1-meg files in ZENHOME/log/audit/. The path must
already exist.
## args=('audit/audit.log', 'a', 1000000, 10, None, True)
##
##
[handler_rotatingfile]
class=Products.ZenUtils.configlog.ZenRotatingFileHandler
level=NOTSET
formatter=file
args=('audit.log', 'a', 10485760, 3, None, True)

## TimedRotatingFile
##
## See http://docs.python.org/library/
logging.handlers.html#timedrotatingfilehandler
##
## To store in ZENHOME/log:
class=Products.ZenUtils.configlog.ZenTimedRotatingFileHandler
## To store elsewhere: class=handlers.TimedRotatingFileHandler
##
## Format and example:
## args=(filename, when, interval, backupCount, encoding, delay, utc)
##
## Example of weekly log files for the past year in ZENHOME/log/audit/
## args=('audit/weekly.log', 'midnight', 7, 52, None, True, False)
##
##
##[handler_rotatingfile]
##class=Products.ZenUtils.configlog.ZenTimedRotatingFileHandler
##level=NOTSET
##formatter=file
##args=('audit/weekly.log', 'midnight', 7, 52, None, True, False)

## Console
##
## See http://docs.python.org/library/
logging.handlers.html#streamhandler
##
##

```

```
##[handler_console]
##class=StreamHandler
##level=NOTSET
##formatter=console
##args=(sys.stdout,)

##### String Formatters
##
## These must be uncommented if used by a handler above.
##
## See the very bottom of http://docs.python.org/library/
logging.config.html
##
##

##[formatter_syslog]
##format=zenoss[% (process) d]: % (message) s

[formatter_file]
format=% (asctime) s % (message) s
datefmt=%Y-%m-%d %H:%M:%S

##[formatter_console]
##format=Audit: % (asctime) s % (message) s
##datefmt=%H:%M:%S
```

After editing the `audit_log.conf` file, restart Zope with the command:

```
zenwebserver restart
```

Examples

While enabled, user actions are tracked as specified by the configuration file.

Example 1

```
Sep 12 12:55:10 zenoss[8432] user=hsolo action=SetDeviceClass
kind=Device device=/Devices/Server/Linux/devices/emailsrv05
deviceclass=/Devices/Server/SSH/Linux
```

In this example, user `hsolo` moved device "emailsrv05" from device class `/Server/Linux` to `/Server/SSH/Linux`.

Example 2

```
Sep 12 12:57:19 zenoss[8432] user=lskywalker action=Edit
kind=ThresholdClass maxval=100 minval=0 old_minval=-100
thresholdclass="/Devices/Server/Linux/rrdTemplates/Device/thresholds/
CPU pct"
```

In this example, user `lskywalker` edited values of threshold "CPU pct" by adding a max value of 100 and changing the min from -100 to 0.

Searching File Content

You can use central logging tools or grep to parse the configuration file content. Using Splunk, for example, searching for "device=*/localhost" finds any action on machines named localhost in any device class. Searching for "action=Add kind=User | table user username" creates a table that lists new users and which user added them.

Utility

The `zensendaudit` utility allows you to log custom messages. For example:

```
zensendaudit Hello world.
```

generates the output:

```
Message sent: Hello world.
```

Further, it logs this message to the configured syslog or files:

```
zenoss[9350]: user=admin type=ManualEntry comment="Hello world."
```

Stopping Audit Logging

Audit logging has minimal impact on system performance; however, you can disable audit logging by removing the `ZenPacks.zenoss.AuditLog` ZenPack. If re-installed, the system does not overwrite the existing configuration and example files, so that you can refer to the `audit_log.conf.example` file if needed.

Setting Portlet Permissions

By setting permissions, you determine which users can view and interact with portlets. Permissions settings restrict which Zope Access Control List (ACL) can access each portlet.

Before you can successfully set portlet permissions, you must assign the user a specific Resource Manager role. (You assign roles from the user edit page, from **Advanced > Settings**.) Each user role is mapped to one or more Zope ACL permissions, which allow you to restrict the portlets a permission level can see.

A user's specific portlet permissions are defined in part by Zope ACL permissions, and in part by the role to which he is assigned.

User Role to ACL Mapping

The following table shows how user roles map to ACLs.

| User Roles | ACLPermission |
|---------------------------|-----------------------------|
| ZenUser, ZenOperator | ZenCommon, View |
| ZenManager, Manager | ZenCommon, View, Manage DMD |
| No Role, Administered Obj | ZenCommon |

Setting Permissions

To set portlet permissions:

- 1 From the Navigation menu, select **ADVANCED**. The Settings page appears.

- 2 In the left panel, select **Portlets**. The Portlets page appears.
- 3 For one or more portlets in the Available Portlets list, select the permissions you want to apply.

Figure 126: Portlet Permissions

| Available Portlets | |
|-------------------------------------|------------------------------------|
| Device Issues | Users with ZenCommon permission ▼ |
| Google Maps | Users with View permission ▼ |
| Daemon Processes Down | Users with Manage DMD permission ▼ |
| Production States | Users with ZenCommon permission ▼ |
| Site Window | Users with View permission ▼ |
| Top Level Organizers | Users with View permission ▼ |
| Messages | Users with ZenCommon permission ▼ |
| Watch List | Users with ZenCommon permission ▼ |
| <input type="button" value="Save"/> | |

- 4 Click **Save**.

Troubleshooting: Users Cannot See All Portlets

You may mistakenly block users from being able to access some portlets. Often, this happens when a user has been set to see only particular devices. By default, this user will see only portlets set to the ZenCommon permission level. In effect, this blocks three of six portlets.

To remedy this problem, you can:

- Change the permission levels (on the Portlets page) to ZenCommon, or
- Change the user role to a role higher than "No Role."

Support Bundles

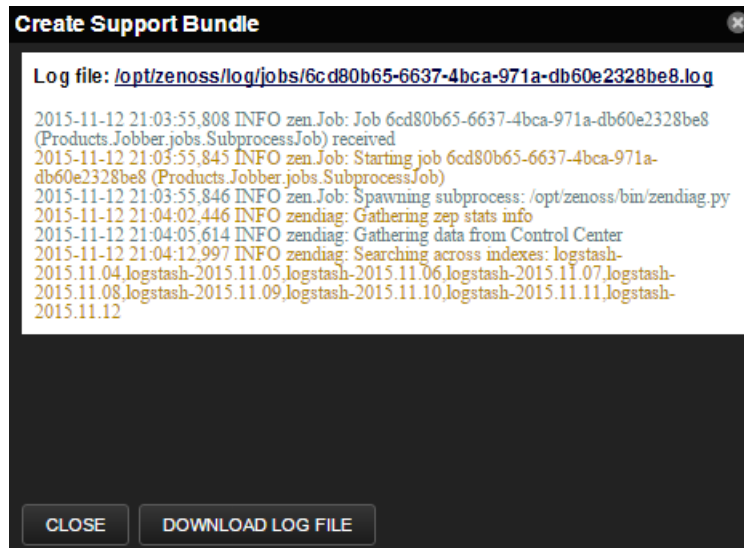
If you contact Zenoss Support about an issue with your Resource Manager system, you may be asked to send a support bundle and/or serviced and docker logs to help the technicians find a solution to your issue. These support bundles contain logs and system information that is valuable to the Support team.

Creating Support Bundles

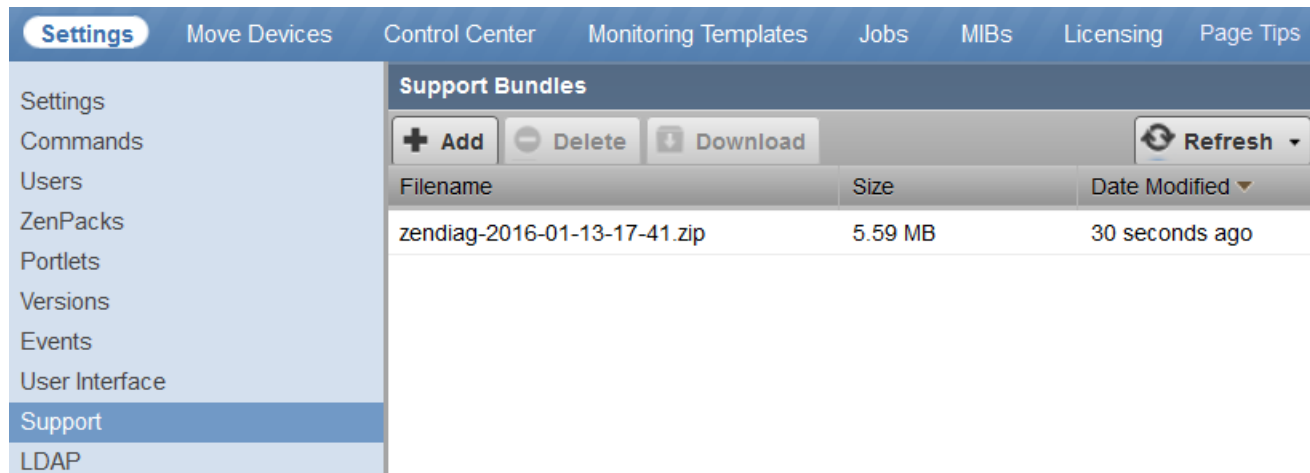
You may be asked to send a support bundle to Zenoss Support to help troubleshoot an issue with your system. This support bundle contains the logs and system status that the Support team needs to solve your issue.

To create a support bundle:

- 1 As an administrator, navigate to **ADVANCED > Support**. The Support Bundles window appears.
- 2 Click **Add**. A window appears indicating that the support bundle will continue to be gathered even after closing the window.
- 3 Click **OK**. The Create Support Bundle window displays. Click **Close**. The support bundle will continue to be gathered. You can monitor the job log for progress. You can continue working with the system during this time.

Figure 127: Create Support Bundle window

- 4 When the completed support bundle is displayed in the Support Bundles window, click on its row and click **Download** to send the support bundle .zip file to your browser's downloads folder.

Figure 128: Support Bundles window

Deleting Support Bundles

From time-to-time you may want to clean up your system by deleting old support bundles that are no longer needed for troubleshooting purposes.

To delete support bundles from your system:

- 1 As an administrator, navigate to **ADVANCED > Support**. The Support Bundles window appears.
- 2 Click the row of the support bundle(s) you want to delete. Be sure that you have selected the bundle(s) you want deleted since there is no confirmation window once you click **Delete**.
- 3 If you are sure of your selections, click **Delete**. The selected bundle(s) will be removed from the system.

Gathering serviced and docker logs

You may need to gather serviced and docker logs separately to help troubleshoot an issue, since these logs cannot be collected by the support bundle process.

To get serviced logs from the system:

- 1 Start a command-line session as `root`.
- 2 Execute the following command to create the `serviced.log` file in the `/tmp` directory.

```
journalctl -u serviced -o cat > /tmp/serviced.log
```

- 3 Execute the following command to create the `docker.log` file in the `/tmp` directory.

```
journalctl -u docker -o cat > /tmp/docker.log
```

Backup and Restore

Having accurate and tested system backups can mitigate problems caused by software or hardware issues. To provide a historical archive, you might back up on a regular schedule. Back up as needed when you want to move data from one instance to another or duplicate an instance for testing or failover purposes.

Because the backup process leverages snapshot functionality, you can take backups without the need to shut down the Resource Manager application container or the docker containers. The services are only momentarily suspended to enable reading the data.

The backup process saves a backup of the current system (all templates and services, the state of all services, and application data) to a compressed archive file (`.tgz`). The archive file can be restored on the same cluster or a similar cluster.

You can back up and restore by using the browser interface or the command-line interface (CLI). Both methods provide comparable results; however, the CLI command offers an option to specify application data directories in `/opt/serviced/var/volumes` to exclude from the backup.

Note The following considerations apply to backup and restore:

- Ensure that you have enough free space to receive and store backups; running low on available disk space results in errors and affects system performance.
 - Any backups that you took using Control Center 1.0.x cannot be restored in Control Center 1.1.x.
 - Restoring from a backup does not remove services that you added after taking the backup. That is, if you take a backup, add a service, and then restore from the backup, the service is not deleted as part of the restore process.
 - There may be occasions when a full backup is not necessary, e.g., a checkpoint prior to installing a ZenPack. In these cases, you can perform a snapshot of the system. In the event that you need to revert back to a snapshot, you can use the rollback feature (before which you will need to stop services that are used in the particular snapshot image). For more information on the snapshot and rollback feature, see [Snapshot and Rollback](#) on page 219.
-

Backing Up Resource Manager using the browser interface

To take a backup of Resource Manager:

- 1 Log in to the Control Center browser interface.

Control Center Applications

| Application | Description | Status | Deployment ID | Resource Pool | Public Endpoints | Actions |
|------------------------|-------------------------|--------|---------------|---------------|-------------------------------------|-------------------------|
| Internal Services | Internal Services | ✓ | Internal | N/A | N/A | N/A |
| Zenoss.resmgr (v5.1.1) | Zenoss Resource Manager | ✓ | Test | default | https://zenoss5.ip-10-111-23-46.443 | ▶ Start ■ Stop 🗑 Delete |

Application Templates

| Application Template | ID | Description |
|------------------------|----------------------------------|-------------------------|
| Zenoss.resmgr (v5.1.1) | b1f8fb3473d412a9a89555ea411d4b72 | Zenoss Resource Manager |

- 2 Click the **Backup / Restore** tab. The Backup / Restore page appears.
- 3 Click **Create Backup**. You will be asked to confirm your selection by clicking **Create Backup** again.
- 4 When the backup is complete, you will see the path of the `.tgz` file displayed. The default directory that receives backups is `/opt/serviced/var/backups`.

Control Center Backup / Restore

| File Name | Modified | Actions |
|--|-------------------------|------------------|
| /opt/serviced/var/backups/backup-2016-01-14-201737.tgz | Jan 14, 2016 2:38:45 PM | 🔄 Restore Backup |

Backing Up Resource Manager using the CLI

As an alternative to backing up Resource Manager from the Control Center browser interface, you can use the command-line interface (CLI).

You can back up your entire system or specify one or more tenant volumes to exclude from the backup. For example, to save resources you might exclude the tenant volume that contains performance data, `hbase-master`. If you want to take a backup to restore on another system, you might exclude the tenant volumes for the database and index, `mariadb-events` and `zeneventserver`.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Choose what to back up:
 - To exclude one or more tenant volumes, skip to Step 3.
 - To include the entire system, execute the following command. The default directory that receives backups is `/opt/serviced/var/backups`. If your directory path is different, specify your path.

```
serviced backup /opt/serviced/var/backups
```

After a successful backup, the system displays the new backup file name. For example:

```
backup-2016-10-04-203717.tgz
```

- 3 Optional: Identify and exclude tenant volumes from the backup:

- a Display the Control Center tenant identifier. By default, Control Center stores application data in `/opt/serviced/var/volumes`. If your directory path is different, specify your path.

```
ls /opt/serviced/var/volumes
```

Example result:

```
cvs0ul2tmvjcit0lr7p0d8bx
```

- b Display directories under the Control Center tenant identifier. Replace *Tenant-ID* with the identifier displayed in the previous step.

```
ls /opt/serviced/var/volumes/Tenant-ID
```

Example result:

Table 7:

| | | | |
|---------------|--------------------------|-------------------|-----------------------|
| hbase-master | hbase-zookeeper-2 | localhost | mariadb-events |
| rabbitmq | var-zenpacks | zeneventserver | zenoss-custom-patches |
| zenoss-export | zenpack-backup | hbase-zookeeper-1 | hbase-zookeeper-3 |
| lost+found | mariadb-model | redis | zencatalogservice |
| zenjobs | zenoss-custom-patches-pc | zenoss-var-ext | |

- 4 Optional: Exclude one tenant volume from the backup; in this example, `hbase-master`. Replace *Tenant-ID* with your tenant identifier.

```
serviced backup /opt/serviced/var/backups --exclude /opt/serviced/var/volumes/Tenant-ID/hbase-master
```

Exclude multiple tenant volumes; in this example, `mariadb-events` and `zeneventserver`. Replace *Tenant-ID* with your tenant identifier.

```
serviced backup /opt/serviced/var/backups --exclude /opt/serviced/var/volumes/Tenant-ID/mariadb-events --exclude /opt/serviced/var/volumes/Tenant-ID/zeneventserver
```

- 5 Optional: Update automated backups to exclude tenant volumes.

Restoring Resource Manager from a backup

You can restore your Resource Manager instance from any given backup `.tgz` file. You can duplicate your instance on a similar cluster for testing or failover purposes by restoring a backup `.tgz` file to a new, similarly configured cluster.

- 1 Log in to the Control Center browser interface.

Control Center Applications

Applications

| Application | Description | Status | Deployment ID | Resource Pool | Public Endpoints | Actions |
|------------------------|-------------------------|---------|---------------|---------------|-------------------------------------|-------------------|
| Internal Services | Internal Services | Running | Internal | N/A | N/A | N/A |
| Zenoss.resmgr (v5.1.1) | Zenoss Resource Manager | Running | Test | default | https://zenoss5.ip-10-111-23-46-443 | Start Stop Delete |

Last Update: a few seconds ago Showing 2 Results

Application Templates

| Application Template | ID | Description |
|------------------------|----------------------------------|-------------------------|
| Zenoss.resmgr (v5.1.1) | b1f8fb3473d412a9a89555ea411d4b72 | Zenoss Resource Manager |

Last Update: a few seconds ago Showing 1 Result

- 2 In the **Applications** table, identify the name of the Resource Manager instance to modify.
- 3 Stop the instance, and verify its subservices are stopped.
 - a In the **Actions** column of the **Applications** table, click **Stop**.

Stop Service

Please choose which services to stop:

- Stop service Zenoss.resmgr
- Stop service Zenoss.resmgr and 52 child services

Cancel Stop Service Stop Service and 52 Children

- b In the **Stop Service** dialog, click **Stop Service and 52 Children**.
 - c In the **Applications** column of the **Applications** table, click the name of the stopped instance, and then scroll down to the **Services** table to verify everything is stopped.
- 4 Click the **Backup / Restore** tab. The Backup / Restore page appears listing your backup files.

Control Center Backup / Restore

Backup / Restore

| File Name | Modified | Actions |
|--|-------------------------|----------------|
| /opt/serviced/var/backups/backup-2016-01-14-201737.tgz | Jan 14, 2016 2:38:45 PM | Restore Backup |

Last Update: a few seconds ago Showing 1 Result

- 5 Click **Restore Backup** next to the backup file you want to use to restore your Resource Manager application. You will be asked to confirm your selection by clicking **Restore**.

Restore

The restore process cannot be reversed.

Cancel Restore

- 6 When the restore is finished, click the **Applications** tab, then click **Start** next to the Resource Manager instance you just restored.

Snapshot and Rollback

There are times when taking a backup of your Resource Manager system is not practical, but you need a way to safeguard potential risk when changing your system. In these cases it may be more beneficial to take a snapshot of the system and then rollback if your changes cause a failure or other degradation of the application.

The main advantage of taking a snapshot over a backup is that you do not have to stop any services when taking a snapshot. You can take a snapshot whenever you need a savepoint for your docker image, e.g., before committing container changes or installing a ZenPack). If you want to back out of your change, you can use the rollback functionality to go back to the snapshot image.

When you take a snapshot, a tag is created for the docker image of each service and a backup of the distributed file system is taken. You do not have to stop any services to create the snapshot, but if you do need to rollback, you will need to stop the particular services that are used by the snapshot image.

Note Rolling back from a snapshot does not remove services that you added after taking the snapshot. That is, if you take a snapshot, add a service, and then roll back, the service remains on the system; it is not deleted as part of the roll back.

Since snapshots are meant as a savepoint and are not meant to be used as an historical backup, there is a default time to live (TTL) value of 12 hours for any snapshot. However, if you decide that you want to keep a snapshot beyond the TTL, you can tag the snapshot which will prevent it from being deleted after the TTL has expired.

Taking a Snapshot

You need to know the `SERVICEID` of the service you want to take a snapshot of, typically the `zenoss.resmgr` service.

Snapshots can only be created from the command-line interface (CLI). You do not have to stop any services prior to taking the snapshot.

- 1 Log in to the Control Center host as a user with `sudo` and `docker` privileges.
- 2 Find the `SERVICEID` of the service you want to snapshot.

```
serviced service list
```

- 3 Take the snapshot.

```
serviced snapshot add SERVICEID
```

The name of the snapshot is displayed. The format of the name is a combination of the `SERVICEID` and a date stamp. To verify its existence, you can list all the snapshots. (By default, snapshots have a time-to-live (TTL) of 12 hours.)

```
serviced snapshot list
```

- 4 (Optional) If you want to keep the snapshot for longer than the TTL value, you can tag the snapshot using the following command.

```
serviced snapshot tag SNAPSHOTID TAG-NAME
```

To untag a snapshot, execute the following command.

```
serviced snapshot untag SERVICEID TAG-NAME
```

Note Once you untag a snapshot, it will be subject to the TTL value.

Rollback to a Snapshot

You can rollback your application to a previous state that existed when you took a snapshot. You must stop the services that are used by the image when the snapshot was taken. The distributed file system will also be rolled back to the state that existed when the snapshot was taken.

Rollbacks can only be performed on the CLI. You must know the `SNAPSHOTID` you want to rollback to.

- 1 Log in to the Control Center host as a user with `sudo` and `docker` privileges.
- 2 Find the `SNAPSHOTID` of the snapshot to which you want to rollback.

```
serviced snapshot list
```

- 3 Rollback to the snapshot. Using the `--force-restart` flag automatically stops the affected services before rollback and starts them after completion.

```
serviced snapshot rollback SNAPSHOTID --force-restart
```

Adding storage for application data or snapshots

To perform this procedure, you need an unformatted primary partition on the Control Center master host.

This procedure prepares a partition as a device mapper thin pool on a Control Center master host and adds it to the thin pool for applications. The storage can be used for application data or for snapshots.

- 1 Log in to the Control Center master host as `root`, or as a user with superuser privileges.
- 2 Identify the current size of the thin pool for applications.

```
lvs -o lv_size serviced
```

Example result:

```
LSize
150.04g
```

In this example, the current size is 150GB. Record the size for use in a subsequent step.

- 3 Identify the partition to prepare, and record its size.
Replace *Device* with the virtual disk added previously:

```
lsblk -p --output=NAME,SIZE,TYPE Device
```

Example output:

```
NAME          SIZE TYPE
/dev/sdc      150G disk
|--/dev/sdc1  150G part
```

In this example, the partition to prepare is `/dev/sdc1`, and its size is 150GB. Record the size for use in a subsequent step.

- 4 Create a new physical volume.

Replace *Partition* with the partition identified previously:

```
pvcreate Partition
```

On success, the `pvcreate` command displays the name of the new physical volume, which is the same name as the partition.

- 5 Extend the volume group used for applications, `serviced`.

Replace *Physical-Volume* with the name of the physical volume created in the previous step:

```
vgextend serviced Physical-Volume
```

- 6 Add space to the thin pool for applications with the updated volume group.

Replace *New-Space* with the size of the partition, identified in the second step. Be sure to include the units ("G" for gigabytes):

```
lvextend -L+New-SpaceG serviced/serviced-pool
```

The `lvextend` command extends both the data and metadata areas of the thin pool.

- 7 Optional: Add storage space for application data.

The preceding steps expand the amount of storage available for snapshots. To use the additional storage for application data, perform this step.

- a Add the size of the application data storage area, identified in the first step, to the size of the new partition, identified in the second step.
- b Display the Control Center tenant ID.

```
ls /opt/serviced/var/volumes
```

Example result:

```
dovygh6m2gz4zgdzxb2ud70g9
```

- c Update Control Center.

Replace *Tenant-ID* with the identifier displayed in the previous substep, and replace *Total-Space* with the total amount of application data storage space. Be sure to include the units ("G" for gigabytes):

```
serviced-storage resize -d /opt/serviced/var/volumes \  
-o dm.thinpooldev=/dev/mapper/serviced-serviced--pool \  
Tenant-ID Total-SpaceG
```

Working with the Job Manager

The Job Manager runs background tasks, such as discovering a network or adding a device. When you ask the system to perform one of these tasks, it adds a job to the queue. Jobs are run by the `zenjobs` daemon.

Not all actions are performed in the Job Manager. Some jobs are run automatically in the foreground. Others, such as moving devices, depend on user interface configuration settings.

When running jobs in the foreground, do not navigate away from the current page until the action completes.

Viewing the Job Manager

To access the Job Manager:

- 1 From the Navigation menu, select **ADVANCED**. The Settings page appears.
- 2 From the menu, select **Jobs**. The jobs list appears and shows information about all jobs currently in the system.

Figure 129: Job Manager

Settings Control Center Monitoring Templates **Jobs** Page Tips

Background Jobs

[-] Delete [X] Abort Refresh

| Status | Description | Scheduled | Started | Finished | Created By |
|---------|---|------------|------------|------------|------------|
| Success | Create 10.171.100.14 under /Network/Cisco/6500 | 2 days ago | 2 days ago | 2 days ago | admin |
| Failure | Discover and model device 10.171.100.14 as /Networ... | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Create 10.171.100.92 under /Network/Cisco/Nexus/70... | 2 days ago | 2 days ago | 2 days ago | admin |
| Failure | Discover and model device 10.171.100.92 as /Networ... | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Create 10.171.100.88 under /Network/Cisco/Nexus/10... | 2 days ago | 2 days ago | 2 days ago | admin |
| Failure | Discover and model device 10.171.100.88 as /Networ... | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Create 10.171.54.9 under /Network/Cisco | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Discover and model device 10.171.54.9 as /Network/... | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Create 10.171.100.107 under /Network/Cisco/Nexus/6... | 2 days ago | 2 days ago | 2 days ago | admin |
| Failure | Discover and model device 10.171.100.107 as /Netwo... | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Create 10.171.53.1 under /Network/Cisco | 2 days ago | 2 days ago | 2 days ago | admin |
| Success | Discover and model device 10.171.53.1 as /Network/... | 2 days ago | 2 days ago | 2 days ago | admin |

DISPLAYING 1 - 13 of 33 ROWS

Job Log

Log file: [/opt/zenoss/log/jobs/c4583592-2506-4466-b909-854ebfa20258.log](#)

```
2015-08-04 17:59:20,643 INFO zen.Job: Job c4583592-2506-4466-b909-854ebfa20258 (Products.ZenModel.ZDeviceLoader.CreateDeviceJob)
received
2015-08-04 17:59:20,657 INFO zen.Job: Starting job c4583592-2506-4466-b909-854ebfa20258
(Products.ZenModel.ZDeviceLoader.CreateDeviceJob)
2015-08-04 17:59:21,051 INFO zen.Job: Job c4583592-2506-4466-b909-854ebfa20258 finished with result
/zport/dmd/Devices/Network/Cisco/Nexus/7000/devices/10.171.100.92
```

The following information is displayed about the jobs:

- **Status** - Shows the current job status. Status options are Pending (waiting for zenjobs to begin running), Running, Succeeded, and Failed.
- **Description** - Provides a description of the job.
- **Scheduled** - Shows when the job was scheduled to begin.
- **Started / Finished** - Provide information about the time period in which the job ran.
- **Created By** - User that created the job.

The lower section of the page displays the job log for the job selected in the list. You can view job info here, or by viewing the log file.

Stopping and Deleting Jobs

To stop a job, select it in the list, and then click **Abort**. The zenjobs daemon will not run the job.

To remove a job from the system, select it and then click **Delete**.

Configuring Jobs

You can determine, when moving devices, whether the action is performed immediately or as a job. By default, if you select five or more devices, the move action is performed as a job. To adjust this setting:

- 1 Select **Advanced > Settings**.
- 2 Select **User Interface** in the left panel.
- 3 Enter a value for Device Move Job Threshold, and then click **Save**.

Running the zenjobs Daemon

You can stop and start the zenjobs daemon from the command line, and from **Advanced > Settings (Daemons selection)**.

Versions and Update Checks

You can check to see if there are newer versions of Resource Manager available. By default, the system is configured to check for this each day; however, you can verify this at any time from **Advanced > Settings > Versions**. Click **Check Zenoss Version Now** to check for updates.

When the system performs its daily update check (enabled by the Check for Updates Daily option), by default it sends a utilization report containing non-identifying system metrics to Zenoss. Types of information sent about your installation include:

- Device count
- Number of systems, locations, and groups
- Installed ZenPacks
- Event count
- System statistics (such as operating system and version)
- Number of users
- Performance data (such as memory and file system space used)

To view a report of the collected data, click **Show Report**. To download an encrypted version of the utilization report, click **Save Encrypted**.



SNMP Device Preparation

This section provides information about SNMP support and lists Net-SNMP configuration settings that are required by the system.

Net-SNMP

By default, Net-SNMP does not publish the full SNMP tree. Check to see if that is currently the case on a device and configure it correctly.

- 1 Confirm `snmpd` is running:

```
> snmpwalk -v 2c -cpublic <your device name> system
```

- 2 Retrieve the IP table for the device with `snmpwalk`:

```
> snmpwalk -v 2c -cpublic <your device name> ip
```

Typical SNMP View:

```
view systemview included .1 view systemview included .1.3.6.1.2.1.25.1
access notConfigGroup "" any noauth exact systemview none none
```

SNMP v3 Support

Resource Manager provides support for SNMP v3 data collection.

The following configuration properties control the authentication and privacy of these requests:

- **zSnmAuthType**- Use "MD5" or "SHA" signatures to authenticate SNMP requests. If only `zSnmAuthType` and `zSnmAuthPassword` are set, then the message is sent with authentication but no privacy.
- **zSnmAuthPassword**- Shared private key used for authentication. Must be at least 8 characters long.
- **zSnmPrivType**- "DES" or "AES" cryptographic algorithms. If `zSnmPrivType` and `zSnmPrivPassword` are set, then the message is sent with privacy and authentication. You cannot set a `PrivType` and `PrivPassword` without also setting an `AuthType` and `AuthPassword`. If neither `Priv` nor `Auth` values are set, then the message is sent with no authentication or privacy.
- **zSnmPrivKey**- Shared private key used for encrypting SNMP requests. Must be at least 8 characters long.
- **zSnmSecurityName**- Security Name (user) to use when making SNMPv3 requests.

If monitoring SNMPv3 devices, make sure that msgAuthoritativeEngineID (also known as snmpEngineID or Engine ID) is not shared by two devices. It must be unique for each device.

Advanced Encryption Standard

SNMPv3 encryption using the Advanced Encryption Standard (AES) algorithm is supported only if the host platform net-snmp library supports it.

You can determine if your platform supports AES by using the following test:

```
$ snmpwalk -x AES 2>&1 | head -1
```

If the response is:

```
"Invalid privacy protocol specified after -x flag: AES"
```

then your platform does not support AES encryption for SNMPv3.

If the response is:

```
"No hostname specified."
```

Then your platform supports AES.

Community Information

Add these lines to your `snmp.conf` file.

This line will map the community name "public" into a "security name":

```
# sec.name source community
```

```
com2sec notConfigUser default public
```

This line will map the security name into a group name:

```
# groupName securityModel securityName
```

```
group notConfigGroup v2c notConfigUser
```

This line will create a view for you to let the group have rights to:

```
# Make at least snmpwalk -v 1 localhost -c public system fast again.
```

```
# name incl/excl subtree mask(optional)
```

```
view systemview included .1
```

This line will grant the group read-only access to the systemview view.

```
# group context sec.model sec.level prefix read write notif access
notConfigGroup "" any noauth exact systemview none none
```

System Contact Information

It is also possible to set the `sysContact` and `sysLocation` system variables through the `snmpd.conf` file:

```
syslocation Unknown (edit /etc/snmp/snmpd.conf)
```

```
syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)
```

```
# Added for support of bcm5820 cards. pass .1 /usr/bin/ucd5820stat
```

Extra Information

For more information, see the `snmpd.conf` manual page, and the output of the `snmpd -H` command.

```
trapcommunity public
```

```
trapsink default
```

B

Syslog Device Preparation

Forwarding Syslog Messages from UNIX/Linux Devices

Resource Manager has its own syslog server (`zensyslog`). Managed devices should point their syslog daemons to the system.

To do this, edit the `/etc/rsyslog.conf` file and add an entry, where 1.2.3.4 is the `zensyslog` IP:

- 1 Log in to the target device as a super user.
- 2 Open the `/etc/rsyslog.conf` file with a text editor (such as `vi`).
- 3 Enter `*.debug`, and then press the Tab key.
- 4 Enter the host name or IP address of the server. For example:

```
*.debug @192.168.X.X
```

- 5 Save the file and exit the file editor program.
- 6 Restart the Syslog service using the command below:

```
/etc/init.d/syslog restart
```

Forwarding Syslog Messages from a Cisco IOS Router

Here are some links to Cisco commands to turn on syslog. Typically, it is easier to use syslog than SNMP traps from network devices. The most basic IOS command to send syslog messages is:

```
logging 1.2.3.4
```

Other Cisco Syslog Configurations

Here are some additional configurations for other Cisco devices. To set up these configurations follow the following steps using the configurations that follow below.

- 1 Log in to the target router.
- 2 Type the command `enable` at the prompt.
- 3 Once you are prompted for a password, enter the correct password.
- 4 Type the command `config` at the prompt.
- 5 Type the command `terminal` at the configuration prompt.

- 6 At the prompt, Set the Syslog forwarding mechanism. See example below:

```
logging <IP address of the server>
```

- 7 Exit out all the prompts to the main router prompt.

Catalyst

```
set logging server enable set logging server 192.168.1.100 set logging
level all 5 set logging server severity 6
```

Local Director

```
syslog output 20.5 no syslog console syslog host 192.168.1.100
```

PIX Firewalls

```
logging on logging standby logging timestamp logging trap notifications
logging facility 19 logging host inside 192.168.1.100
```

Forwarding Syslog Messages from a Cisco CatOS Switch

To forward a syslog message from a Cisco CatOS switch:

- 1 Log in to the target switch.
- 2 Type the command enable at the prompt.
- 3 Enter the password when prompted.
- 4 Set the Syslog forwarding mechanism; for example:

```
set logging server <IP address of the server>
```

- 5 You can set the types of logging information that you want the switch to provide with the commands below as examples:

```
set logging level mgmt 7 default set logging level sys 7 default set
logging level filesys 7 default
```

Forwarding Syslog Messages using Syslog-ng

Here is an example for FreeBSD and Linux platforms.

- 1 Log in to the target device as a super user.
- 2 Open `/etc/syslog-ng/syslog-ng.conf` file with a text editor (e.g `vi`).
- 3 Add source information to file. See the following examples:

FreeBSD:

```
source src { unix-dgram("/var/run/log"); internal ();};
```

Linux: (will gather both system and kernel logs)

```
source src { internal(); unix-stream("/dev/log" keep-alive(yes) max-connections(100)); pipe("/proc/kmsg"); udp(); };
```

- 4 Add destination information (in this case, the server). For example:

```
log { source(src); destination(zenoss); };
```

C

TALES Expressions

Use TALES syntax to retrieve values and call methods on Resource Manager objects. Several areas accept TALES syntax; these include:

- Command templates
- User commands
- Notifications
- zLinks

Commands (those associated with devices and those associated with events) can use TALES expressions to incorporate data from the related devices or events. TALES is a syntax for specifying expressions that let you access the attributes of certain objects, such as a device or an event.

For additional documentation on TALES syntax, see the TALES section of the [Zope Page Templates Reference](#).

Depending on context, you may have access to a device, an event, or both. Following is a list of the attributes and methods you may want to use on device and event objects. The syntax for accessing device attributes and methods is `${dev/attributename}`. For example, to get the `manageIp` of a device you would use `${dev/manageIp}`. For events, the syntax is `${evt/attributename}`.

A command to ping a device might look like this. (The `${. .}` is a TALES expression to get the `manageIp` value for the device.)

```
ping -c 10 ${device/manageIp}
```

Examples

- DNS Forward Lookup (assumes `device/id` is a resolvable name)

```
host ${device/id}
```

- DNS Reverse Lookup

```
host ${device/manageIp}
```

- SNMP Walk

```
snmpwalk -v 2c -c${device/zSnmpCommunity} ${device/manageIp} system
```

To use these expressions effectively, you must know which objects, attributes, and methods are available, and in which contexts. Usually there is a device that allows you to access the device in a particular context. Contexts related to a particular event usually have event defined.

TALES Device Attributes

The following table lists available device attributes.

| Attribute | Description |
|-----------------------------|---|
| getId | The primary means of identifying a device within the system |
| getManageIp | The IP address used to contact the device in most situations |
| productionState | The production status of the device: Production, Pre-Production, Test, Maintenance or Decommissioned. This attribute is a numeric value, use getProductionStateString for a textual representation. |
| getProductionStateString | Returns a textual representation of the productionState |
| snmpAgent | The agent returned from SNMP collection |
| snmpDescr | The description returned by the SNMP agent |
| snmpOid | The oid returned by the SNMP agent |
| snmpContact | The contact returned by the SNMP agent |
| snmpSysName | The system name returned by the SNMP agent |
| snmpLocation | The location returned by the SNMP agent |
| snmpLastCollection | When SNMP collection was last performed on the device. This is a DateTime object. |
| getSnmpLastCollectionString | Textual representation of snmpLastCollection |
| rackSlot | The slot name/number in the rack where this physical device is installed |
| comments | User entered comments regarding the device |
| priority | A numeric value: 0 (Trivial), 1 (Lowest), 2 (Low), 3 (Normal), 4 (High), 5 (Highest) |
| getPriorityString | A textual representation of the priority |
| getHWManufacturerName | Name of the manufacturer of this hardware |
| getHWProductName | Name of this physical product |
| getHWProductKey | Used to associate this device with a hardware product class |
| getOSManufacturerName | Name of the manufacturer of this device's operating system. |
| getOSProductName | Name of the operating system running on this device. |
| getOSProductKey | Used to associate the operating system with a software product class |
| getHWSerialNumber | Serial number for this physical device |
| getLocationName | Name of the Location assigned to this device |
| getLocationLink | Link to the system page for the assigned Location |
| getSystemNames | A list of names of the Systems this device is associated with |
| getDeviceGroupNames | A list of names of the Groups this device is associated with |

| Attribute | Description |
|-----------------------|---|
| getLastChangeString | When the last change was made to this device |
| getLastPollSnmpUpTime | Uptime returned from SNMP |
| uptimeStr | Textual representation of the SNMP uptime for this device |
| getPingStatusString | Textual representation of the ping status of the device |
| getSnmpStatusString | Textual representation of the SNMP status of the device |

TALES Event Attributes

The following table lists available event attributes.

| Attribute | Description |
|---------------|--|
| agent | Collector name from which the event came (such as zensyslog or zentrap). |
| component | Component of the associated device, if applicable. (Examples: eth0, httpd.) |
| count | Number of times this event has been seen. |
| dedupid | Key used to correlate duplicate events. By default, this is: device, component, eventClass, eventKey, severity. |
| device | ID of the associated device, if applicable. |
| DeviceClass | Device class from device context. |
| DeviceGroups | Device systems from device context, separated by . |
| eventClass | Event class associated with this device. If not specified, may be added by the rule process. If this fails, then will be /Unknown. |
| eventClassKey | Key by which rules processing begins. Often equal to component. |
| eventGroup | Logical group of event source (such as syslog, ping, or nteventlog). |
| eventKey | Primary criteria for mapping events into event classes. Use if a component needs further de-duplication specification. |
| eventState | State of event. 0 = new, 1 = acknowledged, 2 = suppressed. |
| evid | Unique ID for the event. |
| facility | syslog facility, if this is a syslog event. |
| firstTime | UNIX timestamp when event is received. |
| ipAddress | IP Address of the associated device, if applicable. |
| lastTime | Last time this event was seen and its count incremented. |
| Location | Device location from device context. |
| manager | Fully qualified domain name of the collector from which this event came. |
| message | Full message text. |
| nteventid | nt event ID, if this is an nt eventlog event. |
| priority | syslog priority, if this is a syslog event. |
| prodState | prodState of the device context. |

| Attribute | Description |
|----------------|---|
| severity | the severity of the event expressed as a number (0, 1, 2, 3, 4, or 5) |
| severityString | the severity of the event expressed as a string (Clear, Debug, Info, Warning, Error, or Critical) |
| stateChange | Time the MySQLrecord for this event was last modified. |
| summary | Text description of the event. Limited to 150 characters. |
| suppid | ID of the event that suppressed this event. |
| Systems | Device systems from device context, separated by . |

Configuration Properties and Custom Properties

Configuration properties and custom properties also are available for devices, and use the same syntax as shown in the previous sections.

D

Monitoring Large File Systems

By default, Resource Manager uses the Host Resources MIB to monitor file systems. A defect in the implementation of the Host Resources MIB in `net-snmp` causes file systems larger than 16TB to report incorrect utilization metrics, such that you might observe file system utilization values greater than 100%.

Note Resource Manager uses the Host Resources MIB in the `ethernetCsmacd` template (rather than the UCD `dskTable` MIB) by default because most of the systems Resource Manager monitors do not have the UCD `dskTable` MIB enabled.

To work around this deficiency, Resource Manager can instead use the UCD `dskTable` MIB to monitor file system utilization.

Procedure

To use the UCD `dskTable` MIB, you must modify the configuration of your data sources, thresholds, and graphs in your `FileSystem` template:

- 1 Create an SNMP data source named `dskPercent`.
- 2 Set the OID of the new data source to:

```
1.3.6.1.4.1.2021.9.1.9
```

- 3 Modify your thresholds to use the `dskPercent` data point. Remove any calculations that were associated with the Host Resource MIB data point. The `dskPercent` data point is reported as an integer from 0 to 100.
- 4 Modify your graphs to use the `dskPercent` data point and thresholds.
- 5 Enable the UCD `dskTable` MIB on your managed hosts:
 - a Add the following line to your `/etc/snmp/snmpd.conf` file:

```
includeAllDisks 0%
```

- b Restart the `snmpd` daemon.

The UCD `dskTable` MIB associates different indexes with the file systems than the Host Resources MIB. As a result, you must remodel the devices to fetch the UCD `dskTable` indexes, and to begin plotting data on the `dskPercent` graph.



Managing Multi-Realm Networks

Resource Manager supports overlapping IP namespaces through the optional `ZenPacks.zenoss.MultiRealmIP` ZenPack.

Note This ZenPack is not installed when Resource Manager is installed. To download it, visit the [Zenoss Support](#) site.

With this ZenPack, Resource Manager can prefix a realm identifier to the IP addresses of a network, enabling unified monitoring.

There are two primary use cases for using multi-realm IP management.


- A large company that manages multiple locations that have the same network spaces defined across these multiple locations and as a result have created multiple overlapping IP spaces and Resource Manager needs a way to identify each separate IP space in the system.
- Service Providers responsible for monitoring multiple customers where the customers have created independent networks and IP spaces that are unique to their location, but not unique to the Service Provider.

The essential workflow for creating and using IP Realms is that first you need to create the IP realms and then associate these realms with a collector. The associations between IP Realms and actual devices is made automatically by the device's association with the collector. All devices on a collector are associated with the realm for that collector.

Example System

The following diagram lays out an example setup. It has a central Resource Manager server in the 10.10.10.0/24 network. The network local to the server is considered the default network within the system. The default network is treated exactly the same as a Resource Manager system without this ZenPack installed.

There are two other networks shown (r1 and r2) which are behind a firewall and have the same IP space 192.168.0.0/24. Each realm has a distributed collector located within it. The collector can be accessed from the Resource Manager server using a IP translation from the firewall to map the address accessible from in front of the firewall to an address behind the firewall. Remote collectors in a multi-realm setup must be accessible from the central server using SSH.

Figure 130: Example IP Realm

System Setup

Set up Resource Manager following the example system described previously.

Note If you do not have overlapping IP space this example can be created using collectors within the same network. To create the example, add a machine multiple times once per collector, making sure to change the name of the device as it is added. The result is similar to a real realm setup.

Under multi-realm IP networks, device names *must* be unique even though the IP addresses will overlap.

On certain server configurations, if a distributed collector is configured, a "zenpack command failed" error occurs when installing this ZenPack. If you encounter this error, then run the following grant (as MySQL root):

```
grant super on *.* to 'zenoss'@'{FQDN_of_Zenoss_host}' identified by
'zenoss';
```

where the first 'zenoss' is the user account that Resource Manager uses to access MySQL, and the second 'zenoss' is that account's password.

Adding Realms

To add a realm:

- 1 Navigate to **Infrastructure > Networks**.
- 2 From the Add menu, select **Add IP Realm**. Add the realms r1 and r2.

Adding Collectors to Realms

To add a collector to a realm:

- 1 Add the two collectors that are installed in each realm. Distributed collectors now have an **IP Realm** field on their configuration screen set each collector to the appropriate realm configured above.
- 2 Change each collector so that it is in the correct realm.

Adding Devices to Realms

Perform the following to add devices to different realms:

- 1 Now we are ready to add devices to the system. As mentioned previously, adding the same device to the system twice can simulate a multi-realm setup. Add a device called `A.test` making sure that when it is added the collector is set to one of the remote collectors, and not `localhost`.
- 2 Rename the device.
- 3 Add the device a second time using a different remote collector, again not `localhost`.
- 4 After the device is loaded, select Software and follow the network link on one of the interfaces. Notice that the network has been created underneath the realm created earlier. This configuration is at the heart of multi-realm, as networks are discovered they are created within each realm. Monitoring is now happening on each representation of the device from the different collectors in different overlapping realms.

As another test try searching by IP from the top-level search. Two devices will be returned -- one within each realm.

Notes

- If an event contains the unique name of a device then it is straight-forward to assign it to the proper device. If only the IP address is sent the event will be assigned by looking up the IP within the context of the realm.
- If a device is moved between realms it must be remodeled so that its IPs are placed in the proper location.
- The **Network Map** only supports the display the default realm.

Glossary

aggregation pools

A logical bundling of multiple physical network interfaces, commonly known as a port channel. For example, the Per Chassis Ethernet Pools includes all links from all chassis to all fabric interconnects which is used for chassis bandwidth balance comparison. For more examples, see the Aggregation Pools component section of CiscoUCS devices.

bandwidth utilization

The total amount of bandwidth being used by an aggregation pool, a port, a fabric interconnect, a FEX, etc.

component

Object contained by a device. Components include interfaces, OS processes, file systems, CPUs, and hard drives.

integrated infrastructure

A bundle of compute, storage, networking, and virtualization components. Most integrated infrastructures are bought as one from a vendor:

- NetApp FlexPod
- VCE Vblock
- EMC VSPEX

All of these have UCS as the common compute element, Nexus as networking components, and VMware as virtualization.

data point

Data returned from a data source. In many cases, there is only one data point for a data source (such as in SNMP); but there may also be many data points for a data source (such as when a command results in the output of several variables).

data source

Method used to collect monitoring information. Example data sources include SNMP OIDs, SSH commands, and perfmon paths.

device

Primary monitoring object in the system. Generally, a device is the combination of hardware and an operating system.

device class

Special type of organizer used to manage how the system models and monitors devices (through configuration properties and monitoring templates).

discovery

Process by which Resource Manager gathers detailed information about devices in the infrastructure. Results of discovery are used to populate the model.

event

Manifestation of important occurrence within the system. Events are generated internally (such as when a threshold is exceeded) or externally (such as through a syslog message or SNMP trap).

event class

Categorization system used to organize event rules.

event rules

Controls how events are manipulated as they enter the system (for example, changing the severity of an event). Configuration properties configure event rules.

graph

Displays one or more data points, thresholds, or both.

headroom

The unused bandwidth in an aggregation pool, a port, a fabric interconnect (FI), a FEX, etc. For example, if an aggregation pool including 4 ports between a chassis and the FIs has 40 GB of capacity and if the bandwidth use of that pool is 25 GB, then the headroom is 15 GB.

managed resource

Servers, networks, virtual machines, and other devices in the IT environment.

model

Representation of the IT infrastructure. The model tells the system "what is out there" and how to monitor it.

monitoring template

Description of what to monitor on a device or device component. Monitoring templates comprise four main elements: data sources, data points, thresholds, and graphs.

notification

Sends email or pages to system users or groups when certain events occur.

organizer

Hierarchical system used to describe locations and groups within the application. Resource Manager also includes special organizers, which are classes that control system configuration.

out of balance

Indicates that the bandwidth use is quite different among the ports in an aggregation pool. This can often be corrected by reconfiguration, for example, by moving a service profile from one chassis to another.

resource component

Interfaces, services and processes, and installed software in the IT environment.

service definition

A service definition contains the information that Control Center needs to start and manage a service, in JavaScript Object Notation (JSON) format.

service profile

A service profile is a software definition of a server and its LAN and SAN network connectivity, in other words, a service profile defines a single server and its storage and networking characteristics.

service template

A service template contains one or more service definitions, in JavaScript Object Notation (JSON) format.

threshold

Defines a value beyond which a data point should not go. When a threshold is reached, the system generates an event. Typically, threshold events use the `/Capacity` event class.

trigger

Determines how and when notifications are sent. Specifies a rule comprising a series of one or more conditions.