

Service Dynamics

Analytics and Optimization
Installation and Administration

Zenoss, Inc.

www.zenoss.com

Zenoss Service Dynamics Analytics and Optimization Installation and Administration

Copyright © 2013 Zenoss, Inc. All rights reserved. Redistribution or duplication of any portion of this document is prohibited without the express written consent of Zenoss, Inc.

Copyright © 2010 Jaspersoft Corporation. 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.

Flash is a registered trademark of Adobe Systems Incorporated.

Oracle, the Oracle logo, and Java are registered trademarks of the Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

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

Part Number: 22-082013-4.2-v05

1. About Analytics	1
1.1. Introduction	1
1.2. Analytics Architecture	1
1.3. In this Guide	2
2. Installing and Configuring Analytics	4
2.1. Installation Considerations	4
2.2. Required Software and Conditions	4
2.2.1. Compatibility and Support	5
2.2.2. Recommendations	5
2.2.3. Open Network Ports	6
2.3. Installing the Oracle Java JDK	6
2.3.1. Setting the JAVA_HOME Environment Variable	6
2.4. Installing the Zenoss DataStore on the Analytics Server	7
2.5. Installing the Analytics Server	8
2.6. Installing the ZenETL ZenPack on the Resource Manager	9
2.7. Configuring Analytics	10
2.7.1. Configure the Connection to the Analytics Server	10
2.7.2. Configure zeneventserver Settings	11
2.7.3. Configure Database Credentials	11
2.7.4. Export Data Point Data	11
2.7.5. Configure Email Services on the Analytics Server	13
2.8. Removing an Analytics Installation	13
3. Working with Analytics	14
3.1. Starting and Stopping the Analytics Server	14
3.2. Logging in to Analytics	14
3.2.1. Logging in Through the Resource Manager Interface	14
3.2.2. Logging in Directly to Analytics	15
3.3. Creating Batches	16
3.3.1. Configuring Repeat Intervals	18
3.3.2. Setting a Delay on Performance Batches	19
3.4. Searching the Analytics Repository	19
3.4.1. Searching by Name	21
3.4.2. Navigating the Folder Tree	21
3.4.3. Refining with Filters	21
3.4.4. Changing the Search Expression	23
3.4.5. Interacting with the List of Resources	23
4. Analytics Dashboards	25
4.1. About Dashboards	25
4.2. Working with Dashboards	25
4.2.1. Viewing a Dashboard	25
4.2.2. Creating a Dashboard	26
4.2.2.1. Overview	26
4.2.2.2. Creating a Simple Dashboard	27
4.2.2.3. Adding Controls to a Dashboard	28
4.2.2.4. Adding a Custom URL to a Dashboard	28
4.2.2.5. Refining Dashboard Layout	29
4.2.2.6. Dashboard Design Tips	29
4.2.2.7. Screen Sizes	30
4.2.3. Editing a Dashboard	30
4.2.4. Designing Reports for Dashboards	31
5. Running Reports	32
5.1. Running a Simple Report	32
5.2. Running a Flash Chart	33
5.3. Changing Report Criteria	34

Zenoss Service Dynamics Analytics and Optimization Installation and Administration

5.4. Scheduling Reports	35
5.4.1. Viewing Scheduled Jobs	35
5.4.2. Scheduling a Report	36
5.4.3. Changing Job Schedules	40
5.4.4. Stopping a Job from Running	40
5.4.5. Specifying Job Recurrence	40
5.4.5.1. Simple Recurrence	41
5.4.5.2. Calendar Recurrence	41
5.4.6. Running a Job in the Background	41
6. Analytics Report Catalog	43
6.1. Availability Reports	43
6.2. Inventory and Audit Reports	43
6.3. Performance Reports	43
6.4. Top N Reports	43
6.5. Virtual Infrastructure Reports	44
6.6. Dashboard	44
6.7. Resource Management	44
7. Understanding the Reporting Database	46
7.1. Overview	46
7.2. Schema	46
7.3. Table Relationships	47
7.4. Catalog Reports	47
7.4.1. Availability	48
7.4.2. Inventory and Audit	49
7.4.3. Performance	51
7.4.4. Server	53
7.4.5. Virtual Infrastructure	56
7.4.6. Resource Manager Management	62
8. Working with Ad Hoc Reports	63
8.1. About the Ad Hoc Editor	63
8.2. Working with the Ad Hoc Editor	63
8.2.1. Topics and Domains	63
8.2.2. Ad Hoc Editor Tool Bar	64
8.2.3. Tables, Charts and Crosstabs	65
8.3. Designing an Ad Hoc Report	65
8.3.1. Selecting Report Content	66
8.3.2. Formatting Table Appearance	67
8.3.2.1. Edit Report Style	68
8.3.2.2. Change Group and Column Labels	68
8.3.2.3. Resizing Columns	69
8.3.2.4. Adding Spacers	69
8.3.2.5. Editing Data Format	69
8.3.2.6. Displaying Detail View	70
8.3.3. Viewing and Running a Report in the Editor	70
8.4. Ad Hoc Editor Features	70
8.4.1. Creating Custom Fields	71
8.4.2. Using Input Controls and Filters	73
8.4.2.1. Using Input Controls	73
8.4.2.2. Using Filters	74
8.4.2.3. Input Controls and Filters Availability	75
8.4.3. Switching Groups	76
8.4.4. Sorting Tables	76
8.4.5. Working with Charts	77
8.4.5.1. Designing Charts	77

Zenoss Service Dynamics An-
alytics and Optimization In-
stallation and Administration

8.4.6. Working with Crosstabs	80
8.5. Creating a Report from a Domain	81
8.5.1. Choosing Ad Hoc Data from a Domain	82
8.5.2. Using the Choose Ad Hoc Data Wizard	82
8.5.2.1. Data Page	83
8.5.2.2. Filters Page	83
8.5.2.3. Display Page	84
8.5.2.4. Topic Page	84
8.5.3. Creating Topics from Domains	84
8.5.3.1. Access Permissions in Domain Topics	85
8.5.3.2. Saving Domain Settings as a Topic	85
8.5.3.3. Editing a Domain Topic	86
8.6. Configuring System-Level Ad Hoc Options	87
8.6.1. Limits on Queries	87
8.6.2. Data Policies	87
9. Creating Domains	88
9.1. Introduction to Domains	88
9.1.1. Domain Use Cases	89
9.1.2. Terminology	89
9.1.3. Components of a Domain	89
9.1.4. Overview of Creating a Domain	90
9.2. Using the Add New Domain Dialog	90
9.2.1. Data and Design Page	90
9.2.1.1. Data Source Area	91
9.2.1.2. Properties Area	91
9.2.1.3. Design Area	91
9.2.2. Resources Page	92
9.2.3. Domain Validation	92
9.3. Using the Domain Designer	92
9.3.1. Designer Tool Bar	93
9.3.2. Tables Tab	93
9.3.3. Manage Data Source	94
9.3.4. Derived Tables Tab	95
9.3.5. Joins Tab	95
9.3.6. Calculated Fields Tab	96
9.3.7. Filters Tab	97
9.3.8. Display Tab	98
9.3.9. Properties Table	99
9.3.10. Domain Design Validation	101
9.4. Editing a Domain	101
9.4.1. Maintaining Referential Integrity	102
10. Administering the Repository	103
10.1. Managing Folders and Resources	103
10.1.1. Creating a Folder	103
10.1.2. Adding Resources to the Repository	103
10.1.3. Renaming a Folder or Resource	104
10.1.4. Viewing a Report or Dashboard	104
10.1.5. Modifying an Ad Hoc Report or Dashboard	105
10.1.6. Editing a Resource	105
10.1.7. Copying Folders or Resources	105
10.1.8. Deleting Folders or Resources	106
10.2. Access Control	107
10.2.1. Authentication Overview	107
10.2.2. Authorization Overview	108

Zenoss Service Dynamics Analytics and Optimization Installation and Administration

10.2.3. Permissions	108
10.2.3.1. Inheriting Permissions	109
10.2.3.2. Cumulative Permissions	109
10.2.3.3. Administrator Permissions	109
10.2.3.4. Default User Permissions	110
10.2.4. Assigning Permissions	110
10.2.5. Testing User Permissions	112
10.3. Referencing Resources in the Repository	112
11. Troubleshooting	114
11.1. Resolving Issues	114
A. Using Secure Copy to Transfer Files to Analytics	115
A.1. Introduction	115
A.2. Procedure	115

Chapter 1. About Analytics

1.1. Introduction

Zenoss Service Dynamics Analytics and Optimization ("Analytics") provides enterprise reporting for Zenoss Service Dynamics Resource Management ("Resource Manager") instances, enabling improved business intelligence and knowledge management. With Analytics, you have access to:

- Catalog reports that provide insight into your entire infrastructure
- Extensible dashboards that provide operational and executive views
- A robust, ad-hoc reporting feature to facilitate self-management
- A built-in report scheduler
- A wide range of report export formats, including PDF, CSV, Excel, Word, HTML, and JPG.

Analytics reports and features are accessed through an integrated interface. Reports are distributed via email.

1.2. Analytics Architecture

Analytics combines extract, transform, and load (ETL) procedures with a data warehouse and reporting tools to process report data.

Resource Manager uses three data sources that are key for reporting:

- Zope object database (ZODB)
- MySQL event database (EventDB)
- Round robin databases (RRD), comprising a distributed collection of performance metrics

One component of Analytics, the ZenETL ZenPack, extracts data from the Resource Manager data sources to build a data warehouse used for enterprise reporting. The data warehouse schema is driven by meta data provided by ZenETL, providing a flexible platform for report creation.

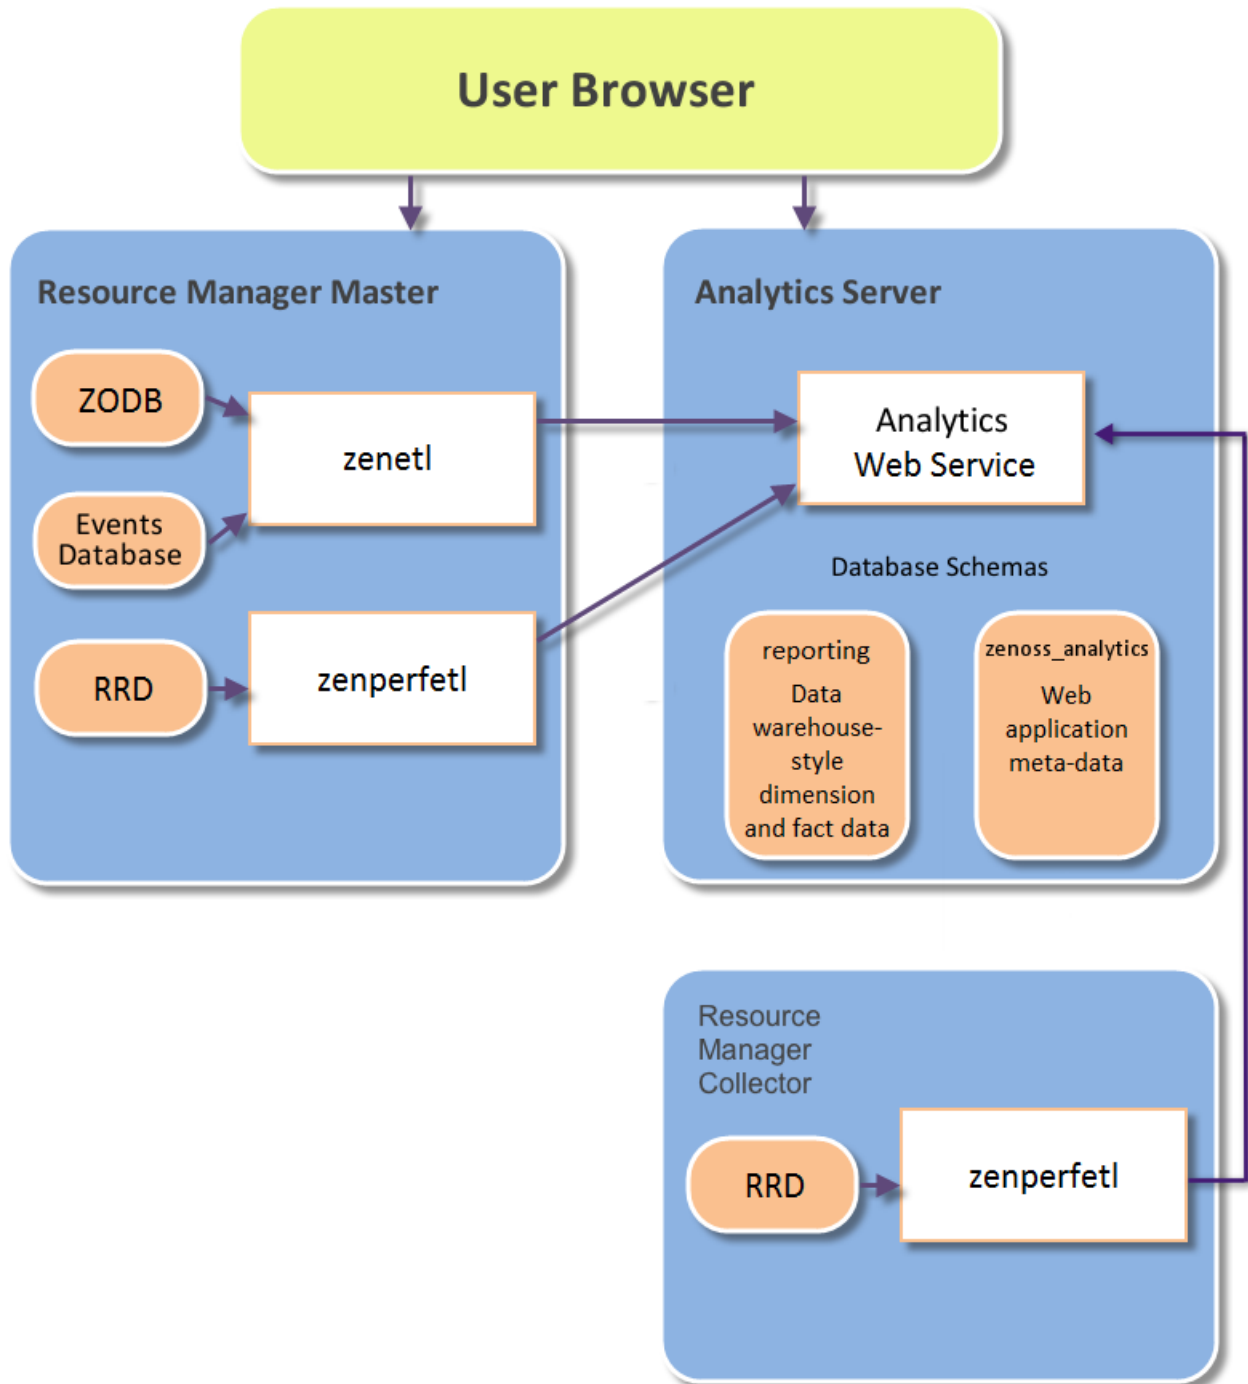
ETL is not a continuous process; it extracts and processes data periodically (according to user-defined parameters) in *batches*. Model, event, and performance data are independently extracted.

On the Resource Manager server, two daemons run to export data from the Resource Manager server to the Analytics server:

- **zenetl** - Extracts, transforms, and loads ZODB model information and events.
- **zenperfetl** - Extracts, transforms, and loads RRD performance information. This daemon also runs on each collector. Typically, this is on the server and at each remote collector.

All RRD Data Points are available for performance data extraction. You can select which data points are extracted by assigning them an alias.

Figure 1.1. Analytics Components



1.3. In this Guide

Read the following chapters for detailed information about installing, configuring and using Analytics:

- *Installing and Configuring Analytics*

- *Working with Analytics*
- *Analytics Dashboards*
- *Running Reports*
- *Analytics Report Catalog*
- *Understanding the Reporting Database*
- *Working with Ad Hoc Reports*
- *Creating Domains*
- *Troubleshooting*

Chapter 2. Installing and Configuring Analytics

The following sections provide information and procedures to help you install and set up Analytics:

- Installation Considerations
- Prerequisite Tasks and Conditions
- Installing the Oracle Java JDK
- Installing the Zenoss DataStore on the Analytics Server
- Installing the Analytics Server
- Installing the ZenETL ZenPack on Resource Manager
- Configuring Analytics
- Removing an Analytics Installation

2.1. Installation Considerations

The Analytics system creates two databases:

- **zenoss_analytics** - Stores report design and execution information.
- **reporting** - Warehouses all data retrieved from Resource Manager. Also referred to as the "etl database."

The reporting database can become large, and load can be significant. For large-scale installations, Zenoss recommends you install the zenoss_analytics database and the Analytics application on one machine, and the reporting database on a separate machine.

2.2. Required Software and Conditions

- These upgrade paths are supported for Analytics:

Table 2.1. Upgrade Paths

If your current Analytics version is:	You can upgrade directly to this version:
4.1.1	4.2.4
4.1.2	4.2.4
4.2.2	4.2.4
4.2.3.1	4.2.4
4.2.3.1.1	4.2.4
4.2.3.1.2	4.2.4

- To run Analytics, you must install:
 - Oracle Java SE 6 Java Development Kit (JDK), this is different than the Java Runtime Environment (JRE) that was installed as part of the Resource Manager installation.

Note

Do not use Java SE 7. It is not compatible with this version of Analytics.

- Zenoss DataStore on the Analytics Server
- The following ZenPacks (included with Resource Manager 4.2.x and Zenoss Enterprise 4.1.1):
 - Distributed Collector
 - vCloud
 - VMware
- Analytics requires that UID 1337 be available. You can verify its availability with the command:

```
getent passwd 1337
```
- Analytics does not support Hypertext Transfer Protocol Secure (HTTPS) communications.
- The Analytics server and any associated Resource Manager instances must be running Network Time Protocol (NTP) so that their clocks remain synchronized. (For more information about NTP, visit <http://www.ntp.org/>.)

2.2.1. Compatibility and Support

- Analytics 4.2.4 is compatible with Service Dynamics 4.1.1 and later versions.
- You must use matched versions of Analytics and the ZenETL ZenPack. For example, you cannot use ZenETL 4.1.2 with Analytics 4.2.4, or ZenETL 4.2.4 with Service Dynamics 4.1.1.
- You must be using a 64-bit system running RHEL 5/CentOS 5 or RHEL 6/CentOS 6.
- The `zenperfetl` daemon that runs on remote collectors makes Web services calls to the Analytics server, and thus must be able to access the server's HTTP address. SSL is not supported for this communication.
- The system is not compatible with Security-Enhanced Linux (SELinux) in enforcing mode. You must disable SELinux to run Analytics.

To disable enforcing mode in SELinux:

1. Edit the `/etc/selinux/config` file.
2. Set the following line:

```
SELINUX=disabled
```

Note

You also can disable enforcing mode temporarily (avoiding the need to reboot) with the command:

```
echo 0 > /selinux/enforce
```

For more information about SELinux, browse to <http://en.wikipedia.org/wiki/SELinux>, or to the SELinux home page at <http://www.nsa.gov/research/selinux/index.shtml>.

2.2.2. Recommendations

- You should install Analytics and its prerequisites on a server separate from the Resource Manager server.

- Analytics is a data warehouse application. Like other database applications, it benefits from fast disc subsystems, fast processors, and abundant RAM.

Zenoss recommends that you meet these minimum system requirements when running Analytics:

- Dual core processor
- 8GB RAM
- 500GB storage

2.2.3. Open Network Ports

The Analytics system comprises multiple components. Communication needs for each of these components differs depending on how you have installed your system (on a single box or multiple boxes, or with our without distributed collectors).

The following table lists Analytics pieces that typically must communicate, and the default network ports that should be open for communication.

From	To	Default Port Numbers
Analytics	Analytics database server	13306
Analytics	Resource Manager	8080
Resource Manager	Analytics	7070
Analytics Web Users	Analytics	7070
zenperfet1 (on each remote collector)	Analytics application server	7070

2.3. Installing the Oracle Java JDK

Analytics requires the Oracle Java 1.6 JDK. It cannot run when just the JRE is installed. To determine if you already have the JDK installed, run the following command:

```
rpm -q jdk
```

If you do not have the JDK installed, perform the following steps:

1. Download the self-installing RPM of Oracle Java SE Development Kit 6u31 from the [Java SE 6 Downloads](#) page. The file to download is `jdk-6u31-linux-x64-rpm.bin`.

Note

Before the file can be downloaded, you must accept the license agreement.

2. Change mode:

```
chmod +x ./jdk-6u31-linux-x64-rpm.bin
```

3. Install Oracle Java JDK:

```
./jdk-6u31-linux-x64-rpm.bin
```

2.3.1. Setting the JAVA_HOME Environment Variable

After installing the Oracle Java JDK, you must set `$JAVA_HOME` for the root user:

1. Locate the path to your JDK directory. Typically, this is `/usr/java/jdkVersionNumber`.
2. Add the following lines to the `/root/.bash_profile` file:

```
export JAVA_HOME=/usr/java/jdkVersionNumber
export PATH=$JAVA_HOME/bin:$PATH
```

3. Reboot your system.

Note

You can test the new path with the command:

```
#echo $JAVA_HOME
```

2.4. Installing the Zenoss DataStore on the Analytics Server

Follow these steps to install the Zenoss DataStore on the Analytics server.

Note

Before installing the Zenoss DataStore on systems running Red Hat Enterprise Linux (RHEL), install the Perl Database Interface Module (DBI) with the following command.

```
yum install perl-DBI
```

Install the Zenoss DataStore:

1. Browse to this URL, and then locate the Zenoss DataStore RPM files:

<https://support.zenoss.com>

Note

Contact your Zenoss representative for site login credentials.

2. Install the Zenoss DataStore. As the root user, enter the following command:

```
rpm -ivh zends-5.5.31-1.Version.RHEL.x86_64.rpm
```

where *Version* is the build version number and *RHEL* is the RHEL or CentOS version (e.g., el5 or el6)

3. Configure the Zenoss DataStore by adding the following lines to the end of the `/opt/zends/etc/zends.cnf` file:

```
[mysqld]
open_files_limit=200000
event_scheduler=ON
```

Note

The event scheduler property is needed on the repository (warehouse) database, but not on the `zenoss_analytics` database.

4. Run the following commands:

```
service zends start
```

```
chkconfig zends on
```

5. Change to the zenoss user:

```
su - zenoss
```

6. Enter the following command:

```
zends -u root
```

7. Enable remote access to the Zenoss DataStore. In the prompt that appears, enter this series of commands:

```
grant all on *.* to 'root'@'%' with grant option;
grant all on *.* to 'reporting'@'%' identified by 'reporting';
flush privileges;
```

Note

The 'reporting' privileges are based on the values set in the `/etc/zenoss_analytics/zenoss_analytics.conf` file on the Analytics server.

2.5. Installing the Analytics Server

To install the Analytics server:

1. As the root user, enter the following command:

```
yum --nogpgcheck localinstall zenoss_analytics-4.2.4-Version.noarch.rpm
```

2. Restart the Zenoss DataStore. (This clears the resource limit values previously set in the `/etc/security/limits.conf` file.)

```
service zends restart
```

3. Modify the database connection setting in the `/etc/zenoss_analytics/zenoss_analytics.conf` file to point to the desired database server or servers.

```
etl.jdbc.hostname=<remote database server>
```

Note

Be sure to modify the admin account information to match the account or accounts you created when setting up the database.

4. Change to the zenoss user:

```
su - zenoss
```

5. Install the database schema with the command:

```
/opt/zenoss_analytics/bin/upgrade_db.py
```

6. Change to the root user:

```
su - root
```

7. Enable chkconfig for the Analytics service:

```
chkconfig zenoss_analytics on
```

8. Start the Analytics service:

```
service zenoss_analytics start
```

2.6. Installing the ZenETL ZenPack on the Resource Manager

After installing the Analytics server, install the ZenETL ZenPack on the Resource Manager master:


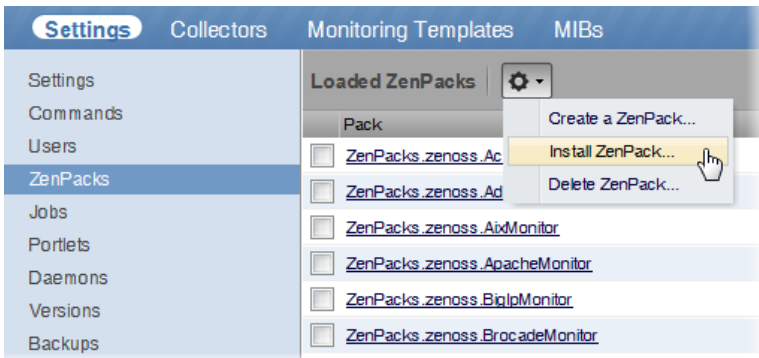

1. From the Resource Manager interface navigation bar, select Advanced, and then select Settings.
2. In the left panel, select ZenPacks.
3. From  (Action menu), select Install ZenPack.

Figure 2.1. Install ZenPack Selection



The Install ZenPack dialog appears.

Figure 2.2. Install ZenPack Dialog



4. Browse to and install the current ZenETL ZenPack (version 4.2 with python 2.7), and then click **OK**.

The file is uploaded to the Resource Manager server and installed.

5. As the zenoss user, restart the system:

```
zenoss restart
```

Note

On distributed collector installations, the updated configuration must be pushed to each hub, and then to each collector.

Optionally, if you are using `daemons.txt` on the master, you should add entries for `zenetl` and `zenperfetl` to the file. If you are using `collectorddaemons.txt` on the collectors, you should add an entry for `zenperfetl` to the file.

When you restart the Resource Manager master, two additional daemons (`zenetl` and `zenperfetl`) are started. The `zenperfetl` daemon is started on each remote collector.

2.7. Configuring Analytics

To configure Analytics, you must:

- Configure the connection to the Analytics server
- Configure `zeneventserver` settings
- (Optionally) Configure MySQL credentials
- (Optionally) Configure additional Resource Manager data points to send to Analytics
- Configure email services
- Open network ports

2.7.1. Configure the Connection to the Analytics Server

Follow these steps to configure Resource Manager to connect to the Analytics server:

1. From the Resource Manager interface navigation bar, select Reports > Configuration.

The Analytics configuration page appears.

2. In the Internal URLs area, enter the URLs that Resource Manager and Analytics will use to communicate directly:

- **Analytics** - Enter the fully qualified domain name (including port number) for the Analytics server. The default port number is 7070.
- **Resource Manager** - Enter the fully qualified domain name (including port number) for the Resource Manager server. (By default, this information is pre-populated.)

3. In the External URLs area, enter the links that will be displayed to users in each of the applications:

- **Analytics** - Enter the complete URL, including port number, to Analytics. The default port number is 7070.
- **Resource Manager** - Enter the complete URL, including port number, to Resource Manager. (By default this information is pre-populated.)

4. Click **Submit**.

The configuration page redisplay with the currently scheduled report jobs.

Figure 2.3. Add Batches

Batch ID	Extractor	State	Begin	End	Attempts
6	EVENTS	COMPLETED	2013-01-14T07:40:11	2013-01-14T15:40:11	1
4	remote-collector	UNSTARTED	2013-01-14T07:40:11	2013-01-14T15:40:11	0
13	MODEL	COMPLETED	2013-01-14T07:40:11	2013-01-14T15:40:11	1
5	localhost	UNSTARTED	2013-01-14T07:40:11	2013-01-14T15:40:11	0

Note

If this is a new installation, then jobs will not initially appear.

2.7.2. Configure zeneventserver Settings

Resource Manager (Version 4.1 and higher) is configured to store only one day of event occurrences. If event extraction falls behind (meaning event batches are older than 24 hours), then event extraction will not properly extract event times for those batches.

You can increase the number of days of event history that can be extracted, from the Resource Manager interface:

1. Select **Advanced**, and then select **Events** from the left panel.

The Event Configuration page appears.

2. Increase the value of **Event Time Purge Interval** (in days), up to 7 days.
3. Click **Save**.

Note

If you increase the **Event Time Purge Interval** value, then you must also increase the number of file descriptors for the Zenoss Datastore for your Resource Manager instance.

2.7.3. Configure Database Credentials

When Analytics is installed it configures a JDBC data source for the reporting database.

To change database settings in Analytics:

1. On the Analytics Getting Started page, click **View Reports**.
2. In the **Refine** area, change the **Reports** item to "Data sources."
3. In the **JDBC Data Sources** area, right-click the **Resource Manager Reporting** data source, and then select **Edit** from the list of options.
4. After editing, click **Test connection** to verify the changes are valid.
5. Click **Save**.

2.7.4. Export Data Point Data

This section provides information about how Analytics and Resource Manager share performance data, and how to manually modify data points.

Resource Manager uses RRDTOol to collect performance data. Data points are collected through the use of templates. (For more information about templates, see *Zenoss Service Dynamics Resource Management Administration*.)

Resource Manager data points represent raw, collected metrics. For example, view the `/Server/Linux` monitoring template. The `ssCpuIdle` data point is collected by the `ssCpuIdle` data source. `ssCpuIdle` represents the percentage of idle time for all CPUs on this box.

Edit the data point, and you can examine its defined aliases. The `cpu__pct` alias represents the percentage of CPU used by the box. The alias value is calculated by using the data point and an RPN expression. Aliases are used to normalize raw data point values across data point types.

Further, view another monitoring template, `/Server/Windows`. Open the `ProcessorTotalProcessorTime` data point in the `ProcessorTotalProcessorTime` data source. You will see that an alias "`cpu__pct`" is defined with no RPN formula. In this case, the raw data point already represents the quantity we want ("Percentage of Consumed CPU Time"). The alias, `cpu__pct`, allows users to compare two different data sources.

The same principle applies to other data points. Some devices give network bandwidth by using kilobytes; others use megabits. Aliases provide a method to normalize these data points to a common denominator.

You can provide a formula to be used during extraction of data. This allows you to convert data points with different units to the same units during extraction. For example, if you are monitoring disk space on a Unix system (bytes) and a Windows system (bits), you can use a formula to convert both to kilobytes and have them use the same alias.

Analytics only extracts "aliases" from Resource Manager. A raw data point is never exported. To export a performance data point to Analytics, examine the data source/data point responsible for the raw value, and ensure that it has an alias that maps to an appropriate alias name whose value is properly normalized. Alias names are mapped to the `meta_metric` table in Analytics. (In previous versions, an alias would map to a `fact_[alias]` table; this mapping is now deprecated.) All performance data is now logically mapped to `merge_raw_v1` and `merge_agg_v1`. These tables have a column named `metric_key`. The table `meta_metric` provides the lookup for a specific performance metric.

```
SELECT f.*
FROM merge_raw_v1 AS f
INNER JOIN meta_metric AS mm ON f.metric_key = mm.metric_key
WHERE mm.metric_name = 'cpu__pct'
limit 5;
```

This query will provide all the performance data for the `cpu__pct` metric. The underlying data is stored in `raw_*` and `agg_*` tables. These tables should not be used directly. Instead, the two `merge_*` tables provide a unified table with which to query all performance metrics.

With respect to event data points, all non-cleared events are brought over to Analytics and are stored in the `fact_event` and `fact_event_time` tables.

Manually Modifying Data Points on the Resource Manager

To export data point data from the Resource Manager (for example, RRD performance metrics), you can manually modify each data point.

1. From the Resource Manager interface navigation bar, select `Advanced > Monitoring Templates` to navigate to the RRD template.
2. In the Data Sources area, expand the data source to see the associated data points.
3. Double-click the data point.

The Edit Data Point dialog appears.

4. Add an alias for the data point.

Tip

You can use the same name as the data point. No formula is required.

5. Click **Add**, and then click **Save**.

2.7.5. Configure Email Services on the Analytics Server

If you want to send the output of scheduled reports as email attachments, you must configure email services. To do this:

1. On the Analytics server, edit the `/opt/zenoss_analytics/webapps/zenoss-analytics/WEB-INF/js.quartz.properties` file.

2. Change the "host" and "from" values in the file.

```
report.scheduler.mail.sender.host=localhost
report.scheduler.mail.sender.username=
report.scheduler.mail.sender.password=
report.scheduler.mail.sender.from=zenoss@localhost
report.scheduler.mail.sender.protocol=smtp
report.scheduler.mail.sender.port=25
```

3. Save the file.

4. As the root user, restart Analytics for the changes to take effect:

```
service zenoss_analytics restart
```

2.8. Removing an Analytics Installation

To remove an installation of Analytics, enter the following commands as the root user:

```
service zenoss_analytics stop
rpm -e zenoss_analytics
```

This command removes program files, but does not remove configuration, log files, data files, or databases.

To remove configuration files, enter:

```
rm -rf /etc/zenoss_analytics
```

To remove data and log files, enter:

```
rm -rf /opt/zenoss_analytics
```

To remove databases:

1. Change to the zenoss user:

```
su - zenoss
```

2. Enter the following command:

```
mysql -u root -e 'drop database reporting; drop database zenoss_analytics';
```

Chapter 3. Working with Analytics

Read the following sections for more information about:

- Starting and stopping the Analytics server
- Logging in to Analytics
- Creating batches
- Searching the Analytics repository

3.1. Starting and Stopping the Analytics Server

Use one of the following commands if you need to stop, start, or restart the Analytics server (such as when troubleshooting problems). While the server is stopped, no data is extracted.

To start Analytics, use the command:

```
service zenoss_analytics start
```

To stop and then start Analytics, use the command:

```
service zenoss_analytics restart
```

Note

If you restart Analytics, there can be a delay of up to 60 seconds as Analytics and Resource Manager synchronize. During this time, single sign-on may not work correctly.

To stop Analytics, use the command:

```
service zenoss_analytics stop
```

Note

While the Analytics server is stopped, no data is extracted.

3.2. Logging in to Analytics

To access Analytics, you can:

- Log in to the Resource Manager interface, using your Resource Manager user name and password
- Log in directly to Analytics, using your Resource Manager user name and password or as an Analytics superuser

3.2.1. Logging in Through the Resource Manager Interface

Access to Analytics through the Resource Manager interface requires Manager, ZenManager, or ReportingUser privileges, assigned through roles. (Manager and ZenManager provide Analytics administrative privileges; ReportingUser provides user-level privileges.) You cannot log in to Analytics solely with Resource Manager administrator privileges.

1. Log in to the Resource Manager interface with your Resource Manager user name and password.
2. Select Reports > Advanced.

The Getting Started page appears.

Figure 3.1. Getting Started

Getting Started



View your
reports



Create a
report



Analyze
your results



Manage the
app

From here, you can choose shortcuts to:

- View your reports
- Create a report
- Analyze report results
- Change Analytics settings

Alternatively, make a selection from one of the Analytics menu options:

- **View** - View the Analytics repository, reports, analyses, or messages.
- **Manage** - Organizations, users, or roles (accessible by administrators)
- **Create** - Create ad hoc reports or dashboards.

Note

If you log out of Analytics after logging in through the Resource Manager interface, you must log out of Resource Manager to log in again through Resource Manager. Further, if the Resource Manager server is restarted, you must log back in to Analytics.

3.2.2. Logging in Directly to Analytics

You can log in directly to Analytics by using your Resource Manager user name and password (with appropriate Resource Manager privileges), or as an Analytics superuser.

1. Browse to the URL of your Analytics server.

The Analytics login page appears.

2. Log in:

- With your Resource Manager user name and password:
 - a. From the Resource Manager Instance list of options, select the Resource Manager instance for which you will provide a user name and password.
 - b. Enter your Resource Manager user name and password in the User ID and Password fields.

Figure 3.2. Analytics Login - Resource Manager Credentials

Login

Zenoss Instance

Organization

User ID

Password

[Show locale & timezone](#)

- As an Analytics superuser:
 - a. From the Resource Manager Instance list of options, select Internal Authentication.
 - b. If the user is located in an Analytics organization other than Resource Manager, enter the organization name. (Typically, you would not enter a value in this field.)
 - c. Enter your Analytics User ID and password in the User ID and Password fields.

Figure 3.3. Analytics Login - Analytics Superuser

Login

Zenoss Instance

User ID

Password

[Show locale & timezone](#)

3. Click **Login**.

3.3. Creating Batches

You create batches to populate the reporting database. Every eight hours, the ETL processes a batch and extracts information from the previous eight-hour window. Three main types of batches match the primary data types (Performance, Model, and Event).

Existing batches and associated status information appear in the batch list. To access the batch list, select Reports > Configuration from the Resource Manager interface.

The batch list appears.

Figure 3.4. Batch List

Batch ID	Extractor	State	Begin	End	Attempts
1	MODEL	COMPLETED	2011-04-26T12:00:00	2011-05-16T12:00:00	1
2	EVENTS	COMPLETED	2011-04-26T12:00:00	2011-05-16T12:00:00	1
3	localhost	COMPLETED	2011-04-26T12:00:00	2011-05-16T12:00:00	1

The list shows each batch and its status (UNSTARTED, EXTRACTING, STAGING, FAILED, or COMPLETED).

To add a batch:

1.



From the batch list, click **+ Add Batches** (Add Batches).

The Add Batch dialog appears.

Figure 3.5. Add Batch

2. Make selections in the dialog:

- **Extractor** - Select MODEL, EVENTS, or localhost. An additional performance extractor is listed for each remote collector.
- **Begin Date** - Optionally adjust the begin date and time.
- **End Date** - Optionally adjust the end date and time. By default, the value is set to the current Analytics server date and time.

Note

Creating a batch that finishes in the future will leave a batch with a status of UNSTARTED. It will not start until that future date is reached.

3. Click **Submit**. The batch is added to the batch list.

Note

RRD stores data samples at various different resolutions. It is configured to store a fixed number of data samples at any resolution. Resource Manager, by default, creates RRD files to store 288 samples at five minute (300 second) intervals, called a *step*. This means that Resource Manager retains, at most, one day of five-minute samples for a data point. To ensure that disk usage does not grow without bound, RRD consolidates the five-minute samples in less granular periods (such as 15 minutes, two hours, one day, or one week).

When you schedule batches, remember that Resource Manager will return data at the finest granularity that exists over the entire period. This means that, for example, if your query is for the last 26 hours, you will not receive 300 second data samples but approximately 1800 second (15 minute) samples. The further back your batch periods reach, the less granular the retrieved data will be.

3.3.1. Configuring Repeat Intervals

When configuring batch intervals (in Resource Manager), you must set an Analytics repeat interval for each Analytics server (rather than for each Resource Manager instance). This means that the same Analytics repeat interval setting will appear across Resource Manager instances.

You can configure repeat intervals separately for each extractor type. By default, all repeat intervals are set to a value of 8 hours.

To edit repeat intervals:

1. Expand the Jobs folder, and then select one of the extractor types.

Figure 3.6. Edit Repeat Interval

Basic	Advanced	Configuration	Page Tips
<ul style="list-style-type: none"> qa-set3-6-64-zen.zenoss.loc <ul style="list-style-type: none"> Jobs <ul style="list-style-type: none"> Performance Jobs Model Jobs Event Jobs 		Name: PerfEtIJob Group: ETL_TRIGGER_GROUP Description: Create an ETL batch for PerfEtIJob Java Class: com.zenoss.reporting.schedule.EtlJob Trigger: PerfEtITrigger Trigger Description: Trigger for ETL job. Next Fire Time: 2013-01-14T23:40:11.870000 Previous Fire Time: 2013-01-14T15:40:11.870000 Start Time: 2013-01-13T23:40:11.870000 Times Triggered: 3 Repeat Interval: 8.0 hours	
		<input type="button" value="Submit"/>	

2. Select a repeat interval (.25 hours to 8 hours), and then click **Submit**.

3.3.2. Setting a Delay on Performance Batches

A configuration option for the `zenperfetl` daemon allows you to delay the execution of a performance batch until it reaches a specified amount of time into the past. This defaults to fifteen minutes, and can be extended if needed to allow the collectors to have sufficient time to gather all desired information and record it in the RRD files before the batch is executed.

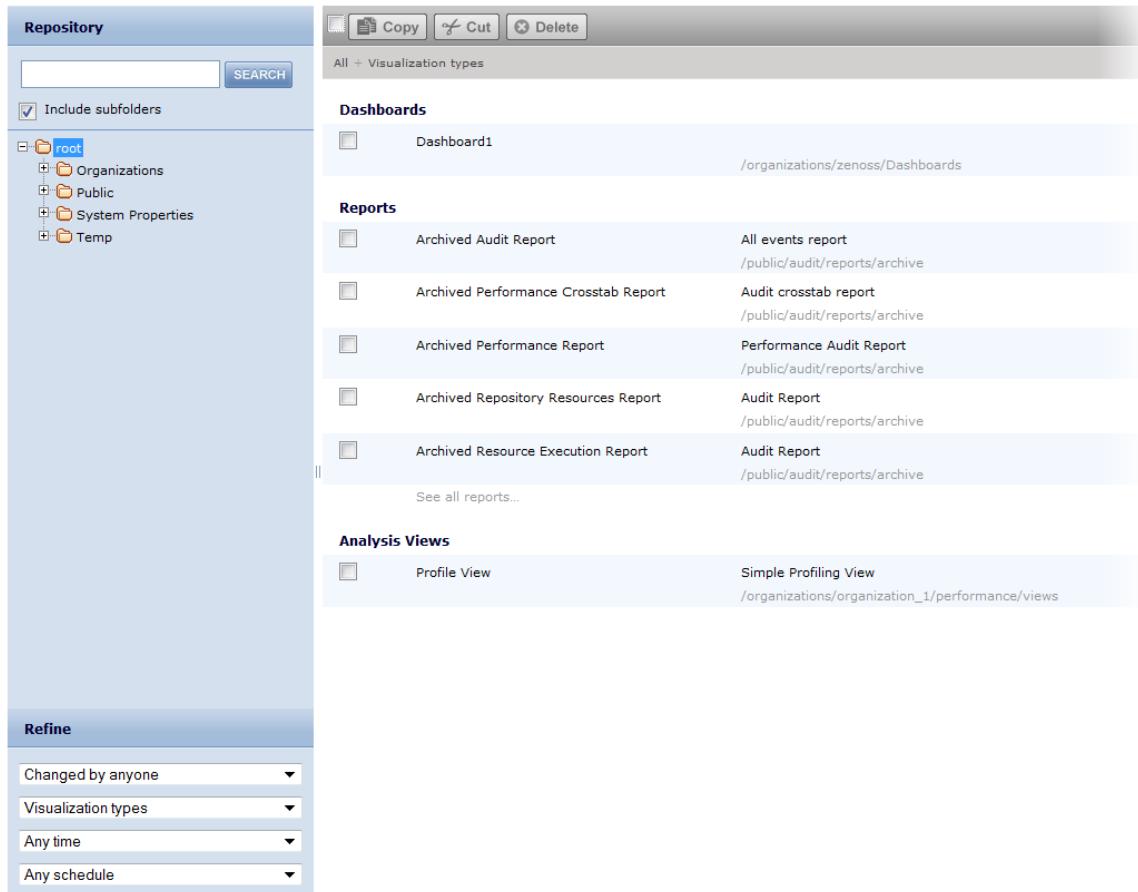
- **perfbatchdelayinminutes** - Run only those performance batches with an end time at least x minutes in the past.

To set a value for this option, select **Advanced > Daemons**, and then click the option to edit the daemon configuration.

3.4. Searching the Analytics Repository

The repository is Analytics' internal storage for all reports, dashboards, analysis views, and related files. The repository is organized as a tree structure of folders containing resources, much like a file system. However, Analytics stores all objects and files in a private database that is otherwise inaccessible, except to administrators.

To access the repository, select **View > Repository**.

Figure 3.7. Analytics Repository View

The Repository page displays search results, as well as lists of reports and analysis views. It features:

- Search controls - Find resources by searching for text in their names or descriptions.
- Folder tree - Use to restrict search results to branches of the tree, or browse individual folder contents.
- Filters for refining results - Provides a quick way to find resources based on other criteria, such as type or access time.
- Bulk operation controls - For selecting and copying, cutting (moving), or deleting multiple resources at one time.
- Search expression and sorting - Interactive status bar that shows search criteria, folders, filters, and sorting criteria.
- List of resources - Resources in the repository that meet the search criteria, grouped by resource type when appropriate. Click or right-click a resource in the list to view it or see available actions.

Repository page controls allow you to refine or enlarge your set of results. Analytics remembers your settings so that your most commonly needed resources remain visible when you return to the page.

The following sections provide more information about:

- Searching by name
- Navigating the folder tree
- Refining with filters

- Changing the search expression
- Interacting with the list of resources

3.4.1. Searching by Name

Use the search field at the top of any Analytics page to begin a search. Search terms are not case sensitive. When you click the search icon, Analytics opens the repository page with the search term and displays the results.

Analytics matches the search term, even word fragments, anywhere it appears in the display name or description string of a results. Analytics does not match folder names or display folders in the results. If you enter multiple words, Analytics searches for resources that match all words, but not necessarily in order. Do not enter quotes around terms or any symbols between terms.

If you enter a search term and click Search on the Repository page, the search uses the current settings on the page. If you enter a search term and click the search icon at the top of any other page, the search uses the default settings:

- Include subfolders
- Start at the top-most visible folder
- Filter for visualization types only
- Sort alphabetically by name

To refine the search, select a folder or filter. To cancel search by name, clear the search term in the field, and update the search results according to the remaining settings.

3.4.2. Navigating the Folder Tree

The folder tree shows the folders in the repository that you have permission to view. Use the icons to expand the tree and view all nested folders. Click a folder name to select it.

There are two ways to navigate the folder tree:

- **Search** - When Include Subfolders is selected, selecting a folder searches the entire branch of the tree, starting with that folder. This mode is most useful in combination with other search settings, allowing you to search by name or by resource type across the entire repository or specific branch. By default, Include Subfolders is selected.
- **Browse** - Clear the Include Subfolders selection to restrict the scope of any search to the resources directly contained in the selected folder. This mode mimics a file system browser, letting you explore folder by folder. This mode works best when you clear the search term and select All types so that the results are exactly the contents of the selected folder.

3.4.3. Refining with Filters

The filters below the folder tree let you refine your search using several other resource attributes. For example, these filters can help you find your most recently viewed reports. You can set each filter independently of the others.

The user filter has the following settings:

Filter Setting	Description
Changed by anyone (default)	Does not take into account user access to the resource.

Filter Setting	Description
Changed by me	Selects only resources that have been last modified by the currently logged in user.
Viewed by me	Selects only resources that have been run and viewed by the currently logged in user. This filter applies not only to visualization types, but also to resources that are included in reports (such as images).

The resource type filter has the following settings:

Filter Setting	Description
All types	Does not take into account the resource type; as a result, every type is displayed.
Visualization types (default)	Default filter that includes reports, dashboards, and analysis views. These are the resources that can be viewed by anyone who has read access.
Reports	Displays only reports.
Report outputs	Displays only the output from reports that were scheduled or run in the background. Report output can be any of the supported export types (such as HTML or PDF).
Dashboards	Displays only dashboards.
Analysis views	Displays only analysis views.
Domains	Displays only domains.
Data sources	Displays only data sources.
Other types	Displays all types except reports, report output, dashboards, views, domains, and data sources. This includes input controls, data types, images, and other repository files.

The access time filter has the following settings. All times are relative to your effective time zone:

Filter Setting	Description
Any time (default)	Does not take into account the last modification time of a resource.
Today	Resources viewed or modified since the previous midnight.
Yesterday	Resources viewed or modified during the previous day ending at midnight.
Past week	Resources viewed or modified during the past 7 days, including today.
Past month	Resources viewed or modified during the past 30 days, including today.

The scheduled report filter has the following settings:

Filter Setting	Description
Any schedule (default)	Does not take into account the existence of scheduled jobs.

Filter Setting	Description
Scheduled	Only reports that have scheduled jobs.
Scheduled by me	Only reports that have jobs scheduled by the currently logged in user.
Not scheduled	Only reports that do not have scheduled jobs, and all other resource types.

3.4.4. Changing the Search Expression

Above the list of resources, the search expression shows you all of the criteria that contribute to the search. It contains the following elements, always in this order from left to right:

- The search term, if any, or the word *All*
- The folder selected, if any
- Any and all filter settings

The search expression provides a quick visual summary of the search criteria for the list of results that appear right below it. The search expression is also interactive, allowing you to easily remove some of the filters. It supports the following actions:

- If there is more than one filter, clicking any of them removes all those to its right.
- If there is a folder selected that is nested more than one level, then you can click any parent folder to select it instead.
- You can click the search term or the word *All* to remove any folders to the right.

Note

To remove the search term, click the X icon in the search field. The search expression then displays *All*.

After any of these actions, the search controls displayed on the left are refreshed, and the results are updated.

To the right of the search expression, the sort criteria lets you change the order of the results. Analytics supports the following sorting:

- Click **Name** to sort alphabetically from A-Z. This is the default sort order.
- Click **Modified Date** to sort by the latest modified time and date (most recent at top).

3.4.5. Interacting with the List of Resources

The list of resources is always the set of results determined by the current search criteria. Because search results can be quite large, the list of resources has two modes:

- **Multiple Types Condensed** – When there are more than two types of resources listed, the number of resources of each type is limited to a certain number, by default 5. When there are more resources of that type, there is a link to see all of that type. You can quickly scroll through the condensed results and find the type of resource you want. Click the *See all...* link to display the results of that type.
- **Single Type Expanded** – When there is a single type of resource in the results, either by circumstance, by clicking a *See all...* link, or because a single-type filter is selected, all resources are listed, with a scroll bar if necessary. If the list of results is still too long, enter or refine the search term, or select additional filters.

Note

When you click See all... for a type that has a single-type filter, for example Domains, that filter is automatically applied. The filter display and the search expression refresh to reflect this. Use the search expression to remove the filter if necessary. When no filter exists for that type, for example input controls, the filter display and search expression are not updated. Use the type filter list of options to change the types displayed.

Once the resource or resources you want are displayed, you can interact with them in several ways:

- Click the name of a report or dashboard to run and view it.
- Right-click the name of a resource to access other operations on the context menu, for example Open in Designer... Items appear on the context menu according to your permissions.
- Select the resource name or click anywhere in the row to select it. (Control-click anywhere in the rows to select multiple resources.) You can drag and drop selected items to move them or press Control while you drag and drop items to copy them. When you have multiple resources selected, you also can use the bulk operation buttons above the results list: Copy, Cut (move), or Delete. You must have the appropriate permissions on the selected items and the target folder to move, copy, or delete resources.

Two icons may appear between a report name and its selection box:

- + indicates that the report has saved options for its input controls. Click + to list the saved options below the report.
- The clock icon indicates that the report has been scheduled to run or that it currently is running in the background. Click this icon to view the list of jobs scheduled for the report.

Chapter 4. Analytics Dashboards

4.1. About Dashboards

A dashboard displays several reports in a single, integrated view. A dashboard also can include input controls (that determine the data displayed in one or more frames), other dashboards, and custom frames that point to other content. By combining different types of related content, you can create appealing, data-rich dashboards that quickly convey trends.

4.2. Working with Dashboards

The following sections provide more information about:

- Viewing a dashboard
- Creating a dashboard
- Editing a dashboard
- Designing reports for dashboards

4.2.1. Viewing a Dashboard

You can view an existing dashboard if you have the proper permissions. By default, the repository includes the / Dashboards folder where dashboards can be saved.

To view a dashboard:

1. In the Search field, enter the dashboard name, and then click **Search**.

The search results appear in the repository.

2. Click the dashboard name in the list of results. (Alternatively, right-click the name, and then select View from the context menu.)

Analytics displays the dashboard. When you mouse over a report in the dashboard, controls appear for that report.

To refresh report content, click . Click  to open a report in a new window.

When viewing a dashboard with input controls:

- An input control may appear as a text field, drop-down, check box, multi-select list box, or calendar icon.
- If one of the frames in a dashboard does not refer to an input control, that frame does not update when you change the input control's value. Only reports that refer to the input control reflect the change.
- If the dashboard does not include a **Submit** button, then any changes to the input control's value take effect immediately.

Note

If a dashboard includes a **Print View** button, then click this to display the dashboard without the Analytics header and footer. Depending on your Web browser, this also opens your browser's Print window.

4.2.2. Creating a Dashboard

Any user can create a dashboard, although user permissions in the repository may limit content and save location.

Read the following sections for more information about creating dashboards:

- Dashboard overview
- Creating a simple dashboard
- Adding controls to a dashboard
- Adding a custom URL to a dashboard
- Refining dashboard layout
- Dashboard design tips
- Screen sizes

4.2.2.1. Overview

A dashboard can include:

- Reports in the repository
- Special content, which includes:
 - Custom URL - Any URL-addressable Web content. Analytics dashboards can point to any Web content, so this feature has nearly-limitless possibilities. For example, you might include a frame that points to the logo on your corporate Web site; when that logo changes, your dashboard will automatically update to reflect the branding change.
 - Free Text - Free-form text entry field. Resizing this type of item changes the size of the font in the label. Use free text items to add titles and instructional text to your dashboard.
 - Controls (single or multiple) - If a report on the dashboard refers to input controls or filters, you can add them to the dashboard. When you add a control, Analytics automatically adds a label for it as well. The input control can be mapped to one or more frames. For example, if multiple reports include the same parameter, the corresponding control is automatically mapped to each of those reports when you add the input control to the dashboard. Controls can also be manually mapped to custom URL frames.

Note

Multiple controls are those used by more than one report. Single controls are those that correspond to parameters included in a single report.

- Standard controls, which include:
 - Submit - Applies the values in the dashboard's input controls to the reports that refer to each input control. Analytics refreshes these reports to display the new set of data. If the dashboard does not include a **Submit** button, then changes to input control values are reflected immediately.
 - Reset - Resets the values of the input controls to the last value selected when the dashboard was saved.
 - Print View - Displays the dashboard without buttons or Analytics' header and footer, and (depending on your browser) opens your browser's Print window.

- Text Label - Label for an input control. When you add an input control to the dashboard, Analytics automatically adds a text label for it. Resizing this type of item changes only the size of the label; the font size in the label is fixed.

4.2.2.2. Creating a Simple Dashboard

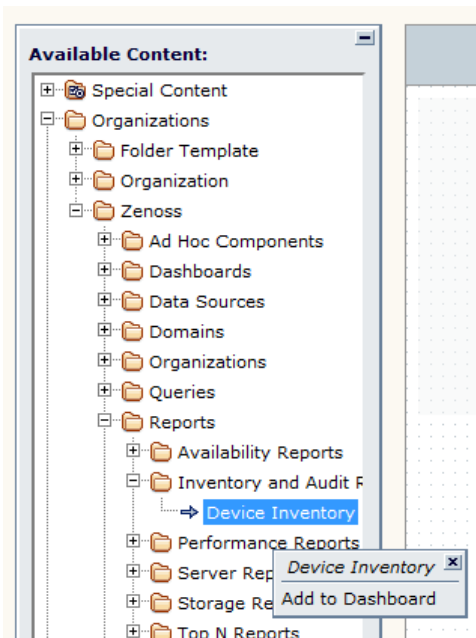
To create a simple dashboard:

1. From the Analytics interface, select Create > Dashboard.

The dashboard designer appears.

2. In the Available Content area, navigate to and select the item you want to add to the dashboard.
3. Select Add to Dashboard.

Figure 4.1. Add Content to Dashboard



The content is added to the dashboard.

4. Hover over the bottom of the right edge of the frame containing the content. When your cursor changes to a resizing icon, click and drag the edge of the frame to resize it as desired.

Note

Analytics allows you to save a dashboard even if its contents "hang over" the edge of the dashboard area. This is because the dashboard area is meant as a guide. If the content fits in the dashboard area, then you can be sure that any viewer will be able to view the dashboard without using scroll bars. For more information, see the section titled "Screen Sizes."

5. Click **Save** to save the dashboard.

Analytics prompts for the name and save location. By default, the dashboard is saved in the /Dashboards folder with the name "New Dashboard." Enter a name and location, and then click **Save**.

Analytics saves the dashboard to the repository.

4.2.2.3. Adding Controls to a Dashboard

Reports refer to input controls that determine the data they return. You can add those controls to a dashboard:

1. Open a dashboard in the dashboard designer.
2. In the Available Content area, expand Special Content, and then expand Single Report Controls.

The input controls associated with the report in the dashboard are listed.

3. Click an input control, and then select Add to Dashboard.

The input control is added, by default, above the report.

Note

You can select and drag one or more input controls and input control labels to any location on the dashboard.

4. Click **Save** to save the dashboard.

4.2.2.4. Adding a Custom URL to a Dashboard

You can create a frame that displays any URL-addressable content. Such mashups can help you leverage data from many sources in a single, integrated view.

To add a custom URL to a dashboard.

1. Open the dashboard in the dashboard designer.
2. In the Available Content list, select Custom URL (in the Special Content folder), and then drag it to the dashboard.

Analytics prompts you to enter the URL of the Web page you want the frame to display.

3. Optionally, select control values to pass to the URL.

Tip

For example, suppose you want to include the URL to news.google.com in your dashboard, and further want it to display the news for a particular country. In this case, you would select the Country option, and then enter "q" in the URL Parameter name field. This maps the dashboard's Country input control to Google's q (query) parameter. Then, when you select a value in the Country input control, Analytics passes that value to news.google.com so that the news displayed in the frame relates to the selected country.

4. Click **OK** to close the custom URL dialog.
5. Optionally, click the custom URL frame, and then select Auto-refresh Interval from the context menu.

The menu expands to display options for how frequently Analytics will refresh the frame's content. Smaller values update the frame more frequently. By default, a frame is never automatically refreshed (Auto-refresh Interval is set to Manual Only).

6. Select a value, and then click **Save**.

The dashboard is saved to the repository.

4.2.2.5. Refining Dashboard Layout

When your overall layout is complete, you can refine the look and feel of your dashboard.

To refine dashboard layout, you can:

- Reduce unused space in the dashboard. Draw a selection rectangle around all of the dashboard frames, and then drag one or more frames.
- Remove unwanted buttons. Hover over the button, and then click it when the hover border appears. From the context menu, select Delete Item.
- Reposition buttons.
- Add and edit title text. In the dashboard selector, select Add Title Text from the context menu.

4.2.2.6. Dashboard Design Tips

When you are creating or editing a dashboard, keep these tips in mind:

- You can drag items from the Available Content list and drop them to add them to the dashboard, or use the Available Content context menu Add to Dashboard option.
- Press the Ctrl key while dragging or resizing items and frames for smoother cursor movement. This disables the default snap-to-grid behavior.
- Use Ctrl-Click to select multiple frames and items in the dashboard area.
- Click the Hide Content icon to hide the Available Content list. Display it again by clicking the Show Content icon.
- Use the Hide Scroll Bars options on the frame context menu when you can size the frame to exactly match its content. This is especially helpful for charts and small crosstabs.
- The context menu changes depending on your selection. For example, it might include the Delete Item or Delete Frame option, depending on whether you selected a button or a frame. If you select multiple items or frames, the context menu only includes options that apply to all the selected items. For example, if you select a frame and a button, the context menu only includes the Delete Items option; when you select multiple frames, the context menu includes several options that can apply to frames, such as Hide All Scroll Bars and Auto-refresh Interval.
- You can use the arrow keys to move selected content one grid space at a time. Press the Ctrl key to move the selected content a single pixel at a time.
- You cannot resize buttons or input controls.
- When you delete a report with input controls from the dashboard, the controls are also deleted, but their labels remain. Delete them manually. If a custom URL frame is mapped to one of the input controls that was deleted, Analytics shows the default URL without passing the parameter.
- Drag an item from the Available Content list and drop it on an existing frame to replace the existing content.
- To create a new version of a dashboard, open it in the designer and click Save As on the Dashboard Selector context menu.
- A dashboard can include other dashboards, so long as this does not create a circular dependency. Put another way, you cannot add a dashboard to itself.
- If a report displayed in a dashboard is moved to a new location in the repository or if it is deleted entirely, it is deleted from the dashboard.

- Select multiple frames to change their sizes all at once. When you drag the edge of one frame, the other frames change size as well.
- The dashboard designer may seem unwieldy if you design for a screen size as big as or larger than your own; in particular, horizontal scrolling can be awkward. Try hiding the content pane, or using a larger monitor. You can also design the dashboard for your resolution, and then select Use Proportional Sizing from the Dashboard Selector context menu.
- Not all Web pages support the technology underlying Analytics' dashboards (IFrame). In such cases, the URL frame's behavior may be inconsistent.
- By default, Analytics assumes that you want to use the HTTP protocol for custom URL frames. However, you can specify that it use the FILE protocol by entering file:// at the beginning of the value in the URL. In this case, Analytics uses the FILE protocol, and looks for the file you specify in the WEB-INF directory. This is helpful for including images.
- If you have multiple reports that refer to the same input control, but you do not want them to be controlled by a single input control, create two dashboards that each refer to the input control; then create a dashboard that includes those dashboards. Now, users can set the input controls separately for each set of reports.

4.2.2.7. Screen Sizes

When you create a dashboard, you can set the screen size to change the dashboard area to match a particular screen resolution. For example, if you are designing a report for end users who use laptops that only display a screen resolution of 800 by 600 pixels, you can set the dashboard area to 800 by 600 pixels to emulate your end users' screen size.

By default, the dashboard designer supports five standard screen resolutions, which are available under the Screen Size option in the Dashboard Selector context menu. When a dashboard uses fixed sizing, its frames do not resize automatically when the window size changes.

In addition to fixed screen resolutions, dashboards support proportional sizing, which calculates frame sizes using percentages of the total window size, rather than using fixed values. Put another way, when a dashboard uses proportional sizing, its frames resize relative to the size of the browser window. Note that frames do not resize automatically in the designer; proportional sizing is only evident in the end user's view of the dashboard, or when you select a new screen size in the designer.

If a dashboard that uses proportional sizing is bigger than the specified screen size, when the dashboard is viewed, Analytics resizes the frames so that they all appear in the window without requiring the user to use the browser's scroll bars. This may result in a change to the shape of the frames.

Note

Zenoss recommends that you design dashboards using fixed sizing mode. Then, switch to proportional sizing before you save.

In proportional sizing:

- You can resize free text items to a smaller size, but you cannot make them bigger.
- The grid turns red when any content hangs over the edge of the dashboard area.

4.2.3. Editing a Dashboard

To edit a dashboard:

1. Select View Repository, and then search for or browse to the dashboard you want to modify.

By default, the repository includes the /Dashboards folder where you can store your dashboards.

2. Right-click the dashboard and select Open in Designer... from the context menu.

The designer appears and displays the dashboard.

3. Edit the dashboard by adding, removing, resizing, or dragging content.

For more information about working with dashboard content, see the section titled "Creating a Dashboard."

4. To save the dashboard, click **Save**. To create a new version of the dashboard, select Save As from the Dashboard Selector context menu, and then specify a new name.

4.2.4. Designing Reports for Dashboards

When defining reports for a dashboard, keep in mind:

- Charts and small crosstabs are best suited to dashboards. However, you can design table reports that work well in the dashboard. Such reports tend to be very narrow and typically refer to input controls to limit the number of rows they return.
- The reports should be small, since dashboards typically contain multiple reports. In particular, your reports should not be too wide, as horizontal room is always at a premium in a dashboard.
- When an Ad Hoc report is displayed on a dashboard, its margins are stripped. You do not need to strip the margins manually when designing the report.
- If you want a single input control on the dashboard to control the data displayed in multiple reports, the reports themselves need parameters with the same name as the input control. For example, you might have a query-based list of employee names that can be used in both sales reports and human resources report. When defining a parameter in a report, give it a meaningful name that can be reused in other reports. Then, when two reports that include this parameter are added to the dashboard, their input controls appear as Special Content in the Available Content list.

Storing such input controls in the repository encourages their reuse when reports are designed and added to the repository.

- If a report's input control is set to display as a radio button set, it is instead displayed as a list of options.
- If you want to pass a value to an external URL, the URL Parameter Name you specify for the input control must match the name of a parameter that URL can accept. In addition, the value of the input control must also be a value the URL can accept. Beyond these requirements, the target URL is likely to have requirements and limitations. For example, the name of the parameter may be case sensitive; in this case, the value you enter in the URL Parameter Name field is also case sensitive.

The input control must pass data that the URL can accept. Otherwise, Analytics may not be able to retrieve the correct data from the external URL.

Chapter 5. Running Reports

Read the following sections to learn more about:

- Running a simple report
- Running a flash chart
- Running a report with input controls or filters
- Scheduling reports

5.1. Running a Simple Report

To run a report:

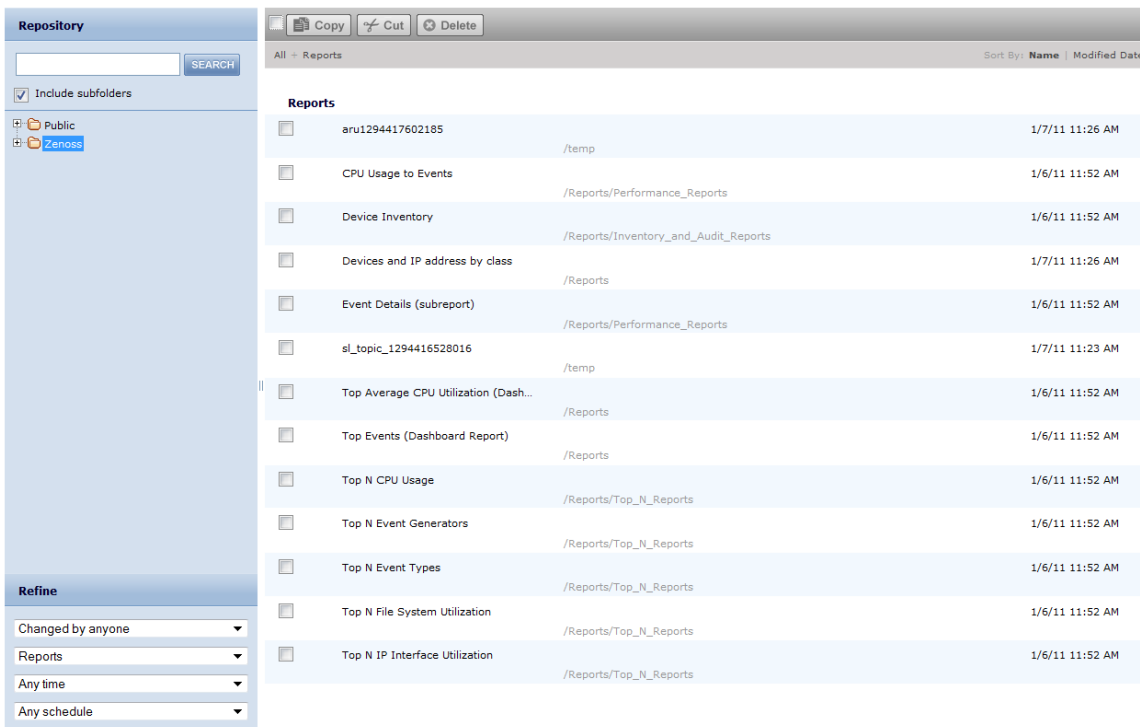
1. From the Analytics Home page, click View Your Reports.

The Repository page appears. Reports that you have permission to view appear.

Note

If you know the name of a report you want to run, you can enter its name in the Search field, and then click Search to locate it in the repository. You also can search on partial names, or for words in a report's description. For more information, see the topic titled "Searching the Repository."

Figure 5.1. Repository: Viewing Reports



Use the scroll bar at the right to view the entire list of reports.

- Click the report to run it. Alternatively, right-click in the report row and then select View from the list of options.

Figure 5.2. View Report



Analytics displays a message indicating it is retrieving the data, and then displays the report. (If needed, you can click **Cancel** to stop the report from running.)

- If desired, select an export icon to view and save the report in another format:




- To return to the list of reports, click the Back icon .

5.2. Running a Flash Chart

To view Flash-based reports from Analytics, Flash must be installed and enabled on your browser. To view these elements in PDF output, you must use a Flash-enabled PDF viewer.

To view a Flash chart, select it in the reports list.

Figure 5.3. Flash Chart




5.3. Changing Report Criteria

When you first view a report, it runs with default values. You can then display the controls and run the report with new values.

1. From the Repository page, run a report.

The report appears, displaying data determined by default values.

2. To run the report with other values, click the Report Options icon: .

The Report Options dialog appears.

3. Select or enter new values, and then click **OK**.

The report runs and appears with data determined by the newly entered values.

Note

You also can click **Apply** in the dialog to refresh the report. This keeps the dialog open, allowing you to further interact with the report and find the data you need.

5.4. Scheduling Reports

The report scheduler enables you to run reports at prescribed times. Jobs that you define in the scheduler run the reports in the background and save the output in the repository. You can view the output there. You also can have the output sent to you and others as email attachments. (For more information, see the section titled "Configuring Email Services" in the chapter "Installing and Configuring Analytics.") Any number of jobs can be scheduled for the same report.


Scheduled reports run as the user who defines the job (see the section titled "Scheduling a Report"). If the report's data source is subject to access permissions, the report contains only the data that the user has permission to see.

Be careful when you schedule a job as an administrative user with no data restrictions. The report will contain all requested data in the data source, and any user that receives the report can view all the data regardless of the user's access restrictions.

To change or delete a schedule, you must be the user who created it or a user with the administrator role.

5.4.1. Viewing Scheduled Jobs

To view the jobs that are currently scheduled for a report, locate the report in the repository. Your options vary depending on whether the report has scheduled jobs:

- In the repository, reports with scheduled jobs have a  icon next to the report name. Click the icon to see the jobs scheduled for the report.
- If a report does not have the icon, right-click it and then select Schedule from the list of options.

The scheduled jobs page appears for the report.

Figure 5.4. Schedule Report Job

Jobs scheduled for /Reports/Performance_Reports /cpu_usage_to_events

ID	Label	Owner	State	Last ran at	Next run time	<input type="checkbox"/>
1 edit	Weekly Report	vmwkst.zenoss.loc/zardemo zenoss	Normal		2011-01-15 18:15	<input type="checkbox"/>

Users see only those jobs they have defined. Administrators see jobs defined by all users.

The list shows information for each job:

- **ID** - Internal ID number of the job, accompanied by a link to edit the job details.
- **Label** - Short description provided when creating the job.
- **Owner** - User who created the job.
- **State** - One of:
 - Normal - Job is scheduled.

- **Running** - Job currently is running.
- **Complete** - Job is complete.
- **Error** - The scheduler encountered an error when scheduling or triggering the job. This does not include cases in which the job is successfully triggered, but an error occurs while it runs.
- **Last ran at** - Last time the report ran.
- **Next run time** - Next time the report will run.


To remove a job, select the option next to the job row, and then click **Remove**.

Buttons on the page are:

- **Back** - Returns to the repository manager or the list of reports.
- **Schedule Job** - Opens the Job Details page to define a new job.
- **Run Now** - Runs the report in the background, allowing you to continue working in Analytics. See the section titled "Running a Job in the Background" for more information.
- **Refresh** - Refreshes the list of jobs. You might want to refresh the list, for example, to see if a job has completed.
- **Remove** - Deletes selected jobs. When a job is deleted while running, Analytics completes the report before deleting the job. Analytics deletes completed jobs automatically.

5.4.2. Scheduling a Report

To schedule a report:

1. From the Analytics interface, select View > Reports.
2. Use the search field or browse the list of reports to find the report you want to schedule.
3. Right-click the report, and then select Schedule from the list of options. If the report has a schedule that you want to add, modify, or delete, click the  icon next to the report name.

The Scheduled Jobs page appears.

4. Click **Schedule Job**.

The Job Details page appears.

Figure 5.5. Job Details

Job | Schedule | Parameters | Output

Report

* Label

Description

5. Enter a name for the job in the Label field, and optionally enter a description.
6. Click **Next**.

The Schedule page appears.

Figure 5.6. Job Schedule

Job | Schedule | Parameters | Output

Time Zone

Start Immediately On


No Recurrence Simple Recurrence Calendar Recurrence


7. Set the schedule attributes:
 - **Time Zone** - Specify the time zone in which you want to define the schedule. The default time zone is that of the server. If you are in a different time zone, then set this field so that the report is scheduled at the time you intend.
 - **Start** - Select Immediately or specify a date in the future. Click the calendar icon to select a future start date and time. If you do not specify recurrence, then the start date and time define when the report will run.
 - **Recurrence** - Choose one of the following:
 - No Recurrence - Run the report once.
 - Simple Recurrence - Schedule the job to recur at an hourly, daily, or weekly interval.
 - Calendar Recurrence - Schedule the job to recur on days of the week or month.
8. Click **Next**.

If the report you are scheduling has input controls that prompt for user input, the Parameters page appears.

Figure 5.7. Job Parameters

Job | Schedule | Parameters | Output

* Report Start Date 

* Report End Date 

Report Top N Limit

Zenoss Instances

Report Device Classes

Groups

Systems

Locations

The content of this page is identical to that on the Report Options dialog. If you have saved values, then they appear in a list of options. The fields that appear are the input controls defined for the report you are scheduling. You can set the input values directly, as well as save them as a named set of values.

9. Choose a set of saved values, or set the input controls as desired.

10. Click **Next**.

The Output page appears.

Figure 5.8. Job Output

Job | Schedule | Parameters | Output

Base output file name

Output description

Output formats PDF HTML Excel CSV DOCX RTF ODT ODS XLSX

Locale

Content Repository

* Folder

Sequential File Names

Timestamp Pattern

Overwrite Files

Email Notification

To

Subject

Message Text

Attach Files

Skip empty reports

11. On the Output page, specify attributes as needed:

- **Base output file name** – The name of the file as it appears in the repository, not including the file extension. The time stamp pattern can be added to the base name, and the format extension is added automatically.
- **Output description** – The output file is stored in a repository object, and this is the optional description for the object.
- **Output formats** – Select at least one. When you select more than one, each format will be stored as a separate file in the repository and attached as a separate file to the email notification.
- **Locale** – Locale for generating the report. The report must support locales (such as a report based on a Domain with language bundles).
- **Folder** – Select the folder where the report output will be saved. This field is required and you must have write permission to the folder. If you email the report and do not want to save it, specify the `/public` folder.
- **Sequential file names** – Appends a time stamp to the names of files created by the job. Useful for the output of recurring jobs or for time-sensitive reports where the output must be dated. When the time stamp is used, the output filename is `<basename>-<timestamp>.<extension>`.

- **Time stamp pattern** – Specify an optional pattern for the time stamp, based on the `java.text.SimpleDateFormat`. However, valid patterns for report output files can only contain letters, numbers, dashes, underscores and periods. The default pattern is `YYYYMMdHHmm` (for example 200906150601). For more information on the valid patterns for this field, refer to:

<http://download.oracle.com/javase/1.4.2/docs/api/java/text/SimpleDateFormat.html>

- **Overwrite files** – If you do not have sequential filenames or if you specify a time stamp pattern that may lead to identical filenames, choose whether newer files overwrite the older ones.
- **To** – For email notification, enter one or more email addresses separated by commas. By default, no mail server is configured in an Analytics installation. To send notifications, the administrator must configure one.
- **Subject** – Subject line of the notification email.
- **Message Text** – Content of the notification email.
- **Attach Files** – Select this option to send the output of the report with the email. Be careful when sending reports containing sensitive data via email.
- **Skip Empty Reports** – In case of an error and the report is empty, select this option to avoid attaching the empty files.

12. Click **Save**.

The job appears in the list of scheduled jobs. (See the section titled "Viewing Scheduled Jobs.")

13. In the list of jobs, verify that the next run time for the new job is the date and time you want.

5.4.3. Changing Job Schedules

When the start date for a job is in the past, you should create a new job rather than changing the outdated one.

To edit a scheduled job:

1. Open the Jobs List page for the report.
2. Click edit in the row of the job you want to change.
3. Make changes on the Job Details, Schedule, Parameters, and Output pages.
4. Click **Save**. Your update applies immediately.
5. In the list of jobs, verify that the next run time for the job is the new date and time you want.

5.4.4. Stopping a Job from Running

To stop a job from running, delete it. To delete a scheduled job:

1. Open the Jobs List page for the report.
2. Select the job you want to delete, and then click **Remove**.

5.4.5. Specifying Job Recurrence

Simple recurrence and calendar recurrence allow you run reports automatically on a regular basis. For each, you specify when the job starts, how often it repeats, and when it ends. Simple recurrence makes the job repeat at an hourly, daily, or weekly interval.

5.4.5.1. Simple Recurrence

Simple recurrence supports these options:

- **Start** - With simple recurrence, the start date and time determines the date and time of all occurrences, according to the chosen interval. For example, if you want to run the job every Monday, set the start date to the first Monday.
- **Occur** - Determines how many times or for how long the job will repeat.
 - Indefinitely - Runs until you delete the job.
 - Until - Runs until a calendar date is reached. Click the calendar to select the date.
 - X times - Runs the specified number of times.
- **Every** - The interval between jobs, in minutes, hours, days, or weeks.

5.4.5.2. Calendar Recurrence

Calendar recurrence enables you to define a job's recurrence at a higher level of detail, including the time of day, days of the week or month, and months of the year.

Calendar recurrence supports these options:

- **End date** - Calendar recurrence runs until a calendar date is reached. Click the calendar to select the date.
- **Minute and Hour** - Set the time of day the job should run. The hours use 24-hour format.

You also can enter multiple minutes or hours, as well as ranges (separated by commas). For example, setting 0, 15, 30, 45 for minutes and 9-17 for hours will run the report every 15 minutes from 9:00 a.m. to 5:45 p.m. Enter an asterisk (*) to run the job every minute or every hour.

- **Days** - Choose the days when the report should run:
 - Every day
 - Week Days - Choose the days of the week.
 - Month Days - Enter dates or date ranges, separated by commas; for example: 1, 15.
- **Months** - Choose the months during which the report will run.

Note

If your server recognizes Daylight Savings Time (DST), then jobs scheduled using simple recurrence may seem to occur one hour later (when DST ends) or one hour earlier (when DST begins). If you want your jobs to recur at the same time of day and respect DST adjustments, use calendar recurrence.

5.4.6. Running a Job in the Background

Running a job in the background lets you generate a report, potentially a long-running one, and keep working in Analytics. It also has the option to export the report directly to any format and save it in the repository. You can also send the generated report by email, which is quick way to share a report with others.

Running a job in the background is equivalent to scheduling the report to run immediately without recurrence.

To run a job in the background:

1. From the Analytics interface, select View > Reports.
2. Use the search field or browse the list of reports to find the report you want to run.
3. Right-click the report, and then select Run in Background from the list of options.
4. If the report you are running has input controls or filters that prompt for user input, then the Parameter page appears. Choose a set of saved values, or set the fields individually.
5. Click **Next**.

The Output page appears. Enter information or make selections, and then click **Save**.

The report runs immediately. When finished, Analytics stores or sends the output as directed on the Output page. While the report runs, you can continue working in Analytics.

Chapter 6. Analytics Report Catalog

This chapter lists the reports included in the Analytics report catalog. Reports in the catalog are divided among these categories:

- Availability
- Inventory
- Performance
- Top "N" Reports
- Virtual Infrastructure
- Dashboard
- Resource Management

6.1. Availability Reports

Report Title	Description
Device Availability by Day	Shows devices with less than 100% availability, by day.
Device Availability by Week	Shows percentage of the total time that the device is reported as available.

6.2. Inventory and Audit Reports

Report Title	Description
Device Inventory	Shows device information grouped by production state, with device class and device comments.
Alerting Rules	Lists alerting rules by user, grouped by Resource Manager server.
Thresholds	Shows minimum and maximum thresholds.

6.3. Performance Reports

Report Title	Description
CPU Usage to Events	Shows the number of events over time, with CPU load as context.
Event Details (subreport)	Shows the events for the device at the time selected.

6.4. Top N Reports

Report Title	Description
Top N CPU Usage	Shows the device, device class, and averaged CPU utilization over a time range.
Top N Event Generators	Shows devices ranked by number of events.
Top N Event Types	Shows event classes ranked by frequency of occurrence.

Report Title	Description
Top N File System Utilization	Shows device file systems, ranked by percentage space used.
Top N IP Interface Volume	Shows device interfaces, ranked by averaged network traffic in and out.

6.5. Virtual Infrastructure Reports

Report Title	Description
VMware Cluster CPU Usage by Time	Graphs cluster usage by aggregated GHz consumed.
VMware Host Metrics by Time	Graphs CPU and memory usage over time.
vCloud vApp Inventory	Lists vApps by Resource Manager server, Cell, Org, and VDC. You must have a vCloud system modeled in your Resource Manager instance to run this report.
VMware Cluster CPU Usage by Time	Graphs cluster usage by aggregated GHz consumed.
VMware Guest Inventory	Shows Guests ordered by cluster and host.
VMware Host CPU Usage by Hour	Graphs CPU usage in percent, averaged per hour over the entire time range.
VMware Host CPU Usage by Time	Graphs CPU usage in GHz by host.
VMware Host Memory Percent Usage by Time	Graphs percent memory utilization.
Top N VMware Host CPU Usage	Shows VMware hosts ranked by CPU utilization.
Top N VMware Host Memory Percent Usage	Shows VMware hosts ranked by memory used.
Top N VMware Host Network Rates	Input and output rates in Mbs, ranked by total rate.

6.6. Dashboard

Report Title	Description
Top Average CPU Utilization (Dashboard Report)	Flash report showing devices ranked by CPU usage percentage.
Top Events (Dashboard Report)	Flash report showing event classes ranked by frequency.

6.7. Resource Management

Report Title	Description
Datapoint Lineage	Shows data points exported into Analytics.

Report Title	Description
Datapoints Not Used	Shows data points not imported into Analytics.

Chapter 7. Understanding the Reporting Database

7.1. Overview

Analytics reports are based on queries to a MySQL reporting database. Data extracted from Resource Manager is translated, and then loaded into the reporting database.

This chapter describes the schema of the Analytics reporting database, providing details about the tables that are queried by the reports included with Analytics. (For more information about the included reports, see the chapter titled "Analytics Report Catalog.")

Zenoss recommends that you read this information if you plan to create domains to support ad hoc reports. Domains provide an abstraction layer between the database schema and the ad hoc report, simplifying the creation of reports. You can learn more about domains in the chapter titled "Creating Domains."

7.2. Schema

At the foundation of the Analytics reporting database are *fact tables* and *dimension tables*.

Fact tables contain metrics that change frequently. Each row of a fact table contains a time stamp and a value for the metric it represents. In the reporting database, fact table names begin with `fact_`, followed by the Resource Manager alias name converted to lower case letters.

Most fact tables are loaded with metrics extracted from RRD files on the Resource Manager collector. (An exception is the `fact_event` table, which is loaded with data extracted from the Resource Manager events database.)

Dimension tables contain persistent or slowly changing entities, and offer a way to group and filter facts.

Dimension tables are dynamically created based on the model data extracted from Resource Manager. Because of this, the dimension tables are dropped and recreated during each load of model data. Dimension table names are prefixed with `dim_`.

The Analytics reporting database schema follows popular conventions of data warehousing. Each fact table is related, through foreign keys, to multiple dimension tables.

Note

For a visual representation of the Analytics reporting database schema, see the diagram appended to this guide.

The schema allows the domain or report author to easily formulate a SQL query by first choosing dimensions, and then joining them to the fact table and using the dimensions to filter the facts.

First you choose your dimensions, and then you join them to the fact table and use the dimensions to filter the facts.

The fact tables in the reporting database contain all of the values extracted from RRD files. Typically, these values are extracted at a resolution of one value every minute or every five minutes. This causes two problems. First, the fact tables contain many rows, which can slow down queries. Second, joining two fact tables (for example, the tables for input and output octets) is difficult because the timestamps are not guaranteed to match.

To solve these problems, the reporting schema includes hourly aggregate tables for each performance fact table. You should use aggregate tables if the report needs one-hour resolution or less (for example, for daily reports).

Aggregate table names are prepended with `agg_`. Aggregate tables follow a similar schema as fact tables, and for the most part can be swapped in and out of the same query with the corresponding fact table.

7.3. Table Relationships

When you create Domains, typically you are interested in one or more fact or aggregate tables and their relationships to dimension tables. All performance fact and aggregate tables follow the same pattern, using these foreign keys:

Key	Associations
<code>device_key</code>	<code>dim_device.device_key</code> or <code>dim_vmware_host.vmware_host_key</code>
<code>component_key</code>	<code>dim_file_system.file_system_key</code> or <code>dim_ip_interface.ip_interface_key</code>
<code>date_key</code>	<code>dim_date.date_key</code>

Dimension tables in the reporting database are not completely de-normalized, making them by definition a snowflake schema. The `dim_device` table uses these foreign keys:

Key	Associations
<code>zenoss_instance_key</code>	<code>dim_zenoss_instance.zenoss_instance_key</code> . Note: Data can be loaded into the reporting database from multiple Resource Manager instances.
<code>device_production_state</code>	<code>dim_production_state.production_state_key</code>

Other dimension tables have columns, named `device_key`, that are foreign keys to `dim_device.device_key`. They are:

- `dim_file_system`
- `dim_group`
- `dim_ip_address`
- `dim_ip_interface`
- `dim_location`
- `dim_system`

7.4. Catalog Reports

This section provides detailed information about the catalog reports provided with Analytics. You can use this information to troubleshoot report problems, or as an example for formulating SQL queries to use as a basis for your Domains.

Reports are grouped according to their containing folder in the resource tree. The following information is provided for each report.

- Report label
- Report title
- Fact and dimension tables used by the report
- SQL query used to generate the report

7.4.1. Availability

Label	Device Availability by Day
Title	Daily Availability
Fact Tables	fct_event
Dimension Tables	dim_date, dim_event_class, dim_device, dim_zenoss_instance, dim_group, dim_system, dim_location

SQL Query

```

SELECT
    dd.date_date,
    d.device_name,
    d.device_class,
    (1 - SUM(unix_timestamp(LEAST(e.event_last_ts, DATE_ADD(dd.date_date, INTERVAL 1 DAY))) -
    unix_timestamp(GREATEST(e.event_ts, dd.date_date))) / 86400.0) as availability
FROM
    dim_date AS dd
    CROSS JOIN fct_event AS e
    INNER JOIN dim_event_class AS ec on ec.event_class_key = e.event_class_key
    INNER JOIN dim_device as d ON d.device_key = e.device_key
    INNER JOIN dim_zenoss_instance as z ON z.zenoss_instance_key = d.zenoss_instance_key
    LEFT JOIN dim_group ON d.device_key = dim_group.device_key
    LEFT JOIN dim_system ON d.device_key = dim_system.device_key
    LEFT JOIN dim_location ON d.device_key = dim_location.device_key
WHERE
    dd.date_date BETWEEN ${report_start_date} AND ${report_end_date}
    AND event_severity >= ${report_event_min_severity}
    AND Date(event_ts) >= dd.date_date
    AND Date(event_last_ts) &lt;= dd.date_date
    AND (ec.event_class_name = '/Status/Ping' OR ec.event_class_name LIKE '/Status/Ping/%')
    AND ${X{IN, d.device_class, report_device_classes}}
    AND ${X{IN, dim_group.device_organizer_name, report_groups}}
    AND ${X{IN, dim_system.device_organizer_name, report_systems}}
    AND ${X{IN, dim_location.device_organizer_name, report_locations}}
    AND ${X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}}
GROUP BY
    d.device_name,
    dd.date_date,
    d.device_class
ORDER BY
    dd.date_date ASC, d.device_class ASC;
    
```

Label	Device Availability By Week
Title	Device Availability
Fact Tables	fct_event
Dimension Tables	dim_date, dim_device, dim_zenoss_instance, dim_group, dim_system, dim_location, dim_event_class

SQL Query

```

SELECT
    dim_date.date_date,
    dim_device.device_name,
    COALESCE(availability.availability, 1.0) as availability,
    dim_date.date_week_of_year,
    dim_date.date_year
FROM
    
```

```

dim_date
CROSS JOIN dim_device
INNER JOIN dim_zenoss_instance as z ON z.zenoss_instance_key = dim_device.zenoss_instance_key
LEFT JOIN dim_group ON dim_device.device_key = dim_group.device_key
LEFT JOIN dim_system ON dim_device.device_key = dim_system.device_key
LEFT JOIN dim_location ON dim_device.device_key = dim_location.device_key
LEFT JOIN (
SELECT
dd.date_key as date_key,
d.device_key as device_key,
(1 - SUM(unix_timestamp(LEAST(e.event_last_ts, DATE_ADD(dd.date_date, INTERVAL 1 DAY))) -
unix_timestamp(GREATEST(e.event_ts, dd.date_date))) / 86400.0) as availability
FROM
dim_date AS dd
CROSS JOIN fct_event AS e
INNER JOIN dim_event_class AS ec on ec.event_class_key = e.event_class_key
INNER JOIN dim_device as d ON d.device_key = e.device_key
WHERE
dd.date_date BETWEEN ${report_start_date} AND ${report_end_date}
AND event_severity >= ${report_event_min_severity}
AND Date(event_ts) >= dd.date_date
AND Date(event_last_ts) <= dd.date_date
AND (ec.event_class_name = '/Status/Ping' OR ec.event_class_name LIKE '/Status/Ping/%')
GROUP BY
d.device_key,
dd.date_key) as availability ON availability.date_key = dim_date.date_key AND
availability.device_key = dim_device.device_key
WHERE
dim_date.date_date BETWEEN ${report_start_date} AND ${report_end_date}
AND ${X}{IN, dim_device.device_class, report_device_classes}
AND ${X}{IN, dim_group.device_organizer_name, report_groups}
AND ${X}{IN, dim_system.device_organizer_name, report_systems}
AND ${X}{IN, dim_location.device_organizer_name, report_locations}
AND ${X}{IN, z.zenoss_instance_fqdn, report_zenoss_instances}

```

7.4.2. Inventory and Audit

Label	Alerting Rules
Title	Alerting Rules
Fact Tables	None
Dimension Tables	dim_user_alerting_rule, dim_zenoss_instance

SQL Query

```

SELECT zenoss_instance_fqdn, zenoss_instance_base_url, user_alerting_rule_id,
user_alerting_rule_address,
user_alerting_rule_next_duration_time,
user_id, user_alerting_rule_criteria,
user_alerting_rule_is_active, user_alerting_rule_next_repeat_time, user_alerting_rule_action,
user_alerting_rule_delay_secs, user_alerting_rule_next_active_time,
user_alerting_rule_url, user_url
FROM dim_user_alerting_rule
JOIN dim_zenoss_instance ON dim_user_alerting_rule.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE ${X}{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND dim_user_alerting_rule.user_alerting_rule_deleted_at is NULL
ORDER BY zenoss_instance_fqdn, user_id, user_alerting_rule_id

```

Label	Device Inventory
-------	------------------

Title	Device Inventory
Fact Tables	None
Dimension Tables	dim_device, dim_zenoss_instance, dim_production_state, dim_group, dim_system, dim_location

SQL Query

```
SELECT *
FROM dim_device
JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
    dim_zenoss_instance.zenoss_instance_key
JOIN dim_production_state ON dim_device.device_production_state = \
    dim_production_state.production_state_value
LEFT OUTER JOIN dim_group ON dim_device.device_key = dim_group.device_key
LEFT OUTER JOIN dim_system ON dim_device.device_key = dim_system.device_key
LEFT OUTER JOIN dim_location ON dim_device.device_key = dim_location.device_key
WHERE $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
    AND dim_device.device_name is not NULL
    AND dim_device.device_name != 'EC2Manager'
    AND dim_device.device_deleted_at is NULL
    AND $X{IN, dim_device.device_class, report_device_classes}
    AND $X{IN, dim_group.device_organizer_name, report_groups}
    AND $X{IN, dim_system.device_organizer_name, report_systems}
    AND $X{IN, dim_location.device_organizer_name, report_locations}
    AND $X{IN, dim_device.device_production_state, report_production_states}
GROUP BY dim_device.zenoss_instance_key, dim_device.device_key
ORDER BY zenoss_instance_fqdn, device_name
```

Label	Top Average CPU Utilization (Dashboard Report)
Title	Top Average CPU Utilization
Fact Tables	fct_cpu_pct
Dimension Tables	dim_date, dim_device, dim_zenoss_instance

SQL Query

```
SELECT
    dim_zenoss_instance.zenoss_instance_fqdn,
    dim_zenoss_instance.zenoss_instance_base_url,
    dim_device.device_key,
    dim_device.device_uid,
    dim_device.device_name,
    AVG(fct_cpu_pct.cpu_pct_value) AS avg_cpu_usage
FROM fct_cpu_pct
INNER JOIN dim_date ON fct_cpu_pct.date_key = dim_date.date_key
INNER JOIN dim_device ON fct_cpu_pct.device_key = dim_device.device_key
INNER JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
    dim_zenoss_instance.zenoss_instance_key
WHERE
    dim_date.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
    AND dim_device.device_name IS NOT NULL
GROUP BY dim_device.device_key
ORDER BY avg_cpu_usage DESC
LIMIT 5
```

Label	Top Events (Dashboard Report)
Title	Most Active Event Classes
Fact Tables	fct_event

Dimension Tables	dim_device, dim_date, dim_zenoss_instance
------------------	---

SQL Query

```
SELECT dim_device.device_key, dim_zenoss_instance.zenoss_instance_fqdn, dim_device.device_uid, \
dim_device.device_name,
myfilter.date_key, myfilter.event_class, myfilter.filter_event_count as event_count, \
SUM(myfilter.filter_event_count) as total_events
FROM
(SELECT fct_event.event_class, dim_device.device_key, dim_device.device_uid, \
dim_device.device_name,
dim_date.date_key, COUNT(*) as filter_event_count
FROM fct_event
JOIN dim_device ON fct_event.device_key = dim_device.device_key
JOIN dim_date ON dim_date.date_key = fct_event.date_key
WHERE dim_date.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
AND fct_event.event_class NOT LIKE '/Change/%'
AND fct_event.event_severity >= $P{report_event_min_severity}
GROUP BY dim_device.device_key, fct_event.event_class
ORDER BY filter_event_count DESC
LIMIT $P{report_topn}
) myfilter
INNER JOIN dim_device ON (dim_device.device_key = myfilter.device_key)
INNER JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
GROUP BY dim_device.device_key, myfilter.event_class
ORDER BY total_events DESC, myfilter.device_name ASC, event_count DESC
```

7.4.3. Performance

Label	CPU Usage to Events
Title	CPU Usage to Event Comparison
Fact Tables	fct_cpu_pct, fct_event
Dimension Tables	dim_device, dim_zenoss_instance, dim_date, dim_group, dim_system, dim_location

SQL Query

```
SELECT
COALESCE(events.zenoss_instance_fqdn,
cpu.zenoss_instance_fqdn) as zenoss_instance_fqdn,
COALESCE(events.zenoss_instance_base_url,
cpu.zenoss_instance_base_url) as zenoss_instance_base_url,
COALESCE(events.device_uid,
cpu.device_uid) as device_uid,
COALESCE(events.device_key,
cpu.device_key) as device_key,
COALESCE(events.device_name,
cpu.device_name) as device_name,
COALESCE(events.date_key,
cpu.date_key) as date_key,
COALESCE(events.date_date,
cpu.date_date) as date_date,
COALESCE(cpu.avg_cpu_usage, 0) as avg_cpu_usage,
COALESCE(events.event_count, 0) as event_count
FROM (
SELECT
z.zenoss_instance_fqdn,
z.zenoss_instance_base_url,
d.device_uid,
```

```

d.device_key,
d.device_name,
dt.date_key,
dt.date_date,
AVG(f.cpu__pct_value) avg_cpu_usage
FROM
dim_device as d
INNER JOIN dim_zenoss_instance as z ON z.zenoss_instance_key = d.zenoss_instance_key
INNER JOIN fct_cpu__pct AS f ON f.device_key = d.device_key
INNER JOIN dim_date dt ON dt.date_key = f.date_key
LEFT JOIN dim_group ON d.device_key = dim_group.device_key
LEFT JOIN dim_system ON d.device_key = dim_system.device_key
LEFT JOIN dim_location ON d.device_key = dim_location.device_key
WHERE dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
AND $X{IN, d.device_class, report_device_classes}
AND $X{IN, dim_group.device_organizer_name, report_groups}
AND $X{IN, dim_system.device_organizer_name, report_systems}
AND $X{IN, dim_location.device_organizer_name, report_locations}
AND $X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
GROUP BY d.device_key, dt.date_key ) as cpu
LEFT JOIN (
SELECT
z.zenoss_instance_fqdn,
z.zenoss_instance_base_url,
d.device_uid,
d.device_key,
d.device_name,
dt.date_key,
dt.date_date,
count(f.device_key) as event_count
FROM
dim_device as d
INNER JOIN dim_zenoss_instance as z ON z.zenoss_instance_key = d.zenoss_instance_key
INNER JOIN fct_event AS f ON f.device_key = d.device_key
INNER JOIN dim_date dt ON dt.date_key = f.date_key
LEFT JOIN dim_group ON d.device_key = dim_group.device_key
LEFT JOIN dim_system ON d.device_key = dim_system.device_key
LEFT JOIN dim_location ON d.device_key = dim_location.device_key
WHERE dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
AND $X{IN, d.device_class, report_device_classes}
AND $X{IN, dim_group.device_organizer_name, report_groups}
AND $X{IN, dim_system.device_organizer_name, report_systems}
AND $X{IN, dim_location.device_organizer_name, report_locations}
AND $X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
AND f.event_severity >= $P{report_event_min_severity}
GROUP BY d.device_key, dt.date_key) as events
ON cpu.device_key = events.device_key AND cpu.date_key = events.date_key

```

Label	Event Details (Subreport)
Title	Device Event Details
Fact Tables	fct_event
Dimension Tables	dim_date, dim_device, dim_zenoss_instance

SQL Query

```

SELECT zenoss_instance_fqdn, zenoss_instance_base_url, dim_device.device_uid, device_name, date_date, \
event_class, event_summary, COUNT(fct_event.event_evid) as event_count
FROM fct_event
INNER JOIN dim_date ON fct_event.date_key = dim_date.date_key
INNER JOIN dim_device ON fct_event.device_key = dim_device.device_key
INNER JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key

```

```
WHERE fct_event.device_key = ${report_device_key}
AND fct_event.date_key = ${report_date_key}
GROUP BY event_class,event_summary
ORDER BY event_count DESC
```

7.4.4. Server

Label	Top N CPU Usage
Title	Top or Bottom CPU Usage
Fact Tables	fct_cpu_pct
Dimension Tables	dim_device, dim_date, dim_group, dim_system, dim_location, dim_zenoss_instance

SQL Query

```
SELECT dim_zenoss_instance.zenoss_instance_fqdn, dim_zenoss_instance.zenoss_instance_base_url, \
dim_device.device_uid, dim_device.device_name, dim_device.device_class, \
AVG(fct_cpu_pct.cpu_pct_value) avg_cpu_usage
FROM fct_cpu_pct
JOIN dim_device ON fct_cpu_pct.device_key = dim_device.device_key
JOIN dim_date ON fct_cpu_pct.date_key = dim_date.date_key
LEFT OUTER JOIN dim_group ON dim_device.device_key = dim_group.device_key
LEFT OUTER JOIN dim_system ON dim_device.device_key = dim_system.device_key
LEFT OUTER JOIN dim_location ON dim_device.device_key = dim_location.device_key
JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE dim_date.date_date BETWEEN ${report_start_date} AND ${report_end_date}
AND ${X}{IN, dim_device.device_class, report_device_classes}
AND ${X}{IN, dim_group.device_organizer_name, report_groups}
AND ${X}{IN, dim_system.device_organizer_name, report_systems}
AND ${X}{IN, dim_location.device_organizer_name, report_locations}
AND ${X}{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
GROUP BY fct_cpu_pct.device_key
ORDER BY avg_cpu_usage ${P!}{report_top_or_bottom}
LIMIT ${P}{report_topn}
```

Label	Top N Event Generators
Title	Top or Bottom Event Generating Devices
Fact Tables	fct_event
Dimension Tables	dim_zenoss_instance, dim_date, dim_device

SQL Query

```
SELECT dim_zenoss_instance.zenoss_instance_fqdn, dim_zenoss_instance.zenoss_instance_base_url, \
dim_device.device_uid,
dim_device.device_name, count(*) AS event_count
FROM fct_event
JOIN dim_zenoss_instance ON dim_zenoss_instance.zenoss_instance_key = \
fct_event.zenoss_instance_key
JOIN dim_date ON fct_event.date_key = dim_date.date_key
JOIN dim_device ON dim_device.device_key = fct_event.device_key
WHERE dim_date.date_date BETWEEN ${report_start_date} AND ${report_end_date}
AND ${X}{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND ${X}{IN, dim_device.device_class, report_device_classes}
GROUP BY dim_device.device_name
ORDER BY event_count ${P!}{report_top_or_bottom}
LIMIT ${P}{report_topn}
```

Label	Top N Event Types
Title	Top or Bottom Event Types
Fact Tables	fct_event
Dimension Tables	dim_date, dim_zenoss_instance, dim_device, dim_group, dim_system, dim_location

SQL Query

```

SELECT fct_event.event_class,COUNT(fct_event.event_class) events
FROM fct_event
JOIN dim_date ON fct_event.date_key = dim_date.date_key
JOIN dim_zenoss_instance ON fct_event.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
JOIN dim_device ON fct_event.device_key = dim_device.device_key
LEFT OUTER JOIN dim_group ON fct_event.device_key = dim_group.device_key
LEFT OUTER JOIN dim_system ON fct_event.device_key = dim_system.device_key
LEFT OUTER JOIN dim_location ON fct_event.device_key = dim_location.device_key
WHERE dim_date.date_date BETWEEN ${report_start_date} AND ${report_end_date}
AND ${X}{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND ${X}{IN, dim_device.device_class, report_device_classes}
AND ${X}{IN, dim_group.device_organizer_name, report_groups}
AND ${X}{IN, dim_system.device_organizer_name, report_systems}
AND ${X}{IN, dim_location.device_organizer_name, report_locations}
AND event_class NOT LIKE '/Change%'
GROUP BY event_class
ORDER BY events ${P!}{report_top_or_bottom}
LIMIT ${P}{report_topn}

```

Label	Top N File System Utilization
Title	Top or Bottom File System Utilization
Fact Tables	fct_usedfilesystemspace__bytes
Dimension Tables	dim_file_system, dim_device, dim_date, dim_zenoss_instance, dim_group, dim_system, dim_location

SQL Query

```

SELECT dim_device.device_key, dim_device.device_uid, dim_device.device_name,
dim_zenoss_instance.zenoss_instance_fqdn,
dim_zenoss_instance.zenoss_instance_base_url,
dim_file_system.file_system_key,
dim_file_system.file_system_name,
dim_file_system.file_system_mount_point,
dim_file_system.file_system_total_bytes,
CEILING(AVG(usedfilesystemspace__bytes_value)) AS avg_fs_used_bytes,
(AVG(usedfilesystemspace__bytes_value) / file_system_total_bytes) AS pct_fs_used
FROM fct_usedfilesystemspace__bytes
INNER JOIN dim_file_system ON (
dim_file_system.file_system_key = fct_usedfilesystemspace__bytes.component_key
AND dim_file_system.device_key = fct_usedfilesystemspace__bytes.device_key)
INNER JOIN dim_device ON dim_device.device_key = fct_usedfilesystemspace__bytes.device_key
JOIN dim_date ON dim_date.date_key = fct_usedfilesystemspace__bytes.date_key
INNER JOIN dim_zenoss_instance ON dim_device.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
LEFT OUTER JOIN dim_group ON dim_device.device_key = dim_group.device_key
LEFT OUTER JOIN dim_system ON dim_device.device_key = dim_system.device_key
LEFT OUTER JOIN dim_location ON dim_device.device_key = dim_location.device_key
WHERE dim_date.date_date BETWEEN ${report_start_date} AND ${P}{report_end_date}
AND dim_file_system.file_system_total_bytes IS NOT NULL
AND dim_file_system.file_system_total_bytes > 0

```

```

AND fct_usedfilesystemspace__bytes.usedfilesystemspace__bytes_value IS NOT NULL
AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND $X{IN, dim_device.device_class, report_device_classes}
AND $X{IN, dim_group.device_organizer_name, report_groups}
AND $X{IN, dim_system.device_organizer_name, report_systems}
AND $X{IN, dim_location.device_organizer_name, report_locations}
GROUP BY device_key, file_system_name
ORDER BY pct_fs_used $P!{report_top_or_bottom}
LIMIT $P{report_topn}

```

Label	Top N IP Interface Volume
Title	Top or Bottom Interface Volume
Fact Tables	agg_inputoctets__bytes
Dimension Tables	dim_date, dim_zenoss_instance, dim_group, dim_system, dim_location

SQL Query

```

SELECT
    d.device_uid,
    d.device_name,
    z.zenoss_instance_fqdn,
    z.zenoss_instance_base_url,
    ip.ip_interface_name,
    ip.ip_interface_description,
        ip.ip_interface_uid,
        ip.ip_interface_speed,
    input_octets.device_key,
    input_octets.component_key,
    input_octets.input_octets,
        input_octets.input_octets_resolution + output_octets.output_octets_resolution \
as total_coverage_secs,
        TIME_TO_SEC(TIMEDIFF(CONCAT(DATE_ADD($P{report_end_date},INTERVAL 1 DAY),' 00:00:0'),
        CONCAT($P{report_start_date},' 00:00:0')) as range_secs,
        DATEDIFF($P{report_end_date}, $P{report_start_date}) + 1 as day_range,
    COALESCE(output_octets.output_octets, 0) as output_octets,
    (input_octets + COALESCE(output_octets.output_octets, 0)) as total_octets
from (
    select
        f.device_key,
        f.component_key,
        sum(3600 * f.inputoctets__bytes_avg) as input_octets,
        sum(f.inputoctets__bytes_resolution) as input_octets_resolution
    from agg_inputoctets__bytes as f
    inner join dim_date as dt on dt.date_key = f.date_key
    inner join (
        select d.device_key
        from dim_device as d
        inner join dim_zenoss_instance as z on d.zenoss_instance_key = z.zenoss_instance_key
        left join dim_group as g on d.device_key = g.device_key
        left join dim_system as s on d.device_key = s.device_key
        left join dim_location as l on d.device_key = l.device_key
        where
            $X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
            AND $X{IN, d.device_class, report_device_classes}
            AND $X{IN, g.device_organizer_name, report_groups}
            AND $X{IN, s.device_organizer_name, report_systems}
            AND $X{IN, l.device_organizer_name, report_locations}
        ) as d on f.device_key = d.device_key
    where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
    group by f.device_key, f.component_key
    ) as input_octets
left join (

```

```

select
  f2.device_key,
  f2.component_key,
  sum(3600 * f2.outputoctets__bytes_avg) as output_octets,
  sum(f2.outputoctets__bytes_resolution) as output_octets_resolution
from agg_outputoctets__bytes as f2
inner join dim_date as dt on dt.date_key = f2.date_key
inner join (
  select d.device_key
  from dim_device as d
  inner join dim_zenoss_instance as z on d.zenoss_instance_key = z.zenoss_instance_key
  left join dim_group as g on d.device_key = g.device_key
  left join dim_system as s on d.device_key = s.device_key
  left join dim_location as l on d.device_key = l.device_key
  where
    $X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
    AND $X{IN, d.device_class, report_device_classes}
    AND $X{IN, g.device_organizer_name, report_groups}
    AND $X{IN, s.device_organizer_name, report_systems}
    AND $X{IN, l.device_organizer_name, report_locations}
) as d on f2.device_key = d.device_key
where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
group by f2.device_key, f2.component_key
) as output_octets on input_octets.component_key = output_octets.component_key
inner join dim_device as d on input_octets.device_key = d.device_key
inner join dim_ip_interface as ip on input_octets.component_key = ip.ip_interface_key
inner join dim_zenoss_instance as z on d.zenoss_instance_key = z.zenoss_instance_key
WHERE ip.ip_interface_name != 'lo'
order by total_octets $P!{report_top_or_bottom}
LIMIT $P{report_topn}

```

7.4.5. Virtual Infrastructure

Label	Top N VMware Host CPU Usage
Title	Top or Bottom VMware Host CPU Usage
Fact Tables	fct_cpu_pct
Dimension Tables	dim_vmware_host, dim_date, dim_zenoss_instance

SQL Query

```

SELECT dim_zenoss_instance.zenoss_instance_fqdn, dim_zenoss_instance.zenoss_instance_base_url, \
  dim_vmware_host.vmware_host_uid, dim_vmware_host.vmware_host_name,
  vmware_host_guest_powered_count as vmware_host_guest_count,
  dim_vmware_host.vmware_cluster_name, dim_vmware_host.vmware_cluster_id,
  AVG(fct_cpu_pct.cpu_pct_value) AS avg_cpu_usage
FROM fct_cpu_pct
join dim_vmware_host on dim_vmware_host.vmware_host_key = fct_cpu_pct.device_key
JOIN dim_date ON fct_cpu_pct.date_key = dim_date.date_key
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE dim_date.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
  AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
  AND $X{IN, dim_vmware_host.vmware_host_id, report_vmware_hosts}
  AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
GROUP BY fct_cpu_pct.device_key
ORDER BY avg_cpu_usage $P!{report_top_or_bottom}
LIMIT $P{report_topn}

```

Label	Top N VMware Host Memory Percent Usage
Title	Top or Bottom VMware Host Memory Percent Usage

Fact Tables	agg_memconsumed
Dimension Tables	dim_date, dim_vmware_host, dim_zenoss_instance

SQL Query

```
SELECT dim_zenoss_instance.zenoss_instance_fqdn, dim_zenoss_instance.zenoss_instance_base_url, \
dim_vmware_host.vmware_host_uid, dim_vmware_host.vmware_host_name,
vmware_host_guest_powered_count as vmware_host_guest_count,
dim_vmware_host.vmware_cluster_name, dim_vmware_host.vmware_cluster_id,
vmware_host_memory as vmware_host_memory_bytes,
(memconsumed_avg / vmware_host_memory) * 102400.0 AS mem_pct_used
FROM dim_date
JOIN agg_memconsumed ON agg_memconsumed.date_key = dim_date.date_key
join dim_vmware_host on dim_vmware_host.vmware_host_key = agg_memconsumed.device_key
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE dim_date.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND $X{IN, dim_vmware_host.vmware_host_id, report_vmware_hosts}
AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
GROUP BY agg_memconsumed.device_key
ORDER BY mem_pct_used $P!{report_top_or_bottom}
LIMIT $P{report_topn}
```

Label	Top N VMware Host Network Rates
Title	Top or Bottom VMware Host Network Usage
Fact Tables	agg_nicrx, agg_nictx
Dimension Tables	dim_date, dim_vmware_host, dim_ip_interface, dim_zenoss_instance

SQL Query

```
SELECT z.zenoss_instance_base_url,
d.vmware_host_uid, d.vmware_host_name,
vmware_host_guest_powered_count as vmware_host_guest_count,
d.vmware_cluster_name, d.vmware_cluster_id,
ip.ip_interface_name,
ip.ip_interface_uid,
input_octets.input_rate_avg,
COALESCE(output_octets.output_rate_avg, 0) as output_rate_avg,
(input_rate_avg + COALESCE(output_octets.output_rate_avg, 0)) as total_rate_avg,
input_octets.input_resolution + output_octets.output_resolution as total_coverage_secs,
TIME_TO_SEC(TIMEDIFF(CONCAT(DATE_ADD($P{report_end_date}, INTERVAL 1 DAY), ' 00:00:0'),
CONCAT($P{report_start_date}, ' 00:00:0'))) as range_secs,
DATEDIFF($P{report_end_date}, $P{report_start_date}) + 1 as day_range
from (
select
f.device_key,
f.component_key,
avg(f.nicrx_avg) as input_rate_avg,
sum(f.nicrx_resolution) as input_resolution
from agg_nicrx as f
inner join dim_date as dt on dt.date_key = f.date_key
inner join (
select d.vmware_host_key from dim_vmware_host as d
where
$X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
AND $X{IN, d.vmware_host_id, report_vmware_hosts}
AND $X{IN, d.vmware_cluster_id, report_vmware_clusters}
) as d on f.device_key = d.vmware_host_key
where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
```

```

        group by f.device_key, f.component_key
    ) as input_octets
    left join (
        select
            f2.device_key,
            f2.component_key,
            avg(f2.nictx_avg) as output_rate_avg,
            sum(f2.nictx_resolution) as output_resolution
        from agg_nictx as f2
        inner join dim_date as dt on dt.date_key = f2.date_key
        inner join (
            select d.vmware_host_key from dim_vmware_host as d
            where
                $X{IN, z.zenoss_instance_fqdn, report_zenoss_instances}
                AND $X{IN, d.vmware_host_id, report_vmware_hosts}
                AND $X{IN, d.vmware_cluster_id, report_vmware_clusters}
        ) as d on f2.device_key = d.vmware_host_key
        where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
        group by device_key, f2.component_key
    ) as output_octets on input_octets.component_key = output_octets.component_key
    inner join dim_vmware_host as d on input_octets.device_key = d.vmware_host_key
    inner join dim_ip_interface as ip on input_octets.component_key = ip.ip_interface_key
    inner join dim_zenoss_instance as z on d.zenoss_instance_key = z.zenoss_instance_key
    WHERE ip.ip_interface_name != 'lo'
    ORDER BY total_rate_avg $P!{report_top_or_bottom}
    LIMIT $P{report_topn}

```

Label	vCloud vApp Inventory
Title	Vcloud Vapp Inventory
Fact Tables	None
Dimension Tables	dim_vcloud_cell, dim_vcloud_vapp, dim_zenoss_instance

SQL Query

```

SELECT
    zenoss_instance_fqdn, zenoss_instance_base_url,
    vcloud_vapp_cloud_name, vcloud_vapp_cloud_href, vcloud_vapp_cloud_description,
    vcloud_vdc_name, vcloud_vdc_href,
    vcloud_admin_org_name, vcloud_admin_org_href,
    vcloud_cell_id, vcloud_cell_uid
FROM dim_vcloud_vapp
JOIN dim_vcloud_cell ON dim_vcloud_vapp.vcloud_cell_key = dim_vcloud_cell.vcloud_cell_key
JOIN dim_zenoss_instance ON dim_vcloud_vapp.zenoss_instance_key = \
    dim_zenoss_instance.zenoss_instance_key
WHERE
    vcloud_vapp_deleted_at IS NULL
    AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
ORDER BY zenoss_instance_fqdn, vcloud_admin_org_name, vcloud_vdc_name, vcloud_vapp_cloud_name

```

Label	VMware Cluster CPU Usage by Time
Title	VMware Cluster CPU Usage by Time
Fact Tables	fct_cpuusagemhz
Dimension Tables	dim_date, dim_vmware_host, dim_vmware_cluster, dim_zenoss_instance

SQL Query

```

SELECT

```



```

dim_vmware_host.vmware_cluster_name,
dim_vmware_cluster.vmware_cluster_uid,
dim_vmware_cluster.vmware_cluster_total_cpu_mhz,
dim_vmware_cluster.vmware_cluster_effective_cpu_mhz,
dim_zenoss_instance.zenoss_instance_fqdn,
dim_zenoss_instance.zenoss_instance_base_url,
fct_cpuusagemhz.cpuusagemhz_ts as timestamp,
sum(fct_cpuusagemhz.cpuusagemhz_value) / 1000.0 as cpuusage_ghz
FROM dim_date
LEFT OUTER JOIN fct_cpuusagemhz ON dim_date.date_key = fct_cpuusagemhz.date_key
JOIN dim_vmware_host ON dim_vmware_host.vmware_host_key = fct_cpuusagemhz.device_key
JOIN dim_vmware_cluster ON dim_vmware_host.vmware_cluster_id = \
dim_vmware_cluster.vmware_cluster_id
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
GROUP BY dim_vmware_host.vmware_cluster_name, timestamp
ORDER BY timestamp

```

Label	VMware Guest Inventory
Title	VMware Guest Inventory
Fact Tables	None
Dimension Tables	dim_vmware_guest, dim_device, dim_zenoss_instance

SQL Query

```

SELECT *
FROM dim_vmware_guest
LEFT OUTER JOIN dim_device ON dim_device.device_key = dim_vmware_guest.device_key
JOIN dim_zenoss_instance ON dim_vmware_guest.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
AND $X{IN, dim_vmware_guest.vmware_host_id, report_vmware_hosts}
AND $X{IN, dim_vmware_guest.vmware_cluster_id, report_vmware_clusters}
ORDER BY zenoss_instance_fqdn, vmware_guest_name

```

Label	VMware Host CPU Usage by Hour
Title	VMware Host CPU Usage by Hour
Fact Tables	fct_cpu__pct
Dimension Tables	dim_date, dim_vmware_host, dim_zenoss_instance

SQL Query

```

SELECT
    HOUR(total_cpu.date_by_hour) the_hour,
    dim_vmware_host.vmware_host_name,
    dim_vmware_host.vmware_host_uid,
    dim_zenoss_instance.zenoss_instance_fqdn,
    dim_zenoss_instance.zenoss_instance_base_url,
    CEILING(total_cpu.hourly_averaged_cpu_pct) avg_cpu_pct
FROM
    (
        SELECT fct_cpu__pct.device_key,
            str_to_date(date_format(cpu__pct_ts, '%Y-%m-%d %H:00:00'), '%Y-%m-%d %H:%i:%s') date_by_hour,
            AVG(cpu__pct_value) hourly_averaged_cpu_pct
        FROM dim_date
        LEFT OUTER JOIN fct_cpu__pct ON dim_date.date_key = fct_cpu__pct.date_key
    )

```

```

WHERE
  date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
GROUP BY date_by_hour, fct_cpu_pct.device_key
) total_cpu
JOIN dim_vmware_host ON dim_vmware_host.vmware_host_key = total_cpu.device_key
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE
  $X{IN, dim_vmware_host.vmware_host_id, report_vmware_hosts}
  AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
  AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
GROUP BY the_hour, dim_vmware_host.vmware_host_name
HAVING the_hour is not NULL

```

Label	VMware Host CPU Usage by Time
Title	VMware Host CPU Usage by Time
Fact Tables	fct_cpuusagemhz
Dimension Tables	dim_date, dim_vmware_host, dim_zenoss_instance

SQL Query

```

SELECT
  dim_vmware_host.vmware_host_name,
  dim_vmware_host.vmware_host_uid,
  dim_zenoss_instance.zenoss_instance_fqdn,
  dim_zenoss_instance.zenoss_instance_base_url,
  fct_cpuusagemhz.cpuusagemhz_ts as timestamp,
  fct_cpuusagemhz.cpuusagemhz_value / 1000.0 as cpuusage_ghz
FROM dim_date
LEFT OUTER JOIN fct_cpuusagemhz ON dim_date.date_key = fct_cpuusagemhz.date_key
JOIN dim_vmware_host ON dim_vmware_host.vmware_host_key = fct_cpuusagemhz.device_key
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
  AND $X{IN, dim_vmware_host.vmware_host_id, report_vmware_hosts}
  AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
  AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
ORDER BY timestamp

```

Label	VMware Host Memory Percent Usage by Time
Title	VMware Host Memory Percent usage by Time
Fact Tables	agg_memconsumed
Dimension Tables	dim_date, dim_vmware_host, dim_zenoss_instance

SQL Query

```

SELECT
  dim_vmware_host.vmware_host_name,
  dim_vmware_host.vmware_host_uid,
  dim_zenoss_instance.zenoss_instance_fqdn,
  dim_zenoss_instance.zenoss_instance_base_url,
  memconsumed_ts as timestamp,
  (memconsumed_avg / vmware_host_memory) * 102400.0 as mem_pct_used
FROM dim_date
JOIN agg_memconsumed ON agg_memconsumed.date_key = dim_date.date_key
JOIN dim_vmware_host ON dim_vmware_host.vmware_host_key = agg_memconsumed.device_key
JOIN dim_zenoss_instance ON dim_vmware_host.zenoss_instance_key = \
dim_zenoss_instance.zenoss_instance_key
WHERE date_date BETWEEN $P{report_start_date} AND $P{report_end_date}

```

```

AND $X{IN, dim_vmware_host.vmware_host_id, report_vmware_hosts}
AND $X{IN, dim_vmware_host.vmware_cluster_id, report_vmware_clusters}
AND $X{IN, dim_zenoss_instance.zenoss_instance_fqdn, report_zenoss_instances}
GROUP BY dim_vmware_host.vmware_host_name, memconsumed_ts
ORDER BY timestamp

```

Label	VMware Host Metrics by Time
Title	VMware Host Metrics by Time
Fact Tables	agg_nicrx, agg_nictx, agg_cpu__pct, agg_memconsumed
Dimension Tables	dim_date, dim_vmware_host, dim_zenoss_instance

SQL Query

```

SELECT  z.zenoss_instance_base_url,
        d.vmware_host_uid, d.vmware_host_name,
        vmware_host_guest_powered_count as vmware_host_guest_count,
        input_octets.timestamp,
        input_octets.input_rate_avg,
        output_octets.output_rate_avg,
        cpu_pct.cpu_pct,
        mem_pct.mem_pct_used
FROM    (
        select
            f.nicrx_ts as timestamp,
            f.nicrx_avg as input_rate_avg
        from  agg_nicrx as f
        inner join dim_date as dt on dt.date_key = f.date_key
        where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
            AND f.device_key = $P{report_vmware_hosts_by_host_key}
        group by timestamp
    ) as input_octets
left join (
        select
            f2.nictx_ts as timestamp,
            f2.nictx_avg as output_rate_avg
        from  agg_nictx as f2
        inner join dim_date as dt on dt.date_key = f2.date_key
        where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
            AND f2.device_key = $P{report_vmware_hosts_by_host_key}
        group by timestamp
    ) as output_octets on input_octets.timestamp = output_octets.timestamp
left join (
        select
            cpu.cpu__pct_ts as timestamp,
            cpu.cpu__pct_avg as cpu_pct
        from  agg_cpu__pct as cpu
        inner join dim_date as dt on dt.date_key = cpu.date_key
        where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
            AND cpu.device_key = $P{report_vmware_hosts_by_host_key}
        group by timestamp
    ) as cpu_pct on input_octets.timestamp = cpu_pct.timestamp
left join (
        select
            memory.memconsumed_ts as timestamp,
            (memconsumed_avg / vmware_host_memory) * 102400.0 as mem_pct_used
        from  agg_memconsumed as memory
        JOIN  dim_vmware_host ON dim_vmware_host.vmware_host_key = memory.device_key
        inner join dim_date as dt on dt.date_key = memory.date_key
        where dt.date_date BETWEEN $P{report_start_date} AND $P{report_end_date}
            AND memory.device_key = $P{report_vmware_hosts_by_host_key}
        group by timestamp
    ) as mem_pct on input_octets.timestamp = mem_pct.timestamp

```

```
inner join dim_vmware_host as d on d.vmware_host_key = $P{report_vmware_hosts_by_host_key}
inner join dim_zenoss_instance as z on d.zenoss_instance_key = z.zenoss_instance_key
ORDER BY timestamp
```

7.4.6. Resource Manager Management

Label	Meta-Data Management: Datapoint Lineage
Title	Fact Meta-Data Report
Fact Tables	None
Dimension Tables	dim_rrd_datapoint, dim_zenoss_instance

SQL Query

```
SELECT dim_zenoss_instance.zenoss_instance_base_url, dim_rrd_datapoint.zenoss_instance_key, \
rrd_datapoint_alias_id,
       rrd_template_id, rrd_template_uid,
       rrd_datasource_id, rrd_datasource_uid, rrd_datasource_enabled,
       rrd_datapoint_id, rrd_datapoint_uid, rrd_datapoint_description,
       TABLE_ROWS as fact_count
FROM dim_rrd_datapoint
JOIN dim_zenoss_instance ON dim_zenoss_instance.zenoss_instance_key = \
dim_rrd_datapoint.zenoss_instance_key
JOIN information_schema.TABLES ON TABLE_NAME = \
CONCAT('fct_', dim_rrd_datapoint.rrd_datapoint_alias_id)
WHERE rrd_datapoint_deleted_at IS NULL
      AND rrd_datapoint_is_alias IS TRUE
ORDER BY rrd_template_id, rrd_datasource_id, rrd_datapoint_id
```

Label	Meta-Data Management: Datapoints Not Used
Title	Datapoints Not Exported Report
Fact Tables	None
Dimension Tables	dim_rrd_datapoint, dim_zenoss_instance

SQL Query

```
SELECT dim_zenoss_instance.zenoss_instance_base_url, dim_rrd_datapoint.zenoss_instance_key,
       rrd_template_id, rrd_template_uid,
       rrd_datasource_id, rrd_datasource_uid,
       id as rrd_datapoint_id, rrd_datapoint_uid, rrd_datapoint_description
FROM dim_rrd_datapoint
JOIN dim_zenoss_instance ON dim_zenoss_instance.zenoss_instance_key = \
dim_rrd_datapoint.zenoss_instance_key
WHERE rrd_datapoint_deleted_at IS NULL
      AND rrd_datapoint_has_aliases IS FALSE
      AND rrd_datasource_enabled IS TRUE
ORDER BY rrd_template_id, rrd_datasource_id, rrd_datapoint_id
```

Chapter 8. Working with Ad Hoc Reports

8.1. About the Ad Hoc Editor

The Ad Hoc Editor lets you create tables, crosstabs, and charts by dragging and dropping elements. For each field, you can:

- Select report types
- Add fields
- Define groups
- Select labels and styles
- Select summaries and data formats

You also can use the editor to explore and analyze your data interactively. The Ad Hoc editor loads data from Topics or Domains, which are repository objects defined to provide a prepared view of a data source. While Topics contain fields used directly in a report, Domains give you the flexibility to filter the data, create input controls, and control the list of available fields.

The following topics provide more information about:

- Working with the Ad Hoc Editor
- Designing an Ad Hoc Report
- Ad Hoc Editor Features
- Creating a Report from a Domain
- Configuring System-Level Ad Hoc Options

8.2. Working with the Ad Hoc Editor

Read the following sections to learn more about working with:

- Topics and domains
- Ad Hoc Editor tools
- Tables, charts, and crosstabs

8.2.1. Topics and Domains

Generally, you will create Topics as JRXML files. A Topic also can be created from a Domain. For more information, refer to the sections titled "Uploading a Topic" and "Saving Domain Settings as a Topic."

Depending on your configuration, Analytics may load a Topic or Domain's entire result set into memory when you edit or run a report. If the data policies are disabled, ensure that each Topic or Domain's query returns a manageable









amount of data, given your environment's load capacity. Alternately, you can change the data loading policy. For more information, refer to the section titled "Data Policies."





8.2.2. Ad Hoc Editor Tool Bar

Figure 8.1. Ad Hoc Editor Toolbar



The tool bar at the top of the editor provides access to many of its functions.

Icon	Name	Description
	Presentation Mode	Click to hide the editor interface and view the report by itself. This mode provides a subset of the Ad Hoc Editor's full feature set. For example, you can change summaries and data types for all report types. You also can change the chart type, legends, and display options in charts, as well as define sorting, order, and column size for tables. If the report includes filters created in the editor, you can change or remove them,
	Save	Select for a menu of save options. You can save the report with its current name and location, or save it with a new name and location by selecting the Save Report As option.
	Run	Click to run the report and view its HTML preview. This view opens in a new window or tab. It provides export options to other formats, including XLS and PDF. When you run a report, you can change its filter and input control values, but you cannot change its layout.
	Undo	Click to undo the most recent action.
	Redo	Click to redo the most recently undone action. Redo is available only immediately after undo; taking any other action disables the Redo icon.
	Undo All	Click to revert the report to its last saved state. After selecting this option, you can click the Redo icon to recreate the steps you took before clicking the Undo All icon.
	Sample Mode	When working with a table or crosstab, click to load the full set of data. By default, Analytics loads a subset of the data in your data source, which improves performance when adding fields, changing grouping, or making other changes that require it to process data. "Sample Data" appears to the right of the Crosstab tab when you are working with a subset of the data. To work with the complete data set, click this icon. The full data set gives you an accurate view off your report, and is helpful for exploring the data. In tables, the subset is helpful for visualizing how your groups will look, and for setting summaries (the summary context menu appears at the bottom of the column).
	Pivot or Switch Group	Click to change the way groups display. Depending on the tab you're viewing, the behavior differs. Crosstab: Pivots the col-

Icon	Name	Description
		umn groups with the row groups. Table: Displays the alternate grouping. Chart: Displays the alternate grouping. For more information, refer to the section titled "Switching Groups."
	Sort	When working with tables, click to view the current sorting (if any) and to select fields by which to sort your data. For more information, see the section titled "Sorting Tables."
	Input Controls	Click to view the input controls applied to this report. In some cases, you can select new values to change the data displayed in the report (for example, the included dates). This icon is available only when input controls are defined for a JRXML-based Topic. For more information see the section titled "Using Input Controls."
	Styles	Click to select a style (fonts and colors) for your report. You cannot undo a style selection. Instead, click the Styles icon and choose a different style.
	Page Properties	Select for a menu of page-level options. You can change whether to display the filter pane or title, or change the page orientation or page size. In tables, you also can hide or show the detail rows when the data is summarized. (This option is available only if the table includes grouped columns.)

8.2.3. Tables, Charts and Crosstabs

The Ad Hoc Editor offers three modes for working with your data:

- Tables

Use tables to view all the data in a report, or to summarize data across one set of related fields. Columns can include properties and measures; each row represents something being measured. You can group rows by their properties.

- Charts

Use charts to compare one or more measures across multiple sets of related fields. Charts summarize data in a graphical representation. With the exception of time series and scatter charts, each chart type compares summarized values for a group.

Depending on your needs, you can configure the chart in several ways: You can group data differently, or use a different chart type.

- Crosstabs

Use crosstabs to summarize data across multiple sets of related fields. Crosstabs are more compact representations than tables. Columns and rows specify properties that are used as dimensions for grouping; cells contain the summarized measurements.

8.3. Designing an Ad Hoc Report

This section discusses basic functionality of the Ad Hoc Editor. Read to learn about:

- Selecting report content
- Formatting table appearance

- Viewing and running a report in the editor

Note

The Ad Hoc Editor also is helpful for exploring and analyzing data; for more information, see the section titled "Exploring and Analyzing Data in the Ad Hoc Editor."

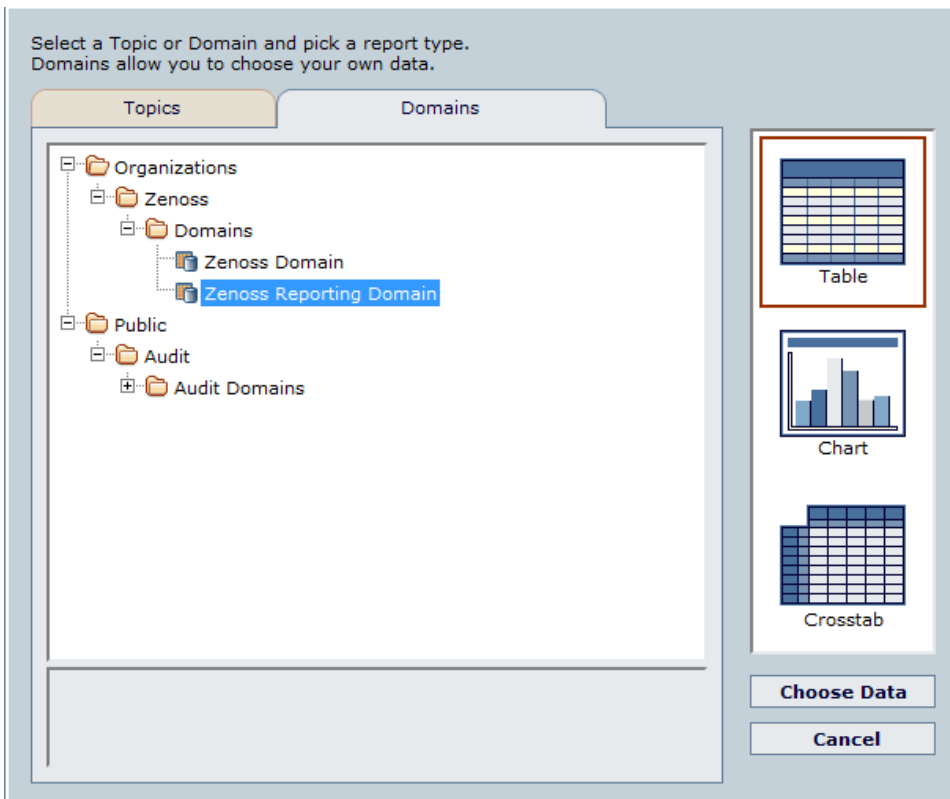
8.3.1. Selecting Report Content

To design a table in the Ad Hoc Editor:

1. From the Analytics Getting Started page, click Create a Report.

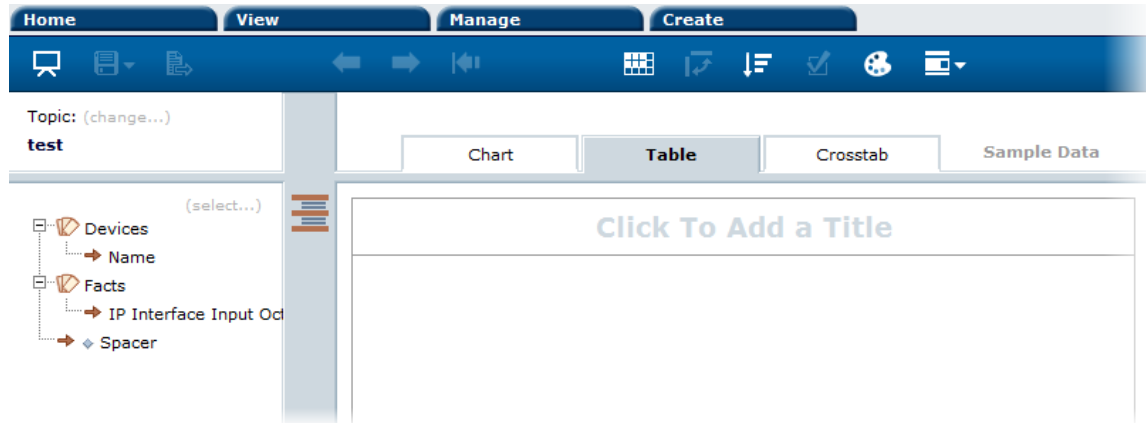
The Topics and Domains window appears.

Figure 8.2. Topics and Domains Window



2. On the Domains tab, expand Organizations in the hierarchy, and then select a Domain.
3. Select Table, and then click **Choose Data**.

The Ad Hoc Editor appears, showing available fields (in the left panel).

Figure 8.3. New Ad Hoc Report


4. At the top of the table (Click to Add a Title), click in the area and enter a title for the report, and then press Enter.
5. In the list of available fields, click ➔ to add desired fields (columns) to the report.

Data for each field appear as you add columns to the table.

Note

To change the list of fields available to the report, click (select...). From the Change field selection dialog, you can select additional fields from the Available list.

To view report data in chart form, click the Chart tab.

6. Select Save the Report from  to save the report.
 - In the Report Name field, enter a report name.
 - Accept the default value for Location Folder, or click **Browse** to select a new location.
 - In the Description field, optionally enter a description for the report.

8.3.2. Formatting Table Appearance

A range of options are available to change the appearance of a table. You can change:

- Report style (colors)
- Group and column labels
- Column width
- Spacing (add "spacers")
- Data format
- Detail view

To format a table in the Ad Hoc Editor:

8.3.2.1. Edit Report Style

To change report style:

1. In the Search field, enter the saved report name, and then click **Search**.

The Search page appears, displaying objects matching the entered text.

2. In the Repository pane, right-click the report, and then select Open in Designer from the options menu.

The report appears.



3.  Click  to select a report style.

Figure 8.4. Styles Window



4. Select a style, and then click **Close**.

8.3.2.2. Change Group and Column Labels

To change the label for a group or column:

1. Right-click the group or column name, and then select Edit Label from the options menu.

Figure 8.5. Edit Label

Zenoss IP Interface Input Data		
Devices	IP Interface Input	Octets
colo2800.zenoss.loc	Name	837
colo2800.zenoss.loc	Use for Sorting...	467
colo2800.zenoss.loc	Add Summary	480
colo2800.zenoss.loc	Edit Label	276
colo2800.zenoss.loc	Delete Label	145
colo2800.zenoss.loc	Remove from Table	360
colo2800.zenoss.loc	Create Filter	796
colo2800.zenoss.loc	Move Left	873
colo2800.zenoss.loc	Move Right	195
colo2800.zenoss.loc		71,301
colo2800.zenoss.loc		7,204
colo2800.zenoss.loc		7,439

The label becomes editable.

2. Enter a new label, and then press Enter.

Tip

If space is at a premium, you can remove labels from the report and make a column only as wide as the data.

8.3.2.3. Resizing Columns

To resize a column:

1. Click a column header.

Borders appear on the column.

2. Move the cursor to the column edge until the resize arrows appear, and then drag the border left or right.

The column resizes to the selected size.

8.3.2.4. Adding Spacers

You can use spacers to create margins. To add a spacer, click ➔ next to the Spacer selection in the list of fields.

8.3.2.5. Editing Data Format

To change the format of data in a column:

1. Right-click the column.

The options menu appears.

2. Click to expand the Change Data Format selection.

A list of available data formats appears.


Figure 8.6. Available Data Formats

Zenoss IP Interface Input Data		
Devices	IP Interface Input Octets	
test-exch2003-1d.zenoss.loc	9,455	<div style="border: 1px solid gray; padding: 5px;"> <i>IP Interface Input Octets</i> ✕ <hr/> Use for Sorting... Add Summary Change Data Format ▲ • -1,235 -1235 (\$1,235) (1,235) <hr/> Create Custom Field ▼ Edit Label Delete Label <hr/> Remove from Table Create Filter <hr/> Move Left </div>
test-exch2003-1d.zenoss.loc	34,845	
test-exch2003-1d.zenoss.loc	39,947	
test-exch2003-1d.zenoss.loc	48,965	
test-exch2003-1d.zenoss.loc	40,403	
test-exch2003-1d.zenoss.loc	40,337	
test-exch2003-1d.zenoss.loc	43,777	
test-exch2003-1d.zenoss.loc	28,343	
test-exch2003-1d.zenoss.loc	29,607	
test-exch2003-1d.zenoss.loc	32,389	
test-exch2003-1d.zenoss.loc	28,551	
test-exch2003-1d.zenoss.loc	20,383	
test-exch2003-1d.zenoss.loc	46,639	
test-exch2003-1d.zenoss.loc	46,276	
test-exch2003-1d.zenoss.loc	30,617	
test-exch2003-1d.zenoss.loc	27,238	
test-exch2003-1d.zenoss.loc	26,965	

3. Select a data format option.



The table refreshes to show the data in the newly selected format.

8.3.2.6. Displaying Detail View

You can choose to display only summarized totals for each group in a table. Select Hide Detail Rows from  to toggle between showing and hiding detail rows.

8.3.3. Viewing and Running a Report in the Editor

When you have created a table, chart, or crosstab, you can save it in the repository for future use. In the Ad Hoc Editor, you can:

- View the report in presentation mode. Click .
- Run the report by clicking . You also can run the report directly from the repository. (See the topic titled "Running a Simple Report.")

8.4. Ad Hoc Editor Features

The Ad Hoc Editor offers a range of features. Read the following sections to learn more about:

- Creating custom fields
- Using input controls and filters

- Switching groups
- Sorting tables
- Working with charts
- Working with crosstabs

8.4.1. Creating Custom Fields

You can create new fields in a report by applying mathematical formulas to an Ad Hoc report's existing numeric fields.

The Ad Hoc Editor supports three types of functions:

- Basic functions

These include addition, subtraction, multiplication, and division. You can use these functions with a constant (such as dividing by 1024.0 to convert from bytes to kilobytes), or you can use them with multiple fields. Select multiple fields by using Ctrl-Click. Column borders change color to indicate that multiple fields are selected. Right-click to open the context menu and create a custom field.

- Special functions

Special functions include round, percent of total, percent of group, and rank. These functions do not take constants, nor do they support multiple fields.

- Date functions

Date functions calculate the difference between two dates in a number of intervals, such as years, weeks, days, and seconds.

A custom field can include only a single function. To use more than one function to create the field you need, create two custom fields, with one building on the other. In this case, pay special attention to the order in which you create the custom fields, as this may affect the results. For example, rounding then multiplying yields different results than multiplying and then rounding.

Because custom fields can build on one another, you can create complex calculations. When a custom field is the basis of another field, you cannot delete it until you delete the one that builds on it.

When working with multiple fields:

- You can select only basic functions and data differences
- You can multiply or add any number of fields, but ordered functions (subtract and divide) and date functions can be used only with two fields at a time.
- For ordered operations, the order in which you select fields matters.

To create a custom field based on one field:

1. In the list of fields, right-click a numeric field, and then select Create Custom Field from the context menu.
2. Depending on the type of function you want to create, click **Basic Functions** or **Special Functions**.
3. If you selected a basic function, select Add #, Subtract #, Multiply #, or Divide #, and then enter a number in the field that appears; this number is the constant to use in the formula.

4. If you selected a basic function, click **Swap Field and Number** to put the constant before the field in the formula. This is useful only in conjunction with ordered functions. You can toggle the order of the field and function from the options menu (Edit Formula Swap Number and Field).

5. Click **Create Field**.

The custom field appears in the list of fields.

To create a custom field based on multiple numeric fields:

1. In the list of fields, Ctrl-Click two or more numeric fields, right-click one of the selected fields, and click Create Custom Field from the context menu.

2. If it is not expanded, click Basic Functions, and select Add, Subtract, Multiply, or Divide.

You can subtract or divide only when exactly two fields are selected.

3. Select Swap Fields from the options menu to change the order of the fields in the formula. This is useful only in conjunction with ordered functions.

4. Click **Create Field**.

The custom field appears in the list of fields.

To create a custom field based on two date fields:

1. In the list of fields, Ctrl-Click two date fields, right-click, and then click Create Custom Field from the options menu.

2. If it is not expanded, click Data Difference, and then select an interval.

You can subtract dates only when exactly two dates are selected.

3. Click Create Field.

The custom field appears in the list of fields.

When creating custom fields, keep these actions and limitations in mind:

- To edit an existing custom field, right-click it and select Edit Formula from the context menu. You can then choose a different function, or enter a different value (if you are editing a basic function).
- When you create a custom field based on a field in the table, the new field appears in the report area; when you create it from a field in the list of fields, it appears at the bottom of the list.
- The percent of group functions are unique among the custom functions in that they calculate values based on the grouping defined in the report. Thus, the following limitations apply:
 - The percent of column group function are not meaningful in charts and tables. In these contexts, the function always returns 100.
 - The percent of group functions cannot be used as a group.
 - The percent of column group parent is meaningless in table; when added to a table, it always returns the value 100.
- Custom fields using the percent of group functions cannot be the basis of other custom fields. This limitation also applies to the Percent of Total function.

- Custom fields using the percent of group functions cannot be used as filters.
- The summary functions for Custom fields using the percent of group functions can't be edited. This limitation also applies to the Percent of Total function.
- The Round function uses the standard method of rounding (arithmetic rounding), in which decimal values greater than 0.5 are rounded to the next largest whole number.
- A custom field's label is determined by the fields, constants, and functions it includes.
- By default, the Ad Hoc Editor supports only two decimal places.
- You cannot delete a custom field that is used in the report area. First, remove the custom field from the report area, and then delete it from the list of available fields.


8.4.2. Using Input Controls and Filters

Topics and Domains use different mechanisms for narrowing the data they return:



- Topics can contain parameterized queries. The parameters can be mapped to input controls that allow you to select the data you want to include.
- Domains (and Domain Topics) can be filtered by selecting fields in the Domain and specifying comparison values. The filters can be configured to allow you to select the data to include.
- Within the Domain design, filters also can be defined. They are never displayed in the editor or when the report is run.

In the Ad Hoc Editor, you also can defined filters, whether you're working with a Domain or Topic.

Input controls and filters interact seamlessly. For example, when a Topic includes input controls that prompt users, and

you also create filters, the Ad Hoc report has both input controls (displayed by clicking ) and filters (displayed in the filters pane). Since this can results in a combination of input controls and filters that do not return data, the system refreshes the editor against both the filters and the input controls. Nevertheless, some combinations of filters and input controls may results in no data being returned.

Note

If your results are empty, but you do not know why, click  (Toggle Filter Pane), and then click . You may find a combination that explains an empty results set, such as a standard filter set against a Country field and a complex filter set against both Country and State.

8.4.2.1. Using Input Controls

In the Ad Hoc Editor, you can display the input controls defined in the Topic as visible to users. You can accept the queries' default values or enter other values.

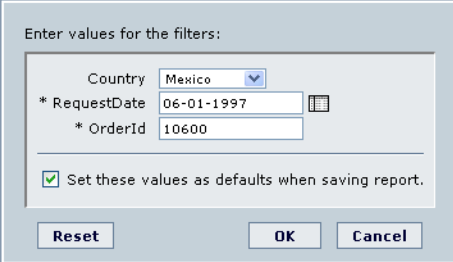
The Ad Hoc Editor indicates that the report has input controls by displaying  as active. Click this icon to select new values or to save values as the new defaults for this report.

To edit the values for a report's input controls:

1. Click .

A window appears and lists the current values for the input controls.

Figure 8.7. Input Controls in the Ad Hoc Editor



2. Select new values.
3. If appropriate, select the Set these values as defaults when saving report option.

When you select this option, the selected values are made the default values. Otherwise, the previous default values are used.

4. Click **OK**.

The report appears. The data that appears is determined by the new values.

8.4.2.2. Using Filters



You can define filters at three levels:

- In the Domain Design
- When creating a report from a Domain (in the Choose Ad Hoc Data wizard).
- In the Ad Hoc Editor (even when the report is based on a Topic).

To create a filter:

1. Right-click a field in the list of available fields, and then click Create Filter.

A new filter appears in the filter pane. If the filter pane was hidden, it appears when you create the filter.

2. Depending on the data type of the selected file, the filter may be multi-select, single-select, or text input.
3. Use the fields in the filter to change its value.
4. Click  to expand the filter and change the filter's operator.
5. Click  to remove the filter.
6. Click the down arrow to hide the filter's details. Click the right arrow to display them again.
7. Click the All check box to select all values currently available in the data set.

Note

The All option does not guarantee that all values will be selected every time the report is run. Instead, the option is a shortcut to help you quickly select all the values currently available in the data set. To ensure that all values appear in the report whenever it is edited or run, remove the filter entirely.

You also can create a filter from groups in tables and crosstabs, or from columns in tables. On the **Chart** tab, you must right-click the field in the list of those available.


Note

Filters are also created when you right-click a group member in a crosstab and select Keep Only or Exclude. The filter pane does not automatically open in this case. When you create a filter against an inner group, the filter that appears may be created as a complex filter; a complex filter cannot be edited but it can be removed. Note that similar complex filters also appear in the Ad Hoc Editor if a Choose Ad Hoc Data wizard filter was created and set to not prompt for values.

8.4.2.3. Input Controls and Filters Availability

Input controls and filters can appear in the editor and when the report is run:

- Input controls can be set to be hidden or displayed:
 - Input controls set to prompt are always displayed in the editor and when the report is run.
 - Input controls that are not set to prompt are always hidden in the editor and when the report is run.
- Filters defined in the domain design are always hidden in the editor and when the report is run.
- Filters created in the Choose Data Wizard can be set to be hidden or displayed.
 - Filters set to prompt are always displayed in the editor and when the report is run.
 - Filters not set to prompt are displayed in the editor but hidden when the report is run. You can remove the filter while in the editor, allowing you to see all data unfiltered.

When a report is run, any input controls or filters that are set to prompt are available. If they do not display directly on the report page, click  to view them.

To determine whether an input control is displayed:


1. Edit the Topic. On the Controls and Resources page, click the input control.
2. Click Next, and then select the Visible option to change whether the input control appears.

If you do not provide a default value for the input control, you are prompted to select a value when creating a report against the Topic.


To determine whether a filter defined in the Choose Ad Hoc Data wizard can be edited:

1. Create a report against a Domain.
2. On the **Filters** tab, create a filter and then click **OK**.
3. To the right of the new filter, click the Prompt? box to de-select this option if you do not want users to edit this filter's value in the editor.

8.4.3. Switching Groups

In the Ad Hoc Editor, the  icon lets you change how the groups are displayed. Depending on the type of report you are viewing, the behavior varies:

- In crosstabs, the system moves the row groups to column groups, and vice versa. This action is known as *pivoting*. This option is available only for crosstabs that have a column group and a row group.
- In charts, the system displays the alternate group. This option is only available for charts that have a column group. When multiple row or column groups are defined, the **Chart** tab displays the outermost group.
- In tables, the system displays the alternate group. This option is only available for tables that have a column group. When multiple row or column groups are defined, the **Table** tab displays the first group in the set (the outermost group).


By default, the table or chart is grouped by the column group. Click  to group it by the row group.

For this feature to be available for tables and charts, the report must have been edited on the **Crosstab** tab to add a column group. Otherwise, there is no alternate group to switch to.

When you design a table or chart and then view the **Crosstab** tab, you only have a row group defined.

Note

Create the table or chart with the most important grouping field, then click Crosstab to add another compelling field as the column group; the field you select is now available as an alternate group in the table and chart.

Click **Table** or **Chart** to return to your initial view, and then click  to display the alternate group. For example, you could create a bar chart that showed your sales by region. By switching to the **Crosstab** tab, you could add Client Industry as the column group. Back in the Chart tab, switch the groups to display the sales data grouped by Client Industry. This comparison can help you identify trends and outliers in your data.

8.4.4. Sorting Tables

In the Ad Hoc Editor, you can sort the rows of a table by one or more fields. You also can sort crosstabs.

To sort a table:

1. Click .

The Sorting window appears. If the table is already sorted, the window shows the fields that are used.

2. To add a field, click **Add Fields**.

The Fields window appears with a list of available fields. It lists only fields that are not already used.

3. Select one or more fields to sort by, and then click **OK**. (Ctrl-click to select multiple fields.)

The Sorting window reappears with the selected fields.

4. To arrange the sorting precedence of the fields, select each field in the Sorting window and click the Up, Down, Top, or Bottom icon.

5. To remove a field, select it and then click X.
6. Click **OK**.

The Table updates to display the rows sorted by the selected fields.

Note

When sorting a table:

You can click a field in the list of fields and select Use for Sorting from the context menu. In this case, the table is sorted by a field that is not in the table. You may want to note the sorting fields in the title.

If a column is already being used and you want to stop using it or change the sorting, right-click the column and then select Change Sorting from the context menu. In either case, the Sorting Window appears and you can define the sorting as above.

To sort a crosstab, click the icon next to a measure's label.

8.4.5. Working with Charts

Charts are graphical representations of your data. The Ad Hoc Editor can create both summarized and unsummarized charts:

- Summarized charts show data grouped and summarized:
 - Bar - Compares values displayed as bars.
 - Pie - Compares values displayed as slices of a circular graph.
 - Line - Compares values displayed as points connected by lines.
 - Area - Compares values displayed as shaded areas.
- Unsummarized charts show each data point individually, and do not support grouping:
 - Scatter - Compares values expressed as pairs of values that determine the position of a point on the chart.
 - Time Series - Compares values based on times or dates associated with the data.

When you create a chart, you generally indicate how to arrange the data by specifying:

- A group field. Similar to the table's grouping option, a chart's group field determines how data are categorized. For example, in a bar chart, the category field determines how the bars are grouped; in a pie chart it determines what each slice represents. Note that time series and scatter charts do not support grouping; instead, each data point is mapped individually. The data are not summarized.
- Measure fields. Measures are the other data in your chart. They determine the length of bars, size of pie slices, location of points (in line charts), and height of areas.

8.4.5.1. Designing Charts

Charts have a different layout and format options than table or crosstabs. Keep these points in mind:

- By default, the Ad Hoc Editor creates a bar chart. You can select a different type of chart at any time.
- The Ad Hoc Editor automatically hides all but the first measure if you change to a pie chart, since pie charts can only show a single measure.

- For each type of chart, the context menu provides these display options:


Type	Description
All Chart Types	Regardless of the type of report, you can: Remove Gradient or Apply Gradient. Determines whether the background color is graded or solid. When the gradient is removed, the background is filled with a solid color. Remove Background or Fill Background. Determines whether the background should be filled. The gradient does not appear when you run the report.
Bar Chart	When creating a bar chart, the context menu provides these options: Flat Bars or 3D Bars. Determines whether to use a flat, two-dimensional appearance or a three-dimensional (3D) appearance. Horizontal Bars or Vertical Bars. Determines whether the bars are parallel to the width or height of the chart. Stacked or Simple Series. Whether multiple measures of each group are depicted as individual bars or as portions of a single bar. A series chart displays multiple measures for each group value.
Pie Chart	When creating a pie chart, you can indicate whether to use a Flat Pie or the default 3D Pie.
Line Chart	When creating a line chart, you can indicate whether to display markers at each data point and lines that join them. Select Hide Lines or Hide Points to remove the indicators from the chart.
Area Chart	When creating an area chart, you can indicate whether the groups should be displayed as transparent overlays or as solid area. Select Stacked Series to display the groups as solid areas arranged vertically, one on top of another; select Simple Series to display groups as transparent overlays. As compared to line charts, area charts emphasize quantities rather than trends.
Scatter	When creating a scatter chart, you can indicate whether to display markers at each data point and lines that join them. Click Hide Points to remove the marker at each data point. Click Show Lines to add lines that join the data points.
Time Series	When creating a time series chart, you can indicate whether to display markers at each data point and lines that join them. Click Show Points to add a marker at each data point. Click Hide Lines to remove the lines that join the data points.


- Before any fields are added to the chart, the Ad Hoc Editor displays a simple placeholder with the legend displaying a single entry: No Data. As you add groups and measures, the mock-up changes to reflect the data in the chart. For example, when you add a group to a pie chart, the pie displays the correct number of pie slices; when you add a measure to a bar chart, the Y-axis units and values show the range of values in the measure you selected. Once you have added at least one group and one measure, the mock-up is replaced by the chart, and you can save or run the report.
- Fields can be added as groups or measures. To add a field as a group, drag it to the group target to the left of the report area. To add a field as a measure, drag it into the report area or right-click it in the list of fields and select an option from the context menu. Note that there is no right arrow for adding fields, and you can add fields more than once; for example, you can add a field as both a measure and a group.
- You can rename the measures displayed in your chart by right-clicking the measure's label in the legend and selecting **Change Legend Label** from the menu that appears. The text you enter replaces the field's name. This option is also available on the chart context menu.

If you change the measure's summary function, the label updates to reflect the change. If you change the summary function back, the last name entered for that measure with that summary function appears.

- You can increase a chart's size to improve its legibility or you can decrease its size if it is intended for use in an Analytics dashboard. When you place your cursor over the right-hand or bottom edge of the chart, the resize arrow appears; click and drag the edge of the chart to increase or decrease its size.
- By default, the page is set to **Actual Size**, which resizes the page to fit your chart's dimensions.



When you place your cursor over  and select **A4 Size** or **Letter Size** from the menu, the Ad Hoc Editor limits the size of the chart to the selected page size. If you resize the chart to be larger, the Ad Hoc Editor snaps the chart back to the maximum allowable size in either dimension (height or width).

- If space is at a premium in your chart, or if you want to use it in a portal, place the cursor over  and click **Strip Margins**. The Ad Hoc Editor removes the margins, giving your chart more horizontal and vertical room. Note that when you strip the margins of a chart set to the **Actual Size** page size, the change is not reflected in the Ad Hoc Editor. You can see the change by setting the page size to Letter or A4.

- A scatter chart uses Cartesian coordinates to display values as points whose locations are determined by two variables you select. One variable determines the horizontal position and the other determines the vertical position. It plots paired values against one another.
- Time series charts plot your data against the passage of time. They can help you understand the nature and underlying context of your data, identify trends, or forecast future results. In a time series chart, the horizontal axis must represent a date or time field, such as Purchase Date. If you add a field to the horizontal axis that has a data type other than a date/time, the Ad Hoc Editor renders the chart as best it can, but the result is likely to be meaningless.
- Unsummarized charts (scatter and time series) can only be run or saved if the fields selected for the X and Y axis are appropriate. For scatter charts, both fields must be numeric. For time series charts, one field must be a date or date/time, and the other must be numeric.
- You can hide legends and labels to allow the chart more space, or you can show them to make the chart easier to understand. In the chart context menu, click **Legends and Labels**, and specify whether to display the legend (generally, the list of measures in the chart; for pie, it is the list of groups), the group title (which identifies the name of the group field), or measure titles (which identifies the measures).
- In bar, line, and area charts that display multiple measures, you can change the order of measures by dragging the legend entry right or left. This reorders the measures in each group. While the order of measures is changed, the order of colors is not. Thus, the measures change color as they are reordered. This can be helpful improving a chart's legibility and visual appeal.
- Right-click a measure in the legend to change its label.
- Remove measures from the chart using the legend or chart context menu. If the chart only displays a single measure, right-click the chart and select **Remove Measure**. If the chart displays multiple measures, the context menu lists them under the Measures menu option. The measures are listed in the order they appear within your groups. On the chart context menu, click Measures [Measure Name] Remove Measure (where [Measure Name] is the name of the measure you want to remove).

Note

The Measures menu also allows you to change the summary function and data format. It is available from the chart context menu as well as from the legend.

- Remove a chart's group by right-clicking the chart and selecting Remove Group from the context menu. You can also right-click a different field in the list of fields, and click Replace Group, or simply drag the new group field and drop it on the group control.


- When working with a chart that displays two measures with very different ranges of values, the chart may be difficult to read or understand. To make the ranges of values more similar, create a custom field that multiplies or divides one of the measures. Consider the case where you are comparing the number of items sold (values ranging from 0-50) and the total sales revenue for each product (values ranging from 0-50,000), you could create a custom field that divided the revenue measure by 1000. Adding this measure in place of the actual revenue measure yields a chart in which the two measures are depicted in closer scale. The new chart is both easier to read and more appealing. Zenoss recommends that you note such changes in the legend label or title.
- To add a field in a specific position in a chart, you can drop it in the legend in the desired location. The new measure appears in the correct location in the chart.

8.4.6. Working with Crosstabs

Crosstabs have different data, layout, and format options than tables or charts. Keep these points in mind:



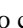



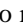
- Fields can be added as row groups, column groups, and measures. To add a field, drag it into the crosstab or right-click it in the list of fields and select an option from the context menu. Note that there is no for adding fields, and you can add fields more than once; for example, you can add a field as both a measure and row group.
- Measure labels are displayed to the left of the first column group. To edit a measure, right-click its label.
- Row and column groups are arranged in hierarchies. Drag the group headings to rearrange the hierarchy; you can also right-click a heading and select a Move option from the context menu or press the cursor keys. Rearranging the groups may change the preview data in the editor.
- When the crosstab includes more than one row group or more than one column group, click the arrow in each heading cell to expand the inner group. Click the arrow again to collapse the inner group once more.

When you collapse a group, its summary is automatically displayed; this prevents invalid crosstab layouts in which there is nothing to display for some totals if the summary has been deleted previously.

- You can pivot a crosstab in two ways:
 - Pivot the entire crosstab. Click  to switch row and column group places.
 - Pivot a single group. To pivot a single row group, right-click it and then select Switch to Column Group. To pivot a single column group, right-click it and then select Switch to Row Group.
- You can use the slice menu to keep or exclude group members in a crosstab. To slice, right-click a group member, and then select:
 - Keep Only to remove all groups except the selected one from the crosstab.
 - Exclude to remove this group from the crosstab.

Use Ctrl-click and Shift-click to select multiple groups to keep or exclude. You can select multiple row groups or multiple column groups; you cannot slice by both row groups and column groups at once. Compare slice to drill-through, drill to details, and filtering.

- Measures are arranged in cells. You can add any number of measures. All the measures appear together in every cell. To rearrange the measures, drag them in the measure label area; you can also right-click a measure label and select a Move option from the context menu, or press the up and down arrow keys on your keyboard.
- All row and column groups are summarized automatically:

- To turn off a group summary, right-click any heading in the group and select Delete Row Summary or Delete Column Summary from the context menu. Note that this option is only available for the outermost group on either axis (that is, either the outermost row group or the outermost column group).
- To select the summary function and data format for a measure, right-click the measure label and select from the context menu. Note that you cannot change the summary function on custom fields that calculate percents (Percent of Total, Percent of Column Group Parent, and Percent of Row Group Parent).
- The summary functions for numeric fields are Maximum, Minimum, Average, Sum, Distinct Count, and Count All. Distinct Count is the number of different items in the row or column; Count All is the total number of items. For instance, if there are 3 widgets of type A and 3 widgets of type B, Distinct Count is 2 and Count All is 6.
- On the Crosstab tab, by default, the editor displays each row and column group in a collapsed state. This means that you can see the totals for the group, but not the measures for its individual members. To see measures for a group's members, right-click the group label and select Expand Members. When a group's members are expanded, select Collapse Members from the same menu to hide the measures. Collapsing an outer group also collapses its inner groups. The Expand Members and Collapse Members options are only available for outer-most groups or inner groups nested in an expanded outer group.
- On the Crosstab tab, by default, the editor only displays a smaller, sample set of the data in your crosstab. Click  to view the full set of data.
- By default, a crosstab is sorted in alphabetical ascending order according to its group member names. Only one measure can be used for sorting at any one time. You can sort it by any of its measures:
 - Click  to change the sort order from the default. The icon changes to  and the crosstab is sorted in ascending numeric order by its measure values.
 - Click  to switch the numeric order from ascending to descending. The icon changes to  and the crosstab's sort order is reversed.
 - Click  to return to the default sort order. The icon changes to  and the crosstab is once again alphabetically sorted by its group names.

When you sort by a measure, the entire crosstab is reordered: both columns and rows reflect are re-sorted within each group. When the crosstab includes more than one row group or more than one column group, the inner groups are also sorted according to your selection.

- Many of the layout and formatting options set manually in tables are set automatically in crosstabs. In particular, row and column sizes are fixed and no spacer is available.
- Hover your cursor over a measure value or total and notice the underline that indicates a hyperlink. Click the hyperlink to open the Table tab of the Ad Hoc Editor in a new window or tab (depending on your browser and its settings). This new table is filtered to only show rows that were rolled up to create the hyperlinked value you selected. Its columns are created from the fields used as measures in the crosstab. This new instance of the editor is completely independent of the original crosstab: you can change it and save it without changing the crosstab.

8.5. Creating a Report from a Domain

Like Topics, Domains are created by data analysts and Analytics administrators for use by others to access data more conveniently. For report makers, Domains differ in that they give more flexibility in choosing fields from the database and allow some filtering of the data before it is included in a report.

Domains also allow you create reports that prompt the reader for input that determines what data is presented. For example, if a Domain includes all sales data for a company, a report based on that Domain could present detailed information grouped by postal code and have an input control to let the reader select the geographic area, such as a US state.

8.5.1. Choosing Ad Hoc Data from a Domain

When you create an Ad Hoc report based on a Domain, the Choose Ad Hoc Data wizard leads you through the steps of choosing fields and filters.

To create a basic report from a Domain:

1. From the Analytics Getting Started page, select Create Ad Hoc Report.

The Ad Hoc Editor opens to the Topics and Domains window.

2. Click the **Domains** tab.

The tab opens with a list of available Domains. The frame at the bottom shows you the description of the Domain currently selected in the list.

3. Select the Domain, click Table, and then click **Choose Data**.

The Choose Ad Hoc Data wizard opens to the Data page.

4. On the Data page, select the fields that appear in the Ad Hoc Editor. You can select whole sets or individual items within a set. A set appears in the Ad Hoc Editor as a set of fields that can be added together to a table.

Note

If a Domain is configured for secure access, you see only sets and items for which you have permission. Other items might not be visible to you, even though another user can see them if he used the same Domain. Conversely, the fields that you can see might include ones that other users cannot see.

5. Click **Next**.

The Choose Ad Hoc Data wizard shows the Filters page. Here, you can set limits on the data that are retrieved and displayed in your table.

6. Click **OK**, and then click **Next**.

The Choose Ad Hoc Data wizard shows the Display page. On this page, you can edit the table and field names that will appear in the Ad Hoc Editor so that they are more appropriate for the report's audience.

7. If you want to create different reports with the same settings, you can save them as a Domain Topic. (See the section titled "Creating Topics from Domains.")

8. To confirm all selections made in the wizard and being creating a report, click **Open Report**.

9. Design your report in the Ad Hoc Editor. (See the sections titled "Selecting the content of a Report" and "Formatting the Appearance of a Table.")

8.5.2. Using the Choose Ad Hoc Data Wizard

This section gives more detail about the four pages of the Choose Ad Hoc Data wizard:

- Data Page - Choose the fields to make available in the Ad Hoc Editor.

- Filters Page - Define a filter on any field, with the option of prompting for user input, or to compare fields.
- Display Page - Change the order and names of fields that will appear in the Ad Hoc Editor.
- Topic Page - Save your settings as a Domain Topic.

You must first select some fields on the Data page, but the other three pages are optional and can be completed in any order. Click Done at any time to begin designing your report.

8.5.2.1. Data Page

Use this page to choose fields and sets of fields to use in your report or make available in your Domain Topic.

- The left-hand list displays the sets of fields in the Domain. Use the controls in the tree to close or expand each set.
- The list on the right-hand side shows you the items you have selected. You can move a field or set back and forth between the lists by dragging the item, double clicking it, or selecting it and clicking an arrow button.
- If you move an individual field out of a set, it appears in a set of the same name. If you do not want sets, use the settings on the Display page.
- You can change Domains on the drop-down list above. Doing so removes all settings and updates the list of fields.
- Some Domains define sets that are not joined. When you select a field from one set, the unjoined sets are not available.

8.5.2.2. Filters Page

The Filters page provides powerful functionality for designing reports within Analytics. You can define filters to limit data that appears in a report and to define whether a report should prompt users for input on a given filter. They are similar to input controls.

- You can define a filter on a field that you do not plan to use in the report. The filter is still applied and only data that satisfies all defined filters will appear in the final report. For example, you can filter data to select a single country, in which case it does not make sense for the Country field to appear as a row, column, or group.
- Select a field in the left-hand list and click **Create Condition** to begin defining a filter on that field. Alternately, you can click the arrow next to the field name or double-click the field name.
- In the Condition Editor, the comparison operator and value operand depend on the data type of the field and on the actual data that exist for the field. Choose the operator and operand for comparison, then click **OK** to define your filter.
- Text fields have both substring comparison operators such as “starts with” or “contains” and whole string matching such as “equals” or “is one of.” When you select a whole string matching operator, the dialog displays a list of all existing values for the chosen field retrieved in real-time from your database. If there are more than 50 values to display, use the search controls to the left and click the search icon to narrow your list of available values. For multiple value matching, double-click the available values to select them. You can perform multiple searches and select values from each list of results.
- The list of Current Conditions shows all the filters you have defined. When the Prompt? option is selected, the selected filter is available to end users when they run the report. On the report page, users can enter a comparison value for this condition; when the user clicks **Apply** or **OK**, the report preview refreshes with data that match the condition. The value defined here is the default. The condition is available as a prompt even if the filtered field does not appear in the report. For example, the final report might present data for a single country, but the country is chosen by the user. Once defined, filter prompts can be modified in the Ad Hoc Editor.
- The note at the bottom of the current conditions reminds you that data rows must match all conditions. In other words, the overall filter applied to your data is the logical AND of all conditions you have defined.

- Click a row in the list of Current Conditions to make it editable again in the Condition Editor. Click **OK** to save your changes. After selecting a row, you can also click **Delete Condition** to remove it from the list.
- To define a filter that compares two fields of the same data type, select the second field using Ctrl-click, then click **Create Condition**. This button is only enabled when two compatible fields are selected. After you click **OK**, you can edit and delete it as with other conditions. Not that such filters cannot be presented to your users when they run the report.
- The maximum number of items that can be displayed in the list of values in the Condition Editor is configurable.

8.5.2.3. Display Page

Use the Display page to change the default name and order of the fields as they should appear in the list of fields in the Ad Hoc Editor. You can always change the field labels and ordering in the Ad Hoc Editor, but setting them here will make them available for reuse in a Domain Topic.

- To change the order of fields, click once anywhere in a row and use the arrow buttons to move it up, down, to the top or bottom. Fields may be moved only within their set, but sets as a whole may be also be moved.
- To change the display name of a field or set, double-click anywhere in the row and type the new name in the text box. A field name becomes the label for the row, column, or measure that you create from the given field.
- Sets and the fields they contain appears in the list of fields in the Ad Hoc Editor. Sets are not used in reports but can be used to add all their fields at once, thereby speeding up report creation.
- If you do not want to use sets in the Ad Hoc Editor, select the check box Show items as simple list in designer above the table of display names. You can now rename the fields and move them in any order.

8.5.2.4. Topic Page

On the Topic page, you can give your settings a name and a description to be saved as a Domain Topic. Thereafter, you can create different reports from this Domain Topic, using all the same fields, filters, and display names you have defined. You can also edit the Domain Topic to change any of the settings you made in the Choose Ad Hoc Data wizard.

- After you enter a name for your Domain Topic, it is saved when you click Create Report from any page in the wizard.
- By default, Domain Topics are saved in the standard Topics folder. This corresponds to the Ad Hoc Components > Topics location in the repository; Topics and Domain Topics in this folder appear on the Topics tab when you start a report. Do not modify this folder name.
- The description text is associated with the Domain Topic in the repository. It is displayed at the bottom of the Topics tab and when browsing the repository. Enter an informative description that helps users identify this Domain Topic.

8.5.3. Creating Topics from Domains

In some circumstances, it is important to create a Topic based on the Domain and data settings you chose. In other circumstances, creating a Topic is not necessary. The main consideration is how reports based on the Domain will be used.

The following table explains your choices.

If you want to:	Then do this:	More Information
Create a single-use report from the Domain with your settings.	Click Open Report after defining your settings in the Choose Ad Hoc Data wizard, and then format and run your report.	When you click Done, your field selections appear in the Ad Hoc Editor and you can create reports from them, but the settings themselves are not saved.

If you want to:	Then do this:	More Information
Run a report repeatedly as is or with prompting for new input, or make very similar reports in the Ad Hoc Editor.	Click Open Report after defining your settings in the Choose Ad Hoc Data wizard, then format and save the report in the Ad Hoc Editor.	Once you save a report, users can run it repeatedly and be prompted for input each time if you defined filters with prompts. Users with permissions may also open it again in the Ad Hoc Editor.
Reuse your field, filter, and display settings on this Domain to create new reports with different formatting.	Before you click Open Report , save your settings as a Domain Topic on the Topic page of the Choose Ad Hoc Data wizard. Start a new report and choose your Domain Topic from the Topics tab.	Domain Topics are saved in JRXML format in the repository. They appear on the Topics tab when you start a report.
Be able to modify one or more of the settings you made in the Choose Ad Hoc Data wizard.	Before you click Open Report , save your settings as a Domain Topic. Find your Domain Topic in the repository and open it in the Edit Domain Topic dialog.	After you edit and save your changes, you can create reports based on the modified Domain Topic.
Run reports from the same repository data source but with different Domain settings, such as derived fields and different joins to make new fields available in a report.	Edit the Domain and save it with a new name. Then start a new report and choose your data from the new Domain.	Domains provide advanced functionality such as table joins and derived fields based directly on a data source. Domains usually require administrator permissions to create and modify.

8.5.3.1. Access Permissions in Domain Topics

If other users run reports from your Domain Topic and the Domain is configured for security, it is important to consider everyone's access permissions.

You might not have access to all of the fields in the Domain nor to all the data in those fields. There may be fields that can be seen only by other users, and in the fields that you can see, there may be data that is hidden from you. When you save your Domain as a Topic, only the fields that you selected appear in the Domain Topic. When you run a report from the Topic, only the data that you can access in those fields appears. When others run reports from the Topic, they see those of your fields to which they have access and the data in those fields that they can access.

For example, in a Domain, user Tomas can access fields B-C and data rows 1-3; Anita can access fields C-E and data rows 2-5. Tomas uses the Choose Ad Hoc Data wizard and saves a Domain Topic based on the Domain. When Tomas and Anita run reports, they see different combinations of fields and the data in them.

Even though Anita has permission to see more fields, they are not available to her because Tomas did not have access to them when he created the Domain Topic. However, Anita does have permission to see more data than Tomas, so when she opens or runs the report, the field she does see has more rows than when Tomas views the report.

8.5.3.2. Saving Domain Settings as a Topic

To save your settings in the Choose Ad Hoc Data wizard as a Domain Topic:

1. While making selections in the Choose Ad Hoc Data wizard, navigate to the Topic page.
2. Enter a Topic name and description.

Important! Do not change the location folder. The default `/adhoc/topics` folder makes the saved Domain Topic available on the Topics tab when you select Create Ad Hoc Report.

3. If your data selections, filter definitions, and display settings are complete, click **Done**. Otherwise, navigate to the other pages to finish making changes.

Your Domain Topic name and description are preserved, and the Domain Topic is created when you subsequently click **Done** anywhere in the Choose Ad Hoc Data wizard. You can also change the name and description by returning to this Topic page anytime before clicking **Done** in the wizard.

The new Topic appears in the Ad Hoc Components Topics folder. While its repository type is Domain Topic, it appears when the Search page is filtered to show reports.

8.5.3.3. Editing a Domain Topic

Editing a Domain Topic modifies the settings that were defined in the Choose Ad Hoc Data wizard when creating it.

To edit the settings in a Domain Topic:

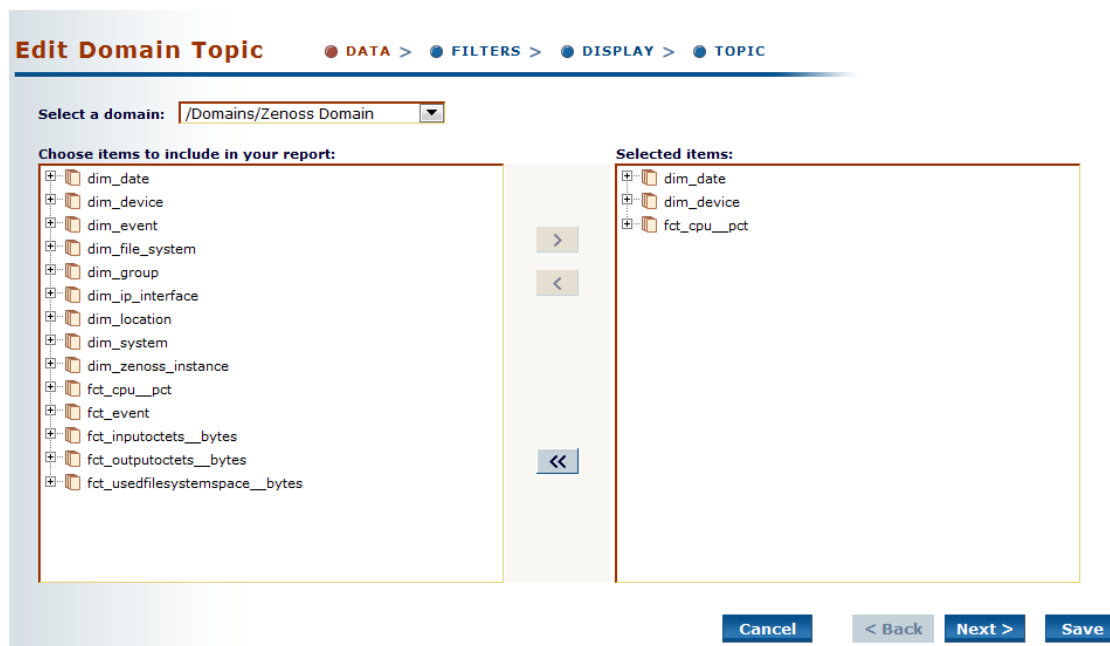
1. From the Analytics interface, select View > Repository, and then search or browse for the Domain Topic you want to modify.

Domain Topics usually are kept in the Ad Hoc Components Topics Folder.

2. Right-click the Domain Topic, and then select Open in Designer from the list of options.

The Domain Topic opens in the Edit Domain Topic Wizard.

Figure 8.8. Edit Domain Topic



3. Follow the guidelines in the section titled "Using the Choose Ad Hoc Data Wizard" to edit the Domain Topic.

Note

In the Select a Domain drop-down menu on the Data page, you can choose a different Domain for the Topic. If you select a different Domain, all the selections that you made with the current Domain are lost when you click **Choose Data**.

4. To save your changes, click **Save** on any page.

Note

Use caution when editing Domain Topics that may have been used to create reports. Users who relied on the Domain Topic might receive unexpected data or errors. It is safer to save your changes as a new Domain Topic.

To save your changes as a new Domain Topic, navigate to the Topic page and enter identifying information for the new Topic, then click **Done**.

8.6. Configuring System-Level Ad Hoc Options

Administrators can control how Analytics handles data for Ad Hoc reports.

8.6.1. Limits on Queries

There is a limit on the number of rows that a report run from the Ad Hoc Editor can contain. There is also a limit on the length of time that a report can take to complete. The limits are defined in configuration parameters set by the administrator:

- **Row Limit** - When the Ad Hoc Editor starts, it sends a query to the current data source to fetch data. Analytics limits the number of rows that it fetches:
 - If the row limit is reached when starting the Ad Hoc Editor, the editor starts but it displays a warning message about the number of rows being returned. If a filter is in effect, you might be able to reduce the number of rows by changing the filter values.
 - If the limit is reached when running an Ad Hoc report, the report stops and a message about the row limit appears. You can redefine the report so that it fetches fewer rows.
- **Time Limit** - When an Ad Hoc report with complicated queries runs, it may take a long time to start returning data. If any one query reaches the time limit before returning data, the query and the report are cancelled. You can redefine the report so that the query completes more quickly, change the limit.

8.6.2. Data Policies

Data policies control whether certain processing occurs in Analytics or in the application storing your data source. These settings apply to reports created or re-saved in the Ad Hoc Editor (that is, changing the settings does not retroactively affect existing reports).

By default, Analytics relies on the database to group, sort, and aggregate data returned by queries based on Domains when this is feasible. It also limits the data it retrieves to the fields that appear in the report. For JDBC data sources, and for non-Domain, non-JDBC data sources, Analytics processes the data in memory by default.

Administrators can configure Analytics to handle Domain-based and JDBC-based Ad Hoc reports differently:

- For Domain-based reports, Analytics can be configured to process the data in memory, instead. Consider factors such as the number of users and report complexity when deciding how Analytics should process data for Ad Hoc reports based on Domains data sources.
- For JDBC-based reports, Analytics can be configured to push the processing down to the database, as is done by default for Ad Hoc reports based on Domains.

Chapter 9. Creating Domains

A Domain is a meta data layer that provides a business view of the data accessed through a data source. Defined in Analytics, it presents data in business terms appropriate to your audience, and can limit the access to data based on the security permissions of the person running the report.

Creating a Domain is the first step in creating an ad hoc report. (The Ad Hoc reporting feature allows you to create your own reports from the Analytics interface. Refer to the chapter titled "Working with Ad Hoc Reports" for more information.) Domains also are used to create Domain Topics.

This chapter covers the process of creating domains and defining their contents. For instructions on creating Domain Topics and reports based on Domains in the Ad Hoc Editor, see the section titled "Creating a Report from a Domain."

The following sections provide more information about creating and editing domains:

- Introduction to Domains
- Using the Add New Domain Dialog
- Using the Domain Designer
- Editing a Domain

9.1. Introduction to Domains

Production databases typically contain data in tables that are optimized for storage and retrieval. Columns with data relevant to users need to be joined across several tables, filtered by business needs, secured against unauthorized access, and labeled with user-friendly names. In relational databases, a view can perform some of these functions. In Analytics, a Domain performs all these functions and more, such as the localization of report text and permissions based on data values.

A Domain is a virtual view created and stored in Analytics, without modifying the data source. Through a Domain, users see columns that have been joined, filtered and labeled for their business needs, and the data values they can access are limited by security policies. Administrators define Domains, and users create reports based on the Domains using the Ad Hoc Editor.

Domains are similar to Topics in that they are used in the Ad Hoc Editor as a basis for designing reports. They are different from Topics in several significant ways:

- Domains can be created directly through the Analytics interface.
- Domain creators can write SQL queries, joins, filter expressions, and security files to specify exactly what data can be accessed, as well as labels and local bundles to specify how the data appears.
- When designing a report based on a Domain, Analytics can optimize database access to allow editing of reports that access huge datasets.
- Report creators can further filter data and prompt for user input later when running the report. This is a powerful feature of Domains. Administrators create Domains, but when users create a report based on a Domain, they can further choose the data from within the Domain. Users can select a subset of columns to appear in the Ad Hoc Editor and give them custom labels. They can filter the data for each column; and more importantly, create input controls on columns so that report readers can modify the filter values.

This last point is a very powerful feature of Domains. Administrators create Domains, but when users create a report based on a Domain, they have to opportunity to further choose the data from within the Domain. Users can select

a subset of columns to appear in the Ad Hoc Editor and give them custom labels. Users can filter the data for each column, and more importantly create input controls on columns so that report readers can modify the filter values.

Once all these settings have been made, users can save the report for others to run, and also save the settings in a Domain Topic so others can design similar reports.

9.1.1. Domain Use Cases

The use of Domains, Domain Topics, and saved reports in Analytics depends on the complexity of your data and your business needs. Domains and Domain Topics can give users great freedom in designing reports while the security features make sure they do not access inappropriate data. On the other hand, administrators could create very targeted Domains and Domain Topics, then use repository access permissions to make sure users cannot modify them.

In the first case, a Domain could contain dozens of tables in several unjoined sets, perform very little filtering, but define a strong data-level security policy. Users of the Domain might only see a single set of tables according to their security permissions. They could then perform their own filtering, relabel columns for their own needs, and save their settings as a Domain Topic for future reuse in Ad Hoc reports. In this case, there might only be one Domain within the organization but many Domain Topics that users have created for specific needs.

In the other case, Domains can be used to perform complex joins, have formulae for calculating new columns, filter data, and select a small set of columns for specific users or specific types of reports. Perhaps reports also need to be internationalized, so the administrator creates the corresponding locale bundles. In this scenario, there might be many specific Domains, and each would have corresponding locale bundles and a single Domain Topic. Users would have a wide variety of Domain Topics to choose from, each for a specific purpose, but no opportunity to access the Domains to perform their own filtering.

The preceding examples illustrate two extreme cases, but there are many scenarios that combine some degree of both. The number of users in your production environment and their level of proficiency will also determine your general use cases. A small number of users who understand their database might be given administrator privileges to define complex Domains as an extension of the Ad Hoc report design process. On the other hand, with large numbers of users or less database proficiency, administrators will want to create Domains that help the users meet their business goals. The table in the section titled "Creating Topics from Domains" gives individual use cases for Domains, Domain Topics, and reports based on Domains.

9.1.2. Terminology

This chapter refers to columns and fields. Conventionally, database tables are composed of columns, and columns are composed of fields in each row. Some Domain operations refer to columns, others to fields, however the two terms refer to the same concept from different perspectives. For example, a calculated field refers to a field in row that is computed from the field values in other columns. But the effect of a calculated field in every row is to create a new calculated column. Similarly, operations such as joins and filters operate on designated field values in a row, but they are defined by the columns that determine the fields involved.

Within a Domain, columns are called items. Because an item may originate from derived tables or calculated fields, it may not correspond to a single column in the database.

9.1.3. Components of a Domain

A Domain is saved as an object in the repository. Like other repository objects, it has a name, optional description and folder location specified at creation time. A Domain references the following components:

- A JDBC or JNDI data source whose tables and columns will be presented through the Domain. Data sources are selected from previously defined data sources in the repository.
- The Domain design that specifies tables and columns in data source, queries for derived tables, joins between tables, calculated fields, and labels to display the columns.

- Optional locale bundles consisting of localized property files to give labels and descriptions in other languages.
- Optional security file that defines row and column-level access privileges based on users or roles.

The Domain design is either created through a dialog or uploaded from an external XML file. Locale bundles and security files are uploaded from external files. The following sections describe the various dialogs for selecting, creating or uploading the components of a Domain.

9.1.4. Overview of Creating a Domain

In practice, you should plan your Domain's features before you begin the creation process. You should decide on the data source and define it in the Analytics repository. You should know the elements of your design: tables and columns to choose, joins to perform, filters to define, sets you need, and item properties to expose to users. You also need to determine whether the users of your Domain need localized resources, such as translated labels and local formats. And finally, you need to determine a security policy for the data retrieved through the Domain.

There are several dialogs involved in creating a Domain in Analytics:

1. As an administrator user, select Add Resource Domain in the repository to open the Add New Domain dialog.
2. On the Data and Resources page, select a data source, and then enter a name and optional description for your domain.
3. Launch the Domain Designer and step through its tabs to define the sets and items of your Domain.
4. Optionally define localization bundles or data-level security on the Resources page of the Add New Domain dialog.
5. Click Save to save the newly created Domain.

The following sections provide detailed information about using these dialogs to create a Domain.

9.2. Using the Add New Domain Dialog

The dialogs entitled Add New Domain and Edit Domain are identical and are used to define all the components of a Domain. There are two pages:

- Data and Design - Specify the domain's data source, its domain design, and the properties of its repository object.
- Resources - Optionally specify a security file and one or more locale bundles.

You must specify at least a data source, a design or valid design file, and a display name on the Data and Design page before you may navigate to the Resources page or Save a Domain.

- To navigate between pages, click the page name at the top of the dialog or the Next and Back buttons at the bottom.
- The **Cancel** button exits the dialog without saving any of your Domain settings. For new Domains, no Domain is created; when modifying an existing Domain, it remains unchanged.
- The **Save** button validates your Domain components, and then saves the Domain. For more information, see the section titled "Domain Validation."

9.2.1. Data and Design Page

The Data and Design page of the Add New Domain and Edit Domain dialogs lets you specify the mandatory components of a Domain, as well as set its repository object properties. The mandatory components of a Domain are:

- A data source selected from the repository

- A Domain design created interactively with the Domain Designer or uploaded from a file

Figure 9.1. Add Domain - Data and Design Page

9.2.1.1. Data Source Area

When working in this area:

- You can use only data sources that are in the repository and to which your user has access privileges.
- Click **Browse** to open a dialog that lets you choose from among all the data sources in your organization. If your data source supports schemas, such as an Oracle RDBMS, you will select schemas when you launch the Domain Designer.
- Domains can have only JDBC-based data sources.
- If you installed the sample data, there are sample data sources in Organization Data Sources and Organization Analysis Components Analysis Data Sources. Both contain JDBC data sources that can be used for a Domain.
- The text field shows the repository path of the chosen data source.

9.2.1.2. Properties Area

When working in this area:

- For new Domains, enter a display name and specify a location in the repository by typing or clicking **Browse**. The default location is the last folder selected in the repository.
- The location for Domains in the sample data is the Organization Domains folder, but you may use any location. Analytics locates Domains by their repository object type, not by their location.
- The description field is optional but recommended. The description is displayed with the Domain in the repository and when users are selecting a Domain to create an Ad Hoc report.

9.2.1.3. Design Area

When working in this area:

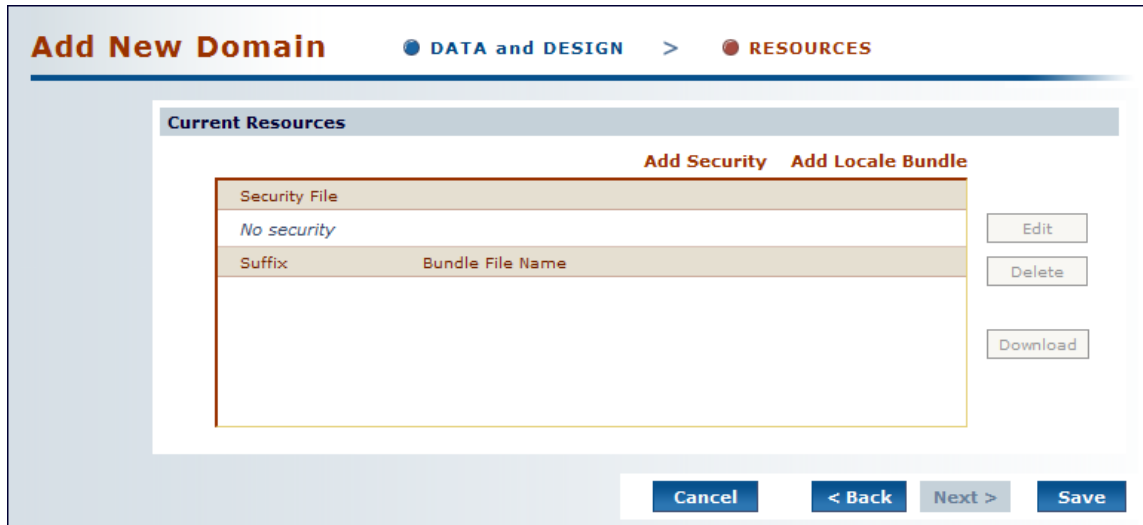
- On the Edit tab, click Launch Domain Designer to open the interactive Domain Designer dialog.

- On the Upload tab, click in the text field or Browse to specify a design file to upload.
- After using the Domain Designer or uploading a design file, both tabs indicate that a design has been created. You can then edit your design in the Domain Designer again or upload another file to replace the current design.

9.2.2. Resources Page

The Resources page of the Add New Domain and Edit Domain dialogs lets you upload security and locale bundle files for your Domain. You can use this page to add, replace, or remove any previously uploaded files for your Domain.

Figure 9.2. Add Domain - Resources Page



9.2.3. Domain Validation

The validation of a Domain ensures that all of its components are consistent among each other. The Add New Domain dialog checks the syntax of files when they are uploaded, but overall consistency must be checked when saving the new or edited Domain.

When you click **Save**, the Add New Domain or Edit Domain dialog performs the following validation:

1. Verify that the tables and columns of the Domain design exist in the data source.
2. In each defined set, verify that all items originate in the same join tree.
3. Verify that all items reference existing columns.
4. Verify that derived tables have valid SQL queries.
5. If a security file has been uploaded, verify that all items and sets in the security file exist in the Domain design.

If validation fails, the Add New Domain dialog remains open and a message appears to help you correct the error. Make the necessary changes to your settings and save again. If the settings are in the uploaded files, you need to edit the files and upload them again.

9.3. Using the Domain Designer

The Domain Designer is an interactive dialog for specifying all the settings in the design of a Domain. Along the top of the Domain Designer are tabs for various aspects of the design:






- Tables – Select all tables whose columns you wish to use in the Domain, either directly or indirectly.
- Derived Tables – Enter queries whose results appear as derived tables available in the Domain.
- Joins – Define inner and outer joins between all the tables and derived tables.
- Calculated Fields – Enter expressions whose results appear as calculated fields.
- Filters – Specify conditions on field values to limit the data accessed through the Domain.
- Display – Give display properties to various tables and columns to expose them sets and items in the Domain.

Once you have selected the tables on the first tab, you may jump to any tab, but you cannot save your settings until you make a selection on the Display tab.

- To navigate between tabs, click the tab name at the top of the dialog or the **Next** and **Back** buttons at the bottom.
- The **Cancel** button exits the dialog without saving any of your design settings. For new Domains, no design is saved; when modifying an existing design, it remains unchanged.
- The **Done** button validates your design then saves it in the Domain, replacing any current design.

9.3.1. Designer Tool Bar

The tool bar buttons operate on the Domain design in its current state, regardless of which tab is selected.

-  **Check Design** - Validates the domain.
-  **Export Design to XML** - Exports the Domain design in its current state to an XML file. Use this feature if you wish to edit your Domain design in an external editor, for example to duplicate settings with copy-paste or to enter a complex formula. Exporting the XML from the Domain Designer avoids having to write it from scratch.
-  **Export Bundle Stub** - Exports a Java properties file for the labels and descriptions of sets and items in your design. Use this button to create a template for your locale bundles after you have defined your sets and items on the Display tab. It gives you the option of automatically generating keys based on the set and item IDs.
-  **Cancel** - Closes the dialog without saving your Domain design and returns to the Data and Design page. Same as clicking **Cancel** at the bottom of the dialog.
-  **Done** - Saves your current selections as the Domain design and returns to the Data and Design page. Same as clicking **Done** at the bottom of the dialog.

9.3.2. Tables Tab

The Tables tab presents all the tables and columns found in the chosen data source so you can select the ones to appear in the Domain. For data bases that support schemas, such as an Oracle RDBMS, you will be prompted to choose one or more schemas, and the Tables tab will present all tables and columns found in those schemas.

Typically, you select the tables that need to be joined, therefore you must understand the logical design of tables in your data source. Additionally, select all tables you wish to reference in your Domain design, even if their columns will not appear directly in the Domain. For example, select the tables containing columns that you will use in a derived table or calculated field.

- Click the tree icons beside the table names to inspect the columns of a table. However, only whole tables may be selected on this tab; column-level selections are made on the Display tab.
- Double-click or drag a table name in the Choose tables list to move it to the Selected tables list. Alternatively, single-click a table name and use the arrow button between the list.
- To remove a table you do not want, double-click it or drag it out of the Selected tables list. You can also click the double-arrow button to clear the list of Selected tables.
- The Inspect new tables and automatically generate option creates joins only if your database has been configured with referential constraints (foreign keys). Otherwise, selecting it has no effect. If applicable, the generated joins will appear on the Join tab.
- The Data Source field at the top shows the name of the data source for your Domain.
- Click Manage to change the chosen data source.

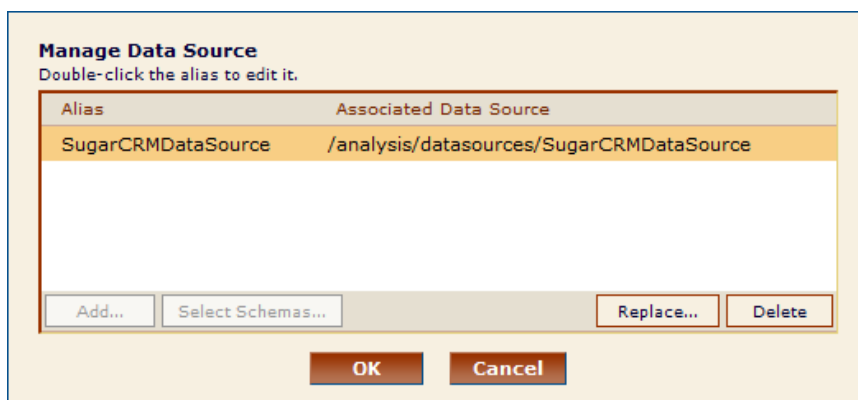
The Tables tab does not detect changes to your database tables and columns in real-time. To update a Domain after making changes to your database structure, click Done to close the Domain Designer, then launch it again:

- New tables and new columns appear in left-hand list; new columns appear under their table name. Because the Tables tab selects entire tables, you must add all new columns together by moving their table name to the right-hand list.
- For deleted tables and columns that were selected in the Domain, you are prompted to remove them from the right-hand list. When you accept, the Domain Designer removes those columns or tables from the Tables tab. However, if you had selected the dropped columns for display, you must manually remove them from the Display tab, otherwise the Domain will have a validation error.

9.3.3. Manage Data Source

Domains assign aliases to data sources. The Manage Data Source dialog on the Tables tab lets you set aliases, as well as change the data source. The default alias for a data source is the display name of its repository object.

Figure 9.3. Manage Data Source Dialog



- You can edit the alias by double-clicking the alias name and typing a new alias. Usually this is not necessary because the data source alias is not displayed to Domain users.
- If your data source supports schemas, such as an Oracle RDBMS, click **Select Schemas** to choose among all available schemas in the data source. All the tables in the schemas you choose will appear in the left-hand list on the Tables tab.

- The **Replace** button lets you select a different data source to be associated with the alias. Use this button to replace the data source with an equivalent one. For example, if your database changes servers, you need to create a new data source object and use it to replace the previous one in your Domain.
- To change the data source and alias, click **Delete** to remove the current one, and then click **Add** to open a dialog to select an available data source from the repository.
- Click **OK** to apply your changes, or click **Cancel** to return to the Tables Tab without changing the data source.

Note

When you change the data source, all previous settings in the Domain Designer that do not conform to the new data source will be lost.

9.3.4. Derived Tables Tab



A derived table in a Domain is defined by a custom query and a selection of the columns in the result. Because Domains are based on JDBC data sources, the query is written in SQL. The result of an SQL query is a table whose structure and contents are determined by the clauses in the query. For example, the WHERE clause may contain conditions that determine the rows of the derived table. Once a derived table is defined, the columns selected from the result are available for use in the Domain design.


- Click **New** to begin defining a derived table. Type a name for your table in the Query ID field.
- Enter a valid SQL (JDBC) query in the Query field. Your query may refer to any table or column available in the data source shown. The syntax for a valid SQL query does not include a closing semi-colon (;).
- Expand the tree icons of the tables in the left-hand list to view column names you may wish to use, but do not click them.
- When your query is complete, click **Run** to test it and see the list of columns in the result.
- By default, all columns in the result are selected. Use control-click to change the selection. If you only want a few columns out of many, it is easier to specify the column names in the SELECT clause of the query.
- Click **Create** to add the derived table with the current selection of columns to the left-hand list of tables and columns. A distinctive icon identifies it as a derived table.

9.3.5. Joins Tab




Joins create associations between tables so that their rows may be presented together in the same report. Multiple joins associate columns across many tables to create powerful data visualizations when used in reports. The number of tables and joins in your Domain depends on your business needs.

To create a join between two tables, each must have a column with the same meaning. For example, a table with data for support cases has a column for the assigned engineer user ID that can be joined with the table of user data that has a user ID column. Analytics supports the four most common join types, all based on equality between values in each column:


-  **Join Inner** - The result will contain only rows where the values in the chosen columns are equal.
-  **Join Left Outer** - The result will contain all the rows of the left-hand table, paired with a row of the right-hand table when the values in the chosen columns are equal and with blanks otherwise.

-  **Join Right Outer** - The result will contain all the rows of the right-hand table, paired with a row of the left-hand table when the values in the chosen columns are equal, and with blanks otherwise.
- **Full Outer Join** - This join type cannot be created directly, it can only be selected by editing an existing join. The result will contain all rows from both tables, paired when the joined columns are equal, and filled with blanks otherwise.

On the Joins tab, the list of selected and derived tables is duplicated in two lists, one to choose the left table, the other to choose the right table. Expand a table in each list, select a column in each table with the same logical meaning and compatible formats, then click one of the join icons. The new join appears below in the list of joins. In some cases, you may need to duplicate a table in order to join it several times without creating a circular join, or in order to join it to itself. You can also duplicate a table so it may be joined with different tables for different uses.

-  **Copy Table** - Copies the selected table and gives it a name with sequential numbering. The copy appears in both lists.
-  **Change Table ID** - Lets you edit the name of the selected table. The new name becomes the ID of the table throughout the Domain, and is updated everywhere it appears in the Domain Designer.
-  **Delete Selected Table** - Removes the table from both lists. If the deleted table was the only instance of a table, removing it on the Joins tab also removes it from the list of selected tables on the Tables tab.

Below the lists of left and right tables, the All Joins tab summarizes the joins you have defined. If you have many joins in your domain, use the Joins on Selected Table tab to see only those defined on the table currently selected in the left-hand list. When you select a join on either tab, you can change the join type by selecting another type in the drop-down list. This list include the Full Outer join type that cannot be created directly.

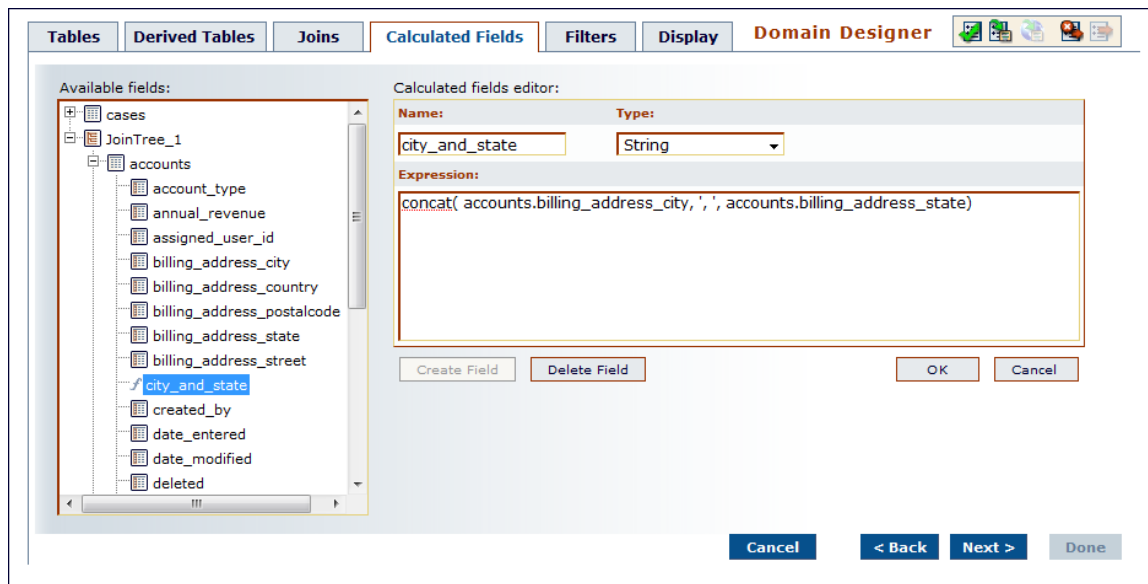
-  **Delete Join** - Removes the selected join definition from the list of joins and from the Domain design.

The result of the Joins tab is a number of join trees that appear on subsequent tabs. For example, if you join tables A and B, B and C, then join tables D and E, there will be two join trees. Columns of table A and table C may appear in the same report because their tables belong to the same join tree. Tables A and D are said to be unjoined; their columns may not be compared or appear in the same report. Tables that are not joined appear individually along with the join trees.

9.3.6. Calculated Fields Tab

A calculated field is defined by an expression that computes a value based on the values of other columns. In order for the values to be coherent, all columns that appear in the expression for a calculated field must be from the same join tree.

- Click **Create Field** to define a calculated field.
- Enter a short name for your calculated field. This name becomes the ID of the field in the Domain; you can later give it a descriptive label and full description.
- Select a data type for your calculated field. The expression you write must compute a value of this type. For most expressions this is also the data type of the columns you wish to use in the expression. Therefore, you need to be familiar with the data types of the various columns in your data source.
- Enter the expression to compute the value of your calculated field. The expression uses the Domain Expression Language fully described in the section titled "DomEL Syntax." To insert a reference to the value of another column, expand the join-tree to find its table and double-click the column name. The column name appears in the expression at the cursor, qualified by its table name. Do not insert column reference from unjoined trees, because the calculated fields editor does not validate expressions as they are written.

Figure 9.4. Domain Designer - Calculated Fields Editor

- Click **OK** to save your new calculated field. The Domain Designer validates the expression and warns you of any errors at this time. If there are errors, use the indications of the error message to help correct the expression. To clear the calculated field editor without saving, click **Cancel**.
- Once validated, a calculated field appears in the table or join tree whose columns are used in the expression. Calculated fields have a distinctive icon for easy recognition in the list of columns.
- An expression that does not use any columns will have a constant value. For example, you might create an integer field named Count that has the value 1 and later has a default summary function to count all occurrences. Constant fields are independent of join trees and automatically appear in a set called Constants.
- Calculated fields may be used to compute other calculated fields. Double-click the calculated field name to insert a reference to it into an expression.
- To view, edit, or delete the definition of a calculated field, cancel any input in the calculated field editor, then click the name of the field in the left-hand list.

9.3.7. Filters Tab

A filter on one or more columns reduces data that is not needed or not wanted in reports based on the Domain. For example, financial reports for the current fiscal year may need data from the previous fiscal year for comparison, but nothing earlier. It is always good practice to filter out irrelevant data to reduce the size of query results and processing time within Analytics.

Also, reports based directly on the Domain can define their own filters. Putting often-used filters in the Domain design avoids the need for each user to define them independently and also reduces the chance for errors.

You can define a filter on a column that you do not plan to expose in the Domain. The filter will still be active and only data that satisfies all defined filters will appear to report users. For example, you may filter data to select a single country, in which case it does not make sense for the column to appear in a report. However, you should clearly document such data restrictions in the description of your Domain, so that users understand what data is accessible through the Domain.

- Select a column in the left-hand list and then click **Create Condition** to begin defining a filter on that column. Alternatively, you can click the arrow next to the column name or double-click the column name.
- To define a filter that compares two columns of the same data type, select the second column with Control-click, and then click **Create Condition**. This button is enabled only when two columns of the same type are selected.
- In the Condition Editor, the choice of comparison operators depends on the data type of the column. For example, string types offer a choice of string search operators and date types have a choice of time comparison operators. The filter value depends on the data type and the comparison operator. For example, if you select a date column with the “is between” operator, the condition editor will display two calendar icons for specifying a date range.
- Text columns have both substring comparison operators such as “starts with” or “contains” and whole string matching such as “equals” or “is one of.” When you select a whole string matching operator, the dialog displays a list of all existing values for the chosen column retrieved in real-time from your database. If there are more than 50 values to display, use the search controls to the left and click to narrow your list of available values. For multiple value matching, double-click the available values to select them. You can perform multiple searches and select values from each list of results.
- After choosing the comparison operator and filter value, click **OK** to define your filter. To clear the condition editor without saving a filter, click **Cancel**.
- The list of Current Conditions shows all the filters you have defined. The note at the bottom of the current conditions reminds you that “data rows much match all conditions.” In other words, the overall filter applied to your data is the logical AND of all conditions you have defined.
- Click a row of Current Conditions to edit it again in the Condition Editor. Click **OK** to save your changes. After selecting a row, you may also click **Delete Condition** to remove it from the list.




9.3.8. Display Tab


The Display tab lets you specify which columns and calculated fields are exposed through the Domain and how they will appear. Typically, only columns that are useful in reports or for further filtering should be selected, and unneeded columns should be omitted to simplify the report creation interface. On the Display tab, you also define display properties for each chosen column, such as a label and description, to further help report creators.

A column along with its display properties is called an item in the Domain. A set is a grouping of items independent of the tables in which the columns originate, however, all items in a set must correspond to columns in the same join tree. Sets are optional; you can create a list of items outside of any sets.


The Join Tree list at the left of the tab contains all the join trees and unjoined tables, including any calculated fields in their defined locations. The list of Sets and Items to the right shows the sets and items that will appear to report creators. You may include any combination of joined and unjoined items in the list of Sets and Items, but when creating a report, users will only be able to include items that originate in the same join tree.

If you only want to display a few of the columns from the join tree, start by creating sets in the right-hand list. Then add items chosen among the columns of the left-hand join tree.




-  Add New Set - Creates a new set or subset in the list of sets and items. If no set is selected it creates a top-level set; if a set or item within a set is selected, it creates a subset.
-  Change Label - Lets you edit the name of the selected set or item. The label name is also updated in the list of properties.
-  Add Selected - Make a selection in the join tree and click this icon: tables are added as sets, and columns are added as items outside of any set. You can also drag-and-drop tables and columns from the join tree to the list of sets and items.

-  Add to Set - Select a destination set in the right-hand list, then make a selection in the join tree and click this icon: tables are added as subsets, and columns are added as items within the destination set. You can also expand sets in the right-hand list and drag-and-drop tables and columns to the destination set.

If you want to display all or most of your tables and columns in the join tree, double-click the join tree name or drag-and-drop it to the list of sets and items. All tables are created as sets, and the columns they contain are created as items within the sets. Then remove any unwanted sets or items individually from the right-hand list as follows:

-  Delete (above Sets and Items) - (above Sets and Items) – Removes the selected set or item from the left-hand list. You can also double-click the object or drag-and-drop it to a blank area of the join tree list. Removed items automatically reappear in the join tree or table where they originated in the left-hand list.

The left-hand list also offers an alternative view mode: click **Table List** just below the icons. The table list displays the list of unjoined tables and derived tables, as they appear on the Join tab. It also activates the following icons:

-  Add All (active only if the right-hand list is empty) - Creates all the tables as sets and all the columns they contain as items in the right-hand list. You can then remove any unwanted sets or items from the right-hand list as described above.
-  Change Table ID - Lets you edit the name of the selected table. The new name becomes the ID of the table throughout the Domain, and is updated everywhere it appears in the Domain Designer.
-  Delete Selected (above Table List) - (above Table List) – Removes the currently selected table from the left-hand list and from the Domain altogether. The table is removed from the Joins tab, and any joins with this table are deleted. Use carefully.

After your list of sets, subsets, and items is defined, refine your Domain display by reordering, renaming, and providing descriptions for them:

- To change the order of sets or items, select the objects and use the arrow buttons to the right of the list to move them up or down. You can reorder items within a set and sets within the list, but items and subsets may not be moved into other sets. To achieve this, remove the item or subset and add it again in the new set. Subsets always appear after the items in a set.
- To rename a set or item, select it then click the **Change Label** icon or double-click its label in the properties table.
- To add a description to a set or item, select it, and then double-click its description in the properties table. Set and item descriptions appear as tooltips in the Ad Hoc Editor and help report creators understand their purpose.

9.3.9. Properties Table

The properties table is located on the Display tab to the right of the list of sets and items. It lets you view and edit the display properties for the currently selected set or item. Display properties determine how the sets and items will appear to users of the Domain. When the left-hand list displays the Table List, the properties of tables and columns can also be viewed.

The following properties are available, depending on the selected object:

- Table – The name of the table and the name of the data source (view-only)
- Field (column) – The name of its table in the data source and its Java data type (view-only)
- Set – The label, description, and internationalization keys (view and edit)
- Item – The label, description, internationalization keys, and default format and summary properties (view and edit)


Table and field properties appear only when the Table List view is enabled. All table and field properties are read-only. For sets and items, properties shown in red are read-only. All others you can double-click to edit, or click ... to edit the description.


The label and description help report creators understand the data represented by the set or item. The internationalization keys are the property names of internationalized strings in locale bundles. The data format and summary properties determine how the item will appear by default in a report. The following table describes each of the properties in detail.

Property	Appears On	Description
ID	Table, Field Set, Item	An identifier used within the Domain. Table and field IDs are based on the names in the data source, but you may change the ID of a table as long as it remains unique. Set and item IDs are a separate namespace in which each ID must be unique, although based on table and field IDs by default. When creating a Domain, you may change any set or item ID as long as it is unique among all set and item IDs. When editing a Domain that has been used to create Topics and reports, you should not change any IDs. For more information, see the section titled "Maintaining Referential Integrity."
Data Source	Table, Field	Alias of the data source for the selected field or table.
Source Table	Table, Field	Name of the selected table or of the field's table in the data source. Does not change when the ID property of a table is modified.
Type	Field	Java type of the selected field.
Label	Set, Item	User-friendly name displayed in the Choose Ad Hoc Data wizard and the Ad Hoc Editor.
Description	Set, Item	User-friendly description displayed as tooltip on the label in the Ad Hoc Editor. The description helps the report creator understand the data represented by this set or item.
Label ID	Set, Item	Internationalization key for the label property; locale bundles associate this key with the localized text for the set or item label.
Descr. ID	Set, Item	Internationalization key for the description property; locale bundles associate this key with the localized text for the set or item description.
Source	Item	References the Domain names of the table and field associated with this item; the syntax is <code>Domain_jointree.Domain_table.datasourcesource_field</code> .
Data Format	Item	Default data format for the item when used in a report. The data format is a numerical or date format, for example (\$1,234) or 01/01/2009. Items with string values do not have a data format.
Summary	Item	Default summary function for the item when used in a report. Numerical items have functions such as sum or average, all others let you count distinct values or count all values.

Labels and descriptions may contain any characters, but the ID property value must be alphanumeric and not start with a digit.

If you want to edit the properties for a large number of sets and items, the following buttons in the tool bar can help:

-  Export Design to XML - Lets you edit the properties of your sets and items within an XML file. For a large number of sets and items, this could be faster than using the properties table on the Display tab.

-  Export Bundle Stub - Generates the internationalization keys if needed and saves them to a file that serves as a template for your locale bundles. Make sure your set and item IDs are finalized because they are used to generate the keys. The generated keys are added to your Domain design and appear in the table of properties.

9.3.10. Domain Design Validation

The Domain Designer must ensure that all tables and columns are consistent with the data source. Validation is important because the Domain design may include derived table queries and calculated field expressions entered by the user.

Validation occurs at the following times:

- When opening the Domain Designer. This will detect any inconsistencies in Domain designs from uploaded files.
- In certain cases, when navigating from tab to tab. This will detect problems on the tab where they occur.
- If you change the data source.
- When you export the design file.
- When you click the Validate tool bar icon.
- When you click Done to exit the Domain Designer.

Validation performs the following steps:

1. Verify that all tables and columns exist in the data source.
2. In each defined set, verify that all items originate in the same join tree.
3. Verify that all items reference existing columns.
4. Verify that derived tables have valid SQL queries.

Unless you clicked the Validate tool bar icon, there is no messages when validation succeeds. However, a message appears when validation fails to help you correct the error.

9.4. Editing a Domain

You can edit a Domain by changing, adding to, and deleting its components.

Important! Be extremely careful when editing Domains that might have been used for reports and Domain Topics. A Domain specifies the data source for the Domain Topics and reports that are based on the Domain. They might fail if the underlying Domain is edited; they will certainly fail if the underlying Domain is deleted. Before proceeding, read the section titled "Maintaining Referential Integrity."

To edit a Domain:

1. Log in to the Analytics interface as an administrator, and then select View > Repository.
2. Browse to or search the repository for the Domain. (Typically, you do this by setting the Domain filter.)
3. Right-click the Domain, and then select Edit from the list of options.

The Domain appears in the Edit Domain dialog.

4. Make changes, and then click **Save** to update the Domain.

Note

After modifying a Domain, you must clear your Ad Hoc cache of all queries based on the Domain. This will remove any data that was based on the old instance of your Domain and avoid inconsistencies in new reports.

9.4.1. Maintaining Referential Integrity

When editing an existing Domain, it is up to the user to maintain the referential integrity between the items in the Domain and any Domain Topics or reports that have been created from the Domain. Referential integrity means that at the time a Domain Topic or report is opened or run, all the items that it references are still defined in the Domain. If you modify a Domain by removing sets or items, you must be sure that no Domain Topics or reports are based on that item. Even if the underlying tables and columns still exist in the database, the sets and items referenced in the Domain must still exist in the Domain.

Domain items are identified by their IDs and the IDs of the sets in which they are located. Changing the ID of an item or moving it to a different set will also make it unavailable to any Topics and reports that referenced the ID. To change the name of an item or set, edit its label property, not its ID. Moving an item must be treated as deleting it from its current location so it can be added elsewhere.

Any report that references a deleted item will no longer run, nor can it be opened in the Ad Hoc Editor again. A Domain Topic that references a deleted item will cause errors when used in the Ad Hoc Editor. You can, however, open a Domain Topic for editing and remove references to deleted items. However, that only allows new reports to use the Domain Topic, it does not fix the broken reports based on the items deleted from the Domain Topic.

The granularity of referential integrity is at the individual set and item level. This means that changes to sets and items that are not used in a given report or Topic do not affect the report or Topic. For example, if you delete an item used by Topic A but not Topic B, then Topic A will fail and reports based on Topic A that included the item will fail. But Topic B and its reports are unaffected, as are any reports based on Topic A that did not include the deleted item.

Chapter 10. Administering the Repository

The repository stores content files, data sources, data types, images, saved reports, and other Analytics resources.

Administering the Analytics repository includes these tasks:

- Creating folders and organizing repository objects
- Controlling access to objects in the repository using roles, users, and object-level permissions
- Managing references to data sources, images, fonts, and other resources on which reports rely

Read this chapter for more information about:

- Managing folders and resources
- Access control
- Referencing resources in the repository

10.1. Managing Folders and Resources

Administrators and users with the proper permissions can create, modify, move, and delete folders and resources within the repository.

The specific roles and permissions of the user determine what actions are available. The following sections explain what permissions are necessary to perform each action. When an action creates a resource, the section also gives the initial permissions assigned to the resource.

Within the repository and search results view, all actions on folders and resources are accessible through their context menu. Right-click on the folder or resource name to see the context menu for that object. Use any combination of search, folder browsing, or filters to display the resources you want to operate on.

10.1.1. Creating a Folder

Any user with write permission on a folder can create new folders within it. By default, a new folder and its future contents inherit all permissions from the parent folder.

1. Log on as a user with write permission to the parent folder.
2. Click **View Repository**, and then locate the folder in which you want to create a folder.
3. Right-click the parent folder, and then select Add Folder from the list of options.
4. Enter a folder name and optional description in the dialog that appears, and then click **Add**.

The folder is created in the repository.

10.1.2. Adding Resources to the Repository

Only administrators can add resources to the repository. Regular users, even those with write permission on a folder, cannot create resources. By default, a new resource inherits all permissions from its parent folder.

1. Log on as an administrator.

2. Click **View Repository**, and then locate the folder in which you want to create a resource.
3. Right-click the parent folder, and then select Add Resource from the list of options. Then, select the type of resource to add.

The most common object types, such as reports and images, appear on the Add Resource menu. Less common object types are listed on the Other menu at the bottom of the Add Resource menu.

4. Enter the information in the resource creation wizard specific to the selected resource.

Some resources are based on uploaded files, others on information you enter in the dialog. All wizards include fields to specify an object name, display name, and description in the repository. The object name is a unique name within the folder. The display name and description appear to users in the repository.

If you are creating a data source, click Test Connection to validate it. If the test fails, review the values that you specified in the other fields and then test the data source again.

5. Click **Save** to create the resource and add it to the repository.

10.1.3. Renaming a Folder or Resource

Any user with write permission on a folder or resource can change its display name or description.

Note

You cannot change the name of an organization's main folder through the repository. The name of the main folder is always the display name of the organization. To change the name of the main folder, change the display name of the organization, as described in section titled "Managing Organizations."

1. Log on as a user with write permission to the folder or resource.
2. Click **View Repository**, and then locate the folder or resource you want to change.
3. Right-click the folder or resource, and then select Properties from the list of options.
4. In the dialog that appears, enter a new display name or description.

Users cannot modify the resource type or resource ID; these are internal fields displayed for information only.

5. Click **OK**.

10.1.4. Viewing a Report or Dashboard

Only reports and dashboards support a viewing mode, which is equivalent to running the report or dashboard with current data. Any user with read permission can view (run) a report or dashboard.

1. Log on as a user with read permission to the report or dashboard.
2. Click **View Repository** and browse or search for the report or dashboard you want to view.
3. Click the name of the report or dashboard to run it. Alternatively, right-click the report or dashboard, and then select View from the list of options.

The report or dashboard begins to run, and an activity monitor appears until the report or dashboard contents are displayed. The longer or more complex the report or dashboard is, the more time it takes to display.

If you selected a long report by accident and it takes too long to display, click **Cancel** in the activity monitor.

To return to the repository or search results page, click **View Repository**.

10.1.5. Modifying an Ad Hoc Report or Dashboard

Users with write permissions can open Ad Hoc reports and dashboards in the Ad Hoc Editor or dashboard designer, respectively. They can then modify the report or dashboard, overwrite the original, or save it as a new one. Ad Hoc reports and dashboards are the only resources whose content can be modified by end-users.

1. Log on as a user with write permission to the report or dashboard.
2. Click View Repository, and then browse or search for the report or dashboard you want to modify.
3. Right-click the Ad Hoc report or dashboard, and then select Open in Designer from the list of options.

Ad Hoc reports open in the Ad Hoc Editor and dashboards open in the dashboard designer.

4. When you save your Ad Hoc report or dashboard, it overwrites the original one. In the Ad Hoc Editor, choose Save As to save a new report and preserve the original.
5. To return to the repository or search results page after saving your work, click **View Repository**.

10.1.6. Editing a Resource

Editing a resource invokes the same wizard used to define the resource when it was created. When editing the definition of a resource, you can reload a file, for example, or change a setting. Editing a resource can also be useful to view its settings, even if you do not want to modify them.

Resources that have been created through the Add Resource menu can only be edited by administrators. Administrators can also edit the report unit of an Ad Hoc report. Regular users, even those with write permission on a resource, cannot edit the definition of repository resources.

1. Log on as an administrator.
2. Click **View Repository**, and then browse or search for the resource you want to edit.
3. Right-click the resource, and then select Edit from the context menu.
4. Use the resource-specific dialog to view or modify its definition or properties.
5. Click **OK** or **Save** to save changes.

10.1.7. Copying Folders or Resources

Any user with read permission can copy folders and resources in the repository, as long as the user has write permission on the target folder where the copy is made. When copying a folder, all the resources and folders it contains are copied recursively as well.

Note

When copying a folder or resource, permissions are not preserved and the new copy inherits all permissions from its parent folder. Administrators must explicitly set permissions again after making copies of folders or resources.

Copying is available through drag-and-drop, as well as from context menus using the copy-paste model. Folders must be copied individually, but resources can be copied in bulk.

To move folders or resources by drag-and-drop:

1. Log on as a user with write permission to the folder or resource.
2. Click **View Repository**, and then locate the folder or resource you want to move.

3. Select the folder or resource. For resources such as reports whose names are active links, check the box beside the name, or click anywhere in the row to select it.

To move more than one resource at a time, select them with Ctrl-Click, or check the box next to each resource.

4. Drag the highlighted folder or resource, and then drop them on a folder for which you have write permission.

While dragging, the mouse pointer changes to show you are moving a folder, resource, or multiple resources. It also changes to show that certain folders such as the top organization folder are never targets. The mouse pointer does not indicate whether you have write permission for the target folder. If you drop objects on a folder for which you do not have write permission, the objects are not moved.

To cancel a move operation, drop the object in a blank area of the repository.

5. If you have write permission to the target folder, the folder or resources are moved, and the repository display updates to reflect the new location according to your current search and filter settings.

To move folders or resources using context menus:

1. Log on as a user with write permission to the folder or resource.
2. Click **View Repository**, and then locate the folder or resources you want to move.
3. Right-click the folder or a single resource, and then select Cut from the list of options.

To move more than one resource at a time, select them with Ctrl-Click, or check the box next to each resource name, and then click Cut above the list of resources.

In both cases, the mouse pointer changes to indicate you have initiated a move operation.

4. Right-click the destination folder, and then select Paste from the list of options.

If Paste does not appear in the options list, then you do not have write permission there. To cancel a move operation, right-click any folder, and then select Cancel. The selected folder or resources will not be cut.

5. After selecting Paste, the folder or resources are moved, and the repository display updates to reflect the new location according to your current search and filter settings.

10.1.8. Deleting Folders or Resources

Users with delete permission on a folder or resource can delete those objects from the repository. In order to delete a folder, the user must have delete permission on all the resources and folders it contains, recursively.

Note

The repository keeps track of which resources are referenced by other resources, and does not allow you to delete them while they are still being referenced. For example an input type that is used by a report or a properties file that is used by a Domain cannot be deleted as long as the report or Domain still reference them.

To find the resources that reference the one you want to delete, you need to look at each report, Ad Hoc topic, or Domain that you suspect of referencing it.

Folders must be deleted individually, but resources can be deleted in bulk:

1. Log on as a user with delete permission to the folder or resource.
2. Click **View Repository**, and then locate the folder or resources you want to delete.
3. Right-click the folder or resource, and then select Delete from the list of options.

To delete more than one resource at a time, select resources with Ctrl-Click and then check the box next to each resource name. Then click **Delete** above the list of resources.

4. Confirm that you want to delete the folder or resource. (There is no undo operation.)

10.2. Access Control

Access control ensures that people using Analytics can access only the data they are allowed to see. Analytics provides access control through an integrated security framework that includes:

- Authentication – Restricts access to identified users and protects that access with passwords. Defines roles for grouping users and assigning permissions.
- Authorization – Controls access to repository objects, pages, and menus based on users and roles.
- Data level security – Defines row and column level permissions to access your data. Row and column level permissions can be defined and enforced in Domains.

10.2.1. Authentication Overview

The first part of security is to define user accounts and secure them with passwords. Users must log in with their user ID and password so that they have an identity within Analytics. Analytics stores user definitions, including encrypted passwords, in a private database. Administrators create, modify, and delete user accounts through the administrator pages.

Analytics also implements roles that can be assigned to any number of users. Roles let administrators create groups or classes of users that are granted certain permissions. For example, administrator privileges are granted by the `ROLE_ADMINISTRATOR` role. A user can belong to any number of roles and receive the privileges from each of them. Analytics stores role definitions in its private database.

Analytics relies on the open source Acegi security framework, which has configurable options for:

- External authentication services such as LDAP (used by Microsoft Active Directory and Novell eDirectory)
- Single sign-on using JA-SIG's Central Authentication Service (CAS)
- Java Authentication and Authorization Service (JAAS)
- Container security (Tomcat, Jetty)
- SiteMinder
- Anonymous user access (disabled by default)

Analytics also supports these encryption and authentication standards:

- HTTP Basic
- HTTP Digest
- X509

The Acegi framework is readily extensible to integrate with custom and commercial authentication services and transports.

Authentication occurs by default through the Web user interface. Analytics can automatically synchronize with an external authentication service. The external users do not need to be created manually in Analytics first. Both users and roles are created automatically in Analytics from their definitions in an external authentication service.

10.2.2. Authorization Overview

With a user's identify and roles established, Analytics controls the user's access in several ways:

Menu Options and Pages	The menus that appear in Analytics depend on the user's roles. For example, only users with the administrator role can see the Manage menu and access the administrator pages. By modifying Analytics' configuration, you can modify access to menus, menu items, and individual pages.
Organization Scope	Users belong to an organization and are restricted to seeing resources in that organization. Organizations have their own administrators, but they see only the users, roles, and resources from their organization. When Analytics is configured with multiple organizations, they are effectively isolated from each other.
Resource Permissions	Administrators can define access permissions on every folder and resource in the repository. Permissions are defined for every role and every user, or they can be left undefined so they are inherited from the parent folder. For example, users may have read-write access to a folder where they create reports, but the administrator can also create standard reports in the same folder that are set to read-only. Permissions are enforced when accessing any resource either directly through the repository interface, indirectly when called from a report, or programatically through the Web services.
Administrator Privileges	Analytics distinguishes between reading or writing a resource in the repository and viewing or editing the internal definition of a resource. For security purposes, granting a user read or write permission on a resource does not allow viewing or editing the resource definition. For example, users need read permission on a data source to run reports that use it, but they cannot view the data source definition, which includes a database password. Only administrators can create, view, or edit the definition of a resource and its internal components.
Data-Level Security	Data-level security defines what data can be retrieved and viewed in a report, based on the user name and roles of the user who runs the report. For example, a management report could allow any user to see the management hierarchy, managers to see the salary information for their direct employees, and only human resource managers to see all salary information.

10.2.3. Permissions

Permissions on folders and resources determine what users see in the repository and the actions they are allowed to perform. In the following table, the actions granted for each permission include all of the actions granted for permissions above it, except for the No Access permission. The actions granted for each permission strictly exclude all of the actions granted for permissions below it.

Permission	Actions Granted on Repository Folders and Resources
No Access	Users cannot see or access the folder or resource, either directly or indirectly.
Read Only	See the folder or resource in any Analytics dialog, see the properties of a folder or resource, copy a folder and all of its readable contents, copy resources individually or in bulk, view (run) a report or dashboard, run a report in the background, schedule a report to run later
Read + Delete	Cut (move) a folder and all of its contents, delete a folder and all of its contents, cut (move) resources individually or in bulk, delete resources individually or in bulk
Read + Write + Delete	Paste into a folder (copy or cut), save a new Ad Hoc report or dashboard in a folder, save the output of a scheduled report in a folder, rename a folder or resource and change its description string, open an Ad Hoc report in the Ad Hoc editor or a dashboard in the designer, modify and overwrite an existing Ad Hoc report or dashboard

Administer	Set the permissions (by role and by user) or a folder or resource
Administer and ROLE_ADMINISTRATOR	Add (create) a resource in a folder, edit a resource (for example, the components of a report unit or a Domain)

Permissions apply to access when browsing or searching the repository, as well as any dialog that accesses the repository, such as when browsing folders to save a report.

Notes:

- Copying does not preserve the permissions on an object. Users may copy a read-only object, paste it into a read-write folder, and then edit the object.
- Copying and cutting (moving) actions can only be completed if the user has Read + Write + Delete access to a folder in which to paste the objects.
- Cutting, deleting, and setting permissions on folders is allowed only if the user has the same permission on all folder contents. Cutting and deleting resources in bulk is allowed only if the user has at least Read + Delete permission on all selected resources.
- Deleting a resource or the contents of a folder is allowed only if no other resources rely on them.

10.2.3.1. Inheriting Permissions

According to the permission architecture, there is a permission setting for every user and role on every folder and resource in the repository. To simplify the definition of permissions, Analytics supports the inheritance of permissions from the parent folder of a folder or resource. If no permission is explicitly defined for a user or role on a given folder or resource, the user or role has the same access permission that is defined on the parent folder. When a permission is defined explicitly, that permission is enforced, regardless of those on the parent folder.

Using this mechanism, administrators can manage large hierarchies of content and keep them secure. When the administrator sets a permission explicitly, that permission for a given user or role is inherited recursively by all of the folder's contents and subfolders, unless they have an explicit definition of their own. Permissions that are assigned on an organization's top folder are inherited across the entire organization. Permissions that are set on the root folder or Organizations folder by the system admin are inherited across multiple organizations.

For example, the system admin can make all organizations read-only by default to ordinary users, and each organization admin can make specific folders writable so that users can store their reports and output.

10.2.3.2. Cumulative Permissions

Because permissions can be assigned to both users and roles, a user belonging to one or more roles may have multiple permissions defined or inherited on any given folder or resource. In fact, every permission must be defined on the root, even if it has the default value of No Access, and therefore every role- and user-based permission on every folder and resource has a setting through inheritance. Therefore, for every folder or resource, every user has a their own user-based permission and the permission assigned to the ROLE_USER role.

Permissions in Analytics are strictly cumulative, meaning that the least restrictive among the set of all permission applies. Even if a more restrictive permission, such as No Access, is set explicitly, the less restrictive permission such as Read-Only applies, regardless of whether it is inherited or set explicitly.

10.2.3.3. Administrator Permissions

The Analytics authorization architecture distinguishes between administrators and all other users. Administrators are defined as users with either ROLE_SUPERUSER, ROLE_ADMINISTRATOR, or both. By design, system administrators with the ROLE_SUPERUSER always have irrevocable Administer access to the entire repository, including to the contents of every organization. The system administrator cannot modify the permissions for ROLE_SUPERUSER,

to prevent being locked out or unable to administer some resource. Therefore, the system admin can set permissions for all other users, on any folder or resource, and in any organization if necessary. In particular, the system administrator can modify permissions for `ROLE_ADMINISTRATOR`, for example to share some resources across all organizations by making them read-only to everyone, including the organization admins.

Organization admins are organization users with the `ROLE_ADMINISTRATOR` role. By default, organization admins have the Administer permissions to everything within their organization, except what the system admin has changed to a lesser permission. However, organization admins cannot change the permissions granted to `ROLE_ADMINISTRATOR`, to prevent them from overriding the settings of the system admin and from locking themselves out of a folder or resource.

10.2.3.4. Default User Permissions

For all other users, the default permission at the root is No Access and any permissions must be explicitly defined. In practice, the default installation of the repository contains sample data with permissions that allow the sample users to access folders and resources. We recommend you familiarize yourself with the permissions mechanism by viewing and setting permissions in the sample data, as described in the following section.

10.2.4. Assigning Permissions

Administrators can assign permissions to access any folder or resource throughout the repository. Users with the Administer permission on a folder can assign permissions to that folder and any contents that inherit the permission. Users granted Administer permission to a resource can only set the permissions on that specific resource.

To set permissions on a folder or resource in the repository:

1. Log in as a user with administrative privileges.
2. Select View > Repository, and then locate the folder or resource.
3. Right-click the folder or resource, and then select Permissions from the list of options.

The Assign Permissions by Role interface shows the existing role-based permissions in effect for the folder or resource. In systems with multiple organizations, the roles displayed only include those in the scope of your user. In the default single organization, the organization admin cannot see the permission for `ROLE_SUPERUSER`.

Permissions that are inherited are indicated by * (asterisk).

Figure 10.1. Assign Permissions by Role



- For each role you want to grant access, select a value from the Permission Level list of options.

Among the possible values, the permission of the parent folder is also indicated by * (asterisk). Setting the permission to the inherited value is the equivalent of removing an explicit permission.

Note

There are two special cases.

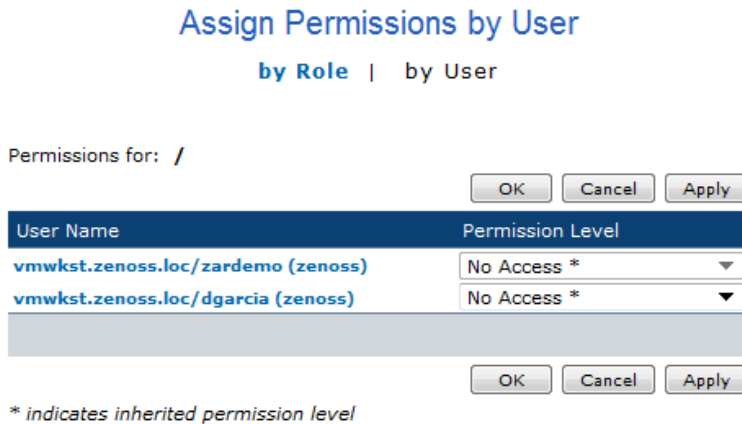
You cannot explicitly set the permission level that is inherited from the parent folder, at least not directly. You need to temporarily change the permission level on the parent folder, then set the explicit permission, and finally set the parent folder's permission back to the original value.

If none of the options in the drop-down are highlighted, the folder and its parent have been independently set to the same value. To reset the permission level so that it once more inherits its parent folder, select a different permissions level and click **Apply**; then select the permission with the asterisk and click **Apply** again.

- Click **Apply** to save your changes before assigning permissions by user. If you have finished, then click **OK** to save and return to the repository view.
- To assign permissions by user, click **by User**.

The Assign Permissions by User interface shows the existing user-based permissions in effect for the folder or resource. In systems with multiple organizations, the users displayed only include those in the scope of your user. In the default single organization, the organization admin cannot see the permission for superuser. Permissions that are inherited are indicated by * (asterisk).

Figure 10.2. Assign Permissions by User



There are no user-based permissions defined by default. In this case, all access permissions have been defined through roles.

- For each user you want to grant access, select a value from the Permission Level list of options.
- Click **Apply** to save your changes and stay on an Assign Permissions page. If you have finished, click **OK** to save and return to the repository.
- Click **Cancel** to return to the repository without changing permissions.

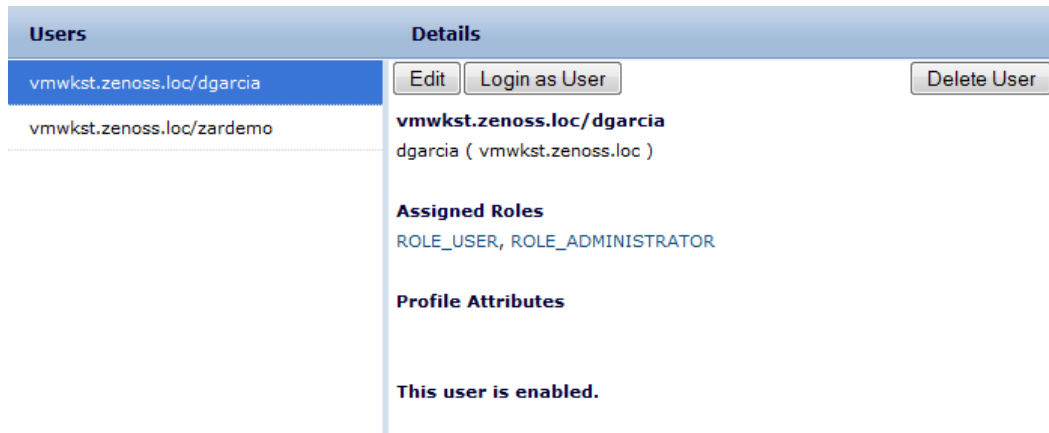
10.2.5. Testing User Permissions

Once you have configured users, roles, and permissions, Zenoss recommends that you test the permissions granted to a few representative users. Testing is also recommended when you add new users, roles, and resources, and when you make any major modifications to your access control configuration.

To test user permissions:

1. Log in as an administrator.
2. Select Manage > Users, and then select the user whose access permissions you want to test.

Figure 10.3. Test User Permissions



3. Click **Login as User**.

The selected user's Home page appears.

4. Select View > Repository, navigate to the desired folder, and verify that Analytics displays the expected folders and resources.

10.3. Referencing Resources in the Repository

All resources in the repository can be referenced by Universal Resource Identifiers (URIs) giving the resource name and the folder path where the resource are located. Because of the hierarchy of organizations, references are relative to the user accessing them. Analytics transforms relative references into actual resource locations in the repository based on the user's organization and the organization's main folder. By default, folder locations are transformed in the following ways:

- For organization members, locations in /public are not transformed, but those within the organization's main folder seen as / (forward slash) by the user are transformed to the actual location, for example /organizations/organization_1.
- For system admins, locations throughout the repository are not transformed. They see the actual repository path names.

This transformation allows some URIs to reference different resources depending on the organization of the user who accesses them. For example, a report may have a public data source and an organization-specific logo as an image. We can set up the report as follows:

- Data source URI specified in the report unit: `/public/CommonDataSource`. Regardless of the user's organization, this always references the same repository resource.
- Logo URI specified in the JRXML: `/images/organizationLogo`. When transformed for each user, the URI will access a location relative to his organization's main folder. The folder and image resource must have the same name in each organization.

If any user in any organization runs this report, and has a logo image in his `/images/organizationLogo` resource, then the report uses the public data source and displays the organization-specific logo.

As shown in the previous example, references found in JRXML resources are transformed before access. In the following cases, references are not transformed in the user context and must work literally:

- The reference to a data source, JRXML, and input controls within a report unit.
- The reference to a data source and OLAP connections in an analysis view.
- The reference to a data source and schemas in an OLAP connection.

Because these references are not transformed at runtime, you must consider these limitations:

- The system admin cannot maintain these three types of resources in an organization level folder, because the organization users would see errors when trying to access the referenced objects. The system admin must log in as a user of the organization and do the maintenance in the context of the organization.
- These three types of resources cannot reference objects in the parent of an organization across organizations. For example, if the data source for a report unit is set by a user in `Org1` into `/dataSources`, a user in another organization cannot access that data source because the reference is not transformed and the literal reference is in another organization.

Chapter 11. Troubleshooting

11.1. Resolving Issues

Refer to this information when resolving Analytics issues.

Table 11.1. Troubleshooting Procedures

Symptom	Diagnostic
The Analytics server does not start or stop.	Run the start or stop script with the DEBUG variable defined. For example: <code>DEBUG=1 service zenoss_analytics start</code>
An event of component type "ANALYTICS" appears in the console.	Follow the directions for resolving the issue in the event.
One of the Zenoss ETL daemons stops.	Restart the daemon with DEBUG level logging (<code>zenetl restart -v10</code> or <code>zenperfetl restart -v10</code>) and send the log file to Zenoss Support.

Appendix A. Using Secure Copy to Transfer Files to Analytics

A.1. Introduction

The ZenETL and ZenPerfETL daemons transfer files to the Analytics server. By default, they use HTTP to transfer files. You can configure the daemons to use Secure Copy (SCP) instead of HTTP.

A.2. Procedure

Follow these steps to configure the ZenETL or ZenPerfETL daemon to use SCP for file transfer:

1. Determine the Linux account under which ZenETL or ZenPerfETL is running. Typically, this is the Zenoss user.
2. Create an SSH key with a blank pass phrase for the account under which the daemon is running.
3. Add the key to an allowed account on the Analytics server.

Note

You can test that SCP is working correctly by trying to manually execute the secure copy of a file.

4. Edit the `zenetl.conf` or `zenperfetl.conf` file and populate these options:

- **scpfolder** - Mandatory. This is the folder on the Analytics server to which you must copy the file. Assuming default installation paths, this line should look like:

```
scpfolder/opt/zenoss_analytics/work/Catalina/localhost/reporting
```

- **scpuser** - Mandatory. This is the user account on the Analytics server with which the SSH key should be associated. Typically, this line should look like:

```
scpuser zenoss
```

- **scpparams** - Optional. Add any SCP parameters that you want included in the secure copy command.

Note

The daemon log file (`zenetl.log` or `zenperfetl.log`) will indicate which method is used to copy files to the Analytics server. If the SCP information is not established, or the SCP fails, then the system will use HTTP transfer.