



Zenoss Analytics Installation and Administration Guide

Release 5.0.x

Zenoss, Inc.

www.zenoss.com

Zenoss Analytics Installation and Administration Guide

Copyright © 2017 Zenoss, Inc. All rights reserved.

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

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

Flash is a registered trademark of Adobe Systems Incorporated.

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

Linux is a registered trademark of Linus Torvalds.

RabbitMQ is a trademark of Pivotal Software, Inc.

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

Sybase is a registered trademark of Sybase, Inc.

Tomcat is a trademark of the Apache Software Foundation.

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

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

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

Part Number: 1100.17.041

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

Contents

About this guide.....	4
Chapter 1: Introduction.....	5
Analytics architecture.....	5
Chapter 2: Installing and configuring Analytics.....	7
Installation considerations.....	7
Required software and conditions.....	7
Installing the OpenJDK.....	9
Setting the JAVA_HOME environment variable.....	9
Installing MariaDB.....	10
Installing the Analytics server.....	11
Installing the user defined function to MariaDB.....	12
Installing the ZenETL ZenPack on the Resource Manager master.....	12
Configuring Analytics.....	14
Removing an Analytics installation.....	23
Chapter 3: Upgrading to Analytics 5.0.x.....	26
Prerequisite tasks.....	26
Exporting the Analytics database.....	26
Dropping the Analytics database.....	27
Saving existing data.....	27
Removing ZenDS (optional).....	27
Installing the OpenJDK.....	27
Setting the JAVA_HOME environment variable.....	28
Installing MariaDB.....	29
Upgrading the Analytics server and Jaspersoft database.....	30
Installing the user defined function to MariaDB.....	31
Importing the reports database.....	31
Installing the ZenETL ZenPack on the Resource Manager master.....	31
Adding aliases.....	32
Starting the Analytics service.....	32
Chapter 4: Working with Analytics.....	33
Starting and stopping the Analytics server.....	33
Logging in to Analytics.....	33
Creating batches.....	35
Configuring the Analytics data retention policy.....	39
Working with the Analytics repository.....	40
Appendix A: Third party licenses.....	57
Appendix B: ERD diagram for Analytics 5.0.x.....	58

About this guide

Zenoss Analytics Installation and Administration Guide provides detailed procedures for installing Zenoss Analytics (Analytics), and for administering the product.

Related publications

Title	Description
<i>Zenoss Analytics Installation and Administration Guide</i>	Describes how to install and administer Analytics.
<i>Zenoss Analytics Release Notes</i>	Describes known issues, fixed issues, and late-breaking information not already provided in the published documentation set.

Additional information and comments

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

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

Introduction

Zenoss Analytics (Analytics) provides enterprise reporting for Resource Manager instances, enabling improved business intelligence and knowledge management. Analytics gives you 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.

Because Jaspersoft provides the back-end technology used to provide the reporting capability, consult the *TIBCO™ JasperReports® Server User Guide, Release 6.0* for detailed information about the user interface of Analytics.

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)
- MariaDB event database (EventDB)
- Centralized performance data storage (HBase cluster)

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.

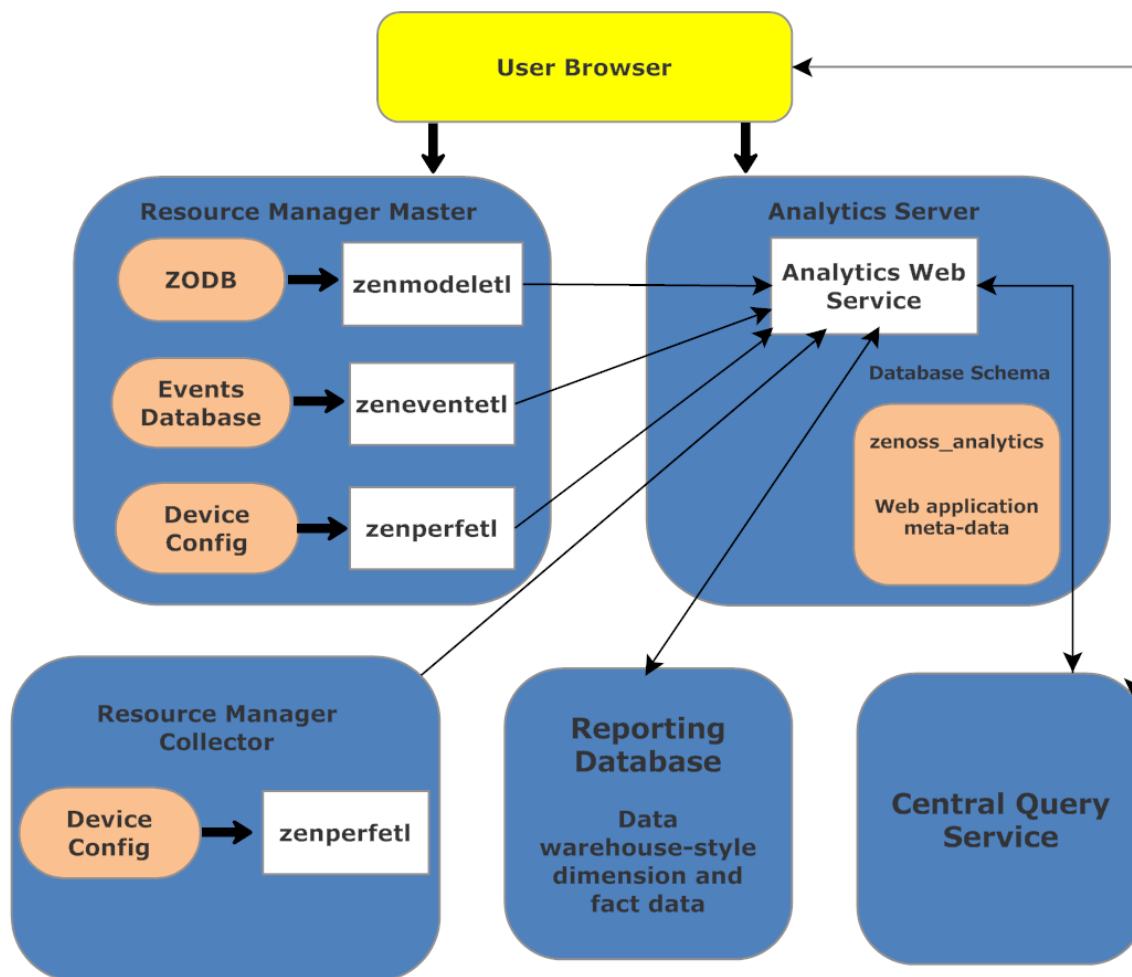
On the Resource Manager server, three daemons run to export data from the Resource Manager server to the Analytics server:

- **zeneventetl** - Extracts event information, which is then sent to the Analytics server for subsequent transform and load.

- **zenmodeletl** - Extracts ZODB model information, which is then sent to the Analytics server for subsequent transform and load.
- **zenperfetl** - Performance data is sent to the Analytics server for subsequent transform and load. This daemon also runs on each collector. Typically, this is run on the Resource Manager server and in the collector instances in the various resource pools.

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

Figure 1: Analytics components



2

Installing and configuring Analytics

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

Note If you are upgrading from Analytics version 4.4.0 to version 5.0.x, skip the installation instructions in this chapter and proceed to [Upgrading to Analytics 5.0.x](#) on page 26.

- [Installation considerations](#)
- [Prerequisite tasks and conditions](#)
- [Installing the OpenJDK](#)
- [Installing MariaDB](#)
- [Installing the Analytics server](#)
- [Installing the ZenETL ZenPack on Resource Manager](#)
- [Configuring Analytics](#)
- [Removing an Analytics installation](#)

Installation considerations

The Analytics system creates the following databases:

- **zenoss_analytics (Analytics database)** - This database stores report design, execution information, and results of the reports, including PDFs. You can also access previous reports here.
- **reporting (etl database)** - This database warehouses all data retrieved from Resource Manager.

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

Required software and conditions

- To run Analytics, you must install the following on the Analytics server:
 - OpenJDK 7, 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 or SE 8. They are not compatible with this version of Analytics.

- Analytics requires that UID 1337 be available. You can verify its availability with the command:

```
getent passwd 1337
```

- 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/>.)
- Ensure that hostname resolves to the loopback interface locally by editing `/etc/hosts` to include an entry for the hostname of the Analytics application server after the listing for `127.0.0.1`.

Compatibility and support

- Analytics 5.0.x is compatible with Resource Manager 5.0.6 and Control Center 1.0.6 and later versions.
- You must use matched versions (first 3 digits of version number) of Analytics and the ZenETL ZenPack. For example, you cannot use ZenETL 4.4.0 with Analytics 5.0.x.
- For new installations, you must be using a 64-bit system running RHEL 7/CentOS 7.
- The `zenperftl` 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.
- Analytics is not compatible with Security-Enhanced Linux (SELinux) in enforcing mode. If SELinux is installed, disable it as follows:

- 1 Determine whether SELinux is installed by running the following command:

```
test -f /etc/selinux/config && grep '^SELINUX=' /etc/selinux/config
```

If the preceding commands return a result, SELinux is installed.

- 2 Set the operating mode to `disabled`, and confirm the setting.

```
EXT=$(date +"%j-%H%M%S")
sudo sed -i.${EXT} -e 's/^SELINUX=.*SELINUX=disabled/g' \
/etc/selinux/config && \
grep '^SELINUX=' /etc/selinux/config
```

- 3 To complete the configuration change, reboot.

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>.

Recommendations

Consider the following recommendations for best performance:

- Install Analytics and its prerequisites on a separate server 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.
- Minimum system requirements for running Analytics are as follows:
 - 8-core processor
 - 8GB RAM
 - 500GB storage

Requirements for open network ports

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

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

From	To	Default Port Numbers
Analytics	Analytics database server	3306
Analytics	Resource Manager	8080
Resource Manager	Analytics	7070
Analytics Web Users	Analytics	7070
zenperfetl (on each remote collector)	Analytics application server	7070

Installing the OpenJDK

Analytics requires the OpenJDK 7 Development package. Analytics cannot be installed when just the JRE is installed.

Before installing the OpenJDK 7 Development package, remove all other versions of Java that are on the system:

```
yum -y remove $(rpm -qa | egrep -i '(jdk|jre|java)')
```

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Install the OpenJDK 7.

```
yum -y install java-1.7.0-openjdk-devel
```

- 3 Execute the following command and choose the OpenJDK 7 selection:

```
alternatives --config java
```

- 4 Verify that the installation succeeded.

```
java -version
```

If the command returns output similar to the following, continue to the next procedure.

```
java version "1.7.0_75"
OpenJDK Runtime Environment (rhel-2.5.4.7.el7_1.x86_64 u75-b13)
OpenJDK 64-Bit Server VM (build 24.75-b04, mixed mode)
```

Setting the JAVA_HOME environment variable

After installing the OpenJDK, you must set `$JAVA_HOME` for the root user.

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Locate and note the path to your OpenJDK directory. Typically, this is `/usr/lib/jvm/java-1.7.0-openjdk-VersionNumber`.
- 3 Navigate to the `/etc/default` directory.

```
echo $JAVA_HOME
```

- 4 Open the `zenoss_analytics` file in an editor.

```
vi zenoss_analytics
```

- 5 Edit the `JAVA_HOME` line in `zenoss_analytics` to point to your OpenJDK location.

```
JAVA_HOME=/usr/lib/jvm/java-1.7.0-openjdk-VersionNumber
```

- 6 Execute the `zenoss_analytics` file.

```
source zenoss_analytics
```

- 7 Test the new path with the following command:

```
echo $JAVA_HOME
```

Installing MariaDB

Before installing MariaDB on systems running Red Hat Enterprise Linux (RHEL), install the Perl Database Interface Module (DBI) with the following command as the root user:

```
yum -y install perl-DBI
```

To install MariaDB 10 on the Analytics server, follow these steps. (The Zenoss DataStore (ZenDS) is not used in Analytics Version 5 and later.)

- 1 Create a custom MariaDB YUM repository. The following example shows information for a CentOS 7 installation. For more information, see <https://downloads.mariadb.org/mariadb/repositories>
 - a As the root user, navigate to the `/etc/yum.repos.d` directory.

```
cd /etc/yum.repos.d
```

- b Create a file called `MariaDB.repo` and enter the following text into the file.

```
# MariaDB 10.0 CentOS repository list
#http://mariadb.org/mariadb/repositories/
[mariadb]
name = MariaDB
baseurl = http://yum.mariadb.org/10.0/centos7-amd64
gpgkey=https://yum.mariadb.org/RPM-GPG-KEY-MariaDB
gpgcheck=1
```

- 2 When the `MariaDB.repo` file is in place, install the MariaDB server:

```
sudo yum -y install mariadb-server
```

Note If you have not accepted the MariaDB GPG key, you will be prompted to do so.

- 3 Add the following configuration settings:

```
echo "[mysqld]" >> /etc/my.cnf
echo "table_open_cache=16K" >> /etc/my.cnf
echo "table_definition_cache=16K" >> /etc/my.cnf
echo "tmp_table_size=2G" >> /etc/my.cnf
echo "max_heap_table_size=2G" >> /etc/my.cnf
echo "join_buffer_size=512K" >> /etc/my.cnf
echo "open_files_limit=200000" >> /etc/my.cnf
echo "tmpdir=/tmp" >> /etc/my.cnf
echo "wait_timeout=86400" >> /etc/my.cnf
echo "innodb_adaptive_hash_index=OFF" >> /etc/my.cnf
echo "innodb_buffer_pool_size=6G" >> /etc/my.cnf
echo "innodb_log_file_size=1892M" >> /etc/my.cnf
echo "innodb_log_buffer_size=128M" >> /etc/my.cnf
```

Note The values for the innodb settings should reflect your system. Set `innodb_buffer_pool_size` to 80 percent of the total memory on a standalone data warehouse server or 50 percent of memory if the Analytics server and database are on the same server. Set `innodb_log_file_size` to 25 percent of the `innodb_buffer_pool_size` value.

- 4 Start the MariaDB server:

```
/etc/init.d/mysql start
```

- 5 Set MariaDB to start on boot:

```
chkconfig mysql on
```

- 6 If the Analytics application and the Analytics database are on the same machine, the installation of MariaDB is complete. However, if you used the recommended configuration of a remote database, you must enable remote access so that the `ana_mariadb` container can connect to it:

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

Installing the Analytics server

Install the Analytics server by using an `.rpm` file on a separate server.

- 1 As the root user, enter the following command. Replace *Version* with the actual RPM version you are installing.

```
rpm -ivh zenoss_analytics-5.0.x-Version.noarch.rpm
```

- 2 Modify the database connection setting in the `/etc/zenoss_analytics/zenoss_analytics.conf` file to point to the desired database server or servers. For best performance, install database servers on a different server than the Analytics server.

```
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.

- 3 Change to the zenoss user:

```
su - zenoss
```

- 4 Install the database schema:

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

- 5 Change to the root user:

```
su - root
```

- 6 Enable chkconfig for the Analytics service:

```
chkconfig zenoss_analytics on
```

- 7 Start the Analytics service:

```
service zenoss_analytics start
```

Installing the user defined function to MariaDB

Before installing the user defined function, ensure that you have installed MariaDB on the *database* server and Analytics on the *application* server.

- 1 As the root user, log into the Analytics application server and locate the following file:

```
/usr/lib64/mysql/plugin/zenoss_percentile.so
```

- 2 Copy `zenoss_percentile.so` to the same directory on the MariaDB database server (`/usr/lib64/mysql/plugin/`).
- 3 On the Analytics application server, change to the zenoss user:

```
su - zenoss
```

- 4 Execute the following script:

```
/opt/zenoss_analytics/bin/setup_zenoss_extensions install
```

Installing the ZenETL ZenPack on the Resource Manager master

Before you install the ZenETL ZenPack (`ZenPacks.zenoss.ZenETL-Version-py2.7.egg`) on the Resource Manager master, install the Analytics server.

- 1 Log in to the Control Center browser interface.

The screenshot shows the Control Center interface. At the top, there's a navigation bar with 'Applications' selected. Below it, the 'Applications' table is displayed. The table has the following data:

Application	Description	Status	Deployment ID	Resource Pool	Public Endpoints	Actions
Internal Services	Internal Services	✓	Internal	N/A	N/A	N/A
Zenoss.resmgr (v5.1.1)	Zenoss Resource Manager	✓	Test	default	https://zenoss5.ip-10-111-23-46.443	▶ Start ■ Stop ⚙ Delete

Below the table, it says 'Last Update: a few seconds ago' and 'Showing 2 Results'. Underneath is the 'Application Templates' section with a table:

Application Template	ID	Description	Actions
Zenoss.resmgr (v5.1.1)	b1f8fb3473d412a9a89555ea411d4b72	Zenoss Resource Manager	⚙ Delete

It also says 'Last Update: a few seconds ago' and 'Showing 1 Result'.

- 2 In the **Applications** table, identify the Resource Manager instance to modify.
- 3 Scroll down to the **Services** table, and then collapse the subservices of the **Infrastructure** and **Zenoss** services.

The screenshot shows the 'Services' table with the following data:

Service	Instances healthy/total	Description	Actions
> Infrastructure			* ▶ Start ■ Stop ⚙ Restart
> Zenoss			▶ Start ■ Stop ⚙ Restart

- 4 Stop the **Zenoss** service, and then verify that its subservices are stopped.
 - a In the **Actions** column of the **Services** table, click **Stop**.
 - b In the **Service** column, click **Zenoss**, and then scroll down to verify all subservices are stopped.
- 5 Create a snapshot.
 - a Log in to the Control Center master host as a user with Control Center CLI privileges.
 - b Create a snapshot.

```
serviced service snapshot Zenoss.resmgr
```

The `serviced` command returns the ID of the new snapshot on completion.

- 6 Restart required services.
 - a Click the **Zenoss** service to expand its subservices.
 - b Restart the following subservices.

- `zeneventserver`
- `Zope`

Note In the **Services** table, the Failing icon (a red circle with an exclamation point) in the **Status** column represents the cumulative result of one or more customized health checks. To view the status of individual health checks, move the pointer over the icon, which displays a pop-up.

- 7 Change directory to the directory in which the ZenPack egg file is located.
For example:

```
cd /tmp/zenpack
```

- 8 Install the ZenPack.

For Resource Manager version 5.0.x:

```
serviced service run zope zenpack install ZenPack-File.egg
```

For Resource Manager version 5.1.x or higher:

```
serviced service run zope zenpack-manager install ZenPack-File.egg
```

Any daemons that a ZenPack provides are packaged in Docker containers and installed as child services of the current instance of Resource Manager.

- 9 In the Control Center browser interface, stop the `zeneventserver` and `zope` services, then start the **Zenoss** service that you stopped previously.

Note You must perform this step exactly in order so that the underlying `zeneventserver` and `zope` services get restarted.

Configuring Analytics

To configure Analytics, complete all required tasks, and optional tasks as needed.

- [Configure the connection to the Analytics server](#)
- [Configure zeneventserver settings](#)
- (Optional) [Configure database credentials](#)
- (Optional) [Configure additional Resource Manager data points to send to Analytics](#)
- (Optional) [Configure SSL communication](#)
- [Configure email services](#)

Configuring the Analytics server connection

Configure Resource Manager to connect to the Analytics server.

- 1 In the Resource Manager browser interface, select **Reports > Configuration**.
The Analytics configuration page appears with the currently scheduled report jobs.
- 2 In the **Internal URLs** area, enter URLs for Resource Manager and Analytics to 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.)

Note If your Resource Manager URLs are not displayed, refresh your browser to clear the cache. Do not proceed until the URLs are displayed correctly.

- 3 In the **External URLs** area, enter the links to 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 In the **Query Service** area, enter the following information and then click **Submit**:
 - 1 Complete URL of the virtual host.

- 2 Access credentials with Manager or higher privileges to the Central Query service of the specified Resource Manager instance. Do not enter the admin account.

For more information about roles and user permissions, see the "Managing Users" chapter in the *Zenoss Resource Manager Administration Guide*.

The configuration page redisplay shows currently scheduled report jobs. If this is a new installation, this screen is empty because no jobs have been created.

Figure 2: Add batches

The screenshot shows the 'Configuration' tab in the Zenoss Resource Manager interface. On the left, there is a navigation pane with 'Jobs' selected. The main area displays a table of scheduled report jobs. The table has columns for Batch ID, Extractor, State, Begin, End, and Attempts. There are four rows of data shown.

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

Configure zeneventserver settings

Resource Manager Version 4.1 and later 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 cannot properly extract event times for those batches.

You can increase the number of days of event history that can be extracted by using the Resource Manager browser interface.

- 1 In the left panel, select **Advanced > Events** .
The Event Configuration page appears.
- 2 Increase the value of Event Time Purge Interval (in days), up to 7 days and click **Save**.
- 3 If you increased the Event Time Purge Interval value, then you must also increase the number of file descriptors for your Resource Manager instance. Perform the following steps on the Analytics database server:
 - a As the root user, view the current open files:

```
lsof | wc -l
```

- b Check the max open files by entering the following command:

```
cat /proc/sys/fs/file-max
ulimit -Hn
ulimit -Sn
```

- c Change the max open files temp by entering the following command:

```
sysctl -w fs.file-max=$((cat /proc/sys/fs/file-max) + 25000)
```

- d Add the following lines to the `/etc/security/limits.conf` file:

```
root hard nofile 999999
root soft nofile 999999
zenoss hard nofile 999999
zenoss soft nofile 999999
```

Configuring database credentials

When Analytics is installed, it configures a JDBC data source for the reporting database. As a best practice, test the database connection.

If you have a split application and database installation, you must change the database settings in Analytics. This task requires Manager or ZenManager privileges. For more information about roles and user permissions, see the "Managing Users" chapter in the *Zenoss Resource Manager Administration Guide*.

- 1 Log in to the Resource Manager browser interface.
- 2 Select **Reports > Advanced**.
- 3 On the Analytics **Getting Started** page, select **View list > Repository**.
- 4 In the Folders area, expand the following folders: **root > Organizations > ZenossData Sources**.
- 5 In the **Repository** area, select **Zenoss Reporting > Edit**.
The Edit Data Source window is displayed.

Figure 3: Edit data source window

zenoss Analytics Library View Manage Create superuser | Log Out

Edit Data Source: Zenoss Reporting

Set Data Source Type and Properties

First, select the type of data source you wish to add, then enter the required property values.

Type:

JDBC Driver:

Host (required):

Port (required):

Database (required):

URL (required):

Hint: jdbc:postgresql://localhost:5432/mydb

User Name:

Password:

Time Zone:

Hint: Do not change the time zone setting unless you know the database timestamp data is incorrect.

- 6 After editing, click **Test connection** to verify the changes, and then click **Save**.

Export data point data

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

In order for performance data to be exported for use in the Analytics system, you must create aliases to the data points on the monitoring templates. For more information about templates, see *Zenoss Resource Manager Administration Guide*.

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.

Note Because metric names are used to name the database tables, periods are not allowed in metric names.

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 a reverse Polish notation (RPN) expression. Aliases are used to normalize raw data point values across data point types. While it may seem like a good idea to collect lots of data by adding many aliases, you must balance the need for data with the load that it would generate on the system.

Note To avoid issues with percentile and projections, limit alias names to 31 characters. Do not use periods in alias names.

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.

You can add aliases to a data point in the following ways:

- [Running the `chkaliases` script](#)
- [Manually modifying the data points on the Resource Manager](#)

Running the `chkaliases` script

To add or remove aliases, run the `chkaliases.zendmd` script from the command line.

The script is contained in the ZenETL ZenPack. It processes the `aliases.dsv` file, which contains alias definitions. Each useful line of the `aliases.dsv` file should contain an alias definition consisting of the following values that are pipe-delimited:

```
device class|template|datapoint|alias|rpn
```

The output of the script produces several codes. Codes in the first column indicate whether the script is adding (+), modifying the RPN (R), or removing (-) an alias. Codes in the second column indicate an inability to find a device class (c), template (t), datapoint (d), or alias (a), or that the alias should not be removed because its defined RPN does not match (r). The (a) and (r) codes only occur during a remove operation.

The initial `aliases.dsv` file contains alias definitions for the following, expressed as percentages:

- Device Level CPU usage (`cpu__pct`)
- Memory utilization (`mem__pct`)
- IP interface component input (`in__pct`)
- IP interface component output (`out__pct`)
- File system utilization (`fs__pct`)

- 1 Log in to the Resource Manager server as the zenoss user:

```
su - zenoss
```

- 2 Navigate to the directory where the script resides by entering the following command. Replace *Version* with the actual version number.

```
cd $ZENHOME/ZenPacks/ZenPacks.zenoss.ZenETL-Version-py2.7.egg/  
ZenPacks/zenoss/ZenETL/etc
```

- 3 To see what would be changed, run the script using the following command:

```
zendmd --script=../bin/chkaliases.zendmd [add | remove]
```

- 4 Review the `aliases.dsv` file (located in the `/etc` directory under the ZenETL ZenPack), and modify it as needed. Re-run the script until the changes are satisfactory.

- 5 When the output is satisfactory, re-run the script using the `--commit` option:

```
zendmd --script=../bin/chkaliases.zendmd [add | remove] --commit
```

On the `/Server/Windows` monitoring template, open the `ProcessorTotalProcessorTime` data point in the `ProcessorTotalProcessorTime` data source. An alias "cpu__pct" is defined with no RPN formula. In this case, the raw data point represents the desired quantity ("Percentage of Consumed CPU Time"). The alias, `cpu__pct`, allows you 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.

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. In Analytics, alias names are mapped to the table `meta_metric` table.

Note In previous versions, an alias mapped to a `fct_[alias]` table; this mapping is now deprecated.

Performance data is logically mapped to tables `raw_v2_%` and then to `hourly_%` and `daily_%`. These tables have a column named `metric_key`. The table `meta_metric` provides the lookup for a specific performance metric.

The following query provides all performance data for the `cpu__pct` metric:

```
SELECT *  
FROM raw_v2_cpu__pct  
limit 5;
```

The underlying data is stored in tables `raw_%`, `hourly_%`, and `daily_%`. Do not use these tables directly. However, if you query against the `raw_v2_%` table, be aware that these tables are very large. Ensure that you include the device key, the component key, and the fact time stamp in your query whenever possible to fully leverage the available index. If the metric is for a device (no component included), the component key is 0.

For event data points, all events are brought over to Analytics and stored in tables `fct_event` and `fct_event_time`.

Manually modifying data points on Resource Manager

To export data point data from a Resource Manager monitoring template, you can manually modify each data point.

- 1 From the Resource Manager interface navigation bar, select **Advanced > Monitoring Templates > desired monitoring template**.
- 2 In the Data Sources area, expand the data source to see the associated data points.
- 3 Double-click the desired data point.
The Edit Data Point dialog appears.

Figure 4: Edit Data Point

- 4 In the bottom section, enter an alias for the data point, click **Add**, and then save.

After saving the data point with an alias, the data point is included in the next extraction batch if. To verify that the data is being brought over, use one of the following methods:

- View the "Meta-Data Management: Datapoint lineage" report provided with Analytics.

Note If you are using RPN formulas for data points, this report cannot be used.

- Log in to the Analytics database and execute the following query:

```
select distinct metric_name from reporting.meta_metric order by
metric_name;
```

After the successful completion of the zenperftl batch for a collector that has devices monitored using the monitoring template that you chose, your alias appears in the returned list.

Configuring SSL communication

When Resource Manager is secured behind SSL, Analytics supports the use of SSL for all types of communication.

Securing Resource Manager is a prerequisite for securing Analytics. For more information about securing Resource Manager, see the *Zenoss Resource Manager Administration Guide*.

To configure secure communication for the Analytics system, provide URLs for the following addresses:

- Internal Analytics server address
 - Resource Manager uses this address to call the web service on the Analytics server to pull the status of batches and scheduled batches.
 - The `zenmodeletl` and `zeneventetl` daemons in Resource Manager use this address to call web services on the Analytics server to work with batches and upload data files by HTTP put.
 - The `zenperfetl` daemon on each collector uses this address to call web services on the Analytics server to work with batches and upload data files by HTTP put.
- Internal Resource Manager server address

The Analytics server uses this address to phone home to Resource Manager and authenticate and authorize a user.

- External Analytics server address
 - When you log into the Resource Manager browser interface, it uses this address to send the credentials to the Analytics server. The password is passed in plain text if you are not using SSL.
 - In the Resource Manager browser interface, when you select **Reports > Advanced**, you are redirected to this address.
- External Resource Manager server address

The Analytics server uses this address to make links in reports that you can click through to a device in Resource Manager.

Installing Apache and SSL

On the Analytics server, install Apache and Mod SSL, and use a self-signed SSL certificate. You can purchase a certificate signed by a third-party certificate authority or generate your own certificate.

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Install Apache.

```
yum install httpd
```

- 3 Install Mod SSL for Apache.

```
yum install mod_ssl
```

Note If the installation, check for an earlier version of `open_ssl` on the system that is being used by other packages. If this is the case, perform a `yum upgrade` to update the rest of the system and to bring in the correct version of `open_ssl`.

- 4 Add Apache to automatically start at the server start.

```
chkconfig --levels 235 httpd on
```

- 5 Start Apache.

```
service httpd start
```

- 6 Check your configuration by visiting both `http://yourserver.com/` and `https://yourserver.com/`.

Removing "Pragma" headers

If you support Internet Explorer 8 (IE8), you must remove the "Pragma" statements from the header of the HTTP files.

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Navigate to the following Apache configuration folder.

```
cd /etc/httpd/conf
```

- 3 Make a backup of the `httpd.conf` file.

```
cp httpd.conf original_httpd.conf_original
```

- 4 Open `httpd.conf` in an editor.

```
vi httpd.conf
```

- 5 Add the following line near the top of the file.

```
Header unset Pragma
```

- 6 Save the file and exit the editor.
- 7 Restart Apache.

```
service httpd restart
```

Configuring Apache proxy and rewrite

Perform the following steps to edit the `ssl.conf` file to add an internal proxy rule for Apache to proxy any request to the Analytics server and to turn on the rewrite engine:

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Navigate to the following Apache configuration folder.

```
cd /etc/httpd/conf.d
```

- 3 Make a backup of the `ssl.conf` file.

```
cp ssl.conf original_ssl.conf_original
```

- 4 Open `ssl.conf` in an editor.

```
vi ssl.conf
```

- The last line of the file should be the closing tag `</VirtualHost>` which is the closing to the `<VirtualHost_default_:443>` tag. Add the following text just above the closing `</VirtualHost>` tag:

```
#Internal proxy rules instructing Apache to proxy any request to the
#Analytics server and data warehouse on 7070
ProxyPass /reports http://127.0.0.1:7070/reports
ProxyPassReverse /reports http://127.0.0.1:7070/reports
ProxyPass /etl http://127.0.0.1:7070/etl
ProxyPassReverse /etl http://127.0.0.1:7070/etl
#Turn on the RewriteEngine
RewriteEngine On
#Redirect any just / over to /reports
RewriteRule ^/+$ https://%{SERVER_NAME}:443/reports/ [R]
```

- Save and close the `ssl.conf` file.
- Restart Apache.

```
service httpd restart
```

Lock Tomcat to localhost only

Only perform the following steps to lock down the Apache Tomcat that Analytics uses at the software level. Alternatively, simply close port 7070 via firewall configuration.

- Log in to the Analytics server as `root`, or as a user with superuser privileges.
- Navigate to the following configuration folder:

```
cd /opt/zenoss_analytics/conf
```

- Make a backup of the `server.xml` file.

```
cp server.xml original_server.xml_original
```

- Open `server.xml` in an editor.

```
vi server.xml
```

- Locate the following section in the file:

```
<Connector port="7070" protocol="HTTP/1.1"
connectionTimeout="20000"
redirectPort="8443"/>
```

- Add in `address="127.0.0.1"` so that the section looks like the following:

```
<Connector port="7070" address="127.0.0.1" protocol="HTTP/1.1"
connectionTimeout="20000"
redirectPort="8443"/>
```

- Save and close the file.
- Restart Tomcat to pick up the changes.

```
service zenoss_analytics stop
service zenoss_analytics start
```

Modifications to Resource Manager

Perform the following steps on the Resource Manager server to set the necessary HTTPS URLs.

- 1 Ensure that you have enabled HTTPS support for your Resource Manager instance. For more information, see the "Enabling HTTPS Support" section in the *Zenoss Resource Manager Administration Guide*.
- 2 Log in to the Resource Manager web interface as a user with Administrator privileges.
- 3 Navigate to **Reports > Configuration**.

Note If you have previously configured the 5 URLs for Analytics, you will see an error message at the top of the screen indicating that the system cannot connect to the Analytics server. This is expected behavior. Click the gear icon in the lower-left corner to open the configuration page again.

- 4 Modify the following URLs with the values indicated:
 - a Internal Analytics: HTTPS URL, including port (443), of the Analytics server.
 - b Internal Zenoss: HTTPS URL, including port (443), of the Resource Manager server.
 - c External Analytics: HTTPS URL, including port (443), of the Analytics server.
 - d External Zenoss: HTTPS URL, including port (443), of the Resource Manager server.
 - e Query Service: HTTPS URL, including port (443), of the Central Query service for this Resource Manager instance.
- 5 Submit the form to save the settings.
- 6 Restart the `zenmodeletl`, `zeneventetl`, and `zenperfetl` daemons on the Resource Manager server so that they receive the new settings.

Configuring email services on the Analytics server

To send the output of scheduled reports as email attachments, complete the following steps to configure email services.

- 1 On the Analytics server, edit the `/opt/zenoss_analytics/webapps/reports/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
```

Removing an Analytics installation

To remove an installation of Analytics that was installed by RPM on a separate server:

- 1 Remove the ZenETL ZenPack from the Resource Manager master:
 - a Log in to the Control Center browser interface.
 - b Stop Resource Manager by clicking **Stop** in the Actions column for `Zenoss.resmgr`.
 - c In the Stop Services dialog, click **Stop Services**.
 - d Display the child services of Resource Manager to ensure they have stopped.

In the **Application** column of the **Applications** table, click `Zenoss.resmgr`, then scroll down to the **Services** table. Stopped services have a grey circle icon in the **Status** column.

- e Create a snapshot of the Resource Manager service by logging in to the Control Center master host as a user with `serviced` CLI privileges.

Execute the following command to create a snapshot.

```
serviced service snapshot Zenoss.resmgr
```

The snapshot ID is displayed upon completion.

- f Restart the following required services.
 - ZooKeeper
 - The modelling and event database service or services:
 - Resource Manager 5.0.0: `MariaDB`
 - Resource Manager 5.0.1 or more recent: `mariadb-events` and `mariadb-model`
 - `memcached`
 - `RabbitMQ`
 - `redis`
 - `zencatalogservice`
 - `zeneventserver`
 - `Zope`

Note In the **Services** table, the Failing icon (a red circle with an exclamation point) in the **Status** column represents the cumulative result of one or more customized health checks. To view the status of individual health checks, move the pointer over the icon, which displays a pop-up.

When the Failing icon is present, a service is unable to support the normal operations of Resource Manager. For this task, the `Zope` health checks includes failing health checks of `zproxy_answering`, which does not affect this procedure.

- g Remove the `ZenPacks.zenoss.ZenETL` ZenPack by executing the following command.

For Resource Manager version 5.0.x:

```
serviced service run zope zenpack uninstall ZenPacks.zenoss.ZenETL
```

For Resource Manager version 5.1.x or higher:

```
serviced service run zope zenpack-manager uninstall
ZenPacks.zenoss.ZenETL
```

- 2 Remove the `zenperfetl`, `zeneventetl`, and `zenmodeletl` services.

```
serviced service rm zenperfetl
serviced service rm zeneventetl
serviced service rm zenmodeletl
```

- 3 In the Control Center browser interface, restart Resource Manager.
- 4 Log in to the Analytics server and change to the root user:

```
su - root
```


5 Stop the Analytics server:

```
service zenoss_analytics stop
```

6 Remove the Analytics program files:

```
rpm -e zenoss_analytics-version-build
```

This command removes the 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 (this will erase all your data), enter the following command as the root user:

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

Upgrading to Analytics 5.0.x

This chapter provides instructions for upgrading Analytics version 4.4.0 to version 5.0.x. The Resource Manager will not be upgraded as part of this process. If Resource Manager is at version 4.x, then it will remain at its 4.x version after the upgrade to Analytics version 5.0.x.

Prerequisite tasks

Before upgrading to Analytics 5.0, ensure that hostname resolves to the loopback interface locally by editing `/etc/hosts` to include an entry for the hostname of the Analytics application server after the listing for `127.0.0.1`.

- 1 Download the following from the <https://support.zenoss.com> site:
 - a RPM file for Analytics 5.0.x, `zenoss_analytics-5.0.x.Version.noarch.rpm`
 - b EGG file for ZenETL 5.0.x ZenPack, `ZenPacks.zenoss.ZenETL-Version-py2.7.egg`
- 2 Wait for all batches in the staging state to complete.
- 3 As the `zenoss` user, stop the `zenmodeletl` and `zeneventetl` daemons on each attached Resource Manager server:

```
zenmodeletl stop
zeneventetl stop
```

- 4 As the `zenoss` user, stop the `zenperfetl` daemons on each attached remote collector and on the Resource Manager server:

```
zenperfetl stop
```

Exporting the Analytics database

Analytics uses JasperReports Server to run, display, schedule, and manage reports against the Analytics data warehouse. The JasperReports Server stores the report information in its Analytics small database. Export this database as part of the upgrade process.

- 1 On the Analytics server as the `root` user, navigate to the `jasperserver` scripts folder:

```
cd /opt/zenoss_analytics/jasperserver/buildomatic/
```

- 2 Run the export script:

Note The name of the backup file must be `backup.zip`.

```
./js-export.sh --everything --output-zip=/path_to_safe_place/
backup.zip
```

3 Stop Analytics:

```
service zenoss_analytics stop
```

Dropping the Analytics database

After you have exported the Analytics database, complete the following steps to drop it.

1 Log in as the zenoss user:

```
su - zenoss
```

2 Enter the following commands:

```
zends -uroot
drop database zenoss_analytics;
```

Saving existing data

If you want to save your existing data in the Analytics database, you must note the location of your `/datadir` directory so that it can be added in the `my.cnf` file when you install MariaDB. If you do not want to save your existing data, skip this procedure.

1 On the Analytics server as `root`, or as a user with superuser privileges, navigate to the `zends.cnf` file in the following folder.

```
cd /opt/zends/etc/
```

2 Open the `zends.cnf` file and note the value of the `datadir` directory. You will need this value when you change configuration settings while installing MariaDB.

Removing ZenDS (optional)

Analytics Version 5 does not use ZenDS. Use the following procedure to remove it without deleting data.

1 Log in to the Analytics server as `root`, or as a user with superuser privileges.

2 Execute the following:

```
rpm -e --nodeps zends
```

Installing the OpenJDK

Analytics requires the OpenJDK 7 Development package. Analytics cannot be installed when just the JRE is installed.

Before installing the OpenJDK 7 Development package, remove all other versions of Java that are on the system:

```
yum -y remove $(rpm -qa | egrep -i '(jdk|jre|java)')
```

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Install the OpenJDK 7.

```
yum -y install java-1.7.0-openjdk-devel
```

- 3 Execute the following command and choose the OpenJDK 7 selection:

```
alternatives --config java
```

- 4 Verify that the installation succeeded.

```
java -version
```

If the command returns output similar to the following, continue to the next procedure.

```
java version "1.7.0_75"  
OpenJDK Runtime Environment (rhel-2.5.4.7.el7_1.x86_64 u75-b13)  
OpenJDK 64-Bit Server VM (build 24.75-b04, mixed mode)
```

Setting the JAVA_HOME environment variable

After installing the OpenJDK, you must set `$JAVA_HOME` for the root user.

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Locate and note the path to your OpenJDK directory. Typically, this is `/usr/lib/jvm/java-1.7.0-openjdk-VersionNumber`.
- 3 Navigate to the `/etc/default` directory.

```
echo $JAVA_HOME
```

- 4 Open the `zenoss_analytics` file in an editor.

```
vi zenoss_analytics
```

- 5 Edit the `JAVA_HOME` line in `zenoss_analytics` to point to your OpenJDK location.

```
JAVA_HOME=/usr/lib/jvm/java-1.7.0-openjdk-VersionNumber
```

- 6 Execute the `zenoss_analytics` file.

```
source zenoss_analytics
```

- 7 Test the new path with the following command:

```
echo $JAVA_HOME
```

Installing MariaDB

Before installing MariaDB on systems running Red Hat Enterprise Linux (RHEL), install the Perl Database Interface Module (DBI) with the following command as the root user:

```
yum -y install perl-DBI
```

To install MariaDB 10 on the Analytics server, follow these steps. (The Zenoss DataStore (ZenDS) is not used in Analytics Version 5 and later.)

- 1 Create a custom MariaDB YUM repository. The following example shows information for a CentOS 7 installation. For more information, see <https://downloads.mariadb.org/mariadb/repositories>
 - a As the root user, navigate to the `/etc/yum.repos.d` directory.

```
cd /etc/yum.repos.d
```

- b Create a file called `MariaDB.repo` and enter the following text into the file.

```
# MariaDB 10.0 CentOS repository list
#http://mariadb.org/mariadb/repositories/
[mariadb]
name = MariaDB
baseurl = http://yum.mariadb.org/10.0/centos7-amd64
gpgkey=https://yum.mariadb.org/RPM-GPG-KEY-MariaDB
gpgcheck=1
```

- 2 When the `MariaDB.repo` file is in place, install the MariaDB server:

```
sudo yum -y install mariadb-server
```

Note If you haven't already accepted the MariaDB GPG key, you will be prompted to do so.

- 3 Add the following configuration settings:

```
echo "[mysqld]" >> /etc/my.cnf
echo "datadir=/opt/zends/data" >> /etc/my.cnf
echo "table_open_cache=16K" >> /etc/my.cnf
echo "table_definition_cache=16K" >> /etc/my.cnf
echo "tmp_table_size=2G" >> /etc/my.cnf
echo "max_heap_table_size=2G" >> /etc/my.cnf
echo "join_buffer_size=512K" >> /etc/my.cnf
echo "open_files_limit=200000" >> /etc/my.cnf
echo "wait_timeout=86400" >> /etc/my.cnf
echo "tmpdir=/tmp" >> /etc/my.cnf
echo "innodb_adaptive_hash_index=OFF" >> /etc/my.cnf
echo "innodb_buffer_pool_size=6G" >> /etc/my.cnf
echo "innodb_log_file_size=1892M" >> /etc/my.cnf
echo "innodb_log_buffer_size=128M" >> /etc/my.cnf
```

Note If you want to keep your existing data, add the `datadir` directory for your system. The value listed above is an example.

Note The values for the `innodb` settings should reflect your system. Set the `innodb_buffer_pool_size` to be 80 percent of the total memory on a standalone data warehouse server or 50 percent of memory if the

analytics server and database are on the same server. The `innodb_log_file_size` should be 25 percent of the `innodb_buffer_pool_size` value.

- 4 If you are setting the value of the `datadir`, you need to change its owner since ZenDS uses the `zenoss` user, but MariaDB uses `mysql`. Change the value of `/opt/zends/data` to your system's `datadir` path.

```
chown -R mysql:mysql /opt/zends/data
```

- 5 Start the MariaDB server:

```
/etc/init.d/mysql start
```

- 6 Set MariaDB to start on boot:

```
chkconfig mysql on
```

- 7 Run `mysql_upgrade` to finish the upgrade of MariaDB, since we are upgrading from version 5.x to version 10.x:

```
mkdir /root/tmpdir
/usr/bin/mysql_upgrade -u root --tmpdir=/root/tmpdir
```

If you previously used third-party tools to access ZenDS using the MySQL JDBC driver, we recommend that you change your JDBC driver to use the "MariaDB Client Library for Java Applications 1.1.8" (client in this case is JDBC driver not the query app client). This can be downloaded from <https://downloads.mariadb.org/client-java/1.1.8/>.

Upgrading the Analytics server and Jaspersoft database

Complete the following steps to upgrade the Analytics server to version 5.0.x and upgrade the Jaspersoft database.

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Back up any customizations that you made in or below the `/opt/zenoss_analytics` folder because they will be overwritten during this upgrade.
- 3 Remove the `/opt/zenoss_analytics` directory:

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

- 4 Using the Analytics RPM you downloaded, upgrade Analytics by issuing the following command:

```
yum -y --nogpgcheck localinstall zenoss_analytics-version.noarch.rpm
```

- 5 Navigate to the `/path_to_safe_place` directory where you saved the `backup.zip` file of your reporting database export.
- 6 Change the permission on the `backup.zip` file so that the `zenoss` user can execute it.

```
chmod +x backup.zip
```

- 7 Switch to the `zenoss` user:

```
su - zenoss
```

- 8 Run the following commands:

```
/opt/zenoss_analytics/bin/upgrade_db.py
/opt/zenoss_analytics/bin/fix-theme.sh /path_to_safe_place/backup.zip
cd /opt/zenoss_analytics/jasperserver/buildomatic
./js-import.sh --input-zip=/path_to_safe_place/fixed.backup.zip
```

- 9 If you made a backup of your customizations in step 2, check the new configuration files installed during this upgrade and incorporate any of your customizations that are still compatible with the new version of Jaspersoft.

Installing the user defined function to MariaDB

Before installing the user defined function, ensure that you have installed MariaDB on the *database* server and Analytics on the *application* server.

- 1 As the root user, log into the Analytics application server and locate the following file:

```
/usr/lib64/mysql/plugin/zenoss_percentile.so
```

- 2 Copy `zenoss_percentile.so` to the same directory on the MariaDB database server (`/usr/lib64/mysql/plugin/`).
- 3 On the Analytics application server, change to the zenoss user:

```
su - zenoss
```

- 4 Execute the following script:

```
/opt/zenoss_analytics/bin/setup_zenoss_extensions install
```

Importing the reports database

Analytics version 5.0.x contains data warehouse schema changes. For help migrating custom reports for a previous version of Analytics, contact Customer Support or Customer Services.

Installing the ZenETL ZenPack on the Resource Manager master

After installing the Analytics server, install the ZenETL ZenPack on the Resource Manager master. You can install it using the command line or from the Resource Manager interface. See the corresponding directions below:

To install using the command line:

- 1 As the zenoss user, enter the following command based on the version you are installing:

```
zenpack --install ZenPacks.zenoss.ZenETL-Version-py2.7.egg
```

- 2 Create aliases to the data points on the monitoring templates that you want to bring into Analytics. See the "Running the chkaliases script" instructions contained in [Export data point data](#) on page 17 for information on using the script found in the ZenETL ZenPack. This step is mandatory and must be completed successfully before proceeding.
- 3 Restart the system:

- a Enter the following commands:

```
zenoss stop
ps -eaf | grep $ZENHOME
```

- b For any process that is reported as running, perform a `kill -9` on those processes.
- c Start the system:

```
zenoss start
```

When you restart the Resource Manager master, three additional daemons (`zeneventetl`, `zenmodeletl`, and `zenperfetl`) are started. Since the ZenETL ZenPack is like any other ZenPack, you must update the remote collectors in order for them to receive the changes. See *Zenoss Service Dynamics Resource Management Extended Monitoring* for directions on updating remote collectors and hubs after a ZenPack install.

Be sure to install the latest RPS for your Resource Manager instance so that you get all the latest fixes. The Distributed Collector ZenPack should be at version 2.5.12 or higher otherwise you will be required to use a `daemons.txt` and a `collectordemons.txt` file to stop `zenmodeletl` and `zeneventetl` on running collectors.

Adding aliases

For information about adding aliases to data points, see [Export data point data](#) on page 17. You can create aliases using the `chkaliases` script or by manually adding them to a data point.

Starting the Analytics service

After Resource Manager has been restarted after the upgrade, you must start the Analytics service:

- 1 Log in to the Analytics server as `root`, or as a user with superuser privileges.
- 2 Enable `chkconfig` for the Analytics service:

```
chkconfig zenoss_analytics on
```

- 3 Start the Analytics service:

```
service zenoss_analytics start
```


4

Working with Analytics

The following sections provide information about working with Analytics:

- [Starting and stopping the Analytics server](#)
- [Logging in to Analytics](#)
- [Creating batches](#)
- [Working with the Analytics Repository](#)

Starting and stopping the Analytics server

To stop, start, or restart the Analytics server, such as when troubleshooting problems, use the following commands. While the server is stopped, no data is extracted.

- To start Analytics, log in to the Analytics server as the `root` user and enter the command:

```
service zenoss_analytics start
```

- To stop and then start Analytics, log in to the Analytics server as the `root` user and enter the command:

```
service zenoss_analytics restart
```

- To stop Analytics, log in to the Analytics server as the `root` user and enter the command:

```
service zenoss_analytics stop
```

Logging in to Analytics

To access Analytics, you can:

- Log in to the Resource Manager interface using your Resource Manager user name and password (single sign on).
- Log in directly to Analytics using your Resource Manager user name and password or as an Analytics superuser.

Logging in through the Resource Manager browser interface

Access to Analytics through single sign on (SSO) on the Resource Manager browser interface requires Manager, ZenManager, or ReportingUser privileges, which are 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. For more information about assigning roles, see the "Managing users" chapter in *Zenoss Resource Manager Administration Guide*.

- 1 Log in to the Resource Manager browser interface with your Resource Manager user name and password.
- 2 Select **Reports > Advanced**.
The Getting Started page appears.

Figure 5: Getting Started

Getting Started

Popular Resources

- How-to videos
- How-to articles
- Online Learning Portal

Recently Viewed Items

Device Availability by Day View	Ad Hoc view
Device Availability by Day Report	Report
Gaps In Analytics Batch Processing View	Ad Hoc view
Gaps In Analytics Batch Processing Report	Report
Current Analytics Batch Status View	Ad Hoc view
Current Analytics Batch Status Report	Report
Analytics ETL Configuration View	Ad Hoc view
Analytics ETL Configuration Report	Report
Analytics Batch Status View	Ad Hoc view
Analytics Batch Status Report	Report

Data Sources
Define connection to a database or other data source. [View tutorial](#)

Domains
Add structure to a data source for use in a Ad Hoc views. [View tutorial](#)

Ad Hoc Views
Visualize your data for analysis and report creation. [View tutorial](#)

Reports
Create and format interactive reports from existing Ad Hoc views. [View tutorial](#)

Dashboards
Combine related reports into custom dashboard layouts. [View tutorial](#)

Admin
Configure your server instance and manage user settings.

From the Getting Started page, you can

- View your reports.
- Create a report.
- Create an ad hoc view.
- Change the Analytics server settings.
-

From the Analytics menu options, you can

- View search results, the Analytics repository, or messages.
- Manage organizations, users, roles, or server settings (accessible by administrators). Do not create organizations in the system because they are linked to Zenoss instances.
- Create ad hoc reports, dashboards, or domains.

Logging directly in 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. (http://<your_server_or_IP_address>:7070 or https://<your_server_or_IP_address>:443 for SSL).
The Analytics login page appears.
- 2 Log in as one of four roles:
 - a With your Resource Manager user name and password:

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

Figure 6: Analytics Login - Resource Manager Credentials

The screenshot shows a login form with the following elements:

- A grey header bar with the word "Login" in blue.
- A "Zenoss Instance" label above a dropdown menu.
- The dropdown menu is open, showing "analytics1-app.zenoss.loc" as the selected option.
- Below the dropdown is an empty input field.
- A "User ID:" label above an empty input field.
- A "Password:" label above an empty input field.
- A "Show locale & time zone" link below the password field.
- A blue "Login" button at the bottom.

- b** As an Analytics superuser:

- 1 From the Resource Manager Instance list of options, select `Internal Authentication`.
- 2 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.)
- 3 Enter your Analytics User ID and password in the `User ID` and `Password` fields.

Figure 7: Analytics Login - Analytics Superuser

The screenshot shows a login form with the following elements:

- A grey header bar with the word "Login" in blue.
- A "Zenoss Instance" label above a dropdown menu.
- The dropdown menu is open, showing "Internal Authentication" as the selected option.
- Below the dropdown is an empty input field.
- A "User ID:" label above an empty input field.
- A "Password:" label above an empty input field.
- A "Show locale & time zone" link below the password field.
- A blue "Login" button at the bottom.

- 3 Click **Login**.

Creating batches

In general, you do not need to create manual batches. Only create a manual batch if a scheduled batch fails. 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. The main types of batches match the primary data types (Event, Model, and Performance).

Existing batches and associated status information appear in the batch list, shown in the following figure.

Figure 8: Batch list

Batch ID	Extractor	State	Begin	End	Attempts
150	EVENTS	COMPLETED	2015-04-21T13:00:00	2015-04-21T21:00:00	1
149	localhost	EXTRACTING_PERF	2015-04-21T19:00:00	2015-04-21T21:00:00	2
148	dyoung-lb3	FAILED	2015-04-21T19:00:00	2015-04-21T21:00:00	3

The list shows each batch and its status: UNSTARTED, EXTRACTING, STAGING, FAILED, COMPLETED, or CANCELLED. The CANCELLED state occurs when the collector goes offline for a long period of time.

To add a manual batch:

- 1 From the Resource Manager browser interface, choose **Reports > Configuration**.
- 2 From the batch list, click **Add Batches** and make selections as follows:
 - a **Extractor** - Choose MODEL, EVENTS, or localhost. For each remote collector, a performance extractor is listed. Extractors only appear after they phone home and self-register. This process requires a few minutes after you connect the Resource Manager to the Analytics server and update the hubs and collectors.
 - b **Begin Date** - Optionally adjust the begin date and time. By default, the value is set to two weeks before the current Analytics server date and time.
 - c **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 leaves a batch with a status of UNSTARTED. The batch will not start until that future date. Setting dates for MODEL is ignored; the current model is used.

- 3 To add the batch to the list, click **Submit**.

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. The default values are:

- Event Jobs - 8 hours
- Model Jobs - 24 hours
- Performance Jobs - 2 hours

To edit the repeat interval:

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

Figure 9: Edit Repeat Interval

The screenshot shows the Zenoss Configuration page for the 'EventEtJob'. The 'Configuration' tab is active. On the left, a tree view shows the job hierarchy: 'b09bf16ac985' > 'Jobs' > 'Event Jobs'. The right pane displays the job configuration details:

- Name: EventEtJob
- Group: ETL_TRIGGER_GROUP
- Description: Create an ETL batch for EventEtJob
- Java Class: com.zenoss.reporting.schedule.EtlJob
- Trigger: EventEtTrigger
- Trigger Description: Trigger for ETL job.
- Next Fire Time: 2015-04-22T21:00:00
- Previous Fire Time: 2015-04-22T13:00:00
- Start Time: 2015-04-21T21:49:34
- Repeat Interval: 8.0 hours (selected in a dropdown menu)

A 'Submit' button is located at the bottom of the configuration pane. At the bottom right of the page, there is a status bar showing '0 Jobs'.

- 2 Select a repeat interval (.25 hours to 24 hours), and then click **Submit**.

Filtering performance batches to a list of metrics

When you schedule a batch for a collector, all aliases that are present for the devices being monitored are extracted. If you want to add a new alias to an existing batch at a later date, you will only get future data when the scheduled batches are run. To collect past data, you need to create a filtered list of metric(s) to collect the past data. The following workflow should be used when adding a new alias to an existing scheduled batch:

- 1 Add the new alias in ZODB. See [Adding aliases](#) on page 32.
- 2 Extract one batch normally, or wait for the next scheduled performance extraction to complete. This ensures that the future batches will contain the new metric and that Analytics has registered the new metric.
- 3 Create a new manual batch to collect past data for the new metric by inserting the metric name(s) into the `meta_batch_metric_assoc` table.

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 should not be changed without consulting Zenoss Support.

To change the value of the delay on performance batches:

- 1 Log in to the Control Center master host as a user with `sudo` privileges.
- 2 Attach to the `zenperfetl` service.

```
serviced service attach zenperfetl
```

- 3 Change to the `opt/zenoss/etc` directory.

```
cd opt/zenoss/etc
```

- 4 Open `zenperfetl.conf` in an editor and uncomment the line containing the setting for `perfbatchdelayinminutes` and set it to the value you want (in minutes).

Changing parallelization of performance loads

To increase loading of performance data, a new configuration has been added which loads performance data in parallel. Based on your I/O performance, you may want to increase or decrease the default number of performance threads that Analytics uses. The default value is 10, which is the equivalent of 5 cores.

To change the value, you must edit the `meta_setting` table directly:

- 1 Log in to the Analytics data warehouse database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Enter the following to lower the `concurrent_threads` to a value of 6. (If you want to increase parallelization, set the value to 12.)

```
update meta_setting set setting_value = '6' where setting_name =
'concurrent_threads';
```

Limiting concurrent performance extractions

Because multiple logical collectors can be installed on a single host, multiple `zenperfetl` daemons run on the same host. Under normal conditions, these daemons begin extraction nearly simultaneously, which can saturate disk I/O for a period of time. To manage this behavior, Analytics tracks the hosts where each of its performance extractors run. If many logical collectors are installed on a single host in your environment, you might need to limit the number concurrent performance extractions that can take place at the same time on the host.

To do so, set a value in column `concurrent_extractions_limit` in table `meta_perf_extractor_limits`. By default, there is no limit.

- 1 Log in to the Analytics data warehouse database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Set the `concurrent_extractions_limit` to a value such as 4. The actual value you set depends on your disk I/O.

```
update meta_perf_extractor_limits set setting_value = '4' where
setting_name = 'concurrent_extractions_limit';
```

Changing extraction logging to debug mode

If you are having issues with extraction, you may want to increase the logging level of `zenoss_analytics.out`. To change the logging level to `DEBUG` mode, perform the following:

- 1 As the `root` user, enter the following command on the Analytics server.

```
DEBUG=1 service service_analytics restart
```

- 2 When you have completed your debug analysis, you can return the logging to the default mode by executing the following command as the `root` user on the Analytics server.

```
DEBUG=0 service service_analytics restart
```

Configuring the Analytics data retention policy

The following table shows the recommended data retention settings for typical reporting and analysis needs. Data retention settings are in table `meta_setting` in the reporting (data warehouse) database. Based on your business needs, you might need to change the retention policy for certain types of data. Data is automatically purged daily by the system, typically at midnight, based on these settings. Updates to the values take effect immediately at the next data purge.

setting_key	setting_name	setting_value	setting_description
1	daily_retention_days	365	Days to keep daily aggregation data.
2	raw_retention_days	45	Days to keep raw fact data.
3	event_retention_days	365	Days to keep event data.
10	hourly_retention_days	90	Days to keep hourly aggregate data.

Column `setting_value` specifies the number of days for which a given type of data is kept.

Values in column `setting_name` are as follows:

- Raw, Hourly, and Daily retention refer to performance data only.
 - Raw indicates the lowest resolution of extracted performance data (typically 5 minute data).
 - Hourly and Daily represent hourly and daily aggregates of that data.
- Event retention applies to all events extracted to Analytics.

- 1 Log in to the Analytics reporting database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Enter the following command:

```
select * from meta_setting where setting_name in
('daily_retention_days', 'raw_retention_days', 'event_retention_days',
'hourly_retention_days');
```

- 3 Issue database updates to change the setting value for setting names.

For example, to change the retention policy for raw performance data from 45 days to 30 days, issue the following database update:

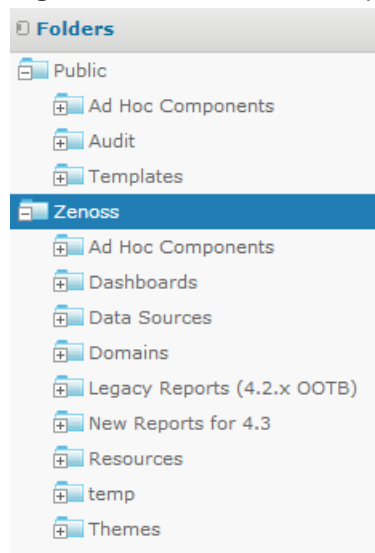
```
update meta_setting set setting_value = '30' where setting_name =
'raw_retention_days';
```

Working with the Analytics repository

The Analytics repository stores all reports, dashboards, analysis views, and related files. The repository is organized as a tree structure of folders containing those components, much like a file system. A set of folders is provided upon installation (such as ad hoc components, dashboards, domains, reports, and resources), as shown in the following figure.

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

Figure 10: Zenoss folder in repository



The offerings in the Zenoss-provided folders might change in a future update, overwriting changes you make to the Zenoss folder. Therefore, it is important to populate a folder specifically for your organization. You can use the contents of the Zenoss folder as a guide to determine which folders you would like to copy to your organization's space in the repository (typically a folder on the same level as the existing Zenoss folder).

To copy a folder:

- 1 Right-click the source folder and select **Copy**.
- 2 Right-click the destination parent folder and select **Paste**.

Searching the repository

After you set up your organization's folder within the repository, you can search for your information for quick access. Select **View > Search Results**. The **Repository Search Results** page displays search results and lists of reports and ad hoc views.

Figure 11: Analytics repository search results

Name	Description	Type	Created Date	Modified Date
Zenoss Reporting	The Zenoss Datawarehouse. Choose this as the datasource for all c	JDBC Data Source	September 10	Today
Service Availability, Performance and MTI		Report	August 26	September 20
Do Not Delete temp folder	Created by Zenoss. See ZEN-8020. Do not delete this report or the	Report	August 29	September 20
Exceptions Dashboard	Dashboard showing all Device CPU, Mempry and Interface Exceptic	Dashboard	September 20	September 20
Last Month Impact Service Availability an	Impact Service Performance and Availability by Service Organizer	Ad Hoc View	August 25	September 20
Device Filesystem Usage by Day View	Filesystem Usage by day	Ad Hoc View	August 26	September 20
Device Filesystem Usage Exceptions Repo	Shows Daily Exceptions on Filesystem usage for selected devices a	Report	September 20	September 20
Gaps In Analytics Batch Processing View	Analysis of potential gaps in analytics data	Ad Hoc View	August 26	September 20
Current Analytics Batch Status View	Any batch unstarted	Ad Hoc View	August 26	September 20
Analytics Batch Status View	Batch Status of all batches	Ad Hoc View	August 26	September 20
Analytics ETL Configuration View	Shows the current state of Analytics ETL Configuration in RM ie whe	Ad Hoc View	August 26	September 20
Impact Service Availability and Performar	Service Impact Availability and Performance	Report	August 26	September 20
IP Interface Performance Exception Repo	Shows Daily Exceptions on IP intreface input and output Utilization I	Report	September 20	September 20
IP Interface Performance View	Table of Input and Output Utilization for IP Interfaces	Ad Hoc View	September 20	September 20
IP Interface Inventory Report	Basic information about IP Interfaces	Report	August 26	September 20
Device Inventory Report	Basic information about Zenoss devices, filterable by any organizer	Report	September 20	September 20
Device Inventory View	Device Inventory Information	Ad Hoc View	September 19	September 20
Device Inventory	Device Information, includes information about CPUs, Filesystems z	Domain	August 26	September 20
Tenant Test	This domain cannot be used to create reports. It is showed here as	Domain	August 26	September 20
Daily Device Availability	Daily Device availability metrics	Domain	August 26	September 20
IP Interface Performance	IP Interface Performance Metrics	Domain	August 26	September 20

On the **Repository Search Results** page, 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.

- Use search controls to find resources by searching for text in their names or descriptions.
- Filters provide a quick way to refine your search results to find resources based on other criteria, such as type or access time.
- The interactive status bar shows search criteria, filters, and sorting criteria.
- List resources in the repository that meet the search criteria. Click or right-click a resource in the list to view it or see available actions.

Searching by name

To begin a search, enter a term the search field at the top of any Analytics page and click the **search** icon. Analytics opens the repository page with the search term and displays the results.

Requirements for search terms and rules for matching are as follows:

- In search terms, do not enter quotes around terms or symbols between terms.
- Search terms are not case sensitive.
- If you enter multiple words, Analytics searches for resources that match all words, but not necessarily in order.
- Analytics matches the search term, even word fragments, wherever the term appears in the display name or description string of the results.
- Analytics does not match folder names or display folders in the results.
- 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 following default settings:
 - Include subfolders
 - Start at the top-most visible folder
 - Sort alphabetically by name

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
All available (default)	Does not take into account the user's access to the resource.
Modified 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 (default)	Does not take into account the resource type; as a result, every type is displayed.
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).
Ad Hoc views	Displays only the output from reports that were created using an ad hoc view into an available data source type (Topics, Domains, or OLAP Connections).
Dashboards	Displays only dashboards.
OLAP views	Displays only OLAP views.
Domains	Displays only domains.
Data sources	Displays only data sources.
Deprecated Reports	Displays only deprecated reports.

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.

Filter Setting	Description
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.
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.

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`
- 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.
- You can click the `search term` or the word `All` to remove any folders to the right.

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).

Working with the list of resources

- 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 `See all...` 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 `See all...`, 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.

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.

Workflow for creating reports

The following sections describe the process to follow for creating your own domains, then creating ad hoc views based on those domains, and finally creating and scheduling reports to be run from those ad hoc views.

Creating domains

A *domain* is a metadata layer that provides a business view of the data accessed through a data source. A domain presents the data in business terms appropriate to your audience, and can limit the access to data based on the security permissions of the user who runs the report. You can use a domain to create reports and ad hoc views. Analytics ships with several out-of-the-box domains, and you can create your own domains based on business needs.

Before you create a domain, review the following considerations and recommendations. For more information about creating domains, see the *JasperReports Server User Guide, Release 6.0*.

- Use the `zenoss_reporting` data source for your domain creation.
- Plan the features of your domain and decide on the elements of your design, such as
 - tables and columns to choose
 - joins to perform
 - filters to define
 - sets you need
 - item properties to expose to users
- Determine a security policy for the data retrieved through the domain.
- Review the [ERD diagram for Analytics 5.0.x](#) on page 58. It shows the relationships between tables in the database. Also included is a sample of dynamically created dimension tables and a flow of performance fact data as it goes through hourly and daily aggregation.

To create a new domain:

- 1 Log in to Analytics as Administrator or superuser.
- 2 Click **CreateDomain..** The **Add New Domain** window is displayed.

- 3 Enter the name, resource ID, and description.
- 4 Choose a **Save Location** by browsing to your organization's Domain folder.
- 5 In the **Data Source** field, browse to the `zenoss_reporting` data source located at `/organizations/zenoss/Data_Sources/zenoss_reporting`.
- 6 In the **Domain Design** section, click **Create with Domain Designer**. The **Domain Designer** window is displayed.
- 7 Enter information as follows:
 - a **Tables** - Select all tables whose columns you want to use in the domain, either directly or indirectly.
 - b **Derived Tables** - Enter queries whose results appear as derived tables in the domain.
 - c **Joins** - Define inner and outer joins between all tables and derived tables.
 - d **Calculated Fields** - Enter expressions whose results appear as calculated fields. For an example of a calculated field, see the **IP Interface Inventory** domain in the `Zenoss/Domains` folder. The `_deleted` calculated fields are based on whether the time stamp is present.
 - e **Pre-filters** - Specify conditions on field values to limit the data that is accessed through the domain. For an example of a pre-filter, see the **Daily Device KPI** domain in the `Zenoss/Domains` folder. You cannot use this domain if you have components. Thus, a pre-filter is defined as `component_key = 0`.
 - f **Display** - Organize the visual aspects of the domain and change the display properties of tables, columns, sets, and items exposed to domain users. Properties are located in the far-right column. **Label** provides the friendly text display, **ID** is the system name that should not be changed unless similar IDs come from different input controls. The **Summary Function** section defines performance metrics as measures; everything else in the system, including time, is defined as a field.
- 8 To validate the design as you create your domain, click **Check Design**. If you have issues and need assistance, click **Export Design**. This action creates an XML file of your schema that you can share with Zenoss Support.
- 9 When you are finished with the **Domain Designer**, click **OK**.
- 10 When the information and design are correct, add the domain to the repository by clicking **Submit**.

Creating ad hoc views

The Ad Hoc editor supports the creation of views for various types of reports: tables, crosstabs, and charts. You interact with the editor to create these views by dragging and dropping elements. You can add and summarize fields, define groups, label and title the report, and format data for each field.

To open the Ad Hoc editor:

- 1 Log in to Analytics as an administrator or superuser.
- 2 Click **Create > Ad Hoc View**. The Select Data screen is displayed.
- 3 Select your data source from the list. Click the icon to navigate through the :

Note You can also navigate through the data source by clicking on the **View as tree** icon. (**root > Organization > Zenoss > Ad Hoc Components**)

- **Topics** - These can be either a JRXML file that is associated with a data source on the server or can be created from a domain with one or more filters applied to it. Using a topic as your source generates an empty view, which allows you to begin adding data to your view right away, without choosing, pre-filtering, or changing display names of the data (all of which are required steps when creating a Domain-based view).
 - **Domains** - They specify tables in the database, join clauses, calculated fields, display names, and default properties, all of which define items and sets of items for creating Ad Hoc views.
- 4 Fill out the New Ad Hoc View screen based on the data you selected. For more details on the functionality and toolbar icons, see the *JasperReports Server User Guide*.
 - 5 Select a view type:


- **Tables** - Used to view values in the database and to summarize the values in columns.
 - **Charts** - Used to compare one or more measures across multiple sets of related fields. Charts summarize data graphically. Some types of charts include bar chart, line chart, and pie chart, among others.
 - **Crosstabs** - Used to aggregate data across multiple dimensions. Crosstabs are more compact representations than tables; they show only computed values, rather than individual database values. Columns and rows specify the dimensions for grouping; cells contain the summarized measurements.
- 6 When completed, save your Ad Hoc View. This view is saved in the repository. You can use this view to create reports or you can save the view as a report itself so that you can embed the data content into a dashboard or see the data in the interactive report viewer.

Note When you create a report from an Ad Hoc view, the report is considered "dependent" on that view. If you later edit the Ad Hoc view, the dependent reports are not updated with the new changes to the view.

Creating and running reports

Zenoss Analytics makes it easy to create and run reports. When you run a report, it opens in the interactive Report Viewer. Using this Viewer, you can personalize and refine the displayed report data. If your report has input controls, you can run the report with one set of data and then another. You can then export the output to another format or you can schedule the report to create an output repeatedly and unattended during off hours or at other times.

Creating a report and editing the output

- 1 Log in to Analytics.
- 2 From the Home page, click **Create** in the Reports section; or from any page, click **Create > Report**.
- 3 From the **Create Report** window, select an ad hoc view and click **OK**.
- 4 If you are prompted for input controls, provide them in the displayed window. If not, the report output appears.
- 5 Edit the report output by clicking the Options  icon in the upper-right corner.

Note By default, a dynamic date range is used to define the date filter. The range is defined by the unit (DAY, WEEK, MONTH) and a numerical unit that can be positive or negative. For example a range of DAY-7 to DAY is equivalent to the last seven days from when the report is run.

- 6 (optional) Change various elements of the output based on the type of report. For example, you can resize columns or filter by a column's value. If you are using a chart, you can change the type of chart you want to display. Refer to the *JasperReports Server User Guide* for even more ways to customize your reports. You cannot change the fundamental type of output (that is, you cannot change a chart into a crosstab, for example).
- 7 When the report output display is satisfactory, save the report.

Scheduling a report job to run

- 1 Log in to Analytics.
- 2 On the Home page, click **View list** in the Reports area; or from any page, click **View > Repository**.
- 3 Use the search field or browse the list of reports to find the report you want to schedule.
- 4 Right-click the report name and select **Schedule** from the context menu. The Schedule Jobs page appears.
- 5 Click **Create Schedule**. The Schedule tab appears. Fill out the information requested and select the recurrence of the job. The default value is None, which will run the report only once. Select **Simple** from the drop-down list to schedule the job to recur at an hourly, daily, or weekly interval or select **Calendar** to schedule the job to recur on certain days of the week or days of the month.
- 6 Set the other values on the appropriate tabs for **Parameters**, **Output Options**, and **Notifications**. Click **Save** to save the schedule. On the Repository view page, you will see a clock icon to indicate that a report has been scheduled to run.

Note The permissions of the user who schedules a job determines the data that the report exposes. Be aware that sensitive data could be exposed if you schedule a job as an administrative user with no data restrictions because the report output will contain all the requested data in the data source. Any user who receives the report output can view all the data regardless of the user's access restrictions.

Entity-Relationship Diagrams for aggregation-related tables

This section describes percentiles and aggregates and the way to understand how projections are created.

Percentiles operate against the raw data, i.e., `raw_v2_tables`, not against an aggregate. Think of the `nth_percentile_configuration` tables as defining what subset of all the raw data has percentiles calculated, i.e., which rows within a specific `raw_v2_table` (specific `device_key` - group filter, specific `fct_ts` values - business hours filter) and which raw tables are actually read at all (specific metrics).

Projections operate against already calculated daily aggregates or against `nth` percentile calculations that have already been performed on `daily_` or `percentile_` tables. You can think of a projection of `nth` percentiles as a dependency, since `nth`-percentile configurations can limit the data produced, the projections on that `nth` percentile have, at a minimum, the same limits. You can't have projections that specify fewer overall restrictions than the data source it depends on.

Projections groups, business hours, and metric configuration can be thought of as being used to decide which existing aggregate table(s) the values to be projected are read from, i.e., they only configure which table's values are projected and what additional group filters to apply to select devices.

Some configuration override scenarios:

- The specified `nth`-percentile metric configuration overrides the projection metric configuration
- The `nth`-percentile configuration overrides the projection configuration for business hour selection on a per metric basis
- If both `nth`-percentile and projection have group filters configured, only the intersection of the two sets is used. Otherwise, the configured group filter from either is used, if any are specified at all.

The following Entity-Relationship Diagrams (ERD) show the relationships between aggregation-related tables in the database.

Figure 12: Configuration ERD of aggregation-related tables

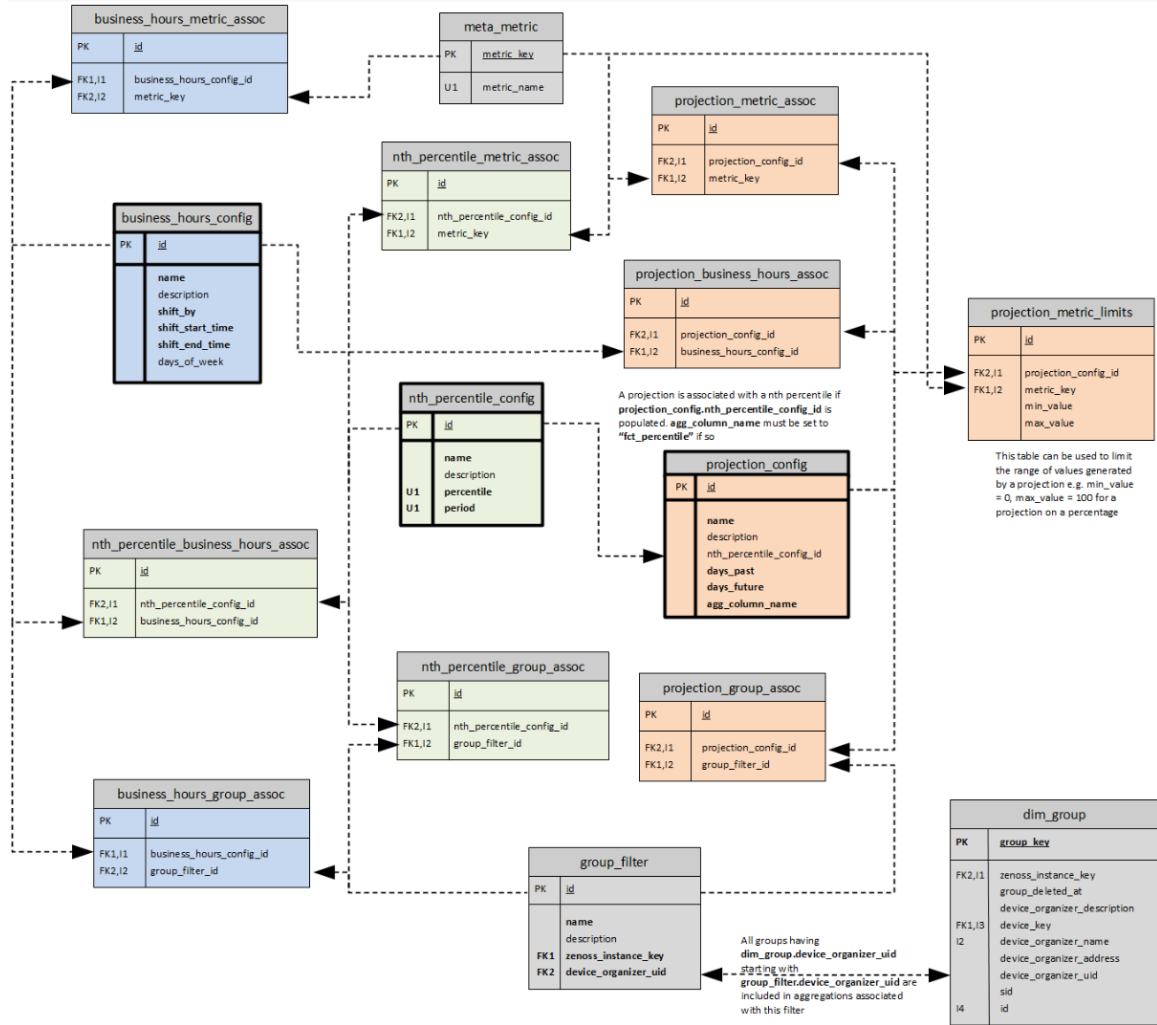
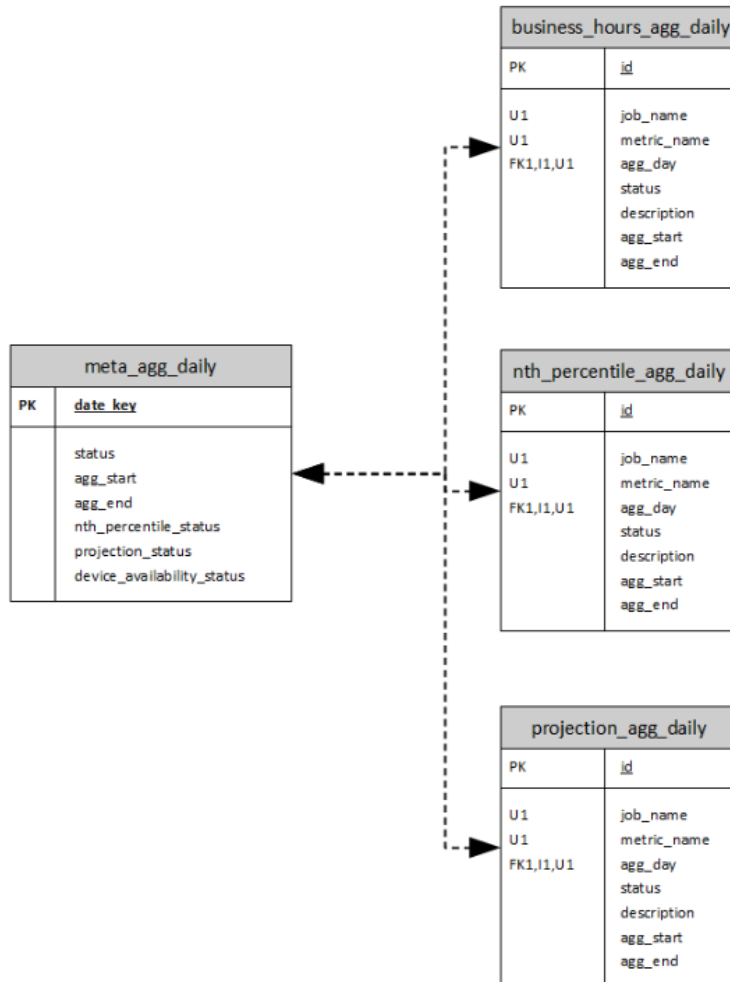


Figure 13: Status tracking ERD of the aggregation execution



Creating business hours aggregation

You can prepare reports showing information relative to the business hours of operation based on a user-specified time offset, which allows for more accurate reporting of how applications and services are performing relative to different user communities. For example, you can aggregate data for a specified time range that has been defined as a "business day" or "shift". Each row is named for the shift or business hours it represents, e.g., austin, london, singapore as locations or shifts like austin_day, austin_swing, austin_night.

The creation of rows for the business hour aggregation is performed manually by inserting records into the reporting database. The tables of interest are described below:

- **business_hours_config**: main table defining the business hours
 - **name**: name of the shift or business hours
 - **description**: up to 255 characters to describe the row (optional)
 - **shift_by**: difference in time between the raw data time stamp and the time specified by the shift_start and shift_end fields. This value can be negative and is in the format hh:mm.
 - **shift_start_time**: time of day when the shift (or business day) begins. The format is hh:mm.

- `shift_end_time`: time of day when the shift (or business day) ends. This value can be less than `shift_start` which indicates that the end of the shift is the following day. The shift begins at the `shift_start` time and goes up to but does not include the `shift_end` time. The format is `hh:mm`.

<code>shift_start_time</code>	<code>shift_end_time</code>	<code>shift_by</code>	Selected range of raw data
08:00	16:00	04:00	12:00 PM through 8:00 PM of the current day
17:00	02:00	00:00	Assuming today is 5/15/2013: 5/15/2013 5:00 PM through 5/16/2013 2:00 AM
06:00	14:00	-11:30	Assuming today is 5/15/2013: 5/14/2013 6:30 PM through 5/15/2013 2:30 AM

- `group_filter`: optional filter used to reduce the set of devices included in the aggregation. If a location is used, such as London, you may only want to include devices in the London data center. This filter uses a `zenoss_instance_key` from `dim_group`, so groups need to be created and populated using the grouping procedure from the user interface. If this filter is not provided for a business hour row (see `business_hours_group_assoc` below), all devices are included in the aggregation.
 - `filter_name`: name to identify the filter
 - `description`: up to 255 characters to describe the filter (optional)
 - `zenoss_instance_key`: key of the zenoss instance, used along with `device_organizer_uid` to identify a device grouping in a particular zenoss instance
 - `device_organizer_uid`: id of the device group, used along with `zenoss_instance_key` to identify a device grouping in a particular zenoss instance
- `meta_metric`: optional filter used to reduce the set of metrics included in the aggregation. If this filter is not provided for a business hour row (see `business_hours_group_assoc` below), all metrics are included in the aggregation.
 - `metric_key`: primary key
 - `metric_name`: name of the metric to aggregate
- `business_hours_group_assoc`: each business hour row may have multiple associated job filters. Each job filter may be used by many business hour rows. For each job filter associated with a business hour row, an association record is created.
 - `business_hours_config_id`: id of the row in the business hours table
 - `group_filter_id`: id of the job filter
- `business_hours_metric_assoc`: each business hour row may have multiple associated metric filters. Each metric filter may be used by many business hour rows. For each metric filter associated with a business hour row, an association record is created.
 - `business_hours_config_id`: id of the row in the business hours table
 - `metric_key`: id of the metric filter

Inserting records for business hours aggregation

Perform the following steps to insert records into the reporting database. The code in this procedure is for demonstration purposes only. This example assumes that you have two data centers (Austin and Bangalore) already defined and the metrics you want to filter on are called `analytics_test1` and `analytics_test2`:

- 1 Log in to the Analytics reporting database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Enter the following commands at the prompt substituting values that apply to your company's situation. The first section of the following inserts groups. If you already have groups, ignore this section and customize the commands to your situation:

```

insert into dim_group
  (group_key, sid, zenoss_instance_key, id, device_organizer_name,
   device_key,
   device_organizer_uid)
values
  (1001, '8525e4e2-4da4-4da5-8d4c-20ec70eaae0b/analytics_test_1', 1,
   'AustinDataCenter',
   'AustinDataCenter', 1, '/zport/dmd/Groups/AustinDataCenter'),
  (1002, '8525e4e2-4da4-4da5-8d4c-20ec70eaae0b/analytics_test_3', 1,
   'AustinDataCenter',
   'AustinDataCenter', 3, '/zport/dmd/Groups/AustinDataCenter'),
  (1003, '8525e4e2-4da4-4da5-8d4c-20ec70eaae0b/analytics_test_2', 1,
   'BangaloreDataCenter',
   'BangaloreDataCenter', 2, '/zport/dmd/Groups/
BangaloreDataCenter'),
  (1004, '8525e4e2-4da4-4da5-8d4c-20ec70eaae0b/analytics_test_4', 1,
   'BangaloreDataCenter',
   'BangaloreDataCenter', 4, '/zport/dmd/Groups/
BangaloreDataCenter');

insert into group_filter
  (name, description, zenoss_instance_key, device_organizer_uid)
values
  ('austin', 'Devices in the Austin data center', 1, '/zport/dmd/
Groups/AustinDataCenter'),
  ('bangalore', 'Devices in the Bangalore data center', 1,
   '/zport/dmd/Groups/BangaloreDataCenter');

insert into meta_metric (metric_name)
values ('analytics_test1'), ('analytics_test2');

insert into business_hours_config
  (name, description, shift_by, shift_start_time, shift_end_time)
values
  ('austin', 'Austin data center', '-06:00', '8:00', '16:00'),
  ('bangalore', 'Bangalore data center', '-09:30', '7:00', '17:00');

insert into business_hours_group_assoc
  (business_hours_config_id, group_filter_id)
values
  (
    (select id from business_hours_config where name='austin'),
    (select id from group_filter where name='austin')
  ),
  (
    (select id from business_hours_config where name='bangalore'),
    (select id from group_filter where name='bangalore')
  );

insert into business_hours_metric_assoc
  (business_hours_config_id, metric_key)
values
  (
    (select id from business_hours_config where name='austin'),
    (select metric_key from meta_metric where
metric_name='analytics_test1')
  )

```

```

),
(
  (select id from business_hours_config where name='austin'),
  (select metric_key from meta_metric where
metric_name='analytics_test2')
),
(
  (select id from business_hours_config where name='bangalore'),
  (select metric_key from meta_metric where
metric_name='analytics_test1')
),
(
  (select id from business_hours_config where name='bangalore'),
  (select metric_key from meta_metric where
metric_name='analytics_test2')
);

```

Creating *n*th percentile calculations

Analytics allows you to create *n*th percentile calculations that can be added to a report. The *n*th percentile is the smallest value in the set of raw data with the property that *n*% of the data values are less than or equal to it. For the calculation to provide the targeted information you need, you must define parameters in reporting database tables.

To create rows for the *n*th percentile calculation, manually insert records into the reporting database. The tables of interest are as follows:

- `nth_percentile_config`

The main table for defining the percentile jobs.

- `percentile` - Specifies the *n*th value to calculate; for example, the 95th percentile.
- `period` - Specifies the number of days of data to use for the *n*th percentile calculation; for example, the prior 30 days. This value is bounded by the `raw_retention_days` value.
- `name` - Not used for the *n*th percentile; only exists for integration with an interface.
- `description` - Not used for the *n*th percentile; only exists for integration with an interface.

- `business_hours_config`

Defines the business hours to calculate within. That is, it specifies the time ranges that define as "business day" or "shift." If not defined, the 24-hour "shift" is used.

- `nth_percentile_business_hours_assoc`

Relates `nth_percentile_config` records with `business_hours_config` records. If not defined, `projection_business_hours_assoc` is used.

- `nth_percentile_config_id` - Specifies the `nth_percentile_config` foreign key.
- `business_hours_config_id` - Specifies the `business_hours` foreign key.

- `group_filter`

Also known as a job filter, it defines an optional filter to use to limit the calculation to certain devices. If not defined, all devices are included in the calculation.

- `nth_percentile_group_assoc`

Relates `nth_percentile_config` records with `group_filter` records. If `projection_group_assoc` is also defined, only the intersection is used.

- `nth_percentile_config_id` - Specifies the `nth_percentile_config` foreign key.
- `group_filter_id` - Specifies the `group_filter` foreign key.
- `meta_metric`

Defines an optional filter to use to limit the calculation to certain metrics. If not defined, all metrics are included.

- `nth_percentile_metric_assoc`

Relates `nth_percentile_config` records with `meta_metric` records. If `nth_percentile_metric_assoc` is not defined for a metric, `projection_metric_assoc` is used.

- `nth_percentile_config_id` - Specifies the `nth_percentile_config` foreign key.
- `metric_key` - Specifies the `meta_metric` foreign key.

For each aggregate table to be projected, the table name from which to read the data is constructed as follows:

```
nth_percentile_<nth_percentile_config.name><business_hours_config.name,
if any><meta_metric.metric_name>
```

The table name to which to write data is constructed as follows:

```
proj_<projection_config.name><nth_percentile_config.name, if
any><business_hours_config.name, if any>_<meta_metric.metric_name>
```

Inserting records for *n*th percentile calculations

This example procedure inserts records into the reporting database and is for demonstration purposes only. For this example, assume that two data centers are already defined, Austin and Bangalore. The metrics on which to filter are called `analytics_test1` and `analytics_test2`.

- 1 Log in to the Analytics reporting database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Enter the following command at the prompt, substituting values that apply to your company's situation:

```
insert into nth_percentile_config
  (name, description, percentile, period)
values
  ('percentile_test_1', 'first percentile test', 90, 10),
  ('percentile_test_2', 'second percentile test', 80, 20);

insert into nth_percentile_business_hours_assoc
  (nth_percentile_config_id, business_hours_config_id)
values
  (
    (select id from nth_percentile_config where percentile = 90 and
     period = 10),
    (select id from business_hours_config where name='austin')
  ),
  (
    (select id from nth_percentile_config where percentile = 80 and
     period = 20),
    (select id from business_hours_config where name='bangalore')
  );
```

```

insert into nth_percentile_group_assoc
  (nth_percentile_config_id, group_filter_id)
values
  (
    (select id from nth_percentile_config where percentile = 90 and
     period = 10),
    (select id from group_filter where name='austin')
  ),
  (
    (select id from nth_percentile_config where percentile = 80 and
     period = 20),
    (select id from group_filter where name='bangalore')
  );

insert into nth_percentile_metric_assoc
  (nth_percentile_config_id, metric_key)
values
  (
    (select id from nth_percentile_config where percentile = 90 and
     period = 10),
    (select metric_key from meta_metric where
     metric_name='analytics_test1')
  ),
  (
    (select id from nth_percentile_config where percentile = 90 and
     period = 10),
    (select metric_key from meta_metric where
     metric_name='analytics_test2')
  ),
  (
    (select id from nth_percentile_config where percentile = 80 and
     period = 20),
    (select metric_key from meta_metric where
     metric_name='analytics_test1')
  ),
  (
    (select id from nth_percentile_config where percentile = 80 and
     period = 20),
    (select metric_key from meta_metric where
     metric_name='analytics_test2')
  );

```

Projecting other aggregates

You can project any numeric metric into the future by defining the number of days of past data to analyze along with the number of days into the future you want to predict. For example, you want to predict the next 30 days of an *n*th percentile calculation based on the last 60 days of data. This prediction returns a value for every day for the next 30 days. To set up this prediction, define a table (`projection_config`) in the reporting database for the projection configuration of the future values.

Projections are performed using a polynomial function, which is beneficial when data follows a curved pattern toward a maximum or minimum value, such as percentage of disk storage used over time.

The `projection_config` table has the following form:

- `name` - Specifies the name to identify the projection; for example, `bangalore_60_30`, `austin_95th`.
- `description` - In up to 255 characters, describes the projection.
- `nth_percentile_config_id` - Specifies the ID of the row in the `nth_percentile_config` table on which to base a projection (only used for *n*th percentile projections).

- `days_past` - Specifies an integer value that indicates the number of past days to use in making the calculation.
- `days_future` - Specifies an integer value that indicates the number of days into the future to make projections.
- `agg_column_name` - Specifies the name of the aggregation column.

Projections run each day at the completion of aggregation and *n*th percentile processing. Tables that result from projection computations have the same structure as aggregation tables, plus a column with the timestamp of the projection. For each aggregate table to be projected, the name of the table from which to read data is constructed as follows:

```
daily_<business_hours_config.name, if any>_<meta_metric.metric_name>
```

The name of the table to which to write data is constructed as follows:

```
proj_<projection_config.name><business_hours_config.name, if any>_<meta_metric.metric_name>
```

Each projection row can have one or more associated group (job) filters and metric filters. A group filter reduces the device sample set by allowing you to create a device group and assigning devices to that group. You can define several group filters for a given projection. Each filter is additive and increases the number of the devices included in the sample set. If no group (job) filters are defined, then all devices are included in the sample set.

Metric filters restrict the metrics that are processed in the projection. If no metric filters are defined, then all metrics are processed. Each metric processed runs in a separate thread. The projection calculations use data from the `raw_v2` tables. To accommodate the maximum flexibility in which data can be used, the metric name in this context is the name of the daily table without the preceding "daily_". For normal aggregated metrics, this equates to the metric name (such as `cpu_pct`). For specially aggregated (computed) data, such as business hours or *n*th percentile, the metric name includes the computed name (such as `austin_cpu_pct`).

Inserting records for daily aggregation projections

This example procedure inserts records into the reporting database and is for demonstration purposes only. For this example, assume that two data centers already defined, Austin and Bangalore. The metrics on which to filter are called `analytics_test1` and `analytics_test2`.

- 1 Log in to the Analytics reporting database as the `root` user or a user with administrative permissions.

```
mysql -u root reporting
```

- 2 Enter the following command at the prompt, substituting values that apply to your company's situation:

```
insert into projection_config
  (name, days_past, days_future, agg_column_name)
values
  ('austin_60_30', 60, 30, 'fct_avg'),
  ('bangalore_60_30', 60, 30, 'fct_avg'),

insert into projection_business_hours_assoc
  (projection_config_id, business_hours_config_id)
values
  (
    (select id from projection_config where name='austin_60_30'),
    (select id from business_hours_config where name='austin')
  ),
  (
    (select id from projection_config where name='bangalore_60_30'),
    (select id from business_hours_config where name='bangalore')
```

```

);

insert into projection_group_assoc
  (projection_config_id, group_filter_id)
values
  (
    (select id from projection_config where name='austin_60_30'),
    (select id from group_filter where name='austin')
  ),
  (
    (select id from projection_config where name='bangalore_60_30'),
    (select id from group_filter where name='bangalore')
  );

insert into projection_metric_assoc
  (projection_config_id, metric_key)
values
  (
    (select id from projection_config where name='austin_60_30'),
    (select metric_key from meta_metric where
metric_name='analytics_test1')
  ),
  (
    (select id from projection_config where name='austin_60_30'),
    (select metric_key from meta_metric where
metric_name='analytics_test2')
  ),
  (
    (select id from projection_config where name='bangalore_60_30'),
    (select metric_key from meta_metric where
metric_name='analytics_test1')
  ),
  (
    (select id from projection_config where name='bangalore_60_30'),
    (select metric_key from meta_metric where
metric_name='analytics_test2')
  );

```

Inserting records for *n*th percentile projections

You can insert records that allow for *n*th percentile calculations. The following sample inserts records into the `projection_config` table for a 95th percentile projection:

```

insert into projection_config
  (name, nth_percentile_config_id, days_past, days_future,
agg_column_name)
values
  ('austin_95th', 1, 60, 30, 'fct_percentile')

```

The value 1 is the ID of the row in the `nth_percentile_config` table where the Austin 95th percentile calculation is defined.

A

Third party licenses

Zenoss Analytics incorporates third party components whose licenses require us to mention their inclusion. For more information, see the following sites:

- Antlr - <http://www.antlr.org/>
- AspectJ weaver - <https://eclipse.org/aspectj/>
- Java Specification Request (JSR) 250 - <https://jcp.org/en/jsr/detail?id=250>
- connector-api-1.5

B

ERD diagram for Analytics 5.0.x

The following page shows the Entity-Relationship Diagram (ERD) for Analytics 5.0.x. You will find the relationships shown between tables in the database as well as a sample of dynamically created dimension tables.

Performance Fact Data One Table Per Metric

meta_metric		
PK	metric_key	INTEGER
U1	metric_name	VARCHAR(40)

raw_v2_device_availability		
PK,FK1,FK2	device_key	INTEGER
PK	component_key	INTEGER
PK	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_resolution	INTEGER
	fct_value	FLOAT

raw_v2_in_pct		
PK	device_key	INTEGER
PK	component_key	INTEGER
PK	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_resolution	INTEGER
	fct_value	FLOAT

* Device Availability calculated from Ping Down Events and when they are Cleared.

Hourly
Aggregation

* Component_key may link to a range of Dimension tables based on what component emits that metric.

hourly_device_availability		
PK,FK1,I1	device_key	INTEGER
PK,FK2,I1	component_key	INTEGER
PK,I1	date_key	DATETIME
PK,I1	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_avg	FLOAT
	fct_resolution	INTEGER
	fct_min	FLOAT
	fct_max	FLOAT
	fct_count	INTEGER

hourly_in_pct		
PK,I1	device_key	INTEGER
PK,I1	component_key	INTEGER
PK,I1	date_key	DATETIME
PK,I1	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_avg	FLOAT
	fct_resolution	INTEGER
	fct_min	FLOAT
	fct_max	FLOAT
	fct_count	INTEGER

Daily
Aggregation

daily_device_availability		
PK,FK1,I1	device_key	INTEGER
PK,FK2,I1	component_key	INTEGER
PK,I1	date_key	DATETIME
PK	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_avg	FLOAT
	fct_resolution	INTEGER
	fct_min	FLOAT
	fct_max	FLOAT
	fct_count	INTEGER

daily_in_pct		
PK,I1	device_key	INTEGER
PK,I1	component_key	INTEGER
PK,I1	date_key	DATETIME
PK,I1	fct_ts	DATETIME
	fct_ts_gmt	DATETIME
	fct_avg	FLOAT
	fct_resolution	INTEGER
	fct_min	FLOAT
	fct_max	FLOAT
	fct_count	INTEGER

Sample of Dynamically Created Dimension Tables

dim_device		
PK	device_key	INTEGER
FK1,I1	zenoss_instance_key	INTEGER
	device_deleted_at	DATETIME
	device_collector	VARCHAR(255)
	device_hw_asset_tag	VARCHAR(255)
	device_hw_serial_number	VARCHAR(255)
	device_last_model	DATETIME
	device_os_manufacturer	VARCHAR(255)
	device_os_uptime	INTEGER
	device_sys_descr	VARCHAR(255)
	device_sys_object_id	VARCHAR(255)
	device_uid	VARCHAR(255)
	device_availability	VARCHAR(8)
	device_comments	VARCHAR(1000)
	device_hw_manufacturer	VARCHAR(255)
	device_is_monitored	TINYINT
	device_manage_ip	VARCHAR(39)
	device_os_product	VARCHAR(255)
	device_production_state	INTEGER
	device_sys_location	VARCHAR(255)
	device_total_memory	INTEGER
	device_class	VARCHAR(255)
	device_first_seen	DATETIME
	device_hw_product	VARCHAR(255)
	device_last_change	DATETIME
	device_name	VARCHAR(255)
	device_os_uname	VARCHAR(255)
	device_rack_slot	VARCHAR(255)
	device_sys_name	VARCHAR(255)
	device_total_swap	INTEGER
	sid	VARCHAR(255)
	id	VARCHAR(255)
	device_production_state_name	VARCHAR(255)

dim_group		
PK	group_key	INTEGER
FK2,I1	zenoss_instance_key	INTEGER
	group_deleted_at	DATETIME
FK1,I3	device_key	INTEGER
I2	device_organizer_name	VARCHAR(255)
	device_organizer_address	VARCHAR(1000)
	device_organizer_uid	VARCHAR(255)
	sid	VARCHAR(255)
	id	VARCHAR(255)

dim_location		
PK	location_key	INTEGER
FK2,I1	zenoss_instance_key	INTEGER
	location_deleted_at	DATETIME
FK1,I3	device_key	INTEGER
I2	device_organizer_name	VARCHAR(255)
	device_organizer_address	VARCHAR(1000)
	device_organizer_uid	VARCHAR(255)
	device_organizer_description	VARCHAR(1000)
	sid	VARCHAR(255)
	id	VARCHAR(255)

dim_system		
PK	system_key	INTEGER
FK2,I1	zenoss_instance_key	INTEGER
	system_deleted_at	DATETIME
FK1,I3	device_key	INTEGER
I2	device_organizer_name	VARCHAR(255)
	device_organizer_address	VARCHAR(1000)
	device_organizer_uid	VARCHAR(255)
	device_organizer_description	VARCHAR(1000)
	device_key	INTEGER
	device_organizer_name	VARCHAR(255)
	sid	VARCHAR(255)
	id	VARCHAR(255)

dim_date		
PK	date_key	INTEGER
I1	date_date	DATE
U1	date_year	TINYINT
U1	date_month	TINYINT
U1	date_day	TINYINT
	date_week_of_year	TINYINT
	date_day_of_week	TINYINT
	weekend	TINYINT
	holiday	TINYINT
	business_day	TINYINT

dim_ip_interface		
PK	ip_interface_key	INTEGER
FK2,I1	zenoss_instance_key	INTEGER
	ip_interface_deleted_at	DATETIME
	ip_interface_description	VARCHAR(255)
I4	ip_interface_is_monitored	TINYINT
	ip_interface_name	VARCHAR(255)
	ip_interface_uid	VARCHAR(255)
	ip_interface_id	VARCHAR(255)
	ip_interface_mac_address	VARCHAR(23)
	ip_interface_speed	INTEGER
	device_key	INTEGER
	ip_interface_index	INTEGER
	ip_interface_mtu	INTEGER
	ip_interface_type	VARCHAR(255)
	sid	VARCHAR(255)
	id	VARCHAR(255)

dim_zenoss_instance		
PK	zenoss_instance_key	INTEGER
U1	zenoss_instance_fqdn	VARCHAR(255)
U2	zenoss_instance_uid	CHAR(36)
	zenoss_instance_base_url	VARCHAR(255)
	zenoss_instance_internal_base_url	VARCHAR(255)
	zenoss_instance_identifier	VARCHAR(255)

dim_file_system		
PK	file_system_key	INTEGER
FK2,I1	zenoss_instance_key	INTEGER
	file_system_deleted_at	DATETIME
I2	file_system_name	VARCHAR(255)
	file_system_uid	VARCHAR(255)
	device_key	INTEGER
	file_system_total_bytes	INTEGER
	file_system_mount_point	VARCHAR(255)
	file_system_type	VARCHAR(255)
	sid	VARCHAR(255)
	id	VARCHAR(255)

Event Fact Data

fct_event_time		
PK,FK1	date_key	INTEGER
PK,FK2,I1	event_evid	CHAR(36)
PK	created_ts	DATETIME
	processed_ts	DATETIME

fct_event		
I6	event_ts	DATETIME
	event_last_ts	DATETIME
	event_source_table	CHAR(7)
	event_dedupid	VARCHAR(255)
	event_evid	CHAR(36)
	event_device	VARCHAR(255)
	event_component	VARCHAR(255)
	event_key	VARCHAR(128)
	event_summary	VARCHAR(128)
	event_message	VARCHAR(4096)
	event_severity	TINYINT
	event_severity_name	CHAR(8)
	event_state	TINYINT
	event_state_name	CHAR(12)
	event_production_state	INTEGER
	event_source_class_key	VARCHAR(128)
	event_group	VARCHAR(64)
	event_count	INTEGER
	event_agent	VARCHAR(64)
	event_monitor	VARCHAR(128)
FK4,FK5	zenoss_instance_key	INTEGER
FK2,I5,I4	device_key	INTEGER
FK1,I1,I4	date_key	INTEGER
FK3	event_class_key	INTEGER
I2	date_last_key	INTEGER
I7	event_clear_evid	CHAR(36)
	event_clear_ts	DATETIME
I5,I3	date_clear_key	INTEGER

dim_event_class		
PK	event_class_key	INTEGER
I1	zenoss_instance_key	INTEGER
	event_class_deleted_at	DATETIME
	event_class_description	VARCHAR(255)
	event_class_transform	VARCHAR(255)
	event_class_action	VARCHAR(255)
	event_class_name	VARCHAR(255)
	event_class_uid	VARCHAR(255)
	event_class_clear_classes	VARCHAR(255)
	event_class_severity	INTEGER
	id	VARCHAR(255)
	sid	VARCHAR(255)

dim_user_role		
PK	user_role_key	INTEGER
I1	zenoss_instance_key	INTEGER
	user_role_deleted_at	DATETIME
I3	id	VARCHAR(255)
	user_role_insight_mapping	VARCHAR(255)
FK1,I2	user_settings_key	INTEGER
	sid	VARCHAR(255)
	user_role_reporting_mapping	VARCHAR(255)

dim_user_settings		
PK	user_settings_key	INTEGER
I1	zenoss_instance_key	INTEGER
	user_settings_deleted_at	DATETIME
	eventConsoleRefresh	TINYINT
	oncallEnd	INTEGER
	passwordHash	VARCHAR(100)
	zenossNetUser	VARCHAR(255)
	defaultAdminRole	VARCHAR(255)
	email	VARCHAR(255)
I2	id	VARCHAR(255)
	oncallStart	INTEGER
	title	VARCHAR(255)
	dashboardState	VARCHAR(255)
	defaultEventPageSize	INTEGER
	escalationMinutes	INTEGER
	netMapStartObject	VARCHAR(255)
	pager	VARCHAR(255)
	zenossNetPassword	VARCHAR(255)
	defaultAdminLevel	INTEGER
	defaultPageSize	INTEGER
	sid	VARCHAR(255)

Data Retention and Logging Settings

meta_setting		
PK	setting_key	INTEGER
	setting_name	VARCHAR(100)
	setting_value	VARCHAR(100)
	setting_description	VARCHAR(250)

Batch Processing

meta_batch		
PK	batch_key	INTEGER
FK1,I1	extractor_key	INTEGER
	batch_attempts	INTEGER
	batch_max_attempts	INTEGER
	batch_begin	DATETIME
	batch_end	DATETIME
	batch_uid	CHAR(36)
	batch_created_at	DATETIME
	batch_state	CHAR(15)
	extract_start	INTEGER
	extract_end	INTEGER
	extract_count	INTEGER
	load_start	INTEGER
	load_end	INTEGER
	staging_count	INTEGER
	load_count	INTEGER
	next_batch_time	DATETIME
	agg_time	DATETIME

meta_extractor		
PK	extractor_key	INTEGER
U1	zenoss_instance_key	INTEGER
U1	extractor_name	VARCHAR(255)
I1	extractor_fqdn	VARCHAR(255)
	extractor_type	CHAR(11)
	extractor_inserted_at	DATETIME

meta_agg		
PK	agg_key	INTEGER
	state	CHAR(10)
	agg_start	INTEGER
	agg_end	INTEGER

meta_etl		
PK	etl_key	INTEGER
	etl_description	VARCHAR(45)
	etl_type	CHAR(5)
	etl_created_at	DATETIME
	fk_batch_key	INTEGER

meta_batch_file		
PK,FK1,I1	batch_key	INTEGER
PK,I2	batch_file_path	VARCHAR(255)
	batch_table_name	VARCHAR(127)
	batch_field_names	VARCHAR(511)
	batch_field_types	VARCHAR(1000)
	batch_file_state	CHAR(10)
	model_subset	TINYINT

meta_etl_log		
PK	etl_log_key	INTEGER
FK1,I1	etl_key	INTEGER
	etl_created_at	DATETIME
	etl_log_comment	LONGVARCHAR
	logger	VARCHAR(250)
	etl_log_level	VARCHAR(20)
	stack_trace	LONGVARCHAR

Application Job Scheduling

QRTZ_TRIGGERS		
PK,I7	TRIGGER_NAME	VARCHAR(200)
PK,I5,I13,I8,I7	TRIGGER_GROUP	VARCHAR(200)
PK,FK1,I5,I14,I13,I12,I11,I13,I12,I11,I10,I9,I8,I7,I6	SCHED_NAME	VARCHAR(120)
FK1,I2,I1	JOB_NAME	VARCHAR(200)
FK1,I3,I2,I1	JOB_GROUP	VARCHAR(200)
I13,I12,I11,I10,I9	NEXT_FIRE_TIME	INTEGER
	PREV_FIRE_TIME	INTEGER
	PRIORITY	INTEGER
I13,I12,I10,I8,I7,I6	TRIGGER_STATE	VARCHAR(16)
	TRIGGER_TYPE	VARCHAR(8)
	START_TIME	INTEGER
	END_TIME	INTEGER
	CALENDAR_NAME	VARCHAR(200)
	MISFIRE_INSTR	TINYINT
I4	JOB_DATA	LONGVARBINARY
I13,I12,I11		

QRTZ_SIMPLE_TRIGGERS		
PK,FK1,I1	TRIGGER_NAME	VARCHAR(200)
PK,FK1,I1	TRIGGER_GROUP	VARCHAR(200)
PK,FK1,I1	SCHED_NAME	VARCHAR(120)
	REPEAT_COUNT	INTEGER
	REPEAT_INTERVAL	INTEGER
	TIMES_TRIGGERED	INTEGER