

IBM Operations Analytics Predictive Insights 1.3.5

Configuration and Administration Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 137.

Contents

Preface	v
Audience	v
Components	v

Chapter 1. Initial configuration of Operations Analytics Predictive Insights. 1

Configuring the User Interface	1
Adding the User Interface (UI) launch scripts	1
Configuring user access to Operations Analytics Predictive Insights	2
Dashboard Application Services Hub	2
Tivoli Integrated Portal	3
Setting the DB2 log file size	5
Configuring DB2 to allow for the processing of many metric groups	6
Scheduling cleanup of the analytics folder	6
Preparing your system for unplanned shut down	7
Adding database drivers	8
Adding database drivers to the Mediation tool	9
Adding database drivers to the Analytics server	10

Chapter 2. Configuring data mediation 11

Starting the Mediation tool	11
Creating an Operations Analytics Predictive Insights project	12
Adding a data source to a project	12
Adding a file system data source	13
Adding a database data source	19
Creating a model	20
Synchronizing with the data source	20
Selecting data.	21
Verifying and updating the time stamp	23
Creating filters	24
Unpegging	26
Validating a model	26
Model configuration examples for different time zone scenarios in source files	27
CSV file names and data are for different time zones	27
CSV file names and data have no time zone identifier	29
CSV files are for multiple time zones	30
Deploying a model	30
Creating a topic	32
Exporting a model	33
Importing a model	33

Chapter 3. Configuring security 35

Configuring security for Dashboard Application Services Hub	35
Securing sensitive cookies	35
Enable Transport Layer Security (TLS)	35
Disable auto-complete on the Dashboard Application Services Hub login panel	36
Configuring security for Tivoli Integrated Portal	36

Securing sensitive cookies	36
Enabling Transport Layer Security	37
Configuring certificates	37
Adding a signed certificate to the User Interface server	37
Adding a signed certificate to the Dashboard Application Services Hub server	38
Configuring LDAP authentication	38
Configuring the Operations Analytics Predictive Insights UI server to use LDAP authentication	39
Configuring OMNIBus Web GUI authentication against an LDAP directory	42
Configuring an SSL connection to the OMNIBus ObjectServer	48

Chapter 4. Configuring integrations with other products 51

Configuring integration with IBM Integration Bus Architecture	51
Setting up the integration script	52
Configuring real time metric collection and analysis.	53
Configuring backlog collection of metrics	55
Configuring integration with IBM Performance Management	58
Before you begin	58
Configuring an Operations Analytics Predictive Insights model for Performance Management data sources	58
Exporting a directory from the Operations Analytics Predictive Insights server	60
Mounting a directory on the Performance Management Server	60
Configuring the File Consumer program	61
Setting up the Event Integration Facility (EIF) Gateway	61
Updating the probe rules file	64
Configuring integration properties.	64
Tuning the Operations Analytics Predictive Insights server	65
Scheduling cleanup of Performance Management data files	66

Chapter 5. Configuring a cloned Operations Analytics Predictive Insights server 67

Chapter 6. Analyzing data 69

Setting the aggregation interval.	69
Configuring purge	70
Extracting data	70
Extracting data in backlog mode	71
Extracting data in switch mode.	72
Extracting data in steady-state mode	73

Extracting data in replay raw mode	73
Extracting data after Operations Analytics Predictive Insights restarts	75
Checking status	75

Chapter 7. Operations Analytics

Predictive Insights administration 77

Filtering alarms	77
Creating a topic	79
Starting Operations Analytics Predictive Insights components	80
Starting DB2	80
Starting the OMNIBus ObjectServer	80
Starting the Operations Analytics Predictive Insights analytics component	81
Starting Tivoli Integrated Portal	81
Starting Dashboard Application Services Hub	82
Checking that all Operations Analytics Predictive Insights components started correctly.	82
Stopping Operations Analytics Predictive Insights components	83
Stopping DB2.	84
Stopping Tivoli Integrated Portal	84
Stopping Dashboard Application Services Hub	84
Stopping the Operations Analytics Predictive Insights Analytics component	85
Backing up and restoring data	85
Changing passwords	86
Changing a Tivoli Integrated Portal user password	86
Changing a Dashboard Application Services Hub user password	86
Changing a Tivoli Integrated Portal administrator password	86
Changing the Tivoli Netcool/OMNIBus ObjectServer user password	87
Changing a Dashboard Application Services Hub administrator password	87

Changing the database user password	88
Changing the DB2 instance owner password	89
Changing the InfoSphere Streams administrator password	89
Changing the database data source password	90

Chapter 8. Event management

administration 93

Configuring alarm forwarding from the Analytics server to an OMNIBus ObjectServer	93
Displaying the Operations Analytics Predictive Insights columns in the Active Event List	94
Generating the OMNIBus interfaces file	94
Customizing the OMNIBus probe rules file.	95
OMNIBus probe tokens	97

Chapter 9. Reference 101

Scripts and utilities	101
admin.sh	101
collect.sh	111
start.sh	112
stop.sh	112
Properties	113
Component properties	113
System properties	124
Log files created	129
Analytics log file rotation	130
REST Interface	130
Aggregations	131
Metrics	132

Notices 137

Trademarks 141

Preface

The purpose of this guide is to help you administrate Operations Analytics Predictive Insights.

Provided in this guide is a reference section describing the Operations Analytics Predictive Insights command line interface (CLI) as well as the available configuration properties. The guide also contains instructions on how to perform common administration tasks.

Audience

The audience for this guide is the network administrator or operations specialist responsible for the administration of Operations Analytics Predictive Insights.

To administrate Operations Analytics Predictive Insights you must have an advanced understanding of the following subjects:

- Administration of the Linux operating system
- Administration of IBM InfoSphere Streams
- Administration of the DB2 database management system
- Tivoli Integrated Portal
- Operations Analytics Predictive Insights

Components

IBM® Operations Analytics Predictive Insights consists of four main components.

The IBM Operations Analytics Predictive Insights components are:

- **The Database component:** is used to store configuration data, metadata and metric data.
- **The Analytic component:** performs data mediation and processes incoming data to discover any anomalies that are present.
- **The UI component:** presents any discovered anomalies through the IBM Dashboard Application Services Hub application or the IBM Tivoli Integrated Portal application.
- **The Mediation Tool:** is used to configure a data source and the data model that Operations Analytics Predictive Insights will monitor.

Operations Analytics Predictive Insights documentation includes the following guides:

- Release notes
- Installation Guide
- Upgrade Guide
- Administration Guide
- Error Messages Guide

Chapter 1. Initial configuration of Operations Analytics Predictive Insights

After you install Operations Analytics Predictive Insights, you must perform configuration tasks to set up the system.

Configuring the User Interface

The Active Event List must be configured to display the Operations Analytics Predictive Insights User Interface.

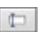
Adding the User Interface (UI) launch scripts

If you are using Operations Analytics Predictive Insights with OMNIbus WebGUI 8.1 FP1 or earlier, you must add the scripts that start the User Interface.

About this task

If you are using Operations Analytics Predictive Insights with OMNIbus WebGUI 8.1 FP2 or later, the launch scripts are automatically added during the installation of Operations Analytics Predictive Insights. In this scenario, you do not need to complete this procedure.

Procedure

1. Log in to your visualization application as an administrative user:
 - For Dashboard Application Services Hub, log in as ncoadmin.
 - For Tivoli® Integrated Portal, log in as tipadmin.
2. Click **Administration > Event management Tools > Menu Configuration**.
3. Select **alerts** and click **Modify**.
4. Select the **ViewChildAlarms...** and **ServiceDiagnosis...** scripts and move to the **Current®** items list. Use the arrow keys to position the scripts in your preferred location. If you want to rename the scripts, click the rename icon . In the **Label** field, update the script label. For example, you might want to change **ServiceDiagnosis...** to **Service Diagnosis**.
5. Click **Save**.

Note: If the Active Event List is already open, you must close and reopen it to see the new menu options.

Results

When you right-click one or more entries in the Active Event List, you see the **ViewChildAlarms...** and **ServiceDiagnosis...** scripts are available in the menu. To select multiple entries, hold down the **Ctrl** key and click the rows in the Active Event List. Select the **ViewChildAlarms...** to drill into an alarm, and select **ServiceDiagnosis...** to start the Operations Analytics Predictive Insights UI. By default, the target metric is displayed for each anomaly and the target metric is shown in bold.

Note: **ServiceDiagnosis...** is available for all analytics alarms and **ViewChildAlarms...** for consolidated analytics alarms.

Configuring user access to Operations Analytics Predictive Insights

To view anomaly information, you must grant users access to the Operations Analytics Predictive Insights User Interface.

Dashboard Application Services Hub

If you installed Operations Analytics Predictive Insights into Dashboard Application Services Hub, you must configure access to the Operations Analytics Predictive Insights User Interface (UI) for users of Dashboard Application Services Hub.

About this task

You must create users in OMNIBus WebGUI and grant the users access to the Operations Analytics Predictive Insights (UI). This topic provides a summary of the steps to create users and add users to groups in OMNIBus WebGUI. For more information, see *Administering users, roles, and groups* and *Creating and editing groups* in the Tivoli Netcool/OMNIBus Knowledge Center.

Procedure

1. Create users in OMNIBus WebGUI.
 - a. Log in to Dashboard Application Services Hub as an administrative user.
 - b. Click **WebSphere Administrative Console** under the **Console Settings** folder.
 - c. In the new page that is opened, click the **Launch WebSphere administrative console** button.

The WebSphere Integrated Solutions Console is launched in a new browser tab.
 - d. Log in using the Dashboard Application Services Hub profile username and password.
 - e. Click **Users and Groups**.
 - f. Click **Manage Users**.
 - g. Click **Create**.
2. Use one of the following methods to grant the WebGUI users access to the Operations Analytics Predictive Insights UI:
 - Add the new users to an existing group that has the netcool_rw role. The netcool_rw role grants users the rights to launch the **Service Diagnosis** and **View Child Alarms** menu options in the Active Event List and the Event Viewer. For example, complete steps a to g to add users to the Netcool_Omnibus_Admin group, which has the netcool_rw role by default.
 - a. In the WebSphere Integrated Solutions Console, click **Users and Groups** to expand this section.
 - b. Click the **Manage Groups** task.
 - c. Click the appropriate group, for example, **Netcool_Omnibus_Admin**.
 - d. Click **Members**.
 - e. Click **Add Users**.
 - f. In the Search for box, enter a user name or use '*' and click **Search**.
 - g. Select the user and click **Add**.
 - Create a group in OMNIBus WebGUI, add the Netcool_rw role to the group, and add users to the group. For information on how to configure groups in OMNIBus WebGUI, see *Creating and editing groups* in the Tivoli

Netcool/OMNIBus Knowledge Center. After you configure the group, complete the following steps to grant the group members rights to access the Operations Analytics Predictive Insights UI. Users can access the UI from the **Service Diagnosis** and **View Child Alarms** menu options in the Active Event List and the Event Viewer.

- a. Log on to the server on which you installed the Operations Analytics Predictive Insights UI as the user that installed the UI.
- b. Navigate to the <Liberty_Install_Home>/UI/bin directory. The default path for <Liberty_Install_Home> is /opt/IBM/scanalytics.
- c. Enter the following command to grant access to the Operations Analytics Predictive Insights User Interface:

```
./addAccess.sh <user | group> <user or group name>
```

For example, to grant access to a group called operators, enter the following:

```
./addAccess.sh group operators
```

Note: The addAccess.sh script authorizes users on the Liberty Profile but does not configure user roles in WebGUI.

What to do next

If you need to remove access for a specific user or group, use the following command:

```
removeAccess.sh <user | group> <user or group name>
```

For example:

```
removeAccess.sh group operators
```

Tivoli Integrated Portal

If you installed Operations Analytics Predictive Insights into a Tivoli Integrated Portal environment, you must configure access to the Operations Analytics Predictive Insights User Interface (UI) for users of Tivoli Integrated Portal.

About this task

Complete this task only if you installed Operations Analytics Predictive Insights into a Tivoli Integrated Portal environment.

Configuring OMNIBus WebGUI users

In order for users to be able to use the Active Event List to monitor events, they must be given permissions to access the WebGUI tools.

About this task

Complete this task only if you installed Operations Analytics Predictive Insights into a Tivoli Integrated Portal environment.

Procedure

- Follow these steps to assign WebGUI permissions to tipadmin:
 1. Log in to Tivoli Integrated Portal as tipadmin.
You can access Tivoli Integrated Portal in your browser, for example:
`https://<hostname>:16311/ibm/console`
 2. Click **Users and Groups > User Roles**.

3. Click **Search** and select **tipadmin**.
4. Select the following options and click **Save**:
 - **ncw_admin**
 - **netcool_rw**
5. Click **Logout**.
6. Log back in as **tipadmin**.
7. In the navigation pane, you can see extra entries that include **Administration** and **Availability**.
8. Operations Analytics Predictive Insights sends events to the Active Event List, which can be found at **Availability > Events > Active Event List (AEL)**.
- Follow these steps to assign WebGUI permissions to all users that are not tipadmin:
 1. Log in to Tivoli Integrated Portal as tipadmin.
 2. Change to **Users and Groups > Manage Users**.
 3. Enter the relevant user ID and click **Search**.
 4. Click the user that is displayed in the table.
 5. Click the **Groups** tab.
 6. Click **Add**.
 7. Click **Search** to see the list of groups.
 8. Assign the user to one of the following groups, as appropriate for the user.
 - Administrative Users: **Netcool_OMNIBus_Admin**
 - Ordinary Users: **Netcool_OMNIBus_User**
 9. Click **Add**.
 10. Click **Close**.

Note: To make the Operations Analytics Predictive Insights action script available you must have a user with **Netcool_OMNIBus_Admin** permissions.

Configuring roles and groups in Tivoli Integrated Portal

Tivoli Integrated Portal users must have the correct roles and groups to manage user access to Operations Analytics Predictive Insights.

About this task

Complete this task only if you installed Operations Analytics Predictive Insights into a Tivoli Integrated Portal environment.

To access the Operations Analytics Predictive Insights UI, users must be members of a group that is assigned an Operations Analytics Predictive Insights role, that is, `predictiveInsightsUser` or `predictiveInsightsAdmin`. The difference between both roles is that the `predictiveInsightsAdmin` user can export data in both streamlating and CSV formats while the `predictiveInsightsUser` can export only in CSV format.

The `predictiveInsightsUser` and `predictiveInsightsAdmin` roles, and the `predictiveInsightsUsers` and `predictiveInsightsAdmins` groups are added automatically when Operations Analytics Predictive Insights is installed.

Procedure

As tipadmin add the ncw_user role to the predictiveInsightsUsers group. For information about how to assign roles to groups, open the Tivoli Integrated Portal online help and change to **Tivoli Integrated Portal > Administrative settings > Working with roles**

Configuring user accounts in Tivoli Integrated Portal

To have access to the Operations Analytics Predictive Insights User Interface, a user must be assigned to the Operations Analytics Predictive Insights roles or groups.

Before you begin

Complete this task only if you installed Operations Analytics Predictive Insights into a Tivoli Integrated Portal environment.

The Tivoli Integrated Portal Administrator can add users through the **Users & Groups** function in the navigation pane.

About this task

For example, to assign a user to the predictiveInsightsUsers group:

Procedure

1. Log in as the Tivoli Integrated Portal admin user. For example, log in as tipadmin.
2. Click **Users & Groups > Manage Groups**.
3. Select the user that you want to give access to the Operations Analytics Predictive Insights UI.
4. Assign the group predictiveInsightsUsers to the selected user.

Setting the DB2 log file size

Instructions on how to set the DB2 log file size to suit your Operations Analytics Predictive Insights installation.

About this task

The DB2 log file size settings described in this section have been optimized for a Operations Analytics Predictive Insights installation with a KPI count of 100,000.

Procedure

1. As the Operations Analytics Predictive Insights database owner, typically db2inst1 log into the Operations Analytics Predictive Insights database server.
2. Update Operations Analytics Predictive Insights database configuration by setting LOGBUFSZ to be 4096.
You can get details on each DB2 configuration parameter by opening the DB2 Knowledge Center and searching for the relevant configuration parameter.
3. Update Operations Analytics Predictive Insights database configuration by setting LOGFILSIZ to be 8102.
4. Update Operations Analytics Predictive Insights database configuration by setting LOGPRIMARY to be 100.

5. Update Operations Analytics Predictive Insights database configuration by setting LOGSECOND to be 75.

What to do next

To optimize IO performance, update the transaction logs to use a disk that has low IO and is separate to the database. For details on how to update the transaction logs, open the DB2 Knowledge Center and search for the section *newlogpath* - *Change the database log path configuration parameter*.

Configuring DB2 to allow for the processing of many metric groups

DB2 must be configured so that it is able to support the processing of metric data for more than fifty groups.

About this task

If Operations Analytics Predictive Insights is used to monitor more than fifty metric groups, that is, your data sources contain in excess of fifty metric groups, you must configure DB2 to be able to support that number.

DB2 is configured to support an increased number of metric groups by changing the maximum number of active applications allowed for that instance.

To configure DB2:

Procedure

1. As the as the Operations Analytics Predictive Insights database owner, typically db2inst1 log into the Operations Analytics Predictive Insights database server.
2. Run the following command:

```
db2 get db cfg for SCAPIDB | grep MAXAPPLS
```

Max number of active applications (MAXAPPLS) = AUTOMATIC(92)

You can see from this the current level of supported active applications.
3. Increase the number of supported active applications to 400 using the command:

```
db2 update db cfg for SCAPIDB using MAXAPPLS 400
```

Scheduling cleanup of the analytics folder

To conserve disk space, you must schedule maintenance jobs to compress the data files in the analytics folder, \$PI_HOME/var/spool/topics, and delete the compressed files later.

About this task

The purpose of this task is to schedule maintenance jobs that:

- After one day, compresses all of the CSV files that are created when data is extracted or merged
- After 60 days, deletes the compressed file

Note: This task describes how to schedule the maintenance job with cron but you can use any scheduling application.

Procedure

1. Log in to the Analytics server as the administrative user, typically scadmin.
2. Enter the following command to edit the crontab file:

```
crontab -e
```

3. Add the following lines to the crontab file:

```
5 0 * * * export PI_HOME=/opt/IBM/scanalytics/analytics;  
/bin/find $PI_HOME/var -type f \( -name *.csv -o -name *.csv.bad -o -name *.ctrl \) -mtime +1  
  
15 0 * * * export PI_HOME=/opt/IBM/scanalytics/analytics;  
/bin/find $PI_HOME/var -type f \( -name *.csv.gz -o -name *.csv.bad.gz -o -name *.ctrl.gz \) -mtime +1
```

If the Analytics component is not installed in the default directory, substitute /opt/IBM/scanalytics/analytics with the correct directory path.

In this example, the maintenance jobs are scheduled to run at 12:05 AM and 12:15 AM each night.

4. Save the crontab file.

Preparing your system for unplanned shut down

It is good practice to prepare your system for the event of a forced restart.

About this task

It is good practice to configure your Operations Analytics Predictive Insights system to use the following reboot commands. The following configuration items ensure the set of hosted Operations Analytics Predictive Insights components restart correctly should your system experience an unplanned shut down.

Procedure

- Before restarting the Database server, prepare for restart of the Database component:

1. Add a database start command in the crontab of the DB2 owner:

```
@reboot source /home/db2inst1/sqllib/db2profile; db2start
```

For example, as db2inst1 run

```
crontab -e
```

and insert the line

```
@reboot source /home/db2inst1/sqllib/db2profile; db2start
```

2. Add a database start command in the crontab of the DB2 DAS owner if applicable:

```
@reboot source /home/dasusr1/das/dasprofile; db2admin start
```

For example, as dasusr1 run

```
crontab -e
```

and insert the line

```
@reboot source /home/dasusr1/das/dasprofile; db2admin start
```

- Before restarting the Analytics server, prepare for restart of the Analytics component:

Add a Streams start command in the crontab of the InfoSphere Streams owner:

```
@reboot source ~/.bashrc; streamtool stopinstance -i spl --force; streamtool  
rminstance -i spl --noprompt
```

For example, as scadmin run

```
crontab -e
```

and insert the line

```
@reboot source ~/.bashrc; streamtool stopinstance -i spl --force; streamtool  
rinstance -i spl --noprompt
```

- Before restarting the OMNIBus server, prepare for restart of the object server:
Add an object server start command in the crontab of the OMNIBus owner:

```
@reboot nohup /opt/IBM/tivoli/netcool/omnibus/bin/nco_objserv
```

For example, as scadmin run

```
crontab -e
```

and insert the line

```
@reboot nohup /opt/IBM/tivoli/netcool/omnibus/bin/nco_objserv
```

- Before restarting the UI server, prepare for restart of the UI component

Dashboard Application Services Hub

If you are using Dashboard Applications Services Hub, on the Jazz for Service Management server, add the following command to the crontab of the user that installed Jazz for Service Management:

```
@reboot source ~/.bashrc; /opt/IBM/JazzSM/profile/bin/startServer.sh server1
```

On the Operations Analytics Predictive Insights UI server, which may be the same as the Jazz for Service Management Server, add the following command to the crontab of the user that installed the UI component:

```
@reboot source ~/.bashrc; sleep 2m; /opt/IBM/scanalytics/UI/bin/pi.sh -start
```

Note: For a distributed configuration where Dashboard Applications Services Hub and the Operations Analytics Predictive Insights UI component are on separate servers, if the Operations Analytics Predictive Insights UI server comes online before the Dashboard Applications Services Hub server, then federation into Dashboard Applications Services Hub may not come up correctly. In this case, a further restart of the Operations Analytics Predictive Insights UI may be necessary:

```
/opt/IBM/scanalytics/UI/bin/pi.sh -restart
```

Tivoli Integrated Portal

If you are using Tivoli Integrated Portal, add the following command to the crontab of the user that installed the application:

```
@reboot /opt/IBM/tivoli/tipv2/profiles/TIPProfile/bin/startServer.sh server1
```

For example, if you are using Tivoli Integrated Portal, as scadmin run

```
crontab -e
```

and insert the line

```
@reboot /opt/IBM/tivoli/tipv2/profiles/TIPProfile/bin/startServer.sh server1
```

Adding database drivers

If a database data source is not a DB2 database, you must add the required database driver to the Mediation tool and the Analytics server. You do not require any database drivers if your data source is a file system.

Adding database drivers to the Mediation tool

To connect the Mediation Tool to a database other than DB2, the driver for which is preinstalled, you must add the appropriate database driver to the Mediation Tool.

Before you begin

By default, the driver properties for the following databases are preconfigured in the Mediation Tool: Oracle, MySQL, Sybase, and Microsoft SQL. After you add the driver for any of these databases to the Mediation Tool, you can connect the Mediation Tool to the database.

If you add a database driver whose properties are not preconfigured in the Mediation Tool, you must complete the steps after the procedure to configure the driver properties.

About this task

After you add a database driver to the Mediation Tool and, if necessary, configure the driver's properties, you can select the driver from the drop-down list when you configure the data source.

Note: The complete JDBC driver that you want to add must be contained in a single JAR file.

Procedure

1. Ensure that the database driver is saved on your local drive.
2. To start the Mediation Tool, go to the `$PI_HOME/bin` directory and enter the following command:

```
./mediationtool.sh
```

On startup, you are asked to specify a workspace location.

This workspace is a folder in which your mediation files are kept. You can keep it local to the Operations Analytics Predictive Insights Mediation Tool installation or in a directory where you have write access.

Note: To ensure that the workspace is not deleted if you uninstall the Mediation Tool, do not set the workspace location to the folder in which the Mediation Tool is installed. By default, the Mediation Tool is installed in `/opt/IBM/scanalytics/mediationtool/eclipse`.

3. Click **Window > Preferences > Predictive Insights > New**.
4. Select the driver location and click **Apply**. Eclipse requests that it is restarted.
5. Click **OK**.

What to do next

If you added a driver other than a Oracle, MySQL, Sybase, or Microsoft SQL driver, you must complete the following steps to configure the driver's properties for the Mediation Tool.

1. Create a `driver.properties` text file in the workspace location that you specified when you started the Mediation Tool. For example, `/eclipse/workspace/driver.properties`.
2. For each database driver, add the following lines to the `driver.properties` file:

```
<driver_name>.caption=Undefined Driver
<driver_name>.url=jdbc://<host>:<port>/<database> <other_properties>
<driver_name>.port=<port>
```

where:

- driver_name is the name of the JDBC database
- driverhost is the host name of the system where the database is located and is mandatory
- port is the port number on which the database listens and is optional
- database is the name of the database and is optional
- other_properties are custom properties for individual drivers

Example:

```
com.microsoft.sqlserver.jdbc.SQLServerDriver.caption = Microsoft JDBC
SQL server
com.microsoft.sqlserver.jdbc.SQLServerDriver.url = jdbc:sqlserver://
{host}:{port};databaseName={database}
com.microsoft.sqlserver.jdbc.SQLServerDriver.port = 1433
```

Note: If a <driver_name>.url field in the driver.properties file contains properties other than host, port, and database, these additional properties appear in the **URL** field in the Mediation Tool. You must click the **Custom** check box beside the **URL** field and edit the URL to remove the additional properties.

3. Restart the Mediation Tool to make the driver available when you connect to the data source.

Adding database drivers to the Analytics server

For extraction to take place from a database other than a DB2 database, you must add the database driver to the \$PI_HOME/lib directory before you start Operations Analytics Predictive Insights.

Procedure

1. Log in to server on which your analytics instance is installed.
2. Ensure that your database drivers are copied to the same server.
3. Place the driver files in the \$PI_HOME/lib directory.

Where \$PI_HOME is the installation location of the analytics component.

Chapter 2. Configuring data mediation

Data mediation is the process of collecting data from various sources and converting the data into a format that Operations Analytics Predictive Insights can read and understand.

You must complete the following steps to configure data mediation.

1. Start the Mediation tool.
2. Create a project.
3. Add one or more data sources.
4. Create a model.
 - a. Synchronize with the data source
 - b. Select the data elements that you want to form your model
 - c. Configure the model by editing the Group Definition, Attributes, and Metrics
 - d. Save and validate the model.
5. Deploy the model.

Each of these steps is described in detail in the following sections.

Starting the Mediation tool

Start the Operations Analytics Predictive Insights Mediation Tool by running the Eclipse application.

Before you begin

You cannot start the Mediation tool successfully from the terminal window that you use to install the Mediation tool. You must start the Mediation tool in a new terminal window so that the environment variables that the Mediation tool requires to run are set.

Procedure

1. As scadmin, log on to the server on which you installed the Operations Analytics Predictive Insights Mediation Tool.
2. Navigate to the directory `$PI_HOME/bin`.
3. Run `./mediationtool.sh`

At startup, you are asked to specify a workspace location. This workspace is a folder in which your mediation files are kept.

Note: If you upgraded from an earlier version of Operations Analytics Predictive Insights and you specify a workspace location that was used by a previous version of the Mediation Tool, you see the following message: Workspace was written with an older version of the product and will be updated. Updating the workspace can make it incompatible with older versions of the product. You can click **OK** to use the existing workspace.

Creating an Operations Analytics Predictive Insights project

Instructions on how to create a new project with the Operations Analytics Predictive Insights Mediation Tool.

About this task

A project is a container that store details of data sources and models configured in the Mediation Tool.

Procedure

1. Open the Operations Analytics Predictive Insights Mediation Tool.
2. When you first open the Mediation tool, you are presented with a welcome screen. Close the welcome tab.
If you want to restore it later, you can select **Help > Welcome**.
3. Click **File > New > Project**. The **New Project** dialog box is opened, and you must choose the wizard you are going to use to create the new project element.
4. In the New Project dialog box, select the **Predictive Insights Wizards > Predictive Insights Project** option and click **Next**.
5. Enter a name for your new project into the **Project name** field.
6. Enter a location for your new project within the **Location** field or use the default.
7. Click **Finish**.
8. Select **Yes** if prompted to associate the project with a **Predictive Insights** perspective.

Results

You now have a new empty project.

Adding a data source to a project

You can choose from two types of data source for a model, a file-based data source or database data source.

About this task

The Operations Analytics Predictive Insights Mediation tool provides a wizard to help you add and configure a data source. A data source and a model configuration combine to form a Operations Analytics Predictive Insights model.

While you must add separate data sources to your project for separate databases and file systems, you must also make sure that each data source only covers one timezone. Multiple timezones per data source are not permitted. If a database or filesystem covers multiple timezones, you must add a data source for each timezone that is covered to include all data.

Procedure

1. In the Operations Analytics Predictive Insights Mediation tool, right-click your new project and select **New > Predictive Insights Data Source**. The **New Predictive Insights Data Source** dialog is displayed.
2. Select the project in which you want the model to be saved.
3. Enter a name for the data source configuration.

The file name must retain the .pamodel file extension.

4. Click **Next**.
5. Select the data source type, which can be either **Database** or **File System**.
6. The **XML Encoding** drop-down menu sets the encoding that is applied to the configuration file you are creating.
7. Choose the data format for the data source. The default format requires that the data source has metric names located in column headers. If the metric names are contained in rows of data, click the radio button beside **Metric names in rows**.
8. Click **Finish**. The Operations Analytics Predictive Insights Mediation tool now displays the connection details for the new data source.

What to do next

You now have a data source item for your model. The appropriate editor is opened in order for you to complete the connection to your database or file system.

If you are adding a file system as a data source, see “Adding a file system data source”

If you are adding a database as a data source, see “Adding database drivers” on page 8 and then “Adding a database data source” on page 19.

Adding a file system data source

When you added a data source for your project, you selected the file system type.

Before you begin

If you installed the Mediation Tool on a different system to the Analytics component, do the following:

- Set up an environment variable to mask the differences between the file system data source path on the server that hosts the Mediation Tool and the file system data source path on the server that hosts the Analytics component.
- Copy CSV files representing at least one full day of data to the server that hosts the Mediation Tool so you can preview the configuration before and during deployment.

The **File Path** parameter, which you set when you add a file system data source, supports the use of environment variables. If the server on which you are running the Mediation Tool is the same as the Analytics server, you do not need to create an environment variable.

One or more path environment variables can be used in the **File Path** element of a Operations Analytics Predictive Insights data source for files extraction. Any environment variables that are entered into the **File Path** field are resolved against the environment of the operating system that hosts the UI or the Extractor instance.

If the Operations Analytics Predictive Insights Mediation tool is running on Windows:

1. Create an environment variable, such as, %MY_AGENT%.
2. Set the environment variable to a valid local path, which points to the location of the files for extraction.

3. Restart the Operations Analytics Predictive Insights Mediation tool in order to make the new variable visible to the UI process.

If the Operations Analytics Predictive Insights Mediation tool is running on Linux:

1. Create an environment variable, such as, **\$MY_AGENT**.
2. Set the environment variable to a valid local path, which points to the location of the files for extraction.
3. Add the environment variable to the user's .bashrc file.
4. Restart the Operations Analytics Predictive Insights Mediation tool in order to make the new variable visible to the UI process.

For the Extractor, which can run only on Linux, as it must run on the same server as the Analytics component, the same applies: The **\$MY_AGENT** variable must be properly exported to make it visible to the next running Extractor instance.

Note: It is assumed that the file system or a copy of that file system exists locally in the directory corresponding to that path.

About this task

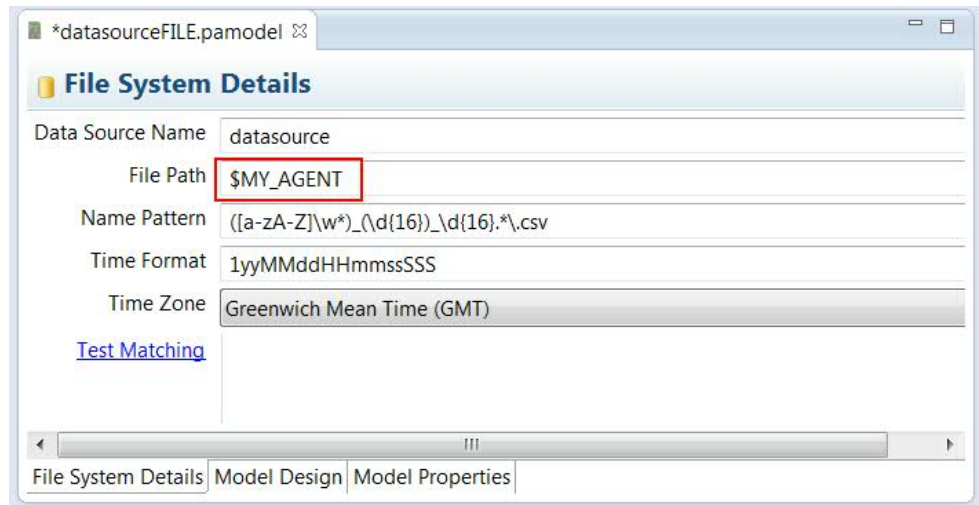
To create a file system data source configuration, you must enter details of that file system in the **File System Details** tab. Note the following when you update the **File System Details** tab:

Note: For information on the rules that are related to file naming, csv file format and file content, see the section “Rules for using CSV data sources” on page 16

- Set the **File Path** to the path that contains the file system data source.

Enter the path in Unix/Linux notation, that is, without drive letter and with forward slashes. Use Unix/Linux notation regardless of whether you are creating your configuration on a Windows server or a Linux server. It is assumed that the file system or a copy of that file system exists locally in the directory corresponding to that path. You can use an environment variable to make the **File Path** setting more portable.

Note: For file based data sources, you can create only one metric group per table within that data source. Creating multiple metric groups per table results in a deadlock situation. If you need to have multiple metric groups for a table, you must deliver the source file data to separate file system directories, and create one data source definition per directory.



For both Windows and Linux, if you are referring to an environment variable in the **File Path** field, use Unix/Linux notation, for example, \$MY_AGENT.

- Use the **Name Pattern** to enter a regular expression that filters the files that are contained in the file path specified. The regular expression must have at least two capturing groups, the second of which matches a timestamp or write time. The default of `([a-zA-Z]\w*)_(\d{16})_(\d{16}).*\'.csv` has three matching groups. The first matching group is expected to be the metric group name, the second is the start time, and the third is the end time; for example, `Disk_1130421121500000_1130421123000000.csv`

Note: For more information on file patterns, see “Example file naming patterns” on page 18

- Set the **Time Format** to a string that matches the format of the date and time used in the CSV file names, for example, `1yyMMddHHmmssSSS`

The formats allowable are defined by the SimpleDateFormat Java™ class. For more information about the SimpleDateFormat Java class, see <http://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html>.

- Use the **Time Zone** to specify which time zone the data source is from. You must consider the time zone of the time stamps in both the source file names and the data:
 - If the time stamps in the file names and data are for the same time zone, select that time zone.
 - If the time stamps in the file names and data are for different time zones, the time stamps in either the file names or the data must include a timezone identifier. In the **Time Zone** field, select the time zone that is not identified in either the file names or the data time stamps. For example, if the time zone in the file names is EST but the time zone is not shown and the time zone in the data is CST and the time zone is shown, set the **Time Zone** field to EST.
 - If the time zone in the file name and data are for different time zones and the time zone is shown in both time stamps, set the time zone to that shown in the file name time stamps.

For examples on how to configure Operations Analytics Predictive Insights for different time zones, see “CSV file names and data are for different time zones” on page 27

All data extracted by Operations Analytics Predictive Insights is loaded into the Operations Analytics Predictive Insights database in UTC time. Therefore, if you can set the time zone of the source data, it is recommended that you set it to UTC for both the file names and data.

If the timezone the data is coming from applies DST, timestamps in both file names and data must include a time offset or timezone. If some of your file names and data timestamps already have timezone or time offset information, you can still set the timezone for the data source, but Operations Analytics Predictive Insights gives priority to the time related information contained in the data.

Only one timezone is supported per model. If the CSV file names or the data within the files represent different time zones, the files for each different time zone must be in separate directories and you must create a separate model for each set of source files.

- Click **Test Matching** to confirm the file set that matches your expression.

Procedure

1. Update the fields within the **File System Details** tab.
2. Click **File > Save**.
3. Check for any issues with your data source by opening the **Problems** tab. As you need to complete further steps to make that data source valid, an error message is displayed within the **Problems** tab. The error message disappears when you complete the data source and model configuration.

Related reference:

Rules for using CSV data sources

Summarizes the set of rules and limitations that apply to the use of CSV files as a data source.

Rules for using CSV data sources

The following rules apply when the data source is a CSV filesystem.

Rules for CSV file naming

- The file name must begin with a table or group reference, followed by a start time, and end in .CSV.
- The start time in the file name must be at or before the first time stamp in the file.
- The end time is optional and can be at or after the last time stamp in the file. If the end time in the file name is at the same time as the start of the next interval, this end time must be after the latest time stamp in the file. For example if the end time of 10:05 is at the start of the next interval, the latest time stamp in the file can be 10:04:59.
- If the time zone of the file name is different to the time zone of the file content, the time zone must be explicitly defined in at least one of either the file name or file content time stamps. The following examples show the time zone defined in the file name:

CPULOAD_2013-07-17-00-00EST_2013-07-17-00-15EST.csv

CPULOAD_2013-07-17-00-00+0300_2013-07-17-00-15+0300.csv

- A sample file name with 15-minute data is:

CPULOAD_2013-07-17-00-00_2013-07-17-00-15.csv

The file name shows that:

- The file contains the CPULOAD source table

- The first timestamp in the file is July 17, 2013 at 00:00
- The last timestamp is earlier than July 17, 2013 at 00:15

Note: For more information about how to create file naming patterns, see “Example file naming patterns” on page 18

Rules for CSV file content

The following rules and limitations apply to CSV file content:

Header line

- The first line of each CSV file must contain a comma-separated list of headers for the columns in the file. For example:
`#timestamp,resourceName,AvgLoad,AvgPercentMemoryUsed`

Column format

- Date and time must be in the same column. Preferably, the time stamp is in UTC or otherwise contains the time zone information that is defined in Java SimpleDateFormat, <http://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html>. By supplying the time zone information in the data, it can be adjusted for Daylight Saving Time. If the date and time is not in the same column, you must preprocess those files to combine those two columns. Alternatively, you can join the two columns as you reexport the CSV files.
- Resource names can be made from several columns (in the Mediation Tool). One part of the resource name must be assigned to an attribute named 'Node'. If a null or empty value is seen for the Node attribute, that row of data is discarded. Null values that are in other parts of the resource key are not recommended, but do not cause the data to be discarded, instead, the resource name is made up of only the non-null parts of the key.

Characters

- Files must be in CSV format. Ensure that any text fields that contain a comma are surrounded in quotation marks, for example, Supplier column contains 'International supplies, Inc.'
- Metric values with a decimal point cannot use a comma as the decimal point, they must use a full stop character, “.”
- If a field is surrounded by double quotation marks, it must not contain double quotation marks, for example, "International "supplies" Inc"

Time

- When processing data in steady state mode, Operations Analytics Predictive Insights expects a single CSV file for each interval. The first interval for each hour starts at 00 minutes and increment at multiples of the interval. For example, if the aggregation interval is 15 minutes, the intervals are 00 to 15 minutes, 15 to 30 minutes and so on. The time stamps in the CSV file cannot span intervals. For example, for the interval 15:00 to 15:15, the earliest time stamp in the file must be 15:00 or later and the latest time stamp in the file must be 15:14:59.999 or earlier.
- When processing data in backlog mode, files can contain data that spans multiple time intervals. However, if the file content spans multiple intervals, then the data must be in chronological order. For performance reasons, ensure that file does not contain more than one day data.

- The end time of the data in the file must be less than the end of the interval, it cannot be equal to the end of the interval. For example, a 5-minute interval with start time of 10:00 and end time of 10:05 can have a last time stamp of 10:04:59.
- Files must be delivered to the source location for the extractor before the latency time set for the extractor expires. The default latency is the same period as the `system.aggregation.interval`. The minimum latency is 1 minute.

Examples of valid CSV file formats

The following example illustrates a CSV file format with data and time stamp in epoch format. Date and time must be in the same column.

```
#timestamp,resourceName,AvgLoad,AvgPercentMemoryUsed
1361523600000,"resource1",6,21.900673
1361523600000,"resource2",0,45.12558
1361523600000,"resource3",12,20.727364
1361523600000,"resource4",5,23.801073
```

The following example illustrates a CSV file format where data and time stamp are a string. Date and time must be in the same column.

```
#Timestamp, ResourceId, Metric_0
201304160815,ResourceId_1,0.0110
201304160815,ResourceId_2,0.0110
201304160815,ResourceId_3,0.0110
```

The following example illustrates a CSV file format where data and time stamp are a string, but not at the first position. The first column shows that fields can contain white space.

```
#Device,Parent Device,Sensor,Location,Time,Value
device 1 complex name,127.0.0.1,Sensor1,11/01/2014 16:35,46
device 1 complex name,127.0.0.1,Sensor1,11/01/2014 16:35,61
device 1 complex name,127.0.0.1,Sensor1,11/01/2014 16:39,46
device 1 complex name,127.0.0.1,Sensor1,11/01/2014 16:40,61
device 1 complex name,127.0.0.1,Sensor1,11/01/2014 16:40,46
```

The following example shows a CSV file where the data is in skinny format. In this format, metric names are contained in rows of data underneath a header row.

```
Timestamp,Node,SubResource,MetricName,MetricValue
2015-04-10 04:16:13.0,server1,subresourceA,M1,7003
2015-04-10 04:16:13.0,server1,subresourceA,M2,6683
2015-04-10 04:16:13.0,server1,subresourceA,M3,1041
2015-04-10 04:16:13.0,server1,subresourceB,M1,7643
```

Example file naming patterns

Naming patterns use regular expressions to identify the files that form part of the data source.

Name Pattern:

The following is an example of a naming pattern:

```
(.+)_(\d{4}\-\d{2}\-\d{2}\-\d{2}\-\d{2})_((\d{4}\-\d{2}\-\d{2})\-\d{2})\-\d{2}).*\..csv
```

The first set of parenthesis denotes the source table name.

`(.+)` = CPULOAD

The second and third set of parenthesis denotes the date and time.

(\d{4}\-\d{2}\-\d{2}\-\d{2}\-\d{2}) = 2013-07-17-00-00

Example - File name pattern 1

Sample file name:

CPUload__20130321.0015+0000__20130321.0030+0000_1363911769995.csv

The file name defines that the file contains data having:

- Source table name CPUload.
- The first timestamp in the file is March 21 2013 at 00:15.
- The last timestamp is earlier than March 21 2013 at 00:30.
- Timezone +0000
- **Name Pattern:** `(.+)__(\d{8}\.\d{4}\+\d{4}).*__(\d{8}\.\d{4}\+\d{4}).*.csv`
- **Time format:** `yyyyMMdd.HHmZ`

Example - File name pattern 2

Sample file name:

ITM1_134654223223__CPUload__20130901.2345+0300__20130902.0000+0300.csv

The file name defines that the file contains data having a source table with the name CPUload.

- **Name Pattern:** `.*__([a-zA-Z]\w*)__(\d{8}\.\d{4}\+\d{4}).*__(\d{8}\.\d{4}\+\d{4}).*.csv`
- **Time format:** `yyyyMMdd.HHmZ`

Adding a database data source

When you added a data source for your project, you selected the database type.

Before you begin

Operations Analytics Predictive Insights supports relational databases that have a JDBC driver. Operations Analytics Predictive Insights must be able to read time stamp, resource key, and metrics from any database set as a data source. The user that you employ to access the database must have permissions to list all tables of interest.

Operations Analytics Predictive Insights does not support normalized databases (normalized tables or star schemas) currently, that is, any database that contains foreign references for resource keys. If your database has this format, you can create a database view, in order for Operations Analytics Predictive Insights to read the database. Contact your database Administrator for information on how to create a database view.

About this task

When you add a database data source configuration, you must enter the details of the database in the **Connection Details** tab. Note the following when you update the **Connection Details** tab:

- If you are connecting to any of the databases, that are listed in the **Driver** drop-down menu, apart from DB2, you must first add the database's driver to the Mediation tool. If you are connecting to a database whose driver is not listed

in the **Driver** drop-down menu, you must first configure the driver's properties and then add the driver to the Mediation tool. For more information, see Adding Database drivers to the Mediation tool.

- Use the **Time Zone** to specify which time zone the data source is from. For example, Asia/Seoul - Korea Standard Time.
- If you are creating a MySQL connection, you must set the **Schema** to the database name. MySQL has no concept of schemas.
- The **Password** is not stored so you need to enter it when you test the connection or when you synchronize with the data source.
- **Prepared Statements** are precompiled SQL statements that Operations Analytics Predictive Insights can use to query the source database. By reducing the need to compile SQL statements, prepared statements save on CPU usage. However, in large data warehouse environments, prepared statements can cause a non-favorable index to be chosen, resulting in reduced disk I/O performance and increased query times.

Procedure

1. Update the fields within the **Connection Details** tab.
2. Click **File > Save**.
3. Check for any issues with your data source by opening the **Problems** tab. As you need to complete further steps to make that data source valid, an error message is displayed within the **Problems** tab. The error message disappears when you complete the data source and model configuration.

Creating a model

The purpose of creating a model configuration is to identify from the available data that which you want Operations Analytics Predictive Insights to analyze.

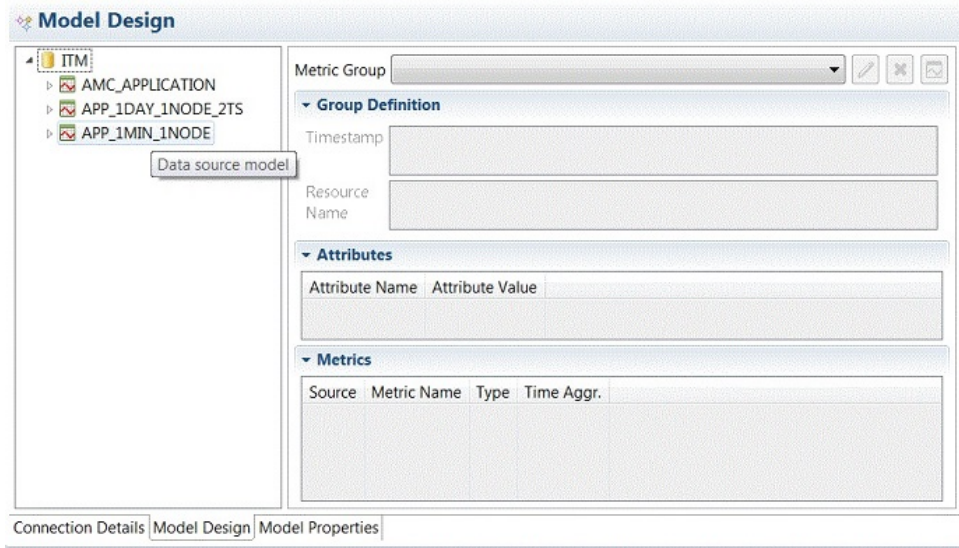
The Operations Analytics Predictive Insights Mediation tool can synchronize with the data source and display the available data. If it is a database, you are synchronizing with, the tool shows the tables that are contained within that database. If it is a file system data source that you are synchronizing with, the tool shows the file names and fields within that file system. You can then pick and choose from the available data items to create your model configuration.

Synchronizing with the data source

The first step in creating your model configuration is to discover the available data.

Procedure

1. Select the data source within the **Predictive Insights Navigator** tab for which you want to create a model configuration. The data source **Connection Details** are displayed.
2. Enter the data source connection **Password**.
3. Select the **Model Design** tab.



The model design tab has two parts, the data source tree view, and the editable properties of the selected source.

4. Right click the data source icon, which is the top level of the data source tree view, and click **Synchronize Schema**. The Operations Analytics Predictive Insights Mediation tool synchronizes with the data source and display all available tables as a list in the **Model Synchronization Preview** dialog.
5. Select the tables that you want to form part of your model.
6. Click **OK**.

The **Model Design** tab now contains the list of tables that you chose.

What to do next

You can expand each table displayed in the **Model Design** tab to see the set of columns they contain. You can also view the table contents by right clicking the table icon and selecting **Show Table Content**. The table contents are displayed in the **Data Extraction Preview** tab.

The **Data Extraction Preview** is provided so you can choose the parts of the various available tables to form your new model configuration.

Selecting data

You must select the items in each data source that you want Operations Analytics Predictive Insights to extract and analyze.

About this task

You use the **Model Design** tab to select the items in each data source that you want Operations Analytics Predictive Insights to extract and analyze.

Procedure



1. From the set of tables you selected and that are displayed within the **Model Design** tab, drag the tables that you want to form part of your model into the **Metric Group** field.
2. For each table you select, update the following properties if necessary:

- a. **Metric Group:** When you drag a table into this field, you are prompted to enter a name for the metric group. You can accept the default name, which is based on the table name, or type a new name. You cannot have two metric groups with the same name within your model.
- b. **Group Definition:** The Timestamp and Resource Key fields are populated with columns by default. You can right-click these fields and remove the columns, and drag columns into these fields.
- c. **Metric Name** You see this field only if you selected **Metric names in rows** for the Data Source Format when you added the data source. The Mediation Tool populates the field with a default column if it finds one that is suitable. You can remove the default column by right-clicking, and drag a replacement column.
- d. **Metric value** You see this field only if you selected **Metric names in rows** for the Data Source Format when you added the data source. The Mediation Tool populates the field with a default column if it finds one that is suitable. You can remove the default column by right-clicking, and drag a replacement column.
- e. **Attributes:** The Attributes field must contain the columns listed in the **Group Definition>Resource Key** field. Otherwise, an error is displayed in the Problems tab. You can rename attributes, but one attribute must be named Node. You can also add Attributes from a table in the data source. To add an attribute, drag a column from the table to the Attributes field or right-click a column in the table and select **Add Attribute**.

Note: Each attribute present in the model can be added to the alarm information displayed in the Operations Analytics Predictive Insights User Interface. By default, the Node attribute is shown in the alarm information. You can set the resource by which alarm consolidation occurs. Consolidation by node is a means of logically grouping alarms under one alarm based on a shared resource. To set node consolidation within your metric group, you must add the resource by which you want to consolidate to the **Attributes** tab and set the **Attribute Name** to **Node**. If multiple anomalies are generated for the same resource in a given interval, consolidation by that resource means that these anomalies are nested under one parent alarm in the Active Event List.

- f. **Metrics:** The set of metrics you want within your model. When you add the metric group, the **Metrics** field is populated with a default set of metrics. The process used to create the default set of metrics depends on the format of the data source.
 - If you selected **Metric names in column headers** for the Data Source Format, the default set of metrics is taken from the metric names in the column headers. You can drag further metrics from a table in the data source into the **Metrics** field.
 - If you selected **Metric names in rows** for the Data Source Format, the Mediation tool queries the metric name column in the first 1,000 rows of data and adds each unique metric name to the **Metrics** field. If the query finds a colon in a metric name, the Mediation Tool displays an error stating that a metric name is invalid because it contains a colon, which is the default data source delimiter. You must change the default data source delimiter to one or more characters that is not in any metric name. To change the default data source delimiter, right click the data source icon, which is the top level of the data source tree view, and click **Set Data Source Delimiter**.

If the data source has metric names that are not found by the query within the first 1,000 rows of data, you must manually add the metric names to the **Metrics** field. To manually add the metric names, do either of the following:

- To add metrics individually, choose the **Add metric** icon . Type a metric name that matches exactly the metric name in the data source.
- To add multiple metrics simultaneously, first add the metric names, delimited with carriage returns, to a .TXT or .CSV file. Then, click the **Add metrics from file** icon  and select the file. Any metric name in the file that exists in the **Metrics** field is not added again.

To remove a metric from the **Metrics** field, right-click the metric and choose **Remove Metric**.

3. Click **File > Save**.

What to do next

1. To confirm that the model contains all the elements that you require, click the **Data Extraction Preview** tab. The **Data Extraction Preview** tab shows a sample of the data that will be extracted based on the model you created. All timestamps in the extraction preview are shown in UTC time. When it begins to extract data, Operations Analytics Predictive Insights converts all source data to UTC.

If the value of a column in the model's Resource key field is blank in any of the sample rows of data retrieved from the data source, the following message is displayed:

Data extraction for {dasourceName}.{table_or_metric_grp} failed

Data for a row that has a blank resource key column is omitted from the preview but you can proceed with the validation and deployment of the model.

2. Validate the model design by checking for the presence of errors, which are listed in the **Problems** tab.
3. For file data sources:
 - a. Check the input directory, which is contained within the directory as specified by the **File Path** parameter, for the presence of any *.bad files. The *.bad files show which rows of the source file will not be accepted by the extractor, for example, text present in a number field.
 - b. Analyze the *.bad files to decide whether the metric group definition needs to be adjusted.
 - c. Remove the *.bad files if you adjusted the model and run the **Data Extraction Preview** again.

Verifying and updating the time stamp

For each metric group, the Operations Analytics Predictive Insights Mediation Tool selects the metric that best matches a time stamp. You must ensure that the correct metric was selected and its properties set correctly.

You can replace the time stamp that was selected for a metric group in the **Timestamp** field within the **Group Definition** section of the **Model Design** tab. To replace the time stamp, right-click the existing metric that is in the **Timestamp** field and select **Remove Timestamp Key**. Then, drag a new metric into the field.

You can verify that the time stamp is configured correctly and, if necessary, update the properties of the time stamp in the **Model Properties** tab. To verify that a metric group time stamp is configured correctly, complete the following steps:

1. Check that the correct metric is in the **Timestamp** field within the **Group Definition** section of the **Model Design** tab.
2. Open the **Model Properties** tab.
3. In the **Model Objects** pane, select the drop-down for the relevant metric group. The **Timestamp** is listed as the first item under the metric group.
4. Click the **Timestamp** and ensure that the **Data Type** and **Time Format** are set correctly. Table 1 provides instructions on the data type to select and the time format to specify for different time stamp formats in the data source.
5. Run a preview of the extraction for the metric group to ensure it does not fail with any time format errors.

Table 1. Data types for different time stamp formats in data source

Time stamp format in data source	Data Type	Time Format
Epoch format in seconds. For example, 1480349643	Integer	Leave blank
Epoch format in milliseconds. For example, 1480349643000	Number	Leave blank
Any other time stamp format	Timestamp	Leave blank
Any time stamp format that cannot be automatically detected by Predictive Insights when the Data Type field is set to Timestamp	String	Enter a string that matches the time format of the time stamp in the data source. For example, if a data source has a time stamp format of 2016-11-24 11:15:00, in the Time Format field, enter the following string format: yyyy-MM-dd HH:mm:ss

Creating filters

Filters contain criteria that Operations Analytics Predictive Insights uses to control the data it extracts from a data source or the data it analyzes after extraction, or both.

About this task

You can configure filters on attributes and resource keys.

Operations Analytics Predictive Insights applies attribute filters before it extracts data from the data source to reduce the amount of data that is extracted.

Operations Analytics Predictive Insights applies resource key filters to the extracted data before it is written to the database and analyzed. Resource key filtering is a useful way of filtering the data that Operations Analytics Predictive Insights analyzes if you are unable to accurately filter the data before it is extracted from the data source. You can use any java regular expression for a resource key filter. If a resource key does not match the filter expression, it is skipped.

Procedure

1. In the **Operations Analytics Predictive Insights Navigator** tab, select the data source for which you want to add a filter.
2. Select the **Model Properties** tab.

3. Under **Model Objects**, expand **Model** and expand a metric group name. Under the metric group name, you see the following model objects: time stamp, resource keys, attributes, and metrics. You can create filters for the resource key and for attributes.
4. To create a new filter, complete the following steps:
 - a. Click **ResourceKey** or expand **attributes** and click an attribute.
 - b. Under **Filter Properties**, enter the filter string in the **Filter expression** field. A filter expression can contain any character but double quotation marks, " ", are ignored.
 - c. To view the new filter, click **Model Objects->Filters**.
 - d. If you want an attribute filter to require an exact match of the filter expression string, under **Model Objects->Filters**, select the filter. Under **Filter Properties**, set the **SQL Direct Match** field to true. Setting this option to true improves the speed of many queries. By default, the **SQL Direct Match** field is set to false and the filter matches anything that contains the filter expression. For more information, see the example.

Note: Setting the **SQL Direct Match** field to true affects attribute filters only. If you set this field to true for a resource key filter, the setting is ignored.

5. To copy an existing filter and apply it to a resource key or attribute, complete the following steps:
 - a. Under **Model Objects->Filters**, click the filter that you want to copy.
 - b. Under **Model Object Properties**, select the value in the **Filter expression** field.
 - c. Right click and choose **Copy**.
 - d. Under **Model Objects**, navigate to the object to which you want to apply a filter.
 - e. Under **Model Object Properties**, select the **Filter expression** field.
 - f. Right click and click **Paste** to paste the expression into the field.

Note: If you select a filter expression under **Model Objects->Filters** and update the filter expression, the filter is updated in any resource key or attribute to which the filter applies.

Example

Table 2. Example filters

Filter	Description
^(?!((ABC_ DEF_).*)	Excludes resources that start with ABC_ or DEF_
^((ABC_ DEF_).*)	Includes resources that start with ABC_ or DEF_ only
East10,West10	<p>If Sql Direct Match is set to false in the filter properties, the filter is a contains filter and matches any string that contains East10 or West10. For example, the filter matches East10, East101, and West1024.</p> <p>If Sql Direct Match is set to true in the filter properties, the filter is an exact match filter that matches East10 or West10 only. For example, it does not match East101 or West102.</p>

Unpegging

The unpegging of metrics results in each received metric value being saved so that it can be compared to the next value. By comparing the current value with the previous value the extractor is able to establish the delta for that metric.

The delta is the difference between the value currently being extracted and the previous value. The delta is often more useful to the algorithm than the actual value.

Unpegging of metrics is implemented using the Mediation Tool.

To unpeg a metric:

1. Select the data source within the **Predictive Insights Navigator** tab for which you want to unpeg metrics.
The **Connection Details** tab is displayed.
2. Select the **Model Properties** tab.
The tab displays a tree view of the available **Model Objects**, which are made up of the time stamp, resource keys, attributes, and metrics of your logical model.
3. Select the metric you wish to unpeg.
4. In the **Model Object Properties** pane, click the **Metric Type** field and select **Accumulation** from the drop-down menu.

Validating a model

It is important to validate your model, as the Operations Analytics Predictive Insights Mediation tool can not deploy a model that has errors or is inconsistent with the database.

Procedure

1. Right click the model you want to validate within the **Predictive Insights Navigator** tab, and click **Open**.
2. If your project is already open, ensure that it is saved, by clicking **Save**.
3. Open the **Problems** tab. If the **Problems** tab is not visible, click **Window > Show View > Problems**.
The **Problems** tab displays all issues that it can see with your tool based on the model you saved.
The Operations Analytics Predictive Insights Mediation tool looks for errors in the model. Errors can be inconsistencies between the new model and the source model, created when you synchronized with the data source.
4. Resolve the problems that are listed.

Model configuration examples for different time zone scenarios in source files

The model configuration scenarios section illustrates how to configure an Operations Analytics Predictive Insights model for specific data source scenarios.

CSV file names and data are for different time zones

The following examples show how to configure Operations Analytics Predictive Insights when the time stamps in the file names and data represent different time zones.

Example 1: Time zone shown in CSV file names but not in data

In this example:

- the time stamp for the CSV source files is in EST as shown in the file name
- the time stamps for the data within the file are known to be in GMT but the time zone is not shown.

File name:

CPULOAD_2013-07-17-00-00EST_2013-07-17-00-15EST.csv

Data:

EndTime,Resource,In_TotalBytes,Out_TotalBytes
2013-07-17-05-00,"ethernet-csv785.tut.com_GigabitEthernet2", 3263768830,9545496800

Configuration in Operations Analytics Predictive Insights:

To configure Operations Analytics Predictive Insights for the time stamps used in the example:

When you add the file system data source in the Mediation tool, set the following fields:

- In the **Name Pattern** field, enter : ([a-zA-Z]\w*)(\d{4}-\d{2}-\d{2}-\d{2}\w{3})_([\d{4}-\d{2}-\d{2}-\d{2}-\d{2}\w{3}]).*\w{3}).*.csv
- In the **Time Format** field, enter: yyyy-MM-dd-HH-mmz. The z at the end of the field indicates that the last 3 characters of the file name are a time zone identifier
- In the **Time Zone** field, set the time zone to Europe/London - Greenwich Mean Time. When the time stamps in the file name and data represent different time zones and the time zone is specified in only one of the time stamps, you set this field to the time zone that is not specified in the time stamp. In this example, the time stamps in the data are in GMT but the time zone is not specified in the time stamp, so you set the timezone to GMT.

Example 2: Time zone shown in data but not in CSV file names

In this example,

- the time stamp in the file name is known to represent EST but the time zone information is not shown in the file name.
- the time stamp in the data is in GMT and time zone is included in the time stamp

File name:

CPULOAD_2015-07-17-00-00_2015-07-17-00-15.csv

Data:

```
EndTime,Resource,In_TotalBytes,Out_TotalBytes
2013-07-17-05-00+0000,"ethernet-csv785.tut.com_GigabitEthernet2",3263768830,9545496800
```

Configuration in Operations Analytics Predictive Insights:

To configure Operations Analytics Predictive Insights for this example, do the following:

In the Mediation tool's **File System Details** tab, do the following:

- In the **Name Pattern** field, enter: `([a-zA-Z]\w*)_(\d{4}-\d{2}-\d{2}-\d{2}-\d{2})_(\d{4}-\d{2}-\d{2}-\d{2}-\d{2})-.*\.csv`
- In the **Time Format** field, enter: `yyyy-MM-dd-HH-mm`.
- In the **Time Zone** field, set the time zone to `America/New_York - Eastern Standard Time`. When the time stamps in the file name and data represent different time zones and the time zone is specified in only one of the time stamps, you set this field to the time zone that is not specified in the time stamp. In this example, the time zone in the file name is not specified but is known to represent EST, so you set the timezone to `America/New_York - Eastern Standard Time`.

In the Mediation tool's **Model Properties** tab, do the following:

- In the **Model Objects** pane, expand the metric group and click **Timestamp**.
- In the Model Object Properties pane, enter the following value in the **Time Format** field: `yyyy-MM-dd-HH-mmZ`. The Z at the end of the field indicates that the time stamps in the data end with the time zone specified in the RFC 822 4-digit format.

Example 3: Time zone shown in both the CSV file name and data

In this example,

- the time stamp in the file name shows the EST time zone.
- the time stamp in the data shows the GMT time zone

File name:

CPULOAD_2013-07-17-00-00EST_2013-07-17-00-15EST.csv

Data:

```
EndTime,Resource,In_TotalBytes,Out_TotalBytes
2013-07-17-05-00GMT,"ethernet-csv785.tut.com_GigabitEthernet2",3263768830,9545496800
```

Configuration in Operations Analytics Predictive Insights:

To configure Operations Analytics Predictive Insights for this example, do the following:

In the Mediation tool's **File System Details** tab, do the following:

- In the **Name Pattern** field, enter: `([a-zA-Z]\w*)_(\d{4}-\d{2}-\d{2}-\d{2}-\d{2})_w{3}_(\d{4}-\d{2}-\d{2}-\d{2}-\d{2})w{3}-.*\.csv`
- In the **Time Format** field enter: `yyyy-MM-dd-HH-mmz`. The z at the end of the field indicates that the last 3 characters of the file name are a time zone identifier

- In the **Time Zone** field, set the time zone to America/New_York - Eastern Standard Time. When different time zones are explicitly shown in the file name and data, set the time zone to that shown in the file name.

In the Mediation tool's **Model Properties** tab, do the following:

- In the **Model Objects** pane, expand the metric group and click **Timestamp**.
- In the Model Object Properties pane, enter the following value in the **Time Format** field: yyyy-MM-dd-HH-mmz. The z at the end of the field indicates that the time stamps in the data end with the time zone specified in the General time zone format.

CSV file names and data have no time zone identifier

The following examples show how to configure Operations Analytics Predictive Insights when there is no time zone identifier in the time stamps in the CSV file names or the data in the files.

Example: CSV file names and data have no time zone identifier

If there is no time zone identifier in the time stamps in the CSV file name or the data within the files, the time stamps in both the file names and the data must be for the same time zone.

File name:

CPULOAD_2013-07-17-00-00.csv

Data:

EndTime,Resource,In_TotalBytes,Out_TotalBytes
2013-07-17-05-00,"ethernet-csv785.tut.com_GigabitEthernet2", 3263768830,9545496800

Configuration in Operations Analytics Predictive Insights:

To configure Operations Analytics Predictive Insights for the time stamps used in the example:

When you add the file system data source in the Mediation tool, set the following fields:

- In the **Name Pattern** field, enter : ([a-zA-Z]\w*)_(\d{4}-\d{2}-\d{2}-\d{2}-\d{2}).*.csv
- In the **Time Format** field, enter: yyyy-MM-dd-HH-mm.
- In the **Time Zone** field, select the appropriate time zone that represents both the file names and the data.

In the Mediation tool's **Model Properties** tab, do the following:

- In the **Model Objects** pane, expand the metric group and click **Timestamp**.
- In the Model Object Properties pane, enter the following value in the **Time Format** field: yyyy-MM-dd-HH-mm.

CSV files are for multiple time zones

This topic describes how to configure Operations Analytics Predictive Insights when a CSV data source has time stamps for multiple time zones in the CSV file names or the data within the files.

If there are multiple time zones in CSV file names, or the data within CSV files, you must do the following:

- Group the CSV files by time zone in separate directories on the Analytics server
- Add a separate file system data source for each directory in the Mediation tool
- Create a separate model for each data source in the Mediation tool

Deploying a model

You must deploy the model to the database so it is usable by Operations Analytics Predictive Insights.

Before you begin

A model is defined by the set of defined data sources that you select to deploy. This can be all of the data sources under a project, or a subset of those data sources. Each data source has its own .pamodel file in the workspace and these will be automatically combined when deployed together. You can edit a model after you deploy it. When you edit a model, you can only add, remove, or rename metric groups. When you redeploy the updated model, the previous version of the model is overwritten.

You must deploy each model that you create to a topic. Topics allow you to segment your data in ways that are relevant to your business. For example, you can use a topic to group data that originates from a specific service, application, or geography. The anomalies generated based on your data can then be grouped and displayed by topic.

For information about creating a new topic see *Creating a topic* in the *Configuring and Administering* Guide.

When you deploy a model to a topic, the previous version of the topic is overwritten. Therefore, if you wish to deploy a model to a topic that contains other models, you must select all models for the specific topic, in the Mediation Tool, and deploy them together. If you are redeploying one or more models to a topic that is already running, you must first stop the topic using the stop command. In addition, if the topic has already consumed data, you need to run data extraction and start model training again after you redeploy the model(s).

Note: If there is only a single topic in your system, the model is deployed to that topic automatically.

Procedure

1. Within the **Predictive Insights Navigator** pane, select the model or set of models that you want to deploy. You can select multiple models only if all of those models exists as part of the one project.
2. Click **Predictive Insights > Deploy Model**.
3. The **Predictive Insights Model Deployment** dialog opens, in which you enter the Operations Analytics Predictive Insights database connection details and password.

Note: The database credentials are automatically populated with the user credentials of scadmin. Make sure that you update the user credentials with the details of your Operations Analytics Predictive Insights database user.

4. Select the **Calculate KPI Count** check box to generate an estimate of the number of KPIs for the model you are deploying. The KPI count is shown as you deploy the model.

Note: The Mediation Tool calculates the KPI Count from a sample of the source data. By default, the sample is taken from data created between 00:00 and 01:00 and 12:00 and 13:00. When you are using a file system data source, ensure that this sample is representative of the complete data set by having at least one full day of data available to the Mediation Tool. If the Mediation Tool is running on a separate server to the Analytics component, you must copy the sample data to the Mediation Tool server.

5. Click **OK**.
6. You are prompted to enter the data source connection password for each database data source. Enter the password for each database data source and click **OK**.

Note: File system data sources do not request a password.

The Operations Analytics Predictive Insights Mediation Tool now validates your model against the actual data source.

Note: An error occurs if you choose to deploy your model without first ensuring you have access through the firewall to the server that contains the database.

7. If you are alerted that there are errors in your model:
 - a. Fix any validation errors that are displayed in the problems tab.
 - b. After fixing errors, deploy your model or models again by selecting them within the **Predictive Insights Navigator** pane and clicking **Predictive Insights > Deploy Model**.
8. If you have more than one topic available, the **Topic Selection** dialog opens. Select the topic with which you want to associate your new model. If you have only one topic available, this topic is selected by default.

Note: You can only deploy separate data sources to the same topic if those sources overlap in time. For example, if data source 1 is a CSV datasource with data from July to September and Datasource 2 is a CSV datasource with data from October to December then you must not deploy these data sources to the same topic.

9. Click **OK**.
10. The Mediation Tool calculates and displays the KPI count. If the KPI count matches the number of KPIs you planned to deploy in the model, choose **Yes** to complete the deployment. Otherwise, choose **No** and modify your model to add or remove metrics, as required.

Creating a topic

You must deploy each model that you create to a topic. A default topic is created as part of the installation of the Analytics component. You can create additional topics to segregate data analysis and presentation.

About this task

Topics allow you to segment your data in ways that are relevant to your business. For example, you can create a topic and use it to group data that originates from a specific service, application, or geography. The anomalies generated based on your data can then be grouped and displayed by topic. Data in different topics are analyzed independently of each other and no analysis takes place across topics. When you deploy a new model, you are asked to choose the topic to which the new model belongs.

Note: For file based data sources, do not set up multiple topics that point to the same source location.

Procedure

1. Log on to the Operations Analytics Predictive Insights Analytics server as the administrative user, typically `scadmin`, and navigate to `$PI_HOME/bin`.

Note: The `$PI_HOME` environment variable is set automatically when the user logs in. However, if your user's shell has been running since before the installation, you might need to exit and log in again.

2. Run the **`create_topic`** command for each topic you wish to create.

The **`create_topic`** command is one of the available commands within the **`admin`** CLI application:

```
./admin.sh create_topic <topic name> <description>
```

Where:

- **<topic name>**: Is the name of the new topic. The topic name must be one word between 3 and 10 characters long. It can contain alphanumeric characters and the underscore character
- **<description>**: The topic description. The description should be enclosed in double quotes if it contains spaces.

For example:

```
./admin.sh create_topic network "Topic for network data"
```

What to do next

If the topic is no longer being used it can be deleted using the **`delete_topic`** command.

Related reference:

“`create_topic`” on page 102

The `create_topic` CLI command.

“`delete_topic`” on page 102

The `delete_topic` CLI command.

Exporting a model

You can reuse a model within another project.

Procedure

1. To export all models within a project, right click the project and click **Export**. To export only a single model, right click that model and click **Export**.
2. Click **General > File System**.
3. Click **Next**
4. Click **Browse** next to the **To directory** field and change to the directory to which you would like to export the model.
5. Select the check box corresponding to the model or models that you would like to export.
6. Click **Finish**.

Importing a model

You can import an existing model file into your project.

About this task

Instructions on how to import a model file into your project within the Operations Analytics Predictive Insights Mediation tool.

Procedure

1. Select your project and then click **File > Import**. The **Import** dialog opens.
2. Click **Predictive Insights > Predictive Insights model**.
3. Click **Next**.
4. Click **Browse** next to the **Data Source File** field and change to the directory that contains the model you would like to import.
5. Select the check box corresponding to the model or models that you would like to import into your project.
6. Click **Finish**.

Chapter 3. Configuring security

After you install Operations Analytics Predictive Insights, you can configure various security options.

Configuring security for Dashboard Application Services Hub

You need to perform steps to configure additional security when Operations Analytics Predictive Insights is deployed with Dashboard Application Services Hub.

Securing sensitive cookies

You can restrict the exchange of cookies to HTTPS sessions only.

Procedure

1. Enter the URL of the WebSphere administrative console: `https://<hostname>:16316/ibm/console`.
where <hostname> is the name or IP address of the Dashboard Application Services Hub server.
2. On the login screen, enter your username and password as configured during installation. The default username set at installation is `smadmin`.
3. On the navigation panel, click **Applications->Application Types->WebSphere enterprise applications**.
4. In the Enterprise Applications table, click **isc**.
5. Under Web Module properties, click **Session Management**.
6. Under General Properties, ensure that **Enable cookies** is checked.
7. Click the **Enable cookies** link.
8. Click the **Restrict cookies to HTTPS sessions** check box. Do not click the **Set session cookies to HTTPOnly ...** check box as this prevents access to the Active Event List.
9. Click **OK** to close the Cookies page and **OK** again to close the Session management page.
10. To save the changes, under Messages, click **Save**.
11. Restart Dashboard Application Services Hub.

Enable Transport Layer Security (TLS)

Read this page to view the steps required to enable TLS.

Procedure

1. Enter the URL of the WebSphere administrative console: `https://<hostname>:16316/ibm/console`.
where <hostname> is the name or IP address of the Dashboard Application Services Hub server.
2. On the navigation panel, click **Security->SSL certificate and key management**.
3. Under Related items, click **ssl configurations**.
4. On the login screen, enter your username and password as configured during installation. The default username set at installation is `smadmin`.

5. In the table, click **NodeDefaultSSLSettings**.
6. Under Additional Properties click **Quality of protection (QOP) settings**.
7. In the Protocol drop down, select **TLS**.
8. To save the changes, under Messages, click **Save**.
9. Restart Dashboard Application Services Hub.

Disable auto-complete on the Dashboard Application Services Hub login panel

How to disable auto-complete on the Dashboard Application Services Hub login panel.

About this task

Procedure

1. As the user that installed Dashboard Application Services Hub, log in to the server.
2. Change to the `<installdir>/config/cells/ /JazzSMNode01Cell/applications/isc.ear/deployments/isc/isclite.war/WEB-INF` directory.
Where `<installdir>` is the directory where Dashboard Application Services Hub is installed.
3. Edit the `customizationProperties.xml` file.
4. Set the value of the `LOGIN.CACHEPASSWORD` property to `false`.
5. Save the file.
6. Restart Dashboard Application Services Hub.

Configuring security for Tivoli Integrated Portal

You need to perform steps to configure additional security when Operations Analytics Predictive Insights is deployed with Tivoli Integrated Portal.

Securing sensitive cookies

You can restrict the exchange of cookies to HTTPS sessions only.

Procedure

1. Enter the URL of the WebSphere administrative console: `https://<hostname>:16316/ibm/console`.
where `<hostname>` is the name or IP address of the Dashboard Application Services Hub server.
2. On the login screen, enter your username and password as configured during installation. The default username set at installation is `smadmin`.
3. On the navigation panel, click **Applications->Application Types->WebSphere enterprise applications**.
4. In the Enterprise Applications table, click **isc**.
5. Under Web Module properties, click **Session Management**.
6. Under General Properties, ensure that **Enable cookies** is checked.
7. Click the **Enable cookies** link.
8. Click the **Restrict cookies to HTTPS sessions** check box. Do not click the **Set session cookies to HTTPOnly ...** check box as this prevents access to the Active Event List.

9. Click **OK** to close the Cookies page and **OK** again to close the Session management page.
10. To save the changes, under Messages, click **Save**.
11. Restart Dashboard Application Services Hub.

Enabling Transport Layer Security

Read this page to view the steps required to enable TLS.

Procedure

1. Enter the URL of the WebSphere administrative console: `https://<hostname>:16316/ibm/console`.
where `<hostname>` is the name or IP address of the Tivoli Integrated Portal server.
2. On the navigation panel, click **Security->SSL certificate and key management**.
3. Under Related items, click **ssl configurations**.
4. On the login screen, enter your username and password as configured during installation. The default username set at installation is `smadmin`.
5. In the table, click **NodeDefaultSSLSettings**.
6. Under Additional Properties click **Quality of protection (QOP) settings**.
7. In the Protocol drop down, select **TLS**.
8. To save the changes, under Messages, click **Save**.
9. Restart Tivoli Integrated Portal.

Configuring certificates

Self-signed certificates are automatically generated during the installation of Operations Analytics Predictive Insights.

In a production environment, it is recommended that you replace the self-signed certificates with certificates that are signed by a third party Certificate Authority. Follow these steps to install certificates on the Operations Analytics Predictive Insights servers.

Adding a signed certificate to the User Interface server

Read this page to view the steps required to add a signed certificate to the Operations Analytics Predictive Insights User Interface Server.

Procedure

1. Go to the `<installdir>/UI/wlp/usr/servers/piserver/resources/security` directory:
`cd <installdir>/UI/wlp/usr/servers/piserver/resources/security`

Where `<installdir>` is the directory where Operations Analytics Predictive Insights is installed. The default installation directory is `/opt/IBM/scanalytics`.

2. Enter the following command to make a backup copy of the `key.jks` keystore file:

```
cp key.jks key.jks.orig
```

3. Enter the following command to generate a certificate request based on the default self-signed certificate:

```
<installdir>/UI/ibm-java-x86_64-70/bin/keytool -certreq -alias default -file certreq.pem -keys
```

4. Obtain a signed certificate from a third party certificate authority. Normally, the certificate authority provides the signed certificate along with a root and intermediate certificate.
5. If the certificate authority provided root and intermediate certificates, enter the following commands in the order shown to import the certificates:


```
<installdir>/UI/ibm-java-x86_64-70/bin/keytool -importcert
      -alias inter -file <intermediate file> -keystore key.jks -storepass emrjtn56rtjg
      <installdir>/UI/ibm-java-x86_64-70/bin/keytool -importcert
      -alias root -file <root file> -keystore key.jks -storepass emrjtn56rtjg
      <installdir>/UI/ibm-java-x86_64-70/bin/keytool -importcert
      -alias default -file <signed cert file> -keystore key.jks -storepass emrjtn56rtjg
```
6. If the certificate authority did not provide root and intermediate certificates, enter the following command to import the chain along with the certificate:


```
<installdir>/UI/ibm-java-x86_64-70/bin/keytool -importcert -alias default -file <signed cert file>
      -keystore key.jks -storepass emrjtn56rtjg
```

Adding a signed certificate to the Dashboard Application Services Hub server

Read this page to view the steps required to add a signed certificate to the Dashboard Application Services Hub server.

To add a signed certificate to the Dashboard Application Services Hub server, see the information at the following links:

- <http://www-01.ibm.com/support/docview.wss?uid=swg21654278>
- http://www.ibm.com/developerworks/websphere/techjournal/1210_lansche/1210_lansche.html

Configuring LDAP authentication

Read this page to view how to configure Operations Analytics Predictive Insights to use LDAP authentication.

Before you configure Operations Analytics Predictive Insights to use LDAP authentication, do the following:

- Verify that the User Interface functions correctly in Dashboard Application Services Hub
- Use groups to map roles so that you only need to move users in to or out of groups to grant or revoke privileges
- If you use HTTPs authentication for your LDAP directory, use the `keytool -importcert` command to import the LDAP server certificate into the liberty trust store. For more information on how to use the command, see “Adding a signed certificate to the User Interface server” on page 37. When importing the certificate, use an alias other than default to identify the LDAP server certificate in the trust store.

Configuring the Operations Analytics Predictive Insights UI server to use LDAP authentication

Read this section to view the steps to configure LDAP authentication for the Liberty container that is installed with the Operations Analytics Predictive Insights UI server.

Configuring the LDAP User Registry

How to configure the LDAP user registry.

About this task

In this procedure, <installdir> is the directory where Operations Analytics Predictive Insights is installed. The default installation directory is /opt/IBM/scanalytics.

Procedure

1. Log in to the server where the Operations Analytics Predictive Insights UI component is installed.
2. Go to the <installdir>/UI/wlp/usr/servers/piserver directory.
3. Create a new file called ldapRegistry.xml and add the appropriate content to the file for the LDAP server that you are using:

If you are using a Tivoli Directory server, add the following:

```
<server>
<ldapRegistry id="ldap"
realm="defaultWIMFileBasedRealm" baseDN="ou=people,ou=internal,O=IBM,C=US"
host="ldap1.ibm.com" port="389" ignoreCase="true"
bindDN="cn=NetcoolReadOnly,cn=ReadOnlyUsers,O=IBM,C=US" bindPassword="***"
ldapType="IBM Tivoli Directory Server" sslEnabled="false" ><idsFilters
userFilter="(&(uid=%v)(objectclass=person))"
groupFilter="(&(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUniqueNames)
objectclass=groupOfURLs))" userIdMap="*:uid"
groupIdMap="*:cn"
groupMemberIdMap="mycompany-allGroups:member;mycompany-allGroups:uniqueMember;
groupOfNames:member;groupOfUniqueNames:uniqueMember"></idsFilters></ldapRegistry>
</server>
```

If you are using a Microsoft Active Directory server, add the following:

```
<server>
<ldapRegistry id="ldap"
realm="defaultWIMFileBasedRealm" host="DC1.ibm.com" port="389"
baseDN="CN=Users,DC=webguil3,DC=ldap,DC=com"
bindDN="cn=manitou,CN=Users,DC=webguil3,DC=ldap,DC=com" bindPassword="*****"
ldapType="Microsoft Active Directory" sslEnabled="false" ><activeFilters
userFilter="(&(sAMAccountName=%v)(objectcategory=user))" groupFilter="(&(cn=%v)(objectc
userIdMap="user:sAMAccountName" groupIdMap="*:cn" groupMemberIdMap="memberOf:member"></activeF
</ldapRegistry></server>
```

4. Edit the values in the text you add to the ldapRegistry.xml file to match your LDAP settings. In particular:
 - The realm specified in the ldapRegistry.xml file must match the realm name specified in the WebSphere Administrative Console. If they do not match, single-sign-on will not work. The default realm is: defaultWIMFileBasedRealm
 - Ensure that the values in the userIdMap and groupIdMap fields are correct. Otherwise, log in will fail.
 - Ensure that the ldapRegistry tag is on a single continuous line. Otherwise, the liberty server may fail parsing it.

Also, ensure that the values you specify match the settings for Dashboard Application Services Hub. To check the settings in Dashboard Application Services Hub

- a. Go the following URL to launch the Webshere Administrative console:
https://<hostname>:16316/ibm/console
- b. Log in as the scadmin user.
- c. Click **Security > Global Security**.
- d. Click **Configure**.
- e. Locate the LDAP repository in the table and click **Base Entry** link to verify the base DN.
- f. Click the **Repository Identifier** and verify the bindDN, host and port.
5. Make a backup copy of the <installdir>/UI/wlp/usr/servers/piserver/server.xml file.
6. Edit the <installdir>/UI/wlp/usr/servers/piserver/server.xml file and make the following changes:
 - a. Comment out or remove the following lines:


```
<feature>usr:omnibusUserRegistry-1.0.0</feature>
<include optional="true" location="{server.config.dir}/omnibusConfig.xml"/>
```
 - b. Add the following feature as a sub element of the <featureManager> tag:


```
<feature>ldapRegistry-3.0</feature>
```

After you add the feature, this section of the file should look similar to the following:

```
<!-- Enable features -->
<featureManager>
  <feature>jsp-2.2</feature>
  <feature>jaxrs-1.1</feature>
  <feature>servlet-3.0</feature>
  <feature>appSecurity-2.0</feature>
  <feature>jndi-1.0</feature>
  <feature>jdbc-4.0</feature>
  <!--<feature>usr:omnibusUserRegistry-1.0.0</feature> -->
  <feature>ldapRegistry-3.0</feature>
</featureManager>
```
 - c. Locate and Remove the comment from the following line:


```
<!-- <include optional="true" location="{server.config.dir}/ldapRegistry.xml"/> -->
```

After you remove the comment, the line should look as follows:

```
<include optional="true" location="{server.config.dir}/ldapRegistry.xml"/>
```
 - d. Save the file.
7. If the Operations Analytics Predictive Insights UI was installed into a Jazz for Service Management instance that has a non-default cookie name configured, update the cookie name in the /opt/IBM/scanalytics/UI/wlp/usr/servers/piserver/ssoConfig.xml file:
 - a. Locate the line:


```
<webAppSecurity ssoDomainNames=".<domain name>" />
```
 - b. Update the line to read:


```
<webAppSecurity ssoDomainNames=".<domain name>" ssoCookieName="<cookie name>" />
```

where <cookie name> is the customized name of the Jazz for Service Management cookie.
8. Enter the following command to restart the Operations Analytics Predictive Insights UI server:


```
<installdir>/UI/bin/pi.sh -restart
```

Granting LDAP users and groups access to Operations Analytics Predictive Insights

To grant users in the LDAP directory access to Operations Analytics Predictive Insights, it is recommended that you assign Operations Analytics Predictive Insights roles to LDAP groups and add LDAP users to the groups.

About this task

Note:

If you create groups called `predictiveInsightsUsers` and `predictiveInsightsAdmins` in your LDAP directory, roles are automatically assigned to these groups and you do not need to complete the following procedure.

Procedure

1. Log in to the Operations Analytics Predictive Insights server.
2. Go to the `<installdir>/UI/bin` directory.
Where `<installdir>` is the directory where Operations Analytics Predictive Insights is installed. The default installation directory is `/opt/IBM/scanalytics`.
3. Run the following command to grant an LDAP user or group access to Operations Analytics Predictive Insights:
`addAccess.sh <user | group> <user or group name> [admin]`

For example, to grant administrator access to a group called `PI_admins`, enter the following:

```
addAccess.sh group PI_admins admin
```

Verifying the LDAP configuration

How to verify that the LDAP authentication and authorization is configured correctly for the Operations Analytics Predictive Insights User Interface server.

About this task

Procedure

1. Go to the following URL: `https://<hostname>:9998/predictiveinsights/jsp/wlp/wlpAnomalySearch.jsp`
where `<hostname>` is the name or IP address of the server on which the Operations Analytics Predictive Insights User Interface component is installed.
2. On the login screen, enter the username and password of an LDAP user. If you log in successfully, the Anomaly Search screen is displayed. If the log in fails:
 - Confirm that the information you added to the `ldapRegistry.xml` file is correct.
 - If you receive an authorization error, ensure that the procedure to grant LDAP users access to Operations Analytics Predictive Insights was completed correctly.
3. After you log in successfully, test that the single sign-on is correct:
 - a. Close all open instances of your browser.
 - b. Launch Dashboard Application Services Hub.
 - c. On the log in screen, enter the username and password of an LDAP user.

- d. Verify that the snowflake icon ❄️ is visible on the Dashboard Application Services Hub menu. If it is not visible, then there is either a single sign-on issue or an issue with the console integration
- e. If the snowflake icon is not visible, open a new browser tab and go to the following URL: `https://<hostname>:9998/predictiveinsights/jsp/wlp/wlpAnomalySearch.jsp`
- f. If you are prompted for a user name or password or receive an authentication error, verify that the realm entered in `ldapRegistry.xml` matches the realm name in the WebSphere Administrative console, under **Federated Repositories**.

Note:

Configuring OMNIbus Web GUI authentication against an LDAP directory

You can configure the Web GUI to authenticate users and groups against an LDAP directory.

Before you begin

- Familiarize yourself with the concept of the VMM realm. See *Web GUI user authentication* in the Tivoli OMNIbus Knowledge Center.
- Ensure that the LDAP directory is running and that it can be accessed from the Web GUI host computer.
- If the previous user repository was the default file-based repository, remove any default users that were created when the file-based repository was added. You need to remove these users to avoid duplicate users across repositories in the realm.
- Obtain the following information about the LDAP directory. You need this information to configure the LDAP directory in the realm.
 - Host name and port number of the primary server that hosts the LDAP directory and the backup server, if applicable. The host names must contain no spaces.
 - Type and version of LDAP directory that is used, for example IBM Tivoli Directory Server V6.2, or Microsoft Active Directory.
 - User ID and password that are used to bind to the LDAP server. This user ID must be unique. For example, `cn=root`. Important: To create users and groups through the Web GUI, the LDAP bind ID must have the appropriate permissions in the LDAP directory. The bind ID must contain no spaces.
 - Subtree of the LDAP directory that you want to be used for authenticating users.

Sample LDAP data

The following configuration tasks use sample data from a subtree in an LDAP directory. When you complete the configuration tasks, replace the sample data with your own.

The LDAP directory is identified as TIVIDS. TIVIDS contains the subtree `ou=NetworkManagement,dc=myco=dc=com`, which contains the users and groups that are authenticated by the Web GUI. In this subtree, the LDAP objects are defined as follows:

- The user prefix is `uid`

- The user suffix is cn=users
- The group prefix is cn
- The group suffix is cn=groups

In the subtree, the administrator user with the user name Administr8or is defined as uid=Administr8or,cn=users,ou=NetworkManagement,dc=myco,dc=com. The administrator user group with the name AdminGroup is defined as cn=AdminGroup,cn=groups,ou=NetworkManagement,dc=myco,dc=com.

Adding the LDAP directory to the realm

To authenticate the users from an LDAP directory, the Web GUI needs to read the LDAP user data. To achieve this, add the LDAP directory to the Virtual Member Manager (VMM) realm as a repository.

Before you begin

You need the following information to configure the LDAP directory in the realm:

- The host name and port number of the primary server that hosts the LDAP directory and the backup server, if applicable. The host names must contain no spaces.
- Type and version of LDAP directory that is used, for example IBM Tivoli Directory Server V6.2, or Microsoft Active Directory.
- The user ID and password that are used to bind to the LDAP server. This user ID must be unique, for example, cn=root. Important: To create users and groups through the Web GUI, the LDAP bind ID must have the appropriate permissions in the LDAP directory. The bind ID must contain no spaces.
- The subtree of the LDAP directory that you want to be used for authenticating users.

About this task

The configuration steps in this task use the sample LDAP directory that is described in “Sample LDAP data” on page 42. Replace the values from this sample with your own.

Procedure

1. Make a backup copy of the JazzSM_HOME/config/cells//wim/config/wimconfig.xml file.
2. Log in as the administrative user, typically scadmin, and launch the WebSphere®:
 - a. From the menu on the left, click **Settings > WebSphere Administrative Console**.
 - b. Click **Launch WebSphere Administrative Console**.
3. Click **Security > Global Security**
4. Under User account repository, select Federated Repositories from the **Available realm definitions** list and then click **Configure**.
5. Click **Add repositories (LDAP, custom, etc)...** and then click **New Repository > LDAP repository**.
6. Add the LDAP directory as a repository to the realm by completing the following fields:

Repository identifier

Type TIVIDS. The repository identifier uniquely identifies the repository within the realm.

Directory type

Select the type of LDAP server. The type of LDAP server determines the default filters that are used by WebSphere Application Server. IBM Tivoli Directory Server users can select either **IBM Tivoli Directory Server** or **Secure Way**, but **IBM Tivoli Directory Server** offers performance. For OpenLDAP directories, select **Custom**.

Primary host name

Type the fully qualified host name of the primary LDAP server. You can enter either the IP address or the domain name system (DNS) name.

Port

Type the port of the LDAP directory. The host name and the port number represent the realm for this LDAP server in a mixed version nodes cell. The default port value is 389, which is not a Secure Sockets Layer (SSL) connection port. Use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port.

Bind distinguished name

Type the bind ID of the LDAP server. If anonymous binds are not possible on the LDAP server, the bind DN is for write-operations or to obtain user and group information. Always provide a bind DN and bind password, unless an anonymous bind can satisfy all of the functions. If the LDAP server is set up to use anonymous binds, you can leave these fields blank.

Bind password

Type the password of the bind ID.

After you completed the fields, click **Apply**, then click **Save**. The LDAP directory is added to the repositories in the realm.

7. Define the repository by completing the following fields:

Distinguished name of a base entry that uniquely identifies this set of entries in the realm

Example: o=TIVIDS. This distinguished name (DN) defines an entry for the LDAP directory in the realm.

Distinguished name of a base entry in this repository

Example: o=NetworkManagement,dc=myco,dc=com. This DN is the root of the subtree in the LDAP directory that you want to use for authentication.

Important: If you leave this field blank, the base entry is mapped to the root of the LDAP directory. All operations are performed at root, which causes errors on most LDAP servers.

After you completed the fields, click **Apply**, then click **Save**.

8. Restart the server.

Results

The users from the subtree of the LDAP directory are replicated in the realm. After you restart the server, the users from the LDAP directory are visible on the

Manage Users page of the administrative console. The DN of the users consists of the user prefix and suffix and the DN of the LDAP directory in the realm, in this case o=TIVIDS, which represents the ou=NetworkManagementdc=myco,dc=com subtree. For example, the administrator user has the DN uid=Administr8or,cn=users,o=TIVIDS.

What to do next

- Configure VMM so that new users and groups that are created in the administrative console are written to the LDAP directory.

Configuring VMM to write to the LDAP directory

After you added the LDAP directory to the realm, make the LDAP directory the repository to which new users and groups are written. The Virtual Member Manager (VMM) component can read from multiple repositories but can write to only a single repository. You can then use the user management functions in the administrative console to create users and groups that are written to the LDAP directory.

You need to configure the mapping between the LDAP objects, such as users and groups, and the entity types of VMM that represent these objects. The entity types are used to map the objects in different LDAP directories to a common object model in VMM.

Before you begin

From your LDAP administrator, obtain the base entries in the LDAP subtree for users and groups. These base entries are the locations in the LDAP subtree where users and groups are created when users and groups are created through the Manage Users page or the Manage Groups page in the administrative console.

Ensure that the LDAP bind ID has write-permissions in the LDAP directory.

About this task

The VMM supported entity types are Group, OrgContainer, and PersonAccount. A Group entity represents a simple collection of entities that might not have any relational context. An OrgContainer entity represents an organization, such as a company or an enterprise, a subsidiary, or an organizational unit, such as a division, a location, or a department. A PersonAccount entity represents a human being. You cannot add or delete the supported entity types because these types are predefined.

The configuration steps in this task use the sample LDAP directory described in “Sample LDAP data” on page 42. Replace the values from this sample with your own.

Procedure

To map the LDAP object types to the entity types in VMM:

1. Log in as the administrative user, typically scadmin,.
2. Open the administrative console:
 - a. From the menu on the left, click **Setting > WebSphere Administrative Console**.
 - b. Click **Launch WebSphere Administrative Console**.
3. Click **Security > Global Security**

4. Under User account repository, select Federated Repositories from the **Available realm definitions** list and then click **Configure**. Then, click **Support entity types**.
5. Click each entity type and type the base entry from the LDAP directory in the **Base entry for the default parent**, as follows, replacing the sample base entries with the entries from your LDAP directory.

Entity type	Sample base entry
Group	cn=groups,ou=NetworkManagement, dc=myco,dc=com
OrgContainer	ou=NetworkManagement,dc=myco,dc=com
PersonAccount	cn=users,ou=NetworkManagement, dc=myco,dc=com

6. Save the configuration for each entity type.
7. Restart the server.

Results

Users and groups that are created in the administrative console are now written to the LDAP directory. It is now good practice to create users and groups only in the administrative console or by using the **tipcli** command-line utility.

What to do next

Assign Web GUI roles to the LDAP users so that they can access Web GUI functions, and so that they can be synchronized with the ObjectServer.

Assigning Web GUI roles to LDAP users and groups

Assign Web GUI roles to the LDAP users so that they have permission to use the Web GUI functions.

About this task

If you assign the roles to groups, the authorizations that are associated with the roles cascade to all users that are members of the groups.

The Web GUI roles do not give the users permission to write to the ObjectServer. This permission is needed for certain Web GUI functions, for example, the Active Event List (AEL) and the Web GUI tools. You set up this permission after you assigned the Web GUI roles.

Write-permission to the ObjectServer can be granted only to Web GUI users that have the ncw_admin role or the ncw_user role. Assign these roles to the users that you want to synchronize to the ObjectServer.

Procedure

To assign Web GUI roles:

1. To assign roles to user groups:
 - a. Click **Console Settings > Group Roles**.
 - b. Complete any combination of the search fields to help locate the groups.
 - c. Select how many groups to display and click Search. A list of groups appears in the grid.

- d. Click the name of the group you want to assign roles to.
 - e. From the **Role(s)** list, select the roles to assign the user group.
 - f. Click **Save**.
2. To assign roles to users:
 - a. Click **Console Settings > User Roles**.
 - b. Complete any combination of the search fields to help locate the users.
 - c. Select how many users to display and click Search. A list of matching users appears in the grid.
 - d. Click the user ID of the user you want to assign roles to.
 - e. From the **Role(s)** list, select the roles to assign the user.
 - f. Click **Save**.

What to do next

Create the LDAP users in the ObjectServer by enabling the user synchronization function.

Synchronizing LDAP users with the ObjectServer

After you defined the LDAP directory and assigned Web GUI roles to the LDAP users, enable the user synchronization function. This function creates the LDAP users in the ObjectServer, so that they can use all functions that write to the ObjectServer. These functions include the Active Event List (AEL) and the the Web GUI tools.

Before you begin

Ensure that the LDAP directory is running. If an ObjectServer was previously added to the realm as a user repository, it needs to be removed. See *Removing user registries* in the Tivoli Netcool/OMNIBus Knowledge Center.

Only Web GUI users that have the ncw_admin role or the ncw_user role can be synchronized. Ensure that you assigned these roles to the required users.

Procedure

To enable user synchronization:

1. Edit the `WEBGUI_HOME/etc/server.init` file and set the **users.credentials.sync** property to TRUE.
2. To change the name of the vmusers user group, assign the required value to the **users.credentials.sync.groupname** property.
3. Specify the intervals at which synchronization occurs:
 - a. Edit the `ncwDataSourceDefinitions.xml` file.
 - b. Set the `maxAge` attribute of the **config** property to the required time in seconds. For example:

```
<config maxAge="time"/>
```

The default is 3600 seconds.

4. Restart the server.
5. If your environment is load-balanced, to enable user synchronization against other nodes in the cluster repeat steps 1 to 4. On each additional node on which you enable user synchronization, change the name of the user group, as

described in step 2 on page 47. On each node of a load balanced environment, the name of the user group that contains the synchronized users must be unique.

Results

The LDAP users and groups are synchronized with the ObjectServers that are configured in the `ncwDataSourceDefinitions.xml` file. In an ObjectServer all synchronized users are assigned to the `vmusers` group (or, whichever name is specified by the `users.credentials.sync.groupname` property). If an ObjectServer does not already contain this user group, it is created automatically. Every 3600 seconds (or whichever refresh interval is specified by the `maxAge` attribute), the `vmusers` group is resynchronized with the ObjectServer.

What to do next

Perform the following tasks:

- To enable synchronized users to connect to the ObjectServer and modify ObjectServer data, for example by using the SQL interactive interface or by running Web GUI tools, assign the following ObjectServer user groups:
 - ISQL
 - ISQLWrite
- To secure your network by using Secure Socket Layer (SSL) encryption, enable SSL communications with the LDAP directory.
- To trigger a synchronization request manually, use the `WEBGUI_HOME/bin/webtop_osresynch` tool. Before you use this tool, configure the WAAPI client. The required `methodName` attribute is `osresync.refreshOSCache`.

Configuring an SSL connection to the OMNIBus ObjectServer

If the OMNIBus ObjectServer requires a Secure Sockets Layer (SSL) connection, you must configure Operations Analytics Predictive Insights probe to connect to the ObjectServer over SSL.

Before you begin

Before you configure the probe to connect to the ObjectServer over SSL, ensure that the following steps are complete:

- You exported the ObjectServer's SSL certificate to a file.
- You specified the directory path to the certificate file when you configured alarm forwarding to OMNIBus during the installation of the Analytics component.

Procedure

1. Edit the `$PI_HOME/probe/omnibus/probes/linux2x86/taspprobe.props` file. Add the following lines to the file:

```
< AuthUserName : <name>
< AuthPassword : <password>
```

where `<name>` and `<password>` are the username and password that the OMNIBus ObjectServer requires to authenticate the probe.
2. Copy the OMNIBus certificate keys to the probe directory. For example:

```
cd $NETCOOL_HOME/etc/security/keys/client/
cp * $PI_HOME/probe/etc/security/keys
```

3. To copy the OMNIbus interfaces file, which contains communication details for the ObjectServer, to the Operations Analytics Predictive Insights probes directory, enter the following commands:

```
cd $NETCOOL_HOME/etc/  
cp interfaces.linux2x86 $PI_HOME/probe/etc/interfaces.linux2x86
```
4. To start the probe, enter the following command:

```
$PI_HOME/bin/start.sh
```

Chapter 4. Configuring integrations with other products

You can set up Operations Analytics Predictive Insights to work with other IBM® products. Read about necessary configuration tasks required to set up the available integrations.

Configuring integration with IBM Integration Bus

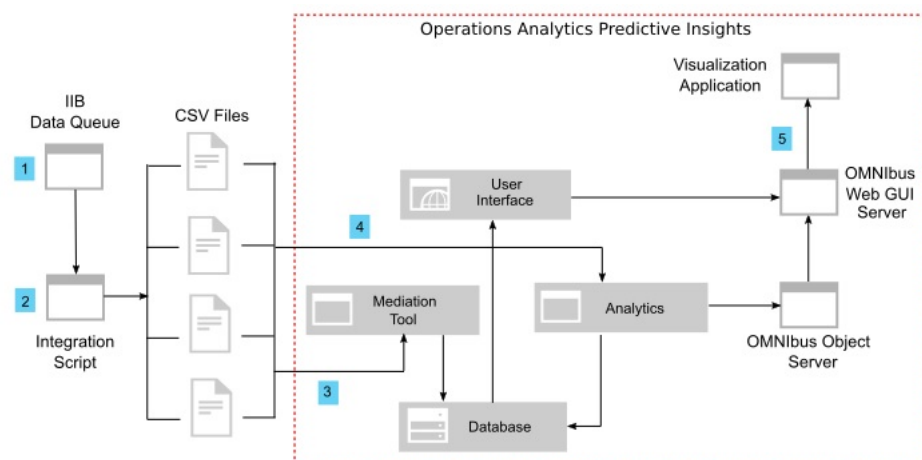
You can integrate Operations Analytics Predictive Insights with IBM Integration Bus to analyze accounting and statistical data that is collected by IBM Integration Bus.

IBM® Integration Bus (IIB) is an enterprise service bus that offers a fast, simple way for systems and applications to communicate with each other. As a result, it can help you achieve business value, reduce IT complexity and save money.

When you integrate Operations Analytics Predictive Insights with IBM Integration Bus, Operations Analytics Predictive Insights analyzes the accounting and statistical data that is collected by IBM Integration Bus. Operations Analytics Predictive Insights generates alarms when it identifies anomalies in the data that it analyzes. Alarms are displayed in the OMNIBus event list from where you can drill down to the Operations Analytics Predictive Insights User Interface to view the details of an anomaly.

Architecture

Read this topic for a brief overview of the component architecture and data flow that is created when you integrate Operations Analytics Predictive Insights and IBM Integration Bus.



1. Configure IBM Integration Bus to write statistical and metric information in XML format to a queue.
2. Run the integration script that is supplied in the integration package to read data directly from the queue and convert it to CSV format so that it can be analyzed by Operations Analytics Predictive Insights.
3. Configure Operations Analytics Predictive Insights to analyze the CSV data.

4. Load the data into Operations Analytics Predictive Insights. If there are more than two weeks of data available, Operations Analytics Predictive Insights can detect anomalies and send them to the OMNIBus event list.
5. You can use a visualization application to display the OMNIBus event list. From the event list, you can drill down to analyze an anomaly in the Operations Analytics Predictive Insights User Interface.

Setting up the integration script

Read this topic to complete the steps to set up the integration script.

About this task

You must copy both the integration script, `IIB_Integration.sh`, and a number of JAR files from the IBM Integration Bus installation directory to the installation directory of the Analytics component. You must also create a directory for the integration script to store the XML and CSV files that it creates.

Procedure

1. As the Operations Analytics Predictive Insights administrative user, typically `scadmin`, copy the `PI-IIB-Accelerator-V1.0.zip` file to the Analytics server.
2. Unzip the `PI-IIB-Accelerator-V1.0.zip` file.
3. Copy the integration script, `IIB_Integration.sh`, from the directory to which you unzipped the package to the `$PI_HOME/bin` directory of your Analytics installation. For example, `/opt/IBM/scanalytics/analytics`.
4. To ensure that you have execute permissions on the `IIB_Integration.sh` file, enter the following command:

```
chmod 755 IIB_Integration.sh
```

5. Copy the `IIBRead.jar` file from the directory to which you unzipped the package to the `$PI_HOME/lib` directory of the Operations Analytics Predictive Insights Analytics installation.
6. Copy the following JAR files from your IBM Integration Bus installation to the `$PI_HOME/lib` directory on the Analytics server:

Note: These files are normally located in the `java/lib` directory of your IBM Integration Bus installation. For example, `/opt/mqm/java/lib`.

7. To verify that the script is installed correctly, run it without specifying any attributes:

```
$PI_HOME/bin/IIB_Integration.sh
```

If the script is installed correctly, it returns the usage information for the script.

8. Create a directory to store the CSV files that are created by the `IIB_Integration.sh` script. For example:

```
mkdir /opt/IBM/CSV
```

This directory is referred to as \$PATH\$ in the remainder of this document. When the IIB_Integration.sh script runs, it writes the XML data to a file called <CurrentTimeandDate>_QueueBackup.xml in this directory.

Configuring real time metric collection and analysis

In this scenario, metrics are collected from IBM Integration Bus at 10-minute intervals and immediately analyzed by Operations Analytics Predictive Insights to check for anomalies in the data.

Configuring a statistics queue on MQ/IIB

To configure accounting and statistics data collection in XML format for IBM Integration Bus, you must define a local WebSphere MQ queue, an appropriate subscription, and enable publication of the data on your integration node.

About this task

Procedure

1. Create a WebSphere MQ queue to store the messages and create a subscription to subscribe to the appropriate topic string. Set the newly created queue as the destination for those messages. For example, to subscribe to archive statistics data from all integration nodes, integration servers and message flows use the following commands:

```
runmqsc <queueManager>
define qlocal(STATS) REPLACE
define sub(STATS_SUB) TOPICSTR('$SYS/Broker/+/StatisticsAccounting/Archive/#') DEST(STATS)
REPLACE
```

The '+' and '#' characters are wildcards that you can use within your topic string. For more information about topic string options and general accounting and statistics configuration, see *Message flow statistics and accounting data* in the IBM Integration Bus Knowledge Center.

2. To turn on accounting and statistics data publication, use the mqsichangeflowstats command. For example, to turn on accounting and statistics data publication for all integration servers and message flows on a specific integration node, enter the following command:

```
mqsichangeflowstats <integrationNodeName> -a -c active -g -j -n advanced -t basic -b basic -o
```

For more information, see *mqsichangeflowstats* in the IBM Integration Bus Knowledge Center.

3. You can set the interval at which the accounting and statistics archive data is collected on the integration node. The default interval is 60 minutes. You can specify an interval in the range 1 - 43200 minutes. For example, to set the data collection interval of an existing integration node to a 15-minute interval, enter the following command:

```
mqsichangebroker <integrationNodeName> -v 15
```

At the end of this interval, the recorded statistics are written to the output directory and the interval is restarted.

For more information, see the mqsicreatebroker or the mqsichangebroker command in the IBM Integration Bus Knowledge Center

Configuring Operations Analytics Predictive Insights data mediation

Mediation is the process of collecting data to analyze.

About this task

The integration pack includes a preconfigured model that specifies which IBM Integration Bus data to analyze. You must install, update, and deploy the preconfigured model. The model has one data source with two metric groups for MessageFlow and Nodes.

Note: To deploy the model, you must first create a topic for the model.

Procedure

1. To create a topic, enter the following command:
`$PI_HOME/bin/admin.sh create_topic <topic name>`
2. On the Analytics server, create a directory in the home directory of the Operations Analytics Predictive Insights administrative user. For example,
`mkdir /home/scadmin/iib_model`
3. Copy the `datasource.pamodel` file from the integration package to the new directory.
4. Edit the `datasource.pamodel` file and customize the value of `$TOPIC$` and `$PATH$`. Replace `$TOPIC$` with the name of the topic you created in step 1. Replace `$PATH$` with the path to the CSV files directory that you created in step 8 of: “Setting up the integration script” on page 52.

```
tasp:TopicConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:tasp="http://www.ibm.com/tivoli/netcool/tasp">
  <name>$TOPIC$</name>
  <DataSourceInstance>

  <config xsi:type="tasp:FileSystemConfig" timeFormat="yyyyMMddHHmmss">
    <name>IIB</name>
    <path>$PATH$</path>
    <pattern>([a-zA-Z]\w*)_(\d{14}).*.csv</pattern>
  </config>
```

For example:

```
<name>ALL</name>
<path>/opt/IBM/CSV</path>
```

5. To configure the topic, enter the following command:
`$PI_HOME/bin/admin.sh set -t=ALL system.aggregation.interval 10 -f`
6. To start the Mediation tool, enter the following command:
`$PI_HOME/bin/mediationtool.sh`
7. Load the `datasource.pamodel` file into the Mediation tool:
8. To deploy the model, from the menu, choose **Predictive Insights > Deploy Model**.

Scheduling the integration script to run

You create a crontab entry to schedule the integration script to run at a recommended interval of 10-minutes.

Procedure

1. On the Analytics server, as the administrative user, typically `scadmin`, edit the crontab file:

```
crontab -e
```

2. Add the following line to the crontab file:

```
* /10 0 * * * . /home/<PI_USER>/bash_profile; cd $PI_HOME/bin && ./IIB_Integration.sh
-h <Host> -p <Port> -c <Channel> -m <QueueManagerName> -q <QueueName> -d <TargetDirectory>
```

where:

- <PI_USER> is the Operations Analytics Predictive Insights administrative user, typically scadmin.
- <Host> is the IP address or FQDN of the MQ/IIB host server.
- <Port> is the channel listener port. Each channel listener runs on a particular port.
- <Channel> is the channel listener. A channel listener program listens for incoming network requests and starts the appropriate receiver channel when it is needed.
- <QueueManagerName> is the associated queue manager.
- <QueueName> is the queue that collects the IIB Statistics.
- <TargetDirectory> is the full path to the location of the generated XML and CSV files.

3. Save the crontab file.

Example

```
10 0 * * * . /home/scadmin/.bash_profile; cd /opt/IBM/scanalytics/analytics/bin && ./IIB_Integrati
```

Loading the data into Operations Analytics Predictive Insights

Operations Analytics Predictive Insights uses the model configuration to determine which data to load.

Procedure

1. To start the topic to which you deployed the model, enter the following command:

```
$PI_HOME/bin/start.sh -t=<topic name>
```

2. To start loading the data into Operations Analytics Predictive Insights:

- If a backlog of CSV files exists in the CSV directory, enter the following command:

```
$PI_HOME/bin/admin.sh run_extractor_instance -s=<start_time> -t=<topic name>
```

where <start_time> is the earliest time stamp in the CSV file names. Specify the start time in the format: yyyyMMdd-hhmm. For example, 20160201-1600.

- If no CSV file backlog exists, enter the following command:

```
$PI_HOME/bin/admin.sh run_extractor_instance -t=<topic name>
```

Configuring backlog collection of metrics

In this scenario, Operations Analytics Predictive Insights collects a backlog of metrics from IBM Integration Bus and performs a once-off analysis of the metrics.

Backing up the XML statistics data from the IIB Queue

About this task

You can choose one of the following methods to backup the XML data from the Statistics Queue on IIB/MQ:

1. Running the IIB_Integration.sh script in queue mode, which removes the XML from the Queue during backup and generates CSV files.
2. Using the IIB rfutil which leaves the XML messages on the Queue.

Procedure

1. If you want to run the integration script to back up the XML statistics data, enter the following command:

```
IIB_Integration.sh -h <Host> -p <Port> -c <Channel> -m <QueueManagerName> -q <QueueName> -d <TargetDirectory>
```

where:

- <Host> is the IP address or FQDN of the MQ/IIB host server.
- <Port> is the channel listener port. Each channel listener runs on a particular port.
- <Channel> is the channel listener. A channel listener program 'listens' for incoming network requests and starts the appropriate receiver channel when it is needed.
- <QueueManagerName> is the associated queue manager.
- <QueueName> is the queue that collects the IIB Statistics.
- <TargetDirectory> is the full path to the location of the generated XML and CSV files.

Note: Each time the IIB_Integration.sh script runs in queue mode, it backs up the XML data to a file called <CurrentTimeandDate>_QueueBackup.xml and CSV files are created.

2. If you want to use the rfhutil utility to back up the XML statistics data, download the utility via the following support pack.

Running the Integration Script

You must run the integration script to convert the XML data to CSV format.

Procedure

To run the integration script, enter the following command:

```
$PI_HOME/bin/IIB_Integration.sh <-f Filename> <-d TargetDirectory>
```

where:

- Filename is the full path to the file that contains the exported XML from the Statistics Queue.
- TargetDirectory is the full path to the location of the generated CSV files

Configuring Operations Analytics Predictive Insights data mediation

Mediation is the process of collecting data to analyze.

About this task

The integration pack includes a preconfigured model that specifies which IBM Integration Bus data to analyze. You must install, update, and deploy the preconfigured model. The model has one data source with two metric groups for MessageFlow and Nodes.

Note: To deploy the model, you must first create a topic for the model.

Procedure

1. To create a topic, enter the following command:

```
$PI_HOME/bin/admin.sh create_topic <topic name>
```

2. On the Analytics server, create a directory in the home directory of the Operations Analytics Predictive Insights administrative user. For example,
`mkdir /home/scadmin/iib_model`
3. Copy the `datasource.pamodel` file from the integration package to the new directory.
4. Edit the `datasource.pamodel` file and customize the value of `$TOPIC$` and `$PATH$`. Replace `$TOPIC$` with the name of the topic you created in step 1. Replace `$PATH$` with the path to the CSV files directory that you created in step 8 of: "Setting up the integration script" on page 52.

```
tasp:TopicConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:tasp="http://www.ibm.com/tivoli/netcool/tasp">
<name>$TOPIC$</name>
<DataSourceInstance>

<config xsi:type="tasp:FileSystemConfig" timeFormat="yyyyMMddHHmmss">
<name>IIB</name>
<path>$PATH$</path>
<pattern>([a-zA-Z]\w*)_(\d{14}).*\..csv</pattern>
</config>
```

For example:

```
<name>ALL</name>
<path>/opt/IBM/CSV</path>
```

5. To configure the topic, enter the following command:
`$PI_HOME/bin/admin.sh set -t=ALL system.aggregation.interval 10 -f`
6. To start the Mediation tool, enter the following command:
`$PI_HOME/bin/mediationtool.sh`
7. Load the `datasource.pamodel` file into the Mediation tool:
8. To deploy the model, from the menu, choose **Predictive Insights > Deploy Model**.

Starting the topic

You must start the topic to which you deployed the model before you can load data in to Operations Analytics Predictive Insights.

Procedure

To start the topic, enter the following command:

```
$PI_HOME/bin/start.sh -t=<topic name>
```

Loading the data into Operations Analytics Predictive Insights

Operations Analytics Predictive Insights uses the model configuration to determine which data to load.

Procedure

To start loading the data into Operations Analytics Predictive Insights, enter the following command:

```
$PI_HOME/bin/admin.sh run_extractor_instance -s=<start_time_date> -e=<end_time_date> -t=ALL
```

where `<start_time_date>` is the earliest time stamp in the CSV file names and `<end_time_date>` is the latest time stamp in the CSV file names. Enter the start and end times in the format: `yyyyMMdd-hhmm`.

Configuring integration with IBM Performance Management

You can integrate Operations Analytics Predictive Insights with IBM® Performance Management to analyze data that is collected by IBM® Performance Management.

Operations Analytics Predictive Insights analyzes the metric data that is collected by the Performance Management application and generates alarms when it identifies anomalies in the data. Anomalies are displayed in the Performance Management Dashboard as individual events associated with the affected application, group, or component. You can then drill down to the Operations Analytics Predictive Insights User Interface to view more details on an anomaly.

Before you begin

Before you start to integrate Operations Analytics Predictive Insights and IBM Performance Management, you must complete the following tasks:

- Install Operations Analytics Predictive Insights. For more information, see the *Installing* documentation in the Operations Analytics Predictive Insights Knowledge Center.
- Install IBM Performance Management. For more information, see the IBM Performance Management Knowledge Center.
- Create a new directory on the Operations Analytics Predictive Insights server to hold the Operations Analytics Predictive Insights Accelerator for IBM Performance Management package, for example: `mkdir ~/PI-APM-Accelerator-v1.0`
- Download and extract the IBM Operations Analytics Predictive Insights Accelerator for IBM Performance Management package, `PI-APM-Accelerator-v1.0.zip`, to the new directory.

For more information, see *Downloading Operations Analytics Predictive Insights*.

Configuring an Operations Analytics Predictive Insights model for Performance Management data sources

How to configure and deploy an Operations Analytics Predictive Insights model for Performance Management data sources.

About this task

A model file, `apm_model.pamodel`, for Performance Management data sources is included in the integration pack. You must edit the model file to specify the details of the two data sources that Operations Analytics Predictive Insights can use to process data generated by Performance Management agents. The two data sources are:

- the pre-fetch database
- a REST service referred to as the Kafka service bus

The file system data source must be a local directory on the Analytics server that you mount via NFS on the Performance Management server. Later, you must install a file consumer program on the Performance Management Server to read metric data from the Kafka service bus and write it to the file system.

Procedure

1. Log in to the Operations Analytics Predictive Insights as the administrative user, typically `scadmin`.

2. Enter the following command to create an Operations Analytics Predictive Insights topic called ALL:
`$PI_HOME/bin/admin.sh create_topic ALL`
3. Enter the following commands to configure the topic:
`$PI_HOME/bin/admin.sh set -t=ALL system consolidate.alarms.correlations.enabled false -f`
`$PI_HOME/bin/admin.sh set -t=ALL system consolidate.alarms.incident.enabled false -f`
`$PI_HOME/bin/admin.sh set -t=ALL system consolidate.alarms.related_events.enabled false -f`
4. Copy the `apm_model.pamodel` file to the `/home/scadmin/apm_model` directory on the Operations Analytics Predictive Insights server.
5. Edit the `apm_model.pamodel` file and customize the following values for your APM environment:
`<driver>com.ibm.db2.jcc.DB2Driver</driver>`
`<url>jdbc:db2://<APM_prefetch_ip_address>:<DB2_port>/WAREHOUS</url>`
`<schema>ITMUSER</schema>`
`<username><APM_prefetch_db_user></username>`
`<password><APM_prefetch_db_user_password></password>`

Where:

`<PI_install_dir>` is the directory where Operations Analytics Predictive Insights was installed, by default `/opt/IBM/scanalytics/analytics`

`<APM_prefetch_ip_address>` is the IP address of a server where the APM prefetch database is installed

`<DB2_port>` is the port used by Performance Management prefetch DB2 database, by default 50000

`<APM_prefetch_db_user>` is the Performance Management prefetch database user, by default `itmuser`

`<APM_prefetch_db_user_password>` is the password for the APM prefetch database user, by default `db2Usrpasswd@08`

6. Go to the end of the file and replace `$APM_KAFKA_DS` with `/var/apm/files`
`<config xsi:type="tasp:FileSystemConfig" tzone="UTC"`
`timeFormat="yyyyMMdd-HH-mm"> <name>KAFKA</name> <path>$APM_KAFKA_DS</path>`
`<pattern>([a-zA-Z]\w*)__.*__(\d{8}-\d{2}-\d{2})UTC.*csv</pattern>`
`</config>`
7. Save the `apm_model.pamodel` file.
8. Make sure that you have write permission to the `/var/` directory and create the following directory on the PI Server:
`mkdir -p /var/apm/files`
`chmod 777 /var/apm/files`
9. To import the `apm_model.pamodel` file in to the Mediation Tool, choose **File ->Import ->Predictive Insights -> Predictive Insights Model**.
10. To deploy the model, from the menu, choose **Predictive Insights->Deploy Model**. If multiple topics exist, select the topic to which you want to deploy the model.

Exporting a directory from the Operations Analytics Predictive Insights server

You must export the `/var/apm/files` directory from the Operations Analytics Predictive Insights server so you can mount it on the Performance Management server. Operations Analytics Predictive Insights uses the mounted directory as a datasource for metric data provided by the Kafka service bus.

Procedure

1. Log in to the Operations Analytics Predictive Insights as root.
2. Add the following line to the `/etc/exports` file to export the `/var/apm/files` directory.

```
/var/apm/files <PM_ip_address>(rw, sync)
```

Where `<PM_ip_address>` is the IP address of the Performance Management Server.

3. Enter the following command to reload the NFS service:

```
/etc/init.d/nfs stop  
/etc/init.d/nfs start
```

4. Enter the following command to start NFS during server startup:

```
chkconfig --level 3 nfs on  
chkconfig --level 5 nfs on
```

5. If you don't want to stop the firewall, add the following rule to iptables:

```
iptables -I INPUT 1 -s <PM_ip_address> -p tcp --dport 2049 -j ACCEPT
```

Note: You can add this rule permanently to the default iptables configuration in `/etc/sysconfig/iptables`:

```
:OUTPUT ACCEPT [0:0]  
-A INPUT -s <PM_ip_address> -p tcp --dport 2049 -j ACCEPT
```

After you add the rule, enter the following command to reload iptables:

```
/etc/init.d/iptables restart
```

Mounting a directory on the Performance Management Server

You must mount the directory that you exported from the Operations Analytics Predictive Insights server on the Performance Management Server.

About this task

Procedure

1. Log in to the Performance Management server.
2. Enter the following command to create a mount point:

```
mkdir -p /mnt/scapi/DefaultInstance
```

3. Enter the following command to mount the directory that you exported from the Operations Analytics Predictive Insights server:

```
mount -o tcp <PI_ip_address>:/var/apm/files /mnt/scapi/DefaultInstance
```

Where `<PI_ip_address>` is the IP address of the Operations Analytics Predictive Insights server.

4. To ensure that the folder is mounted permanently, add the following line to the `/etc/fstab` file:

```
<PI_ip_address>:/var/apm/files /mnt/scapi/DefaultInstance nfs rsize=8192, wsize=8192, timeo=14
```

5. To verify that the mount point has been set correctly, do the following:
 - a. On the Performance Management server, as root, enter the following command to create a test file:

```
echo "text" > /mnt/scapi/DefaultInstance/test
```
 - b. On the Operations Analytics Predictive Insights server, as the administrative user, typically scadmin, enter the following commands to view and delete the test file:

```
cat /var/apm/files/test  
rm -f /var/apm/files/test
```
- If you can create the test file, display its content and delete the file, the mount is configured correctly.

Configuring the File Consumer program

On the Performance Management server, the file consumer program reads metric data from the Kafka service bus and write it to the Operations Analytics Predictive Insights server directory that is mounted on the Performance Management Server.

Procedure

1. On the Performance Management Server, log in as root.
2. Copy the `com.ibm.tivoli.ccm.consumer.scapi.scapifileconsumer-1.0.eba` file from the directory on the Operations Analytics Predictive Insights server where you extracted the `PI-APM-Accelerator-v1.0.zip` file to the `<PM_installdir>/usr/servers/asfrest/dropins/` folder.
where `<PM_installdir>` is the directory where you installed Performance Management. The default installation directory is: `/opt/ibm/wlp`.
3. Enter the following command to restart the **asfrest** process.

```
apm stop asfrest  
apm start asfrest
```

Setting up the Event Integration Facility (EIF) Gateway

The EIF Gateway forwards events that the OMNIBus ObjectServer receives from Operations Analytics Predictive Insights to the Performance Management application.

Installing the EIF gateway

Read this topic to complete the steps to install the EIF gateway.

Procedure

1. Navigate to the directory that contains the extracted IBM Operations Analytics Predictive Insights1.3.5 installation package.
2. Unzip the EIF package file, `im-nco-g-tivoli-eif-4_0.zip` file.
3. Enter the following command to start the Installation Manager wizard:

```
./IBMIM
```
4. In the main Installation Manager window, click **File-> Preferences-> Repositories** and add the `repository.config` file to the repository.
5. Follow the installation wizard and select the default options to complete the installation.

Configuring the EIF gateway

Read this topic for instructions on how to configure the EIF gateway.

Procedure

1. Copy the configuration file from <omnihome>/gates/tivoli_eif/NCO_GATES.props to <omnihome>/etc/NCO_GATES.props.
2. Uncomment the Common Netcool®/OMNIBus properties. Ensure that the path property is correct.
3. Uncomment the Common Gateway properties. Ensure that the path property is correct.
4. Uncomment the TIVOLI EIF Gateway properties. Ensure that the path property is correct.

Loading the ObjectServer customization file

Before you run the EIF gateway for the first time, you must load the ObjectServer customization file.

About this task

The gateway provides an ObjectServer customization file called `tivoli_eifgw_setup.sql`. This SQL file ensures that the ObjectServer is customized for the EIF gateway.

Procedure

To load the SQL commands that are in the <omnihome>/gates/tivoli_eif/tivoli_eifgw_setup.sql file, enter the following command:

```
<omnihome>/bin/nco_sql -server NCOMS -user root -password '' < <omnihome>/gates/tivoli_eif/tivoli_ei
```

Updating the gateway map definition file

The map definition file defines how the gateway transfers data between ObjectServer tables and target databases or applications.

Procedure

1. Edit the \$OMNIHOME/gates/tivoli_eif/tivoli_eif.map file.
2. Replace the contents of the file with the following text:

```
CREATE LOOKUP Status_Table (
    {0 , 'N'})
DEFAULT = 'Y' ;

CREATE MAPPING StatusMap
(
    'situation_name'                = 'Anomaly Detected',
    'situation_origin'              = '@TASPAnomalousResources',
    'server_handle'                  = '',
    'date_reception'                 = '@FirstOccurrence',
    'event_handle'                   = '',
    'extdata1'                       = '@TASPIDentifier',
    'source'                         = 'IBM SCAPI',
    'sub_source'                     = '@TASPAnomalousResources',
    'sub_origin'                     = '',
    'hostname'                       = '',
    'last_modified_time'             = '@InternalLast',
    'adapter_host'                   = '',
    'situation_status'               = Lookup('@Severity','Status_Table'),
    'administrator'                  = '',
    'acl'                             = '',
    'severity'                       = '@Severity',
    'date'                           = '@LastOccurrence',
    'duration'                       = '0',
    'situation_displayitem'          = '@TASPAnomalousMetrics',
```

```

'msg' = '@Summary',
'msg_catalog' = '',
'msg_index' = '0',
'num_actions' = '1',
'credibility' = '0',
'repeat_count' = '0',
'situation_time' = '@LastOccurrence',
'cause_date_reception' = '0',
'integration_type' = 'U',
'situation_type' = 'S',
'appl_label' = 'PI:A:S',
'master_reset' = '',
'situation_thrnode' = '@TASPAAnomalousResources',
'cause_event_handle' = '0',
'from_user' = 'root',
'to_user' = 'root');

```

3. Save the \$OMNIHOME/gates/tivoli_eif/tivoli_eif.map file.

Customizing the EIF configuration file

You must customize the EIF configuration file with the Performance Management server's details.

Procedure

1. Edit the <omnihome>/gates/tivoli_eif/tivoli_eif_config file.
2. Enter the IP address of the Performance Management server in the following line:

```
c1ServerLocation=<PM Server IP>
```

For example:

```
c1ServerLocation=192.168.1.1
```

3. Enter the destination EIF port on the Performance Management server in the following line:

```
c1Port=<port>
```

Starting the EIF gateway

Complete the following steps to start the EIF Gateway.

Procedure

1. Change to the bin directory:

```
cd <omnihome>/bin
```
2. To start the gateway, enter the following command:

```
./nco_g_tivoli_eif &
```
3. To verify that the gateway started successfully, enter the following command to check that the gateway process is running:

```
ps -ef | grep nco_g_tivoli_eif
```

Updating the probe rules file

Read this page for information on how to update the probe rules file for integration with IBM® Performance Management.

Procedure

1. Log in to the Operations Analytics Predictive Insights server as the administrative user, typically scadmin.
2. Back up the <PI_installdir>/probe/omnibus/probes/linux2x86/stdin-tasp.rules file:

```
cp stdin-tasp.rules stdin-tasp.rules.backup
```

3. Edit the stdin-tasp.rules file.
4. Add the following contents to the file before the last }.

```
# Added for APM integration
$PrimaryResource = extract($TASPMetricResourceNameList + ";", "([^;]*)");
@TASPAnomalousResources = $Node
update(@TASPAnomalousResources)

if (match($Node,$PrimaryResource)) {
# Do nothing as non-compound resource key
} else {
# Extract secondary resource from compound resource key
$last_index = length($PrimaryResource) + 1
$node_index = length($Node) + 2
$node_length = length($Node)
$SecondaryResource = substr($PrimaryResource,int($node_index),int($last_index)-int($node_length))
@TASPAnomalousMetrics = @TASPAnomalousMetrics + "_Subsystem[" + $SecondaryResource + "]"
update(@TASPAnomalousMetrics)

}

# Set TASPAnomalousMetrics to Node for consolidated alarms (for clearing purposes)
if (regmatch(@Class, "89360")) {
@TASPAnomalousMetrics = @Node
update(@TASPAnomalousMetrics)
}

# Clear and children of a consolidated alarm to prevent them displaying in APMUI
if (regmatch(@TASPParentIdentifier, ".*TASP.*")) {
@Severity = 0
update(@Severity)
@Type = 2update(@Type)
}
```

Configuring integration properties

How to configure the integration properties on the Performance Management server.

About this task

You must set integration properties to:

- Integrate the Performance Management User Interface with the Operations Analytics Predictive Insights User Interface.
- Set the URL for the Operations Analytics Predictive Insights Service Diagnosis view. This URL will be used to launch in context from a Performance Management event to the Operations Analytics Predictive Insights User Interface.

Procedure

1. Log in to the Performance Management server as root.

2. Enter the following commands to set the URL for the Operations Analytics Predictive Insights Service Diagnosis UI:

Where:

<PI_ip_address> is the IP address of the server where Operations Analytics Predictive Insights is installed.

https://<PM_ip_address>:8091/1.0/monitoring/systemconfig/services/com.ibm.tivoli.ccm.apmui/con

Where <PM_install_dir> is the directory where Performance Management is installed, by default /opt/ibm/wlp.

```
enableSCAPIIntegration = true
```

```
apm stop apmui
apm start apmui
apm stop server1
apm start server1
```

Tuning the Operations Analytics Predictive Insights server

Procedure

Scheduling cleanup of Performance Management data files

After you integrate Operations Analytics Predictive Insights with IBM Performance Management, create cron entries to perform periodic cleanup of data files on the Operations Analytics Predictive Insights server.

Procedure

1. Log in to the Operations Analytics Predictive Insights server as the administrative user, typically `scadmin`.
2. Enter the following command to edit the crontab file:
`crontab -e`
3. Add the following lines to the crontab file:

```
30 0 * * * export PI_HOME=/opt/IBM/scanalytics/analytics;  
/bin/find /var/apm/files/*d -type f -name "*.csv" -mtime +1 -print0 |xargs -0 gzip -f  
45 0 * * * export PI_HOME=/opt/IBM/scanalytics/analytics;  
/bin/find /var/apm/files/*d -type f -name "*.csv.gz" -mtime +7 -print0 |xargs -0 rm -f  
55 0 * * * export PI_HOME=/opt/IBM/scanalytics/analytics;  
/bin/find /var/apm/files/*d -type f -name "*.tmp" -mtime +7 -print0 |xargs -0 rm -f
```
4. Save the crontab file.

Chapter 5. Configuring a cloned Operations Analytics Predictive Insights server

If you clone an Operations Analytics Predictive Insights server, you must reconfigure Operations Analytics Predictive Insights for the new host name.

Before you begin

On the cloned server, change the host name in the `/etc/hosts` file.

About this task

This procedure assumes that all the Operations Analytics Predictive Insights components and prerequisite applications are installed on the cloned server.

Procedure

1. Log in as root.
2. Create a text file with the following content:

```
#!/bin/ksh

NEW_HOSTNAME=server.ibm.com
NETCOOL_DIR=/opt/IBM/tivoli/netcool/
UI_HOME=/opt/IBM/scanalytics/ui
PI_HOME=/opt/IBM/scanalytics/analytics/
DB2_OWNER=db2inst1
STREAMS_USER=scadmin
echo "Update DB2 hostname"

su - "$DB2_OWNER" -c "/home/$DB2_OWNER/sqlllib/instance/db2iset -g DB2SYSTEM=$NEW_HOSTNAME"

echo "Update OMNibus hostname"

su - "$STREAMS_USER" -c "cd $NETCOOL_DIR/etc; \
sed s/omnihost/$NEW_HOSTNAME/g omni.dat.ORIG > omni.dat; \
$NETCOOL_DIR/bin/nc_igen -out $NETCOOL_DIR/etc/interfaces"
cp $NETCOOL_DIR/etc/interfaces $PI_HOME/probe/etc/interfaces.linux2x86

echo 'Update $UI_HOME/wlp/usr/servers/piserver/taspConfig.xml file'
sed -i "s~serverName=\\.\\.\\.~serverName=\\$NEW_HOSTNAME\\$~" $UI_HOME/wlp/usr/servers/piserver/t

echo 'Update omnibus.connection.properties file'
sed -i "s~Tds:[^:]*::~Tds:$NEW_HOSTNAME::~" $PI_HOME/config/omnibus.connection.properties

echo 'Update tasp.connection.properties file'
sed -i "s~//[^:]*::~//$NEW_HOSTNAME::~" $PI_HOME/config/tasp.connection.properties

echo "Update Streams hostname" do

    echo "$NEW_HOSTNAME", build, execution > "/home/$STREAMS_USER/.streams/config/hostfile"

echo "Update ana_topic hostname"

rm -rf /tmp/tmpmodeldeploy
mkdir /tmp/tmpmodeldeploy

HOSTNAME="$NEW_HOSTNAME"

# Need to start db2 so we can update the ana_topic table
```

```

su - "$DB2_OWNER" -c ". /home/$DB2_OWNER/sql/lib/db2profile; db2start"

echo "update ana_topic set machine_name = '$HOSTNAME';" > /tmp/tmpmodeldeploy/update.sql
chmod 777 /tmp/tmpmodeldeploy/update.sql
echo "$PI_HOME/bin/admin.sh sql_loader tasp.connection.properties /tmp/tmpmodeldeploy/update.
chmod 777 /tmp/tmpmodeldeploy/runupdate.sh
su - "$STREAMS_USER" -c "/tmp/tmpmodeldeploy/runupdate.sh" > /tmp/tmpmodeldeploy/topicresult
cat /tmp/tmpmodeldeploy/topicresults
TOPICS=`cat /tmp/tmpmodeldeploy/topicresults`
echo $TOPICS

if [ ! -z "$TOPICS" ]
then
    echo DB connection error

fi

```

3. Update the following lines in the file to match the host name of the server, the directory paths of the Operations Analytics Predictive Insights installation, and the DB2 owner and streams user accounts.

```

NEW_HOSTNAME=server.ibm.com
NETCOOL_DIR=/opt/IBM/tivoli/netcool/
UI_HOME=/opt/IBM/scanalytics/ui
PI_HOME=/opt/IBM/scanalytics/analytics/
DB2_OWNER=db2inst1
STREAMS_USER=scadmin

```

4. Save the file as clone.sh and make the file executable.
5. Run the clone.sh script. When prompted, enter the passwords for the DB2_OWNER and STREAMS_USER accounts that you configured in step 3.
6. To start the Analytics component, log in as the administrative user, typically scadmin, and enter the following command:

```
$PI_HOME/bin/start.sh
```

Chapter 6. Analyzing data

After you configure your model and deploy it to the database, the next step is to start consuming and analyzing data.

Procedure

1. Set the main configuration parameters for Operations Analytics Predictive Insights.
 - a. Set the aggregation interval.
 - b. Configure purge
2. Start Operations Analytics Predictive Insights and run the extractor
3. Monitor the training and extraction status using the AEL

What to do next

The details of how to implement these steps are contained in the following sections.

Setting the aggregation interval

Data is normalized to the same interval so it can be processed by an analytics instance.

Procedure

1. Log in as the administrative user, typically `scadmin`.
2. Change to: `$PI_HOME/bin`.
Where `$PI_HOME` is the installation location of the analytics component.
3. Run the **set** command using the Operations Analytics Predictive Insights administration CLI application by entering the command:
`./admin.sh set -t=<topic name> system.aggregation.interval <desired interval>`
Where
 - `<topic name>` is the name of the topic.
 - `<desired interval>` is the aggregation interval in minutes.

Example

```
./admin.sh set -t=Topic1 system.aggregation.interval 15
```

Consider how often data is arriving when you set the aggregation interval. You must try to match the aggregation interval with the arrival rate of data within your system. The available intervals are 5, 10, 15 and 60 minutes.

Note: With a 60-minute aggregation interval, the data is heavily averaged and will not show small changes. As a result, it is more difficult for Operations Analytics Predictive Insights to perform early detection of anomalies.

If you use a 60-minute aggregation interval, it is recommended that you make the following configuration changes:

- Set the **nofm.error.count.min.size** property to 2
- Set the **nofm.error.count.window.size** property to 3.

With these settings, an alarm is generated if a metric is anomalous for two of the last three intervals, instead of the default, which requires that a metric is anomalous for three of the last six intervals. As a result, Operations Analytics Predictive Insights can detect anomalies that exist for a shorter period of time and generate alarms accordingly.

Related reference:

The set CLI command

The set CLI command.

Configuring purge

Configuring the data purge option. The purge is an optional step to configure how long data is stored in the database. Retaining data for a longer period requires more disk space and can impact UI response times.

About this task

This task describes how to set the details of data purge. The default purge setting is to purge metric data older than 15 days.

Note: When loading backlog data, the timestamps for the alarms may not match the timestamps of the associated data and as such the alarms will be retained while the data is purged. Note that alarms deleted from the Operations Analytics Predictive Insights database will no longer be accessible from your Event Management system and so you should ensure a similar cleanup is performed on your Event Management system.

To configure how long data is stored in the database:

Procedure

1. Change to the `$PI_HOME/bin/` directory.
Where `$PI_HOME` is the installation location of the analytics component.
2. Set the **system.metric.retention.days** property using the **set** command:

```
./admin.sh set -t=<topic_name> system.metric.retention.days <number_of_days>
```


Where `<topic_name>` is the name of the topic and `<number_of_days>` is the number of days you wish to retain data.

Note: For more information about the **set** command, see “set” on page 108.

Extracting data

Use the `run_extractor_instance` command to extract data.

About this task

The data extracted from your datasource is consumed in two ways, it is used to train your analytics model, and when an analytics model is available, the data is then analyzed for any anomalies. There are five types of extraction:

- **Steady state mode:** This form of extraction is the simplest and involves Operations Analytics Predictive Insights reading in current, incoming data from your datasource. The main issue with this method is it takes time for your system to train and create an analytics model for each algorithm, which is required before the algorithm can start to identify anomalies. Operations Analytics Predictive Insights creates a model for an individual algorithm when it

has consumed fifty percent of the number of days data specified for the **<algorithm name>maxTrainingWindowDays** configuration property.

- **Backlog mode:** This form of extraction involves the reading in of recent historical data. If you are intent on processing only backlog data, you must set the start and end time when calling the command **run_extractor_instance** for the first time. Only supply the stop time and not the start time when making any subsequent calls to the command **run_extractor_instance**.
- **Switch:** This form of extraction is a hybrid of steady state mode and backlog mode, the purpose of which is to create an analytics model as quickly as possible using Backlog mode and move on to Steady state mode to begin the analysis of current data.
- **Extract Only:** If you have multiple data sources with a variety of latencies and you wish to extract backlog data, it is good practice to run **run_extractor_instance** in **EXTRACT_ONLY** mode. The result of this will be that the extracted data will be placed in your extracted folder for the topic, and then you can use the **REPLAY_RAW** mode to process this data. This method ensures all data is available for processing by **run_extractor_instance**.
- **Replay Raw:** This form of extraction takes previously extracted data from the Operations Analytics Predictive Insights folder `$PI_HOME/var/spool/topics/<topic>/extracted`, where `<topic>` is the name of your topic. This extraction type is typically used if you have extracted previously in **EXTRACT_ONLY** mode, or if you have decided to cleanup and reprocess your data.

Procedure

1. Log in as the administrative user, typically `scadmin`.
2. Change to: `$PI_HOME/bin`.
Where `$PI_HOME` is the installation location of the analytics component.
3. Run the **run_extractor_instance** as per one of the defined types. The following sections describe how you can use the **run_extractor_instance** command to extract data as per one of the defined types.

Extracting data in backlog mode

Run the **run_extractor_instance** command in backlog mode in order to consume and analyze backlog data.

About this task

A typical scenario might be that you have three months of backlog data you wish to process using Operations Analytics Predictive Insights.

Procedure

1. Log in as the administrative user, typically `scadmin`.
2. Change to: `$PI_HOME/bin`.
Where `$PI_HOME` is the installation location of the analytics component.
3. Run the **run_extractor_instance** command. The start and end time set for the **run_extractor_instance** command defines the total amount of data that will be processed, that is, the total amount of backlog data you want to process.
`./admin.sh run_extractor_instance -t=Topic1 -s=20130801-0000 -e=20130829-1200`
4. Monitor the progress alarms displayed on the Active Event List to observe the status of data consumption.
You will notice alarms of the type:
 - Training is complete

- Extraction is complete

The **run_extractor_instance** command will automatically suspend and resume the extraction of data to allow Operations Analytics Predictive Insights time to rebuild the analytics model. The analytics model is rebuilt regularly in order to keep the model relevant to the current data.

Note: If at any point you have to restart the extraction process because of some issue, run the command specifying the end time and not the start time.

Supplying the start time the first time you run the command means extractions start from your selected time. The extraction process if stopped will pick up from the last timestamp it processed. Allowing the extraction process pick the start time for all but the first running of the command means you will not miss any data, which may occur if you manually set the start time for every time the **run_extractor_instance** command is run.

5. The anomalies generated based on the analyzed data will not be visible in the AEL as they are historical alarms. To see these anomalies using the Operations Analytics Predictive Insights UI, you must use the **Service Diagnosis Anomaly Search**.

Related reference:

"run_extractor_instance" on page 105
The run_extractor_instance CLI command.

Extracting data in switch mode

Run the **run_extractor_instance** command to extract initially in backlog mode and then switch to steady state mode. If you have available backlog data this is the most efficient data extraction type to use.

About this task

A typical scenario might be that you want to extract initially in backlog mode and then switch to steady state mode, for example, extract 28 days of backlog data and then continue to process incoming data.

Procedure

1. Log in as the administrative user, typically scadmin,.
2. Change to: \$PI_HOME/bin.
Where \$PI_HOME is the installation location of the analytics component.
3. Run the **run_extractor_instance** command. The start and end time set for the **run_extractor_instance** command defines the total amount of data that will be processed, that is, the total amount of backlog data you want to process.

```
./admin.sh run_extractor_instance -t=Topic1 -s=20130814-0000 -l=20
```

The start time specified, 20130814-0000 in this example, is the start time of the backlog data. If, for example, you wish to extract 28 days of backlog data, set the start time to be 28 days in the past. If you don't specify an end time, extraction continues indefinitely.

The -l parameter specifies the latency, in minutes, that will be applied when extraction switches to steady state mode. Set this value based on the number of minutes delay in the arrival of data to your datasource. The default latency is set to the same value as the **system.aggregation.interval** property.
4. Monitor the progress alarms displayed on the Active Event List to observe the status of data consumption.

You will notice alarms of the type:

- Training is complete
- Extraction is complete

The **run_extractor_instance** command may suspend and resume the extraction of data to allow Operations Analytics Predictive Insights time to rebuild the analytics model. The analytics model is rebuilt regularly in order to keep the model relevant to the current data.

5. The anomalies generated based on the analyzed data will be visible in the AEL as they are based on current data.

Related reference:

“run_extractor_instance” on page 105

The run_extractor_instance CLI command.

Extracting data in steady-state mode

Run the **run_extractor_instance** command to extract current incoming data only.

About this task

The steady state mode is used when you want to process current incoming data only

Procedure

1. Log in as the administrative user, typically scadmin,.
2. Change to: \$PI_HOME/bin.
Where \$PI_HOME is the installation location of the analytics component.
3. Run the **run_extractor_instance** command.

```
./admin.sh run_extractor_instance -t=Topic1 -l=20
```

The -l parameter specifies the latency, in minutes, that is applied when extraction switches to steady state mode. Set the latency based on the number of minutes delay in the arrival of data to your datasource. The default latency is set to be the same value as the **system.aggregation.interval** property.
4. The anomalies generated based on the analyzed data will be visible in the AEL as they are based on current data.

Related reference:

“run_extractor_instance” on page 105

The run_extractor_instance CLI command.

Extracting data in replay raw mode

Run the **run_extractor_instance** command to replay existing and previously extracted data.

Before you begin

To extract data in **REPLAY_RAW**, you must have run the extractor on a previous occasion.

About this task

Extracting data in replay raw mode is of use if you have had to reconfigure or update your data model, and you want to avoid waiting for the system to process new data.

The following is an example of how this can be done:

Procedure

1. Log in as the administrative user, typically scadmin.
2. Change to: `$PI_HOME/bin`.
Where `$PI_HOME` is the installation location of the analytics component.
3. Run the **stop** command.
`./stop.sh`
4. Make a copy of the files. All previously extracted data for a given topic is stored in the Operations Analytics Predictive Insights folder `$PI_HOME/var/spool/topics/<topic>/extracted`, where `<topic>` is the name of your topic. The files previously extracted may be split between the good and bad subfolders.
5. Run the **cleanup** command.
`./admin.sh cleanup`
6. Deploy your updated data model.
7. Run the **start** command.
`./start.sh`
8. Put the files you copied into the `$PI_HOME/var/spool/topics/<topic>/extracted` directory, where `<topic>` is the name of your topic.
9. Run the **run_extractor_instance** command.
`./admin.sh run_extractor_instance -t=Topic1 -m=REPLAY_RAW -s=20130814-0000`
The start time specified, 20130814-0000 in this example, is the start time of your backlog data (REPLAY_RAW can be considered backlog data). After all the backlog data is consumed, extraction stops. To resume extraction, you must first restart Operations Analytics Predictive Insights. Then run the extractor in steady state mode. For more information, see “Extracting data in steady-state mode” on page 73.

Example

```
./admin.sh run_extractor_instance -m=REPLAY_RAW -s=20140501-0000 -e=20140601-0000
```

The command replays the data between these dates. If data has already been run-through it will be discarded.

```
./admin.sh run_extractor_instance -m=REPLAY_RAW -s=20140501-0000
```

Replay all the data on the OS starting from 1st of May, and ending when all files have been read. Extractor will stop. No more data will be read through. It will not transition to steady state.

Restart Operations Analytics Predictive Insights:

```
cd $PI_HOME/bin
./stop.sh -t=<topicName>
./start.sh -t=<topicName>
```

Resume extraction in steady state mode:

```
./admin.sh run_extractor_instance -m=EXTRACT -s=20140601-0000
```

Related reference:

“run_extractor_instance” on page 105

The run_extractor_instance CLI command.

Extracting data after Operations Analytics Predictive Insights restarts

How to run the **run_extractor_instance** command if your Operations Analytics Predictive Insights system is restarted.

Procedure

1. Log in as the administrative user, typically scadmin.
2. Go to the \$PI_HOME/bin directory.
3. Run the **run_extractor_instance** command in one of the following ways:

- **Restart:** If the extractor stops for some reason, run the following commands:

```
./start.sh  
./admin.sh run_extractor_instance -l=20
```

When you don't specify a start time parameter with the **run_extractor_instance** command, the extractor resumes from the last extracted time.

- **System down time:** If Operations Analytics Predictive Insights experiences down time, a significant backlog of data might build up. If a backlog spans one or more days, you must run the extractor in backlog mode. For more information, see “Extracting data in backlog mode” on page 71.

Related reference:

“run_extractor_instance” on page 105
The run_extractor_instance CLI command.

Checking status

How to check the status of training and of data flow to Operations Analytics Predictive Insights.

About this task

Before Operations Analytics Predictive Insights can begin identifying anomalies it must first create an analytics model. Operations Analytics Predictive Insights provides a set of progress alarms which give you status on the progress of your training, plus the data flow to your Operations Analytics Predictive Insights system.

To check the status of your training and data flow:

Procedure

1. Log in to the same Tivoli Integrated Portal or Dashboard Application Services Hub instance as your Operations Analytics Predictive Insights user interface.
2. Select the **Active Event List**.

The **Active Event List** displays the list of detected anomalies and progress indicator alarms. Progress indicator alarms are raised with a severity of Low.

The following are the set of possible status and progress indicator alarms:

Topic status message

The availability of your data through the created topic or topics is essential before data flow can begin for Operations Analytics Predictive Insights. The display of the following message indicates that data is available for extraction:

- Topic Data Source started for topic: <topic_name>, ready to start data loading

Extraction status messages

Extractors pull the data from your data source allowing Operations Analytics Predictive Insights to first train and then analyze your data. Extraction is done on a per topic basis. The following messages indicate whether extraction is in progress for a given topic:

- Extractor started for topic: <topic_name> and datasource: <datasource_name>
- Extractor stopped for topic: <topic_name> and datasource: <datasource_name>

Training progress messages

Training is the first stage of data analysis. A training or analytics model must be constructed before Operations Analytics Predictive Insights is able to identify anomalies. After the initial training period, data analysis can begin. Operations Analytics Predictive Insights will retrain by default once a day to keep the analytics model current. The delay between each training can be increased for individual algorithms using the **<algorithm name>.retrainingIntervalMinutes** configuration property. The following messages indicate whether training is in progress and its level of completeness:

- Started receiving new data for training to begin.
- Received 25% of necessary data for training to begin.
- Received 50% of necessary data for training to begin.
- Received 75% of necessary data for training to begin.
- Received 100% of necessary data for training to begin.
- Finished Training. New model produced.
- New model training started with data between: <start_date_time> and <end_date_time> baselines and relationships will be updated on completion.

Chapter 7. Operations Analytics Predictive Insights administration

The administration tasks required to manage Operations Analytics Predictive Insights components.

All tasks that you must carry out to administer the Operations Analytics Predictive Insights system are described in the following sections.

Filtering alarms

You can filter the set of alarms output by using the `filtered_alarms.txt` file.

About this task

You can filter the set of alarms that are generated by the Analytics server for any resource, metric group and metric.

Procedure

1. Open the `$PI_HOME/spl/instances/Analytics<topic name>/config/filtered_alarms.txt` file.

Where `<topic name>` is the name of the topic to which you want to apply the filtering .

2. Use a wildcard or regular expression to specify the resource, metric group, or metric and whether the condition means that the alarm is discarded or forwarded.

The format of a condition must be as follows:

```
regex|wild,<resource name>,<metric group name>,<metric name>,<rule type>,<threshold>,<forward|discard|major|minor|warning|critical
```

Where:

- `regex` means use regular expressions.
Regular expression example: `[Dd]isk_space_.*`
- `wild` means use basic wildcard matching - supports '*' and '?' only.
Wildcard example: `*_CPU*`
- `resource name` specifies the name of the resource to filter
- `metric group name` specifies the name of the metric group to filter
- `metric` specifies the name of the metric to filter
- `rule type` sets the type of rule to include in the filter condition and can be one of the following:
 - `actual_expected` checks if both the actual and expected values are less than or equal to the threshold
 - `actual_only` checks only if the actual value is less than or equal to the threshold
 - `expected_only` checks only if the expected value is less than or equal to the threshold
 - `delta` checks if the difference between the actual and expected value is less than or equal to threshold
 - `higher` checks if the actual value is higher than the expected value

- lower checks if the actual value is lower than the expected value
 - threshold specifies the numerical value against which the rule type is compared
 - forward means forward all matching alarms.
 - discard means discard all matching alarms.
 - major means set the severity of all matching alarms to major
 - minor means set the severity of all matching alarms to minor
 - warning means set the severity of all matching alarms to warning
 - critical means set the severity of all matching alarms to critical
 - * or an empty string ensures that the test always passes for regex or wild.
- The first condition that is matched is used.

Note: The `filtered_alarms.txt` file is read dynamically at every interval so you do not need to restart Operations Analytics Predictive Insights after you update the file.

Example

The following are examples of filters:

- `wild,*,*,*,forward`
Forward all alarms -> no other rule will be read after.
- `wild,,,forward`
Forward all alarms -> no other rule will be read after. (same as the previous example - null equals all pass)
- `wild,*,NTPROCSSRGroup,*,forward`
Forward all metrics in the metric group `NTPROCSSRGroup`.
- `wild,*,NTPROCSSRGroup,%,*,forward`
Forward all metrics that begin with the percent symbol (%) AND are in the metric group `NTPROCSSRGroup`.
- `regex,*,NTPROCSSRGroup,%,*,forward`
Same as the previous example, but uses regular expressions instead of wildcards.
- `wild,*,NTPROCSSR*,*,discard`
Discard all alarms that are emitted from metrics in the metric group `NTPROCSSRGroup`.
- `wild,*,*,*Space_Available*,forward`
`wild,*,*,*,discard`
Forward alarms with metric names that contain `Space_Available` only.
- `wild,brayz1*,UNIXDISKGroup,*,forward`
`wild,*,*,*,discard`
Forward alarms with resource names that start with `brayz1` and resource group in `UNIXDISKGroup` only
- `wild,brayz1*,UNIXDISKGroup,Inodes_Used,forward`
`wild,*,*,*,discard`
Forward alarms with resource names that start with `brayz1` and resource group in `UNIXDISKGroup` and metric name `Inodes_Used` only
- `wild,router-sw49.tut.com,ResptimeGroup,Maxresponsetime,delta,1000,discard`
`wild,router-sw49.tut.com,ResptimeGroup,Maxresponsetime,delta,2000,minor`
`wild,router-sw49.tut.com,ResptimeGroup,Maxresponsetime,delta,3000,major`
`wild,router-sw49.tut.com,ResptimeGroup,Maxresponsetime,critical`
Gradually increases the severity of the alarm as the deviation between the actual and expected values increases. If the delta of the actual and expected value is

less than or equal to 1000 the alarm is discarded . If this condition is not met, Operations Analytics Predictive Insights checks the remaining lines and sets the severity to the appropriate value depending on which condition is met. For example, if the delta of the actual and expected values is greater than 2000 but less than 3000, the condition in the third line of the example is met and the alarm severity is set to major.

- `wild,router-nw57.tut.com,ResptimeGroup,Maxresponsetime,actual_only,50,minor`
Set the severity of the alarm to minor if the actual value of the Maxresponsetime metric is less than or equal to 50.
- `wild,router-nw57.tut.com,ResptimeGroup,Maxresponsetime,actual_expected,100,minor`
Set the severity of the alarm to minor if both the actual value and expected value of the Maxresponsetime metric is less than or equal to 100.
- `wild,*,*,*Totalbytes,expected_only,10000000,warning`
If a metric name that ends with Totalbytes has a value less than or equal to 1,000,0000, set the severity of the alarm to warning.

Creating a topic

You must deploy each model that you create to a topic. A default topic is created as part of the installation of the Analytics component. You can create additional topics to segregate data analysis and presentation.

About this task

Topics allow you to segment your data in ways that are relevant to your business. For example, you can create a topic and use it to group data that originates from a specific service, application, or geography. The anomalies generated based on your data can then be grouped and displayed by topic. Data in different topics are analyzed independently of each other and no analysis takes place across topics. When you deploy a new model, you are asked to choose the topic to which the new model belongs.

Note: For file based data sources, do not set up multiple topics that point to the same source location.

Procedure

1. Log on to the Operations Analytics Predictive Insights Analytics server as the administrative user, typically `scadmin`, and navigate to `$PI_HOME/bin`.

Note: The `$PI_HOME` environment variable is set automatically when the user logs in. However, if your user's shell has been running since before the installation, you might need to exit and log in again.

2. Run the **create_topic** command for each topic you wish to create.

The **create_topic** command is one of the available commands within the **admin** CLI application:

```
./admin.sh create_topic <topic name> <description>
```

Where:

- **<topic name>**: Is the name of the new topic. The topic name must be one word between 3 and 10 characters long. It can contain alphanumeric characters and the underscore character
- **<description>**: The topic description. The description should be enclosed in double quotes if it contains spaces.

For example:

```
./admin.sh create_topic network "Topic for network data"
```

What to do next

If the topic is no longer being used it can be deleted using the **delete_topic** command.

Related reference:

“create_topic” on page 102

The create_topic CLI command.

“delete_topic” on page 102

The delete_topic CLI command.

Starting Operations Analytics Predictive Insights components

How to restart your Operations Analytics Predictive Insights components and system after an error or a system update.

About this task

You might want to start or stop your Operations Analytics Predictive Insights system for maintenance.

If you are starting the Operations Analytics Predictive Insights system as a whole, you must start components in the following order:

- DB2®
- OMNIBus
- Operations Analytics Predictive Insights Analytics component
- Dashboard Application Services Hub and the Operations Analytics Predictive Insights User Interface component or Tivoli Integrated Portal

Starting DB2

You must start the DB2 database to start the Operations Analytics Predictive Insights database component.

Procedure

1. As the DB2 owner, for example, db2inst1, run the db2start script.
2. As the DB2 administrator user, for example, dasusr1, run the db2admin start command.

Starting the OMNIBus ObjectServer

When the OMNIBus initial installation is complete and before you install OMNIBus WebGUI, you must start the ObjectServer.

Procedure

1. Change to the /opt/IBM/tivoli/netcool/omnibus/bin directory.
2. To start the ObjectServer, enter the following command:
nohup ./nco_objserv &
3. To confirm that the ObjectServer is running, enter the following command:
./nco_sql -user root
When prompted for a password, enter the OMNIBus root user password or press enter if the password is blank.

- A successful login confirms that the ObjectServer is running.
4. To exit, type quit.

Starting the Operations Analytics Predictive Insights analytics component

Use the start.sh script to start the Operations Analytics Predictive Insights Analytics component.

Procedure

1. Change to the \$PI_HOME/bin directory.
2. As the administrative user, typically scadmin, enter the following command to start the Analytics component:

```
./start.sh [-t=<topic_name>]
```

The command starts the Operations Analytics Predictive Insights processing framework and any topic you specify. If you omit the **-t** parameter, the command starts all topics on the server.
3. To start extracting data from the data source, enter the following command:

```
$PI_HOME/bin/admin.sh run_extractor_instance -t=<topic> -l=<latency>
```

For more information, see run_extractor_instance.

start.sh

The start.sh script is used to start Operations Analytics Predictive Insights.

Purpose

This script has the purpose of starting Operations Analytics Predictive Insights.

```
start.sh [-t=<topic_name>]
```

The script starts the loading of data by Operations Analytics Predictive Insights and starts the analytics instance.

You can use the **-t=<topic_name>** option to specify one topic, so processing of data is started for that one topic.

Starting Tivoli Integrated Portal

If you want to use Tivoli Integrated Portal to display anomalies in the Operations Analytics Predictive Insights user interface, start Tivoli Integrated Portal.

Procedure

1. As the user that installed Tivoli Integrated Portal, change to the Tivoli Integrated Portal directory. For example: /opt/IBM/tivoli/tipv2/profiles/TIPProfile/bin/

Note: Start the Tivoli Integrated Portal as root user only if the IBM Tivoli Network Manager V3.8 is installed on your system.

2. Run the startServer.sh script, passing the script the name of your Tivoli Integrated Portal server: For example,

```
./startServer.sh server1
```

Starting Dashboard Application Services Hub

If you want to use Dashboard Application Services Hub to display anomalies in the Operations Analytics Predictive Insights user interface, start Dashboard Application Services Hub and the Operations Analytics Predictive Insights UI.

Procedure

1. As the user that installed Dashboard Application Services Hub, change to the Dashboard Application Services Hub directory. For example,
`/opt/IBM/JazzSM/profile/bin`
2. Run the `startServer.sh` script, passing the script the name of your Dashboard Application Services Hub server: For example,
`./startServer.sh server1`
3. Log onto the server on which you installed the Operations Analytics Predictive Insights UI as the user that installed it.
4. Navigate to `<Liberty_Install_Home>/UI/bin`

Note: `<Liberty_Install_Home>` defaults to `/opt/IBM/scanalytics`

5. Run the following command
`./pi.sh -start`

Checking that all Operations Analytics Predictive Insights components started correctly

How to confirm that all components are running.

Procedure

1. To confirm that DB2 is running, enter the following command:

```
ps -ef|grep -i db2
```

The command returns output similar to the following:

```
db2fenc1 4057 6594 0 Aug09 ?      00:00:00 db2fmp ( ,0,0,0,0,0,0,0,1,0,8a65b0,14,200014,2,0,1,109ffc0,
0x210000000,0x210000000,1600000,350007,2,77218058
root      6106 1 0 May23 ?      01:06:43 /opt/IBM/db2/V9/bin/db2fmc
dasusr1 6272 1 0 May23 ?      00:00:03 /opt/IBM/dasusr1/das/adm/db2dasrrm
root      6594 1 0 May23 ?      00:00:00 db2wdog 0
db2inst1 6596 6594 0 May23 ?      01:55:51 db2sysc 0
root      6597 6596 0 May23 ?      00:00:01 db2ckpwd 0
root      6598 6596 0 May23 ?      00:00:01 db2ckpwd 0
root      6599 6596 0 May23 ?      00:00:01 db2ckpwd 0
db2inst1 6640 6594 0 May23 ?      00:24:45 db2acd 0 ,0,0,0,1,0,0,0,1,0,8a65b0,14,200014,2,0,1,1ffc0,
0x210000000,0x210000000,1600000,350007,2,a000f
dasusr1 6656 1 0 May23 ?      00:00:02 /opt/IBM/db2/V9/das/bin/db2fmd -i dasusr1 -m
/opt/IBM/db2/V9/das/lib/libdb2dasgcf.so.1
db2fenc1 7134 6594 0 May23 ?      00:01:13 db2fmp ( ,1,0,0,0,0,0,0,1,0,8a65b0,14,200014,2,0,1,22ffc0,
0x210000000,0x210000000,1600000,350007,2,2a002a
dasusr1 7940 1 0 May23 ?      00:00:03 /opt/IBM/db2/V9/das/bin/db2fmd -i dasusr1 -m
/opt/IBM/db2/V9/das/lib/libdb2dasgcf.so.1
db2fenc1 14681 6594 0 May23 ?      00:00:13 db2fmp ( ,0,0,0,0,0,0,0,1,0,8a65b0,14,200014,2,0,1,85ffc0,
0x210000000,0x210000000,1600000,350007,2,88802d
```

2. To confirm that Analytics is running, enter the following command:

```
streamtool lspe -i spl
```

The command returns output similar to the following for each operator:

Id	State	RC	Healthy	Host	ID	JobId	JobName	Operators
0	Running	-	yes	eolas	28679	0	AnalyticsALL	InputDataStream

If the output for any operator does not show yes under the healthy column, investigate the problem.

If you identify the cause of the problem, run `$PI_HOME/bin/stop.sh`, make the change needed to resolve the problem, and run `$PI_HOME/bin/start.sh`.

If you are unable to resolve the problem, use the `collect.sh` script to collect information to send to IBM. For more information, see `collect.sh`.

3. To confirm that the OMNIBus ObjectServer is running, enter the following command:

```
ps -ef | grep nco_objserv
```

If the ObjectServer is running, you see output similar to the following:

```
root 16514 1 0 Feb02 ? 01:19:26 /opt/IBM/tivoli/netcool/omnibus/platform/linux2x86/bin64/nco
```

4. Complete the steps for Tivoli Integrated Portal or Dashboard Application Services Hub as appropriate:

Tivoli Integrated Portal

- a. Enter the command:

```
ps -ef|grep -i tip
```

The returned text looks something like:

```
root      22502      1  0 Aug07 pts/10  00:12:09 /opt/IBM/tivoli/tipv2/java/bin/java
<...> TIPCell TIPNode server1
```

- b. Log in to the Tivoli Integrated Portal console. For example, by navigating to <https://hostname.ibm.com:<port>/ibm/console/logon.jsp>, and confirm the login is successful.

Dashboard Application Services Hub

- a. Enter the following command:

```
ps -ef|grep -i jazzsm
```

The output displayed is similar to the following:

```
root      23237      1  0 Jul01 ?        01:04:04 /opt/IBM/WebSphere/AppServer/java/bin/java
<...> JazzSMNode01Cell JazzSMNode01 server1
```

- b. Log in to the Dashboard Application Services Hub console. For example, go to <https://hostname.ibm.com:<port>/ibm/console/logon.jsp>, and confirm that the login is successful.
- c. To confirm that the Operations Analytics Predictive Insights User Interface component is running, do the following:
 - 1) Log in to the server on which you installed the Operations Analytics Predictive Insights UI component.
 - 2) Change to the <PI_UI_Home>/UI/bin directory. The default <PI_UI_Home> path is /opt/IBM/scanalytics.
 - 3) Enter the following command: `./pi.sh -status`

If the User Interface component is running, you see output similar to the following:

No.	Service	Status	Process ID
1	IBM Websphere Liberty Profile	UP	12046

Stopping Operations Analytics Predictive Insights components

You might want to stop the Operations Analytics Predictive Insights components and system to carry out some system maintenance.

About this task

If you are stopping the Operations Analytics Predictive Insights system as a whole, you must stop components in the following order:

- Tivoli Integrated Portal or Dashboard Application Services Hub
- Streams
- DB2

Stopping DB2

You must stop the DB2 database to stop the Operations Analytics Predictive Insights database component.

About this task

Procedure

1. As the DB2 owner, for example, db2inst1, run the db2stop script.
2. As the DB2 administrator user, for example, dasusr1, run the db2admin stop command.

Stopping Tivoli Integrated Portal

If you are running Tivoli Integrated Portal, you must stop this application to stop the Operations Analytics Predictive Insights user interface.

Procedure

1. As the user that installed Tivoli Integrated Portal, change to the Tivoli Integrated Portal directory.

For example: /opt/IBM/tivoli/tipv2/profiles/TIPProfile/bin/.

Note: Stop the Tivoli Integrated Portal as root user only if the IBM Tivoli Network Manager V3.8 is installed on your system.

2. Run the stopServer.sh script, passing the script the name of your Tivoli Integrated Portal server and your Tivoli Integrated Portal login and password:
./stopServer.sh <tip_server> -username <username> -password <password>

For example:

./stopServer.sh server1 -username tipadmin -password tipadmin

Stopping Dashboard Application Services Hub

If you are running Dashboard Application Services Hub, you must stop this application to stop the Operations Analytics Predictive Insights user interface.

Procedure

1. Log onto the server on which you installed the Operations Analytics Predictive Insights UI as the user that installed it.
2. Navigate to <Liberty_Install_Home>/UI/bin

Note: <Liberty_Install_Home> defaults to /opt/IBM/scanalytics

3. Run the following command

./pi.sh -stop

4. As the user that installed Dashboard Application Services Hub, change to the Dashboard Application Services Hub directory.

For example, /opt/IBM/JazzSM/profile/bin.

5. Run the stopServer.sh script, passing the script the name of your server and your Dashboard Application Services Hub login and password:

./stopServer.sh <dash_server> -username <username> -password <password>

For example:

./stopServer.sh server1 -username smadmin -password smadmin

Stopping the Operations Analytics Predictive Insights Analytics component

Use the stop.sh script to stop the Operations Analytics Predictive Insights Analytics component.

Procedure

1. Change to the \$PI_HOME/bin directory.
2. As the administrative user, typically scadmin, enter the following command to stop the Analytics component:

```
./stop.sh -s
```

The **-s** option stops InfoSphere® Streams. Stopping InfoSphere Streams stops all analytics instances.

stop.sh

The stop.sh script is used to stop Operations Analytics Predictive Insights.

Purpose

This script has the purpose of stopping Operations Analytics Predictive Insights.

```
stop.sh -s [-t=<topic_name>]
```

If you don't specify a parameter, the script stops data processing for all topics but does not stop InfoSphere Streams.

You can use the **-s** parameter to stop Infosphere Streams. If you stop Infosphere Streams, data processing is stopped for all topics.

You can use the **-t=<topic_name>** option to stop an individual topic.

Backing up and restoring data

To prevent the loss of information in the event of a disaster, and for disaster recovery, back up your Operations Analytics Predictive Insights configuration and data from the following locations:

- DB2 database
- \$PI_HOME/states directory
- \$PI_HOME/var/spool/topics/*/extracted directories. Create a rolling backup of, at a minimum, the last 2 weeks of data.

If you need to perform a disaster recovery, restore the DB2 database to recover configuration information including the mediation models. Then, restore the \$PI_HOME/states directory to restore the models created during training. After you reapply any customizations you made to, for example, the OMNIBus probe rules file, the filtered_alarms.txt file, or log settings, your Operations Analytics Predictive Insights system is then ready to resume analyzing data and generating alerts.

Note: The rolling backup of the \$PI_HOME/var/spool/topics/*/extracted directory provides the data to quickly train a new installation of Operations Analytics Predictive Insights if the database recovery and/or a states recovery is unsuccessful.

Changing passwords

If you change the passwords for the users of systems that are integrated with Operations Analytics Predictive Insights, such as, InfoSphere Streams and DB2, you must also update these passwords for Operations Analytics Predictive Insights.

Changing a Tivoli Integrated Portal user password

You change the Tivoli Integrated Portal user password by updating the password in Tivoli Integrated Portal.

About this task

The instructions for changing a Tivoli Integrated Portal user's password are contained in the Tivoli Integrated Portal information center.

Note: Operations Analytics Predictive Insights does not store a copy of this password, so you do not have to make any Operations Analytics Predictive Insights specific changes.


Procedure

1. Change to the Tivoli Integrated Portal Knowledge Center section, *Changing Tivoli Integrated Portal passwords*.
2. Follow the instructions on how to change your own password with the **Change Your Password** portlet.

Changing a Dashboard Application Services Hub user password

You can change the Dashboard Application Services Hub user password by updating the password in Dashboard Application Services Hub.

Procedure

1. Log in to Dashboard Application Services Hub as the user for whom you want to change the password.
2. Click the user icon  in the menu bar.
3. Click **Change Password**.
4. Enter and confirm the new password and click **Set Password**.

Note: Operations Analytics Predictive Insights does not store a copy of this password so you do not have to update the password in Operations Analytics Predictive Insights.

Changing a Tivoli Integrated Portal administrator password

You change the Tivoli Integrated Portal administrator password by updating the password in Tivoli Integrated Portal.

About this task

The instructions for changing a Tivoli Integrated Portal administrator's password are contained in the Tivoli Integrated Portal Knowledge Center.

Note: Operations Analytics Predictive Insights does not store a copy of this password, so you do not have to make any Operations Analytics Predictive Insights specific changes.

Procedure

1. Change to the Tivoli Integrated Portal Knowledge Center section, *Changing Tivoli Integrated Portal passwords*.
2. Follow the instructions on how to change a password with the **Manage Users** portlet.

Changing the Tivoli Netcool/OMNibus ObjectServer user password

If you change the ObjectServer user password in Tivoli Netcool/OMNibus, you must also update the password that is stored in Operations Analytics Predictive Insights and in Dashboard Application Services Hub, if you are using Dashboard Application Services Hub.

Before you begin

For information on how to change the Tivoli Netcool/OMNibus ObjectServer user password in Tivoli Netcool/OMNibus, open the Tivoli Netcool/OMNibus Knowledge Center, select the version of OMNibus that you are using, and search for *Changing the password for the connection to the ObjectServer*.

About this task

Procedure

1. To set the new ObjectServer user password in Operations Analytics Predictive Insights, as the administrative user, typically scadmin, enter the following command on the analytics server:

```
$PI_HOME/bin/admin.sh passwd $PI_HOME/config/omnibus.connection.properties connection.password
```

When prompted to enter the new connection password, enter the new password set for the ObjectServer user.

2. To set the new ObjectServer user password in Dashboard Application Services Hub, as the administrative user, typically scadmin, enter the following command on the UI server:


```
/opt/IBM/scanalytics/UI/bin/changeOmnibus.sh --password=<password>
```

where <password> is the new password set for the ObjectServer.

Changing a Dashboard Application Services Hub administrator password

You can change the Dashboard Application Services Hub administrator password by updating the password in Dashboard Application Services Hub.

Procedure

1. Log in to Dashboard Application Services Hub as administrator.
2. Click the user icon  in the menu bar.
3. Click **Change Password**.
4. Enter and confirm the new password and click **Set Password**.

Note: Operations Analytics Predictive Insights does not store a copy of this password so you do not have to update the password in Operations Analytics Predictive Insights.

Changing the database user password

If you change the password of the UNIX account that is configured as the database user, default scadmin, you must also update the password in Operations Analytics Predictive Insights. The database user password is encrypted in the `$PI_HOME/config/tasp.connection.properties` file. You must also update the password in your visualization application, Dashboard Application Services Hub or Tivoli Integrated Portal.

About this task

To change the DB2 database user password:

Procedure

1. On the Analytics server, log in as the administrative user, typically scadmin,.
2. Change to the `$PI_HOME/bin` directory.
3. Enter the following command syntax:
`./admin.sh passwd $PI_HOME/config/tasp.connection.properties connection.password`

The following message is displayed

Enter connection.user password:

4. Enter the new password that was set for the database user account.
5. Enter the following commands to restart Operations Analytics Predictive Insights.

```
cd $PI_HOME/bin
./stop.sh
./start.sh
```

6. To set the new password in Dashboard Application Services Hub, on the UI server, enter the following command:

```
/opt/IBM/scanalytics/UI/bin/changeDatabase.sh --password=[value] [options]
```

Where options are:

- `--databaseName=<value>` is the name of the Operations Analytics Predictive Insights database.
- `--serverName=<value>` is the server where the Operations Analytics Predictive Insights database is located.
- `--portNumber=<value>` is the port number that is used by the Operations Analytics Predictive Insights database.
- `--user=<value>` is the user name of the Operations Analytics Predictive Insights database user.
- `--show` shows the configuration details for the Operations Analytics Predictive Insights database.

7. To set the new password in Tivoli Integrated Portal, on the UI server, enter the following command:

```
/opt/IBM/tivoli/tipv2/products/tasp/tasp/bin/configure.sh -tipuser <tipuser> -tippassword <tippassword> -driverhome <driverhome> -dburl <jdbcurl> -dbuser <jdbcuser> -dbpassword <jdbcpassword>
```

Where:

- `<tipuser>` is the user name used to connect to Tivoli Integrated Portal.

- <tippassword> is the password used to connect to Tivoli Integrated Portal.
- <driverhome> is where the JDBC driver can be found.
- <dburl> is the JDBC URL for the Operations Analytics Predictive Insights data source.
- <dbuser> is the user name for connecting to the Operations Analytics Predictive Insights data source.
- <dbpassword> is the password for connecting to the Operations Analytics Predictive Insights data source.

Note: After you set the password, you must restart Tivoli Integrated Portal so it uses the updated password. The password change fails if the input values are invalid.

Example:

```
./configure.sh -tipuser tipadmin -tippassword tipadmin007 -type jdbc
-driverhome /opt/IBM/tivoli/tipv2/products/tasp/tasp/lib
-dbur1 jdbc:db2://hostname.ibm.com:50000/TASPDB8
-dbuser dbuser -dbpassword userdb007
```

Changing the DB2 instance owner password

You change the DB2 database instance owner password by updating the password in DB2.

About this task

The instructions on how to change the DB2 instance owner's password are contained in the DB2 information center.

Note: Operations Analytics Predictive Insights does not store a copy of this password, so you do not have to make any Operations Analytics Predictive Insights specific changes.

To change the DB2 database instance owner password:

Procedure

1. Go to the DB2 Knowledge Center and open the section *Maintaining passwords on servers*.
2. Follow the instructions on how to change passwords.

Changing the InfoSphere Streams administrator password

You change the InfoSphere Streams administrator password by updating the password in InfoSphere Streams, and you must also update the password that is stored in the Operations Analytics Predictive Insights system.

About this task

For the default admin user that is created upon installation is scadmin.

To change the InfoSphere Streams admin user password:

Procedure

1. Change to the IBM InfoSphere Streams Knowledge Center and update the Streams administrator password as described for your security service in the section *Setting up user authentication for Streams*.

2. If you enable file extraction, for example, you have a remote system that FTPs CSV files to your Operations Analytics Predictive Insights system, then you must update the password for the FTP on the remote system.

Changing the database data source password

You change the Operations Analytics Predictive Insights database data source password by updating the password within the Mediation tool.

About this task

If a database owner changes the password for a database data source, extraction from that data source is blocked until you update the password for the database data source in the Mediation tool.

The symptoms of a change in database password by the database owner are:

- No extraction or merge files being created.
- TopicDataSourceALL_trace_DataSource.log displays an error: "User or Password invalid".

In the Mediation tool, you see the following symptoms:

- Test Connection fails,
- Show Table Content fails,
- Synchronize Schema fails,
- Data Extraction Preview fails and
- Deploy fails with User ID or Password invalid.

You update the password in the Connection Details tab in the Mediation tool. For a failed Deploy, the password must be updated in the "Password required for data source" pop-up dialog.

Procedure

1. Stop Operations Analytics Predictive Insights. For example:

```
cd $PI_HOME/bin
./stop.sh
```
2. Start the Mediation tool.
3. Open the model that you deployed for extraction.
4. Click the **Connection Details** tab.
5. In the **Password** field, update the password to the correct password.
6. Save the model.
7. Within the **Predictive Insights Navigator** pane, select the model.
8. Click **Predictive Insights > Deploy Model**.
9. The **Predictive Insights Model Deployment** dialog opens, in which you enter the Operations Analytics Predictive Insights database connection details and password.
10. Update the data source password in the **Password required for data source** window.
11. Click **OK**.
12. Start Operations Analytics Predictive Insights. For example:

```
cd $PI_HOME/bin
./start.sh
```


13. Continue extraction from the last extracted time.
`./admin.sh run_extractor_instance`

Chapter 8. Event management administration

Administration of the event management system.

Configuring alarm forwarding from the Analytics server to an OMNIBus ObjectServer

The Analytics server forwards alarms to an OMNIBus ObjectServer so that they can be displayed in the Event Viewer.

Before you begin

You use the `configure_omnibus` script to configure the Analytics component to forward alarms to an OMNIBus ObjectServer. The script also adds columns to an ObjectServer to store and display anomaly information in the OMNIBus event list.

About this task

You must complete this task if:

- You did not set up alarm forwarding from the Analytics component to an OMNIBus ObjectServer when you installed the Analytics component.
- You want to change the OMNIBus ObjectServer to which the Analytics component forwards alarms.
- You have an environment with multiple OMNIBus ObjectServers. In this scenario, you must run the `configure_omnibus` script against each ObjectServer to ensure that the columns that store anomaly information are added to all the ObjectServers.

Note: The Analytics server can forward alarms only to one ObjectServer. If you run the `configure_omnibus` script against multiple ObjectServers, the Analytics server forwards alarms only to the host that you specify for the `-omnihost` parameter the last time you run the script. Therefore, you must run the script and specify the ObjectServer that you want to receive alarms from the Analytics server after you run the script against all the other ObjectServers.

Procedure

1. Change to the `$PI_HOME/bin` directory.

Where `$PI_HOME` is the installation location of the Analytics component.

2. As the administrative user, typically `scadmin`, enter the following command:

```
./configure_omnibus -enable -installdir=<directory> -isuser=<streams user>  
-omniinstance=<instance> -omniuser=<omnibus user> -omnihost=<host> -omniport=<port>  
-omnipwd=<password> -t=<topic>
```

Where:


- `directory` is the installation directory for the Analytics component. For example, `/opt/IBM/scanalytics/analytics`.
- `streams user` is the user name of the InfoSphere Streams user. For example, `scadmin`.
- `instance` is the OMNIBus instance name. For example, `NCOMS`.
- `omnibus user` is the OMNIBus user name. The default user name is `root`.

- host is the OMNIbus ObjectServer host name. For example, hostname.tiv.ibm.com.
- port is the OMNIbus ObjectServer port number. For example, 4100.
- password is the OMNIbus ObjectServer admin password.
- topic is the name of the Operations Analytics Predictive Insights topic. For example, Topic1.

Displaying the Operations Analytics Predictive Insights columns in the Active Event List

Operations Analytics Predictive Insights provides a customized view and filter that displays the Operations Analytics Predictive Insights columns in the Active Event List.

Procedure

1. Log in to Dashboard Application Services Hub or Tivoli Integrated Portal.
2. Open the Active Event List (AEL):
 - If you are using Dashboard Application Services Hub, click the **Incident** icon  and click **Active Event List (AEL)**.
 - If you are using Tivoli Integrated Portal, click **Availability > Events > Active Event List (AEL)**.
3. In both the filter and view drop-down lists, change the value from **Default** to **PredictiveInsights**.

Generating the OMNIbus interfaces file

OMNIbus component communications information is saved in an interfaces file, which is needed by the Tivoli Directory Server communications layer.

Procedure

1. On the OMNIbus server, as the OMNIbus administrator or user, run the **nco_igen** utility to generate the interfaces file:


```
<netcool_home>/bin/nco_igen
```

 where <netcool_home> is typically /opt/IBM/tivoli/netcool
 If OMNIbus is running on a non-Linux architecture, you must pass the **-all** option to the utility. The utility generates an architecture-specific interfaces file for Linux in: <netcool_home>/etc/interfaces directory.
2. As the InfoSphere Streams user, for example scadmin, copy the interfaces file to the \$PI_HOME/probe/etc directory on the Operations Analytics Predictive Insights analytics server.
 Where \$PI_HOME is the installation location of the analytics component.

Customizing the OMNibus probe rules file

The OMNibus probe rules file defines how the OMNibus probe processes event data that it receives from Operations Analytics Predictive Insights. After processing, the probe sends the event data to the OMNibus Object Server as an alert.

The default probe rules file is: `$PI_HOME/probe/omnibus/probes/linux2x86/stdin-tasp.rules`. You can customize the default file for your environment. To ensure that any updates to the probe rules file are read by the system, run a `kill -9` on the probe process after you save the file.

Note: If you want to customize Operations Analytics Predictive Insights to filter alarms generated, as an alternative to updating the probe rules file, you can add filters to the `filtered_alarms.txt` file. For more information, see “Filtering alarms” on page 77.

The following are examples of customizations you can make to the default probe rules file.

Increasing the severity of alarms raised on service-impacting metrics

You can customize the default probe rules file to increase the severity of certain alarms. For example, some metrics that are monitored by Operations Analytics Predictive Insights have a direct user impact. Such metrics include service response time and service availability. To increase the severity of these alarms, update the following section of the probe rules file:

```
# Compute alarm severity depending on whether the potentially anomalous KPI(s)
# contain a service-impacting KPI.
# You will need to update this statement to reflect your TASP implementation.
if (int($TASPCorrelationId) >= 0) {
    if (regmatch($TASPAnomalousMetrics, "Service_Impacting_KPI_1"
        || regmatch($TASPAnomalousMetrics, "Service_Impacting_KPI_2")
    ) {
        @Severity = 4
    } else {
        @Severity = 3
    }
} else if (int($TASPCorrelationId) == -1) {
    # Consolidated alarms - many problems in one alarm.
    @Severity = 4
} else if (int($TASPCorrelationId) == -2) {
    # Data Availability alarms
    @Severity = 5
} else if (int($TASPCorrelationId) == -3) {
    # Information Events should have a lower severity.
    @Severity = 2
} else if (int($TASPCorrelationId) == -4) {
    # System health alarms are very important.
    @Severity = 5
}

# Map the data source severity directly, if it exists.
# Note: resolution events should be set to severity 1, the generic clear will
set them to 0 later
update(@Severity)
```

The `$TASPAnomalousMetrics` element contains the list of metrics that are detected as anomalous by Operations Analytics Predictive Insights. You can use the `$TASPAnomalousMetrics` element to set the severity of the alarm.

You can replace `Service_Impacting_KPI_1` and `Service_Impacting_KPI_2` by the full metric names, as configured by the mediation and visible in the Operations Analytics Predictive Insights User Interface. These include the metric name but not the resource or node name.

You can also set the severity by using the resource names in `$TASPMetricResourceNameList`. You can construct a `$TASPAnomalousMetricResourceNameList` for this purpose. Follow the example for `$TASPAnomalousMetrics` in the default probe rules file.

You must tune your probe rule file according to the configuration of your metrics and groups. For example, if you have metrics with the same name in multiple metric groups, setting the severity on individual metrics can set the severity on a metric of the same name in an incorrect group.

Note: The probe log file for each topic is: `/opt/IBM/scanalytics/analytics/probe/omnibus/log/<topic name>Probe.log`.

Setting a custom alarm summary

By default, Operations Analytics Predictive Insights provides a summary that shows the anomalous metric, its trend and actual and expected values. You can change the details in the probe rules file. For example, you can use the names of the anomalous metrics or resources:

```
# Set a custom summary for alarms impacting a specific service
if (regmatch($TASPAnomalousMetrics, "Service_Impacting_KPI_1")) {
    @Summary = "Anomaly detected on Service 1"
} else {
    @Summary = $Summary
}
update(@Summary)
```

Operations Analytics Predictive Insights can generate alarms with several metrics. Therefore, metrics from two different services can be included in a single alarm. You can decide how you want to handle the scenario, for example, by adding both service names in the summary.

Setting alarm severity based on attributes present in the event

You can set the severity of an alarm based on a particular attribute in the alarm. The following example sets the alarm severity to Critical for any alarm that contains an Infrastructure attribute.

```
if (regmatch($TASPAAttributeNames, "Infrastructure"))
{
    @Severity = 5
}
update(@Severity)
```

Related information:



http://www-01.ibm.com/support/knowledgecenter/SSSHTQ_8.1.0/com.ibm.netcool_OMNIBus.doc_8.1.0/omnibus/wip/probegtwy/task/omn_prb_debuggingrulesfiles.html
Debugging OMNIBus probe rules.

OMNIBus probe tokens

The list of Operations Analytics Predictive Insights probe tokens for OMNIBus.

The following table contains descriptions of the attributes that are provided by Operations Analytics Predictive Insights to the OMNIBus probe. These attributes are available in the rules file.

Table 3. OMNIBus probe list

Element name	Element description	Sample
\$Summary	Contains the summary information about the cause of the alert.	ResponseTime is Higher than expected. Actual: 3000 Expected: 100
\$Identifier	Contains the identifier information of the alert.	id://PI/NETWORK/1_1_3_14145 06900000
\$TASPActualValue	The actual metric value at the time of the anomaly	
\$TASPAgorithmInstanceName	Indicates the algorithm instance name.	AnomalyDetectiongen
\$TASPAgorithmName	Indicates the predictive algorithm name.	com.ibm.pi.analytics.normal _bounds_model
\$TASPAnomalousMetricGroup	The metrics groups that the metrics in this alarm that are tagged as anomalous belong to.	Generatedperf2Group
\$TASPAnomalousMetrics	The metrics in this alarm that are tagged as anomalous.	Metric2
\$TASPAnomalousResources	The resources that the metrics in this alarm that are tagged as anomalous belong to.	ResourceId_6
\$TASPAnomalyTimestamp	Indicates the anomaly time stamp. The time stamp can be set to some time in the future in the case of predictive alarms.	1/13/14 4:00:00 PM
\$TASPCorrelationId	Identifies the correlation, for multivariate alarms.	7522
\$TASPExpectedValue	The expected metric value at the time of the anomaly	
\$TASPIstanceIdentifier	Identifies the Operations Analytics Predictive Insights instance.	PI
\$TASPMetricGroupNameList	Specifies a list of groups names, delimited with 3 semi-colons, for the metrics that were detected as anomalous.	Generatedperf2Group;;; Generatedperf2Group;;; Generatedperf1Group;;; Generatedperf1Group

Table 3. OMNibus probe list (continued)

Element name	Element description	Sample
\$TASPMetricInformation	Provides free text debug information that contains a human-readable summary of the alarm.	Anomaly detected on 4 metrics:ResourceId_6. Metric2 = 11000.0 [0.9669775188394545] ResourceId_8.Metric2 = 174.92455 [0.4090402756931195] ResourceId_1044. Metric0 = 46.99705 [0.415566194969713] ResourceId_1092.Metric0 = 38.38585 [0.181694578856197]
\$TASPMetricNameList	Specifies a list of metric names, delimited with 3 semi-colons, that were detected as anomalous.	Metric2;;;Metric2;;;Metric0 ;;;Metric0
\$TASPMetricNodeList	Specifies a list of node names, delimited with 3 semi-colons, for the metrics that were detected as anomalous. 'Node' is analogous to the 'Node' concept in OMNibus.	ResourceId_6;;;ResourceId_8 ;;;ResourceId_1044;;; ResourceId_1092
\$TASPMetricResourceNameList	Specifies a list of resource names, delimited with 3 semi-colons, for the metrics that were detected as anomalous.	ResourceId_6;;;ResourceId_8 ;;;ResourceId_1044;;; ResourceId_1092
\$TASPMetricResourceTypeList	Specifies a list of resource types, delimited with 3 semi-colons, for the metrics that were detected as anomalous.	gen;;;gen;;;gen;;;gen
\$TASPMetricScoreList	Specifies a list of metric scores, delimited with 3 semi-colons, for the metrics that were detected as anomalous.	0.9669775188394545;0.4090402756931195;;; 0.4175566194969713;0.181694578856197
\$TASPMetricValueList	Specifies a list of metric values, delimited with 3 semi-colons, for the metrics that were detected as anomalous	11000.0;;;174.92455;;; 46.99705;;;38.38585
\$TASPParentIdentifier	Alarms can be consolidated into one alarm. In this case the parent identifier is set to the Identifier of the consolidated alarm. If you use Operations Analytics Predictive Insights OMNibus views, only consolidated alarms are shown, and a right-click menu option lets you display child alarms of the consolidated alarm.	

Table 3. OMNibus probe list (continued)

Element name	Element description	Sample
\$TASPQuantifier	Indicates the quantifier for this alarm.	0.9669775188394545
\$TASPTopic	The name of the topic that includes the anomalous metric.	gen
\$TASPUptime	Indicates the wall-clock time of the last update to this alarm. The wall-clock time could, for example, be used for an analysis of the delay.	1/17/14 1:22:19 PM
\$TASPAActualValue	The value of the metric at the time the anomaly occurs	123
\$TASPDDirection	The direction the metric when compared to its expected value	Lower
\$TASPEExpectedValue	The expected value of the metric at the time the anomaly occurs	52
\$TASPIDentifier	The root identifier for the anomaly	id://PI/NETWORK/1_1_3
\$TASPAAttributeNames	A list of attribute names in the alarm delimited with 3 semicolons	ServiceName;;;Node;;;OSType ;;;dataSourceType
\$TASPAAttributeValues	A list of attribute values in the alarm delimited with 3 semicolons	OnlineBanking;;;router-ne19. tut.com;;;Linux;;;network

Note in particular the two different time stamps provided:

Table 4. Timestamps

Timestamp	Description
\$TASPAAnomalyTimestamp	Timestamp of the predicted anomaly.
\$TASPUptime	Wall-clock time stamp indicating when the alarm was last updated. This may have no relation with the other time stamp, especially if back-processing data.

Chapter 9. Reference

Use this information about Operations Analytics Predictive Insights for reference purposes.

Scripts and utilities

The Operations Analytics Predictive Insights scripts available to the Administrator.

You can run the administration scripts with following steps:

1. Log in as scadmin.
2. Change to the directory: `$PI_HOME/bin`
3. Launch the script; for example, launch the `admin.sh` script by entering the command:
`./admin.sh`

admin.sh

The `admin.sh` script is used manage and maintain the overall system configuration.

To run the `admin.sh` script, do the following:

1. Log in as scadmin.
2. Change to the `$PI_HOME/bin` directory:
3. To run the script, type the following:
`./admin.sh`
4. To display the list of commands that you can use with the `admin.sh` script, type:
`help`

cleanup

The cleanup CLI command.

Purpose

The purpose of this script is to clean up Operations Analytics Predictive Insights so you can replay raw files or re-extract data.

`cleanup [-t=<topic name>]`

Note: Before you can clean up a topic, you must stop all topics on the server where the topic is located. To stop all topics and associated InfoSphere Streams processes on the server, enter the following command `$PI_HOME/bin/stop.sh -s`.

The cleanup command deletes data, metadata, and the configuration model from the database, but retains the schema. It also archives the following three folders to `$PI_HOME/BAK<datestamp>` , for example, `BAK20130808122747`.

- `$PI_HOME/states`
- `$PI_HOME/log`
- `$PI_HOME/var/spool/topics`

The command does not delete alarms from the Objectserver. If you need to delete alarms, log in to Tivoli Integrated Portal or Dashboard Application Services Hub and delete the relevant alarms.

After you clean up a topic, you must redeploy the configuration model from the Mediation Tool and run the following command to start all topics:

```
$PI_HOME/bin/start.sh
```

Parameters

-t= *<topicname>*

Use this option to clean up a specific topic. If you don't specify a topic name, the command performs a cleanup of all topics on the server.

Samples

This sample illustrates how to run the command:

```
./admin.sh cleanup -t=network
```

create_topic

The create_topic CLI command.

Purpose

This command creates a new topic in the system.

```
create_topic <topic name> <description>
```

Parameters

<topic name>

The name for the new topic. The topic name must be one word between 3 and 10 characters long. It can contain alphanumeric characters and the underscore character.

Note: Topic names are created in uppercase, irrespective of the case you use to specify the topic name. For example, if you specify a topic name of "Network", the topic name is created as "NETWORK", which you see when you run the "show_topics" on page 109 command. However, the topic name is not case sensitive when you specify it as part of any other CLI command.

<description>

A description for the new topic. The description must be enclosed in double quotation marks if it contains spaces.

Sample

This sample creates a topic that is named "network" with the description "network data":

```
./admin.sh create_topic network "network data"
```

delete_topic

The delete_topic CLI command.

Purpose

This command deletes a topic.

```
delete_topic <topic name>
```

Parameters

<topic name>

The name of the topic to be deleted. To display a list of topics, use `show_topics`.

Sample

This sample deletes the “network ” topic:

```
./admin.sh delete_topic network
```

deploy_data_model

The `deploy_data_model` command.

Purpose

This command deploys a model to a specified topic. You can use this command if, for some reason, you are unable to deploy a model with the Mediation tool.

```
deploy_data_model -t=<topic name> -m=<model file>
```

Parameters

<topic name>

The name of the topic to deploy the model file to.

<model file>

The name of the model file to deploy.

Sample

This sample deploys the `datasource.pamodel` file to a topic called `network`:

```
./admin.sh deploy_data_model -t=network -m=/workspace/project/datasource.pamodel
```

exit

The `exit` CLI command.

Purpose

The `exit` command closes the administration application. The `exit` command is required only when you run `./admin.sh`.

```
exit
```

Parameters

There are no parameters.

Sample

This sample demonstrates how to exit the administration application:

```
exit
```

help

The `help` CLI command.

Purpose

This command displays help information for the administration application.

```
help
```

Parameters

There are no parameters.

Sample

This sample demonstrates how to display the available help:

```
./admin.sh help
```

history

The history CLI command.

Purpose

This command displays the history of commands that are executed in the current session of the administration application.

The total number of commands that are stored is four.

```
history
```

Parameters

There are no parameters.

Sample

This sample demonstrates how to display the current session history:

Note: To run the history command, you must first launch the admin.sh script and then run the history command from the command line, as illustrated in this sample.

```
./admin.sh
```

```
IBM Operations Analytics Predictive Insights 1.3 - 5725-K26
```

```
(C) Copyright IBM Corp. 2011,2015 All Rights Reserved.
```

```
>> help history
```

```
help: print command history.
```

```
USE: history
```

```
DESCRIPTION:
```

```
Prints the most recently typed user commands
```

```
>>history
```

```
help history
```

```
history
```

passwd

The passwd CLI command.

Purpose

This command is used to update an encrypted password that is saved in a property file. If the file exists, the appropriate property value is replaced. If the file does not exist, a new file is created and appended with the new property.

```
passwd [-c] <filename> <property name>
```

Parameters

-c The file that is specified is new and has to be created.

<filename>

The name of the target file.

<property name>

The name of the target property.

Sample

This sample appends the property, “property”, to a new file called “passwordfile”:

```
passwd -c passwordfile property
```

quit

The quit CLI command.

Purpose

This command quits the administration application.

```
quit
```

Parameters

There are no parameters.

Sample

This sample demonstrates how to exit the administration application:

```
./admin.sh quit
```

run_extractor_instance

The run_extractor_instance CLI command.

Purpose

This command runs an extractor instance. It requires that you deploy a valid configuration model to the Operations Analytics Predictive Insights system with the Operations Analytics Predictive Insights configuration tool.

```
run_extractor_instance -m=<mode> -s=<starttime>  
[ -e=<endtime> -l=<latency> -t=<topicname>]
```

When running the command, remember the timezone differences between the server that you are running the command on and server from which the data is being taken. If the server you are running the command on has a timezone of UTC+5, and the command you entered used `-starttime 20130201-0000` (local); when you convert to UTC, your command runs with `-starttime 20130201-0500` (UTC).

Parameters

-m= *<mode>*

The mode of extraction.

The following are the available modes:

- **EXTRACT** - Extract from data source using the deployed configuration model. If you don't specify a mode, the command uses **EXTRACT** by default.
- **EXTRACT_ONLY** - Only extract to raw files. Do not load into Operations Analytics Predictive Insights. EXTRACT mode does not support the starting and stopping of individual extraction instances. The **APPLY_TO_ALL** option must be used when running in EXTRACT mode.
- **REPLAY_RAW** - Replay previously extracted raw files. This is useful if an updated model is deployed to the Operations Analytics Predictive Insights system. This mode replays previously extracted raw data files and applies the updated model, for example, any different aggregators. On the filesystem, this mode will move files from `$PI_HOME/var/spool/topics/ALL/extracted` to `$PI_HOME/var/spool/topics/ALL/extracted/good` so if you need to REPLAY_RAW again, you'll need to move files back to `$PI_HOME/var/spool/topics/ALL/extracted`.

Note: The `$PI_HOME/bin/admin.sh` cleanup command will archive that data to `$PI_HOME/BAK*` so you may need to retrieve the data from that archive in order to run REPLAY_RAW.

-s= *<starttime>*

The start time of extraction. This is a mandatory parameter the first time the **run_extractor_instance** command is run.

Supplying the start time the first time you run the command means extractions starts from your selected time. The extraction process if stopped will pick up from the last timestamp it processed. Allowing the extraction process pick the start time for all but the first running of the command means you will not miss any data, which may occur if you manually set the start time for every time the **run_extractor_instance** command is run.

The start time is expressed in one of these formats:

- `yyyyMMdd-hhmm` Exact date.
For example, `20130301-10:15`
- `+PxYxMxDtHxMMxS` Time relative to now. x is a positive integer and Y(year) M(month) D(day) H(hour) M(minute) S(second).
For example, `-P2DT4H30M` equates to Two days four hours and 30 minutes ago.

-e= *<endtime>*

The end time of the extraction period. This is an optional parameter. The default is to continue to extract until **stop_extractor_instance** is called.

If backlog data is being loaded and an end time is not specified, extraction will continue until the present time, and then the extractor will enter steady state mode (that is, the reading in of current, incoming data). When the extractor has swapped to steady state mode any further backlog data will be ignored. If the data you want to process is purely backlog data, always specify an end time.

The end time is expressed in one of these formats:

- `yyyyMMdd-hhmm` - Exact date.
- `+PxYxMxDtHxMMxS` - Time relative to now.

-l= *<latency>*

If you specify a latency, the extractor only extracts data that is older than the number of minutes you set for the latency value. The default latency is the value of the **system.aggregation.interval** property.

-t= *<topicname>*

Runs extraction for a specific topic. If multiple topics exist on the server, **-t** is a mandatory parameter.

Extraction

The data extracted from your datasource is consumed in two ways, it is used to train your analytics model, and when an analytics model is available, the data is then analyzed for any anomalies. There are five types of extraction:

- **Steady state mode:** This form of extraction is the simplest and involves Operations Analytics Predictive Insights reading in current, incoming data from your datasource. The main issue with this method is it takes time for your system to train and create an analytics model for each algorithm, which is required before the algorithm can start to identify anomalies. Operations Analytics Predictive Insights creates a model for an individual algorithm when it has consumed fifty percent of the number of days data specified for the **<algorithm name>maxTrainingWindowDays** configuration property.
- **Backlog mode:** This form of extraction involves the reading in of recent historical data. If you are intent on processing only backlog data, you must set the start and end time when calling the command **run_extractor_instance** for the first time. Only supply the stop time and not the start time when making any subsequent calls to the command **run_extractor_instance**.
- **Switch:** This form of extraction is a hybrid of steady state mode and backlog mode, the purpose of which is to create an analytics model as quickly as possible using Backlog mode and move on to Steady state mode to begin the analysis of current data.
- **Extract Only:** If you have multiple data sources with a variety of latencies and you wish to extract backlog data, it is good practice to run **run_extractor_instance** in **EXTRACT_ONLY** mode. The result of this will be that the extracted data will be placed in your extracted folder for the topic, and then you can use the **REPLAY_RAW** mode to process this data. This method ensures all data is available for processing by **run_extractor_instance**.
- **Replay Raw:** This form of extraction takes previously extracted data from the Operations Analytics Predictive Insights folder **\$PI_HOME/var/spool/topics/<topic>/extracted**, where **<topic>** is the name of your topic. This extraction type is typically used if you have extracted previously in **EXTRACT_ONLY** mode, or if you have decided to cleanup and reprocess your data.

show

The show CLI command.

Purpose

This command lists the properties that are configured for a specific topic.

show -t=<topic name>

Parameters

-t= *<topic>*

The topic name. Use “show_topics” on page 109 to display topic names.

Sample

This sample illustrates how to show the properties for the network topic:

```
./admin.sh show -t=network
```

set

The set CLI command.

Purpose

This command sets a new value for a topic property.

```
set -t=<topic name> <property name> <property value>
```

Parameters

-t= *<topic name>*

The topic name.

<property name>

The property name. Use “show” on page 107 to see the list of properties for a topic.

<property value>

The new value for the property specified.

Sample

This sample changes the “system.alarm.autoclear” property to “false” for the “Topic1” topic:

```
./admin.sh set -t=Topic1 system.alarm.autoclear false
```

show_data_model

The show_data_model CLI command.

Purpose

This command creates a file with details of the most recent data model that was deployed to the specified topic.

```
show_data_model <topicname> <output filename>
```

Parameters

<topic name>

The name for the new topic. The topic name must be one word between 3 and 10 characters long. It can contain alphanumeric characters and the underscore character.

Note: Topic names are created in uppercase, irrespective of the case you use to specify the topic name. For example, if you specify a topic name of “Network”, the topic name is created as “NETWORK”, which you see when you run the “show_topics” on page 109 command. However, the topic name is not case sensitive when you specify it as part of any other CLI command.

<output filename>

The name of the file to which you want the command to write the details for the data mode.

Sample

```
./admin.sh show_data_model network /tmp/networktopic.txt
```

show_kpi_counts

The show_kpi_counts CLI command.

Purpose

This command displays the average number of KPIs that Operations Analytics Predictive Insights processed in the last day, week, and month. The average number of KPIs is the total number of KPIs processed in the time period divided by the number of intervals in the time period.

```
show_kpi_counts -t=<topicname>
```

Parameters

-t=<topic name>

The name of the topic for which to display the average KPI count. If you don't specify the **-t** parameter, the command shows the average KPI count for all installed topics.

Sample

```
./admin.sh show_kpi_counts -t=network
```

show_topics

The show_topics CLI command.

Purpose

This command displays installed topics.

```
show_topics [ALL]
```

Parameters

There are no parameters.

Sample

This sample illustrates how to display all installed topics on this Analytics server only:

```
./admin.sh show_topics
```

This sample illustrates how to display all installed topics on every Analytics server in this system:

```
./admin.sh show_topics ALL
```

show_version

The show_version CLI command.

Purpose

This command displays the current version of Operations Analytics Predictive Insights.

```
show_version
```

The following is an example of the result displayed:

1.3.3.1_201509271946

The first 4 digits from the left represent the version number of the product followed by the build date and time. Within the version number, the first 3 digits represents the base product version and the fourth digit identifies the last ifix applied to the base product.

Note: You only see the ifix version if you updated the product version number when you installed the ifix. The instructions on how to update the version number are contained in the installation instructions provided with each ifix.

Sample

This sample illustrates how to display the current version of Operations Analytics Predictive Insights:

```
./admin.sh show_version
```

stop_extractor_instance

The stop_extractor_instance CLI command.

Purpose

This command stops the extraction for a defined period. It requires that you deploy a valid configuration model to the Operations Analytics Predictive Insights system with the Operations Analytics Predictive Insights administration tool.

```
stop_extractor_instance -t=<topicname>
```

Parameters

-t= *<topicname>*

You can use this option to stop extraction for a specific topic. If you don't specify a topic name, the default action is to stop extraction for all topics.

Sample

```
./admin.sh stop_extractor_instance
```

update_memory

The update_memory admin command.

Purpose

This command updates the memory settings and threads for the topic from the values in the specified file.

```
update_memory <topicname> <filename>
```

The memory and thread settings must be calculated based on the number of KPIs being processed by Operations Analytics Predictive Insights and the number of topics you are using. For more information, see *Memory recommendations* under Performance and Sizing on the Operations Analytics Predictive Insights wiki.

Parameters

<filename>

The name of the text file to use as input. The first line of the text file must contain a comma-separated list of headers. The operator header is mandatory and the file can have one or more of the following headers depending on the values you want to update:

- -Xms if you want to update the initial java heap size for an operator.
- -Xmx you want to update the maximum java heap size for an operator.
- Xgcthread if you want to update the number of threads to assign to an operator.

Underneath the headings, each line must have:

- An operator name. Valid operators names are: DataSourceOperator, AnalyticsOperator, SelfMonitorOperator, PostProcessingOperator, OmnibusConfigurator, RossiDecorator, and CorrelationDbWriter.
- The values that you want to update for the operator. If you are updating the initial or maximum java heap size for an operator, after the value, specify m if the value is in megabytes or g if the value is in gigabytes. For example, 6g.

Input file example:

```
operator,-Xmx,-Xms
DataSourceOperator,6g,6g
AnalyticsOperator,11g,11g
RossiDecorator,8g,8g
```

Sample

```
./admin.sh update_memory network /tmp/memorysettings.txt
```

collect.sh

How to use the collect.sh script.

Purpose

Use the collect.sh script to collect information to send to IBM if, for example, you encounter a problem that you cannot resolve.

To run the command:

```
$PI_HOME/collect.sh [-t= <topic name> ] [-o=<output_folder>]
```

For example:

```
$PI_HOME/bin/collect.sh -t=network /tmp/
```

You must then send the output to IBM.

If any other files get copied to this directory, they must be sent along with the collect.log file.

Parameters

<topic name>

The name for the new topic. The topic name must be one word between 3 and 10 characters long. It can contain alphanumeric characters and the underscore character.

Note: Topic names are created in uppercase, irrespective of the case you use to specify the topic name. For example, if you specify a topic name of “Network”, the topic name is created as “NETWORK”, which you see when you run the “show_topics” on page 109 command. However, the topic name is not case sensitive when you specify it as part of any other CLI command.

<directory>

The **<directory>** parameter is optional. This parameter can be used to specify an alternative location to which scripts output is written.

By default the output is written to a subdirectory of the current directory called: `collect_output`.

Samples

This sample illustrates the running of the script, plus how to check the output of the script:

```
cd $PI_HOME/bin
./collect.sh -t=network /tmp/
```

Checking the output:

```
ls -l /collect_output/
total 28
-rw-rw-r-- 1 scadmin scadmin 24696 Mar 30 08:40 collect.log
```

start.sh

The `start.sh` script is used to start Operations Analytics Predictive Insights.

Purpose

This script has the purpose of starting Operations Analytics Predictive Insights.

`start.sh [-t=<topic_name>]`

The script starts the loading of data by Operations Analytics Predictive Insights and starts the analytics instance.

You can use the `-t=<topic_name>` option to specify one topic, so processing of data is started for that one topic.

stop.sh

The `stop.sh` script is used to stop Operations Analytics Predictive Insights.

Purpose

This script has the purpose of stopping Operations Analytics Predictive Insights.

`stop.sh -s [-t=<topic_name>]`

If you don't specify a parameter, the script stops data processing for all topics but does not stop InfoSphere Streams.

You can use the `-s` parameter to stop InfoSphere Streams. If you stop InfoSphere Streams, data processing is stopped for all topics.

You can use the `-t=<topic_name>` option to stop an individual topic.

Properties

The properties that can be used to configure your Operations Analytics Predictive Insights environment.

Component properties

The configuration properties for individual Operations Analytics Predictive Insights components.

Component configuration properties are those properties that are specific to an individual component, such as a specific algorithm or the User Interface. You can display all the component configuration properties by using the following command:

```
admin.sh show -t=<topic>
```

In the output, the names of the component configuration properties begin with the name of the component to which the property applies. For example, in the following sample list of component configuration properties, `robustBounds` relates to the Robust Bounds algorithm and `ui` refers to the User Interface.

```
robustBounds.enabled: true
ui.granularity.default.min: 0
flatLine.maxTrainingWindowDays: 28
granger.enabled: true
nofm.error.count.window.size: 6
relatedEvents.enabled: true
variantInvariant.enabled: true
```

Use the **set** command to set a configuration property:

```
admin.sh set -t=<topic_name> <property> <value>
```

For more information about the **set** command, see “set” on page 108.

Note: Restart Operations Analytics Predictive Insights to activate any changes you make to component configuration properties.

finiteDomain.enabled

The `finiteDomain.enabled` configuration property.

Description

If set to true, enables the Finite Domain algorithm so it can analyze data and generate anomalies. If set to false, disables the Finite Domain algorithm.

Property configuration information

Table 5. finiteDomain.enabled configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

finiteDomain.maxTrainingWindowDays

The `finiteDomain.maxTrainingWindowDays` configuration property.

Description

Specifies the maximum number of days that the Finite Domain algorithm analyzes data for a metric to create a model of the metric's behavior.

Note: The Finite Domain algorithm attempts to train when the smallest `maxTrainingWindowDays` property for any algorithm, excluding `granger.maxTrainingWindowDays`, is reached. However, to train any metric, the algorithm requires that the metric has at least 50% data available relative to the value of the `finiteDomain.maxTrainingWindowDays` property.

Property configuration information

Table 6. *finiteDomain.maxTrainingWindowDays* configuration information

Configuration	Value
Type	INTEGER
Default	14
Min	3
Max	28

finiteDomain.retrainingIntervalMinutes

The `finiteDomain.retrainingIntervalMinutes` configuration property.

Description

Specifies the interval, in minutes, at which the Finite Domain algorithm retrains.

Table 7. *finiteDomain.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

flatLine.enabled

`flatLine.enabled` configuration property.

Description

When set to true, this property enables the Flatline algorithm to analyze your data and generate alarms.

Property configuration information

Table 8. *flatLine.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

flatLine.maxTrainingWindowDays

The flatLine.maxTrainingWindowDays configuration property.

Description

Specifies the maximum number of days that the Flatline algorithm analyzes data for a metric to create a model of the metric's behavior.

Property configuration information

Table 9. flatLine.maxTrainingWindowDays configuration information

Configuration	Value
Type	INTEGER
Default	3
Min	3
Max	28

flatline.retrainingIntervalMinutes

The flatline.retrainingIntervalMinutes configuration property.

Description

Specifies the interval, in minutes, at which the Flatline algorithm retrains.

Table 10. flatline.retrainingIntervalMinutes configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

granger.enabled

granger.enabled configuration property.

Description

When set to “true”, this property enables the Granger algorithm to analyze your data and generate alarms.

Property configuration information

Table 11. granger.enabled configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

granger.maxTrainingWindowDays

The granger.maxTrainingWindowDays configuration property.

Description

Specifies the maximum number of days that the Granger algorithm analyzes data for a metric to create a model of the metric's behavior.

Property configuration information

Table 12. *granger.maxTrainingWindowDay* configuration information

Configuration	Value
Type	INTEGER
Default	3
Min	3
Max	28

granger.retrainingIntervalMinutes

The granger.retrainingIntervalMinutes configuration property.

Description

Sets the interval, in minutes, at which retraining occurs for the Granger algorithm.

Table 13. *granger.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

granger.threads

The granger.threads configuration property.

Description

Sets the number of threads used by the Granger trainer.

Property configuration information

Table 14. *granger.threads* configuration information

Configuration	Value
Type	INTEGER
Default	4
Min	2
Max	64

nofm.error.count.min.size

The nofm.error.count.min.size configuration property.

Description

Minimum number of intervals a metric must be anomalous before Operations Analytics Predictive Insights raises an alarm.

This property must be used with **nofm.error.count.window.size** to establish the error window.

The default configuration is as follows:

```
nofm.error.count.min.size: 3
nofm.error.count.window.size: 6
```

The default configuration generates an alarm, for a given metric, if the data for that metric is seen as anomalous by an analytics instance for three of the last six time intervals. With the default settings, if a metric is seen as anomalous for fewer than three out of the last six intervals, no alarm is generated. The properties can both be set to 1, which results in immediate alarms for all detected anomalies.

When a metric is flagged as anomalous for an interval, and the metric also goes missing, that is, there is no data available for that metric, in a subsequent interval or intervals; the metric is considered anomalous for the interval or intervals where the metric data is missing. Therefore, with the default configuration, if a metric is seen as anomalous for an interval and then it goes missing or is seen as anomalous for two of the following five intervals, an alarm is generated.

Property configuration information

Table 15. *nofm.error.count.min.size* configuration information

Configuration	Value
Type	INTEGER
Default	3
Min	1
Max	30

nofm.error.count.window.size

nofm.error.count.window.size configuration property.

Description

Window size in normalized periods to evaluate anomaly count for alarm filtering. This property must be used with **nofm.error.count.min.size** to establish the error window.

For example:

```
nofm.error.count.min.size: 3
nofm.error.count.window.size: 6
```

The default configuration generates an alarm, for a given metric, if the data for that metric is seen as anomalous by an analytics instance for three of the last six time intervals. With the default settings, if a metric is seen as anomalous for fewer than three out of the last six intervals, no alarm is generated. The properties can both be set to 1, which results in immediate alarms for all detected anomalies.

When a metric is flagged as anomalous for an interval, and the metric also goes missing, that is, there is no data available for that metric, in a subsequent interval or intervals; the metric is considered anomalous for the interval or intervals where the metric data is missing. Therefore, with the default configuration, if a metric is seen as anomalous for an interval and then it goes missing or is seen as anomalous for two of the following five intervals, an alarm is generated.

Property configuration information

Table 16. *nofm.error.count.window.size* configuration property configuration information

Configuration	Value
Type	INTEGER
Default	6
Min	1
Max	30

predominantRange.enabled

The predominantRange.enabled configuration property.

Description

If set to true, enables the Predominant Range algorithm so it can analyze data and generate anomalies. If set to false, disables the Predominant Range algorithm.

Property configuration information

Table 17. *predominantRange.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

predominantRange.maxTrainingWindowDays

The predominantRange.maxTrainingWindowDays configuration property.

Description

Specifies the maximum number of days that the Predominant Range algorithm analyzes data for a metric to create a model of the metric's behavior.

Note: The Predominant Range algorithm attempts to train when the smallest maxTrainingWindowDays property for any algorithm, excluding granger.maxTrainingWindowDays, is reached. However, to train any metric, the algorithm requires that the metric has at least 50% data available relative to the value of the predominantRange.maxTrainingWindowDays property.

Property configuration information

Table 18. *predominantRange.maxTrainingWindowDays* configuration information

Configuration	Value
Type	INTEGER
Default	28

Table 18. *predominantRange.maxTrainingWindowDays* configuration information (continued)

Configuration	Value
Min	3
Max	28

predominantRange.retrainingIntervalMinutes

The `predominantRange.retrainingIntervalMinutes` configuration property.

Description

Specifies the interval, in minutes, at which the Predominant Range algorithm retrains.

Table 19. *predominantRange.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

relatedEvents.enabled

`relatedEvents.enabled` configuration property.

Description

When set to “true”, this property enables the Related Events algorithm, which detects related events.

Property configuration information

Table 20. *relatedEvents.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

relatedEvents.retrainingIntervalMinutes

The `relatedEvents.retrainingIntervalMinutes` configuration property.

Description

Specifies the interval, in minutes, at which retraining occurs for the Related Events algorithm.

Property configuration information

Table 21. *relatedEvents.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440

Table 21. *relatedEvents.retrainingIntervalMinutes* configuration information (continued)

Configuration	Value
Min	1440
Max	40320

robustBounds.enabled

robustBounds.enabled configuration property.

Description

When set to true, this property enables the Robust Bounds algorithm to analyze your data and generate alarms.

Property configuration information

Table 22. *robustBounds.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

robustBounds.maxTrainingWindowDays

The robustBounds.maxTrainingWindowDays configuration property.

Description

Specifies the maximum number of days that the Robust Bounds algorithm analyzes data for a metric to create a model of the metric's behavior.

Note: The Robust Bounds algorithm attempts to train when the smallest maxTrainingWindowDays property for any algorithm, excluding granger.maxTrainingWindowDays, is reached. However, to train any metric, the algorithm requires that the metric has at least 50% data available relative to the value of the robustBounds.maxTrainingWindowDays property.

Property configuration information

Table 23. *robustBounds.maxTrainingWindowDays* configuration information

Configuration	Value
Type	INTEGER
Default	28
Min	7
Max	28

robustBounds.retrainingIntervalMinutes

The robustBounds.retrainingIntervalMinutes configuration property.

Description

Sets the interval, in minutes, at which retraining occurs for the Robust Bounds algorithm.

Table 24. *robustBounds.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

sta.corr.retention.days

sta.corr.retention.days configuration property.

Description

This property specifies the number of days to retain a correlation group in the correlation table in the Operations Analytics Predictive Insights database.

Property configuration information

Table 25. *sta.corr.retention.days* configuration information

Type	INTEGER
Default	15
Min	5
Max	15

sta.corr.decorator.enabled

sta.corr.decorator.enabled configuration property.

Description

When set to true, this property enables the Correlation algorithm. The Correlation algorithm can analyze your data and display, as a related metric, any metric that is correlated with the anomalous metric.

Property configuration information

Table 26. *sta.corr.decorator.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

ui.granularity.default.min

ui.granularity.default.min configuration property.

Description

This configuration property sets the default granularity of data that is displayed in the Operations Analytics Predictive Insights User Interface.

Ideally this property is set so that it matches the granularity that is used by the analytics component. The granularity that is used by the analytics component is set with the **system.aggregation.interval** parameter. With the default value, 0, set for

ui.granularity.default.min, the User Interface automatically uses the value set for **system.aggregation.interval**. To view the value that is set for **system.aggregation.interval**, run the following command:

```
$PI_HOME/bin/admin.sh show -t=<topic name>
```

Where <topic name> is the name of the topic for which you want to display properties.

Note: To list the set of topics, run:

```
$PI_HOME/bin/admin.sh show_topics
```

If the raw data is of a finer granularity than that set for **system.aggregation.interval**, and you want to view the data at that granularity, then adjust this parameter appropriately.

If you have multiple topics configured and those topics have different settings for **ui.granularity.default.min** and if you open KPIs from those topics into a shared chart (for example, through a multi-select of alarms in the Active Event List), then the granularity will become the lowest common multiple of the **ui.granularity.default.min** settings. For example, if you open a KPI from a topic with a **ui.granularity.default.min** setting of 15 minutes with a KPI from a topic with a **ui.granularity.default.min** setting of 10 minutes then the chart will open with a default 30 minute granularity.

Property configuration information

Table 27. *ui.granularity.default.min* configuration information

Configuration	Value
Type	INTEGER
Default	0
Min	5
Max	60

ui.summary.dlg.attr.names

The ui.summary.dlg.attr.names configuration property.

Description

Use this property to specify a comma-separated list of attribute that are visible when you click the information icon in the summary pane in the Operations Analytics Predictive Insights User Interface. For example, to make attributes called "Service", "Node", and "Application" visible, type the following:

```
$PI_HOME/bin/admin.sh set t=<topic name> ui.summary.dlg.attr.names  
"Node,Service,Application"
```

Property configuration information

Table 28. *ui.summary.dlg.attr.names* configuration information

Configuration	Value
Type	STRING
Default	Node

ui.summary.dlg.favorite.attr.name

The ui.summary.dlg.favorite.attr.name configuration property.

Description

Use this property to specify the name of an attribute to display within the summary pane in the Operations Analytics Predictive Insights User Interface. For example, to display an attribute called “Service”, type the following:

```
$PI_HOME/bin/admin.sh set t=<topic name> ui.summary.dlg.favorite.attr.name Service
```

Property configuration information

Table 29. ui.summary.dlg.favorite.attr.name configuration information

Configuration	Value
Type	STRING
Default	

variantInvariant.enabled

variantInvariant.enabled configuration property.

Description

When set to true, this property enables the variant/invariant algorithm to analyze data and generate alarms.

Property configuration information

Table 30. variantInvariant.enabled configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

variantInvariant.maxTrainingWindowDays

The variantInvariant.maxTrainingWindowDays configuration property.

Description

Specifies the maximum number of days that the Variant/Invariant algorithm analyzes data for a metric to create a model of the metric's behavior.

Note: The Variant/Invariant algorithm attempts to start training when the smallest maxTrainingWindowDays property for any algorithm, excluding **granger.maxTrainingWindowDays**, is reached. However, to train any metric, the algorithm requires that the metric has at least 50% data available relative to the value of the variantInvariant.maxTrainingWindowDays property.

Property configuration information

Table 31. variantInvariant.maxTrainingWindowDays configuration information

Configuration	Value
Type	INTEGER
Default	28

Table 31. *variantInvariant.maxTrainingWindowDays* configuration information (continued)

Configuration	Value
Min	7
Max	28

variantInvariant.retrainingIntervalMinutes

The `variantInvariant.retrainingIntervalMinutes` configuration property.

Description

Specifies the interval, in minutes, at which retraining occurs for the Variant/Invariant algorithm.

Table 32. *variantInvariant.retrainingIntervalMinutes* configuration information

Configuration	Value
Type	INTEGER
Default	1440
Min	1440
Max	40320

System properties

The set of system configuration properties for an Operations Analytics Predictive Insights topic.

System configuration properties are those properties that are not specific to individual components, such as an algorithm or the User Interface. You can display all the system properties by using the following command:

```
admin.sh show -t=<topic>
```

In the output, the system properties are prefixed with “system”. For example:

```
system.aggregation.interval: 15
system.alarm.autoclear: true
system.alarm.history.retention.days: 180
system.da.days.before.alarm.cleared: 7
system.da.missing.intervals.before.alarm: 3
```

Use the **set** command to set a configuration property:

```
admin.sh set -t=<topic_name> <property> <value>
```

For more information about the **set** command, see “set” on page 108

Note: Restart Operations Analytics Predictive Insights to activate any changes you make to system configuration properties.

system.aggregation.interval

The `system.aggregation.interval` configuration property.

Description

The aggregation interval that is used by the system. Data is normalized to the same interval so it can be processed by the algorithms. Usually, the aggregation interval needs to be set to the data collection interval, or to the lowest common

multiple of data collection intervals if the system receives data from several data sources. Typical values are 5 minutes, 15 minutes, or 1 hour.

Property configuration information

Table 33. *system.aggregation.interval* configuration information

Configuration	Value
Type	INTEGER
Default	15
Min	5
Max	60

system.alarm.autoclear

The system.alarm.autoclear configuration property.

Description

If the value of this property is set to true, previous analytics alarms, both individual and consolidated, are cleared unless they repeat during the next aggregation interval. However, this property does not affect data availability, progress, or system health alarms, which are never cleared. If this property is set to false, alarms are never cleared by the system and remain in OMNIBus or the third party Event Management system until they are manually cleared by an operator.

Property configuration information

Table 34. *system.alarm.autoclear* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

system.alarm.history.retention.days

system.alarm.history.retention.days configuration property.

Description

This property sets the number of days data to keep in the alarm history table in the Operations Analytics Predictive Insights database. Alarms older than this are purged when the Operations Analytics Predictive Insights Purge job is run.

Property configuration information

Table 35. *system.alarm.history.retention.days* configuration information

Type	INTEGER
Default	180
Min	0
Max	

system.alarm.session.keepalive.hours

system.alarm.session.keepalive.hours configuration property.

Description

Sets the number of hours to reuse an alarm identifier that has been cleared. This property controls the number of alarms displayed in the Active Event List (AEL) when a metric becomes anomalous, clears, and becomes anomalous again. With the default setting, six hours, only one alarm is displayed in the AEL if a metric becomes anomalous, clears, but becomes anomalous again within six hours of the initial time it became anomalous.

For example, if a metric is anomalous at 01:00 and again at 05:00, with the default setting of 6 hours, the AEL has one alarm for the anomalous metric. However, if the **system.alarm.session.keepalive.hours** parameter is set to 3 hours, in this example, the AEL has 2 alarms related to the anomalous metric.

Property configuration information

Table 36. *system.alarm.session.keepalive.hours* configuration information

Type	INTEGER
Default	6
Min	1
Max	72

system.da.days.before.alarm.cleared

system.da.days.before.alarm.cleared configuration property.

Description

The number of days Operations Analytics Predictive Insights data availability will continue to send alarms for a missing resource and metric group before clearing the alarm.

Property configuration information

Table 37. *system.da.days.before.alarm.cleared* configuration information

Configuration	Value
Type	INTEGER
Default	7
Min	1
Max	100

system.da.missing.intervals.before.alarm

system.da.missing.intervals.before.alarm configuration property.

Description

The number of intervals Operations Analytics Predictive Insights can be without data for a resource and metric group before a data availability alarm is raised.

Property configuration information

Table 38. *system.da.missing.intervals.before.alarm* configuration information

Configuration	Value
Type	INTEGER
Default	3
Min	1
Max	1000

system.instance.name

system.instance.name configuration property.

Description

This property sets the SYSTEM instance name.

Property configuration information

Table 39. *system.instance.name* configuration information

Configuration	Value
Type	STRING
Default	PI

system.max.metricgroups

system.max.metricgroups configuration property.

Description

The maximum number of metric groups that can be configured in a model in the Operations Analytics Predictive Insights Mediation Tool.

Property configuration information

Table 40. *system.max.metricgroups* configuration information

Configuration	Value
Type	INTEGER
Default	50
Min	1
Max	50

system.metric.retention.days

system.metric.retention.days configuration property.

Description

This sets the number of days to retain the metric values table AND the alarm table in the Operations Analytics Predictive Insights database. Data and Alarms older than this are purged when the Operations Analytics Predictive Insights Purge job is run.

Property configuration information

Table 41. *system.metric.retention.days* configuration information

Configuration	Value
Type	INTEGER
Default	15
Min	
Max	

system.omnibus.enabled

system.omnibus.enabled configuration property.

Description

Set this property to true if OMNIbus is integrated with Operations Analytics Predictive Insights.

Property configuration information

Table 42. *system.omnibus.enabled* configuration information

Configuration	Value
Type	BOOLEAN
Default	TRUE

system.timeZone

system.timeZone configuration property

Description

This is used to set the time zone used by the Analytics component if the data in a topic is from a different time zone to that of the system on which the Operations Analytics Predictive Insights Analytics component is running.

For example, if the Operations Analytics Predictive Insights Analytics component is running on a machine in London, but the data is from a source in New York.

Property configuration information

Table 43. *system.timeZone* configuration information

Configuration	Value
Type	String
Default	Default

system.topic.allowNegativeNumbers

system.topic.allowNegativeNumbers configuration property

Description

This value of this property determines how Operations Analytics Predictive Insights deals with negative numbers. In some environments, negative numbers are used to represent missing data. When set to false, which is the default value,

negative numbers are treated as missing data. When set to true, negative numbers are processed in the same way as positive numbers.

Property configuration information

Table 44. *system.topic.allowNegativeNumbers* configuration information

Configuration	Value
Type	String
Default	False

Log files created

Read this section for an overview of the log files that are created by the Operations Analytics Predictive Insights components.

Analytics log files

- `$PI_HOME/log/`
- `$PI_HOME/log/streams.spl@<username>/logs/`

where `<username>` is the name of the administrative user, typically `scadmin`, that is used run the `start.sh` and `stop.sh` scripts.

Note: The rotation of the Analytics log files in the `$PI_HOME/log/` directory is automatically managed by `log4j`. For more information, see “Analytics log file rotation” on page 130. There is no automatic rotation of the Streams log files. Each time Streams stops or starts, the Streams log files are copied from `$PI_HOME/log/streams.spl@<username>/logs/` to: `$PI_HOME/log/streams.spl@<username>.<timestamp>/logs`, where `<timestamp>` is the current time. You can manually delete these files if Streams is restarted frequently.

Installation log files

- Installation Manager creates log files in the `/var/ibm/InstallationManager/logs/` directory.
- The Operations Analytics Predictive Insights installation creates a `/tmp/tasp_install_<datestamp>.log` file.
- Problems with the UI installation into JazzSM are logged in the `<JazzSM_Home>/ui/logs/consolecli.log` file, where `<JazzSM_Home>` is typically `/opt/IBM/JazzSM`.

UI log files

- The liberty UI creates log files in the `<UI_HOME>/wlp/usr/servers/piserver/logs` directory, where `<UI_HOME>` is typically `/opt/IBM/scanalytics/UI`.
- JazzSM profile logs are created in the `<JazzSM_Profile_Home>/logs/server1/` directory, where `<JazzSM_Profile_Home>` is typically `/opt/IBM/JazzSM/profile`.

Mediation Tool log files

- On Linux, the Mediation tool logs information in the `$PI_HOME/log/mediationtool.log` file.
- On Windows, the workspace folder has a `.metadata/.log` file.

Database log files

- The DB2 instance owner can view the database log file in the `/sqllib/db2dump/` directory.

Probe log files

- Probe log files are located in `$PI_HOME/probe/omnibus/log/<topic>Probe.log`, where `<topic>` is the name of a topic.

System log files

- The `/var/log/messages*` files can show system problems such as insufficient physical memory available.

Analytics log file rotation

The Log4j properties manage the log file rotation on an Analytics server.

The log4j properties for each topic are located in the `$PI_HOME/sp1/instances/Analytics<topic name>/config/pa_log4j.properties` file. You can customize the properties in this file.

Table 45. Log4j properties for log file rotation

Name	Description
<code>log4j.appender.LOG_FILE.maxFileSize</code>	Used to specify the maximum size for a log file. The default value is 200 MB. When the maximum file size is reached, the file is renamed with <code>.N</code> appended to the file name, where <code>N</code> is an integer from 1 to the value of the <code>log4j.appender.LOG_FILE.MaxBackupIndex</code> property.
<code>log4j.appender.LOG_FILE.MaxBackupIndex</code>	Used to specify the number of backup log files to keep. The default value is 5.

REST Interface

Operations Analytics Predictive Insights has a REST interface. You can use this interface to use Operations Analytics Predictive Insights data in other applications.

The REST interface supports both HTML and JavaScript Object Notation (JSON) formats.

To access the REST interface go to the following URL: `https://<server>:9998/predictiveinsights/rest`, where `<server>` is the server on which the Operations Analytics Predictive Insights User Interface is installed.

Aggregations

You can use the REST interface to retrieve aggregated alarm information for predefined Operations Analytics Predictive Insights datasets.

Datasets

Table 46.

Dataset	Description	Root URL
Alarm Counts by Day	Displays the total number of alarms per day.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/alarms-by-day</code>
Alarm Counts by Hour	Displays the total number of alarms per hour.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/alarms-by-hour</code>
Top Metrics by Alarm Counts	Displays the top metrics ranked by the most number of alarms.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/top-anomalous-metrics</code>
Top Nodes by Alarm Counts	Displays the top nodes ranked by the most number of alarms.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/top-anomalous-nodes</code>
Top Resources by Alarm Counts	Displays the top resources ranked by the most number of alarms.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/top-anomalous-resources</code>
Top Impacting Incidents	Displays the top alarms ranked by the highest number of consolidated alarms they represent.	<code>https://<IP address>:9998/predictiveinsights/rest/aggregations/top-incidents</code>

Parameters

You can pass the following optional parameters in a query request. To pass a single parameter, append the parameter prefixed with a question mark, `?`, to the root URL. To pass multiple parameters, prefix the first parameter with `?`, and each subsequent parameter with an ampersand, `&`.

Table 47. Parameters

Name	Description	Example
topicName	Specifies the topic names to query for data. You can specify multiple topics separated by commas. For example, <code>topicName=topic1,topic2,topic3</code> . If this parameter is not specified then it will query all topics.	<code>https://192.168.1.2:9998/predictiveinsights/rest/aggregations/alarms-by-day?topicName=network</code>

Table 47. Parameters (continued)

Name	Description	Example
startTime	Specifies the start date and time for which to query data. The date and time are accepted in ISO 8601 format and can take the following forms: yyyy-MM-ddTHH:mm:ss.SSSZ for UTC or yyyy-MM-ddTHH:mm:ss.SSSX for time offset. For example, 2013-05-29T00:00:00.000Z for UTC or 2013-05-29T00:00:00.000-05:00 for EST.	https://192.168.1.2:9998/predictiveinsights/rest/aggregations/alarms-by-day?topicName=network&startTime=2015-10-27T00:00:00.000Z
endTime	Specifies the end date and time for which to query data. The date and time are accepted in ISO 8601 format and can take the following forms: yyyy-MM-ddTHH:mm:ss.SSSZ for UTC or yyyy-MM-ddTHH:mm:ss.SSSX for time offset. For example, 2013-05-29T00:00:00.000Z for UTC or 2013-05-29T00:00:00.000-05:00 for EST.	https://192.168.1.2:9998/predictiveinsights/rest/aggregations/alarms-by-day?topicName=network&startTime=2015-10-27T00:00:00.000Z&endTime=2015-10-30T00:00:00.000Z
timeWindow	Specifies a dynamic time window to query data from. Valid values are last24Hours, last7Days, or last30Days. This parameter is ignored when startTime and endTime are specified.	https://192.168.1.2:9998/predictiveinsights/rest/aggregations/alarms-by-hour?last24Hours
num	Specifies the top number of items to return. The default is 10 and the maximum is 1000.	https://192.168.1.2:9998/predictiveinsights/rest/aggregations/alarms-by-hour?last24Hours&num=8
pretty	Returns formatted output when set to true. Valid values are true and false.	

Metrics

Forecast-by-data

You can use the REST interface to generate a forecast for a set of metric values that you provide.

Parameters

You can pass the following optional parameters in a query request.

Table 48. Forecast-by-data parameters

Name	Required	Description	Example
metricData	Yes	Specifies a comma separated list of data for a metric on which the forecast will be based. The metricData list can represent a maximum of 15 days of data at 5-minute intervals, which gives a maximum of 4,320 data elements.	https://192.168.1.2:9998/predictiveinsights/rest/metrics/forecast-by-data?metricData=1.4,2.2,1.8,1.1,2.6,3.4,
pretty	No	Returns formatted output when set to true. Valid values are true and false	

Forecast-by-kpi

You can use the REST interface to generate a forecast for a metric.

Parameters

You can pass the following parameters in a forecast request. To identify the metric to query, you must specify the topic name parameter and either the resourceName and metricName parameters or the resourceId and metricId parameters. To pass the parameters, prefix the first parameter with ?, and each subsequent parameter with an ampersand, &.

Table 49. Forecast-by-kpi parameters

Name	Description
topicName	Specifies the topic that contains the metric to query.
resourceName	Specifies the resource name to query.
metricName	Specifies the metric name to query.
resourceId	Specifies the id of the resource to query.
metricId	Specifies the id of the metric to query.
pretty	Optional parameter that returns formatted output when set to true. Valid values are true and false.

Example

Table 50. Example requests

URL	Description
https://192.168.1.2:9998/predictiveinsights/rest/metrics/forecast-by-kpi?resourceId=392&metricId=8&topicName=network	Generates a forecast for a metric with a resource id of 392 and a metric id of 8 that is within a topic called network.

Table 50. Example requests (continued)

URL	Description
https://192.168.1.2:9998/predictiveinsights/rest/metrics/forecast-by-kpi?resourceName=GigabitLink-n010&metricName=InTotalbytes&topicName=network	Generates a forecast for a metric with a resource name of GigabitLink-n010 and a metric name of InTotalbytes that is within a topic called network.

Baseline

You can use the REST interface to retrieve the baseline for a metric.

Parameters

You can pass the following parameters in a baseline request. To identify the metric to query, you must specify the topicName parameter and either the resourceName and metricName parameters or the resourceId and metricId parameters. To pass the parameters, prefix the first parameter with a question mark, ?, and each subsequent parameter with an ampersand, &.

Table 51. Baseline parameters

Name	Description
topicName	Specifies the topic that contains the metric to query.
resourceName	Specifies the resource name to query
metricName	Specifies the metric name to query
resourceId	Specifies the id of the resource to query
metricId	Specifies the id of the metric to query
StartTime	Optional parameter that specifies the end date and time for which to query data. The date and time are accepted in ISO 8601 format and can take the following forms: yyyy-MM-ddTHH:mm:ss.SSSZ for UTC or yyyy-MM-ddTHH:mm:ss.SSSX for time offset. For example, 2013-05-29T00:00:00.000Z for UTC or 2013-05-29T00:00:00.000-05:00 for EST.
endTime	Optional parameter that specifies the end date and time for which to query data. The date and time are accepted in ISO 8601 format and can take the following forms: yyyy-MM-ddTHH:mm:ss.SSSZ for UTC or yyyy-MM-ddTHH:mm:ss.SSSX for time offset. For example, 2013-05-29T00:00:00.000Z for UTC or 2013-05-29T00:00:00.000-05:00 for EST.
timeWindow	Optional parameter that specifies a dynamic time window to query data from. Valid values are last24Hours, last7Days, or last30Days. This parameter is ignored if startTime and endTime are specified.
pretty	Optional parameter that returns formatted output when set to true. Valid values are true and false

Example query

Table 52. Example requests

URL	Description
<code>https://192.168.1.2:9998/predictiveinsights/rest/metrics/baselines=392&metricId=8&topicName=network</code>	Generates a baseline for a metric with a resource id of 392 and a metric id of 8 that is within a topic called network.
<code>https://192.168.1.2:9998/predictiveinsights/rest/metrics/baselines?resourceName=GigabitLink-n010&metricName=InTotalbytes&topicName=network</code>	Generates a baseline for a metric with a resource name of GigabitLink-n010 and a metric name of InTotalbytes that is within a topic called network.

Example result

```
"x": 1432080900000
"y1": 6.391125448E9
"y2": 2.674418128E9
```

where:

x is the time in epoch format

y1 is the upper value of the baseline at that time

y2 is the lower value of the baseline at that time

Relationships

You can use the REST interface to retrieve the related metric information for each metric that is related to the metric that you query.

Parameters

You can pass the following parameters in a relationships request. To identify the metric to query, you must specify the `topicName` parameter and either the `resourceName` and `metricName` parameters or the `resourceId` and `metricId` parameters. To pass the parameters, prefix the first parameter with a question mark, `?`, and each subsequent parameter with an ampersand, `&`.

Table 53. Relationships parameters

Name	Description
<code>topicName</code>	Specifies the topic that contains the metric to query.
<code>resourceName</code>	Specifies the resource name to query
<code>metricName</code>	Specifies the metric name to query
<code>resourceId</code>	Specifies the id of the resource to query
<code>metricId</code>	Specifies the id of the metric to query

Table 53. Relationships parameters (continued)

Name	Description
time	Optional parameter. If the time parameter is included in the query, the result shows the relationships that existed at that time. If the time parameter is not included in the query, the result shows the latest relationship information that is available. The date and time are accepted in ISO 8601 format and can take the following forms: yyyy-MM-ddTHH:mm:ss.SSSZ for UTC or yyyy-MM-ddTHH:mm:ss.SSSX for time offset. For example, 2013-05-29T00:00:00.000Z for UTC or 2013-05-29T00:00:00.000-05:00 for EST.
pretty	Optional parameter that returns formatted output when set to true. Valid values are true and false

Example query

Table 54. Example requests

URL	Description
<code>https://192.168.1.2:9998/predictiveinsights/rest/metrics/relationships=392&metricId=8&topicName=network</code>	Returns the relationships for a metric with a resource id of 392 and a metric id of 8 that is within a topic called network.
<code>https://192.168.1.2:9998/predictiveinsights/rest/metrics/relationships?resourceName=GigabitLink-n010&metricName=InTotalbytes&topicName=network</code>	Returns the relationships for a metric with a resource name of GigabitLink-n010 and a metric name of InTotalbytes that is within a topic called network.

Example result

```
{ "MetricGroupId": 2, "MetricGroupName": "BpmGroup", "MetricId": 10,
  "MetricName": "InTotalbytes", "Relationships": "[CORRELATION_GROUP]", "ResourceId": 392,
  "ResourceName": "GigabitLink-n010", "TopicName": "network" }, { "MetricGroupId": 2,
  "MetricGroupName": "BpmGroup", "MetricId": 8, "MetricName": "OutTotalbytes",
  "Relationships": "[CORRELATION_GROUP]", "ResourceId": 392, "ResourceName": "GigabitLink-n010",
  "TopicName": "network" }, { "MetricGroupId": 2, "MetricGroupName": "BpmGroup",
  "MetricId": 9, "MetricName": "OutputQueueLength", "Relationships": "[CORRELATION_GROUP]",
  "ResourceId": 392, "ResourceName": "GigabitLink-n010", "TopicName": "network" } ] }
```

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

For trademark attribution, visit the IBM Terms of Use Web site (<http://www.ibm.com/legal/us/>).



Printed in USA