



IBM Analytics, Solutions



IBM Customer Intelligence Administration Guide

Document version : 0.1

CONTENTS

List of Figures.....	vi
1 Solution overview.....	7
1.1 Features.....	7
1.2 Users and benefits	7
1.3 Intended audience.....	7
1.4 Base architecture components	8
1.4.1 IBM Predictive Customer Intelligence (PCI)	8
1.4.2 IBM Open Platform	8
1.4.3 IBM Big Insights 4.1	9
1.4.4 IBM Analytics Accelerator Framework Platform	9
1.5 Design and architecture.....	10
1.6 Typical deployment	10
2 Managing solution components	11
2.1 Managing passwords	11
2.1.1 Example - Managing a DB2 password change	12
2.2 Shutting down and starting up the Customer Insight solution services....	13
2.2.1 Shutting down the Customer Insights solution services	13
2.2.2 Starting up the Customer Insights solution services.....	15
2.3 Managing the Customer Insights datasets.....	16
2.3.1 Starting dataset tasks	17
2.3.2 Stopping.....	20
2.3.3 Checking	20

2.3.4	Viewing log files	21
2.4	Managing reports and dashboards	21
2.4.1	Starting.....	21
2.4.2	Stopping.....	22
2.4.3	Checking	22
2.4.4	Viewing log files	22
2.4.5	List of dashboards and reports	23
2.5	Managing SPSS.....	24
2.5.1	Stopping.....	25
2.5.2	Checking	25
2.5.3	Viewing log files	25
2.6	Managing SPSS Models.....	25
2.6.1	Configuring SPSS Model Execution	25
2.6.2	SPSS Model Logs.....	26
2.7	Managing DBloader jobs	26
2.7.1	Starting DBLoader	26
2.7.2	Checking	27
2.7.3	Viewing log files	27
3	Managing base architecture components.....	28
3.1	Managing PCI components	28
3.1.1	Starting solution services.....	28
3.1.2	Stopping solution services	29
3.2	Managing BigSQL, Hive, Hadoop.....	30
3.2.1	Starting the components.....	30
3.2.2	Stopping the components	31
3.2.3	Checking status of the components	32
3.2.4	Viewing the components log files	32
3.3	Managing Streams	32

3.3.1	Starting.....	33
3.3.2	Stopping.....	33
3.3.3	Checking.....	33
3.3.4	Viewing log files	34
3.4	Managing the Analytics platform.....	34
3.4.1	Managing IBM Open Platform / Big Insights.....	34
3.4.2	Starting.....	34
3.4.3	Checking.....	35
3.4.4	Viewing log files	35
3.5	Managing Postgres, Presto, and Prestogres components	35
3.5.1	Starting the components.....	35
3.5.2	Stopping the components	36
3.5.3	Checking the component status	36
3.5.4	Viewing log files	36
3.5.5	Updating and maintaining configuration properties	37
3.6	Managing Hive, Hadoop and BigSQL	37
3.6.1	Starting.....	37
3.6.2	Stopping.....	37
3.6.3	Checking.....	37
3.6.4	Viewing log files	38
4	Appendices.....	39
4.1	Related information	39

LIST OF FIGURES

Figure 1 Architecture.....	10
Figure 2 Typical deployment.....	10
Figure 3 Password dependencies in solution components	11

1 Solution overview

IBM® Customer Insights for CSPs is a prepackaged and self-contained software solution that integrates the functionality of many IBM software products. The solution contains the following core elements:

- Customer Intelligence dashboards and reports: drives improved marketing and customer care performance
- Predictive Customer Intelligence: advanced and predictive analytics of customer activity across locations, devices, applications, and interests.

1.1 Features

Customer Insights for CSPs comprises of a set of analytics jobs (Scripts, Hive queries, and SPSS models) that are deployed on top of the IBM Analytics Accelerator Framework (AAF) platform to generate datasets for reporting and visualization purposes.

Dataset generation is automated as part of the installation and runs at varying intervals over the underlying datasets. SPSS models that are included in the solution are run against the tables in Hadoop and produce small datasets. The Churn and Net Promoter Score models (NPS) are run on DB2. The Customer Profile models are run on Hadoop.

Customer Insights for CSPs contains a set of dashboards and Cognos reports. Dashboards and reports are run against the datasets in the Analytics Accelerator Framework platform or DB2. Reports are generated for Churn and Net Promoter Score (NPS).

1.2 Users and benefits

The solution provides benefits to a wide range of Telco users as described in the following table.

If you are a	IBM® Customer Intelligence can help you
Business analyst	Design dashboards and reports
Data scientist	Design and build models
Customer service representative	Understand customer behaviour
Solution architect	Design Telco applications
Marketing manager	Understand your customer segments, profiles, behaviour

1.3 Intended audience

This document is intended for people who are administering and maintaining the solution.

1.4 Base architecture components

IBM Customer Insights for CSPs operates with the following base architecture components:

- IBM Predictive Customer Intelligence
 - Cognos Analytics
 - SPSS
 - DB2
- IBM Open Platform (IOP)
- IBM Big Insights 4.1
- IBM Streams

1.4.1 IBM Predictive Customer Intelligence (PCI)

The Customer Insights for CSP solution requires IBM® Predictive Customer Intelligence.

PCI gives you the information and insight that you need to provide proactive service to your customers. PCI version 1.1.1 includes the following components:

- IBM WebSphere Application Server
- SPSS Products (SPSS Modeler Server, Modeler Client, Collaboration and Deployment Services)
- Cognos Analytics
- DB2 Enterprise Server
- IBM Message Queue
- IBM Integration Bus
- Client interfaces

IBM Customer Insights for CSPs uses the following PCI component products:

- SPSS Modeler Server for churn prediction and customer profiling models
- SPSS Collaboration and Deployment Services for deploying models
- Cognos Analytics for Churn, NPS, and visualization of reports and dashboards for Over the Top (OTT) content and User Profiles.
- DB2 for TELCO database and provisioning

1.4.2 IBM Open Platform

IBM Open Platform V4.1 provides a set of open source tools used for data sets and analysis. IBM Customer Insights for CSPs uses the following tools from IBM Open Platform:

- Ambari is an Apache Hadoop open source component and part of the IBM Open Platform. Ambari is a system for provisioning, managing, and monitoring Apache Hadoop clusters.
- Hadoop, Hive, Knox and Parquet for data set storage and encoding

- Sqoop to transfer data, for analysis by SPSS jobs, from Hadoop to DB2.

1.4.3 IBM Big Insights 4.1

IBM Big Insights 4.1.0 is a collection of powerful value-add services that can be installed on top of the IBM Open Platform with Apache Hadoop.

The value-add services in Big Insights include: IBM Big SQL, IBM Big Sheets, IBM Big R, and IBM Text Analytics.

IBM Customer Insights for CSPs uses Big SQL. Big SQL is an IBM DB2-style interface to Hadoop. Big SQL is used by the CEA reports to access data for the reports.

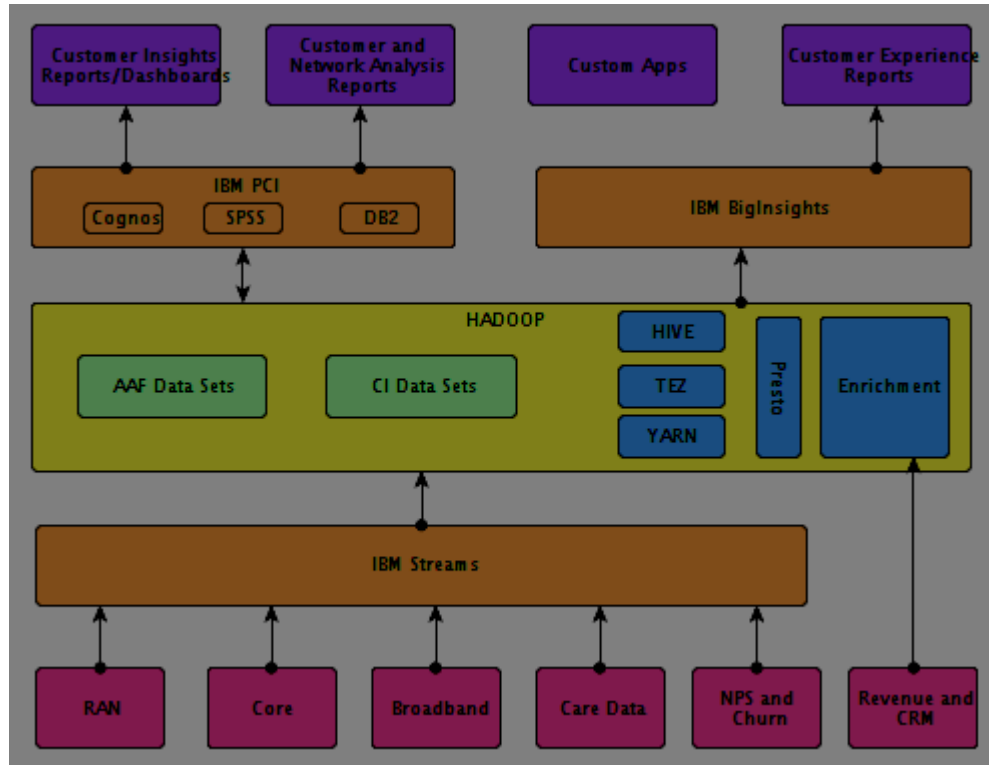
1.4.4 IBM Analytics Accelerator Framework Platform

The IBM Analytics Accelerator Framework (AAF) platform consists of the base data sets and services, upon which Customer Insights for CSPs bases its analysis. IBM AAF Version 1.0.4 is used. The underlying platform version is Analytics Platform 3.1.0.0.1.

1.5 Design and architecture

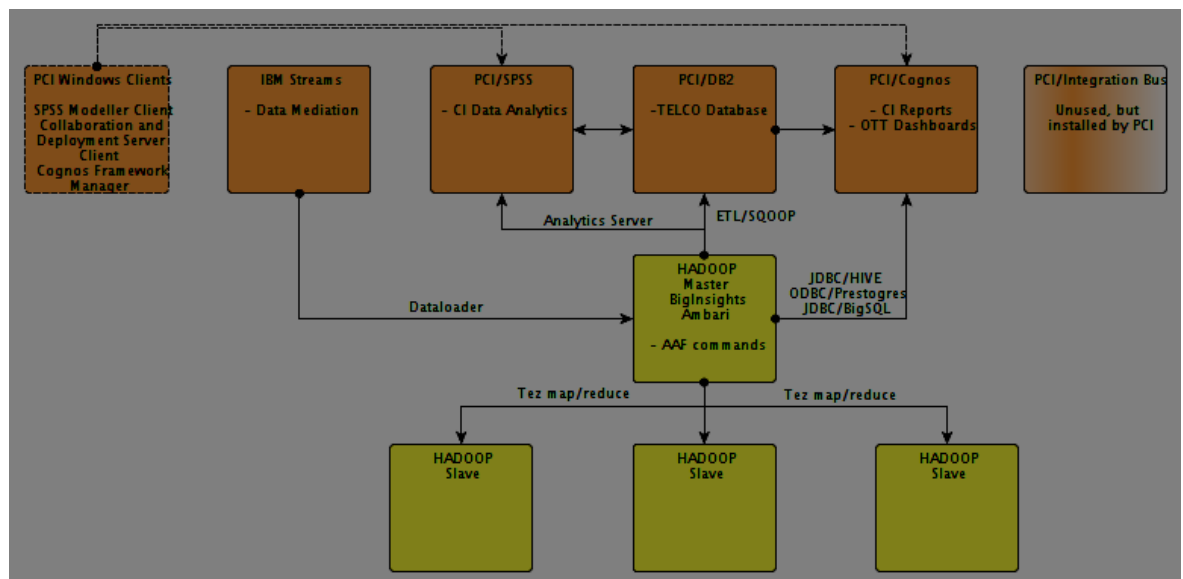
Figure 1 describes the solution design and architecture.

Figure 1 Architecture



1.6 Typical deployment

Figure 2 Typical deployment



2 Managing solution components

Perform the tasks described in this section to keep your solution components running smoothly.


2.1 Managing passwords

Administering the components requires a consistent password management policy. Figure 3 displays the password interdependencies between components.

In a typical deployment, Sqoop, SPSS, Cognos and BigSQL components reference DB2 credentials. Sqoop and BigSQL reference Hive and Hadoop credentials. By updating or changing a DB2 password, other component credentials in the solution must also be changed.

It is recommended that you maintain credentials in the deployment worksheet. A sample worksheet is contained in the Customer Insights for CSPs Deployment Guide.

Figure 3 Password dependencies in solution components



When loading NPS and Churn data, for example, credentials for Hive, DB2 and SPSS are referenced. To view the settings, log in to the Hadoop Master node (aafnode) and go to `/home/boss/CSP_CustomerInsight_DBLoader/conf`. The Hive, DB2 and SPSS

credentials are in the `nps_config.properties` and `churn_config.properties` files.

2.1.1 Example - Managing a DB2 password change

If the DB2 password changes on the PCI DB2 node, the following components need to be reconfigured to use the new DB2 password.

- Database Loader
- SPSS components
- Cognos

To reconfigure the component passwords, complete the following steps.

Updating the Database Loader

Refer to section 3.6.3 in the IBM Customer Insights for CSPs Deployment Guide for steps on how to change the db2 password. After the db2 password in the configuration files is updated complete the steps in section 3.6.5 to base64 encode the file.

Updating the SPSS component password

The following updates must be made to SPSS component credentials.

1. Update C&DS Repository Connection in WebSphere Application Server:
 - Login to the SPSS WebSphere instance.
 - In the left menu expand the 'Resources' menu and navigate to JDBC > Datasources.
 - Select CDS_DataSource.
 - In the menu on the right hand side select JAAS-J2C authentication data located under the 'Related Items' menu.
 - Select CDS_Auth_Alias and update the password to the new DB2 password.
 - Press the OK Button.
 - In the Top Menu select Data sources.
 - Select the checkbox beside the CDS_Datasource and then select the Test connection button.
 - The test should succeed if DB2 is running and the correct credentials were entered.
2. Update the SPSS DB2 credentials in C&Ds Services Client:
 - On the PCI Windows Client Node open the C&Ds application.
 - Navigate to the resource Definitions > Credentials folder.
 - Update the db2 password.
 - In the Churn and NPS jobs select each node in the job. Navigate to the ODBC Data Sources tab and update the credentials.

Updating the Cognos password

1. Update the Cognos Configuration:
 - On the Cognos Analytics PCI node navigate to:
`/opt/ibm/cognos/analytics/bin64`
 - Run `cogconfig.sh`
 - Navigate to Data Access > Content Manager > Content Store.
 - Select the User ID and password row.
 - Select the Pencil button.
 - Enter the New DB2 Password.
 - Select OK.
 - Save the Settings.
 - Start Cognos and ensure you can log on to Cognos.
2. Update the Cognos Reports DB2 password:
 - Logon to the Cognos Analytics web client.
 - On the IBM Analytics page, select Manage and then select Data Servers.
 - Select the DB2 Data Server.
 - Update the DB2 password.
 - Save the Data Server.

2.2 Shutting down and starting up the Customer Insight solution services

The Customer Insight solution is a complex solution with various inter-dependent sub-components. The following sections detail the recommended shutdown and start-up procedures.

2.2.1 Shutting down the Customer Insights solution services

Note: Before you stopping solution services, ensure that the dataset cron jobs are not scheduled to run. Datasets are run on an hourly, daily, or weekly basis using cron jobs at a preconfigured time interval (see section 2.3).

To shutdown solution services, complete the following steps:

1. Stop Infosphere Streams.

In general, an instance is stopped using the command:

```
streamtool stopinstance -i <instancename>
```

Refer to the [AAF Mediation and Installation Guide](#) for information on starting and stopping stream instances. More information is available in the IBM Infosphere Stream knowledge center at:

http://www.ibm.com/support/knowledgecenter/SSCRJU_4.1.1/com.ibm.streams.admin.doc/doc/ibminfospherestreams-instance-running.html

2. Stop the SPSS Collaboration and Deployment Server.
 - Log on to the PCI SPSS node
 - Stop the repository sever by using the WebSphere administration tools
 - Click Start > Control Panel > Administrative Tools > Services. Select IBM SPSS Collaboration and Deployment Server, and click Stop Service.
3. Stop SPSS Modeler Server on the PCI SPSS node.
 - Log in to the PCI SPSS node computer as root and run the following commands:

```
/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh stop  
/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh kill
```
4. Stop the SPSS Analytic Server on the Hadoop master node (aafnode)
 - Select the SPSS Analytic Server service from the Ambari console.
 - Click Service Actions > Stop
5. Stop `bis_daemon` on the Hadoop master node (aafnode)
 - Log in to the aafnode node computer as root.

```
cd /opt/tnf/apps/bis-main/  
/bis-main.sh stop
```
6. Stop Cognos on the PCI/Cognos node (pcibinode)
 - Log in to the pcibinode computer as root.

```
cd /opt/ibm/cognos/analytics/bin64  
/cogconfig.sh -stop
```
7. Stop DB2 on the PCI/DB2 node (on the pcidbnode)
 - Log in to the data node computer as root.
 - In a terminal window, type the following command to change the DB2 instance owner:

```
su - db2inst1
```
 - Enter the following command to stop the DB2 administration server:

```
db2stop
```
8. Stop Presto on the Hadoop Master node (aafnode)
 - Select the Presto service from the Ambari web interface.
 - Click Service Actions > Stop
9. Stop Hive on the Hadoop Master node (aafnode)
 - Select the Hive service from the Ambari web interface.
 - Click Service Actions > Stop
10. Stop Hadoop on the Hadoop Master node (aafnode)

- Select the Hadoop service from the Ambari web interface.
- Click Service Actions > Stop

2.2.2 Starting up the Customer Insights solution services

To startup solution services, complete the following steps:

1. Start Hadoop on the Hadoop Master node (aafnode)
 - Select the Hadoop service from the Ambari web interface.
 - Click Service Actions > Start
2. Stop Hive on the Hadoop Master node (aafnode)
 - Select the Hive service from the Ambari web interface.
 - Click Service Actions > Start
3. Start Presto on the Hadoop Master node (aafnode)
 - Select the Presto service from the Ambari web interface.
 - Click Service Actions > Start
4. Start DB2 on the PCI/DB2 node (on the pcidbnode)
 - Log in to the data node computer as root.
 - In a terminal window, type the following command to change the DB2 instance owner:

```
su - db2inst1
```
 - Enter the following command to start the DB2 administration server:

```
db2start
```
5. Start Cognos on the PCI/Cognos node (pcibinode)
 - Log in to the pcibinode computer as root.

```
cd /opt/ibm/cognos/analytcs/bin64
```

```
export JAVA_HOME=/opt/ibm/cognos/analytcs/jre
```

```
/cogconfig.sh -s
```
6. Start `bis_daemon` on the Hadoop master node (aafnode)
 - Log in to the aafnode node computer as root.

```
cd /opt/tnf/apps/bis-main/
```

```
/bis-main.sh start
```
7. Start the SPSS Analytic Server on the Hadoop master node (aafnode)
 - Select the SPSS Analytic Server service from the Ambari console.
 - Click Service Actions > Start
8. Start SPSS Modeler Server on the PCI SPSS node
 - Log in to the PCI SPSS node computer as root and run the following commands:

```
/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh start
```

9. Start the SPSS Collaboration and Deployment Server

- Log on to the PCI SPSS node
- Start the repository sever by using the WebSphere administration tools
- Click Start > Control Panel > Administrative Tools > Services. Select IBM SPSS Collaboration and Deployment Server, and click Start Service.

10. Start Infosphere Streams.

In general, start an instance using the command:

```
streamtool startinstance -i <instancename>
```

Refer to the [AAF Mediation and Installation Guide](#) for information on starting and stopping stream instances. More information is available in the IBM Infosphere Stream knowledge center at:

http://www.ibm.com/support/knowledgecenter/SSCRJU_4.1.1/com.ibm.streams.admin.doc/doc/ibminfospherestreams-instance-running.html.

2.3 Managing the Customer Insights datasets

Datasets are run on an hourly, daily, or weekly basis using cron jobs. Each dataset installed configures a cron job to run the dataset at a preconfigured time interval. To manage the cron jobs that are configured, run the following commands at the UNIX prompt on the Hadoop Master node (aafnode):

1. Log on as `boss` user to the Hadoop Master node (aafnode).

```
su boss
```

Note: The `boss` user is created when AAF is installed.

2. List of cron jobs by running the command:

```
crontab -l
```

Edit the cronjobs by running the command:

```
crontab -e
```

Datasets require a specific order of execution. Numerous datasets have pre-requisite datasets that must finish executing prior to the dataset running.

Refer to section 3.8.4 and Figure 6 in the IBM Customer Insights for CSPs Deployment Guide for a description of the order of execution of datasets. Phase 1 datasets are executed prior to phase 2. The User Profile dataset must execute after the Customer Profile models.

Note: If in the live environment the cron jobs do not execute in time, the administrator must update the cron schedule to ensure the order of execution. The cron job for the database loader must execute after the datasets in phase 1 and 2 have completed.

The datasets and associated scripts are displayed in the following table.

Dataset	Script
---------	--------

churn_dataset	/opt/tnf/apps/churn-dataset/run_churn_dataset.sh
customer_behaviours	/opt/tnf/apps/customer-behaviours/customer_behaviour/run_all_hourly.sh /opt/tnf/apps/customer-behaviours/customer_behaviour/run_daily_rollup.sh /opt/tnf/apps/customer-behaviours/customer_behaviour/run_weekly_rollup.sh /opt/tnf/apps/customer-behaviours/weighted_interest/run_weighted_interest_daily.sh
customer_experience	/opt/tnf/apps/customer-experience/run_cea_hourly.sh /opt/tnf/apps/customer-experience/run_cea_daily.sh /opt/tnf/apps/customer-experience/run_cea_weekly.sh
customer-profile	/opt/tnf/apps/churn-dataset/run_churn_dataset.sh
customer-profile-data-setup	/opt/tnf/apps/customer-profile-data-setup/run_customer-profile-data-setup.sh
net-promoter-score	/opt/tnf/apps/net-promoter-score/run_nps.sh
OTT	/opt/tnf/apps/ott-applications/run_ott-applications.sh
user-profile	/opt/tnf/apps/user-profile/run_user-profile.sh

2.3.1 Starting dataset tasks

This table contains a summary of commands and steps relating to dataset administration.

On the aafnode node	Do the following
Removing a dataset table group	su boss cd /opt/tnf/apps/bis-main/

	<pre>/bis-dbconfig/install-db.sh <dataset table group> --remove</pre>
Removing a dataset RPM	<pre>su root rpm -qg 'Application/TNF' sudo rpm -e <name></pre> <p>Note: <name> is taken from the name of the dataset that you want to remove and can be found in the list returned by Application/TNF.</p> <p>Verify that the tables no longer exist in Hive by connecting to Hive and checking if you can query from the table.</p> <p>If the table still exists check that bis-main is running.</p>
Installing a dataset RPM	<p>Take the correct dataset rpm for the dataset you wish to install from either Jenkins or build it locally.</p> <p>To build a dataset locally, select the root <code>pom.xml</code> file, right click it, select run as maven build and type clean install pre-site in the goals section. You can locate the build rpm by checking the output in your console window in eclipse.</p> <p>Copy the RPM to your hive machine.</p> <pre>sudo rpm -ivh <rpm></pre>
Installing a table group	<p>Go to <code>opt/tnf/apps/bis-main</code> and run:</p> <pre>su boss ./bis-dbconfig/install-db.sh <dataset table group name></pre> <p>Ensure the analytics platform daemon is running so the tables get created. If not, start the daemon as follows:</p> <pre>cd /opt/tnf/apps/bis-main/ ./bis-main.sh start</pre> <p>Verify that the tables no longer exist in Hive by connecting to Hive and checking if you can query from the table. If the table still exists, check that bis-main is running.</p>
Populating a dictionary table	<p>Dictionary tables are provisioned automatically or through static CSV files.</p> <p>To populate a dictionary table with a CSV file, use the Analytic Platform load script.</p>

	<pre>cd /opt/tnf/apps/bis-main-var/bis- provisioning-tool ./load.sh -t <table_name> -f <csv file name></pre> <p>Note All installed datasets packaged CSV files are located at:</p> <pre>/opt/tnf/apps/bis-main-var/bis- provisioning-tool/csv_files</pre>
Connecting to Hive / BigSQL	<p>Open a terminal window and start Hive:</p> <pre>su boss \$HIVE_HOME/bin/hive</pre> <p>Set your schema to tnf by typing</p> <pre>>use tnf;</pre> <p>Show all tables in the tnf schema by typing:</p> <pre>>show tables;</pre>
Querying the data using a Hive text file	<pre>\$HIVE_HOME/bin/hive -f <file></pre>
Starting the Analytics Platform daemon	<p>Ensure the analytics platform daemon is running so the tables get created. If not, start the daemon:</p> <pre>cd /opt/tnf/apps/bis-main/ ./bis-main.sh start</pre>
Installing the Customer Insights datasets	<p>As the root or sudo user, start the installer from the directory where the installation package is uncompressed.</p> <pre>sudo ./analytics- platform/install_telsol.sh</pre>
Provisioning data	<p>Provisioning data to be loaded must be stored as CSV files in:</p> <pre>/opt/tnf/apps/bis-main- var/bisprovisioning-tool/csv_files/</pre> <p>If you have any data to provision, copy the data files to that directory</p>
Setting up Yarn queues for dataset queries	<p>Go to <code>/opt/tnf/apps/dataset-common/master-config</code></p> <p>Edit the <code>common.properties</code> file</p>

Listing the queues configured in a terminal	Run the command: <code>hadoop queues -list</code>
---	--

2.3.2 Stopping

On the aafnode	Do the following
Stopping all services in Ambari	<p>From the host machine open a browser and open the Ambari web interface</p> <p>Log in to the Ambari server. Type the default user name and password: admin/admin.</p> <p>The Ambari dashboard opens. Use the options to start / stop services.</p>

2.3.3 Checking

On the aafnode	Do the following
Checking the tables are creating	<pre>cd /opt/tnf/apps/bis-main/bis-demon/log</pre> <pre>tail -f demon.log</pre> <p>If the tables are not populating, then check the datasets log file in the following directory, and tail the log:</p> <pre>cd /opt/tnf/apps/<dataset-name>/log</pre> <pre>tail -f tools.log</pre>
Checking the dataset log file	<pre>cd /opt/tnf/apps/<dataset-name>/log</pre> <pre>tail -f tools.log</pre>
Checking Hive query errors	<pre>cd /tmp/boss</pre> <pre>tail -f hive.log</pre>
Checking the permissions on the log folder in the dataset	<p>If the log folder is owned by root, run the following Linux command to change the owner to boss.</p> <pre>sudo chown -R boss:boss</pre> <pre><log folder path></pre>

2.3.4 Viewing log files

If, when you run the dataset population Linux script the log file cannot be created, then check the permissions on the log folder in the dataset.

To check tables are creating, run the following commands

```
cd /opt/tnf/apps/bis-main/bis-demon/log
tail -f demon.log
```

To view the progress of dataset deployment, check the `install_telsol.log` in the deployment directory.

2.4 Managing reports and dashboards

Use Cognos 11 and Cognos Framework Manager to perform administrative tasks related to report and dashboard management.

Cognos is installed on the PCI/Cognos node (pcibinode). To view the Administration console in Cognos Analytics, select **Manage -> Administration** console from the menu.

Reports are deployed to: `/opt/ibm/cognos/analytics/deployment`

Report images are stored at:

`/opt/ibm/cognos/analytics/webcontent/bi/samples/images/`

Churn report images are contained in reports, and automatically saved to:

`/opt/ibm/cognos/analytics/temp/vgs/cache`

To manage report and dashboard content:

1. Start Cognos Analytics.

```
cd /opt/ibm/cognos/analytics/bin64/
export JAVA_HOME=/opt/ibm/cognos/analytics/jre
./cogconfig.sh -s
```

Open Cognos Analytics in your browser:

`http://<hostname>:9300/bi`

2. Open the Administration console, by selecting **Manage -> Administration** console from the menu
3. Select the **Configuration -> Content Administration** to manage import and export of report and dashboard content using the wizard.

2.4.1 Starting

On the pcibinode node	Do the following
	<pre>export JAVA_HOME=/opt/ibm/cognos/analytics/jre export LD_LIBRARY_PATH=\$LD_LIBRARY_PATH:/usr/lib32/</pre>

	<pre>/opt/ibm/cognos/analytics/bin64 /cogconfig.sh -s</pre>
--	---

2.4.2 Stopping

On the pcibinode node	Do the following
	<pre>/opt/ibm/cognos/analytics/bin64 /cogconfig.sh -stop</pre>

2.4.3 Checking

On the pcibinode node	Do the following
	<pre>cd /opt/ibm/cognos/analytics/logs</pre>
Verify deployment of reports and dashboards	<ol style="list-style-type: none"> 1. Start Cognos: <pre>cd /opt/ibm/cognos/analytics /bin64/ export JAVA_HOME=/opt/ibm/ cognos/analytics/jre ./cogconfig.sh -s</pre> 2. Open Cognos Analytics in your browser: <p><a href="http://<hostname>:9300/bi">http://<hostname>:9300/bi</p> <p>Select Team content -> CI -> dashboards to view CI dashboards.</p> <p>Select Team content -> CI -> reports to view CI reports.</p>

2.4.4 Viewing log files

On the pcibinode node	Do the following
	<pre>cd /opt/ibm/cognos/analytics/logs</pre> <p>Check the files: cogconfig_response.csv.*.log and cogserver.log</p>
	<pre>tail -f /opt/ibm/cognos/analytics/logs /cogserver.log</pre>

	tail -f /opt/ibm/cognos/analytics/logs /p2pd_messages.log
--	--

2.4.5 List of dashboards and reports

Dashboard Name	Description
Over the Top Dashboard (OTT)	Displays device and application usage
OTT Overview	Displays a snapshot of fastest growing OTT applications by volume usage
Customer Summary	Displays video, music, and voice and message applications by percentage of customers
Application Summary	Displays top OTT applications by volume and number of users for the past 7 days filtered by device type, or device operating system
Price Plan Summary	Displays top OTT applications by number of customers and application category, filtered by price plan
Application Daily Summary	Displays top 20 applications by number of customers, total volume, and average daily volume
Device Daily Summary	Displays top 20 devices by number of customers, total volume, and average daily volume
Change Summary	Displays percentage customer change by application and device, and percentage volume change by application and device
User Profile Dashboard	Displays customer profile information
Mobility Profile	Display mobility data filtered by age, device and gender
Usage Profile	Displays usage data filtered by device, age and gender
Subscriber Profile	Displays subscriber profile data (device, mobility, usage, experience, location,

	categories, interests, domains) filtered by IMSI
Report Name	Description
Churn report	Churn management report displaying churn overview and propensity to churn score
NPS report	Customer Insight report displaying net promoter score

Customizing the dashboards

Cognos Analytics provides the opportunity to create new dashboards or customize the existing dashboards. Use the Cognos Analytics options and palette to:

- Customize dashboard visualizations
- Create new data modules and data sources

The customizing options are fully described in the IBM Cognos Analytics v11.0 Knowledge Center. For more information, refer to the links below.

http://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.0.0/com.ibm.swg.ba.cognos.cbi.doc/welcome.html

http://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.0.0/com.ibm.swg.ba.cognos.wig_cr.doc/c_gtstd_ica_overview.html

<http://www-01.ibm.com/support/docview.wss?uid=swg27047187>

Managing user access to dashboards

Permissions determine a user's ability to perform a specific action, or access a feature or object. Permissions for a user are a combination of the settings for the user account and the settings for the groups and roles where the user is a member.

When a user is a member of more than one group or role, a deny setting in one group or role takes precedence over a grant setting in a different group or role.

In Cognos Analytics, users see Read, Run, Write, and Full permissions for items. These simple permissions represent combinations of more granular permissions that administrators use to control access.

2.5 Managing SPSS

On the analytics server, review the following SPSS Modeler log:

```
/opt/IBM/SPSS/ModelerServer/17.1/log/messages.log
```

On the aafnode node	Do the following
----------------------------	-------------------------

<p>Verify the SPSS Modeler server is started</p>	<p>On the analytics server log on as the ibmadmin user.</p> <p>Run the following command.</p> <pre>/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh list</pre> <p>If the server is not listed, start the SPSS Modeler server by running the following command:</p> <pre>/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh start</pre>
--	--

2.5.1 Stopping

On the aafnode node	Do the following
<p>Stop the SPSS Modeler Server</p>	<p>On the analytics server log on as the ibmadmin user.</p> <p>Run the following command.</p> <pre>/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh stop</pre> <pre>/opt/IBM/SPSS/ModelerServer/17.1/modelersrv.sh kill</pre>

2.5.2 Checking

On the aafnode node	Do the following
<p>Check the logs</p>	<pre>/opt/IBM/SPSS/ModelerServer/17.1/log/messages.log</pre>

2.5.3 Viewing log files

On the aafnode node	Do the following
<p>View the logs</p>	<pre>/opt/IBM/SPSS/ModelerServer/17.1/log/messages.log</pre>

2.6 Managing SPSS Models

2.6.1 Configuring SPSS Model Execution

The Churn and NPS models are triggered by the Database Loader script. Refer to section 3.6.3 of the IBM Customer Insights for CSPs Deployment Guide for detailed information on updating the configuration options including the SPSS models to be triggered.

The Customer Profile models are triggered by the Customer Profile Data Setup dataset. Refer to section 3.8.4 in the IBM Customer Insights for CSPs Deployment Guide for detailed information on configuring/modifying the Customer Profile model execution.

Logs for the triggering of the models from the Database loader are located in `/home/boss/CSP_CustomerInsight_DBLoader/log` directory.

2.6.2 SPSS Model Logs

Models are executed and trained in Collaboration and Deployment Services. The Collaboration and Deployment services client is a Windows client. To monitor the logs and history of execution of SPSS jobs complete the following steps:

- Open the Collaboration and Deployment Services Client.
- Create a connection to the Collaboration and Deployment Services Server. This should already exist from deploying the SPSS content.
- Navigate to the correct job.

NPS & Churn jobs are located at `/Content Repository/CSP_CustomerInsight_NPS/Jobs`

Customer Profile jobs are located at `/Content Repository/CSP_CustomerInsight_Profile/Jobs`.

- Right click the job and select 'Show Job History'.
- Navigate to the correct job in the Job History tab.
- Expand the latest job.
- Verify the status of the job in the Status field in the Job History tab.
- Expand the Job to view the Log file.

2.7 Managing DBloader jobs

Sqoop is used to transfer data for analysis by SPSS jobs from Hadoop to DB2. Churn and NPS Database Loader configuration files specify Hadoop and DB2 settings

2.7.1 Starting DBLoader

On the aafnode node	Do the following
Churn Database Loader	<pre>cd /home/boss/CSP_CustomerInsight_DBLoader</pre> <pre>./runDBExport.sh churn</pre> <p>This is a one off bulk load command, and does not start a process that runs at a time interval.</p>
NPS Database Loader	<pre>cd /home/boss/CSP_CustomerInsight_DBLoader</pre> <pre>./runDBExport.sh nps</pre>

	This is a one off bulk load command, and does not start a process that runs at a time interval.
--	---

2.7.2 Checking

On the pcidbnode node	Do the following
Monitor the output of the <code>runDBExport.sh</code> command.	<p>Log on to the PCI DB2 node with a user ID that has access to the IBM DB2® database. For example, <code>db2inst1</code>.</p> <p>Connect to the database:</p> <pre>Db2 connect to TELCO;</pre> <p>Select from one of the Churn tables to ensure that one table populated is with data.</p> <p>For example run:</p> <pre>Select count(*) from BBCI.cgr_device</pre>

2.7.3 Viewing log files

On the aafnode node	Do the following
	Check the <code>churn_dbexport.log</code> located in the main <code>CSP_CustomerInsight_DBLoader</code> folder.

3 Managing base architecture components

Perform the tasks described in this section to keep your base architecture components running smoothly.

3.1 Managing PCI components

You must start the IBM® Predictive Customer Intelligence node services in a specific order. Start the node services in the following order:

1. First start the IBM DB2 instance on the IBM Predictive Customer Intelligence data node computer.
2. Then start the IBM® SPSS® services on the IBM Predictive Customer Intelligence Predictive Analytics node computer.
3. Then start the IBM Cognos Analytics services on the IBM Predictive Customer Intelligence Cognos Analytics (CA) node computer.

3.1.1 Starting solution services

On this node	Do the following
Data node (pcidbnode)	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the data node computer as <code>root</code>. 2. In a terminal window, type the following command to change the DB2 instance owner: <code>su - db2inst1</code> 3. Enter the following command to start the DB2 administration server: <code>db2start</code>
Predictive analytics node (pcipaode)	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the Predictive Analytics node computer as <code>root</code>. 2. Go to the WebSphere® Application Server profile bin directory. For example, go to <code>/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin</code>, if you used the default name, or <code>/opt/IBM/WebSphere/AppServer/profiles/CNDSprofile/bin</code>, if you used the solution installer. 3. Enter the following command: <code>./startServer.sh server1.</code> 4. Go to the IBM SPSS Modeler Server directory. For example, go to <code>/opt/IBM/SPSS/ModelerServer/17.1.</code> If you manually installed and accepted the defaults, the path is <code>/usr/IBM/SPSS/ModelerServer/17.1.</code> 5. Enter the following command:

	<pre>./modelersrv.sh start</pre> <p>6. Go to the IBM SPSS Statistics Server directory. For example, go to /opt/IBM/SPSS/StatisticsServer23/bin.</p> <p>7. Enter the following command: ./start_statistics_server.</p> <p>Run the command as ./start_statistics_server -d if you want to run the application as a daemon process.</p>
Integration Bus node	Not used
Cognos Analytics node (pcibinode)	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the pcibinode node computer as root. 2. Go to the Cognos Analytics installation directory. For example, go to /opt/ibm/cognos/analytics/bin64. 3. Enter the following command: ./cogconfig.sh -s

3.1.2 Stopping solution services

You must stop the IBM® Predictive Customer Intelligence node services in a specific order.

Stop the node services in the following order:

1. First stop the services on Cognos Analytics (CA) node
2. Then stop services on the Integration Bus node
3. Then stop the services on Predictive Analytics node
4. Then stop the services on data node

On this node	Do the following
Cognos Analytics node (pcibinode)	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the CA node computer as root. 2. Go to the Cognos Analytics installation directory. For example, go to /opt/ibm/cognos/analytics/bin64. 3. Enter the following command: ./cogconfig.sh -stop
Integration Bus node	Not used in CI

<p>Predictive Analytics node (pcipanode)</p>	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the Predictive Analytics node computer as root. 2. Go to the WebSphere® Application Server profile bin directory. For example, go to <code>/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin</code>, if you used the default name, or <code>/opt/IBM/WebSphere/AppServer/profiles/CNDSprofile/bin</code>, if you used the solution installer. 3. Enter the following command: <code>./stopServer.sh server1</code> 4. Go to the IBM SPSS Modeler Server directory. For example, go to <code>/opt/IBM/SPSS/ModelerServer/17.1</code>. If you manually installed and accepted the defaults, the path is <code>/usr/IBM/SPSS/ModelerServer/17.1</code>. 5. Enter the following command: <code>./modelersrv.sh stop</code> 6. To verify whether any services are still running enter the following command: <code>ps -ef grep statisticsd</code> 7. To stop any services that are still running enter the following command: <code>kill -9 'cat statisticsd.pid'</code>
<p>Data node (pcidbnode)</p>	<p>On Linux operating systems:</p> <ol style="list-style-type: none"> 1. Log in to the data node computer as root. 2. In a terminal window, type the following command to change the DB2 instance owner: <code>su - db2inst1</code> 3. Enter the following command to stop the DB2 administration server: <code>db2stop</code>

3.2 Managing BigSQL, Hive, Hadoop

To manage IBM Open Platform or Big Insights components, use the Ambari Server console. Starting, stopping and checking all services can be performed in the Ambari console. In general, you'll just want to check that all services are running.

Command line options are also possible and are included for reference in the following sections.

3.2.1 Starting the components

This section describes how to start components by using InfoSphere BigInsights commands. Additionally, you can start components by logging into the InfoSphere BigInsights Console, and clicking the Cluster Status tab.

Note: In the following tables, `BIGINSIGHTS_HOME=/opt/ibm/biginsights`. Scripts are located in the `/opt/ibm/biginsights/bin`.

On the aafnode	Do the following
All components	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./start-all.sh</code>
BigSQL	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./start-all.sh bigsql</code>
Hive	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./start.sh hive</code>
Hadoop	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./start.sh Hadoop</code>

3.2.2 Stopping the components

This section describes how to stop components by using InfoSphere BigInsights commands. Before you stop the components in the cluster, stop all applications and reduce other file system activities.

Additionally, you can stop components by logging into the InfoSphere BigInsights Console, clicking the Cluster Status tab.

On the aafnode	Do the following
All components	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./stop-all.sh</code>
BigSQL	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./stop.sh bigsql</code>
Hive	From the management node, run the following command from the <code>\$BIGINSIGHTS_HOME/bin</code> directory: <code>./stop.sh hive</code>

Hadoop	From the management node, run the following command from the \$BIGINSIGHTS_HOME/bin directory: ./stop.sh Hadoop
--------	--

3.2.3 Checking status of the components

On the aafnode	Do the following
All components	From the management node, run the following command from the \$BIGINSIGHTS_HOME/bin directory: \$BIGINSIGHTS_HOME/bin/healthcheck.sh all \$BIGINSIGHTS_HOME/bin/status.sh all
BigSQL	From the management node, run the following command from the \$BIGINSIGHTS_HOME/bin directory: \$BIGINSIGHTS_HOME/bin/healthcheck.sh bigsql \$BIGINSIGHTS_HOME/bin/status.sh bigsql
Hive	From the management node, run the following command from the \$BIGINSIGHTS_HOME/bin directory: \$BIGINSIGHTS_HOME/bin/healthcheck.sh hive \$BIGINSIGHTS_HOME/bin/status.sh hive
Hadoop	From the management node, run the following command from the \$BIGINSIGHTS_HOME/bin directory: \$BIGINSIGHTS_HOME/bin/healthcheck.sh hadoop \$BIGINSIGHTS_HOME/bin/status.sh hadoop

3.2.4 Viewing the components log files

Log files are located at \$BIGINSIGHTS_VAR/bigsql/logs/

Use the log identifier to find the log information. For example, the bigsql-sched.log file is generally located in the /var/ibm/bigsql/logs directory.

If an error occurs check the query logs located at /opt/tnf/apps/dataset-common/log/tools.log

3.3 Managing Streams

For information on managing Streams in the context of the Customer Insights solution, refer to the AAF documentation, referenced in Section 4.1. For information on using the Infosphere Streams streamtool refer to [Creating and Managing Instances](#).

3.3.1 Starting

On the aafnode node	Do the following
Creating an instance	Sample command for a single instance (repeated for each): <pre>streamtool mkinstance -i rapidanalytics_instance_collector -embeddedzk</pre>
Start and stop one by one instance above	<pre>streamtool startinstance -i <instance name> -- embeddedzk streamtool stopinstance -i <instance name> -- embeddedzk</pre>
Configuring probes.csv, probes-tdr.csv, collector_local.xml and corresponding ini files	<ol style="list-style-type: none"> 1. Go to /opt/tnf/cna/cfg 2. Update probes.csv 3. Update probes-tdr.csv 4. Update collector_local.xml
Monitoring loading process	Data loading log files are placed in the folder: /opt/tnf/cna/stat/<dataset described in the probes.csv or probes-tdr.csv>/loader_script <pre>tail -f /opt/tnf/cna/stat/gn/loader_script/loader.log tail -f /opt/tnf/cna/stat/gn/loader_script/dataloader.log</pre>

3.3.2 Stopping

On the aafnode node	Do the following
Stopping instances	Stopping instances

3.3.3 Checking

On the aafnode node	Do the following
Checking instances	Monitoring instances

3.3.4 Viewing log files

On the aafnode node	Do the following
Checking log files	Logging and tracing
Viewing log files	Go to /tmp/logs

[Link to Infosphere Streams 4.1.1 Knowledge Center](#)

3.4 Managing the Analytics platform

For information on managing the analytics platform, refer to the analytics platform documentation referenced in Section 4.1.

3.4.1 Managing IBM Open Platform / Big Insights

Use the Ambari Server console to manage the base architecture components and check that all services are running:

```
<host>:8080 - username=##### password=#####
```

3.4.2 Starting

On the aafnode	Do the following
Running ETL scripts	<ol style="list-style-type: none"> 1. ETL scripts extracts data from hive and place it in a data node (need a hive instance) 2. Edit the properties files 3. Copy db2jcc.jar and db2jcc_license_cu.jar to /usr/iop/4.1.0.0/sqoop/lib and /usr/iop/4.1.0.0/hive/lib <pre>select count(*) from tnf. subscriber_crm subscriber_billing subscriber_care churn_data subscriber_level_cx_score_weekly cgr_device</pre> 4. Ensure there is data in these tables or the churn ETL script will fail 5. Go to /opt/IBM/CSP_CustomerInsight_DBLoader

	<pre>su boss</pre> <p>6. Run the script/</p> <pre>encodePasswordProperties.sh -f conf/churn_config.properties</pre> <p>7. Run the db export</p>
--	---

3.4.3 Checking

On the aafnode node	Do the following
	<p>To manage IOP/BI, use the Ambari Server console at <host>:8080 - username=##### password=#####</p> <p>In general, you'll just want to check that all services are running.</p>

3.4.4 Viewing log files

On the aafnode node	Do the following
	<pre>../bigsql/logs/</pre> <p>Installed services log file paths</p>

3.5 Managing Postgres, Presto, and Prestogres components

Postgres is an object-relational database management system with an emphasis on extensibility and standards-compliance. As a database server, its primary function is to store data securely, and to allow for retrieval at the request of other software applications.

Presto is an open source distributed SQL query engine for running interactive analytic queries against data sources

Prestogres is an open source middleware that allows the user to use Postgres ODBC/JDBC drivers to connect to Presto.

3.5.1 Starting the components

On the aafnode	Do the following
Postgres	<pre>/opt/tnf/apps/bis-main/bis-postgresql/scripts/ctl.sh start</pre>
Presto	<pre>/opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh start</pre>
Prestogres	<pre>/opt/tnf/apps/bis-main/bis-prestogres/scripts/ctl.sh start</pre>

3.5.2 Stopping the components

On the aafnode	Do the following
Postgres	<code>/opt/tnf/apps/bis-main/bis-postgresql/scripts/ctl.sh stop</code>
Presto	<code>/opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh stop</code>
Prestogres	<code>/opt/tnf/apps/bis-main/bis-prestogres/scripts/ctl.sh stop</code>

3.5.3 Checking the component status

On the aafnode	Do the following
Postgres	<code>/opt/tnf/apps/bis-main/bis-postgresql/scripts/ctl.sh status</code>
Presto	<code>/opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh status</code> <p>You can confirm that the installation of the user-defined functions has been successful using the Presto command line utility:</p> <pre>presto presto> use tnf; presto:tnf> select formatDate(DATE '2016-06-22');</pre>
Prestogres	<code>/opt/tnf/apps/bis-main/bis-prestogres/scripts/ctl.sh status</code>

3.5.4 Viewing log files

On the aafnode	Do the following
Postgres	Run the command: <code>tail -f postgresql-*.log</code>
Presto	<code>/opt/tnf/apps/bis-main/bis-prestogres/data/pg_log/</code>
Prestogres	<code>/opt/tnf/apps/bis-main/bis-prestogres/data/pg_log/</code>

3.5.5 Updating and maintaining configuration properties

On the aafnode node	Do the following
<p>To update and maintain configuration properties for applications such as BigSQL, Postgres, Prestogres, go to</p> <pre>/opt/tnf/apps/bis-main-var/bis-demon/cfg</pre>	<p>Review the following files:</p> <ul style="list-style-type: none"> • <code>bigsql_connection.properties</code> • <code>cfg.properties</code> • <code>pg_connection.properties</code> • <code>presto_connection.properties</code>

3.6 Managing Hive, Hadoop and BigSQL

3.6.1 Starting

On the aafnode node	Do the following
Hive	<p>Select the Hive service from the Ambari web interface.</p> <p>Click Service Actions > Start</p>
Hadoop	<p>Select the Hadoop service from the Ambari web interface.</p> <p>Click Service Actions > Start</p>

3.6.2 Stopping

On the aafnode node	Do the following
Hive	<p>Select the Hive service from the Ambari web interface.</p> <p>Click Service Actions > Stop</p>
Hadoop	<p>Select the Hadoop service from the Ambari web interface</p> <p>Click Service Actions > Stop</p>

3.6.3 Checking

The BigInsights web console provides a quick view into the status of the cluster. When you access the web console and select the Cluster Status tab, the cluster status view is displayed.

Another verification option is to run the health check utility that comes with BigInsights. By default, this tool runs through each installed component and checks to see if it is working properly on each node in your cluster. Run the `healthcheck.sh` script at:
`/opt/ibm/biginsights/bin`

3.6.4 Viewing log files

[Installed services log file paths](#)

On the aafnode node	Do the following
Hive	<code>cd /var/log/hive</code>
Hadoop	<code>cd /var/log/hadoop</code>
Big Insights	<code>cd /var/log/*</code>

4 Appendices

4.1 Related information

IBM BigInsights 4.1 documentation
IBM Cognos Analytics documentation
IBM Infosphere Streams documentation
IBM Predictive Customer Intelligence documentation
IBM Now Factory Analytics documentation (AAF documentation – internal wiki)