

IBM Telecom Analytics Solutions
Last updated: 2016-08-18

*Analytics Accelerator Framework for
Wireless Communication Service
Providers Solution Guide*



Note

Before using this information and the product it supports, read the information in "Notices" on page 33.

Product Information

This document applies to IBM Telecom Analytics Solutions Version 1.0.4 and may also apply to subsequent releases.

Licensed Materials - Property of IBM

© **Copyright IBM Corporation 2016.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	v
Chapter 1. Overview	1
Chapter 2. Installation and configuration.	3
Contents of the solution package	3
Overview of the installation and configuration process	3
Installing the Analytics Platform and Mediation node.	5
Installing the Solution Installer	5
Prerequisites	6
Modifying the sudoers file for the Linux user who runs the installation	6
Configuring computer firewall settings required for the installation on Linux.	7
Installing the Solution Installer software	8
Installing the solution content	9
Prerequisite.	9
Deploying the installation files on the nodes	10
Installing solution datasets	11
Deploying reports and content on the Business Intelligence node	12
Creating a data source connection to the BigInsights database	13
Connecting a data source connection to Hadoop	14
Post-installation steps	16
Chapter 3. Using Analytics Accelerator Framework for Wireless Communication	
Service Providers	19
Using the Customer Experience Dashboard.	19
Appendix A. Accessibility features.	23
Appendix B. The data provision tool	25
Appendix C. Accessing the SLM log file	27
Appendix D. Troubleshooting	29
Troubleshooting a problem	29
Troubleshooting resources	29
After the Solution Installer is run, Chef processes are still running after the cleanup.sh script is executed	31
Notices	33
Index	37

Introduction

The Analytics Accelerator Framework for Wireless Communication Service Providers (AAF) solution enables Communication Service Providers to integrate customer-centric network data into their big data initiatives. AAF delivers an open and flexible architecture with adapter-based data ingestion, Hadoop-based data storage, and API-based data access. In addition, AAF bundles a comprehensive toolset for the creation of dashboards, reports, and visualizations.

Audience

This guide is intended for administrators and users of the Analytics Accelerator Framework for Wireless Communication Service Providers solution. It provides information on installation, configuration, and use of the Analytics Accelerator Framework for Wireless Communication Service Providers solution.

Finding information

To find product documentation on the web, including all translated documentation, access IBM® Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products. Some of the components included in the Analytics Accelerator Framework for Wireless Communication Service Providers have accessibility features. For more information, see Appendix A, “Accessibility features,” on page 23.

Analytics Accelerator Framework for Wireless Communication Service Providers HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

Samples disclaimer

Sample files may contain fictional data manually or machine generated, factual data compiled from academic or public sources, or data used with permission of the copyright holder, for use as sample data to develop sample applications. Product names referenced may be the trademarks of their respective owners. Unauthorized duplication is prohibited.

Chapter 1. Overview

IBM Analytics Accelerator Framework for Wireless Communication Service Providers (AAF) provides the technologies, tools, and models to deliver best-in-class, customer-centric, network-based telecommunication datasets into an open source Hadoop-based data store. AAF also provides a comprehensive set of best-in-class analytical technologies to meet Communication Service Providers (CSP) analytical requirements.

Important: Before installing, the AAF IBM Telecom Analytics Solution must be installed.

What are the benefits of AAF?

The solution offers the following benefits:

- Makes customer-centric network data available across CSP organizations.
- Empowers all departments in CSP organizations with the capability to use advanced analytics to drive improved overall business outcomes.
- Integrates with network planning to give smarter prioritization of investment, based on a customer view of service quality.
- Achieves a faster time-to-value than comparable approaches, based on preintegrated capabilities.
- Integrates with network operations to enable a customer-centric prioritization of network issue resolution.

What are the features of AAF?

The solution provides the following features:

- Adapter hooks that provide a proven mechanism to ingest data from both IBM and third-party sources.
- Open APIs that enable CSPs to expose the data to applications across their entire organization.
- Built-in dashboard and visualizations that give a clear starting point and create value from the toolset included with AAF.
- Hadoop compatibility that allows CSPs that have already invested in Hadoop infrastructure to continue to use it with AAF, so no additional investment is necessary.
- Complex event processing capabilities delivered by IBM WebSphere Streams that translate raw network data into widely usable metrics. This gives CSPs deep insights into network usage.

Chapter 2. Installation and configuration

The IBM Analytics Accelerator Framework for Wireless Communication Service Providers solution is for use with IBM Predictive Customer Intelligence.

Analytics Accelerator Framework for Wireless Communication Service Providers consists of the following nodes:

- An Analytics Platform node.
- A Mediation node.
- A Business Intelligence node.

Contents of the solution package

The IBM Analytics Accelerator Framework for Wireless Communication Service Providers solution package contains several parts.

The following parts are included in the solution package:

- Solution Installer. This puts the installation packages on the nodes required for a particular deployment architecture.
- IBM Cognos® Business Intelligence (BI) reports and dashboards.
- The Analytics Platform contains the software for data management and ETL (extract, transfer, and load) and will be installed on the Analytics Platform and Streams nodes.

Overview of the installation and configuration process

This is an overview of the general tasks that you must perform to install and configure the Analytics Accelerator Framework for Wireless Communication Service Providers solution.

The following diagram provides a high-level overview of the installation and configuration process.

Installation


Overview

In a few simple steps



1 DOWNLOAD PACKAGES

Download packages to the installation server and ensure network connectivity between all nodes in the deployment.



2 INSTALL/CONFIGURE THE NODES

Install the Analytics Platform data and Mediation node (or nodes)



Analytics Platform data nodes

Analytics Platform mgmt node




Mediation nodes

You may have multiple streams and hadoop nodes, including a minimum one hadoop management node per three data nodes

3 INSTALL THE INSTALLER


Install the Solutions Installer and configure it to identify the computers where the parts of the solution will reside.



The installer copies the required installation files for the specified servers.

4 INSTALL/DEPLOY CONTENT

Deploy installation files and content on the nodes, as well as set up tools, reports, and connections.



DB2 node

BI node

Predictive Analytics node

Perform the following tasks:

1. Plan where you want the various parts of the solution to reside. For example, you must decide which computers to use for the Analytics Platform data node and the Analytics Platform streams node.

2. Install and configure IBM Analytics Accelerator Framework for Communication Service Providers and the prerequisite software, including PCI 1.1.1.
3. Load data onto the Analytics Accelerator Framework Platform and verify this data.
4. Run the Solution Installer to deploy the installation files to the relevant nodes at the application layer.
5. Install the Solution datasets.
6. Install the TELCO database.
7. Create the TELCO datasource.
8. Install IBM Cognos reports.
9. Install IBM Cognos Dashboards. (Do not include framework steps, as they are not required.)
10. Confirm that the reports and dashboards are connecting to BigInsights and reading the data correctly.

Installing the Analytics Platform and Mediation node

The Analytics Platform is software that is installed on your Analytics Platform nodes. Mediation software is deployed on the mediation node or nodes in your deployment, manages the loading and aggregation of your telecommunication (telco) data into the Apache Hadoop ecosystem, and works in conjunction with the Analytics Platform.

To install and configure the Analytics Platform and mediation node (or nodes), refer to the following documentation available from http://www-01.ibm.com/support/docview.wss?uid=swg27048488#tas_v1r0m0en:

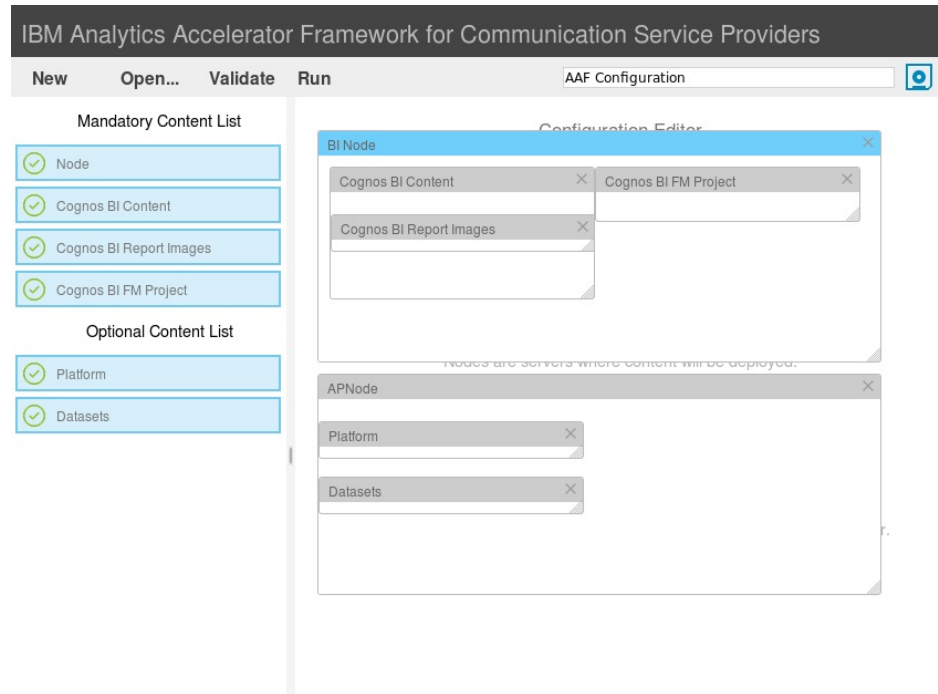
- *Analytics Platform 3.1 Installation and Configuration Guide*
- *CNA 9.1 Mediation Operations Guide*
- *Customer and Network Analytics 9.0 Provisioning Guide*

Installing the Solution Installer

The Solution Installer is used to configure and deploy solution content on the PCI-based nodes in your deployment, including the BI node and the Predictive Analytics node. To use the installer, you must first perform the prerequisites and then the steps for installing the Solution Installer.

Note: The Solution Installer referred to in this section is the Solution Installer for the Analytics Accelerator Framework for Wireless Communication Service Providers solution and not the PCI solution.

The configuration for the installer is shown on following screen:



Prerequisites

Before you install the Solution Installer for the Analytics Accelerator Framework for Wireless Communication Service Providers solution, there are prerequisites that must be met.

Ensure the following conditions are met before you install the solution:

- You have a Red Hat Enterprise Linux operating system that you can install to. The Solution Installer works only on that version of Linux.
- You understand your deployment environment. There are a combination of nodes that you must install to, so you need to determine where you want the various components of the solution to reside.
- You have modified the sudoers file for the user who runs the installation.
- You have configured your computer firewall settings for the installation of Linux.

Modifying the sudoers file for the Linux user who runs the installation

To run the Analytics Accelerator Framework for Wireless Communication Service Providers installation in a Linux environment, you must be either root user or have sudo permission on each node computer, and you must disable requiretty during the installation.

To install with sudo user permissions, you must add the user to the sudoers file.

Procedure

1. Log in as root user.
2. Enter the following command to open the sudoers file for editing:
`visudo -f /etc/sudoers`
3. Locate the following line:

- Defaults requiretty
4. Press the I key to insert text.
 5. Type a number sign (#) in front of Defaults requiretty to comment out the line. For example, the line will appear as #Defaults requiretty
 6. If you run the installer as a user with sudo user permissions, go to the end of the file, and add the following line for your user:
`username ALL=(ALL) NOPASSWD: ALL`
 7. Press the Esc key, and enter :wq to save and close the file.
 8. Repeat these steps on each computer on which you install a Analytics Accelerator Framework for Wireless Communication Service Providers node.

Configuring computer firewall settings required for the installation on Linux

The Solution Installer requires bidirectional communication with each node when you install a server component or server artifact component, and your firewall can interfere with this communication.

To ensure bidirectional communication, you must configure your firewall settings for the installation. Alternatively, you can temporarily disable your firewall during the installation.

A script is provided to open all necessary ports on the installation node. You can either use the `firewall.sh` script, open the ports manually, or disable the firewall settings.

In addition to opening ports on the installation node computer, you must open some ports on each computer to which content from the solution will be copied. You can either disable all firewall settings temporarily or use the `firewall.sh` script, as described in this document.

Disabling firewall settings prior to installation with the Solution Installer (optional)

If you choose to temporarily disable your computer's firewall during the installation of the Analytics Accelerator Framework for Wireless Communication Service Providers solution, ensure that you enable the firewall on each node after the installation is complete.

About this task

Important: If you disable your computer's firewall during the installation, you must still ensure that the required ports are opened on each node computer when you restart your firewall.

Procedure

1. Log on to the computer as the root user or as a user with sudo permissions.
2. Type the following command:
`/etc/init.d/iptables stop`
`/etc/init.d/iptables save`
3. Repeat these steps on each computer on which you install a node and on the server from which you are running the installation.

What to do next

After the installation is complete, on each node, type `/etc/init.d/iptables start` to restart your firewall.

Manually configuring firewall settings on the installation node during installation with the Solution Installer

Disable your firewall settings using the `firewall.sh` script. The script must be run on the computer used by the Solution Installer and each node that will contain content from the solution. You can run the `firewall.sh` script to open the ports on a Linux installation node.

About this task

The `firewall.sh` script opens the following ports on the Solution Installer node:

- 8080 incoming
- 443 incoming
- 9683 incoming
- 22 outgoing

Procedure

1. As the root user or as a user with sudo permissions, log on to the computer that contains the solution installer node.
2. Back up your existing firewall settings by typing the following command:
`/etc/init.d/iptables save.`
3. Run the `firewall.sh` script.

Installing the Solution Installer software

Before you use the Solution Installer to configure and deploy the nodes in a Linux environment, you must install the Solution Installer.

Before you begin

The working directory for solution installation files is `opt/IBM`

About this task

Here is a short video that shows you how to install the Solution Installer:

<http://youtu.be/9u6ZOk8iyt4>

Please note the version numbers shown in the video may have been updated (the video will still be relevant). The steps that are shown in the video are described in the following procedure.

Procedure

1. Log on to the computer where you downloaded the solution package.
2. Decompress the solution package.
3. If the Solution Installer (including the IBM Predictive Customer Intelligence Solution Installer) has been run on the computer previously, run the following command on each of the nodes:
`./cleanupClient.sh`
4. Go to the Solution Installer directory in the following location: .

5. If you haven't already opened the ports that are required by the Solution Installer, you can open them by running the following command:

```
./firewall.sh
```

6. Enter the following command:

```
sh setup.sh username first_name last_name email password
```

This command creates a user (with the details supplied) for accessing the Solution Installer web server.

The web server is started and the URL for the Solution Installer displays in the command line window.

7. A browser window might open automatically. If it doesn't, copy and paste the URL into a web browser and bookmark it.

What to do next

Use the Solution Installer to deploy the installation files to the nodes in your solution.

Installing the solution content

Installing the solution content involves deploying installation files and content on the nodes, installing solution datasets, installing Cognos user-defined functions, deploying reports, and creating a connection to BigInsights.

Prerequisite

Before you install the Analytics Accelerator Framework for Wireless Communication Service Providers solution content, there is a prerequisite that must be met.

Ensure that IBM Predictive Customer Intelligence version 1.1.1 is installed and configured and you have applied all known fix packs to it.

Note: When installing PCI, you need an additional temporary server from which to run the PCI Solution Installer. See also this Technote: <http://www-01.ibm.com/support/docview.wss?uid=swg21612290>

Installing IBM Predictive Customer Intelligence version 1.1.1

To install PCI version 1.1.1, refer to the IBM Predictive Customer Intelligence *Installation Guide* 1.1.1 on IBM Knowledge Center.

Ensure that you adhere to the following sections in this guide:

- Prerequisites and installation preparation
- Solution architecture
- Supported operating systems and hardware requirements
- Prerequisites for using the solution installer
- Install IBM Predictive Customer Intelligence by using the solution installer
- Manually install IBM Predictive Customer Intelligence
- Troubleshooting (if necessary)

Install PCI on the designated PCI nodes in your deployment.

Note: The CI SPSS models require the modeler server to be installed at /usr/IBM/SPSS/ModelerServer/17.1/

Deploying the installation files on the nodes

You use the Solution Installer to deploy the installation files to the appropriate nodes in your solution.

Before you begin

Ensure that you know the URL for accessing the Solution Installer.

Ensure that you know the following information for each of the nodes in your solution:

- Host name or IP address
- A user ID and password that has write permissions to the /opt/ directory.

Each node must also meet the following requirements:

- Has a valid DNS name that is resolvable.
- Does not have Chef server or client software installed on it.

About this task

In this task, you create an installation configuration of your solution by dragging objects onto a canvas, adding components to the objects, and defining properties for the components.

Important: The Solution Installer runs on the latest versions of Chrome and Firefox. It does not run on Internet Explorer and it is untested on Safari.

Procedure

1. In a web browser, go to the Solution Installer.
2. Read the license agreement and click **Accept**.
3. Click **New Configuration** and then click **OK**.
4. From the **Mandatory Content List** pane, select **Node**.
A visual representation of an empty node is displayed on the canvas.
5. In the **Property Editor**, type a name for the node, and in the **Host Name** box, specify the server on which the node will be placed.
For example, you might create a node that is named Cognos Analytics node (or Cognos BI node).
6. Enter a username and password for a user that has sudo access.
If you click **Validate** before all the mandatory content is added, a message indicates that the node is not valid. For example, if you validate the Cognos Analytics node before adding Cognos Analytics (BI) content, then a message indicates that the node is not yet valid.
7. From the **Optional Content** list, drag any content that you want to add into the appropriate node. The reason that this content is optional is because you might already have some of the necessary components installed in your environment.
8. For every node that you want to add content to, define the node and then drag content objects onto it.

9. If you want to save the configuration to use or refer to in the future, click **Save Configuration**.
10. Click **Validate**. The progress of the validation is displayed. If the validation fails, look at the log file to troubleshoot the problem.
11. After all nodes are defined, the contents are specified, and each node is validated, click **Run**. The scripts and files that are used to install and configure the solution are copied to the target nodes.

What to do next

Complete the installation of the solution on each of the nodes.

Installing solution datasets

Use the following procedure to install your solution datasets.

Procedure

1. As the root or sudo user, start the installer from the directory where you uncompressed the installation package:

```
$ sudo ./analytics-platform/install_telsol.sh
```

The following flags may be passed to the script:

Flag	Meaning
--dont_validate_biginsights	The validation of IOP and BigInsights takes a few minutes to run. To save time, this validation step can be skipped by passing this flag.
-h	Prints a message describing the script usage and exits.

2. When installing each use case, the installer will check for an existing installation of that use case and ask if it should be replaced with the version from the installation package. Provide the answer that applies to your situation. For example,

```
-- Installing Customer Behaviours use case...
```

```
Checking for existing installation of Customer Behaviours use case...
```

```
WARNING: Previous installation of Customer Behaviours use case found: customer-behaviours -
Remove? (y/n):
```

Answering y will result in the removal of the existing use case. The question will be asked a second time to confirm removal. If the existing installation is removed, the installation of the new version will proceed automatically after the removal.

3. Some use cases require data provided by the CSP to be loaded (provisioned) into the database before they are run (for example, CRM data). This data should be provided in the form of CSV files. Before attempting to provision your data, the installer will pause to ensure that these files are available in the correct directory.

```
-- Loading provisioned data...
```

```
Provisioning data to be loaded must be stored as CSV files in
/opt/tnf/apps/bis-main-var/bisprovisioning-tool/csv_files/ -
```

if you have any data to provision, please copy it to that directory now and press any key to resume...

If the provisioned data CSV files have not already been copied to that directory, do it now and press any key to

resume the installer which will attempt to load them. It will determine which files should be loaded into which

tables based on the filenames and ask for confirmation that the correct files are being loaded to the correct

tables, e.g.,:

Table to load data to ----- Data file

```
cb_category ----- csv_files/cb_category_20150320.csv
cea_unacceptable_trend_config ----- csv_files/cea_unacceptable_trend_config_00000000.csv
nps_cell_provisioning_table ----- csv_files/nps_cell_provisioning_table_20150331.csv
nps_crm_provisioning_table ----- csv_files/nps_crm_provisioning_table_20150331.csv
nps_device_provisioning_table ----- csv_files/nps_device_provisioning_table_20150331.csv
nps_survey_table ----- csv_files/nps_survey_table_20150429.csv
performance_assign_config ----- csv_files/performance_assign_config_00000000.csv
Load data to tables according to the list? (y/n):
```

Review this information and provide the appropriate response. Upon successful installation, the following message is displayed:

```
-- IBM Telecom Solutions installed completed successfully!
```

The entire session is logged in install_telso1.log in the current directory.

Scheduling datasets

Datasets are run on a scheduled basis using tables that containing customer information.

Most datasets are run on an hourly, daily, or weekly basis, through cron jobs that are added to a boss's crontab as part of the dataset rpm installation. These can be viewed and modified by running crontab -e as a boss on your Analytic Server node. In a default install, here is the scheduling you will see across the various datasets:

Dataset	Schedule
customer-profile-data-setup	0.30, every Monday morning (weekly tables)
ott-applications	02.30, every Monday morning (weekly tables)
customer-profile	02.30, every Monday morning (weekly tables)
user-profile	04.30, every Monday morning (weekly tables)
customer-behaviours	<ul style="list-style-type: none"> • *.45, every hour (hourly tables) • 03.15, every day (daily weighted interest tables) • 03.30, every day (daily other tables) • 00.30, every Monday morning (weekly tables)

Deploying reports and content on the Business Intelligence node

You must deploy the report, dashboards, and images for the Analytics Accelerator Framework for Wireless Communication Service Providers solution.

Before you begin

You will need to create folders for the reports. Prior to installation of the reports, you must create an AAF folder in Team content and then create a reports folder in the AAF folder.

Procedure

1. Log on to the BI node as the **root** user or as a user with **sudo** permissions.
2. As a **root** user, copy AAF_customer-experience-images-1.0.4.zip from /opt/IBM/IS_CSP_Analytics_Accelerator_Framework_1.0.4/ and unzip the contents into /opt/ibm/cognos/analytics/webcontent/bi/samples/images/
3. As a **root** user, copy AAF-customer-experience-report-1.0.4.zip from /opt/IBM/IS_CSP_Analytics_Accelerator_Framework_1.0.4/ to /opt/ibm/cognos/analytics/deployment
4. On a web browser, navigate to <localhost>:9300/bi. This opens the IBM Cognos Analytics page.
5. From the menu on the left-hand side of the page, click **Manage** and then click **Administration Console**. This opens IBM Cognos Administration Console in a new tab.
6. Click **Configuration** and then click **Content Administration**. This opens a page view of the content to be imported and exported.
7. Click the **New Import** symbol from the menu on the right-hand side of the page. This launches the **New Import** wizard.
8. Select AAF-customer-experience-report-1.0.4.zip and click **Next**.
9. On the **Specify a name and description** page, accept the defaults and click **Next**.
10. On the **Select the public folders, directory and library content** page, select all the folders which appear and click **Next**.
11. On the **Specify the general options** page, accept the defaults and click **Next**.
12. On the **Review the summary** page, accept the defaults and click **Next**.
13. On the **Select an action** page, accept the defaults and click **Finish**.
14. On the **Run with options** page, accept the defaults and click **Run**.
15. On the final wizard page, click **OK**. The import should now be visible on the **Content Administration** page.

Creating a data source connection to the BigInsights database

You must create a data source connection from the Business Intelligence node to the IBM BigInsights® database that is on the Analytics Platform node.

About this task

For Cognos dashboards, BigInsights is used to provide access to the Hadoop datasets; for Cognos reports, both BigInsights and Presto can be used to access Hadoop. The latter provides a more efficient interface to Hadoop; however, it does require a separate version of the reports to be installed and a different set of post-installation steps.

Procedure

1. Open a web browser on the BI node.
2. Navigate to the **IBM Cognos Analytics** page, click **Manage**, and then click **Data servers**.

3. Click **Add a dataserver**.
4. Enter the connection information provided by your database administrator. If applicable, select the authentication mode to be used.
 - Type: JDBC
 - Name: BigDataDb2
 - Server: Analytics Platform Node
 - Port: 32051
 - User: bigsql
5. Select the BigSQL database and the TNF schema.
6. Click **Test** to test the connection.
7. Click **Save**.

Connecting a data source connection to Hadoop

This document details the steps required to install the version of Customer Experience Analytics (CEA) reports that use the Presto distributed query engine.

Presto (see also <https://prestodb.io/>) is an alternative to using BigSQL to access the data model for the CEA Cognos reports. To support Presto, there are changes to the framework model and the reports and some user-defined functions are used. These functions have to be installed in Presto. Connecting to Presto is done by using a Postgres ODBC data source; this, in turn, connects to a Prestogres daemon, which forwards the SQL requests to Presto.

The RPM contains the following files:

- `readme.txt`
- `AAF-customer-experience-model-presto-1.0.4.zip` - a zip file that contains the data model for the reports.
- `AAF-customer-experience-images-1.0.4.zip` and `AAF-customer-experience-report-presto-1.0.4.zip` - the CEA reports (built for Presto) and the set of images used by the reports.
- `lib/plugin/hive-hadoop2/cea-prestogres-tools.jar` - the classes that implement the Presto user-defined functions used by the reports.

Prerequisites

The following components are required as prerequisites:

- AAF version 3.1 (or later).
- Install Postgres on the AAF node (this may be included in the AAF packages and installed as part of that installation). If necessary, the packages can be obtained from the Postgres download site:

```
$ yum install http://yum.postgresql.org/9.4/redhat/rhel-6-x86_64/pgdg-redhat94-9.4-1.noarch.rpm
$ yum install postgresql94-server postgresql94-devel postgresql94
$ yum install postgresql94-contrib postgresql94-python postgresql94-plpython
```

- On the Cognos node, install the 32-bit Postgres ODBC driver (`postgresql-odbc-08.04.0200-1.e16.i686`). This can be downloaded from the Postgres download site. Note that the 32-bit driver is required; Cognos does not support 64-bit ODBC drivers. Once downloaded, this package should be installed using the **yum** not **rpm**) command as this will also install any dependent packages:

```
$ yum install postgresql-odbc-08.04.0200-1.e16.i686.rpm
```

Installing the Presto version of the reports

The CEA reports for Presto are delivered as an RPM. Install using the following:

```
$ rpm -ivh cea-prestogres-deployment-1.0.4.rpm
```

Note: This only writes the installation artifacts to a subdirectory `cea-reports-presto` in `bis-main`. Once the package has been installed, the following post-installation steps should be executed.

Post-installation

Execute the following steps:

1. On the AAF node, you should ensure that Postgres is started and running. If it needs to be restarted, you should execute:

```
$ /opt/tnf/apps/bis-main/bis-postgres/scripts/ctl.sh status
$ /opt/tnf/apps/bis-main/bis-postgres/scripts/ctl.sh start
```

You should also ensure that Presto is started and running. If it needs to be restarted, you should execute:

```
$ /opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh status
$ /opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh start
```

2. You may need to initialize the Prestogres connection by executing the following:

```
$ /opt/tnf/apps/bis-main/bis-prestogres/scripts/ctl.sh init
```
3. Install Presto user-defined functions. The user-defined functions are installed on the AAF node. Copy the jar file containing the functions to the Presto plugin directory and restart Presto. The following should be executed as the **boss** user:

```
$ cp lib/plugin/hive-hadoop2/cea-prestogres-tools.jar
$ /opt/tnf/apps/bis-main/bis-presto/plugin/hive-hadoop2
$ /opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh stop
$ /opt/tnf/apps/bis-main/bis-presto/bin/ctl.sh start
```

You can confirm successful installation of the user-defined functions by using the Presto command line utility: `$ presto presto> use tnf; presto:tnf> select formatDate(DATE '2016-06-22');`

4. Configure the ODBC driver on the Cognos node before the reports are imported into Cognos. This step should be executed on the node hosting the Cognos server. Before executing this step, you should confirm the server node and port for the Presto service. The server node is the AAF node, and the port is obtained from the port entry in the Prestogres configuration file `/opt/tnf/apps/bis-main/bis-prestogres/conf/prestogres.conf` on the AAT node. The following (or similar) entry should be made in the file:

`/etc/odbc.ini:`

```
[tnf]
Description=PostgreSQL
Driver=/usr/lib/psqlodbcw.so
Database=default
Servername=AAFNodeName
Port=PrestogresPortNumber
UserName=prestogres
Password= Protocol=7.4
RowVersioning=No
ShowSystemTables=No
UseServerSidePrepare=No
ConnSettings=
```

5. Install the reports by importing the reports zip file using the Cognos Administration user interface. Ensure that the images in the associated zip file

were previously unpacked into the samples directory. This step may not be necessary if the images have already been installed for the standard reports:

```
$ cd /opt/ibm/cognos/analytics/webcontent/bi/samples/images
$ unzip /tmp/AAF-customer-experience-images-1.0.4.zip
```

When installing the reports, default values should be used throughout the dialogue. When complete, you should create an ODBC data source connection, presto, that references the ODBC data source configured in step 2 (tnf in the example). This should be saved and tested. Restart the Cognos server when the installation completes; you may need to set the library paths and Java home directories to include the 32-bit libraries required by Cognos:

```
$ export JAVA_HOME=/opt/ibm/cognos/analytics/jre
$ export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/lib32/
$ /opt/ibm/cognos/analytics/bin64/cogconfig.sh -stop
$ /opt/ibm/cognos/analytics/bin64/cogconfig.sh -s
```

Troubleshooting

Presto logs activity to the directory /opt/tnf/apps/bis-main/bis-prestogres/data/pg_log/

You should review these logs to investigate any issues.

Post-installation steps

For Rapid Analytics 3.0, the AAF reports have been upgraded to use the partitioning defined in the underlying HADOOP datastore. To do this, the Cognos reports have been modified so that the filters include conditions to select the partition in which an object resides. Consequently, some user-defined functions (UDFs) must be added to the BigSQL interface that the reports use to access the data. This topics outlines the post-install steps used to set up these UDFs.

About this task

Three UDFs are defined: TNF.SGM(), TNF.FORMAT(), and TNF.FORMAT_WITH_HOUR(). These require the installation of a Java archive (JAR) containing the Java code and an SQL file containing the signatures of the functions. All of these artifacts are created during the installation of the RPM.

The installation of the RPM creates the following files in the directory:

```
/opt/tnf/apps/cea-cognos-reports:
./bigsql/lib/cea-cognos-reports*.jar
./bigsql/scripts/install.sh
./bigsql/scripts/install.sql
```

The install script, install.sh, uses DB2 to install the UDF signatures. This should be executed as the user **bigsql**.

Procedure

1. Install the UDFs as follows:

```
su bigsql
cd /opt/tnf/apps/cea-cognos-reports/bigsql/scripts/install.sh
sh install.sh
```

Note:

Included in the package is a log4j configuration file, `CommonSgm.properties`. If this is copied to the `/tmp` directory (or BigSQL's home directory), the UDF's activities will be logged in `/tmp/cea-cognos-reports-udf.log`. The properties file can be configured to suit the particular environment.

2. To confirm successful installation, execute the following:

```
$ jqsh
Welcome to JSqsh 4.4
Type "\help" for help topics. Using JLine.
1> \connect tnf 50
Password: *****
[localhost][bigsql] 1> select tnf.sgm(IMSI) from
tnf.SUBSCRIBER_LEVEL_CX_SCORE_WEEKLY;
```

Chapter 3. Using Analytics Accelerator Framework for Wireless Communication Service Providers

This section describes a typical user case that shows how to use the Analytics Accelerator Framework for Wireless Communication Service Providers solution to understand customer experience trends.

Workflow and user roles

This workflow shows how the solution can be used to thoroughly provide a dashboard and key metrics to a communication service provider.

For this example, a customer experience team member is interested in understanding customer experience trends. This person has also been asked to investigate a specific customer complaint. A typical process workflow is described

Using the Customer Experience Dashboard

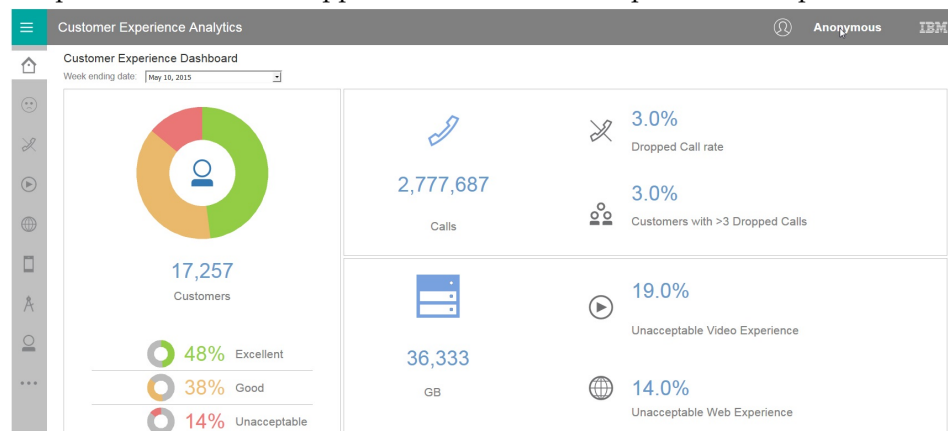
The Analytics Accelerator Framework for Wireless Communication Service Providers solution provides in-depth analytics to help communication service providers evaluate customer experience trends and metrics.


About this task

This task describes how a typical customer experience team member would use the dashboard to evaluate customer trends and metrics, as well as investigate a specific customer complaint. The screen shots shown are for example purposes and will vary depending on your particular implementation.

Procedure

1. Open the **Customer Experience Dashboard** for the past week. The metrics for the past week include dropped call rate and unacceptable web experience.

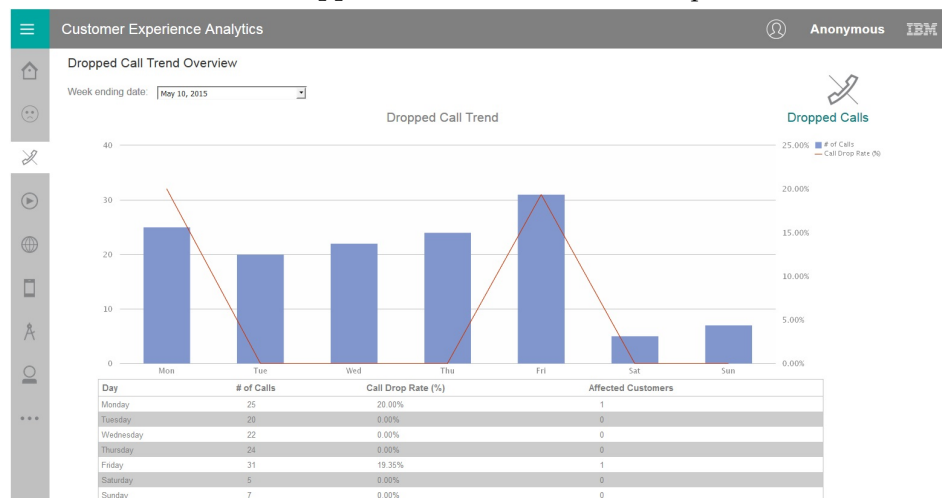


- Click  to view the **Unacceptable Experience Trend** report.





The communications service provider has an SLA of 4.5%. The customer experience team member investigates why the SLA went above that level on Monday and Friday. The team member looks at the reports in the menu on the left-hand side of the screen.


- Click  to view the **Dropped Calls Trend Overview** report.



The bar chart shows the number of dropped calls and the line chart shows a corresponding dropped call rate. The report below the chart includes the number of affected customers. The dropped call rate on Monday and Friday was particularly high and may warrant further investigation.


- Click through the  (**Video Experience**) and  (**Web Experience**) reports on the left-hand menu.

- Click  to view the **Device Data Usage Distribution** report.



The bar chart shows the total volume and average volume per user of the top ten devices by average volume per user. The report below the chart includes the number of unique users. The average volume per user of Nokia Lumia 920 and BlackBerry Q10 users is particularly high, and the total volume of Apple iPhone 4 users is much higher than any other device.

- Click  to view the **Application Data Usage Distribution** report.



The bar chart shows the total volume and average volume per user of the top ten applications by average volume per user. The report below the chart includes the number of unique users. The average volume per user of Sopcast is particularly high, and the total volume of Netflix Video users is higher than any other device.

- Click  to view the **Single Customer Experience Analysis** report.



The report provides a high level overview of a subscriber's activity.



The experience section of the report can be filtered to show only a specific type of experience, such as **Video Experience**.



The cells section of the report can show either **Top Cells by Call** (for voice) or **Top Cells by Volume** (for data).

Appendix A. Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products.

For information about the commitment that IBM has to accessibility, see the IBM Accessibility Center (www.ibm.com/able).

IBM Cognos HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Report output

In IBM Cognos Administration, you can enable system-wide settings to create accessible report output. For more information, see the *IBM Cognos Analytics Administration and Security Guide*. In IBM Cognos Report Studio, you can enable settings to create accessible output for individual reports. For more information, see the *IBM Cognos Report Studio User Guide*. You can access the previously mentioned documents at IBM Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

Appendix B. The data provision tool

The data provision tool loads, deletes, and exports data from the provision tables. The provision tables contain customer relationship management data.

The data provision tool is installed in the `/opt/tnf/apps/bis-main-var/bis-provisioning-tool` folder and consists of the following parts:

load.sh

Loads and updates provision tables.

load-all.sh

Checks all the files in the `.csv` folder and loads them at the same time.

delete.sh

Reads records from input files and deletes records from IBM InfoSphere® BigInsights by using key values.

export.sh

Exports records from IBM InfoSphere BigInsights to text files.

Usage

You must specify key values for all operations. The key field names are obtained by the provision tools from the table metadata. The primary key must be defined in table DDL. In the following examples the primary key field is **tac**:

```
create table device (  
tac varchar(32) primary key, model varchar(64),  
vendor varchar (64)  
);
```

The follow example contains two primary keys:

```
create table device (  
tac varchar(32),  
model varchar(64), vendor varchar (64), primary key (tac, model)  
);
```

To see the parameters for `load.sh` type, the script name with no parameters in a command-line window:

```
Usage: load.sh [options] Options:  
* --file, -f  
CSV file with data * --table, -t  
DB Table name
```

--file, -f

A list of input files that are separated by the pipe (|) character.

--table, -t

The table name. The table name can include a dot separated schema. For example, `load.sh -f device.csv -t device`.

`delete.sh` and `export.sh` take the same parameters.

The data provision tool reads table metadata and converts the input text file into an Apache Parquet or plain text formatted file and populates specified table with the data. It updates existing records and inserts new records.

Creating dictionary tables

Run the `/opt/tnf/apps/bis-main/bis-dbconfig/install-db.sh` script with table group name parameters. The script `/opt/tnf/apps/bis-main/bis-main.sh` creates a dictionary table in the database when it starts.

You cannot modify the data in IBM InfoSphere BigInsights directly. The data provision tool updates the database tables and exports them to IBM InfoSphere BigInsights. During the export stage, the data provision tool must know the storage format. Apache Parquet and plain text formats are supported. Update the storage format in the `/opt/tnf/apps/bis-main/bis-tools/bisprovisioning-tool/cfg/hive-storage.properties` file.

Input files

The data input format is a pipe (|) separated utf-8 encoded plain text file with headers. The header can start with a number sign (#). Header values must match the dictionary table case insensitive field names. For example,

```
TAC|MODEL|VENDOR  
12345678|Zamyang|Ciaobao
```

Appendix C. Accessing the SLM log file

To determine the number of customers in the database, refer to the SLM log file.

The xml files for SLM tagging are generated in the /opt/slm/*.slmtag folder. Information is added until the log file reaches a certain size; then a new log file is created.

Output file location

```
/opt/tnf/apps/ca-slm/log  
73a425370ffb0673c98dde1019e2a74c.slmtag
```

Format of the output

```
<SchemaVersion>2.1.1</SchemaVersion>  
<SoftwareIdentity>  
  <PersistentId>50ae811f644c4a6ea2b28bbb78a6e4a7</PersistentId>  
  <Name>IBM Applied Analytics Framework for Communications Service Providers</Name>  
  <InstanceId>/opt/tnf/apps/ca-slm/</InstanceId>  
</SoftwareIdentity>  
<Metric logTime="2016-07-20T00:01:06+01:00">  
  <Type>RESOURCE_VALUE_UNIT</Type>  
  <SubType>SUBSCRIBER</SubType>  
  <Value>250000</Value>  
  <Period>  
    <StartTime>2016-07-19T00:01:06+01:00</StartTime>  
    <EndTime>2016-07-20T00:01:06+01:00</EndTime>  
  </Period>  
</Metric>  
<Metric logTime="2016-07-21T00:01:03+01:00">  
  <Type>RESOURCE_VALUE_UNIT</Type>  
  <SubType>SUBSCRIBER</SubType>  
  <Value>250000</Value>  
  <Period>  
    <StartTime>2016-07-20T00:01:02+01:00</StartTime>  
    <EndTime>2016-07-21T00:01:02+01:00</EndTime>  
  </Period>  
</Metric>
```

Appendix D. Troubleshooting

This section contains troubleshooting information for the Analytics Accelerator Framework for Wireless Communication Service Providers installation.

Troubleshooting a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem.

Review the following table to help you or customer support resolve a problem.

Table 1. Troubleshooting actions and descriptions

Actions	Description
A product fix might be available to resolve your problem.	Apply all known fix packs, or service levels, or program temporary fixes (PTF).
Look up error messages by selecting the product from the IBM Support Portal, and then typing the error message code into the Search support box (http://www.ibm.com/support/entry/portal/).	Error messages give important information to help you identify the component that is causing the problem.
Reproduce the problem to ensure that it is not just a simple error.	If samples are available with the product, you might try to reproduce the problem by using the sample data.
Ensure that the installation successfully finished.	The installation location must contain the appropriate file structure and the file permissions. For example, if the product requires write access to log files, ensure that the directory has the correct permission.
Review all relevant documentation, including release notes, technotes, and proven practices documentation.	Search the IBM Knowledge Center to determine whether your problem is known, has a workaround, or if it is already resolved and documented.
Review recent changes in your computing environment.	Sometimes installing new software might cause compatibility issues.

If the items in the table did not guide you to a resolution, you might need to collect diagnostic data. This data is necessary for an IBM technical-support representative to effectively troubleshoot and assist you in resolving the problem.

Troubleshooting resources

Troubleshooting resources are sources of information that can help you resolve a problem that you are having with an IBM product.

Support Portal

The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services.

The IBM Support Portal lets you access all the IBM support resources from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution. Familiarize yourself with the IBM Support Portal by viewing the demo videos (https://www.ibm.com/blogs/SPNA/entry/the_ibm_support_portal_videos).

Find the content that you need by selecting your products from the IBM Support Portal (<http://www.ibm.com/support/entry/portal>).

Before contacting IBM Support, you will need to collect diagnostic data (system information, symptoms, log files, traces, and so on) that is required to resolve a problem. Gathering this information will help to familiarize you with the troubleshooting process and save you time.

Service request

Service requests are also known as Problem Management Reports (PMRs). Several methods exist to submit diagnostic information to IBM Software Technical Support.

To open a PMR or to exchange information with technical support, view the IBM Software Support Exchanging information with Technical Support page (<http://www.ibm.com/software/support/exchangeinfo.html>).

Fix Central

Fix Central provides fixes and updates for your system's software, hardware, and operating system.

Use the pull-down menu to navigate to your product fixes on Fix Central (<http://www.ibm.com/systems/support/fixes/en/fixcentral/help/getstarted.html>). You may also want to view Fix Central help.

IBM developerWorks®

IBM developerWorks provides verified technical information in specific technology environments.

As a troubleshooting resource, developerWorks provides easy access to the most popular practices, in addition to videos and other information: developerWorks (<http://www.ibm.com/developerworks>).

IBM Redbooks®

IBM Redbooks are developed and published by the IBM International Technical Support Organization, the ITSO.

IBM Redbooks (<http://www.redbooks.ibm.com>) provide in-depth guidance about such topics as installation and configuration and solution implementation.

Software support and RSS feeds

IBM Software Support RSS feeds are a quick, easy, and lightweight format for monitoring new content added to websites.

After you download an RSS reader or browser plug-in, you can subscribe to IBM product feeds at IBM Software Support RSS feeds (<https://www.ibm.com/software/support/rss>).

Log files

Log files can help you troubleshoot problems by recording the activities that take place when you work with a product.

Error messages

The first indication of a problem is often an error message. Error messages contain information that can be helpful in determining the cause of a problem.

After the Solution Installer is run, Chef processes are still running after the cleanup.sh script is executed

After the cleanup.sh script is run on the Solution Installer node, all Chef processes should be terminated. However, in some cases Chef processes are still running.

If you try to run the Solution Installer setup.sh file again, you might receive a message that required ports are in use, or creation of the Chef user might fail.

The solution to this issue is to reboot the installation node after you run the "cleanup.sh" script.

Note: This information applies only to the Solution Installer node and not the client/target nodes.

Attention: The Solution Installer runs on the latest versions of Chrome and Firefox. It is not tested on Safari and does not run on Internet Explorer.

Notices

This information was developed for products and services offered worldwide.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant 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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 may 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 Software Group
Attention: Licensing
3755 Riverside Dr.
Ottawa, ON
K1V 1B7
Canada

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 measurements 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

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 www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.

Index

A

accessibility 23
Analytics Platform 5, 9

C

configuration 3
 modifying sudoers file 6
configuration files
 db2_connection.properties 25
Customer Experience Dashboard 19
customer relationship data 25

D

data provision tables
 loading 25
data provision tool 25
data source connection
 BigInsights database 13
Dataset cronjobs 12
datasets 11

F

firewall settings 7
 configure with script 8
 disabling 7
firewall.sh 8

I

installation 3
 deploying installation files 10
 modifying sudoers file 6
 overview 3

M

master node 6, 9, 12
Mediation node 5, 9

N

nodes
 Business Intelligence node 13, 16
 requirements 10

P

package
 contents 3
Presto 14

S

servers
 See nodes
SLM log 27
SLM tag 27
solution
 topology 3
solution content 9
solution datasets 11
Solution Installer 5, 6
sudoers file 6

T

troubleshooting 29
 cleanup.sh 31
 resources 29

V

veda component 5, 9