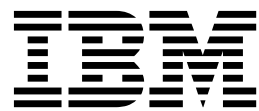IBM Industry Customer Intelligence Solutions
Last updated: 2016-04-07

Customer Experience Analytics for Communication Service Providers Solution Guide

IBM

Product Information

This document applies to IBM Industry Customer Intelligence Solutions Version 1.0.0 and may also apply to subsequent releases.

Contents

Introduction

The IBM Customer Experience Analytics for Communication Service Providers solution gives you the information and insight that you need to provide proactive service to your customers.

Audience

This guide is intended for administrators and users of the Customer Experience Analytics for Communication Service Providers solution. It provides information on installation, configuration, and use of the Customer Experience Analytics for 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 Customer Experience Analytics for Communication Service Providers have accessibility features. For more information, see Appendix A, "Accessibility features," on page 23.

Customer Experience Analytics for 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 Customer Experience Analytics for Communication Service Providers is a prepackaged solution that works with IBM Predictive Customer Intelligence.

Important: Before you install the application, IBM Predictive Customer Intelligence must be installed on your system.

IBM Customer Experience Analytics for Communication Service Providers is a solution that helps measure and improve customer network experience. The solution uses advanced analytics capability to measure network activity across locations, devices, and applications.

The solution provides customer experience team members in the telecommunications industry a comprehensive customer experience analytics dashboard. The dashboard displays customer experience analytics across all services in one simple access view, including metrics such as:

- Unacceptable experience trends
- Dropped calls
- Video experience
- Web experience
- Device usage
- Application usage
- Overall customer experience scores

What need does the solution address?

Customer experience teams have many unanswered questions due to limited data and analytics. Some of these questions include:

- Why are my customers not satisfied or happy?
- How can we measure our overall customer experience score?
- What factors impact customer satisfaction and happiness?
- How can we drill down into specific factors, such as web, video, voice, and messaging services based on a customer usage profile?

What does the solution provide?

IBM Customer Experience Analytics for Communication Service Providers provides the following features:

- An overall customer experience score based on a weighted score for web, video, voice, and messaging services.
- The ability to discover key issues impacting overall customer experience scores, such as technologies, applications, devices, and locations.
- Processing of large volumes of raw usage data from network data sources such as voice or SMS CDRs, and Data Network Records XDRs.

Chapter 2. Installation and configuration

The IBM Customer Experience Analytics for Communication Service Providers solution is for use with IBM Predictive Customer Intelligence.

Customer Experience Analytics for Communication Service Providers consists of the following nodes:

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

Note: The IBM Predictive Customer Intelligence topology also includes a Predictive Analytics node, data node, and an Integration Bus node, which are not used in the Customer Experience Analytics for Communication Service Providers solution.

Contents of the solution package

The IBM Customer Experience Analytics for 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 IBM Cognos Framework Manager models
* 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 Customer Experience Analytics for Communication Service Providers solution.

The following diagram provides a high-level overview of the installation and configuration process. If you are viewing this diagram from the IBM Knowledge Center, you can click on a task to go to the procedure for that task.

1 Download the package to the installation server and ensure network connectivity between all server nodes in the deployment
Configure firewall

Solution installer

2 Install the Solution Installer, and configure it to identify the computers where the parts of the solution will reside. The Solution Installer then copies the files to the computers that you specified.

IBM Predictive Customer Intelligence

Analytics Platform Data Node

Analytics Platform Mgmt Node

PCI Business Intelligence Node

Streams Node

There can be multiple Analytics Platform Data nodes and a minimum of one Analytics Platform Mgmt node per three Analytics Platform Data nodes

There can be multiple Streams nodes

Final steps

3 Configure the various nodes.

Install Analytics Platform package and configure

Deploy Cognos content

Install Streams Mediation package and configure

Perform the following tasks:

1. Install prerequisite software.
2. Download the solution from IBM Passport Advantage®.
3. 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 node and Streams node.
4. Run the Solution Installer to deploy the installation files to the relevant nodes.
5. Set up the IBM Cognos Framework Manager model.
6. Install IBM Cognos reports.
7. Install and configure the Analytics Platform node .
8. Install and configure the Streams node.

Prerequisites

Before you install the Customer Experience Analytics for Communication Service Providers solution, there are prerequisites that must be met.

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

- IBM Predictive Customer Intelligence version 1.1.0 is installed and configured and you have applied all known fix packs to it.
- IBM InfoSphere® BigInsights® 4.1.0 is installed and configured on the Analytics Platform node.
- IBM InfoSphere Streams 4.0.1 is installed and configured on the Streams nodes. For more information, see the installation instructions for IBM InfoSphere Streams 4.0.1 on IBM Knowledge Center.

- 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 administration rights and the ability to copy files between computers.

In the context of the Customer Experience Analytics for Communication Service Providers solution, the BigInsights master node is one of the Analytics Platform nodes.

Installing IBM Predictive Customer Intelligence version 1.1.0

To install PCI version 1.1.0, refer to the IBM Predictive Customer Intelligence *Installation Guide* 1.1.0 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.

Installing IBM InfoSphere BigInsights 4.1.0

Before installing IBM InfoSphere BigInsights 4.1.0, you must first install IBM Open Platform (IOP) with Apache Hadoop. This is an IBM open-source download that provides the components and services to start your Hadoop project, including development tools, a management dashboard, and cluster management tools. This download is available form the IBM Open Platform with Apache Hadoop 4.1 web page.

To prepare your environment and servers, refer to the documentation Preparing to install IBM Open Platform with Apache Hadoop on IBM Knowledge Center.

To install IOP with Apache Hadoop, refer to the documentation Installing the IBM Open Platform with Apache Hadoop software on IBM Knowledge Center.

This documentation will guide you through the key set up and installation steps, including
- reference architecture
- minimum system requirements
- preparing your environment
- default ports

To ensure reliability, you may wish to use a mirror site rather than IBM's hosted repository. If you require a repository, refer to Creating a mirror repository for the IBM Open Platform with Apache Hadoop software. The repository can be installed on the same computer as IOP with Apache Hadoop.

Regarding the information provided in the topic Installing the IBM BigInsights value-added services on IBM Open Platform with Apache Hadoop, ensure the following added value services are installed. These services should be deployed on your designated Analytics Platform nodes:

- BigInsights Home
- BigInsights -- Big SQL Service
- BigInsights -- Data Server Manager Service

Installing IBM IBM InfoSphere Streams 4.0.1

IBM InfoSphere Streams enables you to capture, analyze, and act on business data from a broad range of sources. Using real-time analytic processing (RTAP) technology, InfoSphere Streams can continuously analyze massive volumes of streaming data in real time, with extremely fast response times.

InfoSphere Streams must be installed before you install the IBM Streams mediation package, which is part of the IBM Customer Experience Analytics for Communication Service Providers solution.

To install InfoSphere Streams 4.0.1, refer to Installing IBM InfoSphere Streams on IBM Knowledge Center. Follow the instructions for installing the product on the designated Streams mediation nodes in your deployment. Please note the following:

- InfoSphere Streams provides a dependency checker script that you run on your system before installing the product. This script verifies that your system meets the software requirements for InfoSphere Streams. To run the dependency checker script, enter the **./dependency_checker.sh** command from/opt/tnf/streams/StreamsInstallFiles.
- If required, correct any dependencies and run the dependency checker again. The following message will be displayed once all dependencies are in place:

  ```
  === Summary of Errors and Warnings ===
  CDISI0003I The dependency checker evaluated the system and did not find errors or warnings.
  ```
- If the dependency checker display the following warning, the **ulimit** settings need to be updated and the device rebooted. Please refer to Guidelines for configuring Linux ulimit settings for InfoSphere Streams.

  ```
  Warning: CDISI3044W The system soft limit for the maximum number of processes is low: 1024
  ```
- Install InfoSphere Streams in console mode using the following command:
  ./InfoSphereStreamsSetup.bin -i console

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

To run the Customer Experience Analytics for 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 appears 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 Customer Experience Analytics for Communication Service Providers node.

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.

Disabling firewall settings during installation (optional)

If you choose to temporarily disable your computer's firewall during the installation of the IBM Customer Experience Analytics for 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

If you prefer not to disable your firewall settings or use the firewall.sh script, you can manually configure your firewall settings on the installation node.

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. On the computer that contains the solution installer node, add rules for the ports that are required by the installer by typing the following commands:

```
iptables -I INPUT -p tcp --dport 8080 -j ACCEPT
iptables -I OUTPUT -p tcp --sport 8080 -m state --state
RELATED,ESTABLISHED -j ACCEPT

iptables -I INPUT -p tcp --dport 443 -j ACCEPT
iptables -I OUTPUT -p tcp --sport 443 -m state --state
RELATED,ESTABLISHED -j ACCEPT

iptables -I INPUT -p tcp --dport 9683 -j ACCEPT
iptables -I OUTPUT -p tcp --sport 9683 -m state --state
RELATED,ESTABLISHED -j ACCEPT

iptables -I OUTPUT -p tcp --dport 22 -j ACCEPT
iptables -I INPUT -p tcp --sport 22 -m state --state
RELATED,ESTABLISHED -j ACCEPT
```

Configuring your firewall settings on the installation node computer by using the firewall script

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.

Configuring firewall settings on other computers that will contain content from the solution

On all node computers to which content from the solution will be copied, the ports 443, 9683, and 8080 must be opened for the installation process.

Procedure

1. On each of the computers to which content for the solution will be copied, add rules for the ports that are required by the installer by typing the following commands:

```
iptables -I OUTPUT -p tcp --dport 443 -j ACCEPT
iptables -I INPUT -p tcp --sport 443 -m state --state RELATED,ESTABLISHED -j ACCEPT
iptables -I OUTPUT -p tcp --dport 9683 -j ACCEPT
iptables -I INPUT -p tcp --sport 9683 -m state --state RELATED,ESTABLISHED -j ACCEPT
iptables -I OUTPUT -p tcp --dport 8080 -j ACCEPT
iptables -I INPUT -p tcp --sport 8080 -m state --state RELATED,ESTABLISHED -j ACCEPT
```

2. Type the following command to save your firewall settings:

   `/etc/init.d/iptables save`

3. Type the following command to restart your firewall:

   `/etc/init.d/iptables restart`

4. Verify that your changes are applied by using the following command:

   `iptables -L`

   The added port appears in the list.

5. Repeat these steps on each computer on which you will install an IBM
   Customer Experience Analytics for Communication Service Providers node.

Installing the solution installer

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

About this task

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

http://youtu.be/9u6ZOk8iyt4

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. Go to the solution installer directory in the following location:
   `opt/IBM/IS_CSP_CustomerAnalytics_1.0.1/SolutionInstaller`.
4. If you haven't already opened the ports required by the solution installer, you
   can open them by running the following command:

   `./firewall.sh`

5. Enter the following command:

   `./setup.sh` *username first name last name email*

   Username is the name that can be used to log on to the Chef server. It must be
   all lowercase characters. First and last names are for display purposes only.
   Email is used for notifications from Chef server. The password for the user that
   is created in chef is optional. If it is not included, it will be prompted for at
   login.

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

6. 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.

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 Industry 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 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 BI Node before adding Cognos 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.

Deploying content on the Business Intelligence node

You must deploy the content files (report images and Framework Manager model) for the Customer Experience Analytics for Communication Service Providers solution.

Procedure

1. Log on to the BI node as the **root** user or as a user with **sudo** permissions.

2. Decompress the following files:
   - CSP_CustomerAnalytics_FMProject.zip
   - CSP_CustomerAnalytics_Images.zip

3. Extract the CSP_CustomerAnalytics_BIContent.zip folder and copy it to the *Cognos_Install_location*/Deployment folder on the BI node computer.

4. Copy the CSP_CustomerAnalytics_Images folder to the *Cognos_Install_location*/webcontent/samples folder on the BI node computer. If a folder of this name already exists at the destination location, merge the two folders.

Setting up the Cognos Framework Manager model

Open the project file in IBM Cognos Framework Manager and save it.

Before you begin

The user ID that you log in with when you open Framework Manager must belong to the **db2admins** group.

Procedure

1. On the Business Intelligence node, open Cognos Framework Manager.

2. Click **Open a project**.

3. Select the *File_saved_location*/CSP_CustomerAnalytics_FMModel/ CSP_CustomerAnalytics_FMModel.cpf file.

4. Click **OK**. The model opens and you can view the data model, sources, and package.

5. Publish the package.

Deploying the IBM Cognos BI reports

You use IBM Cognos Administration to deploy the reports.

Before you begin

Ensure that `CSP_CustomerAnalytics_BIContent.zip` is in the *`Cognos_install_location`*`/Deployment` folder on the Business Intelligence node.

Procedure

1. Open a web browser on the Business Intelligence node.
2. Go to the IBM Cognos BI portal URL. For example, go to `http://cognos node/ibmcognos/`.
3. On the **Welcome** page, click **Administer IBM Cognos Content**.
4. Click the **Configuration** tab, and then click **Content Administration**.
5. On the toolbar, click **New Import**.
6. In the **Deployment Archive** pane, select deployment archive (CSP_CustomerAnalytics_BIContent.zip), and click **Next**.
7. In the **Specify a name and description** pane, accept the default or enter a new name, and click **Next**.
8. In the **Select the public folders and directory content** pane, select all of the packages in the table, leave the **Options** as default, ensure that the **Disable after import** option is unselected, and click **Next**.
9. In the **Specify the general options** pane, accept the defaults, and click **Next**.
10. On the **Review the summary** page, click **Next**.
11. On the **Select an action** page, select **Save and run once**, and click **Finish**.
12. On the **Run with options** page, accept the defaults, and click **Run,** and then click **OK**.
13. Select **View the details of this import after closing this window** and click **OK**.
14. In IBM Cognos Administration, click **Home**. The content is available in **Public Folders** > **CSP_CustomerAnalytics_BIContent_Model** > **CSP_CustomerAnalytics** > **CSP_CustomerAnalytics_Analysis_Report**.

Installing the Analytics Platform

The Analytics Platform is software that is installed on your Analytics Platform nodes. It manages the creation of your data sets in the Apache Hadoop ecosystem. Analytics Platform allows deployment of data sets and management of those data sets, which makes up the basis of your solution.

To install and configure the Analytics Platform, refer to the Analytics Platform *Installation and Configuration Guide for CSP Solutions* available from the following location: http://www-01.ibm.com/support/docview.wss?uid=swg27046047

This document describes the installation and configuration of Analytics Platform software and data sets for Communication Service Provider (CSP) solutions. It includes operations that are required to be performed on your Analytics Platform nodes in your deployment.
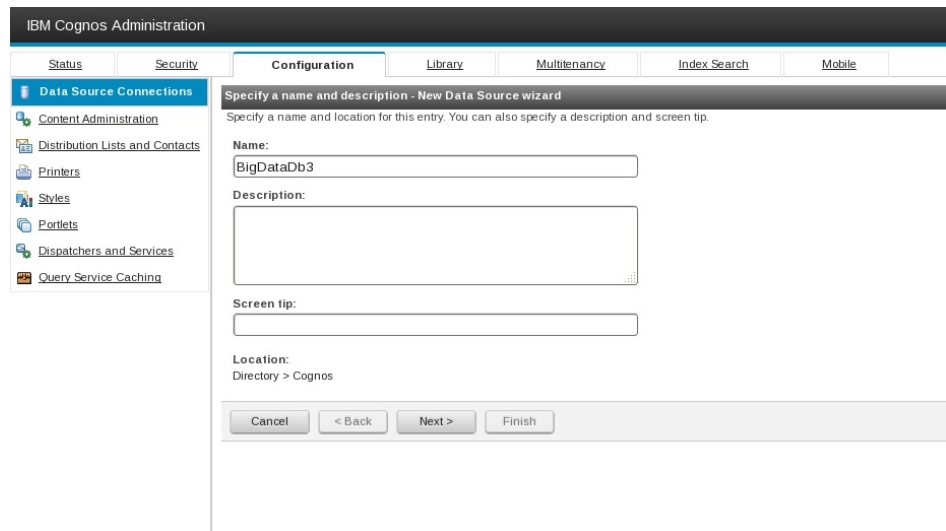
In the context of the Customer Experience Analytics for Communication Service Providers solution, the BigInsights master node is one of the Analytics Platform nodes.

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.
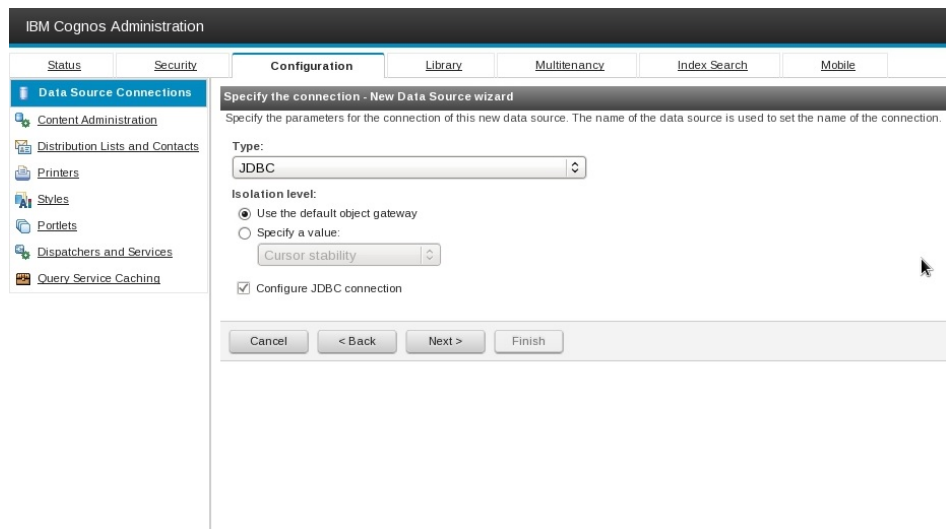
Procedure

1. Open a web browser on the BI node.
2. Go to the IBM Cognos BI portal URL. For example, go to `http://cognos node/ibmcognos`.
3. On the **Welcome** page, click **Administer IBM Cognos Content**.
4. Click the **Configuration** tab, and then click **Data Source Connections**.
5. Click **New Data Source**.
6. In the **Name** field, type `BigDataDb2` (or `BigDataDb3` if you already have a `BigDataDb2` connection), and then click **Next**.
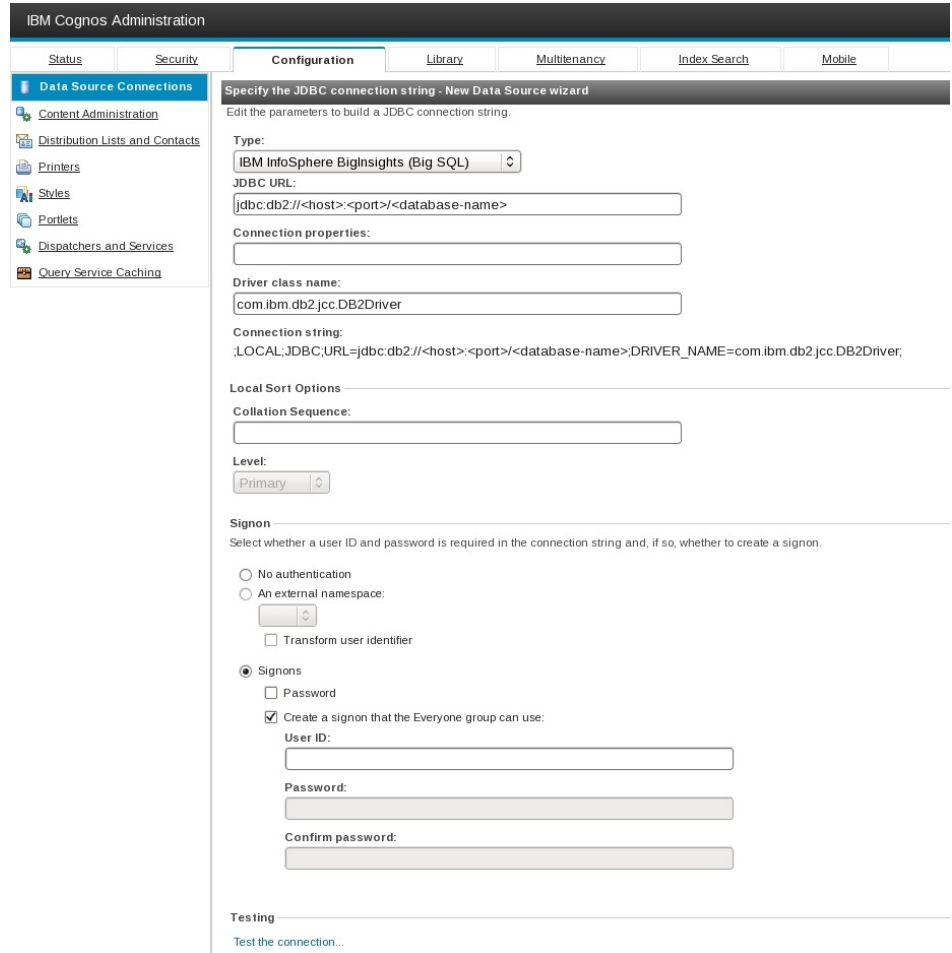

7. From the **Type** list, select `JDBC`. Ensure that **Configure JDBC connection** is selected and click **Next**.

8. From the **Type** list, select `IBM InfoSphere BigInsights (Big SQL)`.

9. Locate the connection information for the **JDBC URL** field in the image shown in step 8.

10. To find the connection settings on the Apache Ambari tool, click **Services**, then click **BigInsights - Big SQL**.

11. Click the **Configs** tab, and note the **db2_port_number**.
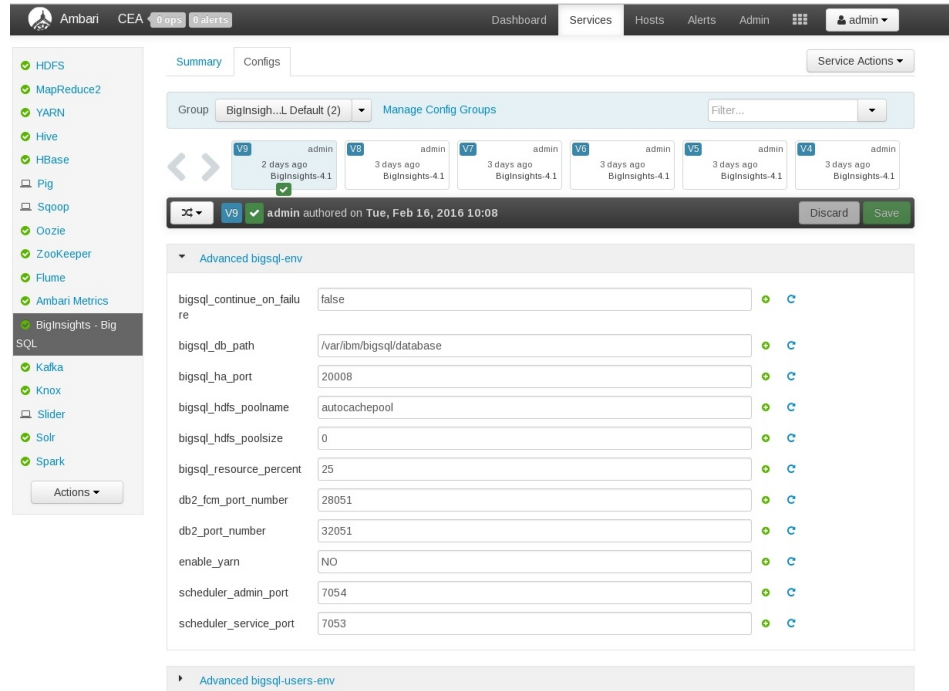
12. Click the **Summary** tab.
13. To locate the server name, click the link for the **Big SQL Head** node.
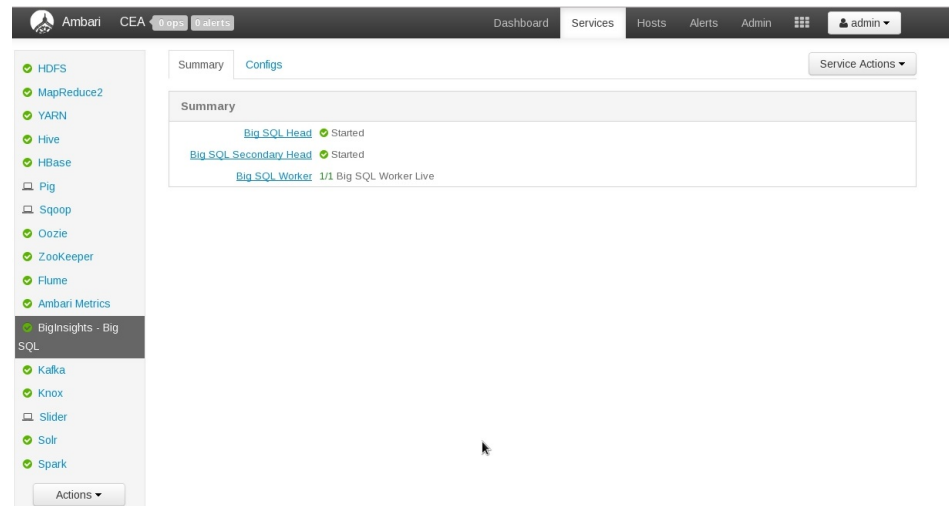
Note: On a multi-node installation, only the head nodes can be queried.

14. The node name appears below the Ambari banner, in this case hadoop2.com:

15. Using the information from steps 10-14, on the **Specify the JDBC connection string** page, specify the server name, port number, and database name.

16. Click **Test the connection**, then click **Test**. On the **View the results** page, the status of the connection tests for the dynamic query mode is **Succeeded**.
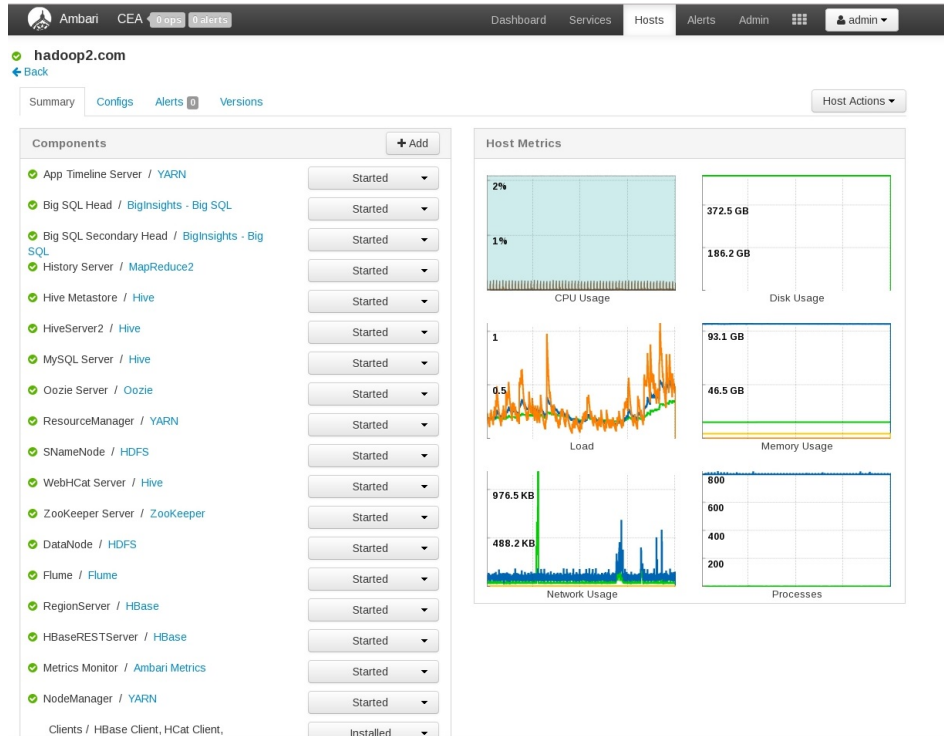
17. Click **Close**, click **Close** again, and then click **Finish**.

Installing the SLM tagging functionality

You can deploy the slm tagging functionality on the Analytics Platform node.

About this task

A RPM Package Manager (RPM) is contained in the `rapidanalytics/build/slm/ci-slm-deployment-1.0.0.rpm` folder. To install this RPM, follow these steps:

Procedure

1. Install the SLM via the RPM installer by typing the following command:

   ```
   rpm -i bis-main-2*.rpm
   ```

   The files will be copied to `/opt/tnf/apps/slm`.

2. Running SLM tagging:
   - A Cron job needs to be created that runs the SLM tagging command daily, just after midnight:

     ```
     java -cp /opt/tnf/apps/slm/lib/log4j-1.2.16.jar:/opt/tnf/apps/slm/lib/* com.tnf.slm.SlmMain
     ```

   - The SLM tagging process appends to a log file in the directory, with the extension `.slmtag`: /opt/tnf/apps/slm/log/<log file>.slmtag

   This log file is used by IBM to ensure that the product is being used according to its license.

Installation and configuration of the Mediation node

Mediation software is deployed on the mediation 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. The mediation software works with the Infosphere Streams software installed on your mediation nodes. It manages the loading and key performance indicator (kpi) calculation of your telco data and prepares the data for analysis by the end user. The volume of data to be analyzed will determine how many mediation nodes are required in your platform.

To install and configure the Mediation node, you will need to reference information on Analytics Platform 2.0.

This information includes the following:
- installation prerequisites
- RPM installation
- Vedia Supercollector installation
- streams domain and instance configuration
- customization
- administration

To perform these operations, refer to the Analytics Platform Mediation *Installation and Configuration Guide for CSP Solutions* available from the Industry Solutions Customer Intelligence Solutions version 1.0 documentation web page.

This document provides information regarding installation, deployment, configuration, and maintenance of the mediation layer of Analytics Platform 2.0. This platform is similar to Analytics Platform 1.0, except that 2.0 is Streams-based and 1.0 is Vedia-based.

The mediation solutions include the following components:
- Supercollector (Vedia-based)
- Lookup Manager (Streams-based)
- ITE Application to produce 5 minute aggregations (Streams-based)
- Raw TDR processing Application (Streams-based)
- DB Loader Application (Streams-based)

Please note that the Analytics Platform Mediation *Installation and Configuration Guide for CSP Solutions* will not cover the installation of the Streams platform, which is a prerequisite for the application deployment. Also, the Hadoop installation and Hadoop-based aggregation are not covered in the document.

Chapter 3. Using Customer Experience Analytics for Communications Service Providers

This section describes a typical user case that shows how to use the Customer Experience Analytics for 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 Customer Experience Analytics for 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.
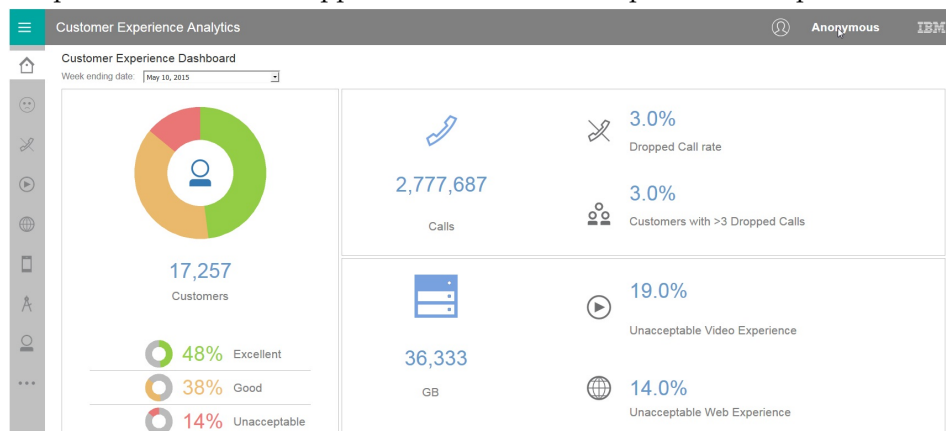
2. Click ☺ to view the **Unacceptable Experience Trend** report.

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.

3. Click ✂ to view the **Dropped Calls Trend Overview** report.
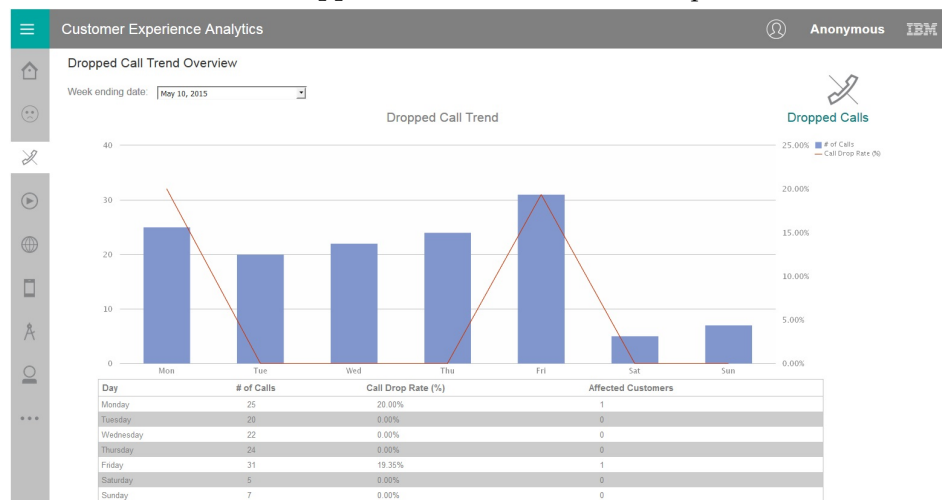
The bar chart shows the number of dropped calls and the line chart shows a corresponding dropped call rate. The report below the chart includes the number of affected customers. The dropped call rate on Monday and Friday was particularly high and may warrant further investigation.

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

5. Click ▢ to view the **Device Data Usage Distribution** report.
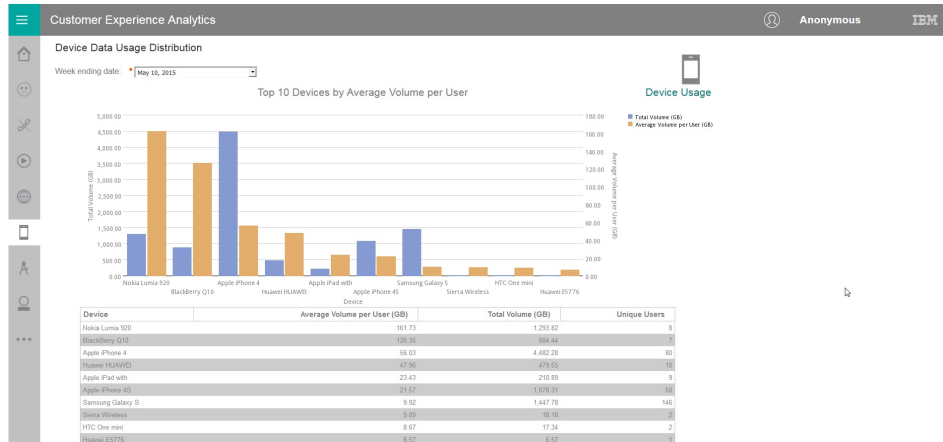
The bar chart shows the total volume and average volume per user of the top ten devices by average volume per user. The report below the chart includes the number of unique users. The average volume per user of Nokia Luria 920 and BlackBerry Q10 users is particularly high, and the total volume of Apple iPhone 4 users is much higher than any other device.

6. Click ⅄ to view the **Application Data Usage Distribution** report.
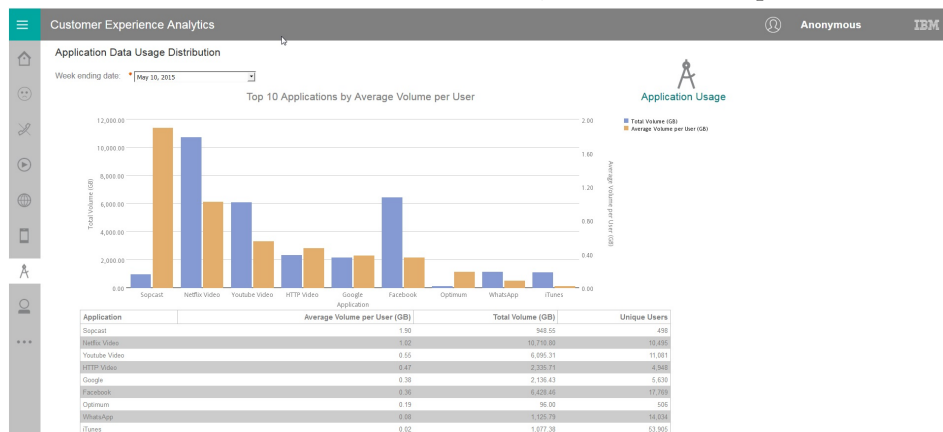
The bar chart shows the total volume and average volume per user of the top ten applications by average volume per user. The report below the chart includes the number of unique users. The average volume per user of Sopcast is particularly high, and the total volume of Netflix Video users is higher than any other device.

7. 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 Business Intelligence 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
```

Configuration files

The data provision tool IBM DB2 database configuration location is `/opt/tnf/apps/bis-main-var/bis-demon/cfg/db2_connection.properties`. For example,

```
db.url=jdbc:db2://localhost:50000/tnfbis
db.user=catalog
db.password=db2
```

where `db.url` is the JDBC URL string of the database, `db.user` is the database user ID, and `db.password` is the database password.

Appendix C. Troubleshooting

This section contains troubleshooting information for the Customer Experience
Analytics for 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.

Troubleshooting the FTP collection of files in Analytics Platform

There are solutions for known issues in the FTP collection of files in Analytics Platform.

On Analytics Platform mediation nodes, the function in Analytics Platform that collects transaction data record (TDR) files for storage in Analytics Platform is `Vedia`.

To verify that the collection of TDRs from the CSP probe or probes is working, note that each probe's collector will be defined in mediation. Use the secure shell (SSH) to connect to the correct collector. The allocation of collectors can be found in the file `Vediamain.cfg` at `/opt/tnf/apps/vedia/cfg/Vediamain.cfg`.

To check FileLoader logs, the format is `logs/Collection_1/FileLoader092708.log`.

When the application is running, messages would appear in the following format. For this example, `Vedia` is comparing TDRs in tmp against files on the probe:

```
[CFileListJob::GetNewFilesList] [TaskTest] Getting list of /var/opt/tnf/tdrs/ *.*
[CFileListJob::GetNewFilesList] [TaskTest] Files: 196, Match: 182, Old: 178
```

The files that are being collected correctly are shown in the following example:

```
[CFileDownloadJob::Process] Task - downloading midas/netrofeed/netro_20130918200012.log.gz
[CFileDownloadJob::Process] Task - download netro_20130918200012.log.gz (9600142)
[CProbe::Store] Save: /opt/tnf/apps/vedia/tmp/Collection_1_FileLoader.dat
[CFileListJob::Save] LastListTime: 20130918200012, Files: 664
[CFileListJob::GetNewFilesList] [Task] Getting list of midas/netrofeed/ *.*
[CFileListJob::GetNewFilesList] [Task] Files: 251, Match: 223, Old: 221
[CFileDownloadJob::Process] Task - downloading midas/netrofeed/netro_20130918200015.log.gz
[CFileDownloadJob::Process] Task - download netro_20130918200015.log.gz (9013063)
```

In contrast, the following is an example of a file that can no longer connect to a probe. Typically, if FTP is not working, the ports are not open, permissions are not set up correctly on the probe, the probe is down, or connectivity between the probe and mediation has dropped.

```
FTP not working
```

TO FIND example of ftp failing

Full queues represent full applications within the chain. The following is an example of a file that has full queues somewhere along the chain. If one

application is full, then the preceding application cannot output to the queue, thus impacting the queues (creating a knock-on effect). The queues will need be examined for skipping.

```
1C8F - 19/09 11:00:55.200 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:00.242 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C91 - 19/09 11:01:00.461 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:05.402 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:10.431 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:15.482 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:20.642 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:25.838 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C91 - 19/09 11:01:30.525 - [CFileListJob::GetNewFilesList]
Skip getting of files list
1C8F - 19/09 11:01:31.063 - [CFileListJob::GetNewFilesList]
Skip getting of files list
```

If skipping is noticed in Vedia logs and does not stop, restart Vedia. Some skipping will occur on restart, but wait up to five minutes for it to clear. It typically takes one to five minutes for the skipping messages to clear from the logs.

Recommendations:

- When collecting files and **Match** is **0**, the regex in the FileLoader configuration is incorrect. Review and correct this in the configuration steps.
- If FileLoader understands the files from the network to be **Old**, contact the person responsible for the production of network files to correct the date and time of the files.
- Occasionally, FTP jobs will fail to collect a file, but FileLoader will retry. Investigate situations where there are many failures, causing files that fail to be collected. Check the configuration and verify that the mediation node can reach the host of the network files (probe). If there is a network problem, contact the person responsible for the network files.
- When the configuration is correct and network connectivity is confirmed, restart the Vedia processes and ensure there are no further reported exceptions.

Queues commonly will fill after a long outage period. Vedia will attempt to collect all files and process them immediately. The number of files is determined by a period of time and is set in Collection.xml. If the look back period is, for example, greater than 20 minutes and there are **Skip getting of files list** messages in the log, then reset the look back period. Carefully tune anything greater than 20. Do not simply increase to 300 minutes, as this can overload Vedia.

The following is configured in Collection.xml:

```
module name="FileLoader"id="FileLoader">
  <probe name="FileLoader"log-file="FileLoader"log--process="2">
  <files>
    <ftp id="Task"host="Data"
    url="ftp://{probe-ip}"port="21"
    user="{probe-user}"password="{probe-password}"
    remotedir="{probe-dir}"storagedir="tmp/data_{probe-no}"
    filename-regex="{probe-files_mask}"
    ftp-timeout="60"ftp-downloads="1"attempt-timeout="1"
    delete-processed="0"process-file-ignore-age="0"
```

```
          max-process-file-age="20"time-shift-="-3600"
          compression_="GNUZIP"
          start-format="{probe-file-date-format}"
          detection-interval="60">
          </ftp>
      </files>
      <output name="LoadedFiles"model="File"limit="50"/>
      </probe>
</module>
```

`process-file-ignore-age` directs `Vedia` to load all the old files. Always set this to **0** in a production environment.

`max-process-file-age` defines, in minutes, how old the files on the probe can be before they will be rejected as being too old.

`detection-interval` indicates how often to check for files, in seconds.

`download-threads` indicates how many threads to download with. Note that, if this is set cumulatively across all collectors to be a higher value than the probe can process, this may impact the probe and it may reject files.

`max-try` indicates how many times to attempt to download a file.

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   12
   transaction data record (TDR)   29
   troubleshooting   29

C
configuration   3
   modifying sudoers file   6
configuration files
   Collection.xml   29
   db2_connection.properties   25
   Vediamain.cfg   29
Customer Experience Dashboard   19
customer relationship data   25

D
data provision tables
   loading   25
data provision tool   25
data source connection
   BigInsights database   13

F
firewall settings   7
   configure with script   8
   disabling   7
   manually configuring   8
   on all computers   8
firewall.sh   8
Framework Manager   11

I
installation   3
   deploying installation files   10
   Mediation node   17
   modifying sudoers file   6
   overview   3
   prerequisites   4

M
master node   4
mediation
   See vedia component
Mediation node   17

N
nodes
   Business Intelligence node   11
   requirements   10

P
package
   contents   3

R
reports   12

S
servers
   See nodes
solution
   topology   3
sudoers file   6

T
troubleshooting   27
   cleanup.sh   31
   resources   27

V
vedia component   12, 17