

Version 11 Release 1
March 15, 2019

IBM Campaign Installation Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 129.

This edition applies to version 11, release 1, modification 0 of IBM Campaign and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1998, 2019.**

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

Contents

Chapter 1. Installation overview 1

Installation roadmap	1
How the installers work	3
Modes of installation	4
Integration of Campaign with eMessage	4

Chapter 2. Planning the Campaign installation 7

Prerequisites	7
IBM SPSS Modeler Advantage Enterprise Marketing Management Edition replaces PredictiveInsight	9
Campaign installation worksheet	9
Installation order for IBM Marketing Software products	12
Planning failover configuration of Campaign	13

Chapter 3. Preparing data sources for Campaign 15

Creating database or schema for Campaign system tables	15
Creating ODBC or native connections	15
Configuring the web application server for your JDBC driver	18
Creating JDBC connections in web application server	18
Information for creating JDBC connections	19
Preparing DB2 on z/OS user datasources for Campaign	21
Preparing Vertica datasources for Campaign	21
Preparing Amazon Redshift user data sources for Campaign	22
Preparing Apache Hadoop Hive user data sources for Campaign	23
Requirements and restrictions for Hive-based Hadoop data sources	24
Overview of the Campaign integration with Hadoop	24
A. Install a Hive ODBC driver	26
B. Configure the Hive ODBC Driver	28
C. Map existing HBase tables to Hive	31
D. Import and configure the BigDataODBCHiveTemplate data source template in Campaign	32
E. Configure SSH on the Campaign listener server	34
F. Map the Hive datasource in Campaign	35
Preparing Cloudera Hadoop Impala based user data sources for Campaign	35
Requirements and restrictions for Impala-based Hadoop data sources	36
A. Install a Cloudera Impala ODBC driver	36
B. Installing the Impala ODBC driver for Windows	37
C. Configure the Impala ODBC Driver	37

C. Import and configure the BigDataODBCImpalaTemplate data source template in Campaign	38
D. Configure SSH on the Campaign listener server	40
F. Configuring connectivity between Campaign listener server and Hadoop cluster	41
G. Map the Impala datasource in Campaign	42
Preparing Cloudera BigData instance using Kerberos authentication	42
Installing and configuring MIT Kerberos for Windows	43
Downloading and installing Kerberos Packages for Linux	44
Seamless Kerberos connection on Windows	45
Seamless Kerberos connection on Linux	46
Configuring the Impala-based Cloudera data sources using Kerberos	47
Configuring the Impala-based Cloudera data sources using Kerberos on Linux	47
Troubleshooting	48

Chapter 4. Installing Campaign 49

Installing Campaign by using the GUI mode	50
Creating an EAR file after running the installer	55
Installing Campaign using the console mode	55
Installing Campaign silently	57
Sample response files	57

Chapter 5. Configuring Campaign before deployment 59

Creating and populating Campaign system tables manually	59
Creating and populating the eMessage system tables manually	60
Registering Campaign manually	61
Registering eMessage manually	62
Setting data source variables in the Campaign startup script (UNIX only)	62
Database and library environment variables	63

Chapter 6. Deploying the Campaign web application 65

Setting the web application session timeout	65
Deploying Campaign on WebSphere Application Server	65
Deploying Campaign on WAS from a WAR file	66
Deploying Campaign on WAS from an EAR file	67
Deploying Campaign on WebLogic	68
Configuring WebLogic to display reports (UNIX)	69
Deploying Campaign on JBoss Application Server	70
Deploying Campaign on Tomcat Application Server	70
Starting the Campaign server	71
Starting the Campaign listener manually	71

Installing the Campaign listener as a Windows service	72
---	----

Chapter 7. Configuring after deployment. 73

Verifying that the Campaign listener is running	73
Setting up the Campaign system user	73
Adding data source properties on the Configuration page.	74
Importing data source templates	75
Duplicating a data source template	75
Campaign configuration properties	75
Mapping user tables in Campaign	77
Verifying the Campaign installation	77
Setting properties for integration with IBM Marketing Software products	77
Additional configuration for enhanced security	77
Disabling the X-Powered-By flag	77
Configuring a restrictive cookie path	78
Starting the Campaign Advance search agent	78

Chapter 8. Configuring multiple partitions for Campaign 81

How multiple partitions work	81
Setting up multiple partitions	82
Partition superuser	83
Configuring data source properties for partitions	83
Setting up system users for Campaign	85
Using IBM Cognos Reports for multiple partitions	85
Assigning roles, permissions, and groups for partitions	85

Chapter 9. Configuring multiple partitions in eMessage 87

Partitions for eMessage: Overview.	87
Roadmap for configuring multiple partitions in eMessage	88
Creating a new partition for eMessage	88
Preparing eMessage system tables for the partition	90
Creating and populating the eMessage system tables manually	90
Configuring a system user to access IBM Marketing Software Hosted Services.	92

Enabling eMessage in Campaign for the new partition	93
Specifying the Recipient List Uploader location for eMessage	93
Restarting system components after configuring eMessage	93
Testing eMessage partition configuration and connections	94

Chapter 10. Marketing Platform utilities 95

alertConfigTool	97
configTool	97
datafilteringScriptTool	101
encryptPasswords	102
partitionTool	103
populateDb	105
restoreAccess	106
scheduler_console_client.	108
eMessage Response and Contact Tracker (RCT) script	109
The MKService_rct script	110

Chapter 11. Uninstalling Campaign 113

Chapter 12. IBM Campaign web application clustering 115

WebSphere clustering guidelines	116
Jboss clustering guidelines	117
Configure ehcache.	118

Chapter 13. Installing a clustered listener environment 121

Supported listener clustering configurations	124
Diagram of listener clustering	124
Shared network location for clustered listeners: campaignSharedHome	125

Before you contact IBM technical support 127

Notices 129

Trademarks	131
Privacy Policy and Terms of Use Considerations	131

Chapter 1. Installation overview

An installation of Campaign is complete when you install, configure, and deploy Campaign. The Campaign Installation Guide provides detailed information about installing, configuring, and deploying Campaign.

Use the Installation Roadmap section to obtain a broad understanding about using the Campaign Installation Guide.

Installation roadmap

Use the installation roadmap to quickly find the information that you need for installing Campaign.

You can use Table 1 to scan the tasks that must be completed for installing Campaign. The **Information** column in the following table provides links to the topics that describe the tasks for installing Campaign:

Table 1. Campaign installation roadmap

Chapter	Information
Chapter 12, "IBM Campaign web application clustering," on page 115	If you are using web application clustering, review this appendix before starting your installation.
Chapter 13, "Installing a clustered listener environment," on page 121	If you are using Campaign listener clustering, review this appendix before starting your installation.
Chapter 1, "Installation overview"	This topic provides the following information: <ul style="list-style-type: none">• "How the installers work" on page 3• "Modes of installation" on page 4• "Integration of Campaign with eMessage" on page 4
Chapter 2, "Planning the Campaign installation," on page 7	This topic provides the following information: <ul style="list-style-type: none">• "Prerequisites" on page 7• "Campaign installation worksheet" on page 9• "Installation order for IBM Marketing Software products" on page 12• "Planning failover configuration of Campaign" on page 13

Table 1. Campaign installation roadmap (continued)

Chapter	Information
Chapter 3, "Preparing data sources for Campaign," on page 15	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • "Creating database or schema for Campaign system tables" on page 15 • "Creating ODBC or native connections" on page 15 • "Configuring the web application server for your JDBC driver" on page 16 • "Creating JDBC connections in web application server" on page 18 • "Preparing DB2 on z/OS user datasources for Campaign" on page 21 • "Preparing Amazon Redshift user data sources for Campaign" on page 22 • "Preparing Apache Hadoop Hive user data sources for Campaign" on page 23
Chapter 4, "Installing Campaign," on page 49	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • "Installing Campaign by using the GUI mode" on page 50 • "Installing Campaign using the console mode" on page 55 • "Installing Campaign silently" on page 57
Chapter 5, "Configuring Campaign before deployment," on page 59	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • "Creating and populating Campaign system tables manually" on page 59 • "Registering Campaign manually" on page 61 • "Setting data source variables in the Campaign startup script (UNIX only)" on page 62
Chapter 6, "Deploying the Campaign web application," on page 65	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • "Setting the web application session timeout" on page 65 • "Deploying Campaign on WebSphere Application Server" on page 65 • "Deploying Campaign on WebLogic" on page 68 • "Deploying Campaign on JBoss Application Server" on page 70 • "Deploying Campaign on Tomcat Application Server" on page 70 • "Starting the Campaign server" on page 71
Chapter 7, "Configuring after deployment," on page 73	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • "Verifying that the Campaign listener is running" on page 73 • "Setting up the Campaign system user" on page 73 • "Adding data source properties on the Configuration page" on page 74 • "Campaign configuration properties" on page 75 • "Mapping user tables in Campaign" on page 77 • "Verifying the Campaign installation" on page 77 • "Setting properties for integration with IBM Marketing Software products" on page 77

Table 1. Campaign installation roadmap (continued)

Chapter	Information
Chapter 8, “Configuring multiple partitions for Campaign,” on page 81	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • “How multiple partitions work” on page 81 • “Setting up multiple partitions” on page 82 • “Assigning roles, permissions, and groups for partitions” on page 85
Chapter 9, “Configuring multiple partitions in eMessage,” on page 87.	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • “Partitions for eMessage: Overview” on page 87 • “Roadmap for configuring multiple partitions in eMessage” on page 88 • “Creating a new partition for eMessage” on page 88 • “Preparing eMessage system tables for the partition” on page 90 • “Configuring a system user to access IBM Marketing Software Hosted Services” on page 92 • “Enabling eMessage in Campaign for the new partition” on page 93 • “Specifying the Recipient List Uploader location for eMessage” on page 93 • “Restarting system components after configuring eMessage” on page 93 • “Testing eMessage partition configuration and connections” on page 94
“configTool” on page 97	<p>This topic provides the following information:</p> <ul style="list-style-type: none"> • Chapter 10, “Marketing Platform utilities,” on page 95 • “eMessage Response and Contact Tracker (RCT) script” on page 109 • “The MKService_rct script” on page 110
Chapter 11, “Uninstalling Campaign,” on page 113	<p>This topic provides information about how to uninstall Campaign.</p>

How the installers work

You must use the suite installer and the product installer when you install or upgrade any IBM® Marketing Software product. For example, for installing Campaign, you must use the IBM Marketing Software suite installer and the IBM Campaign installer.

Make sure that you use the following guidelines before you use the IBM Marketing Software suite installer and the product installer:

- The suite installer and the product installer must be in the same directory on the computer where you want to install the product. When multiple versions of a product installer are present in the directory with the master installer, the master installer always shows the latest version of the product on the IBM Marketing Software Products screen in the installation wizard.
- If you are planning to install a patch immediately after you install an IBM Marketing Software product, make sure that the patch installer is in the same directory as that of the suite and product installers.

- The default top-level directory for IBM Marketing Software installations is /IBM/IMS for UNIX or C:\IBM\IMS for Windows. However, you can change the directory during installation.

Modes of installation

The IBM Marketing Software suite installer can run in one of the following modes: GUI mode, console mode, or silent mode (also called the unattended mode). Select a mode that suits your requirements when you install Campaign.

For upgrades, you use the installer to perform many of the same tasks that you perform during the initial installation.

GUI mode

Use the GUI mode for Windows or the X Window System mode for UNIX to install Campaign by using the graphical user interface.

Console mode

Use the console mode to install Campaign by using the command line window.

Note: To display the Installer screens correctly in console mode, configure your terminal software to support UTF-8 character encoding. Other character encoding, such as ANSI, will not render the text correctly, and some information will not be readable.

Silent mode

Use the silent or unattended mode to install Campaign multiple times. The silent mode uses response files for installation, and does not require user input during the installation process.

Note: Silent mode is not supported for upgrade installations in clustered web application or clustered listener environments.

Integration of Campaign with eMessage

When IBM Campaign is integrated with IBM eMessage, you can use eMessage to conduct highly personalized email marketing campaigns. eMessage provides access to the resources that are hosted by IBM. With eMessage you can design, send, and individually monitor customized messages that are based on the information that is stored in your customer data mart.

In Campaign, use flowcharts to create lists of email recipients and select personalization data for each recipient.

In eMessage, use email design, transmission, and delivery resources that are hosted by IBM to conduct email marketing campaigns.

When you install IBM Campaign, the installer automatically includes the files that are required to support IBM eMessage. The following actions are completed for eMessage:

- eMessage is created as a sub-directory in the Campaign directory structure.

- eMessage configuration properties are listed in the IBM Marketing Platform, but the properties are not active.
- eMessage-specific database tables are created in the Campaign schema. However, the database tables contain only initial data.
- Menus and other features that are specific to eMessage do not appear until you enable and configure eMessage.

Before you can send personalized marketing email, you must request a hosted email account from IBM.

When you request an email account, IBM begins a consultation process to familiarize you with eMessage, connect you to hosted email resources, and establish your reputation as a legitimate email marketer among leading Internet Service Providers (ISPs). Establishing a favorable reputation is critical to successfully delivering your marketing message to your customers and prospects.

For more information about how to enable and configure eMessage and prepare your hosted email account, see the *IBM eMessage Startup and Administrator's Guide*.

eMessage components

eMessage requires specialized components called the Recipient List Uploader (RLU) and the Response and Contact Tracker (RCT).

The RLU is a eMessage plugin component that operates with Campaign to upload addresses, personalization data, and metadata that is associated with a list of email recipients to IBM Marketing Software Hosted Services.

The eMessage RCT retrieves link tracking and email delivery notification data from IBM Marketing Software Hosted Services, and stores the data in the eMessage system tables that reside in the Campaign schema.

The eMessage components operate when you enable and configure IBM eMessage. For more information about enabling eMessage and working with the RLU and RCT, see the *IBM eMessage Startup and Administrator's Guide*.

Default component installation locations

The IBM installer places the RLU on the computer where you installed the Campaign J2EE application. The RLU location is recorded in the **Campaign > partitions > partition1 > eMessage > eMessagePluginJarFile** configuration property.

The installer places the RCT on the computer where you installed the Campaign server.

If the J2EE and server components are on separate computers, run the installer on each machine to install the RLU with the J2EE application and the RCT with the Campaign server.

eMessage components in multiple partitions

One RLU exists for the entire eMessage installation. The installer populates the `eMessagePluginJarFile` configuration property only for the default partition. If you are using multiple partitions in your eMessage installation, you must manually

configure the location of the RLU for all other partitions. The location specified in the `eMessagePluginJarFile` property is the same for all partitions. For more information, see “Specifying the Recipient List Uploader location for eMessage” on page 93.

Only one RCT exists for the entire eMessage installation. eMessage does not need to specify the location of the RCT in a configuration property. Responses that are received by the RCT automatically specify the appropriate local partition for correct response attribution.

Chapter 2. Planning the Campaign installation

When you plan your Campaign installation, you must ensure that you have set up your system correctly, and that you have configured your environment to deal with any failures.

Prerequisites

Before you install or upgrade any IBM Marketing Software product, you must ensure that your computer complies with all of the prerequisite software and hardware.

System requirements

For information about system requirements, see the *Recommended Software Environments and Minimum System Requirements* guide.

Network domain requirements

The IBM Marketing Software products that are installed as a suite must be installed on the same network domain to comply with the browser restrictions that are designed to limit the security risks that can occur with cross-site scripting.

JVM requirements

IBM Marketing Software applications within a suite must be deployed on a dedicated Java™ virtual machine (JVM). IBM Marketing Software products customize the JVM that is used by the web application server. If you encounter errors that are related to the JVM, you must create an Oracle WebLogic or WebSphere® domain that is dedicated to the IBM Marketing Software products.

Knowledge requirements

To install IBM Marketing Software products, you must have a thorough knowledge of the environment in which the products are installed. This knowledge includes knowledge about operating systems, databases, and web application servers.

Access permissions

Verify that you have the following network permissions to complete the installation tasks:

- Administration access for all necessary databases
- Read and write access to the relevant directory and sub-directories for the operating system account that you use to run the web application server and IBM Marketing Software components
- Write permission for all files that you must edit
- Write permission for all directories where you must save a file, such as the installation directory and backup directory if you are upgrading
- Appropriate read, write, and execute permissions to run the installer

Verify that you have the administrative password for your web application server.

The following additional permissions are necessary for UNIX:

- The user account that installs Campaign and Marketing Platform must be a member of the same group as the Campaign users. This user account must have a valid home directory, and have write permissions for that directory.
- All installer files for IBM products must have full permissions, for example, `rwxt-xt-x`.

Points to consider before you install Campaign

For Campaign installation you are required to consider the points underneath before beginning the installation. IBM Campaign 10.0 and subsequent versions install IBM GSKIT with the Campaign installation. IBM GSKIT requires "root" credentials for the installation, based on this requirement, before proceeding with IBM Campaign installation you are required to consider / follow the steps underneath:

Note: If you are installing IBM Campaign or other products in the IBM Marketing Software product suite as "root" / "administrator" users these considerations are not applicable. If you are installing IBM Campaign or other products in the IBM Marketing Software suite then the following points should be considered:

Linux / Unix Installations:

Installing IBM Campaign with non root user:

1. You can add the installation user (operating system user) in sudoers list – As a result, during the IBM Campaign installation you will be prompted for the sudo password. After entering this password IBM GSKIT gets installed.

Note: In this case IBM Campaign cannot be installed in silent mode.

2. You can add the installation user (operating system user) in sudoers list and disable the password prompt. (for example rbhagat is the operating system user for the installation, this user when added in the sudoers list with the sudo password prompt disabled then IBM Campaign installation will get successful along with IBM GSKIT installation successfully.

Note: In this case you can install IBM Campaign in silent mode.

Windows:

In Windows, the IBM Campaign user should be added in the administrators group.

JAVA_HOME environment variable

If a **JAVA_HOME** environment variable is defined on the computer where you install an IBM Marketing Software product, verify that the variable points to a supported version of JRE. For information about system requirements, see the *IBM Marketing Software Recommended Software Environments and Minimum System Requirements* guide.

If the **JAVA_HOME** environment variable points to an incorrect JRE, you must clear the **JAVA_HOME** variable before you run the IBM Marketing Software installers.

You can clear the **JAVA_HOME** environment variable by using one of the following methods:

- Windows: In a command window, enter **set JAVA_HOME=** (leave empty) and press **Enter**.
- UNIX: In the terminal, enter **export JAVA_HOME=**(leave empty) and press **Enter**.

The IBM Marketing Software installer installs a JRE in the top-level directory for the IBM Marketing Software installation. Individual IBM Marketing Software application installers do not install a JRE. Instead, they point to the location of the JRE that is installed by the IBM Marketing Software installer. You can reset the environment variable after all installations are complete.

For more information about the supported JRE, see the *IBM Marketing Software Recommended Software Environments and Minimum System Requirements* guide.

Marketing Platform requirement

You must install or upgrade Marketing Platform before you install or upgrade any IBM Marketing Software products. For each group of products that work together, you must install or upgrade Marketing Platform only once. Each product installer checks whether the required products are installed. If your product or version is not registered with Marketing Platform, a message prompts you to install or upgrade Marketing Platform before you proceed with your installation. Marketing Platform must be deployed and running before you can set any properties on the **Settings > Configuration** page.

IBM SPSS Modeler Advantage Enterprise Marketing Management Edition replaces PredictiveInsight

If you previously used IBM PredictiveInsight, you can no longer use the legacy Model and Score process boxes in IBM Campaign.

After you install IBM Campaign version 9.1.0 or later, the legacy Model and Score process boxes are unconfigured in existing flowcharts. You must manually re-create the predictive models either by using IBM SPSS® Modeler Advantage Enterprise Marketing Management Edition or by using the SPSS Model process box.

Note: Legacy Model and Score process boxes are inactive and cannot be run. However, they can be viewed for configuration details.

For more information, see the *IBM Campaign and IBM SPSS Modeler Advantage Enterprise Marketing Management Edition Integration Guide*.

Campaign installation worksheet

Use the Campaign installation worksheet to gather information about the Campaign database and about other IBM Marketing Software products that are required for the installation of Campaign.

Use the following table to gather information about the empty database that was created for the IBM Campaign system tables. The empty database that you set up for IBM Campaign can have any name. If you use `UA_SYSTEM_TABLES` as the ODBC name for that database, the IBM Campaign system tables will be mapped automatically.

Table 2. Information about the database

Field	Notes®
Database type	
Database name	
Database account user name	
Database account password	
JNDI name	UnicaPlatformDS
ODBC name	

If you are installing on UNIX, obtain the information that is given in the following table:

Table 3. Information about the database for UNIX

Database information	Notes
If your database type is one of the following types, record the database installation directory: <ul style="list-style-type: none">• DB2®• Oracle	
For all database types when Campaign is installed on the Solaris, Linux, or AIX® operating systems, record the location of the lib directory in the database installation.	

You can use the information in Table 3 when you edit the setenv.sh file during the installation and configuration process.

Checklist for IBM Marketing Platform database

The installation wizards for each IBM Marketing Software product must be able to communicate with the Marketing Platform system table database to register the product. Each time that you run the installer, you must enter the following database connection information for the Marketing Platform system table database:

- JDBC connection URL
- Database host name
- Database port
- Database name or schema ID
- User name and password for the database account

Checklist for IBM Marketing Platform deployment on the web application server

Obtain the following information before deploying Marketing Platform:

- Protocol: HTTP, or HTTPS if SSL is implemented in the web application server.
- Host: The name of the machine on which the Marketing Platform will be deployed.
- Port: The port on which the web application server listens.

- **Domain name:** The company domain of each machine where IBM products are installed. For example, example.com. All IBM products must be installed in the same company domain, and you must enter the domain name in all lower case letters.

If there is a mismatch in domain name entries, you might encounter problems when you attempt to use Marketing Platform features or navigate among products. You can change the domain name after the products are deployed by logging in and changing values of the relevant configuration properties in the product navigation categories on the **Settings > Configuration** page.

Checklist for enabling Marketing Platform utilities

If you plan to use the Marketing Platform utilities, obtain the following JDBC connection information before you start to install the Marketing Platform.

- **Path to the JRE.** The default value is the path to the 1.8 version of the JRE that the installer places under your IBM installation directory.
You can accept this default or specify a different path. If you specify a different path, you must point to the 1.8 version of the Sun JRE.
- **JDBC driver class.** The installer automatically provides this, based on the database type that you specify in the installer.
- **JDBC connection URL.** The installer provides the basic syntax, including the host name, database name, and port. You can customize the URL by adding additional parameters.
- **JDBC driver classpath** on your system.

Information about the web component

Obtain the following information for all IBM Marketing Software products that have a web component, which you deploy on a web application server:

- The name of the systems on which the web application servers are installed. You can have one or several web application servers, depending on the IBM Marketing Software environment that you are setting up.
- The port on which the application server listens. If you plan to implement SSL, obtain the SSL port.
- The network domain for your deployment system. For example, mycompany.com.

IBM Site ID

If you are installing an IBM Marketing Software product in one of the countries listed on the Country for Install screen of your product installer, you must enter your IBM Site ID in the space provided. Your IBM Site ID can be found on one of the following documents:

- IBM Welcome letter
- Tech Support Welcome letter
- Proof of Entitlement letter
- Other communications that are sent when you purchase your software

IBM might use the data that is provided by the installed software to better understand how customers use our products and to improve customer support. The data that is gathered does not include any information that identifies individuals. Complete the following actions if you do not want to have such information to be collected:

1. After Marketing Platform is installed, log on to Marketing Platform as a user with administration privileges.
2. Go to **Settings > Configuration**, and set the **Disable Page Tagging** property under the **Platform** category to True.

Installation order for IBM Marketing Software products

When you install or upgrade multiple IBM Marketing Software products, you must install them in a specific order.

The following table provides information about the order that you must follow when you install or upgrade multiple IBM Marketing Software products.

Table 4. Installation or upgrade order for IBM Marketing Software products

For this product or combination:	Install or upgrade in this order:
Campaign (with or without eMessage)	<ol style="list-style-type: none"> 1. Marketing Platform 2. Campaign <p>Note: eMessage is installed automatically when you install Campaign. However, eMessage is not configured or enabled during the Campaign installation process.</p>
Interact	<ol style="list-style-type: none"> 1. Marketing Platform 2. Campaign 3. Interact Design Time environment 4. Interact Run Time environment 5. Interact Extreme Scale Server <p>If you want to install or upgrade the Interact design time environment only, then install or upgrade the Interact design time environment in the following order:</p> <ol style="list-style-type: none"> 1. Marketing Platform 2. Campaign 3. Interact Design Time environment <p>If you want to install or upgrade the Interact runtime environment only, then install or upgrade the Interact runtime environment in the following order:</p> <ol style="list-style-type: none"> 1. Marketing Platform 2. Interact Run Time environment <p>If you want to install the Interact Extreme Scale Server only, then install the Interact Extreme Scale Server in the following order:</p> <ol style="list-style-type: none"> 1. Marketing Platform 2. Interact Run Time environment 3. Interact Extreme Scale Server
Marketing Operations	<ol style="list-style-type: none"> 1. Marketing Platform 2. Marketing Operations <p>Note: If you are integrating Marketing Operations with Campaign, you must also install Campaign. The order of installation for those two products does not matter.</p>
Contact Optimization	<ol style="list-style-type: none"> 1. Marketing Platform 2. Campaign 3. Contact Optimization

Table 4. Installation or upgrade order for IBM Marketing Software products (continued)

For this product or combination:	Install or upgrade in this order:
Interact Advanced Patterns	<ol style="list-style-type: none"> 1. Marketing Platform 2. Campaign 3. Interact 4. Interact Advanced Patterns
IBM SPSS Modeler Advantage Enterprise Marketing Management Edition	<ol style="list-style-type: none"> 1. IBM SPSS Modeler Advantage Enterprise Marketing Management Edition

Planning failover configuration of Campaign

When you install Campaign, you can configure your environment to address any failures of the web application server (front-end) or listener server (back-end).

Definition of terms

Required component	Installation option	Definition
Campaign web application server (front end)	J2EE Application	The web application that provides the Campaign user interface.
Campaign listener server, sometimes called analytic server (back end)	Campaign Server	The startup scripts and components that support the design and execution of flowcharts.

Campaign web application server clustering

You can deploy the Campaign web application server in a clustered environment to support load balancing.

Before you begin the installation, be sure to read Chapter 12, “IBM Campaign web application clustering,” on page 115.

Campaign listener server clustering

Listener (analytic) server clustering, also referred to as back-end clustering, involves the use of multiple listeners for the purpose of high availability and load balancing. Listener clustering is important because flowchart runs happen at the back end.

Before you begin the installation, be sure to read Chapter 13, “Installing a clustered listener environment,” on page 121.

Chapter 3. Preparing data sources for Campaign

Campaign uses data sources to store user and interaction data.

Note: Information about using web application clustering in IBM Campaign is provided in Chapter 12, “IBM Campaign web application clustering,” on page 115.

Complete the following steps to prepare data sources for Campaign.

1. Create a database or a database schema for the Campaign system tables. The empty database that you set up for IBM Campaign can have any name.

Note: Campaign requires customer tables, which should already exist.

2. Create a database user account.

The database user account must have the following rights: CREATE, DELETE, DROP, INSERT, SELECT, and UPDATE.

3. Create ODBC or native connections. If you use UA_SYSTEM_TABLES as the ODBC name for that database, the IBM Campaign system tables will be mapped automatically.
4. Configure the web application server for your JDBC driver.
5. Create JDBC data sources in the web application server.

Creating database or schema for Campaign system tables

The IBM Campaign system tables hold application data. The empty database that you set up for the system tables can have any name. To ensure that the system tables will be mapped automatically, the ODBC name for that database must be UA_SYSTEM_TABLES. If you have multiple partitions, create data sources for each partition where you want to enable IBM Campaign.

Complete the following steps to create a database or database schema for Campaign system tables:

1. Create a system user account for later in the installation process.

Note: The system user account must have the following rights: CREATE, DELETE, DROP, INSERT, SELECT, and UPDATE.

2. Print the “Campaign installation worksheet” on page 9. Obtain the information about your database or schema and database account, and then note this information in the checklist. Continue to fill in the checklist as you perform the rest of the steps in this section, so that you can use this information later in the installation process.

Note: The IBM eMessage system tables are created in the Campaign schema during the installation. However, eMessage is not enabled.

Creating ODBC or native connections

Create ODBC or native connections so that the Campaign server can access the Campaign database. Create ODBC or native connections on the computer where you installed the Campaign server.

The computer on which the Campaign server is installed must have an ODBC or native connection to the following databases:

- The database or schema that contains the Campaign system tables
- The database or schema that contains the customer tables (sometimes called user tables).

Note: If you are configuring just one partition, name the data source UA_SYSTEM_TABLES. If you are configuring more than one partition, do not name any of the data sources UA_SYSTEM_TABLES. For more information, refer *Configuring multiple partitions for Campaign*.

Use the following guidelines to create these ODBC or native connections:

- **For databases on UNIX/SUSE:** Create native connections for DB2 and Oracle databases or ODBC connections for other databases, such as Teradata, Netezza®. The procedure for creating a native data source differs depending on the data source type and the version of UNIX. See the data source and operating system documentation for specific ODBC driver installation and configuration instructions.
- **For databases on Windows:** Create a new ODBC connection in the Control Panel **Administrative Tools > Data Sources (ODBC)** section.
- **For DB2 databases on z/OS®:** The database user account must have SYSADM privileges. The DBADM privilege is not sufficient because it does not provide access to all of the required system objects. Without SYSADM privileges, flowcharts might not run successfully.

Note: Campaign supports certain versions of DB2 on z/OS for customer tables only, not system tables. For more information, see the *Recommended Software Environments and Minimum System Requirements*.

Be sure to record the ODBC name in the “Campaign installation worksheet” on page 9.

Configuring the web application server for your JDBC driver

The web application server where you plan to deploy Campaign must contain the correct JAR file to support JDBC connections. This enables your web application to connect to the system tables. The location of the JAR file must be included in the class path of the web application server.

If you are using WebSphere, Weblogic, Tomcat and if you allow the installer to create the data source when you install the product, you can skip the step in this procedure about setting the classpath. To enable automatic data source creation during installation, on the Datasource Creation panel, select the **Create Campaign Datasource** check box and supply information about your application server. If you do not want to allow the installer to create the data source, then perform this entire procedure.

If you are using WebLogic application server you must perform this entire procedure.

1. Consult the *Recommended Software Environments and Minimum System Requirements* document to determine the correct JAR file for your database type.
2. Obtain the latest vendor-provided Type 4 JDBC driver that is supported by IBM Marketing Software.

- If the driver does not exist on the machine where you plan to deploy Campaign, copy the driver to any location on that machine. Unpack the driver in a path that does not include spaces.
 - If you obtain the driver from a machine where the data source client is installed, verify that the version is supported.
3. Include the full path and file name of the driver in the classpath of the web application server where you plan to deploy Campaign.
 - For all supported versions of WebLogic, set the classpath in the **setDomainEnv** script in the *WebLogic_domain_directory/bin* directory where environment variables are configured.

The driver must be the first entry in the CLASSPATH list of values before any existing values. For example:

UNIX

```
CLASSPATH="/home/oracle/product/<version>/jdbc/lib/ojdbc7.jar:
${PRE_CLASSPATH}${CLASSPATHSEP}${WEBLOGIC_CLASSPATH}
${CLASSPATHSEP}${POST_CLASSPATH}${CLASSPATHSEP}${WLP_POST_CLASSPATH}"
export CLASSPATH
```

Windows

```
set CLASSPATH=c:\oracle\jdbc\lib\ojdbc7.jar;%PRE_CLASSPATH%;
%WEBLOGIC_CLASSPATH%;%POST_CLASSPATH%;%WLP_POST_CLASSPATH%
```

- For all supported versions of WebSphere, set the classpath in the Administration console when you set up the JDBC providers for the IBM Marketing Software products.
 - Please make a note of the database driver class path in the Marketing Campaign installation worksheet, as you must enter the path when you run the installer.
4. Restart the web application server so your changes take effect.

During startup, monitor the console log to confirm that the classpath contains the path to the database driver.

Datasource creation during installation is not supported for JBOSS application Server, so you must perform this entire procedure.

1. Obtain the latest vendor-provided Type 4 JDBC driver for your system table database that is supported by IBM Marketing Software, as described in the *Recommended Software Environments and Minimum System Requirements* guide. Use the following guidelines after you obtain the JDBC driver.
 - If the driver does not exist on the server where you plan to deploy Campaign, obtain it and unpack it on the server. Unpack the drivers in a path that does not include spaces.
 - If you obtain the driver from a server where the data source client is installed, verify that the version is the latest supported by Campaign.
2. Add and register the full path of the driver including the file name, to the module using below commands:
 - For all supported versions of JBoss, add the JDBC driver as module. Use the following procedure to add the JDBC driver as a module.
For example for SQL Server:


```
module add
--name=com.microsoft.sqlserver.jdbc
--resources=<JDBC_Driver_Location>\sqljdbc4.jar
--dependencies=javax.api,javax.transaction.api
```
 - Register this SQL JDBC Driver using the following guidelines: For example:


```
- /subsystem=datasources/jdbc-driver=sql:add(driver-module-
name=com.microsoft.sqlserver.jdbc,driver-name=sql,driver-xa-
```

- ```
datasource-class-
name=com.microsoft.sqlserver.jdbc.SQLServerXADataSource)
- /subsystem=datasources/jdbc-driver=sql:read-resource
- /subsystem=ee/service=default-bindings:write-
attribute(name=datasource, value=undefined)
```
3. Make a note of the database driver class path in the Campaign installation worksheet, as you must enter the path when you run the installer.
  4. Restart the web application server so that your changes take effect.  
During startup, monitor the console log to confirm that the class path contains the path to the database driver.

**If you are using Apache Tomcat, you must perform this entire procedure.**

1. Obtain the latest vendor-provided Type 4 JDBC driver for your system table database that is supported by IBM Marketing Software, as described in the *Recommended Software Environments and Minimum System Requirements* guide. Use the following guidelines after you obtain the JDBC driver.
  - If the driver does not exist on the server where you plan to deploy Campaign, obtain it and unpack it on the server. Unpack the drivers in a path that does not include spaces.
  - If you obtain the driver from a server where the data source client is installed, verify that the version is the latest supported by Campaign.
2. Copy the jar file manually to *\$Tomcat\_Installed\_Location/lib/*, here *\$Tomcat\_Installed\_Location* is the Tomcat root installation directory where you plan to deploy Campaign.
3. Make a note of the database driver class path in the Campaign installation worksheet, as you must enter the path when you run the installer.
4. Restart the web application server so that your changes take effect.  
During startup, monitor the console log to confirm that the class path contains the path to the database driver.

---

## Creating JDBC connections in web application server

Create JDBC connections on each of the web application servers where Campaign is deployed. Campaign uses the JDBC connections to access the required databases.

Use the following list to determine which JDBC connections you must create for the Campaign web application. The steps include the suggested and required JNDI names.

For more information about creating JDBC connections in your web application server, see the WebLogic or WebSphere documentation.

1. Create a connection to the database holding the Marketing Platform system tables. Use UnicaPlatformDS as the JNDI name.

**Important:** UnicaPlatformDS is the required JNDI name.

If you are deploying Campaign in the same JVM as your Marketing Platform deployment, you have probably already set up this connection.

**Note:** If you are using WebLogic to configure a connection to an Oracle or DB2 database, you must also specify the database user name in the form *user=<DBUser>* in the Properties section on the Connection Pool tab.

2. Create a connection to the database holding the Campaign system tables.

You can allow the installer to create this data source (for one partition only) when you install Campaign. To enable automatic data source creation during installation, on the Datasource Creation panel, select the **Create Campaign Datasource** check box and supply information about your application server.

- The installer creates the data source using campaignPartition1DS as the JNDI name.
- Note that, if you are using WebLogic, you must add the JDBC driver to your classpath manually even if you allow the installer to create the data source. The installer does this automatically for WebSphere.
- Automatic DataSource creation during installation is not supported for the JBOSS application server.

If you do not want to allow the installer to create the data source, then you must create this data source.

If there is one partition, the suggested JNDI name is campaignPartition1DS.

If there are multiple partitions,, a best practice is to use campaignPartition1DS for the first connection, campaignPartition2DS for the second, and so on.

**Note:** This practice is provided as an example. You can specify any JNDI names for Campaign system table connections.

3. Record all JNDI names in the “Campaign installation worksheet” on page 9

## Information for creating JDBC connections

Use default values when you create JDBC connections if specific values are not provided. For more information, see the application server documentation.

**Note:** If you are not using the default port setting for your database, make sure that you change it to the correct value.

### WebLogic

Use the following values if your application server is WebLogic:

#### SQLServer

- Database Driver: Microsoft MS SQL Server Driver (Type 4) Versions: 2012, 2012 SP1 and SP3, 2014, 2014 SP1, 2016 SP1
- Default port: 1433
- Driver class: com.microsoft.sqlserver.jdbc.SQLServerDriver
- Driver URL: jdbc:sqlserver://<your\_db\_host>[\\<named\_instance>]:<your\_db\_port>;databaseName=<your\_db\_name>
- Properties: Add user=<your\_db\_user\_name>

#### Oracle

- Driver: Other
- Default port: 1521
- Driver class: oracle.jdbc.OracleDriver
- Driver URL:  
jdbc:oracle:thin:@<your\_db\_host>:<your\_db\_port>:<your\_db\_service\_name>  
Enter the driver URL by using the format that is shown. IBM Marketing Software applications do not allow the use of Oracle's RAC (Real Application Cluster) format for JDBC connections.
- Properties: Add user=<your\_db\_user\_name>

## DB2

- Driver: Other
- Default port: 50000
- Driver class: `com.ibm.db2.jcc.DB2Driver`
- Driver URL: `jdbc:db2://<your_db_host>:<your_db_port>/<your_db_name>`
- Properties: Add `user=<your_db_user_name>`

## WebSphere

Use the following values if your application server is WebSphere:

### SQLServer

- Driver: N/A
- Default port: 1433
- Driver class:  
`com.microsoft.sqlserver.jdbc.SQLServerConnectionPoolDataSource`
- Driver URL: `jdbc:sqlserver://<DBHostName>:1433;databaseName=<DBName>`

In the **Database Type** field, select **User-defined**.

After you create the JDBC Provider and data source, go to the **Custom Properties** for the data source, and add, modify properties as follows.

- `serverName=<your_SQL_server_name>`
- `portNumber =<SQL_Server_Port_Number>`
- `databaseName=<your_database_name>`

Add the following custom property:

**Name:** `webSphereDefaultIsolationLevel`

**Value:** 1

**Datatype:** Integer

### Oracle

- Driver: Oracle JDBC Driver
- Default port: 1521
- Driver class: `oracle.jdbc.OracleDriver`
- Driver URL:  
`jdbc:oracle:thin:@<your_db_host>:<your_db_port>:<your_db_service_name>`  
Enter the driver URL by using the format that is shown. IBM Marketing Software applications do not allow the use of Oracle's RAC (Real Application Cluster) format for JDBC connections.

## DB2

- Driver: JCC Driver Provider
- Default port: 50000
- Driver class: `com.ibm.db2.jcc.DB2Driver`
- Driver URL: `jdbc:db2://<your_db_host>:<your_db_port>/<your_db_name>`

Add the following custom property:



**Name:** webSphereDefaultIsolationLevel

**Value:** 2

**Datatype:** Integer

---

## Preparing DB2 on z/OS user datasource for Campaign

Follow these guidelines to use a DB2 database on z/OS as a data source for Campaign.

Campaign supports DB2 10.1 and 11 on z/OS for customer tables only, not system tables. For information about supported versions, see the *Recommended Software Environments and Minimum System Requirements*.

1. The database user account must have SYSADM privileges. The DBADM privilege is not sufficient because it does not provide access to all of the required system objects. Without SYSADM privileges, flowcharts might not run successfully.
2. Database loaders are supported. However, there is not an out-of-the-box approach for loaders on z/OS. You must set up USS Pipes on z/OS, write a stored procedure to invoke the DSNUTILU and a script to invoke the stored procedure, then configure the loader to invoke the script. For details, see the *IBM Campaign Administrator's Guide*.
3. Set **DB2NotLoggedInitially** and **DB2NotLoggedInitiallyUserTables** to FALSE, under Campaign|partitions|partition[n]|dataSources|datasourcename. For details, see the *IBM Campaign Administrator's Guide*.

---

## Preparing Vertica datasources for Campaign

Follow these guidelines to use a Vertica database as a data source for Campaign.

For information about supported versions, see the *Recommended Software Environments and Minimum System Requirements*.

To download the latest Vertica documentation in zip or tar format:  
[http://www.vertica.com/v-zone/product\\_documentation](http://www.vertica.com/v-zone/product_documentation)

1. Install the HP Vertica client driver on the server where the Campaign analytical server (listener) is installed.  
Obtain the driver here: <http://www.vertica.com/resources/vertica-client-drivers/>

2. Configure the `odbc.ini` file.

Sample `odbc.ini` file entry

[ODBC Data Sources]

```
DriverManagerEncoding=UTF-16
ODBCInstLib=/usr/local/lib/libodbcinst.so
ErrorMessagesPath=/opt/vertica/lib64
LogLevel=4
LogPath=/tmp
```

[Vertica]

```
Description = VerticaDSN ODBC driver
Driver = /opt/vertica/lib64/libverticaodbc.so
Database = exampleDB
```

```
Servename = example.com
UserName = dbadmin
Password = example_password
Port = 5433
```

---

## Preparing Amazon Redshift user data sources for Campaign

Follow these guidelines to use Amazon Redshift as a data source for Campaign.

Campaign supports Amazon Redshift for customer tables. Redshift is not supported for system tables.

Use the Amazon native ODBC driver for connectivity and configure the PostgreSQL template in Campaign for each datasource. See the *IBM Campaign Administrator's Guide* for important information about configuring Campaign|partitions|partition[n]|dataSources|PostgreSQL. Some properties do not have default values, so you must supply them. Other properties might need to be adjusted for your specific configuration. For information about supported versions, see the *Recommended Software Environments and Minimum System Requirements*. For information about Amazon Redshift, see <http://aws.amazon.com/redshift/>

1. Download and configure the Amazon Redshift ODBC driver for Redshift.  
Driver Name : Amazon Redshift ODBC driver (version - 1.4.3.1000-1 – 64-bit) ,  
Unix ODBC driver manager (Version - 2.3.7). See <https://docs.aws.amazon.com/redshift/latest/mgmt/install-odbc-driver-linux.html>
2. Validate connection and authentication details: Use `isql` to test the connection on Linux. Use the Microsoft Windows ODBC Driver Manager for Windows.
3. Import PostgreSQLTemplate.xml using configTool (a utility supplied with IBM Marketing Platform).
4. Fill in the newly added PostgreSQL template with the user database details.
5. Configure the Amazon Redshift user database.
6. Restart the application server and the IBM Campaign listener server.

Follow the steps below to test the connection with Amazon Redshift.

**Note:** See the following link to get the cluster connection string:  
<http://docs.aws.amazon.com/redshift/latest/mgmt/configuring-connections.html#connecting-drivers>

On Windows:

1. Open Windows ODBC Driver Manager.
2. Add a System DSN and select the Amazon Native driver.
3. Provide the connection string and other details.
4. Test the connection using `isql` or the IBM Campaign CxnTest utility.

On Linux:

1. Configure LD\_LIBRARY\_PATH (or LIBPATH on AIX) to include the location of the driver library files.
2. Configure ODBCINI to point to the `odbc.ini` file location.
3. Configure the `odbc.ini` file for the Redshift user database. For example:

```
[ODBC]
Trace = yes
```

```

TraceFile = /tmp/OdbcTrace.log
#[ODBC Data Sources]
AmazonRedshift=Amazon Redshift

[AmazonRedshift]
This key is not necessary and is only to give a description of the data source.
driver_name=Amazon_Redshift
Description=Amazon Redshift ODBC Driver (64-bit) DSN

#[Driver: The location where the ODBC driver is installed to.]
Driver=/opt/amazon/redshiftodbc/lib/64/libamazonredshiftodbc64.so

Required: These values can also be specified in the connection string.
Server==[change-me]
Port=5439
Database==[change-me]
locale=en-US
UID==[change-me]
PWD==[change-me]

```

4. Test the connection using `isql` or the IBM Campaign CxnTest utility.

---

## Preparing Apache Hadoop Hive user data sources for Campaign

Follow the steps below to enable Hive-based Apache Hadoop data sources to be used in IBM Campaign

Campaign supports Apache Hadoop Hive for customer tables only, not system tables. For details about supported versions, see the *Recommended Software Environments and Minimum System Requirements*.

| Task                                                                                              | Description                                                                                                                                                                                                                |
|---------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| "A. Install a Hive ODBC driver" on page 26                                                        | You can install the DataDirect Hive ODBC Driver from Progress Software, the Cloudera Hive ODBC driver from Cloudera, Inc., or the Hortonworks Hive ODBC driver from Hortonworks, Inc.                                      |
| "B. Configure the Hive ODBC Driver" on page 28                                                    | Configuration includes modifying .ini files and setting path values and environment variables. Be sure to follow the appropriate instructions for the driver that you installed.                                           |
| "C. Map existing HBase tables to Hive" on page 31 (OPTIONAL)                                      | This step is required only if you have existing tables that were created in Apache HBase.                                                                                                                                  |
| "D. Import and configure the BigDataODBCHiveTemplate data source template in Campaign" on page 32 | Use the configTool utility to import the template BigDataODBCHive.xml into Campaign.<br><br>Then go to Campaign partitions partition[n] dataSources and create a datasource based on the imported BigDataODBCHiveTemplate. |
| "E. Configure SSH on the Campaign listener server" on page 34                                     | To enable data file transfers between the IBM Campaign listener (analytic) server and the Hive-based Hadoop big data instance, you must configure SCP and SSH seamless login.                                              |

| Task                                                | Description                                                                                |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------|
| "F. Map the Hive datasource in Campaign" on page 35 | Mapping user tables is the process of making external data sources accessible in Campaign. |

You can now use your Hive-based Hadoop big data system as a user data source for IBM Campaign. For example, create a marketing campaign flowchart that uses customer account data from your big data instance to target customers with specific account types and balances.

## Requirements and restrictions for Hive-based Hadoop data sources

The following requirements and restrictions apply to using Hive-based Hadoop data sources with IBM Campaign.

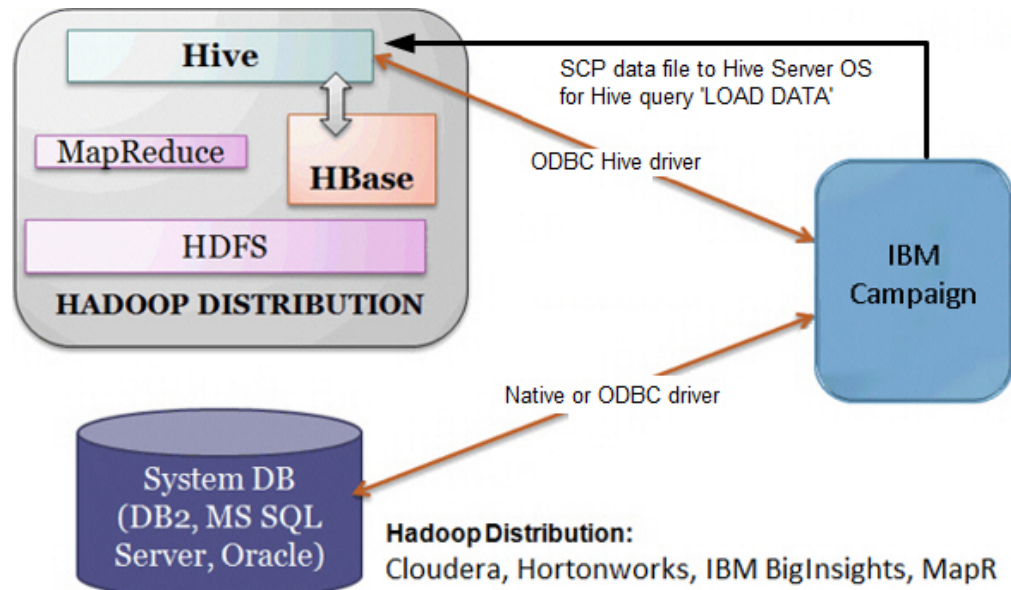
- One of the following drivers is required. The customer is responsible for obtaining the driver.
  - DataDirect Apache Hive ODBC driver from Progress.com: DataDirect Connect64(R) for ODBC Release 7.1.5.
  - Cloudera ODBC Driver for Apache Hive Version 2.5.16 for Apache Hadoop Hive from Cloudera, Inc.
  - Hortonworks 64-bit ODBC driver for HDP 2.3 (v2.0.5) available at <http://hortonworks.com/hdp/addons/>.
- The following Hadoop distributions are supported, with Apache Hive as the connection point: Cloudera, Hortonworks, IBM BigInsights<sup>®</sup> <sup>™</sup>, MapR.
- Minimum supported Hive version: 0.14
- The big data integration is currently supported on Linux RHEL 6.3 or higher.
- Hive-based Hadoop is supported as a user data source only. It is not supported for Campaign system tables.
- The big data integration does not currently support the IBM Campaign Cube, Optimize, or Interact List process boxes or eMessage Landing Pages in an Extract process box.

## Overview of the Campaign integration with Hadoop

Integrating Campaign with Hadoop data sources demands extra software that is not part of the standard Campaign installation requirement. You must be familiar with the terminology and structure of the integration, and be able to access various resources that are available to support the integration.

### Integration architecture: Hive-based Hadoop and Campaign

This diagram shows the integration architecture for IBM Campaign and Hive-based Hadoop user data sources.



The Apache Hive data warehouse software facilitates querying and managing large datasets that reside in distributed storage. Built on top of Apache Hadoop, Hive provides:

- Tools to enable easy data extract/transform/load (ETL)
- A mechanism to impose structure on a variety of data formats
- Access to files stored either directly in Apache HDFS or in other data storage systems such as Apache HBase
- Query execution via MapReduce

Hive defines a simple SQL-like query language, called HiveQL (or HQL), that enables users familiar with SQL to query the data.

You can use the Hue editor (Hadoop UI) to work with your big data instance (for example: connect, view, create tables and databases).

## Terminology and resources for Hive Hadoop

To understand the Hive Hadoop integration with IBM Campaign, refer to this list of definitions and links to online resources.

### Terms

**Apache Hadoop®** is an open-source software framework written in Java for distributed storage and distributed processing of very large data sets on computer clusters built from commodity hardware.

**Apache Hive™** is a data warehouse infrastructure built on top of Hadoop to facilitate querying and managing large datasets that reside in distributed storage. Hive provides a mechanism to project structure onto this data and query the data using an SQL-like language called HiveQL.

**Apache HBase™** is an open source, non-relational, distributed database written in Java. It runs on top of HDFS, providing BigTable-like capabilities for Hadoop.

**Hadoop Distributed File System (HDFS™)** is a distributed file system that runs on commodity hardware. It is designed to reliably store very large files across machines in a large cluster.

**Hue** is a Web interface for analyzing data with Apache Hadoop.

**HiveQL** (or HQL) is the Hive query language.

**MapReduce** is a programming model and an associated implementation for processing and generating large data sets with a parallel, distributed algorithm on a cluster. MapReduce is the heart of Hadoop®. It is this programming paradigm that allows for massive scalability across hundreds or thousands of servers in a Hadoop cluster.

Big data **distributions** of Apache Hadoop: A number of vendors have developed their own distributions of Hadoop, including **Cloudera**, **Hortonworks**, **IBM BigInsights**, and **MapR**.

**User tables** is an IBM Campaign term that indicates any data source that contains an organization's marketing data for access by IBM Campaign flowcharts. Typically, user tables contain data about customers, prospects, and products. For example, customer account data pulled from user tables could be used in a flowchart to target customers with specific account types and balances.

### Online resources

The following links were accurate at the time of publication but are subject to change:

- Cloudera ODBC driver: <https://www.progress.com/products/data-sources/apache-hadoop-hive>
- Hive ODBC driver: <https://www.progress.com/products/data-sources/apache-hadoop-hive>
- Hive: <https://cwiki.apache.org/confluence/display/Hive/Home>
- HiveQL: <https://cwiki.apache.org/confluence/display/Hive/LanguageManual>
- Hive HBase integration: <https://cwiki.apache.org/confluence/display/Hive/HBaseIntegration>
- Hue and Hadoop: <http://gethue.com>
- IBM Campaign: [http://www.ibm.com/support/knowledgecenter/SSCVKV/product\\_welcome\\_kc\\_campaign.dita](http://www.ibm.com/support/knowledgecenter/SSCVKV/product_welcome_kc_campaign.dita)

## A. Install a Hive ODBC driver

This is the first step to integrate IBM Campaign with Hive-based Apache Hadoop data sources. You must install either the DataDirect Hive ODBC Driver from Progress Software or the Cloudera Hive ODBC driver from Cloudera, Inc.

You can install any one of the Hive ODBC drivers that are described below. Be sure to follow the appropriate instructions for the driver that you select.

Install the Hive ODBC driver on the system where you have installed the Campaign Analytical Server. If the Analytical Server is installed in a cluster, install the Hive ODBC driver on each system.

## Installing the DataDirect Hive ODBC driver

This topic explains how to install the DataDirect driver for Apache Hive, which is a fully compliant ODBC driver that supports multiple Hadoop distributions.

KornShell (ksh) must be installed on the IBM Campaign listener (analytic) server.

1. Obtain the Progress DataDirect Connect ODBC driver (Progress DataDirect Connect64(R) for ODBC Release 7.1.5) for Apache Hadoop Hive:  
<https://www.progress.com/products/data-sources/apache-hadoop-hive>
2. Download and install the DataDirect Hive driver on the IBM Campaign listener (analytic server):

**PROGRESS\_DATADIRECT\_CONNECT64\_ODBC\_7.1.5\_LINUX\_64.tar.Z**

```
[DataDirectNew]# gunzip
```

```
PROGRESS_DATADIRECT_CONNECT64_ODBC_7.1.5_LINUX_64.tar.Z
```

```
[DataDirectNew]# tar -xvf
```

```
PROGRESS_DATADIRECT_CONNECT64_ODBC_7.1.5_LINUX_64.tar
```

3. Run the following command to start the installation:

```
>> ksh ./unixmi.ksh
```

4. Follow the prompts to complete the installation.

5. Perform basic testing of the driver:

```
>> ./ddtestlib /opt/Progress/DataDirect/Connect64_for_ODBC_71/lib/
ddhive27.so
```

See “Configuring the DataDirect Hive ODBC driver” on page 28.

## Installing the Cloudera ODBC driver

This topic explains how to install the Cloudera ODBC driver, which is a fully compliant ODBC driver that supports multiple Hadoop distributions.

1. Obtain the Cloudera ODBC Driver for Apache Hive Version 2.5.16 for Apache Hadoop Hive at <http://www.cloudera.com/content/cloudera/en/downloads/connectors/hive/odbc/hive-odbc-v2-5-16.html>. Download and save the Cloudera Hive ODBC driver on the IBM Campaign listener (analytic server). The default location is /opt/cloudera/.

2. Install and configure the Cloudera driver. Installation instructions are downloaded to where you install the driver. By default, the instructions are downloaded to /opt/cloudera/hiveodbc/Cloudera-ODBC-Driver-for-Apache-Hive-Install-Guide.pdf. The PDF file is also available online at [www.cloudera.com](http://www.cloudera.com).

3. Download and install UnixODBC Manager, unixODBC-2.3.2 on the IBM Campaign listener (analytic server). Download UnixODBC Manager from the unixODBC Project at <http://www.unixodbc.org/>. The UnixODBC Manager is available in a gzip, tar format.

The steps to install UnixODBC are available at <http://www.unixodbc.org/download.html>

To create a unixODBC directory that contains the required source files:

- a. Copy the unixODBC-2.3.2.tar.gz file to the Campaign listener (analytic server).
  - b. `gunzip unixODBC-2.3.2.tar.gz`
  - c. `tar xvf unixODBC-2.3.2.tar`
4. On the Campaign listener, run `rpm -ivh ClouderaHiveODBC-2.5.16.1005-1.el6.x86_64.rpm`.

See “Configuring the Cloudera ODBC driver” on page 29.

### **Installing the Hortonworks Hive ODBC driver**

This topic explains how to install the Hortonworks driver for Apache Hive, which is a fully compliant ODBC driver that supports multiple Hadoop distributions.

KornShell (ksh) must be installed on the IBM Campaign listener (analytic) server.

1. Obtain the 64-bit version of the Hortonworks Hive ODBC driver:  
<http://hortonworks.com/hdp/addons>
2. Install the Hortonworks ODBC driver on the IBM Campaign listener (analytic server):

```
rpm -ivh hive-odbc-native-2.0.5.1005-1.el6.x86_64.rpm
```

The default installation location of the Hortonworks Hive ODBC driver is  
`/usr/lib/hive/lib/native`.

For more information on installing the Hortonworks Hive ODBC driver see:  
<http://hortonworks.com/wp-content/uploads/2015/10/Hortonworks-Hive-ODBC-Driver-User-Guide.pdf>.

3. Follow the prompts to complete the installation.

See “Configuring the Hortonworks Hive ODBC driver” on page 30.

## **B. Configure the Hive ODBC Driver**

This is the second step to integrate Campaign with Hive-based Apache Hadoop data sources. You must configure the driver that you installed (either the DataDirect Hive ODBC Driver from Progress Software or the Cloudera Hive ODBC driver from Cloudera, Inc.).

### **Configuring the DataDirect Hive ODBC driver**

If you are using the DataDirect Hive ODBC driver to integrate IBM Campaign with a big data source, you must configure the driver after installing it.

Complete the DataDirect driver installation that is described in “Installing the DataDirect Hive ODBC driver” on page 27.

1. Modify the ODBC.ini file to specify the Hive server information, using the values shown in the following example. Be sure to customize the items in **bold** to match your own configuration.



```
[MapRHive]
Driver=/opt/Progress/DataDirect/Connect64_for_ODBC_71/lib/ddhive27.so
Description=DataDirect 7.1 Apache Hive Wire Protocol
ArraySize=16384
Database=<database-name>
DefaultLongDataBuffLen=1024
EnableDescribeParam=0
HostName=<hostname or ip of Hive server on Hadoop Distribution machine>
LoginTimeout=30
LoginID=<username of Hadoop Distribution machine>
MaxVarcharSize=2147483647
Password=<password of Hadoop Distribution machine>
PortNumber=<port number of Hive server on Hadoop Distribution machine>
RemoveColumnQualifiers=1
StringDescribeType=12
TransactionMode=0
UseCurrentSchema=0
WireProtocolVersion=0
GetTablesWithQuery=1
```

2. Assuming that your ODBC driver is installed in the following location:

```
/opt/Progress/DataDirect/Connect64_for_ODBC_71
```

Make the following changes:

- Ensure that your LD\_LIBRARY\_PATH includes the following path:  
=/opt/Progress/DataDirect/Connect64\_for\_ODBC\_71/lib
- Ensure that your PATH includes the following path: =/opt/Progress/  
DataDirect/Connect64\_for\_ODBC\_71/tools
- Set your ODBCINI variable to point to the correct INI file. For example:  
ODBCINI=/opt/Progress/DataDirect/Connect64\_for\_ODBC\_71/odbc.ini ;  
export ODBCINI
- Set your ODBCINST variable to point to the correct INI file. For example:  
ODBCINST=/opt/Progress/DataDirect/Connect64\_for\_ODBC\_71/  
odbcinst.ini ; export ODBCINST

3. Verify the connectivity of the DataDirect ODBC driver and your Hive-based Hadoop big data system:

```
cd /opt/Progress/DataDirect/Connect64_for_ODBC_71/samples/example
>> ./example
```

4. On the Campaign Analytical Server, set the ODBCINI and CAMPAIGN\_HOME environment variables in the Campaign/bin/setenv file. Run the IBM Campaign odbctest utility to verify connectivity to IBM Campaign:

```
cd <Campaign_Home>/bin
>> ./odbctest
```

“C. Map existing HBase tables to Hive” on page 31

## Configuring the Cloudera ODBC driver

If you are using the Cloudera ODBC driver to integrate IBM Campaign with a big data source, you must configure the driver after installing it.

Complete the Cloudera driver installation that is described in “Installing the Cloudera ODBC driver” on page 27.

1. Modify the ODBC.ini file to specify the Hive server information, using the values shown in the following example. Be sure to customize the items in **bold** to match your own configuration.

**Note:** Do not enable Native Query mode in the configuration.

```
[Cloudera_HIVE]
Driver=/opt/cloudera/hiveodbc/lib/64/libclouderahiveodbc64.so
Description=Hive Cloudera ODBC Driver
Host=<hostname or ip of Hive server on Hadoop Distribution machine>
Port=<port number of Hive server on Hadoop Distribution machine>
Schema=<database-name>
ServiceDiscoveryMode=0
ZKNamespace=
HiveServerType=2
AuthMech=3
ThriftTransport=1
UseNativeQuery=0
UID=cloudera
GetTablesWithQuery=1
```

**Note:** The UID is the user ID of the Cloudera cluster database.

2. Assuming that your ODBC driver is installed in the following location:

/opt/cloudera/hiveodbc/lib/64/libclouderahiveodbc64.so

Make the following change in Campaign/bin setenv.sh: Set LD\_LIBRARY\_PATH (LIBPATH for AIX) to include Cloudera LIB and UnixODBC LIB as follows.

- LD\_LIBRARY\_PATH=/opt/IBM/UCD912\_Install/Campaign/bin:/opt/ibm/db2/V10.1/lib64/;/usr/lib:/lib:/opt/cloudera/hiveodbc/lib/64/;/usr/local/lib/  
export LD\_LIBRARY\_PATH
- ODBCINI=/usr/local/etc/odbc.ini export ODBCINI

3. Open /opt/cloudera/hiveodbc/lib/64/cloudera.hiveodbc.ini

During Cloudera installation, a connectivity option is set by default. To connect Cloudera with ODBC, you must disable the connectivity option.

Change the following lines in cloudera.hiveodbc.ini:

```
Generic ODBCInstLib
iODBC
ODBCInstLib=libiodbcinst.so
```

The updated lines appear as:

```
Generic ODBCInstLib
iODBC
ODBCInstLib=libiodbcinst.so
```

“C. Map existing HBase tables to Hive” on page 31

## Configuring the Hortonworks Hive ODBC driver

If you are using the Hortonworks Hive ODBC driver to integrate IBM Campaign with a big data source, you must configure the driver after installing it.

Complete the Hortonworks driver installation that is described in “Installing the Hortonworks Hive ODBC driver” on page 28.

1. Modify the ODBC.ini file to specify the Hive server information, using the values shown in the following example. Be sure to customize the items in **bold** to match your own configuration.

The default location of odbc.ini file is: /usr/lib/hive/lib/native/hiveodbc/Setup/odbc.ini

```
[ODBC] [ODBC Data Sources] HWDS [HWDS] Description=Hortonworks Hive ODBC Driver (64-bit)
DSN Driver=/usr/lib/hive/lib/native/Linux-amd64-64/libhortonworkshiveodbc64.so
HOST=<hostname or ip of Hive server on Hadoop Distribution machine>
PORT=<port number of Hive server on Hadoop Distribution machine>
Schema=<schema-name>
ServiceDiscoveryMode=0
ZKNamespace=
HiveServerType=2
AuthMech=2
ThriftTransport=1
UseNativeQuery=0
UID=hue
KrbServiceName=<Hive Server 2 Kerberos service name>
KrbRealm=<Hive Server 2 Kerberos realm>
SSL=0
TwoWaySSL=0
ClientCert=
ClientPrivateKey=
ClientPrivateKeyPassword=
```

2. Enable the unixODBC manager and disable Generic ODBCInstLib in the hortonworks.hiveodbc.ini as follows.

The default location of hortonworks.hiveodbc.ini is /usr/lib/hive/lib/native/Linux-amd64-64/hortonworks.hiveodbc.ini.

Comment out the following line by adding a hash character (#) at the beginning: # ODBCInstLib=libiodbcinst.so

Un-comment the following line by removing the hash character(#) at the beginning: ODBCInstLib=/usr/local/unixODBC/lib/libiodbcinst.so

3. Edit the Campaign/bin/setenv.sh as follows.
  - a. Ensure that your LD\_LIBRARY\_PATH includes the following path:  
/usr/local/unixODBC/lib
  - b. Set your ODBCINI variable to point to the correct INI file. For example:  
ODBCINI=/usr/lib/hive/lib/native/hiveodbc/Setup/odbc.ini  
export ODBCINI  
HORTONWORKSHIVEINI=/usr/lib/hive/lib/native/Linux-amd64-64/  
hortonworks.hiveodbc.ini  
export HORTONWORKSHIVEINI

“C. Map existing HBase tables to Hive”

## C. Map existing HBase tables to Hive

This is the third step to integrate IBM Campaign with Hive-based Apache Hadoop data sources. If you do not have existing tables that were created in Apache HBase, you can skip this step.

Complete “B. Configure the Hive ODBC Driver” on page 28.

This step is required only if you have existing tables that were created in Apache HBase. In this case, you must make the existing HBase tables available to Apache Hive by running the CREATE EXTERNAL TABLE query. After you expose the HBase tables to Hive, the tables are then available to IBM Campaign for table mapping within Campaign.

The EXTERNAL keyword lets you create a table and provide a LOCATION so that Hive does not use a default location for the table. This is useful if you already have data generated. An EXTERNAL table points to any HDFS location for its storage, rather than being stored in a folder specified by the configuration property. When dropping an EXTERNAL table, data in the table is NOT deleted from the file system.

For information about Hive HBase integration: <https://cwiki.apache.org/confluence/display/Hive/HBaseIntegration>

1. Open the Hue editor and open the Hive Query Editor.
2. Create and execute the CREATE EXTERNAL TABLE command.

Use the following query as an example, substituting your own table name, field names, and other parameters. This example uses 'CampaignAccounts' as the table name and 'f' as the family name. CREATE EXTERNAL TABLE **HiveExt\_CampaignAccounts**(Acct\_ID INT,Indiv\_ID INT,HHold\_ID INT,Acct\_Type\_Code STRING,Acct\_Status\_Code INT,Acct\_Open\_Date INT,Acct\_Balance STRING,Acct\_Balance\_Last\_Month STRING,Acct\_Balance\_Avg\_6Month STRING,Credit\_Limit STRING,Acct\_Number STRING,Last\_Contact\_Date STRING,Due\_Date STRING) STORED BY 'org.apache.hadoop.hive.hbase.HBaseStorageHandler' WITH SERDEPROPERTIES ('hbase.columns.mapping' = ':key,f:Indiv\_ID, f:HHold\_ID, f:Acct\_Type\_Code, f:Acct\_Status\_Code, f:Acct\_Open\_Date, f:Acct\_Balance, f:Acct\_Balance\_Last\_Month, f:Acct\_Balance\_Avg\_6Month, f:Credit\_Limit, f:Acct\_Number, f:Last\_Contact\_Date, f:Due\_Date') TBLPROPERTIES ('hbase.table.name' = '**CampaignAccounts**');

“D. Import and configure the BigDataODBCHiveTemplate data source template in Campaign”

## D. Import and configure the BigDataODBCHiveTemplate data source template in Campaign

This is the fourth step to integrate IBM Campaign with Hive-based Apache Hadoop data sources.

Complete “C. Map existing HBase tables to Hive” on page 31.

To enable Campaign to communicate with your Hive-based Hadoop system, you must do the following actions:

- Import the BigDataODBCHive.xml template into IBM Campaign. You must import the template only once. Importing a template makes it available for creating data sources.
  - Use the template to create and configure a data source for each Hive implementation that communicates with IBM Campaign.
  - For each data source, configure the HiveQueryMode property in the Campaign configuration.
1. Use the configTool utility to import the BigDataODBCHive.xml template into Campaign.
    - BigDataODBCHive.xml is in <Campaign\_Home>/conf.
    - configTool is in <Marketing\_Platform\_Home>/tools/bin. For more information, see the *IBM Marketing Platform Administrator's Guide* in the IBM Knowledge Center.

The following example imports the template into the default Campaign partition, partition1. Replace <Campaign\_Home> with the complete path to the IBM Campaign installation directory.

```
./configTool -i -p "Affinium|Campaign|partitions|partition1|dataSources"
-f <Campaign_Home>/conf/BigDataODBCHive.xml
```

2. Create a data source based on BigDataODBCHiveTemplate. Do this for each Hive implementation that communicates with Campaign. For example, if you have four implementations (MapR, Cloudera, Hortonworks, BigInsights), create four separate data sources, and configure each one.
  - a. In IBM Campaign, choose **Settings > Configuration**
  - b. Go to Campaign|partitions|partition[n]|dataSources.
  - c. Select BigDataODBCHiveTemplate.
  - d. Supply a New category name that identifies the Hive dataSource, for example **Hive\_MapR** or **Hive\_Cloudera** or **Hive\_HortonWorks** or **Hive\_BigInsights**.
  - e. Complete the fields to set the properties for the new data source, then save your changes.

**Important:** Some properties do not have default values, so you must supply them. Pay special attention to the properties described below. This is only a partial list of the properties included in this template. For complete information, see the *IBM Campaign Administrator's Guide*.

| Configuration property          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASMUserForDBCredentials         | No default value defined. Specify the Campaign system user.                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| DSN                             | DSN Name as specified in the odbc.ini file for the Hive-based Hadoop big data instance.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| HiveQueryMode                   | For data sources that use the DataDirect ODBC driver, use Native.<br><br>For data sources that use the Cloudera ODBC driver or Hortonworks Hive ODBC driver, use SQL.                                                                                                                                                                                                                                                                                                                                     |
| JndiName                        | Not needed for user data source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| SystemTableSchema               | No default value defined. Specify the user of the database that you connect to.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| OwnerForTableDisplay            | No default value defined. Specify the user of the database that you connect to.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| LoaderPreLoadDataFileCopyCmd    | SCP is used to copy data from IBM Campaign to a temp folder called /tmp on the Hive-based Hadoop system. The location must be called /tmp and it must be on the Hive server (the file system location, not the HDFS location). This value can either specify the SCP command or call a script that specifies the command.<br><br>For more information and detailed instructions about how to export data from Campaign to a Hive-based Hadoop system, see the <i>IBM Campaign Administrator's Guide</i> . |
| LoaderPostLoadDataFileRemoveCmd | Data files are copied from IBM Campaign to a temp folder on the Hive-based Hadoop system. You must use the SSH "rm" command to remove the temporary data file.<br><br>For more information and detailed instructions about how to export data from Campaign to a Hive-based Hadoop system, see the <i>IBM Campaign Administrator's Guide</i> .                                                                                                                                                            |
| LoaderDelimiter                 | No default value defined. Specify the delimiter such as comma (,) or semi-colon (;) that separates fields in the temporary data files that are loaded into the big data instance. Tab (/t) is not supported.<br><br>The delimiter value must match the ROW format delimiter that was used when the big data database table was created. In this example, a comma is used: ROW FORMAT DELIMITED FIELDS TERMINATED BY ',';                                                                                  |

| Configuration property        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SuffixOnTempTableCreation     | No default value defined. Use the same character as specified for LoaderDelimiter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| SuffixOnSegmentTableCreation  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SuffixOnSnapshotTableCreation |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SuffixOnExtractTableCreation  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SuffixOnUserBaseTableCreation |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SuffixOnUserTableCreation     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| UseExceptForMerge             | Set to FALSE. Hive does not support the EXCEPT clause, so a setting of TRUE can result in process failures.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| DateFormat                    | All Date strings must use the dash "-" character to format dates. Hive does not support any other characters for dates. Example: %Y-%m-%d %H:%M:%S                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| DateTimeFormat                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DateTimeOutputFormatString    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Type                          | BigDataODBC_Hive                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| UseSQLToRetrieveSchema        | Set to FALSE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| DataFileStagingFolder         | <p>Default location value is set to /tmp. You can change the location value. Example: /opt/campaign/<br/> <b>Note:</b> The value for this folder must have a trailing slash.</p> <p>If you have written shell script to copy the Campaign data file to the Hive server, you need to modify it. Example:</p> <pre>#!/bin/sh scp \$1 root@emm52.in.ibm.com:/opt/campaign/ ssh root@emm52.in.ibm.com "chmod 0666 /opt/campaign/`basename \$1`"</pre> <p>If you are using LoaderPreLoadDataFileCopyCmd, then you need to update the file location. Example:</p> <pre>scp &lt;DATAFILE&gt; &lt;USER&gt;@[hostname]:/opt/campaign/</pre> <p>If you are using LoaderPostLoadDataFileRemoveCmd, then you need to update the file location. Example:</p> <pre>ssh &lt;USER&gt;@[hostname] "rm /opt/campaign/&lt;DATAFILE&gt;"</pre> |

“E. Configure SSH on the Campaign listener server”

## E. Configure SSH on the Campaign listener server

This is the fifth step to integrate IBM Campaign with Hive-based Apache Hadoop data sources.

Complete “D. Import and configure the BigDataODBC HiveTemplate data source template in Campaign” on page 32.

To enable data file transfers between the IBM Campaign listener (analytic) server and the Hive-based Hadoop big data instance, you must configure SCP and SSH seamless login. SSH allows secure connection between two computers; most systems use the OpenSSH client.

1. On the machine that is running the IBM Campaign listener, configure SSH for no password prompt for authentication. Log in as the user who is running the IBM Campaign listener and run the following commands, substituting the **username@IP address** of your Hive server (in this example, the MapR machine):

```
>> ssh-keygen -t rsa
>> ssh mapr@192.0.2.0 mkdir -p .ssh
>> cat .ssh/id_rsa.pub | ssh mapr@192.0.2.0 'cat >> .ssh/authorized_keys'
```

```
>> ssh mapr@192.0.2.0 "chmod 700 .ssh; chmod 640 .ssh/authorized_keys"
```

2. Verify password-less authentication using RSA-based authorized keys. Run each command, substituting the actual **username@IP address**, and verify that it works. You need a local file called test1 for this test to work:

```
>> ssh mapr@192.0.2.0
>> scp test1 mapr@192.0.2.0:/tmp
>> ssh mapr@192.0.2.0 "rm /tmp/test1"
```

“F. Map the Hive datasource in Campaign”

## F. Map the Hive datasource in Campaign

This is the sixth and final step to integrate IBM Campaign with Hive-based Apache Hadoop data sources.

- Complete “E. Configure SSH on the Campaign listener server” on page 34
- You must define audience levels in IBM Campaign before you map user tables.

Mapping user tables makes external data sources accessible in Campaign. A typical user table contains information about your company's customers, prospects, or products, for use in marketing campaigns. You must map any data sources that you configured to make that data accessible to processes in flowcharts.

1. Select **Settings > Campaign Settings > Manage Table Mappings**.
2. In the Table Mappings dialog, click **Show User Tables**.
3. Click **New Table**. The New Table Definition dialog opens.
4. Click **Next**.
5. Select **Map to Existing Table in Selected Database**.
6. Select the **BigDataODBCHive** datasource, then click **Next**.
7. Follow the prompts to map the table, as explained in the *IBM Campaign Administrator's Guide*.

You have completed the installation and configuration process. IBM Campaign flowchart designers can now select Hive-based Hadoop data sources in processes. For example, a Select process can query a Hive-based data source to target customer data for a marketing campaign.

---

## Preparing Cloudera Hadoop Impala based user data sources for Campaign

Follow the steps below to enable Impala based Cloudera Hadoop data sources to be used in IBM Campaign

Campaign supports Cloudera Hadoop using Impala for customer tables only, not system tables. For details about supported versions, see the *Recommended Software Environments and Minimum System Requirements*.

| Task                                | Description                                                                                                                                                                      |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Install Cloudera Impala ODBC driver | You can install the Cloudera Impala ODBC Driver from Cloudera .                                                                                                                  |
| Configure Impala ODBC driver        | Configuration includes modifying .ini files and setting path values and environment variables. Be sure to follow the appropriate instructions for the driver that you installed. |

| Task                                                                       | Description                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Import and Configure the BigDataODBCImpala template datasource in Campaign | Use the configTool utility to import the template BigDataODBCImpalaTemplate.xml into Campaign.<br><br>Then go to Campaign partitions partition[n] dataSources and create a datasource based on the imported BigDataODBCImpalaTemplate. |
| Configure SSH on Campaign listener Server                                  | To enable data file transfers between the IBM Campaign listener (analytic) server and the Hive-based Hadoop big data instance, you must configure SCP and SSH seamless login.                                                          |
| Map the Impala Datasource in camapign                                      | Mapping user tables is the process of making external data sources accessible in Campaign.                                                                                                                                             |

You can now use your Impala based Cloudera big data system as a user data source for IBM Campaign. For example, create a marketing campaign flowchart that uses customer account data from your big data instance to target customers with specific account types and balances.

## Requirements and restrictions for Impala-based Hadoop data sources

The following requirements and restrictions apply to using Impala-based Hadoop data sources with IBM Campaign.

- The following drivers is required. The customer is responsible for obtaining the driver.
- Cloudera ODBC Driver for Impala Version 2.5.41.
- Minimum supported Impala version: 2.9.0.
- Impala based Hadoop user data source integration is supported with Cloudera.
- The big data integration is currently supported on Linux RHEL 6.3 or higher.
- Impala-based Hadoop is supported as a user data source only. It is not supported for Campaign system tables.
- The big data integration does not currently support the IBM Campaign Cube, Optimize, or Interact List process boxes or eMessage Landing Pages in an Extract process box.

### A. Install a Cloudera Impala ODBC driver

This is the first step to integrate IBM Campaign with Impala-based Cloudera data sources. You must install the Cloudera Impala ODBC driver from Cloudera, Inc.

Install the Impala ODBC driver on the system where you have installed the Campaign Analytical Server. If the Analytical Server is installed in a cluster, install the Impala ODBC driver on each system.

Installing the Cloudera Impala ODBC driver. This topic explains how to install the Cloudera Impala ODBC driver, which is a fully compliant ODBC driver that supports multiple Hadoop distributions.



## Procedure

1. Obtain the Cloudera Impala ODBC Driver for version 2.5.41 at <https://www.cloudera.com/>
2. Install and configure the Cloudera driver. Installation instructions are downloaded to where you installed the driver.
3. Download and install UnixODBC Manager, unixODBC-2.3.2 on the IBM Campaign listener.
  - a. Copy the unixODBC-2.3.2.tar.gz file to the Campaign listener (analyticsserver).
  - b. `gunzip unixODBC-2.3.2.tar.gz`
  - c. `tar xvf unixODBC-2.3.2.tar`
4. On the Campaign listener, run `rpm -ivh ClouderaImpalaODBC-2.5.41.1029-1.el6.x86_64.rpm`.

## B. Installing the Impala ODBC driver for Windows

Obtain the Cloudera Impala ODBC Driver for version 2.5.42 at <https://www.cloudera.com/> for Windows and install it.

To install the Cloudera ODBC Driver for Impala:

1. Depending on the configuration of your client application, double-click to run `ClouderaImpalaODBC32.msi` or `ClouderaImpalaODBC64.msi`.
2. Click Next.
3. Select the check box to accept the terms of the License Agreement if you agree, and then click Next.
4. To change the installation location, click the Change button, then browse to the desired folder, and then click OK. To accept the installation location, click Next.
5. Click Install.
6. When the installation completes, click Finish.

## C. Configure the Impala ODBC Driver

This is the second step to integrate Campaign with Impala-based Cloudera data sources. You must configure the driver that you installed (Cloudera Impala ODBC driver from Cloudera, Inc.).

### Before you begin

Complete the Cloudera Impala ODBC driver installation that is described in “Installing the Cloudera ODBC driver”

### Procedure

1. Modify the ODBC.ini file to specify the Impala connector information, using the values shown in the following example. Be sure to customize the items in bold to match your own configuration.

Note: Do not enable Native Query mode in the configuration.

[Cloudera Impala]

Description=Cloudera ODBC Driver for Impala (64-bit) DSN

Driver=/opt/cloudera/impalaodbc/lib/64/libclouderaimpalaodbc64.so

HOST=<Impala Host>

PORT=<Impala port - default is 21050>

Database=<Database name>

Schema=<Database schema>

AuthMech=3

UseSASL=0

UID=<impala user id>

PWD=<password>

SSL=0

TSaslTransportBufSize=1000

RowsFetchedPerBlock=10000

SocketTimeout=0

StringColumnLength=32767

UseNativeQuery=1

Note: The UID is the user ID of the Cloudera cluster database.

2. Assuming that your ODBC driver is installed in the following location  
/opt/cloudera/impalaodbc/lib/64/libclouderaimpalaodbc64.so

Make the following change in Campaign/bin/setenv.sh:

Set LD\_LIBRARY\_PATH (LIBPATH for AIX) to include Cloudera LIB and  
UnixODBC LIB as follows.

```
LD_LIBRARY_PATH=/opt/IBM/Campaign/bin:/opt/ibm/db2/V10.1/lib64/::/
usr/lib:/lib:/opt/cloudera/impalaodbc/lib/64:/usr/local/lib/
```

```
export LD_LIBRARY_PATH
```

ODBCINI=/usr/local/etc/odbc.ini # You can use the ODBC INI from Cloudera  
Impala ODBC Driver install location as well

```
export ODBCINI
```

## **C. Import and configure the BigDataODBCImpalaTemplate data source template in Campaign**

This is the third step to integrate IBM Campaign with Impala-based Cloudera Hadoop data sources.

To enable Campaign to communicate with your Impala-based Hadoop system, you must do the following actions:

- Import the BigDataODBCImpalaTemplate.xml template into IBM Campaign. You must import the template only once. Importing a template makes it available for creating data sources.

1. Use the configTool utility to import the BigDataODBCImpalaTemplate.xml template into Campaign.
  - BigDataODBCImpalaTemplate.xml is in <Campaign\_Home>/conf.
  - configTool is in <Marketing\_Platform\_Home>/tools/bin. For more information, see the *IBM Marketing Platform Administrator's Guide* in the IBM Knowledge Center.

The following example imports the template into the default Campaign partition, partition1. Replace <Campaign\_Home> with the complete path to the IBM Campaign installation directory.

```
./configTool -i -p "Affinium|Campaign|partitions|partition1|dataSources"
-f <Campaign_Home>/conf/BigDataODBCImpalaTemplate.xml
```

2. Create a data source based on BigDataODBCImpalaTemplate.
  - a. In IBM Campaign, choose **Settings > Configuration**
  - b. Go to Campaign|partitions|partition[n]|dataSources.
  - c. Select BigDataODBCImpalaTemplate.
  - d. Complete the fields to set the properties for the new data source, then save your changes.

**Important:** Some properties do not have default values, so you must supply them. Pay special attention to the properties described below. This is only a partial list of the properties included in this template. For complete information, see the *IBM Campaign Administrator's Guide*.

| Configuration property          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASMUserForDBCredentials         | No default value defined. Specify the Campaign system user.                                                                                                                                                                                                                                                                                                                                                                                       |
| DSN                             | DSN Name as specified in the odbci.ini file for the Hive-based Hadoop big data instance.                                                                                                                                                                                                                                                                                                                                                          |
| JndiName                        | Not needed for user data source.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| SystemTableSchema               | No default value defined. Specify the user of the database that you connect to.                                                                                                                                                                                                                                                                                                                                                                   |
| OwnerForTableDisplay            | No default value defined. Specify the user of the database that you connect to.                                                                                                                                                                                                                                                                                                                                                                   |
| LoaderPreLoadDataFileCopyCmd    | <p>SCP is used to copy data from IBM Campaign to any location on the Impala-based Hadoop system. The location must be a valid HDFS location on the Impala server. This value can either specify the SCP command or call a script that specifies the command.</p> <p>For more information and detailed instructions about how to export data from Campaign to a Impala-based Hadoop system, see the <i>IBM Campaign Administrator's Guide</i>.</p> |
| LoaderPostLoadDataFileRemoveCmd | <p>Data files are copied from IBM Campaign to a temp folder on the Hive-based Hadoop system. You must use the SSH "rm" command to remove the temporary data file.</p> <p>For more information and detailed instructions about how to export data from Campaign to a Impala-based Hadoop system, see the <i>IBM Campaign Administrator's Guide</i>.</p>                                                                                            |
| LoaderDelimiter                 | <p>No default value defined. Specify the delimiter such as comma (,) or semi-colon (;) that separates fields in the temporary data files that are loaded into the big data instance. Tab (/t) is not supported.</p> <p>The delimiter value must match the ROW format delimiter that was used when the big data database table was created. In this example, a comma is used: ROW FORMAT DELIMITED FIELDS TERMINATED BY ',' ;"</p>                 |

| Configuration property        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SuffixOnTempTableCreation     | No default value defined. Use the same character as specified for LoaderDelimiter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| SuffixOnSegmentTableCreation  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SuffixOnSnapshotTableCreation |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SuffixOnExtractTableCreation  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SuffixOnUserBaseTableCreation |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SuffixOnUserTableCreation     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| UseExceptForMerge             | Set to FALSE. Hive does not support the EXCEPT clause, so a setting of TRUE can result in process failures.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| DateFormat                    | All Date strings must use the dash "-" character to format dates. Hive does not support any other characters for dates. Example: %Y-%m-%d %H:%M:%S                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| DateTimeFormat                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| DateTimeOutputFormatString    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Type                          | BigDataODBC_Impala                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| UseSQLToRetrieveSchema        | Set to FALSE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| DataFileStagingFolder         | <p>Default location value is set to /tmp. You can change the location value. Example: /opt/campaign/<br/> <b>Note:</b> The value for this folder must have a trailing slash.</p> <p>If you have written shell script to copy the Campaign data file to the Hive server, you need to modify it. Example:</p> <pre>#!/bin/sh ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "mkdir -p \$2" #Operating system location scp \$1 &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt;:\$2 # SCP data file to Impala Server filename=\$(basename "\$1") extension="\${filename##*.}" filename="\${filename%.*}" #Moving Campaign data file with name containing special characters is prohibited, so extract ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "mv \$2/~basename \$1` \$2/\$filename" ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "hadoop fs -put -f \$2/\$filename \$2" ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "hadoop fs -mv \$2/\$filename \$2/\$filename.\$extension" ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "hadoop fs -chmod 0777 \$2/\$filename.\$extension"</pre> <p>If you are using LoaderPreLoadDataFileCopyCmd, then you need to update the file location. Example:</p> <pre>scp &lt;DATAFILE&gt; &lt;USER&gt;@[hostname]:/opt/campaign/</pre> <p>If you are using LoaderPostLoadDataFileRemoveCmd, then you need to update the file location. Example:</p> <pre>ssh &lt;USER&gt;@&lt;HADOOP NODE WITH IMPALA&gt; "hadoop fs -rm \$2/~basename \$1`"</pre> |
| DeleteAsRecreate              | Please set this configuration parameter value to FALSE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| DeleteAsTruncate              | Please set this configuration parameter value to TRUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## D. Configure SSH on the Campaign listener server

This is the fourth step to integrate IBM Campaign with Impala-based Cloudera Hadoop data sources.

To enable data file transfers between the IBM Campaign listener (analytic) server and the Impala-based Hadoop big data instance, you must configure SCP and SSH seamless login. SSH allows secure connection between two computers; most systems use the OpenSSH client.

1. On the machine that is running the IBM Campaign listener, configure SSH for no password prompt for authentication. Log in as the user who is running the IBM Campaign listener and run the following commands, substituting the **username@IP address** of your Impala server (in this example, the MapR machine):

```
>> ssh-keygen -t rsa
>> ssh cloudera@192.0.2.0 mkdir -p .ssh
```

- ```
>> cat .ssh/id_rsa.pub | ssh cloudera@192.0.2.0 'cat >> .ssh/authorized_keys'
>> ssh cloudera@192.0.2.0 "chmod 700 .ssh; chmod 640 .ssh/authorized_keys"
```
2. Verify password-less authentication using RSA-based authorized keys. Run each command, substituting the actual **username@IP address**, and verify that it works. You need a local file called test1 for this test to work:

```
>> ssh cloudera@192.0.2.0
>> scp test1 cloudera@192.0.2.0:/tmp
>> ssh cloudera@192.0.2.0 "rm /tmp/test1"
```

F. Configuring connectivity between Campaign listener server and Hadoop cluster

The loader Script is needed to copy temporary data file at <Campaign_Home>/partitions/partition[n]/tmp to Impala Server which is created during flowchart execution.

1. We are making use of pscp (Putty utility) to connect to Cloudera Hadoop Cluster from IBM Campaign server and upload the data files on Hadoop. Note: Here the PSCP utility needs password in argument in plain text.
2. Configuration for
LoaderPreLoadDataFileCopyCmd:LoaderPreLoadDataFileCopyCmd =
C:\MarketingSoftware11\Campaign\bin\CopytoHadoop.bat <DATAFILE> /tmp/

This is a sample script . You can write your own batch script:

Sample Script of copytoHadoop.bat:

```
=====
REM This script is used to copy a file from IBM Campaign to any location on the
Impala-based Hadoop system. The location must a valid HDFS location
@echo off
REM Operating System location
pscp -pw <password> "%1" <username>@<Hadoop Node with Impala>:"%2"
set filepath="%1"
for /F "delims=" %i in (%filepath%) do set basename="%~nxi"
for /F "delims=" %i in (%filepath%) do set filename="%~ni"
for /F "delims=" %i in (%filepath%) do set fileExtension="%~xi"
REM Moving Campaign data with name containing special character is prohibited,
so extracting
plink -pw <password> -ssh <username>@<Hadoop Node with Impala> "mv %2/%baseName% %2/%fileName%"
plink -pw <password> -ssh <username>@<Hadoop Node with Impala> "hadoop fs -put -f %2/%fileName% %2/%baseName%"
plink -pw <password> -ssh <username>@<Hadoop Node with Impala> "hadoop fs -mv
%2/%fileName% %2/%baseName%"
plink -pw <password> -ssh <username>@<Hadoop Node with Impala> "hadoop fs -chmod 0777 %2/%baseName%"
=====
LoaderPostLoadDataFileRemoveCmd
=
C:\MarketingSoftware10\Campaign\bin\
removeFromHadoop.bat <DATAFILE> /tmp/
```

Data files are copied from IBM Campaign to a temp folder on the Impala-based Hadoop system. You must use the SSH "rm" command to remove the temporary data file.

1. Remove temp data from Hadoop file System.
2. Remove temp data from OS file System.

Note:

1. For LoaderPostLoadDataFileCopyCmd and LoaderPostLoadDataFileRemoveCmd to work, the Hadoop System should have TGT and have access to write, modify and remove files from Hadoop File System.

2. The above configuration is Sample configuration, you can refer *IBM Campaign Installation Guide* for details.

G. Map the Impala datasource in Campaign

This is the Fifth and final step to integrate IBM Campaign with Impala-based Hadoop data sources.

Mapping user tables makes external data sources accessible in Campaign. A typical user table contains information about your company's customers, prospects, or products, for use in marketing campaigns. You must map any data sources that you configured to make that data accessible to processes in flowcharts.

1. Select **Settings > Campaign Settings > Manage Table Mappings**.
2. In the Table Mappings dialog, click **Show User Tables**.
3. Click **New Table**. The New Table Definition dialog opens.
4. Click **Next**.
5. Select **Map to Existing Table in Selected Database**.
6. Select the **BigDataODBCImpala** datasource, then click **Next**.
7. Follow the prompts to map the table, as explained in the *IBM Campaign Administrator's Guide*.

You have completed the installation and configuration process. IBM Campaign flowchart designers can now select Impala-based Hadoop data sources in processes. For example, a Select process can query a Impala-based data source to target customer data for a marketing campaign.

Preparing Cloudera BigData instance using Kerberos authentication

With the v11.1 release, IBM Campaign supports Cloudera based BigData instance using Kerberos authentication. The Campaign Analytical Server connects with the Kerberos enabled Cloudera BigData instance with the valid Kerberos enabled user login or with user impersonation.

Note: In case the kerberos authentication token expires, user will not able to connect to the Bigdata instance. Also, the flowcharts execution will fail due to the authentication issue. User must make sure that the authentication using kerberos is seamless using TGT while Campaign is running.

The Campaign analytical should always be connected with the BigData instance - if the Operating System user is not Kerberos enabled, there is a need to authenticate this user using Kerberos keytab.

Note: In case of kerberos authentication token expired, user will not able to connect to Bigdata instance. Also flowcharts execution will get failed due to authentication issue. User needs to make sure authentication using kerberos is seamless using TGT while Campaign is running.

Note: Kerberos authentication is currently only certified for RHEL Linux and Windows operating system.

Table 5.

Task	Description
Download, Install and Configure Kerberos packages on Campaign listener machine	You can install and configure the required Kerberos packages on Windows and Linux operating systems as mentioned in the respective section. On Windows MIT Kerberos packages should be installed.
Establish seamless connection with Cloudera Hadoop instance using Kerberos on Campaign listener machine	On Windows a script is provided to have a seamless connection with Cloudera Hadoop instance. This Script obtains the Ticket for the Kerberos principal , renews and regenerates the ticket before expiry. On Linux environment user is recommended to write cron jobs to achieve using kinit and k5start utility.
Install Cloudera Impala ODBC driver.	Please follow the steps to install and configure ODBC driver as mentioned in section – (“Preparing Cloudera Hadoop Impala based user data sources for Campaign” on page 35).
Configure Cloudera Impala ODBC Driver using Kerberos.	Configuration include setting hostname, port number, schema, authentication mechanism as Kerberos.
Import and Configure the BigDataODBCImpala template Data Source in Campaign	Please follow the steps to configure BigDataODBCImpala template as mentioned in section (“Preparing Cloudera Hadoop Impala based user data sources for Campaign” on page 35).
Map the Impala Data Source in Campaign	Please follow the steps to map Impala datasource in Campaign as mentioned in section (“Preparing Cloudera Hadoop Impala based user data sources for Campaign” on page 35).

Installing and configuring MIT Kerberos for Windows

Downloading and installing MIT Kerberos for Windows 4.0.1

For information on Kerberos and download links for the installer, see the MIT Kerberos website: <http://web.mit.edu/kerberos/>. To download and install MIT Kerberos for Windows 4.0.1:

1. Download the appropriate Kerberos installer:
 - For a 64-bit machine, use the following download link from the MIT Kerberos website:<http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1-amd64.msi>.
 - For a 32-bit machine, use the following download link from the MIT Kerberos website:<http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1-i386.msi>.
2. To run the installer, double-click the .msi file that you downloaded above.
3. Follow the instructions in the installer to complete the installation process.
4. When the installation completes, click **Finish**.

Setting up the Kerberos configuration file

Settings for Kerberos are specified through a configuration file. You can set up the configuration file as an .ini file in the default location, which is the C:\ProgramData\MIT\Kerberos5 directory. Normally, the C:\ProgramData\MIT\Kerberos5 directory is hidden. For information about viewing and using this hidden directory, refer to *Microsoft Windows documentation*. **Note:** For more information on configuring Kerberos, refer to the *MIT Kerberos documentation*.

To set up the Kerberos configuration file:

1. Obtain a krb5.conf configuration file. You can obtain this file from your Kerberos administrator, or from the /etc/krb5.conf folder on the machine that is hosting the Impala server.
2. Rename the configuration file from krb5.conf to krb5.ini.
3. Copy the krb5.ini file to the C:\ProgramData\MIT\Kerberos5 directory and overwrite the empty sample file.

Setting up the Kerberos credential cache file

Kerberos uses a credential cache to store and manage credentials.

To set up the Kerberos credential cache file:

1. Create a directory where you want to save the Kerberos credential cache file. For example, create a directory named C:\temp.
2. Open the System window:
 - If you are using Windows 7 or earlier, click Start (Windows icon), then right-click Computer, and then click Properties.
 - Or, if you are using Windows 8 or later, right-click This PC on the Start screen, and then click Properties.
3. Click Advanced System Settings.
4. In the System Properties dialog box, click the Advanced tab and then click Environment Variables.
5. In the Environment Variables dialog box, under the System Variables list, click New.
6. In the New System Variable dialog box, in the Variable Name field, type KRB5CCNAME.
7. In the Variable Value field, type the path to the folder you created above, and then append the file name krb5cache. For example, if you created the folder C:\temp, then type C:\temp\krb5cache.
Note: krb5cache is a file (not a directory) that is managed by the Kerberos software, and it should not be created by the user. If you receive a permission error when you first use Kerberos, make sure that the krb5cache file does not already exist as a file or a directory.
8. Click OK to save the new variable.
9. Make sure that the variable appears in the System Variables list.
10. Click OK to close the Environment Variables dialog box, and then click OK to close the System Properties dialog box.
11. To make sure that Kerberos uses the new settings, restart your machine.

Downloading and installing Kerberos Packages for Linux

Enable Kerberos authentication on Cloudera Bigdata for Linux Server.

1. Install Kerberos Packages on Linux, if they don't already exist. Below packages should be installed on IBM Campaign listener machine
 - a. krb5-lib

- b. libkadm
 - c. krb5-devel
 - d. krb5-workstation
2. Contact your Kerberos administrator and configure /etc/krb5.conf on machine where Campaign listener is running.

Seamless Kerberos connection on Windows

You need to run unica_kerbKeyGenRenew.bat script to obtain the Ticket for the Kerberos principal. The purpose of this script is to run continuously and :

1. Renew the TGT before its renewable time.
2. Regenerate the token before expiry.

By default, it renews the ticket every 24 hours and regenerates a new ticket after 7 days. You can also override the default setting for regenerating and renewing the tickets by setting environment variables.

Prerequisite for running the Script

1. Requires MIT Kerberos to be installed and configured.
2. Keytab file is already present for the said principal.
3. The following environment variables are set:
 - a. TGT_EXPIRY_DAYS: number of days after which this TGT will expire, and can no longer be renewed (specify value in days)
 - b. TGT_RENEWAL_HOURS: number of hours within which this TGT needs to be renewed (specify value in hours)
 - c. KRB5CCNAME: path to the cache file
 - d. KEYTAB_FILE_PATH: path to keytab file for this principal
4. From command line run "Where kinit" command should point MIT Kerberos
For example: C:\Kerbores\Campaign\bin>where kinit
C:\Program Files\MIT\Kerberos\bin\kinit.exe
5. Copy the Keytab file to the Campaign System for the principal you want to authenticate.

Procedure to run the Script

1. Make sure the prerequisites are satisfied.
2. Then run the following:
`<campaign_home>/bin/unica_kerbKeyGenRenew.bat <principal>`
 For example: unica_kerbKeyGenRenew.bat impala/
 quickstart.cloudera@CLOUDERA
3. The script will start running, and do the following:
 - a. Generate the TGT.
 - b. Check every 1 minute for renewal and regeneration of the TGT.
 - c. Using the values in TGT_EXPIRY_DAYS and TGT_RENEWAL_HOURS, keep checking for renewal or expiry.
 - d. Before renewal time - renew the TGT.
And before expiry time - regenerate the TGT.

Note:

- a. TGT_RENEWAL HOURS AND TGT_EXPIRY_DAYS values needs to be same as Kerberos Server configuration. Please contact the Kerberos administrator to get the values.

- b. By default Script sets 24 hours for TGT renewal and 7 days for TGT_EXPIRY.

Possible errors:

1. Usage
2. Principal is incorrect
3. Keytab file is not present for the Principal
4. One or more of the prerequisites is not set
5. Keytab file is not valid

Seamless Kerberos connection on Linux

You need to create 2 cron jobs using kinit and k5start utility. These utility perform following operations

1. Renew the TGT before its renewable time. (k5start)
2. Regenerate the token before expiry.(kinit)
1. Renew TGT using the k5start utility. Add this to cron. . Please contact the Kerberos administrator to get the values.

- a. Establish the Cloudera BigData instance authentication using the Kerberos keytab file:

```
k5start -f <KEYTABFILE> -K <minutes> -l <time> [-u <PRINCIPAL>] [-v]
```

You can add this command in the corn job and make sure that the kerberos token is always alive.

- b. See <https://linux.die.net/man/1/k5start>

```
c. k5start -f <KEYTABFILE> -K <minutes> -l <time> [-u <PRINCIPAL>] [-v]
-f keytab
```

Authenticate using the keytab rather than asking for a password. A key for the client principal m

```
-K minutes
```

Run in daemon mode to keep a ticket alive indefinitely. The program reawakens after minutes minut

```
-l time string
```

Set the ticket lifetime. time string should be in a format recognized by the Kerberos libraries f

```
-v
```

Be verbose. This will print out a bit of additional information about what is being attempted and

For example:

```
k5start -f <FilePath>/impala.keytab -K 1 -l 3m impala/quickstart.cloudera@CLOUDERA
```

```
k5start -f <FilePath>/impala.keytab -K 1 -l 3m impala/quickstart.cloudera@CLOUDERA -v
```

2. Regenerate the Token using kinit utility. Add to cron so that ticket gets regenerated before it expires.

```
kinit -k -t <keytab_filename> <principal>
```

Refer <https://linux.die.net/man/1/kinit> .

```
-k [-t keytab_file]
```

requests a ticket, obtained from a key in the local host's keytab file. The name and location of the keytab file may be specified with the -t keytab file option; otherwise the default name and location will be used. By default a host ticket is requested but any principal may be specified. On a KDC, the special keytab location KDB: can be used to indicate that kinit should open the KDC database and look up the key directly. This permits an administrator to obtain tickets as any principal that supports password-based authentication.

e.g kinit -k -t /root/impala.keytab impala/quickstart.cloudera@CLOUDERA

Configuring the Impala-based Cloudera data sources using Kerberos

1. In the ODBC Data Source Administrator, click the Drivers tab and verify that the Cloudera ODBC Driver for Impala appears in the list of ODBC drivers that are installed on your system.
2. In the Create New Data Source dialog box, select Cloudera ODBC Driver for Impala and then click Finish.
3. Use the options in the Cloudera ODBC Driver for Impala DSN Setup dialog box to configure your DSN:
 - a. In the Data Source Name field, type a name for your DSN.
 - b. Optionally, in the Description field, type relevant details related to the DSN.
 - c. In the Host field, type the IP address or hostname.
 - d. In the Port field, type the listening port for the Impala service.
Note: The default port number for the Impala service is 21050.
 - e. In the Database field, type the name of the database schema to use.

In the Authentication area, configure authentication as Kerberos:

1. Realm: Type the Kerberos realm of the Impala server host in the Realm field.
2. In the Host FQDN field, type the fully qualified domain name of the Impala host.
3. In the Service Name field, type the service name of the Impala server. For example, if the principle for the Impala server is `impala/fully.qualified.domain.name@yourrealm.com`, then the value in the service name field is `impala`. If you are unsure of the correct service name to use for your particular Hadoop deployment, contact your Hadoop administrator

Note: Test the connection to make sure Connection is working fine with Kerberos enabled.

Configuring the Impala-based Cloudera data sources using Kerberos on Linux

Configuring ODBC datasource is same as of mentioned under section ("Preparing Cloudera Hadoop Impala based user data sources for Campaign" on page 35). User will need to configure Cloudera Impala ODBC.ini file to authenticate using Kerberos authentication. User will need to add the Kerberos authentication/connection information for the below mentioned parameters in ODBC.ini file.

```
AuthMech=1 #which mention to connect using Kerberos authentication
# Kerberos related settings with sample values.
KrbRealm=CLLOUDERA
KrbFQDN=quickstart.cloudera
KrbServiceName=impala
```

Import and Configure the BigDataODBCImpala template Data Source in Campaign

Please follow the steps to configure BigDataODBCImpala template as mentioned in section ("Preparing Cloudera Hadoop Impala based user data sources for Campaign" on page 35)

Note: 1. All the Campaign users should use the same Kerberos principal for authentication.

2. Please configure dataSource name with a dummy user name and password because Kerberos uses TGT and does not require username and password. Restart the Campaign Web application and Campaign listener.

Map the Impala Data Source in Campaign

Mapping user tables is the process of making external data sources accessible in Campaign

Troubleshooting

If the loaders fail while executing the flowchart:

1. Verify if you have access and permission to the Hadoop file System on which Kerberos is enabled. You can use Hadoop commands to verify the same.
2. For Windows, verify if unica_kerbKeyGenRenew.bat is running on the Campaign Server.
3. Verify odbctest.exe to see if the connection works. If the odbctest.exe fails, implies that there is some issue with the environment / configuration.

Chapter 4. Installing Campaign

You must run the IBM Marketing Software installer to start the installation of Campaign. The IBM Marketing Software installer starts the Campaign installer during the installation process. Make sure that the IBM Marketing Software installer and the product installer are saved at the same location.

Each time that you run the IBM Marketing Software suite installer, you must first enter database connection information for the Marketing Platform system tables. When the Campaign installer starts, you must enter the required information for Campaign.

After installing Campaign, you can create an EAR file for your product, and you can install the reports package for your product. Creating the EAR file and installing the reports package are not mandatory actions.

Important: Before you install Campaign, make sure that the available temporary space on the computer where you install Campaign is more than three times the size of the Campaign installer.

Note: Information about using web application clustering in IBM Campaign is provided in Chapter 12, “IBM Campaign web application clustering,” on page 115.

Installation files

The installation files are named according to the version of the product and the operating system on which they must be installed, except UNIX. For UNIX, different installation files exist for the X Window System mode and the console mode.

The following table displays examples of the installation files that are named according to the product version and the operating system:

Note: For Suse Linux operating system use the corresponding product Linux build except for Campaign and Contact Optimization.

Note: : EAR deployments are not supported on Tomcat application server

Table 6. Installation files

Operating system	Installation file
Windows: GUI and console mode	<i>Product_N.N.N.N_win64.exe</i> , where <i>Product</i> is the name of your product, <i>N.N.N.N</i> is the version number of your product, and Windows 64-bit is the operating system on which the file must be installed.
UNIX: X Window System mode	<i>Product_N.N.N.N_linux.bin</i> , where <i>Product</i> is the name of your product, and <i>N.N.N.N</i> is the version number of your product.
UNIX: Console mode	<i>Product_N.N.N.N.bin</i> , where <i>Product</i> is the name of your product, and <i>N.N.N.N</i> is the version number of your product. This file can be used for installation on all UNIX operating systems.

Table 6. Installation files (continued)

Operating system	Installation file
SUSE Linux	<i>Product_N.N.N.N_linuxsuse64.bin</i> , where <i>Product</i> is the name of your product, and <i>N.N.N.N</i> is the version number of your product. This file can be used for installation on all SUSE Linux operating systems.

Installing Campaign by using the GUI mode

For Windows, use the GUI mode to install Campaign. For UNIX, use the X Window System mode to install Campaign.

Important: Before you use the GUI mode to install Campaign, make sure that the available temporary space on the computer where you install Campaign is more than three times the size of the Campaign installer.

Make sure that the IBM Marketing Software installer and the Campaign installers are in the same directory on the computer where you want to install Campaign.

Complete the following actions to install Campaign by using the GUI mode:

1. Go to the folder where you have saved the IBM Marketing Software installer and double-click the installer to start it.
2. Click **OK** on the first screen to see the Introduction window.
3. Follow the instructions on the installer, and click **Next**. Use the information in the following table to take the appropriate actions on each window in the IBM Marketing Software installer.

Table 7. IBM Marketing Software installer GUI

Window	Description
Introduction	This is the first window of the IBM Marketing Software suite installer. You can open the Campaign installation and upgrade guides from this window. You can also see a link for the installation and upgrade guides for the products whose installers are saved in the installation directory. Click Next .
Response Files Destination	Click the Generate Response File check box if you want to generate response files for your product. Response files store the information that is necessary for the installation of your product. You can use response files for an unattended installation of your product, or to get pre-filled answers if you rerun the installer in GUI mode Click Choose to browse to a location where you want to store the response files. Click Next .
IBM Marketing Software Products	In the Install Set list, select Custom to select the products that you want to install. In the Install Set area, you see all of the products whose installers are in the same directory on your computer. The Description field shows the description of the product that you select in the Install Set area. Click Next .

Table 7. IBM Marketing Software installer GUI (continued)

Window	Description
Installation Directory	<p>In the Specify the installation directory field, click Choose to browse to the directory where you want to install your product.</p> <p>If you want to install the product in the folder where the installers are stored, click Restore Default Folder.</p> <p>Click Next.</p>
Select Application Server	<p>Select one of the following application servers that you configured for Campaign installation:</p> <ul style="list-style-type: none"> • IBM WebSphere • Oracle WebLogic • Jboss • Tomcat <p>Click Next.</p>
Platform Database Type	<p>Select the appropriate Marketing Platform database type.</p> <p>Click Next.</p>
Platform Database Connection	<p>Enter the following information about your database:</p> <ul style="list-style-type: none"> • Database host name • Database port • Database name or System ID (SID) • Database user name • Database password <p>Click Next.</p>
Platform Database Connection (continued)	<p>Review and confirm the JDBC connection.</p> <p>Click Next. The URL can be customized with additional parameters if needed.</p>
Preinstallation Summary	<p>Review and confirm the values that you added during the installation process.</p> <p>Click Install to start the installation process.</p> <p>The Campaign installer opens.</p>

- Follow the instructions on the Marketing Platform installer to install or upgrade Marketing Platform. See *IBM Marketing Platform Installation Guide* for more information.
- In the Installation Complete window, click **Done**. The Marketing Platform installation is complete, and the Campaign installer opens.
- Use the information in the following table to navigate the Campaign installer. In the Platform Database Connection window, enter all the required information and click **Next** to start the Campaign installer.

Table 8. IBM Campaign installer GUI

Window	Description
Introduction	<p>This is the first window of the Campaign installer. You can open the Campaign installation and upgrade guides from this window.</p> <p>Click Next.</p>
Software License Agreement	<p>Carefully read the agreement. Use Print to print the agreement. Click Next after you accept the agreement.</p>

Table 8. IBM Campaign installer GUI (continued)

Window	Description
Installation Directory	Click Choose to browse to the directory where you want to install Campaign. Click Next .
Campaign Components	Select the components that you want to install. When you select the components, the installer displays information about the components. Click Next .
Single or Multiple Listeners	Select Non-clustered listener configuration (single node) Note: If you are doing a clustered listener installation, see Appendix B for installation instructions. Click Next .
Default locale	Select a default locale for your installation. English is selected by default. Click Next .
Campaign Database Setup	Select one of the following options for setting up the Campaign database: <ul style="list-style-type: none"> • Automatic database setup • Manual database setup If you select Automatic database setup , select Run Unicode SQL Script if your system tables are configured for Unicode. Click Next .
Campaign Database Type	Select the appropriate database type. Click Next .
Campaign Database Connection	Enter the following details for the Campaign database: <ul style="list-style-type: none"> • Database host name • Database port • Database system ID (SID) • Database user name • Password Important: If IBM Marketing Software products are installed in a distributed environment, you must use the machine name rather than an IP address in the navigation URL for all of the applications in the suite. Also, if you are on a clustered environment and choose to use ports that are different from the default ports 80 or 443 for your deployment, do not use a port number in the value of this property. Click Next .
JDBC Connection	Review and confirm the JDBC connection. Click Next .

Table 8. IBM Campaign installer GUI (continued)

Window	Description
Campaign Connection Settings	<p>Enter the following connection settings:</p> <ul style="list-style-type: none"> • Network domain name <p>Note:</p> <p>When you add the network domain name, you might see the following message:</p> <p>Warning-Server name includes domain name, final URL contains several occurrences of domain name</p> <p>Select Modify to change the domain name, or click Cancel to cancel the message.</p> <ul style="list-style-type: none"> • Host name • Port number <p>Select the Use secure connection check box if necessary.</p> <p>Click Next.</p>
Platform Utilities Settings	<p>Review and confirm or change the following information for connecting to the Marketing Platform database:</p> <ul style="list-style-type: none"> • JDBC driver class • JDBC connection URL • JDBC driver classpath <p>Click Next to validate your entries and to go to the next window.</p>
Datasource creation	<p>The installer can create the JDBC data source for one partition in your web application server (WebSphere, WebLogic, Jboss or Tomcat). This data source is required to enable the Campaign web application to connect to the Campaign system tables. You have the option to skip this step in the installer and create the data source in the administration console of your web application server after installation is complete.</p> <p>If you select the Create Campaign Datasource check box, the installer creates the data source using information you provide..</p> <ul style="list-style-type: none"> • Your application server must be running. • Enter any name for the datasource. <p>The installer creates the JNDI name automatically (campaignPartition1DS).</p> <ul style="list-style-type: none"> • Enter the directory in which the application server is installed. <p>For WebSphere, this is the directory that contains the profiles directory. You must restart WebSphere after the installer finishes and before you use this data source.</p> <p>For WebLogic, this is the directory that contains the common directory.</p> <p>Click Next to go to the next window.</p>
Application server information (WebSphere)	<p>Applies only if your application server is WebSphere.</p> <ul style="list-style-type: none"> • Enter the application server profile where you plan to deploy Campaign. • Enter the name of the server in the profile. • If security is enabled in WebSphere, enter the administrator user ID and password. <p>Click Next to go to the next window.</p>

Table 8. IBM Campaign installer GUI (continued)

Window	Description
Application server information (WebLogic)	<p>Applies only if your application server is WebLogic.</p> <p>Enter the domain server name and the administrator user ID and password. If SSL is enabled, enter the HTTP port for the domain.</p> <p>Note that you must add the database driver classpath to WebLogic in the WebLogic administration console.</p> <p>Click Next to go to the next window.</p>
Application server information (Tomcat)	<p>Applies only if your application server is Tomcat.</p> <p>Enter the application Tomcat Instance directory where you plan to deploy Campaign.</p> <p>Click Next to go to the next window.</p>
Preinstallation Summary	<p>Review and confirm the values that you added during the installation process.</p> <p>Click Install to start the installation process.</p> <p>The Campaign installer opens.</p>
Installation Complete	<p>Click Done to close the Marketing Platform installer and go back to the IBM Marketing Software installer.</p>

7. In the Installation Complete window, click **Done** to exit the Campaign installer and go back to the IBM Marketing Software installer.
8. Follow the instructions on the IBM Marketing Software installer to finish installing Campaign. Use the information in the following table to take the appropriate actions on each window in the IBM Marketing Software installer.

Table 9. IBM Marketing Software installer GUI

Window	Description
Deployment EAR file	<p>Specify whether you want to create an enterprise archive (EAR) file to deploy your IBM Marketing Software products.</p> <p>Click Next.</p>
Package EAR file	<p>You see this window if you selected Create an EAR file for deployment in the Deployment EAR file window.</p> <p>Select the applications that you want to package in the EAR file.</p>
EAR file details	<p>Enter the following information for your EAR file:</p> <ul style="list-style-type: none"> • Enterprise application identifier • Display name • Description • EAR file path
EAR file details (continued)	<p>Select Yes or No to create an additional EAR file. If you select Yes, you must enter the details for the new EAR file.</p> <p>Click Next to complete the installation of your product.</p>
Deployment EAR file	<p>Specify whether you want to create another EAR file to deploy your IBM Marketing Software products.</p> <p>Click Next.</p>

Table 9. IBM Marketing Software installer GUI (continued)

Window	Description
Installation Complete	<p>This window provides the locations of the log files that are created during installation.</p> <p>Click Previous if you want to change any installation details.</p> <p>Click Done to close the IBM Marketing Software installer.</p>

Creating an EAR file after running the installer

You can create an EAR file after you install IBM Marketing Software products. You might do this to create an EAR file with a desired combination of products.

Note: Run the installer in console mode from the command line.

Use the following procedure if you want to create an EAR file after you install IBM Marketing Software products:

1. If this is the first time you are running the installer in console mode, make a backup copy of the installer's .properties file for each of your installed products.
 - Each IBM product installer creates one or more response files with a .properties extension. These files are in the same directory where you placed the installers. Be sure to back up all files with the .properties extension, including the installer_<product initials><product version number>.properties files and the file for the IBM installer itself, which is named installer.properties.

For example, the properties files for Marketing Platform installer_ump11.1.0.0.properties and the properties file for Optimize installer_uo11.1.0.0.properties.

 - If you plan to run the installer in unattended mode, you must back up the original .properties files, because when the installer runs in unattended mode, it clears these files. To create an EAR file, you need the information that the installer writes in the .properties files during the initial installation.
2. Open a command window and change directories to the directory that contains the installer.
 3. Run the installer executable with this option:
 -DUNICA_GOTO_CREATEEARFILE=TRUE
 On UNIX type systems, run the .bin file rather than the .sh file.
 The installer wizard runs.
 4. Follow the instructions in the wizard.
 5. Before you create more EAR files, overwrite the .properties file or files with the backups you created before you ran in console mode for the first time.
Please note EAR is not supported for Tomcat application server.

Installing Campaign using the console mode

The console mode allows you to install Campaign using the command-line window. You can select various options in the command-line window to complete tasks such as selecting the products to install, or selecting the home directory for the installation.

Before you install Campaign, ensure that you have configured the following.

- An application server profile
- A database schema

To display the installer screens correctly in console mode, configure your terminal software to support UTF-8 character encoding. Other character encodings, such as ANSI, do not render the text correctly, and some information is not readable with these encodings.

1. Open a command-line prompt window and navigate to the directory where you have saved the IBM Marketing Software installer and the Campaign installer.
2. Complete one of the following actions to run the IBM Marketing Software installer.
 - For Windows, enter the following command:
`ibm_ims_installer_full_name -i console`
 For example, `IBM_Marketing_Software_Installer_11.1.0.0_win.exe -i console`
 - For UNIX, invoke the `ibm_ims_installer_full_name.sh` file.
 For example: `IBM_Marketing_Software_Installer_11.1.0.0.sh`
3. Follow the directions that are displayed in the command-line prompt. Use the following guidelines when you have to select an option in the command-line prompt:
 - The default options are defined by the symbol [X].
 - To select or clear an option, type the number that is defined for the option, and then press Enter.

For example, suppose that the following list displays the components that you can install:

- 1 [X] Marketing Platform
- 2 [X] Campaign
- 3 [] Contact Optimization
- 4 [] Distributed Marketing

If you want to install Distributed Marketing, and do not want to install Campaign, enter the following command: **2,4**

The following list would then display the options that you selected:

- 1 [X] Marketing Platform
- 2 [] Campaign
- 3 [] Contact Optimization
- 4 [X] Distributed Marketing

Note: Do not clear the option for Marketing Platform unless you have already installed it.

4. The IBM Marketing Software installer launches the Campaign installer during the installation process. Follow the instructions in the command-line prompt window of the Campaign installer.
5. After you enter quit in the Campaign installer command-line prompt window, the window shuts down. Follow the instructions in the command-line prompt window of the IBM Marketing Software installer to complete the installation of Campaign.

Note: If any error occurs during the installation, a log file is generated. You must exit the installer to view the log file.

Installing Campaign silently

Use the unattended or silent mode to install Campaign multiple times.

Before you install Campaign, make sure that you have configured the following elements:

- An application server profile
- A database schema

When you install Campaign by using the silent mode, response files are used to obtain the information that is required during installation. You must create response files for a silent installation of your product. You can create response files by using one of the following methods:

- Using the sample response files as a template for creating your response files. The sample response files are included with your product installers in the ResponseFiles compressed archive. For more information about sample response files, see “Sample response files.”
- Running the product installers in the GUI (Windows) or X Window System (UNIX) mode or the console mode before you install the product in the silent mode. One response file is created for the IBM Marketing Software suite installer, and one or more response files are created for your product installer. The files are created in the directory that you specify.

Important: For security reasons, the installer does not save database passwords in the response files. When you create response files, you must edit each response file to enter database passwords. Open each response file and search for PASSWORD to find where you must edit the response file.

When the installer runs in the silent mode, it looks for the response files in the following directories sequentially:

- In the directory where the IBM Marketing Software installer is saved
- In the home directory of the user who installs the product

Make sure that all response files are in the same directory. You can change the path where response files are read by adding arguments to the command line. For example: **-DUNICA_REPLAY_READ_DIR="myDirPath" -f myDirPath/installer.properties**

Use the following command for Windows:

- ***ibm_ims_installer_full_name -i silent***

For example:

IBM_Marketing_Software_Installer_11.1.0.0_win.exe -i silent

Use the following command for Linux:

- ***ibm_ims_installer_full_name_operating_system .bin -i silent***

For example:

IBM_Marketing_Software_Installer_11.1.0.0_linux.bin -i silent

Sample response files

You must create response files to set up a silent installation of Campaign. You can use sample response files to create your response files. The sample response files are included with the installers in the ResponseFiles compressed archive.

The following table provides information about sample response files:

Table 10. Description of sample response files

Sample response file	Description
installer.properties	The sample response file for the IBM Marketing Software master installer.
installer_ <i>product initials and product version number</i> .properties	<p>The sample response file for the Campaign installer.</p> <p>For example, installer_uch.n.n.n.properties is the response file of the Campaign installer, where <i>n.n.n.n</i> is the version number.</p>
installer_ <i>report pack initials, product initials, and version number</i> .properties	<p>The sample response file for the reports pack installer.</p> <p>For example, installer_urpcn.n.n.n.properties is the response file of the Campaign reports pack installer, where <i>n.n.n.n</i> is the version number</p>

Chapter 5. Configuring Campaign before deployment

Before you deploy Campaign, create and populate the Campaign and eMessage system tables and register Campaign and eMessage manually.

Note: Information about using web application clustering in IBM Campaign is provided in Chapter 12, “IBM Campaign web application clustering,” on page 115.

Creating and populating Campaign system tables manually

By default, the Campaign installer automatically creates and populates system tables. However, if the system tables are not created and populated automatically during the installation, you must populate the system tables manually. Use your database client to run the Campaign SQL scripts against the appropriate database to create and populate the Campaign system tables.

Note: If you plan to enable eMessage, you must also manually create and populate the eMessage system tables if the installer did not do so automatically. For details, see “Creating and populating the eMessage system tables manually” on page 60.

If you selected the **Campaign System Table DDL Files** option on the Campaign Components page during installation, the IBM installer installs a collection of SQL scripts that you can use to create and populate the Campaign system tables. The SQL scripts are installed in the `ddl` directory under your installation of the Campaign server. If your system tables are configured for Unicode, the appropriate scripts are in the `ddl/unicode` directory under your Campaign installation.

To use the SQL scripts, run your database client to apply the scripts to the database or schema that holds the Campaign system tables. See the documentation for your database software for instructions on running SQL scripts.

The following tables list the SQL scripts that are provided for you to use to create and populate the Campaign system tables manually:

Table 11. Scripts for creating Campaign system tables

Data source type	Script name
IBM DB2	<code>ac_systab_db2.sql</code>
Microsoft SQL Server	<code>ac_systab_sqlsvr.sql</code>
Oracle	<code>ac_systab_ora.sql</code>

Table 12. Scripts for populating Campaign system tables

Data source type	Script name
IBM DB2	<code>ac_populate_tables_db2.sql</code>
Microsoft SQL Server	<code>ac_populate_tables_sqlsvr.sql</code>
Oracle	<code>ac_populate_tables_ora.sql</code>

You can create tables in the Campaign schema if you want to integrate Distributed Marketing with Campaign.

The following tables list the SQL scripts that are provided for you to use to create the Campaign system tables manually:

Table 13. Scripts for creating Campaign system tables to integrate Distributed Marketing with Campaign

Data source type	Script name
IBM DB2	c1b_systab_db2.sql
Microsoft SQL Server	c1b_systab_sqlsvr.sql
Oracle	c1b_systab_ora.sql

To populate the rows in the system tables that are created by the **c1b_systab_DB_type.sql** script, run the **c1b_populate_tables.sql** script.

Creating and populating the eMessage system tables manually

For eMessage, you must create additional system tables in the Campaign schema and then populate the tables with initial data. The Campaign installer automatically creates and populates the eMessage system tables in the Campaign schema if you select the option to create system tables automatically. However, if you do not select the option, you must create and populate the eMessage system tables manually.

Use your database client to run the appropriate script against the Campaign database.

Note: The eMessage installation directory is a subfolder within the Campaign directory.

Scripts for creating eMessage tables

IBM provides the **ace_op_systab** scripts for creating eMessage tables in your local environment.

If your Campaign system tables are configured for Unicode, find the appropriate script in the **ddl/unicode** directory under your eMessage installation.

If your Campaign system tables are not configured for Unicode, use the non-Unicode scripts in the **ddl** directory under your eMessage installation. Use the following table to see the scripts that you must use:

Table 14. Scripts for creating eMessage tables

Data source type	Script name
IBM DB2	ace_op_systab_db2.sql The user tablespace and system temporary tablespace where the system tables reside must each have a page size of 16K or greater.
Microsoft SQL Server	ace_op_systab_sqlsvr.sql
Oracle	ace_op_systab_ora.sql

Scripts for populating eMessage tables

IBM provides the `ace_op_populate_systab` scripts for populating the eMessage tables in your local environment.

The populate scripts are stored in the `ddl` directory of your eMessage installation. IBM provides only one version of the populate scripts because you can use them for either Unicode or non-Unicode tables.

Note: The eMessage installation directory is a subfolder within the Campaign directory.

Use the following table to see the scripts that you must use:

Table 15. Scripts for populating eMessage tables

Data source type	Script name
IBM DB2	<code>ace_op_populate_systab_db2.sql</code>
Microsoft SQL Server	<code>ace_op_populate_systab_sqlsvr.sql</code>
Oracle	<code>ace_op_populate_systab_ora.sql</code>

Registering Campaign manually

During the installation process, if the Campaign installer cannot access the Marketing Platform system tables, you must run the **configTool** utility to register it manually.

You can use the **configTool** and **populateDb** utilities to import and populate the Campaign information in the Marketing Platform system tables.

1. Run the **populateDb** utility by using the following example command as a guideline.

```
populateDb.bat -n Campaign
```

The command imports security roles and permissions into the default partition.

2. **If you are upgrading Campaign**, unregister Campaign by using the following command:

```
configTool -x -p "Affinium|Campaign|about" -f exportedAbout.xml
```

The command exports the Campaign **About** node to the `exportedAbout.xml` file.

3. Run the **configTool** utility by using the following example commands as guidelines:

- **configTool -r Campaign -f "full_path_to_Campaign_installation_directory\conf\campaign_configuration.xml"**
- **configTool -v -i -p "Affinium|suite|uiNavigation|settingsMenu" -f "full_path_to_Campaign_installation_directory\conf\campaign_setup_navigation.xml"**
- **configTool -v -i -p "Affinium|suite|uiNavigation|mainMenu" -f "full_path_to_Campaign_installation_directory\conf\campaign_navigation.xml"**

- `configTool -v -i -p "Affinium|suite|uiNavigation|mainMenu|Analytics" -f "full_path_to_Campaign_installation_directory\conf\campaign_analysis_navigation.xml"`
- `configTool -v -i -p "Affinium|suite|uiNavigation|alerts" -f "full_path_to_Campaign_installation_directory\conf\campaign_alerts.xml"`

The commands import configuration properties and menu items. You must run the utility as many times as there are files.

4. Run the following command after you register Campaign:

```
configtool -i -p "Affinium|Campaign" -f exportedAbout.xml
```

The command imports the Campaign **About** node to the `exportedAbout.xml` file.

Registering eMessage manually

During the installation process, if the eMessage installer cannot access the Marketing Platform system tables, you must run the **configTool** utility to register it manually.

By default, the Campaign installer automatically registers eMessage with the Marketing Platform system tables without enabling eMessage. In some situations, the Campaign installer does not connect with the Marketing Platform system tables to automatically register eMessage.

If the installer does not register eMessage automatically, you must register eMessage manually with the **configTool** utility that is provided with the IBM Marketing Software installation. The **configTool** utility is in the `tools\bin` directory under your Marketing Platform installation.

To register eMessage manually, use the following command to run the **configTool** utility:

```
configTool -r eMessage -f "full_path_to_eMessage_installation_directory\conf\emessage_configuration.xml"
```

The eMessage installation directory is a subdirectory of the Campaign installation directory.

For more information about registering and configuring eMessage, see the *IBM eMessage Startup and Administrator's Guide*.

Setting data source variables in the Campaign startup script (UNIX only)

Data source variables are automatically set by the installer during installation of Campaign. You can change these settings in the `Campaign/bin/setenv.sh` file. You must restart the server every time you modify the `setenv.sh` file.

During Campaign installation, the IBM Installer collects database information from you and uses that information to automatically configure the database and environment variables that are required to create and use the Campaign system tables. Those settings are stored in the `setenv.sh` file in the `bin` directory under your Campaign server installation.

For access to data sources (such as Campaign customer tables) that do not use the same type of database as the system tables, you must manually configure the `setenv.sh` file to add the database and library environment variables that are described in “Database and library environment variables.”

Note that if you modify this file when the Campaign server is already running, you must restart the server before changes to the `setenv` file are recognized. See “Starting the Campaign server” on page 71 for more details.

See Distributed Marketing database information worksheet for the required information to add to the `setenv` file.

Database and library environment variables

Set database and library environment variables that are required for your databases (customer tables and system tables, if you chose Manual Database Setup during installation) and operating system. You can set the database and library variables in the `setenv.sh` file.

Following table describes database names with their syntax and descriptions.

Table 16. Database environment variables

Database	Syntax and Description
DB2	<code>DB2DIR=full_dir_path</code> <code>export DB2DIR</code> DB2 installation directory (for example, <code>/usr/lpp/db2_06_01</code>) <code>. full_path_to_db2profile</code> Sources the database configuration for DB2 users (e.g. <code>/home/db2inst1/sqllib/db2profile</code>). Note the ". " (period and then a space).
Netezza	<code>NZ_ODBC_INI_PATH=full_dir_path</code> <code>export NZ_ODBC_INI_PATH</code> Directory location of the <code>odbc.ini</code> file (e.g. <code>/opt/odbc64v51</code>) <code>ODBCINI=full_path_and_file_name</code> <code>export ODBCINI</code> Full path to the <code>odbc.ini</code> file
Oracle	<code>ORACLE_BASE=full_dir_path</code> <code>export ORACLE_BASE</code> Oracle installation directory <code>ORACLE_HOME=full_dir_path</code> <code>export ORACLE_HOME</code> Oracle's home directory (e.g. <code>/home/oracle/OraHome1</code>).

Table 16. Database environment variables (continued)

Database	Syntax and Description
Teradata	<code>ODBCINI=full_path_and_file_name</code> <code>export ODBCINI</code> Full path to the obdc.ini file

Define the library environment variable as mentioned in the following table, depending on your UNIX operating system.

Table 17. Library environment variable

Operating System	Value
SUSE and RHEL Linux	<code>LD_LIBRARY_PATH</code> For example: <code>LD_LIBRARY_PATH=<Campaign_Home>/bin:<path to DB lib directory>:\$LD_LIBRARY_PATH</code> <code>export LD_LIBRARY_PATH</code> Note: If <code>LD_LIBRARY_PATH_64</code> (for 64-bit linking) is set, remove it. The <code>LD_LIBRARY_PATH</code> variable will be ignored if <code>LD_LIBRARY_PATH_64</code> is set.
AIX	<code>LIBPATH</code> For example: <code>LIBPATH=<Campaign_Home>/bin:<path to DB lib directory>:/usr/lib:\$ORACLE_HOME/lib32:\$ORACLE_HOME/lib</code>

Library directories for Oracle databases

Different versions of Oracle have different naming conventions for their lib directories. Older versions used lib for 32-bit and lib64 for 64-bit. Newer versions use lib32 for 32-bit and lib for 64-bit.

If you install 32-bit Campaign, you must include either `$ORACLE_HOME/lib32` or `$ORACLE_HOME/lib`, whichever one contains the 32-bit Oracle libraries.

If you install 64-bit Campaign, you must include either `$ORACLE_HOME/lib` or `$ORACLE_HOME/lib64`, whichever one contains the 64-bit Oracle libraries.

Note: Do not include paths to both the 32- and 64-bit libraries; include only the path to the library you are using for your version of Campaign.

Chapter 6. Deploying the Campaign web application

You can deploy the Campaign web application by using an EAR file or you can deploy the individual WAR files.

To deploy Campaign, follow the guidelines in this section and then start the Campaign server.

When you ran the IBM installer, you may have included Campaign in an EAR file, or you may choose to deploy the Campaign WAR file. If you included the Marketing Platform or other products in an EAR file, you must follow all the deployment guidelines detailed in the individual install guides for the products included in the EAR file.

You must know how to work with your web application server. Consult your web application server documentation for details such as navigation in the Administration console.

Note: Information about using web application clustering in IBM Campaign is provided in Chapter 12, “IBM Campaign web application clustering,” on page 115.

Setting the web application session timeout

Session timeout determines how long an inactive HTTP session can remain open before the session expires. If required, you can set the web application session timeout for Campaign by using the WebSphere console or the WebLogic console to adjust the value of session timeout in seconds or minutes.

Set the session timeout on your web application server:

- WebSphere: Use the IBM WebSphere Application Server Administrative Console to set the Session timeout in minutes. You can adjust this setting at the server and Enterprise Application levels. See your WebSphere documentation for details.
- WebLogic: Use the WebLogic console to set the Session Timeout in seconds, or adjust the **TimeoutSecs** parameter value for the **session-descriptor** element in the weblogic.xml file.

Deploying Campaign on WebSphere Application Server

You can deploy Campaign runtime environment on supported versions of WebSphere Application Server (WAS) from a WAR file or EAR file.

- Make sure that multiple language encoding is enabled in WAS.
- When you run the Install New Application wizard, ensure that you set the JDK Source Level to 18.

Important: If you are using IBM WebSphere Application Server 8.5.5.x Fix Pack 9 or earlier, you must use the following workaround to disable the annotation scanning for the xstream jar and to allow the application to be deployed correctly.

Go to the `app_server_root/properties` folder in the WebSphere installation. In the `amm.filter.properties` file, add the following line under Ignore-Scanning-Packages.

Deploying Campaign on WAS from a WAR file

You can deploy the Campaign application from a WAR file on WAS.

Complete the following tasks before you deploy Campaign:

- Confirm that your version of WebSphere meets the requirements in the *Recommended Software Environments and Minimum System Requirements* document, including any necessary fix packs or upgrades.
- Confirm that you created the data sources and database provider in WebSphere.
- Please follow the below steps while deploying Campaign.war file on IBM Websphere Application Server 8.5.5.x Fix Pack 9 or below versions

a) Go to <app_server_install_directory>/properties location

Example: /data/webservers/IBM/WebSphere8.5.5/properties

b) Edit "amm.filter.properties" file to include "com.thoughtworks.xstream"

Example:

```
Ignore-Scanning-Packages = javassist, \
organtlr, \
org.apache.avalon, \
org.apache.batik, \
org.apache.bcel, \
org.apache.commons, \
org.apache.log4j, \
org.apache.poi, \
org.apache.regexp, \
org.apache.xalan, \
org.apache.xerces, \
org.apache.xml, \
org.apache.xpath, \
org.ajax4jsf, \
org.dbunit, \
org.richfaces, \
org.jboss.el, \
org.jchrontab, \
org.omg, \
org.seasar, \
org.slf4j, \
com.thoughtworks.xstream
```

By doing this, WebSphere will disable the annotation scanning for the xstream jar

and allow the application deployed properly.

This is also mentioned in here <http://x-stream.github.io/faq.html>

This steps are not required for Websphere version 8.5.0.x or 8.5.5 Fix Pack 10 or higher.

1. Go to the WebSphere Integrated Solutions Console.
2. Complete the following steps if your system tables are in DB2:

- a. Click the data source that you created. Go to the Custom Properties for the data source.
- b. Select the Custom properties link.
- c. Set the value for the **resultSetHoldability** property to 1.
If you do not see the **resultSetHoldability** property, create the **resultSetHoldability** property and set its value to 1.
3. Go to **Applications > Application Types > WebSphere enterprise applications** and click **Install**.
4. In the Preparing for the application installation window, select the **Detailed - Show all options and parameters** check box and click **Next**.
5. Click **Continue** to see the Install New Application wizard.
6. Accept the default settings on the windows of the Install New Application wizard except the following windows:
 - In step 1 of the Install New Application wizard, select the **Precompile JavaServer Pages files** check box.
 - In step 3 of the installation wizard, set the **JDK Source Level** to 18.
 - In step 9 of the installation wizard, set the **Context Root** to /Campaign.
7. In the left navigation panel of WebSphere Integrated Solutions Console, navigate to **Applications > Application Types > WebSphere enterprise applications**.
8. In the Enterprise Applications window, click the Campaign.war file.
9. In the **Web Module Properties** section, click **Session Management** and select the following check boxes:
 - **Override session management**
 - **Enable Cookies**
10. Click **Enable Cookies**, and in the **Cookie name** field, enter a unique cookie name.
11. If you are using version 8 of WebSphere Application Server, select **Servers > WebSphere application server > server 1 > Session management > Enable Cookies** and clear the check box for **Set session cookies to HTTPOnly to help prevent cross-site scripting attacks**.
12. In the **Applications > Enterprise Applications** section of the server, select the WAR file that you deployed.
13. In the **Detail Properties** section, select **Class loading and update detection**.
14. In the **Class loader order** section, select the **Classes loaded with local class loader first (parent last)** option.
15. For **WAR class loader policy**, select **Single class loader for application**.
16. Start your deployment.

Deploying Campaign on WAS from an EAR file

You can deploy Campaign by using an EAR file if you included Campaign in an EAR file when you ran the IBM Marketing Software installer.

- Confirm that your version of WebSphere meets the requirements in the *Recommended Software Environments and Minimum System Requirements* document, including any necessary fix packs or upgrades.
- Confirm that you created the data sources and database provider in WebSphere.
 1. Go to the WebSphere Integrated Solutions Console.
 2. Complete the following steps, if your system tables are in DB2:

- a. Click the data source that you created. Go to the **Custom Properties** for the data source.
- b. Select the Custom properties link.
- c. Set the value for the **resultSetHoldability** property to 1.
If you do not see the **resultSetHoldability** property, create the **resultSetHoldability** property and set its value to 1.
3. Go to **Applications > Application Types > WebSphere enterprise applications** and click **Install**.
4. In the Preparing for the application installation window, select the **Detailed - Show all options and parameters** check box and click **Next**.
5. Click **Continue** to see the Install New Application wizard.
6. Accept the default settings on the windows of the Install New Application wizard except the following windows:
 - In step 1 of the Install New Application wizard, select the **Precompile JavaServer Pages files** check box.
 - In step 3 of the installation wizard, set the **JDK Source Level** to 18.
 - In step 9 of the installation wizard, set the **Context Root** to /Campaign.
7. In the left navigation panel of WebSphere Integrated Solutions Console, navigate to **Applications > Application Types > WebSphere enterprise applications**.
8. In the Enterprise Applications window, select the EAR file that you want to deploy.
9. In the **Web Module Properties** section, click **Session Management** and select the following check boxes:
 - **Override session management**
 - **Enable Cookies**
10. Click **Enable Cookies**, and in the **Cookie name** field, enter a unique cookie name.
11. If you are using version 8 of WebSphere Application Server, select **Servers > WebSphere application server > server 1 > Session management > Enable Cookies** and clear the check box for **Set session cookies to HTTPOnly to help prevent cross-site scripting attacks**.
12. In the **Detail Properties** section, select **Class loading and update detection**.
13. In the **Class loader order** section, select the **Classes loaded with parent class loader first** option.
14. For **WAR class loader policy**, select the **Class loader for each WAR file in application** option.
15. Start your deployment.
For more information about WebSphere Application Server version 8.5, see Welcome to the WebSphere Application Server information center.

Deploying Campaign on WebLogic

You can deploy IBM Marketing Software products on WebLogic.

Use the following guidelines when you deploy Campaign on WebLogic:

- IBM Marketing Software products customize the JVM used by WebLogic. You might need to create a WebLogic instance that is dedicated to IBM Marketing Software products if you encounter JVM-related errors.

- Verify that the SDK selected for the WebLogic domain you are using is the Sun SDK by looking in the startup script (`startWebLogic.cmd`) for the `JAVA_VENDOR` variable. It should be set to: `JAVA_VENDOR=Sun`. If it is set to `JAVA_VENDOR=BEA`, JRockit has been selected. JRockit is not supported. To change the selected SDK, refer to the WebLogic documentation.
- Deploy the IBM Marketing Software products as web application modules.
- On UNIX systems, you must start WebLogic from the console to allow correct rendering of graphical charts. The console is usually the machine on which the server is running. However, in some cases the web application server is set up differently.

If a console is not accessible or does not exist, you can emulate a console using Exceed. You must configure Exceed so that your local Xserver process connects to the UNIX machine in root window or single window mode. If you start the web application server using Exceed, you must keep Exceed running in the background to allow the web application server to continue running. Contact IBM Technical Support for detailed instructions if you encounter problems with chart rendering.

Connecting to the UNIX machine via telnet or SSH always causes problems rendering charts.

- If you are configuring WebLogic to use the IIS plug-in, review the WebLogic documentation.
- If you are deploying in a production environment, set the JVM memory heap size parameters to at least 1024 by adding the following line to the `setDomainEnv` script:

```
Set MEM_ARGS=-Xms1024m -Xmx1024m -XX:MaxPermSize=256m
```

- Under certain circumstances, deploying older legacy interactive channels or interactive channels with large deployment histories can stress the system and require 2048mb or greater of Campaign design time and/or Interact runtime Java heap space.

System administrators can adjust the amount of memory available to the deployment systems via the following JVM parameters:

```
-Xms####m -Xmx####m -XX:MaxPermSize=256m
```

Where the characters `####` should be 2048 or higher (depending on their system load.) Note that a 64-bit application server and JVM are usually necessary for values greater than 2048.

These are the suggested minimum values. Analyze your sizing requirements to determine correct values for your needs.

Configuring WebLogic to display reports (UNIX)

If you install Campaign with Contact Optimization on a UNIX system, you must enable the `java.awt.headless` JVM property so that the WebLogic web application server can display the graphs in Contact Optimization reports.

To enable graph display in Optimize reports in the Weblogic JVM:

1. Shut down the Weblogic server if it is already running.
2. Locate and open the startup script for your Weblogic server (`startWebLogic.sh`) in any text editor.
3. Modify the `JAVA_OPTIONS` parameter to add the following value:
`-Djava.awt.headless=true`
4. Save the startup script, then restart the Weblogic server.

Deploying Campaign on JBoss Application Server

You can deploy IBM Marketing Software products on Jboss application server.

Use the following guidelines when you deploy Campaign on Jboss:

- Verify that your version of JBoss meets the requirements that are described in the *Recommended Software Environments and Minimum System Requirements* document, including any necessary fix packs or upgrades.
- IBM Marketing Software products customizes the JVM used by Jboss. You might need to create a Jboss instance that is dedicated to IBM Marketing Software products if you encounter JVM-related errors
- If you are deploying in a production environment, set the JVM memory heap size parameters to at least 1024 by adding the following line to the standalone.conf in the JBOSS_EAP_HOME/bin

```
e.g : JAVA_OPTS="-Xms1024m -Xmx1024m -XX:MetaspaceSize=96M  
-XX:MaxMetaspaceSize=256m "
```

These are the suggested minimum values. Analyze your sizing requirements to determine correct values for your needs. Depending upon system load **-Xmx** and **MaxMetaspaceSize** value should be adjusted. Note that a 64-bit application server and JVM are usually required for values greater than 2048.

- Modify the JAVA_OPTIONS parameter to add the following value in setenv.bat/sh.

```
set "JAVA_OPTS=%JAVA_OPTS%  
-DUNICA_PLATFORM_CACHE_ENABLED=true  
-Dclient.encoding.override=UTF-8"
```

- Go to the JBoss Enterprise Application management Console e.g http://<Jboss_Host_IP>:<Port>/management
- Deploy Campaign.war/EAR on JBOSS.
- Alternately you can also deploy Campaign.war using jboss-cli.bat/sh.
e.g Execute the command: jboss-cli.bat/sh and connect to Jboss instance using the command - connect <Jboss_Host_IP>:<JBOSS management Port>
deploy <Campaign_Install_Path/Campaign/Campaign.war
- After all these changes restart the Jboss application server.

Deploying Campaign on Tomcat Application Server

You can deploy IBM Marketing Software products on Tomcat application server.

Use the following guidelines when you deploy Campaign on Tomcat:

- IBM Marketing Software products customize the JVM used by Tomcat. You might need to create a Tomcat instance that is dedicated to IBM Marketing Software products if you encounter JVM-related errors
- If you are deploying in a production environment, set the JVM memory heap size parameters to at least 1024 by adding the following line to the setenv.bat/sh
e.g : set "CATALINA_OPTS=%CATALINA_OPTS% -Xms1024m -Xmx1024m -XX:MaxPermSize=512m"

These are the suggested minimum values. Analyze your sizing requirements to determine correct values for your needs. Depending upon system load **-Xmx** value should be adjusted. Note that a 64-bit application server and JVM are usually necessary for values greater than 2048.

- Modify the JAVA_OPTIONS parameter to add the following value in setenv.bat/sh.

```
set "JAVA_OPTS=%JAVA_OPTS%
-DUNICA_PLATFORM_CACHE_ENABLED=true
-Dclient.encoding.override=UTF-8"
```

- You must add Platform and Campaign Data Source in Campaign.xml along with the path of Campaign.war to the Campaign Tomcat instance eg. :

```
<Context docBase="<Campaign_Install_Path>/Campaign/Campaign.war">
<Resource name="<UnicaPlatformDS>" factory="com.unica.manager.tomcat.utils.TomcatDSFactory" autoCommit="true"
maxActive="30" maxIdle="10" maxWait="10000"
username="<db user for Platform schema>" password="<db password>" driverClassName=
"<db specific class name>" url="<db specific jdbc url>"/>

<Resource name="<campaignPartitionDS>" factory="com.unica.manager.tomcat.utils.TomcatDSFactory" autoCommit="true"
maxActive="30" maxIdle="10" maxWait="10000"
username="<db user for Campaign schema>" password="<db password>" driverClassName=
"<db specific class name>" url="<db specific jdbc url>" />
</Context>
```

Note - {{You can encrypt DB password using }} encryptTomcatDBPasswords utility located in <Platform>

- Restart the Tomcat application server.

Starting the Campaign server

Make sure that Marketing Platform and Campaign web applications are deployed and running when you start the Campaign server.

You can start the Campaign server directly, or install it as a service.

Starting the Campaign listener manually

Start the Campaign listener by running the cmpServer.bat file for Windows, or the rc.unica_ac command for UNIX.

Follow the instructions appropriate for your operating system.

Windows

Start the Campaign listener by running the cmpServer.bat file in the bin directory under your Campaign installation. When the unica_aclnr.exe process appears on the Windows Task Manager Processes tab, the server has started successfully.

UNIX

Start the Campaign listener by executing the rc.unica_ac program with a start argument. You must run this command as root. For example:

```
./rc.unica_ac start
```

To determine whether the unica_aclnr process has started successfully, enter this command:

```
ps -ef | grep unica_aclnr
```

To determine the process ID for the server you started, view the unica_aclnr.pid file located in the conf directory under your Campaign installation.

Note: Since Campaign supports listener clustering, a loop is added in the rc.unica_ac which automatically restarts the listener if in case the listener goes down or crashes. Also, process rc.unica_ac still remains after starting listener and the parent process may change. For example, if you are upgrading from Release

8.6.0.4 to Release 9.1.1, the parent process of `unica_aclsnr` can be changed to `rc.unica_ac` instead of `init`, which was the former parent process and the process `rc.unica_ac` still remains after starting the listener.

Installing the Campaign listener as a Windows service

You can install the Campaign listener as a Windows service so that it starts automatically whenever you start Windows.

1. Add the `bin` directory under the Campaign installation directory to the user `PATH` environment variable. If the `PATH` environment variable does not exist for the user, create it.

Make sure that you add this path to the user `PATH` variable, not the system `PATH` variable.

If the Campaign `bin` directory exists in the system `PATH` environment variable, remove it. You do not need it in the system `PATH` environment variable to install the Campaign listener as a service.

2. If you are upgrading from a previous version of Campaign that had the server installed as a service, stop the service.
3. Open a command window and change directories to the `bin` directory under your Campaign installation.
4. Run the following command to install the Campaign listener as a Windows service:

```
unica_aclsnr -a
```

Note: The `-a` option includes auto-restart functionality. If you do not want the service to attempt to restart automatically, use `unica_aclsnr -i`.

The listener is now installed as a service.

Note: Make sure `CAMPAIGN_HOME` has been created as a system environment variable before starting the Campaign listener service.

5. Open the Unica Campaign Listener Service properties dialog box. Click the **Log On** tab.
6. Select **This account**.
7. Enter the user name (system user) and password, and start the services.

Chapter 7. Configuring after deployment

After you deploy Campaign, you must verify that the Campaign listener is running, set up the Campaign system user, set Campaign configuration properties, and verify the Campaign installation.

If you are using the IBM Marketing Software reporting feature, you must complete the tasks that are described in the *IBM Marketing Software Reports Installation and Configuration Guide*.

Note: If you plan to enable IBM eMessage for hosted email, you must use the IBM Marketing Software reporting feature to view standard eMessage performance reports.

Verifying that the Campaign listener is running

The Campaign listener must be running before users can work with any Campaign features. The listener automatically creates a separate `unica_acsvr` process for each login and each active flowchart. For example, if one user logs in and then opens a flowchart, the listener creates two instances of `unica_acsvr.exe`.

Use the following procedure to verify that the Campaign listener is running.

1. Use the appropriate procedure for your operating system:

In Windows, look for `unica_aclsnr.exe` on the **Processes** tab of the Windows Task Manager.

In UNIX, use the `ps` command to look for the Campaign server, as in the following example: `ps -ef | grep unica_aclsnr`.

2. If the listener is not running, restart it:

In Windows, run the `cmpServer.bat` script that is in the `bin` directory under your Campaign installation.

In UNIX, enter the following command at the system prompt: **`rc.unica_ac start`**

For important details about running the listener, including how to start it automatically, see the *IBM Campaign Administrator's Guide*.

Setting up the Campaign system user

Set up the Campaign system user to directly access databases. For multiple partitions in Campaign, create a system user for each partition.

A system user is a IBM Marketing Software user account that is configured for use by IBM applications.

To avoid presenting users with repeated prompts for login credentials you can associate a system user with one or more data sources. Each data source specifies a user name and password. This allows you to provide a user name and password to access a database or other protected resource by referencing the data source. By adding multiple data sources to the configuration of the system user account, you can allow the system user to access multiple databases.

In Campaign, the system user holds login credentials for accessing system tables and other data sources.

Use an existing or new IBM Marketing Software user account to save credentials for the data sources described here.

You set up IBM Marketing Software users and assign user data sources in the **Settings > Users** area of IBM Marketing Software. See the online help in that section for instructions on how to do this.

Set up a user account to hold credentials for the following data sources.

- Campaign system tables (UA_SYSTEM_TABLES)
- All customer (user) tables

On UNIX, for the **Alternate Login** attribute of the system user, enter the UNIX account of a user in a group that shares privileges with the UNIX users of Campaign.

Note: If you have multiple partitions, each partition must have its own system user. The system user cannot be the same across partitions.

Adding data source properties on the Configuration page

Use the appropriate data source template to add data source properties on the Configuration page for each Campaign data source.

The Campaign installer imports the template for the database type that you specified for the Marketing Platform database when you ran the IBM installer.

If you need other data source templates for additional database types, you must import them manually using the Marketing Platform configTool utility. You can import as many templates as you need, to correspond to each different type of database you have.

For example, suppose your installation of Marketing Platform and Campaign used the following databases:

- Oracle - system tables
- DB2 - customer (user) tables
- DB2 - additional customer (user) tables

In this case, you would import the DB2Template.xml template for the two sets of customer (user) tables.

If your Marketing Platform and Campaign system tables databases are of the same database type, the installer automatically imports the template that is used for the system tables; in this example, it imports the Oracle template.

For instructions, see “Importing data source templates” on page 75.

You create a new category from the template, which creates a new set of data source configuration properties. Create as many new categories as needed, one for each data source of that type. In the example above, you would use the Oracle template to create one new category and the DB2 template to create two new categories. See “Duplicating a data source template” on page 75.

After you add the data source properties, set the data source configuration properties in the categories you created from the templates.

For instructions, see “Data source properties” on page 76.

Importing data source templates

The Campaign system tables data source (UA_SYSTEM_TABLES) is supported only on Oracle, DB2, and SQLServer. Import the data source templates for the user tables by using the **configTool** utility to support the database types that are not supported for the Campaign system tables.

The Campaign data source templates are located in the conf directory under your Campaign installation.

To import and export templates, you use the Marketing Platform configTool utility, located in the tools/bin directory under your Marketing Platform installation. If you are not familiar with configTool, see “configTool” on page 97 for details on performing this task.

Here is an example of the command you would use to import the Oracle template into the default partition (Windows environment).

```
configTool -i -p "Campaign|partitions|partition1|dataSources" -f  
full_path_to_directory_containing_your_Oracle_template\OracleTemplate.xml
```

Duplicating a data source template

Duplicate the data source templates to create a new set of configuration properties in the data source category.

1. On the Configuration page, navigate to the data source template you want to duplicate.

Unlike other categories, template category labels are in italics and enclosed in parentheses.

2. Click the data source template.

The Create category from Template page displays.

3. Enter a name in the **New category name** field (required).

Note: The data source category name for the Campaign system tables **MUST** be UA_SYSTEM_TABLES.

4. If desired, edit properties within the new category. You can also do this later.
5. Click **Save and Finish**.

The new category appears in the navigation tree.

Campaign configuration properties

You must specify configuration properties on the Configuration page for a basic installation of Campaign. Additionally, you can use the Configuration page to specify the properties that perform important functions that you can optionally adjust.

Data source properties

The following table provides information about the properties that you must specify for each Campaign data source:

Table 18. Properties for each Campaign data source

Property name	Description
ASMUserForDBCredentials	This property must be the same user you already created as the Campaign system user in "Setting up the Campaign system user" on page 73.
DSN	For SQL server, set this property to the DSN (data source name) you created. For Oracle and DB2, set this property to the database name or the SID (service) name.
JndiName	Set this property to the JNDI you created in your application server to connect to this specific data source. For Tomcat, JNDI name should be prefixed with "java:/comp/env/" for example: campaignPartition1DS should be mentioned as java:/comp/env/campaignPartition1DS
SystemTableSchema	Not required for SQL server. For other data sources, set this property to the user of the database to which you are trying to connect.
OwnerForTableDisplay	Not required for SQL server. For other data sources, set this property to the user of the database to which you are trying to connect.

The data sources are the Campaign system table database and any customer (user) databases you plan to use with Campaign.

Note: The data source category name for the Campaign system tables **must** be UA_SYSTEM_TABLES.

For details on setting the values, see the context help for these properties or the *IBM Marketing Platform Administrator's Guide*.

In addition to creating and setting data source properties, you must set the following properties on the Configuration page for a basic installation of Campaign.

- Campaign > unicaACListener > serverHost
- Campaign > unicaACListener > serverPort
- For the default partition, set values as desired in the categories under Campaign > partitions > partition1.

When you change a property, you must restart the Campaign listener for the changes to take effect.

Mapping user tables in Campaign

Mapping user tables is the process of making external data sources accessible in Campaign. A typical user table contains information about your company's customers, prospects, or products. The data source can be a database table or ASCII flat file. You must map any data sources that you configured to make that data accessible to processes in flowcharts.

For instructions on how to map user tables, see the *Campaign Administrator's Guide*.

Note: User tables are different from system tables. Most Campaign system tables are mapped automatically during initial installation and configuration if you use the system table data source name UA_SYSTEM_TABLES. If a connection issue forces you to map system tables manually, you must log out of Campaign and log in again after you map the tables.

Verifying the Campaign installation

If you performed all of the steps to install and configure Campaign, deploy the Campaign web application, and configure Campaign after deployment, you are ready to verify your installation.

If you have not already done so, log in to IBM Marketing Software as a user that exists in the Campaign Administrators role (such as asm_admin). In **Settings > Users**, assign at least one security role (for example, Global Policy/Admin) to the new user. After the new user is assigned the role, you can log in to Campaign as that new user.

Use the following procedure to verify the installation.

1. Log in to IBM Marketing Software.
2. Verify that all system tables exist in the **Settings > Campaign Settings > Manage Table Mappings** window.
3. Create a campaign and create a flowchart in that campaign.

Setting properties for integration with IBM Marketing Software products

Campaign integrates with various IBM applications. If required, you can specify configuration properties to set up an integration between Campaign and other IBM products.

Additional configuration for enhanced security

The procedures in this section describe some additional configurations in the web application server. They are optional, but you might want to perform them for enhanced security.

Disabling the X-Powered-By flag

If your organization is concerned that the X-Powered-By flag in the header variable is a security risk, you can disable it using this procedure.

1. If you are using Weblogic, in the administration console, under *domainName* > **Configuration > Web Application**, set **X-Powered-By Header** to **X-Powered-By Header will not be sent**.

2. If you are using WebSphere, perform the following steps.
 - a. In the WebSphere administration console, navigate to **Servers > Server Types > WebSphere application servers > *server_name* > Web Container Settings > Web container**.
 - b. Under **Additional Properties** select **Custom Properties**.
 - c. On the Custom Properties page, click **New**.
 - d. On the Settings page, create a custom property named `com.ibm.ws.webcontainer.disablePoweredBy` and set the value to `false`.
 - e. Click **Apply** or **OK**.
 - f. Click **Save** on the console task bar to save your configuration changes.
 - g. Restart the server.

Configuring a restrictive cookie path

In your web application server, you can restrict cookie access to a specific application for enhanced security. If not restricted, cookies are valid across all the deployed applications.

1. If you are using WebLogic, perform the following steps.
 - a. Extract the `weblogic.xml` file from the WAR or EAR package where you want to add the restrictive cookie path.
 - b. Add the following to the `weblogic.xml` file, where *context-path* is the context path for the deployed application. For IBM Marketing Software applications the context path is typically `/unica`.


```
<session-descriptor>
  <session-param>
    <param-name>CookiePath</param-name>
    <param-value>/context-path</param-value>
  </session-param>
</session-descriptor>
```
 - c. Rebuild the WAR or EAR file.
2. If you are using WebSphere, perform the following steps.
 - a. In the WebSphere administration console, navigate to the **Session Manager > Cookie** tab.
 - b. Set the **Cookie Path** to the context path of your application.
For IBM Marketing Software applications the context path is typically `/unica`.

Starting the Campaign Advance search agent

You can use the Campaign Advance search agent to search for flowcharts that you have created in Campaign. You can search for flowcharts only if the Advanced search agent is running when the a flowchart is created or modified.

IBM Marketing Platform and IBM Campaign web applications must be deployed and running before you start the Campaign Advance search agent.

To start the Campaign Advance search agent manually, complete the following steps:

1. For Windows, go to the bin directory under your Campaign installation and run `advSrchAgent.bat`.
The `advSrchAgent.bat` file keeps running at the prompt.
2. For Unix, go to the bin directory under your Campaign installation and run `advSrchAgent.sh`.

The advSrchAgent.sh file keeps running at the prompt.

For a Unix operating system, run the following command to start the Campaign Advance search agent.

```
./advSrchAgent.sh partition1
```

Usage:

```
./advSrchAgent.sh <start | stop> <partition_name> [cluster_home] [-u <user_name>] [-p <password>]
```

Where:

<start | stop> : Start or stop the agent.

<partition_name> : Name of the partition from which flowchart files are to be exported.

[cluster_home] : Home directory of clustered environment, if clustering is enabled.

[-u <user_name>] : User name for user who has read access to all flowchart files in the partition

[-p <password>] : Password of the user specified

Chapter 8. Configuring multiple partitions for Campaign

In the Campaign family of products, partitions provide a way to secure the data that is associated with different groups of users. When you configure Campaign or a related IBM Marketing Software application to operate with multiple partitions, users can see each partition as a separate instance of the application. There is no indication that other partitions exist on the same computer.

How multiple partitions work

For the IBM Marketing Software applications that operate with Campaign, you can configure the application within the partitions where you have configured an instance of Campaign. Application users within each partition can access the Campaign functions, data, and customer tables that are configured for Campaign in the same partition.

Partition benefits

Multiple partitions are useful for setting up strong security between groups of users, because each partition has its own set of Campaign system tables. Multiple partitions cannot be used when groups of users must share data.

Each partition has its own set of configuration settings, and Campaign can be customized for each group of users. However, all partitions share the same installation binaries. With the same binaries for all partitions, you can minimize the installation and upgrade efforts for multiple partitions.

Partition user assignment

Access to partitions is managed through membership in Marketing Platform groups.

With the exception of the partition superuser (`platform_admin`), each IBM user can belong to one partition. A user who requires access to multiple partitions must have a separate IBM user account for each partition.

If there is only one Campaign partition, users do not have to be explicitly assigned to that partition to have access to Campaign.

Partition data access

In a multi-partition configuration, partitions have the following security characteristics:

- A user cannot access a partition if the user is not a member of a group that is assigned to a partition.
- A user of one partition cannot view or change data in another partition.
- Users cannot navigate the Campaign file system above the root directory of their partition from the browse dialog boxes in Campaign. For example, if two partitions, named `partition1` and `partition2` exist, and you belong to the group that is associated with `partition1`, you cannot navigate the directory structure for `partition2` from a dialog box.

Setting up multiple partitions

Configure multiple partitions in Campaign to isolate and secure data for different groups of Campaign users. As each partition has its own set of configuration properties, you can customize Campaign for each group of users.

Before you configure additional partitions in Campaign, complete the following tasks for each partition that you are configuring:

1. Create database or schema for the Campaign system tables
2. Create ODBC or native connections
3. Create JDBC connections in the web application server

Complete the following tasks to set up multiple partitions for Campaign:

1. Create the system tables database and any other data sources required for your partitions. Then configure the JDBC and ODBC or native connections required to access your data sources.
2. For each partition, create a different schema in your database to hold the system tables. Use the database-specific scripts packaged with Campaign to create and populate the system tables.
3. Complete the following tasks to create a directory structure for each additional partition:

Note: Save a clean copy of the original partition1 directory to serve as a backup.

- a. In the partitions directory under your Campaign installation, create a duplicate of the default partition1 directory for each partition that you are adding, ensuring that you include all sub-directories.
 - b. Name each partition directory with a unique name. Use these exact same names for the partitions when you later create their configuration trees on the Configurations page. If you want to create a second partition and created the directory named Campaign/partitions/partition2, when you create the configuration tree on the Configurations page, you must use the name "partition2" to identify this partition.
 - c. Delete any files that exist in the duplicate partition sub-directories.
4. Complete the following tasks to use the **partitionTool** utility with the **-s** option to clone the default partition:

Note: If you decide not to use this option, you must stop the web application server where Marketing Platform is deployed before performing this procedure.

- a. Set the JAVA_HOME environment variable, either in the setenv script located in the tools/bin directory of your Marketing Platform installation, or in the command line window in which you run the partitionTool utility.
- b. Open a command line window and run the utility from the tools/bin directory under your Marketing Platform installation. Use the appropriate command and options (as described in the *Marketing Platform Administrator's Guide*) to achieve your desired results. **partitionTool -c -s partition1 -n partition2**
- c. Repeat this procedure for each new partition you need to create.
- d. When you are finished, stop and restart the application server where Marketing Platform is deployed to see the groups that were created.

Note: For complete information on using this utility, see “partitionTool” on page 103.

5. Complete the following tasks to create the partition structure on the Configuration page by using the New partitionTemplate for each new partition:

- a. On the Configuration page, navigate to **Campaign > partitions** and click (*partitionTemplate*).

If you do not see the (*partitionTemplate*) property in the list, use the configTool utility to import the partition template, using a command similar to the following:

```
configTool -i -p "Affinium|Campaign|partitions"  
-f <CAMPAIGN_HOME>/conf/partitionTemplate.xml
```

Replace *CAMPAIGN_HOME* with the actual path to your Campaign installation.

The configTool utility is found in the tools directory under your IBM Marketing Platform installation. See the *IBM Marketing Platform Administrator's Guide* for information on using this utility. In the right pane, you see the partitionTemplate pane with an empty **New category name** field.

- b. Enter the name for the new partition, using the same name that you used when creating the partition's directory structure in the file system in “Setting up multiple partitions” on page 82.
- c. Click **Save Changes**. You see the new partition structure, with the same categories and properties as in the partition template.

Partition superuser

To administer security across all users of Marketing Platform, a user account must exist that has access to all security settings and user accounts in the system.

By default, this user account is platform_admin. This user account does not belong to any one partition; instead, it has access to all user accounts in all partitions.

The IBM administrator can create additional users with the same level of access. To be a partition superuser, an account must have Admin access to the Marketing Platform and have full access to the Users, User Groups, and User Permissions pages. The partition superuser does not require access to product-specific security pages, such as Campaign security pages.

Configuring data source properties for partitions

You must configure data source properties for each partition that you create. Use the appropriate data source template to create data source properties.

The Campaign installer imports the template for the database type that you specified for the Marketing Platform database when you ran the IBM installer.

If you need other data source templates for additional database types, you must import them manually using the Marketing Platform **configTool** utility. You can import as many templates as you need to correspond to each different type of database you have.

For example, suppose your installation of Marketing Platform and Campaign used the following databases:

- Oracle - system tables
- DB2 - customer (user) tables

- DB2 - additional customer (user) tables

In this case, you would import the DB2Template.xml template for the two sets of customer (user) tables.

If your Marketing Platform and Campaign system tables databases are of the same database type, the installer automatically imports the template used for the system tables; in this example, it imports the Oracle template.

Note: When you create a new partition, you must use the **configTool** utility to import the data source templates for the system tables and user tables.

For instructions, see “Importing data source templates” on page 75.

You create a new category from the template, which creates a new set of data source configuration properties. Create as many new categories as needed, one for each data source of that type. In the example above, you would use the Oracle template to create one new category and the DB2 template to create two new categories. See “Duplicating a data source template” on page 75.

After you add the data source properties, set the data source configuration properties in the categories you created from the templates.

For instructions, see “Campaign configuration properties” on page 75.

Complete the following tasks to configure data source properties for each partition:

1. Add data source configuration properties to the Configuration page for each Campaign data source by using the appropriate data source template. The Campaign installer imports the template for the database type that you specified for the Marketing Platform database when you ran the IBM installer. If you need other data source templates for additional database types, you must import them manually using the Marketing Platform configTool utility. You can import as many templates as you need, to correspond to each different type of database you have. For example, suppose your installation of Marketing Platform and Campaign used the following databases:
 - Oracle - system tables
 - DB2 - customer (user) tables
 - DB2 - additional customer (user) tables

In this case, you would import the DB2Template.xml template for the two sets of customer (user) tables. If your Marketing Platform and Campaign system tables databases are of the same database type, the installer automatically imports the template used for the system tables; in this example, it imports the Oracle template. For instructions, see “Importing data source templates” on page 75.

2. Create a new category from the template, which creates a new set of data source configuration properties. Create as many new categories as needed, one for each data source of that type. In the example above, you would use the Oracle template to create one new category and the DB2 template to create two new categories. See “Duplicating a data source template” on page 75.
3. Set the data source configuration properties for each Campaign data source. See “Campaign configuration properties” on page 75 for more information.

Setting up system users for Campaign

You can associate a system user with one or more Marketing Platform data sources to avoid presenting users with repeated prompts for login credentials. Each data source specifies a user name and password. You can provide a user name and password to access a database or other protected resources by referencing the data source. By adding multiple data sources to the configuration of the system user account, you can allow the system user to access multiple databases.

IBM Marketing Software applications might require a system user account configured with the following attributes:

- Login credentials for accessing system tables or other data sources.
- Specific permissions to create, modify, and delete objects within the system.

For more information about setting up a new user and assigning a data source to a user, see the *IBM Marketing Platform Administrator's Guide*.

Complete the following actions to set up system users for Campaign:

1. Use an existing or a new user account to save credentials for the following data sources:
 - Campaign system tables
 - All customer (user) tables
2. On UNIX, for the **Alternate Login** attribute of the system user, enter the UNIX name for a user in a group that shares privileges with the UNIX users of Campaign.

Note: If you have multiple partitions, you must create a system user for each partition.

Using IBM Cognos Reports for multiple partitions

If you want to use IBM Cognos® Reports for multiple partitions in Campaign, eMessage, or Interact, you must configure IBM Cognos reports packages for each partition.

For instructions, see the *IBM Marketing Software Reports Installation and Configuration Guide*.

Assigning roles, permissions, and groups for partitions

Before you can use the partitions that you have configured for Campaign, you must assign roles to the user with administrator role in each partition. You must also assign a group to each partition.

Use the **partitionTool** utility to create a default admin user for each partition that you create.

Assign roles to each partition Admin user - The **partitionTool** utility creates a default Admin user for each partition you create. On the Users page, assign at least one security role (for example, Global Policy/Admin) to your new user. Once the new user has been assigned the role, you can log into the Campaign partition as that new user.

If you plan to enable IBM eMessage in multiple Campaign partitions, you must configure a corresponding eMessage partition for each Campaign partition. For

more information about creating additional partitions for eMessage, see Chapter 9, “Configuring multiple partitions in eMessage,” on page 87.

Chapter 9. Configuring multiple partitions in eMessage

Configure multiple partitions in eMessage to isolate and secure data for different groups of eMessage users. As each partition has its own set of configuration properties, you can customize eMessage for each group of users.

Installing eMessage creates a default partition for eMessage in the Marketing Platform. You can configure more partitions for eMessage. Each partition that you create for eMessage operates in conjunction with a partition created for Campaign.

Note: To configure multiple partitions in eMessage you must configure corresponding partitions in Campaign. Adding new partitions for eMessage requires making changes in the Marketing Platform configurations for eMessage and Campaign.

Important: After you change the eMessage and Campaign configurations, you must restart the web application server that hosts Campaign and restart the Response and Contact Tracker (RCT). You might also need to restart the Campaign listener.

Be sure to back up your existing configurations before you make changes.

Partitions for eMessage: Overview

Creating partitions for eMessage allows you to isolate and secure data for different groups of users. Each partition appears to users as a separate instance of eMessage, with no indication that other partitions exist on the same system. Because each partition has its own set of configuration properties, you can customize eMessage for each group of users.

Users within each partition can access only those functions, data, and customer tables configured for the partition. For example, if you create partitions named `partition1` and `partition2`, eMessage users working in `partition1` can select email recipients from customer tables configured in `partition1`, but not in `partition2`. IBM does not recommend creating multiple partitions if users need to share data.

When you work with multiple partitions, you should understand characteristics that are unique to partitions for eMessage and how partitions for eMessage relate to partitions for Campaign. You should also be familiar with the workflow for creating and configuring multiple partitions for eMessage.

Characteristics of partitions for eMessage

Note the following when you create and configure new partitions for eMessage.

- You create partitions for eMessage differently than you do for Campaign.
You create new partitions for eMessage using a partition template available in the eMessage configuration properties in Marketing Platform.
- The name of each eMessage partition must exactly match the name of its corresponding Campaign partition.
- Each partition that you create for eMessage must be able to connect to IBM Marketing Software Hosted Services.

You must request a separate IBM Marketing Software Hosted Services account for each partition. IBM will provide a user name and password connected to the account. You must configure a Marketing Platform data source that can provide these access credentials automatically when eMessage connects to IBM Marketing Software Hosted Services.

For more information about how to request an account, see the *IBM eMessage Startup and Administrator's Guide*.

Relationship to partitions in Campaign

Each partition in eMessage operates in conjunction with a specific partition created in the Marketing Platform for Campaign. The Campaign partition provides the following.

- Campaign schema that contains the eMessage system tables
- File structure for Campaign in the partition, including directories used by eMessage to create and process recipient lists
- Configuration properties associated with creating recipient lists and enabling eMessage within the partition

Because eMessage operates with Campaign within a specific partition, the eMessage and Campaign partition structures must specify the same name. The partition names must match exactly.

Roadmap for configuring multiple partitions in eMessage

To create a partition in eMessage, use the exact same name as the partition in Campaign that exists in the Marketing Platform configurations.

Before you create a new partition for eMessage, confirm that you have satisfied all the eMessage prerequisites for the partition in Campaign and eMessage.

Complete the following steps to create a new partition for eMessage.

1. "Creating a new partition for eMessage"
2. "Preparing eMessage system tables for the partition" on page 90
3. "Configuring a system user to access IBM Marketing Software Hosted Services" on page 92
4. "Enabling eMessage in Campaign for the new partition" on page 93
5. "Specifying the Recipient List Uploader location for eMessage" on page 93
6. "Restarting system components after configuring eMessage" on page 93
7. "Testing eMessage partition configuration and connections" on page 94

Creating a new partition for eMessage

Installing eMessage creates a default partition for eMessage in Marketing Platform. You can create multiple partitions for eMessage to isolate and secure data for different groups of users.

Before you create and configure partitions for eMessage, you must meet the following requirements for eMessage and Campaign:

- Before you create multiple partitions in eMessage, complete the following tasks for eMessage:

- Contact IBM Support to request an account and credentials for each partition. Each partition requires a separate IBM Marketing Software Hosted Services account and access credentials. For more information, see the *IBM Startup and Administrator's Guide*.
- Create a system user that can access the eMessage system tables that you will create in the Campaign schema for the partition.
You can update the system user that you created for the Campaign partition so that it can also access the eMessage system tables.
- Before you create multiple partitions in eMessage, complete the following tasks in Campaign:
 - Create a partition in Campaign to operate with the partition that you are creating for eMessage. Note the name of the partition.
 - Create Campaign system tables in the Campaign partition.
 - Configure a system user to access the system tables in the partition.
 If you have multiple partitions, each partition must have its own system user. The system user cannot be the same across partitions.

The IBM installer registers the eMessage configuration properties and a default partition during initial installation. The default partition includes a template that you can copy to create additional partitions.

Complete the following actions to create a new partition for eMessage:

1. Navigate to eMessage > partitions > (*partition*) and duplicate the partition template.
2. Name the new partition.

Note: eMessage does not support deleting a partition after you have created it.


Identifying the partition template

The Configuration page displays the eMessage partition template in the navigation tree for the default partition. You can identify the partition template in the tree because its label is in italics and enclosed in parentheses.

Naming the new partition

The following restrictions apply when you name a new partition.

- The name must be unique among categories that are siblings in the tree (that is, among categories that share the same parent category).
- The partition name may not start with a period. Also, the following characters are not allowed in partition names.



Note: Because eMessage operates with Campaign within a specific partition, the partitions for eMessage and Campaign must specify the same partition name.

Preparing eMessage system tables for the partition

For each partition that you create for eMessage, you must create, populate and configure the eMessage system tables that are used by the partition in the Campaign schema.

Complete the following actions to prepare the eMessage system tables for a partition:

1. Create the eMessage system tables. In your database client, run the SQL scripts for *creating* system tables against the Campaign database.
2. Populate the tables that you created. Use the database client to run the scripts for *populating* tables against the Campaign database. For more information about the SQL scripts, see the reference table for script names and locations at “Creating and populating the eMessage system tables manually” on page 60.
3. Set the following configuration properties in the eMessage configuration for the partition, specifying the user name and platform data source that you configured for the partition's Campaign system user:
 - eMessage > partitions > partition [n] < dataSources > systemTables > asmUserForDBCredentials
 - eMessage > partitions > partition [n] < dataSources > systemTables > amDataSourceForDBCredentials

eMessage accesses the partition's system tables through a system user configured in Marketing Platform. A Marketing Platform data source added to the system user provides the required access credentials. Because the eMessage system tables exist in the Campaign schema for the partition, you can use the system user that you created to access the Campaign schema to access the partition's eMessage system tables.

4. Update the following properties in the partition's configuration properties:
 - eMessage > partitions > partition [n] < dataSources > systemTables > type
 - eMessage > partitions > partition [n] < dataSources > systemTables > schemaName
 - eMessage > partitions > partition [n] < dataSources > systemTables > jdbcBatchSize
 - eMessage > partitions > partition [n] < dataSources > systemTables > jdbcClassName
 - eMessage > partitions > partition [n] < dataSources > systemTables > jdbcURI

See the Marketing Platform online help for each property to learn more about setting the configuration properties. For additional information about these configuration properties and configuring eMessage, see the *IBM eMessage Startup and Administrator's Guide*.

Creating and populating the eMessage system tables manually

For eMessage, you must create additional system tables in the Campaign schema and then populate the tables with initial data. The Campaign installer automatically creates and populates the eMessage system tables in the Campaign

schema if you select the option to create system tables automatically. However, if you do not select the option, you must create and populate the eMessage system tables manually.

Use your database client to run the appropriate script against the Campaign database.

Note: The eMessage installation directory is a subfolder within the Campaign directory.

Scripts for creating eMessage tables

IBM provides the ace_op_systab scripts for creating eMessage tables in your local environment.

If your Campaign system tables are configured for Unicode, find the appropriate script in the ddl/unicode directory under your eMessage installation.

If your Campaign system tables are not configured for Unicode, use the non-Unicode scripts in the ddl directory under your eMessage installation. Use the following table to see the scripts that you must use:

Table 19. Scripts for creating eMessage tables

Data source type	Script name
IBM DB2	ace_op_systab_db2.sql The user tablespace and system temporary tablespace where the system tables reside must each have a page size of 16K or greater.
Microsoft SQL Server	ace_op_systab_sqlsvr.sql
Oracle	ace_op_systab_ora.sql

Scripts for populating eMessage tables

IBM provides the ace_op_populate_systab scripts for populating the eMessage tables in your local environment.

The populate scripts are stored in the ddl directory of your eMessage installation. IBM provides only one version of the populate scripts because you can use them for either Unicode or non-Unicode tables.

Note: The eMessage installation directory is a subfolder within the Campaign directory.

Use the following table to see the scripts that you must use:

Table 20. Scripts for populating eMessage tables

Data source type	Script name
IBM DB2	ace_op_populate_systab_db2.sql
Microsoft SQL Server	ace_op_populate_systab_sqlsvr.sql
Oracle	ace_op_populate_systab_ora.sql

Configuring a system user to access IBM Marketing Software Hosted Services

eMessage components must be able to access the IBM Marketing Software Hosted Services without requiring manual entry of login credentials. To establish automatic login, define a system user in Marketing Platform that can provide the required access credentials.

To simplify user administration and troubleshooting, you can modify an existing system user to access hosted services and local system tables. You can configure a single system user to provide credentials for multiple systems. For example, modifying the configuration of the Campaign system user creates a single user that can automatically access IBM Marketing Software Hosted Services and the eMessage system tables in the Campaign schema.

The credentials required to access IBM Marketing Software Hosted Services are the user name and password that IBM provides for your hosted messaging account. The credentials that you use depend on whether you are connecting to the IBM US data center or the data center that IBM maintains in the United Kingdom. Consult with IBM to determine which data center you use.

For specific information about how to configure a system user to communicate with IBM Marketing Software Hosted Services, see the *IBM eMessage Startup and Administrator's Guide*.

For general information about how to create system users and data sources, see the *IBM Marketing Platform Administrator's Guide*.

Configuring partition access to IBM Marketing Software Hosted Services

IBM eMessage components in the partition must be allowed to automatically provide valid login credentials when attempting to communicate with IBM Marketing Software Hosted Services. To do this, you must add the IBM Marketing Software Hosted Services login credentials to a Marketing Platform user. This user becomes the eMessage system user.

You can add the platform data source containing the IBM Marketing Software Hosted Services credentials to the eMessage system user. This user can be the same system user that accesses the Campaign system tables in the partition.

The steps for configuring system users for a partition are the same as those followed during the initial eMessage installation, which created the first partition. For details about how to add IBM Marketing Software Hosted Services login credentials to a system user, see the *IBM eMessage Startup and Administrator's Guide*.

The credentials required to access IBM Marketing Software Hosted Services are the user name and password that IBM provided during the initial startup process.

Important: For each additional partition, you must request a separate user name and password from IBM.

Enabling eMessage in Campaign for the new partition

To allow users of the new eMessage partition to access the eMessage features in Campaign, enable eMessage in the Campaign partition by updating the eMessageInstalled configuration property for the corresponding Campaign partition.

For example, the eMessage mailing tab does not appear in the Campaign interface until you enable eMessage in the Campaign configuration.

You enable eMessage in the partition by updating the eMessageInstalled configuration property for the corresponding Campaign partition.

In the Marketing Platform configurations, navigate to **Campaign | partitions | partition[n] | server | internal** and set the eMessageInstalled property to **yes**.

Specifying the Recipient List Uploader location for eMessage

For each partition where you enable eMessage, specify the location of the Recipient List Uploader (RLU). The RLU uploads the output list tables data and associated metadata to the remote services hosted by IBM.

During initial installation, the IBM installer automatically adds the location of the RLU to the configuration for default partition (partition1). However, when you add new partitions to your environment, you must manually configure all new partitions to reference the correct location. Because there is only one RLU for each eMessage installation, all partitions access the same plugin file in the local file system of the machine that hosts the Campaign web application.

1. In the configuration for **partition1** in your Campaign installation, navigate to Campaign > partitions > partition1 > eMessage > eMessagePluginJarFile.

The value for this property is the full path to the plugin file (emessageplugin.jar) that operates as the RLU.

For example: C:\IBM\Unica\eMessage\plugin\emessageplugin.jar

2. Copy the value for the eMessagePluginJarFile property.
3. Navigate to the eMessagePluginJarFile for the new partition and enter the path you copied from **partition1**.

All partitions must use the same location for the RLU.

Restarting system components after configuring eMessage

After you change the eMessage and Campaign configurations, you must restart the Campaign web application server, the Response and Contact Tracker (RCT), and the Campaign listener.

1. Restart the web application server for Campaign.

For instructions, see the documentation for your web application server.

To verify that the server has started, log in to your IBM Marketing Software installation, access Campaign, and confirm that you can open an existing mailing.

2. Restart the Response and Contact Tracker (RCT).

To restart the RCT manually, run the rct script in the bin directory under your eMessage installation, as follows: rct start

If RCT is configured to run as a service, restart the RCT service. The first time that you restart the RCT as a service, you must manually restart the RCT afterward.

For more information, see “eMessage Response and Contact Tracker (RCT) script” on page 109.

3. Restart the Campaign listener:

- On Windows, run the `cmpServer.bat` file in the `bin` directory under your Campaign installation.
- On UNIX, run the following command as root: `./rc.unica_ac start`

Testing eMessage partition configuration and connections

Use the scripts that are provided by eMessage to verify the partition configuration and its connection to IBM Marketing Software Hosted Services. You must also confirm if you can access the mailing interface from the partition.

Important: Before you begin, if you have changed the Campaign or eMessage configurations, confirm that you have restarted the web application server that hosts Campaign and that you have restarted the Response and Contact Tracker.

For more information on how to test the partition, see the *IBM eMessage Startup and Administrator's Guide*.

Chapter 10. Marketing Platform utilities

This section provides an overview of the Marketing Platform utilities, including some details that apply to all of the utilities and which are not included in the individual utility descriptions.

Location of utilities

Marketing Platform utilities are located in the `tools/bin` directory under your Marketing Platform installation.

List and descriptions of utilities

The Marketing Platform provides the following utilities.

- “`alertConfigTool`” on page 97 - registers alerts and configurations for IBM Marketing Software products
- “`configTool`” on page 97 - imports, exports, and deletes configuration settings, including product registrations
- “`datafilteringScriptTool`” on page 101 - creates data filters
- “`encryptPasswords`” on page 102 - encrypts and stores passwords
- **`encryptTomcatDBPasswords`** - encrypt the database passwords that the Tomcat Application server uses internally
- “`partitionTool`” on page 103 - creates database entries for partitions
- “`populateDb`” on page 105 - populates the Marketing Platform database
- “`restoreAccess`” on page 106 - restores a user with the `platformAdminRole` role
- “`scheduler_console_client`” on page 108 - lists or starts IBM Marketing Software Scheduler jobs that are configured to listen for a trigger.

Prerequisites for running Marketing Platform utilities

The following are prerequisites for running all Marketing Platform utilities.

- Run all utilities from the directory where they are located (by default, the `tools/bin` directory under your Marketing Platform installation).
- On UNIX, the best practice is to run the utilities with the same user account that runs the application server on which Marketing Platform is deployed. If you run a utility with a different user account, adjust the permissions on the `platform.log` file to allow that user account to write to it. If you do not adjust permissions, the utility is not able to write to the log file and you might see some error messages, although the tool should still function correctly.

Authentication of utilities

Utilities such as `configTool` and other IBM Marketing Software back end utilities are designed to be used by system administrators and require physical access to the host servers for them to be invoked. For this reason, authentication for these utilities has been designed to be independent of the UI authentication mechanism. Access to these utilities is available to users with Marketing Platform administrator privileges. Access to these utilities is expected to be locally defined in Marketing Platform and authenticated against the same.

Troubleshooting connection issues

All of the Marketing Platform utilities except `encryptPasswords` interact with the Marketing Platform system tables. To connect to the system table database, these utilities use the following connection information, which is set by the installer using information provided when the Marketing Platform was installed. This information is stored in the `jdbc.properties` file, located in the `tools/bin` directory under your Marketing Platform installation.

- JDBC driver name
- JDBC connection URL (which includes the host, port, and database name)
- Data source login
- Data source password (encrypted)

In addition, these utilities rely on the `JAVA_HOME` environment variable, set either in the `setenv` script located in the `tools/bin` directory of your Marketing Platform installation, or on the command line. The Marketing Platform installer should have set this variable automatically in the `setenv` script, but it is a good practice to verify that the `JAVA_HOME` variable is set if you have a problem running a utility. The JDK must be the Sun version (not, for example, the JRockit JDK available with WebLogic).

Special characters

Characters that are designated as reserved characters in the operating system must be escaped. Consult your operating system documentation for a list of reserved characters and how to escape them.

Standard options in Marketing Platform utilities

The following options are available in all Marketing Platform utilities.

`-l logLevel`

Set the level of log information displayed in the console. Options are `high`, `medium`, and `low`. The default is `low`.

`-L`

Set the locale for console messages. The default locale is `en_US`. The available option values are determined by the languages into which the Marketing Platform has been translated. Specify the locale using the ICU locale ID according to ISO 639-1 and ISO 3166.

`-h`

Display a brief usage message in the console.

`-m`

Display the manual page for this utility in the console.

`-v`

Display more execution details in the console.

alertConfigTool

Notification types are specific to the various IBM Marketing Software products. Use the alertConfigTool utility to register the notification types when the installer has not done this automatically during installation or upgrade.

Syntax

```
alertConfigTool -i -f importFile
```

Commands

-i -f *importFile*

Import alert and notification types from a specified XML file.

Example

- Import alert and notification types from a file named Platform_alerts_configuration.xml located in the tools\bin directory under the Marketing Platform installation.

```
alertConfigTool -i -f Platform_alerts_configuration.xml
```

configTool

The properties and values on the Configuration page are stored in the Marketing Platform system tables. You can use the configTool utility to import and export configuration settings to and from the system tables.

When to use configTool

You might want to use configTool for the following reasons.

- To import partition and data source templates that are supplied with Campaign, which you can then modify and duplicate by using the Configuration page.
- To register (import configuration properties for) IBM Marketing Software products, if the product installer is unable to add the properties to the database automatically.
- To export an XML version of configuration settings for backup or to import into a different installation of IBM Marketing Software.
- To delete categories that do not have the **Delete Category** link. You do this by using configTool to export your configuration, then manually deleting the XML that creates the category, and by using configTool to import the edited XML.

Important: This utility modifies the usm_configuration and usm_configuration_values tables in the Marketing Platform system table database, which contains the configuration properties and their values. For best results, either create backup copies of these tables, or export your existing configurations by using configTool and back up the resulting file so you have a way to restore your configuration if you make an error when you use configTool to import.

Syntax

```
configTool -d -p "elementPath" [-o]
```

```
configTool -i -p "parent ElementPath" -f importFile [-o]
```

```
configTool -x -p "elementPath" -f exportFile
```

```
configTool -vp -p "elementPath" -f importFile [-d]
```

```
configTool -r productName -f registrationFile [-o] configTool -u  
productName
```

Commands

-d -p "elementPath" [o]

Delete configuration properties and their settings, specifying a path in the configuration property hierarchy.

The element path must use the internal names of categories and properties. You can obtain them by going to the Configuration page, selecting the wanted category or property, and looking at the path that is displayed in parentheses in the right pane. Delimit a path in the configuration property hierarchy by using the | character, and surround the path with double quotation marks.

Note the following.

- Only categories and properties within an application can be deleted by using this command, not whole applications. Use the -u command to unregister a whole application.
- To delete categories that do not have the **Delete Category** link on the Configuration page, use the -o option.

When you use -d with the -vp command, the configTool deletes any child nodes in the path you specify if those nodes are not included in the XML file you specify.

-i -p "parentElementPath" -f importFile [o]

Import configuration properties and their settings from a specified XML file.

To import, you specify a path to the parent element under which you want to import your categories. The configTool utility imports properties under the category you specify in the path.

You can add categories at any level below the top level, but you cannot add a category at same level as the top category.

The parent element path must use the internal names of categories and properties. You can obtain them by going to the Configuration page, selecting the required category or property, and looking at the path that is displayed in parentheses in the right pane. Delimit a path in the configuration property hierarchy by using the | character, and surround the path with double quotation marks.

You can specify an import file location relative to the tools/bin directory or you can specify a full directory path. If you specify a relative path or no path, configTool first looks for the file relative to the tools/bin directory.

By default, this command does not overwrite an existing category, but you can use the -o option to force an overwrite.

-x -p "elementPath" -f exportFile

Export configuration properties and their settings to an XML file with a specified name.

You can export all configuration properties or limit the export to a specific category by specifying a path in the configuration property hierarchy.

The element path must use the internal names of categories and properties, which you can obtain by going to the Configuration page, selecting the wanted category or property, and looking at the path that is displayed in parentheses in the right pane. Delimit a path in the configuration property hierarchy by using the | character, and surround the path with double quotation marks.

You can specify an export file location relative to the current directory or you can specify a full directory path. If the file specification does not contain a separator (/ on UNIX, / or \ on Windows), configTool writes the file to the tools/bin directory under your Marketing Platform installation. If you do not provide the xml extension, configTool adds it.

-vp -p "*elementPath*" -f *importFile* [-d]

This command is used mainly in manual upgrades, to import configuration properties. If you applied a fix pack that contains a new configuration property, and you then upgrade, importing a configuration file as part of a manual upgrade process can override values that were set when the fix pack was applied. The -vp command ensures that the import does not override previously set configuration values.

Important: After you use the configTool utility with the -vp option, you must restart the web application server on which Marketing Platform is deployed so the changes are applied.

When you use -d with the -vp command, the configTool deletes any child nodes in the path you specify if those nodes are not included in the XML file you specify.

-r *productName* -f *registrationFile*

Register the application. The registration file location can be relative to the tools/bin directory or can be a full path. By default, this command does not overwrite an existing configuration, but you can use the -o option to force an overwrite. The *productName* parameter must be one of those names that are listed above.

Note the following.

- When you use the -r command, the registration file must have <application> as the first tag in the XML.
Other files can be provided with your product that you can use to insert configuration properties into the Marketing Platform database. For these files, use the -i command. Only the file that has the <application> tag as the first tag can be used with the -r command.
- The registration file for the Marketing Platform is named Manager_config.xml, and the first tag is <Suite>. To register this file on a new installation, use the populateDb utility, or rerun the Marketing Platform installer as described in the *IBM Marketing Platform Installation Guide*.

- After the initial installation, to re-register products other than the Marketing Platform, use configTool with the -r command and -o to overwrite the existing properties.

The configTool utility uses product names as parameters with the commands that register and unregister products. With the 8.5.0 release of IBM Marketing Software, many product names changed. However, the names that are recognized by configTool did not change. The valid product names for use with configTool are listed below, along with the current names of the products.

Table 21. Product names for configTool registration and unregistration

Product name	Name used in configTool
Marketing Platform	Manager
Campaign	Campaign
Distributed Marketing	Collaborate
eMessage	emessage
Interact	interact
Contact Optimization	Optimize
Marketing Operations	Plan
Opportunity Detect	Detect
IBM SPSS Modeler Advantage Enterprise Marketing Management Edition	SPSS
Digital Analytics	Coremetrics

-u *productName*

Unregister an application that is specified by *productName*. You do not have to include a path to the product category; the product name is sufficient, and it is required. The process removes all properties and configuration settings for the product.

Options

-o

When used with -i or -r, it overwrites an existing category or product registration (node).

When used with -d, you can delete a category (node) that does not have the **Delete Category** link on the Configuration page.

Examples

- Import configuration settings from a file named Product_config.xml in the conf directory under the Marketing Platform installation.

```
configTool -i -p "Affinium" -f Product_config.xml
```
- Import one of the supplied Campaign data source templates into the default Campaign partition, partition1. The example assumes that you placed the Oracle data source template, OracleTemplate.xml, in the tools/bin directory under the Marketing Platform installation.

```
configTool -i -p "Affinium|Campaign|partitions|partition1|dataSources" -f OracleTemplate.xml
```


- Export all configuration settings to a file named `myConfig.xml` in the `D:\backups` directory.
`configTool -x -f D:\backups\myConfig.xml`
- Export an existing Campaign partition (complete with data source entries), save it to a file named `partitionTemplate.xml`, and store it in the default `tools/bin` directory under the Marketing Platform installation.
`configTool -x -p "Affinium|Campaign|partitions|partition1" -f partitionTemplate.xml`
- Manually register an application named `productName`, by using a file named `app_config.xml` in the default `tools/bin` directory under the Marketing Platform installation, and force it to overwrite an existing registration of this application.
`configTool -r product Name -f app_config.xml -o`
- Unregister an application named `productName`.
`configTool -u productName`
- Run the following command to enable `encodeCSV` feature:
`configTool -vp -p "Affinium|Plan|umoConfiguration" -f MarketingOperations_Home\conf\Plan_encodeProperty_11.1.xml`

datafilteringScriptTool

The `datafilteringScriptTool` utility reads an XML file to populate the data filtering tables in the Marketing Platform system table database.

Depending on how you write the XML, you can use this utility in two ways.

- Using one set of XML elements, you can auto-generate data filters based on unique combinations of field values (one data filter for each unique combination).
- Using a slightly different set of XML elements, you can specify each data filter that the utility creates.

See *IBM Marketing Platform the Administrator's Guide* for information about creating the XML.

When to use datafilteringScriptTool

You must use `datafilteringScriptTool` when you create new data filters.

Prerequisites

The Marketing Platform must be deployed and running.

Using datafilteringScriptTool with SSL

When the Marketing Platform is deployed using one-way SSL you must modify the `datafilteringScriptTool` script to add the SSL options that perform handshaking. To modify the script, you must have the following information.

- Truststore file name and path
- Truststore password

In a text editor, open the `datafilteringScriptTool` script (`.bat` or `.sh`) and find the lines that look like this (examples are Windows version).

```
:callexec
```

```
"%JAVA_HOME%\bin\java" -DUNICA_PLATFORM_HOME="%UNICA_PLATFORM_HOME%"
```

```
com.unica.management.client.datafiltering.tool.DataFilteringScriptTool %*
```

Edit these lines to look like this (new text is in **bold**). Substitute your truststore path and file name and truststore password for myTrustStore.jks and myPassword.

```
:call exec
```

```
SET SSL_OPTIONS=-Djavax.net.ssl.keyStoreType="JKS"
```

```
-Djavax.net.ssl.trustStore="C:\security\myTrustStore.jks"
```

```
-Djavax.net.ssl.trustStorePassword=myPassword
```

```
"%JAVA_HOME%\bin\java" -DUNICA_PLATFORM_HOME="%UNICA_PLATFORM_HOME%"  
%SSL_OPTIONS%
```

```
com.unica.management.client.datafiltering.tool.DataFilteringScriptTool %*
```

Syntax

```
datafilteringScriptTool -r pathfile
```

Commands

-r *path_file*

Import data filter specifications from a specified XML file. If the file is not located in the tools/bin directory under your installation, provide a path and enclose the *path_file* parameter in double quotation marks.

Example

- Use a file named collaborateDataFilters.xml, located in the C:\unica\xml directory, to populate the data filter system tables.

```
datafilteringScriptTool -r "C:\unica\xml\collaborateDataFilters.xml"
```

encryptPasswords

The encryptPasswords utility is used to encrypt and store either of two passwords that Marketing Platform uses internally.

The two passwords that the utility can encrypt are as follows.

- The password that the Marketing Platform uses to access its system tables. The utility replaces an existing encrypted password (stored in the jdbc.properties file, located in the tools\bin directory under your Marketing Platform installation) with a new one.
- The keystore password used by the Marketing Platform when it is configured to use SSL with a certificate other than the default one supplied with the Marketing Platform or the web application server. The certificate can be either a self-signed certificate or a certificate from a certificate authority.

When to use encryptPasswords

Use encryptPasswords as for the following reasons.

- When you change the password of the account used to access your Marketing Platform system table database.
- When you have created a self-signed certificate or have obtained one from a certificate authority.

Prerequisites

- Before running `encryptPasswords` to encrypt and store a new database password, make a backup copy of the `jdbc.properties` file, located in the `tools/bin` directory under your Marketing Platform installation.
- Before running `encryptPasswords` to encrypt and store the keystore password, you must have created or obtained a digital certificate and know the keystore password.

Syntax

```
encryptPasswords -d databasePassword
```

```
encryptPasswords -k keystorePassword
```

Commands

-d *databasePassword*

Encrypt the database password.

-k *keystorePassword*

Encrypt the keystore password and store it in a file named `pfile`.

Examples

- When the Marketing Platform was installed, the login for the system table database account was set to `myLogin`. Now, some time after installation, you have changed the password for this account to `newPassword`. Run `encryptPasswords` as follows to encrypt and store the database password.

```
encryptPasswords -d newPassword
```
- You are configuring an IBM Marketing Software application to use SSL and have created or obtained a digital certificate. Run `encryptPasswords` as follows to encrypt and store the keystore password.

```
encryptPasswords -k myPassword
```

partitionTool

Partitions are associated with Campaign policies and roles. These policies and roles and their partition associations are stored in the Marketing Platform system tables. The `partitionTool` utility seeds the Marketing Platform system tables with basic policy and role information for partitions.

When to use partitionTool

For each partition you create, you must use `partitionTool` to seed the Marketing Platform system tables with basic policy and role information.

See the installation guide appropriate for your version of Campaign for detailed instructions on setting up multiple partitions in Campaign.

Special characters and spaces

Any partition description or user, group, or partition name that contains spaces must be enclosed in double quotation marks.

Syntax

```
partitionTool -c -s sourcePartition -n newPartitionName [-u  
admin_user_name] [-d partitionDescription] [-g groupName]
```

Commands

The following commands are available in the `partitionTool` utility.

-c

Replicates (clones) the policies and roles for an existing partition specified using the `-s` option, and uses the name specified using the `-n` option. Both of these options are required with `c`. This command does the following.

- Creates a new IBM Marketing Software user with the Admin role in both the Administrative Roles policy and the global policy in Campaign. The partition name you specify is automatically set as this user's password.
- Creates a new Marketing Platform group and makes the new Admin user a member of that group.
- Creates a new partition object.
- Replicates all the policies associated with the source partition and associates them with the new partition.
- For each replicated policy, replicates all roles associated with the policy.
- For each replicated role, maps all functions in the same way that they were mapped in the source role.
- Assigns the new Marketing Platform group to the last system-defined Admin role created during role replication. If you are cloning the default partition, `partition1`, this role is the default Administrative Role (Admin).

Options

-d *partitionDescription*

Optional, used with `-c` only. Specifies a description that appears in the output from the `-list` command. Must be 256 characters or less. Enclose in double quotation marks if the description contains spaces.

-g *groupName*

Optional, used with `-c` only. Specifies the name of the Marketing Platform Admin group that the utility creates. The name must be unique within this instance of Marketing Platform

If not defined, the name defaults to `partition_nameAdminGroup`.

-n *partitionName*

Optional with `-list`, required with `-c`. Must be 32 characters or less.

When used with `-list`, specifies the partition whose information is listed.

When used with `-c`, specifies the name of the new partition, and the partition name you specify is used as the password for the Admin user. The partition name must match the name you gave the partition in when you configured it (using the partition template on the Configuration page).

-s *sourcePartition*

Required, used with `-c` only. The name of the source partition to be replicated.

-u *adminUserName*

Optional, used with `-c` only. Specifies the user name of the Admin user for the replicated partition. The name must be unique within this instance of Marketing Platform.

If not defined, the name defaults to *partitionNameAdminUser*.

The partition name is automatically set as this user's password.

Examples

- Create a partition with the following characteristics.

- Cloned from partition1
- Partition name is myPartition
- Uses the default user name (myPartitionAdminUser) and password (myPartition)
- Uses the default group name (myPartitionAdminGroup)
- Description is "ClonedFromPartition1"

```
partitionTool -c -s partition1 -n myPartition -d "ClonedFromPartition1"
```

- Create a partition with the following characteristics.

- Cloned from partition1
- Partition name is partition2
- Specifies user name of customerA with the automatically assigned password of partition2
- Specifies group name of customerAGroup
- Description is "PartitionForCustomerAGroup"

```
partitionTool -c -s partition1 -n partition2 -u customerA -g customerAGroup -d "PartitionForCustomerAGroup"
```

populateDb

The `populateDb` utility inserts default (seed) data in the Marketing Platform system tables.

The IBM Marketing Software installer can populate the Marketing Platform system tables with default data for Marketing Platform and for Campaign. However, if your company policy does not permit the installer to change the database, or if the installer is unable to connect with the Marketing Platform system tables, you must insert default data in the Marketing Platform system tables using this utility.

For Campaign, this data includes security roles and permissions for the default partition. For Marketing Platform, this data includes default users and groups, and security roles and permissions for the default partition.

Syntax

```
populateDb -n productName
```

Commands

```
-n productName
```

Insert default data into the Marketing Platform system tables. Valid product names are Manager (for Marketing Platform) and Campaign (for Campaign).

Examples

- Insert Marketing Platform default data manually.

```
populateDb -n Manager
```
- Insert Campaign default data manually.

```
populateDb -n Campaign
```

restoreAccess

The `restoreAccess` utility allows you to restore access to Marketing Platform if all users with `PlatformAdminRole` privileges have been inadvertently locked out or if all ability to log in to the Marketing Platform has been lost.

When to use restoreAccess

You might want to use `restoreAccess` under the two circumstances described in this section.

PlatformAdminRole users disabled

It is possible that all users with `PlatformAdminRole` privileges in Marketing Platform might become disabled in the system. Here is an example of how the `platform_admin` user account might become disabled. Suppose you have only one user with `PlatformAdminRole` privileges (the `platform_admin` user). Assume the Maximum failed login attempts allowed property in the **General | Password settings** category on the Configuration page is set to 3. Then suppose someone who is attempting to log in as `platform_admin` enters an incorrect password three times in a row. These failed login attempts cause the `platform_admin` account to become disabled in the system.

In that case, you can use `restoreAccess` to add a user with `PlatformAdminRole` privileges to the Marketing Platform system tables without accessing the web interface.

When you run `restoreAccess` in this way, the utility creates a user with the login name and password you specify, and with `PlatformAdminRole` privileges.

If the user login name you specify exists in Marketing Platform as an internal user, that user's password is changed.

Only a user with the login name of PlatformAdmin and with PlatformAdminRole privileges can universally administer all dashboards. So if the platform_admin user is disabled and you create a user with restoreAccess, you should create a user with a login of platform_admin.

Improper configuration of NTLMv2 authentication

If you implement NTLMv2 authentication with improper configuration and can no longer log in, use restoreAccess to restore the ability to log in.

When you run restoreAccess in this way, the utility changes the value of the Platform | Security | Login method property to Marketing Platform. This change allows you to log in with any user account that existed before you were locked out. You can optionally specify a new login name and password as well. You must restart the web application server on which Marketing Platform is deployed if you use the restoreAccess utility in this way.

Password considerations

Note the following about passwords when you use restoreAccess.

- The restoreAccess utility does not support blank passwords, and does not enforce password rules.
- If you specify a user name that is in use, the utility resets the password for that user.

Syntax

```
restoreAccess -u loginName -p password
```

```
restoreAccess -r
```

Commands

-r

When used without the *-u loginName* option, reset the value of the Platform | Security | Login method property to Marketing Platform. Requires restart of the web application server to take effect.

When used with the *-u loginName* option, create a PlatformAdminRole user.

Options

-u *loginName*

Create a user with PlatformAdminRole privileges with the specified login name. Must be used with the *-p* option.

-p *password*

Specify the password for the user being created. Required with *-u*.

Examples

- Create a user with PlatformAdminRole privileges. The login name is tempUser and the password is tempPassword.

```
restoreAccess -u tempUser -p tempPassword
```

- Change the value of the login method to IBM Marketing Platform and create a user with PlatformAdminRole privileges. The login name is tempUser and the password is tempPassword.

```
restoreAccess -r -u tempUser -p tempPassword
```

scheduler_console_client

Jobs configured in the IBM Marketing Software Scheduler can be listed and kicked off by this utility, if they are set up to listen for a trigger.

What to do if SSL is enabled

When the Marketing Platform web application is configured to use SSL, the JVM used by the scheduler_console_client utility must use the same SSL certificate that is used by the web application server on which the Marketing Platform is deployed.

Take the following steps to import the SSL certificate

- Determine the location of the JRE used by the scheduler_console_client.
 - If JAVA_HOME is set as a system environment variable, the JRE it points to is the one used by the scheduler_console_client utility.
 - If JAVA_HOME is not set as a system environment variable, the scheduler_console_client utility uses the JRE set either in the setenv script located in the tools/bin directory of your Marketing Platform installation, or on the command line.
- Import the SSL certificate used by the web application server on which the Marketing Platform is deployed to the JRE used by scheduler_console_client. The Sun JDK includes a program called keytool that you can use to import the certificate. Consult the Java documentation for complete details on using this program, or access the help by entering -help when you run the program.
- Open the tools/bin/schedulerconsoleclient file in a text editor and add the following properties. These differ depending on the web application server on which Marketing Platform is deployed.
 - For WebSphere, add these properties to the file.
 - Djavax.net.ssl.keyStoreType=JKS
 - Djavax.net.ssl.keyStore="*Path to your key store JKS file*"
 - Djavax.net.ssl.keyStorePassword="*Your key store password*"
 - Djavax.net.ssl.trustStore="*Path to your trust store JKS file*"
 - Djavax.net.ssl.trustStorePassword="*Your trust store password*"
 - DisUseIBMSSLSocketFactory=false
 - For WebLogic, add these properties to the file.
 - Djavax.net.ssl.keyStoreType="JKS"
 - Djavax.net.ssl.trustStore="*Path to your trust store JKS file*"
 - Djavax.net.ssl.trustStorePassword="*Your trust store password*"

If the certificates do not match, the Marketing Platform log file contains an error such as the following.

```
Caused by: sun.security.provider.certpath.SunCertPathBuilderException:  
unable to find valid certification path to requested target
```


Prerequisites

The Marketing Platform must be installed, deployed, and running.

Syntax

```
scheduler_console_client -v -t trigger_name user_name
```

```
scheduler_console_client -s -t trigger_name user_name
```

Commands

-v

List the scheduler jobs configured to listen for the specified trigger.

Must be used with the -t option.

-s

Send the specified trigger.

Must be used with the -t option.

Options

-t *trigger_name*

The name of the trigger, as configured in the scheduler.

Example

- List jobs configured to listen for a trigger named trigger1.

```
scheduler_console_client -v -t trigger1 myLogin
```
- Execute jobs configured to listen for a trigger named trigger1.

```
scheduler_console_client -s -t trigger1 myLogin
```

eMessage Response and Contact Tracker (RCT) script

Use this script to run and check the status of the Response and Contact Tracker (RCT).

This script is located in the bin directory under your eMessage installation. The eMessage directory is a sub-directory in the Campaign directory.

In UNIX or Linux environments run the script as `rct.sh`.

In Windows, run the script from the command line as `rct.bat`.

Syntax

```
rct [ start | stop | check ]
```

Commands

start

Start the RCT

stop

Stop the RCT

Options

check

Check the status of the connection between the RCT and IBM Marketing Software Hosted Services.

Examples

- To start the RCT on Windows.

```
rct.bat start
```

- To stop the RCT on Windows.

```
rct.bat stop
```

- In a Linux environment, to determine if the RCT is connected to IBM Marketing Software Hosted Services.

```
rct.sh check
```

Depending on the status of your system, the output of this command might look like this:

```
C:\<IMS_HOME>\Campaign\eMessage\bin>rct check
Testing config and connectivity for partition partition1
Succeeded | Partition: partition1 - Hosted Services Account ID:
asm_admin
```

The MKService_rct script

The MKService_rct script adds or removes the Response and Contact Tracker (RCT) as a service. Adding the RCT as a service restarts the RCT every time you restart the computer where you have installed the RCT. Removing the RCT as a service prevents the RCT from restarting automatically.

This script is located in the bin directory under your eMessage installation.

In UNIX or Linux environments run MKService_rct.sh. with a user that has root permissions or permissions to create daemon processes.

In Windows, run the script from the command line as MKService_rct.bat.

Syntax

```
MKService_rct -install
```

```
MKService_rct -remove
```

Commands

-install

Add the RCT as a service

-remove

Remove the RCT service

Examples

- To add the RCT as a Windows service.
MKService_rct.bat -install
- To remove the RCT service on UNIX or Linux.
MKService_rct.sh -remove

Chapter 11. Uninstalling Campaign

Run the Campaign uninstaller to uninstall Campaign. When you run the Campaign uninstaller, the files that were created during the installation process are removed. For example, files such as configuration files, installer registry information, and user data are removed from the computer.

When you install IBM Marketing Software products, an uninstaller is included in the `Uninstall_Product` directory, where *Product* is the name of your IBM product. On Windows, an entry is also added to the **Add or Remove Programs** list in the Control Panel.

If you manually remove the files in your installation directory instead of running the uninstaller, the result might be an incomplete installation if you later reinstall an IBM product in the same location. After uninstalling a product, its database is not removed. The uninstaller only removes default files that are created during installation. Any file that is created or generated after installation is not removed.

Note: On UNIX, the same user account that installed Campaign must run the uninstaller.

1. If you have deployed the Campaign web application, undeploy the web application from WebSphere or WebLogic.
2. Stop the Campaign listener.
3. Shut down WebSphere or WebLogic.
4. Stop the processes that are related to Campaign.
5. If the `ddl` directory exists in the product installation directory, run the scripts that are provided in the `ddl` directory to drop tables from the system table database.
6. Complete one of the following steps to uninstall Campaign:
 - Click the Campaign uninstaller that exists in the `Uninstall_Product` directory. The uninstaller runs in the mode in which you installed Campaign.
 - In a command-line window, navigate to the directory where the uninstaller exists, and run the following command to uninstall Campaign by using the console mode:

`Uninstall_Product -i console`

- In a command-line window, navigate to the directory where the uninstaller exists, and run the following command to uninstall Campaign by using the silent mode:

`Uninstall_Product -i silent`

When you uninstall Campaign by using the silent mode, the uninstallation process does not present any dialogs for user interaction.

Note: If you do not specify an option for uninstalling Campaign, the Campaign uninstaller runs in the mode in which Campaign is installed.

Chapter 12. IBM Campaign web application clustering

You can configure the IBM Campaign web application (J2EE application) in a cluster by completing additional tasks during the IBM Campaign installation.

You can integrate a Campaign clustered web application environment with Interact, Distributed Marketing, Marketing Operations and eMessage.

Overview

The Campaign web application server components support load balancing when the web application is deployed in a clustered environment.

In Campaign, certain objects, such as offers, are cached in the web application server to improve response times in the user interface. In a load balanced environment, offers created on one web application server are available to users whose sessions connect through another server if you have configured `campaign_ehcache.xml`.

The objects that are cached include: offers, offer templates and attributes, campaigns, sessions, folders, attributes, initiatives, and segments.

A clustered configuration can be used to minimize disruptions in application availability in the event of a web application server failure. Because the Campaign system tables are stored in an external database, it is not necessary to replicate data to the standby server. If the primary web application server fails, the Campaign web application must be started on the standby server. When started, the Campaign web application retrieves all current data from the Campaign system table database.

Main steps

To install IBM Campaign in a clustered web application environment, follow the instructions from chapters 2 through 8 and supplement those instructions with information from this chapter.

When you install IBM Campaign in a cluster, there are many ways to configure the installation. The following steps describe the basic process:

1. Run the installers on one system, typically, the administration server or whatever the equivalent is for your application server type.
2. Create and share a file directory to hold the upload files for all of the IBM Campaign installations.
3. Create the EAR or WAR file and deploy it on the cluster and web server.
4. Configure each system to share the IBM Marketing Platform system tables and the IBM Campaign system tables.
5. Configure each system to use the shared file directory.
6. Determine which server in the cluster will send notifications. Then, suppress the notification process on all of the other servers.
7. Configure `campaign_ehcache.xml` for distributed caching of templates and offer folders, as explained in the *IBM Campaign Installation Guide*. To specify how often the cache is refreshed, you can adjust the Campaign cache configuration

settings, such as `offerTemplateDataTTLSeconds`. However, this adjustment can adversely affect performance because smaller values result in more frequent cache updates. Configuration settings are described in the *IBM Campaign Administrator's Guide*.

Note: If you have problems after deploying IBM Campaign in a WebSphere cluster, copy the file `commons-lang.jar` to the directory `/data/webservers/IBM/WAS85ND/lib/ext`.

WebSphere clustering guidelines

If you are installing IBM Campaign on a cluster in WebSphere, complete these extra steps in addition to the steps for installing IBM Campaign on WebSphere.

Preparing the data sources

In addition to any other instructions provided in this guide, complete the following tasks for your data sources:

- The IBM Campaign database must be on a server that is accessible to all servers in the cluster, but it does not need to be on a server in the cluster.
- When you configure the JDBC provider, specify the cluster as the scope.

Installing the products

When you run the installer, be sure that you install Marketing Platform and IBM Campaign only once, on a server that is designated as the administration server for the cluster. The administration server is accessible to all servers in the IBM Campaign cluster.

You do not need to install the software on each cluster member. Instead, you install it on the administration server, create the EAR or WAR file, and then deploy the EAR or WAR file on each cluster member.

Note: The IBM WebSphere application server that is bundled with Campaign does not support clustering. In a clustered environment, use IBM WebSphere Application Server Network Deployment.

Pre-deployment steps

Before you deploy IBM Campaign, complete the following task:

- Share the top-level directory in which IBM Campaign is installed. For example, if you installed Campaign in `C:\CampaignCluster\IBM_IMS\Campaign`, you would share the entire `CampaignCluster` directory.

Deployment steps

In addition to the instructions in the deployment chapter (Chapter 6, “Deploying the Campaign web application,” on page 65), complete the following tasks:

1. Map modules to servers. When you set the options in the **Select Installation Options** wizard for WebSphere, select the cluster and the web server when mapping modules to servers.
2. Extra instructions for the Generic JVM property: Configure the Generic JVM property on each server in the cluster.

The path that you specify in the <CAMPAIGN_HOME> and other properties must point to the shared installation directory.

Set the following extra parameters for clusters:

- Set the following parameter to ensure that eMessage ETL is not triggered by all clustered nodes:

-Dcampaign.emsgetl.disabled=true

For the one node that should trigger ETL eMessage data for all other nodes, set the parameter to **false**.

- Set the following parameter to ensure that Interact ETL is not triggered by all clustered nodes:

-Dcampaign.interactetl.disabled=true

For the one node that should trigger ETL Interact data for all other nodes, set the parameter to **false**.

Post-deployment steps

If you are using a plug-in for load balancing, complete the following configuration steps:

- For IBM Campaign to work effectively in a clustered environment, users must remain on a single node for their entire session. This option for session management and load balancing is known as session affinity. The documentation for your application server provides details about configuring your installation to use session affinity.

Note: When a node fails on a system that uses this configuration option, every user session on that node also fails. The load balancer must not switch users to another available node because user authentication applies only to a single node in IBM Campaign. Users are prompted to log in again. In some cases, an unexpected error and corresponding data loss can occur.

- Log in to IBM Campaign. Select **Settings > Configuration**, and configure the following URL parameters so that all references to the IBM Campaign server use the proxy host and port:

Campaign | navigation | serverURL

Jboss clustering guidelines

For Installing IBM Campaign on Jboss cluster refer (Chapter 12, “IBM Campaign web application clustering,” on page 115).

Preparing the data sources

You will need to add/create JNDI in the same JBOSS domain which you are planning to use in cluster. Please follow below steps to create JNDI.

- Add required database modules in JBOSS. For example, oracle database -

```
jboss-cli.bat --connect --controller=JBOSS_DOMAIN_MASTER:MASTER_PORT --  
commands="module add --name= oracle.jdbc --resources=<Database_Driver_Path\ojdbc8.jar> --  
dependencies=javax.api,javax.transaction.api"
```
- JNDI's can be added by modifying domain.xml of both master and slave.

Deployment steps

In addition to the instructions in the deployment chapter (“Deploying Campaign on JBoss Application Server” on page 70), complete the following tasks:

1. Launch JBOSS management console.
2. Navigate to "Deployments" tab. Select the ServerGroup.
3. Click on Add button under Deployments section
4. A pop up will be shown, browse the application WAR/EAR and complete the deployment process.

Tomcat Clustering Guidelines

For Installing IBM Campaign on Tomcat cluster refer (Chapter 12, "IBM Campaign web application clustering," on page 115).

Note: EAR file is not supported for Tomcat cluster.

For Deployment Steps, you can refer ("Deploying Campaign on Tomcat Application Server" on page 70)

Post deployment, JNDI name should be prefixed with "java:/comp/env/" for example: campaignPartition1DS should be mentioned as "java:/comp/env/campaignPartition1DS" while creating Campaign UA_SYSTEM_TABLE Data Source.

Configure a load Balancer like nginx for redirecting request to all the nodes of Tomcat cluster.

Configure ehcache

The ehcache is an open source Java distributed cache for caching, Java EE, and light-weight containers. Every node in the cluster can use the same campaign_ehcache.xml file or you can set a different campaign_ehcache.xml file for each node. For installations in a cluster, edit the campaign_ehcache.xml file so you do not have to restart the computer if you change any templates or offer folders.

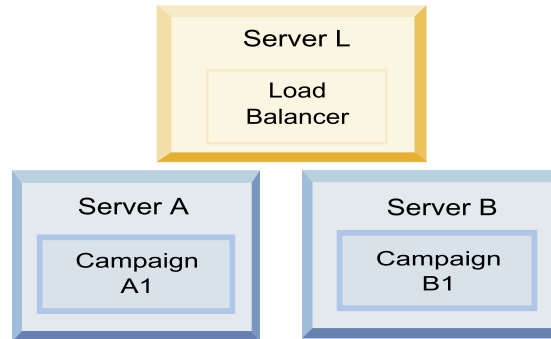
Important:

- If your installation is upgraded from a previous version, the campaign_ehcache.xml file is overwritten. Add and edit the following sections in the file.
- You must also configure caching for user sessions (userSessionCache).

Use one of the following procedures to configure the ehcache file.

Configuring ehcache with Remote Method Invocation (RMI)

IBM Campaign systems that have the following topography typically use RMI:



Go to the `<IBM_IMS_HOME>\<CAMPAIGN_HOME>\conf` directory and open the `campaign_ehcache.xml` file in a text editor. Then, make the following edits:

- Uncomment the following section of the file.

You must customize the settings for machineA and machineB to reflect your environment. Provide a pipe-separated list of all servers in the cluster and use fully qualified host names.

```
<!--<cacheManagerPeerProviderFactory
class="net.sf.ehcache.distribution.RMICacheManagerPeerProviderFactory"
properties="peerDiscovery=manual, rmiUrls=//
<machineA>:40000/campaignApplicationCache|//
<machineB>:40000/campaignApplicationCache|//
<machineA>:40000/userSessionCache|//
<machineB>:40000/userSessionCache"/>
-->
```

- Uncomment the following section of the file.

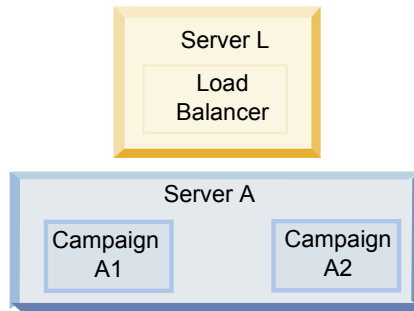
```
<!--
<cacheEventListenerFactory
class="net.sf.ehcache.distribution.RMICacheReplicatorFactory"
properties="replicateAsynchronously=true, replicatePuts=true,
replicateUpdates=true, replicateUpdatesViaCopy=true,
replicateRemovals=true"/>
<cacheEventListenerFactory
class="com.unicacorp.Campaign.cache.CampaignCacheEventListenerFactory" />
-->
```

- Uncomment the following section of the file.

```
<!--
<cacheEventListenerFactory
class="net.sf.ehcache.distribution.RMICacheReplicatorFactory"
properties="replicateAsynchronously=true, replicatePuts=true,
replicateUpdates=true, replicateUpdatesViaCopy=true,
replicateRemovals=true"/>
-->
```

Configuring ehcache with multicast

IBM Campaign systems that have the following topography typically use multicast:



Go to the `<IBM_IMS_HOME>\<CAMPAIGN_HOME>\conf` directory and open the `campaign_ehcache.xml` file in a text editor. Then, make the following edits:

- Uncomment the following section of the file.

You must customize the settings for `multicastGroupAddress` and `multicastGroupPort` to reflect your environment.

```
<!--<cacheManagerPeerProviderFactory
class="net.sf.ehcache.distribution.RMICacheManagerPeerProviderFactory"
properties="peerDiscovery=automatic, multicastGroupAddress=230.0.0.1,
multicastGroupPort=4446, timeToLive=32"/>

<cacheManagerPeerListenerFactory
class="net.sf.ehcache.distribution.RMICacheManagerPeerListenerFactory"/>
-->
```

- Uncomment the following section of the file.

```
<!--
<cacheEventListenerFactory
class="net.sf.ehcache.distribution.RMICacheReplicatorFactory"
properties="replicateAsynchronously=true, replicatePuts=true,
replicateUpdates=true, replicateUpdatesViaCopy=true,
replicateRemovals=true"/>
<cacheEventListenerFactory
class="com.unicacorp.Campaign.cache.CampaignCacheEventListenerFactory" />
-->
```

Chapter 13. Installing a clustered listener environment

Follow these instructions if you want to install a clustered listener configuration. A listener cluster is a set of two or more listeners that act as a unit to provide load balancing and high availability. IBM Campaign listener clusters are Active-Active; meaning each node serves requests using the load-balanced approach. Each Campaign listener provides an interface between front-end clients (such as the Campaign web application) and back-end analytical server processes.

1. Confirm that you have met the prerequisites listed separately in “Supported listener clustering configurations” on page 124. For example, it is important that a shared file location is already configured and that you have prepared separate machines for each node in the cluster.
2. Follow the instructions below to install a clustered listener configuration.

Step	Description
A. Start the installer.	Go to the folder where you saved the IBM Marketing Software installer and run the installer. This invokes all of the product installers that are in that location (Marketing Platform, Campaign).
B. Install Marketing Platform if necessary.	If you have not already done so, complete the screens that install Platform, and click Done on the Installation Complete window. If you need help, see Chapter 4, “Installing Campaign,” on page 49.

Step	Description
C. Install IBM Campaign in a clustered listener configuration, optionally including the first listener.	<p>The IBM Campaign installer opens. In this installer, you must configure IBM Campaign for a clustered configuration and, optionally, install the first listener in the cluster as well. One listener in the cluster may be installed on the Campaign web application server; however, each subsequent listener must be installed on its own standalone server. To install Campaign in a clustered listener configuration, complete the following screens:</p> <ul style="list-style-type: none"> • Introduction • Software License Agreement • Installation Directory • Campaign Components: Select the desired options and be sure to select Campaign Server, which is the option to install the listener. • Single or Multiple Listeners: Select Clustered listener configuration (two or more nodes) • Complete the following screens to install the first listener: <ul style="list-style-type: none"> – Shared Network File Location. The clustered listener configuration requires that certain files and folders are shared and accessible to all of the listeners in a cluster. Use this screen to specify the path to the shared file location. Use a mapped Microsoft Windows server path (such as Z:\Campaign_Shared) or a mounted UNIX path (such as /mnt/Campaign_Shared). This location is known as the campaignSharedHome. – Listener Node Properties. For each listener node that you install in a cluster, you must specify properties such as a unique name for the node, the node's network host and port, and so on. – Master Listener Priority. The priority determines which node in the listener cluster is the master listener and which node to use in the event of failover. – Load Balancing Weight. The weight determines how much listener traffic the node can support to share processing with the other nodes. Specify any value other than zero, because a value of zero prevents the node from handling any listener connections. <p>From this point in the installation process, the remaining screens are essentially the same as for a single-node install. When you complete the Pre-installation Summary screen, click Install to complete the installation of Campaign and the first listener node in the cluster.</p> <p>The Campaign installer runs with the specified options.</p>
D. Deploy and package the EAR file	<p>For the first installer node in the cluster, which you might be installing on the Campaign web application server, the installer displays screens to deploy and package the EAR file, the same as for a single-node installation.</p> <p>Continue with the process of deploying and running Campaign in your web application server, and running the listener on the Campaign server, as described in Chapter 5, "Configuring Campaign before deployment," on page 59.</p>

Step	Description
<p>E. Install the second node in the cluster.</p> <p>Important: Each listener node must be installed on a separate machine.</p>	<p>If you have not already done so, copy the IBM Marketing Software master installer and the Campaign installation files onto the server where your next Campaign installer node will run, and launch the master installer. Complete the screens as described here:</p> <ul style="list-style-type: none"> • In the master installer, provide the information necessary to connect to the Marketing Platform database, exactly the same way as for the first listener you installed. Each listener in the same cluster must use the same master installer configuration. <p>When the Campaign installer appears, complete the screens as described here:</p> <ul style="list-style-type: none"> • Introduction • Software License Agreement • Installation Directory • Campaign Components: Select only Campaign Server, because you are only installing the listener on this system. • Single or Multiple Listeners: Select Clustered listener configuration (two or more nodes) • Complete the following screens to install the second listener: <ul style="list-style-type: none"> – Shared Network File Location. The clustered listener configuration requires that certain files and folders are shared and accessible to all of the listeners in a cluster. Use this screen to specify the path to the shared file location. Use a mapped Microsoft Windows server path (such as Z:\Campaign_Shared) or a mounted UNIX path (such as /mnt/Campaign_Shared). This location is known as the campaignSharedHome. Note: The value that you enter here must be the same for every listener in the cluster. – Listener Node Properties. For each node that you install in a cluster, you must specify properties such as a unique name for the node, the node's network host and port, and so on. – Master Listener Priority. The priority determines which node in the listener cluster is the master listener, and which node to use in the event of failover. – Load Balancing Weight. The weight determines how much listener traffic the node can support for sharing the processing with the other nodes. Specify any value other than zero, because a value of zero prevents the node from handling any listener connections. <p>From this point in the installation process, the remaining screens are essentially the same as for a single-node install. When you complete the Pre-installation Summary screen, click Install to complete the installation of Campaign and the first listener node in the cluster.</p> <p>The Campaign installer runs with the options that you specified.</p>
<p>F. Restart the web app server and listeners.</p>	<p>When the installation is complete, restart the web application server and then start all listeners which are installed.</p>

Step	Description
G. Install each subsequent node in the cluster.	Repeat the installation steps that you performed for the second listener node for each additional node you want to install. Remember that each node must be installed on a separate system from the other nodes. Restart the web application server and the listener on each node as you complete its installation.
H. Optionally, adjust configuration settings.	Log in to Campaign to view or modify any of the configuration settings. See the <i>IBM Campaign Administrator's Guide</i> for instructions.

You have now installed the Campaign clustered listeners.

Supported listener clustering configurations

This topic pertains to a clustered listener configuration.

Prerequisites and requirements for IBM Campaign listener cluster configurations:

- There can be only one listener per physical host machine.
- All back-end machines for clustered listeners must be running the same operating system type.
- All back-end machines for clustered listeners must have the same version of IBM Campaign installed.
- The shared network location (campaignSharedHome) must be in place and must be accessible from each physical host machine on which you plan to install a listener node. You must set this up prior to installing the listener nodes.

Diagram of listener clustering

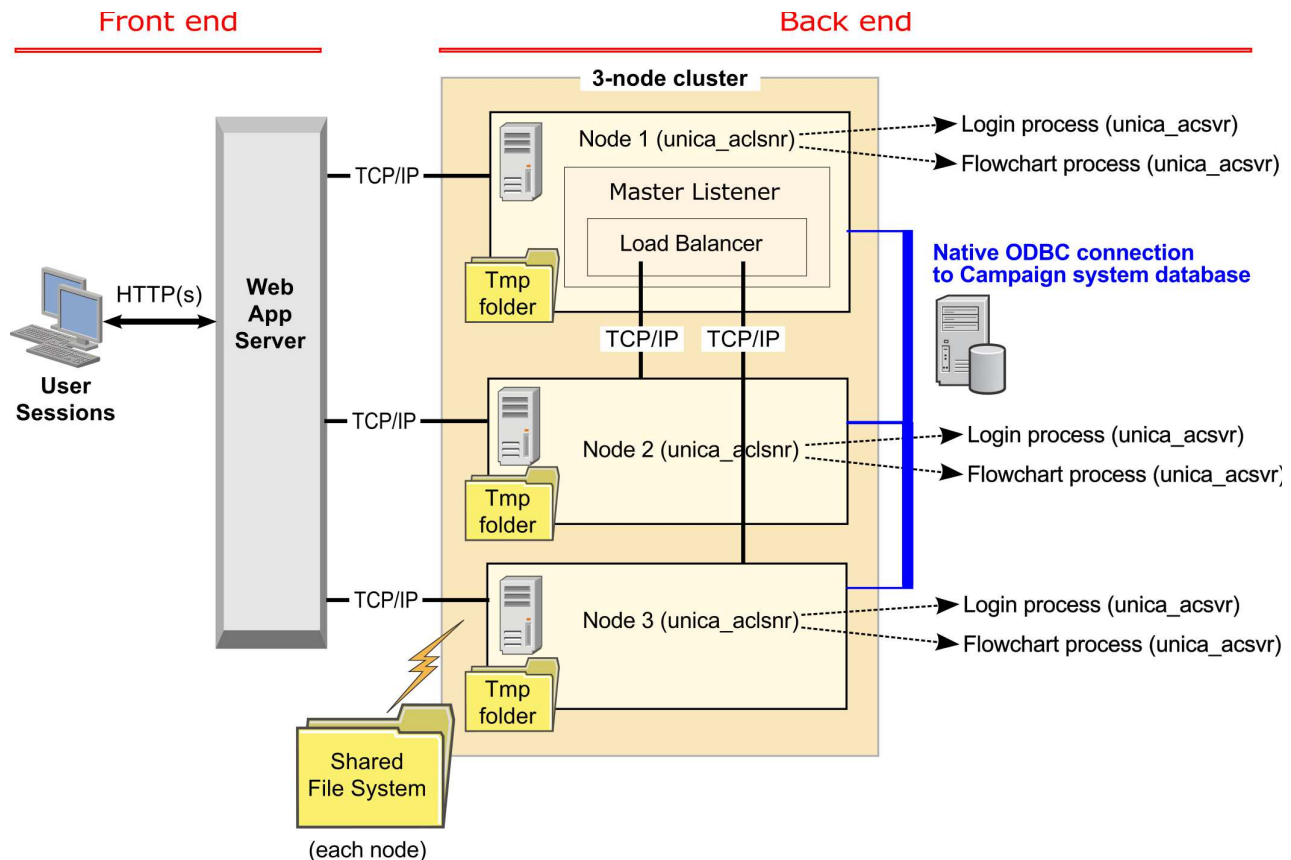
This diagram illustrates a three-node listener cluster configuration.

Note: The following summary provides a high-level overview of the components. Details are provided in separate topics.

A cluster consists of multiple listener nodes. Each node (unica_aclsnr) is on a separate physical machine, and each node has a native ODBC connection to the Campaign system database. As with a single-node configuration, each unica_aclsnr process spawns additional back-end processes for logins and flowcharts.

Each node also has a connection to a back-end user database (not shown in diagram).

In a clustered configuration, one node acts as the master listener. The job of the master listener is to perform load balancing by distributing incoming requests to each node. The Campaign web application sends client requests over TCP/IP, and the load balancer component communicates with clustered nodes over TCP/IP. All of the nodes share a network file system, so they can access shared files. Additionally, each node has its own local tmp folder and its own set of files that are not shared.



Shared network location for clustered listeners: campaignSharedHome

A clustered listener configuration for IBM Campaign requires that certain files and folders must be shared and accessible to all of the listeners in a cluster. Therefore, you must have a shared file system in place.

Requirements

- The common area can be a machine or location that is accessible from all of the other machines in the listener cluster.
- Each listener in a cluster must have full access to the shared files and folders.
- Best practice is for all of the listeners to be on the same network, and for the shared home to also be on that network, to avoid latency issues.
- To avoid any single point of failure, the shared file system should use mirrored RAID or an equivalent method of redundancy.
- If you are installing a single-listener configuration, a shared file system is a best practice, in case you decide to implement listener clusters in the future.

Shared files and folders

In a clustered configuration, all of the listeners share the folder structure shown below. The shared location (`<campaignSharedHome>`) is specified at installation time and is configurable at **Campaign | campaignClustering | campaignSharedHome**. The shared partition contains all of the logs, campaigns, templates, and other files.

```

campaignSharedHome
|--->/conf
|   |-----> activeSessions.udb
|   |-----> deadSessions.udb
|   |-----> etc.
|--->/logs
|   |-----> masterlistener.log
|   |-----> etc.
|--->/partitions
|   |-----> partition[n]
|       |-----> {similar to <Campaign_home> partition folder structure}

```

Files and folders that are not shared

Each IBM Campaign listener has its set of folders and files that are not shared, under *<Campaign_home>*. Campaign_home is an environment variable that represents the IBM Campaign application installation directory. This variable is set in cmpServer.bat (Windows) or rc.unica_ac.sh (UNIX). Partitions are specific to the local listener. Each local partition folder contains a tmp folder for temporary files during flowchart runs and a conf folder for the table manager cache file.

```

Campaign_home
|--->/conf
|   |-----> config.xml
|   |-----> unica_aclsnr.pid
|   |-----> unica_aclsnr.udb
|   |-----> etc.
|--->/logs
|   |-----> unica_aclsnr.log
|   |-----> etc.
|--->/partitions
|   |-----> partition[n]
|       |-----> /tmp
|       |-----> /conf
|       |-----> {other files specific to the partition}

```

Before you contact IBM technical support

If you encounter a problem that you cannot resolve by consulting the documentation, your company's designated support contact can log a call with IBM technical support. Use these guidelines to ensure that your problem is resolved efficiently and successfully.

If you are not a designated support contact at your company, contact your IBM administrator for information.

Note: Technical Support does not write or create API scripts. For assistance in implementing our API offerings, contact IBM Professional Services.

Information to gather

Before you contact IBM technical support, gather the following information:

- A brief description of the nature of your issue.
- Detailed error messages that you see when the issue occurs.
- Detailed steps to reproduce the issue.
- Related log files, session files, configuration files, and data files.
- Information about your product and system environment, which you can obtain as described in "System information."

System information

When you call IBM technical support, you might be asked to provide information about your environment.

If your problem does not prevent you from logging in, much of this information is available on the About page, which provides information about your installed IBM applications.

You can access the About page by selecting **Help > About**. If the About page is not accessible, check for a `version.txt` file that is located under the installation directory for your application.

Contact information for IBM technical support

For ways to contact IBM technical support, see the IBM Product Technical Support website: (http://www.ibm.com/support/entry/portal/open_service_request).

Note: To enter a support request, you must log in with an IBM account. This account must be linked to your IBM customer number. To learn more about associating your account with your IBM customer number, see **Support Resources > Entitled Software Support** on the Support Portal.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not 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 Corporation
B1WA LKG1
550 King Street
Littleton, MA 01460-1250
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some 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.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating

platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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.

Privacy Policy and Terms of Use Considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. A cookie is a piece of data that a web site can send to your browser, which may then be stored on your computer as a tag that identifies your computer. In many cases, no personal information is collected by these cookies. If a Software Offering you are using enables you to collect personal information through cookies and similar technologies, we inform you about the specifics below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user's user name, and other personal information for purposes of session management, enhanced user usability, or other usage tracking or functional purposes. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

Various jurisdictions regulate the collection of personal information through cookies and similar technologies. If the configurations deployed for this Software Offering provide you as customer the ability to collect personal information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for providing notice and consent where appropriate.

IBM requires that Clients (1) provide a clear and conspicuous link to Customer's website terms of use (e.g. privacy policy) which includes a link to IBM's and Client's data collection and use practices, (2) notify that cookies and clear gifs/web beacons are being placed on the visitor's computer by IBM on the Client's behalf along with an explanation of the purpose of such technology, and (3) to the extent required by law, obtain consent from website visitors prior to the placement of cookies and clear gifs/web beacons placed by Client or IBM on Client's behalf on website visitor's devices

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Online Privacy Statement at: <http://www.ibm.com/privacy/details/us/en> section entitled "Cookies, Web Beacons and Other Technologies."



Printed in USA