

IBM SPSS Analytic Server  
Version 2

*Installation and Configuration Guide*



**Note**

Before using this information and the product it supports, read the information in "Notices" on page 31.

**Product Information**

This edition applies to version 2, release 0, modification 0 of IBM SPSS Analytic Server and to all subsequent releases and modifications until otherwise indicated in new editions.

---

# Contents

<b>Chapter 1. Overview</b>	<b>1</b>
Architecture	2
What is new for installers in version 2.	2
<b>Chapter 2. Installation.</b>	<b>5</b>
Prerequisites	5
High-availability cluster prerequisites	5
Installation	6
<b>Chapter 3. Configuration</b>	<b>9</b>
Using the configuration tool	9
Create a user registry	12
Enabling HCatalog data sources	13
Enabling Support for Essentials for R.	15
Configuring IBM SPSS Modeler for use with IBM SPSS Analytic Server	16

<b>Chapter 4. Migration</b>	<b>19</b>
Using the migration tool	20
<b>Chapter 5. Making changes to your system.</b>	<b>23</b>
Using the administration tool	23
Manually editing configuration properties	24
<b>Chapter 6. Uninstalling</b>	<b>27</b>
<b>Chapter 7. Troubleshooting</b>	<b>29</b>
<b>Notices</b>	<b>31</b>
Trademarks	33



---

## Chapter 1. Overview

IBM® SPSS® Analytic Server is a solution for big data analytics that combines IBM SPSS technology with big data systems and allows you to work with familiar IBM SPSS user interfaces to solve problems on a previously unattainable scale.

### Why big data analytics matters

Data volumes collected by organizations are growing exponentially; for example, financial and retail businesses have all customer transactions for a year (or two years, or ten years), telco providers have call data records (CDR) and device sensor readings, and internet companies have the results of web crawls.

Big data analytics is needed where there exists:

- A large volume of data (terabytes, petabytes, exabytes), especially when it is a mixture of structured & unstructured data
- Rapidly changing/accumulating data

Big data analytics also assists when:

- A large number (thousands) of models are being built
- Models are frequently built/refreshed

### Challenges

The same organizations that collect large volumes of data often have difficulty actually making use of it, for a variety of reasons:

- The architecture of traditional analytic products are not suited to distributed computation, and
- Existing statistical algorithms are not designed to work with big data (these algorithms expect the data to come to them, but big data is too costly to move), thus
- Performing state of the art analytics on big data requires new skills and intimate knowledge of big data systems. Very few analysts have these skills.
- In-memory solutions work for medium-size problems, but do not scale well to truly big data.

### Solution

Analytic Server provides:

- A data-centric architecture that leverages big data systems, such as Hadoop Map/Reduce with data in HDFS.
- A defined interface to incorporate new statistical algorithms designed to go to the data.
- Familiar IBM SPSS user interfaces that hide the details of big data environments so that analysts can focus on analyzing the data.
- A solution that is scalable to any size problem.

---

## Architecture

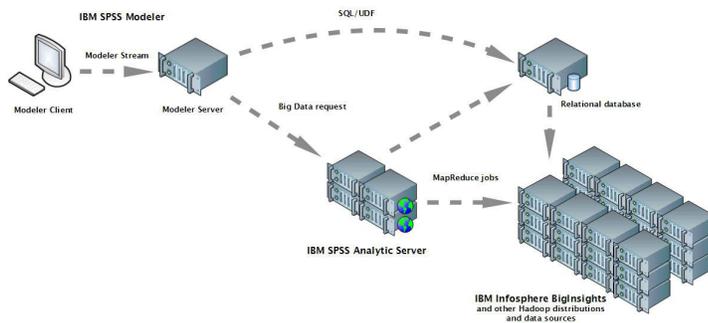


Figure 1. Architecture

Analytic Server sits between a client application and Hadoop cloud. Assuming that the data resides in the cloud, the general outline for working with Analytic Server is to:

1. Define Analytic Server data sources over the data in the cloud.
2. Define the analysis you want to perform in the client application. For the current release, the client application is IBM SPSS Modeler.
3. When you run the analysis, the client application submits an Analytic Server execution request.
4. Analytic Server orchestrates the job to run in the Hadoop cloud and reports the results to the client application.
5. You can use the results to define further analyses, and the cycle repeats.

---

## What is new for installers in version 2

### Installation and configuration

Product installation now uses the IBM Installation Manager to provide a seamless experience when you install multiple IBM products.

The steps to configure Analytic Server are handled in a new configuration tool that is run after all product extensions are installed.

### Platform

Support for new operating systems and Hadoop distributions.

#### Operating systems

Analytic Server now runs on Microsoft Windows.

#### Hadoop distributions

Analytic Server can be configured to run with Hortonworks Data Platform on Microsoft Windows.

#### HCatalog adapters

Data sources can be defined for Apache Cassandra, MongoDB, and Oracle NoSQL. There are more configuration steps to enable these data sources.

## High-availability clusters

Analytic Server can be installed and configured on a high-availability cluster using IBM WebSphere® Liberty Network Deployment.



---

## Chapter 2. Installation

---

### Prerequisites

Before installing Analytic Server, review the following information.

#### System requirements

For the most up-to-date system requirements information, use the Detailed system requirements reports at the IBM Technical Support site: <http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/softwareReqsForProduct.html>. On this page:

1. Type SPSS Analytic Server as the product name and click **Search**.
2. Select the wanted version and scope of report, then click **Submit**.

#### Administrative privileges

The user that installs Analytic Server must have write access to the Hadoop Distributed File System.

**Note:** When installing Analytic Server for use with BigInsights Adaptive MapReduce, the user that installs Analytic Server must be the same user that run the BigInsights Adaptive MapReduce services.

#### IBM Installation Manager

A current version of IBM Installation Manager must be installed and configured to use a repository that contains Analytic Server installation files. If you have an older version of IBM Installation Manager, you are required to update it to a supported version as part of the installation.

If IBM Installation Manager is not present on the system, download and install IBM Installation Manager from IBM Corp. support site: <http://www.ibm.com/support>.

**Important:** On Microsoft Windows, install IBM Installation Manager from an Administrator account. It is possible to do a group installation as another user with: `./groupinstc -acceptLicense`. See the IBM Installation Manager documentation for details. [http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im\\_family\\_welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html)

#### Hive/HCatalog

If you plan to use NoSQL data sources, then configure Hive and HCatalog for remote access. Also ensure that `hive-site.xml` contains a `hive.metastore.uris` property in the form `thrift://<host_name>:<port>` that points to the active Thrift Hive Metastore server. Refer to your Hadoop distribution documentation for details.

#### Metadata repository

By default, Analytic Server installs and uses an Apache Derby database. Alternatively, you can configure Analytic Server to use an existing DB2 installation. The database must have an encoding of UTF-8. IBM DB2 10.5 Workgroup Server Edition is available on the Analytic Server media. For more information: [http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html).

#### High-availability cluster

See “High-availability cluster prerequisites” for the prerequisites to installing Analytic Server on a high-availability cluster.

### High-availability cluster prerequisites

There are extra prerequisites for installing Analytic Server on a high-availability cluster.

#### Systems and DNS configuration

The systems that are used in the environment need to be able to communicate with one another. Only systems that belong to a domain should have the domain suffix. For example, only systems that are registered with the "IBM.COM" domain should have the "ibm.com" suffix for the Fully Qualified Domain Name (FQDN).

If a system not on the IBM.COM domain does not have an FQDN suffix of "ibm.com", then the FQDN for the system must be entered in DNS with the systems that have the DNS server in the TCP configuration. Otherwise, all FQDNs of all systems that are involved must be in the hosts files of all systems involved.

### Metadata repository

Analytic Server must use DB2 when installed in a high-availability cluster.

### Zookeeper 3.4.5

A minimum of three Zookeeper instances are required. See the Zookeeper installation and configuration documentation: [http://zookeeper.apache.org/doc/r3.4.5/zookeeperStarted.html#sc\\_RunningReplicatedZooKeeper](http://zookeeper.apache.org/doc/r3.4.5/zookeeperStarted.html#sc_RunningReplicatedZooKeeper).

**Note:** These steps must be performed on each system where Zookeeper is installed.

### Liberty Network Deployment 8.5.5.4

The high-availability environment consists of a Liberty collective with one controller, two replicas, and at least four cluster members. See the WebSphere Application Server Network Deployment documentation for details on configuring a Liberty collective: [http://www-01.ibm.com/support/knowledgecenter/SSAW57/mapfiles/product\\_welcome\\_wasnd.html](http://www-01.ibm.com/support/knowledgecenter/SSAW57/mapfiles/product_welcome_wasnd.html).

**Important:** All members of the collective must be up and running in order to be enabled for use with Analytic Server.

### Load balancer

Your high availability cluster should have a load balancer that supports session affinity, sometimes also known as sticky sessions. Analytic Server identifies sessions with the cookie "request-token". This identifies a session for the duration of a user login for use in application-controlled session affinity. Please consult the documentation of your particular load balancer for the details of how it supports session affinity.

### Visual C++ 2008 (Windows only)

If you have a cluster member running Analytic Server on Windows, you will need to install the Visual C 2008 Runtime before installing Analytic Server.

---

## Installation

Analytic Server application files are installed on the host system with IBM Installation Manager. Installation files can be found on the physical media or downloaded from IBM Passport Advantage.

#### 1. Start IBM Installation Manager:

GUI mode

```
{IBM Installation Manager installation directory}/eclipse/IBMIM
```

Command-line mode

```
{IBM Installation Manager installation directory}/eclipse/tools/imcl -c
```

**Note:** The user that runs IBM Installation Manager is the user that runs the Analytic Server process. This user requires access to the Hadoop environment.

2. If the installation repository is not configured, specify the repository path. For example, specify the mounted installation disk or a location on the host file system, the network, or an HTTP address.
3. On the main menu, select **Install**.

4. Select the package to install. The option available depends upon the version of Analytic Server you purchased.

**IBM SPSS Analytic Server**

This package is suitable for most systems.

**IBM SPSS Analytic Server High Availability**

This package is required if you plan to run Analytic Server in a high-availability cluster. Differences in the configuration of IBM SPSS Analytic Server High Availability are documented in place.

5. Read the license agreement and accept its terms.
6. Specify the package group and the installation directory.
  - A new package group is required for IBM SPSS Analytic Server install.
  - Specify the installation directory for shared resources. You can specify the shared resources directory only the first time that you install a package.
7. Select features to be installed.
8. Review summary information and proceed with the installation. The application files are installed in the specified directory.
9. Further configuration is necessary. For more information on a fresh configuration, see the topic: "Using the configuration tool" on page 9. For information on migrating an existing configuration to a new installation, see the topic: Chapter 4, "Migration," on page 19. If you are running IBM Installation Manager in GUI mode, the option to start the configuration or migration tool is presented on the last step.

If problems occur during the installation, you can troubleshoot them using IBM Installation Manager logs. The log files can be accessed from the main menu of IBM Installation Manager.

To automate installation, IBM Installation Manager can be run in the silent mode. For more information, see IBM Installation Manager documentation: [http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im\\_family\\_welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html).



---

## Chapter 3. Configuration

After installation, you must configure Analytic Server before you can use the product.

**Note:** The following conventions are used for Analytic Server file paths.

- {AS\_ROOT} refers to the location where Analytic Server is deployed; for example, /opt/IBM/SPSS/AnalyticServer/{version}.
- {AS\_SERVER\_ROOT} refers to the location of the configuration, log, and server files; for example, /opt/IBM/SPSS/AnalyticServer/{version}/ae\_wlpsrver/usr/servers/aeserver.
- {AS\_HOME} refers to the location on HDFS that is used by Analytic Server as a root folder.

### Configuration steps

1. [Optional] Follow the steps to enable HCatalog data sources if you want to run Analytic Server on NoSQL or XML data sources.

**Note:** You can perform a simple configuration Analytic Server now and enable HCatalog data sources later using the administration tool.

2. Run the configuration tool. The tool automates a number of otherwise manual steps. If you used the GUI version of IBM Installation Manager, the configuration tool launches automatically, otherwise you need to run the **configtool** script in {AS\_ROOT}/bin.
3. Create a registry of users who can be added to the Analytic Server system.
4. Configure IBM SPSS Modeler for use with Analytic Server.
5. [Optional] Enable Support for Essentials for R if you want Analytic Server to run R code.
6. Make sure that Analytic Server is running. You can start the process by running the **start** script in {AS\_ROOT}/bin.
7. Open a web browser and enter the address `http://<host>:<port>/analyticserver/admin/ibm`, where <host> is the address of the Analytic Server host, and <port> is the port that Analytic Server is listening on. By default this is 8080. This URL opens the login dialog for the Analytic Server console. Log in as the Analytic Server administrator defined in the configuration tool.

---

### Using the configuration tool

The configuration tool automates a number of otherwise manual steps. The configuration tool must have access to the HDFS environment.

If you used the GUI version of IBM Installation Manager, the configuration tool launches automatically. Follow the prompts to configure your system.

You can also launch the configuration tool by executing the **configtool** script in {AS\_ROOT}\bin. By default, the tool detects whether to launch in GUI or console mode. You can force console mode by passing it the argument **-c**, or silent mode by passing it the path to a file that specifies the desired configuration. See the `configtool.response` file in {AS\_ROOT}\bin for a sample.

### Welcome

The Welcome step determines what path you take through the Configuration tool.

#### Typical

The Typical path automates many of decisions and makes simplifying assumptions about your

system. Select this option if you want to quickly configure a non-production environment. You can always change your configuration in the Administration tool or migrate your configuration to another environment later.

In a Typical configuration, the tool will prompt you for the location of the Hadoop namenode, and progressively ask for more information when needed.

In a Typical configuration of a High Availability installation, the configuration tool will prompt you for information about your IBM WebSphere® Application Server Liberty Network Deployment environment and your Zookeeper ensemble.

### Custom

This allows you perform more advanced configuration tasks. This option is required if you want to set up NoSQL or relational databases for use with Analytic Server right now.

The sections below provide more details about the custom path.

## Hadoop cluster

The Hadoop cluster steps configure Analytic Server for use with your Hadoop distribution. The configuration tool can automatically discover your distribution, or you can manually select the necessary files.

### Automatically collect files from your Hadoop distribution

The tool will attempt to detect the Hadoop distribution, and progressively ask for more information when needed.

1. You will always need to provide host name of the Hadoop cluster.
2. If your cluster is not set up with SSH and public key authentication for the user running the configuration tool, you will need to provide valid credentials so that the configuration tool can gather files from the Hadoop cluster.

**Restriction:** Only RSA SSH encryption is supported.

3. If your Hadoop cluster has been installed in a non-standard location, you may need to provide the installation location.

### Manually collect files from your Hadoop distribution

Identify the Hadoop distribution you are using. The configuration tool will identify the files you need to gather from the Hadoop cluster. Copy these files to the specified location and you can continue the configuration steps.

**Note:** The `hadoop-lzo-*.jar` JAR file is required only when the `"com.hadoop.compression.lzo.LzoCodec"` value is specified for the `io.compression.codecs` parameter in the `core-site.xml` file.

## Hcatalog/Hive

The Hcatalog/Hive steps are only shown as part of a Custom configuration.

### Automatically connect and detect your HCatalog / Hive configuration

The tool will attempt to detect your version of HCatalog / Hive and collect the necessary files to configure Analytic Server, and progressively ask for more information when needed.

1. You will always need to provide host name of the server where Hive is installed.
2. If the server is not set up with SSH and public key authentication for the user running the configuration tool, you will need to provide valid credentials so that the configuration tool can gather files from Hive.

**Restriction:** Only RSA SSH encryption is supported.

## Manually select your HCatalog / Hive configuration

Identify the version of Hive you have installed. Optionally identify the NoSQL databases you are using. The configuration tool will identify the files you need to gather. Copy these files to the specified location and you can continue the configuration steps.

## Skip HCatalog / Hive configuration

You can configure HCatalog / Hive later.

## JDBC

The JDBC steps are only shown as part of a Custom configuration. Copy the drivers for the databases you are using to the specified location and continue the configuration steps.

Table 1. Supported databases

Database	Supported versions	JDBC driver jars	Vendor
DB2 for Linux, UNIX, and Windows	10.5, 10.1, 9.7	db2jcc.jar	IBM
DB2 z/OS	11, 10	db2jcc.jar, db2_license_cisuz.jar	IBM
Greenplum	5, 4.2.x	postgresql.jar	Greenplum
Netezza	7, 6.x	nzjdbc.jar	IBM
Oracle	12c, 11g R2 (11.2)	ojdbc6.jar, orai18n.jar	Oracle
SQL Server	2014, 2012, 2008 R2	sqljdbc4.jar	Microsoft
Sybase IQ	16.x, 15.4, 15.2	jconnect70.jar	Sybase
Teradata	14, 14.1, 15	tdgssconfig.jar, terajdbc4.jar	Teradata

## Application Server

The Application Server steps are dependent upon the type of installation.

- In a non-High Availability installation, you can configure secure communication (SSL/TLS) and set the ports on which Analytic Server listens.
- In a High Availability installation, the configuration tool will prompt you for information about your IBM WebSphere® Application Server Liberty Network Deployment environment.

## Zookeeper

The Zookeeper step is only required as part of the configuration for a High Availability installation. Provide the list of servers in your Zookeeper ensemble.

## Metadata repository

The Metadata repository steps are dependent upon the type of installation.

- In a non-High Availability installation, you have a choice of Derby or DB2. If you choose DB2, it must be installed and available prior to running the configuration tool. See “Prerequisites” on page 5 for more information.
- A configuration of a High Availability installation must always use DB2 as the metadata repository, and the DB2 instance must be installed and available prior to running the configuration tool.

## Analytic Server

Identify the principal that serves as the Analytic Server administrator. This principal will be able to create, delete, and manage tenants.

### Summary

The Summary step allows you to review the configuration steps before completing configuration.

### Post-configuration

The Post-configuration step will provide information that the configuration was successful. In the event of a failure, review the log file in the logs directory.

---

## Create a user registry

The default tenant created during installation and configuration uses WebSphere Liberty profile to authenticate users, which in turn can reference an external LDAP server or define a basic registry of users and groups.

**Note:** After you edit the `{AS_SERVER_ROOT}/server.xml` file to configure WebSphere Liberty Profile to authenticate users, add those users to the Analytic Server system. See the *IBM SPSS Analytic Server Administrator's Guide* for details on tenant management.

### LDAP

The LDAP registry allows you to authenticate users with an external LDAP server such as Active Directory or OpenLDAP. Here is an example of an `ldapRegistry` for OpenLDAP.

```
<ldapRegistry
  baseDN="ou=people,dc=aeldap,dc=org"
  ldapType="Custom"
  port="389"
  host="server"
  id="OpenLDAP"
  bindDN="cn=admin,dc=aeldap,dc=org"
  bindPassword="{xor}Dz4sLG5tbGs="
  searchTimeout="300000m"
  recursiveSearch="true">
  <customFilters
    id="customFilters"
    userFilter="(&!(uid=%v)(objectClass=inetOrgPerson))"
    groupFilter="(&(cn=%v)(|(objectclass=organizationalUnit)))"
    groupMemberIdMap="posixGroup:memberUid"/>
</ldapRegistry>
```

For more examples of configurations, see the templates folder `{AS_ROOT}/ae_wlpserver/templates/config`.

### Basic

The Basic Registry allows you to define a database of users and groups within the `{AS_SERVER_ROOT}/server.xml` file. Passwords can be encoded to obfuscate their values with the `securityUtility` tool, which is in `{AS_ROOT}/ae_wlpserver/bin`.

The Basic registry is useful in a sandbox environment, but is not recommended for a production environment.

```
<basicRegistry id="basic" realm="ibm">
  <user name="user1" password="{xor}Dz4sLG5tbGs="/>
  <user name="user2" password="Pass"/>
  <user name="user3" password="Pass"/>
</basicRegistry>
```

```

<user name="user4" password="Pass"/>
<user name="admin" password="{xor}KzosKw==" />
<group name="Development">
  <member name="user1"/>
  <member name="user2"/>
</group>
<group name="QA">
  <member name="user3"/>
  <member name="user4"/>
</group>
<group name="ADMIN">
  <member name="user1"/>
  <member name="admin"/>
</group>
</basicRegistry>

```

---

## Enabling HCatalog data sources

Analytic Server provides support for XML data and NoSQL databases through Hive/HCatalog.

The configuration tool prepares your Analytic Server installation for use with Hive/HCatalog, but there are some manual steps. You need to copy some JAR files to the Hive `{HIVE_HOME}/auxlib/` directory, because Analytic Server implements some interfaces required by HCatalog. Analytic Server cannot use storage handlers for Hive directly in HCatalog because HCatalog has a slightly different API.

The following sections detail the necessary files, based upon what data sources you want to use.

**Important:** These files must be copied to the `{HIVE_HOME}/auxlib/` directory, and you must restart the Hive metastore, prior to running the configuration tool or administration tool in order for Analytic Server to make use of them.

**Important:** Hortonworks Data Platform on Windows does not recognize the `{HIVE_HOME}/auxlib/` directory. To enable HCatalog, you can:

1. Copy the necessary files, as described below, to the `{HIVE_HOME}/lib/` directory.
2. Remove the `{HIVE_HOME}/auxlib/` directory.
3. Update the `hive-site.xml` and `hive-config.cmd` files to remove any references to the `{HIVE_HOME}/auxlib/` directory.
4. Restart the Hortonworks Data Platform cluster for the changes to take effect.

See <http://hortonworks.com/community/forums/topic/hive-auxlib-issues/> for more information.

### All Hive/HCatalog installations

1. The following files can be found in the `hive_auxlib` archive in the `{AS_ROOT}/add-ons` directory.

```
com.ibm.spss.com.hcatalogstoragehandler-{version}.jar
```

2. Copy the latest version of `hivexmlserde-{version}.jar` from the Maven Central Repository at <http://search.maven.org/#search%7Cga%7C1%7Chivexmlserde>.
3. [Optional] If you plan to use Hive complex types (maps, arrays, structures) copy these JAR files. They can be found in the HCatalog installation. Prior to Hive 0.11.0, this is in `{HCATALOG_HOME}/share/hcatalog`; starting with Hive 0.11.0, this is in `{HIVE_HOME}/hcatalog/share/hcatalog`.

```

hcatalog-core-{version}.jar (for hadoop 1.2)
hive-hcatalog-{version}.jar (for hadoop 2.x)
jackson-core-asl-1.8.8.jar
jackson-mapper-asl-1.8.8.jar

```

4. If you are using Hive version 0.14.0, then download the following JAR file. <https://github.com/dvasilen/Hive-HCatalog-Compatibility/raw/master/release/hive-hcatalog-compatibility-0.14.0.jar>

## XML data

1. [Optional] If you plan to work with compressed XML files, copy the following JAR files found in the `hive_auxlib` archive in the `{AS_ROOT}/add-ons` directory.  
`com.ibm.spss.hivexmlserdecomp-2.0.0.0.jar`

## Apache Accumulo

1. The following files can be found in the HCatalog installation. Prior to Hive 0.11.0, this is in `{HCATALOG_HOME}/share/hcatalog`; starting with Hive 0.11.0, this is in `{HIVE_HOME}/hcatalog/share/hcatalog`.  
`commons-io-{version}.jar`
2. The Apache Accumulo storage handler can be found at <https://github.com/dvasilen/accumulo-hive-storage-manager/tree/master>. This is the master branch; from here you should switch to the branch for your version of Accumulo and Hadoop. The JAR file can be found in the `/release` directory and has a name of the form:  
`accumulo-hive-storage-handler_{hadoopVersion}-<accumuloVersion>.jar`

3.

The following files can be found in the Apache Accumulo installation.

```
accumulo-trace-{version}.jar
accumulo-fate-{version}.jar
accumulo-core-{version}.jar
accumulo-server-{version}.jar
accumulo-start-{version}.jar
```

## Apache Cassandra

1. The following files can be found in the Apache Cassandra 1.2.9 release at <http://www.apache.org/dyn/closer.cgi?path=/cassandra/1.2.9/apache-cassandra-1.2.9-bin.tar.gz>  
`apache-cassandra-1.2.9.jar`  
`apache-cassandra-clientutil-1.2.9.jar`  
`apache-cassandra-thrift-1.2.9.jar`
2. The Apache Cassandra storage handler can be found at <https://github.com/dvasilen/Hive-Cassandra>. This is the master branch; from here you should switch to the branch for your version of Cassandra and Hadoop. The JAR file can be found in the `/release` directory and has a name of the form:  
`hive-<hiveVersion>-hadoop-<hadoopVersion>-cassandra-<cassandraVersion>.jar`

**Note:** For Hadoop versions prior to 2.0, there is no hadoop portion of the file or branch name.

**Note:** Analytic Server can communicate with Cassandra 2.0, but its JAR files are compiled for Java 7. Thus, you need to copy the Cassandra 1.2.9 JAR files because they are compiled for Java 6.

## HBase

There are no extra files to copy for HBase.

## MongoDB

1. The following files should be built locally for your Hadoop distribution from the "master" branch of the "mongodb/mongo-hadoop" project. See <https://github.com/mongodb/mongo-hadoop> for details.  
`mongo-hadoop-core-{version}.jar`  
`mongo-hadoop-hive-{version}.jar`
2. The following file can be found at <https://github.com/downloads/mongodb/mongo-java-driver/mongo-2.10.1.jar>  
`mongo-2.10.1.jar`

## Oracle NoSQL

1. The following file can be found at <https://github.com/vilcek/HiveKVStorageHandler2/blob/master/dist/HiveKVStorageHandler.jar>  
HiveKVStorageHandler.jar
2. The following file can be found in the Apache Hive installation  
avro-\*<version>.jar

---

## Enabling Support for Essentials for R

Analytic Server supports scoring R models and running R scripts.

**Restriction:** Support for R is currently limited to Analytic Server installations on Linux.

To configure support for R after a successful Analytic Server installation:

1. Install R Engine on the server that hosts Analytic Server.

```
wget http://cran.r-project.org/src/base/R-3/R-3.1.0.tar.gz
tar -xzf R-3.1.0.tar.gz
cd R-3.1.0
./configure --enable-R-shlib
make
```

**Important:** Install the R Engine in a location accessible to all users, as R runs on the Hadoop cluster as a different user from the Analytic Server user.

**Important:** In a High Availability installation, the R Engine must be installed in the same location on every cluster member.

2. Install Essentials for R on the server that hosts Analytic Server by running the installer file, following the instructions on the screen. Essentials for R is available for download (<https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=swg-tspssp>).
3. Deploy R Engine and R Component to Hadoop.

- a. If the Analytic Server and all the Hadoop nodes have the same version of the operating system and the same processor architecture:

- 1) Create an archive for the R Engine by using the following script

```
#!/usr/bin/env bash
echo Creating R.zip...
cd /tmp
rm -r -f R
rm -f R.zip
mkdir R
export R_HOME=/home/hdpadmin/APPS/R/R-3.1.0
cp -r $R_HOME/* ./R/
cp $(ldd ./R/bin/exec/R ./R/bin/Rscript|cut -d\ -f3|grep \.so\.|sort -u) ./R/lib/
cp -P /usr/lib64/libgfortran.so.3 ./R/lib
cp /usr/lib64/libgfortran.so.3.0.0 ./R/lib
rm -r ./R/doc ./R/src ./R/include ./R/tests
zip -r R.zip R
```

**Note:** There are two spaces between `-d\` and `-f3` in the script. The PDF document gives the appearance of two spaces, but it is a single wide space. When you copy and paste this script from the PDF, insert the second space before you run the script.

**Note:** The lines that copy Fortran libraries are dependent upon the version of Fortran installed. If, for example, the 1.0 version is installed, those lines read as follows. Edit as needed for your system.

```
cp -P /usr/lib64/libgfortran.so.1 ./R/lib
cp /usr/lib64/libgfortran.so.1.0.0 ./R/lib
```

- 2) Copy the archive R.zip to the {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/configuration/linux/app\_64 directory.
- b. If the versions of the operating systems for the Analytic Server and Hadoop nodes are different, install the R engine and Essentials for R on each Hadoop node. Install them into a directory with the same path and name that the R engine is installed on the Analytic Server. If you install Essentials for R on the Hadoop node, specify the directory in which R is installed and skip the step of specifying the location of ../ext\_64/bin.
- c. Run the **hdfsUpdate** script in the {AS\_ROOT}/bin directory to propagate the changes to HDFS.
4. On a High Availability installation, you must also run the **deployRemote** script in the {AS\_ROOT}/bin directory to propagate the changes to the Analytic Server instances on each cluster member.
5. You must also install Essentials for R on the machine that hosts SPSS Modeler Server.

---

## Configuring IBM SPSS Modeler for use with IBM SPSS Analytic Server

In order to enable SPSS Modeler for use with Analytic Server, you need to make some updates to the SPSS Modeler Server installation.

1. Configure SPSS Modeler Server to associate it with an Analytic Server installation.
  - a. Edit the options.cfg file in the config subdirectory of the main server installation directory, and add or edit the following lines:

```
as_ssl_enabled, {Y|N}
as_host, "{AS_SERVER}"
as_port, PORT
as_context_root, "{CONTEXT-ROOT}"
as_tenant, "{TENANT}"
as_prompt_for_password, {Y|N}
as_kerberos_auth_mode, {Y|N}
as_kerberos_krb5_conf, {CONF-PATH}
as_kerberos_krb5_spn, {AS-SPN}
```

### **as\_ssl\_enabled**

Specify Y if secure communication is configured on Analytic Server; otherwise, N.

### **as\_host**

The IP address of the server that hosts Analytic Server.

### **as\_port**

The port on which Analytic Server is listening (by default this is 8080).

### **as\_context\_root**

The Analytic Server context root (by default this is analyticserver).

### **as\_tenant**

The tenant the SPSS Modeler Server installation is a member of (the default tenant is ibm).

### **as\_prompt\_for\_password**

Specify N if the SPSS Modeler Server is configured with the same authentication system for users and passwords as that used on Analytic Server; for example, when using Kerberos authentication. Otherwise, specify Y.

When running SPSS Modeler in batch mode, you add `-analytic_server_username {ASusername} -analytic_server_password {ASpassword}` as arguments to the `clem` command.

### **as\_kerberos\_auth\_mode**

Specify Y to enable Kerberos SSO from SPSS Modeler.

### **as\_kerberos\_krb5\_conf**

Specify the path to the Kerberos configuration file that Analytic Server should use; for example, `\etc\krb5.conf`.

### **as\_kerberos\_krb5\_spn**

Specify the Analytic Server Kerberos SPN; for example, HTTP/ashost.mydomain.com@MYDOMAIN.COM.

- b. Restart the SPSS Modeler Server service.

In order to connect to an Analytic Server installation that has SSL/TLS enabled, there are some further steps to configuring your SPSS Modeler Server and client installations.

- a. Navigate to `http{s}://{HOST}:{PORT}/{CONTEXT-ROOT}/admin/{TENANT}` and log on to the Analytic Server console.
- b. Download the certification file from the browser and save it to your file system.
- c. Add the certification file to the JRE of both your SPSS Modeler Server and SPSS Modeler Client installations. The location to update can be found under the `/jre/lib/security/cacerts` subdirectory of the SPSS Modeler installation path.
  - 1) Make sure the `cacerts` file is not read-only.
  - 2) Use the `keytool` program Modeler ships with – this can be found in the `/jre/bin/keytool` subdirectory of the SPSS Modeler installation path.

Run the following command

```
keytool -import -alias <as-alias> -file <cert-file> -keystore "<cacerts-file>"
```

Note that `<as-alias>` is an alias for the `cacerts` file. You can use any name you like as long as it is unique to the `cacerts` file.

So an example command would look like the following.

```
keytool -import -alias MySSLCertAlias -file C:\Download\as.cer  
-keystore "c:\Program Files\IBM\SPSS\Modeler\{ModelerVersion}\jre\lib\security\cacerts"
```

- d. Restart your SPSS Modeler Server and SPSS Modeler Client .
2. [optional] Install IBM SPSS Modeler - Essentials for R , if you plan to score R models in streams with Analytic Server data sources. IBM SPSS Modeler - Essentials for R is available for download (<https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=swg-tspssp>).



---

## Chapter 4. Migration

Analytic Server supports a number of scenarios in which you migrate configuration settings from an existing Analytic Server installation to a new installation.

### Upgrade to a new version of Analytic Server

If you have an existing installation of Analytic Server 1.0.1 and have purchased 2.0, then you can migrate your 1.0.1 configuration settings to your 2.0 installation. In this scenario, the migration tool will guide you down the appropriate path.

**Restriction:** If you have version 1.0 installed, you must first migrate from version 1.0 to 1.0.1 and then from version 1.0.1 to 2.0.

**Restriction:** Your 1.0.1 and 2.0 installations cannot coexist on the same Hadoop cluster. If you configure your 2.0 installation to use the same Hadoop cluster as your 1.0.1 installation, the 1.0.1 installation will no longer function.

### Create a copy of your existing Analytic Server

If you have an existing installation of Analytic Server 2.0, you can replicate configuration settings in a new 2.0 installation. This is useful, for example, when you have an existing development environment and need to set up test and production environments.

**Note:** The user that installs the new installation of Analytic Server should be the same user that installed the existing installation of Analytic Server. If this user belongs to a group on the Hadoop cluster with which the old Analytic Server installation is configured, they must belong to a group with the same name on the Hadoop cluster with which the new Analytic Server installation is configured.

### Migration steps, 1.0.1 to 2.0

1. Install the new installation of Analytic Server according to the instructions in “Installation” on page 6.
2. At the end of the installation process, choose to migrate a configuration. This will automatically launch the migration tool. Complete the following steps before returning to the migration tool.
  - a. Copy the `configcollector.zip` archive in your new installation to `{AS_ROOT}\tools` in your old installation.
  - b. Extract the copy of `configcollector.zip`. This creates a new `configcollector` subdirectory in your old installation.
  - c. Run the configuration collector tool in your old installation by executing the **configcollector** script in `{AS_ROOT}\tools\configcollector`. Note the location of the resulting compressed (.zip) file.
  - d. Copy the analytic root. For example, if you installed version 1.0.1 with an `analytic-root` value of `/user/aeuser/analytic-root` and want to migrate to an `analytic-root` location of `/user/aeuser/analytic-root_new`, then you should run the following copy script.

```
hadoop distcp /user/aeuser/analytic-root /user/aeuser/analytic-root_new
```
3. You will need to provide the migration tool with the location of the new analytic root.
4. At the end of the migration tool, follow any post-configuration tool steps according to the instructions in Chapter 3, “Configuration,” on page 9.

**Note:** If you configured R for use with the existing Analytic Server installation, you will need to follow the steps to configure it with the new Analytic Server installation.

## Migration steps, 2.0 to 2.0

1. Extract the configcollector.zip archive in {AS\_ROOT}\tools. This creates a new configcollector subdirectory.
2. Run the configuration collector tool in your existing installation by executing the **configcollector** script in {AS\_ROOT}\tools\configcollector. Note the location of the resulting compressed (.zip) file.
3. Install the new installation of Analytic Server according to the instructions in "Installation" on page 6.
4. At the end of the installation process, choose to migrate a configuration. This will automatically launch the migration tool.
5. At the end of the migration tool, follow any post-configuration tool steps according to the instructions in Chapter 3, "Configuration," on page 9.

**Note:** If you configured R for use with the existing Analytic Server installation, you will need to follow the steps to configure it with the new Analytic Server installation.

---

## Using the migration tool

The migration tool automates a number of otherwise manual steps for migrating your existing configuration to a new installation. The migration tool must have access to the HDFS environment.

If you used the GUI version of IBM Installation Manager, the migration tool launches automatically. Follow the prompts to configure your system.

You can also launch the migration tool by executing the **migrationtool** script in {AS\_ROOT}\bin. By default, the tool detects whether to launch in GUI or console mode. You can force console mode by passing it the argument **-c**.

### Welcome

The Welcome step requires you to provide the compressed (.zip) file containing the Analytic Server configuration you are migrating.

### Hadoop cluster

The Hadoop cluster steps configure Analytic Server for use with your Hadoop distribution.

#### Same Hadoop cluster as original installation

This is a "destructive" copy in which the new installation takes over ownership of the Hadoop cluster from the old installation. This is appropriate if you have no use for the old Analytic Server installation.

#### Different Hadoop cluster

The new installation is connected to a different Hadoop cluster from the old installation.

The screens that follow this initial choice are similar to that of a Typical path through the configuration tool.

### Hcatalog/Hive

The Hcatalog/Hive steps are only shown if your original installation has HCatalog / Hive data sources.

- In a Copy migration to the same Hadoop cluster, you have the option of turning off Hive.
- In a Copy migration to a different Hadoop cluster, you are given the full set of options from the Custom path through the configuration tool for configuring HCatalog / Hive.

## JDBC

In a 1.0.1 to 2.0 migration, if you plan to use database data sources, you need to add JDBC drivers to the new installation through the administration tool.

In a 2.0 to 2.0 migration, any JDBC drivers used in the original installation are migrated to the new installation.

## Application Server

The Application Server steps are dependent upon the type of the new installation.

- In a non-High Availability installation, these steps are skipped. You can configure secure communication (SSL) and set the ports on which Analytic Server listens through the administration tool.
- In a High Availability installation, the migration tool will prompt you for information about your IBM WebSphere® Application Server Liberty Network Deployment environment.

**Note:** If the original installation was High Availability, these steps are prepopulated with the values from the original installation.

## Zookeeper

The Zookeeper step is only required as part of the configuration for a High Availability installation. Provide the list of servers in your Zookeeper ensemble.

**Note:** If the original installation was High Availability, these steps are prepopulated with the values from the original installation.

## Metadata repository

The Metadata repository steps are dependent upon the type of the new installation.

- In a non-High Availability installation, you have a choice of Derby or DB2. If you choose DB2, it must be installed and available prior to running the migration tool. See “Prerequisites” on page 5 for more information.
- A configuration of a High Availability installation must always use DB2 as the metadata repository, and the DB2 instance must be installed and available prior to running the configuration tool.

**Note:** If the original installation was High Availability, these steps are prepopulated with the values from the original installation.

## Analytic Server

The migration tool uses the original configuration to determine the principal that serves as the Analytic Server administrator. This principal will be able to create, delete, and manage tenants. You can change the principal through the administration tool.

## Summary

The Summary step allows you to review the steps before completing migration.

## Post-migration

The Post-migration step provides information that the migration was successful. In the event of a failure, review the log file in the logs directory.



---

## Chapter 5. Making changes to your system

---

### Using the administration tool

The administration tool automates a number of otherwise manual steps for updating your configuration. The admin tool must have access to the HDFS environment.

You can launch the administration tool by executing the **admintool** script in {AS\_ROOT}\bin. By default, the tool detects whether to launch in GUI or console mode. You can force console mode by passing it the argument **-c**.

### Summary

The Summary panel is always shown prior to making any permanent changes to the system. When you complete a set of steps under an Administrative area, the final step takes you to the Summary panel. From here, you can select another Administrative area to modify or click Apply to apply the changes to the system.

### Hadoop cluster

The Hadoop cluster steps allow you to reconfigure Analytic Server for use with your Hadoop distribution.

**Note:** The administration tool will allow you to update the version of your Hadoop distribution, if the Hadoop vendor supports an in-place version update. If you are changing Hadoop vendors, you can migrate the settings for your existing Analytic Server installation to a new installation that is configured to use the new Hadoop distribution. For more information, see Chapter 4, “Migration,” on page 19.

#### Simple refresh

This is useful if your Hadoop distribution version has changed and you simply need to refresh the JAR files used by Analytic Server. In a High Availability installation, the Simple refresh will also find new cluster members and quarantine servers added to the collective.

**Important:** New members of the collective must be up and running in order to be enabled for use with Analytic Server.

#### Guided refresh

This is useful if other properties of your Hadoop cluster have changed. The screens that follow the Guided refresh are similar to those in the configuration tool.

### Hcatalog/Hive

The Hcatalog/Hive steps allow you to reconfigure Analytic Server for use with your Hive installation, or enable your existing Analytic Server installation with a new Hive installation.

If you have an existing Hive installation, there are two choices for updating your configuration.

#### Simple refresh

This is useful if your Hive version has changed and you simply need to refresh the JAR files used by Analytic Server.

#### Guided refresh

This is useful if other properties of your Hive installation have changed. The screens that follow the Guided refresh are similar to those in the configuration tool.

## JDBC

Copy the drivers for the databases you are using to the specified location and continue the administration steps.

## Application Server

The Application Server steps are dependent upon the type of installation.

- In a non-High Availability installation, you can configure secure communication (SSL/TLS) and set the ports on which Analytic Server listens.
- In a High Availability installation, the administration tool allows you to note any changes in your WebSphere® Application Server Liberty Network Deployment environment.

**Note:** In order to add or remove a quarantine server, you need to edit the **ae.cluster.quarantine.server.name** configuration property. For more information, see “Manually editing configuration properties.”

**Note:** The Administration Tool allows you to update the location of a Controller Server or a Replica of a Controller Server for a single Liberty ND environment. Analytic Server uses the Controller or Replica to communicate with and facilitate deployment to the Liberty ND environment. This feature cannot be used to change to a different Liberty ND environment or a different cluster. Use the migration tool to change the Liberty ND environment.

## Zookeeper

The administration tool allows you to modify the list of servers in your Zookeeper ensemble.

## Metadata repository

The Metadata repository steps allow you to modify your DB2 instance. If you change the schema, the new schema you specify must already exist in the database.

If you configured Analytic Server to use Derby as the metadata repository and wish to switch to DB2, use the migration tool.

## Analytic Server

The Analytic Server steps allow you to change the principal that serves as the Analytic Server administrator. This principal will be able to create, delete, and manage tenants.

## Post-administration

The Post-administration step provides information that the configuration changes were successful. In the event of a failure, review the log file in the logs directory.

---

## Manually editing configuration properties

There are some configuration properties that can be changed by manually editing configuration properties.

### config.properties

The `config.properties` file can be found in the `{AS_SERVER_ROOT}/configuration` directory. There are two ways to edit this file.

- The administration tool, when launched in GUI mode, allows you to make changes to this file.

- Otherwise, you can manually edit the file.

In a High Availability installation, you can push changes to `config.properties` through the administration tool. In a non-High Availability installation, you will need to restart Analytic Server.

#### **join.small.data.size**

The maximum amount of data (in bytes) that Analytic Server will attempt to join in a Map side algorithm.

#### **mapred.child.java.opts**

Controls the JVM heap sizes for map and reduce tasks that run in Hadoop. Set this value to as large as the nodes in the cluster can handle.

#### **mapred.job.reuse.jvm.num.tasks**

Reuse JVMs for map and reduce tasks. When the overhead for starting a JVM is high, increase the value of this property. For example, when the value is set to 1, then each map and each reduce task is run in a new JVM. When the value is set to 5, then a new JVM is created for every 5 tasks.

#### **mmr.threads**

The number of threads to use for in-memory map reduce (M3R) jobs.

**Note:** The properties prefixed by "mmr" control in-memory map reduce, and are important to any installation that makes heavy use of small files. You may need to experiment with these values to optimize your system.

#### **mmr.upper.bound.threshold**

This is the maximum amount of data that will be processed by M3R. Larger amounts of data will be processed by Hadoop.

#### **mmr.synthetic.block.size**

The default block size if the file is on the local file system.

#### **mmr.diskspill.threshold**

Intermediate M3R data is written on disk when it exceeds this threshold.

#### **mmr.diskspill.output.buffer**

The mapper output buffer. It must be less than `mmr.diskspill.threshold`.

#### **mmr.taskparallel.targets.threshold**

If the ratio of targets/cores is less than this threshold, the jobs are processed by M3R.

#### **node.max.jobs**

In a High Availability installation, the maximum number of Analytic Server jobs a cluster member can be running at any given time.

#### **ae.cluster.quarantine.server.name**

In a High Availability installation, the name of the server designated as the quarantine server.

### **jvm.options**

The `jvm.options` file can be found in the `{AS_SERVER_ROOT}` directory. Adding or modifying the following lines to `jvm.options` sets the heap size for jobs run on the server that hosts Analytic Server; that is, not Hadoop. This is important if running small (M3R) jobs, and you may need to experiment with these values to optimize your system.

```
-Xms512M  
-Xmx2048M
```

In a non-High Availability installation, you will need to restart Analytic Server. In a High Availability installation, there is no simple way to push changes to `jvm.options` to the collective; you need to manually update `jvm.options` on each cluster member.



---

## Chapter 6. Uninstalling

1. Run IBM Installation Manager (GUI or command line), select the option to uninstall IBM SPSS Analytic Server, and follow the prompts. IBM Installation Manager can also be run in the silent mode. For more information, see IBM Installation Manager documentation: [http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im\\_family\\_welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSDV2W/im_family_welcome.html).

**Note:** This operation does not remove any data from the Analytic Server folder on HDFS. To remove the entire Analytic Server home folder, remove the HDFS directory where the Analytic Server user artifacts (inputs and outputs) are stored. This is the directory specified as `root.folder.path` in `{AS_SERVER_ROOT}/configuration/config.properties`.

**Note:** This operation does not remove any DB2 schemas. Consult the DB2 documentation for information on manually removing schemas.



---

## Chapter 7. Troubleshooting

This section describes some common installation and configuration issues and how you can fix them.

### Issues specific to certain Hadoop distributions

#### Analytic Server fails to configure with BigInsights Adaptive MapReduce

Source the `<biginsights_home>/conf/biginsights-env.sh` file in the `.bashrc` file for the user performing the Analytic Server install. Alternatively, when prompted for the directory containing the Hadoop executable, specify the location `<biginsights_home>/HAManager/data/soam`.

### High availability clusters

#### `java.net.SocketTimeoutException: Read timed out`

Change the Liberty ND timeout environment variable as follows:

```
export LIBERTYND_READ_TIMEOUT=<milliseconds>
```

where `<milliseconds>` is the number of seconds to use for the JMX read timeout.

#### `java.io.IOException: CWWKX7202E: The timeout value 60 (seconds) for command ./server start expired`

Add the following to the Controller Server `server.xml`

```
<!-- Increase start and stop server timeout to accommodate slow hardware -->
<serverCommands startServerTimeout="120" stopServerTimeout="120"/>
```

#### `java.lang.OutOfMemoryError: Java heap space`

Add the following lines to `jvm.options` on every member of the HA cluster.

```
-Xms512M
-Xmx2048M
```

#### Zookeeper transaction data becomes unmanageable

Set the `autopurge.purgeInterval` parameter in `zoo.cfg` to 1 to enable automatic purges of the Zookeeper transaction log.

#### Analytic cluster service loses contact with Zookeeper

Review and modify the `tickTime`, `initLimit`, and `syncLimit` parameters in `zoo.cfg`. For example:

```
# The number of milliseconds of each tick
tickTime=2000
# The number of ticks that the initial
# synchronization phase can take
initLimit=30
# The number of ticks that can pass between
# sending a request and getting an acknowledgement
syncLimit=15
```

See the Zookeeper documentation for details: <https://zookeeper.apache.org/doc/r3.3.3/zookeeperAdmin.html>

#### Analytic Server jobs do not resume

There are two common situations in which Analytic Server jobs do not resume.

1. When an Analytic Server job fails because a cluster member fails, the job is normally restarted automatically on another cluster member. If the job does not resume, check to ensure there are at least 4 cluster members in the High Availability cluster.
2. When you quiesce a cluster member, then all of the Analytic Server jobs on that server are normally resumed on another cluster member. In order to ensure that the jobs resume, set `-Dcom.spss.ae.remoteclient.failover.threshold=100` and use remote mode.



---

## 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.  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement 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 want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group  
ATTN: Licensing  
200 W. Madison St.  
Chicago, IL; 60606  
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 is for planning purposes only. The information herein is subject to change before the products described become available.

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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_. All rights reserved.

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

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://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](http://www.ibm.com/legal/copytrade.shtml).

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.







Printed in USA