

IBM SPSS Analytic Server  
Version 1

*Installation Instructions*





---

# Contents

## IBM SPSS Analytic Server installation

### instructions . . . . . 1

Supported platforms. . . . . 1

Prerequisites . . . . . 1

Installing Analytic Server with the installer . . . . . 1

Manual installation on Hadoop . . . . . 2

    config.properties . . . . . 4

Adding JDBC drivers . . . . . 4

WebSphere Liberty configuration . . . . . 5

    Basic registry . . . . . 5

    LDAP registry configuration . . . . . 6

Authentication using a Kerberos security provider. . . 6

Enabling High Availability (HA) mode in Hadoop 2.0 7

Enabling HCatalog data sources . . . . . 8

Enabling Support for Essentials for R . . . . . 9

Getting users started . . . . . 10

Update Analytic Server . . . . . 10

Uninstall Analytic Server server . . . . . 11

Problem determination . . . . . 11

    Logging . . . . . 11

    Version information . . . . . 12

    Log collector . . . . . 12



---

# IBM SPSS Analytic Server installation instructions

---

## Supported platforms

Table 1. Environments

Product	Release	Processor	Word-size
Red Hat Enterprise Linux	6.x	x64	64-bit
Red Hat Enterprise Linux	5.x	x64	64-bit
SuSE Linux Enterprise Server	11	x64	64-bit
SuSE Linux Enterprise Server	10	x64	64-bit

---

## Prerequisites

- One of the following distributions:
  - Apache Hadoop 0.20.2
  - Apache Hadoop 1.0.3
  - IBM BigInsights 2.0
  - Cloudera CDH3
  - Cloudera CDH4 MRv1
  - Hortonworks Data Platform 1.x
- In order for the installer to copy the necessary files from the Hadoop distribution, you must have ssh and scp installed and be able to connect by using public key authentication without requiring a password to login. Otherwise, you must copy the files manually. In order to validate whether your ssh connection is working:
  1. Log in to the Hadoop Name Node Server by using public key authentication.
  2. Run the following, replacing {HADOOP\_SERVER} with your Hadoop server.

```
export HADOOP_SERVER_USERNAME='whoami'  
export HADOOP_SERVER={HADOOP_SERVER}  
ssh -q -o "BatchMode=yes" -o "ConnectTimeout 5"  
  $HADOOP_SERVER_USERNAME@$HADOOP_SERVER "echo 2>&1"  
  && echo HOST_SSH_OK || echo HOST_SSH_NOK
```
  3. If it returns "HOST\_SSH\_OK", it was successful. Otherwise, it is not configured correctly.

---

## Installing Analytic Server with the installer

The installer automates a number of otherwise manual steps. The installer must have access to the HDFS environment. To install Analytic Server:

1. Run the installer file `install.bin`.
  - a. Review the license agreement and accept the agreement to continue.
  - b. Choose an installation path.
  - c. Follow the prompts to install the files onto your system.
  - d. Configure Analytic Server options. The Administrator account name (DEFAULT: admin) is the user name with "super user" administrator privileges for the Analytic Server application through the console.

- e. Keystore Password. This keystore is used to store the encryption key that is used when encrypting database data source passwords that are stored in the Analytic Server database.
- f. Automatically Detect Hadoop Version

**Yes** Specify the Hadoop Namenode Server (DEFAULT: localhost). If Hadoop is on the same machine as the one on which you install this product, then use localhost as the server name. The installer then determines what version of Hadoop is installed and gathers the required files from the Hadoop distribution. Note: You must have ssh and scp installed and able to connect using public key authentication without requiring a password to log in.

**No**

- 1) Select the Hadoop server version from a list.
- 2) Specify the location of the required files from the Hadoop distribution, or manually copy the files later.
- 3) If your distribution is Cludera CDH3, specify the location of an Oracle Java Runtime Environment 1.6, or manually configure later.

- g. Configure Hadoop parameters

**Hadoop user name**

The UNIX user name that Analytic Server runs as.

**Password**

The password for the Hadoop user name.

**Analytic Server root directory**

This path is where Analytic Server sets up its workspace and stores executable files that are distributed to Hadoop nodes.

**Namenode host**

Name or IP address of the HDFS namenode host.

**Namenode port**

The port number for the namenode host.

**Job tracker**

Name or IP address of the Hadoop job tracker host.

**Job tracker port number**

The port number for the job tracker.

2. Configure WebSphere Liberty to register users. See “WebSphere Liberty configuration” on page 5 for more information.
3. Start the server by running `{AS_ROOT}/bin/start.sh`, where `{AS_ROOT}` refers to the folder where the Analytic Server is deployed; for example, `/opt/IBM/SPSS/AnalyticServer/1.0`.
4. Open a web browser and enter the address `http://localhost:8080`. This URL opens the login dialog for the Analytic Server console. Log in as the "super user" defined above. See the *IBM® SPSS® Analytic Server Administrator's Guide* for details on tenant management.

If the installer is unable to connect to the Hadoop Server or a valid Hadoop Version is not found, follow the manual steps.

---

## Manual installation on Hadoop

1. Run `install.bin`. This creates the Analytic Server database and sets the admin user credentials.
  - a. `{AS_ROOT}` refers to the folder where the Analytic Server is deployed; for example, `/opt/IBM/SPSS/AnalyticServer/1.0`.
  - b. `{AS_SERVER_ROOT}` refers to the location of the configuration, log, and server files; for example, `/opt/IBM/SPSS/AnalyticServer/1.0/ae_wlpserver/usr/servers/aeserver`.
  - c. `{AS_HOME}` refers to the folder on HDFS that is used by Analytic Server as a root folder.

2. Edit the configuration file {AS\_SERVER\_ROOT}/configuration/config.properties. See “config.properties” on page 4 for more information. Configure the Hadoop settings (**hdfs.namenode.ur1**, **hdfs.user**, **hdfs.password**) based on your system's specific Hadoop configuration.
3. Analytic Server runs on top of IBM WebSphere Liberty, which is configured by default to run on port 8080.
4. Copy the correct Hadoop JAR files into {AS\_SERVER\_ROOT}/apps/AE\_BOOT.war/WEB-INF/lib. These JAR files are based on your installation, and can be found on the Hadoop cluster namenode where Hadoop is installed; for example, in the folders /usr/lib/hadoop and /usr/lib/hadoop/lib.

- Apache Hadoop 0.20.2

hadoop-0.2\*-core.jar

- Apache Hadoop 1.0.3

commons-configuration-\*.jar

commons-io-\*.jar

commons-lang-\*.jar

hadoop-core-1.\*.jar

jackson-core-asl-\*.jar

jackson-mapper-asl-\*.jar

jersey-core-\*.jar

jersey-server-\*.jar

- Cloudera CDH3

aspectjrt-1.6.5.jar

aspectjtools-1.6.5.jar

core-3.1.1.jar

guava-r09\*.jar

hadoop-fairscheduler-0.20.2-\*.jar

hadoop-core-0.20.2-\*.jar

hadoop-tools-0.20.2-\*.jar

jackson-core-asl-1.5.2.jar

jackson-mapper-asl-1.5.2.jar

jasper-compiler-5.5.12.jar

jasper-runtime-5.5.12.jar

Cloudera CDH3 update 5 and later requires the additional JAR files

jersey-core-1.8.jar

jersey-json-1.8.jar

jersey-server-1.8.jar

- Cloudera 4.0.1

aspectjrt-1.6.5.jar

aspectjtools-1.6.5.jar

commons-io-2\*.jar

guava-\*.jar

hadoop-fairscheduler-2.\*.jar

hadoop-core-2.\*.jar

hadoop-tools-2.\*.jar

jackson-core-asl-1\*.jar

jackson-mapper-asl-1\*.jar

jasper-compiler-5\*.jar

jasper-runtime-5\*.jar

- Cloudera 4.1.\*

avro-\*.jar

avro-compiler-\*.jar

commons-cli-\*.jar

commons-configuration-\*.jar

commons-io-2\*.jar

commons-lang-\*.jar

guava-\*.jar

hadoop-auth-\*.jar

hadoop-common-2.\*.jar

hadoop-core-2.\*.jar

hadoop-hdfs-2.0.0-cdh4.1.2.jar

jackson-core-asl-\*.jar

jackson-mapper-asl-\*.jar

jersey-core-\*.jar

jersey-server-\*.jar

protobuf-java-\*.jar

slf4j-api-\*.jar

slf4j-log4j12-\*.jar

- IBM BigInsights 2.0

commons-configuration-\*.jar

commons-lang-\*.jar

commons-io-2\*.jar

hadoop-core-1.\*.jar

jackson-core-asl-\*.jar

jackson-mapper-asl-\*.jar

jersey-core-\*.jar

jersey-server-\*.jar

- Hortonworks Data Platform 1.x

commons-configuration-\*.jar

commons-lang-\*.jar

commons-io-2\*.jar

hadoop-core-1.\*.jar

jackson-core-asl-\*.jar

jackson-mapper-asl-\*.jar

jersey-core-\*.jar

jersey-server-\*.jar

5. Cloudera CDH3 requires an Oracle JVM 1.6. Update the setenv.sh script in {AE\_ROOT}/bin by updating JAVA\_HOME, changing the default path to an Oracle JVM.
6. If you are using Apache Hadoop 1.0.3 or later, replace commons-io-1.4.jar in {AS\_SERVER\_ROOT}/apps/AE\_BOOT.war/WEB-INF/lib with commons-io-2.x.jar.
7. If the Hadoop version is 2.0 or later then copy {AS\_ROOT}/cloudera/cloudera4/hadoopmapreduce\_1.1-1.0.0.0.jar to {AS\_SERVER\_ROOT}/apps/AE\_BOOT.war/WEB-INF/lib and remove hadoopmapreduce-1.0.0.0.jar.
8. Make sure that HDFS is running and accessible. Copy Analytic Server executable files to HDFS and initialize the Analytic Server workspace by running {AS\_ROOT}/bin/hdfsUpdate.sh.
9. Start the server by running {AS\_ROOT}/bin/run.sh
10. Enter http://localhost:8080 in your browser for the Analytic Server console. Log in as the admin.username defined in config.properties. See the *IBM SPSS Analytic Server Administrator's Guide* for details on tenant management.
11. To stop the server run {AS\_ROOT}/bin/stop.sh
12. The server log can be found at: {AS\_SERVER\_ROOT}/logs/

## config.properties

The config.properties file contains a number of Analytic Server settings; the following settings correspond to the settings that would be specified through the installer; confirm that the default settings are acceptable when manually installing the product.

### root.folder.path

This path is where the Analytic Server user artifacts (inputs and outputs) are stored.

### admin.consumer=ibm

The default administrator account tenant. Tenants provide a high-level division of roles, permissions, projects, and data sources.

### admin.username=admin

The default administrator account name. This user name has initial administrator privileges for the Analytic Server application through the console.

### hdfs.namenode.url=hdfs://hadoop:54310/user/hdpadding

The URL for the Hadoop namenode.

### hdfs.user=hdpadding

The Hadoop user name. The UNIX user name that Analytic Server runs as.

### hdfs.password=<ENTER PASSWORD>

The password for the user that is specified in the hdfs.user setting.

### hdfs.classpath.folder=/user/hdpadding/classpath

The Analytic Server root directory. This path is where Analytic Server sets up its workspace and stores executable files that are distributed to Hadoop nodes.

### mapred.job.tracker=hadoop:54311

Name or IP address and port number of the Hadoop job tracker host.

### encryption.keystore.password

This keystore is used to store the encryption key that is used when encrypting database data source passwords that are stored in the Analytic Server database.

---

## Adding JDBC drivers

In order to support database data sources, you must add the JDBC drivers to the Analytic Server.

1. Stop the Analytic Server by running {AS\_ROOT}/bin/stop.sh
2. Copy the required JDBC driver jars to {AS\_ROOT}/ae\_wlpsserver/usr/servers/aeserver/apps/AE\_BOOT.war/WEB-INF/lib



3. Update Analytic Server by running {AS\_ROOT}/bin/hdfsUpdate.sh
4. Start Analytic Server by running {AS\_ROOT}/bin/start.sh

Table 2. Supported databases

Database	Supported versions	JDBC driver jars	Vendor
DB2 for Linux, UNIX, and Windows	9.5, 9.7, 10.0	db2jcc.jar	IBM
DB2 z/OS	10	db2jcc.jar, db2_license_cisuz.jar	IBM
Teradata	13.1, 14	tdgssconfig.jar, terajdbc4.jar	Teradata
SQL Server	2012, 2008 R2	sqljdbc4.jar	Microsoft
Netezza	6.x, 7	nzjdbc.jar	IBM
Oracle	12g, 11g R2	ojdbc6.jar, orai18n.jar	Oracle

## WebSphere Liberty configuration

WebSphere Liberty Profile is a lightweight implementation of IBM WebSphere. Analytic Server can use WebSphere application security to authenticate users. This is set up in the server.xml for the server Analytic Server is deployed. To enable application security in Liberty, the appSecurity-1.0 feature must be included in the feature manager:

```
<featureManager onError="FAIL">
...
<feature>appSecurity-1.0</feature>
...
</featureManager>
```

To enable SSL on a server, the SSL feature must be included in the server.xml file:

```
<featureManager>
  <feature>ssl-1.0</feature>
</featureManager>
```

You can find detailed information about WebSphere security at: [ftp://ftp.software.ibm.com/software/webserver/appserv/library/v85/was85base\\_security.pdf](ftp://ftp.software.ibm.com/software/webserver/appserv/library/v85/was85base_security.pdf).

## Basic registry

The Basic Registry allows the administrator 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 securityUtil tool, which is located 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"/>
  <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>
</basicRegistry>
```

```

</group>
<group name="ADMIN">
  <member name="user1"/>
  <member name="admin"/>
</group>
</basicRegistry>

```

## LDAP registry configuration

The LDAP Registry provides the administrator a way 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.

---

## Authentication using a Kerberos security provider

Before you can configure Kerberos, you must obtain the following information from your Hadoop administrator:

1. Kerberos realm; for example, ASSSO.COM
2. Kerberos Key Distribution Center(KDC) host name; for example, kdc.assso.com
3. Name node Kerberos principal; for example, hdfs/namenode.assso.com@ASSSO.COM.
4. MapReduce node Kerberos principal; for example, mapred/jobtracker.assso.com@ASSSO.COM.

Then, you must configure the krb5.conf file at /etc/krb5.conf; for example:

```

[libdefaults]
default_realm = ASSSO.COM
default_tkt_enctypes = rc4-hmac des-cbc-md5
default_tgs_enctypes = rc4-hmac des-cbc-md5
dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
forwardable = yes

[realms]
ASSSO.COM = {
  kdc = kdc.assso.com:88
  default_domain = assso.com
}

[dmain_realm]
.assso.com = ASSSO.COM
assso.com = ASSSO.COM

```

After you create the krb5.conf file, modify the config.properties file as follows:

1. In the Analytic Server modules section:

- Add the `hdfsauth` and `kerberossecurityprovider` modules
  - Remove the `wssecurityprovider` module.
2. Add the following configuration properties:
 

```
#Kerberos authentication parameters
hadoop.security.authentication=kerberos
dfs.namenode.kerberos.principal=hdfs/namenode.asso.com@ASSSO.COM
mapreduce.jobtracker.kerberos.principal=mapred/jobtracker.asso.com@ASSSO.COM
java.security.krb5.conf=/etc/krb5.conf
```

where

**hadoop.security.authentication**  
Hadoop security authentication. Specify `kerberos` to enable Kerberos security provider.

**dfs.namenode.kerberos.principal**  
Kerberos principal that is used for the keytab file, which is used start the Name node.

**mapreduce.jobtracker.kerberos.principal**  
Kerberos principal that is used for the keytab file, which is used start the job tracker.

**java.security.krb5.conf**  
Kerberos configuration file location.
  3. Configure the Liberty LDAP user repository in `service.xml`. Refer to “WebSphere Liberty configuration” on page 5. All the users that are specified in the Liberty user repository must be matched with Kerberos user accounts and the same LDAP setting must be used in Kerberos server.
  4. By default Analytic Server uses the `.temp` directory under the user’s home directory as a temporary directory, but if you want to configure the temporary directory in a different location, then follow these instructions.
    - a. Edit `config.properties` and uncomment the following configuration setting.
 

```
#as.temp.folder=/.temp
```

Change the setting as needed to the absolute path of the temporary directory. No changes are needed if the `/.temp` directory is used.
    - b. Change permissions to this folder to allow all users access to this folder; for example, `hadoop fs -chmod 777 /.temp`.
  5. Give read permission to the `classpath` and `configuration` folders for all Kerberos users.
    - a. Open the `config.properties` file and note the settings for the parameters **`hdfs.classpath.folder`** and **`component.framework.bin.path`**.
    - b. Set the read permissions as follows.
 

If **`hdfs.classpath.folder`**=`/user/hdpadmin/classpath` and **`component.framework.bin.path`**=`/user/hdpadmin/configuration`, then run

```
hadoop fs -chmod -R 755 /user/hdpadmin
hadoop fs -chmod -R 755 /user/hdpadmin/classpath
hadoop fs -chmod -R 755 /user/hdpadmin/configuration
```

---

## Enabling High Availability (HA) mode in Hadoop 2.0

Analytic Server supports running Hadoop custom modes, like High Availability (HA), by providing the cluster client API configuration into the Analytic Server configuration folder. Providing the Hadoop cluster client API configuration to Analytic Server is mandatory only if there are custom settings, like HA. These settings are required so that the Analytic Server API can balance the requests if a cluster namenode fails.

In order to configure the custom mode after a successful Analytic Server installation:

1. Obtain the cluster client configuration files (`hdfs-site.xml` and `core-sites.xml`), usually located on a namenode machine in `/etc/hadoop/conf`.

2. Copy these files into {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/configuration/hadoop-conf.
3. If the Hadoop cluster is configured in HA mode, make sure the Analytic Server configuration property (in the config.properties file) hdfs.namenode.url is pointing to the HDFS service name and /user/Username; for example, hdfs://nameservice1/user/hdpadmin.
4. If you have set up HA for the job tracker service, then update the **mapred.job.tracker** property in config.properties to point to the HDFS service name.
5. Update the Hadoop file system by executing the command:  

```
{AS_ROOT}/bin/hdfsUpdate.sh
```
6. Start the Analytic Server by executing the command:  

```
{AS_ROOT}/bin/run.sh
```

---

## Enabling HCatalog data sources

In order to configure Analytic Server for use with HCatalog databases after a successful Analytic Server installation:

1. Add a hcataloginput entry to the list of Analytic Server modules in the {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/configuration/config.properties file. For example:

```
ae.modules=securityprovidermanager,\
  localClient,\
  componentframework,\
  ...
  hcataloginput
```

2. Uncomment or add the following lines to config.properties

```
hive.metastore.local=false
hive.metastore.uris=thrift://hostname:portnum
where
```

### hostname

The name of the machine that hosts the Thrift server

### portnum

The port number that is used in the HCatalog installation script

3. Make sure that the following files are available in the {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/apps/AE\_BOOT.war/WEB-INF/lib directory and also copy these files to the HDFS directory /user/{ae\_admin}/classpath.

The following JAR files are for HCatalog 0.4.0 and Hive 0.9.0. You must harvest the corresponding HCatalog, Hive, and dependent JAR files as appropriate for other versions.

- The following file can be copied from the HCatalog server installation.

```
hcatalog-0.4.0.jar
```

- The following files can be copied from the Hive server installation.

```
hive-exec-0.9.0.jar
hive-metastore-0.9.0.jar
libfb303-0.7.0.jar
slf4j-api-1.6.1.jar
slf4j-log4j12-1.6.1.jar
```

- The following files can be copied from the Hadoop 1.X installation or from <http://jackson.codehaus.org/>.

```
jackson-core-asl-1.8.8.jar
jackson-mapper-asl-1.8.8.jar
```

The following JAR files are for HCatalog 0.5.0.

- The following files can be copied from <http://code.google.com/p/guava-libraries/wiki/Release13>.
- ```
guava-13.0.1.jar
```

4. If you plan to use Hive complex types (maps, arrays, structures) copy these JAR files to your Hive /usr/local/hive/hive-0.9.0/auxlib/ directory:

```
hcatalog-0.4.0.jar
jackson-core-as1-1.8.8.jar
jackson-mapper-as1-1.8.8.jar
```

5. The HCatalog data sources based on the compressed sequence files require Hadoop native libraries to be available on the Analytic Server machine. Copy the Hadoop native libraries found at {HADOOP}/lib/native/Linux-amd64-64 to a directory on the Analytic Server machine and edit the **LIB\_PATH** variable in {AS\_ROOT}/bin/start.sh to include the path to that directory. For example:

```
export LIB_PATH=$AE_BASE/ae_wlpserver/usr/servers/aeserver/configuration/lib_32:
$AE_BASE/ae_wlpserver/usr/servers/aeserver/configuration/lib_64:
<hadoop_native_libraries_directory>:
```

**Note:** If your Analytic Server installation already has newer versions of these JAR files, you should not copy the older versions. For example, if the Analytic Server installer has already copied jackson-core-as1-1.8.0 from a newer Hadoop distribution, you should not copy the jackson-core-as1.1.7.3 JAR file.

---

## Enabling Support for Essentials for R

Analytic Server supports scoring R models and running R scripts.

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

1. Install R Engine on the server that hosts Analytic Server by using the following steps:

```
mkdir /home/hdadmin/APPS/R
cd /home/hdadmin/APPS/R
wget http://cran.r-project.org/src/base/R-2/R-2.15.2.tar.gz
tar -xzf R-2.15.2.tar.gz
cd R-2.15.2
./configure --enable-R-shlib
make
```

2. Install Essentials for R on the server that hosts Analytic Server by running the installer file install.bin, following the instructions on the screen. The installer:
  - a. Updates the R Engine installation on Analytic Server; it adds the "R plug-in", and
  - b. Updates the {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/configuration/ext\_64/bin directory to add a native library and a configuration file to the pasw.rstats module.

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/hdadmin/APPS/R/R-2.15.2
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
```

- 2) Copy the archive R.zip to the {AS\_ROOT}/ae\_wlpserver/usr/servers/aeserver/configuration/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, and 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, simply specify the directory in which R is installed and skip the step of specifying the location of `../ext_64/bin`.
- c. Run `{AS_ROOT}/bin/hdfsUpdate.sh` to propagate the changes to HDFS.

**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 should read as follows.

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

**Note:** 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.

You must also install Essentials for R on the machine that hosts SPSS Modeler. See the *IBM SPSS Modeler updates and extensions for IBM SPSS Analytic Server installation instructions* for details.

---

## Getting users started

Tell users to navigate to `http://<host>:<port>/admin/<tenant>` and enter their username and password to log on to the Analytic Server console.

**<host>**

The address of the Analytic Server host

**<port>**

The port that Analytic Server is listening on

**<tenant>**

In a multi-tenant environment, the tenant you belong to. In a single-tenant environment, the default tenant is **ibm**.

In order to access IBM SPSS Analytic Catalyst, navigate to `http://<host>:<port>/catalyst.html` and enter their username and password to log on.

---

## Update Analytic Server

When you reinstall with the installer, you are able to update the Analytic Server options, keystore Password, and Configure Hadoop parameters.

### Updating Hadoop Server Version

From time to time it may become necessary to update the version of the Hadoop server that Analytic Server is running against, for example upgrading Cloudera 3.x to Cloudera 4.x. When this occurs, it is necessary to update the Analytic Server with the client jars from new Hadoop server version.

**Note:** These steps assume that no Hadoop related configuration properties have changed, such as: Namenode host/port, Job tracker host/port or the Hadoop username/password.

1. Stop the Analytic Server by executing the following command.

```
{AS_ROOT}/bin/stop.sh
```

2. Follow the vendor's instructions for upgrading the Hadoop server, and ensure it is running.
3. Copy the required client jars into a directory local to the Analytic Server; this directory will be referred to as `{HADOOP_SRC_FOLDER}`. The list of required jars depends upon the new version of the Hadoop server and can be found in "Manual installation on Hadoop" on page 2.

4. Execute the `install_hadoop_files` shell script found in `{AS_ROOT}/bin/hadoop_setup`, providing `{HADOOP_SRC_FOLDER}` as an argument. For example: `./install_hadoop_files.sh /opt/hadoop_src`.
  - a. Select if you would like to remove any installed Hadoop files (This is recommended).
  - b. Select the version of Hadoop you want to use.
5. Update the Hadoop file system with Analytic Server by executing the following command.  
`{AS_ROOT}/bin/hdfsUpdate.sh`
6. Restart the Analytic Server by executing the following command.  
`{AS_ROOT}/bin/run.sh`

---

## Uninstall Analytic Server server

1. Run `{AS_ROOT}/_uninstall/uninstaller.bin`.

**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:** If the uninstaller is not able to launch because of a `Java OutOfMemoryException` then the maximum heap size should be increased. The `lax.nl.java.option.java.heap.size.max` property in `{AS_ROOT}/_uninstall/uninstaller.bin.lax` defines the maximum heap size in bytes for the installer that will be invoked. This number is always specified in bytes, not in kilobytes or megabytes, and is analogous to the VM parameter `-mx` or `Xmx`. The default is 50331648 (48 MB). We suggest changing this to 134217728 (128 MB) and then launching the uninstaller.

---

## Problem determination

Analytic Server provides several helpful tools for problem determination.

### Logging

Analytic Server creates customer log files and trace files that are helpful for diagnosing problems. With the default Liberty installation, you can find the log files in the `{LIBERTY_PROFILE_FOLDER}/logs` directory.

The default logging configuration produces two log files that roll over on a daily basis.

#### **ae.log**

This file contains the high-level summary of informational warning and error messages. Check this file first when server errors occur that cannot be resolved by using the error message that is displayed in the User interface.

#### **ae\_trace.log**

This file contains all the entries from `ae.log`, but adds more information that is primarily targeted to IBM support and development for debugging purposes.

Analytic Server uses Apache LOG4J as its underlying logging facility. Using LOG4J, the logging can be dynamically adjusted by editing the `{AS_SERVER_ROOT}/configuration/log4j.xml` configuration file. You may be asked to do this by Support to help diagnose problems, or you may want to modify this to limit the number of log files kept around. Changes to the file are detected automatically within a few seconds so the Analytic Server does not need restarted.

For more information about `log4j` and the configuration file, see documentation at the official Apache website at <http://logging.apache.org/log4j/>.

## Version information

You can determine what version of Analytic Server is installed by checking the {AS\_ROOT}/properties/version folder. The following files contain version information.

### **analytic\_engine.swtag**

Contains detailed product information.

### **version.txt**

Version and build number for the installed product.

## Log collector

When problems cannot be resolved by directly reviewing the log files, you can bundle all the logs and send them to IBM support. There is a utility that is provided to make collecting all the necessary data simpler.

Using a command shell, run the following commands

```
cd <AS_ROOT>/tools/support/logcollector  
run >sh ./logcollector.sh
```

These commands create a compressed file under <AS\_ROOT>/tools/support/logcollector. The compressed file contains all the log files and product version information.







Printed in USA