IBM SPSS Analytic Server
Version 1.0.1

Installation Instructions

IBM

Product Information

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

Contents

IBM SPSS Analytic Server installation instructions

Prerequisites

- Review the Analytic Server system requirements by following these instructions.
  1. Go to the Detailed system requirements for a specific product page.
  2. Type Analytic Server as the product name and click the search button.
  3. Select the version and scope of report, then click **Submit**.
- 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 provide a username and password, or copy the files manually.

Note: It is not recommended to run multiple instances of Analytic Server on a single machine.

Ports

Analytic Server uses the following ports by default.
- 8080 for http
- 8443 for https
- 9114 for BD
- 9080 for Liberty

Installing Analytic Server with the installer

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

Note: If you want to migrate settings from an existing environment, see "Migrating from a previous version of Analytic Server" on page 7.

To install Analytic Server:
1. The installer should be run by the user that will run the Analytic Server process. This user should have access to the Hadoop environment.
2. 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

      1) 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. If you are able to connect using public key authentication without requiring a password to log in, you do not need to select **Login Required**. Otherwise, provide the username and password for an account that has permissions to copy files from your Hadoop distribution to the Analytic Server installation.

Note: On PureData systems, you must always select **Login Required** and provide credentials.

2) If your distribution is Cloudera CDH3 or MapR, specify the location of an Oracle Java Runtime Environment 1.6, or manually configure later.

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 Cloudera CDH3 or MapR, specify the location of an Oracle Java Runtime Environment 1.6, or manually configure later.

g. Configure Hadoop parameters

Hadoop user name
> The 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.

3. Configure WebSphere Liberty to register users. See the *IBM® SPSS® Analytic Server Administrator's Guide* for more information.

4. If you are running SUSE 10 or RedHat 5, update the following libraries:
   - On RedHat 5, copy `libstdc++.so.6` from `gcc-4.3.4/lib64` to `{AS_ROOT}/ae_wlpserver/usr/servers/aeserver/configuration/lib_64`
   - On SUSE 10, copy `libgomp.so.1` and `libstdc++.so.6` from `gcc-4.3.4/lib64` to `{AS_ROOT}/ae_wlpserver/usr/servers/aeserver/configuration/lib_64`

5. 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/{version}`.

6. 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. The installer should be run by the user that will run the Analytic Server process. This user should have access to the Hadoop environment.
2. 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/{version}`.
   b. {AS_SERVER_ROOT} refers to the location of the configuration, log, and server files; for example, `/opt/IBM/SPSS/AnalyticServer/{version}/ae_wlpserver/usr/servers/aeserver`.
   c. {AS_HOME} refers to the folder on HDFS that is used by Analytic Server as a root folder.
3. Edit the configuration file `{AS_SERVER_ROOT}/configuration/config.properties`. See "config.properties" on page 5 for more information. Configure the Hadoop settings (**hdfs.namenode.url**, **hdfs.user**, **hdfs.password**) based on your system's specific Hadoop configuration.
4. Analytic Server runs on top of IBM WebSphere Liberty, which is configured by default to run on port 8080.
5. 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.*

     ```
     hadoop-0.2*-core.jar
     ```

   - Apache Hadoop 1.*

     ```
     commons-configuration-*.jar          jackson-mapper-asl-*.jar
     commons-io-*.jar                     jersey-core-*.jar
     commons-lang-*.jar                   jersey-server-*.jar
     hadoop-core-1.*.jar
     jackson-core-asl-*.jar
     ```

     Note: 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.*.jar`.

   - Cloudera CDH3

     ```
     aspectjrt-1.6.5.jar                  hadoop-tools-0.20.2-*.jar
     aspectjtools-1.6.5.jar               jackson-core-asl-1.5.2.jar
     core-3.1.1.jar                       jackson-mapper-asl-1.5.2.jar
     guava-r09*.jar                       jasper-compiler-5.5.12.jar
     hadoop-fairscheduler-0.20.2-*.jar    jasper-runtime-5.5.12.jar
     hadoop-core-0.20.2-*.jar
     ```

     Note: Additionally remove hadoopmapreduce1.1-*.jar, hcatalogstoragehandler1.1-*.jar, and accumulo-hive-storage-handler*.jar from `{AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib` and copy `{AS_ROOT}/cloudera/cloudera3/*.jar` to `{AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib`.

     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
     ```

     Note: Cloudera CDH3 requires Oracle Java Runtime Environment 1.6. Update the setenv.sh script in {AS_ROOT}/bin by updating JAVA_HOME, changing the default path to an Oracle JVM.

   - Cloudera CDH4

     ```
     aspectjrt-1.6.5.jar                  hadoop-tools-2.*.jar
     aspectjtools-1.6.5.jar               jackson-core-asl-1*.jar
     commons-io-2*.jar                    jackson-mapper-asl-1*.jar
     guava-*.jar                          jasper-compiler-5*.jar
     hadoop-fairscheduler-2.*.jar         jasper-runtime-5*.jar
     hadoop-core-2.*.jar
     ```

- Cloudera CDH4.1.*

```
avro-*.jar                        hadoop-core-2.*.jar
avro-compiler-*.jar               hadoop-hdfs-2.0.0-cdh4.1.2.jar
commons-cli-*.jar                 jackson-core-asl-*.jar
commons-configuration-*.jar       jackson-mapper-asl-*.jar
commons-io-2*.jar                 jersey-core-*.jar
commons-lang-*.jar                jersey-server-*.jar
guava-*.jar                       protobuf-java-*.jar
hadoop-auth-*.jar                 slf4j-api-*.jar
hadoop-common-2.*.jar             slf4j-log4j12-*.jar
```

  Note: Additionally remove hadoopmapreduce-*.jar, hcatalogstoragehandler-*.jar, and accumulo-hive-storage-handler*.jar from {AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib and copy {AS_ROOT}/cloudera/cloudera4/*.jar to {AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib.

  Note: Additionally copy commons-io-1.4.jar in {AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib to {AS_ROOT}/Apache_0.20_SRC.

- Hortonworks Data Platform 1.*

```
commons-configuration-*.jar       jackson-core-asl-*.jar
commons-lang-*.jar                jackson-mapper-asl-*.jar
commons-io-2*.jar                 jersey-core-*.jar
hadoop-core-1.*.jar               jersey-server-*.jar
```

- IBM BigInsights 2.* (Basic or Enterprise Edition); IBM PureData System for Hadoop H1001

```
commons-configuration-*.jar       jackson-core-asl-*.jar
commons-lang-*.jar                jackson-mapper-asl-*.jar
commons-io-2*.jar                 jersey-core-*.jar
hadoop-core-1.*.jar               jersey-server-*.jar
```

- IBM BigInsights 2.1 with Platform Symphony 6.1.x

```
activation.jar              jackson-core-asl-*.jar
axis-1.4.2.jar              jackson-mapper-asl-*.jar
commons-configuration-*.jar     JavaSoamApi.jar
commons-discovery.jar       jersey-core-*.jar
commons-ego.jar               jersey-server-*.jar
commons-io-2*.jar                 pmr-hadoop-1.1.1.jar
commons-lang-*.jar                pmr-hadoop-patch-1.1.1.jar
commons-soam.jar            wsdl4j.jar
geronimo-jaxrpc_1.1_spec-2.1.jar
```

  For Platform Symphony, take the following additional steps:

  a. Copy the following library files to {AS_SERVER_ROOT}/configuration/lib_64

```
libVEMApiCommon.so
libz.so.1
sec_ego_default.so
libjnativesoamapi_6.1.0_64.so
libcom_platform_log4cxx_097_3.so.9
libsoambase.so
libcom_platform_ACE_544_03.so.5.4.4
libcom_platform_ACE_544_03_SSL.so.5.4.4
libicui18n.so.32
libicuuc.so.32
libicudata.so.32
libsoam_resources.so
libxml2.so.2
libsec.so
libvem.so
```

  b. Create the directory {AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib/conf, then copy {AS_ROOT}/bin/hadoop_setup/symphony_6.1.x_pmr-site.xml to {AS_SERVER_ROOT}/apps/AE_BOOT.war/WEB-INF/lib/conf/pmr-site.xml.

  c. Modify {AS_SERVER_ROOT}/configuration/config.properties, to add :apps/AE_BOOT.war/WEB-INF/lib/conf to the end of the line that starts with **ae.jvm.settings**.

d. Copy the {AS_ROOT}/bin/hadoop_setup/AE_HADOOP_LIB/conf folder into {AS_SERVER_ROOT}/configuration/

e. Modify the {AS_ROOT}/bin/start.sh file, to add the line export EGO_CONFDIR=$AE_BASE/ae_wlpserver/usr/servers/aeserver/configuration/conf following the line export AE_BASE='pwd'.

- MapR 3.* (M3, M5)

```
amazon-s3.jar
google-collect-1.0.jar
hadoop-0.20.2-dev-core.jar
jackson-core-asl-1.5.2.jar
jackson-mapper-asl-1.5.2.jar
maprfs-*.jar
zookeeper-3.3.6.jar
mapr-clusters.conf
```

Note: MapR requires Oracle Java Runtime Environment 1.6. Update the setenv.sh script in {AS_ROOT}/bin by updating JAVA_HOME, changing the default path to a Sun JVM.

For MapR, take the following additional steps.

a. Copy all files from /opt/mapr/hadoop/hadoop-0.20.2/lib/native/Linux-amd64-64/*.* to {AS_SERVER_ROOT}/configuration/lib64. Copy the source of any symbolic links in this directory. You can determine which file are symbolic linka by executing ls -la in the directory.

b. Copy the /opt/mapr/conf/mapr-clusters.conf file to {AS_SERVER_ROOT}/configuration/conf

c. Update the setenv.sh script in {AS_ROOT}/bin to add export MAPR_HOME={AS_SERVER_ROOT}/configuration/

d. Update config.properties to add the following lines.

```
hdfs.namenode.url=maprfs:///
mapred.job.tracker=maprfs:///
```

6. If you are running SUSE 10 or RedHat 5, update the following libraries:

- On RedHat 5, copy libstdc++.so.6 from gcc-4.3.4/lib64 to {AS_ROOT}/ae_wlpserver/usr/servers/aeserver/configuration/lib_64

- On SUSE 10, copy libgomp.so.1 and libstdc++.so.6 from gcc-4.3.4/lib64 to {AS_ROOT}/ae_wlpserver/usr/servers/aeserver/configuration/lib_64

7. If you plan to authenticate users through a Kerberos security provider, follow the instructions for configuring Kerberos in the *IBM SPSS Analytic Server Administrator's Guide*.

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/hdpadmin`
   The URL for the Hadoop namenode.

hdfs.user=hdpadmin
   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/hdpadmin/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.

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 3.

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`

Migrating from a previous version of Analytic Server

If you install a new version of Analytic Server on the same machine as a previous version, you can migrate settings to the new version.

Note: Migration is only supported the first time installing a new version of Analytic Server. If you perform a reinstall to the same location then it is treated as an update install.

Note: Manual configuration outside the installer cannot be migrated. This will require manually configuring the new version of Analytic Server.

Steps to migrate:

1. Stop the previous version of Analytic Server by executing the following command.

   `{AS_ROOT_previousVersion}/bin/stop.sh`

2. Launch the installer and follow the installer steps.

3. When prompted for migrating, select `Yes`.

4. Enter the directory where the previous version of Analytic Server was installed.

5. The installer will attempt to collect any previous settings.

6. You will have the option to modify these if they are found. If the keystore is found you will not have the option.

7. If any settings are not found then you will be prompted to fill these out.

Return to a previous version of Analytic Server

1. Stop the current version of Analytic Server by executing the following command.

   `{AS_ROOT_currentVersion}/bin/stop.sh`

2. Uninstall the current version of Analytic Server by executing the following command.

   `{AS_ROOT_currentVersion}/bin/hdfsUninstall.sh`

3. Update the Hadoop file system with the previous version of Analytic Server by executing the following command.

   `{AS_ROOT_previousVersion}/bin/hdfsUpdate.sh`

4. Restart the Analytic Server by executing the following command.

   `{AS_ROOT_previousVersion}/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.

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 wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
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.

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

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.

IBM ®