

HadoopJavaVersions

Deprecated since 2.7
Please check [Hadoop Java Versions](#) instead.

Hadoop Java Versions

Version 2.7 and later of Apache Hadoop requires Java 7. It is built and tested on both OpenJDK and Oracle (HotSpot)'s JDK/JRE.
Earlier versions (2.6 and earlier) support Java 6.

Tested JDK

Here are the known JDKs in use or which have been tested:

Version	Status	Reported By
oracle 1.7.0_15	Good	Cloudera
oracle 1.7.0_21	Good (4)	Hortonworks
oracle 1.7.0_45	Good	Pivotal
openjdk 1.7.0_09-icedtea	Good (5)	Hortonworks
oracle 1.6.0_16	Avoid (1)	Cloudera
oracle 1.6.0_18	Avoid	Many
oracle 1.6.0_19	Avoid	Many
oracle 1.6.0_20	Good (2)	LinkedIn, Cloudera
oracle 1.6.0_21	Good (2)	Yahoo!, Cloudera
oracle 1.6.0_24	Good	Cloudera
oracle 1.6.0_26	Good(2)	Hortonworks, Cloudera
oracle 1.6.0_28	Good	LinkedIn
oracle 1.6.0_31	Good(3, 4)	Cloudera, Hortonworks

1. Hadoop works well with update 16 however there is a bug in JDK versions before update 19 that has been seen on HBase. See [HBASE-4367](#) for details.
2. If the grid is running in secure mode with MIT Kerberos 1.8 and higher, the Java version should be 1.6.0_27 or higher in order to avoid [Java bug 6979329](#).
3. Hortonworks has certified JDK 1.6.0_31 under RHEL5/CentOS5, RHEL6/CentOS6, and SLES11, with Hadoop 1.x, HBase, Pig, Hive, HCatalog, Oozie, Sqoop, and Ambari.
4. Hortonworks has certified JDK 1.6.0_31 and Oracle 1.7.0.21 under RHEL5/CentOS5, RHEL6/CentOS6, and SLES11, with Hadoop 2.2.0, HBase 0.96, Pig, Hive, HCatalog, Oozie, Sqoop, and Ambari.
5. Hortonworks has certified openjdk 1.7.0_09-icedtea on RHEL6 with Hadoop 2.2.0, HBase 0.96, Pig, Hive, HCatalog, Oozie, Sqoop, and Ambari.

Compressed Object pointers and Java 6

The Sun JVM has 32-bit and 64-bit modes. In a large cluster the [NameNode](#) and [JobTracker](#) need to run in 64-bit mode to keep all their data structures in memory. The workers can be set up for either 32-bit or 64-bit operation, depending upon preferences and how much memory the individual tasks need.

Using the Compressed Object References JVM feature (-XX:+UseCompressedOops) reduces memory consumed and increases performance on 64 bit Sun JVMs. This feature was first introduced in 1.6.0_14 but problems were been reported with its use on versions prior to 1.6.0_20. Several have reported success using it on 1.6.0_21 and above. It is the default in 1.6.0_24 and above on 64 bit JVMs, and appears now to be stable.

Useful tips for discovering and inspecting Sun JVM configuration flags are in the following blog post: [inspecting-hotspot-jvm-options](#)

OpenJDK 7

OpenJDK has been used to qualify Hadoop 2.2 and the rest of the Hortonworks bundle on RHEL6. No problems were noted.

Oracle JRockit

Oracle's JRockit JVM is not the same as the Sun JVM: it has very different heap and memory management behavior. Hadoop has been used on JRockit, though not at "production" scale.

1. [Problems spawning jobs](#)
2. [One of the tests doesn't like JRockit](#)
3. [Log4J configuration issues](#)

IBM JDK

Anyone who has information about compatibility of Hadoop 2.x with IBM JDK is encouraged to edit this wiki.

An older version of Hadoop (0.20.2) has been tested comprehensively and works with IBM Java 6 SR 8. IBM Java can be [downloaded here](#).

A request for help from JVM/JDK developers

We would strongly encourage anyone who produces a JVM/JDK to test compiling and running Hadoop with it. It makes for a fantastic performance and stress test. As Hadoop is becoming a key back-end datacenter application, good Hadoop support matters.