

Hadoop MapReduce Releases

Table of contents

1 Download.....	3
2 News.....	3
2.1 23 May, 2012: Release 2.0.0-alpha available	3
2.2 16 May, 2012: Release 1.0.3 available	3
2.3 3 Apr, 2012: Release 1.0.2 available	3
2.4 10 Mar, 2012: Release 1.0.1 available	4
2.5 27 Feb, 2012: release 0.23.1 available	4
2.6 27 December, 2011: release 1.0.0 available	4
2.7 10 December, 2011: release 0.22.0 available	5
2.8 17 Oct, 2011: release 0.20.205.0 available	5
2.9 5 Sep, 2011: release 0.20.204.0 available	5
2.10 11 May, 2011: release 0.20.203.0 available	5
2.11 23 August, 2010: release 0.21.0 available	6
2.12 22 April, 2009: release 0.20.0 available	6
2.13 24 February, 2009: release 0.19.1 available	6
2.14 29 January, 2009: release 0.18.3 available	6
2.15 21 November, 2008: release 0.19.0 available	6
2.16 3 November, 2008: release 0.18.2 available	6
2.17 17 September, 2008: release 0.18.1 available	7
2.18 22 August, 2008: release 0.18.0 available	7
2.19 19 August, 2008: release 0.17.2 available	7
2.20 23 June, 2008: release 0.17.1 available	7
2.21 20 May, 2008: release 0.17.0 available	7

2.22	5 May, 2008: release 0.16.4 available	7
2.23	16 April, 2008: release 0.16.3 available	7
2.24	2 April, 2008: release 0.16.2 available	7
2.25	13 March, 2008: release 0.16.1 available	8
2.26	7 February, 2008: release 0.16.0 available	8
2.27	18 January, 2008: release 0.15.3 available	8
2.28	2 January, 2008: release 0.15.2 available	8
2.29	27 November, 2007: release 0.15.1 available	8
2.30	26 November, 2007: release 0.14.4 available	8
2.31	29 October 2007: release 0.15.0 available	8
2.32	19 October, 2007: release 0.14.3 available	9
2.33	4 September, 2007: release 0.14.1 available	9

1 Download

- **1.0.X** - current stable version, 1.0 release
- **1.1.X** - current beta version, 1.1 release
- **2.X.X** - current alpha version
- **0.22.X** - does not include security
- **0.20.203.X** - legacy stable version
- **0.20.X** - legacy version

Releases may be downloaded from Apache mirrors.

[Download a release now!](#)

On the mirror, all recent releases are available.

Third parties may distribute products that include Apache Hadoop and derived works, under the Apache License. Some of these are listed on the [Distributions wiki page](#).

2 News

2.1 23 May, 2012: Release 2.0.0-alpha available

This is the first (alpha) version in the hadoop-2.x series.

This delivers significant major features over the currently stable hadoop-1.x series including:

- HDFS HA for NameNode (manual failover)
- YARN aka NextGen MapReduce
- HDFS Federation
- Performance
- Wire-compatibility for both HDFS and YARN/MapReduce (using protobufs)

Please see the [Hadoop 2.0.0-alpha Release Notes](#) for details.

2.2 16 May, 2012: Release 1.0.3 available

This is a bug fix release for version 1.0.

Bug fixes and feature enhancements in this minor release include:

- 4 patches in support of non-Oracle JDKs
- several patches to clean up error handling and log messages
- various production issue fixes

Please see the [Hadoop 1.0.3 Release Notes](#) for details.

2.3 3 Apr, 2012: Release 1.0.2 available

This is a bug fix release for version 1.0.

Bug fixes and feature enhancements in this minor release include:

- Snappy compressor/decompressor is available
- Occasional deadlock in metrics serving thread fixed
- 64-bit secure datanodes failed to start, now fixed
- Changed package names for 64-bit rpm/debs to use ".x86_64." instead of ".amd64."

Please see the [Hadoop 1.0.2 Release Notes](#) for details.

2.4 10 Mar, 2012: Release 1.0.1 available

This is a bug fix release for version 1.0. This release is now considered stable, replacing the long-standing 0.20.203.

Bug fixes in this minor release include:

- Added hadoop-client and hadoop-minicluster artifacts for ease of client install and testing
- Support run-as-user in non-secure mode
- Better compatibility with Ganglia, HBase, and Sqoop

Please see the [Hadoop 1.0.1 Release Notes](#) for details.

2.5 27 Feb, 2012: release 0.23.1 available

This is the second alpha version of the hadoop-0.23 major release after the first alpha 0.23.0. This release has significant improvements compared to 0.23.0 but should still be considered as alpha-quality and not for production use.

hadoop-0.23.1 contains several major advances from 0.23.0:

- Lots of bug fixes and improvements in both HDFS and MapReduce
- Major performance work to make this release either match or exceed performance of Hadoop-1 in most aspects of both HDFS and MapReduce.
- Several downstream projects like HBase, Pig, Oozie, Hive etc. are better integrated with this release

See the [Hadoop 0.23.1 Release Notes](#) for details.

2.6 27 December, 2011: release 1.0.0 available

After six years of gestation, Hadoop reaches 1.0.0! This release is from the 0.20-security code line, and includes support for:

- security
- HBase (append/hsynch/hflush, and security)
- webhdfs (with full support for security)
- performance enhanced access to local files for HBase
- other performance enhancements, bug fixes, and features

Please see the complete [Hadoop 1.0.0 Release Notes](#) for details.

2.7 10 December, 2011: release 0.22.0 available

This release contains many bug fixes and optimizations compared to its predecessor 0.21.0. See the [Hadoop 0.22.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#)

See [Common News](#) for more details.

2.8 17 Oct, 2011: release 0.20.205.0 available

This release contains improvements, new features, bug fixes and optimizations. This release includes rpms and debs, all duly checksummed and securely signed.

See the [Hadoop 0.20.205.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#).

Notes:

- This release includes a merge of append/hsynch/hflush capabilities from 0.20-append branch, to support HBase in secure mode.
- This release includes the new webhdfs file system, but webhdfs write calls currently fail in secure mode.

2.9 5 Sep, 2011: release 0.20.204.0 available

This release contains improvements, new features, bug fixes and optimizations. This release includes rpms and debs for the first time.

See the [Hadoop 0.20.204.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#).

Notes:

- The RPMs don't work with security turned on. (HADOOP-7599)
- The NameNode's edit log needs to be merged into the image via
 - put the NameNode into safe mode
 - run dfsadmin savenamespace command
 - perform a normal upgrade

2.10 11 May, 2011: release 0.20.203.0 available

This release contains many improvements, new features, bug fixes and optimizations. It is stable and has been deployed in large (4,500 machine) production clusters.

See the [Hadoop 0.20.203.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.11 23 August, 2010: release 0.21.0 available

This release contains many improvements, new features, bug fixes and optimizations. It has not undergone testing at scale and should not be considered stable or suitable for production. This release is being classified as a minor release, which means that it should be API compatible with 0.20.2.

See the [Hadoop 0.21.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.12 22 April, 2009: release 0.20.0 available

This release contains many improvements, new features, bug fixes and optimizations.

See the [Hadoop 0.20.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.13 24 February, 2009: release 0.19.1 available

This release contains many critical bug fixes, including **some data loss issues**. The release also introduces an **incompatible change** by disabling the file append API ([HADOOP-5224](#)) until it can be stabilized.

See the [Hadoop 0.19.1 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.14 29 January, 2009: release 0.18.3 available

This release contains many critical bug fixes.

See the [Hadoop 0.18.3 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.15 21 November, 2008: release 0.19.0 available

This release contains many improvements, new features, bug fixes and optimizations.

See the [Hadoop 0.19.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.16 3 November, 2008: release 0.18.2 available

This release contains several critical bug fixes.

See the [Hadoop 0.18.2 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.17 17 September, 2008: release 0.18.1 available

This release contains several critical bug fixes.

See the [Hadoop 0.18.1 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.18 22 August, 2008: release 0.18.0 available

This release contains many improvements, new features, bug fixes and optimizations.

See the [Hadoop 0.18.0 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.19 19 August, 2008: release 0.17.2 available

This release contains several critical bug fixes.

See the Hadoop 0.17.2 Notes for details.

2.20 23 June, 2008: release 0.17.1 available

This release contains many improvements, new features, bug fixes and optimizations.

See the Hadoop 0.17.1 Notes for details.

2.21 20 May, 2008: release 0.17.0 available

This release contains many improvements, new features, bug fixes and optimizations.

See the Hadoop 0.17.0 Release Notes for details.

2.22 5 May, 2008: release 0.16.4 available

This release fixes 4 critical bugs in release 0.16.3.

2.23 16 April, 2008: release 0.16.3 available

This release fixes critical bugs in release 0.16.2.

2.24 2 April, 2008: release 0.16.2 available

This release fixes critical bugs in release 0.16.1.

Note:

HBase has been removed from this release. HBase releases are now maintained at <http://hadoop.apache.org/hbase/>

2.25 13 March, 2008: release 0.16.1 available

This release fixes critical bugs in release 0.16.0.

Note:

HBase releases are now maintained at <http://hadoop.apache.org/hbase/>

2.26 7 February, 2008: release 0.16.0 available

This release contains many improvements, new features, bug fixes and optimizations. See the release notes (above) for details.

Note:

When upgrading an existing HDFS filesystem to a 0.16.x release from an earlier release, you should first start HDFS with 'bin/start-dfs.sh -upgrade'. See the [Hadoop Upgrade](#) page for details.

2.27 18 January, 2008: release 0.15.3 available

This release fixes critical bugs in release 0.15.3.

2.28 2 January, 2008: release 0.15.2 available

This release fixes critical bugs in release 0.15.1.

2.29 27 November, 2007: release 0.15.1 available

This release fixes critical bugs in release 0.15.0.

2.30 26 November, 2007: release 0.14.4 available

This release fixes critical bugs in release 0.14.3.

2.31 29 October 2007: release 0.15.0 available

This release contains many improvements, new features, bug fixes and optimizations. Notably, this contains the first working version of [HBase](#). See the release notes (above) for details.

Note:

When upgrading an existing HDFS filesystem to a 0.15.x release from an earlier release, you should first start HDFS with 'bin/start-dfs.sh -upgrade'. See the [Hadoop Upgrade](#) page for details.

2.32 19 October, 2007: release 0.14.3 available

This release fixes critical bugs in release 0.14.2.

2.33 4 September, 2007: release 0.14.1 available

New features in release 0.14 include:

- Better checksums in HDFS. Checksums are no longer stored in parallel HDFS files, but are stored directly by datanodes alongside blocks. This is more efficient for the namenode and also improves data integrity.
- Pipes: A C++ API for MapReduce
- Eclipse Plugin, including HDFS browsing, job monitoring, etc.
- File modification times in HDFS.

There are many other improvements, bug fixes, optimizations and new features. Performance and reliability are better than ever.

Note:

When upgrading an existing HDFS filesystem to a 0.14.x release from a 0.13.x or earlier release, you should first start HDFS with 'bin/start-dfs.sh -upgrade'. See the [Hadoop 0.14 Upgrade](#) page for details.