

Apache Hadoop Releases

Table of contents

1 Download.....	2
2 Release Notes.....	3

1 Download

Hadoop is released as source code tarballs with corresponding binary tarballs for convenience. The downloads are distributed via mirror sites and should be checked for tampering using GPG or SHA-256.

Version	Release Date	Tarball	GPG	SHA-256
3.0.0-beta1	03 October, 2017	source	signature	checksum file
		binary	signature	checksum file
2.8.1	08 June, 2017	source	signature	0748C0E2 519382F2..
		binary	signature	B5BE5275 78EF2C85..
2.7.4	04 August, 2017	source	signature	D52B8CE8 446F4C10..
		binary	signature	8F791BFC FA5BB7C7..
2.6.5	08 October, 2016	source	signature	3A843F18 73D9951A..
		binary	signature	001AD18D 4B6D0FE5..

To verify Hadoop releases using GPG:

1. Download the release `hadoop-X.Y.Z-src.tar.gz` from a [mirror site](#).
2. Download the signature file `hadoop-X.Y.Z-src.tar.gz.asc` from [Apache](#).
3. Download the [Hadoop KEYS](#) file.
4. `gpg --import KEYS`
5. `gpg --verify hadoop-X.Y.Z-src.tar.gz.asc`

To perform a quick check using SHA-256:

1. Download the release `hadoop-X.Y.Z-src.tar.gz` from a [mirror site](#).
2. Download the checksum `hadoop-X.Y.Z-src.tar.gz.mds` from [Apache](#).
3. `shasum -a 256 hadoop-X.Y.Z-src.tar.gz`

All previous releases of Hadoop are available from the [Apache release archive](#) site.

Many third parties distribute products that include Apache Hadoop and related tools. Some of these are listed on the [Distributions wiki page](#).

2 Release Notes

2.1 03 October 2017: Release 3.0.0-beta1 available

This is the first beta release in the 3.0.0 release line. It consists of 576 bug fixes, improvements, and other enhancements since 3.0.0-alpha4. This is planned to be the final alpha release, with the next release being 3.0.0 GA.

Please note that beta releases are API stable but come with no guarantees of quality, and are not intended for production use.

Users are encouraged to read the [overview of major changes](#) coming in 3.0.0. The beta1 [release notes](#) and [changelog](#) detail the changes since 3.0.0-alpha4.

2.2 04 August, 2017: Release 2.7.4 available

This is the next release of Apache Hadoop 2.7 line.

Please see the [Hadoop 2.7.4 Release Notes](#) for the list of 264 bugs fixes and optimizations since the previous release 2.7.3.

2.3 07 July 2017: Release 3.0.0-alpha4 available

This is the fourth alpha release in the 3.0.0 release line. It consists of 814 bug fixes, improvements, and other enhancements since 3.0.0-alpha3. This is planned to be the final alpha release, with the next release being 3.0.0-beta1.

Please note that alpha releases come with no guarantees of quality or API stability, and are not intended for production use.

Users are encouraged to read the [overview of major changes](#) coming in 3.0.0. The alpha4 [release notes](#) and [changelog](#) detail the changes since 3.0.0-alpha3.

2.4 08 June, 2017: Release 2.8.1 available

This is a security release in the 2.8.0 release line. It consists of 2.8.0 plus security fixes.

Users on 2.8.0 are encouraged to upgrade to 2.8.1.

Please note that 2.8.x release line continues to be not yet ready for production use. Critical issues are being ironed out via testing and downstream adoption. Production users should wait for a subsequent release in the 2.8.x line.

2.5 26 May 2017: Release 3.0.0-alpha3 available

This is a security release in the 3.0.0 release line. It consists of alpha2 plus security fixes, along with necessary build-related fixes. Users on 3.0.0-alpha1 and 3.0.0-alpha2 are encouraged to upgrade to 3.0.0-alpha3.

Please note that alpha releases come with no guarantees of quality or API stability, and are not intended for production use.

Users are encouraged to read the [overview of major changes](#) coming in 3.0.0. The [alpha3 release notes](#) and [changelog](#) detail the changes since 3.0.0-alpha2.

2.6 22 March 2017: Release 2.8.0 available

Apache Hadoop 2.8.0 contains a number of significant features and enhancements. For major features and improvements, please refer: [overview of major changes](#) coming in 2.8.0. For details of 2917 fixes, improvements, and new features since the previous 2.7.0 release, please check: [release notes](#) and [changelog](#)

Please note that this release is not yet ready for production use. Critical issues are being ironed out via testing and downstream adoption. Production users should wait for a 2.8.1/2.8.2 release.

2.7 25 January, 2017: Release 3.0.0-alpha2 available

This is the second alpha in a series of planned alphas and betas leading up to a 3.0.0 GA release. The intention is to "release early, release often" to quickly iterate on feedback collected from downstream users.

Please note that alpha releases come with no guarantees of quality or API stability, and are not intended for production use.

Users are encouraged to read the [overview of major changes](#) coming in 3.0.0. The [alpha2 release notes](#) and [changelog](#) detail 857 fixes, improvements, and new features since the previous 3.0.0-alpha2 release.

2.8 08 October, 2016: Release 2.6.5 available

A point release for the 2.6 line.

Please see the [Hadoop 2.6.5 Release Notes](#) for the list of 79 bug fixes and patches since the previous release 2.6.4.

2.9 03 September, 2016: Release 3.0.0-alpha1 available

This is the first alpha in a series of planned alphas and betas leading up to a 3.0.0 GA release. The intention is to "release early, release often" to quickly iterate on feedback collected from downstream users.

Please note that alpha releases come with no guarantees of quality or API stability, and are not intended for production use.

Users are encouraged to read the [overview of major changes](#) coming in 3.0.0. The [full set of release notes](#) and [changelog](#) detail all the changes since the previous minor release 2.7.0.

2.10 25 August, 2016: Release 2.7.3 available

A point release for the 2.7 line.

Please see the [Hadoop 2.7.3 Release Notes](#) for the list of 221 bug fixes and patches since the previous release 2.7.2.

2.11 11 February, 2016: Release 2.6.4 available

A point release for the 2.6 line.

Please see the [Hadoop 2.6.4 Release Notes](#) for the list of 46 bug fixes and patches since the previous release 2.6.3.

2.12 25 January, 2016: Release 2.7.2 (stable) available

A point release for the 2.7 line.

Please see the [Hadoop 2.7.2 Release Notes](#) for the list of 155 bug fixes and patches since the previous release 2.7.1.

2.13 17 December, 2015: Release 2.6.3 available

Apache Hadoop 2.6.3 is a point release in the 2.6.x release line, and fixes a few critical issues in 2.6.2.

Please see the [Hadoop 2.6.3 Release Notes](#) for details.

2.14 28 October, 2015: Release 2.6.2 available

Apache Hadoop 2.6.2 is a point release in the 2.6.x release line, and fixes a few critical issues in 2.6.1.

Please see the [Hadoop 2.6.2 Release Notes](#) for details.

2.15 23 September, 2015: Release 2.6.1 available

Apache Hadoop 2.6.1 is a point release in the 2.6.x release line, and fixes a lot of critical issues in 2.6.0.

Please see the [Hadoop 2.6.1 Release Notes](#) for details.

2.16 06 July, 2015: Release 2.7.1 (stable) available

A point release for the 2.7 line. This release is now considered stable.

Please see the [Hadoop 2.7.1 Release Notes](#) for the list of 131 bug fixes and patches since the previous release 2.7.0. Please look at the 2.7.0 section below for the list of enhancements enabled by this first stable release of 2.7.x.

2.17 21 April 2015: Release 2.7.0 available

Apache Hadoop 2.7.0 contains a number of significant enhancements. A few of them are noted below.

- IMPORTANT notes
 - This release drops support for JDK6 runtime and works with JDK 7+ only.
 - This release is not yet ready for production use. Critical issues are being ironed out via testing and downstream adoption. Production users should wait for a 2.7.1/2.7.2 release.
- Hadoop Common
 - [HADOOP-9629](#) - Support Windows Azure Storage - Blob as a file system in Hadoop.
- Hadoop HDFS
 - [HDFS-3107](#) - Support for file truncate
 - [HDFS-7584](#) - Support for quotas per storage type
 - [HDFS-3689](#) - Support for files with variable-length blocks
- Hadoop YARN
 - [YARN-3100](#) - Make YARN authorization pluggable
 - [YARN-1492](#) - Automatic shared, global caching of YARN localized resources (beta)
- Hadoop MapReduce
 - [MAPREDUCE-5583](#) - Ability to limit running Map/Reduce tasks of a job
 - [MAPREDUCE-4815](#) - Speed up FileOutputCommitter for very large jobs with many output files.

Please see the [Hadoop 2.7.0 Release Notes](#) for details.

2.18 18 November, 2014: Release 2.6.0 available

Apache Hadoop 2.6.0 contains a number of significant enhancements such as:

- Hadoop Common
 - [HADOOP-10433](#) - Key management server (beta)
 - [HADOOP-10607](#) - Credential provider (beta)
- Hadoop HDFS
 - Heterogeneous Storage Tiers - Phase 2
 - [HDFS-5682](#) - Application APIs for heterogeneous storage
 - [HDFS-7228](#) - SSD storage tier
 - [HDFS-5851](#) - Memory as a storage tier (beta)
 - [HDFS-6584](#) - Support for Archival Storage

- [HDFS-6134](#) - Transparent data at rest encryption (beta)
- [HDFS-2856](#) - Operating secure DataNode without requiring root access
- [HDFS-6740](#) - Hot swap drive: support add/remove data node volumes without restarting data node (beta)
- [HDFS-6606](#) - AES support for faster wire encryption
- Hadoop YARN
 - [YARN-896](#) - Support for long running services in YARN
 - [YARN-913](#) - Service Registry for applications
 - [YARN-666](#) - Support for rolling upgrades
 - [YARN-556](#) - Work-preserving restarts of ResourceManager
 - [YARN-1336](#) - Container-preserving restart of NodeManager
 - [YARN-796](#) - Support node labels during scheduling
 - [YARN-1051](#) - Support for time-based resource reservations in Capacity Scheduler (beta)
 - [YARN-1964](#) - Support running of applications natively in Docker containers (alpha)

Please see the [Hadoop 2.6.0 Release Notes](#) for details.

2.19 19 November, 2014: Release 2.5.2 available

Apache Hadoop 2.5.2 is a point release in the 2.5.x release line, and fixes a few critical issues in 2.5.1

2.20 12 September, 2014: Release 2.5.1 available

Apache Hadoop 2.5.1 is a point release in the 2.5.x release line, and fixes a few release issues with 2.5.0.

2.21 11 August, 2014: Release 2.5.0 available

Apache Hadoop 2.5.0 is a minor release in the 2.x release line.

The release includes the following major features and improvements:

- Authentication improvements when using an HTTP proxy server.
- A new Hadoop Metrics sink that allows writing directly to Graphite.
- Specification for Hadoop Compatible Filesystem effort.
- Support for POSIX-style filesystem extended attributes.
- OfflineImageViewer to browse an fsimage via the WebHDFS API.
- Supportability improvements and bug fixes to the NFS gateway.
- Modernized web UIs (HTML5 and Javascript) for HDFS daemons.
- YARN's REST APIs support submitting and killing applications.
- Kerberos integration for the YARN's timeline store.
- FairScheduler allows creating user queues at runtime under any specified parent queue.

Users are encouraged to try out 2.5.0

Please see the [Hadoop 2.5.0 Release Notes](#) for details.

2.22 30 June, 2014: Release 2.4.1 available

Apache Hadoop 2.4.1 is a bug-fix release for the stable 2.4.x line.

There is also a security bug fix in this minor release.

- CVE-2014-0229: Add privilege checks to HDFS admin sub-commands refreshNamenodes, deleteBlockPool and shutdownDatanode.

Users are encouraged to immediately move to 2.4.1.

Please see the [Hadoop 2.4.1 Release Notes](#) for details.

2.23 27 June, 2014: Release 0.23.11 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.11 Release Notes](#) for details.

2.24 07 April, 2014: Release 2.4.0 available

Apache Hadoop 2.4.0 contains a number of significant enhancements such as:

- Support for Access Control Lists in HDFS
- Native support for Rolling Upgrades in HDFS
- Usage of protocol-buffers for HDFS FSImage for smooth operational upgrades
- Complete HTTPS support in HDFS
- Support for Automatic Failover of the YARN ResourceManager
- Enhanced support for new applications on YARN with Application History Server and Application Timeline Server
- Support for strong SLAs in YARN CapacityScheduler via Preemption

Please see the [Hadoop 2.4.0 Release Notes](#) for details.

2.25 20 February, 2014: Release 2.3.0 available

Apache Hadoop 2.3.0 contains a number of significant enhancements such as:

- Support for Heterogeneous Storage hierarchy in HDFS.
- In-memory cache for HDFS data with centralized administration and management.
- Simplified distribution of MapReduce binaries via HDFS in YARN Distributed Cache.

Please see the [Hadoop 2.3.0 Release Notes](#) for details.

2.26 11 December, 2013: Release 0.23.10 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.10 Release Notes](#) for details.

2.27 15 October, 2013: Release 2.2.0 available

Apache Hadoop 2.2.0 is the **GA** release of Apache Hadoop 2.x.

Users are encouraged to immediately move to 2.2.0 since this release is significantly more stable and is guaranteed to remain compatible in terms of both APIs and protocols.

To recap, this release has a number of significant highlights compared to Hadoop 1.x:

- YARN - A general purpose resource management system for Hadoop to allow MapReduce and other other data processing frameworks and services
- High Availability for HDFS
- HDFS Federation
- HDFS Snapshots
- NFSv3 access to data in HDFS
- Support for running Hadoop on Microsoft Windows
- Binary Compatibility for MapReduce applications built on hadoop-1.x
- Substantial amount of integration testing with rest of projects in the ecosystem

A couple of important points to note while upgrading to hadoop-2.2.0:

- HDFS - The HDFS community decided to push the symlinks feature out to a future 2.3.0 release and is currently disabled.
- YARN/MapReduce - Users need to change ShuffleHandler service name from mapreduce.shuffle to mapreduce_shuffle.

Please see the [Hadoop 2.2.0 Release Notes](#) for details.

2.28 23 September, 2013: Release 2.1.1-beta available

Apache Hadoop 2.1.1-beta is a bug-fix version of the beta release of Apache Hadoop 2.x.

Please see the [Hadoop 2.1.1-beta Release Notes](#) for details.

2.29 25 August, 2013: Release 2.1.0-beta available

Apache Hadoop 2.1.0-beta is the **beta** release of Apache Hadoop 2.x.

Users are encouraged to immediately move to 2.1.0-beta since this release is significantly more stable and has completely whetted set of APIs and wire-protocols for future compatibility.

In addition, this release has a number of other significant highlights:

- HDFS Snapshots
- Support for running Hadoop on Microsoft Windows
- YARN API stabilization
- Binary Compatibility for MapReduce applications built on hadoop-1.x
- Substantial amount of integration testing with rest of projects in the ecosystem

Please see the [Hadoop 2.1.0-beta Release Notes](#) for details.

2.30 23 August, 2013: Release 2.0.6-alpha available

This release delivers a number of critical bug-fixes for hadoop-2.x uncovered during integration testing of previous release.

Please see the [Hadoop 2.0.6-alpha Release Notes](#) for details.

2.31 1 Aug, 2013: Release 1.2.1 (stable) available

A point release for the 1.2 line. This release is now considered stable.

Please see the [Hadoop 1.2.1 Release Notes](#) for the list of 18 bug fixes and patches since the previous release 1.2.0.

2.32 8 July, 2013: Release 0.23.9 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.9 Release Notes](#) for details.

2.33 6 June, 2013: Release 2.0.5-alpha available

This release delivers a number of critical bug-fixes for hadoop-2.x uncovered during integration testing of previous release.

Please see the [Hadoop 2.0.5-alpha Release Notes](#) for details.

2.34 5 June, 2013: Release 0.23.8 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.8 Release Notes](#) for details.

2.35 13 May, 2013: Release 1.2.0 available

This is a beta release for version 1.2.

This release delivers over 200 enhancements and bug-fixes, compared to the previous 1.1.2 release. Major enhancements include:

- DistCp v2 backported

- Web services for JobTracker
- WebHDFS enhancements
- Extensions of task placement and replica placement policy interfaces
- Offline Image Viewer backported
- Namenode more robust in case of edit log corruption
- Add NodeGroups level to NetworkTopology
- Add "unset" to Configuration API

Please see the [Hadoop 1.2.0 Release Notes](#) for more details.

2.36 25 April, 2013: Release 2.0.4-alpha available

This release delivers a number of critical bug-fixes for hadoop-2.x uncovered during integration testing.

Please see the [Hadoop 2.0.4-alpha Release Notes](#) for details.

2.37 18 April, 2013: Release 0.23.7 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.7 Release Notes](#) for details.

2.38 15 February, 2013: Release 1.1.2 available

Point release for the 1.1.X line. Bug fixes and improvements, as documented in the [Hadoop 1.1.2 Release Notes](#).

2.39 14 February, 2013: Release 2.0.3-alpha available

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

This release delivers significant major features and stability over previous releases in hadoop-2.x series:

- QJM for HDFS HA for NameNode
- Multi-resource scheduling (CPU and memory) for YARN
- YARN ResourceManager Restart
- Significant stability at scale for YARN (over 30,000 nodes and 14 million applications so far, at time of release)

This release, like previous releases in hadoop-2.x series is still considered *alpha* primarily since some of APIs aren't fully-baked and we expect some churn in future. Furthermore, please note that there are some API changes from previous hadoop-2.0.2-alpha release and applications will need to recompile against hadoop-2.0.3-alpha.

Please see the [Hadoop 2.0.3-alpha Release Notes](#) for details.

2.40 7 February, 2013: Release 0.23.6 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.6 Release Notes](#) for details.

2.41 1 December, 2012: Release 1.1.1 available

Point release for the 1.1.X line. Bug fixes and improvements, as documented in the [Hadoop 1.1.1 Release Notes](#).

2.42 28 November, 2012: Release 0.23.5 available

A point release for the 0.23.X line. Bug fixes to continue stabilization.

Please see the [Hadoop 0.23.5 Release Notes](#) for details.

2.43 15 October, 2012: Release 0.23.4 available

A point release for the 0.23.X line. Most notably this release now includes support for upgrading a 1.X HDFS cluster with append/synch enabled.

Please see the [Hadoop 0.23.4 Release Notes](#) for details.

2.44 13 October, 2012: Release 1.1.0 available

This is a beta release for version 1.1.

This release has approximately 135 enhancements and bug fixes compared to Hadoop-1.0.4, including:

- Many performance improvements in HDFS, backported from trunk
- Improvements in Security to use SPNEGO instead of Kerberized SSL for HTTP transactions
- Lower default minimum heartbeat for task trackers from 3 sec to 300msec to increase job throughput on small clusters
- Port of Gridmix v3
- Set MALLOC_ARENA_MAX in hadoop-config.sh to resolve problems with glibc in RHEL-6
- Splittable bzip2 files
- Of course it also has the same security fix as release 1.0.4.

Please see the [Hadoop 1.1.0 Release Notes](#) for details.

2.45 12 October, 2012: Release 1.0.4 available

This is a Security Patch release for version 1.0.

There are four bug fixes and feature enhancements in this minor release:

- Security issue CVE-2012-4449: Hadoop tokens use a 20-bit secret
- HADOOP-7154 - set MALLOC_ARENA_MAX in hadoop-config.sh to resolve problems with glibc in RHEL-6
- HDFS-3652 - FSEditLog failure removes the wrong edit stream when storage dirs have same name
- MAPREDUCE-4399 - Fix (up to 3x) performance regression in shuffle

Please see the [Hadoop 1.0.4 Release Notes](#) for details.

2.46 9 October, 2012: Release 2.0.2-alpha available

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

This delivers significant enhancements to HDFS HA. Also it has a significantly more stable version of YARN which, at the time of release, has already been deployed on a 2000 node cluster.

Please see the [Hadoop 2.0.2-alpha Release Notes](#) for details.

2.47 17 September, 2012: Release 0.23.3 available

This release contains YARN and MRv2 but does not have Name Node High Availability

Please see the [Hadoop 0.23.3 Release Notes](#) for details.

2.48 26 July, 2012: Release 2.0.1-alpha available

This release contains important security fixes over hadoop-2.0.0-alpha.

Please see the [Hadoop 2.0.1-alpha Release Notes](#) for details.

2.49 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.

2.50 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.

2.51 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.

2.52 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.

2.53 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.

2.54 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.

2.55 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](#)

Notes:

The following features are not supported in Hadoop 0.22.0.

- Security.
- Latest optimizations of the MapReduce framework introduced in the Hadoop 0.20.security line of releases.
- Disk-fail-in-place.
- JMX-based metrics v2.

Hadoop 0.22.0 features

- HBase support with `hflush` and `hsync`.
- New implementation of file append.
- Symbolic links.
- BackupNode and CheckpointNode.
- Hierarchical job queues.
- Job limits per queue/pool.
- Dynamically stop/start job queues.
- Advances in new mapreduce API: Input/Output formats, ChainMapper/Reducer.
- TaskTracker blacklisting.
- DistributedCache sharing.

2.56 11 Nov, 2011: release 0.23.0 available

This is the alpha version of the hadoop-0.23 major release. This is the first release we've made off Apache Hadoop trunk in a long while. This release is alpha-quality and not yet ready for serious use.

hadoop-0.23 contains several major advances:

- HDFS Federation
- NextGen MapReduce (YARN)

It also has several major performance improvements to both HDFS and MapReduce.

See the [Hadoop 0.23.0 Release Notes](#) for details.

2.57 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.

2.58 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

2.59 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](#).

2.60 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](#).

2.61 26 February, 2010: release 0.20.2 available

This release contains several critical bug fixes.

See the [Hadoop 0.20.2 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.62 14 September, 2009: release 0.20.1 available

This release contains several critical bug fixes.

See the [Hadoop 0.20.1 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.63 23 July, 2009: release 0.19.2 available

This release contains several critical bug fixes.

See the [Hadoop 0.19.2 Release Notes](#) for details. Alternatively, you can look at the complete [change log for this release](#) or the [Jira issue log for all releases](#).

2.64 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](#).

2.65 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](#).

2.66 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](#).

2.67 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](#).

2.68 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](#).

2.69 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](#).

2.70 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](#).

2.71 19 August, 2008: release 0.17.2 available

This release contains several critical bug fixes.

See the [Hadoop 0.17.2 Notes](#) for details.

2.72 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.

2.73 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.

2.74 5 May, 2008: release 0.16.4 available

This release fixes 4 critical bugs in release 0.16.3.

2.75 16 April, 2008: release 0.16.3 available

This release fixes critical bugs in release 0.16.2.

2.76 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/>

2.77 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/>

2.78 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.

2.79 18 January, 2008: release 0.15.3 available

This release fixes critical bugs in release 0.15.3.

2.80 2 January, 2008: release 0.15.2 available

This release fixes critical bugs in release 0.15.1.

2.81 27 November, 2007: release 0.15.1 available

This release fixes critical bugs in release 0.15.0.

2.82 26 November, 2007: release 0.14.4 available

This release fixes critical bugs in release 0.14.3.

2.83 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.

2.84 19 October, 2007: release 0.14.3 available

This release fixes critical bugs in release 0.14.2.

2.85 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.