

Welcome to Hadoop!

Table of contents

1 News.....	2
1.1 2 June, 2006: release 0.3.0 available.....	2
1.2 12 May, 2006: release 0.2.1 available.....	2
1.3 5 May, 2006: release 0.2.0 available.....	2
1.4 2 April, 2006: release 0.1.0 available.....	2
1.5 6 February, 2006: nightly builds.....	2
1.6 3 February, 2006: Hadoop code moved out of Nutch.....	2
1.7 30 March, 2006: Hadoop project approved.....	2

1. News

1.1. 2 June, 2006: release 0.3.0 available

This includes many fixes, improving performance, scalability and reliability and adding new features. For details see the [change log](#). The release can be obtained from [a nearby mirror](#).

1.2. 12 May, 2006: release 0.2.1 available

This fixes a few bugs in release 0.2.0, listed in the [change log](#). The release can be obtained from [a nearby mirror](#).

1.3. 5 May, 2006: release 0.2.0 available

We are now aiming for monthly releases. There have been many bug fixes and improvements in the past month. MapReduce and DFS are now fairly reliable on clusters with up to 200 nodes. See the [change log](#) for details. The release can be obtained from [a nearby mirror](#).

1.4. 2 April, 2006: release 0.1.0 available

This is the first Hadoop release. The release is available [here](#).

1.5. 6 February, 2006: nightly builds

Hadoop now has nightly builds. This automatically creates a [downloadable version of Hadoop every night](#). All unit tests must pass, or a message is sent to the developers mailing list and no new version is created. This also updates the [javadoc](#).

1.6. 3 February, 2006: Hadoop code moved out of Nutch

The Hadoop code has now been moved into its own Subversion tree, renamed into packages under `org.apache.hadoop`. All unit tests pass, but little else has yet been tested.

1.7. 30 March, 2006: Hadoop project approved

The Lucene PMC has elected to split the Nutch MapReduce and distributed filesystem code into a new project named Hadoop.