

# Releases

## Table of contents

1 Download.....	2
2 Release Notes.....	2
3 News.....	2
3.1 25 July, 2007: release 0.13.1 available.....	2
3.2 8 June, 2007: release 0.13.0 available .....	2
3.3 8 June, 2006: FAQ added to Wiki .....	2
3.4 5 May, 2006: release 0.2.0 available .....	2
3.5 6 February, 2006: nightly builds .....	3
3.6 3 February, 2006: Hadoop code moved out of Nutch .....	3

## 1. Download

Releases may be downloaded from Apache mirrors.

[Download a release now!](#)

On the mirror, all recent releases are available, but are not guaranteed to be stable. For stable releases, look in the stable directory.

## 2. Release Notes

Release notes for Hadoop releases are available in Jira.

[Browse release notes now!](#)

## 3. News

### 3.1. 25 July, 2007: release 0.13.1 available

This fixes critical bugs in the 0.13.0 release.

### 3.2. 8 June, 2007: release 0.13.0 available

In addition to new features, this release contains significant performance and reliability improvements. For example, MapReduce sort performance on large clusters has [doubled](#).

**Note:**

When upgrading an existing HDFS filesystem to this release, you should first start HDFS with `bin/start-dfs.sh -upgrade` which makes it possible to revert the upgrade. For more information on this issue, see [HADOOP-702](#).

### 3.3. 8 June, 2006: FAQ added to Wiki

Hadoop now has a [FAQ](#). Please help make this more complete!

### 3.4. 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.

### **3.5. 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](#).

### **3.6. 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.