

Welcome to HBase!

1. News



November 19th 2010, [Hadoop HUG in London](#) is all about HBase

November 15-19th 2010, [Devoxx](#) features HBase Training and multiple HBase presentations

October 12th 2010, HBase-related presentations by core contributors and users at [Hadoop World 2010](#)

October 11th 2010, [HUG-NYC: HBase User Group NYC Edition](#) (Night before Hadoop World)

October 5th, 2010: HBase 0.89.20100924, the third in our series of 'developer preview' releases, is now available; [Download!](#). See [HBase Versions](#) for more on this developer series of releases.

August 9th, 2010: HBase 0.89.20100726, the second in our series of 'developer preview' releases, is now available; [Download!](#). See [HBase Versions](#) for more on this developer series of releases.

July 30th, 2010: HBase 0.20.6 released; [Download!](#)

June 25th, 2010: HBase 0.89.20100621, the first in a series of 'developer preview' releases, is now available; [Download!](#). See [HBase Versions](#) for more on this developer series of releases.

June 24th, 2010: HBase 0.20.5 released; [Download!](#)

June 30th, [HBase Contributor Workshop](#) (Day after Hadoop Summit)

May 10th, 2010: HBase graduates from Hadoop sub-project to Apache Top Level Project

[Old news...](#)

2. Overview

HBase is the [Hadoop database](#). Use it when you need random, realtime read/write access to your Big Data. This project's goal is the hosting of very large tables -- billions of rows X millions of columns -- atop clusters of commodity hardware.

HBase is an open-source, distributed, versioned, column-oriented store modeled after Google's [Bigtable: A Distributed Storage System for Structured Data](#) by Chang et al. Just as Bigtable leverages the distributed data storage provided by the Google File System, HBase provides Bigtable-like capabilities on top of Hadoop. HBase includes:

- Convenient base classes for backing Hadoop MapReduce jobs with HBase tables
- Query predicate push down via server side scan and get filters
- Optimizations for real time queries
- A high performance Thrift gateway
- A REST-ful Web service gateway that supports XML, Protobuf, and binary data encoding options
- [Cascading](#), [hive](#), and [pig](#) source and sink modules
- Extensible [jruby](#)-based (JIRB) shell
- Support for exporting metrics via the Hadoop metrics subsystem to files or Ganglia; or via JMX

Welcome to HBase!

HBase 0.20 has greatly improved on its predecessors:

- No HBase single point of failure
- Rolling restart for configuration changes and minor upgrades
- Random access performance on par with open source relational databases such as MySQL

3. Getting Started

Start by installing HBase on a single machine or a very small cluster.

1. [Learn about](#) HBase by reading the documentation.
2. [Download](#) HBase from the release page.

4. Getting Involved

HBase is an open source volunteer project under the Apache Software Foundation. It is a subproject of [Hadoop](#). We encourage you to learn about the project and contribute your expertise. Here are some starter links:

1. See our [How to Contribute to HBase](#) page.
2. Give us [feedback](#): What can we do better?
3. Join the [mailing list](#): Meet the community.
4. Get the latest at up on the [hbase blog](#).
5. Come hang with us on [IRC](#).
6. Become an [HBase Fan](#) up on facebook.
7. Or up on twitter, [HBase](#).