

Welcome to HBase!

1. 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, 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](#) 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

This most recent version of HBase, 0.20.0, 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

2. News

2.1. HBase 0.20.0 released! September 8th, 2009

HBase 0.20.0 is faster, stronger, slimmer, and sweeter tasting than any previous HBase release. Get it off the [Releases](#) page.

2.2. HBase at ApacheCon in Oakland: November 2-6th, 2009

The Apache Foundation will be celebrating its 10th anniversary in beautiful Oakland by the Bay. Lots of good talks and meetups including an HBase presentation by a couple of the lads. <http://dev.us.apachecon.com/c/acus2009/>

2.3. HBase at Hadoop World in NYC: October 2nd, 2009

A few of us will be talking on Practical HBase out east at [Hadoop World: NYC](#).

2.4. HUG7 and HBase Hackathon: August 7th-9th, 2009 at StumbleUpon in SF

Sign up for the [HBase User Group Meeting #7](#) or for the [Hackathon](#) or for both (all are welcome!).

Welcome to HBase!

2.5. June, 2009 -- HBase at HadoopSummit2009 and at NOSQL

See the [presentations](#)

2.6. March 3rd, 2009 -- HUG6

[HBase User Group 6](#)

2.7. January 30th, 2009 -- LA Hbackathon

[HBase January Hackathon Los Angeles](#) at [Streamy](#) in Manhattan Beach

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. Come hang with us on [IRC](#).
5. Become an [HBase Fan](#) up on facebook.
6. Or up on twitter, [HBase](#).