

Home

This is the home of the Hadoop space.

[Apache Hadoop](#) is a framework for running applications on large clusters built of commodity hardware. The Hadoop framework transparently provides applications for both reliability and data motion. Hadoop implements a computational paradigm named Map/Reduce, where the application is divided into many small fragments of work, each of which may be executed or re-executed on any node in the cluster. In addition, it provides a distributed file system (HDFS) that stores data on the compute nodes, providing very high aggregate bandwidth across the cluster. Both MapReduce and the Hadoop Distributed File System are designed so that node failures are automatically handled by the framework.



Old-wiki

The content of the deprecated old hadoop wiki (<https://wiki.apache.org/hadoop>) has been moved to [HADOOP2](#).

Roadmap

[Hadoop Roadmap](#)

Releases

See the [Apache Hadoop Releases](#) page for prior releases.