



cloudera

# Dell | Cloudera Hadoop Solution

Leverage an enterprise-ready distribution of Apache Hadoop.

## Streamlined to solve business problems

When you move forward with the Dell™ | Cloudera™ Hadoop™ Solution, you have the confidence that comes with Cloudera's distribution including Apache Hadoop (CDH), one of the world's leading distributions of Hadoop in commercial and non-commercial environments.

CDH delivers a streamlined path for putting Apache Hadoop to work solving business problems in production. Ideal for enterprises seeking a stable, tested Hadoop solution without proprietary vendor lock-in, CDH is the bridge between the insights of organizations using Hadoop in production and the continuous stream of innovations from the Apache community.

Fully documented and ready to go, CDH consists of 100-percent open source Apache Hadoop plus the comprehensive set of open source software components needed to use Hadoop, which can be activated on an as-needed basis. CDH is thoroughly tested and certified to integrate with a wide range of operating systems and hardware, databases and data warehouses, and business intelligence and extract, transform and load (ETL) systems. This broad compatibility helps your organization take advantage of Hadoop while leveraging your existing tools and resources.

## Simplicity and certainty

Due to the explosive growth in data volumes and complexity, companies of all sizes are adopting the open source Apache Hadoop platform to store, process, and derive value from all of their data. There are many advantages to Hadoop, particularly in the areas of flexibility and economics. But just like any open source technology, implementation into a production environment presents a unique set of challenges and questions.

How can you guarantee system stability? Which components of the open source platform do you need and how can you ensure that they will integrate correctly? How well is the technology documented? Who do you call for support? The uncertainty caused by questions like these drives some companies to invest in proprietary vendor solutions—despite the higher price tag and risk of vendor lock-in.

Cloudera created its distribution of Hadoop for this very purpose—to remove the uncertainty and barriers that sometimes dissuade enterprises from deploying open source Hadoop in production.

## A complete solution

The Dell | Cloudera Hadoop Solution gives you everything you need to tackle your big-data challenges—including software, hardware, and services. The solution is based on the Dell™ PowerEdge™ C2100 rack server and the Dell™ PowerConnect™ 6248 Ethernet Switch—components that are proven in cloud and big-data deployments.

The Dell PowerEdge C2100 server is designed to maximize space, power, and cost efficiency in data centers where memory and storage density are critical. It provides memory and disk capacity for

MapReduce, Web analytics, database, and cloud computing environments. The Dell PowerConnect 6248 48-port Gigabit Ethernet Layer 3 switch delivers significant rack density and advanced core-switching capabilities.

In addition, Dell and Cloudera work cooperatively to provide services and support. This gives you the assurance that your solution is backed by expert hardware and software support that can be tailored to your specific needs.

## Streamlined path to success with Hadoop

- Gain all of the components needed for use in production; implement components on an as-needed basis.
- Remove uncertainty with a fully documented distribution.
- Protect existing investments and the integration work is done for you.

## Stable and reliable Hadoop without vendor lock-in

- Leverage 100% pure Apache Hadoop, hardened for production environments.
- Increase your confidence with a distribution thoroughly tested by expert engineers.
- Scale across tens of thousands of nodes in enterprise environments.

## Community driven

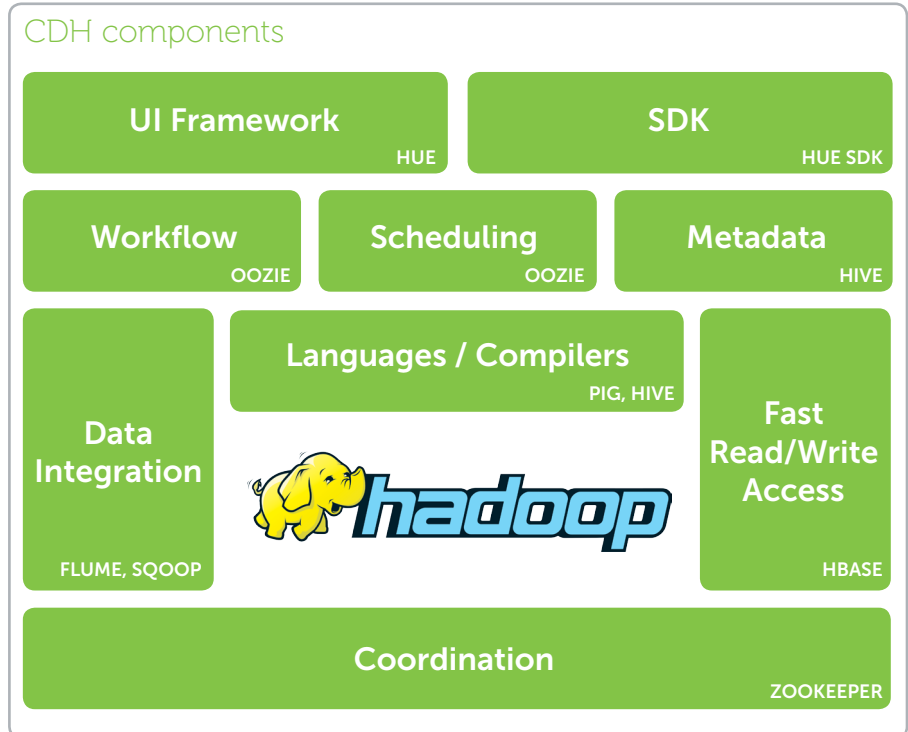
- Incorporate only open source components—no forks or proprietary underpinnings.
- Gain access to a continuous stream of innovations from the Hadoop community.
- Leverage a proven distribution that is free to download.

**“Cloudera’s distribution and support have been instrumental in helping Groupon deliver on our goal to be a technology leader.” – Groupon**

What CDH delivers	
A complete, packaged set of all components needed to use Apache Hadoop in production	CDH contains everything you need for a successful implementation, and you can deploy the different components as you need them. All of the integration work is done for you and the entire solution is fully documented. By taking the guesswork out of building out your Hadoop stack, CDH gives you a streamlined path to success in solving real business problems.
100% pure Apache Hadoop, hardened for production environments	With CDH, Cloudera maintains a vigorous commitment to upholding the open source standard. From the core file system to the UI framework, CDH contains absolutely no code forks or proprietary underpinnings and is free to download. The entire distribution is QA tested by Cloudera and proven at scale across tens of thousands of nodes and varying workloads in enterprise environments. CDH is stable, reliable Hadoop without vendor lock-in.
The Cloudera Ready certification framework	CDH has been tested and certified to integrate with a wide range of hardware, operating systems, and data analysis tools—lowering risk, easing deployment, and protecting your investments in your existing data analysis and business intelligence solutions.

CDH components	
Apache Hadoop	Reliable, scalable distributed storage and computing
Apache Hive	SQL-like language and metadata repository
Apache Pig	High-level language for expressing data analysis programs
Apache HBase	Hadoop database for random, real-time read/write access
Apache Zookeeper	Highly reliable distributed coordination service
Apache Whirr	Library for running Hadoop in the cloud
Flume	Distributed service for collecting and aggregating log and event data
Hue	Browser-based desktop interface for interacting with Hadoop
Oozie	Server-based workflow engine for Hadoop activities
Sqoop	Integrating Hadoop with RDBMS

CDH 3.5 distribution details	
Component versions	
Apache Hadoop	v0.20.2 + 923
Apache Hive	v0.7.0 +27
Apache Pig	v0.8.0 +20
Apache HBase	v0.90.1 +15
Apache Zookeeper	v3.3.2 +12
Apache Whirr	v0.3.0 +5
Flume	v0.9.3 +15
Hue	v1.2.0 +54
Oozie	v2.3.0 +31
Sqoop	v1.2 +24
Supported operating system versions	
Red Hat	RHEL 5, RHEL 6
CentOS	CentOS 5
Ubuntu	Lucid, Maverick
SUSE	SUSE 11
Supported build infrastructure and cloud platforms	
Build Infrastructure	Apache Maven
Cloud Platforms	Rackspace, Amazon EC2, Softlayer



© 2011 Dell Inc. All rights reserved. Dell, the DELL logo, the DELL badge and PowerEdge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others. This document is for informational purposes only. Dell reserves the right to make changes without further notice to the products herein. The content provided is as-is and without expressed or implied warranties of any kind.

To learn more, visit:  
[www.Dell.com/Hadoop](http://www.Dell.com/Hadoop)

PowerEdge C

