

Brief

Derive new business insights

HP with Hortonworks Data Platform for Apache Hadoop

Store, manage, and process Big Data at scale using a modern architecture leveraging SUSE Linux Enterprise Server, designed to reduce complexity and increase value and performance.



Key use cases

- **Finance** – Fraud detection, anti-money laundering, and risk management
- **Government** – Fraud detection, anti-money laundering, and risk management
- **Telecom** – Broadcast monitoring, churn prevention, and advertising optimization
- **Manufacturing** – Supply chain optimization, defect tracking, RFID correlation, and warranty management
- **Energy** – Weather forecasting and natural resource exploration
- **Healthcare** – Drug development, scientific research, evidence-based medicine, and healthcare outcomes analysis

Channel the Big Data deluge

Enterprises today deal with more data than ever before, and storing, accessing, and understanding it creates big challenges.

Apache Hadoop is an open-source project that provides an innovative new platform for consolidating, combining, and understanding large-scale data. It was designed to provide fast, reliable analysis of structured and complex data at Internet scale. As a result, many enterprises deploy Hadoop alongside existing IT systems, which allows them to combine old and new data sets in powerful new ways.

HP and SUSE provide several performance-optimized configurations for deploying Hortonworks Data Platform (HDP) clusters of varying sizes, significantly reducing complexity and increasing value and performance.

HDP includes the full suite of essential Hadoop capabilities and allows you to deploy Hadoop wherever you want it—from the cloud or on-premises, as a cluster, or as an appliance—with accessibility from heterogeneous platforms.

As the only 100 percent open source Hadoop data platform available, HDP also allows you to fully leverage your existing investments.

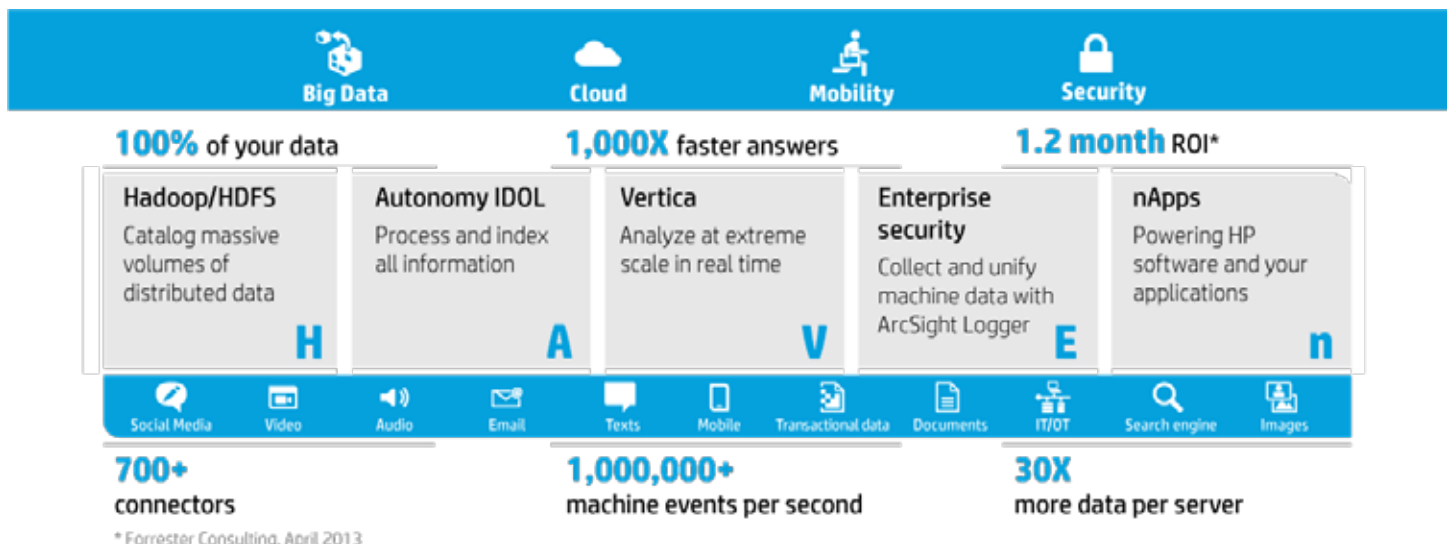
Reference architecture

HP ProLiant DL and SL servers, along with SUSE Linux Enterprise Server, provide optimum computational performance for Hadoop.

When price/performance for high I/O bandwidth workloads is of primary importance, compute power-optimized systems such as the HP ProLiant DL380p Server are ideal candidates. Typical applications in this tier are:

- Extract/transform/load
- Real-time streaming
- NoSQL
- In-memory processing
- Machine learning
- Search indexing

The ProLiant DL380p Server has the high CPU clock rate and large memory configuration required for these workloads, with a data footprint that ranges from hundreds of gigabytes to multiple terabytes.



When low cost per terabyte is the prime factor, storage-optimized systems, such as HP ProLiant SL4540 Servers, are an ideal fit. Typical workloads are:

- Web log analytics
- Sentiment and click stream analysis
- Applications such as Hive and MapReduce

The HP ProLiant SL4540 Server platform is an ideal candidate for this tier as it is slanted toward higher storage density to tackle the large data volume and is also very good for mirroring cold data for backup/disaster recovery.

Our solution partners



HP recommends using a 64-bit operating system to avoid constraining the amount of memory that can be used on worker nodes. SUSE is a leader in high-performance computing, with more than half of the world's largest super computer clusters. HP supports the 64-bit version of SUSE Linux Enterprise Server, which offers superior file system, performance and scalability characteristics, plus the comprehensive, certified support of the Hortonworks and HP-supplied software used in Big Data clusters.

In addition, the SUSE YES Certified Program ensures your HP system platform is compatible, certified, and supported with SUSE Linux Enterprise Server. By using SUSE Linux Enterprise Server in a Hortonworks Hadoop solution, you get a completely supported stack of hardware, operating system, and enterprise-grade data platform. Plus, SUSE's extensive ISV portfolio of software, including Hortonworks, means you have access to even more supported combinations of analytics applications, like HP Vertica, HP Autonomy, and more.

HP services and support

To help you jump-start your Hadoop solution development, HP offers several Big Data services, including Factory Express and Technical Services Consulting. With the purchase of Factory Express services, your Hadoop cluster will arrive racked and cabled, with software installed and configured per an agreed upon custom statement of work. Technology Services Consulting offers specialized Hadoop design, implementation, and installation and setup services. HP offers a variety of support levels to meet your needs.

Get started today

Find out more details on the solution configuration by reading about the [reference architecture](#).

Learn more at

hp.com/go/suse
hp.com/go/hortonworks
suse.com/hp
hortonworks.com/partner/hp

Sign up for updates
hp.com/go/getupdated



Share with colleagues



Rate this document

