



Highlights

- Support the use of unstructured, semi and structured data in-motion and at-rest with the Hortonworks Data Platform (HDP) providing Massively Parallel Processing (MPP) and concurrency on Hadoop
 - Gain flexibility and move CAPEX to OPEX with an enterprise grade, secure and encrypted environment hosted on the IBM Cloud
 - Benefit from the IBM common SQL engine that understands commonly used ANSI SQL syntax from other vendors and products driving data federation and the ability to do real-time ad-hoc and complex queries
 - Support your data scientists with the IBM Data Science Experience (DSX) and integration with Apache Spark to build, train, deploy and manage AI models, and prepare and analyze data
-

IBM Hosted Analytics with Hortonworks

Simplify and modernize your existing data warehouse and analytics environment.

Data lakes have rapidly evolved and been adopted to augment, simplify and modernize existing data warehouse and analytic environments. Hadoop-based data lakes successfully support new technologies such as artificial intelligence (AI), Internet-of-Things (IoT), mobile and cloud that require a new level of data management and analytics. The challenge is finding an enterprise-ready, scalable, secure platform with the tools and services to meet growing analytic needs.

An integrated solution

IBM Hosted Analytics with Hortonworks provides an enterprise-ready, fully-provisioned data management and analytic environment designed to meet growing big data needs. Made up of the Hortonworks Data Platform (HDP), IBM® Db2® Big SQL (Big SQL) and IBM Data Science Experience (DSX), this solution includes hardware setup, provisioning, security, encryption and essential services. It is also hosted on the IBM Cloud.

Solution components

IBM Hosted Analytics with Hortonworks is ideal for organizations interested in moving to the cloud or Hadoop, advancing machine learning and real-time and ad-hoc queries for data scientists, or for offloading data warehouse workloads.

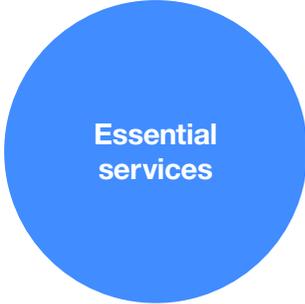


A leading SQL engine for Hadoop

Reduce cost while creating a performance data virtualization layer for federation and complex queries, so you can run data warehousing workloads on Hadoop.



**IBM Db2
Big SQL**



**Essential
services**

Essential services

Several services are included in the base offer:

- Support, troubleshooting and patching
- Changes to the environment configuration and authentication
- Directory synchronization changes

A leading data science platform

With the right tools, data scientists improve data productivity, collaboration and yield better analytic insights across the organization



**IBM data
science
experience**



**Subscription
services**

Subscription services

Additional add-on services are available such as:

- Patch application for IBM-deployed software.
- Design and coaching assistance in data migration
- Consulting services for data integration and migration from disparate sources.
- Directory integration and synchronization
- Set-up for specific networking requirements



Hortonworks data platform

HDP is the industry's only true secure, enterprise-ready open source Apache Hadoop distribution based on a centralized architecture (YARN).

Benefits of the IBM Hosted Analytics with Hortonworks solution

Each component of the IBM Hosted Analytics with Hortonworks solution offers unique benefits. Together, they create an ideal environment for your data management and analytics needs. Key capabilities include:

- **Cloud flexibility.** Avoid the cost of investing in and managing your hardware infrastructure by moving to the cloud - shifting from a CAPEX to an OPEX model. This solution is hosted on IBM Cloud (IBM Softlayer®) for optimal scalability but still allows you to maintain ownership of your data and install third party software. Multiple configurations are available, so you also have the flexibility to evolve the platform to accommodate new requirements.
- **Improved querying** Hortonworks Data Platform (HDP) is an Apache Hadoop-based, enterprise-grade, petabyte-scalable, open source platform for massively parallel processing (MPP). When this hosted analytics platform is combined with Big SQL, they enable virtualization and federation across your business. This, in turn, helps facilitate better insights through the ability to perform real-time and complex queries.
- **Better analytics** IBM Data Science Experience (DSX) helps users improve productivity and yield better analytic insights across the organization with a collaborative project space. Data scientists can choose from multiple tools such as R-Studio, Spark and Python as well as Jupyter and Zeppelin notebooks, using them to drive machine learning and harness valuable insights. In addition, DSX delivers an enterprise-ready administration console with security management, so your data has even greater protection.
- **Thoughtful integration** Db2 Big SQL lets you access, query and summarize data from any platform, including databases, data warehouses, NoSQL databases and more. It can concurrently exploit Apache Hive, HBase and Spark using a single database connection—even a single query. Moreover, it allows you to use pre-existing SQL skills on Hadoop and is compatible with Oracle SQL dialects for easier app migration.

Configuration options

IBM Hosted Analytics with Hortonworks offers the choice of three configurations so that you can select the option that best fits your needs. Optional VM, Bare Metal configurations and expansion with add-on services are available.

Multiple available configurations	Plan		Configuration	Nodes
VM for cost-efficiency	Development (non-HA, VM)	HDP	16 cores, 64GB, 4TB storage	6 HDP/Big SQL + 3 DSX + 1 IDM
		DSX	32 cores, 128GB, 4TB storage	• 2 Management
		Edge	16 cores, 64GB, 2TB storage	• 3 Data
		IDM	2 cores, 8GB, 100GB	• 1 Edge • 3 DSX (Dev) • 1 IDM
Bare Metal for high performance	Enterprise Performance (Bare Metal)	HDP (data intensive)	16cores, 128GB, 32TB storage	13 HDP/BIG SQL + 3 DSX + 2 IDM
		HDP (compute intensive)	36 cores, 384GB, 16TB storage	• 6 Management
		DSX	36 cores, 384GB, 4TB storage	• 5 Data
		Edge	16 cores, 64GB, 2TB storage	• 2 Edge
		IDM	4 cores, 16GB, 1TB	• 3 DSX • 2 IDM
Options to add nodes to base HDP and DSX configurations	Enterprise Performance (Bare Metal)	HDP (data intensive)	16cores, 128GB, 32TB storage	13 HDP/BIG SQL + 3 DSX + 2 IDM
		HDP (compute intensive)	36 cores, 384GB, 16TB storage	• 6 Management
		DSX	36 cores, 384GB, 4TB storage	• 5 Data
		Edge	16 cores, 64GB, 2TB storage	• 2 Edge
		IDM	4 cores, 16GB, 1TB	• 3 DSX • 2 IDM

Experience IBM Hosted Analytics with Hortonworks for yourself

To learn more about the IBM Hosted Analytics with Hortonworks solution and discover the impact it can have on your organization, visit [our website](#). You can view additional information and reach out to an IBM representative who can help guide you as you consider implementing the solution in your business.



©Copyright IBM Corporation 2018

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
June 2018

IBM, the IBM logo, [ibm.com](#), and [db2](#) are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [www.ibm.com/legal/copytrade.shtml](#)

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

All client examples cited or described are presented as illustrations of the manner in which some clients have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions. Contact IBM to see what we can do for you.

It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Statement of Good Security Practices: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful, comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Statements regarding IBM’s future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.



Please Recycle