

Got a big data headache?

Enterprise-grade Hadoop can ease the pain

Data scientists translate business problems into data analysis. In today's world of high-volume, fast-moving data and complex integration and architecture challenges, that's not easy. Here's how IBM® BigInsights™ for Apache Hadoop helps put data scientists' minds at ease with analytics and performance designed for enterprise-scale workloads.

Analytics

The IBM BigInsights Data Scientist module accelerates data science with advanced analytics to extract valuable insights from Hadoop. Big R statistical analysis and distributed frames allow data scientists to use the entire Hadoop cluster. Stable machine learning algorithms are optimized for Hadoop. Text analytics extract insight from unstructured data with existing tooling so analytic applications don't have to be developed from scratch.

Spark

BigInsights 4.0 includes Apache Spark 1.2.1, which is a key capability for advanced Hadoop and Data Scientist users. Spark helps data scientists to do in-memory distributed computation, driving dramatic performance increases. BigInsights 4.0 with Spark accelerates emerging capabilities for streaming, SQL, machine learning & graph processing. It simplifies developer experience, leveraging Java, Python & Scala languages.

Business insight

The BigInsights Analyst module lets data scientists use their existing skills to find data across the organization and visualize it without extra coding. IBM BigSheets is a spreadsheet-style data manipulation and visualization tool that gives business users direct access to data through a recognizable interface. IBM-designed Big SQL offers HDFS caching and high availability benefits as well as query optimization—without forcing data scientists to learn a new skillset.

Data accuracy Speed & scale

The IBM BigInsights Enterprise Management module helps ensure the scalability, performance and security of Hadoop clusters. For example, multi-tenant scheduling and multi-instance support enhance scalability and performance by allowing multiple installs of BigInsights on the same cluster with data isolation and resource sharing.



Ready to accelerate data science? You need BigInsights for Apache Hadoop.

Get the no-cost BigInsights Quick Start Edition at ibm.com/software/data/infosphere/biginsights/quick-start or check out IBM BigInsights on Cloud at ibm.com/software/products/en/ibm-biginsights-on-cloud. To learn more about Hadoop for the enterprise, visit ibm.com/hadoop