

Apache Hadoop 3 Improves Big Data Workloads

Platform Evolves to Meet Modern Requirements

WHITE
PAPER



Sponsored by:





Table of Contents

Apache Hadoop: Background and History	3
High-Performance Computing for Enterprise Machine Learning and Deep Learning	4
Flexibility and Agility	5
Enhanced Storage Options.....	6
Scalability and Availability.....	7
Workload and Resource Management	8
Evolving to Match Workloads.....	9
About Ventana Research.....	10



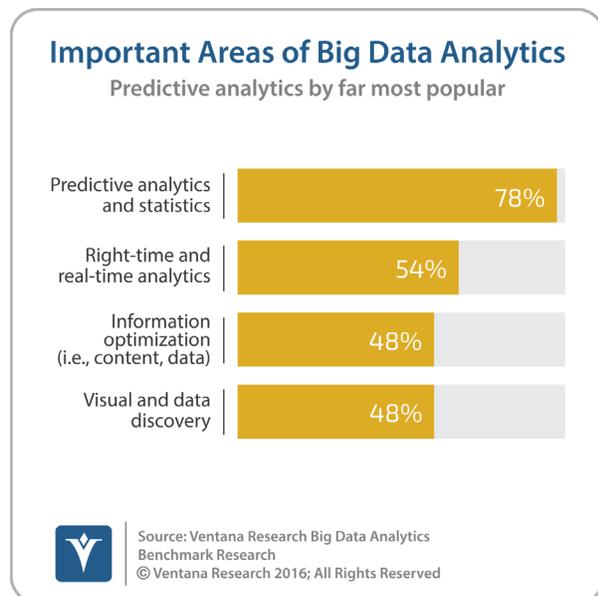
Apache Hadoop: Background and History

In the spring of 2006, the Apache Software Foundation released Hadoop, a distributed computing framework for managing and analyzing very large amounts of data in a scalable and reliable way. The open-source software was designed to run on clusters of servers ranging from a few nodes to thousands of nodes, allowing users to pool computing power to enable the processing of workloads far more cost-effectively than had been possible earlier.

It's been more than 10 years since the project began, and in that time an active open-source community has helped mature Apache Hadoop significantly. The community has contributed many enhancements including high availability, governance and analytical processing improvements.

The core of Hadoop consists of four modules: Hadoop Common, comprised of the common utilities that support the other Hadoop modules; the Hadoop Distributed File System (HDFS), which provides high-throughput access to application data; YARN, a framework for job scheduling and cluster resource management; and MapReduce, a YARN-based system for parallel processing of large data sets. Since the project's inception, other Hadoop-related Apache projects such as Apache HBase, Apache Hive, Apache Spark and many more have been developed, but this paper focuses on the core of Hadoop and key enhancements to its modules.

The size and variety of the workloads being handled by Hadoop continue to grow. Moreover, many of these workload categories are viewed as mission-



critical and require round-the-clock availability. In addition, the availability of this large data processing framework has enabled more sophisticated analyses such as artificial intelligence, machine learning and deep-learning algorithms – analyses that organizations increasingly consider to be vital. In our Big Data Analytics benchmark research, more than three-fourths of organizations (78%) said the most important type of big data analytics is the ability to apply the predictive analytics these algorithms yield.

Initially, Hadoop workloads consisted primarily of long-running batch-oriented tasks. Today's Hadoop workloads are typically a complex mix of longer-running tasks and shorter interactive queries that can be difficult to balance. Those who are performing interactive queries expect responses in seconds regardless of other processing being done. In the face of this evolving data management



situation, organizations today struggle to manage and reduce costs while meeting service level agreements.

High-Performance Computing for Enterprise Machine Learning and Deep Learning

Hadoop has enabled organizations to store and analyze large volumes of data, and the third major release, Hadoop 3, includes a number of improvements that address an array of workload-related issues. Applying Hadoop to the data science function is a common use case. The more data organizations gather and analyze, the more crucial it becomes that the data is processed efficiently and effectively. Currently, only sophisticated algorithms provide organizations with the means to maximize the value of the mass of data they collect. Data scientists use artificial intelligence, machine learning and deep-learning algorithms to sift through these mountains of data to find patterns that help organizations improve their operations.



The more data organizations gather and analyze, the more crucial it becomes that data is processed efficiently and effectively.

Of course, these algorithms require significant amounts of processing power. Fortunately, this processing today can be addressed using graphics processing units (GPUs), which provide far more processing power than CPUs. GPUs were originally designed to accelerate the creation and manipulation of computerized images. They have a highly parallel structure, which makes them well-suited for processing large blocks of data in parallel. While CPUs have tens of cores or fewer, GPUs typically have hundreds of cores.

GPUs' parallel processing capability significantly accelerates certain types of analyses, including those involving deep-learning networks. These networks require recursive analyses during the modeling process that can take days, weeks or even months to perform using CPUs. However, with GPUs the same analyses can be completed in days instead of months, hours instead of days and minutes instead of hours.

Hadoop 3 brings the power of GPUs to its workloads. Hadoop 2 does not support GPUs. Prior to the new release, if an organization wanted to take advantage of GPUs, the processing of workloads would need to be performed on separate clusters with GPU resources. There was no way to manage a mixed set of GPU and CPU resources.



This enhancement to resource-management capabilities means that GPU resources as well as pooling and isolation can be scheduled on the same cluster as other Hadoop workloads. GPUs are expensive; sharing and isolation enables organizations to apply GPUs to more workloads, helping to justify their cost.

These pooling and isolation capabilities have had an enormous impact on a number of industries. In the automotive industry, for example, the enterprise machine learning and deep learning models being used for self-driving cars require immense computation power. The cameras, radar systems and sensors on



**USE CASE:
AUTOMOTIVE
INDUSTRY**

The enterprise machine learning and deep learning models for self-driving cars require immense computation power.

a single car can generate a terabyte of data per hour. Thus, data collected from a fleet of 100 vehicles over a year to serve as a training data set can easily be hundreds of petabytes. The neural networks used to process this data require multiple iterations with various parameters to determine the best-fit model. In the absence of GPUs these calculations could take years to process.

Flexibility and Agility

Modern software architectures benefit significantly from container technology, which packages a complete runtime environment in a single unit or “container.” The container includes the operating system, the executable files and all the necessary libraries and dependencies. When applications are self-contained in this way, they can be built or

modified and rolled out within minutes, which results in faster time-to-market for new services.

Containers run independently, so they provide isolation from other programs and services. This allows organizations to be both agile and flexible; because each containerized service includes all dependencies, it is easy to introduce new services. Containerization also increases stability because these new services are less likely to interfere or conflict with existing services.

Hadoop 3 supports both Hadoop and non-Hadoop containerized workloads, which provides several benefits: MapReduce, Spark and Hive jobs can be run as containerized workloads, and testing and development of new versions of services don’t disrupt existing versions. Data scientists benefit from a containerized approach, which prevents conflicts among libraries such as R and Python and facilitates better management of workloads on the cluster.

For example, the highly dynamic digital advertising market demands fast-paced development and deployment of new models. Trending topics can arise at any



moment and alter the effectiveness of ad recommendations. Models often need to be updated in minutes. This is possible with a containerized approach because models can be deployed quickly and isolated from any existing models. The new models can be tested without impacting existing ones and then can be put into production once approved.

These new features in Hadoop 3 will also help support big data analytic products that include components running on the Hadoop cluster. In short, Hadoop 3 makes combining Hadoop and non-Hadoop workloads on the cluster easier, faster and safer.

Enhanced Storage Options

Massive amounts of data create inevitable storage challenges, but erasure coding, a new storage option in Hadoop 3, reduces typical storage requirements. As originally implemented, HDFS stores data in triplicate, providing redundancy so the system can continue to operate in the event of a failure. But tripled data storage is expensive. In erasure coding, similar to RAID 5 or 6 parity techniques, as data is stored across the cluster, parity checks are stored separately so they can be used to reconstruct the data at a later time. This method maintains the same level of data recoverability as triplicate storage but does so more efficiently, requiring half as much storage as Hadoop's original mechanism. The resulting savings can benefit both development and data science modeling workloads that need large amounts of historical data. Furthermore, erasure coding can be used

alongside other options, thus enabling storage tiering, or storing data for different types of workloads using different techniques.



USE CASE: HEALTHCARE

Healthcare providers can use storage tiering to deal with the government-mandated variety of data and record retention requirements they face.

However, erasure coding does require some additional overhead; because the data must be reconstructed from the encodings rather than simply read from another copy, it takes somewhat longer to rebuild. Organizations can continue to use triplicate storage for the most demanding or "hot" workloads and erasure coding for less frequently accessed or "cold" data. Organizations can use these techniques in combination to balance performance and storage costs.

For example, healthcare providers can use storage tiering to deal with the variety of data and record retention requirements they face. Although healthcare record requirements often vary by jurisdiction and government

agency, most require that data be retained for years. With storage tiering, organizations can minimize storage costs by using erasure coding to store older



data that is accessed less frequently. This sort of method can also be used in internet of things workloads that involve a mix of streaming event data accessed frequently when it is new and historical data that is accessed less frequently as it ages.

Furthermore, additional file connectors expand storage options. Our research shows organizations increasingly rely on the cloud for data storage, with a majority planning eventually to store data in the cloud. Hadoop 3 offers expanded cloud storage options. Earlier versions of Hadoop introduced support for Amazon’s S3 object storage as an alternative to HDFS; this third major release also includes Microsoft Azure Data Lake Storage as well as support for Alibaba Aliyun Object Storage System.

As more workloads move to the cloud, these additional options mean that organizations can take advantage of the storage mechanism that best suits the context and their needs. Our benchmark research on The Internet of Things finds that nearly four out of 10 (39%) IoT workloads rely on the cloud to collect event data, reinforcing the potential benefit of the cloud storage options that Hadoop 3 provides.

Scalability and Availability

Hadoop’s core value is its massive scalability, delivered via distributed computing and parallel processing. Hadoop 3 offers the ability to scale clusters to include thousands of nodes. In order to ensure high availability, distributed computing



**USE CASE:
MINING OR
PETROLEUM
INDUSTRY**

Continued operations mean preventing outages and saving millions of dollars in costs.

includes – and indeed requires – some measure of redundancy. Though most of Hadoop was designed for high availability via redundancy, the NameNode piece of HDFS lagged in this respect.

NameNode provides a directory of all files in a Hadoop cluster as well as their locations. NameNode originally was a single point of failure in a Hadoop configuration; Hadoop 2 expanded it to two nodes for better availability. In this release, enhancements to NameNode significantly improve scalability and availability. This release supports multiple standby NameNodes, further increasing availability and ensuring that even if one NameNode is down the cluster can continue to operate with standbys.

High availability is necessary in an array of mission-critical applications. If the clusters processing sensor data for predictive maintenance of large pieces of



equipment in the mining or petroleum industry fail, it could result in millions of dollars in costs associated with downtime.

Workload and Resource Management

The workloads associated with Hadoop present many resource-management challenges and potential conflicts. Often, there is conflict between long-running, resource-intensive jobs and shorter interactive queries, as well as between Hadoop and non-Hadoop workloads running on the same cluster and competing for the same resources.

Hadoop 3 includes enhancements that address these and other workload- and resource-management concerns. Hadoop's second major release introduced YARN, a global resource manager and scheduling service. In the new release, YARN supports opportunistic containers and distributed scheduling to properly schedule and balance resources for containerized workloads. Opportunistic containers are workloads that can be dispatched to a node for processing even when resources are not immediately available. The workload will be executed when resources become available. These containers are lower-priority than other workloads, and distributed scheduling guarantees that they are allocated across the cluster.

Managing distributed workloads across a cluster involves both inter-queue and intra-queue management. Hadoop 2 allowed only for inter-queue preemption, enabling reprioritization across queues. That means that once a job was assigned to a queue, users could not reprioritize work within the queue. Hadoop 3 includes intra-queue preemption for finer-grained control over prioritization. As a result, users can arrange jobs in a queue based on user- and application-based concerns, which means more control over the utilization and balancing of resources and workloads within the cluster.

In addition to opportunistic containers and intra-queue prioritization, Hadoop 3 includes enhancements to intra-node balancing, which allows the redistribution of data across disks within a node. Data is typically distributed evenly across disks, but when new disks are added to a node, distribution between old and new disks can become skewed. The new release also introduces YARN resource types, which enables first-class support for GPUs and containers. These resources also allow YARN to support and manage other types of resources in the future.

Taken together, these workload and resource management enhancements make for better service levels in mixed-workload environments such as data lakes combined with data marts. While data lakes frequently involve long-running batch processes for loading or uploading data, data marts typically involved successive interactive queries in which users expect rapid response times. Hadoop 3's new features allow users to prioritize the shorter data mart queries over the longer-running processes, interrupting them if necessary.

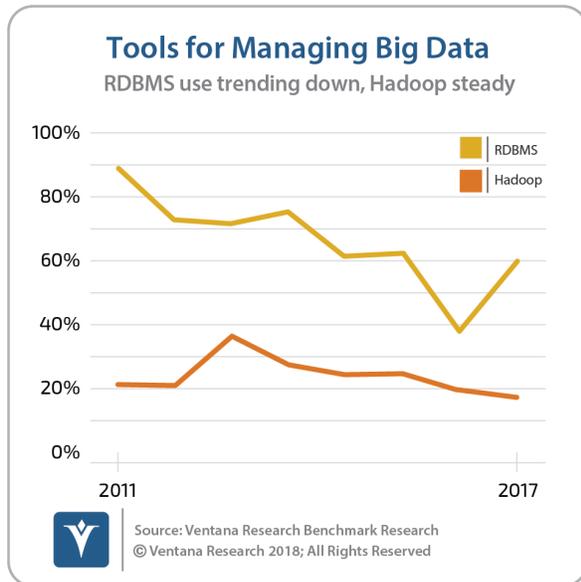


Evolving to Match Workloads

Hadoop continues to evolve to meet the requirements and challenges that accompany mission-critical, large-scale data processing. Because it is open-source, community needs and vision drive contributions and enhancements.

Organizations today require scalability and availability in these mission-critical systems. They need Hadoop to support enterprise machine learning and deep learning workloads. Organizations also need a variety of deployment options and

architectures, as well as efficient use of computing infrastructures in order to minimize the total cost of ownership.



According to our research, Hadoop use is firmly established in organizations. We have been studying the big data market since 2010 and have conducted numerous research projects in that time. We find that organizations' use of RDBMS for big data has declined significantly; we attribute that decline to an increasing reliance on Hadoop and other big data technologies.

These technologies continue to grow and mature, and improvements make them

even more reliable and compatible with enterprise information architectures. The enhancements included in Hadoop 3 provide significant additional value to any organization looking to deploy big data systems in mission-critical ways across a wider range of use cases.



About Ventana Research

Ventana Research is the most authoritative and respected benchmark business technology research and advisory services firm. We provide insight and expert guidance on mainstream and disruptive technologies through a unique set of research-based offerings including benchmark research and technology evaluation assessments, education workshops and our research and advisory services, Ventana On-Demand. Our unparalleled understanding of the role of technology in optimizing business processes and performance and our best practices guidance are rooted in our rigorous research-based benchmarking of people, processes, information and technology across business and IT functions in every industry. This benchmark research plus our market coverage and in-depth knowledge of hundreds of technology providers means we can deliver education and expertise to our clients to increase the value they derive from technology investments while reducing time, cost and risk.

Ventana Research provides the most comprehensive analyst and research coverage in the industry; business and IT professionals worldwide are members of our community and benefit from Ventana Research's insights, as do highly regarded media and association partners around the globe. Our views and analyses are distributed daily through blogs and social media channels including [Twitter](#), [Facebook](#) and [LinkedIn](#).

To learn how Ventana Research advances the maturity of organizations' use of information and technology through benchmark research, education and advisory services, visit www.ventanaresearch.com.