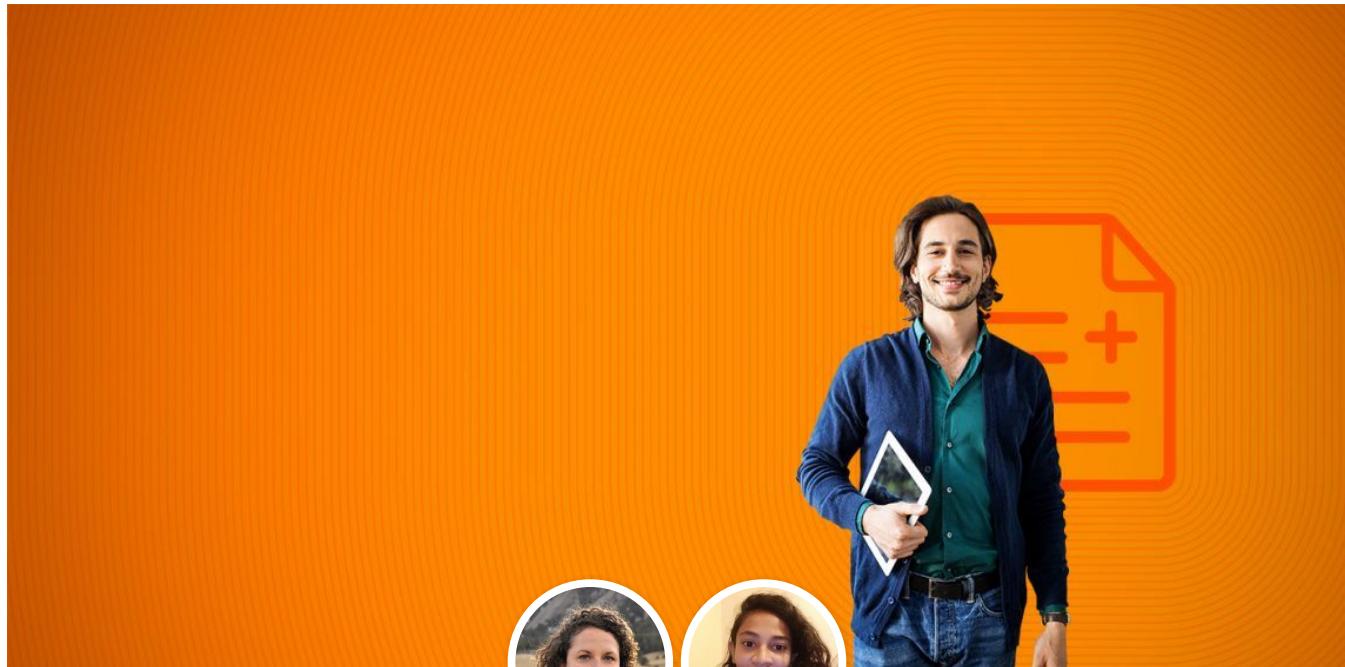


[Machine Learning & Big Data Blog](#)

# Hadoop vs Kubernetes: Will K8s & Cloud Native End Hadoop?



June 18, 2021

6 minute read

[Chrissy Kidd, Shanika Wickramasinghe](#)

Apache Hadoop is one of the leading solutions for distributed [data analytics](#) and [data storage](#). However, with the introduction of other distributed computing solutions directly aimed at data analytics and general computing needs, Hadoop's usefulness has been called into question.

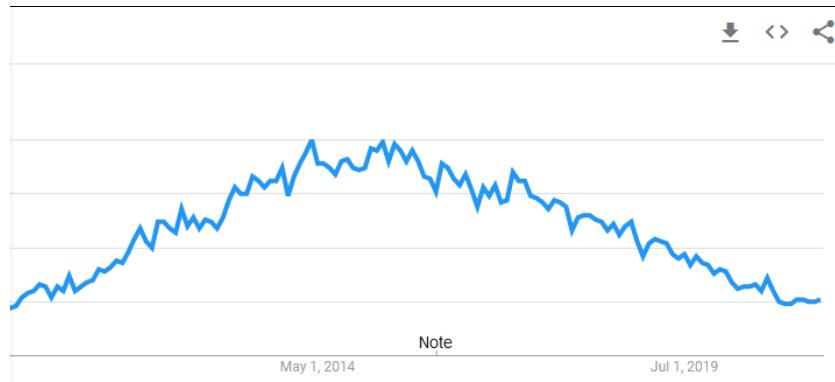
There are many debates on the internet: is Hadoop still relevant? Or, is it dead altogether?

In reality, Apache Hadoop is not dead, and many organizations are still using it as a robust dat A small blue upward-pointing arrow icon.

## Hadoop Guide

On Hadoop reached its **peak popularity** from 2014 to 2017. After that, we see a clear decline in searches for Hadoop. However, this alone is not a good indicator of the current landscape. After all, Hadoop can be integrated into a complete analytics solution.

## Beginning Kubernetes: Knowledge & Tutorials for Getting Started



In this article, we will learn more about Hadoop, its usability, and whether it will be replaced by rapidly evolving technologies like Kubernetes and Cloud-Native development.

[Download e-book >](#)

*(This article is part of our [Hadoop Guide](#). Use the right-hand menu to navigate.)*

# What is Hadoop?

**Hadoop** is an open-source framework that is used to store and process massive datasets efficiently. It is a reliable and scalable distributed computing platform that can be used on commodity hardware.

Hadoop distributes its data storage and analytics workloads across multiple nodes (computers) to handle the work parallelly. This leads to faster, highly efficient, and low-cost data analytics capabilities.

## Hadoop modules

Hadoop consists of four main modules that power its functionality:

**HDFS.** Hadoop Distributed File System is a file system that can run on low-end hardware while providing better throughput than traditional file systems. Additionally, it has built-in fault tolerance and the ability to handle large datasets.

**MapReduce** is a big data processing engine that supports the parallel processing of large data sets. It is the default processing engine available on Hadoop. Currently, however, Hadoop also provides support for other engines such as [Apache Tez](#).

## Beginning Kubernetes: Knowledge & Tutorials for Getting Started

common provides a common set of libraries that can be used in modules.

is behind the popularity of Apache Hadoop.

distributed computing model allows it to handle limitless

**Data Safety.** Hadoop automatically creates and manages data backups. So, you can simply recover your data from a backup in case of a failure.

**Cost.** Hadoop's ability to run on commodity hardware enables organizations to easily deploy a data analytics platform using it. It also eliminates the need for expensive and specialized hardware.

**Availability.** Hadoop is designed to handle failures at the application layer—which means it provides [high availability](#) without relying on hardware.

With its flexibility and scalability, Hadoop quickly gained the favor of both individual [data engineers/analysts](#) and corporations. This flexibility extends to types of data Hadoop can collect:

### Structured and unstructured

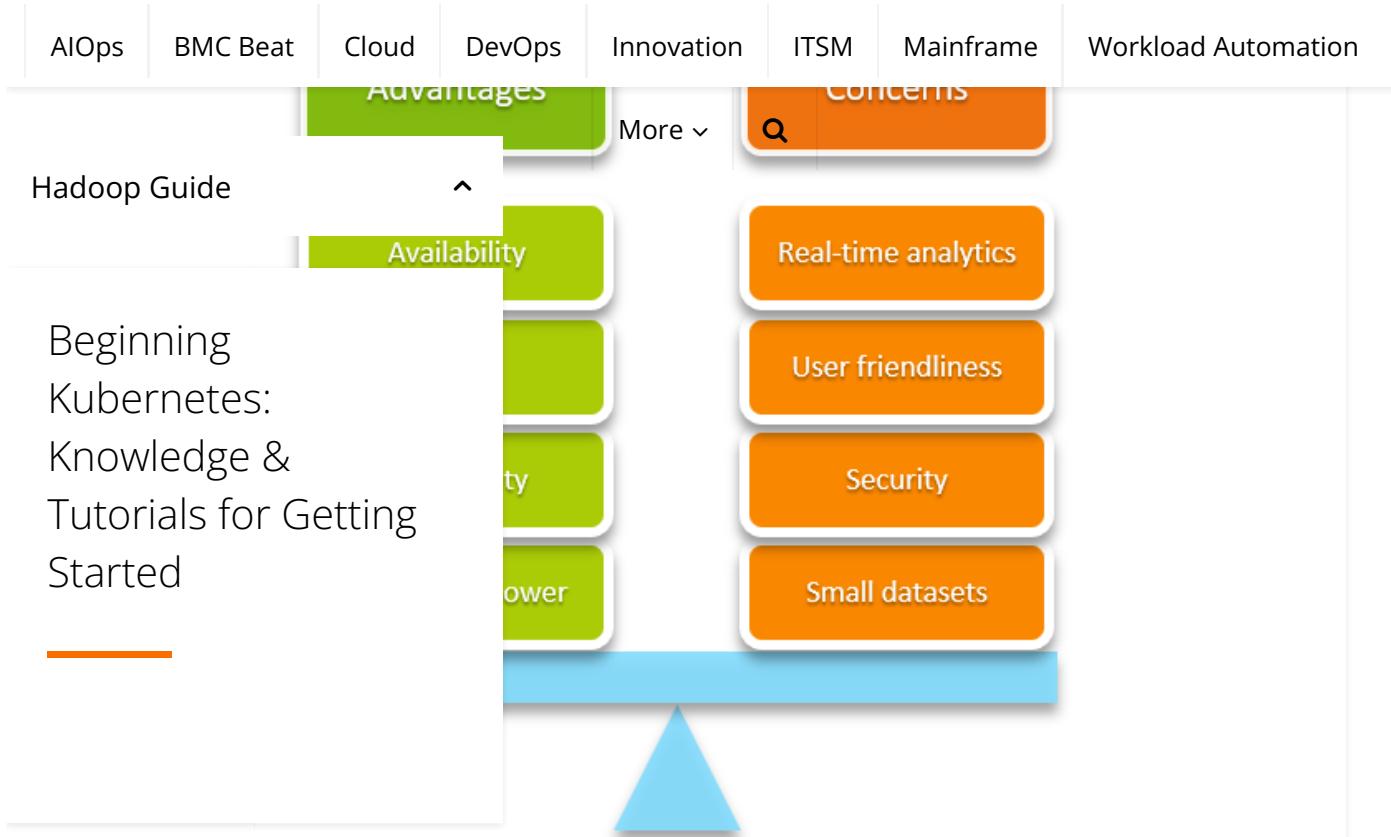
Through different inputs like social media, [streamed data](#), internal collections, etc.

Then Hadoop checks all these data sets and determines the usefulness of each data set. All this is done *without* having to go through the process of converting data into a single format.

Another feature that elevates Hadoop is its storage capability.

Once a large data set is accumulated, and the required data is extracted, we can simply store the unprocessed data with Hadoop endlessly. This enables users to reference older data easily, and the storage costs are also minimal since Hadoop is running on commodity hardware.





## Drawbacks of Hadoop

Apache Hadoop gained prominence thanks to all the above features.

However, as technology advances, new options have emerged, challenging Hadoop and even surpassing it in certain aspects. This, along with the inherent limitations of Hadoop, means it has indeed lost its market lead.

So, what are some drawbacks of Hadoop?

### Inefficient for small data sets

Hadoop is designed for processing big data composed of huge data sets. It is very inefficient when processing smaller data sets. Hadoop is not suited and cost-prohibitive when it comes to quick analytics of smaller data sets.

Another reason: Although Hadoop can combine, process, and transform data, it does not provide an easy way to **output** the necessary data. This limits the available options for business intelligence teams for visualizing and reporting on the processed data sets.

### Security concerns

hadoop which is a technology that is difficult to maintain by itself.

## Hadoop Guide

### Beginning Kubernetes: Knowledge & Tutorials for Getting Started

#### ESS

one of the leading programming languages with a large user base. It is the best language for data analytics, and it can be complex to learn.

Configurations and usage—the user must have thorough knowledge of the language to properly use and **debug** the cluster.

## Not suitable for real-time analytics

Hadoop is designed with excellent support for **batch processing**. However, with its limitations in processing **real-time data** and not providing native support for real-time analytics, Hadoop is ill-suited for real-time analytics.

[Download e-book](#)

## Hadoop alternatives

So, what other options to Hadoop are available? While there is no single solution to replace Hadoop outright, there are newer technologies that can reduce or eliminate the need for Hadoop.

## Apache Spark

Apache Spark is one solution, provided by **the Apache team** itself, to replace MapReduce, Hadoop's default data processing engine. Spark is the new data processing engine developed to address the limitations of MapReduce.

Apache claims that Spark is nearly 100 times faster than MapReduce and supports in-memory calculations. Moreover, it supports real-time processing by creating micro-batches of data and processing them.

The support of Spark for modern languages enables you to interact using your preferred **programming languages**. Spark offers excellent support for data analytics using languages such as Python, Java, and Scala.

# Beginning Kubernetes: Knowledge & Tutorials for Getting Started

• **Flink**. Flink is another processing engine with the same higher performance in some workloads as it is designed to handle unbounded and bounded data streams.

# Is Hadoop still relevant? & cloud-native replace

Even with newer and faster data process engines, Hadoop still limits users to its tools and technologies. It's better to use HDFS and HDFS with Java-based tools. But, what if you need to integrate other tools and technologies? It's best for your specific data storage and analytics needs?

The solution is using [Kubernetes](#) as the orchestration engine to manage your cluster.

With the ever-growing popularity of containerized cloud-native applications, Kubernetes has become the leading orchestration platform to manage any [containerized application](#). It offers features such as:

- Convenient management
- Networking
- Scaling
- High availability

*(Explore our [comprehensive Kubernetes Guide](#).)*

Consider this scenario: you want to move to cheap cloud storage options like [Amazon S3](#) buckets and managed data warehouses like [Amazon Redshift](#), [Google BigQuery](#), [Panoply](#). This is not possible with Hadoop.

Kubernetes, meanwhile, can easily plug them into Kubernetes clusters to be accessed by the

## Beginning Kubernetes: Knowledge & Tutorials for Getting Started

the Spark, Presto, Flink)  
(Jupyter, NLTK, TensorFlow, PyTorch, MATLAB)  
Kubernetes cluster  
many tools, frameworks, or programming languages you're  
's most suitable for your use case—you're no longer limited

*work together.)*

## Portability of Kubernetes

Another factor that uplifts Kubernetes is its portability. Kubernetes can be easily configured to be distributed across multiple locations and run on multiple cloud environments. With containers, users can easily move between development and production environments to facilitate data analytics in any location without major modifications.

By combining Kubernetes with rapid DevOps and [CI/CD pipelines](#), developers can easily create, test, and deploy data analytics, ML, and AI applications virtually anywhere.

## Support of Kubernetes for Serverless Computing

Kubernetes has further eliminated the need to manage infrastructure separately with the support for [serverless computing](#). Serverless computing is a rising technology where the cloud platform automatically manages and scales the hardware resources according to the needs of the application.

Some container-native, open-source, and function-as-a-service computing platforms like [fn](#), [Apache OpenWhisk](#), and [nuclio](#) can be easily integrated with Kubernetes to run serverless applications—eliminating the need for technologies like Hadoop.

Some frameworks, like nuclio, are specifically aimed at automating [data science pipelines](#) with serverless functions.

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

---

## How Hadoop handles large data sets cheaply

is also designed to address a specific need—handling large data sets efficiently on commodity hardware.

Cloud computing technologies have given rise to new requirements and use cases. Containerization technologies, like Kubernetes and serverless computing, offer new ways to manage and scale workloads.

As a result, it's important to identify and utilize the correct technology stack for your organization's needs.

## Related reading

[BMC Machine Learning & Big Data Blog](#)

[Data Architecture Explained: Components, Standards & Changing Architectures](#)

[3 Download e-book > \[Data Pipelines\]\(#\)](#)

[Snowflake 101: Intro to the Snowflake Data Cloud](#)

[Big Data: A Big Introduction](#)

[Data Ethics for Companies](#)



# Beginning Kubernetes: Knowledge & Tutorials for

Hadoop Guide



Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

take a deep  
ng platform



*These postings are my own and do not necessarily represent BMC's position, strategies, or opinion.*

See an error or have a suggestion? Please let us know by emailing [blogs@bmc.com](mailto:blogs@bmc.com).

#Hadoop

[Download e-book >](#)



## BMC Brings the A-Game

BMC works with 86% of the Forbes Global 50 and customers and partners around the world to create their future. With our history of innovation, industry-leading automation, operations, and service management solutions, combined with unmatched flexibility, we help organizations free up time and space to become an Autonomous Digital Enterprise that conquers the opportunities ahead.

[Learn more about BMC >](#)



[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

### You may also like

[Spark's Machine Learning Pipeline: An Introduction](#)[What Is a Neural Network? An Introduction with](#)[How to Add Subplots in Matplotlib](#)[What Is a Canonical Data Model? CDMs Explained](#)[Data Storage Explained: Data Lake vs Warehouse vs Database](#)[The Data Visualization Beginner's Guide](#)

### About the author

#### Chrissy Kidd

Chrissy Kidd is a writer and editor who makes sense of theories and new





Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

### About the author

kramasinghe

nasinghe is a software engineer by profession and a  
rmation Technology. Her specialties are Web and Mobile  
hanika considers writing the best medium to learn and  
edge. She is passionate about everything she does, loves  
to travel, and enjoys nature whenever she takes a break from her busy  
work schedule. You can connect with her on LinkedIn.

View

Download e-book ▾



[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

[Download e-book ▾](#)

[AIOps](#)[BMC Beat](#)[Cloud](#)[DevOps](#)[Innovation](#)[ITSM](#)[Mainframe](#)[Workload Automation](#)[More ▾](#)[Hadoop Guide](#)

Beginning  
Kubernetes:  
Knowledge &  
Tutorials for Getting  
Started

---

[Trust Center](#)[Accessibility Tools](#)[Cookie Consent](#)

©Copyright 2024 BMC Software, Inc. All rights reserved. BMC, the BMC logo, and other BMC marks are trademarks or registered trademarks of BMC Software, Inc. in the U.S. and other countries.

[Download e-book ▾](#)

Use of this site signifies your acceptance of BMC's **Terms of Use**. BMC, the BMC logo, and other BMC marks are trademarks or registered trademarks of BMC Software, Inc. in the U.S. and other countries.

