



Technical Report

## NetApp ONTAP SAN Solution for Hadoop

Karthikeyan Nagalingam (NetApp), Faiz Abidi (NetApp), Aditya Sirna (Hortonworks), and Archana Talamani (Hortonworks)  
Jan 2019 | TR-4697

In partnership with



### Abstract

Organizations collect and analyze large amounts of raw data. To harness the maximum value of this big data, they must transform their raw data into valuable business information. Apache Hadoop and its growing ecosystem of products enable organizations to extract valuable insights from large volumes of diverse data that cannot be analyzed with relational databases. The NetApp® ONTAP® SAN solution for Hadoop provides the high reliability and scalability required by enterprise organizations to ingest, protect, and manage that data.

## TABLE OF CONTENTS

<b>1</b>	<b>Challenges of Big Data Analytics</b>	<b>4</b>
<b>2</b>	<b>NetApp ONTAP SAN Solution for Hadoop</b>	<b>4</b>
2.1	ONTAP 9 Data Management Software	4
2.2	NetApp OnCommand Management Software	5
2.3	NetApp FlexClone Volumes	5
2.4	Hadoop Cluster	5
<b>3</b>	<b>NetApp ONTAP SAN Solution for Hadoop Benefits</b>	<b>6</b>
3.1	Customers Who Will Benefit from This Solution	7
<b>4</b>	<b>NetApp ONTAP SAN Solution for Hadoop Overview</b>	<b>7</b>
4.1	Components	8
4.2	Data Protection and Job Completion	8
4.3	Performance	10
<b>5</b>	<b>Hortonworks Overview</b>	<b>10</b>
5.1	Hortonworks Data Platform Overview	10
<b>6</b>	<b>Hortonworks HDP v3.0.0 certified</b>	<b>11</b>
<b>7</b>	<b>Solution Architecture</b>	<b>12</b>
7.1	Hardware Requirements	13
7.2	Software Requirements	13
<b>8</b>	<b>Performance Testing</b>	<b>16</b>
8.1	Tests Conducted	16
	<b>Conclusion</b>	<b>18</b>
	<b>Where to Find Additional Information</b>	<b>18</b>
	<b>Acknowledgments</b>	<b>19</b>
	<b>Version History</b>	<b>19</b>

## LIST OF TABLES

Table 1)	HDP v3.0.0 verified components	12
Table 2)	Hardware components	13
Table 3)	Software components	13

## LIST OF FIGURES

Figure 1)	NetApp ONTAP SAN solution for Hadoop	4
-----------	--------------------------------------	---

Figure 2) NetApp ONTAP SAN solution for Hadoop overview. ....	8
Figure 3) Cabling topology. ....	14
Figure 4) Storage architecture.....	15
Figure 5) Duplicate Hadoop cluster using NetApp FlexClone technology. ....	15
Figure 6) Performance results. ....	16

# 1 Challenges of Big Data Analytics

Big data requires big analytics, big bandwidth, and big content technologies. Organizations leverage big data by collecting and analyzing large amounts of raw data, such as point-of-sale data, credit card transactions, log files, and machine and security data. To harness the maximum value of big data, organizations must transform their raw data into valuable business information. This essential business driver requires tools that can process very large amounts of both structured and unstructured data. It also requires data management software that is capable of hybrid web-scale deployments, is highly available and resilient, and provides the data management capabilities that enterprises require.

Apache Hadoop and its growing ecosystem of products enable organizations to extract valuable insights from large volumes of diverse data that cannot be analyzed with relational databases. With these insights, people across the organization can ask the right questions and get better answers, supporting more informed decisions that help promote business transformation.

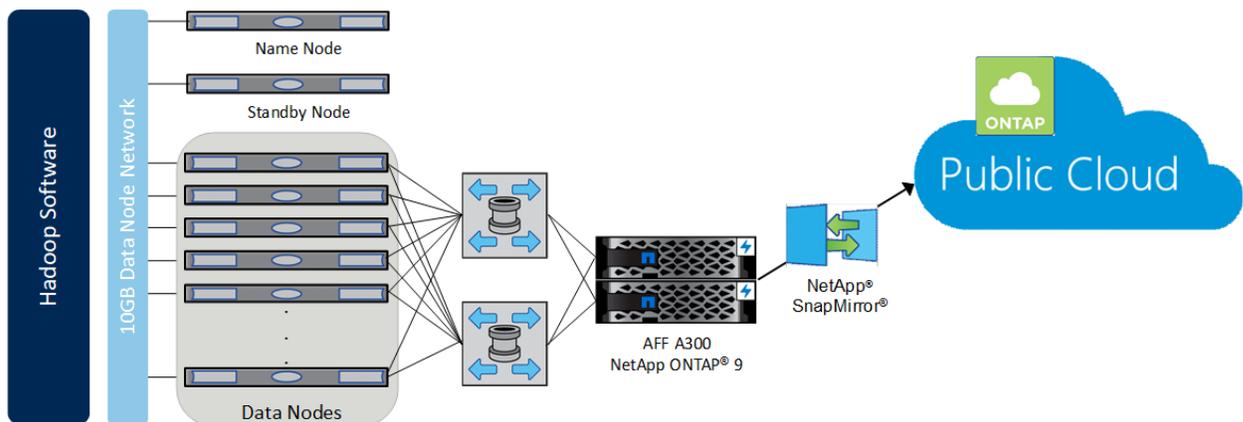
However, because initial Apache Hadoop/Spark deployments often rely on commodity servers with internal drives, infrastructure resilience and agility issues prevent organizations from realizing the full benefits of these deployments. For example, a single disk failure can degrade performance of the entire cluster. Managing disk replacements is continual and error-prone. In addition, triple file replication and failure redistribution models increase network costs and complexity.

For many enterprise customers, Hadoop is a shared-nothing, massively parallel data-processing platform that emphasizes cost effectiveness and replication of data availability. Most enterprises have used traditional databases to manage their rapidly growing need to process data and extract business value. In some cases, however, the costs of these solutions have been so high that they cannot afford to either house or process the data. In other cases, the solutions are not suitable for new data types such as unstructured data or key value pairs.

# 2 NetApp ONTAP SAN Solution for Hadoop

The NetApp ONTAP SAN solution for Hadoop, shown in Figure 1, features enterprise storage that is independent of the compute servers to offer an enterprise-class deployment with lower cluster downtime, higher data availability, and linear scalability for handling data and business growth.

Figure 1) NetApp ONTAP SAN solution for Hadoop.



## 2.1 ONTAP 9 Data Management Software

This solution leverages the robust and proven features of NetApp ONTAP 9 software, the enterprise data management software that powers the NetApp engineered systems of AFF and FAS, as well as software-

only NetApp ONTAP Cloud and NetApp ONTAP Select. ONTAP supports independent scaling of compute and storage resources, reducing costs and maximizing resource utilization. In addition, ONTAP provides integrated data protection to safeguard operations with near-instant backup and recovery using the highly efficient NetApp Snapshot™, SnapRestore®, and SnapCenter® technologies.

## 2.2 NetApp OnCommand Management Software

ONTAP provides a complete and powerful set of tools to manage the enterprise data infrastructure. The NetApp OnCommand® management software family covers the entire range of data management activities. NetApp OnCommand Unified Manager provides a single-pane view of the entire data infrastructure, no matter where the data is located. It allows enterprises to easily monitor the health, availability, capacity, performance, and data protection status of their NetApp clustered storage and provides alerts and vital information for proactive management. OnCommand Unified Manager is tightly integrated, delivers comprehensive storage performance monitoring, and assists in proactive management.

## 2.3 NetApp FlexClone Volumes

FlexClone® volumes are space-efficient, writable data copies that can be created almost instantly, anywhere in the Data Fabric or hybrid cloud where a Snapshot copy exists. FC and iSCSI are supported protocols. FlexClone technology makes it possible for developers, QA engineers, and software testers to work with real production data. FlexClone volumes can be easily created and deleted, and any changes made to them have zero impact on the parent production data. If necessary, they can even be split from the parent production data and promoted to production. FlexClone volumes can be considered as free data copies.

As shown in Figure 5, the NetApp ONTAP SAN solution for Hadoop duplicates the Hadoop cluster by using NetApp FlexClone technology.

## 2.4 Hadoop Cluster

A traditional Hadoop cluster consists of a few basic components:

- A NameNode, which manages the Hadoop Distributed File System (HDFS) namespace.
- A secondary NameNode, more accurately referred to as the checkpoint node, which manages the on-disk representation of the NameNode metadata.
- A ResourceManager node, which manages all jobs submitted to the Hadoop cluster, facilitates job task scheduling, and provides security.
- Hundreds of slave nodes, which provide both NodeManager and DataNode functionality. These nodes perform all of the real work done by the cluster. As DataNodes, they store all of the data blocks that make up the file system, and they serve I/O requests. As NodeManagers, they perform the job tasks assigned to them by the ResourceManager.

HDFS is the default cluster in most of the Hadoop distributions, and it provides data storage for data analytics as well as I/O operations. HDFS allows every DataNode to access analytics data across the cluster. The NameNode houses and manages the metadata, so whenever a DataNode must read or write an HDFS block, the NameNode informs the DataNode where a block exists or where one should be written.

## Hadoop Distributed File System

HDFS is a feature-rich, integrated file system that offers cluster-aware data protection, fault tolerance, and the ability to balance workloads. HDFS maintains multiple copies of data to provide data integrity from disk or node failures. The HDFS mirroring scheme is set on a file-by-file basis. The standard or default practice in Hadoop is to set all data files to a replication count of three. Another innovation of

HDFS is the ease with which the replication or mirroring level can be adjusted. To change the replication count of a file, you simply issue the `hadoop fs -setrep` command.

In addition to data protection, HDFS mirroring makes sure that MapReduce jobs complete successfully, even during hardware failure, providing data high availability (HA).

### Triple-Mirror Replication Trade-Offs

Although HDFS is distributed, feature rich, and flexible, both throughput and operational costs are associated with server-based replication for ingest, redistribution of data following recovery of a failed disk or DataNode, and space utilization.

The costs of triple replication include:

- For each usable terabyte (TB) of data, 3TB of raw capacity are required.
- Server-based triple replication creates a significant load on the servers. In simple terms, at the beginning of an ingest task, before any data blocks are written to a DataNode, the NameNode creates a list of DataNodes on which all three of the first block replicas will be written. This process forms a pipeline and allocates blocks to those nodes. The first data block is written out to drives. That DataNode forwards the same block to the next node in the pipeline, and the preceding procedure is repeated. Finally, the second DataNode sends the block to the third DataNode, and the write to storage follows the same I/O path as previous writes. This process results in a significant load on server resources.
- Server-based triple replication also creates a significant load on the network. For ingest, three network trips are required to accomplish a single block write. All compete for and consume the same network, CPU, memory, and storage resources allocated to process the Hadoop analysis tasks. The NetApp ONTAP SAN solution for Hadoop offers a significantly improved alternative.

## 3 NetApp ONTAP SAN Solution for Hadoop Benefits

The core benefit of the NetApp ONTAP SAN solution for Hadoop is an enterprise-grade storage platform for implementing a data lake built on Hadoop that provides everything customers expect from NetApp and a few bonuses. These features include:

- Ready-to-deploy Hadoop solution
- Only two copies of data required while providing high data availability
- Linear scalability for handling data and business growth
- Transparent RAID operations
- Transparent rebuild of failed media
- Fully online serviceability
- Faster deployment with network boot
- FlexClone volumes that can be created almost instantly
- Performance for the next generation of DataNodes (SMP) and storage networks (FC or iSCSI)
- Proven, tested, supported, and fully compatible system design
- Proven storage hardware based on enterprise-grade designs, offering enterprise-grade performance, uptime, AFF with end-to-end NVMe, and support for 60% or more Hadoop workloads
- Lower total cost of ownership (TCO)
- Creation of a data pipeline from edge to core to full cloud integration

### 3.1 Customers Who Will Benefit from This Solution

Vast amounts of data represent both an opportunity and a threat to enterprises. The NetApp ONTAP SAN solution for Hadoop is the first solution that is affordable enough to reduce the threat of drowning in data while still realizing the business potential that might be discovered in the data.

The customers who will benefit from this solution are enterprises that have experienced frustration operating a Hadoop cluster and are new to using open-source computing for business-critical applications. Hadoop application throughput scales proportionately with the number of DataNodes (NodeManagers). However, as clusters of DataNodes grow, so does the operational expense (opex) burden of running these larger clusters. Opex must scale with throughput.

For customers who are new to Hadoop, there is a steep learning curve. Very few enterprise applications are built to run on massively parallel clusters, so there is much to learn. However, the NetApp ONTAP SAN solution for Hadoop provides an operational model for a Hadoop cluster that, after it is set up, does not require much additional attention. The cluster is more stable and is easier to maintain, and it allows customers to concentrate on getting answers to their business needs. This solution flattens the operational learning curve of Hadoop.

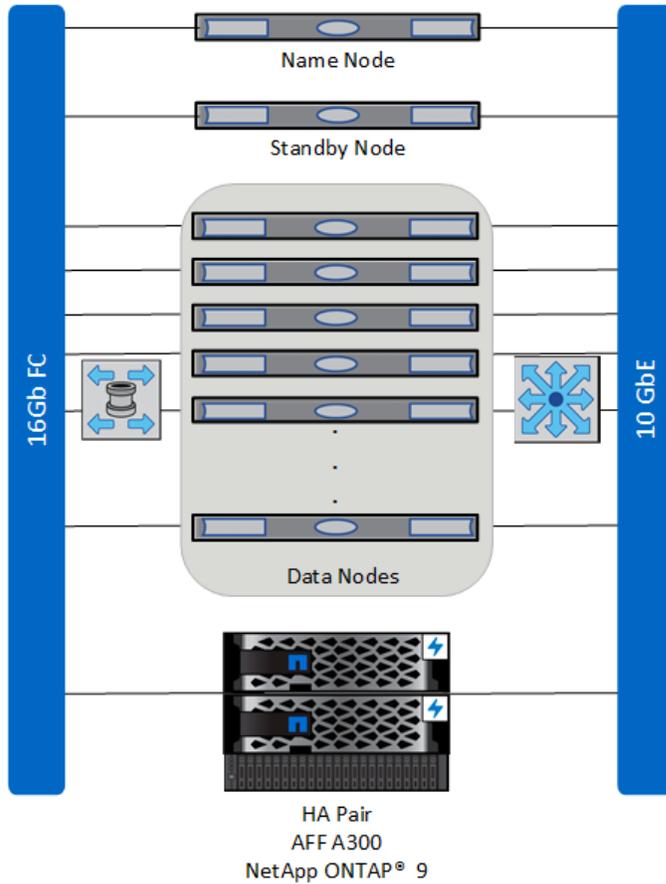
## 4 NetApp ONTAP SAN Solution for Hadoop Overview

This section provides an overview of the NetApp ONTAP SAN solution for Hadoop and covers the following information:

- Solution components
- Data protection and job completion
- Use cases with a replication count of two
- Performance

Figure 2 depicts an overview of the NetApp ONTAP SAN solution for Hadoop.

Figure 2) NetApp ONTAP SAN solution for Hadoop overview.



## 4.1 Components

The solution has the following hardware components:

- NetApp AFF A300 storage HA pair with SSD shelf
- 2 x fabric switches for the FC protocol
- 2 x management switches for the iSCSI protocol
- Minimum of 4 worker nodes and 2 master nodes

See Table 1 for a detailed list of the solution's hardware components.

The solution has the following software components:

- NetApp ONTAP data management software
- Apache Hadoop distribution such as Cloudera, Hortonworks, MapR, Apache Hadoop, and Apache Spark

See Table 2 for a detailed list of software components.

## 4.2 Data Protection and Job Completion

By default, with HDFS, for every block of data loaded into HDFS, two copies are written to two other nodes in the cluster. This replication count of three is the standard practice. Traditionally, Hadoop

deployments use the standard “just a bunch of disks” (JBOD) configuration of SATA DAS disks colocated in a DataNode for storage. HDFS is one of the first file systems to implement a replication strategy on a file-by-file basis. This flexibility allows customers to adjust the protection level to suit particular datasets.

HDFS replication provides data protection, job completion in the event of disk failure or DataNode failure, and load balancing. Replication for an entire cluster can be set by using a single parameter, `dfs.replication`, which has a default value of three. The higher the HDFS replication count, the more copies of the data must be stored, but the more resilient and balanced the access is to that data.

Another important reason for having replicated copies is for job completion. If job tasks on a given DataNode fail for any reason, the ResourceManager must reassign the failed tasks to other nodes in the cluster. Eventually, the ResourceManager might blacklist the node, and the NameNode might recognize the node as being in a failed state. If task reassignment is not possible, then the MapReduce job cannot complete.

However, when a customer deploys with the default replication count of three, for every 1PB of data loaded into a Hadoop cluster, another 2PB must be copied across the network fabric to the other nodes, and then that data must be written to storage. High-replication-count files consume a lot of space and take a long time to ingest because 2PB of the 3PB of data must be written over the network.

## Reduced Overhead

The NetApp ONTAP SAN solution for Hadoop is a more efficient solution that significantly reduces the overhead of data protection, making HDFS more efficient and faster (certainly on data ingest). Many big users of Hadoop already take the risk of loading data at a replication count of one because it is faster. Then they either leave the count at that level or later increase it to a safer level (for example, two) to make sure that node outages do not impede job completion or result in data loss.

Offloading some or even all of the data protection aspects of HDFS onto enterprise-grade, industry-proven ONTAP storage controllers is the single most important feature of the NetApp ONTAP SAN solution for Hadoop.

The NetApp ONTAP SAN solution for Hadoop enables customers to safely run with a replication count of two, because the data is protected by using NetApp RAID DP® on the aggregates that are used for data storage. With RAID DP, as many as two SSDs (or disks) can fail in a single aggregate without any loss of data or interruption in data availability.

## Higher Replication Count

The MapReduce job performance in some use cases might benefit from a replication count higher than two, especially when multiple jobs use the same data run simultaneously. In those cases, higher replication counts might offer greater parallelization of job tasks while leveraging HDFS data locality. Higher replication counts also enable organizations to better leverage the Hadoop speculative execution feature, which can help balance task execution across resources in a cluster. Speculative execution enables the ResourceManager to assign instances of a relatively slow-running task to other NodeManager nodes. In this scenario, when the first instance of the task finishes, the other instance is killed, and its results are discarded. This feature can benefit MapReduce performance in clusters in which unused node resources exist, but it might be bad for performance in very busy clusters. In these cases, higher replication counts are supported by the NetApp ONTAP SAN solution for Hadoop.

## Replication Count of Two

Given all the pros and cons of various replication counts discussed previously regarding the NetApp ONTAP SAN solution for Hadoop, a replication count of three is no longer a requirement for data protection. We recommend a replication count of two, and that value represents a good balance between performance and data protection for a very large number of Hadoop use cases. As with all database and

analytics systems, MapReduce jobs often benefit from tuning at both the system and job code levels. The replication count is one factor that can be adjusted as needed.

### 4.3 Performance

Hadoop was designed to support workloads that are dominated by reads, and those reads are supposed to be very large: 128MB by default. This fact is true on a logical level, but not on a physical level. Hadoop has a performance profile that is very similar to that of an Oracle database running over NFS.

MapReduce jobs do not process I/O directly, but request the I/O from another kernel process (the BlockReader). This process either reads from local disk or issues a read request over the HDFS cluster fabric to retrieve the block from another node. The list of candidate nodes is provided by the NameNode.

The process is critical to understanding the value that the NetApp SAN solution for Hadoop brings to Hadoop performance. Hadoop depends more on IOPS than on streaming I/O bandwidth. From an I/O transaction perspective, a 64K I/O is not much more work than an 8K I/O. Hadoop is sensitive to the IOPS throughput, and a standard Hadoop JBOD layout puts a single file system on a single SATA disk. This file system contains HDFS blocks, which are mostly simple 128MB files. The striping achieved by MapReduce occurs when several jobs access several file systems at once. At any point, a given MapReduce job must read 64KB at a time from a single 7.2K RPM JBOD disk. The IOPS rate of this single disk limits the throughput of that job.

## 5 Hortonworks Overview

In 2011, Rob Bearden partnered with Yahoo to establish Hortonworks with 24 engineers from the original Hadoop team, including founders Alan Gates, Arun Murthy, Devaraj Das, Mahadev Konar, Owen O'Malley, Sanjay Radia, and Suresh Srinivas.<sup>1</sup> Under the leadership of Arun Murthy, Chief Products Officer, this core product team is enriched with enterprise software talent from the likes of Oracle, IBM, HP, VMware, and others to help make HDP and HDF meet the enterprise-grade requirements you expect. Hortonworks is headquartered in Santa Clara, California, and the business model is based on open-source software support subscriptions, services, solutions, training and consulting services. Hortonworks operates in 19 countries with approximately 1,110 employees. There are 29 out of 114 Apache Hadoop committers from Hortonworks and 208 committer seats across 20+ Apache projects, and they focus on the data access, security, operations, and the governance needs of the enterprise Hadoop market.

Hortonworks is a leading innovator in the industry, creating, distributing, and supporting, enterprise-ready open data platforms and modern data applications. Its mission is to manage the world's data. It has a single-minded focus on driving innovation in open-source communities such as Apache Hadoop, NiFi, and Spark. Hortonworks, along with its 1,600+ partners, provides the expertise, training, and services that allow you to unlock transformational value for your organizations across lines of business. Hortonworks's connected data platforms power modern data applications that deliver actionable intelligence from all data: data in motion and data at rest.

### 5.1 Hortonworks Data Platform Overview

HDP is the industry-secure, enterprise-ready open-source Apache Hadoop distribution based on centralized architecture (YARN). HDP addresses the complete needs of data at rest, powers real-time customer applications, and delivers robust big data analytics that accelerate decision making and innovation.

---

<sup>1</sup> Source: [www.hortonworks.com](http://www.hortonworks.com)

HDP products are divided into six categories: data management, data access, data governance and integration, security, operations, and cloud.

## Data Management

YARN and HDFS are the cornerstone components of HDP for data at rest. Whereas HDFS provides the scalable, fault-tolerant, cost-efficient storage for your big data lake, YARN provides the centralized architecture that enables you to process multiple workloads simultaneously. YARN provides the resource management and pluggable architecture for enabling a wide variety of data access methods.

## Data Access

HDP includes a versatile range of processing engines that empower you to interact with the same data in multiple ways at the same time. Therefore, applications for big data analytics can interact with the data in the best way, from batch to interactive SQL or low latency access with NoSQL. Emerging use cases for data science, search, and streaming are also supported with Apache Spark, Storm, and Kafka.

## Data Governance and Integration

HDP extends data access and management with powerful tools for data governance and integration. These tools provide a reliable, repeatable, and simple framework for managing the flow of data in and out of Hadoop. This control structure, along with a toolset to ease and automate the application of schemas or metadata on sources, is critical for successful integration of Hadoop into your modern data architecture. Hadoop has engineering relationships with many leading data management providers to enable their tools to work and integrate with HDP. Data governance and integration are provided by Atlas, Falcon, Oozie, Sqoop, Flume, and Kafka components.

## Security

Security is woven and integrated into HDP in multiple layers. Critical features for authentication, authorization, accountability, and data protection are in place to help secure HDP across these key requirements. Consistent with this approach throughout all of the enterprise Hadoop capabilities, HDP also enables you to integrate and extend your current security solutions to provide a single, consistent, secure umbrella over your modern data architecture. The authentication, authorization, and data protection are provided by Knox and Ranger components.

## Operations

Operations teams deploy, monitor, and manage a Hadoop cluster within their broader enterprise data ecosystem. Apache Ambari simplifies this experience. Ambari is an open-source management platform for provisioning, managing, monitoring, and securing the HDP. It enables Hadoop to fit seamlessly into your enterprise environment. Ambari and Zookeeper manage a Hadoop cluster.

## Cloud

Cloudbreak, as part of HDP and powered by Ambari allows simplified provisioning and Hadoop cluster management in any cloud environment including; Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform, and OpenStack. It optimizes your use of cloud resources as workloads change.

## 6 Hortonworks HDP v3.0.0 certified

NetApp ONTAP SAN is a certified solution with HDP v3.0.0. To obtain certification, we ran the QATS test suite in both a non-secured and a secured cluster that included the components as shown in Table 1.

Table 1) HDP v3.0.0 verified components

Component	Status
Accumulo	Passed
Atlas	Passed
HDFS	Passed
Druid	Passed
HBase	Passed
Hive (LLAP included)	Passed
Kafka	Passed
Knox	Passed
MapReduce2	Passed
Oozie	Passed
Phoenix	Passed
Pig	Passed
Ranger	Passed
Spark2 (hive included)	Passed
Storm	Passed
Sqoop	Passed
Superset	Passed
Tez	Passed
Yarn	Passed
ZooKeeper	Passed

The above testing was done in a six-node HDP cluster using NetApp AFF8080 and FC protocol, but other AFF product and storage protocols such as FCoE and iSCSI are also supported.

## 7 Solution Architecture

This section describes the following components of the NetApp ONTAP SAN solution for Hadoop's architecture:

- Hardware requirements
- Software requirements
- Cabling technology
- Storage architecture

## 7.1 Hardware Requirements

Table 1 lists the hardware components for the NetApp ONTAP SAN solution for Hadoop.

Table 2) Hardware components.

Component	Product or Solution	Details
Storage	NetApp AFF A300 storage array with ONTAP 9.3 (tested) or later	<ul style="list-style-type: none"> <li>• 2 controllers (HA pair) with 24 x 900GB, SSD</li> <li>• 1 hot spare per disk shelf</li> <li>• 1 data aggregate (23 drives shared to both controllers) per controller</li> <li>• Test cases included 2 aggregates and 2 volumes</li> <li>• 2 x 4 x 16Gb FC protocol</li> <li>• RAID DP</li> </ul>
Servers	PRIMERGY RX2540 M1 Intel Xeon CPU E5-2670 v3 at 2.30GHz	<ul style="list-style-type: none"> <li>• 10 servers each with two 2.4GHz (6-core) or 2.3GHz (8-core) processors (40 cores)</li> <li>• 12 DIMMs (up to 192GB), up to 1600MHz (512GB/node)</li> <li>• 1 x 1Gbps Ethernet ports, 2 x 16GbE FC ports</li> </ul>
Networking	<ul style="list-style-type: none"> <li>• 10GbE nonblocking network switch</li> <li>• 1GbE network switch</li> </ul>	A Cisco Nexus 5000 switch was used for testing. Any compatible 10GbE network switch can also be used.
FC switch	16GbE switch	Brocade 5000 switches were used with a multipath configuration.

## 7.2 Software Requirements

Table 2 lists the software components for the NetApp ONTAP SAN solution for Hadoop.

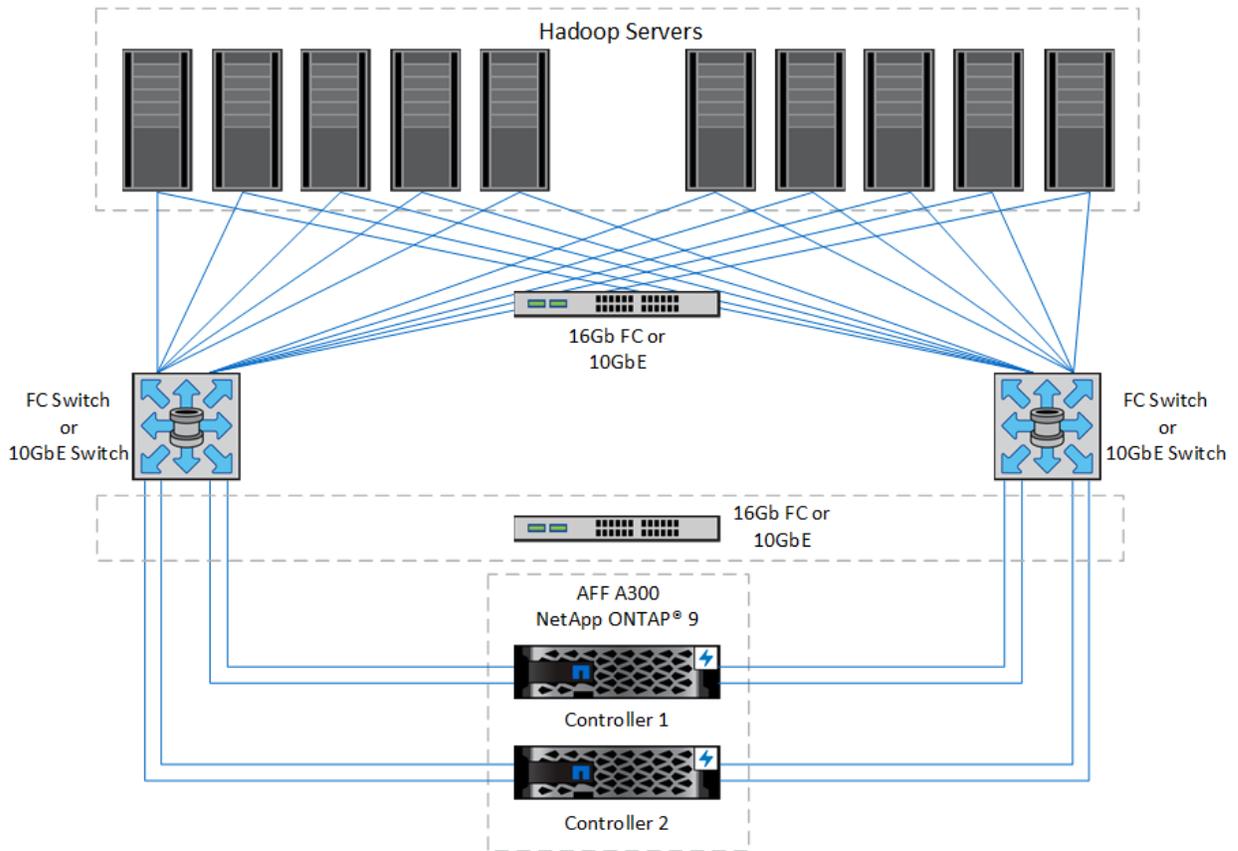
Table 3) Software components.

Component	Product or Solution	Details
Server operating system	Red Hat Enterprise Linux server supported by Hadoop distributors. In our case, we used 7.4 (x86_64) or later.	Hadoop typically requires a Linux distribution.
Hadoop distribution used in the testing	Cloudera Distribution for Hadoop	Cloudera Manager 5.14
	Hortonworks Data Platform 2.6.5 (tested)	Apache Ambari 2.6

## Cabling Topology

In this setup, we connected all the worker nodes and master nodes through 16Gb FC switches as well as 10GbE connections to the NetApp AFF A300 storage controller. We used the 10GbE network connection for iSCSI traffic and Hadoop cluster communication.

Figure 3) Cabling topology.



## Storage Architecture

In our validation of the NetApp ONTAP SAN solution for Hadoop, shown in Figure 4, we used 24 SSDs, which are shared to two aggregates. Each controller has one aggregate. We created eight volumes from each aggregate, and each volume has one LUN. We mapped one LUN from each controller to each worker node (for example, LUN 1 and LUN 9). These LUNs are mapped to the Hadoop worker node on two mount points, and then we created an `xfs` file system on top of them. The `xfs` file system is provisioned to HDFS in the Hadoop/Spark cluster through the `dfs.datadir` parameter.

To duplicate the Hadoop cluster, we provided NetApp SAN LUNs for both the master and worker nodes. In the master nodes, we used NetApp storage LUNs to hold the NameNode, standby node images, and edit logs. In the worker nodes, we kept all of the HDFS blocks in NetApp SAN LUNs. The operating system and temp files are kept in the local hard disks. For the testing, we used the same Hadoop cluster to access the duplicated (cloned) volumes and changed the original mount point to the cloned volume mounted mount points. NetApp FlexClone technology makes the duplication of the Hadoop cluster very quick regardless of the size of the Hadoop dataset. The Hadoop updated datasets consume only additional disk space. Both the original and duplicated Hadoop cluster use the same data while reading the data for the Hadoop jobs.

Figure 4) Storage architecture.

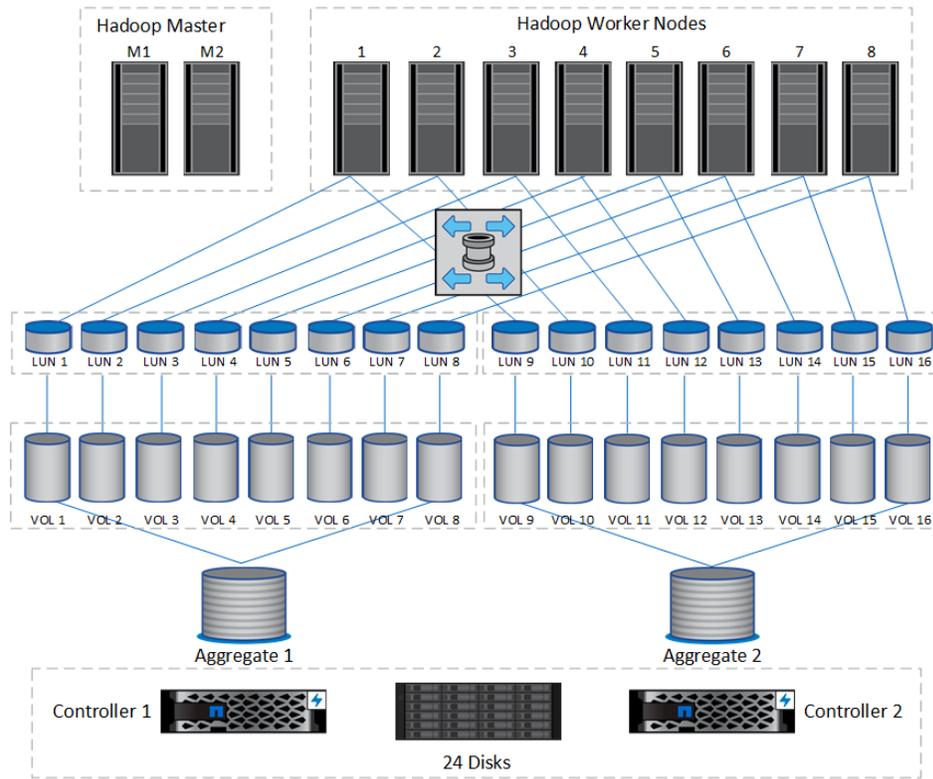
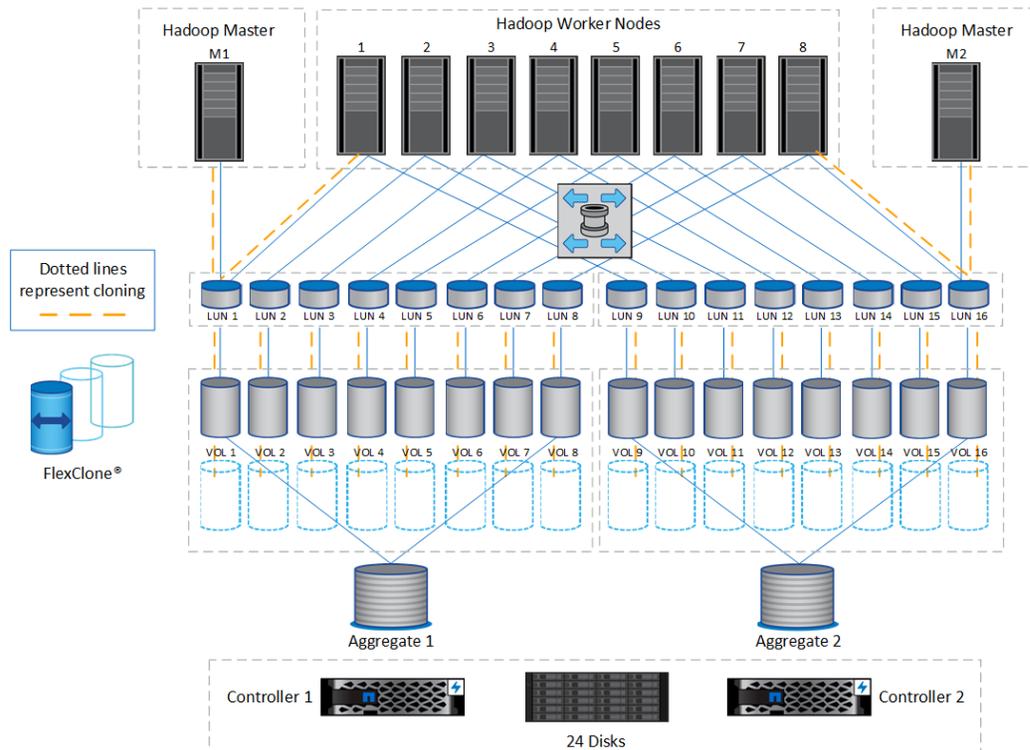


Figure 5 represents the duplication of the Hadoop cluster using FlexClone technology.

Figure 5) Duplicate Hadoop cluster using NetApp FlexClone technology.



## 8 Performance Testing

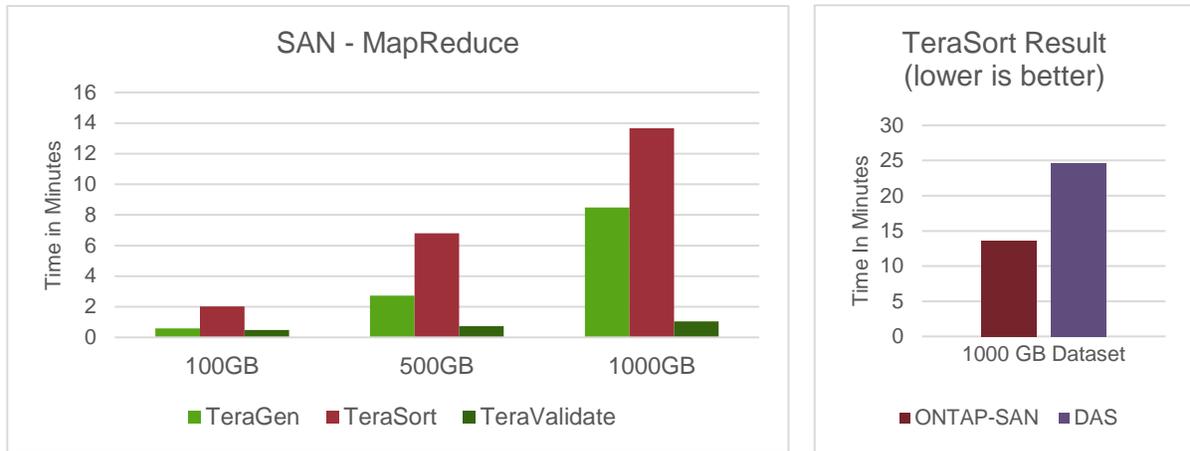
To measure the performance of the NetApp ONTAP SAN solution for Hadoop, we used the TeraGen, TeraSort, and TeraValidate tools. The following subsections present the results of our tests.

### 8.1 Tests Conducted

#### Performance

Based on our validation, we found that a 1TB dataset TeraSort operation completed in 13 minutes. This performance is very good for a Hadoop cluster running with SAN specifically on FC. We also did a performance comparison with DAS in TeraSort with a 1000GB dataset. Figure 6 shows the performance results.

Figure 6) Performance results.



#### Efficiency

Based on our validation with a small dataset, our results indicated a storage savings of close to 30% from built-in compression. See the following test sample.

```

stlaff300-land2::> volume efficiency show -vserver hadoopsvm
Vserver      Volume          State    Status    Progress          Policy
-----
hadoopsvm    a300nc1         Enabled  Idle      Idle for 1687:10:17
hadoopsvm    mongodb_root   Enabled  Idle      Idle for 3368:37:42
hadoopsvm    pocfg           Enabled  Idle      -
hadoopsvm    vol1_1         Enabled  Idle      Idle for 2688:38:55
hadoopsvm    vol1_10        Enabled  Idle      Idle for 128:04:03
hadoopsvm    vol1_11        Enabled  Idle      Idle for 3368:18:56
hadoopsvm    vol1_12        Enabled  Idle      Idle for 3368:18:44
hadoopsvm    vol1_2         Enabled  Idle      Idle for 267:23:48
hadoopsvm    vol1_3         Enabled  Idle      Idle for 267:31:40
hadoopsvm    vol1_4         Enabled  Idle      Idle for 264:41:07
hadoopsvm    vol1_5         Enabled  Idle      Idle for 127:28:17
hadoopsvm    vol1_6         Enabled  Idle      Idle for 127:10:35
hadoopsvm    vol1_7         Enabled  Idle      Idle for 127:36:38
hadoopsvm    vol1_8         Enabled  Idle      Idle for 127:18:09
hadoopsvm    vol1_9         Enabled  Idle      Idle for 127:48:30
  
```

```

hadoopsvm vol2_1          Enabled  Idle      Idle for 3368:17:36
                                           auto
hadoopsvm vol2_10         Enabled  Idle      Idle for 128:39:54 auto
hadoopsvm vol2_11         Enabled  Idle      Idle for 3368:05:21
                                           auto
hadoopsvm vol2_12         Enabled  Idle      Idle for 3368:04:56
                                           auto
hadoopsvm vol2_2          Enabled  Idle      Idle for 126:43:28 auto
hadoopsvm vol2_3          Enabled  Idle      Idle for 128:21:22 auto
hadoopsvm vol2_4          Enabled  Idle      Idle for 127:10:32 auto
hadoopsvm vol2_5          Enabled  Idle      Idle for 263:22:35 auto
hadoopsvm vol2_6          Enabled  Idle      Idle for 263:18:05 auto
hadoopsvm vol2_7          Enabled  Idle      Idle for 126:53:42 auto
hadoopsvm vol2_8          Enabled  Idle      Idle for 127:28:17 auto
hadoopsvm vol2_9          Enabled  Idle      Idle for 127:49:20 auto

```

27 entries were displayed.

stlaff300-land2::>.

stlaff300-land2::> df -h -S -vserver hadoopsvm

Filesystem	compressed	used	total-saved	%total-saved	deduplicated	%deduplicated
/vol/a300nc1/	740GB	1103GB	778GB	41%	37GB	2%
/vol/mongodb_root/	0B	652KB	0B	0%	0B	0%
/vol/pocfg/	1219GB	2630GB	1277GB	33%	57GB	1%
/vol/vol1_1/	35MB	436GB	1384MB	0%	1348MB	0%
/vol/vol1_10/	131GB	306GB	131GB	30%	93MB	0%
/vol/vol1_11/	0B	8484KB	0B	0%	0B	0%
/vol/vol1_12/	0B	8524KB	0B	0%	0B	0%
/vol/vol1_2/	128GB	302GB	134GB	31%	6187MB	1%
/vol/vol1_3/	88GB	347GB	89GB	21%	1358MB	0%
/vol/vol1_4/	80GB	354GB	82GB	19%	2372MB	1%
/vol/vol1_5/	96GB	341GB	96GB	22%	157MB	0%
/vol/vol1_6/	99GB	338GB	99GB	23%	144MB	0%
/vol/vol1_7/	95GB	342GB	95GB	22%	193MB	0%
/vol/vol1_8/	100GB	337GB	100GB	23%	167MB	0%
/vol/vol1_9/	97GB	339GB	98GB	22%	125MB	0%
/vol/vol2_1/	616KB	437GB	158MB	0%	157MB	0%
/vol/vol2_10/	129GB	303GB	133GB	31%	4094MB	1%
/vol/vol2_11/	0B	1512KB	0B	0%	0B	0%
/vol/vol2_12/	0B	1500KB	0B	0%	0B	0%
/vol/vol2_2/	120GB	316GB	120GB	28%	127MB	0%
/vol/vol2_3/	127GB	307GB	130GB	30%	3112MB	1%
/vol/vol2_4/	130GB	307GB	130GB	30%	137MB	0%
/vol/vol2_5/	127GB	307GB	130GB	30%	3308MB	1%
/vol/vol2_6/	129GB	304GB	133GB	30%	3892MB	1%
/vol/vol2_7/	132GB	301GB	136GB	31%	3934MB	1%

```
/vol/vol2_8/          307GB          129GB          30%          3769MB          1%
126GB                29%
/vol/vol2_9/          307GB          129GB          30%          3308MB          1%
126GB                29%
27 entries were displayed.

stlaff300-land2::>
```

## Conclusion

Organizations collect and analyze large amounts of raw data. To harness the maximum value of this big data, they must transform their raw data into valuable business information. The NetApp ONTAP SAN solution for Hadoop provides the high reliability and scalability required by enterprise organizations to ingest, protect, and manage that data. This solution features enterprise storage that is independent of the compute servers to offer an enterprise-class deployment with lower cluster downtime, higher data availability, and linear scalability.

Customers who are new to Hadoop face a steep learning curve. However, the NetApp ONTAP SAN solution for Hadoop provides an operational model for a Hadoop cluster that, after it is set up, does not require much additional attention. The cluster is more stable and is easier to maintain, and it allows customers to concentrate on getting answers to their business needs.

The core benefit of the NetApp ONTAP SAN solution for Hadoop is an enterprise-grade storage platform for implementing a data lake built on Hadoop that provides everything customers expect from NetApp and a few bonuses such as:

- Ready-to-deploy Hadoop solution
- Only two copies of data required while providing high data availability
- Built in storage efficiency
- Linear scalability for handling data and business growth
- Proven, tested, supported, and fully compatible system design
- Lower total cost of ownership (TCO)
- FlexClone volumes that can be created almost instantly

The NetApp ONTAP SAN solution for Hadoop enables customers to safely run with a replication count of two, because the data is protected by using NetApp RAID DP on the aggregates that are used for data storage. With RAID DP, as many as two SSDs (or disks) can fail in a single aggregate without any loss of data or interruption in data availability.

Vast amounts of data represent both an opportunity and a threat to enterprises. The NetApp ONTAP SAN solution for Hadoop is the first solution that is affordable enough to reduce the threat of drowning in data while still realizing the business potential that might be discovered in the data.

## Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp ONTAP 9 Software Setup Guide  
<http://www.netapp.com/us/products/management-software/oncommand/unified-manager.aspx>
- Hortonworks website  
<https://hortonworks.com>
- Cloudera website  
<https://www.cloudera.com>

## Acknowledgments

I would like to thank the following NetApp and Hortonworks experts for their input and assistance:

- Prasad Menon, Partner Certification Manager, Hortonworks
- Vivek Somani, Product Release Manager, Hortonworks
- Ali Bajwa, Principal Partner Solutions Engineer, Hortonworks
- John Ryan, Big Data Global Alliances Director, NetApp
- Nilesh Bagad, Senior Product Manager, NetApp

## Version History

Version	Date	Document Version History
Version 1.0	June 2018	Initial release.
Version 1.1	January 2019	HDP certification added.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright Information**

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### **Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4697-0618