



Technical Report

NetApp Hybrid Data Protection Solutions for Hadoop and Spark

Customer Use Case-Based Solutions

Karthikeyan Nagalingam and Nilesh Bagad, NetApp
January 2018 | TR-4657

Abstract

This document provides Hadoop data protection solutions by using Hadoop native commands, NetApp® FAS/AFF storage systems, NetApp ONTAP® Cloud, NetApp Private Storage (NPS), FlexClone® technology, and the In-Place Analytics Module for Hadoop (previously named the NetApp NFSConnector). These solution architectures enable customers to choose an appropriate data protection solution for their environment. NetApp designed these solutions based on interaction with customers and their use cases.

TABLE OF CONTENTS

1	Solution Overview	4
1.1	Why Hadoop Data Protection?	4
1.2	Data Protection Challenges for Hadoop and Spark Customers	4
2	NetApp Data Fabric Architecture for Big Data	5
2.1	Proven Data Fabric Customer Use Cases	5
3	Hadoop Data Protection and NetApp In-Place Analytics Module	7
4	Overview of Hadoop Data Protection Use Cases	8
4.1	Use Case 1: Backing Up Hadoop Data	8
4.2	Use Case 2: Backup and DR from Cloud to On-Premises	8
4.3	Use Case 3: Enabling Dev/Test on Existing Hadoop Data	8
4.4	Use Case 4: Data Protection and Multicloud Connectivity	8
5	Use Case 1: Backing Up Hadoop Data	9
5.1	Scenario	9
5.2	Requirements and Challenges	9
5.3	Customer's Existing Backup Solution	9
5.4	Backup Solutions	10
6	Use Case 2: Backup and Disaster Recovery from Cloud to On-Premises	12
6.1	Scenario	13
6.2	Requirements and Challenges	13
6.3	Solution	13
7	Use Case 3: Enabling Dev/Test for Hadoop on Existing Hadoop Data	14
7.1	Scenario	14
7.2	Requirements and Challenges	14
7.3	Solution	14
8	Use Case 4: Data Protection and Multicloud Connectivity	15
8.1	Scenario	15
8.2	Requirements and Challenges	15
8.3	Solution	16
9	Conclusion	16
	Where to Find Additional Information	17
	Acknowledgements	17

LIST OF FIGURES

Figure 1) NetApp Data Fabric.....	5
Figure 2) In-Place Analytics.	6
Figure 3) Data Fabric building blocks.....	7
Figure 4) Hadoop and In-Place Analytics Module.....	8
Figure 5) Original backup solution.....	9
Figure 6) Backup solution A.....	10
Figure 7) Backup solution B.....	11
Figure 8) Backup solution C.....	12
Figure 9) Backup and DR on-premises.....	13
Figure 10) Backup and DR from cloud to on-premises solution.....	14
Figure 11) Hadoop cluster for dev/test.....	15
Figure 12) Data protection and multcloud connectivity.....	16

1 Solution Overview

This document provides Hadoop data protection solutions using Hadoop native commands, FAS/AFF storage systems, ONTAP Cloud, NPS, FlexClone technology, and the NetApp In-Place Analytics Module for Hadoop (previously known as the NetApp NFSConnector). This document provides the following detailed information:

- Why we need data protection for Hadoop environments and discussion about the current customer challenges
- The NetApp Data Fabric vision and its building blocks and services
- How those building blocks can be used to architect flexible Hadoop data protection workflows
- The pros and cons of several architectures based on real-world customer use cases. Each use case provides the following components:
 - Customer scenario
 - Requirements and challenges
 - Solution
 - Summary of the solutions

1.1 Why Hadoop Data Protection?

In a Hadoop and Spark environment, the following concerns must be addressed:

- **Software or human failures.** Human error in software and in carrying out operations can lead to faulty behavior in the Hadoop data and can cause unexpected results from the job. In this case, we need to protect the data to avoid failures. For example, as the result of an update for a traffic signal analysis application, a new feature breaks properly analyzing traffic signal data in the form of plain text. The software still analyzes JSON and other non-plain text formats, resulting in the real-time traffic control analytics system producing prediction results that are missing data points. This situation can cause faulty details that might lead to accidents at the traffic signals. Data protection can address this issue by providing the capability to quickly roll back to the previous working application version.
- **Size and scale.** The size of the analytics data grows day by day due to ever-increasing numbers of data sources and volume. Social, mobile, analytics, and cloud are the main sources of data in the current big data market, which increases very rapidly, so this data needs to be protected to make sure of accurate analytics operations.
- **Hadoop's native data protection.** Hadoop has a native command to protect the data, but this command does not provide consistency of data during backup. It only supports directory-level backup. The snapshots created by Hadoop are read-only and cannot be used to reuse the backup data directly.

1.2 Data Protection Challenges for Hadoop and Spark Customers

A common challenge for Hadoop/Spark customers is to reduce the backup time without negatively affecting performance at the production cluster during data protection.

Customers also need control over their on-premises and cloud disaster recovery (DR) sites. This control typically comes from having enterprise-level management tools.

The Hadoop and Spark environments are complicated because not only is the data volume huge and growing, but also the rate this data arrives is increasing. This situation makes it difficult to rapidly create efficient, up-to-date development/test and QA environments from the source data.

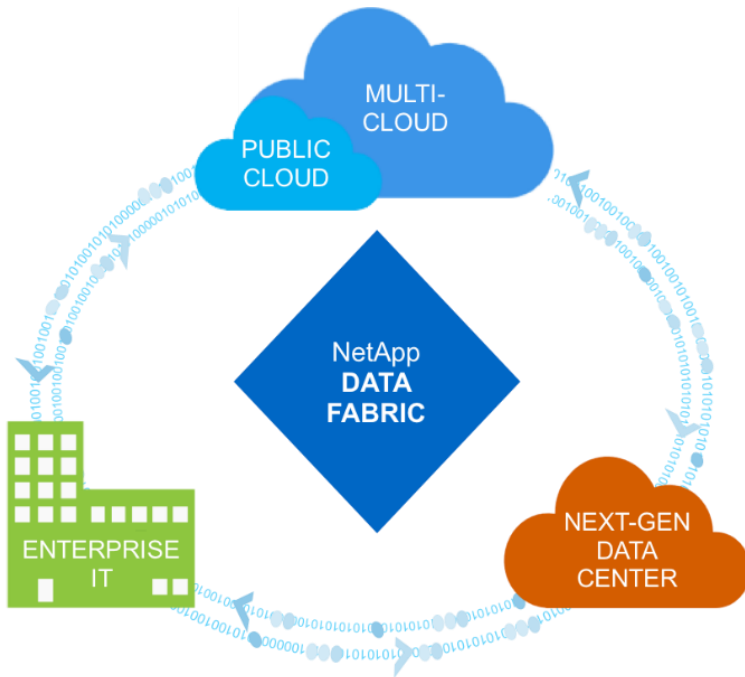
NetApp recognizes these challenges and offers the solutions presented in this paper.

2 NetApp Data Fabric Architecture for Big Data

The NetApp Data Fabric simplifies and integrates data management across cloud and on-premises environments to accelerate digital transformation.

The NetApp Data Fabric delivers consistent and integrated data management services and applications (building blocks) for data visibility and insights, data access and control, and data protection and security, as shown in Figure 1.

Figure 1) NetApp Data Fabric.



2.1 Proven Data Fabric Customer Use Cases

NetApp Data Fabric provides the following nine proven use cases for customers:

- Accelerate workloads and analytics
- Accelerate DevOps transformation
- Build cloud hosting infrastructure
- Integrate cloud data services
- Protect and secure data
- Optimize unstructured data
- Gain data center efficiencies
- Deliver data insights and control
- Simplify and automate

This document covers two of the nine use cases (along with their solutions):

- Accelerate workloads and analytics
- Protect and secure data

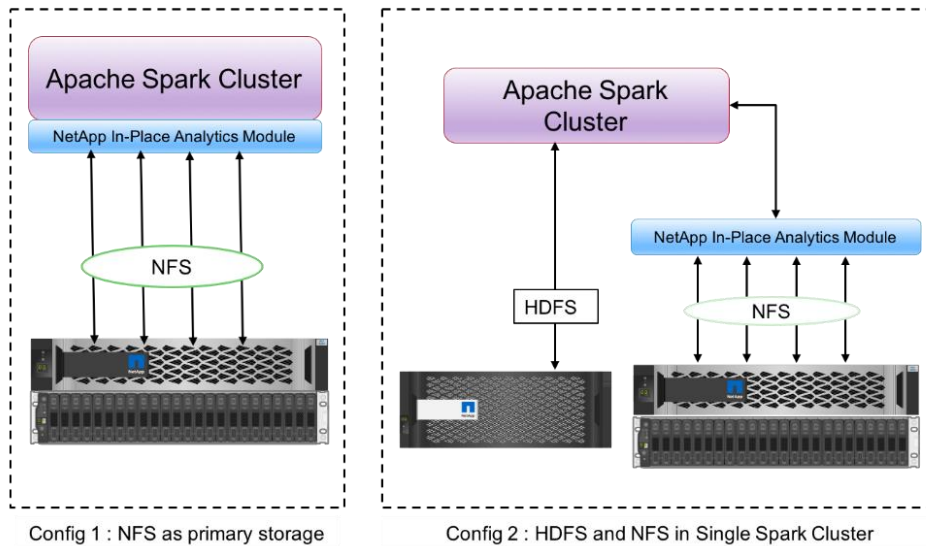
In-Place Analytics Module

The NetApp In-Place Analytics Module enables customers to run big data analytics jobs on their existing or new NFSv3 data without moving or copying the data. It avoids multiple copies of data and eliminates syncing the data with a source. For example, in the financial sector, the movement of data from one place to another place must meet legal obligations, which is not an easy task. In this scenario, the In-Place Analytics Module analyzes the financial data from its original location. Another key benefit is that using the In-Place Analytics Module simplifies protecting Hadoop data by using native Hadoop commands and enables data protection workflows leveraging NetApp's rich data management portfolio.

Figure 2) In-Place Analytics.

In Place Analytics

Enable big data analytics on NFSv3 data



The In-Place Analytics Module provides two kinds of deployment options for Hadoop/Spark clusters:

- By default, the Hadoop/Spark clusters use Hadoop Distributed File System (HDFS) for data storage and the default file system. The In-Place Analytics Module can replace the default HDFS with NFS storage as the default file system, enabling direct analytics operations on NFS data.
- In another deployment option, the In-Place Analytics Module supports configuring NFS as additional storage along with HDFS in a single Hadoop/Spark cluster. In this case, the customer can share data through NFS exports and access it from the same cluster along with HDFS data.

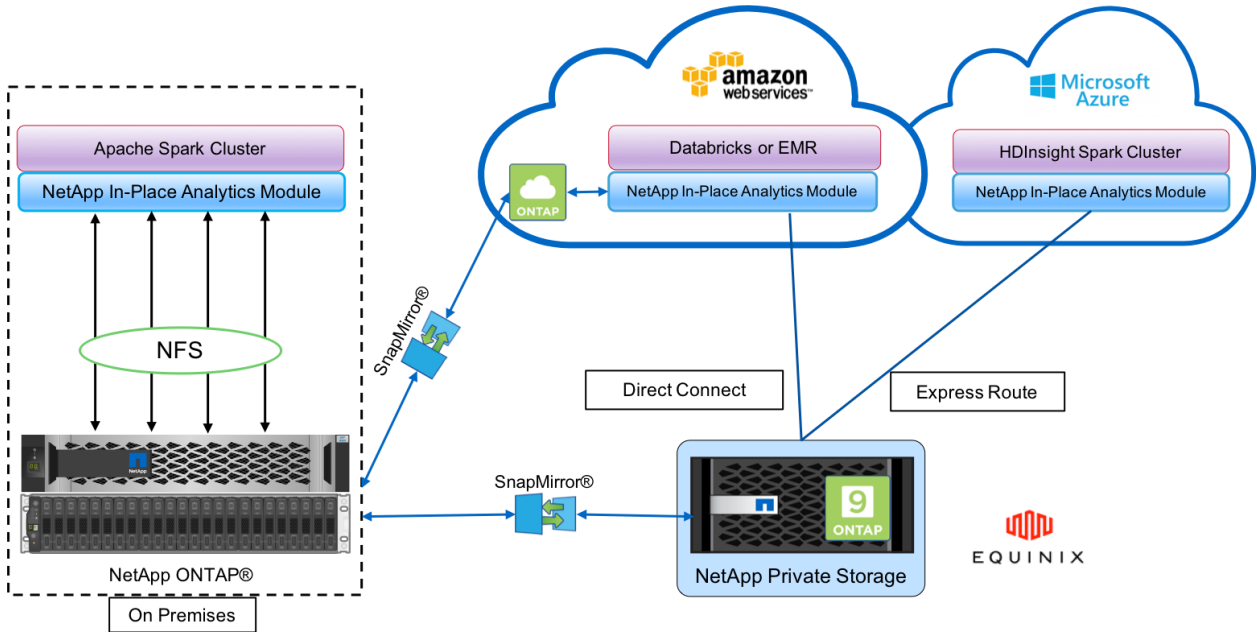
The key benefits of using the NetApp In-Place Analytics Module are:

- Analyzes the data from its current location, which prevents the time- and performance-consuming task of moving analytics data to Hadoop infrastructure such as HDFS.
- Reduces the number of replicas from three to one.
- Enables users to decouple the compute and storage to scale them independently.
- Provides enterprise data protection by leveraging the rich data management capabilities of ONTAP.
- Is certified with the Hortonworks data platform.
- Enables hybrid data analytics deployments.
- Reduces the backup time by leveraging dynamic multithread capability.

NetApp Data Fabric Building Blocks for Big Data

The NetApp Data Fabric integrates data management services and applications (building blocks) for data access, control, protection, and security.

Figure 3) Data Fabric building blocks.




The building blocks in Figure 3 include:

- **NetApp In-Place Analytics Module.** Provides access to NFS data to the Hadoop and Spark clusters.
- **NetApp Private Storage (NPS).** Allows customers to protect and govern their own data and access it from cloud instances through direct connect (AWS) or express route (Microsoft Azure).
- **ONTAP Cloud.** Software-defined storage based on ONTAP running in an Amazon Web Services (AWS) or Azure instance.
- **NetApp SnapMirror®.** Provides data protection capabilities between on-premises and ONTAP Cloud or NPS instances.
- **Cloud service providers.** AWS, Microsoft Azure, and IBM Cloud.
- **PaaS.** Cloud-based analytics services such as Elastic Mapreduce (EMR) and Databricks in AWS as well as Microsoft Azure HDInsight.

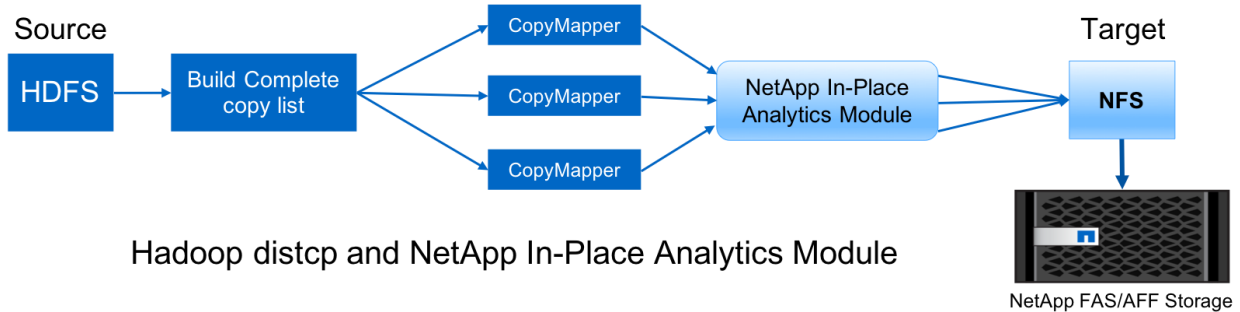
3 Hadoop Data Protection and NetApp In-Place Analytics Module

Hadoop distcp is a native tool used for large inter- and intracluster copying. The Hadoop distcp basic process shown in Figure 4 is a typical backup workflow using Hadoop native tools such as MapReduce to copy Hadoop data from an HDFS source to a corresponding target. The NetApp In-Place Analytics Module enables customers to set NFS as the target destination for the Hadoop distcp tool to copy the data from HDFS source into an NFS share through MapReduce. The NetApp In-Place Analytics Module acts as an NFS driver for the distcp tool.

Figure 4) Hadoop and In-Place Analytics Module.



Hadoop distcp Basic Process



Hadoop distcp and NetApp In-Place Analytics Module

4 Overview of Hadoop Data Protection Use Cases

This section provides a high-level description of the data protection use cases, which constitute the focus of this paper. Sections 5 through 8 provide more details for each use case, such as the customer problem (scenario), requirements and challenges, and solution.

4.1 Use Case 1: Backing Up Hadoop Data

For this use case, the NetApp In-Place Analytics Module helped a large financial institution reduce the long backup window time from more than 24 hours to just under a few hours.

4.2 Use Case 2: Backup and DR from Cloud to On-Premises

By using NetApp Data Fabric building blocks, a large broadcasting company was able to fulfill its requirement of backing up cloud data into its on-premise data center depending on different modes of data transfers, such as on demand, instantaneous, or based on the Hadoop/Spark cluster load.

4.3 Use Case 3: Enabling Dev/Test on Existing Hadoop Data

NetApp solutions helped an online music distributor to rapidly build multiple space efficient Hadoop clusters in different branches to create reports and run daily dev/test tasks by using scheduled policies.

4.4 Use Case 4: Data Protection and Multicloud Connectivity

A large service provider used NetApp Data Fabric to provide multicloud analytics to its customers from different cloud instances.

5 Use Case 1: Backing Up Hadoop Data

5.1 Scenario

In this scenario, the customer has a large on-premises Hadoop repository and wants to back it up for DR purposes. However, the customer's current backup solution is costly and is suffering from a long backup window of more than 24 hours.

5.2 Requirements and Challenges

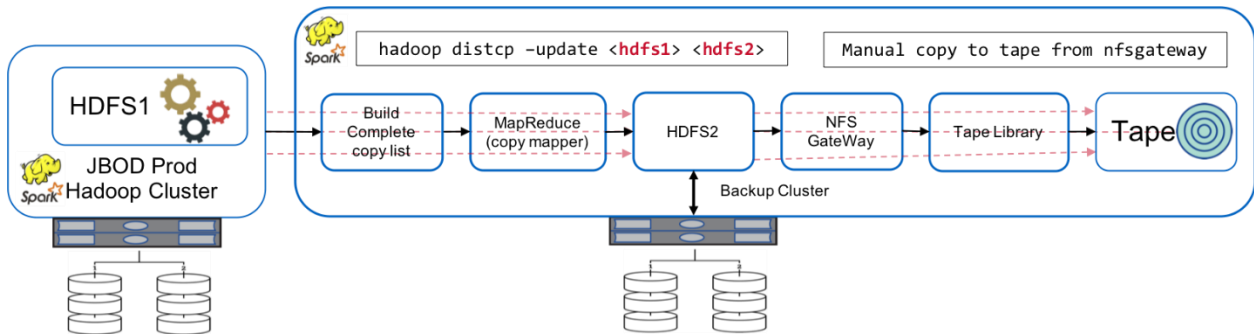
The main requirements and challenges for this use case include:

- Software backward compatibility:
 - The proposed alternative backup solution should be compatible with the current running software versions used in the production Hadoop cluster.
- To meet the committed SLAs, the proposed alternative solution should achieve very low recovery point objectives (RPOs) and recovery time objectives (RTOs).
- The backup created by the NetApp backup solution can be used in the Hadoop cluster built locally in the data center as well as the Hadoop cluster running in the DR location at the remote site.
- The proposed solution must be cost effective.
- The proposed solution must reduce the performance impact on the currently running, in-production analytics jobs during the backup times.

5.3 Customer's Existing Backup Solution

Figure 5 shows the original Hadoop native backup solution.

Figure 5) Original backup solution.



The production data is protected to tape through the intermediate backup cluster:

1. HDFS1 data is copied to HDFS2 by running the `hadoop distcp -update <hdfs1> <hdfs2>` command.
2. The backup cluster acts as an NFS gateway, and the data is manually copied to tape through the Linux `cp` command through the tape library.

The benefits of the original Hadoop native backup solution include:

- The solution is based on Hadoop native commands, which saves the user from having to learn new procedures.
- The solution leverages industry-standard architecture and hardware.

The disadvantages of the original Hadoop native backup solution include:

- The long backup window time exceeds 24 hours, which makes the production data vulnerable.
- Significant cluster performance degradation during backup times.
- Copying to tape is a manual process.
- Backup solution is expensive in terms of the hardware required and the human hours required for manual processes.

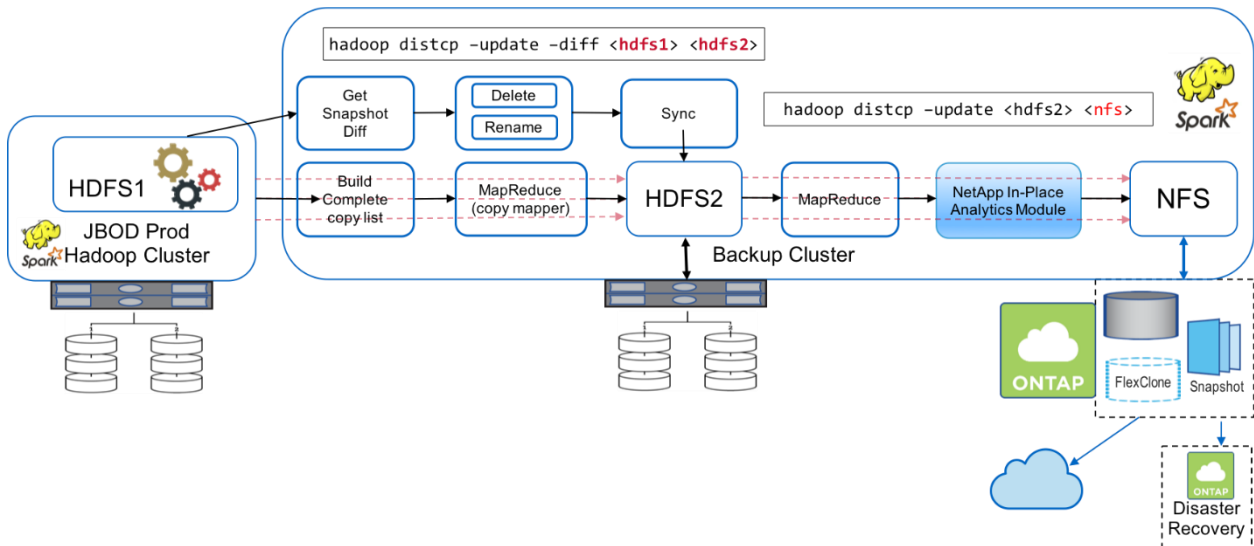
5.4 Backup Solutions

Based on these challenges and requirements and taking into consideration the existing backup system, three possible backup solutions were suggested. The following subsections describe each of these three different backup solutions, labeled solution A through solution C.

Solution A

Solution A adds the NetApp In-Place Analytics Module to the backup Hadoop cluster, which allows secondary backups to NetApp NFS storage systems, eliminating the tape requirement, as shown in Figure 6.

Figure 6) Backup solution A.



The detailed tasks for solution A include:

- The production Hadoop cluster has the customer's analytics data in the HDFS that requires protection.
- The backup Hadoop cluster with HDFS acts as an intermediate location for the data. JBOD (just a bunch of disks) provides the storage for HDFS in both the production and backup Hadoop clusters.
- Protect the Hadoop production data from the production cluster HDFS to the backup cluster HDFS by running the `Hadoop distcp -update -diff <hdfs1> <hdfs2>` command.

Note: The Hadoop snapshot is used to protect the data from production to the backup Hadoop cluster.

- The NetApp ONTAP storage controller provides an NFS exported volume, which is provisioned to the backup Hadoop cluster.
- By running the `Hadoop distcp` command leveraging MapReduce and multiple mappers, the analytics data is protected from the backup Hadoop cluster to NFS by using the NetApp In-Place Analytics Module.

- After the data is stored in NFS on the NetApp storage system, NetApp Snapshot™, SnapRestore®, and FlexClone technologies are used to back up, restore, and duplicate the Hadoop data as needed.

Note: Hadoop data can be protected to cloud as well as DR locations by using SnapMirror technology.

The benefits of solution A include:

- Hadoop production data is protected from the backup cluster.
- HDFS data is protected through NFS enabling protection to cloud and DR locations.
- Improves performance by offloading backup operations to the backup cluster.
- Eliminates manual tape operations
- Allows for enterprise management functions through NetApp tools.
- Requires minimal changes to the existing environment.
- Is a cost-effective solution.

The disadvantage of this solution is that it requires a backup cluster and additional mappers to improve performance.

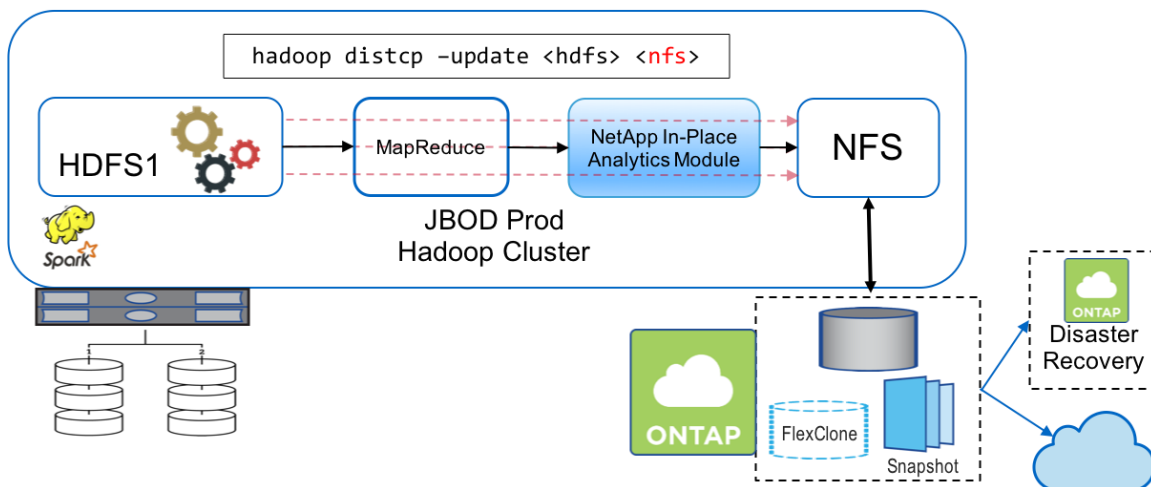
The customer recently deployed solution A due to its simplicity, cost, and overall performance.

Note: In this solution, SAN disks from ONTAP can be used instead of JBOD. This option offloads the backup cluster storage load to ONTAP; however, the downside is that SAN fabric switches are required.

Solution B

Solution B adds the In-Place Analytics Module to the production Hadoop cluster, which eliminates the need for the backup Hadoop cluster, as shown in Figure 7.

Figure 7) Backup solution B.



The detailed tasks for solution B include:

- The NetApp ONTAP storage controller provisioned the NFS export to the production Hadoop cluster. The Hadoop native `hadoop distcp` command protects the Hadoop data from the production cluster HDFS to NFS through the NetApp In-Place Analytics Module.
- After the data is stored in NFS on the NetApp storage system, Snapshot, SnapRestore, and FlexClone technologies are used to back up, restore, and duplicate the Hadoop data as needed.

The benefits of solution B include:

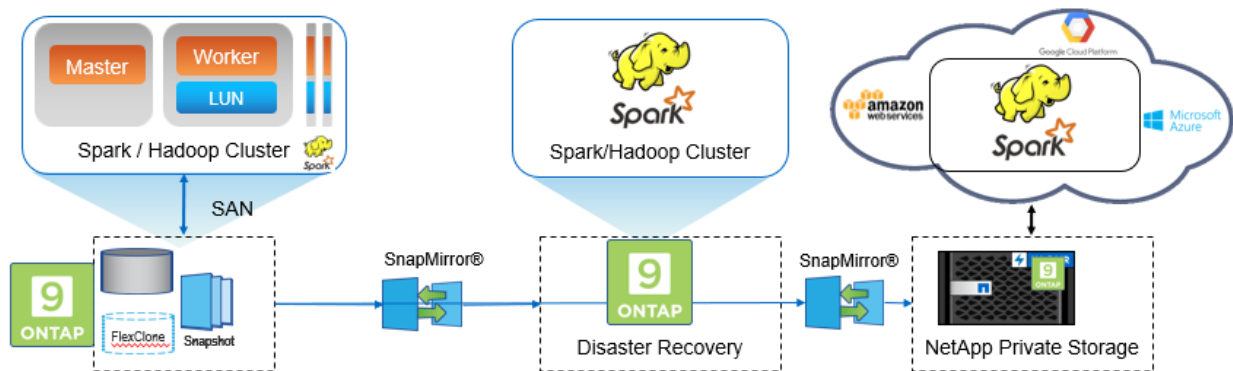
- The production cluster is slightly modified for the backup solution, which simplifies implementation and reduces additional infrastructure cost.
- A backup cluster for the backup operation is not required.
- HDFS production data is protected in the conversion to NFS data.
- The solution allows for enterprise management functions through NetApp tools.

The disadvantage of this solution is that it's implemented in the production cluster, which can add additional administrator tasks in the production cluster.

Solution C

In solution C, the NetApp SAN volumes are directly provisioned to the Hadoop production cluster for HDFS storage, as shown in Figure 8.

Figure 8) Backup solution C.



The detailed steps for solution C include:

- NetApp ONTAP SAN storage is provisioned at the production Hadoop cluster for HDFS data storage.
- NetApp Snapshot and SnapMirror technologies are used to back up the HDFS data from the production Hadoop cluster.
- There is no performance impact to production for the Hadoop/Spark cluster during the Snapshot copy backup process because the backup is at the storage layer.

Note: Snapshot technology provides backups that complete in seconds independent of the size of the data.

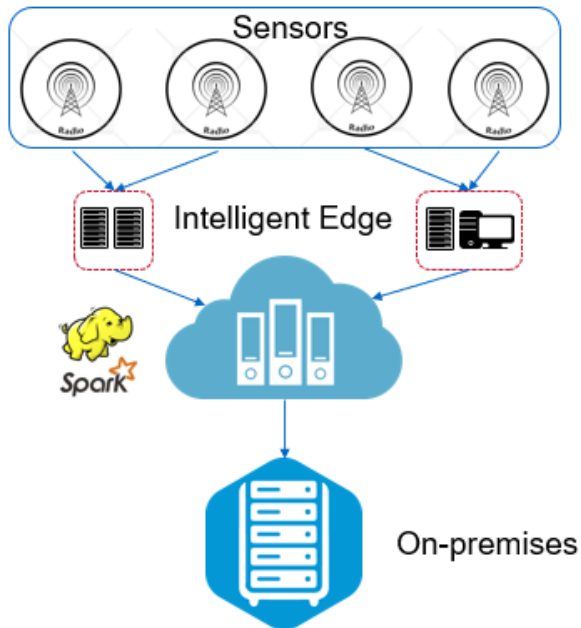
The benefits of solution C include:

- Space-efficient backup can be created by using Snapshot technology.
- Allows for enterprise management functions through NetApp tools.

6 Use Case 2: Backup and Disaster Recovery from Cloud to On-Premises

This use case is based on a broadcasting customer that needs to back up cloud-based analytics data to its on-premises data center, as illustrated in Figure 9.

Figure 9) Backup and DR on-premises.



6.1 Scenario

In this scenario, the IoT sensor data is ingested into the cloud and analyzed by using an open source Apache Spark cluster within Amazon Web Services. The requirement is to back up the processed data from the cloud to on-premises.

6.2 Requirements and Challenges

The main requirements and challenges for this use case include:

- Enabling data protection should not cause any performance impact on the production Spark/Hadoop cluster in the cloud.
- Cloud sensor data needs to be moved and protected to on-premises in an efficient and secure way.
- Flexibility to transfer data from the cloud to on-premises under different conditions, such as on demand, instantaneous, and during low-cluster load times.

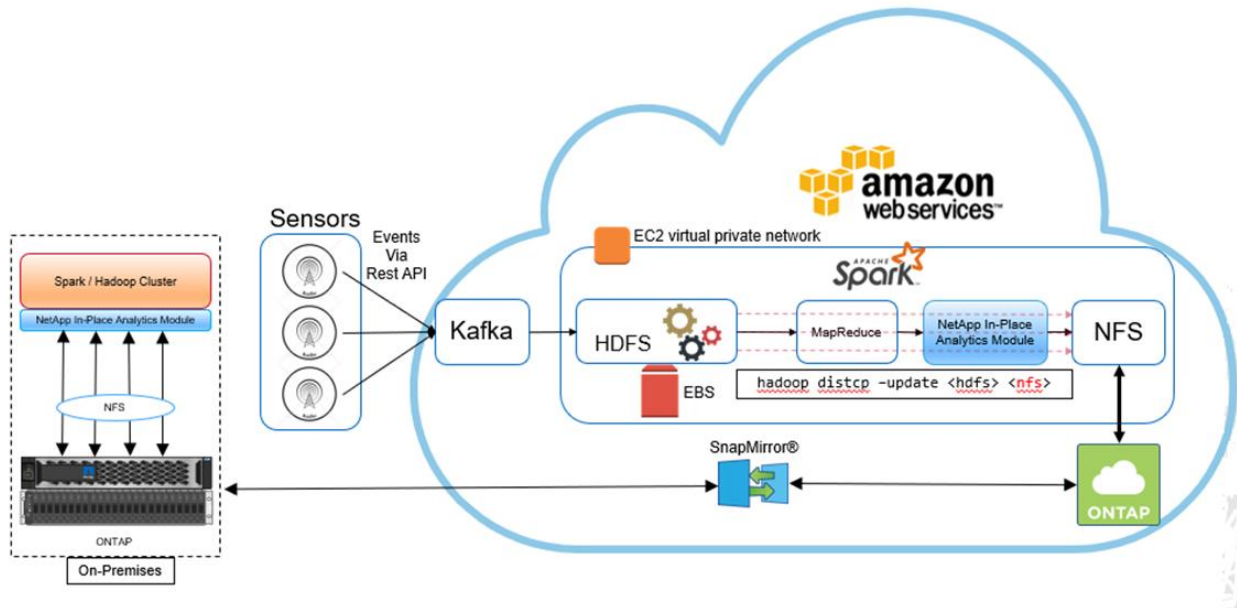
6.3 Solution

The customer uses AWS Elastic Block Store (EBS) for its Spark cluster HDFS storage to receive and ingest data from remote sensors through Kafka. Consequently, the HDFS storage acts as the source for the backup data.

To fulfill these requirements, NetApp ONTAP Cloud is deployed in AWS, and an NFS share is created to act as the backup target for the Spark/Hadoop cluster.

After the NFS share is created, the NetApp In-Place Analytics Module is leveraged to copy the data from the HDFS EBS storage into the ONTAP NFS share. After the data resides in NFS in ONTAP Cloud, SnapMirror technology can be used to mirror the data from the cloud into on-premises storage as needed in a secure and efficient way.

Figure 10) Backup and DR from cloud to on-premises solution.



7 Use Case 3: Enabling Dev/Test for Hadoop on Existing Hadoop Data

In this use case, the customer's requirement is to rapidly and efficiently build new Hadoop/Spark clusters based on an existing Hadoop cluster containing a large amount of analytics data for dev/test and reporting purposes in the same data center as well as remote locations.

7.1 Scenario

In this scenario, multiple Spark/Hadoop clusters are built from a large Hadoop data lake implementation on-premises as well as at DR locations.

7.2 Requirements and Challenges

The main requirements and challenges for this use case include:

- Create multiple Hadoop clusters for dev/test, for QA, or for any other purpose that requires access to the same production data. The challenge here is to clone a very large Hadoop cluster multiple times instantaneously and in a very space-efficient manner.
- Sync Hadoop data to dev/test and reporting teams for operational efficiency.
- Distribute the Hadoop data by using the same credentials across production and new clusters.
- Use scheduled policies to efficiently create QA clusters without affecting the production cluster.

7.3 Solution

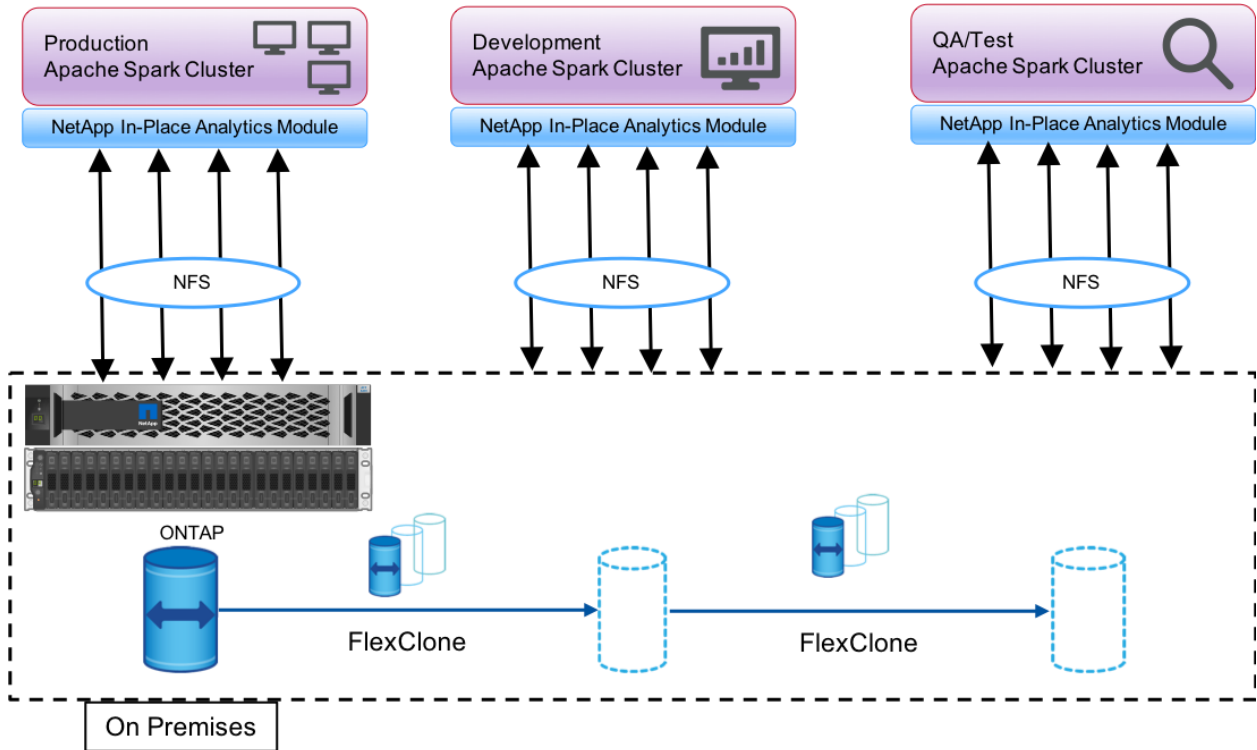
NetApp FlexClone technology is leveraged to answer the requirements just described. FlexClone technology is the read/write copy of a Snapshot copy. It reads the data from parent Snapshot copy data and only consumes additional space for new/modified blocks. It is fast and space-efficient.

First, a Snapshot copy of the existing cluster was created by using NetApp consistency group (CG) Snapshot copies within NetApp System Manager or storage admin prompt. The CG Snapshot copies are

application-consistent group Snapshot copies, and the FlexClone volume is created based on CG Snapshot copies. It is worth mentioning that a FlexClone volume inherits the parent volume's NFS export policy. After the Snapshot copy is created, a new Hadoop cluster must be installed for dev/test and reporting purposes, as shown in Figure 11. The NetApp In-Place Analytics Module accesses the cloned NFS volume from the new Hadoop cluster through NetApp In-Place Analytics Module users and group authorization for the NFS data.

To have proper access, the new cluster must have the same UID and GUID for the users configured in the NetApp In-Place Analytics Module “Users and Groups” configuration.

Figure 11) Hadoop cluster for dev/test.



8 Use Case 4: Data Protection and Multicloud Connectivity

This use case is relevant for a cloud service partner tasked with providing multicloud connectivity for customers' big data analytics data.

8.1 Scenario

In this scenario, IoT data received in AWS from different sources is stored in a central location in NPS. The NPS storage is connected to Spark/Hadoop clusters located in AWS and Azure enabling big data analytics applications running in multiple clouds accessing the same data.

8.2 Requirements and Challenges

The main requirements and challenges for this use case include:

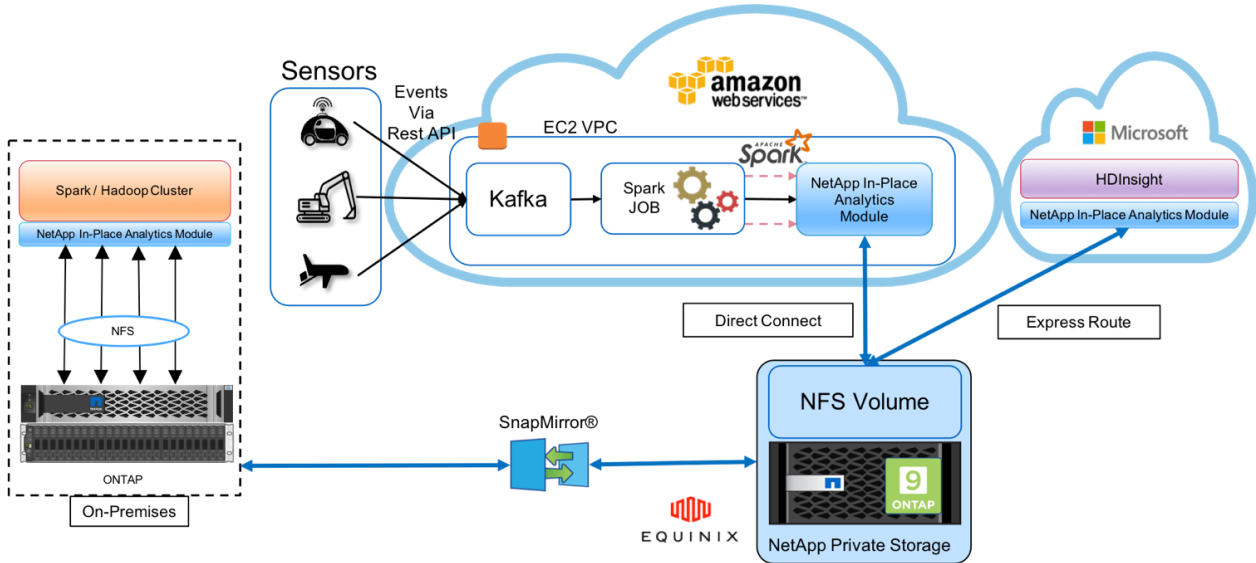
- Customers want to run analytics jobs on the same data using multiple clouds.

- Data must be received from different sources such as on-premises and cloud through different sensors and hubs.
- The solution must be efficient and cost-effective.
- The main challenge is to build a cost-effective and efficient solution that delivers hybrid analytics services between on-premises and across different clouds.

8.3 Solution

Figure 12 illustrates the data protection and multicloud connectivity solution.

Figure 12) Data protection and multicloud connectivity.



As shown in Figure 12, data from sensors is streamed and ingested into the AWS Spark cluster through Kafka. The data is stored in an NFS share residing in NPS, which is located outside of the cloud provider within an Equinix data center. Because NetApp NPS is connected to Amazon AWS and Microsoft Azure through Direct Connect and Express Route connections, respectively, customers can leverage the NetApp In-Place Analytics Module to access the data from both Amazon and AWS analytics clusters. This approach solves having cloud analytics across multiple hyperscalers.

Consequently, because both on-premises and NPS storage runs NetApp ONTAP OS, SnapMirror can mirror the NPS data into the on-premises cluster, providing hybrid cloud analytics across on-premises and multiple clouds.

For the best performance, NetApp typically recommends using multiple network interfaces and direct connection/express routes to access the data from cloud instances.

9 Conclusion

This section provides a summary of the use cases and solutions provided by NetApp to fulfill various Hadoop data protection requirements. By using the NetApp Data Fabric, customers can:

- Have the flexibility to choose the right data protection solutions leveraging NetApp's rich data management capabilities and integration with Hadoop native workflows.
- Reduce their Hadoop cluster backup window time by almost 70%.
- Eliminate any performance impact resulting from Hadoop cluster backups.

- Provide multicloud data protection and data access from different cloud providers simultaneously to a single source of analytics data.
- Create fast and space-efficient Hadoop cluster copies by using FlexClone technology.

Where to Find Additional Information

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

- NetApp Big Data Analytics Solutions
<https://www.netapp.com/us/solutions/applications/big-data-analytics/index.aspx>
- NetApp Storage Solutions for Apache Spark
<http://www.netapp.com/us/media/tr-4570.pdf>
- Apache Hadoop on Data Fabric Enabled by NetApp
<http://www.netapp.com/us/media/tr-4529.pdf>
- NetApp Insight™ presentation: Hybrid Cloud Data Protection Solutions for Hadoop

Acknowledgements

- Paul Burland, Sales Rep, ANZ Victoria District Sales, NetApp
- Hoseb Dermanilian, Business Development Manager EMEA, Big Data Analytics and Video Surveillance, NetApp
- Lee Dorrier, Director MPSG, NetApp
- David Thiessen, Systems Engineer, ANZ Victoria District SE, NetApp

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 © 2018 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-4657-0118