



NetApp and Dremio's next generation hybrid iceberg lakehouse solution

NetApp artificial intelligence solutions

NetApp
August 18, 2025

Table of Contents

- NetApp and Dremio’s next generation hybrid iceberg lakehouse solution 1
 - The NetApp and Dremio Next-Generation Hybrid Iceberg Lakehouse Solution 1
 - Solution Overview 1
 - NetApp Overview 1
 - Dremio overview 2
 - What value does the Dremio and NetApp Hybrid Iceberg Lakehouse solution deliver to customers? 2
 - Technology Requirements 3
 - Deployment Procedure 4
 - Solution verification overview 14
 - Customer Use Cases 21
 - NetApp ActiveIQ Use Case 21
 - Auto Parts Sales Customer Use Case. 22
 - Conclusion 23
 - Where to find additional information 23

NetApp and Dremio's next generation hybrid iceberg lakehouse solution

The NetApp and Dremio Next-Generation Hybrid Iceberg Lakehouse Solution

In this document, we discuss the deployment details of Dremio with different sources of data from NetApp storage controllers, including ONTAP S3, NAS, and StorageGRID. During the deployment, we used the TPC-DS benchmarking tool to execute 99 SQL queries across various sources. The document also explores customer use cases within NetApp, as well as a use case involving an auto parts sales customer.

Solution Overview

The Hybrid Iceberg Lakehouse solution provides unique benefits to address customer challenges faced by data lake customers. By Leveraging the Dremio Unified Lakehouse platform and NetApp ONTAP, StorageGRID, and NetApp Cloud solutions, companies can add significant value to their business operations. The solution not only provides access to multiple data sources, including NetApp sources, but also enhances overall analytical performance and helps companies drive business insights that leads to business growth.

NetApp Overview

- NetApp's offerings, such as ONTAP and StorageGRID, enable the separation of storage and computing, enabling optimal resource utilization based on specific requirements. This flexibility empowers customers to independently scale their storage using NetApp storage solutions
- By leveraging NetApp's storage controllers, customers can efficiently serve data to their vector database using NFS and S3 protocols. These protocols facilitate customer data storage and manage the vector database index, eliminating the need for multiple copies of data accessed through file and object methods.
- NetApp ONTAP provides native support for NAS and Object storage across leading cloud service providers like AWS, Azure, and Google Cloud. This broad compatibility ensures seamless integration, enabling customer data mobility, global accessibility, disaster recovery, dynamic scalability, and high performance.

StorageGRID

Our industry-leading object storage storageGRID offers a powerful policy engine for automated data placement, flexible deployment options, and unmatched durability with layered erasure coding. It has a scalable architecture supporting billions of objects and petabytes of data in a single namespace. The solution enables hybrid cloud integration, allowing data tiering to major cloud platforms. It has been recognized as a leader in the 2019 IDC Marketscape Worldwide Object-Based Vendor Assessment.

Additionally, storageGRID excels in managing unstructured data at scale with software-defined object storage, geo-redundancy, and multi-site capabilities. It incorporates policy-based information lifecycle management and offers cloud integration features like mirroring and search. It has various certifications, including Common Criteria, NF203 Digital Safe Component, ISO/IEC 25051, KPMG, and Cohasset Compliance Assessment.

In summary, NetApp storageGRID provides powerful features, scalability, hybrid cloud integration, and

compliance certifications for efficient management of unstructured data at scale.

NetApp ONTAP

NetApp ONTAP is a robust storage solution that offers a wide range of enterprise features. It includes Snapshot, which provides application-consistent and tamper-proof instant backups. SnapRestore enables near-instant restore of backups on demand, while SnapMirror offers integrated remote backup and disaster recovery capabilities. The solution also incorporates Autonomous Ransomware Protection (ARP), ensuring data security with features like multi-administrator verification, data-at-rest encryption with FIPS certification, in-transit data encryption, multifactor authentication (MFA), and role-based access control (RBAC). Comprehensive logging, auditing, onboard and external key management, secure purge, and secure management of multiple tenants further enhance data security and compliance.

NetApp ONTAP also features SnapLock, which provides regulatory-compliant data retention with high levels of integrity, performance, and retention at a low total cost of ownership. It is fully integrated with NetApp ONTAP 9 and offers protection against malicious acts, rogue administrators, and ransomware.

The solution encompasses NSE/NVE encryption for in-flight and data-at-rest encryption, multifactor admin access, and multi-admin verification. Active IQ provides AI-informed predictive analytics and corrective action, while QoS ensures quality of service workload control. The management and automation integration is intuitive through SysMgr/GUI/CLI/API. FabricPool enables automatic data tiering, and the solution offers efficiency through inline data compression, deduplication, and compaction. NetApp guarantees meeting workload efficiency goals at no cost to the customer.

NetApp ONTAP supports various protocols, including NVMe/FC, FC, NVMe/TCP, iSCSI, NFS, SMB, and S3, making it a unified storage solution. Overall, NetApp ONTAP provides extensive enterprise features, robust security, compliance, efficiency, and versatility to meet diverse storage needs.

Dremio overview

Dremio is the Unified Lakehouse Platform for self-service analytics and AI. The Dremio Unified Analytics Platform brings users closer to the data with lakehouse flexibility, scalability, and performance at a fraction of the cost of legacy data warehouse solutions. Dremio enables "shift-left" analytics to eliminate complex and costly data integration and ETL, delivering seamless enterprise-scale analytics with no data movement. Dremio also features:

- Easy-to-use self-service analytics enabled through a universal semantic layer and a tightly integrated, highly performant SQL query engine, making it easier to connect, govern, and analyze all data, both in the cloud and on-premises.
- Dremio's Apache Iceberg-native lakehouse management capabilities simplify data discovery, and automate data optimization, delivering high-performance analytics with Git-inspired data versioning.
- Foundationally built on open source and open standards, Dremio enables companies avoid lock-in and remain positioned for innovation. Enterprise companies trust Dremio as the easiest-to-use lakehouse platform with the best price performance across all workloads.

What value does the Dremio and NetApp Hybrid Iceberg Lakehouse solution deliver to customers?

- **Improved Data Management and Accessibility:** Dremio is well-known for its data lakehouse platform that enables organizations to query data directly from their data lakes at high speed. NetApp, on the other hand, is a leading provider of cloud data services and data storage solutions. The joint offer provides customers with a comprehensive solution for storing, managing, accessing, and analyzing their enterprise's data efficiently and efficiently.

- **Performance Optimization:** With NetApp's expertise in data storage and Dremio's capabilities in data processing and data optimization, the partnership offers a solution that improves the performance of data operations, reduces latency, and increases speed to business insight. Dremio has even delivered performance benefits to NetApp's own internal IT analytical infrastructure.
- **Scalability:** Both Dremio and NetApp offer a solution that is designed to scale. The joint solution provides customers with highly-scalable data storage, data management, and analytics environments. In a Hybrid Iceberg Lakehouse environment, the Dremio SQL query engine paired with NetApp StorageGRID delivers unparalleled scalability, concurrency, and query performance, capable of handling the analytical needs of any business.
- **Data Security and Governance:** Both companies have a strong focus on data security and governance. Together, they offer robust security and data governance features, ensuring that data is protected and that data governance requirements are met. Features such as role-based and fine-grain access controls, comprehensive auditing, end-to-end data lineage, unified identity management, and SSO with an extensive compliance and security framework ensure companies' analytical data environments are secure and governed.
- **Cost Efficiency:** By integrating Dremio's data lake engine with NetApp's storage solutions, customers can reduce costs associated with data management and data movement. Organizations are also able to move from legacy data lake environments to a more modern lakehouse solution composed of NetApp and Dremio. This Hybrid Iceberg Lakehouse solution delivers high-speed query performance and market-leading query concurrency that lowers TCO and reduces time to business insight.

Technology Requirements

The hardware and software configurations outlined below were utilized for validations performed in this document. These configurations serve as a guideline to help you set up your environment. However, please note the specific components may vary depending on individual customer requirements.

Hardware requirements

Hardware	Details
NetApp AFF Storage array HA Pair	<ul style="list-style-type: none"> • A800 • ONTAP 9.14.1 • 48 x 3.49TB SSD-NVM • Two S3 Buckets: Dremio metadata and customer data.
4 x FUJITSU PRIMERGY RX2540 M4	<ul style="list-style-type: none"> • 64 CPUs • Intel Xeon Gold 6142 CPU @ 2.60GHz • 256 GM Physical Memory • 1 x 100GbE network port
Networking	<ul style="list-style-type: none"> • 100 GbE

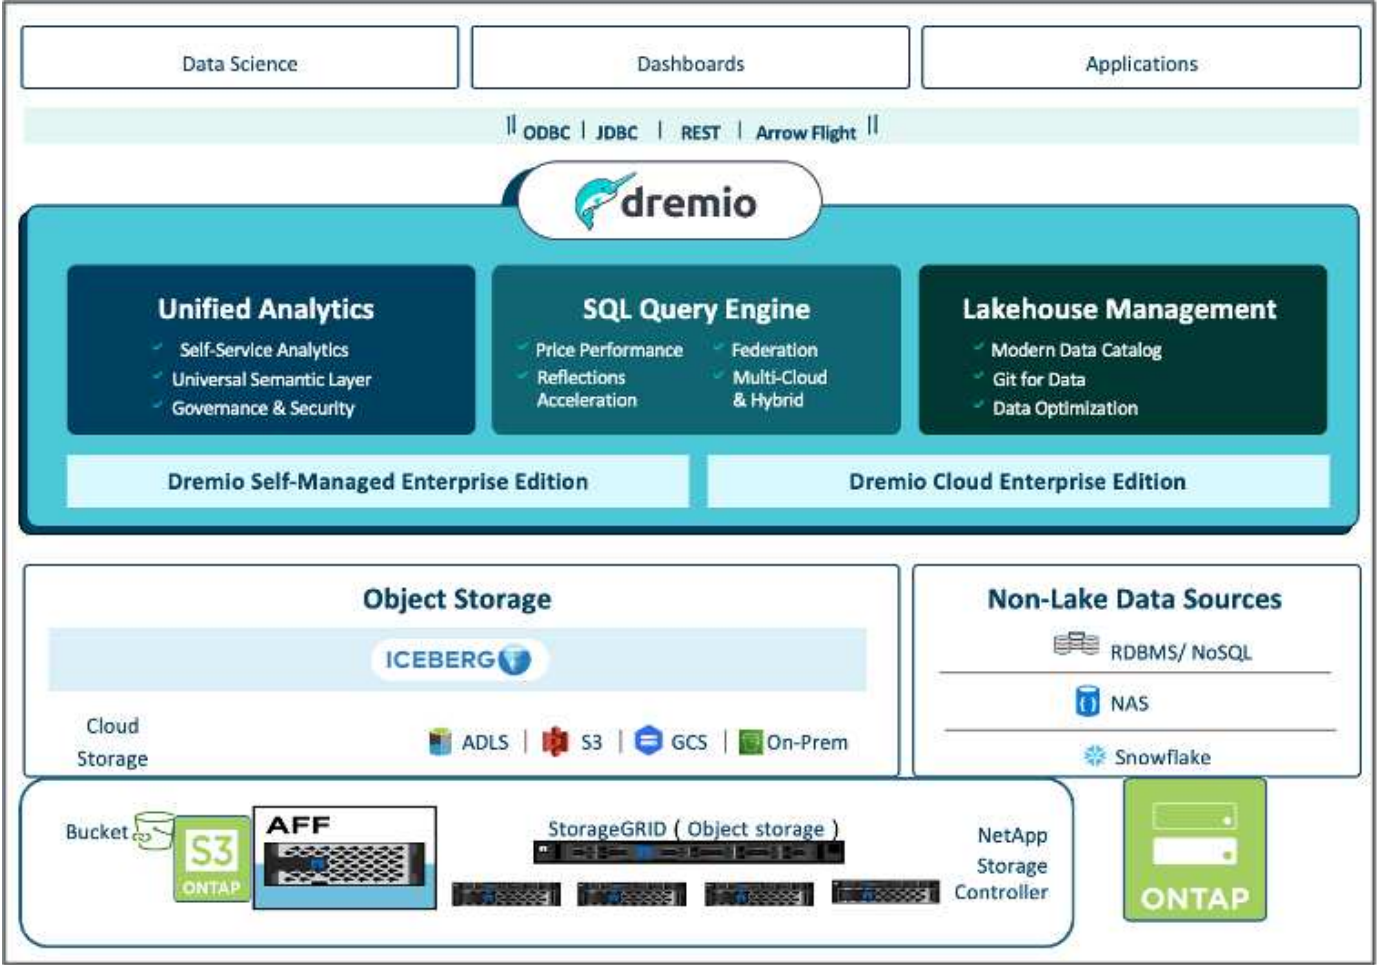
Hardware	Details
StorageGRID	<ul style="list-style-type: none"> * 1 x SG100, 3xSGF6024 * 3 x 24 x 7.68TB * Two S3 Buckets: Dremio metadata and customer data.

Software requirements

Software	Details
Dremio	<ul style="list-style-type: none"> • version - 25.0.3-202405170357270647-d2042e1b • Enterprise Edition
On-Prem	<ul style="list-style-type: none"> • 5 node Dremio cluster • 1 master coordinator and 4 executors

Deployment Procedure

In this reference architecture validation, we used a Dremio configuration composed of one coordinator and four executors



NetApp setup

- Storage system initialization
- Storage virtual machine (SVM) creation
- Assignment of logical network interfaces
- NFS, S3 configuration and licensing

Please follow the steps below for NFS (Network File System):

1. Create a Flex Group volume for NFSv4 or NFSv3. In our setup for this validation, we have used 48 SSDs, 1 SSD dedicated to the controller's root volume, and 47 SSDs spread across for NFSv4]]. Verify that the NFS export policy for the Flex Group volume has read/write permissions for the Dremio servers network.

1. On all Dremio servers, create a folder and mount the Flex Group volume onto this folder through a Logical Interface (LIF) on each Dremio server.

Please follow the steps below for S3 (Simple Storage Service):

1. Set up an object-store-server with HTTP enabled and the admin status set to 'up' using the "vserver object-store-server create" command. You have the option to enable HTTPS and set a custom listener port.
2. Create an object-store-server user using the "vserver object-store-server user create -user <username>" command.
3. To obtain the access key and secret key, you can run the following command: "set diag; vserver object-store-server user show -user <username>". However, moving forward, these keys will be supplied during the user creation process or can be retrieved using REST API calls.
4. Establish an object-store-server group using the user created in step 2 and grant access. In this example, we have provided "FullAccess".
5. Create two S3 buckets by setting its type to "S3". One for Dremio configuration and one for customer data.

Zookeeper setup

You can use Dremio provided zookeeper configuration. In this validation, we used separate zookeeper. we followed the steps mentioned in this weblink <https://medium.com/@ahmetfurkandemir/distributed-hadoop-cluster-1-spark-with-all-dependencies-03c8ec616166>

Dremio setup

We followed this weblink to install Dremio via tar ball.

1. Create a Dremio group.

```
sudo groupadd -r dremio
```

2. Create a dremio user.

```
sudo useradd -r -g dremio -d /var/lib/dremio -s /sbin/nologin dremio
```

3. Create Dremio directories.

```
sudo mkdir /opt/dremio
sudo mkdir /var/run/dremio && sudo chown dremio:dremio /var/run/dremio
sudo mkdir /var/log/dremio && sudo chown dremio:dremio /var/log/dremio
sudo mkdir /var/lib/dremio && sudo chown dremio:dremio /var/lib/dremio
```

4. Download the tar file from <https://download.dremio.com/community-server/>

5. Unpack Dremio into the /opt/dremio directory.

```
sudo tar xvf dremio-enterprise-25.0.3-202405170357270647-d2042e1b.tar.gz
-C /opt/dremio --strip-components=1
```

6. Create a symbolic link for the configuration folder.

```
sudo ln -s /opt/dremio/conf /etc/dremio
```

7. Set up your service configuration (SystemD setup).

1. Copy the unit file for the dremio daemon from /opt/dremio/share/dremio.service to /etc/systemd/system/dremio.service.
2. Restart system

```
sudo systemctl daemon-reload
```

3. Enable dremio to start at boot.

```
sudo systemctl enable dremio
```

8. Configure Dremio on coordinator. See Dremio Configuration for more information

1. Dremio.conf


```

root@hadoopmaster:/usr/src/tpcds# cat /opt/dremio/conf/dremio.conf

paths: {
  # the local path for dremio to store data.
  local: "${DREMIO_HOME}"/dremiocache"

  # the distributed path Dremio data including job results,
  downloads, uploads, etc
  #dist: "hdfs://hadoopmaster:9000/dremiocache"
  dist: "dremioS3:///dremioconf"
}

services: {
  coordinator.enabled: true,
  coordinator.master.enabled: true,
  executor.enabled: false,
  flight.use_session_service: false
}

zookeeper: "10.63.150.130:2181,10.63.150.153:2181,10.63.150.151:2181"
services.coordinator.master.embedded-zookeeper.enabled: false
root@hadoopmaster:/usr/src/tpcds#

```

2. Core-site.xml

```

root@hadoopmaster:/usr/src/tpcds# cat /opt/dremio/conf/core-site.xml
<?xml version="1.0" encoding="UTF-8"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>
<!--
  Licensed under the Apache License, Version 2.0 (the "License");
  you may not use this file except in compliance with the License.
  You may obtain a copy of the License at

      http://www.apache.org/licenses/LICENSE-2.0

  Unless required by applicable law or agreed to in writing, software
  distributed under the License is distributed on an "AS IS" BASIS,
  WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
  implied.
  See the License for the specific language governing permissions and
  limitations under the License. See accompanying LICENSE file.
-->

<!-- Put site-specific property overrides in this file. -->

```

```

<configuration>
  <property>
    <name>fs.dremioS3.impl</name>
    <value>com.dremio.plugins.s3.store.S3FileSystem</value>
  </property>
  <property>
    <name>fs.s3a.access.key</name>
    <value>24G4C1316APP2BIPDE5S</value>
  </property>
  <property>
    <name>fs.s3a.endpoint</name>
    <value>10.63.150.69:80</value>
  </property>
  <property>
    <name>fs.s3a.secret.key</name>
    <value>Zd28p43rgZaU44PX_ftT279z9nt4jBSro97j87Bx</value>
  </property>
  <property>
    <name>fs.s3a.aws.credentials.provider</name>
    <description>The credential provider type.</description>
    <value>org.apache.hadoop.fs.s3a.SimpleAWSCredentialsProvider</value>
  </property>
  <property>
    <name>fs.s3a.path.style.access</name>
    <value>false</value>
  </property>
  <property>
    <name>hadoop.proxyuser.dremio.hosts</name>
    <value>*</value>
  </property>
  <property>
    <name>hadoop.proxyuser.dremio.groups</name>
    <value>*</value>
  </property>
  <property>
    <name>hadoop.proxyuser.dremio.users</name>
    <value>*</value>
  </property>
  <property>
    <name>dremio.s3.compat</name>
    <description>Value has to be set to true.</description>
    <value>true</value>
  </property>
  <property>
    <name>fs.s3a.connection.ssl.enabled</name>

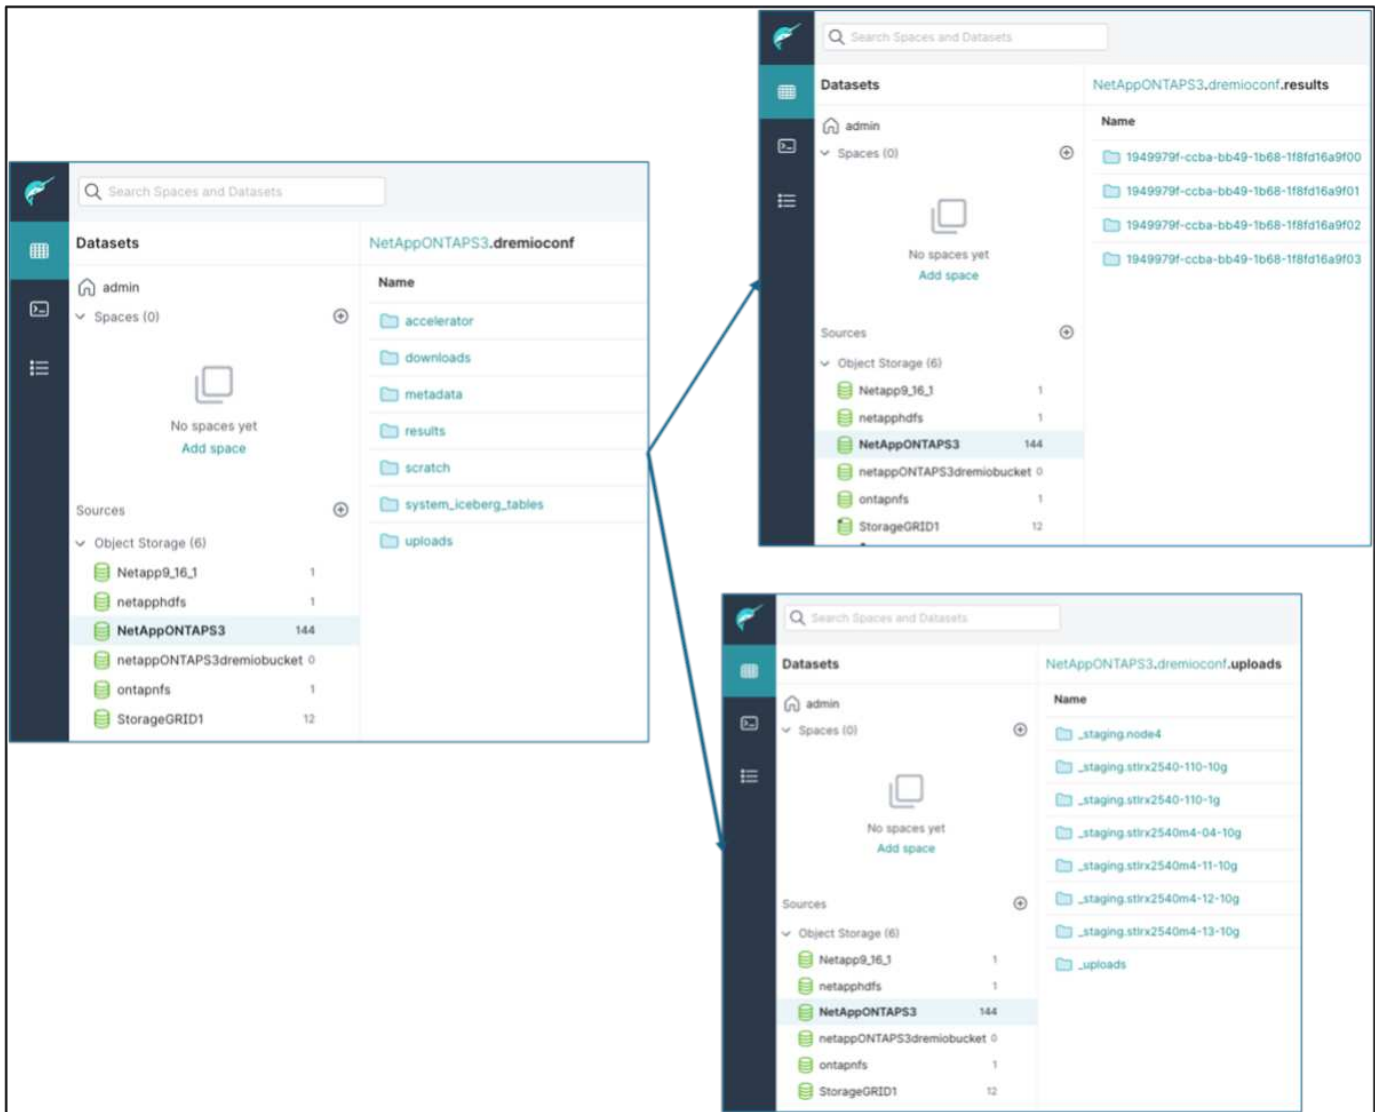
```

```

    <description>Value can either be true or false, set to true
    to use SSL with a secure Minio server.</description>
    <value>>false</value>
  </property>
</configuration>
root@hadoopmaster:/usr/src/tpcds#

```

9. The Dremio configuration is stored in NetApp object storage. In our validation, the "dremioconf" bucket resides in an ontap S3 bucket. The below picture shows some details from "scratch" and "uploads" folder of the "dremioconf" S3 bucket.



1. Configure Dremio on executors. In our setup, we have 3 executors.

1. dremio.conf

```

paths: {
  # the local path for dremio to store data.
  local: "${DREMIO_HOME}"/dremiocache"

  # the distributed path Dremio data including job results,
  downloads, uploads, etc
  #dist: "hdfs://hadoopmaster:9000/dremiocache"
  dist: "dremioS3:///dremioconf"
}

services: {
  coordinator.enabled: false,
  coordinator.master.enabled: false,
  executor.enabled: true,
  flight.use_session_service: true
}

zookeeper: "10.63.150.130:2181,10.63.150.153:2181,10.63.150.151:2181"
services.coordinator.master.embedded-zookeeper.enabled: false

```

2. Core-site.xml – same as coordinator configuration.



NetApp recommends StorageGRID as its primary object storage solution for Datalake and Lakehouse environments. Additionally, NetApp ONTAP is employed for file/object duality. In the context of this document, we have conducted tests on ONTAP S3 in response to a customer request, and it successfully functions as a data source.

Multiple sources setup

1. Configure ONTAP S3 and storageGRID as a s3 source in Dremio.

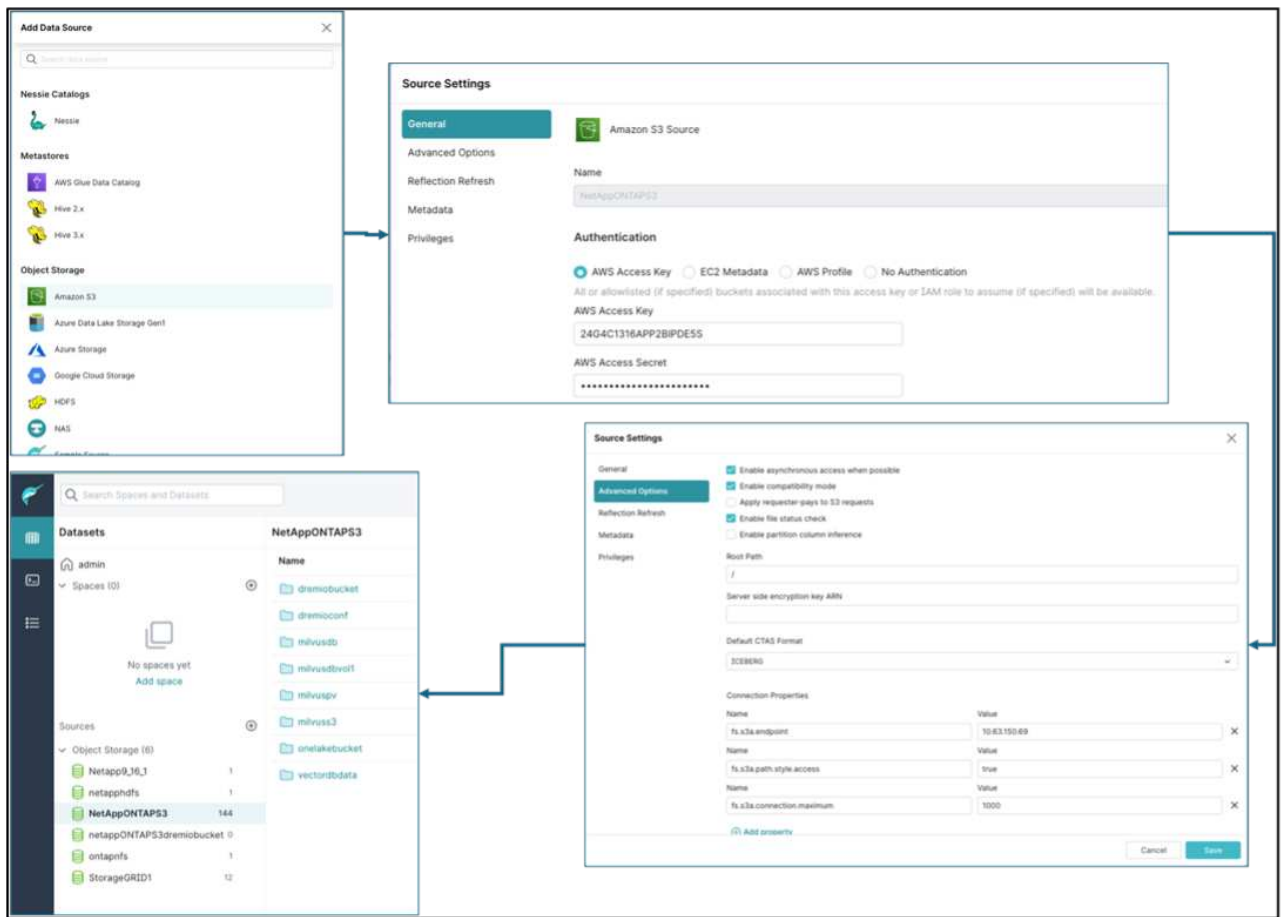
1. Dremio dashboard → datasets → sources → add source.
2. In the general section, please update AWS access and secret key
3. In the advanced option, enable compatibility mode, update connection properties with the below details. The endpoint IP/Name from NetApp storage controller either from ontap S3 or storageGRID.

```

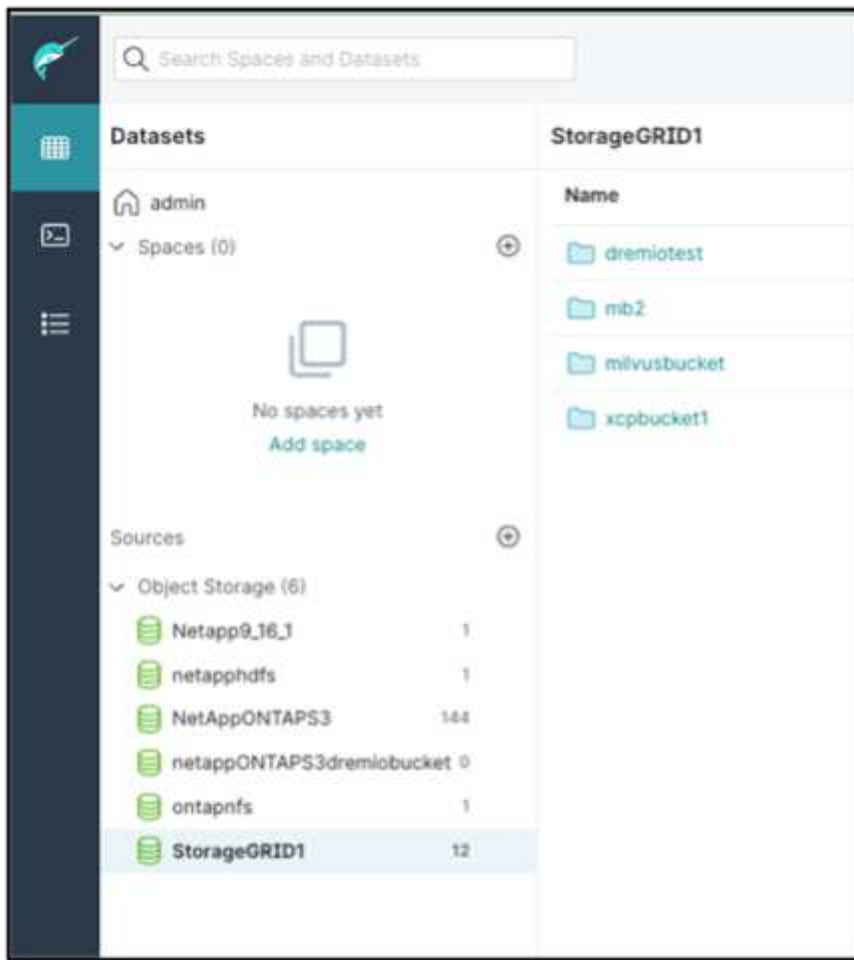
fs.s3a.endpoint = 10.63.150.69
fs.s3a.path.style.access = true
fs.s3a.connection.maximum=1000

```

4. Enable local caching when possible, Max Percent of total available cache to use when possible = 100
5. Then view the list of buckets from NetApp object storage.



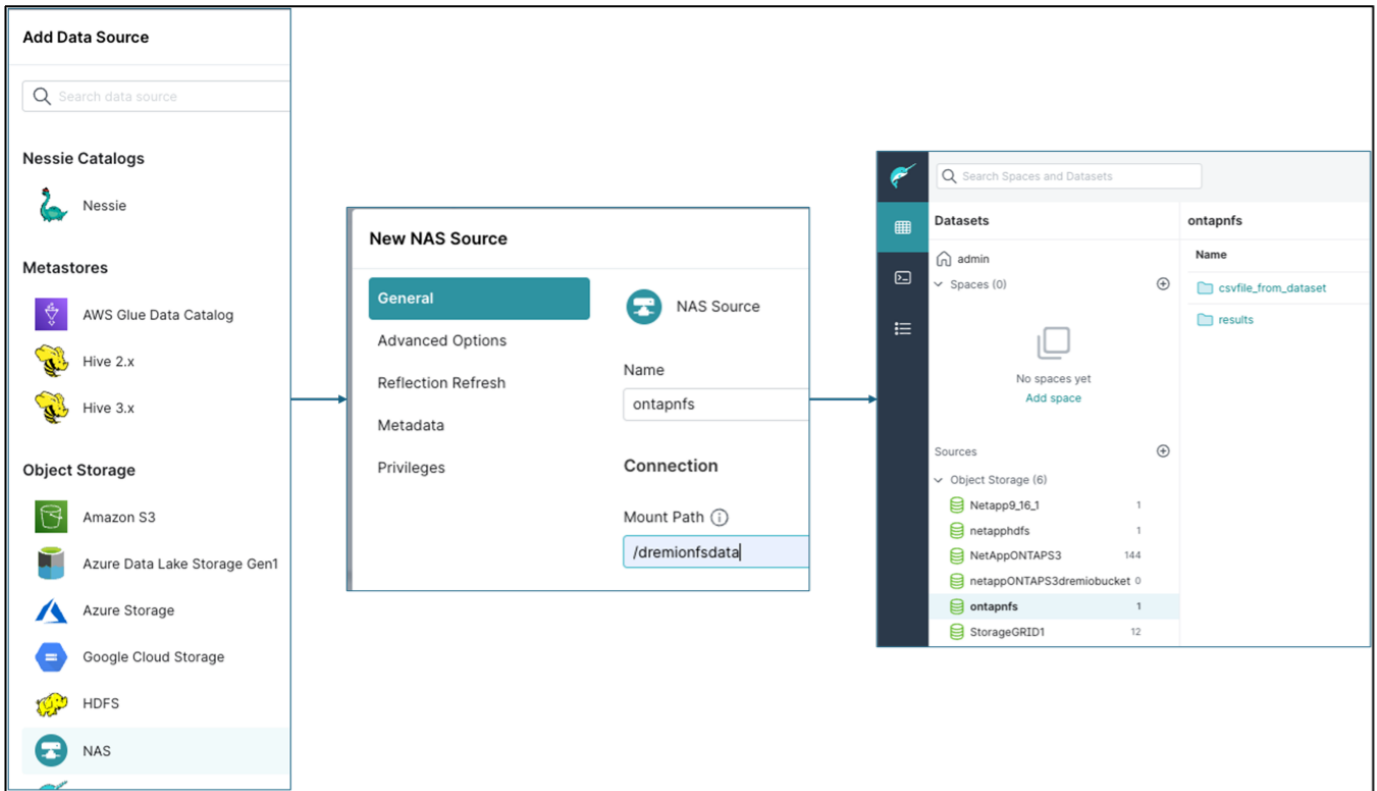
6. Sample view of storageGRID bucket details



2. Configure NAS (specifically NFS) as a source in Dremio.

1. Dremio dashboard → datasets → sources → add source.

2. In the general section, enter the name and NFS mount path. Please make sure the NFS mount path is mounted on the same folder on all the nodes in the Dremio cluster.



+

```

root@hadoopmaster:~# for i in hadoopmaster hadoopnode1 hadoopnode2
hadoopnode3 hadoopnode4; do ssh $i "date;hostname;du -hs
/opt/dremio/data/spill/ ; df -h //dremionfsdata "; done
Fri Sep 13 04:13:19 PM UTC 2024
hadoopmaster
du: cannot access '/opt/dremio/data/spill/': No such file or directory
Filesystem                Size      Used Avail Use% Mounted on
10.63.150.69:/dremionfsdata 2.1T    921M   2.0T   1% /dremionfsdata
Fri Sep 13 04:13:19 PM UTC 2024
hadoopnode1
12K /opt/dremio/data/spill/
Filesystem                Size      Used Avail Use% Mounted on
10.63.150.69:/dremionfsdata 2.1T    921M   2.0T   1% /dremionfsdata
Fri Sep 13 04:13:19 PM UTC 2024
hadoopnode2
12K /opt/dremio/data/spill/
Filesystem                Size      Used Avail Use% Mounted on
10.63.150.69:/dremionfsdata 2.1T    921M   2.0T   1% /dremionfsdata
Fri Sep 13 16:13:20 UTC 2024
hadoopnode3
16K /opt/dremio/data/spill/
Filesystem                Size      Used Avail Use% Mounted on
10.63.150.69:/dremionfsdata 2.1T    921M   2.0T   1% /dremionfsdata
Fri Sep 13 04:13:21 PM UTC 2024
node4
12K /opt/dremio/data/spill/
Filesystem                Size      Used Avail Use% Mounted on
10.63.150.69:/dremionfsdata 2.1T    921M   2.0T   1% /dremionfsdata
root@hadoopmaster:~#

```

Solution verification overview

In this section, we have executed SQL test queries from multiple sources to verify the functionality, test and verify the spillover to NetApp storage.

SQL query on Object storage

1. Set the memory to 250GB per server in dremio.env


```

root@hadoopmaster:~# for i in hadoopmaster hadoopnode1 hadoopnode2
hadoopnode3 hadoopnode4; do ssh $i "hostname; grep -i
DREMIO_MAX_MEMORY_SIZE_MB /opt/dremio/conf/dremio-env; cat /proc/meminfo
| grep -i memtotal"; done
hadoopmaster
#DREMIO_MAX_MEMORY_SIZE_MB=120000
DREMIO_MAX_MEMORY_SIZE_MB=250000
MemTotal:          263515760 kB
hadoopnode1
#DREMIO_MAX_MEMORY_SIZE_MB=120000
DREMIO_MAX_MEMORY_SIZE_MB=250000
MemTotal:          263515860 kB
hadoopnode2
#DREMIO_MAX_MEMORY_SIZE_MB=120000
DREMIO_MAX_MEMORY_SIZE_MB=250000
MemTotal:          263515864 kB
hadoopnode3
#DREMIO_MAX_MEMORY_SIZE_MB=120000
DREMIO_MAX_MEMORY_SIZE_MB=250000
MemTotal:          264004556 kB
node4
#DREMIO_MAX_MEMORY_SIZE_MB=120000
DREMIO_MAX_MEMORY_SIZE_MB=250000
MemTotal:          263515484 kB
root@hadoopmaster:~#

```

2. Check the spill over location (`${DREMIO_HOME}/dremiocache`) in `dremio.conf` file and storage details.

```

paths: {
  # the local path for dremio to store data.
  local: "${DREMIO_HOME}"/dremiocache"

  # the distributed path Dremio data including job results, downloads,
  uploads, etc
  #dist: "hdfs://hadoopmaster:9000/dremiocache"
  dist: "dremioS3:///dremioconf"
}

services: {
  coordinator.enabled: true,
  coordinator.master.enabled: true,
  executor.enabled: false,
  flight.use_session_service: false
}

zookeeper: "10.63.150.130:2181,10.63.150.153:2181,10.63.150.151:2181"
services.coordinator.master.embedded-zookeeper.enabled: false

```

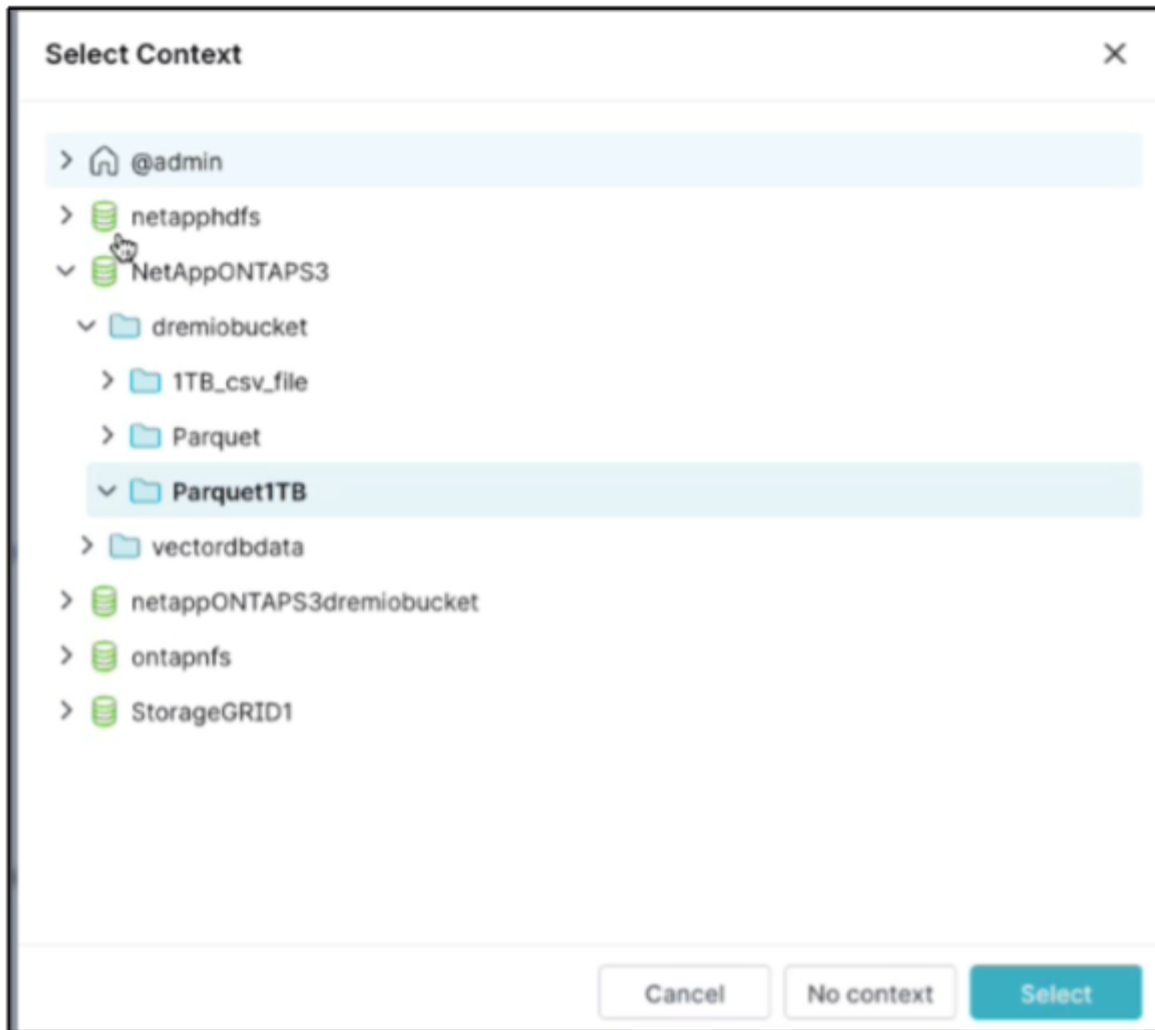
3. Point the Dremio spill over location to NetApp NFS storage

```


root@hadoopnode1:~# ls -ltrh /dremiocache
total 4.0K
drwx----- 3 nobody nogroup 4.0K Sep 13 16:00 spilling_stlrx2540m4-12-
10g_45678
root@hadoopnode1:~# ls -ltrh /opt/dremio/dremiocache/
total 8.0K
drwxr-xr-x 3 dremio dremio 4.0K Aug 22 18:19 spill_old
drwxr-xr-x 4 dremio dremio 4.0K Aug 22 18:19 cm
lrwxrwxrwx 1 root root 12 Aug 22 19:03 spill -> /dremiocache
root@hadoopnode1:~# ls -ltrh /dremiocache
total 4.0K
drwx----- 3 nobody nogroup 4.0K Sep 13 16:00 spilling_stlrx2540m4-12-
10g_45678
root@hadoopnode1:~# df -h /dremiocache
Filesystem                                Size  Used Avail Use% Mounted on
10.63.150.159:/dremiocache_hadoopnode1  2.1T  209M  2.0T   1%
/dremiocache
root@hadoopnode1:~#

```

4. Select the context. In our test, we ran the test against TPCDS generated parquet files residing in ONTAP S3. Dremio Dashboard → SQL runner → context → NetAppONTAPS3→Parquet1TB



1. Run the TPC-DS query67 from the Dremio dashboard



1. Check that the job is running on all executor. Dremio dashboard → jobs → <jobid> → raw profile → select EXTERNAL_SORT → Hostname

Raw Profile

04-xx-04 - FILTER

04-xx-05 - WINDOW

04-xx-06 - EXTERNAL_SORT

Thread	Setup Time	Process Time	Wait Time	Max Batches	Max Records	Peak Memory	Hostname	Record Processing Rate	Operator State	Last Schedule Time
04-00-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540-110-10g	0	CAN_CONSUME	16:35:54
04-01-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540m4-04-10g	0	CAN_CONSUME	16:35:54
04-02-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540m4-12-10g	0	CAN_CONSUME	16:35:54
04-03-06	0.017s	0.000s	0.000s	0	0	128KB	stnx2540m4-13-10g	0	CAN_CONSUME	16:35:54
04-04-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540-110-10g	0	CAN_CONSUME	16:35:54
04-05-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540m4-04-10g	0	CAN_CONSUME	16:35:54
04-06-06	0.027s	0.000s	0.000s	0	0	128KB	stnx2540m4-12-10g	0	CAN_CONSUME	16:35:54
04-07-06	0.000s	0.000s	0.000s	0	0	128KB	stnx2540m4-13-10g	0	CAN_CONSUME	16:35:54

1. When the SQL query running, you can check the split folder for data caching in NetApp storage controller.

```

root@hadoopnode1:~# ls -ltrh /dremiocache
total 4.0K
drwx----- 3 nobody nogroup 4.0K Sep 13 16:00 spilling_stlrx2540m4-12-10g_45678
root@hadoopnode1:~# ls -ltrh /dremiocache/spilling_stlrx2540m4-12-10g_45678/
total 4.0K
drwxr-xr-x 2 root daemon 4.0K Sep 13 16:23 1726243167416

```

2. The SQL query completed with spill over

Job ID	User	Dataset	Query Type	Queue	Start Time	Duration	SQL
19383301-5cd8-0a48-1e38-e25b414d000	admin	store_sales_partition	JDBC Client	High Cost User Q...	08/22/2024, 19:42:54	00:08:25 R	/SF1K_Q67/ select * from (select i_category, i_class, i_brand, i_product_name, d_year, d_qoy, d_m...
193844f3-285d-e07c-5277-48a86168d200	admin	store_sales_partition	JDBC Client	High Cost User Q...	08/22/2024, 18:00:44	00:08:23 R	/SF1K_Q67/ select * from (select i_category, i_class, i_brand, i_product_name, d_year, d_qoy, d_m...
1938650f-0f9a-e205-6ea3-673aaa3c7a00	admin	store_sales_partition	JDBC Client	High Cost User Q...	08/22/2024, 16:09:20	00:08:26 R	/SF1K_Q67/ select * from (select i_category, i_class, i_brand, i_product_name, d_year, d_qoy, d_m...
19387983-2031-164f-cd9e-57c6c287bd00	admin	store_sales_partition	UI (run)	High Cost User Q...	08/22/2024, 14:42:04	00:07:26 R	/SF1K_Q67/ select * from (select i_category, i_class, i_brand, i_product_name, d_year, d_qoy, d_m...
19387b04-3dc3-34db-13a5-d7550fa4a00	admin	store_sales_partition	UI (run)	High Cost User Q...	08/22/2024, 14:22:51	00:07:48 R	/SF1K_Q67/ select * from (select i_category, i_class, i_brand, i_product_name, d_year, d_qoy, d_m...

3. Job completion summary.

Summary

Status: COMPLETED

Total Memory: 287.16 GB

CPU Used: 02h:18m:52s

Query Type: UI (run)

Start Time: 08/26/2024 12:35:53

Duration: 08m:25s

Wait on Client: <1s

User: admin

Queue: High Cost User Queries

Input: 21.32 GB / 563.2M Rows

Output: 6.92 KB / 100 Rows

Total Execution Time 08m:25s (100%)

Pending	2ms (0.00%)
Metadata Retrieval	22ms (0.00%)
Planning	140ms (0.03%)
Queued	30ms (0.01%)
Execution Planning	116ms (0.02%)
Starting	569ms (0.11%)
Running	8m:24s (99.83%)

Submitted SQL

```
1 /*SF1K_Q67*/
2 select *
3 from (select i_category
4         ,i_class
5         ,i_brand
6         ,i_product_name
7         ,d_year
8         ,d_qoy
9         ,d_moy
```

Queried Datasets

- store_sales_partition
NetAppONTAPS3.dremiobucket.Parquet1TB
- date_dim
NetAppONTAPS3.dremiobucket.Parquet1TB
- store
NetAppONTAPS3.dremiobucket.Parquet1TB

Show more >

Scans

- store_sales_partition
- date_dim
- store
- item

4. Check the spilled data size

EXTERNAL_SORT 04-06



Runtime	1.68m (100%)
Startup	49.09ms (0.05%)
Processing	39.62s (39.36%)
IO Wait	1.02m (60.6%)

Overview/Main

Batches Processed:	104333
Records Processed:	387.6M
Peak Memory:	199 MB
Bytes Sent:	44 GB
Number of Threads:	180

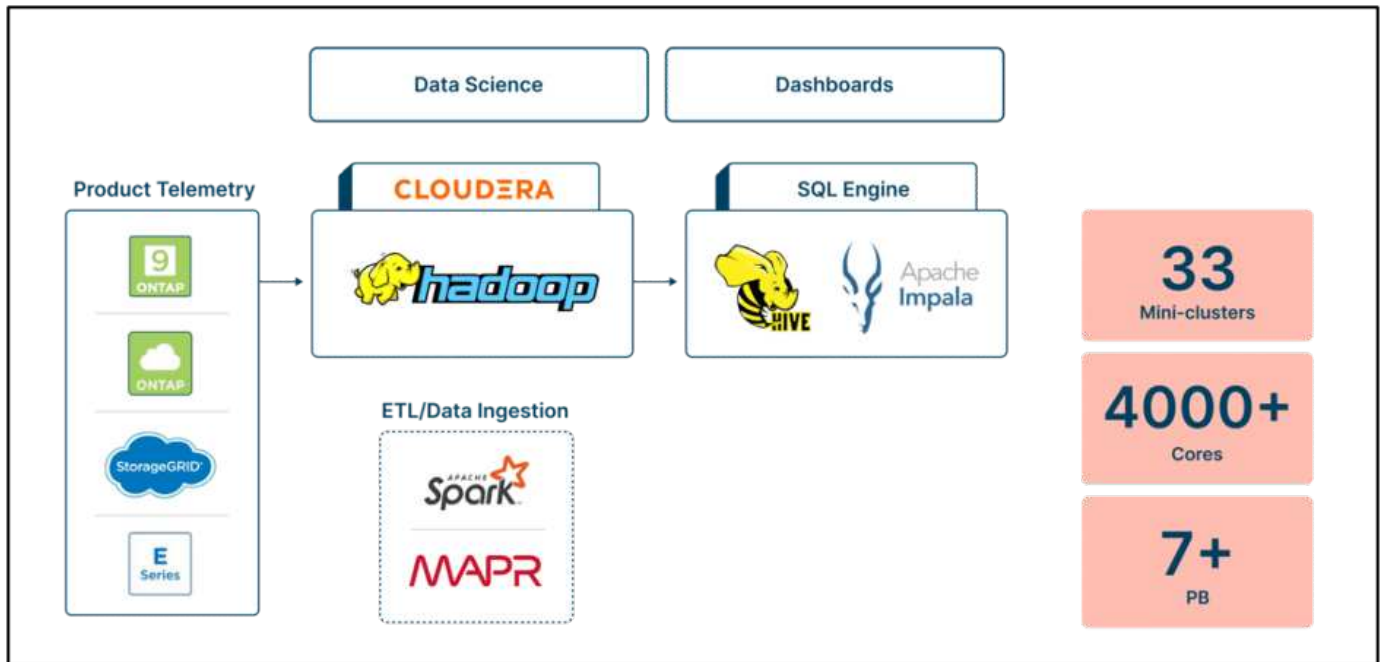
Operator Statistics

Merge Time Nanos:	0ns
Spill Count:	360
Spill Time Nanos:	37.68m
Total Spilled Data Size:	20,339,702,765
Batches Spilled:	97,854

The same procedure applicable for NAS and StorageGRID Object storage.

Customer Use Cases

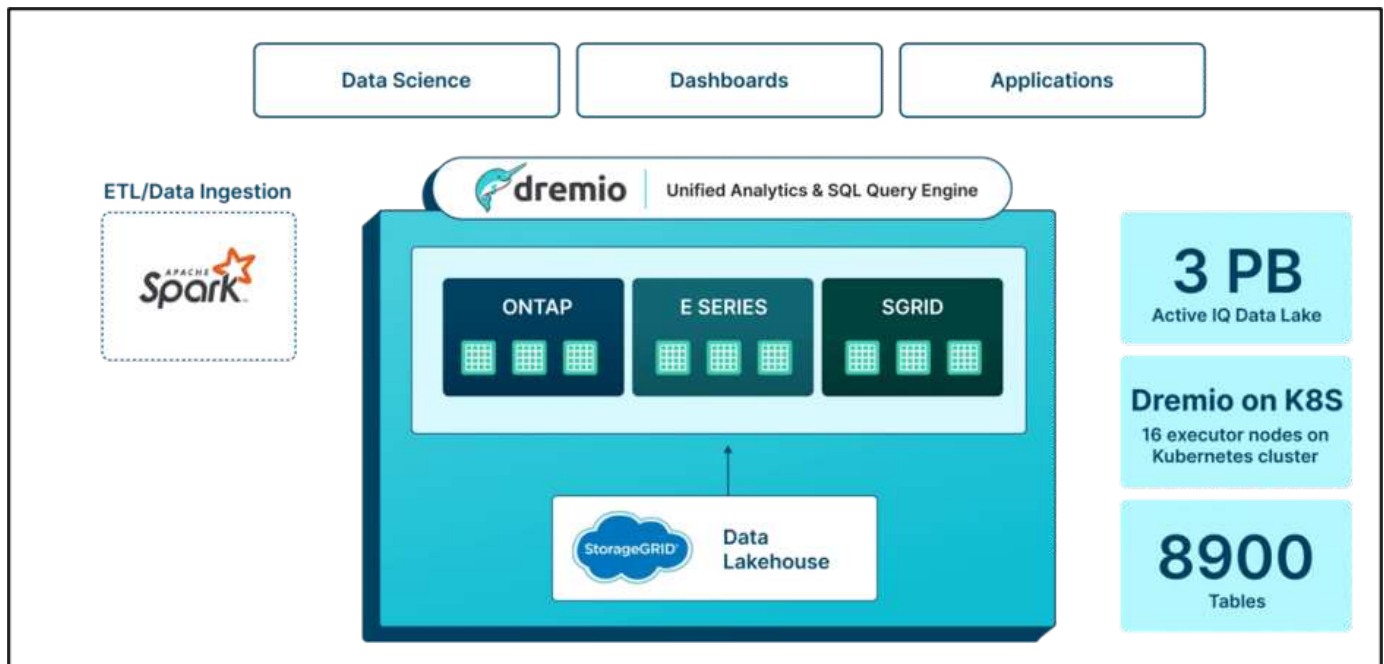
NetApp ActiveIQ Use Case



Challenge: NetApp's own internal Active IQ solution, initially designed for supporting numerous use cases, had evolved into a comprehensive offering for both internal users and customers. However, the underlying Hadoop/MapR-based backend infrastructure posed challenges around cost and performance, due to the rapid growth of data and the need for efficient data access. Scaling storage meant adding unnecessary computing resources, resulting in increased costs.

Additionally, managing the Hadoop cluster was time-consuming and required specialized expertise. Data performance and management issues further complicated the situation, with queries taking an average of 45 minutes and resource starvation due to misconfigurations. To address these challenges, NetApp sought an alternative to the existing legacy Hadoop environment and determined a new modern solution built on Dremio would reduce costs, decouple storage and compute, improve performance, simplify data management, offer fine-grained controls, and provide disaster recovery capabilities.

Solution:



Dremio enabled NetApp to modernize its Hadoop-based data infrastructure in a phased approach, providing a roadmap for unified analytics. Unlike other vendors that required significant changes to data processing, Dremio seamlessly integrated with existing pipelines, saving time and expenses during migration. By transitioning to a fully containerized environment, NetApp reduced management overhead, improved security, and enhanced resilience. Dremio's adoption of open ecosystems like Apache Iceberg and Arrow ensured future-proofing, transparency, and extensibility.

As a replacement for the Hadoop/Hive infrastructure, Dremio offered functionality for secondary use cases through the semantic layer. While the existing Spark-based ETL and data ingestion mechanisms remained, Dremio provided a unified access layer for easier data discovery and exploration without duplication. This approach significantly reduced data replication factors and decoupled storage and computing.

Benefits:

With Dremio, NetApp achieved significant cost reductions by minimizing compute consumption and disk space requirements in their data environments. The new Active IQ Data Lake is comprised of 8,900 tables holding 3 petabytes of data, compared to the previous infrastructure with over 7 petabytes. The migration to Dremio also involved transitioning from 33 mini-clusters and 4,000 cores to 16 executor nodes on Kubernetes clusters. Even with significant decreases in computing resources, NetApp experienced remarkable performance improvements. By directly accessing data through Dremio, query runtime decreased from 45 minutes to 2 minutes, resulting in 95% faster time to insights for predictive maintenance and optimization. The migration also yielded a more than 60% reduction in compute costs, more than 20 times faster queries, and more than 30% savings in total cost of ownership (TCO).

Auto Parts Sales Customer Use Case.

Challenges: Within this global auto parts sales company, executive and corporate financial planning and analysis groups were unable to get a consolidated view of sales reporting and were forced into reading the individual line of business sales metrics reports and attempting to consolidate them. This resulted in customers making decisions with data that was at least one day old. The lead times to get new analytics insights would typically take more than four weeks. Troubleshooting data pipelines would require even more time, adding an additional three days or more to the already long timeline. The slow report development process as well as report performance forced the analyst community to continually wait for data to process or load, rather than enabling them to find new businesses insights and drive new business behavior. These troubled environments were composed of numerous different databases for different lines of businesses, resulting in numerous data silos. The slow and fragmented environment complicated data governance as there were too many ways for

analysts to come up with their own version of the truth versus a single source of truth. The approach cost over \$1.9 million in data platform and people costs. Maintaining the legacy platform and filling data requests required seven Field Technical Engineers (FTEs) per year. With data requests growing, the data intelligence team could not scale the legacy environment to meet future needs

Solution: Cost-effectively store and manage large Iceberg tables in NetApp Object Store. Build data domains using Dremio's semantic layer, allowing business users to easily create, search, and share data products.

Benefits to customer:

- Improved and optimized existing data architecture and reduced time to insights from four weeks to just hours
- Reduced troubleshooting time from three days to only hours
- Decreased data platform and management costs by more than \$380,000
- (2) FTEs of Data Intelligence effort saved per year

Conclusion

In conclusion, this technical report has provided comprehensive deployment details of q Hybrid Iceberg Lakehouse with Dremio in conjunction with various data sources from NetApp storage controllers, including ONTAP S3, NAS, and StorageGRID. The deployment process was successfully executed, and the TPC-DS benchmarking tool was utilized to perform 99 SQL queries across the different data sources. The report has also explored customer use cases within NetApp, demonstrating the versatility and effectiveness of Dremio in meeting diverse business requirements. Additionally, a specific use case involving an auto parts sales customer was examined, highlighting the practical application and benefits of leveraging Dremio for data analytics and insights.

Overall, this document serves as a valuable resource for understanding the deployment and usage of Dremio with NetApp storage controllers, showcasing its capabilities and potential for driving data-driven decision-making and optimization in various industries.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Zookeeper installation

<https://medium.com/@ahmetfurkandemir/distributed-hadoop-cluster-1-spark-with-all-dependencies-03c8ec616166>

- Dremio

<https://docs.dremio.com/current/get-started/cluster-deployments/deployment-models/standalone/standalone-tarball/>

- Configuring Dremio with storageGRID

<https://docs.netapp.com/us-en/storagegrid-enable/tools-apps-guides/configure-dremio-storagegrid.html#configure-dremio-data-source>

- NetApp use case

Copyright information

Copyright © 2025 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.