

DELL EMC ISILON MULTITENANCY FOR HADOOP BIG DATA ANALYTICS

ABSTRACT

The Dell EMC[®] Isilon[®] scale-out storage platform provides multitenancy through access zones that segregate tenants and their data sets. An access zone presents a portion of an Isilon cluster as a secure virtual storage region with a unique HDFS root directory for the zone's tenant. With NFS, SMB, and HDFS access to each zone, an Isilon cluster delivers a scalable multitenant storage solution for Hadoop and other analytics applications.

December 2016

The information in this publication is provided “as is.” DELL EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any DELL EMC software described in this publication requires an applicable software license.

DELL EMC², DELL EMC, the DELL EMC logo are registered trademarks or trademarks of DELL EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2016 DELL EMC Corporation. All rights reserved. Published in the USA. <12/16> <white paper> <H13564>

DELL EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

DELL EMC is now part of the Dell group of companies.

Table of Contents

EXECUTIVE SUMMARY	4
PROBLEMS IN ENTERPRISE DATA ANALYTICS.....	4
Multitenancy Requirements.....	5
Overview of Isilon Scale-out NAS for Big Data Analytics	5
How OneFS Meets Multitenancy Requirements.....	6
SEGREGATING TENANTS AND DATA SETS WITH ACCESS ZONES.....	6
Using OneFS Access Zones	8
Setting a Unique HDFS Root Directory for Each Tenant.....	8
HANDLING MIXED WORKFLOWS IN AN ACCESS ZONE.....	8
Monitoring HDFS in an Access Zone with Apache Ambari.....	9
Working with Different Versions of Hadoop Simultaneously.....	9
Managing Concurrent Hadoop Jobs and Other Resources with YARN.....	9
MANAGING DATA SETS	10
Providing Multiprotocol Data Access to Users and Applications	10
Storing and Analyzing Different Data Sets	10
Storage Pools For Different Data Sets	10
Isolating Data Sets to Automate Shelf-Life Management	11
Associating Networking Resources with Access Zones	11
SECURING ACCESS ZONES FOR ANALYTICS MULTITENANCY	15
Multitenancy for Directory Services	16
Access Control	16
WORM for Data at Rest	17
Auditing	17
USE CASES.....	17
Moving from Silos to Efficiency and Accessibility	17
Sharing Data Among Tenants	18
CONCLUSION	18

EXECUTIVE SUMMARY

The notion of multitenancy plays a key role in the formation of an enterprise data hub or shared infrastructure in which business units pool their data for storage and analysis. Multitenancy means that the storage platform can, among other requirements, segregate business units and data sets while working with different data access protocols and analytics applications.

By supporting multitenancy, the Dell EMC Isilon storage platform lays the foundation for an organization's business units to become data-driven enterprises by rapidly extracting value from their data with Hadoop or other analytics applications.

A key big data analytics platform, Apache Hadoop comprises the Hadoop Distributed File System, or HDFS, a storage system for vast amounts of data, and MapReduce, a processing paradigm for data-intensive computational analysis. Hadoop is supplemented by projects that improve distributed data analytics; the projects include Hive, Pig, Sqoop, Flume, Oozie, Whirr, HBase, Zookeeper, and Spark.

One benefit of multitenancy is analytics as a service, including Hadoop as a service, whether for internal business units or external customers. Another advantage is that with multitenancy and its native HDFS support, the Isilon storage platform sets your information technologies on a course toward a dynamic next-generation architecture that can handle emerging workloads and big data.

With an Isilon storage cluster, multitenancy is implemented through secure access zones. An access zone segregates a portion of the cluster and presents it as a secure virtual storage region. In an access zone, you can connect to directory services like Active Directory, control the access of tenants, consolidate SMB file shares and NFS exports, and set up an autonomous HDFS root directory for each tenant. Isilon access zones work with most of the Hadoop projects and applications as well as YARN, which separates resource management from computational processing to take Hadoop beyond MapReduce.

The distributed Isilon OneFS operating system implements the server-side operations of the HDFS protocol on every node in a cluster. With its native support for HDFS, a Dell EMC Isilon cluster provides a scalable, secure multitenant storage solution for Hadoop and other big data analytics applications.

PROBLEMS IN ENTERPRISE DATA ANALYTICS

Most organizations have yet to implement a single, scalable storage system that exposes all of their key data to analytics applications. Their data is often stuck in silos, inefficient scale-up storage systems that are costly to manage and difficult to scale. The inability to scale out such traditional storage systems hinders the progression to a platform that makes a data-driven enterprise possible.

There are other problems that undermine enterprise data analytics. How can you store data in a highly scalable way that makes it available not only to various analytics applications but also to different business units? Some enterprises aim to solve this problem by implementing HDFS as a mass storage system for all their enterprise data. But using HDFS as an enterprise data hub creates a range of problems that ultimately undermine multitenancy for all business units.

Although HDFS works with more analytics applications beyond MapReduce, including HAWQ, Impala, Spark, Tez, Storm, and GemFire XD—it can impede seamless access by applications outside the Apache Hadoop family, such as analytics applications that use NFS or SMB. Serving the data to other applications often requires an extract, transform, load (ETL) workflow that radically complicates analytics workflows and stifles insights—largely because it takes too long time to load the data. The same problem cuts the other way: Application data that is stored by an NFS or SMB workflow in a storage silo must be extracted and loaded into HDFS for analysis increasing the time to actionable results.

Although these problems take several forms in enterprise environments, many of them hinge on the notion of an enterprise data lake, a shared infrastructure in which business units pool their data for storage and analysis:

- 1) Efficiency in storing data for analysis. Storage silos—coupled with different business units demanding different analytics applications—complicate storage and analysis, bringing costly inefficiencies to data management, data storage, and analytics. A worse effect, however, is that such inefficiencies radically slow down the time it takes to extract valuable insights from data.
- 2) Availability—Are data sets available to all business units and analytics applications without a complicated, costly workflow? For different tenants to be able to analyze a variety of data sets, the data must be made easily available for analysis without time-consuming, costly ELT workflows that do not scale well to handle multiple large data sets. Besides avoiding ELT workflows, there's also the related question of supporting common data access protocols such as NFS and SMB.

- 3) Data ingestion—How do you easily ingest the data into the storage system without time-consuming transformations or migrations? Many enterprises begin to implement big data analytics only to find that data ingestion takes so long that it delays attempts to analyze the data. The cost of an ETL operation can spiral over time because the operation needs to be repeated: As the data expands or changes, it must be loaded again.
- 4) Security—How do you secure data from others in a shared environment? An enterprise content pool or data lake becomes a highly complex undertaking, and with complexity, you see a rise in security risks. If an organization cannot adequately address security for the whole content pool, the organization will split the lake into segments to isolate the access to each segment by business role or need to know—in effect re-creating the information silos that the data lake was intended to eliminate.
- 5) Performance—How do you allocate resources so that multiple business units can run different analytics applications, including resource-intensive frameworks like YARN and MapReduce, against different data sets simultaneously without undermining the performance of concurrent jobs? Building an enterprise data hub with a storage architecture that cannot be scaled out to handle big data will eventually create performance bottlenecks—bottlenecks that get worse as data sets expand. Some traditional scale-up storage architectures, for instance, require manual data migrations among storage components to balance performance for applications. Performance bottlenecks levied by the storage system can reduce performance for concurrent analytics jobs.

These problems in enterprise data analytics give rise to a set of requirements grouped around a loose, nebulous concept: multitenancy. Implicit in the notion of multitenancy is the creation a single, shared pool of stored data, known as an enterprise data lake, that consolidates costly, inefficient storage silos into an efficient, highly scalable storage platform that can be used by multiple data sets, groups of user, and applications.

Multitenancy Requirements

For the purposes of this paper, the term *multitenancy* translates into the following requirements:

- Support many tenants—business units or distinct groups of users—and their data sets on a shared infrastructure.
- Isolate tenants and their data sets with robust security.
- Authenticate tenants with multiple identity management systems, such as LDAP, NIS, and Active Directory.
- Handle mixed workflows—application workloads of different types, such as Hadoop and statistical computing languages.
- Work with different versions of applications simultaneously.
- Work with different versions of Hadoop simultaneously.
- Run many Hadoop jobs concurrently.
- Store a variety of data sets, including unstructured and structured data, from NoSQL to traditional SQL and everything in between.
- Work with multiple data access protocols simultaneously to support tenants who use different applications or different connection methods, such as NFS, SMB, and HDFS.

These requirements combine with the problems outlined in the previous section to form a matrix of issues that a storage solution for enterprise analytics must address.

Overview of Isilon Scale-out NAS for Big Data Analytics

The Dell EMC Isilon multi-protocol scale-out platform provides a single volume that scales out to manage petabytes of data and simplifies the implementation of big data analytics, especially in the form of Hadoop.

The distributed Isilon OneFS operating system combines the resources of an Isilon cluster's nodes into a cohesive storage unit to present a global namespace as a single file system. For multitenancy, the global namespace can be segregated into virtual storage regions with their own security context.

An Isilon cluster provides Hadoop clients with direct access to big data through native support for the Hadoop File System (HDFS). OneFS implements the server-side operations of the HDFS protocol on every node in a cluster to handle calls to NameNodes and read/write requests to DataNodes. Every Isilon node can simultaneously serve as a NameNode and a DataNode.

Storing data in an Isilon scale-out cluster streamlines the entire analytics workflow. Isilon's native HDFS interface eliminates extracting the data from a storage system and loading it into an HDFS file system. You can also store data for analysis, whether with Hadoop or other analytics applications, by using your existing workflows and standard protocols, including SMB, NFS, and Swift. Similarly, Isilon's support for multiple data access protocols eliminates the need to export the results of an analytics job.

OneFS increases the ease and flexibility with which you analyze data while shortening the time it takes to extract value from data. In addition to liberating Hadoop from the bounds of a complex, purpose-built Hadoop storage system, Isilon provides secure multitenancy to enterprise analytics workflows.

How OneFS Meets Multitenancy Requirements

Dell EMC Isilon securely brings Hadoop analytics to each group in an enterprise through its multitenancy capabilities. All groups can be represented equally, with their own space to store and process data for analytics. Multitenancy securely extends the opportunity for solving problems with analytics to a wider set of users—in effect, democratizing Hadoop.

The Dell EMC Isilon scale-out data lake foundation fulfills the following multitenancy requirements:

- Isolates tenants and their data sets with access zones.
- Handles mixed workflows in each access zone, including Hadoop, R, HAWQ, PIG, HIVE, and other applications.
 - Simultaneously runs multiple Hadoop distributions—including Cloudera, Pivotal HD, Apache Hadoop, and Hortonworks—against the same data set at the same time. OneFS also supports HDFS 1.0 and 2.0 simultaneously without migrating data or modifying metadata.
 - Runs many Hadoop jobs concurrently and lets you allocate resources with YARN.
- Stores different data sets in the same access zone, including Impala, video, images, semi-structured data, and unstructured data.
- Stores your analytics data with existing workflows and protocols like NFS, HTTP, and SMB instead of spending time importing and exporting data with HDFS.
- Secures each access zone with Kerberos authentication and other security mechanisms to protect the confidentiality of data.
 - Provides multitenancy for identity management, authentication, and access control.
 - Manages tenants and their data with quotas, storage pools, and other enterprise features.

Segregating Tenants and Data Sets with Access Zones

OneFS access zones lay the foundation for secure multitenancy. An access zone performs two primary functions:

1. An access zone establishes a virtual security context that segregates tenants.
2. An access zone creates a virtual storage region that isolates data sets.

As a virtual security context, an access zone connects to directory services, authenticates users, and controls access to a segment of the file system.

As a virtual storage region, an access zone isolates system connections, directories, and files as well as the application of some enterprise management features like event auditing.

The following sample listing shows how an access zone encapsulates directory services, authentication, auditing, an HDFS root directory, and other features in a security context that contains a virtual storage area:

```
test-cluster-1# isi zone zones list -v

Name: z1 ⓘ
Path: /ifs/z1
```

```
Cache Size: 9.54M

Map Untrusted:

Auth Providers: lsa-ldap-provider:krbldap,
                lsa-krb5-provider:SOUTH.EXAMPLE.COM ❷

NetBIOS Name:

All Auth Providers: Yes

User Mapping Rules: -

Home Directory Umask: 0077

Skeleton Directory: /usr/share/skel

Audit Success: create, delete, rename, set_security, close ❸

Audit Failure: create, delete, rename, set_security, close

HDFS Authentication: all

HDFS Root Directory: /ifs/z1 ❹

WebHDFS Enabled: Yes

HDFS Ambari Server:

HDFS Ambari Namenode:

Syslog Forwarding Enabled: No

Syslog Audit Events: create, delete, rename, set_security

Zone ID: 2
```

- ❶ The name of this access zone.
- ❷ The directory services with which OneFS authenticates users and groups in the access zone.
- ❸ The audit settings that are turned on for the access zone.
- ❹ The unique HDFS root directory for the access zone. Each access zone can have a different HDFS root directory that serves only the zone's tenants.

Although you could isolate stored resources and tenants by using only file permissions and access tokens, access zones are much easier to manage in an enterprise multitenant environment than permissions. Permissions and group memberships remain a viable method of controlling access to data within an access zone. Permissions and group memberships also remain a useful method for controlling access when there is no requirement for multitenancy.

When a Hadoop user connects to an Isilon cluster, OneFS checks the directory services to which the user's access zone is connected for an account for the user. If OneFS finds an account that matches the user's login name, OneFS authenticates the user. During authentication, OneFS creates an access token for the user. The token contains the user's full identity, including group memberships, and OneFS uses the token later to check access to directories and files.

When OneFS authenticates users with different directory services, OneFS maps a user's account from one directory service to the user's accounts in other directory services within an access zone—a process known as user mapping. A Windows user account managed in Active Directory, for example, is mapped by default to a corresponding UNIX account with the same name in NIS or LDAP. With a single token, a user can, if permitted, access files that were stored by a Windows computer over SMB and files that were stored by a UNIX computer over NFS or HDFS.

Using OneFS Access Zones

One purpose of an access zone is to define a list of identity management systems that apply only in the context of a zone that contains SMB shares, NFS exports, or an HDFS root directory. As such, a key use case for an access zone is consolidating two or more file servers into a OneFS cluster and then analyzing the data with Hadoop.

Another key use case is consolidating multiple Hadoop data sets into a single storage system but continuing to expose each data set with a unique root directory and then limiting access to only those who need it. You can set up a cluster to work with multiple identity management systems, NFS namespaces, SMB namespaces, Swift object storage namespaces, and HDFS namespaces and then segment them into access zones. In other words, an access zone can simultaneously contain an HDFS namespace and a namespace for Swift, NFS, and SMB files. Within the access zone, you can authenticate users and groups by simultaneously using Active Directory, NIS, and LDAP.

By placing authentication, access control, directory structures, and enterprise management features in the context of an access zone, OneFS lets you target the following high-level use cases with relative ease:

- Isolate directories and files in virtual storage regions without the overhead of managing individual volumes
- Consolidate SMB file servers and shares into a single storage cluster and then segment the shares by tenant or directory
- Consolidate NFS storage silos into a single cluster and then segment exports by tenant, directory, or data set
- Consolidate Multiple HDFS workflows on a single cluster
- Store different user databases on a single storage system while segmenting each database
- Migrate a scale-up storage system to an Isilon cluster to ease management, reduce operational expenses, and improve storage efficiency
- Set upper limits on resource usage by managing quotas, collections of disks, network interfaces, and other factors
- Avoid collisions of user identifiers when different users have the same UID
- Provision and manage a different Hadoop root directory for different tenants

Setting a Unique HDFS Root Directory for Each Tenant

Perhaps most importantly for Hadoop, each access zone can contain a unique HDFS root directory for the tenant in that zone. For HDFS, the ability to enforce unique HDFS root directories in each access zone allows you to point multiple Hadoop compute clusters at the same HDFS storage system. Access zones can enforce a one-to-one mapping that is transparent to the clusters of compute clients.

It is easy to create an access zone, as the following example shows. This command sets the HDFS root directory for the default access zone, which is named System.

```
isiloncluster1-1# isi zone zones modify System
```

```
--hdfs-root-directory /ifs/isiloncluster1/zone1/hadoop
```

By using this command, you can set a different HDFS root directory for each access zone that you create.

Handling Mixed Workflows in an Access Zone

Every access zone works with the following protocols: NFS, SMB, Swift, and HDFS. With multiprotocol support, an access zone can handle a mix of workflows and analytics applications, including not only Hadoop projects but also other analytics applications like the statistical computing language R. The Swift protocol for object storage empowers you to automate the collection of petabytes of data and store them in an Isilon data lake for later analysis.

Within the Apache Hadoop set of projects and related applications, OneFS works with the following:

- YARN
- Pig

- Hive
- WebHDFS
- Sqoop2
- Solr Sqoop
- Spark
- Flume
- Sentry
- Hue
- Oozie
- Whirr
- HBase
- Impala
- Zookeeper

The latest version of OneFS supports the most common versions of the HDFS protocol, including 2.2, 2.3, and 2.4. The OneFS HDFS implementation also works with Ambari, Kerberos authentication, and Kerberos impersonation.

With Kerberos impersonation, OneFS lets another system—such as Oozie, Hue, Impala, or Sentry—act as a proxy-user or super-user to perform actions on behalf of another user. In addition, OneFS works with so-called metaschedulers that manage Hadoop jobs, such as Oracle Grid Engine and Omero from CSR4.

Monitoring HDFS in an Access Zone with Apache Ambari

Apache Ambari advances the management and monitoring of Hadoop clusters. The OneFS Ambari agent collects information about HDFS connections and Hadoop services in each access zone and sends the information to an external Ambari server. By connecting to an Ambari server with a web browser, you can monitor the status of the HDFS connections and manage such services as YARN and Oozie. The OneFS Ambari agent works with Ambari server versions 1.5.1.110 and 1.6.0.

Working with Different Versions of Hadoop Simultaneously

OneFS interoperates with most major Hadoop distributions—including Cloudera, Pivotal HD, Apache Hadoop, and Hortonworks—against the same data set at the same time. OneFS supports HDFS 1.0 and 2.0 simultaneously without migrating data or modifying metadata.

A Dell EMC Isilon cluster is platform agnostic for compute, and there is no vendor lock-in: You can run most of the common Hadoop distributions with an Isilon cluster, including Apache Hadoop, Hortonworks, Cloudera, and Pivotal HD.

Clients running different Hadoop distributions or versions can connect to the cluster simultaneously. For example, you can point both Cloudera and Pivotal HD at the same data on your Isilon cluster and run MapReduce jobs from both distributions at the same time.

Managing Concurrent Hadoop Jobs and Other Resources with YARN

Yarn allocates resources for Hadoop applications. Because OneFS supports versions 2 of HDFS, which introduced YARN, OneFS works with YARN.

A sub-project of Apache Hadoop, YARN separates resource management from computational processing to expand interactional patterns beyond MapReduce for data stored in HDFS. With the introduction of YARN's more general processing platform, the JobTracker became the ResourceManager and the TaskTracker, the NodeManager. With YARN, a single ResourceManager works with the NodeManager process, which runs on all the compute nodes. The NodeManager executes tasks submitted by users in isolated containers. The ResourceManager schedules the tasks on the NodeManagers.

As a compute-side component, YARN works seamlessly with OneFS. To serve multitenancy scenarios, each tenant can run a NodeManager and a ResourceManager for each OneFS access zone.

Managing Data Sets

The following sections describe how an Isilon cluster can set up pools of storage and networking resources to support multitenancy.

Providing Multiprotocol Data Access to Users and Applications

An Isilon cluster supports the workflows of all your tenants or departments. Instead of running HDFS copy operations to move your data to Hadoop clients in your compute grid, you can continue to store data using existing workflows while achieving faster time to results with data-in-place analytics. A Dell EMC Isilon cluster provides multiprotocol data access with SMB, NFS, and Swift as well as HDFS. OneFS authenticates tenants with the directory services that the workflows are already tied to, such as LDAP for NFS workflows, Active Directory for SMB workflows, and MIT Kerberos for HDFS workflows.

Why is supporting existing workflows so important to advancing data analytics within all the business units in your organization? Consolidation in an enterprise data lake empowers a business to analyze its data without a heavy dependence on IT, reduce complexity, and achieve faster time to results. Instead of relying on Hadoop application developers to painstakingly deploy HDFS and move data to it, business personnel can use their existing workflows—especially SMB and NFS—to collect and manage the data that they want to analyze. The implications of this shift not only help prepare a business for the future of data-driven analysis, but also open up data analysis to members of the core business.

For example, if your technical support staff has collected log data from a year's worth of support cases and stored the log files on an Isilon cluster, support managers can analyze the data in place to identify patterns of recurring problems to help predict when a problem might arise and address it before it spirals out of control.

A 2012 article in "MIT Sloan Management Review" sums up this shift: "Advanced organizations are moving analytics from IT into their core business and operational functions. As big data evolves, a new information ecosystem is also evolving, a network that is continuously sharing information, optimizing decisions, communicating results and generating new insights for businesses."¹

With its secure multitenancy analytics capabilities, an Isilon scale out data lake lays the foundation for such an ecosystem. For more information on the benefits of data-in-place analytics, please see the white paper [Dell EMC Isilon Scale-Out NAS for In-Place Hadoop Data Analytics](#).

Storing and Analyzing Different Data Sets

By supporting HDFS, Swift, and other common data access protocols, OneFS can store data sets originating from different systems and applications. Once stored, applications can analyze the data through HDFS, NFS, or another protocol without a system administrator having to migrate or copy the data.

Return Path, an email intelligence company, provides an example. The company struggled to integrate unique data sets composed of hundreds of millions of small email files into its NFS and Hadoop environments. Analyzing the data was hindered by having to make a separate copy of the data set. To solve these problems, Return Path consolidated the data in an Isilon cluster. The Isilon cluster's built-in support for HDFS not only increased Return Path's ability to scale their storage system to meet changing data sets, but also fostered analytics by enabling Return Path's Hadoop system to directly access and analyze the data sets.

By consolidating all its data repositories into an Isilon cluster—in effect, creating an enterprise data lake—Return Path makes the data sets accessible to its Hadoop analytics team, product development teams, and external customers. Before consolidating the data on Isilon, see [Email intelligence firm captures competitive advantage](#) with Hadoop analytics and Dell EMC Isilon.

Similarly, WGSN INstock, a subsidiary of Top Right Group, uses an Isilon cluster to handle multiple data streams in a multitenant context. Doing so has opened up opportunities to provide analytical services both internally and externally. For more information, see [Supporting rapid launch of game-changing market intelligence service for retailers with Hadoop](#).

Storage Pools For Different Data Sets

OneFS includes several ways to create pools of storage that group data sets by attributes. With Isilon SmartPools™ technology, for instance, you can create node pools, file policies, and tiers of storage. Node pools segment nodes into groups so that you can align a

¹ Davenport, Thomas H., Paul Barth, and Randy Bean. "How 'Big Data' Is Different," MIT Sloan Management Review. Fall 2012.

dataset with its performance requirements. File policies can isolate files by type, path, size, and other attributes. Tiers optimize the storage of data by need, such as a frequently used high-speed tier or a rarely accessed archive.

With a storage pool in place, OneFS SmartConnect™ zones can associate a set of analytics compute clients with each data set to optimize the performance of the compute jobs that analyze each set.

You can also build node pools and policies to streamline an analytics workflow. Because you can combine Dell EMC Isilon nodes from the S-Series, the X-Series, and the NL-Series into a cluster, you can set up pools of nodes to accelerate access to important working sets while placing inactive data in more cost-effective storage.

For example, you can group S-Series nodes with 900 GB SAS drives and 400 GB SSDs per node in one pool, while you put NL-Series nodes with 3 TB SATA drives in another. As each type of node has a different capacity-to-performance ratio, you can assign different data sets to node pools that meet the data set's capacity or performance requirements.

Isolating Data Sets to Automate Shelf-Life Management

A policy can move older files, as identified by the last-modified or last-accessed date, to a data storage target of NL-Series disks to reserve the capacity in the S-Series node pool for newer data. When a policy moves a file or directory from one disk pool to another, the file's location in the file system tree remains the same, and you do not need to reconfigure clients to access the file in the new location. A file pool policy can apply to directories or files, and it can filter files by file name, path, type, and size, as well as other file attributes.

In this way, node pools can isolate datasets to increase performance for important data and decrease storage costs for immaterial data.

Associating Networking Resources with Access Zones

OneFS contains a virtual racking feature that can associate network interfaces with an access zone. The rack associates the clients' IP addresses with a pool of DataNodes so that when a client connects to the cluster, OneFS assigns the connection to one of the DataNodes in the pool. Such a network topology can look something like this:

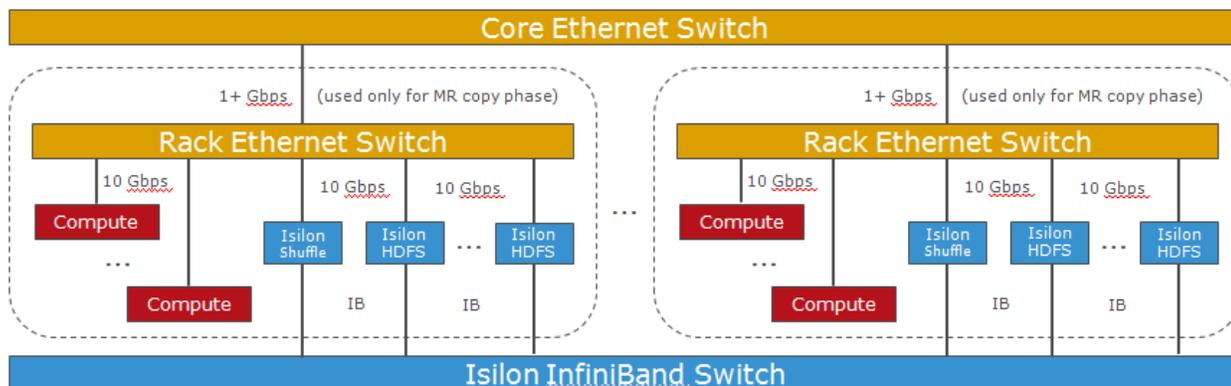


Figure 1. Sample network topology illustrating OneFS's virtual racking feature.

A Hadoop client connects over HDFS to the DataNodes with interfaces that are assigned to the pool. After you add a pool with Dell EMC Isilon SmartPools, you can change the IP address allocation for clients that connect to the cluster.

Here is a series of example commands that illustrate how you can associate networking resources with access zones:

```
isiloncluster1-1# isi networks list pools -v ①
subnet0:pool0
    In Subnet: subnet0
    Allocation: Static
    Ranges: 1
```

```
10.111.129.115-10.111.129.126 ②
```

```
Pool Membership: 4
```

```
1:10gige-1 (up)
```

```
2:10gige-1 (up)
```

```
3:10gige-1 (up)
```

```
4:10gige-1 (up)
```

```
Aggregation Mode: Link Aggregation Control Protocol (LACP)
```

```
Access Zone: z5 ③
```

```
SmartConnect:
```

```
Suspended Nodes : None
```

```
Auto Unsuspend ... 0
```

```
Zone : subnet0-pool0.
```

```
isiloncluster1.lab.example.com ④
```

```
Time to Live : 0
```

```
Service Subnet : subnet0
```

```
Connection Policy: Round Robin
```

```
Failover Policy : Round Robin
```

```
Rebalance Policy : Automatic Failback
```

- ① The command to show a verbose listing of the network pools.
- ② The range of IP addresses assigned to this network pool.
- ③ The access zone to which this pool of network resources is assigned.
- ④ The name of the SmartConnect zone that is associated with this network pool.

To create a new access zone and associate a pool of IP address with it, you can run the following series of commands:

```
isiloncluster1-1# mkdir -p /ifs/isiloncluster1/zonel  
isiloncluster1-1# isi zone zones create  
--name zonel --path /ifs/isiloncluster1/zonel ①  
isiloncluster1-1# isi networks create pool  
--name subnet0:pool1 --ranges=10.111.129.127-10.111.129.138  
--ifaces 1-4:10gige-1 --access-zone zonel ②  
--zone subnet0-pool1.isiloncluster1.lab.example.com ③  
--sc-subnet subnet0 --dynamic  
Creating pool 'subnet0:pool1': OK
```

Saving: OK

- ❶ The command to create a new access zone. In this example, the access zone is named *zone1*.
- ❷ This command creates a new network pool named *subnet0:pool1* and associates it with the access zone named *zone1*.
- ❸ This argument associates the network pool with a SmartConnect zone.

Then you can use the `Isilon racks` command to assign a new IP address pool to DataNode connections:

```
isiloncluster1-1# isi hdfs racks modify /rack0
--add-ip-pools subnet0:pool1

isiloncluster1-1# isi hdfs racks list

Name      Client IP Ranges      IP Pools
-----
/rack0 0.0.0.0-255.255.255.255 subnet0:pool0
                                subnet0:pool1
-----

Total: 1
```

You can use the same approach to distribute pools of IP addresses and the network interfaces of Isilon nodes across access zones. Here, for example, is a listing that illustrates the distribution of sets of IP addresses across four access zones, including the default system access zone:

```
sv-hdfs-lisi networks list pools -v
subnet0:pool0 - Default ext-1 pool ❶

    In Subnet: subnet0

    Allocation: Static

    Ranges: 1

        10.7.177.208-10.7.177.210

    Pool Membership: 3

        1:ext-1 (up)
        2:ext-1 (up)
        3:ext-1 (up)

    Aggregation Mode: Link Aggregation Control Protocol (LACP)

    Access Zone: System (1) ❷

    SmartConnect:

        Suspended Nodes : None

        Auto Unsuspend ... 0

        Zone : N/A
```

Time to Live : 0
Service Subnet : N/A
Connection Policy: Round Robin
Failover Policy : Round Robin
Rebalance Policy : Automatic Failback

subnet0:pool-z1 ③

In Subnet: subnet0

Allocation: Static

Ranges: 1

10.7.177.211-10.7.177.213

Pool Membership: 3

1:ext-1 (up)

2:ext-1 (up)

3:ext-1 (up)

Aggregation Mode: Link Aggregation Control Protocol (LACP)

Access Zone: z1 (2) ④

SmartConnect:

Suspended Nodes : None

Auto Unsuspend ... 0

Zone : N/A

Time to Live : 0

Service Subnet : N/A

Connection Policy: Round Robin

Failover Policy : Round Robin

Rebalance Policy : Automatic Failback

subnet0:pool-z2 ⑤

In Subnet: subnet0

Allocation: Static

Ranges: 1

10.7.177.214-10.7.177.216

Pool Membership: 3

1:ext-1 (up)

2:ext-1 (up)

3:ext-1 (up)

Aggregation Mode: Link Aggregation Control Protocol (LACP)

Access Zone: z2 (3) ⑥

SmartConnect:

Suspended Nodes : None

Auto Unsuspend ... 0

Zone : N/A

Time to Live : 0

Service Subnet : N/A

Connection Policy: Round Robin

Failover Policy : Round Robin

Rebalance Policy : Automatic Failback

① This pool of network resources is named subnet0:pool0. It is the default pool.

② The pool is associated with the default access zone.

③ This pool of network resources is named subnet0:pool-z1.

④ The pool is associated with the access zone named z1.

⑤ This pool of network resources is named subnet0:pool-z2.

⑥ This pool is associated with the access zone named z2.

Securing Access Zones for Analytics Multitenancy

As credit card companies, medical researchers, and financial institutions analyze data with Hadoop to detect fraud, improve health care, and create innovative products, the stored data poses a security and compliance problem: The Hadoop File System lacks the enterprise security features that compliance regulations require. As a result, big data analysts might be imperiling the integrity, confidentiality, and availability of their analytics data.

By securing data and tenants with the following capabilities, OneFS helps solve security and compliance problems for enterprise data hubs:

- Role-based access control for system administration
- Identity management
- Authentication
- Fine-grained access control to the file system
- Cross-protocol permissions and ACL policies
- User and ID mapping to associate each user with one ID
- WORM
- Auditing of SMB events

- Auditing of RBAC administrative changes

With role-based access control (RBAC) for administration, OneFS RBAC lets you manage administrative access by role. You can create separate administrator roles for security, auditing, storage, and backup. Then you can further tailor the RBAC role by assigning privileges to execute administrative commands.

After a security administrator with the correct RBAC privileges sets up an access zone for multitenancy, OneFS does three principal things in the access zone to provide secure multitenancy:

- Connects to directory services, such as Active Directory, NIS, and LDAP, which are also known as identity management systems. A directory service provides a security database of user and group accounts along with their passwords and other account information.
- Authenticates users and groups. Authentication verifies a user's identity and triggers the creation of an access token that contains information about a user's identity.
- Controls access to directories and files at the level of the file system. OneFS compares the information in an access token with the permissions associated with a directory or a file to allow or deny access at a granular level.

For more information about how OneFS manages identities, authenticates users, and controls access to resources, see [OneFS Multiprotocol Security Untangled](#).

Multitenancy for Directory Services

Multiprotocol data access demands multitenancy for services that manage user accounts, authenticate users, and control access to data. OneFS works with the following directory services to authenticate users and control access to files in an access zone:

- Microsoft Active Directory (AD), a popular directory service that uses domain controllers to authenticate users and authorize access to resources; to work with UNIX and Linux systems, Active Directory includes optional support for LDAP attributes with an implementation of RFC 2307.
- Lightweight Directory Access Protocol (LDAP), an application protocol for querying and modifying directory services.
- Network Information Service (NIS), a client/server directory service for distributing system information such as user names.
- Local users and local groups: unlike Active Directory, LDAP, and NIS, which are external systems, the local provider is an internal component of OneFS.
- File provider for working with accounts in `/etc/spwd.db` and `/etc/group` files; the file provider, which is also a OneFS component, lets you copy UNIX user and group account information from other systems.

Access Control

In an access zone, OneFS authorizes users and groups and controls access across different protocols by using POSIX mode bits, NTFS ACLs, or an optimal merging of them. The result is a consistent, predictable permissions model across data access protocols, including HDFS, which preserves the intended security settings for files, directories, and objects in the file system.

OneFS automatically converts the ACLs on Windows files to Linux permissions for Hadoop. Without the automated conversion of OneFS, you would have to use Hadoop tools to copy the files and then explicitly set the permissions for each directory or file—a tedious, error-prone task that is often skipped, leaving sensitive data available to all Hadoop users.

OneFS solves this problem by providing a single global namespace for all the data access protocols, including SMB, NFS, and HDFS. The ACLs that are defined through SMB are automatically converted to the NFS/HDFS permissions model. The result: OneFS secures analytics data and controls access to it as intended. For more information, see [Dell EMC Isilon Multiprotocol Data Access with a Unified Security Model](#).

The OneFS security system includes ACL policies that are, by default, set to address compliance regulations such as HIPAA, FISMA, and PCI DSS. The ACL policies include preserving ACEs that explicitly deny access to specific users and groups. The policies also let you tune the cluster to meet your access control objectives.

To map identities from multiple external directory services to a single, unique user ID in each access zone, OneFS includes a user mapping service that manages the identities and tokens of users and groups by access zone. For more information, see [Identities, Access Tokens, and the OneFS User Mapping Service](#).

WORM for Data at Rest

SmartLock is a OneFS feature that, after it is activated with a license, can lock down directories with read-many, write-once storage, commonly known as WORM. You can selectively apply WORM to directories in an access zone. The WORM settings are honored by all the data access protocols that OneFS supports, including HDFS.

The changes made to WORM directories by RBAC accounts can then be tracked to help fulfill compliance regulations for monitoring and auditing. For more information on SmartLock and WORM, see [Automated Data Retention with Dell EMC Isilon SmartLock](#).

Auditing

The rapidly growing data sets that populate a data lake often contain sensitive information like intellectual property, confidential customer data, and company records. Auditing helps detect fraud, inappropriate entitlements, unauthorized access attempts, and other anomalies, reducing the risk of exposing sensitive information. OneFS provides several auditing mechanisms to ensure the availability, integrity, and confidentiality of data:

- Support for SNMP versions 1, 2c, and 3 to remotely monitor hardware components, CPU usage, switches, and network interfaces for integrity.
- A virtual appliance, called InsightIQ, to monitor and analyze the performance of an Isilon cluster to forecast capacity and maintain availability.
- A RESTful application programming interface to automate monitoring and retrieve statistics.
- Auditing of system configuration events to track changes by administrators.
- SMB, NFS, and HDFS protocol monitoring to track user access and record file events.

The SMB event monitoring and auditing integrates with Varonis® DatAdvantage® with other SIEM applications. The events are logged on the node that an SMB client connects to and then stored in a file in `/ifs/.ifsvar/audit/logs`.

For more information on how OneFS securely stores data for analytics to meet such compliance regulations as PCI DSS, FISMA, and HIPAA, see [Compliance and Security for Hadoop Scale-Out Data Lakes](#).

Use Cases

The following use cases show how OneFS can serve as a multitenant enterprise data hub for big data analytics.

Moving from Silos to Efficiency and Accessibility

A university with disparate silos of stored data had two interrelated problems:

1. The small IT staff found it difficult to manage the silos and to protect the data in them consistently. Performance suffered.
2. Statistical researchers could not easily and consistently access the data that they wanted to analyze. The data was spread out. Trying to consolidate the data for analysis redirected the researchers' focus from analyzing the data with Hadoop to dealing with technology problems.

By putting an Isilon cluster in place and using it as a central storage hub, both problems were solved. The IT department gained a highly reliable and flexible storage infrastructure that protected data from loss and eliminated the need to migrate data between storage silos. Performance improved.

For the statisticians, the solution immediately delivered the accessibility to the data that they required for their research.

Sharing Data Among Tenants

By default, the data in one access zone cannot be accessed by users in another access zone. In certain cases, however, you may need to make the same data set available to more than one Hadoop compute cluster. Fully qualified HDFS paths can render a data set available across two or more access zones.

With fully qualified HDFS paths, the datasets do not cross access zones. Instead, the Hadoop jobs can access the data sets from a common shared HDFS namespace. For instance, you would start with two independent Hadoop compute clusters, each with its own access zone on an Isilon cluster. Then you can add a third access zone on the Isilon cluster with its own IP addresses and HDFS root. The zone can contain a dataset that is shared with the other Hadoop compute clusters.

Conclusion

A Dell EMC Isilon scale-out cluster lays the architectural foundation to create a data lake—a storage strategy to capture data from disparate workflows and make it available to analytics applications. When business units pool their data for storage and analysis, however, the data lake must be able to segregate the business units and their data sets. An Isilon cluster combines storage efficiency with data security to isolate data sets with a multitenant architecture that works seamlessly with Hadoop and other analytics applications.