



White Paper

FileOrbis - Dell EMC PowerScale, Dell EMC ECS Integration

Executive Summary

FileOrbis is a Secure Content Management and Enterprise File Sharing with Integrated Security solution.

Dell EMC PowerScale is the industry's leading scale-out NAS platform. Dell EMC PowerScale provides a scalable, high-performance, modular storage architecture that enables you to innovate with your data.

Dell EMC ECS, the leading object-storage platform from Dell EMC, has been engineered to support both traditional and next-generation workloads alike. Deployable in a software-defined model or as a turnkey appliance, ECS boasts unmatched scalability, manageability, resilience, and economics to meet the demands of modern business.

FileOrbis' data classification capabilities based on data type, content, source and many others. When integrated with intelligent tiering (SmartPools), Write Once Read Many (WORM / SmartLock) and storage based snapshots (SnapshotIQ) capabilities available in Dell EMC PowerScale, and S3 versioning available in Dell EMC ECS object platforms provide unmatched data management, collaboration, and governance solution for enterprise users in file granularity regardless of their physical location and device that they are on. Enterprise IT having full control of the corporate's most valuable asset in the form of data, combined solution provides greater flexibility in accessing, managing data along with tight security driven by corporate security policies.

Joint Solution Components

- FileOrbis
- Dell EMC PowerScale
- Dell EMC ECS

Joint Solution Benefits

- Content-aware Tiering (SmartPools / PowerScale)
- Content-aware Write Once Read Many (SmartLock / PowerScale)
- User-driven File Restores Based on Snapshots (SnapshotIQ / PowerScale)
- User-Initiated File Restores Based on S3 Versioning (S3 Versioning / ECS)
- Amazon S3 Protocol Access to File Classification Tags (PowerScale)
- Rich, "File System Like" ACL Support For ECS and GeoDrive
- Unified Storage Maps with WebDAV (PowerScale, ECS)

Contents

- White Paper 1
- FileOrbis - Dell EMC PowerScale, Dell EMC ECS Integration 1
- Executive Summary..... 1
- 1. Challenge 3
 - 1.1. About Dell EMC PowerScale..... 3
 - 1.2. About Dell EMC ECS 3
 - 1.3. About FileOrbis 4
- 2. Integrated Solution Overview..... 5
- 3. Dell EMC PowerScale Integration Capabilities 5
 - 3.1. Content Based Tiering (SmartPools) 5
 - Example 1: Fast Tier Rule..... 6
 - Example 2: Sensitive Data File Pool Policy..... 10
 - Example 3: HR Archive File Pool Policy..... 14
 - 3.2. Content Aware SmartLock 16
 - 3.3. User-Initiated File Restores Based on Snapshots 22
 - 3.4. Amazon S3 Protocol Access To File Classification Tags 24
 - 3.5. Content Based Batch Operations for Files and Metadata 24
- 4. Dell EMC ECS Integration Capabilities 25
 - 4.1. User-Initiated File Restores Based on Versions..... 27
 - 4.2. GeoDrive..... 28
- 5. Other Integration Points..... 29
 - 5.1. WebDAV Maps 29
 - 5.2. Full Text Search and File Disposal..... 31
- Appendix: Related Resources 33

1. Challenge

Due to the increase in complexity in the IT landscape and the use of different multi-vendor services and applications, the challenge of effectively managing these services and applications has become quite demanding and complicated. IT professionals are expected to follow best practices, solve every incident, keep themselves up to date with new technical improvements, and excel at managing new tools like never before.

Like any aspect of the business, controlling costs plays a significant role, especially in IT operations. The cost of having a just-enough-competent IT support team is skyrocketing, let alone hiring highly specialized individuals. This situation presents a massive challenge for IT operations managers because they must ensure they get the right talent to help provide the needed support as expected while managing the budget wisely. Periodic follow-ups and tracking team efforts and collaboration is a necessity but time-consuming. IT directors are in a constant search for a way to centralize this process so that they can lead efficiently while fulfilling the needs of enterprise IT as expected.

1.1. About Dell EMC PowerScale

The power of the Dell EMC PowerScale is the OneFS operating system powering the industry's leading scale-out NAS platform. Apart from unlocking the potential within your unstructured data, OneFS enables you to store, manage, protect, secure and analyze your data while running a wide variety of applications.

With built-in interoperability, OneFS-based solutions are simple to manage at any scale, and capacity can be provisioned in minutes to your cluster. A single volume, single filesystem, and single namespace enable you to consolidate your data and eliminate storage silos. Regardless of the number of nodes in your cluster, a OneFS powered solution allows you to store and manage many petabytes of data with a single admin.

With support for protocols like NFS, SMB, S3, and HDFS, you can simultaneously run applications that require file and object protocols on the same dataset, which helps you maximize the value of your data in this Data First world.

1.2. About Dell EMC ECS

Organizations require options for consuming public cloud services with the reliability and control of a private-cloud infrastructure. Dell EMC ECS is a software-defined, cloud-scale, object storage platform that delivers S3, Atmos, CAS, Swift, NFSv3, and HDFS storage services on a single, modern platform.

Simple RESTful API access for storage services is being embraced by developers. Use of HTTP semantics like GET and PUT simplifies the application logic required when compared with traditional, but familiar, path-based file operations. In addition, Dell

EMC ECS's underlying storage system is strongly consistent, which means it can guarantee an authoritative response. Applications that are required to guarantee authoritative delivery of data are able to do so without complex code logic by using Dell EMC ECS.

1.3. About FileOrbis

FileOrbis is an on-premise/on-cloud content management system equipped with unique operation and control features allowing you to:

- **Enforce security** scans and **controls** on your files.
- **Conduct content** and sensitive **data analysis** on your files.
- **Share** your files with internal and external users.
- **Manage** permissions and **access** for your files.
- **Access** your files from **everywhere**.

FileOrbis is a framework for merging file servers, user profile folders, user-specific areas, and network disks. FileOrbis provides an access channel for all your file systems as well as an integrated management system.

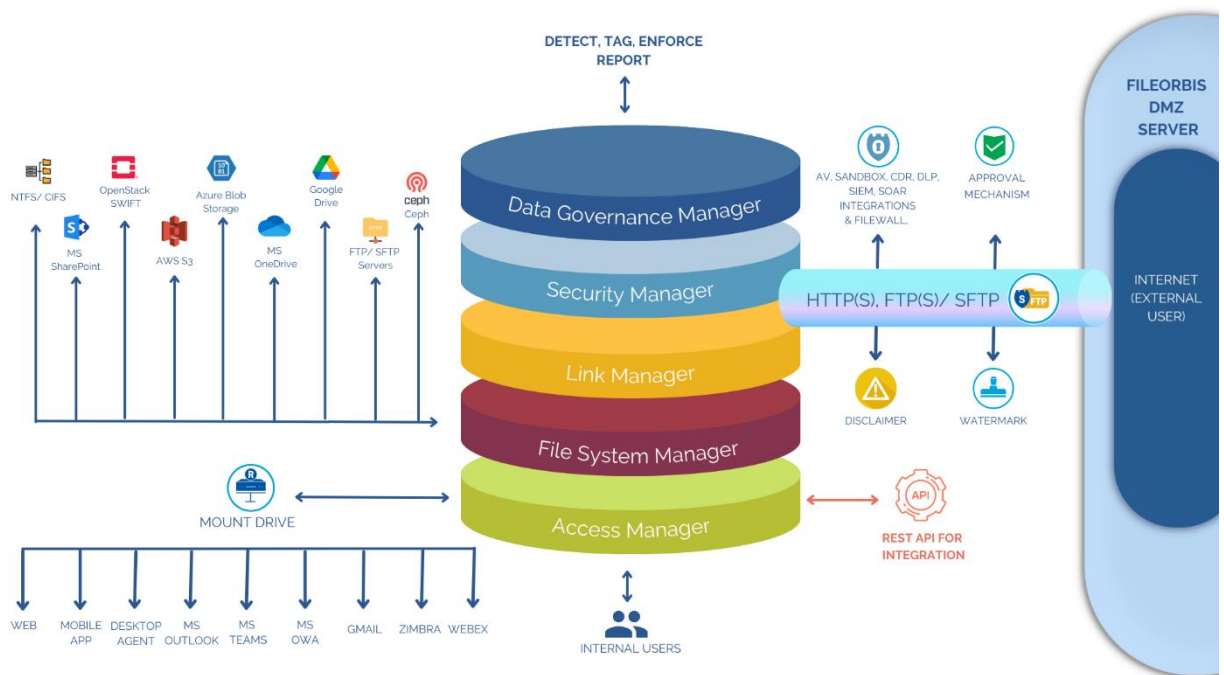


Figure 1.3.a – FileOrbis Topology

FileOrbis does not approach content management from a singular viewpoint, such as access, authorization management, file sharing, logging, etc., but rather assists businesses with all content management-related duties. By integrating with security solutions such as antivirus, sandbox, DLP, and storage solutions of Dell EMC PowerScale and Dell EMC ECS. FileOrbis enables you to run all file environments more effectively and in conjunction with other complementing solutions.

2. Integrated Solution Overview

"Integrated" and "ease of management" are two pillars that led us to build a solution wherein two components, one in the front-end services stack and another in the back-end storage stack, work together and leverage each other's capabilities. The combined solution provides granular management at the file level for unstructured data sets, and automation by leveraging each platform's capabilities. Both, Dell EMC PowerScale and Dell EMC ECS platforms are integrated in the back-end storage services stack for different sets of capabilities.

3. Dell EMC PowerScale Integration Capabilities

FileOrbis is integrated with Dell EMC PowerScale through OneFS API for data management and access.

3.1. Content Based Tiering (SmartPools)

In FileOrbis management GUI, you can create rules that mark and associate files with custom file attributes based on data type, properties, and content. Once files matching certain rules have been processed by the FileOrbis rule engine, attributes designated can become part of files in the backing OneFS filesystem in the form of OneFS native extended file attributes.

OneFS File Pool Policies using these extended file attributes in conjunction with user and group attributes, time, and data location properties as criteria can provide sophisticated and granular control at the file level over file data sets. With this integration, setting the right storage tier and movement and archiving between different storage tiers for the right data can become possible at individual file granularity and with automation. It is also possible to set the right protection and performance profiles for individual files, whichever tier they reside in.

In the next section, we will try to explain and illustrate how the outcomes can be achieved through some real-life examples. There are many use cases that can be addressed with this integration. Some of these are listed below for reference.

- Compliance driven protection
 - IDs, Credit Card Numbers, IBANs, Gender, IPv4, Email Addresses, Phone Numbers, Blood Types, Custom (regex), etc.
 - Request Source (Right SLA for the right source)
- Source type-based SLAs
 - Example: Data coming from SAP is accelerated and protected at a higher level
- Source address-based SLAs
 - Subnet, IP range, individual IPs
 - Example: Data coming from archive servers are placed in the archive tier
- Data type-based SLAs

- Example: Media files are archived, office documents are accelerated and protected with a higher protection
- Source identity-based SLAs
 - Adjust the data SLA based on user and/or group identity
 - User defined classification
 - User defined data profile

Example 1: Fast Tier Rule

In this scenario, a C-level user needs fast access to their files. In order to meet the requirement, when a file is created or uploaded by a C-level user through FileOrbis, the respective files should be marked for placement on the flash-based tier in the Dell EMC PowerScale cluster.

Figure 3.1.a shows how easily a workflow rule on FileOrbis Web Management GUI can be created by providing a name and a type of operation. The "upload" operation is selected for this specific example. There are many types of operations available for other actions, such as "Create File", "Create Folder", "File Edit, Create Copy, Delete, Move, Rename, Download, Preview, Upload, Share, Share Update, Create Archive, Extract Archive, Recycle Bin Delete, Recycle Bin Move, Recycle Bin Restore, Link Create, Link Remove, Link Update, Link Create Folder, Link Download, Link Preview, Link Upload.

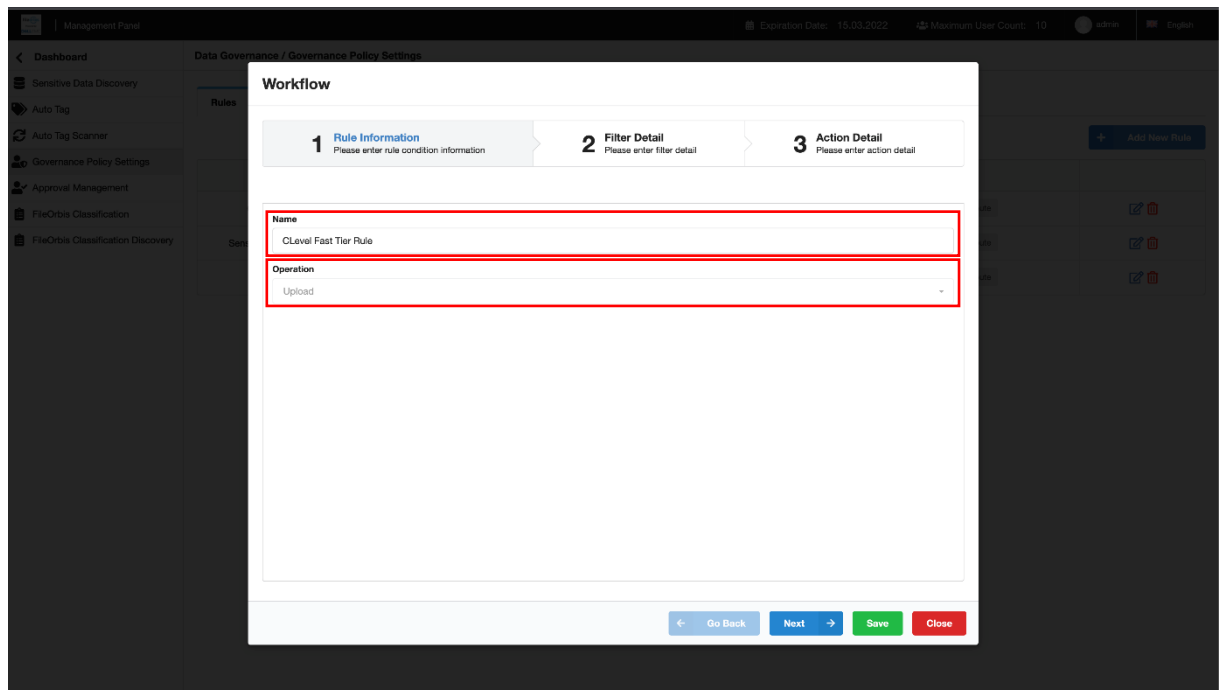


Figure 3.1.a – C-Level Fast Tier Rule Creation

Figure 3.1.b below shows how to apply a user/group and a source filter to the rule.

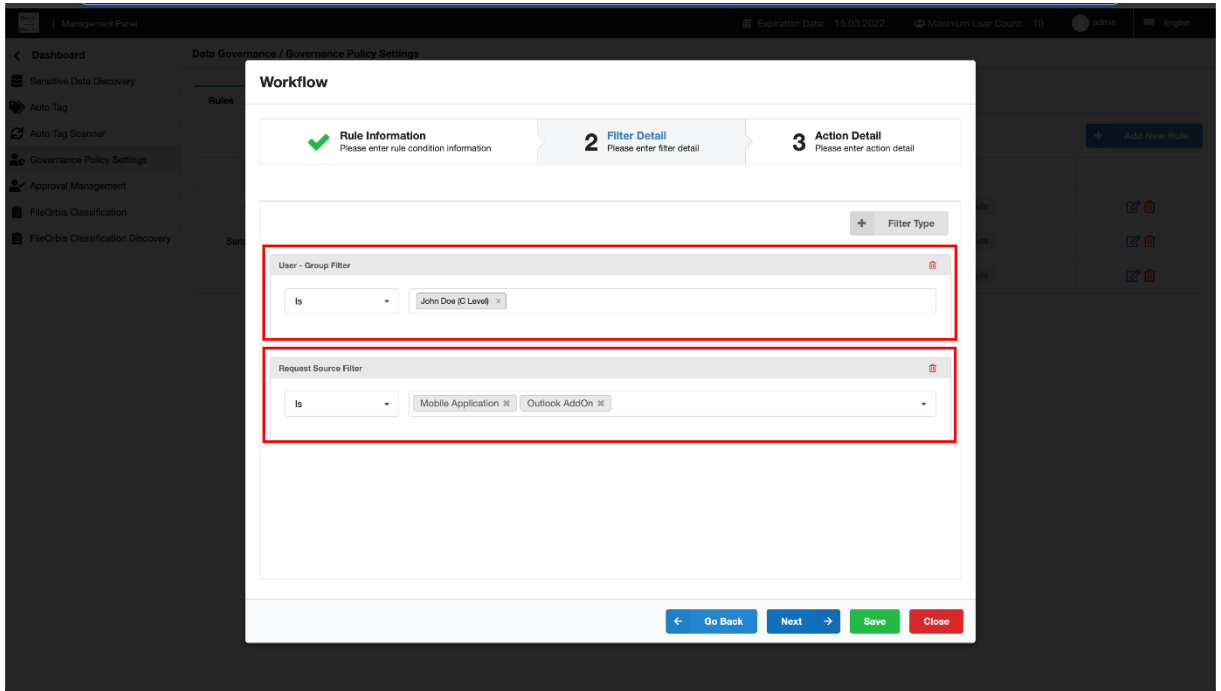


Figure 3.1.b – Providing Filter Details

Finally, Figure 3.1.c below shows adding OneFS extended attributes to the file with key=Clevel and val=Speedup values as an action when the backing file system is OneFS,

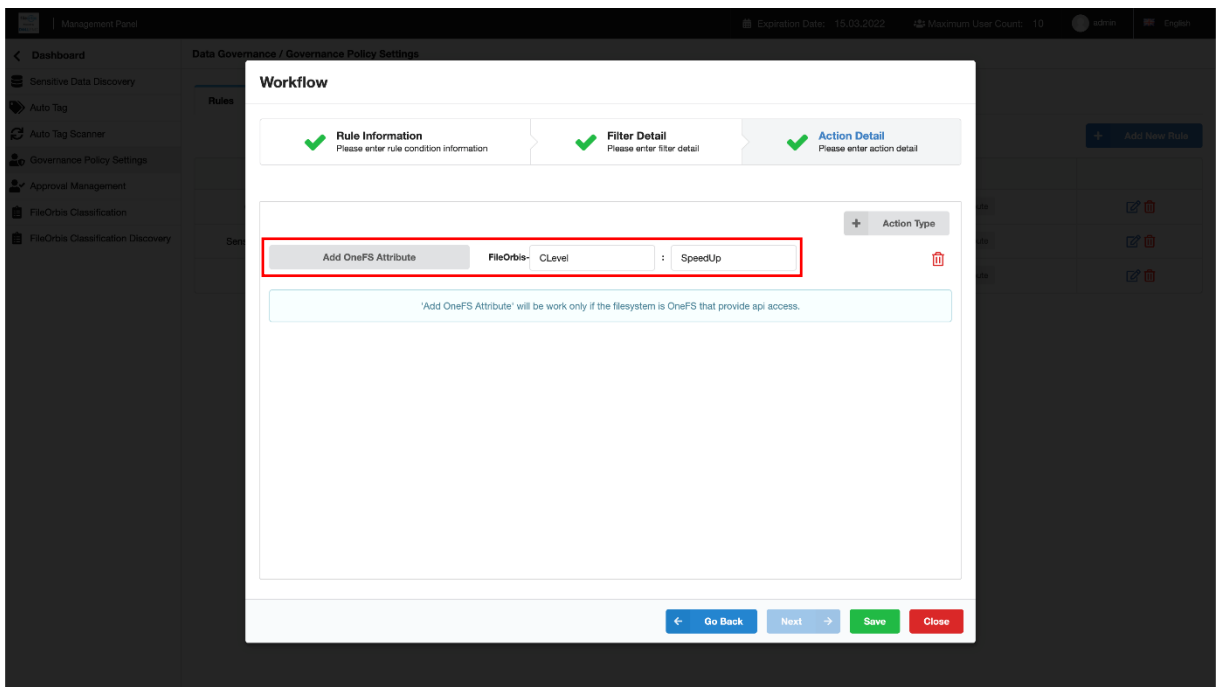


Figure 3.1.c – Creating Custom File Attributes in OneFS

By clicking the save button, the rule has been created in three simple steps.

How files will look like on FileOrbis user web interface for the respective C-level user can be observed in figure 3.1.d. It shows two files uploaded by a C-level user.

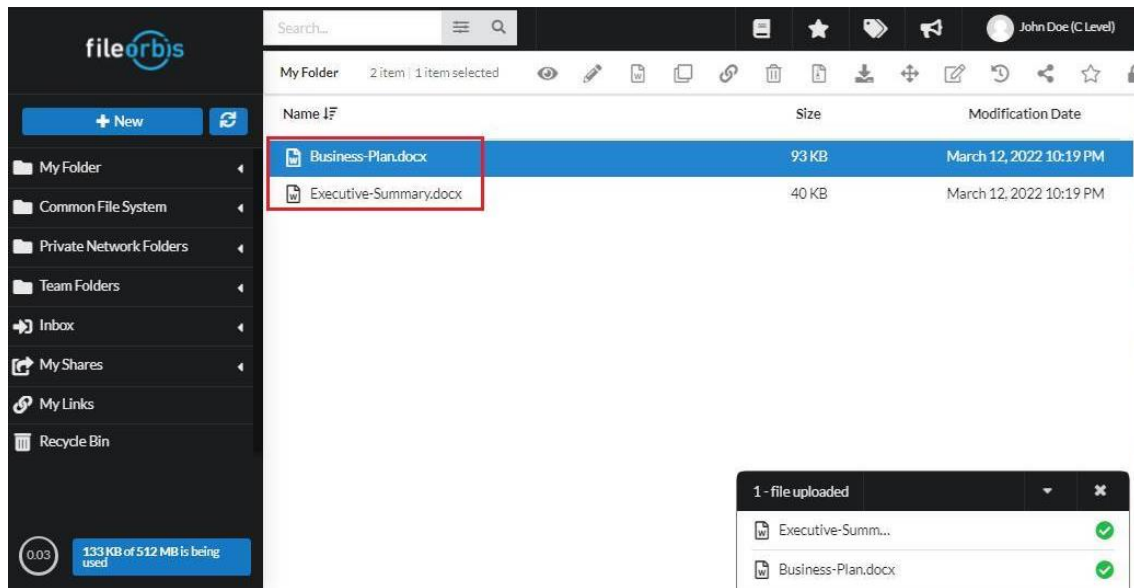


Figure 3.1.d – C-Level File Upload Operation

Figure 3.1.e below shows Dell EMC PowerScale management interface for file pool policies. Please notice that file pool policies have been created in relation to the rules created in FileOrbis by the storage administrator. This is a one-time activity that requires very minimal administrative work on the storage side. The beauty of this integration is that the storage administrator is in control of how the content should be treated on storage based on its value and SLA requirements, versus FileOrbis administrator, who might have very little insight into the dynamics of the underlying storage system while trying to do that. What was previously not possible for the storage administrator to control the data based on its content, value, and properties is now possible without storage administrator gaining access to or running scans on the actual content.

Logged in as admin | [Review recent events](#) | [Log out](#) | [Help](#) ?

Cluster Name: FOCLS(OneFS Version: 9.3.0.0) Node 1

Dashboard ▾ Cluster management ▾ File System ▾ Data protection ▾ Access ▾ Protocols ▾

Storage pools

Summary **File pool policies** SmartPools CloudPools SmartPools settings CloudPools settings

✓ Policy updated

File pool policies [Create a file pool policy](#)

Select a bulk action ▾

<input type="checkbox"/>	Order	Policy name	CloudPools state	Description	Actions
<input type="checkbox"/>	↓	CLevel FPP	No access		View/Edit Delete
<input type="checkbox"/>	↑ ↓	Sensitive Data FPP	No access		View/Edit Delete
<input type="checkbox"/>	↑	HR Archive FPP	No access		View/Edit Delete
<input type="checkbox"/>		Default policy		This policy applies to all files not selected by higher-priority policies.	View/Edit

Policy templates

Template name	Description	Actions
Archive	Move older files to older storage.	View/Use template
VMware Files	Set VMware files for random-access.	View/Use template
ExtraProtect	Protect a subset of files at a higher requested protection.	View/Use template
PoolByPath	Assign files to the performance pool based solely on their path.	View/Use template

Figure 3.1.e – OneFS File Pool Policies

Figures 3.1.f and 3.1.g below illustrate when the view/edit button for the first policy named "CLevel FPP" has been clicked.

Edit file pool policy details Help ?

Description

CloudPools state **No access**

CloudPools state details Policy has no CloudPools actions

*Policy name

Description

Select files to manage

Specify criteria for determining which files will be managed by this policy:

*File matching criteria

IF condition

Figure 3.1.f

The criteria for matching the file pool policy can be seen in figure 3.1.g, which is shown below.

Figure 3.1.g

Figure 3.1.h shows extended attributes in OneFS over CLI.

```
FOCLS-1# ls
Business-Plan.docx      Executive-Summary.docx
FOCLS-1# lsextattr user Business-Plan.docx
Business-Plan.docx      FileOrbis-CLevel
FOCLS-1# lsextattr user Executive-Summary.docx
Executive-Summary.docx FileOrbis-CLevel
FOCLS-1# getextattr user FileOrbis-CLevel Executive-Summary.docx
Executive-Summary.docx SpeedUp
FOCLS-1# getextattr user FileOrbis-CLevel Business-Plan.docx
Business-Plan.docx      SpeedUp
FOCLS-1#
```

Figure 3.1.h – CLI output

Example 2: Sensitive Data File Pool Policy

In this scenario, a use case would be, if a file contains equal to or more than 10 credit card numbers and has classification filters "top secret" or "secret" set by FileOrbis, which are indicators of a file being sensitive, the protection level is increased accordingly.

Figure 3.2.a shows that we start by creating a rule for FileOrbis management. We name the rule "Sensitive Data Protection Rule" and the operation type is set to upload.

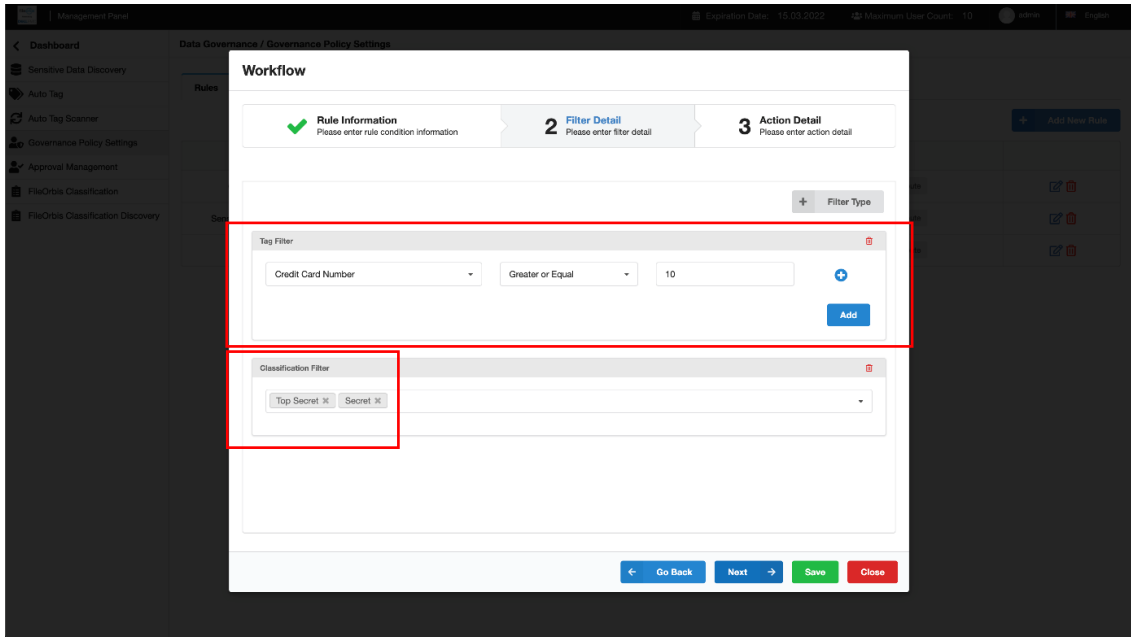


Figure 3.2.a – Providing Filter Details

In the next window, as shown by Figure 3.2.b below, FileOrbis will mark the files containing sensitive data with key=ProtectionLevel, val=High which will be transformed into extended file attributes in OneFS.

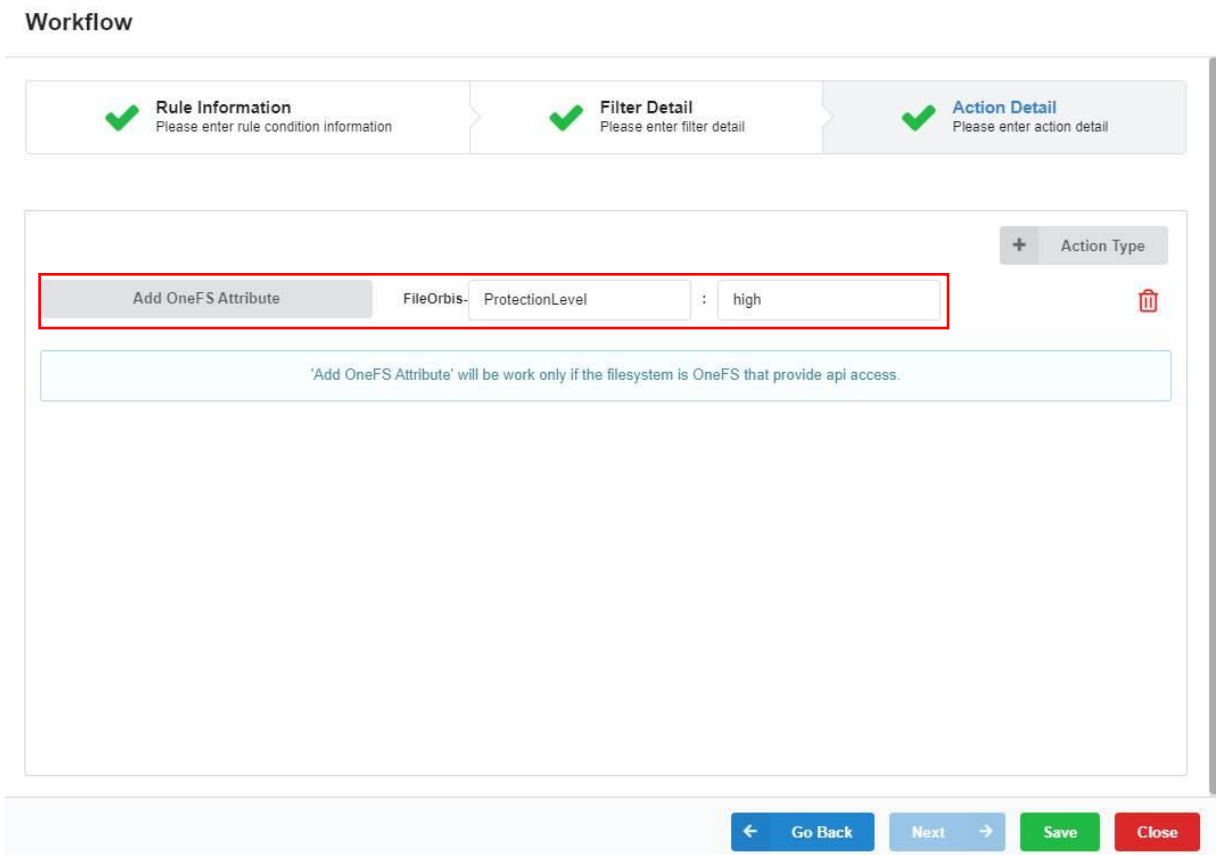


Figure 3.2.b - Creating Custom File Attributes

As shown in FileOrbis user web interface, in Figure 3.2.c below, the file containing sensitive information has been uploaded by a user.

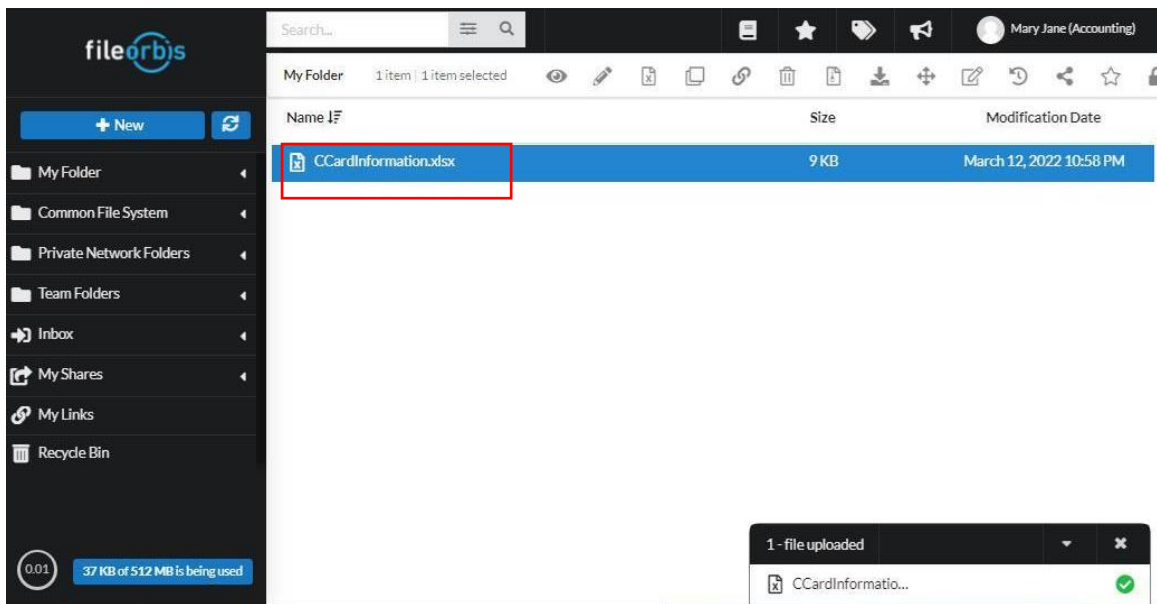


Figure 3.2.c – Sensitive file uploaded

Below, details from the respective File Pool Policy, namely "Sensitive Data FPP" in Dell EMC PowerScale management interface is shown. You will see that the protection of sensitive data will be changed to "[6x] Mirrored over 6 nodes" when an extended file attribute key-value pair is matched.

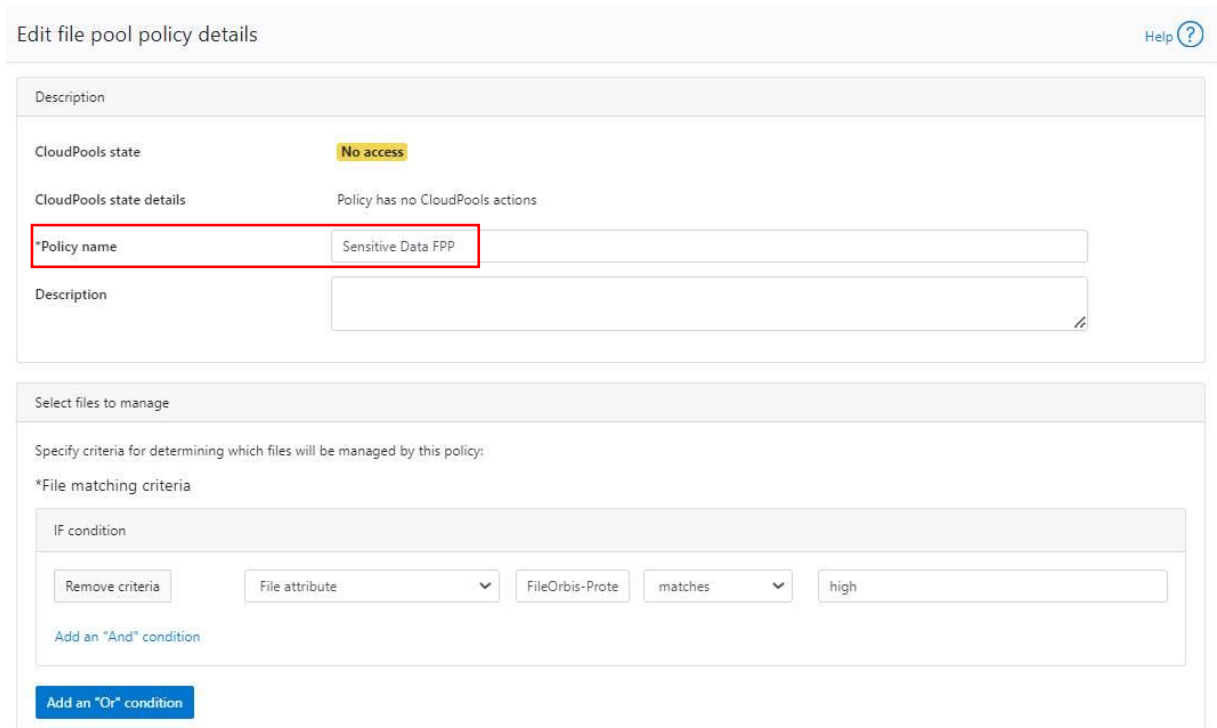


Figure 3.2.d

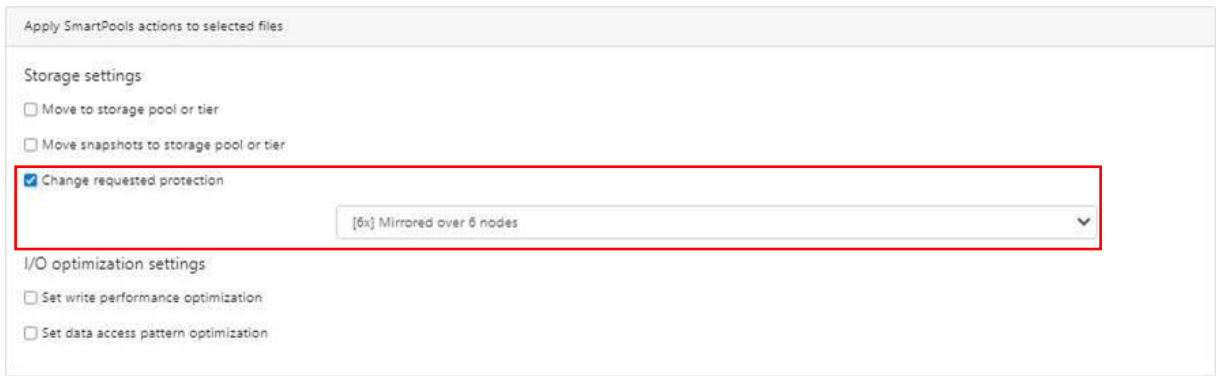


Figure 3.2.e

Figure 3.2.f also shows CLI output for extended file attributes set in OneFS and how the file is protected by 6 copies after a SmartPools job is run.

```
FOCLS-1# ls
CCardInformation.xlsx
FOCLS-1# lsextattr user CCardInformation.xlsx
CCardInformation.xlsx  FileOrbis-ProtectionLevel
FOCLS-1# getextattr user FileOrbis-ProtectionLevel CCardInformation.xlsx
CCardInformation.xlsx  High

FOCLS-1# isi get -DD CCardInformation.xlsx |more
POLICY  W  LEVEL  PERFORMANCE COAL  ENCODING  FILE  IADDRS
6x      6      6x      concurrency on  UTF-8    CCardInformation.xlsx <1,15,1699840:512>, <2,5,6498816:512>, <3,15,36984
> ct: 1647117950 rt: 0
*****
* IFS inode: [ 1,15,1699840:512, 2,5,6498816:512, 3,15,369846272:512, 4,2,2120192:512, 5,5,1940992:512, 6,15,2131456:512 ]
*****
*
* Inode Version:      8
* Dir Version:       2
* Inode Revision:    7
* Inode Mirror Count: 6
* Recovered Flag:    0
* Restripe State:    0
* Link Count:        1
* Size:              9606
* Mode:              0100700
* Flags:             0x130004e0
* SmartLinked:       False
* Physical Blocks:   12
* Phys. Data Blocks: 2
* Compression Ratio: 1:1 (100% of logical size)
* Protection Blocks: 10
* LIN:               1:0005:1b50
* Logical Size:      16384
* Shadow refs:       0
* Do not dedupe:     0
* In CST stats:      True
* Last Modified:     1647117950.590653000
* Last Inode Change: 1647117950.960231000
* Create Time:       1647117950.572482000
* Rename Time:       0
* Write Caching:     Enabled
* Parent Lin         1:0003:1922
* Parent Hash:       226391
* Snapshot IDs:      None
* Min Snapshot ID:   12
* Max Snapshot ID:   HEAD
* Last Paint ID:     11
* IFS Domain IDs:    {2.0100 (Snapshot) }
* Domain IDs:        None
* LIN needs repair:  False
* Manually Manage:
*   Access           False
*   Packing           False
*   Protection        False
* Protection Policy: 6x
* Target Protection: 6x
```

```

PROTECTION GROUPS

lbn 0: 6x
  4,16,127287296:8192#2
  (sparse)#14
  5,5,269451264:8192#1
  5,5,269672448:8192#1
  (sparse)#14
  1,5,130424832:8192#2
  (sparse)#14
  3,16,330940416:8192#1
  3,16,331415552:8192#1
  (sparse)#14
  6,15,212303872:8192#2
  (sparse)#14
  2,16,171089920:8192#1
  2,16,171900928:8192#1
  (sparse)#14

```

Figure 3.2.f – CLI Output

Example 3: HR Archive File Pool Policy

In this scenario, a file has been uploaded by an individual member of a group and contains personal information such as salary, gender, blood type, date, or email. FileOrbis detects the filters and tags these files as CV. Once the file has been marked, it will be moved to the archive tier upon the next successful SmartPools job run.

Figure 3.3.a shows the creation of the respective rule in FileOrbis. HR Archive Rule is created, and an operation type is set to upload.

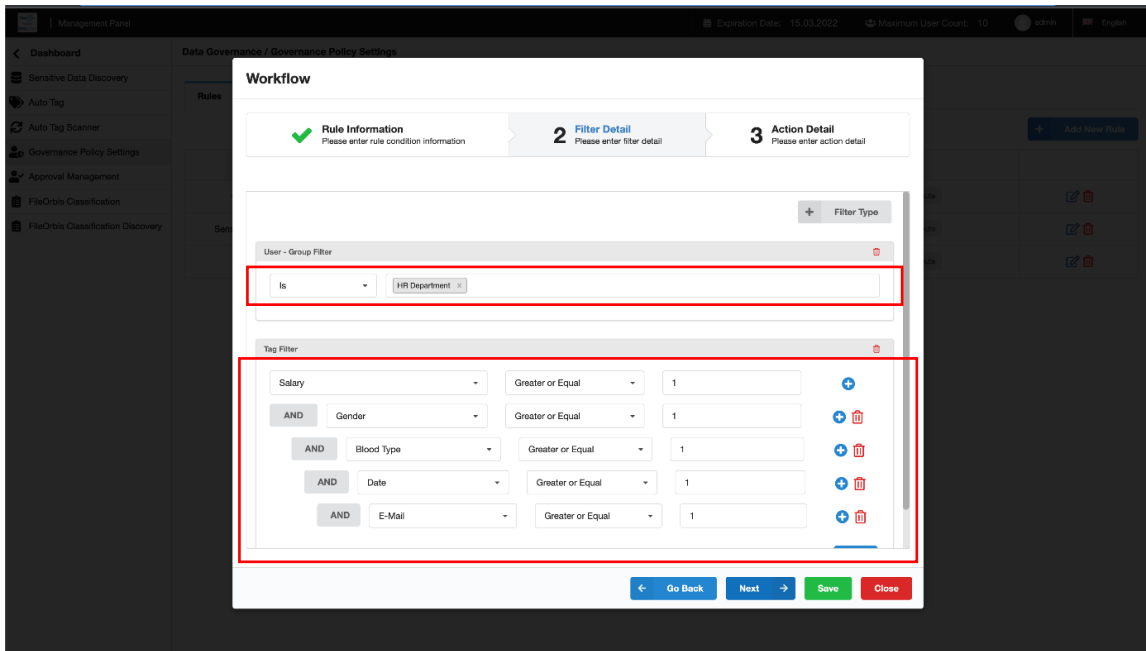


Figure 3.3.a - Providing Filter Details

The OneFS integration part takes place in the next step. A file containing personal information will be marked with a key=CV, and a val=True both of which will be transformed into extended file attributes in OneFS as shown in Figure 4.3.b

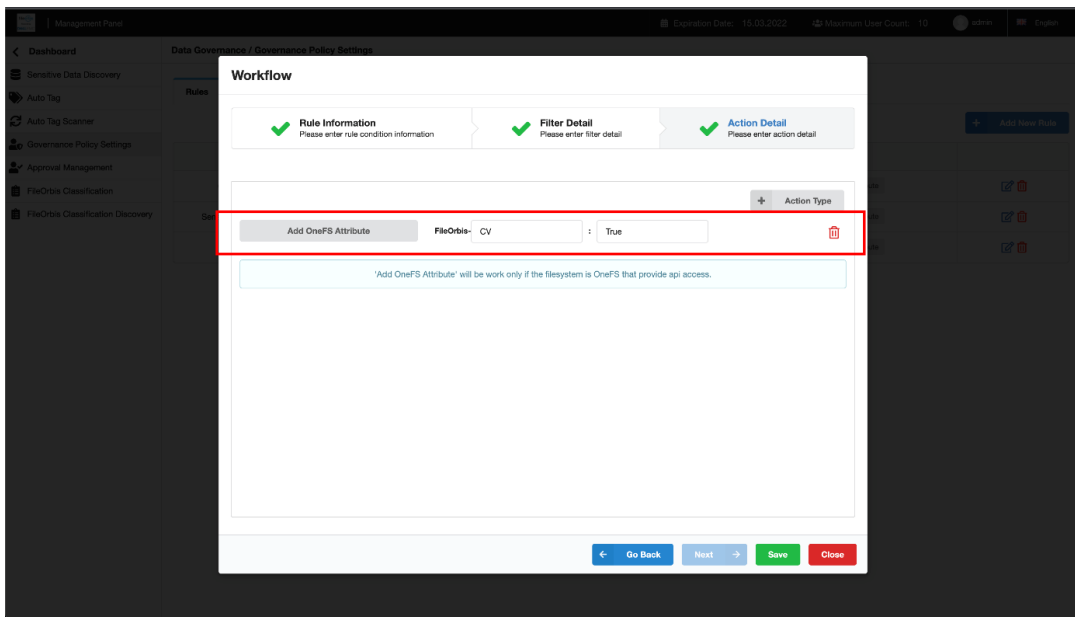


Figure 3.3.b - Creating Custom File Attributes

Figures 3.3.c and 3.3.d below show the respective "HR Archive FFP" file pool policy details, which will move the file to the archive tier when the file matching criteria have been met by the extended attributes on the file, during the next SmartPools job run.

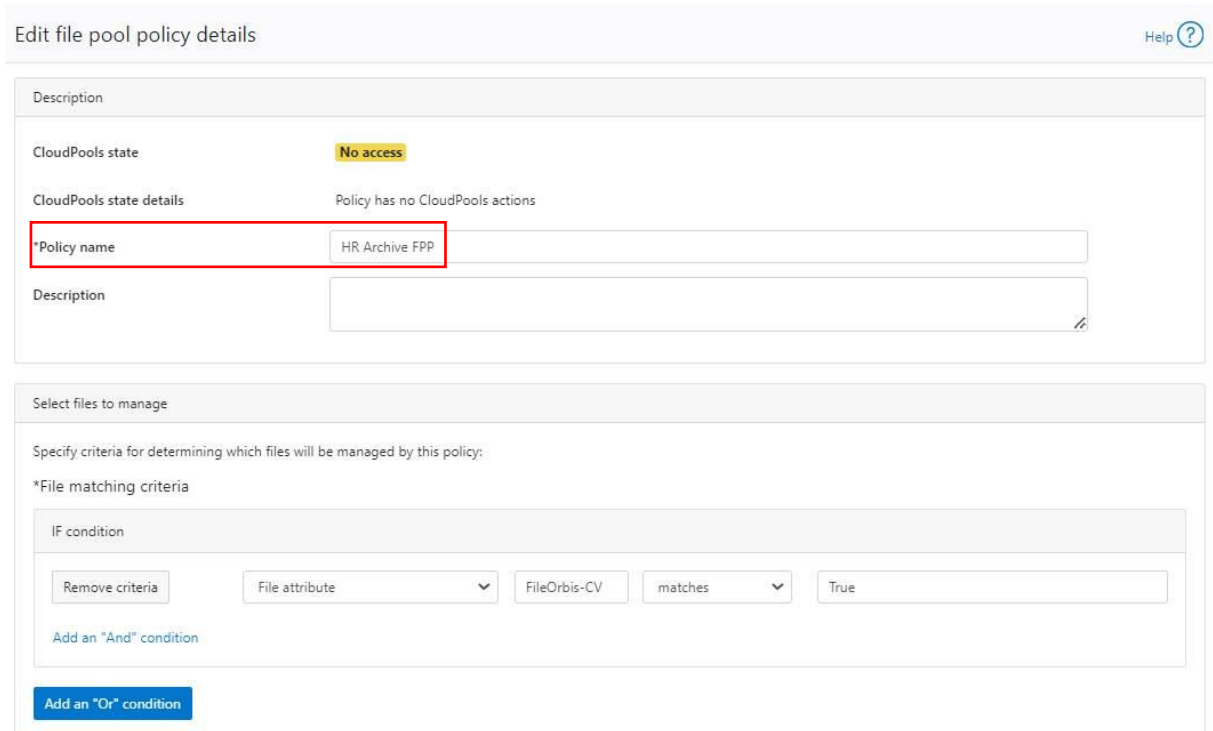


Figure 3.3.c

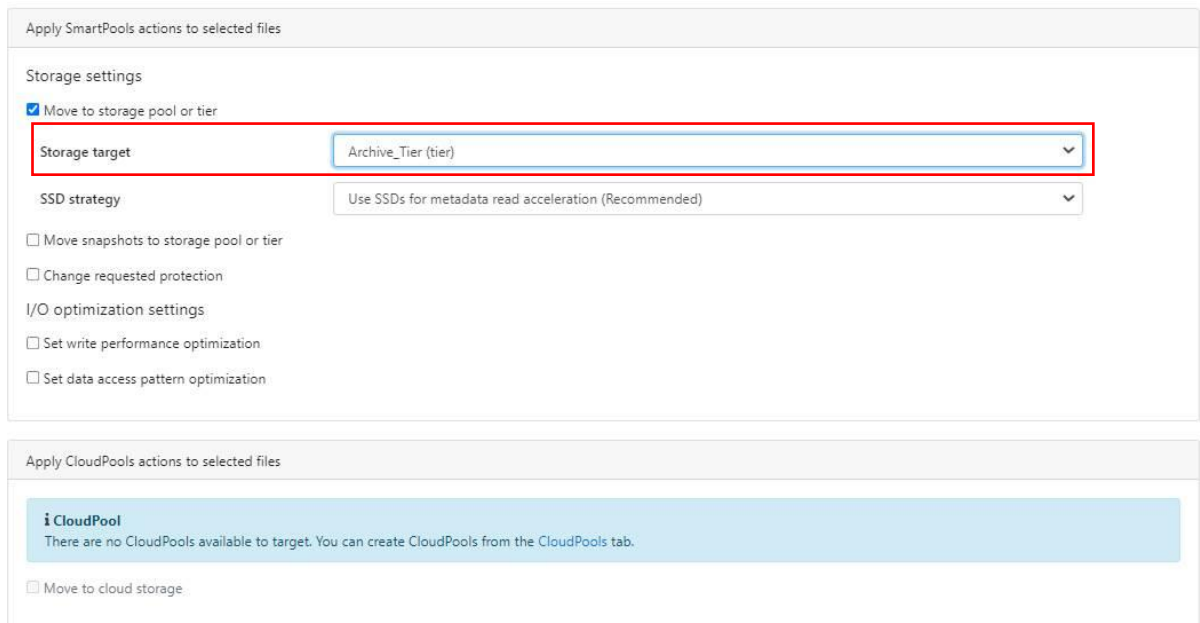


Figure 3.3.d

Figure 3.3.e below shows CLI output for extended file attributes in OneFS.

```
FOCLS-1# ls
ChristopherMorgan_CV.docx
FOCLS-1# lsextattr user ChristopherMorgan_CV.docx
ChristopherMorgan_CV.docx      FileOrbis-CV
FOCLS-1# getextattr user FileOrbis-CV ChristopherMorgan_CV.docx
ChristopherMorgan_CV.docx      True
```

Figure 3.3.e

3.2. Content Aware SmartLock

The SmartLock feature in Dell EMC PowerScale provides an automated data retention solution that is simple to implement and manage. It is also reliable, flexible enough to support multiple use cases without requiring investment in additional hardware and software components, and secure enough to meet the needs of today's IT's strict compliance requirements.

With this integration, the following two key aspects of Write Once Read Many (WORM) protection for file data sets are possible:

- Compliance proof, storage back-end sustained WORM capability,
- Intelligent, file-based WORM committing through file classifications.

Dell EMC PowerScale SmartLock with FileOrbis is achieved through the use of OneFS API. For this integration to take effect, first, a SmartLock domain needs to be created in OneFS.

Dashboard ▾ Cluster management ▾ File system ▾ Data protection ▾ Access ▾ Protocols ▾

SmartLock

WORM

Write once read many (WORM) Domains [+ Create domain](#)

Path	Default	Actions
/ifs/SmartLockDomain	--	View / Edit

Figure 3.2.a

A previously created SmartLock domain, which is also shared in OneFS, is added as a CIFS share with API integration in the FileOrbis management GUI.

File System Settings

File System

Type: Dell EMC OneFS ▾

Access

Name: SmartLock Domain

Path: (highlighted with a red box)

Use FileSystem ACL:

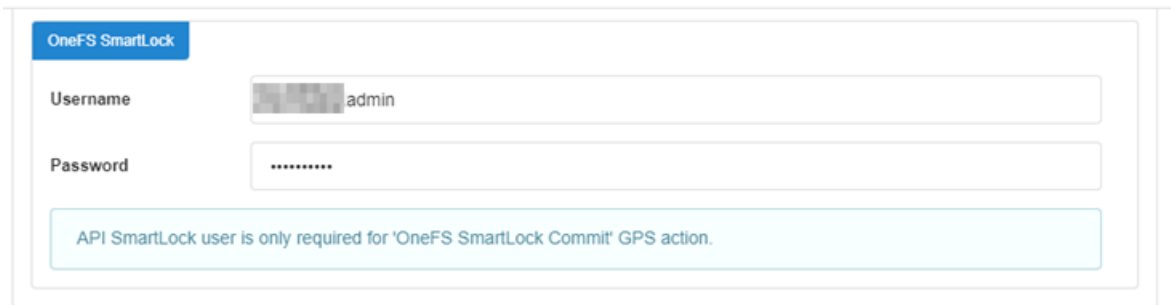
API Access

Url: (highlighted with a red box)

Figure 3.2.b

In OneFS SmartLock section, the username and password for a user who has sufficient permissions should be provided.

File System Settings



OneFS SmartLock

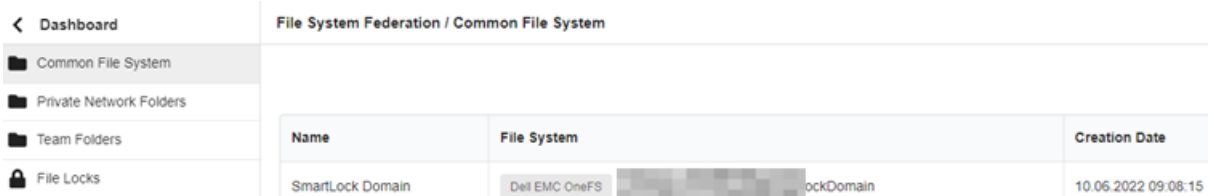
Username

Password

API SmartLock user is only required for 'OneFS SmartLock Commit' GPS action.

Figure 3.2.c

As a result, the specified SmartLock domain has been added as a Common File System to FileOrbis.



Name	File System	Creation Date
SmartLock Domain	Dell EMC OneFS [redacted] lockDomain	10.06.2022 09:08:15

Figure 3.2.d

In the Data Governance module, a custom tag should be created for files to be committed to the WORM state. The custom tag is named Annual Reports, and the required regex value is set to clear text Report 2021 for simplicity in this example. Sophisticated regular expressions can be used in this field as required.

fileorbis | Management Panel

Dashboard | Data Governance / Auto Tag

Sensitive Data Discovery

Auto Tag

Auto Tag Scanner

Governance Policy Settings

Approval Management

FileOrbis Classification

FileOrbis Classification Discovery

Predefined Auto Tags

Blood Type	<input checked="" type="checkbox"/>	
IPv6	<input checked="" type="checkbox"/>	
IBAN	<input checked="" type="checkbox"/>	
Gender	<input checked="" type="checkbox"/>	
GSM Number	<input checked="" type="checkbox"/>	
Address Comprehensive	<input checked="" type="checkbox"/>	
Timeout	<input checked="" type="checkbox"/>	
Encrypted	<input checked="" type="checkbox"/>	

Custom Auto Tag

Custom Tag Name	Regex Value
Annual Reports	REPORT 2021

Figure 3.2.e

To apply the tag to commit files to WORM state automatically a Data Governance Policy should be created.

Dashboard | Data Governance / Governance Policy Settings

Rules Incident Viewer

+ Add New Rule

Rule Name	Operation	Filters	Activity	
SmartLock - Annual Reports 2021	Upload	Tag Filter	OneFS SmartLock Commit	

Figure 3.2.f

Workflow

The screenshot shows a three-step workflow process. Step 1, 'Rule Information', is the active step and contains a text input field for 'Name' with the value 'SmartLock - Annual Reports 2021' and a dropdown menu for 'Operation' set to 'Upload'. Step 2, 'Filter Detail', and Step 3, 'Action Detail', are inactive. At the bottom right, there are four buttons: 'Go Back', 'Next', 'Save', and 'Close'.

Figure 3.2.g

When uploading a file, if the file contains “Report 2021” in the content, the policy will take effect.

Workflow

The screenshot shows the 'Filter Detail' step of the workflow. Step 1, 'Rule Information', is completed and marked with a green checkmark. Step 2, 'Filter Detail', is the active step and shows a 'Tag Filter' configuration. It includes a '+ Filter Type' button, a 'Tag Filter' header with a close icon, and a form with three fields: 'Annual Reports' (dropdown), 'Greater or Equal' (dropdown), and '1' (text input). An 'Add' button is located below the form. Step 3, 'Action Detail', is inactive. At the bottom right, there are four buttons: 'Go Back', 'Next', 'Save', and 'Close'.

Figure 3.2.h

When the policy is enforced, matching files will be committed to WORM with a 5-year retention period and cannot be deleted until that retention period has expired.

Workflow

The workflow configuration interface consists of three main sections, each with a green checkmark icon:

- Rule Information:** Please enter rule condition information
- Filter Detail:** Please enter filter detail
- Action Detail:** Please enter action detail

The Action Detail section is currently active and displays the following configuration:

- Action Type:** + Action Type
- OneFS SmartLock Commit:** Retention Time: 60 Months
- Note:** 'OneFS SmartLock Commit' will be work only if the filesystem is OneFS that provide api access.

Navigation buttons at the bottom: Go Back, Next, Save, Close.

Figure 3.2.i

The retention time can be observed when looking at file properties in the FileOrbis user GUI.

The screenshot shows the FileOrbis user GUI with a file list and a details pane. The file list shows two files:

Name	Size
IC-Annual-Marketing-Report-2021.docx	36 KB
IC-Simple-Year-End-Report-2021.docx	40 KB

The details pane for the selected file shows the following information:

- Name: IC-Simple-Year-End-Report-...
- Location:
- Creation Date: June 23, 2022 11:08 AM
- Modification Date: June 23, 2022 11:08 AM
- Size: 40 KB
- Type: Common File System
- Dell EMC OneFS**
- SmartLock Retention Date: June 23, 2027 11:08 AM**

Figure 3.2.j

3.3. User-Initiated File Restores Based on Snapshots

SnapshotIQ is designed to help you provide highly efficient, cost-effective, and low-recovery objective data protection. Once a baseline snapshot has been established, only changes to blocks that make up a file are reflected in updates to the current version of snapshots. This allows highly efficient snapshot storage utilization. Since snapshots are an integral part of the OneFS file system, there is no need to pre-allocate dedicated snapshot reserve space.

Users can list, preview, download, and restore the previous versions of files in OneFS snapshots through the FileOrbis GUI. To enable this functionality, first, snapshots should be enabled in OneFS. As part of the configuration, when adding CIFS shares to FileOrbis configuration, snapshot configuration should be provided.




Figure 3.3.a

The gear icon in FileOrbis GUI, used to configure snapshot settings. In the Versions section the type should be selected “Snapshot” and “.snapshot” should be populated in the Path field.

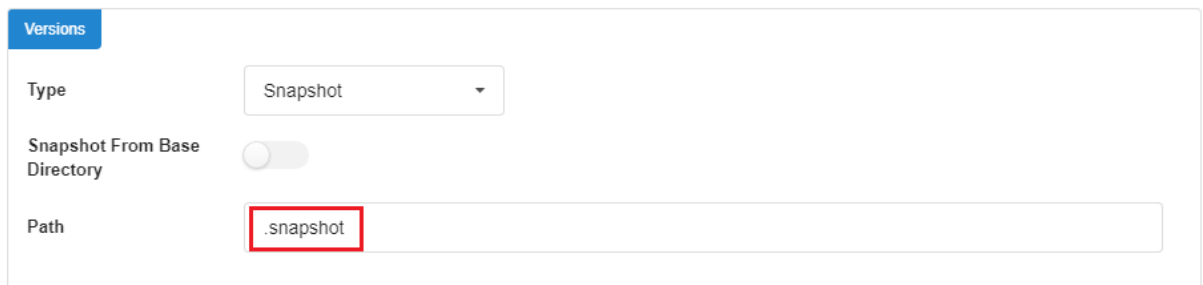


Figure 3.3.b

In the FileOrbis GUI, when a user right-clicks a file and selects “See Versions” all the versions of the file are listed. Any of those versions can be downloaded, previewed, and restored.

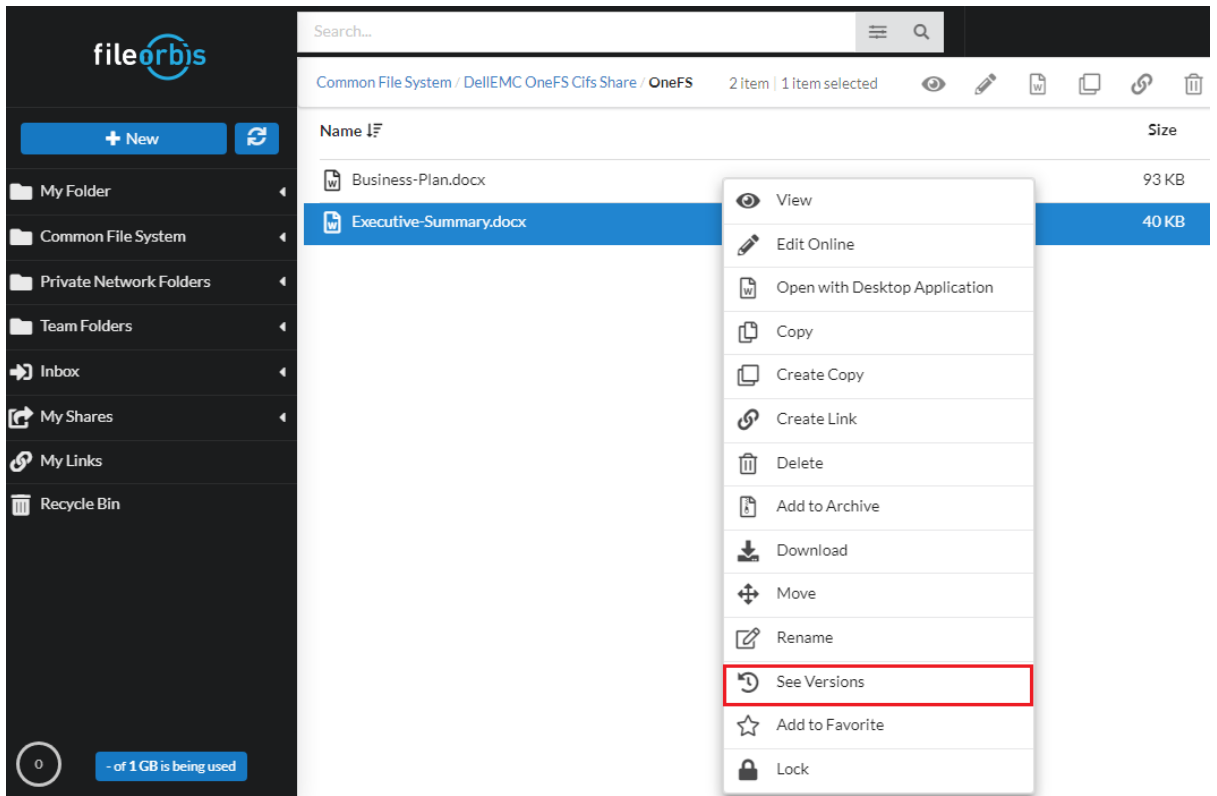


Figure 3.3.c

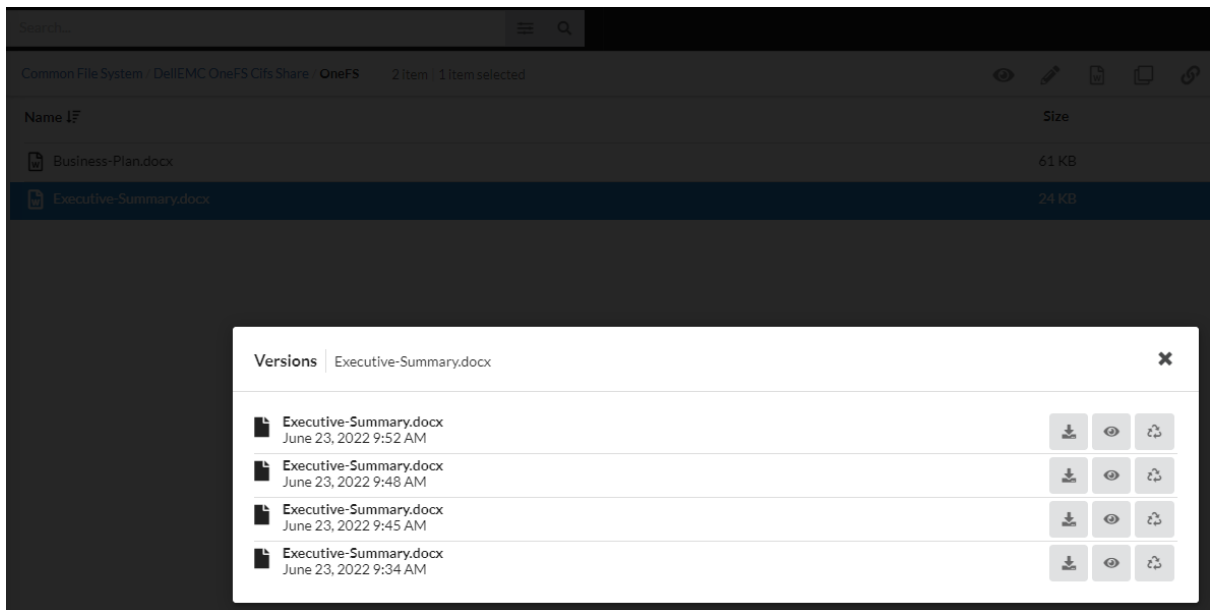


Figure 3.3.d

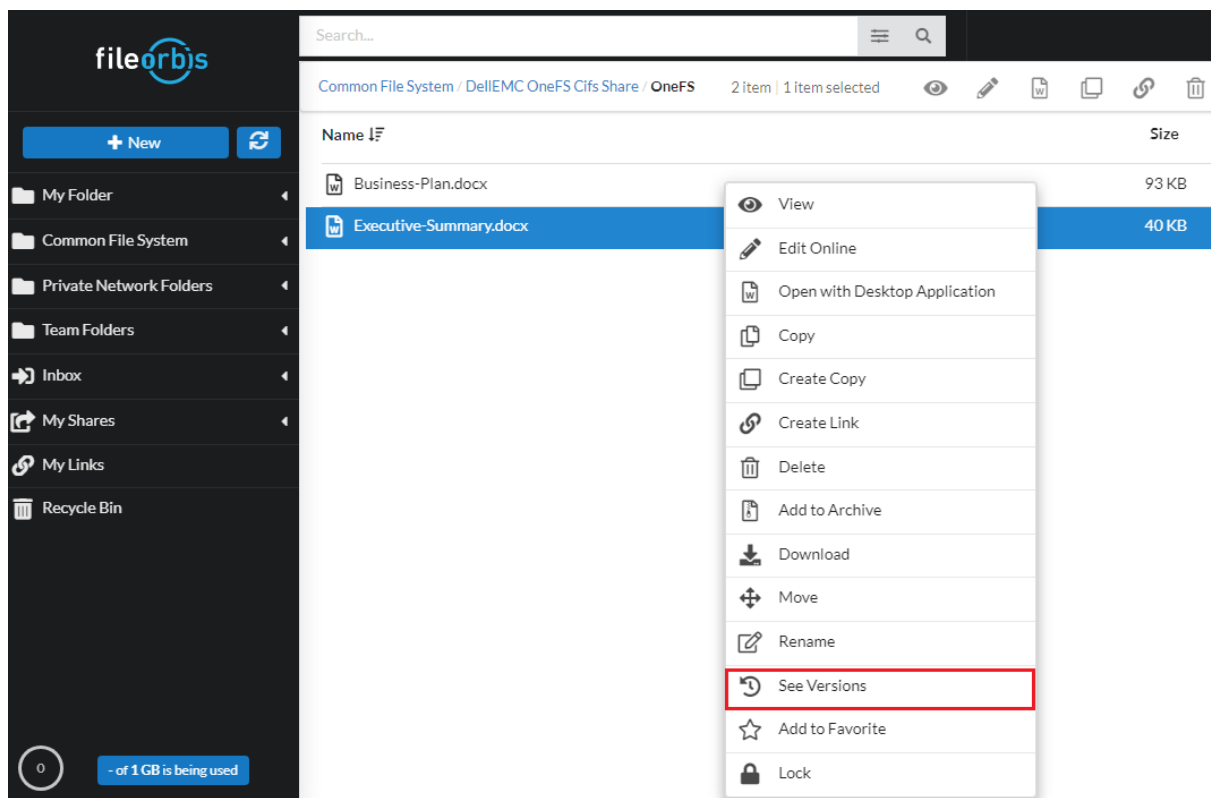


Figure 3.3.e

3.4. Amazon S3 Protocol Access To File Classification Tags


With the integration of the DataGovernance Manager component of FileOrbis and Extended File Attributes in Dell EMC PowerScale OneFS, it is possible to encapsulate files in OneFS filesystem with FileOrbis data classification tags, as explained earlier.

FileOrbis data classification tags are stored as extended file attributes in OneFS. Extended file attributes in OneFS filesystem are also used to store and retrieve custom object metadata. As such, this enables cloud applications using the Amazon S3 protocol to access the same key/value pairs as object custom metadata. Cloud applications can use these key-value pairs to make calculated decisions, do faster searches based on metadata rather than content, and classify data that has already been classified by content.

3.5. Content Based Batch Operations for Files and Metadata

Batch rules are sets of instructions or algorithms that are used to process large volumes of data efficiently. Batch processing is a technique used in computing whereby large amounts of data are processed in batches or groups, rather than being processed individually or in real-time.

The business units in an organization own the data, but it is the job of IT to store, back up, and make the data available to the right people at the right time based on what the organization needs during its own life cycle. IT can use its infrastructure more



efficiently and agilely by automating data type information. It's like puzzle-solving. You just need to arrange the pieces.

Your Dell EMC PowerScale storage system may have millions of files with years of data. You want to use tag data to categorize these files by content and maximize efficiency. Users may rarely read, rename, or update these repositories. Thus, while previous integrations allowed action-based work, file pools with static data cannot. Assume these are old files. Even if the user account associated with these files has been removed from the active directory, the file may still be used by the company. It may also contain GDPR-sensitive personal data. How can we decide without knowing what's in these files? If action-based work isn't possible, use batch rules. Who owned the old files? Who can access those files? The FileOrbis Data Governance Manager module and Dell PowerScale OneFS SmartPools file pool policies allow us to assign the correct custom name attribute to that file without considering the user, even if they no longer have an active directory account.

- The same scenarios as in section 3.1 apply, such as storing six copies of a file or moving it to the SSD tier for quick access or the SATA tier for archiving.
- You can commit a cold file to WORM status with SmartLock
- You can move a file's metadata to the SSD tier to help you find it more quickly in a sea of files.

With Dell EMC PowerScale OneFs and FileOrbis, business logic-based storage optimization using file metadata generates impressive results. This Dell EMC PowerScale OneFs-specific integration automates batch operations to proactively and content-based manage all organization data. When content-based decisions can be made for retrospective bulk data, you will have a content strategy that fits organizations and their infrastructures better, is more scalable and agile, and addresses all data.

4. Dell EMC ECS Integration Capabilities

Dell EMC ECS system can be added as a Common File System to FileOrbis using the S3 protocol. It brings file server grade hierarchy management and permission capabilities to Dell EMC ECS when ECS is the back-end.

An object user created on Dell EMC ECS is defined as an authenticated user for the bucket which is mapped to FileOrbis.

Name * ⓘ

s3user LOCK USER

Namespace *

ns1

Object Access

S3 / Atmos ⓘ

Show Secret Key

..... DELETE

GENERATE & ADD SECRET KEY

Figure 4.a

Basic > Required > Optional

Name * ⓘ

DellEMC

Arn

arn:aws:s3:::DellEMC ⓘ

Namespace *

ns1

Replication Group *

rg1

Bucket Owner (Used for Non IAM Access) * ⓘ

s3user Set current user as Bucket Owner

Figure 4.b

After that the bucket on Dell EMC ECS will be mapped to FileOrbis as a Common File System using the S3 protocol it will look like as in the Figure 5.c below.

File System Settings

The screenshot shows the 'File System Settings' interface. It is divided into two main sections: 'File System' and 'Access'.
Under 'File System', the 'Type' is set to 'Amazon S3'.
Under 'Access', the following fields are visible:
- 'Name': Dell EMC ECS S3
- 'Endpoint': https:// [redacted]
- 'Region': eu-central-1
- 'Bucket': DellEMC
- 'Root Path': [empty]
- 'Cache Directory': C:\ProgramData [redacted]

Figure 4.c

4.1. User-Initiated File Restores Based on Versions

Thanks to S3's versioning feature, every time a change is made to the file in the bucket, the previous version of the file is kept as a version. All these versions can be listed from the FileOrbis interface, and any version can be downloaded, previewed, and restored as the current version of the file.

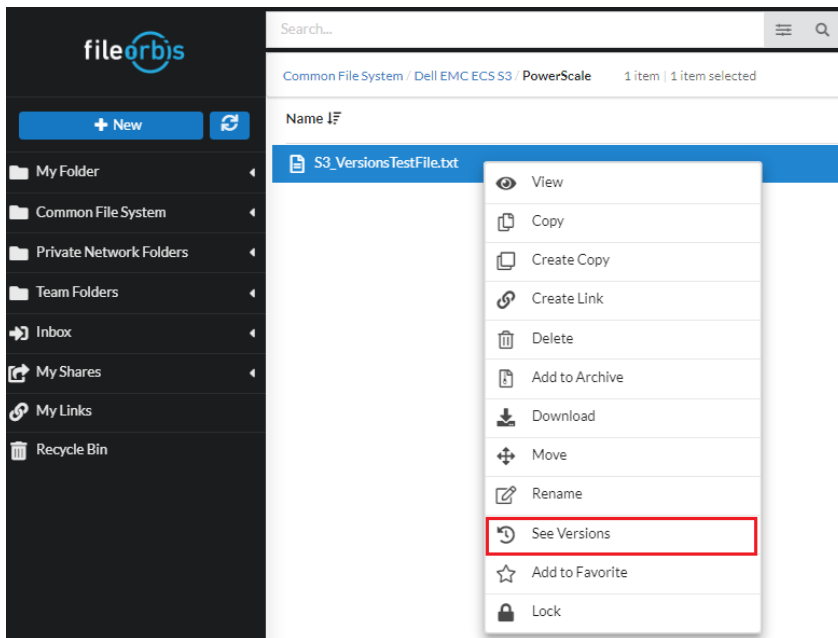


Figure 4.1.a

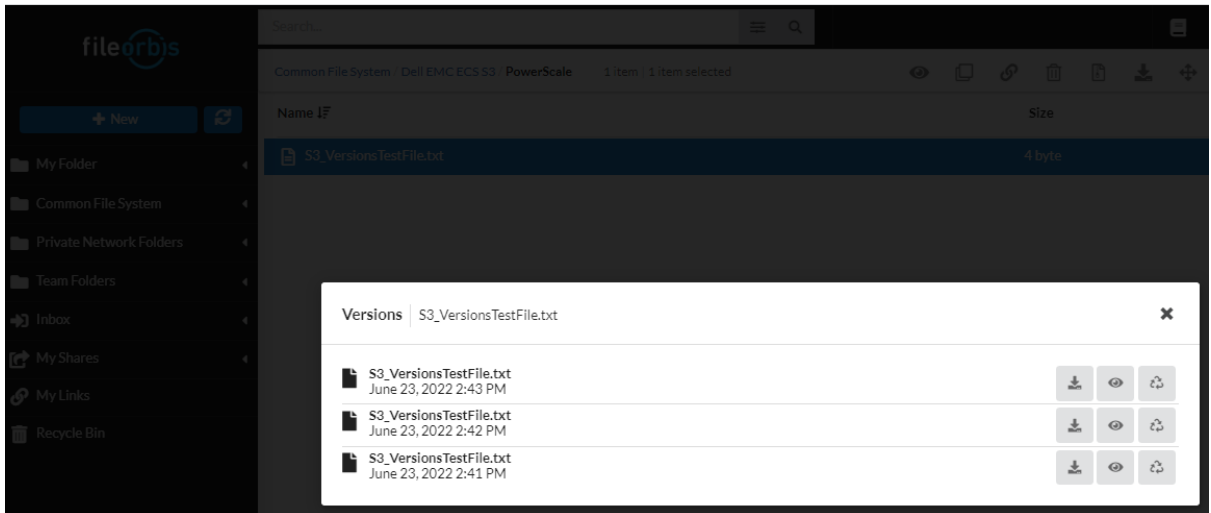
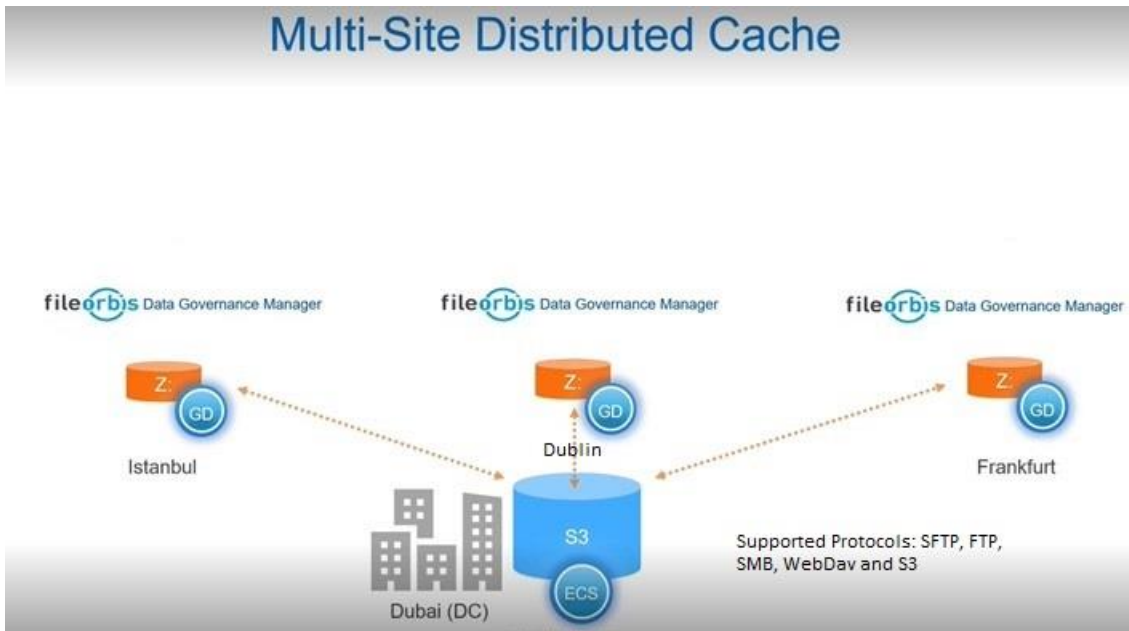


Figure 4.1.b

4.2. GeoDrive

Dell EMC ECS bucket is presented as a local drive in Windows via GeoDrive, therefore users and applications can store data in the cloud using Windows-native apps without investing in any new IT infrastructure. GeoDrive uses a disk cache to store files temporarily, and can upload them to the cloud asynchronously or in real time. There are many use cases that can be addressed using this infrastructure. This architecture also allows FileOrbis to use global file locking and versioning, as described in the previous section.



With single sign-on capability and complete integration with Active Directory, FileOrbis provides data access interfaces such as SFTP, FTP, and WebDav to Dell EMC ECS.

All file transfers through GeoDrive are proactively analysed for security issues, thanks to the platform's API-level integration with existing security tools such as AV, CDR, DLP, and our true type file screening capable Firewall module. With the unique and secure sync and share features, FileOrbis users can share their files both internally and externally using Dell EMC ECS GeoDrive as a gateway to their corporate's central Dell EMC ECS seamlessly. The combined solution allows for in-depth auditing of file events including originating IP addresses. Reports can be imported for use with SIEM and SOAR applications, and downloaded as Excel documents.

5. Other Integration Points

5.1. WebDAV Maps

A user can access storage space on Dell EMC OneFS or Dell EMC ECS that is associated with his or her FileOrbis user profile directly, either from an internal or external network using WebDAV over HTTPS. FileOrbis user profile folders can be mapped to the user's computer as a network location.

The point of this integration is that regardless of whether the back-end type is Dell EMC PowerScale, Dell EMC ECS, or both, the same hierarchy of files and folders that might have been spread across multiple back-ends, is presented to the end users in a unified manner. In addition, rich ACL support can be enabled without reliance on the back-end system. In Dell EMC ECS case, which is an object store without inherent rich file system ACL support, this integration fulfills the need for rich ACL-based permission management.

To enable this functionality, the WebDAV protocol should be enabled in the FileOrbis Management GUI, as can be seen in Figure 5.1.a below.

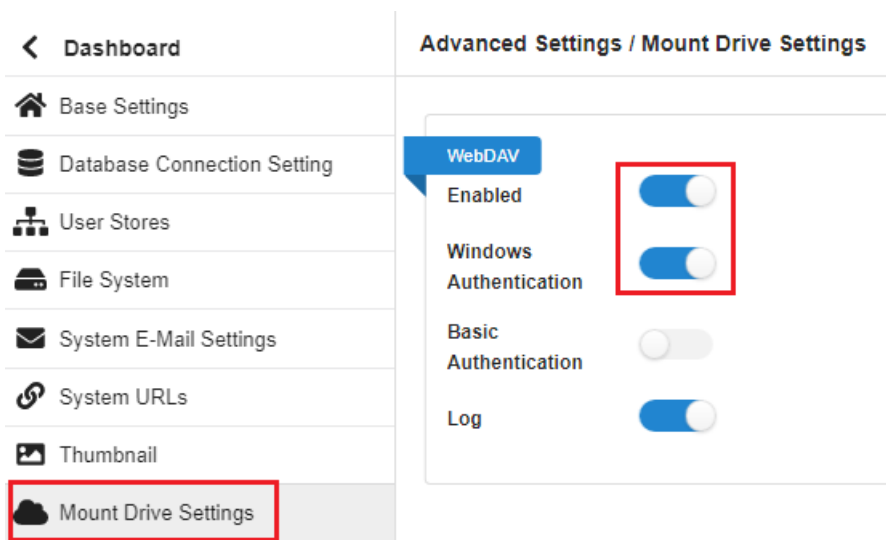


Figure 5.1.a

The web client service that supports the WebDAV protocol should be set to start automatically on the end user's PC.




Figure 5.1.b

WebDAV link is obtained and copied from the FileOrbis User GUI.



Figure 5.1.c




Figure 5.1.d

When setting up the mapping to a drive letter, you can paste a WebDAV link into the Folder field.

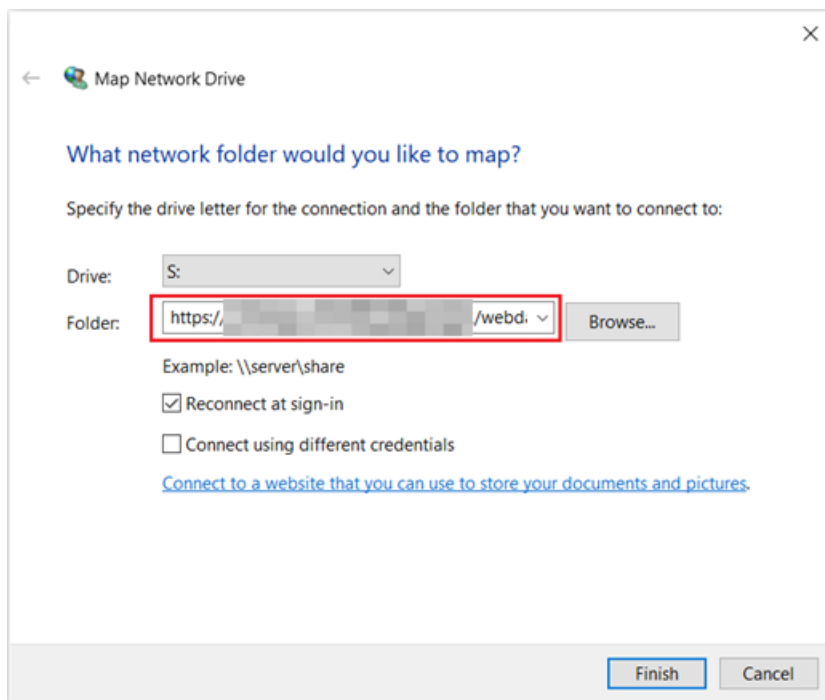


Figure 5.1.e




Figure 5.1.f

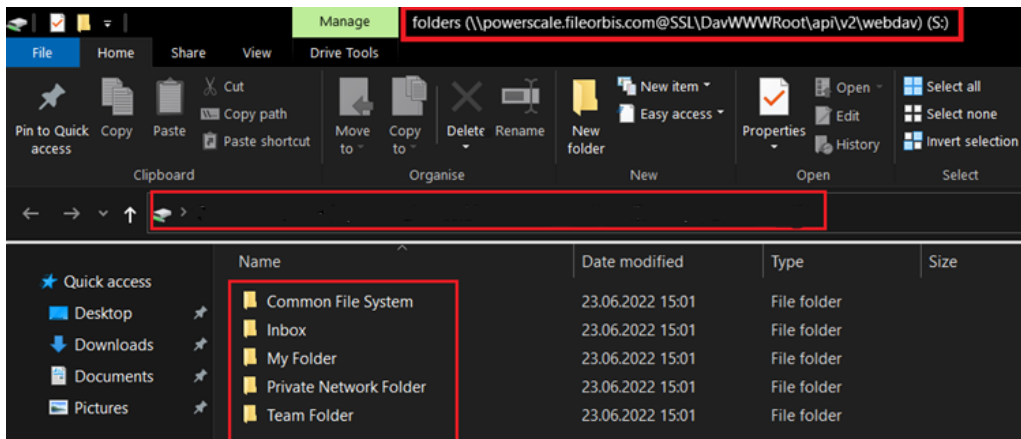


Figure 5.1.g





Figure 5.1.h

5.2. Full Text Search and File Disposal

Full-text search lets users find information quickly without the file name or location. Manually searching each file or folder without full text search is inefficient and time-consuming in large data sets. Full-text search searches the repository for files with



keywords or phrases. Time and data savings boost workplace productivity. Dell EMC PowerScale and FileOrbis can find hidden patterns in large data sets for analysis and decision-making.

To avoid being accessed without authorization, files should be properly discarded. For organizations that handle sensitive data like personal or financial data, improper file disposal can lead to data breaches and other security issues. Disposing of old files also frees up server storage, improving performance and lowering costs. The Dell EMC PowerScale system and FileOrbis can help IT regularly delete unnecessary user files to keep file servers efficient, secure, and effective.

Appendix: Related Resources

Below a list of documents and some assets that are referenced in this paper that may help to explain the solution components individually.

[PowerScale OneFS Technical Overview](#)

[ECS: Overview and Architecture](#)

[Whitepaper: Enterprise File Sharing with Integrated Security Solutions](#)

[Whitepaper: Next Generation File Sharing & Remote Access to Files and Folders](#)

[Whitepaper: FileOrbis' Governance Capabilities and Data Storage Policies](#)

[Datasheet: Secure Content Management](#)