Miria User Documentation

Miria 4.0

Publication Number: MIRIA-USR-PDF-EN-0223-REV1

Publication Date: February 2023

Atempo SAS
23 Avenue Carnot
91300 Massy - France

Contents

CHAPTER 1 - About Miria User Documentation 1

CHAPTER 2 - Web User Interface Overview 2

Features 2
Permissions 2
Miria Web Interface 2

CHAPTER 3 - Connect to Miria Web Interface 6

CHAPTER 4 - Move Data 7

Archive Data 7
Copy, Move, Synchronize Data 8
Retrieve Data 9
Relaunching an Easy Move Job 9

CHAPTER 5 - Organize Repositories 11

Repository Types 11
Repository Organization 11
Folders and Sub-Folders 11
Archived Directories and Files 11
Manage Repositories 11
Add Folder 12
Rename Object 12
Delete Object 12
Permissions on the Repositories 13
Manage Metadata 14
Applying Metadata to Repositories 15
Manage Instances 16
Retrieve Instance 17
Instance Details 19
Session Details 21

CHAPTER 6 - Search for Repositories 22

Browse the File Tree 22
Use Go To Path 22
Search by Criteria 22
Search by Metadata 25

CHAPTER 7 - Manage Video Assets 28

Preview Video Assets 28
Play Video Asset Preview 28

Video Player Overview 28
Partially Retrieve Video Asset 29

CHAPTER 8 - Monitor Data 31
Show a Job Menu 32

CHAPTER 9 - Manage the Infrastructure 33
Add a Storage Manager and Container 33
File Storage Container 34
File Storage One to One 35
SnapStor 37
Virtual Storage 38
Media Storage 38
Easy move Amazon S3 40
Google Cloud Storage 42
Microsoft Azure Blob Block 43
Scality 45
Quantum Active Scale 47
Seagate Lyve Cloud 49
Cloudian HyperStore 50
Add Server and Agent(s) 52
Explore Agent 53
Add a New NAS Platform 53
Add a Storage Platform 54
Add a Shared File System 55
Activate My File System 55
Available options 56
Edit a Platform 57
Platforms Permissions 58

CHAPTER 10 - Organize and Configure the Migration 60
Organize Projects 60
Project description 60
Task description 60
Configure the Migration 60
Create a New Project 60
Create a New Task 61
Edit a Project 63
Edit a Task 63
Start Migration Task(s) 64
Duplicate a Task 64
Disable a Project or a Task 64

CHAPTER 11 - Manage the Migration 66

Step 1: First Synchronization 66
Step 2: Incremental Synchronization 67
Step 3: Final Synchronization (Cutover) 68

CHAPTER 12 - Organize and Configure a Backup Project 72

Organize Projects 72
Project description 72
Task description 73
Configure the Backup 73
Create a New Project 73
Create a New Backup Task 73
Create a New Tiering Task 75
Edit a Project 76
Edit a Task 76
Start Task(s) 76
Duplicate a Task 77
Disable a Project or a Task 77

CHAPTER 13 - Manage Users 78

Add a User 78
Edit a User 79
Add a User Group 79
81
Configure SAP 81
Set Password Policies 82
Add a LDAP Rule 83
Add a LDAPS Rule 83
Edit a Rule 83
Set the Two-Factor Authentication 83
Configuring a Two-Factor Authentication Set as Mandatory 84
Configuring a Two-Factor Authentication Set as Optional 85

CHAPTER 14 - Advanced Tasks 88

Task Types 88
Internal Management Tasks 88
Data Movement Tasks 89
Creating a Task 91
Creating a New Task from the Tasks Tab 91
Duplicating an Existing Task 91
Organizing tasks 91
Manage projects 92
Move a task 92

How Tasks Run 92
Testing a Task 94
Launching a Task Manually 94
Canceling a run 94
General Configuration Parameters for Tasks 95
General Parameters Common to All Tasks 95
Common Tabs 97
Task-specific Parameters 100
Automatic Catalog Ingest Task 100
Automatic Deletion Task 103
Automatic Retention Tasks 104
Automatic Storage Repack Task 105
Database Backup Task 106
Maintenance tasks 106
Archive Report Task 109
Volume Management on Storage Managers Task 111
XML Ingest Tasks 111
CHAPTER 15 - Parameters 119
Default Settings 120
Security Settings 120
Interface Settings 122
Jobs Settings 123
Media Settings 143
Platforms rights 144
Admin and monitoring rights 144
Repositories rights 145
E-mail Settings 146
Settings Templates 146
Retentions 148
Run Locks 149
Run Timeframes 149
Metadata 150
SMTP server 152
Media Manager 153
Media Manager Overview 154
Media Manager Libraries Tab 158
Media Manager Configuration Tab 159
Media Manager Media Tab 162
Media Manager Logs Tab 170
Syslog server 171
Deduplication Domains 172

To create a Deduplication Domain 173
Volume Report 174
Repository 175
Storage manager 175
Dedup. Domain 175
Custom Media Rules 176

APPENDIX Recycling Triggered by Volume on Storage 178

APPENDIX Replications in between two S3 object storages 180

CHAPTER 1 - About Miria User Documentation

Miria is a data management product developed by Atempo. It is designed to manage media and digital assets.

This documentation describes how you can use the Web Interface to perform the following data-move operations:

- Archiving data.
- Retrieving data.
- Copying data.
- Synchronizing data.
- Moving data.
- Monitoring data.
- Previewing video.

CHAPTER 2 - Web User Interface Overview

This chapter provides a detailed presentation of the Web Interface and its features to manage repositories and perform data-move operations.

Features

- **Easy-to-use** The Web Interface makes it easy for you to archive, copy, synchronize and retrieve data.
- **Classification** Move data in personal or in shared projects repositories to which Miria assigns the appropriate metadata information.
- **Advanced search across repositories** The Web Interface is equipped with advanced search capabilities within your own user repository and across all the projects to make data retrieval easy.
- **High performance** Miria handles the data flow to archive or retrieve continuously or in a deferred manner after optimization of requests.

Permissions

To manage project repositories, the administrator defines the level of permissions granted to each user. Next to your personal user repository, the Web Interface displays the repositories you are allowed to access.

The following repository operations require specific permissions:

- Opening and viewing repositories other than your personal user repository.
- Performing data-move operations.
- Creating, renaming, or deleting repositories.
- Moving archived folders within repositories.
- Moving, renaming, or deleting archived files and directories.
- Modifying file extensions.

Miria Web Interface

The Web Interface is divided into a navigation pane on the left and a central page (Figure 1).

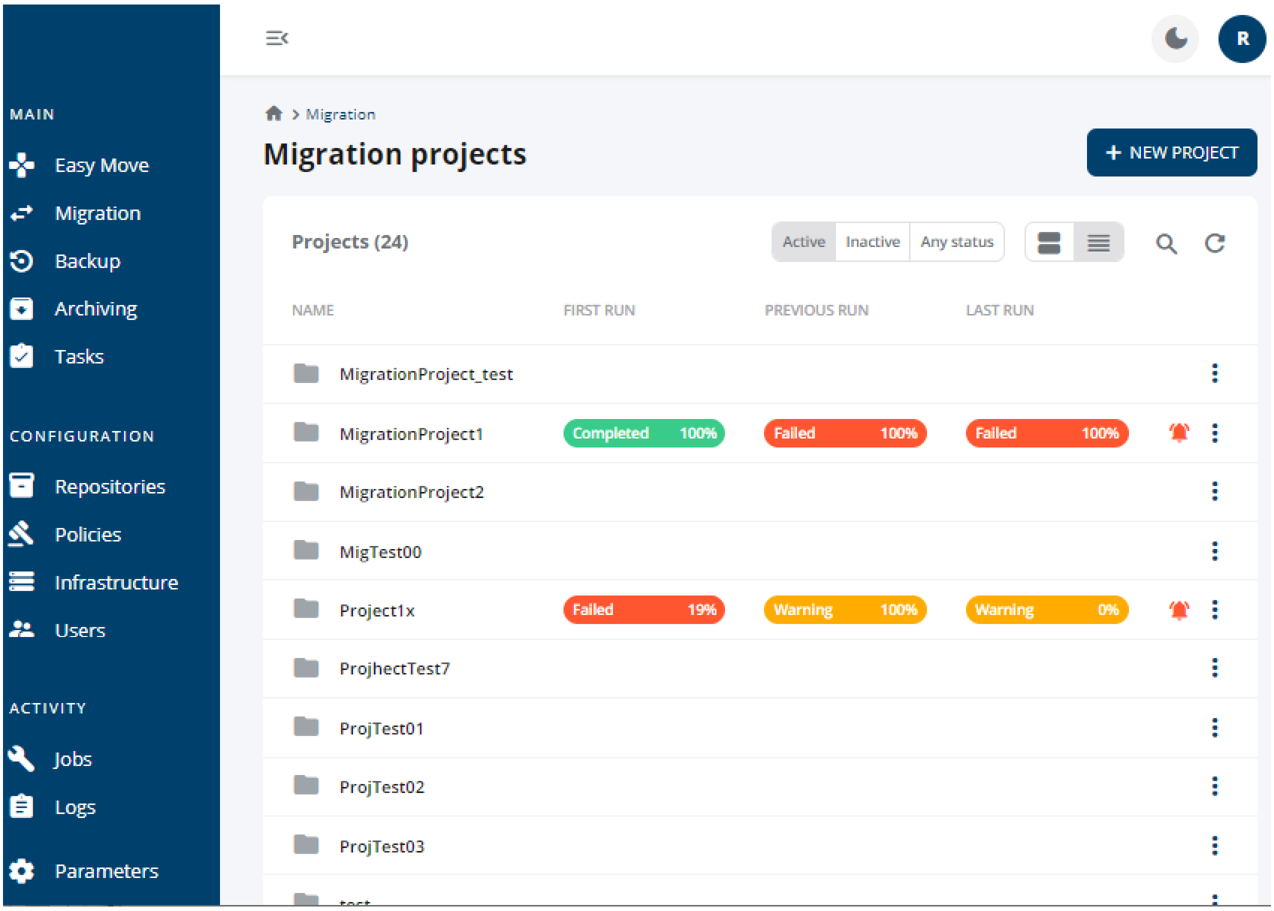

Migration projects

Projects (24)

NAME	FIRST RUN	PREVIOUS RUN	LAST RUN
MigrationProject_test			
MigrationProject1	Completed 100%	Failed 100%	Failed 100%
MigrationProject2			
MigTest00			
Project1x	Failed 19%	Warning 100%	Warning 0%
ProjhectTest7			
ProjTest01			
ProjTest02			
ProjTest03			
test			

Figure 1: Web Interface

The navigation pane gives access to 12 different tabs:

- The **Easy Move** tab lets you:
  - Define a source and target to archive, retrieve, copy, move, or synchronize data.
  - Perform operations directly in the platform or repository tree (add, rename, delete etc.).
  - Visualize Easy Move jobs history.
  - Launch an Easy Move job you have previously launched through the **Job History** page. By suggesting a list of the last 100 operations, Easy Move gives you the opportunity to relaunch a job with the same execution conditions than the original one. This way, you can perform again an Easy Move operation, but faster.
- The **Migration** tab, enables end-to-end management of the migration project. It allows you to:
  - Configure your migrations.
  - Supervise operations.
  - View and send Activity reports by email.
  - View Statistics reports.
- The **Backup** tab allows you to backup and restore large file storages, Object, Cloud and NAS.
- The **Archiving** tab allows you to archive and restore large file storages, Object, Cloud and NAS.

- The **Tasks** tab enables you to create and manage tasks. Knowing that a task defines the scope of a Miria job, the source and destination of the data that it processes, its scheduling, and many other options.
- The **Repository** tab lets you:
  - Display volume information about each repository as well as its owner and type.
  - Access your own user repository or a project.
  - Search in projects and repositories.
  - Include metadata to search for objects in repositories.
- The **Polices** tab enables you to define which storage manager container(s) to use and how long each container retains the data.
- The **Infrastructure** tab lets you see and manage different storage (e.g. agents, local filesystem, shared filesystems, etc.).
- The **Users** tab enables you to create and manage users in order to access and, for instance, manage permissions.
- The **Jobs** tab provides detailed information to each job.
- The **Logs** tab enables you yo opens logs file, that contain an history of any change that happened on a job.
- The **Parameters** tab is divided into different sections that you can manage:
  - **Default settings.** Access the settings and manage them.
  - **Settings templates.** Creates templates to define settings by default.
  - **License.** Access your license information.
  - **Retentions.** Manage the retentions.
  - **Run locks.** Manage the run locks.
  - **Run Timeframes.** Create and manage the run timeframes.
  - **Metadata.** Create and manage metadata.
  - **SMTP server.** Specify an SMTP server.
  - **Syslog server.** Connect Miria to a Syslog server, and export logs towards it.
- **Profil.** Button that includes product information and lets you manage your user credentials.
- **Dark mode.** Button that enables you to change the interface into dark mode (Figure 2). When you do so, the dark mode button turns into a light mode button.

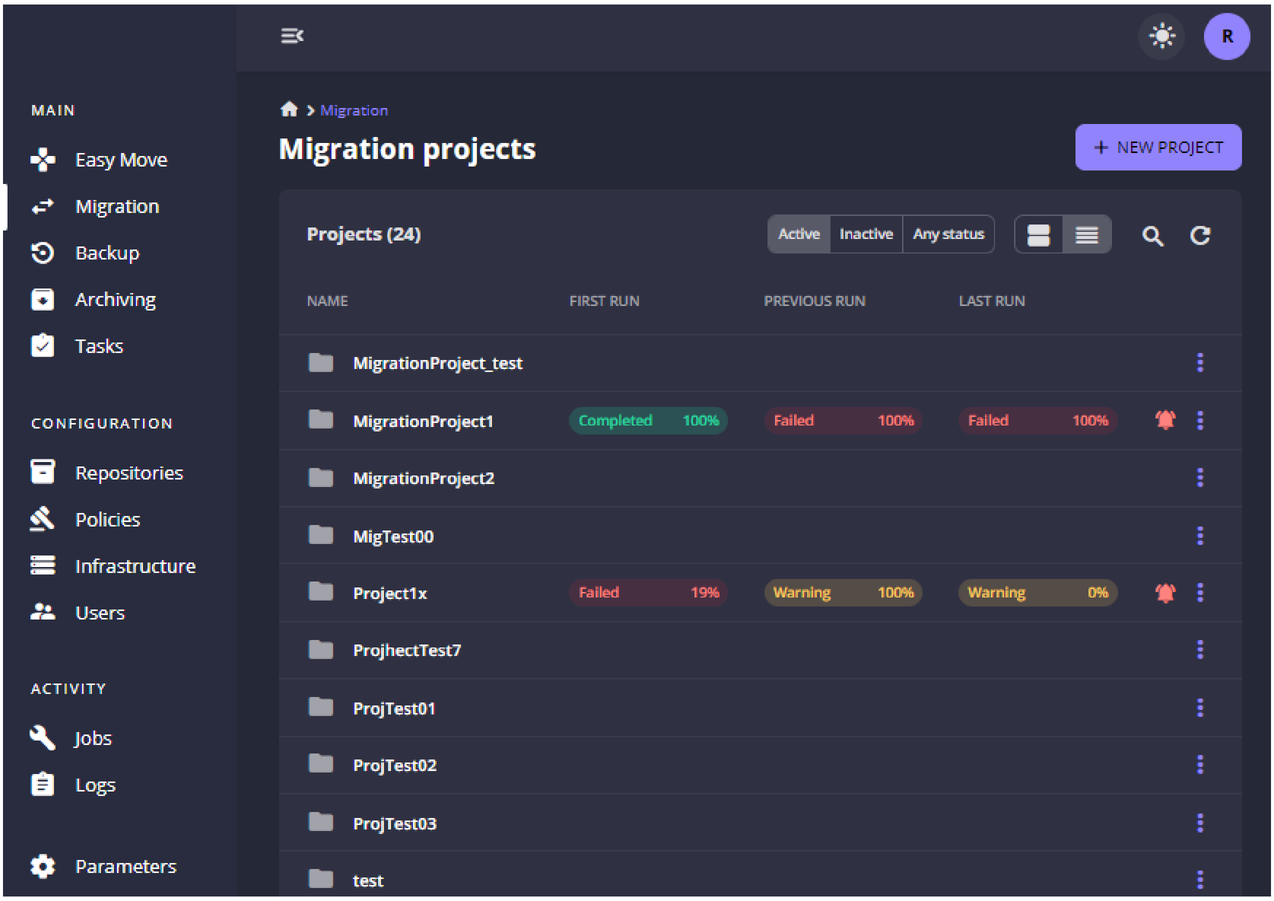

Migration projects

Projects (24)

NAME	FIRST RUN	PREVIOUS RUN	LAST RUN
MigrationProject_test			
MigrationProject1	Completed 100%	Failed 100%	Failed 100%
MigrationProject2			
MigTest00			
Project1x	Failed 19%	Warning 100%	Warning 0%
ProjhectTest7			
ProjTest01			
ProjTest02			
ProjTest03			
test			

Figure 2: Dark mode

CHAPTER 3 - Connect to Miria Web Interface

1. Open a web navigator.
2. Type the server URL: `<Miria server IP Address>/webapp-en`. If you want to display the interface in French or Chinese you can enter `webapp-fr` or `webapp-zh`.The login window is displayed.
3. Enter your authentication information:
   - **Username** Name of the user as defined in Miria Administration Console.
   - **Password** Password associated with the specified username.
4. Click **Login**.

CHAPTER 4 - Move Data

The **Easy Move** interface is divided into a source area and a target area (Figure 3).

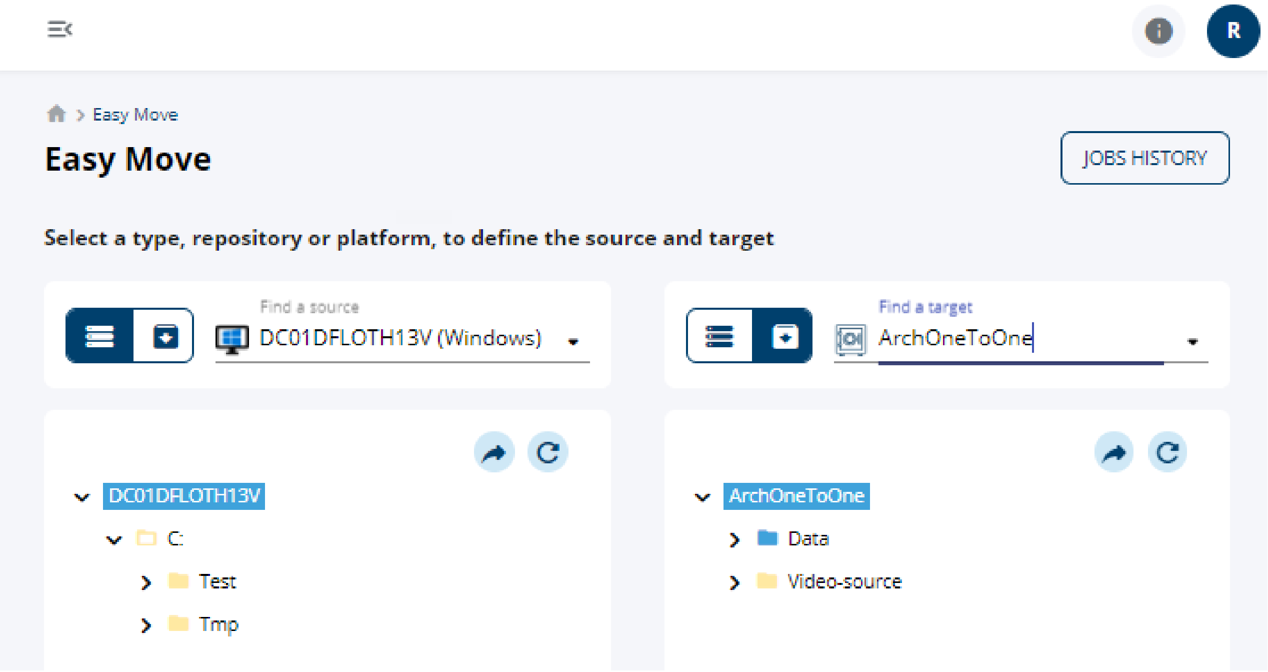

Figure 3: Easy Move interface with a source and target selection

Move operations always go from a source on the left to a target on the right. You can perform the following actions:

- **Archiving** Moving data to a storage location where it remains available for future retrieval.
- **Retrieval** Transfer data from the storage location to your workstation or any other computer where you can view or edit it.
- **Copying** Copy data from one platform to another. The data will be present on both the source and target platforms.
- **Moving** Move data from one platform to another platform. The data will no longer be present on the source platform and is available on the target platform.
- **Synchronizing** Compare a source data directory with a target data directory. An object present on the source but missing from the target is copied into the target.
- **Relaunching a job** Through the **Jobs history** button, you can access **Easy Move** jobs history and relaunch a job, without having to reconfigure it.

This chapter explains how to perform each one of these tasks with Easy Move.

Archive Data

1. In the Web Interface, click the **Easy Move** tab.
2. In the left field, select the source type as **Platform** and find your source (Figure 4).

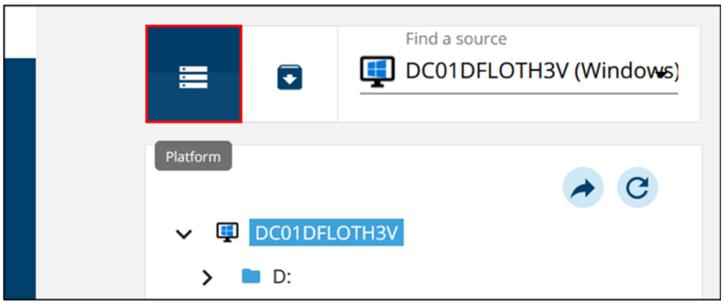

Figure 4: Platform selected as source

3. In the right field, select the target type as **Repository** and find your target.
4. Explore the file tree of the platform and select the object you want to move.
5. Explore the file tree of the repository and select the destination folder or directory. You can also add a new folder.
6. Move the object:
   - Click **Add**.

   Or
   - Drag the object from the source to the desired location in the target.
7. Repeat steps 3, 4 and 5 for each move you want to perform.
8. Validate the basket.
9. If the administrator has configured metadata, enter a value for the metadata of your choice and click **Continue**. The archiving job is launched.
10. Click the **Jobs** tab to verify the job progress.

Copy, Move, Synchronize Data

1. In the Web Interface, click the **Easy Move** tab.
2. Select the source and target types as **Platform** and find your source and target (Figure 5).

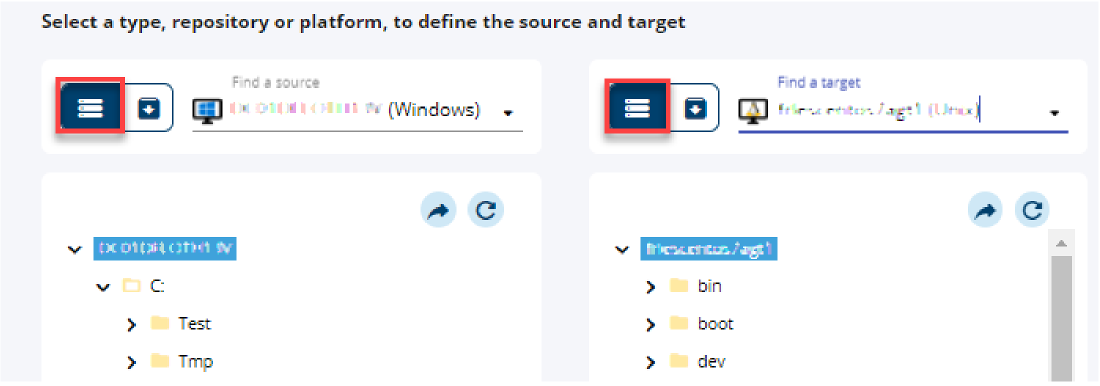
Select a type, repository or platform, to define the source and target

Figure 5: Platform selection

3. Explore the file tree of the source platform and select the object you want to move.
4. Explore the file tree of the target platform and select the destination directory.
5. Move the object:
   - Click **Add**.

Or

- Drag the object from the source to the desired location in the target.

6. Repeat steps 3, 4 and 5 for each move you want to perform.
7. Choose the task to perform: **Copy**, **Move** or **Synchro**.
8. Validate the basket.
9. Click the **Jobs** tab to verify the job progress.

Retrieve Data

1. In the Web Interface, click the **Easy Move** tab.
2. In the left field, select the source type as **Repository** and find your source (Figure 6).

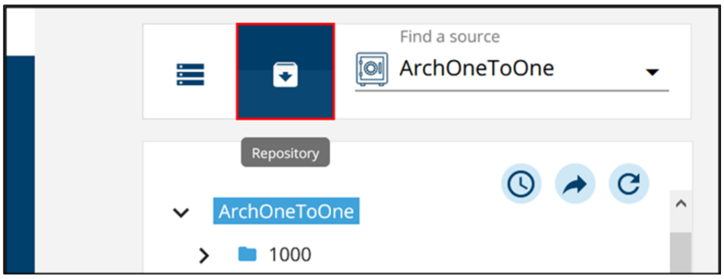

Figure 6: Repository selection

3. In the right field, select the target type as **Platform** and find your target.
4. Explore the file tree of the repository and select the object you want to retrieve.
5. If needed, click the **Time line** 🕓 button to use time navigation and retrieve data at a specific date and time.
6. Explore the file tree of the platform and select the destination directory. You can also add a new directory.
7. Move the object:
   - Click **Add**.

   Or

   - Drag the object from the source to the desired location in the target.
8. Repeat steps 3, 4 and 5 for each move you want to perform.
9. Validate the basket. The retrieval job is launched.
10. Click the **Jobs** tab to verify the job progress.

Relaunching an Easy Move Job

1. In the Web Interface, click the **Easy Move** tab.
2. Click the **Jobs history** button. The **Jobs** pane opens (Figure 7). The **Easy Move** filter is selected.

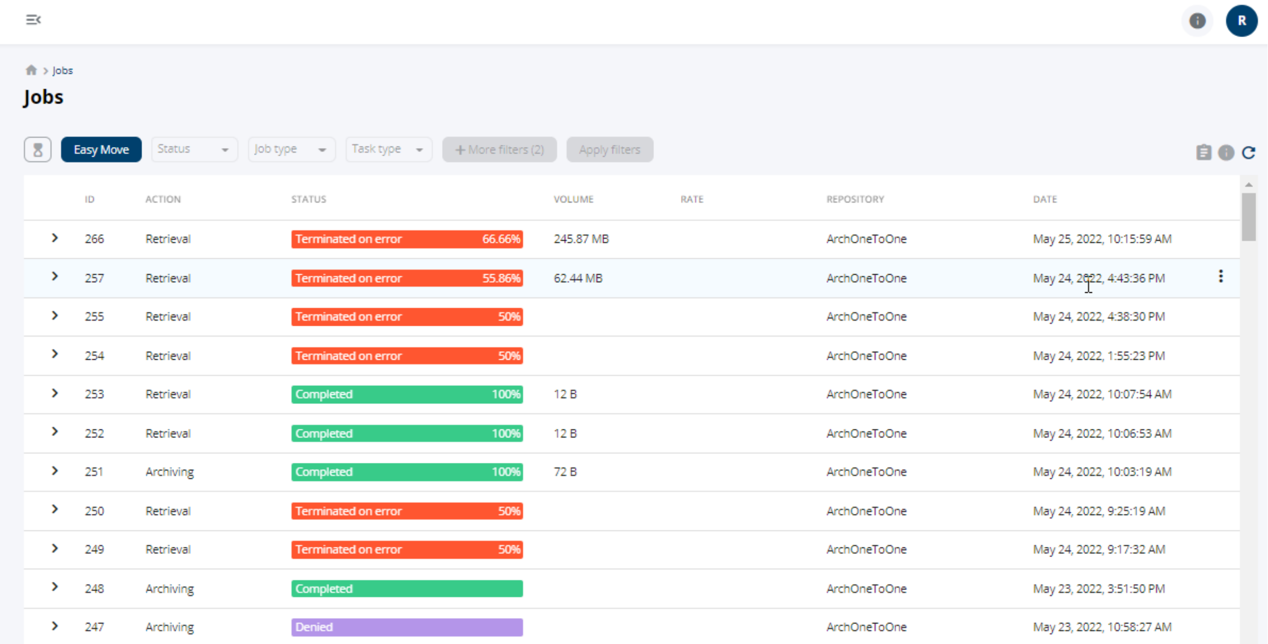

ID	ACTION	STATUS	VOLUME	RATE	REPOSITORY	DATE
266	Retrieval	Terminated on error 66.66%	245.87 MB		ArchOneToOne	May 25, 2022, 10:15:59 AM
257	Retrieval	Terminated on error 55.86%	62.44 MB		ArchOneToOne	May 24, 2022, 4:43:36 PM
255	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 4:38:30 PM
254	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 1:55:23 PM
253	Retrieval	Completed 100%	12 B		ArchOneToOne	May 24, 2022, 10:07:54 AM
252	Retrieval	Completed 100%	12 B		ArchOneToOne	May 24, 2022, 10:06:53 AM
251	Archiving	Completed 100%	72 B		ArchOneToOne	May 24, 2022, 10:03:19 AM
250	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 9:25:19 AM
249	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 9:17:32 AM
248	Archiving	Completed			ArchOneToOne	May 23, 2022, 3:51:50 PM
247	Archiving	Denied			ArchOneToOne	May 23, 2022, 10:58:27 AM

Figure 7: Easy move jobs history

3. Select one of the row. A menu appears.
4. Click on **Relaunch**. This window opens (Figure 8):

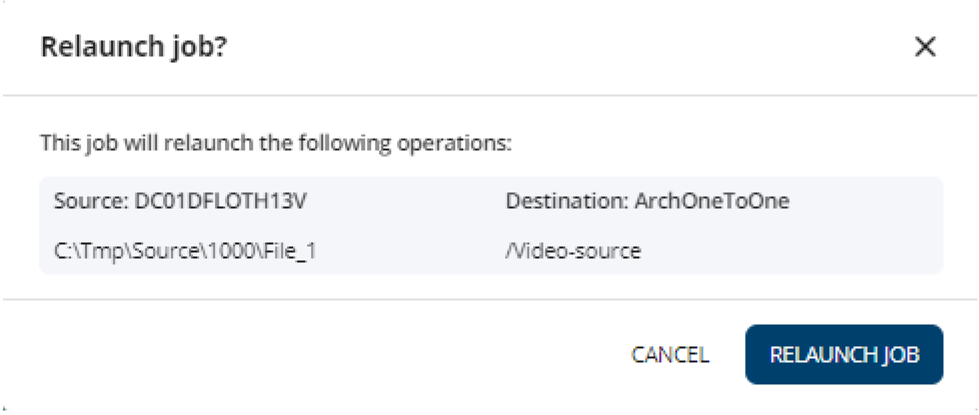

Relaunch job?

This job will relaunch the following operations:

Source: DC01DFLOTH13V	Destination: ArchOneToOne
C:\Tmp\Source\1000\File_1	/Video-source

Figure 8: Relaunch a job

5. Click **Relaunch job**. The job is relaunched.

CHAPTER 5 - Organize Repositories

A repository is a centralized storage location to preserve and manage data for long-term retention. This chapter explains how to structure data within repositories.

Repository Types

There are two types of repositories:

- **Archive**. The process of archiving involves moving data from one location to another location for long-term retention. Data archives are classified in folders and sub-folders that you can organize at your convenience.
- **Backup**. A backup consists of copying data to a repository with the intent to keep the original data in its current location. Backups are useful if the original data are lost or corrupted and you want to restore it to a certain point in time.

Repository Organization

When you log in to the Web Interface, it displays the projects and repositories that are shared with you and other users.

- **Projects** Miria provides an organization tool to group repositories into projects. They can be accessed by multiple users.
- **Repositories** A repository contains folders and sub-folders, in which directories and files are archived. The administrator may define one or multiple repositories that belong to a user and can only be accessed and managed by its owner.

Folders and Sub-Folders

Folders and sub-folders are used to organize archived data in a file tree structure. When you explore a repository, they are represented by a folder icon: 📁 .

You can both manually create and organize folders to classify your archived data. Folders can also be generated by automatic archiving tasks set up by the Administrator, or by external tools.

Archived Directories and Files

Archived directories and files are located in repository folders. When you explore a repository, the archived files are represented by an icon indicating their application. You recognize directories by the directory icon: 📁.

Manage Repositories

When exploring the repositories, various object options are accessible.

When holding the mouse pointer over an object, you can perform multiple operations using the **More actions** ⋮ button (Figure 9):

- Archive data directly in the **Easy Move** interface.
- View and retrieve the instance of a file.
- View and manage metadata on an object.
- Perform a search by criteria on an object.
- Retrieve data at a specific date by using the time navigation.

Alternatively, you can select and drag folders and directories to change their location in the repository.

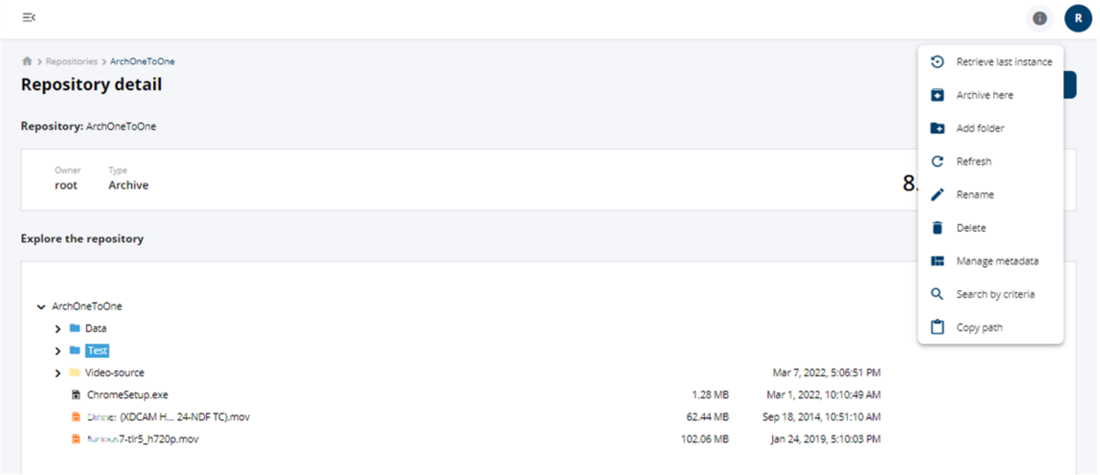

Repositories > ArchOneToOne

Repository detail

Repository: ArchOneToOne

Owner	Type
root	Archive

Explore the repository

ArchOneToOne		
Data		
Test		
Video-source		Mar 7, 2022, 5:06:51 PM
ChromeSetup.exe	1.28 MB	Mar 1, 2022, 10:10:49 AM
Dinner (XDCAM H... 24-NDF TC).mov	62.44 MB	Sep 18, 2014, 10:51:10 AM
[illegible]7-tlr5_h720p.mov	102.06 MB	Jan 24, 2019, 5:10:03 PM

Figure 9: Repository operations

The number of options depends on the selected object (file, directory, or folder) and the permissions you have on that object.

Add Folder

1. In the Web Interface, click the **Repositories** tab.
2. Click on a repository.
3. Explore the repository to highlight the object and click the **Add folder** button.
4. Type a folder name.
5. Confirm with **Enter**.

Rename Object

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repository to highlight the object and click the **Rename** button.
4. Type a new name for the object.
5. Confirm with **Enter**.

Delete Object

1. In the Web Interface, click the **Repositories** tab.

2. Select a repository.
3. Explore the repository to highlight the object and click the **Delete** 🗑 button.
4. Confirm to delete the object.

Permissions on the Repositories

The permissions enable you to grant or deny permissions to individual users, user groups, or overall groups to perform actions on a repository.

To access to the permissions, click the **Repositories** tab, in the repositories list, click the ⋮ button of one of them and select **Permissions**. Here you have access to all the permissions created.

To create a new permission on a repository

1. Click the button **+ NEW PERMISSION**. A pop-up appears.
2. Click on the drop down list and select a user or a users group. Click **NEW PERMISSION** to validate.
3. Click **Save changes**.
   The new permission is created, you can now set it.

The interface is divided in two parts:

- **Users and Groups**:
  Lists all the users and groups for which permissions has been created. Select any of them for which you want to set permissions.
- **Permissions**:
  Shows all the available permissions for the corresponding user or users group.

The permissions that you can set, depending on the user or users group, are the following:

- Open
- Archive
- Retrieve
- Add a folder
- Rename a folder
- Delete a folder
- Move a folder
- Rename an object
- Delete an object
- Move an object
- Modify file extension
- Manage metadata
- Move a folder to another repository
- Administration
- Validate retention jobs

To do so, select each time either **Inherit**, **Deny** or **Allow**. Or select one of those options next to the first line **Apply to all**.

> **Note**: The denial of a permission at any level, takes precedence over acceptance.
> If you select **Inherit**, the values will be those previously set or set by default in the settings.

Manage Metadata

Metadata are information about data that is applied to objects during the archiving process or that has been added manually. You can directly view and manage associated metadata from the contextual menu of your repository.

To view and manage metadata

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repository to highlight an object and select ⋮ > **Manage metadata**. The associated metadata are displayed (Figure 10).

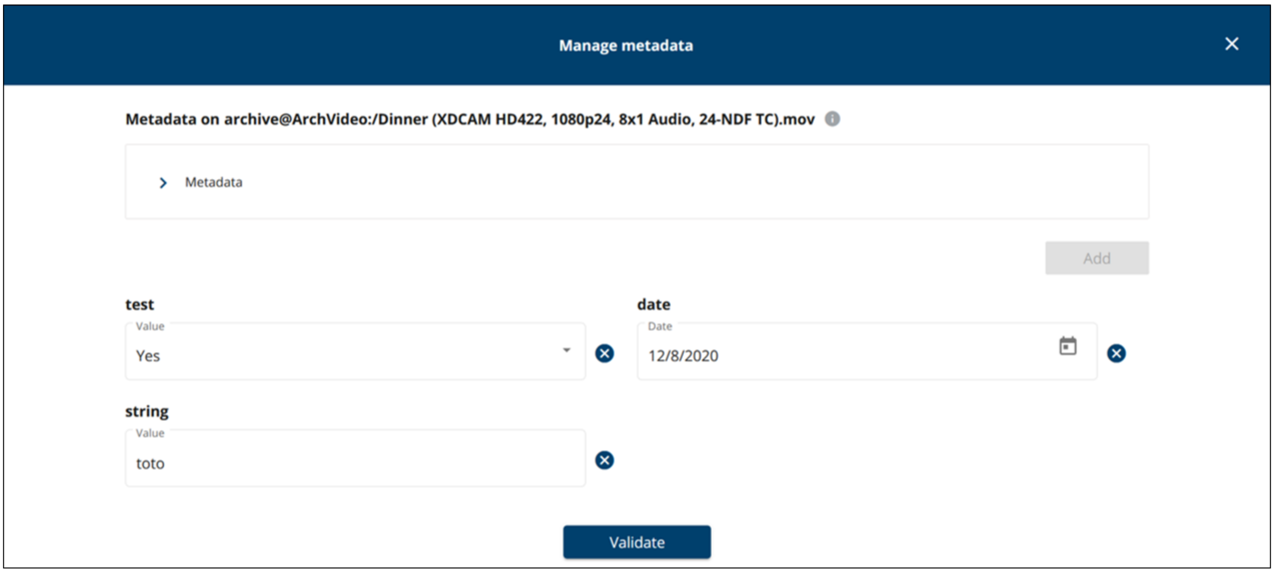

Manage metadata

Metadata on archive@ArchVideo:/Dinner (XDCAM HD422, 1080p24, 8x1 Audio, 24-NDF TC).mov

test — Value: Yes

date — Date: 12/8/2020

string — Value: toto

Figure 10: In this example, 3 metadata types are associated to the object

4. If needed, modify or remove the associated metadata.
5. If needed, select other metadata (Figure 11):
   a. Expand the tree and select a metadata.
   b. Click **Add**. Repeat this for each metadata you want to include.

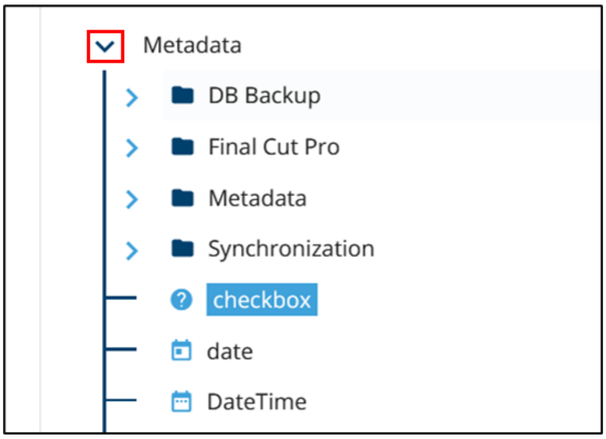

Figure 11: Metadata types

6. Click **Validate** to save the modifications.

For more details about how to manage metadata, see the section Metadata.

Applying Metadata to Repositories

This discussion assumes that you have already created the repositories.

Once you have created the range of metadata that will be available for use within the Miria instance, you can assign these metadata a value and associate them with repositories, repositories folders, objects, or instances.

You can create associations between metadata and repositories manually, using either of these methods:

- Associate metadata with objects when they are selected.
- Associate metadata with repositories, or with objects and instances that have already been archived.

To define setting of metadata values on objects for archiving

1. Click the **Easy Move** tab.
2. Select a platform or a repository to archive, and a target.
3. Click the **+ADD** button.
4. Click **Validate the basket**.
5. Click **Yes** to validate your basket. A new pop-up appears, enter the values for the metadata that are set as **Mandatory**. You have to complete the Metadata that were set as mandatory. You have to do so every time you perform an archiving.
   See Metadata to know how to set a metadata on mandatory.

To apply metadata to repositories and objects

1. Select the **Repositories** tab.
2. Click on one of the repositories to open it.
3. Select the repository, or one of the folders or objects in it by clicking the ⋮ button.
4. Click **Manage metadata**.

5. Select a metadata and click **ADD**. You can select as many metadata as you want.
6. Select a value for each metadata.
7. Click **Validate** to add them to the repository or to the object(s) selected.

To apply metadata to a specific instance of an object

1. Select the **Repositories** tab.
2. Click on one of the repositories to open it.
3. Click the ⋮ button of one of the objects and select **Instance**.
   The list of instances is displayed, ordering by date the past versions of the object.
4. Click an instance ▦ button **Manage metadata**. The manage metadata pane opens. It enables you to view and manage existing metadata on this specific instance, but also to add new metadata.
5. Select metadata in the arborescence and click **ADD** to add a new one.
6. Set a value for each metadata that you add, or change the value of the existing metadata if needed.
7. Click **Validate** to finish the procedure.

Manage Instances

Objects can be archived several times under the same name and the same location. These versions are called instances. You can view details about each instance archived in the Miria Web Interface.

To obtain information on archived objects and storage details

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repositories to highlight an object and select ⋮ > **Instance**. The following information is displayed:
   - **Archiving date/Name** Specifies the instance archiving date and time and the storage manager associated with the instance.
   - **Type** Identifies the archived object.
   - **Date of deletion** Indicates the deletion date from the file system.
   - The Instance Details are displayed on the right.
4. Expand the archiving instance to see the storage that hosts the archived files and directories. You can click the storage to access the Session Details (Figure 12).

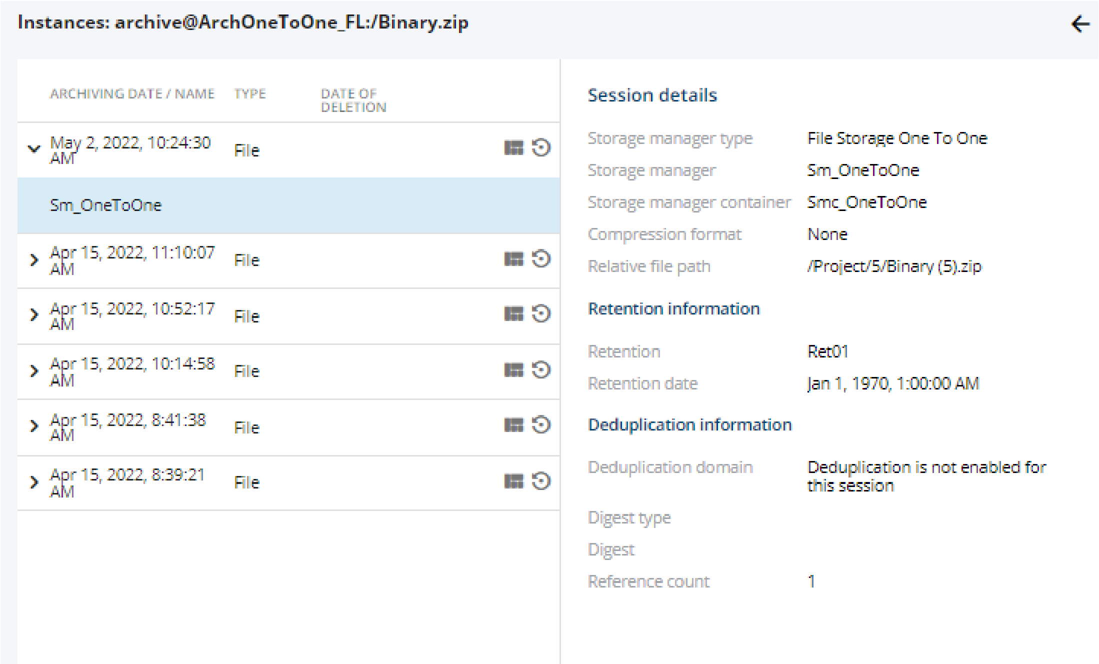

Instances: archive@ArchOneToOne_FL:/Binary.zip

ARCHIVING DATE / NAME	TYPE	DATE OF DELETION
May 2, 2022, 10:24:30 AM	File	
Sm_OneToOne		
Apr 15, 2022, 11:10:07 AM	File	
Apr 15, 2022, 10:52:17 AM	File	
Apr 15, 2022, 10:14:58 AM	File	
Apr 15, 2022, 8:41:38 AM	File	
Apr 15, 2022, 8:39:21 AM	File	

Session details

Storage manager type	File Storage One To One
Storage manager	Sm_OneToOne
Storage manager container	Smc_OneToOne
Compression format	None
Relative file path	/Project/5/Binary (5).zip

Retention information

Retention	Ret01
Retention date	Jan 1, 1970, 1:00:00 AM

Deduplication information

Deduplication domain	Deduplication is not enabled for this session
Digest type	
Digest	
Reference count	1

Figure 12: The storage manager is listed below the instance

The **Instances** interface lets you:

- Perform a retrieve in the **Easy Move** interface by clicking the **Retrieve** button.
- Preview an archived video asset file and perform a partial retrieve.
- Manage and view metadata.
- Return to the repositories by clicking the arrow in the top right corner.

Retrieve Instance

The easiest way to retrieve an archived file is to retrieve its latest instance. However, you can also retrieve any file instance by selecting it from a list of instances.

To retrieve an instance

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repository, choose the object you want to retrieve and:
   - Select ⋮ > **Retrieve last instance**. The **Easy Move** interface is displayed and indicates the path to retrieve the last instance. Continue with step 7 of this procedure.

   Or

   - Select ⋮ > **Instance**. The **Instances** interface is displayed. You can expand an instance to see more details.
4. Select the archiving instance to retrieve. The corresponding instance details are displayed to the right of the instance.
5. If needed, click the **Manage metadata** button to view or modify metadata.
6. Click the **Retrieve** button.
7. Select whether you want to retrieve at new, original, or default location on the window that just opened (Figure 13).

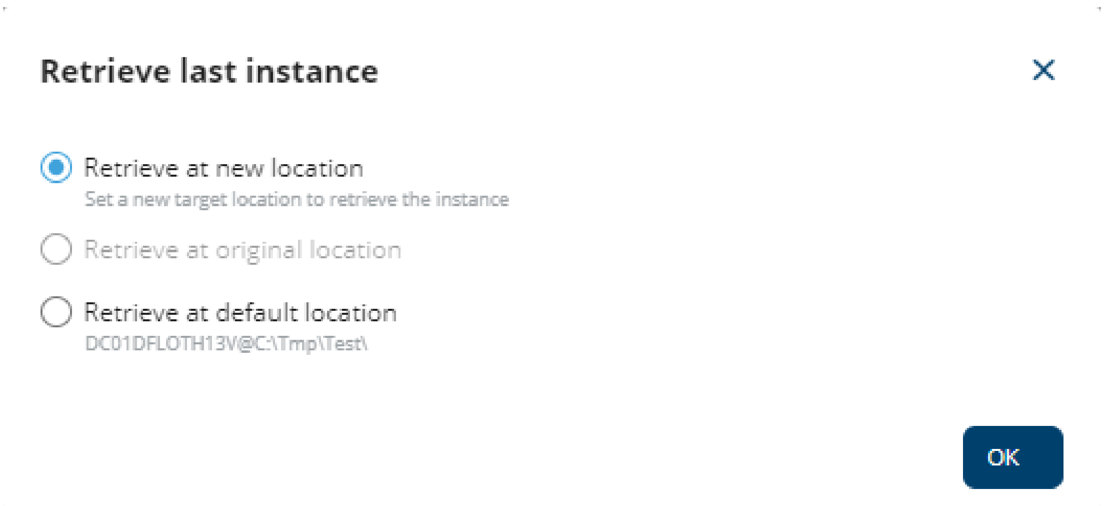

Retrieve last instance

● Retrieve at new location
Set a new target location to retrieve the instance

○ Retrieve at original location

○ Retrieve at default location
DC01DFLOTH13V@C:\Tmp\Test\

Figure 13: Retrieve last instance

8. Select **Retrieve at date** to get to choose a new location.

   The **Easy Move** interface is displayed and indicates the path to retrieve the instance (Figure 14).

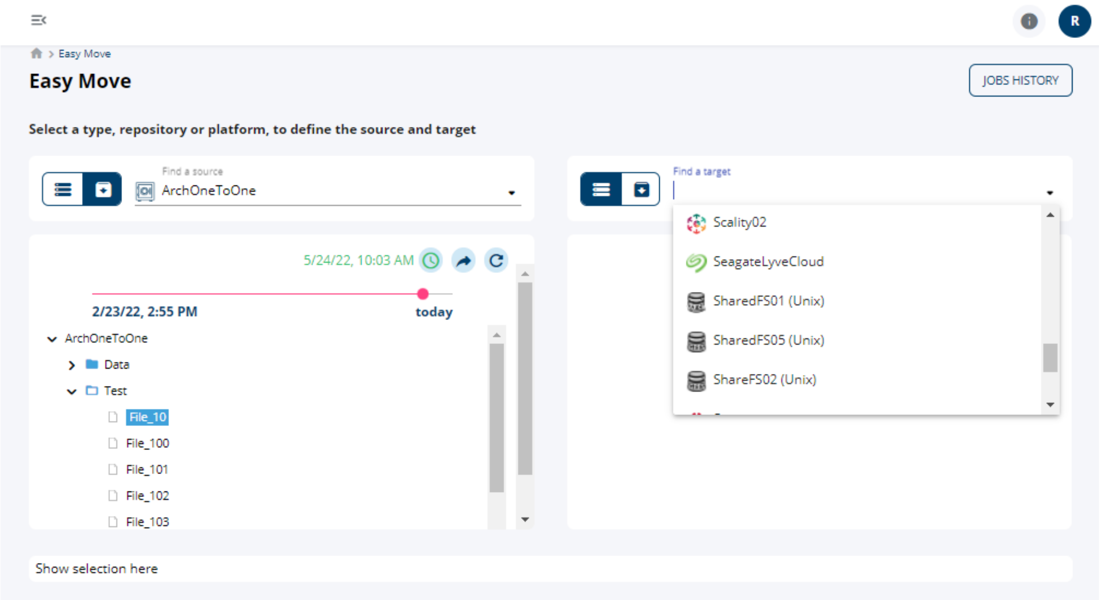

Easy Move

Select a type, repository or platform, to define the source and target

Find a source: ArchOneToOne

5/24/22, 10:03 AM

2/23/22, 2:55 PM — today

- ArchOneToOne
  - Data
  - Test
    - File_10
    - File_100
    - File_101
    - File_102
    - File_103

Find a target:

- Scality02
- SeagateLyveCloud
- SharedFS01 (Unix)
- SharedFS05 (Unix)
- ShareFS02 (Unix)

Show selection here

Figure 14: The selected path is directly indicated in the source repository

9. In the right field, select the target type as **Platform** to choose your target.
10. Move the object:
    - Click **Add**.

    Or

    - Drag the object from the source to the desired location in the target.
11. Validate the basket. The retrieval job is launched.
12. Click the **Jobs** tab to verify the job progress.

Instance Details

Check the following overview (Table 1) for information about the instance details.

Table 1: Instance details properties

Property	Description
Archiving date	Date and time the file or directory was sent to the archive
MIME type	Specifies the MIME type of the archived file which is a two-part identifier of file formats, for example: `audio/mpeg`, `video/quicktime`, or `text/plain`. It is a more reliable indicator of file format than the file extension. This property displays only if the administrator has configured Miria to recover it.
File size	Size of the file or directory.
Creation date	Date the file or directory was created, independent of the archive.
Last update	Last time the file or directory was modified, independent of the archive.
Last access	Last time the file or directory was consulted without modification, independent of the archive.
Owner user	User who owns the file or directory.
Owner group	Group of the user who owns the file or directory.
Original location	Original location of the file on the user’s machine. If the file was archived from a network drive, this displays the path from the network drive (e.g., `X:\Data\MyData`).
Global path	Global path of the file original location.
File path on source	File original location on the primary storage.
Digest type during archiving	Type of digest that is calculated on the file at the moment of archiving on the storage when performing multiple writing. It verifies whether the file content is the same on all storage (i.e., `[None]`, `MD5`, `SHA-1`, `SHA-256`, `SHA-384`, or `SHA-512`).
Digest during archiving	Value of the above digest, if this was used.
Link target	Path of the target file/directory to which the link points on the source disk. Displays only if the archived object is a symbolic link.

Table 1: Instance details properties

Property	Description
Alternate stream	Indicates whether the alternate streams (i.e., file attributes, rights, etc.) are archived in this instance. Possible values are **Yes** or **No.**

Session Details

Check the following overview (Table 2) for information about the session details.

Table 2: Session details properties

Property	Description
Storage manager type	Type of storage manager used (e.g., Miria File Storage One To One etc.).
Storage manager	Name of the storage manager as configured in Miria.
Storage manager container	Name of the storage manager container as configured in Miria.
Compression format	See the Miria Administrator's Documentation for details on compression.
Relative file path	Defines the location that is relative to the current directory or folder.
Retention	Name of the retention period associated with the repository where the file is stored.
Retention date	Date and time of the end of retention period.
Deduplication domain	Name of the deduplication domain as configured in Miria.
Digest type	Type of the deduplication digest that is calculated at archiving time to verify that the new file to be archived is unique (`SHA-1`, `SHA-256`, `SHA-384`, `SHA-512`, or `Filename/Size`).
Digest	Value of the deduplication digest, if this was used.
Reference count	Number of times a file with the same digest has been archived.

CHAPTER 6 - Search for Repositories

If you have a great number of objects, you may find it difficult to locate a folder or file that you want to retrieve. You can browse a project folder or click directly on a repository to explore its content, presenting as a file tree. Or, you can use the search feature that lets you include criteria or metadata to locate your object.

This topic explains how to search for objects in the Web Interface, using the file tree, search criteria and metadata.

Browse the File Tree

If you know where the objects you are looking for are located, you can navigate directly to the file tree:

1. Open the repository in which your object is located.
2. Click the different folders to access them.
3. Select your object.

Use Go To Path

Even when a repository contains a great number of folders, sub-folders, and objects, you can quickly reach a specific object if you know its path in the repository.

1. Click the **Repositories** tab.
2. Select a repository.
3. Click the **Go to path** ➦ button.
4. Enter the absolute path of the object in the archive (e.g., `archive@ArchOneToOne:/Test/test-tlr5_h720p.mov`)
5. Select **Go to path**.

Search by Criteria

Search criteria are based on data contents and file properties, such as name, size, and date of creation (Figure 15).

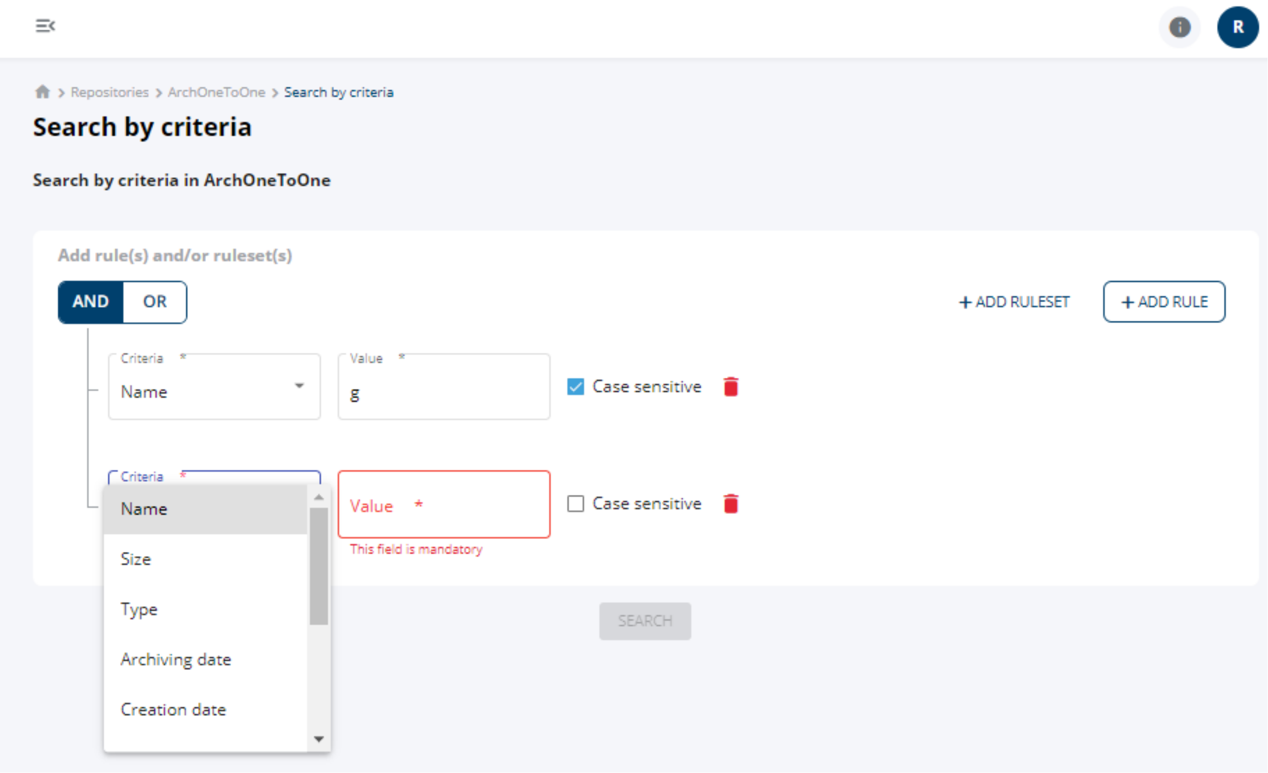

Search by criteria

Search by criteria in ArchOneToOne

Add rule(s) and/or ruleset(s)

Criteria *: Name | Value *: g | ☑ Case sensitive

Criteria * | Value * | □ Case sensitive

This field is mandatory

Figure 15: Different search criteria selected

You can perform a search by criteria within:

- An entire repository.
- An archived folder.
- An archived directory.

To run a search by criteria

1. Click the **Repositories** tab.
2. Select a repository or a project.
3. In the **Repositories** interface, you can either:
   - In the upper right corner, click the**Search by criteria** button to search the entire repository.

   Or

   - In the **Explore the repository** section, highlight a folder or directory and select ⋮ > **Search by criteria**.
4. Click **+Add rule** to add a new criteria, or **+Add ruleset** to add a group of criterias.
5. In the field **criteria**, select from the drop down menu the kind of criteria you want to search and complete the field **Value**. Here is the list of the different types of criteria:
   - **Name** Enter the name of the object. You can use `*` and `?` as wildcards and choose your search to be case-sensitive. For example:
     – If you want to find the `mydocument.doc` file, you can enter either the string **document** or ***document***. In both cases, the result set contains objects such as `mydocument.doc`, `hisdocument2.doc` or `documentA`.

   - If you enter the string **document?**, the result set contains only objects whose names contain the term `document`, plus one character (e.g., `documentA` or `document`").
   - If you enter the string **document***, the result set contains only objects whose names start with `document`, plus any number of characters (e.g., `document`, `document5.doc`, or `documentation.doc`), but not files such as `mydocument.doc`.
   - If you enter the string **document?.***, the result set contains objects whose names contain the term `document`, plus one character (i.e., the dot), plus any number of characters (e.g., `document5.doc`, `documentA.doc`, or `document5.pdf`), but not files such as `documentation.doc`.

- **Size** Choose the size of the object.
- **Type** Select object types.
- **Archiving date** Specify the date if you know when the object was last archived.
- **Creation date** Specify the date if you know when the object was created.
- **Access date** Specify the last access date if you know it.
- **Modification date** Specify the date if you know when the object was modified.
- **User owner** Enter the name of the User owner.
- **User ID** Choose a value corresponding to the User ID.
- **Group ID** Choose a value corresponding to the Group UD.

6. Click **Search**. The results are displayed below (Figure 16).
   In the first column, you have the objects, with their instances. In the two other columns, you have information about the file size and the archiving date of each instance.

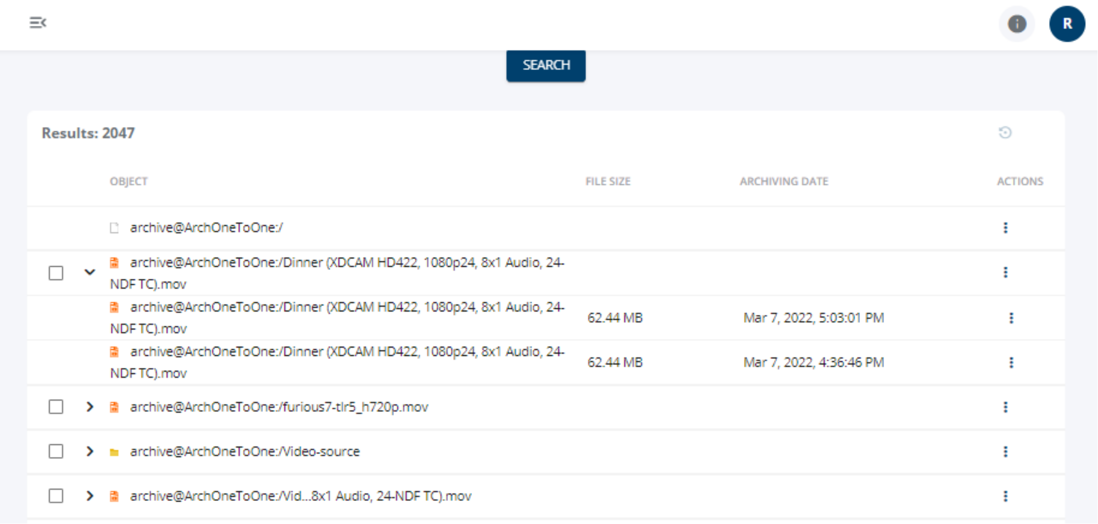

Results: 2047

		OBJECT	FILE SIZE	ARCHIVING DATE	ACTIONS
		archive@ArchOneToOne:/			⋮
□	⌄	archive@ArchOneToOne:/Dinner (XDCAM HD422, 1080p24, 8x1 Audio, 24-NDF TC).mov			⋮
		archive@ArchOneToOne:/Dinner (XDCAM HD422, 1080p24, 8x1 Audio, 24-NDF TC).mov	62.44 MB	Mar 7, 2022, 5:03:01 PM	⋮
		archive@ArchOneToOne:/Dinner (XDCAM HD422, 1080p24, 8x1 Audio, 24-NDF TC).mov	62.44 MB	Mar 7, 2022, 4:36:46 PM	⋮
□	>	archive@ArchOneToOne:/furious7-tlr5_h720p.mov			⋮
□	>	archive@ArchOneToOne:/Video-source			⋮
□	>	archive@ArchOneToOne:/Vid...8x1 Audio, 24-NDF TC).mov			⋮

Figure 16: Results

This results list offers you several options, you can:

- **Retrieve an object**
  Select an object from the result list and click **Retrieve last instance**. See Retrieve Instance for more details.
- **See an object location**
  Select an object from the result list and click **See in repository details**.

- **Perform multiple retrieval**
  Select multiple rows of objects and click on the button **Multiple retrievals** that just got activated on the upper right side of your screen.
- **Retrieve at date**
  Select an instance, and click **Retrieve at date**. The**Easy Move** tab opens at the exact date of the instance (Figure 17).

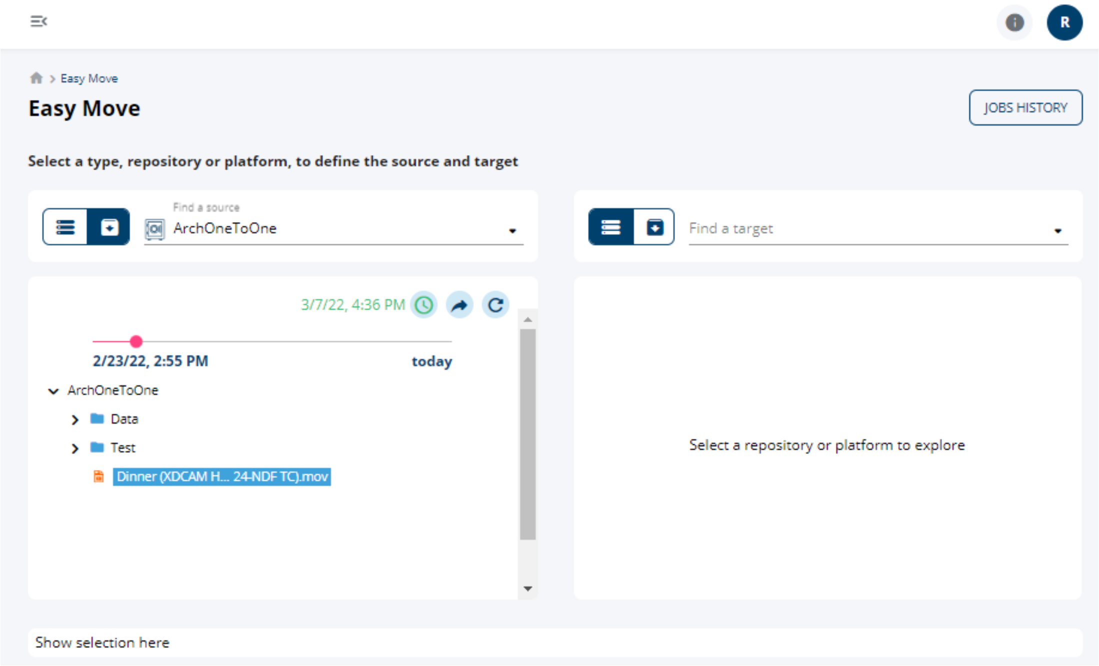
Easy Move
Easy Move
JOBS HISTORY
Select a type, repository or platform, to define the source and target
Find a source
ArchOneToOne
Find a target
3/7/22, 4:36 PM
2/23/22, 2:55 PM
today
ArchOneToOne
Data
Test
Dinner (XDCAM H... 24-NDF TC).mov
Select a repository or platform to explore
Show selection here

Figure 17: Easy Move tab

Search by Metadata

If the repositories you are allowed to access have metadata associated with them, you can use that metadata as keywords to find archived files and directories. Criteria used in the most recent search operations are remembered.

To run a search by metadata

1. Click the **Repositories** tab.
2. In the upper right corner of the **Repositories** interface, click the **Search by Metadata** button.
3. Navigate in the metadata tree and select a metadata.
4. Click **Add**. Repeat this for each metadata you want to include (Figure 18).

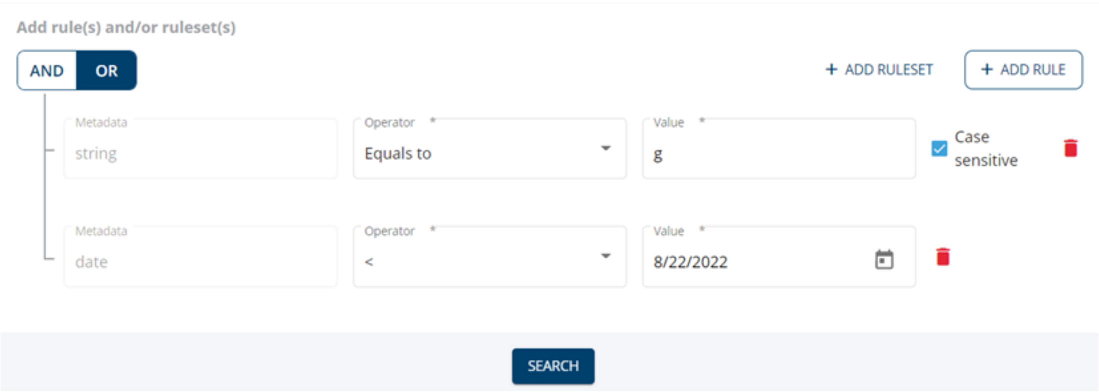

Add rule(s) and/or ruleset(s)

Metadata: string | Operator *: Equals to | Value *: g | ☑ Case sensitive

Metadata: date | Operator *: < | Value *: 8/22/2022

Figure 18: Multiple metadata selected

5. Define the metadata values. For some metadata (e.g. **String**), you can activate case sensitivity besides the value.
6. Choose a search operator:
   - **OR** The result includes all the objects meeting at least one of the metadata criteria. This search operator is selected by default.

   Or

   - **AND** The result includes the objects meeting all metadata.
7. Click the button **Add ruleset** to add a group of criterias or metadatas on which can be applied a new operator **and**, or **or**(Figure 19).

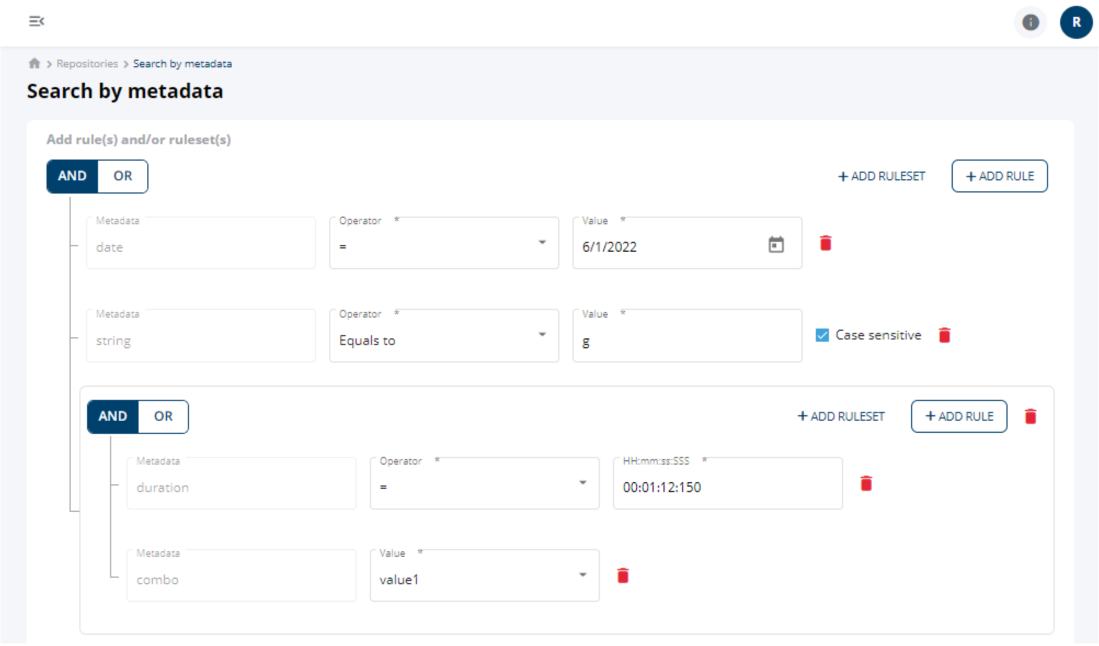

Repositories > Search by metadata

Search by metadata

Add rule(s) and/or ruleset(s)

Metadata: date | Operator *: = | Value *: 6/1/2022

Metadata: string | Operator *: Equals to | Value *: g | ☑ Case sensitive

Metadata: duration | Operator *: = | HH:mm:ss:SSS *: 00:01:12:150

Metadata: combo | Value *: value1

Figure 19: Add a ruleset

8. Click **Search**. The list displays a maximum of 1 000 results categorized as follows:
   - **Objects** Displays objects for metadata applied to already archived objects.
   - **Jobs** Displays archiving jobs for metadata applied at archiving time.
   - **Instances** Displays instances for metadata applied at archiving time.
   - **Repositories** Displays repositories for metadata applied to the repository.
9. From the **Results** list, click a category to expand.

Depending on the category, you can perform different operations:

- Click the **Manage metadata** ▦ button to view and manage metadata.
- Click the **Retrieve**⟲ button to retrieve data in the **Easy Move** interface. Multiple objects can be selected from a single repository only.
- Click the **Go to path** ↪ button to view details on repositories, jobs, and instances.

CHAPTER 7 - Manage Video Assets

Miria is equipped with the following features specifically designed for managing video assets:

- **Video previews** Check video content in the player embedded in the Web Interface.
- **Partial retrieve** Retrieve only a sequence of a much larger media file.

These features require a specific license key to be activated. Contact your administrator for more information. The files must have been archived with the *Video proxy transcoding format* and the *Video proxy location* advanced settings enabled. These advanced settings must be set by the administrator.

Preview Video Assets

Prior to retrieving an archived video asset, you can do a preview to check the content. Previews are a low-resolution version of the video asset. During the archiving process, assets are generated using video content without the sound.

Play Video Asset Preview

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repository to highlight the video file and select ⋮ > **Instance**. The **Instances** interface is displayed.
4. Select the instance from the list and click the **Preview** ▶ button.

Video Player Overview

The following overview (Table 3) describes the video player fields and buttons:

Table 3: Video player overview

Feature	Description
Container	Format of the video file.
Codec	Codec used to generate the low-resolution version of the file.
▶	Plays or pauses the video.
] [	Sets the start or the end of the video sequence that you want to retrieve.
I< >I	Enables to go to the previous or next video frame.

Feature	Description
x1	Enables to set speed control of the video.
⛶	Enables to visualize in full screen.
Current time	Specifies the time elapsed up to the millisecond.
Duration	Specifies the total duration of the video up to the millisecond.
SMPTE	Identifies as the timecode standard for labeling frames of video or film. SMPTE timecode format is *`hour:minute:second:frame`* (e.g. one hour becomes *`01:00:00:00`*). The frame rate signifies the frames per second that appear in the media. This rate corresponds to the frame number, shown in the final part of the timecode.
Mark IN	Indicates the start of the video sequence that you want to retrieve.
Mark OUT	Indicates the end of the video sequence that you want to retrieve.
IN - OUT	Displays the duration of the video sequence.
Same as input	Lets you set the start time code of the retrieved file identical to that of the original, archived file.

Partially Retrieve Video Asset

The embedded player lets you partially retrieve video assets of `.mxf` or `.mov` formats, based on a time code. This is useful in case you only need a portion of the media file.

Step 1: Define the sequence to retrieve

1. In the Web Interface, click the **Repositories** tab.
2. Select a repository.
3. Explore the repository to highlight the video file and select ⋮ **> Instance**. The **Instances** interface is displayed.
4. Select the instance to preview and select the **Preview** ▷ button.
5. Play the video up to the first frame of the section that you want to retrieve and click the **Mark IN** button. The **Mark IN** field displays the start time code.
6. Play the video up to the last frame of the sequence that you want to retrieve and click the **Mark OUT** button. The **Mark OUT** field displays the end time code.
7. To the right of the **IN - OUT** field, click the **Play** ▶ button to check the sequence to retrieve. You can use the time navigation bar to cut the sequence to the desired length (Figure 20).

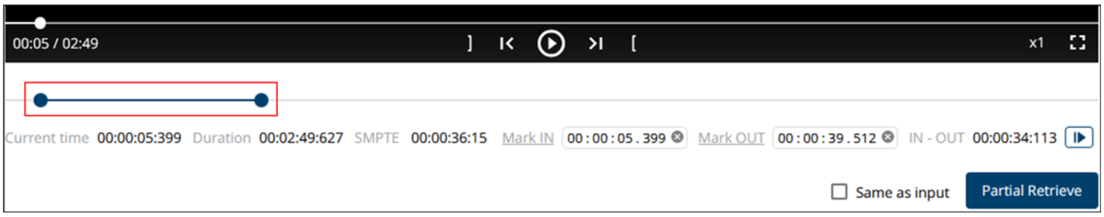

Current time 00:00:05:399 Duration 00:02:49:627 SMPTE 00:00:36:15 Mark IN 00 : 00 : 05 . 399 Mark OUT 00 : 00 : 39 . 512 IN - OUT 00:00:34:113

□ Same as input

Figure 20: Example of a sequence to retrieve

8. If needed, activate **Same as input** to set the start time code of the retrieved file to that of the original archived file.
9. Click the **Partial retrieve** button. The **Easy Move** interface is displayed (Figure 21).

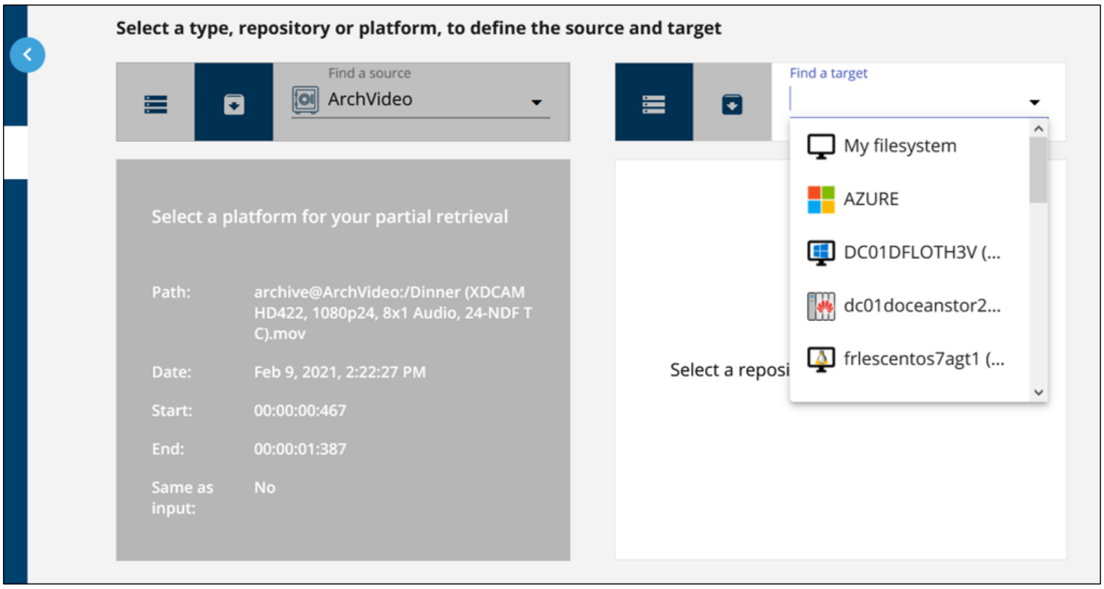

Select a platform for your partial retrieval

Path: archive@ArchVideo:/Dinner (XDCAM HD422, 1080p24, 8x1 Audio, 24-NDF TC).mov

Date: Feb 9, 2021, 2:22:27 PM

Start: 00:00:00:467

End: 00:00:01:387

Same as input: No

Figure 21: Partial retrieve in the Web Interface

Step 2: Launch the job

1. In the right field, select your target from the list.
2. Browse the target directory and click **Add**.
3. Validate the basket. The retrieval job is launched. The selected sequence is retrieved in the appropriate directory.
4. Click the **Activity** tab to verify the job progress.

CHAPTER 8 - Monitor Data

When an action (archiving, retrieval etc.) is launched, Miria creates a job and assigns it a job number. You can then monitor the progress on each job (Figure 22).

> In the Web Interface, click the **Jobs** tab.

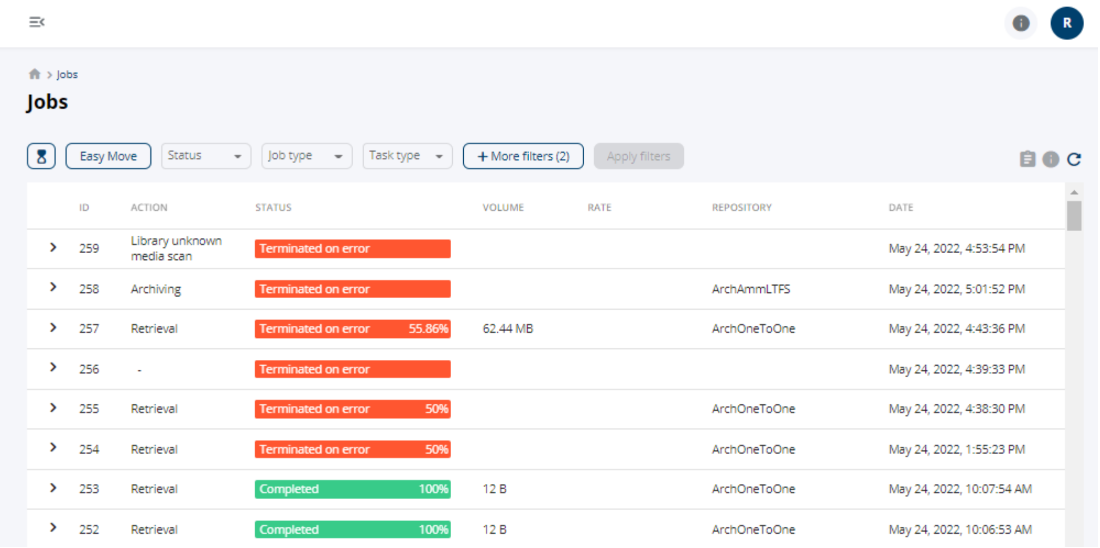

	ID	ACTION	STATUS	VOLUME	RATE	REPOSITORY	DATE
›	259	Library unknown media scan	Terminated on error				May 24, 2022, 4:53:54 PM
›	258	Archiving	Terminated on error			ArchAmmLTFS	May 24, 2022, 5:01:52 PM
›	257	Retrieval	Terminated on error 55.86%	62.44 MB		ArchOneToOne	May 24, 2022, 4:43:36 PM
›	256	-	Terminated on error				May 24, 2022, 4:39:33 PM
›	255	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 4:38:30 PM
›	254	Retrieval	Terminated on error 50%			ArchOneToOne	May 24, 2022, 1:55:23 PM
›	253	Retrieval	Completed 100%	12 B		ArchOneToOne	May 24, 2022, 10:07:54 AM
›	252	Retrieval	Completed 100%	12 B		ArchOneToOne	May 24, 2022, 10:06:53 AM

Figure 22: List of jobs and sub-jobs

The **Jobs** interface enables to:

- Expand a job to view its sub-jobs.
- View all jobs and corresponding status:
  - **Creation in progress** The request has been validated and the corresponding job is being created.
  - **In queue** Miria is waiting for a prior operation to complete. If the status does not change, contact your administrator and quote the job number for reference.
  - **Running** The operation is in progress.
  - **Suspended** The operation is suspended.
  - **Canceled** The administrator has refused the operation.
  - **Denied** The operation was denied. Contact your administrator and quote the job number for reference.
  - **Completed** The operation was successfully completed.
  - **Terminated on error** The operation terminated because at least one file in the job could not be processed. Contact your administrator and quote the job number for reference.
  - **Rejected** The operation was rejected.
  - **Invalid License** The license is invalid. Contact your administrator.
  - **Terminated on warning** The operation terminated because at least one file in the job issued a warning (non-existent file, modified files, file directory change etc.).
- Access Easy Move jobs history by clicking the **Easy Move** filter.

Show a Job Menu

1. In the Web Interface, click the **Jobs** tab.
2. Click the ⋮ button on a job or sub-job row in the list.

   The following options appears:

- For all jobs or sub jobs:
  - **Show details.** Enables to view general information and volume details.
  - **Open logs.**Opens logs file, that contain an history of any change that happened on a job.
  - **Cancel.** Allows you to cancel a job running. The job is then displayed in the list of jobs with a **Canceled** status.
- For all sub jobs:
  - **Settings.** Gives you access to the advanced settings, see Default Settings for further details.
- For Archiving jobs:
  - **View metadata.** Opens a window that lists the metadata associated with the job.
- For different jobs types, like task, drive unmount, device scan, media scratch :
  - **Download report.** Generates and download a report of the job data.
- For Easy Move jobs:
  - **Relaunch.** To relaunch a job with the same configuration.

CHAPTER 9 - Manage the Infrastructure

The infrastructure may consist of different storage types (e.g. agents, local file system, shared file systems) that have been configured as repository platforms (Figure 23).

> To access the infrastructure information in the Web Interface, click the **Infrastructure** tab.

The following tiles are available:

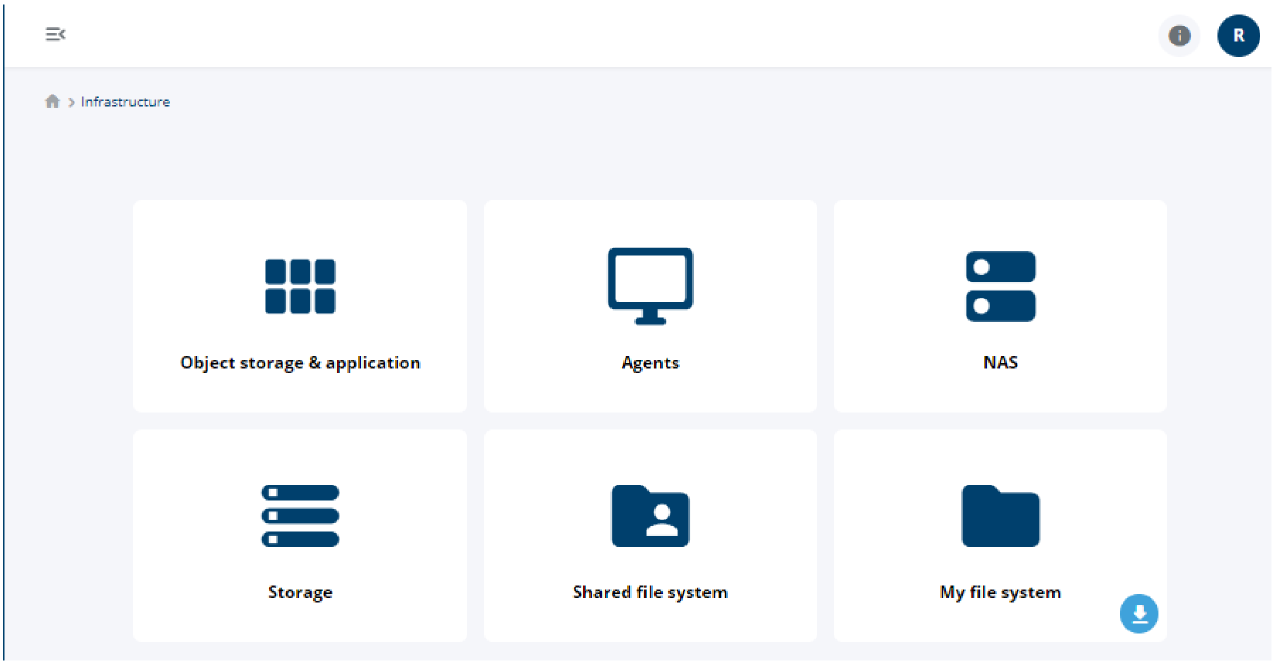

Figure 23: Storage types sorted by list in the **Infrastructure** interface

- **Object storage & Application** View storage manager and storage manager container
- **Agent** Access, explore and verify the status of your own agent.
- **NAS** View NAS.
- **Storage** View storage.
- **Shared file system** View shared file system.
- **My file system** Activate and access My file system to perform data-move operation from and to the local file system.

Add a Storage Manager and Container

To configure the infrastructure, the Object storage & Application entry is used to create storage managers and containers.

A storage manager is the storage definition, which can be tape, cloud (e.g., AWS, Google), disk and object storage. It manages the data migration from a primary storage to a secondary storage. Once you have created a storage manager, you must create a storage manager container. The storage manager container defines the location where the data are archived within the storage manager.

This chapter outlines how to add a storage manager and container for File Storage variants and, as an example, Amazon S3, Google Cloud Storage and Microsoft Azure Blob Block. For other third-party storage managers, see Third Party Storage Managers in Miria Administration documentation.

File Storage Container

When you add a File Storage Container, data are organized by job. Each job corresponds to only one file (container) on the destination file system.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **File Storage Container** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
6. Select a storage platform. This is the name of the destination machine that hosts the files. This machine must be declared as a platform in Miria.
7. Activate **UTF8 Support** if you want to support UTF8 character encoding.
8. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
9. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for File Storage and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
3. Select the directory where the data are archived on the storage manager. This is an absolute path.
4. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.

The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.

5. Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
6. Activate **Alternative path** if you want to archive the data on other machines on the network. You must then declare the paths of the mounting points as alternative paths. Click **Add new path** and complete the following parameters:
   a. **Agent** Name of the Miria platform where the data to archive is located. Select the agent from the list.
   b. **Path** Absolute path of the directory where you want your data to be archived on the platform specified in the Agent field.
   c. **User and password** Credentials of a user that has access permissions to this path.
   d. **Enable On/Off** Disable temporarily the alternative path for an agent.
7. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
8. Click **Create** to add the storage manager container.

File Storage One to One

When you add a File Storage One to One, the archived data are organized in a file tree structure in the same way as the data on the source file system. One file/directory on the source file system corresponds to the same file/directory on the destination file system. Data can be accessed outside of Miria.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **File Storage One to One** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
6. Select a storage platform. This is the name of the destination machine that hosts the archived files. This machine must be declared as a platform in Miria.
7. Activate **UTF8 Support** if you want to support UTF8 character encoding.
8. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
9. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for File Storage and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads.** Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
3. Select the directory where the data are archived on the storage manager. This is an absolute path.
4. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.

   The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.
5. Set the data compression value. This specifies whether the archived data must be compressed in the storage and defines the compression type:
   - **None** No compression.
   - **ADAZip** Optimized internal compression format. Files compressed with this format have a .adazip extension.
6. **Immutable disk repository** This option makes the files immutable. They cannot be modified, deleted, or renamed. No link can be created to these files.
   For more details, regarding immutability flags, please have a look at the XFS CTL Linux Man Page : 3-xfsctl
   To get to know about prerequisites for this option, please refer to the Installation guide, Chapter 1 Preparing to install, in the part Immutable Disk Repository.

> **Important**: As immutability is only supported on Linux XFS and Ext3/4, the option will be grayed if the One to One is on Windows OS.

7. If needed, set volume management. This option enables to define a quantity of disk space that is always kept free on the destination volume (e.g., to permit sharing this volume with other applications). You can define its value either as a percentage of the disk space or as a number of GB. By default, archiving to a File Storage container uses all available disk space on the target volume.
8. Activate **Alternative path** if you want to archive the data on other machines on the network. You must then declare the paths of the mounting points as alternative paths. Click **Add new path** and complete the following parameters:
   a. **Agent** Name of the Miria platform where the data to archive is located. Select the agent from the list.
   b. **Path** Absolute path of the directory where you want your data to be archived on the platform specified in the Agent field.
   c. **User and password** Credentials of a user that has access permissions to this path.
   d. **Enable On/Off** Disable temporarily the alternative path for an agent.

9. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
10. Click **Create** to add the storage manager container.

SnapStor

When you add a SnapStor storage manager, the archived data are organized in a file tree structure in the same way as the data on the source file system *and* located in a snapshot created after the archiving task. An instance of a file/directory in the archive corresponds to the same file/directory in a snapshot. A file or a directory can be retrieved individually.

> **Note**: The SnapStor storage manager only supports Windows agents and Qumulo, Dell Isilon, Huawei OceanStor or GPFS Shared File System as storage platforms.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **SnapStor** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
    - **Online** Default value if you want to perform an archiving.

    Or

    - **Suspended** This status is useful for maintenance operations.
6. Select a storage platform. This is the name of the destination machine that hosts the archived files. This machine must be declared as a platform in Miria.
7. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Snapstor and click the + button to add a container.
2. Complete the properties of the storage manager container:
    a. **Storage container name**.
    b. **Deduplication domain**. A new domain can be created, by clicking the + button.
    c. **Archiving run lock**. A new one can be created, by clicking the +button.
    d. **Threads.** Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
3. Select the directory where the data are archived on the storage manager. This is an absolute path.
4. Configure stream options.
5. Set export share options to be able to remotely access the archived data on a SnapStor storage through several SMB shares and /or NFS exports. When used in an archiving or backup task, SMB shares or NFS exports are created on the storage if they exist on the host source. Permissions applied on the created SMB shares or NFS exports are the same permissions than those used by the SMB shares or NFS exports on the task source host.
6. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).

7. Click **Create** to add the storage manager container.

Virtual Storage

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Virtual Storage** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
6. Select the operating mode:
   - **Load balancing** The archived data are distributed among several storage manager containers to achieve better performance. When enabled, the storage manager container used for archiving is a logical container composed of several physical containers.

   Or

   - **Failover** The archived data are sent to the backup storage manager containers if the primary container fails. When enabled, the storage manager container used for archiving is a logical container composed of several physical containers.
7. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for File Storage and click the ⊕ button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
3. Click **Add container** to select a storage manager container from the list. Then click **Add**.
4. Click **Create** to add the storage manager container.

Media Storage

To add a storage manager for a media (Media Manager or Optical Disk Archive), you must create the application in Miria Administration console. See also Creating a Media Manager Application in the Administration documentation.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select:
   - **Media Manager**.

   Or

- **Optical Disk Archive**.

4. Click **Next**.
5. Enter the name of the storage manager.
6. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
7. Select the application to be linked to this storage manager.
8. Select a user or group that will be notified by email when a response to a media request is needed. This can be the case when an archiving job requires a scratch media, or a retrieval job requires media that are offline or in prevent use mode.
9. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
10. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Media Manager and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**.*(Media Manager only)* Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
3. Select a library in which Miria stores the archived data.
4. Select a media type only if the library may contain media of several types (e.g., LTO-6 and LTO-7). You can check **Show only WORM media** to display only the WORM media and choose one.
5. Select a scratch media group in which the media needed for archiving is selected. One media group (default) is created automatically and selected by default. This parameter is mandatory if you do not specify a library.
6. Select a barcode. By default, the first blank or scratch media available in the scratch media group is selected. Use the wildcard characters *, ?, and | to compare the barcodes of the media.

   Examples:

   - A005?? selects any media with a six character barcode beginning with the string A005. You might use this, for example, to select media from A00500 to A00599.
   - *L4 selects any media with a barcode ending in L4. You might use this, for example, to select only media of LTO4 type.
   - 162*|WV* selects any media with barcode beginning either with the string 162, or with the string WV.

7. Set a media rule.
8. *(Media Manager only)* Specify the media format.
9. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.

   The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.
10. Set other configuration options:
    - **Metadata** The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
    - **Prevent spanning** Prevent an archived file from being written on a media if it is too large to fit the remaining space.
    - **Custom block size** By default, the media is subdivided into blocks of 128 KB. The block size must be a multiple of 16. The maximum size accepted is 4 MB.
    - **Log level** Set the level of the logs that are displayed.
    - *Also for Optical disk archive* Choose whether to set Pack write and the format of the media.
11. *(Media Manager only)* Set LTFS delivery protocols. If you do not have to comply with specific protocols, do not modify the default parameters.
12. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
13. Click **Create** to add the storage manager container.

Easy move Amazon S3

Amazon Simple Storage Service (S3) is an Internet storage solution designed to make web-scale computing easier for customers.
Details of what is replicated in S3 are described in the appendix: Replications in between two S3 object storages ".

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Amazon S3** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or
   - **Suspended** This status is useful for maintenance operations.
6. Enter the network address of the Amazon S3 storage service (e.g., S3.amazonaws.com).
7. If needed, set a default proxy platform. This platform handles the data movement on behalf of the usual agent or agents pool.
8. Set storage manager options:
   - **HTTP Proxy** Enables the Proxy HTTP to communicate with a remote S3 storage.

   - **Transfer acceleration** Sends data to the nearest Amazon S3 node and acknowledges reception. Then, Amazon S3 sends the data to the actual final destination.

9. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
10. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the Amazon S3 storage manager and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
   e. **Available as source**. If needed, make the storage manager container available as source. Then enter the name you want to give to the platform associated.
3. Enter Amazon S3 account information:
   a. **Access Key ID** String that uniquely identifies the Amazon S3 account.
   b. **Secret Access Key** Password associated with the Access Key ID
   c. **Bucket name** Logical path under which the data are stored into the Amazon S3 storage. Refer to your Amazon S3 storage configuration.
4. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.
5. Set the data compression value. This specifies whether the data must be compressed in the storage and defines the compression type.
6. Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
7. Choose whether to activate MD5 checksum on the S3 archiving transfer.
8. Set the retention mode for object lock. See also Data Immutability with S3 Object Lock in Administration documentation.
   - **Governance mode** Users with specific Identity and Access Management (IAM) permissions can overwrite or delete protected object versions during the retention period.

   Or

   - **Compliance mode** No users can overwrite or delete protected object versions during the retention period. To delete objects that have this configuration, you must close the account that they are associated with.
9. Set lifecycle rules and complete following information accordingly:

   a. **Name** Name of the life cycle rule that defines the data migration. This name is any unique string of your choice (e.g., ada_smc_amazon, RuleForArchiving, 1toglacier_200todelete, etc.). When you launch the first job, Miria uses this name to create a rule on the Amazon S3 bucket.
   b. **Transition days** Number of days at the end of which Amazon S3 will transfer the objects to Glacier or Deep Archive. By default, Amazon S3 performs the transfer at 00:00 the same day. A value of 1 indicates that the transfer is performed at 00:00 the next day, and so on.
   c. **Retrieval mode** Define the retrieval mode between standard, bulk and expedited. See also AWS documentation.
   d. **Copy lifetime** Number of copy lifetime.

10. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
11. Click **Create** to add the storage manager container.

Google Cloud Storage

Google Cloud Storage is an Internet storage solution designed to make web-scale computing easier for customers.

The integration between Miria and the Google Cloud Storage technology enables you to store data into a Google Cloud Storage compatible storage.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Google Cloud Storage** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
6. If needed, set a default proxy platform. This platform handles the data movement on behalf of the usual agent or agents pool.
7. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
8. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Google Cloud Storage and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.

   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
   e. **Available as source**. If needed, make the storage manager container available as source. Then enter the name you want to give to the platform associated.
3. Enter the Google Cloud Storage account information:
   a. **Authentication File** Select the JSON key of your Google Cloud Storage service account. Keep the JSON file carefully, it cannot be retrieved from Miria because it is encrypted. See also Generate a Google Cloud JSON key in Administration documentation.
   b. **Bucket Name** Unique name of the bucket that Miria will create in Google Cloud to store the files.
4. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.
5. Set the data compression value. This specifies whether the archived data must be compressed in the storage and defines the compression type.
6. Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
7. Choose whether to activate MD5 checksum on the S3 archiving transfer.
8. Set lifecycle rules:
   - Click the ⊕ button to add a new rule. Then enter the number of days at the end of which GCS will transfer the objects. By default, GCS performs the transfer at 00:00 the same day. A value of 1 indicates that the transfer is performed at 00:00 the next day, and so on.
   - If no rules are defined, Miria will use the GCS Standard Storage Class.
9. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
10. Click **Create** to add the storage manager container.

Microsoft Azure Blob Block

Microsoft Azure Blob Block is an Internet storage solution designed to make web-scale computing easier for customers.

The integration between Miria and the Microsoft Azure Blob Block technology enables you to store data into a Microsoft Azure Blob Block cloud compatible storage (REST interface).

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Microsoft Azure Blob Block** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.

6. Enter the network address of the Microsoft Azure Blob Block storage service (e.g., AZURE-ACCOUNT.blob.core.windows.net).
7. If needed, set a default proxy platform. This platform handles the data movement on behalf of the usual agent or agents pool.
8. Enable **HTTP Proxy** to be able to communicate with a remote S3 storage.
9. Configure an alternative access if you want to add multiple storage manager accesses.
10. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
    - Set a High Water Mark value in GB.
    - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
    - Set a Low Water Mark in GB.
11. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the Microsoft Azure Blob Block storage manager and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
   e. **Available as source**. If needed, make the storage manager container available as source. Then enter the name you want to give to the platform associated.
3. Enter the Microsoft Azure Blob Block account information:
   a. **Account name** String that uniquely identifies the Microsoft Azure Blob Block account.
   b. **Access Key** Key associated with the Account Name. This key can be retrieved from the Account page via Settings > Access Key menu.
   c. **Container name** Unique name of the container created by Miria in Microsoft Azure Blob Block and where Miria stores the files.
4. Set an access tier for data storage:
   - **Hot** Store data that is accessed frequently.
   - **Cool** Store data that is infrequently accessed and stored for at least 30 days.
   - **Archive** Store data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.
5. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.
6. Set the data compression value. This specifies whether the archived data must be compressed in the storage and defines the compression type.

7. Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
8. Choose whether to activate MD5 checksum on the S3 archiving transfer.
9. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
10. Click **Create** to add the storage manager container.

Scality

When you add a Scality Object Storage, you have to complete the following steps:

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Scality** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

Or

   - **Suspended** This status is useful for maintenance operations.
6. **Default Network Address.** Network Address of the Scality storage service.

You can specify several network addresses so that Miria can connect to another node if the first node is not available.

You must enter the node network addresses, separated by semi-colons, without spaces in between (e.g., `S3-SM.archives.atempo.com;128.221.200.56;128.221.200.57`).

Each node may use one of these syntaxes:

- <`name`> or <`address`>
  Miria uses a non-secure connection (HTTP, with the default port `80`) to the node.
- <`name`>:s or <`address`>:s
  Miria uses a secure connection (HTTPS, with the default port `443`) to the node.
- <`name`>:<`port_number`> or <`address`>:<`port_number`>
  Miria uses a non-secure connection (HTTP) with a specific port number.
- <`name`>:<`port_number`>s or <`address`>:<`port_number`>s
  Miria uses a secure connection (HTTPS) with a specific port number.

These four lines are examples of a network address:

- `s3.scality.com`
- `s3.scality.com:s`
- `s3.scality.com:1523s`
- `s3.scal1;s3.scal2:s;s3.scal3:1523s`

7. **Default proxy platform** This platform handles the data movement on behalf of the usual agent or agents pool.

   Choose the proxy platform to be used by default.

8. **Alternative access** Configure it if you want to add multiple storage manager accesses.

9. **Volume management** This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.
10. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Scality and click the + button to add a container.
2. Complete the properties of the storage manager container:

   a. **Storage container name**

   b. **Deduplication domain** A new domain can be created, by clicking the + button.

   c. **Archiving run lock** A new one can be created, by clicking the + button.

   d. **Threads** Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.

   e. **Available as source** If you select this option, you have to enter the name of the platform associated to the storage manager container.
3. Set the configuration of the storage manager container:

   a. **Access Key ID** String that uniquely identifies the Scality account.

   b. **Secret Access Key** Password associated with the Access Key ID.

   c. **Bucket name** Logical path under which the data are stored into the Scality storage. Refer to your Scality storage configuration.

   d. **Digest on storage** This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived. The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.

   e. **Data compression** This specifies whether the data must be compressed in the storage and defines the compression type.

   f. **Metadata** Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.

   g. **MD5 checksum** Choose whether to activate MD5 checksum on the S3 archiving transfer.
4. Set the retention mode on the object lock. If you enable it, you have two options:

   a. **Enable governance mode** Users with specific Identity and Access Management (IAM) permissions can overwrite or delete protected object versions during the retention period.

   b. **Enable compliance mode** No users can overwrite or delete protected object versions during the retention period. To delete objects that have this configuration, you must close the account that they are associated with.
5. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
6. Click **Create** to add the storage manager container.

Quantum Active Scale

When you add a Quantum ActiveScale Object Storage, you have to complete the following steps:

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Quantum ActiveScale** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

Or

   - **Suspended** This status is useful for maintenance operations.

6. **Network Address** Network Address of the Quantum Active Scale storage service.

You can specify several network addresses so that Miria can connect to another node if the first node is not available.

You must enter the node network addresses, separated by semi-colons, without spaces in between (e.g., `S3-AS.archives.atempo.com;128.221.200.56;128.221.200.57`).

Each node may use one of these syntaxes:

– <*name*> or <*address*>
  Miria uses a non-secure connection (HTTP, with the default port `80`) to the node.
– <*name*>:s or <*address*>:s
  Miria uses a secure connection (HTTPS, with the default port `443`) to the node.
– <*name*>:<*port_number*> or <*address*>:<*port_number*>
  Miria uses a non-secure connection (HTTP) with a specific port number.
– <*name*>:<*port_number*>s or <*address*>:<`port_number`>s
  Miria uses a secure connection (HTTPS) with a specific port number.

These four lines are examples of a network address:

– `s3.activescale.com`
– `s3.activescale.com:s`
– `s3.activescale.com:1523s`
– `s3.qsa1;s3.qsa2:s;s3.qsa3:1523s`

7. **Default proxy platform** This platform handles the data movement on behalf of the usual agent or agents pool.

Choose the proxy platform to be used by default.

8. **Connection settings** Select an HTTP REST IP rule; whether Round-robin DNS or TCP/IP latency.
9. **Alternative access** Configure it if you want to add multiple storage manager accesses.
10. **Volume management** This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
    - Set a High Water Mark value in GB.
    - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
    - Set a Low Water Mark in GB.

11. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Quantum ActiveScale and click the + button to add a container.
2. Complete the properties of the storage manager container:
    a. **Storage container name**.
    b. **Deduplication domain** A new domain can be created, by clicking the + button.
    c. **Archiving run lock** A new one can be created, by clicking the + button.
    d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
    e. **Available as source** If you select this option, you have to enter the name of the platform associated to the storage manager container.
3. Set the configuration of the storage manager container:
    a. **Access Key ID** String that uniquely identifies the Quantum ActiveScale account.
    b. **Secret Access Key** Password associated with the Access Key ID.
    c. **Bucket name** Logical path under which the data are stored into the Quantum ActiveScale storage. Refer to your Quantum ActiveScale storage configuration.
    d. **Digest on storage** This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived. The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.
    e. **Data compression** This specifies whether the data must be compressed in the storage and defines the compression type.
    f. **Metadata** Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
4. **Object lock** Set the retention mode on the object lock. If you enable it, see Data Immutability with S3 Object Lock in Administration documentation.
    a. **Enable compliance mode** No users can overwrite or delete protected object versions during the retention period. To delete objects that have this configuration, you must close the account that they are associated with.
5. **Lifecycle rules** Set them and complete following information accordingly:
    a. **Name** Name of the Lifecycle rule that defines the data migration. This name is any unique string of your choice (e.g., ada_smc_amazon, RuleForArchiving, 1toglacier_200todelete, etc.). When you launch the first job, Miria uses this name to create a rule on the Quantum ActiveScale bucket.
    b. **Transition days** Number of days at the end of which Quantum ActiveScale will transfer the objects to Glacier or Deep Archive. By default, Quantum ActiveScale performs the transfer at 00:00 the same day. A value of 1 indicates that the transfer is performed at 00:00 the next day, and so on.
    c. **Retrieval copy lifetime** Number of retrieval copy lifetime.
6. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).

7. Click **Create** to add the storage manager container.

Seagate Lyve Cloud

When you add a Lyve Cloud Object Storage, you have to complete the following steps:

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Seagate Lyve Cloud** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

Or

   - **Suspended** This status is useful for maintenance operations.

6. **Default Network Address.** Enter the network address of the Lyve Cloud storage service.

You can specify several network addresses so that Miria can connect to another node if the first node is not available.

You must enter the node network addresses, separated by semi-colons, without spaces in between (e.g., `S3-SM.archives.atempo.com;128.221.200.56;128.221.200.57`).

Each node may use one of these syntaxes:

– <*`name`*> or <*`address`*>
  Miriauses a non-secure connection (HTTP, with the default port `80`) to the node.
– <*`name`*>:s or <*`address`*>:s
  Miria uses a secure connection (HTTPS, with the default port `443`) to the node.
– <*`name`*>:<*`port_number`*> or <*`address`*>:<*`port_number`*>
  Miria uses a non-secure connection (HTTP) with a specific port number.
– <*`name`*>:<*`port_number`*>s or <*`address`*>:<`port_number`>s
  Miria uses a secure connection (HTTPS) with a specific port number.

These four lines are examples of a network address:

– `s3.lyvecloud.com`
– `s3.lyvecloud.com:s`
– `s3.lyvecloud.com:1523s`
– `s3.lvc1;s3.lvc2:s;s3.lvc3:1523s`

7. **Default proxy platform**. This platform handles the data movement on behalf of the usual agent or agents pool.

Click Selectthe up and down arrow to choose the proxy platform to be used by default.

8. **Alternative access** Configure this pane if you want to add multiple storage manager accesses.
9. **Volume management** This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
   - Set a High Water Mark value in GB.
   - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
   - Set a Low Water Mark in GB.

10. Click Create to add the storage manager.

Step 2: Add a storage manager container

1. Select the storage manager for Lyve Cloud and click the + button to add a container.
2. Complete the properties of the storage manager container:
    a. **Storage container name**.
    b. **Deduplication domain**. A new domain can be created, by clicking the + button.
    c. **Archiving run lock**. A new one can be created, by clicking the + button.
    d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
    e. **Available as source**. If you select this option, you have to enter the name of the platform associated to the storage manager container.
3. Set the configuration of the storage manager container:
    a. **Access Key ID** String that uniquely identifies the Seagate Lyve Cloud account.
    b. **Secret Access Key**Password associated with the Access Key ID.
    c. **Bucket name** Logical path under which the data are stored into the Seagate Lyve Cloud storage. Refer to your Seagate Lyve Cloud storage configuration.
    d. **Digest on storage**This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived. The more complex the hash, the slower the calculation times. A complex hash decreases performance. It is recommended to use SHA-256 for the best compromise between performance and security.
    e. **Data compression** This specifies whether the data must be compressed in the storage and defines the compression type.
    f. **Metadata**Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
    g. **MD5 checksum**
4. Set the retention mode on the object lock. If you enable it, you have two options:
    a. **Enable governance mode** Users with specific Identity and Access Management (IAM) permissions can overwrite or delete protected object versions during the retention period.
    b. **Enable compliance mode** No users can overwrite or delete protected object versions during the retention period. To delete objects that have this configuration, you must close the account that they are associated with.
5. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
6. Click **Create** to add the storage manager container.

Cloudian HyperStore

Cloudian HyperStore is an Internet storage solution designed to resolve your storage issues. Create HyperStore nodes wherever you need more storage.

Step 1: Add a storage manager

1. Click the **Infrastructure** tab, then **Object Storage & Application**.
2. Click **New storage manager**.
3. Select **Cloudian HyperStore** and click **Next**.
4. Enter the name of the storage manager.
5. Choose the appropriate status:
   - **Online** Default value if you want to perform an archiving.

   Or

   - **Suspended** This status is useful for maintenance operations.
6. Enter the network address of the Cloudian HyperStore service (e.g., `cloudian.hyperstore.myS3storage.com`).
7. If needed, set a default proxy platform. This platform handles the data movement on behalf of the usual agent or agents pool.
8. Set the **connection settings** Select an HTTP REST IP rule; whether Round-robin DNS or TCP/IP latency.
9. Set the **Alternative access**. Configure it if you want to add multiple storage manager accesses.
10. If needed, set volume management. This option enables volumes to trigger the configuration and activation of retention. See also Recycling Triggered by Volume on Storage.
    - Set a High Water Mark value in GB.
    - Set a Task High Water Mark in GB only if you want to activate scheduled monitoring.
    - Set a Low Water Mark in GB.
11. Click **Create** to add the storage manager.

Step 2: Add a storage manager container

1. Select the Cloudian HyperStore storage manager and click the + button to add a container.
2. Complete the properties of the storage manager container:
   a. **Storage container name**.
   b. **Deduplication domain**. A new domain can be created, by clicking the + button.
   c. **Archiving run lock**. A new one can be created, by clicking the + button.
   d. **Threads**. Number of threads the storage manager container can manage. Select a numeric value from 1 to 128. When selecting the number of threads, take into account the data mover capacity.
   e. **Available as source**. If needed, make the storage manager container available as source. Then enter the name you want to give to the platform associated.
3. Complete the Cloudian HyperStore account configuration:
   a. **Access Key ID** String that uniquely identifies the Cloudian HyperStore account.
   b. **Secret Access Key** Password associated with the Access Key ID
   c. **Bucket name** Logical path under which the data are stored into the Cloudian HyperStore storage. Refer to your Cloudian HyperStore storage configuration.
4. Set the digest type. This ensures that the object you retrieve has not been modified on the storage and that it is identical to the object that was archived.
5. Set the data compression value. This specifies whether the data must be compressed in the storage and defines the compression type.

6. Choose whether to activate metadata. The metadata associated with the archived object are sent as URL encoded strings. You can send only up to 2 KB of metadata. A log indicates the skipped metadata above 2 KB.
7. Choose whether to activate MD5 checksum on the S3 archiving transfer.
8. Set the retention mode for object lock. See also Data Immutability with S3 Object Lock in Administration documentation.
   - **Enable governance mode** Users with specific Identity and Access Management (IAM) permissions can overwrite or delete protected object versions during the retention period.

   Or

   - **Enable compliance mode** No users can overwrite or delete protected object versions during the retention period. To delete objects that have this configuration, you must close the account that they are associated with.
9. Set a prefix if you want to define how the files are organized on the storage manager container (i.e., what are their paths on the disk).
10. Click **Create** to add the storage manager container.

> **Note**: At object level, legal hold flag is not available.

Add Server and Agent(s)

The infrastructure is composed of a server and agent(s). The server manages all the data movers and can be installed on a dedicated physical or virtual machine.

An agent can be:

- A data mover.
- A source or target platform.

To add an agent

1. Click on **Infrastructure** tab, then the **Agent** tile. The list of agents and agent pools is displayed.
2. In the top right corner, click **Add** and select **New Agent**.
3. In the **Choose Agents to add** window, select agent(s) you want to add.
4. Click **Add Agents**.

To add an agent pool

An agent pool is a logical grouping of several agents. To create an agents pool, you must first declare each agent.

1. Click on **Infrastructure** tab, then the **Agent** tile. The list of agents and pools is displayed.
2. In the top right corner, click **Add** and select **New Pool**.
3. Enter a name for your pool.
4. Select the protocol linked to the agent used.
5. Choose the agents you want to add.
6. Click **Create**. The agents pool is added to the list.
7. To edit an agents pool, select it from the list and click the ✎ button.

Explore Agent

1. In the Web Interface, click the **Infrastructure** tab.
2. Click the **Agent** card to verify the status of your agent (Active/Inactive).
3. In the **Actions** column, click the 👁 button. The hierarchical tree structure of directories and files is displayed (Figure 24).

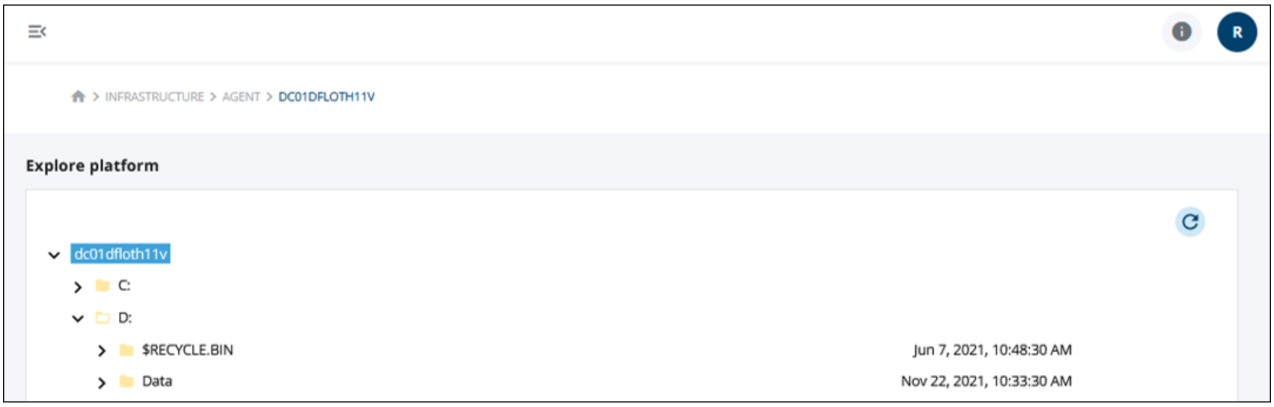

INFRASTRUCTURE > AGENT > DC01DFLOTH11V

Explore platform

dc01dfloth11v
- C:
- D:
  - $RECYCLE.BIN Jun 7, 2021, 10:48:30 AM
  - Data Nov 22, 2021, 10:33:30 AM

Figure 24: Tree structure example of a platform that has been declared as an agent

4. Select an object and perform one of these operations:
   - Add a new directory .
   - Copy a directory or file. You will be directed to the Easy Move interface where you can select a target platform.
   - Rename an object.
   - Refresh the tree structure.

Add a New NAS Platform

NAS platforms allow you to add the source and target for your migration. For an example, see also Configure an Isilon Storage.

Step 1: Connection

1. Click the **Infrastructure** tab, then the **NAS** tile. The list of NAS platforms is displayed.
2. In the top right corner, click **Add NAS**.

Step 2: Configuration

1. Choose the type of NAS (Isilon, Qumulo, etc.) or **Other** for standard CIFS/NFS file servers.
2. For Isilon, Nutanix, OceanStor and Qumulo, activate **Advanced Storage Integration** to enable Snapshot and FastScan.
3. Select the protocol relating to the agent: CIFS or NFS.
4. In the **Datamovers** section, select the agent used for the data movement.
5. Click **Next**.

Step 3 : Advanced Storage Integration (optional)

> **Note**: If you choose Isilon, Nutanix, OceanStor or Qumulo and have activated **Advanced Storage Integration**, one more step is available.

1. In the **Network** section, enter the NAS API credentials. Fields vary depending on the type of NAS.
2. If needed, check **Ignore SSL Check Certificate**.
3. (Nutanix only) Enter the CVM server information.
4. If needed, in the **Options** section, enable **Snapshot** .
5. If needed, enable **FastScan**, enter the maximum number of FastScan and choose if you want to use regular scanning if the last snapshot is missing. Check the following list for the maximum number of FastScan operations (parallel operations) recommended per NAS:
   - Isilon: 3
   - Nutanix: 10
   - OceanStor: 32
   - Qumulo: No limit
6. Click **Next**.

Step 4: Options and summary

1. In the **Stream Option** section, you can enter stream options depending on the platform type.
2. Click **Next**.
3. Read the **Summary** of the NAS.
4. Enter the name of the NAS.
5. If you chose CIFS protocol, enter the user name and the password of your windows account.
6. Click **Add NAS**.

Add a Storage Platform

You can declare an archiving platform to use a cloud or an object storage offering access as a source. To create such a platform, you must configure a storage manager and storage manager container.

The storage manager is the storage definition, which can be tape, cloud (e.g., AWS, Google), disk and object storage. The storage manager container is the path to the data of the storage manager (e.g., Bucket name, etc.).

The option **Available as source** can be enabled for certain platforms when configuring the storage manager container. See also Add a Storage Manager and Container.

To add a storage platform:

1. Click the **Infrastructure** tab, then the **Storage** tile. The list of storage platforms is displayed.
2. In the top right corner, click **Add new storage**.
3. Select a storage manager container and click **Next**.
4. Enter the name of the platform associated to the storage manager container.

5. Click **Create**.

Add a Shared File System

Allows for usage of a shared file system as a source and/or a target for migration.

Step 1: Connection

1. Click the **Infrastructure** tab, the **Shared file system**. The list of shared file systems is displayed.
2. In the top right corner, click **Add new shared file system**.

Step 2: Configuration

1. Choose the type of your Shared File System.
2. (Optional) Enable **Advanced Storage Integration** for Snapshot and FastScan .
3. In the **Datamovers** section, select the agent used for the data movement.
4. Click **Next**.

Step 3 : Advanced Storage Integration (optional)

If you activated **Advanced Storage Integration**, this step is available.

1. If needed, enable **Snapshot**.
2. If needed, enable **FastScan**, enter the maximum number of FastScan and choose if you want to use regular scanning if last snapshot is missing. Check the following list for the maximum number of FastScan operations (parallel operations) recommended per shared file system:
   - GPFS: 10
   - Lustre: no limit
   - WekaFS
   - StorNext
3. Click **Next**.

Step 4: Options and summary

1. In the **Stream Option** section, you can enter stream options depending on the platform type.
2. Click **Next**.
3. Read the **Summary**.
4. Enter the name of the shared file system.
5. Click **Create**.

Activate My File System

To browse the local file system and perform data-move operations, you need to activate **My file system**. A Web Connector application is used that needs to be installed on your workstation.

The Web Connector is able to handle one request at a time and communicates through port number 8089. You must ensure that this port number is free on your workstation.

Step 1: Download and install the Web Connector

1. In the Web Interface, click the **Infrastructure** tab.
2. Click the **My file system** card.
   - If the Web Connector has not been installed, download the WebConnector. If you decide to install the Web Connector, follow the procedure.
   - If the Web Connector is already installed, the URL opens in a new tab. You can go directly to step 2.
3. Select and download the installer that matches your operating system (Figure 25).

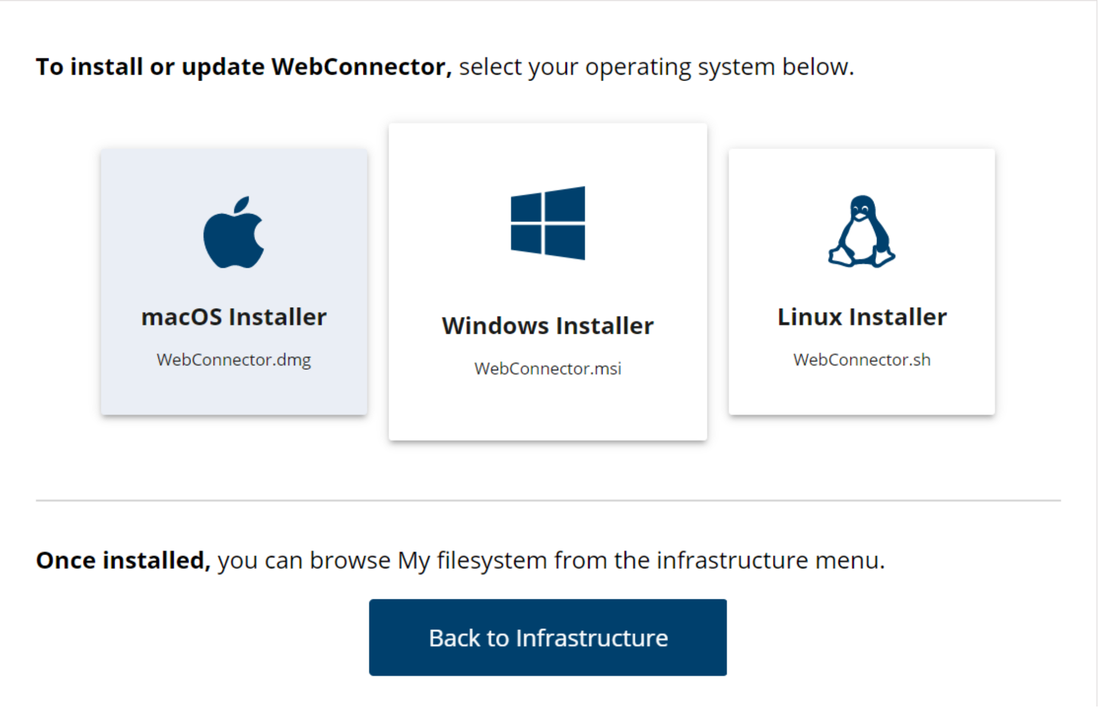

To install or update WebConnector, select your operating system below.

Once installed, you can browse My filesystem from the infrastructure menu.

Figure 25: Installer selection

4. Run the installation.

Step 2: Verify connection status of My file system

1. Once the WebConnector is installed, click **OK** to go back to the **Infrastructure** page.
2. Click the **My file system** card. The file system can now be explored.
3. Select **My file system** in the **Easy Move** interface as a source or target to perform data-move operations.

Available options

Ports Range :

When you have a shared application server, it is not possible for several users to use the same local port, only one user will be able to connect through one port.

Miria enables you to have several users logged all at the same time by using a Ports Range. This means that when you launch Web Connector, the Web browser scans the ports until it finds an available one.

By default, Miria works with TCP port range between 25000 and 26000. The administrator can customize this ports range.

If the administrator set the ports, they all have to be compatible with the Web Connector.

See the Miria's administrator guide for details.

Timeout :

The Web Connector will stop to listen on local port when user disconnects from the Web User Interface or after a timeout on inactivity of 6 hours.

The administrator can change this default timeout. See the Miria's administrator guide for details.

Edit a Platform

1. Click the **Infrastructure** tab, then select a storage type (e.g., NAS, shared file system)
2. From the list of platforms, click the ✎ button. The current platform configuration is displayed (Figure 26).

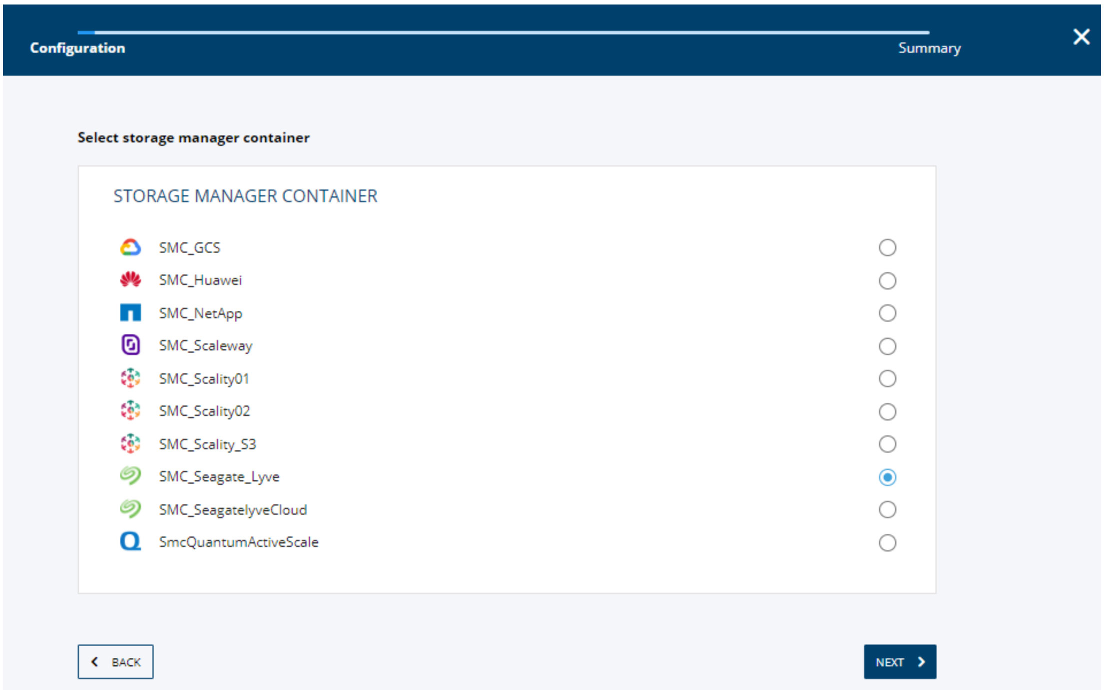

Select storage manager container

STORAGE MANAGER CONTAINER

SMC_GCS ○
SMC_Huawei ○
SMC_NetApp ○
SMC_Scaleway ○
SMC_Scality01 ○
SMC_Scality02 ○
SMC_Scality_S3 ○
SMC_Seagate_Lyve ●
SMC_SeagatelyveCloud ○
SmcQuantumActiveScale ○

Figure 26: Example displaying a Nutanix NAS configuration that can be edited in the wizard

3. Update the platform configuration and complete the wizard to save modifications.

Platforms Permissions

The permissions enable you to grant or deny permissions to individual users, user groups, or overall groups to perform actions on a platform.
You can manage permissions on the following:

- NAS
- Agents
- Storages
- Shared file system

To access to the platforms permissions

1. Click the **Infrastructure** tab.
2. Select one of the tiles according to which kind of platform you want to set permissions on: either **Agents**, **NAS**, **Storage** or **Shared file system**.
   A list of the platforms appears.
3. Click the ⋮ button of one of them and select **Permissions**. Here you have access to all the permissions created.

To create a new permission on a platform

1. Click the button **+ NEW PERMISSION**. A window appears.
2. Click on the drop down list and select a user or a users group. Click **NEW PERMISSION** to validate.
3. Click **Save changes**.
   The new permission is created, you can now set it.

The interface is divided in two parts:

- **Users and Groups**:
  Lists all the users and groups for which permissions has been created. Select any of them for which you want to set permissions.
- **Permissions**:
  Shows all the available permissions for the corresponding user or users group.

The permissions that you can set, depending on the user or users group, are the following:

- Open
- Add a folder
- Rename an object
- Delete an object
- Move an object
- Copy
- Synchronize

To do so, select each time either **Inherit**, **Deny** or **Allow**. Or select one of those options next to the first line **Apply to all**.

> **Note**: The denial of a permission at any level, takes precedence over acceptance.
> If you select **Inherit**, the values will be those previously set or set by default in the settings.

To allow viewing or browsing of a platform, the admin and monitoring rights settings about platforms must also be set to **Monitoring** or **Administration**.

CHAPTER 10 - Organize and Configure the Migration

This chapter outlines how to organize a project and define the tasks within a migration project.

Organize Projects

Miria is organized by projects. It is possible to manage multiple projects with different tasks within.

Project description

A project allows to group several tasks together to provide global statistics and management on .

> Access the **Project Overview** by clicking on a project.

The **Project Overview** enables the ability to:

- Consult information on the tasks inside the project.
- Create a new task.

Task description

Tasks are automatic jobs that can be scheduled or started manually. In general, a project involves several tasks.

> Access the **Task details** by clicking on a task.

The **Task details** window presents:

- The global progress of a task.
- Information about runs.
- Information about Snapshots.

> **Note**: You can obtain a task report. To do so, click the ⋮ button of a task run, and select **Download report**.

Configure the Migration

Once the infrastructure is configured, you can organize the migration. To start a migration, the first step is to create a project and then a set of tasks.

Create a New Project

1. Connect to the Web Interface.
2. In the **Migration** tab, click **New project**.
3. Enter the desired project name and click **New project**

Your project is created.

Create a New Task

Step 1: Create a new task

1. Connect to the Web Interface.
2. In the **Migration** tab, select the project to create a task within it.
3. In the upper right corner, select ⋯ > **New task**. The task configuration wizard appears (Figure 27).

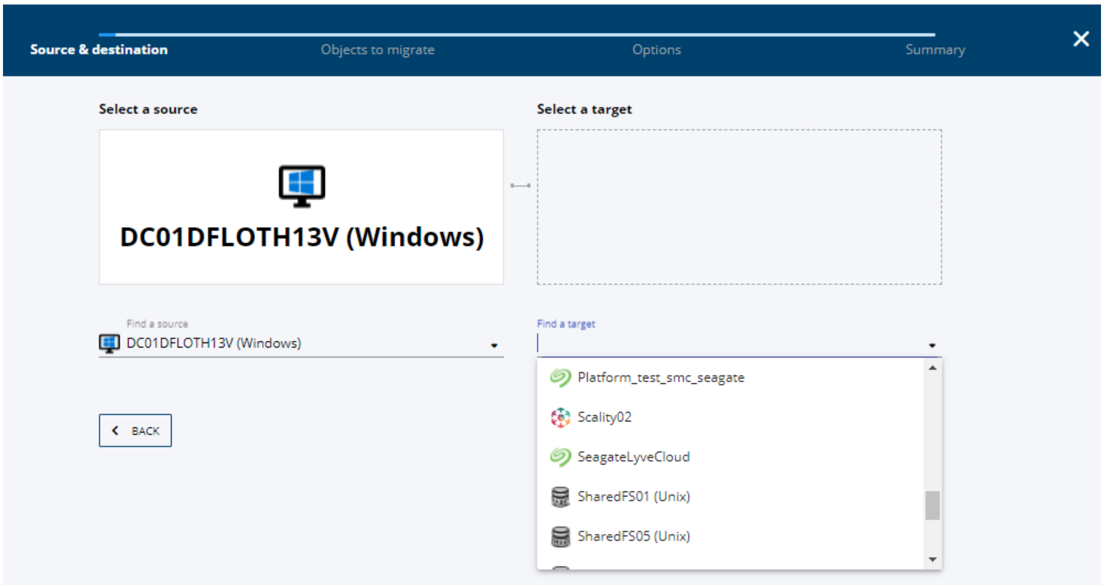

Find a source: DC01DFLOTH13V (Windows)

Find a target:

Figure 27: The search feature functions as a drop-down menu to select a source and destination

Step 2: Select a source and a target

1. Select a source.
2. Select a target. Once the migration is finished, the target storage takes over the data protection role. For example, this can be a server, NAS, shared file system, or cloud.
3. Click **Next**.

Step 3: Migrate objects

1. Select your objects to migrate.
2. Select a target for your objects.
3. Click on **Add**. Repeat this procedure for any additional source and target paths of objects that also require migration.
4. Choose the objects (files and directories) that need to be included or excluded.
5. Click **Next**.

Step 4: Set additional options

1. Select the synchronization type required for the workflow: **As Settings**, **Echo**, **Subscribe**, **Contribute**, **Combine**.
2. Set snapshot options. A snapshot is the state of a file system at a particular point in time.

- **Snapshot** A snapshot of the source directory is taken at the start of the task. This snapshot is then used as the synchronization reference. Any modifications made by the users on the source directory do not impact the synchronization process. The snapshot is used by the current task, and is deleted at the end of the task.
- **FastScan** The snapshot includes the Snapshot and FastScan features. The snapshot is not deleted at the end of the task. It is used by the next task execution to determine the objects that have changed, and must be archived.

3. To define a schedule, enable **Scheduling**:
   a. Select the day of the week and the frequency for the task to run.
   b. Select the hours and minutes of the task to run.
4. To split the current task into multiple tasks, enable **Split tasks**.
5. To activate the multithreading feature, enable **Thread** and write a numerical value from 1 to 128. This parameter sets the maximum number of parallel input / output streams. The higher its value, the better the performance of the source pool when transferring data.

   The optimum value for the Number of Threads parameter depends on the number of cores available to each of the data movers that make up the source platform pool. To calculate this value, divide by 2 the number of cores of data movers that are part of the pool.

> **Example**:
> For a pool where each data mover has 8 cores, set this value to 4. For a pool where the number of cores is different between data movers, divide by 2 the number of cores of the data mover with the smallest number of cores. If one of the data movers has 16 cores and another has 32, set this value to 8. The recommended value here is the number of data mover cores divided by two.

6. To parallelize jobs, activate **Jobs Parallelization**. This option enables Miria to accelerate data movements. This is useful for large volumes of data as it allows, for example, the ability to write several files simultaneously on the target file system.
   By default, the task scans the entire storage to identify the items to be migrated, lists them, and creates a selection of the files. By default, the task cannot trigger the migration job until the selection step is complete. A storage scan can be time consuming if the number of files to be migrated is very high and / or if the deduplication is activated. Using job parallelization changes the default behavior.
7. When job parallelism is enabled, set:
   a. The maximum number of simultaneous parallel jobs: in general, 2 times the number of data movers that make up the platform pool.
   b. The number of minutes you want before starting a new job.
   c. The maximum volumes in GB you want before starting a new job. The default value is 1024 GB.
   d. The maximum number of files you want before starting a new job. The default value is 250,000 files.
   e. To set only one of the limits (time, size, or number of files), enter `0` in the field that the Synchronization task should ignore. You must define at least one of the three fields.

> **Note**: Each of these limits define when a new job will be started. The selection operation continues in parallel with the copy job until the defined limit is reached again, which in turn will cause a new copy job. This continues until you reach the maximum number of parallel copies allowed.
>
> The limitation(s) on file selection should be sufficient to allow the job to select enough

> files, but not too much to prevent the data mover from being idle when selecting files. Experience determines the best compromise between these priorities.
>
> Three limits are considered during the synchronization job. They are presented in the following order:
>
> 1. Volume of the selection
> 2. Number of files in the selection
> 3. Time : even if the selection has not reached volume or number, the copy job is triggered at the end of the set time.

8. Choose the copy mode:
   - The following options are available to copy data:
     - **Copy operating system rights** Miria copies all the file and directory data and alternate streams.
     - **Traverse symbolic links** Miria creates the symbolic links while copying the symbolic link data to the appropriate destination for both files and directories.
     - **Manage hardlinks** (Linux/Unix only over NFS) Miria copies the directory structure from the source file system to the destination. The hardlink structure on the source is preserved and rebuilt at the destination location. It is recommended to use this option in **Echo** mode. The task scans and creates the hardlinks, for which the explicit mode is required. This excludes the **Subscribe** and **Combine** modes, where the synchronization is applied from the destination to the source.

   Or

   - Copy only permissions (ACL), extended attributes (XATTR) and alternate data streams (ADS).
9. To check migration integrity, select a hash algorithm. An integrity check is possible in both normal and cut-over mode and enables monitoring and producing reports on the migration task.
10. Click **Next**.

Step 5: Summary

1. Read the **Summary** of the task.
2. Enter a name for the task
3. Associate the task to a project.
4. Click **New task**.

Edit a Project

1. Click the tab, then the ⋮ button in a project.
2. Click **Edit project**.
3. Enter a new name for the project.
4. Click **Update**.

Edit a Task

1. Click the tab, then select a project.
2. In the , select a task and click the ⋮ button.
3. Select **Edit task.** The task configuration wizard is displayed.

4. Modify the section(s) you want to edit and click **Update** to save changes.

Start Migration Task(s)

A project may consist of multiple tasks. Miria provides the ability to start all tasks at the same time or to start tasks independently.

Start an individual task

1. Select the project.
2. In the **Tasks** section, click the ⋮ button of a task.
3. Choose one of the following options:
    a. **Start task**.
    b. **Start task in Test Mode**.
    c. **Start task with Pre-Cutover**.

Start all tasks

1. Select the project.
2. In the upper right corner, click **Start** and choose:
    - **Start all tasks**.
    - **Start all tasks in test mode**.
    - **Start all tasks with pre-cutover**.

Duplicate a Task

1. Click the tab, then select a project.
2. In the **Project overview**, select a task and click the ⋮ button.
3. Select **Duplicate task**.
4. Enter the name of the task to be duplicated.
5. Click **OK**. The new task is created within the project.

Disable a Project or a Task

1. Click the tab:
    - Click the ⋮ button in a project and select **Disable**.

    Or

    - Select a project, then click the ⋮ button in a task and select **Disable**.
2. Click the **Settings** tab, then **Hidden projects and tasks** to locate and restore disabled projects and tasks.

CHAPTER 11 - Manage the Migration

This chapter lists the 3 steps of a migration project. To manage a migration, tasks must be started at each step of the migration project:

- If a task is scheduled, it will start automatically. See Step 4: Set additional options.
- If a task is not scheduled, the task will need to be run. See Start Migration Task(s).

Step 1: First Synchronization

The first step in a migration project consists of synchronizing the initial transfer of data between the source storage and the target.

Before starting the synchronization, define the user workflows on the source storage and the target storage during the data migration.

- If production on the target storage is halted, the task can start in **Echo** mode.
- If production on the target storage starts immediately, it is recommended to use **Contribute** mode.

See also the Synchronization Types section in Miria Administration Documentation.

The synchronization will behave differently depending on whether snapshot is enabled or not.

With Snapshot:

1. Miria will make a reference snapshot of the source storage (Figure 28).
2. Perform a full scan of the snapshot to validate it.
3. Synchronize data from the source to the target storage.

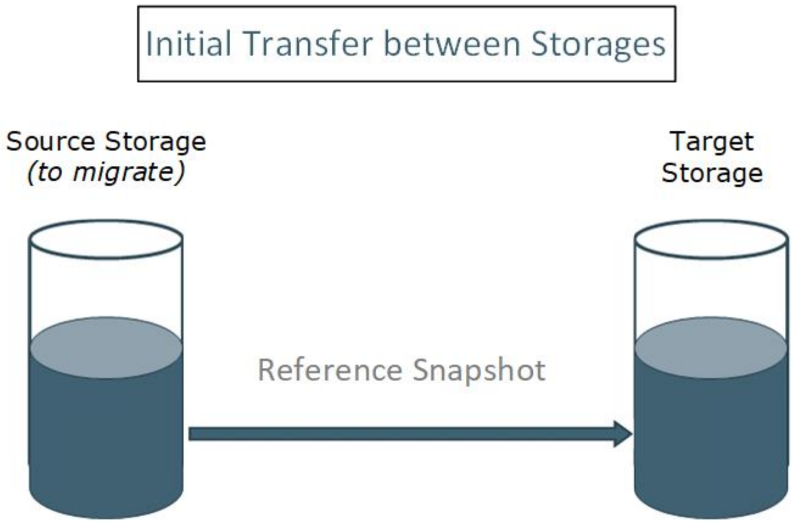

Figure 28: Diagram of a reference snapshot from a source to a target storage

Without Snapshot:

If snapshot is not available on the source storage:

1. Scan the entire file system.

2. Synchronize data to the target storage. The first synchronization without a snapshot can take days/weeks depending on the volume, size and number of files.

When performing the first synchronization with or without snapshot, it is recommended to split the file system by path on several tasks. This enables to scan the file system in parallel.

To start the first synchronization:

1. Open the project or the task you want to migrate.
2. Click Start Migration Task(s).
3. Click the **Activity** tab to verify the job progress.

Step 2: Incremental Synchronization

The incremental synchronization is the process of running a task to migrate daily updates or changes since the first synchronization task was run (Figure 29). It can take several days or weeks depending on your data volume.

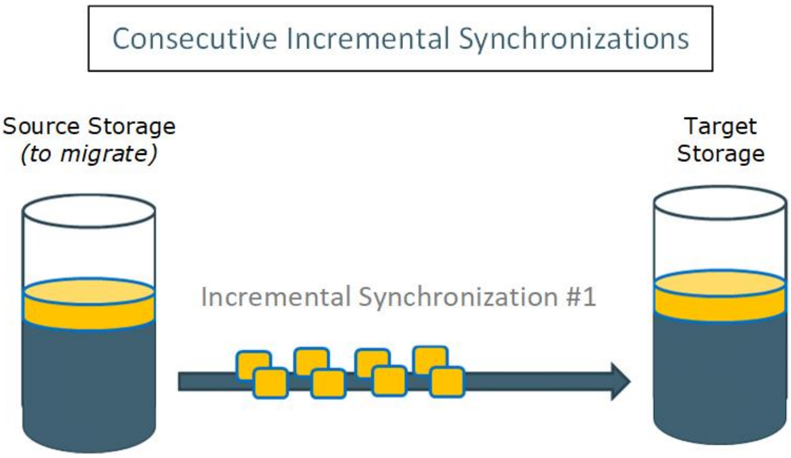
Consecutive Incremental Synchronizations
Source Storage
(to migrate)
Target
Storage
Incremental Synchronization #1

Figure 29: Incremental synchronization from a source to a target storage

With Snapshot:

Each synchronization follows 5 sequential stages:

1. Second snapshot of the source storage.
2. FastScan retrieves the list of objects that have been modified since the previous snapshot.
3. Miria validates the list returned by FastScan through comparison between source and target storage. Miria for Migration produces a change list.
4. Detected modifications are propagated (creation, deletion in **Echo** mode, etc.) to the target storage.
5. If the synchronization is successful, the previous snapshot is deleted and the current snapshot will be the reference point for the next task.

> **Note**: Adding additional data movers will reduce the time needed for migration.

These incremental synchronizations enable the operations to converge towards the final synchronization.

With FastScan:

Scanning begins:

1. First run:
   - If parallelization rules are set:
     - Miria scans until triggering the parallelization rules defined, and spawns a new copy, sync, archive, or backup job.
     - The scan keeps going until triggering once again the parallelization rules, and so on till the maximum parallel jobs defined into the task definition are triggered.
   - If no parallelization rules are defined:
     - All of the source platform is scanned.
     - Once the scanning is over, the data movement is initiated.
2. FastScan incremental runs:
   - Changes list are consolidated against source storage API.
   - Miria scans against this change list.
   - If parallelization rules are set:
     - Miria scans until triggering the parallelization rules defined and spawns a new copy, sync, archive, or backup job.
     - The scan keeps going until triggering a new time the parallelization rules, and so on till the maximum parallel jobs defined into the task definition are triggered.
   - If no parallelization rules are set:
     - Miria scans all.
     - Once scanning is over, the data movement is initiated.

To manage the incremental synchronization:

1. Open the project or the task you already run in the Step 1: First Synchronization.
2. Click Start Migration Task(s).
3. Click the **Activity** tab to verify the job progress.
4. Start the tasks as many times as necessary.

Step 3: Final Synchronization (Cutover)

The final synchronization is performed to complete the data migration to the target storage while the production is stopped on the source storage (Figure 30). It ensures that the source and target storage are fully synchronized. During the final synchronization, Miria does a full scan of the reference snapshot without creating or deleting one.

The final synchronization task is supported only in **Echo** mode. When launching the task, all new data on the target storage (which does not exist on the source storage) will be deleted.

> **Note**: During the final synchronization, the source storage must be on read-only mode in accordance with the customer's constraints. This should have been predefined in the preparation and planning phase and the Statement of Work.

Cutover times are dependent on the size of the data sets being migrated and the daily change rate.

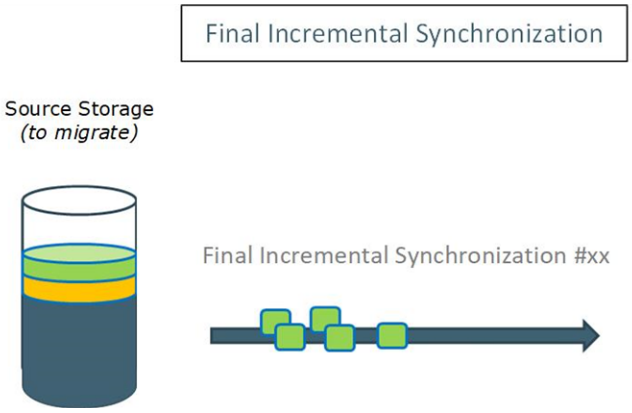
Final Incremental Synchronization
Source Storage
(to migrate)
Target
Storage
Final Incremental Synchronization #xx

Figure 30: Final synchronization process

Step 1: Start the final synchronization

1. Stop production on the source storage or set to read only.
2. Open the project or the task you have already run.
3. Click the ⋮ button of a task.
4. Select **Start task with Pre-Cutover**.

Once the data on the target storage is fully synchronized, proceed with the copy of the ACLs and alternate streams in the next step. This allows to:

- Replicate the rights and permissions on the source storage.
- Restore the modification dates of the folders on the target storage.

Step 2: Copy permissions, extended attributes and alternate streams

1. Select the task for which the data synchronization is run and click the ⋮ button.
2. Select **Edit task**.
3. In the Task creation wizard, click **Next** until the task options are displayed.
4. In the **Copy mode** section, select the option to copy only permissions (ACL), extended attributes (XATTR) and alternate data streams (ADS) (Figure 31).

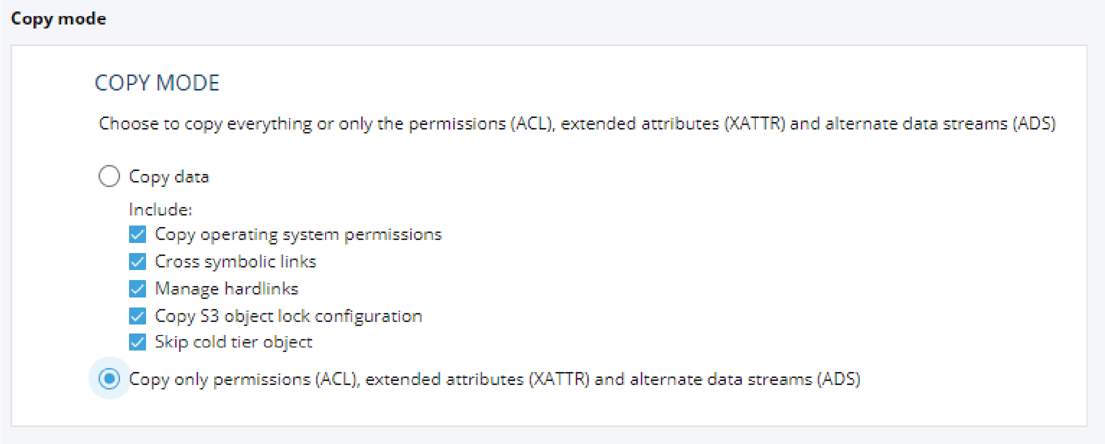
Copy mode

COPY MODE

Choose to copy everything or only the permissions (ACL), extended attributes (XATTR) and alternate data streams (ADS)

○ Copy data

Include:

☑ Copy operating system permissions
☑ Cross symbolic links
☑ Manage hardlinks
☑ Copy S3 object lock configuration
☑ Skip cold tier object

● Copy only permissions (ACL), extended attributes (XATTR) and alternate data streams (ADS)

Figure 31: Copy mode options

5. Click **Next** and then **Update task** to save modifications.
6. In the task overview, click the ⋮ button and select **Start task**.

Once the migration task is successfully completed, production on the target storage can start.

CHAPTER 12 - Organize and Configure a Backup Project

The **Backup** entry in the Web Interface is used to create and manage backup projects and tasks.

> Click the **Backup** tab.

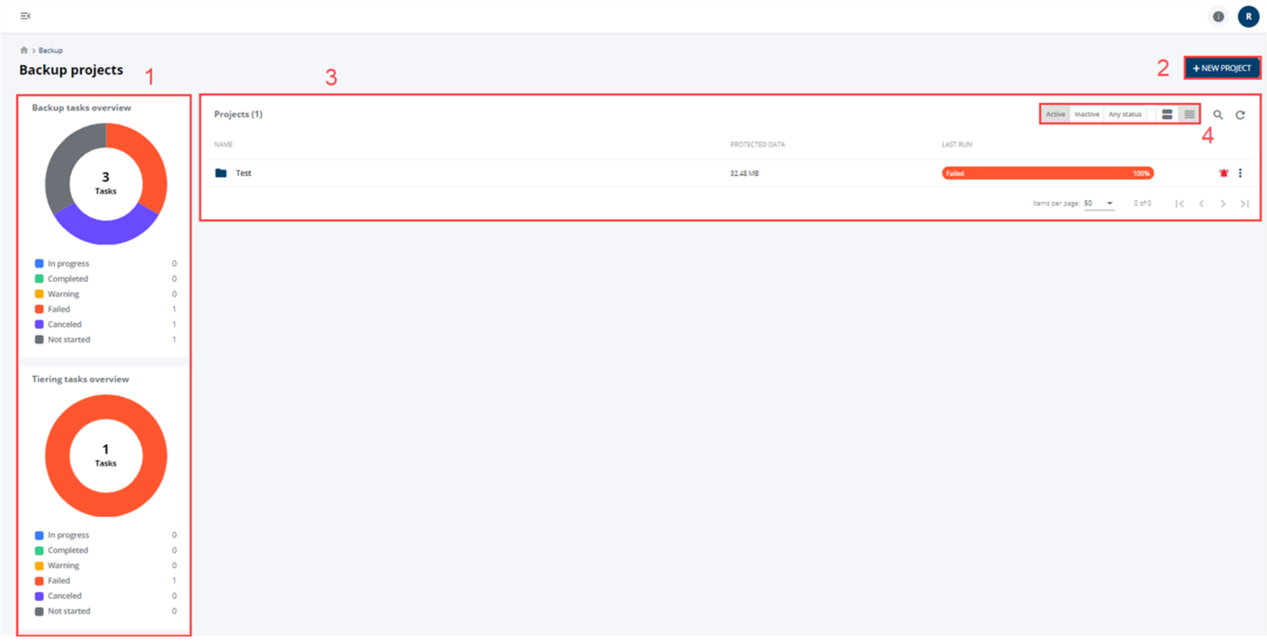

Backup tasks overview

🟦 In progress	0
🟩 Completed	0
🟨 Warning	0
🟥 Failed	1
🟪 Canceled	1
⬛ Not started	1

Tiering tasks overview

🟦 In progress	0
🟩 Completed	0
🟨 Warning	0
🟥 Failed	1
🟪 Canceled	0
⬛ Not started	0

NAME	PROTECTED DATA	LAST RUN
Test	32.48 MB	Failed 100%

Figure 32: 1. Backup and tiering tasks; 2. Add new project; 3. Current projects; 4. Project representation

The following information can be displayed on backup projects and tasks (Figure 32):

- Status of backup and tiering tasks.
- Backup projects, the volume of protected data and the number of tasks created within a project.
- Projects can appear as individual tiles or in a list.

Organize Projects

Miria is organized by projects. It is possible to manage multiple projects with different tasks within.

Project description

A project allows to group several tasks together to provide global statistics and management on .

> Access the **Project Overview** by clicking on a project.

The **Project Overview** enables the ability to:

- Consult information on the tasks inside the project.
- Create a new task.

Task description

Tasks are automatic jobs that can be scheduled or started manually. In general, a project involves several tasks.

> Access the **Task details** by clicking on a task.

The **Task details** window presents:

- The global progress of a task.
- Information about runs.
- Information about Snapshots.

> **Note**: You can obtain a task report. To do so, click the ⋮ button of a task run, and select **Download report**.

Configure the Backup

This section outlines how to create a backup project and add tasks within it.

Create a New Project

1. Click the tab, then the **New project** button.
2. Enter the desired name and click **Create**.

Your project is created.

Create a New Backup Task

The automatic backup task is aimed at automating the backup of specified directories and files. This task enables you to schedule both full and incremental backups into a single task.

The DDN IntelliFlash FastScan NAS helps you to reduce the time needed for an incremental backup.

> **Note**: Before creating a backup task, you must have configured a policy.

Step 1: Create a new task

1. In the **Backup** tab, select the project to create a task within it.
2. In the upper right corner, select **Add** > **New backup task**. The task configuration wizard is displayed (Figure 33).

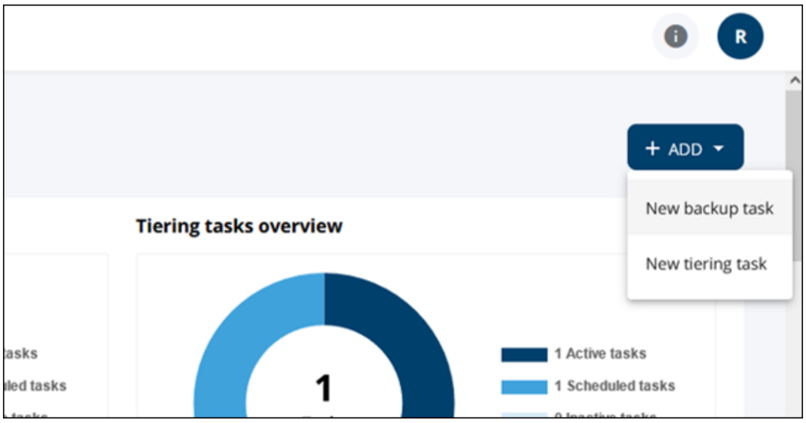

Figure 33: Add a new backup task

Step 2: Select a source and a target

1. Enter the name of the backup task.
2. Select a source platform. This is the root of the path from which the backup job is to search for files and directories on the platform.
3. Select a target and set the following options:
   a. Select a repository name or create a new one in which the data will be backed up.
   b. Activate **Full path auto-generated** if the data must be backed up at the root of the selected repository. Miria controls the path. The task replicates the file tree of the source directory in the target repository.
   c. Activate **Backup repository** if the data must be backed up in the repository of your choice.
4. Select a policy to apply.
5. Click **Next**.

Step 3: Back up objects

1. Explore the source platform and add objects to backup.
2. Repeat step 1 for each object to backup. Added objects can easily be removed from **Your selection** by clicking the button.
3. Include or exclude objects if the backup selection needs to be refined.
4. Click **Next**.

Step 4: Schedule full and incremental backups

1. Enable a backup schedule (full and/or incremental).
2. Set following options for tasks to be launched:
   - Day(s) of the week.
   - Intervals at which the task is launched in the month (e.g., the first Sunday or third Friday).
   - Time(s) of the day.

   Example 1: A task must be launched at 04:45: Choose 4 Hours and 45 Minutes. If you select Every 5 minutes, the task runs every five minutes during the hours that you specify under Hours.

Example 2: A task must be run between 03:00 and 05:00 and relaunched every five minutes: Select 3 and 4 Hours and Every 5 minutes. This option is useful when you have source material that is constantly changing.

3. Click **Next**.

Step 5: Set additional options

1. If needed and if snapshot management is available on the source, activate snapshots.
2. Enable **Jobs Parallelization** to accelerate data movements. This feature is particularly useful for large volumes of data as it enables you, for instance, to back up to multiple tape drives simultaneously. If enabled, set following options:
   a. Maximum number of jobs that can run in parallel.
   b. Create a job after:
      – Period of time assigned to the task to browse the file system and build up a file selection.
      – Maximum size that the file selection can reach. The default value is 1,024 GB.
      – Maximum number of files that the file selection can reach. The default value is 250,000 files.

   If only one of the limits (time, size, or number of files) is set, enter 0 in the field that the task must ignore. You can set to 0 only one of these fields. The limit(s) that you set to the file selection should be enough for the task to select a sufficient number of files to feed a job, but not too much as to leave drives idle if writing to tape. See also Parallel Jobs - Use Case in Miria Administration documentation.
3. Enable **Commands** to be able to enter the full path of any scripts you want to launch before or after the task is run.
4. Enable **Retention of deleted objects** to be able to define an additional retention for backed up objects which were deleted at the source.
5. Click **Next**.

Step 6: Summary

1. Read the summary of the task.
2. Click **Create** to add the backup task.

Create a New Tiering Task

The automatic tiering task allows you to replicate data backed up on one storage to another storage. For tiering to work, the storage for the filter and the storage for the policy must point to the same agent.

Step 1: Create a new task

1. In the **Backup** tab, select the project to create a task within it.
2. In the upper right corner, select **Add** > **New tiering task**. The task configuration wizard is displayed.

Step 2: Select a source and a target

1. Enter the name of the tiering task.
2. Select a repository.

3. Select the data you want to replicate. This can be a complete repository, a directory or a folder.
4. In the **Filter** section, click the **New filter** button.
5. Select the storage manager containers where the data are located. This can be a File Storage One to One or a File Storage Container.
6. Select a policy to apply.
7. Click **Next**.

Step 3: Set task options

1. Select the task mode:
   - Full

   Or
   - Incremental
2. Enable **Scheduling** to specify the days of the week, hours and recurrence of tiering tasks to be run.
3. Enable **Jobs Parallelization** to accelerate data movements.

Step 4: Summary

1. Read the summary of the task.
2. Click **Create** to add the tiering task.

Edit a Project

1. Click the tab, then the ⋮ button in a project.
2. Click **Edit project**.
3. Enter a new name for the project.
4. Click **Update**.

Edit a Task

1. Click the tab, then select a project.
2. In the , select a task and click the ⋮ button.
3. Select **Edit task**. The task configuration wizard is displayed.
4. Modify the section(s) you want to edit and click **Update** to save changes.

Start Task(s)

A project may consist of multiple tasks. Miria provides the ability to start all tasks at the same time or to start tasks independently.

To start a backup task

1. Select the project.
2. In the **Tasks** section, click the button of a task.
3. Choose one of the following options:
   - Start task in full mode.
   - Start task in incremental mode.

- Start task in test full mode.
- Start task in test incremental mode.

To start all backup tasks

1. Click the button of a project.
2. Choose one of the following options:
   - Start all tasks in full mode.
   - Start all tasks in incremental mode.
   - Start all tasks in test full mode.
   - Start all tasks in test incremental mode.

Only backup tasks can be launched in full or incremental mode.

To start a tiering task

1. Select the project.
2. In the **Tasks** section, click the button of a task.
3. Choose one of the following options:
   - Start task.

   Or

   - Start task in test mode.

Duplicate a Task

1. Click the tab, then select a project.
2. In the **Project overview**, select a task and click the ⋮ button.
3. Select **Duplicate task**.
4. Enter the name of the task to be duplicated.
5. Click **OK**. The new task is created within the project.

Disable a Project or a Task

1. Click the tab:
   - Click the ⋮ button in a project and select **Disable**.

   Or

   - Select a project, then click the ⋮ button in a task and select **Disable**.
2. Click the **Settings** tab, then **Hidden projects and tasks** to locate and restore disabled projects and tasks.

CHAPTER 13 - Manage Users

You must create a user to access and, for instance, manage permissions. A user unknown to Miria does not have the permission to access the software.

Users are created either:

- **Manually** Declare each user individually and enter all the user parameters and permissions manually.
- **Automatically** Define a Reference User as a pattern. Auto-creation of users is only possible with the LDAP access modes. The first time a user logs in, a user is created with the profile and permissions of the Reference User.

Add a User

1. Click the **Users** tab, then the **Users** tile. The list of users is displayed.
2. Click the **+ New user** button in the top right corner. The user creation wizard is displayed (Figure 34).

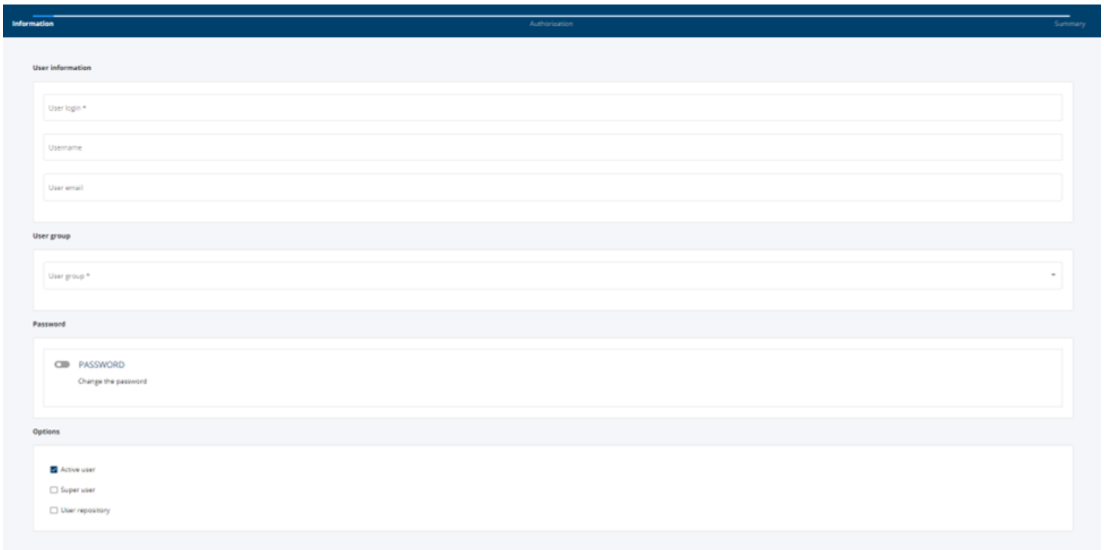

Figure 34: The user creation wizard

3. Enter the user information:
   - **User login** Name by which Miria knows the user. The user must use this name to connect. You cannot use the backslash (\) character.
   - **Username** (Optional) Name of the person to whom the user login is assigned. For example, the person who logs in as *ntillb* is Norbert Tillbury.
   - **User email** (Optional) Email address that can be used to notify the administrator of any action by this user.
4. In the **User group** list, select the user group to which the user belongs. The user can belong to only one user group.
5. Activate the **Password** button to assign the user a password. By default, you can enable empty passwords by letting the button turned off.

6. Define the user options:
   - **Active user** If this option is enabled, the user has the permission to connect to Miria. If it is disabled, the connection is denied.

     > **Note**: If you disable the **Active user** option after users have already archived data, they can no longer access Miria, but any data previously archived is not deleted from the system.

   - **Super user** The user can log as the administrator and have full administration rights over the application. By contrast, standard users have only the right to perform operations on their own repositories.
   - **User repository** Creates a personal repository for the new user as soon as the form is validated. Only this user or a super user has access to this personal file.
7. Click **Next**. A summary of the user configuration is displayed.
8. Click **Create** to confirm the user creation.

Edit a User

1. Click the **Users** tab, then the **Users** tile. The list of users is displayed.
2. From the list of users, click the ✎ button. The User configuration wizard is displayed (Figure 35).

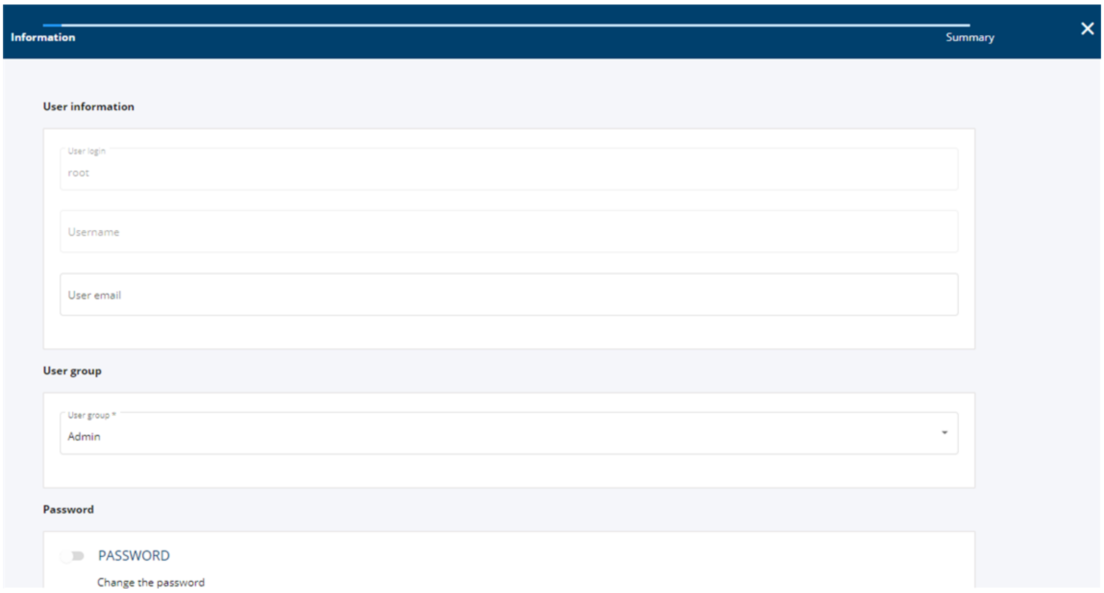

User information

User login
root

Username

User email

User group

User group *
Admin

Password

PASSWORD
Change the password

Figure 35: Example displaying a user configuration that can be edited

3. Update the user configuration and complete the wizard to save modifications.

Add a User Group

User Groups inherit the setting values from the default settings.

Conversely, if you specify a setting on a user group, the value applies to that user group, but also to all of these objects that are lower in the hierarchy.

There are three types of user groups:

- User group, which can contain only users. See To add a user group for its creation procedure.
- Overall group, which is a group of groups. It can contain user groups and users, but not other overall groups. See To add a user group for its creation procedure.
- LDAP group, which is a kind of overall group that represents an existing LDAP group on an LDAP server. It enables you to assign permissions to users of this group. See To add a user group for its creation procedure.

To add a user group

1. Click the **Users** tab, then the **Groups** tile.
2. Click the **+ Add** button at the top right and select **Add user group**. The user creation wizard is displayed.
3. Enter the group name.
4. Click **Next**.
5. Select the group authorizations:
   - **None**.
   - **Monitoring** Enables the group to read information.
   - **Administration** Enables the group to edit information.
6. Click **Next**. A summary of the group configuration is displayed.
7. Click **Create** to confirm the group creation.

To add an overall group

1. Click the **Users** tab, then the **Groups** tile.
2. Click the **+ Add** button at the top right and select **Add overall group**. The user creation wizard is displayed.
3. Enter the group name.
4. Click the **Members** list to display the entire list of users and user groups to add to your overall group.
5. Select each item that you want to add.
6. Click the **+** button to validate your selections.
7. Click **Next**. A summary of the group configuration is displayed.
8. Click **Create** to confirm the group creation.

To add an LDAP group

1. Click the **Users** tab, then the **Groups** tile.
2. Click the **+ Add** button at the top right and select **Add LDAP group**. The user creation wizard is displayed.
3. Enter the group name.
4. (Optional) Activate **Secure mode** to select a certificate.
   - Choose the path to the certificate to be used to connect to the LDAP server with SSL.
   - Check the **Ignore SSL check certificate** box if you do not wish to establish SSL verification of the machine.
5. Define the group configuration:

- **Server type** In the drop-down list, choose the type of server between **Active Directory** and **LDAP**.
- **Server address** Enter the IP address or the name of the LDAP server. This server must host the LDAP directory that contains the users you want to import.
- **User** and **Password** (Optional for Active Directory server) Enter a username and password to authenticate to the LDAP server. If both fields are filled in, you must fill in the User Base DN field before selecting the Base DN. If both fields are empty, an anonymous connection will be used.
- **Base DN** Select the root directory of a server.

  Each entry stored in LDAP databases requires a unique identification or Distinguished Name (DN). The top hierarchy in an LDAP directory tree is called the Base DN.

6. Click **Next** and define the advanced configuration options:
   - **User key** Automatically pre-filled according to the type of server. It contains the attribute name to be used to retrieve the user's name. This is the attribute to be used by default when autocreating the LDAP user. Its value is `sAMAccountName` for Active Directory and `uid` for LDAP.
   - **Internal user key** Enter a specific attribute if you want the user's name to be different from the one used to connect to LDAP. During autocreation, the value of this attribute will be used instead of the value of the attribute entered in the user key.
   - **Bind DN** Automatically pre-filled according to the type of server.
   - **Group base DN** Mask used to authenticate to the LDAP server. It is used to reformat connection credentials during authentication. Its value is different depending on the type of server:
     - For Active Directory: `[DomainName]\{LoginName}`. The prefix `[DomainName]` is not mandatory, and it will be replaced by the domain name used (if it exists). For example: `TEMQLDAP\{LoginName}`.
     - For LDAP: `uid={LoginName}`.
   - **User base DN** Select the group or domain containing all users of the domain. This field is mandatory for LDAP servers.
   - **Group** Select a group from the level at which users are to be searched. The Group must have a Distinguished Name. This field is mandatory for all types of servers.
7. Click **Next**. A summary of the group configuration is displayed.
8. Click **Create** to confirm the group creation.

Configure SAP

A security authentication path (SAP) defines an authentication authority that is in charge of determining whether a user has the permission to access Miria.

These are the three types of authority authentication:

- Miria internal security system.

This is the default authentication authority that corresponds to the Free Login access mode. With this system, user names and passwords are stored in the Miria database.

- LDAP server.

  If the authentication is delegated to an LDAP server, the passwords are not stored in Miria.
- LDAPS server.

  This is the same as LDAP server, but adds encryption between the Miria server and the LDAP server.

You can declare several authentication authorities, and sort them in order of priority. Miria checks user access with the first authentication path, then with the second path if the access is denied with the first one, etc.

Only unique usernames are supported. Configurations in which the same user name exists in several domains associated with different passwords are not supported.

Set Password Policies

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. Click the ✎ button of the local rule. The rule creation wizard is displayed.
3. Enter the rule name.
4. Set a password policy. If enabled, complete these parameters:
   - Number of figures.
   - Number of lower cases
   - Number of upper cases.
   - Numbers of special characters.
   - Length of the password.

If the admin change the password policy, the user will be asked to set a new password that comply with the new policy, at the next connection.

You can define only one path of Local type, but several paths of LDAP or LDAPS types. See Add a LDAP Rule and Add a LDAPS Rule for their creation procedure.

Reorder Rules

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. If needed, drag and drop the SAPs to define its order of priority. The path that to be checked first must be on top of the list.

> **Note**: If you decide that user authentication must be delegated to an LDAP server, and place LDAP on top of the list, the users connecting to Miria must use their LDAP credentials.

Test Rules

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. Click the button next to the list to test the authentication paths.
3. In the window that displays, enter the name and password of the user for whom you want to test the path.
4. Click **Continue** to validate the testing.

The ● icon turns green on the successful paths, whereas it turns red on the paths that fail. The icon remains gray for the paths that are not tested.

Add a LDAP Rule

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. Click the **+ Add** button at the top right and select **Add LDAP rule**. The rule creation wizard is displayed.
3. Enter the rule name.
4. Select the LDAP group member corresponding to the LDAP server group that has the permission to access Miria. All users belonging to this group are auto-created at first login using LDAP access mode. The LDAP Group must have been previously created.
5. Check **Enable auto-creation of users** if you want to create a reference user.
    a. Click the list and select the user that you want to use as a model. The auto-created user belongs to the same user groups and has the same advanced settings and permissions as the reference user.
6. Click **Next**. A summary of the rule configuration is displayed.
7. Click **Create** to confirm the rule creation.

Add a LDAPS Rule

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. Click the **+ Add** button at the top right and select **Add LDAPS rule**. The rule creation wizard is displayed.
3. Enter the rule name.
4. Select the LDAPS group member corresponding to the LDAPS server group that has the permission to access Miria. All users belonging to this group are auto-created at first login using LDAPS access mode. The LDAPS Group must have been previously created.
5. Check **Enable auto-creation of users** if you want to create a reference user.
    a. Click the list and select the user that you want to use as a model. The auto-created user belongs to the same user groups and has the same advanced settings and permissions as the reference user.
6. Click **Next**. A summary of the rule configuration is displayed.
7. Click **Create** to confirm the rule creation.

Edit a Rule

1. Click the **Users** tab, then the **Access** tile. The list of rules is displayed.
2. From the list of rules, click the ✎ button. The current rule configuration is displayed.
3. Update the rule configuration and complete the wizard to save modifications.

Set the Two-Factor Authentication

In the Web User Interface, you can set the Two-Factor Authentication:

1. Click the top right corner of your screen on the user profile icon. This menu appears (Figure 36):

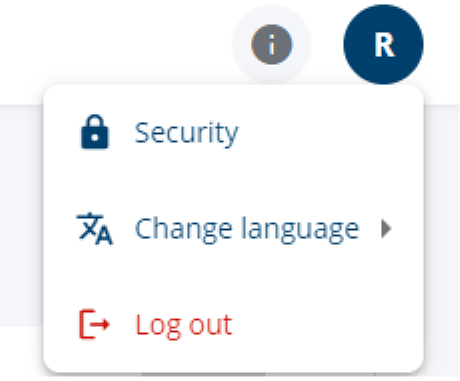

Figure 36: User profile menu

2. Select **Security**.
3. Click the **Two-Factor Authentication** tab, you now have two possibilities; the Two-Factor Authentication was set as mandatory or as optional. These options are defined in the settings by the administrator.

> **Note**: If you are logged in LDAP mode and if Two-Factor Authentication is forbidden, **Security** is hidden in the menu.

Configuring a Two-Factor Authentication Set as Mandatory

In this case, this is how the tab is displayed (Figure 37):

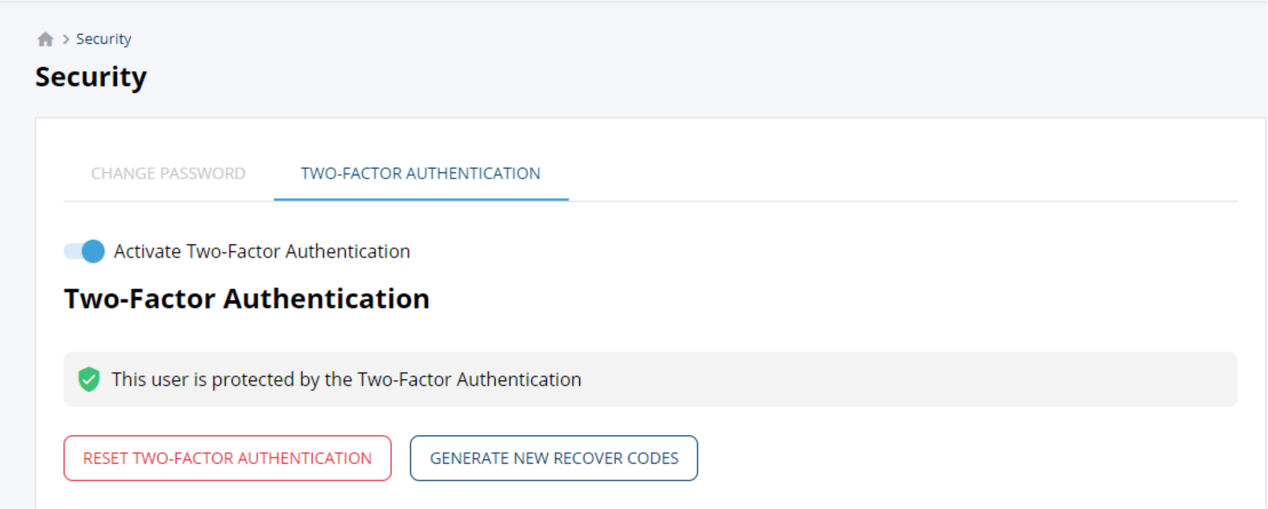

Security
Security
CHANGE PASSWORD
TWO-FACTOR AUTHENTICATION
Activate Two-Factor Authentication
Two-Factor Authentication
This user is protected by the Two-Factor Authentication
RESET TWO-FACTOR AUTHENTICATION
GENERATE NEW RECOVER CODES

Figure 37: Mandatory Two-Factor Authentication

You cannot change the toggle, it remains activated.
But you can execute other actions:

- Generate new recover codes.
- Reset the configuration.

> **Note**: When Two-Factor Authentication is mandatory, if you reset the configuration, you have to configure it again. It can be straight away, or on your next log in.

To reset Two-Factor Authentication

1. Click **Reset Two-Factor Authentication.**
2. Enter the security code that is generated on the authentication application and valid for thirty seconds.
3. Reconfigure the Two-Factor Authentication (Figure 38):

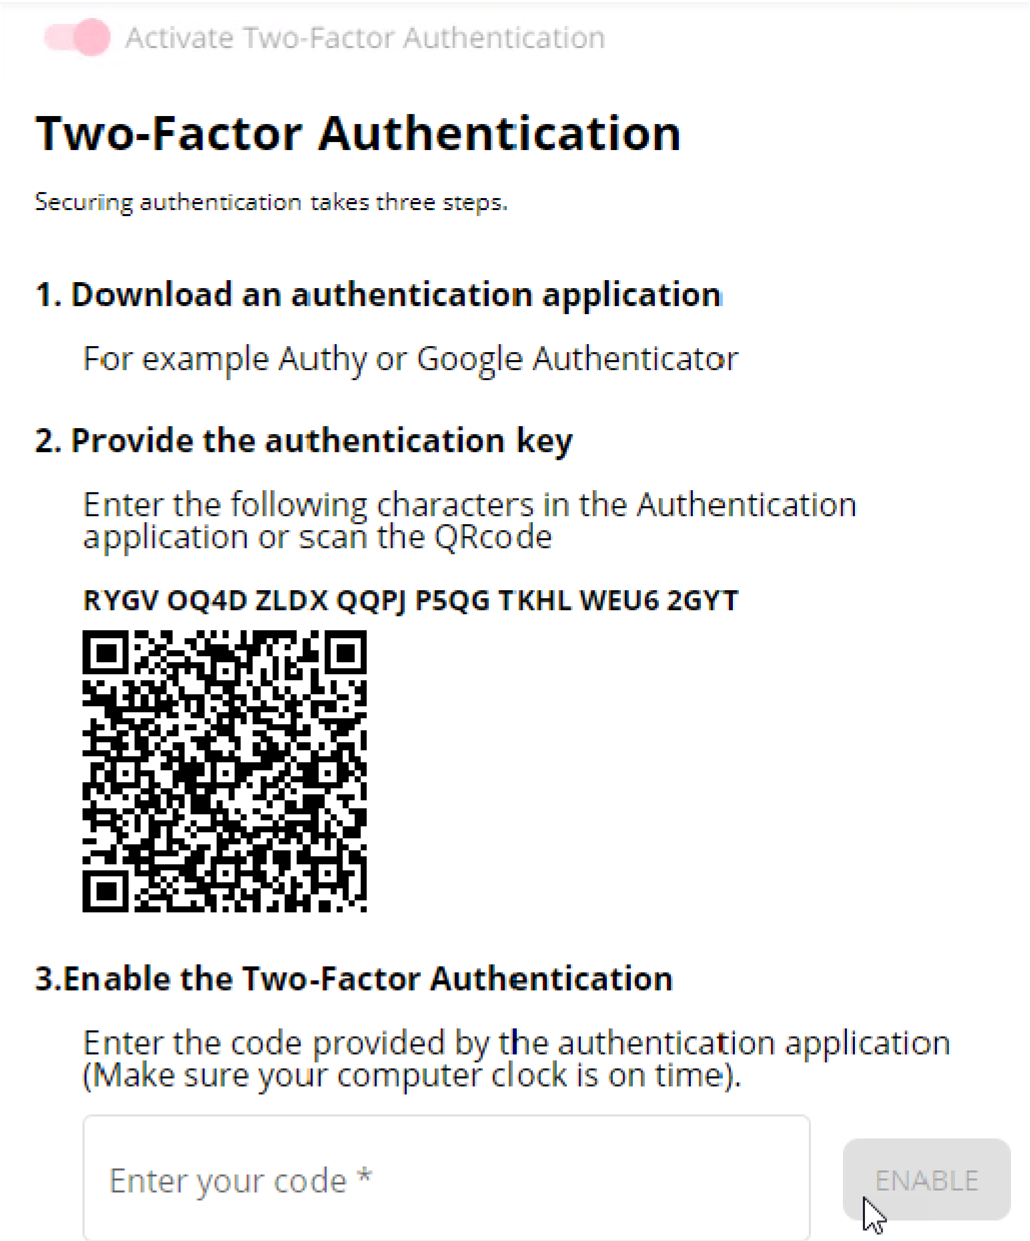

Activate Two-Factor Authentication

Two-Factor Authentication

Securing authentication takes three steps.

1. Download an authentication application

For example Authy or Google Authenticator

2. Provide the authentication key

Enter the following characters in the Authentication application or scan the QRcode

RYGV OQ4D ZLDX QQPJ P5QG TKHL WEU6 2GYT

3.Enable the Two-Factor Authentication

Enter the code provided by the authentication application (Make sure your computer clock is on time).

Figure 38: Securing authentication

4. Scan the QR code and enter the code returned by the authenticator.
   The Two-Factor Authentication is set.

Configuring a Two-Factor Authentication Set as Optional

In this case, the Two-Factor Authentication is not mandatory, you can choose by yourself to enable this option or not. You have more possibilities.

To activate the Two-Factor Authentication

1. Click the toggle to define this option as enabled.
2. Scan the QR code on your Authentication App.
3. Enter the code returned by the authenticator.

   You can always choose to deactivate the Two-Factor Authentication as it is not set as mandatory.

Checking the status when you are allowed to administrate the users

In the **Users** tab, you can check the status of the Two-Factor Authentication option for each user (Figure 39). But you can also change the parameters for any user.

You can select the user tile and check the status of the activation for each user:

- Red: the Two-Factor Authentication is enabled and set.
- Orange: the Two-Factor is enabled but not configured by the user yet.

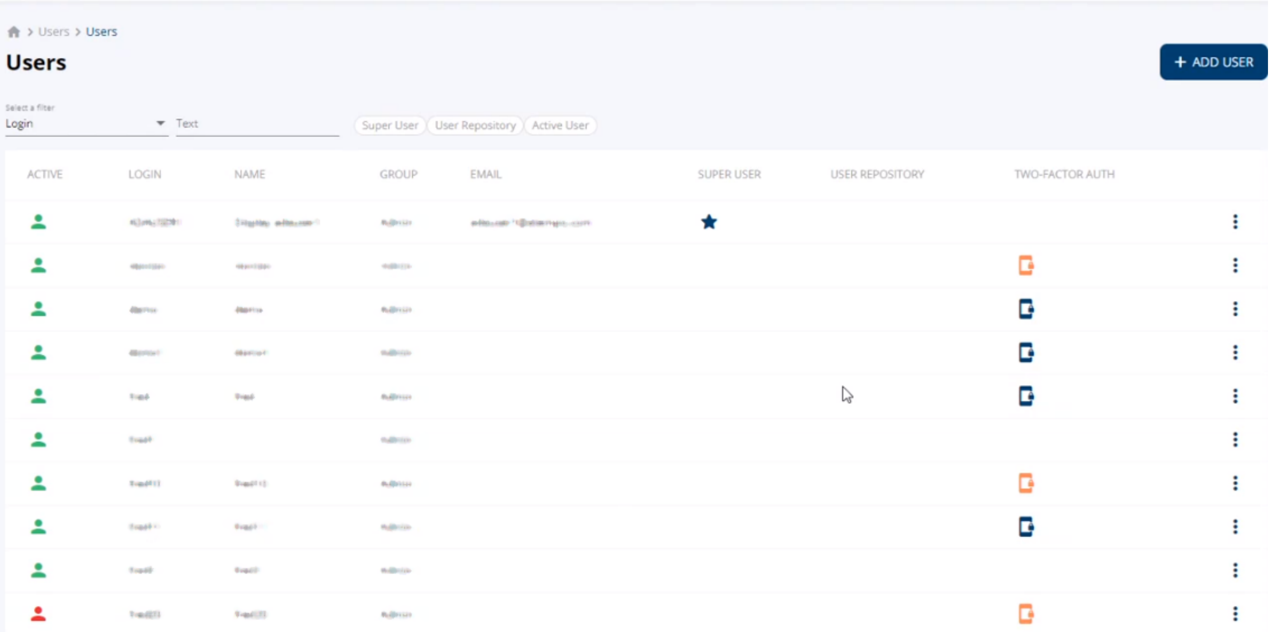

ACTIVE	LOGIN	NAME	GROUP	EMAIL	SUPER USER	USER REPOSITORY	TWO-FACTOR AUTH	
●	[redacted]	[redacted]	[redacted]	[redacted]	★			⋮
●	[redacted]	[redacted]	[redacted]				🟧	⋮
●	[redacted]	[redacted]	[redacted]				■	⋮
●	[redacted]	[redacted]	[redacted]				■	⋮
●	[redacted]	[redacted]	[redacted]				■	⋮
●	[redacted]		[redacted]					⋮
●	[redacted]	[redacted]	[redacted]				🟧	⋮
●	[redacted]	[redacted]	[redacted]				■	⋮
●	[redacted]	[redacted]	[redacted]					⋮
●	[redacted]	[redacted]	[redacted]				🟧	⋮

Figure 39: Users Two-Factor Authentication status

If you are a SuperUser, you can disable or reset the Two-Factor Authentication for a user without any code needed. To do so, select one of the rows. This menu appears (Figure 40):

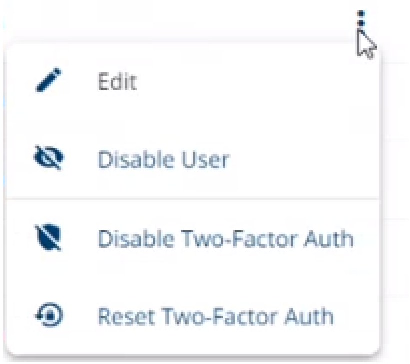

Figure 40: Rows menu

> **Note**: You can also configure Two-Factor Authentication on external users, for example AD or LDAP users.

CHAPTER 14 - Advanced Tasks

Tasks are automatic jobs that you can schedule or start manually. A task defines the scope of a Miria job, the source and destination of the data that it processes, its scheduling, and many other options.

Miria embeds basic maintenance tasks. Thus there is no need for you to create them for the product to be operational; however, you may customize basic tasks or create new tasks with specific characteristics.

Task Types

There are two categories of tasks:

- **Internal management tasks**
  The internal management tasks are automatic tasks (i.e., they launch jobs based on a schedule).
- **Data movement tasks**
  The data movement tasks let you select and move the data to and from the repository. You can launch these tasks either manually or automatically.

Internal Management Tasks

The following table (Table 4) describes the template that you can use to create internal management tasks:

Table 4: Internal management task template

Task	Description
Automatic deletion	Deletes the source data that matches a set of constraints once you have archived the source data. A deletion task only deletes the source data archived through an automatic archiving task. See Automatic Deletion Task for details.
Automatic retention	Keeps archives tidy by permitting the deletion of objects that are no longer needed from the Miria database. It only applies to archives fed by automatic archiving tasks, which replace the archived file with a stub after archiving. See Automatic Retention Tasks for details.

Table 4: Internal management task template

Task	Description
Maintenance	Scans the database and deletes useless data. You can choose the items to delete (e.g., expired instances, sub-job metadata, jobs, events, etc). A Maintenance default task is scheduled to perform all maintenance operations, except job and event deletion, on every first Sunday of the month at 12:00 P.M.; however, you can reconfigure it according to your needs, or create a new one. See Maintenance tasks for details.
Retention	Task automatically launched by the Maintenance task. Create this task only if you need to configure behaviors that will apply to all retention operations launched by the Maintenance task (e.g., to send an email each time it is performed). See Creating a Task for details.
Storage proxy maintenance	Archiving and retrieval jobs in client mode create temporary files in a cache space on a storage proxy. This space is normally deleted after the jobs complete, but if a job terminates on error or is canceled, some files may remain. The Storage proxy maintenance task scans the directories of the storage manager containers and deletes any temporary files that archiving and retrieval jobs in client mode have left. It requires no configuration.
Volume management on storage manager	Checks the space used on one or all of the storage managers. A high and a low water marks are defined in the storage manager configuration. See Recycling Triggered by Volume on Storage and Volume Management on Storage Managers Task for details.
Miria Database backup for PostgreSQL	Backs up the Miria database. See Database Backup Task for details.
Archive Report	Generates a report on the archive volume. See Archive Report Task for details.

Data Movement Tasks

The Data movement tasks can be launched either manually or automatically. They enable you to select and move the data to and from the repository.

Data Movement Manual Tasks

These are the templates that you can use to create a manually launched task:

- Archiving

- Copy
- Delete
- Move
- Retention
- Retrieval
- Synchronization

Create a manually launched task only if you need to configure a behavior (e.g., to send an email each time the task is performed) that will apply to all manual operations of that kind of tasks.
For the tasks created based on manual templates, you do not define any schedule.

See Creating a Task for details.

Data Movement Automatic Tasks

The automatic tasks launch jobs that are based on a schedule. For the tasks that you create based on automatic templates (Table 5), you can either set the Scheduler to launch them automatically or launch them manually.

This table describes the templates that you can use to create data movement automatic tasks:

Table 5: Data movement automatic tasks template

Task	Description
Automatic Archiving	Launches archiving jobs based on a schedule.
Automatic catalog ingest	Launches Catalog Ingest jobs to import external LTFS media into the Miria database.
XML Ingest	XML ingest tasks use XML files to run enriched automatic archiving. The XML files enable the customizing of archiving workflows and ingest interfaces. In addition to the archiving task scheduling, the XML ingest tasks enable you to perform these operations: • Archiving of specific files rather than whole directories. • Attributing metadata to files and folders. • Archiving into different archives and folders with a single run of the task. You are responsible for creating the XML files that launch the task. If they are not written directly with the Atempo format, you can translate them to the proper format applying a style sheet that you can request from Atempo Professional Services. XML Ingest Tasks for details.

Table 5: Data movement automatic tasks template

Task	Description
Interplay Dispatcher	Task available only for users of the Miria for Avid Interplay application. You must obtain a specific license from Atempo. See the Partner Applications Documentation for details on the Miria for Avid Interplay application. See Interplay Dispatcher Migration Task for details on configuring the Interplay Dispatcher task.

Creating a Task

The User Web Interface provides three ways to create a new task:

Creating a New Task from the Tasks Tab

To create a task

1. From the Web User Interface left pane, select **Tasks**.
2. Click the **+ NEW** button at the top right and select a kind of task among the list.
3. Set the configuration parameters set in General Configuration Parameters for Tasks

Duplicating an Existing Task

You can duplicate an existing task and then customize it to exactly fit to your new needs.

To duplicate a task

1. From the Web User Interface left pane, select **Tasks**. The tasks list opens.
2. Click the ⋮ button of one of them and select **Duplicate**.
3. Enter the name of the new task and click **OK**.
   The duplicated task displays in the tasks list. You can modify its parameters to suit your needs.

Organizing tasks

You can classify the tasks into projects. This help you navigate when there are many tasks. The tasks are displayed in a tree-like structure with the projects as a kind of folder, into which you can drag the existing tasks (Figure 41).

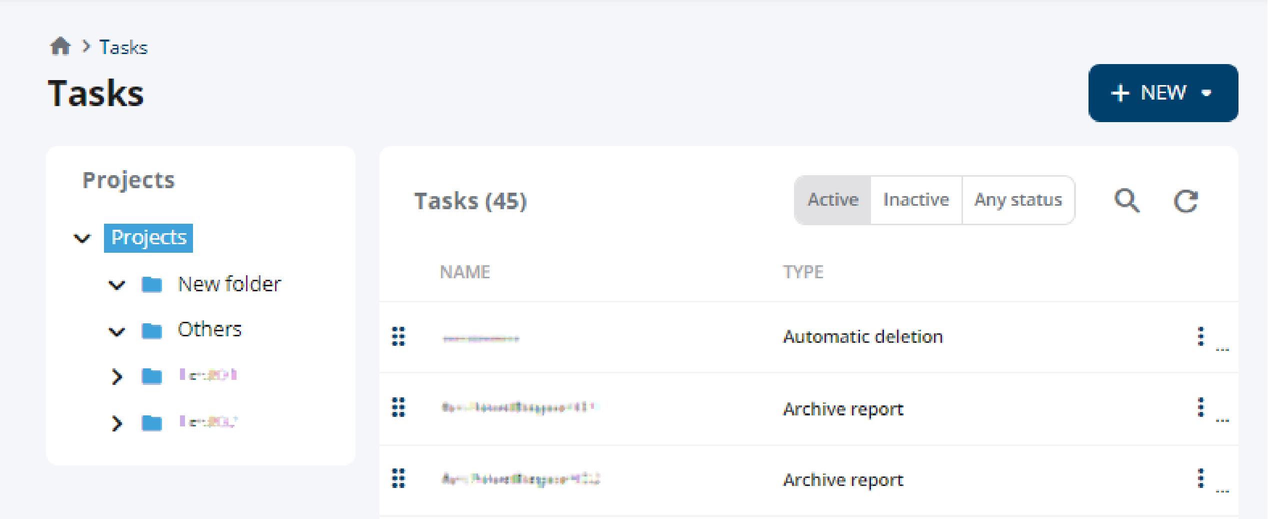

Tasks

Tasks (45)

NAME	TYPE
[redacted]	Automatic deletion
[redacted]	Archive report
[redacted]	Archive report

Figure 41: Tasks view

Manage projects

Here are the different options the projects pane gives you:

- Projects:
  - **Add folders.** Create a new folder and name it.
  - **Refresh.** Give the arborescence an update.
- Folders:
  - **Refresh.** Give a folder an update.
  - **Rename.** Change a folder's name.
  - **Delete.** Delete a folder.

Move a task

When you click on a folder, the tasks it contains appear on the right side of your screen.

You can drag and drop any task to move it from one folder to another.

This topic describes the running phases of Miria tasks.

How Tasks Run

Although you can launch tasks manually, characteristically the Scheduler launches them at regular scheduled intervals.

> **Note**: This process of selection alone can be quite lengthy. For example, if whole file deduplication has been configured on the archiving, the task must read every file selected to calculate the deduplication hash code on it. By default, the task waits until it has finished selecting files. It then gathers them into a single job and launches the job. If you have configured the parallelization parameters, the task launches several jobs before it has finished the selection of the files.

In the List of Jobs window the task displays on several lines (Figure 42).

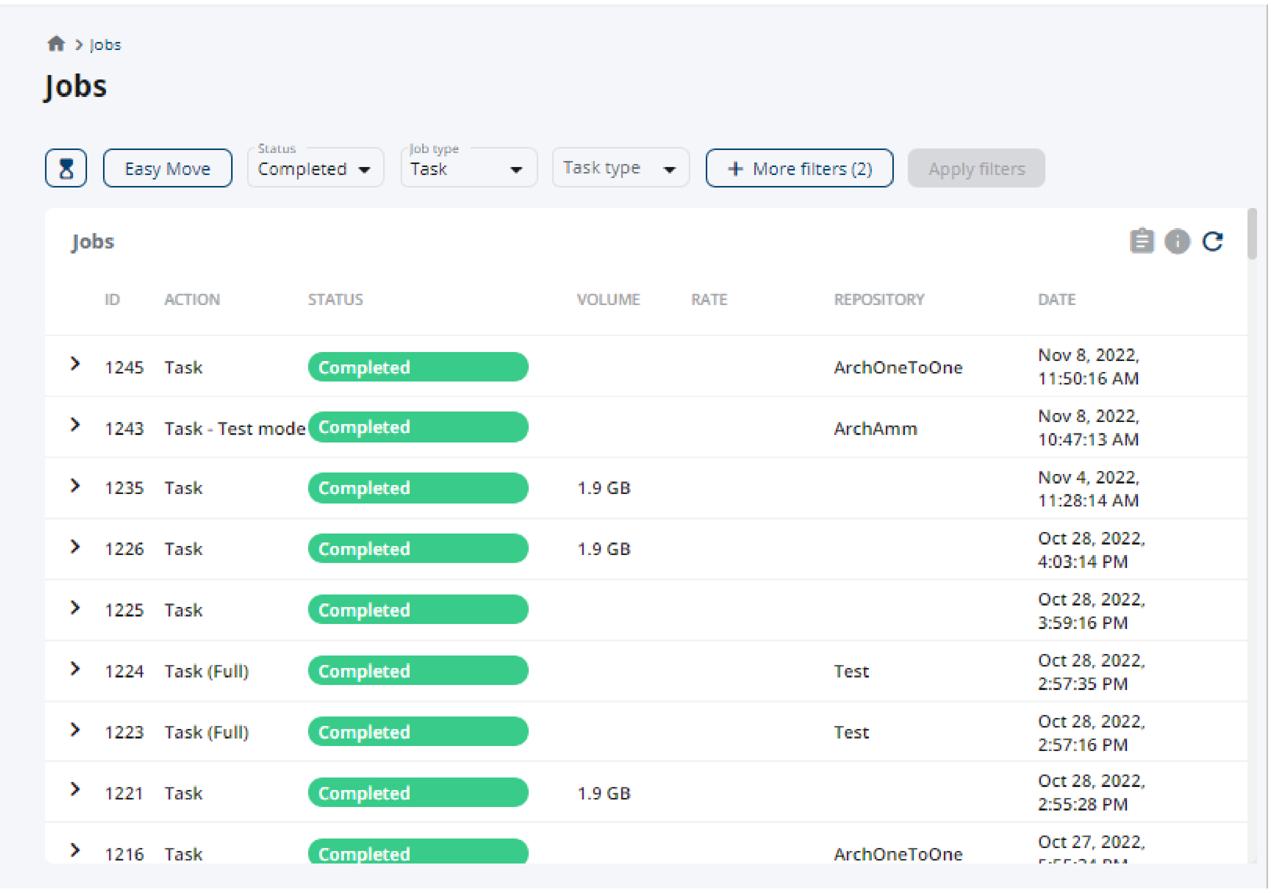

Jobs

Jobs

Easy Move | Status: Completed | Job type: Task | Task type | + More filters (2) | Apply filters

Jobs

	ID	ACTION	STATUS	VOLUME	RATE	REPOSITORY	DATE
›	1245	Task	Completed			ArchOneToOne	Nov 8, 2022, 11:50:16 AM
›	1243	Task - Test mode	Completed			ArchAmm	Nov 8, 2022, 10:47:13 AM
›	1235	Task	Completed	1.9 GB			Nov 4, 2022, 11:28:14 AM
›	1226	Task	Completed	1.9 GB			Oct 28, 2022, 4:03:14 PM
›	1225	Task	Completed				Oct 28, 2022, 3:59:16 PM
›	1224	Task (Full)	Completed			Test	Oct 28, 2022, 2:57:35 PM
›	1223	Task (Full)	Completed			Test	Oct 28, 2022, 2:57:16 PM
›	1221	Task	Completed	1.9 GB			Oct 28, 2022, 2:55:28 PM
›	1216	Task	Completed			ArchOneToOne	Oct 27, 2022,

Figure 42: Job List

Depending on the number of files to select, the use of deduplication, the size of the files, etc., the selection process can take long enough that the automatic task restarts before the previous iteration of the same task is finished, and before all the files have actually been archived. In order to prevent overlapping scheduled tasks from treating the same files, the default behavior is for only one instance of a task to run at any one time. Here is an example:

- The Scheduler launches an automatic archiving task, scheduled to run at midnight every 24 hours.
- The task takes 48 hours to select all the files to be archived.
- Twenty-four hours after the first launch of the task, the Scheduler is ready to relaunch it, but the first task is still collecting files.

If a second instance of the same task were to launch at its scheduled time, it would start to select the same files that the first instance has not finished selecting.

For this reason, the second instance of the task is not permitted to run until all the jobs launched by the first instance have completed. The first instance of the task remains in Running status until the last of its jobs is finished.

Any post-processing, such as the sending of e-mail notifications, thus occurs when both the task itself, and all of its jobs, are complete.

Testing a Task

Once you have created a task, you can preview whether it is configured correctly and if the files and directories included in the task are appropriate. Testing a task creates a task job, but does not archive or delete any data. The List of Jobs displays the task job.

To test a task

1. Click the **Tasks** tab in the left pane of your screen.
   The tasks view opens, and a list of tasks displays.
2. Click the ⋮ button and select **Start task in test mode**. A pop up opens.
3. Click **Yes** to validate. In the jobs list
4. Go to the jobs tab, in the jobs list, you can see your task. Its status is indicated. You can see the details, open the logs or download a report from this task.

Launching a Task Manually

You can manually launch all tasks configured in Miria. Therefore, you can launch a task at any time without waiting for the task to reach its start date as entered in the scheduler.

To launch a task manually

1. Select the Tasks tab in the left pane. The tasks view opens.
2. Click the ⋮ button of one of the tasks in the tasks list.
3. Select **Start task**.
4. Click **Yes** on the confirmation dialog box.
   The task creates and launches the corresponding job.

Canceling a run

You can cancel a run in the same manner as a job.

To cancel a current run

1. Select the **Jobs** tab in the left pane. The jobs view opens.
2. Click the ⋮ button of one of the current runs in the jobs list.
3. Select **Cancel**.
4. Click **Yes** on the confirmation dialog box.
   The run is then displayed in the History section of the List of Jobs with a Canceled status.

For all task types except Automatic Retention and Maintenance, Miria takes into account the Cancel request every five minutes. The effect of a Cancel request is to half the task selection process. No new job is created. Any jobs that the task has already launched continue running until completed, unless you cancel them too.

General Configuration Parameters for Tasks

This topics describes the configuration parameters and tabs that are common to all types of tasks.

General Parameters Common to All Tasks

From the Tasks List

In the tasks interface, choose a task from the task list and click the ⋮ button to access to the following options (This topic describes the running phases of Miria tasks.):

Table 6: Tasks list menu

Parameter	Description
Edit Task	When creating a new task, you have first to select a task type. See Creating a Task
Start Task	Immediate manual launch of the task. See Launching a Task Manually for details. **Note**: For backup tasks, you can start them in full or incremental mode.
Start the task in test mode	Creates a task job for preview, but does not archive or delete any data. See Testing a Task for details.
Duplicate task	Duplicates a task with the same configuration.
Disable	Lets you deactivate the task temporarily, without requering to delete it.

Table 6: Tasks list menu

Parameter	Description
Email	Gives you access to the list of users that are notified by email when a scheduled task is run. • **Add +** click this button in the upper right corner of your screen. Select user(s) or groups, and click on one of those options to set on them: • **None** they won't receive any email. • **Email** they will receive an email every time the task runs. • **On error only** they will receive an email only if the task ends with errors. • • above the users list, there is a toggle **Email only if an action was performed**. A task can run without actually doing anything (e.g., a scheduled task on an empty directory). If this toggle is selected, and if you selected the email option in the creation or edition of the user(s)/group(s), the email is sent only if the task actually accomplished an action. See Task-specific Parameters Receiving E-mail Notifications About an Automatic Archiving Task for details on obtaining an email notification on the status of an automatic archiving task.
Settings	Gives you access to the settings. You can change them directly here, or apply a template. Tasks inherit their settings only from the default settings. The only exception is the Archiving Policy setting which is inherited from the repository. If there is no archiving policy specified for the repository, then the task inherits its Archiving Policy setting from the default settings. If an setting is defined on a specific task through this tab, the value applies to the task on which you define it, overriding the inherited values. For the tasks, these are the available settings: Jobs, Email, and Security. Other types of settings are irrelevant. The Object Groups pop-up list enables you to select one of the object groups already existing in Miria. The task then inherits the subset of advanced settings that you have defined for the group and that are relevant to tasks. Modify the advanced settings for this task individually by clicking the Value field next to the setting that you want to modify and selecting a value from the list. See Default Settings and Settings Templates. **Note**: For tasks that do not require configuration, such as manual tasks, this tab is grayed out.

Table 6: Tasks list menu

Parameter	Description
Report history	You can display and download the volume reports from the Report History button.

When Creating or Editing a Task

Table 7: General parameters common to all tasks

Parameter	Description
Task type	When creating a new task, you have first to select a task type. See Creating a Task
Task name	Name that identifies the task within Miria. **Important**: When naming a new automatic task, do not use these terms (in either lowercase or uppercase) as they are reserved for the template name: - ARCHIVING - RETRIEVAL - COPY - MOVE - DELETE - SYNCHRO - RETENTION

Common Tabs

The configuration tab

The Configuration tab enables you to set the configuration parameters for the individual Miria tasks, as described in Task-specific Parameters.

For tasks that do not require configuration, such as manual tasks, this tab is grayed out.

The Scheduling tab

The schedule allows you to define the regular times at which the automatic tasks must be started.

The Scheduling tab enables you to define the regular times at which automatic tasks are to be launched. It is active for all tasks, except basic archiving, retrieval, copy tasks...

An error message displays if no occurrence in the month, days, and hours, minutes are set when the Scheduler is activated.

To set the scheduling:

1. Click the toggle **ENABLE SCHEDULING**.
2. Select a week day and an occurrence in the drop down menu. It specifies the day(s) of the week when automatic tasks must be started and at which frequency.

3. Select a time and check **Every 5 minutes**, or **Choose an interval**. It specifies the time when the automatic tasks must be started.

> **Note**: For backup tasks, you can schedule both full, and incremental modes.

Options Tab

The options tab enables you to associate run timeframes and run locks with the task, as well as any pre- or post-processing scripts.

The pre- and post-processing scripts must be located on the agent defined in the task.

This table describes the parameters of the Advanced tab:

Table 8: General parameters common to all tasks

Parameter	Description

Commands	This option enables to enter commands manually to add custom settings to the archiving job. Click the toggle to enable it. Two fields appear:
	1. **Pre-processing**:
	Field to be filled in to enable you to launch a script before the launch of the associated job.
	It must always contain the full path of script to run and its interpreter.
	2. **Post-processing**:
	Field to be filled in to enable you to launch a script after the associated job finished, through the following keywords:
	• {Job_Number}: Job ID associated to the task.
	• {Db_Name}: name of the environment on which you are working.
	• {Tpl_Status}: retcode of the job, if the retcode is different than 1, there is an error.
	They will be replaced automatically upon execution. It must always contain the full path of script to run and its interpreter.
	Example:
	`D:\miria\Binary\Bin\ada_perl.exe D:\miria\Custom\custom_action.pl -job_id{Job_Number} -db_name {Db_Name} -retcode {Tpl_Status}`
	Will be replaced by:
	`D:\miria\Binary\Bin\ada_perl.exe D:\miria\Custom\custom_action.pl -job_id 12345 -db_name miria -retcode 1`

Run options.	

Table 8: General parameters common to all tasks

Parameter	Description
Run timeframe	Period during which a task is permitted to run or prevented from running. This feature enables you to prevent tasks from running at times when you know there is heavy use of network resources for other operations. Click to select a run timeframe among those configured in Miria. To associate a task with a run timeframe, you must have configured it in the list of Run Timeframes interface. See Run Timeframes
Run lock	Limits the number of Miria tasks of any kind that can run simultaneously. Click to select a run lock among those configured in Miria. To associate a task with a run lock, you must have configured the Run Lock in the list of Run Locks interface. See Run Locks for details.
Next task	Select the next task among the drop-down list.
Maximum number of simultaneous runs	Number of times the current task can be run simultaneously. You can modify this field only for manual archiving and retrieval tasks. A user can launch one manual archiving and then launch a second before the first has completed. If you do not want to limit the number, leave this field at 0. For all types of tasks other than manual, this parameter is set to 1 and you cannot modify it. The Maximum number of simultaneous runs parameter differs from the Run Lock notion, in that it concerns only the task that is being configured. You can invoke a run lock on tasks of disparate types to prevent more than a specified number of tasks of any kind from running concurrently.
Wait when the maximum number of runs is reached	Select this box if you do not want additional tasks to be canceled (e.g., if the Maximum number of simultaneous runs is 3, the fourth task is not canceled.) If the box is selected, Miria waits for the first three tasks to complete before running the fourth task.

The Report History Tab

The Report History tab displays information on the generated reports (e.g., by the Archive Report task).

This table describes the columns of the List of available reports:

Parameter	Description
Date	Time and hour at which the report was generated.
Name	Name of the report in the `Archive_date-time.ext` format.
Size	Size of the generated report that is in PDF format.
Download button	Opens the report in the associated browser. Then, you can download the PDF file from the server to your local machine. For the Archive Report task, you can also obtain a volume report from the `ADA\Report` directory using the Windows Explorer.
Delete button	Click this button to delete the selected report from the server.

Task-specific Parameters

Some tasks require particular configurations (e.g., it may be necessary to specify the platforms and directories that the task must scan for files to process, or you may want to set constraints on the files, such as age or size).

For all automatic tasks, see General Configuration Parameters for Tasks in the table General parameters common to all tasks, to know how you can name them.

Define these settings in the Configuration tab of the task Properties pane. Each task has its own requirements, so the tab displays different fields depending on the task that you are configuring.

See General Configuration Parameters for Tasks for details on the configuration parameters that apply to all tasks.

Automatic Catalog Ingest Task

This table describes the fields that you can complete to configure the Automatic Catalog Ingest task:

Parameter	Description
Source.	
Ingest Type	Select Media/LTFS from the list.
Storage Manager Container	**Only for Media Manager storage manager container.** Select from the list, an ingest Storage Manager Container into which you want to ingest the media.

Parameter	Description
Retention	Select from the list a retention that you want to apply to the data imported from the media. Click the **+** button to add a retention.
Filters.	
Library	Library in which Miria stores the archived data. The library alias (if any) displays between parenthesis. If you have completed the Barcode Selection filed, this field is ignored. This parameter is mandatory if you do not specify a Scratch Media Group or if you have not completed the Barcode Selection filed. Click the Browse button to select the library. Click the Minus (-) button to reset the field.
Media Type	Type of the media that you want to use for this storage manager container. Complete this field only if the library may contain media of several types (e.g., `LTO-6` and `LTO-7`), and that you want to use only one type. If you have completed the Barcode Selection filed, this field is ignored. If you have not completed the Barcode Selection filed, this field becomes optional as it defines the media type identifier. If this field remains undefined, all the orphan media belonging to the selected library are ingested into the selected archive. Click the Select button to display the list of compatible media types. You can select either a Media or a Class.

Parameter	Description
Define filters on barcodes	Optional. Select **Include** from the list to ingest media with a specific barcode range. Select **Exclude** to exclude a barcode range from ingestion. Enter the barcode range in the text field in the form of a pattern, using these wildcards: • The * means any alphanumeric character any number of times. • The ? means any alphanumeric character once. • The | separates several possible pattern options. The expression must contain at least one * or ?. For example: • A005?? includes/excludes any media with a six character barcode beginning with the string A005. You might use this, for example, to select media from A00500 to A00599. • *L4 includes/excludes any media with a barcode ending in L4. You might use this, for example, to select only media of LTO4 type. • 162*|WV* includes/excludes any media with either a barcode beginning with the string 162, or a barcode beginning with WV.
Target.	
Existing repository	Imports the media into a repository which already exists in Miria. Select **Existing repository** and click the ⋮button to choose a repository from the arborescence. The data will be imported at the repository root in a folder named after the media barcode. If the selected repository is not associated with a Media Manager storage manager container, you will have to complete the Storage Manager Container field.

Parameter	Description
New repository	Imports the media into a new repository. **Project**: Enter a repository project name or click the ⋮ button to select one from an arborescence. Then, if you don't select a reference repository, the data will be imported at the organization root in a repository created automatically and named after the media barcode.
Reference repository	**Optional.** Click the ⋮ button to select from a list the repository that will be used as a template to create the new one.

Automatic Deletion Task

These tasks delete the source data that matches a set of constraints once the source data has been put in repositories. Only the source data put in a repository through an automatic archiving task is deleted.

This table describes the parameters displayed in the Source & Target tab:

Parameter	Description
Source.	
Storage platform	Select the storage platform in which you want to delete data.
Location of data to delete	Enter the path of the data that Miria can delete.
Target.	
Repository name	Select the target repository
Full path auto-generated	Select the repository for which source data is deleted once it has been archived.

Parameter	Description
Archive repository	Possibility to store archived objects in archive repositories.

Automatic Retention Tasks

If you delete a stub from a file system, and if there is no retention set on its corresponding instance(s) in the repository, these instances can remain in the database indefinitely. The automatic retention task provides a method to assign a retention period to these orphaned objects. Then, the first automatic maintenance task to run on the repository after the retention period has expired can delete them.

The automatic retention task works in this way:

- It first runs a check to ensure that all stubs in the file system correspond to object instances in the Miria database.
- If it finds an object in the database that no longer has a corresponding stub in the file system, and the object database instances have no defined retention period (the associated retention was set to Without), then the retention period defined in the automatic retention task is applied.
  If the automatic retention task does not have any retention periods defined, the retention is set to expire on the current date.
- Then, the next maintenance task to run after the expiration of the retention period eliminates the orphaned database instances.

> **Important**: If a stub is renamed, it is considered as lost, and the retention task processes the corresponding instance in the database.

This table describes the Source & Target tab parameters for an automatic retention task:

Parameter	Description
Source.	
Storage platform	Displays the platforms configured in Miria. Select the platform that hosts the data to put in repositories.
Data for applying retention	Enter the root of the path from which the automatic retention task runs its check to see that all object instances in the database have corresponding stubs in the file system. You can enter multiple paths.
Target.	
Repository name	Select the target repository

Parameter	Description
Full path auto-generated	Opens a list of all repositories to which you have rights as a logged-in user. Select a repository to associate with the task and validate by clicking the check mark.
Archive repository	Absolute destination path. Select the destination location to associate with the automatic retention task.

This table describes the Options tab parameters for an automatic retention task:

Parameter	Description
Retention	Opens the List of Retention Periods window that displays the retention periods configured in Miria. Select the appropriate retention period and select the check mark. If you select None, the next maintenance task deletes files. The Apply on Stub and Apply on Object options enables you to refine the retention by applying a retention period to a file from a repository in these circumstances (they can be activated individually or together): • **Apply on Stub.** When the stub is no longer present on the file system. • **Apply on Object.** When the object is no longer present on the file system.

Automatic Storage Repack Task

The repack task is used to defragment a media storage. When you launch an archiving job, Miria creates a .pax file containing all the files that will be put in repositories.

When a file from the repository is deleted, it is no longer referenced in Miria's database but still exists in the .pax file. The purpose of the repack task is to recover all the files that are still referenced in Miria. A new .pax file is then created which will contain all the files except those which have been deleted from the repository.

The volume of the repack task is defined by the following parameters:

Parameter	Description
Nb. Expected Objects	Total number of media present on the storage manager container.
Nb Objects	Number of media selected for repack.
Volume Expected	Total storage volume.
Volume	Actual volume of selected media.

> **Note**: When you create an automatic storage repack task, there are a few specific actions that you need to do:
> - Define the Source Storage Manager Container Name (Only a FileStorageContainer can be selected for a repack task. This container will be used to calculate the fill rate of each media).
> - Enter the Data Fragmentation Ratio percentage. This threshold will be used to select the media to repack.
> For example: A .pax file with a size of 150 MB contains File 1 (100 MB) and File 2 (50MB). When File 2 is deleted from the repository, there is (150-100) = 33% of free space to recover on this media. This percentage is the data fragmentation ratio and allows you to select all the media that have 33% or higher of free space to recover inside the .pax file.
> - (Optional) Activate Delete source media after repack. The media selected for the repack job will be deleted if the repack job is completed without errors.
> - (Optional) Set the maximum number of jobs running in parallel. By default, the task creates one job per media.

Database Backup Task

The Miria database backup task backs up the Miria database.

It is the PostgreSQL database backup task.
See Miria administration guide for more details.

Maintenance tasks

The maintenance task enables you to delete both the repositories objects that have reached the end of their retention period and their associated jobs and events.

The preset Maintenance task performs all maintenance operations, except job and event deletion, on every first Sunday of the month at 12:00 P.M.; however, you can reconfigure it according to your needs or create new maintenance tasks with different parameters.

Parameter	Description
Maintenance.	

Parameter	Description
Check running jobs	Click this toggle to process the jobs that have been running for more than 24 hours: • Jobs on creation or running switch to the Terminated on error status. • Jobs on queue restart or switch to the Terminated on error status, if they do not start.
Delete sub-job metadata	To delete the metadata associated with a sub-job. As the data is , the metadata is linked to an instance and already present in the database. It is recommended that you delete job metadata regularly to optimize database space.
Delete in recycler	To delete permanently the repositories folders that you removed from a repository.
Check expired instances	To search the database for all expired instances, and create a retention job for each repository. A retention job deletes expired repository instances from the storage and the database.
Check expired SML instances	Click this toggle to delete the SML (Storage Manager Layer) archiving instances that are past their retention period. This is only relevant when using the Miria (Messaging) software.
Delete temporary tables	To delete some temporary database tables that Miria may create while operating. It is recommended that you delete such tables regularly to optimize database space.
Delete archiving job details	To delete details, messages, or alarms that are past their retention period.
Delete copy / synchronization job details	To delete details, messages, or alarms that are past their retention period.
Delete repository tiering job details	To delete details, messages, or alarms that are past their retention period.
Delete statistics on jobs	To delete details, messages, or alarms that are past their retention period.

Parameter	Description
Delete jobs	To delete jobs that are past their retention period: 1. Click the toggles for the types of jobs (see the list below), that you want to delete. 2. Click the Job and Event retention toggle. 3. Choose how long you want to keep the jobs In the Job retention fields. The default value is one month. 4. Choose how long you want to keep the events related to the jobs in the Event retention field. This value must be less than or equal to the job retention. The default value is one month. Type of jobs that you can delete: • AER collection • Repository tiering • Archiving • Catalog ingest • Change LTFS volume lock • Change LTFS volume name • Copy • Create directory • Delete • Device scan • Drive diagnostic • Drive performance • Drive unmount • Ejection request • Library unknown media scan • Media deletion • Media duplication • Media mount • Media recycling • Media scratch • Media verify • Move • Rename • Retention • Retrieval • Storage manager integrity check • Synchronization • Synchro. digest • Task

Parameter	Description
Delete events	To delete events that are past their retention period: 1. Click the toggles for the types of events (see the list below), that you want to delete. 2. Click the event retention toggle. 3. Choose how long you want to keep the events related to the jobs in the Event retention field. The default value is one month. Type of events that you can delete: • AUDIT TRAIL • CRITICAL • DEBUG • DEBUG STACK • ERROR • FATAL • INFO • STACK • STOP/DIE • SUCCESS • WARNING

Archive Report Task

The Archive Report task generates a volume report for a selected repository path.

Volume Reports

Configuring the Archive Report Task

This table describes the Configuration tab fields that you must complete to configure an Archive report task:

Parameter	Description
Source.	
Repository name	Select from the drop down list a name for the repository.
Repository path	Click the ⋮ button to select the absolute path of the selected repository for which the volume report is generated.
Layout. Enables you to modify the PDF file layout.	
Format	Select the A4 or the US letter format.

Parameter	Description
Orientation	Select the portrait or landscape orientation.
Heading logo	Image that displays as the output file header. You can either: • Leave the field empty. The PDF file does not display any image. • Enter the `{default}` keyword. The PDF file displays the default image included in the Miria distribution. **Or** • Click the ⋮ button to choose a heading logo in the arborescence. and browse for your own image. You can use any image that is in .jpg, .tif, .png, or .gif formats.
Header title	String that displays as the report title. You can enter a free text plus these keywords: • {Job_Number} • {SubJob_Number} • {Archive_Name} • {Archive_Path} • {Archive_Global_Path} • {Archive_Comment} The PDF file displays the associated value for each keyword that you have defined.
Report information. Enables you to modify the PDF file contents.	
Media details	Details on the media. These are the valid values: • **None.** Default value. The report does not display any media information. • **Media By Job.** The report displays the names of media, associated with the jobs that have put the objects in repositories. • **Folders By Media.** The report displays the associated folders for each media that you have defined in the source pane and that is involved in the archiving of objects.

Parameter	Description
Metadata	Archive metadata. These are the valid values: - **None.** Default value. The report does not display any metadata. The Metadata pane displays grayed out and you cannot access it. - **Job.** The report displays the metadata collected from the jobs that have put the objects in repositories. - **Object.** The report displays the metadata collected from the objects or instances. - **Object and Job.** The report displays the metadata collected from the objects or instances and the jobs.
Folder details	Details on the folder. These are the valid values: - **No.** Default value. The report displays only the source path and the cumulated Volume or number of files. The Directory Rules pane displays grayed out and you cannot access it. - **Yes.** The report displays: – Recursively, all the folders that you have defined in the source pane. – For each folder, the volume and the number of files.

Volume Management on Storage Managers Task

The Storage Manager Name field of the Volume management on storage managers task displays all the storage managers configured in Miria. You can select either all of them or only one at a time. To run the task on more than one, but not all the storage managers, configure a separate task for each.

The task scans the storage manager. If the *Use Volume Level to Trigger Retention Job* parameter is enabled on the storage manager, the task checks the *Task High Water Mark* value. If this value is attained or exceeded, the task deletes files on the storage manager until the *Low Water Mark* value is reached, or until there are no more eligible files to delete.

See Recycling Triggered by Volume on Storage.

XML Ingest Tasks

In contrast to basic automatic repository tasks, which do not permit association of metadata, the XML ingest task uses XML files to run enriched automatic archiving. It reads the XML files in a designated directory and parses them according to the rules set out in an XML schema definition (.xsd file) provided with Miria (`ada_ingest.xsd`). The ada_ingest.xsd file describes the structure to follow to create a valid XML file to be used by Miria.

During the installation of the Miria server, the setup installs the .xsd file and a sample XML ingest file in the Perl subdirectory (e.g., on Windows, they are located in `C:\Miria\Binary\Bin\Perl\Miria\XML`). The sample XML file is `ADA_Ingest_sample.xml.`

You can also get both these files from a browser on the Miria server using this syntax:

```
http://<Miria_server_name>:<port>/xml/ADA_Ingest.xsd
http:// <Miria_server_name>:<port>/xml/ADA_Ingest_sample.xml
```

For more information on the different parameters in the sample file, see the second table below.

In addition to the scheduling of repository tasks, the XML ingest task permits you:

- To put specific files in repositories rather than whole directories.
- To attribute metadata to files and folders.
- To put data into different repositories and folders with a single run of the task.

You are responsible for creating the XML files that launch the task. If they are not written directly with the Atempo format, you can translate them to the proper format using a stylesheet that you can request from Atempo Professional Services.

It is assumed that all repository objects invoked in the XML files have been created in Miria. Metadata and repositories must already exist within Miria. Source directories of files for repositories must exist within the file system.

This table describes the parameters displayed in the Configuration tab of a XML ingest task:

Parameter	Description
Configuration.	
Storage platform	The list displays all storage platforms configured in Miria. Select the one to be used for the task. This is the machine hosting the XML files to be read. If the XML files must run archiving tasks on machines other than this one, the storage platform must have access rights to the other platforms.

Parameter	Description
Directory	Enter the root path which contains the XML file(s) that the XML ingest task must read and parse for the launching of the repository task. **This must always be a path, never a file name.** You can use the Select button to the right of the field to select the path, which is local to the machine that you specified in the platform. The XML ingest task processes all the files having the .xml extension.
Recursive scan	If you select this box, the XML ingest task also processes all the files having the `.xml` extension in the subdirectories of the directory entered in the previous field.
Translator	Absolute path and name of the translator .xslt file. Choose the file by clicking the Select button. If this field is populated, the XML files are not parsed according to the XML schema definition that Miria provides in the `ada_ ingest.xsd` file. Instead, an `.xslt` file translates your XML files into XML files that conform to the ada.xsd. Like for the XML files themselves, it is your responsibility to create and supply the .xslt translator. On request, Atempo Professional Services can also create it.

Parameter	Description
XML processing report	Once the XML files have been correctly parsed and read, and the associated tasks have been launched without error, the used XML files in the XML file directory must be moved in a way that prevents rescanning the next time the task is launched. These are the ways in which you can perform this modification: • If the Processing Report field is not populated, the XML files are simply renamed in their original directory to prevent their rescanning. The file extension is changed from `.xml` to `.out`. The next scan ignores the files having the `.out` extension. • If the Processing Report field is populated, but the Translator field is not populated, then the XML files are moved from their original XML file directory to the Result Directory specified in the corresponding field. In the new directory, they are also renamed in this way: `originalFilename_jobNumber_ADA.xml.` • If both the XML Processing and the Translator fields are populated, then two file sets are moved into the Result Directory. The first set is as described above, and contains the translated files, named `originalFilename_jobNumber_ADA.xml`. The `_ADA` suffix indicates their conformity with the `ada_ingest.xsd`. The second set contains the original, untranslated XML files, named `originalFilename_jobNumber.xml` (without the `_ADA` suffix). **Result Directory.** Path to which the XML files that have been used in an XML ingest task are to be moved after the task has

Parameter	Description
	completed correctly.

Parameter	Description
XML error report	In some cases the XML ingest task does not complete correctly. This happens when: • The XML ingest file is malformed. • The XML ingest file does not conform to the XSD file. • If you used the Translator, it could be incorrect in itself, or not translate correctly to a form that ada_ingest.xsd can read. • The files specified for repository do not exist. • Repositories requested in the XML files do not exist. • The metadata to be associated with the files put in repositories do not exist. • The storage platform machine might not have access rights to all the machines that the XML ingest task must scan for files to put in repositories. In these cases, the Event logs display error messages to help you analyze the error cause. Additionally, if the XML ingest task rejects the XML files due to malformation or non-conformity, you must modify the rejected XML files in the XML file directory so as they are not rescanned the next time the XML ingest task runs. Perform this modification in either of these two ways: • If the Error Report field is not populated, the XML files are simply renamed in the original ingest folder. The file extension is changed from .xml to .err. The next scan ignores the files having the .err extension. • If the Error Report check box is selected, the XML files are moved from

Parameter	Description
	their original directory in the XML file directory path to the Reject Directory specified in the corresponding field. In the new directory they are also renamed in this way: `originalFilename_jobNumber.xml` (without the `_ADA` suffix). Thus, you can easily rename or return the files to the ingest directory after you have corrected them. **Reject Directory.** Path to which the XML files that have been used in an XML ingest task are to be moved after the files have been rejected.

This table describes the parameters displayed in the sample XML ingest file:

Parameter	Description
ada_ingest_v1	Allows you to configure the ingest.
StartComboAtIndex	By default, the ComboBox in Miria goes from 1 to N, while in many client software combos start from 0 to N. The StartComboAtIndex parameter allows to automatically increment the index of a combo box. **Example**: If you generate an xml by programming from a third party application, you will put index 1 for Value 2 in the metadata associated to the ada_file_ingest or ada_folder_ ingest file. Miria StartComboAtIndex parameter will automatically replace "Value 2" by "Value 1".
ada_file_ingest	Allows you to create and configure files.
ada_folder_ingest	Allows you to create folders and add metadatas. It is not possible to browse directories in the XML ingest, in this case you must have an ada_file_ingest tag for each file in the directory.

Parameter	Description
metadata_folder_action	Allows you to add, merge or delete metadatas on the existing folder. To have the exact values of this parameter, see the DTD in Binary/Bin/Perl/ADA/XML/ingest.xsd.

CHAPTER 15 - Parameters

In the Web Interface, click the **Parameters** tab.

The parameters interface is divided in different sections (CHAPTER 15):

- **Default settings** Access the settings and manage them.
- **Settings templates** Create templates to define settings by default.
- **License** Access your license information.
- **Retentions** Manage the retentions.
- **Run lock** Manage the run locks.
- **Run Timeframes** Create and manage the run timeframes.
- **Metadata** Create and manage metadata.
- **SMTP server** Specify an SMTP server.
- **Syslog server** Connect Miria to a Syslog server, and export logs towards it.
- **Media Manager** Manage libraries and media.

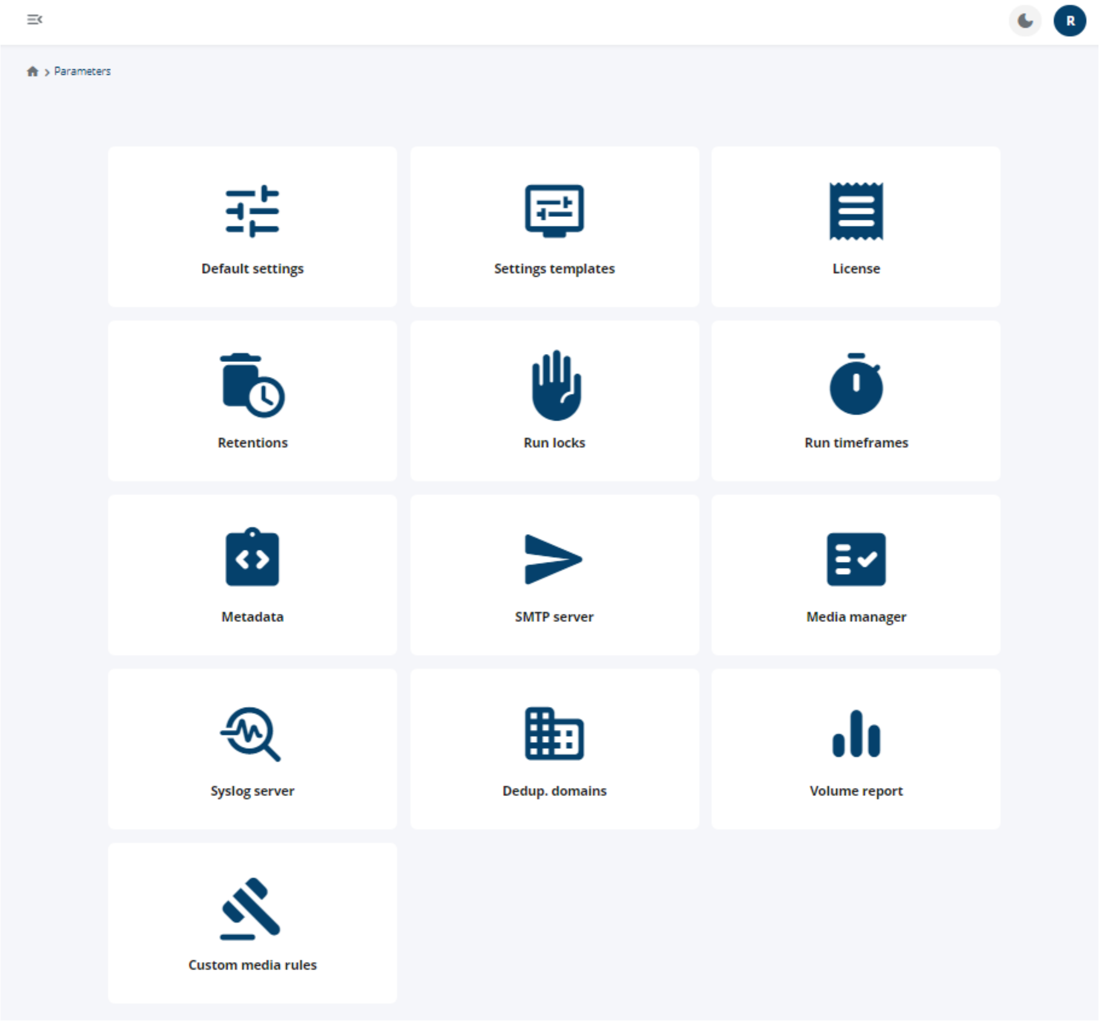

Figure 43: Parameters

Default Settings

Security Settings

This table describes the security settings in Miria:

Setting	Possible Values [default]	Description
Free login without password	No|Yes	Permits you to login without a password even if you are not accessing Miria from a Trusted Domain.
User and Domain are case sensitive	No|Yes	Determines whether capital and small letters are differentiated at the moment of login. For example, if this setting is set to `Yes`, the passwords "dagobar" and "Dagobar" are considered different.
Deprecated ***Ignore User during an HSM query***	No|Yes	*When using HSM, these are the valid options:* • *Yes. Enables you to retrieve files by double-clicking the file stub, even if you are unknown in Miria. The rationale is that if you have Active Directory permissions to access the stub automatically, you have the same permissions on the archived file represented by the stub.* • *No. Makes it mandatory if you are in an HSM environment to be created and known within Miria to be able to retrieve an archived file/directory through its stub.*

Setting	Possible Values [default]	Description
Calculate digest during selection	`Quick Digest`, `MD5`, `SHA-1`, `SHA-256`, `SHA384`, `SHA-512`, `xxHash32`, `xxHash64`	For use when archiving does not take place immediately after selection of files to archive (e.g., when the selection is performed through the graphical interface). Secures against alteration of files between the moment they are selected for archiving and the moment they are effectively archived. Calculates a digest on the files at the moment of their selection. Recalculates digest at the moment the archiving is launched. Comparison of checksum permits the detection of modification of files. Quick Digest default option uses a calculation based on the date of last modification of the files and their sizes. You can perform this calculation rapidly and thus give fairly good protection against file alteration. The other digest calculation algorithms are more secure, but slower to use.
Calculate digest during archiving	[None], `MD5`, `SHA-1`, `SHA256`, `SHA-384`, `SHA-512`, `xxHash32`, `xxHash64`	For use when archiving simultaneously on multiple storage manager containers. It permits verification of consistency between the different writings. The initial digest on files is calculated at the moment of archiving on the first storage. When archiving on the second, third etc. storage, a digest is calculated, and compared to the initial digest.
Advanced administration of archive by its owner	No|Yes	Lets the Super User enable the owner of a project archive to grant the administration rights of his project archive to another standard user. If this setting is set to `Yes`, the Administration permission is displayed for the project archive if you are connected as the owner. You must grant the owner of the project archive administration rights over his own archive before he enables other users to administer it.

Setting	Possible Values [default]	Description
Two-Factor Authentication mode	Allowed, Forbidden, Required. By default, Allowed is selected.	For a matter of safety, when this option is activated, you have to connect through two steps. The password, and the TOTP code (Time-based One Time Password). The TOTP code is generated when you scan the QR code that displays for the first connection. For the next connections, you don't have to scan a QR code anymore. The TOTP code generates automatically. Each TOTP code has a time limit of thirty seconds. **This time limit requires the server to be on time, and synchronized with an external NTP servers.** You can configure the Two-Factor Authentication and choose if it is possible to activate it or not: • **Allowed**: choose whether you want to activate it or not. • **Forbidden**: you cannot activate the Two-Factor Authentication. • **Required**: when first log in, you have to configure the Two-Factor Authentication to connect to Miria.

Interface Settings

This table describes the Interface Settings in Miria:

Setting	Possible Values [default]	Description
Language	[English], French, Japanese, Chinese, Traditional Chinese, Korean	Sets the language for the interface.

Setting	Possible Values [default]	Description
Specific settings on	User groups, Users, Platforms, Archives, Tasks, Applications	Enables or disables the advanced settings on classes of these administration objects: • G: User Groups • U: Users • P: Platforms • A: Archives • T: Tasks • APP: Applications By default, all are enabled. When disabled, the Default settings tab is not displayed in the configuration window for the object, any Settings positioned at the level of the object are ignored, and there is no blocking of inheritance at that level.

Jobs Settings

This topic describes the Miria jobs settings.

Global Settings

Table 9:

Setting	Possible Values [default]	Description
Disable scheduler to prevent new jobs/tasks from running	No|Yes	Enables preventing the scheduler from launching new jobs or tasks. When set to Yes, the scheduler still runs as a background process, but it does not launch any new job or task. This option is useful, for example, for administrative procedures that require stopping a process (e.g., to apply a retention task needs to stop all retrieval tasks).

Table 9:

Setting	Possible Values [default]	Description
Usage Rules for AMM Drives	[Automatic], Force Local Platform, Storage platform	With Media Manager storage managers, it enables you to specify what drives to use for archiving jobs by selecting the hosts to which they are attached. This option is useful, for example, in contexts where a company network is very dispersed, and you want to limit network transfers over large distances. • **Automatic.** With the default setting, media drives are allocated automatically for archiving tasks. The system makes use of the hosts and drives that are available and there is no constraint or guarantee regarding which host is used. • **Force Local Platform.** This setting enables you to force Media Manager to use only the local host and drives attached to it. Archiving is done only from the local machine. If Media Manager or the desired drives are not active on that machine, the job returns an error. • **Storage Platform.** Lets you lock the archiving to a particular host. If that host is unavailable, the job returns an error. Opens a window with a list of the platforms and agent pools configured in your Miria installation. Select a platform and validate by clicking save. This option does not permit selection of a NAS.
Storage proxy	[No Default], List of existing proxies	Specifies the storage proxy that you want to use. A storage proxy is useful when you: • Archive or retrieve in Client mode. • Archive data with Dell EMC Celerra/BlueArc HSM Filter Driver. You must first create a One-to-One storage manager container proxy.
Video proxy location	[No Default], List of existing storage proxies	Specifies the location to use to store the low-resolution version of the archived video assets. You must first create on the Miria server a One-to-One storage manager container proxy. Set this setting if you want to preview video assets before retrieving them.

Archiving Settings

Setting	Possible Values [default]	Description
Archiving mode	Explicit Explicit Strict [Class] Incremental Class	This setting concerns only archiving that is launched manually by dragging the files. Tasks do not take into account this setting. It specifies how directory contents are analyzed when you select them for archiving in the graphical interface. These are the available archiving modes: • **Explicit.** The graphical user interface scans the directory contents and compiles a preliminary file list. Then, the GUI sends this list to the Miria agent or server, which archives each file on the list. • **Explicit Strict.** The graphical user interface scans the directory contents and compiles a preliminary file list. Then, the GUI sends this list to the Miria agent or server, which archives each file on the list. These two scenarios may happen: • No action to perform after archiving has been defined. Then the GUI performs a test on the files over which it has the read permissions. • An action to perform after archiving (e.g., file deletion) has been defined. Then the GUI performs a test on the files over which it has both the read and delete permissions. • **Class.** When the Miria server or agent starts the archiving job, it directly scans the directories in the background. This mode leaves the graphical interface free and prevents it from freezing if a directory contains a great number of files to be archived. • **Incremental Class.** Same as the Class mode, except that only the files that have been modified since last archiving are rearchived. A file is considered modified if either its size, modification date, or digest has changed.

Setting	Possible Values [default]	Description
Archiving mode	Explicit Explicit Strict [Class] Incremental Class	These are the limitations for the Class mode: • Since scanning takes place in the background once the archiving has launched, the expected volume of data is not known beforehand. • Similarly, the full content of the scanned directories (i.e., the list of all the files and subdirectories to be archived), is not known before the archiving completes. • Class mode requires a Miria agent. It is therefore not compatible with the client archiving mode, which makes no use of an agent. • Archiving subject to validation cannot operate in Class mode. The interface handles the validation process. The agent cannot scan directories in the background if it does not know which of the files in the directories have been validated for inclusion in the archiving job.
Archiving type	Agent [Client/ Agent] Client	Specifies the type of archiving. These are the options: • **Agent.** Performs agent archiving. The data to be archived must be located on a machine hosting a Miria agent. • **Client/Agent.** Performs client or agent archiving, depending on the platform where the data is located. If it is an agent declared as a platform in Administration Console, agent archiving is favored. • **Client.** Performs client archiving.
Parallel Class Archiving	No|Yes	Permits creation of multiple streams for simultaneous archiving of several files or directories. By default, all the classes (files and directories selected for archiving) in an archiving job are grouped into a single sub-job. If this option is set to Yes, there are as many sub-jobs as archiving classes.

Setting	Possible Values [default]	Description
Maximum number of archiving sub-jobs running in parallel by job	[0] Numerical value between 0 and 128	Select a numerical value between 0 and 128. 0 sets the maximum parallel running archiving by jobs to unlimited.
Data Mover Mode	Allow platform charset Archive operating system rights Archive stub file Archive symbolic links Detect deleted object Detect the change of the object type	Enables one or several data mover archiving options: • **Allow platform charset.** Miria archives files and directories with their native names, except the / (slash) which is replaced by an underscore. By default if this option is not set, Miria controls that file names are compatible with all supported platforms, and replaces all these characters by underscores in the archive: <, >, :, ", /, \, |, ?, and *. . • **Archive operating system rights.** Miria archives all the alternate streams of files and directories. The alternate streams are all attributes, permissions, etc. of an object (i.e., any information that is not the object data). • **Archive symbolic links.** Miria archives the symbolic links as links without archiving their destination files or directories (i.e., Backup mode). • **Archive stub file.** Miria archives the stub files. Note: Limitation. You cannot use the drag-and-drop feature to individually archive a Filter Driver stub; however, you can use this feature to archive the whole Filter Driver directory of stub files. • **Detect deleted object.** • **Detect the change of the object type.**
Archiving subject to validation	No|Yes	Specifies that the archiving jobs must be validated before being run. An e-mail informs you that the jobs are waiting; you must then validate them manually for the archiving operation to be launched.

Setting	Possible Values [default]	Description
Files to exclude from archiving	[No Default], Field to enter file names or file full paths	Enables exclusion of certain files from archiving. You can use the two standard wildcards: * for any number of characters, and ? for a single character. **Example.** To archive a directory without archiving the temporary files it may contain, you can exclude them by entering the *.tmp string in the Value field. Separate the different kind of files to exclude with a carriage return.
Files to archive	[No Default], Field to enter file names or file full paths	Enables inclusion of certain files into an archiving. **Example.** To only archive the Word files in a directory, enter `*.docx` in the Value field. You can use the wildcards `*` and `?`. Separate the different kind of files to archive with a carriage return. Exclusion takes priority over inclusion (e.g., if you exclude all `.tmp` files from an archiving, but then put `2007_*.tmp` in the Files to archive setting, the `2007_*.tmp` files will not be archived).
Directories not to archive	[No Default], Field to enter directory names or directory full paths	Enables exclusion of certain directories from archiving. You can set this setting on either the Default Settings pane or in the Advanced Settings tab when configuring a task or an archiving platform. Enter manually the directory full path or click the selection button. Then, if you are in: • The Default Settings pane, the Platform List window displays the list of all the platforms known by Miria. • An archiving platform or a task, the current platform opens directly in a new window. To obtain the Platform List window, keep pressed the `Ctrl` key while clicking the selection button. You can use the two standard wildcards: `*` for any number of characters, and `?` for a single character. Separate the directories in the list with a carriage return.

Setting	Possible Values [default]	Description
Directories to archive	[No Default], Directory names or directory full paths	Enables inclusion of certain directories in the archiving. Enter manually the directory full path or click the **Browse** button to browse for a directory. The syntax is the name of a directory at the root of the archiving path (e.g., `Dir1`). Separate the directories in the list with a carriage return. If you specify a directory to archive, no directory other than this one is archivable. Exclusion takes priority over inclusion. If running an automatic archiving task, all files at the root of the archiving path are archived, even though a Directory to Archive is specified in the advanced settings.
Windows archiving policy	[No Default], List of existing archiving policies	Enables you to assign a Windows archiving policy as the default for this Miria instance. Select the desired policy from the list in the Value field. As there is no default, an archiving policy must imperatively be defined in the Default Settings for cases where there is no archiving policy set on the object.
Unix/macOS archiving policy	[No Default], List of existing archiving policies	Enables you to assign a Unix archiving policy as the default for this Miria instance. Select the desired policy from the list in the Value field. As there is no default, an archiving policy must imperatively be defined in the Default Settings for cases where there is no archiving policy set on the object.
Custom Media Rule	[No Default], List of existing Custom Rules	Specifies the use of a Custom Media Rule as one of the Media Rule options when using Media Manager. The rule must have been previously defined, and the Media Rule parameter set to Custom in the configuration of the storage manager container, for this to apply.

Setting	Possible Values [default]	Description
Archiving metadata apply on	[Job + Instance] Job only	Specifies whether the metadata set at archiving time is applied only to the archiving job or also to the instance of each archived file. If set at instance level, a search by metadata finds and directly retrieves the corresponding instance of a file.
Collect metadata during archiving	No|Yes	Miria supports over 100 audiovisual and image formats having predefined metadata wrappers. Setting this to Yes causes this metadata to be converted into Miria metadata so that you can use it for searches on archives. The metadata is collected when the archiving is launched, and a Media Metadata sub-job of the archiving is displayed in the List of Jobs interface. The metadata is stored with the archiving instance and not with the data itself; thus, if the file is modified outside of Miria between two archivings, the metadata is different on each instance.
Collect partial retrieval metadata during archiving	No|Yes	Specifies whether partial retrieval metadata must be collected during archiving. If set to Yes, this setting enables you to retrieve only a specific time sequence of an archived media file. With this setting enabled, an archiving job takes longer to complete because some time code conversion operations must be performed.
Video Proxy Transcoding Format	[None] Default QuickTime/H264	Specifies the format to use to generate low-resolution versions of video assets for video preview. These formats are available: • **Default.** MPEG-1 format. • **Quick Time/H264.** This codec is not included in Miria. If you want to use it, download and install it as described in Video Encoding. A license is required to set this setting.

Setting	Possible Values [default]	Description
Collect MIME Type during archiving	No|Yes	The MIME type is a two-part identifier of file formats, for example `audio/mpeg`, `video/quicktime`, or `text/plain`. It helps to determine the program you must use to open the file in a Web browser. It is a more reliable way of determining a format than the filename extension, which can sometimes be incorrect. This setting specifies that the MIME type of a file must be determined when the file is archived and stored with its properties in the database (as one of its inherent Criteria). If the type cannot be determined, a value of application/binary is assigned. By default, this setting is defined to No because determination of the MIME Type does have a cost in terms of performance. Set it to `Yes` if you are likely to perform searches within the archives based on MIME Types.
Deprecated ***Action to perform after archiving***	*[No Action], File Deletion* *Stub (HSM Client)* *Stub (HSM Filter Driver)*	*If File Deletion is selected, Miria deletes files from the agent as soon as they are sent to the server.* *Stub also causes the files to be deleted from the agent. In their place, a stub is created. You can retrieve the archived file by double-clicking the stub. Use the value corresponding to the HSM type (HSM Client or HSM Filter Driver).*
Prevent Directory Spanning Tree Level	[No|0] Integer between 0 and 99	Directory level that you want to archive as a whole on a single media. By default, directory spanning is allowed, and all directories can be split over several media. Use this setting when archiving with Media Manager, to avoid file tree splitting and make retrieval from the media file system easier. For instance, set this setting to 2 to prevent second-level directories or folders (e.g., `/a/ab, /a/ac, /b/ba, /b/bc`, etc.) to span over several media. If the directory is too large to fit on a whole media, then Miriaspans it over several media.

Setting	Possible Values [default]	Description
Statistics on archiving jobs	[Without] Real Time Only Real Time + History	Activates a window to display statistics on a running and / or historic archiving sub-job, accessible from the List of Jobs.

Retrieval Settings

This table describes the retrieval settings in Miria

Setting	Possible Values [default]	Description
Retrieval Mode	[Overwrite objects], Keep existing objects, Rename existing objects	Specifies the retrieval behavior if the retrieved object already exists. These are the valid options: • **Overwrite objects.** Default value. The retrieved object overwrites the existing file. **Important**: Be aware that if the directory where you are to retrieve the object already contains an object having a different type, Miria overwrites the existing object but does not display any warning message. • **Keep existing objects.** Miria does not retrieve the object, so you keep the existing object. • **Rename existing objects.** Rename the existing object, so as to have two file instances. To rename an object, Miria adds a number between brackets to the name of the object present on the disk (e.g., to retrieve the “mytext.doc” object to a disk that already contains an object bearing this name, Miria renames the object on the disk into “mytext (1).doc”).

Setting	Possible Values [default]	Description
Retrieval Type	Agent [Client/Agent] Client	Specifies the type of retrieval. These are the valid options: • **Agent.** Retrieves onto a machine hosting a Miria agent. • **Client/Agent.** Retrieves onto a client or agent, depending on the platform that has requested the retrieval. If it is an agent declared as a platform, retrieval is favored. • **Client.** Retrieves onto the client.
Default Retrieval Root Location	[No Default], text boxes to enter a path	Specifies the path of the directory to which the objects are retrieved by default. Click the Selection button to enter these parameters: • **Platform.** Name of the platform to which you want to retrieve the objects. Click Browse to select a Miria platform. • **Root Location.** Path of the directory to which you want to retrieve the objects. Click Browse to select a directory on the platform.
Retrieval Policy Order	[No Default], List of existing archiving policies	Specifies the archiving policy that indicates the order of priority of the storage manager containers at retrieval. Select the desired policy from the list in the Value field. This setting is useful whenever the data was archived to several storage manager containers (multiple writing), and that you want to choose the container from which to retrieve.
Maximum number of retrieval sub-jobs running in parallel by job	[0] Numerical value between 0 and 128	Select a numerical value between 0 and 128. 0 sets the maximum parallel running retrieval by jobs to unlimited.

Setting	Possible Values [default]	Description
Parallel Retrieval	[Parallel retrieval] By Media for appropriate storages [Splitting by] Job Volume and/or Files	Permits the creation of multiple streams for simultaneous retrieval of several files or directories. The default behavior is for all the data in a retrieval job to be grouped together into a single sub-job. Defines the maximum number of retrieval jobs that can run in parallel. You can select these options: • **Parallel retrieval.** If a storage manages the media, the `Parallel retrieval` source is always displayed in the first place. The `Parallel retrieval` option creates as many sub-jobs as media involved in the retrieval. Miria creates a pre-processing sub-job prior to running the retrieval sub-jobs. • **Splitting.** by: • **Jobs.** Splits the retrieval operation by jobs. • **Volume.** Splits the retrieval operation according to the chosen value in either Volume in GB or Number of files. The splitting and Jobs/Volume/Files options are mutually excluding.
Retrieve with full archive path	No|Yes	Permits to retrieve the file or directory together with its full archive path in the retrieval destination location. Example. If `file.txt` is located in the Orders archive, in the `/2015/July` folder, its path at retrieval is `destination_ path/2015/July/file.txt`.
Create target directory on retrieval	No|Yes	Permits automatic creation of the retrieval destination directory if it does not exist. If you choose the Default Retrieval Root Location option, Miria creates the whole parent tree structure of the retrieved object in the root location.

Setting	Possible Values [default]	Description
Files to exclude from Retrieval	[No Default], Field to enter file names or file full paths	Enables exclusion of certain files from retrieval. You can use the two standard wildcards: `*` for any character any number of times, and `?` for any character once. For example, if you want to retrieve a directory, but do not want to retrieve the temporary files it may contain, you can exclude them by entering the string `*.tmp` in the Value field. Separate the different kind of files to exclude with a carriage return.
Start retrieval only when all media are online or available for use	No|Yes	This Setting is only relevant to configurations using Media Manager and libraries containing physical media. These are the valid options: • **No.** You can launch a retrieval job even if all of the media needed for the retrieval are not present beforehand in the library. If the job encounters missing media, an error message warns you that all needed media are not online. You must then put the media physically into the library. • **Yes.** A retrieval session does not start unless all media needed for its successful completion are already present and available in the library. You can check if this is the case in the List of Requested Media window prior to launching the retrieval job. In cases where you set this setting to Yes and the retrieval selection is larger than one retrieval session (i.e., more than 2048 files or more than 100 GB), the retrieval job starts anyway if all the media needed to retrieve the first session are online.

Setting	Possible Values [default]	Description
Retrieve alternate streams	[Retrieve archived alternate streams] Do not retrieve alternate streams Retrieve with full access rights	Specifies how the alternate streams of archived objects must be retrieved. The alternate streams are all attributes, permissions, etc. of an object (i.e., any information that is not the object data). • **Retrieve archived alternate streams.** Choose this option to retrieve all alternate streams with the actual data (default behavior). • **Do not retrieve alternate streams.** Choose this option if you retrieve Windows objects to a UNIX file system, or UNIX objects to a Windows file system: the alternate streams are OS-specific and cannot be interpreted by a different OS. With this option, the original dates of the objects are not retrieved. The retrieved objects take the current date. • **Retrieve with full access rights.** Choose this option to retrieve files and directories with full access rights granted to all users. When you retrieve an archive folder (i.e., the virtual container created in an archive to store archived files and directories), Miria applies full access rights to it, regardless of the value of this setting.
Statistics on retrieval jobs	[Without] Real Time Only Real Time + History	Activates a window to display statistics on a running and/or historic retrieval sub-job, accessible from the Jobs List interface.

Retention Settings

This table describes the Retention Settings in Miria:

Setting	Possible Values [default]	Description
Retention subject to validation	No|Yes	Specifies whether the retention job must be validated before being run. By default, only the owner of the archive involved or the super user can validate retention jobs, unless the Validate retention job setting has been set.

Setting	Possible Values [default]	Description
Retention extension	[No Default], List of retention periods	Specifies what retention period(s) are available for users to choose from when they validates a retention job.

Duplication Settings

This table describes the Duplication Setting in Miria:

Setting	Possible Values [default]	Description
Statistics on duplication jobs	[Without] Real Time Only Real Time + History	Activates a window to display statistics on a running and/or historic Duplication sub-job, accessible from the List of Jobs interface.

Copy/Move/Synchronization Settings

Setting	Possible Values [default]	Description
Data Mover action	[Copy] Move Synchronization	These are the valid data mover options: • **Copy.** copies files to the destination platform and the files remain present on the source platform. • **Move.** moves the files to the destination platform, and deletes them from the source platform. • **Synchronization.** synchronizes the objects present on a source directory and the objects present on a destination directory. The action that Miria applies to the directories depends on the synchronization type that you have set. When you perform a drag-and-drop from the graphical interface, by default Miria applies the Echo synchronization type .

Setting	Possible Values [default]	Description
Copy type	Agent [Client/ Agent] Client	Specifies the type of copy. These are the copy types: • **Agent.** Performs agent copy. The data to copy must be located on a machine hosting a Miria agent. • **Client/Agent.** Performs client or agent copy, depending on the platform where the data is located. If it is an agent declared as a platform in Administration Console, agent copy is favored. • **Client.** Performs client copy.

Setting	Possible Values [default]	Description
Synchronization Type	Echo Subscribe Contribute Combine	Specifies the type of synchronization. These are the synchronization types: • **Echo.** Applies the synchronization from the source directory to the destination directory. The Echo type takes into account these three operations: • Add: An object present on the source but missing from the destination is copied into the destination. • Update: An object present on both the source and destination is copied into the destination provided that the modification date on source is more recent than the modification date on destination. • Delete: An object present on the destination but missing from the source is deleted from the destination. • **Subscribe.** Applies only to files. Applies the synchronization from the destination directory to the source directory. The Subscribe type takes into account the update operation (i.e., it copies an object into the source provided that the modification date of the object in the destination is more recent). • **Contribute.** Applies the synchronization from the source directory to the destination directory. The Contribute type takes into account these two operations: • Add: An object present on the source but missing from the destination is copied into the destination. • Update: An object present on both the source and destination is copied into the destination provided that the modification date on source is more recent than the modification date on destination. • **Combine.** Applies the synchronization in both ways. The Combine type takes into

Setting	Possible Values [default]	Description
		account these two operations: • Add: An object present on either the source or the destination but missing from the other directory is copied into the directory that doesn't contain the object. • Update: An object present on both the source and destination is copied into the other directory provided that its modification date is more recent.

Setting	Possible Values [default]	Description
Copy mode	Copy alternate streams only (no data) Copy hard links Copy operating system rights Copy S3 object lock configuration Copy symbolic links Skip cold tier objects	Enables one or several copy options for a data mover: • **Copy alternate streams only (no data).** only copies the file and directory permission and alternate streams, without the data. An object alternate streams are all its attributes, permissions, etc., (i.e., any information that is not the object data). • **Copy hard links.** (Only on Linux/Unix over NFS) copies the directory structure from the source filesystem to the destination. The hard link structure on the source is preserved and rebuilt at the destination location. **Note**: It is recommended to use this option in Echo mode. The task scans and creates the hard links, for which the explicit mode is required. This excludes the Subscribe and Combine modes, where the synchronization is applied from the destination to the source. • **Copy operating system rights**. copies all the file and directory alternate streams. An object alternate streams are all its attributes, permissions, etc., (i.e., any information that is not the object data). • **Copy S3 Object Lock configuration.** enables to retain immuability mode and its settings / retention during a compatible S3 to S3 migration. Everything including Legal Hold flag is replicated. **If target bucket does already exist, it is mandatory that Object Lock was upfront enabled.** • **Copy symbolic links.** creates the symbolic links while copying the symbolic link data to the appropriate destination for both files and directories. • **Skip cold tier objects.** enables you to skip the older content that is no longer frequently accessed and this way, to save space and time.

Setting	Possible Values [default]	Description
Number of copy threads	Numeric value from 1 to 128	Number of threads that the data mover can manage for the copy.
Synchronize in Parallel Class	Yes|No	When set to `Yes`, creates multiple streams for the parallel synchronization of several classes. The default behavior is for all the data in a synchronization job to be grouped together into a single sub-job.
Parallel Copy / Synchronization	[No Default], Field to enter time, volumes and number of files.	By default, **Easy Move** creates a single job for all files selected. To improve performance, you can spread all selected files over several jobs. This way, each job will use a different pool's member. For example, if you have a 4 data movers pool, 12 files and that you set a parallelism of 4, each data movers pool will handle 3 files. You can set the number of parallel jobs, by time, by volume and by number of files. It is only applicable to synchronization or copy jobs. This setting can be applied to users, jobs, platforms or archives.
Files to exclude from copy	[No Default], Field to enter file names or file full paths	Enables exclusion of certain files from copy. You can use the two standard wildcards: `*` for any character any number of times, and `?` for any character once. For example, if you want to copy a directory, but do not want to copy the temporary files it may contain, you can exclude them by entering the string `*.tmp` in the Value field. Separate the different kind of files to exclude with a carriage return.
Files to copy	[No Default], Field to enter file names or file full paths	Enables inclusion of certain files into a copy job. For example, to only copy the Word files in a directory, enter `*.docx` in the Value field. You can use the wildcards `*` and `?`. Separate the different kind of files to copy with a carriage return. Exclusion takes priority over inclusion. For example, if you exclude all `.tmp` files from a copy, but then put `2007_*.tmp` in the Files to copy setting, the `2007_ *.tmp` files will not be copied.

Setting	Possible Values [default]	Description
Directories to exclude from copy	[No Default], Field to enter directory names or directory full paths	Enables exclusion of certain directories from copy. Enter manually the directory full path. You can use the two standard wildcards: * for any character any number of times, and ? for any character once. Separate the directories in the list with a carriage return.
Directories to copy	[No Default], Field to enter directory names or directory full paths	Enables inclusion of certain directories in the copy job. Enter manually the directory full path or click the Browse button to browse for a directory. The syntax is the name of a directory at the copy path root (e.g., Dir1). Separate the directories in the list with a carriage return. If you specify a directory to copy, no directory other than this one will be copied. Exclusion takes priority over inclusion.
Statistics on copy jobs	[Without] Real Time Only Real Time + History	Activates a window to display statistics on a running and/or historic Copy sub-job, accessible from the List of Jobs interface.

Media Settings

The Media settings enable you to view media and perform some actions on them.

They are only relevant when using Media Manager or Miria File Storage Container as the storage manager. This table describes the Media Settings in Miria:

Setting	Possible Values [default]	Description
View Media List	No|Yes	Grants permission to view the media list for Media Manager storage managers from the Miria interfaces and Web Services. If you do not enable this setting, you can perform actions on Media Manager media. This setting does not impact storage managers of Atempo File Storage type, for which the media list is always visible.

Setting	Possible Values [default]	Description
View Media History	No|Yes	Grants permission to view the media history which details all the media moves. Available only in the Administration Console.
Close & Reopen	No|Yes	Grants permission to close and reopen a media.
Export Content	No|Yes	Grants permission to export the content of a media. Available only in the Command Line Interface.
Duplicate	No|Yes	Grants permission to duplicate a media.
Eject	No|Yes	Grants permission to eject a media.
Recycle	No|Yes	Grants permission to recycle a media.
Scratch	No|Yes	Grants permission to scratch a media. A scratch media is removed from the Miria database and its fingerprint is deleted.
Delete	No|Yes	Grants permission to delete a media. A deleted media is removed from the Miria database. Available only in the Command Line Interface.

Platforms rights

Enable you to manage platforms rights and to allow them by selecting **Yes**, or to forbid them by selecting **No**.
The platforms rights are the following:

- **Open**
- **Copy**
- **Move**
- **Rename**
- **Delete**
- **Create directory**
- **Synchronize**

Admin and monitoring rights

Grants standard users permission to monitor or administrate some of the Miria components.

These are the valid values:

- **None.** A standard user has default monitoring rights. Standard users are already allowed to administrate the jobs and events corresponding to their archives, as well as the jobs they launch.
- **Monitoring.** Grants a standard user permission to view the component information. The information is read only and you cannot perform any actions.
- **Administration.** Grants standard users permission to administrate the components like a super user. With this permission, standard users can edit properties and perform actions.

These settings are only relevant as default settings and for Users and User Groups. They are not applicable to project archives, platforms, and tasks.

- **All jobs.** Enables you to grant standard users permission to view or administrate all Miria jobs.
- **Jobs on allowed archives.** Enables you to grant standard users permission to view or administrate only the jobs related to archives that they are allowed to open (permission or advanced setting Open on a project archive). Note that these rights are cumulative with the default rights on jobs.

For instance, if your Jobs are set on allowed archives, your Monitoring is allowed to:

- Monitor the jobs related to the archives it is allowed to open.
- Administrate the jobs related to his archives (default right).
- Administrate the jobs he has launched (default right).

If the **All jobs** setting is set to Administration, then, you are allowed to administrate all jobs, regardless of the value of the Jobs on allowed archives setting.

Repositories rights

Enables you to manage repositories rights and to allow them by selecting **Yes**, or forbid them by selecting **No**.

The platforms rights are the following:

- **Open**
- **Archive**
- **Retrieve**
- **Add a folder**
- **Rename a folder**
- **Delete a folder**
- **Move a folder**
- **Rename an object**
- **Delete an object**
- **Move an object**
- **Modify file extension**
- **Manage metadata**
- **Move a folder to another archive**
- **Move an object to another archive**
- **Administration**
- **Validate retention jobs**

E-mail Settings

You can configure Miria for it to send information e-mails.

This table describes the e-mail settings that you can set up:

Setting	Possible Values [default]	Description
Email for Archiving Jobs to Validate	[Job {Job_Number} Awaiting Validation]	Enables you to configure the e-mails that Miria sends automatically to users to inform them about archiving or retrieval jobs they have to validate.
Email for Retention Jobs to Validate	[Job {Job_Number} Awaiting Validation]	Enables you to configure the e-mails that Miria sends automatically to users to inform them about retention jobs they have to validate.
Email for Offline/Prevent Use Media (Retrieval)	[Job {Job_Number} Awaiting {Media_Offline_Number} Offline Media.]	Enables you to configure the e-mails that Miria sends automatically to administrators to inform them that a retrieval job requires media that are currently offline or in prevent use mode.
Email for Scratch Media (Archiving)	[Job {Job_Number} Awaiting Scratch Media]	Enables you to configure the e-mails that Miria sends automatically to administrators to inform them that an archiving job requires introduction of scratch (empty) media in the library.

Settings Templates

You can set a particular settings template and apply it on platform(s), task(s), user(s), group(s), repositories...

To create a new settings template

1. Select the **Parameters** tab, and click on the tile **Settings templates**.
2. Click the **+New settings template** button. A window opens.
3. Enter the name of your settings template and click **New settings template**.
4. Select your template from the list, and click **Edit**. You can now edit your template (Figure 44). Select a new value or keep the one inherited from the template.
   see Default Settings to know how to set each parameter.

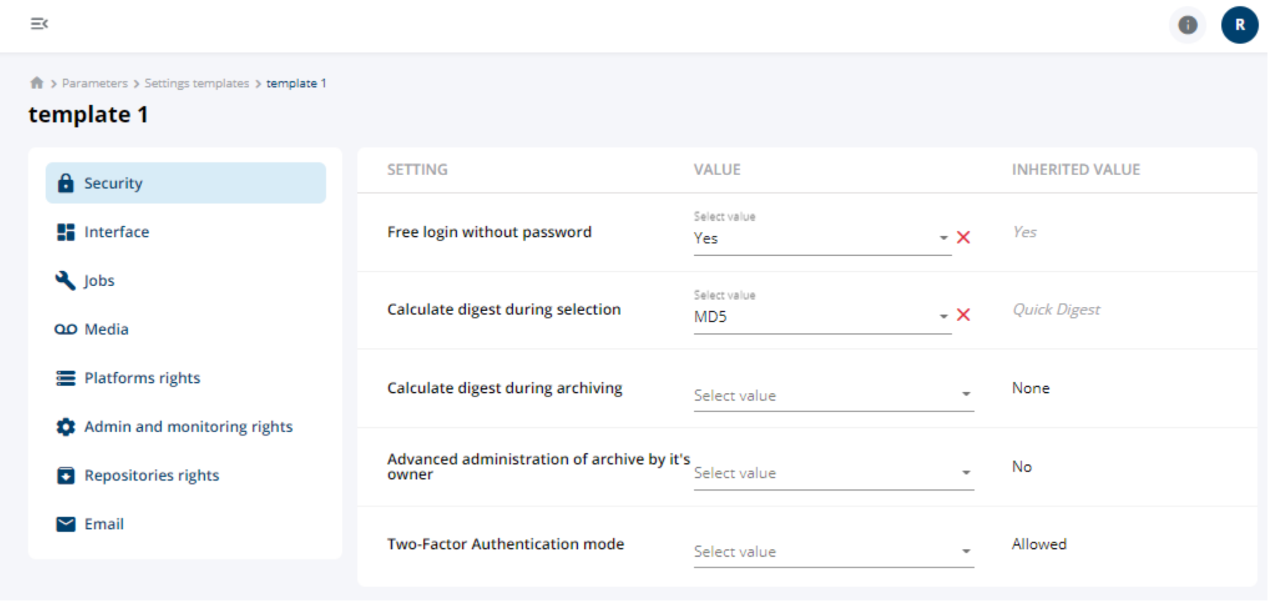
Parameters > Settings templates > template 1
R
Security
Interface
Jobs
Media
Platforms rights
Admin and monitoring rights
Repositories rights
Email

template 1

SETTING	VALUE	INHERITED VALUE
Free login without password	Select value Yes ✕	*Yes*
Calculate digest during selection	Select value MD5 ✕	*Quick Digest*
Calculate digest during archiving	Select value	None
Advanced administration of archive by it's owner	Select value	No
Two-Factor Authentication mode	Select value	Allowed

Figure 44: Settings template

5. Change a settings template any time by selecting it from the list and clicking **Edit**.

To apply a template

1. Go to the page where you want to apply your template, for example on the **Users** page.
2. Select a row and click **Edit**. A new page opens, with all the settings.
3. Select a template in the drop down menu (Figure 45).
   In the Inherited value column, you can see an icon that confirms the value was previously set in the settings template.
   If you select a new value, it will apply only on the item you are editing, and it won't change the settings template.

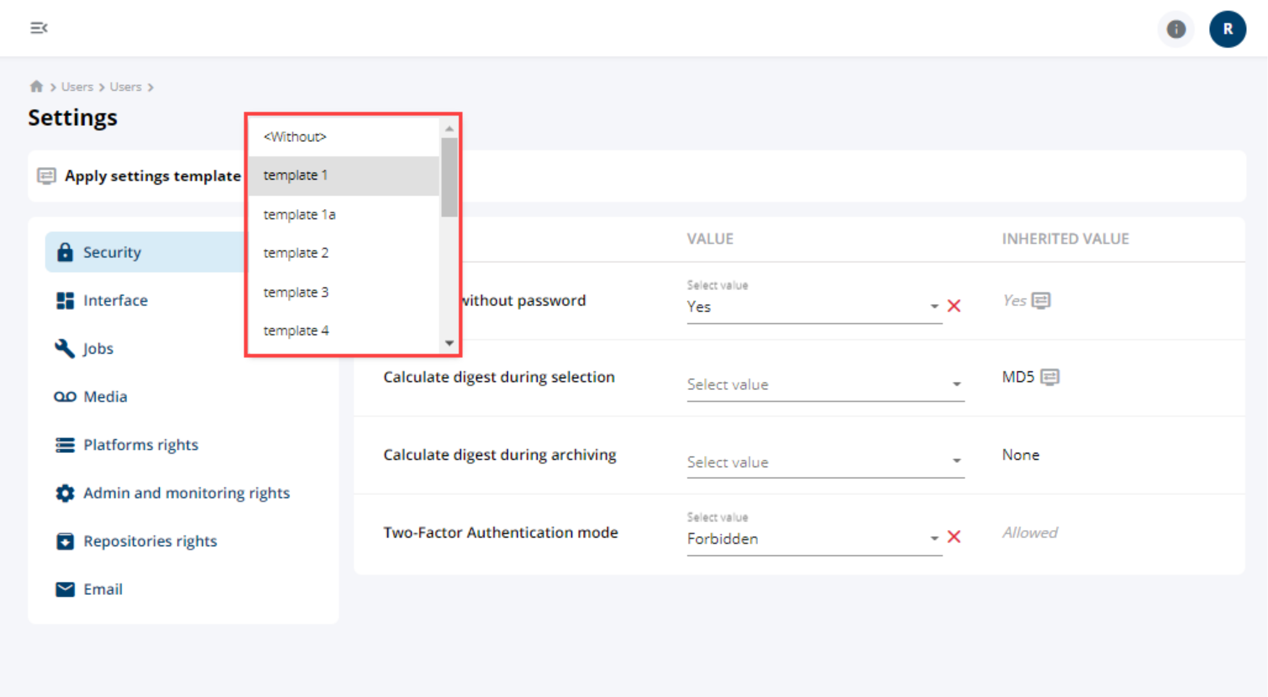

	VALUE	INHERITED VALUE
...ithout password	Select value Yes	Yes
Calculate digest during selection	Select value	MD5
Calculate digest during archiving	Select value	None
Two-Factor Authentication mode	Select value Forbidden	Allowed

Figure 45: Select a template from the drop down menu

Retentions

A retention in Miria is a period of time for which the backed up files are preserved and that you can then search for when objects are archived. The objects are kept in the repository until their defined expiry date is reached or if space is needed on the storage.

To create a retention period

1. From the left pane, select **Parameters** and click the **Retentions** tile.
2. Click the **+New retention** button at the top right.
   A new window opens.
3. Enter the name in the field **Name** to identify the retention in Miria.
4. Use both the **Value** and the **Frequency** fields to enter the duration of the retention period. Select one of these time units:
   - Without.
     This value does not define any retention period and keeps the data forever.
   - Days
   - Weeks
   - Months
   - Years
5. Select the **Keep *x* version(s) after retention ends** check box and enter a value in the field. This value defines the minimal number of instances that Miria must retain in the repository, even in the case that the retention date is reached. This option enables you to always retain the most recent instances and follows these management rules:
   - If this option is disabled, the number of versions to keep is by default set to 0 and all the instances whose retention date is reached will be deleted.

- If the retention date is reached for all the instances, but the number of existing instances for the object is lower than the number of versions to keep, no instance is deleted.
- If the retention date is reached for all the instances, and the number of existing instances for the object is greater than the number of versions to keep (N), the N most recent instances are retained, and all the other instances are deleted.

6. Click **Add** to validate the retention period creation.

> **Note**: Select one of the retentions in the list to edit it. You can change the retention name and set a value and a frequency for the retention.
> You can also choose how many version(s) you want to keep after retention ends.

Run Locks

A run lock limits the number of Miria tasks of any kind that can run simultaneously.

To create a Run Lock

1. Select the **Parameters** tab, and click on the tile **Run Locks**.
2. Click the **+New run lock** button.
   Properties pane of a blank run lock opens.
3. Enter these configuration parameters:
   - **Run Lock Name.** Enter a name to identify the run lock in Miria.
   - **Maximum number of simultaneous runs.**Value, as an integer of the number of tasks of any kind that can run together.
4. Click **Add**.

> **Note**: Select one of the run locks in the list to edit it. You can change the run lock name or the maximum number of simultaneous runs.

Run Timeframes

A Timeframe is a period during which a task is permitted to run or prevented from running. Running Timeframes enables you to prevent tasks from running at times when you know there is heavy use of network resources for other operations.

To create a Run Timeframe

1. In the left pane, select **Parameters** and click the tile **Run Timeframes**.
2. Click **+ New Run Timeframe** at the top right. A window opens and lets you define a name for your run timeframe.
3. Click **Add**. A window opens (Figure 46).

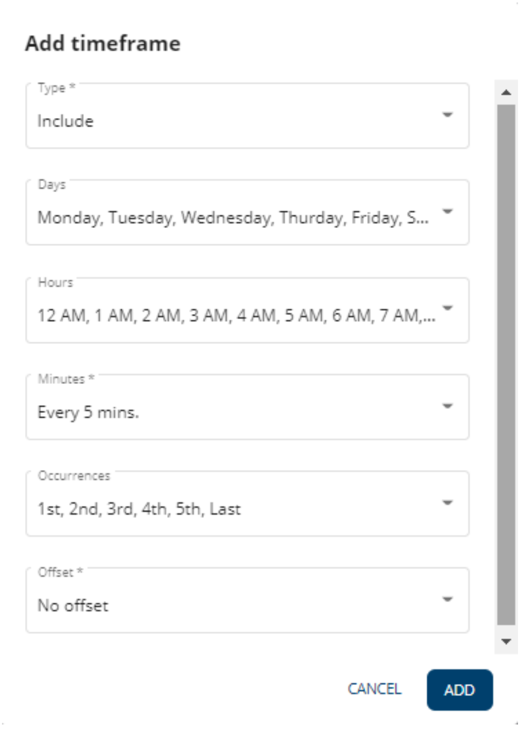

Add timeframe

Type * Include

Days Monday, Tuesday, Wednesday, Thurday, Friday, S...

Hours 12 AM, 1 AM, 2 AM, 3 AM, 4 AM, 5 AM, 6 AM, 7 AM,...

Minutes * Every 5 mins.

Occurrences 1st, 2nd, 3rd, 4th, 5th, Last

Offset * No offset

Figure 46: Add a new timeframe

4. Choose a type of timeframe:
- **Include** Allows a task to run during this timeframe
- **Exclude** Prevents a task from running during this timeframe
5. Select as many days as you want the timeframe to be set on, at which time of the day and how often.
6. Click the **Offset** field to choose if you want an offset and how you want it to be set. Knowing that an offset is a time interval from a defined moment.
7. Click **Add** to create your timeframe.
8. Select your run timeframe and click **+ Add timeframe** to create as many timeframes as you want within it. It enables you to apply different timeframes to a same task and to filter what you want with precision.

Metadata

Metadata are descriptive properties associated with files and assets in repositories for the purpose of classifying them and assisting in their retrieval. They are independent from the file's integral properties such as its name, size, or creation date; the Administrator and/or the user must actively define them and associate them with the file.

For more details about the metadata, see section Manage Metadata.

Select the **Parameters** tab, and then the **Metadata** tile to access the Metadata view. It is divided in two parts:

- **Projects.** Enables you to select a folder and organize them. You can create, rename or delete one.

- **Metadata.** Lists all the metadata from the different folders. Here you can edit each one of them.

To create metadata

1. Click the **+ NEW METADATA** button at the top right of the metadata view.
2. Complete the appropriate fields (Table 10) to define the new metadata.
3. Click **NEW METADATA** validate the metadata creation.

To edit metadata

1. In the metadata list, click the ⋮ button of one of them.
   The **Edit metadata** window opens.
2. You can modify the name, the label, the type, and choose to set it as mandatory or read only (see table).
   If it is already set as read only, you can only modify the field Label.
3. Click **UPDATE MEDATADA** to save the changes.

The following table describes the fields that you can complete to define a new metadata:

Table 10: New metadata settings

Task	Description
Name	**Required**. Descriptive name of your choosing.
Label	**Optional**. It enables you to enter a display name for the metadata, which may be different from its Name parameter. When you perform metadata management tasks, such as assigning a metadata value to a repository or running a search, the Label parameter is displayed as the name of the metadata. If Label is not specified, the metadata Name is displayed instead.
Type	Type of metadata. These are the metadata types: • **Check box.** Boolean, for Yes|No type values. • **Select**. Lets you enter a list of values from which the user can make a selection. • **String.** Lets you enter a free string of characters to find the archived object. • **Date**. Lets you enter a date to find any archived objects that match it. • **Time date**. More detailed than the previous type, this type enables you to find any archived objects that match a particular timestamp. • **Integrer**. Lets you enter a number. • **Duration**. Lets you enter a duration in milliseconds. This type of metadata is useful for media files. • **UUID**. Lets you enter a Universally Unique Identifier with the `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx` format.

Table 10: New metadata settings

Task	Description
Mandatory	Makes the metadata mandatory. When the box is selected, the user cannot launch an archiving job without setting this metadata.
Read only	If you select this box, users can use only this metadata in searches; they cannot change its value.

SMTP server

You can specify an SMTP server and select from a list of configured users to be automatic recipients of the e-mail.

To specify an SMTP server and configure the emails

1. Select the **Parameters** tab, and click on the tile **SMTP Server**.
2. Click the toggle Activate **SMTP Server**.
3. Set the different parameters ().
4. Click **Save changes**.
5. Click the **Send test email** icon.
   A list of all users having e-mail addresses configured in the Users interface opens. It enables you to select the user to receive the test e-mail. If the user is the repository owner, he will also be included in automatically sent e-mails.
6. Click **Send test email** to validate.
   A message informs you that the test email has been sent.

This table describes the parameters that you can set to specify an SMTP server:

SMTP server parameters

Parameter	Description
SMTP server	Hostname or IP of the SMTP server that your company uses (e.g., `smtpserver.my_company.com`.)
Port	Enter a port for the SMTP server.

SMTP server parameters

Parameter	Description
From	By default, this field is empty. You can enter the following variables: • {product} will be replaced by Miria • {host} will be replaced by the Miria server hostname. The messages using this default setting display a sender similar to miria@miria-server.yourdomain.local. No other variable keywords are accepted. You can, however enter a literal string in this field, and the text you supply displays as the e-mails sender. For example, if you enter the Your Atempo Administrator string, it is this string that will be displayed as the e-mail sender.
Message topic	Sets a generic heading for all automatic e-mail messages. This is concatenated with the object set in the advanced settings. See Default Settings for details. For instance, if you enter Automatic Message from Miria Server in this field and Job {Job_Number} Awaiting Validation on the Advanced Setting Email for Archiving/Retrieval Jobs to Validate, the messages display an object similar to Automatic Message from Miria Server - Job 104 Awaiting Validation. You can enter the following variables in this field: • {product} will be replaced by Miria • {host} will be replaced by the Miria server hostname.
Use Authentication. **If you activate this toggle, the following fields appear:**	
User name	Enter the user name you want to use for the authentication.
Password	Enter the password you want to use for the authentication.
Enable SSL/TLS security. **If you activate this toggle, the following field appears:**	
SSL/TLS version	Choose between an SSL or TLS security.

Media Manager

Media Manager is a suite of software components used to manage media in tape libraries and drives. It provides access to media and drives on a specified host or computer for client applications.

The integration between Miria and Media Manager enables you to store the meta data associated with the archived object.

To access the Media Manager interface, select the **Parameters** tab, and then the **Media Manager** tile, the Media Manager interface opens. It is divided in 5 different tabs:

- **Overview** Displays the media distribution, the libraries and all the information about the media manager.
- **Libraries** Allows you to select a library, and to perform actions on its drives, slots, mailboxes, and ejection requests.
- **Configuration** Enables you to configure application, media group, and libraries.
- **Media** Media list that you can apply filters on.
- **Logs** Logs list that you can apply filters on.

Media Manager Overview

Media distribution

Displays the different media, and their assignment statuses. Here are the different assignment statuses (Figure 47):

Table 11: Media status

Field	Description
Assigned	Status of a media known to Media Manager and exclusively used by Miria.
Scratch	Status of a media that you declared to be scratched. If it contains data, Media Manager erases it on the next mount. The media becomes a candidate to be assigned to Miria.
Blank	Status of a media that was identified as empty by the automatic media scan or by a manual mount. A blank media is a candidate to be assigned to Miria.
Unknown	Status of a media which is not known to Media Manager (i.e., it has never been mounted, and is not a scratch media). It can be empty or contain useful data. When mounted by Miria, such media switches to one of these statuses: • If it is empty, it is available to be used by Miria. • If it contains data, the status changes to Orphan.

Table 11: Media status

Field	Description
Orphan	Status of a media which contains data not yet identified by Miria. It will be assigned to Miria as soon as it provides the right fingerprint.
Other	Status of a media that is identified by Media Manager with one of these states: • **Duplicate Fingerprint.** The media fingerprint is the same as another media present in the Media Manager database. You must correct the problem and run a Storage Manager Integrity Check. You cannot scratch, duplicate, or recycle a media having the Duplicate Fingerprint state. • **Scratch Error.** An error occurred while scratching the media. • **Read Error.** An error occurred while reading the media. • **Incompatible Media.** The media is not compatible with any drive in the library.
Free slots	Free slots in a media.

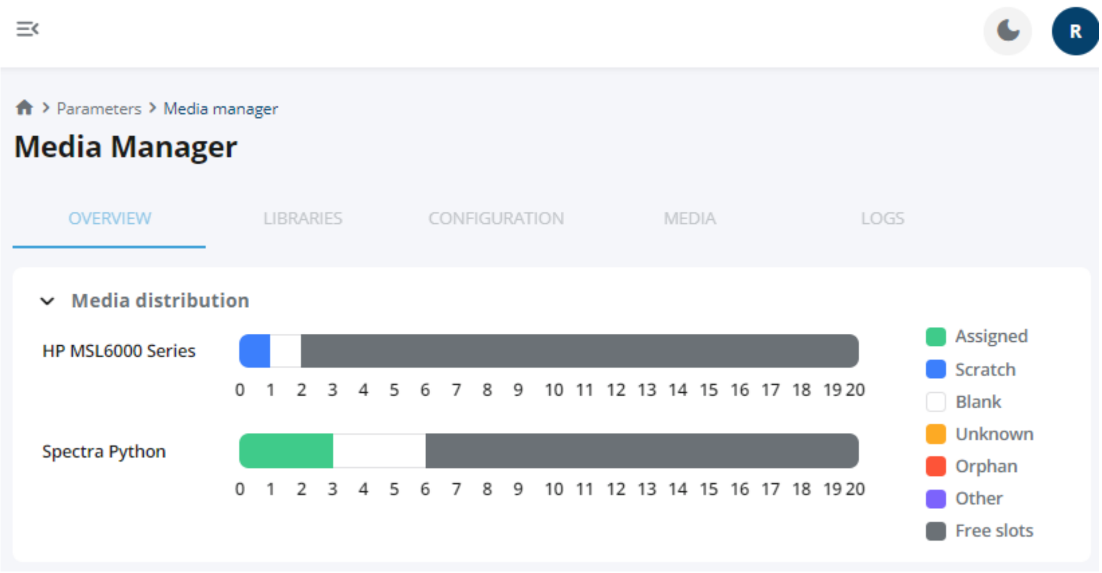

Figure 47: Media distribution

The media manager tree

In the media manager tree, you can choose between these two type of displays (Figure 48):

- **Library view.** Default value.
- **Platform view.** Option value.

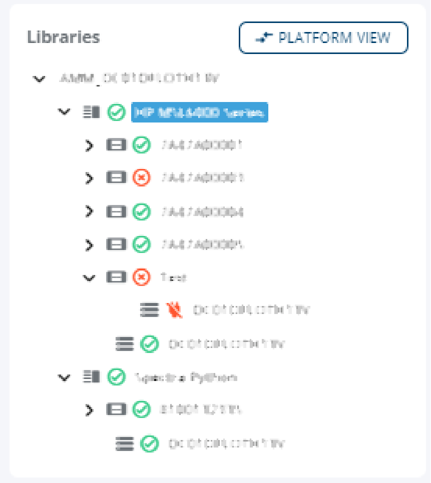

Figure 48: Media manager tree

Library view

In the default view, the tree root level displays the serial number of the library used by the Media Manager.
If an active Media Manager library process is detected, the library line displays a green icon in front of the library name.
A red icon indicates that no library process is detected.
If a red off displays in front of the drive name, this means that the drive is disabled. An orange icon with a wrench means that it is in maintenance mode. To be able to re access the drive, you must enable it.
The third level from the root also relates to the drives, but it describes the Media Manager drive process status on the archiving platform which accesses the drive.

> **Note**: The drive may be enabled while the platform process is disabled or vice-versa. The drive service can continue to run on the platform even if you disabled the drive in Media Manager, or you can stop the process on a drive that remains enabled.

Platform view

In the Platform view, the root element of the tree is the platform (e.g., Minos) on which a Media Manager library or drive process is detected.
At the second level are displayed the libraries and drives whose processes have been detected on the parent platform.
A red off in front of a drive means that the drive is disabled in Media Manager.

Details of the library

Table 12: Library information

Field	Description
Type	There are two available values: - SCSI for SCSI-connected libraries. - ACSCL for ACSLS (Automated Cartridge System Library Software) libraries.
Connection status	Media Manager connection status of the library. Either Online of Offline.
Serial number	Serial number of the library.
Hardware information	Library manufacturer, model, and firmware version.
No. mailboxes enabled	Indicates the number of mailboxes that are enabled. You can disable some of the mailboxes if needed.
Auto-injection	Indicates whether the injection of media in the library is performed automatically. The value can be: - **Yes (ns).** Media Manager triggers media injection automatically by verifying the mailbox contents every n seconds. - **No.** You control media injection and trigger it at your convenience. You can configure the injection mode.
Self-ejecting	Indicates whether the ejection of media from the library is performed automatically. The value can be: - **Yes (ns).** Media Manager triggers media ejection automatically by verifying the ejection requests every seconds. - **No.** You control media ejection and trigger it at your convenience. You can configure the ejection mode.
No. of media / objects	There are 8 different rows that display the number of the following objects: assigned media, scratch media, blank media, unknown media, orphan media, other media, slots, free slots.

Media Manager Libraries Tab

This tab displays a field where you can select a library, which appears under its serial number. You can then click the ⋮ button and select one of the following options:

- **Refresh.**
- **Force inject / eject.**
- **Scan unknown media.**
- **Launch full inventory.**

There are then 4 tabs related to the library you choose which are the following:

- **Drives.**
- **Slots.**
- **Mailboxes.**
- **Ejection requests.**

Drives

In the drives list, there are several columns:

- **Name.**
- **Drive status.** If there is a green icon, it is enabled, an orange icon means it is in maintenance mode, and a red light means it is disabled.
- **Media.** If the drive contains a media, the media name or bar-code displays in the Media column.
- **Activity.**
- **Drive address.**

Click the ⋮ button on one of the drives from the list to perform one of the following actions:

- **Drive diagnostic.** Performs a drive diagnostic.
- **Drive performance.** Runs a drive performance test.
- **Drive history.** Opens the History of drives interface.
  This enables you to review the history of drives moves and statuses on the Media Manager for the drive you selected, and offers a range of filter options.

Mailboxes

This tab gives you access to the mailboxes list, it displays the name of the mailbox and its address.

On the mailbox row, you can click on the ⋮ button to choose whether you want to turn it on / off, or if you want to access the mailbox inventory.

Slots and ejection requests

- **Ejection request.** Lists the media for which an ejection operation has been requested.
- **Slots.** Lists the media present in the library.

Table 13: Ejection request and slots information

Column	Description
Name	Name or bar-code of the media.
Home slot	Library slot in which the media is located.
AMM status	Assignment status of the media.
Volume	Volume of the data already written on the media.
Remaining volume	Volume of the unwritten space remaining on the media.
Fingerprint (Slots only)	The fingerprint is written to the first block of a media when it is assigned and serves to provide logical identification of the media, manage application access to it, and protect it against overwriting.
Ejection map (Ejection requests only)	Mailbox to which the media is moved at ejection.
Ejection planned (Ejection requests only)	Date when the ejection is requested. If the ejection mode is Manual, you must trigger ejection by clicking the Force Eject / Inject button.
WORM (Write Once, Read Many)	Displays the WORM used by a slot.
Write protected	Filters the media list by the selected Write Protected criteria.

On the slots, you can also click the ⋮ button to request an ejection.

Media Manager Configuration Tab

General

Here are displayed some information about the application. you can click the ⋮ button and select the following actions:

- **Edit** Select a storage platform, change the application name, and access to information about the application like the name, the path, the server, the setup or the version.
- **Edit retention** Select a kind of retention. If you select Custom, you can set a number of retention days.
- **New application** Create a new application.

Some information are displayed in this general tab which are displayed in the table below:

Table 14: Media manager information

Field	Description
Application name	Name displayed in the interface.
Application key	A string that identifies the application.
Storage platform	Name of the destination platform that hosts the archived files.
Instance	Show the instance type.
Retention of unassigned Offline Media	Displays the retention settings, whether the retention is set on default, immediate, never or custom. If it is set on custom, or default, you can see on how many days the retention is set.

Application

Access to the applications list, with the application ID, the name of the storage manager, and the number of media used.

Media Group

Access to the media groups list. It displays the media group name, the number of media group used, the number of media group that are free.
For some of the media groups, there is a ⋮ button which enables you to delete a media group.
At the bottom of the list, you can create a new media group. To do so, click the **+New media group** button, a window opens. Enter your media group name, and click **New media group**.

Libraries

Select a library, to access to its parameters displayed in the table below.

Table 15: Library parameters

Field	Description
Name	Name of the library.
Alias	Alias of the library.

Table 15: Library parameters

Field	Description
Auto scan	Indicates whether the timeout is enabled or not. Click the ⋮ button to enable and set the timeout. Here are the options you can select: • **Enable auto scan.** • **Timeout (in minutes).** Time required to start media identification for a library • **Maximum number of drives.**
Default media group	Here are the options when you click the ⋮ button: • **Default (default).** • **Default library (cartgrp_7A47A40000).** • **Custom.** Choose a media group. • **Restart the media manager agent after saving changes.**

Table 15: Library parameters

Field	Description
Injection / Ejection	Click the ⋮ button and choose among the following option: • **Injection triggering.** These are the valid options: ◦ **Manual.** Enables you to control media injection and triggering at your convenience. ◦ **Automatic.** Enables the media injection to be triggered automatically, whenever a media is inserted in the library. Enter the frequency (in seconds) at which Media Manager checks the library contents. • **Ejection triggering.** These are the valid options: ◦ **Manual.**Enables you to control media ejection and triggering at your convenience. ◦ **Automatic.** Enables the media ejection to be triggered automatically, whenever an ejection request is issued. Enter the frequency (in seconds) at which Media Manager checks whether some media are ready to be ejected. • **Restart media manager agent after saving changes.** Click the toggle to restart the Media Manager agent immediately. If you don't, you will still have to restart the Media Manager agent later.
Media cataloging	Select **On** / **Off** and click **Save**.
Media compatibility rule	Choose a rule among the following: • **Default** • **Read / Write only** • **Generating drive only**

Media Manager Media Tab

The media list provides details about all the media present in the libraries managed by Media Manager. These are the type of details provided:

- **Media information.** Related to the hardware (i.e., libraries, slots, drives, and media), the information comes from the Media Manager database.
- **Miria information.** Related to the data on the media, information comes from the Miria database.

Media Information

The following table describes the columns of the Media list:

Table 16: Media list

Column	Description
Tape icon	A lock is displayed on the tape icon when the media is in use and locked by a job. A raised hand is displayed when the media is in prevent use mode.
Name	Name or bar code of the media.
Library	Library that contains the media. If this field is empty, the media is offline.
Drive	Drive that contains the media.
Media group	Media group to which the media belongs.
AMM status	Assignment status of the media.
Home slot	Library slot that contains the media.
Drive address	Location of the drive in the library. This information displays only for media located in a drive.
Fingerprint	The fingerprint is written to the first block of a media when it is assigned and serves to provide logical identification of the media, manage application access to it, and protect it against overwriting.

Options on the media

Media manager rows options, when you click the ⋮ button:

- **Media details.** Opens the Media Details window. This window is very similar to the Media Properties pane.
- **Media status history.** Displays the history of media status.
- **Media move history.** Displays the history of media move.
- **Close / Reopen.** Closes and reopens a media.
- **Duplicate.** Duplicate a media.

- **Recycle.** Recycles a media.
- **Scratch media.** Removes the media from the database and deletes its fingerprint.
- **Verify Media.** Allows you to perform a logical and a physical verification of PAX or LTFS media.
- **Prevent use.** Filters the media list by the selected Prevent Use criteria. These are the valid values:
  - **Yes.** Displays the media that are in prevent use mode.
  - **No.** Displays the media that are not in prevent use mode.
- **Unlock LTFS volume.** Unlocks the LTFS volume.
- **Permanently lock LTFS volume.** Locks permanently the LTFS volume.
- **Change comment.** Adds or edit a comment to describe the media.
- **Request ejection.** Puts the media offline for you to physically remove it from the library.
- **Mount.** Mount the media.

Filters options

Table 17: Filters on the media list

Column	Description
AMM status	Media Manager status: • **Assigned to my storage.** • **Assigned.** Media known to Media Manager and exclusively used by Miria. • **Scratch.** Status of a media that you declared to be scratched. If it contains data, Media Manager erases it on the next mount. The media becomes a candidate to be assigned to Miria. • **Unknown.** Media which is not known to Media Manager (i.e., it has never been mounted, and is not a scratch media). It can be empty or contain useful data. When mounted by Miria, such media switches to one of these statuses: • If it is empty, it is available to be used by Miria. • If it contains data, the status changes to Orphan. You can also change its status manually to scratch. • **Orphan.** Status of a media which contains data not yet identified by Miria. It will be assigned to Miria as soon as it provides the right fingerprint. • **Blank.**Status of a media that was identified as empty by the automatic media scan or by a manual mount. A blank media is a candidate to be assigned to Miria. • **Other.**

Table 17: Filters on the media list

Column	Description
AMM status details	• Media that is identified by Media Manager with one of these states: • **Duplicate Fingerprint.** The media fingerprint is the same as another media present in the Media Manager database. You must correct the problem and run a Storage Manager Integrity Check. You cannot scratch, duplicate, or recycle a media having the Duplicate Fingerprint state. • **Scratch Error.** An error occurred while scratching the media. • **Read Error.** An error occurred while reading the media. • **Incompatible Media.** The media is not compatible with any drive in the library. • **To scratch.** Empties the media, removes its references from the Miria database, and deletes its fingerprint. • **Blank.** Status of a media that was identified as empty by the automatic media scan or by a manual mount. A blank media is a candidate to be assigned to Miria.

Table 17: Filters on the media list

Column	Description
Status	Displays the status of the media and a timestamp indicating when it acquired that status (e.g., Open since 12/22/2016 10:36:02 A.M.). The Status column is sorted by Status and by the date of this Status (i.e., age). These are the statuses that a media can have: • **Closed** • **Closed (No Reopen)** • **Empty** • **New** • **Open** • **Suspended** You can read the media having the status Closed (No Reopen) and retrieve the files written on them; however, you cannot reopen them for new writing. The Reopen button is disabled and the only action that you can perform on them is Request Ejection.

Table 17: Filters on the media list

Column	Description
+ More filters	Advanced filters: • **Media name.** Filters the list media by Media name or Barcode selection. To use the filters, you can: • Filter the bar codes by entering at least one character * or ?. • Use the | character as separator between bar code rules. • **Media type.** Filters the media list by type of media. You can select these values from the Media Type tree: • Disk • ODA • SCSI (e.g., LTO-ULTRIUM, STK, IBM) • Unknown • **Media group.** Filters the media list by the selected media group. • **AMM application.** Filters the media list by the selected Media Manager server instance. • **Library / Drive.** Filters the media list by the selected criteria of library and drive. These are the criteria: • **Not Mounted.** Displays the media outside a drive. • **Mounted.** Displays the media mounted in a drive. • **Offline.** Displays the media located out of the library. • **Online.** Displays the media located inside the library. • Select a library or drive to display the media located in this specific library or drive. • **Storage managers.** Filters the media list by the selected storage manager. • **Storage manager containers.** Filters the media list by the selected storage manager container. • **Cause.** Filters the media list by the selected cause(s).

Table 17: Filters on the media list

Column	Description
	• **Format.** Filters the media list by the selected data format(s). • **Prevent use.** Filters the media list by the selected Prevent Use criteria. These are the valid values: • **Yes.** Displays the media that are in prevent use mode. • **No.** Displays the media that are not in prevent use mode. • **Write protection.** Filters the media list by the selected Write Protected criteria. These are the valid values: • **Yes.** Displays the media that are in write protected mode. • **No.** Displays the media that are not are in write protected mode.
Options on a media	Select the ⋮ button on one of the media row in the list. Two options appear: • **Prevent use.** Put the media on prevent use mode. • **Request ejection.** Puts the media offline for you to physically remove it from the library.

Table 17: Filters on the media list

Column	Description
Options on the media list	Above the media list, there are icons that give you different options: - 🔍 Search for a media - 🕘 Displays all media move history - 🔓 Reopens a previously closed media. You cannot reopen a media with the status Closed (no reopen). - 🔒 Closes a media - ⧉ Opens the duplicate the media window, which enables you to define the duplication parameters and launch media duplication jobs. - ♻ Recycles the media for you to reuse it. - ✎ Scratch media - 🔍 Verify media - ⏏ Request ejection to put the media offline for you to physically remove it from the library. - ⟳ Refresh - ▥ Customize columns These icons enable you to perform actions on multiple media all at once.

Media Manager Logs Tab

This tab displays the media manager logs, here you can filter and manage them.
Here is how the interface is organized:

Table 18: Information and actions on the logs

Column	Description
Filters	You can apply filters on the logs: • **Limit.** Specifies the maximum number of logs that are displayed. • **Libraries.** Select a library. • **Drives.** Select a drive. • **From.** Select a specific date. • **To.** Select a specific limit date.
Columns	The logs information are organized in columns: • **Date.** Date and time when the log occurred, in `dd/mm/yyyy` `hh:mm:ss` format. • **Severity.** Severity of the log. These are the valid values: Information, Error, Notice, Emergency, Alert, Critical, Warning, Debug, and Development. • **Message.** Text description of the log. • **Platform.** Name of the host hosting the logs.
Icons	On the top right of the list, there are the following icons: • 🔍 **Search.** • ⟳ **Refresh.** • ⤓ **Download CSV.** • ▥ **Customize columns.**

Syslog server

You can connect Miria to a Syslog server, and export logs to it. To do so, select the Syslog server tile, in the parameters tab, and click the toggle **Activate export to the Syslog server** and enter the following parameters:

Syslog server parameters

Parameter	Description
Syslog server address	Hostname or IP of the Syslog server. It can be either FQDN, or IP. An hostname can be enough, even if it is not Fully Qualified (FQDN).
Port	Enter a port for the Syslog server.

Syslog server parameters

Parameter	Description
Protocol	Choose from the drop down list one of those protocols: - UDP - TCP - TCPS
Filter by severity	Choose the filter(s) you need among the following: - Select all - DEBUG - INFO - CRITICAL - ERROR - STOP/DIE - WARNING - SUCCESS - STACK - DEBUG STACK - FATAL - AUDIT TRAIL Only the messages with a high severity level (FATAL, CRITICAL, STOP/DIE...) will be exported to the Syslog server.
RFC used for messages format	Choose between the RFC 3164, or the RFC 5424.

> **Note**: The local logs are always there, even if the export is configured.

Deduplication Domains

The deduplication domains are storage zones composed of one or several storage manager containers.
Miria deduplicates a file only if it exists within the grouping of storage manager containers defined in the domain.
Outside that domain, it remains possible to have replicate files.

These are some real use cases, with recommendations on how to configure the deduplication domain:

- **Example 1.** Because the Atempo One to One File storage manager container is filling up quickly, you decide to set up a second container that must be an extension of the first. In this case, you associate both containers with the same deduplication domain. This prevents identical files from being stored redundantly on the two containers.
- **Example 2.** You want to set up two storage manager containers with one being the mirror of the other. In this case, they must not belong to the same deduplication domain. Define

them in separate domains to ensure that if one container is damaged, copies of files are still present in the other.

> **Important**: If you set up a configuration with multiple writing enabled, but both storage manager containers are in the same domain, there is only one writing for every file.

To create a Deduplication Domain

1. From the left pane, select Parameters > Dedup. domains.
2. Click the **+ New deduplication domain** button at the top right. A new window of the new deduplication domain properties opens.
3. Enter the configuration parameters (Table 19).
4. Click the check mark to validate the deduplication domain creation.You must associate the deduplication domain with one or more storage manager containers. This creates the storage zone within which the deduplication is effective.
5. Select the domain in the storage manager container interface, as described in Add a Storage Manager and Container.

This table describes the parameters that you must define to create a deduplication domain:

Table 19: Deduplication domains parameters

Parameters	Description
Deduplication domain name	Name that identifies the deduplication domain within Miria.

Table 19: Deduplication domains parameters

Parameters	Description
Digest type	Enables you to select the digest type. Four options use NSA-designed security hash algorithm cryptographic functions to calculate the digest by reading the file completely. The fifth option calculates the hash without reading the content of the file. These are the available options: • **SHA-1.** Produces a 160-bit digest. • **SHA-256.** Computed with a 32-bit word; produces a 256-bit digest. • **SHA-384.** Produces a 384-bit digest. • **SHA-512.** Computed with a 64-bit word; produces a 512-bit digest. • **File (Name/Size).** Produces a hash sum calculated only on the name and size of a file to be archived. • **File (Name/Size/Modif. Time).** Produces a hash sum calculated only on the name, size, and modification time of a file to be archived. The last two choices address very particular cases and are not recommended for general use. They require a special license key. The more complex the hash, the lower the chance of collision (i.e., two files that are not identical having the same signature); however, calculation times are also slower, which diminishes performance. Atempo recommends using `SHA-256` for the best compromise between performance and collision prevention.
Comment	**Optional.** Descriptive comment that helps you to remember the scope of the domain.

In the deduplication domains list click the ⋮ button and select **Edit** to change those parameters. When clicking the ⋮ button, you can also select **Volume report**. It enables you to view the reporting data for a specific deduplication domain.

Volume Report

Displays volume and deduplication information about the storage manager.

Repository

Displays the repositories list with the number of objects and the volume of data archived in them. You can select a period of time from which you want to see the repositories.

These are the different columns in the list:

- **Nb Files.** Number of files archived in this archive.
- **Nb Directories.** Number of directories archived in this archive.
- **Nb Folders.** Number of archive folders present in this archive.
- **Nb Links.** Number of shared assets links archived in this archive
  (e.g., links that are generated when the archiving of two video edition projects results in the archiving of an asset shared by both projects).
- **Volume.** Volume of all the objects archived in this archive, including all the object instances.

When computing the volume in situations where multiple writing occurs on several storage managers, Miria takes into account only the volume archived (i.e., not the volume actually stored).

Storage manager

Here are the information you can find in the columns from the storage manager list:

- **Name.** Name of the storage manager. The storage managers are displayed by storage manager group.
- **Nb objects.** Number of files archived in this storage manager.
- **Retrievable.** Volume of data at retrieval.
- **Stored.** Volume of data stored. The Retrievable and Stored volumes do not account for compression on storage, if any. In case of multiple writing to several storage managers, Miria takes into account the volume of all writings.
- **% dedup.** Deduplication ratio. The higher the ratio, the more effective the deduplication. When deduplication is not enabled, this ratio is 0%.
- **Media stream volume.** Displays the total volume written on all the media in the storage manager. This is the sum of the sizes of all the tape files. This column only applies to Media Manager or File Storage Container storage managers that use media.
- **Nb media.** Displays the total number of physical media (cartridges) in a Media Manager storage manager, or of virtual media in a File Storage Container storage manager. This column only applies to Media Manager or File Storage Container storage managers that use media.

Dedup. Domain

Here are the information you can find in the columns from the storage manager list:

- **Name.** Name of the deduplication domain. Some storage. When deduplication is not available, the `Deduplication domain is not enabled` message is displayed instead of the domain name.
- **Nb objects.** Number of files archived in this deduplication domain.
- **Retrievable.** Volume of data at retrieval. This is the volume that would be occupied with no deduplication enabled.

- **Stored.** Volume of data stored, accounting for deduplication. When deduplication is not enabled, this volume is similar to the Retrievable Volume. The Retrievable and Stored volumes do not account for compression on storage, if any. In case of multiple writing to several storage managers, Miria takes into account the volume of all writings.
- **% Dedup.** Deduplication ratio. The higher the ratio, the more effective the deduplication. When deduplication is not enabled, this ratio is 0%.
- **Media stream volume.** Displays the total volume written on all the media in the storage manager. This is the sum of the sizes of all the tape files. This column only applies to Media Manager or File Storage Container storage managers that use media.
- **Nb media.**Displays the total number of physical media (cartridges) in a Media Manager storage manager, or of virtual media in a File Storage Container storage manager. This column only applies to Media Manager or File Storage Container storage managers that use media.

Custom Media Rules

A custom media rule enables you to organize data on media according to more complex rules, using an advanced setting.

You can group data together according to these criteria:

- **By user groups.** Each media only contains data belonging to the users in this group.
- **By user.** Each media only contains data belonging to a single user.
- **By platform.** Each media only contains data from a single platform.
- **By project archive.** Each media only contains data from a single project archive.
- **By task.** Each media only contains data from a single task.

Example

You can establish a custom media rule for a user group. You have a group of users in a team called `Research`, who all have personal archives. You can send all these individual archives to the same media, using a Custom Media Rule called `ResearchMedia`.
A user `Gregory`, a member of the `Research` team, runs an archiving task to his personal archive. New media is loaded into the tape drive.
When a second user from the `Research` team, `Elenor`, also runs a personal archiving task, if there is space on the media started by `Gregory’s` archiving task, `Elenor’s` archiving task uses the remaining space on that media. When further members of the `Research` team run personal archiving tasks, the archives are either sent to media already in use (if there is space), or a new media is loaded as needed.

However, users outside the `Research` team do not archive to the same media, so the personal archives of the `Research` team are not mixed up with archives of other groups or types. A custom media rule is a tag name that you associate with the desired object using its advanced settings. This setting applies only if the Media Rule field in the storage manager container definition for the Media Manager storage manager is set to Custom.

Create a custom media rule

1. Click the **+New custom media rule** button on the upper right of your screen.
2. Enter a custom media rule name.
3. Click **Create custom media rule**.

Edit a custom media rule

1. Click the ⋮ button of one of the custom media rule in the list and select edit.
2. Change the custom media rule name and click Update custom media rule name.

Set a custom media rule

1. Go to the list of media on which you need to apply a media rule.
2. Click the ⋮ button in the list.
3. Select **Settings**.
4. Select **Admin and Monitoring rights**.
5. Select a value for the custom media rule.

APPENDIX Recycling Triggered by Volume on Storage

Recycling can also be triggered when a level of volume occupancy is exceeded on the storage. Volume-triggered recycling is used with multiple storage managers. It is non-destructive in that it only deletes data for which Miria has another copy. The purpose of such recycling is to free up space on more expensive, near-line storage by deleting files which also exist on a more economical, deeper storage.

These are the two methods:

- **On demand recycling** Job launches the On demand recycling based on the High Water Mark parameter. When the high water mark in the storage manager is reached, the archiving job stops and a retention job starts running. The retention job deletes files until the data volume reaches the low water mark. The advantage of this approach is that it is on demand. The storage manager is emptied in direct response to your need for space. The disadvantage is that this approach interrupts the archiving job until the retention job completes.
- **Scheduled control of storage occupancy** Volume management on storage managers task performs a monitoring of storage occupancy. At regularly scheduled intervals, this task monitors Miria storage managers. On each storage manager where volume management is enabled, it determines whether the task High Water Mark value is set. If so, it analyzes whether the volume of archived data on the storage manager has attained or exceeded the water mark. If it has, the task triggers a retention job that deletes files until the volume of archived data reaches the Low Water Mark parameter, or until there are no more files eligible for deletion. The advantage of this approach is that it anticipates storage needs and does not interrupt archiving jobs.

These are the options when Volume management is enabled (Table 20):

Table 20: Triggering methods

	On demand recycling	Scheduled monitoring
High Water Mark	**Required** Select the box and set a value in GB. When this value is attained, a retention job is triggered.	Used for jobs and is not needed to launch the Volume management on storage managers task. Do not use this setting as it takes precedence over the Task High Water Mark option if set to a lower value.

	On demand recycling	Scheduled monitoring
Task High Water Mark	Ensure this option is not selected, unless you also want to activate scheduled monitoring.	**Required** Select the box to activate the Volume management on storage managers task on this storage manager. For coherent use of on demand and scheduled monitoring recycling, set the GB value to be between high and low water marks.
Low Water Mark	**Required** Select the box and set a value in GB. The retention job attempts to delete files until the data volume reaches this value. If there are no more eligible files (e.g., it has deleted all files having a second copy and all the files still in the storage manager are single copies), it stops before this water mark is reached.	**Required** Select the box and set a value in GB. The retention job attempts to delete files until the data volume reaches this value. If there are no more eligible files (e.g., it has deleted all files having a second copy and all the files still in the storage manager are single copies), it stops before this water mark is reached.

APPENDIX Replications in between two S3 object storages

What is replicated:

- Filename
- Data
- All user defined metadata (wide open area in S3)
- Only one system defined metadata: content type
- Access Control List (ACL)
- Object Lock values

What is not replicated:

- Last modification time (mtime). It is not possible to re-apply it, but Miria stores the original value coming from source in an user defined metadata at target.
- VersionID: it is not possible to re-apply it, because VersionID is defined on the server side.
- **Everything else that is not listed in the first list above is not replicated.**