Moonwalk
Administration Guide

version 2024.1

document revision 1

Contents

Chapter 1

Overview

1.1 Introduction

Moonwalk is a heterogeneous Data Management System. It automates and manages the movement of data from primary storage locations to lower cost file systems, object stores and cloud storage services. Use cases include storage cost optimization, backup optimization and workload placement.

Files are *migrated* from primary storage locations to secondary storage locations. Files are *demigrated* transparently when accessed by a user or application. Moonwalk also provides functionality to copy and move files, as well as a range of Disaster Recovery options.

What is Migration?

From a technical perspective, file migration can be summarized as follows: first, the file content and corresponding metadata are copied to secondary storage as an MWI file/object. Next, the original file is marked as a 'stub' and truncated to zero *physical* size (while retaining the original *logical* size for the benefit of users and the correct operation of applications). The resulting stub file will remain on primary storage in this state until such time as a user or application requests access to the file content, at which point the data will be automatically returned to primary storage.

Each stub encapsulates the location of the corresponding MWI data on secondary storage, without the need for a database or other centralized component.

1.2 Conventions used in this Book

References to **labels, values and literals** in the software are in *'quoted italics'*.

References to **actions**, such as clicking buttons, are in **bold**.

References to **commands and text typed in** are in `fixed font`.

Notes are denoted: **Note:** This is a note.

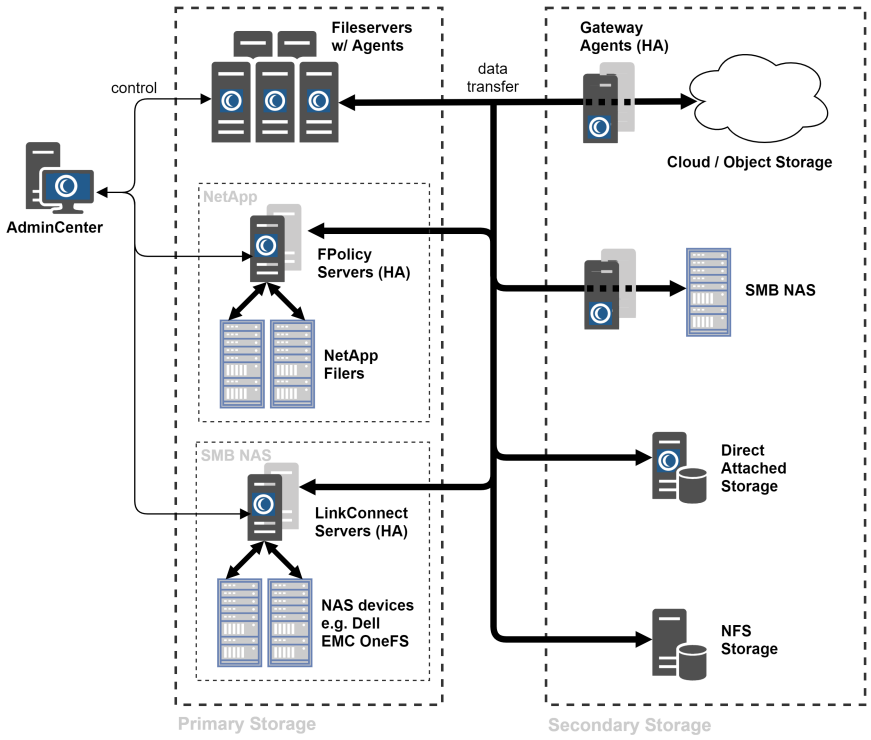

Figure 1.1: Moonwalk System Overview

Important notes are denoted: **Important: Important point here.**

1.3 System Components

Figure 1.1 provides an overview of a Moonwalk deployment. All communication between Moonwalk components is secured with Transport Layer Security (TLS). The individual components are described below.

Moonwalk AdminCenter

AdminCenter is the system's policy manager. It provides a centralized web-based configuration interface, and is responsible for task scheduling, server monitoring and file reporting. It lies outside the data path for file transfers.

Moonwalk Agent

Moonwalk Agent performs file operations as directed by AdminCenter Policies. The Agent is also responsible for retrieving file data from secondary storage upon user / application access.

File operations include migration, move, copy and demigration, as well as a range of operations to assist disaster recovery. Data is streamed directly between agents and storage without any intermediary staging on disk.

When installed in a Gateway configuration, the Agent may function as a plugin container which allows Moonwalk to be extended to enable access to third-party protocols and special devices. Device specific configuration details (such as sensitive encryption keys and authentication details) are isolated from the file servers.

Optionally, Gateways can be configured for High-Availability (HA).

Moonwalk FPolicy Server

FPolicy Server provides migration support for NetApp filers via the NetApp FPolicy protocol. This component is the equivalent of Moonwalk Agent for NetApp filers.

FPolicy Server may also be configured for High-Availability (HA).

Moonwalk LinkConnect Server

LinkConnect Server provides link-based migration support for either Dell EMC OneFS or as an alternative method for migrating files from Windows Server volumes in the case where an agent may not be installed directly on the file server. This component performs a similar role to Moonwalk Agent for SMB shares.

LinkConnect Server may also be configured for High-Availability (HA).

Moonwalk DrTool

Moonwalk DrTool is an additional application that assists in Disaster Recovery.

Note: This functionality is not included with *Starter Edition* licenses.

1.4 AdminCenter Concepts

Moonwalk AdminCenter is the web-based interface that provides central management of a Moonwalk deployment. It is installed as part of the Admin Tools package.

When entering the AdminCenter, the *'Dashboard'* will be displayed – we will come back to the dashboard in §1.5. For now, the remainder of this section will follow the AdminCenter's navigation menu.

1.4.1 Servers

The *'Servers'* page displays the installed and activated agents across the deployment of Moonwalk. Health information and statistics are provided for each server or cluster node. You will use this page when activating the other components in your system.

Click a Server's ellipsis control to:

- view additional server information
- configure storage plugins
- add / retire / restart cluster nodes
- upgrade a standalone server to high-availability
- view detailed charts of recent activity
- edit server-specific configuration (see Appendix F)

1.4.2 Sources

Sources describe volumes or folders to which Policies may be applied (e.g., locations on the network from which files may be Copied, Moved or Migrated). Optionally, a Source may also specify a point-in-time at which to view a versioned object store, or a snapshot of a filesystem.

A Source location is specified by a URI. Platform-specific information for all supported sources is detailed in Chapter 5. A filesystem browser is provided to assist in setting the URI location interactively.

Subdirectory Filtering

Within a given Source, individual directory subtrees may be included or excluded to provide greater control over which files are eligible for policy operations. Excluded directories will not be traversed.

On the Source Details page, the directory tree may be expanded and explored in the *'Subdirectory Filtering'* section. By default, the entire source will be included.

1.4.3 Destinations

Destinations are storage locations that Policies may write files to (e.g., locations on the network to which files are Copied, Moved or Migrated). Platform-specific information for all supported sources is detailed in Chapter 5.

Destinations are specified by entering a URI using the browser panel – if the specified folder does not exist, it will be created at Task execution time.

Optionally, a Destination may be configured to use Write Once Read Many (WORM) semantics for migration operations. This option is useful when the underlying storage device has WORM-like behavior, but is exposed using a generic protocol.

1.4.4 Rules

Rules allow a specific subset of files within a Source or Sources to be selected for processing.

Rules can match a variety of metadata: filename / pathname, size, timestamps / age, file owner, object storage class, and attribute flags. A rule matches if **all** of its specified criteria match the file's metadata. However, rules can be negated or compounded as necessary to perform more complex matches.

You will be able to simulate your Rules against your Sources during Policy creation.

Some criteria are specified as comma-separated lists of patterns:

- wildcard patterns, e.g. `*.doc` (see §A.1 (p.92))
- regular expressions, e.g. `/2004-06-[0-9][0-9]\.log/` (see §A.2 (p.93))

Note that:

- files match if any one of the patterns in the list match
- whitespace before and after each pattern is ignored
- patterns are case-insensitive (see §A.3 (p.93))
- filename patterns starting with '/' match the path from the point specified by the Source URI
- filename patterns NOT starting with '/' match files in any subtree
- literal commas within a pattern must be escaped with a backslash

1.4.5 Policies

A Policy specifies an *operation* to perform on a set of files. Depending on the type of operation, a Policy will specify Source(s) and/or Destination(s), and possibly Rules to limit the Policy to a subset of files.

Each operation has different parameters, refer to Chapter 4 for a full reference.

Optionally, most policies can be configured to log their results to an external database, in addition to the built-in textual logs – see Appendix E.

1.4.6 Tasks

A Task selects one or more Policies for execution.

While a Task is running, its status is displayed in the *'Running Tasks'* panel of the *'Dashboard'*. When Tasks finish they are moved to the *'Recent Tasks'* panel.

Operation statistics are updated in real time as the task runs. Operations will automatically be executed in parallel, see Appendix F for more details.

Scheduling Tasks

While Tasks may be launched manually, it is often more desirable to enable a Task's schedule with a specific start time and optional repeat configuration.

If multiple Tasks are scheduled to launch simultaneously, Policies on each Source are grouped such that only a single traversal of each file system is required.

To launch a task using a webhook or the Management API, refer to Appendix D.

Completion Notification

When a Task finishes running, regardless of whether it succeeds or fails, a completion notification email may be sent as a convenience to the administrator. This notification

email contains summary information similar to that available in the *'Recent Tasks'* panel on the *'Dashboard'*.

To use this feature, either:

- check the *'Notify completion'* option when configuring the Task, or
- click the notify icon on a running task on the *'Dashboard'*

Post-run Actions

Each Task may be configured to perform one or more post-run actions.

Use post-run actions to start further tasks, POST to remote HTTP servers to trigger 3rd party webhooks or start jobs within orchestration platforms (Ansible Tower, Puppet Enterprise, Control-M, Jenkins, etc), run local programs / scripts, or run remote commands via SSH.

Each post-run action may be performed for one or more of the following Task outcomes:

- the Task is stopped by a user
- a fatal Policy error has occurred
- all Policies have completed but some operations failed
- all attempted operations were successful but locked files were skipped
- all operations were successful

Note that when a Task spans multiple Sources – which will result in multiple corresponding entries in the Running / Recent Tasks lists – the outcome of the Task is the most severe of the above outcomes. For example, with two Sources, if all operations succeed on one Source, but the Policies on the other Source complete with errors, the outcome of the entire Task would be *'all Policies have completed but some operations failed'*. Post-run actions will be performed once the Task's Policies have been completed (or stopped) for all Sources.

1.4.7 DB Sources

A DB Source describes a database query which is used as an alternative to Sources and Rules to select the files on which a Policy is to operate. Each result in a database query corresponds to a single file that the Policy should be applied to.

Use a DB Source when the set of source files is either already known and is stored in an SQL database, or when the set of files can be derived using other metadata held in a database by using an SQL query.

For more information, please refer to Appendix E.

1.4.8 Browser

The built-in Browser enables examination of a variety of source / destination locations. Features include:

- Metadata display and sorting
- Optional filtering based on filename pattern or Rules

- Download, upload, and deletion of files
- For versioned object stores, files may be viewed at a specify a point-in-time
- For filesystems, files may viewed within a snapshot

1.4.9 Reports

Reports – generated by Gather Statistics Policies – contain charts detailing:

- a 30-day review of access and change activity
- a long-term trend chart to assist with planning migration strategy
- a breakdown of the most common file types
- optionally, a breakdown of file ownership
- object storage class information (where supported)

Secure links to reports may be shared to other members of your organization without the need to create and distribute user credentials. Either share individual reports or share all reports generated by a Policy on an ongoing basis. To limit access to reports by IP address / subnet, refer to §C.2 (p.97).

1.4.10 Recovery

The *'Recovery'* page provides access to multiple versions of the recovery files produced by each *Create Recovery File From Source/Destination* Policy. Retention options may be adjusted in *'Settings'*.

Refer to Chapter 6 for more information on performing recovery operations.

1.4.11 Settings

The AdminCenter *'Settings'* page allows configuration of a wide range of *global* settings including:

- email notification
- configuration backup (see §3.3 (p.16))
- security roles (see §C.2 (p.97))
- work hours
- AdminCenter logging
- user interface language selection

Additionally, it is possible to:

- view and update the product license
- suspend the scheduler to prevent scheduled Tasks launching while maintenance procedures are being performed
- update the TLS certificate for the AdminCenter web interface

Server-specific settings and plugin configuration are available via the *'Servers'* page.

1.4.12 Help

The *'Help'* page provides version information, as well as links to documentation and support resources. You may also view the global log, or generate a system diagnostic file (`support.zip`) for use when contacting Moonwalk Support.

1.5 AdminCenter Dashboard

The *'Dashboard'* provides a concise view of the Moonwalk system status, current activity and recent task history. It may also be used to run Tasks on-demand via the *'Launch Task. . . '* control.

The *'Notices'* panel, displayed on the expandable graph bar, summarizes system issues that need to be addressed by the administrator. For instance, this panel will guide you through initial setup tasks such as license installation.

The circular *'Servers'* display shows high-level health information for the servers / clusters in the Moonwalk deployment.

For capacity-based licenses, license capacity consumption is shown on the *'Usage'* panel.

1.5.1 Storage Charts

'Primary' and *'Secondary'* storage charts may be read together to gain insight into the impact of currently configured migration policies on primary and secondary storage consumption over time. Each bar indicates an amount of storage space consumed or released. Consumed storage is indicated by a positive bar, while released storage is shown in the negative. Stacked bars indicate the contributions of the different operations by color.

For instance, a Migration Policy consumes secondary storage in order to release primary storage.

By contrast, demigration consumes primary storage immediately, but defers release until later. Specifically, either the primary storage is released by a Quick-Remigrate, or the associated secondary storage is released by a Scrub.

In a complex environment, these charts provide insight into patterns of user-behavior and policy activity.

Click on a bar to zoom in to an hourly breakdown for the chosen day.

1.5.2 Other Charts

The *'Processed'* line chart graphs both the rate of operations successfully performed and data processed over time. Data transfer and bytes Quick-Remigrated (i.e. without any transfer required) are shown separately.

The *'Operations'* breakdown chart shows successful activity by operation type across the whole system over time. Additionally, per-server operations charts are available via the *'Servers'* page – see §1.4.1.

The *'Operations'* radar chart shows a visual representation of the relative operation profile across your deployment. Two figures are drawn, one for each of the two preceding 7-day periods. This allows behavioral change from week to week to be seen at a glance.

1.5.3 Task Control & History

Per-file operation details (including any error messages) may be viewed by clicking a Task's log icon. It is also possible to launch and stop Tasks, update task configuration, or request a completion notification for a task that is already in progress.

Chapter 2

Deployment

Refer to these instructions during initial deployment and when adding new components. For upgrade instructions, please refer to §3.7 (p.21) instead.

For further information about each supported storage platform, refer to Chapter 5.

2.1 Installing Admin Tools

The Moonwalk Admin Tools package consists of the AdminCenter and the DrTool application (not licensed for *Starter Edition* users). Admin Tools must be installed before any other components.

System Requirements

- A **dedicated** server with a supported operating system:
  - Windows Server 2022
  - Windows Server 2019
  - Windows Server 2016
- Minimum 12GB RAM
- Minimum 10GB disk space for log files
- Active clock synchronization (e.g. via NTP)

Setup

1. Run `Moonwalk Admin Tools.exe`
2. Follow the instructions on screen

2.1.1 Initial Configuration

After completing the installation process, Admin Tools must be configured via the AdminCenter web interface. The AdminCenter will be opened automatically and can be found later via the Start Menu.

The web interface will lead you through the process of initial configuration: refer to the *'Notices'* panel on the *'Dashboard'* to ensure that all steps are completed.

Consider configuring additional users or IP restrictions – see §C.2 (p.97).

NFS Browser Agent

If you are planning to use NFS storage, Moonwalk Gateway Agent should also be installed on the same server as Admin Tools. Remember to install this when installing other Agents.

2.2 Installing Agents

Each Agent server may fulfill one of two roles, selected at installation time.

In the *'Fileserver Agent for migration'* role, an agent assists the operating system to migrate and demigrate files. It is **essential** for the agent to be installed on all machines from which files will be migrated.

By contrast, in the *'Gateway Agent'* role, an agent provides access to external devices and storage services. While it does allow access to local disk and mounted SAN volumes, it does not provide local migration source support. Storage plugins will normally be deployed on Gateways.

2.2.1 High-Availability Gateway Configuration

When using Gateway Agents to access third-party devices or storage services using Moonwalk plugins or the `smb` scheme, a high-availability gateway configuration is recommended. Such Gateway Agents must be activated as 'High-Availability Gateway Agents'.

Note: When using a Windows failover cluster to provide Gateway access to a mounted SAN volume, the server must be activated as a 'Windows Failover Cluster' rather than 'High-Availability Gateway Agent'.

High-Availability Gateway DNS Setup

At least two Gateway Agents are required for High-Availability.

1. Add each Gateway Agent server to DNS
2. Create an FQDN that resolves to all of the IP addresses
3. Use this FQDN when activating the HA Servers
4. Use this FQDN (or a CNAME alias to it) in Moonwalk Destination URIs

Example:

- `gw-1.example.com` → `192.168.0.1`
- `gw-2.example.com` → `192.168.0.2`
- `gw.example.com` → `192.168.0.1, 192.168.0.2`

Note: The servers that form the High-Availability Gateway cluster must NOT be members of a Windows failover cluster.

2.2.2 Moonwalk Agent for Windows Servers

System Requirements

- Supported Windows Server operating system:
  - Windows Server 2022
  - Windows Server 2019
  - Windows Server 2016
- Minimum 8GB RAM
- Minimum 2GB disk space for log files
- Active clock synchronization (e.g. via NTP)

Note: When installed in the Gateway role, a **dedicated** server is required, unless it is to be co-located on the Admin Tools server. When co-locating, create separate DNS aliases to refer to the Gateway and the AdminCenter web interface.

Setup

1. Run the `Moonwalk Agent.exe`
2. Follow the instructions to activate the agent via AdminCenter

2.2.3 Moonwalk Gateway Agent for Linux

System Requirements

- A **dedicated** x86_64 server with a supported operating system:
  - Red Hat Enterprise Linux (RHEL) 9
  - Ubuntu Server 22.04 LTS
- Minimum 8GB RAM
- Minimum 2GB disk space for log files
- Active clock synchronization (e.g. via NTP)

Setup – RHEL

In a root terminal:

1. `tar xzf Moonwalk_Gateway_Agent_rhel9_2024_1.tgz`
2. `Moonwalk_Gateway_Agent_rhel9_2024_1/install.sh`
3. Follow the instructions to activate the agent via AdminCenter

Setup – Ubuntu

In a root terminal:

1. `tar xzf Moonwalk_Gateway_Agent_ubuntu22_2024_1.tgz`

2. `Moonwalk_Gateway_Agent_ubuntu22_2024_1/install.sh`
3. Follow the instructions to activate the agent via AdminCenter

2.2.4 Moonwalk FPolicy Server for NetApp Filers

A Moonwalk FPolicy Server provides migration support for one or more NetApp Filers through the FPolicy protocol. This component is the equivalent of Moonwalk Agent for NetApp Filers. Typically FPolicy Servers are installed in a high-availability configuration.

System Requirements

- A **dedicated** server with a supported operating system:
  - Windows Server 2022
  - Windows Server 2019
  - Windows Server 2016
- Minimum 8GB RAM
- Minimum 2GB disk space for log files
- Active clock synchronization (e.g. via NTP)

Setup

Installation of the FPolicy Server software requires careful preparation of the NetApp Filer and the FPolicy Server machines. Instructions are provided in §5.3 (p.42).

2.2.5 Moonwalk LinkConnect Server

A Moonwalk LinkConnect Server provides link-based migration support for one or more Dell EMC OneFS or Windows SMB shares. This component performs a similar role to Moonwalk Agent without the need for software to be installed directly on the NAS or file server.

System Requirements

- A **dedicated** server with a supported operating system:
  - Windows Server 2022
  - Windows Server 2019
  - Windows Server 2016
- Minimum 2GB disk space for log files (on the system volume)
- Minimum 1TB disk space for LinkConnect Cache (as a single NTFS volume)
- RAM: 8GB base, **plus**:
  - 4GB per TB of LinkConnect Cache
  - 0.5GB per billion link-migrated files
- Active clock synchronization (e.g. via NTP)

Setup

Installation of the LinkConnect Server software requires careful configuration of both the NAS / file server and the LinkConnect Server machines. Instructions are provided in §5.4 (p.49) for OneFS and §5.2 (p.38) for Windows file servers. Other devices are not supported.

2.3 LinkConnect Client Deployment

Installation

Having deployed one or more LinkConnect Servers, all Windows clients that will need to access link-migrated files will require the LinkConnect Client Driver to be installed as follows:

1. Ensure the client machine is joined to the Active Directory domain
2. Run `Moonwalk LinkConnect Client Driver.exe`
3. Follow the prompts

Alternatively to ease deployment, the installer may be run in silent mode by specifying `/S` on the command line. Note that when upgrading the driver silently, the updated driver will not be loaded until the next reboot.

Important: Client Driver versions newer than the installed LinkConnect Server version should not be deployed.

Deployment Considerations

Access to NAS / file server shares containing files that have been link-migrated must use the domain credentials of the logged-in Windows desktop session. When a user accesses a link-migrated file, the client driver will transparently redirect the access to the LinkConnect Server if required. This redirected access will use the same logged-in Windows desktop session credentials.

Installation of the client driver will enable remote symlink evaluation in Windows. If remote symlink evaluation was disabled prior to client driver installation (this is the default behavior in Windows 10), the driver will continue to prevent remote symlink access for other symlinks. Do not disable remote symlink evaluation (e.g. by group policy) after installation since doing so will cause the client driver to stop functioning.

Client Driver Removal

In the unlikely event that the LinkConnect client driver must be removed, please complete the following steps:

- Open an Administrator command prompt
- `sc delete mwilcflt`
- `fsutil behavior set symlinkEvaluation R2R:0`
- Reboot

Chapter 3

Usage

3.1 DNS Best Practice

Storage locations in Moonwalk are referred to by URI. Relationships between files must be maintained over a long period of time. It is therefore advisable to take steps to ensure that the FQDNs used in these URIs are valid long-term, even as individual server roles are changed or consolidated.

In a production deployment, always use Fully Qualified Domain Names (FQDNs) in preference to bare IP addresses.

It is recommended to create DNS aliases for each logical storage role for each server. For example, use different DNS aliases when storing your finance department's data as opposed to your engineering department's data – even if they initially reside on the same server.

3.2 Getting Started

3.2.1 Analyzing Volumes

Once the software has been installed, the first step in any new Moonwalk deployment is to analyze the characteristics of the primary storage volumes. The following steps describe how to generate file statistics reports for each volume.

In the AdminCenter web interface:

1. Create Sources for each volume to analyze
2. Create a 'Gather Statistics' Policy and select all defined Sources
3. Create a Task for the 'Gather Statistics' Policy
    - For now, disable the schedule
4. On the *'Dashboard'*, click the *'Launch Task. . . '* control
5. Launch the Task
6. When the Task has finished, view the report(s) on the *'Reports'* page

3.2.2 Migrating Files

Using the information from the reports, create a rule to select files for migration. A typical rule might limit migrations to files modified more than six months ago. The reports' long-term trend charts will indicate the amount of data that will be migrated by a 'modified more than *n* months ago' rule – adjust the age cutoff as necessary to suit your filesystems.

To avoid unnecessary migration of active files, be conservative with your first Migration Rule – it can be updated to migrate more recently modified files on subsequent runs.

Once the Rule has been created:

1. Create a Destination to store your migrated data
    - see Chapter 5 for platform-specific instructions
2. Create a Migration Policy and add the Source(s), Rule and Destination
3. Use the *'Simulate rule matching. . . '* button to explore the effect of your rule
4. Create a Task for the new Policy
5. Launch the task

When the task has completed, check the corresponding *'Recent Tasks'* entries on the *'Dashboard'*. Click on the log icon to review any errors in detail.

Migration is typically performed periodically: configure a schedule on the Migration Task.

3.2.3 Next Steps

Chapter 4 describes all Moonwalk Policy Operations in detail and will help you to get the most out of Moonwalk.

The remainder of this chapter gives guidance on using Moonwalk in a production environment.

3.3 Configuration Backup

This section describes how to backup Moonwalk configuration (for primary and secondary storage backup considerations, see §3.4).

3.3.1 Admin Tools

Backing up the Moonwalk Admin Tools configuration will preserve policy configuration and server registrations as well as per-server settings and storage plugin configuration.

Backup Process

Configuration backup can be scheduled on the AdminCenter's *'Settings'* page. A default schedule is created at installation time to backup configuration once a week.

Configuration backup files include:

- Policy configuration
- Server registrations
- Per-Server settings, including plugin configuration, keys etc.
- Recovery files
- Settings from the AdminCenter *'Settings'* page
- Settings specified when Admin Tools was installed

It is strongly recommended that these backup files are retrieved and stored **securely** as part of your overall backup plan. These backup files can be found at:

`C:\Program Files\Moonwalk\data\AdminCenter\configBackups`

Additionally, log files may be backed up from:

- `C:\Program Files\Moonwalk\logs\AdminCenter\`
- `C:\Program Files\Moonwalk\logs\drtool\`

Restore Process

1. Ensure that the server to be restored to has the same FQDN and IP address as the original server
2. If present, uninstall Moonwalk Admin Tools
3. Run the installer: `Moonwalk Admin Tools.exe`
   - use the same version that was used to generate the backup file
4. On the *'Installation Type'* page, select *'Restore from Backup'*
5. Choose the backup zip file and follow the instructions
6. Optionally, log files may be restored from server backups to:
   - `C:\Program Files\Moonwalk\logs\AdminCenter\`
   - `C:\Program Files\Moonwalk\logs\drtool\`

Note: Restore from any backup (or virtual machine snapshot) that is *more than 6 months old* may fail due to certificate expiry. Contact Moonwalk Support for assistance if this occurs.

3.3.2 Per-Server Logs

Backing up the configuration on each server is not necessary since such configuration is already included in the above Admin Tools backup process. You may optionally backup server logs from the logs location on each server – by default these are located at:

- Windows: `C:\Program Files\Moonwalk\logs\Agent`
- Linux: `/var/opt/moonwalk/moonwalk-agent/log`

Note: Do not attempt to restore logs back into an active installation, since this will interfere with log rotation.

3.4 Storage Backup

Each stub on primary storage is linked to a corresponding MWI file on secondary storage. During the normal process of migration and demigration the relationship between stub and MWI file is maintained.

The recommendations below ensure that the consistency of this relationship is maintained even after files are restored from backup.

3.4.1 Backup Planning

Ensure that the restoration of stubs is included as part of your backup & restore test regimen.

When using Scrub policies, ensure the Scrub grace period is sufficient to cover the time from when a backup is taken to when the restore *and* Post-Restore Revalidate steps are completed (see below).

It is **strongly** recommended to set the global *minimum* grace period accordingly to guard against the accidental creation of scrub policies with insufficient grace. This setting may be configured on that AdminCenter *'Settings'* page.

Important: It will *NOT* be possible to safely restore stubs or MigLinks from a backup set taken more than one grace period ago.

Additional Planning

To complement standard backup and recovery solutions, and to allow the widest range of recovery options, it is recommended to schedule a *'Create Recovery File From Source'* Policy to run after each migration.

3.4.2 Backup Process

Perform these backup steps in the following order:

1. Backup primary storage volumes
2. Backup secondary volumes/devices (if necessary)
    - Allow primary backup to **complete** first
    - Secondary may be backed up less frequently than primary

Usually, backup will be scheduled to run a little while after migration policies have completed.

Note: When adding backup jobs, always recheck the minimum grace period setting for scrub (see above).

3.4.3 Restore Process

If primary *and* secondary volumes are to be restored:

1. Suspend the scheduler in AdminCenter
2. Restore the primary volume
3. Restore the corresponding secondary volume from a **newer** backup set than the primary
4. Run a *'Post-Restore Revalidate'* policy against the primary volume

   - To ensure all stubs are revalidated, run this policy against the **entire** primary volume, NOT simply against the migration source
   - This policy is not required when *only* WORM destinations are in use
5. Restart the scheduler in AdminCenter

If *only* primary is to be restored (including where secondary is cloud storage):

1. Suspend the scheduler in AdminCenter
2. Restore the primary volume
3. Run a *'Post-Restore Revalidate'* policy against the primary volume
   - To ensure all stubs are revalidated, run this policy against the **entire** primary volume, NOT simply against the migration source
   - This policy is not required when *only* WORM destinations are in use
4. Restart the scheduler in AdminCenter

If restoring the primary volume to a different server (a server with a different FQDN), the following preparatory steps will also be required:

1. On the *'Servers'* page, retire the old server (unless still in use for other volumes)
2. Install Agent on the new server
3. Update Sources as required to refer to the FQDN of the new server
4. Perform the restore process as above

3.4.4 Platform-specific Considerations

Windows

Most enterprise Windows backup software will respect Moonwalk stubs and will back them up correctly without causing any unwanted demigrations. For some backup software, it may be necessary to refer to the software documentation for options regarding Offline files.

When testing backup software configuration, test that backup of stubs does not cause unwanted demigration.

Additional backup testing may be required if Stub Deletion Monitoring is required. Please refer to Appendix F for more details.

NetApp Filers

Please consult §5.3.5 (p.45) for information regarding snapshot restore on NetApp Filers.

3.5 Production Readiness Checklist

Backup

1. Check your Moonwalk configuration is adequately backed up – see §3.3
2. Review the *storage* backup and restore procedures described in §3.4
3. Check backup software can backup stubs without triggering demigration

4. Check backup software restores stubs and that they can be demigrated
5. Schedule regular *'Create Recovery File From Source'* Policies on your migration sources – see §4.18 (p.31)

Antivirus

Generally, antivirus software will not cause demigrations during normal file access. However, some antivirus software will demigrate files when performing scheduled file system scans.

Prior to production deployment, always check that installed antivirus software does not cause unwanted demigrations. Some software must be configured to skip offline files in order to avoid these inappropriate demigrations. Consult the antivirus software documentation for further details.

If the antivirus software does *not* provide an option to skip offline files during a scan, Moonwalk Agent may be configured to deny demigration rights to the antivirus software. Refer to Appendix F for more information.

It may be necessary for some antivirus products to exempt the Moonwalk Agent process from real-time protection (scan-on-access). If the exclusion configuration requires the path of the executable to be specified, be sure to update the exclusion whenever Moonwalk is upgraded (since the path will change on upgrade).

Other System-wide Applications

Check for other applications that open all the files on the whole volume. Audit scheduled processes on file servers – if such processes cause unwanted demigration, it may be possible to block them (see Appendix F).

Monitoring and Notification

To facilitate proactive monitoring, it is recommended to:

1. Configure email notifications to monitor system health and Task activity
2. Enable syslog – see Appendix F

Platform Considerations

For further information on platform-specific interoperability considerations, please refer to the appropriates sections of Chapter 5.

3.6 Policy Tuning

Periodically re-assess file distribution and access behavior:

1. Run 'Gather Statistics' Policies
    - Examine reports

2. Examine Server statistics – see §1.4.1 (p.3)
    - For more detail, examine demigrates in file server `agent.log` files

Consider:

- Are there unexpected peaks in demigration activity?
- Are there any file types that should not be migrated?
- Should different rules be applied to different file types?
- Is the Migration Policy migrating data that is regularly accessed?
- Are the Rules aggressive enough or too aggressive?
- What is the data growth rate on primary and secondary storage?
- Are there subtrees on the source file system that should be addressed by separate policies or excluded from the source entirely?

3.7 System Upgrade

When a Moonwalk deployment is upgraded from a previous version, Admin Tools must always be upgraded first, followed by *all* Server components.

Run:

- `Moonwalk Admin Tools.exe`

3.7.1 Automated Server Upgrade

Where possible, it is advisable to upgrade Server agents using the automated upgrade feature by clicking the UPGRADE SYSTEM icon on the *'Servers'* page.

The automated process transfers installers to each server and performs the upgrades in parallel to minimize downtime. If a server fails or is offline during the upgrade, manually upgrade it later. Once the automated upgrade procedure is finalized, the *'Servers'* page will update to display the health of the upgraded servers.

Following the upgrade, resolve any warnings displayed on the *'Dashboard'*.

3.7.2 Manual Server Upgrade

Follow the instructions appropriate for the platform of each server as described below.

Agent for Windows

1. Run `Moonwalk Agent.exe` and follow the instructions
2. Resolve any warnings displayed on the *'Dashboard'*

Gateway Agent for Linux (RHEL)

In a root terminal:

1. `tar xzf Moonwalk_Gateway_Agent_rhel9_2024_1.tgz`
2. `Moonwalk_Gateway_Agent_rhel9_2024_1/install.sh`
3. Resolve any warnings displayed on the *'Dashboard'*

Gateway Agent for Linux (Ubuntu)

In a root terminal:

1. `tar xzf Moonwalk_Gateway_Agent_ubuntu22_2024_1.tgz`
2. `Moonwalk_Gateway_Agent_ubuntu22_2024_1/install.sh`
3. Resolve any warnings displayed on the *'Dashboard'*

NetApp FPolicy Server

1. Run `Moonwalk NetApp FPolicy Server.exe` and follow the instructions
2. Resolve any warnings displayed on the *'Dashboard'*

LinkConnect Server

1. Run `Moonwalk LinkConnect Server.exe` and follow the instructions
2. Resolve any warnings displayed on the *'Dashboard'*

Chapter 4

Policy Operation Reference

This chapter describes the various operations that may be performed on selected files by AdminCenter policies.

4.1 Gather Statistics Operation

Requires: Source(s)
Included in Starter Edition: yes

Generate statistics report(s) for file sets at the selected Source(s). If desired, Rules may be used to specify a subset of files on which to report rather than the whole source.

Optionally include statistics by file owner. By default, owner statistics are omitted which generally results in a faster policy run.

Optionally include object storage archive status information (rather than only object storage class) where supported by the Source's object storage platform.

Gather Statistics Policies can also be configured to export per-file metadata – optionally including ACL information – in JSON or CSV format.

4.2 Migrate Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: yes

Migrate file data from selected Source(s) to a Destination. Stub files remain at the Source location as placeholders until files are demigrated. File content will be transparently demigrated (returned to primary storage) when accessed by a user or application. Stub files retain the original logical size and file metadata. Files containing no data will not be migrated.

Each Migrate operation will be logged as a Migrate, Remigrate, or Quick-Remigrate.

A Remigrate is the same as a Migrate except it explicitly recognizes that a previous version of the file had been migrated in the past and that stored data pertaining to that previous version is no longer required and so is eligible for removal via a Scrub policy.

A Quick-Remigrate occurs when a file has been demigrated and NOT modified. In this case it is not necessary to retransfer the data to secondary storage so the operation can be performed very quickly. Quick-remigration does **not change the secondary storage location** of the migrated data.

Optionally:

- Quick-remigration of files demigrated within a specified number of days may be prevented – this option can be used to avoid quick-remigrations occurring in an overly aggressive fashion
- Policies may be configured to pause during the globally configured work hours
- Sparse files may be skipped – it is often undesirable to migrate files that are highly sparse since sparseness is not preserved by the migration process

If using a capacity-based license, Migrates and Remigrates (but not Quick-Remigrates) consume capacity license quota.

Note: For Sources using a LinkConnect Server, such as Dell EMC OneFS shares, see §4.9 instead.

4.3 Quick-Remigrate Operation

Requires: Source(s), Rule(s)
Included in Starter Edition: yes

Quick-Remigrate demigrated files that do not require data transfer, enabling space to be reclaimed quickly. This operation acts only on files that have not been altered since the last migration.

Optionally:

- Quick-remigration of files demigrated within a specified number of days may be prevented – this option can be used to avoid quick-remigrations occurring in an overly aggressive fashion
- Policies may be configured to pause during the globally configured work hours

Capacity license quota is not consumed.

4.4 Scrub Destination Operation

Requires: Destination (non-WORM)
Included in Starter Edition: yes

Remove unnecessary stored file content from a migration destination. This is a maintenance policy that should be scheduled regularly to reclaim space (and license quota if using capacity-based licensing) .

A grace period must be specified which is sufficient to cover the time from when a backup is taken to when the restore and corresponding Post-Restore Revalidate policy

would complete. The grace period effectively delays the removal of data sufficiently to accommodate the effects of restoring primary storage from backup to an earlier state.

Use of Scrub is usually desirable to maximize storage efficiency. In order to also maximize performance benefits from quick-remigration, it is advisable to schedule migration / quick-remigration policies more frequently than the grace period.

To avoid interactions with Migration policies, Scrub tasks are automatically paused while migration-related tasks are in progress.

Scrub policies may be configured to generate log output only without actually removing files.

Important: Source(s) MUST be backed up within the grace period.

4.5 Post-Restore Revalidate Operation

Requires: Source(s)
Included in Starter Edition: yes

Scan all stubs present on a given Source, revalidating the relationship between the stubs and the corresponding files on secondary storage. This operation is required following a restore from backup and should be performed on the **root** of the restored source volume.

If *only* Write Once Read Many (WORM) destinations are in use, this policy is not required.

Important: This revalidation operation MUST be integrated into backup/restore procedures, see §3.4.1 (p.18).

4.6 Demigrate Operation

Requires: Source(s), Rule(s)
Included in Starter Edition: yes

Return migrated file content back to files on the selected Source(s). This is useful when a large batch of files must be demigrated in advance.

Prior to running a Demigrate policy, be sure that there is sufficient primary storage available to accommodate the demigrated data.

This operation may be used with both Migrated and Link-Migrated files.

4.7 Advanced Demigrate Operation

Requires: Source(s), Rule(s)
Included in Starter Edition: yes

Demigrates files with advanced options:

- **Disconnect files from destination** – remove destination information from demigrated files (both files demigrated by this policy and files that have already been demigrated); it will no longer be possible to quick-remigrate these files
- A **Destination Filter** may optionally be specified in order to demigrate/disconnect only files that were migrated to a particular destination

Prior to running an Advanced Demigrate policy, be sure that there is sufficient primary storage available to accommodate the demigrated data.

4.8 Premigrate Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: yes

Premigrate file data from selected Source(s) to a Destination in preparation for migration. Files on primary storage will not be converted to stubs until a Migrate or Quick-Remigrate Policy is run. Files containing no data will not be premigrated.

This can assist with:

- A requirement to delay the stubbing process until secondary storage backup or replication has occurred
- Reduction of excessive demigrations while still allowing an aggressive Migration Policy.

Premigration is, as the name suggests, intended to be followed by full migration/quick-remigration. If this is not done, a large number of files in the premigrated state may slow down further premigration policies, as the same files are rechecked each time.

By default, files already premigrated to another destination will be skipped when encountered during a premigrate policy.

Optionally:

- Policies may be configured to pause during the globally configured work hours
- Sparse files may be skipped – it is often undesirable to migrate files that are highly sparse since sparseness is not preserved by the migration process

If using a capacity-based license, capacity license quota is consumed.

Note: Most deployments will not use this operation, but will use a combination of Migrate and Quick-Remigrate instead.

4.9 Link-Migrate Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: no

For platforms that do not support standard stub-based migration, Link-Migrate file data from selected Source(s) to a Destination.

Files at the source location will be replaced with Moonwalk-encoded links (MigLinks) which allow client applications to transparently read data without returning files to primary storage. If an application attempts to modify a link, the file will be automatically

returned to primary storage and then modified in-place. Files containing no data will be skipped by this policy.

MigLinks present the original logical size and file metadata.

Since MigLinks remain links when *read* by client applications, there is no analogue of quick-remigration for link-migrate.

Optionally:

- Policies may be configured to pause during the globally configured work hours
- Sparse files may be skipped – it is often undesirable to Link-Migrate files that are highly sparse since sparseness is not preserved by the migration process

If using a capacity-based license, Link-Migrates consume capacity license quota.

4.10 Change Destination Tier Operation

Requires: Source(s), Rule(s), Destination Filter, New Destination
Included in Starter Edition: no

Change migration tier of selected files that are already migrated or link-migrated to a secondary storage destination by copying the secondary storage data to the new Destination and then updating the stubs / MigLinks accordingly. The defunct copy on the original destination will be removed by a subsequent Scrub Policy (scheduled separately and subject to the configured grace period).

This operation is typically used to realize a multi-tier environment, in which files can be aged across storage tiers over time. This is not to be confused with the Retarget Destination operation (§4.11), which caters for the decommissioning of a secondary storage target.

If desired, rules can be used to apply different tiering criteria to different subsets of the data.

If using a capacity-based license, capacity license quota will be transferred to the new destination files – scrubbing the original destination will not recover quota.

This policy may be configured to pause during the globally configured work hours.

Note: Files migrated to WORM destinations cannot be moved to another tier. These files will be skipped.

4.11 Retarget Destination Operation

Requires: Source(s), Destination Filter, New Destination
Included in Starter Edition: no

Permanently retarget stubs or MigLinks to a new migration destination. This operation is intended for use when completely decommissioning an old migration destination – the defunct copy of the file content will *not* be removed from the original destination either by this operation or by subsequent Scrub operations.

Any old data that is no longer referenced by stubs or MigLinks will *not* be transferred to the new destination.

This operation does not affect capacity license quota.

Warning: This policy must be run on the root of all volumes that may contain stubs on all servers prior to finally decommissioning the old migration destination. **Re-run the policies** as necessary until there are no more files to retarget.

This policy may be configured to pause during the globally configured work hours.

4.12 Ingest Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: yes

Ingest files into a cloud or object store Destination. Original filenames and paths are preserved.

Optionally, original file metadata – including a user-specified custom-metadata field – can be attached to the destination objects. The original files' security details can also be attached. Please see the relevant section of Chapter 5 for the specifics of how metadata is stored on a given destination.

An Ingest policy must include a *'Collision Behavior'* option to specify the action to be taken when ingesting files that would overwrite existing objects at the destination. Supported behaviors are:

Collision Behavior	**Description**
Overwrite existing file	Overwrites objects unconditionally
Always append timestamp	Appends an ISO8601-compatible timestamp suffix to ALL object names, e.g. `file.txt_20200401T115959Z`

Since partially ingested datasets are not typically useful, the policy will perform a configurable number of attempts to ingest each object. Such retries are appropriately logged for troubleshooting.

Optionally, for data that will no longer be required on primary storage after ingestion, the original files may be deleted on a file-by-file basis as each file is successfully ingested.

A Content-Disposition header can be automatically added to ingested HTTP objects such that user download via a browser will preserve the original filename (e.g. file.txt without any suffixes).

Where security details are to be attached to ingested objects, the format is source-platform specific. For Windows files, a Security Descriptor (owner, group and ACLs) is recorded in Microsoft SDDL format for each object. For NFS sources, UID, GID and octal permissions are recorded as separate fields.

An upper limit may be specified to prevent very large security details field values being attached to an object. Values exceeding this limit will be replaced with the supplied default value (or omitted entirely if a default is not provided). SDDL default values must be prepended with '`win:`'.

4.13 Restore S3 Archive

Requires: Source(s), Rule(s)
Included in Starter Edition: yes

Enqueue S3 objects for temporary restoration from an archive storage class such as Glacier or Deep Archive. Once restore requests are issued, this policy will wait for the object store to return all of the objects (which may take a significant amount of time).

The performance of restoration operations can be selected, from the cheapest / slowest option (bulk restore), to the fastest / most expensive (expedited restore).

The number of days that restored data will remain available must also be specified – if any objects are already restored or are currently being restored by other processes, this Policy will ensure that availability is extended if necessary to at least the specified number of days. After this period, objects will revert to their archived state.

4.14 Set Azure Access Tier

Requires: Source(s), Rule(s)
Included in Starter Edition: yes

Set the storage tier of objects in Azure storage. Use this Policy to transition objects between 'Hot', 'Cool' and 'Archive' tiers.

Any objects that are transitioned out of the 'Archive' tier will be processed asynchronously – this policy will wait for Azure to transition all of the objects (which may take a significant amount of time).

The performance of archive retrieval operations can be specified, from the cheapest / slowest option, to the fastest / most expensive.

4.15 Copy Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: no

Copy files from Source(s) to a filesystem Destination.

When a migrated stub or MigLink is copied, a full file (not a stub) will be created at the Copy Destination without demigrating at the Source.

Note: Where Source and Destination use radically different filesystems, some filenames may not be representable at the Destination and will be uncopyable. Similarly, Source filenames that are not considered unique by the Destination filesystem – e.g. differing only by case – will conflict with each other. Such conflicts may be addressed by using the *'rename'* overwrite option.

4.16 Move Operation

Requires: Source(s), Rule(s), Destination
Included in Starter Edition: no

Move files from Source(s) to a filesystem Destination.

If possible, the Move Operation will move migrated stubs without demigration – a stub-move – creating a stub at the Move Destination. A stub-move will be performed, rather than a full file move if:

- The Destination supports migration
- The Agent at the Destination is not a Gateway

Moving a link-migrated file always results in a full file move.

Optionally, a Move policy can be configured to force demigration of stubs during transit if desired. However, in the case of Link-Migrated files, destination files are always full files, not links following the operation.

Note: Where Source and Destination use radically different filesystems, some filenames may not be representable at the Destination and will be uncopyable. Similarly, Source filenames that are not considered unique by the Destination filesystem – e.g. differing only by case – will conflict with each other. Such conflicts may be addressed by using the *'rename'* overwrite option.

4.17 Additional Copy and Move Options

- *'Path'* options:
  - **preserve** indicates that the relative path of each file, from the root of its source, should be preserved at the destination
  - **prepend server and volume name** prepends an extra directory component such as `server_volname` to the path
  - **copy full directory structure** copies all directories encountered when traversing a Source, even those that do not contain any files which match the specified Rules
  - **preserve directory metadata** preserves metadata for directories as well as files

- *'Metadata'* options:
  - **preserve timestamps** preserves time and date information
  - **preserve attribute flags** preserves flags such as 'Read-Only', 'Hidden', etc.
  - **preserve Unix ownership and permissions** preserves owner, group and permissions on NFS (this option must first be enabled on the *'Settings'* Page)

- *'Overwrite'* options:
  - **always** – a file moved/copied from a Source will always overwrite any identically named file already existing at the corresponding location on the Destination
  - **never** – never overwrites existing files at the destination
  - **rename** – clashing filenames are disambiguated by renaming the new file. For example, if a file named `letter.doc` is copied to a Destination where a file with this name already exists, the new file would be renamed to `letter[1].doc` (or, if that too exists, `letter[2].doc`, and so forth)
  - **if newer** – only overwrites the file if the Source file is newer than the Destination file

4.18 Create Recovery File From Source Operation

Requires: Source(s), Rule(s)
Included in Starter Edition: no

Generate a disaster recovery file for Moonwalk DrTool by analyzing files at the selected Source(s). DrTool can use the generated file(s) to recover or update source files.

Note: Recovery files generated from *Source* will account for renames.

4.19 Create Recovery File From Destination Operation

Requires: Destination
Included in Starter Edition: no

Generate a disaster recovery file for Moonwalk DrTool by analyzing files at the selected Destination without reference to the associated primary storage files.

Note: Recovery files from *Destination* may not account for renames.

Important: It is strongly recommended to use *'Create Recovery File From Source'* in preference where possible.

4.20 Delete Operation

Requires: Source(s), Rule(s)
Included in Starter Edition: no

Delete files from Source(s).

Important: Deletion of files cannot be undone.

4.21 Platform Transfer Operation

Requires: Source, Rule(s), Destination
Included in Starter Edition: no

This operation is provided for professional services partners under special licensing.

Chapter 5

Source and Destination Reference

The following pages describe the characteristics of the Sources and Destinations supported by Moonwalk. Planning, setup, usage and maintenance considerations are outlined for each storage platform.

Important: Read any relevant sections of this chapter prior to deploying Moonwalk in a production environment.

5.1 Microsoft Windows

5.1.1 Migration Support

Windows NTFS and ReFS volumes may be used as migration sources. To access files over SMB, see §5.19.

Windows stub files can be identified by the 'O' (Offline) attribute in Explorer. Depending on the version of Windows, files with this flag may be displayed with an overlay icon.

Windows volumes may also be used as Destinations (not supported by *Starter Edition* licenses).

Note: If it is not possible to install the Moonwalk Agent directly on the file server, see §5.2 for an alternative solution using Link-Migration.

5.1.2 Planning

Prerequisites

- A license that includes an appropriate entitlement for Windows

When creating a production deployment plan, please refer to §3.5 (p.19).

Cluster Support

Clustered volumes managed by Windows failover clusters are supported. However, the Cluster Shared Volume (CSVFS) feature is NOT supported. As a result, when configuring a 'File Server' role in the Failover Cluster Manager, 'File Server for general use' is the only supported File Server Type. The 'Scale-Out File Server for application data' File Server Type is NOT supported.

When using clustered volumes in Moonwalk URIs, ensure that the resource FQDN appropriate to the volume is specified rather than the FQDN of any individual node.

5.1.3 Setup

Installation

See Installing Agent for Windows §2.2.2 (p.12)

5.1.4 Interoperability

This section describes Windows-specific considerations only and should be read in conjunction with §3.5 (p.19).

Microsoft Storage Replica

Moonwalk supports Microsoft Storage Replica.

If Storage Replica is configured for *asynchronous* replication, a disaster failover effectively reverts the volume to a previous point in time. As such, this kind of failover is directly equivalent to a volume restore operation (albeit to a very recent state).

As with any restore, a Post-Restore Revalidate Policy (see §4.5 (p.25)) should be run across the restored volume within the scrub grace period window. This will ensure correct operation of any future scrub policies by accounting for discrepancies between the demigration state of the files on the (failed) replication source volume and the replication destination volume.

Important: Integrate this process into your recovery procedures prior to production deployment of asynchronous storage replication.

Microsoft DFS Namespaces (DFSN)

DFSN is supported. Moonwalk Sources must be configured to access volumes on individual servers directly rather than through a DFS namespace. Users and applications may continue to access files and stubs via DFS namespaces as normal.

Microsoft DFS Replication (DFSR)

Agents must be installed (selecting the *migration* role during installation) on **EACH** member server of a DFS Replication Group prior to running migration tasks on any of the group's Replication Folders.

If adding a new member server to an existing Replication Group where Moonwalk is already in use, Agent must be installed on the new server first.

When running policies on a Replicated Folder, sources should be defined such that each policy acts upon **only one** replica. DFSR will replicate the changes to the other members as usual.

Read-only (one-way) replicated folders are NOT supported. However, read-only SMB shares can be used to prevent users from writing to a particular replica as an alternative.

Due to the way DFSR is implemented, care should be taken to avoid *writing* to stub files that are being concurrently accessed from another replica.

In the rare event that DFSR-replicated data is restored to a member from backup, ensure that DFSR services on all members are running and that replication is **fully up-to-date** (check for the DFSR 'finished initial replication' Windows Event Log message), then run a Post-Restore Revalidate Policy using the same source used for migration.

Note: If using a capacity-based license, no additional capacity license quota is consumed when stubs are replicated by DFSR.

Retiring a DFSR Replica

Retiring a replica effectively creates two *independent* copies of each stub, without updating secondary storage. To avoid any potential loss of data:

1. Delete the contents of the retired replica (preferably by formatting the disk, or at least disable Stub Deletion Monitoring during the deletion)
2. Run a Post-Restore Revalidate Policy on the remaining copy of the data

If it is strictly necessary to keep both, now independent, copies of the data and stubs, then run a Post-Restore Revalidate Policy on **both** copies separately (not concurrently).

Preseeding a DFSR Replicated Folder Using Robocopy

The most common use of Robocopy with Moonwalk stubs is to preseed or stage initial synchronization. When performing such a preseeding operation:

- for new Replicated Folders, ensure that the 'Primary member' is set to be the original server, not the preseeded copy
- both servers must have Agent installed **before** preseeding
- add a "Process Exclusion" to Windows Defender for `robocopy.exe` (allow a while for the setting to take effect)
- on the source server, preseed by running robocopy with the `/b` flag (to copy stubs as-is to the new server)
- once preseeding is complete and replication is **fully up-to-date** (check for the DFSR 'finished initial replication' Windows Event Log message), it is recommended to run a Post-Restore Revalidate Policy on the original Moonwalk Source

Note: If the process above is aborted, be sure to delete all preseeded files and stubs (preferably by formatting the disk, or at least disable Stub Deletion Monitoring during the deletion) and then run a Post-Restore Revalidate Policy on the original Moonwalk Source.

Robocopy (Other Uses)

Robocopy will, by default, demigrate stubs as they are copied. This is the same behavior as Explorer copy-paste, xcopy etc..

Robocopy with the `/b` flag (backup mode – must be performed as an administrator) will copy stubs as-is.

Robocopy /b is not recommended. If stubs *are* copied in this fashion, the following must be considered:

- for a copy from one server to another, both servers must have Moonwalk Agent installed
- this operation is essentially a backup and restore in one step, and thus inappropriately duplicates stubs which are intended to be unique
  - after the duplication, one copy of the stubs should be deleted immediately
  - run a Post-Restore Revalidate policy on the remaining copy
  - this process will render the corresponding secondary storage files non-scrubbable, even after they are demigrated
- to prevent Windows Defender triggering demigrations when the stubs are accessed in this fashion:
  - always run the robocopy from the source end (the file server with the stubs)
  - add a "Process Exclusion" to Windows Defender for `robocopy.exe` (allow a while for the setting to take effect)

Windows Data Deduplication

If a Windows source server is configured to use migration policies and Windows Data Deduplication, it should be noted that a given file can either be deduplicated or migrated, but not both at the same time. Moonwalk migration policies will automatically skip files that are already deduplicated. Similarly, Windows will skip Moonwalk stubs when deduplicating.

When using both technologies, it is recommended to configure Data Deduplication and Migration based on file type such that the most efficacious strategy is chosen for each type of file.

Note: Microsoft's legacy Single Instance Storage (SIS) feature is not supported. Do not use SIS on the same server as Moonwalk Agent.

Windows Shadow Copy

Windows Shadow Copy – also known as Volume Snapshot Service (VSS) – allows previous versions of files to be restored, e.g. from Windows Explorer.

Users cannot perform self-service restoration of *stubs*. However, an administrator may restore specific stubs or sets of stubs from snapshots by following the procedure outlined below. Be sure to provide this procedure to all administrators.

1. In AdminCenter Browser, navigate to the relevant share using an `smb://` URI
2. Click POINT-IN-TIME to select a snapshot
3. Select the desired files, stubs and folders
4. Click DOWNLOAD

For larger jobs, where a Policy is preferable, the point-in-time feature is also available when creating a Source.

5.1.5 Behavioral Notes

Symbolic Links

Symbolic links (symlinks) will be skipped during traversal of the file system. This ensures that files are not seen – and thus acted upon – multiple times during a single execution of a given policy. If it is intended that a policy should apply to files within a directory referred to by a symbolic link, either ensure that the Source encompasses the real location at the link's destination, or specify the link itself as the Source.

Mount Points / Junctions

Mount Points are always traversed.

By default all other directory junctions will be traversed as expected, unless the junction represents a legacy path, identifiable by its ACL explicitly denying the *Everyone* SID the List Folder permission.

If it is desired to omit traversal of directory junctions, add the following setting to the Server's manual override settings in AdminCenter:

```
Windows.TraverseNonVolumeJunctions=false
```

Mount-DiskImage

On Windows 8 or above, VHD and ISO images may be mounted as normal drives using the PowerShell Mount-DiskImage cmdlet. This functionality can also be accessed via the Explorer context menu for an image file.

A known limitation of this cmdlet is that it does not permit *sparse* files to be mounted (see Microsoft KB2993573). Since migrated image files are always sparse, they must be demigrated prior to mounting. This can be achieved either by copying the file or by removing the sparse flag with the following command:

```
fsutil sparse setflag <file_name> 0
```

5.1.6 Stub Deletion Monitoring

On Windows, the Agent can monitor stub deletions to identify secondary storage files that are no longer referenced in order to maximize the usefulness of Scrub Policies. This feature extends not only to stubs that are directly deleted by the user, but also to other cases of stub file destruction such as overwriting a stub or renaming a different file over the top of a stub.

As of Moonwalk 12.1u2, Stub Deletion Monitoring is disabled by default. To enable it, please refer to Appendix F.

5.2 Microsoft Windows via LinkConnect Server

5.2.1 Link-Migration Support

This section details the configuration of a Moonwalk LinkConnect Server to enable Link-Migration of files from Windows Server SMB shares. This option should be used when it is not possible to install Moonwalk Agent directly on the Windows file server in question. For other cases – where Agent *can* be installed on the server – please refer to §5.1.

Refer to §4.2 (p.23) and §4.9 (p.26) for details of the Migrate and Link-Migrate operations respectively.

Link-Migration works by pairing a Windows SMB share with a corresponding LinkConnect Cache Share. Typically a top-level share on each Windows file server volume is mapped to a unique share (or subdivision) on a LinkConnect Server. Multiple file server shares may use Cache Shares / subdivisions on the same LinkConnect Server if desired.

Once this configuration is completed, Link-Migrate policies convert files on the source Windows Server SMB share to links pointing to the destination files via the LinkConnect Cache Share, according to configured rules.

Link-Migrated files can be identified by the 'O' (Offline) attribute in Explorer. Depending on the version of Windows, files with this flag may be displayed with an overlay icon.

5.2.2 Planning

Prerequisites

- An NTFS Cache Volume of at least 1TB – see §2.2.5 (p.13)
- A Moonwalk license that includes an entitlement for LinkConnect Server.

When creating a production deployment plan, please refer to §3.5 (p.19).

Note: It is recommended that a LinkConnect Server is only associated with one type of SMB device. For example, do not associate a single LinkConnect Server with both Windows and OneFS shares. This is because agent configuration options may need to be tuned differently to best work with the different platforms.

File Server System Requirements

- Windows Server 2016 or higher
- The server must NOT have the Active Directory Domain Services role

Client Requirements

Windows clients require a supported 64-bit Windows operating system:

- Windows 11
- Windows 10
- Windows Server 2022
- Windows Server 2019

- Windows Server 2016

In order to access link-migrated files, the LinkConnect Client Driver must be installed on each client machine – see §2.3 (p.14).

Network

Place the Moonwalk LinkConnect Server on the same subnet and same switch as the corresponding Windows file server(s) to minimize latency.

Additionally, the LinkConnect Server **must** be joined to the same domain as the Windows file server.

Antivirus Considerations

Ensure that Windows Defender or any other antivirus product installed on the LinkConnect Server is configured to **omit** scanning/screening on the LinkConnect Cache Volume **and** any Windows file server SMB shares.

High-Availability for LinkConnect Server

Consider whether High-Availability (HA) is required in your environment (either now *or in the future*). If so, LinkConnect Servers must be deployed in a DFSN configuration from the outset.

LinkConnect Cache Shares are configured for HA by exposing the share name at the domain level using DFSN. If not using HA, it is possible to use either a simple share on a standalone server, or a share exposed at the domain level using DFSN. The latter is always recommended to allow transition to an HA configuration in the future.

Regular Maintenance Activity

Each configured MigLink source will be periodically scanned to perform maintenance tasks such as MigLink ACL propagation and Link Deletion Monitoring (see below).

In an HA configuration, this scanning activity will be performed by a single caretaker node, as can be seen on the AdminCenter Servers page. A standalone LinkConnect Server always performs the caretaker role.

Security Considerations

Files with certain ACLs cannot be link-migrated – they will be skipped during Link-Migrate Policies. If such an ACL is set on a file that has already been link-migrated, the new ACL will NOT be propagated to the LinkConnect Server. Specifically:

- Conditional ACEs will not be link-migrated (consider using Central Access Policies instead where applicable)
- Audit ACEs which track attempted read access will not be link-migrated (Audit ACEs which track e.g. write or delete work as expected)

When using Dynamic Access Control, Central Access Policies must be propagated to the LinkConnect Server as well as the file servers.

Link Deletion Monitoring

Link Deletion Monitoring (LDM) may be enabled on a per-share basis.

Similar to the Stub Deletion Monitoring feature provided by Moonwalk Agents on Windows, LDM identifies secondary storage files that are no longer referenced in order to facilitate recovery of storage space by Scrub Policies. This feature extends not only to MigLinks that are demigrated or directly deleted by the user, but also to other cases such as overwriting a MigLink or renaming a different file over the top of a MigLink.

Unlike SDM, LDM requires a number of maintenance scans to determine that a given secondary storage file is no longer referenced. It should be noted that interrupting the maintenance process (e.g. by restarting the caretaker node or transitioning the caretaker role) will delay the detection of unreferenced secondary storage. For optimal and timely storage space recovery, ensure that LinkConnect Servers can run uninterrupted for extended periods.

Warning: In order to avoid LDM incorrectly identifying files as deleted – leading to unwanted data loss during Scrub – it is critical to ensure that users cannot *move/rename* MigLinks out of the scanned portion of the directory tree within the filesystem. This is typically done by configuring the managed shares at the root of each data volume.

5.2.3 Setup

Installation

Provision a user on the Active Directory domain for the **exclusive** use of your LinkConnect service(s). This user does *not* need to be a member of Domain Admins.

On each LinkConnect Server machine:

1. Add the LinkConnect user to the *local* Administrators group
2. Assign the 'Log on as a service' privilege to this user
3. Run the `Moonwalk LinkConnect Server.exe`
4. Follow the prompts to complete the installation
5. Follow the instructions to activate the installation

Configuring the LinkConnect Server

In the Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the LinkConnect Server.

In the *'Configuration'* panel, select *'LinkConnect Server'*, then use the wizard to add shares.

Windows supports features such as Dynamic Access Control that are not available on all SMB implementations. Support for these features is disabled by default. Such additional security information may be useful in certain advanced recovery situations. Gathering this information may optionally be enabled in the *'Manual Overrides'* panel by entering:

- SMB.SACL.mandatoryLabel=true
- SMB.SACL.scopedPolicy=true
- SMB.SACL.resourceAttribute=true

5.2.4 Additional Shares

Further 'top-level' shares can be added using the *'Configuration'* panel as above.

For all other sub-shares (that is, shares within the directory tree of a registered top-level share, including shares that are simply aliases for top-level shares), simply add 'Full Control' permissions for the LinkConnect user. Be sure to configure the share, not the folder permissions.

5.2.5 Usage

URI format

`smb://{server}/{nas}/{share}/[{path}/]`

Where:

- `server` – FQDN of a LinkConnect Server that is configured to support the file server share
- `nas` – Windows file server FQDN
- `share` – Windows file server SMB share
- `path` – path within the share

Example:

`smb://link.example.com/winserver.example.com/pub/projects/`

5.3 NetApp Filer

This section describes support for NetApp Filers.

5.3.1 Migration Support

Migration support for sources on NetApp Vservers (Storage Virtual Machines) is provided via NetApp FPolicy. This requires the use of a Moonwalk FPolicy Server. Client demigrations can be triggered via SMB or NFS client access.

Please note that NetApp Filers currently support FPolicy for Vservers with FlexVol volumes but not Infinite volumes.

When accessed via SMB on a Windows client, NetApp stub files can be identified by the 'O' (Offline) attribute in Explorer. Files with this flag may be displayed with an overlay icon. The icon may vary depending on the version of Windows on the client workstation.

Note: The `netapp://` scheme described in this section cannot be used in a migration *destination*. To migrate *to* a NetApp filer, it is recommended to use NFS (see also §5.18).

5.3.2 Planning

Prerequisites

- NetApp Filer(s) must be licensed for the particular protocol(s) to be used (FPolicy requires an SMB license)
- A Moonwalk license that includes an entitlement for NetApp FPolicy Server

Moonwalk FPolicy Servers require **EXCLUSIVE** use of SMB connections to their associated NetApp Vservers. This means Explorer windows must not be opened, drives must not be mapped, nor should any UNC paths to the filer be accessed from the FPolicy Server machine. Failure to observe this restriction will result in unpredictable FPolicy disconnections and interrupted service.

When creating a production deployment plan, please refer to §3.5 (p.19).

Filer System Requirements

Moonwalk FPolicy Server requires that the Filer is running:

- Data ONTAP version 9.7+

Network

Each FPolicy Server should have exactly one IP address.

Place the FPolicy Servers on the same subnet and same switch as their corresponding Vservers to minimize latency.

Antivirus Considerations

Ensure that Windows Defender or any other antivirus product installed on FPolicy Server machines is configured to **omit** scanning/screening NetApp shares.

Antivirus access to NetApp files will interfere with the correct operation of the FPolicy Server software. Antivirus protection should still be provided on client machines and/or the NetApp Vservers themselves as normal.

High-Availability for FPolicy Servers

It is strongly recommended to install Moonwalk FPolicy Servers in a High-Availability configuration. This configuration requires the installation of Moonwalk FPolicy Server on a group of machines which are addressed by a single FQDN. This provides High-Availability for migration and demigration operations on the associated Vservers.

Typically a pair of FPolicy Servers operating in HA will service all of the Vservers on a NetApp cluster.

Note: The servers that form the High-Availability FPolicy Server configuration must **not** be members of a Windows failover cluster.

DNS Configuration

All Active Directory Servers, Moonwalk FPolicy Servers, and NetApp Filers, **must** have both forward **and** reverse records in DNS.

All hostnames used in Filer and FPolicy Server configuration must be FQDNs.

5.3.3 Setup

Setup Parameters

Before starting the installation the following parameters must be considered:

- Management Interface IP Address: the address for management access to the **Vserver** (not to be confused with cluster or node management addresses)
- SMB Privileged User: a domain user for the exclusive use of FPolicy

Preparing Vserver Management Access

For each Vserver, ensure that 'Management Access' is allowed for at least one network interface. Check the network interface in OnCommand System Manager — if Management Access is not enabled, *create a new interface* just for Management Access.

Note: Using the same interface for management and data access may cause NetApp firewall problems. It is recommended to use separate interfaces.

Configuring SMB Privileged Data Access

If it has not already been created, create the SMB Privileged User on the domain. Each FPolicy Server will use the same SMB Privileged User for all Vservers that it will manage.

Open a command line session to the cluster management address:

1. Create a new local group
   - `cifs users-and-groups local-group create -group-name <Name> -vserver <vserver fqdn>`
2. Assign ALL available privileges to the local group
   - `cifs users-and-groups privilege add-privilege -user-or-group-name <Group Name> -privileges SeTcbPrivilege SeBackupPrivilege SeRestorePrivilege SeTakeOwnershipPrivilege SeSecurityPrivilege SeChangeNotifyPrivilege -vserver <vserver fqdn>`
3. Add the CIFS Privileged User to this group
   - `cifs users-and-groups local-group add-members -group-name <Name> -member-names <Domain\User> -vserver <vserver fqdn>`
4. Allow a few minutes for the change to take effect (or FPolicy Server operations may fail with access denied errors)

Installation

On each FPolicy Server machine:

1. Close any SMB sessions open to Vserver(s) before proceeding
2. Add the SMB Privileged User to the *local* Administrators group
3. Ensure the SMB Privileged User has the 'Log on as a service' privilege
4. Run the `Moonwalk NetApp FPolicy Server.exe`
5. Follow the prompts to complete the installation
6. Follow the instructions to activate the installation

Configuring the FPolicy Server

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the FPolicy Server. In the *'Configuration'* panel, select *'NetApp FPolicy'*.

First, click on the *'Configuration'* edit icon to provide credentials for *'Management Authentication'*. The same credentials will be used when accessing each Vserver.

Note that encrypted credentials will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding FPolicy Servers.

Next, add the Vserver(s). Once the SMB data FQDN and Management FQDN (if different) have been provided, the Vserver connection will be validated and the FPolicy configuration will be synced to the Vserver.

5.3.4 Usage

SMB shares that will be used in Moonwalk Policies should ideally be configured to **hide** symbolic links. If a different setting is required for other SMB clients, create a new share at the same location just for Moonwalk traversal that *does* hide links. To modify the symlink behavior on a share:

1. Open a command line session to the cluster management address
2. For each share:
   - `cifs share modify -share-name <sharename> -symlink-properties hide -vserver <vserver fqdn>`

As an alternative to the **hide** setting, -symlink-properties may be set to **disable**. In this mode, symlinks are visible but cannot be opened. To avoid access-denied errors, create Policies using a Rule to exclude files with the *'Hidden'* attribute (the Vserver adds this flag to UNIX symlinks when in this mode).

URI Format

`netapp://{FPolicy Server}/{NetApp Vserver}/{SMB Share}/[{path}/]`

Where:

- `FPolicy Server` – FQDN alias that points to all Moonwalk FPolicy Servers for the given Vserver
- `NetApp Vserver` – FQDN of the Vserver's Data Access interface
- `SMB Share` – NetApp SMB share name

Example:

`netapp://fpol-svrs.example.com/vs1.example.com/data/`

5.3.5 Snapshot Restore

Volume Restore

After an entire volume containing stubs is restored from snapshot, a Post-Restore Revalidate Policy must be run, as per the restore procedure described in §3.4 (p.17).

Individual Stub Restore

Users cannot perform self-service restoration of *stubs*. However, an administrator may restore specific stubs or sets of stubs from snapshots by following the procedure outlined below. Be sure to provide this procedure to all administrators.

1. In AdminCenter Browser, navigate to the relevant share using a `netapp://` URI
2. Click POINT-IN-TIME to select a snapshot
3. Select the desired files, stubs and folders
4. Click DOWNLOAD

For larger jobs, where a Policy is preferable, the point-in-time feature is also available when creating a Source.

5.3.6 Interoperability

NDMP Backup

Stub files may be backed up and restored using NDMP Backup products.

Robocopy

Except when following the procedure in §5.3.5, Robocopy **must not** be used with the `/b` (backup mode) switch when copying Moonwalk NetApp stubs.

When in backup mode, robocopy attempts to copy stub files as-is rather than demigrating them as they are read. This behavior is not supported.

Note: The `/b` switch requires Administrator privilege – it is not available to normal users.

5.3.7 Behavioral Notes

Unix Symbolic Links

Unix Symbolic links (also known as symlinks or softlinks) may be created on a Filer via an NFS mount. Symbolic links will not be seen during Moonwalk Policy traversal of a NetApp file system (since only shares which hide symbolic links are supported for traversal). If it is intended that a policy should apply to files within a folder referred to by a symbolic link, ensure that the Source encompasses the real location at the link's destination. A Source URI may NOT point to a symbolic link – use the real folder that the link points to instead.

Client-initiated demigrations via symbolic links will operate as expected.

QTree and User Quotas

NetApp QTree and user quotas are measured in terms of *logical* file size. Thus, migrating files has no effect on quota usage.

Snapshot Traversal

Moonwalk will automatically skip snapshot directories when traversing shares using the `netapp` scheme.

Dynamic Access Control Considerations

If the Netapp Filer is correctly configured to support Dynamic Access Control (Central Access Policy), you may enable Resource Attribute SACL ACEs and Scoped Policy SACL ACEs in the FPolicy Server Configuration.

These settings cause additional security information to be gathered during migration, which may be useful in certain advanced recovery situations.

Important: Do not enable these settings if the filer is not configured for Dynamic Access Control, since this will cause errors while accessing files.

5.3.8 Troubleshooting

Troubleshooting Management Login

- Open a command line session to the **cluster** management address
- `security login show -vserver <vserver-name>`
  - There should be an entry for the expected user for application *'http'* with role *'vsadmin'*
  - If you are using an ONTAP version prior to 9.11, there should additionally be an entry for application *'ontapi'* with role *'vsadmin'*

Troubleshooting TLS Management Access

- Open a command line session to the **cluster** management address
- `vserver context -vserver <vserver-name>`
- `security certificate show`
  - There should be a *'server'* certificate for the Vserver management FQDN (NOT the bare hostname)
  - If using certificate-based authentication, there should be a *'client-ca'* entry
- `security ssl show`
  - There should be an enabled entry for the Vserver management FQDN (NOT the bare hostname)

Troubleshooting Vserver Configuration

If the FPolicy configuration on a Vserver is out of sync with the Moonwalk FPolicy Server, this may be repaired in AdminCenter using the RE-SYNC icon in the Vservers list.

Troubleshooting 'PRIVILEGED_SHARE_NOT_FOUND' Errors

If the FPolicy Server reports privileged share not found, there is a misconfiguration or SMB issue. Please attempt the following steps:

- Check all configuration using troubleshooting steps described above
- Ensure the FPolicy Server has no other SMB sessions to Vservers
  - run `net use` from Windows Command Prompt
  - remove all mapped drives
- Reboot the server
- Retry the failed operation
  - Check for new errors in `agent.log`

Troubleshooting 'The sequence number is already used by another policy' [13115]

The FPolicy Server configuration wizard may report an FPolicy API error of 'The sequence number is already used by another policy' if another FPolicy application is in use with the same Vserver.

Moonwalk always uses FPolicy Sequence Numbers 1 and 2. It may be possible to reconfigure the other application to use a different sequence number. If this is done, both products' policies can be enabled at the same time, with the Moonwalk policy being given greater precedence. As always, care must be taken when running two FPolicy applications at the same time to ensure that their functionality does not conflict. This consideration is outside the scope of this document.

5.4 Dell EMC PowerScale OneFS

This section describes Moonwalk's capabilities when used with OneFS on Dell EMC PowerScale / Isilon platforms.

5.4.1 Link-Migration Support

OneFS does not provide an interface for performing Moonwalk stub-based migration. As an alternative, Moonwalk provides a link-based migration mechanism via a LinkConnect Server. See §4.9 (p.26) for details of the Link-Migrate operation.

Link-Migration works by pairing a OneFS SMB share with a corresponding LinkConnect Cache Share. Typically a top-level share on each OneFS device is mapped to a unique share (or subdivision) on a LinkConnect Server. Multiple OneFS systems may use shares/subdivisions on the same LinkConnect Server if desired.

Once this configuration is completed, Link-Migrate policies convert files on the source OneFS share to links pointing to the destination files via the LinkConnect Cache Share, according to configured rules.

Link-Migrated files can be identified by the 'O' (Offline) attribute in Explorer. Depending on the version of Windows, files with this flag may be displayed with an overlay icon.

Note: If the Link-Migrate operation will not be used, a LinkConnect Server is not necessary. In this case, refer to §5.19.

5.4.2 Planning

Prerequisites

- An NTFS Cache Volume of at least 1TB – see §2.2.5 (p.13)
- A Moonwalk license that includes an entitlement for LinkConnect Server.

When creating a production deployment plan, please refer to §3.5 (p.19).

Note: It is recommended that a LinkConnect Server is only associated with one type of SMB device. For example, do not associate a single LinkConnect Server with both Windows and OneFS shares. This is because agent configuration options may need to be tuned differently to best work with the different platforms.

NAS System Requirements

- Moonwalk LinkConnect Server requires OneFS version 9.1.0 or higher

Client Requirements

Windows clients require a supported 64-bit Windows operating system:

- Windows 11
- Windows 10
- Windows Server 2022

- Windows Server 2019
- Windows Server 2016

In order to access link-migrated files, the LinkConnect Client Driver must be installed on each client machine – see §2.3 (p.14).

Network

Place the Moonwalk LinkConnect Server on the same subnet and same switch as the corresponding OneFS system to minimize latency.

Additionally, the LinkConnect Server **must** be joined to the same domain as the OneFS NAS.

Antivirus Considerations

Ensure that Windows Defender or any other antivirus product installed on the LinkConnect Server is configured to **omit** scanning/screening on the LinkConnect Cache Volume **and** any OneFS SMB shares.

High-Availability for LinkConnect Server

Consider whether High-Availability (HA) is required in your environment (either now *or in the future*). If so, LinkConnect Servers must be deployed in a DFSN configuration from the outset.

LinkConnect Cache Shares are configured for HA by exposing the share name at the domain level using DFSN. If not using HA, it is possible to use either a simple share on a standalone server, or a share exposed at the domain level using DFSN. The latter is always recommended to allow transition to an HA configuration in the future.

Regular Maintenance Activity

Each configured MigLink source will be periodically scanned to perform maintenance tasks such as MigLink ACL propagation and Link Deletion Monitoring (see below).

In an HA configuration, this scanning activity will be performed by a single caretaker node, as can be seen on the AdminCenter Servers page. A standalone LinkConnect Server always performs the caretaker role.

Security Considerations

Files with certain ACLs cannot be link-migrated – they will be skipped during Link-Migrate Policies. If such an ACL is set on a file that has already been link-migrated, the new ACL will NOT be propagated to the LinkConnect Server. Specifically:

- Conditional ACEs will not be link-migrated (consider using Central Access Policies instead where applicable)
- Audit ACEs which track attempted read access will not be link-migrated (Audit ACEs which track e.g. write or delete work as expected)

Link Deletion Monitoring

Link Deletion Monitoring (LDM) may be enabled on a per-share basis.

Similar to the Stub Deletion Monitoring feature provided by Moonwalk Agents on Windows, LDM identifies secondary storage files that are no longer referenced in order to facilitate recovery of storage space by Scrub Policies. This feature extends not only to MigLinks that are demigrated or directly deleted by the user, but also to other cases such as overwriting a MigLink or renaming a different file over the top of a MigLink.

Unlike SDM, LDM requires a number of maintenance scans to determine that a given secondary storage file is no longer referenced. It should be noted that interrupting the maintenance process (e.g. by restarting the caretaker node or transitioning the caretaker role) will delay the detection of unreferenced secondary storage. For optimal and timely storage space recovery, ensure that LinkConnect Servers can run uninterrupted for extended periods.

Warning: In order to avoid LDM incorrectly identifying files as deleted – leading to unwanted data loss during Scrub – it is critical to ensure that users cannot *move/rename* MigLinks out of the scanned portion of the directory tree within the filesystem.

5.4.3 Setup

Installation

Provision a user on the Active Directory domain for the **exclusive** use of your LinkConnect service(s). This user does *not* need to be a member of Domain Admins.

On each LinkConnect Server machine:

1. Add the LinkConnect user to the *local* Administrators group
2. Assign the 'Log on as a service' privilege to this user
3. Run the `Moonwalk LinkConnect Server.exe`
4. Follow the prompts to complete the installation
5. Follow the instructions to activate the installation

Configuring the LinkConnect Server

In the Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the LinkConnect Server. In the *'Configuration'* panel, select *'LinkConnect Server'*, then use the wizard to add shares.

5.4.4 Adding Shares

Further 'top-level' shares can be added using the *'Configuration'* panel as above.

If using sub-shares defined using path expansion variables (e.g. for dynamic home directories such as [HOME → /ifs/data/home%0/%U] or [%U → /ifs/data/homes/%U]), check the *'Path expansion variables in use'* box in the wizard and specify these share expansions to enable the LinkConnect Server to correctly apply them when MigLinks

are accessed. The following variables are supported in expansion paths: %U, %D, %0, %1, and %2.

For all other sub-shares (that is, shares within the directory tree of a registered top-level share, including shares that are simply aliases for top-level shares), simply add permissions for the LinkConnect user:

1. Open the OneFS Storage Administration web console
2. Navigate to Protocols → Windows Sharing (SMB) → SMB Shares
3. Edit the share
    - Add the LinkConnect user as a new member
    - Specify 'Run as root' permission
    - Move the new member to the top of the members list

5.4.5 Usage

URI format

`smb://{server}/{nas}/{share}/[{path}/]`

Where:

- `server` – FQDN of a LinkConnect Server that is configured to support the OneFS share
- `nas` – OneFS FQDN
- `share` – OneFS SMB share
- `path` – path within the share

Example:

`smb://link.example.com/onefs.example.com/pub/projects/`

5.4.6 Policy Limitations

Link-Migration cannot be configured to skip sparse files on this platform due to its limited sparse file support.

5.4.7 Snapshot Support

MigLinks may be restored from OneFS snapshots, providing that the associated secondary storage file has not yet been Scrubbed. This includes restoring an entire snapshot (SnapRevert) as well as copying an individual MigLink from a snapshot. When copying individual MigLinks, depending on the copy method used, the file will either be restored as a regular file (e.g. Explorer copy) or remain as a MigLink (e.g. `cp -a` command at the OneFS console).

5.4.8 SyncIQ

There is a known issue with SyncIQ not correctly copying symbolic links. This issue occurs even when not using Moonwalk. Due to this limitation, it may not be possible to correctly copy MigLinks using SyncIQ – test your OneFS version prior to attempting this operation.

5.5 DataCore Swarm SCSP

5.5.1 Introduction

DataCore Swarm provides a multi-tenanted object storage platform built upon Swarm storage nodes. Swarm may be used as a migration destination only.

This section details the use of Moonwalk with Swarm using SCSP. Use of Swarm with the S3 protocol is described in §5.13.

SCSP traffic may optionally be encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect migrated data *at rest*.

5.5.2 Planning

Before proceeding with the installation, the following will be required:

- Cloud Gateway 3.0.0 or above
- Swarm 7.1.1 or above
- a license that includes an entitlement for Swarm

Policy Limitations

The following Policy limitations apply to this scheme:

- it may not be used as a Link-Migration destination
- it may not be used as the *new* destination for Change Destination Tier policies
- it may not be used as the *new* destination for Retarget Destination policies

Firewall

The TCP port used to access the Swarm Content Gateway via HTTP or HTTPS must be allowed by any firewalls between the Moonwalk Gateway Agent and the Swarm endpoint. For further information regarding firewall configuration see Appendix B.

Named and Unnamed Objects

Migrated files may be stored as either unnamed objects (accessed by UUID), or as named objects residing in a bucket. Bucket creation must be performed ahead of time, prior to configuring Moonwalk.

5.5.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as Domain credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways.

Swarm Destinations require an index to be created prior to use. Once credentials have been supplied, click **Create new index** to create a new index and corresponding migration Destination.

Additional indexes can be added at a later date to further subdivide storage if required. Multiple migration destinations may be created in the same bucket by specifying different partition names.

Important: If multiple Moonwalk deployments are in use migrating to the same Swarm cluster, different indexes are required for EACH AdminCenter.

Metadata Options

Enable *'Include metadata headers'* to store per-file HTTP metadata with the destination objects, such as original filename and location, content-type, owner and timestamps – see §5.5.6 for details. Swarm 8 or above is required to use this option.

Also enable *'Include Content-Disposition'* to include original filename for use when downloading the target objects directly using a web browser.

5.5.4 Legacy URIs

URIs created on previous versions of Moonwalk using the `cloudscaler` scheme will continue to function as expected. Existing destinations should NOT be updated to use the `scsp` scheme. The `cloudscaler` scheme is simply an alias for the `scsp` scheme.

5.5.5 Disaster Recovery Considerations

During migration, each newly migrated file is recorded in the corresponding index. The index may be used in disaster scenarios where:

1. stubs have been lost, and
2. a *Create Recovery File from Source* file is not available, and
3. no current backup of the stubs exists

Index performance is optimized for migrations and demigrations, not for *Create Recovery File from Destination* policies.

Create Recovery File from Source policies are the recommended means to obtain a Recovery file for restoring stubs. This method provides better performance and the most up-to-date stub location information.

It is recommended to regularly run *Create Recovery File from Source* policies following *Migration* policies.

5.5.6 Swarm Metadata Headers

The following metadata fields are supported:

- `X-Alt-Meta-Name` – the original source file's filename (excluding directory path)
- `X-Alt-Meta-Path` – the original source file's directory path (excluding the filename) in a platform-independent manner such that '/' is used as the path separator and the path will start with '/', followed by drive/volume/share if appropriate, but not end with '/' (unless this path represents the root directory)
- `X-Moonwalk-Meta-Partition` – the Destination URI *partition* – if no partition is present, this header is omitted
- `X-Source-Meta-Host` – the FQDN of the original source file's server
- `X-Source-Meta-Owner` – the owner of the original source file in a format appropriate to the source system (e.g. DOMAIN\username)
- `X-Source-Meta-Modified` – the *Last Modified* timestamp of the original source file at the time of migration in RFC3339 format
- `X-Source-Meta-Created` – the *Created* timestamp of the original source file in RFC3339 format
- `X-Source-Meta-Attribs` – a case-sensitive sequence of characters {AHRS} representing the original source file's file flags: *Archive*, *Hidden*, *Read-Only* and *System*
  - all other characters are reserved for future use and should be ignored
- `Content-Type` – the MIME Type of the content, determined based on the file-extension of the original source filename

Note: Timestamps may be omitted if the source file timestamps are not set.

Non-ASCII characters will be be stored using RFC2047 encoding, as described in the Swarm documentation. Swarm will decode these values prior to indexing in Elasticsearch.

5.6 DataCore Swarm (Direct Node Access)

5.6.1 Introduction

The `scspdirect` scheme should only be used when accessing Swarm storage nodes *directly*. Swarm may be used as a migration destination only.

Swarm (SCSP) traffic is *not* encrypted in transit when using this scheme. Optionally, the plugin can employ client-side encryption to protect migrated data *at rest*.

Normally, Swarm will be accessed via a Swarm Content Gateway, in which case the `scsp` scheme must be used instead, see §5.5.

5.6.2 Planning

Before proceeding with the installation, the following will be required:

- Swarm 7.1.1 or above (or CAStor 6.0 or above)
- a license that includes an entitlement for Swarm

Policy Limitations

The following Policy limitations apply to this scheme:

- it may not be used as a Link-Migration destination
- it may not be used as the *new* destination for Change Destination Tier policies
- it may not be used as the *new* destination for Retarget Destination policies

Firewall

The Swarm storage node port must be allowed by any firewalls between the Moonwalk Gateway Agent and the Swarm storage nodes. For further information regarding firewall configuration see Appendix B.

Domains and Endpoints

Swarm storage locations are accessed via a configured endpoint FQDN. Add several Swarm storage node IP addresses to DNS under a single endpoint FQDN (4-8 addresses are recommended). If domains are NOT in use (i.e. data will be stored in the default cluster domain), it is **strongly** recommended that the FQDN be the name of the cluster for best Swarm performance.

Note: In a legacy installation where domains were not previously used, DO NOT create a Swarm domain which matches the FQDN used in existing (or previous) Moonwalk destinations. Such a domain may prevent proper access to the untenanted data already stored in the default cluster domain.

Named and Unnamed Objects

Migrated files may be stored as either unnamed objects (accessed by UUID), or as named objects residing in a bucket. Bucket creation must be performed ahead of time, prior to configuring Moonwalk.

5.6.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options and encryption settings. Note that encrypted keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways.

Swarm Destinations require an index to be created prior to use: click **Create new index** to create a new index and corresponding migration Destination.

Additional indexes can be added at a later date to further subdivide storage if required. Multiple migration destinations may be created in the same bucket by specifying different partition names.

Important: If multiple Moonwalk deployments are in use migrating to the same Swarm cluster, different indexes are required for EACH AdminCenter.

Metadata Options

Enable *'Include metadata headers'* to store per-file HTTP metadata with the destination objects, such as original filename and location, content-type, owner and timestamps – see §5.5.6 for details. Swarm 8 or above is required to use this option.

Also enable *'Include Content-Disposition'* to include original filename for use when downloading the target objects directly using a web browser.

5.6.4 Legacy URIs

URIs created on previous versions of Moonwalk using the `castor` or `swarm` schemes will continue to function as expected. Existing destinations should NOT be updated to use the `scspdirect` scheme. The `castor` and `swarm` schemes are simply aliases for the `scspdirect` scheme.

5.6.5 Disaster Recovery Considerations

Refer to §5.5.5.

5.7 Hitachi Content Platform (HCP)

5.7.1 Introduction

The Hitachi Content Platform may be used as a migration destination only for Moonwalk. Moonwalk accesses HCP clusters using Authenticated Namespaces (ANS) via HTTPS.

5.7.2 Planning

Before proceeding with the installation, the following will be required:

- HCP 7.2 or above
- The HCP system must have at least one namespace configured for use with Moonwalk:
  - HTTPS must be enabled
  - Versioning should be disabled
  - If using retention, allow metadata 'Add, delete and replace'
- An HCP local user with at least [Browse, Read, Write, Delete, Purge] permissions for the namespace
- A license that includes an entitlement for HCP

Firewall

The HTTPS port (TCP port 443) must be allowed by any firewalls between the Moonwalk Gateway Agent and the HCP cluster.

DR Site Replication

For assistance in planning for DR Site Replication, including replicated clusters and Gateways, please contact Moonwalk Support.

5.7.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify proxy options and supply HCP namespace credentials. Once credentials have been supplied, click on the CREATE MIGRATION DESTINATION icon.

Note that encrypted credentials will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways.

5.7.4 Behavioral Notes

Retention and Scrub

When running Moonwalk Scrub Policies, files currently under retention will be automatically skipped.

5.8 Amazon S3

5.8.1 Introduction

Amazon S3 may be used as either a migration or ingest destination.

Additionally, Amazon S3 may be used as a source for copy, move and ingest policies. When copying a dataset that did not originate on a filesystem however, it is important to check that the object naming convention will map suitably to the destination filesystem.

S3 traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

This section strictly pertains to *Amazon* S3. Other supported S3-compatible storage services/devices are documented in separate sections.

5.8.2 Planning

Before proceeding with the installation, the following will be required:

- an Amazon Web Services (AWS) Account
- a license that includes an entitlement for Amazon S3

Dedicated buckets – without versioning enabled – should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The HTTPS port (TCP port 443) must be allowed by any firewalls between the Moonwalk Gateway Agent and the Internet.

5.8.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

S3 source and destination URIs may then be created interactively within the Source and Destination editors respectively.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Transfer Acceleration

Transfer acceleration allows data to be uploaded via the fastest data center for your location, regardless of the actual location of the bucket.

This per-bucket option provides a way to upload data to a bucket in a remote AWS region while minimizing the adverse effects on migration policies that would otherwise be caused by the correspondingly higher latency of using the remote region.

Additional AWS charges may apply for using transfer acceleration at upload time, but for archived data these initial charges may be significantly outweighed by reduced storage costs in the target region. For further details, please consult AWS pricing.

Default Storage Class

This per-bucket option allows files to be uploaded *directly* into the STANDARD, STANDARD_IA or the GLACIER_IR storage class subject to eligibility. Unconfigured buckets use the STANDARD storage class.

Please consult AWS pricing for further details.

Ingesting to other Storage Classes

Ingest Policies may explicitly specify that data is to be ingested into a given storage class. The following table shows supported Storage Class Names.

Storage Class Name	**Amazon Description**
STANDARD	Standard (the default storage class)
STANDARD_IA	Standard-IA (Infrequent Access)
ONEZONE_IA	One Zone-IA
INTELLIGENT_TIERING	Intelligent-Tiering
GLACIER_IR	Glacier Instant Retrieval
GLACIER	Glacier Flexible Retrieval
DEEP_ARCHIVE	Glacier Deep Archive
REDUCED_REDUNDANCY	RRS (not recommended)

Consult Amazon documentation for more information and guidance in choosing the correct storage class for your needs.

Accessing Archived Objects

Objects in the GLACIER or DEEP_ARCHIVE storage classes (or archived by INTELLIGENT_TIERING) must be restored prior to being accessed. To restore such objects using a Moonwalk Policy, see §4.13 (p.28).

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such

as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under <bucket>/<partition>/FILES. Moonwalk-specific metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.8.4 Extended Metadata Fields

If enabled in an Ingest Policy, metadata is stored as described below. For more information about the Ingest operation, see §4.12 (p.28).

Extended metadata fields are also written when the *'Migrate with original filenames'* option is selected for a migration destination bucket.

Header Field	**Content**
x-amz-meta-orig-host	Source server FQDN
x-amz-meta-orig-name	Original filename (without path)
x-amz-meta-orig-modified-time	Modified timestamp
x-amz-meta-orig-created-time	Creation timestamp
x-amz-meta-orig-attribs	Subset of characters {AHRS} representing the original source file's flags
Content-Disposition (optional)	Original name for web browser download
Security Details	*as appropriate*
x-amz-meta-orig-owner	File owner – e.g. Domain\JoeUser
x-amz-meta-orig-sddl	Microsoft SDDL format security descriptor
x-amz-meta-orig-uid	Unix user ID
x-amz-meta-orig-gid	Unix group ID
x-amz-meta-orig-unix-perms	Octal permissions e.g. 00644

Notes:

- headers will be sent in UTF-8 using RFC2047 encoding as necessary to unambiguously represent the original metadata values (in accordance with the HTTP/1.1 specification – see RFC2616/2.2)
- due to Amazon-specific limitations, sequences of adjacent whitespace within x-amz-meta-orig-name may be returned as a single space by some client software
- all timestamps are stored as UTC in RFC3339 format

Custom Metadata

In addition to the fields above, Ingest policies can optionally specify an additional metadata field to be attached to each upload in Header: value format.

Permitted fields include: x-amz-meta-<field> (that is, user metadata), x-amz-grant-<permission>, x-amz-acl, x-amz-tagging and x-amz-storage-class. It is the user's responsibility to ensure that these headers are used in a way that is consistent with the storage service's documentation.

5.9 Cloudian HyperStore

5.9.1 Introduction

Cloudian HyperStore may be used as either a migration or ingest destination and is accessed via the S3 protocol.

S3 traffic may optionally be encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.9.2 Planning

Before proceeding with the installation, the following will be required:

- suitable S3 API credentials
- a license that includes an entitlement for Cloudian HyperStore

Dedicated buckets should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the storage endpoint.

5.9.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

When configuration is complete, click the CREATE MIGRATION DESTINATION icon next to the desired bucket.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under <bucket>/<partition>/FILES. Moonwalk-specific metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.9.4 Compatibility and Limitations

For HyperStore installations that feature an external HTTP proxy load-balancer in front of the storage nodes, ensure that the load-balancer is fully HTTP/1.1 compliant. In particular, Moonwalk requires correct support for HTTP 'Expect: 100-continue' headers.

Moonwalk does not support the following operations for HyperStore destinations:

- Scrub
- Create Recovery File From Destination

Note: The 'Create Recovery File From *Source*' operation is still supported.

5.9.5 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.10 Dell EMC Elastic Cloud Storage

5.10.1 Introduction

Dell EMC Elastic Cloud Storage (ECS) may be used as either a migration or ingest destination and is accessed via the S3 protocol.

S3 traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.10.2 Planning

Before proceeding with the installation, the following will be required:

- suitable S3 API credentials
- a license that includes an entitlement for Dell EMC ECS

Dedicated buckets should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the storage endpoint.

5.10.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

When configuration is complete, click the CREATE MIGRATION DESTINATION icon next to the desired bucket.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

UTF-16 listing order work-around

ECS has been observed to return object information ('file listings') in UTF-16 code-unit order rather than the Amazon-compatible Unicode code-point order. A work-around is enabled by default to process results in this non-standard order.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under <bucket>/<partition>/FILES. Moonwalk-specific metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.10.4 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.11 IBM Cloud Object Storage

5.11.1 Introduction

IBM Cloud Object Storage (COS) may be used as either a migration or ingest destination and is accessed via the S3 protocol.

Additionally, IBM COS may be used as a source for copy, move and ingest policies. When copying a dataset that did not originate on a filesystem however, it is important to check that the object naming convention will map suitably to the destination filesystem.

S3 traffic may optionally be encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.11.2 Planning

Before proceeding with the installation, the following will be required:

- suitable S3 API credentials
- a license that includes an entitlement for IBM COS

Dedicated buckets should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the storage endpoint.

5.11.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

S3 source and destination URIs may then be created interactively within the Source and Destination editors respectively.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Virtual Host Access

IBM Cloud Object Storage supports the virtual-host-style bucket access method as expected for the S3 protocol. For example https://bucket.cos.example.com rather than https://cos.example.com/bucket.

Generally, the *'Use Virtual Host Access'* option should be enabled (the default).

Note: When using Virtual Host Access in conjunction with HTTPS (recommended) it is important to ensure that the endpoint's TLS certificate has been created correctly. For example, if the endpoint FQDN is `cos.example.com`, the certificate must contain Subject Alternative Names (SANs) for both `cos.example.com` **and** `*.cos.example.com`.

UTF-16 listing order work-around

IBM COS has been observed to return object information ('file listings') in UTF-16 code-unit order rather than the Amazon-compatible Unicode code-point order. A work-around is enabled by default to process results in this non-standard order.

Legacy URIs

Older versions of Moonwalk provided IBM COS support via the `s3bluemix://` URI scheme. Sources and Destinations using these URIs will continue to work after upgrade and should NOT be updated to use the `s3cos://` scheme. New Sources and Destinations should use the new scheme.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under `<bucket>/<partition>/FILES`. Moonwalk-specific metadata is stored under `<bucket>/<partition>/HDR` and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.11.4 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.12 Wasabi Object Storage

5.12.1 Introduction

Wasabi Object Storage may be used as either a migration or ingest destination and is accessed via the S3 protocol.

S3 traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.12.2 Planning

Before proceeding with the installation, the following will be required:

- suitable S3 API credentials
- a license that includes an entitlement for Wasabi Object Storage

Dedicated buckets should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the storage endpoint.

5.12.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

When configuration is complete, click the CREATE MIGRATION DESTINATION icon next to the desired bucket.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under <bucket>/<partition>/FILES. Moonwalk-specific metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.12.4 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.13 DataCore Swarm S3

5.13.1 Introduction

DataCore Swarm provides a multi-tenanted object storage platform built upon Swarm storage nodes. Swarm S3 may be used as either a migration or ingest destination and is accessed via the S3 protocol.

This section details the use of Moonwalk with Swarm using the S3 protocol. Use of Swarm with SCSP is described in §5.5.

S3 traffic may optionally be encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.13.2 Planning

Before proceeding with the installation, the following will be required:

- suitable S3 API credentials
- a license that includes an entitlement for Swarm

Dedicated buckets should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the Swarm endpoint.

5.13.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

When configuration is complete, click the CREATE MIGRATION DESTINATION icon next to the desired bucket.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under <bucket>/<partition>/FILES. Moonwalk-specific metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.13.4 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.14 Generic S3 Endpoint

5.14.1 Introduction

Other generic or third-party storage devices and services that support the Amazon S3 protocol may be addressed using the 'Generic S3 Endpoint' feature.

Such endpoints may be used as either migration or ingest destinations.

Additionally, the endpoints may also be used as sources for copy, move and ingest policies. When copying a dataset that did not originate on a filesystem however, it is important to check that the object naming convention will map suitably to the destination filesystem.

S3 traffic may optionally be encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.14.2 Planning

Important: Prior to production deployment, please confirm with Moonwalk Universal that the chosen device or service has been certified for compatibility to ensure that it will be covered by your support agreement.

Prerequisites:

- suitable S3 API credentials
- a license that includes an entitlement for generic S3 endpoints

Dedicated buckets – without versioning enabled – should be used for Moonwalk migration data. However, do not create any S3 buckets at this stage.

Firewall

The S3 port must be allowed by any firewalls between the Moonwalk Gateway Agent and the storage endpoint.

5.14.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as S3 account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create buckets and edit bucket-specific settings.

S3 source and destination URIs may then be created interactively within the Source and Destination editors respectively.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

Omit ISO date from path

Normally, when Moonwalk migrates a file to S3, a timestamp is included in each resulting S3 object key (name). *Amazon* S3 implements a flat, uniform keyspace – there is no concept of a directory structure within an Amazon storage bucket. However, some S3-compatible devices map the keyspace to an underlying directory structure or other non-uniform or hierarchical namespace. On such systems, the inclusion of the timestamp may result in excessive directory creation which may adversely impact performance and/or resource consumption. For such devices, use the *'Omit ISO date from path'* option to omit the timestamp.

Virtual Host Access

The S3 protocol supports a virtual-host-style bucket access method, for example https://bucket.s3.example.com rather than only https://s3.example.com/bucket. This facilitates connecting to a node in the correct region for the bucket, rather than requiring a redirect.

Generally the *'Use Virtual Host Access'* option should be enabled (the default) to ensure optimal performance and correct operation. However, if the generic S3 endpoint in question does not support this feature at all, Virtual Host Access may be disabled.

Note: When using Virtual Host Access in conjunction with HTTPS (recommended) it is important to ensure that the endpoint's TLS certificate has been created correctly. For example, if the endpoint FQDN is `s3.example.com`, the certificate must contain Subject Alternative Names (SANs) for both `s3.example.com` **and** `*.s3.example.com`.

UTF-16 listing order work-around

Some third-party implementations of Amazon's S3 protocol return object information ('file listings') in UTF-16 code-unit order rather than the Amazon-compatible Unicode code-point order. This option allows Moonwalk to correctly process results returned in this non-standard order and thereby allows correct and complete scanning of your S3 buckets. If your object names use non-ASCII characters, it is strongly recommended that you check that this setting is correct for your device prior to use.

Migration Layout

By default, migrated data is stored in Standard migration layout within the object store. Standard layout supports encryption *at rest*.

Alternatively, migrated data may be stored in a manner that preserves original filename information. This layout does **not** support encryption, and is subject to limitations such as path/filename length imposed by the object store. This option is useful in specific circumstances where data at a migration destination must be read *directly* by other applications. Files are stored under `<bucket>/<partition>/FILES`. Moonwalk-specific

metadata is stored under <bucket>/<partition>/HDR and should not be made accessible to other applications.

Note: Buckets configured to preserve original filename information upon migration may not be used as the *new* Destination for Change Destination Tier or Retarget Destination Policies.

5.14.4 Extended Metadata Fields

Please refer to §5.8.4 for S3 metadata field details.

5.15 Microsoft Azure Storage

5.15.1 Introduction

Microsoft Azure may be used as either a migration or ingest destination.

Additionally, Microsoft Azure may be used as a source for copy, move and ingest policies. When copying a dataset that did not originate on a filesystem however, it is important to check that the object naming convention will map suitably to the destination filesystem.

Azure traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.15.2 Planning

Before proceeding with the installation, the following will be required:

- a Microsoft Azure Account
- a Storage Account within Azure – both General Purpose and Blob Storage (with Hot and Cool access tiers) account types are supported
- a Moonwalk license that includes an entitlement for Microsoft Azure

Note: Azure Data Lake Storage (ADLS) Gen 2 is supported as an Ingest Destination and Copy Source but not as a Migration Destination.

Firewall

The HTTPS port (TCP port 443) must be allowed by any firewalls between the Moonwalk Gateway Agent and the Internet.

5.15.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as Azure Storage Accounts. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE CONTAINERS icon to create and view containers.

Azure source and destination URIs may then be created interactively within the Source and Destination editors respectively.

Advanced Encryption Options

The *'Allow unencrypted filenames'* option greatly increases performance when creating Recovery files from an Azure Destination. This is facilitated by recording stub filenames in Azure metadata in unencrypted form, even when encryption at rest is enabled.

Ingesting to a Specific Access Tier

Ingest Policies may explicitly specify that data is to be ingested into a given access tier by providing the tier name ('Hot', 'Cool' or 'Archive') in the *'Storage class'* field.

Consult Azure documentation for more information and guidance in choosing the correct access tier for your needs.

Accessing Archived Objects

Objects in the 'Archive' tier must be transitioned to another tier prior to being accessed. To set the access tier of objects using a Moonwalk Policy, see §4.14 (p.29).

Azure Access Tracking

If Access Tracking has been enabled on a Storage Account, Accessed time will be available for use in Rules and Reports. Remember however, that Azure's Access Tracking feature updates this information daily, not immediately after each access.

5.15.4 Extended Metadata Fields

If enabled in an Ingest Policy, metadata is stored as described below. For more information about the Ingest operation, see §4.12 (p.28).

Header Field	**Content**
`x-ms-meta-originalhost`	Source server FQDN
`x-ms-meta-originalname`	Original filename (without path)
`x-ms-meta-originalmodifiedtime`	Modified timestamp
`x-ms-meta-originalcreatedtime`	Creation timestamp
`x-ms-meta-originalattribs`	Subset of characters {AHRS} representing the original source file's flags
`Content-Disposition` (optional)	Original name for web browser download
Security Details	*as appropriate*
`x-ms-meta-originalowner`	File owner – e.g. Domain\JoeUser
`x-ms-meta-originalsddl`	Microsoft SDDL format security descriptor
`x-ms-meta-originaluid`	Unix user ID
`x-ms-meta-originalgid`	Unix group ID
`x-ms-meta-originalunixperms`	Octal permissions e.g. 00644

Notes:

- headers will be sent in UTF-8 using RFC2047 encoding as necessary to unambiguously represent the original metadata values (in accordance with the HTTP/1.1 specification – see RFC2616/2.2)

- all timestamps are stored as UTC in RFC3339 format

Custom Metadata

In addition to the fields above, Ingest policies can optionally specify an additional metadata field to be attached to each upload in `Header: value` format.

Permitted fields are limited to `x-ms-meta-<field>` and `x-ms-tags`. Tags may be specified as:

```
x-ms-tags: name=value
```

or, to set multiple tags at the same time:

```
x-ms-tags: name1=value1&name2=value2
```

5.16 Google Cloud Storage

5.16.1 Introduction

Google Cloud Storage may be used as either a migration or ingest destination.

Additionally, Google Cloud Storage may be used as a source for copy, move and ingest policies. When copying a dataset that did not originate on a filesystem however, it is important to check that the object naming convention will map suitably to the destination filesystem.

Google Cloud Storage traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect *migrated* data *at rest*.

5.16.2 Planning

Before proceeding with the installation, the following will be required:

- a Google Account
- a Moonwalk license that includes an entitlement for Google Cloud Storage

Firewall

The HTTPS port (TCP port 443) must be allowed by any firewalls between the Moonwalk Gateway Agent and the Internet.

5.16.3 Storage Bucket Preparation

Using the Google Cloud Platform web console, create a new Service Account in the desired project for the **exclusive** use of Moonwalk. Create a P12 format private key for this Service Account. Record the Service Account ID and store the downloaded private key file securely for use in later steps.

Create a Storage Bucket **exclusively** for Moonwalk data.

For Moonwalk use, bucket names must:

- be 3-40 characters long
- contain **only** lowercase letters, numbers and dashes (-)
- not begin or end with a dash
- not contain adjacent dashes

Edit the bucket's permissions to add the new Service Account as a member with the 'Storage Object Admin' role.

5.16.4 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as Google Storage Accounts. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to register previously created buckets.

Google source and destination URIs may then be created interactively within the Source and Destination editors respectively.

5.16.5 Extended Metadata Fields

If enabled in an Ingest Policy, metadata is stored as described below. For more information about the Ingest operation, see §4.12 (p.28).

Header Field	**Content**
`x-goog-meta-orig-host`	Source server FQDN
`x-goog-meta-orig-name`	Original filename (without path)
`x-goog-meta-orig-modified-time`	Modified timestamp
`x-goog-meta-orig-created-time`	Creation timestamp
`x-goog-meta-orig-attribs`	Subset of characters {AHRS} representing the original source file's flags
`Content-Disposition` (optional)	Original name for web browser download
Security Details	*as appropriate*
`x-goog-meta-orig-owner`	File owner – e.g. Domain\JoeUser
`x-goog-meta-orig-sddl`	Microsoft SDDL format security descriptor
`x-goog-meta-orig-uid`	Unix user ID
`x-goog-meta-orig-gid`	Unix group ID
`x-goog-meta-orig-unix-perms`	Octal permissions e.g. 00644

Notes:

- headers will be sent in UTF-8 using RFC2047 encoding as necessary to unambiguously represent the original metadata values (in accordance with the HTTP/1.1 specification – see RFC2616/2.2)
- all timestamps are stored as UTC in RFC3339 format

Custom Metadata

In addition to the fields above, Ingest policies can optionally specify an additional metadata field to be attached to each upload in `Header: value` format.

Permitted fields are limited to those of the form `x-goog-meta-<field>`.

5.17 Alibaba Cloud Object Storage Service (OSS)

5.17.1 Introduction

Alibaba Cloud OSS is used as a migration destination with Moonwalk. Alibaba Cloud is also known as Aliyun.

Aliyun traffic is encrypted *in transit* with TLS. Additionally, the plugin can employ client-side encryption to protect migrated data *at rest*.

5.17.2 Planning

Before proceeding with the installation, the following will be required:

- an Alibaba Cloud account
- a license that includes an entitlement for Alibaba Cloud OSS

Dedicated buckets should be used for Moonwalk migration data. However, do not create any buckets at this stage.

Firewall

The HTTPS port (TCP port 443) must be allowed by any firewalls between the Moonwalk Gateway Agent and the Internet.

5.17.3 Usage

In Moonwalk AdminCenter, navigate to the *'Servers'* page and configure the Server on which the plugin will be enabled. In the *'Configuration'* panel, select the plugin from the *'Enabled Plugins'* or *'Available Plugins'* list as appropriate.

Configure the plugin to specify options such as proxy and encryption, as well as Aliyun account credentials. Note that encrypted credentials and keys will be stored in secured areas on the AdminCenter server (where they will be backed up) and on the corresponding Gateways. Once credentials have been supplied, click on the MANAGE BUCKETS icon to create and view buckets.

When configuration is complete, click the CREATE MIGRATION DESTINATION icon next to the desired bucket.

Partitions may be used to subdivide a bucket into multiple migration destinations. A greater number of smaller migration destinations may be helpful in a recovery scenario where destinations can be recovered in order of priority.

5.18 Built-in NFS Client

5.18.1 Introduction

Moonwalk Agents support NFS version 3 using a built-in NFS client. Both TCP and UDP connections are supported, but TCP is preferred by default.

Files cannot be migrated from NFS sources.

5.18.2 Planning

Requirements:

- A license that includes an entitlement for NFS

NFS servers must be configured to share file systems to the servers running the Moonwalk Agent. Typically, NFS servers only allow connections from servers that have been given permission by hostname.

The NFS client accesses NFS servers using a UID of 0 (root) and GID of 1. It may be necessary to configure root squashing behavior accordingly, to allow UID 0 to access the files/folders of interest. See also Appendix F.

An NFS share point at '/' is not supported.

WORM Devices

Some WORM storage devices present an NFS interface. If using such a device, be sure to set the WORM behavior flag on the Destination in AdminCenter. This will ensure that the agent expects the storage to exhibit this behavior.

NetApp NFS Shares

NetApp filers may not provide an export list to NFS clients via the usual mount protocol. To work around this limitation, NetApp NFS shares created for Moonwalk use must be created to export a top-level directory, e.g. `/data`.

5.18.3 Setup

NFS file transfer does not require the installation of an additional Moonwalk Gateway Agent.

However, an NFS Browser agent should be installed on the Admin Tools machine to allow browsing of the file systems in the AdminCenter interface and Destination-based policies. Refer to §2.1.1 (p.11) for installation instructions.

5.18.4 Behavioral Notes

Symbolic Links

Symbolic links (also known as symlinks or softlinks) will be skipped during traversal of an NFS file system. This ensures that files are not seen – and thus acted upon – multiple times during a single execution of a given policy. If it is intended that a policy should apply to files within a directory referred to by a symbolic link, either ensure that the Source encompasses the real location at the link's destination, or specify the link itself as the Source.

5.19 SMB Protocol Gateway

5.19.1 Introduction

The `smb` scheme allows access to devices using the SMB protocol. Support for the `smb` scheme is provided by Windows Gateway Agents. No plugins are required.

The Migrate operation is not supported for SMB sources. However, support for the Link-Migrate operation with both Dell EMC OneFS and Windows file servers via a Moonwalk LinkConnect Server is detailed in §5.4 and §5.2 respectively.

5.19.2 Planning

Requirements:

- A **Windows** Moonwalk Gateway Agent (see §2.2.2 (p.12)), optionally configured for High-Availability
- A license that includes support for the `smb` scheme

5.19.3 Setup

Ensure the Moonwalk Gateway Agent is run with an appropriate user account that has full access to the NAS device. To set the account to be used by the Moonwalk Agent service:

1. Open Services → Moonwalk Agent
2. Stop the service
3. Right-click and select Properties → Log On tab
4. Check *'This Account'* and enter account name and password
    - If the chosen account is NOT a local Administrator, it **must** be added to the Administrators group before continuing
5. Start the service

5.19.4 Usage

URI Format

`smb://{gateway}/{host}/{share}/[{path}/]`

Where:

- `gateway` – FQDN of Gateway Agent or LinkConnect Server
- `host` – SMB host server
- `share` – SMB share
- `path` – file system path, such as volume and folders

Note: When accessing an SMB share on a device that is configured for Link-Migration, the FQDN of the corresponding LinkConnect Server must be provided. For shares on all other devices, provide the FQDN of any Windows Moonwalk Gateway Agent.

Legacy URIs

Older versions of Moonwalk provided SMB support via the `cifsnas://` URI scheme. Sources and Destinations using these URIs will continue to work after upgrade and should NOT be updated to use the `smb://` scheme. New Sources and Destinations should use the new scheme.

Snapshot Support

Where supported by the SMB host, snapshots may be selected in the Browser or Source editor. Policies run against Windows snapshot Sources will not traverse mount points out of the snapshot.

Chapter 6

Disaster Recovery

6.1 Introduction

The DrTool application allows for the recovery of files where normal backup and restore procedures have failed. Storage backup recommendations and considerations are covered in §3.4 (p.17).

It is recommended to regularly run a *'Create Recovery File From Source'* Policy to generate an up-to-date list of source–destination mappings.

DrTool is installed as part of Moonwalk Admin Tools.

Note: *Starter Edition* licenses do not include DrTool functionality.

6.2 Recovery Files

Recovery files are normally generated by running *'Create Recovery File From Source'* Policies in AdminCenter. To open a file previously generated by AdminCenter:

1. Open Moonwalk DrTool from the Start Menu
2. Go to File → Open From AdminCenter... → Recovery File From Source
3. Select a Recovery file to open

Older versions of Recovery files may be found via the *'Recovery'* page in AdminCenter.

6.3 Filtering Results

In DrTool, click **Filter** to filter results by source file properties. Filter options are described below.

Note: When a Filter is applied, **Save** only saves the filtered results.

Scheme Pattern

In the *'Scheme Pattern'* field, use the name of the Scheme only (e.g. `win`, not `win://` or `win://servername`). This field may be left blank to return results for all schemes.

This field matches against the scheme section of a URI:

- {**scheme**}://{servername}/[{path}]

Server Pattern

In the *'Server Pattern'* field, use the full server name or a wildcard expression.

This field matches against the `servername` section of a URI:

- {scheme}://{**servername**}/[{path}]

Examples:

- `server65.example.com` – will match only the specified server
- `*.finance.example.com` – will match all servers in the 'finance' subdomain

File Pattern

The *'File Pattern'* field will match either filenames only (and search within all directories), or filenames qualified with directory paths in the same manner as filename patterns in AdminCenter Rules – see Appendix A.

For the purposes of file pattern matching, the top-level directory is considered to be the top level of the entire URI path. This may be different to the top-level of the original Source URI.

Using the Analyze Button

Analyze assists in creating simple filters.

1. Click **Analyze**
   - Analyze will display a breakdown by scheme, server and file type
2. Select a subset of the results by making a selection in each column
3. Click **Filter** to create a filter based on the selection

6.4 Recovering Files

Selected Files

To recover files interactively:

- Select the results for which files will be recovered
- Click Edit → Recover File...

All Files

All files may be recovered either as a batch process using the command line (see §6.7) or interactively as follows:

- Click Edit → Recover All Files...

Missing folders will be recreated as required to contain the recovered files.

Recreated files and folders will not have ACLs applied to them so care should be taken when recreating in sensitive areas.

6.5 Recovering Files to a New Location

When recovering to a new location, always use an up-to-date Recovery file generated by a *'Create Recovery File From Source'* Policy.

To rewrite source file URIs to the new location, use the `-csu` command line option to update the prefix of each URI. Once these URI substitutions have been applied (and checked in the GUI) files may be recovered as previously outlined. The `-csu` option is further detailed in §6.7.

Important: DO NOT create stubs in a new location and then continue to use the old location. To avoid incorrect reference counts, only one set of stubs should exist at any given time.

6.6 Updating Sources to Reflect Destination URI Change

Generally, a Retarget Destination Policy – see §4.11 (p.27) – is the most effective way to permanently move migrated data from one destination to another. If, however, the destination URI changes for some other reason, such as an FQDN being updated externally, DrTool may be used to repair the linkage between the source and the destination.

In DrTool, source files may be updated to reflect a destination URI change through use of the `-cmu` command line option – detailed in §6.7.

To apply the destination URI substitution to *existing* files on the source, select *'Update All Source Files...'* from the Edit menu. When given the option, elect to update substituted entries only.

Note: This operation must always be performed using an up-to-date Recovery file generated by a *'Create Recovery File From Source'* Policy.

6.7 Using DrTool from the Command Line

Important: DO NOT create stubs in a new location and then continue to use the old location. To avoid incorrect reference counts, only one set of stubs should exist at any given time.

Use an **Administrator** command prompt. By default `DrTool` is located in:

`C:\Program Files\Moonwalk\AdminTools\drtool\`

Interactive Usage

`DrTool [Recovery file] [extra options]`

Opens the DrTool in interactive (GUI) mode with the desired options and optionally opens a Recovery file.

Batch Usage

`DrTool [<operation> <Recovery file>] [<options>]`

Run the DrTool without a GUI to perform a batch operation on all entries in the input file.

Note: The Recovery file provided as input is usually created by saving (possibly filtered) results to the hard disk from the interactive DrTool GUI.

Basic Command Line Options

- `operation` – is either:
  - `-recoverFiles`
  - `-updateSource`
    - if combined with `-cmu`, only matching entries will be updated
  - `-updateSourceAll`
    - all entries will be updated, even when `-cmu` is specified
  - if `operation` is omitted, the GUI will be opened with any supplied options
- `Recovery file` – the file to open
- `options` (related to the operation are):
  - `-csu {from} {to}` – to change Source URI prefix, this option can be specified multiple times
  - `-cmu {from} {to}` – to change Migrated URI prefix, this option can be specified multiple times

Examples

All the following examples are run from the DrTool directory.

- `DrTool -recoverFiles result.txt` – recover all files from the `result.txt` file
- `DrTool -updateSource result.txt -cmu nfs://oldfqdn/ nfs://newfqdn/` – update existing files to point to a new storage location
- `DrTool -recoverFiles result.txt -csu win://old1/ win://new1/ -csu win://old2/ win://new2/ -cmu nfs://oldfqdn/ nfs://newfqdn/` – recover files to different servers and update the secondary storage location simultaneously

Expert Options

To use the expert options below, include `-expert` as the first parameter.

WARNING: restoring old ACLs into locations that may have changed significantly since the files were originally migrated can lead to unintended and unpredictable results. Use these options with caution.

- `-do-not-restore-acls` – this is the DEFAULT behavior (recreated files and folders will inherit Windows ACLs from the existing folders into which they are created)
- `-merge-acls` – restore files' explicit Windows ACLs and inherit from the parent folder
- `-exact-acls` – Windows ACLs will be restored exactly – if parent folder ACLs are now different, these will be propagated when a future change is made to the parent
- `-restore-acls-disable-inherit` – Windows ACLs will be restored, then inheritance will be disabled (_inherited_ central access policies and resource attributes will be discarded in favor of those inherited at the new destination)

If `-expert-` is specified when running via the GUI (rather than a batch operation), the GUI will prompt which of the above ACL options should be used.

6.8 Querying a Destination

While it is strongly recommended to obtain Recovery files from a *'Create Recovery File From Source'* Policy, where this has been overlooked it is possible to obtain Recovery files from the destination. However, some changes in the source file system, such as renames and deletions, may not be reflected in these results.

Querying the Destination from AdminCenter

Run a *'Create Recovery File From Destination'* Policy, see §4.18 (p.31).

Appendix A

Pattern Matching Reference

This appendix details the specifics of the pattern-matching syntax for filename and owner patterns used in Rules (see §1.4.4 (p.4)).

A.1 Wildcard Patterns

The following wildcards are accepted:

- `?` – matches one character (except '/')
- `*` – matches zero or more characters (except '/')
- `**` – matches zero or more characters, including '/'
- `/**/` – matches *zero* or more directory components

Literal commas within a pattern must be escaped with a backslash.

Examples of Supported Wildcard Patterns:

- `*` – all filenames
- `*.doc` – filenames ending with .doc (including '.doc')
- `?*.doc` – filenames ending with .doc (excluding '.doc')
- `*.do?` – filenames matching *.doc, *.dot, *.dop, etc. but not e.g. *.docx
- `???.*` – filenames beginning with any three characters, followed by a period, followed by any number of characters
- `*\,*` – filenames containing a comma

Examples of Using `*` and `**` in Wildcard Patterns:

- `/*.doc` – matches files ending with `*.doc` *directly* within the Source URI location, but *not* within its subdirectories
- `public/*` – matches all files *directly* within *any* directory named 'public'
- `public/**` – matches all files at *any* depth within *any* directory named 'public'
- `public/**/*.pdf` – matches all .pdf files at *any* depth within *any* directory named 'public'
- `/home/*.archived/**` – matches the contents of any directory ending with '.archived' directly within the home directory (<Source URI>/home)

- `/*/public/**` – matches all files at any depth with *any* directory named 'public' where the public directory is *exactly* one level deep within the Source
- `/*/**/public/**` – matches all files at any depth with *any* directory named 'public' where the public directory is *at least* one level deep within the Source

A.1.1 Directory Exclusion Patterns

In most cases, directory inclusion and exclusion of specific directories is best performed using Subdirectory Filtering (see §1.4.2 (p.4)).

For more complex cases, e.g. where directories to be excluded follow a pattern, wildcard patterns ending with '`/**`' may be used to match all files in a particular tree.

A.2 Regular Expressions

More complex pattern matching can be achieved using regular expressions. Patterns in this format **must** be enclosed in a pair of '/' characters. e.g. `/[a-z].*/`

To assist with correctly matching file path components, the '/' character is **only** matched if used explicitly. Specifically:

- `.` does NOT match the '/' char
- the subpattern `(.|/)` is equivalent to the normal regex '.' (i.e. ALL characters)
- `[^abc]` does NOT match '/' (i.e. it behaves like `[^/abc]`)

Literal commas within a pattern must be escaped with a backslash.

It is recommended to avoid regex matching where wildcard matching is sufficient to improve readability.

Examples of Regular Expressions:

- `/.*/` – all filenames
- `/.*\.doc/` – filenames ending with .doc
- `/~[w|$].+/` – filenames beginning with ~w or ~$ followed by one or more chars
- `/.*\.[0-9]{3}/` – filenames with an extension of three digits
- `/public/.+\.html?/` – .htm and .html files *directly* within *any* 'public' directory
- `/public/(.|/)*/` – equivalent to wildcard pattern `public/**`
- `/public/((.|/)+/)*index.html/` – equivalent to `public/**/index.html`

A.3 Case Sensitivity

By default, all patterns are matched case-insensitively using Unicode case-folding. Specifically, two characters are considered equal if they are the same or if they are equivalent in the Unicode case-folding table. This comparison is independent of language or locale settings. For example 'café', is considered equal to 'CAFÉ'.

Of particular note is that Unicode case-folding will match lowercase 'i' with both the usual ASCII uppercase 'I' *and* the Turkish dotted-i ('İ').

To match using ASCII-only case-insensitivity (i.e. disable Unicode case-folding), prefix pattern with '`(?-u)`'.

To match case-sensitively, prefix pattern with '`(?-i)`'.

For a list of a patterns, these prefixes must be applied to each individual pattern as appropriate.

Appendix B

Network Ports

The default ports required for Moonwalk operation are listed below.

B.1 Admin Tools

The following ports must be free before installing Admin Tools:

- 443 (AdminCenter web interface – configurable during installation)
- 8005

The following ports are used for outgoing connections:

- 4604-4609 (inclusive)
- 443 (to contact the Global Licensing Service if using a capacity-based license)

Any firewall should be configured to allow incoming and outgoing communication on the above ports.

B.2 Agent / FPolicy Server / LinkConnect Server

The following ports must be free before installing Moonwalk server components:

- 4604-4609 (inclusive)

Any firewall should be configured to allow incoming and outgoing communication on the above ports.

NFS

For each file server using Moonwalk Agent to connect to an NFS device, open TCP **and** UDP ports for the RPC Portmapper, Mount service and NFS service. Optionally,

a Moonwalk Gateway Agent may be installed on the AdminCenter machine to facilitate browsing; this machine will then also require access to the above ports.

The Portmapper always resides on port 111. The Mount and NFS ports however are registered with the Portmapper and may change when services are restarted. Please refer to firewall documentation regarding SUN RPC and the Portmapper as well as NFS service documentation for further details. The simplest solution is often to force the Mount and NFS services to use fixed port numbers.

Other Ports

Moonwalk plugins may require other ports to be opened in any firewalls to access storage devices / services from Gateway Agent machines.

Please consult specific device or service documentation for further information.

Appendix C

AdminCenter Security Configuration

C.1 Updating the AdminCenter TLS Certificate

The webserver TLS certificate may be updated using the following procedure:

1. Go to `C:\Program Files\Moonwalk\AdminTools\`
2. Run `Update Webserver Certificate`
3. Provide a PKCS#12 certificate and private key pair

Important: The new certificate MUST appropriately match the original AdminCenter FQDN specified at install time.

C.2 Security Roles and IP Restrictions

The *'Settings'* → *'Security Roles'* page allows for the configuration of multiple security roles to cater for different groups of users in Active Directory. Either full access or read-only access may be granted to an AD group, and access for each group may be limited by IP Address / Subnet.

Webhook clients and Anonymous access for Secure Link Users – used to control access to shared links (e.g. statistics reports) – may also be restricted by IP Address / Subnet.

If Active Directory integration was not configured at install time, you will still be able to restrict by IP Address / Subnet for the Administrator and Secure Link Users roles.

C.3 Password Reset

Normally, the administration password is changed on the *'Settings'* page as needed.

However, should the system administrator *forget* the username or password entirely, the credentials may be reset as follows:

1. Go to `C:\Program Files\Moonwalk\AdminTools\`
2. Run `Reset Web Password`
3. Follow the instructions to provide new credentials

Note: If AdminCenter has been configured to use Active Directory LDAP authentication, then passwords should be changed via Active Directory as normal – this section applies only to local credentials configured during installation.

Appendix D

API Access

D.1 Webhooks

Tasks may be launched from external applications using webhooks. Enable this option on a Task's *'Task Details'* page to generate a Task-specific webhook URI.

Since Moonwalk webhook URIs contain an authentication token, they must be kept secure to prevent unauthorized invocation. However, a webhook URI may be revoked at any time by either disabling the corresponding Task's webhook completely, or by regenerating the webhook URI (thereby revoking the previous URI's token).

Access to webhooks may be restricted by IP Adddress / Subnet – see §C.2 (p.97).

To launch a Task, simply send an empty HTTP POST to the Task's webhook URI. HTTP response status codes are given in the table below.

Status	Meaning
200	Task launched
403	The webhook is invalid or has been revoked
404	The Task no longer exists
500	Failed to launch: check Task configuration

Note: An entry will appear in *'Running Tasks'* on the *'Dashboard'* only for tasks that are successfully launched. If a second launch of the same Task is attempted immediately after a success, the second HTTP response will also be 200, but the Task will not actually be launched a second time (because the Task is already running). This avoids undesirable errors being reported when a webhook invocation is retried by the client due to a communications error.

D.2 Management API

If included in your license, most AdminCenter functions may be invoked via the EMA REST management API. This API can be used to integrate Moonwalk with existing systems in the enterprise for automation, monitoring, statistics analysis and reporting.

EMA authentication keys are created and revoked from the *'Settings'* page. Each key can be configured to provide FULL or READONLY access, and may be restricted by IP Address / Subnet as required.

Given the capability of a Moonwalk system to access filesystems across multiple servers, care should be taken to ensure the security of these keys:

- Store API keys securely
- Restrict access to machines upon which client applications are installed (and their storage)
- Use READONLY API keys where possible
- Use a separate API key for each client application – this will allow them to be revoked separately if necessary
- Always revoke API keys for any applications that are no longer in use
- Use IP Address / Subnet restricted API keys

D.3 Service Probe

To remotely test whether the Moonwalk Webapps service is responding, perform an HTTP GET request on the following resource:

`https://<serverFQDN>[:<port>]/eagle/probe`

For example, to probe with curl:

```
curl -i -k 'https://server.example.com/eagle/probe'
```

The service will respond with 200 OK.

Appendix E

Database Integration

Moonwalk can be integrated with various external SQL databases (such as Oracle Database, Microsoft SQL Server, PostgreSQL, MySQL, MariaDB, etc.) to allow both external selection of files on which to apply Policies, and to export Policy operation results to external workflows. This functionality may used either instead of, or to complement, EMA API-based integrations.

Note: Database integration support is a separately licensed feature.

E.1 Database Connections

Prior to using an external database, its JDBC driver must be installed and configured. To get started, navigate to *'Settings'* → *'Database Connections'* and follow the instructions to install the appropriate driver jar file.

Once the driver is installed, add a connection. You will be required to provide a JDBC URL and any additional key-value configuration options.

Read the sections below to ensure that you understand the configuration requirements for each integration use case.

Once the connection has been configured, use the *'Test'* button in the Actions column to verify connectivity, authentication, etc.

E.2 DB Sources

A DB Source describes an SQL query which may be used as an alternative to Sources and Rules to select the files upon which a Policy is to operate. Each row of the result corresponds to a single file.

Depending on the software that populated the database table(s), the file locations may be split across multiple fields (e.g. SMB server, directory path, filename, etc) or expressed as a single string such as a Windows file path. The DB Source editor provides

a number of methods to easily map a set of result columns to a Moonwalk file URI. Escaping, '\' vs '/' conversions, etc, are handled automatically. Refer to the user interface for a full list of available mappings.

Once a DB Source has been configured, it is recommended to use the *'Test settings'* button to ensure that the SQL query is valid and that results have been mapped to URIs as expected.

E.2.1 JDBC Configuration Considerations

As Policies are executed over a DB Source, the database query may remain open for a long time (because a large number of file operations are being performed). Additionally, the rate of consumption of query results may vary significantly throughout the run. Consequently, it is necessary to ensure that the JDBC driver is suitably configured to support this kind of streaming query.

Specifically:

- The query must stream / 'batch' results – **not** return all results up front
- The query should **not** time out if results are not consumed for a while

MySQL

The MySQL Connector/J driver (com.mysql.cj.jdbc.Driver) provides two options for streaming.

By default, Moonwalk will stream result sets row-by-row by specifying the MySQL-specific setFetchSize(Integer.MIN_VALUE) option.

Alternatively, if the `useCursorFetch` driver option is enabled, cursor-based streaming will be used instead. This method uses more server-side resources since it creates a temporary table to back the query.

MariaDB

When using the MariaDB Connector/J driver (org.mariadb.jdbc.Driver), to prevent time out of connections when streaming, specify an option with key `sessionVariables` and value `net_write_timeout=10800`. The timeout value is expressed in seconds and may be increased further as necessary.

E.2.2 Policy Behavior

Generally, Policies will behave the same when processing regular Sources and DB Sources. The exception to this rule is that, during a Delete Policy on a DB Source, the *size* of deleted files will not be included in Server statistics (although the successful delete operation will be recorded).

E.3 Logging to a Database

Database logging is configured in a Policy by checking the *'Log to database'* option. This feature can be configured to perform one or more of the following:

- Perform a sequence of SQL statements, in a single transaction, before any files are processed (Pre-run SQL Statements)
- Log the results of individual Policy operations
- Perform a sequence of SQL statements, in a single transaction, after all file operations have been performed and logged (Post-run SQL Statements)

A number of variables are available as parameters to the Pre- and Post-run SQL statements – refer to the help information within AdminCenter.

E.3.1 Database-Specific Considerations

When creating a table to receive log information, choose column types carefully based on your database product and configuration. URIs, error reasons and full error traces in particular may be very large, error fields may also be multi-line. Textual fields may contain Unicode. Numerical fields may contain 64-bit values.

Microsoft SQL Server

Typically, the NVARCHAR type should be used in preference to VARCHAR for Unicode support.

MySQL

The CSV engine is generally not suitable for logging since some fields may be NULL (depending on the particular operation result being logged), and NULLs are not allowed by this engine.

Appendix F

Advanced Agent Configuration

Agents may be configured on a per-server basis via the AdminCenter *'Servers'* page.

When the configuration options are saved, they are pushed to the target server to be loaded on the next service restart. In the case of a cluster, all nodes will receive the same updated configuration.

F.1 Logging

Log location and rotation options may be adjusted if required. Debug mode may impact performance and should **only** be enabled following advice from Moonwalk Support.

Additionally Moonwalk can be configured to send UDP syslog messages in either RFC5424 or RFC3164 format. Syslog output is not enabled by default.

F.2 Stub Deletion Monitoring

As described in §5.1.6 (p.37), on Windows file systems, Moonwalk can monitor stub deletion events in order to make corresponding secondary storage files eligible for removal using Scrub Policies.

This feature is not enabled by default. It must be enabled on a per-volume basis either by specifying volume GUIDs (preferred) or drive letters. Volume GUIDs may be determined by running the Windows `mountvol` command or Powershell `Get-WmiObject -Class win32_volume`. For Windows clustered volumes, the cluster volume **must** be specified using a volume GUID.

Note: This feature should not be configured to monitor events on backup *destination* volumes. In particular, some basic backup tools such as Windows Server Backup copy individual files to VHDX backup volumes in a manner which is not supported and so such volumes **must not** be configured for Stub Deletion Monitoring. Of course, deletions may still be monitored on source data volumes.

F.3 Parallelization Tuning

When a Policy is executed on a Source, operations will automatically be executed in parallel. Parallelization parameters may be tuned for each Server if necessary.

F.4 NFS Client

By default, the built-in NFS client accesses NFS servers using a UID of 0 (root) and GID of 1. Wherever possible, NFS exports should be configured to allow such access, rather than reconfiguring the client.

All NFS settings will apply to **ALL** NFS connections from the given agent.

F.5 Deny Demigrations

Applications may be denied the right to demigrate stubs. Such an application – specified either by application binary name or full path – will be unable to access a stub and demigrate the file contents (an error will be returned to the application instead).

Note: Only local applications (applications running directly on the file server) may be blocked.

F.6 Enabled / Available Plugins

Storage plugins may be configured and enabled as necessary for each server. For plugin-specific details, refer to the appropriate section of Chapter 5.

F.7 Manual Overrides

Additional options may be manually entered. Undocumented options should not be entered unless under the direction of a Moonwalk support engineer.

Unknown Content-Encoding behavior

By default, an `UNEXPECTED_CONTENT_ENCODING` error is returned if an attempt is made to copy a file from an object storage Source where the source file has a non-identity Content-Encoding. This restriction avoids copying encoded data between systems in a way that will cause confusion at the destination. To override this behavior, and instead copy the raw bytes *as encoded*, specify a Manual Override:

- Azure: `azure.dca.passthroughUnknownContentEncoding=true`
- GCP: `google.dca.passthroughUnknownContentEncoding=true`
- AWS S3: `s3.passthroughUnknownContentEncoding=true`
- S3 Generic: `s3generic.passthroughUnknownContentEncoding=true`

- Cloudian S3: `s3cloudian.passthroughUnknownContentEncoding=true`
- Dell EMC ECS S3: `s3ecs.passthroughUnknownContentEncoding=true`
- IBM COS S3: `s3bluemix.passthroughUnknownContentEncoding=true`
- Wasabi S3: `s3wasabi.passthroughUnknownContentEncoding=true`
- DataCore Swarm S3: `s3swarm.passthroughUnknownContentEncoding=true`

F.8 Upload Configuration

Under some circumstances it may be necessary to upload a configuration file under the direction of a Moonwalk support engineer. The configuration for this server will be REPLACED in its entirety.

Appendix G

Troubleshooting

Before contacting Moonwalk Support, please review relevant log files and Windows Event Viewer for error messages.

G.1 Log Files

AdminCenter Logs

AdminCenter logs describing each attempted policy operation are accessed through the Recent Tasks panel on the *'Dashboard'*. This is the first place to look when investigating a Policy problem.

The AdminCenter also maintains a *'Global Log'* (accessible from the *'Help'* page) which summarizes Policy start / stop activity.

For other issues, including failure of user-initiated demigrations, it will often be necessary to consult the Agent logs on the servers in question.

Server Statistics

In addition to log files, AdminCenter also provides per-cluster and per-node charts of operation successes and failures on each Server's *'Server Details'* page. This includes information about failed demigrates over time which may be useful in conjunction with a Server's log files to troubleshoot user-initiated demigration issues.

Agent Logs

Location:

- Windows: `C:\Program Files\Moonwalk\logs\Agent`
- Linux: `/var/opt/moonwalk/moonwalk-agent/log`

There are two types of Agent log file. The `agent.log` contains all Agent messages, including startup, shutdown, and error information, as well as details of each individual file operation (migrate, demigrate, etc.). Use this log to determine which operations have been performed on which files and to check any errors that may have occurred.

The `messages.log` contains a subset of the Agent messages, related to startup, shutdown, critical events and system-wide notifications.

Log messages in both logs are prefixed with a timestamp and thread tag. The thread tag (e.g. `<A123>`) can be used to distinguish messages from concurrent threads of activity.

Log files are regularly rotated to keep the size of individual log files manageable. Old rotations are compressed as gzip (.gz) files, and can be read using many common tools such as 7-zip, WinZip, or zless. To adjust logging parameters, including how much storage to allow for log files before removing old rotations, see Appendix F.

Log information for operations performed as the result of an AdminCenter Policy will also be available via the web interface.

DrTool Logs

Location: `C:\Program Files\Moonwalk\logs\drtool`

DrTool operations such as recovering files are logged in this location. DrTool will provide the exact name of the log file in the interface.

G.2 Interpreting Errors

Logged errors are recorded in a tree format which enables user-diagnosis of errors / issues, as well as providing detail for any further investigation by support engineers.

Error trees are structured to show WHAT failed, and WHY, at various levels of detail. This section provides a rough guide to extracting the salient features from an error tree.

Each numbered line consists of the following fields:

- WHAT failed – e.g. a migration operation failed
- WHY the failure occurred – e.g. '`[FILE_NOT_FOUND]`'
- optionally, extra DETAILS about the failure – e.g. the path to a file

As can be seen in the example below, most lines only have a WHAT component, as the reason is further explained by the following line.

A Simple Error

```
ERROR demigrate win://server.test/G/source/data.dat
[0] ERR_DMAGENT_DEMIGRATE_FAILED [] []
 [1] ERR_DMMIGRATESUPPORTWIN_DEMIGRATE_FAILED [] []
  [2] ERR_DMAGENT_DEMIGRATEIMP_FAILED [] []
   [3] ERR_DMAGENT_COPYDATA_FAILED [] []
    [4] ERR_DMSTREAMWIN_WRITE_FAILED [DISK_FULL] [112: There is
        not enough space on the disk (or a quota has been reached).]
```

To expand the error above into English:

- demigration failed for the file: win://server.test/G/source/data.dat
- because copying the data failed
- because one of the writes failed with a disk full error
    - the full text of the Windows error (112) is provided

So, `G:` drive on `server.test` is full (or a quota has been reached).

Errors with Multiple Branches

Some errors result in further action being taken which may itself fail. Errors with multiple branches are used to convey this to the administrator. Consider an error with the following structure:

```
[0] ERR...
 [1] ERR...
  [2] ERR...
   [3] ERR...
    [4] ERR...
     [5] ERR...
      [6] ERR...
   [3] ERR...
    [4] ERR...
     [5] ERR...
```

Whatever ultimately went wrong in line 6 caused the operation in question to fail. However, the function at line 2 chose to take further action following the error – possibly to recover from the original error or simply to clean up after it. This action also failed, the details of which are given by the additional errors in lines 3, 4 and 5 at the end.

Check the Last Line First

For many errors, the most salient details are to be found in the last line of the error tree (or the last line of the first branch of the error tree). Consider the following last line:

```
[11] ERR_DMSOCKETUTIL_GETROUNDROBINCONNECTEDSOCKET_FAILED [COULD_NOT_RE
     SOLVE_HOSTNAME] [host was [svr1279.example.com]]
```

It is fairly clear that this error represents a failure to resolve the server hostname `svr1279.example.com`. As with any other software, the administrator's next steps will include checking the spelling of the DNS name, the server's DNS configuration and whether the hostname is indeed present in DNS.

G.3 Getting Help

Join the free Moonwalk user community forum at:

`https://forum.moonwalkinc.com`

Starter Edition users may purchase official support or a full product upgrade. See:

`https://www.moonwalkinc.com/how-to-buy`

G.4 Contacting Support

If an issue cannot be resolved after reviewing the logs, customers with a current maintenance contract should contact Moonwalk Support at:

`https://servicedesk.moonwalkinc.com/`

Details of your support tier are included in your support agreement.

Appendix H

Glossary

ACL - Access Control List; file/folder/share level metadata encapsulating permissions granted to users or other entities

Caretaker - a specific node within a cluster that performs maintenance tasks that must be run on a single node at a time

CLP - Counterfeit Link Protection

Demigrate - to return migrated file content data to its original location, e.g. in response to user access

DFS - Microsoft's Distributed File System; comprised of DFSN and DFSR

DFSN - DFS Namespace; a Windows mechanism allowing for the presentation of multiple SMB shares as a single logical share

DFSR - DFS Replication; an SMB share-based file replication technology, see also Storage Replica as an alternative

DR - Disaster Recovery

EMA - a REST API providing access to AdminCenter management functions

FPolicy - a component of NetApp Data ONTAP which enables extension of native Filer functionality by other applications

FPolicy Server - a server which connects a NetApp Filer via the FPolicy protocols in order to provide extended functionality

FQDN - Fully Qualified Domain Name, e.g. *server1.example.com*

GID - Group ID, e.g. of a UNIX user group

GUID - Globally unique identifier

HA - High-Availability; specifically the provision of redundant instances of a resource in a manner which ensures availability of service, even in the event of the failure of a particular instance

LDM - Link Deletion Monitoring

Link-Migrate - to transparently relocate file content data to secondary storage, replacing the original file with a MigLink

MigLink - a placeholder for a file that has been Link-Migrated; applications accessing the MigLink will be transparently redirected to the corresponding LinkConnect Server to facilitate data access

Migrate - to transparently relocate file content data to secondary storage without removing the file itself; the existing file becomes a *stub*

MWI file - a file on secondary storage which encapsulates the file content data of a corresponding primary storage *stub* file or *MigLink*

NTP - Network Time Protocol, a protocol for clock synchronization between computer systems over a network

Quick-Remigrate - to quickly return a previously demigrated (but unmodified) file back to its migrated state without the need to re-transfer file content data

Recovery File - a text file describing the relationships between stubs/MigLinks and their corresponding MWI files

Scheduler - the AdminCenter component responsible for starting scheduled Tasks

SDM - Stub Deletion Monitoring

SMB - Server Message Block

Stub - a file whose content data has been transparently migrated to a secondary storage location

Storage Replica - a Windows Server volume replication technology offering synchronous or asynchronous replication modes

Syslog - a protocol used to send system log or event messages, e.g. to a centralized Syslog collector

TLS - Transport Layer Security; a protocol used for establishing secure connections between servers (formerly known as SSL)

UID - User ID, e.g. of an UNIX user

UUID - Universally unique identifier