

ArchiverFS V3.4.x Full User guide (Last updated 17th Oct 2026)

Contents

ArchiverFS V3.4.x Full User guide (Last updated 17th Oct 2026) 0
Introduction to ArchiverFS 1
Basic Concepts 2
Versions 3
Upgrading From a Previous Version 4
Clusters and DFS Shares 4
Choosing Your Archive Storage 5
NAS Devices and Non Windows Storage 6
Installing the Service Account 7
Installing Prerequisites 8
Installing the Archive Share\Destination 9
    Making the Archive Read Only for Users 9
Installing the Software 9
    Upgrading from a previous version 10
Installing Symbolic Link Support 10
Setting up a New Archive Job 11
    Testing Your First Job 13
Link Types 13
Exclusions 14
    Directory path exclusions 14
    File based exclusions 14
Setting up a New Graveyard Job 15
Log Settings 15
Notifications 16
Maintenance Job 16
    Understanding Missing Links 17
File Types 18

Pre/Post Job Tasks 18
Compressing Your Second Line Storage 19
Deduplicating Your Second Line Storage 19
Enabling Long Path Support 19
Dealing With Long Paths 20
Reporting 21
Backing Up Your Configuration 22
Additional Tools 22
Getting Help 23

This guide explains how to install ArchiverFS, set up archive and graveyard jobs, and maintain access to archived files. Work through the installation sections before running your first job, then use the later sections as a reference for day-to-day operation.

Introduction to ArchiverFS

ArchiverFS does not use a database or proprietary data store to hold archived files or record their locations. Instead, it stores them as ordinary files and folders, using NTFS or ReFS storage and preserving the directory structure, permissions, attributes and timestamps where the destination supports them. It is a straightforward way to move old files from a live SAN, NAS or file server to less expensive second-line storage and implement a full file lifecycle management policy.

A SQL database is still used for the software's configuration, such as the jobs you create and the maintenance settings. The important distinction is that your archived content remains in the file system. The 'Backing Up Your Configuration' section explains how to protect both parts.

You also get all of the features that you would expect including the ability to create seamless links which end users can use to access migrated files without the intervention of IT Support resources, without the overhead of big expensive meta-data tracking databases or vendor lock-in of proprietary file storage formats.

A) Symbolic links can provide seamless access to archived files from compatible Windows clients. Folder and file shortcuts are also available where a different way of accessing archived content is more suitable.

B) Unicode support allows ArchiverFS to work with filenames in many languages, including English, Chinese, Hindi and Arabic. Include representative filenames in your initial testing.

C) No 3rd party client software or file system drivers to install on client machines. Unlike some products ArchiverFS does not need any client software to be installed on your workstations. With other products this is sometimes necessary to allow the seamless recall of old files from second line storage. ArchiverFS only uses functionality that is already built into Windows. There

is just a small configuration change to implement on each workstation to enable this functionality.

D) ArchiverFS can manage large archive file systems without maintaining a database entry for each archived file. The storage and network you choose still determine how quickly jobs run and how quickly users can open archived content.

E) Because your archive is made up of normal files and folders, you can use your existing storage and backup tools. There is no proprietary content store to size or manage.

F) File-level security is retained through the permissions replicated from the live system to the archived files and folders. Before using non-Windows storage, check that it can preserve the permissions and other metadata you need. The storage and service-account sections explain what to test.

Basic Concepts

Before installing ArchiverFS, read the installation sections below and the Quick Setup Guide available from the MLtek documentation page. Start with a small test job so you can check the settings and the way users will access archived files.

ArchiverFS works by migrating old and unused files from your live network shares into a dedicated share that you have setup somewhere on your network.

Each job can move old files to a different destination; they do not all have to use the same share. For a new job, choose a dedicated, empty destination folder.

A graveyard job provides the final stage of the file lifecycle. It moves qualifying files out of an archive job's destination and into a separate location, removing their links from the live file system. You can then back up that content and delete it when your retention policy allows.

ArchiverFS preserves the directory structure and is designed to retain file permissions, attributes and Created, Modified and Last Accessed dates. On non-Windows storage, or when using options that relax security replication, test which properties are retained before moving live content.

Second line storage space is typically made up of less expensive storage hardware. Typically this is either several years old and/or its performance will typically be lower than the first line hardware providing the live file system.

This enables organisations to re-use hardware that would otherwise be disposed of, thus significantly increasing the return on investment.



1. Files that have not been used in the period of time specified are moved to long term storage
2. Users open archived files directly through the links left behind. Files that meet your return policy are subsequently moved back to the live file system by the maintenance job.
3. Files that meet your graveyard policy are moved to a separate location, where you can back them up and eventually delete them.

Opening an archived file does not require ArchiverFS to return it to live storage first. The user's computer accesses the archived file directly, so it must be able to reach the archive share and have permission to open the file. This remains true if the archive is in another site or in the cloud.

Versions

	Trial	Express	Standard	Datacenter
Max number of Archive jobs*	5	1	5	Unlimited
Max files processed per run	100	Unlimited	Unlimited	Unlimited
Perpetual buy once use forever licence**	Evaluation	Yes	Yes	Yes
Includes 12 months of assisted support	No	Yes	Yes	Yes
Directory structure retained	Yes	Yes	Yes	Yes
File-system permissions retained where supported	Yes	Yes	Yes	Yes
File attributes retained	Yes	Yes	Yes	Yes
Pre + Post job tasks	Yes	Yes	Yes	Yes
Various stub types, including symbolic links	Yes	Yes	Yes	Yes
Optimised for multi-processor multi-core systems	Yes	Yes	Yes	Yes
100% Developed on the .Net framework	Yes	Yes	Yes	Yes

*The Maintenance job and Graveyard jobs do not count towards the total jobs limit.

Subscription licences provide an unlimited number of archive jobs with an annual allowance for data moved from live to archive storage. The allowance is based on processing, rather than on how much content remains in the archive. See the current MLtek pricing page or your quotation for licence terms and prices.

Upgrading From a Previous Version

If you are upgrading from a version prior to Version 3.2 then please contact MLtek support via Support@MLtekSoftware.com. If you are upgrading from version 3.2 or later simply uninstall the old version of ArchiverFS choosing to 'keep the configuration database' when prompted, then install the new version.

Before uninstalling, back up the configuration database and record the service account, SQL connection and installed build. If you use a custom SQL instance, make sure the new installation reconnects to that instance. After installation, check that your jobs and global settings are present before enabling schedules.

Read the release notes for the exact build you intend to install and all intervening releases. The 3.4x label covers many builds with different fixes and options. Keep a recoverable copy of the previous configuration until the upgrade has been checked.

Clusters and DFS Shares

ArchiverFS supports Windows Server clustered shares and Windows DFS shares. In both cases, use the published service path rather than a path to one individual server.

This is important because jobs and the links they create use UNC paths. Those paths must continue to refer to the correct file system when a cluster fails over or DFS directs users to another server.

For example, a DFS namespace might publish \\domain.local\Files\Data while its underlying folders are held on Server1 and Server2. Add the domain-based namespace path to your job, rather than \\Server1\Data or \\Server2\Data.

For a Windows failover cluster, use the clustered file service name. DFS namespaces and failover clusters are different technologies, but the same principle applies: configure the stable path that users should continue to use.

If DFS replication is involved, allow it to catch up after a large initial archive run. MLtek recommends disabling the Maintenance job for that first run, monitoring the replication queues, and enabling maintenance again once the queues have cleared. This helps avoid maintenance acting on an incomplete view of the live file system.

Choosing Your Archive Storage

The storage you choose as your secondary\archive storage can have a big impact on the performance of your solution. We recommend the following:

1) Choose archive storage that can be presented to the network via a Windows OS.

   For example, you can attach block storage to a Windows server or VM using iSCSI, format the volume with a suitable Windows file system and share it over SMB. The underlying disks can still be provided by a SAN or another non-Windows storage system.

   The important point is that a Windows file server is presenting the archive share and managing its file-system permissions. An NFS share is not the same as a local block volume that Windows can format as NTFS.

   A non-Windows NAS or gateway may work satisfactorily, but SMB access alone does not prove that it supports every Windows permission, audit entry or timestamp. Test the properties your deployment needs, as described in the next section.

2) Choose storage that is fast enough.

   But its archive storage, why does it have to be fast?

   It does not have to match the performance of your live storage, but it does have to be fast enough for users. When users browse folders containing symbolic links, Windows will read information from the archive storage. Slow metadata access at the storage level can therefore make a live folder browsing and traversal feel slow even before a document is opened.

   If your secondary\archive storage is too slow this can slow down users accessing the live file system, they will experience delays when opening folders with lots of Symlinks. If you intend to use Windows shortcuts (.lnk) files as stubs for archived files then this is not so much of an issue.

3) Choose something with at least some redundancy\fault tolerance.

   Things go wrong, things break. If you do not include any fault tolerance in your secondary\archive storage then it can still break and users can lose access to the secondary\archive storage.

   You probably don't have to use the same level of redundancy as the primary\live storage, but having at least some is still a good idea.

4) Using cloud storage? Make sure your internet connection is fast enough.

   This includes BOTH raw bandwidth AND latency.

When planning cloud storage, test access from an ordinary user's computer as well as from the machine ArchiverFS is installed on. Users read archived content directly; it does not pass back through the ArchiverFS server, but you need to ensure both user access and the software's access is fast enough for the intended deployment. Use Pre-Scan to estimate the processing time for the first migration.

An object-storage bucket is not itself a UNC destination. Services such as Wasabi need a suitable gateway or file service to present an SMB share. Test permissions, timestamps, links and restore behaviour, and repeat those tests after gateway upgrades.

NAS Devices and Non Windows Storage

If you want to migrate content to or from non-Windows storage, or between separate domains, allow time for additional planning and testing. ArchiverFS normally tries to carry the live file's permissions with it when the file moves (default setting, but can be changed). The destination must be able to recognise the relevant users and groups SID's and store the required security information.

For example, a NAS that does not recognise your domain identities may be able to store the file contents but be unable to reproduce their permissions correctly. Check domain integration, identity resolution and effective access before using it as an archive destination.

So how do you know if your planned archiving plan will support the retention of the permissions on you live files when they are archived? There is a very easy test you can do.

1) Sign in to the ArchiverFS machine using the account that will run the ArchiverFS service, where your account policy permits this test.
2) Open File Explorer and browse to the proposed archive destination using its UNC path.
3) Right click and create a .txt file.
4) Right click on the text file and click on 'Properties'. Try to grant permissions to the file to the users and groups that are referenced in your live file systems security, e.g. the domain based 'Domain Users' group, the domain based 'Accounts Department' group etc.
5) If you can set these permissions, the destination has passed a useful first check. Next, run a small archive test with representative files and compare permissions, ownership, audit entries, attributes and timestamps. Check access as both an authorised user and a user who should be denied access.
6) If the test fails, check the account's rights and the device's domain integration before changing job settings. Best Efforts Security may allow a job to continue when exact permission replication is not possible, but you should understand and approve the resulting access first.

The table below summarises the checks to make. Storage models and firmware vary, so a test of your actual source and destination is more useful than a blanket compatibility claim.

'Persist Audit Entries' is a separate option on archive jobs. Leaving it disabled (the default setting) can help when a storage device cannot reproduce Windows audit settings. Access permissions and audit settings are different things and require different user access to set\replicate\control.

Storage arrangement	What to check
Windows hosted SMB shares	Service rights, effective user access, link evaluation and representative applications
Domain integrated NAS	Access permissions, ownership, auditing, timestamps and complete archive and return tests
NAS without matching domain identities	Whether source permissions can be represented at all; approve any alternative security model before moving files
SMB gateway to object storage	All metadata checks plus gateway upgrade, tiering, unavailable-target and restore behaviour

Installing the Service Account

Before you run the install you need to create a user account that the ArchiverFS Service will use. The user account must:-

- Have 'Full Control' permissions over all the files you will be archiving.

- Have 'Full Control' permissions over the location(s) you will be archiving too.

- Have the 'Create Symbolic Links' right on the machine ArchiverFS is installed on, any machines it will move content from and any machines it will move content too.

- Have the 'Manage Auditing and Security Log' right on the machine ArchiverFS is installed on if you plan to use the 'Persist File Audit Entries' option, as well as on any machines it will archive content from and any machines it will move content too.

Use a dedicated account for the service. Grant the required rights and access explicitly through the policies used for your servers. Membership of Domain Administrators is not the only way to provide these rights and should not be used simply as a shortcut around an unexplained access problem.

If the account is not a Domain Administrator, make sure both the share permissions and the NTFS file-system permissions allow the required access, in addition to the user rights listed above.

It must have these specified on a) the machine it is running on b) any machines it will move files from c) and machines it will move files too.

You can grant these rights by either adjusting the Local Security Policy on the relevant machines, or by using Active Directory Group policies.

Create Symbolic Links = Open 'Local Security Policy' on the machine ArchiverFS is installed on.

Under Local Policies > User Rights Assignment, open Create symbolic links and add the service account. Alternatively, apply the setting through a suitable Active Directory Group Policy linked to the relevant servers.

Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Local Policies -> User Rights Assignment.

Manage Auditing and Security Log = When ArchiverFS migrates a file it copies all attributes and security access options on it, and then compares the copy to the original file before deleting the original file.

The security access options not only include the NTFS Access Control Lists, but also the auditing options as well. To replicate the auditing settings the chosen account needs the Manage Auditing and Security Log right

Under Local Policies > User Rights Assignment, open Manage auditing and security log and add the service account. You can also apply this setting through the Group Policy used for the relevant servers.

Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment.

Also arrange the service account's logon rights and credential maintenance through your normal administration process. After a password or policy change, check that the ArchiverFSJobEngine service restarts successfully. File-system access and access to the SQL database are separate requirements.

Installing Prerequisites

ArchiverFS requires the pre-installation of the following components before you run the ArchiverFS setup:

a) SQL Server Express 2019 or later, or a custom SQL connection configured using the Custom DB Location instructions in the installation package. Confirm the supported SQL version for your ArchiverFS build before installing a newer SQL release.

b) Microsoft .NET Framework 4.8. Use the prerequisite version specified for your ArchiverFS installation package; .NET Framework 4.8 and 4.8.1 are separate releases.

c) The 64-bit SAP Crystal Reports runtime required by the installation package, if you want to use graphical reports. Use the runtime build specified by MLtek rather than assuming that any Crystal Reports runtime will work.

If you are using the default local SQL Express connection, install the named instance SQLExpress. If you use a differently named or remote instance, follow the Custom DB Location instructions rather than expecting ArchiverFS to discover it automatically.

The original setup route uses a Custom SQL Express installation so that the local Administrators group can be added to the SQL administrators list. Agree the required SQL administration and service-account permissions with your database administrator. Record the server, instance and database used by ArchiverFS for future upgrades or recovery.
Current links to the MLtek documentation, Microsoft prerequisite downloads and SAP runtime information are collected in Useful References at the end of this guide. The older Quick Setup PDF still names .NET 4.6.2; the 3.479913 release notes record the move to .NET 4.8.

Installing the Archive Share\Destination

Set up an empty destination for the old files you want to archive. A Windows-hosted SMB share is the simplest starting point. If the share is hosted by a NAS or gateway, complete the security and metadata tests in the preceding sections first.

For a normal writable archive, a typical setup grants Administrators Full Control and Domain Users Change and Read on the share. At file-system level, the service account needs Full Control and users need the access appropriate to their files. Use the user groups relevant to your organisation rather than granting every user access by default.

If there are entries for 'System', 'Administrator', 'Administrators' or 'Domain Administrators' they should not be removed/changed.

When ArchiverFS runs it will align the permissions on the subfolders it creates in the archive with those in the live file system, thus ensuring content stays secure.

Making the Archive Read Only for Users

If you want users to read archived files but save changes as new copies, restrict them to Read at the archive share level. Keep sufficient access for the ArchiverFS service account to manage the files. The file-system permissions still decide which files each user is entitled to read.

Check other group memberships and alternative shares before testing. Another access route may allow a user to change the same file. Test opening, saving and Save As in the applications people actually use. A read-only share is not immutable storage.

Installing the Software

Once you have carried out the above steps, you are ready to install ArchiverFS. The first step is to install SQL Express.

Prerequisites

ArchiverFS requires the pre-installation of the following components before you run the ArchiverFS setup:

a) SQL Express 2019 or later (Download Here), or the pre-configuration of a custom DB connection (see the 'Custom DB Location' folder that comes with the install package)

b) Microsoft .Net Framework v4.8 (Note v4.8 is distinct from v4.81)
https://dotnet.microsoft.com/en-us/download/dotnet-framework/net48

c) SAP Crystal Reports runtime engine for .NET Framework v4.0 if you want to use the reporting function. (Download the 64 bit version here)

If you are using SQL Express and not using a custom DB connection then SQL Express must be installed using it's default settings (thus creating a named instance called 'SQLExpress') with one exception…

We strongly recommend adding the local 'Administrators' group to the SQL admin users list when installing SQL Express to simplify the setup of the ArchiverFS service account. If you are installing SQL Express 2019 or later you will need to choose a 'Custom' install and leave everything at the default settings until you get to the screen where you can specify the SQL Administrators.

Once the pre-requisites have been completed run the installation routine on the machine you want to install ArchiverFS on, following the instructions on screen. Please make sure you install the program using an account with full administrative rights on the local machine.

Once setup has completed, open ArchiverFS and check that it connects to the intended configuration database (e.g. does the UI open without errors?). Confirm the service account and that the ArchiverFSJobEngine service starts successfully.

If you are upgrading, make sure to follow the steps on the website:

Upgrading from a previous version

The update process is very straight forward..

- Make sure there are no jobs running.
- Uninstall the old version answering 'Yes' when you are asked if you want to keep the configuration database.
- Ideally reboot the machine, but this is not essential if you can't.
- Install the new version. All of your previous configuration will be picked up and all processes will continue as before.

Installing Symbolic Link Support

If you intend to use the seamless symbolic-link action, configure the Windows clients that will follow those links. These normally include your workstations and the ArchiverFS host; a web or application server may also need the setting if it opens archived files on behalf of users.

Older ArchiverFS screens may call these Hard Links or Hard\Symbolic Links. Technically, the cross-share links are 'symbolic links'. A Windows hard link is another name for a file on the same volume, so it is not the mechanism used to point to an archive on another server.

Windows has separate settings for local-to-local, local-to-remote, remote-to-local and remote-to-remote symbolic-link evaluation. Check the current settings and enable the directions required, this will normally be 'Remote To Remote' on workstations and all options on servers.

The deployment example below enables all four directions. Use it only where all four are intended. Apply the approved settings through your usual managed configuration process and test from a normal workstation.

Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem]
"SymlinkLocalToLocalEvaluation"=dword:00000001
"SymlinkLocalToRemoteEvaluation"=dword:00000001
"SymlinkRemoteToRemoteEvaluation"=dword:00000001
"SymlinkRemoteToLocalEvaluation"=dword:00000001

You can inspect the effective settings from an elevated command prompt with fsutil behavior query symlinkevaluation. For example, fsutil behavior set symlinkevaluation R2R:1 enables remote-to-remote evaluation, commonly needed when a link on one network share points to another network share.

If a link will not open, try its archive target directly as the affected user. If direct access fails too, check the target's permissions and network availability before changing link evaluation.

Setting up a New Archive Job

Archive jobs move files from your live file system.

The first step to perform when setting up a new job is to specify the place you want old files to be moved to. This can be any empty folder on your system that is accessible via a UNC path (\\yourserver\share\folder). If you are selecting a DFS share then you must select it via the domain based DFS path.

We strongly recommend against using hidden admin shares as either archive destinations or live file sources e.g. \\Server1\E$.

Once you have set up the destination click on the 'Create Job' button and when the new job screen opens select 'New Archive Job' and click 'begin'

The first step of the wizard involves adding the locations you want to move old files from. All locations must be added as UNC paths. You can also add to and remove from this list later if you like by editing the job once it has been created. Once you have added the locations click on 'Next'.

The second screen is where you want to move old files to. Try to keep this path as short as you can, and again it must be a UNC path. Once done click on 'Next'.

The next screen contains the options that control what the job will do. The most commonly used options are explained below. Start with Pre-Scan enabled so that you can review the proposed selection before moving files.

Action = what exactly do you want the job to do:

> ***Copy (No Links)*** *— Copies files that match the criteria and keeps the originals in the live file system. This is useful for testing destination writes, permissions and metadata, but it uses real archive space. Check any separately enabled cleanup options before running a test.*
>
> ***Move (No Links)*** *— Moves matching files to the destination and removes them from the live file system. It leaves no link behind, so users must use the new location if they still need those files.*
>
> ***Move (Per Folder .Lnk Shortcuts)*** *— Moves matching files and leaves a Windows shortcut in each folder from which files have been archived. Users open the shortcut and browse the archived content. Maintenance can return files according to the return settings. Use a graveyard job when removing managed content from the archive.*

Move (Per File .Lnk Shortcuts) *— Moves matching files and leaves a Windows shortcut for each file. Users click the shortcut to open the archived file. Maintenance can return qualifying files later. Because a shortcut is a separate shell object, test applications that expect to open the original document path.*

Move (Per File Hard\Symbolic Links) *— Moves matching files and leaves a symbolic link with the original filename. Compatible Windows applications can follow the original path to the archived file. Check client link evaluation and application compatibility before using this action across the live file system.*

For a mixed Windows and Mac OS environment, read 'Link Types' before choosing the action. Do not assume that a Windows absolute UNC symbolic link will work transparently on a MacOS client.

Archive Files That Have Not Been Used In — Sets the age threshold for archiving. Check whether the job uses Last Accessed or Last Modified, then choose a period that fits your data. Three years is an example, not a recommendation for every organisation. Background software and Windows timestamp behaviour can affect Last Accessed dates.

Pre-Scan — Scans the file system and reports what would qualify without moving files or creating links. At the end of the run, review the log for the number and size of files that would have been archived, together with any errors or skipped locations. Pre-Scan is not subject to the normal evaluation licence's 100-file limit. It is useful for estimating storage and licensing needs; a later move test is still needed to check links and user access.

Only archive content greater/less than X folders deep = Typically this option will be left at its default values (0 / 100). It allows you to specify the minimum and maximum folder depth for archiving to take place. Useful if you want to target a single root folder, and only process files in subfolders not the root itself.

Delete Empty Folders = When archiving you can set ArchiverFS to remove any empty folders from the live file system.

Archive\Delete Temp Files = If 'Archive Temp Files' is selected then temporary files will be considered for archiving just like any other file, if 'Delete Old Temp Files' is selected then temporary files over 30 days old in the live file system will be deleted.

Temporary-file handling can include files with names beginning with '~', files with a .tmp or .temp file extension, as well as those with the 'Temporary' file attribute set. Review this option carefully: deleting old temporary files is a deletion operation in the live file system, not simply another way to archive them.

Utilisation = This setting directly affects how much load ArchiverFS will place on your file system in terms of disk IO and network traffic. By default it will run as fast as the network and storage allows.

Folder Keywords = If this option is enabled and one or more keyword phrases are entered, then the contents (including all subdirectories and files) of any folders with the phrase in their name will be migrated irrespective of file age or type. While you should test for yourself, we have found that

using "_Archived" as the keyword and adding it to the end of folder names works well. Keywords are case sensitive, partial and must match exactly

Once you have set the required options click on 'next' to move to the scheduling option page.

You can set each job to run once, on specific days of the week, or even on specific days of each month. Once you have set the schedule for this job click 'Next'.

The final options screen lets you enter email addresses that you would like job completion status notifications sent to. For each address you would like a notification sent to enter the SMTP address and pick the level of notification you want for that address, then click 'Add' to add it to the list of addresses.

Once done click 'Next' to move to the final screen.

The final screen just requires you to enter a name for the job and then click the 'Create Job' button to create the job.

Testing Your First Job

Run Pre-Scan first against a small test folder using the action you intend to use in production.

Include files with representative permissions, filenames and lengths of path. Once the run has completed review the resulting log file for any errors or warnings.

You should also check the following: Open the resulting links as ordinary users, including someone who should not have access. Compare the archived files' content, permissions and timestamps with the originals. Finally, modify one of the archived files and save the changes.

Now run the maintenance job to test file return before expanding the job to actual live folders by adding then as additional sources to the archive job (right click -> edit job -> add actual live sources).

Keep the first production run manageable and check the logs afterwards. If you use DFS replication, follow the earlier advice about allowing replication to catch up before maintenance runs.

You may wish to start with a longer threshold date and gradually work towards your desired threshold. For example, if you are targeting a 3 year threshold date then run the job with 10 years as the threshold date, then drop to 5 years for the second run and finally 3 years for a third run.

Link Types

ArchiverFS supports two main types of links, and which one you choose is completely up to you.

Symbolic links — These keep the original filename and allow compatible applications to open the file through its original path. File Explorer may identify the item as a symbolic link. The archive share must remain available, and the client or application identity must be able to access it.

Some non-Windows storage devices handle links or Windows metadata differently. Test the complete workflow on your chosen source, destination and clients rather than relying on the ability to copy a file alone.

MacOS clients: Test the exact SMB server, macOS version and applications in use. Windows .lnk shortcuts can be access by some Mac OS versions with the help of a third-party utility. Any third party utility is a separate product, and its current compatibility needs to be checked.

Windows shortcuts — These are .lnk files. They give users an explicit way to open archived content but are not transparent replacements for the original document. They can break references between documents, such as linked Excel workbooks.

Shortcuts can be useful where seamless symbolic links are unsuitable. Choose per-file shortcuts when users need an individual entry for each document, or per-folder shortcuts when browsing the archive folder by users is the preferred workflow.

Exclusions

Exclusions can be added either globally, or to a specific job. Two types of exclusion are supported:

Directory path exclusions.

If exclusions are added without a trailing "*.*" then they will be treated as a directory exclusion. Directory exclusions check the entered characters against all directory paths. If the entered characters are found in any directory path then that path will be excluded (including any subfolders and files). For example..

1) \\domain.local\dfs\share\subfolder\subfile.ext
2) \\domain.local\dfs\share2\subfolder\subfile2.ext

- If you enter "\share\" as an exclusion then path 1 will be excluded, but not path 2.
- If you enter "\\domain.local\dfs\" then both will be excluded.
- If you use a short phrase without separators, it may match more paths than you intended. Use the more precise path examples above and test the exclusion before running a move job.
- If you enter "\dfs\share" both will be excluded, but if you enter "\dfs\share\" only path 1 will be excluded.

File based exclusions.

If a full directory path is entered in conjunction with "*.*" on the end of it, then all files in the specific folder will be excluded, but subfolders will still be processed.

E.G. "\\domain.local\dfs\share\subfolder*.* "

In this instance all of the files in the \\domain.local\dfs\share\subfolder\ directory will be excluded, but any subfolders and their contents will still be included.

Setting up a New Graveyard Job

Archive jobs move files from live storage to the archive. Graveyard jobs move qualifying archived files to a separate destination, removing their links from the live file system. They select content through the archive jobs you select, rather than through a new list of live source folders.

For example, you might archive files whose selected timestamp is more than three years old and graveyard files whose selected timestamp is more than ten years old. These periods are examples; use the retention policy agreed for your content.

Remember that file timestamps are retained when files move. A ten-year date criteria does not necessarily mean that a file has spent ten years in the archive, it means it hasn't been used by the criteria you have chosen for 10 years (last accessed or last modified date).

This gives an immediate gain in free space at deployment, whilst providing a managed lifecycle for files in the long term. Graveyard jobs are set up in the same way as an Archive job with the following exceptions..

1) You do not choose physical folders to remove content from; you choose a job to attach the graveyard job to.
2) The location you choose to move files to would normally be backed up to long term storage media after completion of the job, and then deleted.

Once a graveyard run completes, check its log and back up the resulting files. Confirm that you can restore them before carrying out any approved deletion. Keep any content that must be retained outside the disposal scope.

Log Settings

ArchiverFS jobs support three different logging levels.

Summary = Minimal information. Records a job starting, any errors encountered and completion statistics.

Simple = Same as Summary logging, but also lists files that have been moved.

Detailed = Includes everything in Summary and Simple logs, plus a lot of other information. Detailed logs are great for troubleshooting but they can be 100's of MB's in size.

Use Detailed logging for a limited troubleshooting period and make sure there is space for the resulting logs. Keep logs long enough to investigate failures, and preserve any required evidence before the maintenance job removes old logs as part of it's housekeeping routines.

Where available, Run Summary gives a convenient overview of job history. Read it alongside the full log: a completed run may still contain errors or skipped files. Recent builds use UTF-16 log files, which is useful to know if you import logs into another system.

Notifications

Before ArchiverFS can send any job status notifications you need to tell it which SMTP server it should connect too to send them. This is done in the 'Notification Settings' screen which can be accessed through the toolbar on the main interface window.

You can also enter a default notification address using this screen. If it is entered then this address will be pre-populated when setting up a new job using the new job wizard.

If mail is not delivered, check the SMTP server, port, relay permissions and the authentication and encryption options specified. A successful archive run does not guarantee that its notification was delivered, the notification is sent after the run completes and records it's final status.

Maintenance Job

When enabled (with the default recommended settings) the maintenance job performs several important functions, the main ones being:

- Synchronises permissions. NTFS file and folder permissions are synchronised between the live file system and migrated files.

- As it checks each archived file it checks if that archived file's symlink\shortcut exists in the live system in its original location. If it does not then the archived files symlink\shortcut is added to a search list for the later live system scan.

- Clears log files. Removes any files older than the setting specified in the log settings screen.

- Returns recently used archived Files. Identifies any files that have been used by users and returns them to the live file system.

- Scans the live file system. ArchiverFS now scans the live file system looking for any missing symlinks\shortcuts that were identified earlier. If one is found then the archived file can be relocated to the correct folder in the archive folder structure to reflect its symlink's\shortcut's new place in the live file system, and then it re-creates the symlink\shortcut to point to the archived file in its new location.

- If any missing links can't be found then they will be re-created (default settings)

The behaviour of the maintenance job can be altered via the maintenance job option screen, accessed through the toolbar on the main interface window.

Enabled — Enables or disables maintenance. The additional scheduling controls can run maintenance on selected days or monthly. Choose a schedule that allows maintenance to complete and review its log after the first runs.

Synchronise Permissions — Keeps the live and archive permission models aligned. Permission synchronisation also depends on the setting for each associated archive job.

Synchronise Shortcut\Stub Times — Checks the Created and Modified timestamps of the shortcuts and stubs left in the live file system and updates them where needed to match the archived files.

Return Files — Returns archived files that meet the configured return policy. The supplied configuration offers returning modified files, with an option to use access instead. If you use Last Accessed, check that backup, antivirus, indexing or other software is not changing the dates and causing unwanted returns.

'Scan Live File System' = As the maintenance job works it builds up a list of archived files whose links cannot be found. With this option enabled the Maintenance Job will scan all of the live locations it knows about (all targets that have been added to all jobs) for the missing files. If it finds the missing links it will re-locate the associated files so their location matches their Shortcuts\Links new live location. If it does not find them then they will be re-created (default behaviour)

'Rename Archived Files if Their Link Has Been Renamed' = If the Maintenance job finds a missing link during it's scan of the live file system it will check to see if the link has been renamed. If it has been renamed then it will rename the archived file to match the new name of the link.

'Delete Archived Files Whose Links Cannot be Located' = If enabled then instead of re-creating missing links if they cannot be found, the files with missing links will be deleted instead. If you wish you can specify a 'recycle' folder where files will be placed instead of being deleted.

Notification Email Address = If this option is enabled and an address specified, then that address will receive notification emails relating to the maintenance job each time it runs, along with a copy of the log file for each run.

Understanding Missing Links

With the normal recreation behaviour, deleting a link from the live file system does not necessarily delete the archived file: maintenance may recreate the missing link. If you enable Delete Archived Files Whose Links Cannot Be Located, you change that policy to removing the archived file instead.

Be particularly careful during server moves, outages or DFS replication delays. Pause the affected jobs until both sides are visible and correctly mapped. Do not use either deletion option simply to tidy up an unexplained set of broken links.

Opening an archived file and returning it are separate operations. Users can work directly with the archive while the file remains there; maintenance later returns files that match the chosen policy. Check the recently-used interval shown in your build and test the trigger with known access and modification dates.

File Types

There are two distinct ways that file extensions can be specified for archiving, at the Global level and at the Job level.

Global Level: This is configured through the 'File Types' screen, accessed through the toolbar on the main interface window. Anything specified in the global screen will apply to all jobs, unless:

- The 'Inherit Global Extension List' option has been deselected at the job level
- A job-specific entry overrides the corresponding entry in the global list.

Job level: You can add file extensions to an individual job. With global inheritance enabled, these entries override matching global entries. If you untick Inherit Global Extension List, the job uses its own list, together with the options below that control unlisted and extensionless files.

At the job level you also have two more options to configure should you wish to.

- Archive All File Types Not In File Types List: With this option unticked a job will not archive files by default. It will only archive them if there is a checked entry for them in the file types list. If this option is ticked then the job will archive files by default, unless there is an unchecked entry for them in the file types list.
- Archive Files With No Extension: With this option selected a job will archive files even if they do not have a file extension.

Executable File Types: Be careful when selecting files such as .exe and .com. A file's timestamps may not reliably show whether an executable is still needed, so age alone is a poor reason to move it. If you intend to archive executables, select the required extensions explicitly and test the effective rules with Pre-Scan. Check the result of Archive All File Types Not In File Types List rather than assuming every executable is excluded.

The global extension list can be exported to a file and imported again. Keep a copy before making a large change, and check how that change affects jobs that inherit the global list.

Pre/Post Job Tasks

It will often be desirable to run some kind of task either before a job starts, or after it completes.

To facilitate this ArchiverFS incorporates the ability to schedule such tasks. This is done using the Pre/Post Job Tasks feature which is accessed through the toolbar on the main interface window.

Any command line task can be started. If you need to execute several tasks, or a task includes complicated command switches (like a Windows Backup job) and it does not execute properly, try placing the command into a .bat file and then executing the .bat file with the pre/post job function.

Test the task using the identity and access available to the service. Use full paths and capture its output so failures can be investigated. If the archive job depends on a task finishing successfully, confirm its completion and failure behaviour with MLtek before relying on it as a prerequisite.

Compressing Your Second Line Storage

ArchiverFS supports file and folder compression through standard NTFS File and folder compression.

It is recommended that if you do want to compress your migrated files then you compress the entire destination by compressing the root folder itself (the actual folder that is shared out to create the destination for a job).

On the server hosting the share, right-click the destination root folder, choose Properties > Advanced, and select Compress contents to save disk space. Apply the change to existing subfolders and files if that is intended. On a large archive this can take time, so plan it as a storage operation.

If you have compressed your old files then new folders and files will be automatically compressed as they are migrated to it.

Compression uses processing time as files are written and read. The overall effect depends on the content and storage; already compressed formats may save little space. Test representative files rather than assuming a fixed saving or performance penalty.

If you decide that you wish to uncompress your old files after it has already been compressed then simply uncompress the root folder as above. This will uncompress all the subfolders and files underneath, and new folders and files will not be compressed when they are migrated.

Deduplicating Your Second Line Storage

Windows Data Deduplication is compatible with ArchiverFS for both live sources and archive destinations. Check support for your Windows release, volume and backup product before enabling it.

Enabling Long Path Support

If you are running Windows 10\Server 2016 or later then you can consider enabling long path support.

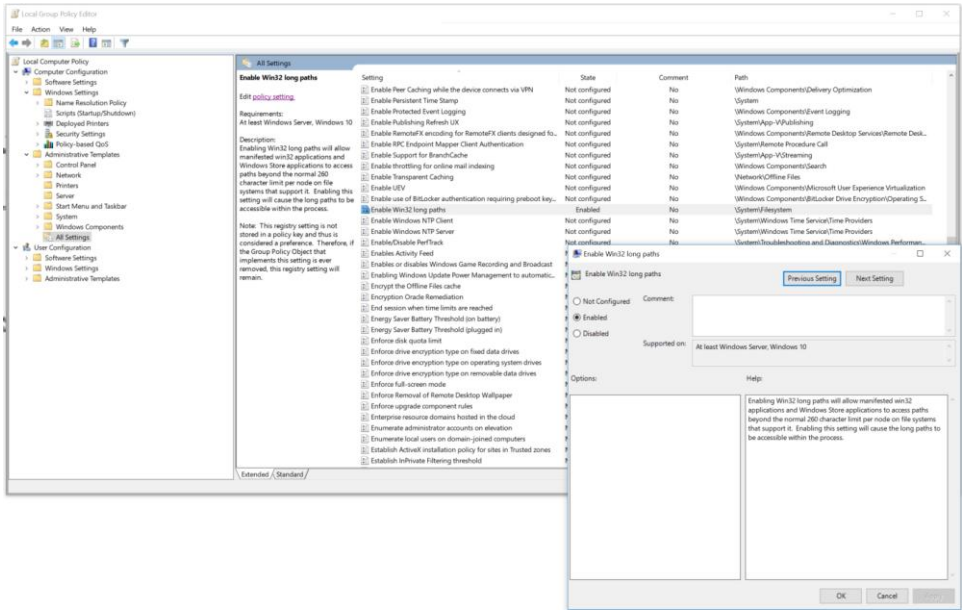
Long-path support needs to work through the whole file access pathway. The ArchiverFS service may be able to process a path that a user's file manager, .lnk shortcut or application cannot open. Test the actual clients and applications before changing the setting.

For this reason we don't currently recommend enabling long path support unless you have been advised to do so, or you are experimenting in a non-live environment.

Support must be enabled at both operating-system and application level. Check the instructions for your installed build before changing the application settings.

We have an article on Long Path Support on our website here:
https://www.mlteksoftware.com/Products/ArchiverFS/FAQ/ArchiverFS_Long_Paths.aspx

Enabling OS Support — Open Group Policy Editor and go to Computer Configuration > Administrative Templates > System > Filesystem. Enable the policy called Enable Win32 Long Paths.



Enabling Application Support — See the 'LongPathEnabled' value in the registry under HKEY_LOCAL_MACHINE\SOFTWARE\MLtek\AMV3.

We strongly recommend contacting MLtek Support for advice before enabling any long path support functionality..

Plan any required service or machine restart as part of the change. Windows processes may cache their long-path setting, so test again after the change has taken effect.

Dealing With Long Paths

A common cause of path problems is the difference between a mapped drive and its full UNC path. A file that users reach through a short drive-letter path may have a much longer network name, and the archive layout can add more characters.

Windows can support extended-length paths in applications that opt in, but that does not guarantee that every shortcut or application will open them. Keep source and destination roots short where possible and test the final archive paths.

Directory symbolic links can provide a shorter route to deeply nested folders. The following example explains the idea; ask MLtek to review the job layout before using it on a live tree.

Imagine you had a file system that contained several directories really deep down in the file system, and those directories contained files over the Windows 260 character limit. With long path support disabled (the default setting) Archive FS will just skip these files. What's needed is a way to shorten the path to the files without having to move the files or change the directory structure they sit in.

This is where directory based Symbolic links come in. The trick is to create a folder fairly high up in the file system which will just contain a few choice directory based symbolic links. These links point

to places further down the file system and effectively shorten the path that ArchiverFS has to travel to reach them.

For example, suppose Share1 on Server1 contains the following deeply nested document path:

\\Server1\Share1\Sales Departments\Regional Figures\North East\Pre Office Move 2006\Regional Subsectors\Central\Jills Annual Reports For Review\Draft Documents Prepared Onsite 2002-08-10\Cost Center Breakdown Of All Expenses By Resource Type\Spreadsheet RAW Data\New York 2002 Q1\Sector Breakdown.xlsx

This path is over the traditional Windows limit. A directory symbolic link can provide a shorter route to part of the tree. On the file server, create the link using local paths for the directory link, as in this example. Adjust D:\Data to the actual folder behind Share1:

mklink /D "D:\Data\Shortcuts\Old Sales Data" "D:\Data\Sales Departments\Regional Figures\North East\Pre Office Move 2006\Regional Subsectors\Central"

The following path will now be available:

\\Server1\Share1\Shortcuts\Old Sales Data\Jills Annual Reports For Review\Draft Documents Prepared Onsite 2002-08-10\Cost Center Breakdown Of All Expenses By Resource Type\Spreadsheet RAW Data\New York 2002 Q1\Sector Breakdown.xlsx

The shorter UNC path is under 260 characters in this example. A support-approved job layout can use that route to reach the files without moving the underlying folders.

Reporting

The reporting function can be opened by clicking on ‘Reports’ on the main screen. It provides several graphical reports which provide an overview of archiving activity.

Before information relating to a job’s archiving activity can be displayed the ‘Collect statistics for reporting’ option for that job must be enabled. This can be done either when setting up a new job, or using the ‘Edit Job’ screen.

If reports are empty, check that statistics collection was enabled for the period in question. If you receive an error message when opening the reporting function, check the specific Crystal Reports runtime version required by the installed ArchiverFS build is installed.

Backing Up Your Configuration

It is easiest to think of the system as two parts that need to be protected: the application configuration and the file systems that contain the live and archived content.

1) The configuration of the software e.g. the jobs that have been setup, the maintenance job settings, notification settings, etc.

   The application configuration is held in the AMDBv3 SQL database. Use a SQL-aware backup or another database-consistent method agreed with your database administrator. A simple copy of files while SQL is running is not proof that the database can be restored.

   Back up the configuration after changing jobs or global settings, and keep a regular backup schedule appropriate to any reporting data. Record custom SQL connections and relevant application settings so that a replacement host can connect to the correct database.

2) Information about the files that have been moved e.g. where they were migrated from, when they were last accessed, etc.

   The file system holds the archived content and the information used to relate it to the live locations. Existing links can still work while the ArchiverFS machine is shut down, provided the source and archive servers remain accessible. Scheduled archiving and maintenance will stop while the service is offline.

Back up both the live tree and the archive, including permissions, timestamps and the links themselves. Check whether your backup product preserves symbolic links or follows their targets: restoring a real file where a link should be can produce a different result. Include graveyard content in the backup plan for as long as it must be retained.

Additional Tools

ArchiverFS includes several standalone tools for specific operations. You can find them in the ArchiverFS group on the Start menu.

Return Archived Files Utility — Manually returns archived files to the live file system. Use it when files must return before the next maintenance run. Check the available live space and test a small selection first. Supporting builds offer extension filters, including files with no extension.

Migrate Archived Files Utility — Moves a job's archived files to a new destination, updates the links in the live file system and, once complete, updates the job's archive destination. Use this when relocating an archive rather than moving its folders in Explorer or changing only the destination field. Build 3.479930 adds a Replicate NTFS Permissions option; choose it deliberately and check the resulting permissions.

'Reset Permissions Tool' = Scans the selected log file and identifies folders and files that ArchiverFS couldn't access due to insufficient permissions. It then attempts to add the selected account (normally the ArchiverFS account) to the ACL of these items, thereby granting access to ArchiverFS.

Getting Help

If you have active Assisted Support and Update entitlement in place or are a new user wanting to evaluate ArchiverFS, please feel free to contact Support@MLtekSoftware.com with any questions you may have.