



ArchiverFS 3.4x User Guide

Contents:

Introduction to ArchiverFS.....	1
Basic Concepts	2
Versions.....	3
Upgrading From a Previous Version	3
Clusters and DFS Shares.....	3
Choosing your Archive Storage.	4
NAS devices, NFS shares, CIFS shares and non-Windows storage.....	5
Installing – Service Account	6
Installing – Prerequisites.....	7
Installing - Running The Install.....	7
Installing - Configuring Your Machines	8
Setting up a New Archive Job	8
Link Types.....	10
Exclusions.....	11
Directory path exclusions.....	11
File based exclusions.....	11
Setting up a New Graveyard Job.....	11
Log Settings	12
Notifications.....	12
Maintenance Job.....	12
File Types.....	13
Pre/Post Job Tasks	14
Compressing Your Second Line Storage.....	14
Enabling Long Path Support.....	15
Dealing With Long Paths	16
Reporting	17
Backing Up Your Configuration	17
Additional Tools	17

Introduction to ArchiverFS

ArchiverFS is unique in that it does not use any sort of 'database' or 'proprietary data store' to store archived files, file metadata or even record the location of archived files. Instead it uses flat files (NTFS\REFS) to store migrated files while maintaining the directory structures, NTFS permissions and file attributes. This makes it incredibly fast, easy to backup, highly cost effective to deploy, very reliable and almost infinitely scalable. It's a great way to move old files from a front line SAN, NAS or file server to second line storage.

The largest file system we know of that uses ArchiverFS is just over 3.5 Peta Bytes (3500TB +), and this is all managed with a single copy running a Datacenter licence with 4vCPU's and 6GB of RAM.

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.

- A) True Symbolic\Hard links to archived files provide seamless access for end users running Windows 7 or later. Other types of link are available to ensure compatibility with older OS's and Mac machines.
- B) Support for every file system character set supported by Unicode, which is most of the main languages in the world including for example English, Chinese (Simplified and Traditional), Hindi, Arabic and many others.
- 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) Practically unlimited file storage with no degradation in performance regardless of how many files are being held in second line storage. ArchiverFS's maximum archive size is limited by nothing more than the available second line disk space. Second line storage sizes of several hundred terabytes are common. Database centred systems are not able to reach anything like this capacity without involving huge costs and complexity.
- E) Massively reduced deployment and ongoing support costs. It's a well known fact that raw disk space is cheaper than database space by several orders of magnitude. A 100GB database centred product can cost 10's of thousands of USD to deploy. ArchiverFS on the other hand costs a fraction of that.
- F) Brick level security. Databases use a different permission structures to NTFS volumes. This means that it is almost impossible for a database centred archiving system to maintain your NTFS file and folder permissions on files that have been migrated to long term storage. Needless to say, ArchiverFS doesn't suffer from this problem. NTFS permissions are maintained through a files entire lifecycle.

Basic Concepts

Before installing ArchiverFS we recommend that you read and follow the Quick Setup Guide.

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 location, they do not all have to migrate to the same share. Indeed, when you set up each job the folder you use for its destination must be empty.

Once files have remained unused in second line storage for the period of time you specify, they can be removed with a graveyard job. This will remove any old files from the selected Archive jobs destination, and transfer them to the location you choose. They can then be backed up using your existing backup solution and deleted.

At all times all files NTFS permissions, file attributes and Created/Modified/Last Accessed dates are preserved. Even the directory structure is preserved from start to finish.

Second line storage space is typically made up of less expensive storage hardware. Typically this is several years old, and whilst it may have significant capacity, 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. Files that are needed by users are returned to the live file system
3. Files that remain unused for the period of time specified after being migrated are graveyarded to a chosen location, backed up and then deleted.

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**	Yes	Yes	Yes	Yes
Includes 12 months of assisted support	No	Yes	Yes	Yes
Directory structure retained	Yes	Yes	Yes	Yes
NTFS permissions retained	Yes	Yes	Yes	Yes
File attributes retained	Yes	Yes	Yes	Yes
Pre + Post job tasks	Yes	Yes	Yes	Yes
Various stub types, including Hard 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 are also available that provide an unlimited number of jobs, but include a quota for the amount of data that can be archived annually and must be renewed annually.

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.

Clusters and DFS Shares

ArchiverFS offers full support for both Windows Server clustered shares, and Windows DFS shares.

However, and this is very important... If you move content to/from any kind of clustered storage you must NOT use the UNC path to an individual node in that cluster.

For example, say you have two Windows servers, Server1 and Server2 and they are node in a DFS pair. You may have [\\Server1\Data](#) and [\\Server2\Data](#), but you must not use these paths. You must specify the true DFS share path e.g. [\\DFSShares\Data](#) for example.

The same principal must also be used with Windows clusters.

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.

This is normally accomplished by using NFS or iSCSI to map the storage to a Windows VM on VMWare or HyperV as a local drive, then using that Windows VM to format the storage and share it to the network as a network share.

Doing this ensures that even if the storage is non-Windows (e.g. a typical NAS) it is a Windows OS that has control of it and it will be using a true Windows NTFS file system.

This is important because even non-Windows OS's that advertise 100% Windows compatible NTFS are rarely if ever 100% Windows NTFS compatible. They are typically more like 99.99% compatible, and if you are moving 100's of thousands of files to it then you will most likely encounter some kind of issue (typically permission related) where something cannot be replicated on them exactly.

- 2) Choose storage that is fast enough.

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

It doesn't have to be fast... but it has to be 'fast enough'. You will have multiple users browsing the file system and if you use Symlinks then each time a user opens a folder on the live system any Symlinks in it cause a read against the secondary\archive storage.

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.

NAS devices, NFS shares, CIFS shares and non-Windows storage

If you would like to migrate content to (or from) a location that is not hosted on a Windows based server (or it is a Windows based server but in a separate domain to your live servers) then you will need to do some additional planning and testing. By default ArchiverFS will try to maintain the permissions on any files in the live file system when the files are moved. If you are archiving from a Windows server to another Windows server and they are both in the same domain then all is well and good.

But, if for example you wish to move files from a Windows server to a NAS box with an NFS or CIFS share and that NAS box is not AD integrated then you will have issues. ArchiverFS will not be able to maintain the permissions on migrated files because the destination device will not be able to deal with Domain\AD user and group objects with respect of file security.

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

- 1) Log into the machine you wish to install ArchiverFS on with the user account that you wish to run ArchiverFS under.
- 2) Open Windows Explorer and go to the location you wish to move old files too.
- 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 successfully set these domain based permissions then ArchiverFS should be able to maintain the permissions absolutely fine.
- 6) If you cannot set these permissions then the device does not support AD\Domain integration and if you use it as storage for your old files then ArchiverFS will not be able to retain the permissions on them. What's more, unless you tick the 'Best Efforts Security' option when setting up any jobs that will use this device the jobs may fail.

Different storage architectures can directly affect the functionality of ArchiverFS. Below is a table which should help highlight the effect different architectures can have.

	DESTINATION	Windows 2003	Windows Server 2008 ->2019	AD Integrated NAS	Non-AD Integrated NAS/CIFS/NFS
SOURCE					
Windows 2003		Permissions = OK Attributes = OK Links = Soft Only	Permissions = OK Attributes = OK Links = Soft Only	Permissions = Test Attributes = Test Links = Soft Only	Permissions = No Attributes = Ok Links = Soft Only
Windows Server 2008 ->2019		Permissions = OK Attributes = OK Links = Hard & Soft	Permissions = OK Attributes = OK Links = Hard & Soft	Permissions = Test Attributes = Test Links = Probably Hard & Soft, Test	Permissions = No Attributes = Ok Links = Probably Hard & Soft, Test
Non-AD Integrated NAS/CIFS/NFS		Permissions = No Attributes = Test Links = Test	Permissions = No Attributes = Test Links = Test	Permissions = No Attributes = Test Links = Test	Permissions = No Attributes = Test Links = Test
AD Integrated NAS		Permissions = Test Attributes = Test Links = Test	Permissions = Test Attributes = Test Links = Test	Permissions = Test Attributes = Test Links = Test	Permissions = No Attributes = Test Links = Test

Installing – 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, any machines it will archive content from and any machines it will move content too.

The easiest way to accomplish this is to make the account a member of the Domain Administrators group.

If the account you are using to run the ArchiverFS Service is not a member of the Domain Administrators group then in addition to requiring suitable NTFS rights to be set up so it can access the required shares with Full Control permissions, it needs to have the two extra rights granted to it.

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.

The relevant setting can be found under Local Policies -> User Rights Assignment. Edit the Create Symbolic Links entry and add the account you created in step 3. Alternatively you can set the settings using the Default Domain Group policy in Active Directory.

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

The relevant setting can be found under Local Policies -> User Rights Assignment. Edit the Manage Auditing and Security Log entry and add the ArchiverFS account that you have set up. Alternatively you can set the settings using Default Domain Group policy in Active directory.

Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Local Policies -> Manage Auditing and Security Log.

Installing - 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 v4.6.2
- 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.

1) Installing - Archive Share

Set up an empty share on your network in the location that you would like to store your old files. The share should ideally be hosted on a computer running on Windows Server 2008 R2\Windows 7 or later. If this is not possible then it should be hosted on a machine running at least Windows 2003 Server.

Set the share permissions on the share to allow 'Full Control' to 'Administrators', and grant 'Change' and 'Read' to the 'Domain Users' group. Then ensure the NTFS permissions will allow normal users 'modify' rights and the ArchiverFS Service Account created previously 'Full Control'.

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 with those in the live file system, thus ensuring content stays secure.

Installing - Running The Install

Once you have carried out the above steps, you are ready to install ArchiverFS.

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.

Installing - Configuring Your Machines

If you intend to use Hard Links (aka Symbolic Links) then you need to enable support for them on your workstations, the machine ArchiverFS is installed on and any Web Servers you intend to move content from.

Even though every Windows OS since Vista supports them, the Hard Link functionality is disabled by default.

To enable the functionality you need to set the following registry keys on each machine specified above.

To ease deployment you can copy and paste the text below into a .reg file which can be imported at logon by logon scripts, or the settings can be distributed via group policies.

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
```

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 set up 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.

Once you have set up the folder 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 too. 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 allows you to set various options that control how the job will run. Most of the options are self explanatory, the most commonly used are listed below.

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

Copy (No Links) -> Copy files that match the criteria specified. Makes no changes to your live file system, useful for testing purposes.

Move (No Links) -> Move files that match the criteria specified. Removes files from your live file system and transfers them to the chosen destination. Leaves no links behind. Migrated files can be backed up and the deleted.

Move (Per Folder .Lnk Shortcuts) -> Move files that match the criteria specified. Removes files from your live file system and transfers them to the chosen destination. Leaves behind Windows shortcuts in each folder that files have been migrated from. Users can click on this shortcut and then browse through the migrated files. Any files that are accessed by users can be automatically returned to the live file system by the overnight maintenance job. Migrated files must be removed with a graveyard job.

Move (Per File .Lnk Shortcuts) -> Move files that match the criteria specified. Removes files from your live file system and transfers them to the chosen destination. Leaves behind a Windows shortcut for each file that is moved. Users can click on the shortcut to open the file. Any files that are used by users can be automatically returned to the live file system by the overnight maintenance job. Migrated files must be removed with a graveyard job.

Move (Per File Hard\Symbolic Links) -> Move files that match the criteria specified. Removes files from your live file system and transfers them to the chosen destination. Leaves behind a Windows Hard\Symbolic link for each file that is moved. Users can click on the link to open the file. Any files that are used by users can be automatically returned to the live file system by the overnight maintenance job. Symbolic links behave exactly like the original file, and have exactly the same name as the original file. Hard\Symbolic links are only available in Windows Vista, 7, 10, 2008, 2008R2, 2012, 2016 and 2019. Files must be removed with a graveyard job.

Archive Files That Have Not Been Used In = Specifies how long it must be since a file was used before it will be moved. Typically this will be around 3 years. However each organisation is different and careful thought should be given to this setting.

Pre-Scan = If a job runs with this option enabled then it will run in 'Pre-Scan' mode. The job will function as normal except that no files will actually be moved to archive storage and no links will be created in the live file system. This option is useful for evaluating how much data is available to be archived with the jobs current settings so that decisions on purchasing the correct licensing and archive storage can be made. Once the job has completed you can review the end of the log file to see information about how much data would have been archived if the job had been run in normal mode. Jobs run in Pre-Scan mode using a Trial\Free licence are not limited by the 100 file Trial\Free limit.

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.

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.

Utilisation = This setting directly affects how much load ArchiverFS will place on your file system in terms of disk IO and network traffic.

Folder Keywords = If this option is enabled and a keyword phrase is entered, then the contents of any folders with the phrase in their name will be migrated irrespective of file age or type. Keyword phrases are case sensitive. 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.

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.

Link Types

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

Hard\Symbolic links = These are completely seamless links that can be left in place of old files which behave exactly like the original file (Windows does display them as a '.Symlink' file in file type column in Windows Explorer, but the icon, file name, path, etc are all maintained). Only machines running Windows Vista, 7, 10, 2008, 2008R2, 2012, 2016 and 2019 or later can access these links so you need to consider the OS's in use on your workstations and/or servers before using Symbolic\Hard links.

Some storage devices that aren't 100% NTFS compatible can also cause issues with Hard\Symbolic Links. Please contact us for assistance with this if you are planning on using an OS other than Windows to present your second tier storage to the network.

Windows Shortcuts = These are standard .lnk Windows shortcuts, and they are compatible with just about everything, you can even get them working with Mac's. The downside is that they are not seamless and will for example break links between documents like Excel spreadsheets. They can also have some issues with long paths over the 255 character limit even with long path support enabled.

On the plus side they typically do not run into the same storage related issues Hard\Symbolic links can, so if you have to present your second line storage to the network using a non-Windows OS they might be a good choice.

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 enter ‘subfile’ then both will be excluded
- 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 to second tier storage. Graveyard jobs move files out of the second line, to a location you specify so they can be backed up and then deleted. Graveyard jobs remove files that haven’t been accessed by users in the period of time you specify. This is based on each files last accessed date.

A typical configuration will use Archive jobs to remove files from the live file system that have not been used by users in 3 years, and place them into one or more sets of archive folders. A graveyard job will then be used to remove files that have not been accessed by users in 10 years.

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.

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.

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 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.

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.
- 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. If ArchiverFS cannot locate a Shortcut or Hard Link for an archived migrated file when synchronising permissions, then it can scan the live file system for it. If it is found then the archived file can be relocated to the correct folder in the archive folder structure to reflect their new place in the live file system and the link re-created to point to the archived file in its new location.
- If any missing links can't be found then they will either be re-created.

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

'Enabled' = Enables/Disables the maintenance job and allows you to set the time it will run each day.

'Synchronise Permissions' = Enabled\Disables the synchronising of permissions between the live file systems and your second line storage.

‘Synchronise Shortcut\Stub Times’ = With this option enabled ArchiverFS will verify that the Created\Modified times on Shortcuts and Stubs left behind in place of migrated files match. If they don’t then the times on the Shortcuts\Stubs are updated.

‘Return Files’ = Selecting this option will tell the Maintenance Job to return migrated files that users have modified back to the live file system. This behaviour can be modified to return files that users have accessed by selecting the relevant option. Before doing this you must ensure that other products (Antivirus, Backup, etc) are not changing the last accessed date on your files.

‘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 know 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.

‘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.

‘Delete shortcuts in the live file system that point to non-existent archived files’ = If you are archiving to an Amazon Storage Gateway and then migrating to Amazon Glacier, this option should be enabled.

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.

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 level has an entry for a specific file type exists that overrides an entry in global list.

Job level: File extensions can be added at the job level, and if added they will override the equivalent entry at the global level if inheritance is enabled. If you untick the ‘Inherit Global Extension List’ option for a job then the job will now only process files that have a corresponding entry in it’s own job specific extension list.

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 a 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: By default ArchiverFS will NOT migrate executable file types like .exe and .com files. There are very few situations where it makes sense to move these file types from a network share. It's normally safe to assume that if an executable file type is hosted on a network drive it is normally there deliberately and is probably being actively used. What's more, the timestamps on an executable files are rarely if ever updated as users access them so it is extremely difficult to determine if it is safe to move executable file types. If you do want a job to migrate these files then you can do this by un-ticking the 'Archive All File Types' option for that job, and adding in specific entries for the executable file types to the jobs file types 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 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.

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).

Right click on the root folder through windows explorer and click on 'Properties'. Tick the 'Compressed' attribute and then click ok. Depending on the size of your old files this might take some time. It is recommended that you perform this action on the server hosting the share, not from another machine over the network.

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

There is a small performance penalty for the compression as files need to be compressed and uncompressed as they are moved but this is minimal.

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.

Enabling Long Path Support

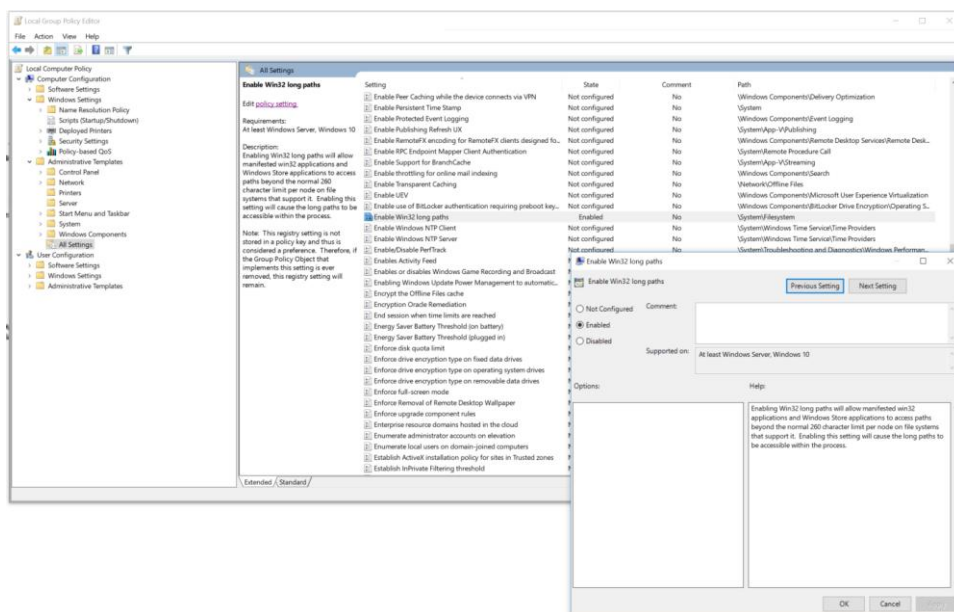
If you are running Windows 10\Server 2016 or later then you can consider enabling long path support.

Be aware that at the time of writing Windows Explorer and many other applications don't currently support long paths, even though the OS itself does. What's more, shortcuts (.lnk files) don't work properly with long paths.

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

If you do wish to enable it then support needs to be enabled in two places, first within the local group policy of the machine ArchiverFS is installed on and secondly within an ArchiverFS Registry setting.

Enabling OS Support = Open 'GPEdit' as an administrator, expand 'Computer Settings', 'Administrative Templates' and click on 'All Settings'. Edit the 'Enable Win32 Long Paths' settings and change it to 'Enabled'.



Enabling Application Support = Open 'Regedit' as an administrator and navigate to "Computer\HKEY_LOCAL_MACHINE\SOFTWARE\MLtek\AMV3".

Once there find a value called 'LongPathEnabled' and change the value from 'false' to 'true'.

Finally, once both steps above have been completed, you will need to restart the machine.

Dealing With Long Paths

As discussed in the previous section, although it is possible to enable long path support in Windows support for long paths amongst many applications including Windows Explorer is severely lacking.

With long path support enabled ArchiverFS can access paths up to 32k characters long, but unless you are using a file manager that also supports long paths like Total Commander or Q-Dir shortcuts and Symbolic links to long path files won't work properly. They will often display an unresolved icon and won't open when clicked on.

Luckily there is a way round this, the use of directory based Symbolic links.

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.

Let's imagine you had a share called 'Share2' on a server called 'Server1' and it contained a lot of paths that were over the windows 260 character limit. Let's go a step further, let's look at an example path to a document.

```
\\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
```

That's over 300 characters and Windows Explorer would refuse to open it. But, if we run the following from an administrative command prompt...

```
MKLINK /D "\\Server1\Share1\Shortcuts\Old Sales Data" "\\Server1\Share1\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 new path is just over 230 characters long and as long as the "[\\Server1\Share1\Shortcuts](#)" directory is within the scope of an archive job these and any other long path files will be discovered and archived as appropriate.

With the creation of a small number of directory based Symbolic links in a shortcut folder you will be able to easily archive even the most stubborn and convoluted directory structures.

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 jobs 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.

Backing Up Your Configuration

It's easiest to think of the configuration being stored in two parts.

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

Everything is stored in the 'AMDBv3' SQL Express database, back that up (a full system backup should grab a copy of the files but it depends on the software you use) and you have everything you need to re-create the application configuration.

It's worth pointing out that you only need to make sure you have a copy of the configuration database after you make a change, you don't need to backup the configuration after running a job. As explained below, all of the information regarding which files have been moved is encoded into the file system itself.

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

All of this information is encoded into the file system. Assuming you had a live server, a machine running ArchiverFS and second line storage, you could shut down the machine running ArchiverFS and your users would have no idea. They'll still be able to access the old files via the links left behind even with the ArchiverFS machine shut down. When you back up your file system, you are backing up everything that is needed for users to access their files.

Additional Tools

ArchiverFS includes several standalone tools that integrate with the main product and let you carry out specific operations. There are available via the ArchiverFS entry on the Start Menu.

'Return Archived Files Utility' = Enables the manual return to the live file system of archived files.

'Migrate Archived Files Utility' = Provides a way to migrate a jobs archived files to a new location.

This utility will move the files, update the links in the live file system and once complete update the jobs settings with the new archive destination.

'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.