

ROBOCOPY

Robust File Copy Utility

Version 1.70

Windows NT® Resource Kits
Copyright © Microsoft Corp. 1995 - 1997

Contents

New features in this version
Introduction
Command line usage
Usage notes
The output log
The run summary

NEW FEATURES IN THIS VERSION

Restartable Copies :

Specify **/Z** on the command line and failed file copies will restart from close to the point of failure, which can save a lot of time for large files. Previous versions of Robocopy would always restart failed copies from the beginning of the file.

%Copied Progress Indications :

By default, the program now regularly displays %copied during each file copy. Specify **/NP** (No Progress) on the command line to turn this feature off.

Wait for sharename creation :

Specify **/TBD** (sharename **To Be Defined**) on the command line to request this. Robocopy will then wait for a sharename to be created by retrying on error 67. Previous versions would always quit on this error.

Force creation of short file names.

Specify **/FAT** on the command line to tell Robocopy to create destination files with names that conform to the traditional 8.3 file name format, rather than long file names. This can be useful when copying to downlevel file systems.

INTRODUCTION

Robocopy is a Win32 console-mode application designed to simplify the task of maintaining an identical copy of a directory tree in multiple locations, either on the same machine, or in separate network locations. Robocopy effectively provides file replication on demand.

Robocopy is robust. If Robocopy encounters recoverable network errors whilst scanning directories or copying files it will simply wait a while and retry the operation. You can control the time between retries and the number of retries to attempt before giving up.

Robocopy is efficient. By default, if a file exists in both the source and destination, Robocopy will only copy the file if the two files have different timestamps, or different sizes. This saves time if the source and destination are separated by a slow network link. Optionally, you may specify that copies are restartable in the event of a copy failure to save even more time if your network links are unreliable.

Robocopy is flexible. You can choose to copy a single directory, or walk a directory tree. You can specify multiple filenames and wildcards to select which source files are candidates for copying. You can exclude source files from being copied by name, path, wildcard, or file attributes. You can exclude directories from being walked by name or path. You can choose to copy only files with the Archive attribute set, and you can choose whether or not to turn off the Archive bit in the source file after copying. The program classifies files by whether or not they exist in the source directory, the destination directory, or both. In the latter case the program further classifies files by comparing time stamps and file sizes between the source file and the corresponding destination file. You have complete control over which of these file classes will be copied. You can also choose to move files rather than copy them. And you can also choose to purge (delete) destination files and directories that no longer exist in the source, and thereby maintain the destination as an exact replica of the source.

Robocopy is informative. Robocopy produces console output (which can be redirected to a disk file for later perusal) which lists the directories processed, which files are copied (and why), network errors, and incompatibilities between the source and destination directory tree structures. Optionally, you can also ask Robocopy to show estimated time of arrival of copied files, list which files are skipped (and why), and highlight differences in the structure of the source and destination trees that might merit further investigation, or require housekeeping. By default, Robocopy displays copy progress indication (% copied) for each file.

Finally, Robocopy runs fine as a scheduled job. Just configure the Schedule service to log on as a user who has appropriate access to the source and destination directories, and specify remote directories as UNC names on the scheduled Robocopy command line.

COMMAND LINE USAGE

Run Robocopy with no command line arguments for brief usage instructions:

```
-----  
ROBOCOPY v 1.70 : Robust File Copy for Windows NT : by kevina@microsoft.com  
-----
```

Started : Wed Aug 28 01:23:45 1996

Usage : ROBOCOPY source destination [file [file]...] [options]

source : Source Directory (drive:\path or \\server\share\path).
destination : Destination Dir (drive:\path or \\server\share\path).
file : File(s) to copy (names/wildcards - default is "*.*").

options :

- /S : copy Subdirectories, but not empty ones.
- /E : copy subdirectories, including Empty ones.

- /T : Timestamp all destination files, including skipped files.

- /R:n : number of Retries on failed copies - default is 1 million.
- /W:n : Wait time between retries - default is 30 seconds.
- /REG : Save /R:n and /W:n in the Registry as default settings.

- /TBD : wait for sharenames To Be Defined (retry error 67).
- /FAT : create destination files using 8.3 FAT file names only.

- /X : report all eXtra files, not just those selected.
- /V : produce Verbose output, showing skipped files.
- /L : List only - don't copy, timestamp or delete any files.

- /A+:[R][A][S][H] : add the given Attributes to copied files.
- /A-:[R][A][S][H] : remove the given Attributes from copied files.

- /XA:[R][A][S][H] : eXclude files with any of the given Attributes
- /A : copy only files with the Archive attribute set.
- /M : like /A, but remove Archive attribute from source files.

- /XF file [file]... : eXclude Files matching given names/paths/wildcards.
- /XD dirs [dirs]... : eXclude Directories matching given names/paths.

- /XC | /XN | /XO : eXclude Changed | Newer | Older files.
- /XX | /XL : eXclude eXtra | Lonely files and dirs.
- /IS : Include Same files.

- /Z : Copies files in restartable mode.
- /NP : No Progress - don't display % copied.
- /ETA : show Estimated Time of Arrival of copied files.
- /MOVE : Move files and dirs (delete from source after copying).
- /PURGE : delete dest files/dirs that no longer exist in source.

USAGE NOTES

Use within a Unix Shell

All Robocopy switches can be specified Unix-style (e.g. you can use **-ETA** instead of **/ETA**), and source and destination directory paths can be specified using the Unix-style delimiter **/** rather than the native Windows-style delimiter ****.

The only restriction is that any argument that starts with a **/** is taken to be a switch if it only contains a single **/**. Thus **//server/share/dir** and **/download/test** are treated as paths, but **/dir** is treated as a switch. This is to avoid any possible confusion between switches and single-level paths subordinate to the root of a drive. To specify such a directory as an argument, use an alternate expression for its path, such as **X:/dir** or **//server/C\$/dir**.

Use with Windows 95 and Windows NT 3.5x

Robocopy is a Unicode application, and will not run under Windows 95, because Windows 95 does not provide full Unicode support.

Also, Robocopy uses the new CopyFileEx() Win32 API, which is specific to Windows NT 4.0, and therefore will not run under Windows NT 3.5x.

Walking a Directory Tree

By default Robocopy will only process the single source directory specified on the command line. To process the whole directory tree beneath the source directory, you should use the **/S** switch or the **/E** switch. Both these options will walk the source tree, the only difference being that **/S** will refrain from creating new empty directories in the destination tree.

Retries

Most file system operations that fail and return an error will cause the program to wait and then retry the operation until it succeeds or the retry limit is reached. By default there will be 30 seconds between retries, and up to one million (1000000) retries. Use **/W:n** to change the wait time in seconds between retries, and **/R:n** to change the retry limit, where **n** is a positive decimal integer, or zero if you do not want retries or wait times between retries. If invalid values are given for **/W** or **/R**, the respective default value is used.

To change the default retry parameters, use the **/REG** switch on a valid Robocopy command that specifies non-default values for **/W** and **/R**. When **/REG** is used, the values you specify for **/W** and **/R** will be stored in the Registry and used as default values for **/W** and **/R** in future Robocopy runs where **/W** and **/R** are *not* specified on the command line. Simply specify **/W** and/or **/R** on the command line to override the stored settings.

Note that certain errors in some operations will not be retried, where practical experience indicates this would be futile. For example, an error of "Network Name Not Found" usually indicates that a remote computer exists, but does not have a sharepoint with the given name. As manual intervention will be required to correct this (by creating a suitable sharename on the remote computer), this error is generally not retried, and the attempted operation fails.

However, in some instances, this might not be the correct action for this error. For example, in a software publishing environment it is quite common to delete a sharename, update the contents of the sharepoint, then recreate the sharename. In this situation you should use the **/TBD** switch, which instructs the program to retry when it encounters a "network name cannot be found." error, on the assumption that we are waiting for the sharename to be defined.

Normally Robocopy will restart failed copies from the beginning of the file. You can override this default behaviour by specifying the **/Z** switch, which requests restartable copies - with **/Z**, failed copies will restart from close the point of failure rather than the beginning of the file. There is one exception to this - if the file's size or timestamp is modified between retries, the file has obviously been changed, and the copy is restarted from the beginning of the file.

Specifying File and Directory Names

By default, Robocopy assumes that any non-switch command-line argument is some form of file name, path, or wildcard. These may be intermingled with switch arguments, but the command line is easier to read if they are grouped together. The actual meaning of such a name depends on where they appear in the command in relation to any **/XF** or **/XD** switches.

The command line is parsed from left to right. There must be two non-switch arguments before any **/XF** or **/XD** switch, and these are taken to be the pathnames of the source and destination directories respectively.

Thereafter any non-switch argument is taken to be an **Include Filespec** - either a filename or wildcard (but not a path) naming one or more files (or sets of files) to include and consider as candidates for copying, until a **/XF** or **/XD** switch is found.

Note that if no **Include Filespecs** are found in the command line, a default of "***.***" (all files) is assumed. Also note that these **Include Filespecs** must be specified as individual arguments separated from other arguments by white space, and not appended to the source or destination directory pathnames as in, for example, the Xcopy command.

/XF (**eXclude Files**) informs the program that subsequent filenames, paths and wildcards are Files to exclude from copying (**Exclude Filespecs** rather than **Include Filespecs**), until a subsequent **/XD** switch is found.

/XD (**eXclude Directories**) informs the program that subsequent filenames and paths are Directories to exclude from copying (**Exclude Dirspecs** rather than **Include** or **Exclude Filespecs**), until a subsequent **/XF** switch is found.

Note the subtle differences in allowed values here :

<u>Arguments</u>	<u>Names</u>	<u>Paths</u>	<u>Wildcards</u>
Source Directory	Yes	Yes	<u>NO</u>
Destination Directory	Yes	Yes	<u>NO</u>
Include Filespecs	Yes	<u>NO</u>	Yes
Exclude Filespecs	Yes	Yes	Yes
Exclude Dirspecs	Yes	Yes	<u>NO</u>

Example : `ROBOCOPY c:\source d:\dest *.* /XF *.cpp /S /XD bin c:\source\unwanted`

This command would cause Robocopy to walk the directory tree whose root is c:\source, except subdirectories named "bin", and the subdirectory c:\source\unwanted. Files whose extensions begin with "c" will be copied, except those whose extension is ".cpp".

Robocopy File Classes

For each directory processed Robocopy constructs a list of files matching the **Include Filespecs**, in both the source and destination directories. The program then cross-references these lists, seeing which files exist where, comparing file times and sizes where possible, and places each selected file in one of the following classes :

<u>File Class</u>	<u>Exists in Source</u>	<u>Exists in Destination</u>	<u>Source/Dest File Times</u>	<u>Source/Dest File Sizes</u>
Lonely	Yes	<u>No</u>	n/a	n/a
Same	Yes	Yes	Equal	Equal
Changed	Yes	Yes	Equal	Different
Newer	Yes	Yes	Source > Dest	n/a
Older	Yes	Yes	Source < Dest	n/a
Extra	<u>No</u>	Yes	n/a	n/a
Mismatched	Yes (file)	Yes (directory)	n/a	n/a

By default, **Changed**, **Newer** and **Older** files will be considered to be candidates for copying (subject to further filtering described below), **Same** files will be skipped (not copied), and **Extra** and **Mismatched** files (and directories) will simply be reported in the output log.

Use the following switches to override this default behaviour :

/XL	eXclude Lonely files
/IS	Include Same files
/XC	eXclude Changed files
/XN	eXclude Newer files
/XO	eXclude Older files

Use the following switch to suppress the reporting and processing of **Extra** files :

/XX	eXclude eXtra files
------------	----------------------------

To just make sure the destination tree includes the current version of all source files, you do not need to specify any of these arguments. Robocopy's default behaviour will be all you need for most situations.

Use **/XO** with caution. If you terminate Robocopy whilst it is copying, any incompletely copied file will almost certainly have a later file time than the source file. If you restart the same copy, Robocopy will see this file as an **Older** file, and will skip it if you use **/XO**. Bear this in mind if you are using **/XO** and need to kill a copy. The most appropriate use for **/XO** is to synchronise two directory trees that can each be updated simultaneously in disjoint areas. Using **/XO** and copying first in one direction, and then the other, will ensure that the latest files are present in both directory trees.

Note that specifying **/XL** restricts the program to copying files from the source directory tree only if a corresponding file of the same name already exists in the destination. This provides a convenient mechanism for maintaining a copy of a selected subset of the source tree.

Also note that specifying **/IS** on its own with no other selection switches forces a total refresh of the destination tree, should you ever need to do this.

File Times and File Names

It is recommended that you make sure that both the source and destination directories reside on an NTFS partition, wherever possible. Copying to downlevel file systems (HPFS or FAT) will work, but you may lose long filenames, and file times may suffer from rounding errors on the destination. This is of course due to the superior capabilities of the NTFS file system.

For example, file time granularity is 100 nanoseconds on NTFS, but only 2 seconds on FAT, so copying a file from NTFS to FAT always causes file times to be rounded to a value that the FAT file system can handle. Accordingly, Robocopy considers file times to be identical if they are within two seconds of each other. Without this 2-second leeway, the program might erroneously classify many otherwise unmodified files as **Older** or **Newer** files, which would result in a great deal of unnecessary copying of these unchanged files.

Sometimes this handling of file times needs to be overridden. For example, assume an NTFS tree is copied temporarily to FAT, then later the FAT tree (and all its rounded FAT file times) is copied to a local NTFS drive. Later, there may be need to recreate the original tree exactly. Refreshing the whole tree would do the job, but it would be inefficient for a large tree. In such a scenario you should use the **/T** switch to force the copying of just file times for **Same** files, rather than the whole file that would be copied if **/IS** was used.

When the destination is on a FAT or HPFS partition, you may also experience problems when copying files and directories with long names, or whose names included extended Unicode characters. To overcome these problems, use the **/FAT** switch. This tells Robocopy to create destination files using system-generated names in the standard 8.3 FAT file system format, rather than trying to create long or extended filenames on downlevel file systems.

Attribute Processing

By default, Robocopy ignores source file attributes when selecting files to copy - any file matching other specified conditions will be copied regardless of its attribute settings.

The **/A** and **/M** switches both modify this behaviour, and cause only those source files with the Archive attribute set to be selected for copying. After copying the Archive attribute of the source file is left unmodified (still set) if **/A** was used, or reset (turned off) if **/M** was used.

Furthermore, the **/XA:[R][A][S][H]** (eXclude Attributes) switch can be used to exclude files from being copied if one or more of the given attributes is set. For example, you could specify **/XA:R** to prevent **Read-only** files from being copied. Similarly, **/XA:SH** would prevent any files with either or both of the **System** or **Hidden** attributes set from being copied.

After a file is successfully copied to the destination, the destination file's attributes are set to match those of the source file, except for the Archive bit, which is always set (turned on). This is to identify newly copied files and make it easy to back them up.

To modify this default behaviour, **/A+:[R][A][S][H]** (Attribute add) and **/A-:[R][A][S][H]** (Attribute subtract) can be used. For example, **/A-:A** would cause the **Archive** attribute to be reset, and **/A+:R** would render all copied files **Read-only** in the destination.

The exact order of attribute operations on newly copied destination files is as follows :

1. Attributes are copied from the corresponding source file.
2. The Archive attribute is set (turned on).
3. Attributes specified by **/A+:[R][A][S][H]** are set (turned on).
4. Attributes specified by **/A-:[R][A][S][H]** are reset (turned off).

Moving Files

Rather than copying files, it is often desirable to move them instead, especially if disk space is at a premium on your network. Robocopy's **/MOVE** switch provides this facility. It causes source files to be deleted from the source directory tree after they have been successfully copied to the destination.

Note that even with **/MOVE** specified, Robocopy will only delete those source files that it successfully copies to the destination. This applies even to skipped **Same** files as there is no absolute guarantee that a skipped source file is identical to its corresponding destination file, even if the file times and sizes are identical, until immediately following a successful copy.

Therefore it is perfectly normal for files and directories to remain in the source tree even after a Robocopy has walked the tree with **/MOVE** specified. The user must decide whether or not it is safe to delete the remaining entries, and, if so, delete them manually.

True Replication

If you require the destination directory tree to be maintained as an **exact** mirror of the source directory tree, and have files and directories deleted from the destination whenever they disappear from the source, you can use the **/PURGE** switch (**at your own risk**).

/PURGE causes Robocopy to delete **ALL Extra** and **Mismatched** destination files and directories. After a **Mismatched** destination entry has been deleted, the corresponding source entry is then treated as a **Lonely** file or directory, and processed accordingly.

Where **/PURGE** results in the deletion of an **Extra** or **Mismatched** destination directory, **the entire directory tree, including all subordinate files and subdirectories, will be deleted.**

You should use /PURGE with extreme caution. If you specify **/PURGE** along with an incorrect (but existing) destination directory, Robocopy **WILL DELETE** lots of data from the destination very quickly. **You use the /PURGE option at your own risk.**

Note that the **/XX** switch excludes **Extra** files from further processing, so **/PURGE** will have no effect if **/XX** is also used.

Scheduling Robocopy

You can use the Windows NT **AT** command, or the Resource Kit **SOON** command, with the Windows NT Schedule Service to create Robocopy jobs that run regularly in the background, to automatically maintain local mirrors of remote directory trees.

By default the Schedule Service logs on as the System Account for the local system, which has no network access. Scheduled jobs run in the same context as the Schedule Service. So, in order to successfully schedule a Robocopy job, some extra configuration is required. There are basically two options :

Option 1) Leave the Schedule Service running in the context of the local System Account, and schedule batch files of the following form :

```
NET USE \\remoteserver\IPC$ /USER:userid password  
ROBOCOPY \\remoteserver\sourcepath \\localserver\destpath ...  
NET USE \\remoteserver\IPC$ /DEL
```

Within each such job, credentials with remote servers are established by connecting to the **IPC\$** (Inter Process Communication) share on remote machines using an appropriate user account, rather than as the local System account. The disadvantage of this approach is that user passwords are stored in the batch files, but they can be protected by storing the batch files on an NTFS partition, and applying appropriate NTFS file security to the batch files.

Option 2) In Control Panel/Services/Schedule Service/Startup configure the Schedule Service to use a "real" user account that you have created in User Manager, by specifying an appropriate User Id and Password in the "Log On As" section. Then you should Stop and Start the Schedule Service to get it running in the new user context. Once the nominated user has been granted appropriate access to source and destination servers, you can now schedule Robocopy jobs to copy files between them.

Finally, as drive mappings can be changed by users, it is generally best to use UNC names for source and destination directories in scheduled Robocopy jobs, as these explicitly specify file locations, and are more reliable. I.e. rather than scheduling a command of the form :

```
ROBOCOPY X:\source Y:\dest ...
```

for increased reliability, you should use commands of the form :

```
ROBOCOPY \\server1\share1\source \\server2\share2\dest ...
```

AT and the Schedule Service are documented in the Windows NT Commands on-line help. **SOON** is documented in the Windows NT Resource Kit on-line documentation.

THE OUTPUT LOG

Robocopy outputs a log of files and directories processed to its console window. This output can be redirected to a disk file if required.

Each log line starts with a brief text tag, which is formatted according to the following rules :

*CAPITALS	indicate an anomaly that should be investigated.
Leading Caps	indicate a file that was selected for copying.
lowercase	indicates a file that was skipped. (lowercase tags will only be seen if /V is used).

Tags that indicate copying are left-aligned, tags that indicate skipping are right-aligned, tags that indicate anomalies are placed further to the left than other tags, and error messages always start in column 1. This arrangement simplifies the task of scanning through even a verbose listing, and makes it easier to spot new downloads, anomalies, and network errors.

If few files are copied, the left margin of the output log will be mostly blank. Copied files and anomalies are easily-spotted non-blank entries in the left margin of the output log.

Disregarding error reporting and retries, at most one line of output is produced per source file or directory. Lines for directories show the number of files matching the **Include Filespecs** in the directory (where known), and the full pathname of the directory. Lines for files indicate what was done with the file, the size of the file, and its name.

One line is also output for each **Mismatched** file and directory, and also for each **Extra** file and directory in the destination. These resemble lines for normal directories and file, except that lines for **Extra** files include the file's full pathname, as an aid to rapid housekeeping.

By default, only **Extra** files that match the **Include Filespecs** on the command line will be reported. The rationale for this is that you are probably not interested in spurious ***.TXT** files in the destination when you are refreshing ***.CPP** files. However, if you really do need a list of **all** extra files in the destination, irrespective of their type, you can obtain one by using **/X**.

By default, no output is produced for skipped files. To obtain a verbose listing which shows **all** source files matching the **Include Filespecs**, including skipped files, use the **/V** switch.

By default, Robocopy will output copy progress information, in the form of a display of the percentage of the file copied so far (% copied). You can use the **/NP** switch to suppress the display of %copied progress information - useful when output is redirected to a disk file.

To see the start time of each file copy, and the estimated time that the copy should complete (based on the observed throughput of previous copies), use the **/ETA** switch. These times are displayed after the file name in the format "hh:mm -> hh:mm" (start time -> ETA).

Finally, to obtain a list the files that **would** be copied, without committing to the overhead of actually doing the copying, you can use the **/L** switch.

The following tags apply to files :

File Tag	Meaning
-----------------	----------------

*MISMATCH	The source file corresponds to a destination directory of the same name. The source file is skipped as it cannot be copied. Housekeeping of the destination tree is recommended.
*EXTRA File	The file exists in the destination but not the source (an Extra file). Some housekeeping of the destination tree may be required.
New File	The file exists in the source but not in the destination (a Lonely file). The file is copied because /XL was not used. To skip this file, use /XL .
lonely	The file exists in the source but not in the destination (a Lonely file). The file is skipped because /XL was used. To copy this file, omit /XL .
Newer	The source file has a later timestamp than the destination file. The file is copied because /XN was not used. To skip this file use /XN .
newer	The source file has a later timestamp than the destination file. The file is skipped because /XN was used. To copy this file, omit /XN .
Older	The source file has an earlier timestamp than the destination file. The file is copied because /XO was not used. To skip this file, use /XO .
older	The source file has an earlier timestamp than the destination file. The file is skipped because /XO was used. To copy this file, omit /XO .
Changed	The source and destination files have identical timestamps, different sizes. The file is copied because /XC was not used. To skip this file use /XC .
changed	The source and destination files have identical timestamps, different sizes. The file is skipped because /XC was used. To copy this file, omit /XC .
Same	The source and destination files have identical timestamps and sizes. The file is copied because /IS was used. To skip this file, omit /IS .
same	The source and destination files have identical timestamps and sizes. The file is skipped because /IS was not used. To copy this file, use /IS .
attrib	At least one source file attribute matches the attributes specified by /XA . The file is skipped because of this. To copy this file, modify or omit /XA .
named	The file is skipped because it was named in the Exclude Filespecs . To process this file, amend the Exclude Filespecs following /XF .

The following tags apply to directories :

Directory Tag Meaning

- (blank)** A normal directory.
- *MISMATCH** This source directory corresponds to a destination file of the same name. It cannot be processed. Housekeeping of the destination is recommended.
- *EXTRA Dir** The directory exists in the destination but not the source (an **Extra** dir). Some housekeeping of the destination tree may be required.
- lonely** The directory exists in the source but not the destination (a **Lonely** dir). The directory is skipped because **/XL** was used. To process this, omit **/XL**.
- named** The directory is skipped because it was named in the **Exclude Dirspecs**. To process this directory, amend the **Exclude Dirspecs** following **/XD**.

THE RUN SUMMARY

Just before Robocopy terminates, it outputs a summary of its activities during the run to its console window (or disk file if redirected) in the following format:

	Total	Copied	Skipped	Mismatch	FAILED	Extras
Dirs :	75	0	75	0	0	0
Files :	960	13	947	0	0	1
Bytes :	19.8 m	190.0 k	19.6 m	0	0	12.5 k
Times :	0:16.914	0:03.504			0:00.000	0:13.410

This summarises the volume of data processed. The first column shows the total number of files and directories processed and the total size of files matching the **Include Filespecs** in the source. The other columns provide a breakdown of these grand totals as follows:

- Copied:** subtotals for directories created and files actually copied.
- Skipped:** subtotals for directories walked but not created, and files skipped.
- Mismatch:** subtotals for **Mismatched** files and directories.
- FAILED:** subtotals for items not processed successfully within the retry limit.
- Extras:** subtotals for items present in the destination but not the source.

The second section of the summary provides timing information for the run. Total time should be self-explanatory. This is broken down as follows:

- Copied:** time spent actually copying files, but excluding retry wait times.
- FAILED:** time spent waiting between retries for failed operations.
- Extra:** time spent scanning directories, and doing everything else.

Large times in the **FAILED** column usually indicate that network problems were experienced.

-----000000-----