
Robocopy.exe

Robust File Copy Utility

Version 1.95

Windows NT[®] 4.0 Resource Kits

Contents

Introduction	2
New Features in Version 1.95	2
Command-line Reference	4
Usage Notes	6
Using Robocopy Within a UNIX Shell	6
Processing a Directory Tree	7
Retrying After a Failed Copy	7
Specifying File and Directory Names	8
Using Robocopy File Classes	9
Using Precise File Times and Long File Names	11
Working with File Attributes	11
Maintaining True Replication	12
Copying NTFS Security Information	13
Moving Files	14
Minimizing Directory Fragmentation	14
Scheduling Robocopy Jobs	15
Output Log	16
Text Tags	17
Run Summary	19
Return Code	20
Feedback	21

Introduction

Robocopy is a 32-bit command-line utility for file replication. This tool helps maintain identical copies of a directory structure on a single computer or in separate network locations.

Important Robocopy version 1.95 requires Windows NT version 4.0 or later.

Using Robocopy, you can copy a single directory, or you can recursively copy a directory and its subdirectories. The program classifies files by whether they exist in the source directory, in the destination directory, or in 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 control which classes of files are copied.

If a file exists in both the source and destination locations, by default Robocopy copies the file only if the two versions have different time stamps or different sizes. This saves time if the source and destination are connected by a slow network link. You can also specify that copies are restarted in the event of a failure, which saves even more time when your network links are unreliable.

Robocopy allows you to do the following:

- Use file names, wildcards, paths, or file attributes to include or exclude source files as candidates for copying.
- Exclude directories by name or by path.
- Delete source files and directories after copying (that is, move rather than copy them).
- Delete destination files and directories that no longer exist in the source.
- Control the number of times the program retries an operation after encountering a recoverable network error.
- Schedule copy jobs to run automatically.

New Features in Version 1.95

Robocopy version 1.95 includes a number of enhancements. For example, you can now do the following:

- Include files with given attributes.
Use the `/IA:[R][A][S][H]` switch to include files with any of the given attributes set. Files with none of these attributes set are excluded.
- Apply time stamps to existing destination files.

The new **/TIMFIX** switch replaces the **/T** switch in previous versions of Robocopy.

Note If you have scripts that use **/T**, you must edit them to use **/TIMFIX** instead.

- Copy NTFS file security information.
The **/SEC** switch copies NTFS file security information for new files and directories. To use **/SEC**, the source and destination volumes must both use the NTFS file system.
- Refresh NTFS file security information.
The **/SECFIX** switch refreshes NTFS file security information for existing files and directories, as well as for new files and directories. **/SECFIX** is a superset of **/SEC**.
- Delete files from the source after copying them.
Similar to **/MOVE**, the **/MOV** switch deletes files from the source after successfully copying them. **/MOV** does not delete directories, however, which preserves the structure of the source directory tree.
- Mirror a directory tree.
The **/MIR** switch creates an exact copy of a directory tree. Using **/MIR** is equivalent to specifying both **/E** and **/PURGE**.
- Exclude files according to their size.
The **/MAX:n** and **/MIN:n** options exclude files larger than **/MAX:n** bytes or smaller than **/MIN:n** bytes, respectively.
- Exclude files according to their age.
The **/MAXAGE:n** and **/MINAGE:n** options exclude files older than **/MAXAGE:n** or newer than **/MINAGE:n**, respectively.
- Create directory tree without copying data.
The **/CREATE** switch copies the directory tree structure and creates zero-length files. No file data is copied. Using this option minimizes directory fragmentation.
- Copy top levels of source directory tree.
Use **/LEV:n** to copy only the top *n* levels of the source directory tree.
- Send output to log file.
/LOG:file redirects output to the specified file, overwriting the file if it already exists.

/LOG+;*file* redirects output to the specified file, appending it to the file if it already exists.

Version 1.95 also includes the following improvements:

- New return code
The range and meaning of Robocopy return code has changed and expanded. For more information, see “Return Code” later in this document.
- Consistent display of run times
All times are now shown in HH:MM:SS format, regardless of the length of the operation.
- Display of copy speed
Copy throughput figures are now calculated and displayed in the run summary.
- Concise usage details
By default, **ROBOCOPY /?** now produces only brief usage details at the command prompt. To display full usage details, type **ROBOCOPY /???**

Command-line Reference

To run Robocopy, use the following syntax at the command prompt:

ROBOCOPY *source destination* [*file* [*file*]...] [*options*]

The following table defines these syntax elements.

Variable	Meaning	Comments
<i>source</i>	Source directory	You can use <i>drive:\path</i> or <i>\\server\share\path</i>
<i>destination</i>	Destination directory	You can use <i>drive:\path</i> or <i>\\server\share\path</i>
<i>file</i>	Names of files to act upon	You can use wildcards (? and *). If no files are listed, Robocopy defaults to all files (*.*)
<i>options</i>	Command-line options you wish to use	Available options are described later in this document.

Tip To view brief usage instructions at the command prompt, run **ROBOCOPY** without specifying any command-line options.

The following table defines the command-line options that you can use with Robocopy.

This switch	Performs this function
/S	Copies subdirectories (excluding empty ones).
/E	Copies all subdirectories (including empty ones).
/LEV:n	Copies only the top <i>n</i> levels of the source directory tree.
/Z	Copies files in restartable mode (that is, restarts the copy process from the point of failure).
/SEC	Copies NTFS security information. (Source and destination volumes must both be NTFS.)
/SECFIX	Applies the NTFS permissions set for source files and directories to existing destination files and directories.
/TIMFIX	Applies a time stamp to all destination files (including skipped files).
/MOV	Moves files (that is, deletes source files after copying).
/MOVE	Moves files and directories (that is, deletes source files and directories after copying).
/PURGE	Deletes destination files and directories that no longer exist in the source.
/MIR	Mirrors a directory tree (equivalent to running both /E and /PURGE).
/A+:[R][A][S][H]	Sets the specified attributes in copied files.
/A-:[R][A][S][H]	Turns off the specified attributes in copied files.
/CREATE	Creates a directory tree structure containing zero-length files only (that is, no file data is copied).
/FAT	Creates destination files using only 8.3 FAT file names.
/IA:[R][A][S][H]	Includes files with the specified attributes.
/XA:[R][A][S][H]	Excludes files with the specified attributes.
/A	Copies only files with the archive attribute set.
/M	Copies only files with the archive attribute set and then resets (turns off) the archive attribute in the source files.
/XF <i>file</i> [<i>file</i>]	Excludes files with the specified names, paths, or wildcards.
/XD <i>dir</i> [<i>dir</i>]	Excludes directories with the specified names, paths, or wildcards.
/XC	Excludes files tagged as <i>Changed</i> .

/XN	Excludes files tagged as <code>Newer</code> .
/XO	Excludes files tagged as <code>Older</code> .
/XX	Excludes files and directories tagged as <code>Extra</code> .
/XL	Excludes files and directories tagged as <code>Lonely</code> .
/IS	Includes files tagged as <code>Same</code> .
/MAX:n	Excludes files larger than <i>n</i> bytes.
/MIN:n	Excludes files smaller than <i>n</i> bytes.
/MAXAGE:n	Excludes files older than <i>n</i> days or specified date. If <i>n</i> is less than 1900, then <i>n</i> is expressed in days. Otherwise, <i>n</i> is a date expressed as YYYYMMDD.
/MINAGE:n	Excludes files newer than <i>n</i> days or specified date. If <i>n</i> is less than 1900, then <i>n</i> is expressed in days. Otherwise, <i>n</i> is a date expressed as YYYYMMDD.
/R:n	Specifies the number of retries on failed copies. (The default is 1 million.)
/W:n	Specifies the wait time between retries. (The default is 30 seconds.)
/REG	Saves /R:n and /W:n in the registry as default settings.
/TBD	Waits for share names to be defined on a “Network Name Not Found” error.
/L	Lists files without copying, deleting, or applying a time stamp to any files.
/X	Reports all files tagged as <code>Extra</code> (including files not selected).
/V	Produces verbose output (including skipped files).
/NP	Turns off copy progress indicator (% copied).
/ETA	Shows estimated time of completion for copied files.
/LOG:file	Redirects output to the specified file, overwriting the file if it already exists.
/LOG+:file	Redirects output to the specified file, appending it to the file if it already exists.

Usage Notes

The following sections provide more details about using the command-line options to tailor Robocopy to match your needs.

Using Robocopy Within a UNIX Shell

You can specify all Robocopy switches in UNIX style (for example, **-ETA** instead of **/ETA**). You can also specify source and destination directory paths using the UNIX delimiter (**/**), rather than the native Windows delimiter (****).

The only restriction is that any argument that begins with a UNIX slash mark (/) is treated as a switch if the argument contains only a single slash mark (/). In other words, /dir is treated as a switch; but //server/share/dir and /download/test are treated as paths.

This avoids any possible confusion between switches and single-level paths subordinate to the root of a drive. To specify a directory as an argument, you must use an alternate expression for its path, such as **X:/dir** or **//server/C\$/dir**.

Note Robocopy uses the CopyFileEx() Win32 API, which is specific to Windows NT version 4.0 and later. Robocopy does not run under the Windows NT version 3.5x or Windows 95 operating system.

Processing a Directory Tree

By default, Robocopy processes only the single source directory specified on the command line. To process the entire directory tree, including subdirectories, use **/S** or **/E**. To limit the copy to only the top *n* levels of the source directory tree, use **/LEV:n**.

Retrying After a Failed Copy

When a copy operation fails and returns an error, Robocopy waits and then retries the operation until it succeeds or until the retry limit is reached. The defaults are a 30-second wait and up to 1 million retries.

Changing Default Retry Parameters

Use the **/W:n** switch to change the wait time (in seconds) between retries, where *n* is a positive decimal integer, or zero (0) if you do not want any wait time between retries. Similarly, use the **/R:n** switch to change the retry limit. If an invalid value is given for **/W** or **/R**, Robocopy uses the default value.

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** are 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. You can always specify **/W** or **/R** to override the stored settings.

Handling “Network Name Not Found” Errors

Robocopy does not automatically retry all failed operations. Certain errors in some operations are not retried, where practical experience indicates that this would be futile. For example, the “Network Name Not Found” error usually indicates that a remote computer exists but does not have a share with the given name. Because manual intervention is required to create a suitable share name on the remote computer, this error usually causes the attempted operation to fail without a retry.

In some situations, however, this might not be the desired action. For example, in a software publishing environment it is common to delete a share name, update the contents of the share, and then recreate the share name. In this scenario, you can use the **/TBD** switch to retry a copy operation when Robocopy encounters a “Network Name Not Found” error.

Restarting from the Point of Failure

Normally, Robocopy restarts failed copies from the beginning of the file. You can override this default behavior with the **/Z** switch, which specifies restartable copies. With **/Z**, failed copies restart from the point of failure rather than the beginning of the file. There is one exception: If the size or time stamp of a file has been modified between retries, Robocopy restarts from the beginning of the file.

Specifying File and Directory Names

By default, Robocopy treats all non-switch command-line arguments as file names, paths, or wildcards. These names can be alternated with switch arguments, but the command line is easier to read when they are grouped together.

The command line is parsed from left to right. The meaning of any non-switch argument depends on where it appears with respect to an **/XF** or **/XD** switch. There must be two non-switch arguments before any **/XF** or **/XD** switch, and these are treated as the path names of the source and destination directories, respectively.

After the source and destination path names, and until Robocopy encounters an **/XF** or **/XD** switch, any non-switch argument is treated as either a file name or a wildcard (but not a path) indicating one or more files to include as candidates for copying.

To specify a directory or file name or path name that includes one or more spaces, you must enclose the entire file or path name in quotation marks. For example:

```
ROBOCOPY "\\sales\reports\Fiscal 97" "C:\finance\Year  
End\sales" /XF "Budget 96*.*"
```

Note Any file names that appear in the command line must be specified as individual arguments separated from other arguments by white space, and not appended to the source or destination directory path names. If no such arguments appear in the command line, Robocopy defaults to all files.

The **/XF** switch specifies that all subsequent file names, paths, and wildcards are files to exclude from copying until Robocopy encounters an **/XD** switch. The **/XD** switch specifies that all subsequent file names and paths are directories to exclude from copying until Robocopy encounters an **/XF** switch.

The following table compares these subtle differences in allowed values:

Argument	Names	Paths	Wildcards
Source directory	Yes	Yes	No
Destination directory	Yes	Yes	No
Files to include	Yes	No	Yes
Files to exclude	Yes	Yes	Yes
Directories to exclude	Yes	Yes	No

The following example shows the correct use of these switches:

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

In this example, Robocopy recursively copies the directory tree (including subdirectories) with the root `c:\source`, excluding the subdirectories `bin` and `c:\source\unwanted`. All files whose extensions begin with `c` are copied, except files with the extension `.cpp`.

Using Robocopy File Classes

For each directory processed, Robocopy constructs a list of files in both the source and destination directories. This list matches the files specified on the command line for copying.

Robocopy then cross-references the lists, determining where files exist and comparing file times and sizes. The program places each selected file in one of the following classes.

File class	In source	In destination	Source/dest file times	Source/dest file sizes
Lonely	Yes	No	n/a	n/a
Same	Yes	Yes	Equal	Equal
Changed	Yes	Yes	Equal	Different
Newer	Yes	Yes	Source > destination	n/a
Older	Yes	Yes	Source < destination	n/a
Extra	No	Yes	n/a	n/a
Mismatched	Yes (file)	Yes (directory)	n/a	n/a

By default, `Changed`, `Newer`, and `Older` files are candidates for copying (subject to further filtering, as described later). `Same` files are not copied. `Extra` and `Mismatched` files and directories are only reported in the output log.

Use the following switches to override this default behavior.

This switch	Performs this function
<code>/XL</code>	Excludes <code>Lonely</code> files and directories.
<code>/IS</code>	Includes <code>Same</code> files.
<code>/XC</code>	Excludes <code>Changed</code> files.
<code>/XN</code>	Excludes <code>Newer</code> files.
<code>/XO</code>	Excludes <code>Older</code> files.
<code>/XX</code>	Suppresses reporting and processing of <code>Extra</code> files.

Note You do not need to specify any of these arguments to include the current version of source files in the destination tree. In most situations, Robocopy includes these files by default.

The `/XL` switch limits Robocopy to copying a file from the source directory tree only when 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. Using the `/IS` switch alone, with no other switches, forces a total refresh of the destination tree.

If you terminate Robocopy during a copy operation, any incompletely copied file has a later time stamp than its corresponding source file. If you restart the same copy operation, Robocopy treats this file as an `Older` file and skips it if you use the `/XO` switch.

The most appropriate use for `/XO` is to synchronize two directory trees so that they can be updated simultaneously in separate areas. To ensure that the latest files are present in both directory trees, copy with `/XO` first in one direction and then in the other.

You can use switches to exclude files based on their size or time stamp. `/MAX:n` excludes files larger than *n* bytes, and `/MIN:n` excludes files smaller than *n* bytes. Similarly, `/MAXAGE:n` excludes files older than *n* and `/MINAGE:n` excludes files newer than *n*, where *n* is either a time span in days (if *n* < 1900), or a date in the form YYYYMMDD (if *n* >= 1900).

Using Precise File Times and Long File Names

Although it is preferable to have both the source and destination directories reside on NTFS partitions, you can use Robocopy to copy to other file systems, such as HPFS or FAT. If you do so, however, file times might be inaccurate due to rounding errors, and long file names might not be preserved.

Handling FAT File Times on NTFS Drives

File-time granularity is 100 nanoseconds on NTFS, but only two seconds on FAT, so copying a file from NTFS to FAT causes file times to be rounded to a value that the FAT file system can manage. Robocopy considers file times to be identical if they are within two seconds of each other. Without this two-second margin, the program might classify many otherwise unmodified files as `Older` or `Newer` files, which would result in unnecessary copying of unchanged files.

Sometimes you need to override this handling of file times. Suppose you copy an NTFS tree to a FAT volume, and then copy the FAT tree (with its rounded file times) to a local NTFS drive. Later, when you want to recreate the original tree exactly, you do not want to refresh the entire tree. In this scenario, you can use the `/TIMFIX` switch, which copies only the file times for `Same` files, rather copying the files themselves.

Handling Long File Names on FAT or HPFS Volumes

When the destination is on a FAT or HPFS volume, you might experience problems copying files and directories with long names or with names that include extended Unicode characters. To overcome these problems, use the `/FAT` switch. This switch prompts Robocopy to create destination files with system-generated names in the standard 8.3 FAT file system format.

Working with File Attributes

By default, Robocopy ignores source file attributes (read-only, archive, system, and hidden) when selecting files to copy. Any file matching other specified conditions is copied regardless of its attribute settings. You can use command-line switches to modify this default behavior, as described in the following sections.

Including or Excluding Files Based on Specified Attributes

The `/IA:[R][A][S][H]` switch includes only files with one or more of the specified attributes set. Files that do not have any of the specified attributes set are excluded from further processing and are not copied.

`/IA` can be useful in data staging applications. For example, you can specify `/IA:R` to copy only read-only files from a test Web server to a live Web server. Developers can remove the read-only attribute from files on the test server while they edit and test the

files. Until the read-only attribute is turned on again, these files are not replicated on the live server.

Another switch, **/XA:[R][A][S][H]**, can be used to exclude files if one or more of the specified attributes is set. For example, you can specify **/XA:R** to prevent read-only files from being copied. Similarly, **/XA:SH** excludes files with either the system or hidden attribute set.

Note Robocopy processes **/IA** before **/XA** when both switches are specified in the same Robocopy command.

The **/A** and **/M** switches allow only those source files with the archive attribute set to be selected for copying. After copying, the archive attribute of the source file is either left set (if **/A** was used) or turned off (if **/M** was used).

Changing File Attributes After Copying

After a file has been successfully copied to the destination, the attributes of the destination file are set by default to match those of the source file. To modify this default behavior, you can use two other switches: **/A+** sets attributes, and **/A-** resets or turns off attributes.

For example, **/A+:A** turns on the archive attribute in destination files (to make it easier to back them up). **/A-:RSH** removes the read-only, system, and hidden attributes from destination files (to make it possible to see and edit them).

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

1. Attributes are copied to the destination file from the source file.
2. Attributes specified by **/A+** are set (turned on).
3. Attributes specified by **/A-** are reset (turned off).

Maintaining True Replication

If you require the destination directory tree to be maintained as an *exact* mirror of the source directory tree, you must delete files and directories from the destination whenever they disappear from the source. To accomplish this, use the **/E** and **/PURGE** switches in combination, or use the **/MIR** switch.

The **/PURGE** switch 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. When **/PURGE** results in the deletion of an `Extra` or

Mismatched destination directory, the entire directory tree, including all subordinate files and subdirectories, is deleted (even if **/S** or **/E** is not specified).

Caution Use **/PURGE** and **/MIR** very carefully. If you inadvertently specify **/PURGE** or **/MIR** with an existing destination directory, Robocopy can quickly delete a large amount of data from the destination.

To prevent specific directory trees or files from being deleted when you use **/PURGE**, you must exclude them by name, using the **/XD** or **/XF** options. **/XD** and **/XF** exclude files and directories before all other processing.

Note Because the **/XX** switch excludes `Extra` files from further processing, the **/PURGE** switch has no effect when **/XX** is also used.

Copying NTFS Security Information

If both the source and destination directories reside on NTFS volumes, you can use Robocopy to copy NTFS security information. Use the **/SEC** switch to copy NTFS file permissions when you copy files to a destination for the first time. For performance reasons, **/SEC** does not copy security information for existing destination directories and files.

To refresh security information for existing destination files and directories without copying file data, use the **/SECFIX** switch. **/SECFIX** is a superset of **/SEC**. **/SECFIX** copies permissions for newly copied files and refreshes security information for existing destination files and directories, even those skipped because they are tagged as `Same`, `Older`, `Newer`, or `Changed`. Files excluded for other reasons (for example, those tagged as `named`) are excluded from all processing, and their security information is not updated.

Note If either the source or the destination volume is not NTFS, **/SEC** and **/SECFIX** are ignored. Only one attempt is made to copy NTFS security information for each file or directory copied.

When you use **/SEC**, you might notice a short delay after Robocopy prints its banner and before it displays arguments on the screen. This delay occurs because the program is checking for the NTFS file system at both ends. Copying NTFS file security information fails if you have insufficient access privileges at either the source or the destination.

For more information about access privileges, see “Scheduling Robocopy Jobs” later in this document.

You can copy security information between computers in separate domains even if they do not have trust relationships. However, this applies security information that is unknown to the destination computer to the destination directory tree. When you view this information on the destination, user accounts are shown as “Account Unknown.” If you copy these permissions back to a computer in the originating domain, they are interpreted correctly again.

Moving Files

In some situations, you may want to move files instead of copying them, especially if disk space is at a premium on your network. You can use the **/MOVE** switch to delete source files from the source directory tree after they have been successfully copied to the destination. **/MOVE** also deletes empty directories from the source directory tree.

Even with **/MOVE** specified, however, Robocopy deletes only those source files that it successfully copies to the destination. There is no 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. For that reason, some files and directories (such as skipped `same` files) might remain in the source tree even when **/MOVE** is specified. You must determine whether to delete the remaining entries, and do so manually.

Note To move files but still preserve the source directory tree structure, use **/MOV** instead of **/MOVE**.

Minimizing Directory Fragmentation

When you copy a directory tree for the first time, the files contained in each directory are copied into it as it is created. The directory entry for each file is created first, and then the file data is copied.

If the tree contains large directories, however, this sequence can cause fragmentation of the directory structure. When disk allocation for the directory itself is expanded, the new directory extent can be separated from the rest of the directory by large amounts of file data, which can reduce performance significantly.

Use the **/CREATE** switch to minimize directory fragmentation when you copy a large tree for the first time. When **/CREATE** is specified, Robocopy creates only the directory tree and zero-length files as placeholders for the actual files. No file data is copied. Immediately following a Robocopy command containing **/CREATE**, run the same command again without **/CREATE** to copy the file data to the destination.

Note When using this technique, do not use **/XO**, because the zero-length files created by the initial **/CREATE** command are tagged as `Older` files on the second run.

Scheduling Robocopy Jobs

You can create Robocopy jobs that run regularly and automatically in the background to maintain local mirrors of remote directory trees. In Windows NT, use the **AT** command, or use the **SOON** command with the Windows NT Schedule service. (The **SOON** command runs `Soon.exe`, a Windows NT Resource Kit utility named Near-Future Command Scheduler.)

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. To run a scheduled Robocopy job, however, Robocopy must connect to network drives, which requires some additional configuration. The following sections describe two ways to accomplish this.

Using the Local System Account

You can leave the Schedule service running in the context of the local system account. You can then schedule batch files by using the following commands:

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

Within each job, credentials with remote servers are validated by connecting to the `IPC$` (Interprocess Communication) share on remote computers using an appropriate user account, rather than the local system account. The disadvantage of this approach is that user passwords are stored in the batch files; however, passwords can be protected by storing the batch files on an NTFS volume, and setting appropriate NTFS permissions on the batch files.

Using a User Account

You can configure the Schedule service to log on as a user account. Once the user account has been granted appropriate access to source and destination servers, you can schedule Robocopy jobs to copy files between them.

To configure the Schedule service to log on as a user account

1. In Services in Control Panel, click **Schedule**, and then click **Startup**.
2. In the **Startup Type** box, click **Automatic** or **Manual**.
3. In the **Log On As** box, select the account in which to run the service, and specify the correct logon password.

Tip Because drive mappings can be changed by users, it is more reliable to use UNC names for source and destination directories in scheduled Robocopy jobs. For example:

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

Note The **AT** command and the Schedule service are documented in Windows NT online Help. The **SOON** command is documented in Rktools.hlp on the *Windows NT Server Resource Kit* or *Windows NT Workstation Resource Kit* CD.

Output Log

After copying, Robocopy displays a log of processed files and directories at the command prompt. This output can be redirected to a file, either by using the standard command-line operators (> or >>), or by using one of the command-line options:

- **/LOG:file** redirects output to the specified file, overwriting the file if it already exists.
- **/LOG+:file** redirects output to the specified file, appending it to the file if it already exists.

Aside from error reporting and retries, Robocopy produces one line of output per source file or directory. Lines for directories show the number of files in that directory specified for copying (where known) and the full path of the directory. Lines for files indicate the name of the file, the size of the file, and what was done with the file.

One line is also generated for each `Mismatched` file and directory and each `Extra` file and directory in the destination. Lines for `Extra` files include the full path and name of the file.

By default, the only `Extra` files reported are those that match files specified for copying on the command line. In most cases, this is more efficient. For example, if you are refreshing CPP files, you probably do not need information about TXT files in the destination. If you want a list of all `Extra` files in the destination, regardless of their type, use the **/X** switch.

No output is produced for skipped files. To obtain a verbose listing that shows all source files specified for copying on the command line, including skipped files, use the **/V** switch.

Robocopy provides copy progress information (% copied) by default. You can use the **/NP** switch to suppress the display of progress information. This can be useful when output is redirected to a file.

To see the start time of each file copy and the estimated time of completion based on the observed throughput of previous copies, use the `/ETA` switch. Times are displayed after the file name in the format `HH:MM -> HH:MM` (start -> finish).

Finally, to obtain only a list of the files that would be copied (without actually copying them), use the `/L` switch.

Text Tags

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

- All capital letters indicate an anomaly that should be investigated.
- Initial capital letters indicate a file that was selected for copying.
- All lowercase letters indicate a file that was skipped (displayed only if the `/V` switch is used).

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

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

Tags for Files

The following table lists the tags that apply to files.

File tag	Meaning
*MISMATCH	The source file corresponds to a destination directory of the same name. The source file is skipped, and housekeeping of the destination tree is recommended.
*EXTRA File	The file exists in the destination but not the source (an <code>Extra</code> file). Some housekeeping of the destination tree may be required.
New File	The file exists in the source but not in the destination (a <code>Lonely</code> file). The file is copied; to skip this file, use <code>/XL</code> .
lonely	The file exists in the source but not in the destination (a <code>Lonely</code> file). The file is skipped; to copy this file, omit <code>/XL</code> .
Newer	The source file has a later time stamp than the destination file. The file is copied; to skip this file, use <code>/XN</code> .
newer	The source file has a later time stamp than the destination file. The file is skipped; to copy this file, omit <code>/XN</code> .

Older	The source file has an earlier time stamp than the destination file. The file is copied; to skip this file, use /XO .
older	The source file has an earlier time stamp than the destination file. The file is skipped; to copy this file, omit /XO .
Changed	The source and destination files have identical time stamps but different file sizes. The file is copied; to skip this file, use /XC .
changed	The source and destination files have identical time stamps but different file sizes. The file is skipped; to copy this file, omit /XC .
Same	The source and destination files have identical time stamps and file sizes. The file is copied; to skip this file, omit /IS .
same	The source and destination files have identical time stamps and file sizes. The file is skipped; to copy this file, use /IS .
attrib	At least one source file attribute matches the attributes specified by /XA . The file is skipped; to copy this file, modify or omit /XA .
named	The file is skipped because it was named following an /XF switch. To process this file, amend the argument following /XF .
large	The file is skipped because it is larger than /MAX:n bytes. To process this file, increase the value of n .
small	The file is skipped because it is smaller than /MIN:n bytes. To process this file, reduce the value of /MIN:n .
too old	The file is skipped because it is older than the date defined by /MAXAGE:n . To process this file, amend n to specify an earlier date.
too new	The file is skipped because it is newer than the date defined by /MINAGE:n . To process this file, amend n to specify a more recent date.

Tags for Directories

The following table lists the tags that apply to directories.

Directory tag	Meaning
(blank)	A normal directory.
*MISMATCH	This source directory corresponds to a destination file of the same name and cannot be processed. Housekeeping of the destination is recommended.
*EXTRA Dir	The directory exists in the destination but not the source (an <code>Extra</code> directory). Some housekeeping of the destination tree may be required.
lonely	The directory exists in the source but not the destination (a <code>Lonely</code> directory). The directory is skipped; to process this directory, omit <code>/XL</code> .
named	The directory is skipped because it was named following an <code>/XD</code> switch. To process this directory, amend the argument following <code>/XD</code> .

Run Summary

Before Robocopy terminates, it displays a summary of its activities at the command prompt (or redirects the information to a file). This output summarizes the volume of data processed in the following format:

	Total	Copied	Skipped	Mismatch	FAILED	Extras
Dirs:	13	0	13	0	0	1
Files:	157	2	155	0	0	3
Bytes:	33.73 m	691.5 k	33.05 m	0	0	162
Times:	0:01:33	0:01:04			0:00:00	0:00:29
Speed:		10952 bps				
Speed:		0.626 mb/min				

The first column of the run summary indicates the total number of files and directories processed, the total size of source files matching the file names specified for copying, and the total time taken by the operation. The remaining columns provide a breakdown of the totals, as described in the following table.

This column	Provides this information
Copied	Number of directories created and files copied
Skipped	Number of directories processed but not created, and number of files skipped
Mismatch	Number of <code>Mismatched</code> files and directories
FAILED	Number of items not successfully processed within the retry limit
Extras	Number of items present in the destination but not the source

The `Times` line of the summary provides timing information for the run, as described in the following table. Times are displayed in HH:MM:SS format.

This column	Provides this information
Copied	Time spent copying files (excluding retry wait times)
FAILED	Time spent waiting between retries for failed operations
Extra	Time spent scanning directories and performing other functions

Note Large times in the FAILED column usually indicate network problems.

Finally, if files were copied, the `Speed` lines show throughput information, calculated by dividing the volume of data copied by the time spent copying. The same data is displayed twice: first in bytes per second (bps), and then in megabytes per minute (mb/min).

Return Code

The return code from Robocopy (version 1.74 and later) is a bit map, defined as follows:

Hex bit value	Decimal value	Meaning if set
0x10	16	Serious error. Robocopy did not copy any files. This is either a usage error or an error due to insufficient access privileges on the source or destination directories.
0x08	8	Some files or directories could not be copied (copy errors occurred and the retry limit was exceeded). Check these errors further.
0x04	4	Some <code>Mismatched</code> files or directories were detected. Examine the output log. Housekeeping is probably necessary.
0x02	2	Some <code>Extra</code> files or directories were detected. Examine the output log. Some housekeeping may be needed.
0x01	1	One or more files were copied successfully (that is, new files have arrived).
0x00	0	No errors occurred, and no copying was done. The source and destination directory trees are completely synchronized.

You can use this information in a batch file to report the most serious anomalies, as follows:

```
if errorlevel 16 echo ***FATAL ERROR*** & goto end
if errorlevel 8 echo **FAILED COPIES** & goto end
if errorlevel 4 echo *MISMATCHES* & goto end
if errorlevel 2 echo EXTRA FILES & goto end
if errorlevel 1 echo Copy successful & goto end
if errorlevel 0 echo --no change-- & goto end
:end
```

Alternatively, full details of the return code could be reported as follows:

```
if errorlevel 16 echo ***FATAL ERROR*** & goto end
if errorlevel 15 echo FAIL MISM XTRA COPY & goto end
if errorlevel 14 echo FAIL MISM XTRA & goto end
if errorlevel 13 echo FAIL MISM COPY & goto end
if errorlevel 12 echo FAIL MISM & goto end
if errorlevel 11 echo FAIL XTRA COPY & goto end
if errorlevel 10 echo FAIL XTRA & goto end
if errorlevel 9 echo FAIL COPY & goto end
if errorlevel 8 echo FAIL & goto end
if errorlevel 7 echo MISM XTRA COPY & goto end
if errorlevel 6 echo MISM XTRA & goto end
if errorlevel 5 echo MISM COPY & goto end
if errorlevel 4 echo MISM & goto end
if errorlevel 3 echo XTRA COPY & goto end
if errorlevel 2 echo XTRA & goto end
if errorlevel 1 echo COPY & goto end
if errorlevel 0 echo --no change-- & goto end
:end
```

Feedback

For questions or feedback concerning this utility, please contact rkinput@microsoft.com.

If you wish to report a problem encountered when using Robocopy, please supply the following information.

- Version number
- Operating system and Service Pack for all computers
- File system
- Complete Robocopy command line or batch file
- Output log
- Directory listings

- Dr. Watson log (for application errors)
- Method of running Robocopy (from the command prompt, using Schedule service, and so on)
- User account information

Copyright © Microsoft Corp. 1995 – 1998