

Native Files

<p><code>{R} ← {X} □MKDIR Y</code></p>	<p>Creates new fully-qualified files/directories Y along with intermediate directories if required, as specified by X. X can be any summation of the following values (maximum value = 3):</p> <ul style="list-style-type: none"> • 0 : directory Y is only created if its path already exists and its base name does not exist. This is the default. • 1 : no action is taken if directory Y already exists. • 2 : directory Y and any part of its path that does not already exist is created. <p>Returns 1 for each file/directory specified in Y that is created successfully, 0 otherwise (if X is 0 or 2 then an error message is also generated if directory Y cannot be created successfully).</p>
<p><code>{R} ← X □NAPPEND Y</code></p>	<p>Appends the ravel of array X to the end of the native file that has tie number Y[1]; optionally, Y[2] can specify the conversion code to use to convert array X (by default, 80 is assumed when using the Unicode version – see <i>Conversion Codes</i>).</p>
<p><code>{R} ← X □NCOPY Y</code></p>	<p>Copies native files/directories from one or more sources specified by Y to destination X. If X specifies an existent directory then each source in Y is copied into that directory, otherwise X specifies the name of the copy. X must specify an existent directory if Y contains multiple names or if the wildcard option (requires variant) is set. For files, further variant options can preserve attributes when copying (<i>PreserveAttributes</i>, Boolean) and define behaviour if the target file already exists (<i>IfExists</i> can be 'Error', 'Skip', 'Replace' or 'ReplaceIfNewer').</p>
<p><code>{R} ← X □NCREATE Y</code></p>	<p>Creates a new native file with name X and file tie number Y; a tie number of 0 allocates the next available tie number to the file. Optionally (requires variant) allows semi-automatic naming of files and for existing files to be overwritten.</p>
<p><code>{R} ← {X} □NDELETE Y</code></p>	<p>Deletes each fully-qualified file/empty directory specified in Y, returning a shy result of the count of top-level entities successfully deleted. A wildcard option (requires variant) can be used. If a file/directory specified in Y does not exist then behaviour is determined by X. X can be any summation of the following values (maximum value = 3):</p> <ul style="list-style-type: none"> • 0 : each file/directory specified in Y must exist. This is the default. • 1 : no action is taken if a file/directory specified in Y does not exist. • 2 : any non-empty directory specified in Y and its contents are deleted.
<p><code>{R} ← X □NERASE Y</code></p>	<p>Erases the tied native file that has name X and file tie number Y.</p>
<p><code>R ← □NEXISTS Y</code></p>	<p>Returns 1 for each of the fully-qualified files/directories specified in Y that exists (can be accessed), 0 otherwise.</p>
<p><code>R ← {X} □NGET Y</code></p>	<p>Reads a text file and returns a 3-element vector in which:</p> <ul style="list-style-type: none"> • [1] is the character data in the file. • [2] is a character vector specifying the file encoding (see <i>File Encoding</i>). • [3] is numeric and is either \emptyset (if no line separator was found in the file) or a vector specifying the first newline separator found in the file (see <i>Line Separators</i>). <p>Y is either a character vector/scalar containing the name of the file to be read or a 2-item vector in which [1] is the filename and [2] is an integer scalar specifying f l a g s:</p> <ul style="list-style-type: none"> • if f l a g s is 0 then R[1] is a character vector. This is the default. • if f l a g s is 1 then R[1] is a nested array of character vectors. <p>X is a character vector specifying the decoding format to use if the specified file does not start with a recognised BOM (see <i>File Encoding</i>). If no BOM is present and no decoding format is specified, then the file is examined and its encoding format is deduced.</p>
<p><code>R ← {X} □NINFO Y</code></p>	<p>Returns an array of the information specified by X about Y (file/directory names or native file tie numbers). X can specify any of the following values, in any order:</p> <ul style="list-style-type: none"> • 0 : Name of Y. This is the default. • 1 : Type of Y. This can be 0 (not known), 1 (directory), 2 (file), 3 (character device) or 4 (symbolic link). On UNIX and Mac OS, can also be 5 (block device), 6 (FIFO) or 7 (socket) • 2 : Size of Y in bytes. • 3 : The time Y was last modified as a timestamp in <code>□TS</code> format. • 4 : The user ID of the owner of Y.

	<ul style="list-style-type: none"> • 5 : The name of the owner of <i>Y</i>. • 6 : Whether <i>Y</i> is hidden. • 7 : Target of symbolic link – only applies when Type=4. • 8 : Current file position (offset) – only applies when the source in <i>Y</i> is a native file tie number. • 9 : The time <i>Y</i> was last accessed as a timestamp in <code>□TS</code> format (if available). • 10 : The time <i>Y</i> was created as a timestamp in <code>□TS</code> format (if available). • 11 : Whether <i>Y</i> can be read. • 12 : Whether <i>Y</i> can be written to. <p>The file/directory name can include wildcard options, recursion through sub-directories and property reporting for symbolic links (requires variant); in this case the result <i>R</i> remains the same shape as <i>X</i> but its depth increases.</p>
<code>{R} ← X □NLOCK Y</code>	Changes the lock status (as defined by <i>X</i>) of part of the native file that has file tie number <i>Y</i> [1]; optionally, <i>Y</i> [2] can define the offset from 0 of the first byte to apply the lock change to (defaults to 0) and <i>Y</i> [3] can specify the number of bytes impacted by the lock change (defaults to the maximum possible file size) (see <i>File Encoding</i>).
<code>R ← □NAMES</code>	Lists the names of all tied native files.
<code>{R} ← X □NMOVE Y</code>	Moves native files/directories from one or more sources specified by <i>Y</i> to destination <i>X</i> . If <i>X</i> specifies an existent directory then each source in <i>Y</i> is moved into that directory. <i>X</i> must specify an existent directory if <i>Y</i> contains multiple names or if the wildcard option (requires variant) is set. For files, further variant options can define behaviour if moving to a different file system (<i>RenameOnly</i> , Boolean) and or if the target file already exists (<i>IfExists</i> can be 'Error' or 'Skip').
<code>R ← □NNUMS</code>	Lists the tie numbers of all tied native files.
<code>R ← {X} □NPARTS Y</code>	Splits the fully-qualified files/directories specified in <i>Y</i> into constituent parts, returning a 3-element vector (or vector of 3-element vectors if <i>Y</i> comprises multiple files/directories) in which: <ul style="list-style-type: none"> • <i>R</i>[1] is the specified path. If <i>X</i> is 1 then the fully-qualified path is derived and returned. • <i>R</i>[2] is the base name, that is, the filename stripped of its path and extension. • <i>R</i>[3] is the file extension, including the leading <code>.</code> character.
<code>{R} ← X □NPUT Y</code>	Writes character data <i>X</i> to a text file <i>Y</i> and returns (shy) the number of bytes successfully written to the file. <i>Y</i> is either a character vector/scalar containing the name of the file to be written or a 2-item vector in which [1] is the filename and [2] is an integer scalar specifying <i>flags</i> : <ul style="list-style-type: none"> • if <i>flags</i> is 0 then, if the file already exists, it will not be overwritten and an error is generated. This is the default. • if <i>flags</i> is 1 then the file will be written irrespective of whether it already exists. • if <i>flags</i> is 2 then <i>X</i> will be appended to <i>Y</i>. <i>X</i> can comprise up to three elements: <ul style="list-style-type: none"> • [1] is a vector of character vectors, each of which represents a line in the file to be written, or a simple character vector. • [2] is a character vector specifying the encoding to use (see <i>File Encoding</i>). Optional – the default is UTF-8-NOBOM. • [3] is numeric and is either \emptyset (same as the default) or a scalar/vector specifying the newline separator to use (see <i>Line Separators</i>). Optional – the default is (13 10) on Microsoft Windows and 10 on other platforms.
<code>R ← □NREAD Y</code>	Reads the content of the native file identified by file tie number <i>Y</i> [1]; <i>Y</i> [2] specifies the conversion code to use (see <i>Conversion Codes</i>), <i>Y</i> [3] specifies the count (see <i>Conversion Codes</i>) and, optionally, <i>Y</i> [4] can define the offset from 0 of the first byte to read.
<code>{R} ← X □NRENAME Y</code>	Renames the tied native file that has file tie number <i>Y</i> to have name <i>X</i> .
<code>{R} ← X □NREPLACE Y</code>	Replaces content in a native file identified by file tie number <i>Y</i> [1] with <i>X</i> ; <i>Y</i> [2] defines the offset from 0 of the first byte to replace and, optionally, <i>Y</i> [3] specifies the conversion code to use (by default, 80 is assumed when using the Unicode version) (see <i>Conversion Codes</i>).

<code>{R} ← X □NRESIZE Y</code>	Changes the size of the native file that has file tie number <code>Y</code> to size <code>X</code> (either by truncating the file or by extending it with undefined additional bytes) .
<code>R ← □NSIZE Y</code>	Returns the size in bytes of the native file that has file tie number <code>Y</code> .
<code>{R} ← X □NTIE Y</code>	Ties the native file that has name <code>X</code> using file tie number <code>Y[1]</code> ; optionally, <code>Y[2]</code> can specify the type of access needed by other users (see <i>Access Codes</i>) .
<code>{R} ← □NUNTIE Y</code>	Unties all native files that have a tie number in vector <code>Y</code> (<code>□NUNTIE 0</code> does not untie any files but flushes all file caches to disk) and returns the tie numbers of native files that have been untied.
<code>{R} ←{X} □NXLATE Y</code>	Associates the native file that has tie number <code>Y</code> with character translation vector <code>X</code> . Note that: <ul style="list-style-type: none"> if <code>X</code> is not specified then the currently-associated translation vector is returned if <code>X</code> has the value <code>(⊔256) - □IO</code> then the translation process is bypassed and raw input/output is provided if <code>Y</code> is set to 0, then the translate vector used by <code>□DR</code> is used Unicode version only: This is only relevant when processing native files that contain characters expressed as indices into <code>□AV</code> .

Access Codes

The access codes used by `□NTIE` are integer values calculated as the sum of:

- the type of access needed from users who have already tied the native file
- the type of access to grant to users who subsequently try to open the file while you have it open

Needed from existing users	
0	read access
1	write access
2	read and write access

Granted to subsequent users	
0	compatibility mode
16	no access (exclusive)
32	read access
48	write access
64	read and write access

Conversion Codes

The conversion codes used by `□NAPPEND`, `□NREAD` and `□NREPLACE` vary according to the installation of Dyalog APL that is used to read the native file; the following two tables show the conversion codes for the Unicode version and Classic version respectively.

Value	Number of Bytes	Result Type	Result Shape
11	count	1 bit Boolean	8 x count
80	count	8 bit character	count
82*	count	8 bit character	count
83	count	8 bit integer	count
160	2 x count	16 bit character	count
163	2 x count	16 bit integer	count
320	4 x count	32 bit character	count
323	4 x count	32 bit integer	count
645	8 x count	64 bit floating	count

Value	Number of Bytes	Result Type	Result Shape
11	count	1 bit Boolean	8 x count
-	-	-	-
82	count	8 bit character	count
83	count	8 bit integer	count
-	-	-	-
163	2 x count	16 bit integer	count
-	-	-	-
323	4 x count	32 bit integer	count
645	8 x count	64 bit floating	count

* Conversion code 82 is permitted in the Unicode Edition for compatibility and causes 1-byte data on file to be *translated* (according to `□NXLATE`) from `□AV` indices into normal (Unicode) characters of type 80, 160 or 320.

File Encoding

The file encoding used by `OPENGET` and `OPENPUT`. The UTF formats can be qualified with -BOM (for example, UTF-8-BOM) or -NOBOM (for example, UTF-16LE-NOBOM) to specify whether a BOM is/should be present; this qualification is always present when returned by `OPENGET`.

Encoding	Description
UTF-8	Data is encoded into UTF-8 format. If -BOM or -NOBOM is not appended, the default is -NOBOM.
UTF-16 UTF-16BE UTF-16LE	Data is encoded into UTF-16 format with either big or little endianness. The default for UTF-16 is the endianness of the host system (BE on AIX platforms, LE on others). If -BOM or -NOBOM is not appended, the default is -BOM.
UTF-32 UTF-32BE UTF-32LE	Data is encoded into UTF-32 format with either big or little endianness. The default for UTF-32 is the endianness of the host system (BE on AIX platforms, LE on others). If -BOM or -NOBOM is not appended, the default is -BOM.
ASCII	Data is encoded into 7-bit ASCII format
Windows-1252	Data is encoded into 8-bit Windows-1252 format
ANSI	ANSI is a synonym of Windows-1252

Line Endings

The line terminators recognised by `OPENGET` and `OPENPUT`.

For `OPENGET`:

- if `R[1]` is simple, then it comprises the contents of the file with all line separators normalised to `UCS-10`.
- if `R[1]` is nested, then it comprises the contents of the file split on the occurrence of any of the line separators.

Value	Char	Description	Notes
13	CR	Carriage Return (U+000D)	Newline separators recognised by <code>OPENGET (R[3])</code> and <code>OPENPUT (X[3])</code> .
10	LF	Line Feed (U+000A)	
13 10	CRLF	Carriage Return followed by Line Feed	
133	NEL	New Line (U+0085)	
11	VT	Vertical Tab (U+000B)	
12	FF	Form Feed (U+000C)	
2028	LS	Line Separator (U+2028)	
2029	PS	Paragraph Separator (U+2029)	

In addition, `OPENPUT`'s NEOL variant specifies how embedded line endings are treated (0 = embedded line ending characters are preserved as is, 1 = every embedded LF is replaced by the host-specific line terminator (the default), and 2 = every embedded line ending character is replaced by the host-specific line terminator).

File Locking

Unlike component files, which can be tied with an exclusive tie or a share tie, native files cannot be tied in different ways. Instead, `OPENLOCK` is used to lock byte ranges within files, thereby managing access between users. There are three possible lock statuses:

- 1 means unlock
- 2 means read (share) lock – multiple read locks can exist over the same byte-range. The presence of a read lock prevents a write lock being obtained
- 3 means write lock – only one write lock can exist for a specific byte-range of a native file. The presence of a write lock prevents a read lock being obtained

The lock status can also, optionally, define a timeout period in seconds; if this period is exceeded before the lock status change has occurred, then a TIMEOUT error is displayed (defaults to no timeout limit).

Different file servers can follow different locking standards – `OPENLOCK` does not standardise this.