NAME

xxd − make a hex dump or do the reverse.

SYNOPSIS

xxd −h[elp]
xxd [options] [infile [outfile]]
xxd −r[evert] [options] [infile [outfile]]

DESCRIPTION

xxd creates a hex dump of a given file or standard input. It can also convert a hex dump back to its original binary form. Like **uuencode**(1) and **uudecode**(1) it allows the transmission of binary data in a 'mail-safe' ASCII representation, but has the advantage of decoding to standard output. Moreover, it can be used to perform binary file patching.

OPTIONS

If no *infile* is given, standard input is read. If *infile* is specified as a `**−**' character, then input is taken from standard input. If no *outfile* is given (or a `**−**' character is in its place), results are sent to standard output.

Note that a "lazy" parser is used which does not check for more than the first option letter, unless the option is followed by a parameter. Spaces between a single option letter and its parameter are optional. Parameters to options can be specified in decimal, hexadecimal or octal notation. Thus **−c8**, **−c 8**, **−c 010** and **−cols 8** are all equivalent.

−a | *−autoskip*
: Toggle autoskip: A single '*' replaces NUL-lines. Default off.

−b | *−bits*
: Switch to bits (binary digits) dump, rather than hex dump. This option writes octets as eight digits "1"s and "0"s instead of a normal hexadecimal dump. Each line is preceded by a line number in hexadecimal and followed by an ASCII (or EBCDIC) representation. The command line switch −p does not work with this mode. Can be combined with −i.

−c cols | *−cols cols*
: Format <*cols*> octets per line. Default 16 (−i: 12, −ps: 30, −b: 6). Max 256. No maximum for −ps. With −ps, 0 results in one long line of output.

−C | *−capitalize*
: Capitalize variable names in C include file style, when using −i.

−d
: show offset in decimal instead of hex.

−E | *−EBCDIC*
: Change the character encoding in the righthand column from ASCII to EBCDIC. This does not change the hexadecimal representation. The option is meaningless in combinations with −r, −p or −i.

−e
: Switch to little-endian hex dump. This option treats byte groups as words in little-endian byte order. The default grouping of 4 bytes may be changed using *−g*. This option only applies to the hex dump, leaving the ASCII (or EBCDIC) representation unchanged. The command line switches −r, −p, −i do not work with this mode.

−g bytes | *−groupsize bytes*
: Separate the output of every <*bytes*> bytes (two hex characters or eight bit digits each) by a whitespace. Specify *−g 0* to suppress grouping. <*Bytes*> defaults to *2* in normal mode, *4* in little-endian mode and *1* in bits mode. Grouping does not apply to PostScript or include style.

−h | *−help*
: Print a summary of available commands and exit. No hex dumping is performed.

−i | *−include*
: Output in C include file style. A complete static array definition is written (named after the input file), unless xxd reads from stdin. Can be combined with −b.

−l len | *−len len*
Stop after writing <*len*> octets.

−n name | *−name name*
Override the variable name output when −i is used. The array is named *name* and the length is named *name*_len.

−o offset
Add <*offset*> to the displayed file position.

−p | *−ps* | *−postscript* | *−plain*
Output in PostScript continuous hex dump style. Also known as plain hex dump style.

−r | *−revert*
Reverse operation: convert (or patch) hex dump into binary. If not writing to stdout, xxd writes into its output file without truncating it. Use the combination *−r −p* to read plain hexadecimal dumps without line number information and without a particular column layout. Additional white-space and line breaks are allowed anywhere. Use the combination *−r −b* to read a bits dump instead of a hex dump.

−R when
In the output the hex-value and the value are both colored with the same color depending on the hex-value. Mostly helping to differentiate printable and non-printable characters. *when* is **never**, **always**, or **auto** (default: auto). When the **$NO_COLOR** environment variable is set, colorization will be disabled.

−seek offset
When used after *−r*: revert with <*offset*> added to file positions found in hex dump.

−s [+][−]seek
Start at <*seek*> bytes abs. (or rel.) infile offset. + indicates that the seek is relative to the current stdin file position (meaningless when not reading from stdin). − indicates that the seek should be that many characters from the end of the input (or if combined with +: before the current stdin file position). Without −s option, xxd starts at the current file position.

−u Use upper-case hex letters. Default is lower-case.

−v | *−version*
Show version string.

CAVEATS

xxd −r has some built-in magic while evaluating line number information. If the output file is seekable, then the line numbers at the start of each hex dump line may be out of order, lines may be missing, or overlapping. In these cases xxd will lseek(2) to the next position. If the output file is not seekable, only gaps are allowed, which will be filled by null-bytes.

xxd −r never generates parse errors. Garbage is silently skipped.

When editing hex dumps, please note that *xxd −r* skips everything on the input line after reading enough columns of hexadecimal data (see option −c). This also means that changes to the printable ASCII (or EBCDIC) columns are always ignored. Reverting a plain (or PostScript) style hex dump with xxd −r −p does not depend on the correct number of columns. Here, anything that looks like a pair of hex digits is interpreted.

Note the difference between
% xxd −i file
and
% xxd −i < file

xxd −s +seek may be different from *xxd −s seek*, as lseek(2) is used to "rewind" input. A '+' makes a difference if the input source is stdin, and if stdin's file position is not at the start of the file by the time xxd is started and given its input. The following examples may help to clarify (or further confuse!):

Rewind stdin before reading; needed because the ‘cat’ has already read to the end of stdin.
% sh −c "cat > plain_copy; xxd −s 0 > hex_copy" < file

Hex dump from file position 0x480 (=1024+128) onwards. The ‘+’ sign means "relative to the current position", thus the ‘128’ adds to the 1k where dd left off.
% sh −c "dd of=plain_snippet bs=1k count=1; xxd −s +128 > hex_snippet" < file

Hex dump from file position 0x100 (=1024−768) onwards.
% sh −c "dd of=plain_snippet bs=1k count=1; xxd −s +−768 > hex_snippet" < file

However, this is a rare situation and the use of ‘+’ is rarely needed. The author prefers to monitor the effect of xxd with strace(1) or truss(1), whenever −s is used.

EXAMPLES

Print everything but the first three lines (hex 0x30 bytes) of **file**.
% xxd −s 0x30 file

Print 3 lines (hex 0x30 bytes) from the end of **file**.
% xxd −s −0x30 file

Note: The results of the examples below are relevant to the xxd.1 man page as of May 2024

Print 120 bytes as a continuous hex dump with 20 octets per line.
% xxd −l 120 −ps −c 20 xxd.1
2e5448205858442031203120224d6179203230323422
20224d616e75616c207061676520666f72207878
64220a2e5c220a2e5c222032317374204d617920
313939360a2e5c22204d616e2070616765206175
74686f723a0a2e5c2220202020546f6e79204e75
67656e74203c746f6e79407363746e7567656e2e

Hex dump the first 120 bytes of this man page with 12 octets per line.
% xxd −l 120 −c 12 xxd.1
00000000: 2e54 4820 5858 4420 3120 224d .TH XXD 1 "M
0000000c: 6179 2032 3032 3422 2022 4d61 ay 2024" "Ma
00000018: 6e75 616c 2070 6167 6520 666f nual page fo
00000024: 7220 7878 6422 0a2e 5c22 0a2e r xxd"..\"..
00000030: 5c22 2032 3173 7420 4d61 7920 \" 21st May
0000003c: 3139 3936 0a2e 5c22 204d 616e 1996..\" Man
00000048: 2070 6167 6520 6175 7468 6f72 page author
00000054: 3a0a 2e5c 2220 2020 2054 6f6e :..\" Ton
00000060: 7920 4e75 6765 6e74 203c 746f y Nugent <to
0000006c: 6e79 4073 6374 6e75 6765 6e2e ny@sctnugen.

Display just the date from the file xxd.1
% xxd −s 0x33 −l 13 −c 13 xxd.1
00000033: 3231 7374 204d 6179 2031 3939 36 21st May 1996

Copy **input_file** to **output_file** and prepend 100 bytes of value 0x00.
% xxd input_file | xxd −r −s 100 > output_file

Patch the date in the file xxd.1
% echo "0000034: 3574 68" | xxd −r − xxd.1
% xxd −s 0x33 −l 13 −c 13 xxd.1
00000033: 3235 7468 204d 6179 2031 3939 36 25th May 1996

Create a 65537 byte file with all bytes 0x00, except for the last one which is ’A’ (hex 0x41).
% echo "010000: 41" | xxd −r > file

Hex dump this file with autoskip.
% xxd −a −c 12 file

00000000: 0000 0000 0000 0000 0000 0000 .....
*
0000fffc: 0000 0000 41 ....A

Create a 1 byte file containing a single 'A' character. The number after '−r −s' adds to the line numbers found in the file; in effect, the leading bytes are suppressed.
% echo "010000: 41" | xxd −r −s −0x10000 > file

Use xxd as a filter within an editor such as **vim(1)** to hex dump a region marked between 'a' and 'z'.
:'a,'z!xxd

Use xxd as a filter within an editor such as **vim(1)** to recover a binary hex dump marked between 'a' and 'z'.
:'a,'z!xxd −r

Use xxd as a filter within an editor such as **vim(1)** to recover one line of a hex dump. Move the cursor over the line and type:
!!xxd −r

Read single characters from a serial line
% xxd −c1 < /dev/term/b &
% stty < /dev/term/b −echo −opost −isig −icanon min 1
% echo −n foo > /dev/term/b

RETURN VALUES

The following error values are returned:

0 no errors encountered.

−1 operation not supported (*xxd −r −i* still impossible).

1 error while parsing options.

2 problems with input file.

3 problems with output file.

4,5 desired seek position is unreachable.

SEE ALSO

uuencode(1), uudecode(1), patch(1)

WARNINGS

The tool's weirdness matches its creator's brain. Use entirely at your own risk. Copy files. Trace it. Become a wizard.

VERSION

This manual page documents xxd version 1.7 from 2024-05.

AUTHOR

(c) 1990-1997 by Juergen Weigert
<jnweiger@informatik.uni−erlangen.de>

Distribute freely and credit me,
make money and share with me,
lose money and don't ask me.

Manual page started by Tony Nugent
<tony@sctnugen.ppp.gu.edu.au> <T.Nugent@sct.gu.edu.au>
Small changes by Bram Moolenaar. Edited by Juergen Weigert.