

Documentation for LZEXE, EXE file packer.

The ecm fork of LZEXE v0.91 is based on the 2025 May free software release under the MIT license. Refer to the LICENSE file for full attribution and usage conditions.

Hardware: PC and compatibles, 80286 or 80386 microprocessor recommended for greater execution speed. Memory required: 128 KiB minimum to run LZEXE.

This document has been compiled on 2025-06-26.

Contents

| | |
|--|----|
| Section 1: Introduction | 4 |
| Section 2: Using LZEXE | 5 |
| 2.1 Running LZEXE | 5 |
| 2.2 Switches for LZEXE application | 5 |
| 2.2.1 Switch /1 | 6 |
| 2.2.2 Switch /2 | 6 |
| 2.2.3 Switch /I | 6 |
| 2.2.4 Switch /J | 6 |
| 2.2.5 Switch /S | 6 |
| 2.2.6 Switch /O | 6 |
| 2.2.7 Switch /L | 7 |
| 2.2.8 Switch /- | 7 |
| 2.2.9 Switch / | 7 |
| 2.2.10 Switch /#XX | 7 |
| Section 3: Tips for use | 8 |
| Section 4: The unlzexe utility | 10 |
| 4.1 unlzexe debugging flags | 10 |
| Section 5: From a technical point of view (for those in the know!) | 11 |
| 5.1 Optimisations in the ecm fork | 12 |
| Section 6: LZEXE version 0.91 and other compressors | 13 |
| Section 7: The future... | 14 |
| Section 8: Warnings and wishes... | 15 |
| Section 9: Evolution of versions | 16 |
| 9.1 LZEXE ecm release 3 (future) | 16 |
| 9.2 LZEXE ecm release 2 | 16 |

| | | |
|----------------------------|-----------|----|
| 9.3 LZEXE ecm release 1 | | 16 |
| 9.4 LZEXE ecm release 0 | | 17 |
| 9.5 LZEXE v0.91 | | 17 |
| Source Control Revision ID | | 18 |

Section 1: Introduction

This software compresses EXE files, that is, EXEcutable files from the 8086 DOS PC world. But you could tell me that a lot of software compresses EXE files better than this one, such as the excellent PKZIP or LHARC. But the advantage of this program is that your EXE files once compressed can be launched! And the depacking is so fast that for virtually all files, this depacking time is negligible! In addition, the depacker does not use any additional disk space or RAM on a virtual disk: it only uses the RAM normally reserved for the unpacked EXE file. In addition, I have greatly optimized my depack algorithm in speed but also in efficiency: EXE files are almost as small as the corresponding ZIP files and much more compact than the old ARC files.

Section 2: Using LZEXE

2.1 Running LZEXE

It's very simple: just type in DOS:

```
LZEXE [switches] filename[.EXE]
```

where 'filename' is the name of the EXE file you want to compress. The .EXE extension is added by default. The compressed file is created in the default directory. The 'switches' part is optional and may hold any number of switches. Switches are indicated by a leading slash. Refer to section 2.2 for the meaning of the switches.

Warning! Some files are EXE only by their name: in fact, for DOS, it is not the .EXE extension that characterizes this type of file, it is the fact that at the beginning there are the letters 'MZ' or 'ZM' followed by a few bytes that indicate the length of the file, the memory size that it occupies, etc. So some people do not hesitate to rename COM files to EXE, and this explains why LZEXE refuses certain EXE files that are just renamed COM files.

But there is a method to make LZEXE accept COM files: just use the LZEXE companion tool COMTOEXE which converts COM files to EXE files.

For added security, LZEXE does not erase your old EXE file: it renames it to *.OLD. In addition, it creates the temporary file LZEXE.TMP which is only renamed to *.EXE after packing is complete.

2.2 Switches for LZEXE application

The following switches are supported:

/1

Choose old v0.91 stub format

/2

Choose new stub format (default)

/I

Inline getbit code alike old stub (faster)

/J

Do not inline getbit code (slower, default)

/S

Stop optimisation, always use LZE3 equivalent

/O

Optimise to drop relocs or segment change (default)

/L

Allow output file that is not smaller than input

/-

No more switches, filename follows

//

Same as /-

/#XX

Force LZ signature letters to XX (for debugging)

2.2.1 Switch /1

Chooses to pack with the old v0.91 stub format. This generates a file largely compatible with the output of LZEXE v0.91 albeit not identicalised to it. The length of the depacker stub and the placement of its variables do exactly match.

2.2.2 Switch /2

Chooses to pack with the new stub format. (This is the default.) The current format is known as LZX0, and comes in 8 variants. This format is not compatible with any older depackers expecting the LZEXE v0.90 or v0.91 formats.

2.2.3 Switch /I

Inline getbit code alike old stub. Inlining the getbit code makes depacking faster at the expense of some code size. In the old format, the code was inlined.

2.2.4 Switch /J

Do not inline getbit code. This is slower than inlining but saves some space. (This is the default.)

2.2.5 Switch /s

Stop optimisation, always use LZE3 equivalent. This disables the dropping of the relocation table if empty, and of the segment change code if the uncompressed image size is below 40 KiB.

2.2.6 Switch /o

Optimise to drop relocs or segment change. Different variants of the LZX0 depacker stub can be used to optimise stub size. This switch enables two different optimisations. The LZEXE application will detect if either or both of these optimisations can be used. (This is the default.)

2.2.7 Switch /L

Allow output file that is not smaller than input. Usually LZEXE will delete its temporary output file if it is at least as large as the input file. With this switch, it will only warn about the condition but still write to the destination file.

2.2.8 Switch /-

No more switches, filename follows. This switch indicates that the next input on the command line is a filename, even if it starts with a dash or slash. This avoids misdetecting such names as switches.

2.2.9 Switch //

Same as /-.

2.2.10 Switch /#xx

Force LZ signature letters to xx (for debugging). The xx can be replaced by any two printable ASCII bytes. The specified bytes are appended to the LZ signature. This overrides the LZ91 or LZX0 signatures usually chosen by the current LZEXE. This switch should not be used lightly and is intended only for debugging LZEXE. It can lead to data corruption if misused.

Section 3: Tips for use

For some files, compression may not work for several reasons:

- The file you specified is not a real EXE. Solution: use COMTOEXE.EXE
- The relocation table is too large. To understand this message, it's necessary to understand the internal structure of an EXE file: such a file can span multiple segments, unlike COM files. This is why it needs a table of values that indicates in which segment branches or subroutine calls are made, for example. And if the file is very long, this table can be very bulky and prevent the packer from working. However, I have allowed 16,000 relocation addresses, which should be sufficient for all EXE files, even the largest.
- The file you specified has already been compressed with LZEXE.

Note that another file packer exists: EXEPACK.EXE from Microsoft. But it is far less efficient than mine, and even if your EXE file is already packed with this program, LZEXE will still be able to pack a lot. But in this case, you'll be presented with a warning message, because there's another LZEXE companion tool: UPACKEXE, which allows you to unpack these files, and thus the gains are even greater.

- The packing was not efficient enough and wasted space on disk: Yes, it can happen, but generally with small EXE files (less than 2 KiB). Otherwise, you can almost always win a few bytes. If the output file is at least as large as the input file, LZEXE will abort the operation and delete the temporary output file, unless the /L switch was specified.
- The EXE file contains internal overlays: these are pieces of program files that are in the EXE file but are loaded only when the main program needs them. LZEXE cannot pack them because it would be necessary to modify the loading routines which are in the main program, and unfortunately these routines depend on the compiler and programmer. In version 0.91, LZEXE warns you about their presence. But in some cases, the difference between the length of the EXE file on disk and the length of the loaded code is minimal (less than 1024 bytes): in this case, you can still choose to pack because these may be "marks" left by some compilers.
- In the ecm fork of LZEXE, a bug in LZEXE v0.91's relocation table format is fixed. The format cannot cope with a relocation at the very beginning of the executable program image, or more than one relocation at the exact same address. If the /1 switch is used, indicating to use LZ91 format, the new LZEXE will detect executables with these and reject them rather than producing a depacker that crashes or corrupts random memory. If the LZX0 format is used these unusual relocations are fully supported.
- Some invalid EXE files are rejected. This includes files with zero pages, zero or negative sized images (too large header size), or relocations addressing space beyond the end of the image.

More serious is that some compressed EXE files will "crash" the machine:

- If the EXE program performs a test on its size or integrity on disk (this is very rare).
- If it contains overlays, which must be loaded afterwards and therefore must occupy fixed positions in the file.
- For programs that run under Windows (from Microsoft): these are not real 8086 DOS EXEs, so they will refuse to work properly when packed.

(This list is not exhaustive.)

Less serious: Some programs have configuration options that modify the EXE file (Turbo Pascal for example). In this case, you must first configure the program, then compress it and keep an uncompressed version so you can edit it.

Section 4: The unzexe utility

The companion tool unzexe can reverse the LZEXE compression and generate an unpacked executable. It can handle all file formats generated by this version of LZEXE, except if the header signature was changed using the /# switch (refer to section 2.2.10).

Note that the exact header size and the order and addressing in the relocation entries may differ as compared to the original file before LZEXE compression. Additional data such as overlays or IDOS iniload's entries cannot be restored either, as they are not retained by LZEXE. The memory allocation and image size should be reproduced exactly. The relocation entries should lead to the same result, but as mentioned may be re-ordered and use different segmented addresses.

Compressing the unzexe output using LZEXE (same version with the same switches) should produce exactly the same compressed file as the unzexe input.

4.1 unzexe debugging flags

The unzexe tool accepts flags from the DEBUG environment variable. (%DEBUG% on DOS, \$DEBUG on Linux.) The following flags are accepted in a mask:

1

Enable all debugging output.

2

If alloc delta field is present (not in current LZX0 format) then display both the old-style (unzexe v0.7) and new-style resulting minimum allocation. These should match.

4

Enable listing all code snippet patterns info. This also allows to determine which optimisations of the LZX0 format stub were used by LZEXE.

8

Enable dump of MZ EXE header of input and output. The display is in both hexadecimal and decimal.

16

Enable listing of input and output file sizes. The display is in both hexadecimal and decimal.

Section 5: From a technical point of view (for those in the know!)

The compression algorithm I made is based on the famous Lempel Ziv method using a "circular" buffer (ring buffer) and a method of finding repetitions of byte sequences by trees. The coding of the position and length of the strings that repeat is optimized by an additional algorithm inspired by the Huffman coding method. Unpacked literal bytes are sent as is in the file. An additional compression algorithm (like "Adaptive Huffman" (see LHARC) or with Shanon-Fano trees (see PKZIP)) would have required a longer unpacking time and above all a more efficient depacker that is more complex and longer, which could actually make the compressed EXE file longer.

The depacker is located at the end of the EXE file and is 395 bytes in size long for version 0.90 and 330 bytes for version 0.91. In the ecm fork the depacker is between 208 and 305 bytes in size. The depacker must:

- Check the CRC to make sure no one has modified it (useful against viruses). If yes, display the message: 'CRC Error'. This option was removed in LZEXE v0.91 because it unnecessarily lengthens the file EXE and the depack time. Furthermore, the CRC check was only performed on the depacker.
- Move to the top of RAM, then move the compressed data to leave some space for the EXE image.
- Depack the code, checking that it is correct, and also adjust segments if you exceed 64 KiB (which caused me problems in terms of speed).

In the LZX0 format of the ecm fork, the segment change handling can be omitted in case the original executable image was smaller than 40 KiB, saving some space for otherwise unused code in the depacker.

- Depack the relocation table, and update the relocation addresses in the EXE image. This is where LZEXE v0.91 was changed: the relocation table is much better compressed.

In the LZX0 format of the ecm fork, the relocation table changed slightly. This supports a wider range of relocation entries.

Further, LZX0 format can be used with the relocation table handling omitted in case the original EXE file had an empty relocation table.

- Before the subsequent transfer, set up AX (FCB drive validity) and DS, ES (both point to the PSP) the same way we received them from the DOS's program loader.
- Launch the program by updating CS,IP,SS,SP

That's all!!!

This depacker is a small masterpiece of 8086 assembler programming in itself: needless to say, it took quite a long time to develop. But the packer also posed quite a few problems for me, particularly when it came to updating all the pointers that the depacker uses later.

5.1 Optimisations in the ecm fork

I optimised the depacker some. In addition to the universal optimisations, LZX0 format can select one of eight variants of the depacker:

- With or without relocation table.
- With or without segment change handling.
- With getbit handling inlined or in a function.

The `/O` switch enables the optimisation of the relocation table and segment change code, whereas the `/S` switch stops these optimisations. The `/J` switch disables inlining getbit whereas the `/I` switch enables it.

Section 6: LZEXE version 0.91 and other compressors

PKARC (latest version): LZEXE performs much better, as "crunching" (aka shrinking for PKZIP) is an outdated algorithm...

PKZIP v0.92: LZEXE does better in almost all cases.

PKZIP v1.02: On large files, LZEXE does better. otherwise, the the difference is quite small.

LHARC v1.01: It does better than LZEXE with "freezing" on small files.

LARC: LZEXE does better.

Important Notes:

- You can't really compare what LZEXE does with other packers since in the EXE files compressed by my software there is also a depacker that launches the EXE by itself. The other compressors can make "self-extracting" files, but they depack to disk, are slow and add several tens of KiB to the compressed files (except for LARC and LHARC which only add 1 or 2 KiB, but which only depack on disk unfortunately).
- In almost all cases, the compressors I mentioned will not be able to further pack a file already reduced with LZEXE, which shows its efficiency. Only LHARC manages to save a few bytes.

UPX: `upx-uc1 --8086 --lzma` does about 6% better than LZEXE. Like LZEXE it produces files with an online depacker stub in the compressed executable file.

Section 7: The future...

- I'm also thinking about an automatic documentation depacker like README.COM or LIST.COM, which would be very convenient, and maybe I'll make one.
- Finally, I hope to make a "universal" packer like PKZIP or LHARC, which is slower than LZEXE when it comes to unpacking, but performs much better than them.

Section 8: Warnings and wishes...

I hope that LZEXE and the EXE files compressed by it will be widely broadcast which will encourage me to make other versions more quickly...

I decline all responsibility in the event of loss of information caused by LZEXE. But rest assured, the algorithms are reliable and I don't think there are many bugs.

Warning! I do not recommend compressing and distributing commercial software protected by copyright: the authors may be displeased...

But if you're making a FREEWARE, SHAREWARE, or even a commercial program, nothing's stopping you from compressing it with LZEXE, and I even recommend it:

- Your EXE files will be smaller, and your compiled programs will look like they're made in assembler. What will your competitors say when they see programs that do the same thing as theirs but are 30% smaller? Plus, you'll be able to fit more programs on floppy disk (and hard disk), because we always need more mass storage...
- Compression is an excellent form of coding that can prevent unsavory people from modifying messages or seeing your secret algorithms unless they disassemble the depacker, which might not be very easy, I'm telling you!

There you go, hoping that this software will be useful to you and that it does not have too many bugs!

Section 9: Evolution of versions

9.1 LZEXE ecm release 3 (future)

- Fix UPACKEXE if a relocation with offset 0FFFFh appears in the exepacked relocation table.
- Translate and update this manual.
- unlzexe dot mode: If second parameter is a lone dot, only calculate everything without writing to any file.

9.2 LZEXE ecm release 2

- Add formats without inlined getbit function. This is purely a speed vs size tradeoff.
- Add debugging outputs to unlzexe, selected using the DEBUG environment variable.
- Make unlzexe able to restore the minimum allocation size of the original unpacked executable file, even if alloc delta field absent.
- Add formats without segment change handling.
- Add format without relocation table.
- Flush the pipeline for depacker's Self Modifying Code.

9.3 LZEXE ecm release 1

- Do not calculate percentages with the 8087 FPU.
- Port all Pascal source texts to allow building with FreePascal, i8086 DOS target.
- Add alloc delta field to depacker stub so unlzexe works better.
- Detect invalid relocations (past image).
- Add new relocation table format which allows zero-difference relocations.
- Add new depacker, optimise it, and pass along the value in AX from the DOS program loader to the depacked program.
- Check for nonzero positive image size.
- If unknown LZ signature is found, display it.
- Make unlzexe compile for either 8086 DOS (gcc ia16) or amd64 Linux (gcc).
- Add unlzexe from LZEXE v0.91e, the former is Public Domain.

- Port packer to NASM.
- In comtoexe add a stub to set up a zero word on top of the stack so that programs depending on it won't break.
- Detect and warn about invalid amount last page bytes field (> 512) and display a hint on IDOS iniload if the low byte of this field equals 0EBh.
- Detect and reject zero in EXE header amount pages field.
- Translate most of the main LZEXE tool to the english language.

9.4 LZEXE ecm release 0

- Reject zero-difference relocations when packing a file with the LZ91 format.
- Display remaining file size and percentage.
- Port depacker to NASM.

9.5 LZEXE v0.91

- The hyphen ‘-’ was not accepted in file names, this has been corrected.
- LZEXE signals the presence of internal overlays.
- LZEXE indicates whether the file has already been packed with EXEPACK by Microsoft.
- A bug in version 0.90 caused EXE files to take up too much memory: this has been fixed.
- Relocation table compression has been improved.
- CRC check has been removed.
- Decompressor size reduced from 395 to 330 bytes.

Source Control Revision ID

hg bc51fd4e3e49, from commit on at 2025-06-26 22:12:45 +0200

If this is in ecm's repository, you can find it at
<https://hg.pushbx.org/ecm/lzexe/rev/bc51fd4e3e49>