

# [MS-PATCH]: LZX DELTA Compression and Decompression

---

## Intellectual Property Rights Notice for Open Specifications Documentation

- **Technical Documentation.** Microsoft publishes Open Specifications documentation for protocols, file formats, languages, standards as well as overviews of the interaction among each of these technologies.
- **Copyrights.** This documentation is covered by Microsoft copyrights. Regardless of any other terms that are contained in the terms of use for the Microsoft website that hosts this documentation, you may make copies of it in order to develop implementations of the technologies described in the Open Specifications and may distribute portions of it in your implementations using these technologies or your documentation as necessary to properly document the implementation. You may also distribute in your implementation, with or without modification, any schema, IDL's, or code samples that are included in the documentation. This permission also applies to any documents that are referenced in the Open Specifications.
- **No Trade Secrets.** Microsoft does not claim any trade secret rights in this documentation.
- **Patents.** Microsoft has patents that may cover your implementations of the technologies described in the Open Specifications. Neither this notice nor Microsoft's delivery of the documentation grants any licenses under those or any other Microsoft patents. However, a given Open Specification may be covered by Microsoft's Open Specification Promise (available here: <http://www.microsoft.com/interop/osp>) or the Community Promise (available here: <http://www.microsoft.com/interop/cp/default.msp>). If you would prefer a written license, or if the technologies described in the Open Specifications are not covered by the Open Specifications Promise or Community Promise, as applicable, patent licenses are available by contacting [iplg@microsoft.com](mailto:iplg@microsoft.com).
- **Trademarks.** The names of companies and products contained in this documentation may be covered by trademarks or similar intellectual property rights. This notice does not grant any licenses under those rights.
- **Fictitious Names.** The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted in this documentation are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

**Reservation of Rights.** All other rights are reserved, and this notice does not grant any rights other than specifically described above, whether by implication, estoppel, or otherwise.

**Tools.** The Open Specifications do not require the use of Microsoft programming tools or programming environments in order for you to develop an implementation. If you have access to Microsoft programming tools and environments you are free to take advantage of them. Certain Open Specifications are intended for use in conjunction with publicly available standard specifications and network programming art, and assumes that the reader either is familiar with the aforementioned material or has immediate access to it.

## Revision Summary

Date	Revision History	Revision Class	Comments
04/04/2008	0.1		Initial Availability.
06/27/2008	1.0		Initial Release.
08/06/2008	1.01		Revised and edited technical content.
09/03/2008	1.02		Revised and edited technical content.
12/03/2008	1.03		Updated IP notice.
03/04/2009	1.04		Revised and edited technical content.
04/10/2009	2.0		Updated technical content and applicable product releases.
07/15/2009	3.0	Major	Revised and edited for technical content.
11/04/2009	3.0.1	Editorial	Revised and edited the technical content.
02/10/2010	3.1.0	Minor	Updated the technical content.
05/05/2010	3.1.1	Editorial	Revised and edited the technical content.
08/04/2010	4.0	Major	Significantly changed the technical content.
11/03/2010	4.0	No change	No changes to the meaning, language, or formatting of the technical content.
03/18/2011	4.0	No change	No changes to the meaning, language, or formatting of the technical content.
08/05/2011	4.0	No change	No changes to the meaning, language, or formatting of the technical content.

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
1.1	Glossary	5
1.2	References	5
1.2.1	Normative References	5
1.2.2	Informative References	5
1.3	Overview	5
1.4	Relationship to Protocols and Other Structures	6
1.5	Applicability Statement	6
1.6	Versioning and Localization	6
1.7	Vendor-Extensible Fields	6
<b>2</b>	<b>Structures</b>	<b>7</b>
2.1	Concepts	7
2.1.1	Bitstream	7
2.1.2	Window Size	7
2.1.3	Reference Data	7
2.1.4	Repeated Offsets	8
2.1.5	Match Lengths	9
2.1.6	Position Slot	9
2.2	Header	10
2.2.1	Chunk Size	10
2.2.2	E8 Call Translation	10
2.3	Block	11
2.3.1	Block Header	11
2.3.1.1	Block Type	12
2.3.1.2	Block Size	12
2.3.2	Block Data	12
2.3.2.1	Uncompressed Block	12
2.3.2.2	Verbatim Block	13
2.3.2.3	Aligned Offset Block	13
2.4	Huffman Trees	14
2.5	Encoding the Trees and Pre-Trees	14
2.6	Compressed Token Sequence	15
2.6.1	Converting Match Offset into Formatted Offset Values	16
2.6.2	Converting Formatted Offset into Position Slot and Position Footer Values	17
2.6.3	Converting Position Footer into Verbatim Bits or Aligned Offset Bits	18
2.6.4	Converting Match Length into Length Header and Length Footer Values	19
2.6.5	Converting Length Header and Position Slot into Length/Position Header Values	20
2.6.6	Extra Length	20
2.6.7	Encoding a Match	20
2.6.8	Encoding a Literal	21
2.7	Decoding Matches and Literals (Aligned and Verbatim Blocks)	21
<b>3</b>	<b>Structure Examples</b>	<b>23</b>
<b>4</b>	<b>Security Considerations</b>	<b>24</b>
<b>5</b>	<b>Appendix A: Product Behavior</b>	<b>25</b>
<b>6</b>	<b>Change Tracking</b>	<b>26</b>



# 1 Introduction

Delta compression is a technique in which one set of data can be compressed within the context of a reference set of data that is supplied both to the compressor and decompressor. Delta compression is commonly used to encode updates to similar existing data sets so that the size of compressed data can be significantly reduced relative to ordinary non-delta compression techniques. Expanding a delta-compressed set of data requires that the exact same reference data be provided during decompression.

## 1.1 Glossary

The following terms are defined in [\[MS-GLOS\]](#):

### little-endian

The following terms are defined in [\[MS-OXGLOS\]](#):

### header

### Lempel-Ziv Extended (LZX)

### Lempel-Ziv Extended Delta (LZXD)

### stream

The following terms are specific to this document:

**MAY, SHOULD, MUST, SHOULD NOT, MUST NOT:** These terms (in all caps) are used as described in [\[RFC2119\]](#). All statements of optional behavior use either MAY, SHOULD, or SHOULD NOT.

## 1.2 References

### 1.2.1 Normative References

We conduct frequent surveys of the normative references to assure their continued availability. If you have any issue with finding a normative reference, please contact [dochelp@microsoft.com](mailto:dochelp@microsoft.com). We will assist you in finding the relevant information. Please check the archive site, <http://msdn2.microsoft.com/en-us/library/E4BD6494-06AD-4aed-9823-445E921C9624>, as an additional source.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <http://www.rfc-editor.org/rfc/rfc2119.txt>

[UASDC] Ziv, J. and Lempel, A., "A Universal Algorithm for Sequential Data Compression", May 1977, <http://ieeexplore.ieee.org/iel5/18/22696/01055714.pdf>

### 1.2.2 Informative References

[MS-GLOS] Microsoft Corporation, "[Windows Protocols Master Glossary](#)".

[MS-OXGLOS] Microsoft Corporation, "[Exchange Server Protocols Master Glossary](#)".

## 1.3 Overview

**LZXD** (D for Delta) is an **LZX** variant that is modified to facilitate efficient delta-compression. LZXD provides a mechanism for both compressor and decompressor to refer to a common reference set of data, and relaxes the constraint that match offset be constrained to less than the current position in

the output **stream**, allowing match offset to refer to the logically prepended reference data. This effectively enables the compressed data stream to encode "matches" both from the reference data and from the uncompressed data stream.

#### **1.4 Relationship to Protocols and Other Structures**

LZ77 refers to the well-known Lempel-Ziv 1977 sliding window data compression algorithm, as specified in [\[UASDC\]](#).

LZX is an LZ77-based compressor that uses static Huffman encoding and a sliding window of selectable size. Data symbols are encoded either as an uncompressed symbol, or as a logical (offset, length) pair indicating that length symbols shall be copied from a displacement of offset symbols from the current position in the output stream. The value of offset is constrained to be less than the current position in the output stream, up to the size of the sliding window.

#### **1.5 Applicability Statement**

None.

#### **1.6 Versioning and Localization**

None.

#### **1.7 Vendor-Extensible Fields**

None.

## 2 Structures

LZXD compressed data consists of a **header** that indicates the file translation size, followed by a sequence of compressed blocks. A stream of uncompressed input can be output as multiple compressed LZXD blocks to improve compression, because each compressed block contains its own statistical tree structures.



A block can be one of following types:


- Uncompressed
- Verbatim
- Aligned offset

The structure of these blocks is specified in sections [2.3.2.1](#), [2.3.2.2](#), and section [2.3.2.3](#).

### 2.1 Concepts

#### 2.1.1 Bitstream

An LZXD Bitstream is encoded as a sequence of aligned 16-bit integers stored in the order least-significant-byte most-significant-byte, also known as byte-swapped or **little-endian** words. Given an input stream of bits named a, b, c,..., x, y, z, A, B, C, D, E, F, the output byte stream (with byte boundaries highlighted) **MUST** be as follows:



#### 2.1.2 Window Size

The sliding window size **MUST** be a power of 2, from 217 (128 KB) up to 225 (32 MB). The window size is not stored in the compressed data stream, and **MUST** be specified to the decoder before decoding begins. The window size **SHOULD** be the smallest power of two between 217 and 225 that is greater than or equal to the sum of the size of the reference data rounded up to a multiple of 32,768 and the size of the subject data.

#### 2.1.3 Reference Data

For delta compression, the reference data is a sequence of bytes given to the compressor prior to compressing the subject data. The exact same reference data sequence **MUST** be given to the decompressor prior to decompression. The reference data sequence is treated as logically prepended to the subject data sequence being compressed or decompressed. During decompression, match offsets are negative displacements from the "current position" in the output stream, up to the specified Window Size. When match offset values exceed the number of bytes already emitted in the uncompressed output stream, they are simply pointing into the reference data that is logically prepended to the subject data.

Offset	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
Value	A	B	C	D	E	F	G	H	I	J	a	b	c	D	E	F	a	b	c	e
	Reference Data Sequence										Subject Data Sequence									

In this example, the reference data is 10 bytes long and consists of the sequence "ABCDEFGHIJ". The data to be compressed, or the subject data, is also 10 bytes long (although the data does not have to be the same length as the reference data) and consists of "abcDEFabce". A valid encoded sequence would consist of the following tokens:

'a', 'b', 'c', (match offset -10, length 3), (match offset -6, length 3), 'e'

The first match offset exceeds the amount of subject data already in the window, pointing instead into the reference data portion. The second match offset does not exceed the amount of subject data in the window and instead refers to a portion of the subject data previously compressed or decompressed.

### 2.1.4 Repeated Offsets

LZXD extends the conventional LZ77 format in several ways, one of which is in the use of repeated offset codes. Three match offset codes, named the repeated offset codes, are reserved to indicate that the current match offset is the same as that of one of the three previous matches, which is not itself a repeated offset.

The three special offset codes are encoded as offset values 0, 1, and 2 (for example, encoding an offset of 0 means "use the most recent non-repeated match offset", an offset of 1 means "use the second most recent non-repeated match offset", and so on). All remaining Encoded offset values are displaced by Real offset +2, as is shown in the following table, which prevents matches at offsets WINDOW\_SIZE, WINDOW\_SIZE-1, and WINDOW\_SIZE-2.

Encoded offset	Real offset
0	Most recent real match offset
1	Second most recent match offset
2	Third most recent match offset
3	1 (closest allowable)
4	2
5	3
6	4
7	5
8	6
500	498
X+2	X
WINDOW_SIZE-1 (maximum possible)	WINDOW_SIZE-3



The three most recent real match offsets are kept in a list, the behavior of which is explained as follows:

Let R0 be defined as the most recent real offset

Let R1 be defined as the second most recent offset

Let R2 be defined as the third most recent offset

The list is managed similarly to an LRU (least recently used) queue, with the exception of the cases when R1 or R2 is output. In these cases, R1 or R2 is simply swapped with R0, which requires fewer operations than an LRU queue would.

The initial state of R0, R1, R2 is (1, 1, 1).

Match offset X where...	Operation
X≠R0 and X≠R1 and X≠R2	R2←R1 R1←R0 R0←X
X = R0	None
X = R1	swap R0↔R1
X = R2	swap R0↔R2

### 2.1.5 Match Lengths

The minimum match length (number of bytes) encoded by LZXD is 2 bytes, and the maximum match length is 32,768 bytes. However, no match of any length can span a modulo-32 KB boundary in the uncompressed stream. Match length encoding is combined with match position encoding as described in section [2.6](#).

### 2.1.6 Position Slot

The window size determines the number of window subdivisions, or "position slots", as shown in the following table.

Window size	Position slots required
128 KB	34
256 KB	36
512 KB	38
1 MB	42
2 MB	50
4 MB	66
8 MB	98
16 MB	162

Window size	Position slots required
32 MB	290

## 2.2 Header

### 2.2.1 Chunk Size

The LZXD compressor emits chunks of compressed data. A chunk represents exactly 32 KB of uncompressed data until the last chunk in the stream, which can represent less than 32 KB. In order to ensure that an exact number of input bytes represent an exact number of output bytes for each chunk, after each 32 KB of uncompressed data is represented in the output compressed bitstream, the output bitstream is padded with up to 15 bits of zeros to re-align the bitstream on a 16-bit boundary (even byte boundary) for the next 32 KB of data. This results in a compressed chunk of a byte-aligned size. The compressed chunk could be smaller than 32 KB or larger than 32 KB if the data is incompressible when the chunk is not the last one.

The LZXD engine encodes a byte-aligned little-endian 16-bit compressed chunk size prefix field preceding each compressed chunk in the compressed byte stream. The chunk prefix chain could be followed in the compressed stream without decompressing any data. The next chunk prefix is at a location computed by absolute byte offset location of this chunk prefix plus 2 (for the size of the chunk size prefix field) plus the current chunk size.

### 2.2.2 E8 Call Translation

E8 Call Translation is an optional feature that is sometimes used when the data to compress contains x86 instruction sequences. E8 Translation operates as a pre-processing stage prior to compressing each chunk, and the compressed stream header contains a bit that indicates whether the decoder shall reverse the translation as a post-processing step after decompressing each chunk.

The x86 instruction beginning with a byte value of 0xE8 is followed by a 32-bit little-endian relative displacement to the call target. When E8 Call Translation is enabled, the following pre-processing step is performed on the uncompressed input prior to compression (assuming little-endian byte ordering):

Let `chunk_offset` refer to the total number of uncompressed bytes preceding this chunk.

Let `E8_file_size` refer to the caller-specified value given to the compressor or decoded from the header of the compressed stream during decompression.

For each 32 KB chunk of uncompressed data (or less than 32 KB if last chunk to compress):

```

if ( ( chunk_offset < 0x40000000 ) && ( chunk_size > 10 ) )
    for ( i = 0; i < (chunk_size - 10); i++ )
        if ( chunk_byte[ i ] == 0xE8 )
            long current_pointer = chunk_offset + i;
            long displacement =  chunk_byte[ i+1 ] |
            chunk_byte[ i+2 ] << 8 |
            chunk_byte[ i+3 ] << 16 |
            chunk_byte[ i+4 ] << 24;
            long target = current_pointer + displacement;
            if ( ( target >= 0 ) && ( target < E8_file_size+current_pointer) )
                if ( target >= E8_file_size )
                    target = displacement - E8_file_size;
            endif

```

```

chunk_byte[ i+1 ] = (byte)( target );
chunk_byte[ i+2 ] = (byte)( target >> 8 );
chunk_byte[ i+3 ] = (byte)( target >> 16 );
chunk_byte[ i+4 ] = (byte)( target >> 24 );
endif
    i += 4;
endif
endfor
endif

```

After decompression, the E8 scanning algorithm is the same, but the translation reversal is:

```

long value =  chunk_byte[ i+1 ]      |
chunk_byte[ i+2 ] << 8  |
chunk_byte[ i+3 ] << 16 |
chunk_byte[ i+4 ] << 24;

if ( ( value >= -current_pointer ) && ( value < E8_file_size ) )
if ( ( value >= 0 )
displacement = value - current_pointer;
else
displacement = value + E8_file_size;
endif
chunk_byte[ i+1 ] = (byte)( displacement );
chunk_byte[ i+2 ] = (byte)( displacement >> 8 );
chunk_byte[ i+3 ] = (byte)( displacement >> 16 );
chunk_byte[ i+4 ] = (byte)( displacement >> 24 );
endif

```

The first bit in the first Chunk in the LZXD bitstream (following the 2-byte Chunk Size prefix described below) indicates the presence or absence of two 16-bit fields immediately following the single bit. If the bit is set, E8 translation is enabled using the 32-bit value derived from the two 16-bit fields as the `E8_file_size` provided to the compressor when E8 translation was enabled. Note that `E8_file_size` is completely independent of the length of the uncompressed data. E8 call translation is always disabled after the 32,768th chunk (after 1 GB of uncompressed data).

Field	Comments	Size
E8 translation	0-disabled, 1-enabled	1 bit
Translation size high word	Only present if enabled	0 or 16 bits
Translation size low word	Only present if enabled	0 or 16 bits

## 2.3 Block

### 2.3.1 Block Header

An LZXD Block represents a sequence of compressed data that is encoded with the same set of Huffman trees, or a sequence of uncompressed data. There can be one or more LZXD Blocks in a compressed stream, each with its own set of Huffman trees. Blocks do not have to start or end on a chunk boundary; blocks can span multiple chunks, or a single chunk can contain multiple blocks. The number of chunks is related to the size of the data being compressed, while the number of blocks is related to how well the data is compressed. The **Block Type** field indicates which type of

block follows, and the **Block Size** field indicates the number of uncompressed bytes represented by the block. Following the generic Block header, there is a type-specific header that describes the remainder of the block.

Field	Comments	Size
Block Type	See valid values in section <a href="#">2.3.1.1</a>	3 bits
Block Size MSB	Block size high 8 bits of 24	8 bits
Block Size byte 2	Block size middle 8 bits of 24	8 bits
Block Size LSB	Block size low 8 bits of 24	8 bits

### 2.3.1.1 Block Type

Each block of compressed data begins with a 3-bit field indicating the block type, followed by the **Block Size** and then type-specific **Block Data**. Of the eight possible values, only three are valid types.

Bits	Value	Meaning
001	1	Verbatim block
010	2	Aligned offset block
011	3	Uncompressed block
other	0, 4-7	Invalid

### 2.3.1.2 Block Size

The **Block Size** field indicates the number of uncompressed bytes that are represented by the block. The maximum **Block Size** is 224-1 (16MB-1 or 0x00FFFFFF). The **Block Size** is encoded in the bitstream as three 8-bit fields comprising a 24-bit value, most significant to least significant, immediately following the **Block Type** encoding.

## 2.3.2 Block Data

### 2.3.2.1 Uncompressed Block

Following the generic Block header, an uncompressed block begins with 1 to 16 bits of zero padding to align the bit buffer on a 16-bit boundary. At this point, the bitstream ends, and a byte stream begins. Following the zero padding, new 32-bit values for R0, R1, and R2 are output in little-endian form, followed by the uncompressed data bytes themselves. Finally, if the uncompressed data length is odd, one extra byte of zero padding is encoded to re-align the following bitstream.

Field	Comments	Size
Padding to align following field on 16-bit boundary	Bits have value of zero	Variable, 1..16 bits

Then, the following fields are encoded directly in the byte stream, not the bitstream of byte-swapped 16-bit words:

Field	Comments	Size
R0	LSB to MSB (little-endian dword)	4 bytes
R1	LSB to MSB (little-endian dword)	4 bytes
R2	LSB to MSB (little-endian dword)	4 bytes
Uncompressed raw data bytes	Can use direct memcpy	1..224-1 bytes
Padding to re-align bitstream	Only if uncompressed size is odd	0 or 1 byte

Then the bitstream of byte-swapped 16 bit integers resumes for the next **Block Type** field (if there are subsequent blocks).

The decoded R0, R1, and R2 values are used as initial Repeated Offset values to decode the subsequent compressed block if present.

### 2.3.2.2 Verbatim Block

The following table lists the fields of a Verbatim block that follow the generic Block header.

Entry	Comments	Size
Pre-tree for first 256 elements of main tree	20 elements, 4 bits each	80 bits
Path lengths of first 256 elements of main tree	Encoded using pre-tree	Variable
Pre-tree for remainder of main tree	20 elements, 4 bits each	80 bits
Path lengths of remaining elements of main tree	Encoded using pre-tree	Variable
Pre-tree for length tree	20 elements, 4 bits each	80 bits
Path lengths of elements in length tree	Encoded using pre-tree	Variable
Token sequence (matches and literals)	Described in section <a href="#">2.6</a> .	Variable

### 2.3.2.3 Aligned Offset Block

An Aligned offset block is identical to the Verbatim block except for the presence of the Aligned Offset Tree preceding the other trees.

Entry	Comments	Size
Aligned offset tree	8 elements, 3 bits each	24 bits
Pre-tree for first 256 elements of main tree	20 elements, 4 bits each	80 bits
Path lengths of first 256 elements of main tree	Encoded using pre-tree	Variable
Pre-tree for remainder of main tree	20 elements, 4 bits each	80 bits
Path lengths of remaining elements of main tree	Encoded using pre-tree	Variable
Pre-tree for length tree	20 elements, 4 bits each	80 bits
Path lengths of elements in length tree	Encoded using pre-tree	Variable

Entry	Comments	Size
Token sequence (matches and literals)	Described in section <a href="#">2.6</a>	Variable

## 2.4 Huffman Trees

LZXD uses canonical Huffman tree structures to represent elements. Huffman trees are well known in data compression and are not described here. Because an LZXD decoder uses only the path lengths of the Huffman tree to reconstruct the identical tree, the following constraints are made on the tree structure.

For any two elements with the same path length, the lower-numbered element **MUST** be further left on the tree than the higher-numbered element. An alternative way of stating this constraint is that lower-numbered elements **MUST** have lower path traversal values; for example, 0010 (left-left-right-left) is lower than 0011 (left-left-right-right).

For each level, starting at the deepest level of the tree and then moving upward, leaf nodes **MUST** start as far left as possible. An alternative way of stating this constraint is that if any tree node has children, then all tree nodes to the right of it with the same path length **MUST** also have children.

A non-empty Huffman tree **MUST** contain at least two elements. In the case where all but one tree element has zero frequency, the resulting tree **MUST** minimally consist of two Huffman codes, "0" and "1".

LZXD uses several Huffman tree structures. The Main Tree comprises 256 elements that correspond to all possible 8-bit characters, plus 8 \* **NUM\_POSITION\_SLOTS** elements that correspond to matches. **NUM\_POSITION\_SLOTS** refers to the position slots required, as specified in section [2.1.6](#). The value of **NUM\_POSITION\_SLOTS** depends on the specified window size as described in section [2.1.6](#). The Length Tree comprises 249 elements. Other trees, such as the Aligned Offset Tree (comprising 8 elements), and the Pre-Trees (comprising 20 elements each), have a smaller role.

## 2.5 Encoding the Trees and Pre-Trees

Because all trees used in LZXD are created in the form of a canonical Huffman tree, the path length of each element in the tree is sufficient to reconstruct the original tree. The main tree and the length tree are each encoded using the method described here. However, the main tree is encoded in two components as if it were two separate trees, the first tree corresponding to the first 256 tree elements (uncompressed symbols), and the second tree corresponding to the remaining elements (matches).

Because trees are output several times during compression of large amounts of data (multiple blocks), LZXD optimizes compression by encoding only the delta path lengths between the current and previous trees. In the case of the very first such tree, the delta is calculated against a tree in which all elements have a zero path length.

Each tree element can have a path length from 0 to 16 (inclusive) where a zero path length indicates that the element has a zero frequency and is not present in the tree. Tree elements are output in sequential order starting with the first element. Elements can be encoded in one of two ways: If several consecutive elements have the same path length, then run-length encoding is employed; otherwise, the element is output by encoding the difference between the current path length and the previous path length of the tree, mod 17. To represent a canonical Huffman tree, specify the path lengths of each of the elements in the tree. The following table specifies how to interpret a code.

Code	Operation
0-16	$\text{Len}[x] = (\text{prev\_len}[x] - \text{code} + 17) \bmod 17$
17	Zeroes = getbits(4) Len[x] = 0 for next (4 + Zeroes) elements
18	Zeroes = getbits(5) Len[x] = 0 for next (20 + Zeroes) elements
19	Same = getbits(1) Decode new code Value = $(\text{prev\_len}[x] - \text{code} + 17) \bmod 17$ Len[x] = Value for next (4 + Same) elements

Codes 17, 18, and 19 are used to represent consecutive elements that have the same path length. Zeroes, Same, and Value are variables created for the purpose of this sample code and getbits(n) is a function that fetches the next n bits from the bitstream. "Decode new Code" is used to parse the next Code from the bitstream, which will have a value of 0-16.

Each of the 17 possible values of  $(\text{len}[x] - \text{prev\_len}[x]) \bmod 17$ , plus three additional codes used for run-length encoding, are not output directly as 5-bit numbers, but are instead encoded via a Huffman tree called the pre-tree. The pre-tree is generated dynamically according to the frequencies of the 20 allowable tree codes. The structure of the pre-tree is encoded in a total of 80 bits by using 4 bits to output the path length of each of the 20 pre-tree elements. Once again, a zero path length indicates a zero frequency element.

Code	Operation
Length of tree code 0	4 bits
Length of tree code 1	4 bits
Length of tree code 2	4 bits
...	...
Length of tree code 18	4 bits
Length of tree code 19	4 bits

The "real" tree is then encoded using the pre-tree Huffman codes.

## 2.6 Compressed Token Sequence


The compressed token sequence (bitstream) contains the Huffman-encoded matches and literals using the Huffman trees specified in the Block header. Decompression continues until the number of decompressed bytes corresponds exactly to the number of uncompressed bytes indicated in the Block header.

The representation of an unmatched literal character in the output is simply the appropriate element index 0..255 from the Main Huffman Tree.

The representation of a match in the output involves several transformations, as shown in the following diagram. At the top of the diagram are the match length (2..257) and the match offset (0..WINDOW\_SIZE-4). The match offset and match length are split into subcomponents and

encoded separately. For matches of length 257..32768, the token indicates match length 257 and then there is an additional Extra Length value encoded in the bitstream following the other **Match** subcomponent fields.

The following figure shows the match subcomponents.



**Figure 1: Diagram of match encoding subcomponents**

### 2.6.1 Converting Match Offset into Formatted Offset Values

The match offset, range 1.. (WINDOW\_SIZE-4), is converted into a formatted offset by determining whether the offset can be encoded as a repeated offset, as shown in the following pseudocode. It is acceptable to not encode a match as a repeated offset even if it is possible to do so.

```

if offset == R0 then
    formatted offset ← 0
else if offset == R1 then
    formatted offset ← 1
else if offset == R2 then
    formatted offset ← 2
else
    formatted offset ← offset + 2
endif
  
```



## 2.6.2 Converting Formatted Offset into Position Slot and Position Footer Values

The formatted offset is subdivided into a position slot and position footer. The position slot defines the most significant bits of the formatted offset in the form of a base position as shown in the table on the following page. The position footer defines the remaining least significant bits of the formatted offset. As the following table shows, the number of bits dedicated to the position footer grows as the formatted offset becomes larger, meaning that each position slot addresses a larger and larger range.

The number of position slots available depends on the window size. The number of bits of position footer for each position slot is fixed and is shown in the following table.

Position slot number	Base position	Footer bits	Base position plus footer range
0 (R0)	0	0	0
1 (R1)	1	0	1
2 (R2)	2	0	2
3 (offset 1)	3	0	3
4 (offset 2..3)	4	1	4-5
5 (offset 4..5)	6	1	6-7
6 (offset 6..9)	8	2	8-11
7 (..etc..)	12	2	12-15
8	16	3	16-23
9	24	3	24-31
10	32	4	32-47
11	48	4	48-63
12	64	5	64-95
13	96	5	96-127
14	128	6	128-191
15	192	6	192-255
16	256	7	256-383
17	384	7	384-511
18	512	8	512-767
19	768	8	768-1023
20	1024	9	1024-1535
21	1536	9	1536-2047
22	2048	10	2048-3071

Position slot number	Base position	Footer bits	Base position plus footer range
23	3072	10	3072-4095
24	4096	11	4096-6143
25	6144	11	6144-8191
26	8192	12	8192-12287
27	12288	12	12288-16383
28	16384	13	16384-24575
29	24576	13	24576-32767
30	32768	14	32768-49151
31	49152	14	49152-65535
32	65536	15	65536-98303
33	98304	15	98304-131071
34	131072	16	131072-196607
35	196608	16	196608-262143
36	262144	17	262144-393215
37	393216	17	393216-524287
38	524288	17	524288-655359
39	655360	17	655360-786431
40	786432	17	786432-917503
41	917504	17	917504-1048575
42	1048576	17	1048576-1179647
..etc..	..etc..	17 (all)	..etc..
288	33292288	17	33292288-33423359
289	33423360	17	33423360-33554431

### 2.6.3 Converting Position Footer into Verbatim Bits or Aligned Offset Bits

The position footer can be further subdivided into verbatim bits and aligned offset bits if the current block type is "aligned offset". If the current block is not an aligned offset block, there are no aligned offset bits, and the verbatim bits are the position footer.

If aligned offsets are used, the lower 3 bits of the position footer are the aligned offset bits, while the remaining portion of the position footer is the verbatim bits. In the case where there are less than 3 bits in the position footer (for example, formatted offset is  $\leq 15$ ), it is not possible to take the "lower 3 bits of the position footer" and therefore there are no aligned offset bits, and the verbatim bits and the position footer are the same.

In situations where it is determined that there are a relatively larger number of Position Footers with identical lower 3 bits, Aligned Offset Block could be used to reduce the number of bits required to represent the Position Footer component in the match encoding.

Verbatim block could be used when the lower 3 bits of the Position Footer are relatively evenly distributed.

The following is pseudocode for splitting the position footer into verbatim bits and aligned offset.

```

if block_type is aligned_offset_block then
  if formatted_offset <= 15 then
    verbatim_bits ← position_footer
    aligned_offset ← null
  else
    aligned_offset ← position_footer
    verbatim_bits ← position_footer >> 3
  endif
else
  verbatim_bits ← position_footer
  aligned_offset ← null
endif

```

### 2.6.4 Converting Match Length into Length Header and Length Footer Values

The match length is converted into a length header and a length footer. The length header can have one of eight possible values, from 0..7 (inclusive), indicating a match of length 2, 3, 4, 5, 6, 7, 8, or a length greater than 8. If the match length is 8 or less, there is no length footer. Otherwise, the value of the length footer is equal to the match length minus 9. The following is pseudocode for obtaining the length header and footer.

```

if match_length <= 8
  length_header ← match_length-2
  length_footer ← null
else
  length_header ← 7
  length_footer ← match_length-9
endif

```

Match length	Length header	Length footer value
2	0	None
3	1	None
4	2	None
5	3	None
6	4	None
7	5	None
8	6	None
9	7	0
10	7	1

Match length	Length header	Length footer value
...	...	...
256	7	247
257 or larger	7	248

### 2.6.5 Converting Length Header and Position Slot into Length/Position Header Values

The Length/Position header is the stage that correlates the match position with the match length (using only the most significant bits), and is created by combining the length header and the position slot, as follows:

$$\text{len\_pos\_header} \leftarrow (\text{position\_slot} \ll 3) + \text{length\_header}$$

This operation creates a unique value for every combination of match length 2, 3, 4, 5, 6, 7, 8 with every possible position slot. The remaining match lengths greater than 8 are all lumped together, and as a group are correlated with every possible position slot.

### 2.6.6 Extra Length

If the match length is 257 or larger, the encoded match length token (or match length, as specified in section 2.6) value is 257, and an encoded Extra Length field follows the other match encoding components, as specified in section 2.6.7, in the bitstream.

Prefix (in binary)	Number of Bits to Decode	Base Value to Add to Decoded Value
0	8	257
10	10	257 + 256
110	12	257 + 256 + 1024
111	15	257

If the encoded match length token is equal to 257, it indicates length of the match is  $\geq 257$ . If this is the case, look for the Extra Length field after the other match encoding components in the bitstream. Then look at the prefix of the Extra Length field. If the prefix is 0, decode the next 8 bits and add 257 to get the match length. If the prefix is 10, decode the next 10 bits and add 257 + 256 to the decoded value to get the match length. If the prefix is 110, decode the next 12 bits and add 257 + 256 + 1024 to the decoded value to get the match length. If the prefix is 111, decode the next 15 bits and add 257 to the decoded value to get the match length.

### 2.6.7 Encoding a Match

The match is finally output as part of the compressed bitstream in up to five components, in the following order:

1. Main Tree element at index ( $\text{len\_pos\_header} + 256$ ).
2. If  $\text{length\_footer} \neq \text{null}$ , then output Length Tree element  $\text{length\_footer}$ .
3. If  $\text{verbatim\_bits} \neq \text{null}$ , then output  $\text{verbatim\_bits}$ .

4. If `aligned_offset_bits != null`, then output element `aligned_offset` from the aligned offset tree.
5. If match length 257 or larger, output appropriate Extra Length prefix and value.

### 2.6.8 Encoding a Literal

A literal byte that is not part of a match is encoded simply as a Main Tree element index 0..255 corresponding to the value of the literal byte.

## 2.7 Decoding Matches and Literals (Aligned and Verbatim Blocks)

Decoding is performed by first decoding an element from the Main Tree and then, if the item is a match, determining which additional components are required to decode to reconstruct the match. The following is pseudocode for decoding a match or an uncompressed character.

```

main_element = main_tree.decode_element()

if (main_element < 256 ) /* is a literal character */

window[ curpos ] ← (byte) main_element
curpos ← curpos + 1

else /* is a match */

length_header ← (main_element - 256) & 7

if (length_header == 7)
match_length ← length_tree.decode_element() + 7 + 2
else
match_length ← length_header + 2 /* no length footer */
endif

position_slot ← (main_element - 256) >> 3

/* check for repeated offsets (positions 0,1,2) */
if (position_slot == 0)
match_offset ← R0
else if (position_slot == 1)
match_offset ← R1
swap(R0 ↔ R1)
else if (position_slot == 2)
match_offset ← R2
swap(R0 ↔ R2)
else /* not a repeated offset */
offset_bits ← footer_bits[ position_slot ]

if (block_type == aligned_offset_block)
if (offset_bits >= 3) /* this means there are some aligned bits */
verbatim_bits ← (readbits(offset_bits-3)) << 3
aligned_bits ← aligned_offset_tree.decode_element();
else /* 0, 1, or 2 verbatim bits */
verbatim_bits ← readbits(offset_bits)
aligned_bits ← 0
endif

formatted_offset ← base_position[ position_slot ]
+ verbatim_bits + aligned_bits

```

```

else /* block_type == verbatim_block */
verbatim_bits ← readbits(offset_bits)
formatted_offset ← base_position[ position_slot ] + verbatim_bits
endif

match_offset ← formatted_offset - 2

/* update repeated offset LRU queue */
R2 ← R1
R1 ← R0
R0 ← match_offset

endif

/* check for extra length */

if (match_length == 257)
if (readbits( 1 ) != 0)
if (readbits( 1 ) != 0)
if (readbits( 1 ) != 0)
extra_len = readbits( 15 )
else
extra_len = readbits( 12 ) + 1024 + 256
endif
endif
else
extra_len = readbits( 10 ) + 256
endif
else
extra_len = readbits( 8 )
endif

match_length ← 257 + extra_len

endif

/* copy match data */
for (i = 0; i < match_length; i++)
window[curpos + i] ← window[curpos + i - match_offset]

curpos ← curpos + match_length

endif

```

### 3 Structure Examples

The LZXD bitstream is to be interpreted as a sequence of aligned 16-bit integers stored in the order least-significant-byte most-significant-byte (little-endian words).

The only exception is the uncompressed data bytes stored in the uncompressed block interpreted as a sequence of bytes. The following is an example of a sample encoding sequence of a simple 3-byte text input "abc" encoded as uncompressed block type.

Bits to Decode	Value of Decoded Bits	Interpretation
16	0x0014	Chunk Size: 20 bytes
1	0	E8 Translation: disabled
3	3 (binary 011)	Block Type: uncompressed
24	0x000003	Block Size: 3 bytes
4	binary 0000	Padding to word-align following
4 bytes	0x00000001 (little-endian dword)	R0: 1
4 bytes	0x00000001 (little-endian dword)	R1: 1
4 bytes	0x00000001 (little-endian dword)	R2: 1
3 bytes	0x61, 0x62, 0x63	Uncompressed bytes: "abc"
1 byte	0x00	Padding to restore word-alignment

This is the raw hexadecimal compressed byte sequence of the encoded fields:

```
14 00 00 30 30 00 01 00 00 00 01 00 00 00 01 00 00 00 61 62 63 00
```

## 4 Security Considerations

None.



## 5 Appendix A: Product Behavior

The information in this specification is applicable to the following Microsoft products or supplemental software. References to product versions include released service packs:

- Microsoft® Exchange Server 2003
- Microsoft® Exchange Server 2007
- Microsoft® Exchange Server 2010
- Microsoft® Office Outlook® 2003
- Microsoft® Office Outlook® 2007
- Microsoft® Outlook® 2010

Exceptions, if any, are noted below. If a service pack or Quick Fix Engineering (QFE) number appears with the product version, behavior changed in that service pack or QFE. The new behavior also applies to subsequent service packs of the product unless otherwise specified. If a product edition appears with the product version, behavior is different in that product edition.

Unless otherwise specified, any statement of optional behavior in this specification that is prescribed using the terms SHOULD or SHOULD NOT implies product behavior in accordance with the SHOULD or SHOULD NOT prescription. Unless otherwise specified, the term MAY implies that the product does not follow the prescription.

## 6 Change Tracking

No table of changes is available. The document is either new or has had no changes since its last release.

## 7 Index

### A

[Applicability](#) 6

### C

[Change tracking](#) 26

[Common data types and fields](#) 7

### D

[Data types and fields - common](#) 7

Details

[common data types and fields](#) 7

### E

[Examples](#) 23

### F

[Fields - vendor-extensible](#) 6

### G

[Glossary](#) 5

### I

[Implementer - security considerations](#) 24

[Informative references](#) 5

[Introduction](#) 5

### L

[Localization](#) 6

### N

[Normative references](#) 5

### O

[Overview \(synopsis\)](#) 5

### P

[Product behavior](#) 25

### R

References

[informative](#) 5

[normative](#) 5

[Relationship to protocols and other structures](#) 6

### S

[Security - implementer considerations](#) 24

Structures

[overview](#) 7

### T

[Tracking changes](#) 26

### V

[Vendor-extensible fields](#) 6

[Versioning](#) 6