

```

*****
26894 Sat Apr 20 02:07:53 2013
new/usr/src/uts/common/fs/zfs/lz4.c
Integrated r91 LZ4.
*****
1 /*
2  * LZ4 - Fast LZ compression algorithm
3  * Header File
4  * Copyright (C) 2011-2013, Yann Collet.
5  * BSD 2-Clause License (http://www.opensource.org/licenses/bsd-license.php)
6  *
7  * Redistribution and use in source and binary forms, with or without
8  * modification, are permitted provided that the following conditions are
9  * met:
10 *
11 *   * Redistributions of source code must retain the above copyright
12 *   notice, this list of conditions and the following disclaimer.
13 *   * Redistributions in binary form must reproduce the above
14 *   copyright notice, this list of conditions and the following disclaimer
15 *   in the documentation and/or other materials provided with the
16 *   distribution.
17 *
18 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
19 * "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
20 * LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
21 * A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
22 * OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
23 * SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
24 * LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
25 * DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
26 * THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
27 * (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
28 * OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
29 *
30 * You can contact the author at :
31 * - LZ4 homepage : http://fastcompression.blogspot.com/p/lz4.html
32 * - LZ4 source repository : http://code.google.com/p/lz4/
33 * Upstream release : r91
34 */

36 #include <sys/zfs_context.h>

38 static int real_LZ4_compress(const char *source, char *dest, int isize,
39                             int osize);
39 static int real_LZ4_uncompress(const char *source, char *dest, int osize);
40 static int LZ4_compressBound(int isize);
41 static int LZ4_uncompress_unknownOutputSize(const char *source, char *dest,
42                                             int isize, int maxOutputSize);
43 static int LZ4_compressCtx(void *ctx, const char *source, char *dest,
44                             int isize, int osize);
45 static int LZ4_compress64kCtx(void *ctx, const char *source, char *dest,
46                               int isize, int osize);

48 /*ARGSUSED*/
49 size_t
50 lz4_compress(void *s_start, void *d_start, size_t s_len, size_t d_len, int n)
51 {
52     uint32_t bufsiz;
53     char *dest = d_start;

55     ASSERT(d_len >= sizeof (bufsiz));

57     bufsiz = real_LZ4_compress(s_start, &dest[sizeof (bufsiz)], s_len,
58                               d_len - sizeof (bufsiz));

60     /* Signal an error if the compression routine returned zero. */

```

```

61     if (bufsiz == 0)
62         return (s_len);

64     /*
65     * Encode the compressed buffer size at the start. We'll need this in
66     * decompression to counter the effects of padding which might be
67     * added to the compressed buffer and which, if unhandled, would
68     * confuse the hell out of our decompression function.
69     */
70     *(uint32_t *)dest = BE_32(bufsiz);

72     return (bufsiz + sizeof (bufsiz));
73 }
    unchanged_portion_omitted

94 /*
95 * LZ4 API Description:
96 *
97 * Simple Functions:
98 * real_LZ4_compress() :
99 *   isize : is the input size. Max supported value is ~1.9GB
100 *   return : the number of bytes written in buffer dest
101 *   or 0 if the compression fails (if LZ4_COMPRESSMIN is set).
102 *   note : destination buffer must be already allocated.
103 *   destination buffer must be sized to handle worst cases
104 *   situations (input data not compressible) worst case size
105 *   evaluation is provided by function LZ4_compressBound().
106 *
107 * real_LZ4_uncompress() :
108 *   osize : is the output size, therefore the original size
109 *   return : the number of bytes read in the source buffer.
110 *   If the source stream is malformed, the function will stop
111 *   decoding and return a negative result, indicating the byte
112 *   position of the faulty instruction. This function never
113 *   writes beyond dest + osize, and is therefore protected
114 *   against malicious data packets.
115 *   note : destination buffer must be already allocated
116 *
117 * Advanced Functions
118 *
119 * LZ4_compressBound() :
120 *   Provides the maximum size that LZ4 may output in a "worst case"
121 *   scenario (input data not compressible) primarily useful for memory
122 *   allocation of output buffer.
123 *
124 *   isize : is the input size. Max supported value is ~1.9GB
125 *   return : maximum output size in a "worst case" scenario
126 *   note : this function is limited by "int" range (2^31-1)
127 *
128 * LZ4_uncompress_unknownOutputSize() :
129 *   isize : is the input size, therefore the compressed size
130 *   maxOutputSize : is the size of the destination buffer (which must be
131 *   already allocated)
132 *   return : the number of bytes decoded in the destination buffer
133 *   (necessarily <= maxOutputSize). If the source stream is
134 *   malformed, the function will stop decoding and return a
135 *   negative result, indicating the byte position of the faulty
136 *   instruction. This function never writes beyond dest +
137 *   maxOutputSize, and is therefore protected against malicious
138 *   data packets.
139 *   note : Destination buffer must be already allocated.
140 *   This version is slightly slower than real_LZ4_uncompress()
141 *
142 * LZ4_compressCtx() :
143 *   This function explicitly handles the CTX memory structure.

```

```

134 *      ILLUMOS CHANGES: the CTX memory structure must be explicitly allocated
135 *      by the caller (either on the stack or using kmem_zalloc). Passing NULL
136 *      isn't valid.
137 *
138 * LZ4_compress64kCtx() :
139 *      Same as LZ4_compressCtx(), but specific to small inputs (<64KB).
140 *      isize *Must* be <64KB, otherwise the output will be corrupted.
141 *
142 *      ILLUMOS CHANGES: the CTX memory structure must be explicitly allocated
143 *      by the caller (either on the stack or using kmem_zalloc). Passing NULL
144 *      isn't valid.
145 */

147 /*
148 * Tuning parameters
149 */

151 /*
152 * COMPRESSIONLEVEL: Increasing this value improves compression ratio
153 *      Lowering this value reduces memory usage. Reduced memory usage
154 *      typically improves speed, due to cache effect (ex: L1 32KB for Intel,
155 *      L1 64KB for AMD). Memory usage formula : N->2^(N+2) Bytes
156 *      (examples : 12 -> 16KB ; 17 -> 512KB)
157 */
158 #define COMPRESSIONLEVEL 12

160 /*
161 * NOTCOMPRESSIBLE_CONFIRMATION: Decreasing this value will make the
162 *      algorithm skip faster data segments considered "incompressible".
163 *      This may decrease compression ratio dramatically, but will be
164 *      faster on incompressible data. Increasing this value will make
165 *      the algorithm search more before declaring a segment "incompressible".
166 *      This could improve compression a bit, but will be slower on
167 *      incompressible data. The default value (6) is recommended.
168 */
169 #define NOTCOMPRESSIBLE_CONFIRMATION 6

171 /*
172 * BIG_ENDIAN_NATIVE_BUT_INCOMPATIBLE: This will provide a boost to
173 *      performance for big endian cpu, but the resulting compressed stream
174 *      will be incompatible with little-endian CPU. You can set this option
175 *      to 1 in situations where data will stay within closed environment.
176 *      This option is useless on Little_Endian CPU (such as x86).
177 */
178 /* #define      BIG_ENDIAN_NATIVE_BUT_INCOMPATIBLE 1 */

180 /*
181 * CPU Feature Detection
182 */

184 /* 32 or 64 bits ? */
185 #if (defined(__x86_64__) || defined(__x86_64) || defined(__amd64__) || \
186     defined(__amd64) || defined(__ppc64__) || defined(__WIN64) || \
187     defined(__LP64__) || defined(__LP64))
188 #define LZ4_ARCH64 1
189 #else
190 #define LZ4_ARCH64 0
191 #endif

193 /*
194 * Limits the amount of stack space that the algorithm may consume to hold
195 *      the compression lookup table. The value '9' here means we'll never use
196 *      more than 2k of stack (see above for a description of COMPRESSIONLEVEL).
197 *      If more memory is needed, it is allocated from the heap.
198 */
199 #define STACKLIMIT 9

```

```

201 /*
202 * Little Endian or Big Endian?
203 * Note: overwrite the below #define if you know your architecture endianness.
204 */
205 #if (defined(__BIG_ENDIAN__) || defined(__BIG_ENDIAN) || \
206     defined(__BIG_ENDIAN) || defined(__ARCH_PPC) || defined(__PPC__) || \
207     defined(__PPC) || defined(PPC) || defined(__powerpc__) || \
208     defined(__powerpc) || defined(powerpc) || \
209     ((defined(__BYTE_ORDER__) && (__BYTE_ORDER__ == __ORDER_BIG_ENDIAN__)))
210 #define LZ4_BIG_ENDIAN 1
211 #else
212 /*
213 * Little Endian assumed. PDP Endian and other very rare endian format
214 * are unsupported.
215 */
216 #endif

218 /*
219 * Unaligned memory access is automatically enabled for "common" CPU,
220 * such as x86. For others CPU, the compiler will be more cautious, and
221 * insert extra code to ensure aligned access is respected. If you know
222 * your target CPU supports unaligned memory access, you may want to
223 * force this option manually to improve performance
224 */
225 #if defined(__ARM_FEATURE_UNALIGNED)
226 #define LZ4_FORCE_UNALIGNED_ACCESS 1
227 #endif

229 /* #define      LZ4_FORCE_SW_BITCOUNT */

231 /*
232 * Compiler Options
233 */
234 #if __STDC_VERSION__ >= 199901L /* C99 */
235 /* "restrict" is a known keyword */
236 #else
237 /* Disable restrict */
238 #define restrict
239 #endif

241 #define GCC_VERSION (__GNUC__ * 100 + __GNUC_MINOR__)

243 #ifdef _MSC_VER
244 /* Visual Studio */
245 /* Visual is not C99, but supports some kind of inline */
246 #define inline _forceinline
247 #if LZ4_ARCH64
248 /* For Visual 2005 */
249 #pragma intrinsic(_BitScanForward64)
250 #pragma intrinsic(_BitScanReverse64)
251 #else /* !LZ4_ARCH64 */
252 /* For Visual 2005 */
253 #pragma intrinsic(_BitScanForward)
254 #pragma intrinsic(_BitScanReverse)
255 #endif /* !LZ4_ARCH64 */
256 #endif /* _MSC_VER */

258 #ifdef _MSC_VER
259 #define lz4_bswap16(x) _byteswap_ushort(x)
260 #else /* !_MSC_VER */
261 #define lz4_bswap16(x) (((unsigned short int) (((x) >> 8) & 0xffu) | \
262     (((x) & 0xffu) << 8))
263 #endif /* !_MSC_VER */

265 #if (GCC_VERSION >= 302) || (__INTEL_COMPILER >= 800) || defined(__clang__)

```

```

266 #define expect(expr, value)    (__builtin_expect((expr), (value)))
267 #else
268 #define expect(expr, value)    (expr)
269 #endif

271 #define likely(expr)    expect((expr) != 0, 1)
272 #define unlikely(expr) expect((expr) != 0, 0)

274 /* Basic types */
275 #if defined(_MSC_VER)
276 /* Visual Studio does not support 'stdint' natively */
277 #define BYTE    unsigned __int8
278 #define U16    unsigned __int16
279 #define U32    unsigned __int32
280 #define S32    __int32
281 #define U64    unsigned __int64
282 #else /* !defined(_MSC_VER) */
283 #define BYTE    uint8_t
284 #define U16    uint16_t
285 #define U32    uint32_t
286 #define S32    int32_t
287 #define U64    uint64_t
288 #endif /* !defined(_MSC_VER) */

290 #ifndef LZ4_FORCE_UNALIGNED_ACCESS
291 #pragma pack(1)
292 #endif

294 typedef struct _U16_S {
295     U16 v;
296 } U16_S;
unchanged portion omitted

500 /* Compression functions */

502 /*ARGSUSED*/
503 static int
504 LZ4_compressCtx(void *ctx, const char *source, char *dest, int isize,
505     int osize)
506 {
507 #if HEAPMODE
508     struct refTables *srt = (struct refTables *)ctx;
509     HTYPE *HashTable = (HTYPE *) (srt->hashTable);
510 #else
511     HTYPE HashTable[HASHTABLESIZE] = { 0 };
512 #endif

514     const BYTE *ip = (BYTE *) source;
515     INITBASE(base);
516     const BYTE *anchor = ip;
517     const BYTE *const iend = ip + isize;
518     const BYTE *const oend = (BYTE *) dest + osize;
519     const BYTE *const mflimit = iend - MFLIMIT;
520 #define matchlimit (iend - LASTLITERALS)

522     BYTE *op = (BYTE *) dest;

524     int length;
525     int len, length;
525     const int skipStrength = SKIPSTRENGTH;
526     U32 forwardH;

529     /* Init */
530     if (isize < MINLENGTH)
531         goto _last_literals;

```

```

533     /* First Byte */
534     HashTable[LZ4_HASH_VALUE(ip)] = ip - base;
535     ip++;
536     forwardH = LZ4_HASH_VALUE(ip);

538     /* Main Loop */
539     for (;;) {
540         int findMatchAttempts = (1U << skipStrength) + 3;
541         const BYTE *forwardIp = ip;
542         const BYTE *ref;
543         BYTE *token;

545         /* Find a match */
546         do {
547             U32 h = forwardH;
548             int step = findMatchAttempts++ >> skipStrength;
549             ip = forwardIp;
550             forwardIp = ip + step;

552             if unlikely(forwardIp > mflimit) {
553                 goto _last_literals;
554             }

556             forwardH = LZ4_HASH_VALUE(forwardIp);
557             ref = base + HashTable[h];
558             HashTable[h] = ip - base;

560         } while ((ref < ip - MAX_DISTANCE) || (A32(ref) != A32(ip)));

562         /* Catch up */
563         while ((ip > anchor) && (ref > (BYTE *) source) &&
564             unlikely(ip[-1] == ref[-1])) {
565             ip--;
566             ref--;
567         }

569         /* Encode Literal length */
570         length = ip - anchor;
571         token = op++;

573         /* Check output limit */
574         if unlikely(op + length + (2 + 1 + LASTLITERALS) +
575             (length >> 8) > oend)
576             return (0);

578         if (length >= (int)RUN_MASK) {
579             int len;
580             *token = (RUN_MASK << ML_BITS);
581             len = length - RUN_MASK;
582             for (; len > 254; len -= 255)
583                 *op++ = 255;
584             *op++ = (BYTE)len;
585         } else
586             *token = (length << ML_BITS);

588         /* Copy Literals */
589         LZ4_BLINDCOPY(anchor, op, length);

591         _next_match:
592         /* Encode Offset */
593         LZ4_WRITE_LITTLEENDIAN_16(op, ip - ref);

595         /* Start Counting */
596         ip += MINMATCH;
597         ref += MINMATCH; /* MinMatch already verified */

```

```

607         ref += MINMATCH;          /* MinMatch verified */
608         anchor = ip;
609         while likely(ip < matchlimit - (STEP_SIZE - 1)) {
610             UARCH diff = AARCH(ref) ^ AARCH(ip);
611             if (!diff) {
612                 ip += STEP_SIZE;
613                 ref += STEP_SIZE;
614                 continue;
615             }
616             ip += LZ4_NbCommonBytes(diff);
617             goto _endCount;
618         }
619 #if LZ4_ARCH64
620         if ((ip < (matchlimit - 3)) && (A32(ref) == A32(ip))) {
621             ip += 4;
622             ref += 4;
623         }
624 #endif
625         if ((ip < (matchlimit - 1)) && (A16(ref) == A16(ip))) {
626             ip += 2;
627             ref += 2;
628         }
629         if ((ip < matchlimit) && (*ref == *ip))
630             ip++;
631         _endCount:
632
633         /* Encode MatchLength */
634         length = (int)(ip - anchor);
635         len = (ip - anchor);
636         /* Check output limit */
637         if unlikely(op + (1 + LASTLITERALS) + (length >> 8) > oend)
638             if unlikely(op + (1 + LASTLITERALS) + (len >> 8) > oend)
639                 return (0);
640         if (length >= (int)ML_MASK) {
641             if (len >= (int)ML_MASK) {
642                 *token += ML_MASK;
643                 length -= ML_MASK;
644                 for (; length > 509; length -= 510) {
645                     len -= ML_MASK;
646                     for (; len > 509; len -= 510) {
647                         *op++ = 255;
648                         *op++ = 255;
649                     }
650                 }
651                 if (length > 254) {
652                     length -= 255;
653                     if (len > 254) {
654                         len -= 255;
655                         *op++ = 255;
656                     }
657                 }
658                 *op++ = (BYTE)length;
659                 *op++ = (BYTE)len;
660             } else
661                 *token += length;
662                 *token += len;
663
664         /* Test end of chunk */
665         if (ip > mflimit) {
666             anchor = ip;
667             break;
668         }
669         /* Fill table */
670         HashTable[LZ4_HASH_VALUE(ip - 2)] = ip - 2 - base;
671
672         /* Test next position */
673         ref = base + HashTable[LZ4_HASH_VALUE(ip)];
674         HashTable[LZ4_HASH_VALUE(ip)] = ip - base;

```

```

654         if ((ref > ip - (MAX_DISTANCE + 1)) && (A32(ref) == A32(ip))) {
655             token = op++;
656             *token = 0;
657             goto _next_match;
658         }
659         /* Prepare next loop */
660         anchor = ip++;
661         forwardH = LZ4_HASH_VALUE(ip);
662     }
663
664     _last_literals:
665     /* Encode Last Literals */
666     {
667         int lastRun = iend - anchor;
668         if (op + lastRun + 1 + ((lastRun + 255 - RUN_MASK) / 255) >
669             oend)
670             return (0);
671         if (lastRun >= (int)RUN_MASK) {
672             *op++ = (RUN_MASK << ML_BITS);
673             lastRun -= RUN_MASK;
674             for (; lastRun > 254; lastRun -= 255) {
675                 *op++ = 255;
676             }
677             *op++ = (BYTE)lastRun;
678         } else
679             *op++ = (lastRun << ML_BITS);
680         (void) memcpy(op, anchor, iend - anchor);
681         op += iend - anchor;
682     }
683
684     /* End */
685     return (int)(((char *)op) - dest);
686 }
687
688 unchanged_portion_omitted
689
690 /* Decompression functions */
691
692 /*
693  * Note: The decoding functions real_LZ4_uncompress() and
694  * LZ4_uncompress_unknownOutputSize() are safe against "buffer overflow"
695  * attack type. They will never write nor read outside of the provided
696  * output buffers. LZ4_uncompress_unknownOutputSize() also insures that
697  * it will never read outside of the input buffer. A corrupted input
698  * will produce an error result, a negative int, indicating the position
699  * of the error within input stream.
700 */
701
702 static int
703 LZ4_uncompress_unknownOutputSize(const char *source, char *dest, int isize,
704     int maxOutputSize)
705 real_LZ4_uncompress(const char *source, char *dest, int osize)
706 {
707     /* Local Variables */
708     const BYTE *restrict ip = (const BYTE *) source;
709     const BYTE *const iend = ip + isize;
710     const BYTE *ref;
711
712     BYTE *op = (BYTE *) dest;
713     BYTE *const oend = op + maxOutputSize;
714     BYTE *const oend2 = op + osize;
715     BYTE *cpy;
716
717     unsigned token;
718
719     size_t length;
720     size_t dec32table[] = {0, 3, 2, 3, 0, 0, 0, 0};

```

```

924 #if LZ4_ARCH64
925     size_t dec64table[] = {0, 0, 0, (size_t)-1, 0, 1, 2, 3};
926 #endif

949     /* Main Loop */
950     for (;;) {
951         /* get runlength */
952         token = *ip++;
953         if ((length = (token >> ML_BITS)) == RUN_MASK) {
954             size_t len;
955             for (; (len = *ip++) == 255; length += 255) {
956                 length += len;
957             }
958         }
959         /* copy literals */
960         cpy = op + length;
961         if unlikely(cpy > oend - COPYLENGTH) {
962             if (cpy != oend)
963                 /* Error: we must necessarily stand at EOF */
964                 goto _output_error;
965             (void) memcpy(op, ip, length);
966             ip += length;
967             break; /* EOF */
968         }
969         LZ4_WILDCOPY(ip, op, cpy);
970         ip -= (op - cpy);
971         op = cpy;

973         /* get offset */
974         LZ4_READ_LITTLEENDIAN_16(ref, cpy, ip);
975         ip += 2;
976         if unlikely(ref < (BYTE * const) dest)
977             goto _output_error;
978         /* Special case
979          * A correctly formed null-compressed LZ4 must have at least
980          * one byte (token=0)
981          * Error: offset create reference outside destination
982          * buffer
983          */
984         if (unlikely(ip == iend))
985             goto _output_error;

987         /* get matchlength */
988         if ((length = (token & ML_MASK)) == ML_MASK) {
989             for (; *ip == 255; length += 255) {
990                 ip++;
991             }
992             length += *ip++;
993         }
994         /* copy repeated sequence */
995         if unlikely(op - ref < STEPSIZE) {
996             #if LZ4_ARCH64
997                 size_t dec64 = dec64table[op-ref];
998             #else
999                 const int dec64 = 0;
1000             #endif
1001             op[0] = ref[0];
1002             op[1] = ref[1];
1003             op[2] = ref[2];
1004             op[3] = ref[3];
1005             op += 4;
1006             ref += 4;
1007             ref -= dec32table[op-ref];
1008             A32(op) = A32(ref);
1009             op += STEPSIZE - 4;
1010             ref -= dec64;

```

```

1007     } else {
1008         LZ4_COPYSTEP(ref, op);
1009     }
1010     cpy = op + length - (STEPSIZE - 4);
1011     if (cpy > oend - COPYLENGTH) {
1012         if (cpy > oend)
1013             /*
1014              * Error: request to write beyond destination
1015              * buffer
1016              */
1017             goto _output_error;
1018         LZ4_SECURECOPY(ref, op, (oend - COPYLENGTH));
1019         while (op < cpy)
1020             *op++ = *ref++;
1021         op = cpy;
1022         if (op == oend)
1023             /*
1024              * Check EOF (should never happen, since last
1025              * 5 bytes are supposed to be literals)
1026              */
1027             goto _output_error;
1028         continue;
1029     }
1030     LZ4_SECURECOPY(ref, op, cpy);
1031     op = cpy; /* correction */
1032 }

1034 /* end of decoding */
1035 return (int)(((char *)ip) - source);

1037 /* write overflow error detected */
1038 _output_error:
1039 return (int)-(((char *)ip) - source);
1040 }

1042 static int
1043 LZ4_uncompress_unknownOutputSize(const char *source, char *dest, int isize,
1044     int maxOutputSize)
1045 {
1046     /* Local Variables */
1047     const BYTE *restrict ip = (const BYTE *) source;
1048     const BYTE *const iend = ip + isize;
1049     const BYTE *ref;

1051     BYTE *op = (BYTE *) dest;
1052     BYTE *const oend = op + maxOutputSize;
1053     BYTE *cpy;

1055     size_t dec32table[] = {0, 3, 2, 3, 0, 0, 0, 0};
1056     #if LZ4_ARCH64
1057     size_t dec64table[] = {0, 0, 0, (size_t)-1, 0, 1, 2, 3};
1058     #endif

1096     /* Main Loop */
1097     /*LINTED E_CONSTANT_CONDITION*/
1098     while (1) {
1099         while (ip < iend) {
1100             unsigned token;
1101             size_t length;

1103             /* get runlength */
1104             token = *ip++;
1105             if ((length = (token >> ML_BITS)) == RUN_MASK) {
1106                 int s = 255;
1107                 while ((ip < iend) && (s == 255)) {
1108                     s = *ip++;

```

```

948         length += s;
949     }
950 }
951 /* copy literals */
952 cpy = op + length;
953 if ((cpy > oend - MFLIMIT) ||
954     (ip + length > iend - (2 + 1 + LASTLITERALS))) {
1076 if ((cpy > oend - COPYLENGTH) ||
1077     (ip + length > iend - COPYLENGTH)) {
955     if (cpy > oend)
956         /* Error: writes beyond output buffer */
957         goto _output_error;
958     if (ip + length != iend)
959         /*
960          * Error: LZ4 format requires to consume all
961          * input at this stage (no match within the
962          * last 11 bytes, and at least 8 remaining
963          * input bytes for another match + literals
1084          * input at this stage
964          */
965         goto _output_error;
966     (void) memcpy(op, ip, length);
967     op += length;
968     /* Necessarily EOF, due to parsing restrictions */
969     break;
970 }
971 LZ4_WILDCOPY(ip, op, cpy);
972 ip -= (op - cpy);
973 op = cpy;

975 /* get offset */
976 LZ4_READ_LITTLEENDIAN_16(ref, cpy, ip);
977 ip += 2;
978 if (unlikely(ref < (BYTE * const) dest))
1099 if (ref < (BYTE * const) dest)
979     /*
980      * Error: offset creates reference outside of
981      * destination buffer
982      */
983     goto _output_error;

985 /* get matchlength */
986 if ((length = (token & ML_MASK)) == ML_MASK) {
987     while (likely(ip < iend - (LASTLITERALS + 1))) {
1108         while (ip < iend) {
988             int s = *ip++;
989             length += s;
990             if (s == 255)
991                 continue;
992             break;
993         }
994     }
995     /* copy repeated sequence */
996     if unlikely(op - ref < STEPSIZE) {
997 #if LZ4_ARCH64
998         size_t dec64 = dec64table[op-ref];
999 #else
1000         const int dec64 = 0;
1001 #endif
1002         op[0] = ref[0];
1003         op[1] = ref[1];
1004         op[2] = ref[2];
1005         op[3] = ref[3];
1006         op += 4;
1007         ref += 4;
1008         ref -= dec32table[op-ref];

```

```

1009         A32(op) = A32(ref);
1010         op += STEPSIZE - 4;
1011         ref -= dec64;
1012     } else {
1013         LZ4_COPYSTEP(ref, op);
1014     }
1015     cpy = op + length - (STEPSIZE - 4);
1016     if (unlikely(cpy > oend - (COPYLENGTH + (STEPSIZE - 4)))) {
1017         if (cpy > oend - LASTLITERALS)
1018             if (cpy > oend - COPYLENGTH) {
1138                 if (cpy > oend)
1018                     /*
1019                      * Error: last 5 bytes must be literals
1140                      * Error: request to write outside of
1141                      * destination buffer
1020                      */
1021                     goto _output_error;
1022                 LZ4_SECURECOPY(ref, op, (oend - COPYLENGTH));
1023                 while (op < cpy)
1024                     *op++ = *ref++;
1025                 op = cpy;
1026                 if (op == oend)
1027                     /*
1028                      * Check EOF (should never happen, since
1029                      * last 5 bytes are supposed to be literals)
1030                      */
1031                     goto _output_error;
1032                 continue;
1033             }
1034             LZ4_WILDCOPY(ref, op, cpy);
1156             LZ4_SECURECOPY(ref, op, cpy);
1035             op = cpy; /* correction */
1036         }

1038     /* end of decoding */
1039     return (int)((char *)op) - dest);

1041     /* write overflow error detected */
1042     _output_error:
1043     return (int)-(((char *)ip) - source));
1044 }

```

unchanged_portion_omitted