Conceptual :-> NTFS $DATA runlist for huge compressed files, with multiple Extension MFT Records.

25 Apr 00 02:22 **Message 1 of 1**

ntfsd member 1727
xxxxxx@lists.osr.com

Join Date:
Posts To This List: 6

Conceptual :-> NTFS $DATA runlist for huge compressed files, with multiple Extension MFT Records.

Hi,

We encountered something surprising & new (to us) while parsing the NTFS runlists (extentlists) of a few bigger MS-Access files (although should very well happen with others) that were compressed.

Problem:
  a.. We found that run-elements making up a single Compression Unit (16 clusters) in the on-disk runlist, got split across two MFT Extension Records of the file.
    1.. So the question is, whether there was a single runlist across all the $DATA entries in the various extension records of a file?
    2.. We generally believed that different $DATA entries contain independent runlists, since each of them started with an Lcn Address & not an Offset.
    3.. What would the address field of a run-element succeeding such a split compression unit contain: an Lcn Address or an address offset?
  b.. Secondly, in another case we found the Termination Extent (a run-element with no address field, telling us just the no. of saved clusters in the compression unit after the compression & the end of compression unit) of a Compression Unit to be split itself across two extension records.
    1.. What could this mean? The first termination extent doesn't make 16 clusters with the previous run-element(s), so what does it signify?
    2.. Together with the second termination extent in the next extension record it does make 16, but then, why two???

Any inputs, most welcome & desperately needed.

Any references, web-links ... anything absolutely.

Detailed View & Analysis:

Conceptually speaking, the problem may be localized to those files which have many Extension MFT Records. Typically the files which fall into this category are either very, very huge uncompressed files, on fragmented volumes (This would make the $DATA Attribute's extentlist very long & take multiple extension records).
Or, we might have moderately sized compressed files on volumes with small cluster sizes (This would again inflate the runlist, as there will be an entry for every Compression Unit ~ 16 clusters).

The problem on my end is with the second case above - Compressed.

We found that one of our files which had multiple extension records (around 11), had its $DATA Attribute in all of them.
  1.. To the best of our present understanding, we expect all these $DATA Attributes to house a runlist within them. Till here everything was ok.
  2.. These individual runlists we believed would be independent. We thought, they would start & terminate in the same $DATA Attribute, within the same extension record.
  3.. This apparently doesn't seem to be at least, absolutely true. We have encountered cases where the a group (typically two) of run-elements in a runlist (we are talking compressed files) referring to a single Compression Unit, were found to be split across two extension records. E.g. in the enclosed raw image of dumps of extension records below (The addresses on left. The mft record contents in italics. The start of $DATA Attribute & its runlist underlined & red.) .....

  01B9FDA8 00 00 00 00 46 49 4C 45 2A 00 03 00 F3 AD 50 00 00 00 00 00
01 00 00 00 30 00 ....FILE*...??P.........0.   The record
header starts with FILE.
  01B9FDC2 01 00 48 02 00 00 00 04 00 00 D4 00 00 00 00 00 01 00 01 00
02 00 01 0D 00 00 ..H.......?.............
  01B9FDDC 80 00 00 00 10 02 00 00 01 00 00 00 01 00 00 00 00 00 00 00
00 00 00 00 14 04 ?.......................   The $DATA
starts here.
  01B9FDF6 00 00 00 00 00 00 48 00 04 00 00 00 00 00 00 00 3B 00 00 00
00 00 00 E8 3A 00 ......H.........;......?:.
  01B9FE10 00 00 00 00 00 E8 3A 00 00 00 00 00 00 00 96 13 00 00 00 00 00
21 05 5D 7B 01 0B .....?:......-......!.]{..
  01B9FE2A 31 05 09 8E 02 01 0B 11 05 34 01 0B 31 04 C8 71 FD 01 0C 31
06 84 8E 02 01 0A 1. Z.....4..1.?q?..1."Z.   The runlist
starts here.
  01B9FE44 11 05 34 01 0B 31 03 4C 71 FD 11 03 07 01 0A 31 06 F9 8E 02
01 0A 11 05 34 01 ..4..1.Lq?.....1.?Z.....4.
  01B9FE5E 0B 31 08 D6 70 FD 01 08 11 01 08 11 0B 35 01 04 11 01 0B 11
09 08 01 06 11 05 .1.?p?.......5...... .....
  01B9FE78 14 01 0B 31 05 92 8F 02 01 0B 31 08 77 70 FD 01 08 11 0C 10
01 04 31 07 F9 8F ...1.'....1.wp?.......1.?.

```
  01B9FE92  02 01 09 31 07 1A 70 FD 01 09 11 02 07 11 05 0D 01 09 31 06
52 90 02 01 0A 31  .. 1..p?. ....... 1.R....1
  01B9FEAC  07 B7 6F FD 01 09 11 01 07 11 06 09 01 09 31 08 B9 90 02 01
08 11 06 34 01 0A  .?o?. ..... . 1.?......4..
  01B9FEC6  31 05 19 6F FD 01 0B 31 06 33 91 02 01 0A 11 05 34 01 0B 31
01 9E 6E FD 11 05  1..o?..1.3'.....4..1.zn?          (Extension
Record 1)
  01B9FEE0  09 01 0A 31 07 A5 91 02 01 09 11 04 34 01 0C 31 03 2C 6E FD
01 0D 31 03 20 92   ..1.?'.. ..4..1.,n?..1. '
  01B9FEFA  02 01 0D 11 04 34 01 0C 31 03 B3 6D FD 01 0D 31 03 99 92 02
01 0D 11 04 34 01  .....4..1.?m?..1.T'.....4.
  01B9FF14  0C 31 01 36 6D FD 11 03 05 01 0C 31 04 11 93 02 01 0C 11 04
34 01 0C 31 01 BE  .1.6m?.....1.."".....4..1.?
  01B9FF2E  6C FD 11 03 05 01 0C 31 04 89 93 02 01 0C 31 05 7A 6C FD 11
02 0D 01 09 11 07  l?.....1.?"...1.zl?.... ..
  01B9FF48  02 01 09 31 06 F7 93 02 01 0A 11 05 34 01 0B 31 03 DC 6B FD
11 02 0A 01 0B 31  .. 1.?".....4..1.?k?.....1
  01B9FF62  05 66 94 02 01 0B 11 06 34 01 0A 31 06 68 6B FD 01 0A 31 06
E4 94 02 01 0A 11  .f".....4..1.hk?..1.?"....
  01B9FF7C  04 34 01 0C 31 01 EE 6A FD 11 02 05 01 0D 31 03 59 95 02 01
0D 11 03 34 01 0D  .4..1.?j?.....1.Y......4..
  01B9FF96  31 02 75 6A FD 11 01 06 01 0D 31 03 21 96 02 01 0D 11 03 44
01 0D 31 03 9C 69  1.uj?.....1.!-.....D..1.oi
  01B9FFB0  FD 11 03 07 01 0A 31 04 19 97 02 01 0C 11 03 44 01 0D 31 04
A6 68 FD 01 0C 31  ?.....1..-.....D..1.?h?..1
  01B9FFCA  05 16 98 02 01 0B 11 05 44 01 0B 31 03 AA 67 FD 01 0D 31 05
12 99 02 01 0B 11  ..~.....D..1.?g?..1..T....
  01B9FFE4  05 44 00 0C 10 11 0C 10 FF FF FF FF 00 00 00 00 09 9A 02 01
0B 11 04 44 01 0C  .D......????.... s.....D..            The
runlist ends here.
```

The last run-element (run) above in the runlist (in blue above) viz.: "11 05 44", refer to an extent of 5 clusters. However they don't make up a complete compression unit (like the last two run-elements in purple: "31 05 12 99 02" & "01 0B"). The other half of this compression unit is found in the next extension record (in blue: 01 0B)

.....

```
  01B9FD9A  00 00 D6 01 00 00 00 00 00 00 00 00 00 00 00 00 00 46 49
4C 45 2A 00 03 00  ..?.............FILE*...
  01B9FDB4  6B E0 4A 00 00 00 00 00 01 00 00 00 30 00 01 00 08 03 00 00
00 04 00 00 D4 00  k?J.........0...........?.
  01B9FDCE  00 00 00 00 01 00 01 00 02 00 34 01 00 00 80 00 00 00 D0 02
00 00 01 00 00 00  ..........4...?...?.......
  01B9FDE8  01 00 00 00 15 04 00 00 00 00 00 00 FD 09 00 00 00 00 00 00
40 00 00 00 00 00  ............? ......@.....            (Extension
Record 2)
  01B9FE02  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00  ..........................
  01B9FE1C  01 0B 21 02 A7 7C 11 03 06 01 0B 31 05 09 9A 02 01 0B 11 04
```

```
44 01 0C 31 04 B6  ..!.?|.....1. s.....D..1.?


  4.. So ... what would this mean?
  5.. A single runlist all through the various data attributes? In which
case, is the first address after the split case (in red: A7 7C etc.), a
Lcn Address Offset or a proper Lcn Address?
  6.. Or, multiple runlists each with the first address being an Lcn
Address and thereafter all Offsets? If so, how come one compression
unit, supposed to belong to the same runlist, getting split into two
records?




  01B9FE36  65 FD 01 0C 31 05 06 9B 02 01 0B 11 04 44 01 0C 31 05 BA 64
FD 01 0B 31 05 02  e?..1..>.....D..1.?d?..1..
  01B9FE50  9C 02 01 0B 11 04 44 01 0C 31 05 C8 63 FD 01 0B 31 05 F4 9C
02 01 0B 11 04 44  o.....D..1.?c?..1.?o.....D
  01B9FE6A  01 0C 31 02 CD 62 FD 11 02 06 01 0C 31 04 E9 9D 02 01 0C 11
04 44 01 0C 31 02  ..1.?b?.....1.?......D..1.
  01B9FE84  D5 61 FD 11 02 07 01 0C 31 04 10 9F 02 01 0C 11 04 44 01 0C
31 01 AE 60 FD 11  ?a?.....1..
```

Posting Rules

You **may not** post new threads
You **may not** post replies
You **may not** post attachments

You must login to OSR Online AND be a member of the ntfsd list to be able to post.

All times are GMT -5. The time now is 17:30.

Contact Us - **Osr Online Homepage** - **Top**