AFF4-L: A scalable open logical evidence container

By
Dr. Bradley Schatz

From the proceedings of
The Digital Forensic Research Conference
DFRWS 2019 USA
Portland, OR (July 15th - 19th)

Evimetry

Digital forensics at wire speed

AFF4-L
A scalable open logical evidence container

Dr. Bradley Schatz

Director, Schatz Forensic

V1.1 – Techno Security Myrtle Beach 2019

This seminar

- Background on AFF4
- Logical imaging
- AFF4 logical imaging
- AFF4 *deduplicated* logical imaging
- Evaluation

Background

Sealed Digital Evidence Bags (Schatz, 2006):

Cross-container referencing for forensic images

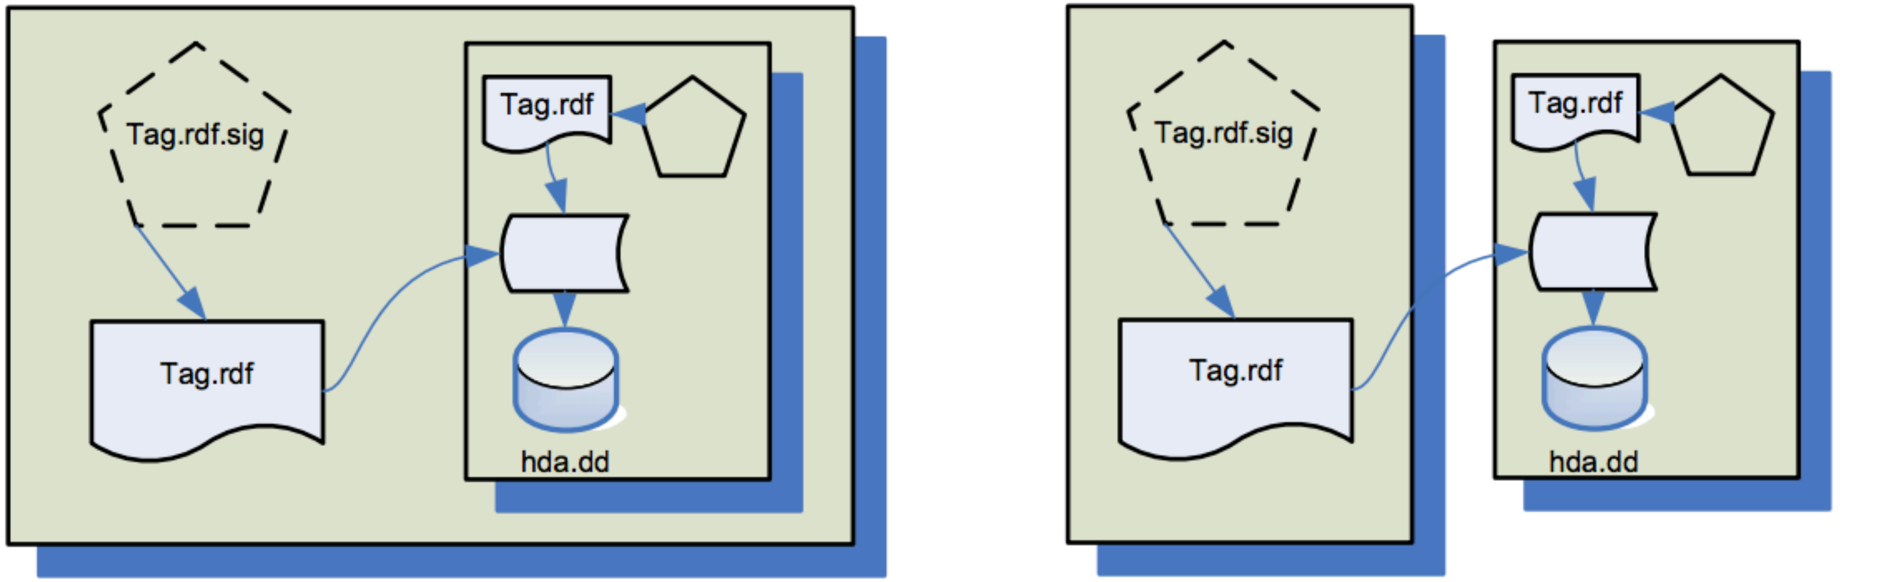

Embedded DEB

Referenced DEB

The Advanced Forensic Format (Garfinkel, 2006):

A vendor neutral next-generation format

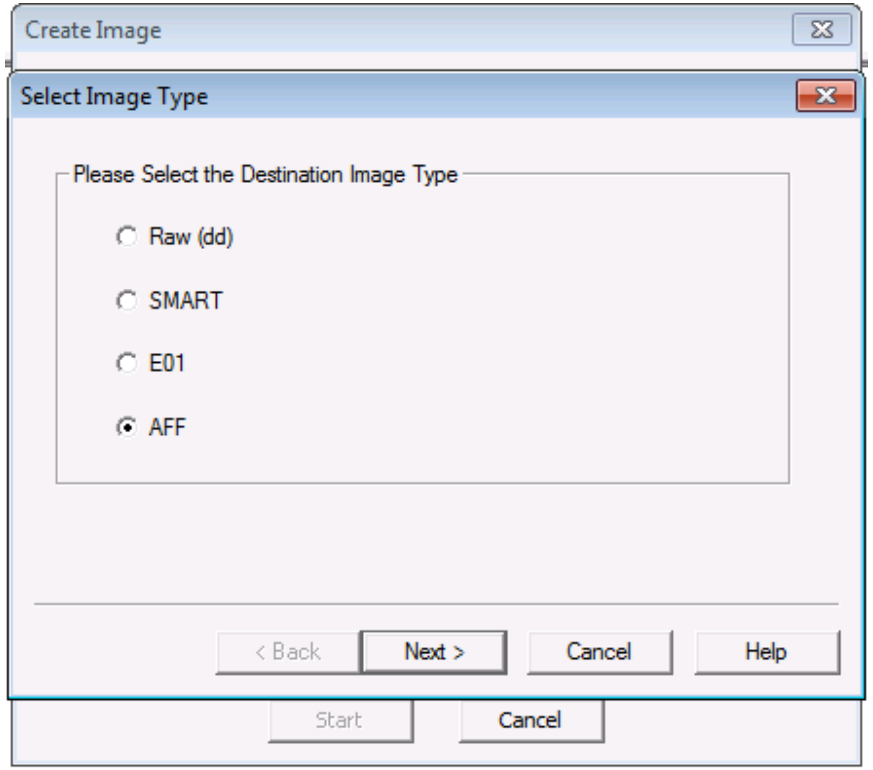

Create Image

Select Image Type

Please Select the Destination Image Type

○ Raw (dd)

○ SMART

○ E01

● AFF

- Good
  - Well defined format
  - Open source
  - Extensible Name/Value pair metadata storage
- Bad
  - **Large compressed chunk sizes (16M by default) slow w/ NTFS MFT**

AFF4 (Cohen, Garfinkel, Schatz 2009)

Virtualisation + Efficient compressed block store + Arbitrary linked data

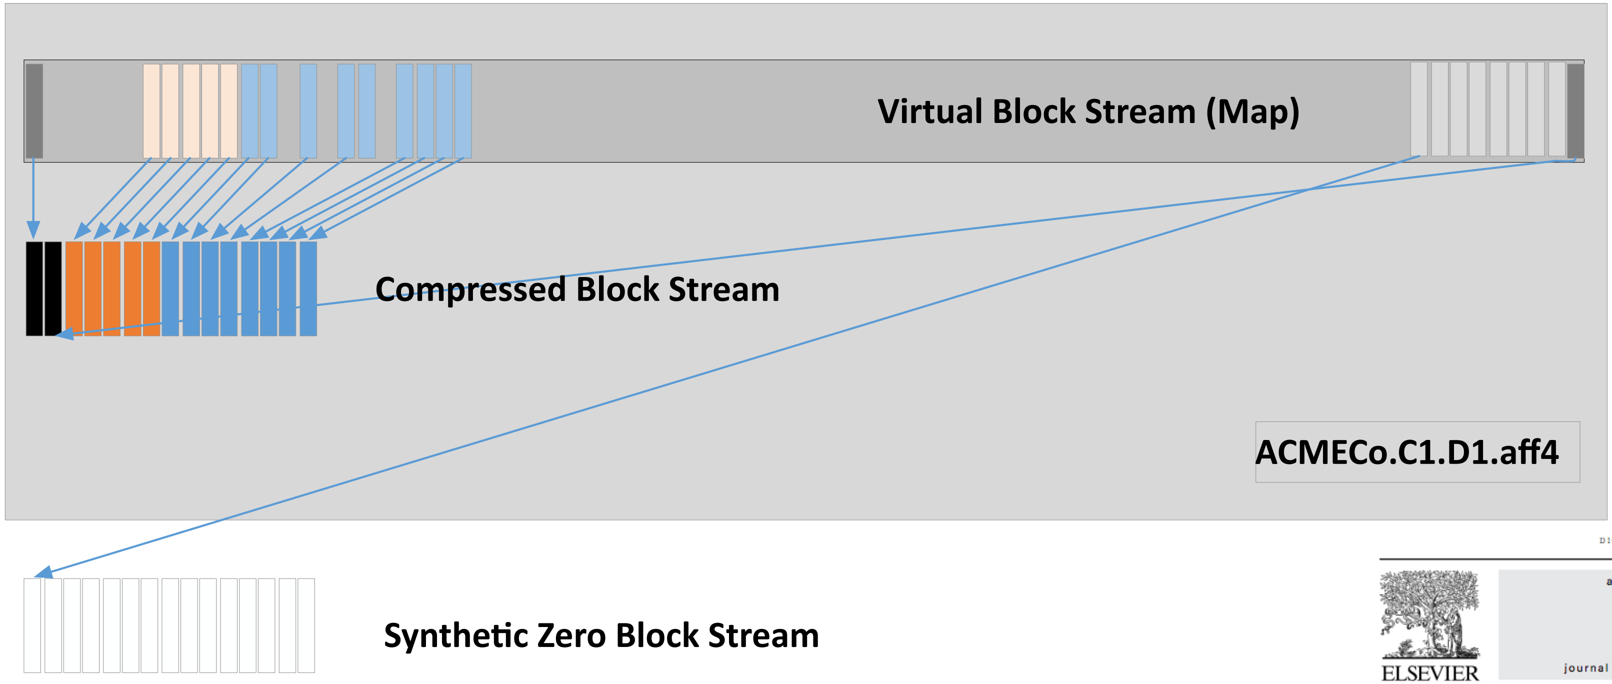

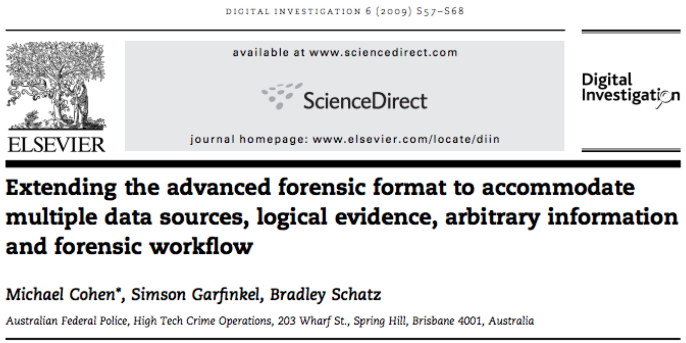

DIGITAL INVESTIGATION 6 (2009) S57–S68

available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/diin

ELSEVIER

Digital Investigation

Extending the advanced forensic format to accommodate multiple data sources, logical evidence, arbitrary information and forensic workflow

Michael Cohen, Simson Garfinkel, Bradley Schatz*

Australian Federal Police, High Tech Crime Operations, 203 Wharf St., Spring Hill, Brisbane 4001, Australia

The Compressed Block Stream (in-ZIP)

```
<aff4://c215ba20-5648-4209-a793-1f918c723610>
        a                          aff4:ImageStream ;
        aff4:chunkSize             "32768"^^xsd:int ;
        aff4:chunksInSegment       "2048"^^xsd:int ;
        aff4:compressionMethod     <http://code.google.com/p/snappy/> ;
        aff4:hash                  "fbac22cca549310bc5df03b7560afcf490995fbb"^^aff4:SHA1 ,
"d5825dc1152a42958c8219ff11ed01a3"^^aff4:MD5 ;
        aff4:imageStreamHash       "7c909ad458a90ca083cf2d10848fb3aaee7d9ac008605f85aef1ac2db8249973ac7b6716f3250edb80219ff628d6fb
4873c33c59de0a3e6c7657e234e7ba0db3"^^aff4:SHA512 ;
        aff4:imageStreamIndexHash  "c663bc90d996d2c9699e00dc1ea2c55b3724f1eaca2b92119bb7c764aad222eed321cb00ee67899c027f6837a3bd8f
789a96adb6e9df51629b3cac0b6f9f0722"^^aff4:SHA512 ;
        aff4:size                  "3964928"^^xsd:long ;
        aff4:stored                : ;
        aff4:target                <aff4://fcbfdce7-4488-4677-abf6-08bc931e195b> ;
        aff4:version               "1"^^xsd:int .
```

```
aff4://685e15cc-d0fb-4dbc-ba47-48117fc77044
  Length      Date    Time    Name
---------  ---------- -----   ----
       43  11-07-2016 13:40   container.description
       36  11-07-2016 13:40   version.txt
  3047794  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000
     1936  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.md5
     2420  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.sha1
     1452  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.index
   114884  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/map
      152  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/idx
        6  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/mapPath
     6580  11-07-2016 13:40   information.turtle
---------                     -------
  3175303                     10 files
```

The Virtual Block Stream (in-ZIP)

```
<aff4://c215ba20-5648-4209-a793-1f918c723610>
        a                          aff4:ImageStream ;
        aff4:chunkSize             "32768"^^xsd:int ;
        aff4:chunksInSegment       "2048"^^xsd:int ;
        aff4:compressionMethod     <http://code.google.com/p/snappy/> ;
        aff4:hash                  "fbac22cca549310bc5df03b7560afcf490995fbb"^^aff4:SHA1 ,
"d5825dc1152a42958c8219ff11ed01a3"^^aff4:MD5 ;
        aff4:imageStreamHash       "7c909ad458a90ca083cf2d10848fb3aaee7d9ac008605f85aef1ac2db8249973ac7b6716f3250edb80219ff628d6fb
4873c33c59de0a3e6c7657e234e7ba0db3"^^aff4:SHA512 ;
        aff4:imageStreamIndexHash  "c663bc90d996d2c9699e00dc1ea2c55b3724f1eaca2b92119bb7c764aad222eed321cb00ee67899c027f6837a3bd8f
789a96adb6e9df51629b3cac0b6f9f0722"^^aff4:SHA512 ;
        aff4:size                  "3964928"^^xsd:long ;
        aff4:stored                : ;
        aff4:target                <aff4://fcbfdce7-4488-4677-abf6-08bc931e195b> ;
        aff4:version               "1"^^xsd:int .
```

```
aff4://685e15cc-d0fb-4dbc-ba47-48117fc77044
  Length      Date    Time    Name
---------  ---------- -----   ----
       43  11-07-2016 13:40   container.description
       36  11-07-2016 13:40   version.txt
  3047794  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000
     1936  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.md5
     2420  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.sha1
     1452  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.index
   114884  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/map
      152  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/idx
        6  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/mapPath
     6580  11-07-2016 13:40   information.turtle
---------                     -------
  3175303                     10 files
```

Metadata & information is represented in RDF

```
<aff4://c215ba20-5648-4209-a793-1f918c723610>
        a                          aff4:ImageStream ;
        aff4:chunkSize             "32768"^^xsd:int ;
        aff4:chunksInSegment       "2048"^^xsd:int ;
        aff4:compressionMethod     <http://code.google.com/p/snappy/> ;
        aff4:hash                  "fbac22cca549310bc5df03b7560afcf490995fbb"^^aff4:SHA1 ,
"d5825dc1152a42958c8219ff11ed01a3"^^aff4:MD5 ;
        aff4:imageStreamHash       "7c909ad458a90ca083cf2d10848fb3aaee7d9ac008605f85aef1ac2db8249973ac7b6716f3250edb80219ff628d6fb
4873c33c59de0a3e6c7657e234e7ba0db3"^^aff4:SHA512 ;
        aff4:imageStreamIndexHash  "c663bc90d996d2c9699e00dc1ea2c55b3724f1eaca2b92119bb7c764aad222eed321cb00ee67899c027f6837a3bd8f
789a96adb6e9df51629b3cac0b6f9f0722"^^aff4:SHA512 ;
        aff4:size                  "3964928"^^xsd:long ;
        aff4:stored                : ;
        aff4:target                <aff4://fcbfdce7-4488-4677-abf6-08bc931e195b> ;
        aff4:version               "1"^^xsd:int .
```

```
aff4://685e15cc-d0fb-4dbc-ba47-48117fc77044
  Length      Date    Time    Name
---------  ---------- -----   ----
       43  11-07-2016 13:40   container.description
       36  11-07-2016 13:40   version.txt
  3047794  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000
     1936  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.md5
     2420  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.sha1
     1452  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.index
   114884  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/map
      152  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/idx
        6  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/mapPath
     6580  11-07-2016 13:40   information.turtle
---------                     -------
  3175303                     10 files
```

AFF4 Objects = Metadata and/or data streams

```
<aff4://c215ba20-5648-4209-a793-1f918c723610>
        a                          aff4:ImageStream ;
        aff4:chunkSize             "32768"^^xsd:int ;
        aff4:chunksInSegment       "2048"^^xsd:int ;
        aff4:compressionMethod     <http://code.google.com/p/snappy/> ;
        aff4:hash                  "fbac22cca549310bc5df03b7560afcf490995fbb"^^aff4:SHA1 ,
"d5825dc1152a42958c8219ff11ed01a3"^^aff4:MD5 ;
        aff4:imageStreamHash       "7c909ad458a90ca083cf2d10848fb3aaee7d9ac008605f85aef1ac2db8249973ac7b6716f3250edb80219ff628d6fb
4873c33c59de0a3e6c7657e234e7ba0db3"^^aff4:SHA512 ;
        aff4:imageStreamIndexHash  "c663bc90d996d2c9699e00dc1ea2c55b3724f1eaca2b92119bb7c764aad222eed321cb00ee67899c027f6837a3bd8f
789a96adb6e9df51629b3cac0b6f9f0722"^^aff4:SHA512 ;
        aff4:size                  "3964928"^^xsd:long ;
        aff4:stored                : ;
        aff4:target                <aff4://fcbfdce7-4488-4677-abf6-08bc931e195b> ;
        aff4:version               "1"^^xsd:int .

aff4://685e15cc-d0fb-4dbc-ba47-48117fc77044
  Length      Date    Time    Name
---------  ---------- -----   ----
       43  11-07-2016 13:40   container.description
       36  11-07-2016 13:40   version.txt
  3047794  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000
     1936  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.md5
     2420  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.sha1
     1452  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.index
   114884  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/map
      152  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/idx
        6  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/mapPath
     6580  11-07-2016 13:40   information.turtle
---------                     -------
  3175303                     10 files
```

AFF4 Objects = Metadata and/or data streams

```
<aff4://c215ba20-5648-4209-a793-1f918c723610>
        a                          aff4:ImageStream ;
        aff4:chunkSize             "32768"^^xsd:int ;
        aff4:chunksInSegment       "2048"^^xsd:int ;
        aff4:compressionMethod     <http://code.google.com/p/snappy/> ;
        aff4:hash                  "fbac22cca549310bc5df03b7560afcf490995fbb"^^aff4:SHA1 ,
"d5825dc1152a42958c8219ff11ed01a3"^^aff4:MD5 ;
        aff4:imageStreamHash       "7c909ad458a90ca083cf2d10848fb3aaee7d9ac008605f85aef1ac2db8249973ac7b6716f3250edb80219ff628d6fb
4873c33c59de0a3e6c7657e234e7ba0db3"^^aff4:SHA512 ;
        aff4:imageStreamIndexHash  "c663bc90d996d2c9699e00dc1ea2c55b3724f1eaca2b92119bb7c764aad222eed321cb00ee67899c027f6837a3bd8f
789a96adb6e9df51629b3cac0b6f9f0722"^^aff4:SHA512 ;
        aff4:size                  "3964928"^^xsd:long ;
        aff4:stored                : ;
        aff4:target                <aff4://fcbfdce7-4488-4677-abf6
        aff4:version               "1"^^xsd:int .
```

Note the URL encoding

```
aff4://685e15cc-d0fb-4dbc-ba47-48117fc77044
  Length      Date    Time    Name
---------  ---------- -----   ----
       43  11-07-2016 13:40   container.description
       36  11-07-2016 13:40   version.txt
  3047794  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000
     1936  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.md5
     2420  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.blockHash.sha1
     1452  11-07-2016 13:40   aff4%3A%2F%2Fc215ba20-5648-4209-a793-1f918c723610/00000000.index
   114884  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/map
      152  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/idx
        6  11-07-2016 13:40   aff4%3A%2F%2Ffcbfdce7-4488-4677-abf6-08bc931e195b/mapPath
     6580  11-07-2016 13:40   information.turtle
---------                     -------
  3175303                     10 files
```

Evimetry
Digital forensics at wire speed

Rethinking logical imaging

Time for an open logical imaging format

Acquisition challenges increase as we go up the stack

Logical Imaging

- No currently widely adopted standard for interoperability
  - L01, AD1, TGZ, ZIP…
- All approaches preserve less metadata than is desirable
  - e.g. File birth time

Research Goals

- Human interpretability using regular Zip tools
- Efficient access for large logical files
- Arbitrary metadata

Goal: make AFF4 Logical Images viewable in 7Zip and WinRAR

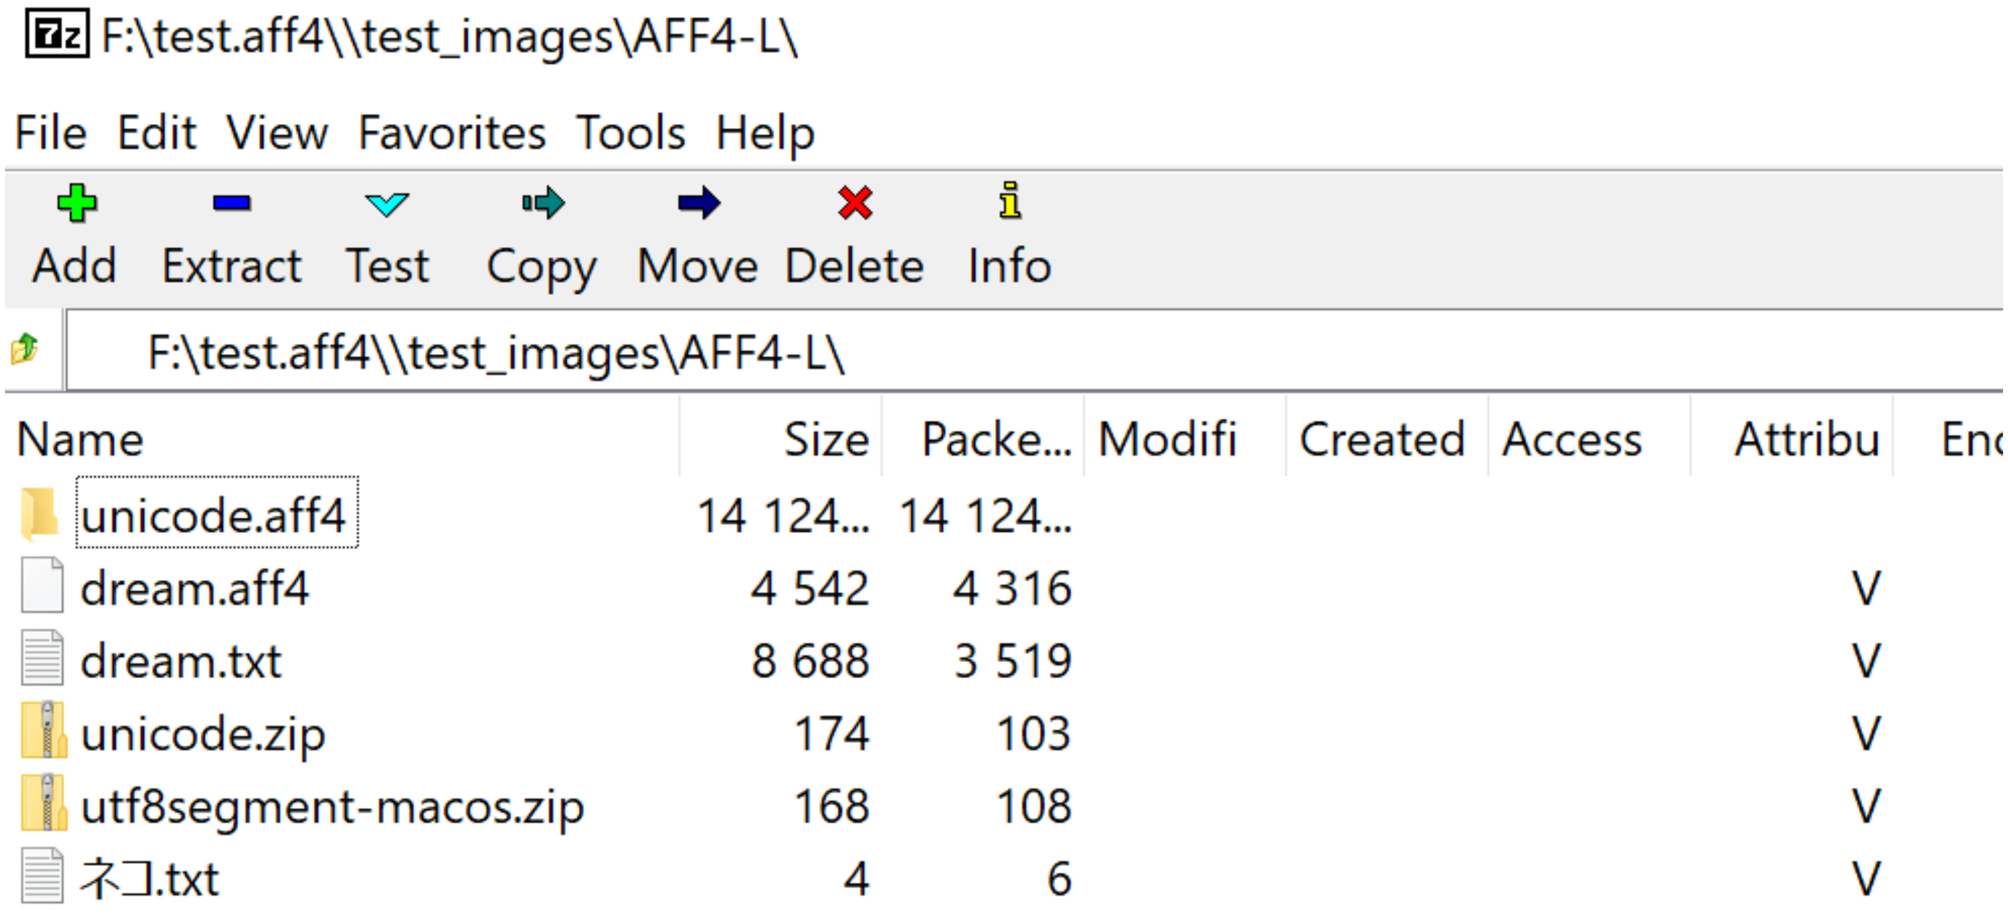

F:\test.aff4\\test_images\AFF4-L\

F:\test.aff4\\test_images\AFF4-L\

Name	Size	Packe...	Modifi	Created	Access	Attribu	Enc
unicode.aff4	14 124...	14 124...					
dream.aff4	4 542	4 316				V	
dream.txt	8 688	3 519				V	
unicode.zip	174	103				V	
utf8segment-macos.zip	168	108				V	
ネコ.txt	4	6				V	

Efficiency concerns: small vs large

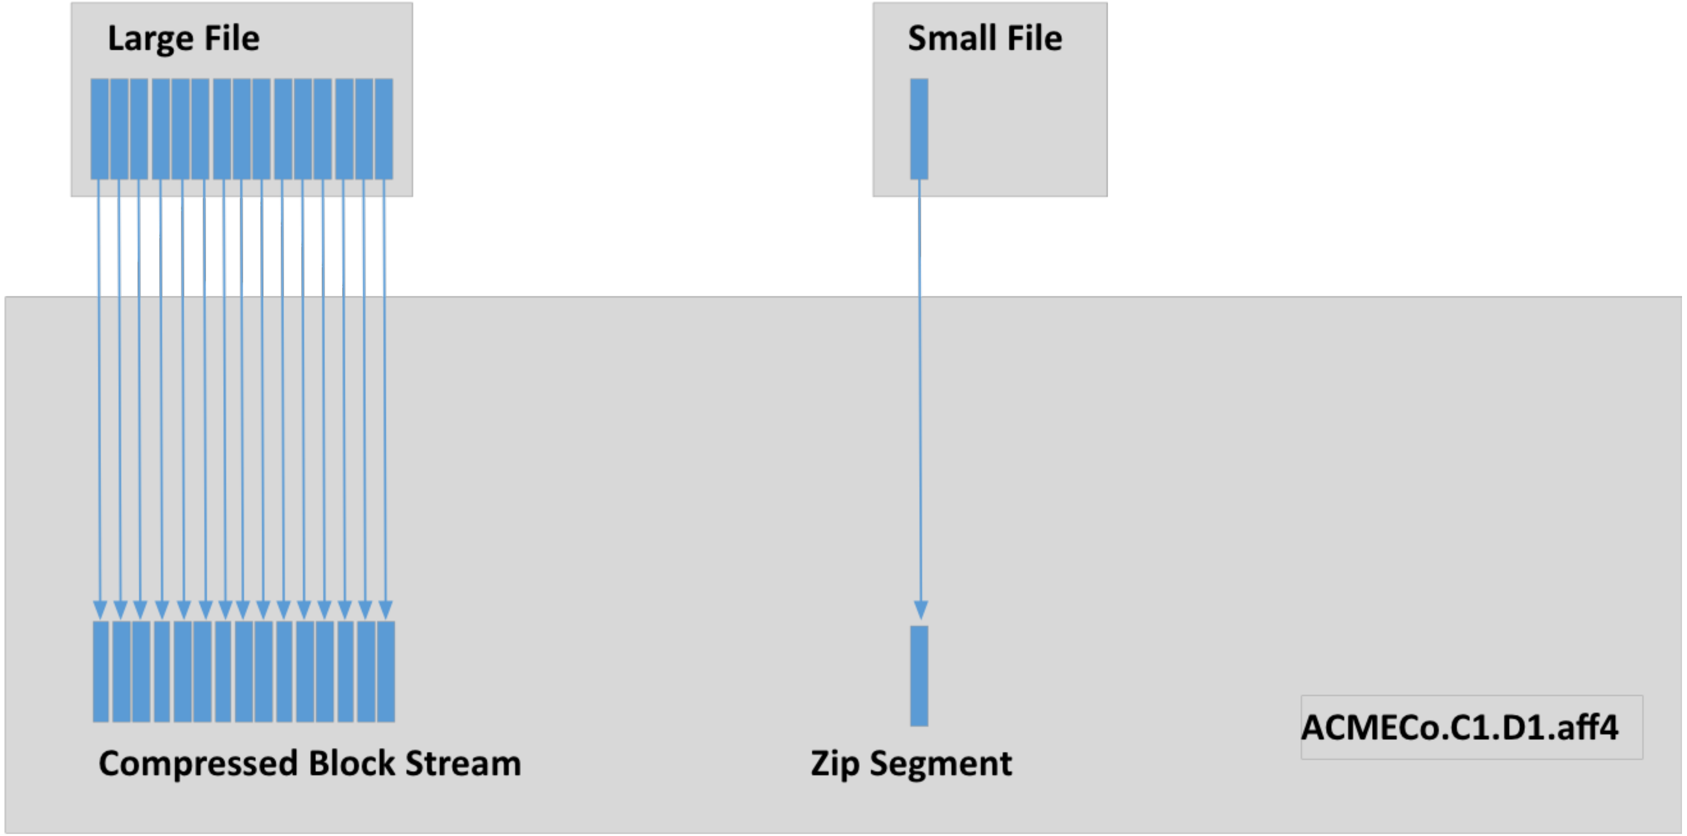

How do we name things in AFF4?

Evolution

- 2009: Used URN for identifying AFF4 objects
  - urn:aff4:f901be8e-d4b2…
  - But the definition of a URN was overly specific
- 2010: Shifted to URL as the identifier of AFF4 objects
  - aff4://f901be8e-d4b2…
  - Percent encoded the "://" in the ZIP layer
  - But the definition of the URL is overly specific for representing logical files
    - Query syntax is valid in some file names "?" as is fragment "#"
  - Also the URL syntax is meant to represent a location and not a name…

- Conclusion
  - We need our own IRI scheme: The AFF4 Resource Identifier

We need our own IRI name scheme: ARN

- ***AFF4-ARN = "aff4://" object-guid-part [“/” host "/" path***
- host and path
  - may contain any Unicode character that is not forbidden in the IRI specification (“/” is used as a path delimiter).
  - Forbidden printable characters
    - <>\^`{|}
    - Percent Encoded

Naming is the primary challenge

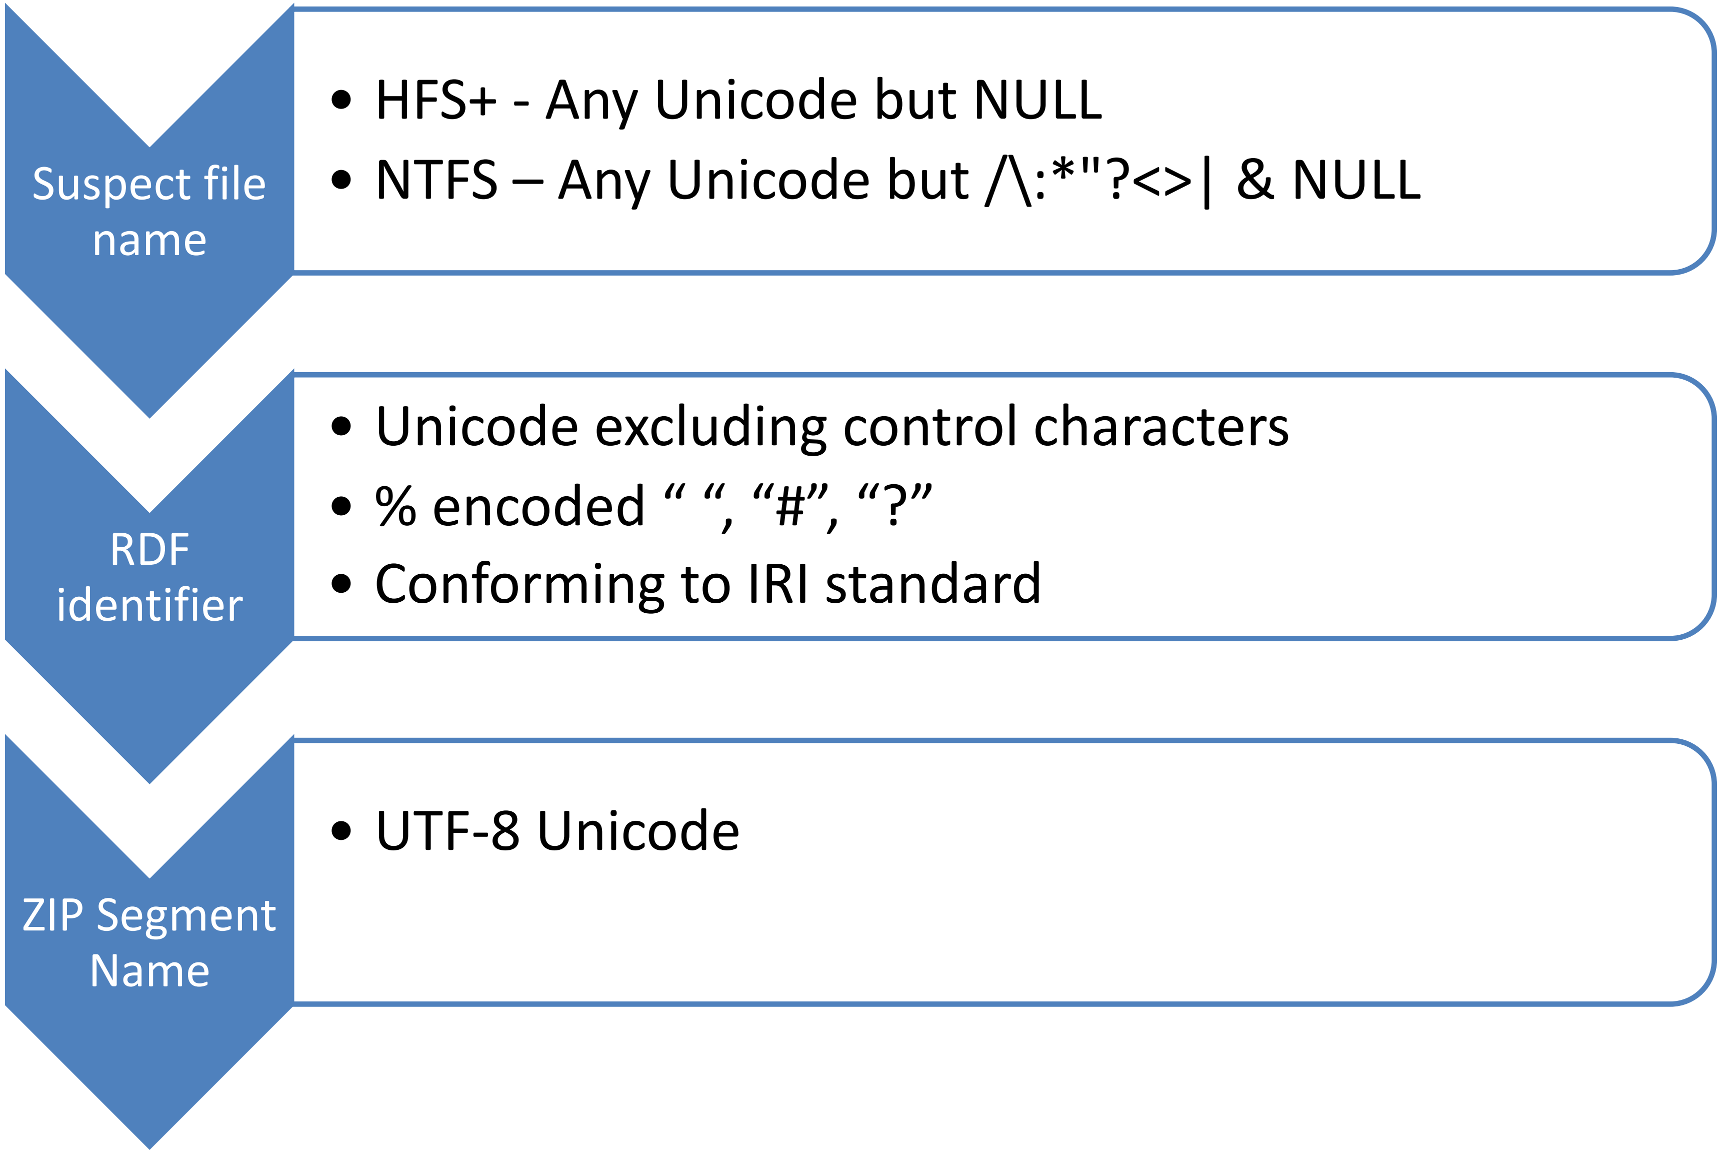

Suspect file name

- HFS+ - Any Unicode but NULL
- NTFS – Any Unicode but /\:*"?<>| & NULL

RDF identifier

- Unicode excluding control characters
- % encoded “ “, “#”, “?”
- Conforming to IRI standard

ZIP Segment Name

- UTF-8 Unicode

Suspect path to ARN Mapping

Similar to the file:// protocol

OS Path	AFF4 Resource Name
c:	aff4://e6bae91b-14d231833e18//c:
c:\	aff4://e6bae91b-14d231833e18//c:/
c:\foo	aff4://e6bae91b-14d231833e18//c:/foo
\\bar\c$	aff4://e6bae91b-14d231833e18/bar/c$
\\bar\c$\foo\ ネコ.txt	aff4://e6bae91b-14d231833e18/bar/c$/foo/ ネコ.txt
/foo/bar	aff4://e6bae91b-14d231833e18//foo/bar
/foo/some file	aff4://e6bae91b-14d231833e18//foo/some%20file

ARN to Zip segment name mapping

AFF4 Resource Name	Zip segment name
aff4://e6bae91b-14d231833e18//c:	/C:
aff4://e6bae91b-14d231833e18//c:/	/C:/
aff4://e6bae91b-14d231833e18//c:/foo	/C:/foo
aff4://e6bae91b-14d231833e18/bar/c$	bar/c$/foo
aff4://e6bae91b-14d231833e18/bar/c$/foo/ネコ.txt	bar/c$/foo/ネコ.txt
aff4://e6bae91b-14d231833e18//foo/bar	/foo/bar
aff4://e6bae91b-14d231833e18//foo/some%20file	/foo/some file

Example RDF

```
</test_images/AFF4-L/dream.txt> a aff4:FileImage,
        aff4:Image,
        aff4:ImageStream ;
    aff4:birthTime "2018-09-17T13:42:20+10:00"^^xsd:datetime ;
    aff4:chunkSize 32768 ;
    aff4:chunksInSegment 1024 ;
    aff4:compressionMethod <http://code.google.com/p/snappy/> ;
    aff4:hash "75d83773f8d431a3ca91bfb8859e486d"^^aff4:MD5, "9ae1b46bead70c322eef7ac8bc36a8ea2055595c"^^aff4:SHA1 ;
    aff4:lastAccessed "2018-09-30T11:18:27+10:00"^^xsd:datetime ;
    aff4:lastWritten "2018-09-17T13:42:20+10:00"^^xsd:datetime ;
    aff4:originalFileName "./test_images/AFF4-L/dream.txt"^^xsd:string ;
    aff4:recordChanged "2018-09-17T13:42:20+10:00"^^xsd:datetime ;
    aff4:size 8688 .

</test_images/AFF4-L/ネコ.txt> a aff4:FileImage,
        aff4:Image,
        aff4:zip_segment ;
    aff4:birthTime "2018-09-18T15:49:51+10:00"^^xsd:datetime ;
    aff4:hash "d3b07384d113edec49eaa6238ad5ff00"^^aff4:MD5, "f1d2d2f924e986ac86fdf7b36c94bcdf32beec15"^^aff4:SHA1 ;
    aff4:lastAccessed "2018-09-30T11:18:34+10:00"^^xsd:datetime ;
    aff4:lastWritten "2018-09-18T15:49:51+10:00"^^xsd:datetime ;
    aff4:originalFileName "./test_images/AFF4-L/ネコ.txt"^^xsd:string ;
    aff4:recordChanged "2018-09-18T15:49:51+10:00"^^xsd:datetime ;
    aff4:size 4 .
```

AFF4 Logical Imaging

Code available now in the pyaff4 github

```
git clone --recurse-submodules
https://github.com/aff4/pyaff4.git
python aff4.py -r --create-logical test.aff4
./test_images/AFF4-L/
Creating AFF4Container: file://test.aff4 <aff4://05e730d3-
f6de-4961-9e9a-a30d5043a562>
      Adding: ./test_images/AFF4-L/
      Adding: ./test_images/AFF4-L/dream.aff4
      Adding: ./test_images/AFF4-L/dream.txt
      Adding: ./test_images/AFF4-L/unicode.aff4
      Adding: ./test_images/AFF4-L/unicode.zip
      Adding: ./test_images/AFF4-L/utf8segment-macos.zip
      Adding: ./test_images/AFF4-L/ネコ.txt
```

Evimetry

Digital forensics at wire speed

Deduplicated logical imaging

Research Goals

- Extend deduplication to AFF4 logical imaging

Logical deduplication structure

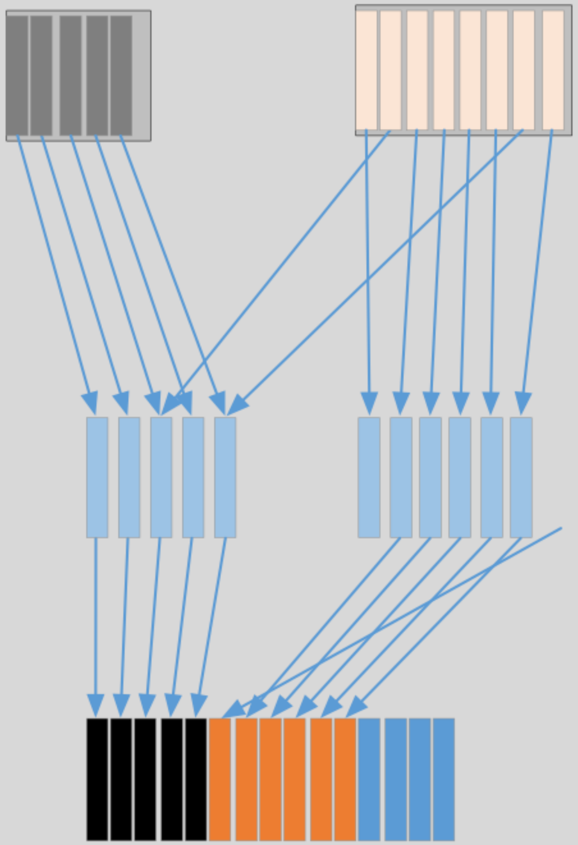

Logical Files (Maps)

Block Hashes
(Metadata)

Data blocks
(Compressed Block Stream)

ACMECo.C1.D1.aff4

Logical deduplication structure

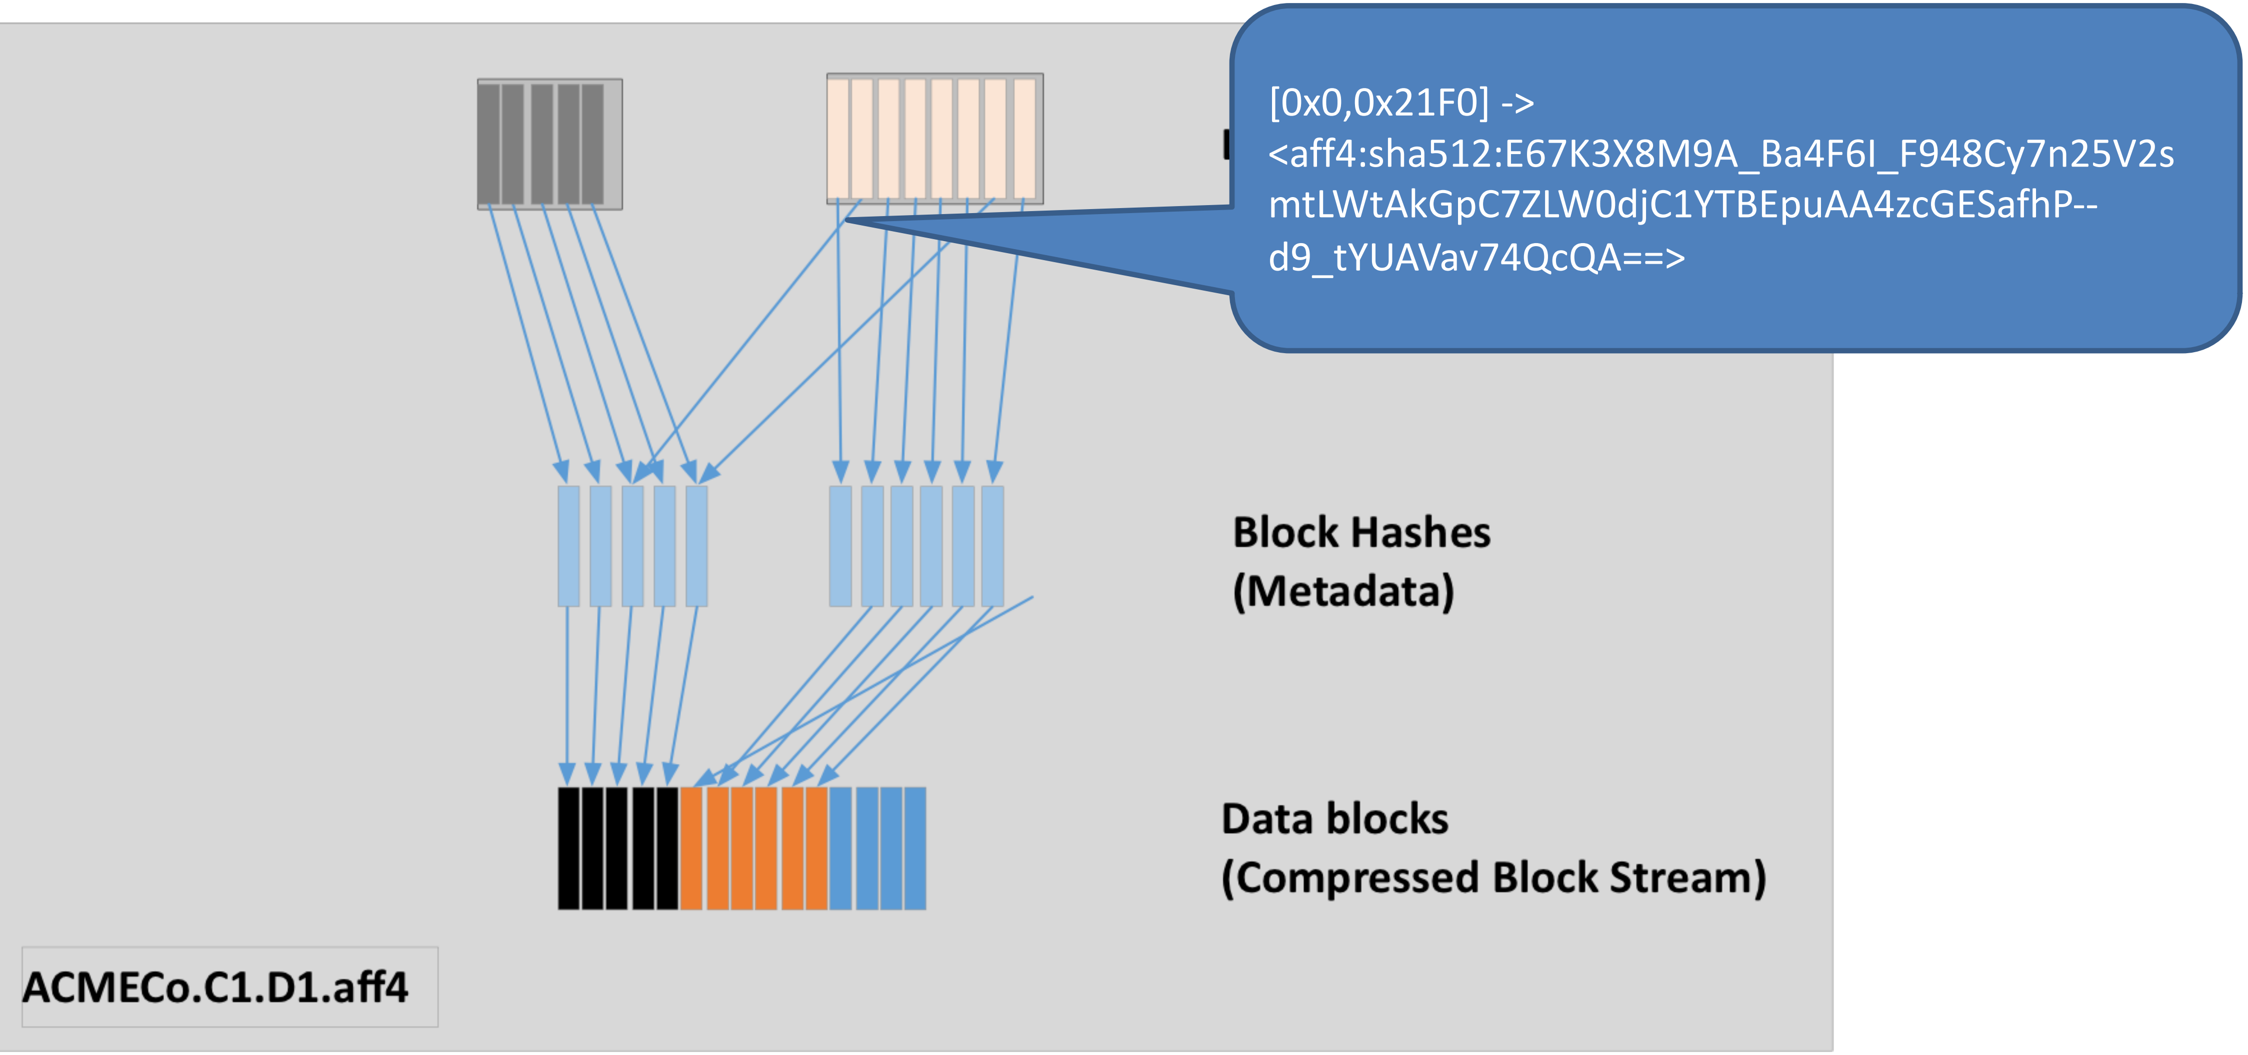

[0x0,0x21F0] ->
<aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2s
mtLWtAkGpC7ZLW0djC1YTBEpuAA4zcGESafhP--
d9_tYUAVav74QcQA==>

**Block Hashes
(Metadata)**

**Data blocks
(Compressed Block Stream)**

ACMECo.C1.D1.aff4

Logical deduplication structure

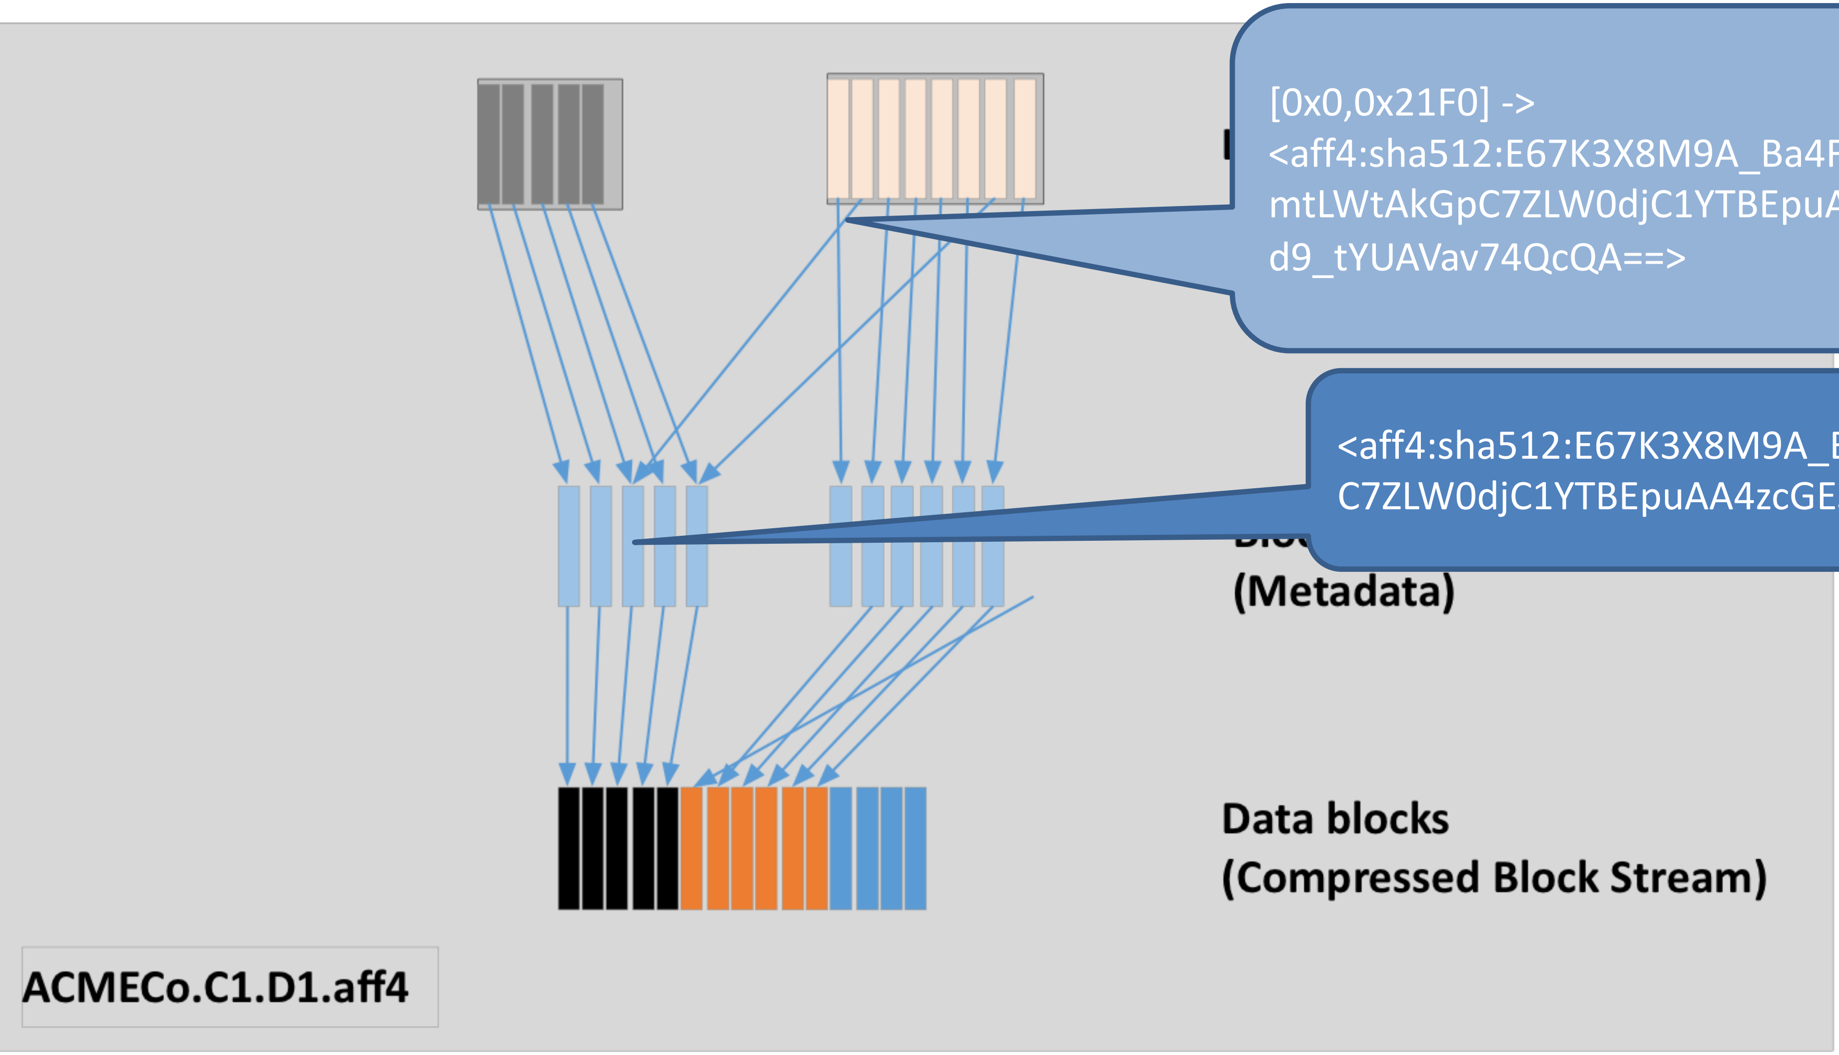

[0x0,0x21F0] -> <aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2smtLWtAkGpC7ZLW0djC1YTBEpuAA4zcGESafhP--d9_tYUAVav74QcQA==>

<aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2smtLWtAkGpC7ZLW0djC1YTBEpuAA4zcGESafhP--d9_tYUAVav74QcQA==>

Logical deduplication structure

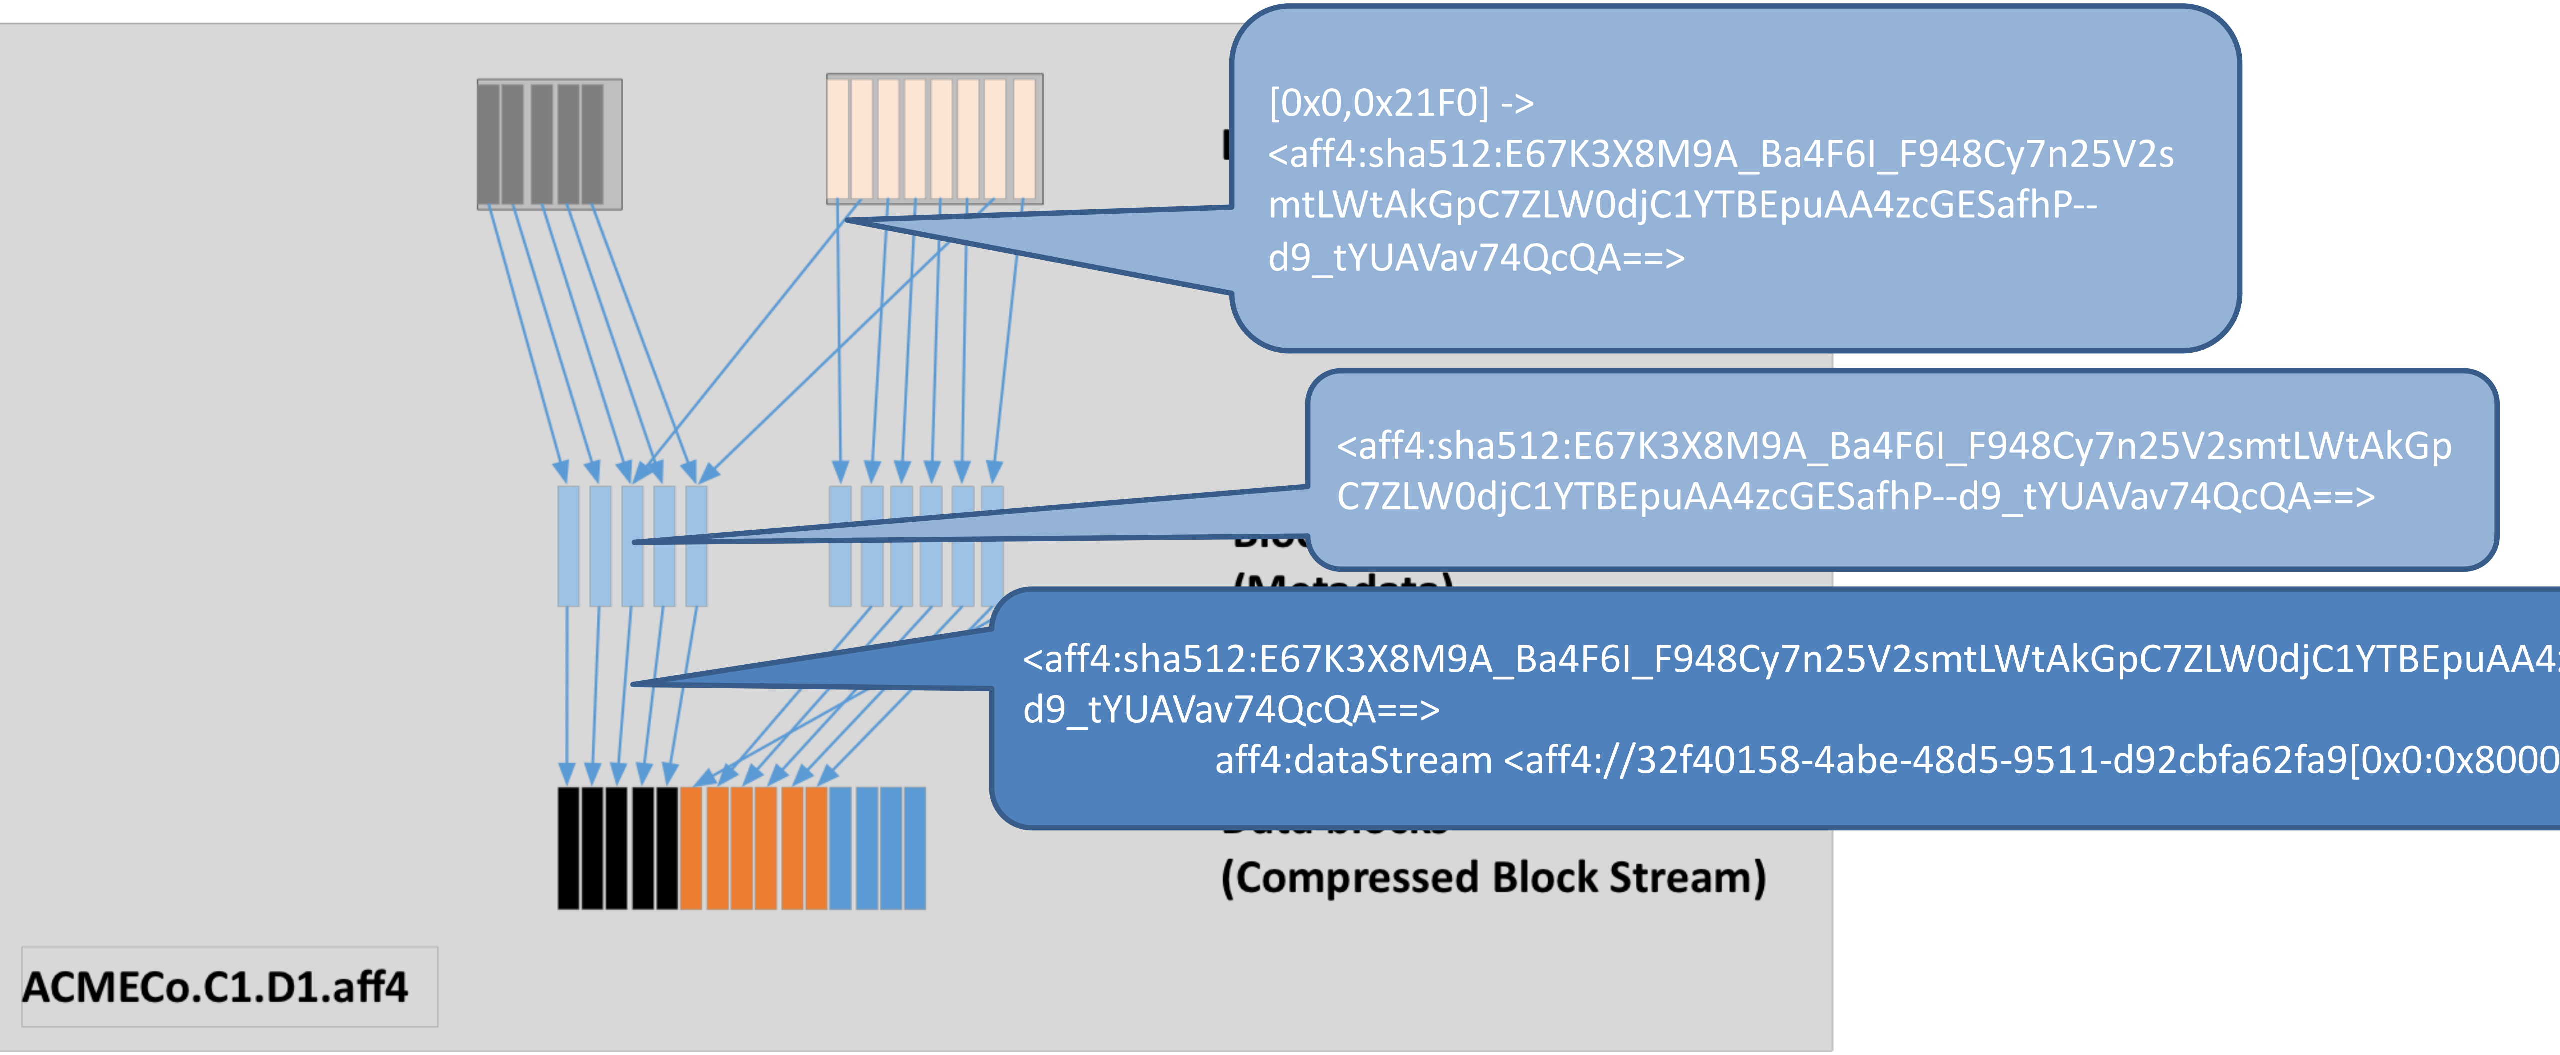

[0x0,0x21F0] ->
<aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2smtLWtAkGpC7ZLW0djC1YTBEpuAA4zcGESafhP--d9_tYUAVav74QcQA==>

<aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2smtLWtAkGpC7ZLW0djC1YTBEpuAA4zcGESafhP--d9_tYUAVav74QcQA==>

<aff4:sha512:E67K3X8M9A_Ba4F6I_F948Cy7n25V2smtLWtAkGpC7ZLW0djC1YTBEpuAA4
d9_tYUAVav74QcQA==>
aff4:dataStream <aff4://32f40158-4abe-48d5-9511-d92cbfa62fa9[0x0:0x8000

Evaluation

Opening large logical images took too long.

Img	Count files	RDF Triples	Container Size (GB)	First Access Latency	Subsequent Access Latency
A	19,463	228,287	1.5	33	33
B	21,835	236,235	1.9	39	39
C	41,298	461,220	3.4	67	67

A: Logical of Server 2012 / system32
B: Logical of Windows 10 / system32
C: Logical of A + B

Use a better RDF encoding?

- RDFHDT "RDF Header Data Triples"
  - A highly compressed indexed RDF encoding.
- Convert the RDF turtle on first open to RDFHDT
- Reuse cached encoding on subsequent opens

Choice of RDF serialization has major impacts for image consumers.

Img	Count files	RDF Triples	Container Size (GB)	Initial Access Latency RDFLib (s)	Initial access latency HDT (s)
A	19,463	228,287	1.5	33/33	3.8/0
B	21,835	236,235	1.9	39/39	4.4/0
C	41,298	461,220	3.4	67/67	8.9/0

A: Logical of Server 2012 / system32
B: Logical of Windows 10 / system32
C: Logical of A + B

Observations

- CASE/UCO uses the same representational approach as AFF4
  - – RDF (JSON-LD) + Python/RDFLib
  - – JSON-LD is equally as slow to load as Turtle
- Ontology design choices made now have far reaching effects on the number of triples stored

Future work

- Logical Imaging
  - Using “/” at the start of the ZIP Segment name violates the spec
  - The ARN mapping rules need to be fleshed out to identify more edge cases
    - Is it overly ambitious keeping the ZIP segment names human readable?
  - Standardisation
- Record based imaging
  - Sub file level imaging (web service calls, etc)
- Deduplication
  - Use different chunking algorithm (CDC, and others)
- AFF4-L as a container/transport for CASE/UCO

Acknowledgements

- NIST
  - Partial funding of the deduplication work
- Michael Cohen & Simson Garfinkel
  - Early AFF4 research collaborators
- Michael Cohen & Joe Sylve
  - Ongoing AFF4 standardisation collaborator

Evimetry

Digital forensics at wire speed

Contact

Dr Bradley Schatz
https://evimetry.com/
bradley@evimetry.com
@blschatz