

Navigating Unmountable Media with the Digital Forensics XML File System

Alex J. Nelson
NIST, Computer Security Division

CurateGear
January 14, 2016

Note: Any mention of a vendor or product is not an endorsement or recommendation.
Logos and trademarks are copyright their respective owners.




Problem:

File systems age out of usability.

- Operating systems can't support every file system that ever was, or will be.
 - HFS (classic Mac) - outdated
 - XTAF (XBox 360) - niche device support
- *How do we interact with disks that aren't or won't be the norm?*
 - Custom file system parsers?
Each custom parsing program has its own interface & software dependencies.

Proposal:

Treat a file listing like a file system.




```
aj@tmpubuntu1510: ~/mnt
aj@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests

./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
aj@tmpubuntu1510:~/mnt$
```


Proposal:


Treat a file listing like a file system.



```
aj@tmpubuntu1510: ~/mnt
aj@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests


./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
aj@tmpubuntu1510:~/mnt$
```



Proposal:


Treat a file listing like a file system.




```
ajn@tmpubuntu1510: ~/mnt
ajn@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests

./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
ajn@tmpubuntu1510:~/mnt$
```



- Navigate file metadata like a file system.




```
ajn@tmpubuntu1510: ~/mnt
ajn@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests

./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
ajn@tmpubuntu1510:~/mnt$
```

Proposal:


Treat a file listing like a file system.




```
ajn@tmpubuntu1510: ~/mnt
ajn@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests

./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
ajn@tmpubuntu1510:~/mnt$
```



- Navigate file metadata like a file system.



```
ajn@tmpubuntu1510: ~/mnt
ajn@tmpubuntu1510:~/mnt$ ls -lR
.:
total 16
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 hfs2dfxml
-rw-rw-r-- 1 ajn ajn 1482 Nov 16 12:39 LICENSE.md
-rw-rw-r-- 1 ajn ajn 3104 Nov 16 12:39 README.md
drwxrwxr-x 2 ajn ajn 4096 Nov 16 12:39 tests

./hfs2dfxml:
total 28
-rw-rw-r-- 1 ajn ajn 1212 Nov 16 12:39 debug_raw_hfs.py
-rw-rw-r-- 1 ajn ajn 18610 Nov 16 12:39 hfs2dfxml.py
-rw-rw-r-- 1 ajn ajn 1 Nov 16 12:39 __init__.py

./tests:
total 4
-rw-rw-r-- 1 ajn ajn 1062 Nov 16 12:39 hfs2dfxml_tests.py
ajn@tmpubuntu1510:~/mnt$
```

- With DFXML *and* disk image, get file contents from a disk without mounting the disk.

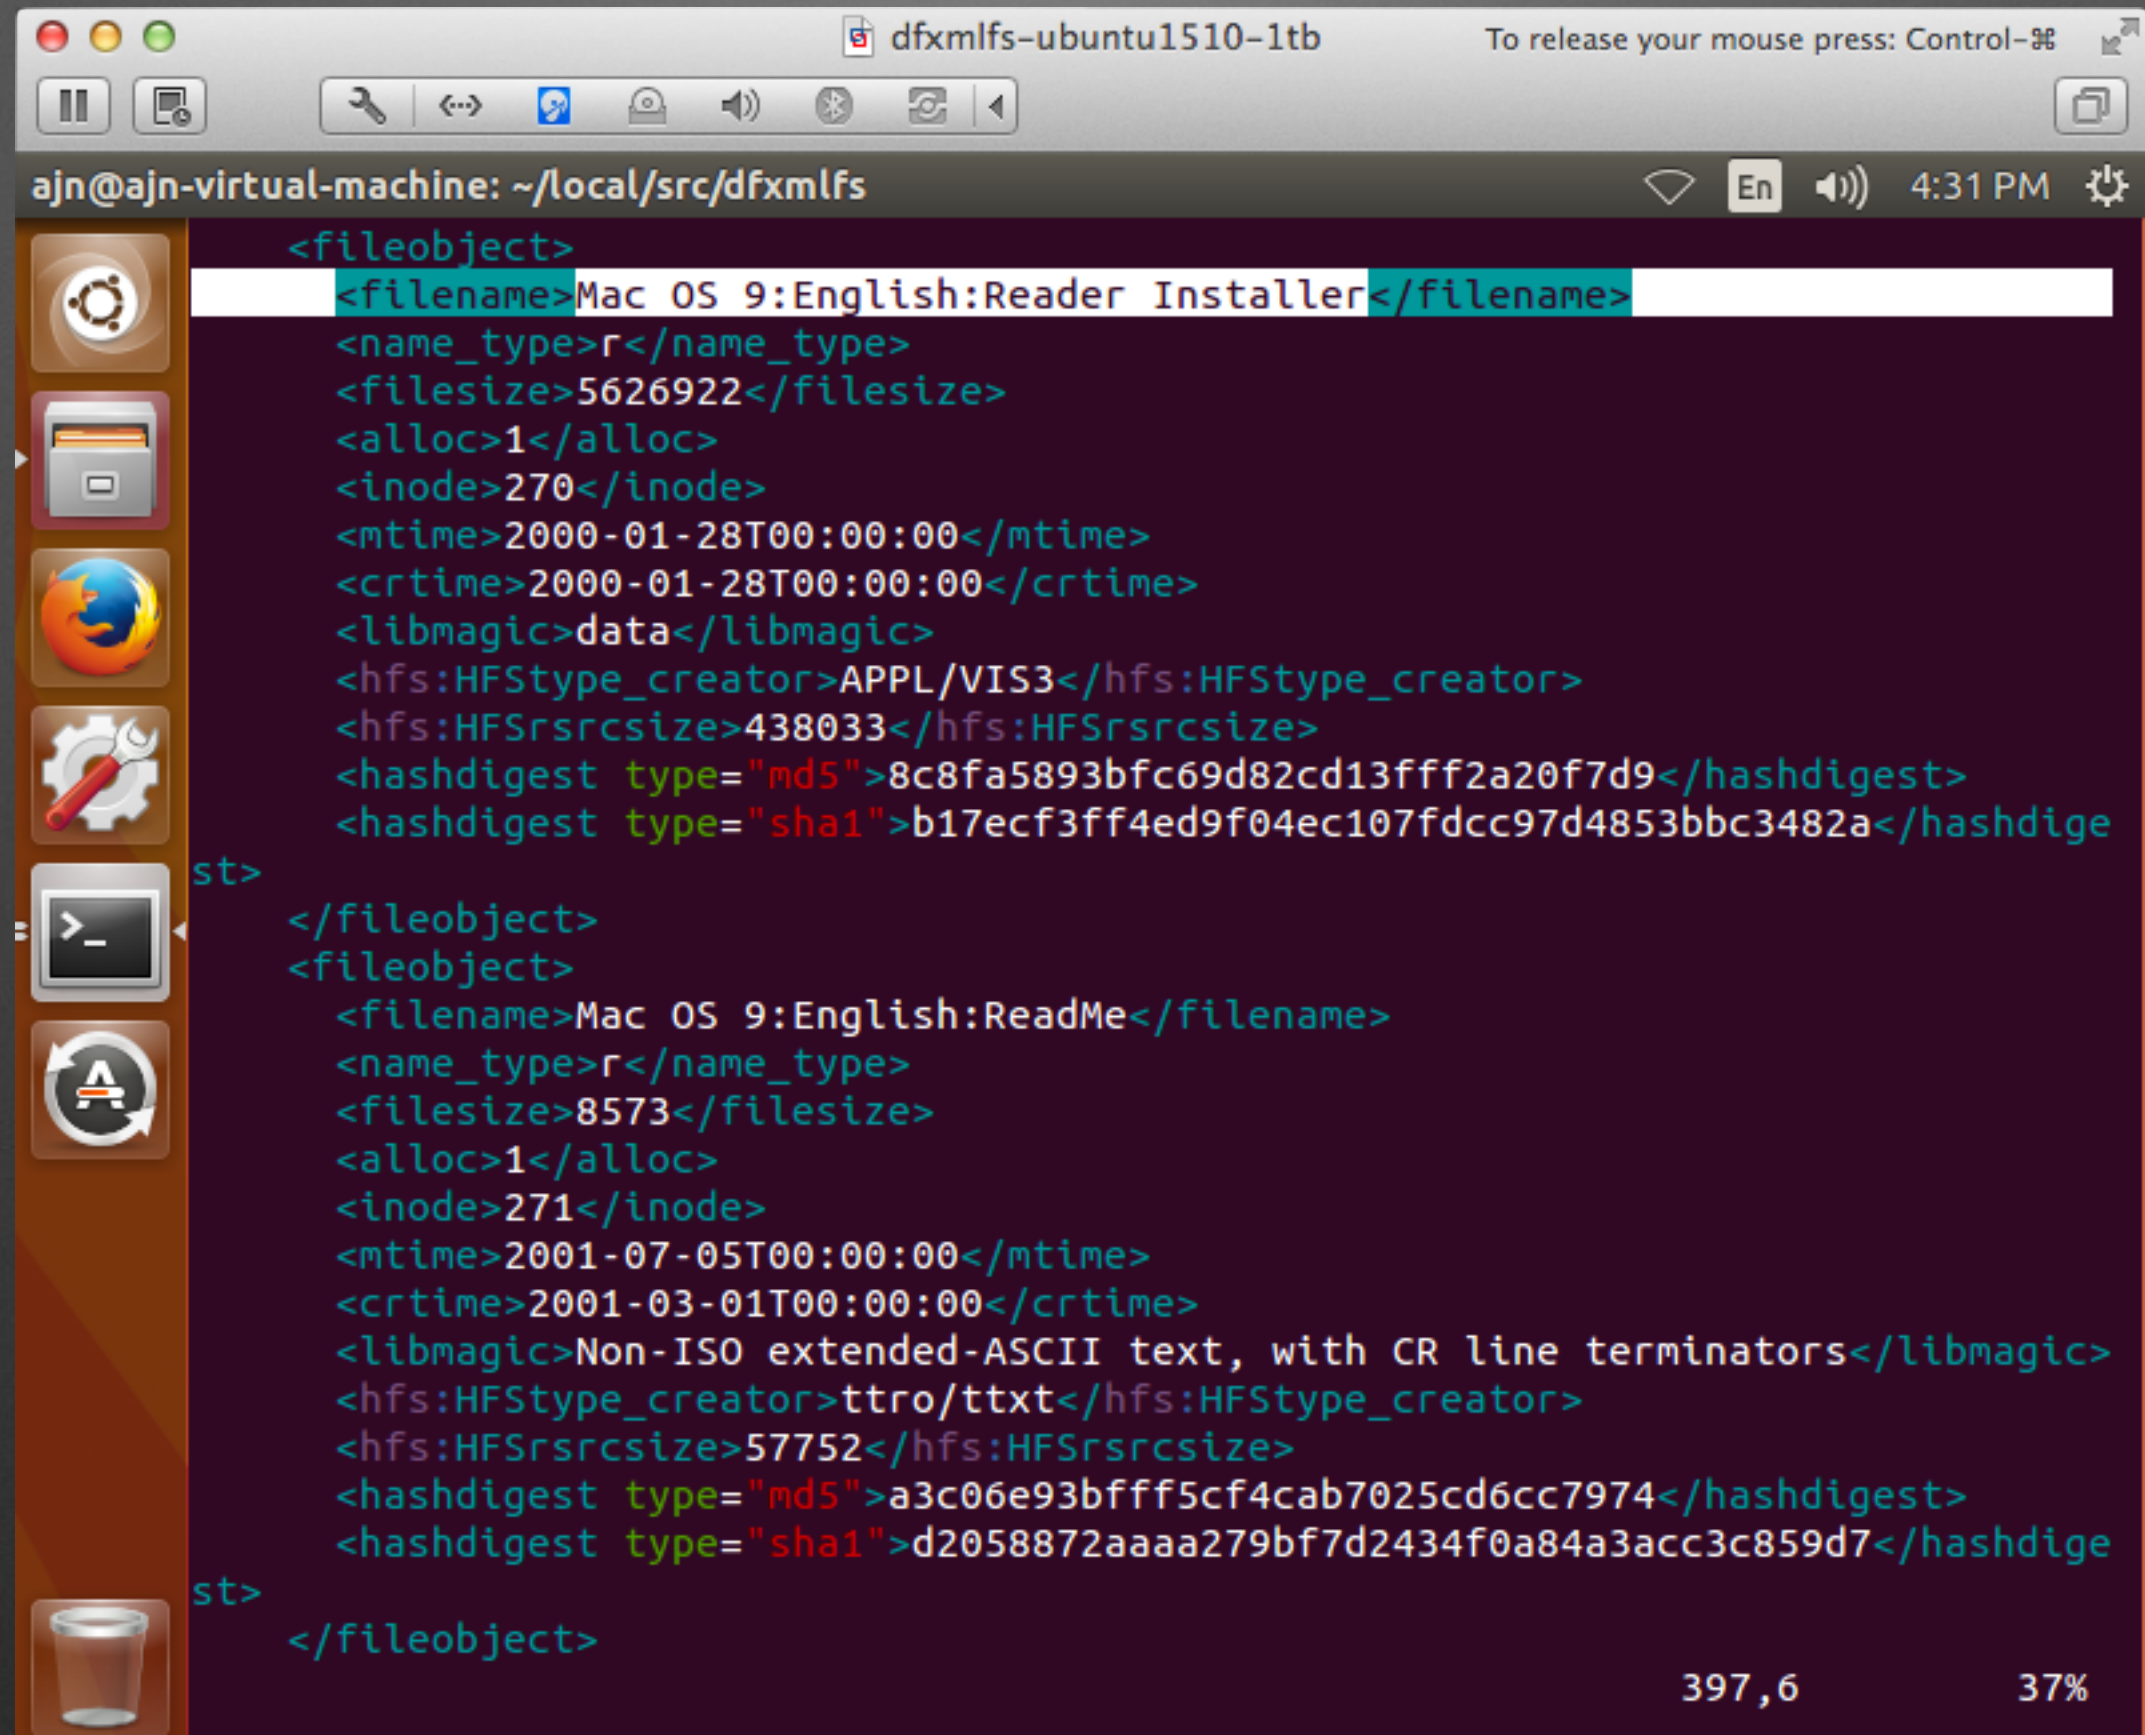
Benefits

- Non-standard disks can be read with normal interfaces, *without writing kernel code*.
 - (Debug prints become DFXML pretty easily.
Validation available with the DFXML schema.)
- Cantankerous parsing software?
 1. Make DFXML on special machine;
 2. Deploy file system on locked-down systems.
- Subsets of file systems can be presented.
 - Curated selections; deleted files; files since last backup;
"Hidden" files; Alternate Data Streams; etc...

Demo

- DFXMLFS implemented using DFXML Objects.py bindings, and Python-FUSE.
 - Tested in OS X, Ubuntu and Fedora.
- Showing two non-standard file systems:
 - HFS via hfs2dfxml
[Dietrich, 2015]
 - XTAF via modified SleuthKit
[Nelson *et al.*, 2014]

DFXML: An HFS disk




The image shows a terminal window titled 'dfxmlfs-ubuntu1510-1tb' with a status bar indicating 'To release your mouse press: Control-⌘'. The prompt is 'ajn@ajn-virtual-machine: ~/local/src/dfxmlfs'. The terminal displays XML data for two file objects. The first object is 'Mac OS 9:English:Reader Installer' with various metadata including size, inode, timestamps, and hashes. The second object is 'Mac OS 9:English:ReadMe' with similar metadata. The terminal also shows a file manager sidebar on the left with icons for applications and a trash can at the bottom.

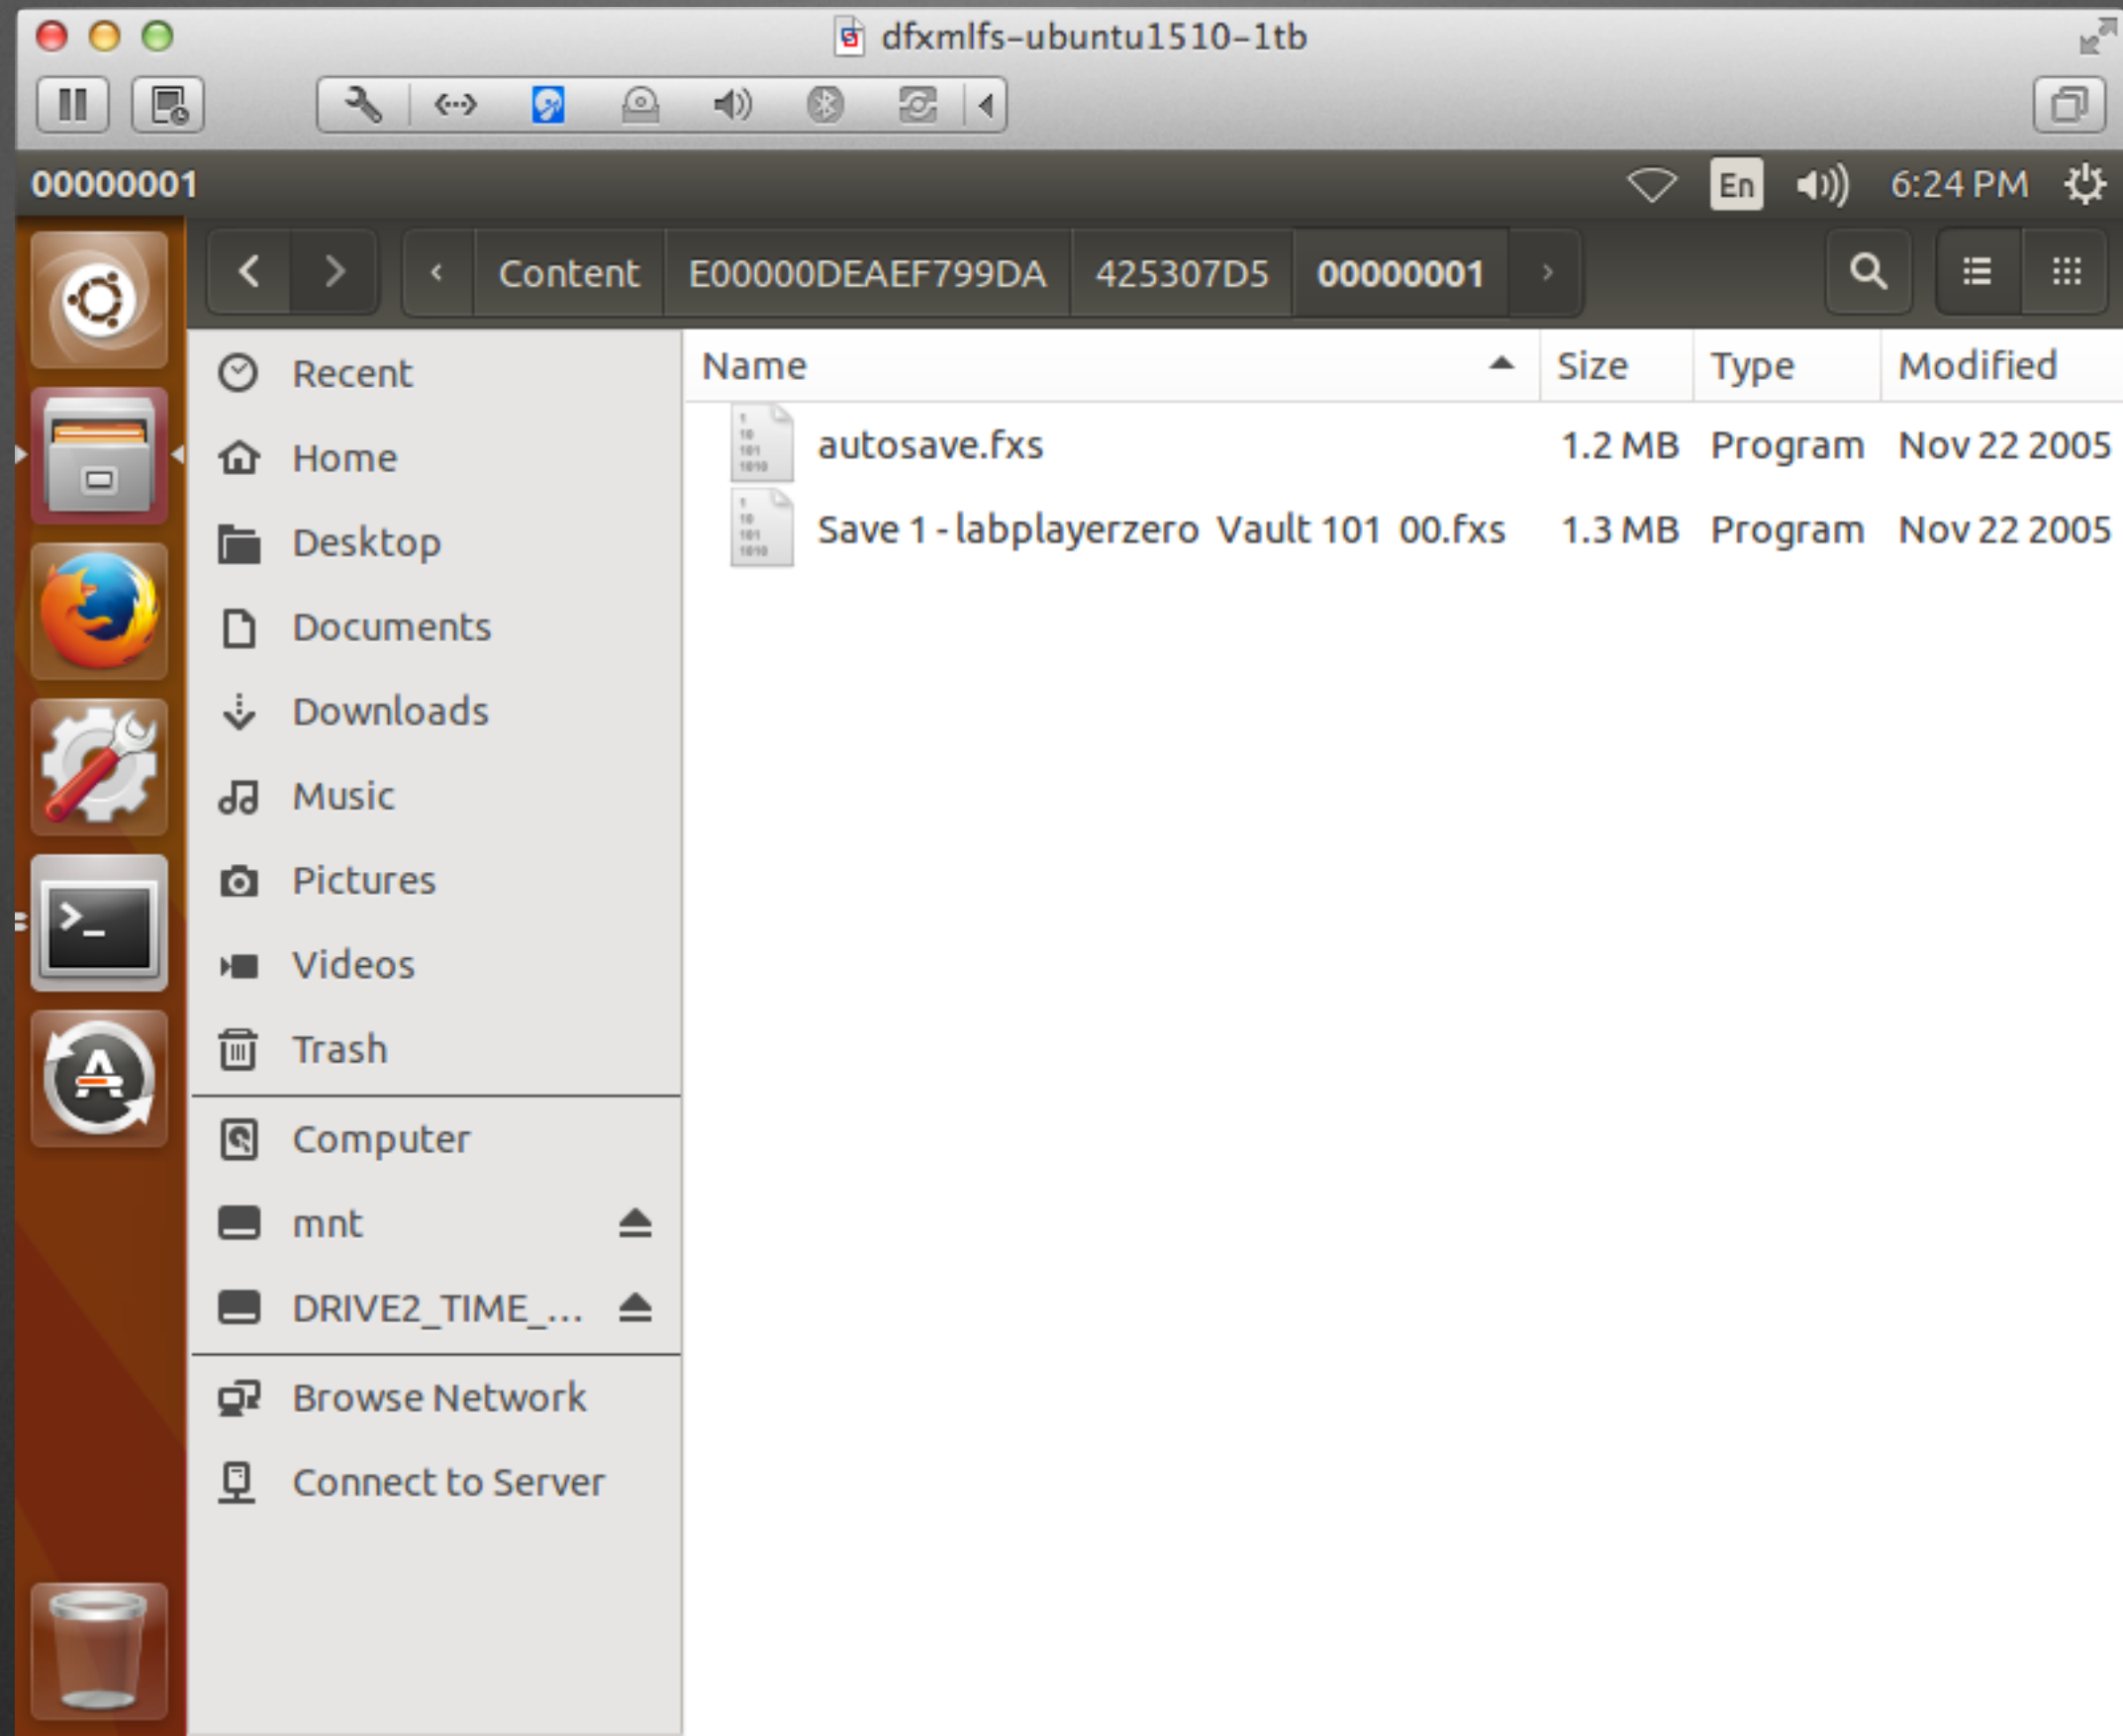
```
<fileobject>
  <filename>Mac OS 9:English:Reader Installer</filename>
  <name_type>r</name_type>
  <filesize>5626922</filesize>
  <alloc>1</alloc>
  <inode>270</inode>
  <mtime>2000-01-28T00:00:00</mtime>
  <ctime>2000-01-28T00:00:00</ctime>
  <libmagic>data</libmagic>
  <hfs:HFStype_creator>APPL/VIS3</hfs:HFStype_creator>
  <hfs:HFSrsrcsize>438033</hfs:HFSrsrcsize>
  <hashdigest type="md5">8c8fa5893bfc69d82cd13fff2a20f7d9</hashdigest>
  <hashdigest type="sha1">b17ecf3ff4ed9f04ec107fdcc97d4853bbc3482a</hashdige
st>
</fileobject>
<fileobject>
  <filename>Mac OS 9:English:ReadMe</filename>
  <name_type>r</name_type>
  <filesize>8573</filesize>
  <alloc>1</alloc>
  <inode>271</inode>
  <mtime>2001-07-05T00:00:00</mtime>
  <ctime>2001-03-01T00:00:00</ctime>
  <libmagic>Non-ISO extended-ASCII text, with CR line terminators</libmagic>
  <hfs:HFStype_creator>ttro/ttxt</hfs:HFStype_creator>
  <hfs:HFSrsrcsize>57752</hfs:HFSrsrcsize>
  <hashdigest type="md5">a3c06e93bfff5cf4cab7025cd6cc7974</hashdigest>
  <hashdigest type="sha1">d2058872aaaa279bf7d2434f0a84a3acc3c859d7</hashdige
st>
</fileobject>
```

397,6 37%

DFXMLFS: An HFS disk



DFXMLFS: An XBox 360 disk



Summary

- If an operating system won't read your disk, it can read DFXML of the disk instead.
- DFXMLFS available at:
<https://github.com/ajnelson-nist/dfxmlfs/>

Feedback is welcome.

Acknowledgements and references

- Acknowledgements:
 - Thanks to Dianne Dietrich for hfs2dfxml.
 - Thanks to Doug White and NSRL team for providing HFS disk images.
- Image credits:
 - Text file image by RRZEicons (Own work) [CC BY-SA 3.0 (<http://creativecommons.org/licenses/by-sa/3.0/>)], via Wikimedia Commons
<https://commons.wikimedia.org/wiki/File%3AText-txt.svg>
 - Floppy disk image by Jay Forsythe (Own work) [CC BY-SA 4.0 (<http://creativecommons.org/licenses/by-sa/4.0/>)], via Wikimedia Commons
https://commons.wikimedia.org/wiki/File%3AGreen_floppy_disk_graphic_SVG.svg
- Paper reference:
 - Nelson et al., DFRWS 2014:
<http://www.dfrws.org/2014/proceedings/DFRWS2014-6.pdf>

Questions?

alexander.nelson@nist.gov