SerNet

Apple Dances Samba

sambaXP 2014 - May 14, 2014

Ralph Böhme

SerNet GmbH, Göttingen - Berlin

Apple dances Samba

SerNet

What's the story ?

OS X Mavericks: new default protocol SMB2

Pre-OS X Mavericks: AFP

OS X happily connects to any SMB server including Samba.

And it works – until it doesn't:

Apple dances Samba ?

SerNet

Challenges and Pitfalls

- Alternate data streams
- Character encoding
- Interoperability with Netatalk/AFP
  - Access to data stored used with Netatalk
  - Locking across SMB and AFP protocols
- Spotlight

SerNet

Goals:

- Develop a new VFS module for a true cross-protocol SMB/AFP fileserver
- Add Apple Spotlight search support to Samba

Code, work in progress:

- https://github.com/slowfranklin/samba/commits/vfs_apple
- https://github.com/slowfranklin/samba/commits/spotlight

Alternate Data Streams (ADS)

SerNet

OS X and Alternate Data Streams

- In OS X every file can have optional data and metadata
  - Legacy Mac OS Finder metadata (fixed size)
  - An additional full fledged file fork called resource fork
  - Extended Attibutes (xattrs)
- OS X SMB client checks whether SMB server supports ADS:
  - No: client packs all extra data in an ._ prefixed AppleDouble file
  - Yes: client sends data as ADS
- OS X SMB server supports ADS

ADS used by OS X:

- Mac metadata stream: `AFP_AfpInfo`
  - Fixed size (60 bytes)
  - Contains Mac metadata like Type, Creator, creation date, attributes (visibility, locked, ...)
- Mac resource fork stream: `AFP_Resource`
  - Any size (like normal file)
- Extended attributes are also send as ADS

SerNet

Samba and ADS

- Samba without ADS and case sensivity: lots of possibly negative name lookups for ._ AppleDouble files which result in repeated directory rescans
- Opening folders with many files: performance goes south
- Samba supports ADS:
  - `vfs_streams_xattr`: size limitation (remember: resource fork), restricted xattr API (no partial reads/writes)
  - `vfs_streams_depot`: works, but incompatible with access from other network protocols

Challenges and Pitfalls • ADS

SerNet

New VFS module `vfs_apple` to the rescue:

- Supports ADS
- Store resource fork (`AFP_Resource`) in filesystem xattr in case the OS and filesystem support this (Solaris/ZFS), otherwise resort to storing it in an AppleDouble helper file
- Store Mac OS metadata (`AFP_AfpInfo`) in an xattr
- Avoids costly lookups for ._ AppleDouble files
- (Optionally) Compatible with Netatalk

Character Encoding

Challenges and Pitfalls • Character Encoding

SerNet

- NTFS illegal characters: / | * “ < > ?

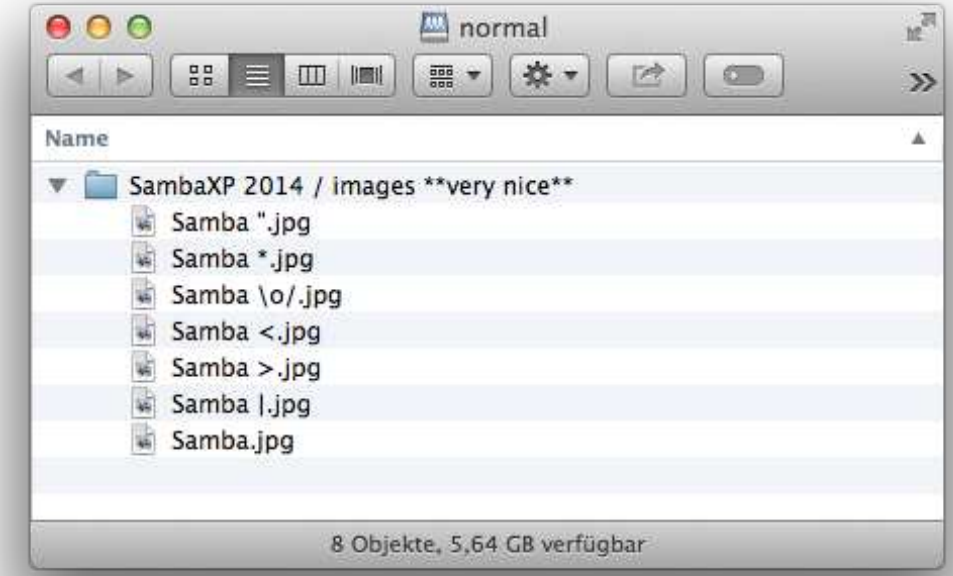

- OS X SMB Client maps them to Unicode private range [1]

SerNet

On the server:

```
$ ls -1 "/Volumes/normal/SambaXP 2014 images very nice
Samba jpg
Samba jpg
Samba jpg
Samba .jpg
Samba jpg
Samba.jpg
Samba jpg
```

, replacement glyph for characters in the Unicode private range.

SerNet

`vfs_apple` to the rescue:

```
$ ls -1 "/Volumes/vfs_apple/SambaXP 2014 : images **very nice**"
Samba <.jpg
Samba >.jpg
Samba |.jpg
Samba ".jpg
Samba *.jpg
Samba.jpg
Samba \o:.jpg
```

- Employs a builtin mapping for `vfs_catia`, characters in their native encoding
- Works in both ways, even Window clients are happy:

Challenges and Pitfalls • Character Encoding

SerNet

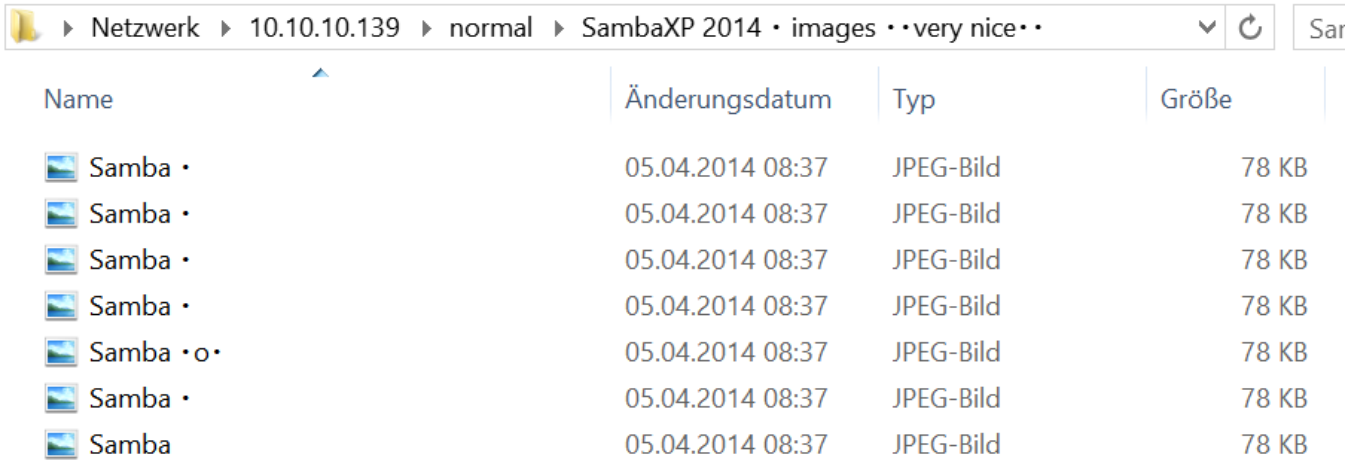

Windows 8 client displaying files with illegal NTFS characters

Windows uses a • as replacement glyph for characters in the Unicode private range, not the 

Interoperability

Challenges and Pitfalls • Interop with Netatalk

SerNet

Interoperability?

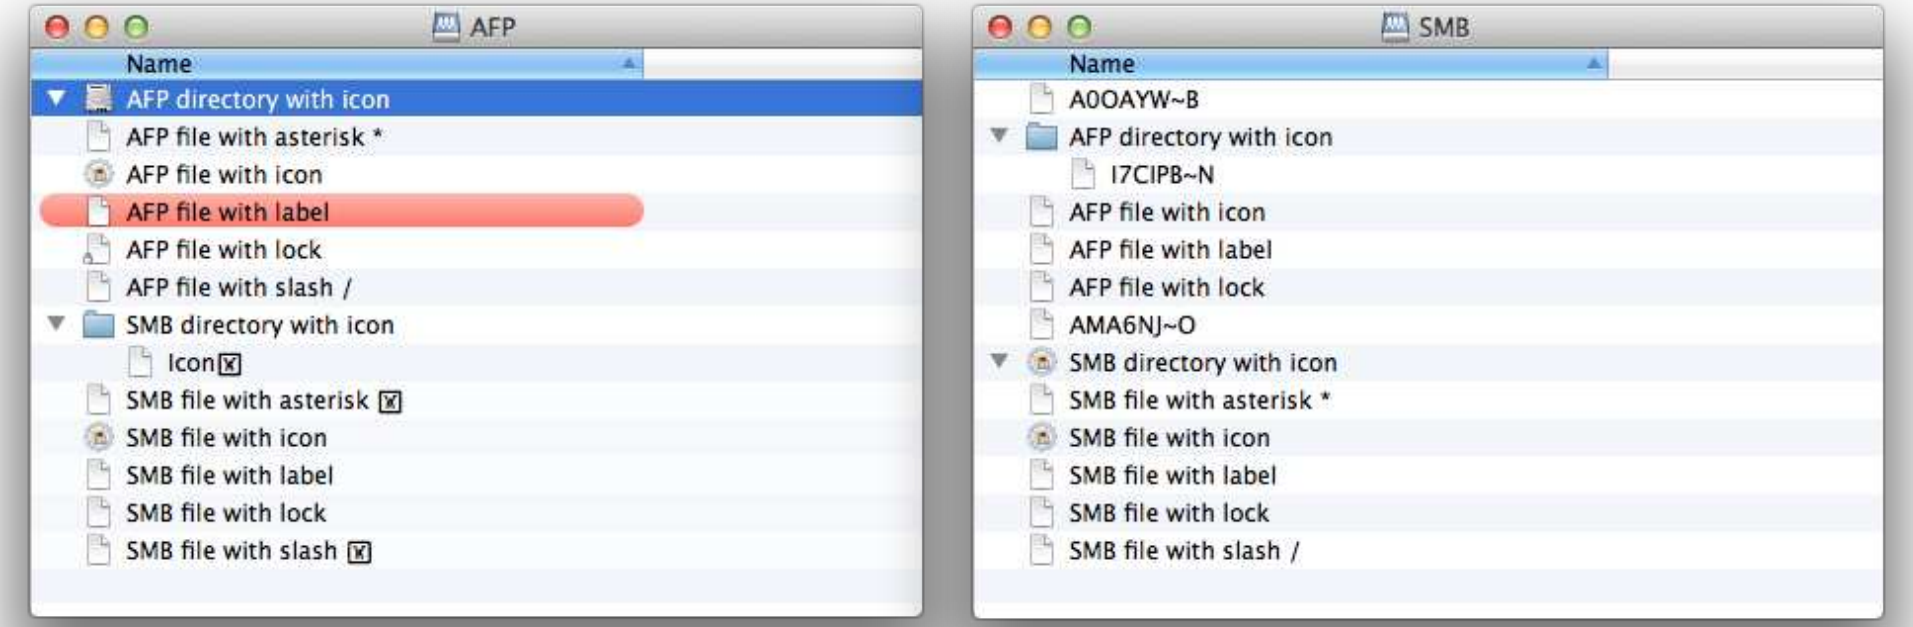

One element displays correctly, can you find it?

`vfs_apple`, while we're at it lets add some sweet topping:

- Map relevant named streams `AFP_AfpInfo` and `AFP_Resource` to corresponding Netatalk data structrures
- Cross-protocol locking: AFP file sharing modes similar to SMB
  - SMB open modes can be checked against AFP sharing modes with a simple fcntl() byte-range lock
  - This is not yet implemented

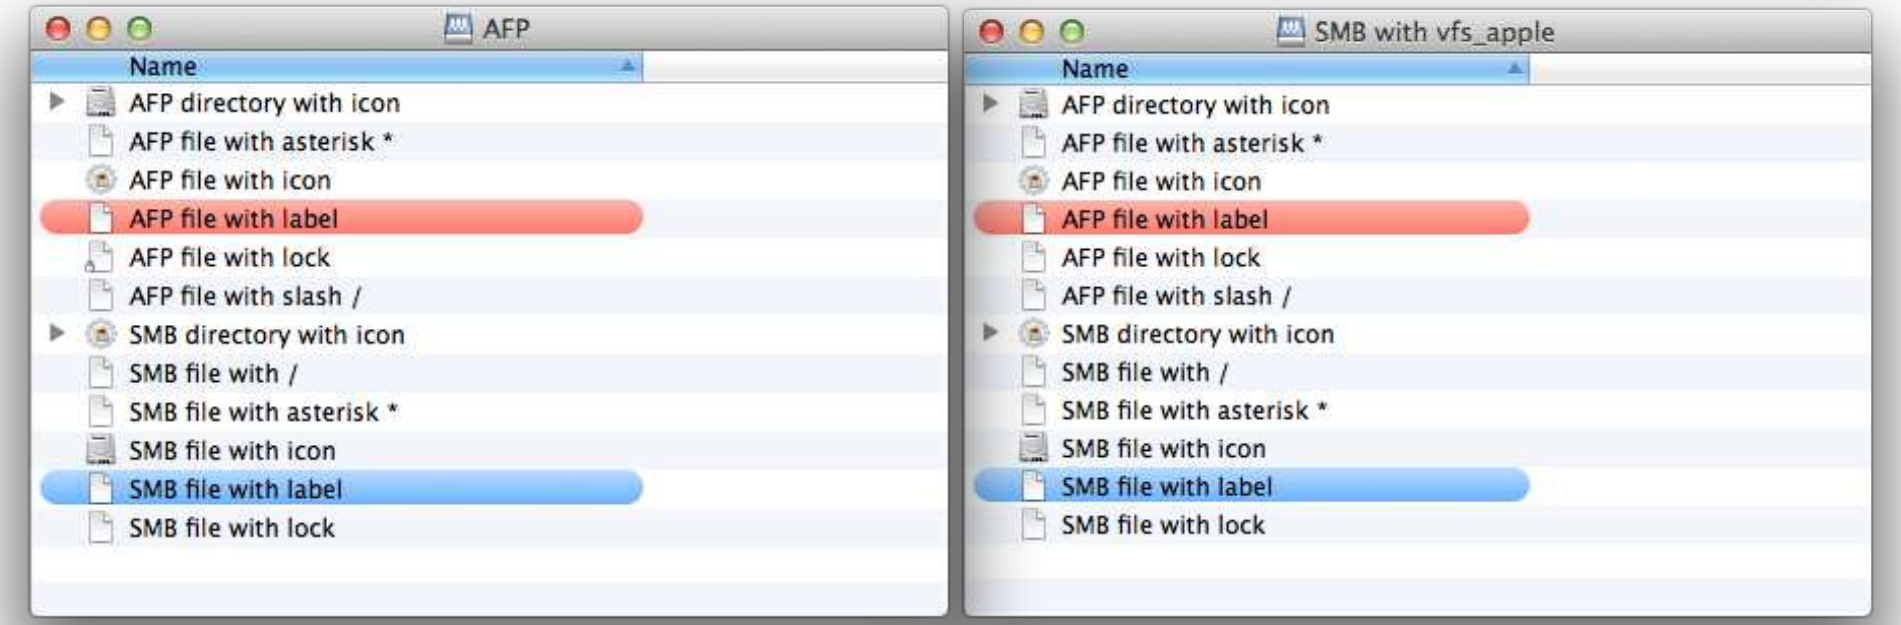

With vfs_apple

Challenges and Pitfalls • Interop with Netatalk

SerNet

Work in progress:

- Displaying directory icon via SMB somtimes not working
- Locked flag (SMB1: read only attribute, SMB2 ?)

Example smb.conf:

```
[SMB with vfs_apple]
    path = /Volumes/vfs_apple
    vfs objects = catia apple streams_xattr
    apple:ressource = file
    apple:metadata = netatalk
    apple:locking = netatalk
    apple:encoding = native
```


blog.tice.de

Challenges and Pitfalls • Spotlight

SerNet

- Apple technology for indexing and searching data
- Introduced 2005 in Mac OS X 10.4, desktop search only
- Since Mac OS X 10.5 search queries can be sent over the network to connected AFP fileservers
- Since Mac OS X 10.7 networked Spotlight supports SMB connections to Apple's own SMB server SMBX
- Added Spotlight support to Netatalk in 2012/2013
- Reverse engineered as Spotlight implementation details are not published by Apple

Challenges and Pitfalls • Spotlight

SerNet

On the wire:

- Over the wire format supports marshalling simple (integer, bool, float) and complex (string, set, key/value dictionary) data
- Marshalled data stream contains full type information
- Therfore doesn't require IDL
- Call it SDR (Spotlight Data Representation)

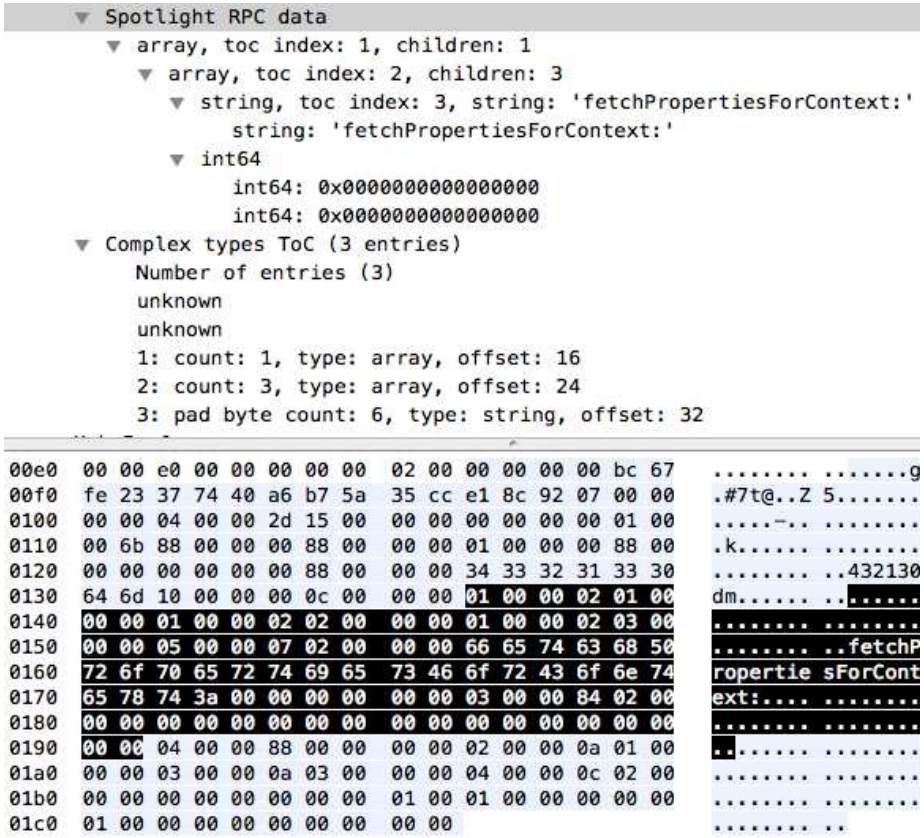

Challenges and Pitfalls • Spotlight

SerN
et

On the wire, SMB:

- SDR (remember: Spotlight Data Representation) wrapped inside DCE/RPC
- Uses `mdssvc` pipe
- Good news: Spotlight on the wire data mostly the same in AFP and SMB
- Added Wireshark dissector, already available, calls into AFP dissector for SDR

```
▶ SMB (Server Message Block Protocol)
▶ SMB Pipe Protocol
▶ Distributed Computing Environment / Remote Procedure Call (DCE/RP
▼ Spotlight metadata search service, cmd
    Operation: cmd (2)
    [Response in frame: 294]
  ▶ Policy Handle
    Unkn1: 0
    Device Id: 822083610
    Unkn3: 23
    Unkn4: 0
    Flags: 1795162113
  ▼ Request Blob
      Length: 136
      Size: 136
      Endianness: Little Endian
      ToC Offset: 88 Bytes, Query length: 120 Bytes
    ▼ Spotlight RPC data
      ▼ array, toc index: 1, children: 1
        ▼ array, toc index: 2, children: 3
          ▼ string, toc index: 3, string: 'fetchPropertiesForCont
              string: 'fetchPropertiesForContext:'
          ▼ int64
              int64: 0x0000000000000000
              int64: 0x0000000000000000
    ▼ Complex types ToC (3 entries)
        Number of entries (3)
        unknown
        unknown
        1: count: 1, type: array, offset: 16
        2: count: 3, type: array, offset: 24
        3: pad byte count: 6, type: string, offset: 32
```

Challenges and Pitfalls • Spotlight

SerNet

On the server:

- Receive SDR data stream
- Unpack query which is in string format. Examples:
  - `*=="main*"cdw||kMDItemTextContent=="main*"cdw`
  - `_kMDItemGroupId=="13"`
  - `kMDItemContentTypeTree=="public.tiff"cd`
- Use a scanner and a parser for translating this to SPARQL and throw it at Gnome Tracker

SPARQL? Tracker?

Gnome Tracker:

- „**Tracker** is a search engine, search tool and metadata storage system.“ [2]
- Crawls and monitors filessystem, stores metadata in a index database
- Robust, mature, capable, well designed, standard on Linux
- Provides a library for querying the store (sync and async)
- Query language is SPARQL

Challenges and Pitfalls • Spotlight

SerNet

SPARQL:

- „SPARQL (a recursive acronym for SPARQL Protocol and RDF Query Language) is an RDF query language, that is, a query language for databases, able to retrieve and manipulate data stored in Resource Description Framework format.“ [3]
- RDF? Nevermind!
- SPARQL supports all relevant logic expressionsin queries, therefor Spotlight search strings can be transformed one-to-one

SPARQL examples:

- ```
  SELECT DISTINCT ?url WHERE { ?obj nie:url ?url
  FILTER(regex(?url, '^file:///Volumes/spotlight/')) . ?
  obj fts:match 'main*'}
  ```

- ```
  SELECT DISTINCT ?url WHERE { ?obj nie:url ?url
  FILTER(regex(?url, '^file:///Volumes/spotlight/')) . ?
  obj nie:mimeType 'image/tiff'}
  ```

Demo Time

SerNet

Spotlight in action

Questions?

SerNet

Thank you!

Ralph Böhme, rb@sernet.de

SerNet GmbH

Bahnhofsallee 1b	**Schützenstr. 18**
37081 Göttingen	**10117 Berlin**
tel +49 551 370000-0	**+49 30 5 779 779 0**
fax +49 551 370000-9	**+49 30 5 779 779 9**

http://www.sernet.de

References

SerNet

[1] http://support.microsoft.com/kb/q117258

[2] https://wiki.gnome.org/Projects/Tracker

[3] http://en.wikipedia.org/wiki/SPARQL