

IBM Spectrum Scale: Support for macOS

—

Ralph Würthner
Software Engineer



Agenda

- **SMB – The preferred Protocol**
- **VFS Module “fruit“**
- **Configuration Model**
- **Known Issues**
- **Limitations**
- **Summary**



SMB – The preferred Protocol

- Starting from OS X 10.9 (Mavericks) Apple made SMB2 the default network file protocol
- Apple Filing Protocol (AFP) is deprecated
- Apple has its own proprietary SMB client and server implementation (independent from Microsoft or Samba)
- Apple specific extensions to SMB2 (a.k.a. SMB AAPL extensions) to increase browsing speed in Apple Finder application
 - Some SMB2 protocol fields carry different information
- macOS resource forks are stored in Alternate Data Streams
- POSIX file owner, group, and mode bits are mapped to S-1-55-88-* ACL entries (not used on Spectrum Scale)

Samba VFS Module “fruit”

- Most Apple specific Samba enhancements are placed in a VFS module
 - `vfs objects = shadow_copy2 syncops fruit streams_xattr gpfs fileid time_audit`
- Features
 - Handling of AAPL SMB2 Create Contexts
 - Provide additional information in SMB2_QUERY_DIRECTORY requests
 - Alternate Data Streams are stored in Extended Attributes
 - Existing Apple DoubleFiles are automatically converted into Alternate Data Streams
- “fruit” will not
 - Improve network throughput or latency when accessing files (other than avoiding access to Apple Double files)
 - Avoid strange access patterns seen with Apple Finder

SMB2 Create Contexts

- Additional flags or attributes included in SMB2_CREATE requests and responses
- In create requests used to specify how a SMB2_CREATE must be processed
- In create responses used to specify how a SMB2_CREATE was processed
- Flexible mechanism to enhance the SMB2 protocol

[-] ExtraInfo SMB2_CREATE_TIMEWARP_TOKEN SMB2_CREATE_DURABLE_HANDLE_REQUEST SMB2_C

```
[-] Chain Element: SMB2_CREATE_TIMEWARP_TOKEN "TWrP"
    Chain Offset: 0x00000020
    [+ Tag: TWrP
        Blob Offset: 0x00000018
        Blob Length: 8
    [-] Data: Timestamp
        Timestamp: Feb 12, 2020 17:39:26.000000000 CET
[+] Chain Element: SMB2_CREATE_DURABLE_HANDLE_REQUEST "DHnQ"
[+] Chain Element: SMB2_CREATE_QUERY_MAXIMAL_ACCESS_REQUEST "MxAc"
[+] Chain Element: SMB2_CREATE_QUERY_ON_DISK_ID "QFid"
```

SMB2 Create Context AAPL – Request from macOS Client

⊖ ExtraInfo SMB2_AAPL_CREATE_CONTEXT SMB2_CREATE_QUERY_MAXIMAL_ACCESS_REQUEST

⊖ Chain Element: SMB2_AAPL_CREATE_CONTEXT "AAPL"

Chain Offset: 0x00000030

⊖ Tag: AAPL

Blob Offset: 0x00000010

Blob Length: 4

Blob Offset: 0x00000018

Blob Length: 24

⊖ Data: AAPL Create Context request

⊖ AAPL Create Context request

Command code: Server query (1)

Reserved: 0x00000000

⊖ Query bitmask: 0x0000000000000007, Server capabilities, Volume capabilities, Model information

..... .1 = Server capabilities: True

..... .1. = Volume capabilities: True

..... .1.. = Model information: True

⊖ Client/Server capabilities: 0x000000000000000f, Supports REaddirAttr, Supports macOS copyfile, UNIX-based, Supports NFS ACE

..... .1 = Supports REaddirAttr: True

..... .1. = Supports macOS copyfile: True

..... .1.. = UNIX-based: True

..... 1... = Supports NFS ACE: True

⊖ Chain Element: SMB2_CREATE_QUERY_MAXIMAL_ACCESS_REQUEST "MxAc"

SMB2 Create Context AAPL – Response from Samba

[-] ExtraInfo SMB2_AAPL_CREATE_CONTEXT SMB2_CREATE_QUERY_MAXIMAL_ACCESS_REQUEST

[-] Chain Element: SMB2_AAPL_CREATE_CONTEXT "AAPL"

Chain Offset: 0x00000050

[-] Tag: AAPL

Blob Offset: 0x00000010

Blob Length: 4

Blob Offset: 0x00000018

Blob Length: 56

[-] Data: AAPL Create Context response

[-] AAPL Create Context response

Command code: Server query (1)

Reserved: 0x00000000

[-] Query bitmask: 0x0000000000000007, Server capabilities, Volume capabilities, Model information

.....1 = Server capabilities: True

.....1. = Volume capabilities: True

.....1.. = Model information: True

[-] Client/Server capabilities: 0x0000000000000007, Supports REaddirAttr, Supports macOS copyfile, UNIX-based

.....1 = Supports REaddirAttr: True

.....1. = Supports macOS copyfile: True

.....1.. = UNIX-based: True

.....0... = Supports NFS ACE: False

[-] Volume capabilities: 0x0000000000000000

.....0 = Supports Resolve ID: False

.....0. = Case sensitive: False

.....0.. = Supports full sync: False

Model string: MacSamba

[+] Chain Element: SMB2_CREATE_QUERY_MAXIMAL_ACCESS_REQUEST "MxAc"

ReadDirAttr Extension

- Return macOS specific information in FileIdBothDirInformation information level
 - Finder Info (→ ShortName)
 - Resource Fork Length (→ ShortName)
 - Access rights, Unix mode (→ EaSize, ReservedField)
- Avoid additional network round trips to query information required by Finder Application
 - SMB2_CREATE → SMB2_GETINFO → SMB2_CLOSE

Configuration Model in Spectrum Scale

- macOS support can be **only enabled** and is enabled for all SMB shares within cluster!
- Enabling requires SMB service to be stopped
- How to enable:

```
mmsmb config change --vfs-fruit-enable
```

- Disabling macOS support requires involving IBM support
 - All GPFS file systems must be scanned and resource forks stored in Alternate Data Streams must be converted back to AppleDouble files
 - May result in a lengthy cluster outage (depends on file system size, number of files to be converted)

Limitations with macOS

- Alternate Data Streams are limited to 50KB per file / 16KB per single stream
 - Size limitation for Extended Attributes
- Alternate Data Streams support is advertised to all SMB clients (not only macOS clients)
 - Keep eye on 50KB/16KB limitation
- No ACL Management is possible
- Spotlight Search is not supported
 - Would require additional services running on Spectrum Scale cluster nodes
- Time Machine is not supported
 - Time Machine requires durable file handles

Known Issues: macOS FileIDs / CNID

- FileIDs are a HFS+ heritage where files are identified via a unique, monotonically-increasing number (Catalog Node ID)
- FileIDs are returned on request to the macOS client
- Returning FileID = 0 or setting “File_ids_off=yes” in `/etc/nsmb.conf` made macOS to ignore the returned FileIDs
- At least since macOS 10.14 behavior changed and macOS requires correct FileIDs otherwise strange access errors can show up
- The upcoming Spectrum Scale release will generate semantically correct FileIDs
→ remove “File_ids_off=yes” setting in `/etc/nsmb.conf`

Known Issues

- No native support for Volume Snapshot Services (VSS)
 - Manual access to GPFS Snapshots used for VSS (`.snapshots/@GMT-<timestamp>`) is not working
 - Limitation within SMB2 protocol implementation of Samba
 - Lookup of `@GMT-<timestamp>` directory entries is not possible
 - Access to other snapshots is possible with upcoming Spectrum Scale release (non `@GMT-<timestamp>` entries)
- When HSM is used offline bit is ignored
 - Apple Finder can cause recall storms when accessing directories containing migrated files

Summary

- Spectrum Scale does now support SMB AAPL extensions
- Major performance improvements when working with Apple Finder & large directories
- Resource Forks are stored in Alternate Data Streams / limited in size

Thank you!

Provide Feedback ×



Tell IBM What You Think

Let us know what you think about IBM Spectrum Scale. It takes only a couple of minutes for you to help us improve our service. [IBM Privacy Policy](#)

Please help us to improve Spectrum Scale with your feedback

- If you get a survey in email or a popup from the GUI, please respond
- We read every single reply