

Samba VFS to libgfapi



Driving the Backroads

Christopher R. Hertel
Samba Team Member



Introductions





Introductions

Me

- ▶ Author
- ▶ Storage Architect
- ▶ Network Engineer
- ▶ SMB Know-It-All
- ▶ Samba Team Member
(since '98)
- ▶ Incurable Idealist



A ruminant mammal (Geekus geekus) with long legs, humped shoulders, and broadly palmated antlers.



Introductions



redhat®

The opinions expressed are my own
and not necessarily those of my
employer, my spouse, my childrenz,
the dog, or "the Voices".





Introductions

You





Introductions

Them



Members of the Samba Team gather at the 10th annual Samba eXPerience conference in Göttingen, Germany.

Introductions


A Note About Me

For good or ill, I have become...

✍ **The** source for SMBv1 documentation,

✍ Purveyor of protocol pedanticism, and

✍ Chronicler of the apopsicle. 🍦



Introductions

Lead author of two SMBv1 protocol specifications, published by Microsoft:



[MS-CIFS]

- 💡 Covers SMB/CIFS as implemented in Windows NT and 98.

[MS-SMB] (revised)

- 💡 Rewritten to reference [MS-CIFS].
- 💡 Details changes from NT4 to W2K through W2K12 & Win 8.



Introductions

Let me get this straight...



Microsoft contracted with



a *Samba Team* member



to produce
publicly available



SMB protocol
specifications!?



Wait... What?



Introductions

Where are we going?

- 🌍 Overview of SMB and Samba
- 🌍 Samba and Gluster and FUSE, oh my
- 🌍 The Samba VFS layer
- 🌍 Building Samba VFS modules
- 🌍 Missing Pieces
- 🌍 Questions

(and answers, if I have any)





Introductions

Where we are not going:

- We won't go into `libgfapi` internals
- I rely on others for `libgfapi` know-how







What's in a Name?

A Rose by Any Other Name
Would
Wither and Die

-- Alan Swann (Peter O'Toole),
My Favorite Year





What's in a Name?

- SMB: Server Message Block**
Originally, a simple network file protocol created by IBM. It grew, over time, to include 75 primary commands and an unbounded number of secondary commands.
- CIFS: Common Internet File System**
A 1990's marketing upgrade. CIFS is just another name for SMB1. Microsoft officially dropped the CIFS name in 2012.
- SMB1:** A new name for SMB, as implemented in Windows NT and still used in Windows today.
- SMB2:** A new protocol intended to eventually replace SMB1. SMB2.0 was included in Windows Vista. SMB2.1 in W2K8r2 and Windows7.
- SMB3:** Just another dialect of SMB2 (specifically SMB2.2). SMB3 adds support for cluster and multipath features.





What's in a Name?

SMB = **S**erver **M**essage **B**lock
DOS, OS/2, and Windows

CIFS = SMB1

SMB1 = SMB on Windows,
starting with NT

SMB2 = New SMB,
introduced in Vista

SMB2.0 = Vista

SMB2.1 = Win7 / W2K8r2

SMB3 = SMB2.2 = Win8 / W2012



The naming is confuzzeling. Keep these charts handy.



What's in a Name?

Samba and CTDB

Samba = SaMBa

An Open Source SMB server,

➔ Full support of SMB1 (and older)

➔ Partial support of SMB2/3


Samba v4 is the current production branch.

CTDB = Cluster Trivial Database

A distributed key⇒value database,


➔ Maintains Samba state across a cluster

➔ Also provides IP failover



What's in a Name?

Samba
is...




- 🔌 An SMB protocol engine, and...
- 🔌 A semantic translation machine:
 - Different locking
 - Different access controls
 - Different timestamps
 - Different just about everythingSamba has to translate from Windows to POSIX and back again.



What's in a Name?

CTDB is...

A Distributed Database



- Tables are replicated on every node
- Persistent tables are updated atomically
- Volatile tables use a token-passing scheme; the “owner” has current data:
 - data will be lost when a node crashes
 - ...and that's okay!

In theory,
CTDB's functionality could be replaced...




What's in a Name?

State of the Art:

Samba and SMB2/3

- ✦ Samba has full support for SMB1
- ✦ Samba can negotiate SMB2
 - 👾 SMB2 features are limited
 - 👾 SMB3 cluster features are missing





What's in a Name?

State of the Art: SMB3 cluster features

 Scale-out

 Transparent Failover

 Client/server negotiated

 CTDB does not yet support failover

CTDB (or its replacement) must support
SMB3 Persistent Handles





What's in a Name?

Summary

What we need from Samba:

- 🌳 SMB3 cluster features

What we need from CTDB:

- 🌳 Persistent handle support

What we need from Gluster:

- 🌳 Internal support for Windows semantics

- 🌿 ACLs, Leases, etc.


- 🌿 Persistent Handles?







Samba and Gluster and FUSE





Samba/Gluster/FUSE

“Samba is a semantic translation machine”

—We said that already

- 🚲 Samba translates from Windows to POSIX and back again
- 🚲 Samba expects POSIX behaviors:
 - 🔵 Cache Coherency
 - 🔵 POSIX byte-range locking support
- 🚲 Samba also wants Linux FS features:
 - 🔵 Extended Attributes
 - 🔵 “POSIX” ACLs
 - 🔵 RichACLs (would be nice)





Samba/Gluster/FUSE

Gluster

...is adaptable, and cool.

- ▶ Can add support for SMB-specific features:
 - Windows ACL support
 - OpLocks and Leases
 - Case-insensitive names
- ▶ Can selectively enforce semantics:
 - Share modes
 - Preallocation
 - Other stuff



There are possibilities to explore here.



Samba/Gluster/FUSE

FUSE


...presents a problem.

- 🔑 A generic mount point
- 🔑 No way to enable/disable per-access method features

Fuse provides a single standard interface, but effectively locks away Gluster internals.

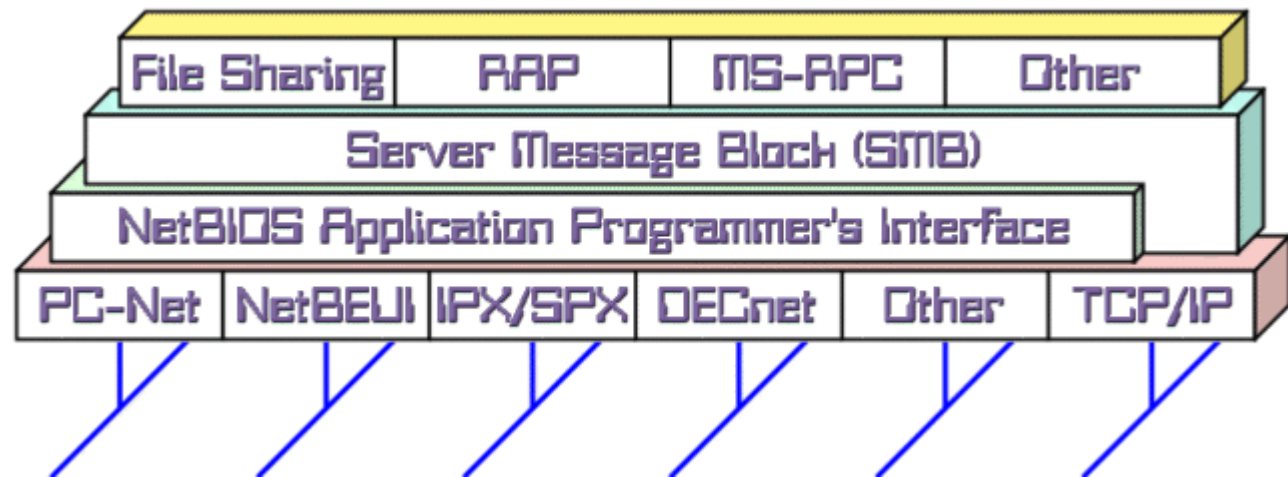
What's a coder to do?







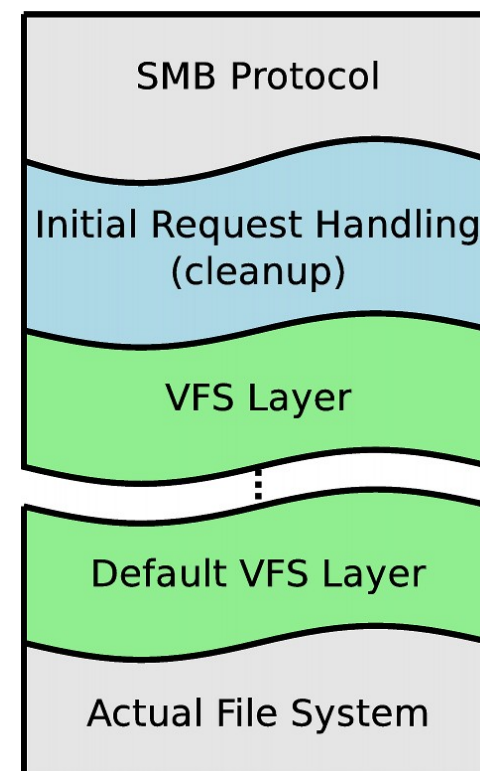
The Samba VFS layer



The Samba VFS Layer


Samba is Built in Layers (conceptually)

- ④ SMB messages received and parsed
- ④ (Non-FS commands are handled elsewhere)
- ④ Call the VFS layer
- ④ The final VFS module talks to the File System
 - ☀ Upper layers may bypass lower layers




The Samba VFS Layer

If there's no real File System,
bypass the lower VFS layers



All VFS calls must be implemented
(possibly returning ENOTSUP)
to avoid errors.





Building a Samba VFS









The Samba VFS Layer

There is a guide:

[Writing a Samba VFS](#) by Richard Sharpe

Unfortunately...

-  The Samba VFS Interface changes as new features are added
-  VFS modules must be kept in sync with Samba releases
-  VFS developers need to coordinate with Samba developers





Example code and such was s'posed to go here.

I'll work out a good place to put the code and we'll start a project...





The End

