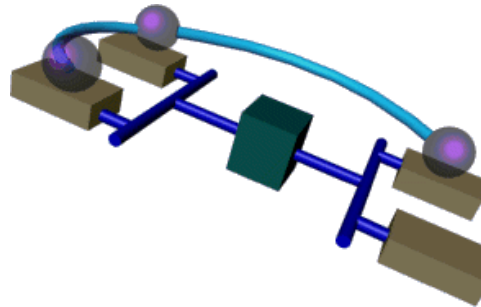


SMR 2.2

and

Real-World Data Networks



Christopher R. Hertel

Storage Architect, CIFS Geek

Founder and CTO

DTC Leading Edge Seminar

December 7, 2011



Introductions





Introductions

YOU

- ▶ Network Storage Administrators
- ▶ Network Storage Developers
- ▶ Students
- ▶ The Curious
- ▶ The Others
(*you know who you are*)



Introductions

ME: Your Friendly Neighborhood CIFS Geek

- ▶ Samba Team member (since 1998-ish)
- ▶ jCIFS Project co-founder
- ▶ CIFS Author (shameless plug)
- ▶ Network Storage Geek
- ▶ Incurable Idealist
- ▶ Etc., etc., ad nauseum



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



Introductions

SMB/CIFS and SMB2



- *The* Microsoft network file protocol
- Originally created by IBM for PC-DOS
- Ported and updated for OS/2, then W/NT
- SMB2 was introduced with Windows Vista
 - V2.1 with Windows 7, W2K8r2 Server
 - V2.2 with Windows 8 (next year)

A de facto (vs. de jure) standard.



Introductions

Terminology (real world)



SMB: Server Message Block protocol

A stateful network file system protocol originally created by IBM in the early 1980s for use with the PC-DOS operating system.

CIFS: Common Internet File System

A “marketing upgrade” to SMB. This new name for SMB was coined in the mid 1990's. The term “CIFS” is now often used as a name for the complete suite of protocols that include and provide support for SMB. Often written “SMB/CIFS”.

SMB2: Server Message Block protocol version 2

A complete rewrite of the SMB protocol, introduced with Windows Vista. SMB2 reduces the top-level command set from 75 commands to 19.



Introductions

Terminology (legal and regulatory world)

CIFS: The Server Message Block file sharing protocol as implemented in Windows NT 3.51, NT 4, and Windows 9x clients.

SMB: The Server Message Block file sharing protocol as implemented in Windows starting with Windows 2000, up to and including current versions of Windows.

SMB2: The Server Message Block protocol, v2 as defined on the previous slide.





The terminology changes depending upon who you talk to, when you talk with them, and the context of the conversation.

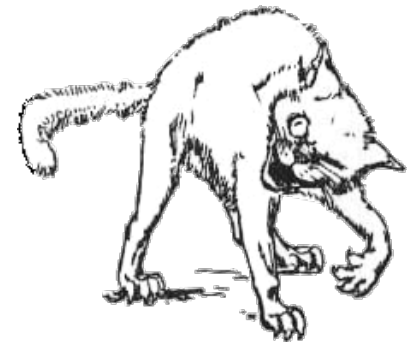




Introductions

The Competition: NFS

-  The POSIX/Unix network file protocol
-  Originally created by Sun
-  Given to the IETF for standardization
-  NFSv4.x specified over the past 10 years
 - pNFS == Parallel I/O (objects)
 - NFS over RDMA



A de jure (vs. de facto) standard.



Introductions

Samba Team:

World-Renowned SMB/CIFS and SMB2 Developers

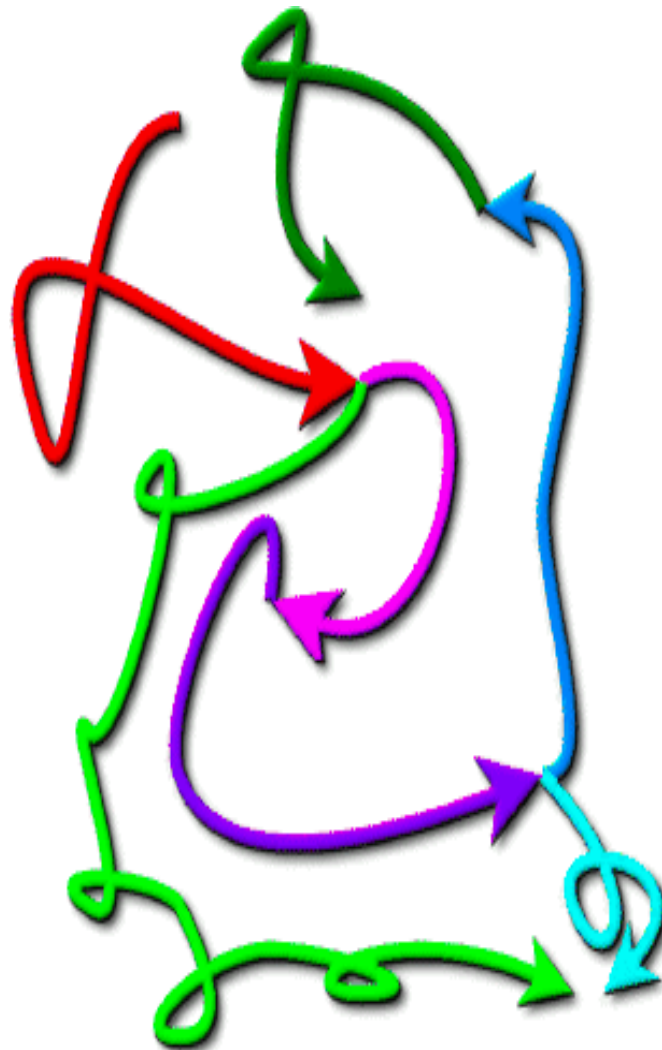


Members of the Samba Team gather at the 10th annual **Samba eXPerience** conference in Göttingen, Germany.



Introductions

Whither Shall We Wander?



- Why do we care?
- Breaking into SMB
- SMB2.2 Features
- SMB vs. NFS
- SMB Internals
- “Real World” SMB



Why do we Care About SMB?





Why do we care about SMB?

There is little to love about SMB

- It is proprietary
Protocol development is under Microsoft's control.
- It is somewhat closed
The Windows NT version (still used today) was not documented until late 2009.
- It is bulky and inefficient
SMB/CIFS supports DOS and OS/2 system calls. 75 primary commands plus pass-through RPC and system calls.



SMB is “Windows on the Wire”.



Why do we care about SMB?

SMB2 is only a slightly better

— It is still proprietary

Protocol development is under Microsoft's control.

— It is still somewhat closed

Specifications are available, including pre-release SMBv2.2 documentation, but these are not “standards”.

— It is leaner and cleaner

No DOS or OS/2 support.

SMB2 has only 19 primary commands.




SMB2 is still “Windows on the Wire”.



Why do we care about SMB?

Don't Laugh Yet...


 SMB has been built into every version of Windows since WfWG3.11

 SMB2 has been included since Vista

- Microsoft is putting lots of effort into SMB2

 All major NAS vendors support SMB

- 3rd party SMB2 adoption is accelerating

 NFSv4 “borrows” features from SMB




There are no hard numbers... but the SMB/CIFS community guestimates 90+% market share.



Why do we care about SMB?

In the Market: It's NFS vs. SMB/SMB2

- ✧ Other proprietary NAS protocols are gone
 - Apple, Novell, etc. now support NFS and SMB
 - ✧ There's a new crop of specialized protocols
 - Object Storage is big, but diverse
 - Clustered storage is another big (crashing?) wave
 - FUSE makes it easy to create new file systems
 - ✧ Specialized protocols are typically platform-specific (though there are exceptions)
 - ✧ Windows 8 will include SMB2.2 client/server
- 



Breaking Into SMB





Breaking Into SMB

At the start, *SMB was documented:*

- 🕒 **1984:** IBM Personal Computer Seminar Proceedings, Volume 2, Number 8
- 🕒 **1986:** OpenNET/Microsoft Networks FILE SHARING PROTOCOL EXTENSIONS, Version 1.9, Microsoft and Intel (XENIX extensions)
- 🕒 **1988:** Microsoft Networks/OpenNet, Document Version 2, Microsoft and Intel (Core)
- 🕒 **1988:** Microsoft Networks SMB File Sharing Protocol Extensions Version 2.0, Document Version 3.3, Microsoft Corporation (LAN Manager 1.0)
- 🕒 **1989:** Microsoft Networks SMB File Sharing Protocol Extensions Version 3.0, Document Version 1.09, Microsoft Corporation (LAN Manager 1.2)
- 🕒 **1990:** Microsoft Networks SMB File Sharing Protocol Extensions Version 3.0, Document Version 1.11, Microsoft Corporation (LAN Manager 2.0)
- 🕒 **1992:** Microsoft Networks SMB File Sharing Protocol Extensions, Document Version 3.4, Microsoft Corporation (LAN Manager 2.1)





Breaking Into SMB

Then things started thinning out.

- **1992:** X/Open CAE Specification, Protocols for X/Open PC Interworking: SMB, Version 2, X/Open Company, Ltd. (Core through LAN Manager 2.0)
- **1996:** Microsoft Networks SMB File Sharing Protocol, Document Version 6.0p, Microsoft (Unfinished draft of NT LAN Manager 0.12 documentation.)
- **1997:** A Common Internet File System (CIFS/1.0) Protocol, IETF INTERNET-DRAFT, Paul J. Leach, Dilip C. Naik (Unfinished draft v2 of NT LAN Manager 0.12 specification.)
- **2002:** Common Internet File System (CIFS) Technical Reference, Revision: 1.0, Storage Networking Industry Association (SNIA)
- **2003:** Implementing CIFS, yours truly, Prentice Hall PTR





Breaking Into SMB

During this time...



[MS]PC-DOS



OS/2



Windows NT



Windows 2000



Windows XP



Windows 2003



Windows Vista

...and we already knew that
the documentation we had
was, in places,



Incorrect



Incomplete



Incomprehensible



Never ascribe to malice that which is adequately explained
by incompetence. — attributed to Napoleon Bonaparte, among others



Breaking Into SMB

This situation made people unhappy.





Breaking Into SMB

Open Source Credentials Notwithstanding...

Microsoft asked a member of the
Samba Team to document SMB/CIFS!



Breaking Into SMB

Thus, ~~SMB~~/CIFS is covered in two documents:

[MS-CIFS]

- ✿ Provides the base specification of the “NT LM 0.12” dialect.
- ✿ A “snapshot in time”.
- ✿ Most of this stuff is still there in current Windows versions. Really.

[MS-SMB]

- ✿ “Extends” [MS-CIFS].
- ✿ Documents changes made to SMB starting in W2K.
- ✿ Still the same “NT LM 0.12” dialect.



Note: The naming is backwards!



Breaking Into SMB

Go here:

<http://www.microsoft.com/openspecifications/>

Over 400 documents have been published, covering:



Authentication



Windows Internals



File Formats



Client-Server Protocols



Server-Server Protocols



Overview docs provide starting points for understanding groups of docs.

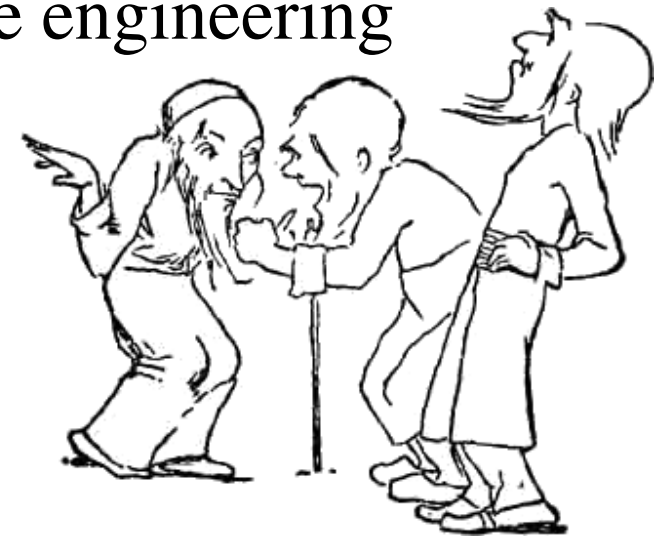


Breaking Into SMB

“We should implement them all.”

— Tridge

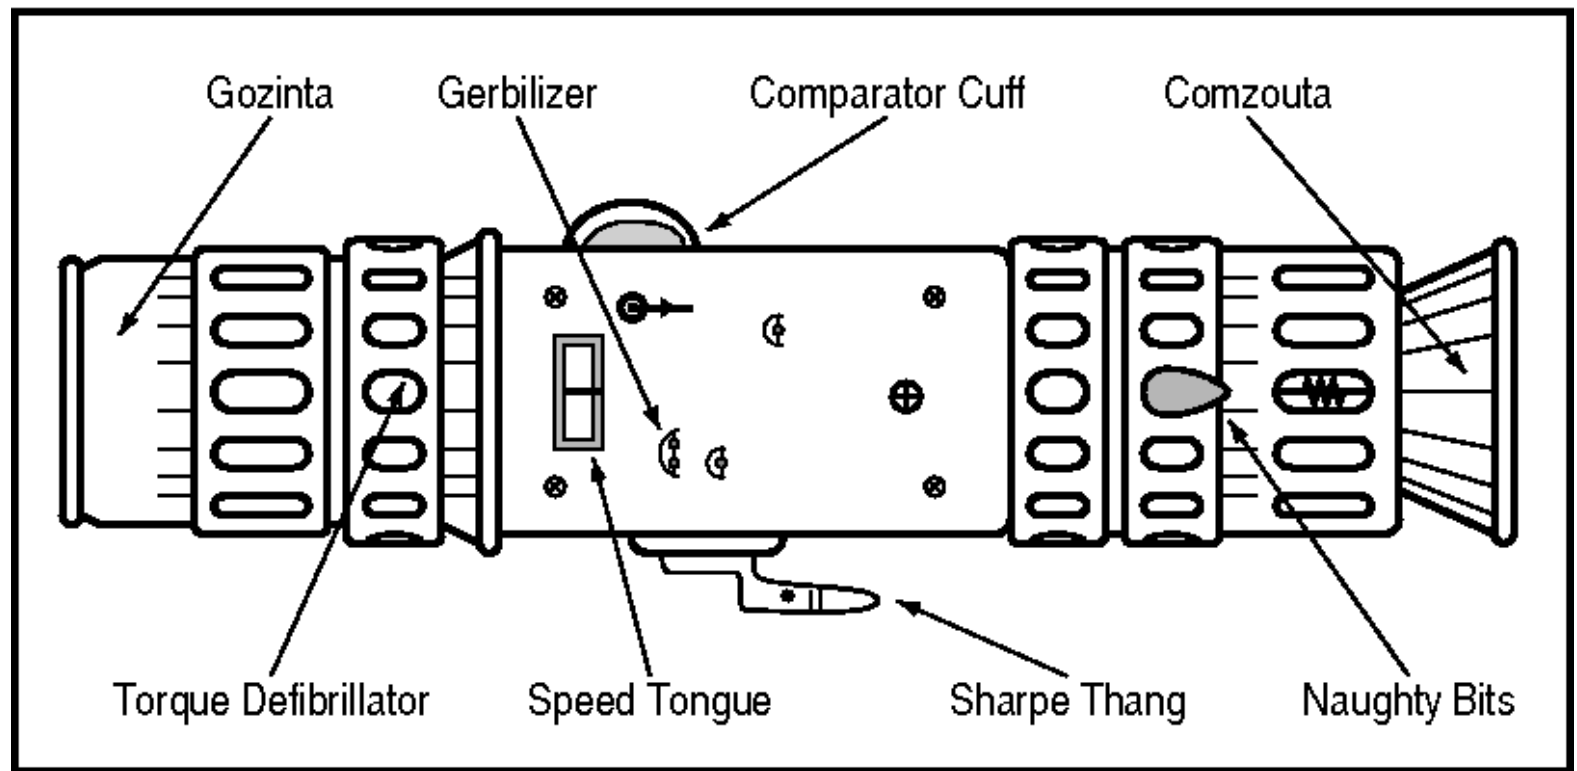
- ✦ There is an opportunity here to leverage both the technology and the installed base.
- ✦ [Preview specifications](#) cover SMB2.2 and other features of Windows 8.
- ✦ This will feed the software engineering ecosystem for years.





SMB2.2

Features





SMB2.2 Features

SMB2.0 was a sleeper:

- 🦊 No user-visible features
- 🦊 Performance improvements were subtle
- 🦊 The user did not even know when SMB2 was being used instead of SMB/CIFS

SMB2.1 offered little more, but 3rd parties started to notice.





SMB2.2 Features

Meanwhile...

- Samba/CTDB added cluster support
- NFS developments included:
 - ✿ NFS over RDMA
 - ✿ Parallel NFS (pNFS)





The competition was moving ahead.





SMB2.2 Features

SMB2.2 Supports:

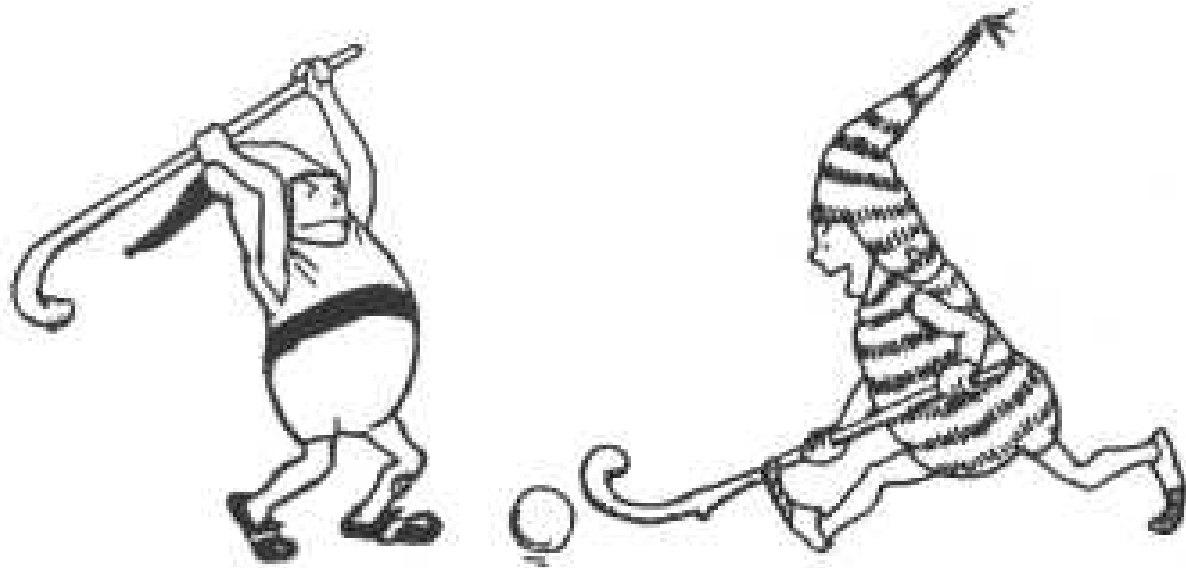
-  Multipath Communication
-  SMB-over-RDMA
-  Scale-up and Failover Clusters
-  Distributed content caching

There is nothing “new” here,
except that it is all in one place
from one vendor.





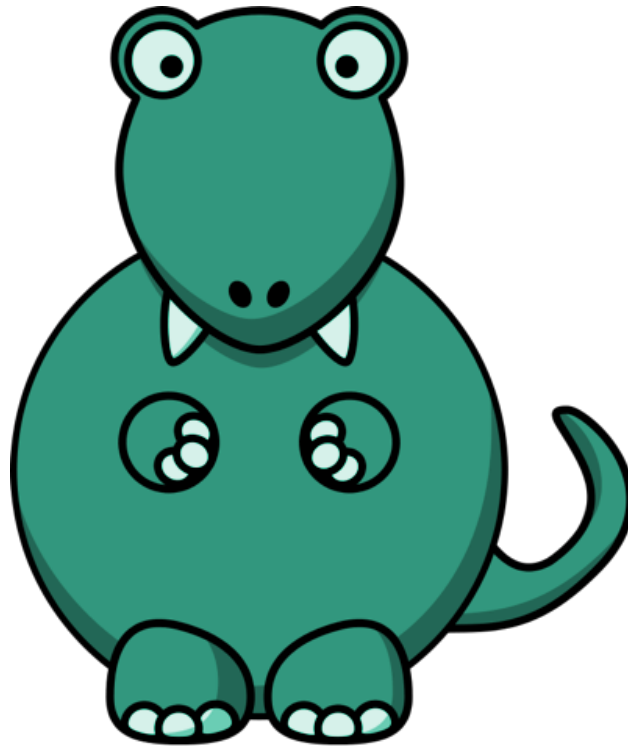
SMB vs. NFS





SMB vs. NFS

The Battle of the Dinosaurs



With SMBv2.2 Microsoft is aiming to conquer datacenter storage, a traditional NFS stronghold.



SMB vs. NFS

SMB1

- Stateful
- Per-user connections
- Simple RPC-style mechanism
- Used as a transport
- Tuned for DOS/OS2 and Windows

NFSv3

- Stateless
- Per-system connections
- RPC-based protocol
- Not a transport
- Generally tuned for Unix/POSIX environments



SMB vs. NFS

SMB2.2

- Stateful
- SMB over RDMA
- Multipath
- Distributed caching and command chaining for improved WAN performance

NFSv4.1

- Semi-stateful
- NFS over RDMA
- Parallel NFS (pNFS)
- Improved authentication and Access Control support



SMB vs. NFS

SMB2.2

- Will simply “be there” in Windows 8 clients and servers
- It will simply “be there” in data centers and on desktops around the world

NFSv4.1

- Despite ten years of open specification development, it isn't “there” yet
- NFS developers appeared jolted by Microsoft's SMB2.2 presentation at the 2011 SNIA SDC



SMB Internals

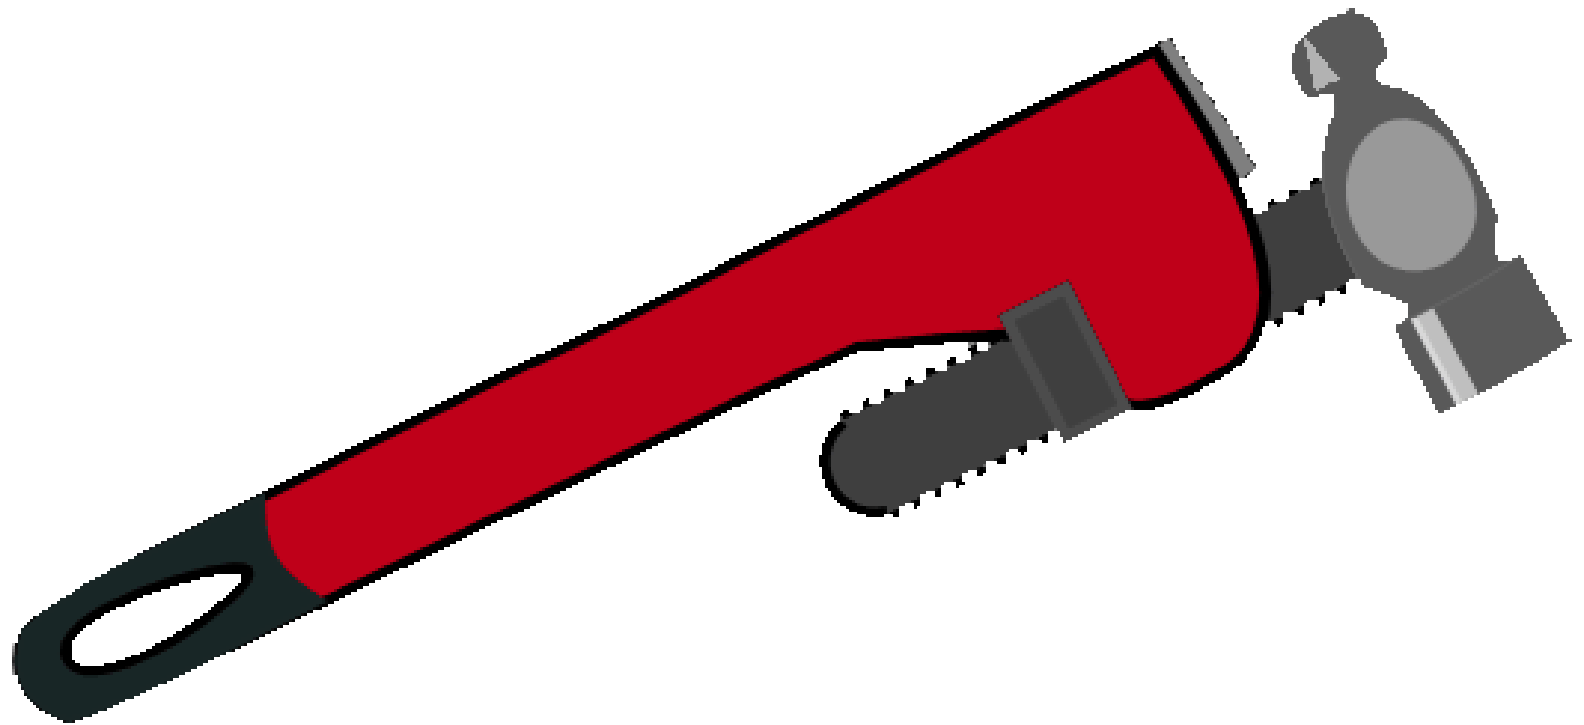





SMB Internals

So... What do SMB and SMB2 look like?

Wireshark and NetMon are your friends...



SMB Internals



File Edit View Go Capture Analyze Statistics Telephony Tools Help

Filter: Expression... Clear Apply

| No. | Time | Source | Destination | Protocol | Info |
|-----|----------|----------------|----------------|----------|---|
| 6 | 0.012509 | 192.168.101.21 | 192.168.101.54 | NBSS | Session request, to PUGSLEY<20> from JCIFS0_2_30<20> |
| 7 | 0.013053 | 192.168.101.54 | 192.168.101.21 | NBSS | Positive session response |
| 8 | 0.013578 | 192.168.101.21 | 192.168.101.54 | TCP | 38427 > netbios-ssn [ACK] Seq=73 Ack=5 Win=5840 Len=0 |
| 9 | 0.015232 | 192.168.101.21 | 192.168.101.54 | SMB | Negotiate Protocol Request |
| 10 | 0.016442 | 192.168.101.54 | 192.168.101.21 | SMB | Negotiate Protocol Response |
| 11 | 0.038718 | 192.168.101.21 | 192.168.101.54 | SMB | Session Setup AndX Request, User: ?\CRH; Tree Connect AndX, Path: \\PUGSLEY\IPC\$ |
| 12 | 0.059613 | 192.168.101.54 | 192.168.101.21 | SMB | Session Setup AndX Response; Tree Connect AndX |
| 13 | 0.061587 | 192.168.101.21 | 192.168.101.54 | LANMAN | NetShareEnum Request |
| 14 | 0.067041 | 192.168.101.54 | 192.168.101.21 | LANMAN | NetShareEnum Response |
| 15 | 0.069211 | 192.168.101.21 | 192.168.101.54 | TCP | 38427 > netbios-ssn [FIN, ACK] Seq=431 Ack=446 Win=7504 Len=0 |
| 16 | 0.070055 | 192.168.101.54 | 192.168.101.21 | TCP | netbios-ssn > 38427 [FIN, ACK] Seq=446 Ack=432 Win=8330 Len=0 |

Frame 11 (248 bytes on wire, 248 bytes captured)

Ethernet II, Src: CameoCom 1f:82:43 (00:40:f4:1f:82:43), Dst: Vmware 40:00:85 (00:50:56:40:00:85)

Internet Protocol, Src: 192.168.101.21 (192.168.101.21), Dst: 192.168.101.54 (192.168.101.54)

Transmission Control Protocol, Src Port: 38427 (38427), Dst Port: netbios-ssn (139), Seq: 124, Ack: 98, Len: 64

NetBIOS Session Service

SMB (Server Message Block Protocol)

SMB Header

Server Component: SMB

[Response in: 12]

SMB Command: Session Setup AndX (0x73)

Error Class: Success (0x00)

Reserved: 00

Error Code: No Error

Flags: 0x18

Flags2: 0x0001

Process ID High: 0

Signature: 0000000000000000

Reserved: 0000

Tree ID: 0

Process ID: 46208

User ID: 0

Multiplex ID: 2

Session Setup AndX Request (0x73)

Word Count (WCT): 13

AndXCommand: Tree Connect AndX (0x75)

Reserved: 00

AndXOffset: 142

```

0010 00 ea f2 36 40 00 40 06 fc 3a c0 a8 65 15 c0 a8 ...6@.@. ....e...
0020 65 36 96 1b 00 8b 03 3a e1 90 00 06 d4 c8 50 18 e6.....P.
0030 16 d0 23 d4 00 00 00 00 00 be ff 53 4d 42 73 00 .#.....SMBs.
0040 00 00 00 18 01 80 00 00 00 00 00 00 00 00 00 00 .....u...
0050 00 00 00 00 80 b4 00 00 02 00 0d 75 00 8e 00 14 .....NY^ .>...0
0060 05 0a 00 01 00 00 00 00 00 18 00 18 00 00 00 00 ...>3...sF..B.
0070 00 14 00 00 00 51 00 45 68 ba ae f9 8e db f2 84 ...../..C.R.H...
0080 17 90 03 a2 bc 4e 59 5e d3 7d 3e 97 1a 1f ee 30 ?...L..n.u.x...
0090 f5 da 3e 3e 33 aa f4 18 99 e0 73 46 d6 16 42 ee f.o.o.....
00a0 18 eb d5 c3 8f 2f 2c 00 43 00 52 00 48 00 00 00 %..\\..P.U.G.S.
00b0 3f 00 00 00 4c 00 69 00 6e 00 75 00 78 00 00 00 L.E.Y..I.P.C.$.
00c0 66 00 6f 00 6f 00 00 00 04 ff 00 00 00 00 00 01 ..7777?
00d0 00 25 00 00 5c 00 5c 00 50 00 55 00 47 00 53 00
00e0 4c 00 45 00 59 00 5c 00 49 00 50 00 43 00 24 00
00f0 00 00 3f 3f 3f 3f 3f 00
  
```

SMB (Server Message Block Protocol) Packets: 17 Displayed: 17 Marked: 0 Profile: Default




SMB Internals

SMB/SMB2 are “stateful”

... so what state is maintained?

In the Microsoft documentation:

- 
- 🍎 A state machine is presented, providing:
 - 🟢 Different context levels in which state is kept
 - 🟡 ...with constructors and destructors
 - 🟢 A set of state variables per context level
 - 🟢 A set of operations that change the state
 - 🟢 Relationships between variables is maintained across documents
 - 🍎 This is in addition to basic packet formats, given in a separate section.






SMB Internals

This page intentionally left blank





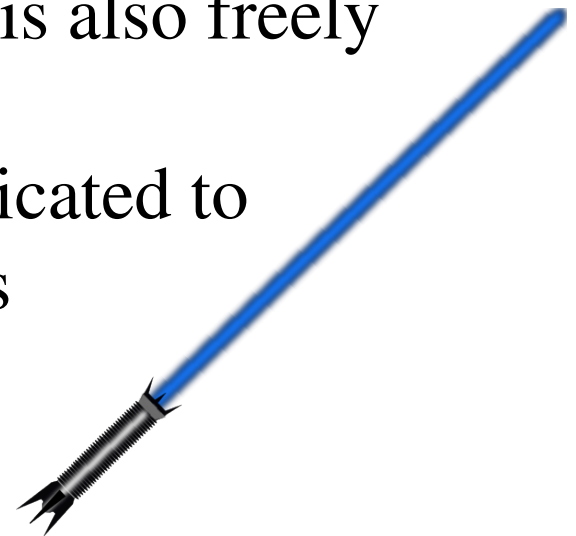
SMB Internals

Read the Docs

-  **Implementing CIFS** is an Open Source book, freely available on the Web
-  **Microsoft's Documentation** is also freely available from their website
-  **CIFS.Org** is a Wiki site dedicated to capturing developer insights

Use the Source

-  **Samba** provides a complete server suite, available for study
-  **jCIFS** is an SMB client implementation in Java, also available for study





“Real-World” SMB





SMB in the “Real World”

We have established:

- ④ SMB is difficult and annoying
 - SMB2 is a little better
- ④ SMB/SMB2 are *de facto* standard protocols
 - ...and market leaders
- ④ Adoption of NFSv4.x is slow
- ④ SMB2.2 is feature-comparable to NFSv4.x



What are you going
to do about it?



SMB in the “Real World”

Consider the Market for Storage Engineers

In Silicon Valley, you can put out a sign that says “NFS coders wanted”, and a line will form. NFS is an open specification, studied in University classes. There are multiple books on the subject.

In comparison, the pool of SMB/CIFS/SMB2 engineers is very small indeed.

Scarcity == Opportunity





The End



