Development Using Samba 4

Jelmer Vernooij
Samba Team / OpenChange Team
jelmer@samba.org

Epitech
November 2007

Agenda

- The SMB protocol
- Samba
  - Quick history
  - Samba 4
- Development Process
- Exported libraries
- OpenChange

The SMB protocol

- Protocol for simple file sharing between DOS machines
- Created in 1983 at IBM by Barry Feigenbaum
- Aimed at low-end pcs, so optimized for speed
- Development continued by Microsoft as of later 80's
- Somewhat documented

Feature Creep

- Used as transport for other protocols:
  - RAP: Remote administration
  - Mailslots: Browsing the network
  - DCE/RPC: Printing, Remote Login, Registry, ...
- Several extensions:
  - Dialects for different Windows versions
  - Various mechanisms for authentication

DCE/RPC

- Generic protocol for remote function calls
- Interface specified in Interface Description Language (IDL)
- Marshalling code generated by IDL compiler
- Used for printing, registry, ...

Samba

Samba

- Implementation of SMB and related protocols for POSIX-systems
- Originally developed in '91 by Andrew Tridgell
- Free Software (GPL)
- Development team of ~30 people
  - 10-15 active contributors

Mapping

- Differences in semantics
  - Unix: case-sensitive
  - Windows: case-insensitive, but case-preserving
  - Strange attributes:
    - Delete-on-close
    - “Streams”

Samba - The Code

- Languages used:
  - Mostly C
  - Some perl (during building only)
  - Scripting in JavaScript
    - Soon to be replaced by Python
- Test driven development (... mostly)

Samba 4

- Originally started in 2003, as an effort to improve the SMB server
- Later became the effort to get an AD DC working
- Strong focus on the right infrastructure
- 14000 commits since April 2004, 4000 in the last year

Samba 4 - On our way to an alpha

- Improvements in the last year
  - Protocol knowledge
  - Protocol Coverage
  - Usability Improvements
  - Development Tool Improvements

Development Process

Test driven development

1)Write test that checks a feature is present

2)Check that test fails

3)Implement the feature

4)Make sure test succeeds

Interoperability

- Reasonably well documented at first
- Competition with NFS original reason for specification
  - RFC1000/RFC1001
  - Renamed to “CIFS”: *Common Internet File System*
- Since '98 mostly undocumented

Specifications

"The only spec I trust is written in C"
- *Andrew Tridgell*

Network Analysis

- Passive
  1. Try operation between two Windows machines
  2. Hardcode result in own implementation
  3. Change input to determine other fields
- Active
  - Try other values for fields and see if that changes the return value
- Not the same thing as reverse engineering
  - "French Café"

Network Analysis - Example

Wireshark

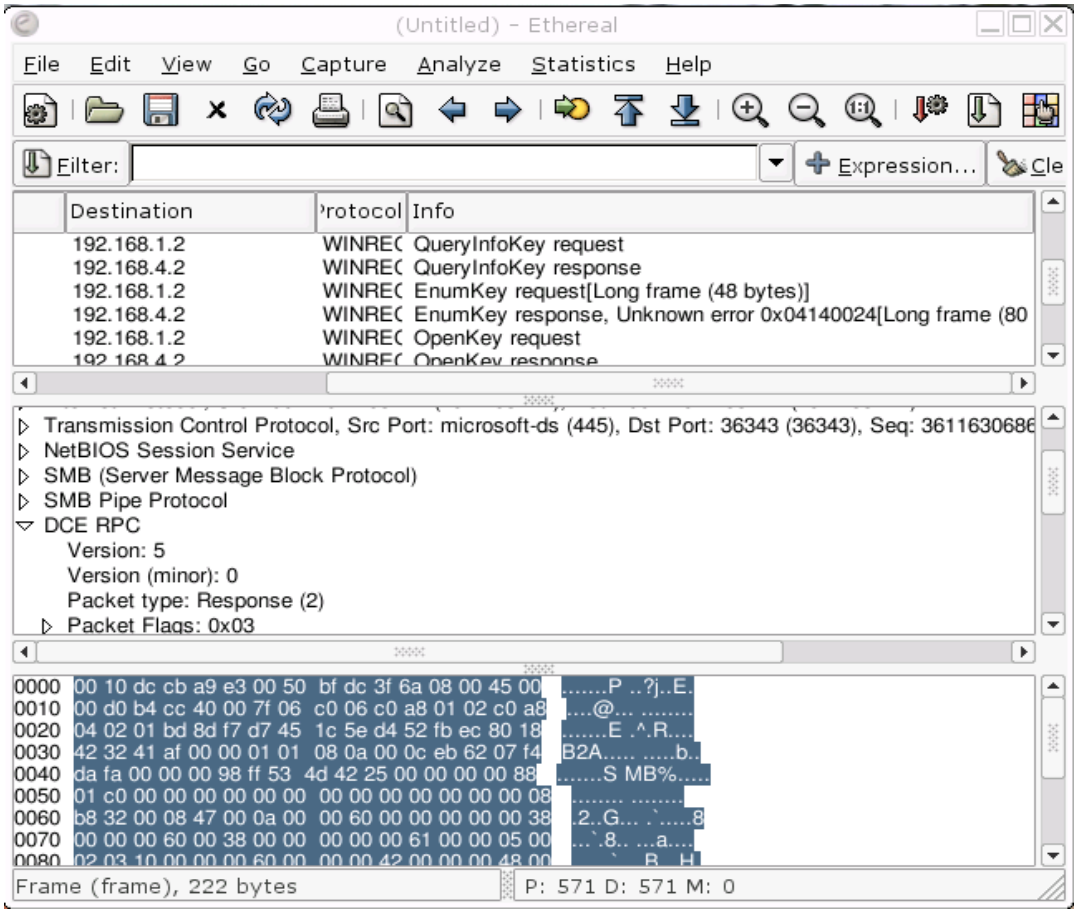

- Most important tool during network analysis
- Already decodes a lot of protocols
- Free Software

Man in the middle

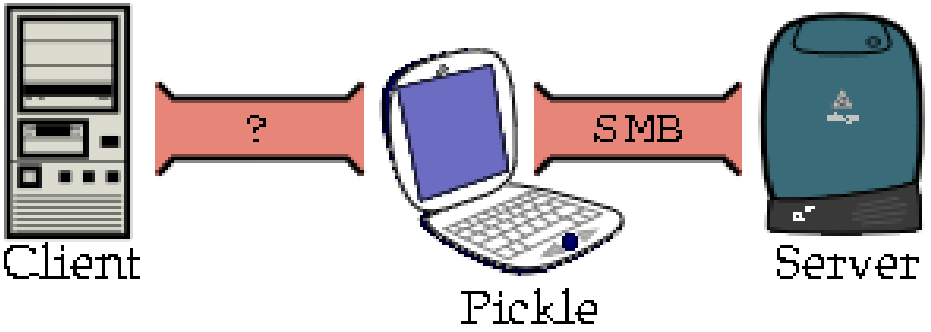

- Intercepts server traffic
- Change it
- Send modified request to client

Gentest

1)Generate random commands

2)Send commands to Windows and Samba

3)Compare results

4)Repeat 1-3 until there is a difference

5)Backtrack

Code Coverage

- Percentage of the code that is run during the testsuite
- 41.6% now!

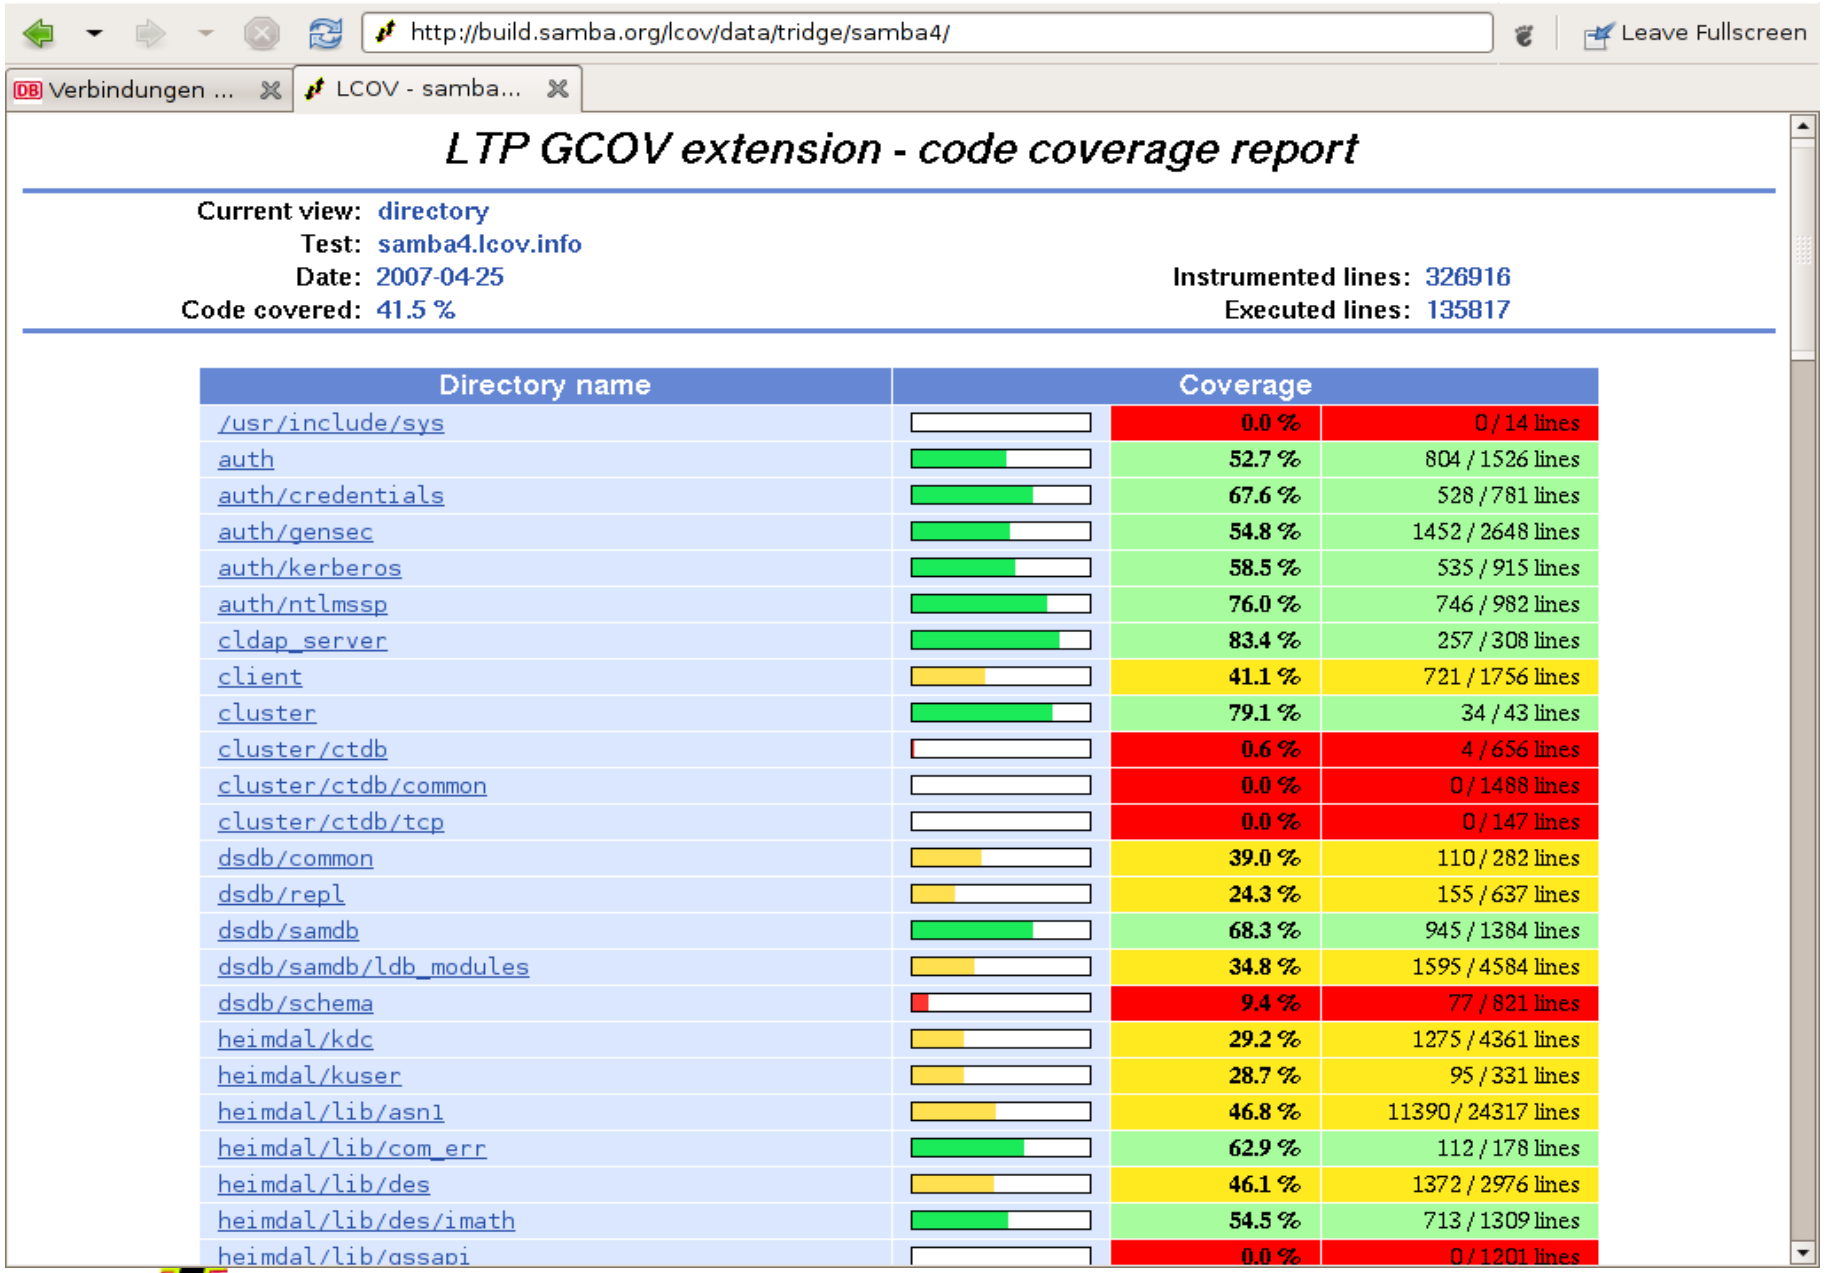

LTP GCOV extension - code coverage report

Current view: directory
Test: samba4.lcov.info
Date: 2007-04-25
Code covered: 41.5 %

Instrumented lines: 326916
Executed lines: 135817

Directory name	Coverage	
/usr/include/sys	0.0 %	0 / 14 lines
auth	52.7 %	804 / 1526 lines
auth/credentials	67.6 %	528 / 781 lines
auth/gensec	54.8 %	1452 / 2648 lines
auth/kerberos	58.5 %	535 / 915 lines
auth/ntlmssp	76.0 %	746 / 982 lines
cldap_server	83.4 %	257 / 308 lines
client	41.1 %	721 / 1756 lines
cluster	79.1 %	34 / 43 lines
cluster/ctdb	0.6 %	4 / 656 lines
cluster/ctdb/common	0.0 %	0 / 1488 lines
cluster/ctdb/tcp	0.0 %	0 / 147 lines
dsdb/common	39.0 %	110 / 282 lines
dsdb/repl	24.3 %	155 / 637 lines
dsdb/samdb	68.3 %	945 / 1384 lines
dsdb/samdb/ldb_modules	34.8 %	1595 / 4584 lines
dsdb/schema	9.4 %	77 / 821 lines
heimdal/kdc	29.2 %	1275 / 4361 lines
heimdal/kuser	28.7 %	95 / 331 lines
heimdal/lib/asn1	46.8 %	11390 / 24317 lines
heimdal/lib/com_err	62.9 %	112 / 178 lines
heimdal/lib/des	46.1 %	1372 / 2976 lines
heimdal/lib/des/imath	54.5 %	713 / 1309 lines
heimdal/lib/gssapi	0.0 %	0 / 1201 lines

Static Analysis (checker)

- Analyses source code without running it
- Developed by IBM
- Run regularly on the build farm

Build Farm

http://build.samba.org/

- Continuously builds Samba on various hosts
- Different operating systems
- Now shows test coverage
- List with most often broken tests

Samba Libraries

- libsmbclient
  - Used by GNOME, KDE for accessing files on remote machines
- DCE/RPC
  - Used by OpenChange

OpenChange

- Built on top of Samba4
  - Uses the DCE/RPC library for client side access and the Samba 4 IDL compiler
  - Hooks into smbtorture

Releases

- TP4: January 2007
- TP5: May 2007
- Alpha1: September
- Alpha2: January 2008 ?

Contribute!

- IRC: #samba-technical on irc.freenode.net
- Mailing list: samba-technical@samba.org
- Wiki: http://wiki.samba.org/