

Samba Status Update

Munich

13. March 2008

Volker Lendecke

SerNet

Samba Team



Volker Lendecke

- Co-founder SerNet - Service Network GmbH
 - Free Software as a successful business model
 - Network Security for the industry and the public sector
 - Samba-Support/Development in Germany
- For almost 20 years concerned with Free Software
- First patches to Samba in 1994
- Consultant for industry in IT questions
- Co-founder emlix GmbH (Embedded Systems)



What is Samba?

- Interoperability between Windows and Unix systems
- Most protocols Windows speaks today
 - SMB (File Sharing), Printing, Browsing, Authentication
 - „Samba makes unix machines show up in Network Neighborhood“
- Samba runs on most Unixes these days
 - Main development platform is Linux
 - Solaris, AIX, HP/UX, Stratus V/OS, Tru64, etc...



Samba 3

- Stable version of Samba, production code
- Solid file and print server
- NT and Active Directory domain member
- NT4-compatible Domain Controller
- Security and access control model based on Unix
- Flexible posix-based VFS modules to change file access semantics
- Current version: 3.0.28a



Samba 4

- Started in 2003 by Andrew Tridgell as a new VFS system with the goal to support cluster file systems
- Complete re-write of Samba
- Target: 100% semantics of Windows
- Main feature today: Active Directory Domain controller
- Most useful component for development: smbtoriture, a very comprehensive SMB test suite



Samba 3.2

- License change: GPLv3
 - This caused the naming change from 3.0 to 3.2
- 3.0 will only be maintained for security updates
- When will 3.2 come?
 - It's long overdue
 - New Release Manager: Karolin Seeger
- 3.2.0pre2 was released March 4, she's actively pushing us to get it out soon



IPv6

- NetBIOS over TCP (RFC1001 / 1002) only specifies IPv4
- RFC1002 name resolution naturally embeds IP-addresses
- With AD and CIFS, DNS has become the primary name resolution
 - Nobody is going to change RFC1002
- Samba 3.2 listens on IPv6 interfaces now
 - Main changes: Interface enumeration and name resolution for Samba as a client (getaddrinfo)



Memory usage

- The thumb-rule is that you need 2-3MB real RAM for an active, connected client
 - Quite a lot for embedded systems (the 50 Euro NAS hard drive)
- Number of clients on large systems mainly limited by RAM usage, there are numbers from competitors that advertise less than 100k per client
- The footprint for an idle client was reduced by 50%



Memory handling

- Smbd used to allocate static buffers for the largest possible SMB request: 128k incoming, 128k outgoing
- 3.2: allocation on demand
- All of smbd internally assumed to have 128k available, *many* internal changes
- Many other buffers now malloc'ed on demand, not by maximum use



No static pstring buffers

- Pstring: Rest in peace :-)
- All file names used to be stored in 1k buffers
- The typical file name is a lot shorter than 1k
- Everything malloc'ed on demand now
- Malloc started to show up high in profiling
 - Changes to talloc to use one pool per SMB request without calling malloc



Malloc implementation

- For some requests, we do have to allocate much memory
- Glibc malloc is not nice to forked processes
 - Smbd parent does quite a bit of memory allocation which is never changed in the child
- Allocations in the child cause write access to parent malloc internal structures, causing copy on write where smbd does not write itself
- Playing around with a malloc of our own



Registry configuration

- Parsing and writing smb.conf files with GUI tools is awkward at best
- Samba 3 has to implement a registry, clients expect to find certain keys to determine the server type
- Registry data model matches exactly smb.conf format, it was designed as a .ini file replacement
- HKLM\Software\Samba\smbconf
- Enabled only if „config backend = registry“ is enabled in the smb.conf text file



Cluster support

- Samba 3.0 was heavily modified to enable sharing of cluster file systems correctly via CIFS
- Product shipped by IBM right now based on a heavily modified 3.0.25
- Ctdb.samba.org has all Samba-related components
- 3.2 contains the relevant changes
- Heavy, serious testing needs to be done on the cluster feature of 3.2, that's why we mark it as experimental for the release



AD LDAP signing

- Policies on W2k3 domain controllers can require that all LDAP access needs to be authenticated and integrity protected
- Cyrus SASL is the standard OpenLDAP way to provide this, but it is hard to configure
- Current OpenLDAP libraries allow to intercept the network traffic
- Samba 3.2 winbind does the SASL integrity wrapping itself



LDB

- Ldb is a lightweight library providing an LDAP-like data model on top of tdb
 - Same idea as SQLite for SQL
 - Main feature: Multi-indexed tdb
- Samba needs to maintain its internal user database
- As a start, group mapping has been converted to ldb
 - In the future, maybe more will be converted



SMB transport encryption

- All SMB authentication variants provide a session key as one result
 - This is used to encrypt the password changes
 - RPC can be encrypted
 - Normal SMB traffic (file content) is plain text
- Samba 3.2 re-uses the RPC encryption routines to encrypt the whole SMB data stream
- Currently only smbclient->smbd
- Now we have to persuade MS to use this :-)



RPC interface generation

- Many important sub-protocols in the SMB world are DCE/RPC based
 - Marshalling of function call arguments has been done with manual code
- „Real“ RPC implementations use IDL
- Samba4 provides the Perl-based IDL compiler IDL
 - Contrary to MIDL (the MS tool), PIDL generates readable code
- Lots of Samba 3.2 code is converted, much ugly code is gone



Questions/comments?

Volker Lendecke, VL@SerNet.DE

SerNet - Service Network GmbH
Bahnhofsallee 1b
37081 Göttingen

Tel: +49 551 370000 0

Fax: +49 551 370000 9

<http://www.SerNet.DE>

<http://Samba.SerNet.DE>

