

## Samba Status Update

SNIA SDC 2011  
Santa Clara, CA

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
- Free Software for >20 years
- First patches to Samba in 1994
- Consultant for industry in IT questions
- Co-founder emlix GmbH (Embedded Systems)



## SerNet

- SLA based support for more than 650 customers
- network security for industrial and public customers
  - firewalls, VPN, certificates, audits
  - based on open standards wherever possible
- Support for many OS: Linux, Cisco IOS, Windows etc.
- Compliant with BSI Grundschutz and ISO 27001 and other international regulations

## SerNet and Samba

- technological leadership of SerNet worldwide
  - involved in almost every big European Samba project
  - 5 out of 6 European developers work for SerNet
  - SerNet distributes up-to-date Samba packages
- samba eXPerience
  - *The* international Samba conference
  - > 150 developers & users from > 15 countries

## 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: 3.6.0

## 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
- Samba 4.0 will merge Samba 3 and Samba 4 code bases

## Samba 3.6.0

- Released Aug 9, 2011
- SMB2 support
- Tighten security defaults (client uses ntlmv2 by default)
- Printer subsystem overhauled
- Internal use of RPC interfaces
- Winbind idmap configuration changed (again :-())
- Tons of bugfixes

## Cluster support

- Clustering is still hot, several solutions being presented every year
- Several Posix-level cluster file systems available:
  - GFS, OCFS2
  - GPFS, Panasas, CXFS, Isilon and many more
- NFS clustering is relatively easy on top of those
- SMB clustering puts very heavy demands on the cluster lock manager

## CTDB

- None of the existing lock managers provided the semantics needed for CIFS clustering
  - Samba requires locks with associated data, a big share of Samba is to implement the correct locking
  - Many lock managers are much too slow
- Ctdb is the „clustered tdb“ lock manager
  - So ctdb is more like a complete HA solution
  - Recent developments remove HA aspects for cooperation with existing HA suites

## Cluster support

- Samba 3.2: First version with cluster support
- Fedora has ctdb and Samba cluster-aware
  - To be shipped as part RHEL on top of GFS2
- SUSE ships ctdb and Samba as part of SLES HA
- Ctdb.samba.org has all Samba-related component
- More cluster file system vendors working on it
  - Make ping\_pong.c work!

## Current development

- CTDB is a full HA solution
  - IP address management, service monitoring
  - Service monitoring is moved out for cooperation with other HA solutions
- CTDB readonly locks / records
  - Reading a huge file from different nodes sucks
  - Brlock checks make records bounce on every read
  - Multiple copies of brlock records will be available, pulled back on locking&x calls

## SMB2

- New protocol introduced with Windows Vista
- Basic SMB2 is semantically very similar to SMB1
  - Main call CreateFile is almost the same
- New protocol features added only to SMB2
- Main advantage: Huge MS client improvements
  - SMB1 abandoned

## SMB2 in Samba

- First prototype implementation in Samba4
  - No docs available back then, so the client library and torture suite was done first to discover the protocol
- Samba 3.6 fully supports SMB2.0
- Enabled with „max protocol = smb2“
- Next:
  - Durable file handles
  - SMB2.1

## RPC servers

- Samba has a home-grown RPC implementation
- All RPC servers are in-process and single threaded
- Large file servers suffer from memory footprint
  - LSA, SAMR, WINREG etc take space in every smbd
  - SPOOLSS possibly the worst memory hog
- Samba 3.6 can offload RPC traffic via Unix sockets
- Samba master can offload RPC to central daemons
  - Smbd workers only do file and RPC transport

## Printer subsystem overhauled

- Windows stores printer data in the registry
- Printer drivers assume the Registry data model
- ntprinters.tdb etc presented as registry
- Some printer drivers break because ntprinters can't store / and \ (yes, this is a bug :-))
- Printer subsystem now **uses** the registry instead of faking it

## Antivirus support

- samba-vscan not really supported
  - Not in-tree due to GPL concerns (linking with proprietary libraries)
- Scannedonly: Only show scanned files, external daemon: <http://olivier.sessink.nl/scannedonly/>
- Open questions:
  - Scan on read/write, scan on close after write?
  - Asynchronous scanning?
  - Standardized interface for GPL compliance

## Async echo responder

- Some file system operations can be very slow
- Single threaded `smbd` sits in syscalls
  - `unlink(100GB file)` can take ages
- Windows clients start to send `SMBEcho`
- „async smb echo handler = yes“
  - Forks a process that responds to `SMBEchos` when the main worker `smbd` hangs
- Not required for `SMB2`, According to `MS` clients use `TCP` keepalives

## Idmap autorid

- Idmap tdb: 1 tdb entry for each SID
  - winbindd\_idmap.tdb is a persistent tdb (slow!)
- Idmap rid: No tdb entries, static configuration
- Idmap hash: No config, but risk of collision
- Idmap autorid: Like idmap rid with dynamic domain prefix allocation
  - Just 1 tdb entry per domain
  - No configuration, as fast as idmap rid

## Performance monitoring

- No statistics from inside Samba so far:
  - SMB requests per second? Latency?
    - Data transfer rates per share/user/connection?
- Performance collection must be cheap
- Intended design:
  - Mmap area per smbd writing data without mutexes
  - Single monitor process consolidates data
  - Little inaccuracies (no mutex) acceptable?

## Samba 3/4

- Samba 3: Good file/print server/domain client, working NT4 PDC
- Samba 4: Starts to become an AD/DC, file server lacks a lot, domain client barely existing
- Samba 4.0 will contain Samba 3 file/print server and winbind, Samba 4 AD components
- Subsystems shared between 3 and 4 (tdb, talloc, etc)

## Active Directory Domain Controller

- The main services beyond NT4: LDAP and Kerberos
- Kerberos is almost standard, Samba4 ships heimdal
- LDAP has been extended extensively for AD
  - DRSUAPI for multi-master replication
  - Many improvements for scalability and consistency
- Samba4's LDAP/DRSUAPI server is right now under heavy development
  - Corporate interests by many vendors

## Current work in Samba4

- Single Domain DCs work fine
  - LDAP, Kerberos, Directory Replication works
- Replication with W2k3/W2k8 works fine
- Migration scenarios: Replicate to and from Windows
- Build system: A normal Samba4 build will build S3 as well to enable cooperating processes
- Upgrade S3 to S4

## Read Only Domain Controller

- New Windows 2008 feature
- Branch offices get selected data sets Read Only
- In 2010, the RODC was developed in Samba4
- Safe alternative to fully replicating Dcs
- Samba 4 is a full AD replication member
  - Bugs in Samba4 could destroy the AD database
  - Windows DCs don't accept changes from Samba

## S3FS

- S3FS: Embedded fileserver
- Starts Samba3 smbd inside Samba4
- Named pipe bridge
  - IPC\$ `\\pipe\<<pipename>` is forwarded to Samba4
  - Forwarding completely asynchronous
- Aims to replace the unfinished Samba4 VFS
- S3FS enables the existing S3 VFS modules to continue working

## S3 utilities

- „smbpasswd -a <username>“ **MUST** continue to work
- Access to AD user database:
  - pdb\_ads („Franky“ approach): Independent implementation, requires running Samba4
  - pdb\_samba4 („s3compat“ approach): Direct access to ldb
- „net“, „smbclient“, „smbstatus“ used from Samba3
- „samba-tool“ for AD-specific things

## winbind

- Filesystem access requires unix users
- Winbind is used for AD members and trusting Dcs
  - S4 winbind has very rudimentary id mapping
- Samba4 AD can work without local unix users
- Services for Unix (SFU) schema will be used for winbind running on S4 Dcs
- Unix ID allocation requires something like the RID allocator

## Naming things: What will Samba4 be?

- Samba 4.0 will contain both Samba4 as well as Samba3 components
- Basic S3 will continue to work as today
  - Support for exotic platforms
  - File/Print/Domain Member will continue to work without any S4 involvement
- AD support will require a slightly advanced Platform
  - For example Python will be required for both building and running it

## 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>

## 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>