Chemnitzer Linuxtage 2019

Volker Lendecke

Samba Team / SerNet

17. März 2019

Who am I?

- Co-Founder of SerNet in Göttingen, Germany
- First Samba patches in 1994
- Early Samba Team member
- Samba infrastructure (tdb, tevent, etc)
- File server
- Clustered Samba
- Winbind
- AD controller is my colleague Stefan Metzmacher's domain
  - Stefan implemented AD multi-master replication in Samba

What is Samba?

- www.samba.org: Samba is the standard Windows interoperability suite of programs for Linux and Unix
- Server- and Client-Implementation of the Server Message Block (SMB) protocol
  - SMB is **the** Windows protocol to share drives across the network
  - Comparable to NFS (NFSv4 copied many SMB features ...)
- Print server for Windows clients
- Active Directory domain member
  - Make Active Directory users and groups available on Linux
- Active Directory domain controller
  - Provide user database for Windows and Unix clients

Samba Release Cycles

- Regular release cycle is six to nine months
- Current release fully supported (4.9)
  - Bug fixes, some new features
- Previous release (4.8)
  - Only bug fixes
- Before previous release (4.7)
  - Security fixes only
- Samba 4.6 went out of security support with 4.9
- 1.5 to 2 years of security support
- https://wiki.samba.org/index.php/Samba_Release_Planning has a nice table

Active Directory

- Microsoft's central user database
  - Successor to NT4-based Security Account Manager (SAM)
  - It's what eDirectory is for the Bindery (Novell anyone?)
- Kerberos KDC with an LDAP database backend
- Multi-Master replicated LDAP database
- Highly specific LDAP schema with custom extensions
  - A lot of internal magic and validity checks
- Authentication server for Challenge-Response based schemes
- DNS database for server lookup

LDAP data store

- For historical reasons, Samba implements its own LDAP server
  - The custom extensions were not well received by OpenLDAP
  - This has changed now, but it's a huge effort to change this
- Samba LDAP stores AD in tdb
  - TDB (Trivial DataBase) is a simple key/value store
  - Mainly for small, highly volatile records (AD is not volatile)
  - Transactions were added later
  - TDB is limited to 4GB (32-bit)
- Howard Chu has developed LMDB for OpenLDAP
  - Highly tuned btree implementation
  - 64-bit, so no size limitations
- Samba 4.9 allows the use of lmdb

Trust Relationships

- Active Directory can scale to huge numbers of users
  - A single administrator account for the whole database
  - Organizations don't always trust each other
- Multiple ADs can be linked by Trust Relationships
  - Domain A trusts Domain B $\Rightarrow$ Users from B can log in to A
- Pure domain member (winbind) has supported trusts for ages
- Samba 4.8 starts to support trusts as an AD DC
  - Some limitations: Domains have to fully trust each other
  - Main goal not yet achived, but we're close

Audit Logging

- For compliance reasons, many users need detailed logs of all authentication-related actions
- Samba 4.9 adds logging in JSON format
- { "timestamp": "2019-03-16T08:29:32.635688+0100",
  "account": "vlendec",
  "sid": "S-1-5-21-3898457107-206185458-3872133680-1000",
  "logonServer": "DC1",
  ... Lots of other information }
- Consumable by many backends

Active Directory Backup

- AD backup is not like other backups
  - Multi-Master replication creates any number of valid copies
  - A single DC that lost its db can be purged and re-joined
  - Normal replication will restore the full database
- Samba 4.9 adds “samba-tool domain backup online”
- Samba 4.10 extends to “samba-tool domain backup online”
- Restore must be done very carefully
  - All Domain Controllers must be switched off
  - One DC is restored from backup
  - All other DCs must be re-joined and replicated
  - Otherwise, replication is destroyed

Clustered Samba

- Samba can provide a single SMB server across multiple nodes
  - Underlying clustered file system
  - GFS, Ceph, Gluster, OCFS, GPFS, StorNext, CXFS, Panasas
- Main task: Helper for smbd to distribute internal databases $\Rightarrow$ Clustered TDB (CTDB)
- CTDB provides a cluster manager
  - Node membership
  - Service monitoring
  - IP address assignment
- All these tasks are becoming modularized
  - Integration into existing clustering solutions gets easier
  - Monitoring and IP management now someone else's problem
- Configuration changed radically with 4.9
  - No longer just arguments to ctdb, we now have a proper config file

SMB2 unix extensions

- SMB implements the semantics Windows clients expect for files
- Cool SMB features:
  - Security: Authentication, Signing, Encryption
  - High performance through multi-channel, SMB over RDMA
  - Cache coherence protocol
  - NFS can do most of it, but only with NFSv4.1
- Linux clients can mount SMB shares
- Not everything works as with a local file system
  - File names are case insensitive: test.txt and Test.txt are the same
  - No symlinks, sockets, device nodes
  - Different locking semantics

SMB2 unix extensions

- Goal: Run home directories off SMB (kill NFS?)
  - SMB1 was extended to do all this: "unix extensions = yes"
  - SMB1 is deprecated, insecure, slow
- Apple switched from AFP (Apple File Protocol) to SMB
  - Semantic extensions for Apple specific features
  - Optimizations for resource fork and Finder metadata
  - $\Rightarrow$ Protocol is extensible
- OS/X is Unix based, why not use those extensions?
  - Not **quite** what we need
  - For example, no symlinks
- SMB2 extensions: Minimum required change
- Status: Mostly works with Jeremy Allison's patches, now working to get them upstream

Questions?

vl@samba.org / vl@sernet.de
http://www.sambaxp.org/