Samba Status

FFG 2016
Köln

Volker Lendecke

Samba Team / SerNet

2016-02-25

SerNet

- SLA based support for more than 650 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
    - SerNet distributes up-to-date Samba packages
  - samba eXPerience
    - May 10-12 in Göttingen, www.sambaxp.org

Samba 4.3

- New FileChangeNotify subsystem
- Profiling code
- Improved security for winbind
- SMB 3.1.1: Better security
- AD DC: Trust support
- Samba KCC improved, but still disabled
- Samba 4.4: Mostly small changes, clustering improvements

Release Cycles

- Regular release cycle is nine months
- Current release fully supported (4.3)
  - Bug fixes, some new features
- Previous release (4.2)
  - Only bug fixes
- Next to previous (4.1)
  - Security fixes only
- Samba 4.0 went out of security fix support with 4.3
- All code from 3.6 continues to live
  - File server, print server, NT-style DC

What is FileChangeNotify?

- MSDN on "Obtaining Directory Change Notifications":
  - An application can monitor the contents of a directory and its subdirectories by using change notifications.
- Client queries a directory handle for changes
- Filters are sent for just specific events:
  - "I'm only interested in new and deleted files"
  - "Please tell me when a file size changes"
  - ...
- API parameter **bWatchSubtree**:
  - If this parameter is TRUE, the function monitors the directory tree rooted at the specified directory.

FileChangeNotify in Samba

- On every change, Samba has to check for interested clients
- Four implementations
- Samba 3.0
  - Timeout-based polling of directories per smbd
- Tridge's Samba4 implementation
  - Tridge figured out how much more of the protocol
  - One big array of all Notify Requests in every smbd
  - Messaging-based notification
  - Ported to Samba 3.2
- Samba 4.0 notify_index.tdb
  - Starts to make notify possible in a cluster
- Samba 4.3 notifyd

FileChangeNotify in 4.3 – notifyd

- notify_trigger: "This function is called a lot ..."
- For every directory component in a new/changed/removed file, we must check for interested clients (**bWatchSubtree**!)
  - This function (notify_trigger) is O(n) in the number of path components
- Notify events must be as cheap as possible
  - FileChangeNotify is asynchronous
  - notify_trigger now delegated to another process (notifyd)
- Samba now has cheap inter-process messaging based on unix domain datagram sockets

notifyd Benefits

- One message per metadata modification
  - Unix domain datagram messages do roughly 150k/sec (on my Laptop)
- Less load on inotify
  - One notify listener instead of every smbd
- Clusterwide file change notify
  - Many cluster file systems do not provide clusterwide inotify
  - inotify works locally, notifyd tells others
- External event sources (Ganesha?)
  - A single unix dgram per event
  - Extremely simple protocol

Profiling code

- Samba measures request counts, request times, VFS calls, latencies, etc
  - Lots of probe points
  - Low performance impact necessary
- How to collect performance data from hundreds of smbd processes
  - Only shared memory is fast enough
  - Samba used (shock, horror...) sysV IPC shared memory
  - Every smbd just incremented counters in a central mmap area
  - Atomicity, NUMA effects were just ignored
- Samba has a very good mmap abstraction: tdb
  - With Samba 4.3 every smbd maintains its own tdb record for profiling
  - Once a second, data is assembled into one record

winbind changes

- Tightened security settings
  - Are we talking to the right DC?
- For the most sensitive authentication requests (NETLOGON SamLogon) RPC is encrypted and authenticated
- Winbind does lots of other calls, many over SMB
  - SMB signing now requrired when talking to a domain controller
  - All AD controllers offer signing
  - Old Samba domains might require "server signing = auto"
- New idmap_script
  - Flexible idmap backend for special configurations
  - Shell script called for idmap requests
  - 4.3 idmap_script is sequential, parallel version available now

SMB 3.1.1

- New dialect introduced with Windows 10
- Improved security
- Both Samba client as well as Samba server support 3.1.1
- Much improved secure negotiation
  - Before 3.1.1, downgrade attacks for protocol features were possible
  - Complete protocol exchange until after successful authentication now checksummed and signed
- New encryption algorithm: AES-GCM-128
  - **MUCH** too slow in software only, so only Samba's second choice
  - Very fast performance with CPU support, Samba needs to support this

smb encryption

- With SMB3 (Windows 8+) transport encryption is standard
- Session key ultimately based on user password
  - With Kerberos and NTLM a lot of key generation magic takes place
- SMB encryption is totally controlled by the server
- Samba parameter: "smb encrypt"
  - "smb encrypt = off": No encryption
  - "smb encrypt = desired": Encryption enforced for SMB3 clients, SMB2 clients allowed unencrypted
  - "smb encrypt = mandatory": Only encryption-capable SMB3 clients allowed
- Two levels of encryption
  - Per Session: Everything encrypted after user login, "smb encrypt" in [global]
  - Per Share: "smb encrypt" in share definition

Active Directory DC

- Improved trusted domain support
  - Inbound and outbound trusts work
  - Transitive trusts works for Kerberos, but not for NTLM
- KCC much improved
  - Samba DC replicates with all other DCs
  - This does not scale in larger networks
  - Microsoft DCs replicate "sparsely"
  - Samba now has an experimental implementation of the MS algorithms for the replica graph
  - Disabled by default

CTDB changes (4.4)

- Parallel recovery: Avoid deadlocks between smbd and ctdbd
- ctdb volatile databases in tmpfs
  - Usually TDB files live in /var on rotating rust
  - locking.tdb and other TDB files can see a LOT of churn, flushd can saturate slow local disks
  - Samba restart recreates them
  - ctdb now can create a tmpfs for volatile tdbs
- Continuous work to separate out tasks from main single-threaded ctdbd

Questions?

vl@samba.org / vl@sernet.de

http://www.sambaxp.org/