Samba 4 AD + Fileserver

SerNet

Linuxtag 2013

Volker Lendecke
SerNet

Samba Team

Volker Lendecke

SerNet

- 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
- 20 years concerned with Free Software
- First patches to Samba in 1994

SerNet

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

SerNet

- technological leadership of SerNet worldwide
  - involved in almost every big European Samba project
  - Half of the 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?

SerNet

- 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 3.6

SerNet

- Last release before 4.0
- SMB2.0 support (except durable file handles)
- Tighten security defaults (client uses ntlmv2 by default)
- Printer subsystem overhauled
- Internal use of RPC interfaces
- Winbind idmap configuration changed (again :-()
- All goodies are part of Samba 4.0

Samba 4

SerNet

- Finally released on December 11, 2012
  - We missed the really cool date 12.12.12...
- 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 new feature: Active Directory Domain controller
- Samba 4.0 merges Samba 3 and Samba 4 code bases

Active Directory Domain Controller

SerNet

- The main services beyond NT4: LDAP, Kerberos and DNS
- Kerberos is almost standard, Samba4 ships heimdal
  - Work is planned to run with a stock MIT KDC
- LDAP has been extended extensively for AD
  - DRSUAPI for multi-master replication
  - Many improvements for scalability and consistency
- DNS: Both internal and bind plugin
  - Not invented here was necessary because many team members could not get the bind config right...

AD: Dirty little secrets

SerNet

- „officially“ supported: Small business server
  - 1 domain, 1 domain controller
- Trusts:
  - Samba can be trusted, but Samba can not trust (is that a bad thing? :-))
- Replication: (multiple Dcs)
  - Directory replication mostly works
  - Sysvol replication not Windows compatible
  - Rsync based manual replication possible (multiple Samba Domain Controllers)

AD: Next steps

SerNet

- Fix the limitations (trusts, replication)
- Sysvol (file system replication):
  - Async rpc server infrastructure, also for witness protocol
  - Lots of NTFS semantics laid down in the protocol
- Trusts:
  - Winbind from source3 needs to be extended to take over from source4
  - Extensions for AD style intra- and inter-forest trusts

SMB2

SerNet

- New protocol introduced with Windows Vista
- Basic SMB2 is semantically very similar to SMB1:
  CreateFile is almost the same
- New protocol features added only to SMB2
- Samba 3.6 supports SMB2.0 excluding durable file handles
  - Samba 4.0 includes durable file handles
  - SMB3.0 negotiated

SMB server TODOs

SerNet

- SMB 2.1:
  - leases
- SMB 3.0:
  - directory leases
  - multi channel
  - RDMA
  - cluster concepts: (scale-out/continuous availability)
  - persistent handles
  - witness protocol
- https://wiki.samba.org/index.php/Samba3/SMB3

SMB: Leases and Directory Leases

SerNet

- Leases ⇒ oplocks done right (content caching)
- Directory Leases ⇒ change notify done right (metadata caching)
  - extend oplocks to cope with SMB oplocks and leases
  - remove the 1:1 relation between open and oplock (locking.tdb)
- add support for oplock keys (empty for SMB 1)
- cleanup / preparation work was already started by Volker
- lease keys and client guid need to be maintained at the SMB layer

SMB: Multi Channel

SerNet

- bind multiple transport connections to one SMB session
- interface discovery:
  - new fsctl (FSCTL QUERY NETWORK INTERFACE INFO)
  - Server can tell the client what interfaces it provides
  - extend current 1:1 relation smbd ↔ TCP connection
- transfer TCP-socket to smbd serving connections with the same ClientGUID in negprot (fd-passing)
  - session bind automatically on correct smbd
  - only one process has the file open for multi-channel sessions
  - we only need to do book-keeping on the SMB level
- (replay/retry counters, channel sequence numbers, ...)
  - the posix/file system level won't notice multi channel

SMB: RDMA

SerNet

- RDMA uses (infiniband, iWarp or RoCE hardware, or software emulation)
- transport abstraction needed (TCP/NBT vs. RDMA)
- buffer abstraction needed in order to do zero-copy transfers
  - SMB VFS READ/WRITE BUFFER SEND/RECV
- Problems with the current libibverbs/librdmacm libraries
  - it's not fork() safe, which is currently required by smbd
  - "FD-passing" is not supported, would be needed for the current planned multi channel design

SMB: Persistent Handles

SerNet

- persistent handles: durable handles with strong guarantees
- server application workload
- need to make some DBs persistent (or by record) (⇒ changes to tdb / ctdb ...)
  - smbXsrv open global
  - locking, brlock
- need new index databases for smbXsrv open

SMB: Witness

SerNet

- List server network interfaces and interface changes
- The witness protocol is implemented as DCERPC service.
  - using ncacn ip tcp as transport
  - heart beat link between a SMB 3.0 client and server cluster.
  - it provides faster planed or unplaned failover
- needs async dcerpc server infrastructure
  - which can be used independently from smbd/samba

Kontakt

SerNet

Volker Lendecke, vl@sernet.de

SerNet GmbH

Bahnhofsallee 1b
37081 Göttingen

Schützenstr. 18
10117 Berlin

tel +49 551 370000-0
fax +49 551 370000-9

+49 30 5 779 779 0
+49 30 5 779 779 9

http://www.sernet.de