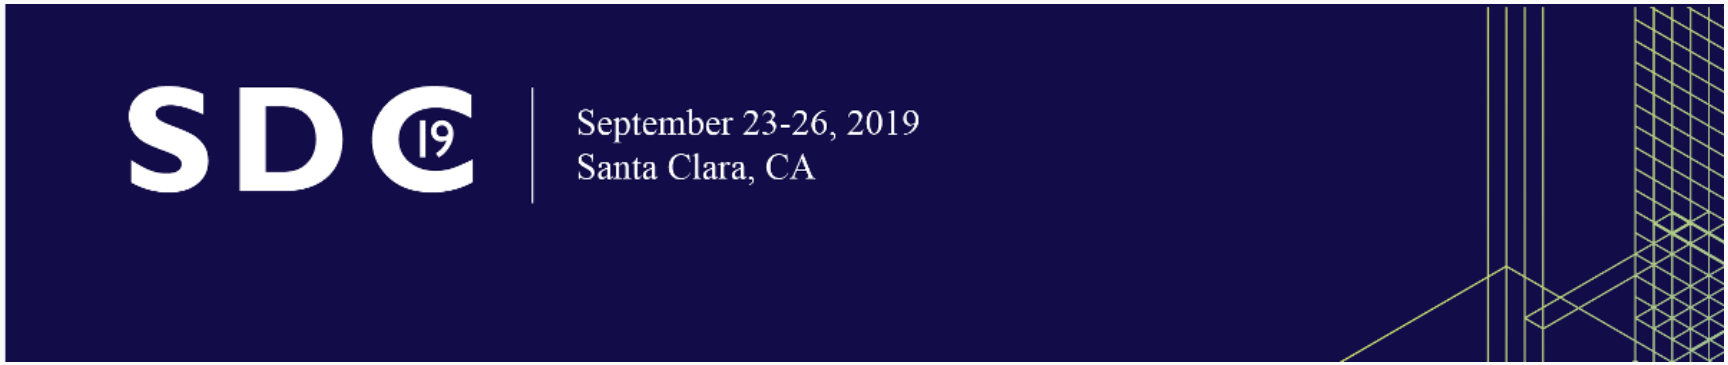

What's new in Samba

Ralph Böhme, Samba Team, SerNet

2019-09-23

Outline

SerNet

Tooling

CI on gitlab

- adopted gitlab for CI about a year ago
  `https://gitlab.com/samba-team/devel/samba`
- used by the Samba team and open to external contributors
- already at 800 merge requests and 4000 CI runs (gitlab lingo: pipelines)
- used for reviewing and testing new code, not for the final merge
- most changes going through gitlab these days
  - sending in patches via mailing list stil welcome for drive-by contributors!
- CI runtime down from hours to to one by adding (more) parallelisation (thanks metze!)

What's new in Infrastructure Code

Move to using GnuTLS with Samba 4.11 for (nearly) all crypto

- AES: AES-CCM, AES-GCM, AES-CFB, AES-CMAC
- MD5, HMAC-MD5
- HMAC-SHA-256 and 512
- RC4 (GnuTLS 3.4.7)
- Left: DES and MD4
- Samba AD DC still uses Heimdal's own crypto

Two RPC server implementations, two client implementations

- duplicate protocol frontends
- duplicate RPC service implementations
- not all are created equal:
  - RPC server used by the fileserver doesn't support async RPC requests
  - needed for Witness, MS-PAR

Plan

- take one RPC frontend (source4), drop the other (source3)
- keep both independent RPC service implementations
- RPC frontend code calls into one of the implementations
- large chunks of preparatory work and cleanups already upstream
- full merge code at the prototype stage, ready for 4.12?

What's new in the Fileserver

Removing old cruft in Samba 4.11

- SMB1 disabled by default, woohoo!
- (NTLMv1 disabled since 2016)
- Lanman authentication and cleartext password deprecated (and will be removed in 4.12)

Scalibility improvements, part 1

- problem: churn on gencache, Samba's internal generic caching subsystem
- internal cache architecture required stabilize run every X modifications
- stabilize was expensive for large, possibly contented caches
- deadlocks reported from highly loaded systems
- solution:
  - add hash-chain loop detection to TDB
  - checksum gencache records
- no stabilize anymore

Scalibility improvements, part 2

- contention on single files or directory being opened by many clients
  - can cause serious performance issues
  - seen case where time to open directory was 2 seconds
- performance improved by a factor X
- see Volker's talk for the gory details and the value of X

What's new in the fileserver, cont.

asynchronous path-based VFS functions:

- Samba has had async handle-based IO functions in the VFS for ages
- path-based VFS functions are a bit more tricky: impersonation
  - by the time the async operation is scheduled the process may have switched to a different user
  - pass user token to worker thread, thread performs impersonation
  - requires kernel support for per-thread credentials
- first async path-based VFS functions in Samba:
  - `SMB_VFS_GET_DOS_ATTRIBUTES_SEND/RECV()`
  - used to fetch metadata from filesystem xattrs
- IO operations are scheduled in a threadpool
- decent improvement enumerating large directories eg on FUSE-based fs
- more work underway, see metze's presentation

SerNet

Simplified support for trusts as domain member

- old code needed to enumerating trusts (recursive!) in winbindd
  - this is broken by design
- now domains are added to internal domain list when users authenticate
- new option `"winbind scan trusted domains = yes | no"`
  - default still `"yes"`
- `wbinfo -m --verbose` out changed to reflect the dynamic view

wbinfo -m –verbose example

wbinfo -m –verbose, after starting winbindd

```
$ wbinfo -m --verbose
Domain Name DNS Domain          Trust Type  Transitive  In   Out
BUILTIN                         Local
TITAN                           Local
WDOM2       wdom2.site          Workstation Yes         No   Yes
```

wbinfo -m –verbose example

wbinfo -m –verbose, after users from trusted domains authenticated

```
$ wbinfo -m --verbose
Domain Name DNS Domain           Trust Type  Transitive  In   Out
BUILTIN                          Local
TITAN                            Local
WDOM2       wdom2.site           Workstation Yes         No   Yes
WDOM1       wdom1.site           Routed (via WDOM2)
WDOM3       wdom3.site           Routed (via WDOM2)
SUBDOM21    subdom21.wdom2.site  Routed (via WDOM2)
SDOM1       sdom1.site           Routed (via WDOM2)
SUBDOM11    subdom11.wdom1.site  Routed (via WDOM2)
```

Upcoming

- prefer AES-GCM over AES-CCM with GnuTLS with 4.12
- new impersonation model, see metze's talk
- SMB3 Multichannel, last stages, hopefully with 4.12?
- SMB-Direct and Persistent Handles:
  - working prototypes, still require serious effort to finalize
- Witness: mostly ready, blocked on new DCE-RPC infrastructure
- SMB2 POSIX Extensions
  - consensus on the spec
  - slowly making progress on client and server
  - getting delayed by other work
  - 2020 should be the year of SMB2 POSIX Extensions!

What's missing

SMB3 Directory Leases:

- decent reduction in the number of SMB requests for metadata heavy workloads
- a must have, but currently no-one is working on it

SMB3 compression

- should be low hanging fruit, any takers?

What's new in Clustered Samba

New configuration style

- configuration changed in 4.9
- new ini-style ctdb.conf config file
- requires manual config migration, a script can help:
  - `ctdb/doc/examples/config_migrate.sh`

Example

```
[logging]
    log level = NOTICE
[cluster]
    recovery lock = /cluster/reclock
```

What's new in Samba AD

Performance and Scalability

Scaling to 100k users and more

- historically RPC services, LDAP server and the KDC ran as single process
- with 4.11 those three now default to a pre-fork process model
- replication improvements, especially around linked attributes
- LDAP database backend (ldb) uses new GUID based indexing
- new experimental LMDB LDB backend
- New LDB <= and >= index mode to improve replication performance
- now ready for 100k users and larger Samba AD domains

Enhanced support for trusts

- trusts are supported in both directions for Kerberos and NTLM authentication
- users/groups of a trusted domain can be added into domain groups
  - group memberships are now expanded on trust boundaries
- no SID filtering, no support for Selective Authentication
  - both sides of the trust need to fully trust each other!
- Samba can still only operate in a forest with just one single domain

The End

Q&A

Thank you!

Questions?

Ralph Böhme <slow@samba.org>

SerNet -> Sponsorbooth