Unforking Samba4: The Success!

Presented by Andrew Bartlett of Catalyst // 2015-05-21

catalyst

open source technologists

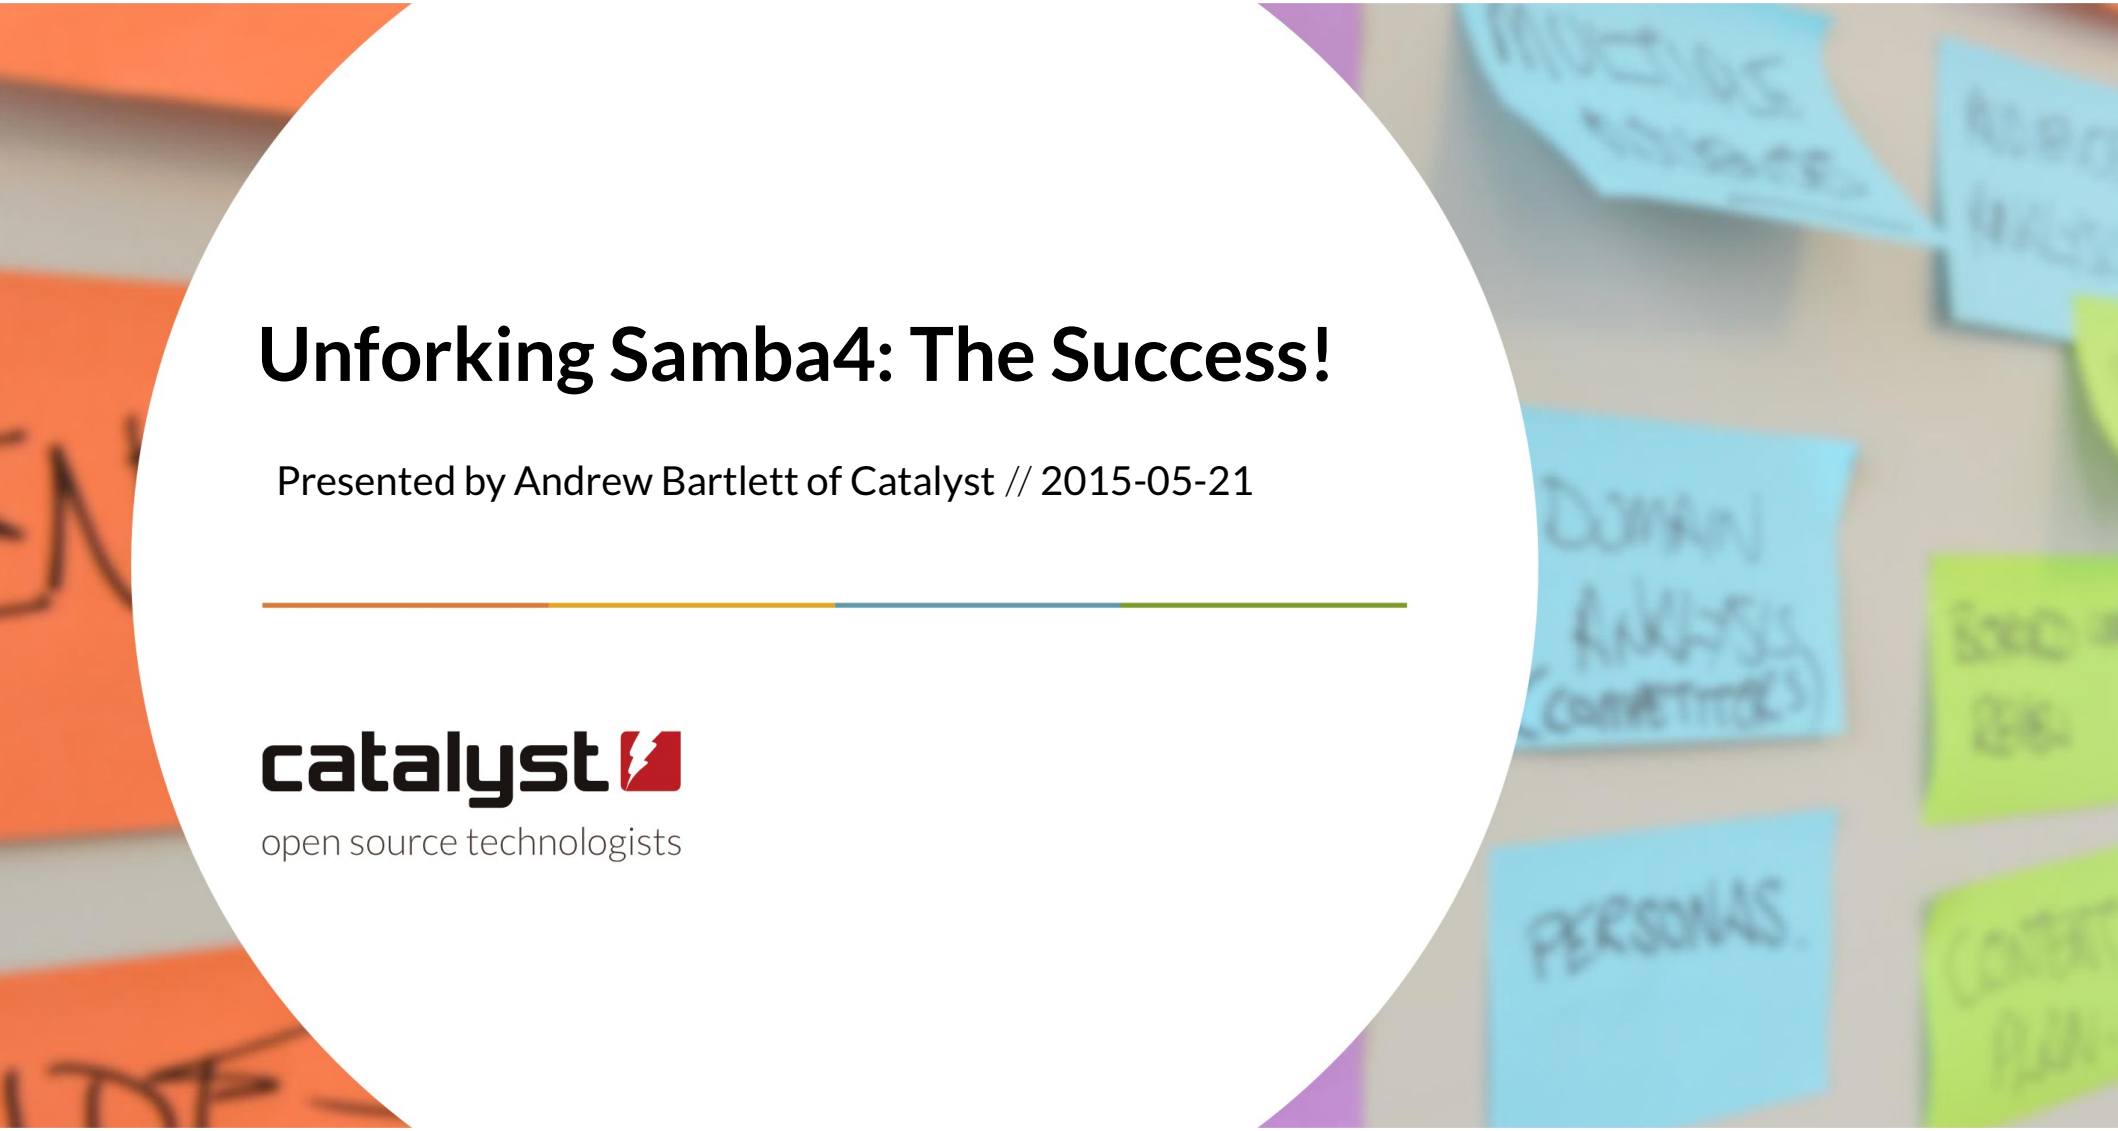

Andrew Bartlett

- Samba Team member for 14 years
- Key developer on the Samba AD DC component
- Based in Wellington NZ
- Thank you to:
  - My employer, Catalyst for their great support
  - Tranquil IT for funding my travel to Europe

The great success

- We released Samba 4.0
  - I wish I had been here for the party!
  - It took time, but we didn't loose sight of the goal
- In doing so, we reunited as a Team
  - Stronger together!
- Taking on new challenges like SMB3 and inter-forest trust

Our roller-coaster ride

- Samba forked
  - We didn't like to say it, but that is the reality
  - Both a social and a technical fork
- Many, many team members worked really hard to undo the damage
  - I will speak mostly about the areas I was involved in
  - Much great work many others
- With Samba 4.0, we finally merged again

How did we get to 4.0? – a timeline

- Technical and social steps
- Merge team motto:
  - “Solving social problems with technical solutions since...”

2004	2008	2010	2011	2012
Samba4 Development starts	Franky proposal Combined GIT tree IDL files merged Named pipe forwarding	waf introduced s3compat net4 » samba-tool	Combined build Single make test 'as is' release s3fs proposed	GENSEC s3fs done 4.0 released!

Beyond 4.0, merge work to 4.2 and beyond

- A decade later, and we still have work to do
  - Will we ever get beyond source3/source4?

2004	2013	2014	2015	2016
Samba4 Development starts	Autoconf removed 4.1 released	Winbindd merge	4.2 released Datagram messaging	What next?

Unlocking possibilities

- Each merge step enables another
- Named pipe forwarding showed this was possible
- Merging the tree stopped version skew
- Merging the IDL avoided pointless diversion
- Merging the build systems enabled a merged test
- Merging loadparm wrappers enabled sharing of more complex code
- Passdb and auth modules provided the glue
- Merging GENSEC enabled merging schannel fully
- Merging winbindd enabled inter-forest trusts

Not the only way it could have been done

- I'm not interested in re-arguing the past
  - But I do have some apologies for my tone and behaviour at points
- I am interested in explaining why we did what we did
- Samba continues to evolve

Named pipe forwarding

- The first and longest-lasting part of the Franky effort
- Allows ncacn_np connections to be answered by the AD DC

Using common IDL and PIDL

- We had two divergent sets of IDL
  - Merged
- We had hand-generated NDR
  - Replaced
- We had different copies of pidl
  - Merged

Authentication

- The most sensitive area of the merge
  - A key part of the original s3compat effort
  - Perhaps single-handedly derailed that merge
- Key requirement:
  - Consistent behaviour
- Key implementation pattern
  - Code merge where possible
  - Plugin-based code replacement otherwise

Common IDL and structures in auth

- Authentication
  - auth_usersupplied_info made common
  - auth4_context made available in common
- Authorization
  - auth_session_info made in common
  - Replaced netr_SamInfo3 in named_pipe_auth.idl
  - Replaced auth_serversupplied_info with auth_session_info (slowly)

NTLMSSP merge

- We had:
  - two NTLMSSP clients
  - two NTLMSSP servers
- We merged the NTLMSSP servers into libcli/auth
- And moved the source4 NTLMSSP client into libcli/auth
- A GENSEC module was built around the new common code

auth_generic – the Trojan horse

- A very poor disguise for GENSEC
- Initially only the rpc_server code
  - Nominally wrapping the NTLMSSP gensec module
  - But written such that it could wrap anything
- Also unified the code in the SMB / SMB2 servers

GENSEC

- GENSEC was merged into common
- Replaced the similar gse layer in the source3 RPC server
  - gse_krb5 became a gensec module
- Removed duplication of code in the SMB / SMB2 file server
- Created a common abstraction
  - over the remaining existing source3 code
  - Able to be replaced by plugin from the source4 code

Full GSSAPI for SMB

- The big 'not incremental' step was to
  - Remove the fake GSSAPI server from source3
  - Replace it with one using gse_krb5
- This is what increased the MIT krb5 minimum to 1.8

auth_samba4

- Much more than a normal auth module
  - Simply loading auth_samba4 causes hook functions to run
  - Forces AD DC mode on the rest of the auth/GENSEC subsystems
- Totally overrides all the GENSEC plugins
  - Allows a difference, forced set of modules to run
- Local group handling and idmap lookup forced via AD DC codepaths
- The 'normal' NTLM functions are only called from winbindd
  - For local user authentication on a RW DC

Regarding auth_netlogond?

- I'm not proud of my behaviour in removing that code
- Moving the NTLM auth to an IPC mechanism may still be possible

PASSDB

- Important so that existing tools keep working
  - smbpasswd
  - net
  - pdbedit
- Also used in winbindd and in smbd
  - Very helpful hook for idmap override
- An important access method for upgrades
  - Samba-tool domain classicupgrade

pdb_samba_dsdb

- Built for the needs of classicupgrade first
  - Offline access was required
    - no DC until provision finished
  - Uses the LDB API (helper functions)
  - Based on pdb_ads by Volker
- Idmap hooks read the local idmap.ldb used in the AD DC
- Get/Set trusted domain credentials

Regarding pdb_ads?

- I'm not proud of my behaviour in removing that code
- pdb_samba_dsdb can use ldapi:// URLs if desired, once the server is running

Build systems

- The combined waf build has been critical
- Removing autoconf was even more important in the long term
  - No more hand-crafted object lists

Testing

- Combined make test
- Tests AD domain member against our AD DC for example
- All run from selftest.pl in selftest/
- Glued together rather than integrated
  - Done early in the process to reduce breakage and improve tests

Test code in smbtorture{3,4}

- Even at the darkest points of the split, tests written in smbtorture4
- The 'merged build' was for building smbtorture4
- But many simple tests still added to smbtorture3
- Blackbox test scripts scattered over the codebase

Test environments

- selftest/target/Samba.pm is the glue
  - selftest/target/Samba3.pm
  - selftest/target/Samba4.pm
- Left over from when we had to be able to test autoconf alone
- Michael Adam did a long over-due rename in 2015

Messaging

- We now use a common datagram-based messaging bus
  - Thanks to Volker Lendecke
- Initial use is for smbcontrol to obtain a talloc report

File server

- File server started from inside samba with exec()
- Python bindings added to the VFS
  - Allows provision to write ACLs to disk
- Unfortunate name of s3fs
  - This happens if you don't check for name conflicts first...

Loadparm

- lib/param imported from source4
- loadparm_init_s3() hook allows using a 'source3' loadparm
- Parameter table merged
  - Initially with #include of a C file!
  - Now properly shared as a normal C file
- Parameter list now autogenerated from XML docs

Winbindd

- With Samba 4.2 we now use the source3 winbindd
- Main task was adding an IRPC listener and forwarder
- May have been possible for 4.0 in hindsight
- Key task for inter-forest trusts
  - But not enough on it's own, but metze doing great work

Netlogon SCHANNEL

- Merged and AES support added
  - Great to have that enabled in both servers at once
- Potential for further merging of NETLOGN servers
- Now a common GENSEC module
  - Removing a layer of wrapping

RPC Binding Handles

- Allows implementation-agnostic RPC clients
  - Even in python!
- Enabled the AES SCHANNEL work to be in common

Still TODO

- NTLMSSP client code
- GSSAPI client and server code
- Loadparm code
  - Registry loadparm in particular
- Smbclient4
- Command-line syntax differences

TODO: Test plans remain mostly separate

- Source3/selftest/tests.py
- Source4/selftest/tests.py
- Selftest/tests.py
- Some cross-over of tests vs environments
  - Tests in source3 run against ad_dc environment

TODO: Remove internal winbind

- We do not need two winbind implementations
- We should remove source4/winbind
  - Once last compatibility issues are fixed
  - Just need to force sync of secrets.tdb on startup

What about the NTVFS file server?

- Kill it
  - Revenge? It was what started this war!
  - Reduce nominal security exposure for vendors?
- Keep it (behind a ./configure option)?
  - Still only protocol level CIFS / SMB1 proxy
- Still a good working model for a NTVFS layer
  - What our competitors at likewise, as I understand it
  - Avoids matching client / server bugs between smbtorture / smbd

Structural Reform

- Continue to de-emphasise source3 / source4
  - Perhaps we should rename some of these parts?
  - I found a mail recently when I argued against that...
- RPC server handlers
  - It would be great if the parse and handler interface was shared
- Continue to find common code and merge it

One team / Branding

- Could we move beyond Samba 4.x as a (confusing) version number?
  - We may need some better brands
  - Unfortunate to discard **samba4** as a brand, as it is still strong
  - Samba AD DC just doesn't resonate in the same way
  - Make the next release Samba 5.x?
- Continue to avoid referring to and thinking of team members as 'samba3' / 'samba4' developers?

Avoiding a repeat in the future

- We forked twice already, and that hurt
  - Samba TNG
  - Samba4
- Avoid long-running feature branches?
  - 'Not required' by git
  - But may be required to keep the team a team
- Deliberately take an interest outside our own areas?
- Recognise and celebrate our diverse users and features!

Conclusion

- We did it!
  - We focussed on the task,
  - united on the goal and worked as a team
- We have much still to do,
  - but having come this far
  - what remains is entirely practical
- Would you like to help?

Catalyst: Using, building and supporting Samba and Beyond

- Worldwide Offices in Wellington NZ, Brighton UK and Australia
- Samba Support and Development
- Samba and Windows integration