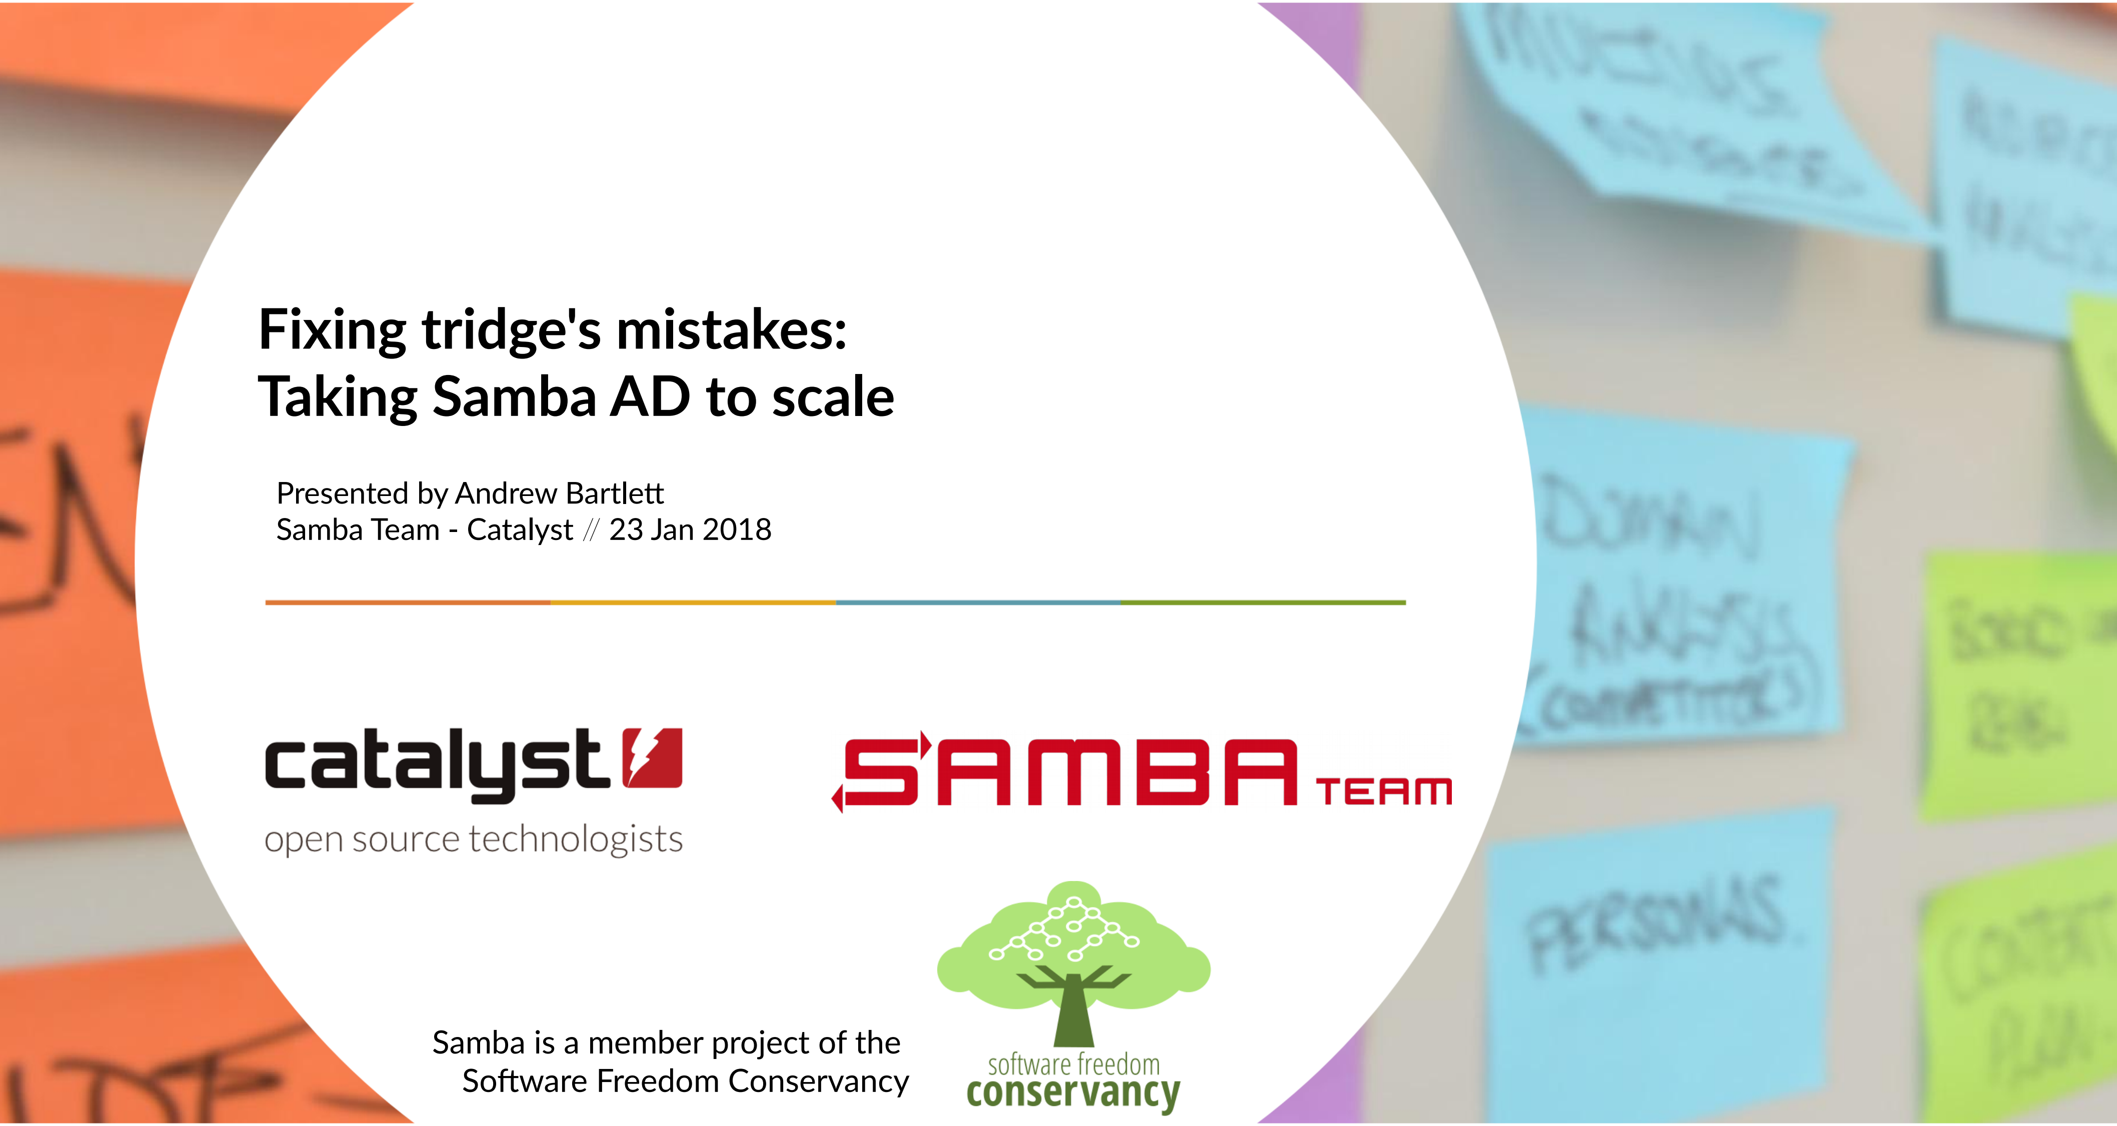

Fixing tridge's mistakes: Taking Samba AD to scale

Presented by Andrew Bartlett
Samba Team - Catalyst // 23 Jan 2018

catalyst
open source technologists

SAMBA TEAM

Samba is a member project of the
Software Freedom Conservancy

software freedom
conservancy

Andrew Bartlett

- Samba Team member since 2001
  - Focus on our Active Directory Domain Controller
- Team lead for the Catalyst Samba Team
  - Based in Wellington, New Zealand
- These views are however mine alone

Samba user scale as an AD DC

- To be clear, in our wildest dreams we never expected Samba's AD to be deployed so large
- In a two-hour benchmark adding users and adding to up to four groups:
  - Samba 4.4: 26,000 users
  - Samba 4.5: 48,000 users
  - Samba 4.6: 55,000 users
  - Samba 4.7: 85,000 users
  - Samba 4.8: 100,000 users

Improved internal LDB format in Samba 4.8

- Old format:
  - TDB key is the DN, casefolded
    - CN=ANDREW,OU=USERS,DC=EXAMPLE,DC=COM
  - But index values are original case!
    - dn: @INDEX:attr:value
    - @IDX: CN=Andrew,OU=Users,DC=example,DC=com
- New format:
  - Using a GUID as the key in the TDB and index record

(not) Taking out full-DB read locks

- Our big ‘ouch’ moment running up to Samba 4.7:
  - We forgot to take out DB-wide locks during an LDB search
- Also had a massive performance cost for a full DB scan
  - As per-record locks were taken out instead
- Regression introduced in 2009:
  - s4-ldb: fixed nested searches inside ldb modules
- More than a a performance issue:
  - Risk of inconsistent DB reads, Replication failure

Supporting more connections on each DC

- Samba 4.6 removes single-process restrictions on NETLOGON
  - Really important for 802.1x backed authentication
- Samba 4.7 will support a multi-process LDAP server
  - Previously all LDAP clients used a single process
  - Actually reduces number of connections you can fit in memory (oops)
- Samba 4.8 adds a prefork mode
  - Great for a big AD DC with many, many clients

The future for scale

- Remove other O(n) and O(n^2) operations
  - Yes, we still have them elsewhere
- Better index handling
  - Our current index code is still very much a first pass
- LMDB from Symas (OpenLDAP)
  - Current work in progress to use LMDB rather than TDB
  - Primary focus is to remove the 4GB limit

Performance tool

- Last year I spoke of plans for a perf tool
- We can now record and re-play traffic
  - Recreate a real-world load
  - Amplify the traffic

Performance graphs
- April 2016 to Dec 2017

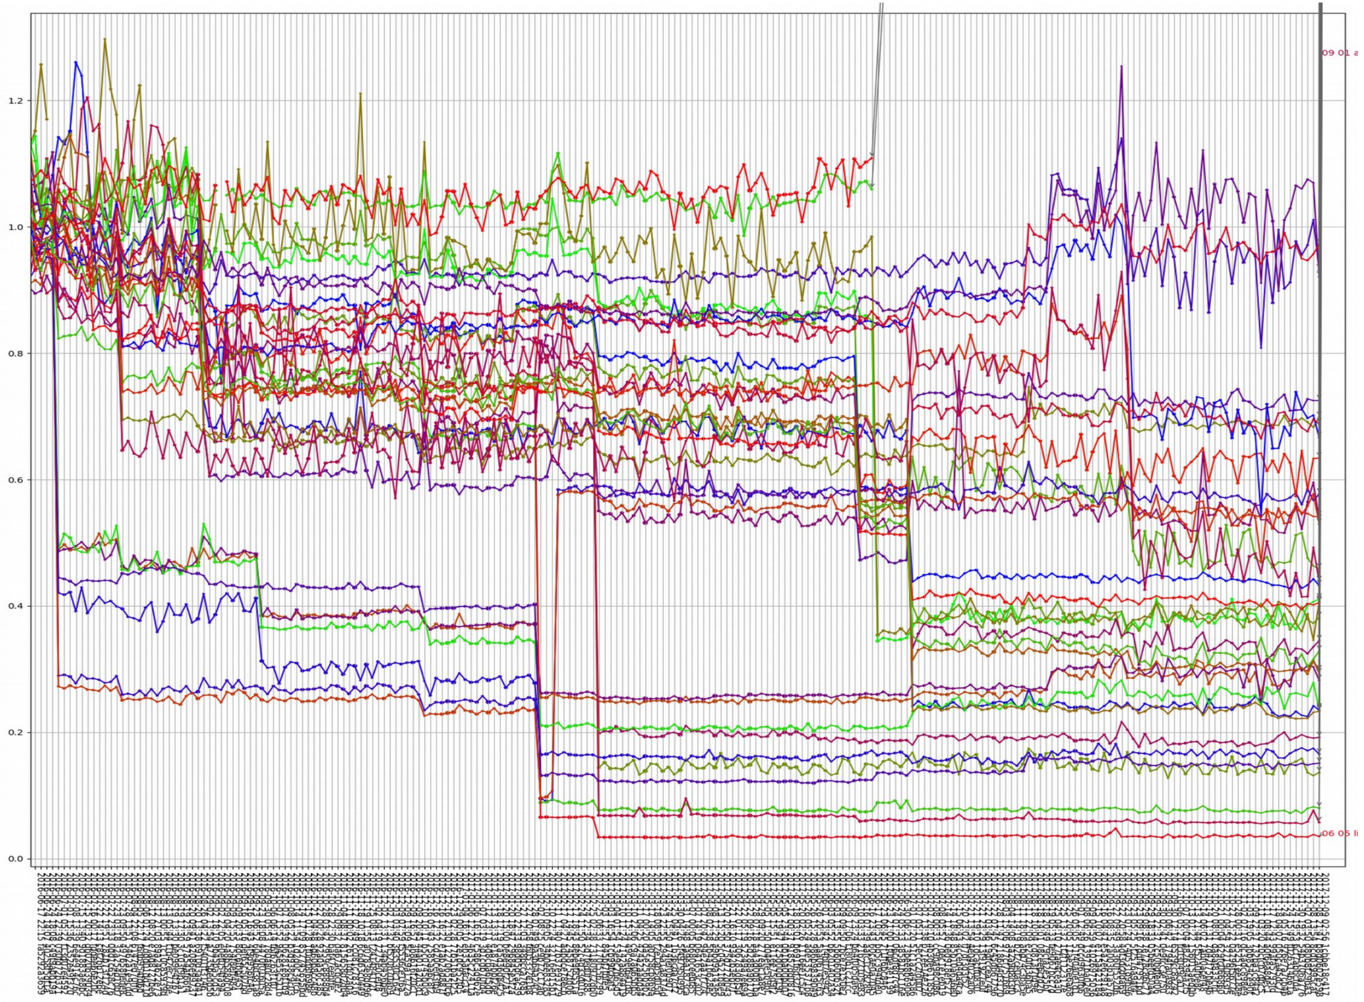

Beyond performance

- Encrypted secrets (4.8)
  - Use a local file key to encrypt secret attributes (could then be network-deployed)
- Unix-compatible passwords (4.7)
  - Store and retrieve passwords to sync with other systems
- Audit Logging (4.7)
  - Output audit logs into JSON
- RODC support (4.7)
  - This was experimental until now

Enterprise Reliability

- Samba AD DC is now the backbone for large companies
  - including for Network Access Control
  - Surprisingly reliable!
- Multi-master configuration removes single points of failure
  - FSMO roles are single, but can be transferred or stolen
- Clearly ready for the enterprise:
  - HP-Enterprise engineers enquiring about putting Samba AD on HP-UX (on IA-64)!

MIT Kerberos

- Blocking Samba being a part of SLES and RHEL
- First release with Samba 4.7 but still in progress
- Very important as Heimdal Upstream only just restarted releases
- I'm hoping to update Heimdal as well
  - 5 year old security code is not a great thing

WARNING: Samba 4.7 AD DC upgrades

- A serious data loss issue has been found by Stefan Metzmacher
- Impact is to Samba 4.6 and earlier installs the upgrade to 4.7
- A new release of Samba 4.7 is expected very soon
- A new ‘provision’ or ‘join’ is safe, even to an old domain

Become an

OFFICIAL CONSERVANCY SUPPORTER!

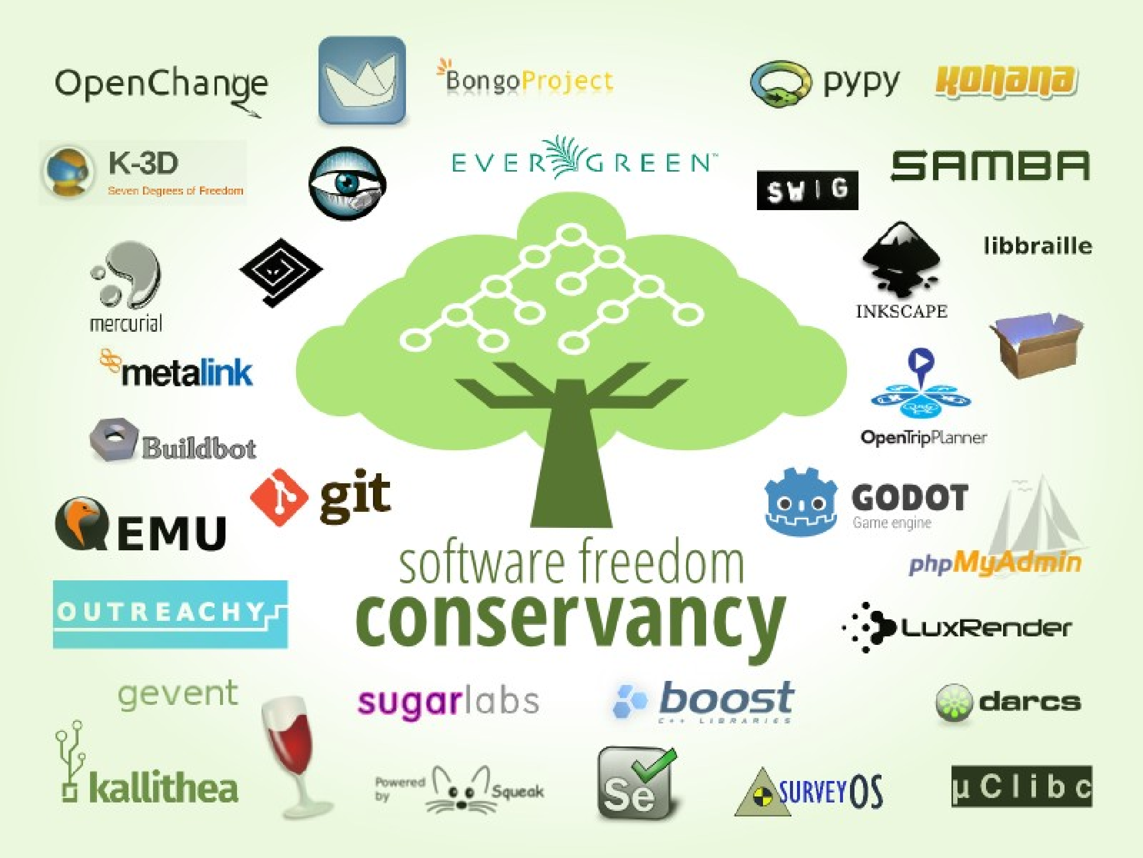

Catalyst's Open Source Technologies – Questions?

Linux
debian
perl
php
Git
PostgreSQL
Lucene
Drupal
WordPress
moodle
totara
mahara
koha
kannel
openstack
CLOUD SOFTWARE
ceph
SAMBA

Want to join Catalyst? - talk to me in the hallway track! - Sysadmin positions open in Wellington!