ReOpenLDAP

– through the thorns to the stars

Leonid Yuriev
LDAPCON-2017

ReOpenLDAP – Nowadays in Russia

6 districts, 24×7

Workload: W10K, R50K

Replication: 4×4 Full-mesh

Size: 100M DNs, 100Gb

Falling into a Black hole…

2013

Company "XYZ" implementing NGDR:
- UDR, 3GPP 23.335
- the Central repository for a "services"

OpenLDAP chosen:
- the Performance (LMDB)
- the Multi-master replication

Troubles and facepalm:
- DB corruptions, slapd segfaults
- No SLA guarantee from Symas Corp

2014

Sunrise manually

2014 Q4

Couple of LMDB bugs

«Dreamcatcher» and LIFO

~5000 warnings, etc…

2015

ReOpenLDAP

Eliminated the most of segfaults

+ memcheker
+ iddqd, idkfa
+ backtrace

+ biglock
+ fast refresh
+ etc...

2016

Multi-master replication

Tests less-flaking

+ syncprov-showstatus, etc…

+ new configure

+ server-side keepalive, etc…

Two years

Multi-master Replication Issues

- Happened in production in 2013…2015
- Reproduced and **fixed** in ReOpenLDAP at 2016-Q1
- Seems present in OpenLDAP

Two categories:

1. Erase of a replication scope
2. Loss of individual changes

Replication Issue: Erase a scope

- Provider **may** generate 'present phase'
  without present-UUID list:
  - is possible that… no TAG_SYNC_ID_SET before TAG_SYNC_REFRESH_PRESENT
  - consumer **will** delete all entries in scope

- Provider **may** generate stripped present-list:
  - wrong error handling in appropriate callbacks
  - leads a list of present-UUIDs to be partial
  - consumer **will** delete missed entries

Replication Issue: Lost a changes

- Provider could generate only partial list-of-changes:
  - the «baseline» CSN should be chosen
  - in case of multi-master cluster this is tricky
  - a recent from one server could hide updates from another(s)
- Racing between updates and the mapping GUID to DNs:
  - replication works with GUIDs, but DIT uses DNs as a keys
  - GUIDs from update-notifications should mapped to DNs
  - but DN could be removed and re-created between such mapping and applying an update.
  - consumer may delete DNs which added recently.
- More over there a set of related bugs:
  - notify-of-modify could be applied to DN which is updated since comparison of the sync-cookies.
  - notify-of-add could be applied after a new version of particular DN was created and deleted, and old (removed before) version of DN will be "revived".
  - notify-of-delete of old DNs could remove a recent one.

TLDR; Some meditation is needed…

ReOpenLDAP – **the reasons to fork**

a lot of **Changes**	• Removal of 5000 warnings, and so on…
Decision freedom	• LIFO for Write-Back cache, and so on…
still **Open source**	• Returning an improvements • Code review and testing

Check out the difference

ReOpenLDAP

Only Linux

–Wall –Werror

444 issues from Coverity
(**with** Contrib/*)

No flaky tests
(except known issues)

Reliable multi-master replication

OpenLDAP

Linux, FreeBSD, Windows, etc…

≈5000 warnings

≈1800 issues from Coverity
(**without** Contrib/*)

≈42 iterations for tests failure
(just try to loop the tests)

Replication may (?) loose changes
and chance (?) erase the scope

ReOpenLDAP – **Objectives and plans**

libmdbx	• Release a new MDBX (incompatible) • New storage backed for it
Perfectionism	• Refactoring, a lot of… • Fix all Coverity™ issues
Packaging	• For all common Linux distros • Especially: AltLinux and E2K “Elbrus”
Compete	• Comparative tests of replication • Performance benchmarking

Thanks you all, especially to Howard Chu !

Leonid Yuriev <leo@yuriev.ru>,
lead of ReOpenLDAP

Since September 2016:
Positive Technologies,
Advanced Research, Lead «Shaman»

Formerly:
Perter-Service R&D, Infowatch, Cronyx

---

https://github.com/leo-yuriev/libmdbx
https://github.com/leo-yuriev/ReOpenLDAP

LDAPCON-2017