

# What's New in OpenLDAP

Howard Chu

CTO, Symas Corp. [hyc@symas.com](mailto:hyc@symas.com)

Chief Architect, OpenLDAP [hyc@openldap.org](mailto:hyc@openldap.org)

2017-10-19



# OpenLDAP Project

- Open source code project
- Founded 1998
- Three core team members
- A dozen or so contributors
- Feature releases every 12-18 months
- Maintenance releases as needed

**IMDB**



# A Word About Symas

- Founded 1999
- Founders from Enterprise Software world
  - *platinum* Technology (Locus Computing)
  - IBM
- Howard joined OpenLDAP in 1999
  - One of the Core Team members
  - Appointed Chief Architect January 2007
- No debt, no VC investments: self-funded



# Intro

- Howard Chu
  - Founder and CTO Symas Corp.
  - Developing Free/Open Source software since 1980s
    - GNU compiler toolchain, e.g. "gmake -j", etc.
    - Many other projects...
  - Worked for NASA/JPL, wrote software for Space Shuttle, etc.

**IMDB**



# Topics

- (1) Recent Releases
- (2) Features Previously in 2.5
- (3) New Features in 2.5
- (4) Work In Progress

IMDB

# (1) Recent Releases

- 2.4 Release Winding Down
  - Feature frozen, bugfix only
  - 3 releases in the past 2 years
    - Commit rate still fairly high
    - Not quite "release early, release often"
    - Scattered fixes, mostly in documentation

**IMDB**



## (2) Features in 2.5

- Multiple Threadpool Queues
- Streamlined Write Waiters
- Offline slapmodify/slapdelete
- LDAP Transactions in primary DBs

**IMDB**

# Features in 2.5


- Multiple Threadpool Queues
  - Significantly reduced lock contention on multi-processor servers
  - Not much visible impact on back-bdb/hdb
  - 25% throughput boost with back-mdb on quad-core server

LMDB



# Features in 2.5

- Multiple Threadpool Queues



# Features in 2.5

- Multiple Threadpool Queues
  - New tests July 2017 on Oracle M8 with 2048 VCPUs and 1.5TB RAM, Solaris 11.3
  - 16 threads, 180,000 searches/sec, 1M DNs, 100 clients and 10 connections each
  - 1024 threads, 64 threadqueues, 64 listeners, 930,000 searches/sec - at this point, profile showed bottleneck was inside the kernel
  - recommend no more than 16 threads per queue and listener

**IMDB**



# Features in 2.5

- Streamlined Write Waiters
  - Take responsibility for select() of blocked writers away from central listener thread
  - Allows higher throughput in the presence of slow clients interspersed with heavy users

**IMDB**

# Features in 2.5

- Offline tools - slapmodify/slapdelete
  - The obvious missing pieces to complement slapcat/slapadd
  - Essential for editing cn=config when slapd not running (or not able to run)

**IMDB**



# Features in 2.5

- LDAP Transactions
  - Completed for back-bdb, -hdb, and -mdb
  - Support in back-ldap exposes a need for a distributed txn story - 2-phase commit at least

IMDB

## (3) New For 2.5

- Syncrepl Lazy Commit
- Non-blocking TLS Handshake
- Non-blocking SASL Interactive Bind
- SASL Channel Binding support for OpenSSL, GnuTLS
- Elliptic Curve support for OpenSSL
- kqueue support for FreeBSD / MacOSX

**IMDB**



# New For 2.5

- New backends
  - WiredTiger, asyncmeta
- New server
  - LDAP load balancer
- New modules
  - RFC6238 TOTP
  - RFC3829 Authzid
  - vc (Verify Credentials)
  - adremap
  - usn
  - autoca

# IMDB

# New For 2.5

- 64 bit Index Hashes
- LDIF parsing API in libldap
- Disable Flag for DBs and overlays
- High resolution operation timestamps
- Store TLS certs/keys directly in cn=config
- Large multival attribute optimization for back-mdb

The IMDB logo, with the letters "IMDB" in a bold, white, sans-serif font, and a large, stylized, white letter "L" to the left of the text.



## (4) Work in Progress

- Faster Stats/syslog for slapd
- 2-phase commit for LMDB and LDAP txns
- Other LMDB enhancements

LMDB

# Work in Progress

- Faster Stats/syslog for slapd
  - glibc syslog() is braindead
    - acquires a mutex to write a msg on a datagram socket, which is already inherently atomic
    - OpenLDAP 2.4.39 8-core server 200,000 queries/sec with no logging
    - With Stats logging enabled, 21,000/sec - ~10x perf loss
    - With streamlined OpenLDAP syslog(), 26,000/sec
  - Multiple other bottlenecks

**IMDB**



# Work in Progress

- Faster Stats/syslog for slapd
  - rsyslogd/syslog-ng are major hogs, use 100% CPU to accept slapd log traffic
    - use our own single-purpose syslogd
  - libc is still a significant hog, 10% slowdown just formatting msgs, skipping actual msg send
    - avoid stdio/sprintf for msg formatting

IMDB

# Work in Progress

- 2-Phase Commit for LMDB and LDAP txns
  - Requirement is unavoidable if we want to support txns across back-ldap/back-meta etc.
  - Update to RFC 5805 txn spec
    - TxnPrepare with abort on timeout or commit on timeout

LMDB



# Work in Progress

- 2-Phase Commit for LMDB and LDAP txns
  - Requirement is unavoidable if we want to support txns across back-ldap/back-meta etc.
  - Update to RFC 5805 txn spec
    - TxnPrepare with abort on timeout or commit on timeout

LMDB


# Work in Progress

- Other Enhancements for LMDB 1.0
  - Incremental backup
  - Headerless overflow pages
  - Raw partition support
  - Optional support for DBs >2GB on 32-bit
  - Page-level encryption

# LMDB



# Questions?



IMDB