

OpenLDAP Retrospective

Howard Chu

CTO, Symas Corp. hyc@symas.com


Chief Architect, OpenLDAP hyc@openldap.org

2018-10-08

IMDB

Release Timeline

OpenLDAP Releases




Performance Model

- All builds are 32bit
 - Earlier versions aren't 64bit clean
- Testing with a 1 million entry database
 - Just under 2GB in size
 - 32bit fopen refuses to read LDIF bigger than 2GB
- The cache sizes used here are slightly smaller than optimal, max process size stays under 1.8GB

IMDB


Search Performance

Search Throughput



LDIF Import Time

LDIF Import Time




In the beginning...

- OpenLDAP started from the last release from the University of Michigan LDAP team
 - v3.3 from 1996
 - The UMich team had been hired by Netscape and went on to create the Netscape Directory Server
- Founded by Kurt Zeilenga in 1998
 - Needed a directory server for his company NetBoolean

LMDB


In the beginning...

Search Throughput



In the beginning...

LDIF Import Time




In the beginning...

- OpenLDAP 1.0.0 released 1998-08-26
 - Essentially UMich LDAP 3.3 plus unofficial patches that had been floating around mailing lists for years
 - 1.0.0 - 1.0.3, 1998-08-26 to 1998-11-06
 - 3 committers
 - Kurt Zeilenga (297)
 - Hallvard Furuseth (20)
 - Stuart Lynne (3)


IMDB

Release 1.0

Search Throughput



Release 1.0




Release 1.0

- Added POSIX(final) threads support
- Various Y2K fixes
- Mainly supported on Linux and FreeBSD
- (Tested here using BDB 2.7.7's db1.85 compat wrapper)

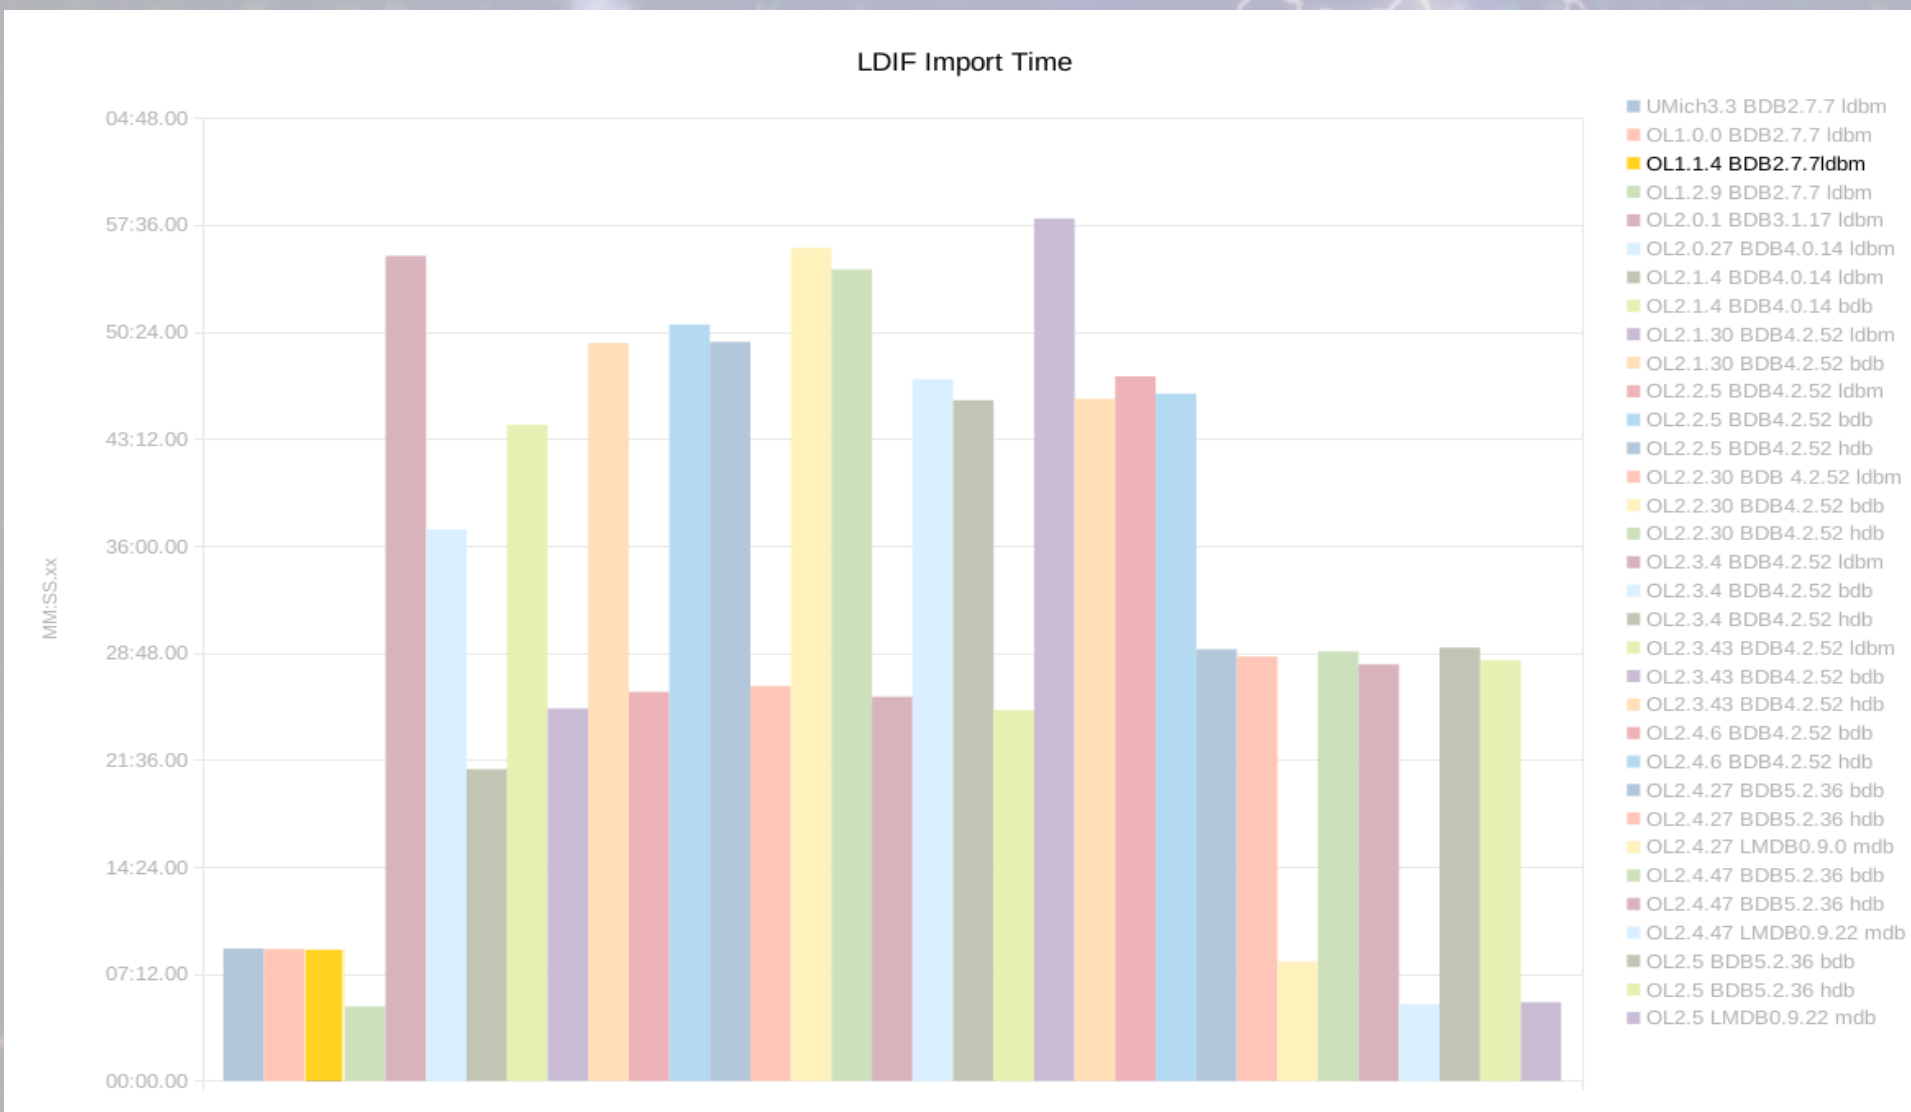
LMDB

Release 1.1

Search Throughput



Release 1.1




Release 1.1

- Added initial Windows NT support
- Adopted autoconf build tools
- Support for BerkeleyDB 2.x
- Seven contributors
 - Kurt Zeilenga (359)
 - Hallvard Furuseth (57)
 - Predrag "Pele" Balorda (26)
 - Dave Storey (13)
 - Bart Hartgers (11)
 - Kurt Spanier (9)
 - Randy Kunkee (2)

IMDB


Release 1.2

Search Throughput



Release 1.2

LDIF Import Time



Release 1.2

- 21 contributors
 - Kurt Zeilenga, Hallvard Furuseth, Julio Sanchez Fernandez, Howard Chu, Gary Williams, Mark Valence, Ben Collins, Juan Gomez, Pierangelo Masarati, Randy Kunkee, Stig Venaas, Kurt Spanier, Dmitry Kovalev, Mark Adamson, Predrag "Pele" Balorda, Ralf Haferkamp, Luke Howard, Will Ballantyne, Bart Hartgers, Bastiaan Bakker, John Quillan
- From here on out the community just keeps growing...

The logo for IMDB, consisting of the letters "IMDB" in a white, bold, sans-serif font with a thick white underline.


Release 1.2

- Significant improvement in import speed
 - multiple indexing fixes
 - removal of DN substring index
- Added creator/modifier attributes

LMDB


OpenLDAP 2.0

Search Throughput



OpenLDAP 2.0

LDIF Import Time



OpenLDAP 2.0

- The main thrust of OpenLDAP 2.0 was support for LDAPv3
- Also added SSL/TLS support, strong authentication support via Cyrus SASL
- Use a threadpool instead of spawning a thread per op
- Introduced back-sql

IMDB

OpenLDAP 2.0

- Significant slowdown in import speed
 - mainly due to schema validation
 - one of the defining differences between OpenLDAP and all of the other UMich-derived projects
 - also due to Unicode / UTF8 support
 - normalization / canonicalization of DirectoryStrings

LMDB


OpenLDAP 2.0

- This is where my/Symas' involvement really begins
 - Symas was building an enterprise resource management system on top of a directory
 - We wanted X.500-style chaining, the ability for a collection of servers to serve separate pieces of a single DIT
 - I wrote back-ldap based on OL1.2; it was released as part of 2.0
 - Had no desire/intention to touch back-ldbm or any database code

LMDB


OpenLDAP 2.1

Search Throughput



OpenLDAP 2.1

LDIF Import Time




OpenLDAP 2.1

- First serious look at performance
 - Revamped memory use
 - Avoid runtime computation of string lengths
- Ported to IBM mainframe with EBCDIC
- Introduced back-bdb, using BDB's transaction support
- Introduced back-meta, back-monitor, back-null, back-perl, backglue
- First OpenLDAP Developers' Day March 2003 in San Francisco

IMDB


OpenLDAP 2.2

Search Throughput



OpenLDAP 2.2

LDIF Import Time




OpenLDAP 2.2

- Initial support for overlays
 - dyngroup, pcache
- Introduced syncrepl
- Introduced back-hdb, a hierarchical variant of back-bdb that supported atomic subtree rename
- Added SLAPI support
- Merged slap* tools support into slapd binary
- Added smbK5pwd overlay

LDAP


OpenLDAP 2.3

Search Throughput



OpenLDAP 2.3

LDIF Import Time



OpenLDAP 2.3

- Introduced back-ldif, cn=config, back-relay
- Added slapadd -q quick mode
- Added more overlays
 - accesslog, auditlog, collect, denyop, dynlist, lastmod, ppolicy, refint, retcode, rwm, seqmod, syncprov, translucent, unique, valsort
- Introduced delta-syncrepl
- Added threaded indexer for back-bdb/hdb




OpenLDAP 2.3

- Discovered critical bug in BDB 4.6
 - reject it in configure tests
- Added epoll support
- Added more contrib modules
- Kurt Zeilenga resigns as Project lead
 - Howard Chu takes over as Chief Architect
 - Quannah Gibson-Mount starts assisting as Release Engineer

IMDB


OpenLDAP 2.4

Search Throughput



OpenLDAP 2.4

LDIF Import Time



OpenLDAP 2.4

- Enabled Lightweight Dispatcher in slapd
 - faster connection management
- Introduced back-mdb
 - designed as a standalone DB library to allow reuse in other projects
 - with the functionality of back-hdb wrapped on top
- Added support for GnuTLS and MozillaNSS
- Expanded cn=config functionality

IMDB


OpenLDAP 2.4

- Introduced back-ndb
 - built on MySQL Cluster backend
 - project killed when Oracle acquired Sun (which had acquired MySQL)
- Introduced back-sock
 - similar in spirit to back-shell, but sending requests to a separate process over Unix domain sockets for isolation / thread safety
- Added more overlays
 - dds, deref, sssvlv
- Added even more contrib modules
 - autogroup, cloak, dupent, kinit, lastbind, noopsrch, nops, nssov, samba4-related modules

LMDB


OpenLDAP 2.5

Search Throughput



OpenLDAP 2.5

LDIF Import Time




OpenLDAP 2.5

- Further analysis of mutex bottlenecks
- Added multiple queue support to threadpool
- Even lighter weight socket writer management
- More to come, but that's another talk

IMDB

Questions?



IMDB