



ApacheDS, 10 years in the making

Emmanuel Lécharny, SYMAS

Emmanuel Lécharny

- Works for **Symas** (www.symas.com)
- **Apache** Member
- Works on **ApacheDS** since 2005
- Also works on Apache **MINA**

The tale of 2 LDAP servers...

„It was the best of times, it was the worst of times, it was the age of **OpenLDAP**, it was the age of **AD**, it was the epoch of **C**, it was the epoch of **Java**“



A bit of history

Day 0

Alex Karasulu (some time after the internet bubble blasted) :

„What is the **crapiest** technology, that **nobody** wants to work on, and is used all over the IT, which will still be around in **20** years ?“

At the beginning...

- First SVN Incubating project ! (rev 131)

[/\[Apache-SVN\]/incubator](#)

Index of /incubator



Files shown: 0
Directory revision: [131](#) (of [1706036](#))
Sticky Revision:

File ^	Rev.	Age	Author	Last log entry
Parent Directory				
 directory/	r131	11 years	jwoolley	adding this per noel's req

- Proposed on Sep, 10, 2003
- Accepted on Nov, 12, 2003
- TLP on Feb, 23, 2005 !

Releases

- 0.8.0 : Jan 13, 2005
 - 1.0.0 : Oct 6, 2006 (EOL)
 - 1.5.0 : Apr 11, 2007 (EOL)
 - 2.0.0-M1 : June 25, 2011
 - 2.0.0-M20 : May 2, 2015
-
- Overall, a release every quarter for 10 years...

Sub-projects

- Studio : Started on september 2006, 27 releases
 - Ldap API : Started on early 2008, 32 releases
 - Mavibot : Started on June 12, 2012, 8 releases
 - Escimo : Started on July 7, 2013, no releases yet
 - Fortress : Accepted on Oct 25, 2014, 1 release
 - Kerby : Accepted on Dec 31, 2014, 1 release
- and
- MINA : TLP on Oct 26, 2006, 5 sub-projects, 73 releases
 - Triplesec : dormant

Code

- Studio : 220 422 SLOCs
 - Ldap API : 212 776 SLOCs
 - Mavibot : 28 793 SLOCs
 - Escimo : 7 125 SLOCs
 - Fortress : 90 716 SLOCs
 - Kerby : 63 677 SLOCs
- and
- ApacheDS : 223 709 SLOCs

Total : 847 218 SLOCs, as of today

Committees

- **Founder : Alex Karasulu**
- **Chairman : Kiran Ayyagari**
- **Former Chairmen : Alex Karasulu, Emmanuel lécharny, Pierre-Arnaud Marcelot**
- **PMC : 13 members (7 active)**
- **Committees : 14 (5 active)**
- **Emeritus : 16**
- **Overall : 43 people involved**

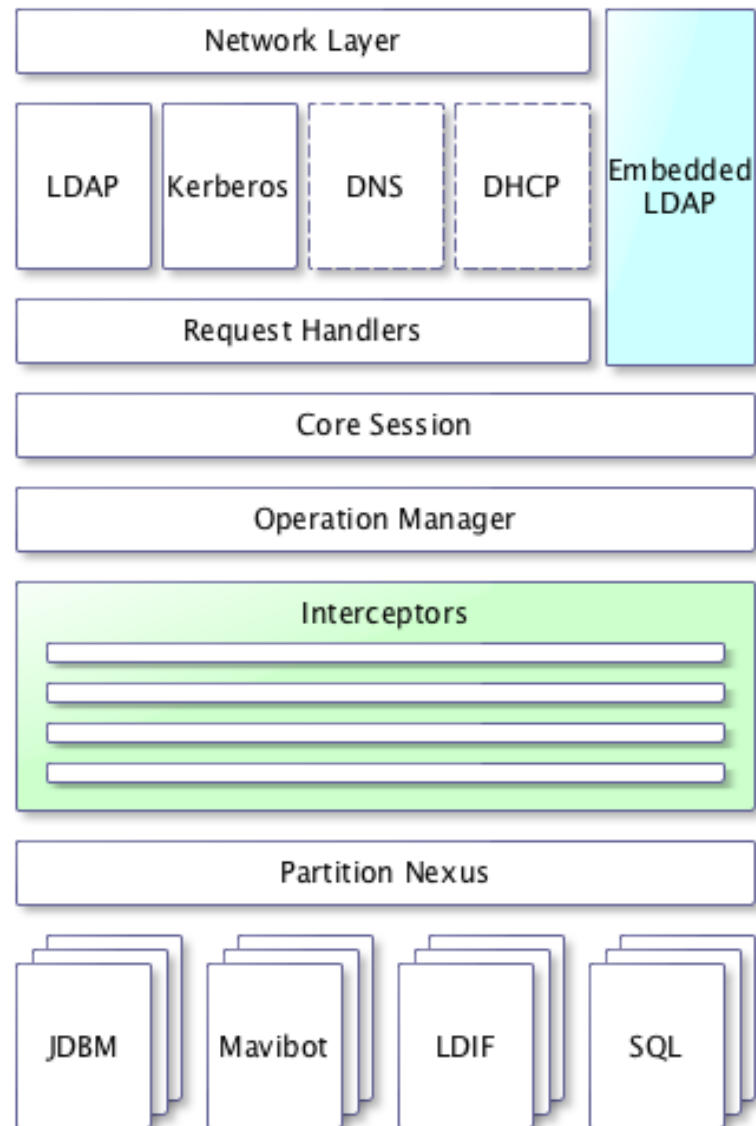
Technicalities

- **CVS → SVN in 2005, Git in 2013 (some projects)**
- **Ant → Maven 1 → Maven 2 → Maven 3**
- **Site : hand crafted → Confluence based → Apache CMS**
- **JIRA**
- **Bamboo → Hudson → Jenkins**
- **Java 4 → Java 5 → Java 6**
- **Many mailing lists**

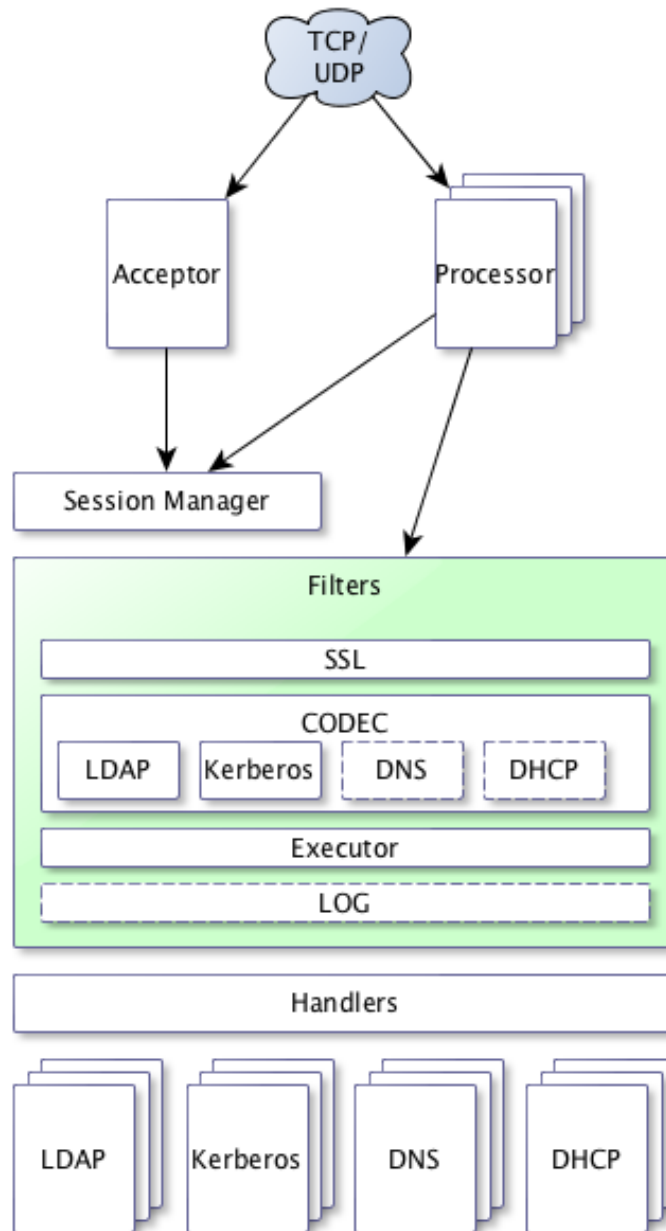


Let's dive in !

ApacheDS Architecture



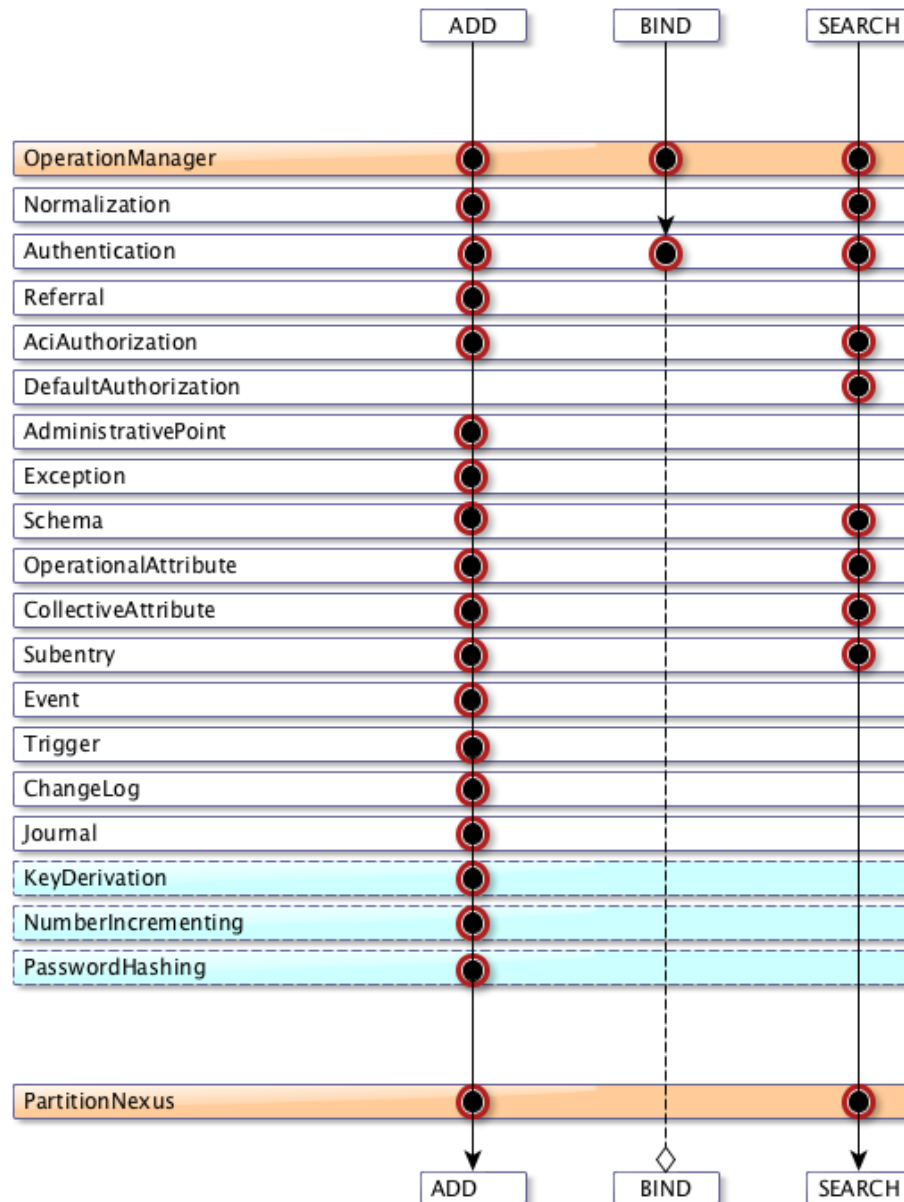
Network Layer



Protocol Layer

- DirectoryBacked
 - LDAP TCP
 - Kerberos TCP/UDP
 - DNS UDP
 - ChangePassword TCP
- No backend
 - NTP UDP

Interceptors

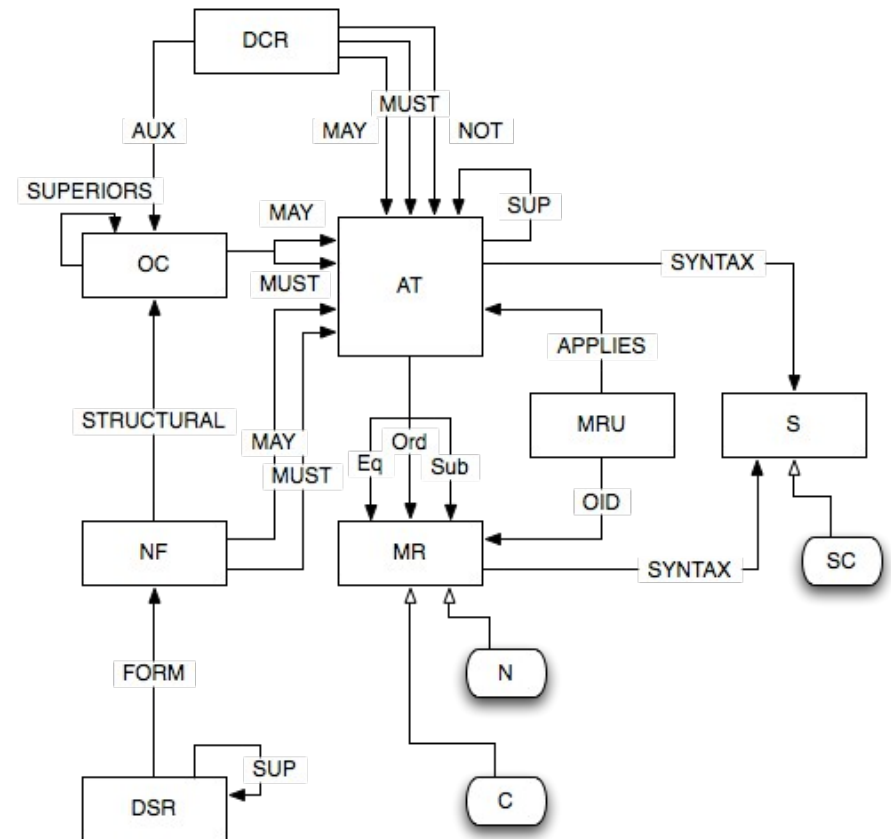


Backends

- **One Master table**
- **System indexes** (RDN, Presence, Alias, One_Alias, Sub_Alias, ObjectClass, EntryCSN, Admin_Role)
- **User indexes**
- **Reverse Index (RDN, Alias)**
- **Multi-Values**
 - **Array of values**
 - **sub-btrees when too many values**

Schema

- 8 LDAP elements + 3 internal
- SchemaManager
- Registries (8)
- LDIF based
- Semi-dynamic...
- Strict mode
- Relaxed mode in the API



Configuration

- LDIF based
- Loaded at startup
- Static (needs a restart when changed :/)
- Can't be replicated (atm)
- OSGi ?



A word on performances

Performances

- First version : 35 search/s...
- As of today : 13 000 search/s
- Still behind many other LDAP servers
 - Java still slower than C
 - Serialization is expensive
 - String prep (RFC 3454, RFC 4518) is costly
 - The current Backend is SLOW
- Mainly, Writes are sloooooooooow...

Performance : ASN.1

- Pre-incubation : SNACC (IBM, not open source)
- 1.0 Snicker : rewrite rules
- 1.5 Twix : State Engine (10x faster)
- 2.0 Complete rewrite, simpler SE (50% faster)
- Kerby : 20% faster, easier to use

Performance : DN

- First version : Antlr
- Second version : hand crafted (5x faster)
- Third Version :
 - Hand crafted parser for simple DN
 - Antlr parser for complex DN
 - 2x faster

Performance : Backend

- Read operations are fast (13 000 searches/s)
- Write operations are embarrassingly slow (8/s with a spinning disk, 50/s with a SSD...)
- Cache helps (a bit)
 - Entry cache
 - Page Cache
- No bulk loading (but comes with Mavibot)



Some cool features

Tests

- Annotations
- Reverter


```
@RunWith(FrameworkRunner.class)
@CreateDS(partitions =
{
    @CreatePartition(
        name = "example",
        suffix = "dc=example,dc=com",
        indexes =
        {
            @CreateIndex(attribute = "objectClass"),
            @CreateIndex(attribute = "dc"),
            @CreateIndex(attribute = "ou"),
            @CreateIndex(attribute = "member")
        },
        contextEntry = @ContextEntry(entryLdif =
            "dn: dc=example,dc=com\n" +
            "objectClass: domain\n" +
            "dc: example"))
    })
@CreateLdapServer(transports =
    { @CreateTransport(protocol = "LDAP") })
@ApplyLdifs(
{
    // Entry # 0
    "dn: cn=Kate Bush,ou=system",
    "objectClass: person",
    "objectClass: organizationalPerson",
    "objectClass: inetOrgPerson",
    "objectClass: strongAuthenticationUser",
    "objectClass: top",
    "userCertificate:: NFZOXw==",
    "cn: Kate Bush",
    "description: this is a person",
    "sn: Bush",
```

Replication

- Syncrepl RFC 4533
- Multi-Master Replication
- Delta-syncrepl to come
- Administrative model to come
- OpenLDAP/ApacheDS replication to come...

Administrative Model

- X.500
 - Access-Control
 - Collective Attributes
 - SubSchema
 - Triggers
- To do : Replication



And they said, “Come, let us build ourselves a city, and a tower whose top is in the heavens; let us make a name for ourselves, lest we be scattered abroad over the face of the whole earth.”

Not a perfect world !

- **Big community** → bump in the road
 - **Conflict of interest**
 - **Inter-personal conflicts**
- **10 years** → some may lost interest
- **Code age starts to show...**
- **10 year old architecture might not fit**

Configuration changes...

- Configuration
 - Property files
 - Spring : welcome to XMHELL...
 - LDIF based, but static...
 - What about a dynamic configuration ?

Backend changes...

- JDBM
 - AL 2/0 : fine !
 - Small code base
 - B-Tree with transaction : required
 - But no support for cross B-tree transaction : killing !
- Mavibot (Wip)
 - MVCC Btree
 - Cross B-tree transaction
 - Fast

Search internals

```
(&(cn=test)(ObjectClass=inetOrgPerson))
```

- **Evaluation :**

```
([200]&([200]:cn=test)([450]:ObjectClass=inetOrgPerson))
```

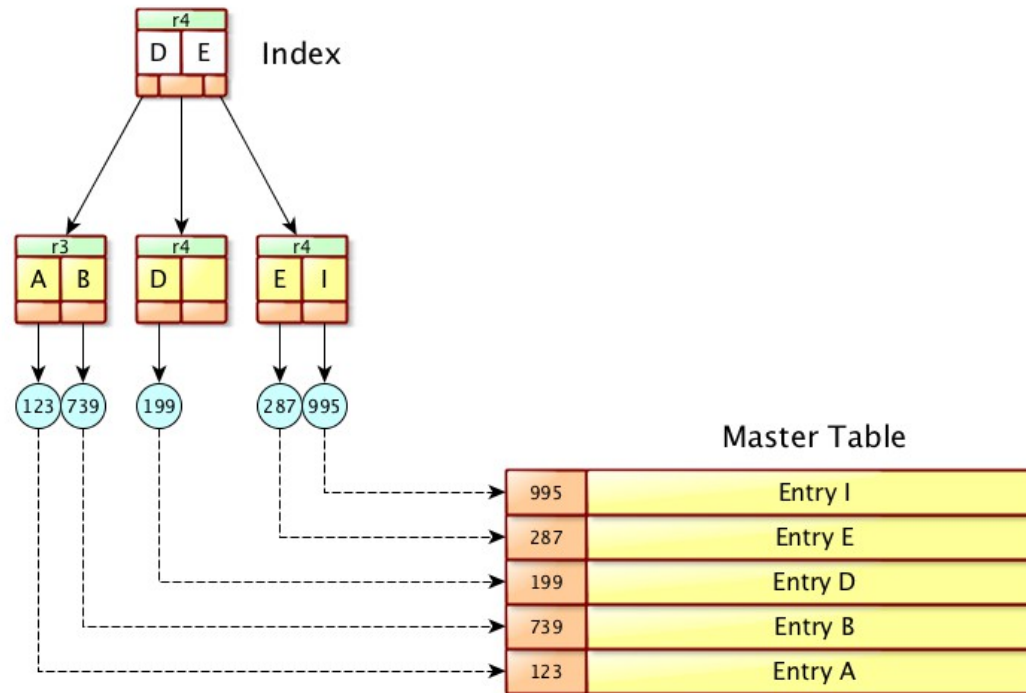
- **Candidates**

{ 145, 277, 137, ... } : 200 elements

- **Processing**

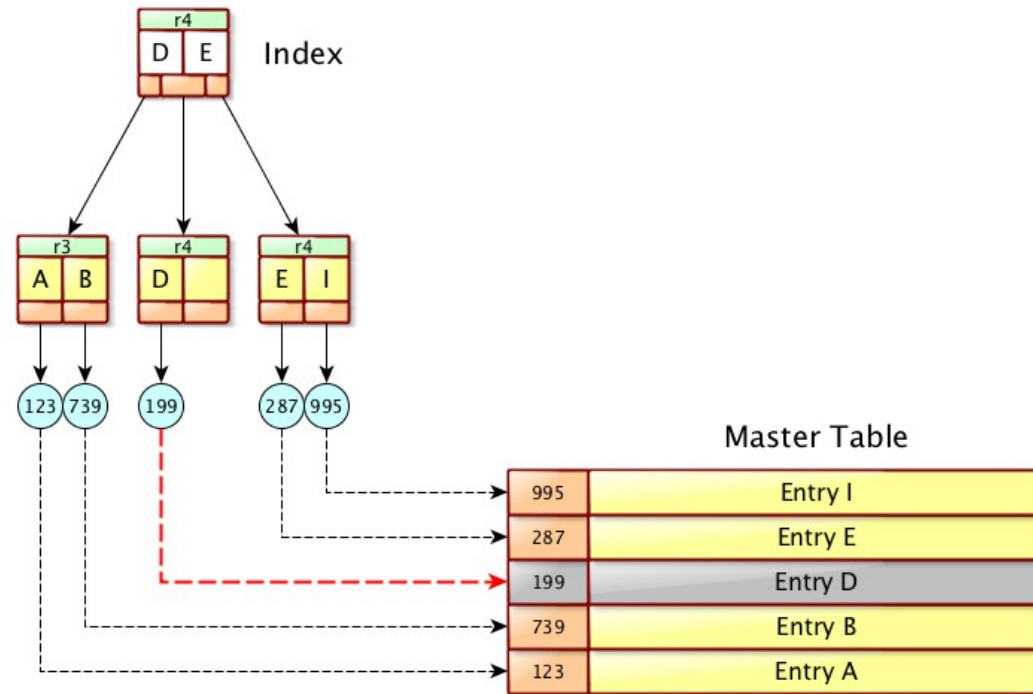
Check each Entry against the filter (200 times) using the **cn** index, reading entries from the **master** table

Search internals (2)



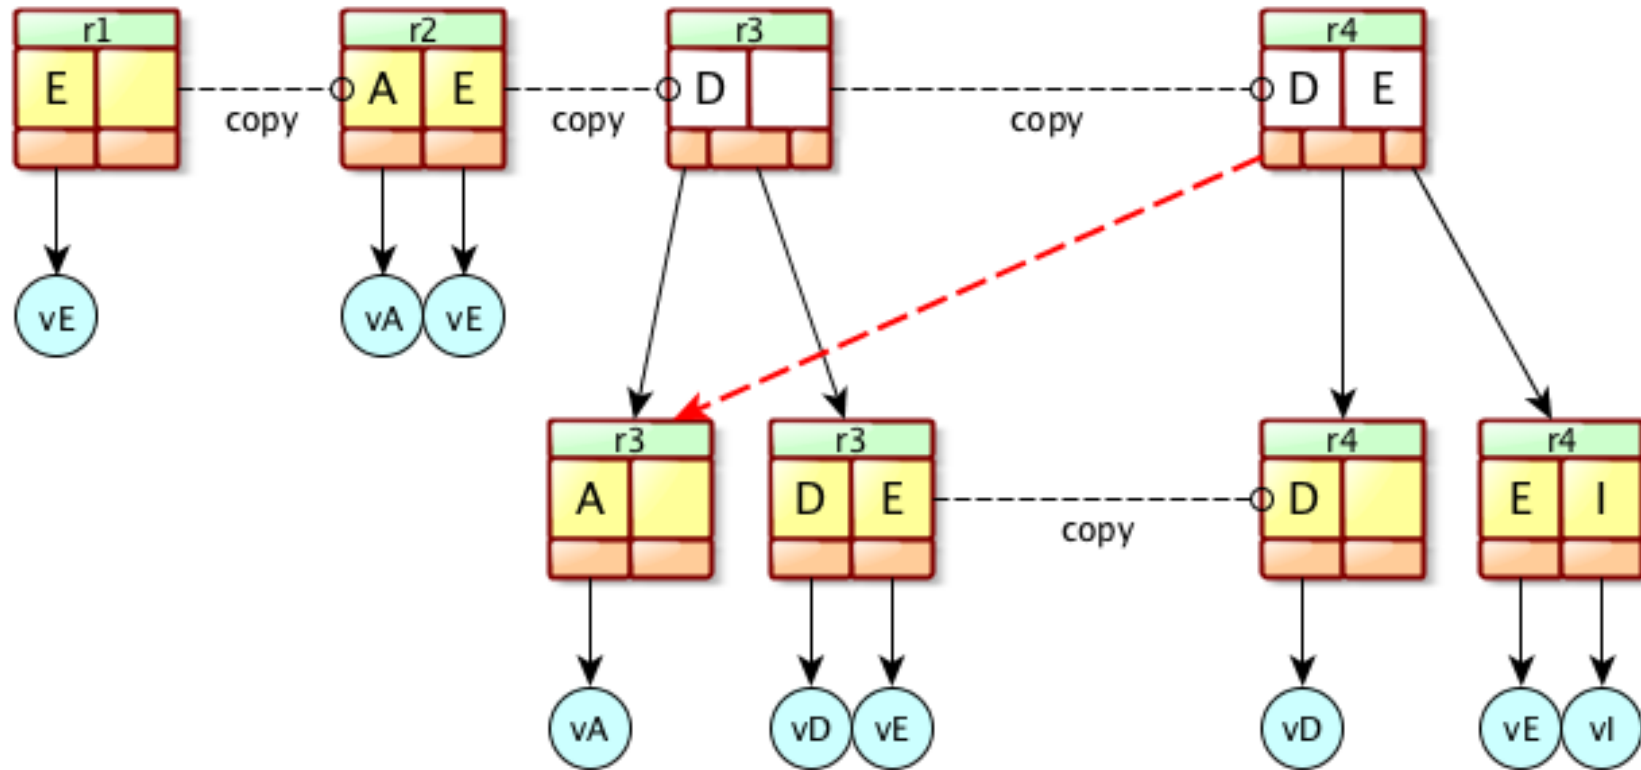
- What if another thread modify the Master table ?

Search internals (3)



- This occurs because updates on the index and the Master Table are not done in a transaction...

Mavibot



- We keep old revisions up to the point we don't need them.

Code changes...

- Java 4 → Java 8
 - Java API breakages
 - Generics (Kind of a nightmare)
 - Annotations (big plus!)
 - Concurrent Data structures (At last !)
 - JVM bugs... (workarounds needed)

Architecture changes...

- OSGi
 - Require some kind of refactoring (well, a lot !)
 - Require some understanding on what's inside
 - 'stop the world' effort :/
 - Bundles : ok
 - Life-Cycle ?
 - Services ??
- At the end, a paradigm change !



**There is more than just
ApacheDS!**

Apache LDAP API

- JNDI sucks (Wanna use JNDI ? **J**ust **N**ot **D**o **I**t...)
 - No schema
 - Not LDAP compliant
 - Not user friendly
 - Not Open Source
 - Dead in the water...
- Apache LDAP API
 - Tries to solve these issues !

Apache LDAP API (2)

- ApacheDS uses the API
- Studio uses the API
- Fortress uses the API
- **You should use the Apache LDAP API !!!**

STUDIO



STUDIO (2)

- LDAP Browser
- LDIF Editor
- Schema Editor
- ApacheDS/OpenLDAP configuration Editor
- ACL Editor
- ApacheDS embedded !

What's next ?

- Mavibot only
 - Memory Mapped File ?
- ACL rewrite (Fortress ?)
- Administrative Model extension
- Stored Procedures revival
- ApacheDS/OpenLDAP cross replication

Today

„... **crapiest** technology...which will still be
around in **20 years** ?“

Mission half-accomplished !!!



Thanks !

QUESTIONS ?