

LucyProposal

Superseded

This proposal launched Apache Lucy as a subproject of Lucene in 2006.

In 2010, Lucy was rebooted and entered the Apache Incubator. The 2010 proposal can be found at [<http://wiki.apache.org/incubator/LucyProposal>].

Name

Lucy

Synopsis

A shared core engine for Lucene, written in C, with Perl and Ruby bindings.

Rationale

Ferret, a Ruby implementation of Lucene written by David Balmain, and KinoSearch, a Perl implementation of Lucene written by Marvin Humphrey, are both driven by C engines. It has proven impractical to implement Lucene in either pure Ruby or pure Perl; however experience has shown that the Perl and Ruby communities are most enthusiastic about using and contributing to native code-bases.

Lucy will provide a C engine which can be shared across both Perl and Ruby implementations. It will provide core, performance-critical functionality, but leave as much up to the higher-level language as possible. KinoSearch's back end will serve as a template, as certain algorithms developed for KinoSearch, in particular its merge algorithm, were designed specifically to work well with the OO models of interpreted languages. Native versions will be preferred for public classes, though that will not always be feasible.

Sharing a C core will make it possible to for Lucy's developers to pool expertise on the aspects of Lucene our communities are least able to contribute to. Not coincidentally, these are often the most challenging aspects of Lucene to work on.

Eventually, it may be possible to bind Lucy to other languages, including Java.

Committers

- David Balmain
- Doug Cutting
- Marvin Humphrey