

LibreS3

Design, challenges and steps toward reusable libraries

Edwin Török

edwin@skylable.com



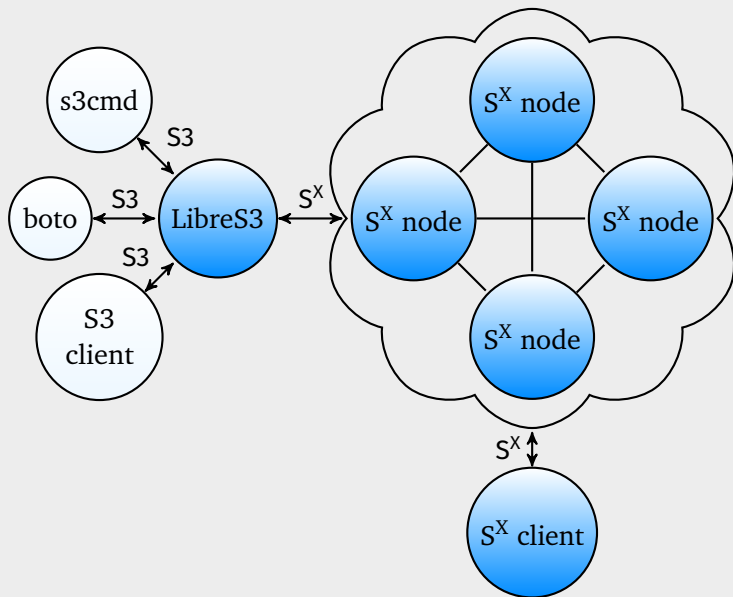
The OCaml Users and Developers Workshop, 2014

Concepts

- Monads
 - ▶ deferred computations
 - ▶ result monad
- S3 server
 - ▶ proprietary: Amazon S3 server
 - ▶ FOSS replacements



S^X cluster, LibreS3 server



Libraries & architecture

- principle: you don't have to choose one particular library/framework
- not novel: existing practices, monadic interface, OASIS, OUnit
- be careful with acquiring resources: `with_resource`
- some code in LibreS3 could be generalized
 - ▶ caching, json/xml parse to/generate from OCaml types, http, I/O
- available as sample code: `any-cache`, `any-http`



any-cache

```
val lookup_exn: ( $\alpha$ , exn) t  $\rightarrow$   
  string  $\rightarrow$  (string  $\rightarrow$   $\alpha$  M.t)  $\rightarrow$   $\alpha$  M.t
```

- Least Recently Used with 2 Queues (PostgreSQL, OpenBSD)
 - ▶ main cache: for often (re)used keys
 - ▶ input cache: for newly used keys
 - ▶ victim cache: using fixed-size weak-array
- Functorized interface, parameter: a result monad
- Functorized interface : monadic concurrency \rightarrow result monad
- Testing: a functor using OUnit, instantiated with Lwt/Async/immediate



any-http

- not novel, very similar to Mirage
- functor interface based on Cohttp API
- implementations: Lwt&Ocsigenserver, Lwt&Cohttp, Async&Cohttp
- use: write REST applications independent of HTTP/monadic concurrency library
- testing: compatibility testing between multiple HTTP implementations



Issues

- good collaboration with upstream maintainers for patches
- double-release of OCaml runtime lock in OCamlNet, ocaml-ssl
 - ▶ debugging: patch OCaml's thread runtime (on mantis)
- `Unix.fork` vs `Lwt_unix.fork`
- `ocamlbuild SIGPIPE`
- Xen Dom0 crash (XSA-90) when running Mirage app
- PUT/DELETE support in Ocsigenserver
- Ocsigenserver default unix worker threads too high (1000)
- debugging monadic code is hard: stacktraces from monadic concurrency incomplete (many tail-recursive calls), needs better logging
- ...



Packaging

- Easy if your user has OPAM, but ...
 - ▶ lack of "new enough" (RHEL/CentOS 6)
 - ▶ newer compiler versions: 3.12.1, 4.00.1
 - ▶ opam switches: 4.01.0+lsb, 4.01.0+musl+static
 - ▶ lack of complete packaging workflow
 - ★ oasis2opam ✓
 - ★ oasis2debian ✓
 - ★ opam2deb ?
 - ★ opam2rpm ?
 - ★ opam2bsd ?
 - ★ opam2fullsrc ?
 - ▶ ...



Summary

- LibreS3: a FOSS implementation of the S3 server protocol
- Reusable libraries: any-cache, any-http
- Bugs, packaging issues

- Future work
 - ▶ any-io, ioconvcomb
 - ▶ opam2release

Additional material: <http://goo.gl/jmF0cn>

