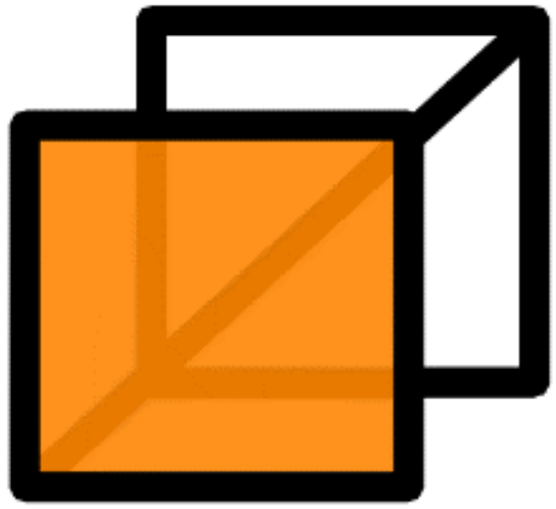

BOX@PSNC

Maciej Brzeźniak, Stanisław Jankowski, Sławomir Zdanowski
HPC Department, PSNC, Poznan

BOX@PSNC - OVERVIEW

- Where we are and why
- Seafile - features, motivation
- Seafile operations and usage experience
- Future

WHERE WE WENT & WHY?

IN SYNC & SHARE BUSINESS

- We started in 2012 by implementing **CryptoBox** under NDS2 project (presented at TNC'13 & '14)

- We ***developed*** *solution a sync & share client with encryption* for Windows (SFTP+AES512/RSA4096)

- We now ***use*** *a solution:* Seafile

EXPERIENCE FROM SYNC&SHARE SOLUTION DEVELOPMENT

- **We've been there…**
  - CryptoBox works for most Windows versions
  - It **outperformed** popular solutions at the time (2014) even while encrypting!

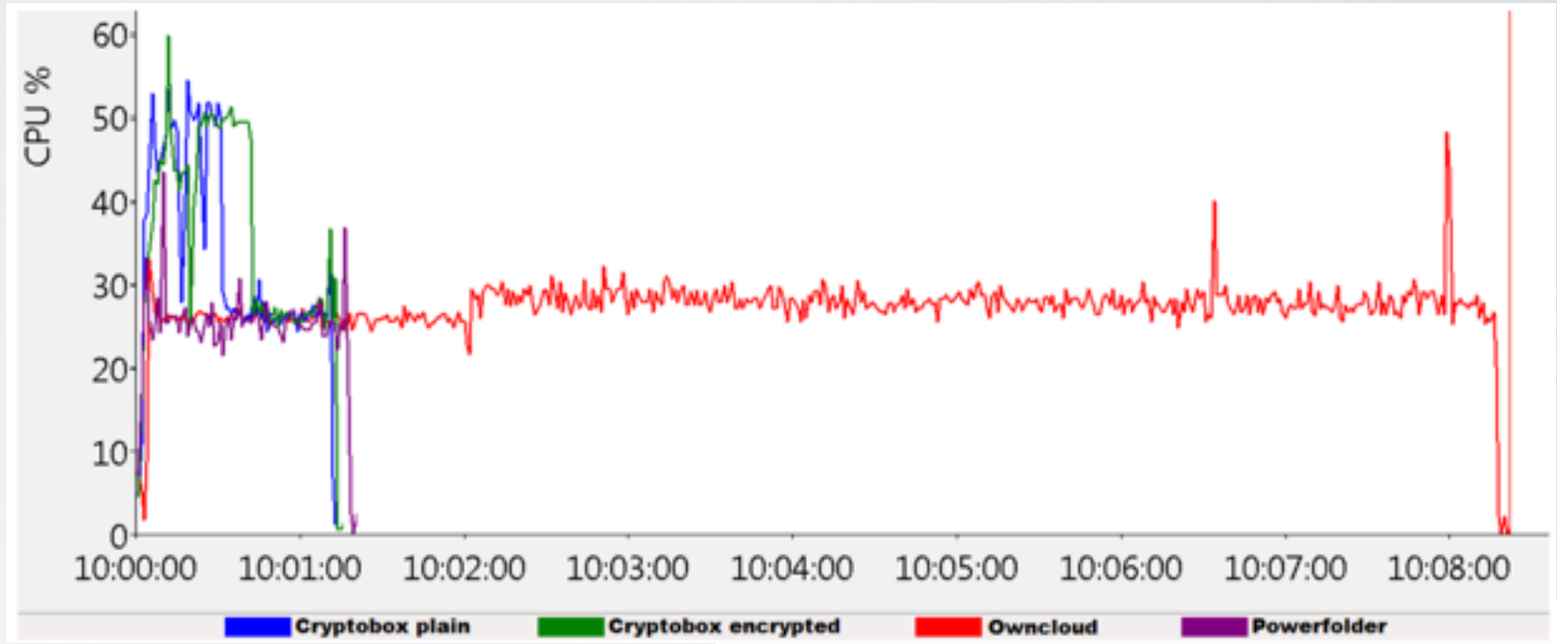

- **but**
  - keeping it all up is too costly: Win, Linux, MacOS, Android, iOS
  - CryptoBox was one of the NDS2 clients: CryptoFS, CryptoDroid, CryptoGUI
  - CryptoBox uses FS, in NDS system metadata are kept in DB: => ***this is not enough scalable*** for sync&share (although some believe in it)

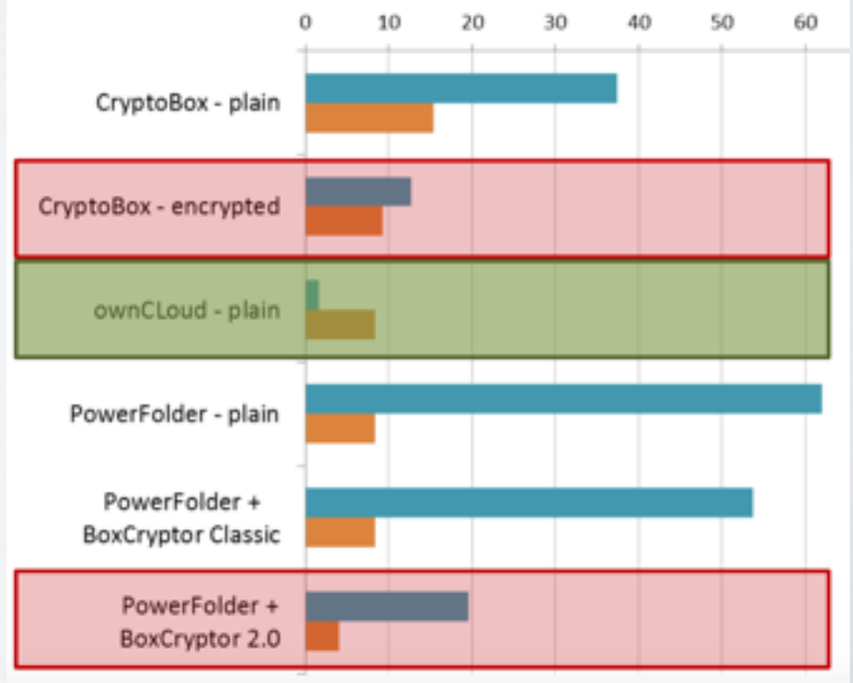

SEAFILE ADVANTAGES

- Fast & reliable & ligthweight sync, share and collaboration
- Speaks to filesystem, NFS and object stores: S3, Swift, Ceph.
- Both data and meta-data can be stored on object stores!
- Minimum data in the DB
- Daemons implemented in C

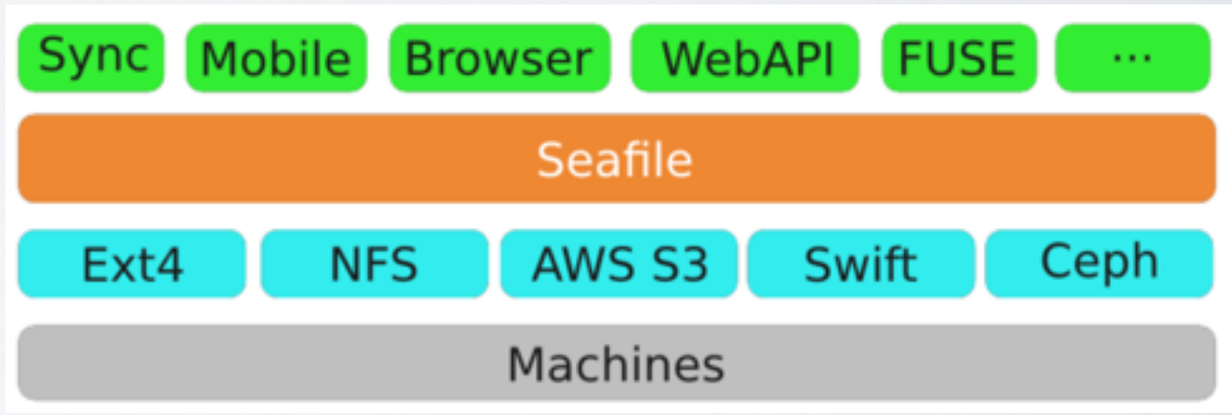

SEAFILE SYNCHRONISATION

- syncs based on library / filesystem snapshots,
  not per-file versioning

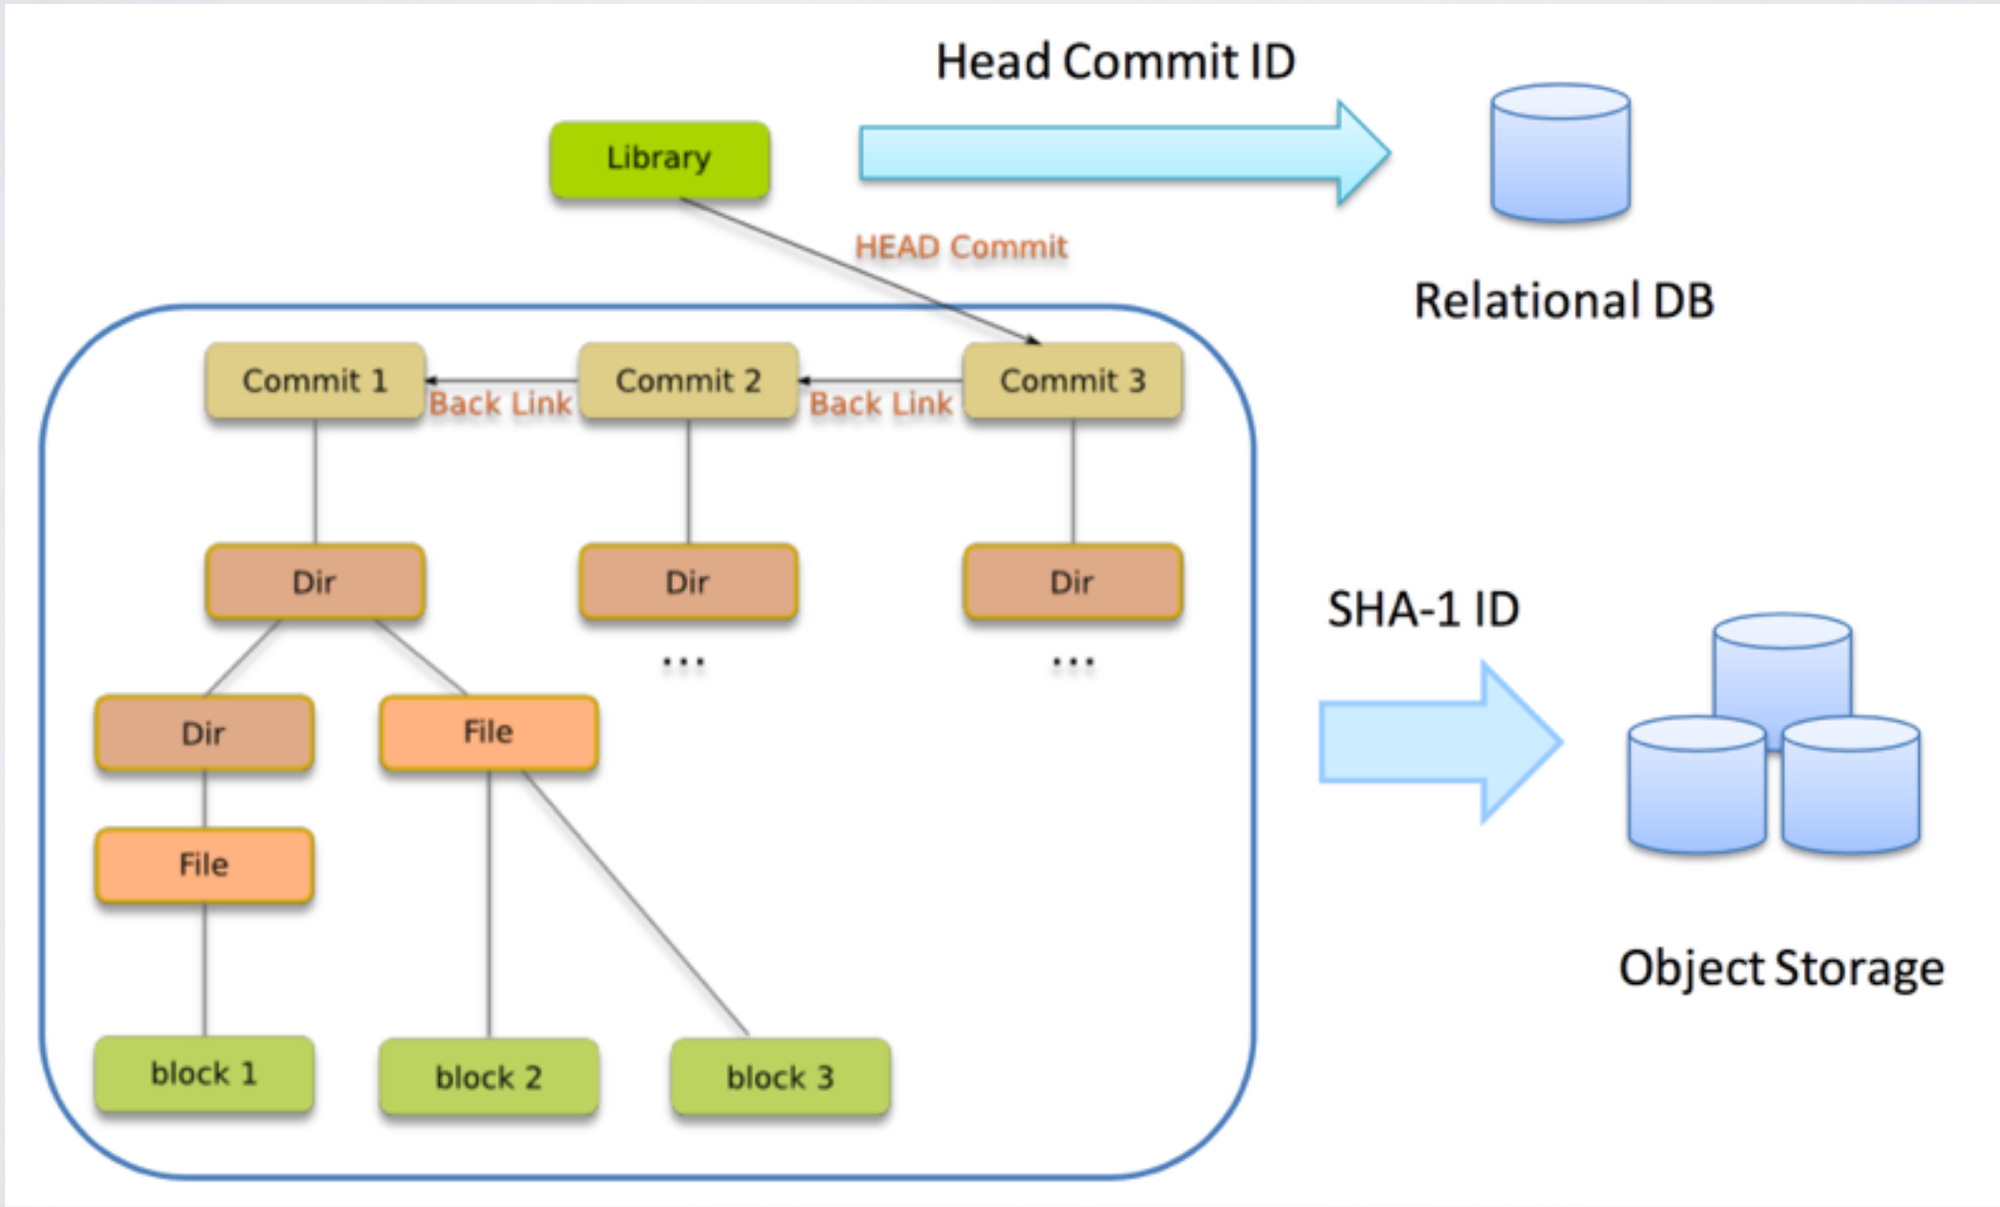

SEAFILE SYNCHRONISATION

- only deltas are included in the commits (snapshots of library)
  data is chunked using Content Defined Chunking algorithm

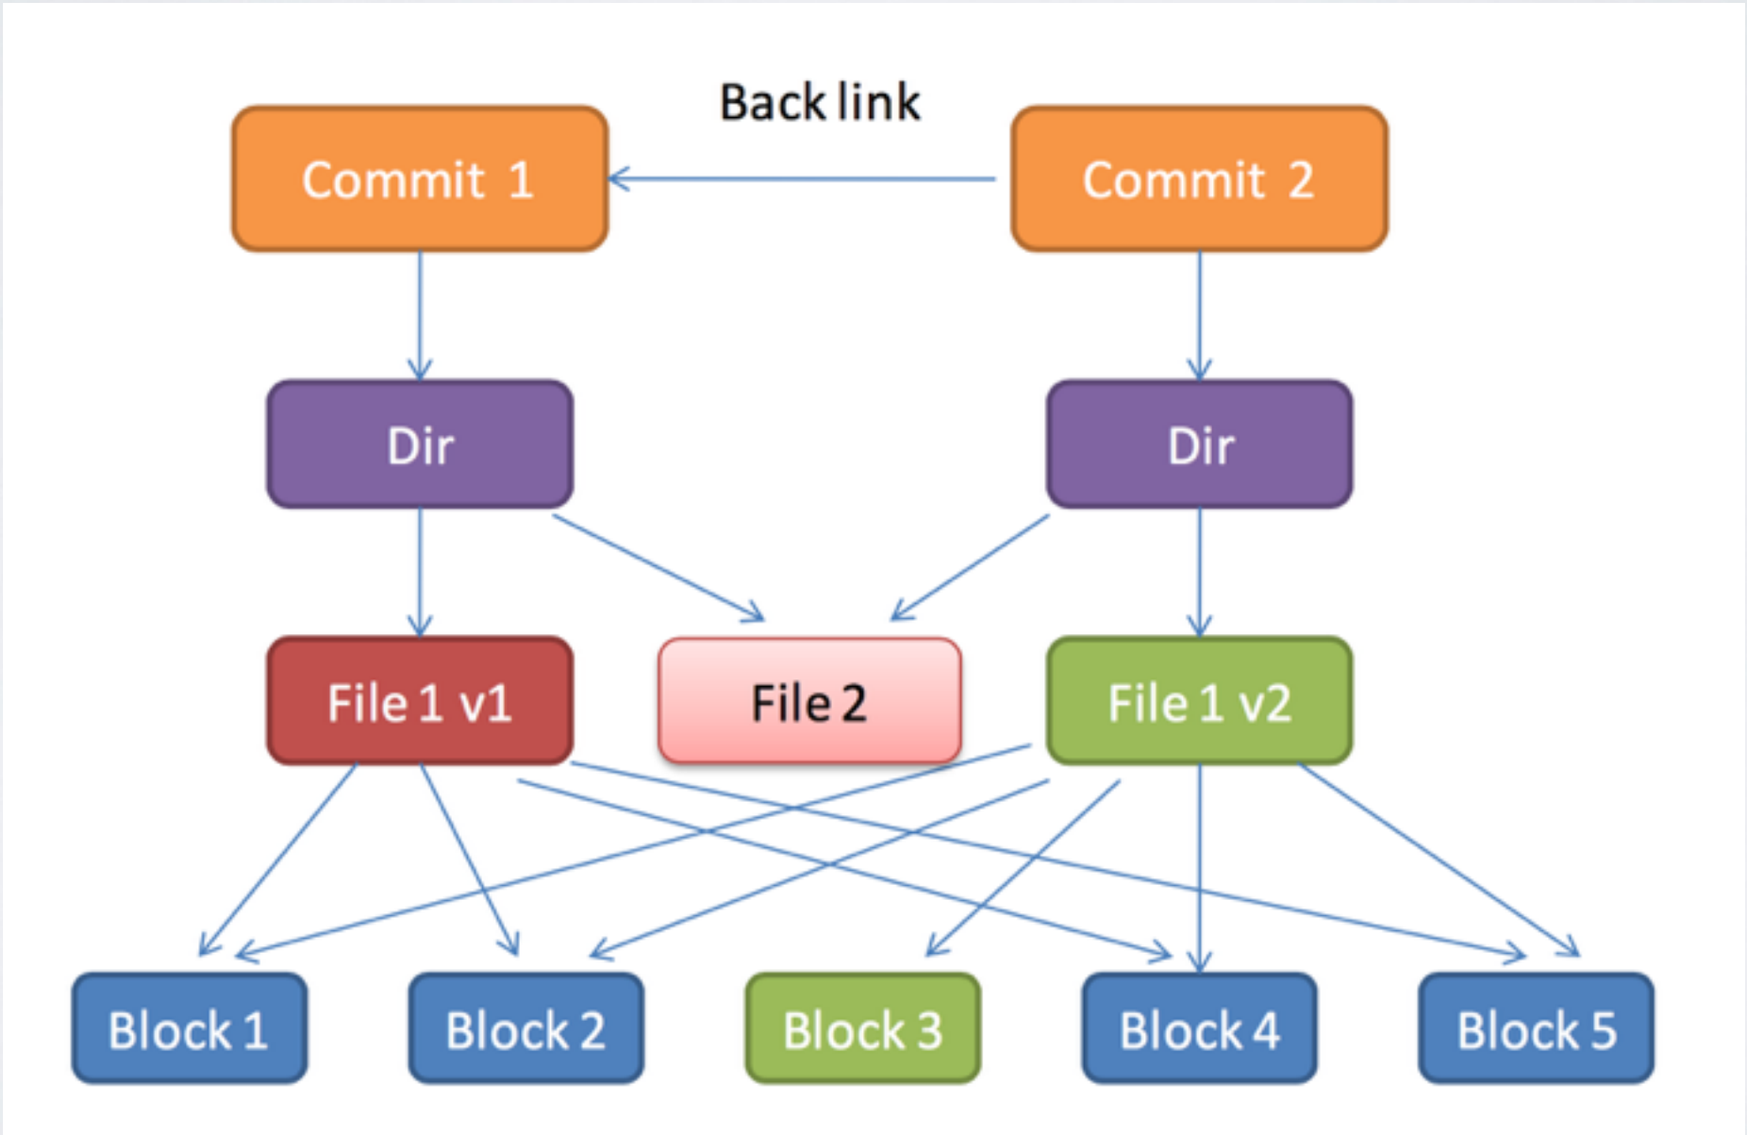

SEAFILE CHALLENGES

- Dedicated infrastructure needed
- Exchanging / sharing data with POSIX filesystems difficult
- Exit path:
  - data and metadata stored in a specific format
  - Seafile provides an export tool in seaf-fsck to export server data without the database

PROS VS CONS -> DECISION WAS…

- One size fits all? NOPE!
- Performance of sync&share
  on top of filesystem + DB
  will always hit the wall
  - the architectural wall :)
- We're not in the corner case ;)

WHERE HAVE WE LANDED

SEAFILE EXPERIENCE

- Limited scale pilot service
  @ PSNC since 12'2014
  (with an ambition to extend)

- 400 users, mainly PSNC'ers
  plus other people: staff at
  universities, project partners

- 500 libraries, ~1 TB space used

- Purpose: understand how
  it works, check user experience,
  compare to other solutions,
  & discover issues?

OPER'S WOOOWS

- **WORKS AS CHARM!**
- Interventions:
  - upgrades: 4.0.5 -> 4.0.6 -> 4.4.6 (seamless)
  - increasing apache limits
  - adding storage space ;) … VMware + LVM
- practically NO LOAD ON SERWER!

  WE RUN IT ON THE TINY **8GB RAM** , 2xVCPU @AMD Opteron 2435

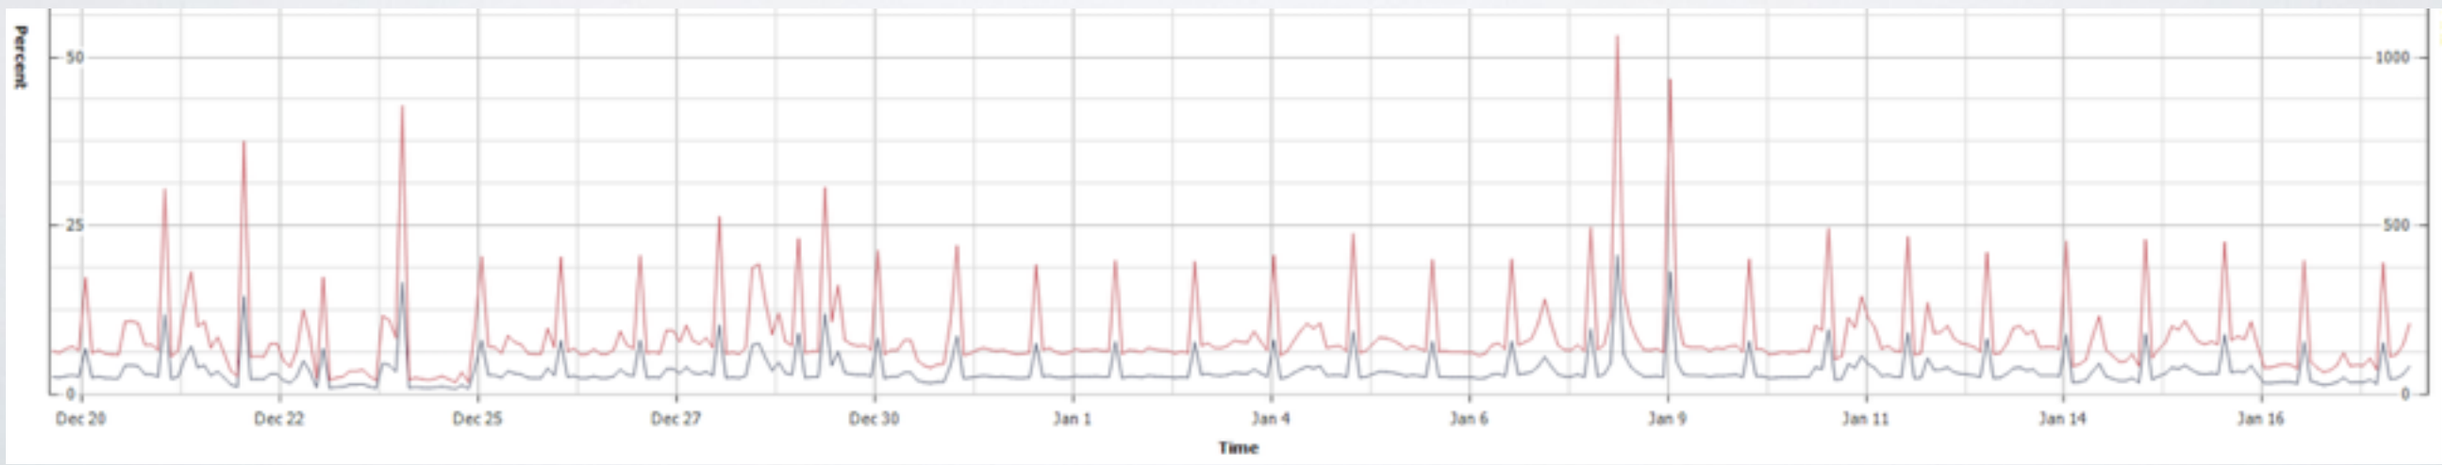

FIRST IMPRESSIONS?

- **USER'S WOOOOWS:**
  - It's REALLY FAST!
  - No load on clients!
  - Many libraries / local dir pairs can be configured
  - Thunderbird Filelink plug-in vastly used & appreciated!
- User issues and workarounds:
  - 500 files can be uploaded at once by Web (more possible with dekstop clients)
  - limit of one-shot dowload (setup to 5GB)

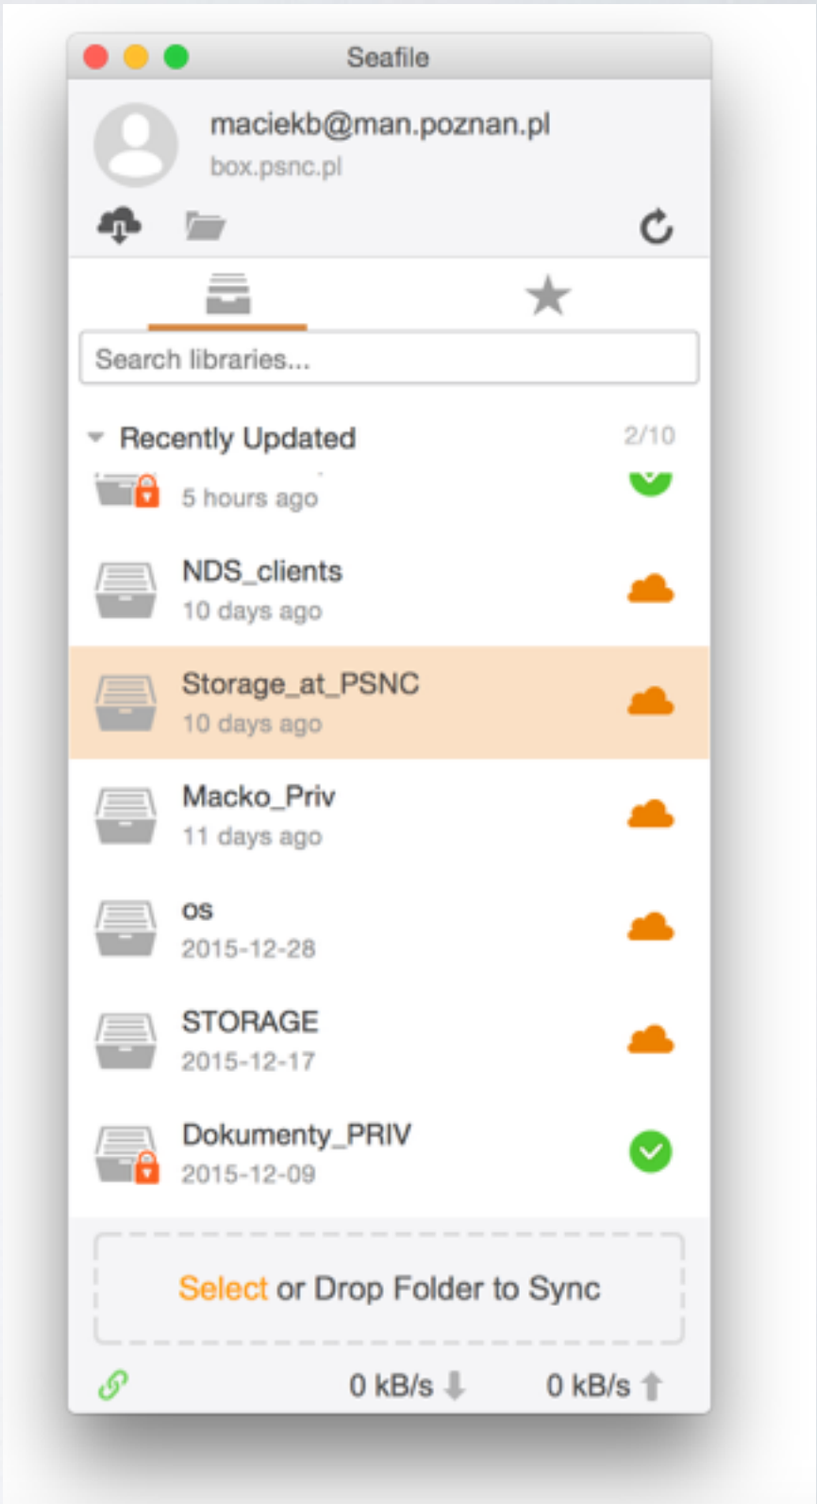

DAILY LIFE WITH SEAFILE

- **USAGE STATS & TRENDS:**
  - Most people use Web but more requests come from desktop clients - *see stats ->*
  - Mobiles (Android) are not generating many requests but are appreciated

    Mobile apps: aren't full sync&share apps

Client type	overall	desktop clients
Win	85%	84%
Linux	13%	11%
OS X	<1%	4%
Android	<1%	<%1
other	<1%	<%0.1

DAILY LIFE WITH SEAFILE?

- **USAGE cont.:**
  - Easy sharing data with others
    - shared libraries
    - public/protected links
  - Syncing own data among devices:
    - *mostly desktop-desktop*
  - Reads dominate writes
  - GETs dominate PUT/POSTS

Link type	created	re-use
upload	21%	4,8x
download	79%	21x

Share types	user	groups
libraries	51%	49%

HTTP	MGETs	MPUTs
on „repo”	36	0,9

WHERE ARE WE GOING

NOW -> FUTURE:

LDAP	SAML/Shibboleth (works) PIONIER.ID (testing)
Filesystem / VMFS / VMware	Filesystem / RBD / Ceph (usage) Object Store: librados (testing)
@PSNC	-> Universities in Poznan -> …
Single server	-> cluster of servers (load balancing, HA)

NOW -> FUTURE:

ceph

LATEST IMPROVEMENTS @SEAFILE

- Real-time backup (HA clustering functionality) — until now only clustering for load-balancing
- More responsive & wide-screen GUI for desktops
- Server-side data encryption: used drives decomissioning safer and easier

CONCLUSIONS:

- Seafile is a solid sync&share solution:
  scalable architecture, good-quality low-level code

- Seafile team looks to be aware of academic community needs: see integration of SAML etc.

- PSNC's experience is promising (however limited scale-based)

- PSNC intends to extend th scope of the service