

SYSTEM OF HPC CONTENT ARCHIVING

Alexander Bogdanov, Prof., D.Sc.
Saint Petersburg State University
bogdanov@csa.ru

Andrei Ivashchenko, postgraduate
Saint Petersburg State University
aiivashchenko@cc.spbu.ru

The problem


Educational process based on

- Lectures
- Presentations
- Publications
- Textbooks
- etc.

We have to

- Store them
- Distribute them
- Provide an access
- Keep them up to date

Data gained from text document



Binary data

Plain text data

Document structure

Metadata

Main goals

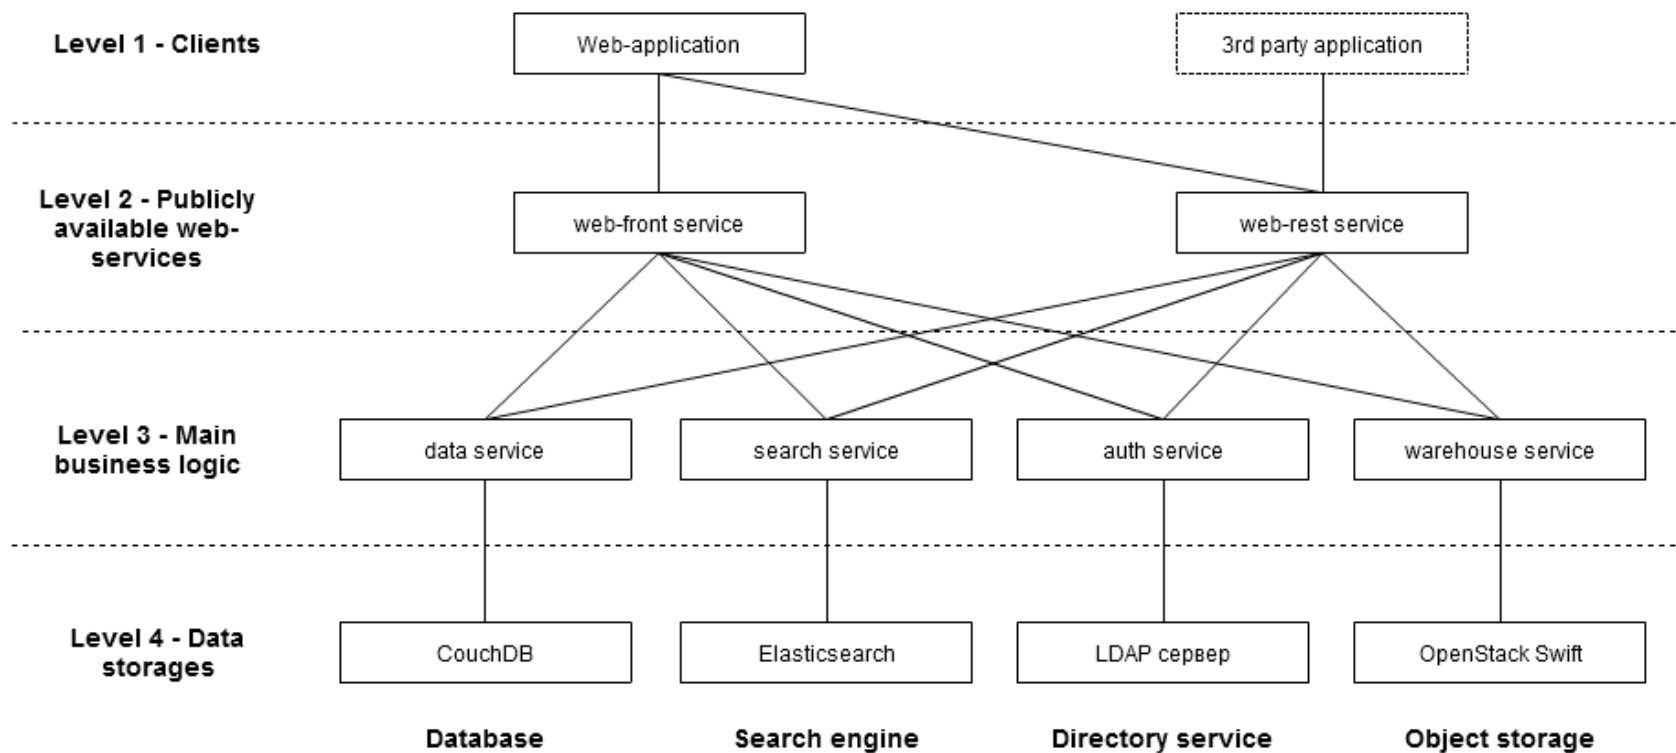
We have to organize:

- Long-term file warehousing
- Basic file operations support
- Full text search based of file contents
- Cluster file analysis for documents

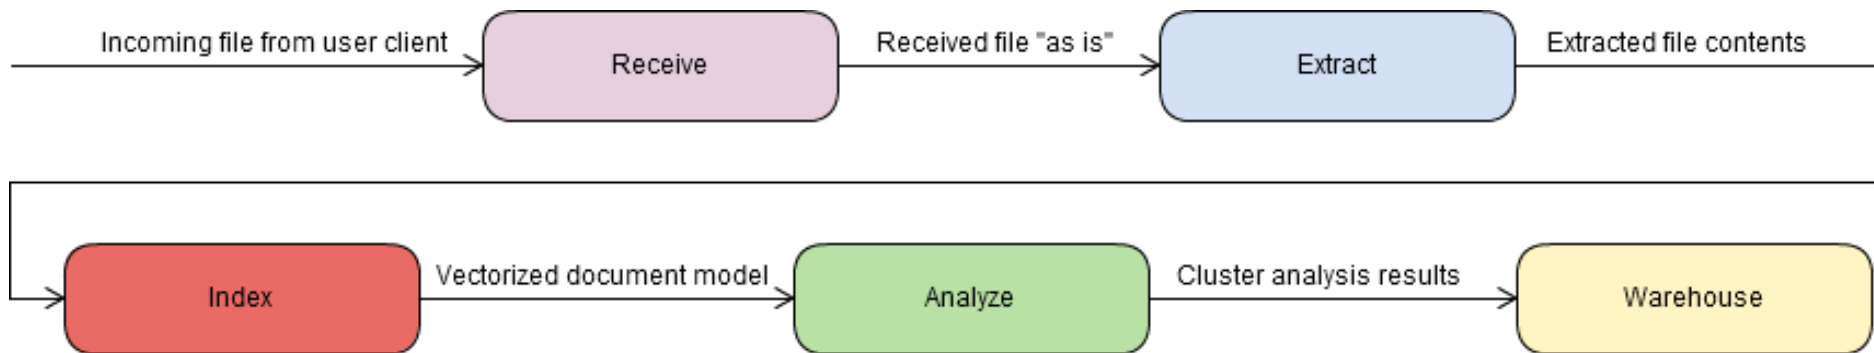
Complicating factors:

- Big amount of files
- Different formats of documents are getting used
- User stores files not systematically

System layer architecture



Data flow



We are dealing with...

Processing results

- Huge amount of small data chunks
- Similar to parent documents


Files

- File size varies a lot, up to ~100MB
- Could be versioned
- Could be duplicated

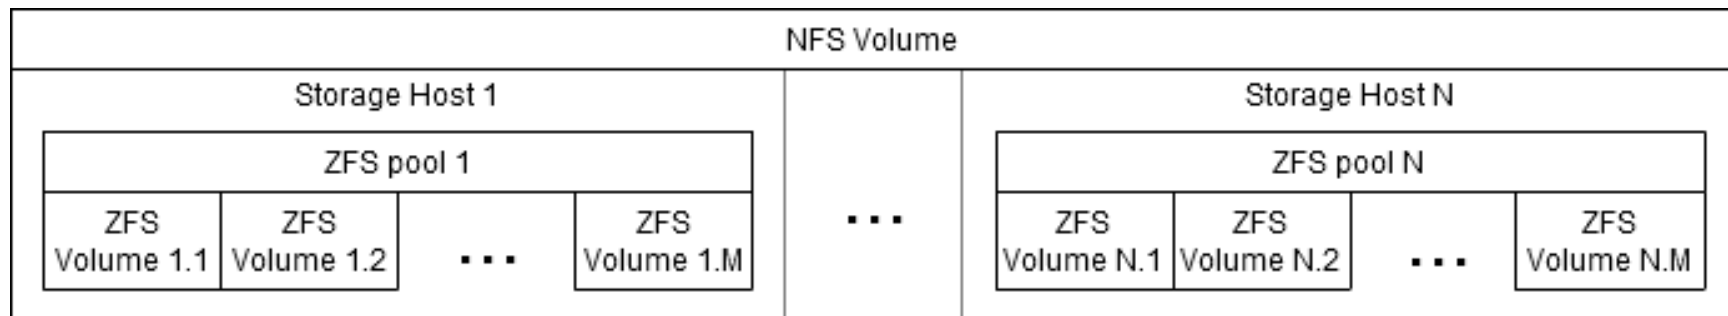
Deduplication support

- OpenDedup
- btrfs
- lessfs
- ZFS

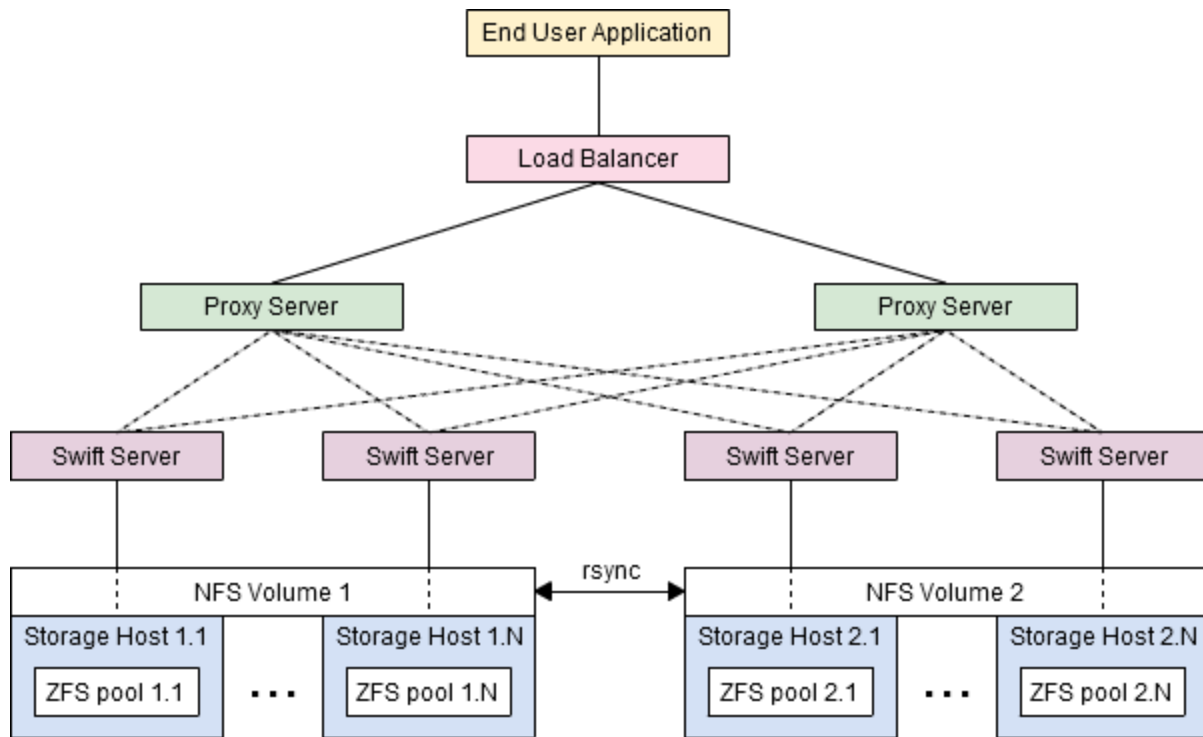
ZFS configuration on a single host



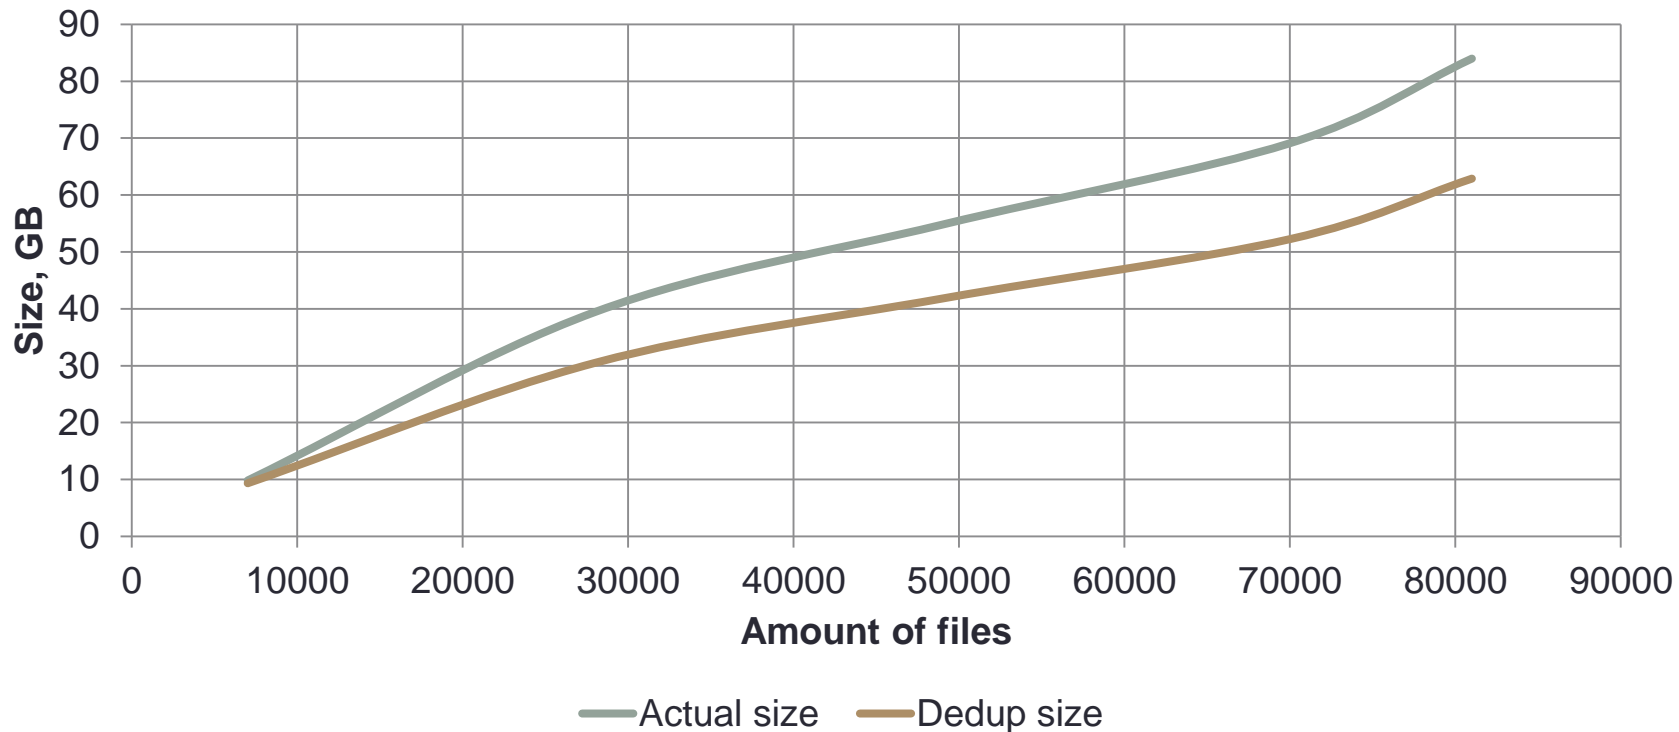
ZFS on a multiple hosts



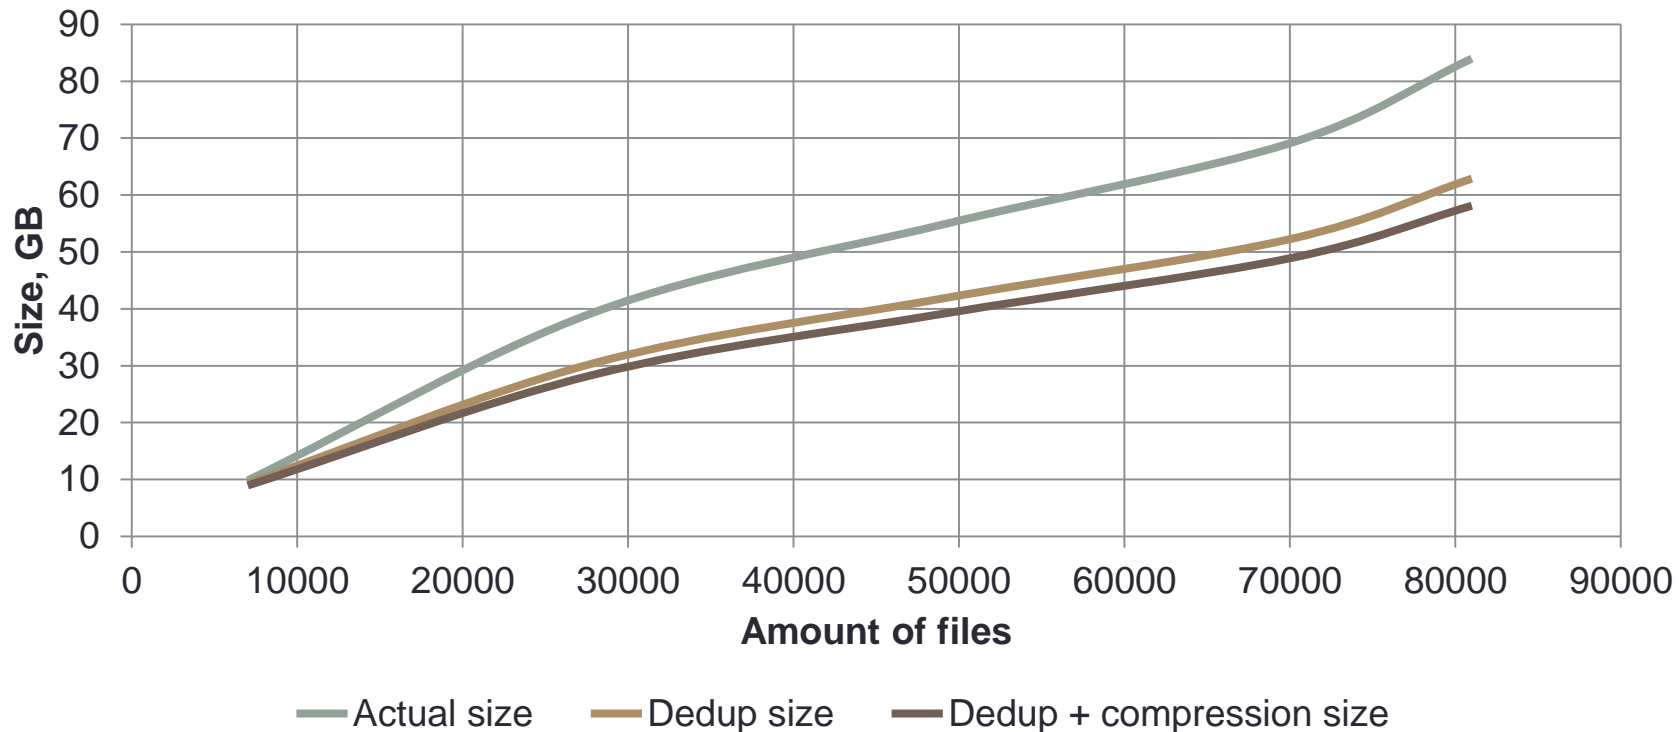
Attaching OpenStack Swift servers



Deduplication performance



Dedup + compression performance



Goods we've got

Common:

- Available on Linux, FreeBSD, Solaris

Brought by ZFS:

- Native data encryption support on Solaris
- Data deduplication
- Data compression

Brought by Swift:

- REST interfaces
- Multitenancy
- Versioning
- Replication
- Scalability

Disadvantages we have to deal with

- Each TB of 64K unique data blocks brings up to 5 GB of block database
- Direct file access procedure through the POSIX interface becomes a little bit complicated
- Internal network traffic overhead
- Swift performance loss

Possible ToDo's

- Usage of native ZFS clone and snapshot functions for file versioning and replication
- Check out Swift-On-File project
- Look for the future GlusterFS 4.0 Release

Thank You!