

S3 Compatibility: Enabling seamless integration with EOS, CERN's Large Scale Disk Storage System

Andreas-Joachim Peters, Elvin Alin Sindrilaru, Mano Segransan, Luca Mascetti

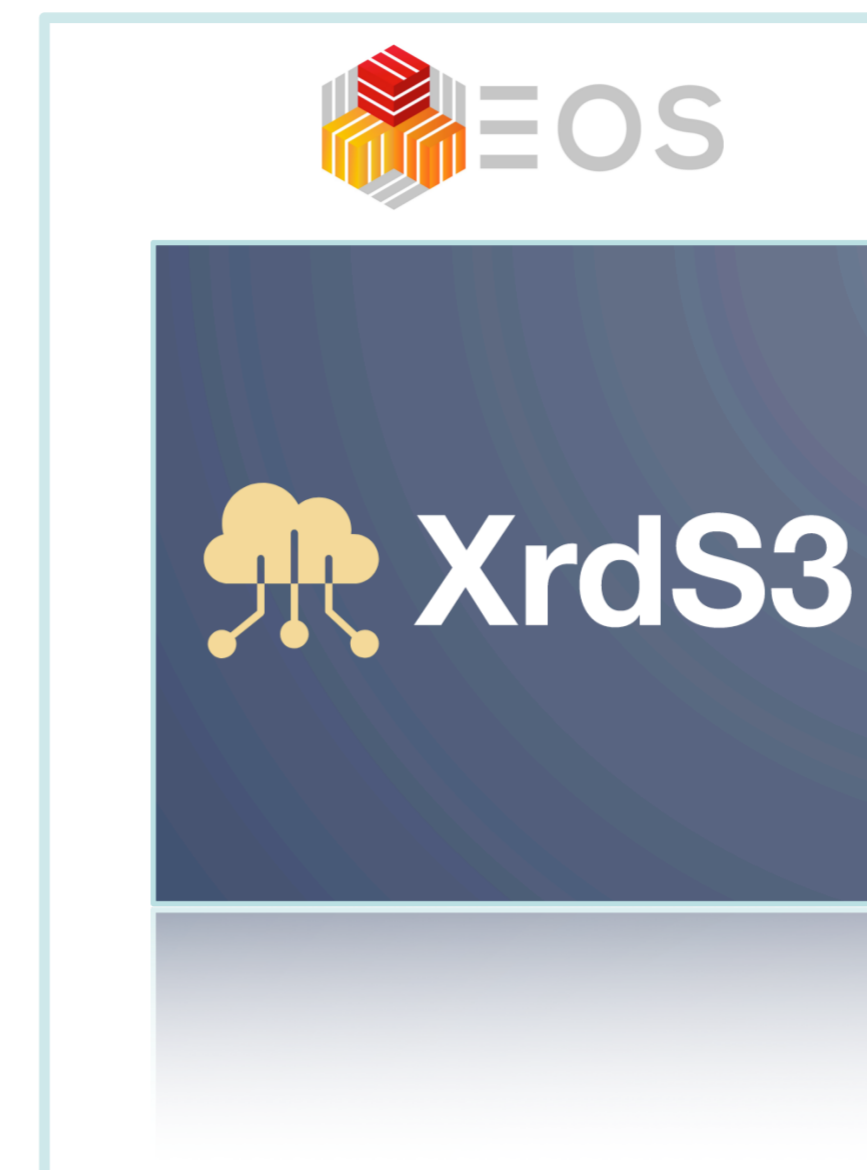
Presenter
Andreas-Joachim Peters
CERN - IT Storage



Why S3?




- **Amazon S3 leading object storage service** on the market
- scalability, data reliability
- security, performance
- **Widely used** for data lakes, websites, mobile applications, backup, archiving
- Supported as **de-facto standard protocol** by many tools
- support of S3 **enables many standard applications**



- **XrdS3** is server-side plug-in in XRootD framework
- unreleased branch XRootD
- **S3 protocol**
- S3 based on HTTP(S) verbs + header
- AWS v4 Signatures supported

EOS Open Storage - EB storage system at CERN providing access via XRootD, POSIX-like & HTTP protocols

Architecture



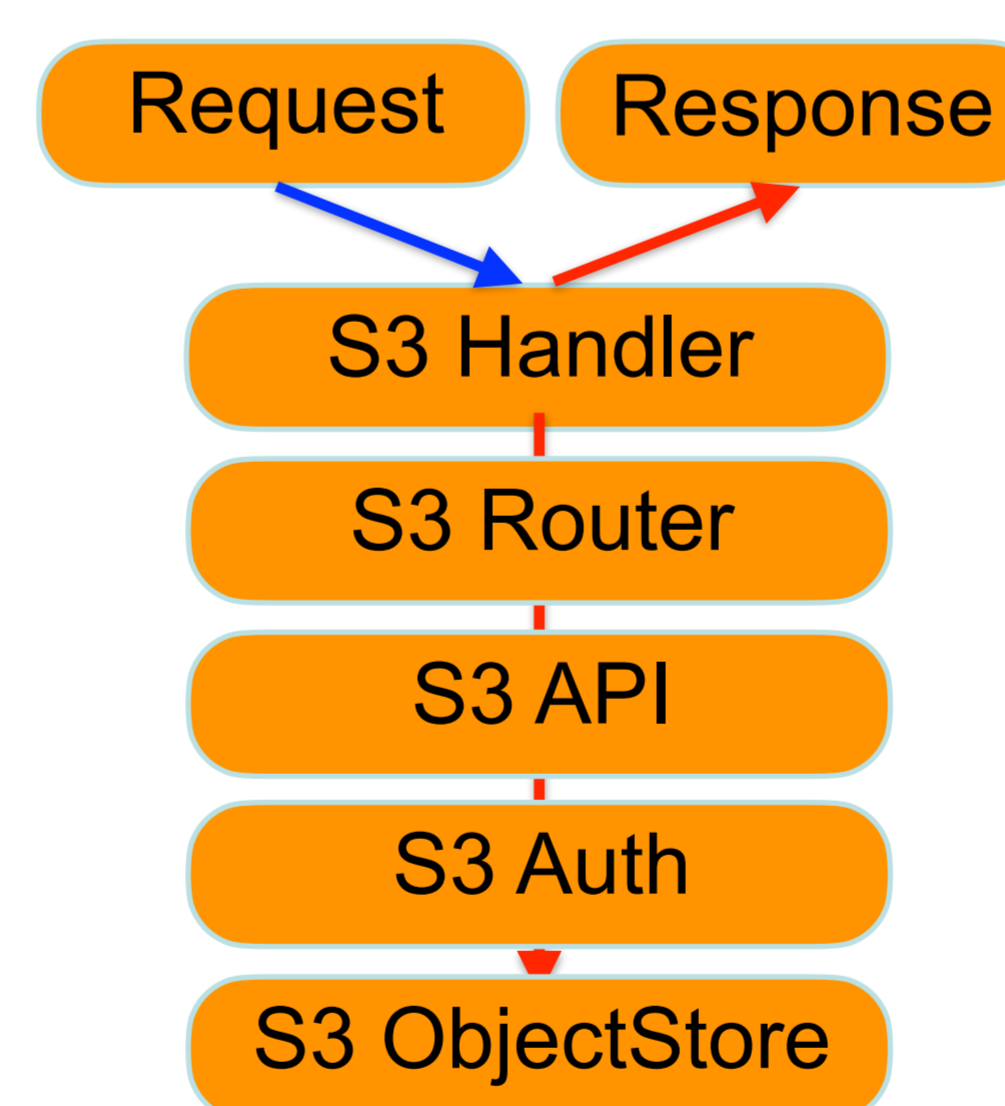
Current Prototype:
XrdS3 requires xattr API currently only available using FUSE



Design Implementation:
XrdS3 could use a single API using only XRootD protocol
- currently hybrid mode, xattr via FUSE, file IO via XRootD client

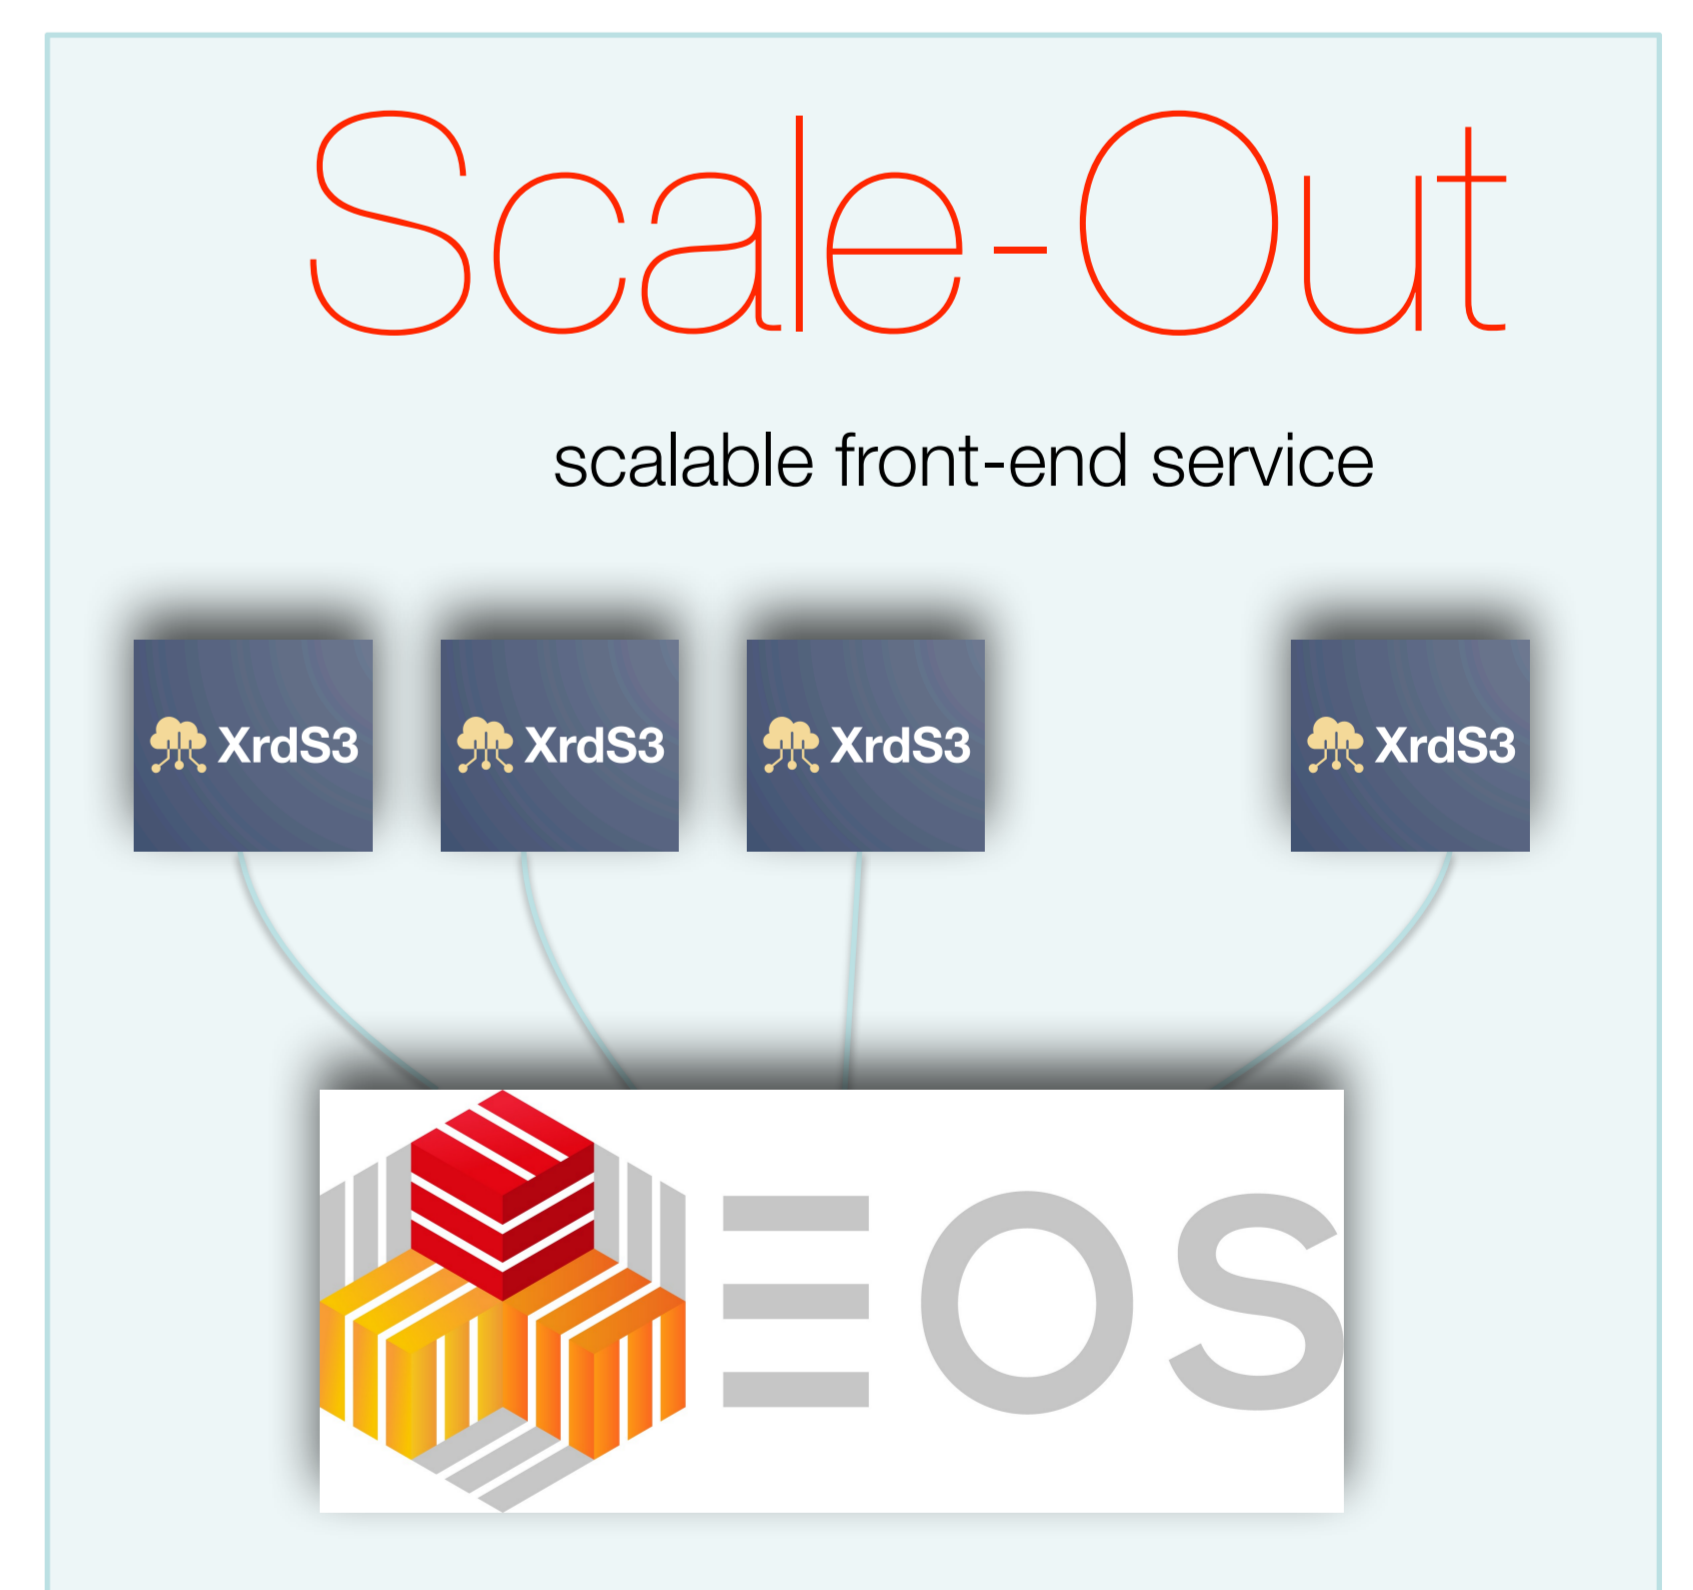
Software Design

- Modern & modular design with clear abstraction layers
- provides core implementation mapping functionality to filesystem
 - files are accessible with all EOS protocols
 - multi-part uploads result in complete files in EOS
 - optimisation for *good behaving* clients to do in-place uploads without the need to concatenate fragments
 - ownership mapped from S3 credential to filesystem (uid,gid)
 - easily extensible
 - future extension to various ObjectStore implementations possible



Scale-Out

scalable front-end service



Configuration

- currently stored in a shared directory between all gateways
- stores access keys & secrets
- maps EOS paths to bucket names

xrds3 - S3 Admin

usage: xrds3 [-h] {config,adduser,deleteuser,addbucket,deletebucket,ls} ...

positional arguments:

argument	description
{config,adduser,deleteuser,addbucket,deletebucket,ls}	
config	Configuration subcommand
adduser	Add a new user
deleteuser	Delete an existing user
addbucket	Add a new bucket
deletebucket	Delete an existing bucket
ls	List users/buckets

```
xrds3 ls --keys
Info: Listing matching keys:
-----
USER          ID          SECRET
-----
foo           bACuDCDS  1e422a60-eacc-4371-becf-3dcaa91473f6
monitops     n1Q3NA9W  22b684a6-bad2-474d-858a-7f3a034b8f3f
woujroba    W4V03WAM  55P00490-p95s-4149-0200-11303490831
100         P4V00CD2  1045500-090c-4317-09c1-130903490831
```

```
xrds3 ls
Info: Listing all user directories:
foo/
├── b-foo          foo          /eos/user/f/foo/          new: []
├── backup        foo          /eos/user/f/foo/s3/backup new: []
├── photos        foo          /eos/user/f/foo/photos   new: []
├── higgs         foo          /eos/cms/higgs/          new: []
├── web           foo          /eos/user/f/foo/www/     new: []
├── opendata      foo          /eos/public/opendata/    new: []
monitops/
├── monit        monitops    /eos/project/m/monit/s3/  new: []
├── kafka        monitops    /eos/project/k/kafaka/s3/ new: []
```

Outlook

- future multi-checksum support in EOS will provide multi-protocol access mixing S3, HTTPS, Posix & XRootD
- Glacier extensions possible to provide an S3 front-end for **CTA**

Contact: Andreas-Joachim.Peters@cern.ch

