




 Hugo Gonzalez Labrador



remember

Log in

CERNBOX Service



CERNBOX

- Owncloud 5 with NFS storage (currently deployed)
 - Desktop sync clients 1.6.4
 - iOS, Android branded mobile apps
- Owncloud 7 with EOS storage backend (currently migrating)
 - Desktop sync clients 1.7.0
 - iOS, Android official mobile apps



Looking for the [source code](#)?



CentOS



Debian



Fedora



openSUSE



SLE



Ubuntu



Looking for the [Android source code](#)?



CERNBOX/NFS deployment

- Setup 100% RH6 on “standard” hardware
- Guaranteed failover (redundant nodes)

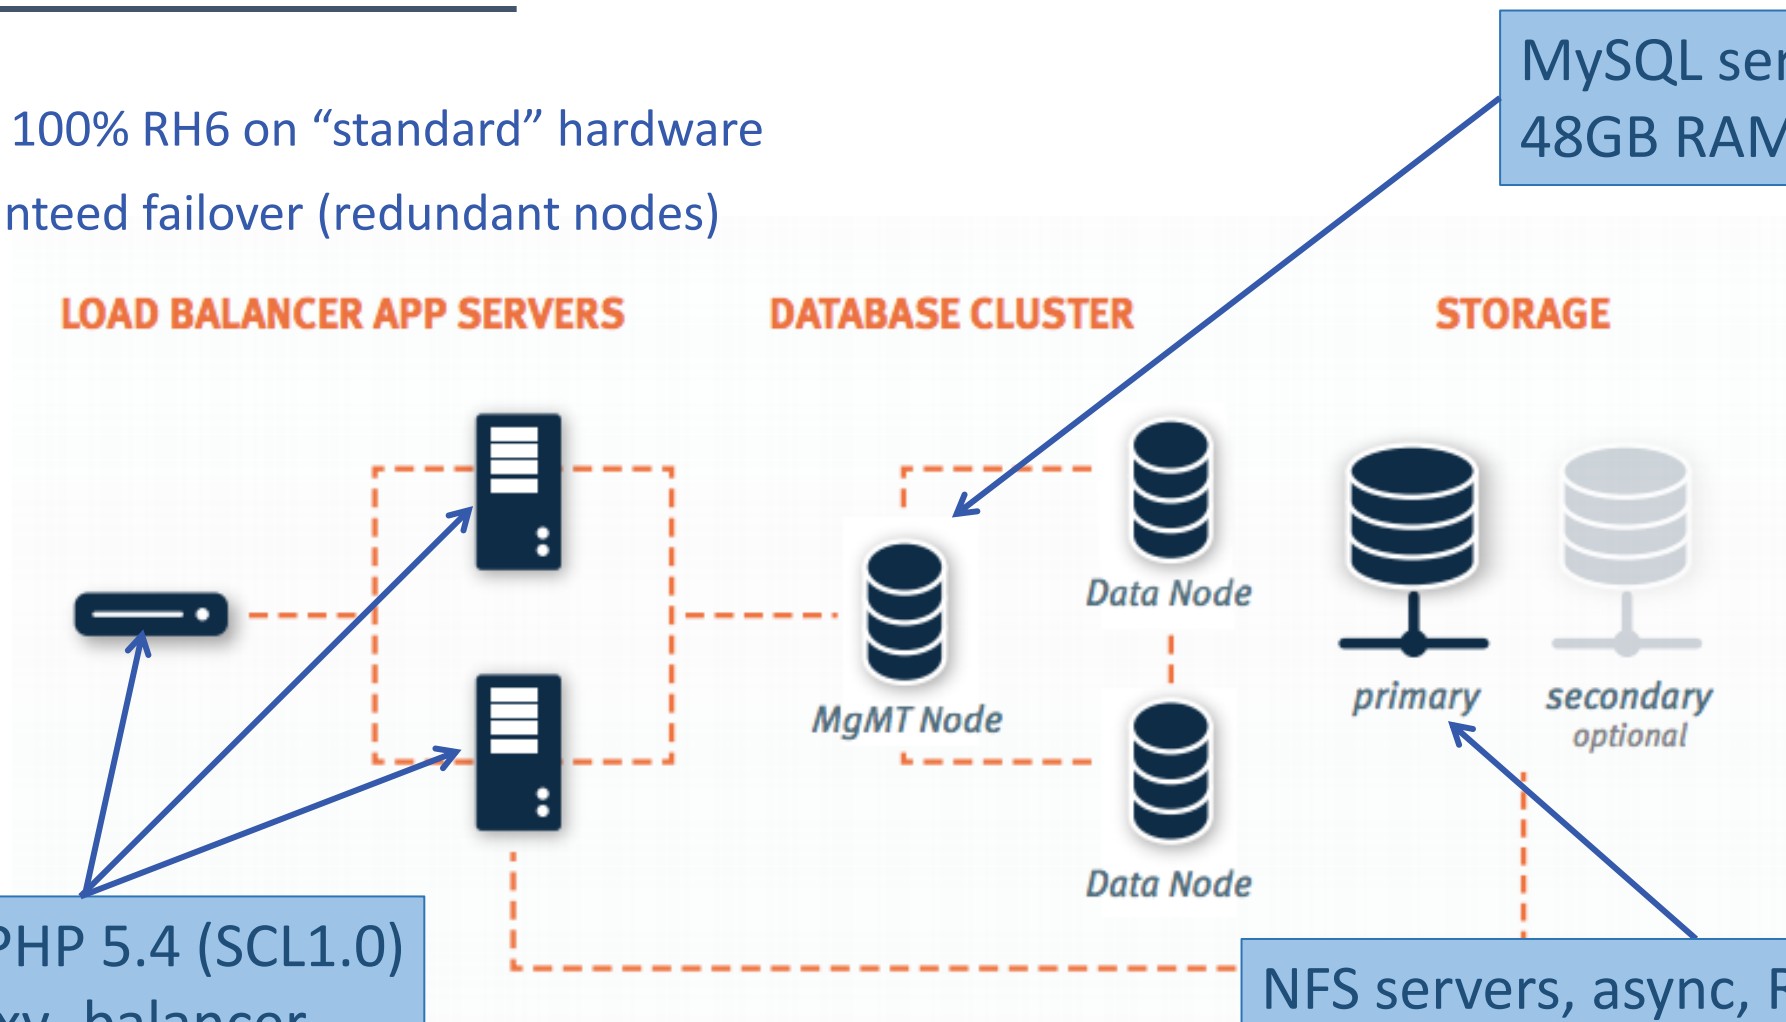


diagram source:
owncloud.com




Apache, PHP 5.4 (SCL1.0)
mod_proxy_balancer
64 core, 64GB RAM

NFS servers, async, RAID JBODs
Initial space: 20 TB

Usage of CERNBOX/NFS

CERNBox Beta 2014	March	April	May	June	...	October	November (until 15th)
users	190 (*)	285	361	429		805	1058
files	191K	907K	1.6M	2.7M		7.6M	9.6M
size	480GB	1TB	1.5TB	1.9TB		3.8TB	4.3TB


Size per user



Avg ~5GB

- <10GB
- >10GB
- up to 100GB

Files per user

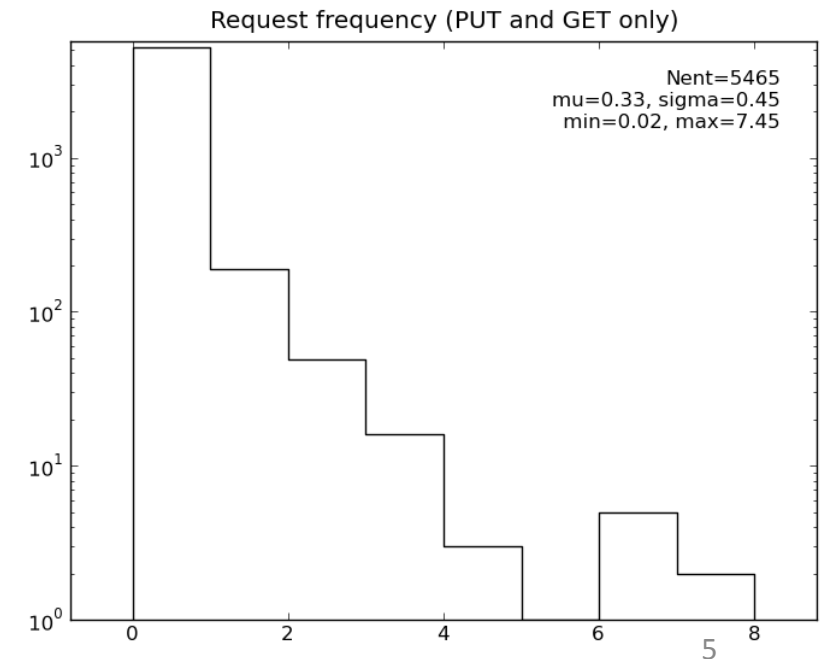
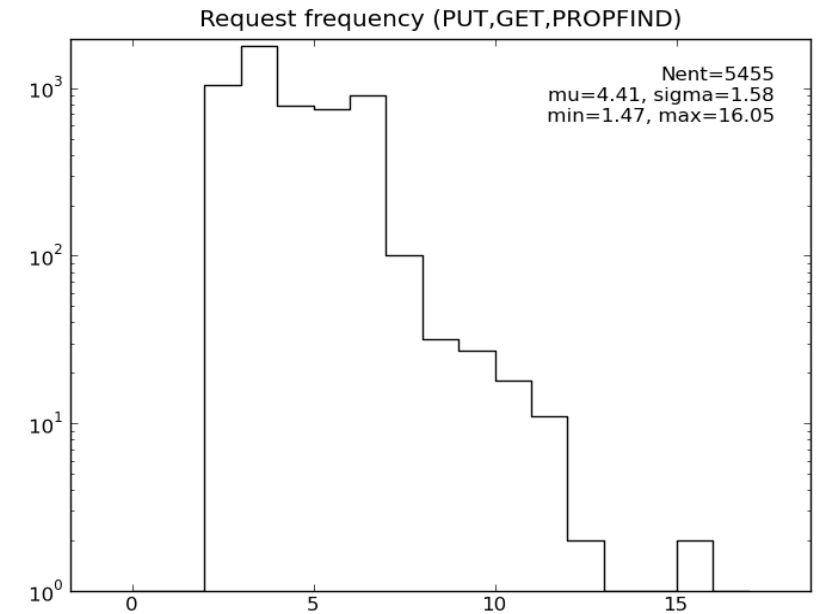




Avg ~10K files

- < 5K
- 5K-20K
- up to 100K

File access patterns

- GET/PUT ratio: 2/1
- File type distribution:
 - 1200 different file extensions!
 - 30% **.c .h .C**
 - 30% **.jpg .png**
 - 15% no extension (UNIX world!)
 - 25% other: **.pdf, .txt, .ppt, .docx, .root, .py, .eps, .tex**
- ~1104 shares
 - 833 via public link (only 48 password protected!)
 - 271 via internal share





Platform ratio



20%



60%



20%

Clients ratio



Mobile access

4%



Desktop access (sync clients)

70%



Web access

26%

It's time to change

NFS



EOS



#	instance	age	host	version	space	used	n(fs)	iops	bw-MB/s	files	directories
#	eosdev	8s	eosdevsrv1.cern.ch	0.3.55	313.23 TB	1.88%	174	9096	12962	4237	322
	dicos-eos02	13s	tw-eos02.grid.sinica.edu.tw	0.3.53	408.01 TB	0.18%	8	529	190	1253	1510
	dicos-eos03	6s	tw-eos03.grid.sinica.edu.tw	0.3.53	816.02 TB	0.19%	16	1074	2972	2566	2879
	eosalice	12s	chachkib.grid.unam.mx	0.3.35	504.08 TB	40.17%	8	561	930	13471527	45454
	eosalice	12s	lxbrf39c01.cern.ch	0.3.35	13.81 PB	72.27%	4474	266610	547200	115748659	93169
	eosatlas	6s	lxbse15c06.cern.ch	0.3.35	35.00 PB	41.83%	12280	702949	1445221	73178061	13670163
	eoscms	7s	lxbrf39c03.cern.ch	0.3.35	27.05 PB	39.83%	9258	566149	1075603	18600525	1597767
	eoscmsarchive	9s	p05508916c58033.cern.ch	0.3.52	44.99 TB	0.22%	15	700	1313	32	19
	eosdev	11s	pcitdss1400.cern.ch	0.3.72	373.07 GB	29.96%	2	0	0	35	35
	eosdevbox	1s	lxfsre13a02.cern.ch	0.3.70	13.85 TB	7.72%	7	494	1245	1805859	126929
	eos..lit.jinr.ru	4s	hydra.jinr.ru	0.3.53	7.87 TB	56.64%	1	78	648	178748	3483
	eoslhcb	4s	eoslhcbstv1.cern.ch	0.3.35	11.54 PB	58.73%	4005	206826	412740	5107824	4124004
	eospps	1s	p05508916e55509.cern.ch	0.3.72	1.13 PB	1.86%	354	18489	43531	4515003	1193660
	eospublic	14s	lxbrf41c07.cern.ch	0.3.35	15.98 PB	17.04%	4332	264305	603977	10474749	330738
	eossaske.sk	13s	eos-head-iep-grid.saske.sk	0.3.53	15.75 TB	2.81%	8	563	1308	215296	10295
	eossivvp.sk	15s	mgmt2.iep.sivvp.sk	0.3.53	22.00 TB	0.00%	1	70	1	1094	30
	eastest.cern.ch	14s	mgm.cern.ch	0.3.53	13.82 GB	100.00%	4	0	0	116	18
	eosuscmsdev	10s	cmsdev31.fnal.gov	0.3.53	436.24 GB	0.05%	1	54	134	88	23
	eosus...fnal.gov	10s	cmssrv222.fnal.gov	0.3.53	4.01 PB	69.12%	70	4696	31586	11360328	224020
	eosus...fnal.gov	8s	cmssrv153.fnal.gov	0.3.53	228.11 GB	4.19%	1	78	90	4300457	260843
	eosuser	12s	lxbst2146.cern.ch	0.3.68	95.31 TB	5.68%	48	3405	9171	4102750	284631
	hp-disk1	8s	hp-disk1.grid.sinica.edu.tw	0.3.35	640.01 TB	37.26%	12	740	8442	497585	182
#				111.40 PB		35079			263566792	21970174

CERNBOX/EOS Storage

- Deployed capacity in **EOSUSER** repo: **650TB** raw disk
 - => 2 replica on top of raid-1 (same as 4 replicas, hyper safe...)
- 8 machines (disk servers)
- Average file footprint 250B (in memory)
- 2 head nodes (1 master plus 1 slave with 48 GB RAM)
 - Migrating to 1 master and 1 slave with 256GB

What's new in CERNBOX/EOS ?

- Awesome **PERFORMANCE**
- **UNLIMITED** user quota
- **INTEGRATION** with **PETABYTE-SIZE** scientific repos
- **COLLABORATION** between different communities
- **NO METADATA IN SQL** database (oc_filecache)
 - => All metadata in EOS xattrs
 - => No rescan of file systems triggered
- **USER-PERMISSIONS** not OC permissions (apache:apache)
- **DIRECT MAPPING** between OC permissions and EOS acls

What's new in CERNBOX/EOS ?

The screenshot shows the CERNBOX/EOS file manager interface. The top navigation bar includes a search bar and the user 'labrador'. The left sidebar shows navigation options: 'All files', 'Shared with you', 'Shared with others', 'Shared by link', and 'Deleted files'. The main content area displays a list of files and folders. The 'uuup' folder is selected, and a share configuration dialog is open. The dialog shows the share name 'uuup' and a list of users with their permissions. The user 'gonzaleh' is highlighted with a red box, showing 'read' and 'write+delete' permissions checked. The 'Share link' option is unchecked.

Name	Size	Modified
KUBA		3 days ago
uuu		3 days ago
uuup		3 days ago
uuup	2 GB	3 days ago
uuup	2 GB	

Share with user or group ...

gonzaleh read write+delete

Share link

The screenshot shows an EOS Console terminal window. The terminal prompt is 'EOS Console [root://eosuser.cern.ch/] |/eos/user/l/labrador/>'. The command 'attr ls uuup' has been executed, and the output is displayed. The output is a list of attributes and their values for the 'uuup' directory. The attributes are: 'sys.acl', 'sys.allow.oc.sync', 'sys.forced.atomic', 'sys.forced.blockchecksum', 'sys.forced.blocksize', 'sys.forced.checksum', 'sys.forced.layout', 'sys.forced.maximumsize', 'sys.forced.nstripes', 'sys.forced.space', 'sys.mtime.propagation', 'sys.owner.auth', 'sys.recycle', and 'sys.versioning'. The values for 'sys.acl' and 'sys.owner.auth' are highlighted with red boxes.

```
EOS Console [root://eosuser.cern.ch/] |/eos/user/l/labrador/> attr ls uuup
sys.acl="u:gonzaleh:rwx+d"
sys.allow.oc.sync="1"
sys.forced.atomic="1"
sys.forced.blockchecksum="crc32c"
sys.forced.blocksize="4k"
sys.forced.checksum="adler"
sys.forced.layout="replica"
sys.forced.maximumsize="10000000000"
sys.forced.nstripes="2"
sys.forced.space="default"
sys.mtime.propagation="1"
sys.owner.auth="krb5:gonzaleh,http:gonzaleh,gsi:gonzaleh,unix:gonzaleh"
sys.recycle="/eos/user/proc/recycle/"
sys.versioning="10"
EOS Console [root://eosuser.cern.ch/] |/eos/user/l/labrador/>
```

What's new in CERNBOX/EOS ?

some_report.txt < 1 kB seconds ago

1 file

14 B

- ↓ 27 seconds ago Restore
- ↓ 30 seconds ago Restore
- ↓ 36 seconds ago Restore

labkode — root@cbox02:/tmp/1000 — ssh — 163x12

```
EOS Console [root://eosuser.cern.ch/] |/eos/user/l/labrador/> file versions some_report.txt
-rw-r----- 2 labrador it 0 Nov 17 16:00 1416236412.003a08b1
-rw-r----- 2 labrador it 7 Nov 17 16:00 1416236418.003a08b2
-rw-r----- 2 labrador it 9 Nov 17 16:00 1416236421.003a08b3
```





User feedback

“I start using the cernbox since I'm a heavy user of Dropbox and I recently reached the limit of free disk space (5Gb). For work it will be great to have at least 50Gb of personal space “

“I would like to have is a free client for Android, which should be much more stable. “

“I find the service perfect to be able to get always the files/sources/documents I need independently of the place and connection. ”



User feedback

- “On my Macbook Air I noticed that the battery was draining much faster than usual. I checked on the activity monitor and CERNBOX was consuming 80-95% of the total energy. “
- “What I would like to do in the future is to combine my private data like my photos for example on my home owncloud server, and my work data on the Cern owncloud server.”
- “I'm very glad that CERN has launched the service using the OwnCloud platform. “
- “I hope that you will be supporting this service officially soon!! ”

Risk factors

- Stability of the sync clients across diff versions
- Resiliency of the system with “exotic” failure modes
 - Risk to lose files!!!
 - Risk to corrupt data!!!
- Product evolution
 - Stability of the sync protocol (controlled evolution)
 - Bulletproof core functionality
 - Market evolution



Service summary

Status:	Beta
Number of users (current, target):	Current 1000 users, targeting 10 000 CERN users
Default and Maximum quota:	No quota
Linux/Mac/Win user ratio:	20/60/20
Desktop clients/Mobile Clients/Web access ratio:	4/70/26
Technology:	Owncloud 7 + EOS + No DB for metadata (EOS xattrs)
Target communities:	CERN staff (Engineers, physicists and administrative staff)
Integration in your current environment (examples):	Synchronizing batch data submitted to GRID to your laptop
Risk factors:	Previous slide
Most important functionality:	Robustness of the sync clients
Missing functionality (if any):	Shibboleth integration, end-to-end checksums, client journaling

Questions ?

BACKUP SLIDES



Example of integration


2. Data results synced automatically

CERNBOX/EOS



FUSE MOUNTED

GRID (~100 000 cores)



User laptop

1. User submit job to grid