



EOS productisation

CERN Openlab Technical Workshop
5-6th November 2015

Xavier ESPINAL
(CERN)

Vedran Vujinović, Mihajlo Zdravković, Branko Blagojević
(COMTRADE)



› **CERN's Large Scale disk-only Storage System**

› **140PB of disk space and 480M files**

› **Keystone of LHC data recording and data processing**

› **Supporting all analysis activities and the majority of the data reconstruction**

› **Massively used by LHC experiments, fixed target experiments and AMS**

Aggregate disk I/O rates




Directories and files creation rates



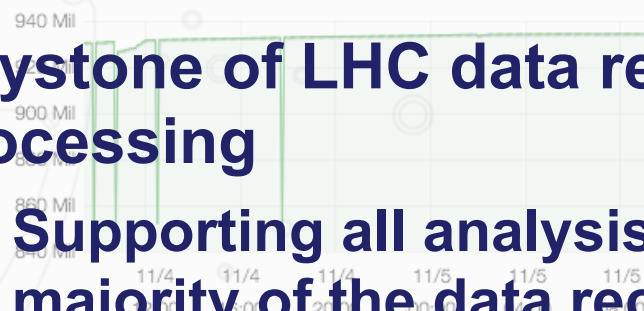
Files in the Namespace

480 Mil

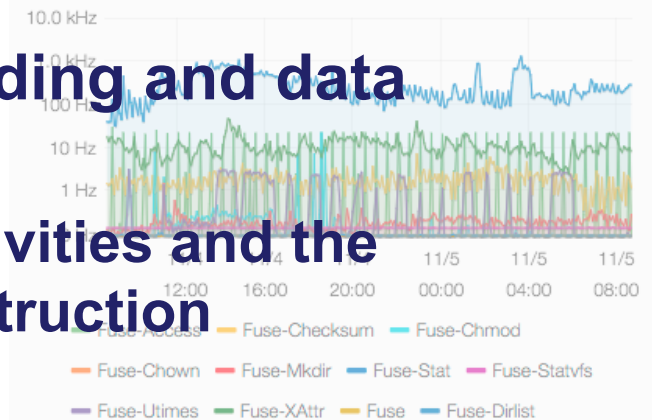
Files opened R/W




Used Files



Fuse operations



Network I/O rates



Current readers

43978

Current writers

10246

EOS productisation

“The project scope is the *evolution* of the EOS system in the direction of *simplified usage, installation and maintenance* and to extend its utilisation by adding new supported platforms. In the initial phase the emphasis will be in providing a robust *installation kit* to allow rapid installation of EOS on an agreed set of platforms. The kit will include the necessary *installation instructions* and *tools for operations* (admin guide) and for user (user guide). A test suite will exercise the native EOS interface (*xroot*) and the *main access protocols* (Fuse, Webdav/HTTP)”

“The project scope is the **evolution** of the EOS system in the direction of **simplified usage, installation and maintenance** of its installation by adding new supported platforms. In the initial phase the emphasis will be in providing a **installation kit** to allow rapid installation of EOS on an agreed set of platforms. The kit will include the new **installation instructions** **tools for operations** (admin guide) and for user (user guide). A test suite will exercise the native EOS interface **xroot** and **main access protocols** (e.g. Webdav/HTTP)”

evolution

simplified usage, installation and maintenance

installation kit

installation instructions

tools for operations

xroot

main access protocols

project status :: installation kit

- › **Streamlined installation of EOS instance:** ✓
 - › Single-node (headnode+storage nodes all-in-one)
 - › Full cluster (2 headnodes + multiple FSTs)
 - › Headnode master-slave configuration

- › **Post-install setup validation (put-stat-get)** ✓

- › **Installation summary** ✓

- › **Uninstallation** ✓

- › **On-failure-rollback** ✓

project status :: installation kit

```
./eossetup.sh EOS_comtrade --slave eosmgm2 --fst eosfst1 --numfs 4 --fst eosfst2,eosfst3 --numfs 6
```

```
# EOS Installation Summary
-----
MGM master=eosmgm1.eos.comtrade.co.yu mode=master-rw
MGM slave=eosmgm2.eos.comtrade.co.yu mode=slave-ro

number of FST nodes=3
space=default : total number of filesystems=16 total capacity=843.37 G
-----
# File Storage Nodes
-----
FST: eosfst1.eos.comtrade.co.yu: number of filesystems=4
status=online-on
heartbeat=3

FST: eosfst2.eos.comtrade.co.yu: number of filesystems=6
status=online-on
heartbeat=1

FST: eosfst3.eos.comtrade.co.yu: number of filesystems=6
status=online-on
heartbeat=2
-----
```

It's required that you open ports in your firewall:

On the MGM:

- 1094: XRootD MGM port
- 1096: XRootD SYNC port
- 1097: XRootD MQ port
- 8000: http port
- 8443: https KRB5 port
- 443: https X509 port

On the FST:

- 1095: XRootD FST port
- 8001: http port

```
[root@eosmgm1 ~]# eos ns
#
# Namespace Statistic
#-----
ALL Files 5 [booted] (0s)
ALL Directories 10
#-----
ALL Compactification status=off waitstart=0 interval=0 ratio-file=0.0:1 ratio-dir=0.0:1
#-----
ALL Replication mode=master-rw state=master-rw master=eosmgm1.eos.comtrade.co.yu conf
igdir=/var/eos/config/eosmgm1.eos.comtrade.co.yu/ config=default active=true mgm=eosmgm2.eos.comtrade.co.yu=ok
mgm:mode=slave-ro mq:eosmgm2.eos.comtrade.co.yu:1097=ok
#-----
ALL File Changelog Size 580 B
ALL Dir Changelog Size 1100 B
#-----
ALL avg. File Entry Size 116 B
ALL avg. Dir Entry Size 110 B
#-----
ALL memory virtual 750.99 MB
ALL memory resident 178.80 MB
ALL memory share 8.24 MB
ALL memory growths 363.07 MB
ALL threads 74
ALL uptime 628
#-----
```




```
[root@eosmgm1 ~]# eos fs ls
#-----
# host (#...) # id # path # schedgroup # geotag # boot # configstatus # drain # active
#-----
...t1.eos.comtrade.co.yu (1095) 1 /var/eos/fs/0 default.0 booted rw nodrain online
...t1.eos.comtrade.co.yu (1095) 2 /var/eos/fs/1 default.1 booted rw nodrain online
...t1.eos.comtrade.co.yu (1095) 3 /var/eos/fs/2 default.2 booted rw nodrain online
...t1.eos.comtrade.co.yu (1095) 4 /var/eos/fs/3 default.3 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 5 /var/eos/fs/0 default.0 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 6 /var/eos/fs/1 default.1 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 7 /var/eos/fs/2 default.2 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 8 /var/eos/fs/3 default.3 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 9 /var/eos/fs/4 default.4 booted rw nodrain online
...t2.eos.comtrade.co.yu (1095) 10 /var/eos/fs/5 default.5 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 11 /var/eos/fs/0 default.0 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 12 /var/eos/fs/1 default.1 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 13 /var/eos/fs/2 default.2 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 14 /var/eos/fs/3 default.3 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 15 /var/eos/fs/4 default.4 booted rw nodrain online
...t3.eos.comtrade.co.yu (1095) 16 /var/eos/fs/5 default.5 booted rw nodrain online
#-----
```

```
[root@eosmgm1 ~]# eos node ls
#-----
# type # hostport # geotag # status # status # txgw #gw-queued # gw-ntx #gw-rate # heartbeatdelta #nofs
#-----
nodesview eosfst1.eos.comtrade.co.yu:1095 online on off 0 10 120 0 4
nodesview eosfst2.eos.comtrade.co.yu:1095 online on off 0 10 120 1 6
nodesview eosfst3.eos.comtrade.co.yu:1095 online on off 0 10 120 2 6
#-----
```

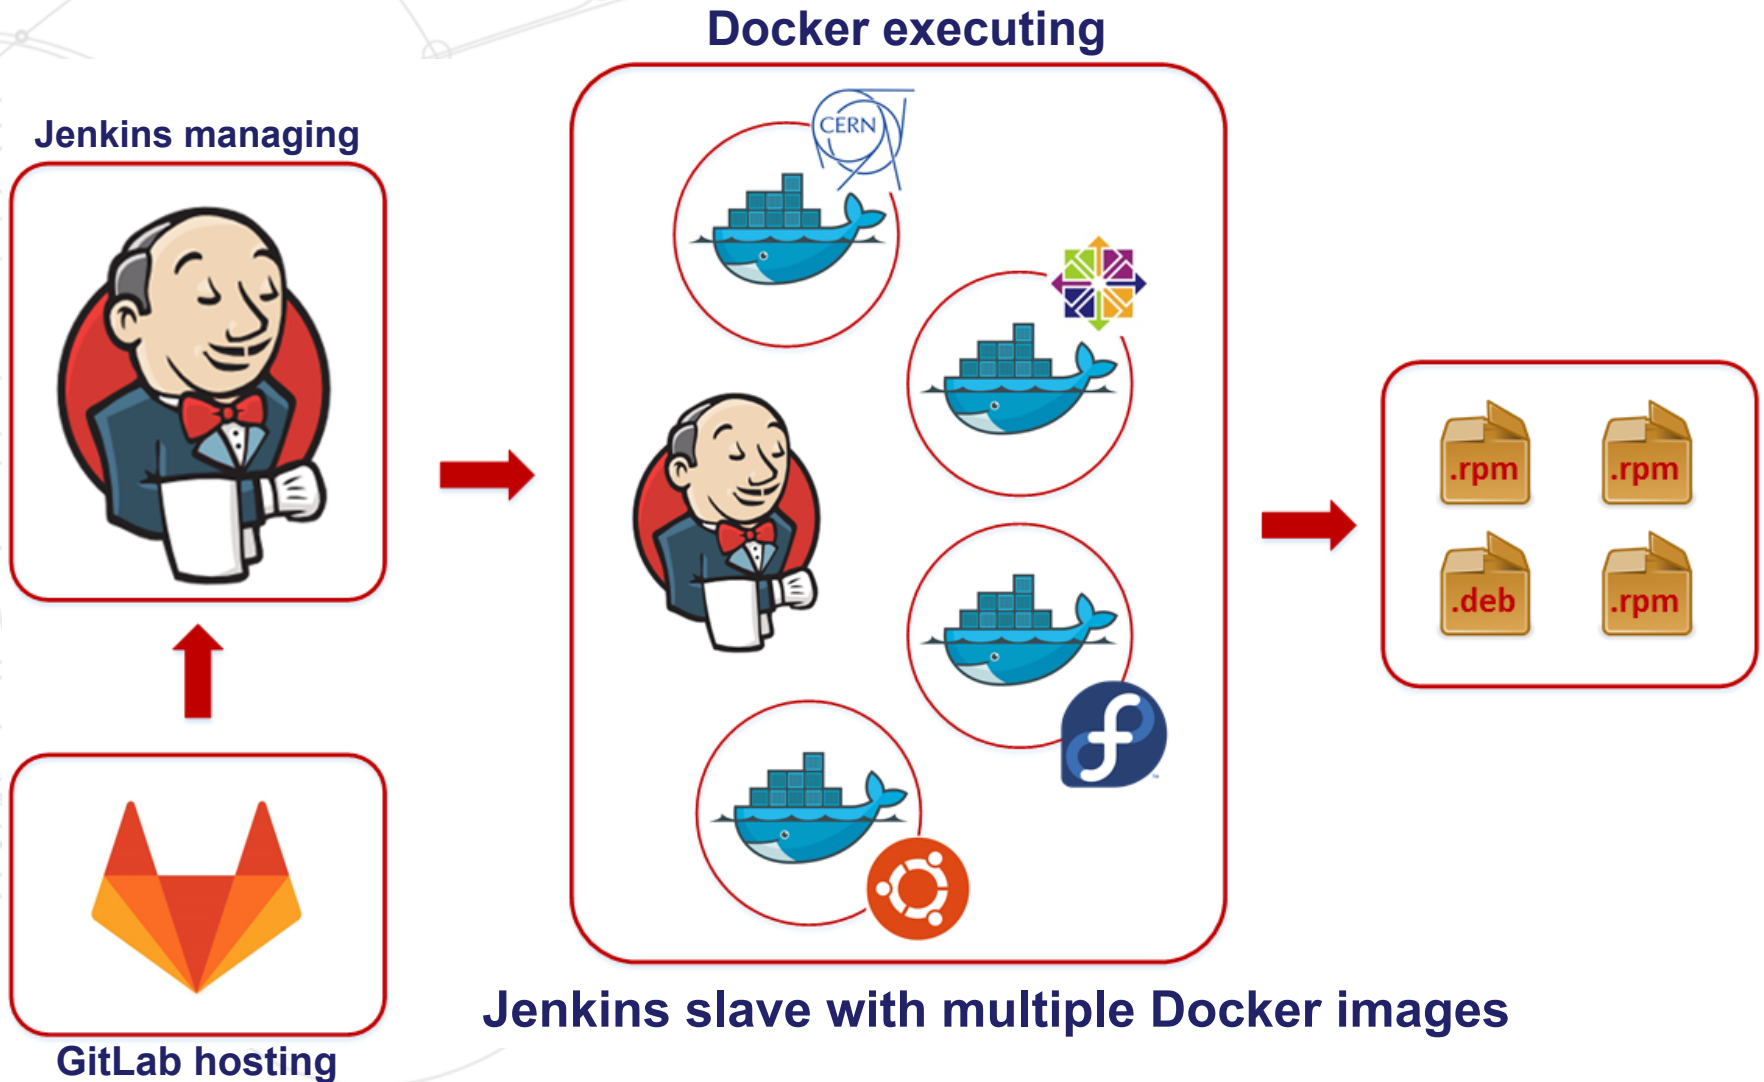
project status :: installation kit

- › **Work in progress:**
 - › **Defining and testing various setup scenarios**
 - › **Extending Linux distribution support**
 - › **CentOS, Fedora, RHEL, SuSe, Ubuntu/Debian, etc.**
 - › **Configuration management systems support**
 - › **Puppet, Chef, Ansible cookbooks**
 - › **Offer Docker based installation possibility**

project status :: documentation

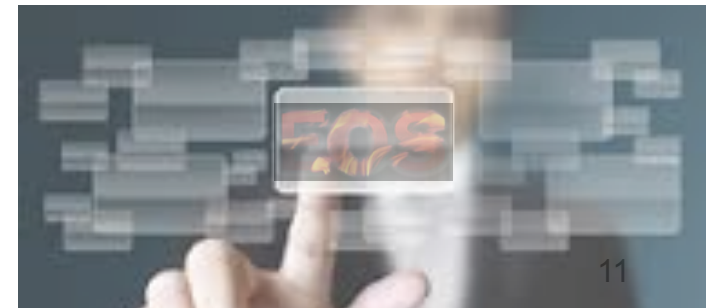

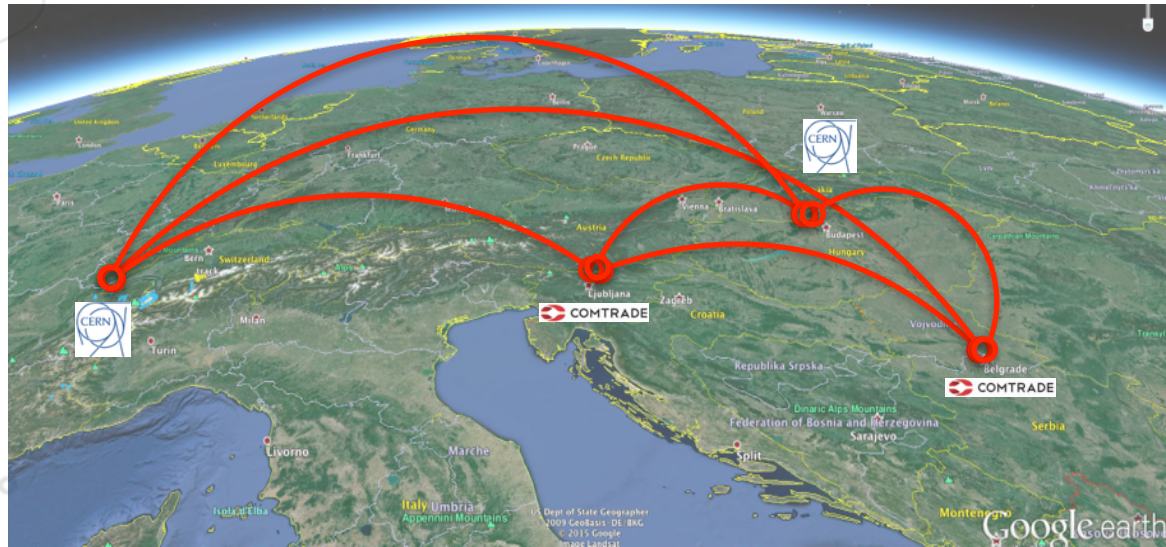
- › **Simple usage manual** 
- › **Gathering all relevant information** ✓
- › **High-level EOS overview** ✓
- › **EOS functional units** ✓
- › **EOS functionalities uncovered** 
- › **EOS whitepaper** 

project status :: build environment



project status :: next steps

- › **Multi-site EOS cluster**
 - › **Geo-scheduling**
 - › **RAIN configuration**
 - › **Erasure coding**
- › **Documentation**
 - › **EOS overview, user, administration guides**
- › **Automated testing framework**
- › **EOS Administration/Monitoring GUI**
- › **Native Microsoft Windows client support**



› **Rapid progress during the first phase**

› **Excellent collaboration**

› **Exciting future - looking forward to the 2nd year**