

# EOS storage at SIVVP SASKE and Heterogeneous LIT JINR clusters

Martin Vaľa<sup>1,2</sup>, Ján Astaloš<sup>4</sup>, Lucia Vaľová<sup>1,3</sup>, Dmitry Belikov<sup>1</sup>

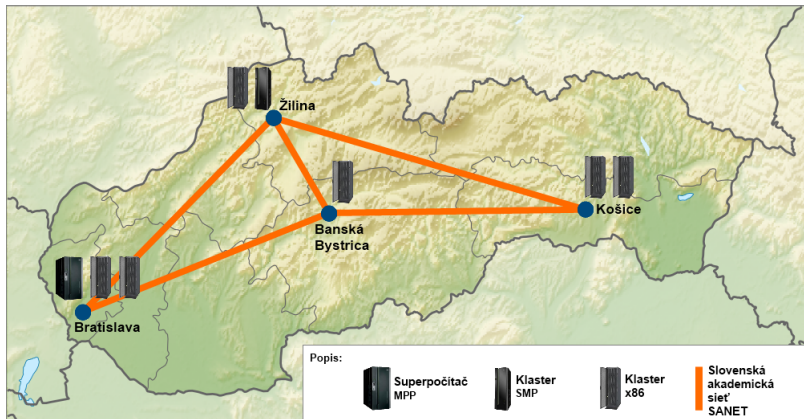
- 1 - JINR : Joint Institute for Nuclear Research
- 2 - SASKE : Slovak Academy of Science at Košice, Slovakia
- 3 - TU : Technical University at Košice, Slovakia
- 4 - Institute of Informatics, Slovak Academy of Sciences

July 2, 2014

- SIVVP SASKE cluster
- Heterogeneous LIT JINR cluster
- EOS storage
  - Introduction
  - Features
  - Deployment
  - Data analysis flow for user

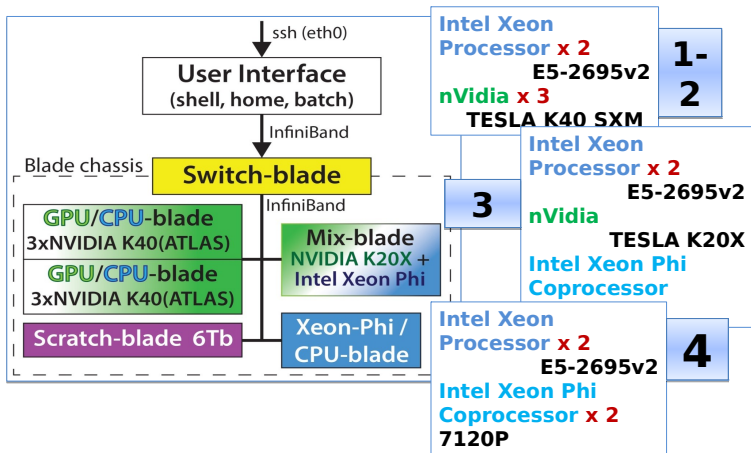
- Hybrid cluster
  - OpenMP
  - OpenMPI
  - OpenCL
  - CUDA
  - PGI Fortran
- SLURM queuing system
- ROOT for Data Analysis
- EOS for federation storage
- Plan for future is to combine these clusters

# SIVVP project structure



- System: Server IBM System x iDataPlex dx360 M3
- CPU cores: 552
- GPU: 192
- Memory/core: 4GB
- Storage: 48TB
- OS: Scientific Linux
- Network: 40GB/s infiniband
- <http://www.sivvp.sk/>

## Heterogeneous cluster (LIT JINR)



# Heterogeneous LIT JINR clusters

- CPU cores: 120
- Xeon Phi cores : 122
- CUDA cores : 19968
- Memory/core: 4.4GB
- Storage: 28 TB
- OS: Scientific Linux
- Network: 40GB/s infiniband

- Kerberos a LDAP
- LXPLUS like system
- Indico
- Redmine
- Nagios
- EOS - Storage system (idea)



- XROOTD storage element
- Fast access to meta-data (without database)
- Web interface
- File access via `http://` and `root://`
- Authentication via KRB5, certificates a SSS
- CERN experiments are migrating to EOS
- Git: <http://eos.cern.ch/cgi-bin/cgit.cgi/eos/>
- Doc: <https://eos.readthedocs.org/en/latest/>

- Space quota
- Disk grouping
- Scratch directory setup possible with cache functionality
- Disk draining
- Data balancing
  - hot data
  - data locality (JINR and SASKE)
- EOS console for users and admins
- FUSE mount
- Support for RW via root:// protocol

# EOS support and future releases features

- EOS support
  - Thanks CERN EOS team
  - Specialy to Andreas-Joachim Peters
- EOS feature in future
  - Local cache for FUSE (similar as AFS or CVMFS)
  - Owncloud support
  - Support for utf8
  - Archive support: Freeze a subtree in the EOS namespace and migrate into archive storage with the possiblity to recall it back into the namespace.

- master/slave configuration
- 8 data servers - 2TB each = 16 TB
- SASKE KRB5 Authentication
- SIVVP extension (soon)
  - 24 TB
  - 11 TB visible from workers only
- ALICE and ATLAS storage disks in Kosice (next year)

- master only configuration
- Data servers
  - 1 file server with 14.4 TB
  - 1 file server with 14.4 TB visible from workers only
- JINR KRB5 Authentication
- CERN experiment federation (idea)
- Storage for NICA experiment (in future)

- eos console (many linux-like commands supported)
  - list files
  - find files
  - share file
  - change permissions
  - dropbox
  - multiple copy of file
- eos web interface
- eos fuse mount

# EOS user data analysis flow [1]

```
[root@lx000 ~]$ cat /etc/yum.repos.d/eos.repo
[eos-beryll]
name=EOS 0.3 Version
baseurl=http://eos.cern.ch/rpms/eos-beryll/slc-6-x86_64/
gpgcheck=0
enabled=1
[root@lx000 ~]# yum install eos-client eos-fuse
Loaded plugins: priorities, protectbase, security
1635 packages excluded due to repository priority protections
0 packages excluded due to repository protections
Setting up Install Process
Package eos-client-0.3.35-1.x86_64 already installed and latest version
Package eos-fuse-0.3.35-1.x86_64 already installed and latest version
Nothing to do
[root@lx000 ~]#
```

# EOS user data analysis flow [2]

```
[mvala@vala ~]$ kinit
Password for mvala@SASKE.SK:
[mvala@vala ~]$ export EOS_MGM_URL="root://eos.saske.sk"
[mvala@vala ~]$ eos
# -----
# EOS Copyright (C) 2011 CERN/Switzerland
# This program comes with ABSOLUTELY NO WARRANTY; for details type
# 'license'.
# This is free software, and you are welcome to redistribute it
# under certain conditions; type 'license' for details.
# -----
EOS_INSTANCE=eossaske.sk
EOS_SERVER_VERSION=0.3.35 EOS_SERVER_RELEASE=1
EOS_CLIENT_VERSION=0.3.35 EOS_CLIENT_RELEASE=1
EOS Console [root://eos.saske.sk] |/> whoami
Virtual Identity: uid=10481 (10481,99) gid=10481 (10481,99) [authz:krb5]
host=159.93.32.4
EOS Console [root://eos.saske.sk] |/>
[mvala@vala ~]$ eos -b whoami
Virtual Identity: uid=10481 (10481,99) gid=10481 (10481,99) [authz:krb5]
sudo* host=159.93.32.4
```



# EOS user data analysis flow [3]

```
[mvala@vala ~]$ eos -b ls -al /eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/
drwxrwsr+ 14 mvala  alice          13 Jan 30 09:56 .
drwxrwsr+  4 mvala  alice           3 Aug 20 09:00 ..
drwxrwsr+  2 mvala  alice           1 Jan 30 09:56 000146746
drwxrwsr+  2 mvala  alice           1 Jan 30 09:56 000146747
drwxrwsr+  2 mvala  alice           1 Jan 30 09:56 000146748
drwxrwsr+  2 mvala  alice           1 Jan 30 09:56 000146801
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146802
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146803
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146805
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146806
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146807
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146817
drwxrwsr+  2 mvala  alice           1 Jan 30 09:57 000146824
drwxrwsr+  2 mvala  alice           1 Jan 30 09:58 000146856
drwxrwsr+  2 mvala  alice           1 Jan 30 09:58 Merged
```

# EOS user data analysis flow [4]

```
[mvala@vala ~]$ eos -b find -f /eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/  
Merged/All/STD2010/DCAZ10/ | grep .root  
/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCn  
sig10/RsnOutput.root  
/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCn  
sig20/RsnOutput.root  
/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCn  
sig30/RsnOutput.root  
/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCn  
sig40/RsnOutput.root  
/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/quali  
tyonly/RsnOutput.root
```

# EOS user data analysis flow [5]

```
[mvala@vala ~]$ eos -b find -f --online --partition /eos/saske.sk/alice/rsn/PHIKK/  
LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DCAZ10/ | grep .root  
path=/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DC  
AZ10/KTPCnsig10/RsnOutput.root partition=n55-iep-grid.saske.sk:/data/01/eos-data  
path=/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DC  
AZ10/KTPCnsig20/RsnOutput.root partition=n54-iep-grid.saske.sk:/data/01/eos-data  
path=/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DC  
AZ10/KTPCnsig30/RsnOutput.root partition=n60-iep-grid.saske.sk:/data/01/eos-data  
path=/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DC  
AZ10/KTPCnsig40/RsnOutput.root partition=n60-iep-grid.saske.sk:/data/01/eos-data  
path=/eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/RSN_20140125/Merged/All/STD2010/DC  
AZ10/qualityonly/RsnOutput.root partition=n61-iep-grid.saske.sk:/data/01/eos-data
```

# EOS user data analysis flow [6]

```
[mvala@vala ~]$ eos -b find -f --partition /eos/saske.sk/users/m/mvala/Test/  
/aaa.txt partition=n57-iep-grid.saske.sk:/data/01/eos-data,n61-iep-grid.sas  
ke.sk:/data/01/eos-data  
/test.txt partition=n57-iep-grid.saske.sk:/data/01/eos-data,n59-iep-grid.sa  
ske.sk:/data/01/eos-data
```

# EOS user data analysis flow [7]

```
[mvala@vala ~]$ root -b
root [0] TFile::Open("root://eos.saske.sk//eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/ESD/
RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCnsig10/RsnOutput.root")
(class TFile*)0x1dc7ec0
[mvala@vala ~]$ wget -q http://eos.saske.sk:8000//eos/saske.sk/alice/rsn/PHIKK/LHC12f1b/E
SD/RSN_20140125/Merged/All/STD2010/DCAZ10/KTPCnsig10/RsnOutput.root
[mvala@vala ~]$ ls -al RsnOutput.root
-rw-rw-r--. 1 mvala mvala 21M Feb 25 16:56 RsnOutput.root
```

# EOS user data analysis flow [8]

```
[mvala@vala ~]$ export EOS_MGM_URL="root://eos.saske.sk"
[mvala@vala ~]$ mkdir ~/eos
[mvala@vala ~]$ eos fuse mount ~/eos/
OK
===> Mountpoint      : /home/mvala/eos/
===> Fuse-Options    : kernel_cache,attr_timeout=30,entry_timeout=30,max_readahead=131072,
max_write=4194304,fsname=eos.saske.sk root://eos.saske.sk/eos/
===> xrootd ra       : 131072
===> xrootd cache    : 393216
===> fuse debug      : 0
===> fuse write-cache : 1
===> fuse write-cache-size : 100000000

[mvala@vala ~]$ cat ~/eos/saske.sk/proc/whoami
Virtual Identity: uid=99 (99) gid=99 (99) [authz:unix] sudo* host=159.93.32.4

[mvala@vala images]$ ls -al ~/eos/saske.sk/users/m/mvala/Test/
total 0
drwx-----.  3 10481 10481  0 Mar 11 17:50 .
drwx-----. 14 10481 10481 12 Apr  9 2013 ..
```

# EOS user data analysis flow [9]

```
[mvala@vala ~]$ export EOS_MGM_URL="root://eos.saske.sk"
[mvala@vala ~]$ mkdir ~/eos
[mvala@vala ~]$ eos fuse mount ~/eos/
OK
===> Mountpoint      : /home/mvala/eos/
===> Fuse-Options    : kernel_cache,attr_timeout=30,entry_timeout=30,max_readahead=131072,max_w
root://eos.saske.sk/eos/
===> xrootd ra       : 131072
===> xrootd cache    : 393216
===> fuse debug      : 0
===> fuse write-cache : 1
===> fuse write-cache-size : 100000000

[mvala@vala ~]$ cat ~/eos/saske.sk/proc/whoami
Virtual Identity: uid=10481 (10481,99) gid=10481 (10481,99) [authz:krb5]
host=159.93.32.4

[mvala@vala ~]$ ls -al ~/eos/saske.sk/users/m/mvala/Test/
total 0
drwx-----.  3 10481 10481  0 Mar 11 17:50 .
drwx-----. 14 10481 10481 12 Apr  9 2013 ..
-rw-r--r--.  2 10481 10481 16 Jun  9 14:45 aaa.txt
-rw-r--r--.  2 10481 10481  8 Mar 11 17:53 test.txt
```

- We would like to thank project Slovak infrastructure for high performance computing (SIVVP), ITMS 26230120002.
- Institute of Informatics, Slovak Academy of Sciences
  - Ján Astaloš
- Technical University support
  - Lucia Vaľová and Ján Jadlovský
- Slovak Academy of Science at Košice
  - Miloslav Straka, Ingrid Kuľková, Viktor Kočan, Richard Bilek, Marcela Brasová , Tibor Kožár and Martin Vaľa



- Heterogeneous calculations team
  - Petr Zrelov, Dmitry Podgainy, Evgeny Aleksandrov, Alexander Ayriyan, Dmitry Belyakov, Adam Gheorghe, Dmitry Golub, Margarita Kirakosyan, Mikhail Matveyev, Alexey Prikhodko, Tatiana Sapozhnikova, Oksana Streltsova, Tatiana Strizh, Shushanik Torosyan, Martin Vaľa, Lucia Vaľová, Tatiana Zaikina, Elena Zemlyanaya and Maxim Zuev

Thank you