

IHEP Site Status

Qiulan Huang

huangql@ihep.ac.cn

Computing Center, IHEP, CAS

HEPIX FALL 2015

Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Monitoring and IT Service
- Summary

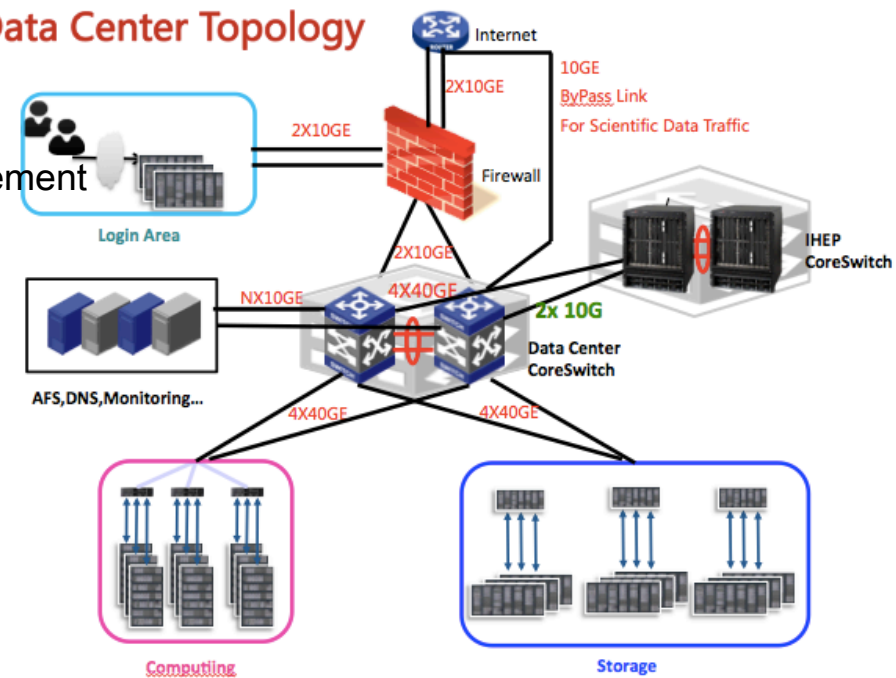


Infrastructure Update

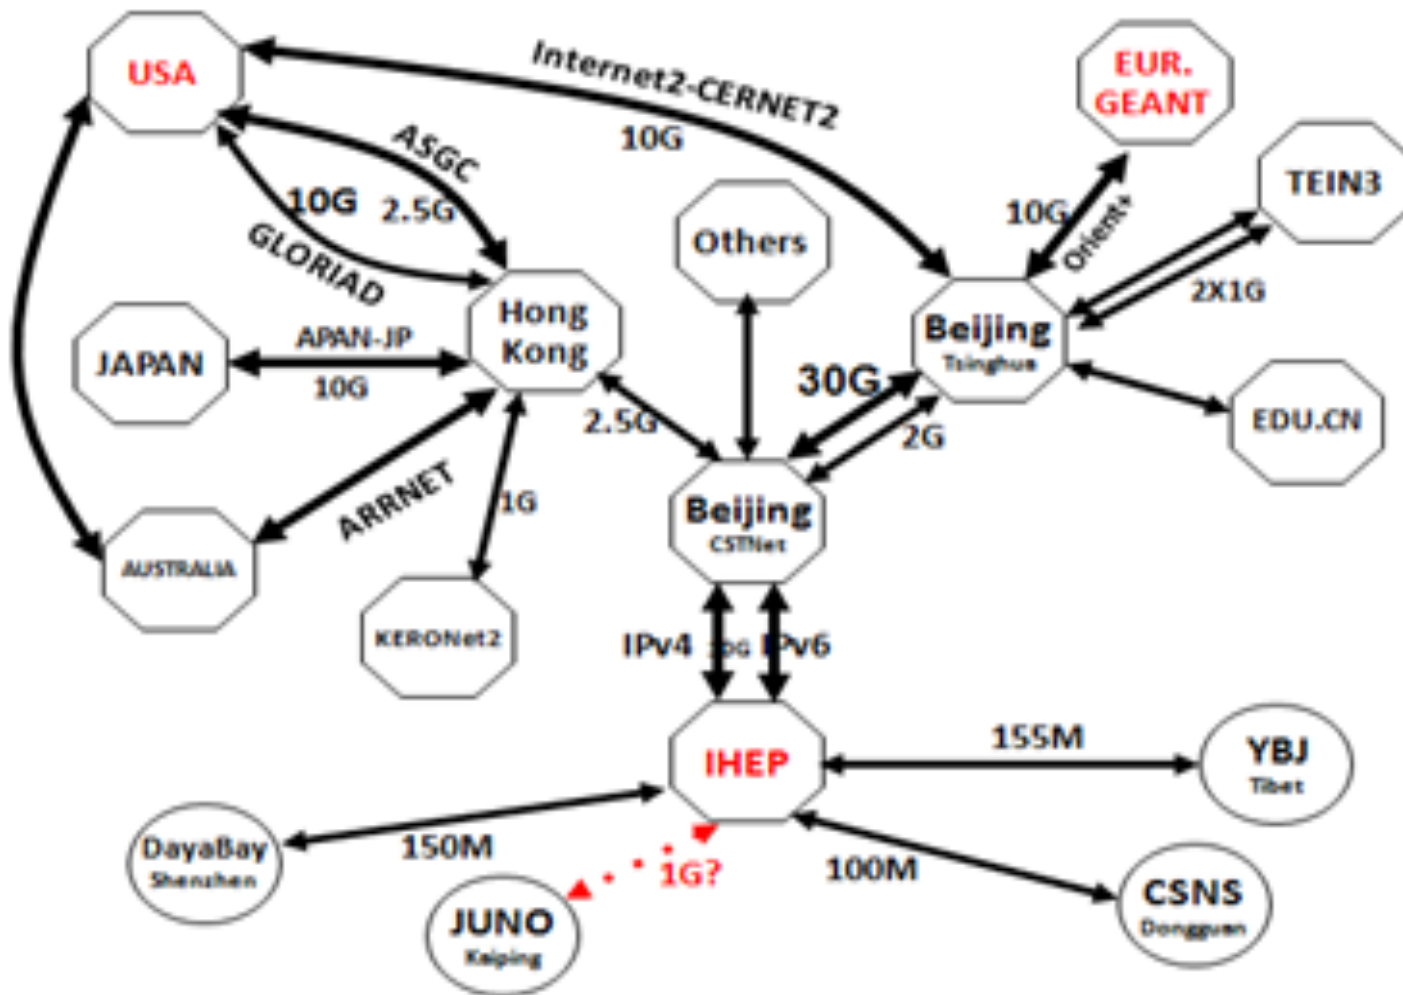
Local cluster

- Data Storage: **4.5PB -> ~5.3PB**
 - Retired **760 TB** disk servers and disk arrays which are 5+ years old
 - Added **~1.5PB**
- New Network for DataCenter
 - Double Core switches & Firewalls
 - High Capability , Stability, Easy Management
 - Two networks
 - Data Network
 - » High Stability & Capability
 - Management Network
 - » High Stability

Data Center Topology



Internet and Domestic Network



- IHEP-EUR.:
• **10Gbps**
- IHEP-USA:
• **10Gbps**
 - From the end of 2013, switch to the link (IHEP-USA), good performance
- IHEP-Univ.:
• **30Gbps**
- Performance improved
 - RTT reduced **10%**
 - Data transfer performance improved **+40%**



Infrastructure Upgrade

- EGI Site: Hardware upgrade
 - Storage: **940TB**(Atlas 400TB, CMS 540TB)
 - 24 Slot diskarray with 4TB disk
 - One box has to array, Each array has 12 disk with raid6
 - Really capacity 80TB per box.
 - Disk server:
 - 2 HP 380g9 and 1 Dell R620 (atlas)
 - 2 HP 380g9 and 2 Dell R620 (cms)
 - **New Worknode: 29 HP blade bl460c gen9**
 - CPU: 2 E5-2680 V3 12 core
 - MEM: 64GB
 - DISK: 1TB SATA



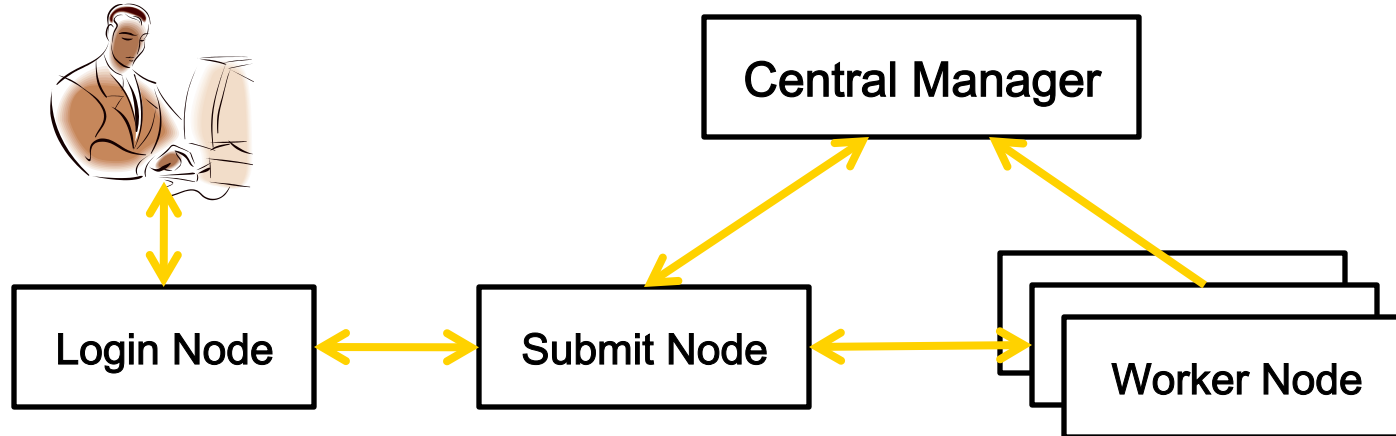
Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Monitoring and IT Service
- Summary



Status of HTCondor



- Version: 8.2.5 at present, 8.4.x very soon
- Submit jobs via login node in remote mode



- Resource: **1080 CPU** cores
- Plan: **10,000+** CPU cores (local cluster), and several remote sites

Resource Restriction

- 2 separated experiments (JUNO and CMS) in 1 pool
- A wrapper of condor_submit guarantees:
 - ✓ “*accounting_group*” is set by user, or else a default value is assigned
 - ✓ the user is a valid user in the “*accounting_group*”
- At the worker node side:




- JUNO nodes run JUNO jobs, CMS nodes run CMS jobs ...



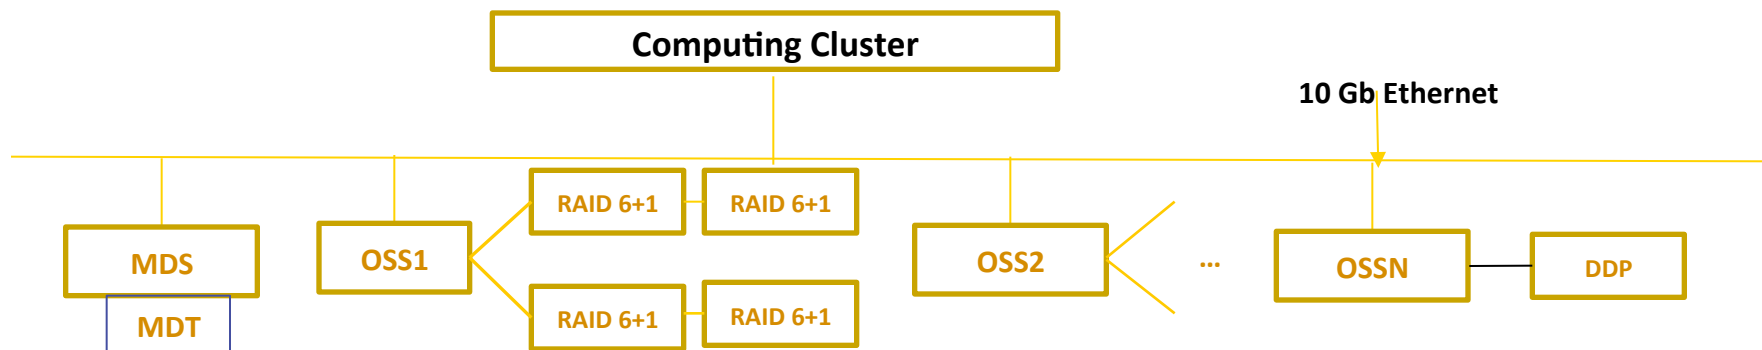
Monitoring

- Integrated with the IHEP cluster monitoring system



Lustre Status

- Total capacity: **4.5PB**, 8 instances, average usage: 67%, max 1.2 PB, min 14 TB, avg 560 TB, 52 OSSs



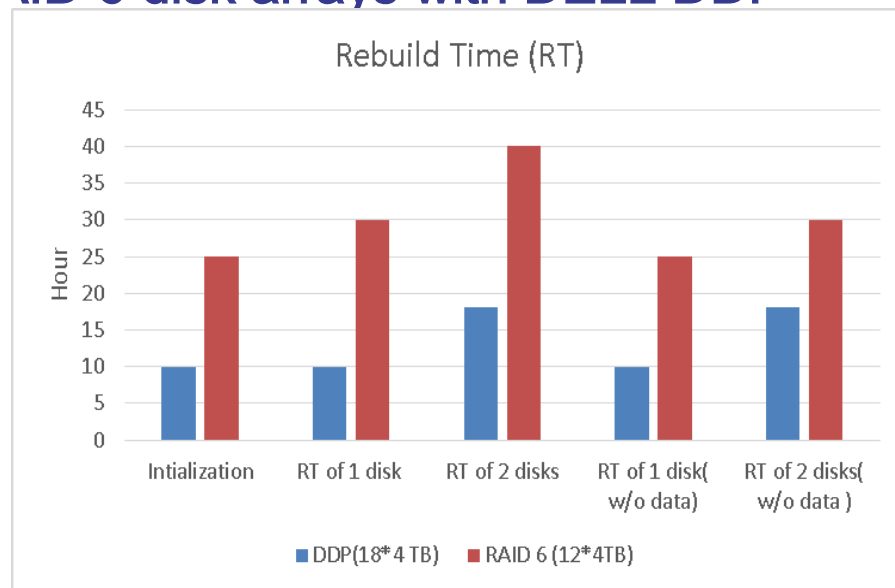
- Simple Architecture: disk arrays directly attached to OSS (Object Storage Server), OSS exports disk storage as Lustre objects
- Average disk space per OSS ~100TB, Read throughput per OSS 800 MB/s



Upgrade of Lustre

- **Hardware:** Replaced traditional RAID 6 disk arrays with DELL DDP

- ✓ rebuild time of 4 TB SATA disk largely reduced
- ✓ risk of simultaneous disk failure reduced
- ✓ Similar technologies such as HUAWEI RAID 2.0 is also under consideration



- **Software**
 - ✓ Lustre servers: Lustre 1.8.8 → Lustre 2.5.3
 - ✓ SL 6 computing nodes: Lustre-client-1.8.9 → Lustre-client-2.5.3
- **Old data is unaffected** and some management CLI changed slightly
- The earliest upgraded instance has been stably running for more than **6 months**

MDS stuck issue

- **MDS stuck by large scale lock timeouts (>600s)**

LustreError: 0:0:(ldlm_lockd.c:344:waiting_locks_callback()) ### lock callback timer expired after 12300s: evicting client at 192.168.61.200@tcp ns: mdt-dybf-MDT0000_UUID lock: ffff8807fd2f5b80/0x6b0381f861a50e6a lrc: 3/0,0 mode: PR/PR res: [0x166b0002:0x3efd81b4:0x0].0 bits 0x13 rrc: 323 type: IBT flags: 0x60200400000020 nid: 192.168.61.200@tcp remote: 0xddf7a5419061b23 expref: 7 pid: 13370 timeout: 5854515000 lvb_type: 0

- **Error -107: Client evictions caused client reconnections frequently**

- **Temporary solution**

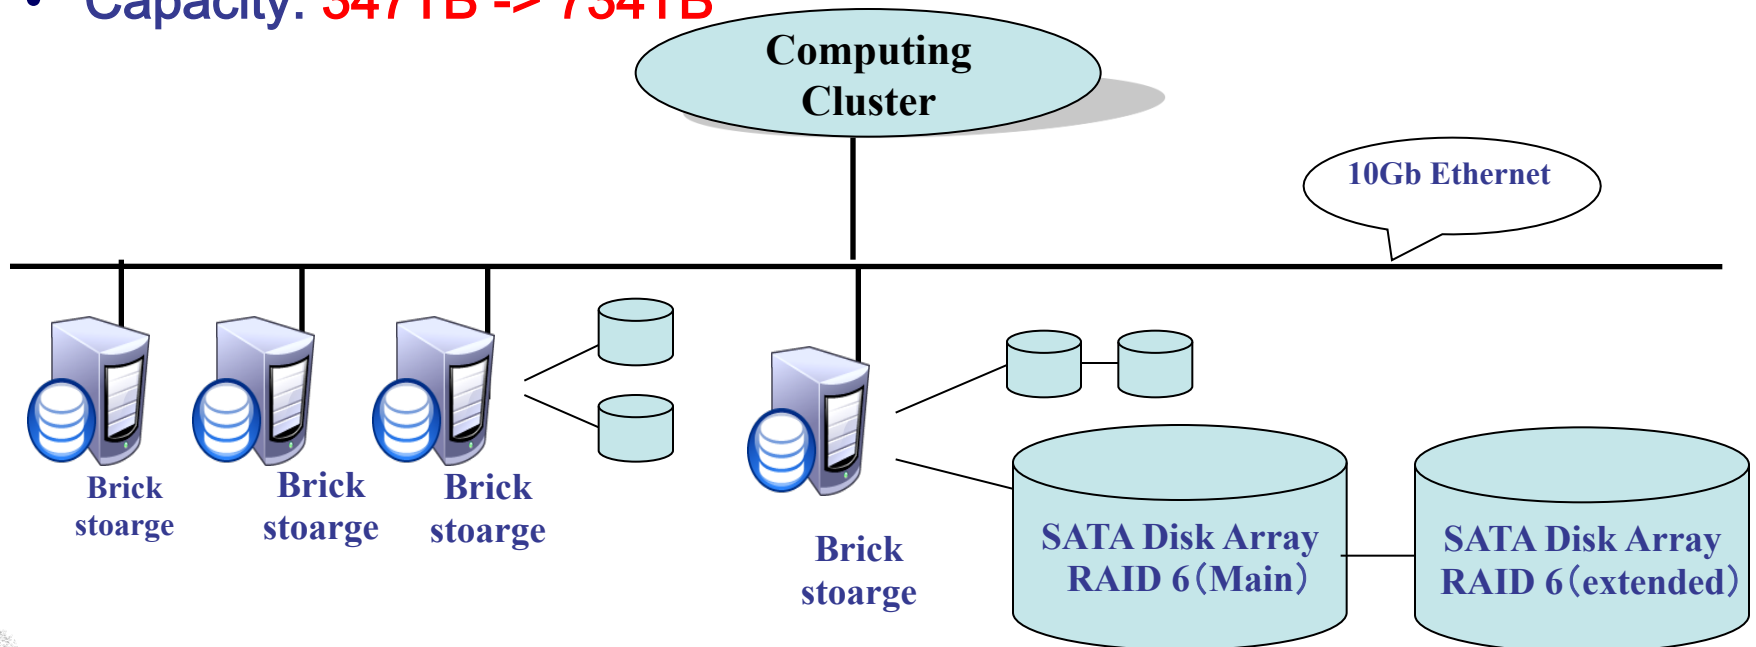
- Remount of MDT to resume the MDS service

- **Welcome any comments**



Gluster Status

- Total capacity: **734TB**, currently 2 instances
- Deployed 3 years old disk arrays(780TB) to a replication volume
- Scale upgrade
 - I/O server: **5 -> 13**
 - Bricks: **43 -> 107**
 - Capacity: **347TB -> 734TB**



AFS Status

- Total capacity: **19TB**
- 2 file servers were added with 9.2TB effective space
 - OS Kernel: SL 6.5 64bit
 - OpenAFS version: 1.6.11
- Upgrade one DB server, which serves more than **6** years
- Add monitoring to check replica consistence in Nagios



IHEPCloud status

- Released in 18 November, 2014
- Built on Openstack Icehouse
- 1 control node
- 20 computing nodes
 - 232 physical CPU cores / 464 Virtual CPU cores, 1.7TB memory, 10.1TB storage
- Active 229 VMs, 1TB memory
- >1000 users, >70 active users



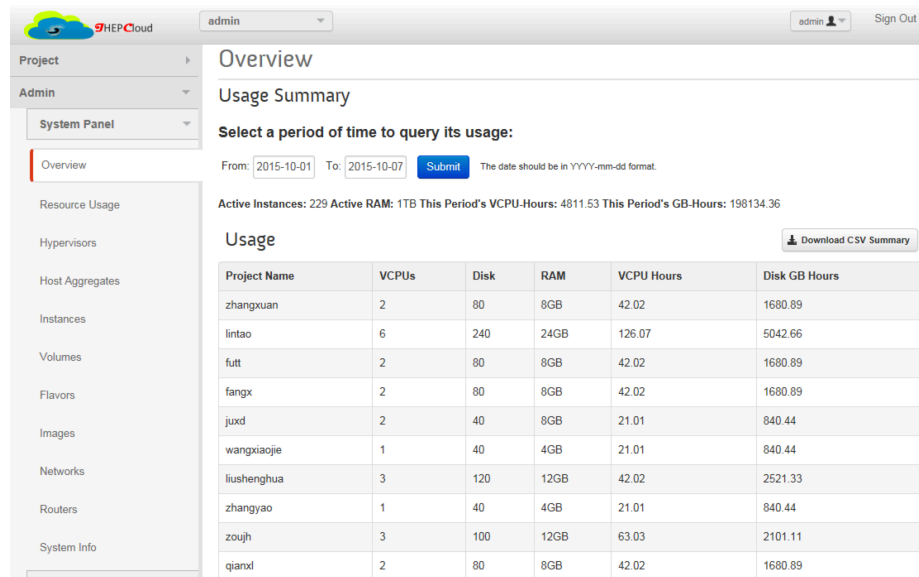
VCPU Usage
Used 284 of 464



Memory Usage
Used 1.0TB of 1.7TB



Disk Usage
Used 9.4TB of 10.1TB



SDN@IHEP

- Goals

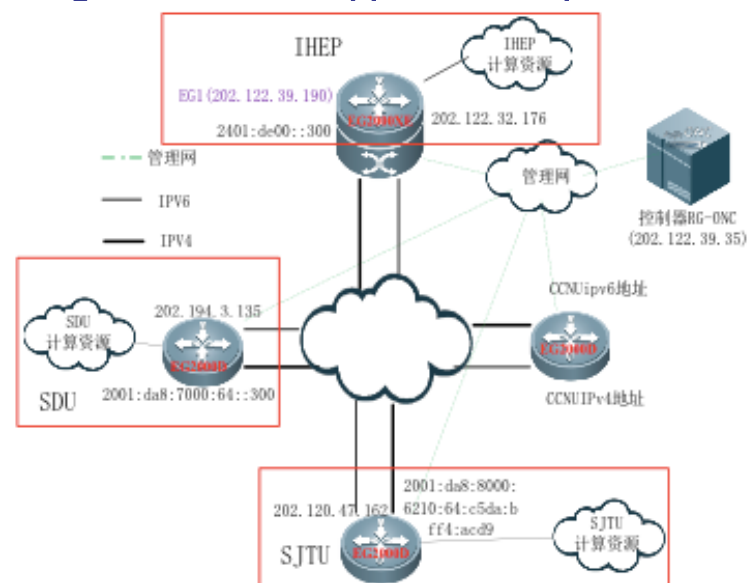
- Improve the data exchange performance around China, based on the current
- Network infrastructure Applications

- Thoughts

- Overlay: use IPv4 & IPv6 network link
- Automatically and Dynamically network path choosing based on the application requirements and network performance status

- Current Status

- **4 Sites Deployed:** SJTU,SDU,IHEP,CCNU
- IPv6 & IPv4 Supported
- Path Choosing automatically
- Stable & Good performance for the stress test
- More performance data will be released soon



Outline

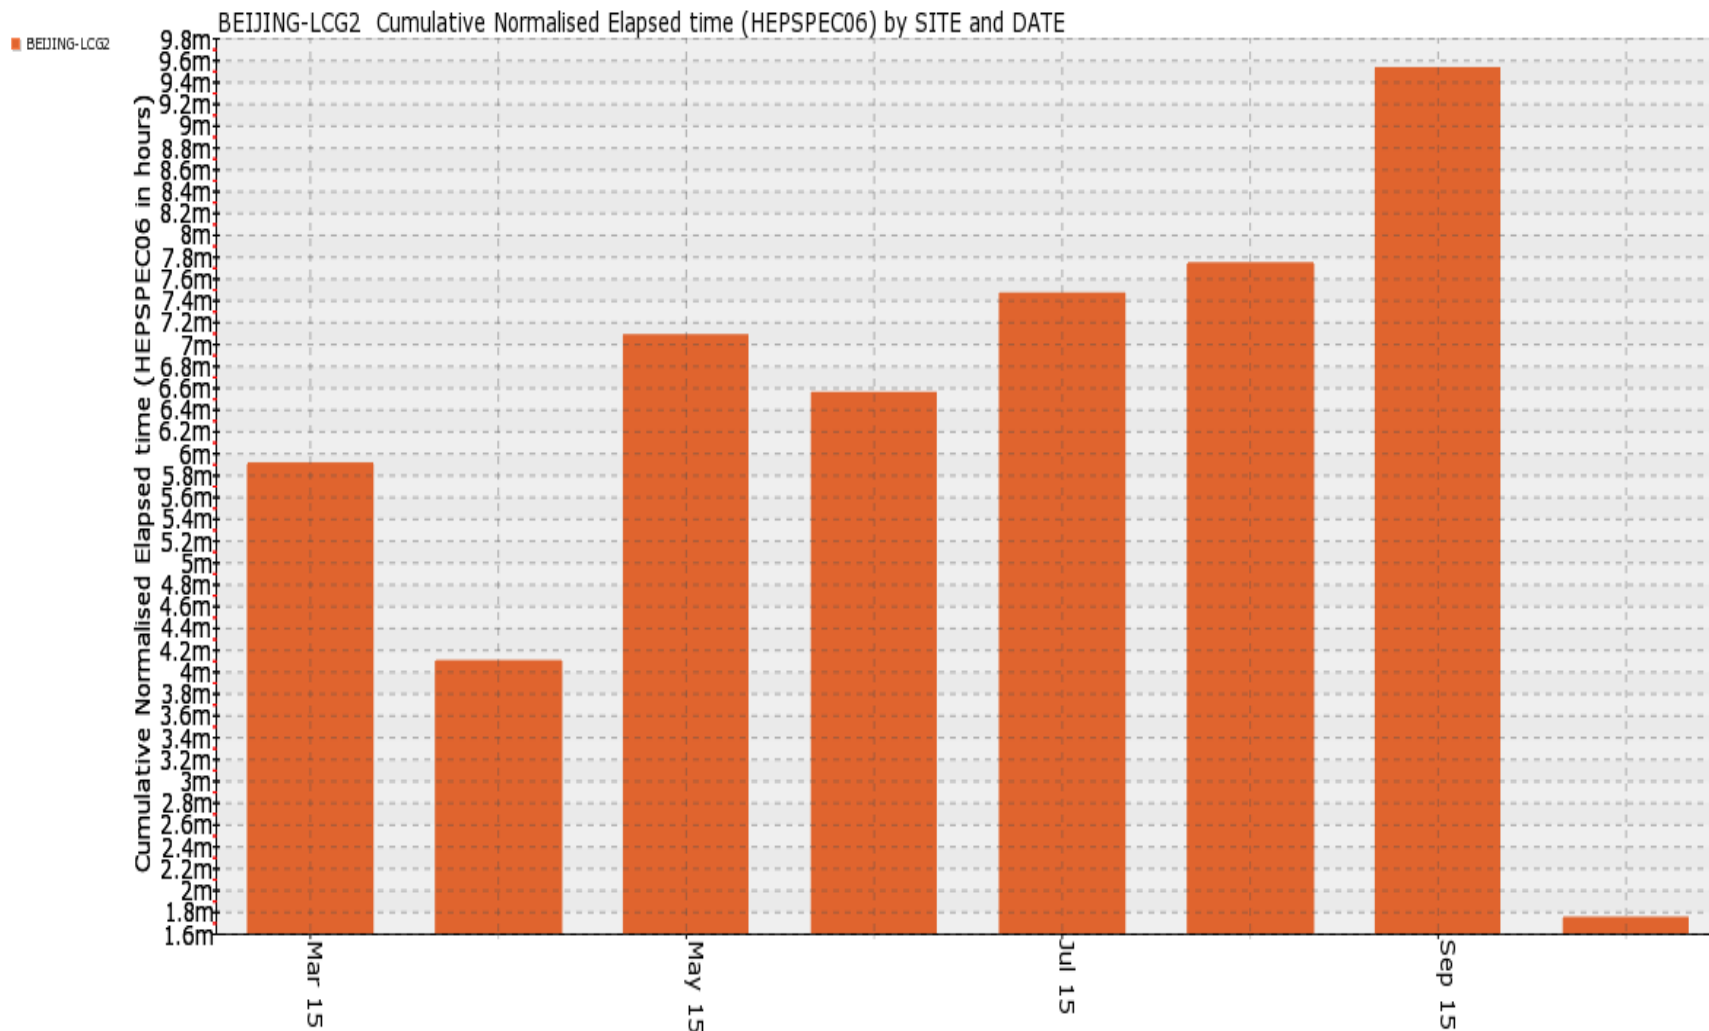
- Infrastructure Update
- Local Cluster Status
- **EGI Site Status**
- Monitoring and IT Service
- Summary



BEIJING-LCG2 Site report

Developed by CESGA EGI View: / nomelap+HEPSPEC06 / 2015:3-2015:10 / SITE-DATE / Lhc (x) / GRBAR-LIN / i

2015-10-07 21:11




BEIJING-LCG2 Site report

Developed by CESGA 'EGI View': / normelap-HEPSPEC06 / 2015:3-2015:10 / SITE-VO / lhc (x) / GRBAR-LIN / 1

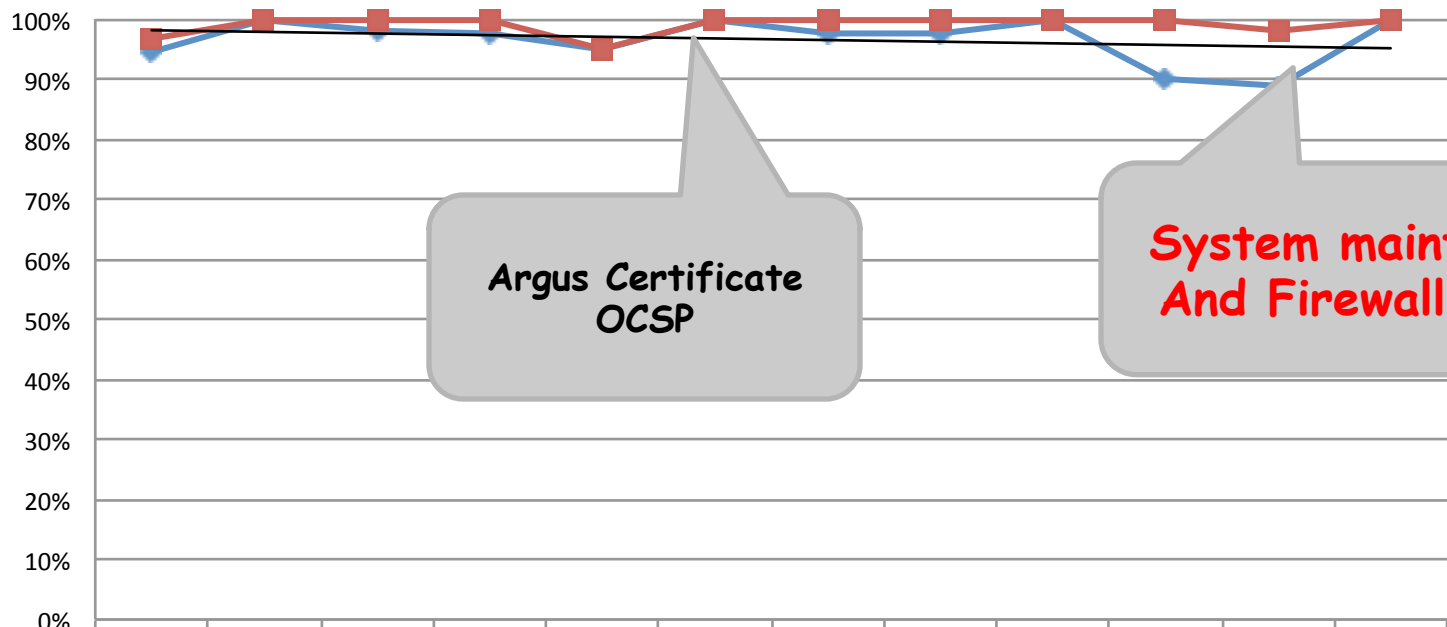
2015-10-07 21:11

BEIJING-LCG2 Normalised Elapsed time (HEPSPEC06) per VO



Reliability and Availability

2015 Availability and Reliability



	10/14	11/14	12/14	01/15	02/15	03/15	04/15	05/15	06/15	07/15	08/15	09/15
◆ Availability	94.53%	99.95%	97.94%	97.71%	95.02%	100.00%	97.81%	97.81%	100.00%	89.95%	88.87%	100%
■ Reliability	96.91%	99.95%	99.95%	99.95%	95.02%	100.00%	100.00%	100.00%	100.00%	99.66%	97.96%	100%



Site Trouble Shooting

- 2015.2.26 Creamce timeout. Argus use oosp to update certificates. Upgrade argus pkg.
- 2015.4.25 SAM test failed only happend in midnight. Sync uid and gid all the server.
- 2015.5.18 SAM test timeout. It was DNS problem. Change DNS to another server.
- 2015.6.1 SAM test failed. Network have high packet lost between China and Europe.
- 2015.8.10 Could not update Cern CA. Connected to Cern was blocked by upper network service.
- 2015.9.10 cvmfs warning. Clean cvmfs cache.



Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- **Monitoring and IT Service**
- Summary



Log Analysis

- Deployed Flume+Elasticsearch+Kibana to collect logs generated by servers and work nodes
 - syslog/Lustre log/AFS log/Gluster monitoring log
 - Logs collected from 1500+ devices in parallel

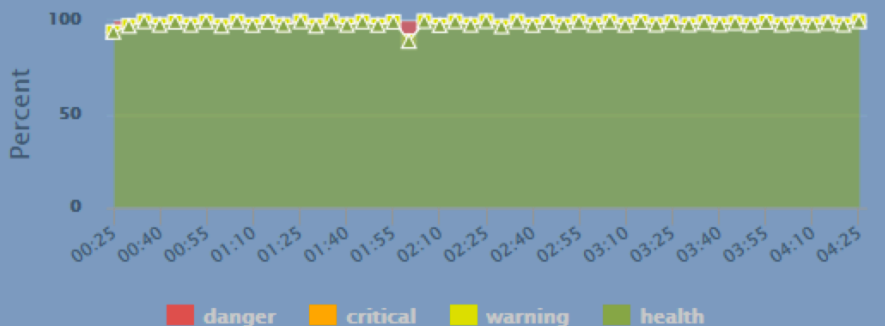
	Elasticsearch cluster	Flume-server
CPU	E5-2640 v3 @ 2.60GHz	E5420 @ 2.50GHz
RAM	64G	16G
OS	Scientific Linux 6.2	Scientific Linux 6.2
Servers	2	1



Dashboard

Summary of the Whole Cluster Healthy Status

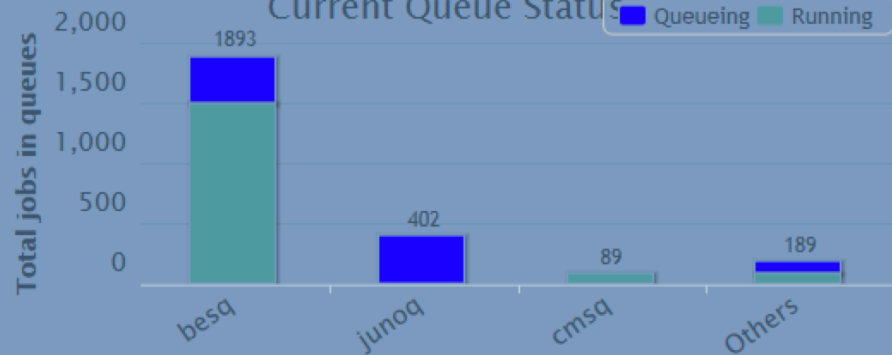
last 4 hours



Highcharts.com


Current Queue Status

Queueing (blue), Running (green)



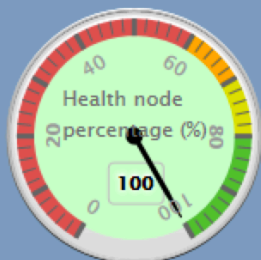
BES Healthy Status

last 10 minutes



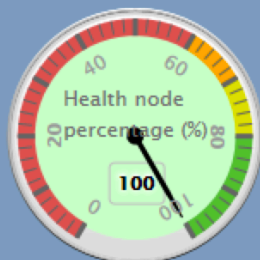
DYW Healthy Status

last 10 minutes



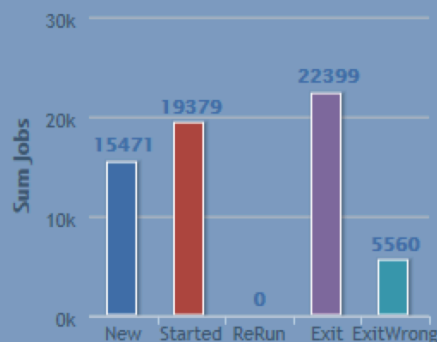
GRID Healthy Status

last 10 minutes

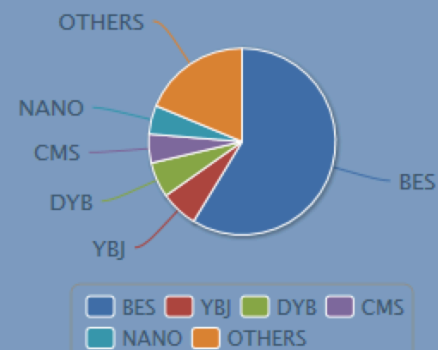


Sum Jobs:

Last Hour

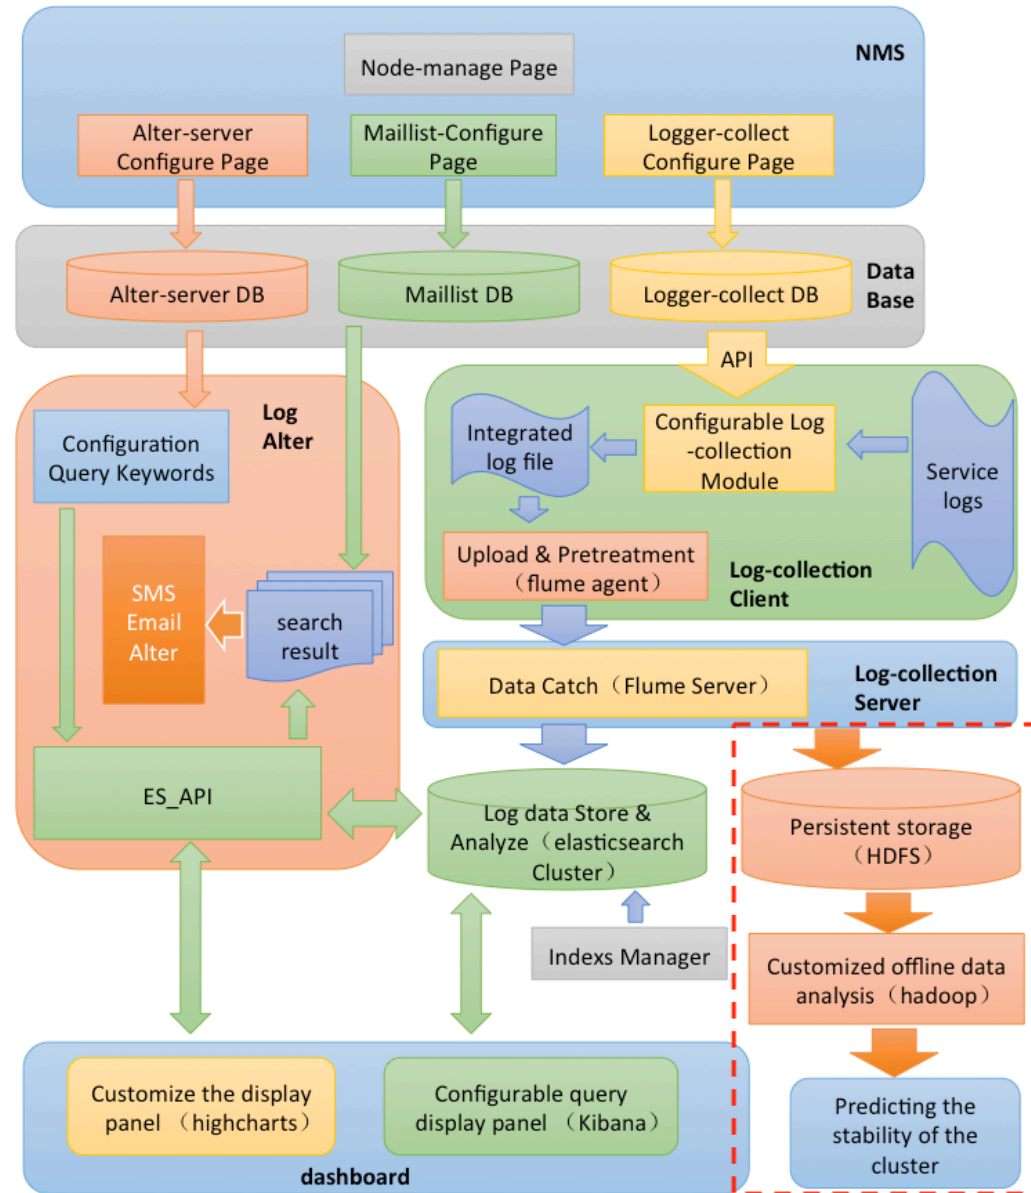


Exit jobs



Next plan

- More information will be collected.
- Trends and detailed analysis will be provided based on data mining in huge and long term logs by Map-Reduce technology (Hadoop)



IHEP IT Services

- **IHEP has officially joined the EDUROAM community since March 2015.**
- Currently, IHEP EDUROAM supports username/password auth. And The CA auth Will be supported soon
 - Please refer to: <http://eduroam.ihep.ac.cn/eduroam>
- **Service Desk (More Functions are in developing)**
 - 24/7 coverage (someone will take the call)
 - Request/Incident/Problem submit/tracking/...
 - Services Board/Services Maintenance/...
 - Authentication with the IHEP SSO (Email account and password)
 - Service Catalogue
 - Computing & Storage
 - Network & Public network services
 - ARP & Database
 - Others
- **Vidyo@IHEP: personal and group video conferencing**
 - Website: vidyoguide.ihep.ac.cn
- **IHEPBox: Upload/download/Sharing Files/Version Control**



Summary

- Our site is running well
- Constructed the new network for Data Center
- Map-reduce will be applied in Log analysis to obtain correct forecasting information
- Provide better and better IT Services for physicists
- Lustre MDS stuck is still under research



Thank you
Any Question?

