# IHEP Site Report

Jingyan Shi (Jingyan. Shi@ihep.ac.cn)
Computing Center
IHEP

# Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Management and Operation
- Summary

# Local Cluster-- Infrastructure

- More cpu/cores added
  - 1800 cpu/cores added
    - 7082 → 8878 cores
  - 200 GPU cards added
    - 148 → 348 gpu cards
- Lustre
  - 672TB retired, 1PB added
  - 3PB → 3.3PB
- Gluster
  - 129 TB added
  - 135TB → 186TB

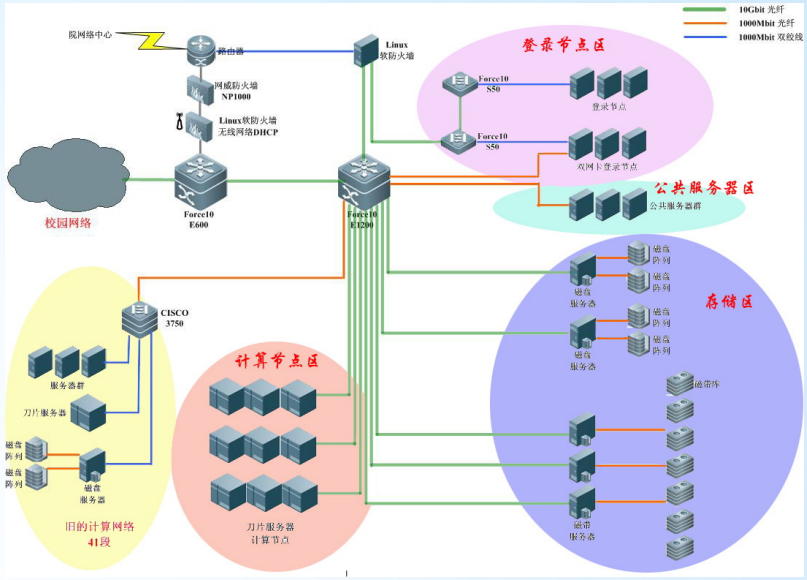

# International and Domestic Links

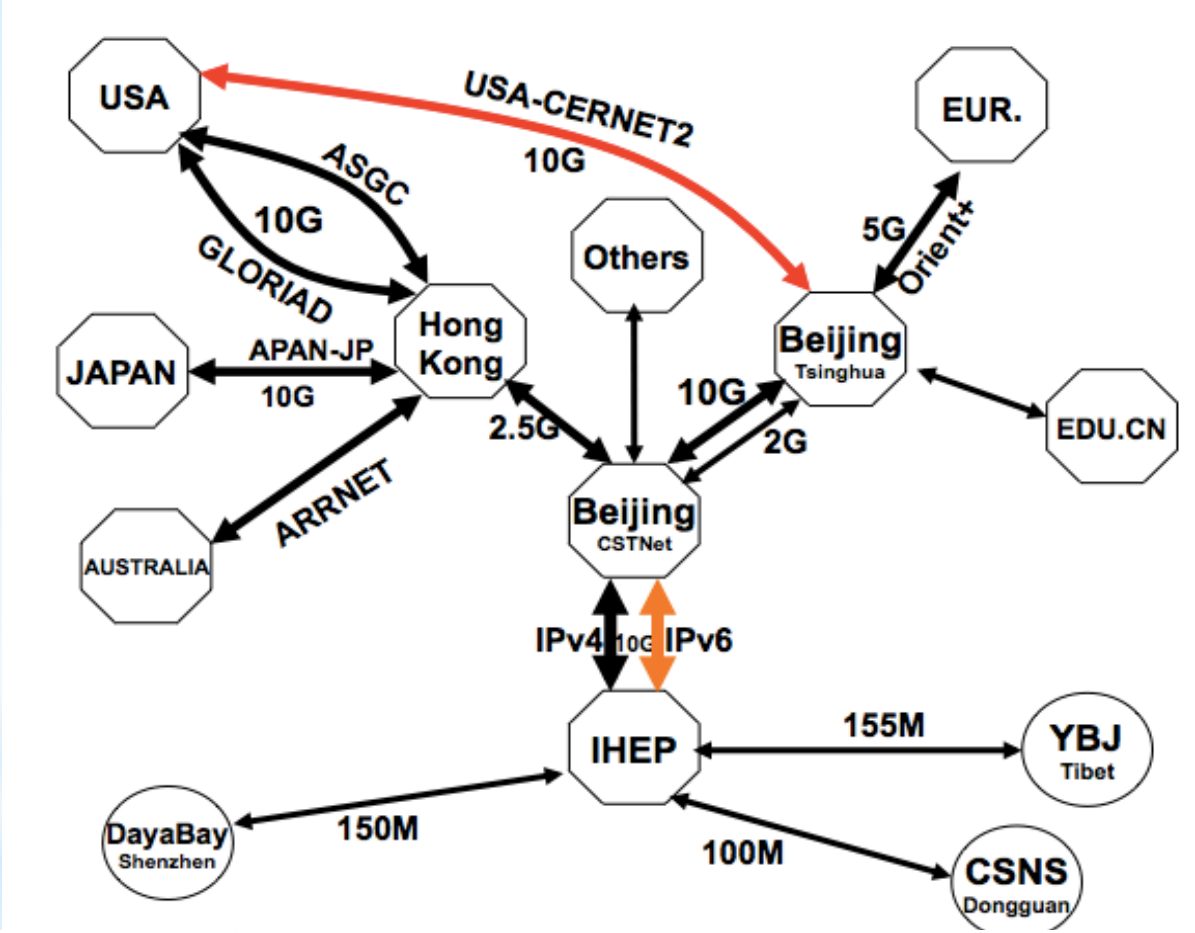


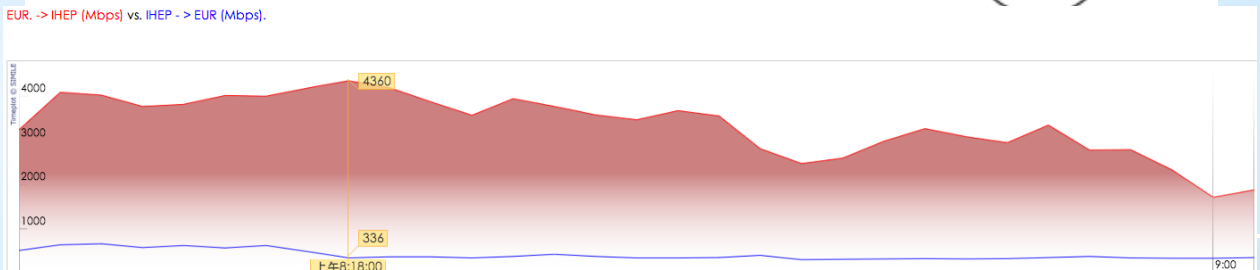

# CA

- Encryption algorithm
  - SHA1->SHA2
  - Smooth transition.
- Ocsp and ipv6
  - Expect to be finished before 2015.
  - Renew
  - Under discussion and seeking a solution

# Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Management and Operation
- Summary

# Local Cluster-- Scheduler

- Torque 2.5.5 + Maui 3.2.6
  - 8878 job slots
  - Including serial jobs, MPI jobs and GPU jobs
- Stable than new version
- More than 66% devices have been out of guaranteed period
- Problem
  - Scheduler block
  - Black-hole for jobs running
  - Zombie jobs

# Local Cluster-- Scheduler

- New tool developed and deployed
  - integrated Monitor results
  - Excluded error nodes
  - Zombie jobs detection
- Error Detection and Recovery are done automatically
- More Stable
  - Job Failure Rate dropped down
  - Torque process running stronger

# Local Cluster-- Scheduler

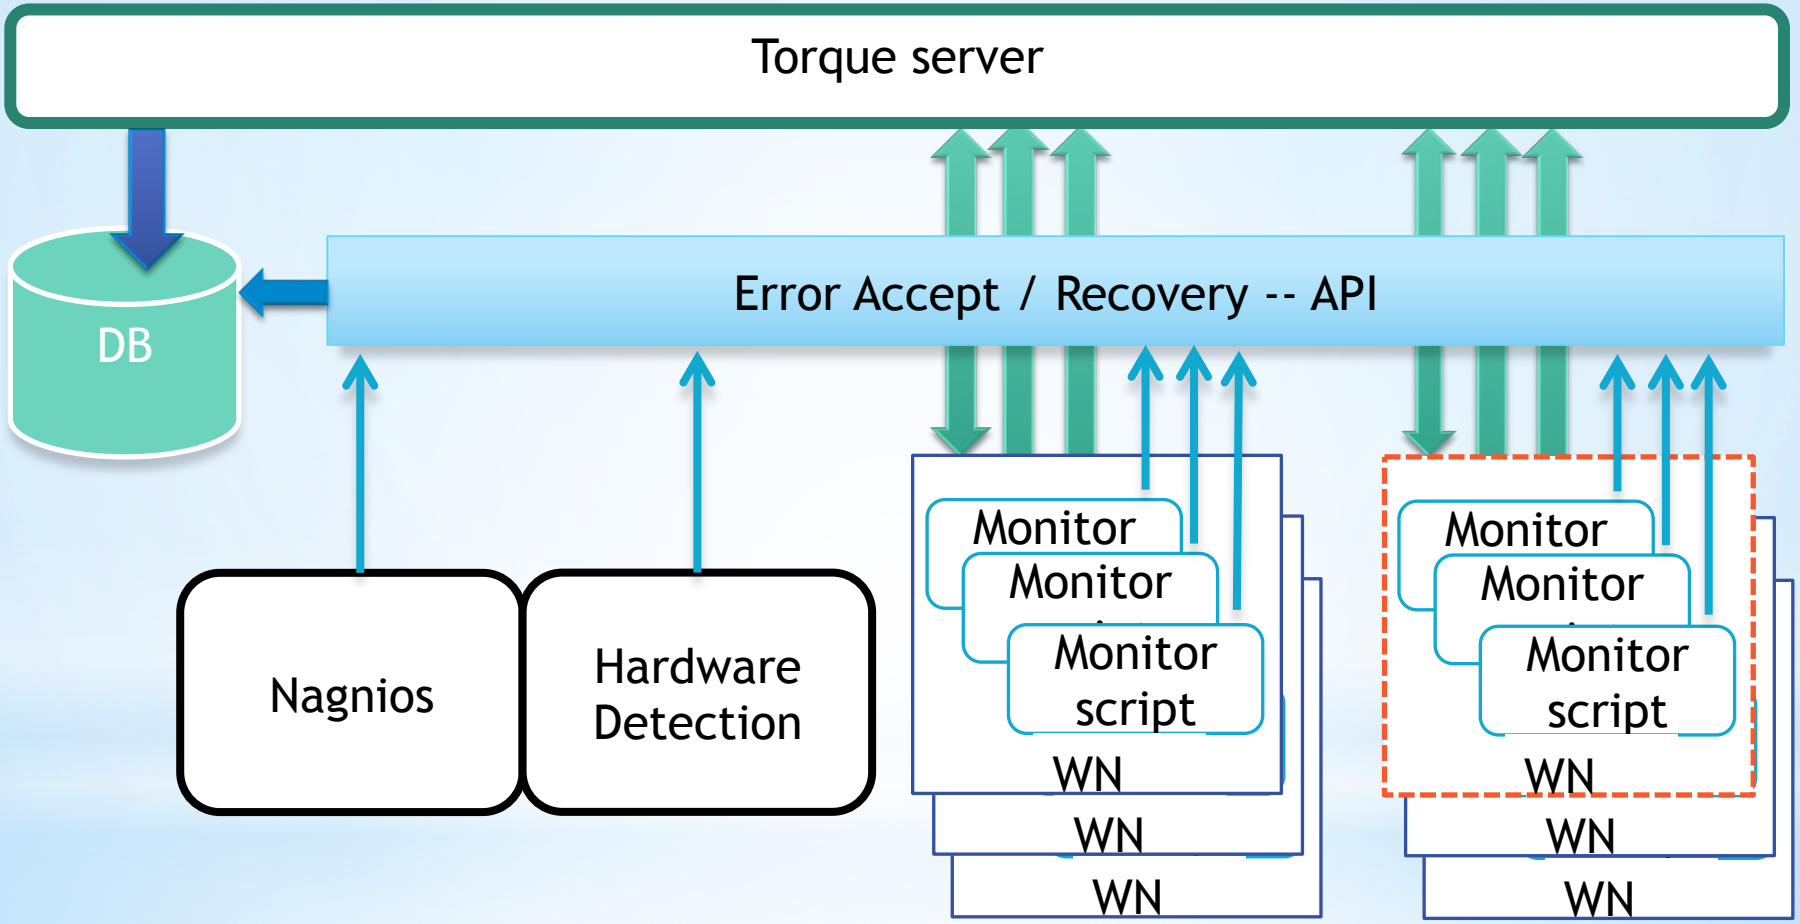

# Resource Utility Tuning

- Several queues shared same cores group
- Parameter “max_user_run” is set for fairness for users
- The amount of queue user fluctuated unexpectedly
  - free cores but jobs queuing

# Resource Utility Tuning

- Established short queue
  - Tune “max_user_run” of short queue dynamically
  - Maximized the cores for short queue
    - Depending on the resource and queue status
    - Two ways of tuning steps
      - increase or decrease gradually
      - increase or decrease sharply depending on number of free cores in queue
- setup threshold for short queue
  - Max_queue_run
  - Min_queue_run

# Resource Utility Tuning

Before

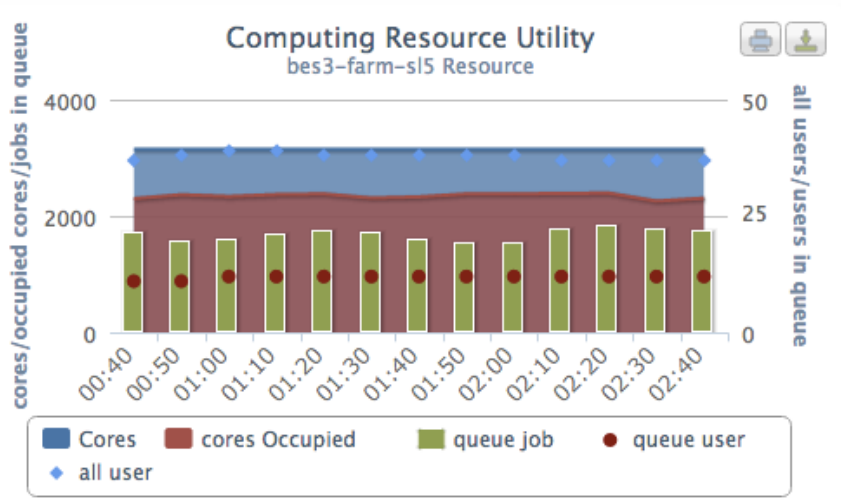


Now

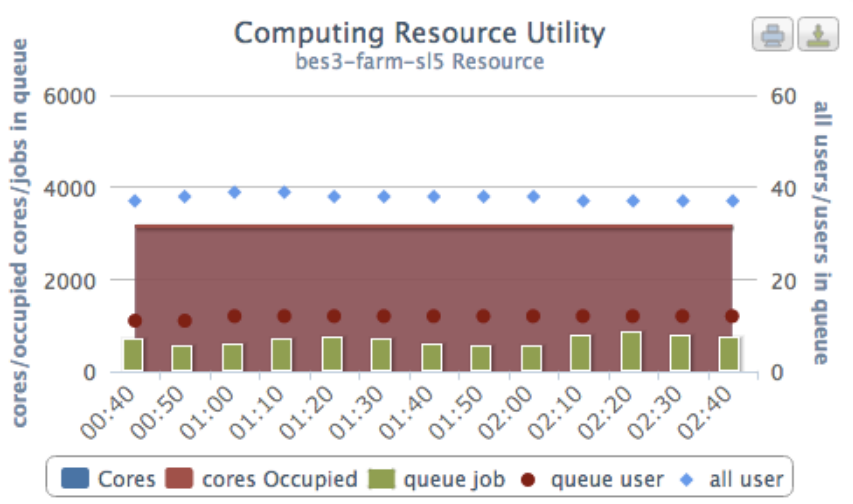

# Storage -- Lustre

- 7 file systems share a similar ( commodity server DAS SATA disk arrays ) architecture, totally 7 MDSs,34 OSSs, 3.3 PB

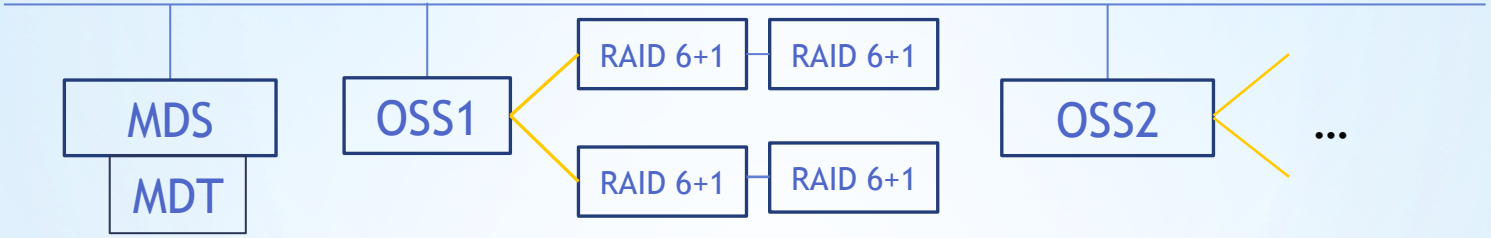


- 5 used for physics data of different experiments, 1 for individual data, 1 for temporary data storage

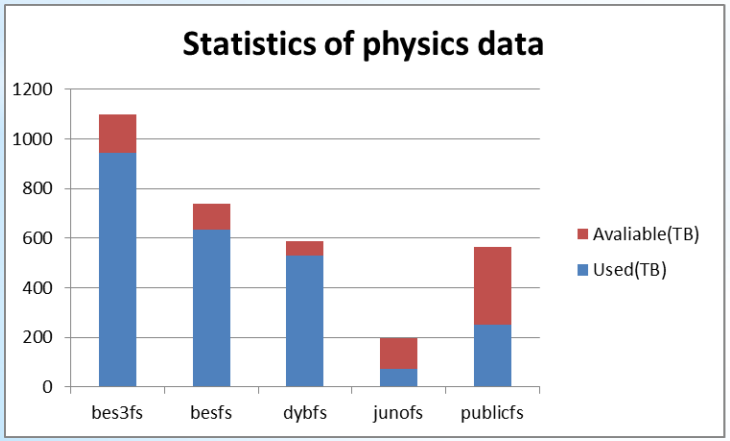


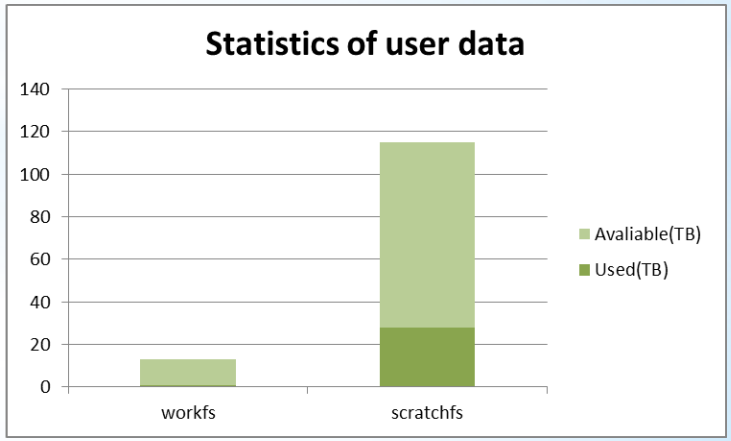

# Long rebuild time of aged disk arrays

- After intensively usage for more than 3 years, failure rate of disks increased dramatically
- With continuously heavy IO load, rebuild was very slow.

| Rebuild Time | Disk Size | RAID size | IO load |
|---|---|---|---|
| 10 days | 3 TB | 24 | High |
| 4 days | 3 TB | 24 | Low |
| 2 days | 3 TB | 12 | Low |
| 1 day | 2 TB | 12 | Low |

- During this period, possibility of another disk failure is very high
- Simultaneous failure happened 3 times during last 6 months
- For the safety of data, we have to suspend the service of related Lustre file system, and wait for the rebuild to finish.

# Solution and Plan

- In the last six month
  - Retired 40 five-year old disk arrays consists of 1 TB SATA disks
  - Resize all the 24-disk-RAID to 12-disk-RAID
  - Migrate some hot data from aged Lustre file system/disk pools to newer ones
- Setup a disk retirement and data flow policy, to guarantee data availability with low disk cost.
- Your suggestions are warmly welcomed.

# Status of Gluster File System

- Finished annual  inspection: doing fsck on each bricks
- Scale upgrade of the File system
  - I/O server: ***4->5***
  - *Bricks:**23->39***
  - Capacity:**186TB->315TB**
- Scalable and dynamic **distributed metadata service** is developed to solve
  - When  data server is busy , “ls” performance lost
  - With bricks increase, “mkdir” , “ rmdir”performance changed worse
  - Directory tree inconsistent
- Optimization and bugs fixed in data rebalance
- Monitoring: Add new nagios monitoring plugins to check **brick target status**

# Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Management & Operation
- Summary

# BEIJING-LCG2 Site report

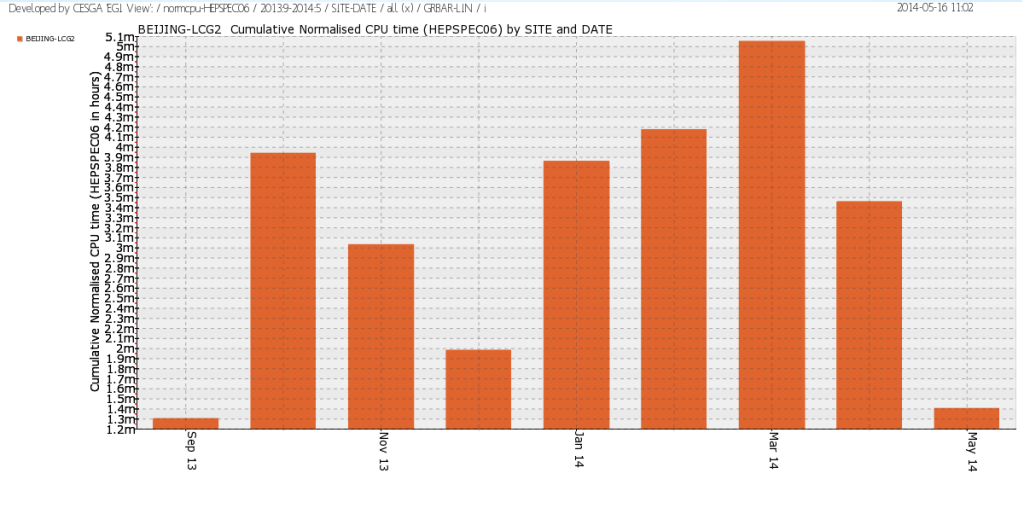


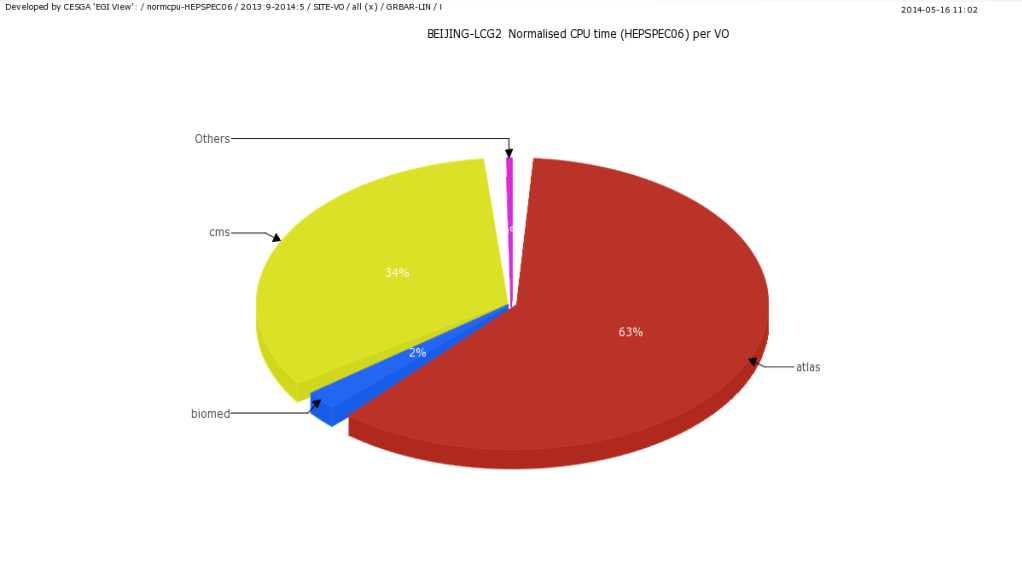

# Reliability and Availability

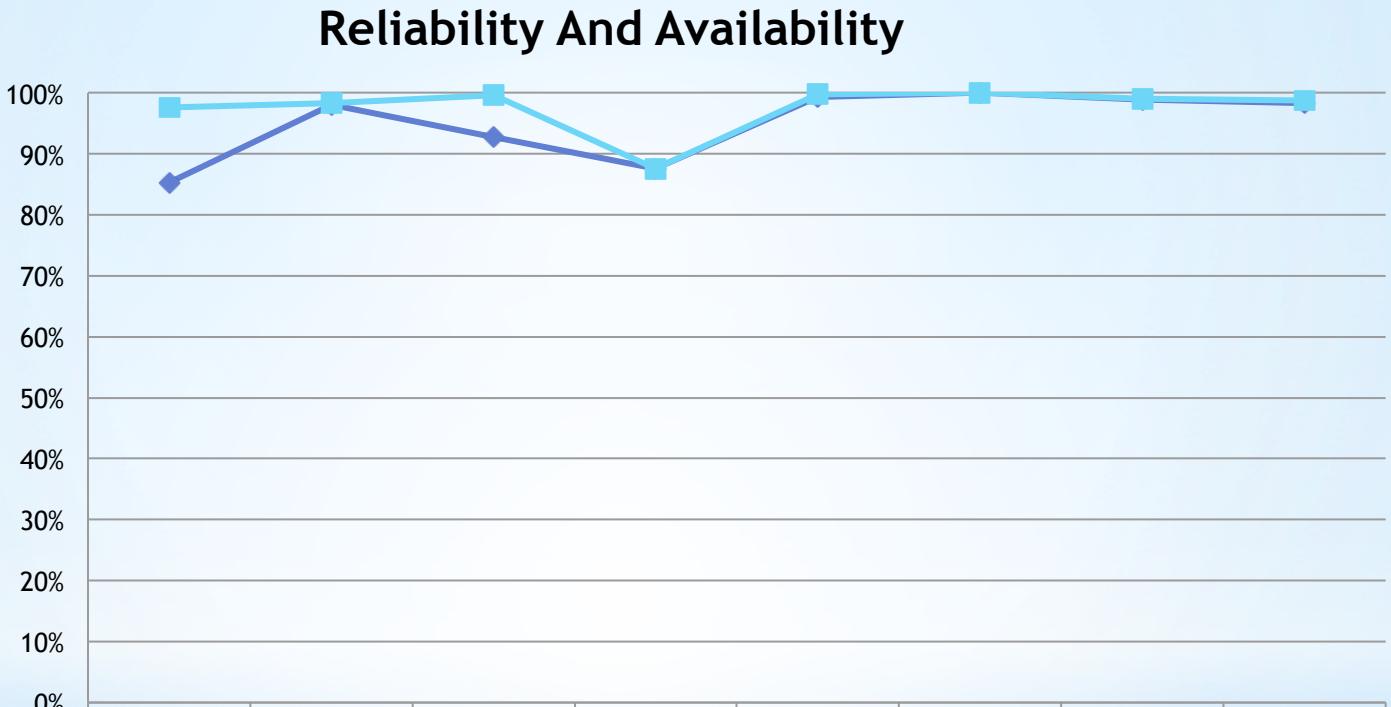


| | 09/13 | 10/13 | 11/13 | 12/13 | 01/14 | 02/14 | 03/14 | 04/14 |
|---|---|---|---|---|---|---|---|---|
| Avalability | 85.23% | 98.00% | 92.80% | 87.59% | 99.26% | 100.00% | 98.90% | 98.39% |
| Reliability | 97.59% | 98.32% | 99.59% | 87.59% | 99.79% | 100.00% | 98.99% | 98.82% |

# Site Update

- EMI3 update
  - Updated: Cream, WN, Argus, BDII, Myproxy
  - To be updated: dCache, DPM, LFC, APEL,
- All Nodes upgraded to SL6.5
- Disks for CMS replaced to 4TB * 24 Array
- Disks for ATLAS will be replaced next month
- All grid server will be replaced this year
  - Old server had been run over 5 years

# Outline

- Infrastructure Update
- Local Cluster Status
- EGI Site Status
- Management and Operation
- Summary

# Management and Operation

- Device database created
- Power accounting for each experiment
- Half of machines are managed by puppet
- ROC is under construction

# Perfsonar@ihep

- Two hosts for perfsonar
  - Perfsonar.ihep.ac.cn for Bandwidth test
  - Perfsonar2.ihep.ac.cn for Latency test
- Network performance tuning is in progress between IHEP and Eur. Sites
  - http://twiki.ihep.ac.cn/twiki/bin/view/InternationalConnectivity/IHEP-CCIN2P3

| Test Members | |
|---|---|
| hcc-ps02.unl.edu | T2_US_Nebraska |
| 193.109.172.190 | SPAIN |
| perfsonar-ps01.gridpp.rl.ac.uk | UK |
| lhc-bandwidth.twgrid.org | lhc-bandwidth.twgrid.org |
| perfsonar-ps.cern.ch | CERN |
| perfsonar.nersc.gov | NERSC |
| perfsonar-ps.cnaf.infn.it | INFN |
| perfsonar1.cc.kek.jp | KEK |
| psonar2.lal.in2p3.fr | psonar2.lal.in2p3.fr |
| heplnx128.pp.rl.ac.uk | heplnx128.pp.rl.ac.uk |
| ps.lhcopn-ps.sara.nl | SARA |
| perfsonar-ps2.ndgf.org | Nordic countries |
| 157.82.112.69 | Japan |
| perfsonar2.ihepa.ufl.edu | perfsonar2.ihepa.ufl.edu |
| sunn-pt1.es.net | 1 |
| 210.72.16.8 | 210.72.16.8 |
| perfsonar-bw.sprace.org.br | perfsonar-bw.sprace.org.br |
| perfsonar-ps-02.desy.de | perfsonar-ps-02.desy.de |
| 193.48.99.79 | IN2P3-20121006 |
| cmsperfsonar01.fnal.gov | FNAL |
| perfsonar-de-kit.gridka.de | KIT |

# Summary

- Devices running out of guarantee period caused a lot of problem
  - Disk array rebuild time
  - unstable WN
- A big challenge

THANK YOU!

QUESTION?