

Brief review of the HEPIX 2011 spring

Darmstadt, 2-6 May

<http://indico.cern.ch/conferenceDisplay.py?confId=118192>

Andrey Y Shevel

The presentation outlook

- HEPiX program
- Site reports
- Main topics from CERN
- GSI
- Batch systems
- Trends
 - Commercialization
 - HPC performance trends
 - Security

Quick HEPiX program overview

- Site reports (15 minutes all reports) [16]
- Networking & Security [4]
- IT infrastructure [10]
- Computing [details] [6]
- Storage and file systems [8]
- Cloud, Grid and virtualization [8]
- Oracle [2]

Discussed problems

- Material life
 - Electric consumption and Cooling
 - Hardware & software Upgrade
 - FTEs support
- Technical & Scientific
 - Virtualization
 - Computing Clouds
 - Network Security
 - Computing Grids
 - IPv6

Site reports ...

- GSI site report
- Fermilab Site Report - Spring 2011 HEPiX
- GridKa Site Report
- Nikhef site report
- CERN site report
- DESY Site Report
- RAL Site Report
- SLAC Site Report
- CC-IN2P3 Site-Report update
- IHEP (China) Site Report
- NDGF Site Report
- Petersburg Nuclear Physics Institute (PNPI) status report
- DLS site report
- BNL Site report
- ASGC site report
- PSI - Site report

Main info from CERN

- CERN:

- Record intensities, record luminosities
- Peaks of more than 6 GB/s to tape
- Frontier instead local databases
- Obligatory web-based security course (with test)

Virtualization

- Working group around CERN
 - How to transfer the virtual images
 - How to revoke virtual image
 - How to trust virtual image
 - How to create virtual image
 - How to keep ...
 - How to ...

CERN Electronic Documents

- <http://cds.cern.ch/> - CERN Document Server (80 contributors for 10 years)
 - Hardware: **2x** Dell Blade, Intel(R) Xeon(R) CPU E5410 @ 2.33GHz (4 cores), 16 GB RAM
- [Http://indico.cern.ch](http://indico.cern.ch) - team: 2 LD staff, 1 fellow, 2 technical students, 1 internship student
- CVS, SVN and Twiki — 2.2 FTE,
 - Hardware: Server DELL Poweredge 1950 2.33GHz, 2 CPU/8 core 48 GB of memory

Programm	2010			2013			2016		
	CPU # cores	On-line PB	Tape PB/a	CPU # cores	On-line PB	Tape PB/a	CPU # cores	On-line PB	Tape PB/a
APPA	100	< 0,1	0,1	200	1,0	2,0	1.000	1,0	2,0
CBM	250	< 0,2	0,2	7.000	15,0	11,0	60.000	15,0	11,0
NUSTAR	100	< 0,1	0,0	300	2,0	2,0	2.000	2,0	2,0
PANDA	250	< 0,2	0,2	8.000	12,0	12,0	66.000	12,0	12,0
ALICE	3.000	1,5	0,0	6.600	3,5	0,0	14.500	8,0	0,0
FOPI	200	0,1	0,2	400	0,2	0,3			
HADES	1.000	1,0	1,5	1.000	2,0	3,0			
Gitter-QCD	2.500	< 0,3	0,3	20.000	0,5	0,5	100.000	1,0	2,0
Theorie	1.000	< 0,2	0,2	5.000	0,5	0,5	20.000	1,0	2,0
rest	500	< 0,3	0,3	2.000	1,5	1,5	10.000	3,0	3,0



Computing

- Selecting a new batch system at CC-IN2P3
 - The candidates were: Condor, PBS-Pro, Torque/Maui, LSF, BQS, SGE, Loadleveler (IBM), SLURM (LLNL)
 - Requirements
 - Scalability, robustness, sharing quality, scheduler, AFS, worker nodes features, interface to the Grid, administration, information repository, accounting, parallel and interactive jobs, software support, procurement cost, maintenance cost
 - Final choice was «SGE» (the second by criteria was «LSF»)


Oracle and commercialization related issues

- Oracle corporation has bought Sun corporation
 - Oracle stopped support for several software packages: Lustre and several others.
 - In result people on HEPiX suggested to ask EU for money to keep Lustre free of charge for anybody (same procedure as Scientific Linux).
- Other commercializations:
 - Similar commercialization we see with «IBM systems journal» and «IBM research journal», «Intel Technology Journal» (earlier they were available for free over Internet)

HPC – Price/Performance

The price/performance of clusters installed at Fermilab for lattice QCD has fallen steadily over the last six years, with a halving time of around 1.5 years, as shown by the solid line in the graph at right.

Product roadmaps provided by vendors of system components make clear that this trend is likely to continue for the next several years.



Security update

- SONY play station network: 77 million account were compromised (personal details, credit cards data).
- Other play sites — almost the same.
- In academic area: SSH remain the primary intrusion vector
 - Passwords+Keys: sniffed/copied and re-used by attackers
 - The vast majority of Linux incidents at CERN results from compromised account **at other sites**

Rootkit checkers

Signature-based:

- rkhunter, chkrootkit, etc.
- Very efficient **against known versions** of rootkits
- Very easy to use
- Not so efficient if the rootkit is open source or maintained
 - e.g. there are at least 10+ Phalanx updates
- snapshot and check:
 - Samhain, Tripwire, Zeppoo, the99lb, etc.
 - **Good results**, but often require **significant work**
 - **Little/no public rootkit checkers would detect DR rookits**
- OSSEC's rootcheck could detect a sample discovered last year
(hidden process + wrong link count on the filesystem)

Good ways to manage security risks at a reasonable cost?

- Strict **access control** of users (multifactor is a significant help)
- Tight **security patching** policy
- **Periodic reinstallation** of the nodes
- Prevent users to **escalate as root** (system hardening)
- Implement **in-depth monitoring** of the system
- Detailed **log trails** (remote syslog)
- Good **incident response capability**
- Ability to manage **incidents** as part of normal operations

Data loss might happen at any time

- Amazon's Cloud Crash Disaster Permanently Destroyed Many Customers' Data
 - «Many days after the crash, Amazon still hasn't gotten its systems fully up and running again.»
 - Read more: <http://www.businessinsider.com/amazon-lost-data-2011-4>
- Need to have spare copies in independent sites.
- The same might happen with your own communicators, laptops, desktops, etc. at any time.

Thank you ! Questions?