

NAF Status

Andreas Haupt, [Yves Kemp](#)
NUC, 9.9.2009

Status

- > Running was smooth
- > Planned downtime due to switch maintenance
 - NAF was up and running quite fast
 - dCache upgrade made some troubles, fixed now
- > Kernel vulnerability found, root-escalation kit exists
 - ... and worked on SL4 systems (vulnerability also on SL5 machines)
 - No z-series Kernel for SL4.7 anymore (remember last NUC)
 - With new Kernel, we had to switch Lustre@SL4 to TCP/Ethernet
 - InfiniBand for SL5 works OK
- > Upgrade to SL4.8 (from SL4.7)
 - Atlas detected problems with their software, fixed.
- > Some HW problems, using maintenance window Thursday:
 - VM host for <experiment>.naf.desy.de
 - tcx040 Atlas WGS
- > Action items solved since last NUC:
 - Added Autoreply for naf@desy.de tickets (to requestor and all CCs) (caveats see last august report)
 - “Waiting and running jobs”-plots added



Status SL5 WNs

- > Using the downtime, we switched ~50% (51 hosts SL5 / 42 hosts SL4) of all the WNs to SL5
 - ... and probably most important also the queue default
- > Some problems seen on the mailing lists, seem solved
 - Reports from the experiments?
- > Queue: Vast majority of waiting jobs asking for SL5 (8.9.9@14:00)
 - Proposal: Have 90% WNs on SL4 fast (before 1.10.09)?
- > DESY-HH Grid: Migration to SL5 64bit finished, no SL4 resources available for NAF experiments.
 - CMS, Atlas and ILC have used them (in that order)



Statistics

CPU days	SL4 only	SL5 only	SL4&SL5
atlas	1300	1808	99,6
cms	1	1462	6,2
ilc	164	71,7	281,7
lhcb	1932	100,7	0
all	3395	3442,7	387,6

- > Experiments: Try to minimize number “SL4 only” column
- > Majority of users use SL5 (only or SL4|SL5)
 - Are there principal reasons for continued use of SL4 only (by a few people)?
- > We are at the moment disfavoring CMS, and favoring LHCb
 - If possible, put in -l os='sl4|sl5'
 - Or raise SL5 share soon: 90% SL5 still this month



Status and plans WGS

WGS list 8.9.2009

tcx030	SL4.8_64	atlas
tcx040	SL4.8_64	atlas
tcx050	SL4.8_64	atlas
tcx060	SL5.3_64	atlas
tcx031	SL4.8_64	cms
tcx041	SL4.8_64	cms
tcx051	SL5.3_64	cms
tcx032	SL4.8_64	ilc
tcx042	SL4.8_64	ilc
tcx052	SL5.3_64	ilc
tcx033	SL4.8_64	lhcb
tcx043	SL4.8_64	lhcb
tcx053	SL5.3_64	lhcb

- > Interactive WGS is still SL4 default
- > CERN is migrating on 5.10.2009
 - What should we do?
- > My proposal: Move on 15.10.2009
 - Move all experiments together
- > Migrate (at least) one WGS per VO from SL4->SL5 on that day
 - Need at least two WGS per VO for load-balancing
- > Keep SL4 WGS for “hopping login”
 - Number t.b.d.
- > No load-balancing for SL4 WGS



Information: Hardware purchases (not only NAF, but related)

> General ideas:

- Emphasis on storage for NAF
- Some investment in backbone hardware
- Rumor: “Haushaltssperre” after the “Bundestagswahl” ...

> Grid: 12 Blade Enclosures (each 16*8 cores, 3GB RAM/core)

- 10 Hamburg, 2 Zeuthen, partially replacement for older HW

> Storage:

- DDN (~450 TB net) + servers for Hamburg, dCache and/or Lustre or something else
- ~200 TB for Atlas dCache + ~50 TB LHCb dCache in Zeuthen

> Network and other infrastructure:

- New switch for Zeuthen
- 10 GBit infrastructure in HH (HBA for poolnodes, switches)
- Minor items (Servers, small filers,...)

> Nothing there yet, but most is on its way to be ordered/already ordered

