

NAF status

Yves Kemp
NUC July 2010

Current issues

> “kswapd” issues (reported already last time(s))

- Instabilities on some nodes. Affected IN: Mainly (only?) ATLAS. Why?
- Similar (identical?) problems also reported from other sites. One fix (swappiness 60->10) did not help us.
- Hope problem will get solved by RedHat in some kernel upgrade
- We reboot machines faster in case of problems
- Bug found in load balancing script: Ignored one special error state.

> InfiniBand errors

- Only occur in Hamburg
- Old firmware on HP-Blade InfiniBand switches: Currently being upgraded
- Large Voltaire switch: One Fabric Board gives errors. Vendor tells us to swap two FB for debugging, Zeuthen had seen similar errors in the past, such an operation solved their problem
- Problem: Taking FB out means: No network for Lustre, means clients will hang up, server could be troubled.... Means complete downtime of the NAF.
- In contact with vendor whether other options exist



Current reorganisation of IN&WN

> Different reasons:

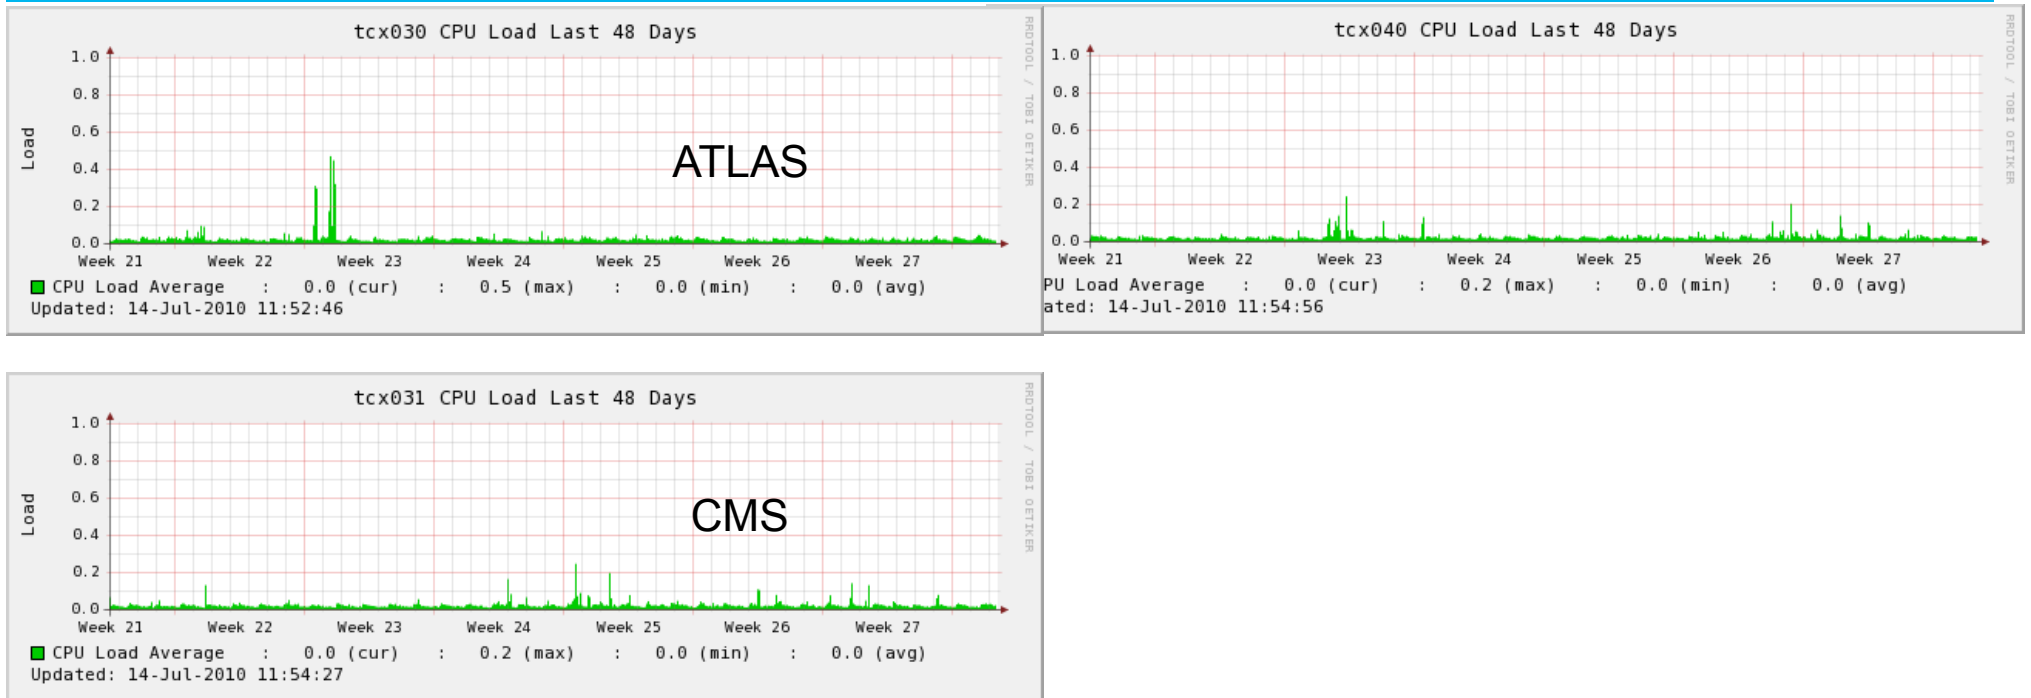
- Firmware of HP-IB switches
- New kernel and starting migration to SL5.5
- Better distribution of IN over different racks (rack failures)
- SL4 -> SL5 migration for some IN

> Prospected situation after the intervention

- ATLAS: tcx040 A, tcx060 B, tcx080 C, tcx120 D
- CMS: tcx061 A, tcx081 C, tcx121 D
- ILC: tcx032 A, tcx062
- BLHCB: tcx033 A, tcx063 B
- COMMON: tcx030 (SL4) A, tcx053 (SL5.Nbeta testing) B
- (A,B,C,D stand for different racks; tcx03-tcx06 are “old” HP-Blades, other are new DELL blades)



SL4



- > User login count sent to atlas & cms support lists.
- > tcx030 will be the only SL4 machine in a couple of days
- > Future of SL4? Needs?
- > BTW: SL6beta Q1 2011



Lustre additions

- > ATLAS, CMS and ILC have gotten each ~16 TB more in Hamburg Lustre space

▪ Tot	Used	Free	
▪ 47T	12T	33T	/scratch/zn/lustre/lhcb
▪ 20T	19T	252G	/scratch/zn/lustre/atlas
▪ 41T	37T	2.0T	/scratch/hh/lustre/atlas
▪ 37T	17T	18T	/scratch/hh/lustre/cms
▪ 33T	20T	12T	/scratch/hh/lustre/ilc

- > Lustre extension: Currently finding a way of integrating dedicated DESY-ATLAS resources into HH Lustre ... ongoing discussion



DRAFT: Template to filled in case of storage access problems in the NAF

> To be sent to naf-helpdesk@desy.de

- 1) Date and time of the problem
- 2) Userid (local access) or Certificate DN (Grid-SE access)
- 3) What files are you accessing. Exact location and full path. Name example files if many are affected, directory is not sufficient.
- 4) Which storage system is used AFS, Lustre, D-Cache-SE?
 - 1) (Leave blank, if you do not know, but be sure to have answered Q3.)
- 5) What access protocol is used, GridFTP/SRM, Dcap, GSIDcap, Posix?
 - 1) (Name the Experiment application and the config, if you do not know the protocol)
- 6) FROM where do you want to access the data (ideally hostname)
- 7) Quote the error message here (Cut & Paste)
- 8) Is the problem persistent?
 - 1) Yes always the same file(s).
 - 2) Yes, but problematic files are varying.



Downtimes

> Short downtime tomorrow

- Reboot of Login VMs
- Ins
- WNs will be rebooted silently

> Longer downtime needed

- HH: Voltaire switch & IB
- Zeuthen: Network and infrastructure work
- Probably 14&15.9. date for latter work, could combine this with HH work



Open Action Items:

id	who	description	status/comments
old-1	NAF/ATLAS	CMT problem	experts contacted
1005-1	NAF	present SL4 work group server usage done	show usage at next NUC
1005-2	ATLAS/CMS	plans for SL4 work group servers done	reduce to one shared SL4 wgs, talk to users
1005-3	NAF	email notification for /scratch monitoring (instance full)	Still open
1005-4	NAF	technical constraints for Lustre extension	ongoing discussion ongoing
1005-5	NAF	automatic Lustre space clean-up	expert meeting foreseen ongoing
1006-1	NUC	CPU, dCache, Lustre resource requirements	needed for next PRC ongoing
1006-2	NAF	SGE queue for downloads/IO?	ongoing
1006-4	NAF	form for storage report on NAF web	Draft proposed
1006-3	NAF/ATLAS	reboot procedure for (ATLAS) WGS	done
1006-5	NUC	ideas for accounting Still open	by e-mail or next NUC or extra meeting
1006-6	NAF	dCache upgrade time line (10 GE in HH)	Ongoing: ~30% done for ATLAS&CMS

