



GSI AF

"CAF" experience at GSI

Kilian Schwarz

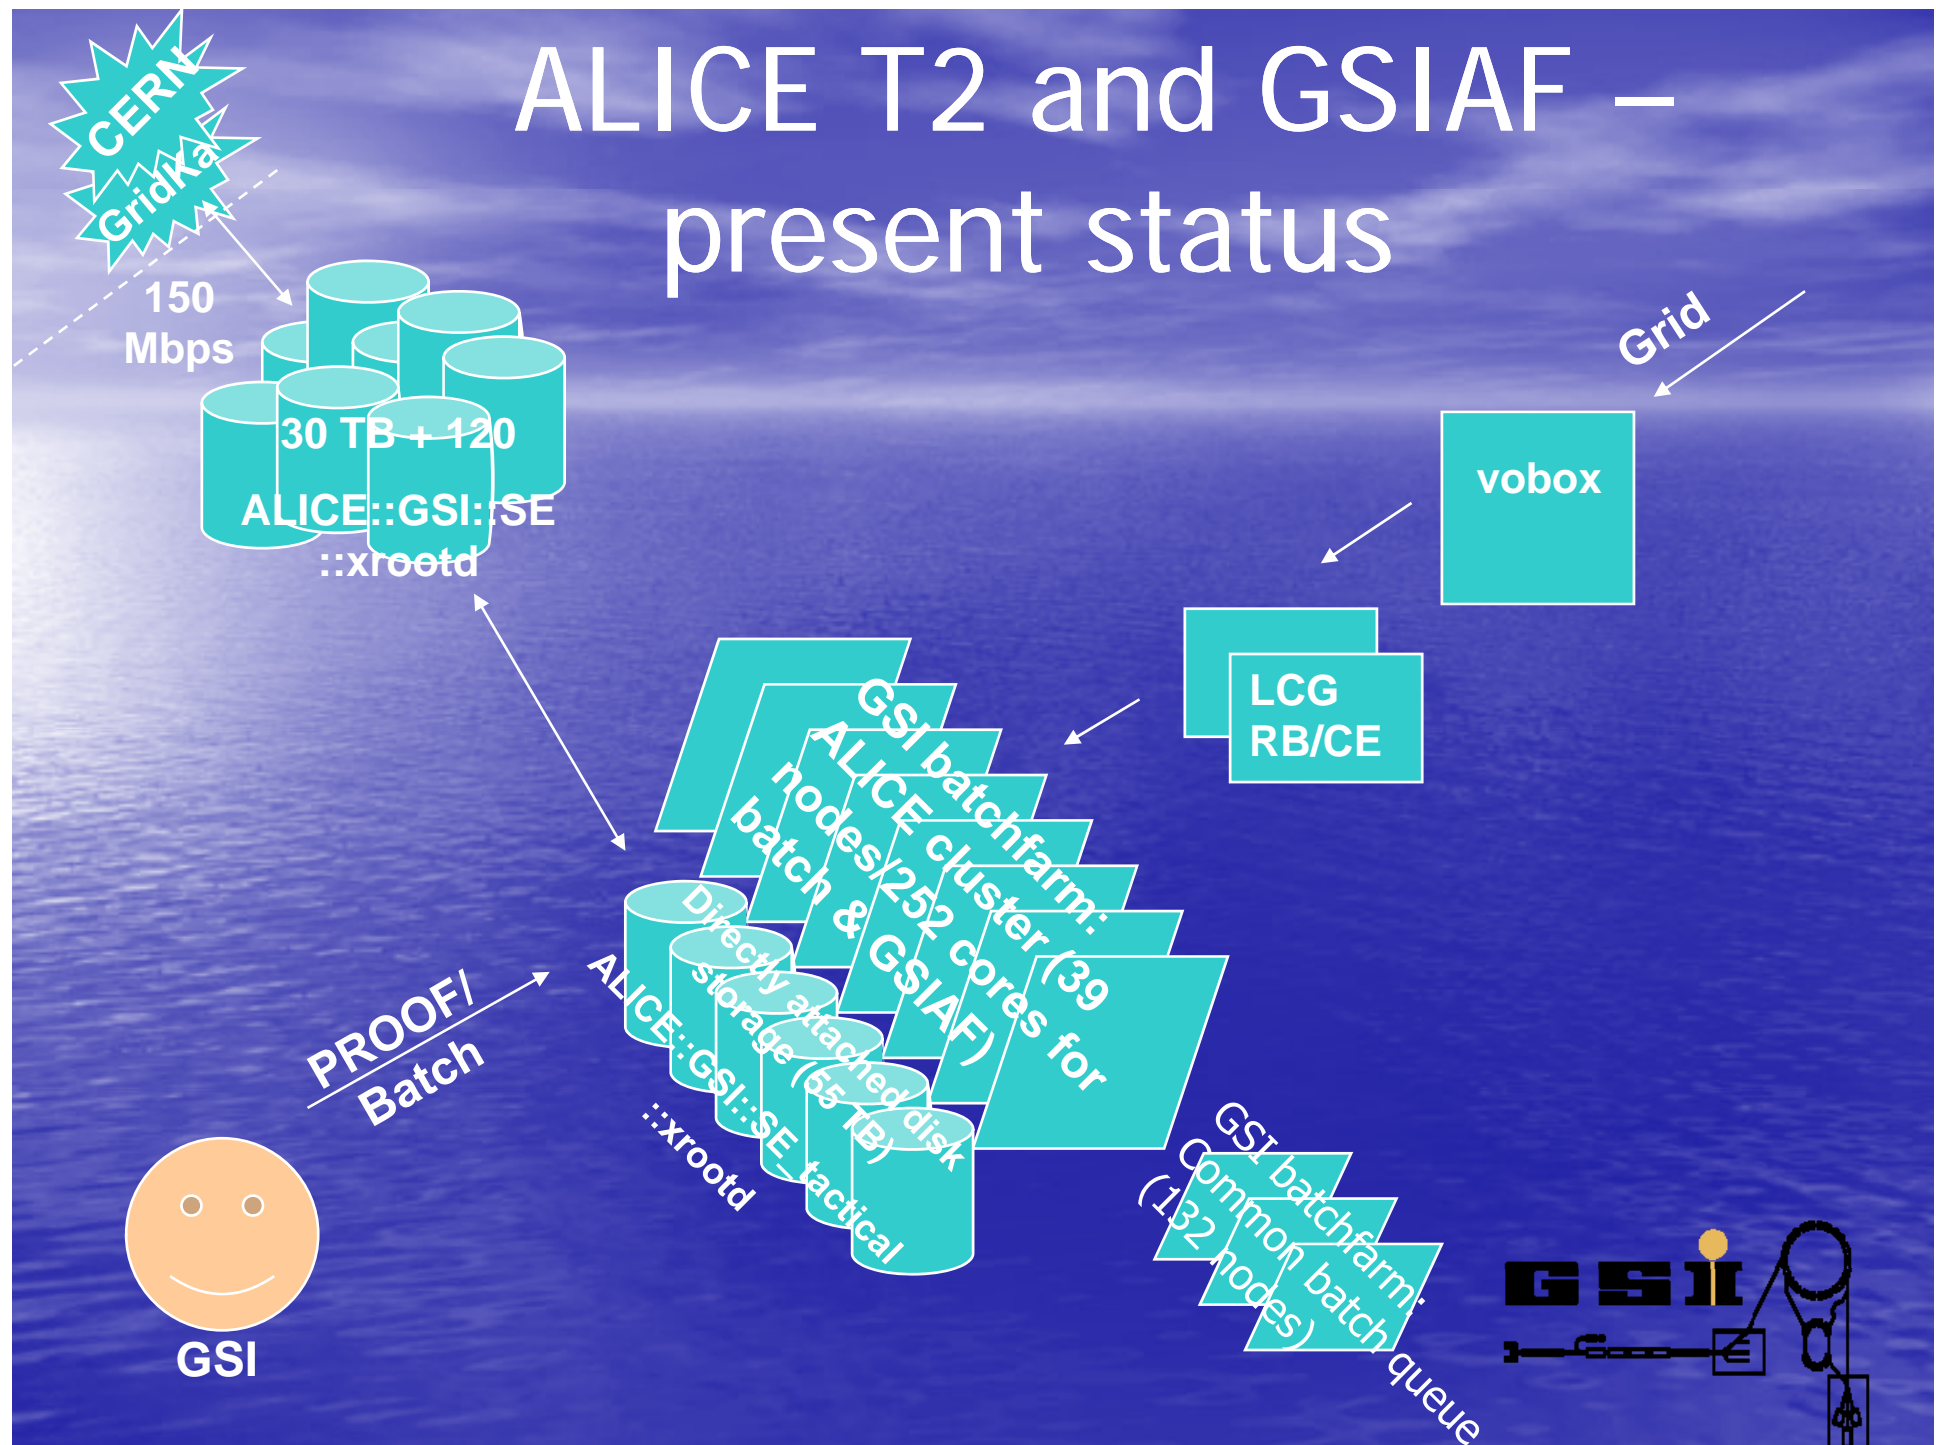
GSIAF



- Present status
- installation and configuration
- usage experience (see talk of M. Ivanov on Friday)
- Plans and outview
- Issues and problems



ALICE T2 and GSI AF – present status



Present Status



- ALICE::GSI:SE::xrootd
- > 30 TB disk on fileserver (8 FS a 4 TB each)
- + 120 TB disk on fileserver
 - 20 fileserver 3U 15*500 GB disks RAID 5
 - 6 TB user space per server
- Batch Farm/GSIAF and ALICE::GSI::SE_tactical::xrootd nodes dedicated to ALICE:
- 15 D-Grid funded boxes: each
 - 2*2core 2.67 GHz Xeon, 8 GB RAM
 - 2.1 TB local disk space on 3 disks + system diskAdditionally 24 new boxes: each
 - 2*4core 2.67 GHz Xeon, 16 GB RAM
 - 2.0 TB local disk space on 4 disks including system
- on all machines: Debian Etch 64bit

installation

- shared NFS dir, visible by all nodes
 - xrootd (version 2.7.0 build 20070926-0000)
 - ROOT (5.17/04)
 - AliRoot (head)
 - all compiled for 64bit
- reason: due to fast software changes
- disadvantage: possible NFS stales
- tried to build Debian packages of the used software, but this process is too long lasting

configuration

- setup: 1 standalone, high end 8 GB machine for xrd redirector and proof master, Cluster: xrd data servers and proof workers, AliEn SE
- so far no authentication/authorization
- via Cfengine
 - platform independent computer administration system (main functionality: automatic configuration).
- xrootd.cf, proof.conf, TkAuthz.Authorization, access control, Debian specific init scripts for start/stop of daemons (for the latter also Capistrano for fast prototyping)
- all configuration files are under version control (subversion)

Cfengine – config files in subversion



```
# -*- cfengine -*-
#
# xrootd configuration
#
# $Id$
#

copy:

# Install xroot/proof configuration files
xrootd::
    $(configroot)/xrootd/xrootd.cf
    dest=/etc/xrootd/xrootd.cf
    server=cfmaster
    type=sum
    mode=0644

xrootd::
    $(configroot)/xrootd/proof.conf
    dest=/etc/xrootd/proof.conf
    server=cfmaster
    type=sum
    mode=0644

xrootd::
    $(configroot)/xrootd/TkAuthz.Authorization
    dest=/etc/xrootd/TkAuthz.Authorization
    server=cfmaster
```

monitoring via MonaLisa

http://grid5.gsi.de:8080

GSIAF Cluster Monitoring with MonALISA - GSIAF Cluster - Microsoft Internet Explorer

Datei Bearbeiten Ansicht Favoriten Extras ?

Zurück Suchen Favoriten

Adresse <http://grid5.gsi.de:8080/stats?page=CLUSTER%2Ftable> Wechseln zu Links

GSIAF Cluster

Machines status

Machine	Machine status		CPU					Memory					
	Online	Ping	load	usr	sys	nice	idle	Total	Used	Cached	Buffers	Free	Total
grid1.gsi.de		0.023ms	0.64	4.532	0.808	0	94.66	4 GB	1.633 GB	1.607 GB	33.64 MB	2.367 GB	
grid13.gsi.de		0.309ms	1.23	23.1	26	0	50.91	1.194 GB	177.8 MB	708.6 MB	9.473 MB	1.021 GB	
grid19.gsi.de		0.895ms	0.05	0.067	0.033	0	99.9	1 GB	328.5 MB	629 MB	38.54 MB	695.7 MB	
grid29.gsi.de													
grid5.gsi.de		0.448ms	0.04	3.675	1.292	0	95.03	1 GB	301.1 MB	693.6 MB	21.99 MB	723.1 MB	
grid8.gsi.de		0.229ms	0.4	33.52	4.725	0	61.75	756.2 MB	530.3 MB	147.2 MB	2.551 MB	225.9 MB	
gsiaf.gsi.de													
lcb06.gsi.de		5.68ms	0.68	10.88	4.051	0	85.07	3.911 GB	607.4 MB	2.402 GB	105.9 MB	3.318 GB	1.946 GB
lxb255.gsi.de		5.86ms	0	0.015	0.023	0	99.96	7.801 GB	927 MB	6.124 GB	31.78 MB	6.896 GB	1.953 GB
lxb256.gsi.de		0.244ms	1.04	10.98	2.546	0.019	86.45	7.801 GB	979.8 MB	882.8 MB	27.23 MB	6.844 GB	15.27 GB
lxb257.gsi.de		0.357ms	1.16	13.17	0.564	0.015	86.25	7.801 GB	1.169 GB	2.416 GB	29.8 MB	6.632 GB	15.27 GB
lxb258.gsi.de		0.205ms	6.03	50.56	4.487	0	44.96	7.801 GB	3.204 GB	3.762 GB	13.07 MB	4.597 GB	14.91 GB
lxb259.gsi.de		0.195ms	3.03	37.71	0.986	0	61.3	7.801 GB	1.84 GB	1.309 GB	28.7 MB	5.961 GB	14.91 GB
lxb260.gsi.de		6.04ms	4.42	50.5	2.676	0	46.82	7.801 GB	2.61 GB	2.127 GB	13.73 MB	5.191 GB	14.91 GB
lxb261.gsi.de		0.225ms	4.26	50.48	2.318	0	47.21	7.801 GB	2.709 GB	4.308 GB	28.6 MB	5.093 GB	14.91 GB
lxb262.gsi.de		0.232ms	3.04	37.91	2.47	0	59.62	7.801 GB	2.373 GB	3.17 GB	27.93 MB	5.428 GB	14.91 GB
lxb263.gsi.de		0.171ms	3.15	37.77	1.257	0	60.97	7.801 GB	1.959 GB	1.794 GB	28.72 MB	5.843 GB	14.91 GB
lxb264.gsi.de		0.317ms	3.36	37.57	1.876	0	60.56	7.801 GB	2.191 GB	2.497 GB	28.63 MB	5.61 GB	14.91 GB
lxb265.gsi.de		0.281ms	3.18	37.99	1.618	0	60.39	7.801 GB	1.878 GB	1.593 GB	25.43 MB	5.923 GB	14.91 GB
lxb266.gsi.de		0.352ms	4.44	51.28	4.252	0	44.46	7.801 GB	2.445 GB	2.871 GB	27.26 MB	5.356 GB	14.91 GB
lxb267.gsi.de		0.346ms	4.26	50.73	3.736	0	45.54	7.801 GB	3.335 GB	4.407 GB	14.57 MB	4.466 GB	14.91 GB
lxb268.gsi.de		0.318ms	4.39	50.5	2.275	0	47.22	7.801 GB	2.716 GB	3.255 GB	27.95 MB	5.085 GB	14.91 GB
lxb281.gsi.de		0.25ms	0.03	0	0.006	0	99.99	15.7 GB	705.6 MB	1.609 GB	27.12 MB	15.01 GB	29.81 GB
lxb283.gsi.de		0.204ms	3.25	37.36	0.609	0	62.03	15.7 GB	2.702 GB	4.021 GB	27.75 MB	13 GB	29.81 GB

Fertig Internet

Start 5 Internet ... X-Win32-Kon... X-Win32 - St... root@grid5:~ DE 09:44

GSI AF usage experience

- see talk of M. Ivanov on Friday at 10:00
"Analysis Experience at GSI AF "
- real life analysis work of staged data by GSI ALICE group (1-4 concurrent users)
- 2 user tutorials for GSI ALICE users (10 students each training)
- GSI AF and CAF were used as PROOF clusters during PROOF course at GridKa School 2007

GridKa School 2007

GridKa School 2007 - Agenda - Microsoft Internet Explorer

Datei Bearbeiten Ansicht Favoriten Extras ?

Zurück Suchen Favoriten

Adresse <http://gks07.fzk.de/Agenda.html> Wechseln zu Links

	room number 157	Introduction to the Simple API for grid applications SAGA Tutorial	room: Aula
		room number 162	
20:00	"Kneipenabend"		

Thursday, September 13

9:00-9:45	Grid Access Manager (U. Biewer, IBM) abstract		
9:45-10:30	Research and Education on Microsoft HPC (Dr. Ingo Dahm, Microsoft) A Live Demonstration with Wolfram gridMathematica (Maryam Karbalai, Additive GmbH) MatLabMMA.zip abstract		
10:30-11:00	Break		
11:00-11:30	Grid Applications in D-Grid (W. Gentzsch, D-Grid) abstract		
11:30-12:00	Grid Business Models (J. Altmann, IU Bruchsal) abstract		
12:00-12:30	dCache (O. Synge, DESY) abstract		
12:30-14:00	Break		
14:00-19:00	Unicore course (R. Breu, FZ Juelich, and team) abstract	dCache course (O. Synge, DESY, and team) abstract Installing dCache with dCache Yaim dCache Clients	PROOF (J.F. Grosse-Oetringshaus and Andreas Joachim Peters, CERN) abstract info tar-ball
	room: Aula	room number 157	RGLite demo (Anar Manafov, GSI) abstract room number 162
20:00	School dinner		


Start Internet ... Microsoft Po... X-Win32-Kon... X-Win32 ... Rechner DE 12:59

plans and outview

- Study coexistence of interactive and batch processes (PROOF analysis on staged data and Grid user/production jobs) on the same machines. Develop possibility to increase/decrease the number of batch jobs on the fly to give advantage to analysis. Currently at GSI each PROOF worker is an LSF batch node
- optimise I/O. Various methods of data access (local disk, file servers via xrd, mounted lustre cluster) are being investigated systematically.
- extend GSI T2 and GSIAF according to promised ramp up plan

systematic I/O tests still under investigation

Preliminary results



lustre: 1 job: 200 events/sec up to 8 jobs: 1200 events/sec
more tests are on the way (also PROOF analysis on lustre)



ALICE T2 – ramp up plans

<http://lcg.web.cern.ch/LCG/C-RRB/MoU/WLCGMoU.pdf>

Germany, GSI, Darmstadt	Pledged	Planned to be pledged				
	2006	2007	2008	2009	2010	
CPU (kSI2K)	100	260	660	860	1100	
Disk (Tbytes)	30	80	200	260	340	
Nominal WAN (Mbits/sec)	100	100	1000	1000	1000	

GSI AF will grow with a similar rate

issues and problems

- instability of PROOF (or our setup). During the last user tutorial (10 users) we had to restart the PROOF master 3 times within 5 hours. It looks like a Reset of a messed up PROOF session can lead under certain circumstances to the end of the xrootd redirector.
- our first set of workers provide once in a while lengthy response times which is why PROOF refuses to connect and skips them as workers. This leads sometimes to hanging login sessions since PROOF does not seem able to skip WNs and continue at this stage. Reason: probably missing NIS cache
- helpful would be if in the xrootd log there would appear some key word like "error", e.g. if something goes wrong. This would facilitate finding problems like above significantly
- xrootd versioning in ROOT/AlIEn/SLAC and software dependencies (AliRoot/ROOT/xrootd/AlIEn)

POSIX

- GSI users want to know what files we have on the GSIAF cluster and want to be able to deal with their files in a POSIX like way.
- XrootdFS based and Fuse tested in collaboration with Andreas Petzold (TUD)
 - so far it seems to work only with individual data servers, not on the whole cluster

issues and problems

- how to bring data to GSIAF ?
- suggested method:
- see
<http://alien.cern.ch/twiki/bin/view/AliEn/HowToUseCollections>
 - create a collection of files
 - mirror the collection (and the files in it) to the new SE.
- But during last exercise of the 231614 files in /alice/sim/2007/LHC07c/pp_minbias/ only 42191 seem to have landed at GSI

issues and problems

- how to bring data to GSIAF ? Since GSIAF is a working AliEn SE we would like to access the stored files directly on the SE (local disks) via the same names the files would have in the AliEn FC without additional staging from the closest SE.
- Do PROOF Datasets (see talk of Jan Fiete) actually work, now ? Is this a solution to our issue ?

