

Availability of ALICE Grid resources in Germany

Kilian Schwarz
GSI Darmstadt
ALICE Offline Week

Overview

- ◆ installed resources for ALICE Grid computing in Germany (especially GridKa) are shown and how they compare to the pledged resources as well as the resources being actually used by Grid jobs.
- ◆ CPU: the pledged resources are not fully used. An analysis within this talk discusses the reasons.
- ◆ possible methods to improve the situation
- ◆ summary and conclusions

resources at GridKa (CPU)

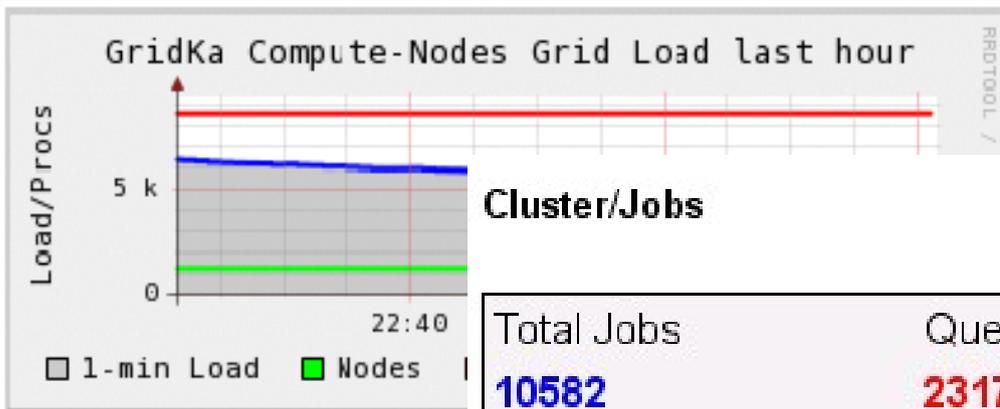
→ all requested resources are installed and ready to be used !!!

GridKa Compute-Nodes Grid (tree view)

CPU's Total: **8597**
Hosts up: **1267**
Hosts down: **51**

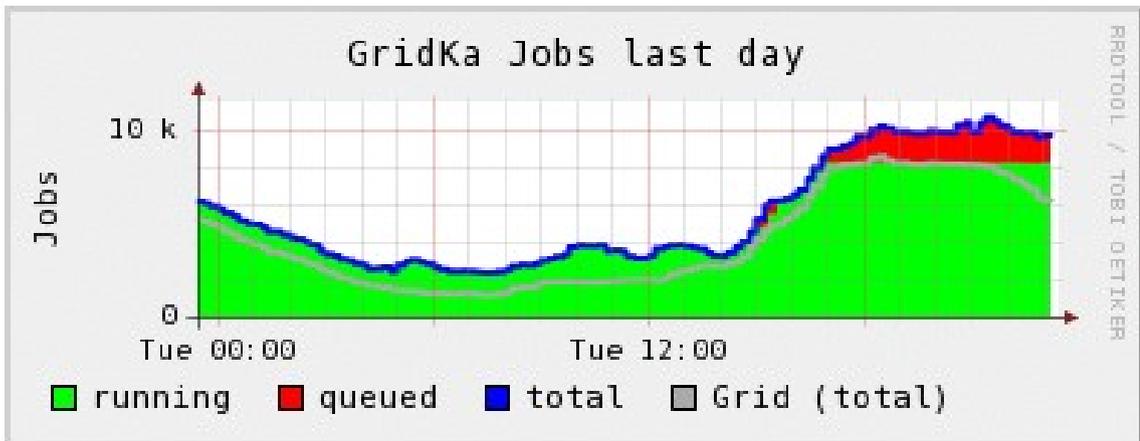
Avg Load (15, 5, 1m):
69%, 69%, 69%

Localtime:
2009-10-27 23:21



Total Jobs	Queued	Running
10582	2317	8264

nominal ALICE share:
38% of CPUs



Jobs finished today: 27501 Jobs finished this hour: 907

GridKa CPU Share

statistics
September
2009

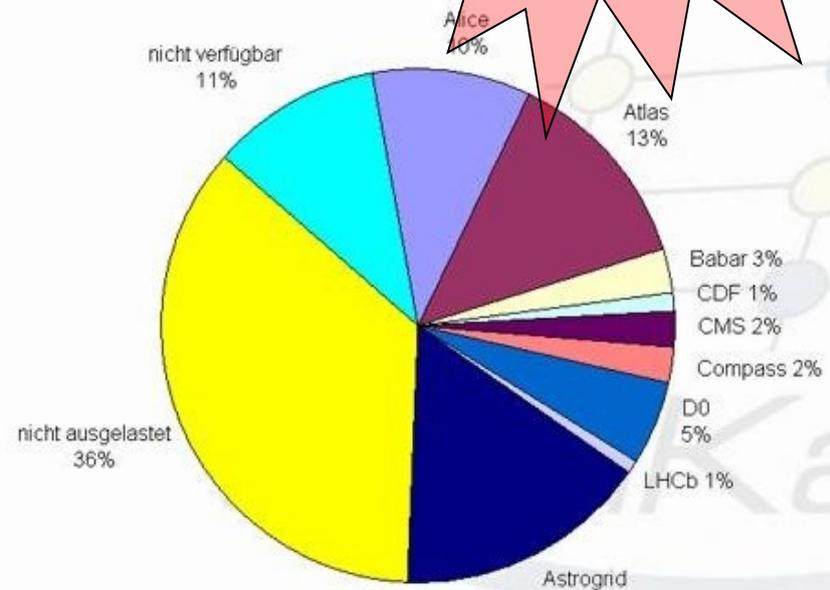
http://www.gridka.de - neu berechnete Fairshare-Werte - Iceweasel

Neu berechnete Fair-Share-Werte:

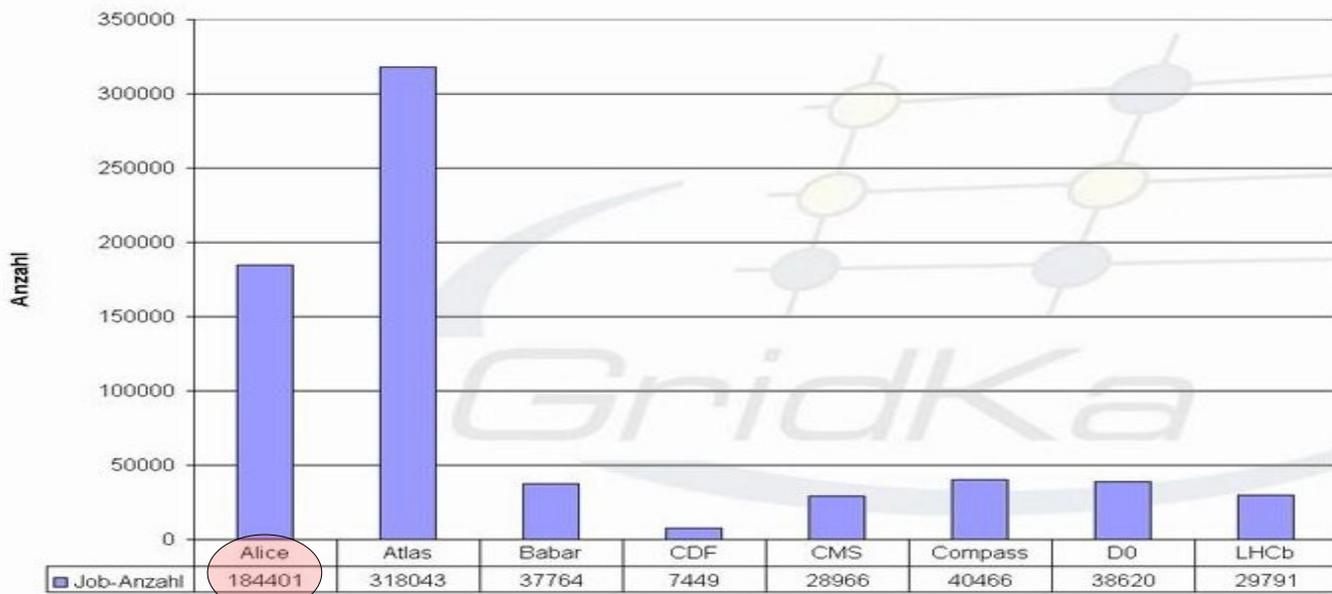
Seit 8.2.2006 werden die neu berechneten Fair-Share-Werte mit Korridor täglich ins PBS eingetragen
letzte Änderung: Wed Oct 7 04:02:35 2009

Wed Oct 7 04:02:35 2009	alice	atlas	babar	cdf	cms	compass	dzero	lhcb
nominell %	38.13	25.91	3.81	2.10	14.01	1.56	8.17	6.32
verbraucher Wall-Zeit % half_life=180 Tage	25.36	31.48	3.09	3.71	7.81	2.60	16.04	9.91
Gewicht Vorschlag TAB	57.33	21.33	4.70	1.19	25.13	0.93	4.16	4.03
Korridor 1/2 - 2	57.33	21.33	4.70	1.19	25.13	0.93	4.16	4.03
nomiert % Korridor	48.26	17.95	3.96	1.00	21.15	0.78	3.50	3.39
verbrauchte Wall-Zeit % half_life=1 Tag	1.95	40.88	3.03	1.1				
Usage/Perc Korridor	87.07	4918.85	1656.09	2196.0				

usage/perc = verbrauchte Wall-Zeit mit Halbwert (nomiert)



Anzahl Jobs pro Experiment



→ If CPUs are not requested by ALICE they can be used easily by other experiments

storage at GridKa

19. FZK - DCACHE	ALICE::FZK::DCACHE	423.8 TB	6.773 TB	417 TB	1.598%	69,917	SRM						07.10.2009
20. FZK - DCACHE_SINK	ALICE::FZK::DCACHE_SINK	2.728 PB	54.96 GB	2.728 PB	0.002%	1,380	SRM						07.10.2009
21. FZK - DCACHE_TAPE	ALICE::FZK::DCACHE_TAPE	2.728 PB	42.83 TB	2.687 PB	1.533%	63,169	SRM						07.10.2009
22. FZK - SE	ALICE::FZK::SE	322.3 TB	38.17 TB	284.1 TB	11.84%	553,762	File						07.10.2009

- ◆ **new SE to follow: ALICE::FZK::TAPE**
- ◆ full xrootd based storage solution in the sense of WLCG (SRM interface, gridftp interface, tape backend, many more features)
- ◆ **The 1.5 PB disk space installed at GridKa are to a large extend empty !!! Tapes are basically not used at all !!!**
- ◆ **If the storage devices are not used the disks stay empty. They can not be used by other experiments easily.**

resources at GSI

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	

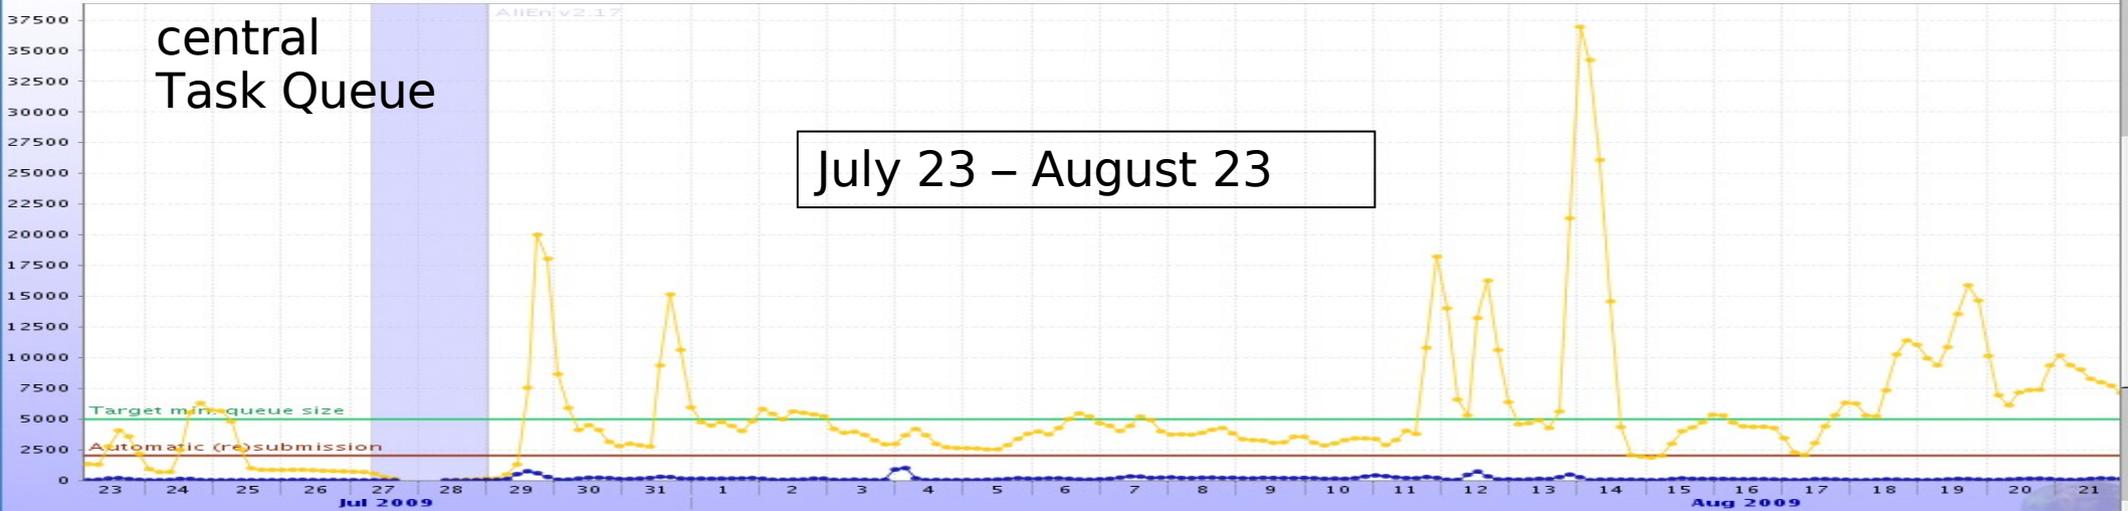
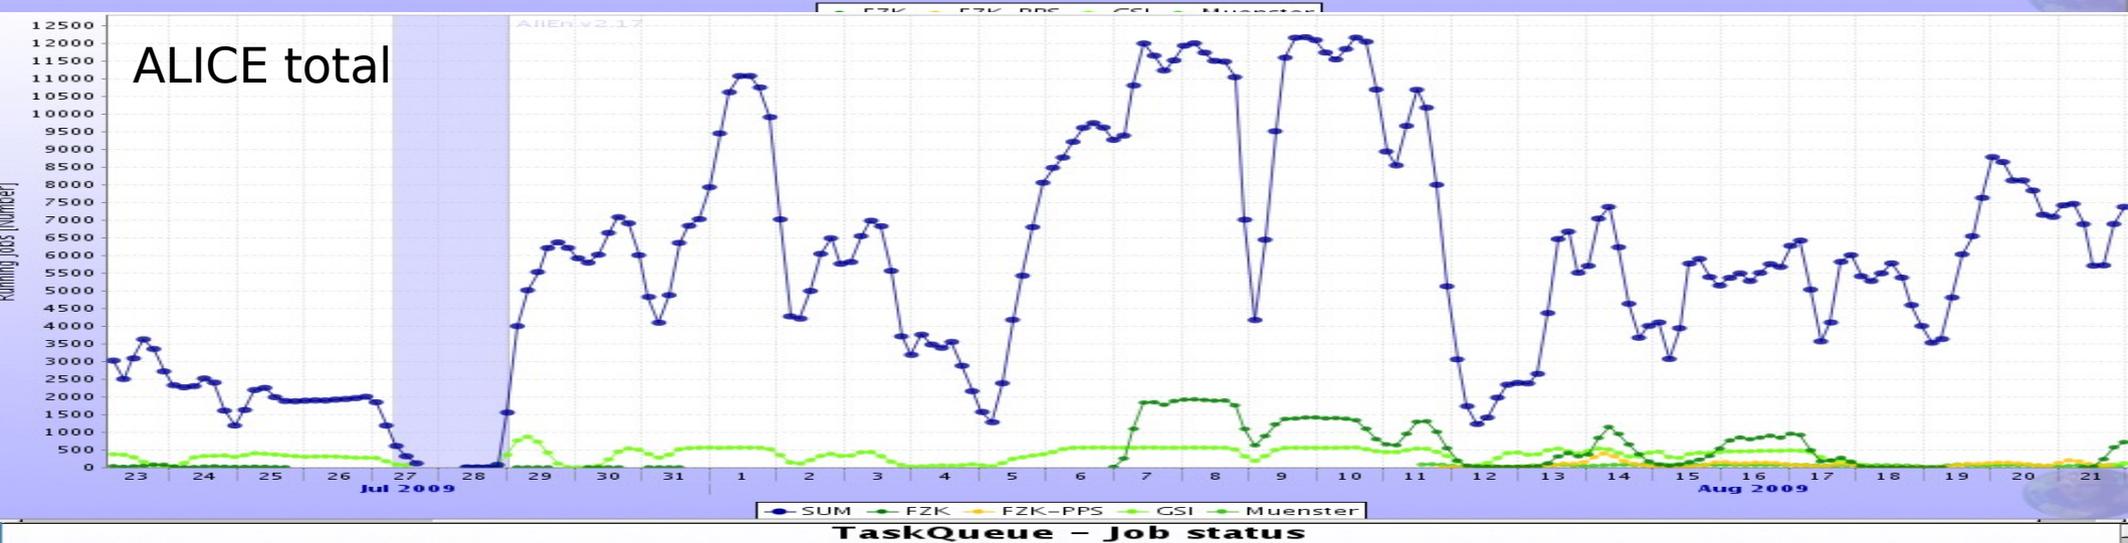
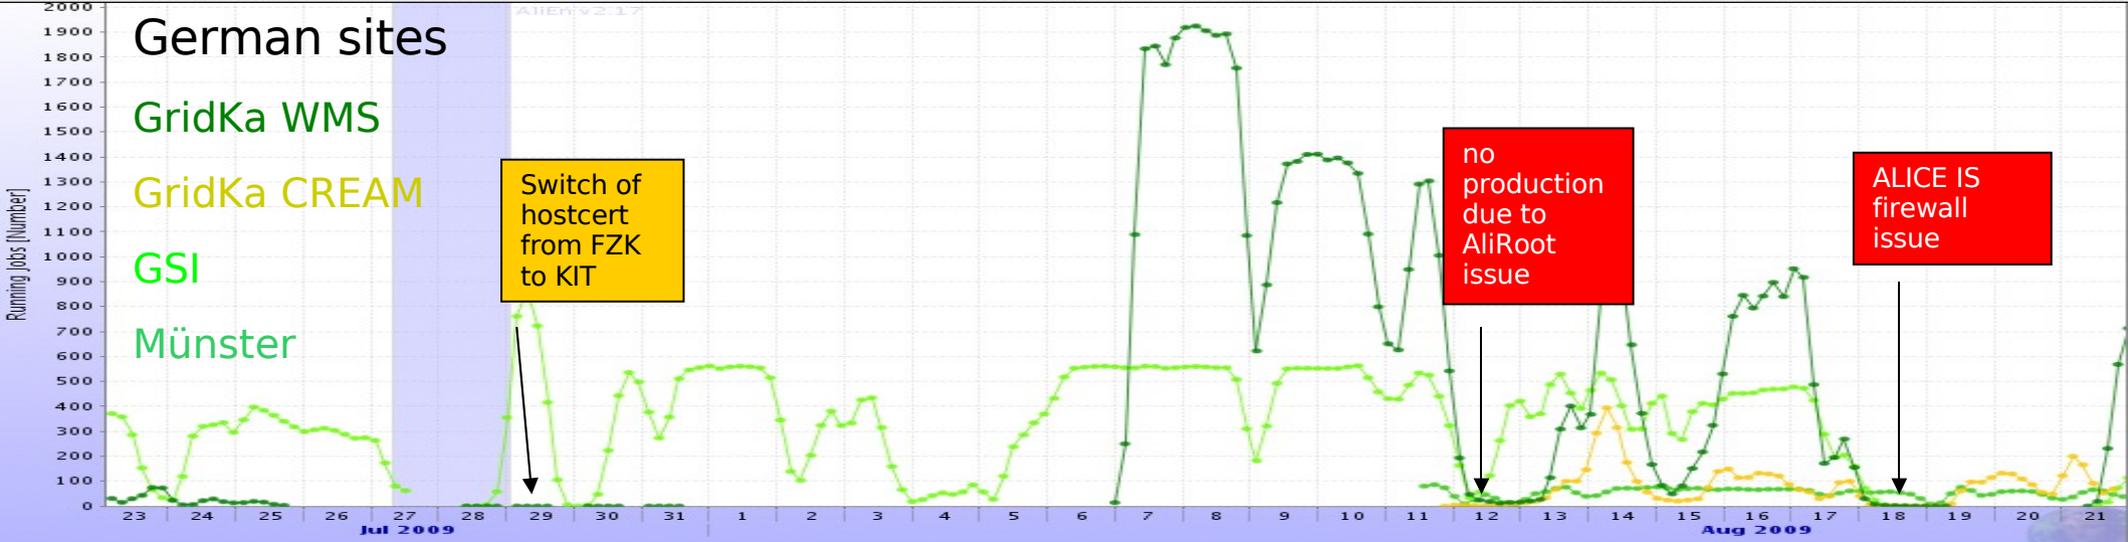
- ◆ CPUs: installed
- ◆ Disk: installed

***analysis of consumed
CPU resources by ALICE
Grid jobs at German
sites***

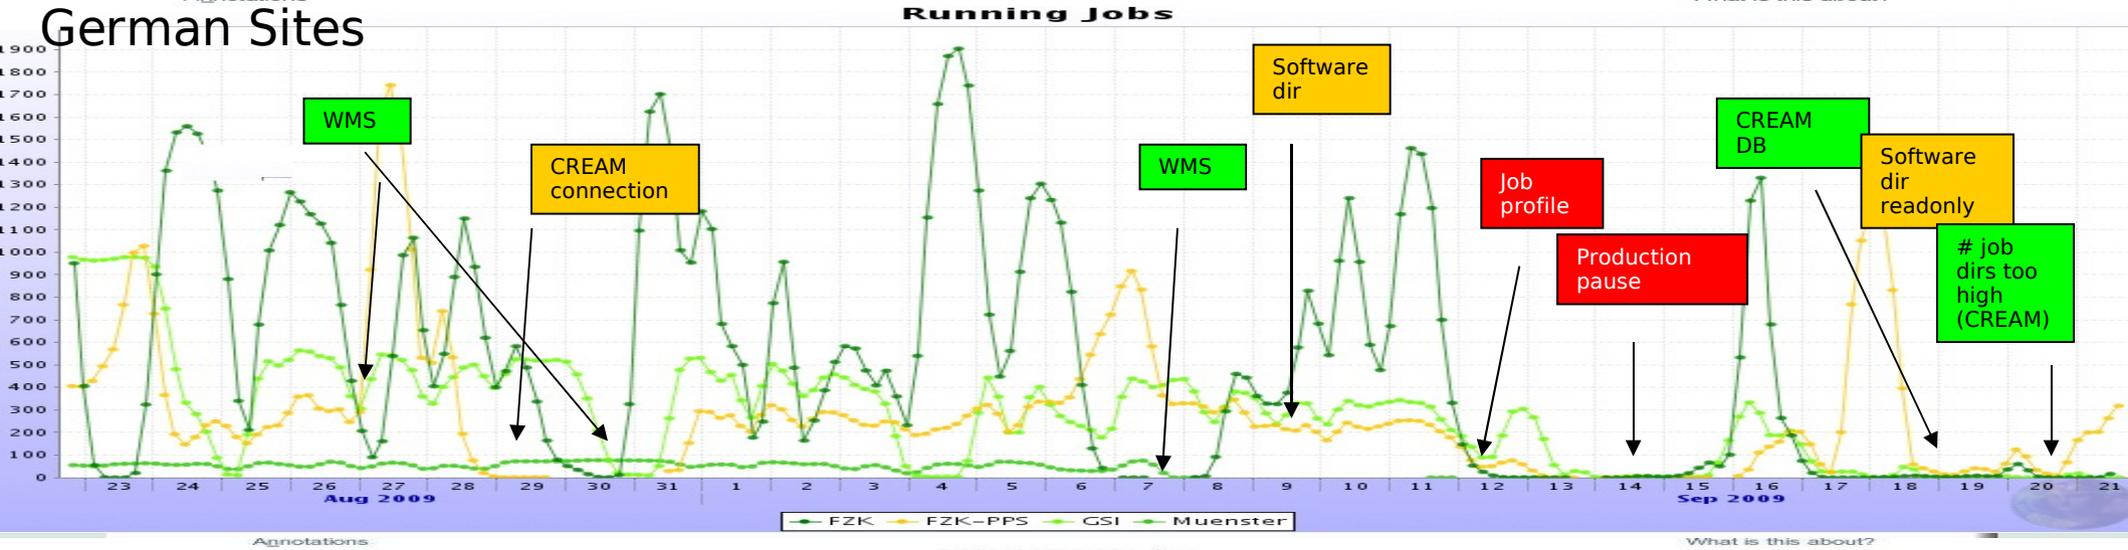
the following 3 slides show 3 run
periods of 1 month each within
the time span July to October

-

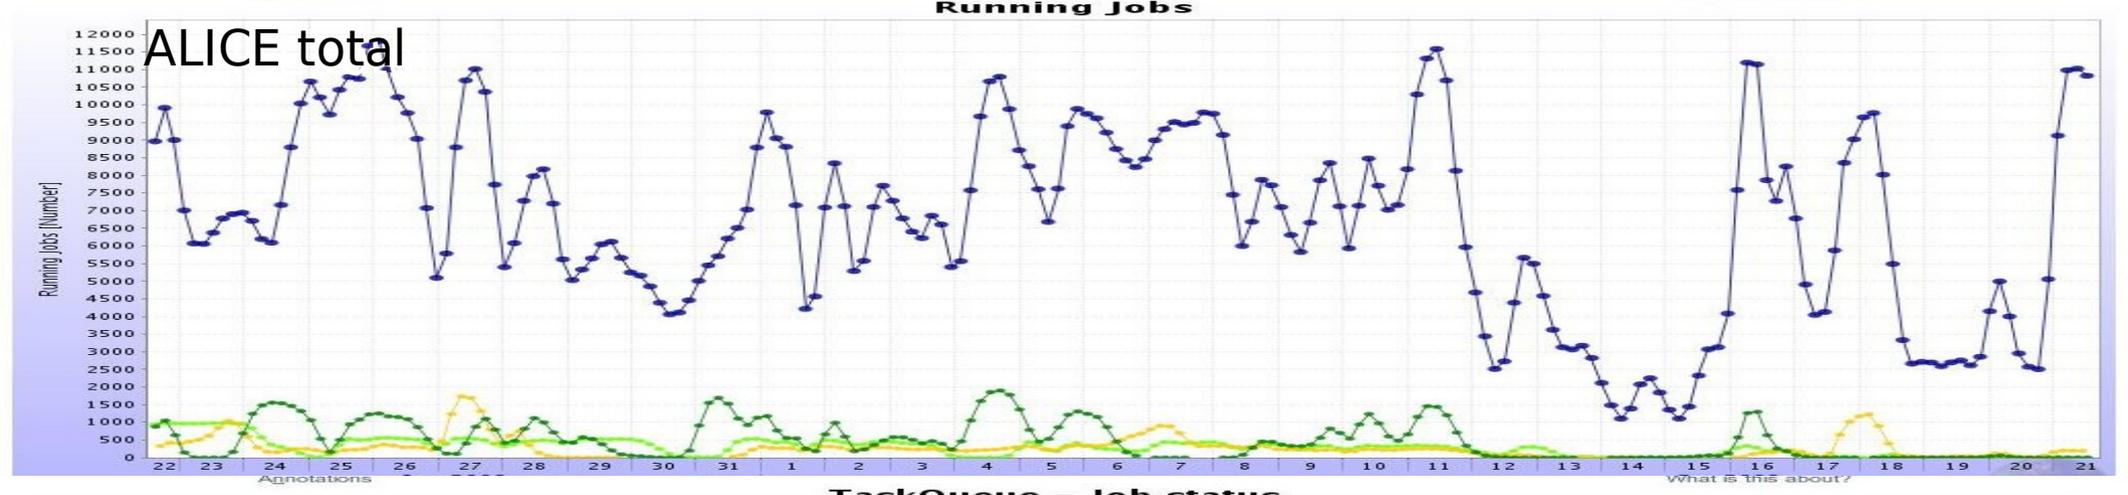
issues at GridKa are shown in the plots



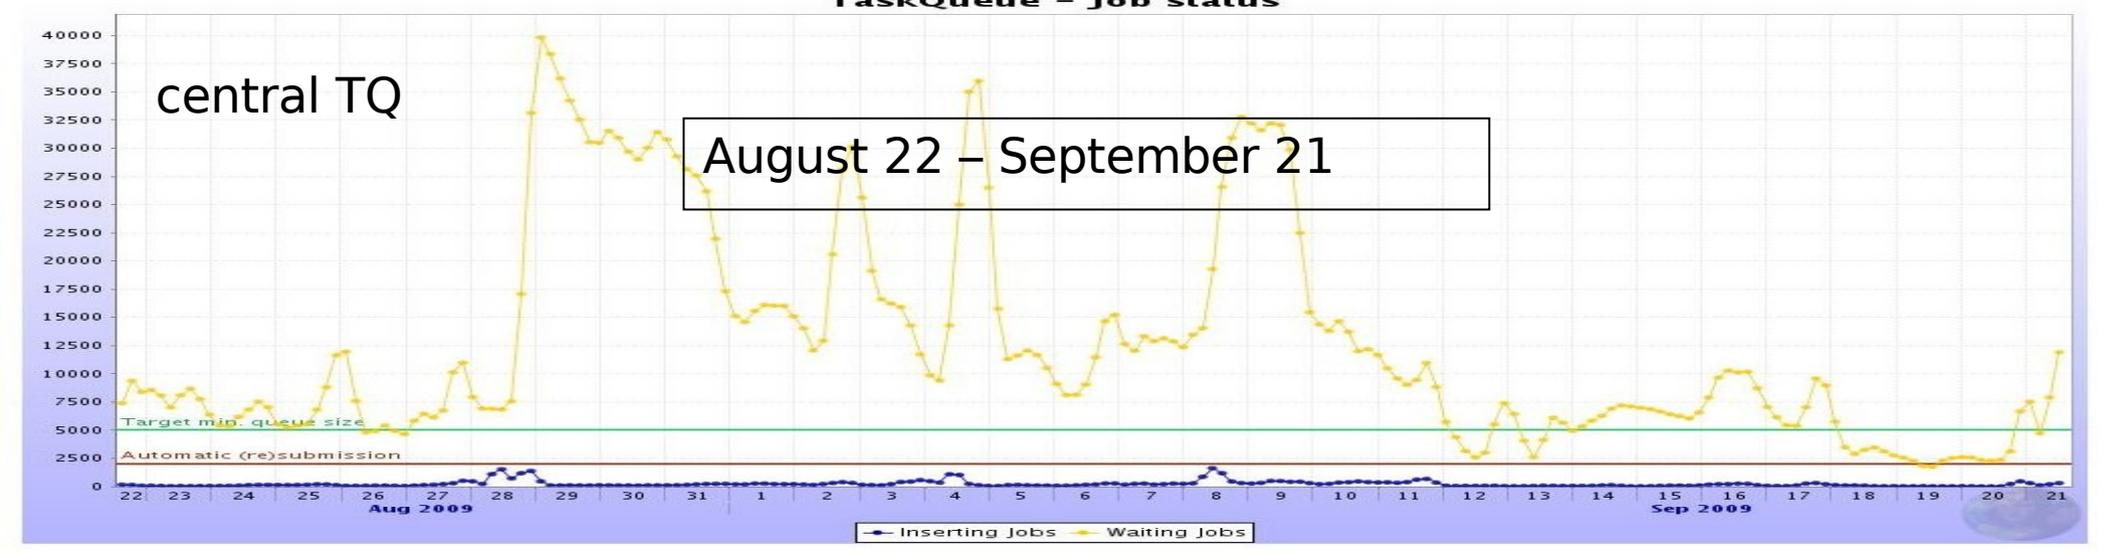
German Sites

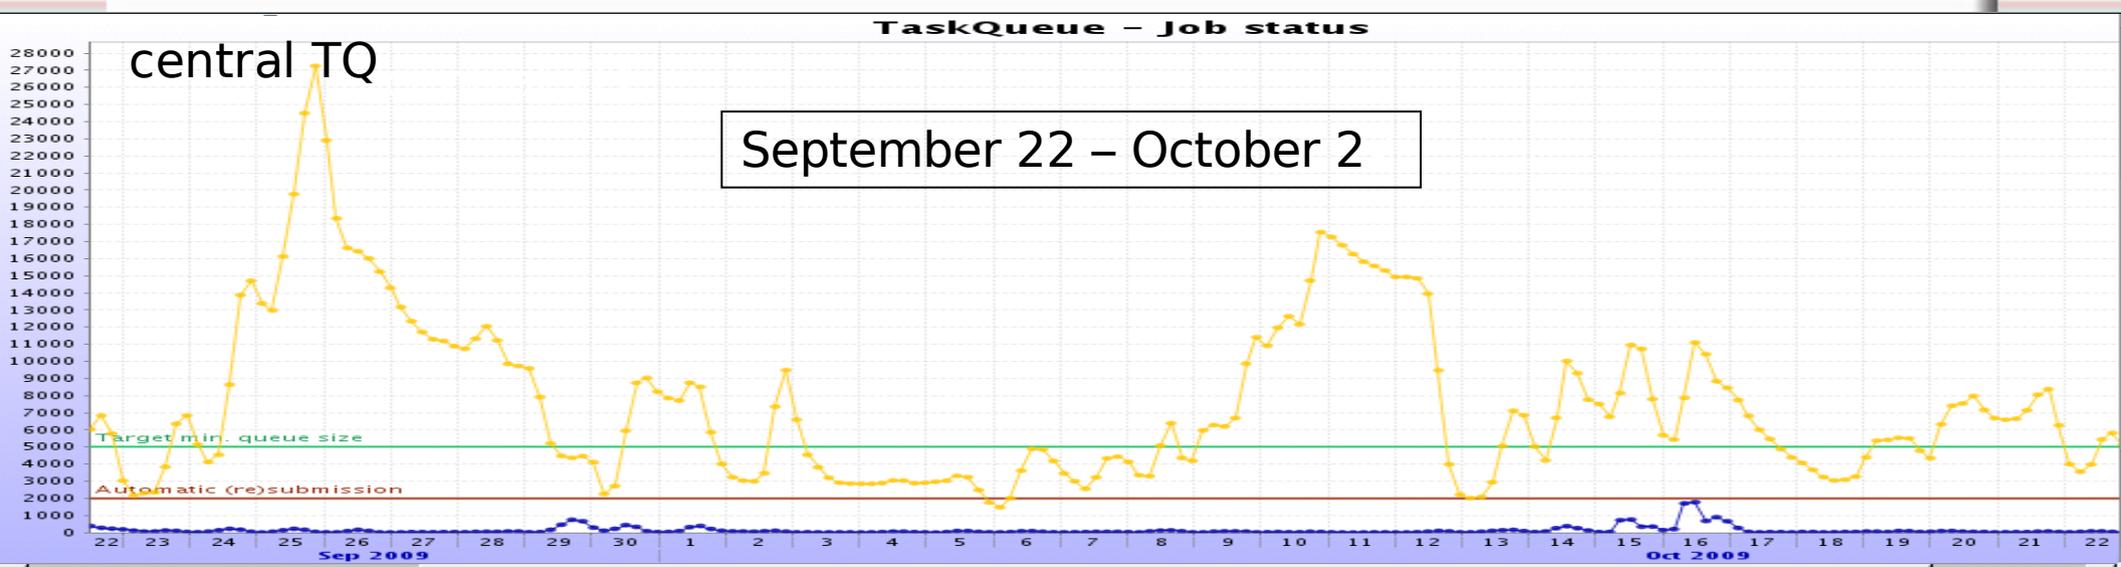
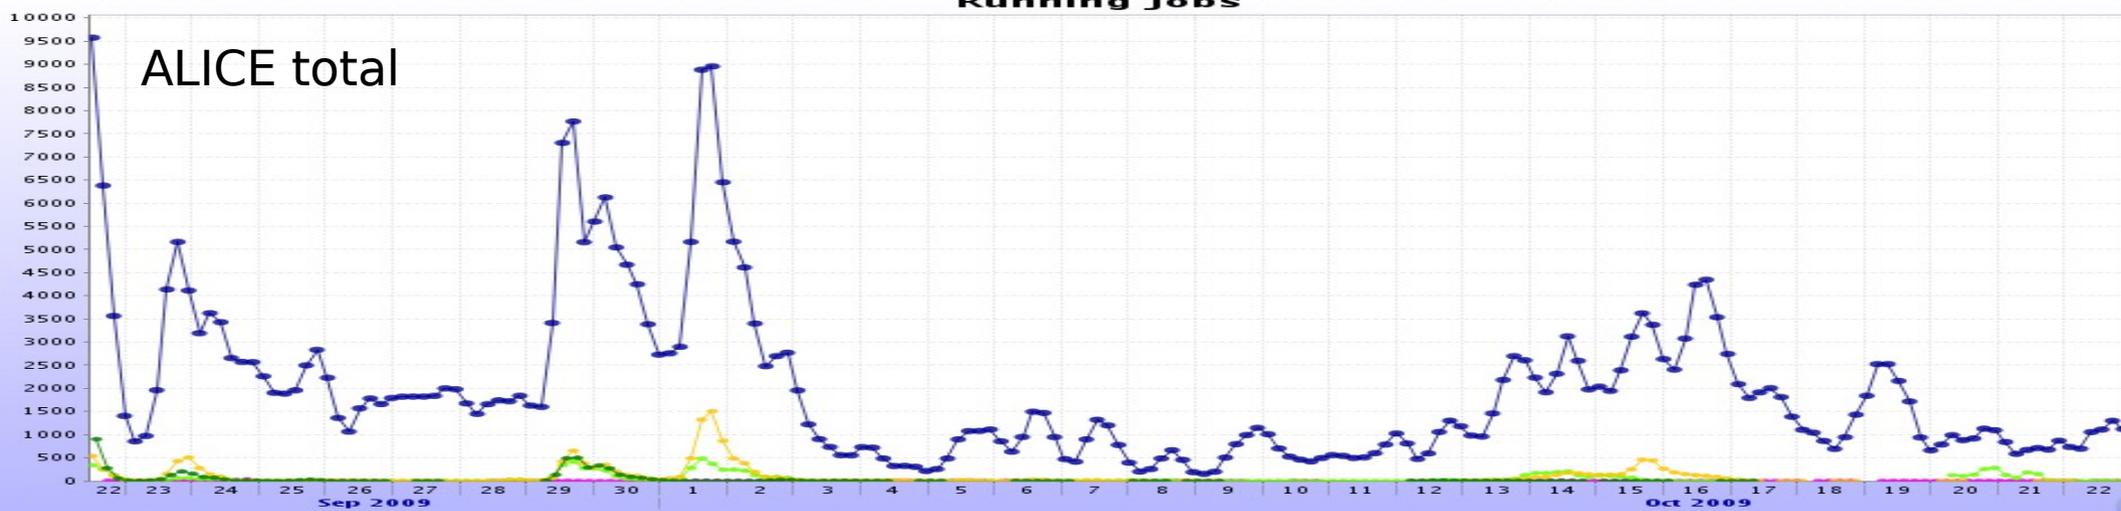


ALICE total



central TQ





pledged resources

Report on ALICE groups' activity (01.08.2009 - 31.08.2009)

Group	Pledged	Delivered		Occupancy	Missing KSI2K	Efficiency	Job statistics		
	KSI2K	CPU	Wall	Wall/Pledged	Pledged - Wall	CPU/Wall	Assigned	Completed	Efficiency
1. Brasil	0	-	-	-	-	-	-	-	-
2. CERN	11600	740.2	835.6	7.203%	10764	88.58%	65733	47610	72.43%
3. China	0	-	-	-	-	-	-	-	-
4. Czech Republic	154	311	335.7	218%	-	92.66%	14958	6339	42.38%
5. Germany	5892	1857	2063	35.02%	3828	90.04%	302776	128262	42.36%
6. Greece	80	48.31	53.64	67.05%	26.36	90.06%	15888	961	6.049%
7. HLT	0	-	-	-	-	-	-	-	-
8. Hungary	100	45.69	48.57	48.57%	51.43	94.07%	16118	1551	9.623%
9. IN2P3	4419	1862	2211	50.04%	2207	84.23%	244531	112482	46%
10. India	350	73.72	101.7	29.06%	248.3	72.48%	8997	1931	21.46%
11. INFN	2448	984.2	1382	56.46%	1065	71.21%	304801	111870	36.7%
12. Japan	0	59.62	67.3	-	-	88.59%	4230	1673	39.55%
13. Mexico	22	6.704	7.356	33.44%	14.64	91.13%	714	253	35.43%
14. Nordic Countries	2160	1301	1566	72.53%	593.4	83.11%	144664	75722	52.34%
15. Other	0	-	-	-	-	-	-	-	-
16. Pakistan	0	-	-	-	-	-	-	-	-
17. Poland	410	394.3	471.4	115%	-	83.65%	38123	10076	26.43%
18. RDIG	1059	622.4	693.8	65.51%	365.2	89.72%	54210	27494	50.72%
19. Republic of Korea	183	12.19	13.38	7.31%	169.6	91.1%	920	372	40.43%
20. Romania	1379	1290	1791	129.9%	-	72.05%	211630	103418	48.87%
21. Slovakia	80	37.4	40.06	50.07%	39.94	93.35%	3984	2337	58.66%
22. South Africa	0	-	-	-	-	-	-	-	-
23. Spain	246	52.99	58.03	23.59%	188	91.32%	3002	528	17.59%
24. The Netherlands	318	67.07	69.82	21.96%	248.2	96.05%	3266	687	21.03%
25. UK	336	30.71	33.31	9.913%	302.7	92.19%	2131	879	41.25%
26. Ukraine	250	28.15	30.07	12.03%	219.9	93.6%	11224	1039	9.257%
27. US	2371	153.9	196.6	8.29%	2174	78.31%	79809	4137	5.184%
Total	33857	9981	12070		22509		1531709	639621	

analysis

- ◆ GridKa various issues (MW, GridKa, ALICE) – during first run period issues were more GridKa related, in second run period issues were equally distributed, in last run period the issues are more ALICE related
- ◆ in the last run period almost no region did well compared to pledged resources due to missing Grid jobs
- ◆ big oscillations in job pattern also due to complicated interplay between AliEn and LCG services
- ◆ GSI runs more stable
- ◆ before July 27 and end of September: no official statement but number of Jobs waiting in TQ on lower end
- ◆ GridKa belongs with up to 15% to the largest producers of DONE jobs
- ◆ in terms of delivered CPUs usually Germany belongs to the largest groups
- ◆ ALICE is not able to fully use the German pledged resources, though, with the largest deficit being at GridKa

Suggestions how to improve things

- ◆ **GridKa** related issues
 - ◆ More manpower
 - ◆ Better communication
 - ◆ Participation in related meetings
 - ◆ Better monitoring
 - ◆ continuous testing via small productions and corresponding analysis jobs
- ◆ In terms of resources: GSI could pledge more resources to the Grid, resource distribution among German Grid sites can be done in a flexible way to ensure that ALICE Grid gets what it needs
- ◆ **Middleware** related issues
 - ◆ Better monitoring (e.g. at CERN exists a technique to monitor overladed WMS)
 - ◆ Also here participation in related meetings
- ◆ **ALICE** related issues
 - ◆ Production free time should be reduced and better communicated
 - ◆ if no central production jobs then user jobs should fill the gap
 - ◆ resource monitoring should be folded with availability of Grid jobs (e.g. jobs matching with site)

current situation

- ◆ actually, a lot of things suggested (see slide before) have already been implemented
- ◆ as soon as production restarts we will see if things improved

AOB

- ◆ missing tape space
 - ◆ GridKa pledged a significant part of the requested T1 tape capacity
 - ◆ we are investigating if more tape space can be provided on short notice

summary and conclusion

- ◆ ALICE is currently not able to use the pledged resources of German Grid sites to a full extend
- ◆ reasons for this are many folded (**site problems**, **Middleware problems**, **missing production jobs**)
 - ◆ **site problems** could be solved with higher response time
 - ◆ **MW problems** could be monitored better (e.g. overloaded WMS following a procedure existing at CERN)
 - ◆ but if **production jobs** are missing the sites can not fulfill the request for pledged resources
- ➔ the delivered CPUs have to be folded with the number of existing jobs (e.g. "jobs match the site resources")

from the hardware point of view all resources are installed and available