Karlsruher Institut für Technologie

Grid resources and services for multiple science disciplines at KIT

Resource Centre Forum @ EGI Technical Forum
Lyon, 2011

Steinbuch Centre for Computing **Dr. Andreas Heiss**

GridKa, SCC and KIT ~~FZK, IWR~~

- KIT ist the merger of the Research Centre Karlsruhe and the University of Karlsruhe.
- Steinbuch Centre for Computing (SCC): the merger of the former Institute for Scientific Computing of the Research Centre Karlsruhe and the compute centre of the University of Karlsruhe.
- GridKa is the regional centre for high energy and astroparticle physics computing, currently supporting 10 international collaborations including the four LHC experiments.
- GridKa is operated by SCC

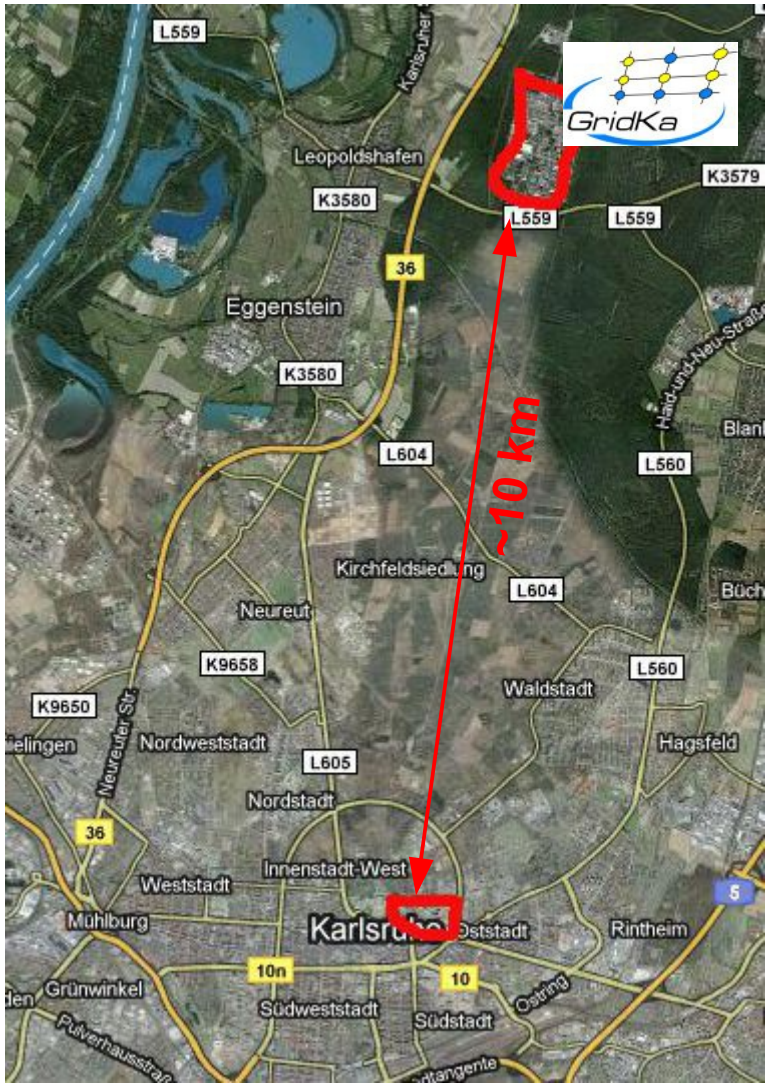

Grid projects and user communities

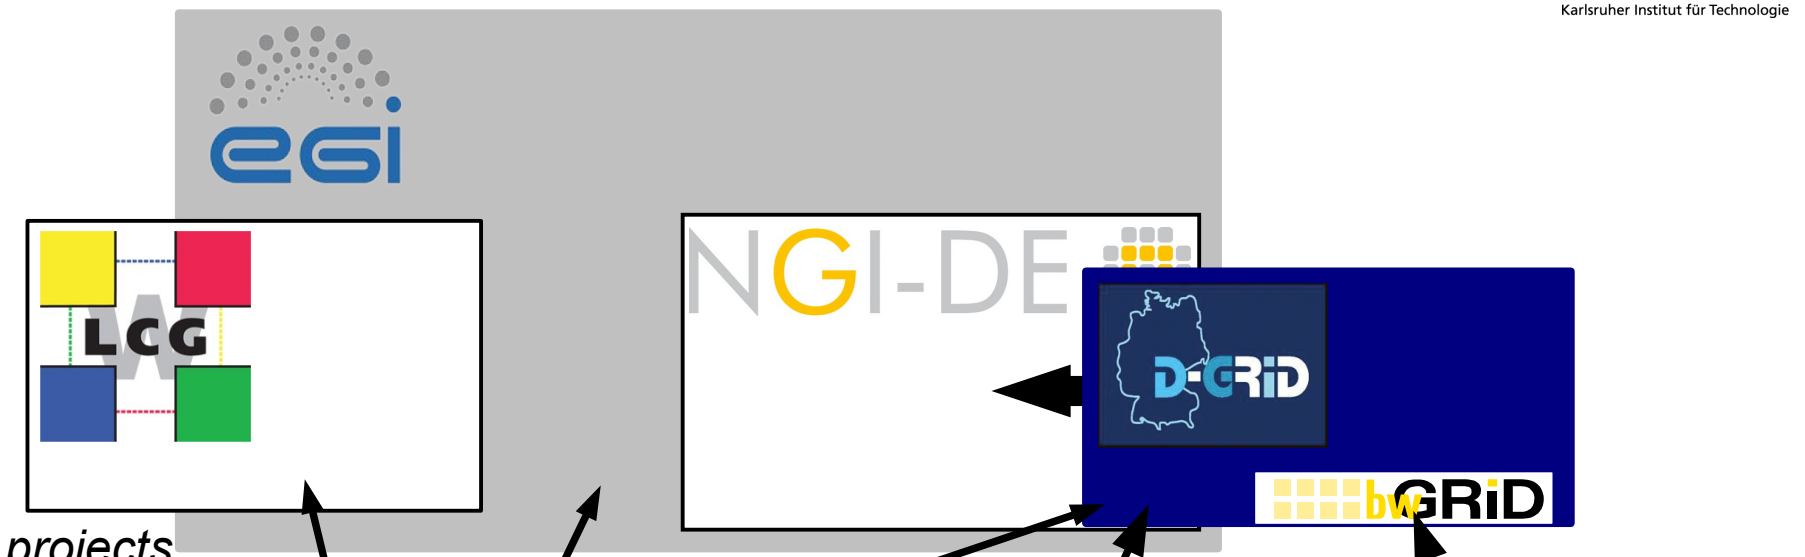

GridKa

HEP and astroparticle physics, WLCG Tier-1

Opus^{IB}

Infiniband-Cluster for Grid and local use

bwGRiD

bwGRID Infiniband Cluster and storage

Access via gLite, Globus Toolkit 4 and Unicore

bwGRID partners

Bundesministerium
für Bildung
und Forschung

Karlsruhe Institute of Technology

ulm university universität ulm

Universität Stuttgart

UNIVERSITÄT MANNHEIM

RUPRECHT-KARLS-
UNIVERSITÄT
HEIDELBERG

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

ZDV Zentrum für Datenverarbeitung
im IKM Informations-, Kommunikations- und Medienzentrum

ALBERT-LUDWIGS-
UNIVERSITÄT FREIBURG

Hochschule Esslingen
University of Applied Sciences

bwGRID users

Karlsruher Institut für Technologie

Discipline (# of projects):
Astrophysics (5)
Biology (21)
Chemistry (27)
Economic science (9)
Informatics (7)
Mathematics (1)
Physics (17)
Political science (2)
Social science (2)

e.g. neuro science:

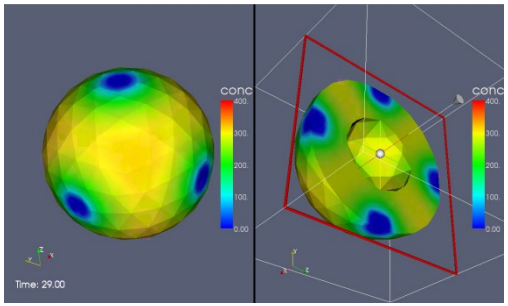

e.g. quantum chemical analysis:

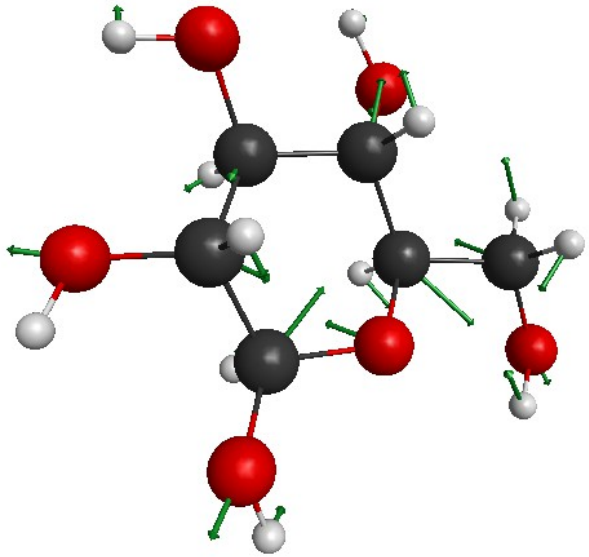

e.g. geophysics

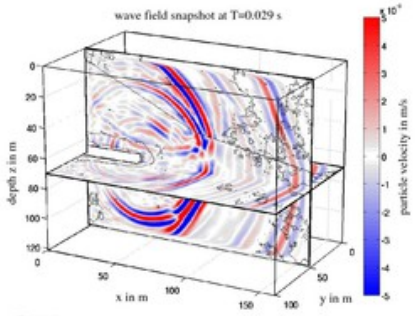

GridKa today: resources and services for HEP and Astroparticle physics

- Started in 2002
- GridKa serves all 4 of the LHC experiments as a 'Tier-1' centre.
- GridKa supported non-LHC experiments:

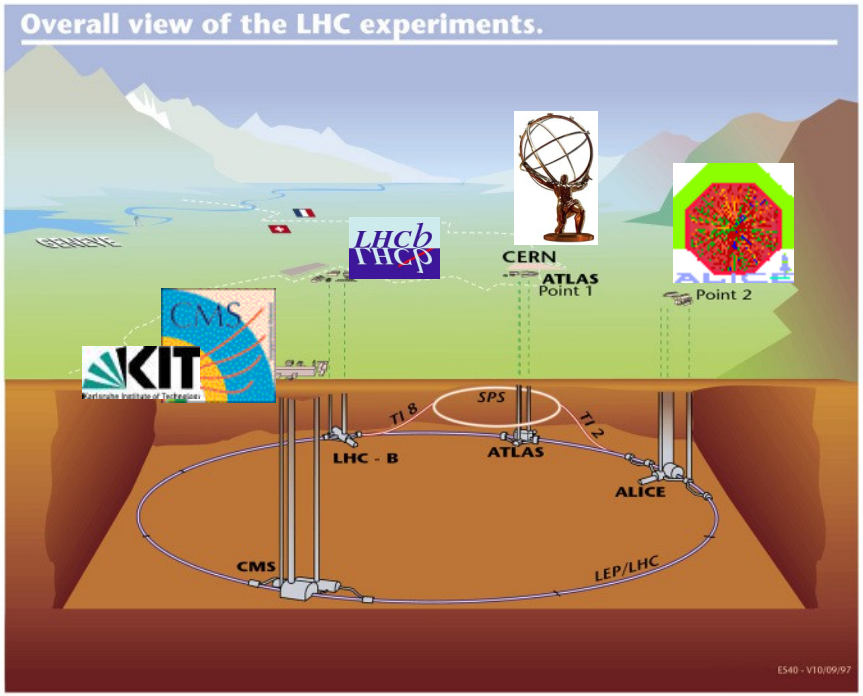

GridKa today

- 1270 compute nodes
  (13700 CPU cores, 107 kHEPSPEC'06)
- 240 file and database servers (dCache, xrootd, 3D DBs, gridftp, NFS, ...)
- ~90 servers for (Grid) services
- 12 VO boxes and login machines
- ~60 infrastructure servers (monitoring, installataion and configuration management, ...)

GridKa today

- **High I/O hierarchical storage systems**
- dCache (4 instances)
  - 6500 TB disk
  - ~ 2 GB/s yearly average read rate
  - ~ 500 MB/s yealy av. write
  - peaks of >6 GB/s
  - > 60 PB of data read last year
- xrootd (2 instances)
  - 2700 TB disk
  - ~ 500 MB/s yearly average I/O
  - peaks > 1GB/s
- 3 tape libs, > 25000 slots

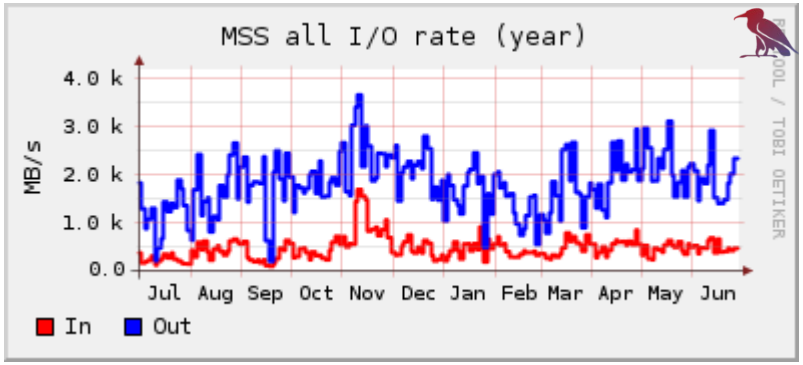

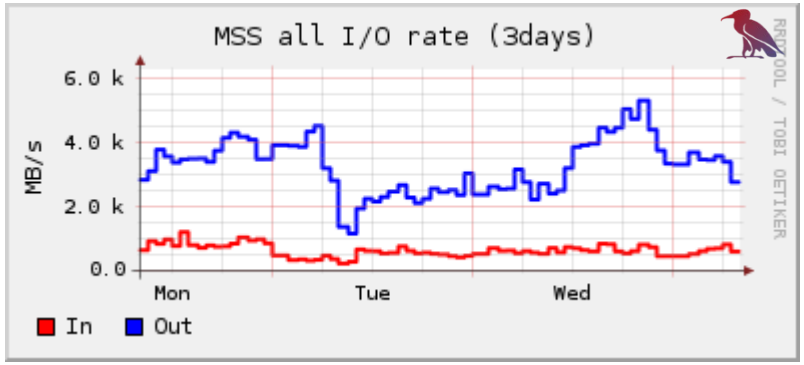

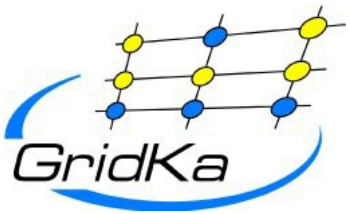

CPU usage
per VO
[CPU h]

Astro physics
HEP
Astro particle physics

D-Grid VOs	Kürzel	Januar	Februar	März	April	Mai	Juni	Total
AeroGrid	ae	0	0	0	0	0	0	**0**
AstroGrid-D	ad	218745	31461	464	281	530	3816	**255297**
BauVOGrid	ba	0	0	0	0	0	0	**0**
BioInterfaces	if	0	0	0	0	0	0	**0**
BisGrid	bi	0	0	0	0	0	0	**0**
Biz2Grid	bz	0	0	0	0	0	0	**0**
BWGrid	bw	0	0	0	24	0	0	**24**
C3Grid	c3	0	0	0	0	0	0	**0**
dgcms	cm	11874	42304	19701	5088	68918	87812	**235696**
dgops	op	0	0	0	0	0	0	**0**
dgsi	si	0	0	0	0	0	0	**0**
dgtest	dt	0	0	0	0	0	0	**0**
Education	ed	0	0	0	0	0	0	**0**
FinGrid	fi	0	0	0	0	0	0	**0**
GapSLC	gs	0	0	0	0	0	0	**0**
GDIGrid	gd	0	0	0	0	0	0	**0**
HEPCG	hp	0	0	0	0	0	0	**0**
InGrid	in	0	0	0	0	0	0	**0**
InterLogGrid	il	0	0	0	0	0	0	**0**
Kern-D-Grid	kg	0	0	0	0	0	0	**0**
Lifescience	ls	0	0	0	0	0	0	**0**
MediaGrid	mg	0	0	0	0	0	0	**0**
MediGrid	md	0	0	0	0	0	0	**0**
MosGrid	ms	0	0	0	0	0	0	**0**
m3hpc	mh	0	0	0	0	0	0	**0**
OptiNum	on	0	0	0	0	0	0	**0**
PartnerGrid	pa	0	0	0	0	0	0	**0**
PneumoGrid	pn	0	0	0	0	0	0	**0**
PTGrid	pt	0	0	0	0	0	0	**0**
ProGrid	pr	0	0	0	0	0	0	**0**
SoftComp	sc	0	0	0	0	0	0	**0**
TextGrid	tx	0	0	0	0	0	0	**0**
ValueGrids	va	0	0	0	0	0	0	**0**
Viola	vi	0	0	0	0	0	0	**0**
Wisent	ws	0	0	0	0	0	0	**0**
VOs außerhalb D-Grid		1266149	4925346	5831233	6117705	6219152	5424913	29784461
Alice		195048	1429227	659859	1046136	268292	869939	**4468501**
Atlas		512690	1383541	2673025	2673561	3162570	2724052	**13129439**
Babar		45067	141659	8043	58573	16019	4995	**274356**
CDF		18735	0	5	1	2107	57763	**78611**
CMS		171325	244400	940332	595238	601038	189411	**2741744**
Compass		4214	24976	69439	36259	92925	64694	**292507**
Dzero		76478	189538	238352	193454	359601	185729	**1243153**
LhcB		58090	1189002	845611	1073438	1151195	692235	**5009570**
auger		168518	309411	348679	431617	354152	518712	**2131089**
belle		2	928	7	30	195830	85419	**282216**
datlas		15945	12664	47880	9399	15423	31964	**133275**
dech		0	0	0	1	0	42	**44**
dteam		2	3	4	5	5	4	**22**
enmr		0	0	0	0	0	0	**0**
ops		35	49	52	95	0	51	**282**

GridKa and OpusIB usage

	VO	Kürzel	# jobs	CPU-h	# jobs	CPU-h
D-Grid	AstroGrid-D	ad	2.338	51.200	225.354	255.297
	bwGRiD	bw	66	50	80	24
	dgtest	dt	4.586	35	2.276	0?
	Education	kg	24	0	1	0?
	Kern-D-Grid	kg	259	14	174	0?
	MediGrid	md	24.329	2.978	--	--
andere	Euforia		91	1	--	--
	gEclipse		4	0	--	--
			OpusIB		GridKa	

Different types of user communities

- The less demanding
  - uses no or almost no storage
  - low I/O jobs (e.g. Monte Carlo only)
  - managed job submission (no e.g. user analysis)
  - has low expectations on short-term availability of job slots, data and services
  - Not much efforts or special services necessary

- 'Power user communities'
  - uses Grid mainly as a data Grid
    - complex data management
    - high I/O jobs
    - disk + tape
  - high expectations on availability of services and data
  - mix of different job types, also user analysis

GridKa Services: gLite

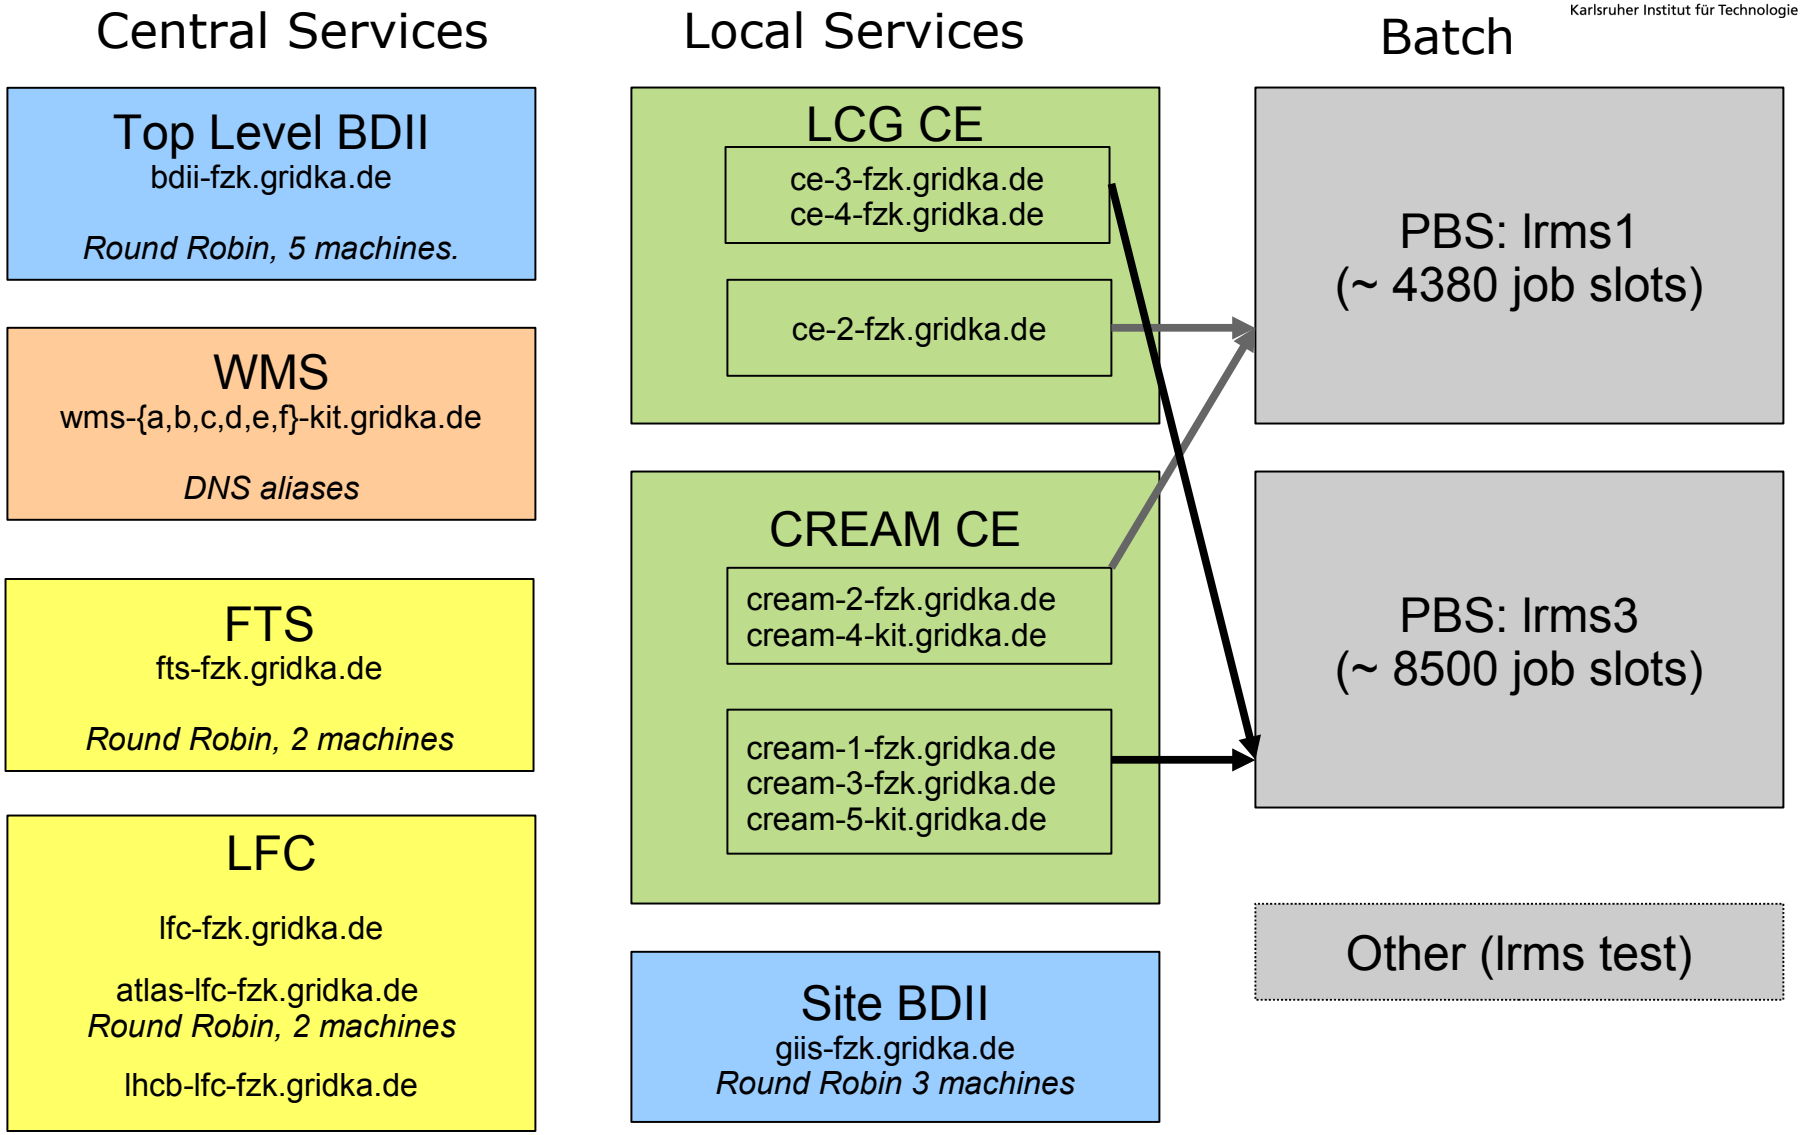

Grid Services

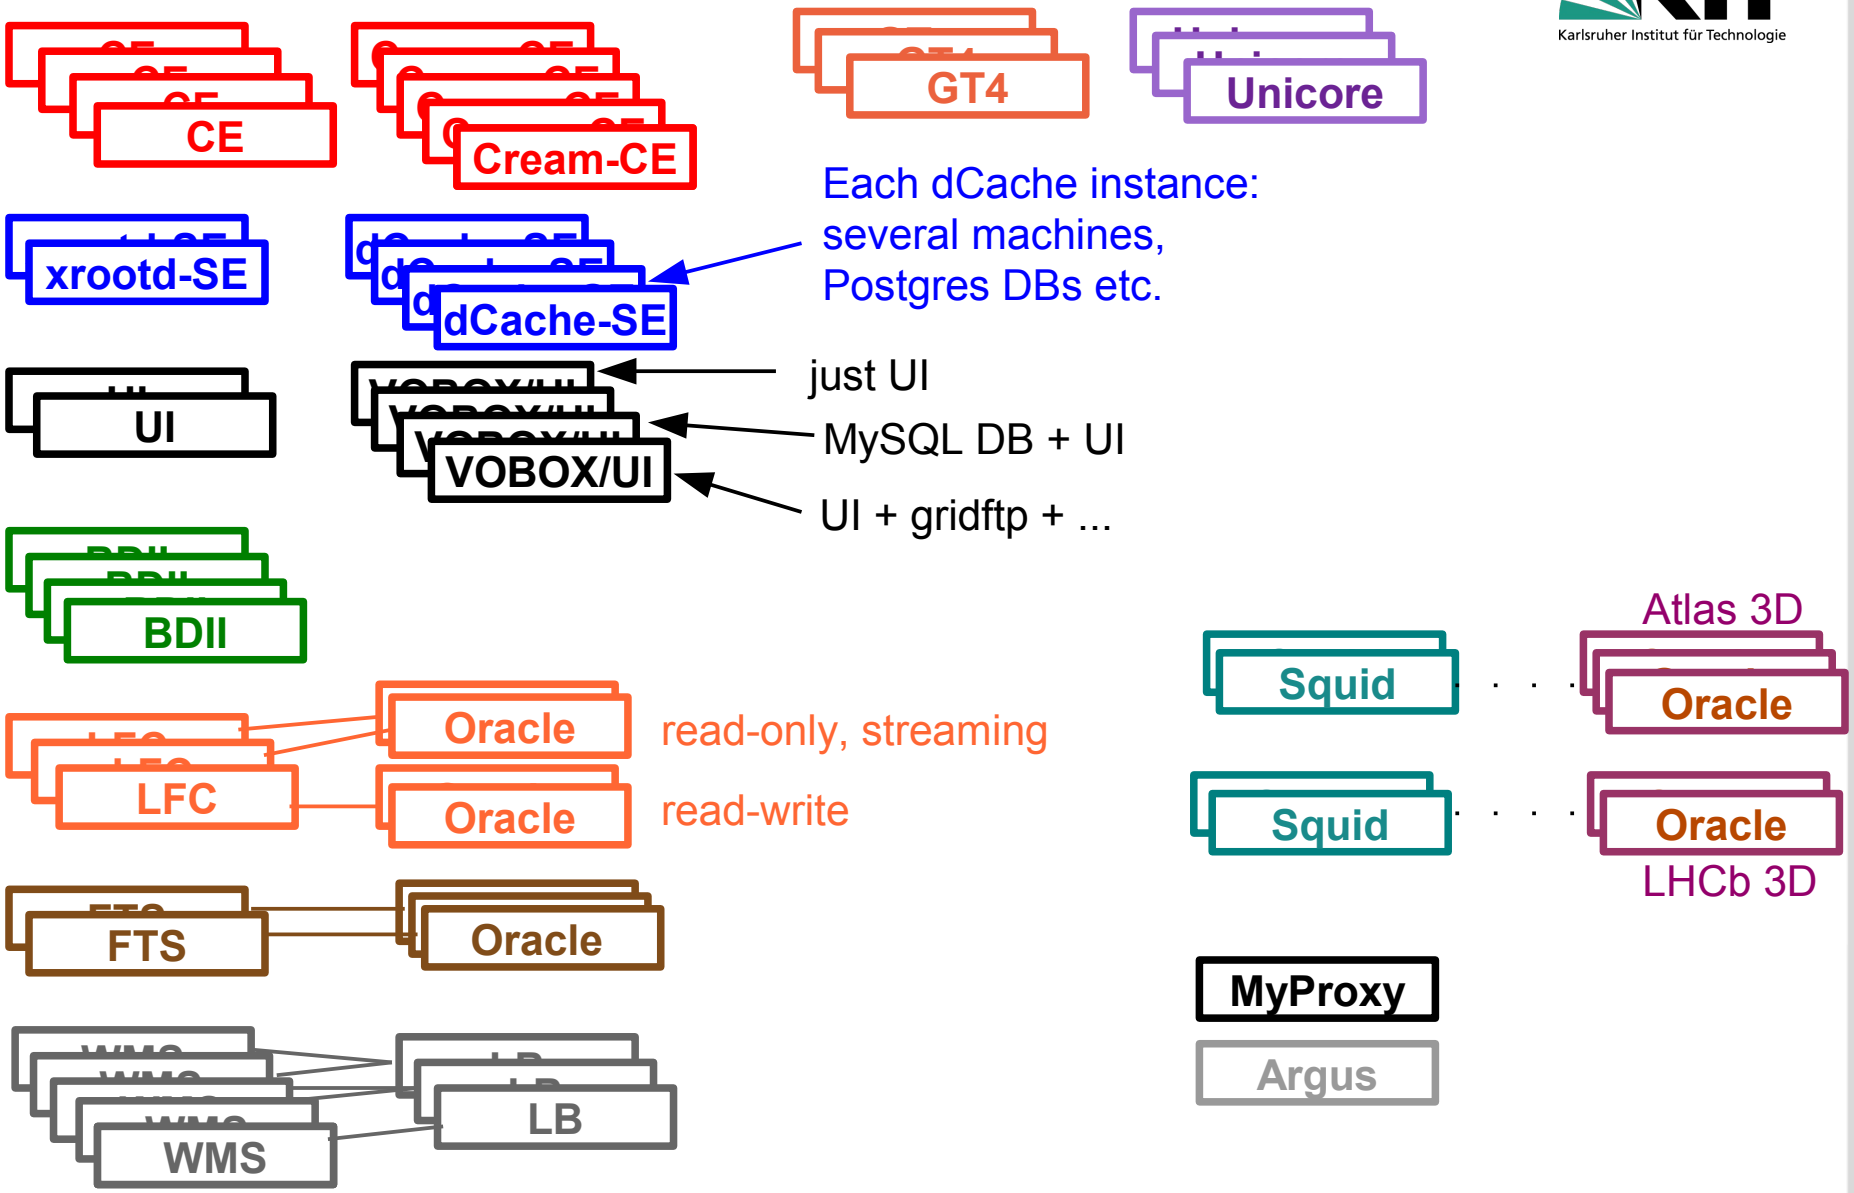

Storage systems

-  dCache is our workhorse for high-I/O capable storage
  - Easy setup for small user communities with low I/O requirements
  - Setup for 'power user communities' needs good planning and detailed knowledge about the computing model and access patterns, required storage classes etc.
    - collaboration with experiments' computing people essential
  - Very stable system (if setup is well planned)
  - New protocols (NFS4.1) at the horizon.
    - More versatile, suitable for other use cases

Current and medium-term challenges

- Changing computing models of the LHC communities
  - e.g. dropping of the Monarc model
  - A chance to move to more common services and tools!
    - WLCG TWGs!
- New and evolving technologies, e.g.
  - SSDs
    - no more development of magnetic disc technology!
    - will prices per TB continue to fall like in the last years?
    - SSDs enable new storage and DB applications
      - physicists will want it
  - Many-core CPUs
  - Virtualization and Cloud
    - common solution at large sites?

Requirements for the support of multiple disciplines

- Reliable and scalable basic Grid services
- Common computing models and services
  - As little as possible VO specific services
  - Use of Grid as a system of distributed redundant services (where possible), not distributed single point of failures
- Common support and collaboration tools
  - e.g. GGUS, GOCDB etc.

=> As many commonalities as possible for basic services and tools!

Outlook

- Support of new user communities is on the roadmap of SCC
  - make use of experiences made with LHC experiments
  - reuse well known and stable services (gLite, dCache) as building blocks
  - Close collaboration between site people and (computing) people from the user communities.
  - Collaboration between sites to exchange technology and experiences would be beneficial.

Thank you!

Questions?