

Grid @ KIT (www.kit.edu)

Frank Schmitz
Forschungszentrum Karlsruhe
Institut for Scientific Computing (IWR)
Hermann-von-Helmholtz-Platz 1
76344 Eggenstein-Leopoldshafen
Germany
www.CampusGrid.de



University of La Serena, Chile, 28.5.2007

Frank.Schmitz@iwr.fzk.de

Karlsruhe Institute of Technology (KIT)

University

11	fakulties
120	institutes
4,000	employees
18,500	students
250 mio €	budget

KIT...

Forschungszentrum science

10	programs (research area)
21	large institutes
4,000	employees
310	institutes
910	services
810	mio € budget

infrastructure

New research groups starting

The Universität Karlsruhe/Forschungszentrum Karlsruhe puts into practice the success in the Exzellece Initiative (German government, two universities in Munich and KIT):

In an internal competition, it has started roundabout 20 new research groups. They explore scientific fields relevant for the future, which have not really been examined so far.

At the same time, University and Research Centre work on ideas and structures for the Karlsruhe Institute of Technology (KIT), which are to be fixed in a contract at the end of the year. Thus, the Fridericiana strengthens her position in Germany, which Professor Dr. Jürgen Mittelstraß, theorist of science from Konstanz, has described in a speech: Not all Universities are equal - a few jut out, for example the Universität Karlsruhe.

GridKa, the Tier-1 computer center for LHC

LHC=large hadron collider (CERN)

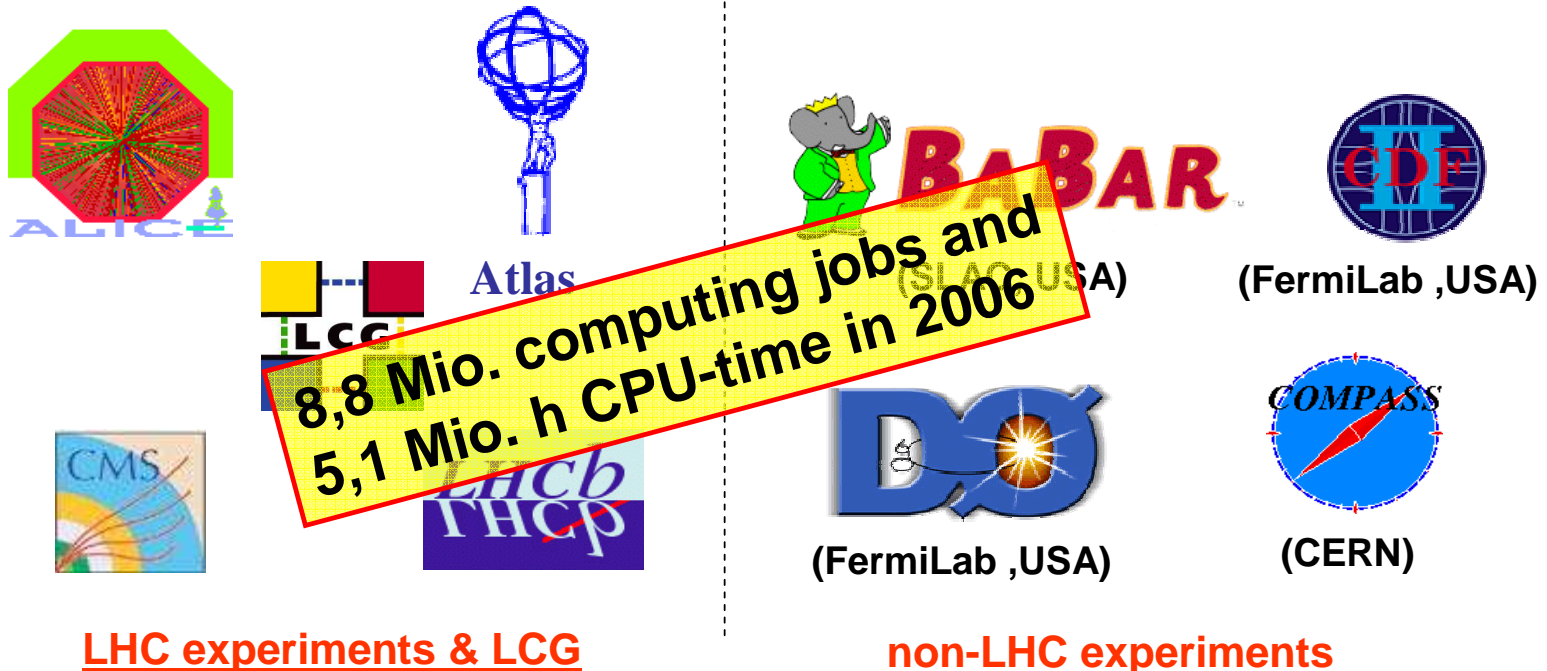
Grid based structure in a worldwide environment for high energy physics

- gLite
- storage resource manager (SRM) based on dCache
- Resource broker for world wide usage of LHC resources
- PBSpro for resource management
- Certification Authority (CA)


GridKa, the Tier-1 computer center for LHC

Deliver yearly
12 PB measured data

Have already
"real" data



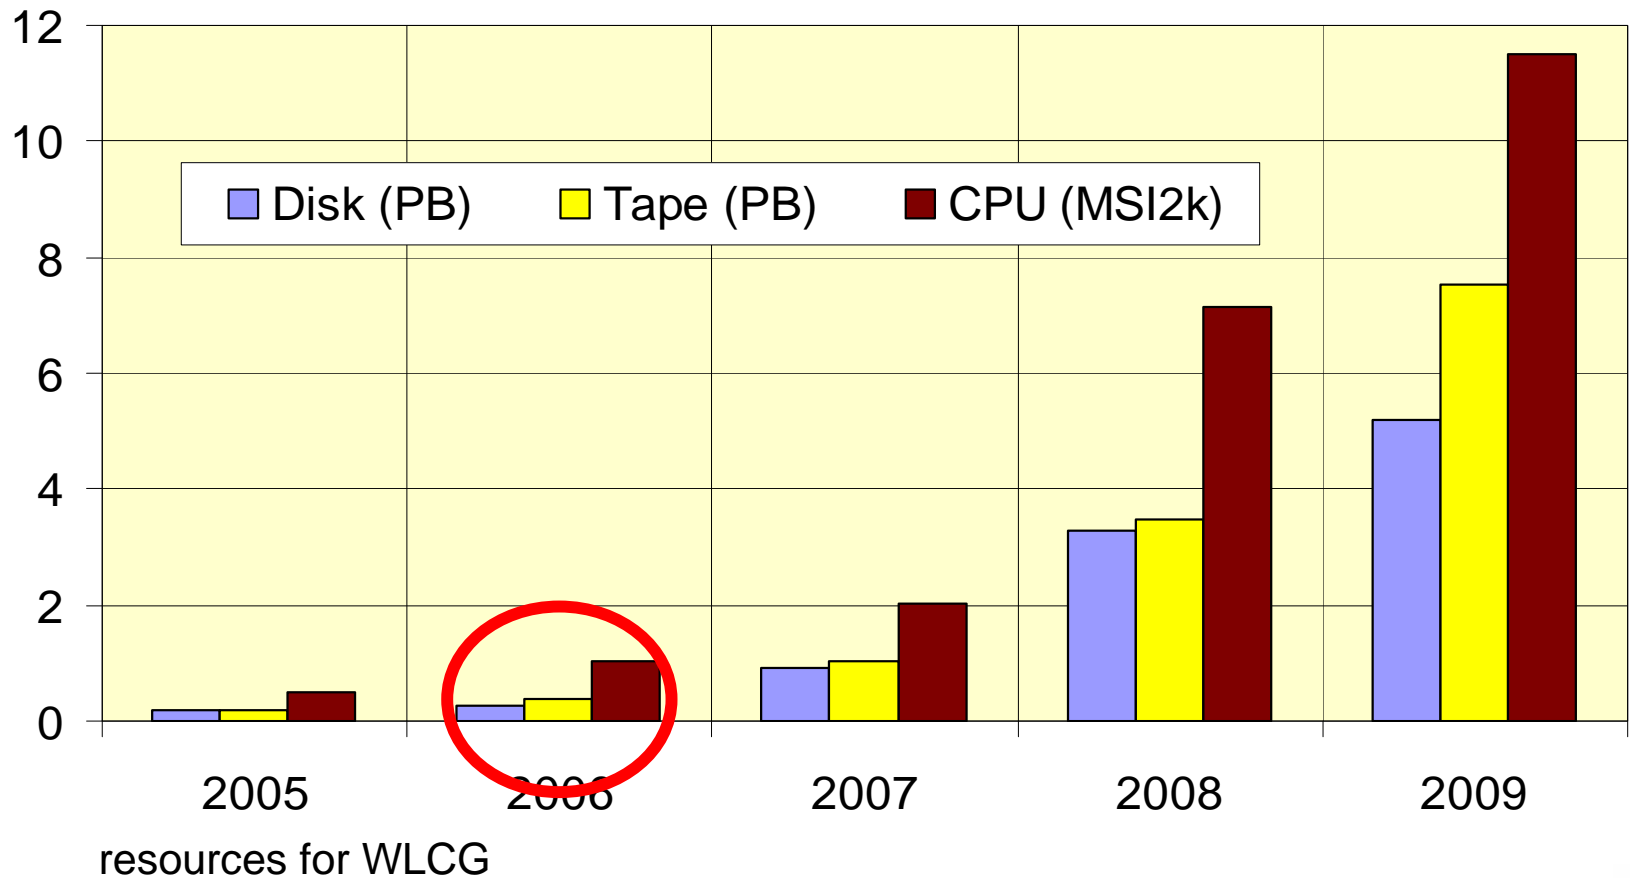
- Aachen (4) ●
- Bielefeld (2) ●
- Bochum (2) ●
- Bonn (3) ●
- Darmstadt (1) ▲
- Dortmund (1) ●
- Dresden (2) ●
- Erlangen (1) ●
- Frankfurt (1) ●
- Freiburg (2) ●
- Hamburg (1) ▲
- Heidelberg (1) ▲ (6) ●
- Karlsruhe (2) ●
- Mainz (3) ●
- Mannheim (1) ●
- München (1) ● (5) ▲
- Münster (1) ●
- Rostock (1) ●
- Siegen (1) ●
- Wuppertal (2) ●



22 Research Centers
44 Working Groups
350 Scientists

- University
- ▲ other research institutions
- () Number of working groups

FZK has signed the WLCG MoU and is able to deliver resources for the LHC experiment up to 2022 as the Tier-1 center



Motivation for the project CampusGrid

- heterogeneous IT-environment: vector-, SMP-, cluster-, blade-systems, SAN, NAS, Unix, Linux, Windows, Solaris, SuperUX,
- global view by the user
- only one user management (ADS from Microsoft)
- one job management
- metacomputing (MPI, ..), “real-time” applications
- access data for visualization at the local workstation
- global accounting
- seamless integration into different projects and middleware concepts like gLite, LCG, D-Grid, Unicore, GT4,

Ideas and solution

- testing of different global file system solutions à StorNextFS from Quantum/ADIC seems to be the best
- SRM for world wide access
- integration of InfiniBand, iSCSI and FC-SAN
- Globus Toolkit 4 (GT4), gLite/LCG and UNICORE as the middleware solution in the project
- Resource Broker (local solution!)
- security à Kerberos 5 integration
- accounting à should be solved by D-Grid and GridKa

Parts of the CampusGrid Project




watercooled Infiniband cluster with 32 SUN V20z and 64 FSC RX220 nodes (more than 190 cores and >1 TByte main memory)



2 Infinicon 9100, Infiniband switches, MPI latency 4.0 μ s between nodes

Parts of the CampusGrid Project

SX-8R vector computer



Usage of CampusGrid

- at the moment 260 jobs, 10 running in the vector environment, 25 running under AIX, 150 Linux jobs waiting for CPU resources
- MPI and OpenMP applications
- more than 50 processors Power4 of IBM using AIX as operating system
- 16 vector processors having 0.5 TByte of main memory
- some applications are using more than one week of CPU-time
- 40 TByte of disk space available
- write performance to disk reaches more than 1.3 GByte/s

Vector environment

- Two Nodes SX-8R
 - each with 8 CPUs
 - 35.2 GFlop/s per CPU
 - each with 256 GByte Memory
- Storage
 - 25 TB FibreChannel Storage with gStorageFS
 - 4x NEC S2500 with a total of 120 disks (300 GB) plus some Hotspares
 - RAID 50
 - redundant FC SAN
 - 2x 1.1 TB local FibreChannel storage direct attached to SX-8R

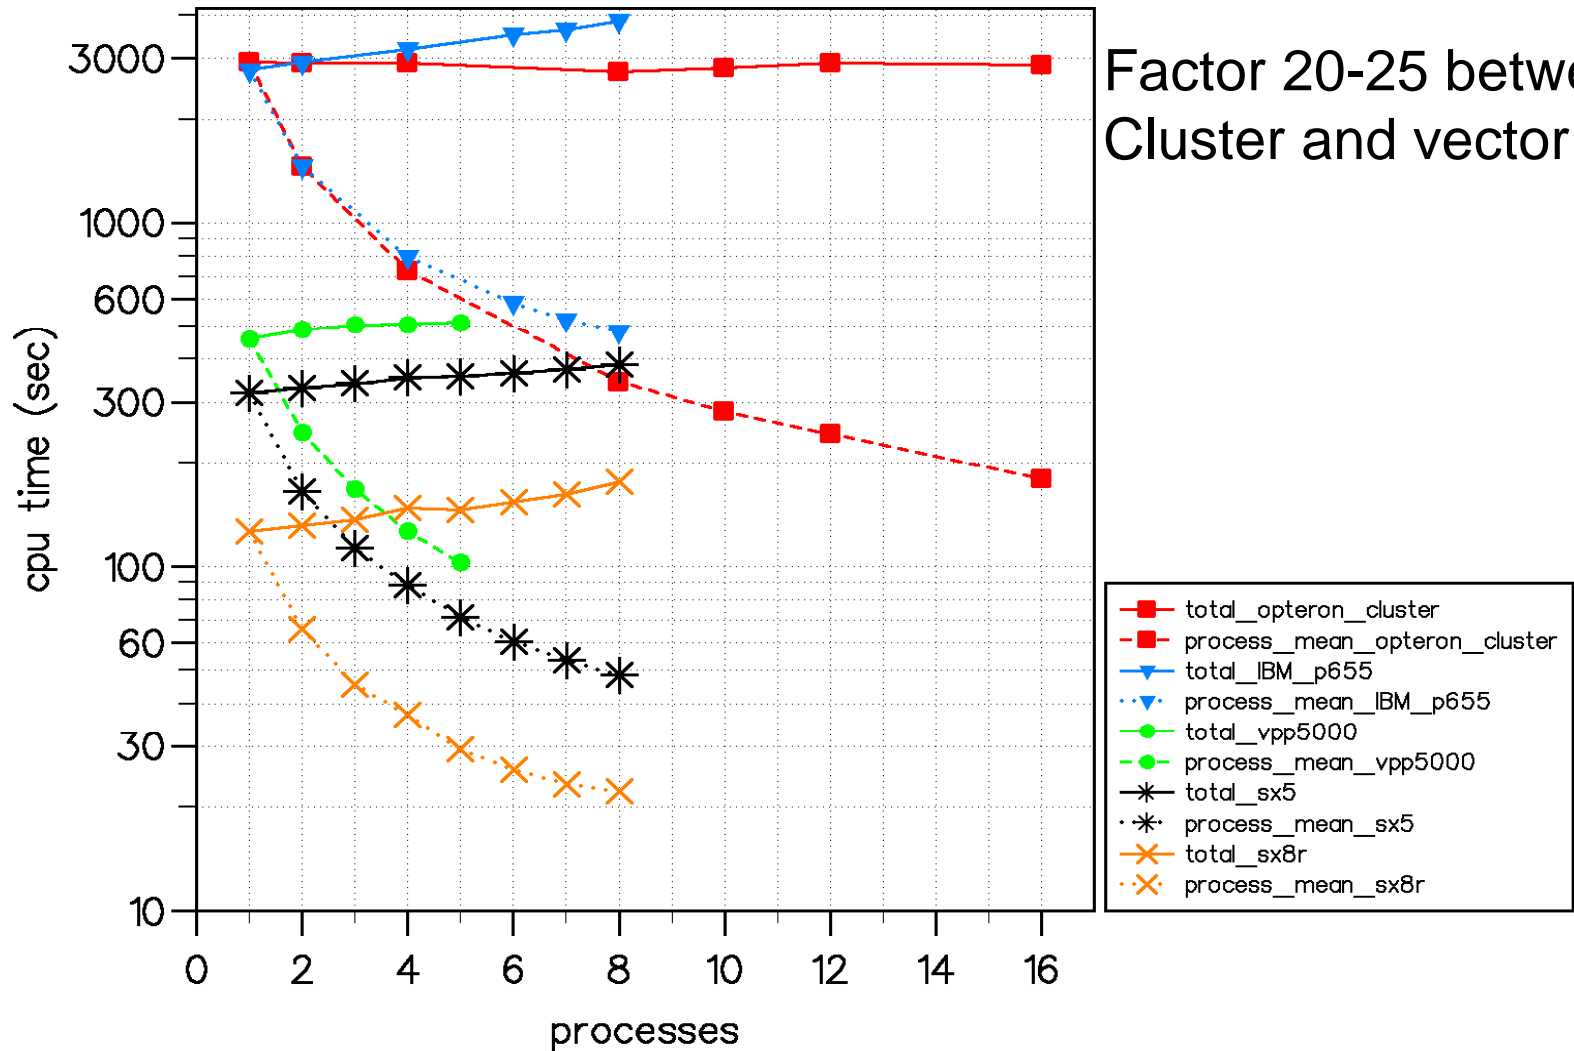
Performance: I/O

1. SX-8R:

- „big“ GFS with 8 LUNs, each 2.1 TB
writing 5 TB with 5 Tasks (F90 program):
mean throughput = 1.42 GB/s
- writing from both nodes doesnt change throughput
- SFS on the same 8 LUNs
maximum write throughput with 8 tasks = 1.77 GB/s
maximum read throughput = 2.2 GB/s

2. IA32 frontend system:

- 350 MByte/s for large files



D-Grid (the German Grid initiative)

- starting September 2005 to build a sustainable Grid infrastructure
- six Community Grid projects and the D-Grid Integration Project (DGI)
- funded by BMBF, the Federal Ministry of Education and Research
- astronomy, climate research, high energy physics, engineering research, medical research, humanities
- gLite, GT4, UNICORE
- SRM/dCache for the data access
- heterogeneous CPU access (AIX, Linux, Solaris, SuperUX planned)

Other Grid-projects

- EGEE (enabling Grids for e-science)
- GGUS (global Grid user support)
- g-Eclipse
- Int.EU.grid (interactive European Grid)
- ISSeG (integrated site security for Grids)

Thank you!

Questions?