

Site Report

Manfred Alef, Jos van Wezel

Grid Computing Centre Karlsruhe (GridKa)

Forschungszentrum Karlsruhe
Institut für Wissenschaftliches Rechnen
Hermann-von-Helmholtz-Platz 1
D-76344 Eggenstein-Leopoldshafen

<http://www.fzk.de>, <http://www.gridka.de>

firstname.lastname@iwr.fzk.de

GridKa (Grid Computing Centre Karlsruhe):

- **German LCG tier-1 centre:**
 - ◆ Alice, Atlas, CMS, LHCb

- **German regional grid computing centre:**
 - ◆ Non-LHC HEP experiments:
BaBar, CDF, Compass, D0
 - ◆ D-Grid

GridKa (Grid Computing Centre Karlsruhe):

→ Hardware currently available:

- ◆ 980 worker nodes = 1960 CPUs = 2650 cores (SL4 i386)
- ◆ 1.9 TB disk (GPFS, dCache)
- ◆ 1.4 TB tape
- ◆ 140 water-cooled 19" cabinets
- ◆ 4x 10 Gbit/s links
(internet + 3x OPN: CERN, CNAF, IN2P3)

Storage:

→ **Tape systems:**

- ◆ Continue to use LTO on separate 'tape SAN'
- ◆ Currently 24 drives in 2 libraries (IBM, GRAU)
- ◆ Procurement started for LTO4 (+12 drives)

Storage:

→ Disk storage:

- ◆ Disk outage because of:
 - Off-spec chip in controller:
fixed with pull-up resistor on Cu FC links.
 - No FC path failover under certain conditions:
firmware bug, waiting to be fixed.

Storage:

→ Disk storage:

- ◆ First tests commenced with 10GE file servers:
 - Next generation storage at GridKa
 - Chelcio PCI-E and PCI-X cards
 - Quad core CPU

Worker Nodes:

- **+ 168 worker nodes (since April):**
 - ◆ Intel S5000VCL
 - ◆ 2x Intel Xeon 5160 (3 GHz, Woodcrest)
 - ◆ 6 GB RAM
 - ◆ 200 GB S-ATA
 - ◆ ...

Worker Nodes:

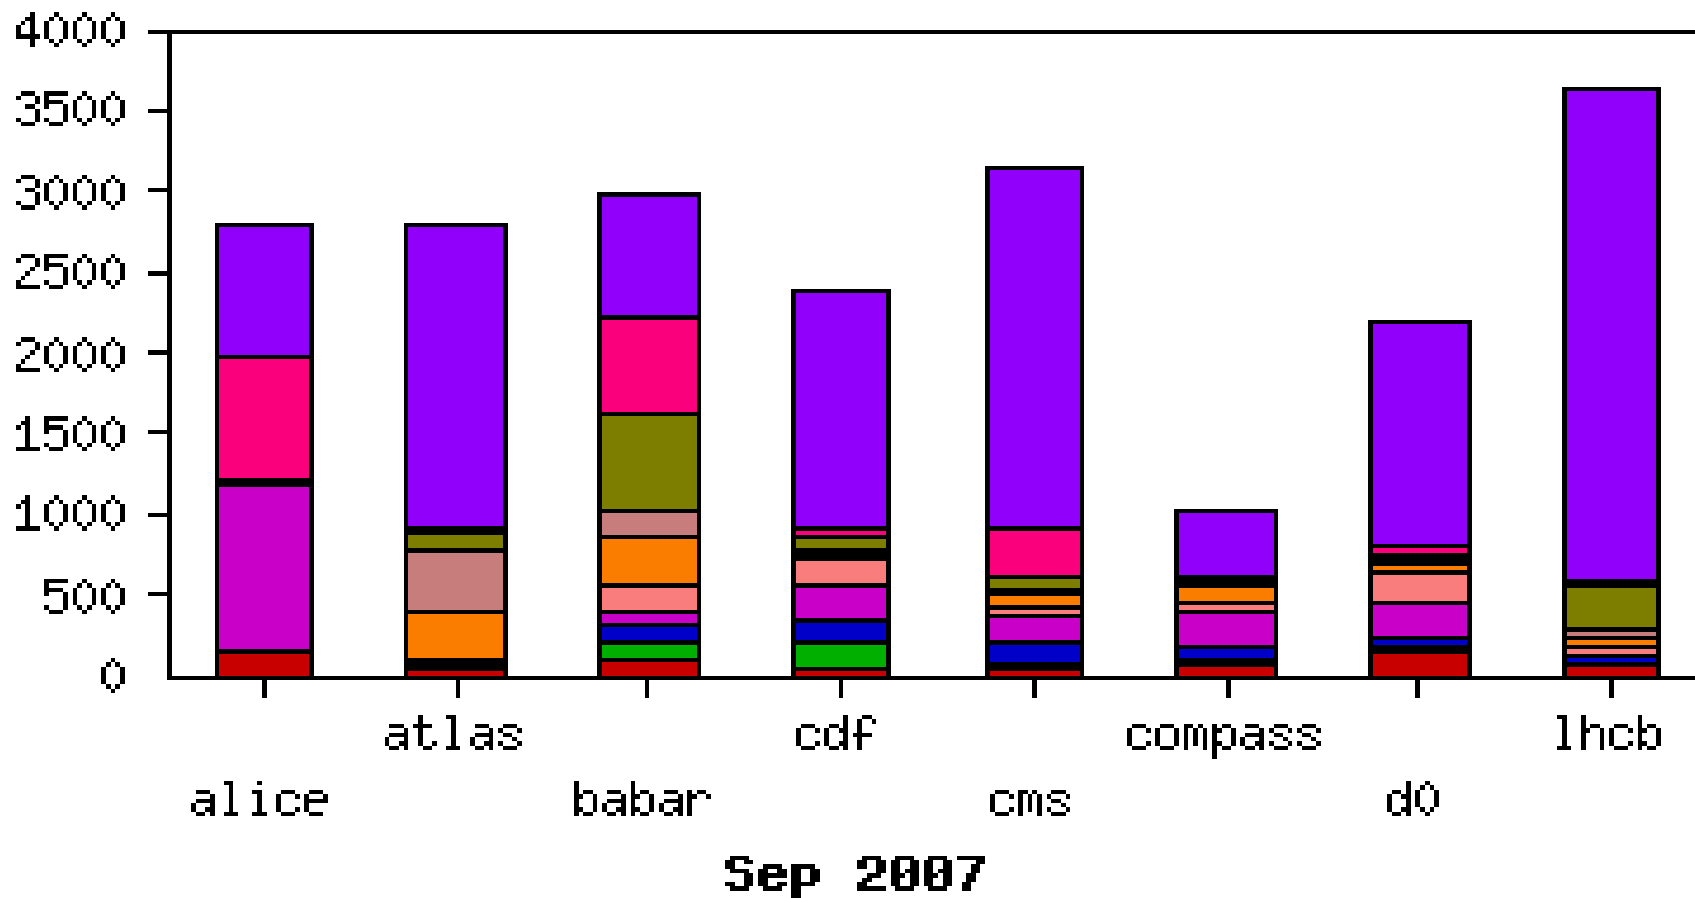
- **+ 320 worker nodes (delivery + installation in progress):**
 - ◆ TYAN S5380G2NR
 - ◆ 2x Intel Xeon E5345 (2.33 GHz, Clovertown)
 - ◆ 16 GB RAM
 - ◆ 2x 250 GB S-ATA (Seagate Barracuda ES, 16 MB cache)

- **Next procurement started ("D-Grid"):**
 - ◆ Quad-core CPUs (Intel Xeon, or AMD Barcelona)


Worker Nodes:

- **500 GB hard disk space – why?**
 - ◆ Local homedirectories for poolaccounts:
Atlas: 15-20 GB local disk space per job?
 - ◆ Virtual memory usage ...

Virtual Memory Usage of Batch Jobs at GridKa



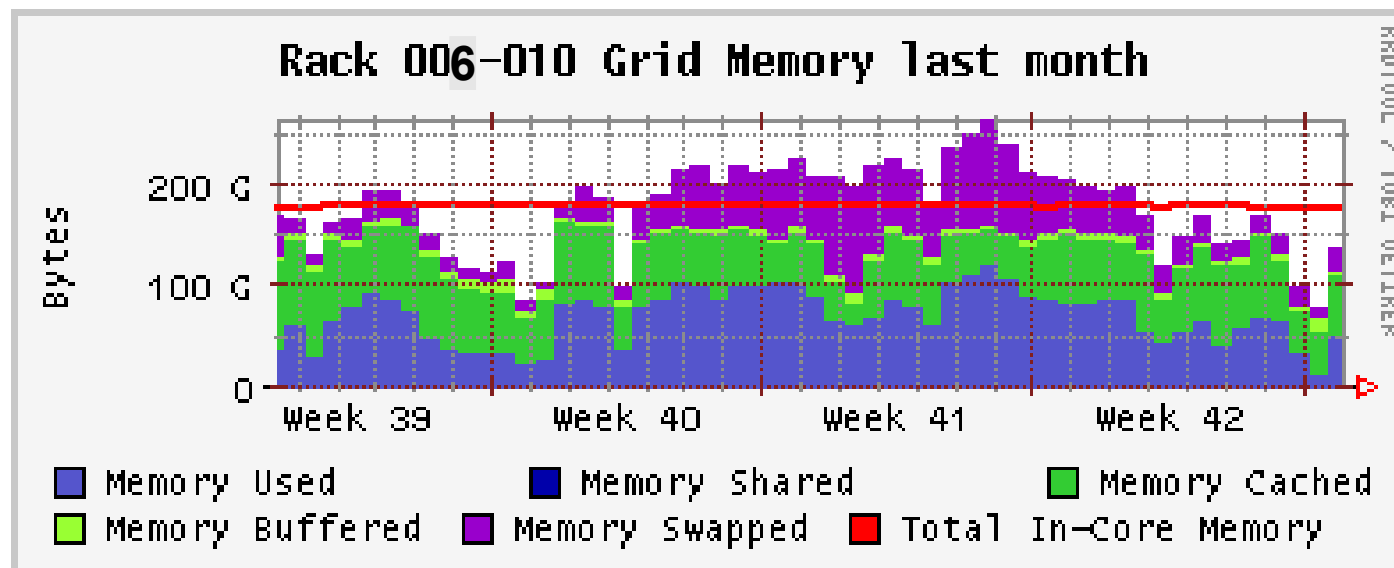
Compiled from
PBS logfiles



Memory Utilization of Worker Nodes at GridKa


Oldest WNs:

- 2 CPU cores
- 1 GB RAM



Latest WNs:


- 4 CPU cores
- 4-6 GB RAM




Ganglia plots

Memory Utilization of Worker Nodes at GridKa

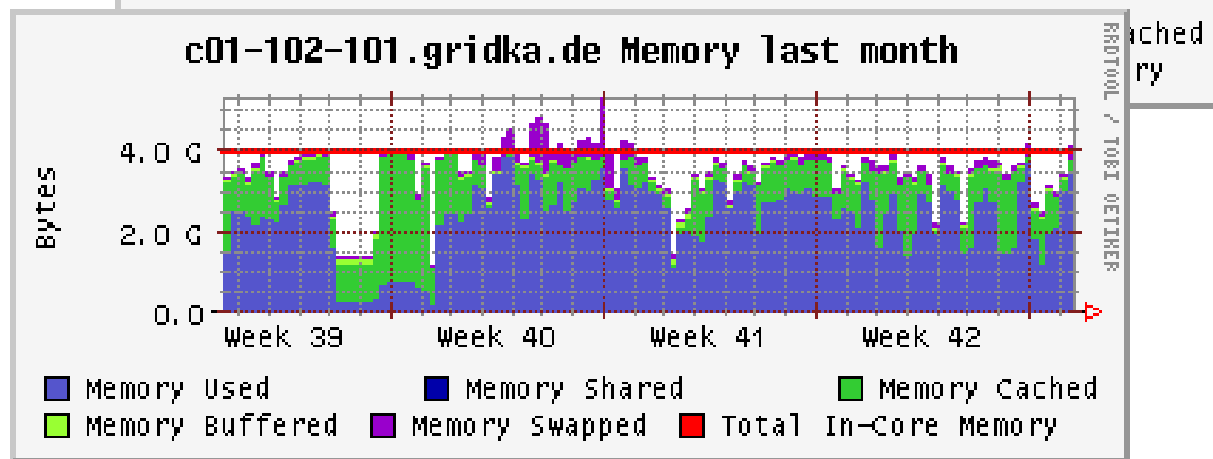
2 CPU cores, 1 GB RAM



2 CPU cores, 2 GB RAM



4 CPU cores, 4 GB RAM



Ganglia plots

Worker Nodes:

- **Virtual memory usage, swapping:**
 - ◆ Currently 0.5 ... 1.5 GB RAM + 2 GB swap (per core)
 - ◆ New generation of WNs: 2 GB RAM + 4 GB swap per core
 - ◆ Future requirements?
 - ◆ x86_64 issues?
 - ◆ More than 1 job slot per core?

Comments, Questions?