

# ***CMS Status Report***




## PCP for CMS DC04 at FZK

**Beware of acronyms!**

The German CMS Production Team:


Klaus Rabbertz, Jens Rehn, Serge Sushkov, Joanna Weng



# Data Challenge 04




- ★ Start: January 2004
- ★ **5% test** of expected computing requirements for CMS running in 2007
  - 1) Data distribution and access
  - 2) Data calibration
  - 3) **Distributed user analysis jobs**
    - ➔ Jobs on data on multiple Tiers



# *Pre-Challenge Production*




- ★ Simulation of events necessary for DC04
- ★ Anticipated: **50 M** events until December
- ★ Requested from physics working groups: **69 M** events
- ★ German participation: **10 %**
  - ➔ 5 M events for GridKa
- ★ Current total rate: **12 M / month**
- ★ **But:** No pile-up yet, only „old“ Geant3 simulation



# Assignment Scheme




- ★ RefDB: Central database at CERN
  - ➔ Distribution and tracking of MC production (via e-mail)
- ★ CMKIN, CMSIM, OSCAR:
  - ➔ CMS simulation software (n-tuples, Geant3, Geant4)
- ★ McRunJob:
  - ➔ Local tracking and job submission (Python scripts)
- ★ SRB: Storage Resource Broker:
  - ➔ Data and software distribution



# *Last three months ...*



- ★ 23.06. GridKa validated for CMKIN production
- ★ 01.07. Start of Pre-Challenge Production
- ★ 02.07. First CMKIN assignment
- ★ 18.07. GridKa validated for CMSIM production
- ★ 18.07. First CMSIM assignment
- ★ Today: 26th assignment running
  - ➔ 2.6 M events generated (CMKIN), 3.8 % over all Tiers
  - ➔ 3.2 M events simulated (CMSIM), 9.6 % over all Tiers



# Problems with ... filesystems



## ★ CMS: filesystems cms1 – cms3

1) cms1: 1.8 TB, used for start of production, **ok**


2) cms2: 1.5 TB, **very problematic**, last incident:  
1. October

→ used for checks only

3) cms3: 6.5 TB, available since 13.07., after small  
start-up problems runs **very stable**

→ Migrated production to cms3 on 28.08.

## ★ Home filesystem sometimes not working



# Problems with ... mails



- ★ CMS uses mails from compute nodes for production surveillance
  - 1) Sending mails from within NAT ok, but
    - ➔ Rejected as **SPAM** at CERN
  - 2) Set-up change at GridKa requested: June 2003
    - ➔ Completion announced: 08.09.2003
    - ➔ Really finished: 26.09.2003




# *Problems with ... firewall*



- ★ CMS employs SRB for file import/export
  - 1) Need one primary port: 5543
  - 2) Need range for multithreading: 20000 - 20200
    - ➔ Blocked by firewall despite our request because ...
    - ➔ **CMS installation recipe was wrong:**
      - Primary port 5544
      - Range 20000 - 20200 is default ... **it's not!**







# Problems with ... /tmp



- ★ Maximal Walltime: 96 h
- ★ /tmp clean-up after: 240 h
- ★ Left-over files:
  - 1) Fill up /tmp (e.g. CDF: 15 GB from crashed job)
    - ➔ Such a „free“ node may empty the queues
  - 2) Confuse jobs (e.g. CMS: diff. jobs read from same directory /tmp/job.pid by accident)
    - ➔ Wrong output copied back
    - ➔ CMS clean-up will be improved



# Problems with ... I/O



- ★ Very bad Walltime to CPUtime ratio!
- ★ Currently only jobs w/o pile-up for CMS
- ★ I/O will be a major issue for user analyses

The screenshot shows the xpbs1.1.12 interface with three main sections: HOSTS, QUEUES, and JOBS.

**HOSTS Table:**

Server	Max	Tot	Que	Run	Hld	Mat	Trn	Ext	Status	PEsInUse
pbs	440	828	392	436	0	0	0	0	Active	-/-

**QUEUES Table (Listed By Host(s): pbs):**

Queue	Max	Tot	Ena	Str	Que	Run	Hld	Mat	Trn	Ext	Type	Server
test_2_60	0	0	yes	yes	0	0	0	0	0	0	Execution	pbs
default	300	0	yes	yes	0	0	0	0	0	0	Execution	pbs
short	2	0	yes	yes	0	0	0	0	0	0	Execution	pbs
long	139	147	yes	yes	6	141	0	0	0	0	Execution	pbs
test	16	0	yes	yes	0	0	0	0	0	0	Execution	pbs
extralong	229	681	yes	yes	386	295	0	0	0	0	Execution	pbs

**JOBS Table (Listed By Queue(s): extralong@pbs):**

Job id	Name	User	PEs	CputUse	WalltUse	S	Queue
325631.pbs	DsSel-1	marc	1	01:14:19	53:37:17	R	extralong@pbs pbs 325631.pbs
325632.pbs	DsSel-2	marc	1	01:27:40	53:17:17	R	extralong@pbs pbs 325632.pbs
325633.pbs	DsSel-3	marc	1	00:55:58	50:10:04	R	extralong@pbs pbs 325633.pbs
325634.pbs	DsSel-4	marc	1	01:22:52	49:51:33	R	extralong@pbs pbs 325634.pbs
325635.pbs	DsSel-5	marc	1	01:30:00	49:25:15	R	extralong@pbs pbs 325635.pbs
325636.pbs	DsSel-6	marc	1	01:06:10	49:25:14	R	extralong@pbs pbs 325636.pbs
325640.pbs	DsSel-10	marc	1	01:28:09	47:27:27	R	extralong@pbs pbs 325640.pbs
325644.pbs	DsSel-14	marc	1	01:15:41	47:17:35	R	extralong@pbs pbs 325644.pbs
325645.pbs	DsSel-15	marc	1	01:07:29	47:14:22	R	extralong@pbs pbs 325645.pbs
325646.pbs	DsSel-16	marc	1	00:44:10	47:13:31	R	extralong@pbs pbs 325646.pbs
325647.pbs	DsSel-17	marc	1	00:41:27	46:38:17	R	extralong@pbs pbs 325647.pbs

A red circle highlights the 'CputUse' and 'WalltUse' columns in the JOBS table. A handwritten note '~1h/50h' is written next to the circle.




# Problems with ... I/O




- ★ Too many I/O intensive jobs block each other
  - ➔ occupy slots in queues,
  - ➔ spoil GridKa performance and
  - ➔ prevent the other experiments to run their jobs

Impact of I/O  
int. jobs from  
BaBar




Similar for D0  
this Tuesday



# *Other Problems*




- ★ CMKIN jobs finish too fast (~5 sec.)
  - ➔ NFS overload => increased submission delay (10 s)
- ★ Limits for max. running jobs per queue, determined automatically?
  - ➔ Sometimes prevents jobs from running
- ★ sshd blocks login, restart necessary
- ★ Grid User Support, portal there, but is it in use?
  - ➔ Got automatic acknowledge for request, nothing more yet




# Good News






- ★ Very nice Ganglia interface to monitor GridKa
- ★ It is possible to run  $> 390$  CMSIM jobs in par.







# Example of impact of CMS production on GridKa




We (CMS) could run **more** assignments

Impact of CMS assignment



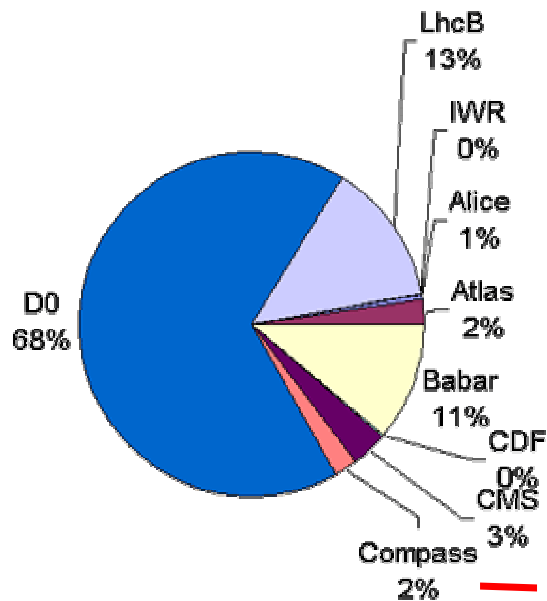


# CPU Usage Jul -> Sep




## July

CPU-Time  
insgesamt:  
163 248,31 Stunden



## August

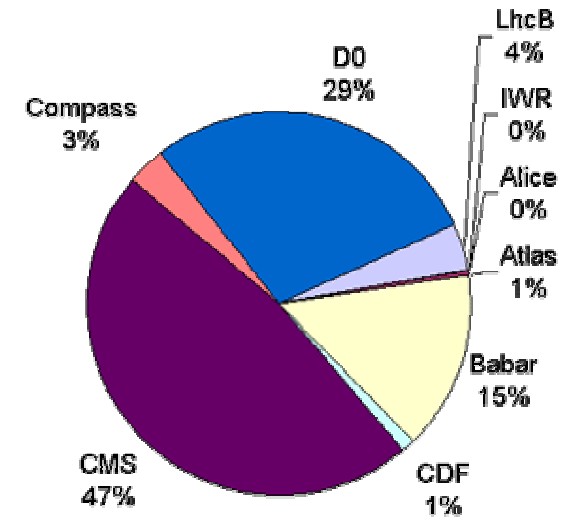
CPU-Time  
insgesamte CPU-Zeit:  
124 951,23 Stunden




Color coding changed!?

## September

CPU-Time  
insgesamte CPU-Zeit:  
172 431,53 Stunden






# *Documentation Requests*



- ★ Web archive for rdccg-user mailing list
  - ➔ New web pages announce GridKa-user mailing list ... which doesn't exist, yet?
- ★ List of user accounts and corresponding e-mail addresses
  - ➔ No cat /etc/passwd plus Google search necessary
- ★ Documentation of fileserver sharing between the experiments
- ★ Documentation of fairshare settings/points






# *Other Requests*



- ★ More efficient handling of jobs with extremely bad Walltimes/CPUtime ratios necessary
- ★ Better timing of announcements for non-emergency interventions
  - ➔ Not e.g. for tapes: Announced: 23.09. 23:23  
Scheduled: 24.09.
- ★ One hour autologoff of ssh sessions is quite annoying from outside Karlsruhe



# Summary



- ★ Encountered many problems and resolved quite some fraction of them.
- ▶ It's a stony path to go.
- ▶ Nevertheless we're taking up speed.

Many thanks to the GridKa people,  
special thanks go to Manfred Alef.