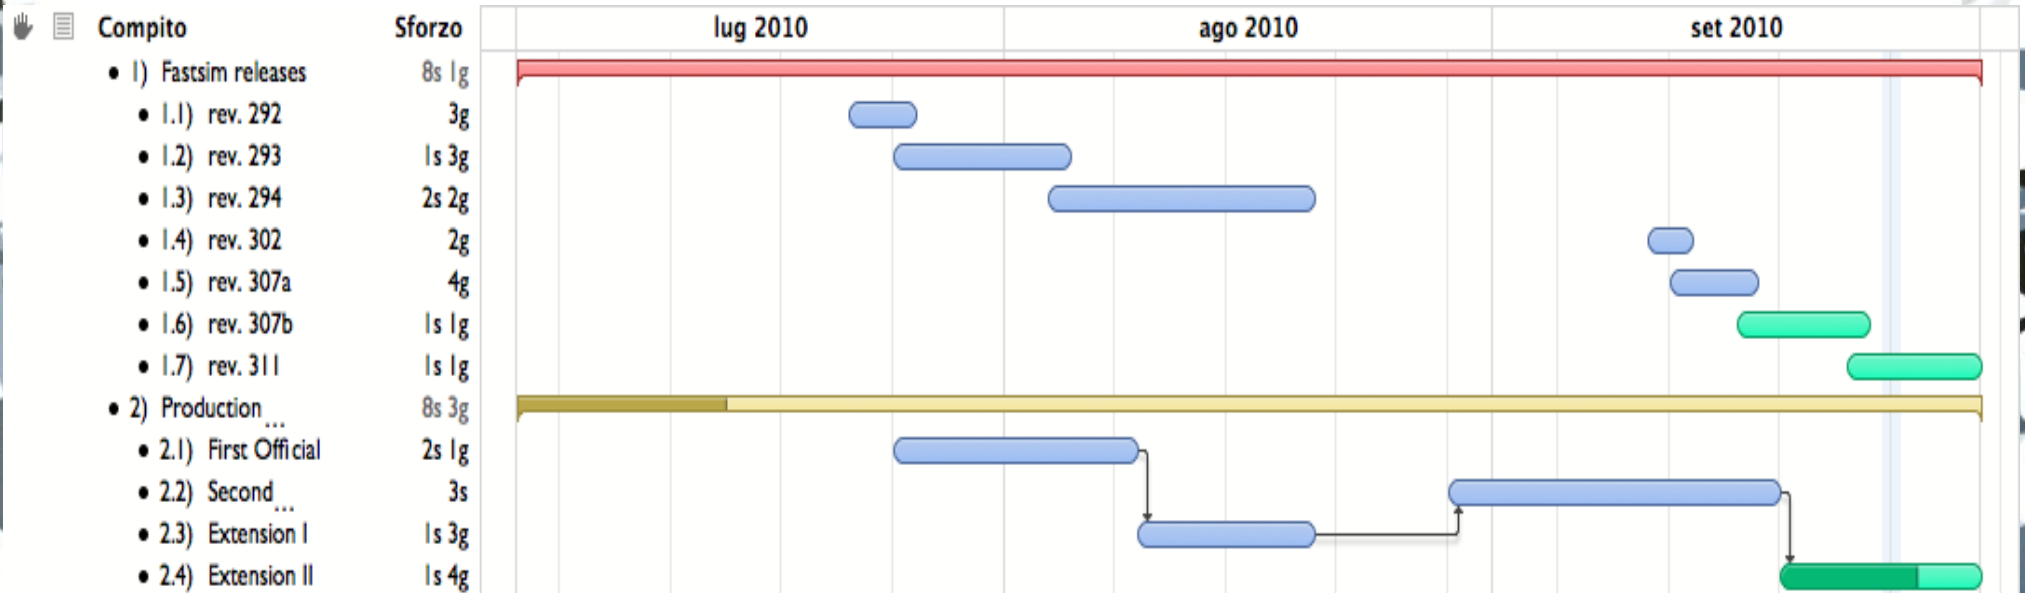


Production status

Armando Fella
for the Distributed Computing Group

Production info

- The time plan
 - first production period: July 26th - August 9th
 - extension to August 20th
 - second production period: August 30th - September 19th
 - extension to October 1st
- Six different sw revisions have been deployed during production period



Production data

- **Production requests**

- DG_Study, mixing bkg nopair: 3.5×10^9
- DG_Study, mixing all bkg: 5.7×10^8
- DG_Study/Physics, mixing all bkg: 5.4×10^8
- **Total:** **4.6×10^9**
- *Physics, mixing all bkg:* 6.2×10^{11} (Postponed)

- **“Preferred data” produced (rev 307 + rev 311, 2010_September)**

- DG_Study, mixing bkg nopair: 95%
- DG_Study, mixing all bkg: 50%
- DG_Study/Physics, mixing all bkg: 100%

Distributed production status

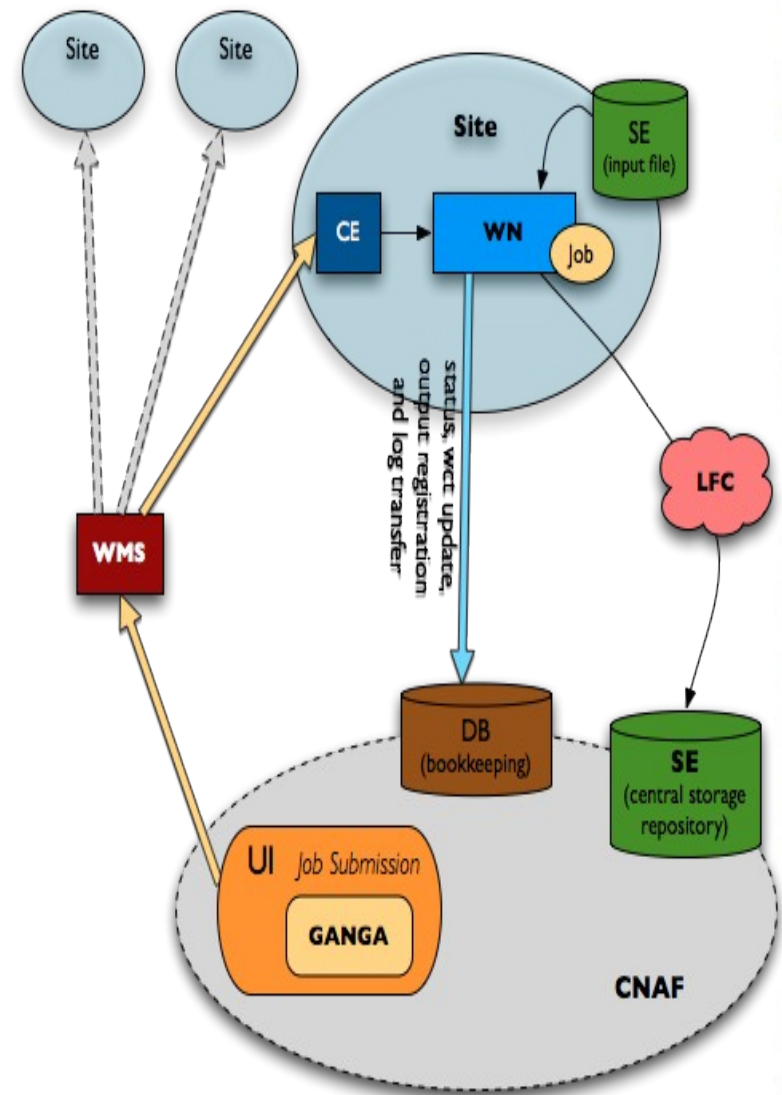
- Amount of produced data:
 - Job done: ~ 135K, #Evt: 8.6×10^9
 - ~7 times jobs and ~5.4 times evts wrt February prod.
 - Job failed: ~13K, Failure rate: 10% (mostly due to grid infr. problems)
 - ~2 times failure rate wrt February prod
- Production series vs sw revision vs evt:
 - 2010_July: 292, 293, 294, 302, 307a: 5.4×10^9 evt
 - 2010_September: 307b, 311 (preferred data): 3.1×10^9 evt

Production Tools status

- Several new features and improvements
 - New Shifter portal
 - New Elog system
 - New Physics request portal
 - Improved job script code
 - Improved web monitor
 - Improved Grid submission procedure
- see next presentation by Luca Tomassetti

Distributed infrastructure evolution

- **Improvements wrt Feb. prod:**
- Grid job submission optimized, quicker bulk submission via Ganga
- We are using n.g. CE (CREAM)
- Job data handling fail over procedures improved (lcg-cr > lcg-cp > globus-url-copy)
- The job output replica on hosting site has been inhibited
- Setup sites procedures optimized: reduced the transfers time (bkg and test_release)
- SBK DB upgraded and optimized to be able to host new features backend



“On Shift” step by step

- Elba WS shifter session, recruitment
- The shifter guide has been produced
- “Hands on” training meeting
- Offline training week
- **Feedback from and collaboration with shift takers have been of great worth for prod tools tuning and procedures improvement**
- **Thanks to all shift takers for patience in managing with roller coaster time plan**

The shift takers

1ST WEEK - FROM 26 JULY TO 01 AUGUST

First shift [from 08:00 to 16:00 (CET)] : WebUI Managers

Second shift [from 16:00 to 24:00 (CET)] : WebUI Managers

Third shift [from 24:00 to 08:00 (CET)] : WebUI Managers

2ND WEEK - FROM 02 AUGUST TO 08 AUGUST

First shift [from 08:00 to 16:00 (CET)] : Alejandro Perez

Second shift [from 16:00 to 24:00 (CET)] : Dana Lindemann

Third shift [from 24:00 to 08:00 (CET)] : Bertrand Echenard

3RD WEEK - FROM 30 AUGUST TO 05 SEPTEMBER

First shift [from 08:00 to 16:00 (CET)] : Alberto Cervelli

Second shift [from 16:00 to 24:00 (CET)] : Leonid Bumistrov

Third shift [from 24:00 to 08:00 (CET)] : Michael Lazos

4TH WEEK - FROM 06 SEPTEMBER TO 12 SEPTEMBER

First shift [from 08:00 to 16:00 (CET)] : Alessandro Rossi

Second shift [from 16:00 to 24:00 (CET)] : Alex Martin

Third shift [from 24:00 to 08:00 (CET)] : Adam Simpson

5TH WEEK - FROM 13 SEPTEMBER TO 19 SEPTEMBER

First shift [from 08:00 to 16:00 (CET)] : Nicola Gagliardi

Second shift [from 16:00 to 24:00 (CET)] : Enrico Feltresi

Third shift [from 24:00 to 08:00 (CET)] : Darren Swersky

Thanks again!

Distributed prod:

Tier-1 sites:

INFN-T1 - CNAF Bologna, Italy
IN2P3-CC - Lyon, France
RAL-LCG2 - Oxford, UK

Tier-2 sites:

UKI-LT2-QMUL - London, UK
UKI-SOUTHGRID-RALPP - London, UK

GRIF - Orsay, Paris, France


SLAC - Stanford, CA, USA
CIT_CMS_T2B - Caltech, Los Angeles, CA, USA
VICTORIA-LCG2 - UVIC, Victoria, CA

INFN-PISA - Pisa, Italy
INFN-LNL-2 - Legnaro, Italy
INFN-BARI - Bari, Italy
INFN-NAPOLI-ATLAS - Napoli, Italy
INFN-CAGLIARI - Cagliari, Italy
INFN-TORINO - Torino Italy

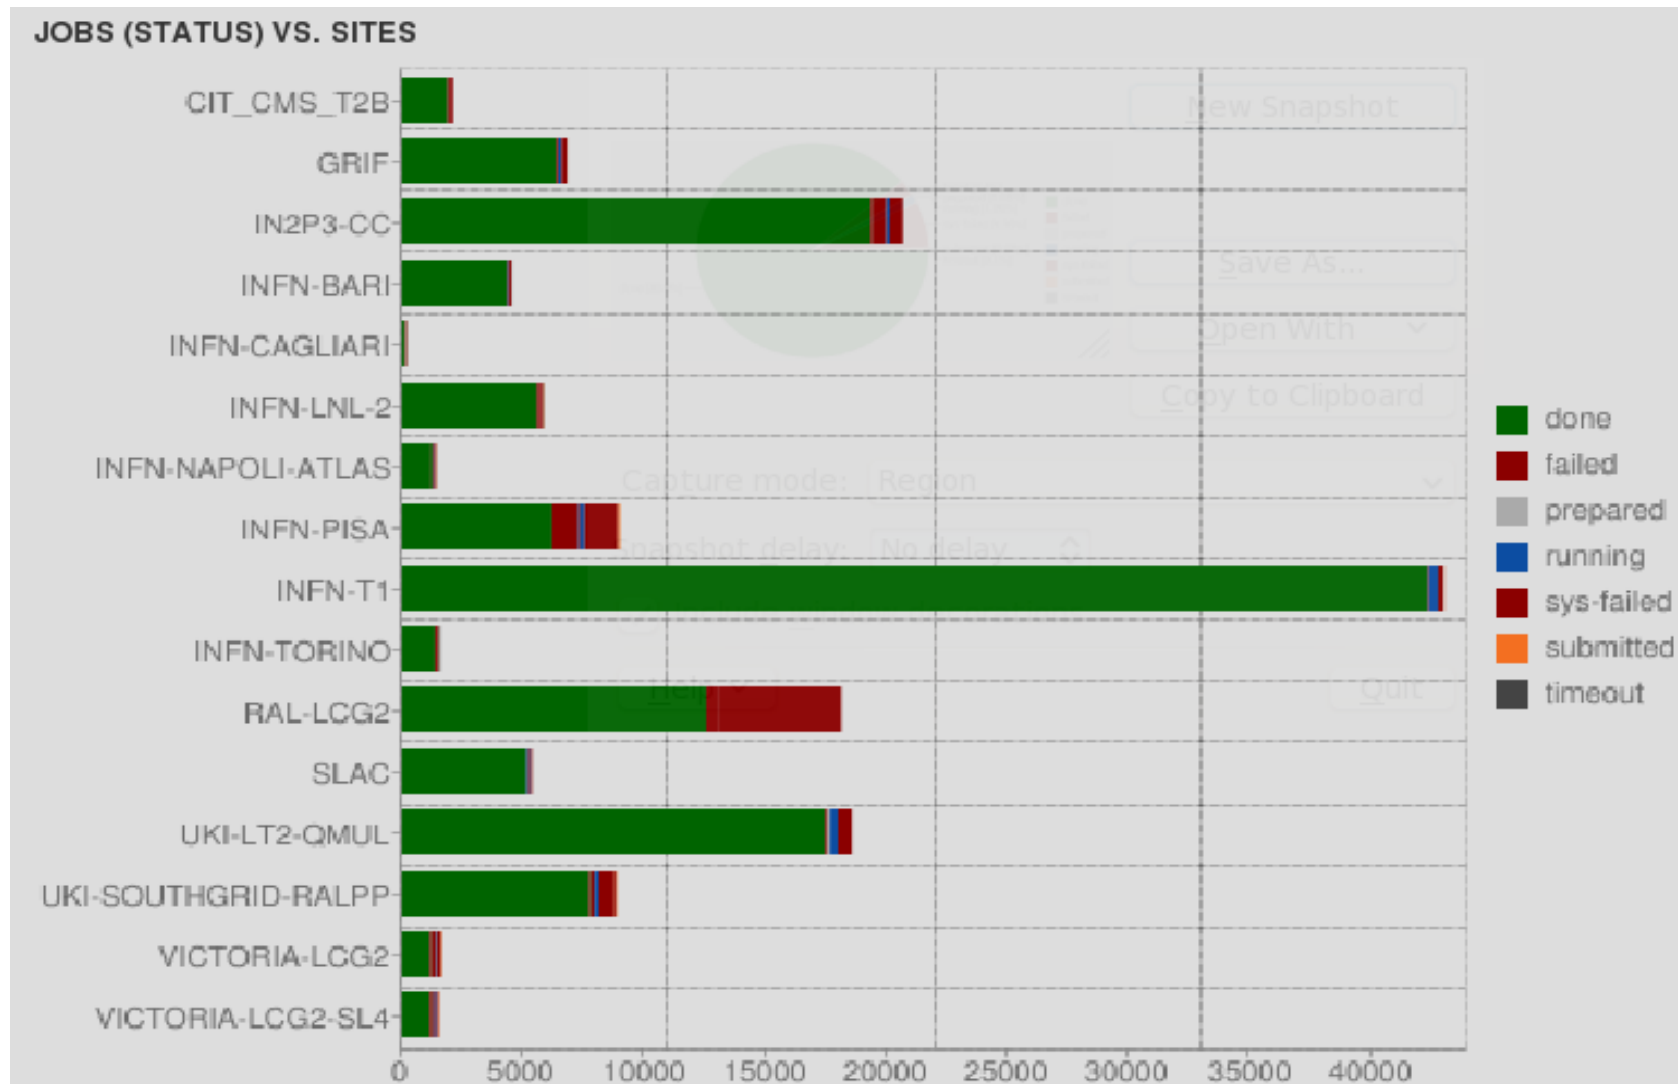
- Site usage:
- **15 sites on 3 Grid flavours:**
 - **EGI, OSG, WestGrid**

ALL SITES		
Status	# of jobs	events
done	134 846	8 572 681 500
failed	3 367	231 820 000
running	1 166	78 910 000
submitted	334	20 870 000
sys-failed	9 683	675 170 000
timeout	136	9 500 000
Total	149 532	9 588 951 500

Distributed prod: evt/site



Distributed prod: jobs/site



Sites collaboration

CNAF
Caltech
SLAC
McGill
UVIC
Queen Mary
RAL T1, T2
CCIN2P3, GRIF
INFN-Bari
INFN-LNL
INFN-Napoli
INFN-Ferrara
INFN-Pisa
INFN-Torino
INFN-Cagliari

Armando Fella
Michael Thomas, Frank Porter, Piti Ongmongkolkul
Steffen Luiz, Wei Yang
Steven Robertson
Ashok Agarwal
Adrian Bevan, Christopher Wilson, Alex Martin
Fergus Wilson, Chris Brew
Nicolas Arnaud
Giacinto Donvito, Vincenzo Spinoso
Gaetano Maron, Alberto Crescente, Sergio Fantinelli
Silvio Pardi, Alessandra Doria
Luca Tomassetti, Eleonora Luppi, Marco Ronzano
Alberto Ciampa, Enrico Mazzoni, Dario Fabiani
Stefano Bagnasco, Stefano Lusso, Riccardo Brunetti
Daniele Mura

- A stable and conscious collaboration has been achieved
- Effort impact for site setup and configuration decreased
- New site inclusion procedure improved
- **Well done!! Thanks again**

SuperB Grid integration

- The Virtual Organization has been included into the Global Grid User Support system:

<https://gus.fzk.de/pages/home.php>

- The VO now is included into - Grid view – LCG Grid Site availability system:

http://gridview.cern.ch/GRIDVIEW/same_index.php

CNAF disk space

- **Size Used Avail Use% Mounted on**

15T 14T 1.1T 93% /storage/gpfs_babar6

- User AWG: ~3TB
- Old production: ~1TB
- BaBar stuff: 10TB


43T 31T 13T 72% /storage/gpfs_superb

- 2010_July occupancy: 15TB
- 2010_September occupancy: 10TB
- Old production occupancy: 6TB

CNAF status

- CNAF efficiency increased (x2) wrt Feb. prod.
- CNAF fail rate almost zero

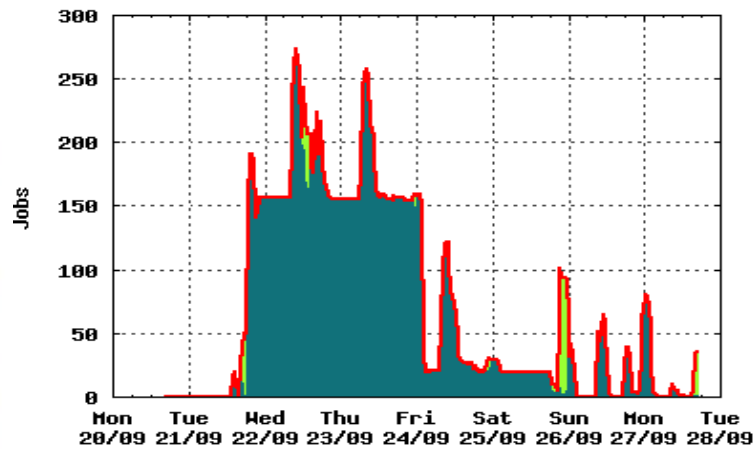
superb_prod users




CNAF status

- Overload reading data from gpfs_superb

superb queue users




superb_prod users




- After short discussion with CNAF experts, need to add 2 gpfs server machines

Network utilization



'Weekly' Graph (30 Minute Average)



Max In:987.4 Mb/s (98.7%) Average In:347.0 Mb/s (34.7%) Current In:533.7 Mb/s (53.4%)
 Max Out:265.0 Mb/s (26.5%) Average Out:38.7 Mb/s (3.9%) Current Out:13.8 Mb/s (1.4%)

Todo list

- 1) il ganga con multiple sottomissioni continua a fare il cambio di stato di tutti i job delle sottomissioni precedenti. Dovremmo inibire del tutto le funzioni di monitoring mantenendo le funzioni di raccolta dei gjid. Possibile soluzione farla remove dalla dir ad ogni lancio di script php e non solo all'inizio.
- 2) aggiungere al base_job_lcg la gestione dei signal con la funzione trap:
 - trap traponint INT TERM EXIT
 - traponint()
 - { echo "raccolto un segnale \$SIGNAL"
 - 1) aggiornamento dello stato del job via rest
 - 2) trasferimento del log via lcg-cr o via rest nel db
 - exit 0
 - } vedere anche <http://www.davidpashley.com/articles/writing-robust-shell-scripts.html>
- 3) aggiungere il check di riempimento della request con popup js all'utente in cui si da la possibilita' di andare a riempire la request successiva, in aggiunta si potrebbe dare la possibilita' di selezionare piu' request contemporaneamente.
- 4) fare lo status verboso che funziona
- 5) togliere i campi analysis nella fast_log (forse anche in fast_output)
- 6) nel base_ganga.gpi aggiungere le request qui di seguito piu' altre piu' furbe; non mettere il ce:coda direttamente, ma far fare la scelta a lui con le wilde card per risolvere il problema delle shallow resubmission dovute o a temporanea unavailability del sito o a problemi di proxy expiration.
- 7) aggiornare ganga all'ultima versione.
- 8) mettere in collegamento il db con l'aggiornamento di quel file, check sulla rev, cosi' da controllare che la rev in uso sia quella scritta nel file cfg
- 9) problema di velocita' di ganga: i submitted sono running
- 10) quando lo status va in sysfailed il timestamp si aggiorna a quello di running ed invece deve aggiornare solo l'ultimo time
- 11) Post production: