CMS experience on EDG testbed

A.Fanfani

Dept. of Physics and INFN, Bologna

on behalf of CMS/EDG Task Force

- Introduction
- Use of EDG middleware in the CMS experiment:
  - CMS/EDG Stress test
  - Other Tests

Introduction

- Large Hadron Collider
- CMS (Compact Muon Solenoid) Detector
- CMS Data Acquisition
- CMS Computing Model

Large Hadron Collider LHC

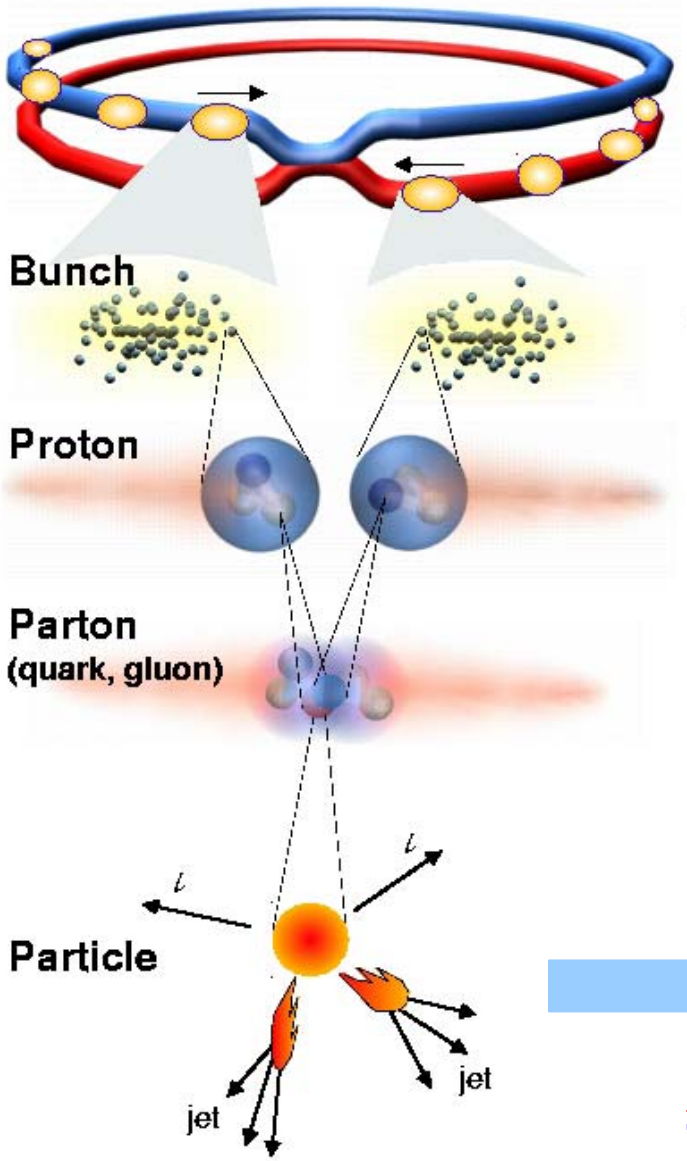

Proton- Proton Collision

Beam energy : 7 TeV
Luminosity : 10^{34} cm^{-2} s^{-1}
Data taking : > 2007

bunch-crossing rate: 40 MHz

~20 p-p collisions for each bunch-crossing
p-p collisions ≈ 10^9 evt/s (Hz)

CMS detector

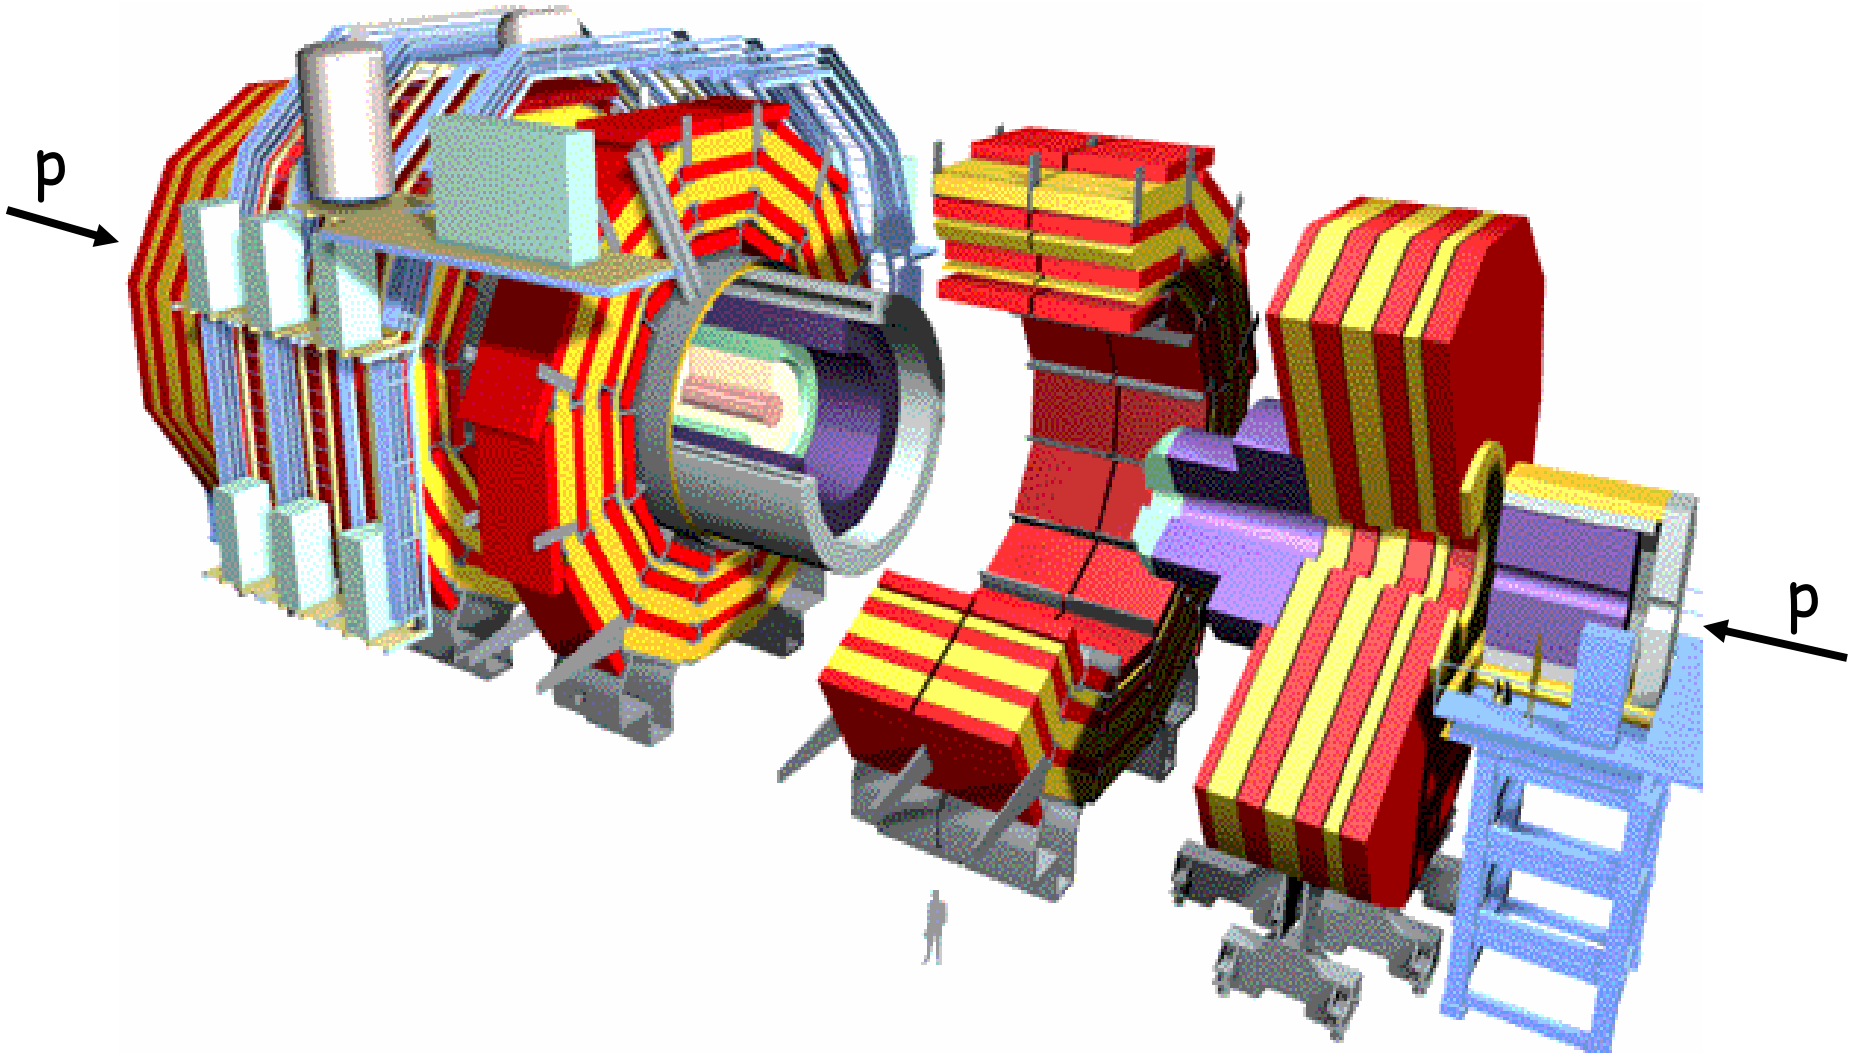

CMS Data Acquisition

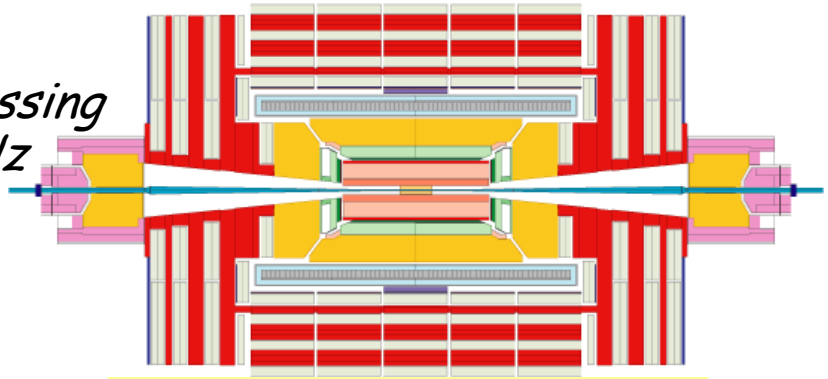

Bunch crossing 40 MHz

1event is ~ 1MB in size

~ GHz (~ PB/sec)

Level 1 Trigger - special hardware

75 KHz (75 GB/sec)

High Level Trigger - PCs

100 Hz
(100 MB/sec)

Online system

multi-level trigger to:
- filter out not interesting events
- reduce data volume

data recording

data

Offline analysis

CMS Computing

- **Large scale distributed Computing and Data Access**
  - Must handle PetaBytes per year
  - Tens of thousands of CPUs
  - Tens of thousands of jobs
  - heterogeneity of resources :
    hardware, software, architecture and Personnel

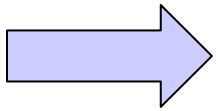

CMS Computing Hierarchy

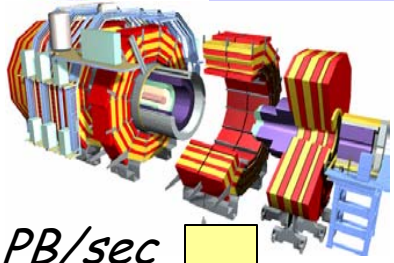

~ PB/sec

~ 100MB/sec

1PC* → PIII 1GHz

Online system

Offline farm

Tier 0

CERN Computer center
~10K PCs*

Tier 1

France Regional Center ~2K PCs

Italy Regional Center

Fermilab Regional Center

• • •

~ 2.4 Gbits/sec

Tier 2

Tier2 Center ~500 PCs

Tier2 Center

Tier2 Center

• • •

~ 0.6 – 2. Gbits/sec

Tier 3

workstation

InstituteA

InstituteB

~ 100-1000 Mbits/sec

CMS Production and Analysis

- The main computing activity of CMS is currently related to the simulation, with Monte Carlo based programs, of how the experimental apparatus will behave once it is operational

- The importance of doing simulation

large samples of simulated data are needed to :

  - optimise the detectors and investigate any possible modifications required to the data acquisition and processing
  - better understand the physics discovery potential
  - perform large scale test of the computing and analysis models

This activity is know as "CMS Production and Analysis"

CMS MonteCarlo production chain

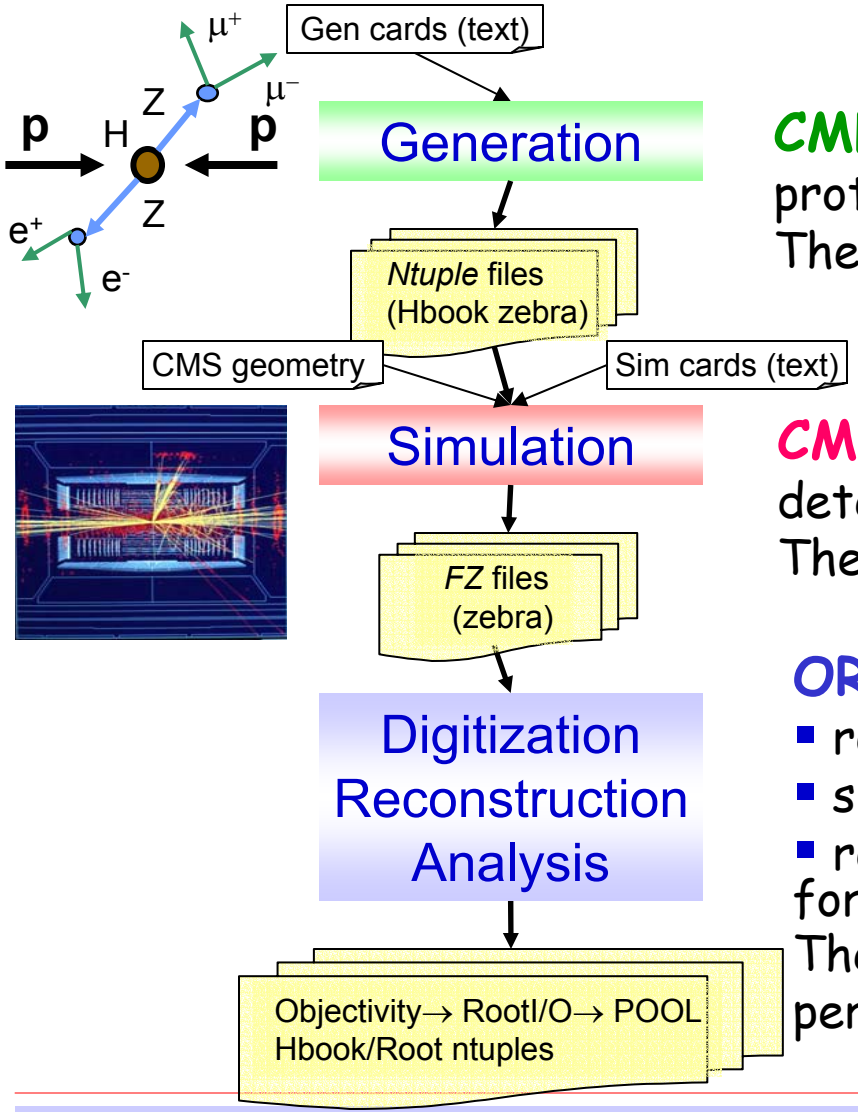

CMKIN: MonteCarlo Generation of the proton-proton interaction, based on PYTHIA.
The ouput is a random access zebra file (*ntuple*).

CMSIM: Simulation of tracking in the CMS detector, based on GEANT3.
The ouput is a sequential access zebra file (*FZ*).

ORCA:

- reproduction of detector signals (*Digis*)
- simulation of trigger response
- reconstruction of physical information

for final analysis
The replacement of Objectivity for the persistency will be POOL.

CMS Tools for "Production"

- **RefDB**
  - Contains production requests with all needed parameters to produce a physic channel and the details about the production process.
    It is a SQL Database located at CERN.
- **IMPALA**
  - Accepts a production request
  - Produces the scripts for each single job that needs to be submitted
  - Submits the jobs and tracks the status
- **MCRunJob**
  - Evolution of IMPALA: modular (plug-in approach)
- **BOSS**
  - tool for job submission and real-time job-dependent parameter tracking. The running job standard output/error are intercepted and filtered information are stored in BOSS database. The remote updator is based on MySQL .

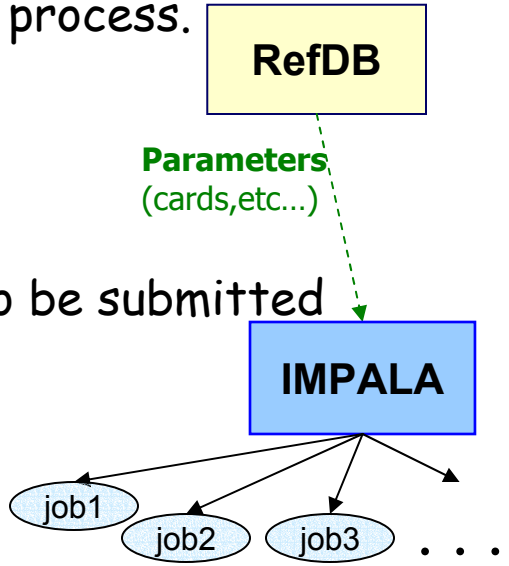

CMS/EDG Stress Test

- Test of the CMS event simulation programs in EDG environment using the full CMS production system
- Running from November 30th to Xmas
  (tests continued up to February)
- This was a joint effort involving CMS, EDG, EDT and LCG people

CMS/EDG Stress Test Goals

- Verification of the portability of the CMS Production environment into a grid environment;
- Verification of the robustness of the European DataGrid middleware in a production environment;
- Production of data for the Physics studies of CMS, with an ambitious goal of ~ 1 million simulated events in a 5 weeks time.

CMS/EDG Strategy

- Use as much as possible the High-level Grid functionalities provided by EDG:
  - Workload Management System (Resource Broker),
  - Data Management (Replica Manager and Replica Catalog),
  - MDS (Information Indexes),
  - Virtual Organization Management, etc.
- Interface (modify) the CMS Production Tools to the Grid provided access method
- Measure performances, efficiencies and reason of job failures to have feedback both for CMS and EDG

CMS/EDG Middleware and Software

- **Middleware** was: **EDG** from version 1.3.4 to version 1.4.3
  - Resource Broker server
  - Replica Manager and Replica Catalog Servers
  - MDS and Information Indexes Servers
  - Computing Elements (CEs) and Storage Elements (SEs)
  - User Interfaces (UIs)
  - Virtual Organization Management Servers (VO) and Clients
  - EDG Monitoring
  - Etc.

- **CMS software** distributed as rpms and installed on the CE
- **CMS Production tools** installed on UserInterface

User Interface set-up

CMS Production tools installed on the EDG User Interface

- IMPALA
  - Get from RefDB parameters needed to start a production
  - "JDL" files are produced along with the job scripts
- BOSS
  - BOSS will accept and pass on a JDL file to the Resource Broker
  - Additional info is stored in the BOSS DB:
    - Logical file names of input/output files
    - Name of the SE hosting the output files
    - Outcome of the copy and registration in the RC of files
    - Status of the replication of files

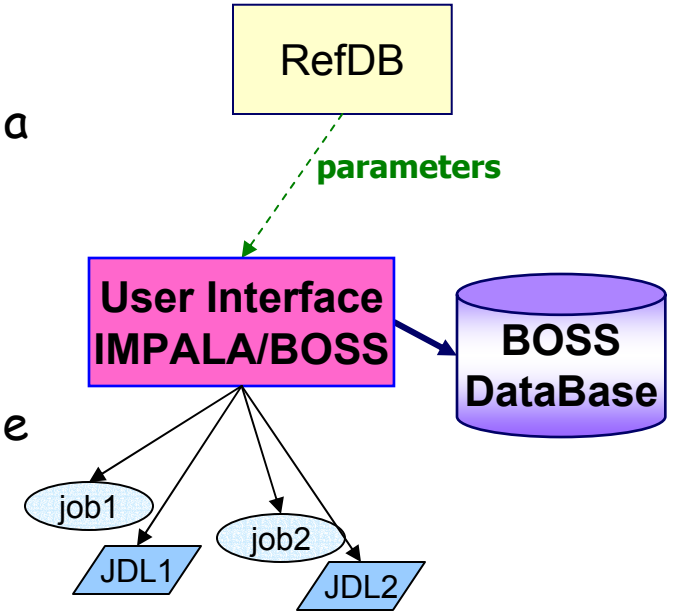

CMS production components interfaced to EDG middleware

- Production is managed from the EDG User Interface with IMPALA/BOSS

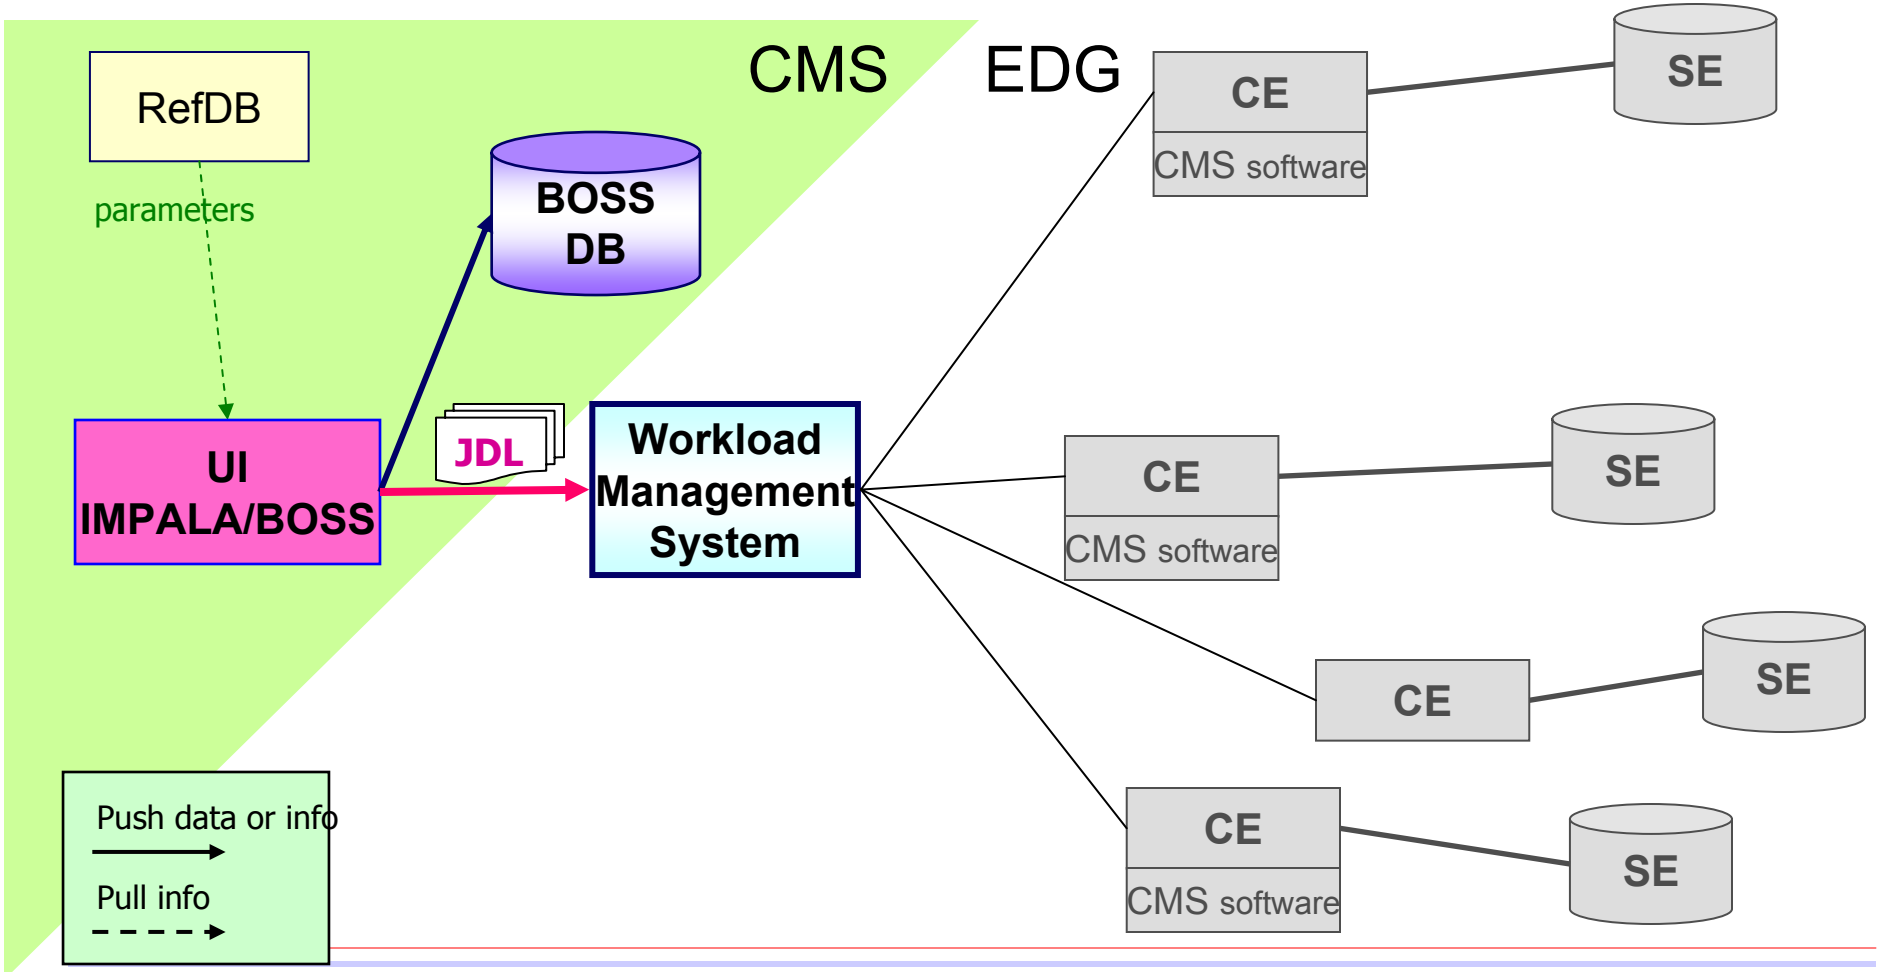

CMS jobs description

- **CMS official jobs for "Production" of results used in Physics studies**

Dataset *eg02_BigJets*

	size/event	time*/event
CMKIN	~ 0.05MB	~ 0.4-0.5 sec
CMSIM	~ 1.8 MB	~ 6 min

* PIII 1GHz
512MB
~ 46.8 SI95

- Production in 2 steps:

1. **CMKIN** : MC Generation for a physics channel (dataset)

   125 events ~ 1 minute ~ 6 MB ntuples — "Short" jobs

2. **CMSIM** : CMS Detector Simulation

   125 events ~ 12 hours ~ 230 MB FZ files — "Long" jobs

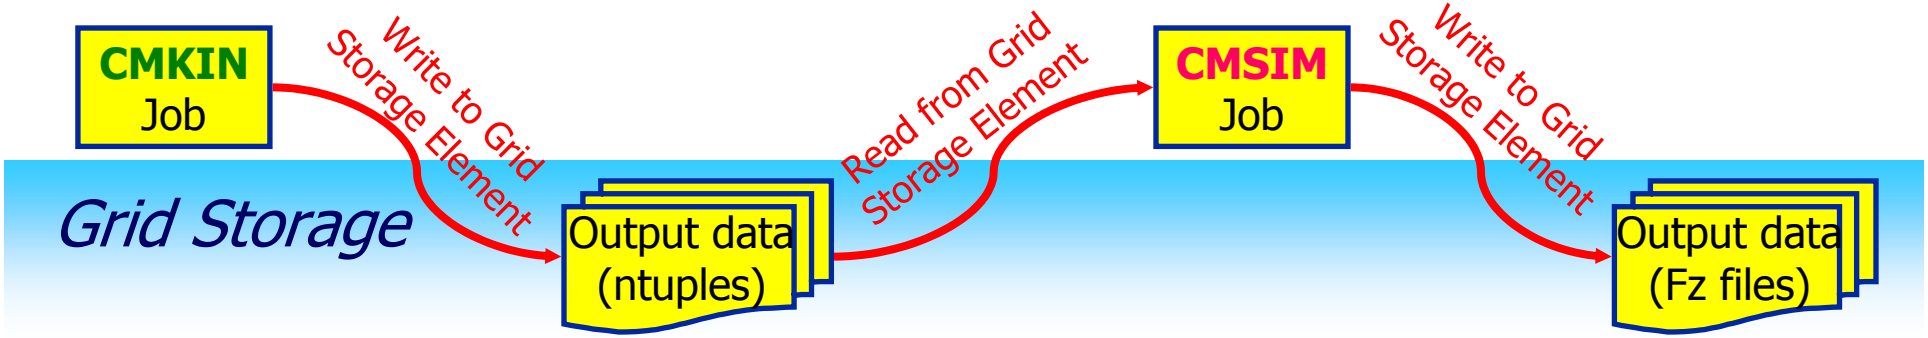

CMKIN Workflow

IMPALA creation and submission of **CMKIN** jobs:

- Resource Broker sends jobs to Computing resources (CEs) having CMS software installed
- Output ntuples are saved on Close SE and registered into ReplicaCatalog with a Logical File Name (LFN)
- the LFN of the ntuple is recorded in the BOSS Database

CMS production of CMKIN jobs

- CMKIN jobs running on all EDG Testbed sites with CMS software installed

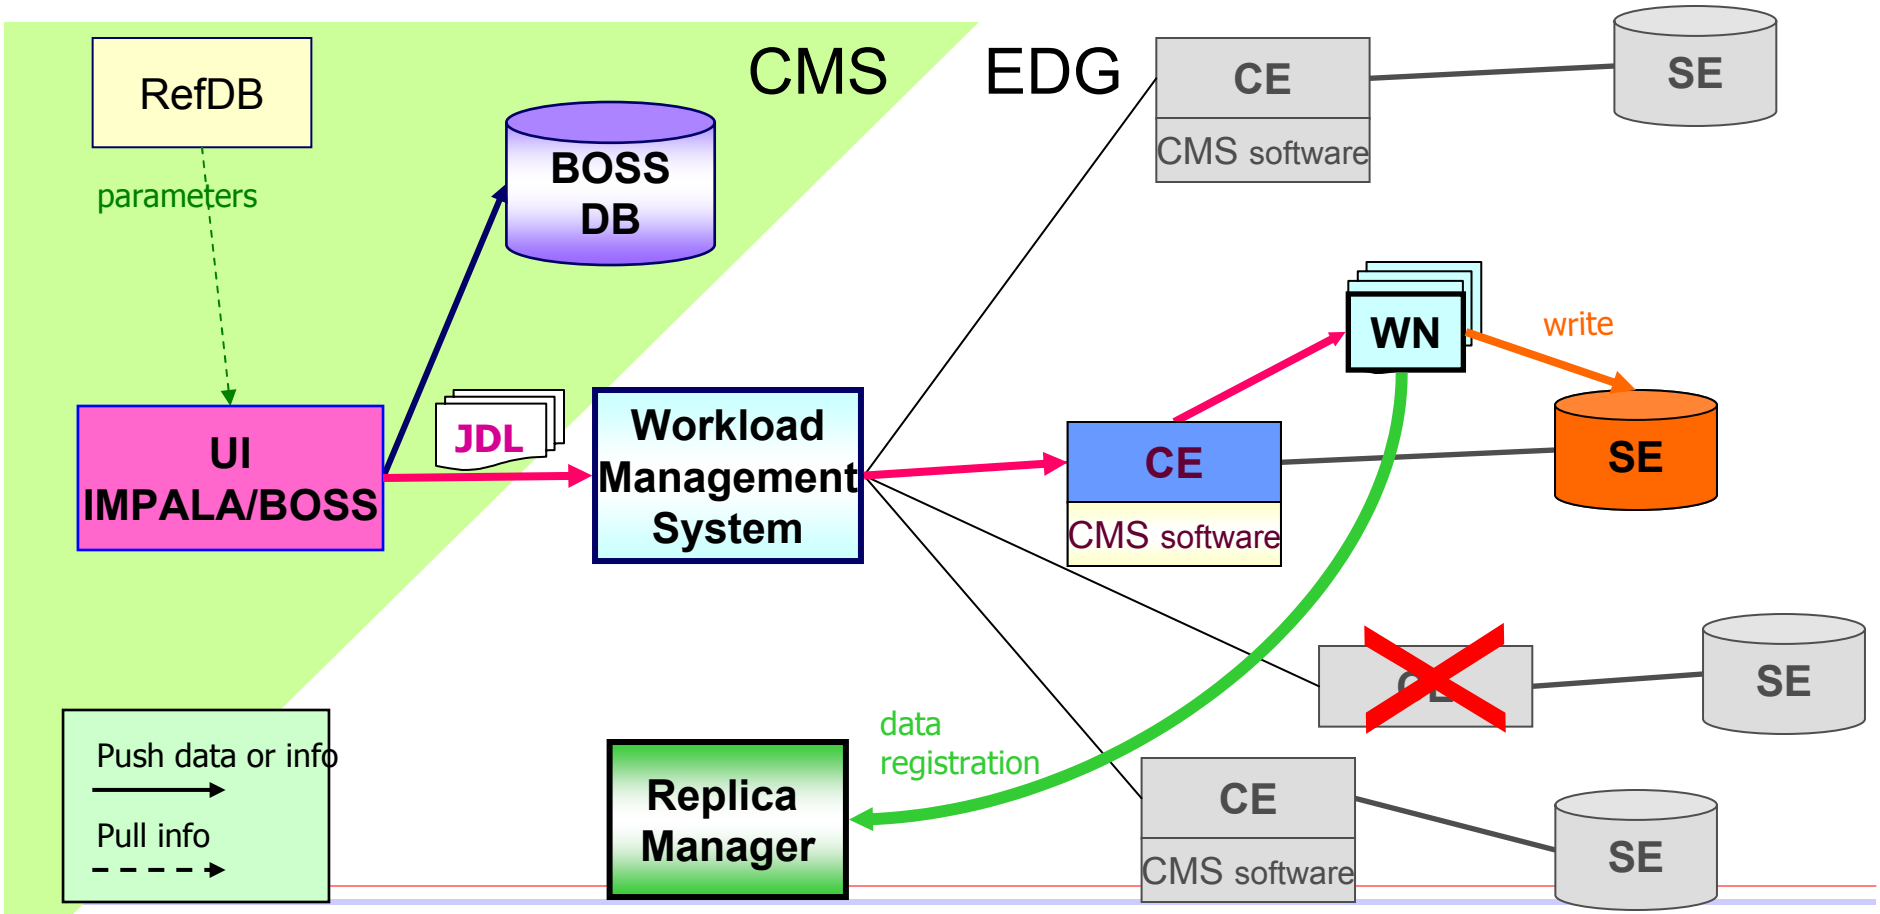

CMSIM Workflow

IMPALA creation and submission of **CMSIM** jobs:

- Computing resources are matched to the job requirements:
  - Installed CMS software, MaxCPUTime, etc.
  - CE near to the input data that have to be processed
- FZ files are saved on Close SE or on a predefined SE and registered in the Replica Catalog
- the LFN of the FZ file is recorded in the BOSS DB

CMS production of CMSIM jobs

- CMSIM jobs running on CE close to the input data

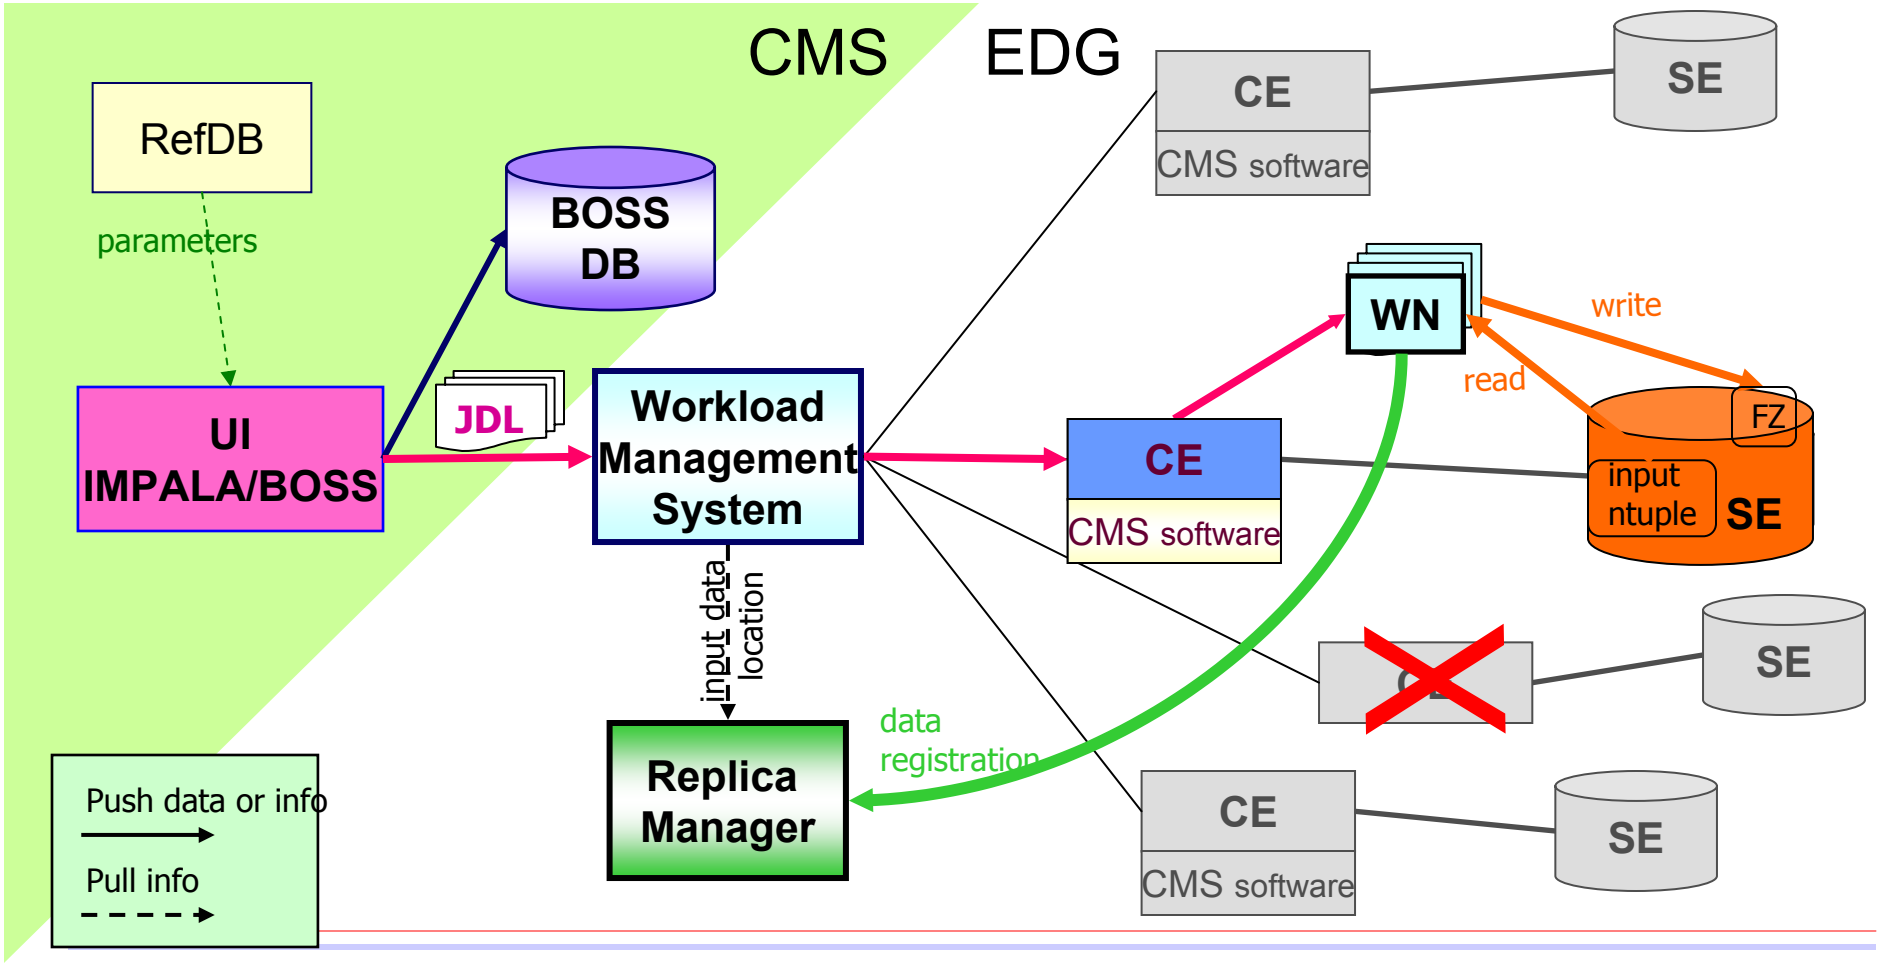

Data management

- Two practical approaches:
  1. FZ files are directly stored at some dedicated SE
  2. FZ files are stored on the "close SE" and later replicated to CERN

  test the creation of replicas of files : 402 FZ files (~ 96GB) were replicated

- All sites use disk for the file storage, but :

Mass Storage

- CASTOR at CERN: FZ files replicated to CERN are also automatically copied into CASTOR
- HPSS in Lyon: FZ files stored in Lyon are automatically copied into HPSS

monitoring CMS jobs

- Job monitoring and bookkeeping: BOSS Database, EDG Logging & Bookkeeping service

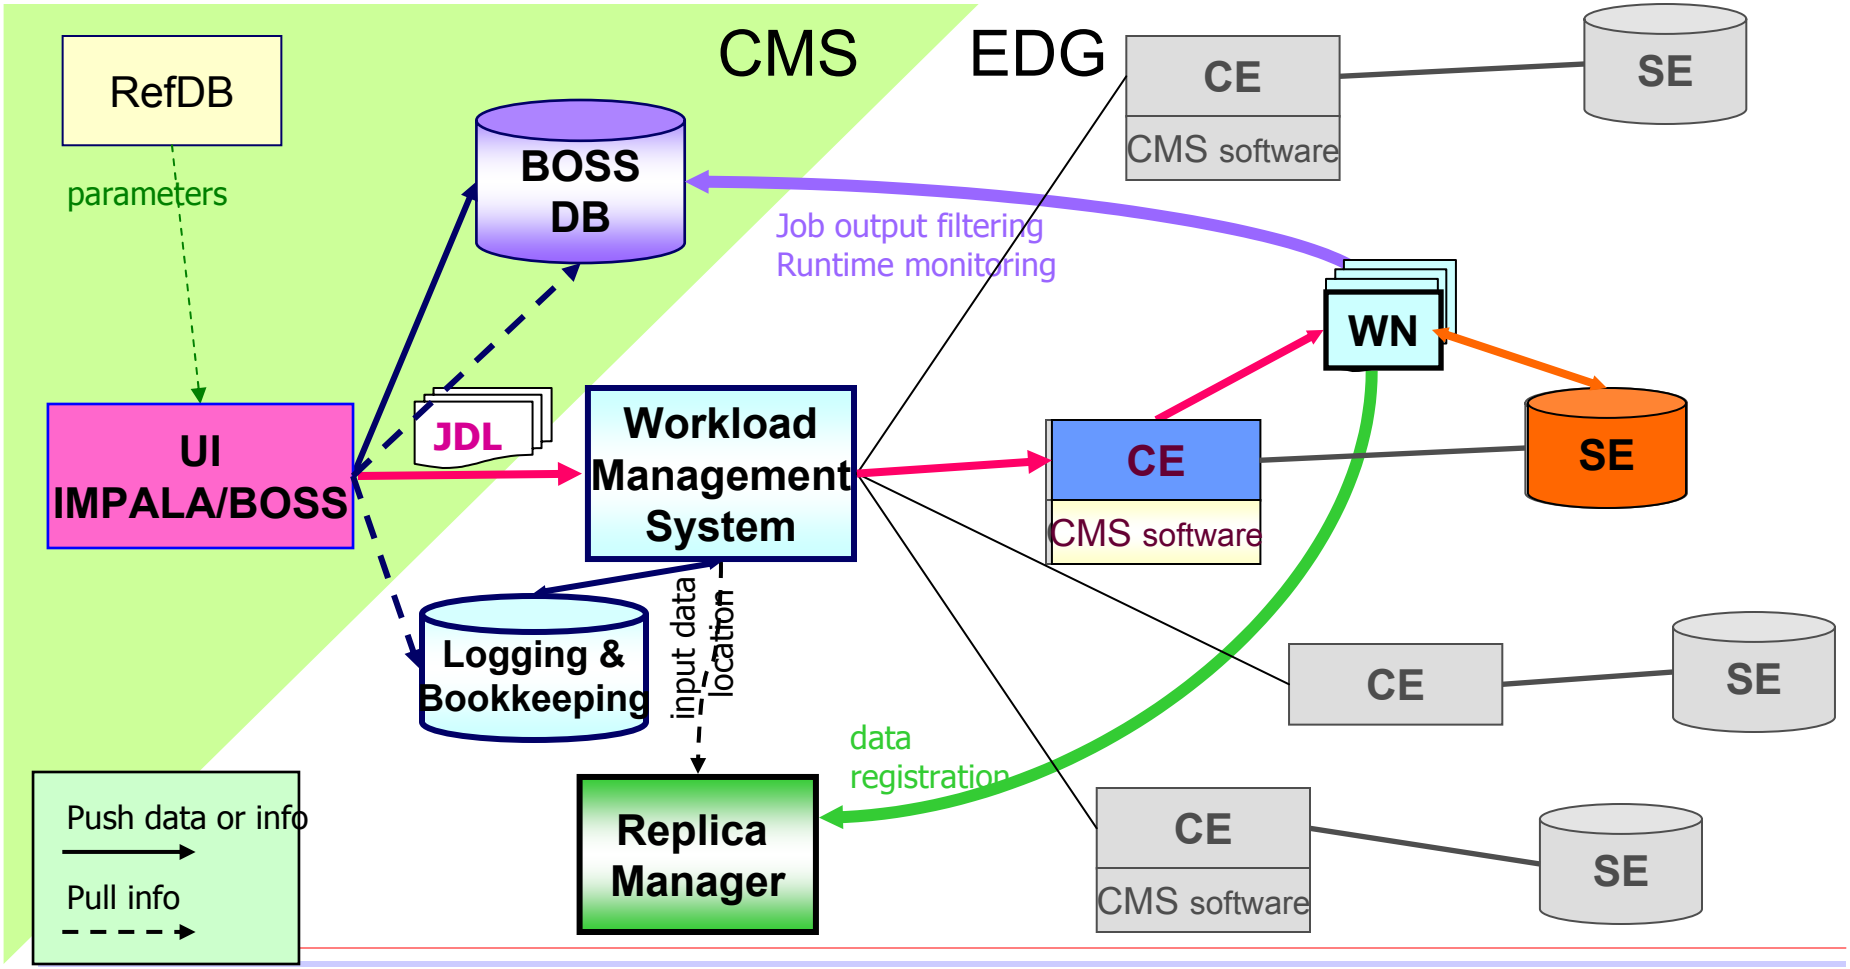

Monitoring the production

Information about the job: nb. of events, executing host, …
from BOSS database *(boss SQL)*

Job status
from L & B
(dg-job-status)

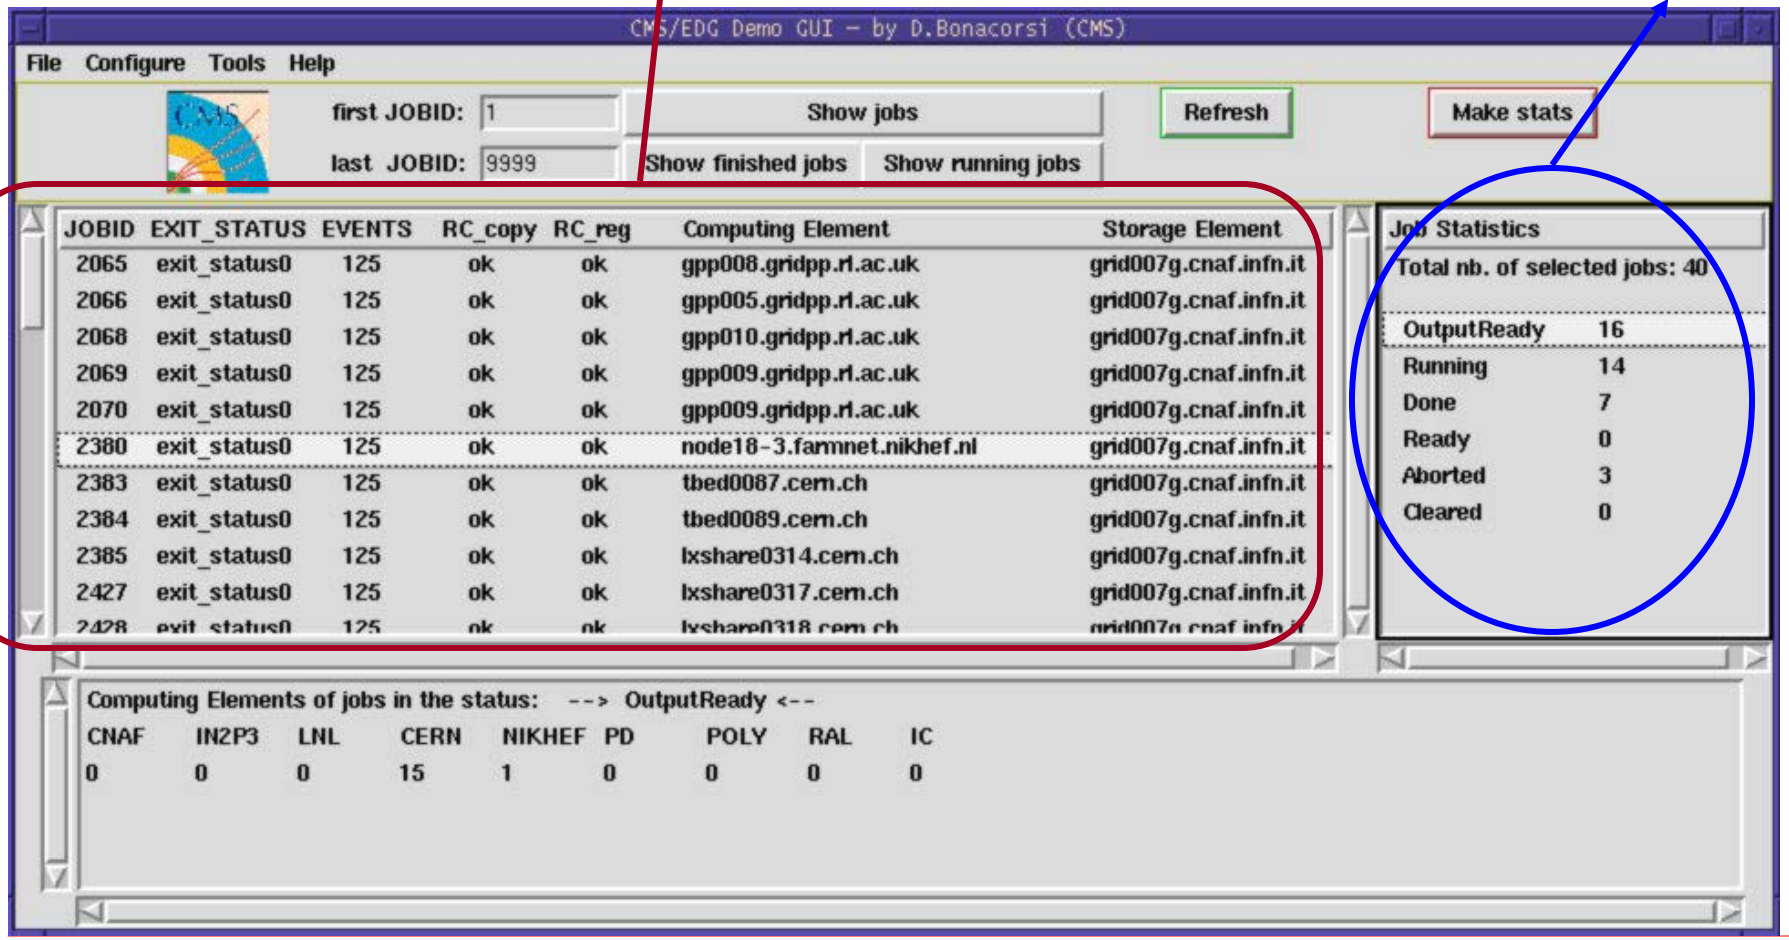

Monitoring

Offline monitoring :

- Two main sources of information:
  - EDG monitoring system (MDS based)
    - MDS information is volatile and need to be archived somehow
    - collected regularly by scripts running as cron jobs and stored for offline analysis
  - BOSS database
    - permanently stored in the MySQL database

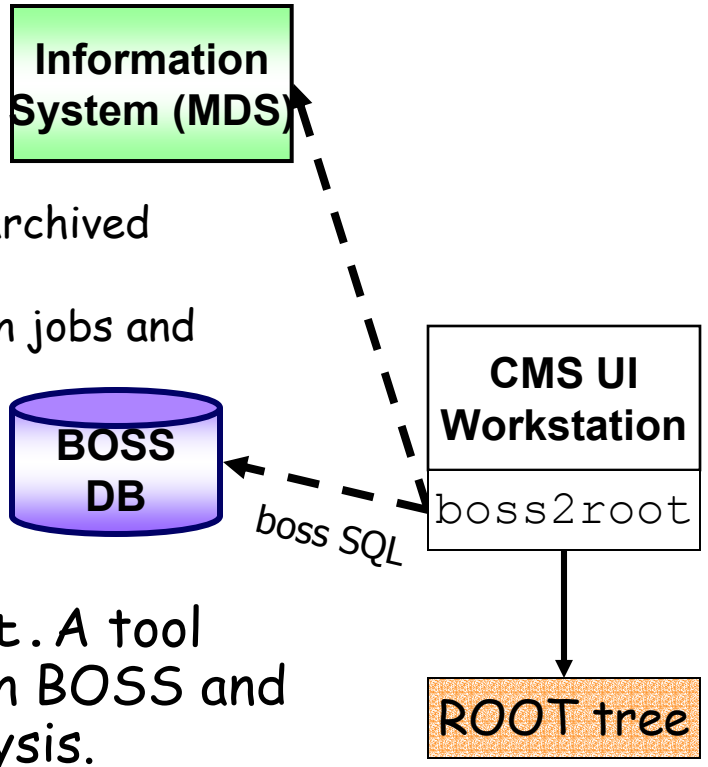

- Both sources are processed by **`boss2root`**. A tool developed to read the information saved in BOSS and store them in ROOT tree to perform analysis.

Online monitoring :

with Nagios, web based tool developed by the DataTag project

Organisation of the Test

- Four UI's controlling the production:
  - Bologna / CNAF
  - Ecole Polytechnique
  - Imperial College
  - Padova
  - reduces the bottleneck due to the BOSS DB
- Several resource brokers (each seeing all resources):
  - CERN (dedicated to CMS) (EP UI)
  - CERN (common to all applications) (backup!)
  - CNAF (common to all applications) (Padova UI)
  - CNAF (dedicated to CMS) (CNAF UI)
  - Imperial College (dedicated to CMS and BABAR) (IC UI)

  - reduces the bottleneck due to intensive use of the RB and the 512-owner limit in Condor-G
- Replica catalog at CNAF
- Top MDS at CERN
- II at CERN and CNAF
- VO server at NIKHEF

EDG hardware resources

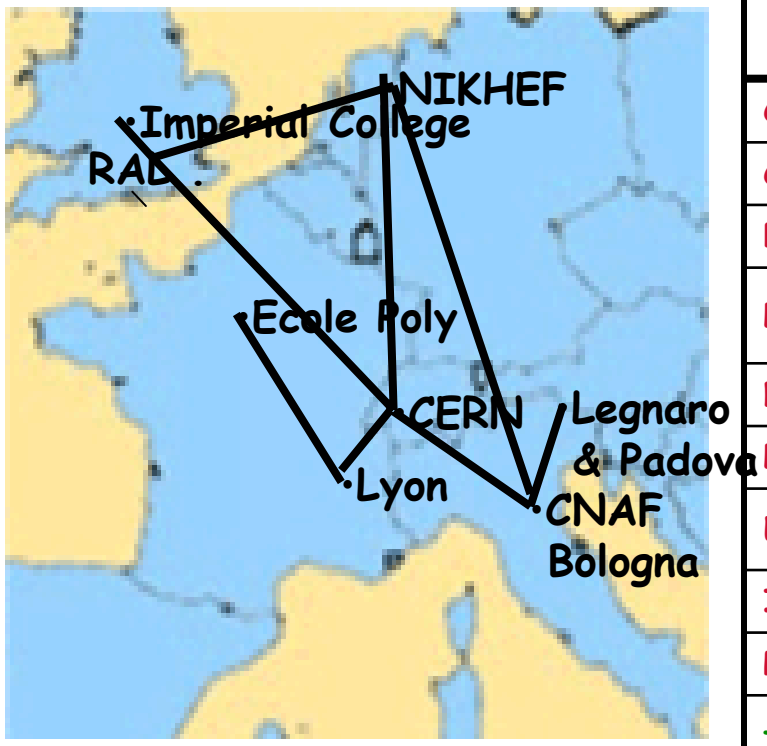

Site	Number of CPUs	Disk Space GB	Availability of MSS
CERN (CH)	122	1000* (+100)	yes
CNAF (IT)	20 + 20*	1000*	
RAL (UK)	16	360	
Lyon (FR)	shared 120 (400)	200	yes
NIKHEF (NL)	22	35	
Legnaro (IT)*	50	1000*	
Ecole Polytechnique (FR)*	4	220	
Imperial College (UK)*	16	450	
Padova (IT)*	12	680	
Totals	402 (400)	3000* + (2245)	

*Dedicated to CMS Stress Test

distribution of job: executing CEs

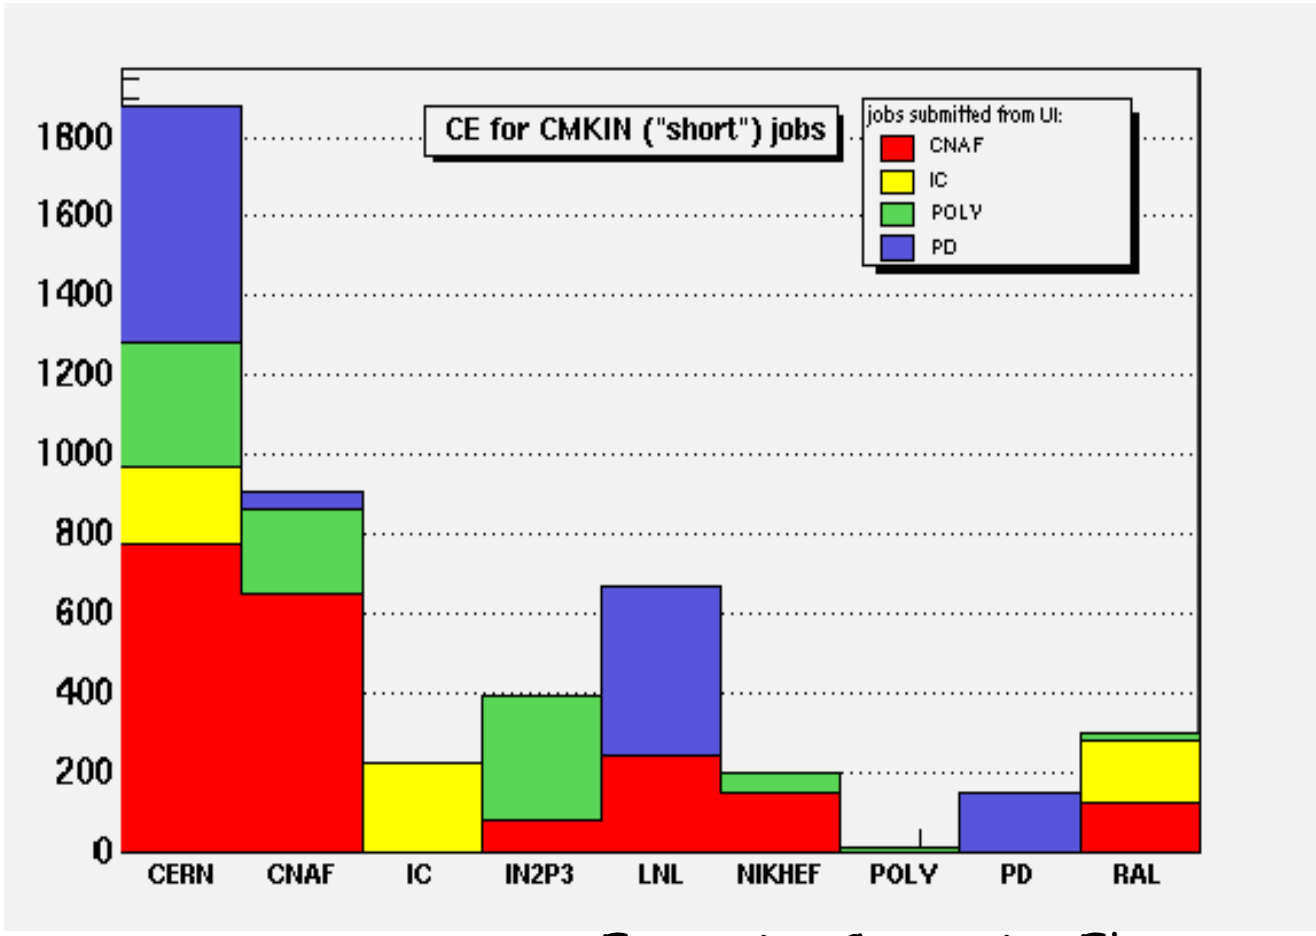

CMS/EDG Production

CMKIN "short" jobs

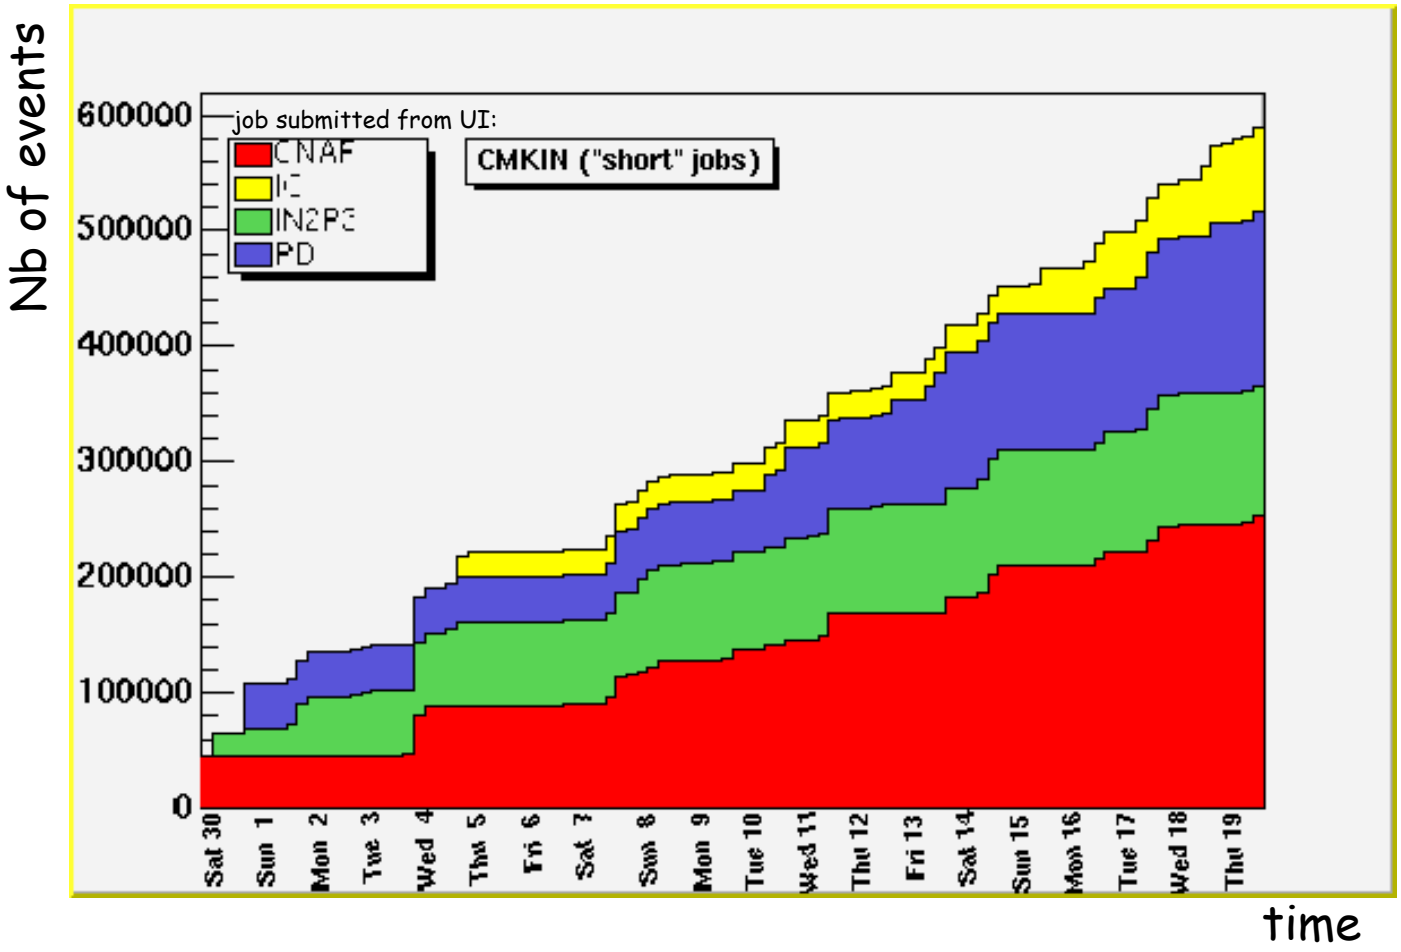

CMS/EDG Production

CMSIM "long" jobs

- ~260K events produced
- ~7 sec/event average
- ~2.5 sec/event peak (12-14 Dec)

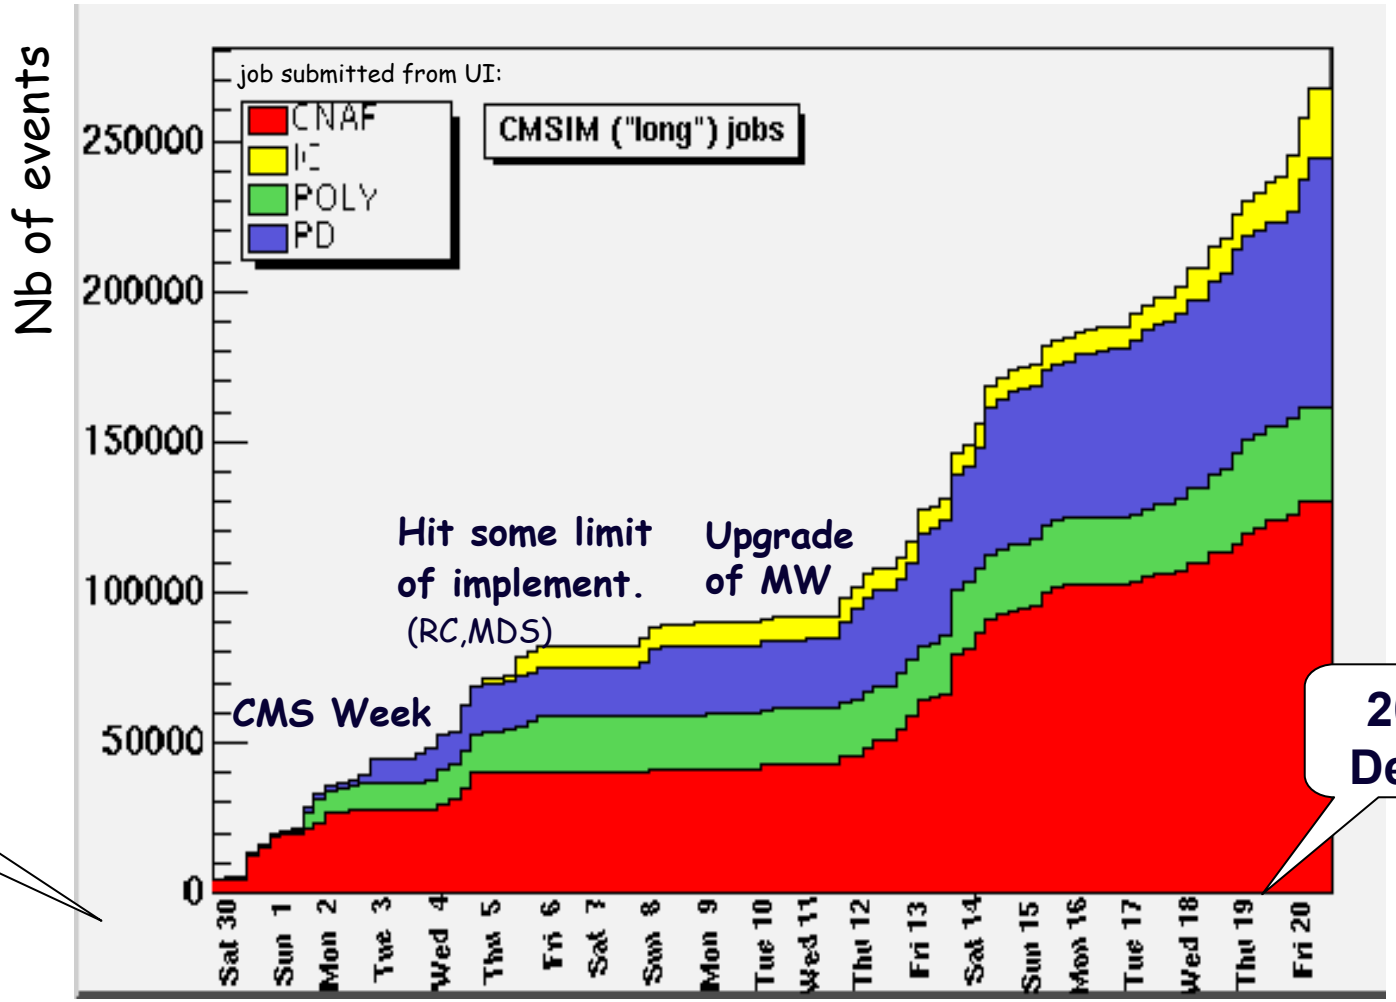

Total no. of events

UI Submitting	Total no. of CMKIN evts	% of total	Total no. of CMSIM evts	% of total
CNAF	253625	43	130250	48
PD	151750	26	82625	31
IC	73125	12	23375	9
POLY	114250	19	32125	12
Total	**592750**		**268375**	

- each job with 125 events
- 0.05 MB/event (CMKIN)
- 1.8 MB/event (CMSIM)

$\Rightarrow$ Total number of successful jobs: ~ 7000

$\Rightarrow$ Total size of data produced: ~ **500 GB**

Summary of Stress Test

CMKIN jobs		
Status	EDG Evaluation	"CMS Evaluation"
Finished Correctly	5518	4742
Crashed or bad status	818	958
Total number of jobs	6336	5700
Efficiency	87%	83%

"Short" jobs

EDG Evaluation:
- All submitted jobs are considered
- Successful jobs are those correctly finished for EDG

"CMS Evaluation":
- only jobs that had a chance to run are considered
- Successful jobs are those with the output data properly stored

CMSIM jobs		
Status	EDG Evaluation	"CMS" Evaluation
Finished Correctly	1678	2147
Crashed or bad status	2662	935
Total number of jobs	4340	3082
Efficiency	39%	70%

"Long" jobs

Total EDG Stress Test jobs = 10676 , successful =7196 , failed = 3480

EDG reasons of failure (categories)

CMKIN jobs	
Status	**Totals**
Crashed or bad status	818
Reasons of Failure for Crashed jobs	
No matching resource found	509
Generic Failure:MyProxyServer not found in JDL expr.	102
Running forever	74
Failure while executing job wrapper	37
Other failures	96

"Short" jobs

"Long" jobs

CMSIM jobs	
Status	**Totals**
Crashed or bad status	2662
Reasons of Failure for Crashed jobs	
Failure while executing job wrapper	1476
No matching resource found	722
Globus failure: Globus down/submit to globus failed	144
Running forever	116
Globus failure	90
Other failures	114

main sources of trouble (I)

- The Information service (MDS and Information Index) weakness
  → "*No matching resources found*" error :

  - As the query rate increase the top MDS and II slow down dramatically. Since the RB relies on the II to discover available resources, the MDS instability caused job to abort due to lack of matching resources.

  - ✓ Work-around: Use a cache of the information stored in a Berkeley database LDAP back-end (*from EDG version 1.4*).
    The rate of aborted jobs due to information system problems was reduced from 17% to 6%

main sources of trouble (II)

- Problems in the job submission chain related to the Workload Management System
  - → "*Failure while executing job wrapper*" error:
    (the most relevant failure for "long" jobs)

    - Failures in downloading/uploading the *Input/Output Sandboxes* files from RB to WN
      Due for example to problems in the gridftp file transfer, network failures, etc….
    - The standard output of the script where the user job is wrapped around was empty. This is transferred via *Globus GASS* from the CE node to the RB machine in order to check if the job reached the end.
      There could be many possible reasons (i.e. home directory not available on WN, glitches in the *GASS* transfer, race conditions for file updates between the WN and CE node with PBS etc..)

    - ✓ Several fixes to reduce this effect (if necessary transfer the stdout also with gridftp, PBS specific fixes,…) *(from EDG1.4.3)*

main sources of trouble (III)

- Replica catalog limitation of performances
  - limit of the number of lengthy named entries in one file collection
    ⇒ several collections used
  - The catalog respond badly to a high query/writing rate, with queries hanging indefinitely.
    ⇒ a very difficult situation to deal with since the jobs hung while accessing and stayed in *"Running"* status forever, and thus requiring manual intervention from the local system administrators

The efficiency of copy the output file into SE and register it into RC :

	ε (copy)	ε (register)	ε (copy & register)	
CMKIN	97%	86%	83%	← small output file, higher writing rate into RC
CMSIM	84%	93%	78%	← bigger output file, slower writing rate into RC

Total number of files written into RC : ~ 8000

- Some instability of the Testbed due to a variety of reasons (from hardware failures, to network instabilities, to mis-configurations)

Tests after the StressTest

- Including fixes and performance enhancements mainly to reduce the rate of failures in the job submission chain

CMKIN jobs	
Status	**EDG Evaluation**
Finished Correctly	1014
Crashed or bad status	57
Total number of jobs	1071
Efficiency	95%

"Short" jobs

Increased efficiency
in particular for long jobs
(Limited statistic wrt Stess Test)

CMSIM jobs	
Status	**EDG Evaluation**
Finished Correctly	653
Crashed or bad status	264
Total number of jobs	917
Efficiency	71%

"Long" jobs

Main results and observations

- RESULTS
  - Could distribute and run CMS software in EDG environment
  - Generated ~250K events for physics with ~10,000 jobs in 3 week period
- OBSERVATIONS
  - Were able to quickly add new sites to provide extra resources
  - Fast turnaround in bug fixing and installing new software
  - Test was labour intensive (since software was developing and the overall system was fragile)
    - ***WP1 At the start there were serious problems with long jobs- recently improved***
    - ***WP2 Replication Tools were difficult to use and not reliable, and the performance of the Replica Catalogue was unsatisfactory***
    - ***WP3 The Information System based on MDS performed poorly with increasing query rate***
    - ***The system is sensitive to hardware faults and site/system mis-configuration***
    - ***The user tools for fault diagnosis are limited***
  - EDG 2.0 should fix the major problems providing a system suitable for full integration in distributed production

Other tests: systematic submission of CMS jobs

- Use CMS jobs to test the behaviour/response of the grid as a function of the jobs characteristics
- No massive tests in a production environment
- systematic submission over a period of ~ 4 months (march-june)

characteristics of CMS jobs

- CMS jobs with different CPU and I/O requirements, varying:
  - Kind of application : CMKIN and CMSIM jobs
  - Number of events: 10, 100 , 500
  - Cards file : define the kind of events to be simulated
    datasets "ttbar", "eg02BigJets", "jm_minbias"

  18 different kind of jobs

- Measure the requirements of these jobs in term of:
  - Resident Set Size
  - Wall Clock Time
  - Input size
  - Output size

i.e. →

Definition of Classes and strategy for job submission

- Definition of classes of jobs according to their characteristics:

	Class G1	Class G2	Class G3	Class G4
Time (h)	0-0.08	0.08-1.78	1.78-7.72	7.72-46.31
RSS (MB)	0-10	10-100	10-100	10-100
Input (MB)	0-10	0-10	0-10	10-30
Output (MB)	0-50	0-50	50-200	200-1000

Not demanding CMKIN jobs

CMSIM jobs with increasing requirements

- Submission of the various kind of jobs to the EDG testbed:
  - use of the same EDG functionalities as described for the StressTest (Resource Broker, Replica Catalog, etc.....)
  - 2 Resource Broker were used (Lyon and CNAF)
  - several submission for each kind of jobs:
    - submission in bunches of 5 jobs
    - submission spread over a long period

Behaviour of the classes on EDG

- Comparison the Wall ClockTime and Grid Wall Clock Time

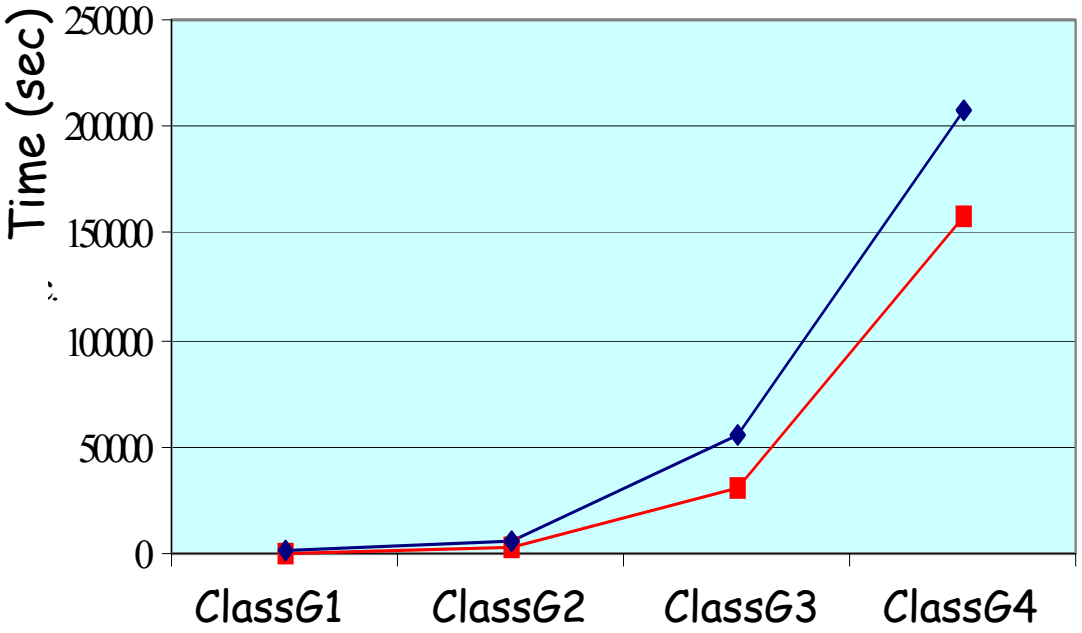

	GWCT (sec)	WCT (sec)	Overhead
ClassG1	1624.5	74.8	2072%
ClassG2	5467.4	3142.6	74%
ClassG3	55178.0	30265.0	82%
ClassG4	208092.2	157336.4	32%

- Report the failure rate for each class

	Failure rate %
ClassG1	26%
ClassG2	47%
ClassG3	53%
ClassG4	86%

Comments

- The behaviour of the identified classes of jobs on EDG testbed is:

increasing complexity ↓

	Time Overhead %	Failure rate %
ClassG1	2072%	26%
ClassG2	74%	47%
ClassG3	82%	53%
ClassG4	32%	86%

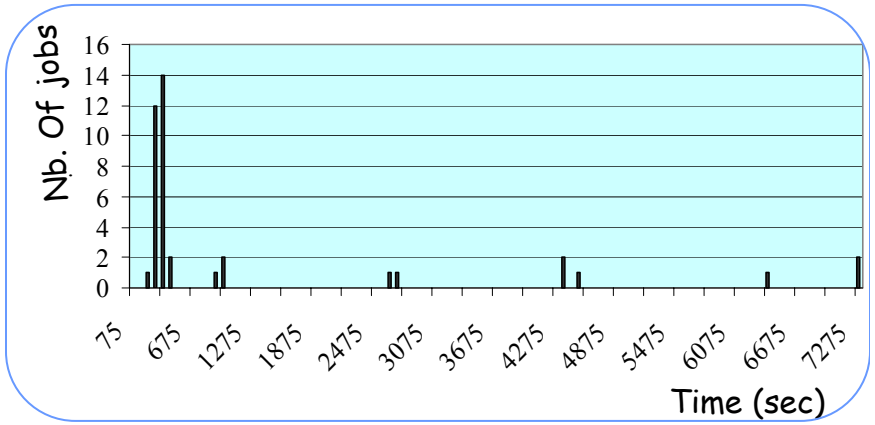

- The best class is G2 with an execution time ranging from 5 mins to ~2 hours
- Very short jobs have a huge overhead
  → Mean time affected by few jobs with strange pathologies
- The failure rate increases dramatically as the CPU time needed increases.
  → Instability of the testbed: i.e. there where frequent operational intervention on the RB which caused loss of jobs. Jobs lasting more then 20 hours have very little chances to survive

Conclusions

- HEP Applications requiring GRID Computing are already there
- All the LHC experiments are using the current implementations of many Projects
  - Need to test the scaling capabilities (Testbeds)
  - Robustness and reliability are the key issues for the Applications
- LHC experiments look forward for EGEE and LCG deployments