

Bookkeeping and submission tools prototype

L. Tomassetti

on behalf of distributed computing group

Outline

- General Overview
- Bookkeeping database
- Submission tools (for simulation productions)
- Framework Design and Grid Services
- Latest developments, future work
- Planning

Overview

- The production tools consist of
 - a bookkeeping database
 - a set of submission tools to the Grid
 - a Web Portal to manage and monitor submissions
- Development started in 2009

Overview

- Initial version devoted to FullSim, just one database table and a 'simple' submission script to LSF @CNAF (summer 2009)
- Then, FastSim included and initial Web interface provided – LSF @CNAF only (Winter 2009)
- First 'tagged' version in Spring 2010 with FastSim submissions to the Grid
- In Summer 2010 requests and shift interface added

Overview

- The underlying idea has been to **minimize** the platform requirements in terms of:
 - hardware and human resources,
 - custom developments,
 - Grid service customization and configuration.
- Keywords: **standard** and **simple**
- **The prototype framework aimed at supporting the Detector TDR activities**

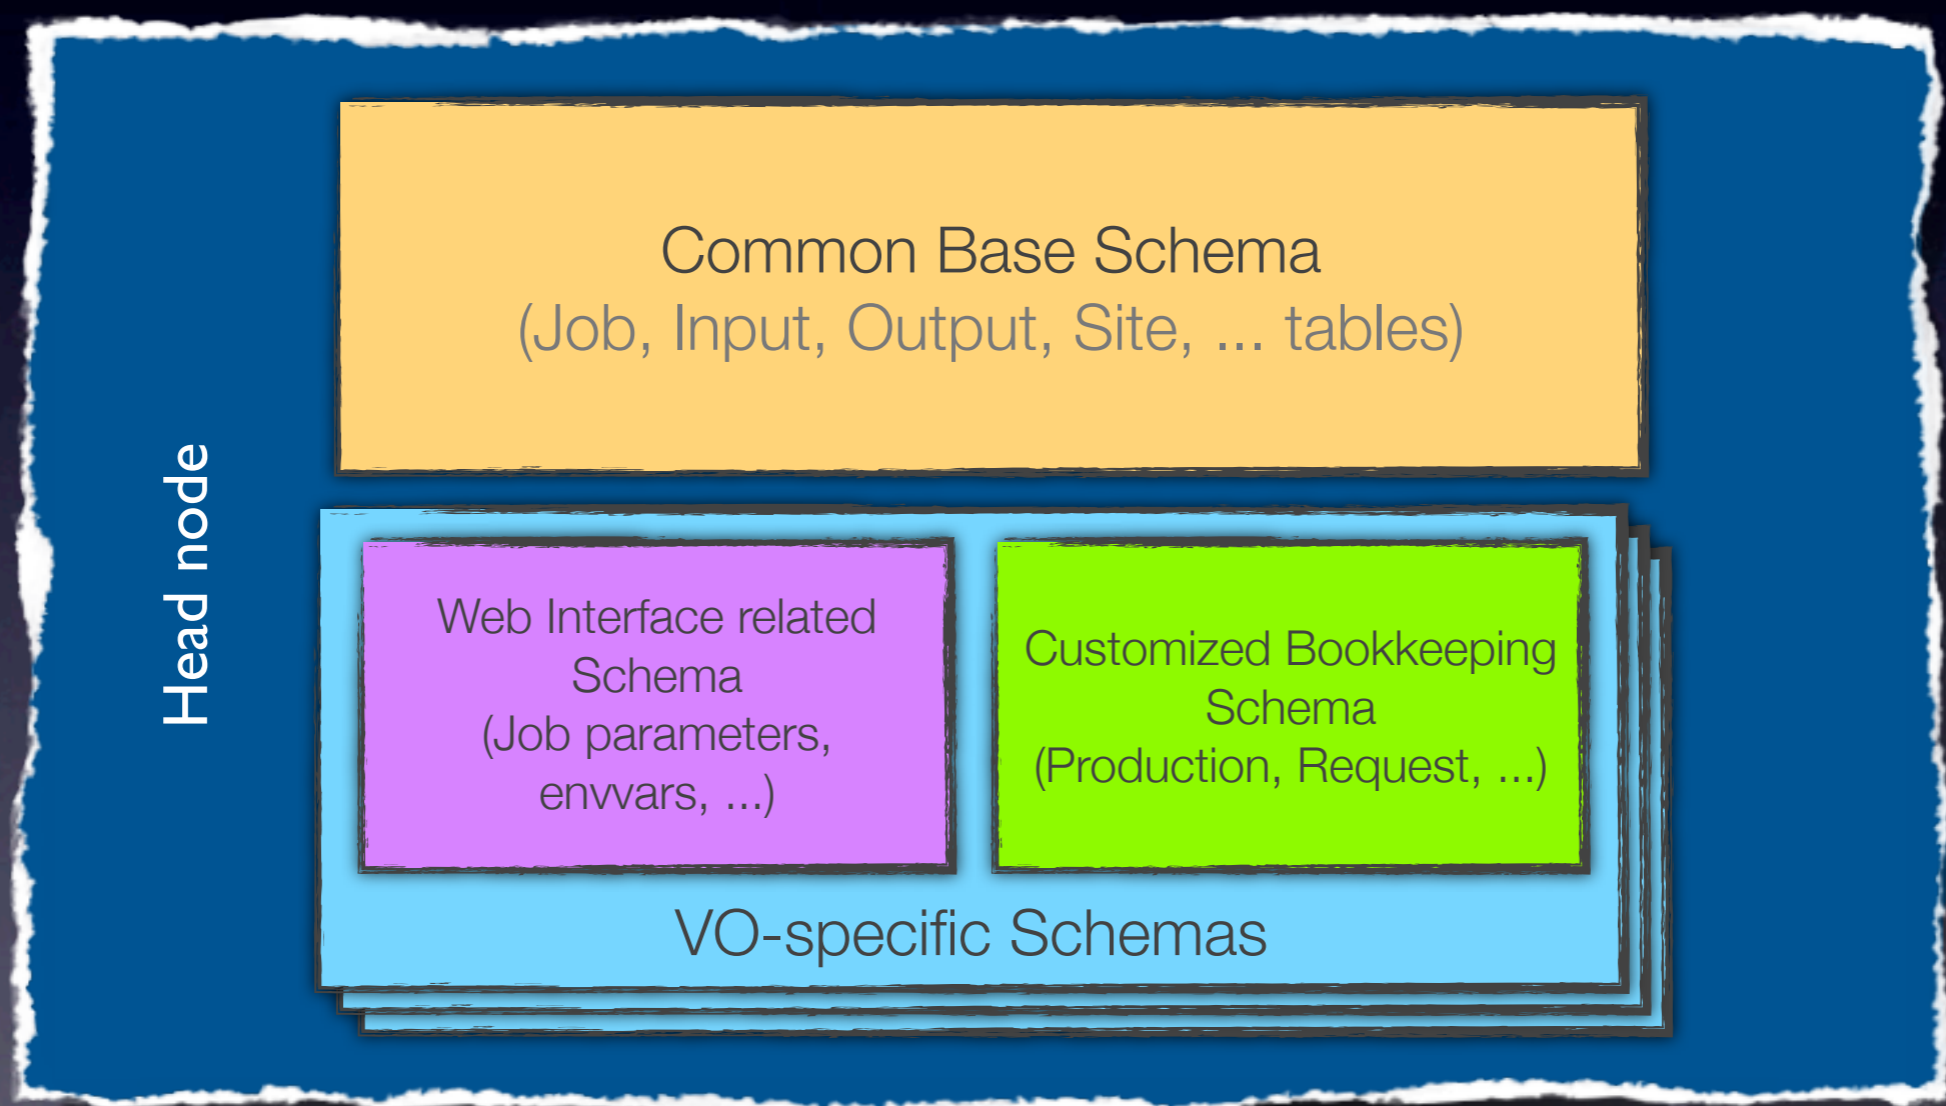
Overview

- Requirements:
 - superbvo.org enabled and software installed at sites
 - Input files $O(10GB)$, Output files $O(1GB)$, RAM $O(1GB)$
 - 1 Grid UI, 1 framework head node (it is the same at present)
 - 1 SE to store the results (central data repository)

Bookkeeping database

- Job description metadata (job parameters, ...)
- Status and system metadata
- Back-end for all the framework sub-services
- Jobs running on remote sites communicate to the database via a RESTful interface (GSI proxy + SSL)
- Database schema **partially** configurable

Bookkeeping database



Bookkeeping database

- At present, relational data model is in use
- MySQL 5.0.77 (SL5 default)
- Upgrading test to 5.5.10 in next weeks
- Evaluation of No-SQL databases for specific purposes in progress (no plans to abandon relational model)
- rDBMS agnostic design – both schema and interface (tests with PostgreSQL in progress)

Grid services

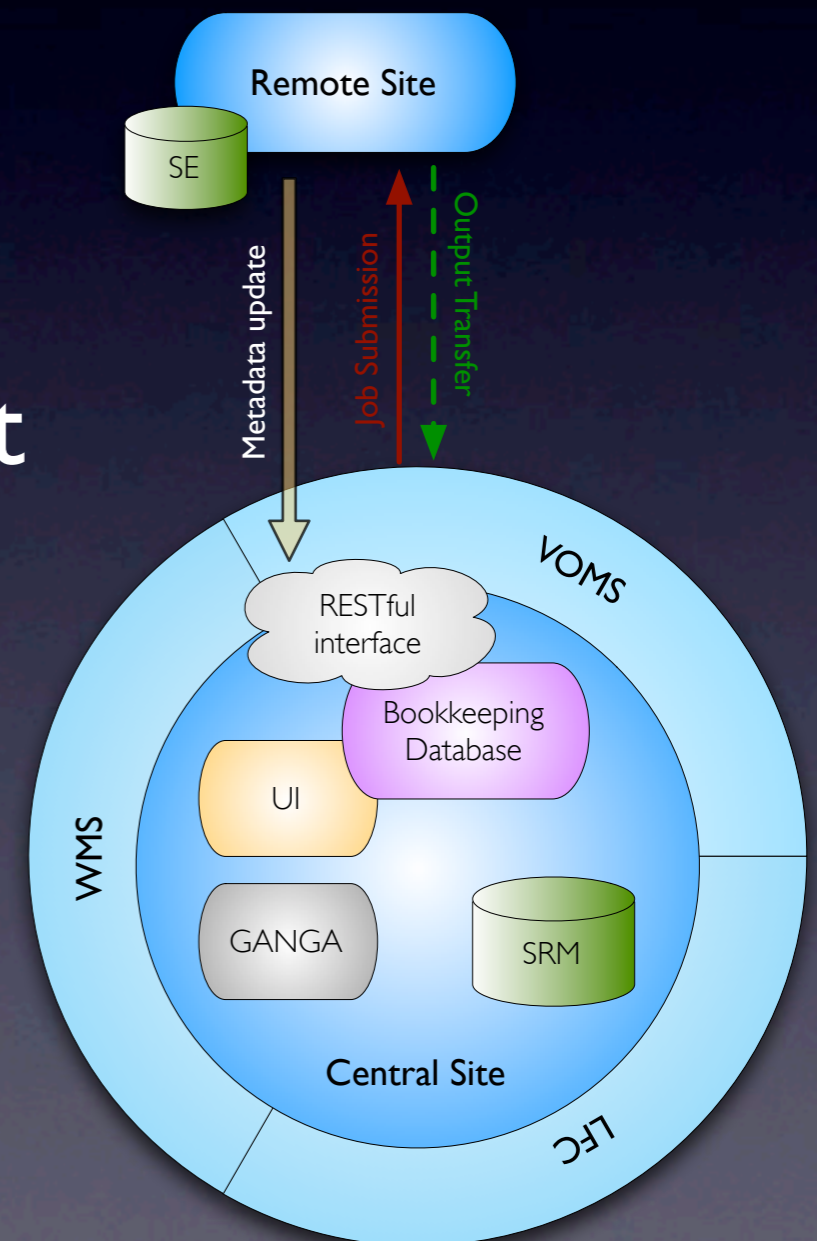
- **Workload Management System (WMS)**
Job brokering service and middleware **interoperability** element
- **Virtual Organization Membership Service (VOMS)**
Authentication, Authorization and Accounting system
- **LHC Computing Grid File Catalog (LFC, LCG-Utils)**
File metadata catalog and data-handling solution
- **GANGA**
Job management system (submissions)
- **Storage Resource Manager (SRM V2.2)**
Data access protocol/layer

Grid services ↔ conf

- The **single head node** requires: gLite UI, Framework installation (**GANGA, MySQL, Apache**)
- LFC server: VO shared, standard, per site name space
- VOMS: default conf. (Production | Software manager, standard user)
- WMS: CE-queue filter and middleware interop.
- LCG-Utils: Input (site SE to WN) / Output (WN to Repository) transfers

Production Framework

- Central EGI service site
- Web Interface
- Job and metadata management tool for submissions
- Database system to store the distributed production metadata (**bookkeeping**)

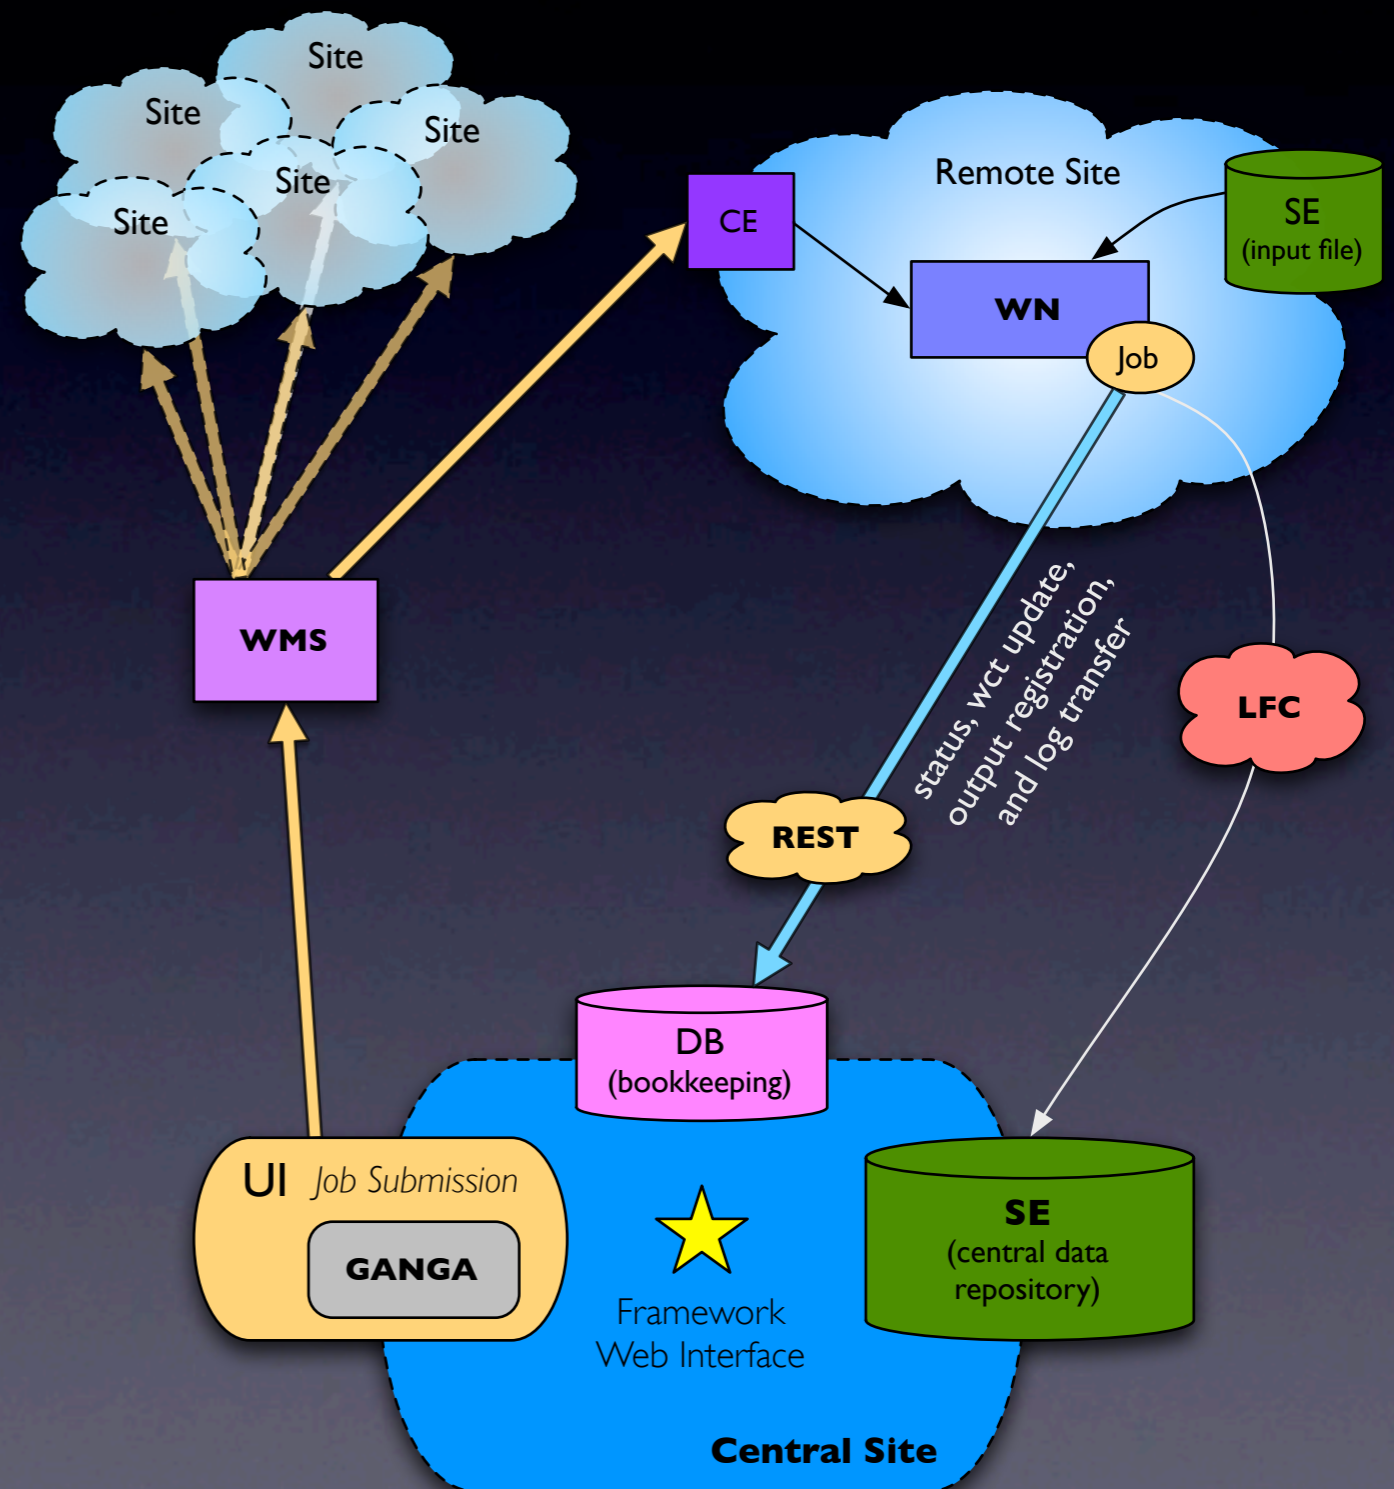


Job structure

- Job for the Grid is a script (Bash or python)
- It takes care of environment preparation, database communications, file transfers, failover conditions, ...
- VO executable launched within the script
 - FastSim (PacProduction) only at present
 - FullSim (Bruno) to be included

Job workflow

- **Pre-Production**
Input files and sw transfers to site SEs
- **Job preparation**
with the Framework Web Interface;
launch of submission script from UI
or AUTOMATIC from Portal (**NEW FEATURE**)
- **GANGA bulk submission** via WMS
- **Output file(s) transfer**
- **Metadata update**



Job script

- Failover chain on output transfers:
lcg-cr → lcg-cp [rf] → lcg-cp --nobdii [rf]
→ globus-url-copy [rf]
- RESTful communication
client-side: curl, server-side: php layer
proxy authentication using mod_ssl and
mod_gridsite

```
curl -v --cert $X509_USER_PROXY --key $X509_USER_PROXY --capath /etc/grid-security/certificates  
https://bbr-serv09.cr.cnaif.infn.it:8080/restcert1/index.php
```

Web Portal

- Production initialization

Setup for a new production,
bookkeeping database initialization

- Monitor system

Parametric search engine on job metadata,
logfile analysis sub-system, job status, job analysis tool, ...

- Report generator

Production report generator with summaries
on parameters, site usage and status graphs, ...

- Submission interface for shift takers (parametric)

automatic submission to all
the available sites,
automatic request selection
based on progress

- Submission interface for experts (custom)

Fine-grained selection
of job submission
parameters

- Request definition

Interface for specifying
new production requests
(set of job parameters)

- eLog system

Web Portal

PRODUCTION FORM

General Submit

Production Series: 2010_September Submit to checked sites Show all requests Refresh

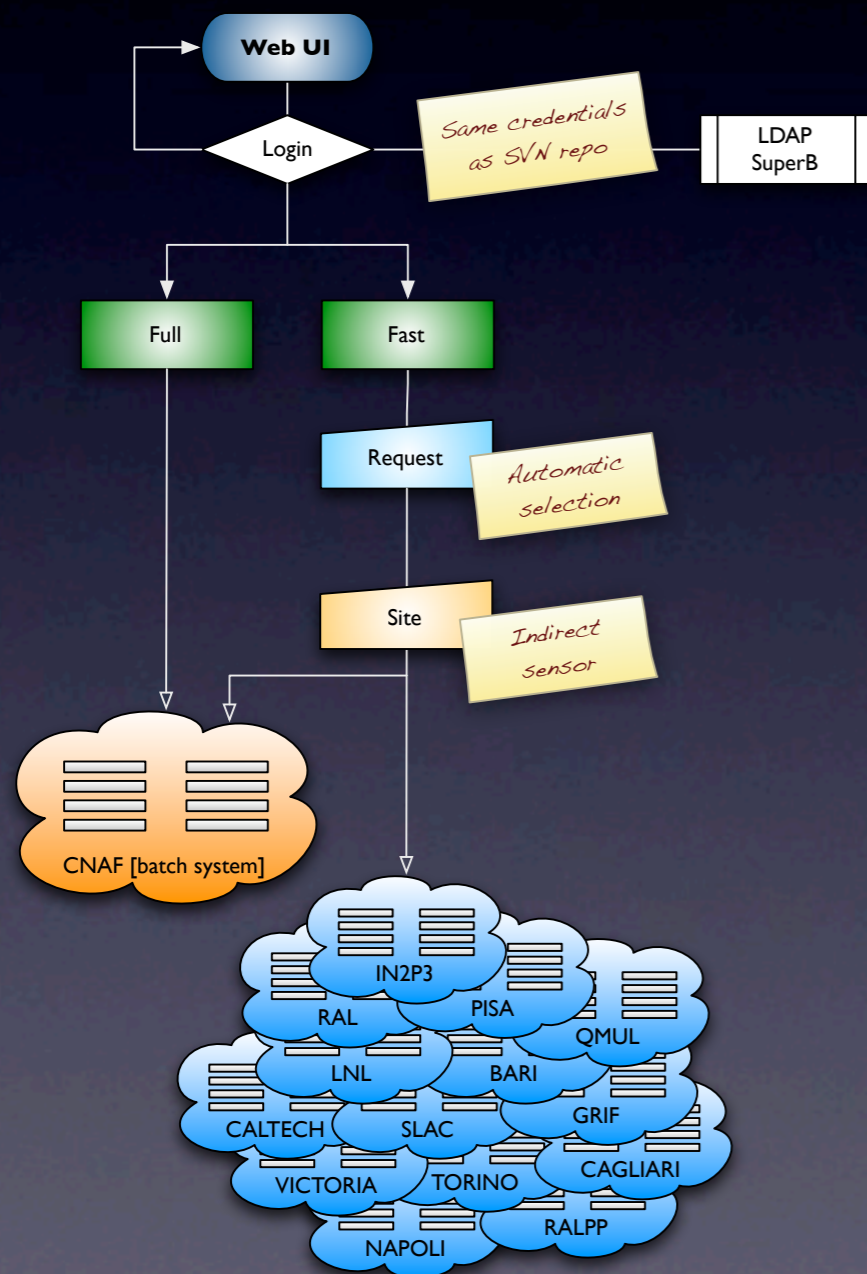
[+ show/hide details](#)

JOB DETAILS (REQUESTS)

| S | Analysis | Job Parameters | Graph Pending - Running - Failed - Done [TOT] | events per job | jobs d/nf/r | events d/nf/r | |
|---|---------------------------------------|--|--|----------------------|---|---|--|
| - | HadRecoilCocktail | | | | | | |
| ✓ | HadRecoilCocktail | • DG_4 • B0B0bar_Btag-HD_Cocktail • Bkg_NoPair.tcl | 0 - 0 - 488 - 4717 [5205] | 80000 | 4 717 (100.4%) 4 717 (100.4%) 4 700 | 377 360 000 (100.4%) 377 360 000 (100.4%) 376 000 000 | |
| ✓ | HadRecoilCocktail | • DG_4 • B+B-_Btag-HD_Cocktail • Bkg_NoPair.tcl | 0 - 0 - 1007 - 5918 [6925] | 80000 | 5 918 (100.3%) 5 918 (100.3%) 5 900 | 473 440 000 (100.3%) 473 440 000 (100.3%) 472 000 000 | |
| + | SLRecoilCocktail | | | | | | |
| + | Generics | | | | | | |
| + | KplusNuNu | | | | | | |
| + | BtoKNuNu | | | | | | |
| + | BtoKstarNuNu | | | | | | |
| Name | Last Activity | Last Submission | Load Pending - Running - Failed - Done [TOT] | Last Sub. P-R-F-D | Threshold P-F-R | Status | Action |
| <input checked="" type="checkbox"/> CIT_CMS_T2B | 2010-10-03 00:11:38 [568:14:09] | 2010-10-02 18:56:47 [573:29:00] | 0 - 0 - 101 - 124 [225] | 0-0-0-50 | [30] [10] [350] | 🟢 [50] | Submit Test |
| Reset GRIF | 2010-10-04 16:58:00 [527:27:47] | 2010-10-01 18:11:16 [598:14:31] | 0 - 0 - 345 - 3355 [3700] | 0-0-37-13 | [30] [10] [551] | 🟢 [50] | Test |
| <input checked="" type="checkbox"/> IN2P3-CC | 2010-10-06 09:19:14 [487:06:33] | 2010-10-02 20:18:44 [572:07:03] | 0 - 0 - 707 - 10143 [10850] | 0-0-2-98 | [50] [10] [700] | 🟢 [100] | Submit Test |

| ALL SITES | | |
|--------------|----------------|-----------------------|
| Status | # of jobs | events |
| done | 161 671 | 10 292 680 000 |
| failed | 4 171 | 285 730 000 |
| running | 3 | 210 000 |
| sys-failed | 11 391 | 790 780 000 |
| timeout | 136 | 9 500 000 |
| Total | 177 372 | 11 378 900 000 |

Submissions

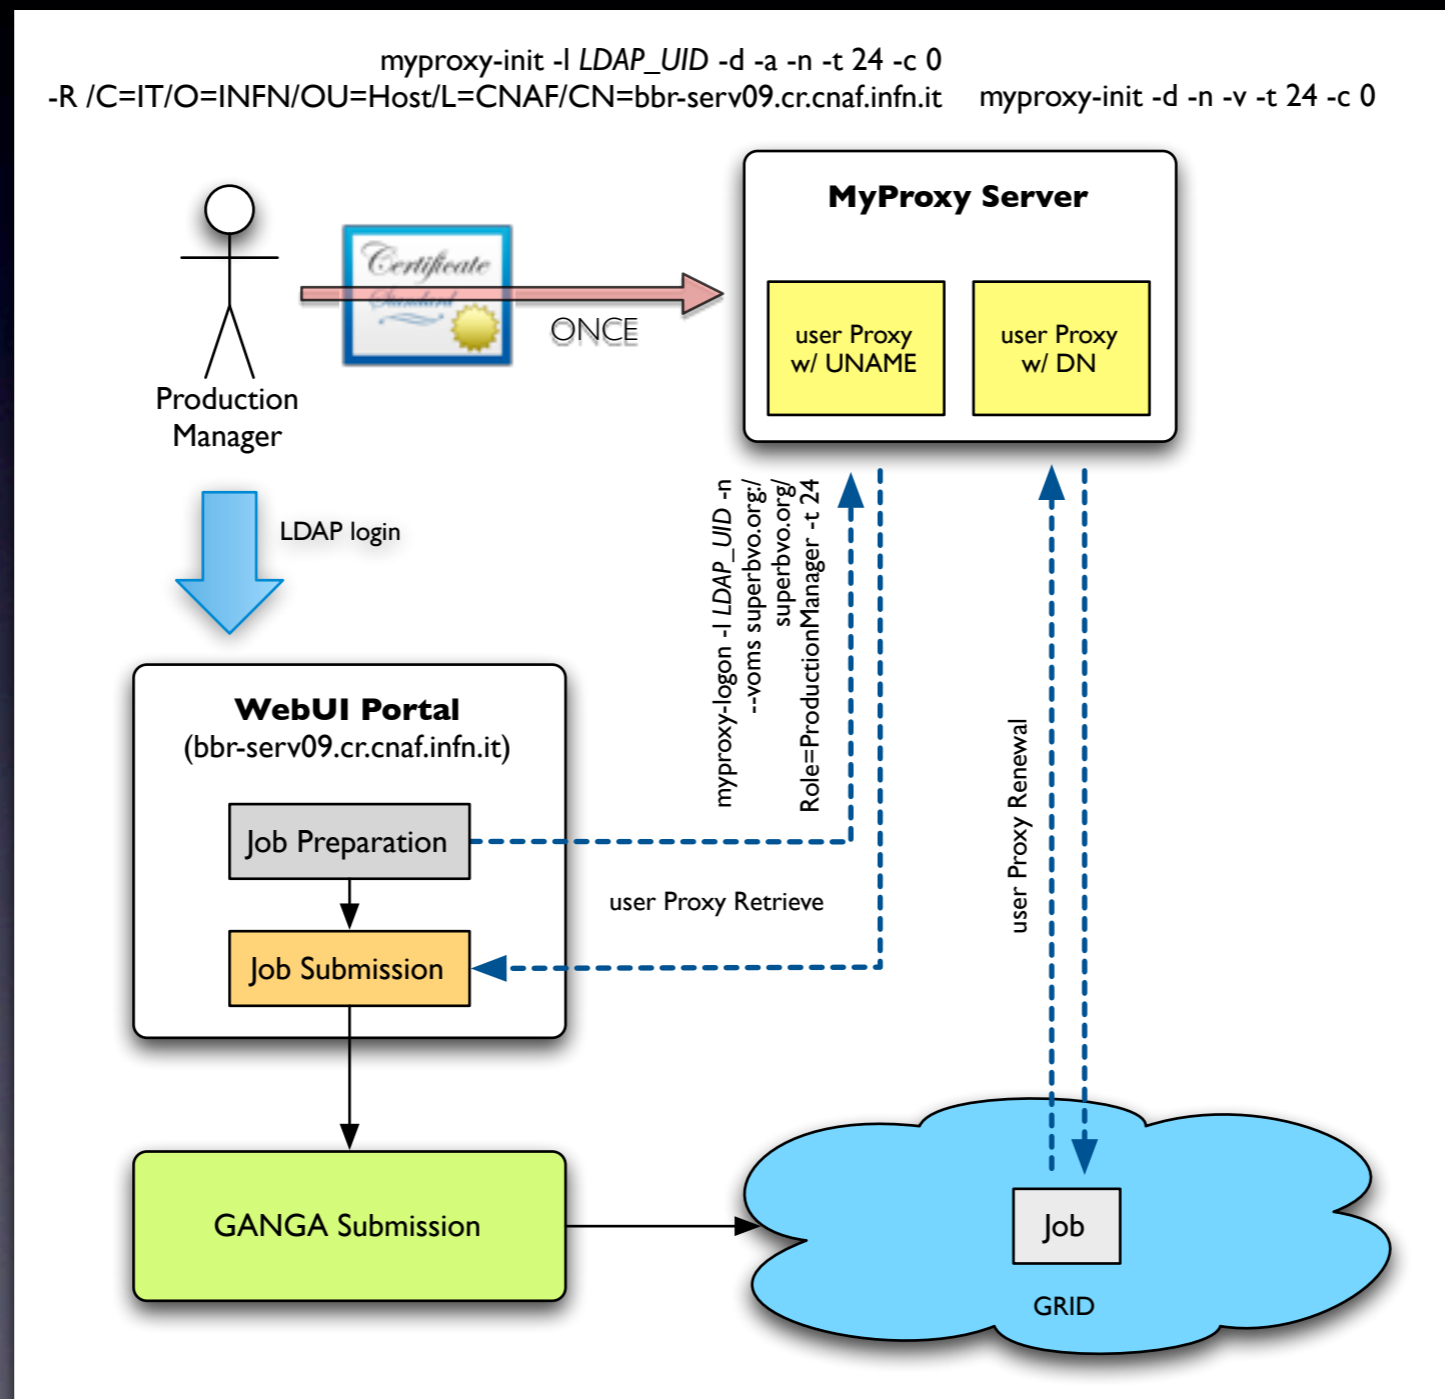


- Productions managed with the customized Framework
- Set of users with valid X509 certificates and grid-proxies
- Automatic preparation of job scripts, configuration files and bookkeeping entries
- Manual submission (**AUTOMATIC** submission via myProxy)
- INFN Tier-I site hosts the head node

Automatic Submission

- issue: proxy certificate must be personal; apache (httpd user) should get (securely) a user Proxy of the submitting user
- MyProxy Delegation feature in addition to Renewal already in use
- Each user (VOMS Production Manager role) must store into MyProxy server his Proxy in two ways (until personal certificate expiration):
 - with a Username (== LDAP uid), w/o password, retrievable by bbr-serv09 only (apache user) [delegation]
 - with DN, w/o password, retrievable by his own (expiring) proxy [renewal]
- apache user can retrieve the proxy stored in MyProxy by using his LDAP uid and then execute the submission scripts using the retrieved user proxy - user proxy is then destroyed

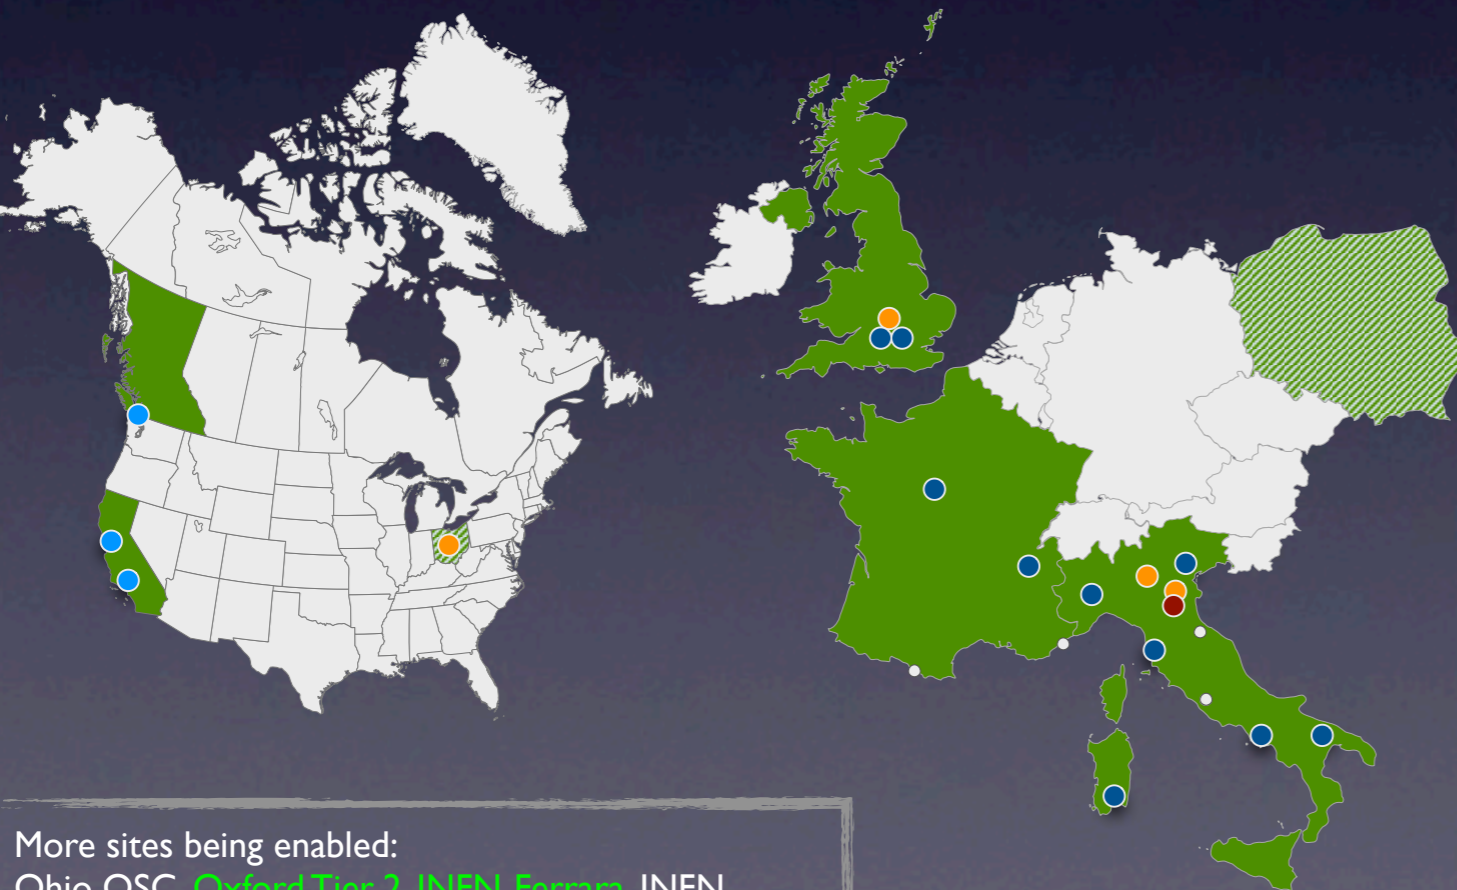
Automatic Submission



Remote sites

- 15 (+5) sites, two 'flavors' (EGI, OSG)

- Tier-1
INFN-T1 (Bologna, Italy),
IN2P3-CC (Lyon, France),
RAL-LCG2 (Oxford, UK)
- Tier-2
UKI-LT2-QMUL (London, UK),
UKI-SOUTHGRID-RALPP (London, UK),
UKI-SOUTHGRID-OX-HEP (Oxford, UK),
GRIF (Orsay, France),
SLAC (Stanford, USA),
CIT-CMS-T2B / CIT-HEP-CE (Caltech, USA),
OSC (Ohio, USA),
VICTORIA-LCG2 (Victoria, Canada),
INFN-BARI (Bari, Italy),
INFN-CAGLIARI (Cagliari, Italy),
INFN-FERRARA (Ferrara, Italy),
INFN-LNL-2 (Legnaro, Italy),
INFN-MILANO (Milano, Italy),
INFN-NAPOLI (Napoli, Italy),
INFN-PISA (Pisa, Italy),
INFN-TORINO (Torino, Italy)
- **PL-GRID** (Poland)



More sites being enabled:
Ohio OSC, **Oxford Tier-2**, **INFN-Ferrara**, INFN-
Milano, PL-Grid, ...

Results

- September 2010: large production cycle
- $> 11 \times 10^9$ events
(about 180 kjobs in 4 weeks)
- Peak rate of 7000 simultaneous jobs running on the Grid (3.5k on average)
- about 8% failure rate
executable errors (0.5%), site misconfiguration (2.0%), proxy expiration (4.0%), central site overload during file transfers (1.5%)
- about 195 years of WCT

Latest Development

- Automatic Submission (to the Grid)
- Bookkeeping database tuning:
use of triggers to speed-up statistic calculations
- Use of template engine for code/layout separation (in progress)
- JQuery upgrade, datatable plug-in, updated Google API for charting
- Python script, v2 and v3 (testing phase)
- Target site per request (replica at CNAF)

Future Work

- Bookkeeping database replication (beginning with master/slave configuration)
- Monitor integration with LB service (through direct queries) both on demand and with a daemon
- Integration of packaged FullSim
- Enabling FullSim request interface + submission to Grid

Planning

- Do we have new feature requests for the production tools?
- Next FastSim productions...
- Next FullSim productions...
[do we really need Grid exploitation at present?]

Conclusions

- The SuperB productions proved that the prototype framework is **reliable and efficient**
- Customization of the Web interface, the bookkeeping database, and the site requirements have been the key points to achieve the goal

Collaboration

- A. Fella
- E. Luppi
- M. Manzali
- L. Tomassetti
- students from Ferrara
- M. Favaro
- G. Fontana
- C. Luzzi
- M. Ronzano
- L. Vettorello
- E. Vianello

Thank You