

Distributed resource and production system updates

Armando Fella & Cristian De Santis

VI SuperB Collaboration Meeting
Frascati, 11-14 December 2012



Presentation Outline

- Distributed resources and management status
- Production system update
- Future work

Distributed resource status

- Full distributed resource testing phase is in progress
 - ~1/4 of the sites results misconfigured
 - Testing, fixing, reinstalling, contacting is slowly going on
- VO enabling operations at remote sites almost completed
 - McGill resources fully configured
 - Caltech site tuning is in progress
 - SLAC dedicated resources no more accessible due to SuperB cancellation
 - See resource summary in backup slides [**]
- Resource management system moved to standard EGI procedures and frameworks
 - Via SuperB Grid resource dashboard
 - EGI global resource: <https://operations-portal.egi.eu/voDashboard?vo=superbvo.org>
 - OSG: <https://jira.opensciencegrid.org/secure/Dashboard.jspa>
 - GGUS ticket system
 - <https://ggus.eu/pages/support.php> (superb-vo-admin@lists.infn.it related SuperB emailing list)

SuperB dashboard

Welcome Ammarah Fella



Superbvo.org

EVERYTHING >> MONITORING > HISTORY > CRITICAL

ADMINISTRATION

EVENTS

TICKETS

OPERATIONS PORTAL TOOL



<input type="checkbox"/> NGL_IT	37
<input type="checkbox"/> OSG	48
<input type="checkbox"/> NGL_UK	15
<input type="checkbox"/> ROC_Canada	7
<input type="checkbox"/> NGL_FRANCE	6
<input type="checkbox"/> NGL_PL	2

Flapping detection

4 status changes during last 48 hours



monitoring > ngl (6)

NGI	WARNING	CRITICAL	UNKNOWN	OPENED GDS TICKETS	OPENED HOTPANS
<input type="checkbox"/> NGL_FRANCE					

monitoring > site (4)

SITE	NGI	WARNING	CRITICAL	UNKNOWN	OPENED GDS TICKETS	OPENED HOTPANS
<input type="checkbox"/> UGP2-BES	NGL_FRANCE					
<input type="checkbox"/> UGP2-CC	NGL_FRANCE					
<input type="checkbox"/> GEF	NGL_FRANCE					
<input type="checkbox"/> UGP2-LPC	NGL_FRANCE					

- GDS Site Information
- Add a GDS Ticket
- Add a hotpan

<input type="checkbox"/> NGL_IT						
<input type="checkbox"/> NGL_PL						
<input type="checkbox"/> NGL_UK						

monitoring > site (4)

SITE	NGI	WARNING	CRITICAL	UNKNOWN	OPENED GDS TICKETS	OPENED HOTPANS
<input type="checkbox"/> UMI-SOUTHGRID-SMPP	NGL_UK					
<input type="checkbox"/> BAL-UG2	NGL_UK					
<input type="checkbox"/> UMI-SOUTHGRID-OC-HEP	NGL_UK					
<input type="checkbox"/> UMI-IT2-QMUL	NGL_UK					

<input type="checkbox"/> OSG						
<input type="checkbox"/> ROC_Canada						

SuperB software management system in distributed environment: CVMFS

- CVMFS in production at CNAF since November 1st
 - Few problems detected and solved during the very first days
 - VO_SUPERBVO_ORG_SW_DIR=/cvmfs/superbvo.org/
 - IGI + SuperB (P. Veronesi + P.Franchini)
- CVMFS setup for test purpose at Pisa and Victoria
- Ashok Agarval from UVIC joined SuperB
 - CVMFS R&D work started at Victoria University
 - CVMFS client installation on Nimbus Cloud[*] cluster
- Setup of a geographical redundancy of the service is a constraint to move to production state --> Torino 1q 2013
- The CVMFS installation campaign on distributed resources may be able to start on January 2013

- [*]<http://www.nimbusproject.org/> Nimbus Infrastructure is an open source EC2/S3-compatible Infrastructure-as-a-Service implementation specifically targeting features of interest to the scientific community such as support for proxy credentials, batch schedulers, best-effort allocations and others.

CNAF resource and service status

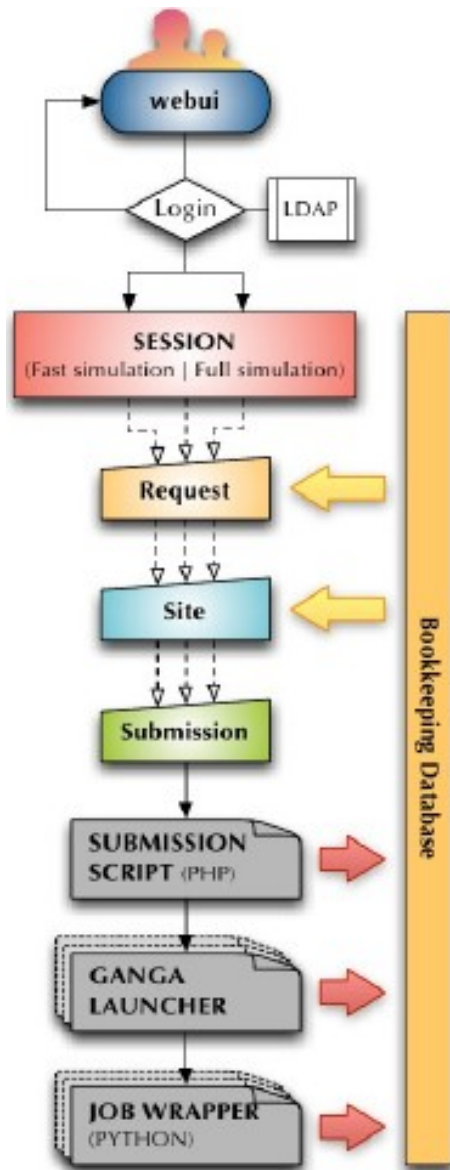
- The assigned disk space for 2012 is 200TB, currently 127TB have been allocated
 - **73TB will be added within December**
 - **AMS will use SuperB disk space (~8TB) until end of December**
 - **disk assignment to CNAF served VOs is in progress**
- **Four new blade machines** will be added to the SuperB machine set in one month, courtesy of CNAF.
 - Postponed to end of December
- SuperB will use the new machines for the following services:
 - **DB clusterization**
 - **Dirac system**
 - **Production head node replacement**

[**]Resource summary

Site	Min (cores)	Max (cores)	Disk (TB)	SRM layer	Grid Org.	Site contacts
RAL(T1)	200	1000	25	Castor	EGI	F. Wilson, C. Brew
Ralpp	50	500	5	dCache	EGI	F. Wilson, C. Brew
Queen Mary	300	2000	150	StoRM	EGI	A. Martin, C. Walker
Oxford Univ.	50	200	1	DPM	EGI	K. Mohammad, E. MacMahon
IN2P3-CC(T1)	500	1000	16	dCache	EGI	N. Arnaud, O. Dadoun
Grif	50	300	2	DPM	EGI	N. Arnaud, O. Dadoun
in2p3-lpsc	50	100	2	DPM	EGI	J.S. Real
in2p3-ires	50	100	2	DPM	EGI	Y. Patois
CNAF(T1)	500	1000	180	StoRM	EGI	A. Fella, P. Franchini
Pisa	50	500	0.5	StoRM	EGI	A. Ciampa, E. Mazzoni, D. Fabiani
Legnaro	50	100	1	StoRM	EGI	G. Maron, A. Crescente, S. Fantinel
Napoli	500	2000	15	DPM	EGI	S. Pardi, A. Doria
Bari	160	260	0.5	StoRM/Lustre	EGI	G. Donvito, V. Spinoso
Ferrara	10	50	0.5	StoRM	EGI	L. Tomassetti, A. Donati
Cagliari	10	50	1	StoRM	EGI	D. Mura
Perugia	10	50	1	StoRM	EGI	L. Fano'
Torino	50	100	2	DPM	EGI	S. Bagnasco, R. Brunetti
Frascati	30	100	2	DPM	EGI	E. Vilucchi, G. Fortugno, A. Martini
Milano	50	100	2	StoRM	EGI	N. Neri, L. Vaccarossa, D. Rebatto
Catania*	?	?	?	StoRM	EGI	G. Platania
Slac	400	400	10	NFS	OSG	S. Luiz, W. Yang
Caltech*	200	400	4.5	NFS	OSG	D. Kcira
Fnal	50	400	1	dCache	OSG	M. Slyz
Victoria	50	100	5	dCache	EGI	A. Agarwal
McGill	100	200	1	StoRM	EGI	S. Robertson, S.K. Nderitu
Cyfronet	100	500	10	DPM	EGI	L. Flis, T. Szeplenie, J. Chwastowski
Total	3570	11510	440			

* VO enabling procedure in progress

Production System



- Relies on *sbk5* PostgreSQL book-keeping database
- New WebUI version will be available at: https://bbr-serv09.cr.cnaf.infn.it:8443/~webui_sbk5
 - access distributed resources for fastsim e fullsim
 - new functionalities (production management, job re-submission)
- New job wrapper features

Production system update – BK DB (1)

Book-keeping database (*sbk5*) modified to implement new production system features on FullSim session:

- **Request:** new request and production request definition [TESTED]
 - ▶ Number of events per job removed from production request and included in request parameters (*hstore* field) used to generate request name
 - ▶ Requests having same physical parameters can differ if the number of events per job is different
 - ▶ Old production request names re-defined accordingly
 - ▶ Old productions having two different number of events per job split into two new productions

Production system update – BK DB (2)

- **Jobs:** new fields and constraints to implement functionalities to re-submit failed job in bulk mode (**job re-submission**) [ONGOING]
 - ▶ Implemented one-to-one correspondence between job and input file to re-submit failed job n times
 - ▶ Introduced new fields:
 - `input_id`: realizes a one-to-one correspondence between job and input file
 - `resub_id`: counts number of resubmission for a give job (0 for first submission)
 - ▶ Re-submitted jobs will have same `prod_series`, `request_name`, `input_id`
 - ▶ Same PK (`pros_series`, `runnum`)
 - ▶ Drawbacks: now resubmitted jobs will have different `runnum` which is used as seed for simulations (in future an independent seed will be defined)

Production system update – Job wrapper

Changes to improve performances and implement re-submission mechanism:

- **Submission:** changes to implement **job re-submission** functionality [ONGOING]
 - ▶ New submission script which will contain information concerning input ids
 - ▶ New configuration script which will contain information concerning parametric input file submission
- **Stage-in:** implemented new fail-over/optimization mechanism based on site access time scoring for the stage-in phase [TO BE TESTED]

Production system update – Web-UI (1)

Web-UI code (PHP, JavaScript, Smarty) modified in order to handle new features and *sbk5* updates on FullSim session

A lot of debugging work carried out by means of FireBug tool

Fixed minor bugs and improved exception handling

In detail:

- **Production:** new production creation interface with a robust software version and parameters management [TESTED]
- **Production Requests:**
 - ▶ possibility to create different production requests having the same physical parameters by changing the number of events per job [TESTED]
 - ▶ Simplification: disabled automatic “Request Settings” (required events, number of jobs) consistency check
 - ▶ new parametric input mode (PARFILE) for *job re-submission* [ONGOING]

Production system update – Web-UI (2)

- **Expert Submission Interface:**

- ▶ New “Re-submission Mode” checkbox [TO BE TESTED]
- ▶ “Job Details” form section updated keeping previous functionality for production initialization [TESTED]
- ▶ Improved Ajax exception handling
- ▶ New job wrapper configuration file creation to handle re-submission mode [TO BE COMPLETED]

- **Shift Submission Interface:** minor changes [TESTED]

- **Job Monitor:** minor changes [TESTED]

- **Submission Monitor:** minor changes [TESTED]

Production interface

- A production can be created by means of new form which allows to insert a complete and well formed tuple in Soft_Ref table
- The user must specify the release working directory (e.g. /storage/gpfs_superb/prod/test_releases_local/build/sl53_x86_64_FullSim/SuperB/FullSim/<release_tag>
- The system:
 - validates the directory before running the "svn info" command
 - executes the "svnversion" and "svn info" to obtain version and revision
 - uses version and revision to complete Soft_Ref values.
[soft_release, soft_name, soft_version, soft_revision, soft_lfn, exe_path, working_dir]

SoftRef upgrade

- New Session definition:

```
array(2) {
```

```
  ["FullSim"]=>array(9) {
```

```
    ["soft_name"]=>string(5) "Bruno"
```

```
    ...
```

```
    ["soft_lfn"]=>string(96) "lfn:/grid/superbvo.org/production/FullSim/test_release/<version>_SL-5.3_x86_64_rev<revision>.tjz" } }
```

- New methods in Soft_Ref class:
 - public static function validateWorkDir(\$work_dir)
 - check the structure of \$work_dir
 - returns a boolean value.
 - public static function GetGeoDir(\$release_work_dir)
 - returns the name(s) of file(s) that match \$SoftRefSession['GeoName'] in an array
 - ...

Future Work

- Finalizing and testing the new re-submission mechanism (Jan '13)
- Completing distributed resources reviewing (Jan '13)
- Porting of the production system head node on new machines @ CNAF (Jan '13)
- HA clustering solution for book-keeping database (Feb '13)
- Documentation writing
- Credits: C.De Santis, A.Fella, E.Luppi, M.Manzali, A.Mesin, L.Tomassetti