

Production Tools for Shift Takers

Luca Tomassetti
University of Ferrara & INFN

SuperB Distributed Computing Group



Foreword

- The distributed computing development team is almost hopelessly undermanned:
 - E. Luppi and L. Tomassetti have heavy teaching duties until end of June (since March 2010)
 - A. Fella is without contract (since 31/03/2010)
 - M. Manzali joined the group recently (still under training)
 - Students from Ferrara almost gone:
 - C. Luzzi, G. Fontana graduated in Dec. 2009 and March 2010
 - M. Ronzano going to graduate in July 2010

Foreword

- The Production Tools for Shifts is not ready yet

It has been fully designed

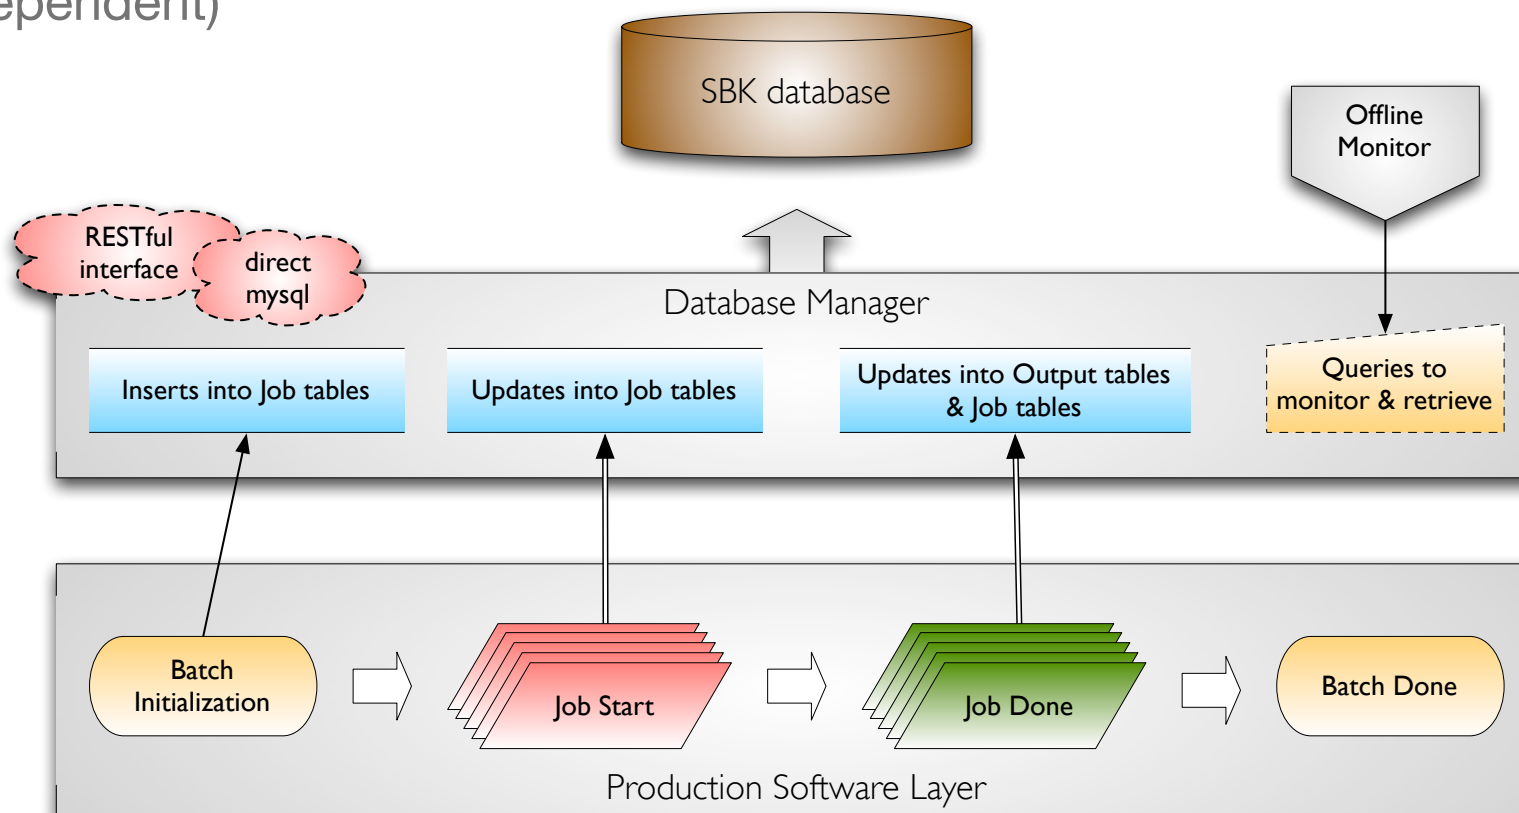
It is under development

It will be fully tested after the FastSim release

Summary

- Production Tools:
 - WebUI for experts / for shift takers
 - User's Production requests
 - Site availability sensors
 - Job preparation
 - Job submission
 - Tests and elog
- Shift takers:
 - requirements
 - duties
 - training
- Tentative Schedule

- **Web based User Interface**
- PHP, javascript, JQuery framework for AJAX functionalities
- database classes for communication with bookkeeping database (RDBMS independent)



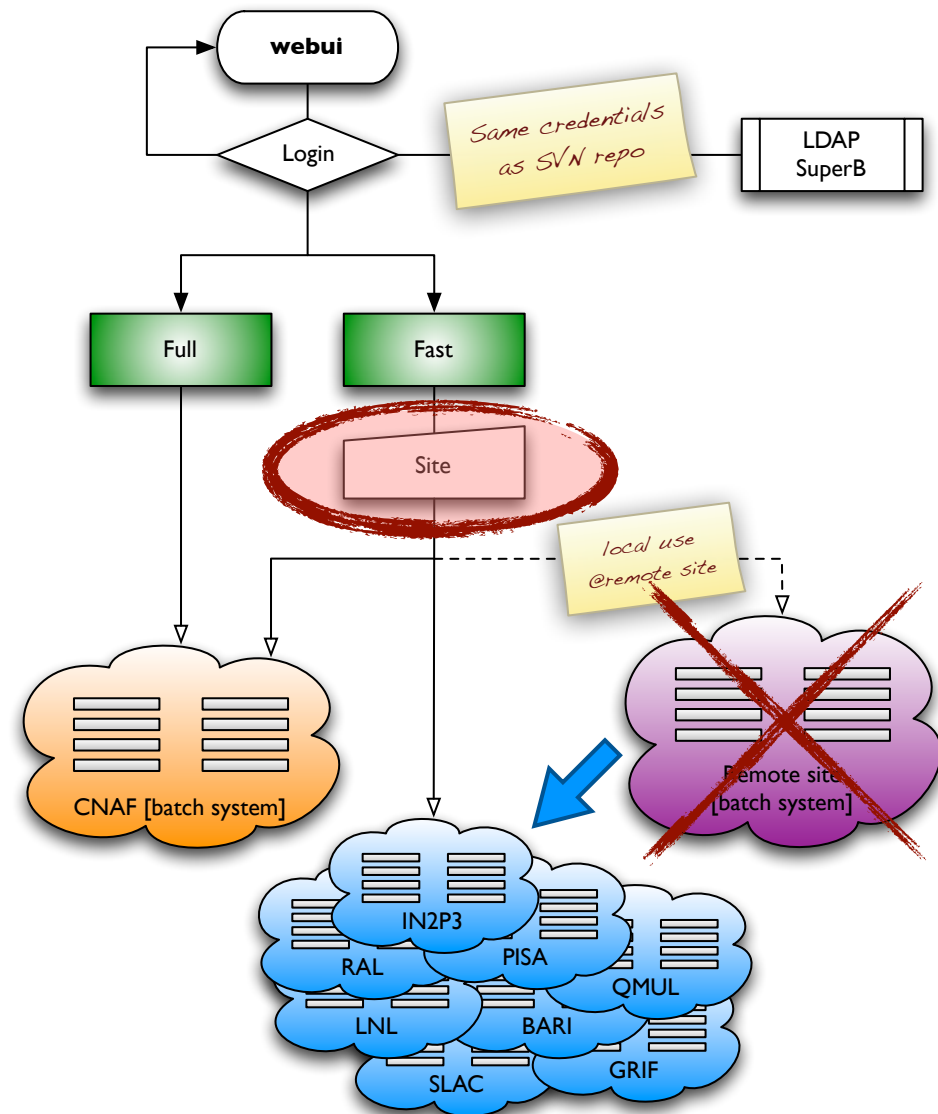
- **Production Submission Workflow**

- basic authentication and authorization based on common LDAP directory service

- FullSim / FastSim differentiation

- Batch System / Distributed (Grid) submission

- Site selection will be automatically managed by the shift interface



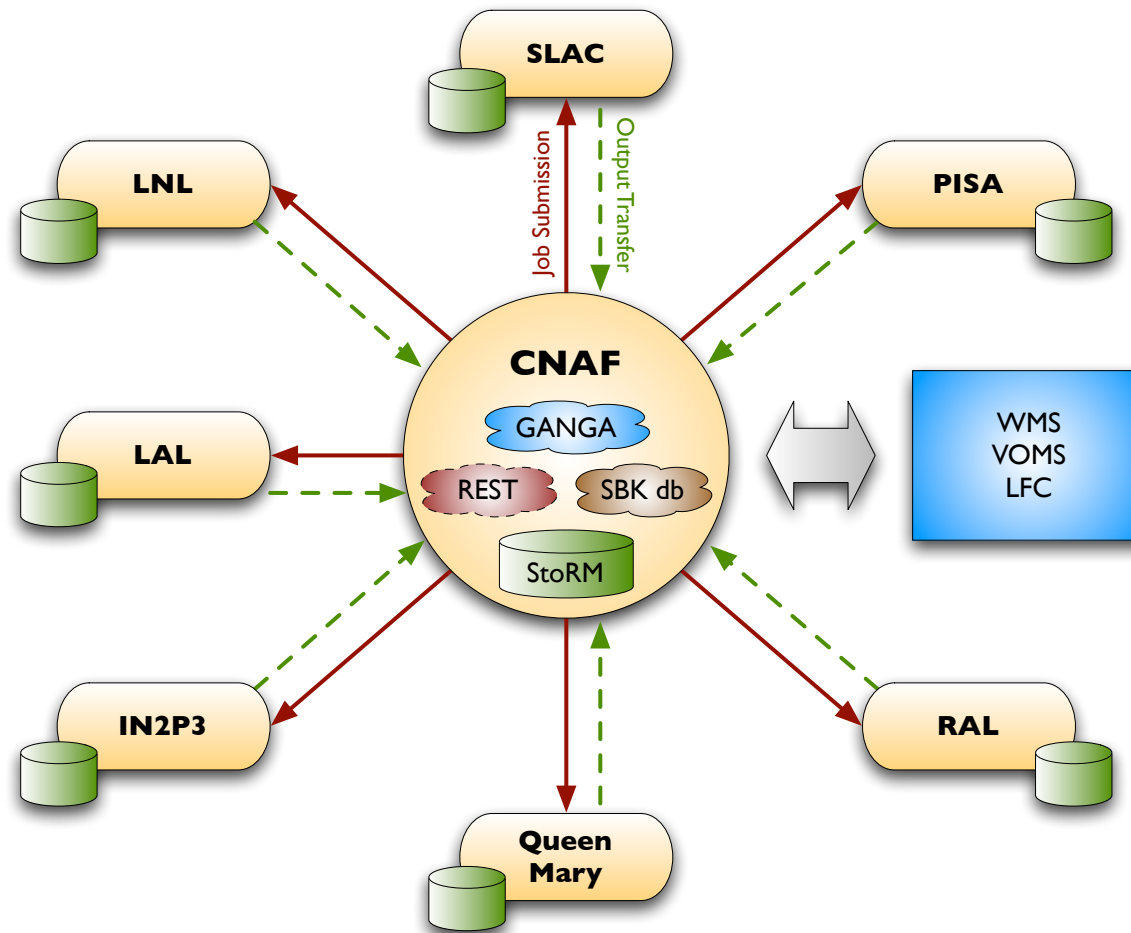
- **FastSim**

- generator, geometry, bkg mixing files (tcl, dec, ...) must be in the software
- Production root ⇒ permissions
 - local user + 'grid' user

The screenshot displays the FastSim WebUI interface. At the top, there are three dropdown menus for background selection: 'No Backgrounds', 'All Backgrounds', and 'All Bkg, No Pair' (selected). Below these are two more dropdown menus: 'GRIF', 'GRISU-UNINA', 'IN2P3-CC' (selected), 'INFN-BARI', 'INFN-FERRARA', 'INFN-LNL-2', 'INFN-PERUGIA', 'INFN-PISA', 'INFN-T1', 'SLAC', 'UKI-LT2-QMUL', 'UKI-SOUTHGRID-RALPP', and 'UNINA-EGEE'. In the center, a large list of generators is shown, with 'BOB0bar_generic' selected. To its right, a list of geometries is shown, with 'DG_3' selected. At the bottom, the 'Submission mode' section shows 'Submission Type: Grid' and 'Submission Site: IN2P3-CC'. The 'Job Details' section shows '# of Runs: 200', '# of Events (per job): 10000', 'Generator: BOB0bar_gen', 'Geometry: DG_3', and 'Bkg: All Bkg, No P'.

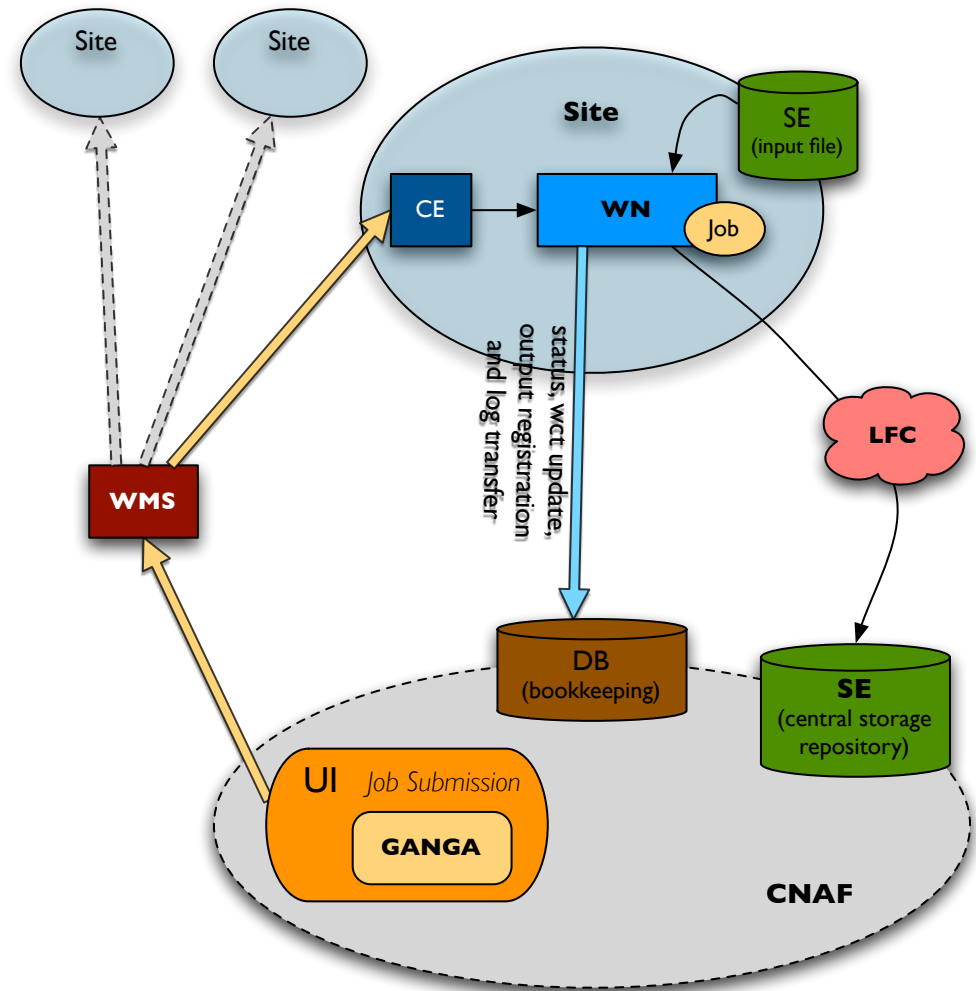
Production tools

- **FastSim Distributed**



Production tools

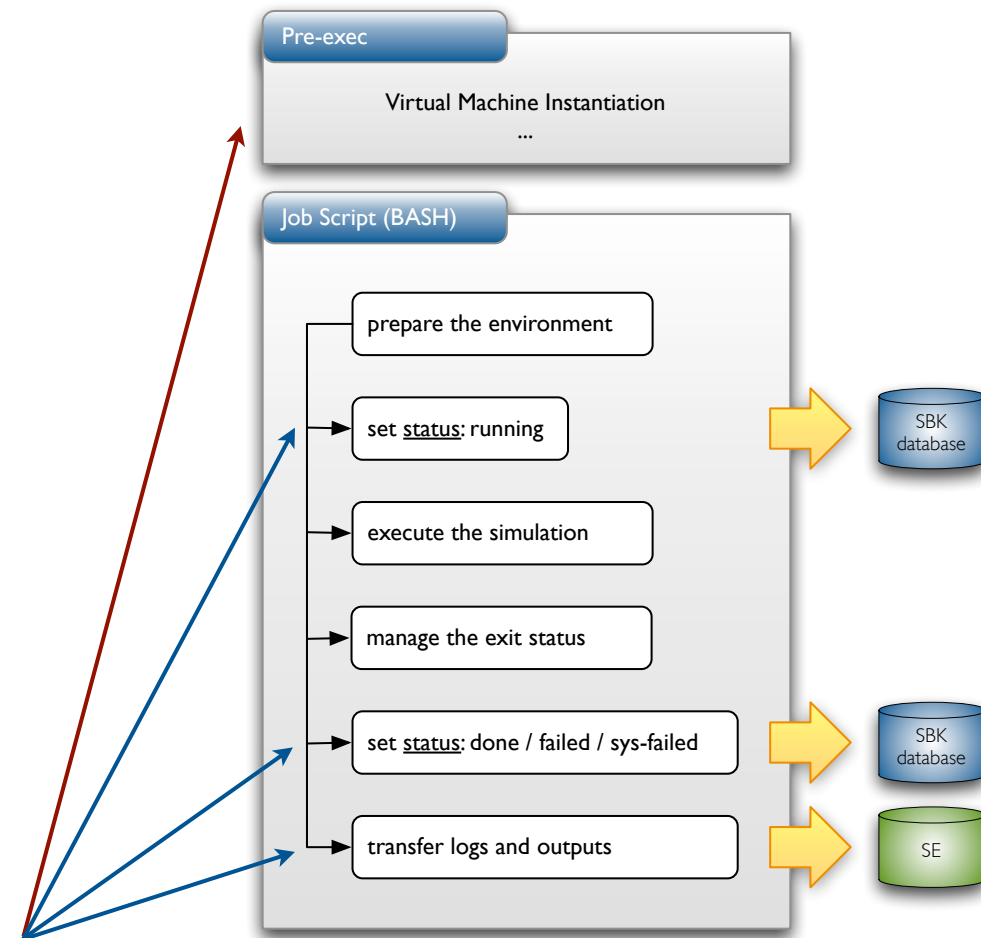
- **FastSim Distributed**
input files and test release have been transferred via LCG-Utils to remote SE
- Submission performed by GANGA
- WMS routes the job(s) to the remote site
- The job is scheduled by the remote CE to their WN



Production tools

Job scripts

- **bash script** which executes the simulation application
- prepared with proper parameters (from WebUI)
- base script (one per batch) with runnum as parameter
- environment + db interaction + exit status management + log and output transfer
- differentiation: LSF / Grid



- DB interface:
 - local jobs
direct mysql connection
 - distributed jobs
RESTfull interface



CURL
only update

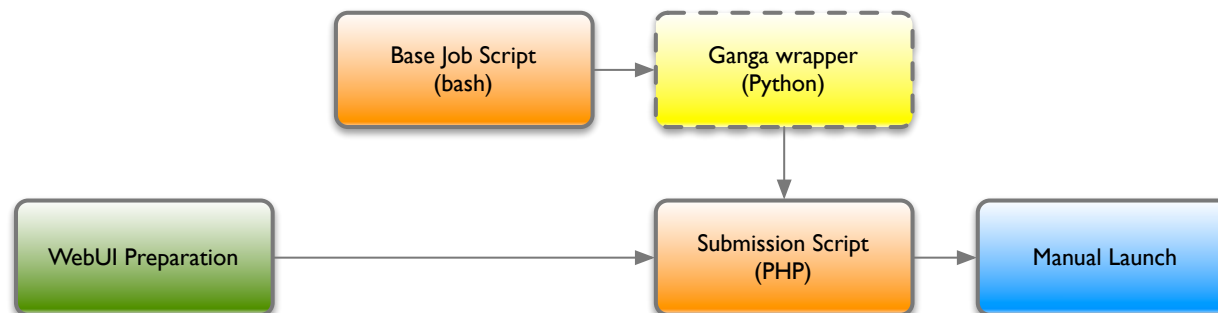
- File transfer:
 - local jobs
cp
 - distributed jobs
lcg-cr, lcg-cp, grid-ftp



site dependent
failover policy

Replica on
remote site

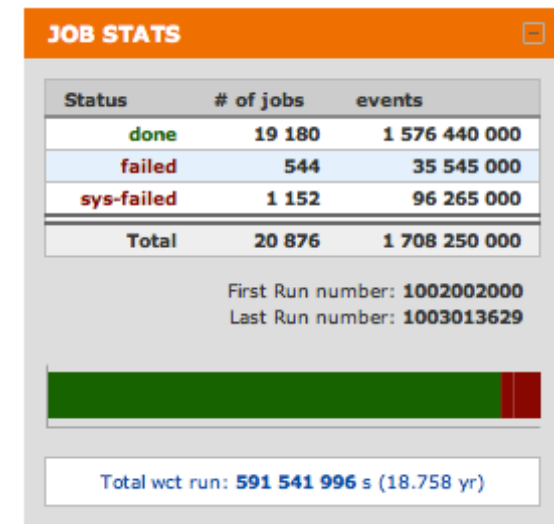
- Wrappers / launchers around the job scripts
- PHP *or* PHP + GANGA (Python) *or* PHP + low-level Grid commands
The one that the production user launches from CNAF (bbr-serv09)
- Allows multi-job / multi-set submissions
- Provides initialization (inserts) of the bookkeeping database
- LSF Batch system (PHP) script provides resubmit functionalities



Production tools

Monitor&Report

- Basic Monitor based on bookkeeping database
- Parametric list of jobs
- Detailed view of jobs
- Log file accessible from UI
- Check on output file existence



Runnum start end	# of events	Generator	Geometry	Soft release	Job Status Reason	(Grid) Job Id	Site	tcl	wct (s)
100200 100301		B0B0bar_gen	DG_3	ALL	done ALL		INFN-T1	MixSuperbBk	
259 jobs found matching the search criteria.									
1003013599	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072757	INFN-T1	MixSuperbBkg_NoPair.tcl	19213.24
1003013598	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072756	INFN-T1	MixSuperbBkg_NoPair.tcl	18830.00
1003013597	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072755	INFN-T1	MixSuperbBkg_NoPair.tcl	19634.58
1003013596	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072754	INFN-T1	MixSuperbBkg_NoPair.tcl	23934.02
1003013595	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072753	INFN-T1	MixSuperbBkg_NoPair.tcl	19002.69

Production tools

Monitor&Report

- Basic Monitor based on bookkeeping database
- Parametric list of jobs
- Detailed
- Log file a
- Check or

The screenshot displays a web-based monitoring interface. A central window titled "Job's Log" shows the contents of a log file for a specific job. The log file path is `/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/log/run1003011099.log`. The log content includes a header for "PacProductionApp" and a series of environment variable assignments and their values, such as `set RUNNUM 1003011099` and `set NEVENTS 10000`. A "JOB STATS" panel on the right shows a table of events with columns for "events" and "wct (s)". Below the log window, there are search filters and a table of job results.

Job's Log Content:

```

***** PacProductionApp ../PacProduction/MixSuperbBkg_NoPair.tcl
*****

set RUNNUM 1003011099           ;# using value from envvar
set NEVENTS 10000              ;# using value from envvar
set TIMER No                    ;# not set, using default
set DG DG_3                     ;# using value from envvar
set GENERATOR B0B0bar_generic   ;# using value from envvar
set PRODSERIES 2010_February_Generics ;# using value from envvar
set PRODTYPE FastSim           ;# not set, using default
set PRODROOT /scratch/supb05535369.ccw11127/tmp/gangajob_QlZ18494/sb_1267984408 ;#
using value from envvar
# UFILTER not set
set BKGROOT
/scratch/supb05535369.ccw11127/tmp/gangajob_QlZ18494/sb_1267984408/input_files ;#
using value from envvar
    
```

JOB STATS Panel:

events	wct (s)
1 576 440 000	
35 545 000	
96 265 000	
1 708 250 000	

number: 1002002000
number: 1003013629

996 s (18.758 yr)

Job Search Results:

259 jobs found matching the search criteria.

Runnum	# o	# o	# o	# o	# o	# o	# o	# o	# o	# o
start	end	start	end	start	end	start	end	start	end	start
100200										
100301										
1003013599	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072757	INFN-T1	MixSuperbBkg_NoPair.tcl	19213.24	
1003013598	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072756	INFN-T1	MixSuperbBkg_NoPair.tcl	18830.00	
1003013597	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072755	INFN-T1	MixSuperbBkg_NoPair.tcl	19634.58	
1003013596	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072754	INFN-T1	MixSuperbBkg_NoPair.tcl	23934.02	
1003013595	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072753	INFN-T1	MixSuperbBkg_NoPair.tcl	19002.69	

Production tools

Monitor&Report

- Basic Monitor based on bookkeeping database

- Parametric

- Detailed view

- Log file access

- Check on completion

Runnum	# of events
100200	
100301	

Job's Output

channel_type	location_path	file_name	size	
DstD0ToKspipi	/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/1003011099	DstD0ToKspipi.root	1.453 mb	✓
HadRecoil	/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/1003011099	HadRecoil.root	4.690 mb	✓
S2b	/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/1003011099	S2b.root	0.596 mb	✓
SemiLepKplusNuNu	/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/1003011099	SemiLepKplusNuNu.root	0.062 mb	✓
TwoBody	/storage/gpfs_superb/prod/2010_February_Generics/FastSim/DG_3/B0B0bar_generic/1003011099	TwoBody.root	0.081 mb	✓

obs	events
9 180	1 576 440 000
544	35 545 000
1 152	96 265 000
0 876	1 708 250 000

Run number: 1002002000
Run number: 1003013629

541 996 s (18.758 yr)

wct (s)

uperbBk

show output graph

259 jobs found matching the search criteria.

1003013599	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072757	INFN-T1	MixSuperbBkg_NoPair.tcl	19213.24
1003013598	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072756	INFN-T1	MixSuperbBkg_NoPair.tcl	18830.00
1003013597	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072755	INFN-T1	MixSuperbBkg_NoPair.tcl	19634.58
1003013596	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072754	INFN-T1	MixSuperbBkg_NoPair.tcl	23934.02
1003013595	10000	B0B0bar_generic	DG_3	V0.2.1 221	done	1072753	INFN-T1	MixSuperbBkg_NoPair.tcl	19002.69

Production tools

Monitor&Report

- Reports on sites, channels and last status changes (only FastSim)

GRIF			IN2P3-CC			INFN-BARI		
Status	# of jobs	events	Status	# of jobs	events	Status	# of jobs	events
done	2 021	177 650 000	done	4 457	383 930 000	done	29	2 850 000
failed	41	8 200 000	failed	254	15 930 000	failed	67	3 350 000
sys-failed	28	650 000	sys-failed	39	2 640 000	sys-failed	404	31 300 000
Total	2 090	186 500 000	Total	4 750	402 500 000	Total	500	37 500 000

INFN-LNL-2			INFN-PISA			INFN-T1		
Status	# of jobs	events	Status	# of jobs	events	Status	# of jobs	events
done	1 014	100 070 000	done	2 505	183 310 000	done	3 763	246 660 000
sys-failed	86	7 430 000	failed	82	3 015 000	sys-failed	18	540 000
Total	1 100	107 500 000	Total	2 616	189 700 000	Total	3 781	247 200 000

SLAC			UKI-LT2-QMUL			UKI-SOUTHGRID-RALPP		
Status	# of jobs	events	Status	# of jobs	events	Status	# of jobs	events
done	699	69 900 000	done	3 474	317 945 000	done	1 218	94 125 000
failed	100	5 050 000	sys-failed	266	21 955 000	sys-failed	282	28 375 000
Total	799	74 950 000	Total	3 740	339 900 000	Total	1 500	122 500 000

- Reports on sites, channels and last status changes (only FastSim)

done JOBS					
Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_generic	MixBaBarBkg_NoPair.tcl	410	10 250 000	6 199 292
DG_BaBar	B0B0bar_generic	MixBaBarBkg_NoPair.tcl	404	10 100 000	6 299 508
DG_BaBar	ccbar	MixBaBarBkg_NoPair.tcl	215	10 750 000	2 357 493
DG_BaBar	uds	MixBaBarBkg_NoPair.tcl	203	19 490 000	2 703 634
DG_3	B+B-_generic	MixSuperbBkg_NoPair.tcl	526	5 260 000	9 798 367
DG_3	B0B0bar_generic	MixSuperbBkg_NoPair.tcl	516	5 160 000	9 817 195
DG_3	ccbar	MixSuperbBkg_NoPair.tcl	257	5 140 000	3 259 682
DG_3	uds	MixSuperbBkg_NoPair.tcl	252	10 080 000	3 782 590
DG_4	B+B-_generic	MixSuperbBkg_NoPair.tcl	520	5 200 000	9 292 856
DG_4	B0B0bar_generic	MixSuperbBkg_NoPair.tcl	623	6 230 000	11 393 686
DG_4	ccbar	MixSuperbBkg_NoPair.tcl	330	6 600 000	3 976 354
DG_4	uds	MixSuperbBkg_NoPair.tcl	252	10 080 000	3 679 825
DG_3	B+B-_generic	PacProduction.tcl	2 092	104 600 000	79 818 019
DG_3	B0B0bar_generic	PacProduction.tcl	2 060	103 000 000	79 992 260
DG_3	ccbar	PacProduction.tcl	1 129	112 900 000	24 731 284
DG_3	uds	PacProduction.tcl	1 556	311 200 000	51 541 216
DG_4	B+B-_generic	PacProduction.tcl	2 085	104 250 000	80 612 976
DG_4	B0B0bar_generic	PacProduction.tcl	2 025	101 250 000	84 969 268
DG_4	ccbar	PacProduction.tcl	1 101	110 100 000	27 916 869
DG_4	uds	PacProduction.tcl	2 624	524 800 000	89 399 621
Total			19 180	1 576 440 000	591 541 996

- Reports on sites, channels and last status changes (only FastSim)

List of the last 50 done JOBS							
Timestamp	runnum	Site	DG	Generator	tcl	status	
2010-03-10 09:10:24	1003010682	UKI-LT2-QMUL	DG_4	B+B-_generic	MixSuperbBkg_NoPair.tcl	done	
2010-03-10 09:09:46	1003010609	UKI-LT2-QMUL	DG_4	B+B-_generic	MixSuperbBkg_NoPair.tcl	done	
2010-03-10 09:09:46	1003010637	UKI-LT2-QMUL	DG_4	B+B-_generic	MixSuperbBkg_NoPair.tcl	done	
2010-03-10 09:07:11	1002016717	UKI-SOUTHGRID-RALPP	DG_3	uds	PacProduction.tcl	done	
2010-03-10 09:06:13	1002014818	UKI-SOUTHGRID-RALPP	DG_4	ccbar	PacProduction.tcl	done	
2010-03-09 11:16:53	1002013408	UKI-SOUTHGRID-RALPP	DG_4	uds	PacProduction.tcl	done	
2010-03-09 11:14:47	1002012521	UKI-SOUTHGRID-RALPP	DG_4	ccbar	PacProduction.tcl	done	

Production efforts

- February 2010:
Full Simulation (two flavours), Fast Simulation (two flavours)

Not taking into account the Production Tools development/testing time
3 people: **A. Fella**, M. Ronzano, **L. Tomassetti**
~~24/24 with some kind of shifts~~
3 exhausting weeks!

- **Summer 2010** (July?): Fast Simulation only?
production operation team involved in the 'execution' by taking shifts
- Production tools under upgrade to allow easy submissions;
more automation, sensors, utilities for testing
- Expert version still running (some adjustments to be done)

WebUI for shifts

- Simplified job preparation
- Collection of user's requests and validation by the production manager
- Selection of jobs to be submitted based on requests
- Automatic preparation of scripts to launch jobs to ALL available sites
- Sensor(s) to determine working and available sites
- 'Single-User', 'Single-script' submission to the GRID

User's Requests

- Extension of the bookkeeping database to keep track of user's requests for production
- Instead of using the wiki page, users should fill a form with their requests (analysis, geometry, generator, requested event, bkgmix?)
- The production manager should then validate the requests by giving approval and filling some extra fields:
either the number of events per job and the number of jobs
or the time needed to produce an event (to calculate the previous)
- Users must be registered with the SuperB LDAP directory in order to login to the WebUI and submit requests

FAST REQUESTS FORM

Production

Production Series:
2010_Summer

Analysis

Current Analysis:

Analysis	Available Analysis:
BtoTauNu	X
Generics	X

BtoKNUu

Add

REQUESTS

Requests

Insert new Request:

Analysis: BtoTauNu Generator: B+B-_generic Geometry: DG_BaBar Req Events: 0

Select Analysis for Requests:
all analysis

Current Requests:

Valid	Analysis	Generator ID	Geometry Type	Req Events	User ID	Operations
	BtoTauNu	B+B-_tau_DX	DG_BaBar	10000000	manzali	
	BtoTauNu	B+B-_tau_DX	DG_SuperB	1000000	manzali	
	Generics	B+B-_generic	DG_BaBar	50000000	manzali	
	Generics	B+B-_generic	DG_SuperB	500000000	manzali	
	Generics	B0B0bar_generic	DG_BaBar	50000000	manzali	
	Generics	B0B0bar_generic	DG_SuperB	500000000	manzali	
	Generics	ccbar	DG_BaBar	50000000	manzali	
	Generics	ccbar	DG_SuperB	500000000	manzali	
	Generics	uds	DG_BaBar	100000000	manzali	
	Generics	uds	DG_SuperB	1000000000	manzali	

User's Requests

Insert / View / ...

Modify a Request ✕

All form fields are required.

Prod_series: 2010_Summer

Analysis: Generics

Generator ID: B+B-_generic

Geometry Type: DG_SuperB

Note:

Req Events: 500000000

Min Run Num:

Events per job:

Number of jobs:

Max WCT: 57600

Exp event time:

Priority:

Status: open

User's Requests

Modify / Validate

User's Requests

- The Web interface, will use the sensor(s) to enable the submission to the available sites
- The system will select the type of job to submit among the validated requests until the requested number of events is reached
- Priority to requests can be given by the production manager
- Possibility to override the automatic selection of site / job (optional)
- Other policies/strategies can still be implemented (please, give comments...)

Sensors


- Indirect:
 - keep memory of the last submission to each site (using the db),
 - disable submission to these sites (red light),
 - query the bookkeeping database to count the pending jobs,
 - allow new submission (of half a bunch) when queue is below a threshold (yellow light),
 - allow new submission when queue is empty (green light)
- Direct (optional):
 - submit test jobs to each site,
 - monitor their status,
 - mark the site as available/not available
(and include/exclude the site from the undirect sensor)

Submission procedure

- When at least one site is available, the submission button (or the site submission button) will be enabled
- Pushing the button, the system will present a summary of the jobs that will be submitted (job details, site)
- After confirmation, the system will prepare all the job scripts + one wrapper script per site + a general launcher, presenting a final report with the shell command to use for real submission
- The shift taker must then use a ssh shell @cnafr to execute the general launcher

SUBMISSION FORM

General Submit

Production Series: 2010_Summer Status:  Submit

Prod Details

Data taken from database

Production Series:	2010_Summer
Production root:	/storage/gpfs_superb/prod/
TAG:	V0.2.3
ARCH:	Linux26SL4_j386_gcc346
Release working Directory:	/home/BABAR/brownd/SuperB/FastSim/V0.2.3/workdir/
LDAP uid:	LTomassetti

WebUI for shifts

Job preparation

Single sites

Name	Last Check	Status	
CIT_CMS_T2B	2010-05-24 16:33:06		<input type="button" value="Submit"/>
GRIF	2010-05-24 16:33:06		<input type="button" value="Submit"/>
GRISU-UNINA	2010-05-24 16:33:06		<input type="button" value="Submit"/>
IN2P3-CC	2010-05-24 16:33:06		
INFN-BARI	2010-05-24 16:33:06		
INFN-FERRARA	2010-05-24 16:33:06		
INFN-LNL-2	2010-05-24 16:33:06		<input type="button" value="Submit"/>
INFN-NAPOLI-ATLAS	2010-05-24 16:33:06		<input type="button" value="Submit"/>
INFN-PERUGIA	2010-05-24 16:33:06		<input type="button" value="Submit"/>
INFN-PISA	2010-05-24 16:33:06		<input type="button" value="Submit"/>
INFN-T1	2010-05-24 16:33:06		
RAL-LCG2	2010-05-24 16:33:06		<input type="button" value="Submit"/>
SLAC	2010-05-24 16:33:06		
UKI-LT2-QMUL	2010-05-24 16:33:06		<input type="button" value="Submit"/>
UKI-SOUTHGRID-RALPP	2010-05-24 16:33:06		
UNINA-EGEE	2010-05-24 16:33:06		<input type="button" value="Submit"/>
VICTORIA-LCG2	2010-05-24 16:33:06		

WebUI for shifts

Job preparation

Development status

- Main tasks:
 - Indirect sensor (70%)
 - Automatic site/request selection (50%)
 - Script preparation (70%)
 - Job script optimization (20%)
 - FastSim configuration structure (10%)
[see D. Brown talk on tuesday afternoon]

Operation team

- Currently, 10 people joined the operation team:
 - Matteo Manzali from Ferrara (expert, developer)
 - Alberto Cervelli and Benjamin Oberhof from Pisa
 - Leonid Burmistrov and Alejandro Perez from IN2P3
 - Alessandro Rossi from Perugia
 - Enrico Feltresi and Nicola Gagliardi from Padova
 - Alex Martin and Michael Lazos from Queen Mary
 - Students of Vladimir Aushev (?)
- More people is needed!

Shift takers

- A mailing list has already been created:

superb-shift-prod@lists.infn.it

- It will be used to:
 - coordinate the initial training and solve the ‘configuration’ problem
 - communicate between the operation team and experts during shifts

Shift takers duties

- Shift takers will be responsible of the production during her/his shift and will take care of submissions and monitoring.
- They will contact the on-call experts in case of problems/failures (after troubleshooting, using an ELOG, via email/skype/phone)
- Shift period: 8 hours (8 - 16, 16 - 24, 24 - 8)
- Schedule to be defined!
General rule: 2 shifts per week per shift taker
The first week will be covered by Luca and Armando

Shift takers duties

- Shift takers will be requested to submit the jobs by monitoring the production status and pushing a button in the interface (easy way!)
- They will carefully check the status of the sites by submitting – from time to time – a set of testing jobs through a specific web interface and collecting the results for analysis by production experts
- They will launch the jobs by running the scripts prepared through the web interface

Shift takers requirements

- Shift takers must:
 - be registered with the SuperB LDAP service (in order to get access to the WebUI)
 - have (or obtain) a personal certificate issued by INFN or a qualified institution (in order to get access to the grid via a grid proxy)
- see:
http://mailman.fe.infn.it/superbwiki/index.php/SuperB_registration_procedure
- Experience with fast simulation and/or grid submission is preferable but not required

Training

- Instruction Manual for the production tools will be given to the operation team
- Online help (in the WebUI) will be provided
- A detailed procedure for shifts will be available in the WebUI site
- Training using a test prod_series with the production release and all sites will be part of a stress-test of the system.
- Contacts: Luca and Armando

Tentative Schedule

- Development will finish on 21st (or 28th) June
 - Dependencies on FastSim configuration structure (two-ways)
 - tcl files for generics and signal modes (analysis)
 - tcl files for background inclusion
 - Discussion with FastSim people is needed
 - Previous releases cannot be used with the new interface & old interface will not be able to use the new release
 - Initial testing will include a couple of shift takers in order to stress-test the system for a week
- Training may begin on 29th June
 - Phone-Video Conference
 - VNC + individual test