

Distributed analysis functional testing using GangaRobot in the ATLAS experiment

F. Legger¹, B. Caron^{2,3}, J. Elmsheuser¹, M. Úbeda García⁴, A. W. Gordon², M. K. Jha⁵, D. C. Van der Ster⁶ for the ATLAS collaboration

¹Ludwig-Maximilians University, Munich, ²University of Alberta, ³TRIUMF, ⁴Royal Institute of Technology, Stockholm, ⁵INFN Bologna, ⁶CERN

Abstract

Automated distributed analysis tests are necessary to ensure smooth operations of the ATLAS grid resources. In this work we present the recent developments of the ATLAS GangaRobot, a tool designed to perform regular tests of all grid sites by running arbitrary user applications with varied configurations at predefined time intervals. Specifically the GangaRobot creates and submits several real ATLAS user applications to the various grid sites using the distributed analysis framework GANGA, a front end for easy grid job definition and management. Success or failure rates of these test jobs are individually monitored. Test definitions and results are stored in a database and made available to users and site administrators through a web interface, the ATLAS Site Status Board (SSB) and the Service Availability Monitor (SAM). The test results provide on the one hand a fast way to identify systematic or temporary site problems, and on the other hand allow for an effective distribution of the workload on the available resources.

ATLAS data analysis

Requirements:

- Large data volume: several PetaBytes/year
- High demand of computing resources: ~ 100k CPUs
- Data availability 24/7 at high bandwidth

➔ **The ATLAS computing model: Job goes to data**

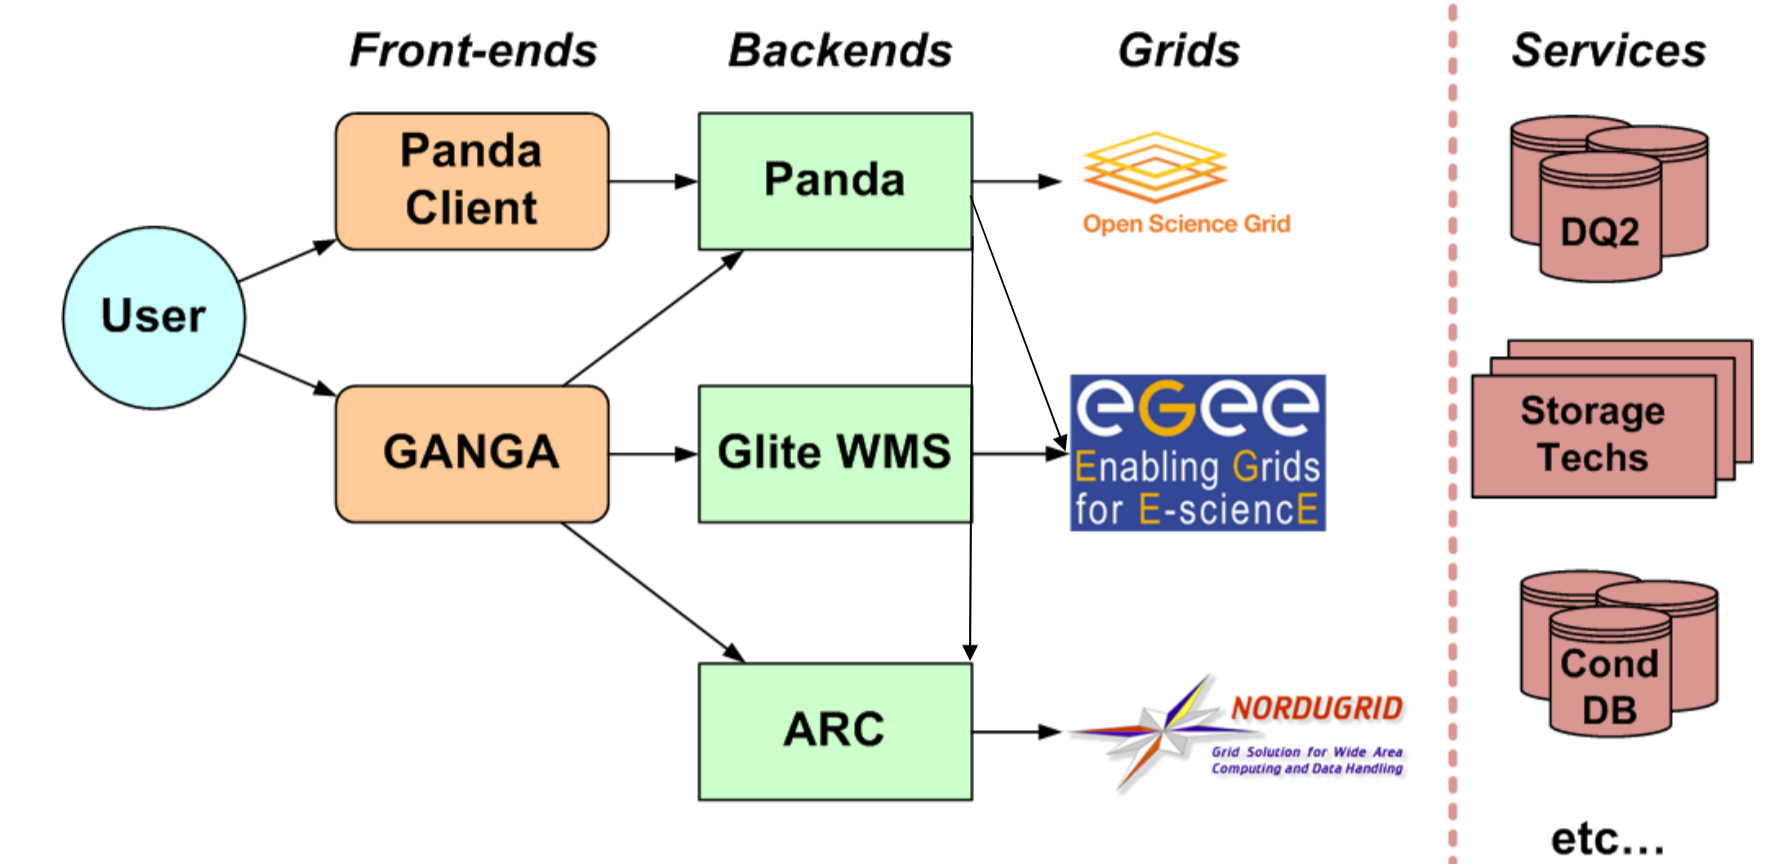
Typical users analysis:

- Athena analysis on both Monte Carlo (MC) and collision data, using various types of input: AODs, DPDs, ESDs;
- TAG-based analysis, which require direct access to data using a metadata index;
- calibration and alignment studies, requiring access to the conditions database;
- private MC production;
- ROOT analysis.

Distributed analysis in ATLAS

Grid computing:

- Centralized MC production and data reconstruction
- Centralized data management by DDM/DQ2 Tools
- Centralized distributed data analysis
- Grid infrastructure:
 - EGEE, NorduGrid, Open Science Grid
- Clients for job submission:
 - pathena (Panda client), GANGA



Most wanted user request:

“What runs on my laptop should run on the Grid!”

→ To ensure that this complex interplay of very different software/hardware environments is working:

- Validation of distributed analysis with automated tests:
 - **Functional tests: GangaRobot**
 - **Stress tests: HammerCloud** → see: 'HammerCloud: A Stress Testing System for Distributed Analysis' (PS36-4-094) presented by D. Van der Ster

The GangaRobot/HammerCloud framework

GangaRobot/HammerCloud allows for rapid definition and execution of test jobs, with hooks for pre- and post-processing:

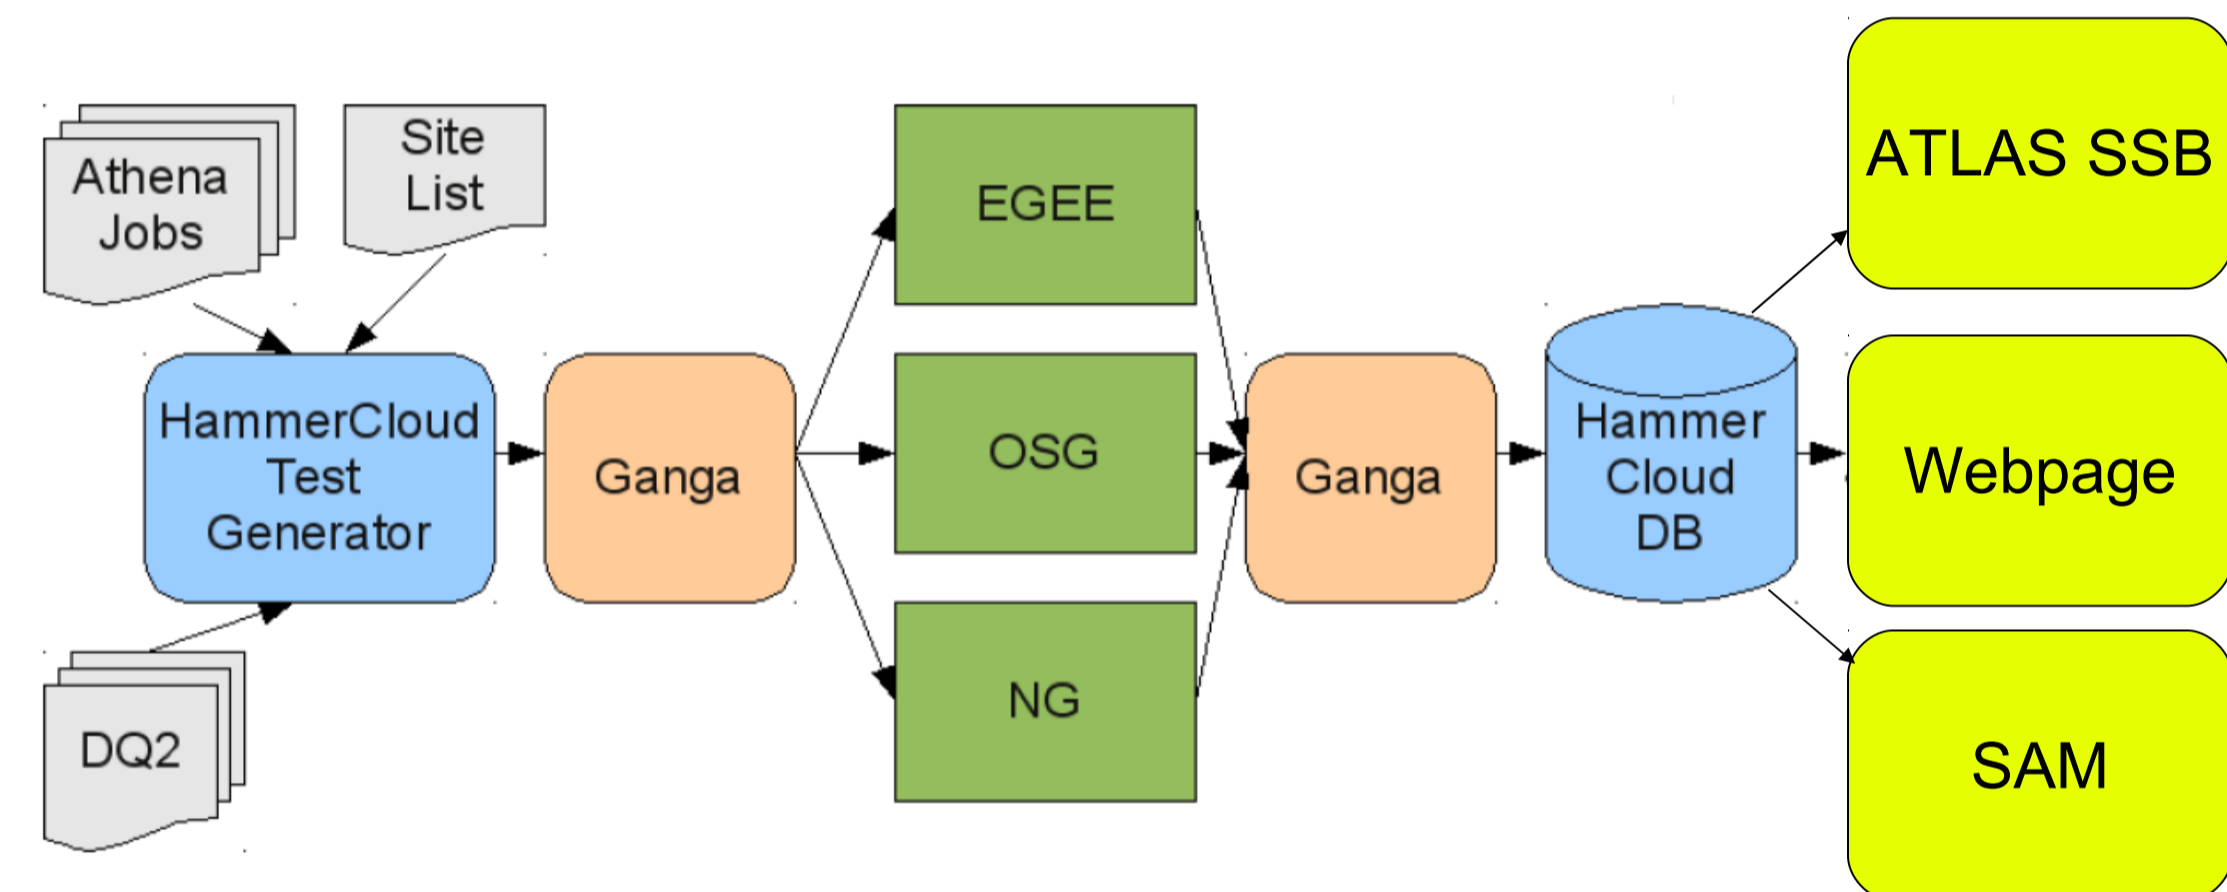
- Tests defined by operator:
 - Athena version, analysis code, input datasets, sites to be tested
- Jobs are submitted using **GANGA**:
 - OSG/Panda, EGEE/LCG
- Job status is periodically monitored:
 - Results locally recorded
- Results are published to:
 - **GANGA Runtime Info System**, to avoid failing sites
 - **Service Availability Monitor (SAM)**
 - **ATLAS Site Status Board (SSB)**
 - **HammerCloud/GangaRobot website**, monitored by DA shifters
 - GGUS tickets sent for failures

Functional tests with GangaRobot:

- Runs several types of analysis jobs (MC, data)
- Tests several input file access modes (copy-to-scratch Filestager, dcap/rfio)
- Tests database access (SQUID, Frontier)

Interface:

- MySQL database
- Web pages generated with Django



Functional tests with GangaRobot

Administrator creates a test template for each functional test.

Select test template to change	id	Category	Active	Description	Period	Lifetime	Inputtype	Clouds	Sites
<input type="checkbox"/>	34	functional	<input checked="" type="checkbox"/>	DIPMakler 15.6.12 Panda Frontier/Squid test - sample dataset	6	1	PANDA	DE_PANDA	ANALY_CV, ANALY_FREIBURG, 13 more...
<input type="checkbox"/>	34	functional	<input checked="" type="checkbox"/>	DIPMakler 15.6.12 LCG Frontier/Squid test - sample dataset	6	1	DQ2_LOCAL	DE	FIX-LCG-DATADISK, IP2-ML-DATADISK, 11 more...
<input type="checkbox"/>	31	functional	<input checked="" type="checkbox"/>	DIPMakler 15.6.12 LCG Frontier/Squid test	6	1	DQ2_LOCAL	ES, DE, 5 more...	FIX-FCO, FIX-LCG-DATADISK, 29 more...
<input type="checkbox"/>	30	functional	<input checked="" type="checkbox"/>	DIPMakler 15.6.12 PANDA default data access Frontier/Squid test	6	1	PANDA	US, DE, PANDA, 10 more...	ANALY_BNL_ATLAS_1, ANALY_NETZ, 83 more...
<input type="checkbox"/>	22	functional	<input checked="" type="checkbox"/>	15.6 LCG Frontier/liquid CLODED	6	1	DQ2_LOCAL	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...
<input type="checkbox"/>	21	functional	<input checked="" type="checkbox"/>	15.6 LCG Frontier/liquid	6	1	DQ2_LOCAL	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...
<input type="checkbox"/>	20	functional	<input checked="" type="checkbox"/>	15.6 LCG Frontier/liquid	6	1	DQ2_LOCAL	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...
<input type="checkbox"/>	19	functional	<input checked="" type="checkbox"/>	15.6 LCG Frontier/liquid	6	1	DQ2_LOCAL	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...
<input type="checkbox"/>	12	functional	<input checked="" type="checkbox"/>	15.6.9 Panda Filestager	6	1	FILE_STAGER	US, DE, PANDA, 10 more...	ANALY_BNL_ATLAS_1, ANALY_NETZ, 83 more...
<input type="checkbox"/>	11	functional	<input checked="" type="checkbox"/>	15.6.9 Panda	6	1	PANDA	US, DE, PANDA, 10 more...	ANALY_BNL_ATLAS_1, ANALY_NETZ, 83 more...
<input type="checkbox"/>	10	functional	<input checked="" type="checkbox"/>	15.6.9 LCG DQ2 Filestager	6	1	FILE_STAGER	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...
<input type="checkbox"/>	9	functional	<input checked="" type="checkbox"/>	15.6.9 LCG DQ2 Local	6	1	DQ2_LOCAL	IT, DE, 8 more...	INPN-TL_MCCDISK, INPN-FRASCATI_MCCDISK, 99 more...

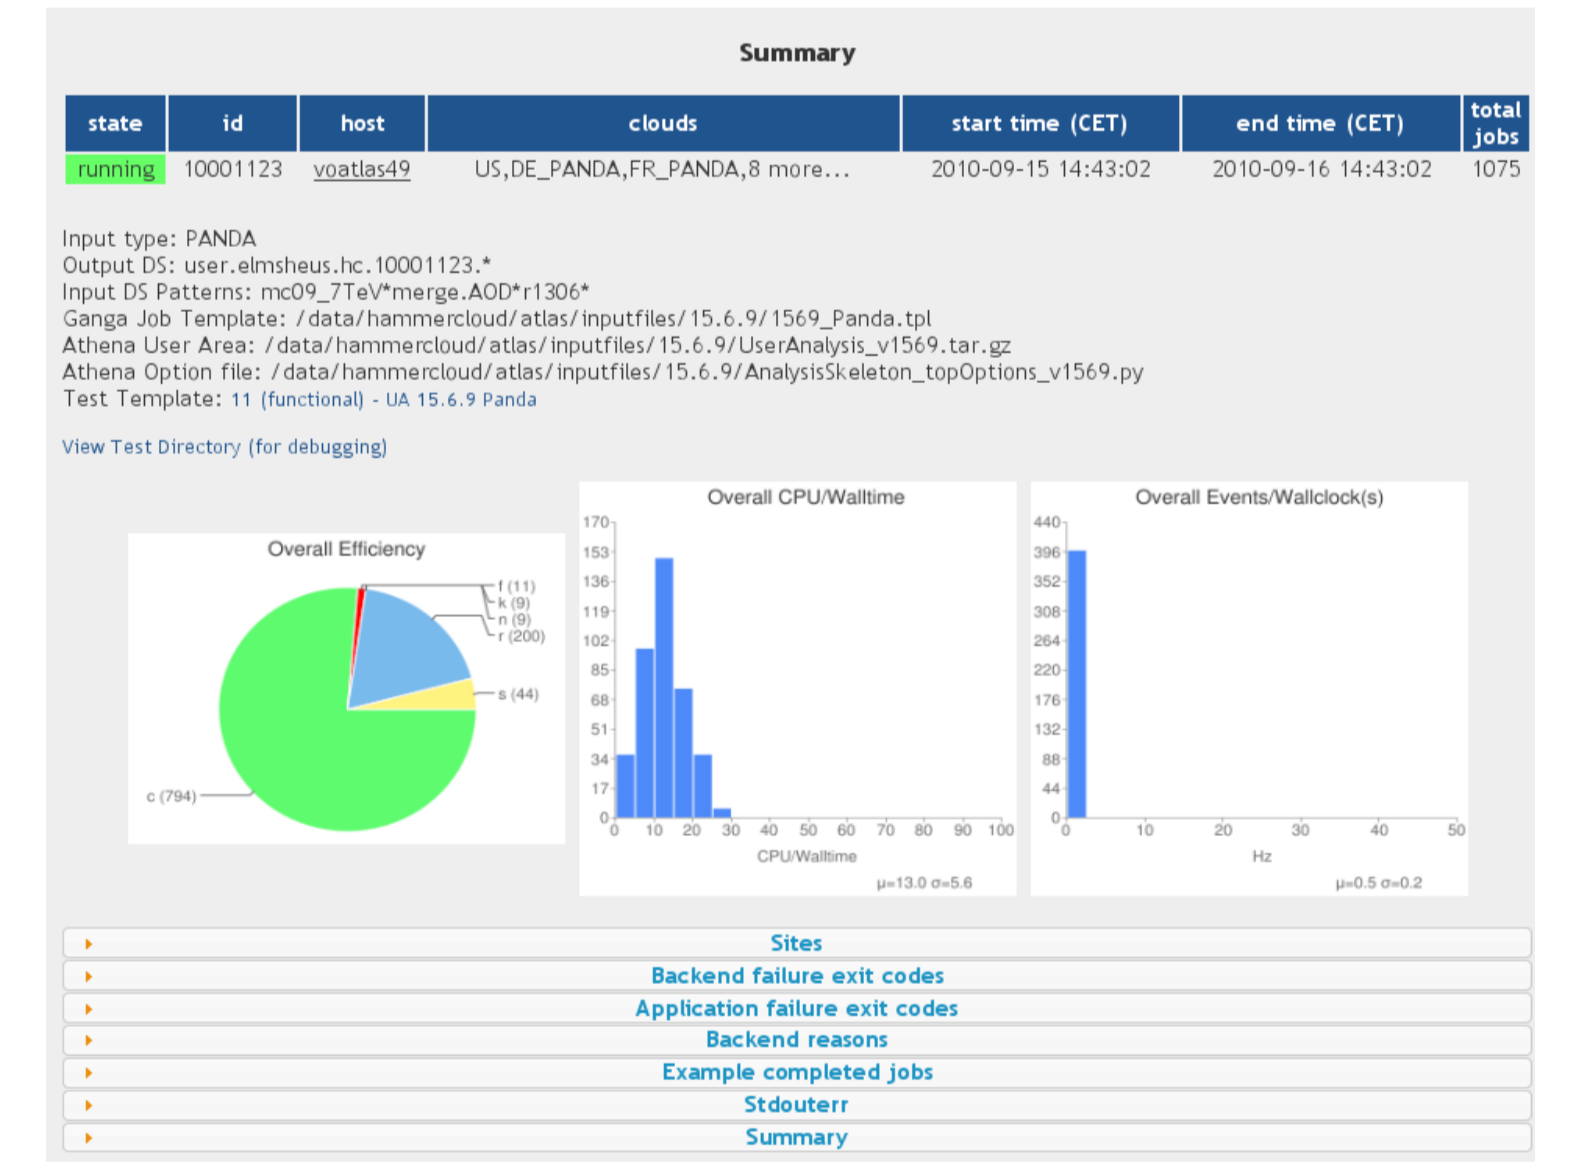
Functional tests are automatically created from active test templates and scheduled to run 24 hours. Functional tests are configured to always have 1 job running at each site

More at: <http://hammercloud.cern.ch/atlas/>

Currently 8 functional tests with different configurations are run daily. Each test submits ~2000 jobs. Completed jobs are normally more than 80% of the total.

state	id	host	template	start time (CET)	end time (CET)	clouds	sites	submit jobs	run jobs	comp jobs	fail jobs
running	10001123	voatlas73	DIPMakler 15.6.12 LCG Frontier/Squid test - sample dataset	2010-09-15 10:00:02	2010-09-16 10:00:02	DE, NL, 4 more...	ES, DE, NL, 4 more...	15	25	28	7
running	10001124	voatlas73	DIPMakler 15.6.12 PANDA default data access Frontier/Squid test	2010-09-15 14:43:04	2010-09-16 14:43:04	DE, PANDA, 9 more...	ANALY_BNL_ATLAS_1, ANALY_NETZ, 83 more...	12	14	64	39
running	10001123	voatlas49	15.6.9 Panda	2010-09-15 14:43:02	2010-09-16 14:43:02	US, 9 more...	ANALY_ALBERTA_WG1, ANALY_ANASIS_81 more...	46	194	802	12
running	10001120	voatlas73	DIPMakler 15.6.12 PANDA Frontier/Squid test - sample dataset	2010-09-15 14:30:02	2010-09-16 14:30:02	DE, PANDA, 9 more...	ANALY_CV, ANALY_FREIBURG, 13 more...	12	20	86	4
running	10001123	voatlas73	DIPMakler 15.6.12 LCG Frontier/Squid test - sample dataset	2010-09-15 12:56:01	2010-09-16 12:56:01	DE	FIX-LCG-DATADISK, CIPROBET-LCG-DATADISK, 10 more...	16	14	100	0
running	10001120	voatlas49	15.6.9 Panda Filestager	2010-09-15 12:06:01	2010-09-16 12:06:01	US, 9 more...	ANALY_ALBERTA_WG1, ANALY_ANASIS_81 more...	76	128	1214	90
running	10001118	voatlas73	DQ2 Filestager	2010-09-15 11:25:01	2010-09-16 11:25:01	IT, DE, US, 7 more...	AUSTRALIA-ATLAS_MCCDISK, BEIJING-LCG_MCCDISK, CA-ALBERTA-VESTIGID, TL-MCCDISK, 75 more...	22	22	256	28
submitting	10001127	voatlas73	DQ2 Local	2010-09-15 16:19:01	2010-09-16 16:19:01	IT, DE, US, 7 more...	AUSTRALIA-ATLAS_MCCDISK, BEIJING-LCG_MCCDISK, CA-ALBERTA-VESTIGID, TL-MCCDISK, 75 more...	0	0	0	0

An overview page is available for each test, reporting various metrics and debug informations.



Results

The **site efficiency** on functional tests is defined as:

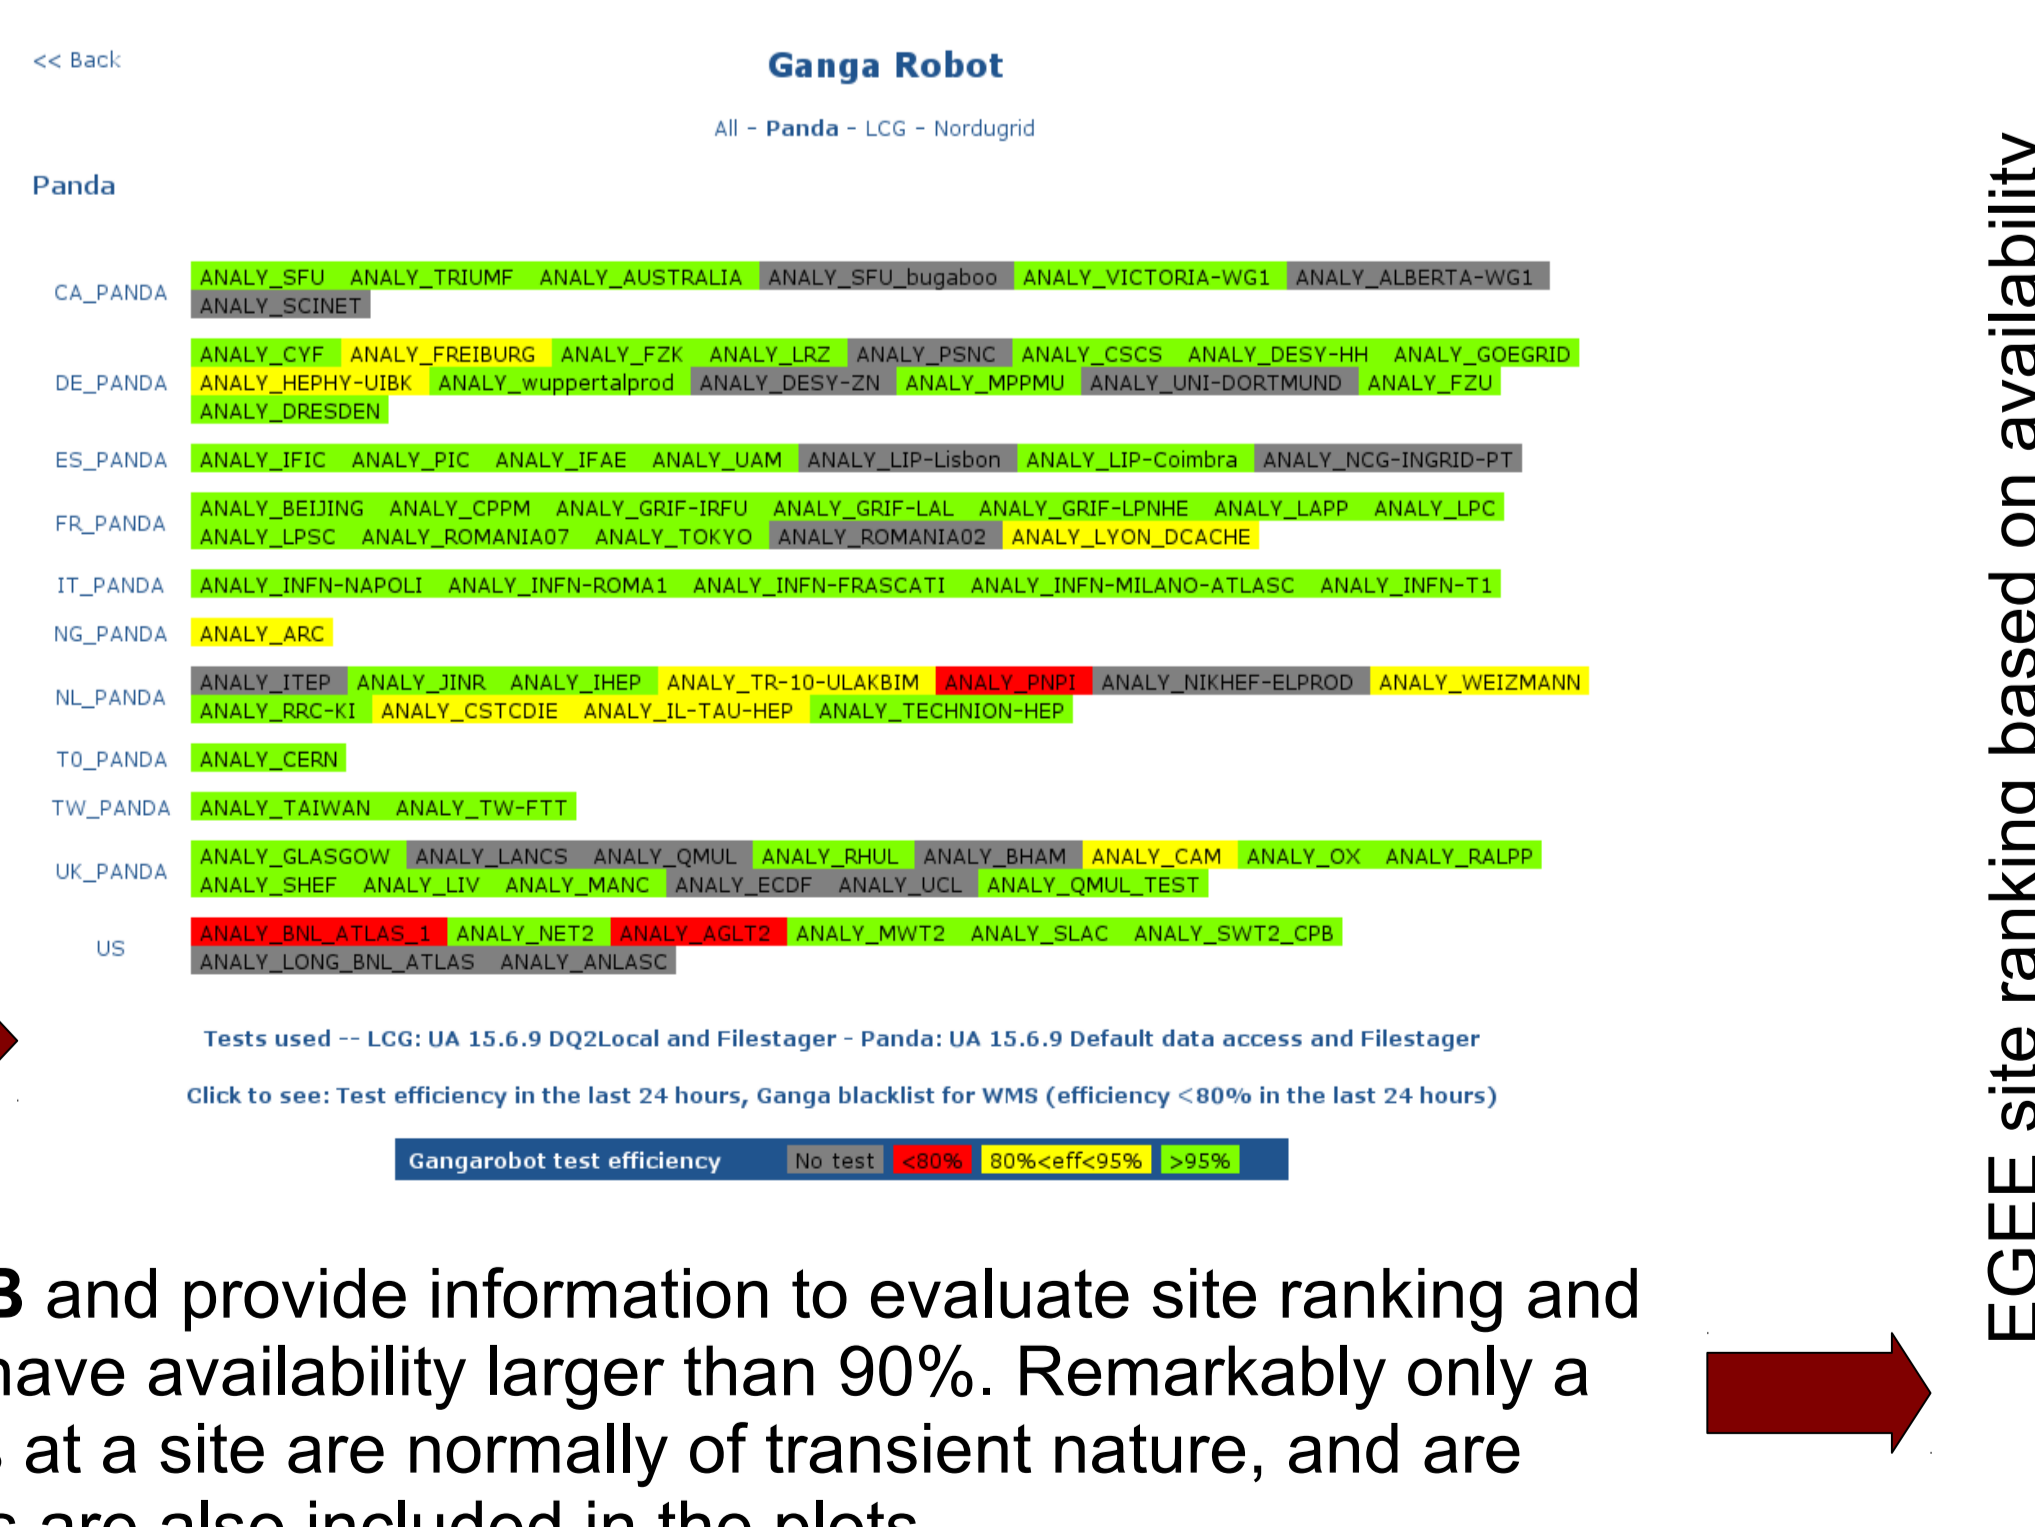
$$\frac{\# \text{ successful jobs in the last 24 hours}}{\# \text{ successful + failed jobs in the last 24 hours}}$$

Automatic blacklisting for EGEE sites

GangaRobot results are used to automatically prevent user analysis jobs to be submitted to temporarily failing sites. A site is automatically blacklisted if the efficiency drops below 80% in the last 24 hours.

A summary of the efficiency per site is displayed on the HammerCloud/GangaRobot website.

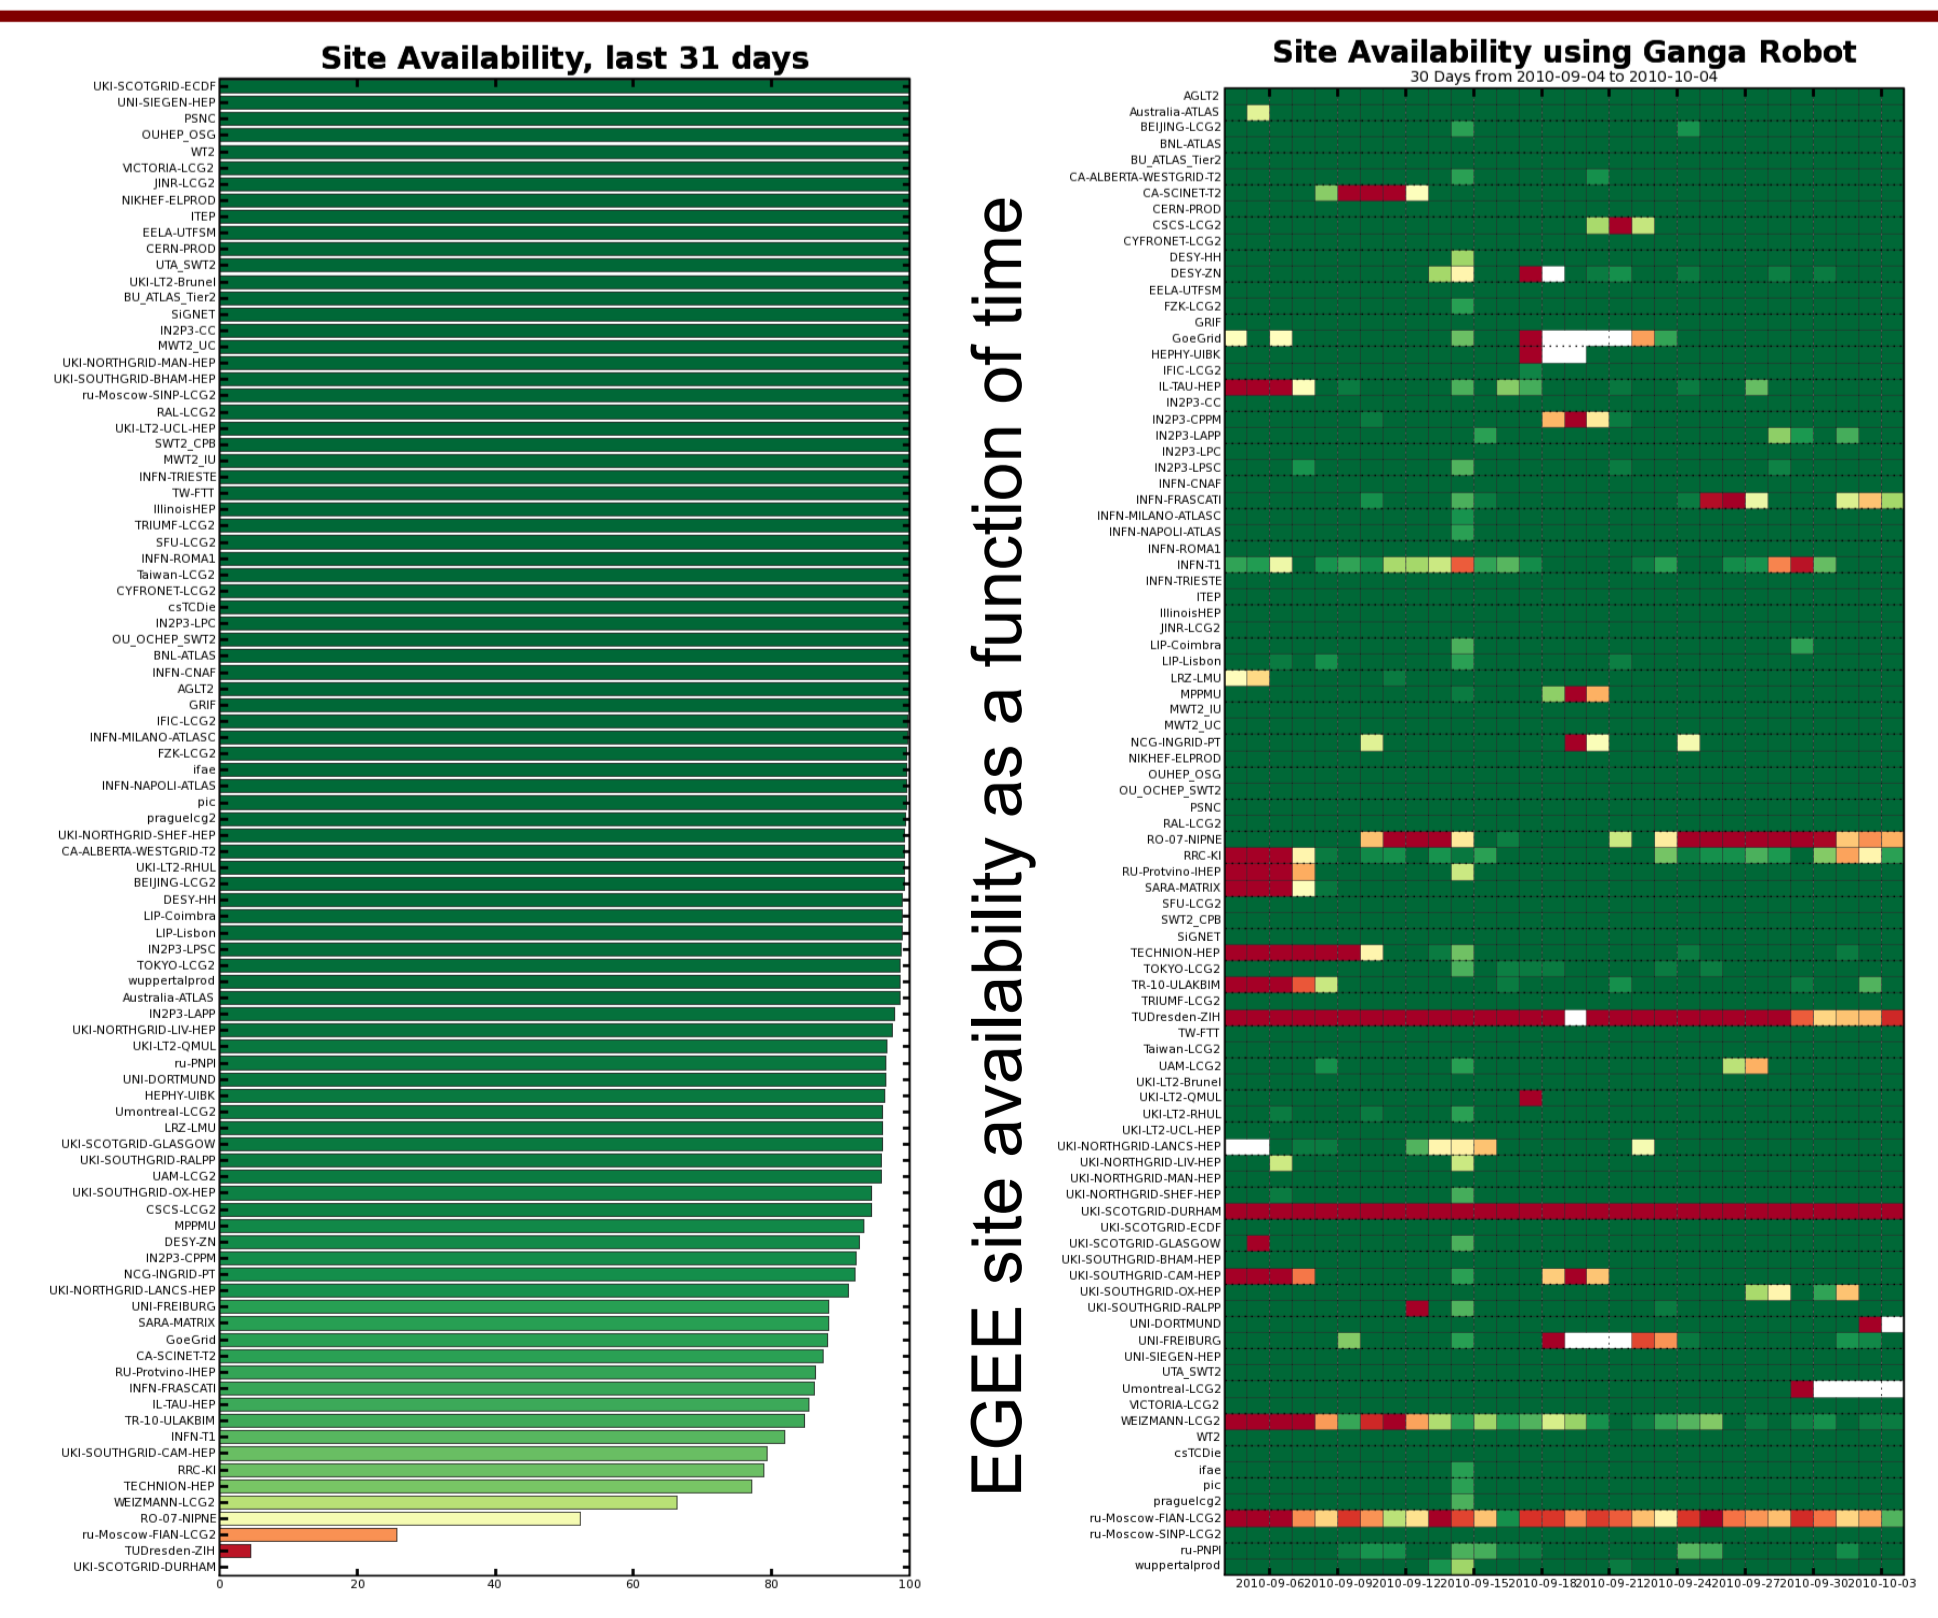
GangaRobot results are exported to **SAM** and the **ATLAS SSB** and provide information to evaluate site ranking and availability as a function of time. In general, most EGEE sites have availability larger than 90%. Remarkably only a handful of sites have availabilities smaller than 80%. Problems at a site are normally of transient nature, and are solved in a couple of days. Note that scheduled site downtimes are also included in the plots.



EGEE site ranking based on availability

Plots generated by SAM:

<http://dashb-atlas-sam.cern.ch/>



EGEE site availability as a function of time

Conclusion

The new Hammercloud/GangaRobot framework is running since May 2010. It provides an easy interface for the creation and submission of functional/stress tests. Functional test results are stored in a database, published on the HammerCloud website, and exported to further Grid information services such as SAM and the ATLAS SSB. The site efficiency on functional tests is used to blacklist sites temporarily failing the tests.