

Experiences and Design of the New EGI Metrics Portal

Iván Díaz
Álvaro Simón
Carlos Fernández

CESGA



Introduction

Base Technologies

Python

Django

The Portal

Functionality

Conclusions

- An important tool for managing the EGI project.
- Aggregates quality and health metrics from multiple and diverse sources.
- These measure both infrastructure status and the activity of affiliated staff and users.

Metrics sources include:

- GOCDB
- Accounting Portal
- Top BDII
- GGUS
- ...?

- Automatic Metrics are derived automatically from a information source.
 - Number of solved tickets in a month.
 - Nearline disk storage on a NGI.
 - Number of jobs sent.
- Since automatic data can be incomplete or wrong, users must flag them as "verified" explicitly.
- If users know a metric to be wrong, they can also input a value manually.

- Manual Metrics are compiled manually by users due to their intrinsic nature or lack of source support.
 - Number of papers published by members of an NGI.
 - Number of users of a community or portal.
 - Number of regional instances of EGI operational tools.
- Manual metrics have normally a fixed numeric format, but comments can be added for more detail/clarifications.

- Based on legacy code from accounting portal:
 - PHP based, non optimally modularized.
 - Lack of expandability.
 - Not geared to fast code changes.
 - Code base still comparatively small (vs. Acct Portal).
 - There was a perceived added complexity in the requirement capture and development of the following stages.
- The above qualities favored a change of framework and methodology.

- Good for rapid prototyping.
- Functional and object oriented.
- Very flexible (list comprehensions, closures, first-order functions, reflection).
- “Batteries included”.

- Connector scripts and other admin code already written in Python.
- Existing coding experience.
- Improved code maintenance compared with Perl or PHP.
- Reduced development effort and risks.

- A Python framework for web applications.
 - Separates concerns like DB access, URL mapping, view templating...
 - Based on the DRY (Don't Repeat Yourself) philosophy.
 - Extensive third-party and community support. Very detailed and extensive documentation.
 - Suited for rapid prototyping.
 - Support for model-based development and low-impact model changes and migration.

The motivation behind modular and OO development

- The access to data in the DB is separated from the underlying query:
 - Support for non-SQL DB backends.
 - Lazy access and caching.
- URLs are decoupled from the view and HTTP GET format.
- View data gathering is separated from template application.
- Makes easy the distribution of tasks between several servers and differential scaling.

- DRY (Don't Repeat Yourself) states that neither data nor code should be repeated in the codebase.
- It is specially useful in web services (where repetition is common) and requires framework support.
- Although present in some form in every high-level language, Django pursues total adherence.
- Related with separation on concerns.
- It opposes boiler-plate and generated code, in a sense ultimately promotes data as code.
- Example: Metrics forms created on-the-fly from the model.

Strong third party and community support with generic feature-laden extensions for:

- Authentication.
- Database and model migration.
- Historic records.
- Debugging.

New views only require:

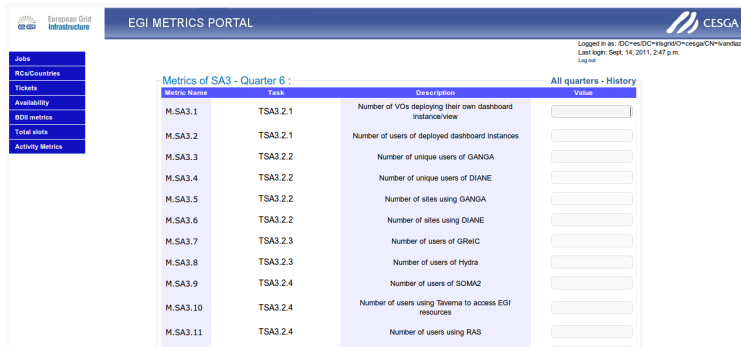
- Database and model support for the data shown.
- A view function and a template.
- A generic regular-expression mapping between the view function and a URL.

Django also allows for rapid model changes. Using the South migration facility:

- The developer modifies, adds or removes a field from a model or adds a new model.
- A command invocation creates a South migration automatically with the differences.
- The migration is applied, and the underlying database is modified (most times).
- The new changes are useable in 1-2 minutes.
- There is coding support for more complicated migrations (e.g. data migrations).

The portal offers:

- A public query interface to graph various metrics data (like in the old portal).
- Secure forms to visualize and edit metrics data.
- Aggregated views for project quarter and NGI metrics.



The screenshot shows the EGI Metrics Portal interface. At the top, there is a navigation bar with the EGI logo, "European Grid Infrastructure", "EGI METRICS PORTAL", and the CESGA logo. A user is logged in as "IDC=es:IDC=Hsgr10@cesga/CN=ivandaz" with a last login of "Sept. 14, 2011, 2:47 p.m." and a "Log out" link.



On the left side, there is a vertical menu with the following items: Jobs, RCs/Countries, Tickets, Availability, BDI metrics, Total slots, and Activity Metrics.

The main content area displays "Metrics of SA3 - Quarter 6" and "All quarters - History". Below this is a table with the following columns: Metric Name, Task, Description, and Value.

Metric Name	Task	Description	Value
M.SA3.1	TSA3.2.1	Number of VOs deploying their own dashboard instance/view	<input type="text"/>
M.SA3.2	TSA3.2.1	Number of users of deployed dashboard instances	<input type="text"/>
M.SA3.3	TSA3.2.2	Number of unique users of GANGA	<input type="text"/>
M.SA3.4	TSA3.2.2	Number of unique users of DIANE	<input type="text"/>
M.SA3.5	TSA3.2.2	Number of sites using GANGA	<input type="text"/>
M.SA3.6	TSA3.2.2	Number of sites using DIANE	<input type="text"/>
M.SA3.7	TSA3.2.3	Number of users of GREIC	<input type="text"/>
M.SA3.8	TSA3.2.3	Number of users of Hydra	<input type="text"/>
M.SA3.9	TSA3.2.4	Number of users of SOMA2	<input type="text"/>
M.SA3.10	TSA3.2.4	Number of users using Taverna to access EGI resources	<input type="text"/>
M.SA3.11	TSA3.2.4	Number of users using RAS	<input type="text"/>

- Only needed for metric input pages, the data pages are public.
- Due to the existence of several authentication methods in EGI, supports a chain of authentication options in this order:
 - EGI certificate authentication. Users with a valid EGI certificate are automatically authenticated without further interaction.
 - SSO LDAP authentication. It delegates authentication on SSO and does not store passwords locally.
 - Local DB authentication. Usually for admin users, others should be handled by the above cases.

- Manual metrics are entered exclusively by users without any computation.
- Nevertheless, the format of the data is checked.
- Data for all project quarters until the current can be visualized and edited (this policy can change in a future).


EGI METRICS PORTAL


Jobs

RCs/Countries

Tickets

Availability

BDII metrics

Total slots

Activity Metrics

Metrics of SA3 - Quarter 6 :

Metric Name	Task	Description	Value
M.SA3.1	TSA3.2.1	Number of VOs deploying their own dashboard instance/view	<input type="text"/>
M.SA3.2	TSA3.2.1	Number of users of deployed dashboard instances	<input type="text"/>
M.SA3.3	TSA3.2.2	Number of unique users of GANGA	<input type="text"/>
M.SA3.4	TSA3.2.2	Number of unique users of DIANE	<input type="text"/>
M.SA3.5	TSA3.2.2	Number of sites using GANGA	<input type="text"/>
M.SA3.6	TSA3.2.2	Number of sites using DIANE	<input type="text"/>
M.SA3.7	TSA3.2.3	Number of users of GREIC	<input type="text"/>
M.SA3.8	TSA3.2.3	Number of users of Hydra	<input type="text"/>
M.SA3.9	TSA3.2.4	Number of users of SOMA2	<input type="text"/>
M.SA3.10	TSA3.2.4	Number of users using Taverna to access EGI resources	<input type="text"/>
M.SA3.11	TSA3.2.4	Number of users using RAS	<input type="text"/>

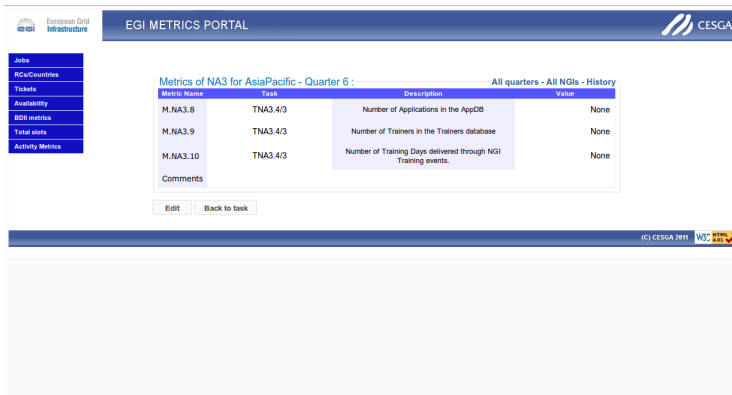
Logged in as: (DC=es,DC=Higrid)O=cesga,CN=ivandiaz

Last login: Sept. 14, 2011, 2:47 p.m.

[Log out](#)

All quarters - History













- Metrics associated to NGIs, traditionally filled by NGI staff each quarter.
- Also viewable and editable by quarter.



The screenshot displays the EGI Metrics Portal interface. At the top left, there is the EGI logo and the text 'European Grid Infrastructure'. The main header reads 'EGI METRICS PORTAL'. On the right side of the header is the CESA logo. A left-hand navigation menu includes links for 'Jobs', 'RCs/Countries', 'Tickets', 'Availability', 'BDII metrics', 'Total slots', and 'Activity Metrics'. The main content area shows a table titled 'Metrics of NA3 for AsiaPacific - Quarter 6 :'. To the right of the title is a link 'All quarters - All NGIs - History'. The table has four columns: 'Metric Name', 'Task', 'Description', and 'Value'. It lists three metrics, all with a value of 'None'. Below the table is a 'Comments' section and two buttons: 'Edit' and 'Back to task'. At the bottom right of the page, there is a footer with '(C) CESA 2011' and a 'HTML 5.01' icon.

Metric Name	Task	Description	Value
M.NA3.8	TNA3,4/3	Number of Applications in the AppDB	None
M.NA3.9	TNA3,4/3	Number of Trainers in the Trainers database	None
M.NA3.10	TNA3,4/3	Number of Training Days delivered through NGI Training events.	None

- These metrics are precalculated by the portal, which accesses the data sources and computes values for them.
- Automatic values can be verified by users or entered manually in case of data or human errors.
- Any modification in the values sets the “manual” flag.

M.SA1.OperationalSecurity.4	TSA1.2	Number of suspended sites for security issues		12
M.SA1.Service.Validation.2	TSA1.3	Number of staged rollout releases undertaken & rejected		32
M.SA1.Size.1	TSA1.1	Total number of production resource centers that are part of EGI.		348 
M.SA1.Support.1	TSA1.7	Number of operational tickets created/solved per month		570-912 
M.SA1.Support.2	TSA1.7	Average & median ticket solution time (business days) per month		30.41-4.370 
M.SA1.Support.4	TSA1.7	Average ticket response time (business days) per month		7.024 
M.SA1.Support.5	TSA1.7	Number of tickets solved by TPM (first line support)		0 
M.SA1.Support.6	TSA1.7	Ticket assignment time to Support Units by TPM (first line support)		0.2993 

- Provides wiki-like access to old versions of metrics.
- Records the time, date and author of the modifications.
- Allows undo functionality by saving old versions.

Modification history of SA1 - Quarter 6

#	Revision	Username	Comment
5	Aug. 29, 2011, 11:06 a.m.	/DC=es/DC=irisgrid/O=cesga/CN=alvarosimon	
4	Aug. 22, 2011, 10:59 a.m.	/DC=es/DC=irisgrid/O=cesga/CN=ivandiaz	
3	Aug. 11, 2011, 10:33 a.m.	/DC=es/DC=irisgrid/O=cesga/CN=alvarosimon	
2	Aug. 11, 2011, 10:32 a.m.	/DC=es/DC=irisgrid/O=cesga/CN=alvarosimon	
1	Aug. 11, 2011, 10:31 a.m.	/DC=es/DC=irisgrid/O=cesga/CN=alvarosimon	

M.SA1.Accounting	TSA1.5	Number of sites adopting new messaging for Usage Record publication	None
M.SA1.Operation.2	TSA1.7	Number of sites suspended	None

- All Quarters View
 - Shows all recorded quarters for common or NGI metrics of a task.
- All NGIs View
 - Shows all recorded NGI metrics for a task and quarter, grouped by NGI.
- These views do not show NGIs or Projects quarters not in the DB.

- The new Metrics Portal reduces development effort and cycles, and minimize development cycles by basing itself in a flexible and powerful framework and by externalizing common portal requirements when possible.
- In this way, the development effort is centered in adaptation and customization for the EGI project, and the removal of change inertia on the set of metrics of the project.
- The portal allows the introduction and modification of metrics with ease and transparency, even new tasks and their associated metrics could be created in a short time.

- Tasks like authentication or database retrieval are externalized to simplify the portal and avoid security issues.
- DRY and separation of concerns allow change resilience. Thus models, the name of tasks, URL format or the appearance could be changed with minimal effort.