



The gLite Data Management Continuous Integration and Testing Process

Alejandro Álvarez (CERN)
DMS

Overview

- Introduction
- Previous status
- Deployment & Tests
- SAKET
- Future work
- Conclusions



Introduction

- As in any other software product, developers change code daily
- But, additionally, even if they don't, we still can get changes at project level
 - New releases of software we depend on (VOMS, BDII, ...)
- One change can break nothing, everything or just one thing
 - The later we find out, the harder to fix
 - Who? What? How?
- If these problems pop out at certification time, the release process is delayed

Introduction

- So, it is useful to frequently build, deploy and test
 - Errors and conflicts can be detected the day after a change occurred
 - Easier to identify
- Also, the developers will get feedback quickly
 - Like a regression bug coming back
 - Or an dependency changing
- Release candidates will be more reliable and up to date
 - Certification process will likely be faster

Introduction

- So we need
 - A build platform
 - ETICS
 - A test platform
 - We opted for ETICS and use the same platform as for build
 - A test environment
 - This is, external service providers (SE, LFC, BDII)
 - Tests
 - The existing ones, if possible
 - Something to orchestrate the process and generate the reports

Previous status

- ETICS was already used for building, but not widely for testing
- cert-tb-cern had everything we needed
 - Some minor issues were fixed
- Yaimgen did the deployment and testing launching
 - Custom tool for deployment
 - But not completely automated
- The tests were out there, but not consistent
 - Each set had its own “director” script
 - Plain text output: painful to debug
 - No common naming policies or anything

Steps taken

Automated deployment

- Yaimgen had to be adapted to run with no manual intervention
 - No confirmation messages or anything
 - Users and passwords through command line
- The argument values are passed through properties
 - Client → ETICS → Yaimgen
- Yaimgen generates XML output if requested
 - Easy to parse
 - XSLT does the “magic” to have HTML content
- Still can be used manually

Tests

- A common naming convention and structure was agreed inside DM
 - Tests are classified as Functional, Regression and Performance
 - Plus unstable for new ones
 - Named as prefix-test_name (dpns-mkdir, fts-bug3365)
- A bash script initializes the environment
- A common wrapper is used for all of them
 - Creates the proxy, set up the environment (using the bash script), ...
 - The tests are still stand-alone
- So, as a side effect of the continuous integration, we have achieved test consistency at PT level
- Tests are packaged during build time

SAKET

- ETICS builds and gives us a report and a repository
- Yaimgen deploys and triggers the tests
 - XML logs result from this process
- But we lack a tool to orchestrate the whole process
 - Summarizing the results
- Therefore we created SAKET

SAKET

- Swiss Army Knife for ETICS Testing
- Python application that
 - Orchestrates builds and tests
 - Generates a XML report with the results
 - Stored and submitted by mail
 - Human readable thanks to XSLT



SAKET Archive



Day	Time	Duration	Build	Deployment	Test
February 2011					
Monday	21 01:02:34	04:01:54	71%	40%	0%
Sunday	20 02:25:24	05:25:07	71%	40%	0%
Saturday	19 02:31:11	06:30:19	71%	40%	0%
Friday	18 06:17:27	09:16:23	100%	67%	27%
Thursday	17 03:51:29	07:51:19	100%	73%	27%
Wednesday	16 00:28:08	03:27:49	57%	33%	21%
Tuesday	15 01:36:31	05:36:20	22%	17%	0%
Monday	14 04:14:57	03:29:34	33%	13%	0%
Friday	11 03:35:53	03:56:18	0%	0%	0%
Thursday	10 07:57:42	07:20:13	78%	22%	10%
Wednesday	09 12:58:49	04:27:48	44%	39%	10%
Monday	07 06:30:02	06:05:12	78%	65%	6%
Thursday	02 08:22:27	11:02:17	100%	70%	25%

SAKET Archive (C) Members of the EGEE Collaboration
 Alejandro Alvarez <aalvarez@cern.ch>

[Documentation](#)

Status of DM HEAD (2011-02-18 06:17:27)



Project config.	Platform	Build	Test component	Deployment	Test
glite_3_2_cert	sl5_ia32_gcc412	Success (101 / 101)	glite-SE_dpm_mysql	Failed	-
			glite-SE_dpm_disk	Failed	-
			glite-LFC_mysql	Failed	-
			glite-LFC_oracle	Failed	-
	sl5_x86_64_gcc412	Success (101 / 101)	glite-SE_dpm_mysql	Success	Success (38 / 38)
			glite-SE_dpm_disk	Success	Not defined
			glite-LFC_mysql	Success	Success (61 / 61)
			glite-LFC_oracle	Success	Success (61 / 61)

org.glite.data

org.glite.data.HEAD

Project config.	Platform	Build	Test component	Deployment	Test
glite_3_2_cert	sl5_ia32_gcc412	Success (54 / 54)	-	-	-
	sl5_x86_64_gcc412	Success (54 / 54)	glite-lcutil	Success	Failed (47 / 48)

SAKET Archive (C) Members of the EGEE Collaboration
 Alejandro Alvarez <aalvarez@cern.ch>

[Documentation](#)

[INFO] Using configuration file /opt/glite/tests/fts/FTS-config
/opt/glite/tests/fts/FTS-config [SUCCESS]

fts [SUCCESS]

Success rate: 100% (61/61)
Start time: Wed Apr 06 02:36:31 2011
End time: Wed Apr 06 03:20:22 2011
Duration: 0h 43m 51s

fts-CERNCERN-bulk-submission [SUCCESS]
fts-CERNCERN-submission [SUCCESS]
fts-CERNCERN-submission-same [SUCCESS]
fts-CERNDASY-submission [SUCCESS]
fts-DESY-bulk-submission [SUCCESS]
fts-DESYCERN-submission [SUCCESS]
fts-DESYDESY-submission-same [SUCCESS]
fts-basic [SUCCESS]
fts-bug23762_info_provider_script_does_not_check_sqlplus_return_code [SUCCESS]
fts-bug25776_properties_xml_incorrect_file_permissions [SUCCESS]
fts-bug32360_glite_sd2cache_error [SUCCESS]
fts-bug35997_fts_cli_if_invalid_value_passed_with_s_option_default_used [SUCCESS]
fts-bug41682_glite-transfer-submit_wrong_handling_of_dest_option [SUCCESS]
fts-bug42579_only_active_channels_shall_be_published_in_bdii [SUCCESS]
fts-bug45167_provide_cli_for_group_management [SUCCESS]
fts-bug46431_transfer_status_should_report_not_found [SUCCESS]
fts-bug52416-ExistsExceptionTreatment [SUCCESS]
fts-bug62394_agent_not_start_pid_left_behind [SUCCESS]
fts-cancel [SUCCESS]
fts-cancel-s [SUCCESS]
fts-channel-add [SUCCESS]
fts-channel-add-s [SUCCESS]
fts-channel-audit [SUCCESS]
fts-channel-audit-s [SUCCESS]
fts-channel-drop [SUCCESS]
fts-channel-drop-s [SUCCESS]
fts-channel-list [SUCCESS]
fts-channel-list-s [SUCCESS]
fts-channel-managers [SUCCESS]
fts-channel-set [SUCCESS]

lfc-cli-delcom [SUCCESS]

lfc-cli-dren [SUCCESS]

lfc-cli-fchmod [FAILED]

```
# /opt/glite/tests/logdm/functional/lfc/lfc-cli-fchmod
```

```
Exit code: 65280
```

```
lxetr-s1564-02.cern.ch
Wed Apr  6 01:39:36 CEST 2011
Test LFC CLI chmod for files
VO=dteam
SE=lxbra2506v1.cern.ch
File=/etc/group
LCG GFAL INFOSYS=lxbra2305.cern.ch:2170
SCENARIO: Directory doesn't exist
Strong check if directory nonexists
List of nonexist directory
OK

1. Change access mode for nonexist directory/file
OK

SCENARIO: check mode for existing file
SHORT=1, MODE=4
Create main directory: /grid/dteam/LFC-1302046657/test-chfile
SE=lxbra2506v1.cern.ch, LFN=/grid/dteam/LFC-1302046657/test-chfile/test file
Create guid and lfn:/grid/dteam/LFC-1302046657/test-chfile/test_file

Noncorrect create GUID, exit
-TEST FAILED-
```

```
/grid/dteam/LFC-1302046657/test-chfile: No such file or directory
/grid/dteam/LFC-1302046657/test-chfile: No such file or directory
/grid/dteam/LFC-1302046657/test-chfile: No such file or directory
Using grid catalog type: lfc
Using grid catalog : lxetr-s1564-02.cern.ch
Checksum type: None
lq_cr: Invalid argument
```

lfc-cli-fren [FAILED]

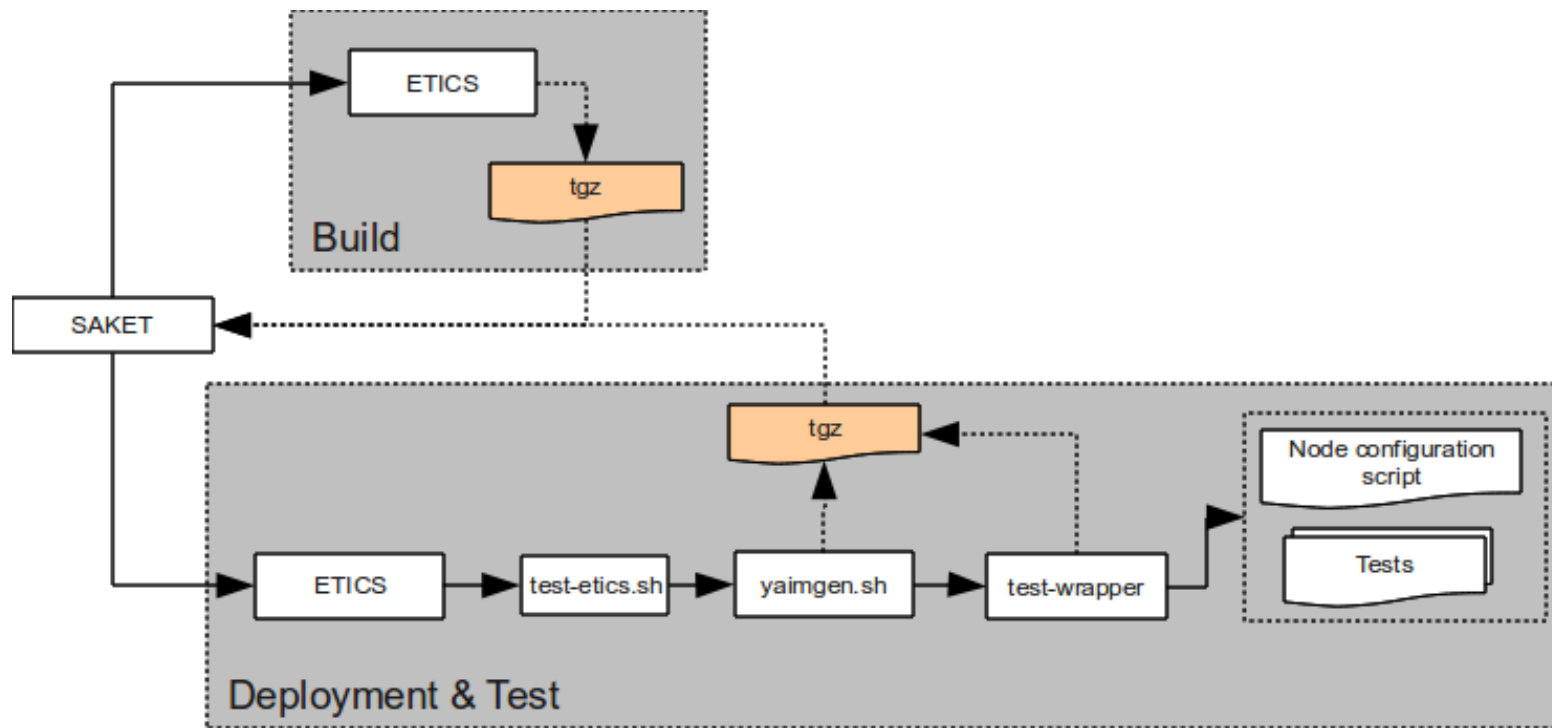
```
# /opt/glite/tests/logdm/functional/lfc/lfc-cli-fren
```

```
Exit code: 65280
```

```
Wed Apr  6 01:39:39 CEST 2011
Test LFC CLI rename file. lfc-rename <oldfilename> <newfilename>
VO=dteam
FILE=/etc/group
SE=lxbra2506v1.cern.ch
LCG GFAL INFOSYS=lxbra2305.cern.ch:2170
Scenario: Directory exists
1. Create main directory: /grid/dteam/LFC-1302046657/renamefile
SCENARIO: Directory and file doesn't exist
Strong check if directory nonexists
List of nonexist directory
OK
```

Workflow

- SAKET
 - Submits build jobs to ETICS
 - ETICS is completely responsible for this
 - If success, the associated tests are sent
 - In ETICS deployed nodes, Yaimgen performs the deployment and calls the test wrapper
 - All the reports and some logs are copied to a location where ETICS retrieves them
 - Parses the build and test logs
 - Submits the summary mail and stores the report in the AFS area



SAKET Configuration

- Builds and tests are defined in a hierarchical structure
 - At project level we can define the project configuration, but a specific component can override it
- Test user password, oracle account...
- Mail recipient, subject,...
- Storage location (if any)
- Timeout

Current combinations

- FTS, FTA, FTM
 - SL4 32 bits
 - SL5 32 and 64 bits
- DPM, LFC
 - SL5 32 and 64 bits
- GFAL/lcg_util
 - SL5 32 and 64 bits
 - Deployment and tests only in 64 bits
- Plus EMI configurations

Current DM workflow

- Changes occur
 - Developers commit changes in the code
 - A new project configuration may be released
- At night, builds and tests are executed
- In the morning, developers and integrators will have a report of the head status
 - If one of the last changes broke something, it will be visible
 - Once the source of the problem is identified, the developer or integrator commits the change
- When a milestone is reached, the head is tagged
 - We already know that version works with the latest project configuration and it passes the tests
 - Certification becomes just a confirmation after locking

Success cases

- Other PT adopted SAKET (BDII)
- BDII 5.1 changed the user to ldap
 - Quickly identified the issue and fixed done in yaim.dpm
- edg-mkgridmap in EMI

Pitfalls

- Occasionally some test-bed machines fail
 - Or repositories, or network connection...
 - It must be stable, or it will make error identification harder
- Limited number of machines
 - Or, to the same effect, high loaded
 - Big issue at the beginning, but not now
 - Thanks to SA2!
 - Still, if more people start doing nightly tests, it may be a problem again
- A deployment process might hang
 - If a test hangs, the wrapper kills it after a timeout, no big deal
 - No equivalent mechanism for deployment steps
 - SAKET will give up after some time, but just a “Execution error” will appear

Future work

- Decouple completely from Yaimgen
 - Working on it
- Multinode support
 - Ready in ETICS
- New platforms support
 - Scientific Linux 6
 - Debian 6 Squeeze
- Upgrade testing
- Automate certification
 - Some bugs still may need manual intervention

Conclusions

- Continuous integration and testing has been used daily for several months by Data Management PT
- It has successfully helped to identify bugs or dependency problems quickly
 - Not only in our own components
- Release candidates and new tests are more reliable at certification time, improving the process

Also, thanks to...

- SA2
 - ETICS, cert-tb-cern
 - Lots of help during the implementation of this process
- Yaimgen developers
- Test maintainers



Thank you

**EMI is partially funded by the European Commission under Grant Agreement
INFSO-RI-261611**