

Feedback on the transition EGEE-EGI middleware rollout

Mail sent by M. David on the 23 Feb and 24 Feb.

To: project-egEE-roc-managers@cern.ch early-adopters@cern.ch

What are the foreseen changes, if any, going from a roc like structure to a NGI structure, in this task in particular?

Since the staged rollout will be conducted inside the production infrastructure, with services properly tagged as beta, so as to distinguish and test those newer version, what issues, if any, do you foresee for the service you are responsible for, or others?

From the operational point of view do you foresee any issues integrating beta services in your production site?

PPS or Early Adopter sites

Dmitry Ozerov
DESY-PPS

Dmitry Ozerov <ozarov@mail.desy.de>

Date: 02/23/2010 11:38:14 AM

Hi,

DESY participate in the staged rollout procedure serving for the glite3.2 LFC_mysql, VOBOX and WMS(once it will be released in glite3.2). Last week we decommissioned the DESY-PPS site and put the lfc(local) and vobox to the production site DESY-HH.

We (site) don't know of any future of this activity in the NGI time, so we do this service certainly till the end of EGEE and if at the end of April we have no agreement with our NGI - we stop it.

Pity that we didn't know about the meeting in Amsterdam, being one of the largest site in EGEE and having long experience running services, it may be that we would bring some inputs to discussion.

Cheers,
Dima.

Michel Jouvin
GRIF

Michel Jouvin <jouvin@lal.in2p3.fr>

Date: 02/24/2010 09:21:01 PM

Mario,

This is a very busy time, sorry for not answering earlier and thanks for the reminder!

Historically I was always very skeptical with PPS because it was more a second certification than a real beta test of new versions because PPS could not emulate the production environment. I always advocated the "early adopter" approach, something we practiced for many years before it became the official process for EGEE.

When the early adopter/staged rollout program was announced I joined it as an enthusiast supporter. I'd say that I have not yet participated to many staged rollout, I think the previous ones beginning of this month was one of the first one... But for me this is exactly the right process and in fact it proved to be quite useful.

In the transition from EGEE to EGI, I think this process must remain an EGI rather than NGI responsibility. There is no reason for each NGI to do the validation on its own and the software quality will improve much quicker if there is a wide-variety of participants.

I am not sure to understand why the ROC to NGI transition will impact the process if this remains a central one. As far as I know, ROCs are not involved in this activity (I personally have not contact with the French ROC about this) and I think it should remain as short a circuit as possible between sites and MU.

I read your doc file, I don't really have the time to react in details but I have few remarks:

- You tend to insist about tagging the "early rolled out" services. Personally I'll argue against that. The whole thing with staged rollout is to put the service in production, really, so that it is used by normal jobs. This is the site responsibility to do it in an ordered manner and to be sure it can roll back the service in case of a severe problem. Their "Quality Level" should be "production" or they will not be selected by production jobs. I'll put more insistance on site selection than service tagging. Participating sites should really demonstrate that their internal process guarantee they will be able to react to potential problems. We have

done this for years at GRIF with many different services including DPM, WMS, BDII, GIP publisher, lcg_util and I don't think any user had to complain about this. Because we have a large team, with a significant expertise, and a tool that allows quick and easy rollback (Quattor).

- About publishing service version with a patch id into the BDII, I have nothing against this if this is done automatically. Service version are published by the info-service-provider based on some RPM versions in general. Every static information put by a site will not be dependable... and probably not be up-to-date.

- I am not sure to understand how a service-based version rather than an overall version number will impact the early adoption/staged rollout. For me this is a non issue.

Based on my short experience with staged rollout, I'd suggest to move tracking of this activity to something "better" than Savannah (even though Savannah can remain the issue tracker). As currently organized this is a bit too fragmented : for each product, there is one Savannah task per site open, making difficult for each site testing the same product to interact. And also making difficult to contribute to something you are not an official tester but have some useful information to provide. A tool more wiki-based or blog-based would probably more appropriate.

Hope this is a useful contribution. Please, feel free to ask if you need some clarification or if there is an important question I missed. Unfortunately, I will not be able to attend the meeting next week as I'm in holidays.

Cheers,

Michel

Eugenia Kovalenko
RU-Moscow-KIAM-PPS

E.Kovalenko <kei@keldysh.ru>
Date: 02/25/2010 09:14:14 AM

Hi Mario,

Our site RU-Moscow-KIAM-PPS participates presently in the staged roll-out procedure serving glite-SE_dpm_mysql and glite-UI for glite3.1.

We do not plan to take part in the new EGI project. So we will support these services up to the end of EGEE-III.

At the end of April we are going to decommission site RU-Moscow-KIAM-PPS as well as our production site RU-Moscow-KIAM-LCG2.

We were pleased to work in PPS team.

Cheers,
Eugenia.

Angela Poschlad
FZK-LCG2

Poschlad, Angela <angela.poschlad@kit.edu>
Date: 02/26/2010 11:55:27 AM

We don't expect any changes in the rollout of Middleware. E.g. we expect to run PPS also in EGI/NGI. PPS is mainly used for pilot services (currently argus) and tests of new updates before putting it into production. So, also updates in staged rollout are first installed in PPS and after successful deployment the next day rolled out into production.

The staged rollout is inherent part of the future setup at KIT/GridKa. The middleware types GridKa currently is interested in are CREAM, ARGUS, SCAS, gLExec and the VOBx.

ROCs

Angela Poschlad
ROC-DECH

The staged rollout has to be strengthened within the region and more sites should participate in it. The aim for NGI-DE is to have (at least) one site for each service supporting the staged rollout. This will be discussed in future meetings within the region. Since the ROC cannot order anything to be done in EGI/NGI we are concentrating to have sites volunteer for this and raise a general interest for this activity.

Tiziana Ferrari
ROC-IT

we don't expect major changes related to the transition from ROC to NGI. Still ITROC is concerned about the small number of sites playing the role of Early Adopters. Currently the Italian sites playing this role are "T3" sites in terms of size (or smaller than this), probably this applies to other ROCs. Generally speaking, I doubt that the current EA infrastructure allows for scalability tests and correctly reflects the fabric setup in production at major sites. So incentives and more commitment are needed in general.

According to our experience, this problem can be partly circumvented by

involving the larger sites in the early phases of the software release lifecycle, for example by involving them in alpha testing of the middleware, especially when major releases are expected. This proved to be particularly valuable for services such as SRM (StoRM in our case), CREAM, WMS. If alpha testing is conducted rigorously with the involvement of production sites, the EA phase is smoother.

About your statement "Presently there are several sites which are in the Pre-Production service, which are expected to become the EA in the production, to perform the stage-rollout.", I'm afraid that the involvement of PPS sites as EA sites won't improve the overall commitment to and quality of the rollout process, as (at least in Italy) just a minor amount of resources is available from PPS. At least in Italy we're undergoing the process of retiring PPS sites and the related services.

- > since the staged rollout will be conducted inside the production
- > infrastructure, with services properly tagged as beta, so as to
- > distinguish and test those newer version, what issues, if any, do you
- > foresee for the service you are responsible for, or others?

This tagging mechanism shouldn't be used by the VOs to redirect their workflows somewhere else (to avoid the usage of the EA nodes), otherwise we will undermine the overall concept of early adoption.

However, it could be handy to:

- make sure that failures of EA nodes do not affect the overall site availability/reliability.
- group EA sites/nodes under a single monitoring (e.g. myEGEE) view to check the overall status of the EA process (still, the EA sites/nodes should continue to be flagged as "production" in GOCDB)

COMMENT TO DOCUMENT. The document says: "For all services: services should be in the production infrastructure, but tagged somehow in the information system and operational tools as "beta"."

Why in the BDII and not the GOCDB? GOCDB would be the ideal place for this kind of flag.

COMMENT TO DOCUMENT. "For this kind of proposal to work, the service version HAS to be properly published in the information system, and there should be a requirement for it's correctness in the validation/verification process.

There is still the issue of how this propagates into the GOCDB for example." I'm not sure why this information should propagate from one place to the other, what is the use case you have in mind? see more comments about this below...

COMMENT TO DOCUMENT. "If the GOCDB gets some of it's data from Gstat" Why GOCDB should take info from gstat... can you clarify? I don't see a manual configuration on GOCDB as a problem. We said that

most of the EA infrastructure should be stable, didn't we?

In any case, I think we should just have a single source of authoritative information, without any propagation of info from one tool/service to the other. I see GOCDB as a simpler way to go right now, while we will probably have to wait for GLUE2.0 in production.

> I would like to know if/which of you will be present in the Amsterdam
> meetings

unfortunately not, the original invitation forwarded by Maite didn't advertise the middleware meeting, and it was unclear to me if this meeting were open or not, and I booked my flights before you sent your mail.

Hope this helps. Thanks for kicking this off.

Tiziana