



EGI-InSPIRE

DEPLOYING SOFTWARE INTO THE EGI PRODUCTION INFRASTRUCTURE

EU MILESTONE: MS409

Document identifier:	EGI-MS409-478-v7.docx
Date:	05/07/2011
Activity:	SA1
Lead Partner:	LIP
Document Status:	Draft
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/478

Abstract

This document gives an update on the processes and procedures of software provisioning into the EGI production infrastructure. This includes the EGI Unified Middleware Distribution (UMD) and Operational Tools provided by JRA1 activity. The UMD will be composed by middleware components from the EMI project, namely ARC, dCache, gLite and UNICORE, plus Globus from the Initiative for Globus in Europe (IGE) project as well as other software components provided by the community.

I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Mário David	LIP/SA1	27 April 2011
Reviewed by	Moderator: Reviewers:		
Approved by	AMB & PMB		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	27 April 2011	TOC, references	Mário David/LIP
2	30 May 2011	Almost complete draft	Mário David/LIP
3	30 May 2011	Comments from Michaela Barth	Michaela Barth/KTH
4	1 June 2011	1 st complete draft	Mário David/LIP
5	12 June 2011	SA1 Internal review and comments	Mário David/LIP
6	4 July 2011	Reviewers comments	Aleš Křenek/CESNET, Mark Santcroos/UvA
7	11 July 2011	Slight additions: Reviewers comments	Mário David/LIP

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed: <https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>.



VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting ‘grids’ of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today’s production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

VIII. EXECUTIVE SUMMARY

This document describes the processes and procedures that have been put in place during the first project year, to provision the software for the European Grid Infrastructure. It updates the MS402 [R1] document.



The major categories of software in use or expected to be used by EGI-InSPIRE are given below:

1. UMD provided by EGI-InSPIRE is composed of the following sets of components:
 - a. EMI release containing gLite, ARC, UNICORE and dCache middleware components.
 - b. Globus middleware provided by IGE.
 - c. Trust Anchor packages containing the Certification Authority root certificates, provided by EGI-InSPIRE.
2. Operational Tools provided by EGI-InSPIRE JRA1 activity: Regional Dashboard, SAM/Nagios monitoring tools, accounting portal, etc. .

The staged rollout process that was in place in the EGEE project was targeted at the gLite components. This process has been adapted and extended to take into account other middleware components and the Operational Tools that are in use in EGI.

This document is organized as follows. Section 1 gives a brief introduction to the software provisioning to EGI. Section 2 describes the staged rollout during the first project year as well as the designed workflow. Section 3 describes the Early Adopters concept. In section 4, some metrics for the first project year, are presented. Section 5 describes the communication channels between the several involved parties. Section 6 gives an overview of the objectives for the second project year and in section 7 some conclusions are drawn.



TABLE OF CONTENTS

1	INTRODUCTION	6
2	STAGED ROLLOUT PROCESS	7
2.1	First project year	7
2.2	Software capabilities and interoperability between different components	8
2.3	Software versioning scheme	9
2.4	Tools supporting the Staged Rollout	10
2.5	Workflow and process management	10
2.6	Naming conventions	13
3	EARLY ADOPTERS	14
3.1	Tools to manage Early Adopter teams	14
4	METRICS	16
4.1	First project year results	16
5	COMMUNICATION CHANNELS	19
5.1	Case 1: External Technology Provider (currently EMI and IGE)	19
5.2	Case 2: Internal Technology Providers	19
5.3	Case 3: EGI members - SA2, TSA1.3 and Early Adopter teams	19
6	YEAR TWO ROADMAP	21
7	CONCLUSIONS	22
8	REFERENCES	23



1 INTRODUCTION

The software provision to the European Grid Infrastructure (EGI) comprises two processes:

1. Verification of Quality Criteria
2. Staged Rollout

The first process is described in MS508 [R4] and is handled by EGI TSA2.3 task. The description of the second process is the aim of this document.

The staged-rollout process was adapted from the EGEE-III with the objective to ensure that middleware in the EGI is swiftly and safely deployed in the production infrastructure.

In EGI, the staged-rollout process has been extended from process that was in place in the EGEE project for gLite components, to a wider variety of software components and products from different Technology Providers.

The main software components currently deployed in EGI are from gLite, dCache and ARC. The European Middleware Initiative (EMI) will provide UNICORE components as well and the Initiative for Globus in Europe (IGE) will provide Globus components.

As such, a wide variety of deployed middleware will increase the heterogeneity of the infrastructure in the short term, including the different middleware stacks. The aim of EMI will be to partially harmonize gLite, dCache, ARC and UNICORE components in the medium term.

The operational tools used in the infrastructure will follow a similar procedure to the middleware components, and will be the main driver for the integration and interoperability of the different middleware components.

This document is targeted primarily to Early Adopter teams of software components and also to all other resource infrastructure administrators including the EGI Operation Unit (SA1). The process and workflow has been developed in close collaboration with EGI SA2. Other interested parties are the Technology Providers.



2 STAGED ROLLOUT PROCESS

The EGI staged rollout is a procedure by which certified and verified updates of the supported middleware, are first released to and tested by Early Adopter sites before being made available to all sites through the production repositories.

The certification step is performed by the Technology Provider and consists typically of several checks from packaging policies to deployment in a controlled test infrastructure, and some level of testing.

The verification phase is performed by EGI task TSA2.3 following a previously approved Quality Criteria definition, between the EGI task TSA2.2 and Technology Providers. It consists in checking a set of generic criteria, such as available and updated documentation to specific criteria that depends on the capability under test. It may include deployment in a small test infrastructure, and some level of testing.

The staged rollout procedure goes a step forward. It permits testing an update in a production environment that is also more heterogeneous than what is possible during the certification and verification phases. It allows for potential issues with an update to be discovered at a small scale and potential workarounds to be added to the release notes if need be. In some cases an update may be rejected.

The staged rollout serves to increase the confidence in the quality of the updates that pass this test, such that the vast majority of the sites should experience smooth updates. Resource Infrastructure Providers (sites), participate in the staged rollout for services that they have a particular interest in, with the proviso that they may need to debug issues with a particular update and are required to report their findings.

The description of the New Software Release Workflow is detailed in MS508 [R4] (NSRW) includes all stages from software release by a Technology Provider to a full release for deployment in the EGI production infrastructure.

The remainder of this document describes the part of the workflow corresponding to the staged rollout phase.

2.1 *First project year*

The staged rollout process has had a smooth evolution from the one implemented during the EGEE project. It involved the change of EGEE procedures into the new software release workflow, using the tools provided within EGI and deprecating the ones previously used.

One of the main points during this transition was to have a gradual and smooth transition in the procedures that the Early Adopter sites have to follow. On the other hand, the transition has had a larger impact on the coordination and management of the whole process, e.g. partially using the tools and processes inherited from EGEE together with tools and processes/workflows devised in EGI.

In the first iterations all involved parties (EGI Operation Unit and Technology Providers), agreed that Technology Providers should not interact with the EGI RT system. As such, the initial workflow [R1] was re-designed so that the process starts with a Global Grid User Support ticket (GGUS)¹ that is opened by the external Technology Provider, followed by all steps described in the next sections.

¹ The Global Grid User Support is the EGI helpdesk system.



Therefore all communications between EGI and Technology Providers are carried through GGUS tickets (see section 5.1), while allowing a single well determined point of exchange of information, public availability, traceability and easier extraction of metrics. This well determined communication channel does not exclude other means of more informal communication, such as e-mail or private communications.

Although some software components already went through the EGI Software Provision workflow, the gLite middleware components still go through the old process, i.e. using the CERN Savannah patches (<https://savannah.cern.ch/patch/?group=jra1mdw>) to track the verification and staged rollout phases.

Although gLite 3.1 and 3.2 components are now supported by EMI until their end of life, it is not foreseen to adopt the new procedure for those components. Only the components that are part of new EMI and IGE releases are foreseen to undergo the new workflow.

At the time of writing, EMI and IGE have done the first release and several components are under the verification process [R4]. After this stage, all accepted software components will undergo the staged rollout process.

2.2 Software capabilities and interoperability between different components

The EGI infrastructure will have a large set of different middleware components deployed within it.

The main players (EGI, EMI, IGE) have defined the concept of software capabilities depending on the product used, targeting the specificity of each component.

Due to the high heterogeneity of Unified Middleware Distribution software, the components have been grouped in nine major categories as defined in the “UMD Roadmap” [R16]:

1. **Information Capabilities:** Grid information model schema, messaging and service discovery capabilities.
2. **Compute Capabilities:** Products for job execution (parallel, sequential and interactive), job scheduling and workflows.
3. **Storage Capabilities:** Products that target data transfer and storage capabilities as well as file encryption/decryption, file access and management.
4. **Data Capabilities:** Product for data management, data and metadata catalogs.
5. **Instrumentation Capabilities:** Products for remote instrumentation management, access and monitoring.
6. **Security Capabilities:** Grid authentication, authorization, access and Virtual Organizations management capabilities.
7. **Virtualization Capabilities:** Products for Virtual machine management, deployment and distribution.
8. **Operations Capabilities:** Products for monitoring and accounting capabilities.
9. **Client Capabilities:** Client Application Programming Interface and Command Line Interface.

Each software product is in general composed of several capabilities. As such, it is the responsibility of the EGI-InSPIRE Technology Unit, to provide requirements about any given capability. For example, the Compute capability should have integration to several Local Resource Management Systems (LRMS), with several parallel programming environments, etc..

Since a given capability may be provided by more than one product (or by more than one Technology Provider), the infrastructure should allow to accommodate any of such products in a seamless way. We can take as an example the Compute Capability, there are three EMI products offering this capability (the ARC-CE, CREAM and UNICORE), while GRAM is provided by IGE, as such any of these products should be deployable in the infrastructure and be monitored (SAM/Nagios framework), publish accounting records, properly registered in the GOCDB, etc..



Regarding the interoperability and integration of products with origin in different middleware stacks, the present status is described next.

Presently, as a result of the EGEE and WLCG projects, only gLite is fully integrated into all the operational tools, whilst ARC has been fully integrated into all official EGI provided operational tools, but still needs to be better supported in other operational tools like gridmap or MyEGI. Globus and UNICORE operational integration is still ongoing. Comprehensive integration is a short-term objective of the project.

In a second phase, it is expected that site administrators and user communities will provide requirements for the interoperability between different middleware stacks, and that the EGI infrastructure will be integrated with new types of resources, such as digital libraries and repositories, desktop grids, High Performance Computing, etc.

2.3 Software versioning scheme

The Unified Middleware Distribution (UMD) software versioning scheme can be categorized as follows [R1]:

1. **Emergency release** (when needed): it fixes critical functionality problems and/or serious security vulnerabilities. It is backwards compatible.
2. **Revision release** (at most once every two weeks): it provides bug fixes and is backwards compatible.
3. **Minor release** (at most once per month): it provides new functionality and is backwards compatible.
4. **Major release** (at most once per year): It offers new functionality, but is not necessarily backwards compatible. It may include new services.

A given UMD major release will contain baseline major versions of a set of components. These major versions are subject to the agreement between EGI and the Technology Providers.

EGI-InSPIRE defines the schedules of all types of UMD releases. Technology Providers may release new versions of their software components at their pace, as such EGI-InSPIRE may buffer several software components to suit a previously agreed schedule for those components to undergo the Software Provisioning process and later its release into the production infrastructure.

Any middleware component can be updated only up to a minor release within any major UMD release. The major releases of any given component may only be included in the next major UMD release, depending on the roadmap.

It is foreseen that all categories of component updates will undergo the staged-rollout process, but the time-lines and the extensiveness of the staged-rollout will vary according to the category. The sole possible exception is an emergency release, which may skip staged-rollout under exceptionally well documented circumstances that are evaluated on a case by case basis.

The components classified as “Community Contributions” and “Operational Tools” will follow a similar procedure.

EGI-InSPIRE will accept only certified and validated updates provided by the software providers. The validated components will undergo the staged-rollout procedure. If successful, they can then be widely deployed into the production infrastructure.

If bugs or issues are found during the staged-rollout phase in a given component for which some solution or workaround is proposed, the fix(es) should be communicated and implemented by the respective Technology Provider.

2.4 Tools supporting the Staged Rollout

This section describes the tools and their use in the Staged Rollout process:

- EGI Single Sign On (SSO) - <https://www.egi.eu/sso/> : User authentication and authorization for several EGI tools and web applications. Also hold several sets of groups for specific purposes:
 - One group per Early Adopter team members, see section 2.4: Naming Conventions.
 - **early-adopters-admins**: group of managers of the Early Adopter teams.
 - **early-adopters-<arc|glite|globus|unicore|opstools>**: groups of all Early Adopters members.
 - **staged-rollout**: group of Staged Rollout managers.
- EGI Mail manager - <https://mailman.egi.eu/mailman/listinfo> : mailing lists, the set of people in each mailing list are taken from the SSO groups:
 - **early-adopters-arc@mailman.egi.eu**: for ARC early adopter team members.
 - **early-adopters-glite@mailman.egi.eu**: for gLite early adopter team members.
 - **early-adopters-globus@mailman.egi.eu**: for Globus early adopter team members.
 - **early-adopters-opstools@mailman.egi.eu**: for Ops tools early adopter team members.
 - **early-adopters-unicore@mailman.egi.eu**: for UNICORE early adopter team members.
- **EGI Request Tracker** [R11] - <https://rt.egi.eu/rt/index.html>: A queue called “**staged-rollout**” has been created to track through tickets the various stages of the Staged Rollout process.
- **EGI Wiki** [R8] - <https://wiki.egi.eu/wiki/Staged-rollout-procedures>: Staged Rollout process description and workflow.
- **EGI Repositories** [R13] - <http://repository.egi.eu/>: provide access to the software packages that are part of the UMD distributions, during the various stages of the software lifecycle. Details are given in [R5].
- **EGI Document Database** [R12] - <https://documents.egi.eu/public/DocumentDatabase>: Document repository where Staged Rollout reports and summaries are stored, and publicly available.
- **Early Adopters Web Application** [R9] - <https://www.egi.eu/earlyAdopters/table>: web application to manage Early Adopter teams, creation, modification, assignment to software components, etc..
- **Software components table** [R10] - <https://www.egi.eu/earlyAdopters/teams>: table with all software components and the number of Early Adopter team committed to each component.

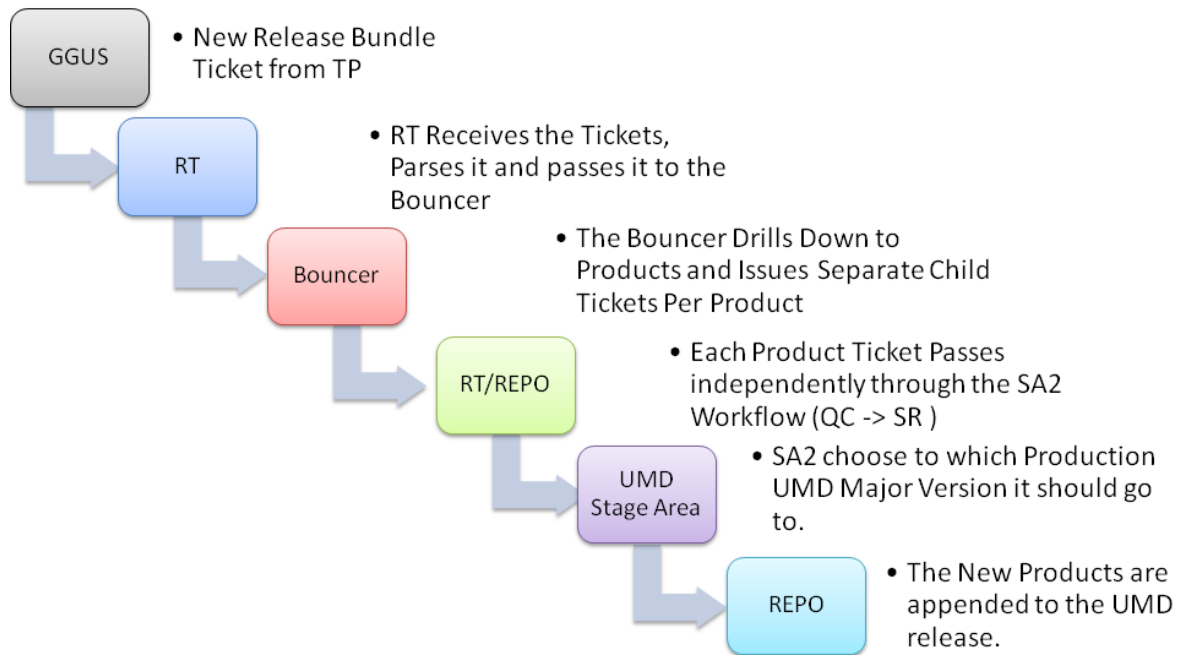
2.5 Workflow and process management

The Software Provisioning activity in EGI [R4] governs two processes to ensure software quality and its correct integration into EGI UMD repository acting as a filter to accept or to reject new middleware products or components into EGI's UMD repository.

The Software Provisioning workflow is shown in Figure 1. It starts when the Technology Provider makes a new software component release available and creates a ticket in the Global Grid User Support system (GGUS).

The GGUS ticket is filled with complete information about the new software release such as: release notes, documentation, installation and configuration notes and the list of packages contained in the release.

Figure 1: Software release workflow implemented in the EGI RT and repository portal.



The ticket is assigned to the EGI Technology Unit (SA2), triggering the creation of a ticket in the EGI RT system [R11], in a dedicated queue named “**sw-rel**”, and is the main tool to handle the Software provisioning workflow.

Upon the RT ticket creation, the Software Delivery phase (implemented by an external RT module called “**Bouncer**”) takes place and the new Technology Provider release is processed automatically. The end result of the “**Bouncer**” is a list of product release descriptions specific for each platform and architecture (*Product-Platform-Architecture: PPA*), creating one child ticket (in the RT queue “**sw-rel**”) per combination of Product-Platform-Architecture.

Each Product-Platform-Architecture combination is tracked in it's own child ticket during the Verification and Stage Rollout phases. As such, several PPA combinations may be processed in parallel and independently from each other. The RT ticket state changes from “**Unverified**” to “**InVerification**” when the Verification process starts.

If the outcome of the Verification phase is “**Accepted**”, the RT ticket state changes from “**InVerification**” to “**StagedRollOut**”. This state transition triggers the start of the staged rollout phase with the creation of a child ticket in another specialized RT queue named “**staged-rollout**”.

It is assumed that the staged rollout repository has a fixed URL for any given major release of any Technology Provider (EMI, IGE), This is publicly available in the repository and EGI wiki for Early Adopters. This state change triggers the following actions, in bold are the actors that perform each of the actions:

1. The software packages are moved from the “**InVerification**” repository to the “**StageRollOut**” repository.
2. The “**sw-rel**” ticket creates a new child ticket in the “**staged-rollout**” queue. All the Staged Rollout process occurs in this queue/ticket.
 - 2.1. It is advisable to change to the RT ticket tab “**Jumbo**”.
 - 2.2. The “**Owner**” of the “**staged-rollout**” ticket is “**Nobody**”.
 - 2.3. The “**Status**” of the ticket is “**New**”.
 - 2.4. Change the ownership of the ticket to the “**Staged Rollout Manager**” responsible for that MW component: ARC, gLite (including dCache), UNICORE, Globus, Operational Tools.
 - 2.5. (**Staged Rollout Manager**) The staged rollout manager is notified and takes the actions described below.
 - 2.5.1. The staged-rollout ticket should contain the links to:
 - 2.5.1.1. Release notes, contains information about install/upgrade and configuration.
 - 2.5.1.2. Bugs or issues fixed.
 - 2.5.1.3. Documentation
 - 2.5.1.4. Staged rollout repository (needed by the Early Adopter teams).
 - 2.6. The status of the ticket is “**Open**”.
3. (**Staged Rollout Manager**) The queue “**staged-rollout**” has a custom field (dropped down list) where they can select all the Early Adopter teams committed to that staged rollout test.
4. (**Staged Rollout Manager**) When the “**Save changes**” button is pressed an automatic notification mail is sent to all Early Adopter teams. The mail notification should contain the links to the release notes, installation and/or upgrade information, configuration, documentation and the repository URL. All Early Adopter teams are added to the “**Administrative Cc**” field.
5. (**EA teams**) Each Early Adopter team has to acknowledge the reception of the notification within 1 working day. Either by replying to the mail sent by the RT or directly in the ticket with: `<accept|reject> <NGI>-<Site-name>`.
6. (**Staged Rollout Manager**) Will check the ticket and if there are no Early Adopters accepting the staged rollout test, it will pool the “**early-adopters-XXX.mailman.egi.eu**” mailing lists, and other (s)he sees fit to get other Early Adopter teams to perform the test.
7. (**EA teams**) The Early Adopter team has the option to put the service node into downtime in the Grid Operation Centre DataBase (GOCDB), special “**beta**” tag. This tag may be set only during the staged rollout phase. If/when the component is release into production, this tag (or downtime), should be removed from the GOCDB.
8. (**EA teams**) performs the staged rollout: installation or upgrade, configuration, and some tests as they see fit.
9. (**EA teams**) If the Early Adopter team finds problems or issues, the staged rollout managers and other Early Adopter teams can clarify them within the ticket or, a GGUS ticket should be opened. If a GGUS ticket is opened, this ticket should be assigned to the DMSU support unit, which will then route it to the Technology Providers as they see fit.
10. (**EA teams**) A link to the GGUS ticket should be created in the staged rollout ticket.
11. (**EA teams**) The service should ideally be exposed to production load/environment and users. This period lasts between 5 to 7 days, but may be extended or shortened depending on the cases or components under test.
12. (**EA teams**) Each team should fill the staged rollout report after the previous point, and send it to the “**Staged Rollout Managers**”:
13. (**EA teams**) The report should contain as much information as possible. More specifically the correctness of the release notes, tests that have been preformed, and possible metrics when the service is exposed to production (like number of jobs per day, or number of transfers, what VOs

are configured for that service, etc.). The name of the file should follow the naming conventions described in section 2.4.

14. (**Staged Rollout Manager**) The staged rollout managers:
 - 14.1. Collect all reports and produce a summary report.
 - 14.2. Create a document in the EGI Document DataBase server with a given ID, which will contain all reports for the staged rollout of that component, as well as the summary report.
 - 14.3. Insert a link in the RT ticket to the Document server ID of the reports.
 - 14.4. Select from the custom field “**Outcome**” the outcome: < **accept** | **reject** >.
 - 14.5. Set the ticket status to “**resolved**”.
15. (**Staged Rollout Manager**) When the staged rollout child ticket is set to “**resolved**” the parent ticket is notified, and it gets the outcome and the document server ID containing the reports.

2.6 Naming conventions

The Early Adopter team names are SSO groups containing the members of each team:

- **ea-<NGI>-<Site-name>**

The custom fields provided in the “**staged-rollout**” queue are:

- Drop down list containing all Early Adopter teams, on the EGI SSO: possibility to select several teams.
- The button “**Save Changes**” notifies those teams and they will be added to the “**Administrative CC**” field.
- Outcome of the staged rollout, drop down box with: <**ACCEPT**|**REJECT**>.

The title of the ticket is of the form:

- **Staged Rollout <SW stack-MajorVersion> <COMPONENT> <VERSION> <OS> <ARCH>**
- Examples:
 - **Staged Rollout CA 1.38**
 - **Staged Rollout EMI-1.0 SE-DPM_mysql 1.8.0-1 SL5 x86_64**

Two persons per main software stack should ideally compose the staged rollout managers group:

- gLite (including dCache).
- ARC.
- UNICORE.
- Globus.
- Operational Tools.

The name of the file containing the staged rollout report should have the form:

- **ea-<NGI>-<Site-name>-<MW stack>-<component>-<version>.doc(odt)**

The file name containing the summary of the staged rollout of a given component should have the form:

- **summary-<MW stack>-<component>-<version>.doc(odt)**

All reports and the summary will be uploaded to <https://documents.egi.eu> with a given document ID that will be referenced in the respective RT ticket.

The description in the document database should have the following naming convention:

- **Staged rollout <MW stack> <component> <version>**



3 EARLY ADOPTERS

An Early Adopter (EA) is a production site committed to perform the staged-rollout test for one or more software products. In general, the Early Adopter sites deploy these new versions of the software in their production services.

The deployment layout is always a final decision of the site managers and depends on the specificity of the service at the site.

It is assumed that any software component that was accepted in the verification phase by the EGI Technology Unit (SA2) has production quality and can be safely deployed in production without disruption of the service it provides. Thus the Early Adopter sites are just the first Resource Infrastructure Providers to be notified and to deploy a new version of a software component. They are also requested to report back about the overall process outcome.

The Early Adopter site manager defines the deployment configuration that is applicable to the respective site. It must be taken into account that the site “Reliability and Availability” should not be affected in case of problems with the staged rollout components. For that purpose, a new special tag called “beta” can be set for that service during the staged-rollout period.

3.1 Tools to manage Early Adopter teams

The management of Early Adopter teams, and their commitments to software components was initially performed in a wiki page. This has become unmanageable with the steady grow of new teams, as such a web application has been developed and deployed by the EGI IT team under the EGI.eu domain².

Figure 2 shows the management web application for Early Adopter teams. The application is

Figure 2: Early Adopters management application.

NGI, site	contacts	component	software	os/arch	
Australia, Australia-ATLAS	Tim Dyce Tom Fifield	BDII-site	Glite 3.2	SL5/64	manage
		BDII-top	Glite 3.2	SL5/64	
		APEL	Glite 3.2	SL5/64	
		AMGA	Glite 3.2	SL5/64	
Canada, CA-McGill-CLUMEQ-T2	Simon K. Nderitu	STORM	Glite 3.2	SL5/64	manage
CERN, CERN	Steve Traylen	VOMS oracle	Glite 3.2	SL5/64	manage
		FTS/FTM/FTA	Glite 3.2	SL5/64	
NGI AEGIS, AEGIS01-IPB-SCL	Antun Balaz Dusan Vudragovic	CREAM	Glite 3.2	SL5/64	manage
		Torque Utils	Glite 3.2	SL5/64	



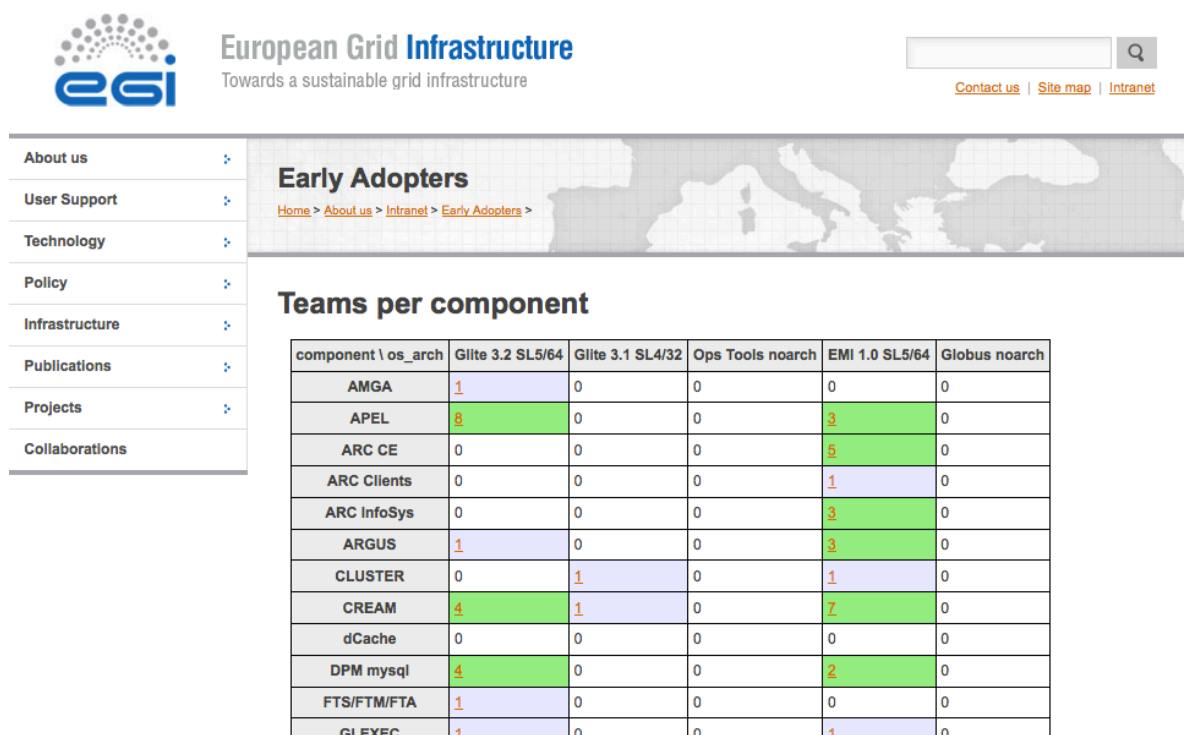
interfaced with the EGI SSO, see section 2.6. It permits the creation and modification of Early Adopter teams and their membership as well as the software components and its assignment to the teams.

The web application displays the following information for each Early Adopter team:

- Their National or Regional Grid Initiative – NGI.
- The site name.
- The team membership.
- Their e-mail contacts.
- The list of components they have committed for the staged rollout test, along with the corresponding software stack (Technology Provider), operating system and architecture.

A second web application, shown in Figure 3, was also deployed and interfaced to the one described previously, containing the list of software components and showing in each cell the number of Early Adopter teams committed to each component.

Figure 3: Software components and number of Early Adopter teams committed to each component.



² The EGI.eu domain is hosted at CESNET, the Czech academic network operator.

4 METRICS

This section reports several metrics. The time span is the first EGI-InSPIRE project year: May 2010 – April 2011. The metrics definition and results have been taken from the SA1 quarterly reports.

The set of metrics and a brief description is given next:

1. **M.SA1.ServiceValidation.1**: Total number of component that have been tested and the number that have been rejected during staged rollout.
2. **M.SA1.ServiceValidation.2**: Number of staged rollout tests performed by all Early Adopter teams combined.
3. **M.SA1.ServiceValidation.3**: The total number of Early Adopters teams.

4.1 First project year results

The results of the first metric “**M.SA1.ServiceValidation.1**” (shown in Table 1 and Figure 4), are the number of components that underwent the staged rollout test and that where accepted or rejected, per project quarter. It is also shown the number of components that have been passed into production without the staged rollout test.

Table 1: Number of components tested in staged rollout per project quarter, (Accepted, Rejected, No test). Period 1 May 2010 -- 30 April 2011.

Project Quarter	Number of components		
	Accepted	Rejected	No test
Q1	9	1	7
Q2	8	2	7
Q3	9	2	0
Q4	26	2	0
Total	52	7	14

It can be seen that the number of rejected components is low (1 or 2) and constant throughout the reported period, in particular it did not increase in the fourth quarter when the number of tested components more than doubled. The objective is that, no components are rejected. The number of components that go into the process varies with the SW providers releases and updates.

The number of accepted components is also constant throughout the first three quarters, and a large increase in the last quarter. This was due to the approach of the EMI release, where several gLite 3.1 and 3.2 patches have been finalized and released. A total of 59 components have been tested.

The last column shows the number of components that have been set into production without the staged rollout test. The value in the two first quarters was 7 components each, from which only one corresponded to a high priority update. During the third quarter, a decision has been agreed between EGI and the gLite consortium, not to release components that do not undergo the staged rollout test, thus, the values of zero for the last two project quarters.

The second metric (**M.SA1.ServiceValidation.2**) is the number of staged rollout tests undertaken by each NGI per project quarter, is shown in Table 2 (and in Figure 5). For each NGI there maybe one or more Early Adopter teams, so the integrated result is taken. The number of components that go into the process varies with the SW providers releases and updates.

It can be seen that the number of staged rollout tests exceed that of the number of tested components, this is due to the fact that some of the components were tested by more than one Early Adopter team.

Table 2: Number of staged rollout tests per NGI and per Project Quarter. Also shown the total number and percentage.

NGI	Total %	Total	Q1	Q2	Q3	Q4
Australia	2.5	2	0	1	1	0
Canada	1.3	1	0	0	0	1
CERN	5.0	4	1	0	1	2
Croatia	2.5	2	0	1	0	1
France	2.5	2	1	1	0	0
Germany	8.8	7	3	2	1	1
Greece	3.8	3	0	0	0	3
Italy	16.3	13	3	2	2	6
Macedonia	1.3	1	0	0	1	0
Netherlands	2.5	2	0	0	0	2
Poland	5.0	4	0	1	0	3
Portugal	13.8	11	0	4	1	6
Serbia	3.8	3	1	1	0	1
Slovenia	8.8	7	1	1	2	3
Spain	12.5	10	2	0	2	6
Switzerland	3.8	3	0	0	1	2
UK	6.3	5	0	0	2	3
Total	100	80	12	14	14	40

It should be noted that Australia and Canada are not partners of EGI, but the Early Adopter teams (one for each country) are WLCG Tier2 sites. All other NGIs, except CERN, are formally funded for the TSA1.3 activity.

Figure 5 shows a bar chart with the number of staged rollout tests per NGI and accumulated for the four project quarters. The staged rollout test of any given component is independent from the test of any other component, thus in several occasions there were several components under test as well as in different stages of the process.

The third metric (**M.SA1.ServiceValidation.3**), quantifying the evolution of the number of Early Adopter teams is shown in Table 1. The number was not recorded for the second project quarter. It shows a steady increase from 26 to 45 Early Adopter teams during the first year.

Table 3: Number of Early Adopter teams at the end of each project quarter.

	Project Quarter			
	Q1	Q2	Q3	Q4
Number of EAs	26	NA	40	45

Additionally we report the coverage of components by Early Adopter teams. The present status is as follows:

- A total of 37 EMI products, 30 are covered and 7 missing Early Adopter teams.
- A total of 26 gLite 3.2 and 3.1 products, 22 are covered and 4 missing Early Adopter teams.
- Operational Tools, has 1 component and is covered by 1 Early Adopter team.

A direct target estimate is not given, except that it should be as many as possible. There is an indirect target that may be taken from this metric, which is the full coverage with redundancy of all components. TSA1.3 has made a recent call so that some Early Adopter teams move from gLite to EMI components in view of the recently released EMI 1.0.

Figure 4: Number of components "Accepted", "Rejected" "Not tested" in staged rollout per project quarter.





Figure 5: Number of staged rollout tests per NGI and per quarter.



5 COMMUNICATION CHANNELS

There are several processes and players in the overall Software Provisioning workflow. This section will detail for all those processes the communication channels that are or will be used.

5.1 Case 1: External Technology Provider (currently EMI and IGE).

The EGI Software Provisioning workflow starts when the Technology Provider opens a GGUS ticket when there is a new release of their software. This GGUS ticket will create an RT ticket in the queue “sw-rel”. The acceptance or rejection of a given release of the software is communicated by EGI in the RT ticket that will be feed into the original GGUS ticket, thus the Technology Provider is notified of the outcome of the process. If the release is rejected, a description is given in the GGUS ticket coming from the RT ticket.

During each of the two main phases of the EGI Software Provisioning workflow: Verification and Staged Rollout, any problems or issues that are found in the new software release are reported in GGUS tickets. These tickets should be record as at least “High Priority” and assigned to the Technology Provider by the second level support team called “Deployed Middleware Support Unit” (DMSU) [R3]. The GGUS ticket URL should be referenced in the RT ticket.

Depending on the problem reported, the new software release may be rejected, or the Technology Provider may produce workarounds that are not showstoppers for the release. The release may be accepted for General Availability (into production) given the known issues are published and publicised in the final release notes.

5.2 Case 2: Internal Technology Providers

The EGI Software Provisioning workflow in the case of Internal Technology Providers, such as the SAM/Nagios monitoring framework (in the category of Operational Tools), and the EGI Trust Anchor packages [R15], start with a creation of a RT ticket in the “sw-rel” queue when a new release is available. In this case GGUS is by-passed and all communication between Verifiers and after Staged Rollout Early Adopter teams, and the Technology Providers occur inside the RT ticket.

Nonetheless, if bugs or problems affecting the new software release are found, they should be recorded in GGUS tickets, and referenced in the RT ticket.

As in the case of external software, the release may be rejected and a description will be recorded in the RT ticket.

5.3 Case 3: EGI members - SA2, TSA1.3 and Early Adopter teams

The communication between the Early Adopter teams, the “staged-rollout” managers and the EGI Technology Unit (SA2), is expected to occur within the RT ticket, during the whole process of the Software Provisioning workflow, until the component is released into the production infrastructure.

There might be cases where the release is approved for production, even with some workaround that have to be documented in a known issues page as well as in the final release notes.

After the new software release gains General Availability, all problems and issues found in the production infrastructure should be reported through GGUS tickets. These tickets will be routed to the second level DMSU support team.

The announcement of any given new version of a software component for production, is done through the operations portal broadcast tool to the following groups or mailing lists:

1. Production Site Administrators (taken from the GOCDB).



2. WLCG Tier-1 contacts (static CERN maintained list).
3. inspire-sa2@mailman.egi.eu: EGI SA2 team including the repository administrators, the DMSU support unit, the staff responsible of defining the middleware quality criteria and of their verification.
4. NGI Operations Centre managers: noc-managers@mailman.egi.eu .

An RSS feed is foreseen to be implemented, where site administrators and other parties may subscribe to received news/notifications of coming releases.



6 YEAR TWO ROADMAP

At the time of writing, EMI and IGE have just made their first release of the middleware. EGI SA2 together with TSA1.3 have performed dry runs of the overall workflow testing the RT queue “**sw-rel**” and “**staged-rollout**” as well as the movement of packages through the EGI repositories. A few EMI software components were used in these dry runs.

The tests have permitted to find weak points and issues that needed to be mitigated, both on the side of the EGI RT as well as in the EMI Release Candidate (RC3) components. A complete report documents the issues and problems that have been found, see [R17]. On the other hand, the real deployment and configuration of components was not performed.

The SA2 verification process is currently ongoing for a large set of EMI components, so the real challenge is to verify if the system that has been design and implemented for most of the first project year is robust enough to support the full workflow with dozens of components in parallel.

The highest priority in the following weeks is to solve some of the issues that have already been found in the real usage of the workflow with respect to the verification phase. The staged rollout process follows, and the primary goal besides having the software components tested, are to have a detailed account of all the problems that may arise in this phase and mitigate them. The Early Adopter teams have an important contribution to give in order to improve the process where needed.

Another important point is to have full coverage, redundancy and more then one deployment scenario of all components both of EMI and IGE, as well as to have staged rollout managers for missing major middleware stacks; UNICORE and Globus. The objective is to have a minimum of two Early Adopter teams per component.

Involvement of the user communities and Virtual Organizations is also an important aspect. Users are involved to some extent in the development and testing phases of some components such as compute elements, storage elements and a few other capabilities, in close collaboration with the Technology Providers.

Since the staged rollout test is preformed mainly in production services, users and Virtual Organizations are naturally involved in this phase. A further step is to pool a wider range of user communities through the User Community Board (UCB), for a deeper involvement in this phase and to bring the sites that support them to be Early Adopters.

Last but not least, gLite 3.2 components are supported until the release of EMI2.0 in April 2012, as such several Early Adopter teams have not moved their commitment and will continue to test gLite 3.2 components.



7 CONCLUSIONS

During the first project year, TSA1.3 task had two major challenges: To keep the staged rollout process, inherited from the EGEE III project, smoothly running while designing and implementing the new software provisioning workflow for EGI.

The first objective has been achieved, and the second one is mostly accomplished. The only thing lacking is its proof in the real “production run”.

The number of Early Adopter teams has increased steadily and the coverage of EMI and IGE software components is almost complete.

The Early Adopter teams are generally responsive and perform the staged rollout tests in time. Issues and problems found by the Early Adopter teams are readily answered by the Technology Providers, and in some cases, software components have been rejected. In general, the Technology Provider suggests the rejection of a given component in the case of a more serious issue or bug, and in most cases a new version is produced in a short timeframe (a few days).

The new software provision workflow has been used by some internal Technology Providers (SAM/Nagios framework and Trust Anchor distribution), even if it was not yet completely implemented. This has permitted to partially test the process. On the other hand, dry runs with some early EMI software components, allowed the identification of further problems and their mitigation.

The recent release of EMI containing dozens of software components triggered the real usage of the workflow. Already, some more issues have been identified that will be mitigated for the next releases.

8 REFERENCES

R 1	MS402 - Deploying Software into the EGI production infrastructure: https://documents.egi.eu/public/ShowDocument?docid=53
R 2	MS501 - Establishment of the EGI Software Repository and associated support tools: https://documents.egi.eu/public/ShowDocument?docid=46
R 3	MS502: DMSU Operations Procedures: https://documents.egi.eu/public/ShowDocument?docid=69
R 4	MS508: Software Provisioning Process: https://documents.egi.eu/public/ShowDocument?docid=505
R 5	MS506: EGI software repository architecture and plans: https://documents.egi.eu/public/ShowDocument?docid=503
R 6	MS505 - Service Level Agreement With A Software Provider: https://documents.egi.eu/public/ShowDocument?docid=212
R 7	D5.2 UMD Roadmap: https://documents.egi.eu/public/ShowDocument?docid=272
R 8	EGI WIKI: Staged Rollout workflow https://wiki.egi.eu/wiki/Staged-rollout-procedures
R 9	Early Adopters portal: https://www.egi.eu/earlyAdopters/table
R 10	SW components table: https://www.egi.eu/earlyAdopters/teams
R 11	EGI Request Tracker: https://rt.egi.eu/rt/index.html
R 12	EGI Document Database: https://documents.egi.eu/public/DocumentDatabase
R 13	EGI Repositories: http://repository.egi.eu/
R 14	CERN savannah "jra1mdw" project: https://savannah.cern.ch/projects/jra1mdw/
R 15	EGI Trust Anchor procedures: https://wiki.egi.eu/wiki/EGI_IGTF_Release_Process
R 16	UMD Roadmap – D5.2: https://documents.egi.eu/public/ShowDocument?docid=272
R 17	Findings of Dry Runs Using EMI Components: https://documents.egi.eu/public/ShowDocument?docid=479