



WestminsterResearch

<http://www.wmin.ac.uk/westminsterresearch>

Automatic deployment of interoperable legacy code services.

Gabor Kecskemeti¹, Yonatan Zetuny¹, Tamas Kiss¹, Gergely Sipos², Peter Kacsuk², Gabor Terstyanszky¹ and Stephen Winter¹

¹ School of Informatics, University of Westminster

² 2MTA SZTAKI Laboratory of Parallel and Distributed Systems
H-1518 Budapest, P.O. Box 63, Hungary

This is an electronic version of a paper presented at the CoreGRID Workshop on Grid Systems, Tools and Environments (WP7 Workshop) (in conjunction with GRIDS@Work), 12-14 Oct 2005, Sophia Antipolis, France.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners. Users are permitted to download and/or print one copy for non-commercial private study or research. Further distribution and any use of material from within this archive for profit-making enterprises or for commercial gain is strictly forbidden.

Whilst further distribution of specific materials from within this archive is forbidden, you may freely distribute the URL of WestminsterResearch.
(<http://www.wmin.ac.uk/westminsterresearch>).

In case of abuse or copyright appearing without permission e-mail wattsn@wmin.ac.uk.

Automatic Deployment of Interoperable Legacy Code Services

G. Kecskemeti¹, Y. Zetuny¹, T. Kiss¹, G. Sipos², P. Kacsuk², G. Terstyanszky¹, S. Winter¹

¹*Centre of Parallel Computing, Cavendish School of Computer Science,
University of Westminster, 115 New Cavendish Street, London W1W 6UW,*

²*MTA SZTAKI Laboratory of Parallel and Distributed Systems
H-1518 Budapest, P.O. Box 63, Hungary*

Abstract

The Grid Execution Management for Legacy Code Architecture (GEMLCA) enables exposing legacy applications as Grid services without re-engineering the code, or even requiring access to the source files. The integration of current GT3 and GT4 based GEMLCA implementations with the P-GRADE Grid portal allows the creation, execution and visualisation of complex Grid workflows composed of legacy and non-legacy components. However, the deployment of legacy codes and mapping their execution to Grid resources is currently done manually. This paper outlines how GEMLCA can be extended with automatic service deployment, brokering, and information system support. A conceptual architecture for an Automatic Deployment Service (ADS) and for an x-Service Interoperability Layer (XSILA) are introduced explaining how these mechanisms support desired features in future releases of GEMLCA.

1. Legacy Code Services for the Grid

The Grid requires special Grid enabled applications capable of utilising the underlying middleware and infrastructure. Most Grid projects so far have either developed new applications from scratch, or significantly re-engineered existing ones in order to be run on their platforms. This practice is appropriate in this context, where the applications are mainly aimed at proving the concept of the underlying architecture. However, as the Grid becomes stable and commonplace in both scientific and industrial settings, a demand will be created for porting a vast legacy of applications onto the new platform. Companies and institutions can ill afford to throw such applications away for the sake of a new technology, and there is a clear business imperative for them to be migrated onto the Grid with the least possible effort and cost. Grid computing is now progressing to a point where reliable Grid middleware and higher level tools will be offered to support the creation of production level Grids. A high-level Grid toolkit should definitely include components for turning legacy applications into Grid services.

The Grid Execution Management for Legacy Code Architecture (GEMLCA) [1] enables legacy code programs written in any source language (Fortran, C, Java, etc.) to be easily deployed as a Grid Service without significant user effort. GEMLCA does not require any modification of, or even access to, the original source code. A user-level understanding, describing the necessary input and output parameters and environmental values such as the number of processors or the job manager required, is all that is needed to port the legacy application binary onto the Grid.

In order to offer a user friendly application environment, and support the creation of complex Grid applications from building blocks, GEMLCA is integrated with the workflow oriented P-GRADE Grid portal [2]. Using the integrated GEMLCA – P-GRADE portal solution users can create complex Grid workflows from legacy and non-legacy components, map them to the available Grid resources, execute the workflows, and visualise and monitor their execution.

A drawback of the current solution is the static mapping of legacy components onto resources. Before creating the workflow the legacy application has to be deployed on the target site, and during workflow creation, but prior to its submission, the user has to specify the resource where the component will be executed. It would be desirable to allocate resources dynamically at run-time and to automatically deploy a legacy component on a different site in order to achieve better performance.

Figure 1 illustrates how GEMMLCA can be extended with these functionalities. Instead of mapping the execution of workflow components statically to the different Grid sites, the abstract workflow graph created by the user is passed to a resource broker together with quality of service (QoS) requirements. The broker contacts an information service and tries to map different components of the workflow to different resources and pre-deployed services. If user QoS requirements cannot be fulfilled with the currently deployed services, or if the required service is not deployed on any of the resources, the broker contacts the automatic deployment service in order to deploy the code on a different site. As the sites can belong to different Grids with different middleware, policy and security standards, the deployer service should resolve these interoperability problems.

Unfortunately no currently existing information system, resource broker or deployment service can be directly used and integrated with GEMMLCA to solve these problems. Significant research, extension and improvement of existing solutions are necessary. In this paper we concentrate on a subset of this complex architecture and propose a solution for the Automatic Deployment Service (ADS) and for an x-Service Interoperability Layer (XSILA).

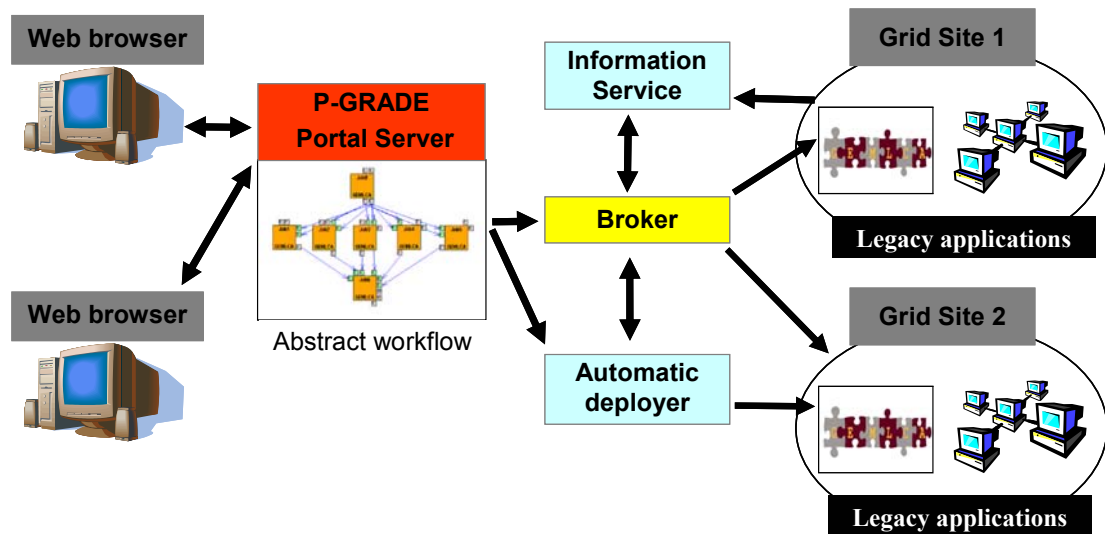


Figure 1 Brokering, information system and automatic deployment support in GEMMLCA

2. Automatic Deployment Service in GEMMLCA

In the current GEMMLCA architecture legacy code services are deployed and mapped manually to Grid resources at workflow construction time. As a pre-requisite to extending GEMMLCA with QoS based brokering and load-balancing capabilities, services have to be automatically deployed or migrated from one site to another. This section describes the challenges faced when deploying services, and proposes a general architecture for an Automatic Deployment Service.

2.1 Deployment Scenarios

There are several research efforts identifying and implementing solutions for scenarios when automatic deployment of services is important [3]. Each scenario can be derived from the following two basic cases:

1. *Deploying new Grid services.* This scenario means the deployment of a new Grid service onto a target site by the service developer. Dependencies have to be detected and resolved by the automatic service deployment tool, and the service container has to be prepared accordingly in order to prevent misbehaviour.
2. *Migrating existing Grid services.* This scenario occurs when migrating an already deployed Grid service to a different site where a dependency description is available. However, even within the same Grid, this description could be in a different format than is required, depending on the selected service container. An automated deployment tool should provide a transformation between different dependency descriptions. Where the description is not appropriate, dependencies have to be investigated like in the previous scenario.

Based on these two basic scenarios the following examples illustrate where automatic service deployment is important in a Grid environment:

- *Automatic selection services.* An already deployed service can't process any more request as its hosting container is overloaded. The service has to be migrated to a site with lower load, and some of its requests have to be redirected to the newly deployed service.
- *Grid systems integration.* Joining different Grids can be more efficient when some services are installed on both of them. Migration of a service in this situation may result in lower communication overhead. In this case a translation is needed between the different site description languages, and deployment specific information has to be provided. Following this, the system has to install the proper environment on the Grid receiving the service in order to carry out the migration.
- *Refining existing services.* Some services (usually data retrieval solutions) provide very generic information to their users, irrelevant to their real, usually restricted, needs. In this case users have to filter this information in order to retrieve what is relevant for them. To avoid high network traffic this filtering can be implemented and deployed as a new service on the site where the general service resides.

2.2 Deployment Service Architecture

In order to support the previously described scenarios a layered deployment service architecture has been identified. Figure 2 shows this architecture, and illustrates how it is utilised when migrating an already deployed service to a target site. The migration process and the tasks of the different layers of the architecture are the following:

1. The Grid sites register themselves in an information system. The registration contains basic site descriptions.
2. In order to be migrated from site A to an appropriate target site, the service contacts the Automatic Deployment Service.
3. The deployment service queries the information system in order to access site descriptions, and also generates the description of the service to be migrated. The classifier module [4] tests the description of the service against the site descriptions, and generates a set of sites that are the most capable of hosting the service. All the descriptions, with the help of ontologies, are transformed into a meta-description suitable for classification [5]. Following this, the dependency checker investigates the capabilities of the selected sites. The capabilities should

be identified with a black box method as the source code is not available in GEMMLCA. In a black-box method, dependencies are detected using an observer execution environment. The service uses generic test data that affects all of its features in order to gather runtime dependencies, such as the files accessed, network connections used, or environment variables needed to be set up. The generated descriptions are stored in the information system for further use.

4. Based on the information received from the dependency checker the comparator prepares some metrics (cost and time requirements of the deployment based on the descriptions), and selects the site with the lowest deployment cost (Site B in our example) [6].
5. In order to make *Site B* compatible with *Site A* from the service's point of view, the dependency installer prepares several installation scripts and environment configuration files/setup scripts. These scripts have to take care of all third party software necessary for the service. The established network connections have to be simulated with a proxy. This proxy has to be prepared on both sites.
6. The deployer prepares a sandbox [7] on *Site B* in order to separate the execution of the service from others. The sandboxing technique used can be various; e.g. a basic *chrooted* environment, some Java security model based solution, or a virtualisation technique (Xen, VirtualPC, VMware). The deployer interfaces with the actual sandboxing technique to create a new sandbox, and then the installation scripts, created in the previous step, are executed in it.
7. The deployer notifies *Site A*, and negotiates the transfer of the service between the sites. The negotiator can detect the available and accessible transfer services on each site. It also has the capability to act as an intermediate layer between the source and the destination, if it is necessary. The service has to be registered with the new host environment in an execution environment specific way without restarting it (the state information of services should not be modified) [8]).

After the transfer is completed between the two sites the service becomes available on the new site.

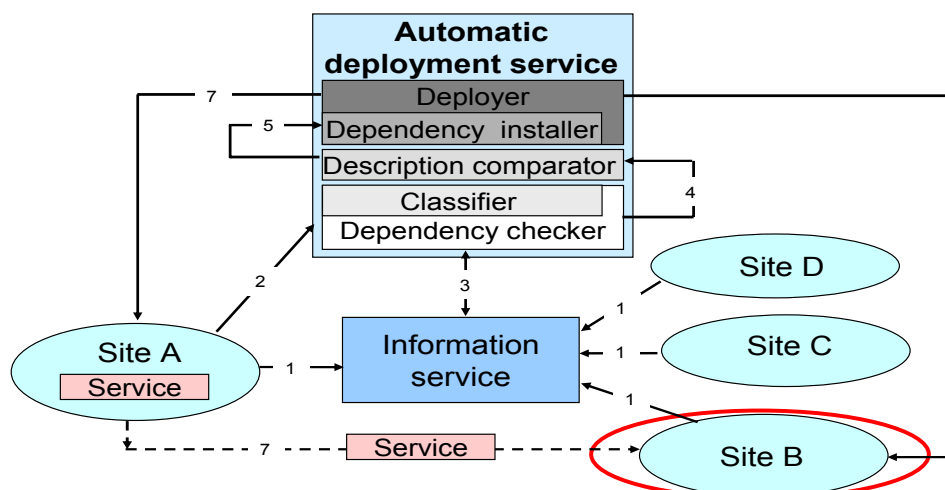


Figure 2 Automatic Deployment Service Architecture

3. x-Service Interoperability

The ADS presented in the previous section offers solution for automatic deployment of services within the same administrative domain. However, interoperability issues have to be taken into consideration when bridging different Grid domains. The aim of

our Grid services interoperability research is to build on existing policy and security solutions and standards that are managed independently by different Grid sites, and to develop an architecture that is capable of bridging isolated Grids in a flexible, scalable and dynamic manner. As a result of this work, GEMLCA is significantly extended to enable the deployment, creation, invocation and management of Grid services between multi-domain Grid environments, thus enabling a dynamic integration of different Grid sites.

A general interoperability architecture, the x-Service Interoperability Layer (XSILA), has been specified in order to handle interoperability issues between Grid clients and Grid services when they are in different domains ("x" refers to any kind of Grid or Web service in this context). Extending the ADS with XSILA enables the automatic deployment of a Grid service into different domains. XSILA serves as a bridge between the different Grids, and makes the deployment to a different domain transparent for the ADS by redirecting the communication between the ADS and the services through XSILA, as illustrated on Figure 3. The architecture is composed of five layers:

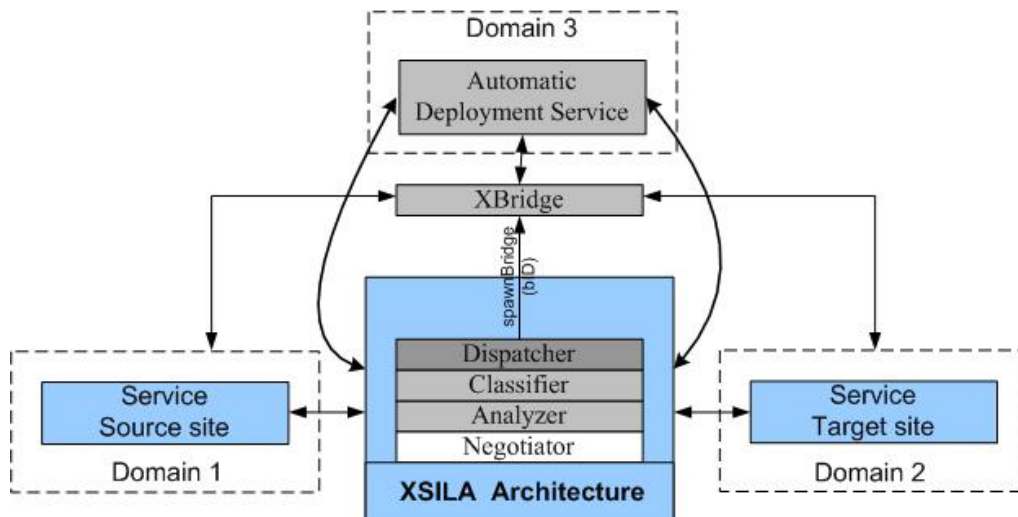


Figure 3 ADS and the x-Service Interoperability

1. *Negotiator Layer* - collects interoperability properties, such as access mechanisms, policies, and security mechanisms of the involved domains.
2. *Analyzer Layer* - analyses the properties collected by the negotiator layer, defines the differences between domains, and prepares a list of interoperability requirements based on these differences.
3. *Classifier Layer* - classifies the interoperability requirements into interoperability classes. It utilizes a mapping engine to create correlation between the demands of each domain.
4. *Dispatcher Layer* - uses the mapping produced by the classifier layer to spawn a Bridge Service that contains the generated mappings. Each dispatched bridge includes a unique identifier which is then can be used by a client to access the service.
5. *Bridge Layer* - encompasses one or more Bridge Services that are spawned by the Dispatcher Layer. Each Bridge Service is intended to resolve a particular interoperability problem. The Bridge service is discarded once a communication is no longer required.

4. Conclusion and Further Work

Deploying legacy applications on the Grid without reengineering the code is crucial for the wider scientific and industrial take-up of Grid technology. GEMMLCA provides a general solution in order to convert legacy applications as black-boxes into OGSA compatible Grid services, without any significant user effort.

Current GEMMLCA implementations fulfil this objective, and the integrated GEMMLCA - P-GRADE Portal solution offers a user friendly Web interface and workflow support on top of this. However, GEMMLCA should be further developed and extended with additional features, like information system support, brokering, load balancing or automatic deployment and migration of services, in order to offer a more comprehensive solution for Grid users.

This paper presented an Automatic Deployment Service Architecture that enables the automatic deployment and migration of GEMMLCA Grid services to different sites within the same Grid domain. The combination of this architecture with the x-Service Interoperability Layer extends deployment and migration capabilities to different domains. Adding these features to GEMMLCA enables service developers to deploy their services automatically on the target site, or to migrate the service to a different site, spanning multiple Grid domains when required, if execution is more efficient there.

The implementation of these architectures and their integration with GEMMLCA is currently work in progress. Also, the investigation has already started how it could be integrated and extended with existing information system and brokering solutions in order to realise the full GEMMLCA-based Grid presented in Figure 1 of this paper.

References

- [1] T. Delaittre, T. Kiss, A. Goyeneche, G. Terstyanszky, S. Winter, P. Kacsuk: GEMMLCA: Running Legacy Code Applications as Grid Services, To appear in "Journal of Grid Computing" Vol. 3. No. 1.
- [2] Cs. Nemeth, G. Dozsa, R. Lovas, P. Kacsuk, "The P-GRADE Grid portal", In: Computational Science and Its Applications - ICCSA 2004: International Conference, Assisi, Italy, 2004, LNCS 3044, pp. 10-19.
- [3] J. B. Weissman, S Kim, D. England. Supporting the Dynamic Grid Service Lifecycle, Technical Report, University of Minnesota, 2004,
- [4] Mike James : Classification Algorithms, Wiley, 1985, ISBN: 0-471-84799-2
- [5] M. Cannataro, C. Comito: A Data Mining Ontology for Grid Programming, Conf. Proc of the 1st Workshop on Semantics in Peer-to-Peer and Grid Computing at the Twelfth International World Wide Web Conference, 20 May 2003, Budapest, Hungary
- [6] P. Watson, C. Fowler. An Architecture for the Dynamic Deployment of Web Services on a Grid or the Internet, Technical Report, University of Newcastle, February, 2005.
- [7] M. Smith, T. Friese, B. Freisleben. Towards a service-oriented ad hoc grid, Conf. Proc of the ISPDC/HeteroPar Conference, 2004
- [8] A. Ting, W. Caixia, X. Yong. Dynamic Grid Service Deployment, Technical Report, March, 2004