

4.1 WP 1 Installation and Evaluation of the Globus Toolkit

Version 1.7.4 – 4 July 2000

4.1.1 General Description of the Package

In order to create and deploy usable Grids, a framework providing basic services and functionalities is necessary.

The toolkit provided by the Globus project has been identified as a possible candidate for this “software infrastructure” for several reasons:

- The Globus “layered bag of services” model is suitable to support a wide variety of applications and programming models: developers of tools or applications can select only the specific tools they need
- Services are distinct and have well-defined interfaces, therefore they can be integrated into applications or tools in an incremental fashion
- The Globus toolkit has already been used with success in several projects of different disciplines. Other Grid projects (i.e. PPDG, NASA IPG) have chosen Globus as basic infrastructure
- The current research activities of the Globus team (i.e. data grids, advance reservation, end-to-end QoS, ...) are considered of much interest for this project
- Contacts with the Globus core team are already in place, and good and productive relations may be established.

The goal of this work package is to evaluate Globus as a distributed computing framework: Globus mechanisms will be evaluated for robustness, ease of use, effectiveness and completeness, in order to find which services can be useful for the middleware work packages of the INFN Grid project and of the DataGrid project (the European Grid project).

Modifications and extensions needed to correct any Globus shortcomings will be identified.

Work in this work package will be coordinated with the similar activities that are going to be performed also by other members of the European Grid project (i.e. Cern), in order to share the acquired experience, to avoid the duplication of the same activities, etc...

4.1.2 Work plan

Activities in this work package will investigate the functionalities of the Globus toolkit, to evaluate which services can be used, what is necessary to integrate/modify, what is missing and therefore must be found somewhere else, or must be developed.

In particular Globus services for security, information, resource management, data migration, fault monitoring, execution environment management will be investigated.

Some activities have already been started, and there is already an experience with the Globus toolkit by some members contributing to this work package. Therefore it is foreseen that the described activities can be performed in a short period (6 months).

The results of this evaluation will be considered to plan and schedule the future work. For example, if the activities of this work package will prove that Globus can be used as a framework providing basic GRID functionalities, it will be necessary to widely deploy Globus in all INFN sites in order to build the infrastructure for the INFN GRID, it could be necessary to modify and extend the software to correct the identified shortcomings, etc...

Task 1.1 Globus deployment

In order to widely deploy Globus in a successful way, it is very important to reduce the complexity for installation and maintenance, limiting the manpower required for these operations. Local administrators, responsible to install Globus on local resources, will be provided with specific tools, documentation, operational procedures, etc...

A group of "central" Globus administrators will implement and maintain a central software repository (for example using AFS): they will install Globus for all the considered platforms, test and install new software releases, install the required patches, develop tools and documentation to enable local administrators to quasi-automatically install the software on the local resources, etc...

Local administrators will only be responsible to deploy Globus on the local machines. These activities will require limited manpower, thanks to the tools and the documentation provided by the "central" team.

Task 1.2 Security

The Grid Security Infrastructure (GSI) service, focusing on authentication, will be evaluated.

The single sign-on mechanisms for all Grid resources provided by GSI will be tested between different administrative domains, each one using different security solutions, mechanisms and policies (such as one-time passwords and kerberos).

Activities for this task will also include the integration of some existing applications with the GSI authentication mechanisms, using the gss-api.

Since the Grid Security Infrastructure is based on public key technology, that relies upon certificates, a "local" Certification Authority responsible to issue these certificates must be implemented. The relations (i.e. CA hierarchy) between this Certification Authority and other Certification Authorities (for example those ones implemented by other partners of the DataGrid project) must be investigated, defined and implemented.

Task 1.3 Information Service

The Grid Information Service (GIS), formerly known as Metacomputing Directory Service (MDS) will be analyzed.

Using LDAP as standard protocol, GIS provides information on system components of the dynamic GRID environment (i.e. architecture type, amount of memory, load of a computer, network bandwidth and latency, ...) in a common interface. This information can then be accessed and used by various components (for example by a resource broker, for resource discovery).

Activities for this task have already been started, in particular with the implementation of a national GIS server.

Tests on performance and scalability (increasing the number of agents that produce useful information and the number of components that access and use that information) will be performed. The related results will be used to decide and establish the most suitable architecture for reliable and redundant access to the GIS in the multi-site INFN environment (How many GIS servers have to be considered ? Where do these servers have to be installed ? How the information can be replicated ? ...).

Tools and interfaces (in particular web user interfaces) addressed to Grid users and administrators must be developed, to support the various operations related to the Globus GIS service: browsing the information servers, adding, deleting, modifying information, etc...

Tests will also be performed on the possibility and suitability to generalize the GIS schema, integrating the default information with other required information (i.e. agents that automatically provide the GIS with specific information on resource characteristics and status, information on replicas of data, ...).

Task 1.4 Resource Management

The GRAM (Globus Resource Allocation Management), that is the Globus service for resource allocation and process management, will be analyzed and evaluated.

In the Globus resource manager architecture, each GRAM is responsible for a set of resources operating under the same site-specific allocation policy, often implemented by an existing resource management system.

Activities in this task will be done in an incremental fashion, starting with simple scenarios.

As first step, the Globus functionalities for job submission on remote resources that don't rely upon resource management systems will be evaluated.

The functionalities of GRAM in conjunction with different (local) resource management systems (in particular LSF and Condor) will be investigated, to evaluate if GRAM could be a possible uniform interface to different resource managers. Demonstrations showing submitting jobs via Globus which run on remote systems using different underlying resource managers, complete with file staging, forwarding of output, etc... without the need for direct human intervention will be performed.

For these tests the INFN Condor WAN pool will be interfaced to Globus as a "single" resource.

It will be necessary to evaluate how the resource requests expressed in terms of Globus RSL (Resource Specification Language) are translated in the resource description languages of the specific resource management systems considered.

It will also necessary to evaluate if and how each local GRAM is able to provide the Grid Information Service with information on characteristics and status of the local resources.

Task 1.5 Data Access and Migration

The GASS (Global Access to Secondary Storage) APIs, that allow programs that use the C standard I/O library to open and subsequently read and write files located on remote computers, will be analyzed, considering the possible file access patterns. The capabilities, performance, ease of use of the GASS services will be evaluated and compared with other solutions, such as distributed file systems. Some existing applications will be adapted to use these APIs, to verify if GASS mechanisms can be a viable and suitable solution to access remote data.

Moreover, tests will be performed to evaluate the functionalities, performance and robustness of the Globus tools for data migration. In particular the GASS service and GSIFTP (an FTP which works with the Globus authentication system and allows high speed file transfers with fully adjustable TCP window sizes) will be evaluated in LAN and WAN environments, and compared with other tools and services, such as FTP.

Task 1.6 Fault Monitoring

The Globus HBM (HeartBeat Monitor) service, providing simple mechanisms for monitoring the health and status of a distributed set of processes, will be analyzed, evaluating its scalability, accuracy, completeness, overhead and flexibility.

This service for fault detection will be evaluated not only for controlling Globus system processes (as it has been originally designed), but it will also be analyzed if it is suitable for monitoring other processes implementing particular services (i.e. http server daemons) and also for application processes associated with “user” computations.

Using the provided API, specific data collectors will be implemented: in the beginning these data collectors should simply detect and notify about failures of processes that have identified themselves to the HBM. Fault recovery mechanisms, such as automatic restart of crashed daemon processes, will be implemented later.

Task 1.7 Execution Environment Management

Tools are needed to set up the environment in which code must actually execute: binaries have to be staged, runtime libraries have to be located and/or instantiated, environment variables and other state needed must be established, etc...

The Globus Executable Management (GEM) service will be tested, in order to evaluate whether it is able to provide some of these functionalities, required when a code migration strategy (the application is moved where the computation will be performed) must be considered.

Table 4.1 – List of the Deliverables for WP1

Deliverable	Description	Delivery Date
D1.1	Tools, documentation and operational procedures for Globus deployment	6 Months
D1.2	Final report on suitability of the Globus toolkit as basic Grid infrastructure	6 Months

Table 4.2– List of the Milestones for WP1

Milestone	Description	Date
M1.1	Basic development grid infrastructure for the INFN GRID	6 Months

4.1.3 Resources

Only the resources required for the described activities, that will last for the year 2000, are listed in the following tables. As described above, the resources necessary for the next years will be better identified after this first evaluation of the Globus toolkit.

Table 4.3 – Personnel Contribution to the WP1 Tasks (man-months)

Site	Name	Total
Bari	Maria D'Amato	1
	Art. 15	2.5
Bologna	Franco Semeria	2
Catania	Roberto Barbera	1
	Ernesto Cangiano	1
	Carlo Rocca	1
CNAF	Andrea Chierici	2
	Luigi Fonti	1
	Francesco Giacomini	1.5
	Alessandro Italiano	1
	Giulia Vita Finzi	2
Lecce	Massimo Cafaro	1
	Lucio Depaolis	1
	Enrico M. V. Fasanelli	2
LNL	Massimo Biasotto	1
Milano	Giuseppe Lo Biondo	2
	Francesco Prelz	1
Padova	Massimo Sgaravatto	2.5
Pisa	Giuseppe Bagliesi	2
	Andrea Controzzi	3
	Flavia Donno	2
	Andrea Sciaba`	2
	Zhen Xie	2
Roma 1	Daniela Anzellotti	1.5
	Marco De Rossi	0.5
	Ettore Majorana	1
	Cristiano Palomba	2.5
	Davide Rossetti	1.5
	Alessandro Spanu	1.5
	Enzo Valente	1.5
Torino	Cosimo Anglano	1
	Antonio Forte	1
	Luciano Gaido	1
Total		50.5

Table 4.4 – Table of Materials for WP1 in K€

Site	Description	Year 2000	Total
Bari	QUANTUM		
Bologna	QUANTUM		
Catania	QUANTUM		
CNAF	QUANTUM		
Lecce	QUANTUM		

LNL	QUANTUM		
Milano	QUANTUM		
Padova	QUANTUM		
Pisa	QUANTUM		
Roma1	QUANTUM		
Torino	QUANTUM		
Total			

Table 4.5 – Table of Consumables for WP1

Site	Description	Year 2000	Total
Bari			
Bologna			
Catania			
CNAF	License for Netscape LDAP Server– 1000 keys	10 MLire	
Lecce			
LNL			
Milano			
Padova			
Pisa			
Roma 1			
Torino			
Total			

Table 4.6 – Table of Travels for WP1 in in K€

Site	Description	Year 2000	Total
Bari	Italian		
	International		
Bologna	Italian		
	International		
Catania	Italian		
	International		
CNAF	Italian		
	International		
Lecce	Italian		
	International		
LNL	Italian		
	International		
Milano	Italian		
	International		
Padova	Italian		
	International		
Pisa	Italian		
	International		
Roma 1	Italian		
	International		
Torino	Italian		
	International		
Total			