



IST-2001-32133

GridLab - A Grid Application Toolkit and Testbed

Production testbed

Author(s):	Luděk Matyska, Miroslav Ruda, Martin Kuba, Petr Holub
Document Filename:	GridLab-5-D5.4-0010-Testbed
Work package:	Testbed
Partner(s):	Masaryk University
Lead Partner:	Masaryk University
Config ID:	GridLab-5-D5.4-0010-1.0
Document classification:	DELIVERABLE

Abstract: This document provides basic information about the final state of the Gridlab testbed.



Contents

1 Overview	2
2 GridLab Testbed Overview	2
2.1 Security and Accounts	2
3 Testbed Management	3
3.1 Testbed monitoring	3
3.2 Portals	3
4 Testbed components	4
4.1 Information services	5
4.2 Distributed components	5
5 Required Software on Nodes	5
6 Conclusions	5
A Current machines in testbed	7
B Current hardware architectures and operating systems available on the testbed	8
C Services currently monitored by testbed monitoring	9

1 Overview

This document describes the GridLab testbed environment in its final state. It is an evolution of the Deliverables 5.2 *Prototype testbed* [2] and 5.3 *Production testbed* [3]. The document summarizes current testbed status together with changes and differences from previous versions.

2 GridLab Testbed Overview

The GridLab testbed purpose has not changed since the prototype—it provides a stable and reliable Grid environment for the *development* of GAT and other Grid Services within the GridLab project. The stability of resources and services and timely provided information about non-functional nodes and/or services was the major goal of the GridLab Testbed Operation Centre.

Proof of concept production runs are also anticipated on this testbed, e.g. as part of demonstrations of the GridLab results. More information about the conditions for use of the GridLab testbed can be found in previous deliverables, in particular Deliverable *Prototype Testbed* [2].

Testbed is built on top of existing grids, most machines in testbed are not dedicated to GridLab testbed. Problems related to building and managing such environment has been described in Deliverable *Grid of the Grids* [5]. Current list of machines and machine architectures is in Appendix.

The GridLab testbed is fully available (without any specific restrictions) for all the GridLab project participants. New users are added by the GridLab Testbed Operation Centre (GTOC), established in Brno, at the Masaryk University. People outside the GridLab project can also use the testbed to some extent after Steering Committee approval. This is the case of, e.g., LSU and KISTI users. Currently, the GridLab testbed has 80 registered users.

2.1 Security and Accounts

As described in previous versions of testbed description ([3, 2]) and in *Grid of the Grids* deliverable ([5]), the only authentication mechanism accepted on GridLab testbed is GSI. The GTOC maintains a list of generally accepted Certification Authorities, whose certificates are accepted by most sites. Currently, GridLab testbed accepts CAs as specified by the *EU Grid PMA consortium*. The list of accepted CAs is maintained on the web. Each site can restrict this list or add new CAs (this has been used, e.g., during demonstrations when adhoc users without certificates issued by some approved CA were temporarily allowed on the GridLab testbed).

An access to the individual computing nodes requires use of local accounts. The Grid-wide user identity is provided by certificates, whose subjects are mapped to local login names via a `grid-mapfile` on each local resource. In GridLab testbed, user certificates map to already established user accounts, or to GridLab specific login names `glabXXX` created by local system administrators. Dynamic maintenance of user accounts by GridLab Authorization Service (GAS) was tested on small subset of testbed machines, but no dynamic account management is used. Central user information is maintained by the GTOC (using *Perun*, a tool developed at Masaryk University [1] primary for the national Grid) in a database, which is used to generate master `grid-mapfile` used by all GridLab sites. Site administrators are notified whenever a change in this file occurs. While the local sites have authority to remove any particular user from the local `grid-mapfile`, currently all sites accept all registered users (provided they have certificates signed by an appropriate CA).

GridLab testbed does not use any testbed-wide accounting system.

3 Testbed Management

While the GTOC is responsible for the general GridLab testbed management, it does not have the authority to manage local sites. They are deployed, configured, and maintained by local administrators, who have no direct formal responsibility to the GTOC team. Therefore, an essential part of the testbed management within the GridLab project is the continuous monitoring and event driven as well as periodic reporting of resource and service state to appropriate administrators.

3.1 Testbed monitoring

The GridLab testbed monitoring provides accurate and up to date information about state of all individual sites, resources, and services. It in fact links different nodes into a coherent reliable testbed. The testbed monitoring relies on a whole monitoring system, developed as part of the GridLab project and described in the deliverable *Production Testbed Monitoring Tool* [6], where complete design and a list of features can be found.

The GridLab testbed monitoring system consists of:

- machine oriented tests (one site at a time tests)
- gridlab web service tests (one service at a time tests)
- matrix tests (tests involving two or more sites or services synchronously)
- unit tests (specific tools, middleware components, applications, ...)
- nightly builds of gridlab software (testing the whole environment at a site)
- migrating worm (specific application-like tests, providing a *view from within*)

All these tests run with a specific periodicity and results are stored in a database and also provided through the GridLab Administrative portal. All the monitoring results are also available through a specific web service (which is also tested as one of the GridLab services), this allows to use monitoring information by other middleware components. In particular, the GRMS (Grid Resource Management System) uses the monitoring web service interface to obtain more accurate information about resources status compared to what is available in information services.

3.2 Portals

The primary access to the testbed is provided by the portals. They are described in the Deliverable 5.9 *Production portals* [4], from the GridLab testbed operation the most important sections of these portals are

- Administration portal, which is in fact a part of the general GridLab portal (<https://portal.gridlab.org>), which provides access to all the testbed monitoring, including the possibility to view a history of state of individual site/service, to trigger particular tests and to follow the GridLab monitoring worm.
- GridSphere based user portal (the second part of the GridLab portal at <https://portal.gridlab.org>), which allows jobs submission and monitoring
- Administrative section (which is still part of the GridLab testbed web pages, (<http://www.gridlab.org/WorkPackages/wp-5/testbed.html>), offering among other the step by step description of adding a new node into the testbed

4 Testbed components

GTOC is responsible for coordination of software installation on all machines. Currently, the following software is installed:

- Basic grid middleware
 - Globus 3.2 (pre-webservices part) developed by Globus Alliance
 - GSI-SSH developed by NCSA
- Application toolkit and applications
 - GAT—Grid Application Toolkit developed by WP-1
 - Cactus with GAT Thorn developed by WP-2
 - Triana developed by WP-3
- Information and monitoring services
 - MDS version 2.4, extended by WP-10
 - Igrid, developed also by WP-10
 - Mercury, developed by WP-11
 - Delphoi, developed by WP-7
 - Testbed webservice, developed by WP-5
- Data Management Services developed by WP-8
 - Data Movement service
 - Data Browsing service
 - Replica Catalog
 - Metadata service
- Grid Resource Management System (GRMS), developed by WP-9

GTOC is maintaining basic central services for the GridLab testbed. They include:

- GIIS (Grid Index Information Service) server
- Igrid index server
- Testbed user database and master `grid-mapfile`
- Testbed status monitoring
- Testbed status service
- GridLab CA
- GridLab Portal
- Central CVS repository
- Backup GRMS

Several central services are managed directly by their authors:

- Grid Resource Management System service is managed by the WP-9. It runs on machine `rage1.man.poznan.pl`. Backup server is maintained by GTOC on machine `loni.ics.muni.cz`.
- Data Management Services are taken care by WP-8 and are running in ZIB. Backup copy of these services is planned to be installed at the GTOC before the end of the project (after last stable version of this service is released).
- Adaptive service (Delphoi) is taken care by WP-7 which is running the central service on `fs0.das2.cs.vu.nl`.
- Prototype of Grid Authorization service developed and managed by WP-6 is running on `rage1.man.poznan.pl`.
- Message Box service—developed and managed by WP-12

4.1 Information services

Currently, MDS (from Globus) is used alongside with the newly developed Grid information system (Igrid) to provide information services. The Igrid is expected to replace completely the MDS before end of the project. Information published to MDS and Igrid is used by GRMS for scheduling of jobs and can be also made available to users via MDS/Igrid portlets on GridLab portal.

4.2 Distributed components

Information and monitoring services has providers installed on all testbed machines. These providers are managed by the local administrators, however their general configuration is provided by the GTOC. These providers are not monitored directly, but the information provided by them is checked and discrepancies trigger some corrective actions.

5 Required Software on Nodes

On each testbed machine, a set of prescribed software modules must be installed. The guide “Rules for software installed on GridLab testbed” lists all these modules, together with basic installation and configuration instructions. The guide can be found at <http://www.gridlab.org/WorkPackages/wp-5/documents/software.pdf>.

For the heterogeneous Grid, it is impossible to require a specific path for each software package. Therefore, only the GridLab configuration file, `/etc/gridlab.conf` is strictly required and must be made available on all GridLab testbed machines. This file contains information necessary to provide uniform environment for all users, i.e., is sets variables that specify locations of all required software. Each submitted job is expected to read this file to get the uniform environment. Current list of variables can be found in Deliverable 5.10 “Grid of Grids”.

6 Conclusions

The final GridLab testbed provides robust environment for all the services and tools developed as part of the GridLab project. While highly heterogeneous, with nodes spanning not only Europe, but other continents as well, at least a subset is always available for experiments and further development. The production-level quality heavily depends on extensive Grid status monitoring, that provides timely information about any service or node failures or other problems. These

states are propagated to local administrators via a notification services, with a copy of the information available at the GTOC site. This way, local administrators are immediately aware of problems and can repair them, even as they are formally independent from the GTOC team. The GridLab testbed does not expect all the services or nodes to be always available. On the contrary, it has been built with the understanding that in any reasonable sized Grid there is always some failed node or service. The high robustness is achieved via a combination of redundancy in critical services and extensive Grid status monitoring, whose results are immediately available to all interested (and responsible) parties. This way, the final GridLab testbed does prove that a production-level quality Grid can be deployed and managed even in very heterogeneous environment without a single administrative responsibility of individual nodes.

A Current machines in testbed

Currently, the testbed does contain the following machines:

- skirit.ics.muni.cz (MU, Brno, Czech Republic)
- ibun.ics.muni.cz (MU, Brno, Czech Republic)
- glamdring.zcu.cz (Pilsen, Czech Republic)
- fs0.das2.cs.vu.nl (VU, Amsterdam, The Netherlands)
- origin.aei.mpg.de (AEI, Golm/Potsdam, Germany)
- peyote.aei.mpg.de (AEI, Golm/Potsdam, Germany)
- sr8000.lrz-muenchen.de (Leibniz-Rechenzentrum (LRZ), Munich, Germany)
- hitcross.lrz-muenchen.de (Leibniz-Rechenzentrum (LRZ), Munich, Germany)
- cluster3.zib.de (ZIB, Berlin, Germany)
- onyx3.zib.de (ZIB, Berlin, Germany)
- grape.man.poznan.pl (PSNC, Poznan, Poland)
- rage1.man.poznan.pl (PSNC, Poznan, Poland)
- eltoro.pcz.pl (IMI, Politechnika Czestochowska, Czestochowa, Poland)
- ia64.icis.pcz.pl (IMI, Politechnika Czestochowska, Czestochowa, Poland)
- pegasos.icis.pcz.pl (IMI, Politechnika Czestochowska, Czestochowa, Poland)
- n0.hpcc.sztaki.hu (SZTAKI, Budapest, Hungary)
- bouscat.cs.cf.ac.uk (Cardiff University, Cardiff, Wales)
- gridentry.uni-paderborn.de (PC, Paderborn, Germany)
- sierra0.unile.it (ISUFI/HPCC, Lecce, Italy)
- packcs-e0.scai.fraunhofer.de (Fraunhofer, Germany)
- pclab120.telecom.ece.ntua.gr (NTUA, Athens, Greece)
- mike4.lsu.edu (LSU, Baton Rouge, Louisiana , USA)
- helix.bevc.lsu.edu (LSU, Baton Rouge, Louisiana , USA)
- venus.gridcenter.or.kr (KISTI, Korea)

In total, 24 machines with very different size and computing capacity are part of the GridLab testbed.

B Current hardware architectures and operating systems available on the testbed

- Intel (PIII, Xeon) based Linux clusters, managed by PBS, SGE, Condor
- Itanium2 based Linux machine
- Opteron based Linux machine
- SGI Origin and Onyx, running SGI Irix
- IBM Power4+ based machine, running IBM AIX
- PowerPC, 1 GHz G4 based, running Linux
- AlphaServer running Digital UNIX
- Hitachi SR8000-F1 running HI-UX/MPP

C Services currently monitored by testbed monitoring

The testbed status monitoring currently checks for the state of the following services

- services which must be provided by each node:
 - iStore—registration in iGrid central server called by iStore.
 - iServe—tests whether iGrid local part iServe is running
 - iGrid jobmans—tests whether any jobmanagers are listed in iStore
 - GIIS—registration into the main Grid Index Information Service
 - GRIS—Grid Resource Information Service from Globus
 - GridFTP—grid-enabled GSI-secured FTP server from Globus
 - Gatekeeper—job submission service from Globus
 - GSISsh—GSI-secured shell for remote command-line access
 - jobmanagers—all types of job management systems provided by each node
 - extensions to GRIS developed by WorkPackage-10
 - webservice for writing into GRIS developed by WorkPackage-10
 - Delphoi—Adaptive software developed by WorkPackage-7
 - Mercury—GMA-compliant monitoring system developed by WorkPackage-11
 - GRMS—Grid Resource Management System developed by WorkPackage-9
 - CAs—all required Certificate Authorities' certificates installed
 - required software—as defined in `/etc/gridlab.conf`
 - grid-mapfile—all testbed users have an account
 - mpicc—MPI (Message Passing Interface) C compiler
 - MPI-C—programs compiled with mpicc can be really submitted as jobs
 - mpif77—MPI Fortran compiler
 - MPI-F77—programs compiled with mpif77 can really be submitted as jobs
 - GRMS-MPI—MPI program submitted by GRMS
- web services with one or two instance per testbed
 - Grid Resource Management System service—developed by WP-9
 - Adaptive service—developed by WP-7
 - Data Movement service—developed by WP-8
 - Data Browsing service—developed by WP-8
 - Replica Catalog—developed by WP-8
 - Metadata service—developed by WP-8
 - Authorization service—developed by WP-7
 - Message Box service—developed by WP-12
 - Testbed Status service—developed by WP-5

References

- [1] Aleš Křenek and Zora Sebestiánová. Perun—Fault-Tolerant Management of Grid Resources, Cracow Grid Workshop'04.
- [2] Gridlab WP5 deliverable D5.2, Prototype testbed.
- [3] Gridlab WP5 deliverable D5.3, Production testbed.
- [4] Gridlab WP5 deliverable D5.9, Production portals.
- [5] Gridlab WP5 deliverable D5.10, Grid of the Grids.
- [6] Gridlab WP5 deliverable D5.7, (Production) Testbed Monitoring Tools.