



DataGrid

WP1 SOFTWARE HOWTO

WP01: Grid Workload Management

Document identifier:	DataGrid-01-TEN-0123-0_0
Date:	13/02/2002
Work package:	WP01: Grid Workload Management
Partner(s):	INFN
Lead Partner:	INFN
Document status	DRAFT

Deliverable identifier:

Abstract:

This note discusses briefly what is necessary to install and deploy, to use and exploit the services provided by WP1, considering various scenarios.

Delivery Slip

	Name	Partner	Date	Signature
From	Massimo Sgaravatto	INFN Padova	13/02/2002	
Verified by				

Document Log

Issue	Date	Comment	Author
0_0	13/02/2002	First draft	Massimo Sgaravatto

Document Change Record

Issue	Item	Reason for Change

Files

Software Products	User files
MS Word 2000	DataGrid-01-TEN-0123-0_0.doc
Adobe Acrobat 5.0	DataGrid-01-TEN-0123-0_0.pdf



CONTENT

1. INTRODUCTION..... 4

1.1. OBJECTIVES OF THIS DOCUMENT 4

1.2. APPLICABLE DOCUMENTS AND REFERENCE DOCUMENTS..... 4

1.3. TERMINOLOGY 4

2. OVERVIEW OF THE EDG-WP1 SOFTWARE 6

3. DEPLOYING WP1 SOFTWARE 8

3.1. SCENARIO 1 8

3.1.1. *Installing and configuring a submitting machine*..... 8

3.1.2. *Installing and configuring RB and JSS*..... 9

3.1.3. *Installing and configuring a LB server*..... 9

3.1.4. *Installing and configuring a CE*..... 9

3.2. SCENARIO 2 10

3.2.1. *Installing and configuring a submitting machine*..... 10

3.2.2. *Installing and configuring RB and JSS*..... 10

3.2.3. *Installing and configuring II* 10

3.2.4. *Installing and configuring a LB server*..... 11

3.2.5. *Installing and configuring a CE*..... 11

3.3. SCENARIO 3 11

3.3.1. *Installing and configuring a submitting machine*..... 11

3.3.2. *Installing and configuring RB and JSS*..... 11

3.3.3. *Installing and configuring II* 12

3.3.4. *Installing and configuring a LB server*..... 12

3.3.5. *Installing and configuring a CE*..... 12

1. INTRODUCTION

1.1. OBJECTIVES OF THIS DOCUMENT

This note discusses briefly what is necessary to install and deploy, to use and exploit the services provided by WP1. Three different scenarios have been considered, starting with a simple use-case, and then considering more complicated requirements.

1.2. APPLICABLE DOCUMENTS AND REFERENCE DOCUMENTS

Applicable documents

Reference documents

[R1] WP1 - WMS Software Administrator and User Guide

<http://edms.cern.ch/document/335068>

[R2] <http://marianne.in2p3.fr/datagrid/testbed1/globus/>

1.3. TERMINOLOGY

Definitions

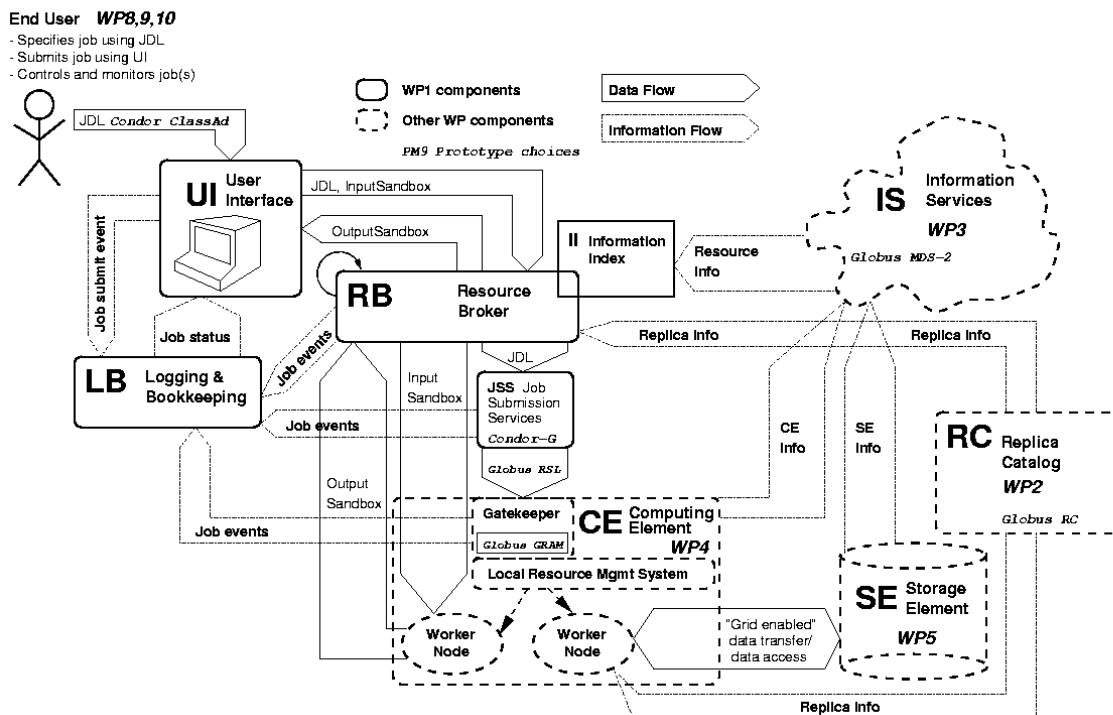


Glossary

CE: Computing Element
EDG: European DataGrid
II: Information Index
JDL: Job Description Language
JSS: Job Submission Service
LB: Logging and Bookkeeping
LRMS: Local Resource Management System
MDS: Metacomputing Directory Service
RB: Resource Broker
RC: Replica Catalog
SE: Storage Element
UI: User Interface
VO: Virtual Organization
WMS: Workload Management System

2. OVERVIEW OF THE EDG-WP1 SOFTWARE

The following figure provides an overview of the WP1 software architecture, and shows the interactions between WP1 software components and the services provided by the other EDG workpackages.



Therefore users “interact” with the WP1 Workload Management System (WMS) via a User Interface (UI).

In particular users can submit jobs, monitor their status, retrieve job outputs, cancel jobs previously submitted to the system.

When a user submits a job, the characteristics and requirements for this job are expressed using a Job Description Language (JDL). Users can choose the computing resource (Computing Element, CE) where to submit the job, or they can leave this decision to the Resource Broker (RB). In the latter case the RB chooses the “best” CE where to submit the job, querying the Information Index (II), the MDS GRIS’es, and possibly a Replica Catalog (RC).

The initial, very simple scheduling strategy is to move the jobs where the data are, taking into account also the other requirements and/or ranks specified in the JDL expression, concerning the Computing Elements, and checking also in which resources that user is allowed to submit jobs.

When the CE where to execute the job has been chosen by the RB, the job is then “transferred” to the Job Submission Service (JSS), which is responsible for the actual submission to the CE, for

performing the other job management operations (issued on requests of the RB), and for monitoring the status of the various jobs.

Note that in both cases (a]: the user decides where to submit the job, and b]: the RB decides where to submit the job) the “logical path” is always UI → RB → JSS → CE, even if the user chose “by-hand” the computing element where to submit the job (in this case the RB doesn’t have to perform any brokering, but it has just to update some internal registries).

The various elements of the WP1 WMS (UI, RB, JSS, and also the Globus jobmanager) push information concerning jobs to a LB (Logging and Bookkeeping) Server. Job status information can then be retrieved from the UI, by querying this LB server.

The WP1 RB is designed to work both at the personal level and the VO level. The decision of the best strategy to pursue will come with the testbed evaluation.

For each RB, there must be a correspondent instance of JSS. Multiple UI’s can “use” the same RB, and a single RB can submit jobs to multiple CE’s.

3. DEPLOYING WP1 SOFTWARE

In this section we briefly describe what and how is necessary to install, to use and exploit the WP1 software, considering different configurations, starting with simple use cases, and then considering more complicated scenarios.

The WP1 software runs and has been tested on platforms running *Globus Toolkit 2.0 Beta Release 21* on top of *Linux RedHat 6.2*.

We discuss here about three possible scenarios:

- Scenario 1: User submits a job, choosing by hand the CE where the job must be executed
- Scenario 2: User submits a job, letting the RB chooses the CE where the job must be executed. No data requirements or requirements on the target storage element specified in the JDL expression
- Scenario 3: User submits a job, letting the RB chooses the CE where the job must be executed. Data requirements and/or requirements on the target storage element specified in the JDL expression

3.1. SCENARIO 1

As first use case, let's consider a user submitting a job to the WP1 Workload Management System, choosing by hand the CE where this job must be executed (therefore using the option `-r` with the `dg-job-submit` command).

For this purpose it is necessary:

- A submitting machine, used to submit jobs and perform the other job management operations
- A machine where to install RB and JSS
- A machine where to install LB server (it could be the same machine used for RB and JSS)
- A CE, where the job execution takes place

3.1.1. Installing and configuring a submitting machine

The installation and the configuration of a submitting machine are described in detail in section 4.4 of [R1].

In short, it is necessary to install:

- 3 WP1 RPM's (*workload-profile.X.Y.Z-K.i386.rpm*, *userinterface-profile.X.Y.Z-K.i386.rpm*, *userinterface-X.Y.Z-K.i386.rpm*)
- Some Globus GT2 RPM's

- Python interpreter
- SWIG

3.1.2. Installing and configuring RB and JSS

Section 4.2 of [R1] discusses about the installation and configuration of the machine which has to host RB and JSS.

In short it is necessary to install:

- 5 WP1 RPM's: *workload-profile.X.Y.Z-K.i386.rpm*, *jobsubmission-X.Y.Z-K.i386.rpm*, *jobsubmission-profile-X.Y.Z-K.i386.rpm* (which install RB and JSS) and *locallogger-X.Y.Z-K.i386.rpm* and *locallogger-profile-X.Y.Z-K.i386.rpm* (for LB local-logger services)
- Globus 2 (the Globus gridftp server package must also be installed and configured on this host)
- PostgreSQL
- Condor-G
- ClassAd library

3.1.3. Installing and configuring a LB server

Section 4.2 of [R1] discusses also the installation and configuration of LB server. Here it is explained that to implement a LB server it is necessary to install:

- 2 WP1 RPM's (*lbserver-X.Y.Z-K.i386.rpm* and *lbserver-profile-X.Y.Z-K.i386.rpm*)
- Globus 2
- MySQL

3.1.4. Installing and configuring a CE

In the DataGrid project, a CE is represented by a queue of a Local Resource Management System (LRMS), configured as Globus resource (that is a "front-end" host has to be configured as Globus gatekeeper machine). Moreover some specific "information providers", which provide the Information Service with information about characteristics and status of the CE, has to be considered (see section 3.2.5 for more details).

However for this simple use case, the information providers are not needed: it is just necessary to consider a LRMS supported by Globus.

Other requirements that have to take into account:

- The home directory for the user "mapped" to the Grid user who submitted the job must be shared between the gatekeeper node and all the worker node (constraint due to Globus GASS)
- In the gatekeeper node, Globus 2 must be installed. The default Globus 2 jobmanager must be replaced with a specific jobmanager (included in the EDG Globus distribution, [R2]) which includes some specific customisations required by the LB service)

- In the gatekeeper node, 2 WP1 RPM's for the local-logger services (*locallogger-X.Y.Z-K.i386.rpm* and *locallogger-profile-X.Y.Z-K.i386.rpm*) must be installed
- In each worker node the Globus 2 RPM's necessary to install the *globus-url-copy* command must be deployed
- There must be at least outbound connectivity from each worker node (necessary to transfer the input and output sandboxes)

3.2. SCENARIO 2

As second scenario, let's consider a user who submits a job, letting the RB chooses the CE where the job must be executed. Let's assume that no data requirements (*InputData* attribute) and no requirements on the target storage element (*OutputSE* attribute) have been specified in the JDL expression.

For this use-case it is necessary:

- A submitting machine, used to submit jobs and perform the other job management operations
- A machine where to install RB and JSS
- A machine where to install II (it could be the same machine used for RB and JSS)
- A machine where to install LB server (it could be the same machine used for RB and JSS)
- A CE, where the job execution takes place

3.2.1. Installing and configuring a submitting machine

The same requirements described in section 3.1.1 for Scenario 1 have to be considered.

3.2.2. Installing and configuring RB and JSS

See section 3.1.2 (same as Scenario 1)

3.2.3. Installing and configuring II

Section 4.3 of [R1] explains what is necessary to implement a II.

In short it is necessary to install and configure:

- Some Globus 2 RPM's
- 3 WP1 RPM's (*workload-profile.X.Y.Z-K.i386.rpm*, *informationindex.X.Y.Z-K.i386.rpm* and *informationindex-profile.X.Y.Z-K.i386.rpm*)

Configuring the II means registering the "address" of the GRIS'es and/or GIIS'es which define the CE's that must be considered by the RB when choosing the best resource where to submit a job.

II should be installed in a machine “close” to the machine hosting the RB. Possibly a same machine should be used for both services.

3.2.4. Installing and configuring a LB server

The same considerations reported in section 3.1.3 are valid here.

3.2.5. Installing and configuring a CE

Besides the considerations reported in section 3.1.4 (for Scenario 1), in the gatekeeper node of the CE it is necessary to install some specific information providers, which provide the information system with information on characteristics and status of local farms. This information is then used in the matchmaking process, when the RB has to find the best CE where to submit a specific job.

The information published by these information providers refer to a specific schema defined in the EDG project (not the default Globus schema) and have been implemented by EDG WP4 so far for LSF, PBS and BQS¹.

Therefore for this particular scenario just these 3 LRMS’s can be considered

3.3. SCENARIO 3

As last scenario, let’s consider a user who submits a job, letting the RB chooses the CE where the job must be executed. Unlike scenario 2, in this case data requirements (*InputData* attribute) and/or requirements on the storage element that must be used to “collect” the output data (*OutputSE* attribute) can be specified in the JDL expression.

For this use-case it is necessary:

- A submitting machine, used to submit jobs and perform the other job management operations
- A machine where to install RB and JSS
- A machine where to install II (it could be the same machine used for RB and JSS)
- A machine where to install RC
- A machine where to install LB server (it could be the same machine used for RB and JSS)
- A CE, where the job execution takes place

3.3.1. Installing and configuring a submitting machine

The same requirements described in section 3.1.1 for Scenario 1 have to be considered.

3.3.2. Installing and configuring RB and JSS

¹ Support for BQS has been provided by Lyon EDG WP6 people



See section 3.1.2 (same as Scenario 1 and Scenario 2)

3.3.3. Installing and configuring II

The same considerations reported in section 3.2.3 for Scenario 2 apply here.

In this case, II should be populated also by information on SE's, besides information on CE's.

3.3.4. Installing and configuring a LB server

The same considerations reported in section 3.1.3 are valid here.

3.3.5. Installing and configuring a CE

See section 3.2.5: the same requirements described for scenario 2 can be applied here.