

EGEE

GRID BATCH JOB ENVIRONMENT VARIABLES

A PROPOSAL FOR STANDARDIZING THE WORKING
ENVIRONMENT FOR A LCG/EGEE JOB

Document identifier:	EGEE-SA1-TEC-GridBatchJobEnvironmentVariable-v2-1.doc
Date:	29/08/2005
Activity:	SA1: European Grid Support, Operation Management
Document status:	DRAFT
Document link:	https://edms.cern.ch/document/630962

Abstract: This document presents a proposal for establishing a minimal set of middleware-independent environmental variables a grid job should find on every site. The definition and intended use of each one of the variables are presented.

Copyright (c) Members of the EGEE Collaboration. 2004.

See <http://public.eu-egee.org/partners/> for details on the copyright holders.

EGEE (“Enabling Grids for E-science in Europe”) is a project funded by the European Union. For more information on the project, its partners and contributors please see <http://www.eu-egee.org> .

You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: "Copyright (C) 2004. Members of the EGEE Collaboration. <http://www.eu-egee.org> ".

The information contained in this document represents the views of EGEE as of the date they are published. EGEE does not guarantee that any information contained herein is error-free, or up to date.

EGEE MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

Document Log

Issue	Date	Comment	Author/Partner
1-0	17/05/2005	First draft	David Bouvet
2-0	01/08/2005	Updated version including initial feedback	Fabio Hernandez (fabio@in2p3.fr)
2-1	26/08/2005	More feedback included.	Fabio Hernandez (fabio@in2p3.fr)

Document Change Record

Issue	Item	Reason for Change

CONTENT

1. INTRODUCTION	4
1.1. PURPOSE	4
1.2. APPLICATION AREA	4
1.3. REFERENCES	4
1.4. DOCUMENT EVOLUTION PROCEDURE	4
2. RATIONALE	5
2.1. ENVIRONMENT VARIABLE FOR NON-GRID BATCH JOBS	5
2.2. ENVIRONMENT VARIABLE DEFINITION FOR GRID BATCH JOBS	5
2.3. VARIABLE DEFINITION	6
2.3.1. HOME	6
2.3.2. TMPDIR	6
2.3.3. GRID_WORKDIR	6
2.3.4. GRID_SITENAME.....	7
2.3.5. GRID_HOSTNAME.....	7
2.3.6. GRID_CEID.....	7
2.3.7. GRID_LOCAL_JOBID.....	7
2.3.8. GRID_GLOBAL_JOBID.....	7
2.3.9. GRID_USERID	8
2.3.10. GRID_NAME	8

1. INTRODUCTION

1.1. PURPOSE

While some effort has been put on standardizing the working environment for jobs at the site level, no work has been reported on the working environment for grid jobs. This document presents a proposal for establishing a set of environment variables a grid job should find on every site, the meaning of each variable and the context where it can be used.

An initial version of this work was presented at the [HEPIX Spring 2005](#) meeting and during the [LCG/EGEE Operations Workshop](#) in May 2005 in Bologna.

1.2. APPLICATION AREA

This document may be of interest to end-users of a grid infrastructure, to grid middleware developers, to site administrators and to operators of grids.

1.3. REFERENCES

[This subsection provides a complete list of all documents referenced elsewhere in the document.]

[R1] http://www.opengroup.org/onlinepubs/009695399/basedefs/xbd_chap08.html	« IEEE Std 1003.1, 2004 POSIX Part 1: Base definitions, Amendment 8 »
[R2] http://www.opengroup.org/onlinepubs/009695399/utilities/xcu_chap03.html	« IEEE Std 1003.1, 2004 POSIX Part 2: Shell and Utilities, Amendment 1 »

1.4. DOCUMENT EVOLUTION PROCEDURE

This document will evolve as a better understanding of the required set of variables and their meaning is acquired.

2. RATIONALE

Porting grid applications to different grids or even to different sites within the same grid may be a difficult task. Given the high heterogeneity inherent to the grid, grid jobs do not have a uniform way to obtain information on the environment it must run in.

We propose in this document a minimal set of middleware-independent environment variables that grid jobs can use to discover the environment it has been scheduled to run in. The intended benefits are:

- To end-users, it will enable thme to write grid applications which better respond to the environment the application runs in,
- To site-administrators, it will enable them to express information on how grid applications should use the resources of their site.

2.1. ENVIRONMENT VARIABLE FOR NON-GRID BATCH JOBS

The IEEE Std 1003.1, 2004 Edition [R1, R2] defines a set of environmental variables that every POSIX compliant batch system must set for a job. This set includes some non-batch-related variables like *HOME*, *PATH*, *PWD*, *SHELL*, *TMPDIR*, *USER* and some batch-specific variables like *PBS_JOBID*, *PBS_JOBNAME*, *PBS_QUEUE*, etc.

In this document we propose to use the existing standard for non-grid-specific variables when appropriate and to introduce a set of grid-specific environment variables for conveying information related to the grid environment.

2.2. ENVIRONMENT VARIABLE DEFINITION FOR GRID BATCH JOBS

This section presents an overview the proposed variables and their meaning. The naming convention was chosen to be independent of the grid middleware for making easier the portability of grid applications between grids using different middleware.

Variable type	Definition	Name
POSIX	Home directory of job user on WN	HOME
	Directory to create temporary files	TMPDIR
Grid-related	Job working directory on WN	GRID_WORKDIR
	Site name on which the job runs	GRID_SITENAME
	Full hostname of the worker node which the job runs on	GRID_HOSTNAME
	Computing Element and queue names on which the job runs (same as <i>GlueCEUniqueID</i> in Information Provider)	GRID_CEID
	Job identifier at the site batch system level	GRID_LOCAL_JOBID
	Job identifier at the grid level	GRID_GLOBAL_JOBID
	Job submitter identifier	GRID_USERID

Table 1: Environment variables for grid batch jobs.

2.3. VARIABLE DEFINITION

This sections presents the meaning and intended use of the proposed environmental variables. As stated above, when existing, well defined POSIX variables should be used for the usage intended in the standard. For instance, if defined, *TMPDIR* should be preferred to middleware-specific environmental variables for creating temporary files.

The variables are classified in two categories: REQUIRED and OPTIONAL. The REQUIRED category means that the variable must be specified by the site and the grid application must use it for its intended purpose. The variables of the OPTIONAL category may or may not be defined: if they are defined, they must be used by the grid application for its intended purpose.

2.3.1. HOME (POSIX)

Category: REQUIRED

Description: This variable represents the full pathname of the user's home directory on the worker node the job runs on.

Example: /home/cms005.

Comments: depending on the site configuration, this directory may be local to the worker node or on a shared file system. It may also be persistent or volatile, in which case, its content is deleted after the end of the job.

In addition, as a grid application developer, one should not assume that the grid job will start execution on this directory. See variable GRID_WORKDIR below.

Open Issues: how to express the shared/non-shared and persistent/volatile information to the job? Should we use other environment variables?

2.3.2. TMPDIR (POSIX)

Category: REQUIRED

Description: This variable is initialized to the full path of a directory made available for applications that need a disk space to create temporary files. Grid applications must use the value of this variable for creating any temporary file they may need.

Example: /tmp/scratch-for-job-1234

Comments: The disk area pointed to by this variable should be considered volatile, i.e. its contents will be deleted after the job finishes its execution. In addition, this directory may be local to the worker node or over a networked file system. Typically, the permissions associated to this disk space allow only the job to read and write files on it.

Open Issues: there are some grid applications that make an intensive use of this temporary disk space. As a consequence, they cannot efficiently run on a worker node where this directory is on a networked file system. How to inform the job about the local/remote nature of this space? Should we use another environmental variable? In addition, how to inform the job on the disk space available for it under this directory?

2.3.3. GRID_WORKDIR

Category: OPTIONAL

Description: This variable contains the full pathname of the working directory the job is started in.

Example: /scratch/atlas0011293.ccw10092

Comments: Files created in this directory may be deleted after the end of the job. If this variable is not defined, the grid application can obtain the directory where it is started by querying the operating system about the current working directory.

2.3.4. GRID_SITENAME

Category: REQUIRED

Description: This variable contains the name of the grid site on which the job runs. In the case of the LCG/EGEE grid, it should be defined with the value of the attribute *siteName* in the site information system.

Example: IN2P3-CC

Comments: This variable is useful for a grid job to know in what site it has been scheduled for execution. From this information, the job owner can find either manually or in an automated way the contact information for the site (user support contact, grid security contact, grid operations contact, etc.). For instance, for the LCG/EGEE grid, this information is stored in the GOC database and can be retrieved from the site identifier by querying the database.

2.3.5. GRID_HOSTNAME

Category: OPTIONAL

Description: This variable contains the fully qualified name of the host the application has been scheduled for execution.

Example: ccwl0006.in2p3.fr

Comments: This variable is provided for the convenience of the application developer and, if defined, must contain the value returned by the UNIX command `/bin/hostname -f`.

2.3.6. GRID_CEID

Category: REQUIRED

Description: This variable contains the identifier of the computing element on which the job has been scheduled for execution. Its value is dependent on the grid middleware. For instance, in the case of the LCG/EGEE grid, it must contain the value of the attribute *GlueCEUniqueID* in the site information system. It has the form `<CE-full-hostname>:2119/jobmanager-<lrms>-<queue-name>`

Example: cclcgceli02.in2p3.fr:2119/jobmanager-bqs-cms_long

Comments: None.

2.3.7. GRID_LOCAL_JOBID

Category: REQUIRED

Description: This variable contains the job identifier for the batch system local to the execution site.

Example: lcg0509104420-07243

Comments: This information may be useful to the job owner for reporting problems with grid jobs to the execution site's operations team.

2.3.8. GRID_GLOBAL_JOBID

Category: REQUIRED

Description: This variable is initialized to the identifier of the job at the grid level. It may take several forms depending on the grid middleware. The example given below is a grid job identifier as assigned by the LCG/EGEE Resource Broker.

Example: `https://lxn1188.cern.ch:9000/HPMN2WVHurMlji-Fnqba0A`

Comments: None.

2.3.9. GRID_USERID

Category: REQUIRED

Description: This variable contains the distinguished name of user's certificate who submitted the job.

Example: `/O=GRID-FR/C=FR/O=CNRS/OU=CC-LYON/CN=David Bouvet/Email=dbouvet@in2p3.fr`

Comments: Grid applications should not make assumptions on the form and contents of the user identifier. For instance, the e-mail address is not always included in it.

2.3.10. GRID_NAME

Category: REQUIRED

Description: This variable contains the name of the grid flavour the job has been submitted to.

Example: LCG or EGEE (?)

Comments: This information may be useful for virtual organizations using several grids for running their applications. Please note that this is not the name of the middleware used to build the grid infrastructure, but an identifier of the grid itself.