



EGI-InSPIRE

UMD QUALITY CRITERIA COMPUTE CAPABILITIES V3

Document identifier:	EGI-COMPUTE-QC-V3.doc
Date:	10/04/2012
Document Link:	https://documents.egi.eu/document/718

Abstract

This document describes the Quality Criteria that all software of the UMD distribution must meet.



Copyright notice

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

Document Log

Issue	Date	Comment	Author/Partner
v0.1	02/11/2010	First draft	Enol Fernández
v1.0	03/11/2010	Changed Management, Traceability and Monitoring section	Enol Fernández
v1.1	03/11/2010	Added Probe description in GEN_MON_1	Enol Fernández
v1.2	11/11/2010	Some formatting update	Enol Fernández
v1.3	31/01/2011	Better test specification	Enol Fernández
1.4	09/02/2011	Review of criteria	Enol Fernández
2 DRAFT 1	24/06/2011	Preparation of new release	Enol Fernández
2	02/08/2011	Reorganisation, added new criteria.	Enol Fernández
3 DRAFT 1	13/10/2011	First draft of release 3	Enol Fernández
3 DRAFT 2	24/01/2012	Second draft of release 3	Enol Fernández



TABLE OF CONTENTS

1	Job Execution	5
1.1	Job Execution Interface	5
	JOBEXEC_IFACE_1	5
1.2	Job Submission tests	6
	JOBEXEC_JOB_1	6
	JOBEXEC_JOB_2	7
	JOBEXEC_JOB_3	8
1.3	Execution Manager Support	9
	JOBEXEC_EXECMNGR_1	9
	JOBEXEC_EXECMNGR_2	10
	JOBEXEC_EXECMNGR_3	11
1.4	Availability/Scalability	12
	JOBEXEC_AVAIL_1	12
	JOBEXEC_AVAIL_2	13
	JOBEXEC_AVAIL_4	14
2	Parallel Job	15
2.1	Submission of parallel jobs	15
	PARALLEL_JOB_1	15
	PARALLEL_JOB_2	16
	PARALLEL_JOB_3	17
2.2	MPI support	18
	PARALLEL_MPI_1	18
	PARALLEL_MPI_2	19
2.3	OpenMP support	20
	PARALLEL_OMP_1	20
	PARALLEL_OMP_2	21
3	Interactive Job Management	22
	INTERACTIVE_JOB_1	22
	INTERACTIVE_JOB_2	23
	INTERACTIVE_JOB_3	24
	INTERACTIVE_JOB_4	25
4	Job Scheduling	26
4.1	Job Scheduling Interface	26
	JOBSCH_IFACE_1	26
4.2	Job Execution Capability Support	27
	JOBSCH_EXEC_1	27
	JOBSCH_EXEC_2	29
4.3	End-to-end job submission tests	30
	JOBSCH_JOB_1	30
	JOBSCH_JOB_2	31
	JOBSCH_JOB_3	32
	JOBSCH_JOB_4	33
	JOBSCH_JOB_5	34
	JOBSCH_JOB_6	35
	JOBSCH_JOB_7	36
	JOBSCH_JOB_8	37
4.4	gLite WMS	38
	JOBSCH_WMS_1	38
	JOBSCH_WMS_2	39



JOBSCH_WMS_3.....	40
4.4.1 Security Advisories.....	41
JOBSCH_WMS_SEC_1.....	41
4.4.2 Bugs.....	42
JOBSCH_WMS_BUG_1.....	42
4.5 Service availability, monitoring and error handling.....	43
JOBSCH_SERVICE_1.....	43
JOBSCH_SERVICE_2.....	44
JOBSCH_SERVICE_3.....	45
JOBSCH_SERVICE_4.....	46
JOBSCH_SERVICE_5.....	47
5 References.....	48

1 JOB EXECUTION

1.1 Job Execution Interface

Currently, there are different interfaces considered for the Job Execution Capability, although not interoperable several of them co-exist in the EGI Infrastructure. The implementations must support, at least, one of the interfaces listed.

Job Execution Interface	
ID	JOBEXEC_IFACE_1
Description	Job Execution Appliances must support (at least one of) the interfaces currently in production in the EGI Infrastructure or identified by the UMD Roadmap
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Implementation of one of the Job Execution Interface as defined in the UMD Roadmap. Ideally, a complete test suite of the Job Execution interfaces supported by the appliance. The test suite must include tests for all the documented functions, and for all functions, check both correct and invalid input and with valid and invalid credentials.
Test Description	<p>Pre-condition Valid user credentials.</p> <p>Test Test all interface functionality, with correct/incorrect input and with valid and invalid credentials.</p> <p>Expected Outcome Log of all the operations performed. All the documented functions work as documented. Errors/exceptions should be generated as documented.</p>
Pass/Fail Criteria	<p>The Job Execution Appliance that claims to support an interface must pass complete tests for that interface (provided by the TP or by the verification team). If the API is not completely supported, this must be documented. The test suite must be executed without errors.</p> <p>At least one of the following interfaces must be supported:</p> <ul style="list-style-type: none"> • ARC-CE gridFTP [R 10] • CREAM [R 11] • EMI-ES [R 12] • Globus GRAM5 [R 13] • OGSA BES [R 15] • UNICORE UAS [R 16]
Related Information	UMD Roadmap [R 1]
Revision Log	V2: unification of several criteria regarding interfaces into this one. V3: removed DRMAA as possible interface.

1.2 Job Submission tests

The following tests propose example job descriptions using the gLite JDL format for the specification of jobs. These examples are just used for illustrative purposes. Each appliance should execute the tests using their native format.

Simple Job	
ID	JOBEXEC_JOB_1
Description	Execute a simple job in the appliance.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Support for the submission of a job with no input or output files.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission of simple job: Executable = /bin/sleep; Arguments = "120";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: merged JOBEXEC_*_JOB_1 into this criterion.

Simple Job with input/output files	
ID	JOBEXEC_JOB_2
Description	Execute a simple job in the appliance that uses both input and output files.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Support for the submission of a job with input or output files.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system) Non-empty files “myfile”</p> <p>Test Job submission for job with input and output files: Executable = "/bin/ls"; Arguments = "-l"; StdOutput = "std.out"; StdError = "std.err"; InputSandbox = {"myfile"}; OutputSandbox = {"std.out", "std.err"};</p> <p>Expected Outcome Job finishes correctly; output contains the listing of the directory including the input file with correct size. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: merged JOBEXEC_*_JOB_2 into this criterion.

Cancel Job	
ID	JOBEXEC_JOB_3
Description	Cancel a previously submitted job.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Support for the cancellation of a job. Job cancelling must be possible for all different states that the job may be, e.g. cancel the job when it's running or cancel the job when it's already done.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job Submission and then cancellation. Possible description for job: Executable = "/bin/sleep"; Arguments = "20m";</p> <p>Expected Outcome Job is submitted and then cancelled correctly. Unique Identifier for the submitted jobs, status log of the job. The job must be removed from the execution manager.</p>
Pass/Fail Criteria	Pass if the appliance is able to cancel jobs for any previous state of the job. If the job is in the execution manager system, it should be completely removed, especially if it's running.
Related Information	
Revision Log	V2: merged JOBEXEC_*_JOB_3 into this criterion. Added clarification

1.3 Execution Manager Support

These QC refer to the interaction of the Job Execution Capability with the underlying execution manager (usually a LRMS) for the work items submitted.

Not Invasive Deployment	
ID	JOBEXEC_EXECMNGR_1
Description	Job Execution Appliances should not introduce any modifications to the underlying execution manager or to the operations of the resources.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Description of all needed, if any, modifications on the local resources in order to deploy the Job Execution Appliance.
Pass/Fail Criteria	Any modifications must be documented, especially invasive ones. Modifications to consider are: <ul style="list-style-type: none"> • Installation of additional software at the WN is permitted as long as no extra services are run permanently at the WN. • Require the deployment of extra (shared) filesystems • Modification of the local submission mechanism of jobs (e.g. require the modification of prologue/epilogue scripts of the batch system) • Require the creation of extra user accounts or add special privileges to a specific account. • Require inbound or outbound connectivity
Related Information	
Revision Log	V2: added inbound, outbound connectivity. Relax Pass/Fail criteria

Job Management	
ID	JOBEXEC_EXECMNGR_2
Description	Job Execution Appliances must support the creation and management of work items to an execution manager.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	<p>Appliances must be able to:</p> <ul style="list-style-type: none"> • create new jobs • retrieve the status of the jobs submitted by the appliance • cancel jobs • optionally, hold and resume jobs <p>The Appliance may perform these operations for individual jobs or for set of jobs in order to improve its performance (e.g. for retrieving the status instead of querying each of the individual jobs, do a single query for all jobs submitted for the appliance)</p>
Test Description	<p>Pre-condition Configured system</p> <p>Test Create new job(s) in execution manager</p> <p>Expected Outcome New job(s) is created in the execution manager; id of job(s) returned</p>
	<p>Pre-condition Previously submitted job(s)</p> <p>Test Cancel job(s) in execution manager</p> <p>Expected Outcome Job(s) is cancelled successfully.</p>
	<p>Pre-condition Previously submitted job(s)</p> <p>Test Query status of previously submitted job(s)</p> <p>Expected Outcome Job (s) status is correctly fetched</p>
Pass/Fail Criteria	<p>Pass if the Appliance correctly manages jobs in the underlying execution manager. Tests must be executed (and pass) for each of the execution managers the appliance supports. All appliances should provide support for, at least one, of the following systems:</p> <ul style="list-style-type: none"> • Torque/PBS • LSF • SGE/OGE • Slurm <p>Optionally, the appliance may support a <i>fork</i> execution manager (spawning processes in the appliance host)</p>
Related Information	
Revision Log	V2: Major rewrite of criterion specification.

Information Retrieval	
ID	JOBEXEC_EXECMNGR_3
Description	Job Execution Appliances must be able to collect information from the underlying execution manager.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Support for the information retrieval from execution manager. Information should be returned as a valid GlueSchema representation.
Test Description	<p>Pre-condition Configured system</p> <p>Test Get information from execution manager</p> <p>Expected Outcome Representation of the current information from the execution manager is generated.</p>
Pass/Fail Criteria	<p>Pass if the Appliance produces information for each of the supported execution managers. The information must include all mandatory attributes of the Computing Element related entities in GlueSchema. All appliances should provide support for, at least one, of the following systems:</p> <ul style="list-style-type: none"> • Torque/PBS • LSF • SGE/OGE • Slurm <p>Optionally, the appliance may support a <i>fork</i> execution manager (spawning processes in the appliance host)</p>
Related Information	Information Capabilities QC
Revision Log	

1.4 Availability/Scalability

Service Redundancy	
ID	JOBEXEC_AVAIL_1
Description	More than one Job Execution Capability implementation should be able to access a single execution manager concurrently.
Mandatory	YES
Applicability	Job Execution Appliances
Input from Technology Provider	Documentation on how to use more than one appliance instance accessing the same execution manager (if any special consideration must be taken into account) Test of concurrent access to same execution manager from at least two instances.
Test Description	<p>Pre-condition More than one appliance instance configured to use the same execution manager</p> <p>Test Submission of jobs to all configured appliances</p> <p>Expected Outcome Jobs are executed without problems; they are not mixed up in any situation.</p>
Pass/Fail Criteria	Pass if the documentation specifies the configuration steps for using more than one instance in the same execution manager. Tests passes correctly
Related Information	
Revision Log	V2: Required documentation, changed ID

Self Disabling Mechanism	
ID	JOBEXEC_AVAIL_2
Description	The Job Execution Capability should detect high load conditions and self-disable the job submission in order to maintain the quality of the service.
Mandatory	NO
Applicability	Job Execution Appliances
Input from Technology Provider	Self-disable mechanism under high-load scenarios. Ideally, stress test for the service that triggers a self-disabling mechanism.
Test Description	<p>Pre-condition Correctly configured service.</p> <p>Test Introduce high load into machine, submit job.</p> <p>Expected Outcome High load situation is detected, job submission request is not allowed and message is sent to client.</p>
Pass/Fail Criteria	Pass if the test executes as expected. The high load level should be configurable (e.g. CPU load > x, swap usage > y...)
Related Information	
Revision Log	Changed ID

Timely Job Status Updates	
ID	JOBEXEC_AVAIL_4
Description	Job Execution Appliances should be able to report the job status within a reasonable time frame since the events that originate those statuses even in situations of high load
Mandatory	NO
Applicability	Job Execution Appliances
Input from Technology Provider	Appliance must be able to report the status of the submitted jobs without big delays from the event that originates the status change (e.g. mark the job as running/done once the job enters the running/done status in the local batch system). Ideally TP provides a test for the service that asserts that the appliance is able to report immediately the job statuses under high load conditions (big number of concurrent jobs changing status)
Pass/Fail Criteria	Appliances <i>should</i> be able to report the status immediately after the event that generated the status change.
Related Information	
Revision Log	

2 PARALLEL JOB

2.1 Submission of parallel jobs

The following tests propose example job descriptions using the gLite JDL format for the specification of jobs. These examples are just used for illustrative purposes. Each appliance should provide the tests using their native format.

Simple parallel job submission	
ID	PARALLEL_JOB_1
Description	Job Execution Appliances that also provide the Parallel Job Capability must allow users to submit a job requesting more than one execution slot.
Mandatory	YES
Applicability	Job Execution Appliances with Parallel Job Capability.
Input from Technology Provider	Support for the submission of parallel job, requesting more than 1 slot.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission: Executable = "/bin/sleep"; CPUNumber = 4; Arguments = "20";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job. Correct number of slots are allocated</p>
Pass/Fail Criteria	Test is executed correctly. Mapping of slots to machines/cores not relevant for the test.
Related Information	#1391: Support for parallel jobs in JDL.
Revision Log	V2: Unified PARALLEL_JOB_1, 3 & 4 into this criterion.

Single machine parallel job submission	
ID	PARALLEL_JOB_2
Description	Job Execution Appliances that also provide the Parallel Job Capability should allow users to submit a job requesting more than one execution slot in a single machine.
Mandatory	NO
Applicability	Job Execution Appliances with Parallel Job Capability.
Input from Technology Provider	Support for the submission of parallel job, requesting more than 1 slot in a single machine and for a complete machine.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission: Executable = "/bin/sleep"; NodeNumber = 1; SMPGranularity = 4; Arguments = "20";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job. Correct number of slots are allocated in a single machine</p>
	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission: Executable = "/bin/sleep"; NodeNumber = 1; SMPGranularity = 4; WholeNode = True; Arguments = "20";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job. Complete machine with the requested slots is allocated.</p>
Pass/Fail Criteria	Test is executed correctly.
Related Information	
Revision Log	V2: Unified PARALLEL_JOB_2 & 5.

Fine grained mapping parallel job submission	
ID	PARALLEL_JOB_3
Description	Job Execution Appliances that also provide the Parallel Job Capability should allow users to submit a job requesting a combination of slots per physical machine.
Mandatory	NO
Applicability	Job Execution Appliances with Parallel Job Capability.
Input from Technology Provider	Support for the submission of parallel job requesting specific configurations of slots in several machines.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission: Executable = "/bin/sleep"; NodeNumber = 5; SMPGranularity = 2; Arguments = "20";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job. Correct number of slots is allocated.</p>
Pass/Fail Criteria	Test is executed correctly for different combinations (e.g.: N processes in N different hosts, N processes in a single host, N processes per host in K hosts, K number of complete hosts with at least N slots)
Related Information	
Revision Log	V2: Unified PARALLEL_JOB_2 & 5.

2.2 MPI support

Precompiled MPI job Execution	
ID	PARALLEL_MPI_1
Description	Parallel Job Appliances must support the execution of MPI jobs.
Mandatory	YES
Applicability	Parallel Job Appliances.
Input from Technology Provider	Support for the submission of a MPI job with pre-existing binary.
Test Description	<p>Pre-condition Valid User proxy and valid delegation in the service. MPI Binary</p> <p>Test Submission of a MPI job requesting more than one execution slot with MPI Binary included in input sandbox of job or already installed in the system (description of job depending on Job Execution interface)</p> <p>Expected Outcome Job is submitted and executed without errors; the requested slots are allocated. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test is provided and passes for all the MPI implementations supported. Support for Open MPI and MPICH2 should be included
Related Information	User requirements: #672: MPI support
Revision Log	

MPI job Execution from source.	
ID	PARALLEL_MPI_2
Description	Parallel Job Appliances must support the execution of MPI jobs that are compiled at submission time.
Mandatory	YES
Applicability	Parallel Job Appliances.
Input from Technology Provider	Support for the submission of a MPI job compiled from source during its execution.
Test Description	<p>Pre-condition Valid User proxy and valid delegation in the service. Source code for MPI application.</p> <p>Test Submission of a MPI job requesting more than one execution slot with MPI source code included in input sandbox of job (description of job depending on Job Execution interface). Prior to the execution of the application, the source must be compiled with the available compiler at the site.</p> <p>Expected Outcome Job is submitted, compiled and executed without errors; the requested slots are allocated. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test is provided and passes for all the MPI implementations supported. Support for Open MPI and MPICH2 should be included
Related Information	User requirements: #672: MPI support
Revision Log	

2.3 OpenMP support

Precompiled OpenMP job Execution	
ID	PARALLEL_OMP_1
Description	Parallel Job Appliances must support the execution of OpenMP jobs.
Mandatory	YES
Applicability	Parallel Job Appliances.
Input from Technology Provider	Support for the submission of an OpenMP job with pre-existing binary.
Test Description	<p>Pre-condition Valid User proxy and valid delegation in the service. OpenMP Binary</p> <p>Test Submission of an OpenMP job requesting more than one execution slot with OpenMP Binary included in input sandbox of job (description of job depending on Job Execution interface)</p> <p>Expected Outcome Job is submitted and executed without errors; the requested slots are allocated. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test is provided and passes for all the OpenMP implementations supported.
Related Information	
Revision Log	

OpenMP job Execution from source	
ID	PARALLEL_OMP_2
Description	Parallel Job Appliances must support the execution of OpenMP jobs that are compiled at submission time.
Mandatory	YES
Applicability	Parallel Job Appliances.
Input from Technology Provider	Support for the submission of an OpenMP job that gets compiled at the remote site.
Test Description	<p>Pre-condition Valid User proxy and valid delegation in the service. Source code for OpenMP application.</p> <p>Test Submission of an OpenMP job requesting more than one execution slot with OpenMP source code included in input sandbox of job (description of job depending on Job Execution interface). Prior to the execution of the application, the source must be compiled with the available compiler at the site.</p> <p>Expected Outcome Job is submitted, compiled and executed without errors; the requested slots are allocated. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test is provided and passes for all the OpenMP implementations supported.
Related Information	
Revision Log	

3 INTERACTIVE JOB MANAGEMENT

Interactive login	
ID	INTERACTIVE_JOB_1
Description	Login interactively to a remote site using grid credentials
Mandatory	NO
Applicability	Interactive Job Management (Interactive Login)
Input from Technology Provider	Tool for providing interactive login to remote machine using any of the supported authn/authz in the UMD Roadmap.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Interactive login to remote site</p> <p>Expected Outcome Login is performed and a shell is provided.</p>
Pass/Fail Criteria	Pass if the tool is able to perform the remote logins correctly using the grid credentials
Related Information	gsissh, glogin UMD Roadmap Interactive Job Management [R 1]
Revision Log	

Interactive Job Perusal	
ID	INTERACTIVE_JOB_2
Description	Provide a mechanism for getting files produced by a job running in a remote site.
Mandatory	NO
Applicability	Interactive Job Management (Interactive Job Steering)
Input from Technology Provider	Mechanism that is able to retrieve the files produced by a job during its runtime. The provided service should be configurable to retrieve the files at periodic intervals of time. Files to retrieve <i>should</i> be configurable.
Pass/Fail Criteria	Pass if the provided service is able to retrieve at periodic intervals job output files during the job execution.
Related Information	WMS Job Perusal UMD Roadmap Interactive Job Management [R 1]
Revision Log	

Interactive Job Monitoring	
ID	INTERACTIVE_JOB_3
Description	Provide a mechanism for streaming files produced by a job running in a remote site.
Mandatory	NO
Applicability	Interactive Job Management (Interactive Job Steering)
Input from Technology Provider	Mechanism that is able to stream the files produced by a job during its runtime. Ideally, the files to stream should be configurable. By default the standard output and error of the job should be used.
Pass/Fail Criteria	Pass if the provided service is able to stream the job output files during the job execution.
Related Information	globus-job-get-output, i2glogin UMD Roadmap Interactive Job Management [R 1] #1385: Interactive jobs monitoring
Revision Log	

Interactive Job Steering	
ID	INTERACTIVE_JOB_4
Description	Provide a mechanism for steering a job running in a remote site.
Mandatory	NO
Applicability	Interactive Job Management (Interactive Job Steering)
Input from Technology Provider	Mechanism that is able to stream the files produced by a job during its runtime and to control the job execution (i.e. stream the job's standard input from the user location to the remote site).
Pass/Fail Criteria	Pass if the provided service is able to control the job execution by creating a communication channel that forwards output/error and input streams between the user and the remote job
Related Information	i2glogin UMD Roadmap Interactive Job Management [R 1]
Revision Log	

4 JOB SCHEDULING

4.1 Job Scheduling Interface

The Job Scheduling Capabilities does not have a standard interface. Any implementation of this capability can support on of the Job Execution interfaces proposed by the OGF (DRMAA, BES) or proprietary interfaces (gLite WMS)

Job Scheduling Interface	
ID	JOBSCH_IFACE_1
Description	Job Scheduling Appliances must support one of the interfaces currently in use or identified by the UMD Roadmap
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Implementation of one of the Job Scheduling Interfaces as defined in the UMD Roadmap. Ideally, a complete test suite of the Job Execution interfaces supported by the appliance. The test suite must include tests for all the documented functions, and for all functions, check both correct and invalid input and with valid and invalid credentials.
Test Description	<p>Pre-condition Valid user credentials.</p> <p>Test Test all interface functionality, with correct/incorrect input and with valid and invalid credentials.</p> <p>Expected Outcome Log of all the operations performed. All the documented functions work as documented.</p>
Pass/Fail Criteria	<p>The Job Scheduling Appliance that claims to support an interface must pass complete tests for that interface (provided by the TP or by the verification team). If the API is not completely supported, this must be documented. The test suite must be executed without errors.</p> <p>At least one of the following interfaces must be provided:</p> <ul style="list-style-type: none"> • gLite WMS [R 17] • OGF DRMAA [R 14] • OGSA BES [R 15]
Related Information	UMD Roadmap Job Scheduling Capability
Revision Log	V2: Merged all the interface related criteria into this.

4.2 Job Execution Capability Support

Remote Job Management	
ID	JOBSCH_EXEC_1
Description	Job Scheduling Appliances must support the creation and management of work items to an Job Execution Appliance
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	<p>Appliance must be able to:</p> <ul style="list-style-type: none"> • create new jobs • retrieve the status of the jobs submitted by the appliance • cancel jobs • optionally, hold and resume jobs <p>The Appliance may perform these operations for individually for each submitted job or for set of jobs in order to improve its performance (e.g. for retrieving the status instead of querying each of the individual jobs, do a single query for all jobs submitted at a given appliance)</p>
Test Description	<p>Pre-condition Configured system</p> <p>Test Create new job(s) in job execution appliance</p> <p>Expected Outcome New job(s) is created in the job execution appliance; id of job(s) returned</p>
	<p>Pre-condition Previously submitted job(s)</p> <p>Test Cancel job(s) in job execution appliance.</p> <p>Expected Outcome Job(s) is cancelled successfully.</p>
	<p>Pre-condition Previously submitted job(s)</p> <p>Test Query status of previously submitted job(s)</p> <p>Expected Outcome Job (s) status is correctly fetched</p>
Pass/Fail Criteria	<p>Pass if the Appliance correctly manages jobs in the job execution appliances. Tests must be executed (and pass) for each of the job execution appliances supported.</p> <p>At least one of the following interfaces must be supported:</p> <ul style="list-style-type: none"> • ARC-CE gridFTP [R 10] • CREAM [R 11] • EMI-ES [R 12] • Globus GRAM5 [R 13] • OGF DRMAA [R 14] • OGSA BES [R 15] • UNICORE UAS [R 16]



Related Information	UMD Roadmap Job Execution QC
Revision Log	V2: Major rewrite of criterion specification.

Remote Resource Information	
ID	JOBSCH_EXEC_2
Description	Job Scheduling Appliances must be able to use the resource descriptions using the current Information Model and Information Discovery interfaces.
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Appliances must handle resources described with the current Information Model (GlueSchema1.3 and optionally GlueSchema2) and Information Discovery (LDAPv3) interfaces.
Test Description	<p>Pre-condition Configured system</p> <p>Test Fetch information from Information Discovery Appliance.</p> <p>Expected Outcome Information is fetched correctly; resources described are added to the list of possible resources to use.</p>
Pass/Fail Criteria	Pass if the Appliance correctly fetches information from Information Discovery appliances and is able to use the resources described by GlueSchema v1.3 and/or GlueSchema v2.
Related Information	Information Capabilities in the UMD Roadmap [R 1]
Revision Log	

4.3 End-to-end job submission tests

The following tests propose example job descriptions using the gLite JDL format for the specification of jobs. These examples are just used for illustrative purposes. Each appliance should execute the tests using their native format.

Simple Job	
ID	JOBSCH_JOB_1
Description	Execute a simple job.
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Support for the submission of a job with no input or output files.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission of simple job: Executable = /bin/sleep; Arguments = "120";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

Simple Job with input/output files	
ID	JOBSCH_JOB_2
Description	Execute a simple job that uses both input and output files.
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Support for the submission of a job with input or output files.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system) Non-empty file "myfile"</p> <p>Test Job submission for job with input and output files: Executable = "/bin/ls"; Arguments = "-l"; StdOutput = "std.out"; StdError = "std.err"; InputSandbox = {"myfile"}; OutputSandbox = {"std.out", "std.err"};</p> <p>Expected Outcome Job finishes correctly; output contains the listing of the directory including the input file with correct size. Unique Identifier for the submitted jobs, status log of the job.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic.

Cancel Job	
ID	JOBSCH_JOB_3
Description	Cancel a previously submitted job.
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Support for the cancellation of a job. Job cancelling must be supported for the different states that the job may be, e.g. cancel the job when it's running or cancel the job when it's already done.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job Submission and then cancellation. Possible description for job: Executable = "/bin/sleep"; Arguments = "20m";</p> <p>Expected Outcome Job is submitted and then cancelled correctly. Unique Identifier for the submitted jobs, status log of the job. Job is removed from remote Job Execution Appliance.</p>
Pass/Fail Criteria	Pass if the appliance is able to cancel jobs for any previous state of the job. If the job is already submitted to a Job Execution Appliance, it should be completely removed from it, especially if it's running.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

Parallel Job	
ID	JOBSCH_JOB_4
Description	Execute a parallel job.
Mandatory	NO
Applicability	Job Scheduling Appliances with Parallel Job Support.
Input from Technology Provider	Support for the submission of a job with input or output files.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job Submission or parallel job. Possible description for job: Executable = "/bin/sleep"; CPUNumber = 2; Arguments = "20";</p> <p>Expected Outcome Job finishes correctly. Unique Identifier for the submitted jobs, status log of the job. Correct number of slots is allocated at the remote site.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

Job List Match	
ID	JOBSCH_JOB_5
Description	List the available resources for a given job.
Mandatory	YES
Applicability	Job Scheduling Appliances
Input from Technology Provider	Support for the list match of a job.
Test Description	<p>Pre-condition Valid user credentials and delegation in the service.</p> <p>Test Job list match for job with requirements and rank expressions, for example:</p> <pre>Executable = "/bin/sleep"; Requirements = other.GlueCEStateStatus = "Production"; Rank = -other.GlueCEStateEstimatedResponseTime;</pre> <p>Expected Outcome List of available resources for execution (with correct rank) is returned.</p>
Pass/Fail Criteria	The Job Scheduling Appliance must return a list of available resources for the execution of any given job. Optionally, a <i>rank</i> defined by the user is returned by each of the resources.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

Parametric Job Submission		
ID	JOBSCH_JOB_6	
Description	Execute a parametric job.	
Mandatory	NO	
Applicability	Job Scheduling Appliances with support for parametric jobs.	
Input from Technology Provider	Support for the submission of parametric jobs.	
Test Description	Pre-condition Valid user credentials (and delegation if needed in the system) Test Job submission of job with numeric parameters (e.g. Parameters = 10000;ParameterStart = 1000; ParameterStep = 10;). Expected Outcome Job is executed correctly. List of JobIds for the parametric jobs and each of the subjobs is obtained; all states of the jobs must be logged correctly.	
	Pre-condition Valid user credentials (and delegation if needed in the system) Test Job submission of job with a list of parameters (e.g. Parameters={A, B, C,...}). Expected Outcome Job is executed correctly. List of JobIds for the parametric jobs and each of the subjobs is obtained; all states of the jobs must be logged correctly.	
	Pass/Fail Criteria	Pass if the test passes correctly.
	Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling	

Job Collection Submission	
ID	JOBSCH_JOB_7
Description	Execute a job collection
Mandatory	NO
Applicability	Job Scheduling Appliances with support for job collections.
Input from Technology Provider	Support for the submission of job collections.
Test Description	<p>Pre-condition Valid user credentials (and delegation if needed in the system)</p> <p>Test Job submission for job collection.</p> <p>Expected Outcome Job is executed correctly. List of JobIds for the job collections and each of the subjobs is obtained; all states of the jobs must be logged correctly.</p>
Pass/Fail Criteria	Pass if the test passes correctly.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

DAG Submission	
ID	JOBSCH_JOB_8
Description	Execute a DAG job.
Mandatory	NO
Applicability	Job Scheduling Appliances with support for DAGs.
Input from Technology Provider	Support for the submission of DAGs.
Test Description	<p>Pre-condition Valid user credentials and delegation in the service.</p> <p>Test Job submission for DAG.</p> <p>Expected Outcome Job is executed correctly. List of JobIds for DAG and each of the subjobs is obtained; all states of the jobs must be logged correctly.</p>
Pass/Fail Criteria	Pass if the test passes correctly. DAGs must be able to use any of the Job Execution Interfaces supported by the Job Scheduling Appliance. Explicit test this possibility.
Related Information	
Revision Log	V2: moved specific WMS criteria to generic to all Job Scheduling

4.4 gLite WMS

This section includes criteria applicable to the gLite WMS system.

Proxy Renewal		
ID	JOBSCH_WMS_1	
Description	The WMS must manage the user credentials and renew them if necessary.	
Mandatory	YES	
Applicability	gLite WMS Job Scheduling Appliances.	
Input from Technology Provider	Support for the proxy renewal mechanism for long running jobs.	
Test Description	Pre-condition Valid user credentials with short duration (e.g. 30 min) and delegation in the service. Credentials Renewal service available. Test Submit job that takes longer to complete than the credential lifetime (e.g. 1 hour) Expected Outcome Job executes successfully. The scheduling services should perform a proxy renewal and state it in the log messages (if there is an error, log it also). Output of the job, and status messages stating the renewal of the user credentials.	
	Pre-condition Valid user credentials with short duration, e.g. 30 min, no renewal service. Test Submit job that takes longer to complete than the credential lifetime (e.g. 1 hour) Expected Outcome Job does not complete successfully. Log of operations and status of the job updated with information about the error (no renewal possible)	
	Pass/Fail Criteria	Will Pass if the proxy renewal is done, or if there is an error logged stating the problem. Will fail if there is no clear information about the process.
	Related Information	
Revision Log		

Job Resubmission	
ID	JOBSCH_WMS_2
Description	Any job failures (due to resource malfunctioning or the job itself) must be resubmitted with a configurable amount of retries.
Mandatory	NO
Applicability	gLite WMS Job Scheduling Appliances.
Input from Technology Provider	Support for the resubmission mechanism of the WMS.
Test Description	Pre-condition Valid user credentials and delegation in the service. Test Job submission that fails due to simulated remote resource malfunctioning. Expected Outcome Job is resubmitted to other resource. Log of all failures and a complete trace of the job.
	Pre-condition Valid user credentials and delegation in the service. Test Job submission for job that always fails (e.g. exit code 1) Expected Outcome Job is resubmitted until resubmission attempts reach the configured limit. Log of all failures and a complete trace of the job.
Pass/Fail Criteria	Job failures due to resource malfunctioning and not to the job itself must be resubmitted to other resources, with a configurable amount of repetitions. In the case of job failures due to the job itself must be resubmitted with a configurable amount of repetitions. In both situations, status must reflect clearly what is the cause of resubmission, new resource selected and attempt number
Related Information	Requirements gathered in MS305 related to resubmission of jobs, and information provided in error messages.
Revision Log	V2: originally JOBEXEC_WMS_JOB_9

JDL Acceptance Limits	
ID	JOBSCH_WMS_3
Description	The service should accept JDLs without size restrictions
Mandatory	NO
Applicability	gLite WMS Job Scheduling Appliances.
Input from Technology Provider	A test to submit a job and check if it is accepted or rejected, specially for big JDLs.
Test Description	<p>Pre-condition Valid user credentials and delegation in the service.</p> <p>Test Submission of job descriptions (specially large)</p> <p>Expected Outcome Normal job submission if everything is correct; an error message if any problem arises.</p>
Pass/Fail Criteria	Will Pass if JDL is correct, and submits the job or if there is a report on a known syntax error in the jdl. Will Fail if a wrong Jdl is accepted or if it crashes
Related Information	Requirements gathered in MS305 related to resubmission of jobs, and information provided in error messages.
Revision Log	V2: originally JOBEXEC_WMS_JOB_10

4.4.1 Security Advisories

Security Advisory 1502	
ID	JOBSCH_WMS_SEC_1
Description	Steal of proxies is possible without leaving trace.
Mandatory	YES
Applicability	gLite WMS Job Scheduling Appliances.
Input from Technology Provider	Test that assures the problem described in the SVG Advisory 1502 (proxy stealing) is fixed.
Pass/Fail Criteria	Fix for Advisory-SVG-2011-1502 is provided. A test that proves that the fix is provided should be also present.
Related Information	Advisory-SVG-2011-1502 (https://wiki.egi.eu/wiki/SVG:Advisory-SVG-2011-1502)
Revision Log	

4.4.2 Bugs

Long Proxy Chain Support	
ID	JOBSCH_WMS_BUG_1
Description	Long proxy chains should be supported without no issues.
Mandatory	YES
Applicability	gLite WMS Job Scheduling Appliances.
Input from Technology Provider	Support for long proxy chains such as the ones created when using myproxy (C=[...]/CN=proxy/CN=proxy/CN=proxy/CN=proxy)
Test Description	<p>Pre-condition Valid authorized user credentials with long proxy chain.</p> <p>Test Delegation of proxy into service.</p> <p>Expected Outcome Delegation is performed without issues.</p>
Pass/Fail Criteria	No authorization errors (for authorized users) given when using long proxy chains.
Related Information	GGUS Ticket: #73035
Revision Log	

4.5 Service availability, monitoring and error handling.

Error Messages	
ID	JOBSCH_SERVICE_1
Description	Error messages provided by the service should be clear and facilitate the solution of those errors by users or service administrators
Mandatory	NO
Applicability	Job Scheduling Appliances.
Input from Technology Provider	Include in documentation, a list of possible errors and possible solution/cause for it. For errors that may reach the user, this list has to be exhaustive.
Pass/Fail Criteria	Will pass if the list of errors is documented and includes information about: <ul style="list-style-type: none"> • Error code • Error message (if applicable) • Error source (internal module or remote resource (specify it explicitly)) • Cause of error (syntax error, module malfunctioning, configuration problem, network error, other (specify it explicit)) • Type (critical, informative) • Possible solution
Related Information	Requirements gathered in MS305 related to resubmission of jobs, and information provided in error messages.
Revision Log	

Service Information	
ID	JOBSCH_SERVICE_2
Description	Job Scheduling Appliances must be able to generate information about the provided service that can be used in a Information Discovery Appliance.
Mandatory	NO
Applicability	Job Scheduling Appliances.
Input from Technology Provider	Support for information generation about the service status.
Test Description	<p>Pre-condition Configured system, Information Discovery appliance available.</p> <p>Test Generate service information and publish to Information Discovery Appliance. Access Info Discovery Appliance.</p> <p>Expected Outcome Information is produced and can be accessed through the Information Discovery Appliance.</p>
Pass/Fail Criteria	Test is provided and executed as expected.
Related Information	Requirements gathered in MS305 related to resubmission of jobs, and information provided in error messages.
Revision Log	

Self Disabling Mechanism	
ID	JOBSCH_SERVICE_3
Description	The Job Scheduling Capability should detect high load conditions and self-disable the job submission in order to maintain the quality of the service.
Mandatory	NO
Applicability	Job Scheduling Appliances
Input from Technology Provider	Support for self-disabling mechanism under high load conditions. Ideally, stress test for the service that triggers a self-disabling mechanism.
Test Description	<p>Pre-condition Correctly configured service.</p> <p>Test Introduce high load into machine, submit job.</p> <p>Expected Outcome High load situation is detected, job submission request is not allowed and message is sent to client.</p>
Pass/Fail Criteria	Pass if the test executes as expected. The high load level should be configurable (e.g. CPU load > x, swap usage > y...)
Related Information	User requirements: #698: WMS stability and performance #702: Stability of UMD services and improvements
Revision Log	V2: Changed ID (from JOBSCH_SERVICE_4 to JOBSCH_SERVICE_3)

Job Submission Peaks	
ID	JOBSCH_SERVICE_4
Description	Job Scheduling Appliances should be able to handle high job submission rates of several hundreds jobs in short intervals.
Mandatory	NO
Applicability	Job Scheduling Appliances
Input from Technology Provider	Appliance should be able to handle a high number of jobs submitted in a short time interval (e.g. 500 jobs / minute). Ideally, test the service to assert that this is provided
Pass/Fail Criteria	Appliances should be able to handle job bursts of several hundreds of jobs in short intervals.
Related Information	User requirements: #698: WMS stability and performance
Revision Log	

Timely Job Status Updates	
ID	JOBSCH_SERVICE_5
Description	Job Scheduling Appliances should be able to report the job status within a reasonable time frame since the events that originate those statuses even in situations of high load
Mandatory	NO
Applicability	Job Execution Appliances
Input from Technology Provider	Appliance must be able to report the status of the submitted jobs without big delays from the event that originates the status change (e.g. mark the job as running/done once the job enters the running/done status in the local batch system). Ideally TP provides a test for the service that asserts that the appliance is able to report immediately the job statuses under high load conditions (big number of concurrent jobs changing status)
Pass/Fail Criteria	Appliances <i>should</i> be able to report the status immediately after the event that generated the status change.
Related Information	User requirements: #698: WMS stability and performance.
Revision Log	

5 REFERENCES

R 1	UMD roadmap: https://documents.egi.eu/public/ShowDocument?docid=100
R 2	Web Services Data Access and Integration – The Relational Realisation (WS-DAIR) Specification, Version 1.0
R 3	Web Services Data Access and Integration – The XML Realization (WS-DAIX) Specification, Version 1.0
R 4	OGSA-DAI: http://www.ogsadai.org.uk/
R 5	gLite LFC: https://twiki.cern.ch/twiki/bin/view/EGEE/GliteLFC
R 6	AMGA: http://amga.web.cern.ch/amga/
R 7	AMGA WSDL: http://amga.web.cern.ch/amga/soap_wsdaire.html
R 8	AMGA streaming API: http://amga.web.cern.ch/amga/protocol.html
R 9	AMGA Metadata Queries: http://amga.web.cern.ch/amga/queries.html
R 10	A. Konstantinov, ARC Computational Job Management Component – A-REX, NORDUGRID-TECH-14
R 11	CREAM: http://grid.pd.infn.it/cream/
R 12	EMI-ES: https://twiki.cern.ch/twiki/bin/view/EMI/EmiExecutionService
R 13	GRAM5: http://www.globus.org/toolkit/docs/latest-stable/execution/gram5/
R 14	OGF DRMAA: http://www.drmaa.org/
R 15	OGSA Basic Execution Service v1.0: http://www.ogf.org/documents/GFD.108.pdf
R 16	UNICORE UAS: http://www.unicore.eu/unicore/architecture/service-layer.php#anchor_uas
R 17	gLite WMS: http://web.infn.it/gLiteWMS/
R 18	SAGA-CORE-WG: A Simple API for Grid Applications (SAGA) v1.0 (GFD.90)
R 19	SAGA (A Simple API for Grid Applications): http://saga.cct.lsu.edu/
R 20	Instrument Element: http://www.dorii.eu/resources/adaptation:middleware:IE
R 21	DORII (Deployment of Remote Instrumentation Infrastructure) Project: http://www.dorii.eu/
R 22	GlueSchema Specification v1.3: http://glueschema.forge.cnaf.infn.it/Spec/V13

R 23	GlueSchema Specification v2.0: http://www.ogf.org/documents/GFD.147.pdf
R 24	JMS (Java Message Service Specification) 1.1: http://www.oracle.com/technetwork/java/jms/index.html
R 25	AMQP (Advanced Message Queuing Protocol): http://www.amqp.org/confluence/display/AMQP/Advanced+Message+Queuing+Protocol
R 26	Nagios Config Generator: https://tomtools.cern.ch/confluence/display/SAM/NCG
R 27	My EGI portal: https://tomtools.cern.ch/confluence/display/SAM/MyEGI
R 28	SAM Probes Documentation: https://tomtools.cern.ch/confluence/display/SAM/Probes
R 29	Accounting Portal: http://accounting.egi.eu/
R 30	GridSite Delegation Protocol: http://www.gridsite.org/wiki/Delegation_protocol
R 31	Globus Delegation Service: http://www.globus.org/toolkit/docs/4.0/security/delegation/
R 32	European Policy Management Authority for Grid Authentication (EuGridPMA): http://www.eugridpma.org/
R 33	ARGUS Authorization Service: https://twiki.cern.ch/twiki/bin/view/EGEE/AuthorizationFramework
R 34	XACML: http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf
R 35	Hydra encrypted file storage: https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS
R 36	gLite FTS: https://twiki.cern.ch/twiki/bin/view/EGEE/GLiteFTS
R 37	SRM v2.2: http://www.ggf.org/documents/GFD.129.pdf
R 38	S2 Test: http://s-2.sourceforge.net/
R 39	SRM-Tester: https://sdm.lbl.gov/twiki/bin/view/Software/SRMTester/WebHome
R 40	Lcg-utils: http://grid-deployment.web.cern.ch/grid-deployment/documentation/LFC_DPM/lcg_util/
R 41	Lcg-utils test suite: http://glite.cvs.cern.ch/cgi-bin/glite.cgi/org.glite.testsuites.ctb/UI/tests/test-lcg-utils.sh?view=markup
R 42	Open Cloud Computing Interface WG, OGF, http://www.ggf.org/gf/group_info/view.php?group=occi-wg
R 43	Virtualization Management (VMAN), DMTF http://www.dmtf.org/standards/vman
R 44	StratusLab http://stratuslab.eu/



R 45	StratusLab MarketPlace Technical Note TN-Marketplace (V3.0)
-------------	---