

Subject: **JDL Attributes**

Author: **Fabrizio Pacini (fabrizio.pacini@datamat.it)**

Partner: **Datamat SpA**

Diffusion:

Information: **DataGrid-01-NOT-0101-0_7-Note**

1. INTRODUCTION

The JDL is a fully extensible language, hence it is allowed to use whatever attribute for the description of a job. Anyway only a certain set of attributes that we will refer as “supported attributes” from now on, is taken into account by the Workload Management System components in order to schedule a submitted job.

The supported attributes can be grouped into two main categories:

- Resources attributes
- Job attributes

Resource attributes are those that have to be used to build expressions of the **Requirements** and **Rank** attributes in the job class-ad and to be effective, i.e. to be actually used for selecting a resource, have to belong to the set of characteristics of the resources that are published in the GIS (aka MDS).

Job attributes represent instead job specific information and specify in some way actions that have to be performed by the RB to schedule the job. Some of these attributes are provided by the user when he/she edits the job description file while some other (needed by the RB) are inserted by the UI before submitting the job.

A small subset of the attributes that are inserted by the user are mandatory, i.e. necessary for the RB to work correctly and can be split in two categories:

- *Mandatory*: the lack of these attributes does not allow the submission of the job
- *Mandatory with default value*: the UI provides default value for these attributes if they are missing in the job description.

Next sections of this note provide a list of JDL supported attributes specifying their characteristics in relation to what just discussed.

2. JOB ATTRIBUTES PROVIDED BY THE USER

In the following Table 1 the column **M** indicates those attributes that are mandatory. Default values (indicated in the **with default** column) are assigned by the UI on the basis of what specified in a UI configuration file.

Attribute	M	With default	Meaning
Executable	✓		<p>Executable/command name. The user can specify an executable that lies already on the remote CE. The absolute path, possibly including environment variables, of this file should be specified. The other possibility is to provide a local executable name, which will be staged on the CE. In this case only the file name has to be specified as executable. The absolute path on the local file system executable should be then listed in the <i>InputSandbox</i> attribute expression. It is important to remark that if the job needs for the execution some command line arguments, they have to be specified through the <i>Arguments</i> attribute.</p>
Arguments			<p>This is a string containing all the job command line arguments. E.g. an executable <i>sum</i> that has to be started as:</p> <pre>\$ sum N1 N2 -out result.out</pre> <p>is described by:</p> <pre>Executable = "sum"; Arguments = "N1 N2 -out result.out";</pre> <p>If you want to specify a quoted string inside the <i>Arguments</i> then you have to escape quotes with the <code>\</code> character. E.g. when describing a job like:</p> <pre>\$ grep -i "my name" *.txt</pre> <p>you will have to specify:</p> <pre>Executable = "/bin/grep"; Arguments = "-i \"my name\" *.txt";</pre> <p>Analogously, if the job takes as argument a string containing a special character (e.g. the job is the <i>tail</i> command issued on a file whose name contains the quotes character, say <i>file1&file2</i>), since on the shell line you would have to write:</p> <pre>\$ tail -f file1\&file2</pre>

Attribute	M	With default	Meaning
			<p>in the JDL you'll have to write:</p> <pre>Executable = "/usr/bin/tail"; Arguments = "-f file1\\\\"&file2";</pre> <p>i.e. a \ for each special character.</p> <p>In general, special characters such as &, , >, < are only allowed if specified inside a quoted string or preceded by triple \.</p> <p>The character "" cannot be specified in the JDL.</p>
InputData			<p>A list of:</p> <ul style="list-style-type: none"> - logical file names and/or - physical file names <p>This attribute refers to data used as input by the job; these data are stored in SEs and published in replica catalogues.</p> <p>Listed names have to be prefixed with "LF:" and "PF:" to indicate that they are respectively: logical file names and physical file names. E.g.:</p> <pre>InputData = {"LF:<LFN1>", "PF:<PFN>", "LF:<LFN2>"};</pre>
StdInput			<p>Standard input of the job. It can be:</p> <ul style="list-style-type: none"> - just a file name (staging required) - absolute path (available on the CE) <p>The same mechanism as described for the <i>Executable</i> attribute can be applied.</p>
StdOutput			<p>Standard output of the job. The user has to specify just the file name. To have this file staged back on the submitting machine he/she has to list the file name also in the <i>OutputSandbox</i> attribute expression and use the dg-job-get-output command.</p>

Attribute	M	With default	Meaning
StdError			Standard error of the job. The user has to specify just the file name. To have this file staged back on the submitting machine he/she has to list the file name also in the <i>OutputSandbox</i> attribute expression and use the dg-job-get-output command.
OutputSE			URI of the Storage Element where to store the output data. Once specified, this attribute is used by the RB to choose a CE being "close" to this SE comparing it with the <i>CloseSE</i> attribute published in the GIS. E.g.: <i>OutputSE = "grid001.cnaf.infn.it";</i>
InputSandbox			List of files on the UI local disk needed by the job for running. The listed files are staged from the UI to the remote CE. Wildcards and environment variables are admitted in the specification of this attribute. File names can to be provided as absolute paths or relative paths starting from the cwd. This attribute is also used to accomplish executable and standard input staging from the submitting machine to the remote CE where job execution takes place. It is important to note that since <code>globus-url-copy</code> (the Globus command used for the <i>InputSandbox</i> files staging) in general doesn't preserve the <code>x</code> flag, the script specified as Executable in the JDL (on which <code>chmod +x</code> is done automatically by the WP1 JobWrapper), should perform a <code>chmod +x</code> for all the files needing execution permission, that are transferred within the <i>InputSandbox</i> of the job.
OutputSandbox			List of files, generated by the job, which have to be retrieved. The listed files are transferred on the UI local file system by mean of the dg-job-get-output command. Wildcards are admitted in the specification of this attribute. The list shall contain file names (neither absolute nor relative paths are admitted).

Attribute	M	With default	Meaning
ReplicaCatalog	✓ (*)		<p>Replica Catalogue Identifier, i.e. something in the following format:</p> <pre><protocol>://<full hostname> :<port>/<Replica Catalog DN>.</pre> <p>where the Replica Catalogue DN also comprises the mandatory logical collection field <i>lc</i>. I.e. it is something like:</p> <pre>lc=<Logical collection>, rc=<replica catalogue>, dc=...</pre> <p>Hereafter is reported an example of Replica Catalogue address:</p> <pre>ldap://sunlab2f.cnaf.infn.it:2010/lc=test0, rc=WP2 INFN Test Replica Catalog, dc=sunlab2g, dc=cnaf, dc=infn, dc=it</pre> <p>(*) This attribute is mandatory iff the <i>InputData</i> attribute has been also specified and contains at least one LFN.</p>
DataAccessProtocol	✓ (*)		<p>This is the protocol or the list of protocols that the application is able to "speak" for accessing <i>InputData</i> on a given SE. The RB matches indeed this attribute with the <i>SEProtocol</i> attribute of published in the IS. E.g.:</p> <pre>DataAccessProtocol = {"file", "gridftp"};</pre> <p>(*) This attribute is mandatory iff the <i>InputData</i> attribute has been also specified.</p>
Rank	✓	-other.Estimated TraversalTime	<p>A ClassAd Floating-Point expression that states how to rank queues that have already met the <i>Constraints</i> expression. Essentially, rank expresses a preference. A higher numeric value equals a better rank. The RB will give to the job the queue with the highest rank. Default value for this attribute is:</p> <pre>-other.EstimatedTraversalTime</pre> <p>The default value is configurable through the UI configuration file <i>UI_ConfigENV.cfg</i></p>
Requirements	✓	TRUE	<p>Boolean ClassAd expression that uses C-like operators. It represents job requirements on resources. To have a job scheduled to run on a</p>

Attribute	M	With default	Meaning
			<p>given queue, this Requirements expression must evaluate to true on the given queue. Default value for this attribute is TRUE.</p> <p>The default value is configurable through the UI configuration file <i>UI_ConfigENV.cfg</i></p>
Environment			<p>This is a list of string representing environment settings that have to be performed on the submitting machine and are needed by the job to run properly. Each item of the list is an equality "VAR_NAME=VAR_VALUE". E.g.:</p> <p><i>Environment =</i> <i>{"JOB_LOG=/tmp","CNF_PATH=/opt/edg/etc"};</i></p>
RetryCount			<p>It is a positive integer.</p> <p>The <i>RetryCount</i> attribute allows setting the number of submission retries for a job upon failure due to some grid component (i.e. not to the job itself). <i>RetryCount</i> has to be a positive number and the actual number of submission retries for a job is represented by the minimum value between <i>RetryCount</i> itself and the value of the <i>RB_submission_retries</i> parameter in the RB configuration file. The resubmission is tried for all the CEs satisfying the job requirements.</p>

Table 1

3. JOB ATTRIBUTES PROVIDED BY THE UI

Attribute	Meaning
dg_jobId	<p>Grid-wide unique job identifier assigned by the UI to the job before submission. Format of the ob identifier is <code><LBname>/<UIname>/<time><PID><RND>?<RBname></code> where</p> <ul style="list-style-type: none"> - <i>LBname</i> is the LB server name and port (protocol is https) - <i>UIname</i> is the UI machine IP address or FQDN - <i>time</i> is the current time on the submitting machine in <i>hhmmss</i> format - <i>PID</i> is the UI process (dg-job-submit) identifier - <i>RND</i> is a random number generated at each job submission - <i>RBname</i> is the RB hostname and port
CertificateSubject	<p>Subject of the X509 user credentials used for submitting the job. The user's proxy certificate is searched in the file indicated by X509_USER_PROXY environment variable. If the variable is not set the default is:</p> <p><code>/tmp/x509up_u<UID></code></p> <p>This attributes is used for the authorization check by the RB that matches it against the list of users authorized to submit job to the CE (the <i>AuthorizedUsers</i> resource attribute published in the IS) and with the one it takes form the credentials exchanged during the authentication hand-shake done with the UI.</p>
UserContact	<p>This is a valid e-mail address where the job status changes notifications have to be sent. This attribute is set by the UI when the user issues the dg-job-submit command with <i>-notify</i> option.</p>
SubmitTo	<p>Value for this attribute has to be the DataGrid-wide unique identifier of a resource published in the IS. This attribute is set by the UI when the user issues the dg-job-submit command with <i>-resource</i> option and makes the RB directly submit the job to the specified resource completely skipping the matchmaking process. The accepted format for a CE identifier is:</p> <p><code><full-hostname>:<port-number>/jobmanager-<service>-<queuename></code></p>

Attribute	Meaning
	<p>where supported services are currently: <i>lsf, pbs, bqs</i>.</p> <p>It is important to remark that the <i>SubmitTo</i> is a <i>job attribute</i> that can only be inserted by the UI. Indeed if <i>SubmitTo</i> is found in the JDL, it is discarded and not passed to the RB. The user has to rely on the <i>--resource</i> option of dg-job-submit to specify direct submission to a specific CE.</p> <p>It is important to note that when the <i>--resource</i> option is used, the RB does not generate the <i>BrokerInfo</i> file also if data requirements have been specified in the JDL, so jobs submitted using this option should not rely on the <i>BrokerInfo</i> file information when running on the CE.</p> <p>A way for performing direct submission to a given CE and at the same time having the <i>BrokerInfo</i> file generated by RB and shipped to the CE is not using the <i>--resource</i> option and specify the following requirements in the JDL:</p> <p><i>Requirements = other.CEId == <Ce_identifier>;</i></p>

Table 2

4. RESOURCES ATTRIBUTES

In this section (Table 3, Table 4, Table 5 and Table 6) are reported the Computing Element, Close Storage Element, Storage Element and Storage Element Protocol entities attributes. For completeness all resource attributes published in the MDS have been included in the list, anyway some of them (they have been greyed in the text) shall not be used by the user to build the `Requirements` or `Rank` expression since they are automatically taken into account by the RB for carrying out the match-making algorithm. It is also reminded that resource attributes, when inserted in the `Requirements` or `Rank` expression have to be prefixed with “*other.*” in order to allow a correct matchmaking.

CE Attribute	Meaning
LRMSType ^(§)	<p>Defines the type of the local resource management system (e.g. LSF, Condor, PBS...).</p> <p>(§) This attribute is defined only when the Computing Element is a queue of a LRMS.</p>
LRMSVersion ^(§)	<p>The version of the local resource management system.</p> <p>(§) This attribute is defined only when the Computing Element is a queue of a LRMS.</p>
QueueName ^(§)	<p>Defines the name of the queue in the LRMS.</p> <p>(§) This attribute is defined only when the Computing Element is a queue of a LRMS.</p>
GlobusResourceContactString	<p>This attribute represents the Globus resource contact string that identifies this Globus resource (e.g. <i>pcgrid01.pd.infn.it:2119/jobmanager-lsf</i>).</p>
CEId	<p><i>CEId</i> is a string, univocally identifying the CE published in the Grid Information Space.</p> <p>The <i>CEId</i> format is:</p> <p><code><full-hostname>:<port-number>/jobmanager-<service>-<queue-name></code></p> <p>where supported services are currently: <i>lsf</i>, <i>pbs</i>, <i>bgs</i> (i.e. this value can be obtained “combining” the <i>GlobusResourceContactString</i> and <i>QueueName</i> attributes).</p> <p>We assume that WP4 will provide the Grid Information Space with this appropriate value.</p>
GRAMVersion	the GRAM version.

CE Attribute	Meaning
Architecture	the architecture of the machine or of the machines associated to the queue (we assume that all the machines "belonging" to the queue have the same architecture). E.g.: <i>INTEL, SPARC etc.</i>
OpSys	the operating system type and version of the machine or of the machines associated to the queue (assuming that all these machines run the same operating system). E.g.: <i>RH 6.2, SOLARIS 2.6 etc.</i>
MinPhysicalMemory	Minimum available physical memory (expressed in Mbytes) among the hosts associated to the Computing Element. If the CE is a "single host", this value represents its actual physical memory.
MinLocalDiskSpace	This attribute represents the minimum local disk footprint (that is the "working directory" where the job computation will take place) available to a running job running on a worker node (expressed in Mbytes). If more than one node is associated to the CE, we assume that all these worker nodes make available the same local disk space. It is also assumed that this advertised local disk footprint is actually available to a running job, even in case that more than one process is running on a given "worker" node.
TotalCPUs	the number of total CPUs associated to the resource.
FreeCPUs	the total number of free processors associated to the resource, processors able to run, in that moment, jobs submitted to the resource.
NumSMPs	number of SMP processors associated to the resource.
MinSPUProcessors	This is the minimum number of SPU processors (for SMP hosts).



Note

Doc. Identifier:
DataGrid-01-NOT-0101_07

Date: 03/12/2002

CE Attribute	Meaning
MaxSPUProcessors	This is the maximum number of SPU processors (for SMP hosts).
TotalJobs	the number of jobs submitted to the resource, jobs that have not already been completed.
RunningJobs	The number of jobs submitted to the resource that are currently running.
IdleJobs	the number of jobs submitted to the resource, jobs that are not running since they are waiting for available resources.
MaxTotalJobs	the maximum number of jobs (running and idle) allowed for the resource.
MaxRunningJobs	the maximum number of running jobs allowed for the resource.
WorstTraversalTime	Worst traversal time (in seconds) for jobs submitted to the Computing Element.
EstimatedTraversalTime	Scaled value of the last traversal time (in seconds), i.e. $(\text{Last job traversal time}) * (\text{queue length}) / (\text{queue length when that job arrived})$
Active	This is a boolean attribute indicating if the Computing Element is active. For example if the CE is a queue it indicates if it is ready or not to dispatch jobs to the executing machines.
RunWindow	the time windows that define when the resource is active, (for a queue: ready to dispatch jobs to the executing machines). This attribute may appear zero or more times for a Computing Element entity.
Priority	the priority of the resource.
MaxCPUTime	the maximum CPU time (in seconds) allowed for jobs submitted to the resource.
MaxWallClockTime	the maximum wall clock time (in seconds) allowed for jobs submitted to the resource.
MinSIOO	It is the minimum value of the SpecInt2000

CE Attribute	Meaning
	benchmark among the processors associated to this CE. If the CE is a "single processor", this value represents its actual performance.
MaxSI00	It is the maximum value of the SpecInt2000 benchmark among the processors associated to this CE. If the CE is a "single processor", this value represents its actual performance.
AverageSI00	It is the average of the SpecInt2000 benchmark of the nodes associated to this CE. If the CE is a "single processor", this value represents its actual performance.
AuthorizedUser	This is the subject of a X509 user certificate, representing a user authorized to submit job to the CE. This attribute may appear zero or more times for a ComputingElement entity.
RunTimeEnvironment	It is a tag defining a run time environment/package/software installed on the Computing Element. In case, the version of this package/environment is included in this string. This attribute may appear zero or more times for a ComputingElement entity.
AFSAvailable	Boolean attribute defining if AFS is installed on the Computing Element.
OutboundIP	Boolean. It indicates if outbound connectivity is allowed (e.g. all the worker nodes associated to the CE can "initiate" a data transfer, sending and/or receiving data to/from a remote Internet node).
InboundIP	Boolean. It indicates if inbound connectivity is allowed (e.g. a remote Internet node can "initiate" a data transfer, sending and/or receive data to/from any worker node associated to the CE).

Table 3 Computing Element attributes

Close SE Attribute	Meaning
CloseSE	This is the string that univocally identifies the Storage Element close enough to the computing element. This corresponds to the <i>SEId</i> attribute of the SE.
CEId	This is the string that univocally identifies the Computing Element close enough to the storage element.
MountPoint	The mount point of this SE from the considered CE (defined only if "local access" is supported). E.g. MountPoint = "/disk1";

Table 4 Close Storage Element attributes

SE Attribute	Meaning
SEId	This is a string that univocally identifies the Storage Element (it is the hostname for PM9).
CloseCE	This is the string that univocally identifies the Computing Element close enough to this Storage Element (this corresponds to the <i>CEId</i> attribute of the <i>CE</i>). This attribute may appear zero or more times for a StorageElement entity.

Table 5 Storage Element attributes

SE Protocol Attribute	Meaning
SEId	This is a string that univocally identifies the Storage Element (it is the hostname for PM9).
SEProtocol	This attribute defines the access protocol for the storage element (e.g. GridFtp, RFIO, etc...).
Port	the port number associated to the considered protocol.

Table 6 Storage Element Protocol attributes