

NorduGrid's Extended Globus RSL (xRSL) vs JSDL, a comparison

Oxana Smirnova, Marko Niinimäki
April 2005; revised and edited in April-May 2005

Remarks

- Based on
 - JSDL specs draft 0.9.5-03 (May 02, 2005)
(<https://forge.gridforum.org/projects/jsdl-wg/document/draft-ggf-jsdl-spec/en/18>)
 - xRSL reference (March 21, 2005)
(<http://www.nordugrid.org/documents/xrsl.pdf>)
 - Discussions in EGEE 3rd conference, Athens April 2005.
- JSDL elements are outlined according to the nesting suggested in the specification; xRSL elements are not outlined.
- JSDL elements of “complex” type (those having children) are not applicable in xRSL, are marked as “N/A”.
- JSDL substitutes relational operators with range values (realized as XML tags); xRSL uses operators explicitly.
- Allowed multiplicity of elements:
 - if not specified, a single instance is REQUIRED,
 - {0..n} – none is required, many allowed,
 - {0..1} – not more than one allowed,
 - {1..n} – at least one is REQUIRED.
- Clarifications and examples of the use of JSDL Application tag would be most welcome. It should be noticed that the multiplicity of the Application tag is {0..1} and it should encode a lot of information like “chemistry package X in parallel execution environment Y”.
- JSDL “POSIXApplication” is considered an extension in the JSDL specs; it is the best match for an xRSL application.
- The role of the xRSL runTimeEnvironment (RTE in what follows) and its correspondence to JSDL elements have been extensively discussed. Contrary to a widespread opinion, RTE is not another kind of an application: it is an attribute of an application, thus being situated on another level. RTE requirement to a POSIX application is the same as a memory limit requirement: an application must not be scheduled to a resource that does not meet these requirements.
- JSDL “DataStaging” section is still underdeveloped and not very satisfactory defined; difficult to compare/evaluate.
- It is very important to be able to specify alternative requirements like “operating system LINUX OR operating system Solaris required for the job”. From the present JSDL draft it is not clear whether it is possible.
- “JobRefinement” and “Profile” elements could perhaps be used to express alternative requirements as above, however, they are absent from the latest JSDL draft.
- A very important feature of xRSL: possibility to concatenate many job definitions in one request. The present JSDL draft does not describe such possibility.
- Conclusion: core elements (executable, arguments, stdin/stdout/stderr) are present in both structures, can be easily converted. If JSDL is capable of defining several job requests and can be extended by including the missing elements (which seems to be trivial), NorduGrid/ARC will benefit from switching from xRSL to JSDL.

JSDL	xRSL
1 Numerical operators	
Not supported, can be used by XML	Not supported
2 Structure	
2.1 Document layout	
<p>Structured layout with mandatory and optional blocks:</p> <pre><JobDefinition> <JobDescription> <JobIdentification ... />? <Application ... /> <Resource ... />* <DataStaging ... />* </JobDescription> </JobDefinition></pre>	<p>No structure, flat unordered layout. If an element appears several times, the later values overwrite the earlier ones.</p>
2.2 Element types	
<p>Data types:</p> <ul style="list-style-type: none"> • xsd:string • xsd:normalizedString, • xsd:positiveInteger • xsd:boolean • xsd:NCName • xsd:anyType • xsd:token • xsd:any##other • xsd:complex • jsdl:rangeValue 	<p>Data type not explicitly enforced, de-facto used:</p> <ul style="list-style-type: none"> • string • URL • positiveInteger • time
<p>Processor architectures: MUST be supported</p> <ul style="list-style-type: none"> • sparc • powerpc • x86 • parisc • mips • ia64 • arm • other 	N/A
<p>Filesystem types: MUST be supported</p> <ul style="list-style-type: none"> • swap (memory paging) • tmp (not available after job end) • spool (may persist) • normal (deleted only by user) 	N/A
<p>Operating System Types: may be supported, see CIM_OperatingSystem model</p> <ul style="list-style-type: none"> • large enumeration (currently 53 items) 	N/A
<p>File Creation Flags:</p> <ul style="list-style-type: none"> • Overwrite • DontOverwrite • Append • Prepend 	N/A

Range Value: a complex type, allows to define open- and/or close-ended ranges, as well as exact values; MUST be supported	N/A
--	-----

3 Element set

3.1 Job structure

• JobDefinition	& (ampersand) {1..n}
- JobDescription	N/A
- Description {0..1}	-
- -	rsl_substitution {0..n}

3.2 Job identity

• JobIdentification {0..1}	N/A
- JobName {0..1}	jobName {0..n}
- JobAnnotation {0..1}	-
- JobProject {0..n}	-

3.3 Application requirements

• Application {0..1}	N/A
- ApplicationName {0..1}	
- ApplicationVersion {0..1}	
- Description {0..1}	
• POSIXApplication {0..1} - extension	N/A
- Executable {0..1}	executable {1..n}
- Argument {0..n}	arguments {0..n}
- Input {0..1}	stdin {0..n}
- Output {0..1}	stdout {0..n}
- Error {0..1}	stderr {0..n}
- -	join {0..n}
- -	gmLog {0..n}
- WorkingDirectory {0..1}	-
- Environment {0..n}	environment {0..n}
- TileSize {0..1}	-
- ProcessCount {0..1}	-
- WallTimeLimit {0..1}	-
- FileSizeLimit {0..1}	-
- CoreDumpLimit {0..1}	-
- DataSegmentLimit {0..1}	-
- LockedMemoryLimit {0..1}	-
- MemoryLimit {0..1}	memory {0..n}
- OpenDescriptorsLimit {0..1}	-
- PipeSizeLimit {0..1}	-
- StackSizeLimit {0..1}	-
- CPUTimeLimit {0..1}	cpuTime {0..n}
- ProcessCountLimit {0..1}	-
- VirtualMemoryLimit {0..1}	-
- ThreadCountLimit {0..1}	-
- UserName {0..1}	-
- GroupName {0..1}	-
- -	runTimeEnvironment {0..n}
- -	executables {0..n}
- -	startTime {0..n}
- -	lifeTime {0..n}

- -	notify {0..n}
- -	rerun {0..n}
- -	dryRun {0..n}
- -	jobReport {0..n}

3.4 Resource requirements

• Resource {0..1}	N/A
- CandidateHosts {0..1}	N/A
▪ HostName {1..n}	cluster {0..n}
- -	queue {0..n}
- CPUArchitecture {0..1}	architecture {0..n}
- CPUSpeed {0..1}	-
- CPUTime {0..1}	cpuTime {0..n}
- -	gridTime {0..n}
- CPUCount {0..1}	count {0..n}
- ResourceCount {0..1}	-
- PhysicalMemory {0..1}	memory {0..n}
- VirtualMemory {0..1}	-
- NetworkBandwidth {0..1}	-
- -	nodeAccess {0:n}
- DiskSpace {0..1}	disk {0..n}
- FileSystem {0..n}	N/A
▪ MountPoint {0..1}	-
▪ MountSource {0..1}	-
▪ DiskSpace {0..1}	disk {0..n}
▪ FileSystemType {0..1}	-
▪ Description {0..1}	-
- ExclusiveExecution {0..1}	-
- OperatingSystem {0..1}	N/A
▪ OperatingSystemType {0..1}	-
▪ OperatingSystemVersion {0..1}	-
▪ Description {0..1}	-
- -	middleware {0..n}

3.5 Data staging

• DataStaging {0..n}	inputFiles {0..n}
	outputFiles {0..n}
- FileName	-
- FileSystemName {0..1}	-
- CreationFlag	-
- DeleteOnTermination {0..1}	-
- -	ftpThreads {0..n}
- -	replicaCollection {0..n}
- Source {0..1}	inputFiles/outputFiles source
▪ URI	inputFiles/outputFiles source
▪ -	cache {0..n}
- Target {0..1}	inputFiles/outputFiles target
▪ URL	inputFiles/outputFiles target