

<sup>1</sup> GWD-R  
DRMAA-WG  
drmaa-wg@ogf.org

Peter Tröger, Hasso-Plattner-Institute (editor)  
Daniel Templeton, Cloudera (editor)  
March 2011

## <sup>2</sup> **Distributed Resource Management Application API Version 2** <sup>3</sup> **(DRMAA) - Draft 1**

### <sup>4</sup> **Status of This Document**

<sup>5</sup> Group Working Draft Recommendation (GWD-R)

<sup>6</sup> (See footnote)<sup>1</sup>

### <sup>7</sup> **Obsoletes**

<sup>8</sup> This document obsoletes GFD-R.022 [7], GFD-R-P.130 [10], and GWD-R.133 [8].

### <sup>9</sup> **Copyright Notice**

<sup>10</sup> Copyright © Open Grid Forum (2005-2011). Some Rights Reserved. Distribution is unlimited.

### <sup>11</sup> **Trademark**

<sup>12</sup> All company, product or service names referenced in this document are used for identification purposes only  
<sup>13</sup> and may be trademarks of their respective owners.

### <sup>14</sup> **Abstract**

<sup>15</sup> This document describes the *Distributed Resource Management Application API Version 2 (DRMAA)*, which  
<sup>16</sup> provides a generalized API to *Distributed Resource Management (DRM)* systems in order to facilitate the  
<sup>17</sup> development of portable application programs and high-level libraries for such systems. DRMAA defines  
<sup>18</sup> interfaces for a tightly coupled, but still portable access by abstracting the fundamental functions available  
<sup>19</sup> in the majority of DRM systems. The scope is limited to job submission, job control, and retrieval of job  
<sup>20</sup> and machine monitoring information.

<sup>21</sup> This document acts as root specification for the abstract API concepts and the behavioral rules that must be  
<sup>22</sup> fulfilled by a DRMAA-compliant implementation. The programming language representation of the abstract  
<sup>23</sup> API concepts must be formulated by a separate *language binding specification* derived from this document.

<sup>24</sup> The intended audience for this specification are DRMAA language binding designers, DRM system vendors,  
<sup>25</sup> high-level API designers and meta-scheduler architects. End users are expected to rely on product-specific  
<sup>26</sup> documentation for the DRMAA API implementation in their particular programming language.

---

<sup>1</sup> This is the non-normative annotated version of the specification with line numbers. It includes historical information concerning the content and why features were included or discarded by the working group. It also emphasizes the consequences of some aspects that may not be immediately apparent. This document is only intended for internal working group discussions.

## 27    Contents

28	1	Introduction . . . . .	3
29	1.1	Notational Conventions . . . . .	3
30	1.2	Language Bindings . . . . .	4
31	1.3	Slots and Queues . . . . .	4
32	1.4	Multithreading . . . . .	5
33	2	Namespace . . . . .	5
34	3	Common Type Definitions . . . . .	5
35	4	Common Data Structures and Enumerations . . . . .	6
36	4.1	OperatingSystem enumeration . . . . .	7
37	4.2	CpuArchitecture enumeration . . . . .	8
38	4.3	ResourceLimitType enumeration . . . . .	9
39	4.4	JobTemplatePlaceholder enumeration . . . . .	10
40	4.5	Queue structure . . . . .	11
41	4.6	Version structure . . . . .	11
42	4.7	Machine structure . . . . .	11
43	4.8	JobInfo structure . . . . .	13
44	5	Common Exceptions . . . . .	16
45	6	The DRMAA Session Concept . . . . .	18
46	6.1	SessionManager Interface . . . . .	18
47	7	Working with Jobs . . . . .	21
48	7.1	The DRMAA State Model . . . . .	21
49	7.2	JobSession Interface . . . . .	24
50	7.3	DrmaaCallback Interface . . . . .	27
51	7.4	JobTemplate Structure . . . . .	27
52	7.5	Job Interface . . . . .	36
53	7.6	JobArray Interface . . . . .	38
54	8	Working with Advance Reservation . . . . .	39
55	8.1	ReservationSession Interface . . . . .	39
56	8.2	ReservationTemplate structure . . . . .	40
57	8.3	Reservation Interface . . . . .	43
58	9	Monitoring the DRM System . . . . .	44
59	9.1	MonitoringSession Interface . . . . .	45
60	10	Annex A: Complete DRMAA IDL Specification . . . . .	47
61	11	Security Considerations . . . . .	52
62	12	Contributors . . . . .	52
63	13	Intellectual Property Statement . . . . .	53
64	14	Disclaimer . . . . .	53
65	15	Full Copyright Notice . . . . .	54
66	16	References . . . . .	54

## 67 1 Introduction

68 This document describes the *Distributed Resource Management Application API Version 2 (DRMAA)* in-  
69 terface semantics in a generalized way by using the *OMG Interface Definition Language (IDL)* [4] syntax for  
70 a language-agnostic description. Based on this abstract specification, *language binding* standards have to  
71 be designed that map the described concepts into a library interface for a particular programming language  
72 (e.g. C, Java, Python). While this document has the responsibility to ensure consistent API semantics over  
73 all possible DRMAA implementations, the language binding has the responsibility to ensure source-code  
74 portability for DRMAA applications on different DRM systems.

75 An effort has been made to choose an API layout that is not unique to a specific language. However, in some  
76 cases, various languages disagree over some points. In those cases, the most meritorious approach was taken,  
77 irrespective of language.

78 There are other relevant OGF standards in the area of job submission and monitoring. An in-depth com-  
79 parison and positioning of the obsoleted DRMAA1 specification was provided by another publication [11].

80 The DRMAA specification is based on the following stake holders:

- 81 • *Distributed resource management system / DRM system / DRMS*: Any system which supports the con-  
82 cept of distributing computational jobs on execution resources through the help of a central scheduling  
83 entity. Examples are multi-processor systems controlled by a operating system scheduler, cluster sys-  
84 tems with multiple machines controlled by a central scheduler software, grid systems, or cloud systems  
85 with a job concept.
- 86 • *DRMAA implementation resp. DRMAA library*: The implementation of a DRMAA language binding  
87 specification with the functional semantics described in this document. The resulting artifact is ex-  
88 pected to be a library that is deployed together with the DRM system that is wrapped by the particular  
89 implementation.
- 90 • *(DRMAA-based) application*: Software that utilizes the DRMAA implementation for gaining access to  
91 one or multiple DRM systems in a standardized way.
- 92 • *Submission host*: A execution resource in the DRM system that runs the DRMAA-based application.
- 93 • *Execution host*: A execution resource in the DRM system that can run a job submitted through the  
94 DRMAA implementation.

### 95 1.1 Notational Conventions

96 In this document, IDL resp. programming language elements and definitions are represented in a **fixed-width**  
97 font.

98 The key words “MUST” “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD  
99 NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” are to be interpreted as described in RFC 2119 [1].

Parts of the specification which are normative for derived language binding specifications only are graphically  
marked as shaded box.

## 100 1.2 Language Bindings

A language binding specification derived from this document MUST define a mapping between the IDL constructs and specific programming language constructs, with focus on source code portability for the resulting DRMAA-based applications.

A language binding SHOULD NOT rely itself completely on the OMG language mapping standards available for many programming languages, since they have a huge overhead of irrelevant CORBA-related mapping rules. Therefore, language binding authors must carefully decide if a binding decision reflects a natural and simple mapping of the intended purpose for the DRMAA interfaces. The binding SHOULD reuse OMG value type mappings (e.g. IDL `long long` to Java `long`), and SHOULD define custom mappings for the other types. The language binding MUST use the described concept mapping in a consistent manner for its overall API layout.

Due to the usage of IDL, all method groups for a particular purpose (e.g. job control) are described in terms of interfaces, and not classes. The mapping to a class concept depends on the specific language-mapping rules.

It may be the case that IDL constructs do not map directly to any language construct. In this case it MUST be ensured that the chosen mapping retains the intended semantic of the DRMAA interface definition.

Access to scalar attributes (`string`, `boolean`, `long`) MUST operate in a pass-by-value mode. A language binding must ensure that this behavior is always fulfilled. For non-scalar attributes, the language binding MUST specify a consistent access strategy for all these attributes – either pass-by-value or pass-by-reference – according to the use cases of language binding implementations.

This specification tries to consider the possibility of a Remote Procedure Call (RPC) scenario in a DRMAA-conformant language mapping. It SHOULD therefore be ensured that the programming language type for an IDL `struct` definition supports the serialization and comparison of instances. These capabilities should be accomplished through whatever mechanism is most natural for the specific programming language.

A language binding MUST define a way to declare an invalid struct member value resp. scalar value (`UNSET`). In case, a definition per data type needs to be provided. The `UNSET` value for a boolean data type MUST translate to `False`.

Unclear if  
UNSET for  
numeric  
values could be  
zero.

101

102 (See footnote)<sup>2</sup>

## 103 1.3 Slots and Queues

104 DRMAA supports the notion of slots and queues as resources of a DRM system. A DRMAA application  
 105 can request them in advance reservation and job submission. However, slots and queues SHALL be opaque  
 106 concepts from the viewpoint of a DRMAA implementation, meaning that the requirements given by the  
 107 application are just passed through to the DRM system. This is reasoned by the large variation in interpreting  
 108 that concepts in the different DRM systems, which makes it impossible to define a common understanding  
 109 on the level of the DRMAA API.

<sup>2</sup> The concept of a `UNSET` value was decided on a conf call (Aug 25th 2010). Boolean in C should use custom enumeration (TRUE, FALSE, INVALID) or pointer to static values. A numerical `UNSET` in C should use a magic number, since all long attributes are unsigned, it could be `MIN_INT`. With Python, just use `None`. For Java, Dan has an idea.

110 (See footnote)<sup>3</sup>111 

## 1.4 Multithreading

112 High-level APIs such as SAGA [3] are expected to utilize DRMAA for asynchronous operations, based on the  
 113 assumption that re-entrancy is supported by DRMAA implementations. For this reason, implementations  
 114 SHOULD ensure the proper functioning of the library in case of re-entrant library calls. A DRMAA library  
 115 SHOULD allow a multithreaded application to use DRMAA interfaces without any explicit synchronization  
 116 among the application threads. DRMAA implementers should document their work as thread safe if they  
 117 meet the above criteria. Providers of non-thread-safe DRMAA implementations should document all the  
 118 interfaces that are thread unsafe and provide a list of interfaces and their dependencies on external thread  
 119 unsafe routines.

120 

## 2 Namespace

121 The DRMAA interfaces and structures are encapsulated by a naming scope, which avoids conflicts with  
 122 other API's used in the same application.

123 `module DRMAA2 {`

Language binding authors MUST map the IDL module encapsulation to an according package or namespace  
 concept and MAY change the module name according to programming language conventions.

124 (See footnote)<sup>4</sup>125 

## 3 Common Type Definitions

126 The DRMAA specification defines some custom types to express special value semantics not expressible in  
 127 IDL.

```
128     typedef sequence<string> OrderedStringList;
129     typedef sequence<string> StringList;
130     typedef sequence<Job> JobList;
131     typedef sequence<Queue> QueueList;
132     typedef sequence<Machine> MachineList;
133     typedef sequence<Reservation> ReservationList;
134     typedef sequence< sequence<string,2> > Dictionary;
135     typedef string AbsoluteTime;
136     typedef long long TimeAmount;
137     native ZERO_TIME;
138     native INFINITE_TIME;
139     native UNSET;
```

---

<sup>3</sup> As one example, queues can be either treated as representation of execution hosts (Sun Grid Engine) or as central waiting line located at the scheduler (LSF).

<sup>4</sup> Comparison to DRMAA v1.0: The IDL module name was changed to DRMAA2, in order to intentionally break backward compatibility of the interface.

<sup>140</sup> **OrderedStringList:** An unbounded list of strings, which supports element insertion, element deletion, and iteration over elements while keeping an element order.

<sup>141</sup> **StringList:** An unbounded list of strings, without any demand on element order.

<sup>142</sup> **JobList:** An unbounded list of `Job` instances, without any demand on element order.

<sup>143</sup> **MachineList:** An unbounded list of `Machine` instances, without any demand on element order.

<sup>144</sup> **QueueList:** An unbounded list of `Queue` instances, without any demand on element order.

<sup>145</sup> **ReservationList:** An unbounded list of `Reservation` instances, without any demand on element order.

<sup>146</sup> **Dictionary:** An unbounded dictionary type for storing key-value pairs, without any demand on element order.

<sup>147</sup> **AbsoluteTime:** Expression of a point in time, at least with a resolution to seconds.

<sup>148</sup> **TimeAmount:** Expression of an amount of time, at least with a resolution to seconds.

<sup>149</sup> **ZERO\_TIME:** A constant value of type `TimeAmount` which expresses a zero amount of time.

<sup>150</sup> **INFINITE\_TIME:** A constant value of type `TimeAmount` which expresses an infinite amount of time.

A language binding MUST replace these type definitions with semantically equal reference or value types in the according language. This may include the creation of new complex language types for one or more of the above concepts. The language binding MUST define a consistent mapping on module level, and a mechanism for obtaining the RFC822 string representation from a given `AbsoluteTime` resp. `TimeAmount` instance.

<sup>153</sup> (See footnote)<sup>5</sup>

## <sup>154</sup> 4 Common Data Structures and Enumerations

<sup>155</sup> DRMAA defines a set of data structures commonly used by different interfaces to express information for resp. from the DRM system. A DRMAA implementation is allowed to extend the specified structures, if explicitly noted in the description of the particular structure (e.g. as with `JobInfo`). Behavioral aspects of such extended attributes are out of scope for DRMAA. Implementations SHALL only extend data structures in the way specified by the language binding.

A language binding MUST define a consistent mechanism to realize implementation-specific structure and enumeration extension, without breaking the portability of DRMAA-based applications that rely on the original version of the structure. Object oriented languages MAY use inheritance mechanisms for this purpose.

Language bindings SHOULD define numerical values for all constants and enumeration members, in order to foster binary portability of DRMAA-based applications. Instances of these structures SHALL be treated in a "call-by-value" fashion, meaning that the collection of struct member values is handed over as one to the called interface method.

<sup>5</sup> The PartialTimestamp functionality from DRMAA 1.0 was completely removed. Absolute date and time values are now expressed as RFC822 conformant data items with stringification support (conf. call Mar 31st 2009). String list for job identifiers are replaced by Job object lists (F2F meeting July 2009)

160 (See footnote)<sup>6</sup>

161 **4.1 OperatingSystem enumeration**

162 DRMAA supports the identification of an operating system installation on execution resources in the DRM  
 163 system. The **OperatingSystem** enumeration is used as data type both in the advanced reservation and the  
 164 DRM system monitoring functionalities. It defines a set of standardized identifiers for operating system  
 165 types. The list is a shortened version of the according CIM Schema [6]. It includes only operating systems  
 166 that are supported by the majority of DRM systems available at the time of writing:

```
167 enum OperatingSystem {
168     HPUX, LINUX, IRIX, TRUE64, MACOS, SUNOS, WIN, WINNT, AIX, UNIXWARE,
169     BSD, OTHER_OS};
```

170 **AIX:** AIX Unix by IBM.

171 **BSD:** All operating system distributions based on the BSD kernel.

172 **LINUX:** All operating system distributions based on the Linux kernel.

173 **HPUX:** HP-UX Unix by Hewlett-Packard.

174 **IRIX:** The IRIX operating system by SGI.

175 **MACOS:** The MAC OS X operating system by Apple.

176 **SUNOS:** SunOS resp. Solaris operating system by Sun / Oracle.

177 **TRUE64:** True64 Unix by Hewlett-Packard, or DEC Digital Unix, or DEC OSF/1 AXP.

178 **UNIXWARE:** UnixWare system by SCO group.

179 **WIN:** Windows 95, Windows 98, Windows ME.

180 **WINNT:** Microsoft Windows operating systems based on the NT kernel

181 **OTHER\_OS:** An operating system type not specified in this list.

182 Implementations SHOULD NOT add new operating system identifiers to this enumeration, even if they are  
 183 supported by the underlying DRM system.

184 The operating system information is only useful in conjunction with version information (see Section 9.1),  
 185 which is also the reporting approach taken in most DRM systems. Examples:

- 186 • The Apple MacOS X operating system commonly denoted as "Snow Leopard" would be reported as  
 187 "MACOS" with the version structure ["10", "6"]
- 188 • The Microsoft Windows 7 operating system would be reported as "WINNT" with the version infor-  
 189 mation ["6", "1"], which is the internal version number reported by the Windows API.
- 190 • All Linux distributions would be reported as operating system type "LINUX" with the major revision  
 191 of the kernel, e.g. ["2", "6"].

---

<sup>6</sup> Comparison to DRMAA 1.0: The binding of job template attribute names and exception names to strings was removed from the main specification. Language bindings such as for the C programming languages have to define their own mapping. It is recommended to keep string identifiers from DRMAA 1.0 as far as possible.

192 • The Solaris operating system is reported as "SUNOS", together with the internal version number, e.g.  
 193 [ "5", "10" ] for Solaris 10.

194 The DRMAA `OperatingSystem` enumeration can be mapped to other high-level APIs. Table 1 gives a  
 195 non-normative set of examples.

DRMAA <code>OperatingSystem</code> value	JSDL <code>jsdl:OperatingSystemTypeEnumeration</code> value
HPUX	HPUX
LINUX	LINUX
IRIX	IRIX
TRUE64	Tru64_UNIX, OSF
MACOS	MACOS
SUNOS	SunOS, SOLARIS
WIN	WIN95, WIN98, Windows_R_Me
WINNT	WINNT, Windows_2000, Windows_XP
AIX	AIX
UNIXWARE	SCO_UnixWare, SCO_OpenServer
BSD	BSDUNIX, FreeBSD, NetBSD, OpenBSD
OTHER_OS	Other

Table 1: Mapping example for the DRMAA `OperatingSystem` enumeration

## 196 4.2 CpuArchitecture enumeration

197 DRMAA supports identifying the processor instruction set architecture on execution resources in the DRM  
 198 system. The `CpuArchitecture` enumeration is used as data type both in the advanced reservation and the  
 199 DRM system monitoring functionalities. It defines a set of standardized identifiers for processor architecture  
 200 families. The list is a shortened version of the according CIM Schema [6], It includes only processor families  
 201 that are supported by the majority of DRM systems available at the time of writing:

```
202 enum CpuArchitecture {
 203   ALPHA, ARM, CELL, PARISC, X86, X64, IA64, MIPS, PPC, PPC64,
 204   SPARC, SPARC64, OTHER_CPU};
```

205 **ALPHA:** The DEC Alpha / Alpha AXP processor architecture.

206 **ARM:** The ARM processor architecture.

207 **CELL:** The Cell processor architecture.

208 **PA-RISC:** The PA-RISC processor architecture.

209 **X86:** The IA-32 line of the X86 processor architecture family, with 32bit support only.

210 **X64:** The X86-64 line of the X86 processor architecture family, with 64bit support.

211 **IA-64:** The Itanium processor architecture.

212 **MIPS:** The MIPS processor architecture.

213 **PPC:** The PowerPC processor architecture, all models with 32bit support only.

214 **PPC64:** The PowerPC processor architecture, all models with 64bit support.

215 **SPARC:** The SPARC processor architecture, all models with 32bit support only.

216 **SPARC64:** The SPARC processor architecture, all models with 64bit support.

217 **OTHER\_CPU:** A processor architecture not specified in this list.

218 The DRMAA `CpuArchitecture` enumeration can be mapped to other high-level APIs. Table 2 gives a  
219 non-normative set of examples.

220 The reporting and job configuration for processor architectures SHOULD operate on a "as-is" base, if sup-  
221 ported by the DRM system. This means that the reported architecture should reflect the current operation  
222 mode of the processor with the running operating system. For example, X64 processors executing a 32-bit  
223 operating system typically report themselves as X86 processor.

DRMAA <code>CpuArchitecture</code> value	JSDL <code>jsdl:ProcessorArchitectureEnumeration</code> value
ALPHA	other
ARM	arm
CELL	other
PA-RISC	parisc
X86	x86_32
X64	x86_64
IA-64	ia64
MIPS	mips
PPC	powerpc
PPC64	powerpc
SPARC	sparc
SPARC64	sparc
OTHER	other

Table 2: Mapping example for DRMAA `CpuArchitecture` enumeration

### 224 4.3 `ResourceLimitType` enumeration

225 Modern DRM systems expose resource constraint capabilities from the operating system for jobs on the  
226 execution host. The `ResourceLimitType` enumeration represents the typical *ulimit(3)* parameters [5] in  
227 different DRM systems. All parameters relate to the operating system process representing some job on the  
228 execution host.

```
229 enum ResourceLimitType {
230     CORE_FILE_SIZE, CPU_TIME, DATA_SEG_SIZE, FILE_SIZE, OPEN_FILES,
231     STACK_SIZE, VIRTUAL_MEMORY, WALLCLOCK_TIME };
```

232 **CORE\_FILE\_SIZE:** The maximum size of the core dump file created on fatal errors of the process, in  
233 Kibibyte. Setting this value to zero SHOULD disable the creation of core dump files on the execution  
234 host.

235 **CPU\_TIME:** The maximum accumulated time in seconds the process is allowed to perform computations  
236 on all processors in the execution host.

237 **DATA\_SEG\_SIZE:** The maximum amount of memory the process can allocate on the heap e.g. for object  
238 creation, in Kibibyte.

239 **FILE\_SIZE:** The maximum file size the process can generate, in Kibibyte.

240 **OPEN\_FILES:** The maximum number of file descriptors the process is allowed to have open at the same  
241 time.

242 **STACK\_SIZE:** The maximum amount of memory the process can allocate on the stack, e.g. for local  
243 variables, in Kibibyte.

244 **VIRTUAL\_MEMORY:** The maximum amount of memory the process is allowed to allocate, in Kibibyte.

245 **WALLCLOCK\_TIME:** The maximum wall clock resp. real time in seconds the job is allowed to exist in  
246 any of the "Started" or "Queued" states (see Section 7.1).

Explanations  
need approval  
by the group.  
Does WALL-  
CLOCK\_TIME  
really also  
include  
queued time  
?

247  
248 (See footnote)<sup>7</sup>

#### 249 4.4 JobTemplatePlaceholder enumeration

250 The `JobTemplatePlaceholder` enumeration defines constant macros to be used in string attributes of a  
251 `JobTemplate` instance.

252 

```
enum JobTemplatePlaceholder {
    HOME_DIRECTORY, WORKING_DIRECTORY, HOST_NAME, USER_NAME, PARAMETRIC_INDEX };
```

254 A `HOME_DIRECTORY` placeholder SHOULD be only allowed at the beginning of a `JobTemplate` attribute value.  
255 It denotes the remaining portion as a directory / file path resolved relative to the job users home directory  
256 at the execution host.

257 A `WORKING_DIRECTORY` placeholder SHOULD be only allowed at the beginning of a `JobTemplate` attribute  
258 value. It denotes the remaining portion as a directory / file path resolved relative to the jobs working  
259 directory at the execution host.

260 The `HOST_NAME` placeholder SHOULD be usable at any position within an attribute value that supports place  
261 holders. It SHALL be substituted by the full-qualified name of the execution host were the job is executed.

262 The `USER_NAME` placeholder SHOULD be usable at any position within an attribute value that supports place  
263 holders. It SHALL be substituted by the job users account name on the execution host.

264 The `PARAMETRIC_INDEX` placeholder SHOULD be usable at any position within an attribute value that  
265 supports place holders. It SHALL be substituted by the parametric job index in a `JobSession::runBulkJobs`  
266 call (see Section 7.2.6). If the job template is used for a `JobSession::runJob` call, `PARAMETRIC_INDEX` should  
267 be substituted with a constant implementation-specific value.

268 (See footnote)<sup>8</sup>

<sup>7</sup> "Pipe size" was not added, since there is no use case in DRM systems with a job concept. "Max user processes" was omitted because it operates on the notion of users, which is not an explicit concept in DRMAA.

<sup>8</sup> Placeholders for other job template attributes were rejected, in order to avoid circular dependencies (Conf. call Oct 20th 2010)

269 **4.5 Queue structure**

270 The **Queue** structure denotes a job waiting queue in the DRM system. Queue is an opaque concept from the  
 271 perspective of the DRMAA application (see Section 1.3). The **Queue** struct contains read-only information.  
 272 Implementations MAY extend this structure with implementation-specific attributes.

```
273   struct Queue {
274     string name;
275     TimeAmount maxWallclockTime;
276   };
```

277 **4.5.1 name**

278 This attribute contains the name of the queue as reported by the DRM system. The format of the queue  
 279 name is implementation-specific. The naming scheme SHOULD be consistent for all strings returned.

280 **4.5.2 maxWallclockTime**

281 This attribute contains the maximum amount of wallclock time allowed for jobs submitted to the queue.  
 282 The attribute value is UNSET when there is no restriction. If this value is not UNSET, then any job submitted  
 283 to this queue SHOULD enter one of the "Terminated" states when the wallclock time limit is reached.

Termination condition  
must be approved by the  
group

285 **4.6 Version structure**

286 The **Version** structure denotes versioning information for an operating system, DRM system, or DRMAA  
 287 implementation.

```
288   struct Version {
289     string major;
290     string minor;
291   };
```

292 Both the **major** and the **minor** part are expressed as strings, in order to allow specific extensions with  
 293 character combinations such as "rev". Original version strings containing a dot, e.g. Linux "2.6", SHOULD  
 294 be interpreted as having the major part before the dot, and the minor part after the dot. The dot character  
 295 SHOULD NOT be added to the **Version** attributes.

296 **4.7 Machine structure**

297 The **Machine** structure describes the properties of a particular execution host in the DRM system. Im-  
 298 plementations MAY extend this structure with implementation-specific additional information. It contains  
 299 read-only information. An implementation resp. its DRM system MAY restrict jobs in their resource util-  
 300 ization even below the limits described in the **Machine** structure. The limits given here MAY be imposed  
 301 by the hardware configuration, or MAY be imposed by DRM system policies.

```
302   struct Machine {
303     string name;
304     long sockets;
```

```
305     long coresPerSocket;
306     long threadsPerCore;
307     double load;
308     long physMemory;
309     long virtMemory;
310     OperatingSystem machineOS;
311     Version machineOSVersion;
312     CpuArchitecture machineArch;
313 }
```

#### 314 4.7.1 name

315 This attribute describes the name of the machine as reported by the DRM system. The format of the  
316 machine name is implementation-specific, but MAY be a DNS host name. The naming scheme SHOULD be  
317 consistent for all strings returned.

#### 318 4.7.2 sockets

319 This attribute describes the number of processor sockets resp. CPUs usable for jobs on the machine from  
320 operating system perspective. The attribute value MUST be greater than 0. In the case where the correct  
321 value is unknown to the implementation, the value MUST be set to 1.

#### 322 4.7.3 coresPerSocket

323 This attribute describes the number of cores per socket usable for jobs on the machine from operating system  
324 perspective. The attribute value MUST be greater than 0. In case where the correct value is unknown to  
325 the implementation, the value MUST be set to 1.

#### 326 4.7.4 threadsPerCore

327 This attribute describes the number of threads that can be executed in parallel by a job on one core in the  
328 machine. The attribute value MUST be greater than 0. In case where the correct value is unknown to the  
329 implementation, the value MUST be set to 1.

#### 330 4.7.5 load

331 This attribute describes the 1-minute average load on the given machine, similar to the Unix *uptime* com-  
332 mand. The value has only informative character, and should not be utilized by end user applications for job  
333 scheduling purposes. An implementation MAY provide delayed or averaged data here, if necessary due to  
334 implementation issues. The implementation strategy on non-Unix systems is undefined.

#### 335 4.7.6 physMemory

336 This attribute describes the amount of physical memory in Kibibyte available on the machine.

#### 337 4.7.7 virtMemory

338 This attribute describes the amount of virtual memory in Kibibyte available for a job executing on this  
339 machine. The virtual memory amount is defined as the sum of physical memory installed plus the configured

340 swap space for the operating system. The value is expected to be used as indicator whether or not an  
 341 application is able to get its memory allocation needs fulfilled on a particular machine. Implementations  
 342 SHOULD derive this value directly from operating system information, without further consideration of  
 343 additional memory allocation restrictions such as address space range or already running processes.

344 **4.7.8 machineOS**

345 This attribute describes the operating system installed on the described machine, with semantics as specified  
 346 in Section 4.1.

347 **4.7.9 machineOSVersion**

348 This attribute describes the operating system version of the machine, with semantics as specified in Section  
 349 4.1.

350 **4.7.10 machineArch**

351 This attribute describes the instruction set architecture of the machine, with semantics as specified in Section  
 352 4.2.

353 **4.8 JobInfo structure**

354 The `JobInfo` structure describes job information that is available for the DRMAA-based application.

```
355     struct JobInfo {
356         string jobId;
357         Dictionary resourceUsage;
358         long exitStatus;
359         string terminatingSignal;
360         string annotation;
361         JobState jobState;
362         any jobSubState;
363         OrderedStringList allocatedMachines;
364         string submissionMachine;
365         string jobOwner;
366         string queueName;
367         TimeAmount wallclockTime;
368         long cpuTime;
369         AbsoluteTime submissionTime;
370         AbsoluteTime dispatchTime;
371         AbsoluteTime finishTime;};
```

372 The structure is used in two occasions - first for the expression of information about a single job, and second  
 373 as filter expression when retrieving a list of jobs from the DRMAA implementation.

374 In both usage scenarios, the structure information has to be understood as snapshot of the live DRM system.

375 Multiple values being set in one structure instance should be interpreted as "occurring at the same time".

376 In real implementations, some granularity limits must be assumed - for example, the `wallclockTime` and  
 377 the `cpuTime` attributes might hold values that were measured with a very small delay one after each other.

378 In the use case of job information monitoring, it is assumed that the DRM system has three job information  
 379 states: running, buffered, purged. Only information for jobs that are still running or are still held in the  
 380 buffer of finished job information will be reported completely. In this case, the information SHOULD reflect  
 381 the current status of the job as as close as possible to the time of the call.

382 If jobs have been purged out to accounting, different attributes might not contain valid data. Implementations  
 383 MAY decide to return only partially filled `JobInfo` instances due to performance restrictions in the  
 384 communication with the DRM system.

385 For additional DRMS-specific information, the `JobInfo` structure MAY be extended by the DRMAA imple-  
 386 mentation (see Section 4).

387 (See footnote)<sup>9</sup>

#### 388 4.8.1 `jobId`

389 For monitoring: Returns the stringified job identifier assigned to the job by the DRM system.

390 For filtering: Returns the job with the chosen job identifier.

#### 391 4.8.2 `resourceUsage`

392 For monitoring: Returns resource consumption information for the given job. The dictionary keys are  
 393 implementation-specific.

394 For filtering: Returns the jobs that have the dictionary key-value pairs as subset of their own.

395 Standardize  
resource  
usage key  
names ??

#### 396 4.8.3 `exitStatus`

397 For monitoring: The process exit status of the job, as reported by the operating system. If the job is not in  
 398 one of the terminated states, the value should be `UNSET`.

399 For filtering: Return the jobs with the given `exitStatus` value. Jobs without exit status information should  
 400 be filtered out by asking for the appropriate states.

#### 401 4.8.4 `terminatingSignal`

402 For monitoring: This attribute specifies the UNIX signal that reasoned the ending of the job. Implementa-  
 403 tions should document the extent to which they can gather such information in the particular DRM system  
 404 (e.g. with Windows hosts).

405 For filtering: Returns the jobs with the given `terminatingSignal` value.

---

<sup>9</sup> In comparison to DRMAA 1.0, the `JobInfo` value type was heavily extended for providing more information (solves issue #2827). `JobInfo::hasCoreDump` is no longer supported, since the information is useless without according core file staging support, which is not implementable in a portable way. (conf. call Jun 9th 2010)

Some DRM systems (SGE / Condor at least) support the automated modification of job template attributes after submission, and therefore allow to fetch the true job template attributes at run-time from the job. The monitoring for such data was intentionally not included in DRMAA (mailing list July 2010).

## 406 4.8.5 annotation

407 For monitoring: Gives a human-readable annotation describing why the job is in its current state or sub-state.  
408 The support for this information is optional.

409 For filtering: This attribute is ignored for filtering.

## 410 4.8.6 jobState

411 For monitoring: This attribute specifies the jobs current state according to the DRMAA job state model  
412 (see Section 7.1).

413 For filtering: Returns all jobs in the specified state. If the given state is simulated by the implementation  
414 (see Section 7.1), the implementation SHOULD raise an `InvalidArgumentException` explaining that this  
415 filter can never match.

## 416 4.8.7 jobSubState

417 For monitoring: This attribute specifies the jobs current DRMAA implementation specific sub-state (see  
418 Section 7.1).

419 For filtering: Returns all jobs in the specified sub-state. If the given sub-state is not supported by the imple-  
420 mentation (see Section 7.1), the implementation SHOULD raise an `InvalidArgumentException` explaining  
421 that this filter can never match.

## 422 4.8.8 allocatedMachines

423 This attribute expresses the set of machines that are utilized for job execution. Implementations MAY  
424 decide to give the ordering of machine names a particular meaning, for example putting the master node in  
425 a parallel job at first position. This decision should be documented for the user. For performance reasons,  
426 only the machine names are returned, and SHOULD be equal to the according `Machine::name` attribute in  
427 monitoring data.

428 For monitoring: This attribute lists the set of names of the machines to which this job has been assigned.

429 For filtering: Returns the list of jobs which have a set of assigned machines that is a superset of the given  
430 set of machines.

## 431 4.8.9 submissionMachine

432 This attribute provides the machine name of the submission host for this job. For performance reasons,  
433 only the machine name is returned, and SHOULD be equal to the according `Machine::name` attribute in  
434 monitoring data.

435 For monitoring: This attribute specifies the machine from which this job was submitted.

436 For filtering: Returns the set of jobs that were submitted from the specified machine.

## 437 4.8.10 jobOwner

438 For monitoring: This attribute specifies the job owner as reported by the DRM system.

439 For filtering: Returns all jobs owned by the specified user.

## 440 4.8.11 queueName

441 For monitoring: This attribute specifies the queue in which the job was queued resp. started (see Section  
442 1.3).

443 For filtering: Returns all jobs that were queued resp. started in the specified queue.

## 444 4.8.12 wallclockTime

445 For monitoring: Accumulated time the job spent in "Queued" or "Started" states. Implementations MAY  
446 determine this value by subtracting the current time resp. `finishTime` by the `dispatchTime` of the job.

447 For filtering: Returns all jobs that have consumed at least the specified amount of wall clock time.

## 448 4.8.13 cpuTime

449 For monitoring: This attribute specifies the amount of CPU time consumed by the job. This value includes  
450 only time the job spent in `JobState::RUNNING` (see Section 7.1).

451 For filtering: Returns all jobs that have consumed at least the specified amount of CPU time.

## 452 4.8.14 submissionTime

453 For monitoring: This attribute specifies the time at which the job was submitted. Implementations SHOULD  
454 use the submission time recorded by the DRM system, if available.

455 For filtering: Returns all jobs that were submitted at or after the specified submission time.

## 456 4.8.15 dispatchTime

457 For monitoring: The time the job first entered a "Started" state (see Section 7.1). On job restart or re-  
458 scheduling, this value does not change.

459 For filtering: Returns all jobs that entered a "Started" state at, or after the specified dispatch time.

## 460 4.8.16 finishTime

461 For monitoring: The time the job first entered a "Terminated" state (see Section 7.1).

462 For filtering: Returns all jobs that entered a "Terminated" state at or after the specified finish time.

463

## 464 5 Common Exceptions

465 The exception model specific error information that can be returned by a DRMAA implementation on  
466 method calls.

```
467 exception AuthorizationException {string message;};
468 exception DefaultContactStringException {string message;};
469 exception DeniedByDrmException {string message;};
470 exception DrmCommunicationException {string message;};
471 exception TryLaterException {string message};
```

Same discussion as above

```

472     exception SessionManagementException {string message;};
473     exception TimeoutException {string message;};
474     exception InternalException {string message;};
475     exception InvalidArgumentException {string message;};
476     exception InvalidSessionException {string message;};
477     exception InvalidStateException {string message;};
478     exception OutOfMemoryException {string message;};
479     exception UnsupportedAttributeException {string message;};
480     exception UnsupportedOperationException {string message;};

```

481 If not defined otherwise, the exceptions have the following meaning:

482 **AuthorizationException:** The user is not authorized to perform the given function.

483 **DefaultContactStringException:** The DRMAA implementation could not use the default contact string  
484 to connect to DRM system.

485 **DeniedByDrmException:** The DRM system rejected the operation due to security issues.

486 **DrmCommunicationException:** The DRMAA implementation could not contact the DRM system. The  
487 problem source is unknown to the implementation, so it is unknown if the problem is transient or not.

488 **TryLaterException:** The DRMAA implementation detected a transient problem with performing the  
489 operation, for example due to excessive load. The application is recommended to retry the call.

490 **SessionManagementException:** A problem was encountered while trying to create / open / close /  
491 destroy a session.

492 **TimeoutException:** The timeout given in one the waiting functions was reached without successfully  
493 finishing the waiting attempt.

494 **InternalException:** An unexpected or internal error occurred in the DRMAA library, for example a system  
495 call failure. It is unknown if the problem is transient or not.

496 **InvalidArgumentException:** From the viewpoint of the DRMAA library, a function parameter is invalid  
497 or inappropriate for the particular function call.

498 **InvalidSessionException:** The session used for the function is not valid, for example since it was closed  
499 before.

500 **InvalidStateException:** The function call is not allowed in the current state of the job.

501 **OutOfMemoryException:** This exception can be thrown by any method at any time when the DRMAA  
502 implementation has run out of free memory.

503 **UnsupportedAttributeException:** The optional attribute is not supported by the DRMAA implemen-  
504 tation.

505 **UnsupportedOperationException:** The function is not supported by the DRMAA implementation. One  
506 example is the registration of an event callback function.

507 .

The DRMAA specification assumes that programming languages targeted by language bindings typically

We might want to introduce InvalidTemplate for separating input parameter issues

support the concept of exceptions. If a destination language does not support them (like ANSI C), the language binding specification SHOULD map error conditions to an appropriate consistent concept. A language binding MAY chose to model exceptions as numeric error code return values, and return values as additional output parameters of the operation. In this case, the language binding specification SHOULD specify numeric values for all DRMAA error constants.

The representation of exceptions in the language binding MUST support a possibility to express an exception cause as textual description. Implementations MAY use this text to express DRMS-specific error conditions that are outside of the DRMAA scope.

Object-oriented language bindings MAY decide to derive all exceptions from one or multiple exception base classes, in order to support generic catch clauses. Whenever it is appropriate, language bindings SHOULD replace DRMAA exceptions by their semantically equivalent native exception from the application runtime environment.

Language bindings MAY decide to introduce a hierarchical ordering of the DRMAA exceptions through class derivation. In this case, any new exceptions added for aggregation purposes SHOULD be prevented from being thrown, for example by marking them as abstract.

The `UnsupportedAttributeException` may either be raised by the setter function for the attribute or by the job submission function. A consistent decision for either one or the other approach MUST be made by the language binding specification.

508 (See footnote)<sup>10</sup>

## 509 6 The DRMAA Session Concept

510 DRMAA relies on an overall session concept, which supports the persistency of job and advance reservation  
 511 information over multiple application runs. This supports short-lived applications that need to work with  
 512 DRM system state spanning multiple application runs. Typical examples are job submission portals or  
 513 command-line tools. The session concept is also intended to allow implementations to perform DRM system  
 514 attach / detach operations at dedicated points in the application control flow.

### 515 6.1 SessionManager Interface

```
516 interface SessionManager{
517     readonly attribute string drmsName;
518     readonly attribute Version drmaaVersion;
519     readonly attribute boolean reservationSupported;
520     JobSession createJobSession(in string sessionId,
521                               in string contactString);
522     ReservationSession createReservationSession(in string sessionId,
523                                               in string contactString);
524     MonitoringSession createMonitoringSession (in string contactString);
```

<sup>10</sup> Comparison to DRMAA 1.0: The `InconsistentStateException` was removed, since it is semantically equal to the `InvalidStateException` (conf. call Jan 7th 2010) The former `HoldInconsistentStateException`, `ReleaseInconsistentStateException`, `ResumeInconsistentStateException`, and `SuspendInconsistentStateException` from DRMAA v1.0 are now expressed as single `InvalidStateException` with different meaning per raising method. (F2F meeting July 2009)

```

525     JobSession openJobSession(in string sessionId);
526     ReservationSession openReservationSession(in string sessionId);
527     void closeJobSession(in JobSession s);
528     void closeReservationSession(in ReservationSession s);
529     void closeMonitoringSession(in MonitoringSession s);
530     void destroyJobSession(in string sessionId);
531     void destroyReservationSession(in string sessionId);
532     StringList getJobSessions();
533     StringList getReservationSessions();
534 }

```

535 The **SessionManager** interface is the main interface for establishing communication with a given DRM system.  
 536 By the help of this interface, sessions for job management, monitoring, and/or reservation management  
 537 can be maintained.

538 An implementation MUST allow the application to keep multiple DRMAA job session, reservation session  
 539 and monitoring session instances open at the same time. The implementation SHOULD take care of according  
 540 race conditions.

541 Job and reservation sessions maintain persistent state information (about jobs resp. reservations created)  
 542 between application runs. State data SHOULD be persisted by the library implementation or the DRMS  
 543 itself (if supported) after closing the session through the according method in the **SessionManager** interface.

544 The re-opening of a session MUST be possible on the machine where the session was originally created.  
 545 Implementations MAY also offer to re-open the session on another machine.

546 The state information SHOULD be kept until the job resp. reservation session is explicitly reaped by the  
 547 according destroy method in the **SessionManager** interface. If an implementation runs out of resources for  
 548 storing the session information, the closing function SHOULD throw a **SessionManagementException**. If  
 549 an application ends without closing the session properly, the behavior of the DRMAA implementation is  
 550 undefined.

551 (See footnote)<sup>11</sup>

### 552 6.1.1 drmsName

553 A system identifier denoting a specific type of DRM system, e.g. "LSF" or "GridWay". It is intended  
 554 to support conditional code blocks in the DRMAA application that rely on DRMS-specific details of the  
 555 DRMAA implementation. Implementations SHOULD NOT make versioning information of the particular  
 556 DRM system a part of this attribute value.

### 557 6.1.2 drmaaVersion

558 A combination of minor / major version number information for the DRMAA implementation. The major  
 559 version number MUST be the constant value "2", the minor version number SHOULD be used by the

<sup>11</sup> Comparison to DRMAA 1.0: The concept of a factory from GFD.130 was removed (solves issue #6276). Version 2.0 of DRMAA supports restartable sessions by the newly introduced SessionManager interface. It allows creating multiple concurrent sessions for job submission (solves issue #2821), which can be restarted by their generated session name (solves issue #2820). Session.init() and Session.exit() functionalities are moved to the according session creation and closing routines. The descriptions were fixed accordingly (solves issue #2822). The AlreadyActiveSession error was removed. (F2F meeting July 2009) The drmaaImplementation attribute from DRMAA 1.0 was removed, since it was redundant to the drmsInfo attribute. This one is now available in the new SessionManager interface. (F2F meeting July 2009).

560 DRMAA implementation for expressing its own versioning information.

561 6.1.3 `reservationSupported`

562 The attribute indicates if advance reservation is supported by the DRMAA implementation. If `False`, all  
563 methods related to advance reservation will raise an `UnsupportedOperationException` if being used.

564

565 (See footnote)<sup>12</sup>

New, needs  
group ap-  
proval

566 6.1.4 `createJobSession` / `createReservationSession` / `createMonitoringSession`

567 The method creates a new session instance of the particular type for the application. On successful completion  
568 of this method, the necessary initialization for making the session usable MUST be completed. Examples are  
569 the connection establishment from the DRMAA library to the DRM system, or the prefetching of information  
570 from non-thread-safe operating system calls, such as `getHostByName`.

571 The `contactString` parameter is an implementation-dependent string that SHALL allow the application to  
572 specify which DRM system instance to use. A contact string represents a specific installation of a specific  
573 DRM system, e.g. a Condor central manager machine at a given IP address, or a Grid Engine ‘root’ and  
574 ‘cell’. Contact strings are always implementation dependent and therefore opaque to the application. If  
575 `contactString` has the value `UNSET`, a default DRM system SHOULD be contacted. The manual configura-  
576 tion or automated detection of a default contact is implementation-specific.

577 The `sessionName` parameter denotes a specific name to be used for the new session. If a session with such  
578 a name was created before, the method MUST throw an `InvalidArgumentException`. In all other cases,  
579 including if the provided name has the value `UNSET`, a new session MUST be created with a unique name  
580 generated by the implementation. A `MonitoringSession` instance has no persistent state, and therefore  
581 does not support the name concept.

582 If the DRM system does not support advance reservation, than `createReservationSession` SHALL throw  
583 an `UnsupportedOperationException`.

584 6.1.5 `openJobSession` / `openReservationSession`

585 The method is used to open a persisted `JobSession` resp. `ReservationSession` instance that has previously  
586 been created under the given `sessionName`. The implementation MUST support the case that the session  
587 have been created by the same application or by a different application running on the same machine. The  
588 implementation MAY support the case that the session was created resp. updated on a different machine.  
589 If no session with the given `sessionName` exists, an `InvalidArgumentException` MUST be raised.

590 If the session described by `sessionName` was already opened before, implementations MAY return the same  
591 `JobSession` resp. `ReservationSession` instance.

592 If the DRM system does not support advance reservation, `openReservationSession` SHALL throw an  
593 `UnsupportedOperationException`.

<sup>12</sup>This attribute is intended to avoid test calls for checking if advance reservation is supported by the implementation

594 6.1.6 `closeJobSession / closeReservationSession / closeMonitoringSession`

595 The method MUST do whatever work is required to disengage from the DRM system. It SHOULD be callable  
 596 only once, by only one of the application threads. This SHOULD be ensured by the library implementation.  
 597 Additional calls beyond the first SHOULD lead to a `NoActiveSessionException` error notification.

598 For `JobSession` resp. `ReservationSession` instances, the according state information MUST be saved to  
 599 some stable storage before the method returns. This method SHALL NOT affect any jobs or reservations in  
 600 the session (e.g., queued and running jobs remain queued and running).

601 If the DRM system does not support advance reservation, `closeReservationSession` SHALL throw an  
 602 `UnsupportedOperationException`.

603 6.1.7 `destroyJobSession / destroyReservationSession`

604 The method MUST do whatever work is required to reap persistent session state and cached job state  
 605 information for the given session name. If `JobSession` resp. `ReservationSession` instances for the given  
 606 name exist, they MUST become invalid after this method was finished successfully. Invalid sessions MUST  
 607 throw `InvalidSessionException` on every attempt of utilization. This method SHALL NOT affect any  
 608 jobs resp. reservations in the session in their operation, e.g. queued and running jobs remain queued and  
 609 running.

610 If the DRM system does not support advance reservation, `destroyReservationSession` SHALL throw an  
 611 `UnsupportedOperationException`.

612 6.1.8 `getJobSessions / getReservationSessions`

613 This method returns a list of `JobSession` resp. `ReservationSession` names that are valid input for a  
 614 `openJobSession` resp. `openReservationSession` call.

615 If the DRM system does not support advance reservation, `getReservationSessions` SHALL throw an  
 616 `UnsupportedOperationException`.

## 617 7 Working with Jobs

618 A DRMAA job represents a single computational activity that is executed by the DRM system on a execution  
 619 host, typically as operating system process. The `JobSession` interface represents all control and monitoring  
 620 functions commonly available in DRM systems for such jobs as a whole, while the `Job` interface represents the  
 621 common functionality for single jobs. Sets of jobs resulting from a bulk submission are separately represented  
 622 by the `JobArray` interface. `JobTemplate` instances allow to formulate conditions and requirements for the  
 623 job execution by the DRM system.

## 624 7.1 The DRMAA State Model

625 DRMAA defines the following job states:

```
626 enum JobState {
 627   UNDETERMINED, QUEUED, QUEUED_HELD, RUNNING, SUSPENDED, REQUEUED,
 628   REQUEUED_HELD, DONE, FAILED};
```

629 **UNDETERMINED:** The job status cannot be determined. This is a permanent issue, not being solvable  
 630 by querying again for the job state.

631 **QUEUED:** The job is queued for being scheduled and executed.

632 **QUEUED\_HELD:** The job has been placed on hold by the system, the administrator, or the submitting  
 633 user.

634 **RUNNING:** The job is running on a execution host.

635 **SUSPENDED:** The job has been suspended by the user, the system or the administrator.

636 **REQUEUED:** The job was re-queued by the DRM system, and is eligible to run.

637 **REQUEUED\_HELD:** The job was re-queued by the DRM system, and is currently placed on hold.

638 **DONE:** The job finished without an error.

639 **FAILED:** The job exited abnormally before finishing.

640 If a DRMAA job state has no representation in the underlying DRMS, the DRMAA implementation MAY  
 641 never report that job state value. However, all DRMAA implementations MUST provide the `JobState`  
 642 enumeration as given here. An implementation SHOULD NOT return any job state value other than those  
 643 defined in the `JobState` enumeration.

644 The status values relate to the DRMAA job state transition model, as shown in Figure 1.

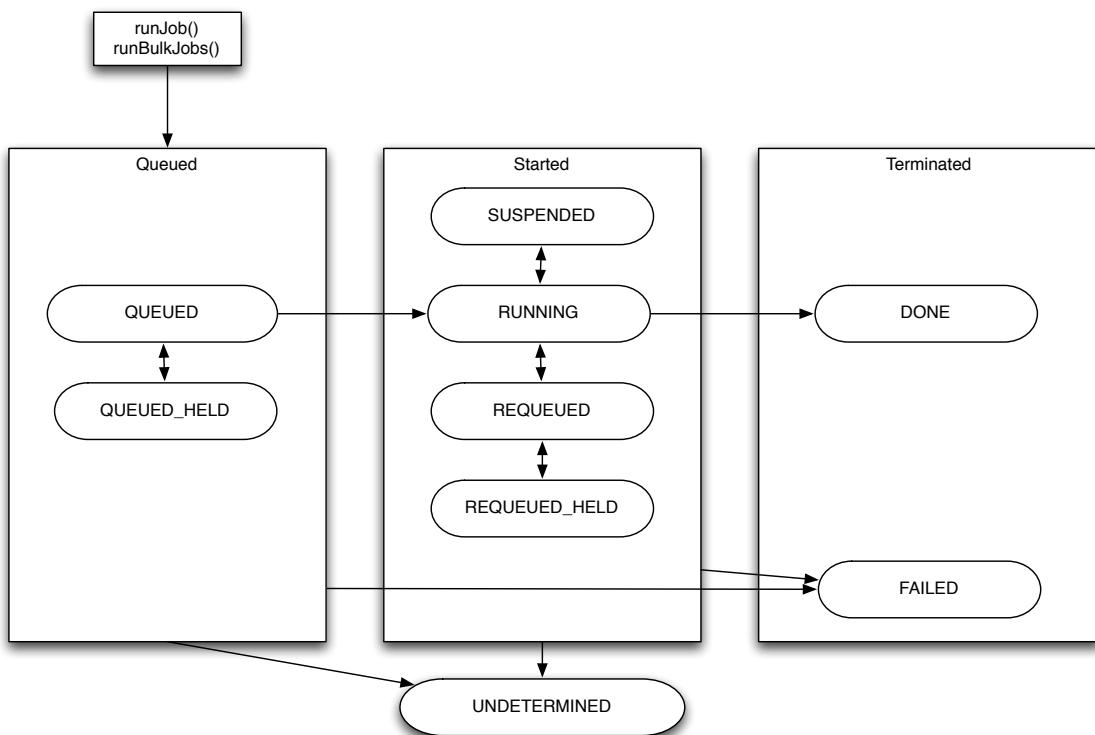


Figure 1: DRMAA Job State Transition Model

645 The transition diagram in Figure 1 expresses the classification of possible job states into "Queued", "Started",  
 646 and "Terminated". This is relevant for the job waiting functions (see Section 7.2 and Section 7.5), which  
 647 operate on job state classes only. The "Terminated" class of states is final, meaning that further state  
 648 transition is allowed.

649 Implementations SHALL NOT introduce other job transitions (e.g. from `RUNNING` to `QUEUED`) beside the  
 650 ones stated in Figure 1, even if they might happen in the underlying DRM system. In case, implementations  
 651 MAY emulate the necessary intermediate steps for the DRMAA-based application.

652 Every job state information in DRMAA can be extended by a `subState` property, which expresses specialized  
 653 information about the state coming from the implementation or the particular DRM system. The possible  
 654 values of this attribute are implementation-specific, but should be documented properly. In case of additional  
 655 job state information from the DRM system, such as extra states for staging phases or details on the hold  
 656 reason, implementations SHOULD map them to the `subState` information. Implementations of the DRMAA  
 657 API SHOULD define a DRMS-specific data structure for the sub-state information that can be converted to  
 658 / from the data type defined by the language binding.

The IDL definition declares the sub state attributes as type `any`, expressing the fact that the language binding MUST map the data type to a generic language type (e.g. `void*`, `Object`) that maintains source code portability across DRMAA implementations and still accepts an `UNSET` value.

659 The DRMAA job state model can be mapped to other high-level API state models. Table 3 gives a non-  
 660 normative set of examples.

Complete and  
re-check job  
state map-  
ping

DRMAA JobState	SAGA JobState [3]	OGSA-BES Job State [2]
UNDETERMINED	N/A	N/A
QUEUED	Running	Pending (Queued)
QUEUED_HELD		
RUNNING	Running	Running (Executing)
SUSPENDED	Suspended	Running (Suspended)
REQUEUED		
REQUEUED_HELD		
DONE	Done	Finished
FAILED	Cancelled, Failed	Cancelled, Failed

Table 3: Example Mapping of DRMAA Job States

661 (See footnote)<sup>13</sup>

<sup>13</sup> Comparison to DRMAA 1.0:

The differentiation between the system hold, user hold, and system / user hold job states was removed (conf. call Jan 20th 2009). There is only one hold state now. A job can now change its state from one of the SUSPENDED states to the QUEUED\_ACTIVE state (conf. call Jan 20th 2009, solves issue #2788). The job state UNDETERMINED is now clearer defined. It expressed a permanent issue, meaning that the job state will not change by just waiting. Temporary problems in the detection of the job state are now expressed by the TryLaterException (conf. call Feb 5th 2009, solves issue #2783). The description of the FAILED state was extended to support a more specific differentiation between different job failure reasons. The new subState feature allows the DRMAA implementation to provide better information, if available. There was no portable way of standardizing extended failure information in a better way. (conf. call May 12th 2009, solves issue #5875) The different suspend job states from DRMAA1 (user suspended, system suspended, user / system suspended) are now combined into one suspend state. DRM systems with the need to express the different suspend reasons can use the new sub-state feature (conf. call Mar 5th 2010).

## 662 7.2 JobSession Interface

663 A job session instance acts as container for job instances controlled through the DRMAA API. The session  
 664 methods support the submission of new jobs, the monitoring and the control of existing jobs. The relationship  
 665 between jobs and their session MUST be persisted, as described in Section 6.1.

```
666 interface JobSession {
667     readonly attribute string contact;
668     readonly attribute string sessionName;
669     readonly attribute boolean notificationSupported;
670     JobList getJobs(in JobInfo filter);
671     Job runJob(in JobTemplate jobTemplate);
672     JobArray runBulkJobs(
673         in JobTemplate jobTemplate,
674         in long beginIndex,
675         in long endIndex,
676         in long step);
677     Job waitAnyStarted(in JobList jobs, in TimeAmount timeout);
678     Job waitAnyTerminated(in JobList jobs, in TimeAmount timeout);
679     void registerEventNotification(in DrmaaCallback callback);
680 };
```

681 (See footnote)<sup>14</sup>

## 682 7.2.1 contact

683 This attribute contains the **contact** value that was used in the `SessionManager::createJobSession` call  
 684 for this instance (see Section 6.1). If no value was originally provided, the default contact string from the  
 685 implementation MUST be returned. This attribute is read-only.

## 686 7.2.2 sessionName

687 This attribute contains the **sessionName** value that was used in the `SessionManager::createJobSession`  
 688 or `SessionManager::openJobSession` call for this instance (see Section 6.1). This attribute is read-only.

<sup>14</sup> Comparison to DRMAA 1.0: The original separation between `synchronize()` and `wait()` was replaced by a complete new synchronization semantic in the API. DRMAA2 has now two methods, `waitStarted()` and `waitTerminated()`. The first waits for any state that expresses that the job was started, the second for any terminal status. Both methods are available on session level (wait for any of the given jobs to start / end) or on single job level (solves issue #5880 and #2838). The function returns always a `Job` object, in order to allow chaining, e.g. `job.wait(JobStatus.RUNNING).hold()`. The session-level functions implement the old DRMAA `wait(SESSION_ANY)`. The old `synchronize()` semantics are no longer directly supported - instead, the DRMAA application should use a looped `Job.wait... / JobSession.waitAny...` call. The result is a more condensed and responsive API, were the application can decide to keep the user informed during synchronization on a set of jobs. DRMAA library implementations should also become easier to design, since the danger of multithreading side effects inside the DRMAA API is reduced by this change. As a side effect, `JOB_IDS_SESSION_ANY` and `JOB_IDS_SESSION_ALL` are no longer needed. The special consideration of a partial failures during `SESSION_ALL` wait activities is also no longer necessary (F2F meeting July 2009). The `JobSession` now allows to fetch also information about jobs that were not submitted through DRMAA (conf. call June 23th 2010).

## 689 7.2.3 notificationSupported

690 The attribute indicates if event notification is supported by the DRMAA implementation for the job session.  
 691 If **False**, then `registerEventNotification` will raise an `UnsupportedOperationException` if being used.

New, needs group approval

## 693 7.2.4 getJobs

694 This method returns a sequence of jobs that belong to the job session. The `filter` parameter allows to  
 695 choose a subset of the session jobs as return value. The attribute semantics for the `filter` argument are  
 696 explained in Section 4.8. If no job matches or the session has no jobs attached, the method MUST return  
 697 an empty sequence instance. If `filter` is `UNSET`, all session jobs MUST be returned.

698 Time-dependent effects of this method, such as jobs no longer matching to filter criteria on evaluation time,  
 699 are implementation-specific. The purpose of the filter parameter is to keep scalability with a large number  
 700 of jobs per session. Applications therefore must consider the possibly changed state of jobs during their  
 701 evaluation of the method result.

## 702 7.2.5 runJob

703 The `runJob` method submits a job with the attributes defined in the job template parameter. It returns a  
 704 `Job` object that represents the job in the underlying DRM system. Depending on the job template settings,  
 705 submission attempts may be rejected with an `InvalidArgumentException`. The error details SHOULD  
 706 provide further information about the attribute(s) responsible for the rejection.

707 When this method returns a valid `Job` instance, the following conditions SHOULD be fulfilled:

- 708 • The job is part of the persistent state of the job session.
- 709 • All non-DRMAA and DRMAA interfaces to the DRM system report the job as being submitted to  
 710 the DRM system.
- 711 • The job has one of the DRMAA job states.

## 712 7.2.6 runBulkJobs

713 The `runBulkJobs` method creates a set of parametric jobs, each with attributes defined in the given  
 714 job template. Each job in the set is identical, except for the job template attributes that include the  
 715 `JobTemplatePlaceholder::PARAMETRIC_INDEX` macro (see Section 7.4).

716 If any of the resulting parametric job templates is not accepted by the DRM system, the method call MUST  
 717 raise an `InvalidArgumentException`. No job from the set SHOULD be submitted in this case.

718 The first job in the set has an index equal to the `beginIndex` parameter of the method call. The smallest valid  
 719 value for `beginIndex` is 1. The next job has an index equal to `beginIndex + step`, and so on. The last job  
 720 has an index equal to `beginIndex + n * step`, where `n` is equal to `(endIndex - beginIndex) / step`. The  
 721 index of the last job may not be equal to `endIndex` if the difference between `beginIndex` and `endIndex` is not  
 722 evenly divisible by `step`. The `beginIndex` value must be less than or equal to the `endIndex` value, and only  
 723 positive index numbers are allowed, otherwise the method SHOULD raise an `InvalidArgumentException`.

724 Implementations MAY provide custom ways for the job to determine its index number.

725 The `runBulkJobs` method returns a `JobArray` (see Section 7.6) instance that represents the set of `Job` objects  
 726 created by the method call under a common array identifier. For each of the jobs in the array, the same  
 727 conditions as for the result of `runJob` SHOULD apply.

The largest valid value for `endIndex` MUST be defined by the language binding.

728 (See footnote)<sup>15</sup>

#### 729 7.2.7 `waitAnyStarted` / `waitAnyTerminated`

730 This method blocks until any of the jobs referenced in the `jobs` parameter entered one of the "Started" resp.  
 731 "Terminated" states (see Section 7.1). If the input list contains jobs that are not part of the session, the call  
 732 to `waitAnyStarted` SHALL fail with an `InvalidArgumentException`.

733 The `timeout` argument specifies the desired behavior when a result is not immediately available. The con-  
 734 stant value `INFINITE_TIME` may be specified to wait indefinitely for a result. The constant value `ZERO_TIME`  
 735 may be specified to return immediately. Alternatively, a number of seconds may be specified to indicate  
 736 how long to wait for a result to become available. If the invocation exits on timeout, an `TimeoutException`  
 737 SHALL be raised.

738 In a multi-threaded environment with multiple `JobSession::waitAny...` calls, only one of the active thread  
 739 SHOULD get the status change notification for a particular job, while the other threads SHOULD continue  
 740 waiting. If there are no more queryable jobs left in the session, all remaining waiting threads SHOULD fail  
 741 with an `InvalidStateException`. If thread A is waiting for a specific job with `Job::wait...`, and another  
 742 thread, thread B, waiting for that same job or with `JobSession::waitAny...`, than B SHOULD receive the  
 743 notification that the job has finished, thread A SHOULD fail with an `InvalidStateException`. Waiting for  
 744 a job state is a read-only operation.

745 An application waiting for some condition to happen in *all* jobs of a set is expected to perform looped calls  
 746 of these waiting functions.

747 (See footnote)<sup>16</sup>

#### 748 7.2.8 `registerEventNotification`

749 This method is used to register a `DrmaaCallback` interface (see Section 7.3) implemented by the DRMAA-  
 750 based application. If the callback functionality is not supported by the DRMAA implementation, the method  
 751 SHALL raise an `UnsupportedOperationException`. Implementations MAY support the registration of  
 752 multiple callback methods.

A language binding specification MUST define how the reference to an interface-compliant method can be given as argument to this method.

<sup>15</sup> There was a discussion (mailing list Jan 2011) about having specialized job templates for bulk submission, with support for the start / end index and a slots limit. We rejected that, since job templates are intended for re-usage.

<sup>16</sup> People typically ask for the `waitAll()` counterparts of these functions. Since they are so easy to implement in the application itself, we could not see any benefit in adding them. Due to the intended long-blocking operation, the DRM system would not be able to offer any better (meaning much faster) implementation to be wrapped by DRMAA.

753 

### 7.3 DrmaaCallback Interface

754 The DrmaaCallback interface allows the DRMAA library resp. the DRM system to inform the application  
 755 about relevant events from the DRM system in a asynchronous fashion. One expected use case is  
 756 loseless monitoring of job state transitions. The support for such callback functionality is optional, but all  
 757 implementations MUST define the DrmaaCallback interface type as given in the language binding.

```
758     interface DrmaaCallback {
759         void notify(in DrmaaNotification notification);
760     };
761
762     struct DrmaaNotification {
763         DrmaaEvent event;
764         Job job;
765         JobState jobState;
766     };
767
768     enum DrmaaEvent {
769         NEW_STATE, MIGRATED, ATTRIBUTE_CHANGE
770     };

```

771 The application callback interface is registered through the `JobSession::registerEventNotification`  
 772 method (see Section 7.2). The DrmaaNotification structure represents the notification information from  
 773 the DRM system. Implementations MAY extend this structure for further information (see Section 4). All  
 774 given information SHOULD be valid at least at the time of notification generation.

775 The DrmaaEvent enumeration defines standard event types for notification:

776 **NEW\_STATE** The job entered a new state, which is described in the jobState attribute of the notification  
 777 structure.

778 **MIGRATED** The job was migrated to another execution host, and is now in the given state.

779 **ATTRIBUTE\_CHANGE** A monitoring attribute of the job, such as the memory consumption, changed  
 780 to a new value. The jobState attribute MAY have the value UNSET on this event.

781 DRMAA implementations SHOULD protect themselves from unexpected behavior of the called application.  
 782 This includes indefinite delays or unexpected exceptions from the callee. An implementation SHOULD  
 783 also disallow any library calls while the callback function is running, to avoid recursion scenarios. It is  
 784 RECOMMENDED to raise `TryLaterException` in this case.

785 Scalability issues of the notification facility are out of scope for this specification. Implementations MAY  
 786 decide to support non-standardized throttling configuration options.

787 (See footnote)<sup>17</sup>

788 

### 7.4 JobTemplate Structure

789 In order to define the attributes associated with a job, a DRMAA application uses the JobTemplate structure.  
 It specifies any required job parameters and is passed to the DRMAA JobSession instance when job  
 execution is requested.

---

<sup>17</sup> We intentionally did not add `subState` to the notification information, since this would make callback interface implementations specific for the DRM system, without any chance for creating a portable DRMAA application.

```

790     struct JobTemplate {
791         StringList attributeNames;
792         string remoteCommand;
793         OrderedStringList args;
794         boolean submitAsHold;
795         boolean rerunnable;
796         Dictionary jobEnvironment;
797         string workingDirectory;
798         string jobCategory;
799         StringList email;
800         boolean emailOnStarted;
801         boolean emailOnTerminated;
802         string jobName;
803         string inputPath;
804         string outputPath;
805         string errorPath;
806         boolean joinFiles;
807         string reservationId;
808         string queueName;
809         long minSlots;
810         long maxSlots;
811         long priority;
812         OrderedStringList candidateMachines;
813         long minPhysMemory;
814         OperatingSystem machineOS;
815             CpuArchitecture machineArch;
816             AbsoluteTime startTime;
817             Dictionary drmsSpecific;
818             AbsoluteTime deadlineTime;
819             Dictionary stageInFiles;
820             Dictionary stageOutFiles;
821             Dictionary softResourceLimits;
822             Dictionary hardResourceLimits;
823             string accountingId;
824     };

```

825 The DRMAA job template concept makes a distinction between *mandatory* and *optional* attributes. Mandatory  
826 attributes MUST be supported by the implementation in the sense that they are evaluated on job  
827 submission. Optional attributes MAY be evaluated on job submission, but MUST be provided as part of the  
828 `JobTemplate` structure in the implementation. If an unsupported optional attribute has a value different to  
829 `UNSET`, the job submission MUST fail with a `UnsupportedAttributeException`. DRMAA applications are  
830 expected to check for the availability of optional attributes before using them.

831 Implementations MUST set all attribute values to `UNSET` on struct allocation. This ensures that both the  
832 DRMAA application and the library implementation can determine untouched attribute members. If not  
833 described differently in the following sections, all attributes SHOULD be allowed to have the `UNSET` value  
834 on job submission.

835 An implementation SHALL NOT extend the `JobTemplate` structure with implementation-specific attributes,  
 836 but SHOULD support according keys in the `drmsSpecific` attribute (see Section 7.4.9).

837 An implementation MAY support `JobTemplatePlaceholder` macros in more occasions than defined in this  
 838 specification.

A language binding specification SHOULD define how a `JobTemplate` instance is convertible to a string for printing, through whatever mechanism is most natural for the implementation language. The resulting string MUST contain the values of all set properties.

The initialization to `UNSET` SHOULD be realized without additional methods in the DRMAA interface, if possible. The according approach MUST be specified by the language binding.

Which attributes should allow the new `HOST_NAME` and `USER_NAME` place holders?

839

840 (See footnote)<sup>18</sup>

#### 841 7.4.1 attributeNames

842 The `attributeNames` list of strings SHALL enumerate the names of the required and of the supported  
 843 optional job template attributes.

844

This is especially intended for languages which do not provide an inherit notion of struct introspection and therefore map job template attribute access to getter / setter functions.

This doesn't make sense anymore, since job templates are now value types.

845 The support for this attribute is mandatory.

#### 846 7.4.2 remoteCommand

847 This attribute describes the command to be executed on the remote host. In case this parameter contains  
 848 path information, it MUST be seen as relative to the execution host file system and is therefore evaluated  
 849 there. The implementation SHOULD NOT relate the value of this attribute to binary file management or  
 850 file staging activities. The behavior with an `UNSET` value is implementation-specific.

851 The support for this attribute is mandatory.

<sup>18</sup> Comparison to DRMAA 1.0: `JobTemplate` is now a value type, meaning that it maps to a struct in C. This removes the need for DRMAA-defined methods for construction and destruction of job templates. An eventual RPC scenario for DRMAA gets easier with this approach, since it is closer to the JSDL concept of a job description document.

Supported string placeholders for job template attributes are now listed in the `JobTemplatePlaceholder` enumeration, and must be filled with values by the language binding. Invalid job template settings are now only detected on job submission, not when the attribute is set.

Implementation-specific job template extensions were decided to be no longer supported, which hopefully fosters portable DRMAA-based source code. Implementation-specific job template settings are now covered by the `drmsSpecific` dictionary. This more generic approach also makes the old `nativeOptions` obsolete, so it was removed. Implementations therefore should support all relevant native settings explicitly as keys in the `drmsSpecific` dictionary. (conf. call May 26th 2010).

DRMAA1 supported the utilization of new DRM features through an old DRMAA implementation, based on the "nativeSpecification" field. A conf call (Jul 14th 2010) voted for dropping this intentionally. Implementations instead should be creative with their supported key names.

## 852 7.4.3 args

853 This attribute contains the list of command-line arguments for the job(s) to be executed.

854 The support for this attribute is mandatory.

## 855 7.4.4 submitAsHold

856 This attribute defines if the job(s) should be submitted as QUEUED or QUEUED\_HELD (see Section 7.1). Since  
857 the boolean UNSET value defaults to False, jobs are submitted as non-held if this attribute is not set.

858 The support for this attribute is mandatory.

## 859 7.4.5 rerunnable

860 This flag indicates if the submitted job(s) can safely be restarted by the DRM system, for example on a  
861 node failure or some other re-scheduling event. Since the boolean UNSET value defaults to False, jobs are  
862 submitted as not rerunnable if this attribute is not set. This attribute SHOULD NOT be used by the  
863 implementation to let the application denote the checkpointability of a job.

864 The support for this attribute is mandatory.

865 (See footnote)<sup>19</sup>

## 866 7.4.6 jobEnvironment

867 This attribute holds the environment variable key-value pairs for the execution machine(s). The values  
868 SHOULD override the execution host environment values if there is a collision.

869 The support for this attribute is mandatory.

## 870 7.4.7 workingDirectory

871 This attribute specifies the directory where the job resp. the bulk jobs are executed. If the attribute value  
872 is UNSET, the behavior is implementation dependent. Otherwise, the attribute value MUST be evaluated  
873 relative to the file system on the execution host. The attribute value MUST be allowed to contain either the  
874 JobTemplatePlaceholder::HOME\_DIRECTORY or the JobTemplatePlaceholder::PARAMETRIC\_INDEX place-  
875 holder (see Section 4.4).876 The workingDirectory attribute should be specified by the application in a syntax that is common at the  
877 host where the job is executed. Implementations MAY perform according validity checks on job submission.  
878 If the attribute is set and no placeholder is used, an absolute directory specification is expected. If the  
879 attribute is set and the job was submitted successfully and the directory does not exist on the execution  
880 host, the job MUST enter the state JobState::FAILED.

881 The support for this attribute is mandatory.

---

<sup>19</sup> The differentiation between rerunnable and checkpointable was decided on a conf call (Aug 25th 2010)

## 882 7.4.8 jobCategory

883 DRMAA facilitates writing DRM-enabled applications even though the deployment properties, in particular  
884 the configuration of the DRMS, cannot be known in advance.

885 Through the `jobCategory` string attribute, a DRMAA application can specify additional needs of the job(s)  
886 that are to be mapped by the implementation or DRM system itself to DRMS-specific options. It is intended  
887 as non-programmatic extension of DRMAA job submission capabilities. The mapping is performed during  
888 the process of job submission. Each category expresses a particular type of job execution that demands  
889 site-specific configuration, for example path settings, environment variables, or application starters such as  
890 MPIRUN.

891 A valid input SHOULD be one of the returned strings in `MonitoringSession::drmsJobCategoryNames` (see  
892 Section 9.1), otherwise an `InvalidArgumentException` SHOULD be raised.

893 A non-normative recommendation of category names is maintained at:

894 <http://www.drmaa.org/jobcategories/>

895 In case the name is not taken from the DRMAA working group recommendations, it should be self-  
896 explanatory for the user to understand the implications on job execution. Implementations are recommended  
897 to provide a library configuration facility, which allows site administrators to link job category names with  
898 specific product- and site-specific configuration options, such as submission wrapper shell scripts.

899 The interpretation of the supported `jobCategory` values is implementation-specific. The order of precedence  
900 for the `jobCategory` attribute value, the `drmsSpecific` attribute value, or other attribute values  
901 is implementation-specific. It is RECOMMENDED to overrule job template settings with a conflicting  
902 `jobCategory` setting, and overrule a given `jobCategory` with a conflicting `drmsSpecific` setting.

903 The support for this attribute is mandatory.

## 904 7.4.9 drmsSpecific

905 This dictionary allows the application to pass DRMS-specific native options as key-value pairs during job  
906 submission. In contrast to the usage of predefined configuration sets with the `jobCategory` attribute, this  
907 supports passing DRMS-specific options directly. The interpretation of keys and values in this dictionary is  
908 implementation-specific. Valid key strings should be documented by the implementation.

909 The order of precedence rules is described in the `jobCategory` section above.

910 The support for this attribute is mandatory.

## 911 7.4.10 email

912 This attribute holds a list of email addresses that should be used to report DRM information. Content and  
913 formatting of the emails are defined by the implementation resp. the DRM system. If the attribute value is  
914 UNSET, no emails SHOULD be sent to the user running the job(s), even if the DRM system default behavior  
915 is to send emails on some event.

916 The support for this attribute is optional. If an implementation cannot configure the email notification  
917 functionality of the DRM system, or if the DRM system has no such support, the attribute SHOULD NOT  
918 be supported in the implementation.

This became  
an optional  
attribute,  
since we  
mandate the  
'switch off'  
semantic in  
case of UNSET

919

920 (See footnote)<sup>20</sup>

921 7.4.11 emailOnStarted / emailOnTerminated

922 This flag indicates if the given email address(es) SHOULD get a notification when the job (or any of the  
 923 bulk jobs) entered one of the "Started" resp. "Terminated" states. Since the boolean UNSET value defaults  
 924 to **False**, the notification about state changes SHOULD NOT be sent if the attribute is not set.

925 The support for this attribute is optional. It SHALL only be supported if the **email** attribute is supported  
 926 in the implementation.

927 7.4.12 jobName

928 The job name attributes allows the specification of an additional non-unique string identifier for the job(s).  
 929 The implementation MAY truncate any client-provided job name to an implementation-defined length.

930 The support for this attribute is mandatory.

931 7.4.13 inputPath / outputPath / errorPath

932 This attribute specifies standard input / output / error stream of the job as a path to a file. If the attribute  
 933 value is UNSET, the behavior is implementation dependent. Otherwise, the attribute value MUST be evaluated  
 934 relative to the file system of the execution host in a syntax that is common at the host. Implementations  
 935 MAY perform according validity checks on job submission. The attribute value MUST be allowed to contain  
 936 any of the **JobTemplatePlaceholder** placeholders (see Section 4.4). If the attribute is set and no placeholder  
 937 is used, an absolute file path specification is expected.

938 If the **outputPath** or **errorPath** file does not exist at the time the job is about to be executed, the file  
 939 SHALL first be created. An existing **outputPath** or **errorPath** file SHALL be opened in append mode.

940 If the attribute is set and the job was submitted successfully and the file cannot be created / read / written  
 941 on the execution host, the job MUST enter the state **JobState::FAILED**.

942 The support for this attribute is mandatory.

943 7.4.14 joinFiles

944 Specifies whether the error stream should be intermixed with the output stream. Since the boolean UNSET  
 945 value defaults to **False**, intermixing SHALL NOT happen if the attribute is not set.

946 If this attribute is set to **True**, the implementation SHALL ignore the value of the **errorPath** attribute and  
 947 intermix the standard error stream with the standard output stream as specified by the **outputPath**.

948 The support for this attribute is mandatory.

949 7.4.15 stageInFiles / stageOutFiles

950 Specifies what files should be transferred (staged) as part of the job execution. The data staging operation  
 951 MUST be a copy operation between submission host and one resp. all execution hosts. File transfers between  
 952 execution hosts are not covered by DRMAA.

---

<sup>20</sup> The blockEmail attribute in the **JobTemplate** was replaced by the UNSET semantic for the email addresses. (conf. call July 28th 2010).

953 For each key-value pair in the dictionary, the key defines the source path of one file or directory, and the value  
 954 defines the destination path of one file or directory for the copy operation. For `stageInFiles`, the submission  
 955 host acts as source, and the execution host(s) act as destination. For `stageOutFiles`, the execution host(s)  
 956 acts as source, and the submission host act as destination.

957 All values MUST be evaluated relative to the file system on the host in a syntax that is common at that  
 958 host. Implementations MAY perform according validity checks on job submission. Paths on the execution  
 959 host(s) MUST be allowed to contain any of the `JobTemplatePlaceholder` placeholders. Paths on the sub-  
 960 mission host MUST be allowed to contain the `JobTemplatePlaceholder::PARAMETRIC_INDEX` placeholder  
 961 (see Section 4.4). If no placeholder is used in the values, an absolute path specification on the particular  
 962 host SHOULD be assumed by the implementation.

963 Jobs SHOULD NOT enter `JobState::DONE` unless all staging operations are finished. The behavior in  
 964 case of missing files is implementation-specific. The support for wildcard operators in path specifications is  
 965 implementation-specific.

966 The support for this attribute is optional.

967

968 (See footnote)<sup>21</sup>

Needs final  
approval by  
the group.

#### 969 7.4.16 reservationId

970 Specifies the identifier of the advance reservation associated with the job(s). The application is expected  
 971 to create an advance reservation through the `ReservationSession` interface, the resulting `reservationId`  
 972 (see Section 8.3) then acts as valid input for this job template attribute. Implementations MAY support an  
 973 reservation identifier from non-DRMAA information sources as valid input.

974 The support for this attribute is mandatory.

#### 975 7.4.17 queueName

976 This attribute specifies the name of the queue the job(s) should be submitted to. In case this attribute  
 977 value is `UNSET`, and `MonitoringSession::getAllQueues` returns a list with a minimum length of 1, the  
 978 implementation SHOULD use the DRM systems default queue.

979 The `MonitoringSession::getAllQueues` method (see 9.1) supports the determination of valid queue names.  
 980 Implementations SHOULD allow these queue names to be used in the `queueName` attribute. Implementa-  
 981 tions MAY also support queue names from other non-DRMAA information sources as valid input. If no  
 982 default queue is defined or if the given queue name is not valid, the job submission MUST lead to an  
 983 `InvalidArgumentException`.

984 If `MonitoringSession::getAllQueues` returns an empty list, this attribute MUST be only accepted with  
 985 the value `UNSET`.

986 Since the meaning of "queues" is implementation-specific, there is no implication on the effects in the DRM  
 987 system when using this attribute. As one example, requesting a number of slots for a job in one queue has no

<sup>21</sup> Comparison to DRMAA 1.0: New job template attributes for file transfers were introduced. They allow to express a set of file staging activities, similar to the approach in LSF and SAGA. They replace the old `transferFiles` attribute, the according `FileTransferMode` data structure and the special host definition syntax in `inputPath` / `outputPath` / `errorPath` (different conf. calls, SAGA F2F meeting, solves issue #5876)

988 implication on the number of utilized machines at run-time. Implementations therefore SHOULD document  
989 the effects of this attribute accordingly.

990 The support for this attribute is mandatory.

991 7.4.18 minSlots / maxSlots

992 This attribute expresses the minimum / maximum number of slots requested per job (see also Section 1.3).  
993 If the value of `minSlots` is UNSET, it SHOULD default to 1. If the value of `maxSlots` is UNSET, it SHOULD  
994 default to the value of `minSlots`.

995 Implementations MAY interpret the slot count as number of concurrent processes being allowed on one  
996 machine. If this interpretation is taken, and `minSlots` is greater than 1, than the `jobCategory` SHOULD  
997 also be demanded on job submission, in order to express the nature of the intended parallel job execution.

998 The support for this attribute is mandatory.

999 7.4.19 priority

1000 This attribute specifies the scheduling priority for the job. The interpretation of the given value incl. an  
1001 UNSET value is implementation-specific.

1002 The support for this attribute is mandatory.

1003 7.4.20 candidateMachines

1004 Requests that the job(s) should run on any subset (with minimum size of 1), or all of the given machines.  
1005 If the attribute value is UNSET, it should default to the result of the `MonitoringSession::getAllMachines`  
1006 method. If this resource demand cannot be fulfilled, an `InvalidArgumentException` SHOULD be raised  
1007 on job submission time. If the problem can only be detected after job submission, the job should enter  
1008 `JobState::FAILED`.

1009 The support for this attribute is mandatory.

1010 7.4.21 minPhysMemory

1011 This attribute denotes the minimum amount of physical memory in Kibibyte expected on the / all execution  
1012 host(s). If this resource demand cannot be fulfilled, an `InvalidArgumentException` SHOULD be raised  
1013 at job submission time. If the problem can only be detected after job submission, the job SHOULD enter  
1014 `JobState::FAILED` accordingly.

1015 The support for this attribute is mandatory.

1016 7.4.22 machineOS

1017 This attribute denotes the expected operating system type on the / all execution host(s). If this resource de-  
1018 mand cannot be fulfilled, an `InvalidArgumentException` SHOULD be raised on job submission time. If the  
1019 problem can only be detected after job submission, the job SHOULD enter `JobState::FAILED` accordingly.

1020 The support for this attribute is mandatory.

1021 (See footnote)<sup>22</sup>

1022 **7.4.23 machineArch**

1023 This attribute denotes the expected machine architecture on the / all execution host(s). If this resource  
 1024 demand cannot be fulfilled, an **InvalidArgumentException** SHOULD be raised on job submission time. If  
 1025 the problem can only be detected after job submission, the job should enter **JobState::FAILED**.

1026 The support for this attribute is mandatory.

1027 **7.4.24 startTime**

1028 This attribute specifies the earliest time when the job may be eligible to be run.

1029 The support for this attribute is mandatory.

1030 **7.4.25 deadlineTime**

1031 Specifies a deadline after which the implementation resp. the DRM system SHOULD change the job state  
 1032 to any of the "Terminated" states (see Section 7.1).

1033 The support for this attribute is optional.

1034 **7.4.26 softResourceLimits / hardResourceLimits**

1035 This attribute specifies the soft resp. hard limits on resource utilization of the job(s) on the execution host(s).  
 1036 The valid dictionary keys and their value semantics are defined in Section 4.3. An implementation MAY  
 1037 map the settings to an *ulimit(3)* on the operating system, if available.

1038 The support for this attribute is optional. If only a subset of the attributes from **ResourceLimitType** is  
 1039 supported by the implementation, and some of the unsupported attributes are used, the job submission  
 1040 SHOULD raise an **InvalidArgumentException** expressing the fact that resource limits are supported in  
 1041 general.

1042 Conflicts of these attribute values with any other job template attribute or with referenced advanced reser-  
 1043 vations are handled in an implementation-specific manner. Implementations SHOULD try to delegate the  
 1044 decision about parameter combination validity to the DRM system, in order to ensure similar semantics in  
 1045 different DRMAA implementations for this system.

1046

1047 (See footnote)<sup>23</sup>

<sup>22</sup> Requesting a particular operating system version is not supported by the majority of DRM systems (conf call Jul 28th 2010)

<sup>23</sup> In comparison to DRMAA 1.0, resource usage limitations can now be expressed by two dictionaries and an according standardized set of valid dictionary keys (LimitType). The idea is to allow a direct mapping to *ulimit(3)* semantics, which are supported by the majority of DRM system today. A separate run duration limit is no longer needed, since this is covered by the new CPU\_TIME limit parameter. (conf. call Jun 9th 2010).

Unclear what happens from DRMAA perspective if a soft limit is violated. We have no signals.

## 1048 7.4.27 accountingId

1049 This attribute denotes a string that can be used by the DRM system for job accounting purposes. Implementations SHOULD NOT utilize this information as authentication token, but only as additional identification  
 1050 information beside the implementation-specific authentication (see Section 11).

1052 The support for this attribute is optional.

## 1053 7.5 Job Interface

1054 Every job in the `JobSession` is expressed by an own instance of the `Job` interface. It allows to instruct the  
 1055 DRM system for a job status change, and to query the status attributes of the job in the DRM system.

```
1056 interface Job {
1057     readonly attribute string jobId;
1058     readonly attribute JobSession session;
1059     readonly attribute JobTemplate jobTemplate;
1060     void suspend();
1061     void resume();
1062     void hold();
1063     void release();
1064     void terminate();
1065     JobState getState(out any jobSubState);
1066     JobInfo getInfo();
1067     Job waitStarted(in TimeAmount timeout);
1068     Job waitTerminated(in TimeAmount timeout);
1069 };
```

1070 (See footnote)<sup>24</sup>

## 1071 7.5.1 jobId

1072 This attribute provides the string job identifier assigned to the job by the DRM system. It is intended as  
 1073 performant alternative for fetching a complete `JobInfo` instance for this information.

## 1074 7.5.2 session

1075 This attribute offers a reference to the `JobSession` instance that represents the session used for the job  
 1076 submission creating this `Job` instance.

---

<sup>24</sup> In comparison to DRMAA v1.0, DRMAA2 replaces the identification of jobs by strings with `Job` objects. This enables a tighter integration of job meta-data and identity, for the price of reduced performance in (so far not existing) DRMAA RPC scenarios. The former DRMAA `control()` with the `JobControlAction` structure is now split up into dedicated functions (such as `hold()` and `release()`) on the `Job` object.

Even though the DRMAAv2 surveys showed interest in interactive job support, this feature was intentionally left out. Reasons are the missing support in some major DRM systems, and the lack of a relevant DRMAA-related use case (conf. call Jan 7th 2010)

Issue #5877 (support for direct job signaling) was rejected, even though there was an according request from the SAGA WG. Issue #2782 (change attributes of submitted, but pending jobs) was rejected based on group decision.

## 1077 7.5.3 jobTemplate

1078 This attribute provides a reference to a `JobTemplate` instance that has equal values to the one that was  
 1079 used for the job submission creating this `Job` instance.

## 1080 7.5.4 suspend / resume / hold / release / terminate

1081 The job control functions allow modifying the status of the single job in the DRM system, according to the  
 1082 state model presented in Section 7.1.

1083 The `suspend` method triggers a transition from `RUNNING` to `SUSPENDED` state. The `resume` method triggers  
 1084 a transition from `SUSPENDED` to `RUNNING` state. The `hold` method triggers a transition from `QUEUED` to  
 1085 `QUEUED_HELD`, or from `REQUEUED` to `REQUEUED_HELD` state. The `release` method triggers a transition from  
 1086 `QUEUED_HELD` to `QUEUED`, or from `REQUEUED_HELD` to `REQUEUED` state. The `terminate` method triggers a  
 1087 transition from any of the "Started" states to one of the "Terminated" states. If the job is in an inappropriate  
 1088 state for the particular method, the method MUST raise an `InvalidStateException`.

1089 The methods SHOULD return after the action has been acknowledged by the DRM system, but MAY  
 1090 return before the action has been completed. Some DRMAA implementations MAY allow this method  
 1091 to be used to control jobs submitted externally to the DRMAA session, such as jobs submitted by other  
 1092 DRMAA sessions in other DRMAA implementations or jobs submitted via native utilities. This behavior is  
 1093 implementation-specific.

## 1094 7.5.5 getState

1095 This method allows the gather the current status of the job according to the DRMAA state model, together  
 1096 with an implementation specific sub state (see Section 7.1). It is intended as performant alternative for  
 1097 fetching a complete `JobInfo` instance for state checks. The timing conditions are described in Section 4.8.

1098 (See footnote)<sup>25</sup>

## 1099 7.5.6 getInfo

1100 This method returns a `JobInfo` instance for the particular job under the conditions described in Section 4.8.

## 1101 7.5.7 waitStarted / waitTerminated

1102 This method blocks until the job entered one of the "Started" resp. "Terminated" states (see Section 7.1).  
 1103 The `timeout` argument specifies the desired behavior when a result is not immediately available. The con-  
 1104 stant value `INFINITE_TIME` may be specified to wait indefinitely for a result. The constant value `ZERO_TIME`  
 1105 may be specified to return immediately. Alternatively, a number of seconds may be specified to indicate  
 1106 how long to wait for a result to become available. If the invocation exits on timeout, an `TimeoutException`  
 1107 SHALL be raised. If the job is in an inappropriate state for the particular method, the method MUST raise  
 1108 an `InvalidStateException`.

<sup>25</sup> The `getState()` function now also returns job subState information. This is intended as additional information for the given DRMAA job state, and can be used for expressing the hold state differentiation from DRMAA 1.0 (conf. call Mar 31st 2009).

## 1109 7.6 JobArray Interface

1110 The following section explains the set of methods and attributes defined in the `JobArray` interface. Any  
 1111 instance of this interface represent an *job array*, a common concept in many DRM systems for a job set created  
 1112 by one operation. In DRMAA, `JobArray` instances are only created by the `runBulkJobs` operation (see  
 1113 Section 7.2). `JobArray` instances differ from the `JobList` data structure due to their potential for representing  
 1114 a DRM system concept, while `JobList` is a DRMAA-only concept mainly realized by the language binding  
 1115 sequence support. Implementations SHOULD realize the `JobArray` functionality as wrapper for DRM system  
 1116 job arrays, if possible. If the DRM system has only single job support or incomplete job array support with  
 1117 respect to the DRMAA-provided functionality, implementations MUST realize the `JobArray` functionality  
 1118 on their own, for example based on looped operations with a list of jobs.

```
1119 interface JobArray {
1120     readonly attribute string jobArrayId;
1121     readonly attribute JobList jobs;
1122     readonly attribute JobSession session;
1123     readonly attribute JobTemplate jobTemplate;
1124     void suspend();
1125     void resume();
1126     void hold();
1127     void release();
1128     void terminate();
1129 };
```

Completely  
new, needs  
group ap-  
proval

1130  
 1131 (See footnote)<sup>26</sup>

### 1132 7.6.1 jobArrayId

1133 This attribute provides the string job identifier assigned to the job array by the DRM system. If the DRM  
 1134 system has no job array support, the implementation MUST generate a system-wide unique identifier for  
 1135 the result of the successful `runBulkJobs` operation.

### 1136 7.6.2 jobs

1137 This attribute provides the static list of jobs that are part of the job array.

1138 (See footnote)<sup>27</sup>

### 1139 7.6.3 session

1140 This attribute offers a reference to a `JobSession` instance that represents the session which was used for the  
 1141 job submission creating this `JobArray` instance.

<sup>26</sup> We are aware of the fact that some systems (e.g. LSF at the time of writing) do not support all DRMAA control operations offered for JobArrays. Since we intended to avoid optional DRMAA operations wherever we could, the text here mandates the implementation to simulate the `JobArray` support on its own. For example, looping over all jobs in the array and calling "suspend" for each one is trivial to implement and fulfills the same purpose.

<sup>27</sup> We were asked for offering a filter support similar to `JobSession` here. This was rejected by discussion on the list (Jan 2011), since the number of jobs returned here is normally comparatively short. In this case, the DRM system cannot provide any benefit over the looped check in the application itself.

## 1142 7.6.4 jobTemplate

1143 This attribute provides a reference to a `JobTemplate` instance that has equal values to the one that was  
 1144 used for the job submission creating this `JobArray` instance.

1145 (See footnote)<sup>28</sup>

## 1146 7.6.5 suspend / resume / hold / release / terminate

1147 The job control functions allow modifying the status of the job array in the DRM system, with the same  
 1148 semantic as with the counterparts in the `Job` interface (see Section 7.5). If one of the jobs in the array is in  
 1149 an inappropriate state for the particular method, the method MUST raise an `InvalidStateException`.

1150 The methods SHOULD return after the action has been acknowledged by the DRM system for all jobs in  
 1151 the array, but MAY return before the action has been completed. Some DRMAA implementations MAY  
 1152 allow this method to be used to control job arrays created externally to the DRMAA session, such as job  
 1153 arrays submitted by other DRMAA sessions in other DRMAA implementations or job arrays submitted via  
 1154 native utilities. This behavior is implementation-specific.

## 1155 8 Working with Advance Reservation

1156 Advance reservation is DRM system concept that allows the reservation of execution resources for jobs to be  
 1157 submitted. DRMAA encapsulates such functionality of DRM systems with the interfaces and data structures  
 1158 described in this chapter.

1159 DRMAA implementations for DRM systems that do not support advance reservation still MUST imple-  
 1160 mented the described interfaces, in order to keep source code portability for DRMAA-based applications.

## 1161 8.1 ReservationSession Interface

1162 Every `ReservationSession` instance represents a set of advance reservations in the DRM system. Every  
 1163 `Reservation` instance SHALL belong only to one `ReservationSession` instance.

```
1164 interface ReservationSession {
1165   readonly attribute string contact;
1166   readonly attribute string sessionName;
1167   Reservation getReservation(in string reservationId);
1168   Reservation requestReservation(in ReservationTemplate reservationTemplate);
1169   ReservationList getReservations();
1170 };
```

1171 If the DRM system does not support advance reservation, all methods in this interface SHALL throw an  
 1172 `UnsupportedOperationException`.

---

<sup>28</sup> The use case from SAGA perspective is that the user can easily resubmit the same job - just changing for example some command line parameter, but leaving the remainder fixed (mail by Andre Merzky, July 29th 2010).

## 1173 8.1.1 contact

1174 This attribute contains the `contact` value that was used in the `createReservationSession` call for this  
 1175 instance (see Section 6.1). If no value was originally provided, the default contact string from the implemen-  
 1176 tation MUST be returned. This attribute is read-only.

## 1177 8.1.2 sessionName

1178 This attribute contains the name of the session that was used for creating or opening this `Reservation`  
 1179 instance (see Section 6.1). This attribute is read-only.

## 1180 8.1.3 getReservation

1181 This method returns a `Reservation` instance that belongs to the session instance and has the given  
 1182 `reservationId`. If no reservation matches, the method SHALL raise an `InvalidArgumentException`. Time-  
 1183 dependent effects of this method are implementation-specific.

## 1184 8.1.4 requestReservation

1185 The `requestReservation` method SHALL request an advance reservation in the DRM system with at-  
 1186 tributes defined in the provided `ReservationTemplate`. On a successful reservation, the method returns a  
 1187 `Reservation` instance that represents the advance reservation in the underlying DRM system.

1188 The method SHALL raise an `InvalidArgumentException` if the reservation cannot be performed by the  
 1189 DRM system. It SHOULD further provide detailed information about the rejection cause in the extended  
 1190 error information (see Section 5).

1191 In case some of the conditions are not fulfilled after the reservation was successfully created, for example due  
 1192 to execution host outages, the reservation itself SHOULD remain valid.

## 1193 8.1.5 getReservations

1194 This method returns the list of reservations successfully created so far in this session, regardless of their start  
 1195 and ending time. The list of `Reservation` instances is only cleared in conjunction with the destruction of  
 1196 the actual session instance through `SessionManager::destroyReservationSession` (see also Section 6.1).

## 1197 8.2 ReservationTemplate structure

1198 In order to define the attributes associated with an advance reservation, the DRMAA application creates  
 1199 an `ReservationTemplate` instance and requests the fulfilment through the `ReservationSession` methods  
 1200 in the DRM system.

```
1201 struct ReservationTemplate {
1202     StringList attributeNames;
1203     string reservationName;
1204     AbsoluteTime startTime;
1205     AbsoluteTime endTime;
1206     TimeAmount duration;
1207     long minSlots;
1208     long maxSlots;
```

```

1209     OrderedStringList candidateMachines;
1210     long minPhysMemory;
1211     OperatingSystem machineOS;
1212     CpuArchitecture machineArch;
1213     Dictionary drmsSpecific;
1214 }

```

1215 Similar to the `JobTemplate` concept (see Section 7.4), there is a distinction between *mandatory* and *optional* attributes. Mandatory attributes MUST be supported by the implementation in the sense that they  
 1216 are evaluated in a `ReservationSession::requestReservation` call. Optional attributes MAY NOT be  
 1217 evaluated in a particular implementation, but MUST be provided as part of the `ReservationTemplate`  
 1218 structure in the implementation. If an optional attribute is not evaluated by the particular implementation,  
 1219 but has a value different to `UNSET`, the call to `ReservationSession::requestReservation` MUST fail with  
 1220 a `UnsupportedAttributeException`. DRMAA applications are expected to check for the availability of  
 1221 optional attributes by the `ReservationTemplate::attributeNames` list.

1223 Implementations MUST set all attribute values to `UNSET` on struct allocation. This ensures that both the  
 1224 DRMAA application and the library implementation can determine untouched attribute members. If not  
 1225 described differently in the following sections, all attributes SHOULD be allowed to have the `UNSET` value  
 1226 when `ReservationSession::requestReservation` is called.

A language binding specification SHOULD model the `ReservationTemplate` representation the same way as  
 the `JobTemplate` interface (see Section 7.4), and therefore MUST define the realization of implementation-  
 specific attributes, printing, and the initialization of attribute values.

Complete section needs group approval

1227

### 1228 8.2.1 attributeNames

1229 The `attributeNames` list of strings SHALL enumerate the names of the required and the supported optional  
 1230 reservation template attributes.

This doesn't make sense anymore, since reservation templates are now value types.

1231

This is especially intended for languages which do not provide an inherit notion of struct introspection and  
 therefore map template attribute access to getter / setter functions.

1232 The support for this attribute is mandatory.

### 1233 8.2.2 reservationName

1234 A human-readable reservation name. If this attribute is omitted then the name of the reservation SHALL be  
 1235 automatically defined by the implementation. The implementation MAY truncate any application-provided  
 1236 job name to an implementation-defined length.

1237 The support for this attribute is mandatory.

1238 8.2.3 `startTime / endTime / duration`1239 The time frame in which resources should be reserved. Table 4 explains the different possible parameter  
1240 combinations and their semantic.

<code>startTime</code>	<code>endTime</code>	<code>duration</code>	Description
UNSET	UNSET	UNSET	The implementation resp. DRM system is free to choose a time frame for the reservation.
Set	UNSET	UNSET	Invalid, SHALL leave to a <code>InvalidAttributeException</code> on the reservation attempt.
UNSET	Set	UNSET	Invalid, SHALL leave to a <code>InvalidAttributeException</code> on the reservation attempt.
Set	Set	UNSET	Perform reservation attempt to get resources in the specified time frame.
UNSET	UNSET	Set	Perform reservation attempt the get resources at least for the time amount given in <code>duration</code> .
Set	UNSET	Set	Implies <code>endTime = startTime + duration</code>
UNSET	Set	Set	Implies <code>startTime = endTime - duration</code>
Set	Set	Set	If <code>endTime - startTime</code> is larger than <code>duration</code> , perform a reservation attempt were the demanded <code>duration</code> is fulfilled at the earliest point in time after <code>startTime</code> , and without extending <code>endTime</code> . If <code>endTime - startTime</code> is smaller than <code>duration</code> , the reservation attempt SHALL leave to a <code>InvalidAttributeException</code> . If <code>endTime - startTime</code> and <code>duration</code> are equal, <code>duration</code> SHALL be ignored.

Table 4: Parameter combinations for the advance reservation time frame. If `duration` is not supported, it should be treated as UNSET.1241 The support for `startTime` and `endTime` is mandatory. The support for `duration` is optional.1242 8.2.4 `minSlots`1243 The minimum number of requested slots (see also Section 1.3). If the attribute value is UNSET, it should  
1244 default to 1.

1245 The support for this attribute is mandatory.

1246 8.2.5 `maxSlots`1247 The maximum number of requested slots (see also Section 1.3). If this attribute is not specified, it should  
1248 default to the value of `minSlots`.

1249 The support for this attribute is mandatory.

1250 8.2.6 `candidateMachines`1251 Requests that the reservation must be created on any subset of the given list of machines. If this attribute  
1252 is not specified, it should default to the result of `MonitoringSession::getAllMachines` (see Section 9.1).

1253 The support for this attribute is optional.

## 1254 8.2.7 minPhysMemory

1255 Requests that the reservation must be created with machines that have at least the given amount of physical  
 1256 memory in Kibibyte.

1257 The support for this attribute is optional.

## 1258 8.2.8 machineOS

1259 Requests that the reservation must be created with machines that have the given type of operating system,  
 1260 regardless of its version, with semantics as specified in Section 4.1.

1261 The support for this attribute is optional.

1262 (See footnote)<sup>29</sup>

## 1263 8.2.9 machineArch

1264 Requests that the reservation must be created with machines that have the given instruction set architecture,  
 1265 with semantics as specified in Section 4.2.

1266 The support for this attribute is optional.

## 1267 8.2.10 drmsSpecific

1268 This dictionary attribute allows the application to pass DRMS-specific native options for the advance reservation  
 1269 as key-value pairs. The interpretation of keys and values in this dictionary is implementation-specific,  
 1270 implementations MAY even ignore them. Valid key strings should be documented by the implementation.  
 1271 The order of precedence for the drmsSpecific attribute value and other, maybe conflicting, attribute values  
 1272 is implementation-specific. Implementations MAY decide to overrule reservation template settings with the  
 1273 ones defined by the drmsSpecific attribute.

1274 The support for this attribute is mandatory.

## 1275 8.3 Reservation Interface

1276 The **Reservation** interface represents attributes and methods available for an advance reservation successfully  
 1277 created in the DRM system.

```
1278 interface Reservation {
1279     readonly attribute string reservationId;
1280     readonly attribute ReservationSession session;
1281     readonly attribute ReservationTemplate reservationTemplate;
1282     OrderedStringList reservedMachines;
1283     AbsoluteTime reservedStartTime;
1284     AbsoluteTime reservedEndTime;
1285     void terminate();
1286 };
```

---

<sup>29</sup> Requesting a particular operating system version is not supported by the majority of DRM systems (conf call Jul 28th 2010)

1287 (See footnote)<sup>30</sup>

1288 **8.3.1 reservationId**

1289 The **reservationId** is an opaque string identifier for the advance reservation. If the DRM system has  
 1290 identifiers for advance reservations, this attribute SHOULD provide the according stringified value. If not,  
 1291 the DRMAA implementation MUST generate value this is unique in time and extend of the DRM system.

Relationship  
to  
ReservationT  
?

1292

1293 **8.3.2 session**

1294 This attribute references the **ReservationSession** which was used to create the advance reservation instance.

1295 **8.3.3 reservationTemplate**

1296 This attribute provides a reference to a **ReservationTemplate** instance that has equal values to the one  
 1297 that was used for the advance reservation creating this **Reservation** instance.

Could that  
be UNSET ?

1298 **8.3.4 reservedMachines**

1299

1300 This attribute describes the set of machines which was reserved under the conditions described in the  
 1301 according reservation template.

1302 **8.3.5 reservedStartTime**

1303

1304 This attribute describes the start time for the reservation described by this instance.

Could that  
be UNSET ?

1305 **8.3.6 reservedEndTime**

1306

1307 This attribute describes the end time for the reservation described by this instance.

Could that  
be UNSET ?

1308 **8.3.7 terminate**

1309 This method terminates the advance reservation in the DRM system represented by this **Reservation**  
 1310 instance.

Needs ad-  
ditional ex-  
planation of  
expected be-  
havior

1311 **9 Monitoring the DRM System**

1312 The DRMAA monitoring facility supports four basic units of monitoring:

1313 • Properties of the DRM system as a whole (e.g. DRM system version number) that are independent  
 1314 from the particular session resp. contact string,

<sup>30</sup> The reason for not having a separate **ReservationInfo** struct is that there are only three relevant attributes for this structure, and that all of them have static semantics. There is, therefore, no need for refetching reservation information several times, which is the case with **JobInfo**. Because of this, the according information can be a part of the **Reservation** interface itself.

- Properties of the DRM system that depend on the current contact string (e.g. list of machines in the currently accessed Grid Engine cell)
- Properties of individual queues known from a `getAllQueues` call
- Properties of individual machines available with the current contact string (e.g. amount of physical memory in a chosen machine)

The `MonitoringSession` interface in DRMAA supports the monitoring of execution resources in the DRM system. This is distinct from the monitoring of jobs running in the DRM system, which is covered by the `JobSession` resp. `Job` interface.

## 9.1 MonitoringSession Interface

The `MonitoringSession` interface represents a set of stateless methods for fetching information about the DRM system and the DRMAA implementation itself. It MAY be used to implement DRM system monitoring tools like `qstat`.

```
interface MonitoringSession {
    readonly attribute Version drmsVersion;
    ReservationList getAllReservations();
    JobList getAllJobs(in JobInfo filter);
    QueueList getAllQueues(in StringList names);
    MachineList getAllMachines(in StringList names);
    readonly attribute StringList drmsJobCategoryNames;
};
```

All returned data SHOULD be related to the current user running the DRMAA-based application. For example, the `getAllQueues` function MAY be reduced to only denote queues that are usable or generally accessible for the DRMAA application and user performing the query.

Because no guarantee can be made as to future accessibility, and because of cases where list reduction may demand excessive overhead in the DRMAA implementation, an unreduced or partially reduced result MAY be returned on all methods returning lists. The behavior of the DRMAA implementation in this regard should be clearly documented. In all cases, the list items MUST all be valid input for job submission or advance reservation through the DRMAA API.

### 9.1.1 drmsVersion

This attribute provides the DRM-system specific version information. While the DRM system type is available from the `SessionManager::drmsName` attribute (see Section 6.1), this attribute provides the according version of the product. Applications are expected to use the information about the general DRM system type for accessing product-specific features, e.g. with the `JobTemplate::drmsSpecific` dictionary. Applications are not expected to make decisions based on versioning information from this attribute - instead, the value should only be utilized for informative output to the end user.

### 9.1.2 getAllReservations

This method returns the list of all DRMS advance reservations accessible for the user running the DRMAA-based application. In contrast to a `ReservationSession::getReservations` call, this method SHOULD return also reservations that were created outside of DRMAA (e.g. through command-line tools) by this user.

1354 The returned list MAY also contain reservations that were created by other users if the security policies of  
 1355 the DRM system allow such global visibility. The DRM system or the DRMAA implementation is at liberty,  
 1356 however, to restrict the set of returned reservations based on site or system policies, such as security settings  
 1357 or scheduler load restrictions.

1358 This method SHALL raise an `UnsupportedOperationException` if advance reservation is not supported by  
 1359 the implementation.

1360 **9.1.3 getAllJobs**

1361 This method returns the list of all DRMS jobs visible to the user running the DRMAA-based application. In  
 1362 contrast to a `JobSession::getJobs` call, this method SHOULD also return jobs that were submitted outside  
 1363 of DRMAA (e.g. through command-line tools) by this user. The returned list MAY also contain jobs that  
 1364 were submitted by other users if the security policies of the DRM system allow such global visibility. The  
 1365 DRM system or the DRMAA implementation is at liberty, however, to restrict the set of returned jobs based  
 1366 on site or system policies, such as security settings or scheduler load restrictions.

1367 Querying the DRM system for all jobs might result in returning an excessive number of `Job` objects. Impli-  
 1368 cations to the library implementation are out of scope for this specification.

1369 The method supports a `filter` argument for fetching only a subset of the job information available. Both  
 1370 the return value semantics and the filter semantics SHOULD be similar to the ones described for the  
 1371 `JobSession::getJobs` method (see Section 7.2).

Language bindings SHOULD NOT try to solve the scalability issues by replacing the sequence type of the  
 return value with some iterator-alike solution. This approach would break the basic snapshot semantic  
 intended for this method.

1372 (See footnote)<sup>31</sup>

1373 **9.1.4 getAllQueues**

1374 This method returns a list of queues available for job submission in the DRM system. All `Queue` instances  
 1375 in this list SHOULD be (based on their `name` attribute) a valid input for the `JobTemplate::queueName`  
 1376 attribute (see Section 7.4). The result can be an empty list or might be incomplete, based on queue, host,  
 1377 or system policies. It might also contain queues that are not accessible for the user (because of queue  
 1378 configuration limits) at job submission time.

1379 The `names` parameter supports restricting the result to `Queue` instances that have one of the names given in  
 1380 the argument. If the `names` parameter value is `UNSET`, all `Queue` instances should be returned.

1381 **9.1.5 getAllMachines**

1382 This method returns the list of machines available in the DRM system as execution host. The returned list  
 1383 might be empty or incomplete based on machine or system policies. The returned list might also contain  
 1384 machines that are not accessible by the user, e.g. because of host configuration limits.

1385 The `names` parameter supports restricting the result to `Machine` instances that have one of the names given  
 1386 in the argument. If the `names` parameter value is `UNSET`, all `Machine` instances should be returned.

<sup>31</sup> The non-argumentation about the scalability problem was the final result of a clarification attempt. We hand this one over to the implementors. (conf call Jul 14th 2010)

## 1387 9.1.6 drmsJobCategoryNames

1388 This method provides the list of of valid job category names which can be used for the `jobCategory` attribute  
 1389 in a job template. The semantics are described in Section 7.4.8.

## 1390 10 Annex A: Complete DRMAA IDL Specification

1391 The following text shows the complete IDL specification for the DRMAAv2 application programming interface.  
 1392 The ordering of IDL constructs here has no normative meaning, but ensures the correct compilation  
 1393 with a standard CORBA IDL compiler for syntactical correctness checks. This demands only some additional  
 1394 forward declarations to resolve circular dependencies.

```
1395 module DRMAA2 {

1396   enum JobState {
1397     UNDETERMINED, QUEUED, QUEUED_HELD, RUNNING, SUSPENDED, REQUEUED,
1398     REQUEUED_HELD, DONE, FAILED};

1399   enum OperatingSystem {
1400     HPUX, LINUX, IRIX, TRUE64, MACOS, SUNOS, WIN, WINNT, AIX, UNIXWARE,
1401     BSD, OTHER_OS};

1402   enum CpuArchitecture {
1403     ALPHA, ARM, CELL, PARISC, X86, X64, IA64, MIPS, PPC, PPC64,
1404     SPARC, SPARC64, OTHER_CPU};

1405   enum ResourceLimitType {
1406     CORE_FILE_SIZE, CPU_TIME, DATA_SEG_SIZE, FILE_SIZE, OPEN_FILES,
1407     STACK_SIZE, VIRTUAL_MEMORY, WALLCLOCK_TIME };

1408   enum JobTemplatePlaceholder {
1409     HOME_DIRECTORY, WORKING_DIRECTORY, HOST_NAME, USER_NAME, PARAMETRIC_INDEX };

1410   enum DrmaaEvent {
1411     NEW_STATE, MIGRATED, ATTRIBUTE_CHANGE
1412   };

1413   typedef sequence<string> OrderedListString;
1414   typedef sequence<string> StringList;
1415   typedef sequence<Job> JobList;
1416   typedef sequence<Queue> QueueList;
1417   typedef sequence<Machine> MachineList;
1418   typedef sequence<Reservation> ReservationList;
1419   typedef sequence< sequence<string,2> > Dictionary;
1420   typedef string AbsoluteTime;
1421   typedef long long TimeAmount;
1422   native ZERO_TIME;
1423   native INFINITE_TIME;
1424   native UNSET;
```

```
1425     struct JobInfo {
1426         string jobId;
1427         Dictionary resourceUsage;
1428         long exitStatus;
1429         string terminatingSignal;
1430         string annotation;
1431         JobState jobState;
1432         any jobSubState;
1433         OrderedStringList allocatedMachines;
1434         string submissionMachine;
1435         string jobOwner;
1436         string queueName;
1437         TimeAmount wallclockTime;
1438         long cpuTime;
1439         AbsoluteTime submissionTime;
1440         AbsoluteTime dispatchTime;
1441         AbsoluteTime finishTime;};
1442
1443     struct JobTemplate {
1444         StringList attributeNames;
1445         string remoteCommand;
1446         OrderedStringList args;
1447         boolean submitAsHold;
1448         boolean rerunnable;
1449         Dictionary jobEnvironment;
1450         string workingDirectory;
1451         string jobCategory;
1452         StringList email;
1453         boolean emailOnStarted;
1454         boolean emailOnTerminated;
1455         string jobName;
1456         string inputPath;
1457         string outputPath;
1458         string errorPath;
1459         boolean joinFiles;
1460         string reservationId;
1461         string queueName;
1462         long minSlots;
1463         long maxSlots;
1464         long priority;
1465         OrderedStringList candidateMachines;
1466         long minPhysMemory;
1467         OperatingSystem machineOS;
1468         CpuArchitecture machineArch;
1469         AbsoluteTime startTime;
1470         Dictionary drmsSpecific;
1471         AbsoluteTime deadlineTime;
1472         Dictionary stageInFiles;
```

```
1472     Dictionary stageOutFiles;
1473     Dictionary softResourceLimits;
1474     Dictionary hardResourceLimits;
1475     string accountingId;
1476 };
1477
1478 struct ReservationTemplate {
1479     StringList attributeNames;
1480     string reservationName;
1481     AbsoluteTime startTime;
1482     AbsoluteTime endTime;
1483     TimeAmount duration;
1484     long minSlots;
1485     long maxSlots;
1486     OrderedStringList candidateMachines;
1487     long minPhysMemory;
1488     OperatingSystem machineOS;
1489     CpuArchitecture machineArch;
1490     Dictionary drmsSpecific;
1491 };
1492
1493 struct DrmaaNotification {
1494     DrmaaEvent event;
1495     Job job;
1496     JobState jobState;
1497 };
1498
1499 struct Queue {
1500     string name;
1501     TimeAmount maxWallclockTime;
1502 };
1503
1504 struct Version {
1505     string major;
1506     string minor;
1507 };
1508
1509 struct Machine {
1510     string name;
1511     long sockets;
1512     long coresPerSocket;
1513     long threadsPerCore;
1514     double load;
1515     long physMemory;
1516     long virtMemory;
1517     OperatingSystem machineOS;
1518     Version machineOSVersion;
1519     CpuArchitecture machineArch;
1520 };
1521
```

```
1516     exception AuthorizationException {string message;};
1517     exception DefaultContactStringException {string message;};
1518     exception DeniedByDrmException {string message;};
1519     exception DrmCommunicationException {string message;};
1520     exception TryLaterException {string message;};
1521     exception SessionManagementException {string message;};
1522     exception TimeoutException {string message;};
1523     exception InternalException {string message;};
1524     exception InvalidArgumentException {string message;};
1525     exception InvalidSessionException {string message;};
1526     exception InvalidStateException {string message;};
1527     exception OutOfMemoryException {string message;};
1528     exception UnsupportedAttributeException {string message;};
1529     exception UnsupportedOperationException {string message;};

1530     interface DrmaaCallback {
1531         void notify(in DrmaaNotification notification);
1532     };

1533     interface ReservationSession {
1534         readonly attribute string contact;
1535         readonly attribute string sessionName;
1536         Reservation getReservation(in string reservationId);
1537         Reservation requestReservation(in ReservationTemplate reservationTemplate);
1538         ReservationList getReservations();
1539     };

1540     interface Reservation {
1541         readonly attribute string reservationId;
1542         readonly attribute ReservationSession session;
1543         readonly attribute ReservationTemplate reservationTemplate;
1544         OrderedStringList reservedMachines;
1545         AbsoluteTime reservedStartTime;
1546         AbsoluteTime reservedEndTime;
1547         void terminate();
1548     };

1549     interface JobArray {
1550         readonly attribute string jobArrayId;
1551         readonly attribute JobList jobs;
1552         readonly attribute JobSession session;
1553         readonly attribute JobTemplate jobTemplate;
1554         void suspend();
1555         void resume();
1556         void hold();
1557         void release();
1558         void terminate();
1559     };

```

```

1560     interface JobSession {
1561         readonly attribute string contact;
1562         readonly attribute string sessionId;
1563         readonly attribute boolean notificationSupported;
1564         JobList getJobs(in JobInfo filter);
1565         Job runJob(in JobTemplate jobTemplate);
1566         JobArray runBulkJobs(
1567             in JobTemplate jobTemplate,
1568             in long beginIndex,
1569             in long endIndex,
1570             in long step);
1571         Job waitAnyStarted(in JobList jobs, in TimeAmount timeout);
1572         Job waitAnyTerminated(in JobList jobs, in TimeAmount timeout);
1573         void registerEventNotification(in DrmaaCallback callback);
1574     };
1575
1576     interface Job {
1577         readonly attribute string jobId;
1578         readonly attribute JobSession session;
1579         readonly attribute JobTemplate jobTemplate;
1580         void suspend();
1581         void resume();
1582         void hold();
1583         void release();
1584         void terminate();
1585         JobState getState(out any jobSubState);
1586         JobInfo getInfo();
1587         Job waitStarted(in TimeAmount timeout);
1588         Job waitTerminated(in TimeAmount timeout);
1589     };
1590
1591     interface MonitoringSession {
1592         readonly attribute Version drmsVersion;
1593         ReservationList getAllReservations();
1594         JobList getAllJobs(in JobInfo filter);
1595         QueueList getAllQueues(in StringList names);
1596         MachineList getAllMachines(in StringList names);
1597         readonly attribute StringList drmsJobCategoryNames;
1598     };
1599
1600     interface SessionManager{
1601         readonly attribute string drmsName;
1602         readonly attribute Version drmaaVersion;
1603         readonly attribute boolean reservationSupported;
1604         JobSession createJobSession(in string sessionId,
1605                                     in string contactString);
1606         ReservationSession createReservationSession(in string sessionId,
1607                                         in string contactString);

```

```
1605     MonitoringSession createMonitoringSession (in string contactString);  
1606     JobSession openJobSession(in string sessionId);  
1607     ReservationSession openReservationSession(in string sessionId);  
1608     void closeJobSession(in JobSession s);  
1609     void closeReservationSession(in ReservationSession s);  
1610     void closeMonitoringSession(in MonitoringSession s);  
1611     void destroyJobSession(in string sessionId);  
1612     void destroyReservationSession(in string sessionId);  
1613     StringList getJobSessions();  
1614     StringList getReservationSessions();  
1615 };  
1616 };
```

## 1617 11 Security Considerations

1618 The DRMAA API does not specifically assume the existence of a particular security infrastructure in the  
1619 DRM system. The scheduling scenario described herein presumes that security is handled at the point of job  
1620 authorization/execution on a particular resource. It is assumed that credentials owned by the application  
1621 using the API are in effect for the DRMAA implementation too.

1622 It is conceivable an authorized but malicious user could use a DRMAA implementation or a DRMAA enabled  
1623 application to saturate a DRM system with a flood of requests. Unfortunately for the DRM system this  
1624 case is not distinguishable from the case of an authorized good-natured user that has many jobs to be  
1625 processed. For temporary load defense, implementations **SHOULD** utilize the `TryLaterException`. In case  
1626 of permanent issues, the implementation **SHOULD** raise the `DeniedByDrmException`.

1627 DRMAA implementers should guard against buffer overflows that could be exploited through DRMAA  
1628 enabled interactive applications or web portals. Implementations of the DRMAA API will most likely  
1629 require a network to coordinate subordinate DRMS, however the API makes no assumptions about the  
1630 security posture provided the networking environment. Therefore, application developers should further  
1631 consider the security implications of "on-the-wire" communications.

1632 For environments that allow remote or protocol based DRMAA clients, the implementation **SHOULD** offer  
1633 support for secure transport layers to prevent man in the middle attacks.

## 1634 12 Contributors

1635 **Roger Brobst**  
1636 Cadence Design Systems, Inc.  
1637 555 River Oaks Parkway  
1638 San Jose, CA 95134  
1639 Email: [rbrobst@cadence.com](mailto:rbrobst@cadence.com)  
1640

1641 **Daniel Gruber**  
1642 Univa  
1643

1644 **Mariusz Mamonski**

1645

1646 **Daniel Templeton (Corresponding Author)**

1647 Cloudera

1648

1649 **Peter Tröger (Corresponding Author)**

1650 Hasso-Plattner-Institute at University of Potsdam

1651 Prof.-Dr.-Helmert-Str. 2-3

1652 14482 Potsdam, Germany

1653 Email: peter@troeger.eu

1654

1655

1656 We are grateful to numerous colleagues for support and discussions on the topics covered in this document,  
1657 in particular (in alphabetical order, with apologies to anybody we have missed):

1658 Guillaume Alleon, Ali Anjomshoaa, Ed Baskerville, Harald Böhme, Nadav Brandes, Matthieu Cargnelli, Karl  
1659 Czajkowski, Piotr Domagalski, Fritz Ferstl, Paul Foley, Nicholas Geib, Becky Gietzel, Alleon Guillaume,  
1660 Andreas Haas, Tim Harsch, Greg Hewgill, Rayson Ho, Eduardo Huedo, Dieter Kranzmüller, Krzysztof  
1661 Kurowski, Peter G. Lane, Miron Livny, Ignacio M. Llorente, Martin v. Löwis, Andre Merzky, Ruben S.  
1662 Montero, Greg Newby, Steven Newhouse, Michael Primeaux, Greg Quinn, Hrabri L. Rajic, Martin Sarachu,  
1663 Jennifer Schopf, Enrico Sirola, Chris Smith, Ancor Gonzalez Sosa, Douglas Thain, John Tollefsrud, Jose R.  
1664 Valverde, and Peter Zhu.

Add missing contact details

## 1665 13 Intellectual Property Statement

1666 The OGF takes no position regarding the validity or scope of any intellectual property or other rights that  
1667 might be claimed to pertain to the implementation or use of the technology described in this document or the  
1668 extent to which any license under such rights might or might not be available; neither does it represent that  
1669 it has made any effort to identify any such rights. Copies of claims of rights made available for publication  
1670 and any assurances of licenses to be made available, or the result of an attempt made to obtain a general  
1671 license or permission for the use of such proprietary rights by implementers or users of this specification can  
1672 be obtained from the OGF Secretariat.

1673 The OGF invites any interested party to bring to its attention any copyrights, patents or patent applications,  
1674 or other proprietary rights which may cover technology that may be required to practice this recommendation.  
1675 Please address the information to the OGF Executive Director.

## 1676 14 Disclaimer

1677 This document and the information contained herein is provided on an "As Is" basis and the OGF disclaims  
1678 all warranties, express or implied, including but not limited to any warranty that the use of the information  
1679 herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular  
1680 purpose.

## 1681 15 Full Copyright Notice

1682 Copyright © Open Grid Forum (2005-2011). Some Rights Reserved.  
1683 This document and translations of it may be copied and furnished to others, and derivative works that  
1684 comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and  
1685 distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice  
1686 and this paragraph are included on all such copies and derivative works. However, this document itself  
1687 may not be modified in any way, such as by removing the copyright notice or references to the OGF or  
1688 other organizations, except as needed for the purpose of developing Grid Recommendations in which case  
1689 the procedures for copyrights defined in the OGF Document process must be followed, or as required to  
1690 translate it into languages other than English.  
1691 The limited permissions granted above are perpetual and will not be revoked by the OGF or its successors  
1692 or assignees.

## 1693 16 References

- 1695 [1] Scott Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119 (Best Current  
1696 Practice), March 1997. URL <http://tools.ietf.org/html/rfc2119>.
- 1697 [2] I. Foster, A. Grimshaw, P. Lane, W. Lee, M. Morgan, S. Newhouse, S. Pickles, D. Pulsipher, C. Smith,  
1698 and M. Theimer. OGSA Basic Execution Service v1.0 (GFD-R.108), nov 2008.
- 1699 [3] Tom Goodale, Shantenu Jha, Hartmut Kaiser, Thilo Kielmann, Pascal Kleijer, Andre Merzky, John  
1700 Shalf, and Christopher Smith. A Simple API for Grid Applications (SAGA) Version 1.1 (GFD-R-P.90),  
1701 jan 2008.
- 1702 [4] Object Management Group. Common Object Request Broker Architecture (CORBA) Specification,  
1703 Version 3.1. <http://www.omg.org/spec/CORBA/3.1/Interfaces/PDF>, jan 2008.
- 1704 [5] The IEEE and The Open Group. The Open Group Base Specifications Issue 6 IEEE Std 1003.1.  
1705 <http://www.opengroup.org/onlinepubs/000095399/utilities/ulimit.html>.
- 1706 [6] Distributed Management Task Force (DMTF) Inc. CIM System Model White Paper CIM Version 2.7,  
1707 jun 2003.
- 1708 [7] Hrabri Rajic, Roger Brobst, Waiman Chan, Fritz Ferstl, Jeff Gardiner, Andreas Haas, Bill Nitzberg,  
1709 Daniel Templeton, John Tollefsrud, and Peter Tröger. Distributed Resource Management Application  
1710 API Specification 1.0 (GFD-R.022), aug 2007.
- 1711 [8] Hrabri Rajic, Roger Brobst, Waiman Chan, Fritz Ferstl, Jeff Gardiner, Andreas Haas, Bill Nitzberg,  
1712 Daniel Templeton, John Tollefsrud, and Peter Tröger. Distributed Resource Management Application  
1713 API Specification 1.0 (GWD-R.133), jun 2008.
- 1714 [9] Eric Rescorla, Brian Korver, and Internet Architectures Board. Guidelines for Writing RFC Text on  
1715 Security Considerations. RFC 3552 (Best Current Practice), July 2003. URL <http://tools.ietf.org/html/rfc3552>.
- 1717 [10] Peter Tröger, Daniel Templeton, Roger Brobst, Andreas Haas, and Hrabri Rajic. Distributed Resource  
1718 Management Application API 1.0 - IDL Specification (GFD-R-P.130), apr 2008.

1719 [11] Peter Tröger, Hrabri Rajic, Andreas Haas, and Piotr Domagalski. Standardised job submission and  
1720 control in cluster and grid environments. *International Journal of Grid and Utility Computing*, 1:  
1721 134–145, dec 2009. doi: {<http://dx.doi.org/10.1504/IJGUC.2009.022029>}.