GWD-R
DRMAA-WG
drmaa-wg@ogf.org

Peter Tröger, Hasso-Plattner-Institute
(Corresponding Author)
Roger Brobst, Cadence Design Systems
Daniel Gruber, Univa
Mariusz Mamoński, PSNC
Daniel Templeton, Cloudera
July 2011

Distributed Resource Management Application API Version 2 (DRMAA) - Draft 8

Status of This Document

Group Working Draft Recommendation (GWD-R)

(See footnote)[1]

Obsoletes

This document obsoletes GFD-R.022 [8], GFD-R-P.130 [10], and GWD-R.133 [9].

Document Change History

Date	Notes

Copyright Notice

Trademark

Abstract

This document describes the *Distributed Resource Management Application API Version 2 (DRMAA)*. It defines a generalized API to *Distributed Resource Management (DRM)* systems in order to facilitate the development of portable application programs and high-level libraries.

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

[1] 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 in only intended for internal working group discussions.

Notational Conventions

In this document, IDL language elements and definitions are represented in a `fixed-width` font.

The key words "MUST" "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [2].

Memory quantities are expressed in *kilobyte (KB)*. 1 kilobyte equals 1024 bytes.

Parts of this document are only normative for DRMAA language binding specifications. These sections are graphically marked as shaded box.

(See footnote)[2] .

[2]The usage of kikibyte as memory quantity unit, as well as the usage of bytes as in JSDL, was rejected by the group (conf call Apr. 13th 2011)

Contents

1 Introduction . . . 5
1.1 Basic concepts . . . 5
1.2 Slots and Queues . . . 6
1.3 Language Bindings . . . 6
1.4 Job Categories . . . 7
1.5 Multithreading . . . 8
2 Namespace . . . 8
3 Common Type Definitions . . . 8
4 Enumerations . . . 10
4.1 OperatingSystem enumeration . . . 10
4.2 CpuArchitecture enumeration . . . 11
4.3 ResourceLimitType enumeration . . . 12
4.4 JobTemplatePlaceholder enumeration . . . 13
4.5 DrmaaCapability . . . 14
5 Extensible Data Structures . . . 15
5.1 QueueInfo structure . . . 15
5.2 Version structure . . . 16
5.3 MachineInfo structure . . . 16
5.4 SlotInfo structure . . . 18
5.5 JobInfo structure . . . 18
5.6 ReservationInfo structure . . . 22
5.7 JobTemplate structure . . . 23
5.8 ReservationTemplate structure . . . 31
5.9 DrmaaReflective Interface . . . 35
6 Common Exceptions . . . 36
7 The DRMAA Session Concept . . . 38
7.1 SessionManager Interface . . . 38
8 Working with Jobs . . . 42
8.1 The DRMAA State Model . . . 42
8.2 JobSession Interface . . . 44
8.3 DrmaaCallback Interface . . . 47
8.4 Job Interface . . . 48
8.5 JobArray Interface . . . 50
8.6 The DRMAA_INDEX_VAR environment variable . . . 52
9 Working with Advance Reservation . . . 52
9.1 ReservationSession Interface . . . 53
9.2 Reservation Interface . . . 54
10 Monitoring the DRM System . . . 55
10.1 MonitoringSession Interface . . . 55
11 Complete DRMAA IDL Specification . . . 57
12 Security Considerations . . . 62
13 Contributors . . . 63
14 Intellectual Property Statement . . . 64

15 Disclaimer 64
16 Full Copyright Notice 65
17 References 65

1 Introduction

The *Distributed Resource Management Application API Version 2 (DRMAA)* specification defines an interface for tightly coupled, but still portable access by abstracting the fundamental functions available in the majority of DRM systems. The scope is limited to job submission, job control, reservation management, and retrieval of job and machine monitoring information.

This document acts as root specification for the abstract API concepts and the behavioral rules of a DRMAA-compliant implementation. The programming language representation of the API is defined by a separate *language binding specification*.

There are other relevant OGF standards in the area of job submission and monitoring. An in-depth comparison and positioning of the obsoleted first version of the DRMAA [9] specification was provided by another publication [11]. This document was created in close collaboration with the OGF SAGA and the OGF OCCI working group.

First-time readers are recommended to complete reading this section. After that, they should jump to Section 7 for getting an overview of the supported functionality in DRMAA. Section 11 can be always consulted in parallel for a global overview on the API layout.

1.1 Basic concepts

The DRMAA specification is based on the following stakeholders:

- *Distributed resource management system / DRM system / DRMS*: Any system that supports the concept of distributing computational tasks on execution resources through the help of a central scheduling entity. Examples are multi-processor systems controlled by a operating system scheduler, cluster systems with multiple machines controlled by a central scheduler software, grid systems, or cloud systems with a job concept.
- *(DRMAA) implementation / (DRMAA) library*: The implementation of a DRMAA language binding specification, with the functional behavior as described in this document. The resulting artifact is expected to target one DRM system.
- *(DRMAA-based) application*: Software that utilizes the DRMAA implementation for gaining access to one or multiple DRM systems in a standardized way.
- *Submission host*: A resource in the DRM system that runs the DRMAA-based application. A submission host MAY also be able to act as execution host.
- *Execution host*: A resource in the DRM system that can run a submitted job.
- *Job*: A computational activity submitted by the DRMAA-based application to a DRM system, with the help of the DRMAA implementation. A job is expected to run as one or many operating system processes on one or many execution hosts.

Table 1 defines the conceptual mapping of DRMAA to the GLUE 2.0 Information model [1]. Since the DRMAA API design is derived from existing DRM system functionality and terminology, not all GLUE concepts are applicable here, such as the expression of ID's as URI's, the SI metric model, the representation of date information, or the endpoint concept.

DRMAA	Reference	GLUE 2.0	Reference [1]
DRM system	Section 1.1	Manager	Section 5.9
Execution host	Section 1.1	ExecutionEnvironment + ComputingManager	Section 6.4 / 6.6
Socket	Section 5.3.3	Physical CPU	Section 6
Core	Section 5.3.4	Logical CPU	Section 6
Job	Section 1.1	ComputingActivity	Section 6.9
Job category	Section 1.4	ApplicationEnvironment	Section 6.7
`UNSET` value	Section 1.3	Placeholder values for unknown data	Appendix A

Table 1: Mapping of DRMAA concepts to GLUE 2.0

1.2 Slots and Queues

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

(See footnote)[3]

1.3 Language Bindings

The interface semantics are described with the *OMG Interface Definition Language (IDL)* [5] syntax. Based on this language-agnostic specification, *language binding* standards have to be designed that map the abstract concepts into a library interface for a particular programming language (e.g. C, Java, Python). While this document has the responsibility to ensure consistent API semantics for all possible DRMAA implementations, the language binding has the responsibility to ensure source-code portability for DRMAA applications on different DRM systems.

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

A language binding specification derived from this document MUST define a mapping between the IDL

[3] 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). Slots have a meaning in GLUE, but we intentionally stick with the opaque concept approach.

constructs and the constructs of its targeted programming language. The focus MUST be on source code portability for the DRMAA-based application in the particular language.

A language binding SHOULD NOT rely completely on the OMG IDL language mapping standards available for many programming languages, since they have a significant overhead of CORBA-related mapping rules that are not relevant here. The language binding MUST use its initially defined type system mapping in a consistent manner for the complete 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. Language bindings MAY map the DRMAA IDL interfaces to classes.

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. For non-scalar attributes, the language binding MUST specify a consistent access strategy for all these attributes, for example pass-by-value or pass-by-reference.

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

A language binding MUST define a way to declare an invalid value (`UNSET`). In case, a definition per data type needs to be provided. Evaluating an `UNSET` boolean value MUST result in a negative result, e.g. for `JobTemplate::emailOnStarted`. Invalid strings MAY be modelled according to the GLUE 2.0 scheme [1], were an `UNSET` string contains the value "UNDEFINEDVALUE". Invalid integers MAY be also modelled according to GLUE 2.0 scheme, were an `UNSET` integer is expressed as "all nines".

(See footnote)[4]

1.4 Job Categories

DRMAA facilitates writing DRM-enabled applications even though the deployment properties, in particular the configuration of the DRMS, cannot be known in advance. This is realized by a set of standardized attributes that can be specified for job submission or advance reservation.

One of these attributes is the job category, which allows to give an indication about the nature of the job at execution time. Examples are parallel MPI jobs, OpenMP jobs, jobs targeting specific accelerator hardware, or jobs demanding managed runtime environments (e.g. Java).

Job categories typically map to site-specific reservation or submission options. Each category expresses a particular type of job execution that demands site-specific configuration such as path settings, environment variables, or application starters. This mapping to site-specific conditions SHOULD take place at submission time of the job or advance reservation.

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

[4] The concept of a UNSET value was decided on a conf call (Aug 25th 2010). Boolean in C can 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.

http://www.drmaa.org/jobcategories/

Implementations SHOULD use these recommended names. In case the name is not taken from this list, it should be self-explanatory for the user to make her understand the implications on job execution.

Implementations MAY provide a library configuration facility, which allows a site administrator to link job category names with specific product- and site-specific configuration options.

The order of precedence between the job category and other attributes is implementation-specific. It is RECOMMENDED to overrule explicit job / reservation settings with the implicit settings resulting from a conflicting job category.

For bulk job submissions, the category is expected to be valid for each of the jobs created.

(See footnote)[5]

1.5 Multithreading

High-level APIs such as SAGA [4] are expected to utilize DRMAA for their own asynchronous operation, based on the assumption that re-entrancy is supported by the DRMAA implementation. For this reason, implementations SHOULD ensure the proper functioning of the library in case of re-entrant library calls without any explicit synchronization among the application threads. DRMAA implementers should document their level of thread safety.

2 Namespace

The DRMAA interfaces and structures are encapsulated by a naming scope, to avoid conflicts with other APIs used in the same application.

```
module DRMAA2 {
```

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

(See footnote)[6]

3 Common Type Definitions

The abstract DRMAA specification defines some custom types to express special value semantics not available in original IDL:

```
  typedef sequence<string> OrderedStringList;
  typedef sequence<string> StringList;
  typedef sequence<Job> JobList;
  typedef sequence<QueueInfo> QueueInfoList;
```

[5] There was a discussion on supporting the specification of multiple categories at the same time. Since this would put more burden on the implementation in terms of conflict resolving, we avoided that intentionally. This allows to map categories simply to some additional job submission command line arguments, similar to the old nativeSpecification thing.

[6] Comparison to DRMAA v1.0: The IDL module name was change to DRMAA2, in order to intentionally break backward compatibility of the interface.

```
typedef sequence<MachineInfo> MachineInfoList;
typedef sequence<SlotInfo> OrderedSlotInfoList;
typedef sequence<Reservation> ReservationList;
typedef sequence< sequence<string,2> > Dictionary;
typedef string AbsoluteTime;
typedef long long TimeAmount;
native ZERO_TIME;
native INFINITE_TIME;
native NOW;
```

OrderedStringList: An unbounded list of strings, which supports element insertion, element deletion, and iteration over elements while keeping an element order.

StringList: An unbounded list of strings, without any demand on element order.

JobList: An unbounded list of `Job` instances, without any demand on element order.

QueueInfoList: An unbounded list of `QueueInfo` instances, without any demand on element order.

MachineInfoList: An unbounded list of `MachineInfo` instances, without any demand on element order.

OrderedSlotInfoList: An unbounded list of `SlotInfo` instances, which supports element insertion, element deletion, and iteration over elements while keeping an element order.

ReservationList: An unbounded list of `Reservation` instances, without any demand on element order.

Dictionary: An unbounded dictionary type for storing key-value pairs, without any demand on element order.

AbsoluteTime: Expression of a point in time, with a resolution at least to seconds.

TimeAmount: Expression of an amount of time, with a resolution at least to seconds.

ZERO_TIME: A constant value of type `TimeAmount` that expresses a zero amount of time.

INFINITE_TIME: A constant value of type `TimeAmount` that expresses an infinite amount of time.

NOW: A constant value of type `AbsoluteTime` that represents the point in time at which it is evaluated by some function.

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 mechanism for obtaining the RFC822 string representation from a given `AbsoluteTime` or `TimeAmount` instance.

(See footnote)[7]

[7] 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)

4 Enumerations

Some methods and attributes in DRMAA expect enumeration constants as input. The specified enumerations SHOULD NOT be extended by an implementation or language binding.

Language bindings SHOULD define numerical values for all enumeration members.

(See footnote)[8]

4.1 OperatingSystem enumeration

DRMAA supports the identification or demanding of an operating system installation on execution hosts. The enumeration defines a set of standardized identifiers for operating system types. The list is a shortened version of the according CIM Schema [7]. It includes only operating systems that are supported by the majority of DRM systems available at the time of writing:

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

AIX: AIX Unix by IBM.

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

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

HPUX: HP-UX Unix by Hewlett-Packard.

IRIX: The IRIX operating system by SGI.

MACOS: The MAC OS X operating system by Apple.

SUNOS: SunOS or Solaris operating system by Sun / Oracle.

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

UNIXWARE: UnixWare system by SCO group.

WIN: Windows 95, Windows 98, Windows ME.

WINNT: Microsoft Windows operating systems based on the NT kernel

OTHER_OS: An operating system type not specified in this list.

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

The operating system information is only useful in conjunction with version information (see Section 5.2), which reflects the reporting approach taken in most DRM systems. Examples:

- The Apple MacOS X operating system commonly denoted as “Snow Leopard” would be reported as “MACOS” with the version structure [“10”,“6”]

[8] Enumeration member value definitions are expected from the binding in order to foster binary portability of DRMAA-based applications.

- The Microsoft Windows 7 operating system would be reported as "WINNT" with the version information ["6","1"], which is the internal version number reported by the Windows API.
- All Linux distributions would be reported as operating system type "LINUX" with the major revision of the kernel, such as ["2","6"].
- The Solaris operating system is reported as "SUNOS", together with the internal version number, e.g. ["5","10"] for Solaris 10.

The DRMAA `OperatingSystem` enumeration can be mapped to other high-level specifications. Table 2 gives a non-normative set of examples.

DRMAA `OperatingSystem`	JSDL `jsdl:OperatingSystemTypeEnumeration`	GLUE v2.0
HPUX	HPUX	
LINUX	LINUX	OSFamily_t:linux
IRIX	IRIX	
TRUE64	Tru64_UNIX, OSF	
MACOS	MACOS	OSFamily_t:macosx
SUNOS	SunOS, SOLARIS	OSFamily_t:solaris
WIN	WIN95, WIN98, Windows_R_Me	OSFamily_t:windows
WINNT	WINNT, Windows_2000, Windows_XP	OSFamily_t:windows
AIX	AIX	OSName_t:aix
UNIXWARE	SCO_UnixWare, SCO_OpenServer	
BSD	BSDUNIX, FreeBSD, NetBSD, OpenBSD	

Table 2: Mapping example for the DRMAA `OperatingSystem` enumeration

4.2 CpuArchitecture enumeration

DRMAA supports identifying the processor instruction set architecture on execution hosts. The `CpuArchitecture` enumeration is used as data type in job submission, advance reservation and system monitoring. It defines a set of standardized identifiers for processor architecture families. The list is a shortened version of the according CIM Schema [7], It includes only processor families that are supported by the majority of DRM systems available at the time of writing:

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

ALPHA: The DEC Alpha / Alpha AXP processor architecture.

ARM: The ARM processor architecture.

CELL: The Cell processor architecture.

PARISC: The PA-RISC processor architecture.

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

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

IA64: The Itanium processor architecture.

MIPS: The MIPS processor architecture.

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

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

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

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

OTHER_CPU: A processor architecture not specified in this list.

The DRMAA `CpuArchitecture` enumeration can be mapped to other high-level APIs. Table 3 gives a non-normative set of examples.

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

(See footnote)[9]

DRMAA `CpuArchitecture`	JSDL `jsdl:ProcessorArchitectureEnumeration`	GLUE v2.0
ALPHA	other	
ARM	arm	
CELL	other	
PARISC	parisc	
X86	x86_32	Platform_t:i386
X64	x86_64	Platform_t:amd64
IA64	ia64	Platform_t:itanium
MIPS	mips	
PPC	powerpc	Platform_t:powerpc
PPC64	powerpc	Platform_t:powerpc
SPARC	sparc	Platform_t:sparc
SPARC64	sparc	Platform_t:sparc

Table 3: Mapping example for DRMAA `CpuArchitecture` enumeration

4.3 ResourceLimitType enumeration

Modern DRM systems expose resource constraint capabilities from the operating system also for jobs. The `ResourceLimitType` enumeration represents the most common *setrlimit* parameters [6] supported in DRM systems. In general, resource limitations aim at the level of jobs. If a job is instantiated as multiple processes, the behavior is implementation-specific.

(See footnote)[10]

[9] This kind of reporting is the only one that makes sense from the application point of view.

[10] The June 2011 face-to-face meeting had hard discussion on the relation between operating system processes, jobs, and slots. It was decided that slot is a truly opaque concept, which means that you cannot do resource contraints on something that is implementation-specific. Therefore, the spec semantics must focus on jobs only, and leave the interpretation to the DRM system / DRMAA implementation.This leads to some intentional fuzzying of descriptions for ResourceLimitType members.

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

CORE_FILE_SIZE: The maximum size of the core dump file created on fatal errors of the job, in kilobyte. Setting this value to zero SHOULD disable the creation of core dump files on the execution host.

CPU_TIME: The maximum time in seconds the job is allowed to perform computations. The value SHOULD be interpreted as sum for all processes belonging to the job. This value MUST only include time the job is spending in `JobState::RUNNING` (see Section 8.1).

DATA_SEG_SIZE: The maximum amount of memory the job can allocate on the heap e.g. for object creation, in kilobyte.

FILE_SIZE: The maximum file size the job can generate, in kilobyte.

OPEN_FILES: The maximum number of file descriptors the job is allowed to have open at the same time.

STACK_SIZE: The maximum amount of memory the job can allocate on the stack, e.g. for local variables, in kilobyte.

VIRTUAL_MEMORY: The maximum amount of memory the job is allowed to allocate, in kilobyte.

WALLCLOCK_TIME: The maximum wall clock time in seconds that all processes of a job are allowed to exist. The time amount MUST include the time spent in RUNNING state, and MAY also include the time spent in SUSPENDED state (see Section 8.1). The limit value MAY also be used for job scheduling decisions by the DRM system or the implementation.

(See footnote)[11]

4.4 JobTemplatePlaceholder enumeration

The `JobTemplatePlaceholder` enumeration defines constant macros to be used in string attributes of a `JobTemplate` instance (see Section 5.7).

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

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

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

The `PARAMETRIC_INDEX` placeholder SHOULD be usable at any position within an attribute value that supports place holders. It SHALL be substituted by the parametric job index when `JobSession::runBulkJobs`

[11] "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.

The understanding of wall clock time and CPU time was decided in the Apr 6th and 13th 2011 and June 29th 2011 conf call. Condor and Grid Engine also add SUSPEND time to wallclock time, but LSF does not.

is called (see Section 8.2.7). If the job template is used for a `JobSession:runJob` call, `PARAMETRIC_INDEX` SHOULD be substituted with a constant implementation-specific value.

(See footnote)[12]

4.5 DrmaaCapability

The `DrmaaCapability` enumeration expresses DRMAA features and data attributes that may or may not be supported by a particular implementation. Applications are expected to check the availability of optional capabilities through the `SessionManager::supports` method (see Section 7.1.5).

```
enum DrmaaCapability {
  ADVANCE_RESERVATION, RESERVE_SLOTS, CALLBACK, BULK_JOBS_MAXPARALLEL,
   JT_EMAIL, JT_STAGING, JT_DEADLINE, JT_MAXSLOTS, JT_ACCOUNTINGID,
  RT_STARTNOW, RT_DURATION, RT_MACHINEOS, RT_MACHINEARCH
};
```

ADVANCE_RESERVATION: Indicates that the implementation supports advance reservation through the interfaces (`ReservationSession` and `Reservation`).

RESERVE_SLOTS: Indicates that the advance reservation functionality is targeting slots. If this capability is not given, the advance reservation is targeting whole machines as granularity level.

CALLBACK: Indicates that the implementation supports event notification through a `DrmaaCallback` interface in the application.

BULK_JOBS_MAXPARALLEL: Indicates that the `maxParallel` parameter in the `JobSession::runBulkJobs` method is considered and supported by the implementation.

JT_EMAIL: Indicates that the optional `email`, `emailOnStarted`, and `emailOnTerminated` attributes in job templates are supported by the implementation.

JT_STAGING: Indicates that the optional `JobTemplate::stageInFiles` and `JobTemplate::stageOutFiles` attributes are supported by the implementation.

JT_DEADLINE: Indicates that the optional `JobTemplate::deadlineTime` attribute is supported by the implementation.

JT_MAXSLOTS: Indicates that the optional `JobTemplate::maxSlots` attribute is supported by the implementation.

JT_ACCOUNTINGID: Indicates that the optional `JobTemplate::accountingId` attribute is supported by the implementation.

RT_STARTNOW: Indicates that the `ReservationTemplate::startTime` attribute accepts the `NOW` value.

RT_DURATION: Indicates that the optional `ReservationTemplate::duration` attribute is supported by the implementation.

RT_MACHINEOS: Indicates that the optional `ReservationTemplate::machineOS` attribute is supported by the implementation.

[12] Placeholders for other job template attributes were rejected, in order to avoid circular dependencies (Conf. call Oct 20th 2010). Any extended semantic of placeholders in comparison to DRMAA1 was rejected, since the support in the DRM system didn't change. (conf call Apr. 20th 2011)

RT_MACHINEARCH: Indicates that the optional `ReservationTemplate::machineArch` attribute is supported by the implementation.

5 Extensible Data Structures

DRMAA defines a set of data structures commonly used in the API to express information for and from the DRM system. A DRMAA implementation MAY extend these structures with *implementation-specific attributes*. Behavioral aspects of such extended attributes are out of scope for DRMAA. Implementations MAY even ignore the attribute values in some situations.

A language binding MUST define a consistent mechanism to realize implementation-specific structure 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. Instances of extended structures SHALL still be treated in a "call-by-value" fashion.

Implementations SHALL only extend data structures in the way specified by the language binding. The introspection of supported implementation-specific attributes is offered by the `DrmaaReflective` interface (see Section 5.9). Implementations SHOULD also support native introspection functionalities if defined by the language binding.

Language bindings MAY define how the native introspection capabilities of a language or it's runtime environment can be used. These mechanisms MUST work in parallel to the `DrmaaReflective` interface.

(See footnote)[13]

5.1 QueueInfo structure

DRMAA defines queues as opaque concept for an implementation, which allows different mappings to DRMS concepts (see Section 1.2). The DRMAA `QueueInfo` struct therefore contains only the name of the queue, but can be extended by the implementation as described above. All such structure instances are read-only.

```
struct QueueInfo {
  string name;
};
```

5.1.1 name

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

[13] Comparison to DRMAA 1.0: The binding of job template attribute names and exception names to strings was removed. Language bindings have to define their own mapping, if needed.
One example for native language introspection support could be attributes.
There was a discussion to remove the attribute ignorance possibility for implementations, in order to have a defined error when unknown attributes are used. This was rejected on the Apr. 13th conf call, since applications do not need the error as indication for missing attribute support. Instead, they should use the given introspection capabilities.

5.2 Version structure

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

```
struct Version {
  string major;
  string minor;
};
```

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

Implementations SHOULD NOT extend this structure with implementation-specific attributes.

(See footnote)[14]

5.3 MachineInfo structure

The `MachineInfo` structure describes the properties of a particular execution host in the DRM system. It contains read-only information. An implementation or its DRM system MAY restrict jobs in their resource utilization even below the limits described in the `MachineInfo` structure. The limits given here MAY be imposed by the hardware configuration, or MAY be be imposed by DRM system policies.

```
struct MachineInfo {
  string name;
  boolean available;
  long sockets;
  long coresPerSocket;
  long threadsPerCore;
  double load;
  long physMemory;
  long virtMemory;
  OperatingSystem machineOS;
  Version machineOSVersion;
  CpuArchitecture machineArch;
};
```

5.3.1 name

This attribute describes the name of the machine as reported by the DRM system. The format of the machine name is implementation-specific, but MAY be a DNS host name. The naming scheme SHOULD be consistent among all machine struct instances.

[14] We could see no use case in doing implementation-specific extensions here, so this structure is not considered in DrmaaReflective. Another reason is that versioning information may be used for control flow decisions. Therefore, it should be portable in any case.

5.3.2 available

This attribute expresses the usability of the machine for job execution at the time of querying. The value of this attribute SHALL NOT influence the validity of job templates referencing `MachineInfo` instances. DRM systems and their DRMAA implementation MAY allow to submit jobs intended for machines unavailable at this time.

(See footnote)[15]

5.3.3 sockets

This attribute describes the number of processor sockets (CPUs) usable for jobs on the machine. The attribute value MUST be greater than 0. In the case where the correct value is unknown to the implementation, the value MUST be set to 1.

5.3.4 coresPerSocket

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

5.3.5 threadsPerCore

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

5.3.6 load

This attributes describes the 1-minute average load on the given machine. Implementations MAY use the same mechanism as the Unix *uptime* command. The value has only informative character, and should not be utilized by end user applications for job scheduling purposes. An implementation MAY provide delayed or averaged data here, if necessary due to implementation issues. The implementation strategy on non-Unix systems is undefined.

(See footnote)[16]

5.3.7 physMemory

This attribute describes the amount of physical memory in kilobyte installed in this machine.

5.3.8 virtMemory

This attribute describes the amount of virtual memory in kilobyte available for a job executing on this machine. The virtual memory SHOULD be defined as the sum of physical memory installed, plus the configured swap space for the operating system. The value is expected to be used as indicator whether or not an application is able to get its memory allocation needs fulfilled on a particular machine. Implementations

[15] These jobs are expected to be queued until the machine becomes available again.

[16] In July 2011, there was a short debate on the list if this value should be normalized by the library to ¡0,1¿. It was rejected, since DRMAA should just forward given information from the DRM / OS, for which the maximum value is typically not known.

SHOULD derive this value directly from operating system information, without further consideration of additional memory allocation restrictions, such as address space ranges or already running processes.

5.3.9 machineOS

This attribute describes the operating system installed on the machine, with values as specified in Section 4.1.

5.3.10 machineOSVersion

This attribute describes the operating system version on the machine, with values as specified in Section 4.1.

5.3.11 machineArch

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

5.4 SlotInfo structure

DRMAA defines slots as opaque concept for an implementation, which allows different mappings to DRMS concepts (see Section 1.2). The DRMAA `SlotInfo` structure describes the amount of reserved slots on a machine. Implementations SHALL NOT extend this structure with implementation-specific attributes. All such structure instances are read-only.

(See footnote)[17]

```
struct SlotInfo {
  string machineName;
  string slots;
};
```

5.4.1 machineName

The name of the machine. Strings returned here SHOULD be equal to the `MachineInfo::name` attribute in the matching `MachineInfo` instance.

5.4.2 slots

The number of slots reserved on the given machine. Depending on the intepretation of slots in the implementation, this value MAY be always one.

5.5 JobInfo structure

The `JobInfo` structure provides detailed information about the characteristics of a (bulk) job.

```
struct JobInfo {
  string jobId;
  long exitStatus;
```

[17] We could see no use case in realizing implementation-specific extensions here, so this structure is not considered in DrmaaReflective.

```
    string terminatingSignal;
    string annotation;
    JobState jobState;
    any jobSubState;
    OrderedSlotInfoList allocatedMachines;
    string submissionMachine;
    string jobOwner;
    long slots;
    string queueName;
    TimeAmount wallclockTime;
    long cpuTime;
    AbsoluteTime submissionTime;
    AbsoluteTime dispatchTime;
    AbsoluteTime finishTime;
  };
```

It is used in two occasions - first for the representation of information about a single job, and second as filter expression when retrieving a list of jobs.

In both usage scenarios, the structure information has to be understood as snapshot of the live DRM system. Multiple values being set in one structure instance should be interpreted as "occurring at the same time". In real implementations, some granularity limits must be assumed - for example, the `wallclockTime` and the `cpuTime` attributes might hold values that were measured with a very small delay one after each other.

In the filtering case, the value `UNSET` for an attribute MUST express wildcard semantics, meaning that this part of `JobInfo` is ignored for filtering.

DRMAA makes no assumption on the `JobInfo` availability for jobs in a a "Terminated" state (see Section 8.1). Implementations SHOULD allow to fetch information about such jobs, complete or incomplete, for a reasonable amount of time. For such terminated jobs, implementations MAY also decide to return only partially filled `JobInfo` instances.

(See footnote)[18]

For additional DRMS-specific information, the `JobInfo` structure MAY be extended by the DRMAA implementation (see Section 5).

(See footnote)[19]

5.5.1 jobId

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

[18] We want to tackle performance restrictions in the communication with the DRM system by this.

[19] 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) resourceUsage is no longer supported, since this should be modelled with implementation-specific attributes (conf call Apr 13th 2011).

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).

A comment attribute was rejected (conf call May 11th).

Several conf. calls in 2011 ended up in the conclusion that data reaping cannot be clarified by DRMAA. There are too many completely different use cases in local and distributed systems.

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

5.5.2 exitStatus

For monitoring: The process exit status of the job, as reported by the operating system on the execution host. The value MAY be `UNSET`. If the job contains of multiple processes, the behavior is implementation-specific.

For filtering: Return the jobs with the given `exitStatus` value.

(See footnote)[20]

(See footnote)[21]

5.5.3 terminatingSignal

For monitoring: This attribute describes the UNIX signal that reasoned the end of the job. Implementations should document the extent to which they can gather such information in the particular DRM system.

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

5.5.4 annotation

For monitoring: Gives a human-readable annotation describing why the job is in its current state or sub-state. Implementations MAY decide to offer such description only in specific cases, so it MAY also be `UNSET`.

For filtering: This attribute is ignored for filtering.

5.5.5 jobState

For monitoring: This attribute reports the jobs current state according to the DRMAA job state model (see Section 8.1).

For filtering: Returns all jobs in the specified state. If the given state is emulated by the implementation (see Section 8.1), the implementation SHOULD raise an `InvalidArgumentException` explaining that this filter can never match.

5.5.6 jobSubState

For monitoring: This attribute reports the current implementation-specific sub-state for this job (see Section 8.1).

For filtering: Returns all jobs in the specified sub-state. If the given sub-state is not supported by the implementation, it MAY raise an `InvalidArgumentException` explaining that this filter can never match.

(See footnote)[22]

[20] Jobs without exit status information should be filtered out by asking for the appropriate states.

[21] June 29th 2011 conf call decided to explicitely decline any relationship between job status and exit code, since there is no common behavior in DRM systems. For this reason, exit status is allowed to be UNSET, without giving any further reasons. It is, however, expected that many implementations will put this on UNSET in the non-terminal job states.

[22] As the jobSubState is an opaque object, any invalid usage may lead to a crash of the library. For this reason, the JUne 29th 2011 conf call decided to use only MAY here, in order to reflect the potentially missing reflection capabilities in the languages.

5.5.7 allocatedMachines

This attribute expresses a set of machines that is utilized for job execution. Each `SlotInfo` instance in the attribute value describes the utilization of a particular execution host, and of a set of slots related to this host.

Implementations MAY decide to give the ordering of machine names a particular meaning, for example putting the master node of a parallel job at first position. This decision should be documented for the user.

For monitoring: The attribute lists the machines and the slot count per machine allocated for the job. The slot count value MAY be `UNSET`. The machine name value MUST be set.

For filtering: Returns all jobs that fulfill the following condition: The job is executed on a superset of the given list of machines, and got at least the given number of slots on the particular machine. The `slots` value per machine MUST be allowed to have an `UNSET` value. In this case, only the machine condition SHALL be checked.

5.5.8 submissionMachine

This attribute provides the name of the submission host for this job. The machine name SHOULD be equal to the according `MachineInfo::name` attribute in monitoring data.

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

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

5.5.9 jobOwner

For monitoring: This attribute reports the job owner as recorded in the DRM system.

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

5.5.10 slots

For monitoring: This attribute reports the number slots that were allocated for the job. The value SHOULD be in between `JobTemplate::minSlots` and `JobTemplate::maxSlots`.

For filtering: Return all jobs with the specified number of reserved slots.

5.5.11 queueName

For monitoring: This attribute reports the name of the queue in which the job was queued or started (see Section 1.2).

For filtering: Returns all jobs that were queued or started in the queue with the specified name.

5.5.12 wallclockTime

For monitoring: The accumulated wall clock time, with the semantics as defined in Section 4.3.

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

5.5.13 cpuTime

For monitoring: The accumulated CPU time, with the semantics as defined in Section 4.3.

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

5.5.14 submissionTime

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

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

5.5.15 dispatchTime

For monitoring: The time the job first entered a "Started" state (see Section 8.1). On job restart or rescheduling, this value does not change.

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

5.5.16 finishTime

For monitoring: The time the job first entered a "Terminated" state (see Section 8.1).

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

5.6 ReservationInfo structure

The structure provides information about an existing advance reservation, as reported by the DRM system.

```
struct ReservationInfo {
  string reservationId;
  string reservationName;
  AbsoluteTime reservedStartTime;
  AbsoluteTime reservedEndTime;
  StringList usersACL;
  long reservedSlots;
  OrderedSlotInfoList reservedMachines;
};
```

The structure is used for the expression of information about a single advance reservation. Information provided in this structure SHOULD NOT change over the reservation lifetime. However, implementations MAY reflect the altering of advance reservations outside of DRMAA sessions.

For additional DRMS-specific information, the `ReservationInfo` structure MAY be extended by the implementation (see Section 5).

5.6.1 reservationId

Returns the stringified identifier assigned to the advance reservation by the DRM system.

5.6.2 reservationName

This attribute describes the reservation name that was stored by the implementation or the DRM system for the reservation. It SHOULD be derived from the `reservationName` attribute in the originating `ReservationTemplate`.

5.6.3 reservedStartTime

This attribute describes the start time for the reservation. If the value is `UNSET`, it expresses an unrestricted start time (i.e., *minus infinity*) for this reservation.

5.6.4 reservedEndTime

This attribute describes the end time for the reservation. If the value is `UNSET`, the behavior is implementation-specific.

(See footnote)[23]

5.6.5 usersACL

The list of the users that are permitted to submit jobs to the reservation. The formatting of user identities is implementation-specific, but SHOULD be consistent with the user information representation in job templates and reservation templates.

5.6.6 reservedSlots

This attribute describes the number of slots reserved by the DRM system. The value SHOULD range in between `ReservationTemplate::minSlots` and `ReservationTemplate::maxSlots`.

5.6.7 reservedMachines

This attribute describes the set of machines that were reserved under the conditions described in the according reservation template. Each `SlotInfo` instance in this list describes the reservation of a particular machine and of a set of slots related to this machine. The sum of all slot counts in the sequence SHOULD be equal to `ReservationInfo::reservedSlots`.

5.7 JobTemplate structure

A DRMAA application uses the `JobTemplate` structure to define characteristics of a job submission. The template instance is passed to the DRMAA `JobSession` instance when job execution is requested.

```
struct JobTemplate {
  string remoteCommand;
  OrderedStringList args;
  boolean submitAsHold;
  boolean rerunnable;
  Dictionary jobEnvironment;
  string workingDirectory;
  string jobCategory;
```

[23] Mai 18th 2011 conf call rejected to treat UNSET as unrestricted end time (i.e. "plus infinity") here.

```
    StringList email;
    boolean emailOnStarted;
    boolean emailOnTerminated;
    string jobName;
    string inputPath;
    string outputPath;
    string errorPath;
    boolean joinFiles;
    string reservationId;
    string queueName;
    long minSlots;
    long maxSlots;
    long priority;
    OrderedStringList candidateMachines;
    long minPhysMemory;
    OperatingSystem machineOS;
    CpuArchitecture machineArch;
    AbsoluteTime startTime;
    AbsoluteTime deadlineTime;
    Dictionary stageInFiles;
    Dictionary stageOutFiles;
    Dictionary resourceLimits;
    string accountingId;
  };
```

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

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.

The DRMAA job template concept makes a distinction between *mandatory* and *optional* attributes. Mandatory attributes MUST be supported by the implementation in the sense that they are evaluated on job submission. Optional attributes MAY be evaluated on job submission, but MUST be provided as part of the `JobTemplate` structure in the implementation. If an unsupported optional attribute has a value different to `UNSET`, the job submission MUST fail with a `UnsupportedAttributeException`. DRMAA applications are expected to check for the availability of optional attributes before using them (see Section 4.5).

An implementation MUST support `JobTemplatePlaceholder` placeholders at the occasions defined in this specification. They MAY also allow their usage in other attributes.

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.

(See footnote)[24]

5.7.1 remoteCommand

This attribute describes the command to be executed on the remote host. In case this parameter contains path information, it MUST be interpreted as relative to the execution host file system. The implementation SHOULD NOT use the value of this attribute to trigger file staging activities. Instead, the file staging should be performed by the application explicitly.

The behavior of the implementation with an `UNSET` value in this attribute is undefined.

The support for this attribute is mandatory.

5.7.2 args

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

The support for this attribute is mandatory.

5.7.3 submitAsHold

This attribute defines if the job(s) should have `QUEUED` or `QUEUED_HELD` (see Section 8.1) as initial state after submission. Since the boolean `UNSET` value defaults to `False`, jobs are submitted as non-held if this attribute is not set.

The support for this attribute is mandatory.

5.7.4 rerunnable

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

The support for this attribute is mandatory.

(See footnote)[25]

[24] 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.

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 should use according implementation-specific attributes for this.

GridEngine does not support to request a number of slots per machine - of course in a default installation, since you can do everything in GridEngine ... This is the reason for not having such an attribute.

[25] The differentiation between rerunnable and checkpointable was decided on a conf call (Aug 25th 2010). Checkpointability indication was intentionally left out, since there is no common understanding in the DRM systems (conf call Apr. 27th, 2011).

5.7.5 jobEnvironment

This attribute holds the environment variable settings to be configured on the execution machine(s). The values SHOULD override the execution host environment settings.

The support for this attribute is mandatory.

5.7.6 workingDirectory

This attribute specifies the directory where the job or the bulk jobs are executed. If the attribute value is `UNSET`, the behavior is undefined. If set, the attribute value MUST be evaluated relative to the file system on the execution host. The attribute value MUST be allowed to contain either the `JobTemplatePlaceholder::HOME_DIRECTORY` or the `JobTemplatePlaceholder::PARAMETRIC_INDEX` placeholder (see Section 4.4).

The `workingDirectory` attribute should be specified by the application in a syntax that is common at the host where the job is executed. Implementations MAY perform according validity checks on job submission. If the attribute is set and no placeholder is used, an absolute directory specification is expected. If the attribute is set and the job was submitted successfully and the directory does not exist on the execution host, the job MUST enter the state `JobState::FAILED`.

The support for this attribute is mandatory.

5.7.7 jobCategory

This attribute defines the job category to be used (see Section 1.4). A valid input SHOULD be one of the strings in `JobSession::jobCategories` (see Section 8.2.3), otherwise an `InvalidArgumentException` SHOULD be raised.

The support for this attribute is mandatory.

5.7.8 email

This attribute defines a list of email addresses that SHOULD be used when the DRM system sends status notifications. Content and formatting of the emails are defined by the implementation or the DRM system. If the attribute value is `UNSET`, no emails SHOULD be sent to the user running the job(s), even if the DRM system default behavior is different.

The support for this attribute is optional, expressed by the `DrmaaCapability::JT_EMAIL` flag. If an implementation cannot configure the email notification functionality of the DRM system, or if the DRM system has no such functionality, the attribute SHOULD NOT be supported in the implementation.

(See footnote)[26]

5.7.9 emailOnStarted / emailOnTerminated

The `emailOnStarted` flag indicates if the given email address(es) SHOULD get a notification when the job (or any of the bulk jobs) entered one of the "Started" states. `emailOnTerminated` fulfills the same purpose for the "Terminated" states. Since the boolean `UNSET` value defaults to `False`, the notification about state changes SHOULD NOT be sent if the attribute is not set.

[26] The blockEmail attribute in the JobTemplate was replaced by the UNSET semantic for the email addresses. (conf. call July 28th 2010). This became an optional attribute, since we mandate the 'switch off' semantic in case of `UNSET`.

The support for these attributes is optional, expressed by the `DrmaaCapability::JT_EMAIL` flag.

5.7.10 jobName

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

The support for this attribute is mandatory.

5.7.11 inputPath / outputPath / errorPath

This attribute specifies standard input / output / error stream of the job as file path. If the attribute value is `UNSET`, the behavior is undefined. If set, the attribute value MUST be evaluated relative to the file system of the execution host. Implementations MAY perform validity checks for the path syntax on job submission. The attribute value MUST be allowed to contain any of the `JobTemplatePlaceholder` placeholders (see Section 4.4). If the attribute is set and no placeholder is used, an absolute file path specification is expected.

If the `outputPath` or `errorPath` file does not exist at the time of job execution start, the file SHALL automatically be created. An existing `outputPath` or `errorPath` file SHALL be opened in append mode.

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

The support for this attribute is mandatory.

5.7.12 joinFiles

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

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

The support for this attribute is mandatory.

5.7.13 reservationId

Specifies the identifier of the existing advance reservation to be associated with the job(s). The application is expected to generate this ID by creating an advance reservation through the `ReservationSession` interface. The resulting `reservationId` (see Section 9.2.1) then acts as valid input for this job template attribute. Implementations MAY support a reservation identifier from non-DRMAA information sources as valid input. The behavior on conflicting settings between the job template and the granted advance reservation is undefined.

The support for this attribute is mandatory.

5.7.14 queueName

This attribute specifies the name of the queue the job(s) should be submitted to. In case this attribute value is `UNSET`, the implementation SHOULD use the DRM systems default queue. If no default queue is defined or if the given queue name is not valid, the job submission MUST lead to an `InvalidArgumentException`.

The `MonitoringSession::getAllQueues` method (see Section 10.1) supports the determination of valid queue names. Implementations SHOULD allow at least these queue names to be used in the `queueName` attribute. Implementations MAY also support queue names from non-DRMAA information sources as valid input.

If `MonitoringSession::getAllQueues` returns an empty list, this attribute MUST be only allowed to have the value `UNSET`.

Since the meaning of "queues" is implementation-specific, there is no DRMAA-defined effect when using this attribute. Implementations therefore should document the effects of this attribute in their targeted environment.

The support for this attribute is mandatory.

(See footnote)[27]

5.7.15 minSlots

This attribute expresses the minimum number of slots requested per job (see also Section 1.2). If the value of `minSlots` is `UNSET`, it SHOULD default to 1.

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

The support for this attribute is mandatory.

(See footnote)[28]

5.7.16 maxSlots

This attribute expresses the maximum number of slots requested per job (see also Section 1.2). If the value of `maxSlots` is `UNSET`, it SHOULD default to the value of `minSlots`.

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

The support for this attribute is optional, as indicated by the `DrmaaCapability::JT_MAXSLOTS` flag.

(See footnote)[29] .

5.7.17 priority

This attribute specifies the scheduling priority for the job. The interpretation of the given value is implementation-specific.

The support for this attribute is mandatory.

[27] As one example, requesting a number of slots for a job in one queue has no implication on the number of utilized machines at run-time.

[28] The hint regarding number of concurrent processes intentionally does not speak about processes per host - this would create semantics for our opaque slot concept.

[29] Torque does not support maxSlots on job submission, conf call on May 11th decided to keep it as optional feature. Expected use cases are billing limitations and parallel job scalability considerations

5.7.18 candidateMachines

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

The support for this attribute is mandatory.

5.7.19 minPhysMemory

This attribute denotes the minimum amount of physical memory in kilobyte that should be available for the job. If the job gets more than one slot, the interpretation of this value is implementation-specific. If this resource demand cannot be fulfilled, an `InvalidArgumentException` SHOULD be raised at job submission time. If the problem can only be detected after job submission, the job SHOULD enter `JobState::FAILED` accordingly.

The support for this attribute is mandatory.

5.7.20 machineOS

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

The support for this attribute is mandatory.

(See footnote)[30]

5.7.21 machineArch

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

The support for this attribute is mandatory.

5.7.22 startTime

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

The support for this attribute is mandatory.

5.7.23 deadlineTime

Specifies a deadline after which the implementation or the DRM system SHOULD change the job state to any of the "Terminated" states (see Section 8.1).

The support for this attribute is optional, as expressed by the `DrmaaCapability::JT_DEADLINE`.

[30] Requesting a specific operating system version beyond the type is not supported by the majority of DRM systems (conf call Jul 28th 2010)

5.7.24 stageInFiles / stageOutFiles

This attribute specifies what files should be transferred (staged) as part of the job execution. The data staging operation MUST be a copy operation between the submission host and a execution host. File transfers between execution hosts are not covered by DRMAA.

The attribute value is formulated as dictionary. For each key-value pair in the dictionary, the key defines the source path of one file or directory, and the value defines the destination path of one file or directory for the copy operation. For `stageInFiles`, the submission host acts as source, and the execution host act as destination. For `stageOutFiles`, the execution host acts as source, and the submission host act as destination.

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

Relative path specifications for the submission host should be interpreted starting from the current working directory of the DRMAA application at the time of job submission. The behavior for relative path specifications on the execution is implementation-specific. Implementations MAY use *JobTemplate::workingDirectory*, if defined, as starting point on the execution host.

Jobs SHOULD NOT enter `JobState::DONE` unless all staging operations are finished. The behavior in case of missing files is implementation-specific. The support for wildcard operators in path specifications is implementation-specific. Any kind of recursive or non-recursive copying behavior is implementation-specific.

If the job category (see Section 1.4) implies a parallel job (e.g., MPI), the copy operation SHOULD target the execution host of the parallel job master as destination. A job category MAY also trigger file distribution to other hosts participating in the job execution.

The support for this attribute is optional, expressed by the `DrmaaCapability::JT_STAGING` flag.

(See footnote)[31]

5.7.25 resourceLimits

This attribute specifies the limits on resource utilization of the job(s) on the execution host(s). The valid dictionary keys and their value semantics are defined in Section 4.3.

The following resource restrictions should operate as soft limit, meaning that exceeding the limit SHOULD NOT influence the job state from a DRMAA perspective:

- `CORE_FILE_SIZE`
- `DATA_SEG_SIZE`
- `FILE_SIZE`

[31] 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)

- OPEN_FILES
- STACK_SIZE
- VIRTUAL_MEMORY

The following resource restrictions should operate as hard limit, meaning that exceeding the limit MAY terminate the job. The termination MAY be performed by the DRM system. It MAY also be done by the job itself if it reacts on a signal from the DRM system or the execution host operating system:

- CPU_TIME
- WALLCLOCK_TIME

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

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

(See footnote)[32]

5.7.26 accountingId

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 untested identification information in addition to the implementation-specific authentication (see Section 12).

The support for this attribute is optional, as described by the `DrmaaCapability::JT_ACCOUNTINGID` flag.

5.8 ReservationTemplate structure

In order to define the characteristics of a reported advance reservation, the DRMAA application creates an `ReservationTemplate` instance and submits it through the `ReservationSession` methods.

```
struct ReservationTemplate {
  string reservationName;
  AbsoluteTime startTime;
  AbsoluteTime endTime;
  TimeAmount duration;
  long minSlots;
```

[32] 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).

This distinguishing between different reactions on limit violation was restricted to the job entering, or not entering, the FAILED state. All further effects (e.g., no more *open()* calls possible) are out of scope for DRMAA, since they relate to operating system behavior on execution host (conf call May 4th 2011).

The attribute is mandatory, since the missing general support for resource limits can be simply expressed by raising InvalidArgumentException for all types.

```
    long maxSlots;
    string jobCategory;
    StringList usersACL;
    OrderedStringList candidateMachines;
    long minPhysMemory;
    OperatingSystem machineOS;
    CpuArchitecture machineArch;
  };
```

Similar to the `JobTemplate` concept (see Section 5.7), there is a distinction between *mandatory* and *optional* attributes in the `ReservationTemplate`. Mandatory attributes MUST be supported by the implementation in the sense that they are evaluated in a `ReservationSession::requestReservation` method call. Optional attributes MAY NOT be evaluated by the particular implementation, but MUST be provided as part of the `ReservationTemplate` structure in the implementation. If an optional attribute is not evaluated, but has a value different to `UNSET`, the method call to `ReservationSession::requestReservation` MUST fail with an `UnsupportedAttributeException`.

Implementations MUST set all attribute values to `UNSET` on struct allocation. This ensures that both the DRMAA application and the library implementation can determine untouched attribute members.

A language binding specification SHOULD model the `ReservationTemplate` representation the same way as the `JobTemplate` interface, and therefore MUST specify the realization of implementation-specific attributes, printing, and the initialization to `UNSET`.

5.8.1 reservationName

A human-readable reservation name. The implementation MAY truncate or alter any application-provided name in order to adjust it to DRMS-specific constraints. The name of the reservation SHALL be automatically defined by the implementation if this attribute is `UNSET`.

The support for this attribute is mandatory.

5.8.2 startTime / endTime / duration

The time frame in which resources should be reserved. Table 4 explains the different possible parameter combinations and their semantic.

The support for `startTime` and `endTime` is mandatory. The support for `duration` is optional, as described by the `DrmaaCapability::RT_DURATION` flag. Implementations that do not support the described "sliding window" approach for the SET / SET / SET case SHOULD express this by NOT supporting the *duration* attribute.

Implementations MAY support `startTime` to have the constant value `NOW` (see Section 3), which expresses that the reservation should start at the time of reservation template approval in the DRM system. The support for this feature is declared by the `DrmaaCapability::RT_STARTNOW` flag.

5.8.3 minSlots

This attribute expresses the minimum number of slots requested per job (see also Section 1.2). If the value of `minSlots` is `UNSET`, it SHOULD default to 1.

startTime	endTime	duration	Description
UNSET	UNSET	UNSET	Invalid, SHALL leave to an InvalidArgumentException.
Set	UNSET	UNSET	Invalid, SHALL leave to an InvalidArgumentException.
UNSET	Set	UNSET	Invalid, SHALL leave to an InvalidArgumentException.
Set	Set	UNSET	Attempt to reserve resources in the specified time frame.
UNSET	UNSET	Set	Attempt to reserve resources at least for the time amount given in duration.
Set	UNSET	Set	Implies endTime = startTime + duration
UNSET	Set	Set	Implies startTime = endTime - duration
Set	Set	Set	If endTime - startTime is larger than duration, perform a reservation attempt where the demanded duration is fulfilled at the earliest point in time after startTime, and without extending endTime ("sliding window" approach). If endTime - startTime is smaller than duration, the reservation attempt SHALL leave to an InvalidArgumentException. If endTime - startTime and duration are equal, duration SHALL be ignored.

Table 4: Parameter combinations for the advance reservation time frame. If duration is not supported, it should be treated as UNSET.

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

The support for this attribute is mandatory.

(See footnote)[33]

5.8.4 maxSlots

This attribute expresses the maximum number of slots requested per job (see also Section 1.2). If the value of maxSlots is UNSET, it SHOULD default to the value of minSlots.

Implementations MAY interpret the slot count as number of concurrent processes being allowed to run. If this interpretation is taken, and maxSlots is greater than 1, than the jobCategory MAY also be demanded on job submission, in order to express the nature of the intended parallel job execution.

The support for this attribute is mandatory.

(See footnote)[34]

[33] The hint regarding number of concurrent processes intentionally does not speak about processes per host - this would create semantics for our opaque slot concept.

[34] Conf call June 29th 2011: For maxSlots ¿ 1, the demand for a job category is intentionally only MAY. This is reasoned by the fact that in most DRM systems, advance reservation is a concept that is independent to the jobs that are later used in this reservation. So you are just requesting just some container, and you do not have to specify at this moment what kind of jobs do you want to run using this reservation (e.g. OpenMP, OpenMPI, MPICH). However, some systems need that information, so we leave it to the implementation how to deal with that.

5.8.5 jobCategory

This attribute defines the job category to be used (see Section 1.4). A valid input SHOULD be one of the strings in `JobSession::jobCategories` (see Section 8.2.3), otherwise an `InvalidArgumentException` SHOULD be raised.

The support for this attribute is mandatory.

5.8.6 usersACL

The list of the users that would be permitted to submit jobs to the created reservation. If the attribute value is `UNSET`, it should default to the user running the application.

The support for this attribute is mandatory.

5.8.7 candidateMachines

Requests that the reservation SHALL be created for the given set of machines. Implementations and their DRM system MAY decide to reserve only a subset of the given machines. If this attribute is not specified, it should default to the result of `MonitoringSession::getAllMachines` (see Section 10.1).

The support for this attribute is mandatory.

(See footnote)[35]

5.8.8 minPhysMemory

Requests that the reservation SHALL be created with machines that have at least the given amount of physical memory in kilobyte. Implementations MAY interpret this attribute value as filter for candidate machines, or as memory reservation demand on a shared execution resource.

The support for this attribute is mandatory.

(See footnote)[36]

5.8.9 machineOS

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

The support for this attribute is optional, the availability is indicated by the `DrmaaCapability::RT_MACHINEOS` flag.

(See footnote)[37]

[35] May 18th 2011 conf call identified the subset reservation feature to be only available in some of the systems, so it is no promise here.

[36] May 18th 2011 conf call identified the different understandings of memory reservation.

[37] May 18th 2011 conf call identified support in DRM systems to be mainly given by additional configuration only.

5.8.10 machineArch

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

The support for this attribute is optional, the availability is indicated by the `DrmaaCapability::RT_MACHINEARCH` flag.

(See footnote)[38]

5.9 DrmaaReflective Interface

The `DrmaaReflective` interface allows an application to determine the set of supported implementation-specific attributes. It also standardizes the read / write access to such attributes when their existence is determined at run-time by the application.

For the second class of non-mandatory attributes, the *optional* ones, applications are expected to use the DRMAA capabiliy feature (see Section 4.5).

```
interface DrmaaReflective {
  readonly attribute StringList jobTemplateImplSpec;
  readonly attribute StringList jobInfoImplSpec;
  readonly attribute StringList reservationTemplateImplSpec;
  readonly attribute StringList reservationInfoImplSpec;
  readonly attribute StringList queueInfoImplSpec;
  readonly attribute StringList machineInfoImplSpec;
  readonly attribute StringList notificationImplSpec;

  string getInstanceValue(in any instance, in string name);
  void setInstanceValue(in any instance, in string name, in string value);
  string describeAttribute(in any instance, in string name);
};
```

5.9.1 jobTemplateImplSpec

This attribute provides the list of supported implementation-specific `JobTemplate` attributes.

5.9.2 jobInfoImplSpec

This attribute provides the list of supported implementation-specific `JobInfo` attributes.

5.9.3 reservationTemplateImplSpec

This attribute provides the list of supported implementation-specific `ReservationTemplate` attributes.

5.9.4 reservationInfoImplSpec

This attribute provides the list of supported implementation-specific `ReservationInfo` attributes.

[38] May 18th 2011 conf call identified support in DRM systems to be mainly given by additional configuration only.

5.9.5 queueInfoImplSpec

This attribute provides the list of supported implementation-specific `QueueInfo` attributes.

5.9.6 machineInfoImplSpec

This attribute provides the list of supported implementation-specific `MachineInfo` attributes.

5.9.7 notificationImplSpec

This attribute provides the list of supported implementation-specific `DrmaaNotification` attributes.

5.9.8 getInstanceValue

This method allows to retrieve the attribute value for `name` from the structure instance referenced in the `instance` parameter. The return value is the stringified current attribute value.

5.9.9 setInstanceValue

This method allows to set the attribute `name` to `value` in the structure instance referenced in the `instance` parameter. In case the conversion from string input into the native attribute type leads to an error, `InvalidArgumentException` SHALL be thrown.

5.9.10 describeAttribute

This method returns a human-readable description of an attributes purpose, for the attribute referenced by `name` and `instance`. The content and language of the result value is implementation-specific.

6 Common Exceptions

The exception model specifies error information that MAY be returned by a DRMAA implementation on method calls. Implementations MAY also wrap DRMS-specific error conditions in DRMAA exceptions.

```
exception DeniedByDrmsException {string message;};
exception DrmCommunicationException {string message;};
exception TryLaterException {string message;};
exception SessionManagementException {string message;};
exception TimeoutException {string message;};
exception InternalException {string message;};
exception InvalidArgumentException {string message;};
exception InvalidSessionException {string message;};
exception InvalidStateException {string message;};
exception OutOfResourceException {string message;};
exception UnsupportedAttributeException {string message;};
exception UnsupportedOperationException {string message;};
```

The exceptions have the following general meaning, if not specified otherwise in a method description:

DeniedByDrmsException: The DRM system rejected the operation due to security issues.

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

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

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

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

InvalidArgumentException: From the viewpoint of the DRMAA library, an input parameter for the particular method call is invalid or inappropriate. If the parameter is a structure, the exception description SHOULD contain the name(s) of the problematic structure attribute(s).

InvalidSessionException: The session used for the method call is not valid, for example since the session was closed before.

InvalidStateException: The operation is not allowed in the current state of the job.

OutOfResourceException: The implementation has run out of operating system resources, such as buffers, main memory, or disk space.

UnsupportedAttributeException: The optional attribute is not supported by this DRMAA implementation.

UnsupportedOperationException: The method is not supported by this DRMAA implementation.

The DRMAA specification assumes that programming languages targeted by language bindings typically support the concept of exceptions. If a destination language does not support them (like ANSI C), the language binding specification SHOULD map error reporting to an appropriate alternative concept.

A language binding MAY chose to model exceptions as numeric error codes. 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 additional error cause as textual description. This is intended as specialization of the general error information.

Object-oriented language bindings MAY decide to derive all exception classes from one or multiple base classes, in order to support generic catch clauses.

Language bindings MAY decide to introduce a hierarchical ordering of DRMAA exceptions based on 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.

Language bindings SHOULD replace a DRMAA exception by some semantically equivalent native exception from the application runtime environment, if available.

The `UnsupportedAttributeException` may either be raised by a setter function for an attribute, or by the job submission function. This depends on the language binding design. A consistent decision for either one or the other approach MUST be declared by the language binding specification.

(See footnote)[39]

7 The DRMAA Session Concept

DRMAA relies on a session concept for most parts of the API, in order to support the persistency of job and advance reservation information in multiple runs of short-lived applications. Typical examples are job submission portals or command-line tools. The session concept also allows implementations to perform DRM system attach / detach action at dedicated points in the application control flow.

7.1 SessionManager Interface

```
interface SessionManager{
  readonly attribute string drmsName;
  readonly attribute Version drmsVersion;
  readonly attribute string drmaaName;
  readonly attribute Version drmaaVersion;
  boolean supports(in DrmaaCapability capability);
  JobSession createJobSession(in string sessionName,
                              in string contact);
  ReservationSession createReservationSession(in string sessionName,
                                              in string contact);
  JobSession openJobSession(in string sessionName);
  ReservationSession openReservationSession(in string sessionName);
  MonitoringSession openMonitoringSession (in string contact);
  void closeJobSession(in JobSession s);
  void closeReservationSession(in ReservationSession s);
  void closeMonitoringSession(in MonitoringSession s);
  void destroyJobSession(in string sessionName);
  void destroyReservationSession(in string sessionName);
  StringList getJobSessionNames();
  StringList getReservationSessionNames();
  void registerEventNotification(in DrmaaCallback callback);
};
```

The `SessionManager` interface is the main interface of a DRMAA implementation for establishing communication with the DRM system. By the help of this interface, sessions for job management, monitoring, and/or reservation management can be maintained.

Job and reservation sessions maintain persistent state information (about jobs and reservations created) between application runs. State data SHOULD be persisted in the DRMS itself. If this is not supported, the DRMAA implementation MUST realize the persistency. The data SHOULD be saved when the session is closed by the according method in the `SessionManager` interface.

The state information SHOULD be kept until the job or reservation session is explicitly reaped by the according destroy method in the `SessionManager` interface. If an implementation runs out of resources

[39] 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)

for storing session information, the closing function SHOULD throw an `OutOfResourceException`. If an application ends without closing the session properly, the behavior is unspecified.

The `contact` parameter in some of the interface methods SHALL allow the application to specify which DRM system instance to use. A contact string represents a specific installation of a specific DRM system, e.g., a Condor central manager machine at a given IP address, or a Grid Engine 'root' and 'cell'. Contact strings are always implementation-specific and therefore opaque to the application. If `contact` has the value `UNSET`, a default DRM system SHOULD be contacted. The manual configuration or automated detection of a default contact string is implementation-specific.

The re-opening of a session MUST work on the machine where the session was originally created. Implementations MAY also offer to re-open the session on another machine, if the state information is accessible.

An implementation MUST allow the application to have multiple open sessions of the same or different type at the same time. This includes the proper coordination of parallel calls to session methods that share state information.

A `SessionManager` instance SHALL be available as singleton at DRMAA application start. Language bindings MAY realize this by mapping the session manager methods to global functions.

(See footnote)[40]

7.1.1 drmsName

A read-only system identifier denoting the DRM system targeted by the implementation, e.g., "LSF" or "GridWay". Implementations SHOULD NOT make versioning information of the particular DRM system a part of this attribute value.

The value is only intended as informative output for application users.

7.1.2 drmsVersion

This attribute provides the DRM-system specific version information.

The value is only intended as informative output for application users.

7.1.3 drmaaName

This attribute contains a string identifying the vendor of the DRMAA implementation.

The value is only intended as informative output for application users.

[40] 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).

7.1.4 drmaaVersion

This attribute provides the minor / major version number information for the DRMAA implementation. The major version number MUST be the constant value "2", the minor version number SHOULD be used by the DRMAA implementation for expressing its own versioning information.

7.1.5 supports

This method allows to test if the DRMAA implementation supports a feature specified as optional. The allowed input values are specified in the `DrmaaCapability` enumeration (see Section 4.5). This method SHOULD throw no exceptions.

7.1.6 createJobSession / createReservationSession

The method creates a new job / reservation session instance. On successful completion of this method, the necessary initialization for making the session usable MUST be completed. Examples are the connection establishment from the DRMAA library to the DRM system, or the prefetching of information from non-thread-safe operating system calls.

The `sessionName` parameter denotes a unique name to be used for the new session. If a session with such a name already exists, the method MUST throw an `InvalidArgumentException`. In all other cases, including if the provided name has the value `UNSET`, a new session MUST be created with a unique name generated by the implementation.

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

7.1.7 openJobSession / openReservationSession

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

If the session referenced by `sessionName` is already opened, implementations MAY return this job or reservation session instance.

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

7.1.8 openMonitoringSession

The method opens a stateless `MonitoringSession` instance for fetching information about the DRM system. On successful completion of this method, the necessary initialization for making the session usable MUST be completed. One example is the connection establishment from the DRMAA library to the DRM system.

7.1.9 closeJobSession / closeReservationSession / closeMonitoringSession

The method MUST perform the necessary action to disengage from the DRM system. It SHOULD be callable only once, by only one of the application threads. This SHOULD be ensured by the library implementation.

Additional calls beyond the first one SHOULD lead to a `InvalidSessionException` error notification.

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

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

A language binding MAY define implicit calls to `closeJobSession`, `closeReservationSession`, or `closeMonitoringSession`, for example when session objects are destroyed. It MAY also add a `close` method to `JobSession`, `ReservationSession`, or `MonitoringSession` with the same functionality as described here. However, the `SessionManager` still MUST contain all methods as described in this specfication.

(See footnote)[41]

7.1.10 destroyJobSession / destroyReservationSession

The method MUST do whatever work is required to reap persistent or cached state information for the given session name. It is intended to be used when no session instance with this particular name is open. If session instances for the given name exist, they MUST become invalid after this method was finished successfully. Invalid sessions MUST throw `InvalidSessionException` on every attempt of utilization. This method SHALL NOT affect any jobs or reservations in the session, e.g., queued and running jobs remain queued and running.

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

7.1.11 getJobSessionNames

This method returns a list of `JobSession` names that are valid input for the `openJobSession` method.

(See footnote)[42]

7.1.12 getReservationSessionNames

This method returns a list of `ReservationSession` names that are valid input for the `openReservationSession` method.

If the DRM system does not support advance reservation, the method SHALL always throw an `UnsupportedOperationException`.

(See footnote)[43]

7.1.13 registerEventNotification

This method is used to register a `DrmaaCallback` interface (see Section 8.3) offered by the DRMAA-based application, which can be called by the implementation. If the callback functionality is not supported by the

[41] Conf call June 29th 2011: The closing of stateless monitoring sessions was intentionally kept, in order to allow an orderly shut down of the monitoring connection.

[42] June 29th 2011 conf call decided to make the method names explicit enough to see the return type.

[43] June 29th 2011 conf call decided to make the method names explicit enough to see the return type.

DRMAA implementation, this method SHALL raise an `UnsupportedOperationException`. Implementation can check for the support through the `DrmaaCapability::CALLBACK` flag (see Section 4.5). Implementations with callback support SHOULD allow to perform multiple registration calls that just update the callback target.

If the argument of the method call is `UNSET`, the currently registered callback MUST be unregistered. After such a method call returned, no more events SHALL be delivered to the application. If no callback target is registered, such a method call SHOULD return immediately without an error.

A language binding specification MUST define how the reference to an interface-compliant method can be given as argument to this method. It MUST also clarify how to pass an `UNSET` callback method reference.

8 Working with Jobs

A DRMAA job represents a single computational activity that is executed by the DRM system. There are three relevant method sets for working with jobs: The `JobSession` interface represents all control and monitoring functions available for jobs. The `Job` interface represents the common control functionality for one existing job. Sets of jobs resulting from a bulk submission are controllable as a whole by the `JobArray` interface.

8.1 The DRMAA State Model

DRMAA defines the following states for jobs:

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

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

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

QUEUED_HELD: The job has been placed on hold by the system, the administrator, or the submitting user.

RUNNING: The job is running on an execution host.

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

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

REQUEUED_HELD: The job was re-queued by the DRM system, and is currently placed on hold by the system, the administrator, or the submitting user.

DONE: The job finished without an error.

FAILED: The job exited abnormally before finishing.

If a DRMAA job state has no representation in the underlying DRMS, the DRMAA implementation MAY never report that job state value. However, all DRMAA implementations MUST provide the `JobState`

enumeration as given here. An implementation SHOULD NOT return any job state value other than those defined in the `JobState` enumeration.

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

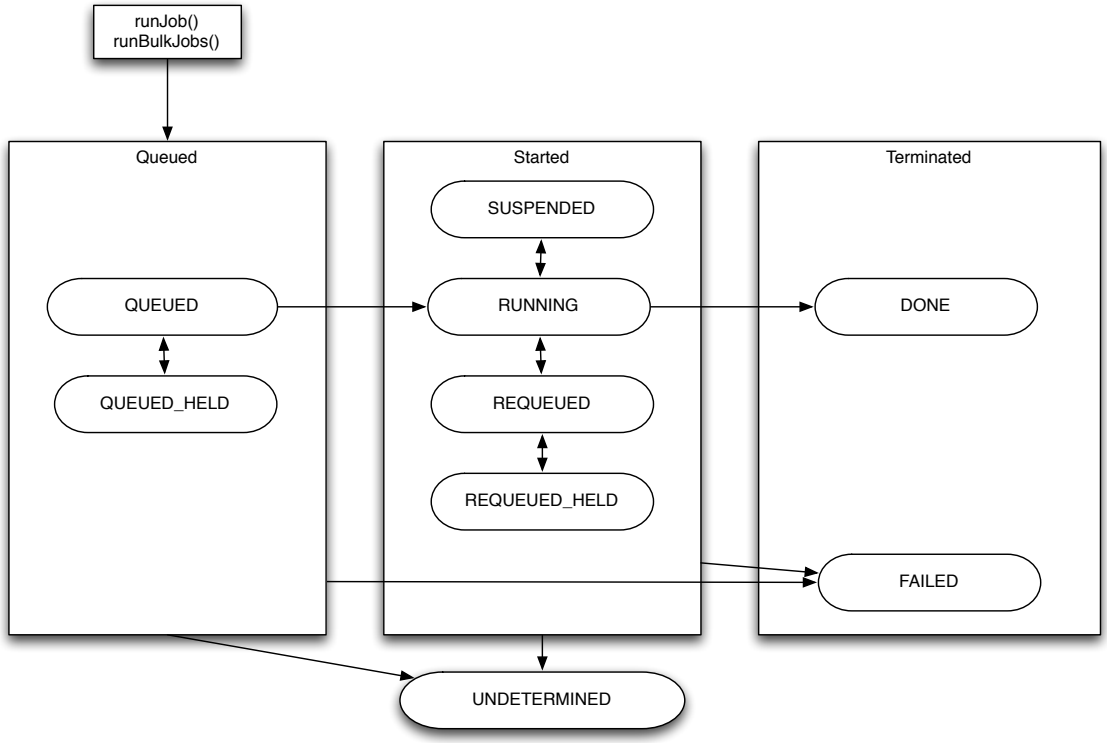

Figure 1: DRMAA Job State Transition Model

The transition diagram in Figure 1 expresses the classification of possible job states into "Queued", "Started", and "Terminated". The "Terminated" class of states is final, meaning that no further state transition is allowed.

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

When an application requests job state information, the implementation SHOULD also provide the `jobSubState` value (see Section 5.5.6) to explain DRM-specific details about the job state. The value of this attribute is implementation-specific, but should be documented properly. Examples are extra states for staging phases or details on the hold reason. Implementations SHOULD define a DRMS-specific data structure for the sub-state information that can be converted to / from the data type defined by the language binding.

The IDL definition declares the `jobSubState` attribute 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 keeps source code portability across DRMAA implementations, and accepts an `UNSET` value.

The DRMAA job state model can be mapped to other high-level API state models. Table 5 gives a non-normative set of examples.

DRMAA JobState	SAGA JobState [4]	OGSA-BES Job State [3]
UNDETERMINED	N/A	N/A
QUEUED	Running	Pending (Queued)
QUEUED_HELD	Running	Pending (Queued)
RUNNING	Running	Running (Executing)
SUSPENDED	Suspended	Running (Suspended)
REQUEUED	Running	Running (Queued)
REQUEUED_HELD	Running	Running (Queued)
DONE	Done	Finished
FAILED	Cancelled, Failed	Cancelled, Failed

Table 5: Example Mapping of DRMAA Job States

(See footnote)[44]

8.2 JobSession Interface

A job session instance acts as container for job instances controlled through the DRMAA API. The session methods support the submission of new jobs and the monitoring of existing jobs. The relationship between jobs and their session MUST be persisted, as described in Section 7.1.

```
interface JobSession {
  readonly attribute string contact;
  readonly attribute string sessionName;
  readonly attribute StringList jobCategories;
  JobList getJobs(in JobInfo filter);
  JobArray getJobArray(in string jobArrayId);
  Job runJob(in JobTemplate jobTemplate);
  JobArray runBulkJobs(
      in JobTemplate jobTemplate,
```

[44] 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).

REQUEUED and REQUEUED_HELD maps to RUNNING in BES, since BES does not allow a transition between Running and Pending (mailing list, APr. 2011)

```
        in long beginIndex,
        in long endIndex,
        in long step,
        in long maxParallel);
    Job waitAnyStarted(in JobList jobs, in TimeAmount timeout);
    Job waitAnyTerminated(in JobList jobs, in TimeAmount timeout);
  };
```

(See footnote)[45]

8.2.1 contact

This attribute reports the `contact` value that was used in the `SessionManager::createJobSession` call for this instance (see Section 7.1). If no value was originally provided, the default contact string from the implementation MUST be returned. This attribute is read-only.

8.2.2 sessionName

This attribute reports the session name, a value that resulted from the `SessionManager::createJobSession` or `SessionManager::openJobSession` call for this instance (see Section 7.1). This attribute is read-only.

8.2.3 jobCategories

This method provides the list of valid job category names which can be used for the `jobCategory` attribute in a `JobTemplate` instance. Further details about job categories are described in Section 1.4.

8.2.4 getJobs

This method returns the set of jobs that belong to the job session. The `filter` parameter allows to choose a subset of the session jobs as return value. The semantics of the `filter` argument are explained in Section 5.5. If no job matches or the session has no jobs attached, the method MUST return an empty set. If `filter` is `UNSET`, all session jobs MUST be returned.

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

[45] 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).

(See footnote)[46]

8.2.5 getJobArray

This method returns the `JobArray` instance with the given ID. If the session does not / no longer contain the according job array, `InvalidArgumentException` SHALL be thrown.

(See footnote)[47]

8.2.6 runJob

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

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

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

8.2.7 runBulkJobs

The `runBulkJobs` method creates a set of parametric jobs, each with attributes as defined in the given job template instance. Each job in the set has the same attributes, except for the job template attributes that include the `JobTemplatePlaceholder::PARAMETRIC_INDEX` macro (see Section 4.4).

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

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

Jobs can determine their index number at run time by the mechanism described in Section 8.6.

The `maxParallel` parameter allows to specify how many of the bulk job's instances are allowed to run in parallel on the utilized resources. Implementations MAY consider this value if the DRM system supports such functionality, otherwise the parameter MUST be silently ignored. If given, the support MUST be expressed by the `DrmaaCapability::BULK_JOBS_MAXPARALLEL` capability flag (see Section 4.5). If the parameter value is `UNSET`, no limit SHOULD be applied.

[46] We are aware of the fact that the automated reaping of terminated jobs in some DRM systems might change this methods result. However, there was no way to demand some standardized behavior for that.

[47] June 2011 conf. call decided to not support JobArray filtering in the session at this point. The face-to-face meeting in June 2011 identified that DRM systems typically do not support the identification of bulk jobs in the system, so it would be hard to implement the according reporting function.

The `runBulkJobs` method returns a `JobArray` (see Section 8.5) instance that represents the set of `Job` objects created by the method call under a common array identity. For each of the jobs in the array, the same conditions as for the result of `runJob` SHOULD apply.

The largest (syntactically) allowed value for `endIndex` MUST be defined by the language binding.

Further restrictions on the maximum `endIndex` MAY be implied by the implementation.

(See footnote)[48]

8.2.8 waitAnyStarted / waitAnyTerminated

The `waitAnyStarted` method blocks until any of the jobs referenced in the `jobs` parameter entered one of the "Started" states. The `waitAnyTerminated` method blocks until any of the jobs referenced in the `jobs` parameter entered one of the "Terminated" states (see Section 8.1). If the input list contains jobs that are not part of the session, the method SHALL fail with an `InvalidArgumentException`.

The `timeout` argument specifies the desired waiting time for the state change. The constant value `INFINITE_TIME` MUST be supported to get an indefinite waiting time. The constant value `ZERO_TIME` MUST be supported to express that the method call SHALL return immediately. A number of seconds can be specified to indicate the maximum waiting time . If the method call returns because of timeout, an `TimeoutException` SHALL be raised.

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

(See footnote)[49]

8.3 DrmaaCallback Interface

The `DrmaaCallback` interface allows the DRMAA library or the DRM system to inform the application about relevant events in an asynchronous fashion. One expected use case is continuous monitoring of job state transitions. The implementation MAY decide to not deliver all events occurring in the DRM system. The support for such callback functionality is optional, indicated by the `DrmaaCallback::CALLBACK` flag. Also, all implementations MUST define the `DrmaaCallback` interface type as given in the language binding, regardless of the support for these functions.

```
interface DrmaaCallback {
  void notify(in DrmaaNotification notification);
};
```

[48] 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.

The May 4th 2011 conf call identified Grid Engine, Torque and LSF as the only systems having support for maxParallel. The feature was determined as critical enough for still adding it, therefore the ignorance rule and the MAY semantics are applied.

[49] 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 their intended long-blocking operation, the DRM system would no be able to offer any better (meaning much faster) implementation to be wrapped by DRMAA.

A section on synchronization of multi-threaded parallel wait calls was removed. This would complicate DRMAA implementations, since synchronization does not map to the obvious state polling approach. An optimization like this would be classically a task of application-oriented APIs - so, Andre has to solve it.

```
struct DrmaaNotification {
  DrmaaEvent event;
  string jobId;
  string sessionName;
  JobState jobState;
};

enum DrmaaEvent {
  NEW_STATE, MIGRATED, ATTRIBUTE_CHANGE
};
```

The application implements a `DrmaaCallback` interface as pre-condition for using this functionality. This interface is registered through the `SessionManager::registerEventNotification` method (see Section 7.1). On notification, the implementation or the DRM system pass a `DrmaaNotification` instance to the application. Implementations MAY extend this structure for further information (see Section 5). All given information SHOULD be valid at least at the time of notification generation.

The `DrmaaNotification::jobState` attribute expresses the state of the job at the time of notification generation.

The `DrmaaEvent` enumeration defines standard event types for notification:

NEW_STATE The job entered a new state, which is described in the `jobState` attribute.

MIGRATED The job was migrated to another execution host, and is now in the state described by `jobState`.

ATTRIBUTE_CHANGE A monitoring attribute of the job, such as the memory consumption, changed to a new value. The `jobState` attribute MAY have the value UNSET on this event.

DRMAA implementations SHOULD protect themselves from unexpected behavior of the called application. This includes indefinite delays or unexpected exceptions from the callee on notification processing. The implementation SHOULD prevent a nested callback at the time of occurrence, and MAY decide to deliver the according events at a later point in time.

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

(See footnote)[50]

8.4 Job Interface

Every job in the `JobSession` is represented by an own instance of the `Job` interface. It allows one to instruct the DRM system for a job status change, and to query the properties of the job in the DRM system. Implementations MAY provide `Job` objects for jobs created outside of a DRMAA session.

```
interface Job {
  readonly attribute string jobId;
```

[50] 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.

The `DrmaaNotification` structure intentionally avoids to reference a `Job` object - instead, all relevant lookup information (session name + job ID) is provided. This demands only non-interface data types to be understandable in the callback target. Also, it hopefully helps to support scalability of high-frequent event callbacks.

```
    readonly attribute string sessionName;
    readonly attribute JobTemplate jobTemplate;
    void suspend();
    void resume();
    void hold();
    void release();
    void terminate();
    JobState getState(out any jobSubState);
    JobInfo getInfo();
    Job waitStarted(in TimeAmount timeout);
    Job waitTerminated(in TimeAmount timeout);
  };
```

(See footnote)[51]

8.4.1 jobId

This attribute reports the stringified job identifier assigned by the DRM system. This method is expected to be used as fast alternative to the fetching of a complete `JobInfo` instance.

8.4.2 sessionName

This attribute reports the name of the `JobSession` that was used to create the job. If the session name cannot be determined, for example since the job was created outside of a DRMAA session, the attribute SHOULD be `UNSET`.

(See footnote)[52]

8.4.3 jobTemplate

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

For jobs created outside of a DRMAA session, implementations MUST also return a `JobTemplate` instance here, which MAY be empty or only partially filled.

[51] 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.

[52] June 29th 2011 conf call decided to return session names instead of session objects. This keeps the consistent approach that instantiated session objects represent a live 'connection' to the DRMS. Connecting to the referenced session is then a separate explicit step in the application. It also supports better that people create instances from jobs created outside of a DRMAA session.

8.4.4 suspend / resume / hold / release / terminate

The job control functions allow modifying the status of the single job in the DRM system, according to the state model presented in Section 8.1.

The `suspend` method triggers a transition from `RUNNING` to `SUSPENDED` state.

The `resume` method triggers a transition from `SUSPENDED` to `RUNNING` state.

The `hold` method triggers a transition from `QUEUED` to `QUEUED_HELD`, or from `REQUEUED` to `REQUEUED_HELD` state.

The `release` method triggers a transition from `QUEUED_HELD` to `QUEUED`, or from `REQUEUED_HELD` to `REQUEUED` state.

The `terminate` method triggers a transition from any of the "Started" states to one of the "Terminated" states.

If the job is in an inappropriate state for the particular method call, it MUST raise an `InvalidStateException`.

The methods SHOULD return after the action has been acknowledged by the DRM system, but MAY return before the action has been completed. Some DRMAA implementations MAY allow these methods to be used to control jobs submitted externally to the DRMAA session. Examples are jobs submitted by other DRMAA sessions, in other DRMAA implementations, or jobs submitted via native utilities. This behavior is implementation-specific.

8.4.5 getState

This method allows the application to get the current status of the job according to the DRMAA state model, together with an implementation specific sub state (see Section 8.1). It is intended as fast alternative to the fetching of a complete `JobInfo` instance. The timing conditions are described in Section 5.5.

(See footnote)[53]

8.4.6 getInfo

This method returns a `JobInfo` instance for the particular job, under the conditions described in Section 5.5.

8.4.7 waitStarted / waitTerminated

The `waitStarted` method blocks until the job entered one of the "Started" states. The `waitTerminated` method blocks until the job entered one of the "Terminated" states (see Section 8.1). All other behavior MUST work as described in Section 8.2.8.

8.5 JobArray Interface

An instance of the `JobInfo` interface represents a set of jobs created by one operation. In DRMAA, `JobArray` instances are only created by the `runBulkJobs` method (see Section 8.2). `JobArray` instances differ from the

[53] 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).

`JobList` data structure due to their potential for representing a DRM system concept, while `JobList` is a DRMAA-only concept realized by language binding support.

Implementations SHOULD realize the `JobArray` functionality as wrapper for DRM system job arrays, if available. If the DRM system has only single job support or incomplete job array support with respect to the DRMAA-provided functionality, implementations MUST offer the `JobArray` functionality on their own, for example based on looped activities with a list of jobs.

```
interface JobArray {
  readonly attribute string jobArrayId;
  readonly attribute JobList jobs;
  readonly attribute string sessionName;
  readonly attribute JobTemplate jobTemplate;
  void suspend();
  void resume();
  void hold();
  void release();
  void terminate();
};
```

(See footnote)[54]

8.5.1 jobArrayId

This attribute reports the stringified job identifier assigned to the job array by the DRM system. If the DRM system has no job array support, the implementation MUST generate a system-wide unique identifier for the result of the `runBulkJobs` method.

8.5.2 jobs

This attribute provides the list of jobs that are part of the job array, regardless of their state.

(See footnote)[55]

8.5.3 sessionName

This attribute states the name of the `JobSession` that was used to create the bulk job represented by this instance. If the session name cannot be determined, for example since the bulk job was created outside of a DRMAA session, the attribute SHOULD have an `UNSET` value.

(See footnote)[56]

[54] We are aware of the fact that some systems (e.g., LSF at the time of writing) do not support all DRMAA control methods offered for job arrays. Since we intended to avoid optional DRMAA methods wherever we could, the text here mandates the implementation to simulate the array 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.

[55] 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.

The dissappearance of terminated jobs is intentionally not specified (see discussion above for `getJobs`).

[56] June 29th 2011 conf call decided to return session names instead of session objects. This keeps the consistent approach that instantiated session objects represent a live 'connection' to the DRMS. Connecting to the referenced session is then a separate explicit step in the application. It also supports better that people create instances from bulk jobs created outside of a DRMAA session.

8.5.4 jobTemplate

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

(See footnote)[57]

8.5.5 suspend / resume / hold / release / terminate

The job control functions allow modifying the status of the job array in the DRM system, with the same semantic as in the `Job` interface (see Section 8.4.4). If one of the jobs in the array is in an inappropriate state for the particular method, the method MAY raise an `InvalidStateException`.

The methods SHOULD return after the action has been acknowledged by the DRM system for all jobs in the array, but MAY return before the action has been completed for all of the jobs. Some DRMAA implementations MAY allow this method to be used to control job arrays created externally to the DRMAA session. This behavior is implementation-specific.

(See footnote)[58]

8.6 The DRMAA_INDEX_VAR environment variable

DRMAA implementations SHOULD implicitly set an environment variable with the name `DRMAA_INDEX_VAR` for each job submitted to the DRM system.

An expected implementation strategy would be the transparent addition of an environment variable specification in the job submission. Such a definition SHOULD NOT be visible for the application as part of the job template. If the application defines its own `DRMAA_INDEX_VAR` environment variable, it SHOULD override the implementation-defined value.

The environment variable MUST contain the name of the DRM system environment variable that holds the parametric job index. Examples are `TASK_ID` in GridEngine, `PBS_ARRAYID` in Torque, or `LSB_JOBINDEX` in LSF. By doing an indirect fetching of the environment variable value, jobs are enabled to get their own parametric index regardless of the DRM system type. For DRM systems that do not provide such an environment variable, `DRMAA_INDEX_VAR` SHOULD not be set.

9 Working with Advance Reservation

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

DRMAA implementations for a DRM system that does not support advance reservation MUST still implement the described interfaces, in order to keep source code portability for DRMAA-based applications. All methods related to advance reservation MUST raise an `UnsupportedOperationExeption` in this case. Support for advance reservation is expressed by the `DrmaaCapability::ADVANCE_RESERVATION` flag (see Section 4.5).

[57] 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).

[58] We were asked to make explicit that some of these functions may not be atomic. However, this holds for most methods, and is not supported to be a part of the API standard.

9.1 ReservationSession Interface

Every `ReservationSession` instance acts as container for advance reservations in the DRM system. Every `Reservation` instance SHALL belong only to one `ReservationSession` instance.

```
interface ReservationSession {
  readonly attribute string contact;
  readonly attribute string sessionName;
  Reservation getReservation(in string reservationId);
  Reservation requestReservation(in ReservationTemplate reservationTemplate);
  ReservationList getReservations();
};
```

9.1.1 contact

This attribute reports the `contact` value that was used in the `createReservationSession` call for this instance (see Section 7.1). If no value was originally provided, the default contact string from the implementation MUST be returned. This attribute is read-only.

9.1.2 sessionName

This attribute reports the name of the session that was used for creating or opening this `Reservation` instance (see Section 7.1). This attribute is read-only.

9.1.3 getReservation

This method returns the `Reservation` instance that has the given `reservationId`. Implementations MAY support the access to reservations created outside of a DRMAA session scope, under the same regularities as for the `MonitoringSession::getAllReservations` method (see Section 10.1.1). If no reservation matches, the method SHALL raise an `InvalidArgumentException`. Time-dependent effects of this method are implementation-specific.

9.1.4 requestReservation

The `requestReservation` method SHALL request an advance reservation in the DRM system as described by the `ReservationTemplate`. On a successful reservation, the method returns a `Reservation` instance that represents the advance reservation in the underlying DRM system.

If the current user is not authorized to create reservations, `DeniedByDrmsException` SHALL be raised. If the reservation cannot be performed by the DRM system due to invalid `ReservationTemplate` attributes, or if the demanded combination of resources is not available, `InvalidArgumentException` SHALL be raised. The exception SHOULD provide further details about the rejection cause in the extended error information (see Section 6).

Some of the requested conditions might be not fulfilled after the reservation was successfully created, for example due to execution host outages. In this case, the reservation itself SHOULD remain valid. A job using such a reservation may spend additional time in one of the non-`RUNNING` states. In this case, the `JobInfo::jobSubState` information SHOULD inform about this situation.

(See footnote)[59]

9.1.5 getReservations

This method returns the list of reservations successfully created so far in this session, regardless of their start and end time. The list of `Reservation` instances is only cleared in conjunction with the destruction of the actual session instance through `SessionManager::destroyReservationSession` (see Section 7.1).

9.2 Reservation Interface

The `Reservation` interface represents attributes and methods available for an advance reservation successfully created in the DRM system. Implementations MAY offer `Reservation` instances for advance reservations created outside of a DRMAA session.

```
interface Reservation {
  readonly attribute string reservationId;
  readonly attribute string sessionName;
  readonly attribute ReservationTemplate reservationTemplate;
  ReservationInfo getInfo();
  void terminate();
};
```

9.2.1 reservationId

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

9.2.2 sessionName

This attribute states the name of the `ReservationSession` that was used to create the advance reservation instance. If the session name cannot be determined, for example since the reservation was created outside of a DRMAA session, the attribute SHOULD have an `UNSET` value.

(See footnote)[60]

9.2.3 reservationTemplate

This attribute provides a reference to a `ReservationTemplate` instance that has equal values to the one that was used to create this reservation. For reservations created outside of a DRMAA session, implementations MUST also return a `ReservationTemplate` instance, which MAY be empty or only partially filled.

[59] In DRMAA 2.0 we do not have an explicit state model for advance reservations, as the reservation state can be easily deducted by comparing current time with reservation start and end time. For this reason, we use the subState approach for informing the user about the described situation.

[60] June 29th 2011 conf call decided to return session names instead of session objects. This keeps the consistent approach that instantiated session objects represent a live 'connection' to the DRMS. Connecting to the referenced session is then a separate explicit step in the application. It also supports better that people create instances from reservation created outside of a DRMAA session.

9.2.4 getInfo

This method returns a `ReservationInfo` instance under the conditions described in Section 5.6. The method SHOULD throw `InvalidArgumentException` if the reservation is already expired (i.e., its end time passed), or if it was terminated before.

9.2.5 terminate

This method terminates the advance reservation represented by this `Reservation` instance. All jobs submitted with a reference to this reservation SHOULD be terminated by the DRM system or the implementation, regardless of their current state.

10 Monitoring the DRM System

The monitoring support in DRMAA focusses on the investigation of resources and on global data maintained by the DRM system. Session-related information is available from the `JobSession` and `ReservationSession` instances, respectively.

10.1 MonitoringSession Interface

The `MonitoringSession` interface provides a set of stateless methods for fetching information about the DRM system and the DRMAA implementation itself.

```
interface MonitoringSession {
  ReservationList getAllReservations();
  JobList getAllJobs(in JobInfo filter);
  QueueInfoList getAllQueues(in StringList names);
  MachineInfoList getAllMachines(in StringList names);
};
```

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 report queues that are usable or generally accessible for the DRMAA application and the user performing the query.

Because of cases where such a list reduction may demand excessive overhead in the DRMAA implementation, an unreduced or only partially reduced result MAY also be returned. The behavior of the DRMAA implementation in this regard should be clearly documented. In all cases, the list items MUST be valid input for job submission or advance reservation through the DRMAA API, but MAY lead to later exceptions.

10.1.1 getAllReservations

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

The DRM system or the DRMAA implementation is at liberty to restrict the set of returned reservations based on site or system policies, such as security settings or scheduler load restrictions. The returned list MAY contain reservations that were created by other users. It MAY also contain reservations that are not usable for the user.

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

10.1.2 getAllJobs

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

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

The method supports a `filter` argument for fetching only a subset of the job information available. Both the return value semantics and the filter semantics SHOULD be similar to the ones described for the `JobSession::getJobs` method (see Section 8.2).

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

(See footnote)[61]

10.1.3 getAllQueues

This method returns a list of queues available for job submission in the DRM system. The names from all `QueueInfo` instances in this list SHOULD be a valid input for the `JobTemplate::queueName` attribute (see Section 5.7.14). The result can be an empty list or might be incomplete, based on queue, host, or system policies. It might also contain queues that are not accessible for the user at job submission time because of queue configuration limits.

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

10.1.4 getAllMachines

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

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

[61] 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)

11 Complete DRMAA IDL Specification

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

```
module DRMAA2 {

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

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

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

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

  enum JobTemplatePlaceholder {
    HOME_DIRECTORY,WORKING_DIRECTORY,PARAMETRIC_INDEX };

  enum DrmaaEvent {
    NEW_STATE, MIGRATED, ATTRIBUTE_CHANGE
  };

  enum DrmaaCapability {
    ADVANCE_RESERVATION, RESERVE_SLOTS, CALLBACK, BULK_JOBS_MAXPARALLEL,
     JT_EMAIL, JT_STAGING, JT_DEADLINE, JT_MAXSLOTS, JT_ACCOUNTINGID,
    RT_STARTNOW, RT_DURATION, RT_MACHINEOS, RT_MACHINEARCH
  };

  typedef sequence<string> OrderedStringList;
  typedef sequence<string> StringList;
  typedef sequence<Job> JobList;
  typedef sequence<QueueInfo> QueueInfoList;
  typedef sequence<MachineInfo> MachineInfoList;
  typedef sequence<SlotInfo> OrderedSlotInfoList;
  typedef sequence<Reservation> ReservationList;
  typedef sequence< sequence<string,2> > Dictionary;
  typedef string AbsoluteTime;
  typedef long long TimeAmount;
```

```
1638    native ZERO_TIME;
1639    native INFINITE_TIME;
1640    native NOW;

1641    struct JobInfo {
1642      string jobId;
1643      long exitStatus;
1644      string terminatingSignal;
1645      string annotation;
1646      JobState jobState;
1647      any jobSubState;
1648      OrderedSlotInfoList allocatedMachines;
1649      string submissionMachine;
1650      string jobOwner;
1651      long slots;
1652      string queueName;
1653      TimeAmount wallclockTime;
1654      long cpuTime;
1655      AbsoluteTime submissionTime;
1656      AbsoluteTime dispatchTime;
1657      AbsoluteTime finishTime;
1658    };

1659    struct ReservationInfo {
1660      string reservationId;
1661      string reservationName;
1662      AbsoluteTime reservedStartTime;
1663      AbsoluteTime reservedEndTime;
1664      StringList usersACL;
1665      long reservedSlots;
1666      OrderedSlotInfoList reservedMachines;
1667    };

1668    struct JobTemplate {
1669      string remoteCommand;
1670      OrderedStringList args;
1671      boolean submitAsHold;
1672      boolean rerunnable;
1673      Dictionary jobEnvironment;
1674      string workingDirectory;
1675      string jobCategory;
1676      StringList email;
1677      boolean emailOnStarted;
1678      boolean emailOnTerminated;
1679      string jobName;
1680      string inputPath;
1681      string outputPath;
1682      string errorPath;
```

```
1683      boolean joinFiles;
1684      string reservationId;
1685      string queueName;
1686      long minSlots;
1687      long maxSlots;
1688      long priority;
1689      OrderedStringList candidateMachines;
1690      long minPhysMemory;
1691      OperatingSystem machineOS;
1692      CpuArchitecture machineArch;
1693      AbsoluteTime startTime;
1694      AbsoluteTime deadlineTime;
1695      Dictionary stageInFiles;
1696      Dictionary stageOutFiles;
1697      Dictionary resourceLimits;
1698      string accountingId;
1699    };

1700    struct ReservationTemplate {
1701      string reservationName;
1702      AbsoluteTime startTime;
1703      AbsoluteTime endTime;
1704      TimeAmount duration;
1705      long minSlots;
1706      long maxSlots;
1707      string jobCategory;
1708      StringList usersACL;
1709      OrderedStringList candidateMachines;
1710      long minPhysMemory;
1711      OperatingSystem machineOS;
1712      CpuArchitecture machineArch;
1713    };

1714    struct DrmaaNotification {
1715      DrmaaEvent event;
1716      string jobId;
1717      string sessionName;
1718      JobState jobState;
1719    };

1720    struct QueueInfo {
1721      string name;
1722    };

1723    struct Version {
1724      string major;
1725      string minor;
1726    };
```

```
1727  struct MachineInfo {
1728    string name;
1729    boolean available;
1730    long sockets;
1731    long coresPerSocket;
1732    long threadsPerCore;
1733    double load;
1734    long physMemory;
1735    long virtMemory;
1736    OperatingSystem machineOS;
1737    Version machineOSVersion;
1738    CpuArchitecture machineArch;
1739  };

1740  struct SlotInfo {
1741    string machineName;
1742    string slots;
1743  };

1744  exception DeniedByDrmsException {string message;};
1745  exception DrmCommunicationException {string message;};
1746  exception TryLaterException {string message;};
1747  exception SessionManagementException {string message;};
1748  exception TimeoutException {string message;};
1749  exception InternalException {string message;};
1750  exception InvalidArgumentException {string message;};
1751  exception InvalidSessionException {string message;};
1752  exception InvalidStateException {string message;};
1753  exception OutOfResourceException {string message;};
1754  exception UnsupportedAttributeException {string message;};
1755  exception UnsupportedOperationException {string message;};

1756  interface DrmaaReflective {
1757    readonly attribute StringList jobTemplateImplSpec;
1758    readonly attribute StringList jobInfoImplSpec;
1759    readonly attribute StringList reservationTemplateImplSpec;
1760    readonly attribute StringList reservationInfoImplSpec;
1761    readonly attribute StringList queueInfoImplSpec;
1762    readonly attribute StringList machineInfoImplSpec;
1763    readonly attribute StringList notificationImplSpec;
1764
1765    string getInstanceValue(in any instance, in string name);
1766    void setInstanceValue(in any instance, in string name, in string value);
1767    string describeAttribute(in any instance, in string name);
1768  };

1769  interface DrmaaCallback {
1770    void notify(in DrmaaNotification notification);
1771  };
```

```
1772   interface ReservationSession {
1773     readonly attribute string contact;
1774     readonly attribute string sessionName;
1775     Reservation getReservation(in string reservationId);
1776     Reservation requestReservation(in ReservationTemplate reservationTemplate);
1777     ReservationList getReservations();
1778   };

1779   interface Reservation {
1780     readonly attribute string reservationId;
1781     readonly attribute string sessionName;
1782     readonly attribute ReservationTemplate reservationTemplate;
1783     ReservationInfo getInfo();
1784     void terminate();
1785   };

1786   interface JobArray {
1787     readonly attribute string jobArrayId;
1788     readonly attribute JobList jobs;
1789     readonly attribute string sessionName;
1790     readonly attribute JobTemplate jobTemplate;
1791     void suspend();
1792     void resume();
1793     void hold();
1794     void release();
1795     void terminate();
1796   };

1797   interface JobSession {
1798     readonly attribute string contact;
1799     readonly attribute string sessionName;
1800     readonly attribute StringList jobCategories;
1801     JobList getJobs(in JobInfo filter);
1802     JobArray getJobArray(in string jobArrayId);
1803     Job runJob(in JobTemplate jobTemplate);
1804     JobArray runBulkJobs(
1805         in JobTemplate jobTemplate,
1806         in long beginIndex,
1807         in long endIndex,
1808         in long step,
1809         in long maxParallel);
1810     Job waitAnyStarted(in JobList jobs, in TimeAmount timeout);
1811     Job waitAnyTerminated(in JobList jobs, in TimeAmount timeout);
1812   };

1813   interface Job {
1814     readonly attribute string jobId;
1815     readonly attribute string sessionName;
```

```
    readonly attribute JobTemplate jobTemplate;
    void suspend();
    void resume();
    void hold();
    void release();
    void terminate();
    JobState getState(out any jobSubState);
    JobInfo getInfo();
    Job waitStarted(in TimeAmount timeout);
    Job waitTerminated(in TimeAmount timeout);
  };

  interface MonitoringSession {
    ReservationList getAllReservations();
    JobList getAllJobs(in JobInfo filter);
    QueueInfoList getAllQueues(in StringList names);
    MachineInfoList getAllMachines(in StringList names);
  };

  interface SessionManager{
    readonly attribute string drmsName;
    readonly attribute Version drmsVersion;
    readonly attribute string drmaaName;
    readonly attribute Version drmaaVersion;
    boolean supports(in DrmaaCapability capability);
    JobSession createJobSession(in string sessionName,
                                in string contact);
    ReservationSession createReservationSession(in string sessionName,
                                                in string contact);
    JobSession openJobSession(in string sessionName);
    ReservationSession openReservationSession(in string sessionName);
    MonitoringSession openMonitoringSession (in string contact);
    void closeJobSession(in JobSession s);
    void closeReservationSession(in ReservationSession s);
    void closeMonitoringSession(in MonitoringSession s);
    void destroyJobSession(in string sessionName);
    void destroyReservationSession(in string sessionName);
    StringList getJobSessionNames();
    StringList getReservationSessionNames();
    void registerEventNotification(in DrmaaCallback callback);
  };

};
```

12 Security Considerations

The DRMAA API does not specifically assume the existence of a particular security infrastructure in the DRM system. The scheduling scenario described herein presumes that security is handled at the point of

interaction with the DRM system. It is assumed that credentials owned by the application using the API are in effect for the DRMAA implementation too, so that it acts as stakeholder for the application.

An authorized but malicious user could use a DRMAA implementation or a DRMAA-enabled application to saturate a DRM system with a flood of requests. Unfortunately for the DRM system, this case is not distinguishable from the case of an authorized good-natured user who has many jobs to be processed. For temporary load defense, implementations SHOULD utilize the `TryLaterException`, if possible. In case of permanent issues, the implementation SHOULD raise the `DeniedByDrmsException`.

DRMAA implementers SHOULD guard their product against buffer overflows that can be exploited through DRMAA enabled interactive applications or portals. Implementations of the DRMAA API will most likely require a network to coordinate subordinate DRM system requests. However, the API makes no assumptions about the security posture provided by the networking environment. Therefore, application developers SHOULD also consider the security implications of "on-the-wire" communications in this case.

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

13 Contributors

The DRMAA working group is grateful to numerous colleagues for support and discussions on the topics covered in this document, in particular (in alphabetical order, with apologies to anybody we have missed):

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

Special thanks must go to Andre Merzky, who participated as SAGA working group representative in numerous DRMAA events.

This specification was developed by the following core members of the DRMAA working group at the Open Grid Forum:

Roger Brobst
Cadence Design Systems, Inc.
555 River Oaks Parkway
San Jose, CA 95134
United States
Email: rbrobst@cadence.com

Daniel Gruber
Univa GmbH
c/o Rüter und Partner
Prielmayerstr. 3 80335 München
Germany

Email: dgruber@univa.com

Mariusz Mamoński
Poznań Supercomputing and Networking Center
ul. Noskowskiego 10
61-704 Poznań
Poland
Email: mamonski@man.poznan.pl

Daniel Templeton
Cloudera Inc.
210 Portage Avenue
Palo Alto, CA 94306
United States
Email: daniel@cloudera.com

Peter Tröger (Corresponding Author)
Hasso-Plattner-Institute at University of Potsdam
Prof.-Dr.-Helmert-Str. 2-3
14482 Potsdam
Germany
Email: peter@troeger.eu

14 Intellectual Property Statement

15 Disclaimer

16 Full Copyright Notice

17 References

[1] Sergio Andreozzi, Stephen Burke, Felix Ehm, Laurence Field, Gerson Galang, Balazs Konya, Maarten Litmaath, Paul Millar, and JP Navarro. GLUE Specification v. 2.0 (GFD-R-P.147), mar 2009.

[2] Scott Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119 (Best Current Practice), March 1997. URL `http://tools.ietf.org/html/rfc2119`.

[3] I. Foster, A. Grimshaw, P. Lane, W. Lee, M. Morgan, S. Newhouse, S. Pickles, D. Pulsipher, C. Smith, and M. Theimer. OGSA Basic Execution Service v1.0 (GFD-R.108), nov 2008.

[4] Tom Goodale, Shantenu Jha, Hartmut Kaiser, Thilo Kielmann, Pascal Kleijer, Andre Merzky, John Shalf, and Christopher Smith. A Simple API for Grid Applications (SAGA) Version 1.1 (GFD-R-P.90), jan 2008.

[5] Object Management Group. Common Object Request Broker Architecture (CORBA) Specification, Version 3.1. http://www.omg.org/spec/CORBA/3.1/Interfaces/PDF, jan 2008.

[6] The IEEE and The Open Group. The Open Group Base Specifications Issue 6 IEEE Std 1003.1. http://www.opengroup.org/onlinepubs/000095399/utilities/ulimit.html.

[7] Distributed Management Task Force (DMTF) Inc. CIM System Model White Paper CIM Version 2.7, jun 2003.

[8] Hrabri Rajic, Roger Brobst, Waiman Chan, Fritz Ferstl, Jeff Gardiner, Andreas Haas, Bill Nitzberg, Daniel Templeton, John Tollefsrud, and Peter Tröger. Distributed Resource Management Application API Specification 1.0 (GFD-R.022), aug 2007.

[9] Hrabri Rajic, Roger Brobst, Waiman Chan, Fritz Ferstl, Jeff Gardiner, Andreas Haas, Bill Nitzberg, Daniel Templeton, John Tollefsrud, and Peter Tröger. Distributed Resource Management Application API Specification 1.0 (GWD-R.133), jun 2008.

[10] Peter Tröger, Daniel Templeton, Roger Brobst, Andreas Haas, and Hrabri Rajic. Distributed Resource Management Application API 1.0 - IDL Specification (GFD-R-P.130), apr 2008.

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