

# Information Modeling in OGSA Position Paper

## Status of This Document

This document provides information to the Grid community on the proposed direction for information modeling of OGSA resources. It does not define any standards or technical recommendations. Distribution is unlimited.

## Copyright Notice

Copyright © Global Grid Forum (2006). All Rights Reserved.

## Trademarks

OGSA is a trademark of the Global Grid Forum.

## Abstract

Resources in the grid need to advertise their capabilities and activities in the grid need to consume those resources. This position paper proposes a way to model resources' capabilities and requirements in the OGSA. It builds on the wealth of existing systems management information already modeled and instantiated in systems today. An example is included.

Contents

1.	Introduction .....	3
2.	Notational Conventions.....	3
3.	Model for OGSA Resources .....	3
4.	Example .....	5
4.1	Systems management instance information .....	5
4.2	Advertised Capabilities.....	6
4.3	Activity Requirements .....	7
4.4	Matching requirements with advertised capabilities.....	7
4.5	XML .....	7
4.6	UML.....	10
5.	Security Considerations.....	10
6.	Contributors .....	10
7.	Glossary.....	10
8.	Intellectual Property Statement.....	10
9.	Disclaimer .....	10
10.	Full Copyright Notice .....	10
11.	References.....	11

## 1. Introduction

Resources in the grid need to advertise their capabilities and activities in the grid need to consume those resources. This position paper proposes a way to model resources' capabilities and requirements in the OGSA. It builds on the wealth of existing systems management information already modeled and instantiated in systems today. An XML example to advertise resource capabilities and express activity requirements with a mapping to systems management information is included. A UML rendering will be added later.

## 2. Notational Conventions

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 [BRADNER1].

## 3. Model for OGSA Resources

Resources are typically modeled for three reasons: (1) to manage resource information in the system, (2) to advertise the capabilities of resources in the system, and (3) to express a set of requirements such as those needed by a job that is to run in the system.

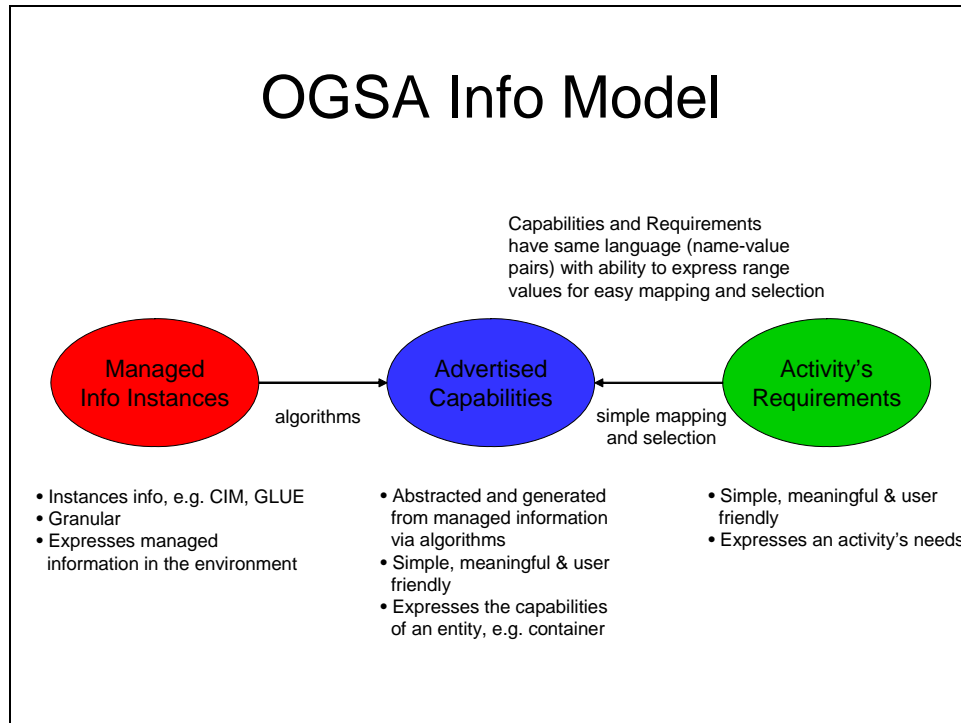
Today, a system's resources are typically modeled so systems (including network) management applications can manage (provision, configure, manage, and de-provision) those resources. Information modeled for systems management applications tends to be very granular and detailed. Examples of models of managed information are GLUE and CIM. Examples of managed information are:

- The processor(s) for each computer with attributes like ProcessorFamily, Version, MaxClockSpeed, CurrentClockSpeed, DataWidth, AddressWidth, Load, CPUStatus, ExternalBusClockSpeed.
- The operating system for each computer (node) and processor with attributes like OSType, Version, LastBootUpTime, LocalDateTime, CurrentTimeZone, NumberOfProcesses, MaxNumberOfProcesses, MaxProcessesPerUser, TotalSwapSpaceSize, TotalVirtualMemorySize, FreeVirtualMemory, FreePhysicalMemory.
- Each computer (node) with attributes like Name and NameFormat (e.g. DNS style hostname), Load.
- Relationships between this information to express that a given computer has n processors, has 1 or more installed operating systems, and one of those installed operating systems is actually the running operating system.
- Values of attributes are typically single values and specific

In the Grid, jobs (or in the more general sense – activities) need to run where its resource requirements are satisfied or can be provisioned to satisfy those requirements. To make that determination, systems in the grid need to advertise the capabilities of its resources. Activities that are run in the grid may be parallelized applications and run on multiple nodes in the grid. So requirements may need to be expressed for multiple resources of the same type (e.g. Activity A needs a minimum of three processors with n CPU seconds (or m seconds of wall clock time) available) as well as total CPU seconds or wall clock time for that activity. Therefore, these requirements and capabilities need to be expressed in terms that are meaningful to the activity's writer (e.g. user friendly). Values may be specific or may be a range of values.

How managed information is modeled for systems management and how resources are modeled in the Grid to express requirements and capabilities leads one to conclude that the systems management type of model for information modeling of resources is not what's needed in the

OGSA. Instead, a higher level, less granular, and more user friendly information model is needed to express an activity's requirements and to advertise a resource's capabilities. However, since managed information exists in systems today, it is beneficial to determine a resource's advertised capabilities from that existing managed information. Figure 1 proposes the relationship between a system's managed information and a Grid's need to advertise resource capabilities and an activity's requirements to consume resources.



**Figure 1. Model Overview for OGSA Resources**

Capabilities are abstracted and generated from managed information via algorithms. These capabilities are updated as the managed information in the system changes. An activity's requirements are mapped against the advertised capabilities to determine which resources it will consume and hence where that activity will execute. These capabilities and requirements have the same language for simple mapping and selection of resources. The language is expressed in XML as name-value pairs where values can be expressed as ranges as well as single values.

There are two categories of capabilities that can be advertised. The first category is capabilities that are typically consumed by an activity such as an application where that activity's requirements are expressed in a JSDL document. These capabilities typically have static values, e.g. total physical memory or operating system type. The second category is capabilities that are typically consumed by some system component such as a job manager. These capabilities may be static or may change as activities are executed, e.g. available physical memory, load.

## 4. Example

This example shows the managed information in a system, the capabilities advertised from that system that an activity can consume, and the activity's requirements that are matched against the advertised capabilities. In this example, the activity is some application with a JSDL document, and the capabilities advertised are those that may be typically matched against a JSDL document. Capabilities that a system component may want to consume are not advertised in this example.

### 4.1 Systems management instance information

The system management information consists of the following:

Computer A has DNS-style hostname, Name=computerA.acme.com. It is classified as a non-dedicated system (ability to run applications, store data, act as a router or gateway, etc) with attribute Dedicated=0. Likewise, Computer C has a DNS-style hostname, Name=computerC.acme.com and the same characteristics as Computer A.

Computer B has DNS-style hostname, Name=computerB.acme.com. It is classified as a router with attribute Dedicated=4.

Computer A has 1 processor and Computer C has 2 processors. Each processor instance has the following managed information:

- ProcessorFamily=118 (AMD Athlon 64) for Computer C and ProcessorFamily=185 (Intel Pentium M) for Computer A
- UpgradeMethod=16 (Socket 754)
- MaxClockSpeed=3200 MHz
- CurrentClockSpeed=3000 MHz
- DataWidth=8 bit
- AddressWidth=32 bit
- Load=50% (Computer A – processor 1), Load=73% (Computer C – processor 1), Load=30% (Computer C – processor 2)
- Version=3.1.c
- ID=975403-2 (Computer A – processor 1), ID=984567-3 (Computer C – processor 1), ID=967021-8 (Computer C – processor 2)
- CPUStatus=1 (CPU enabled)
- ExternalBusClockSpeed=800 MHz (front side bus)
- Characteristics=2 (64 bit capable)

Computers A and C each have operating systems Windows XP Service Pack 2 and Linux installed. Computers A and C are currently configured and running the Linux operating system. Managed information for the current running Linux operating system is:

- OSType=36 (Linux)
- Version=2.1.f (major.minor.revision)
- LastBootUpTime=20060501
- CurrentTimeZone=-6 (CentralUS)
- NumberOfUsers=4
- NumberOfProcesses=67
- MaxNumberOfProcesses=0 (no maximum)
- TotalVirtualMemorySize=amt of total RAM + amt of paging space (SizeStoredInPagingFiles)=12000000 KB
- FreeVirtualMemory=amt of free RAM + amt of free paging space (FreePhysicalMemory + FreeSpaceInPagingFiles)=4100000 KB
- FreePhysicalMemory=100000 KB
- TotalVisibleMemorySize (amount of physical memory allocated to this OS)=4000000 KB (4GB)
- SizeStoredInPagingFiles=8000000 KB
- FreeSpaceInPagingFiles=4000000 KB
- MaxProcessMemorySize (max bytes allocated to a process)=1500000 KB

- MaxProcessesPerUser=32

Computer B has 1 processor of type Intel running Windows XP operating system. The details for Computer B are not stated here because that computer is classified as a router and will not run activities. So for this example, the advertised capabilities are minimal (in reality, it does advertise more than the minimal capabilities listed in this example).

Some values of this managed information are static for the life of the instance (e.g. hostname, operating system type), and some values change over time (e.g. processor load, free memory (physical, virtual, page space)).

This instance managed information will not be presented in XML in this example because different models have different XML representations. The representation of this managed information is not the focus of this position paper – what matters in this example is that this managed information exists and can be used to generate a resource's advertised capabilities.

## 4.2 Advertised Capabilities

Computers A and C advertise their capabilities as a stated below. These capabilities are algorithmic generated from granular detailed managed information. Some capabilities may map one-to-one with the managed information, e.g. CPUSpeed is the same as MaxClockSpeed (only a difference in name for usability, in this case; in other cases the name may be preserved). Some capabilities are computed from one or more pieces of managed information, e.g. PhysicalMemory for a machine with 2 processors that is advertised as available for 1 or more activities is the computation (TotalVisibleMemorySize for processor 1 + TotalVisibleMemorySize for processor 2).

Computer B is dedicated to routing and will not execute work (activities). Computers A and C can execute work (activities) because there is not advertised work restriction.

- Hostname=computerA.acme.com
  - Processor
    - Type=Intel
    - CPUSpeed=3200
  - OS=Linux
    - PhysicalMemory=3000000
    - VirtualMemory=12000000
    - MaxProcessesPerUser=32
- Hostname=computerC.acme.com
  - Processor
    - CPUCount=2
    - Type=AMD
    - CPUSpeed=3200
  - OS=Linux
    - PhysicalMemory=3000000
    - VirtualMemory=12000000
    - MaxProcessesPerUser=32

Computer B advertises (minimally for brevity in this example) its capabilities as:

- Hostname=computerB.acme.com
  - Dedicated=router
  - Processor
    - Type=Intel
  - OS=WindowsXP
    - PhysicalMemory=8000000

- VirtualMemory=20000000

### 4.3 Activity Requirements

Activities 1 and 2 list their resource requirements needed to execute. These resource requirements are mapped against advertised capabilities to determine which resources this activity will consume and hence where the job will run.

- Activity1
  - Processor=AMD or Intel
    - CPUSpeed >2500 and <3500
  - OS=Linux
    - PhysicalMemory=2000000
- Activity2
  - Processor=Intel
    - CPUCount<=2
    - CPUSpeed>3000
  - OS=WindowsXP

### 4.4 Matching requirements with advertised capabilities

A system component, such as a job manager, will match an activity's requirements against resources' advertised capabilities to determine the possible place(s) that activity may execute.

Activity 1 can execute on either Computer A or Computer C.

Activity 2 can execute on Computer A if that computer is re-provisioned to run the WindowsXP operating system.

### 4.5 XML

**NOTE: The XML is not complete, and most likely there are errors. Two XML formats are provided to drive the discussion of which is the preferred method of typing and extensibility. We need to agree on a single format.**

#### 4.5.1 Using XML typing and extensibility

Advertised capabilities for Computer A

```
<node>
  <name>computerA.acme.com</name>
  <processor>
    <type>Intel</type>
    <CPUSpeed>3200</CPUSpeed>
  </processor>
  <OS>
    <type>Linux</type>
    <physicalMemory>3000000</physicalMemory>
    <virtualMemory>12000000</virtualMemory>
    <maxProcessesPerUser>32</maxProcessesPerLimit>
  </OS>
</node>
```

Advertised capabilities for Computer C

```
<node>
```

```
<name>computerC.acme.com</name>
<processor>
  <CPUCount>2</CPUCount>
  <type>AMD</type>
  <CPUSpeed>3200</CPUSpeed>
</processor>
<OS>
  <type>Linux</type>
  <physicalMemory>3000000</physicalMemory>
  <virtualMemory>12000000</virtualMemory>
  <maxProcessesPerUser>32</maxProcessesPerUser>
</OS>
</node>
```

#### Advertised capabilities for Computer B

```
<node>
  <name>computerB.acme.com</name>
  <dedicated>router</dedicated>
  <processor>
    <type>Intel</type>
  </processor>
  <OS>
    <type>WindowsXP</type>
    <physicalMemory>8000000</physicalMemory>
    <virtualMemory>12000000</virtualMemory>
  </OS>
</node>
```

#### Requirements for Activity 1

```
<activity>
  <name>activity1</name>
  <processor>
    <type> AMD or Intel </type>
    <CPUSpeed> >2500 & <3500 </CPUSpeed>
  </processor>
  <OS>
    <type>Linux</type>
    <physicalMemory>2000000</physicalMemory>
  </OS>
</activity>
```

#### Requirements for Activity 2

```
<activity>
  <name>activity1</name>
  <processor>
    <type>Intel</type>
    <CPUCount> <=2 </CPUCount>
    <CPUSpeed> >3000 </CPUSpeed>
  </processor>
  <OS>
    <type>WindowsXP</type>
  </OS>
</activity>
```

## 4.5.2 Using typing outside of XML

### Advertised capabilities for Computer A

```
<resource Name=computerA.acme.com>
  <resource Name=Processor>
    <capability Name=Type Value=Intel>
    <capability Name=CPUSpeed Value=3200>
  </resource>
  <resource Name=OS>
    <capability Name=Type Value=Linux>
    <capability Name=PhysicalMemory Value=3000000>
    <capability Name=MaxProcessesPerUser Value=32>
  </resource>
</resource>
```

### Advertised capabilities for Computer C

```
<resource Name=computerC.acme.com>
  <resource Name=Processor>
    <capability Name=CPUCount Value=2>
    <capability Name=Type Value=AMD>
    <capability Name=CPUSpeed Value=3200>
  </resource>
  <resource Name=OS>
    <capability Name=Type Value=Linux>
    <capability Name=PhysicalMemory Value=3000000>
    <capability Name=MaxProcessesPerUser Value=32>
  </resource>
</resource>
```

### Advertised capabilities for Computer B

```
<resource Name=computerB.acme.com>
  <capability Name=Dedicated Value=router>
  <resource Name=Processor>
    <capability Name=Type Value=Intel>
  </resource>
  <resource Name=OS>
    <capability Name=Type Value=WindowsXP>
    <capability Name=PhysicalMemory Value=8000000>
    <capability Name=VirtualMemory Value=20000000>
  </resource>
</resource>
```

### Requirements for Activity 1

```
<resource Name=activity1>
  <resource Name=processor>
    <capability Name=Type Value=AMD or Value=Intel>
    <capability Name=CPUSpeed Value>2500 & Value<3500 >
  </resource>
  <resource Name=OS>
    <capability Name=Type Value=Linux>
    <capability Name=PhysicalMemory Value=2000000>
  </resource>
</resource>
```

## Requirements for Activity 2

```
<resource Name=activity2>
  <resource Name=processor>
    <capability Name=Type Value=Intel>
    <capability Name=CPUCount Value<=2>
    <capability Name=CPUSpeed Value>3000>
  </resource>
  <resource Name=OS>
    <capability Name=Type Value=WindowsXP>
  </resource>
</resource>
```

## 4.6 UML

TBD.

## 5. Security Considerations

Use of managed information in models is subject to an activity's authentication and authorization to use those resources. That authentication and authorization is outside the scope of this proposal. It is the responsibility of the security subsystem and protocols.

## 6. Contributors

TBD.

## 7. Glossary

TBD.

## 8. Intellectual Property Statement

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

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

## 9. Disclaimer

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

## 10. Full Copyright Notice

Copyright (C) Global Grid Forum (2006). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assignees.

## 11. References

*Note that only permanent documents should be cited as references. Other items, such as Web pages or working groups, should be cited inline (i.e., see the Global Grid Forum, <http://www.ggf.org>). References should conform to a standard such as used by IEEE/ACM, MLA, Chicago or similar. Include an author, year, title, publisher, place of publication. For online materials, also add a URL.*

[BRADNER1] Bradner, S. Key Words for Use in RFCs to Indicate Requirement Levels, RFC 2119. March 1997.