

## **SAGA API Extension: Service Discovery API**

### Status of This Document

This document provides information to the grid community, proposing a standard for an extension to the Simple API for Grid Applications (SAGA). As such it depends upon the SAGA Core API Specification [1]. This document is supposed to be used as input to the definition of language specific bindings for this API extension, and as reference for implementors of these language bindings. Distribution of this document is unlimited.

### Copyright Notice

Copyright © Open Grid Forum (2007). All Rights Reserved.

### Abstract

This document specifies a Service Discovery API extension to the Simple API for Grid Applications (SAGA), a high level, application-oriented API for grid application development. This Service Discovery API is motivated by a number of Use Cases collected by the OGF SAGA Research Group in GFD.70 [2], and by requirements derived from these Use Cases, as specified in GFD.71 [3]). It allows users to find services with minimal prior knowledge.

## **Contents**

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	Notational Conventions . . . . .	3
1.2	Security Considerations . . . . .	3
<b>2</b>	<b>SAGA Service Discovery API</b>	<b>4</b>

2.1	Introduction . . . . .	4
2.2	Specification . . . . .	5
2.3	Specification Details . . . . .	6
2.4	Examples . . . . .	11
<b>3</b>	<b>Intellectual Property Issues</b>	<b>13</b>
3.1	Contributors . . . . .	13
3.2	Intellectual Property Statement . . . . .	13
3.3	Disclaimer . . . . .	14
3.4	Full Copyright Notice . . . . .	14
	<b>References</b>	<b>15</b>

## 1 Introduction

Most of the SAGA use cases [2] exhibit a need for service discovery - though it is sometimes described as resource discovery. For example the DiVA entry says:

DiVA infrastructure must; a) Discover available components on distributed resources. The list of available components must be searchable by different attributes. This overlaps the needs of RealityGrid.

and:

On startup, the application must gather a list of available “components”. Typically this is done by consulting a local configuration file to find the locations of the binaries (or bytecode files) associated with each component as well as their names and interface definitions. For DiVA, we would like to support the discovery of remote modules as well by contacting information services on other machines or a broker that locates components on all machines in a given Virtual Organization. From the application programmers point of view, they want to be presented with a searchable database of components (regardless of location) that can be queried and sorted based on criteria such as “name,” “location,” interface definition, etc... Organization

as an Relational Database or LDAP directory or even a flat-file is unimportant. The API should be able to hide these details as a query for components that satisfy the search criteria is presented.

This API extension is tailored to provide exactly this functionality, at the same time keeping coherence with the SAGA Core API look & feel, and keeping other Grid related boundary conditions (in particular middleware abstraction and authentication/authorization) in mind.

## 1.1 Notational Conventions

In structure, notation and conventions, this documents follows those of the SAGA Core API specification [1], unless noted otherwise.

## 1.2 Security Considerations

As the SAGA API is to be implemented on different types of Grid (and non-Grid) middleware, it does not specify a single security model, but rather provides hooks to interface to various security models – see the documentation of the `saga::context` class in the SAGA Core API specification [1] for details.

A SAGA implementation is considered secure if and only if it fully supports (i.e. implements) the security models of the middleware layers it builds upon, and neither provides any (intentional or unintentional) means to by-pass these security models, nor weakens these security models' policies in any way.

## 2 SAGA Service Discovery API

### 2.1 Introduction

The SAGA Service Discovery API provides a mechanism to locate services.

**Fixme:** *Check this next para* The next para probably needs updating - SMF

This SAGA API extension inherits the `object`, `async` and `monitorable` interfaces from the SAGA Core API [1].

#### 2.1.1 Service Model

The API is based upon the GLUE model of a service. This makes use of the GLUE service model as summarised in figure 1.

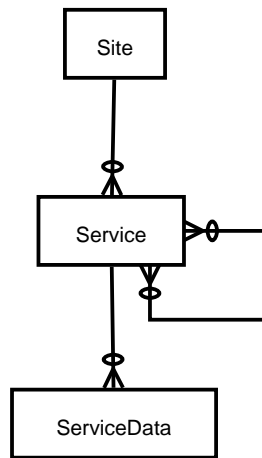


Figure 1: ER diagram of Service Model

The attributes are not shown as these are more subject to change. The figure indicates that a site may host many **Services** and a **Service** has multiple **ServiceData** entries associated with it. **ServiceData** entries have two attributes: Key and Value. This allows any set of keyword value pairs to be associated with an instance of a **Service**. In addition a **Service** has a many-to-many relationship with itself. This allows the model to describe groupings of services.

### 2.1.2 Classes

The SAGA Service Discovery API consists of a `discoverer` class with a method `list_services()`. This returns a set of objects of the `service` class. In addition the method `get_attribute_names()` returns the set of attribute names. The `service` class has a method `get_url()` – which is all that most people will use to obtain the address registered for the service. In the case of a Web Service this will be the service endpoint. It also has methods to extract properties of the service. These might be used by those who wish to generate a web page of services. These methods are `get_attribute_values()`, `get_services()`, `get_vo_names()`, `get_key_names()` and `get_key_values()`.

## 2.2 Specification

---

```
package saga.sd
{

    class discoverer : implements saga::object
    {
        CONSTRUCTOR (out discoverer dis);
        DESTRUCTOR (in discoverer dis);

        list_services(in string service_filter,
                     in string vo_filter,
                     in string key_filter,
                     out array<service> services);

        get_attribute_names (out array<string> attribute_names);
    }

    interface service : extends saga::object
    {
        get_url (out string url);

        get_attribute_values (in string attribute_name,
                             out array<string> attribute_values);

        get_services (out array<service> services);

        get_vo_names (out array<string> vo_names);

        get_key_names (out array<string> key_names);
    }
}
```

---

```

        get_key_values      (in    string      key_name,
                           out    array<string> key_values);
    }

```

---

## 2.3 Specification Details

### class discoverer

The **discoverer** object is the entry point for service discovery. Apart from the constructor and destructor it has two operations: **list\_services** which returns the list of services matching the specified filter strings and **get\_attribute\_names** which returns the full set of potentially available attribute names.

There are three filter strings: **service\_filter**, **vo\_filter** and **key\_filter** which act together to restrict the set of services returned.

Each of the filter strings uses SQL 92 syntax as if it were part of a **WHERE** clause acting to select from a single table that includes columns corresponding to each key name in the key/value pairs. If the programming language permits it, empty strings may be replaced by a representation of **NULL**.

Three strings are used rather than one as it clarifies the description of the functionality, avoids problems with key values being themselves existing GLUE attributes and facilitates implementation.

Only the following operators are permitted in the expressions: **IN**, **LIKE**, **AND**, **OR**, **NOT**, **=**, **>=**, **>**, **<=**, **<**, **<>** in addition to column names, parentheses, column values as single quoted strings, numeric values and the comma. An implementation should try to give an informative error message if the filter string does not conform however it is sufficient to report in which filter string the syntax error was found.

The **LIKE** operator matches patterns:

**'%xyz'** matches all entries with trailing xyz

**'xyz%'** matches all entries with leading xyz

**'%xyz%'** matches all entries with xyz being substring

Column names are not case sensitive but values are.

For matching on multivalued attributes it is sufficient that one attribute in the information system matches.

### Service Filter

Column names in the `service_filter` are dependent upon the GLUE service definition. Only those attributes considered useful to service discovery are supported. For GLUE 1.2 these are:

**type** type of service

**name** name of service - not necessarily be unique

**uid** unique identifier of service

**site** name of site

**url** the endpoint to contact the service - will normally be used with the LIKE operator

**service** for associated services

Some examples are:

- `type = 'org.glite.security.voms'`
- `site IN ('INFN-CNAF', 'RAL-LCG2')`
- `type = 'ResourceBroker' AND Site LIKE '%INFN%'`

### VO Filter

Column names in the `vo_filter` string are

**vo** Virtual Organization - will often be used with the IN operator

Some examples are:

- `VO IN ('cms', 'atlas')`
- `VO = 'dteam'`

## Data Filter

Column names in the the `data_filter` string are taken from the service data key/value pairs.

If values are specified as numeric values and not in single quotes the service data will be converted from string to numeric for comparison.

Some examples are:

- `source = 'RAL-LCG2' OR destination = 'RAL-LCG2'`
- `RunningJobs >=1 AND RunningJobs <= 5`

---

### - CONSTRUCTOR

Purpose: create a new discoverer object  
 Format: CONSTRUCTOR (out discoverer dis);  
 Inputs: -  
 Outputs: dis: new discoverer object  
 Throws: NotImplemented  
 NoSuccess  
 Notes:

### - DESTRUCTOR

Purpose: Destructor for discoverer object.  
 Format: DESTRUCTOR (in discoverer dis)  
 Inputs: dis: object to be destroyed  
 Outputs: -  
 Throws: -  
 Notes: -

### - list\_services

Purpose: return the set of services that pass the set of specified filters  
 Format: list\_services (in string service\_filter,  
 in string vo\_filter,  
 in string key\_filter,  
 out array<service> services);  
 Inputs: service\_filter: filter on the basic service and site attributes  
 and on associated services  
 vo\_filter: filter on VOs associated with the service  
 data\_filter: filter on key/value pairs associated with the  
 service  
 Outputs: -  
 Throws: NotImplemented

```

        BadParameter
        NoSuccess
    Notes:  - if any filter has an invalid syntax a
            'BadParameter' exception is thrown.

- get_attribute_names
  Purpose: return the set of attribute names
  Format:  get_attribute_names (out array<string> attribute_names);
  Inputs:  -
  Outputs: attribute_names:      set of available attribute names
  Throws:  NotImplemented
          BadParameter
          NoSuccess
  Notes:   -

```

---

## interface service

The service only has a set of getter methods for the user to obtain details of that service

**FiXme:** *multivalued attributes* We have assumed that all attribute values and key values are multi-valued this is not the case currently in GLUE but it allows for GLUE to evolve without messing up the specification - is this reasonable?

---

```

- get_url
  Purpose: return the URL to contact the service
  Format:  get_url                (out string url);
  Inputs:  -
  Outputs: url:                    URL to contact the service
  Throws:  NotImplemented
          NoSuccess
  Notes:   -

- get_attribute_values
  Purpose: return the set of attribute values
  Format:  get_attribute_names    (in string attribute_name,
                                out array<string> attribute_values);

```

---

Inputs: attribute\_name  
Outputs: attribute\_values: set of attribute values  
Throws: NotImplemented  
BadParameter  
NoSuccess  
Notes: - if the attribute\_name parameter is not in the set of  
expected attribute\_names a 'BadParameter'  
exception is thrown.

- get\_services

Purpose: return the set of associated services  
Format: get\_services (out array<service> services);  
Inputs: -  
Outputs: services: set of associated service objects  
Throws: NotImplemented  
NoSuccess  
Notes: -

- get\_vo\_names

Purpose: return the set of vo names  
Format: get\_vo\_names (out array<string> vo\_names);  
Inputs: -  
Outputs: vo\_names: set of associated vo\_names  
Throws: NotImplemented  
NoSuccess  
Notes: -

- get\_key\_names

Purpose: return the set of key names  
Format: get\_key\_names (out array<string> key\_names);  
Inputs: -  
Outputs: key\_names: set of associated key\_names  
Throws: NotImplemented  
NoSuccess  
Notes: -

- get\_key\_values

Purpose: return the set of key values  
Format: get\_key\_names (in string key\_name,  
out array<string> key\_values);  
Inputs: key\_name  
Outputs: key\_values: set of key values  
Throws: NotImplemented  
NoSuccess  
Notes: -

## 2.4 Examples

**Fixme:** *Service handle* This code passes in a session handle to the constructor - the main spec needs updating to match. Maybe we will have a pair of constructors as the session is not always needed it depends partly on the security model of the underlying info system.

This C++ example shows how SAGA service discovery model can be used to retrieve services from the underlying information system. All the “ResourceBroker” services with a name of “CERN-PROD-rb” and owned by either “Atlas” or “DTeam” are requested. The service objects returned from the `list_services` call are then queried for attributes and key/values using its getter methods. It would be more common to issue a sufficiently precise query so that any service returned would be suitable and then call `get_url` on the first service returned.

### Code Example

```

1  using namespace std;
2
3  // ... code for setting up the session omitted ...
4
5  saga::discoverer d (session_handle);
6
7  vector<std::string> attrib_names = d.get_attribute_names();
8
9  string svc_filter =
10     "Type = 'ResourceBroker' AND NAME = 'CERN-PROD-rb'";
11  string vo_filter =
12     "VO IN ('atlas', 'dteam')";
13
14  vector<saga::service> slist =
15     d.list_services (svc_filter, vo_filter, NULL);
16
17  cout << "Total number of services found = " << slist.size() << endl;
18
19  for (int i = 0; i < slist.size(); i++)
20  {
21      cout << "SERVICE #" << i << endl;
22      cout << "-----" << endl;
23      for (int j = 0; j < attrib_names.size(); j++)
24      {
25          attrib_values = slist[i].get_attribute_values(attrib_names[j]);
26          for (int k = 0; k < attrib_values.size(); k++)
27              cout << attrib_names[j] << " : " << attrib_values[k] << endl;
28      }
29  }

```

```
29     cout << "-----" << endl;  
30 }
```

**Fixme:** *Explain* Need to provide C code to go in the box below  
and explain it here

This C example ....

Code Example

```
1  Replace by code  
2  (indented by at least 2) ...
```

## 3 Intellectual Property Issues

### 3.1 Contributors

This document is the result of the joint efforts of several contributors. The authors listed here and on the title page are those committed to taking permanent stewardship for this document. They can be contacted in the future for inquiries about this document.

**Steve Fisher****s.m.fisher@rl.ac.uk**

Rutherford Appleton Lab

Chilton

Didcot

Oxon

OX11 0QX

UK

**A Paventhan****paventhan@rl.ac.uk**

Rutherford Appleton Lab

Chilton

Didcot

Oxon

OX11 0QX

UK

The initial version of the presented SAGA SD API was drafted by members of the SAGA Research Group. Members of this group did not necessarily contribute text to the document, but did contribute to its current state. Additional to the authors listed above, we acknowledge the contribution of the following people, in alphabetical order:

**FiXme:** *Credits* Need to credit the right people

### 3.2 Intellectual Property Statement

The OGF 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 OGF Secretariat.

The OGF 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 OGF Executive Director.

### 3.3 Disclaimer

This document and the information contained herein is provided on an "As Is" basis and the OGF 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.

### 3.4 Full Copyright Notice

Copyright (C) Open Grid Forum (2007). 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 OGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the OGF 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 OGF or its successors or assignees.

## References

- [1] T. Goodale, S. Jha, H. Kaiser, T. Kielmann, P. Kleijer, A. Merzky, J. Shalf, and C. Smith. A Simple API for Grid Applications (SAGA). Grid Forum Document GFD.xx, 2007. Global Grid Forum.
- [2] A. Merzky and S. Jha. A Collection of Use Cases for a Simple API for Grid Applications. Grid Forum Document GFD.70, 2006. Global Grid Forum.
- [3] A. Merzky and S. Jha. A Requirements Analysis for a Simple API for Grid Applications. Grid Forum Document GFD.71, 2006. Global Grid Forum.

## List of Corrections

FiXme: Check this next para . . . . .	4
FiXme: multivalued attributes . . . . .	9
FiXme: Service handle . . . . .	11
FiXme: Explain . . . . .	12
FiXme: Credits . . . . .	13