

## The SAGA C++ API Programmer's Guide

The Simple API for Grid Applications (SAGA) is a collection of C++ API packages defined on top of an overall programmatic appearance. The SAGA API packages try to cover the bulk of the requirements needed for grid programming like, for example, job submission, remote file transfer and replica management, to name a few. This guide aims for the developer who wants to learn the basics of how to use SAGA and C++ to develop applications that can make use of a broad set of existing distributed and grid computing infrastructures.<sup>1</sup>

### Status of This Document

This guide is still work in progress.

---

<sup>1</sup>editor

<sup>1</sup>In order to understand the intent and also the limitations of the SAGA approach, it is useful to read the Wikipedia entry on "Leaky Abstractions" ([http://en.wikipedia.org/wiki/Leaky\\_abstraction](http://en.wikipedia.org/wiki/Leaky_abstraction)), and also Joel Spolky's column on the same topic (<http://www.joelonsoftware.com/articles/LeakyAbstractions.html>).

## Contents

## 1 Introduction

What are Grids? How can you use them? Why would you need an API for doing so? What makes SAGA different from other Grid APIs?

These are some of the questions which we will try to answer in this introduction. Don't expect a complete treatise on the subject though; in stead, we will refer the interested reader to the literature where appropriate. Also, this introduction can safely be skipped by readers which are familiar with Grid programming, or even just with using Grids.

The text assumes that the reader has some familiarity with the concept of distributed computing and programming of distributed applications.

### Grid Computing

In "What is a Grid?" [?], a three point checklist is provided. Although by no means rigorous or complete, it provides a starting point. Following that, Grids are:

1. coordinated resources that are not subject to centralized control,
2. use standard, open, general-purpose protocols and interfaces,
3. deliver non-trivial qualities of service.

Compared to 'ordinary' distributed computing, scalability seems thus to be a prevalent property of Grids as these points above are arguable properties which mostly arise from scalability issues. **Jha: unclear...**

### Grid Standardization

A diverse Grid Computing community has emerged over the past few years; the Open Grid Forum (OGF) [?] serves as a focus point for the Grid community groups and Grid standardization efforts. The main topics of standardization are many fold, and cover, amongst others,

- service architectures
- service interfaces
- wire protocols
- information models
- Grid related markup languages

- user interfaces

For this document, as might be obvious, the last point, which includes Application Programming Interfaces (APIs), is of special interest. It should be noted that OGF's standardization mostly targets Grid middleware developers and Vendors, and that the OGF's API specifications are the top layer of OGF's standards stack.

## Grid APIs

Traditionally, most distributed (and Grid) middleware systems come with a 'native' API. Best known for such an API is Globus, a pioneering Grid Middleware project: each version of Globus services was/is accompanied by an extensive API which allows programmers (end-user application or otherwise), to make use of these services, and provides access to Globus service client libraries, Globus protocol implementations, Globus security tools, Globus markup languages, etc.

Despite the success of Globus, its approach exhibits a number of problems: applications written against a specific version of the Globus API are not (easily) portable to other versions of the API, and not portable at all to other Grid middleware, such as Genesis-II, GridSam, or Unicore. Furthermore, typical middleware APIs expose a specific set of features available in that middleware. These do not necessarily match the Grid abstractions required by the various applications; higher level services which provide these abstractions, are then required to implement their own specific APIs.

That situation is, to some extent, comparable to the 80's and 90's, where many vendors of parallel computers shipped their specific proprietary communication library along with their product, to enable the end users to make use of the distributed compute power from within a single (distributed) application. Before the establishment of MPI as a single dominant standard for this inter-node communication, the creation of *portable* distributed applications was an extremely difficult and tedious effort. Only the adoption of MPI by basically all cluster vendors, and the *native support* for MPI on most platforms, could finally alleviate that problem<sup>2</sup>.

## SAGA

The Simple API for Grid Applications (SAGA), a proposed OGF standard, is in several ways a comparable effort to MPI: SAGA tries to establish a *single*

---

<sup>2</sup>Note that MPI serves a specific problem space of tightly coupled massive parallel applications. Loosely coupled parallel applications, for example, are served by a different set of distributed communication concepts, e.g., Peer2Peer systems, or component models, etc.

API for Grid application programmers, which is ideally shipped with all Grid middleware, so that the programmer can focus on the application logic, instead of having to deal with tedious and complex middleware details.

But SAGA tries to go further: instead of defining a fixed single common denominator, SAGA is by definition an extensible and modular effort. By being extensible both horizontally and vertically, SAGA should be able to adapt to a variety of use cases and user communities. Additionally, our SAGA implementation can be extended by additional middleware bindings, to support the widest possible portability for applications.

### Horizontal Extensibility of SAGA

The SAGA approach is extremely modular: a stable and finite core set of SAGA Look & Feel packages is accompanied by a variable set of functional API packages. That latter set is expected to grow over time, and to cover future emerging Grid programming paradigms.

#### **FIXME: graphics goes here**

At the moment, the SAGA API covers the following functional packages:

|                 |                               |
|-----------------|-------------------------------|
| <b>jobs:</b>    | job creation and management   |
| <b>files:</b>   | interaction with file systems |
| <b>replica:</b> | management of logical files   |
| <b>streams:</b> | BSD socket oriented IPC       |
| <b>rpc:</b>     | remote procedure calls        |

All these packages are provided by our SAGA implementation. Additional packages, which are in the process of being defined, include:

|                  |  |
|------------------|--|
| <b>adverts:</b>  | persistent storage of application level information <sup>2</sup> |
| <b>sd:</b>       | service discovery <sup>2</sup>                                   |
| <b>messages:</b> | message based IPC  |
| <b>cpr:</b>      | checkpoint and recovery  |
| <b>dais:</b>     | database access and integration                                  |

### Vertical Extensibility of SAGA

The experiences in various OGF user communities show that the Grid programming models used in different application domains vary widely. In particular, it

---

<sup>2</sup>These packages are also provided in our implementation.

seems impossible to completely standardize information models, data models, and data formats, without losing the abstractive power of a high level application oriented API.

**FIXME: graphics goes here**

For that reason, SAGA tries to stay independent of these issues, and, at the same time, offers (a) the flexibility to natively accommodate a variety of data types and structures, and (b) the ability to additionally define higher level API packages, which address application domain specific services and programming models. If and how these additional high level packages are standardized is then up to the community creating these packages. SAGA, however, maintains a uniform and consistent API for a very wide variety of users and use cases, without being too specific and inflexible.

### **Implementation Extensibility**

Any Grid API is only as good and powerful as its underlying middleware is. How does this statement hold for SAGA, which is by design not bound to a specific Grid middleware? It holds true in fact: although the syntax and semantics of the SAGA API calls remain stable over the various Grid middlewares SAGA is running on top of, the specific performance characteristics, and in fact the exact set of SAGA calls provided, may vary from case to case.

That may seem like a serious drawback, and probably is, but an API can do only so much to emulate missing middleware features. In any case, to simplify the runtime portability of applications, our SAGA implementation allows binding to a variety of Grid middlewares *at runtime*, i.e., without the need to even relink your application.

**FIXME: graphics goes here**

That functionality is provided by a plugin type mechanism, which loads middleware *adaptors* into SAGA as required. These adaptors translate the SAGA API calls to the specific Grid middleware actions. For more details on adaptors and their maintenance and configurations, please refer to the installation manual.

## **1.1 Getting Started**

The reader should by now have a fair idea of the target scope of the SAGA API: it is designed to simplify the programming of novel applications which are to be run on a variety of Grid middlewares.

The next chapter, the 'Quick Start Guide', should be enough to get you go-

ing, and to compile and run small SAGA applications. The chapters after that will discuss the SAGA Look & Feel in more detail, and also dive into the various functional packages we provide. The chapters are independent, no specific reading order is implied, unless noted otherwise.

## 2 Quick Start Guide

So, you are one of the impatient readers who dread long dry user manuals? No problem: read this chapter, and you will have a good starting point to use SAGA. After all, the **S** in **SAGA** stands for **Simple**! The remainder of the document provides considerable more details about the API, but you can defer the later chapters until you have indeed the need for more information.

It is important to understand that the API consists of two parts: the Look & Feel, and the API packages. The packages are what you are most likely interested in, because they provide the means to interact with the Grid: they start jobs, copy files, perform remote procedure calls, etc. The Look & Feel provides the syntactic and semantic expressiveness to control *how* these actions are expressed (syntax) and performed (semantics). For example, the session management in the Look & Feel part tells you how to specify the security constraints of your actions, while the task model from the Look & Feel determines how you can express synchronous versus asynchronous actions.

That all sounds rather theoretical: let us dive into some examples! The first one allows you to copy a file<sup>3</sup>:

### File copy

```
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      // do a file copy
6      saga::url u(argv[1]);
7      saga::filesystem::file f (u);
8      f.copy (saga::url (argv[2]));
9  }
```

Isn't that simple? Three lines, and your file is copied! `file.copy()` is provided by the `file` class in the `saga::filesystem` package. That is a *functional* SAGA call. So, how will *non-functional* properties of that call enter the picture? Let's try to have the same call asynchronously:

### File copy (async version)

```
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      // run a file copy asynchronously
```

<sup>3</sup>For the sake of brevity, we leave error and sanity checks out of the examples

```
6   saga::url u(argv[1]);
7   saga::filesystem::file f (u);
8   saga::task t = f.copy <saga::task::Async> (saga::url (argv[2]));
9
10  // do something else
11
12  // wait for the copy task to finish
13  t.wait ();
14 }
```

That is almost the same. In particular, the signature for the `file.copy()` method is the same. It is, however, now qualified as `saga::task::Async`, and returns a `saga::task` instance. That (stateful) task represents the asynchronous call. You can think of the synchronous qualification as the default, which can be left out.

Another Look & Feel package deals with monitoring, and allows, for example, to get notifications if the task finishes. Yet another Look & Feel package allows to specify security credentials for the copy operation, and so on and so forth.

So, that is the general picture: SAGA packages provide you with means to interact with the Grid environment, and the Look & Feel packages determine how these actions are executed, syntactically and semantically, altering the functional method's performance.

The remainder of the document will give more details, first about the Look & Feel packages of SAGA, and then about the functional packages. Each section starts with a quick introduction (which should be enough to get you going in most cases), followed by a reference section, and by a more detailed discussion where appropriate. Finally, a number of advanced topics are discussed, which are probably not of your interest initially, but they are certainly useful concepts as soon as your application reaches a certain complexity itself.

Before diving into the API itself, we will describe the general context of the SAGA API in a couple of pages, including a short description of the Grid standardization landscape, and some details about our specific implementation of the SAGA API Specification. **Hasn't this already been done in the previous chapter?**

## 3 Building and Running your Application

### 3.1 Building your Application

For Unix-like systems, a configure/make based build system is provided. That system can also be used to compile your SAGA application, and is (briefly) described below.

Furthermore, your SAGA installation offers a tool named `saga-config`, which can also be used to determine the relevant compile parameters.

SAGA uses a gnu-make based build system. It includes a number of makefiles throughout the source tree, and the `$$SAGA_ROOT/make/` directory. All these *make includes*<sup>4</sup> get installed into `$$SAGA_LOCATION/share/saga/make/`<sup>5</sup>.

When compiling your application, these make includes may provide an easy starting point. All you need to do is to set `SAGA_LOCATION` to have it point to your SAGA installation root, and to include `saga.application.mk`, which will add all rules required to build a SAGA application. The following Makefile stub should get you started:

```

----- Makefile for SAGA applications -----
1  SAGA_SRC  = $(wildcard *.cpp)
2  SAGA_BIN  = $(SRC:%.cpp=%)
3
4  include $(SAGA_LOCATION)/share/saga/make/saga.application.mk

```

This stub loads the make rules, etc., needed to build the application. If the application needs additional include directories or libraries, use the following syntax *after* the make includes:

```

----- Makefile: setting compiler/linker flags -----
5  SAGA_CPPFLAGS += -I/opt/super/include
6  SAGA_LDFLAGS  += -L/opt/super/lib -lsuper

```

Of course it is possible to build SAGA applications with custom Makefiles. The SAGA make includes can still, however, be used to obtain the SAGA specific compiler:

<sup>4</sup>A make include is a makefile building block

<sup>5</sup>For details on building and installing SAGA, please refer to the SAGA installation manual.

```

----- Custom Makefile -----
1  SRC      = $(wildcard *.cpp)
2  OBJ      = $(SRC:%.cpp=%.o)
3  BIN      = $(SRC:%.cpp=%)
4
5  CXX      = g++
6  CXXFLAGS = -c -O3 -pthread -I/opt/mpi/include
7
8  LD       = $(CXX)
9  LDFLAGS  = -L/usr/lib/ -lc
10
11 include $(SAGA_LOCATION)/share/saga/make/saga.engine.mk
12 include $(SAGA_LOCATION)/share/saga/make/saga.package.file.mk
13
14 all: $(BIN)
15
16 $(OBJ): %.o : %.cpp
17         $(CXX) $(CXXFLAGS) $(SAGA_CXXFLAGS) -o $@ $<
18
19 $(BIN): % : %.o
20         $(LD) $(LDFLAGS) $(SAGA_LDFLAGS) -o $@ $<

```

SAGA\_CXXFLAGS and SAGA\_LDFLAGS contain only those options and settings which are required to use SAGA. You may want to use 'make -n' to print what the resulting make commands are, in order to debug eventual incompatibilities between the SAGA compiler and linker flags, and your own ones.

Yet another option is to use the output of saga-config for makefiles:

```

----- saga-config used in Makefile -----
1  SRC      = $(wildcard *.cpp)
2  OBJ      = $(SRC:%.cpp=%.o)
3  BIN      = $(SRC:%.cpp=%)
4
5  CXX      = gcc
6  CPPFLAGS = -I/opt/mpi/include
7  CXXFLAGS = -c -O3 -pthread
8
9  LD       = $(CXX)
10 LDFLAGS  = -L/opt/mpi/lib/ -lmpi
11
12 CPPFLAGS += $(shell $(SAGA_LOCATION)/bin/saga-config --cppflags)
13 CXXFLAGS += $(shell $(SAGA_LOCATION)/bin/saga-config --cxxflags)
14 LDFLAGSS += $(shell $(SAGA_LOCATION)/bin/saga-config --ldflags)
15
16 .default: $(BIN)
17

```

```
18 $(BIN): % : %.cpp
19
20 $(OBJ): %.o : %.cpp
21     $(CXX) $(CPPFLAGS) $(CXXFLAGS) -o $@ $<
22
23 $(BIN): % : %.o
24     $(LD) $(LDFLAGS) -o $@ $<
```

## Linkage Options

The SAGA libraries come in two flavours: *standard* and *lite*. The standard version is what you have observed in the examples above: your application is linked against the `libsaga_engine`, and against the available/required packages (`libsaga_package_abc`). The adaptor libraries (`libsaga_adaptor_xyz`) are not linked to the application, but are loaded at runtime.

That standard version offers the most flexibility for your application, and allows it to adapt linkage to a wide range of use cases. The runtime adaptor loading allows you to adapt your application to Grid middleware variations at runtime. That comes at a cost: at runtime, the shared library dependencies have to be resolved, and SAGA has to be correctly configured to find the adaptor libraries at runtime. The former can be resolved by linking your application statically, but the SAGA runtime configuration cannot be avoided.

That is the reason why we provide the second, 'lite' version of the SAGA library: it is a *single* shared library which contains the SAGA engine, all available packages, and a set of adaptors. These adaptors are thus *not* loaded at runtime, but are linked at link-time. Other adaptors can, however, still be loaded at runtime, as before.

That way, no SAGA runtime configuration is required at all, as all dependencies are resolved on link-time. The SAGA installation manual provides more information on the selection of adaptors to be included into the (`libsaga_lite`). Note that this comes at a cost as well: your application binary will be larger, as all the adaptors are linked against it, whether they are needed or not.

The SAGA make system supports linking against the (`libsaga_lite`) like this:

```
----- Makefile for SAGA-Lite -----
1  SAGA_SRC = $(wildcard *.cpp)
2  SAGA_BIN = $(SRC:%.cpp=%)
3
4  SAGA_USE_LITE = yes
5
```

```
6 | include $(SAGA_LOCATION)/share/saga/make/saga.application.mk
```

## 3.2 Running your Application

Your SAGA application should have only few runtime dependencies. First of all, it needs to find the shared libraries you used, and the shared libraries your SAGA installation is linked against. That may include libraries required by the SAGA adaptors, e.g., globus libraries, openssl, etc. Depending on your operating system, you may need to set the environment variables `LD_LIBRARY_PATH`, `DYLD_LIBRARYPATH`, or similar.

After starting, SAGA reads a couple of initialization files to determine your specific set of adaptors, and several configuration options for these adaptors. In general, you should not have the need to touch these ini files – SAGA configuration and installation is supposed to be performed by the system administrator, not the end user. In most cases, SAGA should be able to find these ini files automatically (the installation prefix is statically compiled into SAGA. If that fails, SAGA searches the following locations (in this order) for the main `saga.ini` file:

```
/etc/saga.ini
$SAGA_LOCATION/share/saga/saga.ini
$HOME/.saga.ini
$PWD/.saga.ini
$SAGA_INI
```

Those ini files point to the individual adaptor ini files, which then allow SAGA to load these adaptors.

While running, the SAGA library can print log messages, of varying verbosity. That output is controlled by the environment variable `SAGA_VERBOSE`. The values are as follows:

```
1: Critical
2: Error
3: Warning
4: Info
5: Debug
6: Blurb
```

When `SAGA_VERBOSE` is not set, the library is silent.

## Part I

# General SAGA Concepts

A number of concepts and paradigms are used in all packages of the SAGA API. The SAGA specification names these concepts as the '*SAGA Look & Feel*'. This section describes the most important concepts.

As this section does not actually provide the reader with the means to perform any useful remote operation in Grids, the reader may want to skip this part and continue to read Part ??, and come back to this part as needed.

## 4 Error Handling

### 4.1 Quick Introduction

SAGA error handling is exception-based: a rather flat hierarchy of 14 exception types inherit the properties of the general `saga::exception` class. That allows to inspect the exception for its exact type, and for the detailed error message(s) associated with this exception.

**HINT:**

As our SAGA implementation is adaptor-based, a single API call could actually return more than one exception at the same time. That is rendered by returning the most specific exception, whose error message then also contains the error messages from all the other exceptions, usually one per adaptor.

### 4.2 Reference

Prototype: `saga::exception` class and derivatives

```
1
2 namespace saga
3 {
4   class exception : public std::exception
5   {
6   public:
7     ~exception() throw() {}
8
9     // Gets the message associated with the exception
10    char const* get_message() const throw()
11
12    // an alias for get_message
13    char const* what() const throw()
14
15    // get type of exception
16    saga::error get_error () const
17
18    // Gets the SAGA object associated with exception.
19    saga::object get_object () const throw()
20 };
21
22
23
24 // parameter related exceptions
25 class parameter_exception : public saga::exception
```

```
26 class incorrect_url      : public saga::parameter_exception
27 class bad_parameter      : public saga::parameter_exception
28
29 // state related exceptions
30 class state_exception    : public saga::exception
31 class already_exists     : public saga::state_exception
32 class does_not_exist    : public saga::state_exception
33 class incorrect_state    : public saga::state_exception
34 class timeout            : public saga::state_exception
35
36 // security related exceptions
37 class security_exception : public saga::exception
38 class permission_denied  : public saga::security_exception
39 class authorization_failed : public saga::security_exception
40 class authentication_failed : public saga::security_exception
41
42 // general exceptions
43 class no_success         : public saga::exception
44 class not_implemented   : public saga::exception
45
```

### 4.3 Details

Error handling in SAGA is completely exception-based<sup>6</sup>. The following exceptions exist (note that the exception class names and the `saga::error` enums, returned by `exception.get_error()` have identical names):

#### NotImplemented

A method is specified in the SAGA API, but is not provided by this specific SAGA implementation<sup>7</sup>.

#### IncorrectURL

This exception indicates that a URL argument could not be handled. For example, this implementation may be unable to handle the specified protocol, or the access to the specified entity, via the given protocol, is impossible.

#### BadParameter

This exception indicates that at least one of the parameters of the method call is ill-formed, invalid, out of bounds or, otherwise, not usable. The

<sup>6</sup>The SAGA specification allows for an additional `error_handler` interface to be implemented by SAGA objects—that seemed unnecessary for object-oriented languages such as C++.

<sup>7</sup>In particular, the method is not provided by any of the available adaptors, see Section ??.

error message gives specific information on what parameter caused the exception, and why.

**AlreadyExists**

This exception indicates that an operation cannot succeed because an entity to be created already exists, and cannot be overwritten. Explicit flags on the method invocation may allow the operation to succeed, e.g., if they indicate that `Overwrite` is allowed.

**DoesNotExist**

This exception indicates that an operation cannot succeed because a required entity is missing. Explicit flags on the method invocation may allow the operation to succeed, e.g., if they indicate that `Create` is allowed.

**IncorrectState**

This exception indicates that the object a method was called upon is in a state where that method cannot possibly succeed. A change of state might allow the method to succeed with the same set of parameters.

**PermissionDenied**

An operation failed because the identity used for the operation did not have sufficient permissions to successfully perform the operation.

**AuthorizationFailed**

An operation failed because none of the available contexts of the used session could be used to access the given resource. In contrast to the `PermissionDenied`, this exception usually indicates an error on the administrative level, which usually cannot be fixed by the end user.

**AuthenticationFailed**

An operation failed because none of the available session contexts could successfully be used for authentication. That exception usually indicates invalid or outdated security credentials.

**Timeout**

This exception indicates that a remote operation was not completed successfully, because the network communication or the remote service timed out. This exception is never thrown if a timed `wait()` or similar method times out, as that is not an error condition.

**NoSuccess**

This exception indicates that an operation failed for any other reason. The exception message may, or may not, contain some more details about the cause of the error.

Even for simple SAGA operations, the implied Grid interactions can be fairly complex and may invoke multiple remote operations. It may thus happen that multiple exceptions apply during the execution of the method. In such cases, the **most specific exception** is thrown (the list of exceptions above is ordered, with the most specific exception up front).

Also, as the above Reference section shows, the exceptions are ordered in an exception hierarchy. The application can thus try to catch whole classes of exceptions, e.g., all security related exceptions, or all state related exceptions. The code example below shows, however, how the lower level exceptions can be accessed for debugging purposes:

## Code Example

```
1  try
2  {
3      saga::url u ("/path/which/does/not/exist");
4      saga::filesystem::file f (u);
5  }
6  catch ( saga::state_exception const & e )
7  {
8      switch ( e.get_error () )
9      {
10         // handle does not exist
11         case saga::DoesNotExist:
12         {
13             std::cout << "file does not exist"
14                 << std::endl;
15             exit (-1);
16             break;
17         }
18
19         // generic handler for state related problems
20         default:
21         {
22             std::cout << "some other saga state exception caught: "
23                 << std::string (e.what ())
24                 << std::endl;
25             exit (0);
26             break;
27         }
28     }
29 }
```

## 5 Using Data Buffers

### 5.1 Quick Introduction

Various classes (e.g., `saga::file` and `saga::stream`) in the SAGA API expose I/O operations, i.e., chunks of binary data can be written to, or read from these classes. Other classes (such as `saga::rpc`) handle binary data as parameters. In order to unify the application management of these data, SAGA introduces the `saga::buffer` class, which is essentially a simple container class for a byte buffer, plus a number of management methods. Various subclasses of the `saga::buffer` exist, and, as described below, users are allowed, and actually encouraged, to build their own ones.

The C++ rendering of SAGA distinguishes between mutable and non-mutable buffers: non-mutable buffers are used for write-type operations, and cannot be changed by the SAGA implementation; mutable buffers are for read-type operations, and can be changed (i.e., new data can be added to the buffer).

### 5.2 Reference

Prototype: `saga::buffer`

```
1 namespace saga
2 {
3     class const_buffer
4         : public saga::object
5     {
6     public:
7         const_buffer (void const * data,
8                     saga::ssize_t size);
9         ~const_buffer (void);
10
11         saga::ssize_t get_size (void) const;
12         void const * get_data (void) const;
13         void close (double timeout = 0.0);
14     };
15
16     class mutable_buffer
17         : public saga::const_buffer
18     {
19     public:
20         typedef void buffer_deleter_type(void* data);
21
22         typedef TR1::function <buffer_deleter_type> buffer_deleter;
23         static void default_buffer_deleter (void * data);
24     }
```

```
25     mutable_buffer (saga::ssize_t size = -1);
26     mutable_buffer (void * data,
27                   saga::ssize_t size);
28     mutable_buffer (void * data,
29                   saga::ssize_t size,
30                   buffer_deleter cb);
31     ~mutable_buffer (void);
32
33     void set_size (saga::ssize_t size = -1);
34     void set_data (void * data,
35                   saga::ssize_t size,
36                   buffer_deleter cb = default_buffer_deleter);
37     void * get_data (void); // non-const version
38 };
39 }
```

## 5.3 Details

Although the concept of an I/O buffer is very simple, and the prototype shown above is rather straight forward, the semantic details of the SAGA buffer are relatively rich. That holds true in particular for the memory management of the buffer data segment. For the interested reader, the `saga::buffer` section in the SAGA Core API specification contains quite some detail on that issue.

### 5.3.1 Buffer Memory Management

In general, buffers can operate in two different memory management modes: the data segment can be user-managed (i.e., application-managed), or SAGA-managed (i.e., implementation-managed). The constructors allow the application to pass a memory area on buffer creation: if that buffer is given, and not-NULL, then the SAGA implementation will use that buffer, and will never re- or de-allocate that memory (memory management is left up to the application). On the other hand, if that memory area is not given, or given as NULL, then the SAGA implementation will internally allocate the required amount of memory, which **MUST NOT** be re- or de-allocated by the application (memory management is left to the SAGA implementation).

Although the latter version is certainly convenient for the end user, it comes with a potential performance penalty: data from the implementation allocated buffer may sometimes need an additional memcopy into application memory. If that is the case, it is up to you to decide what memory management mode works best for your application use case.

**HINT:**

The most performant case is most of the time to re-use a single (or a small set of) application allocated memory buffer(s) over and over again. Note that you can use a larger size memory segment for a small buffer by giving a smaller size parameter to the constructor.

Example `saga::buffer` usage

```
1 // allocate a 'large' buffer statically
2 char mem[1024];
3
4 // create a saga buffer object for reading (mutable)
5 saga::mutable_buffer buf (mem, 512);
6
7 // open a file
8 saga::url u ("/etc/passwd");
9 saga::filesystem::file f (u);
10
11 // read data into buffer - the first 512 bytes get fill
12 f.read (buf);
13
14 // seek the buffer, so that the next read goes into the
15 // second half of the buffer
16 buf.set_data (mem + 512, 512);
17
18 // now read again
19 f.read (buf);
20
21 // the complete buffer should be filled now
22 // print what we got
23 std::cout << mem << std::endl;
```

### 5.3.2 Const versus Mutable Buffers

On write-like operations, the SAGA implementation has no need to change the buffer's data segment in any way: it only needs to read the data, and to copy them to whatever entity the write operations happens upon. The implementation can thus treat the buffer as `const`, which allows a number of optimizations and memory access safeguards to be employed.

On the other hand, read-like operations will usually require the SAGA implementation to write, or even to (re-)allocate the buffers memory segment. In such cases, `const` safeguards cannot be employed.

**HINT:**

It is encouraged the use of `const_buffer` instances for `write`-like operations, and of `mutable_buffer` instances for `read`-like operations.

In order to simplify memory management and to provide optimal memory access safeguards, the SAGA C++ bindings distinguish between `const_buffer` and `mutable_buffer` classes. Both types can be used for `write`-like operations, but only `mutable_buffer` instances can be used for `read`-like operations.

**HINT:**

It is possible to case `mutable_buffer` instances to `const_buffer`, which allows to re-use buffers for all I/O operations and, at the same time, allows the implementation to use `const` checking.

**FIXME: Is that true?**

## 6 Using Attributes

### 6.1 Quick Introduction

Attributes in SAGA are handled via the `saga::attribute` interface shown above. That interface allows to set, query, and to inspect specific attributes on those SAGA objects that implement the interface.

### 6.2 Reference

```
----- Prototype: saga::attribute -----
1 namespace saga
2 {
3     namespace attributes
4     {
5         // common attribute values
6         char const * const common_true = "True";
7         char const * const common_false = "False";
8     }
9
10    class attribute
11    {
12    public:
13        typedef std::vector <std::string>          strvec_type;
14        typedef std::map <std::string, std::string> strmap_type;
15
16        std::string get_attribute (std::string key) const;
17        void set_attribute (std::string key,
18                           std::string val);
19
20        strvec_type get_vector_attribute (std::string key) const;
21        void set_vector_attribute (std::string key,
22                                   strvec_type val);
23
24        void remove_attribute (std::string key);
25
26        strvec_type list_attributes (void) const;
27        strvec_type find_attributes (std::string pat) const;
28
29        bool attribute_exists (std::string key) const;
30        bool attribute_is_readonly (std::string key) const;
31        bool attribute_is_writable (std::string key) const;
32        bool attribute_is_vector (std::string key) const;
33        bool attribute_is_removable (std::string key) const;
34    };
```

```
35     }
```

### 6.3 Details

A prominent example is the `saga::job::description` class:

```
Code Example
1  saga::job::description jobdef;
2
3  std::vector <std::string> args;
4  args.push_back ("2");
5
6  jobdef.set_attribute      ("Executable", "/bin/sleep");
7  jobdef.set_vector_attribute ("Arguments",  args);
8
9
10 saga::job job = js.create_job (jobdef);
```

That example, as simple as it is, already shows most of what the attribute interface offers—that interface is really simple! For some classes, the set of available attributes is fixed, and will never change. That is also the case for the job description from the example above. For other classes, such as for logical files (`saga::logical_file`), the application programmer can specify arbitrary attributes – their interpretations are, of course, up to the application again.

Note that SAGA defines the known available attributes as static strings. Thus, the above example is more safely written as:

```
Code Example
1  saga::job::description jobdef;
2
3  std::vector <std::string> args;
4  args.push_back ("2");
5
6  namespace sja = saga::job::attributes;
7
8  jobdef.set_attribute      (sja::description_executable, "/bin/sleep");
9  jobdef.set_vector_attribute(sja::description_arguments, args);
10
11 saga::job job = js.create_job (jobdef);
```

| Type   | Format                                   | Example                  |
|--------|--|--------------------------|
| String | as in <code>printf ("%s", val);</code>   | Hello World              |
| Int    | as in <code>printf ("%lld", val);</code> | 123                      |
| Float  | as in <code>printf ("%lld", val);</code> | 1.234E-4                 |
| Time   | as in <code>printf ("%lld", val);</code> | Mon Oct 20 11:31:54 1952 |
| Bool   | True or False                            | True                     |
| Enum   | literal value of the enum                | Done                     |

Table 1: Available types and the prescribed formatting for the attribute value

### 6.3.1 Attribute Types

Although all attributes in SAGA are string-based, we distinguish between different types of attributes. The available types and the prescribed formatting for the attribute values are summarized in Table ??.

Trying to set an attribute which is, for its type, incorrectly formatted, will result in a 'BadParameter' exception.

## 7 Using URLs

### 7.1 Quick Introduction

URLs (and URIs, see below) are a dominant concept for referencing application-external resources. As such, they are also widely used in the Grid world, and in SAGA. The `saga::url` class helps to manage such URLs.

### 7.2 Reference

```
----- Prototype: saga::url -----
1 namespace saga
2 {
3     class url
4         : public saga::object
5     {
6     public:
7         url (void);
8         url (std::string const & urlstr);
9
10        ~url (void);
11
12        url & operator= (std::string const & urlstr);
13
14        std::string get_string (void) const;
15
16        std::string get_url (void) const;
17        void set_url (std::string const & url);
18
19        std::string get_scheme (void) const;
20        void set_scheme (std::string const & scheme);
21
22        std::string get_host (void) const;
23        void set_host (std::string const & host);
24
25        int get_port (void) const;
26        void set_port (int port);
27
28        std::string get_fragment (void) const;
29        void set_fragment (std::string const & fragment);
30
31        std::string get_path (void) const;
32        void set_path (std::string const & path);
33
34        std::string get_userinfo (void) const;
35        void set_userinfo (std::string const & userinfo);
```

```
36     saga::url translate    (std::string const & proto);
37 };
38
39
40
41     std::ostream& operator<< (std::ostream      & os,
42                             saga::url      const & u);
43     std::istream& operator>> (std::istream      & is,
44                             saga::url      & u);
45     bool          operator== (saga::url      const & lhs,
46                             saga::url      const & rhs);
47     bool          operator<  (saga::url      const & lhs,
48                             saga::url      const & rhs);
49 }
```

## 7.3 Details

In the time of the Internet, there are probably only few people in the first and second world who are not aware of the concept of URLs. However, many miss the finer details of the *Uniform Resource Locators*, and for a good reason: URLs are *designed* to hide a number of complexities from the user. In order to make efficient use of the SAGA API, and of many Grid concepts in general, we will need a basic understanding of several key elements of URLs<sup>8</sup>.

### 7.3.1 The Structure of URLs

RFC 3986[?] defines the generic syntax for URIs (and thus for URLs). A URL consists of four parts:

```
<scheme> : <hierarchical part> [ ? <query> ] [ # <fragment> ]
```

The `scheme` defines how the other parts of the URL are interpreted. The remaining parts define what resource the URL points to, and possibly how to access that resource.

The `hierarchical part` usually contains the following information: `userinfo`, `host`, `port`, and `path`. This information specifies **where** the resource is to be found.

The `saga::url` class defines setters and getters for most of these individual

---

<sup>8</sup>Technically, a URL is a URI that, “*in addition to identifying a resource, [provides] a means of locating the resource by describing its primary access mechanism (e.g., its network location).*” (“*Cool URLs don’t change*”, <http://www.w3.org/Provider/Style/URI>) This document, however, uses the terms URL and URI interchangeably.

URL elements. When setting one of these elements, SAGA will make sure that the result is a valid URL—otherwise the setter will decline the operation with a `BadParameter` exception:

Code Example

```
1  saga::url u ("ftp://remote.host.net:1234/data/old.dat");
2
3  u.set_host ("local.host/net");
4  u.set_path ("/data/new.dat");
```

Line 3 would result in a `BadParameter` exception, as the host contains an invalid character. Line 4 in the example above would change the URL to

```
ftp://remote.host.net:1234/data/new.dat
```

### 7.3.2 URLs in Grids

Now, in Grids we face the problem that different URLs may point to the same resource. For example, the following two URLs may refer to the same file:

```
http://data.silo.net/data/joe_doe/recent/abc.dat
ftp://data.silo.net/pub/joe_doe/recent/abc.dat
```

How can the SAGA API handle these differences transparently? In short: it can't<sup>9</sup>, however, it can help. For example, it allows to translate URLs from one scheme to the other. The code snippet below *may* be able to repeat the printout from above<sup>10</sup>:

Code Example

```
1  // Translating URLs
2  saga::url in("http://data.silo.net/data/joe_doe/recent/abc.dat");
3  saga::url out = in.translate ("ftp://");
4
5  std::cout << in.get_string () << std::endl;
6  std::cout << out.get_string () << std::endl;
```

Also, SAGA allows you to use the special scheme `any://` as a placeholder. If specified, the SAGA implementation tries to guess the correct scheme, and does the potentially required URL translation in the background.

<sup>9</sup>For a detailed discussion see section 2.11 of the SAGA specification [?].

<sup>10</sup>Please note that, at the moment, no adaptor implements `url::translate()`.

**HINT:**

Please note that URL translation is a non-trivial and error prone process, so it is not a good idea to rely on it if you want to keep your application portable (which is a MUST in Grids!). Wherever possible stick to the known schemes, and keep your transactions inside a single scheme and name space.

## Part II

# Using the SAGA-API Packages

As described in the introduction, the SAGA functional packages provide the programmer with the means to interact with Grid resources. Here, we describe these packages one by one, starting with a short overview for each, then a reference section (as C++ declarations of the respective classes and methods), and some details and examples for each of the classes.

## 8 Using the File Package

### 8.1 Quick Introduction

The SAGA filesystem package provides an abstraction to remote file systems. It inherits the SAGA namespace package, so the reader may want to study Section ?? as well.

File systems provide more than just a namespace: they also provide access to the *contents* of files: the `saga::filesystem` package extends the namespace package by adding `read()`, `write()` and `seek()` to the entries (i.e., to `saga::filesystem::file`):

```

                                     Accessing a file
1  // allocate a 'large' buffer statically
2  char mem[1024];
3  saga::mutable_buffer buf (mem, 512);
4
5  // open a file
6  saga::url u ("/etc/passwd");
7  saga::filesystem::file f (u);
8
9  // read data into buffer - the first 512 bytes get fill
10 f.read (buf);
11
12 // seek file and buffer
13 buf.set_data (mem + 512, 512);
14 f.seek (123, saga::filesystem::Current);
15
16 // read again
17 f.read (buf);
18
19 // print what we got
20 std::cout << mem << std::endl;
21
22 return (0);
```

The `saga::buffer` class represents the memory the data are written into. The buffers' memory management can be controlled, but the above example leaves everything to the SAGA implementation—for details, see Section ??.

Note that there are more I/O methods available in the filesystem package. They are, however, mostly provided for optimization, but are really *required* to make remote file I/O operations useful.

Conceptually, the file remains a namespace entry with added read, write and seek capabilities. For more information, see the details below, and also the section "*Namespaces*" (Sec. ??).

## 8.2 Reference

```

----- saga::filesystem enums -----
1  namespace saga
2  {
3      namespace filesystem
4      {
5          enum flags
6          {
7              Unknown      = /* -1, */  saga::name_space::Unknown    ,
8              None         = /*  0, */  saga::name_space::None      ,
9              Overwrite    = /*  1, */  saga::name_space::Overwrite  ,
10             Recursive    = /*  2, */  saga::name_space::Recursive  ,
11             Dereference  = /*  4, */  saga::name_space::Dereference ,
12             Create       = /*  8, */  saga::name_space::Create     ,
13             Exclusive    = /* 16, */  saga::name_space::Exclusive  ,
14             Lock         = /* 32, */  saga::name_space::Lock       ,
15             CreateParents = /* 64, */  saga::name_space::CreateParents ,
16             Truncate     =   128,
17             Append       =   256,
18             Read         = /* 512, */  saga::name_space::Read      ,
19             Write        = /*1024, */  saga::name_space::Write     ,
20             ReadWrite    = /*1536, */  saga::name_space::ReadWrite ,
21             Binary       =   2048
22         };
23
24         enum seek_mode
25         {
26             Start  = 1,
27             Current = 2,
28             End    = 3
29         };
30     }
31 }

```

```

----- Prototype: saga::filesystem::iovec -----
1  namespace saga
2  {
3      namespace filesystem
4      {
5          class const_iovec
6              : public saga::const_buffer
7          {

```

```

8     public:
9         const_iovec (void const * data,
10                    saga::ssize_t size,
11                    saga::ssize_t len_in = -1);
12     ~const_iovec (void);
13
14     saga::ssize_t get_len_in (void) const;
15     saga::ssize_t get_len_out (void) const;
16
17 };
18
19 class iovec
20     : public saga::mutable_buffer
21 {
22     public:
23     iovec (void *      data   = 0,
24           saga::ssize_t size = -1,
25           saga::ssize_t len_in = -1,
26           buffer_deleter cb   = default_buffer_deleter);
27     ~iovec (void);
28
29     void      set_len_in (saga::ssize_t len_in);
30     saga::ssize_t get_len_in (void) const;
31     saga::ssize_t get_len_out (void) const;
32 };
33 }
34 }

```

Prototype: `saga::filesystem::file`

```

1 namespace saga
2 {
3     namespace filesystem
4     {
5         class file
6             : public saga::name_space::entry
7         {
8             public:
9                 file (saga::session const & s,
10                    saga::url          url,
11                    int                 mode = Read);
12                 file (saga::url url,
13                    int                 mode = Read);
14                 file (saga::object const & o);
15                 file (void);
16                 ~file (void);
17
18                 file & operator= (saga::object const & o);
19

```

```

20         saga::off_t  get_size (void);
21
22         saga::ssize_t read      (saga::mutable_buffer buffer,
23                                 saga::ssize_t          length = 0);
24         saga::ssize_t write    (saga::const_buffer  buffer,
25                                 saga::ssize_t          length = 0);
26         saga::off_t  seek      (saga::off_t          offset,
27                                 seek_mode              mode);
28
29         void         read_v    (std::vector <iovec>   buffer_vec);
30         void         write_v   (std::vector <const_iovec>
31                                 buffer_vec);
32
33         saga::ssize_t size_p    (std::string         pattern);
34         saga::ssize_t read_p    (std::string         pattern,
35                                 saga::mutable_buffer buffer);
36         saga::ssize_t write_p   (std::string         pattern,
37                                 saga::const_buffer  buffer);
38
39         std::vector <std::string>
40         modes_e (void);
41         saga::size_t size_e     (std::string         ext_mode,
42                                 std::string         specification);
43         saga::ssize_t read_e    (std::string         ext_mode,
44                                 std::string         specification,
45                                 saga::mutable_buffer buffer);
46         saga::ssize_t write_e   (std::string         ext_mode,
47                                 std::string         specification,
48                                 saga::const_buffer  buffer);
49     };
50 }
51 }

```

----- Prototype: saga::filesystem::directory -----

```

1 namespace saga
2 {
3     namespace filesystem
4     {
5         class directory
6             : public saga::name_space::directory
7         {
8         public:
9             directory (saga::session const & s,
10                      saga::url          url,
11                      int                 mode = ReadWrite);
12             directory (saga::url        url,
13                      int                 mode = ReadWrite);
14             directory (saga::object const & o);

```

```
15     directory (void);
16     ~directory (void);
17
18     directory & operator= (saga::object const & o);
19
20     saga::off_t get_size (saga::url url);
21     bool        is_file  (saga::url url);
22     file        open     (saga::url url,
23                          int         flags = Read);
24     directory   open_dir (saga::url url,
25                          int         flags = ReadWrite);
26     }
27 }
28 }
```

### 8.3 Filesystem Details

The described classes are syntactically and semantically POSIX-oriented [?, ?, ?]. Executing large numbers of simple POSIX-like remote data access operations is, however, prone to latency-related performance problems. To allow for efficient implementations, the presented API borrows ideas from GridFTP and other specifications which are widely used for remote data access. These extensions should be seen as just that: optimizations. Be aware that the SAGA adaptors usually implement the POSIX-like `read()`, `write()` and `seek()` methods, and rarely implement the additional optimized methods (a 'NotImplemented' exception is thrown if these are not implemented). The optimizations included here are:

**Scattered I/O** Scattered I/O operations are already defined by POSIX, as `readv()` and `writenv()`. Essentially, these methods represent vector versions of the standard POSIX `read()/write()` methods; the arguments are, basically, vectors of instructions to execute, and buffers to operate upon. In other words, `readv()` and `writenv()` can be regarded as specialized bulk methods, which cluster multiple I/O operations into a single operation. The advantages of such an approach are that it is easy to implement, it is very close to the original POSIX I/O in semantics, and, in some cases, it is even very fast. The disadvantage is that for many small I/O operations (a common occurrence in SAGA use cases), the description of the I/O operations can be larger than the sent, returned or received data.

**Pattern-Based I/O (FALLS)** One approach to address the bandwidth limitation of scattered I/O is to describe the required I/O operations at a more

abstract level. Regularly, repeating patterns of binary data can be described by the so-called 'FAMiLy of Line Segments' (FALLS) [?]. The pattern-based I/O routines in SAGA use such descriptions to reduce the bandwidth limitation of scattered I/O. The advantage of such an approach is that it targets very common data access patterns (at least those commonly found in SAGA use cases). The disadvantages are that FALLS is a paradigm not widely known or used, and that FALLS is by definition, limited to regular patterns of data, hence the inefficiency for more randomized data access.

FALLS was originally introduced for transformations in parallel computing. There is also a parallel filesystem which uses FALLS to describe the file layout. FALLS can be used to describe regular subsets of arrays with a very compact syntax.

FALLS patterns are formed as 5-tuples: "(from,to,stride,rep,(pat))". The **from** element defines the starting offset for the first pattern unit; **to** defines the finishing offset of the first pattern unit; **stride** defines the distance between consecutive pattern units (start to start); and **rep** defines the number of repetitions of the pattern units. The optional 5th element, **pat**, allows to define nested pat-

terns, where the internal pattern defines the unit the outer pattern is applied to (by default it is one byte). As an example, the following FALLS describe the highlighted elements of the matrix in Fig. ??: "(0,17,36,6,(0,0,2,6))": the inner pattern describes a pattern unit of one byte length (from 0 to 0), with a distance of 2 to the next application, and 6 repetitions. These are the 6 bytes per line which are marked. The outer pattern defines the repeated application of the inner pattern, starting at 0, ending at 17 (end of line), distance of 36 (to the beginning of next but one line), and repetition of 6.

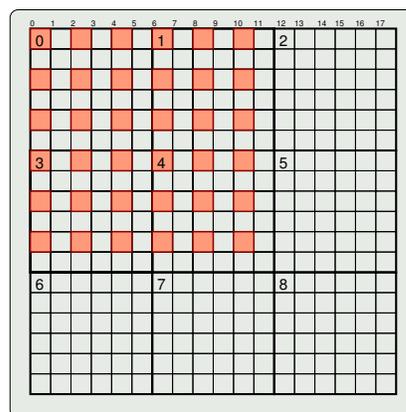


Figure 1: The highlighted elements are defined by "(0,17,36,6,(0,0,2,6))".

**Extended I/O** GridFTP (which was designed for a similar target domain) introduced an additional remote I/O paradigm, that of Extended I/O operations.

In essence, the Extended I/O paradigm allows the formulation of I/O requests using custom strings, which are not interpreted on the client, but on the server side; these can be expanded to arbitrarily complex sets of I/O operations. The type of I/O request encoded in the string is called **mode**. A server may support one or many of these extended I/O modes. Whereas the approach is very flexible

and powerful and has proven its usability in GridFTP, a disadvantage is that it requires very specific infrastructure to function, i.e., it requires a remote server instance which can interpret opaque client requests. Additionally, no client side checks or optimizations on the I/O requests are possible. Also, the application programmer needs to estimate the size of the data to be returned in advance, which in some cases is very difficult.

The three described operations have, if compared to each other, increasing semantic flexibility, and are increasingly powerful for specific use cases. However, they are also increasingly difficult to implement and support in a generic fashion. It is up to the SAGA adaptors and the specific use cases, to determine the level of I/O abstraction that serves the application best and that can be best supported in the target environment.

### **Enum flags**

The enum `flags` are inherited from the `namespace` package. A number of file specific flags are added to it. All added flags are used for the opening of `file` and `directory` instances, and are not applicable to the operations inherited from the `namespace` package.

#### **Truncate**

Upon opening, the file is truncated to length 0, i.e., a following `read()` operation will never find any data in the file. That flag does not apply to directories.

#### **Append**

Upon opening, the file pointer is set to the end of the file, i.e., a following `write()` operation will extend the size of the file. That flag does not apply to directories.

#### **Read**

The file or directory is opened for reading—that does not imply the ability to write to the file or directory.

#### **Write**

The file or directory is opened for writing—that does not imply the ability to read from the file or directory.

#### **ReadWrite**

The file or directory is opened for reading and writing.

#### **Binary**

Some operating systems (notably windows-based systems) distinguish between binary and non-binary modes—this flag mimics that behaviour.

### Class `iovec`

The `iovec` class inherits the `saga::buffer` class, and three additional state attributes: `offset`, `len_in` and `len_out` (with the latter one being read-only). With that addition, the new class can be used very much the same way as the `iovec` structure defined by POSIX for `readv/writev`; the buffer `len_in` is being interpreted as the POSIX `iov_len`, i.e., the number of bytes to read/write.

If `len_in` is not specified, that length is set to the size of the buffer. For application-managed buffers, it is a `BadParameter` error if `len_in` is specified to be larger than size, (see Section ?? for details on buffer memory management). Before an `iovec` instance is used, its `len_in` must be set to a non-zero value; otherwise its use will cause a `BadParameter` exception.

After a `read_v()` or `write_v()` operation completes, `len_out` will report the number of bytes read or written. Before completion, the SAGA implementation will report `len_out` to be `-1`.

## 9 Using the Replica Package

### 9.1 Quick Introduction

Another package inheriting the namespace package is the replica management in `saga::replica`. It introduces logical files, which are namespace entries that have a number of locations (URLs) attached pointing to identical physical copies of the same file. The replica package is again a very simple extension of the namespace package: the class `logical_file` is a namespace `entry` class with a number of additional methods which allow to manage the list of attached URLs:

```
----- Managing replica locations -----
1 // open a logical file
2 saga::url u ("/replica_1");
3 saga::replica::logical_file lf (u, saga::replica::Create
4                               | saga::replica::ReadWrite);
5
6 // Add a replica location, replicate the file
7 lf.add_location (saga::url ("file://localhost/etc/passwd"));
8 lf.replicate (saga::url ("file://localhost/tmp/passwd"),
9              saga::replica::Overwrite);
10
11 // list all locations
12 std::vector <saga::url> replicas = lf.list_locations ();
13
14 for ( unsigned int i = 0; i < replicas.size (); i++ )
15 {
16     std::cout << "replica: " << replicas[i] << std::endl;
17 }
18
19 // remove the first location
20 lf.remove_location (replicas[0]);
```

Managing replicas is fairly simple, because it only covers methods to list, add, delete, and update replica locations on logical files. Additionally, logical files and logical directories can have meta data attached, which are sets of string-based key-value pairs<sup>11</sup>:

```
----- Managing replica meta data -----
1 // open a logical file
2 saga::url u ("/replica_1");
3 saga::replica::logical_file lf (u, saga::replica::Create
4                               | saga::replica::ReadWrite);
```

<sup>11</sup>For brevity, the example does not distinguish between scalar and vector attributes.

```
5 // get all attributes
6 std::vector <std::string> keys = lf.list_attributes ();
7
8 // print the keys and values
9 for ( unsigned int i = 0; i < keys.size (); i++ )
10 {
11     std::string key = keys[i];
12     std::string val;
13
14     if ( lf.attribute_is_vector (key) )
15     {
16         std::vector <std::string> vals = lf.get_vector_attribute (key);
17         val = vals[0] + " ...";
18     }
19     else
20     {
21         val = lf.get_attribute (key);
22     }
23     std::cout << key << " -> " << val << std::endl;
24 }
```

The attentive reader will notice that the attribute management is in accordance with the attribute management part of the SAGA Look & Feel.

## 9.2 Reference

```
----- Prototypes: saga::replica -----
1 namespace saga
2 {
3     namespace replica
4     {
5         namespace metrics
6         {
7             char const * const logical_file_modified
8                 = "logical_file.Modified";
9             char const * const logical_file_deleted
10                = "logical_file.Deleted";
11             char const * const logical_directory_created_entry
12                = "logical_directory.CreatedEntry";
13             char const * const logical_directory_modified_entry
14                = "logical_directory.ModifiedEntry";
15             char const * const logical_directory_deleted_entry
16                = "logical_directory.DeletedEntry";
17         }
18
19         enum flags
```

```

20     {
21         Unknown      = /* -1, */  saga::name_space::Unknown      ,
22         None         = /*  0, */  saga::name_space::None         ,
23         Overwrite    = /*  1, */  saga::name_space::Overwrite    ,
24         Recursive    = /*  2, */  saga::name_space::Recursive    ,
25         Dereference  = /*  4, */  saga::name_space::Dereference  ,
26         Create       = /*  8, */  saga::name_space::Create       ,
27         Exclusive    = /* 16, */  saga::name_space::Exclusive    ,
28         Lock         = /* 32, */  saga::name_space::Lock         ,
29         CreateParents = /* 64, */  saga::name_space::CreateParents ,
30                                     // 128,          reserved for Truncate
31                                     // 256,          reserved for Append
32         Read         =   512,
33         Write        =  1024,
34         ReadWrite    =  1036,
35                                     // 2048,          reserved for Binary
36     };
37
38     class logical_file
39     : public saga::name_space::entry,
40       public saga::attributes
41     {
42     public:
43         logical_file (session const & s,
44                     saga::url      url,
45                     int             mode = Read);
46         logical_file (saga::url     url,
47                     int             mode = Read);
48         logical_file (saga::object const & o);
49         logical_file (void);
50         ~logical_file (void);
51
52         logical_file & operator= (saga::object const & o);
53
54         void add_location (saga::url url);
55         void remove_location (saga::url url);
56         void update_location (saga::url old,
57                              saga::url new);
58         std::vector<saga::url>
59             list_locations (void);
60         void replicate (saga::url url,
61                       int      flags = None);
62     };
63
64     class logical_directory
65     : public saga::name_space::directory,
66       public saga::attributes
67     {
68     public:
69         logical_directory (saga::session const & s,

```

```
70         saga::url          url,
71         int                mode = ReadWrite);
72     logical_directory (saga::url          url,
73         int                mode = ReadWrite);
74     logical_directory (saga::object const & o);
75     logical_directory (void);
76     ~logical_directory (void);
77
78     logical_directory & operator=(saga::object const& o);
79
80     bool is_file      (saga::url url);
81     std::vector<saga::url>
82         find          (std::string      name_pattern,
83         std::vector <std::string> key_pattern,
84         int            flags = Recursive);
85
86     saga::replica::logical_file
87         open          (saga::url          url,
88         int            flags = Read);
89     saga::replica::logical_directory
90         open_dir      (saga::url          url,
91         int            flags = None);
92     };
93 }
94 }
```

### 9.3 Replica Details

## 10 Using the Namespace Package

### 10.1 Quick Introduction

Namespaces are used for a wide variety of computer subsystems: they are used to organize files, to manage domain names, web content, etc. Hierarchical namespaces are prevalent for managing files (indeed, most file systems provide a hierarchical namespace, the directory tree). SAGA provides the `saga::name_space` package to navigate and manipulate such namespaces:

```

----- Navigating a name space -----
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      // open a namespace directory
6      saga::url u (std::string ("file://localhost/") + getenv ("PWD"));
7      saga::name_space::directory d (u);
8
9      // list the contents
10     std::vector <saga::url> entries = d.list ();
11
12     // print the entries, and their type
13     for ( unsigned int i = 0; i < entries.size (); i++ )
14     {
15         std::string type;
16         // get some details for the entry
17         if ( d.is_dir (entries[i]) )
18         {
19             type = "/";
20         }
21         // if a link (symbolic)
22         else if ( d.is_link (entries[i]) )
23         {
24             type = " -> " + d.read_link (entries[i]).get_string();
25         }
26         // print the info
27         std::cout << entries[i] << type << std::endl;
28     }
29
30     return (0);
31 }
```

This is a poor man's `ls`! It needs only five SAGA calls, which are all very simple. Note that a directory has a 'cwd', a Current Working Directory. Calling `cd()`

on a directory changes that. Relative file names are always interpreted with respect to that cwd.

So, the `saga::name_space` package provides two classes: `directory` and `entry` (note that `directory` inherits `entry`). On entries, you can perform `copy()`, `link()`, `move()`, and `remove()`. You can also inspect entries, with `is_dir()`, `is_entry()`, `is_link()` and `read_link()`.

Directories provide the same operations, but with an additional `source` argument, as shown below:

```
entry.copy (target);           // copy entry    to target
dir.copy (target);             // copy dir      to target
dir.copy (source, target);     // copy dir/source to target
```

But the `directory` copy should be recursive, of course:

```
dir.copy (target, saga::name_space::Recursive);
```

Additionally, the `directory` has two open methods:

```
saga::entry e = dir.open (saga::url (entry_name));
saga::directory d = dir.open_dir (saga::url (dir_name) );
```

## 10.2 Reference

```
----- Prototype: saga::namespace::flags -----
1  namespace saga
2  {
3      namespace name_space
4      {
5          enum flags
6          {
7              Unknown          = -1,
8              None              =  0,
9              Overwrite        =  1,
10             Recursive         =  2,
11             Dereference       =  4,
12             Create            =  8,
13             Exclusive         = 16,
14             Lock              = 32,
15             CreateParents     = 64
16         };
17     }
18 }
```

Prototype: `saga::namespace::entry`

```
1 namespace saga
2 {
3     namespace name_space
4     {
5         class entry
6             : public saga::object,
7               public saga::monitorable,
8               public saga::permissions
9         {
10            entry (session const & s,
11                  saga::url      url,
12                  int             mode = None);
13            entry (saga::url      url,
14                  int             mode = None);
15            entry (void);
16            ~entry (void);
17
18            // inspection methods
19            saga::url get_url  (void) const;
20            saga::url get_cwd  (void) const;
21            saga::url get_name (void) const;
22            saga::url read_link (void) const;
23
24            bool      is_dir   (void) const;
25            bool      is_entry (void) const;
26            bool      is_link  (void) const;
27
28            // management methods
29            void      copy     (saga::url target,
30                              int flags = saga::name_space::None);
31            void      link     (saga::url target,
32                              int flags = saga::name_space::None);
33            void      move     (saga::url target,
34                              int flags = saga::name_space::None);
35            void      remove   (int flags = saga::name_space::None);
36            void      close    (double timeout = 0.0);
37        };
38    }
39 }
40
```

Prototype: `saga::namespace::directory`

```
1 namespace saga
2 {
3     namespace name_space
4     {
5         class directory
```

```
6         : public saga::name_space::entry
7     {
8     public:
9         directory (session const & s,
10                saga::url      url,
11                int             mode = None);
12         directory (saga::url    url,
13                int             mode = None);
14         directory (void);
15         ~directory (void);
16
17         void      change_dir (saga::url  target)
18         std::vector <saga::url>
19             list      (std::string pattern = "*",
20                      int             flags  = None) const;
21         std::vector <saga::url>
22             find      (std::string pattern,
23                      int             flags  = Recursive) const;
24
25         saga::url  read_link (saga::url  url) const;
26         bool       exists   (saga::url  url) const;
27         bool       is_dir   (saga::url  url) const;
28         bool       is_entry (saga::url  url) const;
29         bool       is_link  (saga::url  url) const;
30         unsigned   int      get_num_entries
31             (void) const;
32         saga::url  get_entry (unsigned int entry) const;
33         void       copy      (saga::url  source_url,
34                             saga::url  dest_url,
35                             int         flags = None);
36         void       link      (saga::url  source_url,
37                             saga::url  dest_url,
38                             int         flags = None);
39         void       move      (saga::url  source_url,
40                             saga::url  dest_url,
41                             int         flags = None);
42         void       remove    (saga::url  url,
43                             int         flags = None);
44         void       make_dir  (saga::url  url,
45                             int         flags = None);
46         entry      open      (saga::url  url,
47                             int         flags = None);
48         directory  open_dir  (saga::url  url,
49                             int         flags = None);
50     };
51     } // namespace_dir
52 } // namespace saga
53
```

### 10.3 Details

The notion of a namespace is shared by several SAGA packages: the file, the replica and the advert packages. All allow to manage entities which are organized in a hierarchical namespace. The namespace package is abstracting the management of that hierarchy, and leaves only the entity specific operations (File-IO, Replica Management, Advert Creation) to the various packages.

**HINT:**

Namespace operations handle namespace entries as opaque, and are never able to look 'inside'. To do that, use one of the derived packages, which add exactly that functionality, specific to the type of namespace they represent.

The namespace package introduces two classes: `saga::namespace::entry` and `saga::namespace::directory` (which is an entry, i.e., inherits from `entry`). Both classes refer to entities which are specified by a pathname, typically a URL. Several calls additionally allow to refer to entries with wildcards, similar to the shell wildcards known by POSIX (see `glob(7)`)<sup>12</sup>.

The namespace package allows to create and delete entries, to copy, move or link entries, and to inspect entries and directories.

**FIXME: add infos about available attributes and metrics**

---

<sup>12</sup>Namespace entries can also have permissions, just as in most file systems. For details on permissions, see Section ??.

## 11 Using the Job Package

### 11.1 Quick Introduction

For an overwhelming number of use cases, job submission and management is the most important aspect of Grid computing. The `saga::job` package provides the respective functionality. There are four classes:

- `saga::job::description`  
describes the properties of a job to be submitted,
- `saga::job::service`  
represents a service which accepts job descriptions for submission,
- `saga::job::job`  
represents the submitted job, and
- `saga::job::self`  
which represents the application itself.

The most basic code example is the following:

```
Job submission
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      saga::job::service js;
6      saga::job::job j = js.run_job ("localhost", "/bin/sleep 3");
7      j.wait ();
8
9      return (0);
10 }
```

Yes, it is that easy. But what about the job's I/O, can we handle that as well? And what about job state information <F2>? Let's see:

```
Job submission and job I/O
1  saga::job::service js;
2  saga::job::ostream in;
3  saga::job::istream out;
4  saga::job::istream err;
5
6  saga::job::job j = js.run_job ("localhost",
7                                "/bin/date",
```

```
8             in, out, err);
9
10          saga::job::state state = j.get_state ();
11
12          while ( state != saga::job::Done &&
13                 state != saga::job::Failed &&
14                 state != saga::job::Canceled )
15          {
16              char buffer[256];
17
18              // get stdin
19              out.read (buffer, sizeof (buffer));
20              if ( out.gcount () > 0 )
21              {
22                  std::cout << std::string (buffer, out.gcount ());
23              }
24              if ( out.fail () )
25              {
26                  break;
27              }
28              usleep (10000);
29              state = j.get_state ();
30          }
```

That is basically the same as above, but we catch the job's I/O channel as it is getting created. The remainder of the example code (lines 12ff) is just watching the job state and the output stream.

The `run_job()` method is just a shortcut. It actually provides exactly the following functionality, just with less code, as the example below:

```
----- run_job() expanded -----
1  namespace sa = saga::attributes;
2  namespace sja = saga::job::attributes;
3
4  std::string exe ("/bin/date");
5
6  std::vector <std::string> hosts;
7  hosts.push_back ("localhost");
8
9  saga::job::description jd;
10
11  jd.set_attribute      (sja::description_interactive, sa::common_true);
12  jd.set_attribute      (sja::description_executable, exe);
13  jd.set_vector_attribute (sja::description_candidatehosts, hosts);
14
15  saga::job::service js;
```

```
16
17     saga::job::job j = js.create_job (jd);
18
19
20     saga::job::ostream in = j.get_stdin ();
21     saga::job::istream out = j.get_stdout ();
22     saga::job::istream err = j.get_stderr ();
23
24     // job is in 'New' state here, we need to run it
25     j.run ();
26
27     saga::job::state state = j.get_state ();
28
29     while ( state != saga::job::Done &&
30            state != saga::job::Failed &&
31            state != saga::job::Canceled )
32     {
33         char buffer[256];
34
35         // get stdin
36         out.read (buffer, sizeof (buffer));
37         if ( out.gcount () > 0 )
38         {
39             std::cout << std::string (buffer, out.gcount ());
40         }
41         if ( out.fail () )
42         {
43             break;
44         }
45         usleep (10000);
46         state = j.get_state ();
47     }
```

This long way exposes some additional details:

- The `job_description` class allows specifying attributes of the job.<sup>13</sup>
- The `stdio` channels of a job can be obtained before the job is running, to avoid race conditions.
- Only 'Interactive' jobs will provide `stdio` channels.

Table ?? shows the job description attributes available in SAGA. Note that these attributes are closely related to the respective keys of the JSDL standard [?].

<sup>13</sup>Only the 'Executable' attribute is mandatory.

|                     |                 |                        |
|---------------------|-----------------|------------------------|
| Executable          | string          | what executable to run |
| Arguments           | list of strings | the job arguments      |
| CandidateHosts      | list of strings | where to run           |
| SPMDVariation       | string          | MPI type (if any)      |
| TotalCPUCount       | int             | # CPUs, total          |
| NumberOfProcesses   | int             | # processes, total     |
| ProcessesPerHost    | int             | # processes per host   |
| ThreadsPerProcess   | int             | # threads per process  |
| Environment         | list of strings | environment vars       |
| WorkingDirectory    | string          | working directory      |
| Interactive         | bool            | usage type             |
| Input               | string          | stdin source file      |
| Output              | string          | stdout log file        |
| Error               | string          | stderr log file        |
| FileTransfer        | list of strings | file staging           |
| Cleanup             | bool            | cleanup-after-run      |
| JobStartTime        | datetime        | when to start job      |
| TotalCPUTime        | int             | how long will job run  |
| TotalPhysicalMemory | int             | memory the job needs   |
| CPUArchitecture     | string          | architecture to run on |
| OperatingSystemType | string          | OS to run on           |
| Queue               | string          | queue to submit to     |
| JobContact          | list of strings | notification contact   |

Table 2: Job description attributes available in SAGA.

Line 10 in the job I/O example above obtains the job state:

```
saga::job::state state = j.get_state ();
```

That implies that `saga::job` is a stateful object. The job state model is shown in Figure ??, which also shows what method calls cause which state transitions.

## 11.2 Reference

```

----- Prototypes: saga::job -----
1  namespace saga
2  {
3      namespace job
4      {
5          typedef char const * char ccc_type;
6
7          namespace attributes
8          {
9              // job description attributes

```

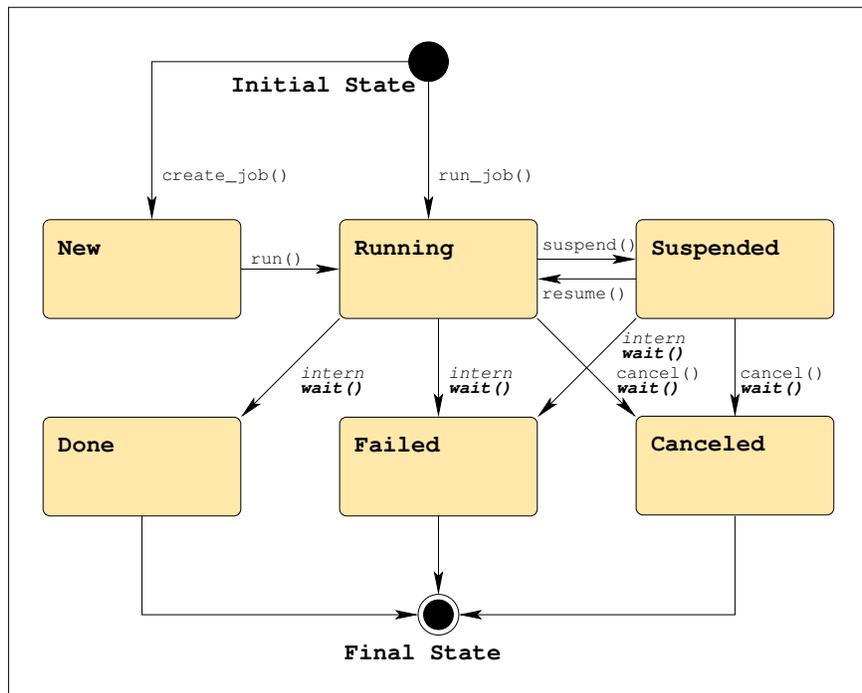


Figure 2: The SAGA job state model.

```

10 ccc_type description_executable = "Executable";
11 ccc_type description_arguments = "Arguments";
12 ccc_type description_environment = "Environment";
13 ccc_type description_workingdirectory = "WorkingDirectory";
14 ccc_type description_interactive = "Interactive";
15 ccc_type description_input = "Input";
16 ccc_type description_output = "Output";
17 ccc_type description_error = "Error";
18 ccc_type description_filetransfer = "FileTransfer";
19 ccc_type description_cleanup = "Cleanup";
20 ccc_type description_jobstarttime = "JobStartTime";
21 ccc_type description_totalcputime = "TotalCPUTime";
22 ccc_type description_totalphysicalmemory = "TotalPhysicalMemory";
23 ccc_type description_cpuarchitecture = "CPUArchitecture";
24
25 ccc_type cpuarchitecture_sparc = "sparc";
26 ccc_type cpuarchitecture_powerpc = "powerpc";
27 ccc_type cpuarchitecture_x86 = "x86";
28 ccc_type cpuarchitecture_x86_32 = "x86_32";
29 ccc_type cpuarchitecture_x86_64 = "x86_64";
30 ccc_type cpuarchitecture_parisc = "parisc";

```

```

31     ccc_type cpuarchitecture_mips           = "mips";
32     ccc_type cpuarchitecture_ia64         = "ia64";
33     ccc_type cpuarchitecture_arm          = "arm";
34     ccc_type cpuarchitecture_other        = "other";
35
36     ccc_type description_operatingsystemtype = "OperatingSystemType";
37     ccc_type description_candidatehosts     = "CandidateHosts";
38     ccc_type description_queue             = "Queue";
39     ccc_type description_jobcontact        = "JobContact";
40     ccc_type description_spmdivariation    = "SPMDVariation";
41     ccc_type description_totalcpucount     = "TotalCPUCount";
42     ccc_type description_numberofprocesses = "NumberOfProcesses";
43     ccc_type description_processesperhost = "ProcessesPerHost";
44     ccc_type description_threadsperprocess = "ThreadsPerProcess";
45 }
46
47 class description
48     : public saga::object,
49     public saga::attributes
50 {
51     public:
52     description (void);
53     ~description (void);
54 };
55
56
57 class service
58     : public saga::object
59 {
60     public:
61     service (saga::session const & s,
62             saga::url          rm = "");
63     service (saga::url          rm = "");
64     ~service (void);
65
66     saga::job::job create_job (description job_desc);
67
68     saga::job::job run_job      (std::string  hostname,
69                                std::string  commandline,
70                                ostream      & stdin_stream,
71                                istream      & stdout_stream,
72                                istream      & stderr_stream);
73
74     std::vector <std::string>
75         list      (void);
76     saga::job::job get_job      (std::string  job_id);
77     saga::job::self get_self    (void);
78 };
79
80 namespace attributes

```

```
81     {
82         // read only job attributes
83         ccc_type jobid          = "JobID";
84         ccc_type executionhosts = "ExecutionHosts";
85         ccc_type created        = "Created";
86         ccc_type started        = "Started";
87         ccc_type finished       = "Finished";
88         ccc_type workingdirectory = "WorkingDirectory";
89         ccc_type exitcode       = "ExitCode";
90         ccc_type termsig        = "Termsig";
91     }
92
93     namespace metrics
94     {
95         // job metrics
96         ccc_type statedetail    = "job.StateDetail";
97         ccc_type signal         = "job.Signal";
98         ccc_type cputimelimit   = "job.CPUTimeLimit";
99         ccc_type memoryuse      = "job.MemoryUse";
100        ccc_type vmemoryuse      = "job.VmemoryUse";
101        ccc_type performance     = "job.Performance";
102    }
103
104    enum state
105    {
106        // job state
107        Unknown = saga::task_base::Unknown, // -1
108        New     = saga::task_base::New,     // 0
109        Running = saga::task_base::Running, // 1
110        Failed  = saga::task_base::Failed,  // 2
111        Done    = saga::task_base::Done,    // 3
112        Canceled = saga::task_base::Canceled, // 4
113        Suspended = 5
114    };
115
116    class job : public saga::task,
117               public saga::attributes,
118               public saga::permissions,
119               public saga::sync
120    {
121    public:
122        // no constructor
123        ~job (void);
124
125        job & operator= (saga::object const & o);
126
127        std::string get_job_id    (void);
128        state       get_state     (void);
129        description get_description (void);
130        ostream     get_stdin     (void);
```

```
131         istream    get_stdout    (void);
132         istream    get_stderr    (void);
133
134         void       suspend       (void);
135         void       resume        (void);
136
137         void       checkpoint     (void);
138         void       migrate        (description job_desc);
139         void       signal         (int         signal);
140     };
141 }
142 }
```

### 11.3 Job Details

The SAGA Job API covers four classes: `saga::job_description`, `saga::job-service`, `saga::job`, and `saga::job_self`. The job description class is nothing more than a container for a well-defined set of attributes, which, using JSDL [?] based keys, defines the job to be started, and its runtime and resource requirements. The job server represents a resource management endpoint which allows the starting and insertion of jobs.

The job class itself is central to the API, and represents an application instance running under the management of a resource manager. The `job_self` class IS-A job. Its purpose is to represent the current SAGA application, which allows for a number of use cases with applications that actively interact with the grid infrastructure, for example, to migrate itself, or to set new job attributes.

The job class inherits the `saga::task` class (Sec. ??), and uses its methods to `run()`, to `wait()` for, and to `cancel()` jobs. The inheritance also allows for the management of large numbers of jobs in task containers. Additional methods provided by the `saga::job` class relate to the `Suspended` state (which is not available on tasks), and provide access to the job's standard I/O streams, and to more detailed status information. The `saga::job` iostreams are available as `saga::job::istream` and `saga::job::ostream` instances, which all inherit from `std::stream`.

#### 11.3.1 Job State Model

The SAGA job state diagram is shown in Figure ?? . It is an extension of the `saga::task` state diagram (Figure ??), and extends the state diagram with a 'Suspended' state, which the job can enter/leave using the `suspend()/resume()` calls.

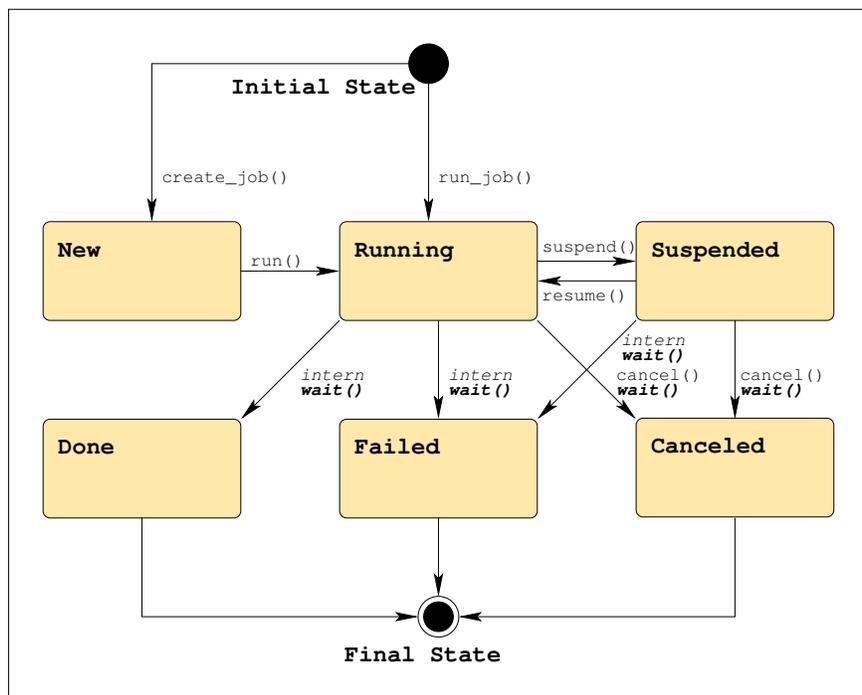


Figure 3: The SAGA job state model extends the SAGA task state model with a 'Suspended' state, and additional transitions.

SAGA maps the native back-end state model onto the SAGA state model. The SAGA state model is simple enough to allow a straightforward mapping in most cases. For some applications, access to the native back-end state model is useful, though—for that reason, an additional metric named 'StateDetail' allows to query the native job state.

State details in SAGA are formatted as follows:

```
'<model>:<state>'
```

with valid models being "BES", "DRMAA", or other implementation specific models. For example, a state detail for the BES state 'StagingIn' would be rendered as 'BES:StagingIn', and would be a substate of **Running**. If no state details are available, the metric is still available, but it has always an empty string value.

### 11.3.2 File Transfer Specifications

The syntax of a file transfer directive for the job description is modeled on the LSF syntax (LSF stands for *Load Sharing Facility*, a commercial job scheduler by Platform Computing). It has the general syntax:

```
local_file operator remote_file
```

Both the `local_file` and the `remote_file` can be URLs. If they are not URLs, but full or relative pathnames, then the `local_file` is relative to the host where the submission is executed, and the `remote_file` is evaluated on the execution host of the job.

The operator is one of the following four:

- '>' copies the local file to the remote file before the job starts. Overwrites the remote file if it exists.
- '>>' copies the local file to the remote file before the job starts. Appends to the remote file if it exists.
- '<' copies the remote file to the local file after the job finishes. Overwrites the local file if it exists.
- '<<' copies the remote file to the local file after the job finishes. Appends to the local file if it exists.

### 11.3.3 Command Line Specification

The `run_job()` method of the `saga::job_service` class accepts a string parameter which constitutes a command line to be executed on a remote resource. The parsing of that command line follows the following rules:

- Elements are either delimited by white space (a space, or a tab), or are delimited by not-escaped double quotes.
- The escape character for double quotes is the backslash (which is also used to escape the backslash itself).
- The first element is used as the executable name; all other elements are treated as job arguments.

The following command line parameter

```
/bin/cp "/tmp/my file with spaces" /data/
```

is thus invoking the executable `/bin/cp`, with two arguments.

### 11.3.4 Job Identifiers

The SAGA JobID is treated as an opaque string in the SAGA API, and is structured as:

```
'[backend url]-[native id]'
```

For example, a job submitted to the host `remote.host.net` via `ssh` (whose daemon runs on port 22), and having the POSIX PID 1234, would get the job id:

```
'[ssh://remote.host.net:22/]-[1234]'
```

The implementation may free the resources used for the job, and hence may invalidate a JobID after a successful wait on the job, and after the application received the job status information at least once.

A JobID may be unknown until the job enters the `Running` state, as the backend will often not assign IDs to jobs which are not yet running. In such cases, the value of the JobID attribute is empty. The job will, however, retain its JobID after it enters a final state.

## 12 Using the Stream Package

### 12.1 Quick Introduction

(FIXME: english)

Inter-process communication (IPC) is historically central to parallel and distributed computing, and numerous approaches exist for the wide variety of use cases. Streams are amongst the most well-known and most simple of these paradigms, and are also provided by SAGA. The concept is very close to that of BSD sockets: a `stream_service` can be listened to (calling `serve()`), and a connection client `stream` results in a `stream` instances on both ends, upon which the application can call `read()` and `write()`:

```
Stream server
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      // reusable io buffer
6      saga::buffer buf;
7
8      // create a stream_server and listen for clients
9      saga::stream::server ss (saga::url (argv[1]));
10
11     while ( 1 )
12     {
13         // wait for incoming client
14         saga::stream::client s = ss.serve ();
15
16         // read data from client
17         s.read (buf);
18     }
19 }
```

```
Stream client
1  #include <saga/saga.hpp>
2
3  int main (int argc, char** argv)
4  {
5      // create client stream and connect to server
6      saga::stream::client s (saga::url (argv[1]));
7
8      s.connect ();
9
10     // send some data
```

```
11     s.write (saga::buffer ("Hello"));
12 }
```

Note that SAGA is silent about the used wire protocol. In particular, server and client are responsible for picking a compatible transfer mechanism. Also, the bootstrapping mechanism is out of scope in the SAGA stream package, i.e., both sides need to agree on the URL for the server endpoint out-of-band. Both issues will be addressed in future SAGA API extensions (message API and advert API extension).

In particular, for stream-based communication, asynchronous operations and notifications are extremely useful programming paradigms. Both are provided within SAGA, more details can be found in the later description of the SAGA Look & Feel (`saga::monitoring` and `saga::task` package).

## 12.2 Reference

```
----- Prototypes: saga::stream -----
1 namespace saga
2 {
3     namespace stream
4     {
5         namespace attributes
6         {
7             char const * const stream_bufsize      = "Bufsize";
8             char const * const stream_timeout     = "Timeout";
9             char const * const stream_blocking    = "Blocking";
10            char const * const stream_compression  = "Compression";
11            char const * const stream_nodelay     = "Nodelay";
12            char const * const stream_reliable    = "Reliable";
13        }
14
15        namespace metrics
16        {
17            char const * const stream_state        = "stream.State";
18            char const * const stream_read        = "stream.Read";
19            char const * const stream_write       = "stream.Write";
20            char const * const stream_exception   = "stream.Exception";
21            char const * const stream_dropped     = "stream.Dropped";
22
23            char const * const server_clientconnect = "server.ClientConnect";
24        }
25
26        enum state
27        {
28            Unknown      = -1,
```

```
29     New          = 1,
30     Open         = 2,
31     Closed       = 3,
32     Dropped      = 4,
33     Error        = 5
34 };
35
36 enum activity
37 {
38     Read          = 1,
39     Write         = 2,
40     Exception     = 4
41 };
42
43 class stream
44     : public saga::object,
45     public saga::attributes,
46     public saga::monitorable
47 {
48     public:
49     stream (saga::session const & s,
50            saga::url          url = saga::url ());
51     stream (saga::url          url);
52     stream (void);
53     stream (saga::object const & o);
54     ~stream (void);
55
56     stream & operator= (saga::object const & o);
57
58     saga::url  get_url      (void) const;
59     saga::context get_context (void) const;
60
61     saga::context connect   (void);
62
63     std::vector <activity>
64         wait      (activity      what,
65                  double        timeout = -1.0);
66     void         close    (double        timeout = 0.0);
67
68     saga::ssize_t read     (saga::mutable_buffer buffer,
69                            saga::ssize_t      length = 0);
70     saga::ssize_t write   (saga::const_buffer  buffer,
71                            saga::ssize_t      length = 0);
72 };
73
74 class server
75     : public saga::object,
76     public saga::monitorable,
77     public saga::permissions
78 {
```

```
79     public:
80         server (saga::session const & s,
81               saga::url url = saga::url ());
82         server (saga::url url);
83         server (void);
84         server (saga::object const & o);
85         ~server (void);
86
87         server & operator= (saga::object const & o);
88
89         saga::url get_url (void) const;
90         saga::stream::stream serve (double timeout = 0.0);
91         void close (double timeout = 0.0);
92     };
93 }
94 }
```

## 12.3 Stream Details

The SAGA streams package allows establishing the simplest possible authenticated socket connection, with hooks to support application level authorization, and encryption schemes. The stream API has the following characteristics:

1. It is not performance-oriented: If performance is required, then it is better to program directly against the API's existing performance-oriented protocols like GridFTP or XIO.
2. It does not attempt to create a programming paradigm that diverges very far from baseline BSD sockets, Winsock, or Java Sockets.

This API greatly reduces the complexity of establishing authenticated socket connections in order to communicate with remotely located components. However, it provides very limited functionality and is thus suitable for applications that do not have very sophisticated requirements (as per 80-20 rule). It is envisaged that as applications become progressively more sophisticated, they will gradually move to more sophisticated native APIs in order to support those needs. Later SAGA versions may offer higher level communication abstractions, such as messages.

### 12.3.1 Endpoint URLs

The SAGA stream API uses URLs to specify connection endpoints. These URLs are supposed to allow SAGA implementations to be interoperable. For example,

the URL

```
tcp://remote.host.net:1234/
```

is supposed to signal that a standard `tcp` connection can be established with host `remote.host.net` on port 1234. No matter what the specified URL scheme is, the SAGA stream API implementation **MUST** have the same semantics on an API level, i.e., behave like a reliable byte-oriented data stream.

### 12.3.2 Usage

Just as for BSD sockets, a stream communication channel is established by creating a serving part (BSD: listening socket), and a client party (BSD: connecting socket). After connecting both parties, reading and writing on the stream allows to exchange data.

Stream example - server side

```
1  {
2    // set up the serving endpoint on port 1234
3    saga::stream::server server (saga::url ("tcp://localhost:1234"));
4
5    // wait for incoming connections
6    saga::stream stream = server.serve ();
7
8    // if one arrived, greet the client
9    stream.write (saga::buffer ("Hello World!", 13));
10 }
```

Stream example - client side

```
1  {
2    // setup stream for the server endpoint
3    saga::stream::stream (saga::url ("tcp://remotehost:1234"));
4
5    // connect to the server
6    stream.connect ();
7
8    // read the greeting message
9    saga::buffer buf (13);
10   stream.read (buf);
11
12   // print it
13   std::cout << buf.get_data () << std::endl;
14 }
```

**FIXME:** add details about async ops, and AAA

## 13 Using the RPC Package

### 13.1 Quick Introduction

Remote Procedure Calls are the second IPC mechanism provided in SAGA. The respective `saga::rpc` package is modeled after the GridRPC standard [?]. The concept could not be simpler: an `rpc` handle instance (`saga::rpc::rpc`) has a single method, `call()`, which invokes the respective remote operation. The call accepts a stack of `In`, `Out`, and `InOut` parameters.

### 13.2 Reference

Prototypes: `saga::rpc`

```
1
2 namespace saga
3 {
4     namespace rpc
5     {
6         enum io_mode
7         {
8             Unknown = -1,
9             In      = 1,
10            Out     = 2,
11            InOut   = In | Out
12        };
13
14        class parameter
15            : public saga::mutable_buffer
16        {
17        public:
18            parameter (void * data = 0,
19                      saga::ssize_t size = -1,
20                      io_mode mode = In,
21                      buffer_deleter cb = default_buffer_deleter);
22
23            ~parameter (void);
24
25            io_mode get_mode() const;
26        };
27
28        class rpc
29            : public saga::object,
30              public saga::permissions
31        {
32        public:
33            rpc (saga::session const & s,
```

```

34         saga::url          name = saga::url ();
35         rpc (std::string  const & name);
36         rpc (void);
37         ~rpc (void);
38
39         void call (std::vector <parameter> parameters);
40         void close (double timeout = 0.0);
41     };
42 }
43 }
44
45

```

### 13.3 Details

The `rpc` class constructor is used to initialize the remote function handle. Be aware that this process may involve connection setup, service discovery, and other remote interactions! In the constructor, the remote procedure to be invoked is specified by a URL, with the syntax:

```
gridrpc://server.net:1234/my_function
```

with the elements responding to:

|                          |   |                     |   |                                     |
|--------------------------|---|---------------------|---|-------------------------------------|
| <code>gridrpc</code>     | - | <code>scheme</code> | - | identifying a grid rpc operation    |
| <code>server.net</code>  | - | <code>server</code> | - | server host serving the rpc call    |
| <code>1234</code>        | - | <code>port</code>   | - | contact point for the server        |
| <code>my_function</code> | - | <code>name</code>   | - | name of the remote method to invoke |

All elements can be empty, which allows the implementation to fall back to invoke a default remote method.

Furthermore, the `rpc` class offers one method, `call()`, which invokes the remote procedure, and returns the respective `Out` and `InOut` parameters.

```

----- Remote Procedure Call Example -----
1  {
2  // initialize the rpc handle
3  saga::rpc::rpc rm(saga::url ("rpc://remote.host.net:31415/get_pi"));
4
5  // initialize the parameter stack
6  saga::rpc::parameter pi (NULL, -1, saga::rpc::Out);
7  std::vector <saga::rpc::parameter> parameters;
8  parameters.push_back (pi);

```

```
9
10 // invoke the remote procedure
11 rm.call (parameters);
12
13 // when completed, the output parameters are available
14 std::cout << "Pi equals " << pi.get_data () << std::endl;
15 }
```

The argument and return value handling is very basic, and reflects the traditional scheme for remote procedure calls, that is, an array of structures act as a parameter stack. For each element of the vector, the `parameter` struct describes its data `buffer`, the `size` of that buffer, and its input/output mode.

The `mode` value has to be initialized for each `parameter`, and the `size` and `buffer` values have to be initialized for each `In` and `InOut` struct. For `Out` parameters, `size` may have the value 0, in which case the `buffer` must be un-allocated. The `buffer` is to be created (e.g., allocated) by the SAGA implementation upon arrival of the result data, with a size that is sufficient to hold all result data. The `size` value is then set by the implementation to be equal to the allocated buffer size.

**HINT:**

This argument handling scheme allows efficient (copy-free) passing of parameters, because the parameter vector are passed by reference.

When an `Out` or `InOut` struct uses a pre-allocated buffer, any returned data exceeding the buffer size are discarded. The application is responsible for specifying correct buffer sizes for pre-allocated buffers; otherwise, the behaviour is in general undefined. For more details on buffer management, see Section ?? about buffer management.

## 14 Using the Advert Package

### 14.1 Quick Introduction

### 14.2 Reference

```
----- Prototypes: saga::advert -----
1 namespace saga
2 {
3     namespace advert
4     {
5         namespace metrics
6         {
7             char const * const advert_modified
8                 = "advert.Modified";
9             char const * const advert_deleted
10                = "advert.Deleted";
11
12             char const * const directory_created_entry
13                 = "directory.CreatedEntry";
14             char const * const directory_modified_entry
15                 = "directory.ModifiedEntry";
16             char const * const directory_deleted_entry
17                 = "directory.DeletedEntry";
18         };
19
20         enum flags
21         {
22             Unknown      = /* -1, */ saga::name_space::Unknown,
23             None         = /* 0, */ saga::name_space::None,
24             Overwrite    = /* 1, */ saga::name_space::Overwrite,
25             Recursive    = /* 2, */ saga::name_space::Recursive,
26             Dereference  = /* 4, */ saga::name_space::Dereference,
27             Create       = /* 8, */ saga::name_space::Create,
28             Exclusive    = /* 16, */ saga::name_space::Exclusive,
29             Lock         = /* 32, */ saga::name_space::Lock,
30             CreateParents = /* 64, */ saga::name_space::CreateParents,
31                         // 256,          reserved for Truncate
32                         // 512,          reserved for Append
33             Read         = 512,
34             Write        = 1024,
35             ReadWrite    = 1536,
36                         //2048,          reserved for Binary
37         };
38
39         class entry
40             : public saga::name_space::entry,
41               public saga::attributes
```

```
42     {
43     public:
44         entry (saga::session const & s,
45              saga::url          url,
46              int                 mode = Read);
47         entry (saga::url        url,
48              int                 mode = Read);
49         entry (saga::object const & other);
50         entry (void);
51         ~entry (void);
52
53         entry & operator= (saga::object const & object);
54
55         void          store_object  (saga::object object);
56         saga::object retrieve_object (void);
57         saga::object retrieve_object (saga::session const & s);
58
59         void          store_string  (std::string str);
60         std::string  retrieve_string (void);
61     };
62
63     class directory
64     : public saga::name_space::directory,
65       public saga::attributes
66     {
67     public:
68         directory (saga::session const & s,
69                  saga::url          url,
70                  int                 mode = ReadWrite);
71         directory (saga::url url,
72                  int                 mode = ReadWrite);
73         directory (saga::object const & other);
74         directory (void);
75         ~directory (void);
76
77         directory & operator= (saga::object const & o);
78
79         advert::entry  open      (saga::url url,
80                                  int        mode = Read);
81         advert::directory open_dir (saga::url url,
82                                     int        mode = Read);
83         std::vector <saga::url>
84             find      (std::string name,
85                       std::vector <std::string> pattern,
86                       int        flags = None)
87     };
88 }
89 }
```

90



### 14.3 Details

## **Part III**

# **Advanced Topics**

## 15 Using Asynchronous Operations

### 15.1 Quick Introduction

All functional SAGA API calls come in three flavours: synchronous, asynchronous, and task versions. The easiest way to understand the relations and differences between them is to consider that *all* functional API calls are represented by stateful **tasks** (e.g., a `file.copy()` API call is a task which can be in **Running** or **Done** state).

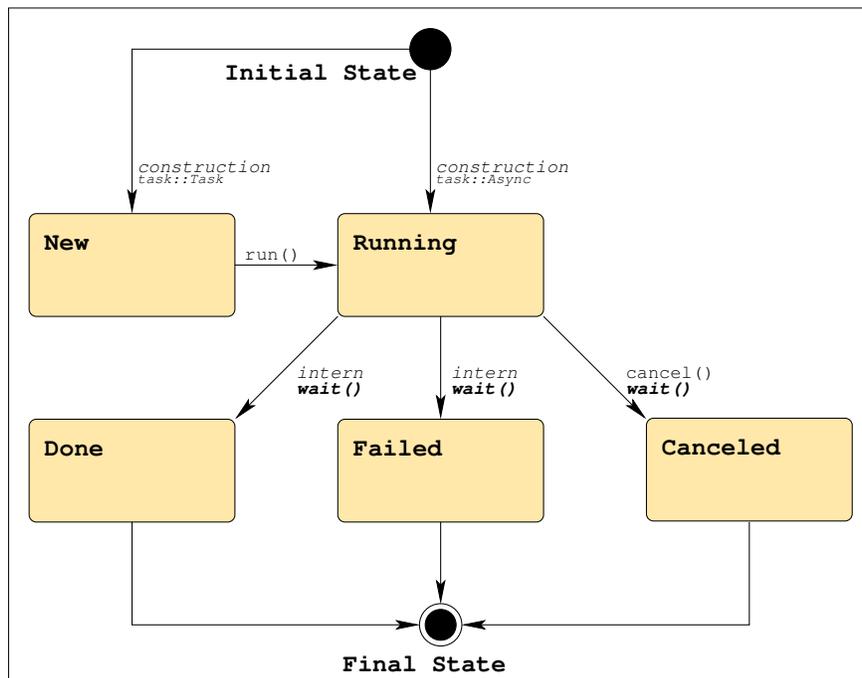


Figure 4: The SAGA task state model.

Figure ?? shows the available task states, and the various possible state transitions. The different SAGA API call invocations are simply representing ways to create the respective tasks in different states: Synchronous calls create tasks in a final state (**Done** or **Failed**); Asynchronous operations create tasks which are **Running**; and the Task version of the calls creates tasks which are **New**, and need to be `run()` to do anything at all:

```

    _____ Asynchronous operations _____
1 | int main (int argc, char** argv)
2 | {

```

```
3 // create a file instance
4 saga::url u (argv[1]);
5 saga::filesystem::file f (u);
6
7 // target dir to copy file to
8 saga::url tgt ("file://localhost/tmp");
9
10 // run file copy in three flavours
11 saga::task t_1 = f.copy <saga::task::Sync> (tgt);
12 saga::task t_2 = f.copy <saga::task::Async> (tgt);
13 saga::task t_3 = f.copy <saga::task::Task> (tgt);
14
15 // the three tasks do the same, but t_1 is already done at
16 // this point, t_2 could still be running, and t_3 did not yet
17 // start.
18
19 // let's get the async version done:
20 t_2.wait ();
21
22 // now, run and finish the Task version:
23 t_3.run ();
24 t_3.wait ();
25
26 // all three copies are done here.
27 }
```

For convenience, the synchronous version is provided simply as:

```
f.copy (tgt);
```

Now, what about the return values? For those, the tasks provide a typed member method, `get_result()`, which can only be called when a task is in a final state:

```
Return values for asynchronous operations
1 int main (int argc, char** argv)
2 {
3 // create a file instance
4 saga::url u (argv[1]);
5 saga::filesystem::file f (u);
6
7 // run method in three flavours
8 saga::task t_1 = f.get_size <saga::task::Sync> ();
9 saga::task t_2 = f.get_size <saga::task::Async> ();
10 saga::task t_3 = f.get_size <saga::task::Task> ();
11
12 // let's get the async version done:
```

```
13     t_2.wait ();
14
15     // now, run and finish the Task version:
16     t_3.run ();
17     t_3.wait ();
18
19     // all three calls are done here.
20
21     ssize_t size;
22     size = t_1.get_result <ssize_t> ();
23     size = t_2.get_result <ssize_t> ();
24     size = t_3.get_result <ssize_t> ();
25
26
27 }
```

The convenience of the synchronous version above is that it returns the return value immediately:

```
    ssize_t size = f.get_size ();
```

## 15.2 Reference

## 15.3 Details

## 16 Using Monitorables and Notifications

### 16.1 Quick Introduction

#### Monitorables and Metrics

**FIXME:** fill in

#### Notifications

Closely related to monitorables and metrics are notifications; they notify the application of certain events. For example, when a task is done, when it changes its state, etc. For those events, the application can create and register custom callbacks:

```

                                     Notifications on task state changes
1  #include <saga/saga.hpp>
2
3  class my_callback : public saga::callback
4  {
5      public:
6          // cb exits the program with appropriate error code
7          bool cb (saga::monitorable mt,
8                  saga::metric      m,
9                  saga::context      c)
10         {
11             if ( m.get_attribute ("state") == "Done" ) {
12                 exit (0);
13             }
14
15             if ( m.get_attribute ("state") == "Failed" ) {
16                 exit (-1);
17             }
18
19             // all other state changes are ignored
20             return true; // keep callback registered
21         }
22     }
23
24     int main (int argc, char** argv)
25     {
26         // run a file copy asynchronously
27         saga::url u (argv[1]);
28         saga::filesystem::file f (u);
29         saga::task t = f.copy <saga::task::Async> (saga::url (argv[2]));

```

```
30
31     // monitor the task state
32     my_callback cb;
33     t.add_callback ("state", cb);
34
35     // make sure the task finishes
36     t.wait ();
37 }
```

The example above registers a private callback to the "state" metric of the task. In fact, many SAGA objects (which implement the monitoring interface, and are thus monitorables) have several metrics, which can be queried by `monitorable.list_metrics()`.

**HINT:**

The return value of the callback determines if the callback stays registered (**true**) or not (**false**). This allows for a *call-once* semantics.

## 16.2 Reference

## 16.3 Details

## 17 Specifying Security Details

### 17.1 Quick Introduction

Tight security is arguably one of the most required features of Grid environments. On the other hand, for the average end user, it is also one of the most annoying features. That is mostly caused by the confusion about how security credentials are to be maintained, when and where they are valid, how to choose between them, etc.

SAGA can alleviate only some of these problems—more standardization work outside of SAGA is required to be able to simplify credential management further.

To understand SAGA's security model, one needs to consider two concepts: *sessions* and *contexts*; a `saga::context` simply represents one specific security credential; a `saga::session` can have multiple of these contexts attached, and is defined by the lifetime of the objects and operations using these credentials.

By default, both `saga::sessions` and `saga::contexts` are invisible: a default session is always implicitly created on the SAGA call, and picks up all the default security credentials your SAGA implementation knows about, such as your default ssh keys, your default globus proxy (if it was initialized before), your default Unicore keyring, etc. In most cases, that should allow writing applications which have no single line of code addressing security explicitly, which are still secure, as these credentials are used on adaptor level, as needed:

```
Secure gsiftp file copy
1  #include <saga/saga.hpp>
2
3  int main ()
4  {
5      saga_url src ("ssh://remote.host.net/tmp/one.dat");
6      saga_url tgt ("ssh://remote.host.net/tmp/two.dat");
7
8      // do a file copy, using the default globus X509 proxy, in the
9      // default session
10     saga::filesystem::file f (src);
11     f.copy (tgt);
12 }
```

On the other hand, it allows simple means of control on which credentials are to be used for a given operation, as in the following example:

Secure gsiftp file copy

```
1  #include <saga/saga.hpp>
2
3  int main ()
4  {
5      saga_url src ("ssh://remote.host.net/tmp/one.dat");
6      saga_url tgt ("ssh://remote.host.net/tmp/two.dat");
7
8      // create a new session, no default contexts are attached
9      saga::session my_session;
10
11     // create a new ssh context
12     saga::context my_context ("ssh");
13
14     // point ssh context to a specific ssh key
15     my_context.set_attribute ("UserCert",
16                             "/home/username/.ssh/id_dsa.special");
17
18     // set all other attributes to their defaults
19     my_context.set_defaults ();
20
21     // create a new file object, in that session
22     saga::filesystem::file f (my_session, src);
23
24     // do a file copy, using ONLY that ssh key
25     f.copy (tgt);
26 }
```

**HINT:**

In most cases, the default session is what you want to use; if that does not work, try to convince your system administrator to configure SAGA so that the default session works!

Objects created from other objects inherit their session. Asynchronous operations are living in the session of their spawning objects. A session does not die when going out of scope, but only when all associated objects and operations die.

## 17.2 Reference

## 17.3 Security Details

## 18 Miscellaneous Issues

### 18.1 Primitive Data Types

Alas, the C++ standard does not define all data types needed within SAGA. For that reason, we provide these types also in the `saga` namespace. Although these types should usually map to the native types on the respective platform the SAGA application is running on, we can only guarantee portability when using the SAGA provided types.

The types provided by SAGA are:

```
----- Primitive types in the saga namespace -----
1 namespace saga
2 {
3     typedef ... char_t;
4
5     typedef ... int8_t;
6     typedef ... uint8_t;
7
8     typedef ... int16_t;
9     typedef ... uint16_t;
10
11    typedef ... int32_t;
12    typedef ... uint32_t;
13
14    typedef ... int64_t;
15    typedef ... uint64_t;
16
17    typedef ... intmax_t;
18    typedef ... uintmax_t;
19
20    typedef ... long_long_t;
21
22    typedef ... size_t;
23    typedef ... ssize_t;
24    typedef ... off_t;
25
26 } // namespace saga
```

Note that the `int64_t` and `uint64_t` are only defined on those platforms which support 64 bit integers. How the individual types are defined is platform-dependent, and should not be relied upon. However, they will often be mapped to the respective Boost types.

## 18.2 Boost, and C++/TR1

Our implementation relies heavily on Boost [?]. Boost is an open source collection of state-of-the-art C++ libraries, is highly portable, and is actually thought to be the playground for upcoming versions of the C++ standard. In fact, all Boost features visible on the surface of our SAGA implementation (i.e., visible to the user of SAGA), are features which can already be found in the 'C++ Technical Recommendation 1' [?], a recent extension of the C++ standard. However, if your compiler does not support the TR1 extensions yet, Boost is required to use SAGA. Please refer to the installation manual for details on Boost.

**FIXME:** correct?