

Manual: Writing a portlet using the Rapid tool

Jos Koetsier, Srihathai Prammanee

March 5, 2009

Abstract

1 Introduction

In many scientific disciplines, the demand for computational power has increased dramatically and is likely to grow even further in the future. New techniques, such as micro array analysis in biology, generate ever larger data sets that need to be analysed. Sciences such as astronomy and physics are developing more accurate instruments, that, coupled with new and more complicated algorithms require ever larger computing resources. These new innovations and techniques mean that for many applications simple computation on a desktop computer can no longer be considered adequate.

Thankfully, the demand for more computing power is mirrored by the availability of newer and faster High Performance Computing resources. For example, the list of top 500 supercomputers in the world shows that available computing power has increased exponentially for the last fifteen years. New developments in Grid Computing enable users to access even more computing power by allowing users to submit computational tasks transparently to not one, but several, often heterogeneous, resources.

Unfortunately, high performance computing comes at the price of increased complexity. Submitting a computational task to a cluster or grid often requires issuing a set of complex commands through a command line interface. This is further complicated by the necessity of some form of security infrastructure to regulate access to often expensive resources. Users need to request certificates, download proxy certificates and understand the security implications of their actions. Accessing computing power is no longer a matter of double-clicking an application on a desktop, but more and more requires an in-depth knowledge of computer science, which not all scientists possess.

In order to alleviate this burden on the researcher, scientific applications are often wrapped in a user interface. In the grid community, the use of portals is often considered the preferred solution mainly due to its distributed nature. No complicated installation of software is necessary; users can log on to the portal from any computer, simply by using the browser supplied by their operating system. An added advantage is that security is dealt with by the portal. Most portals offer some sort of login mechanism and issue privileges using a role based access mechanism, allowing for fine-grained security to be applied.

When considering the state of the art of current portlets that allow users to submit tasks to clusters or grids, roughly two approaches can be identified. First there is the *generic* portlet. This type of portlet, in effect, wraps a command line job submission. For every parameter a text box or drop-down list is supplied which the user needs to fill out. These portlets allow a wide variety of computational tasks to be submitted, but at the cost of increased complexity. This type of portlet is far from ideal, especially because we set out to shield the user from all this complexity in the first place.

The other approach is to write a *custom* portal application for each computational task. We consider this type of portlet preferable to the generic portlet of the first approach, because of its simplicity. Any user can access the power of large compute clusters without the indepth knowledge required to perform a command-line job submission. However, this approach does come at the cost of increased development time spent on programming the portlets themselves. It is often the case that there is simply not enough time and money available to write a portlet for each new project. There is also the additional problem of maintainability. When the portlet is finished, it requires an expert to maintain it, fix bugs and perform upgrades.

The Rapid project is an attempt to address these issues. We propose an *automatic job submission portlet generation system*. This system allows an expert to specify information about a job submission portlet in an XML file. This can be information about file transfers, options to display to the user and the available compute clusters as well as the security information required to access them. The Rapid system takes this XML file and translates it into a fully working JSR 168 compliant job submission portlet.

This approach gives us the benefits of both portlet types we discussed earlier. On the one hand, users get a job submission portlet that shields them from any complexity and on the other hand the time and effort to write and maintain a generic portlet is greatly reduced. Generating a portlet in this way does reduce flexibility somewhat however; not each portlet that can be written by hand can be generated using Rapid. However, we think that using Rapid will be able to meet most users' needs and will allow scientists to make better use of the available computing power.

2 Overview

This section describes the high level concepts required to understand the mechanisms that are used in the Rapid System. The job submission model used by Rapid is loosely based on the Basic Execution Service (BES) and the Job Submission Description Language (JSDL) standards. In these standards a method of job submission is described, which (in its most basic form) consists of the following stages:

1. Copy one or more files from one or more file servers into the host where the computation takes place. (Stage in, or *source* file transfers)
2. Execute and monitor the job as it is executing. This job may consist of several sub-jobs, for instance in the case of so-called parameter-sweeps.
3. Copy one or more files from the execution host (presumably the results

of the computation) to one or more file servers. (Stage out, or *target* file transfers)

4. Visualise the results of the computation.

The Rapid system allows a *template* for such a job to be defined. The user can pre-define a set of source and target file transfers, using methods such as ftp, sftp, http or gsftp. Defining the job itself follows the POSIX model, consisting of a set of environment variables, an executable and a number of parameters.

Obviously, submitting a pre-defined job is not very useful and therefore the Rapid system allows a user to change the pre-defined job template, by introducing a set of variables. These variables are given a unique name and are initialised with a default value as part of the job template. In the user interface of the portlet, the user is allowed to change these variables. The portlet developer can specify which user interface element to use to change a variable. Rapid supports a wide variety of input methods such as file browsers, radio buttons, lists and text entry fields.

When a job is submitted the variables are substituted in place holders within the job template. As an example, suppose we have an executable called 'factorial', which takes one parameter, which is a positive whole number for which this program computes the factorial. We would specify this in the job template as:

```
...
<executable>
  <single><value>factorial</value></single>
</executable>
<parameter index="1">
  <single><value>5</value></single>
</parameter>
...
```

This example always sets the first parameter to the value '5'. If we want the user to be able to change this value, we add a variable which we call 'parameterValue'.

```
...
<variable name="parameterValue">
  <single><value>5</value></single>
</executable>
...
```

This initialises the variable with a default value of '5'. To bind this variable to the parameter, we have to change the job template to:

```
...
<executable>
  <single><value>factorial</value></single>
</executable>
<parameter index="1">
  <single><value>$(parameterValue)</value></single>
</parameter>
...
```

The $\$(\langle name \rangle)$ notation allows a value to be substituted by the variable, when the job is submitted. This notation can be used in most parts of the Rapid XML file.

In the user interface section of the Rapid XML file we can now introduce a user interface element, that enables the user to examine and change the value of a variable. For example a list of possible values can be added, from which the user makes a selection:

```
<page name="configure job">
...
<x:h1>Choose factorial</x:h1>
<variable name="parameterValue">
  <list>
    <item value="1">1</item>
    <item value="5">5</item>
    <item value="20">20</item>
  </list>
</executable>
...
</page>
```

In the following sections we will discuss all elements in the Rapid XML file.

3 The Rapid XML Document

The `<rapid>` element denotes the root of the Rapid XML document and contains two namespaces: the *rapid* namespace: <http://www.nesc.ac.uk/Rapid> and the *XHTML* namespace: <http://www.w3.org/1999/xhtml>. Elements from the Rapid namespace are used to set the job template, initialise file servers and configure user interface elements. The XHTML interface is used for markup, such as tables, fonts and incorporating images.

Element `<rapid>`

Root element of the Rapid XML document.

Child Elements

- `<condor>` defines a new Condor job submission instance. (optional)
- `<sge>` defines a new Sub Grid Engine job submission instance. (optional)
- `<pbs>` defines a new PBS job submission instance. (optional)
- `<fork>` defines a new Fork submission instance. (optional)
- `<local>` defines a local file system. (optional)
- `<ftp>` defines an (s)ftp file system. (optional)
- `<http>` defines an http file system. (optional)
- `<gsiftp>` defines the gsiftp file system. (optional)

- `<initialise>` initialises a new job template.
- `<page>` defines the user interface.
- `<persistence>` defines state preservation. (optional)

Example

```
<?xml version="1.0" encoding="UTF-8"?>
<rapid xmlns="http://www.nesc.ac.uk/Rapid"
       xmlns:x="http://www.w3.org/1999/xhtml">
  <!-- Contents go here -->
</rapid>
```

3.1 Job submission servers

A job submission server is the compute resource to which a compute job is submitted. In the *rapid* namespace, the `<gridsam>`, `<condor>`, `<sge>`, `<pbs>` and `<fork>` elements support GridSAM, Condor, Sun Grid Engine, PBS and simple forking a new process, respectively. They are the immediate child elements of the `<rapid>` root element.

3.1.1 Condor cluster

The `<condor>` element allows Rapid to submit to a Condor cluster.

Element `<condor>`

Attributes

- `name`: a unique identifier

Child Elements

- `<bin>` element specifies the location of Condor execution files 'condor_submit' and 'condor_q'. (optional)
- `<condorconfig>` path to the condor.config file. (optional)
- `<universe>` defines the runtime environment under which Condor cluster should execute a job. (optional)
- `<filesystemname>` indicates through which file system (specified in 3.2) the condor submit host is accessed. Supports SSH and Local.
- `<submitline>` defines an additional lines in a Condor submit file. (zero or more)

Example

```
<condor name="condor">
  <bin>/home/condor/condor/bin</bin>
  <universe>vanilla</universe>
  <submitline>priority = 10</submitline>
  <submitline>image_size = 20</submitline>
  <filesystemname>CondorFS</filesystemname>
  <condorconfig>/etc/condor/condor_config</condorconfig>
</condor>
```

3.1.2 PBS

The `<pbs>` element allows Rapid to submit to a PBS scheduler.

Element `<pbs>`

Attributes

- `name`: a unique identifier

Child Elements

- `<bin>` specifies the location of the PBS binary files 'qstat' and 'qsub'
- `<filesystemname>` indicates through which file system (specified in 3.2) the condor submit host is accessed. Supports SSH and Local.
- `<option>` defines an additional lines in the PBS submit file. (zero or more)

Example

```
<pbs name="pbs">
  <bin>/opt/pbs/bin</bin>
  <option>#PBS -l nodes=1:ppn=$(numberOfCPUS),cput=12:00:00</option>
  <filesystemname>PBSFS</filesystemname>
</pbs>
```

3.1.3 Sun Grid Engine

The `<sge>` element allows Rapid to submit to a Sun Grid Engine cluster.

Element `<sge>`

Attributes

- `name`: a unique identifier

Child Elements

- `<root>`: path to SGE_ROOT. (optional).
- `<filesystemname>`: indicates through which file system (specified in 3.2) the Sun Grid Engine submit host is accessed. Supports SSH and Local.
- `<option>`: lines to add to a SGE submit script. (zero or more)

Example

```
<sge name="sge">
  <root>/home/sge/sge</bin>
  <filesystemname>LocalFS</filesystemname>
  <option>#$ -M admin@nesc.ac.uk</option>
</sge>
```

3.1.4 Fork

The `<fork>` element configures a simple 'fork' submission resource where the job is simply run as a new process.

The `<fork>` Element

Attributes

- **name**: a unique identifier

Child Elements

- The `<filesystemname>` indicates through which file system (specified in 3.2) the process is to be run. Supports SSH and Local.

Example

```
<fork name="fork">
  <filesystemname>LocalFS</filesystemname>
</fork>
```

3.2 File systems

File systems are defined using four elements: `<http>`, `<ftp>`, `<local>` and `<gsiftp>`. These elements are the direct child nodes of `<rapid>` and require a unique 'name' attribute used as a reference throughout the document. All filesystem tags contain a `<url>` child element of which the body specifies the complete URL of the file system. Some elements may encompass additional child nodes, depending on authorisation and authentication. Those file system elements are detailed in the following sections.

3.2.1 HTTP

The `<http>` element is used for both *http* and *https* servers. Neither of them requires authentication.

Element `<http>`

Attributes

- **name**: a unique identifier

Child Elements

- `<url>` URL of the file system

Example

```
<http name="http">
  <url>http://www.nesc.ac.uk/index.html</url>
</http>
<http name="http2">
  <url>https://www.omii.ac.uk/index.html</url>
</http>
```

3.2.2 Local file system

The `<local>` element allows access to a local file system. Similar to `<http>`, there is no authentication. This file system refers to the Portal file system and not the users' file system. When using GridSAM, extra care has to be taken as during the job submission phase a 'local' filesystem refers to the filesystem GridSAM resides on.

Element `<local>`

Attributes

- `name`: a unique identifier

Child Elements

- `<url>` URL of the file system

Example

```
<local name="Local File System">
  <url>file:///home/portaluser</url>
</local>
```

3.2.3 FTP

The `<ftp>` element is used for both *ftp* and *sftp* servers and requires authentication using a username and password.

Element `<ftp>`

Attributes

- `name`: a unique identifier

Child Elements

- `<url>` URL of the file system. Paths are relative to the users' home directory, unless preceded with a `%2f` sequence.
- `<username>` username
- `<password>` password

Example

```
<ftp name="myftpserver">
  <url>sftp://host.university.ac.uk/full/path/</url>
  <username>myname</username>
  <password>mypassword</password>
</ftp>
```

3.2.4 GSIFTP

The `<gsiftp>` element describes a Grid FTP server. For authentication, a valid X509 proxy certificate must be obtained through a MyProxy server. Paths are relative to the users' home directory, unless preceded with a `%2f` sequence.

Element `<gsiftp>`

Attributes

- `name`: a unique identifier

Child Elements

- `<url>` URL of the file system
- `<myproxyhost>` hostname of the myproxy server
- `<myproxyport>` port through which to access the myproxy server
- `<myproxyusername>` username
- `<myproxypassword>` password

Example

```
<gsiftp name="mygisftpserver">
  <url>gsiftp://host.university.ac.uk/full/path/</url>
  <myproxyhost>myproxy.university.ac.uk</myproxhost>
  <myproxyport>7512</myproxyport>
  <myproxyusername>myusername</myproxyusername>
  <myproxypassword>secret</myproxypassword>
</gsiftp>
```

3.3 Initialising a job

The `<initialise>` element is used to define a template for a new job. It can contain four child elements. The `<datastage>` element involves the file staging details and the `<posix>` element is used to set parameters of a job according to the POSIX model. The compute resource to use is set in the `<submitto>` element and, finally, the `<variable>` element can be used to declare variables and set initial values.

3.3.1 Data types

For each value in the `<initialise>` section, a type has to be declared. There are currently four types, specified by the elements `<single>`, `<array>`, `<range>` and `<uuid>`. The `<single>` element is used to indicate simple, single values, the `<array>` element is used to specified multiple values and the `<range>` element is used to specify a range of numbers.. The `<uuid>` element generates a single new Universally Unique IDentifyfier (UUID) for each new job that is generated. The values themselves are specified in the body of `<value>` elements which can appear as direct children of `<single>`, `<array>` or `<range>`.

If multiple values in an array are used, multiple jobs will be created at the job submission stage. If multiple array types are used, the maximum number of values is used to determine the number of job to submit, where the first job will use the first value from each array, the second job will use the second value and so on.

Elements `<single>`, `<array>`

Child Elements

- `<value>` specifies a value (one or more if parent is `<array>`, one for `<single>`).
- `<min>` sets the minimum possible value. Implies `<value>` is a number.
- `<max>` sets the maximum possible value. Implies `<value>` is a number.
- `<regexp>` sets a regular expression the `<value>` elements must conform to.
- `<errormessage>` error message to display if `<value>` elements contain invalid values.

Example

```
<array>
  <value>parameter1</value>
  <value>parameter2</value>
  <value>parameter3</value>
  <regexp>[0-9]+</regexp>
  <errormessage>Please enter a whole, positive number</errormessage>
</array>
```

Element `<range>`

Child Elements

- `<min>` Sets the minimum of the range.
- `<max>` Sets the maximum of the range.
- `<step>` Sets the stepsize to use

Example

```
<range>
  <min>-1.0</min>
  <max>0.6</max>
  <step>0.3</step>
</array>
```

This example generates the range $-1.0, -0.7, -0.4, -0.1, 0.2, 0.5$

Element `<uuid>`

Automatically creates a new Universally Unique Identifier (UUID) when a new job is created.

Example

```
<uuid/>
```

3.3.2 Data staging

Data staging in the Rapid system follows the staging model as defined in JSDL. One *data stage instance* can contain a *source* and *target* transfer. The *source* file transfer copies a file or directory into the job execution engine and is performed before running a job. Once the job execution has finished, the *target* file transfer moves a file or directory out of the job execution engine.

Rapid defines three additional parameters for a data stage instance. First, a filename uniquely identifies the file called in the job execution engine. Second, a creation flag determines how the file will be created in the job execution engine and finally, a deleteontermination flag indicates whether the file will be removed from the execution host once the job is completed.

Element `<datastage>`

The `<datastage>` element simply groups any `datastage` elements and has no attributes.

Child Elements

- `<filesystem>` determines the file system used in either the source or target file staging.
- `<path>` determines the full path of either 'source' or 'target' file stage.
- `<filename>` refers to which file or directory will be staged into the job execution file system.
- `<deleteontermination>` determines whether the file will be deleted from the job execution engine once the job has finished. (optional)
- `<creationflag>` refers to the creationflag of the file. Its value can be either APPEND, OVERWRITE, or DONTOVERWRITE. (optional)

Example

```
<datastage>
  <filesystem type="source" datastage="executable_in">
    <single><value> myftpserver </value></single>
  </filesystem>
  <path type="source" datastage="executable_in">
    <single><value>xxx/yyy</value></single>
  </path>
  <filename datastage="executable_in">
    <single><value>.</value></single>
  </filename>
  <deleteontermination datastage="executable_in">
    <single><value>>true</value></single>
  </deleteontermination>
  <creationflag datastage="executable_in">
    <single><value>OVERWRITE</value></single>
  </creationflag>
</datastage>
```

Element <filesystem>

Attributes

- **datastage**: unique identifier, groups datastage elements together
- **type**: either 'source' for a source transfer or 'target' for a target transfer

Element <path>

Attributes

- **datastage**: unique identifier, groups datastage elements together
- **type**: either 'source' for a source transfer or 'target' for a target transfer

Element <filename>

Attributes

- **datastage**: unique identifier, groups datastage elements together

Element <deleteontermination>

Attributes

- **datastage**: unique identifier, groups datastage elements together

Child Elements

- <single> containing either 'true' or 'false'
- <array> containing either 'true' or 'false'

3.3.3 Posix

Job execution in Rapid follows the Unix POSIX model and all its parameters are grouped under the `<posix>` element.

Element `<posix>`

Child Elements

- `<executable>` denotes the name of the executable to run.
- `<stdin>`, `<stdout>`, and `<stderr>` elements refer to the standard input, output and error, respectively.
- `<workingdir>` element specifies the directory running the executable.
- `<parameter>`
- `<environmentname>`
- `<environmentvalue>`

Example

```
<initialise>      :
  <posix>
    <executable>
      <single><value>./java</value></single>
    </executable>
    <environmentname index="0">
      <single><value>CLASSPATH</value></single>
    </environmentname>
    <environmentvalue index="0">
      <single><value>.:aaa:bbb/ccc/ddd.jar</value></single>
    </environmentvalue>
    <workingdir>
      <single><value>.</value></single>
    </workingdir>
    <parameter index="0">
      <single><value>para0</value></single>
    </parameter>
    <parameter index="1">
      <array>
        <value>para1</value>
        <value>para2</value>
      </array>
    </parameter>
  </posix>
  :
</initialise>
```

Element `<parameter>`

Parameter of the executable.

Attributes

- **index**: index of the parameter, starting at 0.

Example

```
<initialise>
  <parameter index="0">
    <single><value>-al</value></single>
  </parameter>
</initialise>
```

Element <environmentname>

Shell environment value NAME=VALUE. This element refers to the NAME.

Attributes

- **index**: index of the environment value, starting at 0.

Example

```
<initialise>
  <environmentname index="1">
    <single><value>CLASSPATH</value></single>
  </environmentname>
</initialise>
```

Element <environmentvalue>

Shell environment value NAME=VALUE. This element refers to the VALUE

Attributes

- **index**: index of the environment value, starting at 0.

Example

```
<initialise>
  <environmentvalue index="1">
    <single><value>/opt/libs/lib.jar</value></single>
  </environmentname>
</initialise>
```

3.3.4 Variables

Rapid allows variables to be declared under the <variable> element, where the 'name' attribute uniquely identifies the variable name. The value of a variable can be referred to by using the format: \$(VARIABLE) .

Element <variable>

Variable to set.

Attributes

- **name**: Name of the variable.

Example

```
<initialise>
  <variable name="myVar">
    <uuid/>
  </variable>
</initialise>

<page>
  The value of myVar is ${myVar}
</page>
```

3.3.5 Execution node

By default, a job is submitted to the execution node specified under the `<submitto>` element. The value of this element refers to a submission engine, previously defined in section 3.1.

Element `<submitto>`

Example

```
<initialise>
  <submitto>
    <single><value>Manchester</value></single>
  </submitto>      :
</initialise>
```

3.4 Pages

The user interface (UI) is defined in a set of `<page>` elements, that are immediate children of the `<rapid>` root element. When a Rapid portlet starts up, the view associated with the first `<page>` element that is defined within the document is loaded up and displayed. The user can navigate between views by pressing a button, containing a 'navigate' action, as described in section 3.4.2.

Element `<page>`

A page containing markup and rapid user interface elements.

Attributes

- **name**: unique name of the page.

Child Elements

- `<variable>`
- `<button>`

- `<fileupload>`
- `<joblist>`
- `<subjoblist>`
- `<setjob>`
- elements from the XHTML namespace

Example

```
<page name="page1">
  <!-- page layout is defined here -->
</page>
```

Child elements of the `<page>` can be taken from both the *xhtml* and the *rapid* namespace. The markup of the user interface is written using standard XHTML. The logic of the portlet as well as user input is handled by elements from the *rapid* namespace.

3.4.1 User interface elements

The rapid system enables a user to change or display a variable by specifying a `<variable>` element as child element of the `<page>` element. How this is done, is defined in the child elements of then `<variable>` element, which are currently `<list>`, `<checkbox>`, `<radio>`, `<browser>`, `<text>` and `<output>`.

Element `<variable>`

Enables a user to change a variable or display the value of this variable.

Attributes

- **name**: Name of the variable to change or display.

Child Elements

- `<list>`
- `<checkbox>`
- `<radio>`
- `<browser>`
- `<text>`
- `<output>`

Example

```
<variable name="variableone">
  <text/>
</variable>
```

Element `<list>` and `<radio>`

Shows a list of items, a dropdown list or a set of radio buttons.

Attributes

- **refresh**: if 'true', the page will be reloaded if a selection is made.
- **size**: `<list>` only. Length of the list. If the length is 1, this user interface element will be a dropdown list.
- **index**: `<radio>` only. An index from 0. If the value edited is an 'array' value, it refers to the index-th element in the array. If it is a 'range' value, then 0 refers to the minimum value, 1 refers to the maximum value and 2 refers to the step size.

Child Elements

- `<item>` (one or more)

Example 1

```
<list refresh="true" size="1">
  <item value="itemone">select item one</item>
  <item value="itemtwo">select item two</item>
  <item value="itemthree">select item three</item>
</list>
```

Example 2

```
<radio refresh="true" index="2">
  <item value="itemone"/>1.0
  <item value="itemtwo"/>2.0
  <item value="itemthree"/>3.0
</list>
```

Example 2 allows a user to set the step size of a range value.

Element `<checkbox>`

Adds one or more checkboxes. Can only be used if the value set in the `<initialise>` section is of type `<array>`

Attributes

- **refresh**: if 'true', the page will be reloaded if a selection is made.

Child Elements

- `<item>` (one or more)

Example

```
<checkbox refresh="true" size="1">
  <item value="itemone"/>select item <x:h1>one</x:h1>
  <item value="itemtwo"/>select item <x:h1>two</x:h1>
  <item value="itemthree"/>select item <x:h1>three</x:h1>
</checkbox>
```

Element <item>

Adds a choice to a <list>, <radio>, or <checkbox>

Attributes

- **value:** The value to set into the job.

Child Elements

- <list> can contain text

Example

```
<checkbox refresh="true" size="1">
  <item value="itemone"/>select item <x:h1>one</x:h1>
  <item value="itemtwo"/>select item <x:h1>two</x:h1>
  <item value="itemthree"/>select item <x:h1>three</x:h1>
</checkbox>
```

Element <text>

Adds a text box.

Attributes

- **size:** Size of the text box.
- **index.** An index from 0. If the value edited is an 'array' value, it refers to the index-th element in the array. If it is a 'range' value, then 0 refers to the minimum value, 1 refers to the maximum value and 2 refers to the step size.
- **password:** If the variable refers to a password in a filesystem, or is otherwise a value that should not be echoed to the screen, set this attribute to 'true'

Example 1

```
<text size="10" password="true"/>
```

Example 2

```
Specify Minimum Value: <text size="10" index="0"/> <br\>
Specify Maximum Value: <text size="10" index="1"/> <br\>
Specify Step Size: <text size="10" index="2"/>
```

Element `<browser>`

Displays a file browser, allowing the user to navigate and select a file or directory. The value of the variable will be the path of the file selected.

Attributes

- `filesystem` the name of the filesystem this browser uses.. Can refer to a variable.
- `size`: size of the file browser.

Example

```
<variable name="path">
  <browser filesystem="$(filesystem" size="15"/>
</variable>
```

Element `<output>`

Displays the value of a variable. Can replace values by strings, which can be useful if the values are stored in the job as codes.

Attributes

- `default`: String to display if the job value has not been set

Child Elements

- `<replace>` Replace the value by another string. Specify the value in a 'search' attribute and the replacement in the body.
- `<pre>` Used for `<array>` type. Displays a string before each value in the `<array>`. Can contain XHTML elements.
- `<post>` Used for `<array>` type. Displays a string after each value in the `<array>`. Can contain XHTML elements.

Example

```
<variable name="parameterValue">
  <output>
    <replace search="item1">this is item one</replace>
    <replace search="item2">this is item two</replace>
    <replace search="item3">this is item three</replace>
    <pre>before</pre>
    <post>after</post>
  </output>
</variable>
```

Element `<fileupload>`

Allows a user to upload a file to the portal's file system. The file upload takes place immediately when updating the page. This file can be staged further to an execution host by referring to it in a `<datastage>` section. Contains the destination filename in the filesystem of the portal as the body.

Attributes

- `size`: size of the file upload input element (optional)
- `name`: unique name identifying this file upload

Example

```
<fileupload name="myupload" size="10">  
  /tmp/my-file-$(uuid).jpg  
</fileupload>
```

3.4.2 Buttons

Rapid allows the portlet designer to add buttons that perform one or more actions when pressed. Possible actions include navigating between pages, submitting jobs, loading and saving files, deleting jobs from a job monitor and changing the value of a variable.

Element `<button>`

This element defines a button. Its child elements are a list of actions that are performed in order.

Attributes

- `size`: size of the button.
- `display`: string to display on the button.

Child Elements

- `<navigate>`
- `<submit>`
- `<loadfile>`
- `<savefile>`
- `<deletejob>`
- `<navigate>`
- `<setvariable>`

Example

```
<button display="navigate and submit">
  <navigate nextpage="secondpage"/>
  <submit/>
</button>
```

Element <navigate>

Navigates to another page

Attributes

- **nextpage**: next page to load.

Element <submit>

Submits the job. Contains no attributes or child elements.

Elements <loadfile> and <savefile>

The element <loadfile> loads the contents of a file into the value of a variable. Similarly, the element <savefile> enables a user to save the value of a variable to a file.

Attributes

- **filesystem**: The name of the filesystem.
- **path**: Full path of the file.
- **variable**: Name of the variable.

Example

```
...
<variable name="contents">
  <single><value/></single>
</variable>
<variable name="path">
  <single><value>/path/to/my/file.txt</single>
</variable>
...
<page name="pagename">
...
  <variable name="contents">
    ,<text cols="40" rows="20"/>
  </variable>
...
<button display="Load File">
  <loadfile filesystem="myfilesystem" variable="contents" path="$(path)"/>
</button>
```

```
<button display="Save File">
  <savefile filesystem="myfilesystem" variable="contents" path="$(path)" />
</page>
```

Element `<deletejob>`

Allows a user to delete a job from a job monitor. The 'selection' attribute refers to a selection set in a `<selection>` element. See section 3.4.3 for more information.

Attributes

- **selection:** selection indicating the job to be deleted.

Element `<copyvariable>`

Sets the value of a variable to that of another. The type the variable is not changed. .

Attributes

- **from:** name of the variable from which the value is copied.
- **to:** name of the variable to which the value is copied.

Example

```
<button display="submit">
  <navigate nextpage="next page" />
  <submit />
  <copyvariable from="variableone" to="variabletwo" />
</page>
```

3.4.3 Job monitor

Job monitoring can be implemented using the `<joblist>`, `<subjoblist>` and the `<setjob>` elements. The `<joblist>` element iterates through all submitted jobs and allows the user to select a job by clicking on a radiobutton. If a job contains a set of subjobs, the element `<subjoblist>` can be used to iterate through any subjobs selected in a `<joblist>` element. The `<setjob>` element refers to the selection made in the `<joblist>` or `<subjoblist>` element and allows a more detailed view of a particular job to be presented. Adding a `<deletejob>` element allows a user to delete a job from the job monitor.

Elements `<joblist>`, `<subjoblist>` and `<setjob>`

`<joblist>` iterates through all previously submitted jobs. If a job consists of subjobs, `<subjoblist>` can be used to iterate through subjobs. `<setjob>` sets the job selected a `<selection>` element, which is a child element of `<joblist>` or of `<subjoblist>`.

The elements `<datastage>`, `<posix>`, `<submitto>` and `<variable>` elements have all been defined previously. However, if they appear as child elements of `<joblist>` and `<setjob>`, they can only be used to display values of

jobs that have been submitted and therefore do not contain input or output child elements.

Attributes

- **selection:** <setjob> and <subjoblist> only. When used with <setjob> sets the (sub)job that was selected. When used with <subjoblist>, this attribute refers to a selection of a job in <joblist>.

Child Elements

- <submitto>
- <variable>
- <jobid>
- <selection> Used only in conjunction with <joblist> and <subjoblist>.
- <date>
- elements from the XHTML namespace

Element <jobid>

When a job has been submitted it is given a Job ID in the form of a UUID. This Job ID can be displayed using this tag.

Element <status>

The status of a submitted job.

Element <selection>

This tag displays a radiobutton that can be used to select a job.

Attributes

- **selection:** unique name identifying the selection

Element <date>

The date the job has been submitted.

Example

```
<joblist>
  <x:table border="1" width="800"> <x:tr> <x:td>
    <selection name="mySelection"/>
  </x:td> <x:td>
    <date/>
  </x:td> <x:td>
    <jobid/>
```

```

        </x:td> <x:td>
            <status/>
        </x:td> </x:tr>
    </x:table>
</joblist>

<button display="delete job">
    <deletejob selection="mySelection" />
</button>

<subjoblist selection="mySelection">
    <x:table border="1" width="800"> <x:tr> <x:td>
        <selection name="mySubSelection"/>
    </x:td> <x:td>
        <date/>
    </x:td> <x:td>
        <jobid/>
    </x:td> <x:td>
        <status/>
    </x:td> </x:tr>
    </x:table>
</subjoblist>

<setjob selection="mySubSelection">
    <x:h1>executable:
        <posix>
            <executable/>
        </posix>
    </x:h1>
</setjob>

```

3.5 Persistence

In the default persistence model the Rapid system preserves the state of the portlet as part of the current login session. As a consequence, all state information, such as which jobs were submitted, which user they belong to and what the status of each job is, is lost each time a user logs out or if the session expires.

If the state of the portlet is required to be maintained, a database or file must be specified under the `<persistence>` tag, in which the portlet can store this information.

When using persistence it is important that the database exists and is empty. When upgrading Rapid to a new version, ensure that the database is cleared before running the portlet.

Element `<persistence>`

Child Elements

- `<username>` Username of the database to connect to.. The username is specified in the `name` attribute.

- `<password>` Password to use to connect to the database. The password is specified in the `name` attribute.
- `<host>` Host name of the database. The host name is specified in the `name` attribute.
- `<dbms>` Type of Database Management System. Currently supported are `mysql`, `postgres`, `hsqldb`, `h2` and `file`. The type of database is specified in the `name` attribute.
- `<database>` Name of the database to use. If the `<dbms>` tag 'file' was chosen, this will refer to the full path of the file to use. The database is specified in the `name` attribute.

Example 1

```
<persistence>
  <username name="rapiduser"/>
  <password name="password"/>
  <host name="database.nesc.ac.uk"/>
  <dbms name="mysql"/>
  database name="rapidDB"/>
</persistence>
```

Example 2

```
<persistence>
  <username name="admin"/>
  <password name="sa"/>
  <host name=""/>
  <dbms name="file"/>
  database name="/home/portal/rapidDB"/>
</persistence>
```

4 Generating the portlet

In this section we describe the steps necessary to generate and install the portlet.

4.1 Infrastructure

This section reviews the necessary infrastructure for the Rapid system.

4.1.1 GridSAM

If job submissions through GridSAM are required, we need to set the relevant security options to enable authentication between the portal and the GridSAM instance. This is simply done by copying the file 'crypto.properties' as defined within the OMII GridSAM container into the 'configuration' directory of GridSAM.

4.1.2 Portal

Because RAPID generates portlets, we require a JSR 168 compliant portlet container to deploy our portlet into. Our build scripts are currently written specifically for the GridSphere portal (version 3.0 and higher), but as we use the WAR format for the portlets it is relatively straightforward to deploy them in other portlet containers.

4.1.3 Apache Maven and Ant

Before installing a Rapid portlet, Apache Maven (mvn) and Apache Ant are required. Both packages can be downloaded from the Apache website¹. Rapid assumes that the 'ant' and 'mvn' commands can be accessed through the current search path.

4.2 Rapid portlet generation and installation

In order to install a new rapid portlet, download and unpack the RapidPortlet tarball. Within the RapidPortlet directory, all the configuration options are set in the 'configuration' directory. The file called 'crypto.properties' is the same as used by the OMII client and is only necessary if job submissions to GridSAM are used.

Another file in the configuration directory, called rapid.properties, sets properties such as the name and title of the portlet that is to be generated. Finally, the configuration directory should contain a file called 'rapid.xml', which specifies the content of the new portlet,

The portlet is generated by issuing the `mvn package` command from the RapidPortlet directory. The first time this command is executed, a number of jars will be downloaded into the `~/m2` directory, which can take a long time. If the command is issued again, the downloaded jars will be used and the process will be much quicker. When this process is completed, a 'portlet' directory containing a small ant install script and a 'war' file is generated.

The final portlet will be packaged as a 'war' file into the 'portlet' directory together with a small ant install script. The install script works for the GridSphere portal (version 3.0 and higher) only and assumes the `CATALINA_HOME` environment variable has been set. The portlet is installed by entering the `portlet` directory and issuing the `ant deploy` command.

An in-depth example of Rapid portlet installation can be found in Tutorial: Rapid-based Portlet Installation.

¹<http://www.apache.org/>