



DYNAMIC RUNTIME ENVIRONMENTS WITH JANITOR

This Janitor and this document with it is under continuous development. Your comments and suggestions are appreciated.

Michael Glodek, Daniel Bayer, Steffen Möller*

*moeller@inb.uni-luebeck.de

Contents

1	Introduction	5
1.1	Motivation	5
1.2	Overview	5
2	Installation	7
2.1	Configuration	7
2.2	Limitations	10
3	Usage	11
3.1	Janitor with A-REX	11
3.2	Janitor without A-REX	12
3.3	Janitor with A-REX	13
4	Maintenance	15
4.1	Catalog	15
4.1.1	Debian packages - dysfunctional in current implementation	16
4.2	HTML interface of the catalog	18
4.3	Introducing new packages	18
4.3.1	Debian Etch (tar based)	18
4.3.2	Automated transformation of install directory to dRTE	19
4.3.3	Protoypes	19
4.3.4	Example: ATLAS for High Energy Physics	20
5	Technical Motivation and Concepts	21
5.1	Implementation	21
5.1.1	Choice of Perl	21
5.1.2	Modular structure	21
5.2	Concepts	23
5.2.1	States of Runtime Environments	23
5.2.2	Subset of that functionality as implemented today	24
5.2.3	Dependencies on Job IDs	24
5.3	Job states	25
5.4	Integration with AREX	25
5.5	WebService Interface	25
5.5.1	What happens during installation	27
5.5.2	Security Consideration	27
6	Outlook	29
6.1	Representation of dynamic RTEs in the information model	29
6.2	Integration with Workflow Management	29
6.3	Implementation of a Catalog service	29
6.4	Integration with the Virtualization work	29
6.5	Use of RDF	29
6.6	Manual Verification	30
7	Appendix	31
7.1	Useful tutorials and documentations	31

1 Introduction

The *Janitor* is a service for the automated installation of runtime environments for grid computing elements. Its command line interface allows for a direct interaction with site administrators. However, the main stimulus for its development was the idea integrate such a service with the regular handling of compute jobs. For ARC this is performed by the A-REX module.

From the programmer's view, the Janitor is mostly a Perl script and the routines within A-REX to invoke it. The site administrator will also associate with it also the Catalog files, that describe the availability of runtime environments, and the repository of installable runtime environments themselves, which are regular tar archives obeying to particular structure and reside in a separate folder. In order to minimise the latency for the invocation of the Perl script and the associated parsing of files, a Janitor web service was developed, which still is a Perl script.

1.1 Motivation

A major motivation for grid projects is to stimulate new communities to adopt the technology to start sharing their resources. From the current grid user's viewpoint, the admission of users with a very different education will suddenly impose difficulties in the communication between site maintainers. One will not even understand the respective other side's research aims. Hence, the proper installation of non-standard software (Runtime Environments, RTEs) is not guaranteed. And extra time for manual labour plus self-education is scarce.

A core problem remains to distribute a locally working solution, the Know-How, quickly across all contributing sites, i.e. without manual interference. Every scientific discipline has its respective own set of technologies for the distribution of work load. For instance, research in bioinformatics requires access to so many different tools and databases, that few sites, if any, install them all. Instead, the use of web services became a commodity, with all their intrinsic problems as there are bottlenecks and restrictions of repeated access. The EU project KnowARC*, amongst other challenges, with the here presented work extends the NorduGrid's Advanced Research Connector (ARC) grid middleware [?] towards an infrastructure for the automated installation of software packages.

An automation of the software installation, referred to as dynamic Runtime Environments (dRTEs), seems the only approach to use the computational grid to its full potential. Components of workflows shall be spawned as jobs in a computational grid using dRTEs rather than accessing a public web service at one particular machine that is shared amongst all users. The grid introduces an extra level of parallelism that web services cannot provide. The demands for short response times and the heterogeneous education of site-administrators on a grid demand an automatism for the installation of software and databases without manual interference [?].

1.2 Overview

This document starts with a chapter on how to set-up the Janitor locally. It is followed by a chapter that gives further instructions on how to use the Janitor with A-REX and/or without A-REX. Afterwards, in the third chapter, the maintenance of the program will be presented, which is basically covering the method how to prepare new dRTEs. Deeper insights on the design of the Janitor will be given by the forth chapter. The document ends with an outlook to anticipated future developments and opportunities.

*<http://www.knowarc.eu>

Abbreviations

RTE –Runtime Environment

dRTE–dynamic Runtime Environment

RDF –Resource Description Framework (supporting the RTE Catalog)

2 Installation

The Janitor requires the two perl packages listed in table 2.1. To have the WebService interface for the Janitor, the packages listed in table 2.2 need to be installed before the build process is. The Perl modules are available on CPAN and ship with all major Linux distributions.

Table 2.1: Required perl packages for the Janitor. Log4perl is used for the internal logging of the Janitor, while the Redland RDF library is used for accessing the knowledge base (catalog) of Runtime Environments.

liblog-log4perl-perl	<i>Log4perl is a port of the log4j logging package</i>
librdf-perl	<i>Perl language bindings for the Redland RDF library</i>

Table 2.2: Optional libraries for the Janitor. The library libperl-dev provides the required header files to link the WebService to the Perl interpreter.

libperl-dev	<i>Perl library: development files</i>
-------------	--

If you are using regular Debian or Ubuntu packages, then the Janitor can be installed as root by "apt-get install nordugrid-arc1-janitor". Installing it will not drag other components of ARC with it, since the Janitor can be used in its own right – or in conjunction with another grid system, possibly. Packages for Redhat/Fedora and SuSE/OpenSuSE are also provided, the redland library however may not yet be available for those systems.

The Janitor source code is shipped as a part of the regular ARC-NOX source tree. If you are compiling the sources yourself, the default is to have the A-REX grid manager technically prepared to interact with the Janitor. The interaction can be prohibited with the *configure* flags *-disable-janitor-service* for the complete janitor or *-disable-janitor-webservice* for only the Web Service support.

Furthermore many users will want to consider installing the ontology editor Protégè^{*} to easily maintain the knowledge database of installable packages. At the time of writing, no Linux distribution is offering packages for this fine tool. However, the basic editing can also be performed fairly easily without that tool, and everyone is working on simplifying that process.

2.1 Configuration

The current version of the Janitor can be configured using the common file *arc.conf*. It is expected in the configuration directory *etc*. The Janitor is using the environment variable *NORDUGRID_CONFIG* to determine the location of the corresponding file. If that variable is not set, the default location */etc/arc.conf* will be used. The parameter *use_janitor* in the [grid-manager]section has to be set in order to tell A-REX whether to use Janitor or not. By default, Janitor is not used. Use the value "1" to enable Janitor.

Janitor is configured through parameters in the section [janitor]. Table 2.3 describes the available tags for the Janitor's configuration.

^{*}<http://protege.stanford.edu>

Table 2.3: Tags usable in *arc.conf* within the section **janitor**. Tags usable in *arc.conf* within the section **janitor**.

tag	example	description
uid	"root"	The effective uid.
gid	"0"	The effective gid.
registrationdir	" /var/spool/nordugrid/janitor"	Directory where we the current states of jobs are kept.
catalog	" /var/spool/nordugrid/janitor/catalog/knowarc.rdf"	URL of the catalog containing the package information.
downloadaddr	" /var/spool/nordugrid/janitor/download"	Directory for downloads
installationdir	" /var/spool/nordugrid/janitor/runtime"	Directory for installation of packages
jobexpirytime	"7200"	If a job is older than this, it is considered dead and assigned to be removal pending.
rteexpirytime	"36"	If a runtime environment was not used for this time, it will be assigned to be removal pending.
allow_base	"*"	Allow rule for base packages.
deny_base	"debian::etch"	Deny rule for base packages.
allow_rte	"*"	Allow rule for base packages.
deny_rte	" APPS/MATH/ELMER-5.0.2"	Deny rule for base packages.
logconf	" /opt/nordugrid/etc/log.conf"	Location of the logging configuration file for janitor.

The `uid` and the `gid` are defining which effective user id (`uid`) and group id (`gid`) shall be used for the Janitor.

The `registrationdir` describes the directory in which the subdirectories `jobs` and `rtes` will be created. In these directories the states of the jobs and the runtime environments are stored. Please recall that the Janitor does not use a database as a backend, but all communication between invocations are performed via files in those folders.

The knowledge base of installable packages is specified by the parameter `catalog`. Its value can be any kind of URL pointing to a file written in the Resource Description Framework (RDF) format. One should not light-heartedly use a remote address for this purpose. Such a remote source needs to be trusted, since any runtime environment specified in a catalog (if the package description matches constraints by the local site administrator) may possibly be installed by regular grid users.

The specification of the RDF file will be explained in detail in section 4.1. The parameter `downloadaddir` assigns the directory to which the installation files will be saved after they have been downloaded or copied from the repository which was specified by the catalog. Please remember: the URL in `arc.conf` indicates the location of the catalog. And the URLs somehow specified in the catalog specify the location from where to download the runtime environment.

The `installationdir` finally specifies the directory into which all packages will be installed. This directory needs to be available for all computing nodes for the execution of arbitrary programs, most commonly by using it as a shared NFS volume.

If the configuration file furthermore contains the `runtime` tag within the section `grid-manager`, the Janitor will also create a symbolic link in the `runtime` pointing to the configuration script of the installation performed by the Janitor. The tags `jobexpirytime` and `rteexpirytime` are used for an automated cleanup and is defined in seconds. The default value for the `jobexpirytime` is seven days and for the `rteexpirytime` three days. The additional tags `allow_base` `deny_base` `allow_rte` and `deny_rte` are used to include or exclude certain base packages or runtime environments of the catalog. This feature is useful, if the catalog is maintained by a higher organization. But again: you need to trust it.

The path to the `log4perl` configuration file is defined by the tag `logconf`. Examples on how to configure ARC and `log4perl` are provided in the Listings 2.1 and 2.2.

Listing 2.1: Example *arc.conf* settings for janitor.

```

1  [janitor]
2  enabled="1"
3  logconf="/opt/nordugrid/etc/log.conf"
4  registrationdir="/var/spool/nordugrid/janitor"
5  installationdir="/var/spool/nordugrid/janitor/runtime"
6  downloadaddir="/var/spool/nordugrid/janitor/download"
7  jobexpirytime="7200"
8  rteexpirytime="36"
9  uid="root"
10 gid="0"
11 allow_base="*"
12 allow_rte="*"
13
14 [janitor/nordugrid]
15 catalog="/var/spool/nordugrid/janitor/catalog/knowarc.rdf"
```

It should be noted that the `downloadaddir` or the `installationdir` specified in `arc.conf` could be any directory. Those will not be prepared by the package for the Linux distribution but need be created by the administrator manually after the Janitor has been installed. This also holds for the catalog.

When working with several catalogs, then the multiple catalog lines can be placed into the same `arc.conf` file. But every must go into its own block as separated with

janitor/someName

directives.

Listing 2.2: Example *log.conf* settings for janitor.

```

1  # Master Loglevel
2  # [OFF | DEBUG | INFO | WARN | ERROR | FATAL]
3  #log4perl.threshold = OFF
```

```
4
5 log4perl.rootLogger = WARN, DebugLog, MainLog, ErrorLog
6 log4perl.appender.DebugLog = Log::Log4perl::Appender::Screen
7 log4perl.appender.DebugLog.layout = PatternLayout
8 log4perl.appender.DebugLog.layout.ConversionPattern = [%C] %d %p> %m%n
9
10 log4perl.appender.MainLog = Log::Log4perl::Appender::File
11 log4perl.appender.MainLog.Threshold = DEBUG
12 log4perl.appender.MainLog.filename = /var/log/arc/janitor.log
13 log4perl.appender.MainLog.layout = PatternLayout
14 log4perl.appender.MainLog.layout.ConversionPattern = %d %p> %m%n
15
16 log4perl.appender.ErrorLog = Log::Log4perl::Appender::File
17 log4perl.appender.ErrorLog.Threshold = ERROR
18 log4perl.appender.ErrorLog.filename = /var/log/arc/janitor_error.log
19
20 log4perl.appender.ErrorLog.layout = PatternLayout
21 log4perl.appender.ErrorLog.layout.ConversionPattern = %d %p> %m%n
```

2.2 Limitations

The Janitor was designed to be used for UNIX-compatible operating systems and tested for various Linux distributions. It should also be functional on MacOS X and Windows with Cygwin or coLinux. The porting of the Janitor to other platforms has not yet been addressed.

The ARC middleware is not ultimately essential for dynamic Runtime Environments. All the Perl code would be functional with any Grid middleware.

3 Usage

The Janitor can be used either with or without the A-REX service. In case A-REX is used, the invocation of the Janitor will be performed in an automated manner. It is then triggered by incoming jobs that request a particular RTE for their execution. If it is not already installed, but

1. found as a MetaPackage in a Catalog that the site supports
2. with a package that first the BaseSystem of the site

then it will be installed without further manual intervention by the Janitor - triggered by the A-REX that received the compute request.

The Janitor's installation will not be affected by this decision pro or cons and integration with A-REX . Both can be installed in parallel. If A-REX is not allowed to install runtime environments upon demand, such automated installations can still be invoked manually via the Janitor's command line interface.

3.1 Janitor with A-REX

Runtime Environments can be specified using the supported job description languages. The most representative two common languages shall be explained at this point: xRSL and JSDL. Listing 3.1 shows the xRSL example in which two runtime environments are requested.

Listing 3.1: Job submission using the xRSL job description language.

```
1  &
2  (executable = "run.sh" )
3  (arguments = "weka.classifiers.trees.J48" "-t" "weather.arff")
4  ("inputfiles" = ("weather.arff" "" ))
5  ("stderr" = "stderr" )
6  ("stdout" = "stdout" )
7  ("gmlog" = "gmlog" )
8  ("runtimeenvironment" = "APPS/BIO/WEKA-3.4.10")
9  ("runtimeenvironment" = "APPS/BIO/WISE-2.4.1-5")
```

The runtime environment names are composed out a directory name, the package name and the version number.

A comprehensive reference manual of the Extended Resource Specification Language (XRSL) can be found at www.nordugrid.org/documents/xrsl.pdf [?]. Within Listing 3.2 an example using JSDL is provided. The specification of how to assign runtime environments in JSDL is currently only defined within the nordugrid jsdl-arc schema http://svn.nordugrid.org/repos/nordugrid/arc1/trunk/src/services/a-rex/grid-manager/jobdesc/jsdl/jsdl_arc.xsd.

Listing 3.2: Job submission using JSDL.

```
1  <?xml version="1.0" encoding="UTF-8"?>
2  <JobDefinition
3    xmlns="http://schemas.ggf.org/jsdl/2005/11/jsdl"
4    xmlns:posix="http://schemas.ggf.org/jsdl/2005/11/jsdl-posix"
5    xmlns:arc="http://www.nordugrid.org/ws/schemas/jsdl-arc">
6    <JobDescription>
7      <Application>
8        <posix:POSIXApplication>
9          <posix:Executable>bin/sleep</posix:Executable>
10         <posix:Argument>120</posix:Argument>
11       </posix:POSIXApplication>
12     </Application>
13     <DataStaging>
14       <FileName>test.sh</FileName>
15     </Source/>
16     <Target/>
17   </DataStaging>
```

```

18     <DataStaging>
19       <FileName>transferGSI-small</FileName>
20       <Source>
21         <URI>gsiftp://pgs02.grid.upjs.sk:2811/unixacl/transferGSI-small</URI>
22       </Source>
23     </Target/>
24   </DataStaging>
25   <Resources>
26     <arc:RunTimeEnvironment>
27       <arc:Name>APPS/BIO/WISE-2.4.1-5</arc:Name>
28       <arc:Version><Exact>2.4.1</Exact></arc:Version>
29     </arc:RunTimeEnvironment>
30     <arc:RunTimeEnvironment>
31       <arc:Name>APPS/BIO/APPS/BIO/WEKA-3.4.10</arc:Name>
32       <arc:Version><Exact>3.4</Exact></arc:Version>
33     </arc:RunTimeEnvironment>
34   </Resources>
35 </JobDescription>
36 </JobDefinition>

```

3.2 Janitor without A-REX

On Linux systems, the Janitor's standalone commandline tool is available as `/usr/lib/arc/janitor`. Some Linux distributions may prefer `/usr/libexec` or similar paths. The script is only functional as root*. To find that binary directly, you may decide to add that location to your `$PATH` environment variable.

The available commands to the Janitor, implemented as options to the `janitor` script, are listed in the Table 3.1.

The most important commands for the Janitor are **register**, **deploy** and **remove**. To register a job along with a set of runtime environments in the Janitor, the first command **register** followed by a job identifier and a list of runtime environments has to be used. A job is identified by a sequence of numbers. Runtime environments are specified by a string containing the name as it is defined within the Catalog (resp. the runtime directory of the grid-manager). The command **deploy** extracts the necessary dependencies of the desired dRTEs and then downloads and installs the required packages.

In order to remove jobs registered in the Janitor, the command **remove** has to be used. The command only removes the job entry and the lock on the runtime environment. If there are no more locks on the runtime environment it is ok to be deleted also physically from the disk. The demand to pass a job number for the removal of a RTE is irritating at first. This shall prevent the removal of runtime environments that are still being used by jobs in the system. Instead, the janitor is informed about a job's termination and is requested to remove the assignment of that job to the runtime environment. Only those RTEs with no job-assignment are eligible for being swept. RTEs come with an expiry time or the command may be performed via the command line.

Easy command line examples are provided in Listing 3.3. You may also want to inspect the `janitor(8)` man page.

Every command has a certain behaviour for its exit status. Table 3.2 lists the possible outcomes. A value of 0 always indicates that no error occurred.

Listing 3.3: Example *log.conf* settings for janitor.

```

# janitor register 1999 APP/BIO/JASPAR-CORE-1.0 APPS/BIO/APPS/BIO/WEKA-3.4.10
# janitor deploy 1999
# janitor remove 1999

# janitor sweep --force
# janitor setstate REMOVAL_PENDING APP/BIO/JASPAR-CORE-1.0 APPS/BIO/APPS/BIO/WEKA-3.4.10

# janitor search JASPAR WEKA
# janitor list
# janitor info 1999

```

*Should you find that constraint unbearable for your purpose, please investigate the file `rjanitor.cc` in the ARC source tree. It wraps the janitor application and as a C binary can be configured to attract root privileges.

Table 3.1: Overview about the available commands to the Janitor.**janitor** [COMMAND] [JOB-ID] [RTE] ...**Command:**

register	Registers a job and a set of runtime environments in the Janitor database. Requires the parameters [JOB-ID] and a list of [RTE]s.
deploy	Downloads and installs the desired runtime environments. Requires the name of an already registered [JOB-ID].
remove	Removes the placeholder of the job on the runtime environments. If no more jobs are using the runtime environment and the lifespan of the runtime environment has expired, the runtime environment can be removed using the sweep command. Requires the [JOB-ID] to be removed.
sweep	Removes unused runtime environments. No further arguments are required. Using the option --force enforces the removal of all unused runtime environments. Runtime environments having the state FAILED will not be removed.
setstate	Changes the state of a dynamically installed runtime environment. This might be useful in case a runtime environment with a state FAILED shall be removed (new state might be REMOVAL_PENDING). Requires the argument [STATE] followed by a list of [RTE]s.
search	Performs a simple search in the catalog and the manually installed runtime environments (runtime_dir). Requires no [JOB-ID] nor [RTE]s, but only a list of string to be searched for.
list	Lists all information about jobs, automatically installed runtime environments and manually installed runtime environments. No additional parameters have to be passed.
info	Renders information about a job. Requires the parameter [JOB-ID].

Job id:

A unique sequence of numbers. Once Janitor registered a job id, it cannot register a second job having the same job id.

Runtime environments:

Runtime environments are defined by a continuous string. The name of valid runtime environment names can be investigated using the **list** or the **search** commands. They are defined in the catalog or by the directories and scripts of the **runtime_dir** of the **grid-manager**.

Once a dynamic runtime environment is installed, it looks completely indistinguishable from traditionally installed runtime environments. This also means that the general concept to have one installation performed for all compute nodes in the network is kept.

3.3 Janitor with A-REX

The motivation to have a runtime environment available comes from the submitters of the grid jobs that depends on that runtime environment for their execution. The site administrator's sole responsibility is to have the dynamic runtime environment at the site's disposal. No more. With A-REX allowed to initiate the commands to the Janitor, no further interaction from the site administrator is required. An exception may be to confirm the consistency of the system when the machine has crashed and the Janitor may still find jobs assigned to runtime environments that are no longer running.

Table 3.2: Possible exit states of the janitor application**Exit status:**

The exit status of Janitor depends on the used command.

register	0	Registration was successful. No noteworthy occurrences.
	1	Registration was successful but some runtime environments aren't installed yet. Deploy is mandatory.
	2	An error occurred.
deploy	0	Successfully initialized job.
	1	Can't provide requested runtime environments.
remove	0	Successfully removed job or no such job.
	1	Can't provide requested runtime environments.
sweep	0	Always returns this exit code.
setstate	0	Changing the state was successful.
	1	Can not change the state.
search	0	Search successfully finished.
list	0	Successfully retrieved information.
info	0	Successfully retrieved job information.
	1	No such job.
	2	Error while retrieving job information.

Another exception for an active involvement of the site administrator is the initial configuration of the Janitor and the updating of runtime environments that are eligible to be installed.

4 Maintenance

This chapter explains how to maintain the *Catalog* and the *Janitor* themselves. It gives detailed instructions how to create new packages for the Janitor. To administrate the catalog, one can either use an ontology editor like Protégé, as explained in the next section, edit the documents manually. In the last section, a typical use case in maintaining the Janitor will be presented.

4.1 Catalog

The Catalog describes runtime environments and is either served through a web server or is distributed together with the Janitor as a regular file. It is specified by a (Resource Description Framework) RDF file assigned to the Janitor using the tag `catalog` within the configuration file (see ??). The format of the RDF file is defined by an RDF schema file `knowarc.rdfs` which can be found along with an RDF example file `knowarc.rdf` in the Janitor source directory <http://svn.nordugrid.org/repos/nordugrid/arc1/trunk/src/services/janitor/resources/catalog/>.

To edit the catalog, the ontology editor Protégé may be used. With some experience gained, one is likely to prefer a manual editing. Figure 4.1 shows the editor while the MetaPackage APPS/BIO/JASPAR-CORE-1.0 of the example file has been selected.

On the left side of the editor the class browser is placed. Three main classes are prepared: `MetaPackage`, `Note` and `Package`. The RDF file is kept in the RDF:XML format, and when inspecting the entry, one will find a direct correspondence to the data stored in RDF:

```
<rdfs:Class rdf:about="&kb;MetaPackage"
  rdfs:label="MetaPackage">
  <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
  <rdfs:comment>
    Reference to a piece of software that shall be made available -
    somehow, and to the expectation to the grid user. The traditional
    way to achieve the installation is via tarballs. Debian packages may
    be an alternative.
  </rdfs:comment>
</rdfs:Class>
<rdfs:Class rdf:about="&kb;Note"
  rdfs:label="Note">
  <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
</rdfs:Class>
<rdfs:Class rdf:about="&kb;Package"
  rdfs:label="Package">
  <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
  <rdfs:comment>
    Superclass of "TarPackage" and "DebianPackage".
  </rdfs:comment>
</rdfs:Class>
```

The `Metapackage` is a general platform-independent description of a `Package`. It can be understood as a reference to a functionality that should be implemented at the remote site. But the exact instruction on how to install it is not given. The "instruction-level" comes with instances of its subclass `Package`.

Links between resources are referred to as *Properties*. The following specifies the dependencies that any package may have on other resources, i.e. on other packages. The dependency on a base system is expressed by the property *basesystem*, which not to be confused with the resource *BaseSystem*.




Figure 4.1: Example of a RDF catalog file as displayed in the program Protégé.

```
<rdf:Property rdf:about="&kb;depends"
  rdfs:comment="lists dependencies of this package"
  rdfs:label="depends">
  <rdfs:domain rdf:resource="&kb;Package"/>
  <rdfs:range rdf:resource="&rdfs;Resource"/>
```

MetaPackages contain one or more instances of the class **Package**, which are then tangible software packages providing the functionality that the MetaPackage references. The packages may be aiming for different versions of the operating system or be different in the way these are installed, but any package being assigned to the same MetaPackage needs to perform the same functionality. MetaPackages are described by the subclasses of **Note**, which in turn has two subclasses: **BaseSystem** and **Tag**.

The **MetaPackage** is described by the subclasses of **Note**. The class **Note** has two subclasses: **BaseSystem** and **Tag** to describe the **MetaPackage**.

The **BaseSystem** describes the Debian release a **Package** refers to (i.e. here etch or sid), i.e. the name of a common installation or a virtual image. The class **Tag** provides small keywords which can be assigned to **MetaPackages** such that they can be found more easily. **TarPackage** and **DebianPackage** or currently the only subclasses of **Package**.

They are representing the necessary information (i.e. URL or Packagename) for the installation. To provide an overview on how the classes are interacting with each other the Tables 4.1, 4.2, 4.3, 4.4 and 4.5 are pictured.

4.1.1 Debian packages - dysfunctional in current implementation

The problem with Debian packages is that these are available only for the local machine and not immediately also for the whole network. This feature was meant for setups that use virtual machines for the execution of jobs.

Table 4.1: Specification of class Metapackage.

Name	Cardinality	Type
description	single	String
homepage	single	String
instance	multiple	Instance of Package
lastupdated	single	String
name	required single	String
tag	multiple	Instance of Tag

Table 4.2: Specification of class BaseSystem.

Name	Cardinality	Type
description	single	String
distribution	required single	String
name	required single	String
short_description	required single	String
url	required single	String

Table 4.3: Specification of class Tag.

Name	Cardinality	Type
description	single	String
name	required single	String

Table 4.4: Specification of class DebianPackage.

Name	Cardinality	Type
basesystem	required single	Instance of BaseSystem
debconf	multiple	String
depends	multiple	Instance of MetaPackage or Package
package	required multiple	String

Table 4.5: Specification of the class TarPackage.

Name	Cardinality	Type
basesystem	required single	Instance of BaseSystem
depends	multiple	Instance of MetaPackage or Package
environ	multiple	String
url	required multiple	String

Those entries can also be used to help the specification of further dependencies of runtime environments. It would then be left to the responsibility of the system administrator to manually (or assisted with scripts) distribute a series of extra Debian packages throughout the compute nodes and use the catalog entry merely to indicate their presence.

One can as such interpret the catalog to represent an interface between software distributed via Linux distri-




Figure 4.2: Directory structure in the tar files for automated installation.

butions and independently from these via grid communities. It should be noted that the Linux distributions have now all started to accept communities to maintain packages, which may bringt many scientific packages away from being traditional ARC runtime environments towards becoming regular packages of some Linux distribution. For Debian, the Debian-Science and Debian-Med^{*} communities are known to be very open to grid and cloud computing.

4.2 HTML interface of the catalog

The dynamic Runtime Environments stored in the Catalog are presented on the aforementioned dedicated web page[†]. This site also links to both the formal Catalog in RDF syntax and its automated transformation to HTML. The latter mimics the traditional site describing Runtime Environments in the Runtime Environment Registry[‡] in order to minimise issues with an eventual transition to the new system.

That page, in a look resembling the classical description of RTE, collects descriptions for Runtime Environments to encourage human site administrators to install these. This HTML page listing the manually or automatically installable RTEs is prepared by the script `web/list.pl`. This script is meant to be run by a mod-perl enabled Apache. The script itself does not contribute to the core functionality of the Janitor. It only performs the human-readable presentation of a catalog's RDF file to users. In the first lines of the script some variables specific to the site are set. To configure the script these have to be changed [?, p. 9].

4.3 Introducing new packages

This section describes how to add new packages to the Catalog. In the current implementation, only tar based packages are processed by the Janitor. Within the here presented example they are assigned to be used together with Debian Etch. This limitation is only literal. There is no restriction for newer Debian distributions.

4.3.1 Debian Etch (tar based)

At the time of writing, only the tape archive (tar) file format is accepted for dynamic Runtime Environment installation, a well accepted file format throughout the UNIX community. The concept reflects the traditional manual approach towards RTE in ARC, for which one directory is made available to all compute nodes. This section explains the inner structure of the tar files for the representation of dynamic runtime environments. Subdirectories are visualised in Figure 4.2.

The tar file contains two directories, named `control` and `data`:

^{*}<http://debian-med.aliath.debian.org>

[†]<http://dre.knowarc.eu:8080/list.pl>

[‡]<http://gridrer.csc.fi/>

data/ contains all software that the grid-job may need

control/ contains files formally specifying how to deal with the information in the **data/** directory.

Upon installation of such tar-based runtime environments, the content of the data directory is extracted to some directory \$BAR. After this unpacking of the tar file, the Janitor executes the install script provided in the control directory. It is executed within the working directory \$BAR. The job of this script is to perform any necessary post-processing. The Janitor stores the file **control/remove**. It will be executed in the same way as **control/install**, just before the tar-package is removed. In most cases **control/remove** will be empty, implying that the working directory \$BAR shall be removed and no other action is required. Finally, the file **control/runtime** is sourced multiple times by the Grid Manager's job-submit script. After installing the package, the Janitor changes all occurrences of %BASEDIR% in the runtime script to \$BAR. Once the tar file was prepared, it must its entry to a RTE Catalogue [? , p. 10]. But the working directory shall not be moved. All post-processing needs to be performed *in situ*.

From such Catalogs, the Janitor finds all information to install packages that possibly have never been installed on the site before. The offers of RTEs in a Catalog are cross-checked against the local infrastructure and a subset of the available packages will be accepted as "installable". This list of installable RTEs is forwarded to the grid information system.

The remainder actions are regular actions performed upon execution of every grid-job. Upon submission, the file **control/runtime** is sourced multiple times by the Grid Manager's job-submit script. Every ARC runtime environment must specify such a runtime script, new is only its specific location as **control/runtime**. Since the directory \$BAR is not known for the individual preparing the runtime environment, that file will instead use the placeholder %BASEDIR%. After installing the package, the Janitor changes all occurrences of %BASEDIR% in the runtime script to \$BAR. To be offered to computing elements for an installation, the such prepared runtime environment must be announced to a Catalog to which the Janitor on the computing element subscribes [? , p. 10].

4.3.2 Automated transformation of install directory to dRTE

The script 'prepareDRE.pl' was created to help with the transformation of a readily installed software into a dynamic runtime environment. I also prepares a complete catalog file that can be offered individually or next to other catalogs. See the associated man page prepareDRE(8) for details.

4.3.3 Prototypes

In order to have an impression how the tar files are created, several prototypes are provided at <http://dre.knowarc.eu>.

Example: WEKA machine learning Java library

The WEKA package for machine learning [?] and the Java Runtime Environment are available as dynamic Runtime Environments. Further packages for bioinformatics comprise dynamic variants of tools for the analysis of transcription factor binding sites. These are already offered for manual installation via the prior mentioned traditional page representing Runtime Environments for ARC. The **data** directory simply contains a ZIP file which needs to be unzipped in the installation directory. For that reason, the **control/install** script is written as follows:

```
#!/bin/sh
set -e # Makes the script to terminate at the first line it fails.

WEKA_ZIP="weka-3-4-8a.zip"
unzip $WEKA_ZIP
rm -f $WEKA_ZIP
```

The runtime script sets the environment variable of the Java Classpath:

```
#!/bin/sh

WEKA_JAR="weka-3-4-8a/weka.jar"
case "$1" in
0) # Just before job submission
# none
;;
1) # Just before job execution
# Initialize the java environment
CLASSPATH="%BASEDIR%/$WEKA_JAR:$CLASSPATH"
export CLASSPATH
;;
2) # After job termination
# none
;;
*)
return 1
;;
esac
```

The remove script, which will be executed right before WEKA is deinstalled, is empty. The Janitor will delete the whole directory, so there remains nothing to be removed in addition. The remove script may e.g. be used to remove indices or other files in /tmp.

4.3.4 Example: ATLAS for High Energy Physics

To address the concerns of the physicists using ARC, a dynamic runtime environment for the ATLAS software suite was prepared. It extends prior work on an automated installation that is available at <http://guts.uio.no/atlas/12.0.6/>.

The preparation comprised the following steps:

- The file system path specifications in the automated installation scripts were modified using the Janitor path variables.
- A tarball was prepared containing a directory structure as illustrated in Figure 4.2. The data directory was empty, since the automatic installation script downloads the software from a remote server.
- An entry was added to the Catalog file.

What sets High Energy Physics software apart is its sheer size. The package in question takes up more than 5 GB. This was a test illustrating the feasibility of using dynamic RTEs in High Energy Physics. The application of the dRTEs for ATLAS needs to wait for the planned web service extension of the Catalog. With such a service, e.g. a software manager of a big experiment will be able to deploy software packages on production sites simply by creating a tarball and adding an entry to the Catalog.

5 Technical Motivation and Concepts

The current implementation, using the rather advanced concepts of the semantic web, may seem unexpected. They were chosen since they are expected to scale with all the extra demands foreseen for the Janitor.

5.1 Implementation

5.1.1 Choice of Perl

The main language for the implementation of the functionality of the dRTEs' functionality is Perl. And it is solely required (exceptions are the integration with the Grid Manager and the Web service) for the Janitor. The language was perceived a side-issue, the complexity of the data expected was of major concern. This led to the choice of RDF with its intrinsic query engine and Perl offering access to the Redland library.

5.1.2 Modular structure

In the pre-web-service implementation the Catalog remains a static web page. The Perl code is split into multiple modules as depicted in Figure 5.1.2. The modules can be separated into two functional groups. One addresses the retrieval of information from the Catalog's RDF file in the left major branch of the figure. The other addresses the process of fetching and installing the packages.




Figure 5.1: Modules of the Janitor and their dependencies

In order to get a more detailed view on the full functionality of the envisioned system it is suggested to consult the Design Document^{*}.

5.2 Concepts

In the following, some paragraphs fail render it difficult to clearly distinguish between conceptional truths and the state of the current implementation. Please read carefully.

5.2.1 States of Runtime Environments

A major motivation for the managed, manual initiation of dynamic RE installation is the subsequent manual verification of the installed packages – prior to their use in production. It would be nice to see the Janitor and/or the Catalog prepare for reviews by selected users. This has not yet been implemented.

With an automation of the installation, the verification of that process shall be performed externally to that process. At this time, only the automation of the installation has been implemented. To reflect the progress the external verification has made, REs are said to be in states. The current implementation lists installable REs aside the installed REs in the grid information system, in order to stimulate grid clients to submit packages. The here described states will be represented to the clients in upcoming developments.

These states are specific for every compute element (CE) and communicated between the Janitor and the Execution Service. Table 5.1 shows all possible states, while Figure 5.2 displays the transitions between the states that a Runtime Environment may be in during its life time at a particular CE.

State	Description
UNAVAILABLE	The RE is not available for the BaseSystem (see ??) the site uses.
INSTALLABLE	The RE is available for the BaseSystem the site uses and it will be automatically installed once a job requests it.
INSTALLING/a	A job requested the RE and it is currently being installed
INSTALLING/m	The RE-administrator requested the installation of the RE. Its currently being installed.
FAILED	The installation process failed.
INSTALLED/a	The RE is installed dynamically.
INSTALLED/m	The RE ist installed manually by the RE-administrator
BROKEN/m	The RE is installed but failed tests of the RE-administrator
VALIDATED/m	The RE is installed and successfully passed the tests of the RE-administrator
REMOVAL PENDING	The RE is still installed but will be removed as soon as possible. It is not available to new jobs.
REMOVING	The RE is currently being removed.
INSTALLED/s	The RE was installed in the traditional way by the site administrator.
BROKEN/s	The RE was installed in the traditional way and failed validation by the RE-administrator,
VALIDATED/s	The RE was installed in the traditional way and was successfully verified.

Table 5.1: States a Runtime Environment can possibly be in.

The concept of ARC prevails to have a single directory into which to install software to be execute on all compute elements of the site. Otherwise, the installation of a runtime environment would be required to be performed just when the jobs initiates its computation.

^{*}<http://www.knowarc.eu/documents/Knowarc.D1.1-1.07.pdf>

In principle this would be doable, particularly for those runtime environments that are part of a trusted Linux distribution already. But this would also mean that the compute nodes become inhomogeneous, and at the time of writing this was considered undesirable.




Figure 5.2: Relationships between the possible states of Runtime Environments. Red arcs represent human interaction. The distinction between /a, /m and /s states does not need to be visible for all clients.

The manually induced transitions are marked in red, the automated transitions in black. A transition between states can be induced automatically (i.e. by the advent of a job requesting a particular dynamic RE) or manually by the site's administrator or an individual with respective rights to use the Janitor's command line.

5.2.2 Subset of that functionality as implemented today

Upon presentation of a package name to a Catalog, from which details about the package are retrieved, a CE may classify a package to be **INSTALLABLE** if all the dependencies are installable or already **INSTALLED**. The installation can be performed manually (**INSTALLING/m**) or in an automated fashion (.../a). Should the installation process return an error, then the installation has **FAILED**. Once the installation succeeded, the installed package is validated for its correctness. Should that process fail, then the package's state is said to be **BROKEN**.

Automatically installed packages can be removed by the automatism. A manually installed package or one that has failed to be installed, can only be removed upon manual induction. The .../s states represent those Runtime Environments that are installed in the original manual way of RE installation in ARC 0.6.

5.2.3 Dependencies on Job IDs

The Janitor was designed with the job execution in mind. As such, all actions it performs are driven by the need of a job. And those jobs should then be specified, to learn when the demand for a particular runtime environment has ended.

This concept partially conflicts with the idea to use the Janitor as an aid for the manual invocation of an installation – there is no job ID that could be assigned for that process. Today, the administrator is suggested to use a special number, e.g. 0, to indicate such manual installs. This will change in near future.

5.3 Job states

The Janitor manages the states that the runtime environments at a particular compute element are in. However, it is also most important for the Janitor to be aware of the jobs that depend on the installation of a RE. REs still in use should not be removed until the respective job has completed its computations. The installation or removal of REs by the Janitor is perceived as a mere consequence of jobs demanding a RE or not, thus, the communication between the job-manager AREX and the Janitor will be performed on that 'job level'.

The Janitor has two states for jobs: **PREPARED** and **INITIALIZED**. After a job has been successfully registered with the Janitor, its state will be set to **PREPARED**. Invalid jobs are not cached. After the Janitor is requested to deploy the runtime environment, the state of the job will change to **INITIALIZED**. If an unforeseen exception occurs during that process, The Janitor will drop the job from its database and set the affected runtime environments to the state **FAILED**.

5.4 Integration with AREX




Figure 5.3:




Figure 5.4:

5.5 WebService Interface

Default port number: 55555

Client command equal, except assignment of HED.xml
(from /arc1/trunk:12561)

Proposal for SOAP messages:

namespace: dynamicruntime or janitor




Figure 5.5: To be translated and beautificated. SVG file is missing!

Create WSDL files for that Permission concepts: Depending on certificates. Certain certificates may sweep. Defined in service_HED.xml. Evaluated in:??

Listing 5.1: Example *arc.conf* settings for janitor.

```

1 <Request action="SEARCH|SWEEP|LIST|DEPLOY|REMOVE|CHECK|REGISTER">
2   <Initiator jobid="1234"/> <!-- Needed for: CHECK|REGISTER|DEPLOY|REMOVE -->
3   <!-- May contain no jobID in this case a new one will
4     be created and returned via the response message-->
5
6   <Runtimeenvironment type="dynamic"> <!-- Needed for: SEARCH|REGISTER-->
7     <Package name="APPS/BIO/WEKA-3.4.10"/>
8     <Package name="APPS/BIO/WEKA-3.4.11"/>
9   </Runtimeenvironment>
10
11   <!-- SWEEP and LIST only works, if the TLS-administrator
12     identity
13     (which is assigned in the arched configuration file)
14     is
15     to be found by the SecHandler. Both need neither
16     initiator
17     nor runtimeenvironment elements -->
18 </Request>

```

Listing 5.2: Example *arc.conf* settings for janitor.

```

1 <response action="SEARCH|SWEEP|LIST|DEPLOY|REMOVE|CHECK|REGISTER">
2   <initiator jobid="1234"/> <!-- Needed for REMOVE|REGISTER|DEPLOY|CHECK-->
3   <result code="0" message="Sucessfully initailized job."> <!-- -->
4   <jobs> <!--LIST|CHECK-->
5     <job jobid="1234">
6       <created>1234567890</created> <!-- in unix time-->
7       <age>0</age> <!-- in seconds-->
8       <runtimeenvironment>
9         <package>APPS/BIO/WEKA-3.4.10</package>
10      </runtimeenvironment>
11      <state>INITIALIZED</state>
12    </job>
13    <job jobid="4321">
14      <created>1234567891</created>
15      <age>0</age>
16      <package>APPS/BIO/WEKA-3.4.10</package>
17      <state>INITIALIZED</state>
18      <runtimeenvironmentkey>APPS_BIO_WEKA_3_4_10-835614b62c98c4eb6cb03d74d3161b5d</
19        runtimeenvironmentkey> <!-- at least CHECK -->
20      <uses>/nfshome/nowarc/dredesign/src/services/dRE3/perl/spool/runtime/
21        jre_57Tike1UVz/runtime</uses> <!-- at least CHECK -->
22      <uses>/nfshome/nowarc/dredesign/src/services/dRE3/perl/spool/runtime/
23        weka_wHfyytar1E/runtime</uses> <!-- at least CHECK -->
24    </job>
25  </jobs>
26
27  <runtimeenvironment type="local"> <!-- Needed for: LIST|SEARCH -->
28    <package name="APPS/BIO/MUSTANG-3.0-1"/>
29    <package name="APPS/BIO/EXONERATE-2.1.0-1"/>
30  </runtimeenvironment>
31
32  <runtimeenvironment type="dynamic"> <!-- Needed for: LIST -->
33    <package name="APPS/BIO/WEKA-3.4.11">
34      <state>INSTALLED_A</state>
35  </runtimeenvironment>

```

```

32         <lastused>1234567890</lastused>
33         <jobid>1234</jobid>
34     </package>
35     <package name="APPS/BIO/WEKA-3.4.10">
36         <state>INSTALLED_A</state>
37         <lastused>1234567890</lastused>
38         <jobid>1234</jobid>
39         <jobid>4321</jobid>
40     </package>
41 </runtimeenvironment>
42
43 <runtimeenvironment type="installable" <!-- Needed for: LIST -->
44     <package name="APPS/GRAPH/POVRAY-3.6">
45         <description>The Persistence of Vision Raytracer</description>
46         <lastupdate>1234567890</lastupdate>
47     </package>
48     <package name="APPS/BIO/WEKA-3.4.8A">
49         <description>WEKA Machine Learning Software</description>
50         <lastupdate>1234567890</lastupdate>
51     </package>
52 </runtimeenvironment>
53
54 </response>

```

5.5.1 What happens during installation

Register

1. Look up Catalog and find corresponding runtime environments.
2. Check if dependencies are available (installed or installable)
3. Store job in *registrationdir*

Deploy

1. Load job out of *registrationdir*
2. Look up Catalog and find corresponding runtime environments.
3. Check if dependencies are available (installed or installable)
4. Download RTEs into the *downloadaddr*
5. [to be continued or skipped]

5.5.2 Security Consideration

Security is a major concern for grid systems. Any additional feature and especially an automated software installation inherently introduces security threats. This section addresses those and describes the available solutions to limit security risks.

In the current installation, every user authorised to execute a job is also authorised to install a REs. Restrictions are only imposed on the set of dynamic RTEs that are available for installation. Restrictions are imposed by the site administrators on the descriptions that are given by the Catalog that is offering the package. These descriptions may explicitly mention dynamic REs' names, e.g. a regular expression on these, or refer to tags of packages that categorise these. However, the core of these controls lies with the maintainers of the Catalog, who needs to be trusted.

All dynamic REs are installed in separate directories. The provisioning of disk space is the duty of the site administrator. In the current implementation, the installation is completely transparent to the user:

- Dynamic RTEs are not distinguished between *installed* and *installable* in the information system.
- No status information is given at the time a dynamic RE is installed.

Malevolent regular users with respective training in using system exploits to gain root access are likely to find security holes by regularly submitted scripts. The authentication and authentication of users, together with respective logging, is the major defense against such attacks. What is consequently left to be protected against are unwanted side-effects by the installation of software.

The worst case scenario would be the installation of a RE that overwrites system files. With the current implementation, which is based solely on tar files, this is barely possible, unless such is performed by the install scripts that accompany the tar files. **Clarify root vs non-root execution of Janitor.**

The installation of packages from the Debian distribution (or other packages of mainstream Linux distributions) is sought to reduce the complexity and burden in the maintenance for dynamic REs. In the current implementation, Debian packages may be installed only by their transformation into tar files. With the advent of the interface for the virtualisation of the grid infrastructure, it is anticipated to work with native packages of the Debian Linux distribution. The reuse of packages that passed many eyeballs - as it is the case with packages from major Linux distribution - security is further increased or becomes as high as with the operating system underneath virtual clients.

Summarising, there is general concern about the security of grid computing. Dynamic REs introduce new dangers since a manual control at the grid site is substituted by a remote process that is out the direct supervision of a local site administrator. The signing of packages by known and directly or indirectly trusted developers is a good indicator that no malevolent individuals have tampered with the binary. The site administrators can limit the sources of packages and specify packages that are eligible or excluded from installations.

6 Outlook

6.1 Representation of dynamic RTEs in the information model

Dynamic RTEs require an extended representation in the information model. The Application Software description should be able to distinguish installed RTEs from installable RTEs, potentially offer insights on the state that an dRTE is in. This work is planned to be carried out as part of the Glue-2.0 effort of OGF*.

6.2 Integration with Workflow Management

Future development of ARC aims at integrating grid computing with workflow tools for the web services that have a growing user base in bioinformatics. The challenge is to prepare RTEs for programs or databases and to offer such concisely to users of the workflow environments. In the bioinformatics community, such are today offered as web services. This anticipated development instead fosters the dynamic installation on the grid whenever appropriate to allow for special computational demands in high-throughput analyses. Conversely, because of the increased complexity of workflows with respect to the already today not manually manageable number of RTEs, without an automatism for the automated installation of software packages on the grid, the use of workflows in grid computing seems mute.

6.3 Implementation of a Catalog service

A Catalog service is planned to be implemented on top of the ARC HED component. This service will render the currently used locally accessible RDF file externally accessible. Selected users are then allowed to remotely add/edit/remove REs to/from it. The Janitor will access the content of the Catalog through a well-defined Web Service interface.

6.4 Integration with the Virtualization work

The RDF schema nicely prepares for the upcoming virtualisation of worker nodes. Hereto, the **BaseSystem** indicates a virtual image to which further packages, the dynamic REs, would then be added. How exactly the dynamics are integrated will depend on how dynamic the virtualisation of the nodes is. In the simplest scenario, a worker node's CPU will only be occupied by a single virtual machine and that will not be changed. In this case, there is no difference to the setup of the Janitor with today's static setups.

However, if the BaseSystems can be substituted dynamically, then a RE can possibly be offered via multiple BaseSystems. The RDF Schema describes BaseSystems as separate instances and as such differs from the current RE registry. Heuristics that prefer one BaseSystem for another can make direct use of the data that is presented in the schema. The integration of packages from Linux distributions in the description of REs is essential to have a means to decide for the equivalence of manual additions and the functionality that comes with BaseSystem.

6.5 Use of RDF

The RDF is "correctly" used within the Janitor via the Redland libraries. What is still missing is a semantical reasoning on what packages to allow or disallow. For instance: allow all packages associated with

*OGF GLUE: <https://forge.gridforum.org/sf/projects/glue-wg>

bioinformatics. This shall wait a bit longer until more experiences on the community's demand for this service have been made.

6.6 Manual Verification

There is yet no explicit notion of a concept on how users can report on the reliability of individual compute nodes and use such information for the decision making on where to sent their jobs. And with an increasing complexity of software installed, being used across versions, one will rather trust one's very own experiences for individual sites than some external repository or other pieces of information that may be weeks or months old.

Consequently, no means for a manual verification and the communication of results of such have yet been implemented.

7 Appendix

7.1 Useful tutorials and documentations

- **Another document describing Janitor.**

D2.5-1 RDF Based Semantic Runtime Environment (RE) Description And Dynamic RE Management Framework Including Creating Proof Of Concept Bioinformatics REs, Daniel Bayer and Steffen Möller and Frederik Orellana[?]

Bibliography