NGS Uniform Execution Environment

Matthew Viljoen / Jonathan Churchill

STFC RAL, UK

1. Document History

Version	Date	Author	Comments
0.1	26/3/07	M Viljoen	First version
0.2	28/6/07	M Viljoen	Provision for a 3rd field for version number in application script name
0.3	26/9/07	M Viljoen	Updated to work with the GIP which is now added to VDT
0.4	2/11/07	M Viljoen	Now added a GIP UEE plugin, written by Jason Lander
0.5	24/2/10	J Churchill	Updated to latest

2. Executive Summary

The NGS Uniform Execution Environment (UEE) is a solution for invoking the same applications across different sites of the UK National Grid Service, regardless of where the application may be installed at each site and other local settings. By using standard tags as defined in the GLUE Schema[1] published to the NGS BDII server[2], the UEE also provides a mechanism for determining which NGS sites has a particular application installed.

In this way, users can benefit from using the NGS Resource Broker (WMS) to dynamically select which sites to run a particular job, either by writing a JDL file to specify the required application or via the NGS Portal using an existing JSDL file. The UEE enables the environment for running applications to be set up on the nodes, allowing for different configurations across sites, and shielding users from this additional complexity when submitting jobs. In this way, it becomes trivial to transparently send jobs to a heterogeneous Grid such as the NGS where the same application may be deployed on different platforms and architectures.

Finally the UEE also enables a dynamic way of publishing the list of applications that are available on different sites in a human readable format on, for example, the NGS website.

[1] http://glueschema.forge.cnaf.infn.it/

[2] http://www.ngs.ac.uk/site-level-services/bdii

3. Technical Details

3.1. Application Scripts

In the UEE an application on the NGS is invoked using a script which must be installed at the same location across all sites (`/usr/ngs`). Application script names are formed by the name of the application and optional version numbers. The script sets the local execution environment of the application and calls the application on behalf of the user. If the user needs to pass any arguments onto the application, then the script should accept the arguments and pass them onto the application.

Although the absolute filename of each script must be the same across all sites, the contents of the script can be modified to account for different local settings (location of application, environmental variables/values, modules to be loaded etc.). The script can also account for factors such as the platform and architecture where the application is deployed, thus abstracting the user from the non-homogeneity of the NGS when launching jobs across NGS sites using the RB.

The scripts must adopt the following naming convention:

$SCRIPT_LOCATION/<APPLICATIONNAME>[_<VERSION>[_<MAJORVERSION>[_<MINORVERSION>]]]

The filename can contain any combination of letters and numbers. All letters in the script name must be capital letters. If version numbers are included, there should be a softlink to the script from another file without the version numbers. This is to cater for users who do not care about using a particular version of an application, or for users who want to use a default version (maybe the latest version), decided by the application deployer.

As an example, at the time of writing, the following scripts are deployed at RAL:

```
[ngs0015@grid-data ngs0015]$ ls -la /usr/ngs/
total 12
drwxr-xr-x 2 root root 4096 Mar 26 13:37 .
drwxr-xr-x 18 root root 4096 Mar 16 11:56 ..
lrwxrwxrwx 1 root root 16 Mar 26 13:37 GAUSSIAN -> GAUSSIAN G03 C02
lrwxrwxrwx 1 root root 21 Mar 16 14:59 GAUSSIAN G03 C02 -> /usr/local/bin/rung03
lrwxrwxrwx 1 root root 14 Mar 26 13:37 SIESTA -> SIESTA 2 0
-rwxr-xr-x 1 root root 125 Mar 16 14:21 SIESTA_2_0
```

3.2. Publishing the Available Scripts

The list of available scripts deployed in `/usr/ngs` is published to the NGS BDII server as `GlueHostApplicationSoftwareRunTimeEnvironment` tags which are defined in version 1.2 of the GLUE schema. This list is generated periodically by the Generic Information Provider (GIP), included in VDT version 1.6.1 and later. More specifically, the GIP runs a dynamic plugin **`ngs-uee-gip-plugin`** (Latest at http://forge.nesc.ac.uk/projects/ngs/ or see Appendix A for June 2009 version) which must be placed in **`$VDT/lcg/var/gip/plugin`**
This script prefixes the name of the script in /usr/ngs by “NGS-UEE-“ to avoid

confusion with other tags published to `GlueHostApplicationSoftwareRunTimeEnvironment` for example by gridPP sites for LCG VOs

3.3. Listing Tags and Executing Applications

Once sites are publishing tags for their application scripts, they are visible to any gLite/LCG client software configured to work with the NGS BDII server. Tags can be listed by executing the **`lcg-info`** command on the gLite User Interface (UI) at RAL: ngsui03.ngs.ac.uk:

```
[ngs000@ngsui03 ngs000]$ lcg-info --vo ngs.ac.uk --list-ce --attrs 'Tag'
- CE: ngs.leeds.ac.uk:2119/jobmanager-pbs-mpi
  - Tag                  MPICH
                         NGS-UEE 1 0
                         NGS-UEE-MPICH
                         NGS-UEE-AMBER
                         NGS-UEE-AMBER 10
                         NGS-UEE-AMBER 10 0
                         NGS-UEE-AMBER 9
                         NGS-UEE-AMBER 9 0
                         ……
                         NGS-UEE-SIESTA
                         NGS-UEE-SIESTA 2 01
- CE: ngs.oerc.ox.ac.uk:2119/jobmanager-pbs-workq
  - Tag                  MPICH
                         NGS-UEE 1 0
                         NGS-UEE-AUTODOCK
                         NGS-UEE-AUTODOCK 4 0 1
                         NGS-UEE-AUTODOCK 4 2 2 1
                         …...
- CE: vidar.ngs.manchester.ac.uk:2119/jobmanager-pbs-workq
  - Tag                  MPICH
                         NGS-UEE 1 0
                         NGS-UEE-GROMACS
                         NGS-UEE-GROMACS 3 3 1
                         …..
- CE: ngs.rl.ac.uk:2119/jobmanager-lsf-ngs
  - Tag                  MPICH
                         NGS-UEE 1 0
                         ...
                         NGS-UEE-GAUSSIAN03
                         NGS-UEE-GAUSSIAN03 C 02
                         NGS-UEE-GAUSSIAN03 E 01
                         NGS-UEE-GAUSSIAN09
                         NGS-UEE-GAUSSIAN09 A 02MPICH
                         NGS-UEE-SIESTA
                         NGS-UEE-SIESTA 2 01
                         …..
```

An example job that can be submitted to the WMS from a UI to run the SIESTA application on all sites which has the “SIESTA_2_0” application script and application installed is as follows:

siesta.jdl

```
Type = "Job";
JobType = "MpiCh";
Executable = "/usr/ngs/SIESTA 2 01";
CpuNumber = 4;
StdOutput = "h2ob.out";
StdError = "h2ob.err";
InputSandbox = {"h2o.fdf"};
OutputSandbox = {"h2ob.err","h2ob.out"};
RetryCount = 3;
ShallowRetryCount = -1;
Requirements = Member("NGS-UEE-SIESTA 2 01",
other.GlueHostApplicationSoftwareRunTimeEnvironment);
```

3.4. Dynamic Applications Information

The "Applications Software on NGS"[3] page on the NGS website[2] is a dynamic (PHP) page which is populated by querying the BDII server for tags from different sites. This page is a repository of all tags that are currently used across NGS sites, so if a new application is deployed on a site, the deployer can check whether the application has already been installed elsewhere. If it has, the deployer can use the existing tag name (and application script). The "API"s for the UEE scripts are also documented on the NGS website[5] to ensure compatibility between sites.

Application script templates of existing applications on the NGS can be copied from ngs.rl.ac.uk which has a complete set for the use of application deployers.

3.5. Future Work

Once the NGS Application Repository/Job Submission Portal[4] supports the WMS, it need no longer hardcode the sites where different JSDL may be run from its application repository, but rather incorporate a way of selecting sites which publish the appropriate tag in the JSDL.

[3] http://www.ngs.ac.uk/applications

[4] https://portal.ngs.ac.uk/

[5] http://www.ngs.ac.uk/applications/UEE-API

Appendix A - ngs-uee-gip-plugin

```
#! /usr/bin/env perl
#
# $Id: ngs-uee-gip-plugin 1529 2009-06-18 10:09:48Z lecjpl $
#
# Creates GLUE GlueHostApplicationSoftwareRunTimeEnvironment entries
# for UEE scripts in /usr/ngs

use warnings;
use IO::File;
use IO::Dir;

my $Pub prefix="NGS-UEE-";
my $Pub version="NGS-UEE 1 0";

my $UEE dir="/usr/ngs";
exit unless -d $UEE dir;

my @Default tags = ();

my $Default tags file = "$UEE dir" . "/default-tags";
if ( -f $Default tags file ) {
      @Default_tags = map { split /\s+/ }
            ( IO::File->new($Default tags file)->getlines );
}

my $fqdn=$ENV{GLOBUS HOSTNAME} || (qx{/bin/hostname --fqdn} =~
m{(\S+)})[0];

my $dn="GlueSubClusterUniqueID=$fqdn,GlueClusterUniqueID=$fqdn,mds-
vo-name=local,o=grid";

my @tags = sort grep { /^[A-Z0-9]+ ?/ } ( IO::Dir->new($UEE dir)-
>read );
print "dn: $dn\n";
print "GlueHostApplicationSoftwareRunTimeEnvironment:
${Pub version}\n";
for my $t (@Default_tags, @tags) {
      print "GlueHostApplicationSoftwareRunTimeEnvironment:
${Pub prefix}$t\n";
}
print "\n";
```