

Table of Contents

Multiplatform support for Grid applications.....	1
Situation for our LHCb supported applications.....	1
Limitations of compatibility.....	1
Brainstorming on the LHCb process.....	1
Examples.....	2
Summary of actions.....	3

Multiplatform support for Grid applications

Hereafter I call "platform" the concatenation of the OS and the architecture (i.e. slc3_ia32, slc4_ia32 or slc4_amd64)... Yes, I agree the names are not necessarily correct (ia64 exist, sl4 is not necessarily slc4 etc...) but have the advantage of uniqueness ;-). This is what is supported by the LCG-AA. The compiler is added for forming the CMTCONFIG env variable, but currently there is only one compiler per OS (gcc323 for slc3 and gcc34 for slc4)

Situation for our LHCb supported applications

1. Production jobs: all releases can run on any platform provided it has a 32-bit compatibility. The current production releases are available on slc3_ia32 and slc4_ia32. As from now the new releases will be available also on slc4_amd64 (expected for production soon for DaVinci, not for the other applications).
2. Analysis jobs: users have prepared their software on a well defined platform (we don't intend cross-compilation at this stage ;-). Currently the settings on interactive sessions select only the OS, and the platform is always ia32, but we should be able soon to give users the possibility to upgrade to amd64 if they wish (for the supported applications). The CMTCONFIG used for building the software should be transmitted by ganga to DIRAC and should then be dealt with in the expected manner....

There is a need for a negotiation between the job, the RB and the WN in order to run the right version at the right place... I would like to launch brainstorming on this as we better be ready when WNs start to emerge on new platforms.

Limitations of compatibility

1. User job prepared on a given platform (in this case it is sure the application exists for that platform ;-). This should include a JDL requirement such as: `Platform = "slc4_amd64_gcc34"`
2. The application/version is available for a limited number of platforms. One should be able to discover this from the SW installation or availability of tarballs
3. WNs have limited platform capabilities (one or more platforms)

Worker node OS	Platforms
slc3 WN	slc3_ia32
slc4 WN 32-bit	slc3_ia32, slc4_ia32
slc4 WN 64-bit with 32-bit compat	slc3_ia32, slc4_ia32, slc4_amd64
slc4 WN 64-bit no 32-bit compat	slc4_amd64

Brainstorming on the LHCb process

When a pilot starts, it should first determine what are its capabilities, then match jobs depending on this capabilities.

It would be good for production jobs if the JDL could contain also the available platforms (as a list). This can be guessed from the application/version and generate something like:

```
Platform = "slc3_ia32_gcc323, slc4_ia32_gcc34" for the current applications.
```

Once the job is there, the platform that will be used should be selected from the AND of the job and WN platforms, with a priority list: slc4_amd64 / slc4_ia32 / slc3_ia32. Then CMTCONFIG is set and the job behaves as currently (install_project etc...).

In order not to waste pilots and also to be sure pilots are sent to the appropriate platforms, it would be good that LCG/gLite gives the possibility of targeting platforms if needed. For this, it should be assumed that a given CE only gives access to one type of machine, and publishes what it provides. However the guidelines on this topic are not very explicit:

How to publish the OS name [↗](#)

How to publish my machine architecture [↗](#)

The information published there cannot be easily mapped to the platform convention above. In particular the OS version contains the minor version, platform doesn't mention if a 64-bit architecture has 32-bit compatibility (but this may be a must?).

Examples

On a CERN slc3 CE, one gets:

```
GlueHostOperatingSystemName: Scientific Linux CERN
GlueHostOperatingSystemRelease: 3.0.8
GlueHostOperatingSystemVersion: SL
```

On SLC4 at CERN:

```
[lxplus214] ~ > lsb_release -i | cut -f2
ScientificCERNSLC
[lxplus214] ~ > lsb_release -r | cut -f2
4.5
[lxplus214] ~ > lsb_release -c | cut -f2
Beryllium
```

which would translate into

```
GlueHostOperatingSystemName: ScientificCERNSLC
GlueHostOperatingSystemRelease: 4.5
GlueHostOperatingSystemVersion: Beryllium
```

On an SLC4 site in Belgium (BEgrid-UGent), one gets:

```
GlueHostOperatingSystemName: ScientificSL
GlueHostOperatingSystemRelease: 4.4
GlueHostOperatingSystemVersion: Beryllium
```

On CSCS-LCG2 (Manno)

```
GlueHostOperatingSystemName: ScientificCERNSLC
GlueHostOperatingSystemRelease: 4.4
GlueHostOperatingSystemVersion: SL
```

So I guess we can expect a finite but large list of combinations ;-). None of these sites publishes an architecture (GlueHostArchitecturePlatformType hence there is no way to know at the CE level whether it is 32 or 64-bit. It can however be checked at the WN level as described above.

Question: how from an LHCb requirement "slc4_ia32_gcc34" can one match the last 3 CEs above that all have a different Glue description? Is it possible to match "Scientific" and the first "4" only?

Summary of actions

For being able to select the right platform corresponding to a job:

- Define a JDL parameter for specifying the platform in DIRAC jobs (DIRAC WMS)
- Define the platform requirement from the application / version (production tools) (DIRAC production tools)
- Specify the current CMTCONFIG as JDL parameter from ganga through the DIRAC API (ganga, DIRAC API)
- Implement a script for determining the capabilities of a WN (not only setting CMTCONFIG) (Environment scripts)
- Include platform matching (DIRAC WMS)
- Adapt software installation for setting CMTCONFIG (Environment scripts)

For pilot job submission, it looks quite more tricky. If jobs are targetted to a specific platform, one should find an efficient way of matching this with the Glue schema. This needs study and verifying the publication is consistent. Currently the architecture is missing.

- Ask LCG to publish the architecture in the Glue schema
- Work out a matching expression for targetting SLC4 CEs (or excluding them)

-- Main.phicharp - 21 Jun 2007

This topic: LHCb > MultiPlatform

Topic revision: r2 - 2011-06-22 - AndresAeschlimann



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors. Ideas, requests, problems regarding TWiki? Send feedback