
Subject: GRID, PandaROOT, and static binaries

Posted by [Johan Messchendorp](#) on Wed, 13 Feb 2008 19:13:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

Message from Dan

Thanks for all the work done on this. I am really glad to see that serious work is now done on the grid and me and the site admins will do our best to offer you a stable system.

I agree that there are many options to package management on our grid and only trying we will find the best solution.

My main concerns on this are the following:

- static binaries once we can agree on some batch-optimized code would be the best choice by far; even if we have to break our simulation in an execution chain
- would prefer to have one maximum two packages and not a too lengthy chain of dependencies
- would be prudent to not rely on downloads at installation (wget and such)
- it is crucial that installation/compilation errors are properly handled and passed upwards so that the packman (package manager service on grid) knows what's going on and can tell the status of the installation

Please remember that we will not install packages on the grid weekly, but only once they are tested, benchmarked and deemed fit for production (be it a test production).

That means that these packages will not have to download the 'latest' sources and that these packages will be optimized for batch execution. Please feel free to install and test pandaroot on grid, but the strategy for package management should not be based on a weekly test scenario.

Please tell me what you think ?

Cheers,
Dan

P.S. Should we not post these discussions on one of our forums ?

Florian Uhlig wrote:

Dear all

Some comments on Johans comments.

1) keep the external packages separate, with name cbmsoft::r<date of release>. I think only twice per year, a new update of this will appear.

PandaROOT will be updated more frequently (that would be a good reason not to combine the two for the grid). We could consider to

have a "./configure grid" option or so for all the external packages if we need to set specific GRID related stuff. The question remains is

whether we would like to include cernlib in the external packages as well.

We had a discussion about this a couple of meetings ago, and at that time it was decided not

to do this in order to minimize the size of the external package. With the option "grid" in the configure script, we could automatically download the cernlib stuff via the web (wget), similar to what is done with geant4 datafiles. Alternatively, we might consider to provide binaries for GRID applications for available platforms: Linux-i686, Linux-x86_64, ... This would make the installation procedure via packman much easier. I agree to have the external packages as a separate package and I don't like the idea to split this even more. The only difference between the grid and the normal version of the external packages is, that the grid version has everything on board. This was done to have a version which does not use wget because it was not clear if all the machines can access the internet. If this is no point anymore I think we can live with only one version of the external packages. The idea with the binaries was also discussed during the Vienna workshop but that would mean that you need a static binary, due to all the different versions of the OS on the different production sites. To produce such a binary is not straightforward. I worked on that topic for a while but I forgot about it. If this is again an issue I will continue on that topic.

2) For Pandaroot, I propose to use the package name "pandaroot::r<stable version number>" (for production releases) and "pandaroot::revXXXX"

for test revisions. I would think that the easiest is that the release numbers match the the subversion numbering, instead of have pandaroot::0.1 etc.

I fully agree.

The DPMgen and EvtGen stuff, we could compile and build as well within this package, and place the stuff in a "/bin" directory.

At the moment, I have created virtual packages for this, which is not really needed and only brings too many unnecessary dependencies. Also here,

lets do consider to provide binaries instead of the sources with compilation scripts.

Here again you will need a static binary.

Ciao

Florian

--

Dr. Dan PROTOPOPESCU | Tel: +44 141 330-5531
Department of Physics & Astronomy, | FAX: +44 141 330-5889
University of Glasgow, | Mob: +44 794 046-3355
Glasgow, G12 8QQ, Scotland, UK | WWW: Google me!

Subject: Re: GRID, PandaROOT, and static binaries

Posted by [Johan Messchendorp](#) on Wed, 13 Feb 2008 20:18:13 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Dan and all others,

I am still a little bit puzzled about the usage static libraries and binaries. As I understand, one still needs to obtain them for the different hardware configurations and compilers, which means one has to compile and build them on various machines anyhow. In this respect, we might as well provide shared libraries within the packman software packages?!? Furthermore, if we

would go for static binaries, does this mean we need to make one executable which archives ROOT and all Panda related stuff? This would be a major job, no? I certainly would like to keep the macro-oriented nature of the simulation and analysis, which would imply we require one static "(panda)root.exe" which includes all the objects of Pandaroot. I don't know, maybe it would be more efficient to limit the static stuff for some of the software components, such as eventgenerators and cernlib.

From my experience in installing the cbmsoft and pandaroot via packman on the Grid sites, I have to say that for most of the cases it works like a charm. In particular, since the configuration scripts of Florian and Mohammad are really excellent. Also pandaroot is not a problem at all to let it compile on the different sites automatically (thanks to the same persons and cmake). The only problems I encounter were related to cernlib and eventgenerator stuff (hence my comment above). Actually, once we have solved all the problems we find during the software installation on the sites (and there are not many left!), we surely have it easier next time we make an upgrade of the software. From this perspectives, to employ static binaries would be an extra overhead which might not be necessary. Maybe from the batch system point-of-view it might be advantageous, but that we have to quantify (where and what are the problems).

Concerning the number of dependencies etc: I would say lets indeed avoid to have many (sub)packages. Actually, only two packages are of relevance: cbmsoft (grid specific with geant4 data files included to avoid dependence on wget installation and maybe with static cernlib) and pandaroot. cmbsoft depends on "nothing" and pandaroot on "cbmsoft". That's it.

Johan.
