

Installation of PALM on the FLOW cluster at Carl von Ossietzky University of Oldenburg

Björn Witha
July 11th, 2013

Introductory comments

- Please have as well a look at the “official” PALM documentation, especially at the section “Installation”:

<http://palm.muk.uni-hannover.de/>

- So please use both sources of information in parallel, when you are trying to install PALM on FLOW.
- There are two methods to install and run PALM:
 - The ***advanced method*** is the **standard method for installing and running PALM**. It allows using the full functionality of PALM. The installation procedure is using the script `mbuild`. All jobs are submitted from your local computer to FLOW using the script `mr`.
 - The ***simple method*** is only using two scripts: `palm_simple_install` and `palm_simple_run`. It has several restrictions compared to the *advanced (standard) method* and is **only recommended for users who like to do a quick test of PALM!**
- You need a permit to access the PALM repository. If you have no permit so far, please contact the PALM group (raasch@muk.uni-hannover.de) and define a user name under which you like to access the repository. You will then receive a password which allows access under this name.
- In the following, **only the standard method will be described**

Requirements on your local computer

- Linux workstation
- A bash or Korn-(ksh)shell
- A netCDF-library with version number of 3.6.3 or higher (packages *libnetcdf* and *libnetcdf-dev*) (preferably version 4.x - if netCDF 3.x is used, the configuration file `.mrun.config` has to be modified)
- A Fortran90/95 compiler (packages: *gfortran* and *gcc*)
- MPI (Message Passing Interface, packages: *openmpi-common*, *libopenmpi-dev*, *libopenmpi*)
- The revision control system *subversion*

Create passwordless login (1)

A passwordless login between your local computer and FLOW (in both directions!) must be established to ensure a proper functioning of PALM. Test this by connecting to FLOW with:

```
ssh <FLOW-username>@flow.hpc.uni-oldenburg.de
```

and then on FLOW:

```
ssh <username on local computer>@<local IP-address>
```

In both cases you should not be asked for a password.

If a password is required, please follow the following procedure to enable public-key authentication for connections from your local computer to FLOW:

If not existing create a .ssh directory **on both your local computer and FLOW:**

```
mkdir ~/.ssh  
chmod go-rwx ~/.ssh
```

On your local computer, create the key:

```
cd ~/.ssh  
ssh-keygen -b 2048 -t dsa (type <return> when asked for password!)
```

Create passwordless login (2)

Copy the public key to FLOW:

```
scp id_dsa_pub <flow-username>@flow.hpc.uni-  
oldenburg.de:~/.ssh/authorized_keys
```

Log in to FLOW and create the key:

```
cd ~/.ssh  
ssh-keygen -t dsa          (type <return> when asked for password!)
```

Copy the public key to your local computer:

```
scp id_dsa_pub <username on local computer>@<local IP-  
address>:~/.ssh/authorized_keys
```

Append the public key to the authorized_keys file on FLOW:

```
cat id_dsa.pub >> authorized_keys
```

Installation of PALM (1)

Open a shell on both your local computer and FLOW and **create the following set of directories (on your local computer AND FLOW!)**:

```
~/job_queue  
~/palm  
~/palm/current_version  
~/palm/current_version/JOBS
```

The directory `~/palm/current_version` will be called the **working directory** from now on (in principal, the directory name can be freely chosen, but it is recommended to chose “current_version” as a different name will require several modifications in the scripts).

Installation of PALM (2)

Change into the working directory (on your local computer!) and **check out a working copy of the most recent PALM version** (currently as of 11-July-2013 revision 1204):

```
cd ~/palm/current_version
svn checkout --username <your username> /
                svn://130.75.105.7/palm/trunk trunk
```

Replace <your username> with the name you have chosen to access the PALM repository.

You will be prompted for your password. After completion, a subdirectory `trunk` will appear in your working directory. It contains a number of further subdirectories, which contain e.g. the PALM source code (`SOURCE`) and the scripts for running PALM (`SCRIPTS`).

The most recent version may still contain bugs. However, it is recommended as fixed releases are released only every 1-2 years. If you need to install an older (fixed) PALM release, please see the PALM homepage for detailed instructions. You will then have to adjust some scripts manually so that PALM can run on FLOW.

Installation of PALM (3)

To use the PALM scripts, the PATH-variable has to be extended and the environment variable `PALM_BIN` has to be set in your `.bashrc` or `.profile` file (located in your home directory – if not, create a new file starting with the line “# `.bashrc -l`”) by **adding the following two lines (on both your local computer AND FLOW!)**:

```
export PALM_BIN=$HOME/palm/current_version/trunk/SCRIPTS
export PATH=$PALM_BIN:$PATH
```

Now, the name and IP address of your local computer have to be added to the list of hosts in the script `trunk/SCRIPTS/subjob` (in the upper part of the script):

```
# HOSTSPEZIFISCHE VARIABLEN VEREINBAREN BZW. PRUEFEN, OB LOKALER HOST
# UEBERHAUPT ZULAESSIG IST
# Note: One of the entries for "lck" or "lckordi" always has to be
# comment out, because the hostname (node*) is same for both machines
case $local_host in
  (ambiel-lx)          local_addres=134.106.74.48;          local_host=lcfor;;
  (atmos)             local_addres=172.20.25.35;          local_host=lcide;;
  . . .
  (<name of local computer>) local addres=<local IP-address>; local_host=lcfor;;
```


Installation of PALM (4)

To install and run PALM, a configuration file `.mrun.config` is required. There are some default versions in the directory `trunk/SCRIPTS`. To use PALM on FLOW please **copy the following file to your working directory**:

```
cd /palm/current_version  
cp trunk/SCRIPTS/.mrun.config.forwind .mrun.config
```

The configuration file contains all information for compiling the PALM code and for generating and running the PALM executable using MPI.

Please edit this file by inserting your user names on FLOW and on your local computer, as specified in the comments in the file.

Warning: When editing the configuration file, please NEVER use the TAB key. Otherwise, very confusing errors may occur when mrun is executing.

When finished, please log out and then log in again to activate the settings.

Installation of PALM (5)

As a next step, the script **mbuild** has to be run to generate the executables for the local computer and FLOW and to copy the PALM scripts to FLOW:

```
cd /palm/current_version  
mbuild -u -h lcfor  
mbuild -u -h lcflow
```

Please answer all queries with 'y'. If the executables have been generated successfully, the message “*** mbuild finished” will be displayed.

To avoid the re-compilation of the complete source code for each model run, **PALM will be pre-compiled** once on FLOW by again using the script `mbuild`. The subroutines have to be compiled in a certain order. Therefore the so-called make mechanism is used, requiring a Makefile, in which the dependencies are described. This file is found in the subdirectory `trunk/SOURCE` where also the PALM code is stored. The compiled sources (object file `*.o`) are stored in the default directory `palm/current_version/MAKE_DEPOSITORY_parallel`. Please type:

```
mbuild -h lcflow
```

Again, all queries have to be answered with 'y' (the compilation is done for each block in your `.mrun.config` separately. By default, there is only one block: 'lcflow parallel').

The compilation may take several minutes.

If the compilation has been successful, „*** mbuild finished“ will be displayed.

Verification of the installation (1)

Please log out and then log in again.

As a last step, after the compilation has been finished, the PALM installation has to be verified. For this purpose, **a simple test run is carried out**. Please create a job directory for the example run (on your local computer) and copy the parameter file to this directory:

```
cd palm/current_version
mkdir -p JOBS/example_cb1/INPUT
cp trunk/INSTALL/example_cb1_p3d \
    JOBS/example_cb1/INPUT/example_cb1_p3d
```

Submit the example run with the following command (in one line):

```
mrun -z -d example_cb1 -h lcflow -K parallel -X8 -t 600
-r „d3# pr# ts#“
```

Verification of the installation (2)

By typing `qstat` (after having logged in on FLOW) you can check if your job is actually running. If the job does not appear in the list, it has already finished (the test run will be finished in much less than a minute).

When the example job has finished, a job protocol should have appeared in `~/job_queue` (on your local computer), named `lcf_flow_example_cb1`. Please check the content of this file carefully for any error messages (if a job is aborted, you will find error message in the job protocol). Note: it may happen that the job protocol is not transferred back to your local computer – in this case you will find it in `~/job_queue` on FLOW.

Beside the job protocol you will find three monitoring files in the following folder:

```
~/palm/current_version/JOBS/example_cb1/MONITORING
```

<code>example_cb1_cpu</code>	contains the CPU time measures
<code>example_cb1_header</code>	contains general information about the run and the selected parameters
<code>example_cb1_rc</code>	contains the content of <code>example_cb1_header</code> plus the values of certain control parameters at certain temporal intervals

Verification of the installation (3)

The netCDF output files are stored by default in the directory (on FLOW):

```
/data/work/gpfs/fw/<user-  
id>/palm/current_version/JOBS/example_cb1/OUTPUT
```

To verify the results of this example run, please compare the content of the run control file with the example result file:

```
diff JOBS/example_cb1/MONITORING/lcflow_example_cb1_rc \  
trunk/INSTALL/example_cb1_rc
```

You should not find any difference between these two files, except for the run date and time and the PALM revision number displayed at the top of the file header.

If the remaining file contents are identical, the installation is successfully completed!

Installation of new versions / updates (1)

The PALM group announces code revisions by emails sent to the PALM mailing list. If you like to be put on this list, just send an email to raasch@muk.uni-hannover.de. Details about all releases and major revisions can be found in the PALM change log:

<http://palm.muk.uni-hannover.de/wiki/doc/tec/changelog>

There are two possibilities to install new (or different) versions of PALM:

If you have previously checked out the most recent PALM code as recommended above, you can easily **update PALM to the newest version** by changing to the working directory and executing:

```
svn update trunk
```

This updates all files in the subdirectory trunk. The update may fail due to the subversion rules, if you have modified the contents of trunk. In case of any conflicts with the repository, please refer to the subversion documentation on how to remove them. In order to avoid such conflicts, modifications of the default PALM code should be omitted and restricted to the user-interface only (see PALM documentation).

Installation of new versions / updates (2)

Alternatively, you can install new or other releases in a different directory, eg.:

```
mkdir ~/palm/release-3.1c
cd ~/palm/release-3.1c
svn checkout --username <your username> /
  svn://130.75.105.7/palm/tags/release-3.1c trunk
```

You can also check out specific revisions:

```
mkdir ~/palm/revision_xxxx
cd ~/palm/revision_xxxx
svn checkout --username <your username> /
  svn://130.75.105.7/palm/trunk -r <revision number> trunk
```

After updating, please check for any differences between your current `.mrun.config` file and the default configuration file `trunk/SCRIPTS/.mrun.config.forwind` and adjust your file if necessary (note: the paths in `.mrun.config` and `.profile` or `.bashrc` have to be adjusted, as the working directory is no longer `current_version!`).

The scripts and the pre-compiled code must then be updated via:

```
mbuild -u -h lcfor
mbuild -u -h lcflow
mbuild -h lcflow
```

Quick installation guide (checklist) (1)

1. Create directories (on local computer and FLOW):

```
~/job_queue  
~/palm  
~/palm/current_version  
~/palm/current_version/JOBS
```

2. Check out PALM code (on local computer):

```
cd ~/palm/current_version  
svn checkout --username <your username> /  
          svn://130.75.105.7/palm/trunk trunk
```

3. Set environment variables on local computer and FLOW (in .bashrc or .profile):

```
export PALM_BIN=$HOME/palm/current_version/trunk/SCRIPTS  
export PATH=$PALM_BIN:$PATH
```

4. Add name and IP-address of your local PC in the script trunk/SCRIPTS/subjob:

```
(ambiel-lx)      local_addres=134.106.74.48;      local_host=lcfor;;  
(atmos)         local_addres=172.20.25.35;      local_host=lcide;;  
(<name of local computer>) local_addres=<local IP-address>; local_host=lcfor;;
```


Quick installation guide (checklist) (2)

5. Copy the configuration file to working directory and edit it:

```
cd /palm/current_version  
cp trunk/SCRIPTS/.mrun.config.forwind .mrun.config
```

6. Log out and then log in again

7. Generate executables and precompile the code (afterwards: logout and login):

```
mbuild -u -h lcfor  
mbuild -u -h lcflow  
mbuild -h lcflow
```

8. Verify the installation by carrying out the test run and comparing the results:

```
cd palm/current_version  
mkdir -p JOBS/example_cbl/INPUT  
cp trunk/INSTALL/example_cbl_p3d \  
    JOBS/example_cbl/INPUT/example_cbl_p3d
```

```
mrun -z -d example_cbl -h lcflow -K parallel -X8 -t 600 -r „d3#  
pr# ts#“
```

```
diff JOBS/example_cbl/MONITORING/example_cbl_rc \  
    trunk/INSTALL/example_cbl_rc
```