

Introduction to PDC's environment

Elisabet Molin

elimo@pdc.kth.se

PDC Center for
High Performance Computing



KTH, Sweden

What to learn?

- How to login
- Where to store things
- How to run a program



PDC Center for
High Performance Computing

Don't!

- NEVER write your PDC password on a computer which you are not able to touch!

- Never disclose your PDC password. Keep it secret. Keep it safe.
- User account details? Log in information? Don't e-mail your password!
It's yours, all yours!
- Never let anyone else use your PDC account. No, not even me!



Finding Information

- man, -help, -h, apropos ...
- <http://www.pdc.kth.se/education/tutorials/summer-school>
- Flash News
- Mailinglists
- Course assistants
- support@pdc.kth.se



Modules

- Modification of the user's environment; handles PATH, MANPATH . . .

```
module list  
module show modulename  
module avail  
module load modulename  
module swap fromModulename toModulename
```



PDC Center for
High Performance Computing

Common PDC modules are: PrgEnv-pgi, afsws and heimdal. Essential modules are preloaded at login.

Cray compilers and libraries

- Always use the wrappers `cc` (C code), `cc (C++)` or `ftn` (FORTRAN).
- The PrgEnv module loaded (`PrgEnv-pgi`, `PrgEnv-intel` or `PrgEnv-gnu`), decides if the PGI (default), Intel or GNU compiler is used by `cc`, `CC` and `ftn`.

```
module swap PrgEnv-pgi PrgEnv-gnu
```

- The wrappers automatically links with math libraries if their modules are loaded (no flags for linking required).

```
module load xt-libsci fftw
```

- to link your code against `lapack`, `blas`, `scalapack`, `blacs`, `fftw` . . .



Kerberos



Now this Cerberus had three heads of dogs, the tail of a dragon, and on his back the heads of all sorts of snakes.

- Kerberos — system for **authenticating users and services on a network.**
- Kerberos server, trusted by users and services.
- A **Kerberos principal** (`username@NADA.KTH.SE`) is a user's or service's username for a certain **Kerberos realm** (`NADA.KTH.SE`).



PDC Center for
High Performance Computing

Getting your Kerberos ticket

- kinit — proves your identity
- klist — list your Kerberos tickets
- kdestroy — destroy your Kerberos ticket file
- kpasswd — change your Kerberos password

```
kinit -f -l 7d username@NADA.KTH.SE
```

```
klist -Tf
```

```
Credentials cache : FILE:/tmp/krb5cc_500
Principal: username@NADA.KTH.SE
Issued Expires Flags Principal
Mar 25 09:45 Mar 25 19:45 FI krbtgt/NADA.KTH.SE@NADA.KTH.SE
Mar 25 09:45 Mar 25 19:45 afs/pdc.kth.se@NADA.KTH.SE
```



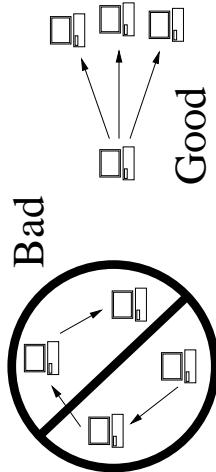
Important!

REMEMBER! You are NOT allowed to write your password on any computer that you are remotely connected to!!! If you expose your password you endanger not only your own work, you are putting all other cluster users at risk.

- Once you have your local **Kerberos ticket** you **never** need to type your password again!
- Get **forwardable** Kerberos tickets (pass the **-f** flag to `kinit`).
- **Kerberos enabled software** knows to forward/use local Kerberos tickets for remote authentication.
- **Forward your forwardable tickets when opening a remote connection.**



PDC Center for
High Performance Computing



Login using Kerberos tickets

- Get a **forwardable** Kerberos ticket on your **local system**:

```
kinit -f -1 7d username@NADA.KTH.SE
```

- **Forward** your forwardable tickets for remote login and secure X-windows redirect:

```
ssh -Y -l username summer-crash.pdc.kth.se
```

- The ssh command has to use **GSSAPI key exchange**.

- Test your ssh using:

```
ssh -vvv -Y -o GSSAPIKeyExchange=yes -o \
GSSAPIAuthentication=yes -l username \
summer-crash.pdc.kth.se
```

- The prompt changes when logged in. Check your tickets on the **remote system**:

```
klist -Tf
```

- **The lifetime of your tickets must exceed the runtime of the program you intend to run!**



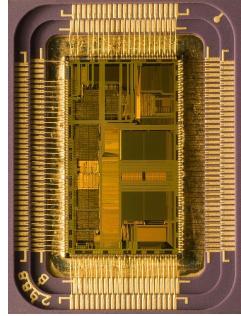
Login from any computer

- You can reach PDC from any computer and network!
- The **Kerberos implementation Heimdal** can be installed on most operating systems (Mac, Windows, Linux)
- An SSH command that knows **GSSAPI key-exchange** can forward Kerberos tickets
- Names of the commands depend on your operating system!
- Follow the instructions for your operating system:
<http://www.pdc.kth.se/resources/software/login-1>



What's a computer cluster?

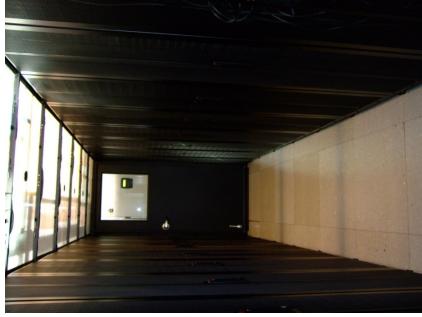
CPU/core



Rack

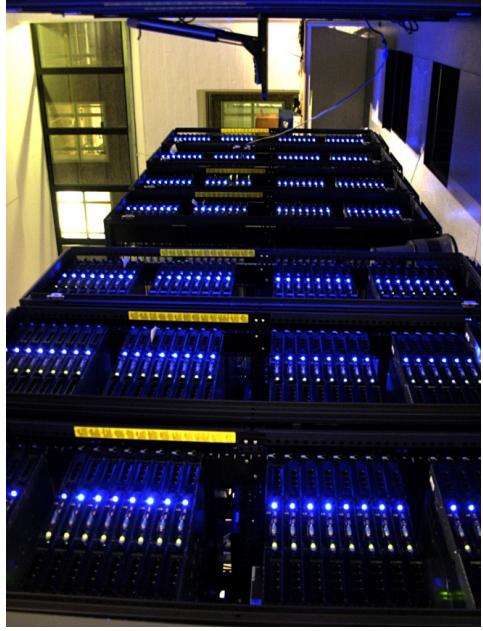


Cluster



Where to run your programs

- Login nodes — **don't run jobs here!**
- Interactive nodes — for test runs; alone or shared among users
- Dedicated nodes (compute nodes) — for running final programs



The Summer Cray (XE6)

- Login node
summer-cray.pdc.kth.se
- 22 interactive nodes — each has 16 cores, 32 GB RAM
- 0 dedicated nodes (i.e. no queue)
- Gemini interconnect



PDC Center for
High Performance Computing



Parallel file system

- Lustre — massively parallel distributed file system
 - **No backup**
 - No personal quota — move data off!
- /cfs/klemming/nobackup/u/username
- where username is your username and u is the first letter in your username.
- /cfs/klemming/scratch/u/username
- **Always start and run your programs in Lustre!**



Your home directory

- Your **home directory** is located in **AFS**.
- When you login in at PDC you arrive in your AFS home directory, `~`.
- Your home directory on the lab computers for the exercises is **the same**.
- The path to your PDC AFS home directory is
`/afs/pdc.kth.se/home/u/user` where `username` is your
username and `u` is the first letter in your username.
- Some pre-existing **files and folders to keep**:
`.bashrc .forward Public Private ...`
- **Oldfiles** contain `~`, as it was yesterday, i.e. a **backup**.
- **Save your code** in AFS to get backup
- **You can not run jobs from AFS on Lindgren**
- You can access AFS from any computer
 - `http://www.pdc.kth.se/resources/software/file-transfer/`



Serial job on summer-Cray

- Get a Kerberos ticket on your local computer
- Login on summer-cray.pdc.kth.se
- Check the lifetime of your ticket
- Go to your directory:

/cfs/klemming/nobackup/u/username

- Compile:
ftn hello_serial.f90 -o hello_serial_f90
cc hello_serial.c -o hello_serial_c
 - Run:
aprun -n 1 ./hello_serial_f90
aprun -n 1 ./hello_serial_c
- **ALWAYS use aprun before running a program!**



MPI job on summer-Cray

- Get a Kerberos ticket on your local computer
- Login on summer-cray.pdc.kth.se
- Check the lifetime of your ticket
- Go to your directory /cfs/klemming/nobackup/u/username
- Compile:

```
ftn hello_mpi.f90 -o hello_mpi_f90  
cc hello_mpi.c -o hello_mpi_c
```
- Run:

```
aprun -n 16 ./hello_mpi_f90  
aprun -n 16 ./hello_mpi_c
```
- **ALWAYS use aprun before running a program!**



PDC Center for
High Performance Computing

OpenMP job on summer-cray

- Get a Kerberos ticket on your local computer
- Login on summer-cray.pdc.kth.se
- Check the lifetime of your ticket
- Go to your directory /cfs/klemming/nobackup/u/username
- Compile:

```
ftn hello_omp.f90 -mp -o hello_omp_f90
cc hello_omp.c -mp -o hello_omp_c
```
- Run:

```
export OMP_NUM_THREADS=16
aprun -n 1 -d 16 ./hello_omp_f90
export OMP_NUM_THREADS=16
aprun -n 1 -d 16 ./hello_omp_c
```
- **ALWAYS use aprun before running a program!**



PDC Center for
High Performance Computing

Lindgren

- System used after the summer school!
- Cray XE6
- Login node lindgren.pdc.kth.se
- 8 interactive nodes
- 1516 dedicated nodes (queue needed!)
- 24 cores ($4 * 6$ cores, 32 GB RAM per node)
- Gemini interconnect



PDC Center for
High Performance Computing



When does my job start?

- Many users continuously submit jobs to Lindgren
- Each job needs a certain **time** and a number of **nodes/cores** to run
- The **queue system (scheduler)** keeps track of the submitted jobs and when, and where, each job can run on the cluster
- Every user belongs to at least one **time allocation** on a cluster. You belong to the time allocation **summer-2011**
- A time allocation states **how many node hours per month you have on a specific cluster.**



Running jobs on Lindgren

- After the Summer School, summer-cray will be gone
- Submit your jobs to the queue on Lindgren
- Get Kerberos tickets on your own computer
- Login on lindgren.pdc.kth.se
- Check the lifetime of your ticket
- Go to your directory /cfs/klemming/nobackup/u/username
- Compile:

```
ftn hello_mpi.f90 -o hello_mpi_f90
cc hello_mpi.c -o hello_mpi_c
```
- Submit your job:

```
qsub ./mpi_job.pbs
```
- mpi_job.pbs is a **job script** that the queue can understand.



PDC Center for
High Performance Computing

Job script

```
#PBS -N hello_mpi_job
# 5 minute wall-clock time will be given to this job
#PBS -l walltime=5:00
# Number of MPI tasks. Should always be a multiple of 24
# i.e. an entire node on Lindgren
#PBS -l mppwidth=24
# Change to the work directory
echo $PBS_O_WORKDIR
cd $PBS_O_WORKDIR
# Run the program, always use aprun!
aprun -n 12 ./hello_mpi
```

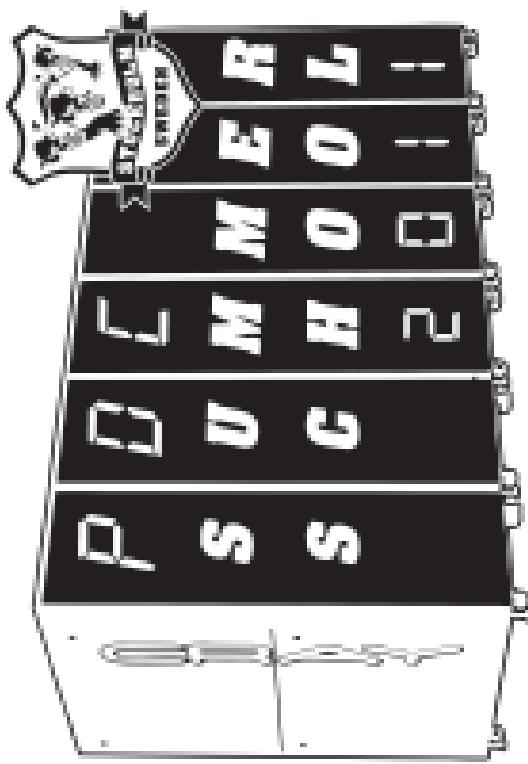
- If you forget **aprun** in the script, the program will start on the login node!



PDC Center for
High Performance Computing

Thank you for listening!

Introduction to High-Performance Computing



To 0x7FFF000000000000 and beyond!



PDC Center for
High Performance Computing