

HP XC4000: Brief Information for new Users

University Karlsruhe (TH), Computing Center

Version 1.0
January 11, 2007

This publication shall facilitate the beginning of work with HP XC4000 at Computing Center of University Karlsruhe. The HP XC4000 system is a distributed memory parallel computer with 750 nodes where each node has two or more AMD Opteron sockets with 2.6 GHz frequency, 16 GB local memory, local disks and network adapters. Each socket comprises 2 cores. Special file server nodes are added to HP XC4000 to provide a fast and scalable parallel file system. All nodes are connected by an InfiniBand 4X DDR interconnect.

If you run into troubles when working on HP XC4000 you should not be afraid to contact the XC-hotline (see 8) or responsible staff members of the Computing Center.




Figure 1: Configuration of HP XC4000

1 Login

The HP XC4000 system is a distributed memory parallel computer with two dedicated login nodes. Both login nodes are equipped with 8 cores, 32 GB main memory, local disks and network adapters. The operating system Linux (Red Hat AS 4.0, kernel version 2.6.x) runs on both nodes (and all other nodes), so that working on a single node of HP XC4000 is comparable with working on a HP workstation.

1.1 Login on a Login Node

There are 2 login nodes on HP XC4000 each with 8 cores. The selection of the login node is done automatically. If you are connecting another time to a login node, the sessions might run on different (login) nodes of HP XC4000. Only the secure shell `ssh` is allowed to login. Other commands like `telnet` or `rlogin` are not allowed for security reasons.

For commercial customers a connection to HP XC4000 can be established by the command
`ssh user-id@hwwxc2.hww.de`

For non-commercial customers a connection to HP XC4000 can be established by the command
`ssh user-id@xc2.rz.uni-karlsruhe.de`

When logging in you must type your password; when logging in the first time the password should be changed.

If you are using OpenSSH (usually installed on Linux based systems) and you want to use a GUI-based application on HP XC4000 like e.g. the debugger DDT, you should use the command

```
ssh -X user-id@xc2.rz.uni-karlsruhe.de
```

with the option `-X`.

2 Copying Data onto HP XC4000

HP XC4000 uses the parallel filesystem Lustre for the HOME-directory (and also for the WORK-directory. For security reasons Lustre can not be exported to your local workstation, i.e. all your data must be copied from HP XC6000 (or your workstation) to HP XC4000.

To achieve a high security level the files must be transferred by the secure copy command `scp` or the command `rsync` that again uses the command `scp`.

The command `rsync` is used similar to the command `scp`. In the following recommendation the HOME-directory (WORK-directory) and all subdirectories of the computer you are logged in are copied recursively to the user typed instead of <userid> on HP XC4000 (XC2) into the directory `xc1_home` (`xc1_work` of the WORK-directory).

```
rsync -av $HOME/ <userid>@xc2:xc1_home  
rsync -av $WORK/ <userid>@xc2:\$WORK/xc1_work
```

To guarantee consistent data neither batch jobs on HP XC6000 nor batch jobs on HP XC4000 should run while the data are copied.

Further informations on the command `rsync` are available by manual pages (`man rsync`).

3 Choosing your User Environment

HP XC4000 supports the use of Modules software to make it easier to configure and modify the user environment. Modules software enables dynamic modification of your environment by the use of modulefiles. A modulefile contains information to configure the shell for an application.

We provide modules for all supported compilers, debuggers, tools and libraries. Modulefiles that are available on HP XC4000 can be viewed by the command `module avail[able]`. To view the modulefiles that are currently loaded in your environment, issue the command `module list`. By default the modulefiles `dot` adding the current directory to your environment variable `PATH`, `intel` loading the Intel C/C++ and Fortran compilers in the latest stable version and `hp-mpi` loading HP MPI in the latest stable version are made available. You can load a modulefile in to your environment to enable easier access to software that you want to use by executing the command `module add <modulefile>`. To unload a modulefile that is currently loaded the command `module rm <modulefile>` must be executed.

If e.g. you want to use the PGI compiler instead of the Intel compiler (this is the default compiler), you only must type `module add pgi`. The command `module rm intel` can be omitted because unloading of an arbitrary loaded compiler will be done automatically within our Module environment when loading another compiler.

4 Compiler and Parallelization Environment

On HP XC4000 different compilers are available for the programming languages Fortran90/95, High Performance Fortran, C and C++. The most important parallelization environments are the Message Passing Interface MPI for jobs running on many nodes and OpenMP for jobs running within one node.

The default compiler is the Intel compiler; its call is `ifort` for Fortran programs, `icc` for C programs and `icpc` for C++ programs. Other compilers from PGI, PathScale and GNU can also be chosen for both Fortran and C/C++; exchanging of the compiler can simply be done by the command `module add {pgi|pathscale|gcc}`.

For the compilation of message passing applications with MPI there are special compiler scripts like `mpif77` for Fortran77 programs, `mpif90` for Fortran95 programs, `mpicc` for C programs and `mpiCC` for C++ programs.

Further informations on compilers and the parallelization environment can be extracted from the HP XC4000 User Guide and the online manuals¹.

5 Batch System SLURM

A few nodes are provided for interactive operations. All other nodes can only be accessed by the batch system SLURM. Batch jobs are submitted by the command `job_submit`. A required parameter is the name of the executable for jobs running on a single processor or rather "`mpirun name_of_executable`" for jobs running on more than one processor or a shellscript; required options are the required resources like number of processors, CPU-time and main memory request. Informations on `job_submit` can be viewed by the command `job_submit -H`. One important parameter of the command `job_submit` is `-m <memory_request>` because it controls the used number of processors within one node. Setting up to 4000 (MB) as memory request your parallel application will use all (four) cores within one node; setting more than 4000 up to 8000 (MB) means that your application will only use two in four cores within one node; setting more than 8000 (MB) means that only one in four cores within one node will be used. In all cases all four cores of a node are accounted! Further actual limitations of the job queues can be displayed by the command `job_info`.

¹http://www.rz.uni-karlsruhe.de/ssck/xc4k_manuals

6 User Guide and Course

All manuals on HP XC4000 are available online either in PDF- or HTML-format. The substantial local adjustments on usage of HP XC4000 are outlined in **HP XC4000 User Guide**².

Together with the start of operation of HP XC4000 a introductory course takes place. All slides of this course will be made available online³.

7 Important Addresses for Online Informations

- **HP XC4000 Website at Computing Center of University Karlsruhe**
<http://www.rz.uni-karlsruhe.de/ssck/hpxc4000>
- **HP XC4000 User Guide**
<http://www.rz.uni-karlsruhe.de/rz/docs/HP-XC/ug/ug4k.pdf>
- **HP XC4000 Introductory Course**
http://www.rz.uni-karlsruhe.de/ssck/xc4k_training
- **Homepage of Scientific Supercomputing Center Karlsruhe (SSCK)**
<http://www.rz.uni-karlsruhe.de/ssck>
- **Homepage of Computing Center of University Karlsruhe**
<http://www.rz.uni-karlsruhe.de>
- **Maintenance and Repair Information**
<http://www.rz.uni-karlsruhe.de/rd/stoerung>

8 Contact

- **XC Hotline**
xc-hotline@uni-karlsruhe.de
Phone: +49 721 608-8011

²<http://www.rz.uni-karlsruhe.de/rz/docs/HP-XC/ug/ug4k.pdf>

³http://www.rz.uni-karlsruhe.de/ssck/xc4k_training