HPC2N

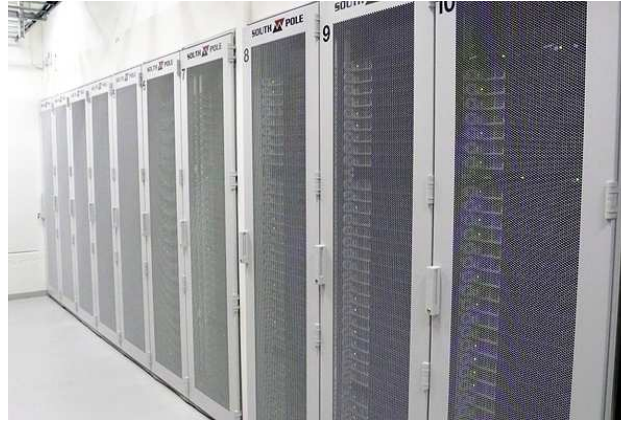## Our Systems and how to use them

Umeå
20 June 2016

Abisko

- 328 nodes/15744 cores

- 10 fat nodes (512 GB RAM), 318 thin nodes (128 GB RAM)

- CPUs: (thin) 4 x AMD Opteron 6238 (Interlagos) 12 core (2.6 GHz)

- CPUs: (fat) 4 x AMD Opteron 6344 (Abu Dhabi) 12 core (2.6 GHz)

- Interconnect: Infiniband QDR, 40Gb/s, Mellanox

- OS: Linux Ubuntu

- Installed 2011, upgraded with extra nodes 2014

Kebnekaise

- 13 racks, Mixed system:
- 432 Intel Broadwell, E5-2690v4 (2x14 cores/node), 128 GB/node
- 36 Intel Broadwell + NVidia K80 GPUs (32 with 2/node, 4 with 4/node)
- 20 Intel Haswell E7-4850v3, (4x14 cores/node), 3 TB/node (fat nodes)
- 36 Intel KNL: 7250 SKU (68 cores, 1.4/1.2 Ghz/AVX), 192 GB/node
- Interconnect: Mellanox Infiniband FDR
- OS: Linux Ubuntu
- Installation: Summer/Fall 2016

Hardware, Storage resources

- GPFS - fast parallel file system / center storage

  ◇ ~2PB

- Tape storage - Six IBM Ultrium 5 tape drives capable of

  ◇ storing 1500GB uncompressed data/tape at the transfer

  ◇ speed of 140MB/s. Total storage of 3PB

Available software

- **Application Software**: Abinit, Ansys, DDT, Espresso, Gamess, Gaussian, Gromacs, HDF5, Matlab, NetCDF, NWChem, Octave, PETSc, Python, R, Siesta, VASP, WRF ...

- **Compilers**: Intel, GCC, PGI, Pathscale

- **Num. and Comm. libraries**: BLACS, FFTW, BLAS, LAPACK, ScaLAPACK, ACML, Intel MKL, ParMETIS ...

- **MPI**: Intel MPI, OpenMPI

- Other software on request

Using our systems

- Register in SUPR (https://supr.snic.se/).
- Accept the SNIC User Agreement
- Apply for an account a HPC2N (https://supr.snic.se/account/)

  Note that you need to be a member of a project first, or

  have applied for a project
- Several users can share a project owned by one PI

  PI can add users in SUPR

Using our systems

- Most programs are accessed by first loading them as a 'module'
- See which modules exists

  `module avail`
- Different versions of software

  `module avail <module name>`
- Example: loading the default intel compilers

  `module load intel`
- Unload the module

  `module unload intel`

The module system will be different on Kebnekaise!

Using our systems

- Large/parallel programs: run through the 'batch system' (SLURM)
- Important! The batch system cannot access AFS (your home)
- Run everything from the parallel file system

  (http://www.hpc2n.umu.se/filesystems/overview)

The Batch System

- Keeps track of available system resources
- Takes care of scheduling jobs of multiple users, running tasks simultaneously
- Enforces local system resource usage and job scheduling policies
- Users submit to a queue (running, idle, blocked)
- Guides at: http://www.hpc2n.umu.se/support/

The Batch System - example

- Submitting: sbatch $< jobscript >$
- Job Script:

```
#!/bin/bash
#SBATCH -A SNICXXX-YY-ZZ
#SBATCH -n 48
#SBATCH --exclusive
#SBATCH --time=01:00:00


module add openmpi/gcc


srun ./my_parallel_program
```

The Batch System - Gromacs example (single node, all cores)

```bash
#!/bin/bash
#SBATCH -A SNIC2016-3-38
#SBATCH --reservation comp_chem
#SBATCH -n 48
#SBATCH --ntasks-per-node=48
#SBATCH --time=01:00:00


module add gromacs/5.1.2


if [ -n "$SLURM_CPUS_PER_TASK" ]; then
    mdargs="-ntomp $SLURM_CPUS_PER_TASK"
else
    mdargs="-ntomp 1"
fi


gmx mdrun -ntmpi $SLURM_NTASKS $mdargs
```
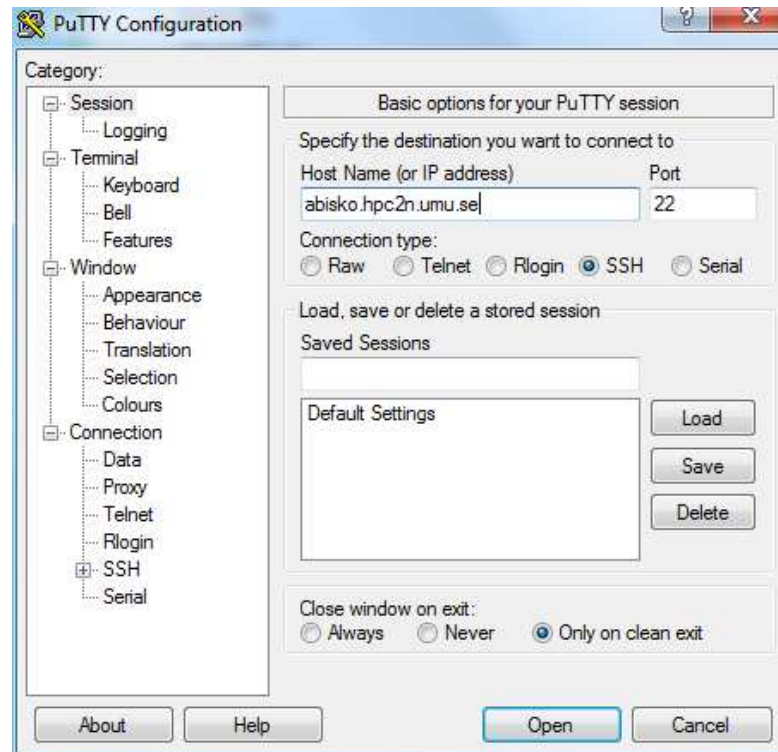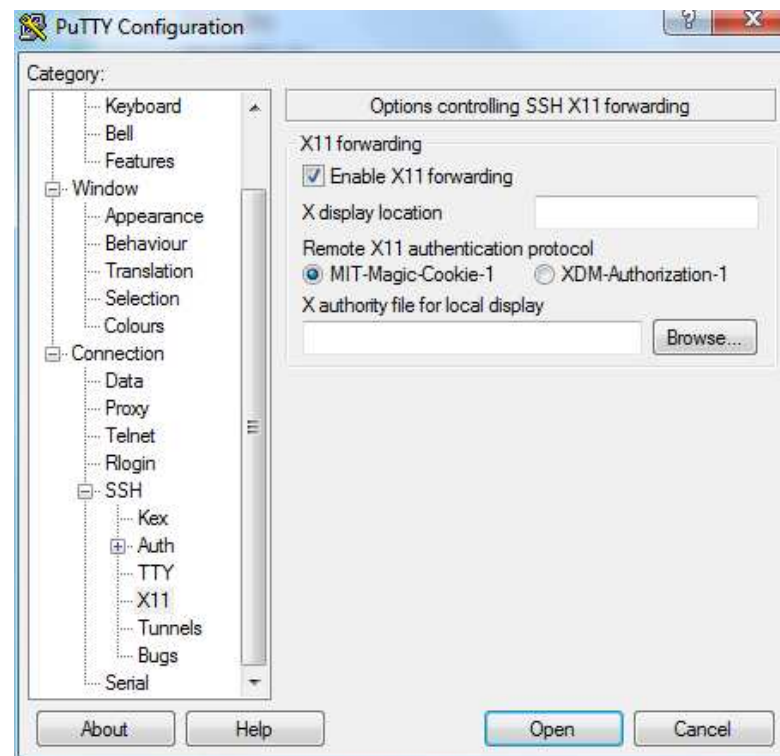
Connecting from a Windows System - PuTTY

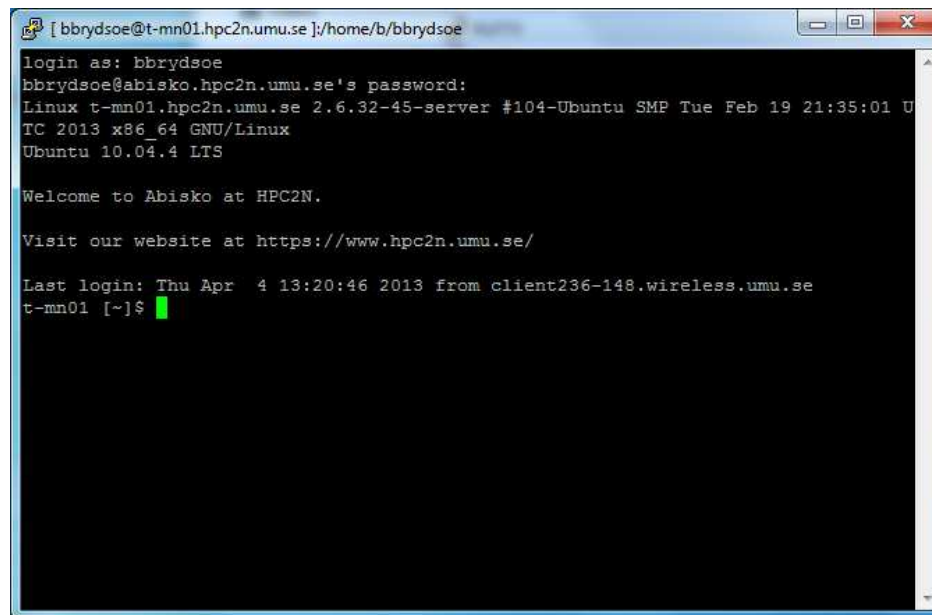• Get the Zip file with both PuTTY, PSCP, and PSFTP. Unzip, run putty.exe

Connecting from a Windows System - PuTTY

• Enabling X11 forwarding (start up the X11 server before connecting)

Connecting from a Windows System - PuTTY

- Logging on:

Connecting from a Windows System - Xming

- Download from the Xming page
  (http://www.straightrunning.com/XmingNotes) or directly from Sourceforge
- Install (default options should work)
- Launch Xming
- You can now launch (for instance) PuTTY SSH client and enable X11
  forwarding as shown earlier.

Questions?

- http://www.hpc2n.umu.se/
- support@hpc2n.umu.se