

e-Science's High Performance Computers (HPC) Instructions

(@NECTEC, Sybyl X, @Chula)

The High Performance Computers (HPC) of e-Science Infrastructure can be access through the Secure Shell Client (SSH) such as Putty and SSHClient for Window Operating System. For Linux and MacBook, the High Performance Computers can be access via the Terminal.

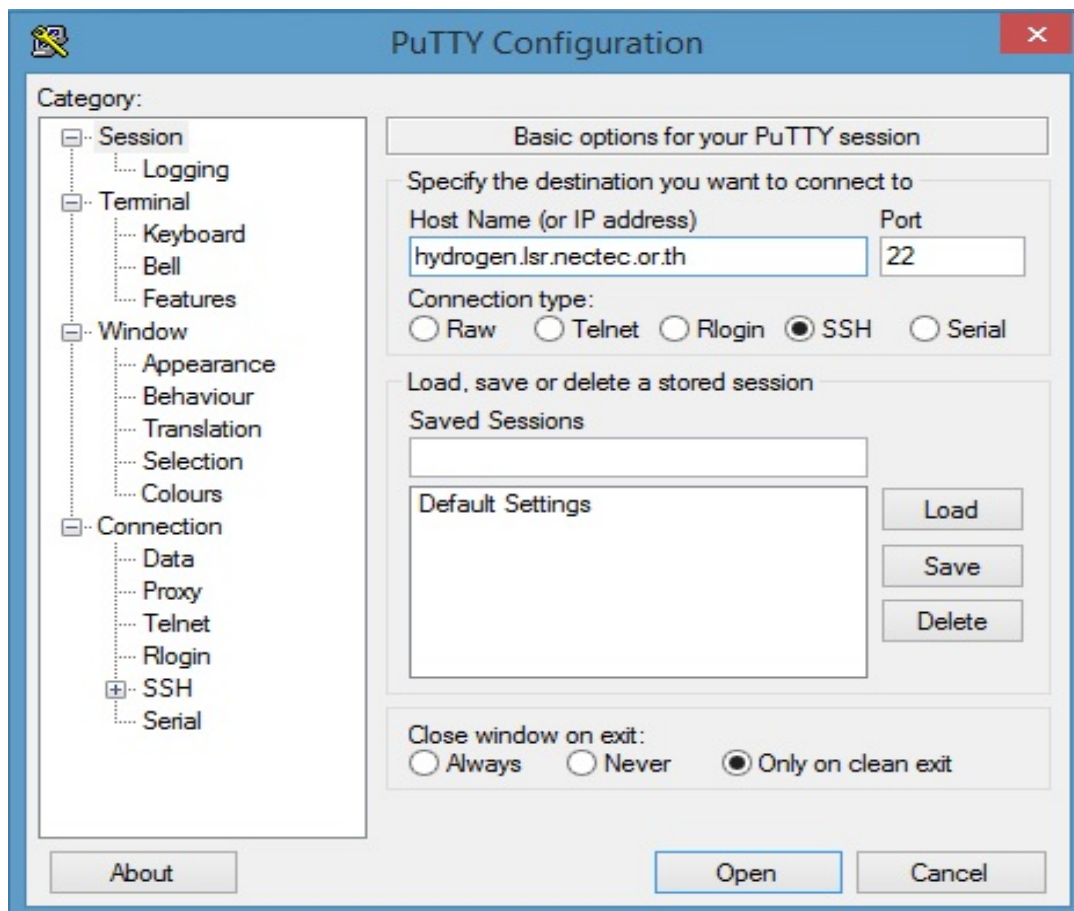
Window OS

Download links

- [Putty](#)
- [SSHClient](#)

Configurations:

Putty:



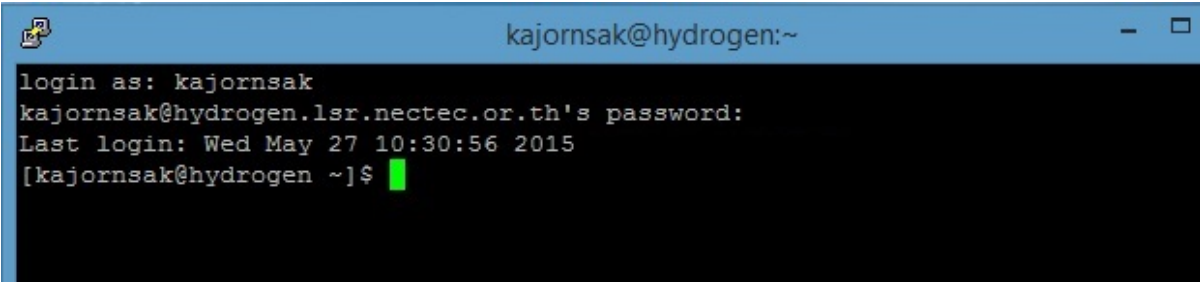
1. Fill in the “Host Name” or the IP address of the HPC Clusters that send by e-Science staff such as followings:

- hydrogen.lsr.nectec.or.th (HPC @NECTEC)
- silicon.lsr.nectec.or.th (for Sybyl X users)
- esci-cu.sc.chula.ac.th (HPC @Chula)

Selecting Connection Type “SSH” with port “22”. Then click “Open”.

Website : <https://twiki.cern.ch/twiki/bin/view/Main/ChulaeScienceManual>

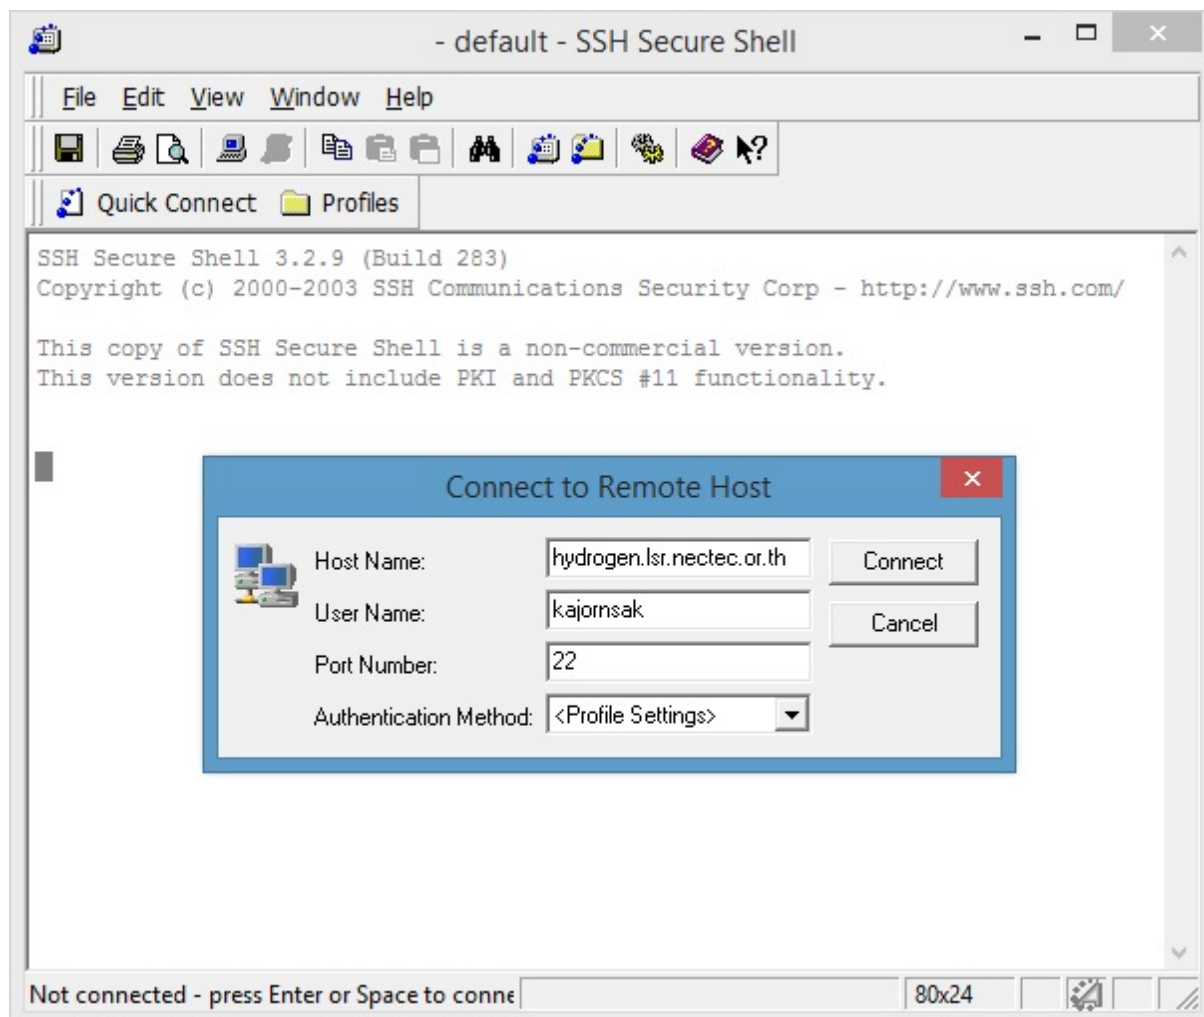
2. Putty will prompt for user log-in with username and initiate password send via email by e-Science staff.



```
kajornsak@hydrogen:~  
login as: kajornsak  
kajornsak@hydrogen.lsr.nectec.or.th's password:  
Last login: Wed May 27 10:30:56 2015  
[kajornsak@hydrogen ~]$
```

Screen captured of Putty after user log-in

SSHClient:

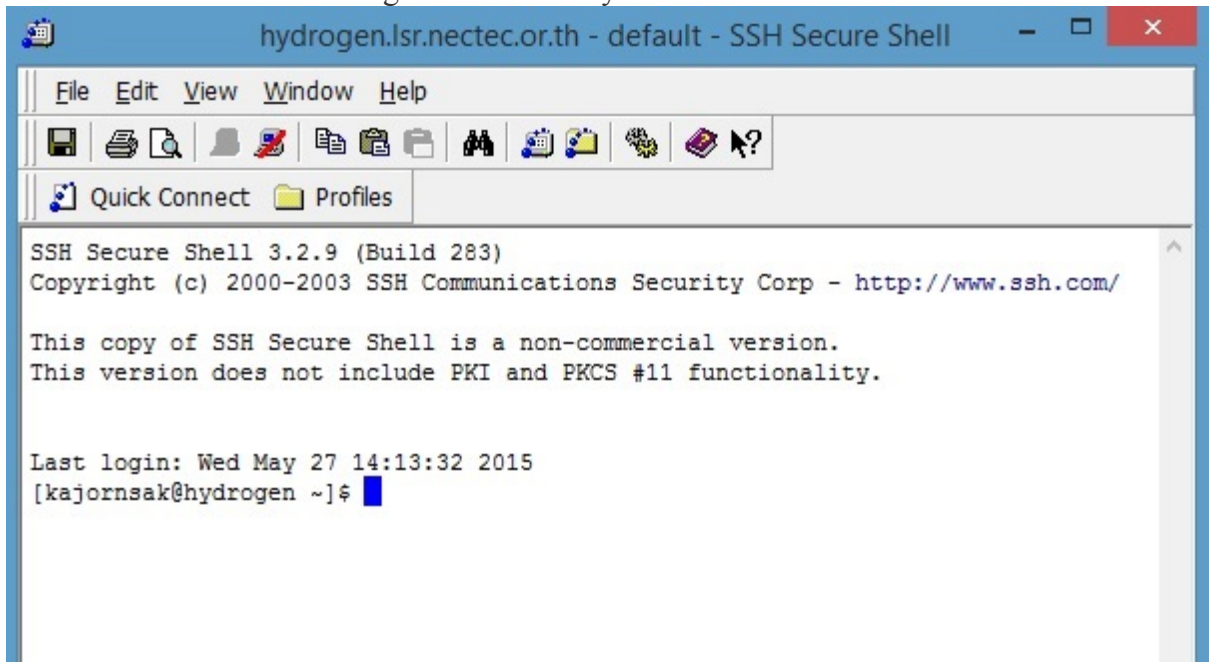


SSHClient's screen

1. Fill in the Host Name Fill in the "Host Name" or the IP address of the HPC Clusters that send by e-Science staff such as followings:
 - hydrogen.lsr.nectec.or.th (HPC @NECTEC)
 - silicon.lsr.nectec.or.th (for Sybyl X users)
 - esci-cu.sc.chula.ac.th (HPC @Chula)

Website : <https://twiki.cern.ch/twiki/bin/view/Main/ChulaeScienceManual>

2. Enter the user name according to email send by e-Science staff. Then click “Connect”



screen of SSHClient after Log-in

LINUX and MAC OS

User can access via SSH by open “Terminal”. Then using the following commands:

HPC Clusters	Terminal Commands
Hydrogen (HPC @NECTEC)	<code>ssh hydrogen.lsr.nectec.or.th</code>
Silicon (for Sybyl X user)	<code>ssh silicon.lsr.nectec.or.th</code>
CU eScience (HPC @Chula)	<code>ssh esci-cu.sc.chula.ac.th</code>

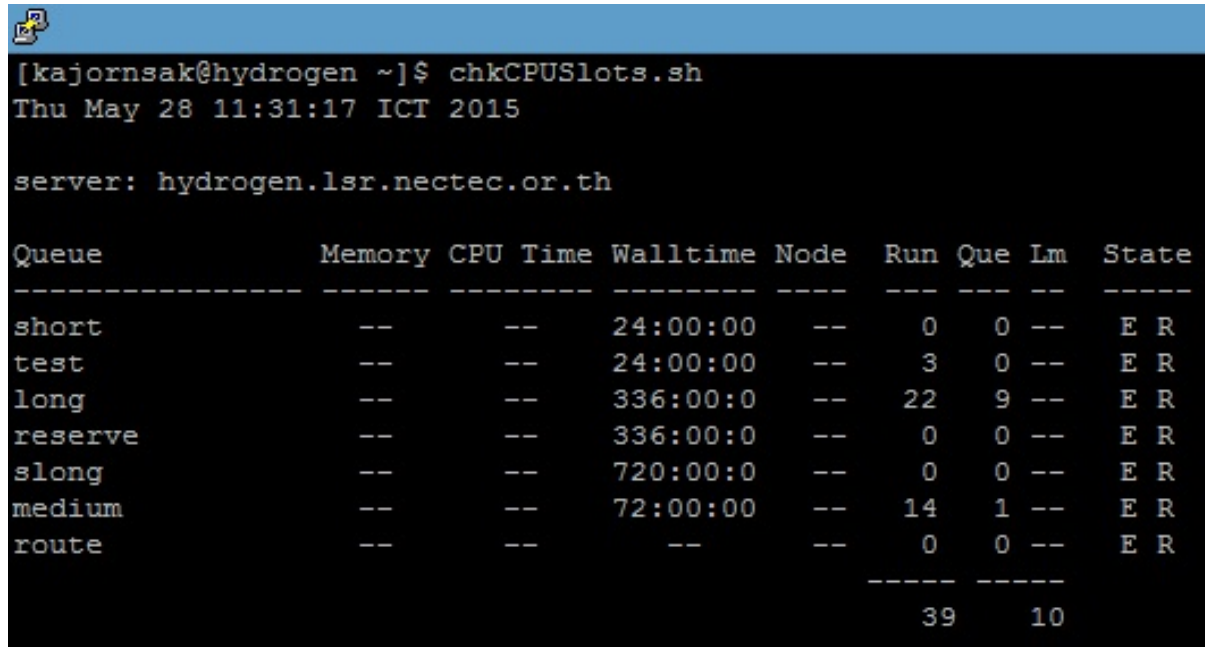
Website : <https://twiki.cern.ch/twiki/bin/view/Main/ChulaeScienceManual>

*Remarks:

- In order to access the HPC Clusters, users need to use the username and password that sent to you by email.
- Users need to change the password after the first log-in by using “*passwd*” command. The Terminal will ask for the old password and new password twice. The new password need to be minimum 6 characters and must be word that are not exist in the dictionary. The password need to be change every 3 months.

Checking Your job's status on HPC Cluster

To check your job's status that submitted on the Hydrogen cluster, type in the “*chkCPUSlots.sh*”. By using this command, the information such as Queue Type, CPU Time, Memory, etc. has shown in the picture below.



```
[kajornsak@hydrogen ~]$ chkCPUSlots.sh
Thu May 28 11:31:17 ICT 2015

server: hydrogen.lsr.nectec.or.th

Queue           Memory CPU Time  Walltime  Node   Run  Que  Lm  State
-----
short           --      --      24:00:00  --     0    0  --  E R
test            --      --      24:00:00  --     3    0  --  E R
long            --      --      336:00:0  --    22    9  --  E R
reserve         --      --      336:00:0  --     0    0  --  E R
slong           --      --      720:00:0  --     0    0  --  E R
medium          --      --      72:00:00  --    14    1  --  E R
route           --      --          --      --     0    0  --  E R
-----
                                39    10
```

Screen caption of resulting in using the “*chkCPUSlots.sh*” command.

Meaning of Columns

Queue = List all the queue type available on the cluster,

Walltime = Allowance duration for each queue type can run before the job will terminate. The display number on the screen present in format of hhh:mm:ss,

Run = The quantity of queues running for each queue type,

Que = The quantity of queues waits to be running of each queue type,

Queue Types

The e-Science HPC Service has aim to support on the computational calculation in various research areas and numerous users. Therefore, the specifications of nodes such as FIFO and Backfill are deploy to serve and gain the highest beneficial regarding to users' job.

Additionally, there are 4 types of queues adopted available for users as following:

1. “test” queue type is the shortest queue, which allow the submitted job to run for 1 day per job (4 cores per 1 node / 10 jobs per user)
2. “short” queue type allows the submitted job to be run for 1 day per job (256 cores for 8 nodes / 24 jobs per user)

3. “medium” queue type allows the submitted job to be run for 3 days per job (128 cores for 4 nodes / 12 jobs for user)
4. “long” queue type allows the submitted job to be run for 14 days per job (128 cores for 4 nodes / 6 jobs per user)
5. “slong” and “reserve” are the special queue type that require users to contact the e-Science staffs. These queue types are for research project that need large quantity of CPU Cores or need to run job that longer than 14 days.

Queue type	Time limit (Day)	Resources limit	Number of job limit
test	1	4 CPU cores on 1 node	10
short	1	256 CPU cores on 8 nodes	24
medium	3	128 CPU cores on 4 nodes	12
long	14	128 CPU cores on 4 nodes	6

***Remarks:**

The job that requires more than 7 days to run, we are advice to divide the job that suitable for queue types. However, user can contact for the system administrator if the job cannot be divide.

Please be advice that the number of CPU Core, the number of node, and the number of jobs allowance per user can be change. For any enquiry, please contact the system administrator/ e-Science HPC Staffs.

How to read the current resource.

Hostname is the name of the HPC cluster that serve users. Our main Computer Cluster server is called “Atom” server with Hostname: “hydrogen.lsr.nectec.or.th.” Our HPC Clusters may have different name, label, and capacity as present on the table below:

Atom Computer Cluster (Hostname: hydrogen.lsr.nectec.or.th)

Name	Label	#CPU Cores	Memory (GB)	#Nodes	Note
sodium	<i>old</i>	12	40	10	Old nodes

radium	<i>himem</i>	32	256	3	High Memory
osmium	<i>new</i>	32	128	8	New nodes
nobelium	<i>jumbo</i>	96	512	1	96 CPU Cores on one node
gadolinium	<i>cope</i>	12	48	1	Co-processor Xeon Phi

Check Cluster Status.

```

===== hydrogen.lsr.nectec.or.th =====
<Hostname>: <CPUs>          <Memory (M)>          <Status>
=====
sodium-0-0: 10/12           3041/40187           (free)
sodium-0-1: 12/12           3437/40187           (job-exclusive)
sodium-0-2: 9/12            2995/40187           (free)
radium-0-0: 12/32           18906/258145         (free)
radium-0-1: 16/32           7000/258145          (free)
radium-0-2: 16/32           117650/258145        (free)
osmium-0-0: 16/32           6919/128873          (free)
osmium-0-1: 32/32           83580/128873         (job-exclusive)
osmium-0-2: 32/32           4360/128745          (job-exclusive)
osmium-0-3: 16/32           22765/128873         (free)
osmium-0-4: 26/32           12359/128873         (free)
osmium-0-5: 32/32           23632/128873         (job-exclusive)
osmium-0-6: 20/32           23806/128873         (free)
osmium-0-7: 4/32            3564/128905          (free)
nobelium-0-0: 56/96         35014/516573         (free)
gadolinium-0-0: 12/12       13936/48236          (job-exclusive)
=====

```

The screen shot of CPU and Memory status available for users.

CPUs, Memory (M) display on the screen shot above has shown the status of CPU and Memory that recently use. The system has display the current use of resource along with the maximum number that the resource allows user to engage.

Status column on the right-hand side of the screen indicate the status of compute nodes. The compute nodes can be “free” or “job-exclusive”.

“free” means that the CPUs Core of compute node is available for users to submit the jobs.

“job-exclusive” means that the CPUs Core of the compute node are fully engaged and cannot engage anymore jobs.

Portable Batch System (PBS) Script

The PBS is a job scheduler mechanism that manage the submitted jobs. We are using TORQUE version of PBS Scheduler, which use the command of PBS Script to executing jobs. The PBS Script has divide into two parts: 1) information of the job, by provide the information on number of nodes, CPU Cores, Memory request for job execution; 2) job execution commands (Linux based command systax)

The important parameters in first part of PBS Script

```
#PBS -N Job_Name (define the name of the job)
#PBS -S /bin/bash (define the shell environment path i.e. bash shell)
#PBS -q test (define the queue types to execute the job. Queue Types are explained on
              the Queue Types section above)
#PBS -l nodes=2:ppn=2 (define the number of nodes and CPU Cores request to
execute the job on each node)
```

Example of PBS Script

Serial Job - the sample PBS Script below present the Serial Job execution that require to run the test queue. This queue type will execute no more than 1 day long.

```
##### PBS Part #####
#PBS -N Serial_Job
#PBS -S /bin/bash
#PBS -l nodes=1:ppn=1
#PBS -q test
##### End Part #####

cd ~/my_work_space
./program.exe > output.file
```

Parallel Job – the sample script below demonstrates the Parallel job execution, which need run with the Message Passing Interface (MPI) protocol such as “mpich”, “mpich2”, “openmpi”, and “mvapich2”. The script has specified to execute job on ABINIT program to compute the result for 3 days. The job request 32 CPU Cores of 4 nodes from New HPC resource group.

```
##### PBS Part #####
#PBS -N Parallel_Job
#PBS -S /bin/bash
#PBS -l nodes=4:ppn=8:new
```



```
#PBS -q medium
#### End Part ####

export MPI_EXE=/opt/openmpi-1.8.5/bin
export ABINIT_EXE=/opt/abinit-7.10.2/bin
cd ~/my_work_space
cat $PBS_NODEFILE > node_list.txt
$MPI_EXE/mpirun -np 32 -machinefile node_list.txt $ABINIT_EXE/abinit
< input.file > output.file
```

*Remarks:

More of the PBS script can be view by using “#PBS -l” command.

Apache Spark Job submission

The job for Big Data Analytics can be submit via PBS Queue as described above to execute on Apache Spark. However, the Spark Job environment require to be build according to match the number of resource request for the job execution on PBS Scheduler. This allow the Spark Job’s users to specify the number of HPC resource node, CPU Cores and Memory require to execute the job. Further reading on this matter can be read on <https://github.com/ekasitk/spark-on-hpc>.

PBS Job Submission

The job can submit by using the “qsub” command with the syntax below:

```
qsub test.pbs
```

```
840.hydrogen.lsr.nectec.or.th
```

840 is the Job ID of the submitted job.

To check the status of submitted queue by using “qstat” command by type in the command the press “Enter” button. Then the screen will display the status of all queues that submitted into columns.

```
qstat
```

Job id	Name	User	Time Use S Queue
824.hydrogen	ReD1000	c00b010	02:15:42 R medium
825.hydrogen	ReD1200	c00b010	02:15:30 R medium
826.hydrogen	ReD1400	c00b010	02:15:15 Q long
831.hydrogen	ReD3000	c00b010	02:08:33 R long
835.hydrogen	ABINIT	c00c00x	00:56:43 R short
840.hydrogen	test.pbs	user	00:01:43 R test

Columns Header for Queue’s status

Job id = Job’s unique Identification number,

Name = The submitted PBS Script file name

User = User name that submitted the job

Time Use = Time that exist on the system

S = Stand for Status (R : Running), (Q : queuing/waiting)

Queue = Queue type specified on the job request. There are options of command qstat such as

qstat -q = to display all the jobs that are in the system's queue.

qstat -a = to display all the jobs in the system. This is include the running jobs and queuing jobs.

PBS Job Deletion

The command uses to delete the PBS Job is "qdel", then follows by the Job ID as display below:

```
qdel 840.hydrogen.lsr.nectec.or.th
```

```
qstat
```

Job id	Name	User	Time Use	S	Queue
824.hydrogen	ReD1000	c00b010	02:15:42	R	medium
825.hydrogen	ReD1200	c00b010	02:15:30	R	medium
826.hydrogen	ReD1400	c00b010	02:15:15	Q	long
831.hydrogen	ReD3000	c00b010	02:08:33	R	long
835.hydrogen	ABINIT	c00c00x	00:56:43	R	short

As the result of using "qdel" above, the Job id : 840 has disappear on the screen.

HPC Resources status checking

To check the HPC Resources status can be done by "pbsnodes -a". The screen will display the server name, state, number of nodes, properties, server type, and many more as display below.

```
pbsnodes -a
```

```
sodium-0-0.ib  
state = free  
np = 12  
properties = base,old  
ntype = cluster  
jobs = 0/43916.hydrogen.lsr.nectec.or.th,  
status = uname=Linux sodium-0-0.local  
el6.x86_64,totmem=49540512kb,ncpus=12,loadave=10.31,state=free
```