

Intel® Manycore Testing Lab (MTL)

Getting Started Guide

Introduction

What are the intended uses of the MTL?

The MTL is prioritized for supporting the Intel Academic Community for the testing, validation and scaling of parallel algorithms and workloads, primarily for courseware delivery, and secondly for research -- based on availability.

The MTL supports a default shared login node (many users), for workload development and a limited number of exclusive (one user) batch nodes for benchmarking. The nodes are located in DuPont, WA, USA, and are directly connected to the Internet via dedicated firewall devices.

The MTL is not configured or setup as a cluster, so MPI is not a supported option for our community. The current system resources therefore do not support a distributed memory programming model.

How do I get an account?

Use of the MTL is a benefit for members of the Intel Academic Community, available free of charge, for the asking. However, users must request an account from the Intel Academic Community; <http://software.intel.com/en-us/academic/>

Where can I get a copy of a VPN client to access MTL?

Cisco® owns the rights for the VPN client – it can be obtained from this site: <http://helpdesk.ugent.be/vpn/en/akkoord.php>

Download the Cisco VPN client (without Config).

Note: these only work with 32-bit clients.

Within the VPN client please specify the following (under Group Authentication):

Host: **192.55.51.80**

Name: <VPN_GROUP_NAME>

Password: <VPN_PASSWORD>

Connection Entry & Description (not significant)

Note: when the VPN tunnel is connected you will not be able to use your machine to connect to anything else, i.e. the web or other systems.

How do I connect?

Once an account is available, use **SSH** to connect to one of the following IP addresses using your favorite terminal connection software, e.g. Putty, F-Secure, etc:

IP-address	System name	Connection Model
36.81. 203.1	acano01	VPN
192.55.51.81	acano01	Direct IP

Note: If you're using a VPN client to connect to the MTL, you must first invoke the VPN client to establish the connection to the Host (named above), before using the appropriate terminal connection model

Intermediate

What tools are available on the MTL?

The login node as well as the rest of the MTL is built upon Linux. The current distribution in use is Redhat Enterprise Linux (RHEL 5.4, kernel 2.6.18-164.el5) but that particular distribution is subject to change and an alternative distribution can be requested.

A reasonable set of login shells are available including the default, bash, as well as tsh and zsh. In addition most of the typical Linux command line tools such as vim, emacs, gcc, make, python, perl, mc, etc are available. VNC(GUI) and screen(tty) interfaces are available for multiplexing your ssh connection as well as preserving the state of user logins.

VNC will need to be routed through the primary ssh connection using port forwarding.

Man pages are available for most commands and are a good starting point when you have questions.

What resources are available on the MTL?

A PBSPro 10.2 batch system, is available and required for exclusive job submission.

For the latest PBSPro v10.2 commands, please. read:
`/opt/docs/PBSProUserGuide10.2.pdf`

The `/home` directories are NFS-mounted from a standard storage server with a total capacity of over 1TB.

The Intel compilers, debugger, MKL, TBB, IPP can also be found under /opt/intel by version number.

Additional system wide tools will be made available, so look for announcements in the MTL forum: <http://software.intel.com/en-us/forums/intel-manycore-testing-lab/>

How do I compile my code using the xx compiler or yy library?

gcc/g++ (v4.1.2) is the default compiler

To compile with the updated version (4.4.3) of the gcc compiler suite (that supports OpenMP v3.0), use the following commands in your makefile:

```
GCC_VERSION = 4.4.3
PREFIX = /opt/gcc/${GCC_VERSION}/bin
CC = ${PREFIX}/gcc
CPP = ${PREFIX}/g++
LD_LIBRARY_PATH = /opt/mpfr/lib:/opt/gmp/lib
```

To compile with the Intel 32-bit compiler, execute the following source command:

```
$ source /opt/intel/Compiler/<version>/<release>/bin/iccvars.sh ia32
```

To compile with the Intel 64-bit compiler, execute the following source command:

```
$ source /opt/intel/Compiler/<version>/<release>/bin/iccvars.sh intel64
```

To use the TBB library, execute the following source command:

```
$ source /opt/intel/Compiler/<version>/<release>/tbb/bin/tbbvars.sh
[ia32,intel64]
```

How do I run a program in an exclusive mode (say for benchmarking)?

Qsub is the command used to submit jobs to the batch system. Using a script to “wrap” the actual program or programs is the most effective way to submit jobs. Here is a very simple script that merely prints the date when submitted to PBS using the following qsub command:

For latest PBSPro v10.2 qsub commands, please. read:

</opt/docs/PBSProUserGuide10.2.pdf> - example subsection

```
$ qsub $HOME/myjob
```

The contents of the 5 line myjob (for OpenMP) could be:

```
#!/bin/sh
#PBS -N myjob
#PBS -j oe
export OMP_NUM_THREADS=16
./hello_world
```

The output of the batch job will be left in a file in the working directory where the qsub command was entered. The file will be named using the PBS job name (e.g. myjob) followed by a suffix built using .o + the job number, e.g. myjob.o28802

Note: Because of the shared environment on the MTL login node (acano01), it's not recommended that repeatable testing (and results) be performed on this node, the exclusive batch system it set up for this purpose.

How do I monitor the programs launched?

The qstat utility will output the status of the jobs submitted to PBS. See the man page for the various formatting options. Here is an example output:

```
$ qstat -a
```

```
acano01:
```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	S	Elap Time
323.acano01	user01	workq	myjob	19290	1	32	--	01:00	R	00:03
324.acano01	user01	workq	myjob	19659	1	32	--	01:00	R	00:02
325.acano01	user01	workq	myjob	--	1	32	--	01:00	Q	--
326.acano01	user01	workq	myjob	--	1	32	--	01:00	Q	--
327.acano01	user01	workq	myjob	--	1	32	--	01:00	Q	--

How do I abort a program after launch?

The qdel command can be used to abort a job submitted with qsub. The job ID is the only parameter required:

```
$ qdel <job_id>
```

Advanced

How can a program be run interactively?

Use the -I option to qsub and you will be dropped into a shell on the first compute node:

```
$ qsub -I
```

How is batch run duration specified?

```
$ qsub "walltime=1:0:0" ./my_script
```

How can a run-away or unresponsive program be stopped?

You can try using `qsig` command.

```
$ qsig -s 9 job_id
```

How can I determine if my batch job will possibly run immediately?

There are a number of exclusive batch nodes on the MTL, but they may be busy running other jobs. You can run the `qfree` script to determine if there are any batch nodes free at this time. If not then you may have to wait for your batch job to be scheduled once you submit it with `qsub`. This `qfree` script will not reserve a batch node for your use, it just reports the current status of all the available MTL batch nodes. Here is an example output:

```
$ qfree
```

```
Number of MTL Batch nodes free: 1
```

```
Number of MTL Batch nodes busy: 1
```

How can I transfer files to the MTL and from the Internet?

You can always initiate a `scp` (`winscp`) session to MTL from your own system and use this to upload and download files to and from your local login system. However, due to firewall restrictions, this will only work from the login server (`acano01`). The MTL login node (`acano01`) is restricted in its ability to directly access the internet for security reasons. Also it's not recommended to transfer large files or data sets (GBs), using `scp` – they should be split into smaller chunks and submitted in parallel.

How do I backup my code and/or data sets?

The MTL does not support any kind of backup of user data. It is suggested that users keep a local copy of their data/code, to mitigate any potential loss of data from the MTL.

Why when I run my workload, do I get resources temporarily unavailable or unable to create thread?

This is possibly due to exceeding the number of processes/threads that are allocated on a per user basis.