

Introduction to Condor

(and using it at Lehigh University)

Ashutosh Mahajan

High Performance Computing Lehigh University

Lehigh, September 28, 2006

Checklist

You should have:

- 1 Login and password for blaze1.cc.lehigh.edu
- 2 Basic Linux Survival Skills (Bliss)
 - 1 Moving around
 - 2 copy, paste, edit, save, delete
 - 3 running, killing jobs etc.
- 3 C/C++/pascal/fortran knowledge

```
cat .bashrc  
rsync -av /home/asm4/condor .
```

What do you want to solve?

¹100 years of cpu time

²50 dedicated processors for 30 days

³50,000

What do you want to solve?



: A **BIG**¹ MW problem

¹100 years of cpu time

²50 dedicated processors for 30 days

³50,000

What do you want to solve?



: A **BIG**¹ MW problem



: A **BIG**² MPI problem

¹100 years of cpu time

²50 dedicated processors for 30 days

³50,000

What do you want to solve?



: A **BIG**¹ MW problem



: A **BIG**² MPI problem



: **MANY**³ small problems

¹100 years of cpu time

²50 dedicated processors for 30 days

³50,000

Buy me a Supercomputer!

- Costs Money.
- Limited Software.
- Limited Hardware.
- Still needs scheduling.

Buy me a Supercomputer!

- Costs Money.
- Limited Software.
- Limited Hardware.
- Still needs scheduling.

HPG

I **DON'T** need a High Performance Computing

Buy me a Supercomputer!

- Costs Money.
- Limited Software.
- Limited Hardware.
- Still needs scheduling.

HPG

I **DON'T** need a High Performance Computing

Just Give me High Throughput

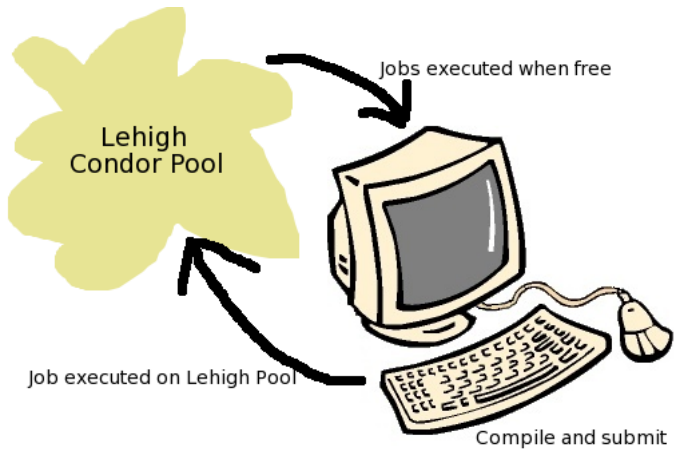
Build a cluster!

- Connect a lot of computers (Lehigh has 1000's).
- Gain access to these computers.
- Start running your jobs.
- Still needs scheduling.


Build a cluster!

- Connect a lot of computers (Lehigh has 1000's).
- Gain access to these computers.
- Start running your jobs.
- Still needs scheduling.

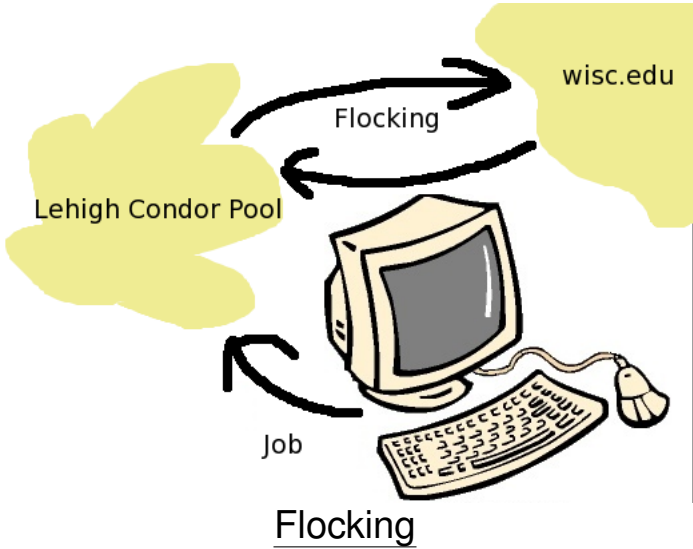
Thanks! but bad0 controls my computer.



Pooling resources



Checkpointing and migration





Condor

High Throughput Computing

Condor is a **lot** more than a scheduler

- Can use non-dedicated resources. [Use a machine only when free]
- Checkpoint and migrate
- Remote I/O
- Flocking to other pools

Fire and forget

Users submit their serial or parallel jobs to Condor, Condor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion.



What can condor run

Anything!

- Simple user compiled programs: c, c++, fortran, python scripts ...
- Parallel programs
- System commands
- Compilation/installation of condor!

\$\$\$oftware

It can even run commercial licen\$e ba\$ed \$oftware.

Condor multiverse

- Vanilla: Simple programs
- Standard: Programs with checkpointing: Needs linking with condor libraries.
- Parallel: Parallel programs. MPI/PVM/anything else!. Needs dedicated machines.
- Java
- Many more.

Condor at Lehigh

- blaze: 128 processors, 64-bit
- fire: 96
- Cor@l: 57
- vega: 825
- More (egenera, smp1 ...)

Starting with condor

Log in

```
ssh hpcxxx@blaze1.cc.lehigh.edu
```

Setting up

```
[hpc40@blaze1 ]$ scp -r /home/asm4/condor .  
[hpc40@blaze1 ]$ cd condor  
[hpc40@blaze1 condor]$ condor -version
```

Setting up

```
[hpc40@blaze1 condor]$ echo $PATH  
$ export PATH=/usr/local/condor/bin:::$PATH
```

Looking at condor

Status

```
[hpc40@blazel1 condor]$ condor_status
[hpc40@blazel1 condor]$ condor_status -total
[hpc40@blazel1 condor]$ condor_status -server
[hpc40@blazel1 condor]$ condor_status -server
[hpc40@blazel1 condor]$ condor_status -submitters
[hpc40@blazel1 condor]$ condor_q asm4 -analyze
[...]$ condor_status -constraint 'RemoteUser ==
"bad0.cc.lehigh.edu"'
```

Looking at condor

Queues and Priorities

```
[hpc40@blazel1 condor]$ condor_q asm4  
[hpc40@blazel1 condor]$ condor_q asm4 -long : less  
[hpc40@blazel1 condor]$ condor_userprio -allusers  
[hpc40...]$ condor_userprio -all -allusers
```

Running simple jobs

Submit in a vanilla universe

```
[hpc40@blaze1 condor]$ cd serial
[hpc40@blaze1 serial]$ make
[hpc40@blaze1 serial]$ cat partition.condor
$ condor_submit partition.condor
$ condor_status -submitters
$ condor_rm hpcxxx
$ condor_submit partition.condor
$ condor_status -submitters
$ condor_rm hpcxxx
```

Running simple jobs

Submit in a standard universe

- 1 needs to be linked with condor libraries
- 2 condor_compile
- 3 rest is same as a vanilla universe

Running MPI jobs

Submit in a parallel universe

```
[hpc40@blaze1 condor]$ cd mpi
[hpc40@blaze1 condor]$ make
[hpc40@blaze1 condor]$ cat pi.condor
[hpc40@blaze1 condor]$ condor_submit pi.condor
$ condor_status -submitters
```

- ssh lehigh-id@vega.cc.lehigh.edu and see condor there.
- <http://www.cs.wisc.edu/condor/> for more information.