# Supercomputing in Plain English

High Throughput Computing

Henry Neeman, Director

Director, OU Supercomputing Center for Education & Research (OSCER)
Assistant Vice President, Information Technology – Research Strategy Advisor
Associate Professor, College of Engineering
Adjunct Associate Professor, School of Computer Science
University of Oklahoma
Tuesday April 7 2015

This is an experiment!

It's the nature of these kinds of videoconferences that
**FAILURES ARE GUARANTEED TO HAPPEN!
NO PROMISES!**

So, please bear with us. Hopefully everything will work out well enough.

If you lose your connection, you can retry the same kind of connection, or try connecting another way.

Remember, if all else fails, you always have the toll free phone bridge to fall back on.

PLEASE MUTE YOURSELF

No matter how you connect, **PLEASE MUTE YOURSELF**, so that we cannot hear you.

At OU, we will turn off the sound on all conferencing technologies.

That way, we won't have problems with echo cancellation.

Of course, that means we cannot hear questions.

So for questions, you'll need to send e-mail.

PLEASE MUTE YOURSELF.

PLEASE MUTE YOURSELF.

PLEASE REGISTER

If you haven't already registered, please do so.

You can find the registration link on the SiPE webpage:

http://www.oscer.ou.edu/education/

Our ability to continue providing Supercomputing in Plain English depends on being able to show strong participation.

We use our headcounts, institution counts and state counts (since 2001, over 2000 served, from every US state except RI and VT, plus 17 other countries, on every continent except Australia and Antarctica) to improve grant proposals.

Download the Slides Beforehand

Before the start of the session, please download the slides from the Supercomputing in Plain English website:

`http://www.oscer.ou.edu/education/`

That way, if anything goes wrong, you can still follow along with just audio.

PLEASE MUTE YOURSELF.

H.323 (Polycom etc) #1

If you want to use H.323 videoconferencing – for example, Polycom – then:

- If you AREN'T registered with the OneNet gatekeeper (which is probably the case), then:
  - Dial `164.58.250.51`
  - Bring up the virtual keypad.
    On some H.323 devices, you can bring up the virtual keypad by typing:
    #
    (You may want to try without first, then with; some devices won't work with the #, but give cryptic error messages about it.)
  - When asked for the conference ID, or if there's no response, enter:
    `0409`
  - On most but not all H.323 devices, you indicate the end of the ID with:
    #

H.323 (Polycom etc) #2

If you want to use H.323 videoconferencing – for example, Polycom – then:

- If you ARE already registered with the OneNet gatekeeper (most institutions aren't), dial:

    ```
    2500409
    ```

Many thanks to James Deaton, Skyler Donahue, Jeremy Wright and Steven Haldeman of OneNet for providing this.

PLEASE MUTE YOURSELF.

Wowza #1

You can watch from a Windows, MacOS or Linux laptop using Wowza from the following URL:

<div align="center">

`http://jwplayer.onenet.net/stream6/sipe.html`

</div>

Wowza behaves a lot like YouTube, except live.

Many thanks to James Deaton, Skyler Donahue, Jeremy Wright and Steven Haldeman of OneNet for providing this.

PLEASE MUTE YOURSELF.

Wowza #2

Wowza has been tested on multiple browsers on each of:

- Windows (7 and 8): IE, Firefox, Chrome, Opera, Safari
- MacOS X: Safari, Firefox
- Linux: Firefox, Opera

PLEASE MUTE YOURSELF.

RTMP

If you have a video player that can handle RTMP, you can watch the Wowza feed that way:

rtmp://stream3.onenet.net/live/mp4:sipe-wowza

Toll Free Phone Bridge

IF ALL ELSE FAILS, you can use our toll free phone bridge:

800-832-0736

* 623 2874 #

Please mute yourself and use the phone to listen.

Don't worry, we'll call out slide numbers as we go.

Please use the phone bridge **ONLY** if you cannot connect any other way: the phone bridge can handle only 100 simultaneous connections, and we have over 500 participants.

Many thanks to OU CIO Loretta Early for providing the toll free phone bridge.

PLEASE MUTE YOURSELF.

Please Mute Yourself

No matter how you connect, **PLEASE MUTE YOURSELF**, so that we cannot hear you.

(For Wowza, you don't need to do that, because the information only goes from us to you, not from you to us.)

At OU, we will turn off the sound on all conferencing technologies.

That way, we won't have problems with echo cancellation.

Of course, that means we cannot hear questions.

So for questions, you'll need to send e-mail.

PLEASE MUTE YOURSELF.

PLEASE MUTE YOURSELF.

Questions via E-mail Only

Ask questions by sending e-mail to:

sipe2015@gmail.com

All questions will be read out loud and then answered out loud.

PLEASE MUTE YOURSELF.

Onsite: Talent Release Form

If you're attending onsite, you **<u>MUST</u>** do one of the following:

- complete and sign the Talent Release Form,

OR

- sit behind the cameras (where you can't be seen) and don't talk at all.

If you aren't onsite, then **PLEASE MUTE YOURSELF.**

TENTATIVE Schedule

Tue Jan 20: Hi Thruput: What the Heck is Supercomputing?

Tue Jan 27: The Tyranny of the Storage Hierarchy

Tue Feb 3: Instruction Level Parallelism

Tue Feb 10: Stupid Compiler Tricks

Tue Feb 17: Shared Memory Multithreading

Tue March 3: Distributed Multiprocessing

Tue March 10: Applications and Types of Parallelism

Tue March 17: **NO SESSION** (OU's Spring Break)

Tue March 24: **NO SESSION** (Henry has a huge grant proposal due)

Tue March 31: Multicore Madness

Tue Apr 7: High Throughput Computing

Tue Apr 14: GPGPU: Number Crunching in Your Graphics Card

Tue Apr 21: Grab Bag: Scientific Libraries, I/O Libraries, Visualization

Thanks for helping!

- OU IT
  - OSCER operations staff (Brandon George, Dave Akin, Brett Zimmerman, Josh Alexander, Patrick Calhoun)
  - Horst Severini, OSCER Associate Director for Remote & Heterogeneous Computing
  - Debi Gentis, OSCER Coordinator
  - Jim Summers
  - The OU IT network team
- James Deaton, Skyler Donahue, Jeremy Wright and Steven Haldeman, OneNet
- Kay Avila, U Iowa
- Stephen Harrell, Purdue U

This is an experiment!

It's the nature of these kinds of videoconferences that **FAILURES ARE GUARANTEED TO HAPPEN! NO PROMISES!**

So, please bear with us. Hopefully everything will work out well enough.

If you lose your connection, you can retry the same kind of connection, or try connecting another way.

Remember, if all else fails, you always have the toll free phone bridge to fall back on.

PLEASE MUTE YOURSELF.

Coming in 2015!

- Linux Clusters Institute workshop May 18-22 2015 @ OU

  http://www.linuxclustersinstitute.org/workshops/
- Great Plains Network Annual Meeting, May 27-29, Kansas City
- Advanced Cyberinfrastructure Research & Education Facilitators (ACI-REF) Virtual Residency May 31 - June 6 2015
- XSEDE2015, July 26-30, St. Louis MO

  https://conferences.xsede.org/xsede15
- IEEE Cluster 2015, Sep 23-27, Chicago IL

  http://www.mcs.anl.gov/ieeecluster2015/
- OKLAHOMA SUPERCOMPUTING SYMPOSIUM 2015, **Sep 22-23 2015** @ OU
- SC13, Nov 15-20 2015, Austin TX

  http://sc15.supercomputing.org/

PLEASE MUTE YOURSELF.

Outline

- What is High Throughput Computing?
- Tightly Coupled vs Loosely Coupled
- What is Opportunistic Computing?
- HTCondor
- Grid Computing

What is High Throughput Computing?

High Throughput Computing

High Throughput Computing (HTC) means getting lots of work done per large time unit (for example, jobs per month).

This is different from **High *Performance* Computing** (HPC), which means getting **a particular job** done in less time (for example, calculations per second).

Throughput vs Performance

- **<u>Throughput</u>** is a side effect of how much time your job takes from when you first submit it until it completes.

- **<u>Performance</u>** is the factor that controls how much time your jobs takes from when it first starts running until it completes.

- Example:

  - You submit a very big job at 1:00am on January 1.

  - It sits in the queue for a while.

  - It starts running at 5:00pm on January 2.

  - It finishes running at 6:00pm on January 2.

  - Its performance is fast; its throughput is slow.

High Throughput on a Cluster?

Is it possible to get high throughput on a cluster?

Sure – it just has to be a cluster that no one else is trying to use!

Normally, a cluster that is shared by many users is fully loaded with jobs all the time. So your throughput depends on when you submit your jobs, how big your jobs are, and even how many jobs you submit at a time.

Depending on a variety of factors, a job you submit may wait in the queue for anywhere from seconds to days.

Tightly Coupled vs Loosely Coupled

Tightly Coupled vs Loosely Coupled

- ***Tightly coupled*** means that all of the parallel tasks have to advance forward in lockstep, so they have to communicate frequently.

- ***Loosely coupled*** means that the parallel tasks can largely or completely ignore each other (little or no communication), and they can advance at different rates.

Tightly Coupled Example

Consider weather forecasting.

You take your simulation domain – for example, the continental United States – split it up into chunks, and give each chunk to an MPI process.

But, the weather in northern Oklahoma affects the weather in southern Kansas.
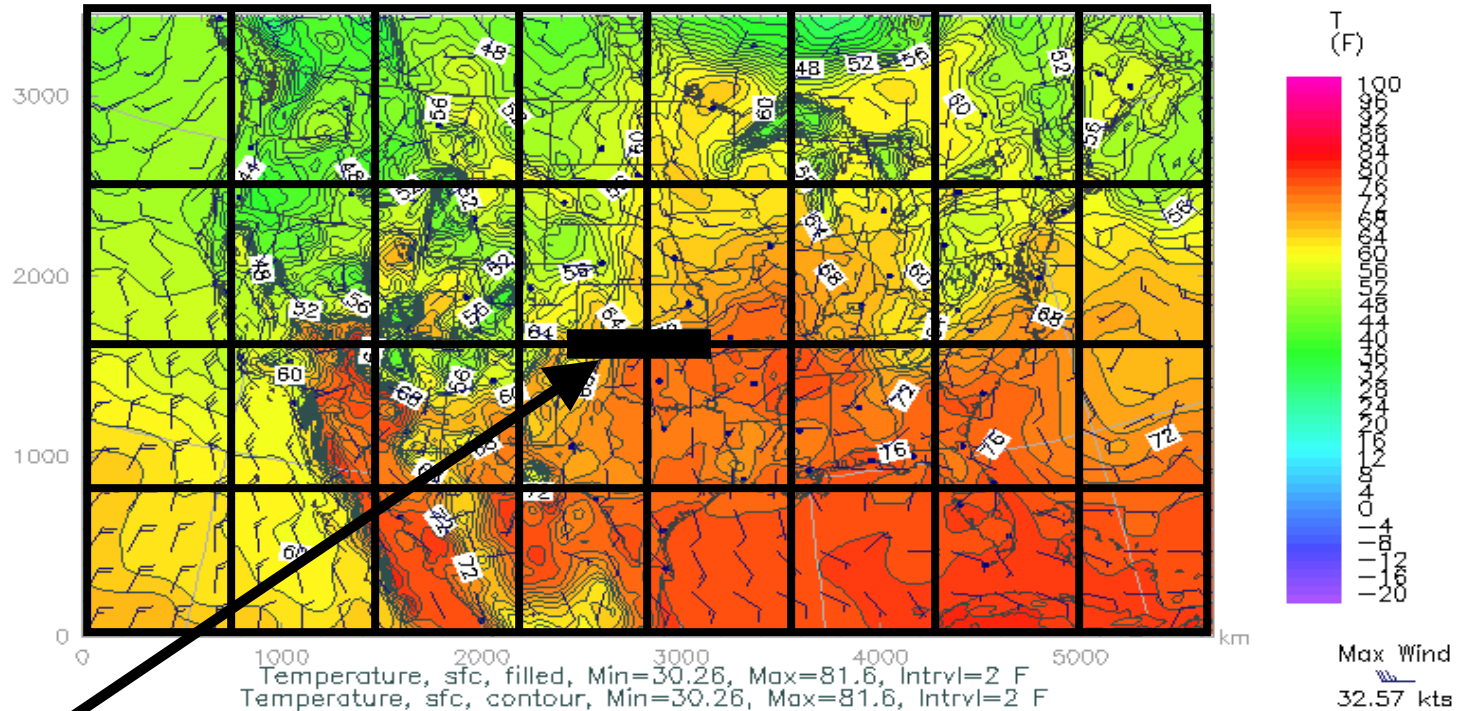
So, every single timestep, the process that contains northern Oklahoma has to communicate with the process that contains southern Kansas, so that the interface between the processes has the same weather at the same time.

Tightly Coupled Example

Thu, 25 May 2006, 8 am CDT (13Z)
Surface Temperature

Temperature, sfc, filled, Min=30.26, Max=81.6, Intrvl=2 F
Temperature, sfc, contour, Min=30.26, Max=81.6, Intrvl=2 F

OK/KS boundary

http://www.caps.ou.edu/wx/p/r/conus/fcst/

CAPS/OU Experimental ADAS Anlys   CONUS, 210x128x50, dx=27 km   05/25/06 08:45 CDT

Supercomputing in Plain English: Hi Thruput
Tue Apr 7 2015

Loosely Coupled Example

An application is known as ***embarrassingly parallel***, or ***loosely coupled***, if its parallel implementation:

1. can straightforwardly be broken up into roughly equal amounts of work per processor, **AND**
2. has minimal parallel overhead (for example, communication among processors).

We **love** embarrassingly parallel applications, because they get **near-perfect parallel speedup**, sometimes with only modest programming effort.

Monte Carlo Methods

Monte Carlo is a city in the tiny European country Monaco.

People gamble there; that is, they play games of chance, which involve randomness.

Monte Carlo methods are ways of simulating (or otherwise calculating) physical phenomena based on randomness.

Monte Carlo simulations typically are embarrassingly parallel.

Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.

BANG!

We want to know, say, the average properties of this phenomenon.

There are infinitely many ways that two particles can be banged together.

So, we can't possibly simulate all of them.

Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.

BANG!

We want to know, say, the average properties of this phenomenon.

There are infinitely many ways that two particles can be banged together.

So, we can't possibly simulate all of them.

Instead, we can **randomly choose a finite subset** of these infinitely many ways and simulate only the subset.

Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.

BANG!

We want to know, say, the average properties of this phenomenon.

There are infinitely many ways that two particles can be banged together.

So, we can't possibly simulate all of them.

The average of this subset will be **close** to the actual average.

Monte Carlo Methods

In a Monte Carlo method, you randomly generate a large number of example cases (***realizations***) of a phenomenon, and then take the average of the properties of these realizations.

When the realizations' average **converges** (that is, doesn't change substantially if new realizations are generated), then the Monte Carlo simulation stops.

This can also be implemented by picking a high enough number of realizations to be sure, mathematically, of convergence.

MC: Embarrassingly Parallel

Monte Carlo simulations are embarrassingly parallel, because each realization is completely independent of all of the other realizations.

That is, if you're going to run a million realizations, then:

1. you can straightforwardly break up into roughly $1M / N_p$ chunks of realizations, one chunk for each of the N_p processes, **AND**

2. the only parallel overhead (for example, communication) comes from tracking the average properties, which doesn't have to happen very often.

Serial Monte Carlo

Suppose you have an existing serial Monte Carlo simulation:

```
PROGRAM monte_carlo
  CALL read_input(…)
  DO realization = 1, number_of_realizations
    CALL generate_random_realization(…)
    CALL calculate_properties(…)
  END DO
  CALL calculate_average(…)
END PROGRAM monte_carlo
```

How would you parallelize this?

Parallel Monte Carlo: MPI

```
PROGRAM monte_carlo_mpi
  [MPI startup]
  IF (my_rank == server_rank) THEN
    CALL read_input(…)
  END IF
  CALL MPI_Bcast(…)
  number_of_realizations_per_process = &
& number_of_realizations / number_of_processes
  DO realization = 1, number_of_realizations_per_process
    CALL generate_random_realization(…)
    CALL calculate_realization_properties(…)
    CALL calculate_local_running_average(...)
  END DO
  IF (my_rank == server_rank) THEN
    [collect properties]
  ELSE
    [send properties]
  END IF
  CALL calculate_global_average_from_local_averages(…)
  CALL output_overall_average(...)
  [MPI shutdown]
END PROGRAM monte_carlo_mpi
```

Parallel Monte Carlo: HTC

Suppose you have an existing serial Monte Carlo simulation:

```
PROGRAM monte_carlo
  CALL read_input(…)
  number_of_realizations_per_job = &
 &    number_of_realizations / number_of_jobs
  DO realization = 1, number_of_realizations_per_job
    CALL generate_random_realization(…)
    CALL calculate_properties(…)
  END DO
  CALL calculate_average_for_this_job(…)
  CALL output_average_for_this_job(…)
END PROGRAM monte_carlo
```

To parallelize this for **HTC**, simply submit **number_of_jobs** jobs, and then at the very end run a little program to calculate the overall average.

What is Opportunistic Computing?

Desktop PCs Are Idle Half the Day

Desktop PCs tend to be active during the workday.

But at night, during most of the year, they're idle. So we're only getting **half their value** (or less).

Supercomputing at Night

Many institutions have lots of desktop PCs that are **idle during the evening and during intersessions**.

Wouldn't it be great to put them to work on something **useful** to the institution?

That is: What if they could pretend to be a big supercomputer **at night**, when they'd **otherwise be idle anyway**?

This is sometimes known as ***opportunistic computing***: When a desktop PC is otherwise idle, you have an opportunity to do number crunching on it.

Supercomputing at Night Example

SETI – the Search for Extra-Terrestrial Intelligence – is looking for evidence of green bug-eyed monsters on other planets, by mining radio telescope data.

SETI@home runs number crunching software as a screensaver on idle PCs around the world (2+ million PCs in 252 countries):

http://setiathome.berkeley.edu/ SETI HOME

There are **many similar projects**:

- folding@home (protein folding)
- climateprediction.net
- Einstein@Home (Laser Interferometer Gravitational wave Observatory)
- Cosmology@home
- …

BOINC

The projects listed on the previous page use a software package named BOINC (**B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing), developed at the University of California, Berkeley:

http://boinc.berkeley.edu/

To use BOINC, you have to insert calls to various BOINC routines into your code. It looks a bit similar to MPI:

```
int main ()
{ /* main */
    …
    boinc_init();
    …
    boinc_finish(…);
} /* main */
```

BOINC: a Big Supercomputer

- As of this morning (Tue Apr 7 2015), according to the BOINC website, the following were active:
  - 288,814 volunteers
  - 674,060 computers
  - 24-hour average: 7.495 PetaFLOPS
    - If this were a single machine, on the most recent Top 500 list, it'd be the 8th fastest supercomputer in the world.
- Note to self: If your science isn't easy to fund, do amazing computing, so someone will fund that ….

HTCondor

HTCondor is Like BOINC

- HTCondor **steals computing time** on existing desktop PCs **when they're idle**.

- HTCondor **runs in background** when no one is sitting at the desk.

- HTCondor allows an institution to get **much more value** out of the hardware that's **already purchased**, because there's little or no idle time on that hardware – all of the idle time is used for number crunching.

HTCondor is Different from BOINC

- To use HTCondor, **you don't need to rewrite your software** to add calls to special routines; in BOINC, you do.

- HTCondor **works great under Unix/Linux**, but less well under Windows or MacOS (more on this presently); BOINC works well under all of them.

- It's **non-trivial to install HTCondor** on your own personal desktop PC; it's straightforward to install a BOINC application such as SETI@home.

Useful Features of HTCondor

- **<u>Opportunistic</u>** computing: HTCondor steals time on existing desktop PCs when they're otherwise not in use.

- HTCondor **<u>doesn't require any changes to the software</u>**.

- HTCondor can **<u>automatically checkpoint</u>** a running job: Every so often, HTCondor saves to disk the state of the job (the values of all the job's variables, plus where the job is in the program).

- Therefore, HTCondor can **<u>preempt</u>** running jobs if more important jobs come along, or if someone sits down at the desktop PC.

- Likewise, HTCondor can **<u>migrate</u>** running jobs to other PCs, if someone sits at the PC or if the PC crashes.

- And, HTCondor can do all of its **<u>I/O over the network</u>**, so that the job on the desktop PC doesn't consume the desktop PCs local disk.

HTCondor Pool @ OU

Previously, OU IT had deployed a large HTCondor pool (795 desktop PCs in dozens of labs around campus).

OU's HTCondor pool provided a significant amount of computing.

The hardware and software cost is zero, and the labor cost is modest.

Also, we had been seeing empirically that lab PCs were available for HTCondor jobs about 80% of the time.

We stopped doing this when we realized that the lab team was overwhelmed and didn't have time to maintain HTCondor.

But, we had a meeting this morning about it ….

HTCondor Limitations

- The Unix/Linux version has **more features** than Windows or MacOS, which are referred to as "clipped."
- Under Windows, you can't checkpoint your jobs (though you can pause them).
- Your code **shouldn't be parallel** to do opportunistic computing (MPI requires a fixed set of resources throughout the entire run), and it shouldn't try to do any funky communication (for example, opening sockets).
- For a Red Hat Linux HTCondor pool, you have to be able to **compile your code** with gcc, g++, g77 or NAG f95 (which is a Fortran90-to-C translator that then calls gcc).
- Also, depending on the PCs that have HTCondor on them, you may have limitations on, for example, how big your jobs' RAM footprint can be.

Running a HTCondor Job

Running a job on HTCondor pool is a lot like running a job on a cluster:

1. You compile your code using the compilers appropriate for that resource.

2. You submit a batch script to the HTCondor system, which decides when and where your job runs, magically and invisibly.

Sample HTCondor Batch Script

```
Universe       = standard
Executable     = /home/hneeman/NBody/nbody_compiled_for_HTCondor
Notification   = Error
Notify_User    = hneeman@ou.edu
Arguments      = 1000 100
Input          = /home/hneeman/NBody/nbody_input.txt
Output         = nbody_$(Cluster)_$(Process)_out.txt
Error          = nbody_$(Cluster)_$(Process)_err.txt
Log            = nbody_$(Cluster)_$(Process)_log.txt
InitialDir     = /home/hneeman/NBody/Run001
Queue
```

The batch submission command is **condor_submit**, used like so:

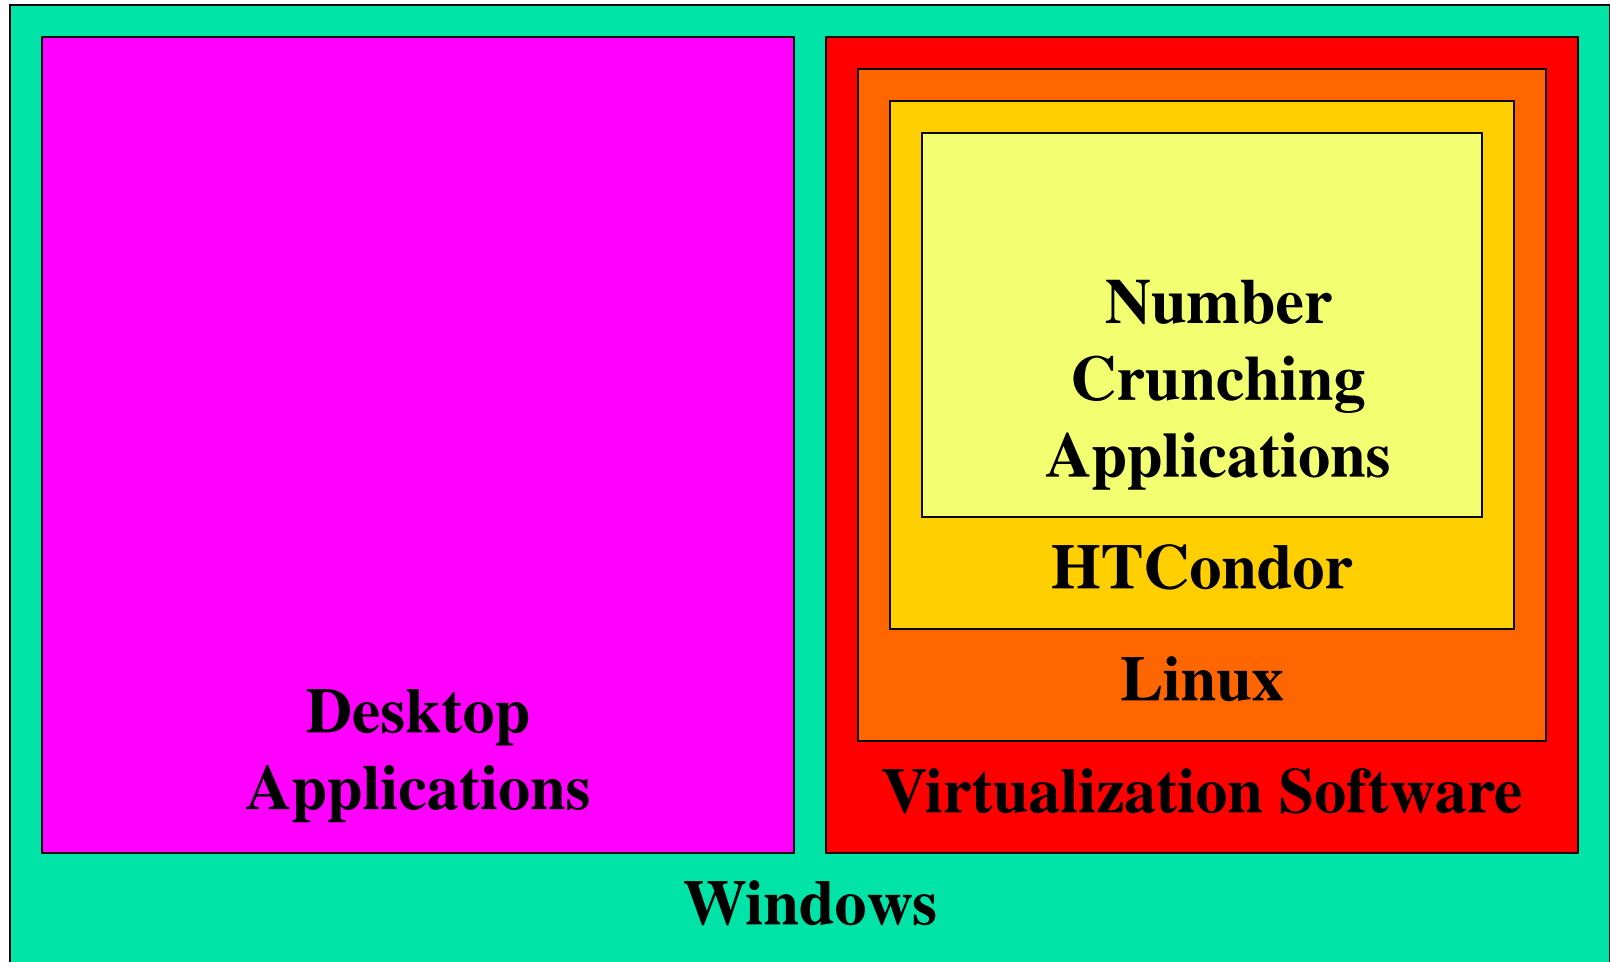**condor_submit   nbody.condor**

Linux HTCondor on Windows PCs?

If OU's HTCondor pool was using Linux, how could it be installed in OU IT PC labs? Don't those run Windows?

Yes.

Our solution was to run Linux inside Windows, as a "virtual machine."

HTCondor inside Linux inside Windows

Advantages of Linux inside Windows

- Under Linux, HTCondor is full featured rather than clipped.
  - Especially, you can do automatic checkpointing of jobs.
- Desktop users have a full Windows experience, without even being aware that HTCondor exists.

Grid Computing

What is Grid Computing?

The term ***grid computing*** is poorly defined, but the best definition I've seen so far is:

"a distributed, heterogeneous operating system."

A ***grid*** can consist of:

- compute resources;
- storage resources;
- networks;
- data collections;
- shared instruments;
- sensor networks;
- and so much more ....

Grid Computing is Like and Unlike ...

IBM's website has a very good description of grid computing:

- *"**Like the Web**, grid computing keeps complexity hidden: multiple users enjoy a single, unified experience.*
- *"**Unlike the Web**, which mainly enables communication, grid computing enables full collaboration toward common ... goals.*
- *"**Like peer-to-peer**, grid computing allows users to share files.*
- *"**Unlike peer-to-peer**, grid computing allows many-to-many sharing – not only files but other resources as well.*
- *"**Like clusters and distributed computing**, grids bring computing resources together.*
- *"**Unlike clusters and distributed computing**, which need physical proximity and operating homogeneity, grids can be geographically distributed and heterogeneous.*
- *"**Like virtualization technologies**, grid computing enables the virtualization of IT resources.*
- *"**Unlike virtualization technologies**, which virtualize a single system, grid computing enables the virtualization of vast and disparate IT resources."*

http://www.thocp.net/hardware/grid_computers.htm

HTCondor is Grid Computing

HTCondor creates a grid out of disparate desktop PCs.

(Actually, they don't have to be desktop PCs; they don't even have to be PCs. You can use HTCondor to schedule a cluster, or even on a big iron supercomputer.)

From a user's perspective, all of the PCs are essentially invisible; the user just knows how to submit a job, and everything happens magically and invisibly, and at some point the job is done and a result appears.

TENTATIVE Schedule

Tue Jan 20: Hi Thruput: What the Heck is Supercomputing?

Tue Jan 27: The Tyranny of the Storage Hierarchy

Tue Feb 3: Instruction Level Parallelism

Tue Feb 10: Stupid Compiler Tricks

Tue Feb 17: Shared Memory Multithreading

Tue March 3: Distributed Multiprocessing

Tue March 10: Applications and Types of Parallelism

Tue March 17: **NO SESSION** (OU's Spring Break)

Tue March 24: **NO SESSION** (Henry has a huge grant proposal due)

Tue March 31: Multicore Madness

Tue Apr 7: High Throughput Computing

Tue Apr 14: GPGPU: Number Crunching in Your Graphics Card

Tue Apr 21: Grab Bag: Scientific Libraries, I/O Libraries, Visualization

Thanks for helping!

- OU IT
  - OSCER operations staff (Brandon George, Dave Akin, Brett Zimmerman, Josh Alexander, Patrick Calhoun)
  - Horst Severini, OSCER Associate Director for Remote & Heterogeneous Computing
  - Debi Gentis, OSCER Coordinator
  - Jim Summers
  - The OU IT network team
- James Deaton, Skyler Donahue, Jeremy Wright and Steven Haldeman, OneNet
- Kay Avila, U Iowa
- Stephen Harrell, Purdue U

Coming in 2015!

Linux Clusters Institute workshop May 18-22 2015 @ OU

> http://www.linuxclustersinstitute.org/workshops/

Great Plains Network Annual Meeting, May 27-29, Kansas City

Advanced Cyberinfrastructure Research & Education Facilitators (ACI-REF) Virtual Residency May 31 - June 6 2015

XSEDE2015, July 26-30, St. Louis MO

> https://conferences.xsede.org/xsede15

IEEE Cluster 2015, Sep 23-27, Chicago IL

> http://www.mcs.anl.gov/ieeecluster2015/

OKLAHOMA SUPERCOMPUTING SYMPOSIUM 2015, **Sep 22-23 2015** @ OU

SC13, Nov 15-20 2015, Austin TX

> http://sc15.supercomputing.org/

OK Supercomputing Symposium 2015

2003 Keynote:
Peter Freeman
NSF
Computer & Information
Science & Engineering
Assistant Director

2004 Keynote:
Sangtae Kim
NSF Shared
Cyberinfrastructure
Division Director

2005 Keynote:
Walt Brooks
NASA Advanced
Supercomputing
Division Director

2006 Keynote:
Dan Atkins
Head of NSF's
Office of
Cyberinfrastructure

2007 Keynote:
Jay Boisseau
Director
Texas Advanced
Computing Center
U. Texas Austin

2008 Keynote:
José Munoz
Deputy Office
Director/Senior
Scientific Advisor
NSF Office of
Cyberinfrastructure

2009 Keynote:
Douglass Post
Chief Scientist
US Dept of Defense
HPC Modernization
Program

2010 Keynote:
Horst Simon
Deputy Director
Lawrence Berkeley
National Laboratory

2011 Keynote:
Barry Schneider
Program Manager
National Science
Foundation

2012 Keynote:
Thom Dunning
Director
National Center for
Supercomputing
Applications

2015 Keynote:
John Shalf
Dept Head CS
Lawrence
Berkeley Lab
CTO, NERSC

2014 Keynote:
Irene Qualters
Division Director
Advanced
Cyberinfarstructure
Division, NSF

FREE!
Wed Sep 23 2015
@ OU

Reception/Poster Session
Tue Sep 22 2015 @ OU
Symposium
Wed Sep 23 2015 @ OU

Thanks for your attention!

Questions?

www.oscer.ou.edu