PURDUE UNIVERSITY

TeraGrid™

Condor at Purdue: Vegetating, Virtualization, and Videos

May 5, 2011

Ben Cotton
Purdue University
bcotton@purdue.edu

ITaP
INFORMATION TECHNOLOGY AT PURDUE

PURDUE UNIVERSITY

But first!

ITaP
INFORMATION TECHNOLOGY AT PURDUE

We've got a lot of cores

- Over 37 kilocores across campus

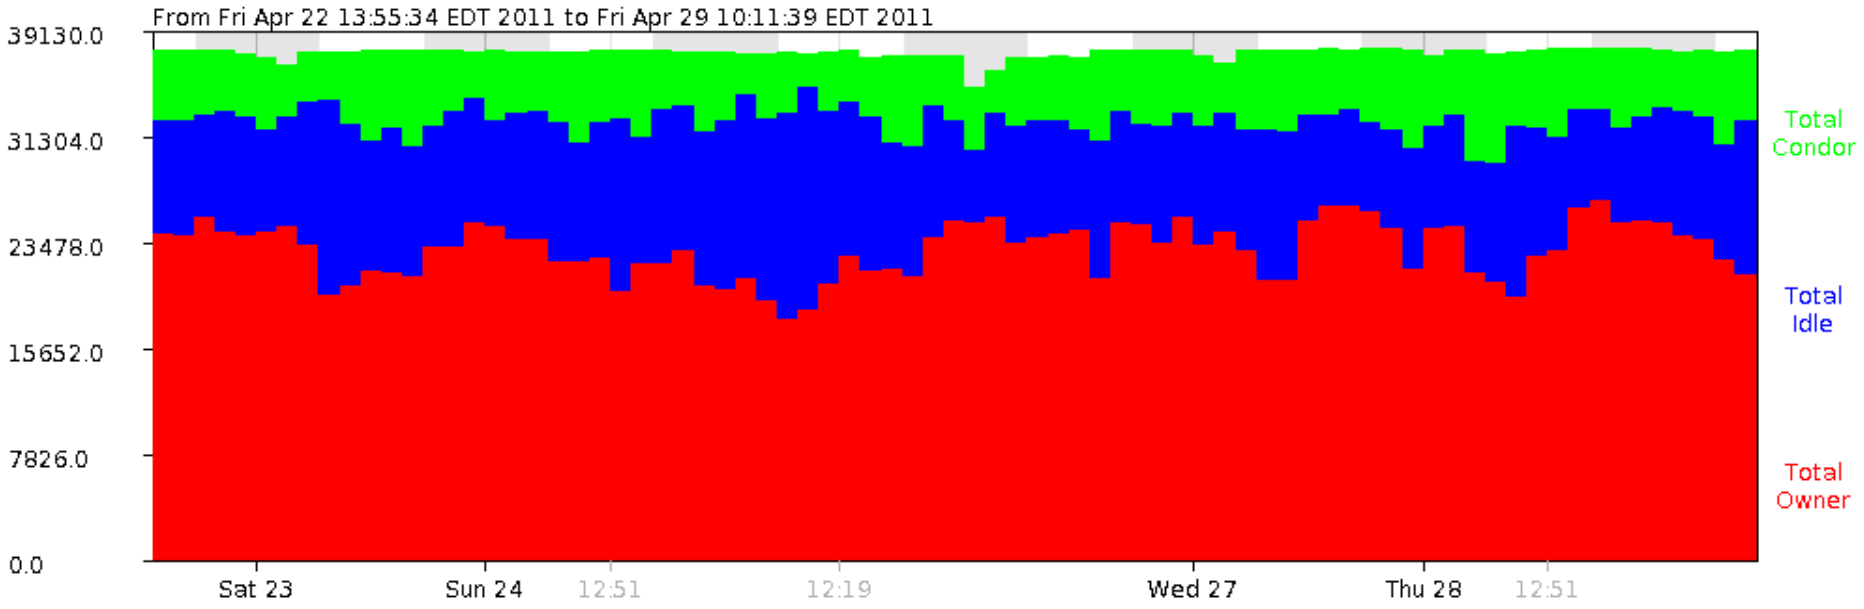

We've got a lot of cores

- Three community clusters (Steele, Coates, Rossmann)
- Two "ownerless" clusters (Radon, Miner)
- CMS Tier-2 cluster
- Other small clusters
- Instructional labs and academic departments

PURDUE UNIVERSITY

Might as well share

DiaGrid

ITaP
INFORMATION TECHNOLOGY AT PURDUE

PURDUE UNIVERSITY

Vegetating

ITaP
INFORMATION TECHNOLOGY AT PURDUE

Vegetating = hibernation

- Financial concerns
  - CIO must find $15 million in *recurring* cost savings
- Environmental concerns
  - Labs get really toasty!
  - Burning coal is not the most environmentally-friendly thing to do
- Let's try hibernation

The path to hibernation

- Condor support
  - Requires version 7.4.3/7.5.3 or greater for waking to work properly
- Network ACLs
  - Allow WoL traffic across subnets
  - Alternative: run a rooster inside each subnet
- Is condor_kbdd.exe running as the logged in user?
  - It had better be!
- Finding willing test subjects
  - This is actually the hard part

Hibernation results?

- To be determined…
  - Working with lab admins to test in several labs
  - Reaching out to departments to find more testers
- …but CIO report estimates ~ $300,000/year in savings
  - Based on hibernating 26k machines 8 hours a day
  - Assumes 100% compliance
  - Ignores increased hardware failure rate
  - Saves less than 1% of budget

But when we have results…

- There's a dashboard for that

PURDUE UNIVERSITY **Dashboards** DiaGrid

Energy Savings

Refresh

Owner	Energy Saved ▼	Energy Used	Avg Watts	Compute
undefined	252,069 kWh	317,857,566 kWh	135 kW	269y 181
Science Admin	3,347 kWh	1,732 kWh	625 W	115d 12h
ITaP - Instructional L	0 Wh	3,169,920,112 kWh	329 kW	1099y 35
crc	0 Wh	36,119,982 kWh	36.97 kW	111y 191
ITaP - Research Clus	0 Wh	53,928,541,149 kWh	850 kW	7241y 8d
dthain	0 Wh	5,647,494 kWh	13.62 kW	47y 117d
ramzi	0 Wh	1,450 kWh	298 W	202d 15h
cpoellab	0 Wh	815 kWh	298 W	113d 19h
Office of the Presider	0 Wh	1,455 kWh	299 W	202d 23h
rjansen	0 Wh	33,063 kWh	875 W	4y 114d
cse	0 Wh	7,016,075 kWh	14.91 kW	53y 261d
Psychology	0 Wh	10,000 kWh	531 W	2y 54d 1
striegel	0 Wh	371,081 kWh	3.88 kW	10y 337d
dchen	0 Wh	5,536 kWh	582 W	1y 30d 2
unknown	0 Wh	4,878,023 kWh	9.96 kW	55y 337d
izaguirr	0 Wh	3,379 kWh	455 W	309d 8h
semrich	0 Wh	25,183 kWh	1.10 kW	2y 222d
mniemier	0 Wh	5,489 kWh	580 W	1y 29d 7
Agronomy	0 Wh	81,929 kWh	2.24 kW	4y 63d 7
Purdue RCAC VM	0 Wh	1,212,870 kWh	12.91 kW	10y 265d
curt	0 Wh	40,053 kWh	1.43 kW	3y 71d 1
cheg	0 Wh	45,278 kWh	1.69 kW	3y 23d 2
ITaP - Other	0 Wh	512,616 kWh	5.49 kW	10y 243d
flynn	0 Wh	460,731 kWh	2.19 kW	24y 12d

Energy Usage by Owner

Time Frame: 3 Hours | Day | Week | Month View as: Area | Line

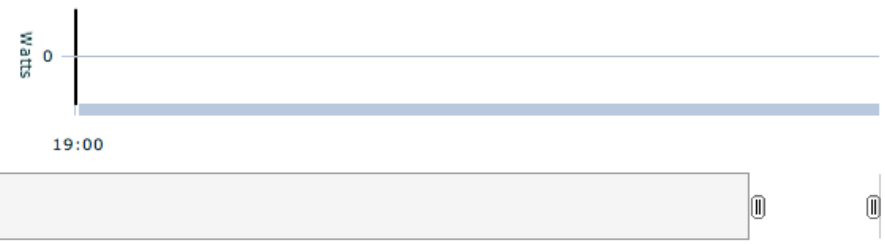

Total Energy Savings

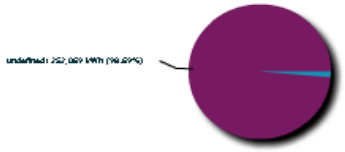

ITaP INFORMATION TECHNOLOGY AT PURDUE

PURDUE UNIVERSITY

Virtualization

- We have a lot of unused Windows cores…
  - Approximately 6 kilocores
- …and not many Windows jobs
  - 0.5% of submitted jobs in the past year

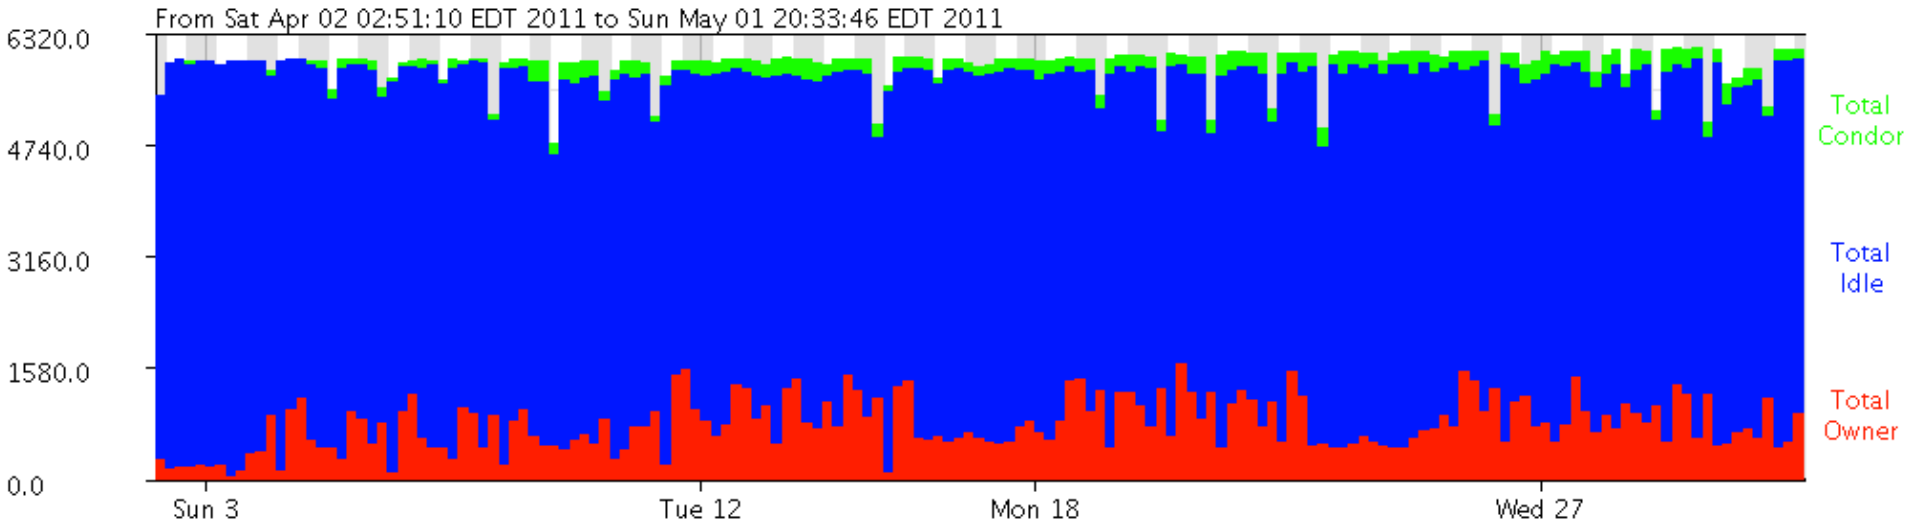

ITaP INFORMATION TECHNOLOGY AT PURDUE

PURDUE UNIVERSITY

The solution?

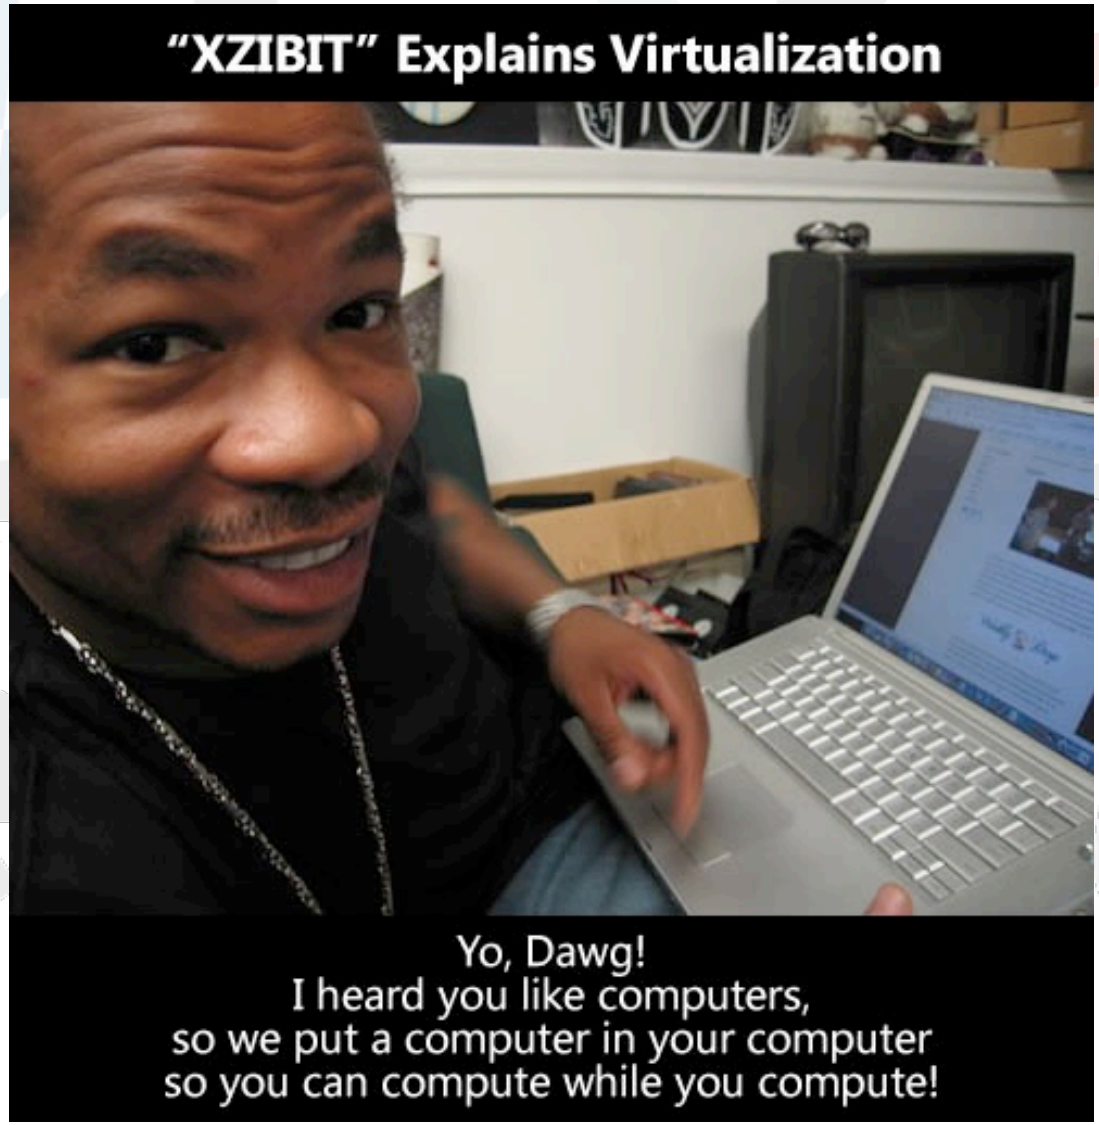

ITaP
INFORMATION TECHNOLOGY AT PURDUE

Condor inside Condor!

- Bare-bones Linux VMs launched as Condor VM Universe jobs
  - Require jobs to be self-contained
  - A great way to do a Denial of Service attack

PURDUE UNIVERSITY

Making life better

- Pre-staging vmdk files conserves bandwidth
  - Job starts require KB of network bandwidth instead of GB
- Replace IPOP with CCB
  - VM-based jobs start faster
  - No extra service required

ITaP INFORMATION TECHNOLOGY AT PURDUE

What's next?

- Wider deployment of VMs
  - Consistent pool of running VMs
  - Standalone (non-Condor) VM management tool
  - Bolt-on application disks
  - Running VMs at partner institutions
- Dynamic VM pool size
  - If a VM is running, the host can't hibernate
- Finding a user to run lots of code over a weekend
  - Management loves sending out press releases

Videos

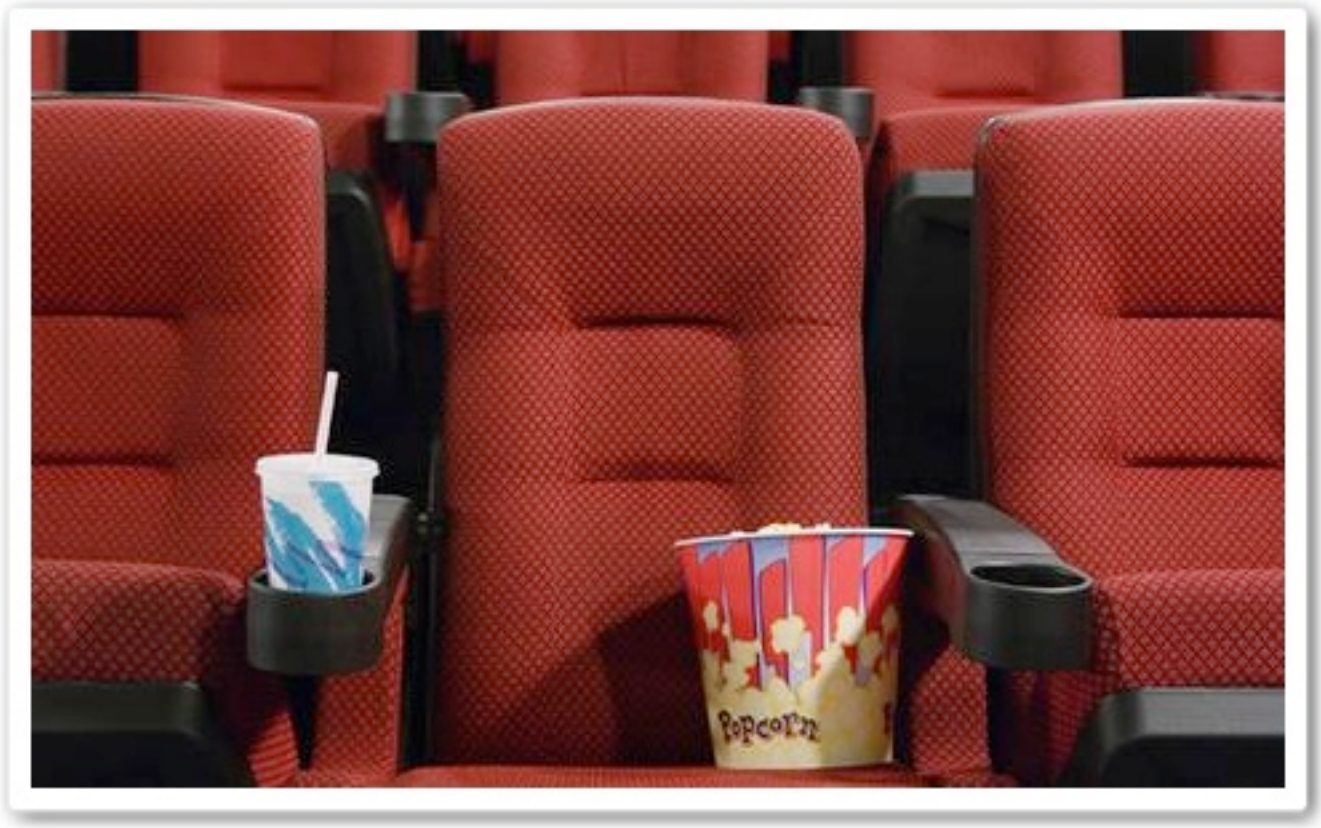

PURDUE UNIVERSITY

Doubletake

Condor Week 2011

ITaP
INFORMATION TECHNOLOGY AT PURDUE

Doubletake – YouTube at Purdue

- Student-submitted video assignments
  - Great for ASL assignments
- Jobs grabbed from web submission and rendered with Condor and ffmpeg
  - Preempts jobs on the two dedicated nodes
  - Can flock to the rest of the Condor pools
- Jobs shoved to website when rendering is complete
  - 10 MB/minute encoding times

PURDUE UNIVERSITY

DoubleTake

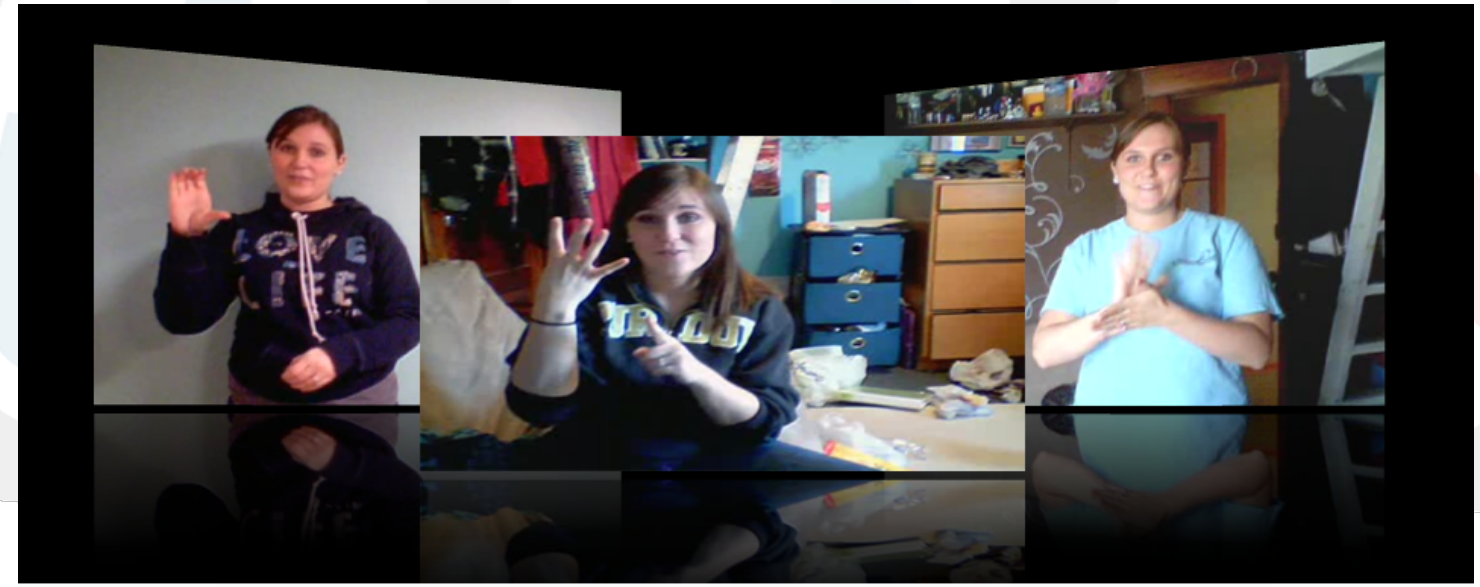

ADD ROW

ADD COLUMN

		✖	✖	✖		
50	Spelling Criteria Demo	1 Lowest	2 Middle	3 High	↕	✖
25	Grammar Grammar	1 Lowest	2 Middle	3 High	↕	✖
25	Fluency Fluency	1 Fluenc	2 Middle	3 High	↕	✖

ITaP
INFORMATION TECHNOLOGY AT PURDUE

The Future


Huh... that's weird.
This doesn't LOOK
like the future...
ICANHASCHEEZBURGER.COM

The Future

- More cores!
  - Another cluster coming in the fall
- More DiaGrid participants
  - Ball State University, University of Delaware
- More VMs
  - Application/domain specific virtual machines?
- Need more users
  - What are we going to do with all of these spare cycles?

We're hiring!

- DiaGrid job opportunities coming
  - Condor expert/enthusiast
  - Experience working with researchers
  - Strong problem solving
  - Parallel programming skills preferred
- www.purdue.edu/jobs
  - (positions are not yet posted)

The End

Questions?