

Supporting Campus Researchers

David Swanson
Holland Computing Center

Talk Outline

- Share a (brief) collection of experiences
- A methodology
- Offer a few generalizations

HCC Context

- University system-wide provider of HPC, HTC
- Facilities in Omaha (10,000 cores, 500 TB) and Lincoln (5,000 cores/slots, 1 PB)
- 30 gbps between centers
- campus grid, OSG
- campus champions

Aaron Dominguez and Ken Bloom

- Coming to campus, call about a Tier2 site
- Would be 50/50 hardware/personnel
- Meeting in Iowa on July 22, 2006
 - (thank you Mrs. Swanson)
 - first face to face meeting with Aaron
- Submit proposal, site visit, selected
- quickly included Carl, Brian, several others

Mutually Beneficial Arrangement

- Researchers buy into infrastructure and support staff (*Priority Access*)
- HCC operates the facility, helps researchers use it (\$50/node/month)
- Opportunistic use by rest of campus
- Continued and growing support as more funded projects develop and subsequently collaborate and contribute in turn

Priority Access

- Climatology (WRF)
- Mechanical Engineering (LS-Dyna)
- Software Engineering (AFOSR)
- NanoScience (EPSCoR)
- AMO Physics
- Proteomics
- Ed Psych

Neethu Shah

- Identifying protein homologues
- *Cluster and Grid Computing course project*
- worked with Brian, used glidein
- now meeting monthly with her research group (Moriyama)
- Poster

Brian Pytlik-Zilig

- Digital Humanities research
- *Course Project MR of large corpus*
- White-board sessions, Kyle, Brian, Adam, Ashu, me
- switch from MR to Condor DAGMAN
- Still under development ... but funded (!)
- Plenary

Bob Powers

- CPASS: Comparison of Protein Active-Site Structures
- *Came asking for help (!)*
- White board sessions, Bob, Jennifer, Ashu, Adam, me, others
- Set up LVS for http transfers, SVN for code
- Poster (Jennifer Copeland)

Shi-Jian Ding

- analyzing Mass Spectra to decipher protein structure
- *Met at UNMC open house*
- Later swapped talks at group meeting (Shi-Jian, several students, Ashu, Adam, me)
- Ashu configured OMSSA, requires SRM
- Poster (Hong Peng)

Steven Massey

- Computing robustness of a given population
- *Met at Starbucks in San Juan with PR physicist*
- Met local HPC staff at EPSCoR meeting, discuss Condor, Campus Grids, Gratia
- Several teleconferences, a few skypes, IM with Jose Medina (*and Caballero!*)
- Yaling used osg-xside to submit 1000s of jobs (thank you Mats Rynge)
- Poster (Yaling Zheng)

HCC Triage

- What are you doing now?
 - research area
 - computing approach
- Is there some way we could help?
 - team approach
 - scale up or scale out

HCC Triage

- Can it be run as an OSG job?
- Campus Grid job?
- Cluster only?

HCC Triage

- What can we: start today?
- ... do in a week?
- ... do this month?
- How do we find a mutual no-loss scenario, with possible big win?
- Are they invested?

No loss is no loss

- If we deliver what we promise, we earn some trust and good will (Matt/CPASS)
- If we help even though it is not directly beneficial to HCC, we earn some trust (Janos/NPOD)
- It is very difficult to predict the most successful projects ... so try them all

Acknowledgements

- NU administration, NRI, Holland Foundation
- NSF, EPSCoR
- OSG, UW, Purdue
- DoE, FNAL
- OR, I2, IS




Extra slides

HW vs SW Scaling

- Now 64 cores/node
- code scaling not increasing at same rate
- we're not a "largest job next" shop

Wall Hours by Number of Cores in the Job (Sum: 49,048,618 Hours)

52 Weeks from Week 11 of 2011 to Week 11 of 2012



- 1 Core (17,673,646)
- 32 Cores (8,284,867)
- Other (4,574,582)
- 128 Cores (3,422,172)
- 16 Cores (2,895,544)
- 64 Cores (2,828,036)
- 256 Cores (1,468,437)
- 105 Cores (1,132,682)
- 8 Cores (1,056,589)
- 384 Cores (645,049)
- 20 Cores (609,278)
- 4 Cores (594,535)
- 12 Cores (580,404)
- 5632 Cores (568,309)
- 136 Cores (519,398)
- 80 Cores (517,065)
- 6 Cores (447,991)
- 24 Cores (428,242)
- 48 Cores (421,768)
- 7 Cores (380,025)

Relative prices

- 256 GB RAM (\$3200)
- 4 6272 procs (\$2200)
- IB card (\$550)

Operating Principles and Policies

- Resources Priority Access, Shared or Opportunistic
- Opportunistic use of Priority Access resources (preempted as necessary) -- this extends to Grid resources
- Shared resources FairShare per Research Group -- very short half-life (1 day)

Operating Principles and Policies

- NU researchers have first priority
- Grid jobs opportunistic
- Students involved at all levels as appropriate