

Research Technologies Monthly Roundtable

*Batch Queueing Systems on
BigRed and Quarry*

George Wm Turner

28-May-2009



INDIANA UNIVERSITY



INDIANA UNIVERSITY

Getting Started

Indiana University Knowledge Base

<http://kb.iu.edu/>

Hint: FireFox Search engine setup

Getting Started on BigRed

<http://kb.iu.edu/data/avjx.html>

At IU, what is BigRed

<http://kb.iu.edu/data/aeo.html>

Getting Started on Quarry

<http://kb.iu.edu/data/avkx.html>

At IU, what is Quarry

<http://kb.iu.edu/data/avju.html>



BigRed & Quarry: similarities and differences

Similar

Linux (SLES vs RHEL)

GPFS, Data Capacitor

Home Directories

Nodes: Head, WildWest, Debug, Compute

Scheduler: Moab

Different

BigRed: Power PPC750 2.5 GHz, 4 cores, 8 GB RAM

“parallel”, Myrinet switch fabric, ~1024 nodes*

resource manager: LoadLeveler

Quarry: Intel Xeon E5375 2.0GHz, 8 cores, 8 GB RAM/16 GB himem

“serial”, GigE ethernet fabric, ~140 nodes

resource manager: TORQUE



Queue structures

How they relate to infrastructure

BigRed

Big Red usage policies

<http://kb.iu.edu/data/avgm.html>

Quarry

What is TORQUE/PBS

<http://kb.iu.edu/data/avmy.html>



Queue structures

How they relate to infrastructure

BigRed: <http://kb.iu.edu/data/avgm.html>

SERIAL	max 2 days, 1 node
NORMAL	max 2 days, 256 nodes
LONG	max 14 days, 32 nodes
DEBUG	max 15 minutes, 4 nodes
IEDC	max 14 days, 256 nodes



Queue structures

How they relate to infrastructure

Quarry: <http://kb.iu.edu/data/avmy.html>

batch	default queue, autosort
serial	max 12 hours, 1 node
normal	max 7 days, 2 nodes
long	max 14 days
himem	max 14 days (**not autosort yet**)
debug	max 30 minutes, 2 nodes
osg	Open Science Grid restricted



Resource managers verses schedulers

Migration between IBM Loadleveler and TORQUE

<http://kb.iu.edu/data/axpz.html>

What is Moab

<http://kb.iu.edu/data/avmu.html>

LoadLeveler

IBM, commercial

hint: ll<tab> then man XYZ

TORQUE/PBS

Open Source, derived from PBS/OpenPBS/PBSpro

hint: man qsub

Moab

commercial, derived from Maui



INDIANA UNIVERSITY

Running

Interactively

Head nodes: 20 minutes, please refrain

WildWest nodes: free for all; wheee!

BigRed:

Quarry: `qsub -l -X`

Mass production batch processing



INDIANA UNIVERSITY

Scheduling strategies

Fairshare

Backfill

Dispatch tactics



INDIANA UNIVERSITY

Queue packing

I just wanna run!



INDIANA UNIVERSITY

Packed Queue

When will I run?

Moab: Scheduler issue

showq, showq -i, showstart, showbf

Common Moab scheduler commands

<http://kb.iu.edu/data/axqa.html>



INDIANA UNIVERSITY

Debugging

It won't run!

WildWest nodes

Submit to DEBUG queue

Quarry: `qsub -l`

BigRed: `call`



INDIANA UNIVERSITY

Help

IU KB

<http://kb.iu.edu/>

Indiana University's Advanced CyberInfrastructure: The Least You Need To Know

<http://rc.uits.iu.edu/documentation/Cyberinfrastructure.pdf>

Where can I get help using IU's supercomputers?

<http://kb.iu.edu/data/alfa.html>

Systems: High Performance Systems (HPS)

<http://rc.uits.iu.edu/hps/index.shtml>

hps-admin@iu.edu

Applications: High Performance Applications (HPA)

<http://rc.uits.iu.edu/hpa/index.shtml>

hpa-admin@iu.edu

StatMath Center

<http://www.indiana.edu/~statmath/>

statmath@iu.edu

IUB phone 812-855-4724, IUPUI phone 317-278-4740

Distributed Data Storage @ IU

<http://storage.iu.edu/>

store-admin@iu.edu

Research Technologies Division

<http://rtinfo.uits.indiana.edu/>

researchtechnologies@iu.edu

Google ☺