2016 COMMUNITY CLUSTER

PURDUE UNIVERSITY

Preston Smith

Director of Research Computing Services

Michael Shuey

Infrastructure Architect

8/30/2016

HPC FACULTY MEETING

COMPUTATION

RESEARCH COMPUTING STRENGTH

Since Steele in 2008, Research Computing has deployed many world-class offerings in computation

PURDUE UNIVERSITY

7 COMMUNITY CLUSTERS

STEELE

7,216 cores, Installed May 2008

Retired Nov. 2013

COATES

8,032 cores, Installed May 2008

24 departments, 61 faculty

Retired Sep. 2014

ROSSMANN

11,088 cores

Installed Sept. 2010

17 departments

37 faculty

Retired Sep. 2015

HANSEN

9,120 cores

Installed Sept. 2011

13 departments

26 faculty

Retiring Oct. 2016

CARTER

10,368 cores

Installed April 2012

26 departments

60 faculty

#54 on June 2012 Top 500

CONTE

9,280 Xeon cores
(69,600 Xeon Phi cores)

Installed August 2013

20 departments

51 faculty (as of Aug. 2014)

#28 on June 2013 Top 500

RICE

13,200 cores

Installed May 2015

17 departments

70 faculty

#105 STEELE 2008

#102 COATES 2009

#150 HANSEN 2010 (est.)

#126 ROSSMANN 2011

#54 CARTER 2012

#28 CONTE 2013

#166 RICE 2015

PURDUE COMMUNITY CLUSTERS TOP500 RANKINGS

PURDUE UNIVERSITY

PARTNERS

238M hours delivered in 2015

185 investors from 36 departments

PURDUE UNIVERSITY

FACULTY DEMOGRAPHICS

Department	Cores
Electrical and Computer Engineering	9816
OSG CMS Tier2	9168
Mechanical Engineering	7008
Aeronautics and Astronautics	5048
Earth, Atmospheric, and Planetary Sciences	3632
Chemistry	1936
Materials Engineering	1504
Chemical Engineering	1144
Biological Sciences	1104
Medicinal Chemistry and Molecular Pharmacology	1104
Mathematics	720
Physics	664
Biomedical Engineering	640
Statistics	520
Nuclear Engineering	492
Civil Engineering	448
Agricultural and Biological Engineering	416
Industrial and Physical Pharmacy	384
Commercial Partners	304
Computer Science	280
Other College of Agriculture	256
Agronomy	240
Forestry and Natural Resources	64

BID PROCESS

Open bid process:

- Quantity approx. 600 nodes
- Included various interconnects (FDR/EDR Infiniband)
- At least 20- core compute nodes
- Memory size for base node 128GB
- SSD boot drive

Prices ranged from $3500-4000 per node

Vendors included Dell, HP, Lenovo, etc.

BID RESULTS

HARDWARE SPECS

Base node: HP XL170

- 20-core node, 2.6 GHz Intel “Haswell” processors (E5-2660v3)
    - Same processor as Rice cluster
    - Still a good balance between cost-effectiveness and overall node price
    - Explored “Broadwell” chips, but these incurred a ~20% drop in clock speed
- 128 GB DDR4 memory (minimum)
    - 512 GB & 1 TB options available shortly after Rice launch
- 480 GB local SSD
    - No longer using spinning disks in nodes
- 25 Gbps Ethernet on all nodes for IP, home directory, Depot access
- EDR Infiniband interconnect
    - 100 Gbps, 2:1 fat tree – evolutionary improvement from Rice interconnect

Scratch system:

- Approximately 1 PB, 25 GB/sec
- SSD acceleration for smaller files

PURDUE UNIVERSITY

IMPLEMENTATION

General implementation schedule:

- Facilities preparation begins immediately following the October 1 retirement of Hansen
- Benchmarking and burn-in late November
- General Availability Dec 1

- It’s cold out and the datacenter is crowded – no install event this year.

PURDUE UNIVERSITY

THE BOTTOM LINE

NEW CLUSTER

Same processors as Rice

- 2x the memory
- 2x the IB bandwidth
- 2.5x the network bandwidth
- SSD local disk

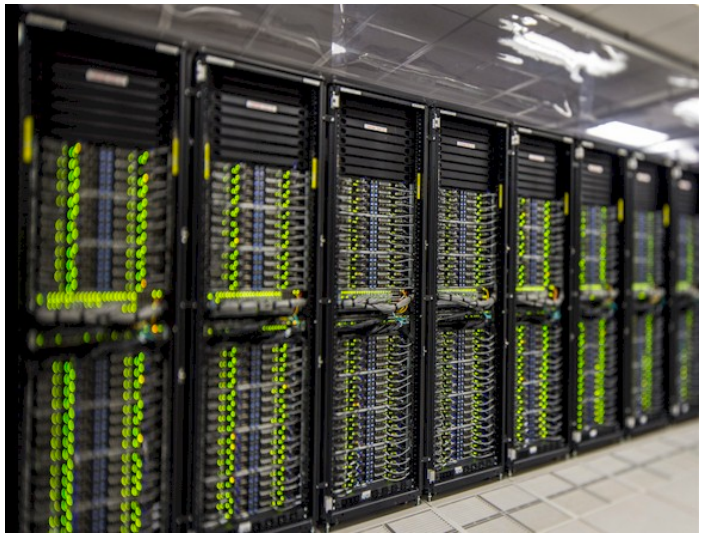

http://2eof2j3oc7is20vt9q3g7tlo5xe.wpengine.netdna-cdn.com/wp-content/uploads/2013/10/bp-cluster.jpg

PURDUE UNIVERSITY

FOR THE MAJORITY:

HPC SYSTEM

Halstead: A traditional HPC system very similar to Rice

The same, familiar model:

- New cluster acquisition every year
- Each a distinct, non-heterogeneous system.
- Service through December of 2021

Great for most science and engineering codes

PURDUE UNIVERSITY

OPTIONS

HARDWARE CHOICES

Base node option, **plus**

- Node with 2 Nvidia GPUs
- Large memory options, with 512G and 1TB configurations

128G Node	512G Node	1T Node	128G Node + GPU
$3,800.00	$6,400.00	$12,200.00	$11,000.00

Some capacity remains in Rice, and is priced to move!

$2250/node

Service through 2020

PURDUE UNIVERSITY

PURDUE
UNIVERSITY

RESEARCH DATA STORAGE

DATA IS HUGE!

AND GROWING

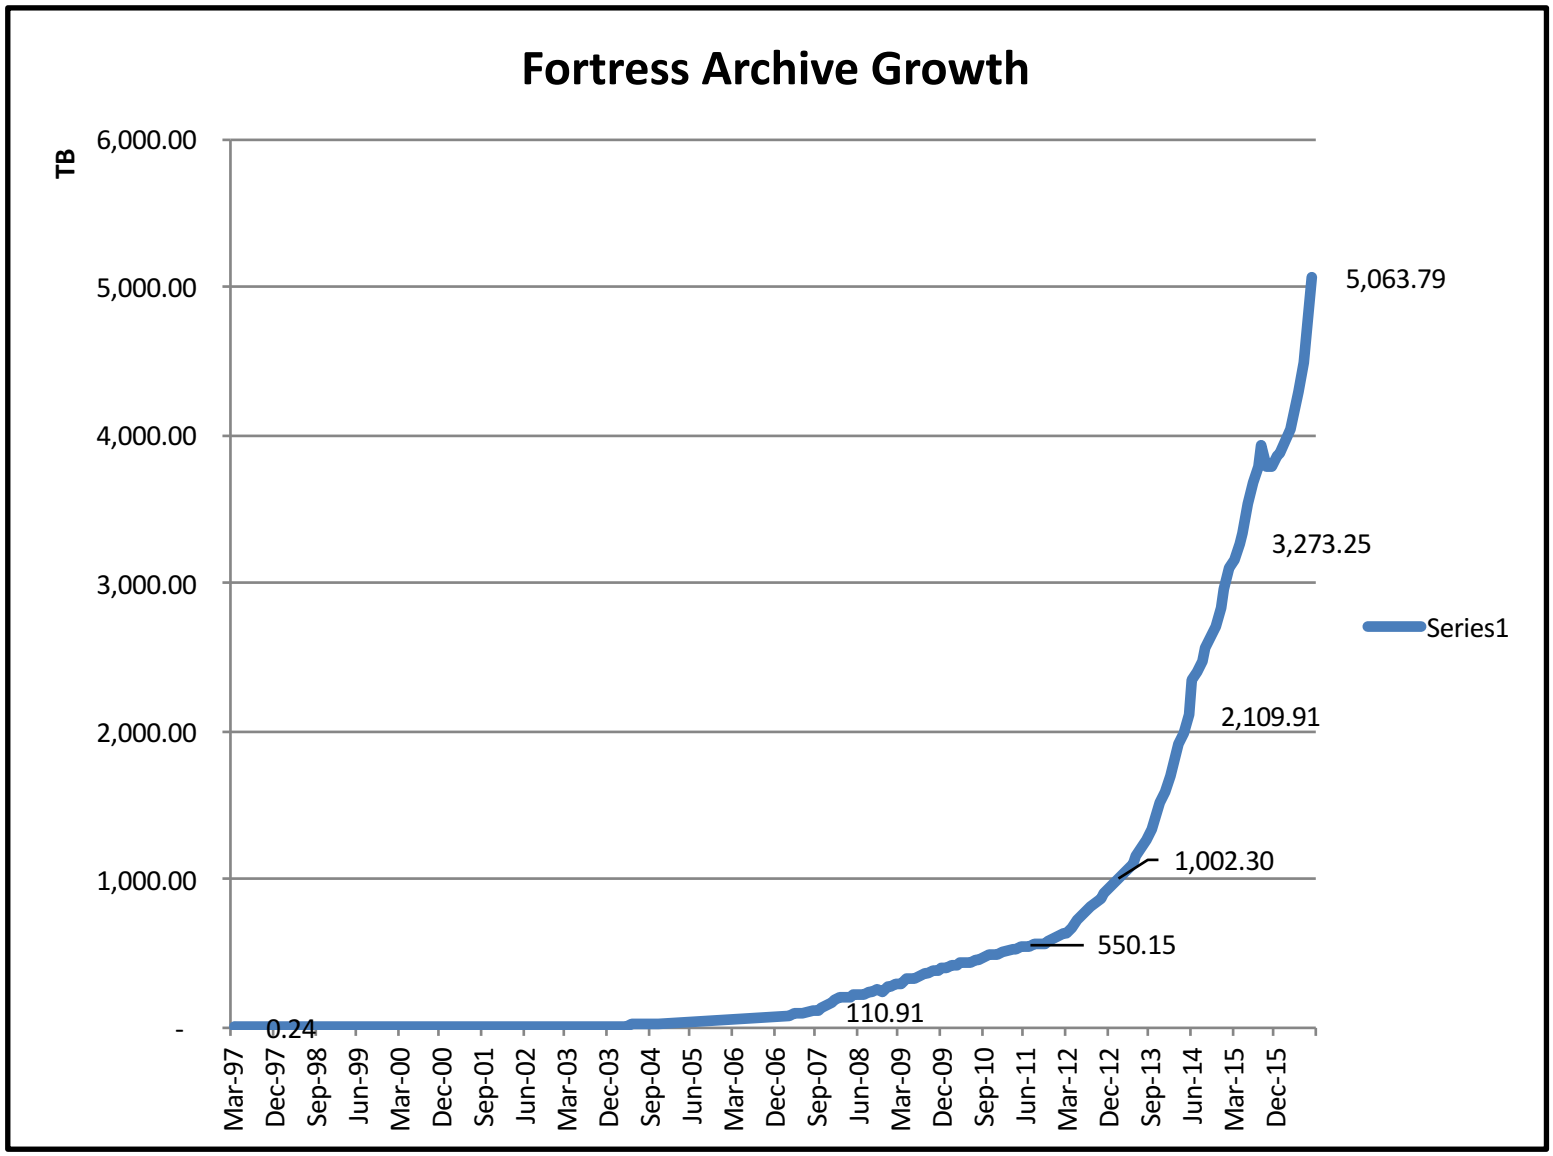

PURDUE UNIVERSITY

NEW FILE SERVICES

Home directories

- September installation
- Entirely devoted to research computing!

Improved scratch for Life Science users on Snyder

- No longer sharing scratch system with Rice
- Optimized for Snyder's distinct usage patterns

RESEARCH DATA DEPOT

SERVICE DETAILS

At $150/TB per year:

- Storage oriented around your research lab, with
  - Snapshots
  - Multi-site copies of your data
  - Disaster protection
  - A scalable, expandable storage resource optimized for HPC
- Access to Globus data transfer service, and endpoint sharing

PURDUE UNIVERSITY

GET ACCESS

GIVE IT A TRY!

To buy 1 or more TB of space,

Or to set up a trial for your lab

Order online:

https://www.rcac.purdue.edu/purchase/

PURDUE UNIVERSITY

GLOBUS

STATISTICS

Data moved in 2015:

300 TB transferred

Average of 23 TB, 50 unique users per month

https://transfer.rcac.purdue.edu

PURDUE UNIVERSITY

PURDUE
UNIVERSITY

RESEARCH SERVICES

OTHER SERVICES YOU MIGHT BE INTERESTED IN

COMPUTATIONAL SUPPORT

EVERYTHING PAST HERE IS SPARE

DOMAINS

Chemistry

Physics

Astrophysics

Earth and Atmospheric Sciences

Computer Science

Chemical Engineering

Electrical and Computer Engineering

Cell and Molecular Biology

Agriculture

APPLICATION SPACES

Molecular Dynamics

Image Processing

Quantum Chemistry

Weather Modeling

Machine Learning

Big Data

Computer Architecture

Finite Element Analysis

Statistics

Bioinformatics

Geospatial

Remote Sensing

Visualization

PURDUE UNIVERSITY

NEED HELP?

- Hard to solve problems with HPC?
- Need help building your software or optimizing your workflow?
- Need to learn what resources are available?

COFFEE BREAK CONSULTATIONS

PURDUE UNIVERSITY

SCHOLAR

HPC FOR INSTRUCTION

- Need to teach students to use HPC or work with big data in a course?
- Scholar cluster is available to any instructor at no cost.

Spring 2016:	EAPS
CS	AAE
STAT	ANSC
CHEM	ME

Just send a CRN

PURDUE UNIVERSITY

VERSION CONTROL

NEED GIT OR SVN?

Local-to-Purdue Github repositories for your lab, managed by you!

PURDUE UNIVERSITY

TRAINING

COMING EVENTS

XSEDE Workshop: MPI

Sept 7-8, 11am – 5pm

Unix 101 – Part I

Sept 12, 14, 16, 1pm-4pm

Unix 101 – Part II

Sept 26, 28, 30, 1pm-4pm

XSEDE Workshop: OpenMP

Oct 4, 11am-5pm

Clusters 101

Oct 17, 19, 21, 1pm-4pm

Software Carpentry

Late Oct / Early Nov

Other sessions in the works!

Talk to us about working directly with your lab!

PURDUE UNIVERSITY

THE END

QUESTIONS?

Questions?

PURDUE UNIVERSITY