



PURDUE
UNIVERSITY

Preston Smith

Director of Research
Services

July 2, 2015

INTRODUCTION
TO RESEARCH
SERVICES

RESEARCH COMPUTING AT PURDUE UNIVERSITY

RCAC Staffing

<https://www.rcac.purdue.edu/about/staff/>

WHO ARE WE?

OVERVIEW

- IT Research Computing (RCAC)
- A unit of ITaP (Information Technology at Purdue) – the central IT organization at Purdue.
- RCAC provides advanced computational resources and services to support Purdue faculty and staff researchers.

Our goal: To be the one-stop provider of choice for research computing and data services at Purdue -Delivering powerful, reliable, easy-to-use, service-oriented computing and expertise to Purdue researchers.



PURDUE
UNIVERSITY

COMMUNITY

CLUSTERS

A BUSINESS MODEL FOR HPC AT
PURDUE UNIVERSITY

THE POWER OF SHARING

- Without a large capital acquisition by the university, providing cutting-edge computing capabilities for researchers was not possible.
- Many faculty were getting funding to acquire and operate HPC resources for themselves
- **Solution:** pool these funds to operate clusters for researchers!
 - The faculty no longer have to devote a grad student to managing their cluster!

VERSION 1:
THE BASIC RULES

- You get out at least what you put in
 - Buy 1 node or 100, you get a queue that guarantees access up to that many CPUs
- But wait, there's more!!
 - What if your neighbor isn't using his queue?
 - You can use it, but your job is subject to preemption if he wants to run.
- You don't have to do the work
 - Your grad student gets to do research rather than run your cluster.
 - Nor do you have to provide space in your lab for computers.
 - ITaP provides data center space, systems administration, application support.
 - Just submit jobs!

SIX COMMUNITY CLUSTERS

STEELE

7,216 cores
Installed May 2008
Re9red Nov. 2013

COATES

8,032 cores
Installed July 2009
24 departments 61
faculty **Re9red**
Sep. 2014

ROSSMANN

11,088 cores
Installed Sept. 2010
17 departments 37
faculty

HANSEN

9,120 cores
Installed Sept. 2011
13 departments 26
faculty

CARTER

10,368 cores
Installed April 2012
26 departments 60
faculty
#175 on June 2013 Top 500

CONTE

9,280 Xeon cores
(69,600 Xeon Phi cores)
Installed August 2013
20 departments
51 faculty (as of Aug. 2014)
#39 on June 2014 Top 500

VITAL STATS

- 165 “owners”
- ~1200 active users
- 259M hours provided in 2014
- Nationally, the gold standard for condo-style computing

- Today, the program is part of many departments’ faculty recruiting process.
 - ***A selling point to attract people to Purdue!***
 - ***Please feel free and have your faculty candidates meet with us during recruitment!***

IMPACT




FACULTY PARTNERS

Department	Cores
Electrical and Computer Engineering	9816
OSG CMS Tier2	9168
Mechanical Engineering	7008
AeronauNcs and AstronauNcs	5048
Earth, Atmospheric, and Planetary Sciences	3632
Chemistry Materials	1936
Engineering Chemical	1504
Engineering Biological Sciences	1144
Medicinal Chemistry and Molecular Pharmacology	1104
MathemaNcs	720
Physics	664
Biomedical Engineering	640
StaNsNcs	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

IMPACT

HPC USERS AND SPONSORED DOLLARS



**NEW MODEL – ORGANIZED
BY COMMON PROFILES**

COMPUTATION

Community Clusters to Cluster Communities

What neighborhoods are in our community?

HPC (Rice): Multiple cores or nodes, probably MPI. Benefit from high-performance network and parallel filesystem. The vast majority of campus -80% of all work!

HTC (Hammer): Primarily single core. CPU-bound. No need for high-performance network.

Life Science/Big Memory (Snyder):
Use entire node to get large amounts of memory. Less need for high-performance network. Needs large, fast storage.

DATA

STORAGE

INFRASTRUCTURE FOR RESEARCH DATA

WHAT IS AVAILABLE
TODAY FOR HPC

DATA STORAGE

- Research computing has historically provided some storage for research data for HPC users:
 - Archive (Fortress)
 - Actively running jobs (Cluster Scratch - Lustre)
 - Home directories

... And Purdue researchers have PURR to package, publish, and describe research data.

THE SERVICE

HPC researchers can at last purchase storage!

A storage service for research to address many common requests:

- 100G available at no charge to research groups
- Mounted on all clusters and exported via CIFS to labs
- *Not scratch*: Backed up via snapshots, with DR coverage
- Data in Depot is owned by faculty member!
- Sharing ability – Globus, CIFS, and WWW
- Maintain group-wide copies of application software or shared data

Well received!

- In less than 7 months, over 105 research groups are participating.
 - *Many are not HPC users!*
- Half a PB in use since November
- A research group purchasing space has purchased, on average, 8.6TB.


WHAT DID WE GET?

Approximately 2.25 PB of IBM GPFS

Hardware provided by a pair of Data Direct Networks SFA12k arrays, one in each of MATH and FREH datacenters

160 Gb/sec to each datacenter

5x Dell R620 servers in each datacenter



WHAT DO WE NEED TO DO?

DESIGN TARGETS

The Research Data Depot Can do:

Depot Requirements	Previous solutions
At least 1 PB usable capacity	>1 PB
40 GB/sec throughput	5 GB/sec
< 3ms average latency, < 20 ms maximum latency	Variable
100k IOPS sustained	55k
300 MB/sec min client speed	200 MB/sec max
Support 3000 simultaneous clients	Yes
Filesystem snapshots	Yes
Multisite replication	No
Expandable to 10 PB	Yes
Fully POSIX compliant, including parallel I/O	No

GUIDING PRINCIPLES

DATA

- It's important to think of Depot as a “data service” – not “storage”
- It is not enough to just provide infrastructure
 - “Here’s a mountpoint, have fun”
- Our goal: enabling the frictionless use and movement of data
 - *Instrument -> Depot -> Scratch -> Fortress -> Collaborators -> and back*
 - Continue to improve access to non-UNIX users

HOW CAN I MANAGE ALL MY DATA?

- Collaborations on multi-disciplinary grant proposals, both internal and external
- Developing customized Data Management Plans
- Organizing your data
- Describing your data
- Sharing your data
- Publishing your datasets
- Preserving your data
- Education on data management best practices

OTHER

SERVICES

BEYOND THE COMMUNITY CLUSTERS

DEVELOPMENTS IN CAMPUS NETWORK

2014 network improvements

- 100 Gb/sec WAN connections
- Research Core
 - 160 Gb/sec core to each resource (up from 40)
 - 20 Gb/sec research core to most of campus
- Campus Core Upgrade

<https://www.rcac.purdue.edu/news/681>

EVERYBODY NEEDS
TO SHARE

GLOBUS

Globus:

Transfer and share large datasets....

.... With dropbox-like characteristics

.... ***Directly from your own storage system!***



GLOBUS

Data moved in 2014:

13 TB in, 19TB out

200k files both directions

55 unique users

In progress:

Globus interface to Fortress

[h" ps://transfer.rcac.purdue.edu](https://transfer.rcac.purdue.edu)

EDUCATION

- Programming practices – Software Carpentry
- Parallel Programming – MPI, OpenMP
- Big Data
- Matlab
- Accelerators – Xeon Phi, OpenACC, CUDA
- UNIX 101
- Effective use of Purdue research clusters

NEED HELP?

COFFEE BREAK CONSULTATIONS



RESEARCH COMPUTING COFFEE BREAK CONSULTATIONS

Meet up with ITaP research computing staff and other researchers who use or are interested in High Performance Computing at Purdue. Join us for informal discussions of scientific computing along with any other topic that might spring up. We'll be meeting at different coffee shops around campus each week.

Check the coffee web page to see this week's location and time.


rcac.purdue.edu/coffee

HPC FOR INSTRUCTION

SCHOLAR

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

Spring 2015: EAPS
CS AGRY
STAT ANSC
CHEM ChemE



NEED A PROGRAMMER?

Bring in our expertise to help your researchers create or modify software to take advantage of the latest technology in advanced computation, web frameworks, data analysis, visualization, sharing, and management.

Our software development effort can be funded through grant awards or contracts based on developer time.