

RESEARCH TECHNOLOGIES

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

rt.iu.edu



RESEARCH TECHNOLOGIES

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Supercomputing Innovation at IU

HPE Innovator's Breakfast

RT is a division of UITS and affiliated with the Pervasive Technology Institute

David Y. Hancock

Primary Investigator – Jetstream Cloud

Director, Advanced Cyberinfrastructure

SC19 – November 20, 2019



7 University Campuses
9 Medical Education Centers
2 Academic Centers

Indiana University, est. 1820

- **\$3.7B** enterprise
- Partnered with **\$6.4B** IU Health system
- **94,000** Students
- **1.1M** credit hours per semester
- **>20,000** degrees per year
- **\$950M** in financial aid
- **\$604M** in research grants
- **20,000** faculty and staff
- **7,873** Acres
- **949** owned and leased buildings, **36M** square feet
- **>690,000** living alumni

*Almost completed
2009 Network Master
Plan for IUB/IUPUI 361
buildings <\$100M
(Orig Budget \$172M)*

Multi-institution Partnerships

Networks



National network collaboration of 400+ Universities and national labs. IU acts as network operations center for the organization.



Indiana state network for higher education with over 100 members. IU acts as its network operations center.



Collaboration of I-Light, IU, Notre Dame, and Purdue providing high-speed, high-availability, feature-rich network to Indiana's higher education institutions.

Libraries



Began as collaboration between the Big Ten Academic Alliance and the University of California system to digitize library print holdings. Now includes over 150 universities. IU serves as primary backup site for repository.

Research Technologies



IU-led collaboration with Texas Advanced Computing Center, University of Chicago, University of Arizona, and University of Texas (SA) to develop cloud based tools for scientific research funded by National Science Foundation.



Collaboration with Texas Advanced Computing Center, Pittsburgh Supercomputing Center and San Diego Supercomputing Center to analyze massive genomic data.



Collaboration of IU, Purdue, and Notre Dame, with public and private partners to facilitate the translation of scientific discoveries in the lab into clinical trials and new patient treatments.

Learning Technologies



Coalition of universities dedicated to collaborative digital education. Founded by IU, Michigan, Colorado State, and Florida. Now includes 25 institutions.

Security



Research and Education Networks Information Sharing and Analysis Center. Includes 620 member universities. One of 19 nationally recognized industry-specific ISACs.



Center for Applied Cybersecurity Research, Est. 2003, integrates applied research in cybersecurity technology, education, and policy guidance.



Shared Cybersecurity Operations Center. Founded by IU, Nebraska, Northwestern, Purdue, and Rutgers.



ResearchSOC provides cybersecurity services to NSF-funded facilities and projects, such as Gemini Observatory, UNAVCO and GAGE, and the National Radio Astronomy Observatory.

Enterprise Systems



Community sourced software suite driving down cost of enterprise systems. IU co-founded and currently serves on board of directors. Includes over 50 university members.



Portal replacement service discovery application developed at IU. Currently used by 85 university campuses.

Service & Support



Application developed at IU to integrate cloud storage systems into a single place. Currently has 8 university subscribers and 6 universities testing.



System to manage software licenses developed at IU. Subscribers include University of Maryland and University of Alabama-Huntsville.

UITS, PTI, and RT

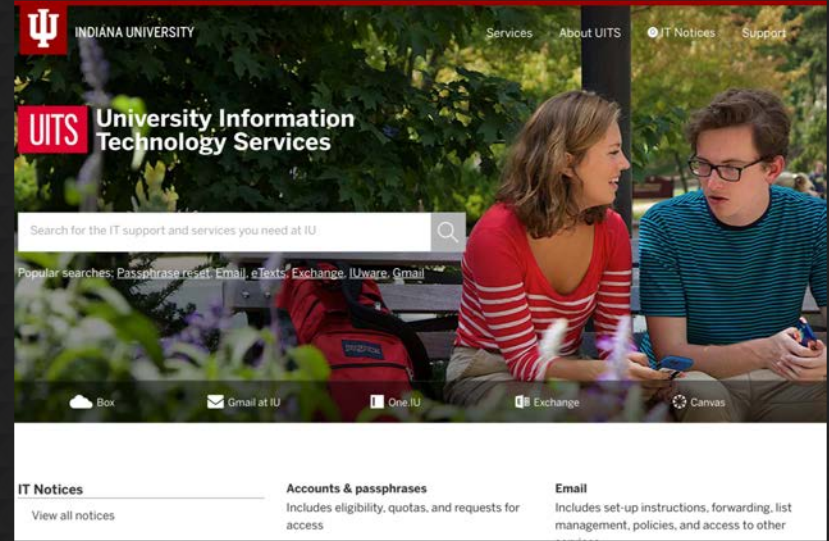
Research Technologies is part of...

UITS (University Information Technology Services) as a distinct division

PTI (Pervasive Technology Institute) as a center

Both report to Brad Wheeler, Vice President for IT & CIO

RT is led by Matt Link; PTI is led by Craig Stewart





RESEARCH TECHNOLOGIES

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

- Mission is to support research
- Deliver a wide range of services:
 - Advanced Visualization
 - High Performance Computing
 - High Performance Storage
 - Research Software & Solutions
 - Education and Outreach
 - HIPAA & FISMA/CMS
- Contact rt.iu.edu
- Our Impact rt.iu.edu/impact/
- Partner with faculty on grants.
 - 52 grants totaling \$51M in the past 10 years.





RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



Partnering in Research

- IU researchers were the first members of ADNI to align and detect SNPs in full genome sequencing of 818 Alzheimer's patients.
- We reduced the runtime of the pipeline by 25% and enabled checkpointing




RT Supports Everyone at IU

153 different departments and over 331 disciplines use IU's advanced research cyberinfrastructure provided and supported by RT

Breadth of impact on IU research, scholarship, and creative communities.

Over 1,000 IU researchers, who use IU supercomputers or other HPC tools, received nearly \$313 million in grant awards in 2018.






History of Computing at IU



One view of the IU System



Grids at Indiana University



What's our “recent” track record?

Supporting research with multiple HPC systems and storage for 20 years

- Sustained investment in computation
- Growth of storage (GPFS, Lustre, Ceph)
- Use of HPSS (Disk/Tape archive) over the entire period

What's this mean for the future?

- Continued support for large-scale, large-memory, and high-throughput workloads by offering multiple systems
- Continued storage growth with no-fee tiers and large scratch space but some chargeback necessary
- Continued involvement with access to national cyberinfrastructure through the eXtreme Science and Engineering Discovery Environment (XSEDE)



Research Desktop (RED): a friendly gateway to HPC

<https://kb.iu.edu/d/apum>

- Making supercomputers more user friendly
 - Research Desktop provides a new way to login and interact with HPC
 - A GUI/desktop instead of a terminal

New: Web browser access and a dedicated service for apps like Jupyter and Rstudio

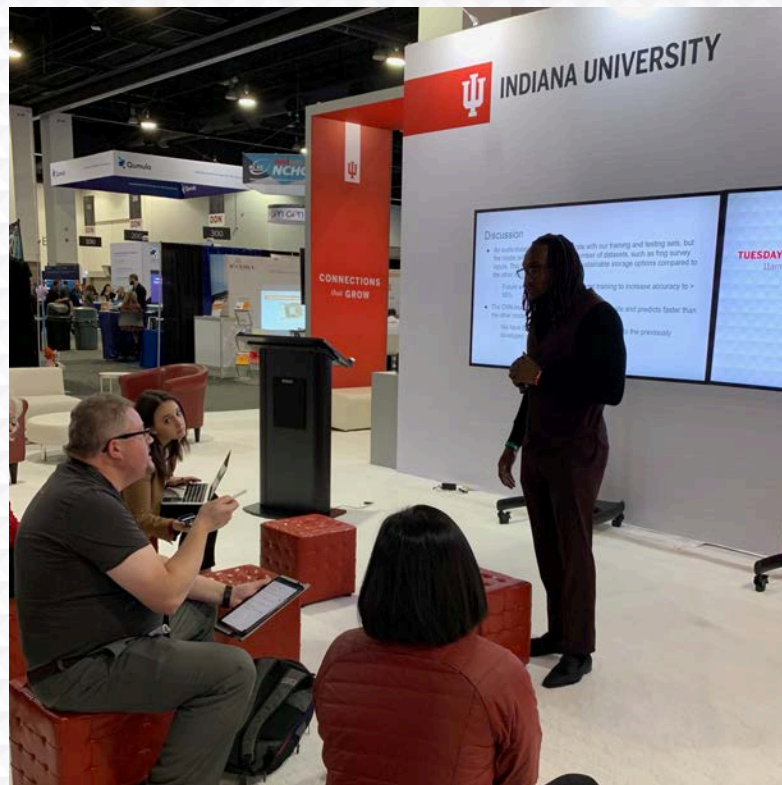
- Research Desktop in the browser:
<https://red.uits.iu.edu>
- IUWare (Thinlinc) client download:
<https://iuware.iu.edu/search?q=thinlinc>



AI and Deep Learning applications

<https://kb.iu.edu/d/avjk>

- Currently 40+ projects on Carbonate Deep Learning test bed
- Support a full DL stack on Python 2 & 3 including: Tensorflow, Pytorch, Theano, Keras, Mxnet, Caffe2, CNTK, Matlab
- Consultation on performance and throughput optimization
- **AI for Everyone** through projects with Jetstream REU students




Future [is here] Compute Resources

<https://kb.iu.edu/d/aoku>

Big Red III (3)

- An upgrade of Big Red II+ to a Cray XC40 (Intel Haswell processors)
- 3X the capacity at nearly 1 petaFLOPS from 936 dual-processor nodes.
- Open for all users of Big Red II (faculty, staff, and graduate students)
- Upgrade completed in September



Big Red 200

IU's next supercomputer

~6 petaFLOPS

Installation in early 2020

Cray Shasta supercomputer

- AMD "Rome" CPUs
- NVIDIA GPUs
- Cray Slingshot Interconnect

How big is Big Red 200?

The new supercomputer is the latest major milestone in IU's decades-long leadership in pushing the boundaries of computing to advance world-class research.

Big Red 200 can process **53 times more data** in memory than Big Red and **10 times more** than Big Red II. For example, it can process **71,000** 3 GB brain scans in RAM simultaneously.

Big Red: **1,300** brain scans processed simultaneously



Big Red II: **7,100**




Big Red 200: **71,000**




Big Red 200 is almost **300 times faster** than Big Red and **6 times faster** than Big Red II...

Big Red: **20.48** teraFLOPS



Big Red II: **1** petaFLOPS



Big Red 200: **5.9** petaFLOPS




... which is almost **6 million** times faster than an iPhone XS.

Future Storage Resources

<https://kb.iu.edu/d/avkm>

Scholarly Data Archive

- Upgraded to HPSS 7.5 on October 6th
- Limits of 25K files and 50TB for users
- Additional capacity available for \$0.025 per GB per year.
- Tape drive upgrades in plan (20TB capacity) for next fiscal year.
- Storage maintained for active accounts.



Foundational Innovations

<https://kb.iu.edu/d/axpd>

Environment

- A shift toward Slurm as the primary scheduler on all systems
- Multiple storage migrations
- \$12M in Data Center Upgrades

Retirement

- Big Red II to be retired December 15th
- Karst to be retired in 2020



Future rambling

Systems

- Continued cluster and on-premise environments
- Hybrid uses of cloud and national CI
- Convergence of cloud and HPC tools
- Incorporating warm-water cooling to be a good citizen

Projects

- Involvement in national CI remains a priority (as partners and leaders).
- Post-XSEDE transition will be challenging.



Questions?



Flickr user Oiluj Samall Zeid - Lejos de Yulín



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



Acknowledgements

NSF Awards 1053575 & 1548562 (XSEDE), 1341711 (Wrangler), and 1445604 (Jetstream).

This document was developed with support from the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

Special thanks to contributors & Jetstream partners:

- Robert Henschel, Craig Stewart, Therese Miller, Jeremy Fischer, J. Michael Lowe, Winona Snapp-Childs, George Turner, Brad Wheeler, Dan Calarco, and Maria Morris.
- TACC, University of Arizona, Johns Hopkins, University of Chicago, Cornell, UT San Antonio

A PARTNER
in new
POSSIBILITIES



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

RT.IU.EDU