

4a – Advanced Computational Systems

Action Item Template Response

General Action Item Information

Lead Division/Office: Research Technologies

Action Item Number: 4a

Action Item Short Name: Advanced Computational Systems

Dependencies with other EP Action Items: 3, 4

Implementation leader (name & email): Dave Hancock (dyhancoc@indiana.edu), Matt Link (mrlink@indiana.edu)

I. INTRODUCTION AND BACKGROUND.

Cyberinfrastructure needs of researchers are ever evolving and expanding. IU has been increasingly successful in doing important research (and getting funding for more), in part on the strength of its local cyberinfrastructure. This success is in part fueled by recruitment to IU of leading computationally intensive researchers. This creates a positive feedback loop in demands on our cyberinfrastructure. Researchers require advanced computational systems to stay innovative and relevant in obtaining grant awards and supporting the needs of the next generation researchers they're training, and the need for additional cycles is accelerating.

IU has moved to the leading edge of R1 research institutions in many areas both by serving local audiences at unprecedented levels and by participating in national infrastructure projects. Examples include the Clinical and Translational Studies Institute, ATLAS experiment (part of the Large Hadron Collider) and the related Open Science Grid, the TeraGrid, and PolarGrid. Maintaining this leadership even at the campus level requires significant capital investment and timely lifecycle funding not currently present. IU's investments from 2006 must be renewed in order to maintain our standing and further attract external research funding. Several proposals within RT/PTI are outstanding, with an excellent chance of funding due to the leveraging of these resources: National Science Foundation (NSF) Science and Technology Center: Center for Remote Sensing of Ice Sheets (IU partner), NASA ICE-Bridge (1-year), NASA VENTURES/ICE-Bridge (5-year funding).

IU has periodically held the lead within US universities in the speed of its largest supercomputers, most notably Big Red. Big Red was the impetus for much change, and provided momentum in our TeraGrid grant award and FutureGrid. Our success with FutureGrid (funded by the NSF Office of Cyberinfrastructure's Track II program) places IU in a group of just six institutions nationally that have won such an award.

I. DESCRIBE YOUR PLANS FOR IMPLEMENTING THIS ACTION.

The overall goal of IU's participation in this initiative is to further develop and maintain advanced computational systems. IU's grant competitiveness is greatly affected by advanced cyberinfrastructure. Implementing the goal to maintain current faculty researcher expectations

requires significant capital investment. Supporting this infrastructure in a safe, secure, and timely manner also requires additional personnel resources. Fundamentally what is required to empower our faculty is the same core lifecycle replacement of central research systems that is provided for university laptops or workstations.

Facilities and Hardware Replacement

- Big Red replacement funding is essential no later than calendar year 2010. The system was installed in May 2006 and both failures and maintenance costs have begun to steeply increase. Installing a new system in mid-2010 allows researchers to continue productive work and ensures we continue on the leading edge of technology. This funding should occur on a regular cycle so that demand does not so outpace supply that users become discouraged. We recommend the acquisition of a 120 TFLOPS system in FY10-11 and a 360 TFLOPS system in FY13-14 to keep pace with technology improvement and demand from researchers. This pace is not quite at pace with Moore's Law, but close.
- Quarry expansion and replacement funds are critical to IU's participation in the Open Science Grid and a research condominium computing service. Further expanding Quarry will provide a basis for researchers to buy into central research systems that are professionally and securely managed without the increased cost of collocation or the spread of departmental research clusters. We recommend increasing the size of Quarry by 50% in FY10-11 and a lifecycle replacement in FY13-14.
- Microsoft's high performance cluster expansion in FY11-12. Microsoft is producing better and better services for high performance computing. We currently have more demand for our small Microsoft cluster than we are able to meet, and expect demand to grow significantly. And for the foreseeable future, it's important to have separate systems. Microsoft HPC server is not graceful about sharing ownership of the hardware it runs on with other operating systems. We expect that in FY11-12 Microsoft's high performance computing software offering will be ready for production environments.
- Research Database Server replacement enables large-scale database research that's uniquely supported by Research Technologies. Since undergoing the HIPAA certification process, UITS has positioned itself to greatly expand services for sensitive data, including electronic protected health information that is essential to supporting the IU School of Medicine. We recommend replacing the existing environment with 6 lightweight database/applications servers and leveraging virtualization technologies where possible.

Software

- New paradigms from cloud computing technologies are changing researcher expectations and may eventually surpass traditional general-purpose HPC campus infrastructure in cost and availability. To allow IU researchers to explore options in the marketplace we believe it's necessary to seed commercial cloud computing funding for researchers.

Personnel

- New cluster administrators and experts are critical to not only maintain our large central research systems, but to also interact with users as experts at a Tier 2/3 level not addressed elsewhere within UITS. For this reason it's necessary to grow our administrative base for central systems to not only support local efforts outlined above but also keep IU in a position to implement the next big grant award. As we expand into commercial cloud environments, researchers will require additional in-depth system programming and administration to effectively manipulate their environments.

II. WHAT ARE THE POLICY AND PRACTICE IMPLICATIONS OF YOUR PLANS?

None.

III. IDENTIFY STAKEHOLDERS.

Faculty and other researchers at IU campuses, including the IU School of Medicine, national and state-level researcher partners such as the TeraGrid, Open Science Grid, Indiana Economic Development Corporation, Purdue University, and IU health affiliates involved in biomedical research in collaboration with IU.