INDIANA UNIVERSITY

3 – Critical Facilities

Action Item Template Response

General Action Item Information

Lead Division/Office: Enterprise Infrastructure & Research Technologies
Action Item Number: 3
Action Item Short Name: Critical Facilities
Dependencies with other EP Action Items: 17
Implementation leader (name & email): Matt Link, mrlink@indiana.edu, for research infrastructure; Rob Lowden, rlowden@iu.edu, for enterprise infrastructure

I. DESCRIBE YOUR PLANS FOR IMPLEMENTING THIS ACTION.

IU President Michael A. McRobbie dedicated the new IU Data Center in Bloomington on 5 November 2009. During the dedication the President said:

"In the decade that we have worked to reach this day, the quality and excellence of our IT infrastructure and facilities has been a base that has enabled IU faculty and staff to secure hundreds of millions of dollars in external support from agencies like the National Science Foundation, the National Institutes for Health, and the Lilly Endowment, among others.

The Data Center that we are dedicating today will provide IU with a resource that will enable us to be even more successful in attracting such funding. Funding agencies now want to know that if they are going to fund large-scale IT infrastructure to support science and research across the nation, that it is housed at the highest standards of reliability and safety.

The Data Center has already proven critical to enhancing our competitiveness and our national role in IT research. In just the last year, more than $14 million dollars of federal research funding have come to IU dependent upon our ability to house computing systems securely in this state-of-the-art facility. In particular, IU's FutureGrid project - an experimental testbed that will define the future of grid and cloud technology - would not have been funded by the National Science Foundation were it not for this building.

The Data Center also serves the needs of the broader university community. For instance, it is home to email systems, electronic student records, business data, library resources, and the Oncourse system, to name just a fraction of the systems required by so many academic and administrative units throughout the university. It protects the systems and data that enable the university to operate on a day-to-day basis, and it serves every IU campus across the state."

These comments made specifically about the Data Center apply equally well to the entire range of critical infrastructure supporting IU's essential information technology services — from core business functions such as grades, registration, and payroll — to university operations, including processing over a million mail messages a day (while rejecting six million more per day as spam), running digital libraries servers, and providing some of the largest research data storage systems and supercomputers in the US. In planning and implementing proper facilities to protect the

infrastructure that enables IU to operate and innovate, we face critical challenges. Severe weather incidents in the past three years demonstrate that it is not sufficient to plan facilities based on longstanding data on 5-, 10-, 50-, and 100-year incidents. Data gathered over decades that once allowed us to say with some confidence: "Based on historical patterns, a weather event of such-and-such a severity should occur only once in a hundred years . . . " have proven inadequate as a planning guide. Indeed, we have had a series of 50-year or 100-year events in the central Indiana region in three of the last four years.

Adding to complexity in planning, growth in demand for electrical capacity, air conditioning, and other facilities basics is certain. The rate of growth in demand is tremendously uncertain. In planning facilities that take years to build and at least a year or more to modify, we face two choices:

- Incur bond debt early in, to invest in facilities that match the maximum possible capacity demands years into the future, paying interest and locking ourselves into current technology when we know technology for such infrastructure gets better every year.

- Do our best to forecast immediate demand, knowing that upgrades may be needed regularly as demand grows.

Indiana University has chosen to invest in facilities in a way that seems prudent: build for what can reasonably be predicted, bearing in mind the need to constantly evaluate needs for facilities expansion on a timescale that enables best use of the university's monies and allows us to build with and implement the best available infrastructure technology when expansion is necessary, and thus simultaneously provide the critical facilities needed by the university, while being good stewards of university and state finances.

The current need for expansion of infrastructure supporting critical facilities

The new Data Center was designed around key goals and financial constraints. A key goal was to ensure that if IU won a major cyberinfrastructure facility grant it would be possible to house such a facility in the new IU Data Center.

The goal of winning a Track II award, and housing it in the Data Center, has been achieved with FutureGrid. However, we are now very near to using the total electrical capacity of the research pod in the Data Center. Wonderful as the new Data Center is, needs for high performance computing at IU have accelerated rapidly - more rapidly than it is possible to change a building. As a consequence of our unprecedented success in securing federal monies to build new supercomputer systems at IU, we now need to plan for additional electrical capacity to support future growth.

Furthermore, the extent to which IU researchers involved in critical and real-time health research depend on 7 x 24 availability of research systems has changed dramatically since the IU Data Center was planned. The Uptime Institute is an IT industry organization that classifies the resiliency of physical infrastructure for critical IT facilities. In particular, it defines a tiered system of facility robustness. The key definitions for IU are as follows:

- Tier I: Basic Site Infrastructure. "A Tier I basic data center has non-redundant capacity components and a single, non-redundant distribution path serving the computer equipment." In this case the Uptime Institute indicates that "The site is susceptible to disruption from both planned and unplanned activities... An unplanned outage or failure of any capacity system, capacity component, or distribution element will impact the computer equipment."

- Tier II: Redundant Site Infrastructure Capacity Components. "A Tier II data center has redundant capacity components and a single, non-redundant distribution path serving the computer equipment." In this case the Uptime Institute indicates that "The site is

susceptible to disruption from both planned activities and unplanned events. An unplanned capacity component failure may impact the computer equipment. An unplanned outage or failure of any capacity system or distribution element will impact the computer equipment."

- Tier III: Concurrently Maintainable Site Infrastructure. "A Concurrently Maintainable data center has redundant capacity components and multiple independent distribution paths serving the computer equipment. Only one distribution path is required to serve the computer equipment at any time. In this case, the Uptime Institute states that "An unplanned outage or failure of a capacity component or distribution element may impact the computer equipment... Planned site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment."

The portions of the IU Data Center designed to house Enterprise Infrastructure were planned to be compliant with Tier III. Tier III design standards provide appropriate goals for critical IT facilities that need very high levels of availability.

The Data Center space housing Research Systems was initially designed with modest redundancy requirements with certain aspects in all tier levels, but it was not designed to meet a specific tier. This was consistent with national practice at the time - understanding the possibility of occasional outages. Since that time two things have changed dramatically.

First, researchers at IU and throughout the nation now depend on IU's research cyberinfrastructure in ways not anticipated at the time the Data Center was planned. UITS and the Pervasive Technology Institute have completed the risk analyses plan and documentation development required under HIPAA to store, manage, and analyze electronic Protected Health Information (ePHI). When the IU Office of Counsel approved these UITS/PTI plans for storage and analysis of health records — including personally identifiable information – on IU supercomputers, the capabilities provided to health researchers at IU and throughout the nation were dramatically expanded. Now researchers can use IU supercomputers to analyze critical and sensitive data in real time, or nearly so. IU is just one of two supercomputer centers in the nation that provide the ability to analyze ePHI. The importance of this capability was driven home early in 2009, when the IU School of Informatics and PTI researcher Alex Vespignani was using Big Red to perform daily predictions of the spread of the H1N1 pandemic for use in planning by the federal government. (NB: this particular work did not involve ePHI). But because IU researchers need the best possible facilities to discover new treatments for diseases, and because the capability to analyze ePHI without the tedious and time-consuming steps of de-identifying all data enables IU health researchers to use systems such as Big Red routinely, the health research community relies on IU's research cyberinfrastructure in ways that are unprecedented.

Second, while the degree to which health researchers depend upon IU's research infrastructure was expanding, weather events were demonstrating that the nationally standard practice of housing research computing systems in facilities that do not meet high availability standards simply wasn't good enough for IU researchers. Multiple severe weather events during 2008 and 2009 resulted in days or weeks of outage for IU's most advanced supercomputers. It is only a matter of good luck that none of these outages occurred while Prof. Vespignani was helping the federal government plan its response in the critical early days of the H1N1 pandemic.

Furthermore, many research groups operate extremely important computing facilities in buildings throughout IU's 8 campuses in what are essentially "Tier I" facilities - facilities that can be described in these terms: "If anything at all goes wrong the systems shut down." UITS has demonstrated tremendous success with its Intelligent Infrastructure service, which allows for extremely flexible and robust housing of departmental services within the IU Data Center. There is similarly great demand for a "condominium cluster" service in which small clusters spread throughout IU would be succeeded, over time, with larger facilities housed in the IU Data Center. The idea behind condominium cluster computing is straightforward. A research group buys

computing equipment that is added to a large condominium facility without giving up ownership of that equipment. Condominium computing enables individual researchers to use resources they own with only seconds notice, while enabling those resources to be used by the broader university community when not in use by the owner. This approach also decreases the cost of systems administration.

Co-location Services

Many departments and researchers throughout IU have expressed an interest in relocating equipment to the IU Data Center. Using the Intelligent Infrastructure, condominium model, or shared services (web, file, print) are preferred models for using robust services hosted in the Data Center and create the best and most efficient use of resources for Indiana University. There are several valid reasons that a department cannot move to those services in the short term. UITS will continue to extend co-location services to those departments to *facilitate the transition to preferred models*. Part of planning for Action 3 is to provide electrical and cooling capacity to support growth in co-location and significant increases in the other models.

In sum, there has been tremendous growth in capacity needs and robustness requirements for the research cyberinfrastructure housed at IU. At the same time, the needs for expansion and modernization of critical facilities for IU's basic business and academic support needs have not stood still. In particular, the following facilities are all in need of upgrades and modernization:

- IUB
  - Telecommunications Hub

- IUPUI
  - ICTC Data Center
  - ICTC Network Operations Center
  - ES & Riley Switch Rooms
  - UL Telecommunications Hub

Green IT

IU and UITS take seriously the commitment to maintaining the quality of our global environment. Continual review and research into new trends and approaches to improved Power Utilization Efficiency (PUE) approaches should be implemented, as appropriate, to minimize the impact on the environment of IU's information technology systems, and reduce the university's overall carbon footprint. In so doing we can ensure facilities that balance the specialized requirements placed upon IU's IT infrastructure with our shared responsibilities to the environment.

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

This plan represents a continuation of IU and UITS policies and plans regarding mission-critical and enterprise IT facilities.

This plan represents a significant change in past UITS practice, and recognizes that IU's research cyberinfrastructure has become so critical that it merits the high availability of a Tier III facility.

This plan also makes concrete the strategy for facilities upgrades outlined above: monitor need constantly, and implement facilities enhancements in a way that meets growth in a "just-in-time" fashion, enabling the optimal use of the university's finances and allowing the use of the newest and best technology available when demand arises.

III. IDENTIFY STAKEHOLDERS.

- UITS
  - Research Technologies
  - Enterprise Systems
  - All
- University
  - Office of Risk Management
  - Architect's Office
  - All researchers, particularly the health research community
- Statewide and nationally
  - All who depend upon or benefit from IU research