Indiana University Cyberinfrastructure Plan (CI Plan) – Updated January 2020

1. Introduction

Indiana University's plans for cyberinfrastructure have been guided since 1998 by two university-wide information technology strategic plans, the first endorsed by the Trustees of Indiana University in 1998, succeeded by a new plan endorsed by the Trustees in 2008. The first strategic plan, subtitled "Architecture for the 21st Century" [1], primarily focused on correctly implementing various technologies and technology stacks. After a decade of focus on bringing the information technology within the university, the 2008 plan, entitled "Empowering People – Indiana University's Strategic Plan for Information Technology" [2] focused on supporting the various roles of people within the University and advancing the university's mission of education, research, and engagement in the life of the state, nation, and world. This plan contains 15 general, high--level recommendations and 72 specific Actions to accomplish those recommendations.

Before going into details of IU's cyberinfrastructure plan as regards networks and data movement, it is worth outlining some key facts about the current cyberinfrastructure and cyberinfrastructure support at IU:

- Indiana University is one university comprising eight campuses with one central information technology organization (University Information Technology Services) serving the entire university. There is thus one set of policies, integrated support, integrated application development, and integrated network and cyberinfrastructure service for the entire university. There are IT staff based in departments and schools of the university; they work very closely, and under one unified university IT policy framework, through a program called 1ITIU.
- IU has had leadership in information technology deployment, development, and use since 1997. This is embodied in the current Empowering People strategic plan in *Recommendation 1: "Indiana University's national and international leadership should be sustained through continued maintenance and advancement of an IT infrastructure that is supported by sound fiscal planning*."
- IT at IU in general, and University Information Technology Services in particular, operates under what we refer to as a 'philosophy of abundance" as defined in Empowering People *Action 5: "IU should pursue strategies that approximate a philosophy of abundance, within reason, towards unmetered availability of basic IT services, support, and infrastructure for creative activity, storage, computation, communication, and other activities fundamental to the work of the university via any appropriate sourcing strategy."* As a result, there are no quotas limiting use of IU's high performance computing systems (usage is regulated on a "fair share" basis). Default quotas are set very high - 10 TB for disk storage on the IU Data Capacitor and 50 TB for storage in the IU Scholarly Data Archive (the tape archival system). Initial default quotas are doubled upon request without question, and higher quotas – up to PBs of tape storage – are allocated to projects that advance the University's mission and/or the research and education activities of the US science and engineering community.

IU is a leader nationally and internationally with the Global Research Network Operations Center (GlobalNOC). This leadership is set as an objective in Recommendation 2 of the Empowering People strategic plan: "Indiana University should ensure that its wired and wireless campus networks continually evolve just ahead of the needs of IU's faculty, staff, and students. The network must provide secure, reliable, effective, and appropriate access to support the missions of the university." Actions related to this Recommendation are as follows:

- Action 8: IU should continue to maintain its networks to accommodate the increase in demand for capacity, speed, security, and stability." (This action can be summarized as "provide excellent network connectivity within and among the IU campuses).
- Action 9: IU should continue to pursue opportunities for strategic partnerships that can provide services for advanced networks to further the missions of the university. (This action can be summarized as: IU should pursue additional high speed network connections to national and international research networks, while also looking to expand the GlobalNOC's activities providing operations and support services for other networks)."
- Action 10: Whenever feasible, IU should develop prudent agreements, partnerships, and other mechanisms (e.g., strategic alliances, negotiations with service providers) that allow faculty, staff, and

students to acquire low- cost access to high--speed home and mobile connectivity.“ (This action can be summarized as: continue pursuing means by which home and mobile connectivity can be enhanced and made less expensive for members of the IU community.)

Research in a few selected areas of information technology is also a part of IU’s strategic plan. Recommendation 15 of Empowering People is *“While Indiana University should advance IT- enabled research across all disciplines, it should also focus on a few highly promising opportunities for which it has a skills, knowledge, and reputational advantage to push the frontiers of IT-enabled research and scholarship.”* The areas selected for particular attention are as follows: data--enabled and data--intensive science; networking; development and implementation of scientific workflows in support of discipline--specific research; security; life sciences and biomedical research

2. IU networks

IU Networks are managed as three components: connectivity to external networks, the IU research network, and the IU enterprise network. In 2018, the Indiana GigaPOP and IU launched Monon400, a next generation Ethernet network bringing 400 Gbps connectivity to the GigaPOP in Indianapolis to Chicago. For the Monon400, Indiana GigaPOP deployed Ciena's Waveserver® Ai 400G stackable interconnect platform and Blue Planet® Manage, Control and Plan (MCP) domain controller with Liquid Spectrum. The resulting programmable infrastructure enables the use of 400G channels connecting Indianapolis to Chicago. The GigaPOP connects in Chicago with multiple 100G and 10G ports. The 100G ports connect the GigaPOP to Internet2 and Indianapolis, enabling a full 100G path from Internet2 to the IU Research Network, which has been built out to 100 Gbps to the Bloomington Data Center (70% of the bandwidth between Indianapolis and Bloomington will be reserved for research traffic). These linkages are described in greater detail in the Facilities, Equipment, and Other Resources document.

2.1 Physical network plant. IU has completed a very detailed ten year (2008-2017) Network Master Plan covering all aspects of networking communications (voice and data). This plan was predicated on existing plans for IU to convert all telephony services to voice over IP (VoIP). This is complete across all 8 campuses.

Also, Indiana University’s wire plant has been systematically replaced for nearly ten years. Over 90% of the two main campuses have at least CAT 6E wiring, with Regional campus upgrades catching up quickly. These upgrades are being carried out by completely replacing the wire plant inside buildings including upgrades to wiring closets, ensuring physical security of the wiring closets, and installing 50 micron mm OM4 fiber between floors. As for outside plant, all buildings at the main campuses are now have redundant connections to the IU enterprise network at 10Gbps.

From a wireless perspective, IU has continuously added access points to address areas with no or less than optimal coverage. Generally speaking, in-building coverage is near 100% in all academic and shared use and open spaces where the community congregates. There is also a large percentage of outdoor coverage between our academic buildings.

IU is three years into a second 10 year Network Master Plan (2018-2027). So far, the edge component upgrade (switches and access points) is well over half way to completion. This portion of the plan incorporates 1,541 switches and 11,329 access points across all 8 campuses. Another 2000 additional access points will be added in the next year in residential areas at the IUB campus.

2.2 Implementation of IPv6. By 2011, Indiana University (Bloomington and Indianapolis campuses) implemented IPv6 on the wired and wireless networks. IPv6 has subsequently been implemented on the smaller regional campuses. Each campus has been assigned a /48 block of addresses from IU's /38 assignment from the Indiana GigaPOP, which allows for 65,536 subnets on each campus, which should be sufficient for the foreseeable future. In terms of security, Indiana University implements a router advertisement guard on wired edge switches and wireless controllers to prevent rogue router advertisements. We implement BGP black hole route injection, so the security office can inject a null route for any specific IPv6 host destination. Additionally, we collect the IPv6 neighbor tables on the routers to ensure that we have an IPv6 address to MAC address correlation for identifying the owner's of hosts related to a security incident. This is done by correlating the MAC address to user information

in our MAC address registration database. The wireless network requires 802.1X authentication, which also allows us to identify a user based on the username they used to authenticate to the network.

2.3 Mutually Agreed Norms for Routing Security (MANRS). IU is a member of MANRS[3] and has adopted the emerging best practices in network routing security for network operators as expressed in the Mutually Agreed Norms for Routing Security standard.

3 Compute and storage capacity. As outlined in depth in the Facilities, Equipment, and Other Resources document, IU has available to all faculty, staff, and students a Cray XC40 supercomputer (Big Red 3), and two IBM NeXtScal nx360 clusters (Karst, Carbonate) with a total of 1.26 PFLOPS capacity. By Summer 2020, a Cray XK7 (Shasta) supercomputer (Big Red 200) capable of 6 PFLOPs peak computational performance, will be installed. IU has several major disk-based file systems (called Geode, Slate, DC-WAN) with a total capacity that exceeds 11 PB (usable) and one archival storage system (Scholarly Data Archive; 79 PB).

4. InCommon implementation

Indiana University is an InCommon Federation entity. IU meets the InCommon Baseline Expectations for Trust in Federation. As an Identity Provider (IdP), IU releases attributes to Research and Scholarship Service Providers *registered by InCommon[4]).* IU information technology experts have published a detailed guide to InCommon implementation, along with a condensed summary of that guide [5]. The default mechanism for access to all IU--delivered research applications is based on the InCommon Federation trust relationships, XSAML certificates, and software that supports mechanisms for authentication via these mechanisms (e.g CILogin, eduroam, Shibboleth).

Staff of University Information Technology Services and the Pervasive Technology Institute also support the use of the following applications by members of the IU community – all of these applications are accessed via InCommon--based authentication:

- *IU Login.* Single sign-on portal for IU internal applications, e.g. Outlook Web Access to Exchange email for IU faculty and staff.
- *Box at IU*. Cloud--based file sharing solution provided by Box.com.
- *Canvas.* Learning management system deployed at IU.
- *EDUCAUSE.* A national organization that promotes information technology in higher education.
- *Google at IU.* Portal to Google G Suite for IU students, faculty, and staff.
- *IUWare.* Software distribution service for IU students, faculty, and staff.
- *mybtaa.* A collaboration portal for the Big10 Academic Alliance.

5. IU Cybersecurity and resilience

IU has robust physical and cybersecurity infrastructure, policies and procedures, and expertise in cybersecurity and research cybersecurity in particular. Information about the physical security and safeguards for the system is in the Facilities and Other Resources document. There are three major security entities at IU: the *University Information Security Office and Security Operations Center (OmniSOC) and* the *Center for Applied Cybersecurity Research (CACR).*

- The University Information Security Office (UISO), part of the Office of the Vice President for Information Technology (OVPIT), provides security analysis, development, education, and guidance related to Indiana University's information assets and information technology environment. The objective is to establish and maintain a resilient and secure infrastructure in which to conduct university business. UISO and UIPO (University Information Policy Office) staff work together in responding to and investigating incidents related to misuse or abuse of IU information technology resources, including computer and network security breaches and unauthorized disclosure or modification of electronic institutional or personal information.
- OmniSOC[6] is a security operations center that provides trusted and actionable intelligence to higher education institutions. Hosted at Indiana University and founded by five institutions - Indiana University, Northwestern University, Purdue University, Rutgers University and the University of

Nebraska-Lincoln - OmniSOC helps members thwart threats through collaboration, threat detection and data sharing, leveraging two decades of experience and capabilities from GlobalNOC, the 24/7 network operations center that provides services to government, research and education networks across the nation. In addition to the GlobalNOC, OmniSOC works in close coordination with the federally chartered Research and Education Networking Information Sharing and Analysis Center (REN-ISAC)[7] at IU, an organization that collects and shares cybersecurity information among its 580 members within the research and higher education communities. OmniSOC collects real-time data from a broad array of devices and systems from each member and stores it in a central repository at IU. Using real-time data from members as well as governmental and corporate security subscriptions, OmniSOC identifies suspicious and malicious activities that require mitigation and provides rapid incident response through human analysis and machine learning.

- The *Center for Applied Cybersecurity Research (CACR)*. CACR is a university-wide research center affiliated with the Indiana University Pervasive Technology Institute and a member of the Indiana University cybersecurity community, which includes the Maurer School of Law, the Kelley School of Business, the School of Informatics and Computing, REN-ISAC, the OmniSOC, the University Information Policy Office, and the University Information Security Office. CACR provides the nation with leadership in applied cybersecurity technology, education, and policy guidance. Fundamental to CACR's mission is to properly balance public needs, homeland security concerns, and individual privacy rights. Indiana University has been named an NSA and DHS National Center of Academic Excellence in both Information Assurance Education and Information Assurance.

All compute and storage hardware have standard technical and administrative safeguards as well as policies and procedures for its computing environment; these include firewalls, automated intrusion prevention, federated identity, multi-factor authentication (DUO) for administrative access, network intrusion detection, and vulnerability scanning. Indiana University implements a unicast reverse path forwarding--based feature on all its LAN-facing interfaces to ensure IP source spoofed packets will not be forwarded beyond the local subnet. Data from intrusion detection systems and other low-level system log data is made available to the OmniSOC, a shared cybersecurity operations center, which makes use of threat intelligence from the Research and Education Network Information Sharing and Analysis Center. IU has an acceptable use policy [8]. Generally, in the event of an incident, the IU Security Office will be responsible for incident response and proactive notifications. Government agencies (police, FBI, etc) will be notified as appropriate, as determined by the IU Security Office (see IU IT policy[9]).

IU has robust and well-defined approaches to privacy [10,11], including policies for complaints[12]. IU adheres to federal and Indiana state standards (e.g. HIPAA, FERPA, FISMA, etc.) for data privacy as appropriate to protected health information, student records, financial records, and other sensitive data[13].

6. Cyberinfrastructure in support of scientific research, collaboration, and scientific workflows

6.1 Support for general scientific workflows. The Cyberinfrastructure Integration Research Center (CIRC) [14] researches, develops, and operates science gateways in collaboration with many clients and partners. CIRC develops and uses open source software to create user tools and environments that aid scientific communities. The cyberinfrastructure software systems are designed and developed strategically, not just from the requirements of scientists but from deeper understanding of core principles of cloud-scale distributed systems.

6.2 Support for bioinformatics workflows. The National Center for Genome Analysis Support (NCGAS) [15] is house primarily at Indiana University. NCGAS provides bioinformatics consulting services to biologists doing research in the areas of transcriptome and genome assembly, phylogenetics, metagenomics/transcriptomics and community genomics. NCGAS also Maintains, support, and deliver genome assembly and analysis software on national cyberinfrastructure systems and provide software tools for download and installation on cyberinfrastructure resources local to campuses and research labs. The group supports Galaxy, Trinity, Genepattern gateway users as well as an extensive software library (> 221 software titles available for download from the NCGAS web site; and >123 software titles available for download from the XSEDE Community Software Repository).

6.3 Support for cyberinfrastrucure and data collaboration support. The Engagement and Performance Operations Center (EPOC) [16] was established in 2018 as a collaborative focal point for operational expertise and analysis and is jointly led by Indiana University (IU) and the Energy Sciences Network (ESnet). EPOC provides researchers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. By considering the full end-to-end data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process.

7.0 References

[1] McRobbie, Michael A., "Information Technology Strategic Plan – Architecture for the 21st Century, Indiana University, 1998, URI: http://hdl.handle.net/2022/6823

[2] Wheeler, Brad; Acito, Frank, Empowering People: Indiana University's Strategic Plan for Information Technology 2009", Indiana University, 2009, URI: http://hdl.handle.net/2022/20037

[3] Mutually Agreed Norms for Routing Security (MANRS), https://www.manrs.org/

[4] Research and Scholarship Adopters for In Common, https://www.incommon.org/federation/research-scholarship-adopters/

[5] Barnett, W., V. Welch, A. Walsh and C.A. Stewart. A Roadmap for Using NSF Cyberinfrastructure with InCommon. 2011. http://hdl.handle.net/2022/13024

[6] OmniSOC, https://omnisoc.iu.edu/

[7] Research & Education networks Information Sharing & Analysis Canter (REN-ISAC), https://www.ren-isac.net/

[8] Acceptable Use Agreement for access to technology and information resources at IU, https://protect.iu.edu/online-safety/acceptable-use.html

[9] Information and Information System Incident Reporting, Management, and Breach Notification. University Policies, https://policies.iu.edu/policies/ispp-26-information-system-incident-reporting/index.html

[10] Protect IU, https://protect.iu.edu/privacy/index.html

[11] Privacy of Electronic Information and Information Technology Resources IT-07, https://policies.iu.edu/policies/it-07-privacy-it-resources/index.html

[12] Privacy Complaints ISPP-27, https://policies.iu.edu/policies/ispp-27-privacy-complaints/index.html

[13] Management of Institutional Data - DM-01, https://policies.iu.edu/policies/dm-01-management-institutional-data/index.html

[14] Cyberinfrastructure Integration Research Center (CIRC), https://pti.iu.edu/centers/cyberinfrastructure-integration/index.html

[15] National Center for Genome Analysis Support (NCGAS), https://ncgas.org/

[16] The Engagement and Performance Operations Center (EPO), http://www.epoc.global