University of Cincinnati
Cyberinfrastructure (CI) Plan
2019-2023

University of Cincinnati
Cyberinfrastructure Plan
2020-2025

EXECUTIVE SUMMARY

The University of Cincinnati has a robust enterprise network that provides network services for greater than 45,000 faculty, staff, and students across several campuses, however, the demand for shared computing power, ever-increasing data volumes, and higher network capacities are all on exponential growth paths. To address the changing needs of the academic and research communities, this plan outlines a comprehensive cyber-infrastructure framework that has the potential to revolutionize science, engineering, and other research disciplines across the University of Cincinnati.

At the core of the plan is the expansion of a friction-free campus networking architecture that facilitates the growth of UC's Research Ecosystem. Key goals include:

- **Modify and Enhance UCScienceNet (UCSN)**, a 100Gb Science DMZ following the model of ESnet and incorporating best practices and lessons learned for monitoring and tuning network performance to provide broader access to central storage and compute resources. Implement the new network design pattern recommended by the NSF EPOC team.
- **Expand High-Performance Compute (HPC) resources**, leveraging the existing baseline HPC infrastructure and expand capacity to the entire campus research community via the Advanced Research Computing initiative.
- **Expand local Research Storage environments,** including a HIPAA compliant storage, data transfer nodes and Globus file transfer mechanism.
- **Continue IPv6 Implementation** and transition plans in conjunction with academic sponsorship.
- **Become an InCommon "Silver Assurance" Identity Provider**, to provide federated login access to online services that require a greater confidence in identity.

The University of Cincinnati is committed to implementing this Cyberinfrastructure Plan to fully embrace the mission and vision of enabling transformative academic and research initiatives.

INFRASTRUCTURE

Current Network

The UC Office of Information Technology (IT@UC) has implemented a network architecture with the goal to provide a resilient, stable network to the university community. The network design consists of redundant core routers each connected via 40Gbps fiber uplinks to distribution layer switches at one of five distributed node rooms. The closet switches connect back to the distribution switches via dual 1Gbps or 10Gbps fiber uplinks providing redundancy to the access layer switches. End users connect via a 10/100/1000Mbps Ethernet connection and share the uplink bandwidth back to the distribution layer switches. This shared uplink has the potential to restrict research capabilities of transferring large data sets from locations in the enterprise network.

Internet connectivity is provided by a Dense Wavelength Division Multiplexing (DWDM) metropolitan optical ring also known as the Cincinnati Educational Research Fiberloop (CERF) ring. This ring

provides redundant 10Gbps connections for the university to OARnet, and Internet2. The CERF ring, which is managed by IT@UC, also provides Internet connectivity for Cincinnati Children's Hospital Medical Center, Xavier University, and Cincinnati State Technical and Community College on separate interfaces. The CERF ring is optimized to prevent any loss of service in the event of a fiber cut.

As previously mentioned, UC utilizes a 10Gb network interface to connect to OARnet through the CERF ring. The 10Gb pipe is a connection shared by the university community for commodity Internet along with any research data traffic housed on the enterprise network. This bottleneck is prohibitive for research data to be transferred between peer institutions. To remove the bottleneck, IT@UC increased the hardware capability of the CERF ring. The enhanced capability enables researcher access to a 100Gb pipe, via the SciNet, connecting the university's main campus to OARnet's 100Gb Internet 2 backbone.

UCScienceNet (UCSN) – UC's Science DMZ for research traffic and data transfer
UCScienceNet (UCSN), a 100Gb Science DMZ , modeled after ESnet's Science DMZ design, incorporates PerfSONAR for monitoring and tuning network performance, enables software-defined networking and OpenFlow capabilities, and provides high-throughput capacity required to achieve STEM research goals and enable multiple disparate high-speed big data transfers across a comprehensive, integrated, cyberinfrastructure.

UCSN consists of hardware deployed specifically for aggregation of high-speed networking. This hardware has characteristics of high-throughput with minimal latency to ensure rapid delivery of large scientific data sets. The hardware employs bandwidth scalable from 40Gb depending on research requirements and 100Gb delivery from the aggregation layer outward to Internet 2 and National Research and Education Networks (NREN).

UCSN, servicing five research intensive locations, provides a friction-free network, creating a true Science DMZ to address the limitations of the existing commodity network. IT@UC in partnership with the Office of Research, provided funding to add additional endpoints to UCSN expanding benefits of a high-speed network to researchers not connected during the initial deployment of UCSN.

NSF EPOC Requirements Deep Dive Report Recommendations and Action Items

The coordinated effort of the NSF EPOC Requirements Deep Drive process produced several action items and recommendations that are being implemented by the UC Cyberinfrastructure team in 2020. The items were brought to the IT@UC R&D governance committee and all were recommended and approved.

1. **New Network Design Pattern**
    - Estimated completion date – January 2021
        - Installing additional 100G circuit providing connectivity to OARnet the regional Internet/Internet2 provider
        - Installation of two Cisco ASR 9901 routers at redundant OARnet entry sites
        - Consolidate routing for ScienceNet and commodity network into these two routers.
        - Data access will be accessible from the ScienceNet via Data Transfer Nodes
        - Control/Management access to the new compute resource will be provided via the commodity network

- The two 100G pathways to OARnet will be configured for redundancy so that the ScienceNet utilizes one circuit as primary and the commodity network uses the other circuit as primary but in the event of single circuit failure, both networks can share the remaining circuit
- The ASR 9901 routers provide deep buffers that will enable better transfer rates across high latency links

2. ***University of Cincinnati*** will explore the addition of local data storage options for university departments that includes a data transfer node, a HIPAA complaint storage solution, and a data transfer mechanism that supports federated identity and high-performance use cases. (e.g. Globus).
3. ***University of Cincinnati*** will deploy additional measurement and monitoring tools, campus wide, with a focus on flow data analysis. Additional perfSONAR nodes, at key areas of interest, are also being explored.
5. ***University of Cincinnati*** and ***OARnet***, will work together to better connect industry and government collaborations via direct peering arrangements.
6. ***University of Cincinnati*** and ***OARnet***, will work together to establish specific network relationships, via peering and other mechanisms, to explore secure transfer of PII/PHI/ePHI information between collaborators in this space.
7. ***University of Cincinnati*** will explore the demand for ITAR/EAR data management via implementation of security frameworks such as NIST 800-53/800-171. They will work with ***OARnet*** and ***EPOC*** to implement solutions.
8. ***University of Cincinnati*** will work with the Department of Physics to better understand data growth needs and requirements beyond the LHC Long Shutdown 2 and the impacts of new data movement tools.
9. ***University of Cincinnati*** will work with Aerospace Engineering to establish a 'visualization' host that is capable of existing on the DMZ, but supports a low-latency graphical use case, as well as identifying other resources that should be exposed via the DMZ infrastructure.

PerfSONAR

The increased interest in quantifying high-speed bandwidth available for research and education networks has led to an initiative to deploy network monitoring tools at key points of the network. We have incorporated a perfSONAR framework, to gather throughput statistics that are relevant to the use cases of researchers on UCSN and produce usability studies from applied use of remote big data transfers.

Data Center

The UC Data Center, managed by IT@UC Enterprise Shared Services, is an enterprise level facility that provides 6700 square feet of managed space for core IT@UC systems, university research systems and UC co-locators. Services provided include: 24-hour badge access and video security systems, enterprise system for infrastructure management and monitoring (DCIM), clean agent fire suppression (HALON) and dry-pipe sprinkler solution, in room enterprise UPS systems, and an Automatic Transfer Switch (ATS) connected to a backup diesel generator. The data centers internal network provides high-speed data transfers between enterprise storage and the university's core systems.

UC has entered into a partnership with the State of Ohio and established a secondary data center at the State of Ohio Computer Center (SOCC) in Columbus. Our SOCC data center, provides real-time

synchronization with data storage systems in our primary data center, replication of data backups, and both active-active and active-standby hardware for critical business continuity and disaster recovery scenarios.

The data center is planning a two-phase upgrade to support the power and cooling requirements necessary to support the ever-increasing demands of the HPC cluster and equipment. Phase I to be completed March 2021 and Phase II, which is much more extensive and establishes a segregated research computing environment, is anticipated to be completed by 2025.

Storage Capacity

Research Scratch and Project storage
Planned (January 2021) Research Scratch and Project storage, funded by a NSF MRI award, will be provided on a Cray ClusterStor E1000 appliance with 1.68 PB RAW, 25 GB/s Read/Write. The parallel file system is integrated into the UC Central HPC switch fabric for maximum, non-blocking I/O performance and security. Performance is 25 GB/s sequential write, 25 GB/s sequential read. The system is scalable to multiple petabytes of capacity and can easily be expanded to meet increasing storage demands. The storage appliance will be available to both UC HPC clusters.

Compute Capacity
The Advanced Research Computing (ARC) facility provides central HPC clusters which are available to all university researchers, their students and collaborators. A Faculty Advisory Committee sets policies and reviews/approves allocations. An external advisory board is in place to provide guidance and strategic direction.

ARCC1: UC's pilot HPC Cluster Omnipath 100 Gb/s interconnect
17 CPU Nodes: Intel Xeon Gold 6148 2.4G, 20C/40T, 10.4GT/s, 27M Cache, Turbo, HT (150W) DDR4 2666 RAM 192GB per node, DDR4 2666

1 GPU Node: NVIDIA Tesla V100 32G Passive GPU with 2 nodes
1 - ZFS Storage Node – 96TB raw storage (initially configured to offer 43TB)

ARCC2: UC's Production HPC Cluster - Mellanox InfiniBand 100 Gb/s interconnect
71 CPU nodes: HPE Apollo 2000, dual AMD EPYC 7452, 32 cores (64 total) 2.3GHz, 256 GB RAM, Gen4 PCIe bus

7 GPU nodes: HPE ProliantDL 385 Gen10+ GPU nodes, dual AMD EPYC 7452, 32 cores (64 total), 1024 GB RAM, dual Nvidia A100-40 GPUs, Gen4 PCIe bus

IPv6 Implementation - NOC
The university has a /48 IPv6 assignment from ARIN and has successfully deployed a pilot IPv6 network in conjunction with academic sponsorship. Our perimeter firewall and infrastructure support services, such as DNS and DHCP, are fully capable of IPv6 support as determined from the IPv6 pilot project. The next phase of IPv6 deployment will be on strategic internet-facing web servers positioned in our data center. This IPv6 build out will be in coordination with our Business Application Services team.

BCP 38 – NOC/Security
BCP 38 or RFC2827 is an internet best practice of employing Network Ingress Filtering to defeat Distributed Denial of Service (DDoS) attacks which employ IP Source Address Spoofing. The university has been an early adopter of Network Ingress Filtering; access lists are deployed on all internal routers allowing only known-source IP addresses from local subnets. Perimeter routers are configured with access lists to deny any incoming packets with spoofed source IP addresses that mimic our class B public IP range.

The following are abbreviated web page results from running the Spoofer Project's Spoofer tool from an internal university computer:

Egress Filtering Depth: - NOC/Security The "tracefilter" test found your host unable to spoof valid, non-adjacent source addresses through even the first IP hop.

Received (Adjacent Netblock Testing) Your host can spoof 0 neighboring addresses (within your /32 prefix). Spoofed probe packets appear to be blocked by your local NAT rather than being rewritten with a public source address. NAT is blocking spoofed traffic rather than rewriting the source address.

Sustainability
IT@UC currently engages Cisco Smartnet for maintenance on all of our core, distribution, and optical equipment. We utilize Aruba maintenance for all of our wireless controllers and access points. All network operations and engineering services are provided by the Network Operations Center.

In addition, any hardware that is scheduled to reach an End of Support status will be incorporated into equipment refresh cycles outlined in our 5-year plan.

IDENTITY & ACCESS MANAGEMENT
The University of Cincinnati is already an InCommon Identity Provider, providing federated access to 77 service providers (30 organizations, 41 higher educational, 6 government). In cooperation with our Office of Information Security, the university is actively pursuing InCommon "Silver Assurance" status to provide federated login access to online services that require a greater confidence in identity.

INTEGRATION WITH STATE, NATIONAL, & INTERNATIONAL PARTNERS
Researchers across the institution leverage many big data research partners which necessitate a need for data transfer requiring high bandwidth capabilities. Among these are the Ohio Supercomputer Center, NSF National Snow and Ice Data Center, Nasa (Goddard, Ames, JPL), Alaska SAR Facility, USGS EROS Data Center, Oak Ridge National Laboratories, NCSA, XSEDE Supercomputers, NIST, CERN, Fermilab and numerous other peer institutions.

The UC Office of Information Technology staff is closely integrated with OARnet, Internet2, industry partners and other research-intensive universities staff to share best practices and resources to accelerate the national research and discovery efforts.