

Lawrence Berkeley National Laboratory

LBL Publications

Title

University of Cincinnati Campus-Wide Deep Dive

Permalink

<https://escholarship.org/uc/item/6t58p052>


Authors

Zurawski, Jason
Addleman, Hans
Schopf, Jennifer
[et al.](#)

Publication Date

2019-11-01

Peer reviewed



University of Cincinnati Campus-Wide Deep Dive

April 26th, 2019



U.S. DEPARTMENT OF
ENERGY
Office of Science



ESnet

ENERGY SCIENCES NETWORK



INDIANA UNIVERSITY

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor The Trustees of Indiana University, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California or The Trustees of Indiana University. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California, or The Trustees of Indiana University.

University of Cincinnati Campus-Wide Deep Dive

Final Report

*University of Cincinnati
Cincinnati, OH
April 26th, 2019*

The Engagement and Performance Operations Center (EPOC) is supported by the National Science Foundation under Grant No. 1826994.

ESnet is funded by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. Benjamin Brown is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory, which is operated by the University of California for the U.S. Department of Energy under contract DE-AC02-05CH11231.

This is a University of California, Publication Management System report number LBNL-2001320¹.

¹<https://escholarship.org/uc/item/6t58p052>

Contents

Disclaimer	2
University of Cincinnati Campus-Wide Deep Dive	3
Final Report	3
Contents	4
Participants and Contributors	8
Report Editors	8
1. Executive Summary	9
2. Process Overview and Summary	11
2.A Campus-Wide Deep Dive Background	11
2.B Campus-Wide Deep Dive Structure	12
2.C University of Cincinnati Campus-Wide Deep Dive Background	13
2.D Organizations Involved	14
3. University of Cincinnati Case Studies	16
3.1 High Energy Physics Case Study	17
3.1.A Science Background	17
3.1.B Collaborators	17
3.1.C Instruments and Facilities	17
3.1.D Process of Science	19
3.1.E Remote Science Activities	20
3.1.F Software Infrastructure	20
3.1.G Network and Data Architecture	22
3.1.H Cloud Services	23
3.1.I Known Resource Constraints	23
3.1.J Parent & Affiliated Organizational Cooperation	24
3.1.K Outstanding Issues	24
3.2 Department of Aerospace Engineering and Engineering Mechanics	25
3.2.1 Fuzzy AI for Predictive Modeling Case Study	26
3.2.1.A Science Background	26
3.2.1.B Collaborators	26
3.2.1.C Instruments and Facilities	26

3.2.1.D Process of Science	27
3.2.1.E Remote Science Activities	28
3.2.1.F Software Infrastructure	28
3.2.1.G Network and Data Architecture	28
3.2.1.H Cloud Services	28
3.2.1.I Known Resource Constraints	29
3.2.1.J Parent & Affiliated Organizational Cooperation	29
3.2.1.K Outstanding Issues	29
3.2.2 Gas Turbine Simulation Lab (GTSL) Case Study	30
3.2.2.A Science Background	30
3.2.2.B Collaborators	30
3.2.2.C Instruments and Facilities	30
3.2.2.D Process of Science	31
3.2.2.E Remote Science Activities	32
3.2.2.F Software Infrastructure	33
3.2.2.G Network and Data Architecture	34
3.2.2.H Cloud Services	34
3.2.2.I Known Resource Constraints	34
3.2.2.J Parent & Affiliated Organizational Cooperation	34
3.2.2.K Outstanding Issues	35
3.3 Human Genetics and Genomics Case Study	36
3.3.A Science Background	36
3.3.B Collaborators	36
3.3.C Instruments and Facilities	37
3.3.D Process of Science	38
3.3.E Remote Science Activities	40
3.3.F Software Infrastructure	40
3.3.G Network and Data Architecture	41
3.3.H Cloud Services	41
3.3.I Known Resource Constraints	41
3.3.J Parent & Affiliated Organizational Cooperation	42
3.3.K Outstanding Issues	42
3.4 University of Cincinnati Corrections Institute (UCCI) Case Study	43
3.4.A Science Background	43
3.4.B Collaborators	43
3.4.C Instruments and Facilities	44

3.4.D Process of Science	44
3.4.E Remote Science Activities	46
3.4.F Software Infrastructure	46
3.4.G Network and Data Architecture	47
3.4.H Cloud Services	47
3.4.I Known Resource Constraints	48
3.4.J Parent & Affiliated Organizational Cooperation	48
3.4.K Outstanding Issues	48
3.5 Division of Statistics and Data Science at the University of Cincinnati Case Study	49
3.5.A Science Background	49
3.5.B Collaborators	49
3.5.B.1 Jet Propulsion Laboratory (JPL), Pasadena, CA	50
3.5.B.1 Cincinnati Children’s Hospital Medical Center (CCHMC), Cincinnati, Ohio	50
3.5.C Instruments and Facilities	50
3.5.D Process of Science	51
3.5.E Remote Science Activities	52
3.5.F Software Infrastructure	53
3.5.G Network and Data Architecture	53
3.5.H Cloud Services	53
3.5.I Known Resource Constraints	54
3.5.J Parent & Affiliated Organizational Cooperation	54
3.5.K Outstanding Issues	54
4. University of Cincinnati Network and Computational Environment	55
4.1 University of Cincinnati Enterprise Network Diagram	56
4.2 University of Cincinnati Science Network Diagram	57
4.3 Computational Environment	58
5. Discussion Summary	60
5.1 High Energy Physics	60
5.2 Department of Aerospace Engineering and Engineering Mechanics	61
5.2.1 Fuzzy AI for Predictive Modeling	61
5.2.2 Gas Turbine Simulation Lab (GTSL)	62
5.2.3 High Fidelity Computations and Models for Propulsion	63
5.3 Bioinformatics	63
5.4 University of Cincinnati Corrections Institute (UCCI)	64

5.5 Division of Statistics and Data Science	65
5.6 University of Cincinnati Network and Computational Environment	67
6. Action Items	73
Appendix A – University of Cincinnati Cyberinfrastructure Plan	74
Background	74
Infrastructure	75
Current Network	75
UCScienceNet (UCSN)	76
PerfSONAR	76
Data Center	76
Storage Capacity	77
Compute Capacity	77
UC’s Advanced Research Computing (ARC)	77
IPv6 Implementation - NOC	78
Sustainability	78
Identity & Access Management	78
Resources & Services	78
Integration with State, National, and International Partners	79

Participants and Contributors

Jim Ablordeppey, University of Cincinnati - Information Technology
Hans Addleman - Indiana University
Hunter Bridewell, University of Cincinnati - Information Technology
Jane Combs, University of Cincinnati - Information Technology
Franklin Curry, University of Cincinnati - Information Technology
Daniel Gerard, University of Cincinnati - Criminal Justice
Jason Gerst, University of Cincinnati - Criminal Justice IT
Brett Harnett, University of Cincinnati - Biomedical Informatics
Emily L. Kang, University of Cincinnati - Math, Statistics and Data Science
Prashant Khare, University of Cincinnati - Aerospace Eng
Amy Koshoffer, University of Cincinnati - Libraries Research & Data Services
Amy Latessa, University of Cincinnati - Information Technology
Dr. Yaping Liu, University of Cincinnati- Cincinnati Children's Hospital Medical
Center, Division of Human Genetics
Annette Ready, University of Cincinnati - Information Technology
Barbara Renner, University of Cincinnati - Information Technology
Kurt Roberts, University of Cincinnati - College of Medicine IT
Brian Ruehl, University of Cincinnati - Information Technology
Anoop Sathyan, University of Cincinnati - Aerospace Eng
Larry Schartman, University of Cincinnati - Mechanical & Materials Eng IT
Paul Schopis - OARnet
Cody Smith - OARnet
Myrinda Smith, University of Cincinnati - Criminal Justice
Mike Sokoloff, University of Cincinnati - Physics
Adam Steele, University of Cincinnati - Information Technology
Phil Taylor, University of Cincinnati - Office of Research
George Turner, Indiana University
Mark Turner, University of Cincinnati - Aerospace
Jason Zurawski, ESnet

Report Editors

Hans Addleman, Indiana University: addlema@iu.edu
Jennifer M. Schopf, Indiana University: jmschopf@indiana.edu
Doug Southworth, Indiana University: dojosout@iu.edu
Jason Zurawski, ESnet: zurawski@es.net

1. Executive Summary

In April 2019, staff members from the Engagement and Performance Operations Center (EPOC) and the The Ohio Academic Resources Network (OARnet) met with faculty and staff from the University of Cincinnati to perform a Campus-Wide Deep Dive. The goal of this interaction was to help characterize a number of scientific and educational drivers for the entire campus, and to enable cyberinfrastructure support staff to better understand the needs of the researchers they support. Material for this event includes both the written documentation from the University of Cincinnati, but also a writeup of the discussion that took place in person on April 26, 2019.

The ongoing challenges that the campus is facing is highlighted in several supporting use cases:

- High Energy Physics
- Department of Aerospace Engineering and Engineering Mechanics
 - Fuzzy AI for Predictive Modeling
 - Gas Turbine Simulation Lab (GTSL)
 - High Fidelity Computations and Models for Advanced Propulsion
- Human Genetics and Genomics
- University of Cincinnati Corrections Institute (UCCI)
- Division of Statistics and Data Science

Each use case highlights specific challenges, but a number of themes emerged that the campus is encouraged to address. In particular, the lack of available local compute and storage resources is impacting several groups and others can greatly benefit from a fully featured Science DMZ infrastructure to facilitate remote data access and transfer.

University of Cincinnati was a previous recipient of an NSF award (CC*II award #1440539) to help support upgrading the campus network, specifically to include a Science DMZ and monitoring equipment. As a part of ongoing support and maintenance of this infrastructure, they are planning to upgrade and augment key components. As part of the overall review, it was determined that there was a need to identify and collaborate with regional or national providers for computational resources in addition to any that may exist locally. Additional challenges with securing sensitive data, cybersecurity, and supporting collaborations were also discussed.

Action items from the meeting included:

1. **University of Cincinnati**, with the assistance of **OARnet** and **EPOC**, will work toward a new research network design pattern, and will attempt to provide a friction free network path to local and remote storage and compute resources.

2. **University of Cincinnati** will explore the addition of local data storage options for university departments that includes a data transfer node, a HIPAA compliant storage solution, and a data transfer mechanism that supports federated identify 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.
4. **University of Cincinnati** will split off the functionality of operating the campus research network from that of the enterprise network. Having dedicated staff for the purpose of engaging with researchers on how to use network infrastructure.
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.

2. Process Overview and Summary

2.A Campus-Wide Deep Dive Background

Over the last decade, the scientific community has experienced an unprecedented shift in the way research is performed and how discoveries are made. Highly sophisticated experimental instruments are creating massive datasets for diverse scientific communities and hold the potential for new insights that will have long-lasting impacts on society. However, scientists cannot make effective use of this data if they are unable to move, store, and analyze it.

The Engagement and Performance Operations Center (EPOC) uses Campus-Wide Deep Dives as an essential tool as part of a holistic approach to understand end-to-end data use. 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.

EPOC supports five main activities

- Roadside Assistance via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers reactively;
- Campus-Wide Deep Dives to work more closely with application communities to understand full workflows for diverse research teams in order to evaluate bottlenecks and potential capacity issues;
- Network Analysis enabled by the NetSage monitoring suite to proactively discover and resolve performance issues;
- Provision of managed services via support through the IU GlobalNOC and the Regional Network Partners;
- Coordinated Training to ensure effective use of network tools and science support.

Whereas the Roadside Assistance portion of EPOC can be likened to calling someone for help when a car breaks down, Campus-Wide Deep Dives offer an opportunity for broader understanding of the longer term needs of a researcher. The Deep Dive process aims to understand the full science pipeline for research teams and suggest alternative approaches for the scientists, local IT support, and national networking partners as relevant to achieve the long-term research goals via workflow analysis, storage/computational tuning, identification of network bottlenecks, etc.

The Deep Dive process is based on an almost 10-year practice used by ESnet to understand the growth requirements of DOE facilities (online at <https://fasterdata.es.net/science-dmz/science-and-network-requirements-review>). The EPOC team adapted this approach to work with individual science groups through a set of structured data-centric conversations and questionnaires.

2.B Campus-Wide Deep Dive Structure

Campus-Wide Deep Dives are basically structured conversations between a research group and relevant IT professionals to understand at a broad level the goals of the research team and how their infrastructure needs are changing over time.

The researcher team representatives are asked to communicate and document their requirements in a case-study format that includes a data-centric narrative describing the science, instruments, and facilities currently used or anticipated for future programs; the advanced technology services needed; and how they can be used. Participants considered three timescales on the topics enumerated below: the near-term (immediately and up to two years in the future); the medium-term (two to five years in the future); and the long-term (greater than five years in the future).

The Case Study document includes:

- **Science Background**—an overview description of the site, facility, or collaboration described in the Case Study.
- **Collaborators**—a list or description of key collaborators for the science or facility described in the Case Study (the list need not be exhaustive).
- **Instruments and Facilities**—a description of the network, compute, instruments, and storage resources used for the science collaboration/program/project, or a description of the resources made available to the facility users, or resources that users deploy at the facility.
- **Process of Science**—a description of the way the instruments and facilities are used for knowledge discovery. Examples might include workflows, data analysis, data reduction, integration of experimental data with simulation data, etc.
- **Remote Science Activities**—a description of any remote instruments or collaborations, and how this work does or may have an impact on the network traffic.
- **Software Infrastructure**—a discussion focused on the software used in daily activities of the scientific process including tools that are used to locally or remotely to manage data resources, facilitate the transfer of data sets from or to remote collaborators, or process the raw results into final and intermediate formats.
- **Network and Data Architecture**—description of the network and/or data architecture for the science or facility. This is meant to understand how data moves in and out of the facility or laboratory focusing on local infrastructure configuration, bandwidth speed(s), hardware, etc.
- **Cloud Services**—discussion around how cloud services may be used for data analysis, data storage, computing, or other purposes. The Case Studies included an open-ended section asking for any unresolved issues, comments or concerns to catch all remaining requirements that may be addressed by ESnet.

- **Resource Constraints**—non-exhaustive list of factors (external or internal) that will constrain scientific progress. This can be related to funding, personnel, technology, or process.
- **Parent Organization**—overview of the sources of funding and cooperation that facilitate the process of science and technology support.
- **Outstanding Issues**—Final listing of problems, questions, concerns, or comments not addressed in the aforementioned sections.

At an in-person meeting, this document is walked through with the research team (and usually cyberinfrastructure or IT representatives for the organization or region), and an additional discussion takes place that may range beyond the scope of the original document. At the end of the interaction with the research team, the goal is to ensure that EPOC and the associated CI/IT staff have a solid understanding of the research, data movement, who's using what pieces, dependencies, and time frames involved in the Case Study, as well as additional related cyberinfrastructure needs and concerns at the organization.. This enables the teams to identify possible bottlenecks or areas that may not scale in the coming years, and to pair research teams with existing resources that can be leveraged to more effectively reach their goals.

2.C University of Cincinnati Campus-Wide Deep Dive Background

In April 2019, EPOC and OARnet organized a Campus-Wide Deep Dive in collaboration with the University of Cincinnati to characterize the requirements for several use cases on campus, including:

- Section [3.1 High Energy Physics Case Study](#)
- Section [3.2 Department of Aerospace Engineering and Engineering Mechanics](#)
- Section [3.3 Human Genetics and Genomics Case Study](#)
- Section [3.4 University of Cincinnati Corrections Institute \(UCCI\) Case Study](#)
- Section [3.5 Division of Statistics and Data Science at the University of Cincinnati Case Study](#)

The University of Cincinnati representatives were asked to communicate and document their requirements in a case-study format (see Section 3). Each Case Study offers a unique view into requirements that the campus can provide, on a number of different time scales ranging from immediate to future needs.

This Case Study is a follow on to the NSF Campus Cyberinfrastructure award, NSF #1541410 , entitled “CC*DNI Engineer: University of Cincinnati (UC) Cyberinfrastructure Engineer and Educator (CI2E)”. This award, which expired in December 2018, was focused on understanding and supporting a number of scientific use cases for the campus.

The CC* project's main objective was to establish a Campus Cyberinfrastructure Engineer and Educator (CI2E) to work side by side with UC's researchers and

students to facilitate the use and integration of local, regional and national cyberinfrastructure components with Science, Technology, Engineering and Mathematics (STEM) research. The position was not able to be retained after the funding ended, which left a critical gap in understanding and supporting the research community on campus.

The face-to-face meeting took place on the University of Cincinnati campus on April 26, 2019, and included in-depth discussion of the Case Studies, campus technology upgrades, and a support structure provided by informational technology and the research community, detailed in Section 5. Next steps are documented in Section 6.

2.D Organizations Involved

The Engagement and Performance Operations Center (EPOC) 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.

The Energy Sciences Network (ESnet) is the primary provider of network connectivity for the U.S. Department of Energy (DOE) Office of Science (SC), the single largest supporter of basic research in the physical sciences in the United States. In support of the Office of Science programs, ESnet regularly updates and refreshes its understanding of the networking requirements of the instruments, facilities, scientists, and science programs that it serves. This focus has helped ESnet to be a highly successful enabler of scientific discovery for over 25 years.

Indiana University (IU) was founded in 1820 and is one of the state's leading research and educational institutions. Indiana University includes two main research campuses and six regional (primarily teaching) campuses. The Indiana University Office of the Vice President for Information Technology (OVPIT) and University Information Technology Services (UITS) are responsible for delivery of core information technology and cyberinfrastructure services and support.

The Ohio Academic Resources Network (OARnet) was created in 1987 by the Ohio Board of Regents, now the Ohio Department of Higher Education, through legislation by the Ohio General Assembly. OARnet was founded to provide Ohio researchers with their first online access to high performance computing resources of the Ohio Supercomputer Center, established in Columbus earlier that same year.

Today, the OARnet network consists of more than 2,240 miles of fiber-optic backbone, with more than 1,500 miles of it operating at ultrafast 100 Gbps speeds. The network blankets the state, providing connectivity to Ohio's colleges and

universities, K-12 schools, public broadcasting stations, academic medical centers, government agencies, and partnering research organizations.

The University of Cincinnati (UC) is a public research university in Cincinnati, Ohio. Founded in 1819 as Cincinnati College, it is the oldest institution of higher education in Cincinnati and has an annual enrollment of over 44,000 students, making it the second largest university in Ohio. It is part of the University System of Ohio.

3. University of Cincinnati Case Studies

The University of Cincinnati presented 6 Case Studies during the Campus-Wide Deep Dive. these are as follows:

- [3.1 High Energy Physics at the University of Cincinnati Case Study](#)
- [3.2 Department of Aerospace Engineering and Engineering Mechanics at the University of Cincinnati](#)
 - [3.2.1 Fuzzy AI for Predictive Modeling Case Study](#)
 - [3.2.2 Gas Turbine Simulation Lab \(GTSL\) Case Study](#)
- [3.3 Human Genetics and Genomics at the University of Cincinnati/Cincinnati Children's Hospital Medical Center Case Study](#)
- [3.4 University of Cincinnati Corrections Institute \(UCCI\) Case Study](#)
- [3.5 Division of Statistics and Data Science at the University of Cincinnati Case Study](#)

3.1 High Energy Physics Case Study

3.1.A Science Background

Michael Sokoloff is a Professor in the Department of Physics. His research involves heavy quark physics and the interplay of the strong and weak nuclear interactions. Of particular interest is his involvement in the LHCb² experiment at CERN³.

LHCb is a high energy physics experiment with a goal of helping to elucidate some of the fundamental questions about the nature of matter and its interactions:

- What are the underlying building blocks of nature?
- What are the symmetries and forces governing the interactions between the underlying building blocks?
- How do these forces produce observable particles (mesons and nucleons) from the underlying building blocks (quarks & gluons)?
- What is the source of the observed asymmetry between matter and antimatter in the universe?

3.1.B Collaborators

The LHCb collaboration has over 850 Members, from 79 institutes, in 18 Countries around the world⁴. Most institutions are European, although the experiment has significant representation from other parts of the world (e.g. U.S., China, Russia, and Brazil). In the U.S., there are several affiliated LHCb sites:

- Los Alamos National Laboratory (LANL)
- Massachusetts Institute of Technology
- Syracuse University
- University of Cincinnati
- University of Maryland
- University of Michigan

3.1.C Instruments and Facilities

LHCb is a detector⁵ at the LHC⁶. Results from any of the LHC detectors are tied to the operation of the LHC itself. All experimental data (real and simulated) are available to all collaborators. As a default, data is available from files stored on the Worldwide LHC Grid (WLCG)⁷.

² <http://lhcb-public.web.cern.ch/lhcb-public/>

³ <https://home.cern>

⁴

https://lhcb.web.cern.ch/lhcb/lhcb_page/collaboration/organization/list_of_members/members_default.pdf

⁵ <http://lhcb-public.web.cern.ch/lhcb-public/en/Detector/Detector-en.html>

⁶ <https://home.cern/science/accelerators/large-hadron-collider>

⁷ <http://wlcg.web.cern.ch>

Data describing collision events in the LHC are produced at CERN. The data streams from CERN to regional centers (often called “Tiers”) around the world where it is archived on tape for long-term storage as well as processing and analysis. Analysis of the data occurs on computers at the participating institutions. Most data read by analysis software is localized, but this is only loosely optimized. Analysis activity often relies on streaming of data between computing centers, bypassing the local storage system. Simulated data is produced by computers at the participating locations and is either stored there for use in analysis. Analysis files created by users range from 10s of MB in size to a few GB. Data files (real and simulated) vary in size from a few GB to 10s of GB.

Present

Currently the facility has entered a period known as a “Long Shutdown” (LS2) after completion of “Run 2” in the fall of 2018. During this period of experimental inactivity, there are several R&D efforts that will take place:

- The experiment is being rebuilt and serviced. Most importantly:
 - All of the charged tracking detectors are being replaced
 - All the front-end electronics is being replaced so it can be read out at 40 MHz
 - The hybrid hardware/software trigger (event selection) is being replaced by a pure software trigger.
- Simulation (e.g. creation of fabricated data sets that simulate a particle collision) will continue, and gradually increase in size and intensity to map the expected output of the detectors once they are upgraded.

2-5 Years

A partial timeline for experimental operation is as follows:

- 2015-2018: Run 2
- 2019-2020: Long Shutdown 2
- 2021-2023: Run 3
- 2024-mid 2026: Long Shutdown 3
- Beyond 2026: “High Luminosity” (HL-LHC)

The computing model changes dramatically in Run 3 and will feature an expected 10% across the board increase in data sizes. LHCb anticipates at least a factor of two increase in data persisted offline.

Beyond 5 Years

The High-Luminosity LHC is expected to come online around 2026. This will result in 10x the event rate, much more data, and more complex events with large event sizes. In order to process, store, and analyze the volume of data, the total computing capacity required by the experiments is expected to be 50-100 times larger than the current capacity, with data storage needs expected to be in the order of exabytes.

The transition from Run 3 to Run 4 requirements will not be as dramatic. The computing requirements for the longer term, compared with Run 3, will increase by no more than a factor of 10.

These requirements are not expected to be satisfied under the flat-budget hardware improvement scenario.

3.1.D Process of Science

Data is recorded from high energy proton-proton collisions at the LHC. Physics conclusions are drawn, primarily, by analyzing highly reduced data sets on computing facilities at home institutions, sometimes using computing facilities hosted at CERN, and potentially using computing facilities in the commercial cloud.

Distributed resources are used to make to make highly reduced data sets that are typically analyzed at home institutions, or NTUPLES, that describe both real and simulated data. The simulated data is produced centrally using WLCG and other resources, such as the Ohio Supercomputer Center and Open Science Grid (OSG) resources made available through the University of Wisconsin. Overall, the experiment generates more simulated data than real data, although this may not be true for all the analyses of interest (as they use some of the largest channels, with billions of signal events).

Present

There will be no new experimental data produced during this phase, but there will be two primary activities to be aware of:

- **Simulation** (e.g. creation of fabricated data sets that simulate a particle collision) will continue, and gradually increase in size and intensity to map the expected output of the detectors once they are upgraded.
- **Reprocessing** existing data sets. Periodically a campaign will begin to read off all raw data, send to participating sites in the WLCG, and re-process again for events. From a network standpoint, this will resemble regular operation of the LHC / WLCG.

2-5 Years

Raw and reconstructed data from selected events for Run 3 (starting to take data in 2021) will be written to permanent storage (tape and disk) at the rate of 5 – 10 GB/s and the current target is 10 GB/s. Assuming 5×10^{E6} seconds per year of running, the project will initially store approximately 50 PB/year.

It is expected that each year will have longer runs. With “reduced” data sets replicated online, and simulated data as well, the experiment will need more than 100 PB/year of additional storage. For 2022 and 2023, the LHCb computing model will require an additional 100 PB of tape storage per year and a similar increase (if not more) in disk storage.⁸

⁸ <https://ipac19.org/wp-content/uploads/2019/05/IPAC19-Conference-Guide-proof-v2.pdf>

In Cincinnati, the Sokoloff group currently stores several terabytes of reduced data. On the time scale of two years, this might increase by a factor of five to ten. There is currently approximately 50 terabytes of usable disk space on local computers. On the time scale of five years, it is likely the team will want to store several hundred terabytes of data in Cincinnati.

Beyond 5 Years

While the models described previously are likely to persist for the next five to seven years, the quantities of data and data transfer will increase by an order of magnitude over this period.

3.1.E Remote Science Activities

All LHC experiments use data from the LHCb experiment at CERN, and their computing resources, as described above. Data can come from any LHC-affiliated partner on the WLCG.

Present

Simulation and reprocessing campaigns will imply a heavy network usage pattern during shutdown.

2-5 Years

There are no expected changes to the patterns, but expect 10% increases in data sizes.

Beyond 5 Years

The present pattern of data exchange between WLCG facilities is expected to continue, but there will be increasing need, especially in HL-LHC operation starting in 2026 for large data flows between UC and additional compute sites such as DOE facilities and commercial clouds.

3.1.F Software Infrastructure

LHCb takes advantage of several software packages related to both analysis and data transfer. Most are open source, affiliated with OSG, and highly specific to the LHC analysis use case.

Present

A selection of software used by the LHCb collaboration at the University of Cincinnati includes:

- NVIDIA drivers: must be maintained to use latest CUDA
- CUDA Toolkit
- Docker
- NVIDIA Docker plugin
- CVMFS: Cern file system; enables use of the LHCb software on Goofy, either directly or in Docker.

- **Anaconda:** Includes a large collection of useful software for us, like several versions of Python, TensorFlow, PyTorch, Jupyter and more.
- **ROOT:** The project may move to just using Anaconda, but for now, is using its own ROOT versions. This is likely to continue until the Conda compilers can be relied on fully.
- **X2Go & Mate desktop:** For fast connections to CERN
- **Network transfer program:** Currently Globus/GridFTP, but that may be losing open source support. Will need to evaluate alternatives which may include Rucio. The department typically move 5 - 10 GB data files from CERN to Cincinnati using Globus/GridFTP tools. For smaller data files or source code files, SCP is also used.
- **LMod:** Environment modules in Lua. Allows the use the “module” command to manage multiple software versions of tools like ROOT.

Ganga and Dirac used to submit many jobs to the WLCG to read data summary tapes (DSTs) and produce NTUPLES. At the moment, the NTUPLE files (in ROOT format) are downloaded to CERN before being shipped to Cincinnati manually.

All of the code for making NTUPLES from DSTs, as well as the code for making DSTs, is developed by the LHCb collaboration and lives in CVMFS. The middleware is generally developed and maintained by the high energy physics community.

The parent networking organization (IT@UC) has, and uses, perfSONAR to monitor performance to external entities.

2-5 Years

Data Transfer tools will change during this time as migration from GridFTP protocols in other parts of LHC (ATLAS, CMS) moves toward Rucio.

It is possible that the use of Ganga and Dirac will change in this time frame. NTUPLE files (in ROOT format) may not have to be curated at CERN, and could be shipped directly to Cincinnati. This would require more sophisticated software infrastructure and careful coordination with the WLCG and probably with the LHCb core computing team.

Beyond 5 Years

Little is known about software in HL-LHC operation, but many patterns are expected to remain similar. Streaming of data to facilitate analysis jobs will stay at current levels, or increase.

3.1.G Network and Data Architecture

The University of Cincinnati Network and Computational infrastructure is described in Section 4.

Present

The Department of Physics is connected to the University of Cincinnati DMZ, UCSciencNet via a 10Gpbs connection. The local network Intermediate Distribution Frame (IDF) houses components for exclusive use of the group for research purposes, as well as affiliated machinery from other researchers:

- 10Gbps copper switch connects all components, and uplinks to campus (enterprise and UCScienceNet).
- LHCb Components:
 - “Conventional” linux node (24 physical cores, 48 logical) with an nVidia P100 GPU plus 2 nVidia Titan V GPUs. Connected to UCScienceNet.
 - Intel Xeon Phi “Ninja” node. Connected to UCScienceNet.
 - Less powerful, “standard” linux node for analysis. (This is not connected to UCScienceNet.)
 - 2U SuperMicro Network-Attached Storage device (NAS) with ~20TB of effective storage (after RAID with mirroring). This is connected to all 4 computational resources. The NAS runs CentOS and presents filesystem as an NFS mount. iSCSI is available, but not used.
 - Connections from the NAS to the compute nodes are 10Gpbs (mixture of optical fiber and copper). Dual NICs are available on certain machines, and can in theory facilitate a 2 x 10Gbps connection.
- Other components:
 - DGX-1 from other physics group. Connected to UCScienceNet.

Machines are administered by the Physics Department, and do not integrate with the greater campus network mechanisms for login/identity. All machines are monitored for access and security. The only complication encountered is maintaining the white lists for Internet2/UCScienceNet access, and keeping several ports open to provide access to software that needs updating.

The Physics department does not have perfSONAR, but has relied on the institutional resources to debug problems.

2-5 Years

The same basic approach to computing for the next 5 years is expected, although with progressively larger datasets. It is not clear if the use of Globus will be precluded due to the future pricing model.

Beyond 5 Years

It is expected that the network architecture and computational framework will undergo significant changes prior to this time. Experimental drivers are unknown, but HL-LHC will imply more data to process, in a shrinking time window between experiments.

3.1.H Cloud Services

Commercial cloud services are not a high priority for the physics group, or LHCb, at this time.

Present

The LHCb group uses WLCG resources, which are a form of Grid/Cloud computing. It would be convenient to have files produced on the grid transferred directly to Cincinnati rather than CERN for trans-shipment. To the extent that CERN uses commercial cloud resources in addition to bespoke WLCG resources, it would be useful to be seamlessly connected to those as well.

2-5 Years

It is expected that some compute resources may move from bespoke nodes to the commercial cloud. To the extent that bespoke resources are used 24/7, they provide the most cost-effective solution. To the extent that they are used substantially less, commercial cloud resources (at discounted rates) become competitive. When elastic resources are required for short periods of time, even some more “standard” commercial rates can be competitive. How this will work in terms of funding agencies and grants is not clear.

Beyond 5 Years

It is expected that computational and storage use will undergo significant changes prior to this time. Experimental drivers are unknown, but HI-LHC will imply more data to process, in a shrinking time window between experiments.

3.1.I Known Resource Constraints

The biggest explicit cost of LHCb computing is that of data storage. The University of Cincinnati does not currently offer a research data storage environment so Physics and other researchers must store their data on local, departmental servers or at external venues, including the Ohio Supercomputer Center (OSC). For the moment, data transfer costs are essentially transparent to the high energy physics (HEP) community – for example, within the United States, most of the cost of operating ESnet (primary long-haul/backbone network between participating facilities) is absorbed funded by the Department of Energy Office of Science. HEP is currently a primary user of ESnet, and it is expected that there will be greater competition for the resources in the future as other scientific programs increase their own data volumes. ESnet will continue to upgrade to accommodate this.

3.1.J Parent & Affiliated Organizational Cooperation

The Sokoloff group collaborated with IT@UC in writing the NSF CC*IIE proposal (award #1440539) that funded construction of UCScienceNet, and is a direct beneficiary of the resources acquired as a result.

3.1.K Outstanding Issues

In the past there has been sporadic network connectivity issues between CERN and the University of Cincinnati that have impacted data transfers. The physics group typically achieves 100MBps (800Mbps) for large files using Globus/GridFTP running parallel TCP streams, but have achieved throughput as high as 280 MBps (2.25 Gbps). Transfers are done out of /tmp typically, to take advantage of RAM and avoid slower local disk. This difference between the average performance of 700Mbps and 2.25Gbps in performance could be related to the use of resources at CERN, which are typically executing jobs for many users concurrently. When using scp on the commercial (non-R&E) internet, the typical transfer speed was at most 2MBps (16Mbps). At this slower speed, individual file transfers could take between one and two hours, with a high probability of failure due to a lost connection at some point in that period.

Most networking issues are solved within the Physics department by local staff, and only issues that impact wide area connectivity are reported or dealt with by IT@UC. IT@UC has, and uses, perfSONAR to monitor performance to external entities in these cases. Local modifications to manage or improve workflows and performance include the implementation of the NAS for local file sharing, and the use of advanced networking tools including Globus GridFTP. As noted, the Globus tools for high speed transfers from CERN to University of Cincinnati may not be used in the future due to the changing model for Globus support and the upcoming changes in the GridFTP protocol. There will need to be investigations in how other options in the future.

3.2 Department of Aerospace Engineering and Engineering Mechanics

Modern aerospace systems are increasingly intricate as technology advances. Successful, safe, and sustainable aerospace engineering requires coordination of many interrelated processes and systems. By encompassing the areas of aeronautics and astronautics, aerospace engineering focuses on the process to analyze, design and implement aerospace systems.

Two use cases from Aerospace engineering participated in the Campus-Wide Deep Dive, Fuzzy AI for Predictive Modeling and the Gas Turbine Simulation Lab (GTSL).

3.2.1 Fuzzy AI for Predictive Modeling Case Study

3.2.1.A Science Background

Kelly Cohen and Anoop Sathyan are working to develop fuzzy logic based artificial intelligence (AI) for various aerospace applications. Fuzzy logic is an approach to AI and computing that bases learning on "degrees of truth" rather than the usual "true or false" (1 or 0) Boolean logic. This nuance leads to a more gentle slope of learning that relies heavily on input and feedback to advance from different situations.

Research in this area is strongly tied to data sets for a specific application. For the immediate term, this is related to aerospace use cases, but the work being performed and has broad applicability to other areas.

Collaborators share data sets specific to a need. There are several pre-processing steps performed to format and prepare the data, and then each set is used to train the Fuzzy Logic AI application. Currently, the entire process is done using local desktop computers, and the datasets involved are on average less than 2GB.

3.2.1.B Collaborators

The collaboration space varies and may change over time. Current collaborators include:

- Cincinnati Children's Hospital
- VegaMX
- China State Shipbuilding Corporation (CSSC)

The team at Cincinnati Children's Hospital is working to capture virtual reality (VR) data to find correlations between an athlete's movements and concussion probability. This data is used to analyze factors involved in concussion which can then be used to train athletes to reduce the chances of such injury.

The team at VegaMX uses satellite image data for monitoring crops grown in particular regions. This data is gathered from public sources that include the U.S. Geological Survey (USGS).

CSSC, a major shipbuilding conglomerate, will be working with the team for the development of an AI that can provide suggestions to the ship's captain to ensure safe navigation in open waters. These suggestions will take into account the positions and velocities of nearby vessels and obstacles along with other weather related parameters.

3.2.1.C Instruments and Facilities

This work is focused heavily on software algorithms and requires access to reasonably small computation and storage work during this phase of research. This will grow over time, as the type and size of data sets grows, and the requirement to process more data at a faster rate, forces the use of high performance computation.

Present

Current work involves the use of personal computational resources (desktops and laptops). This is done due to the relatively small dataset sizes (< 2GB on average), and lack requirements for high performance computation. Long term storage is not a factor given the requirements during development phase.

2-5 Years

Data types are expected to migrate from text toward image and video. As this shift continues, the software will pivot to different types of analysis and processing. This will increase the storage as well as computing requirements. It is likely that a shift to use GPUs will take place in order to speed up processing. Cloud Computing is also being examined as a possible option.

Beyond 5 Years

Data storage and computing requirements are expected to keep increasing.

3.2.1.D Process of Science

The current workflow is focused on research and development more than a production pipeline:

1. Accept training data set from collaborators. Exchange is typically via email, but could also utilize shared cloud storage (Google Drive, Box).
2. Data is pre-processed as needed before being fed into the Fuzzy Logic AI system.
3. Local processing and storage are used for the running of the code against the training data set.
4. Results are shared with collaborators.
5. Modifications to the code, as needed, are performed locally.

Once the system is developed, the basic workflow will remain similar, but could involve more automated ways to handle the pre-processing and processing steps that involve the use of high performance computation and shared storage.

Present

All processing is done on desktop and laptop systems. Multiple cores are used to speed up the computation, but an HPC/HTC environment has not yet been explored.

2-5 Years

Additional storage and computing resources are needed as the application set expands to include image and audio processing. The move to HPC and HTC is expected, as the processing capabilities of a single machine will no longer be sufficient.

Beyond 5 Years

Little is known about the future, but additional growth in storage and processing are expected.

3.2.1.E Remote Science Activities

At the current time, there are no remote activities beyond sharing of input data sets, as well as results with collaborators. This is not expected to change into the future.

3.2.1.F Software Infrastructure

Software development is an active part of this project.

Present

Most algorithms are developed in MATLAB on local machines. Python is also used for some of the codes that handle processing and pre-processing.

2-5 Years

Software use during this time will remain similar. The use of parallel libraries will be explored when the project moves toward using HPC/HTC resources. .

Beyond 5 Years

Little is known about the future, as new languages and development environments are not known at this time.

3.2.1.G Network and Data Architecture

At the current time there are no special network or data requirements that cannot be met by the institutional network infrastructure.

3.2.1.H Cloud Services

Current work is done exclusively in the local environment, with the exception of data sharing which may involve the use of cloud services. A pivot to use more external resources (commercial and R&E sponsored) is expected.

Present

Data sharing, via institutional BOX access, is used for storing and sharing data. No additional cloud processing is used at this time.

2-5 Years

There are plans to explore the use of cloud use, e.g. Google cloud, AWS. Processing of audio and video files will require more computation than is currently available within the local group.

Beyond 5 Years

Most processing and storage will be done via the cloud.

3.2.1.I Known Resource Constraints

There are no known resource constraints at this time.

3.2.1.J Parent & Affiliated Organizational Cooperation

Networking is handled by IT@UC, and there are no requirements to upgrade for this work. Computation is currently handled locally.

3.2.1.K Outstanding Issues

There are no known outstanding issues to report at this time.

3.2.2 Gas Turbine Simulation Lab (GTSL) Case Study

3.2.2.A Science Background

The Gas Turbine Simulation Lab (GTSL) is a part of the Department of Aerospace Engineering that runs engine component and full-engine simulations. The primary research conducted in this lab includes:

- Development of improved computational fluid dynamics (CFD) modeling algorithms, a branch of fluid mechanics that uses numerical analysis and data structures to analyze and solve problems that involve fluid flows.
- Development of improved methods of simulation to better model flows of interest in the jet engine or gas turbine community
- High fidelity unsteady time marching CFD of 3 blade rows of transonic fan
- Multidisciplinary and multi-fidelity optimization of turbomachinery in gas turbine.

3.2.2.B Collaborators

Collaborators include:

- Stanford Linear Accelerator Center (SLAC)
- NASA: both the NASA Glenn Research Center and the NASA Advanced Supercomputing (NAS) facility at NASA Ames Research Center.
- The Ohio Supercomputer Center
- Air Force Research Laboratory, Wright-Patterson Air Force Base, Dayton Ohio
- Industrial partners including GE and Pratt-Whitney

3.2.2.C Instruments and Facilities

High Performance Computing (HPC) facilities are a critical part of the workflow of GTSL. The problems are considered small to medium size, utilizing, for example, approx 20 cores for 5 days, and scale to multi-core and multi-processor environments easily. The optimization problem can be parallel-parallel, and the unsteady problems run on 500-1000 cores if available.

Present

Data varies by use case, using the examples from 3.2.2.A:

- The data sets for CFD algorithms are typically 200MB, but rapid communication is required for debugging while doing code development.
- The directories associated with flow model simulation are 5-40GB, and are generated every few days.
- The high fidelity unsteady time marching CFD of 3 blade rows of transonic fan data sets are 5-15 GB, generated weekly.
- The multidisciplinary/multi-fidelity optimization of turbomachinery in gas turbine data sets are 0.5 GB, and will be generated twice daily.

2-5 Years

More sophisticated modeling associated with HPC and HTC resources will allow more variables to be simulated at once - producing higher resolution models. Data growth during this time is anticipated to match the sophistication of the hardware.

Beyond 5 Years

Little is known about the trajectory of simulation software and hardware at this range, data growth is expected to increase exponentially.

3.2.2.D Process of Science

The primary science activities within GTSL are algorithm development, improved process for simulation, improved understanding of Unsteady Turbomachinery flows including Boundary Layer Ingestion and film cooling, as well as optimization of turbomachinery for jet engines.

Present

GTSL focuses on 4 major activities, each with a different subset of technology requirements.

GTSL is making improvements to CFD modeling algorithms through the simulation and rapid development of high-order concepts. The iterative workflow for this activity typically consists of:

1. Algorithm development
2. Simulation by running an algorithm numerous times on HPC/HTC resources, with different input variables
3. Evaluation of results and modification of algorithms.

The input data sets for CFD algorithms are typically on the order of 200MB, with an output data set sizes that are smaller. Rapid data movement capabilities are required for debugging purposes while doing code development, e.g. the developer needs feedback quickly based on how the algorithm may be reacting over the running of a set of code.

A second area of focus involves improving simulation methods to better model and explore flows of interest in the jet engine or gas turbine community. This work uses large eddy simulation (LES), a mathematical model for turbulence used in CFD.

Workflow for this activity is similar to the prior example:

1. Simulation development and deployment
2. Numerous trials on HPC/HTC resources
3. Capturing and visualizing output
4. Evaluation of results, and modification of simulations

The outcome of this work is to improve understanding of the flow physics via simulation. The data output associated with flow model simulation are 5-40GB, and are generated every few days due to the iterative nature of the work and the time

required to run the simulation. The size is in part due to checkpoint files as well as visualization files that can be larger.

The third area of focus is examining high-fidelity aerodynamics and unsteady time-marching CFD research as it relates to the design of 3 blade transonic fans. This particular work is very data intensive, as the input data set and snapshots during execution must be saved for use during post-processing work. These data sets can be 5-15 GB, and are generated on a weekly basis due to the iterative nature of the work and the time required to run the simulation.

The last area of focus involves multidisciplinary and multi-fidelity optimization of turbomachinery in gas turbines. This research uses geometry generation, CAD, CFD, and finite element methods. Multi-fidelity is required to work with the design methodology, and the final solutions need to be saved and compared. Resulting datasets are 0.5 GB, and are generated twice daily due to the iterative nature of the work and the time required to run the simulation.

In all of the above cases, the “final data” for each topic needs to be stored. Currently, postprocessing is best handled at a local resource so files do not need to be returned over a large distance. Alternatively, a high bandwidth, low latency network could allow for remote interaction with the “right” post-processing software. This has not been tested fully at this stage.

2-5 Years

Most of the above work is performed on local server resources. Transition to HPC class nodes will happen in the next two years. The file size numbers will expect to increase by a factor of four in two years.

Open source codes are expected to replace the current commercial codes. That will open up the usage of the capability by a factor of 10 in this time frame.

It is expected that automated on-the-fly post processing will be available. This way only final scene rendering and statistics need to be returned which should lower the amount of data by a factor of 10. The net amount of network needs would then be the same.

Beyond 5 Years

Grid sizes will be 1000 times greater and whether files are stored or cases rerun will be dictated based on the speed of the network and storage devices. It may be more cost effective to rerun large cases.

3.2.2.E Remote Science Activities

R&D on the simulation software is done within GTSL, but testing and simulation work uses a combination of local and remote resources. This is not expected to

change, although the number and location of remote resources is harder to predict beyond 5 years.

Present

Most work is done local to GTSL using in-house computational resources. In some cases, time can be used on NASA NAS or Ohio Supercomputer Center (OSC).

2-5 Years

As the complexity of simulations and algorithms increase, HPC use will be a stronger requirement. In the absence of local HPC/HTC, GTSL will rely on the resources of collaborators (AFRL, NASA, XSEDE).

Beyond 5 Years

Little is known about this period, but there is an expectation that HPC / HTC resource use will increase during this period of time along with the data volumes produced.

3.2.2.F Software Infrastructure

No data management software is used currently. Post processing is done by commercial CFD and FEM codes, but could use more home-grown or open source software.

Present

In terms of in-house software used on a regular basis, the following flow solvers are used in GTSL:

- FINE Turbo
- MSU Turbo
- CFD++
- FDL3DI
- CFX
- FLUENT
- NCC
- LEWICE

Additionally, GTSL uses a number of visualization platforms:

- Visual3
- TV3
- PV3
- Tecplot
- Sleipnir (Under Development)

Currently, scp is used for the majority of data movement activities. Globus GridFTP are being actively integrated into the workflow.

2-5 Years

While more open source solvers will be used, more data will be reduced remotely and only processed data returned. This will require the use of customized software.

Beyond 5 Years

More customized software is expected to have data analysis co-processes with the data produced. This will keep the size of data transferred fairly constant even as the number and size of simulations grows by at least four orders of magnitude.

3.2.2.G Network and Data Architecture

At the current time there are no special network or data requirements that cannot be met by the institutional network infrastructure.

Present

At the present time, there is a 100Mbps commodity link that serves GTSL and is used by all resources (workstations, servers, etc.). GTSL is working with IT@UC to ensure that there is a 10Gbps clean path in place to external locations (NASA, OSC) in the later half of 2019.

2-5 Years

GTSL will continue to work with IT@UC to upgrade capabilities locally, and to remote locations. It is not expected that all GTSL resources will require 10Gbps, but a data transfer server (for communication with external parties) that has Globus/GridFTP available will simplify some of the external work flows.

Beyond 5 Years

Not much is known about this period, other than data volumes and reliance on external computation will increase.

3.2.2.H Cloud Services

No commercial cloud services are expected to be used in the near or future terms.

3.2.2.I Known Resource Constraints

No constraints exist besides funding and size of potential users. The market for improvements in tools for the jet engine and airplane community is large, but finite.

3.2.2.J Parent & Affiliated Organizational Cooperation

There are no additional details that can be provided for this section.

3.2.2.K Outstanding Issues

The largest constrain that GTSL is dealing with is a shared LAN infrastructure between components that limits capacity to 100Mbps. This is a part of a legacy network connection and hardware that IT@UC is working to upgrade to 1Gbps

locally, and then link to a 10Gbps externally facing connection to facilitate reaching other facilities. At the current time, GTSL does not have a system or storage that can keep up with 10Gbps speeds, but will be trying to upgrade internal capabilities to support this.

The NASA collaboration involves files that are shared with researchers at NASA Glenn Research Center, and exchanged with the NAS facility at NASA Ames Research Center. During this data movement, there are often issues with accounts due to a lack of federated access (must create and maintain local accounts on the storage location). A 3rd party location of data storage would be helpful. This is still a relatively small project, but 8 UC students are involved as well as about 6 external researchers.

3.3 Human Genetics and Genomics Case Study

3.3.A Science Background

The Division of Human Genetics (DHG)⁹, a part of the Cincinnati Children's Hospital Medical Center (CCHMC)¹⁰ and in collaboration with the University of Cincinnati College of Medicine Department of Pediatrics¹¹ hosts the work that Yaping Liu and his staff are performing. Of particular interest are epigenomic and gene regulation mechanisms in cancer and other common complex diseases, along with affiliated work on liquid biopsy, computational biology/bioinformatics, gene-regulation, cell-free DNA, exosomal-DNA and single-cell -omics.

Despite the successes of genome-wide association studies (GWAS), important challenges remain and limit the impact of GWAS on biology and medicine, especially for non-coding variants that are still poorly understood. Identifying causal GWAS single nucleotide polymorphisms (SNPs) and their targeted genes in the relevant cell types is one of the major challenges in the post-GWAS era for common diseases. Recent publications have shown that circulating cell-free DNA (cfDNA) in blood is a promising non-invasive biomarker to longitudinally bring up multi-omics information from relevant cell types that are dying during the disease progression. Here, the goal is to establish cfDNA and other related blood/urine-based biomarkers (e.g. exosomal DNA) as a platform to dissect the gene regulatory circuits behind non-coding GWAS SNPs.

The primary research activities in the Liu lab are currently supported by start-up funding. The research project is a collaboration across multiple physicians and PIs within CCHMC and the UC Medical School, as well as external PIs from the College of Medicine in the University of Pittsburgh and the Broad Institute of MIT and Harvard. Most of the data will be multiple different types of next-generation sequencing results generated internally from the lab and obtained from public databases. The dataset with downstream analysis will be also created during the research and shared with the international community through publications, databases, and websites.

3.3.B Collaborators

The Liu lab has active collaborators within CCHMC that include:

- Louis Muglia, 1-3 people to share the data through HPC resources that are working to process and present output
- Brian Weiss, 1-2 people
- Richard Lu, 1-3 people
- Craig Erickson, 2-5 people

⁹ <https://www.cincinnatichildrens.org/service/h/human-genetics>

¹⁰ <https://www.cincinnatichildrens.org/>

¹¹ <https://med.uc.edu/>

- John Harley, 3-6 people
- Matt Weirauch, 2-4 people
- Jose Cancelas-Perez, 1-2 people

Additional local collaborators within the College of Medicine include:

- Jiajie Diao, 1-2 people to share the data through HPC resources that are working to process and present output

External collaborators that also share data through HPC resources that are working to process and present output:

- College of Medicine in the University of Pittsburgh, PA - Zongqi Xia, 1-3 people
- UCSD, CA - Bing Ren, 2 people
- Broad Institute of MIT and Harvard, MA - Viktor Adalsteinsson, and Gaddy Getz, 5-15 people total

3.3.C Instruments and Facilities

There are a number of instruments and external facilities that are critical to the mission of the Liu lab. Some of these are local, and others are remote. All are capable of producing large data sets that require access to storage and processing.

Present

The Liu lab is currently using the University of Cincinnati Flow Cytometry Core Facility, and a variety of genome sequencers to perform research. These are typically called “next generation” machines due to their operational and data production profile. These include

- Novaseq 6000 (located in the Liu lab)
- HiSeq2500 that is located in the UC Gene Expression Core Facility as well as available via partnerships with external commercial vendors (e.g. BGI, Novogene)

Sequencing requires access to extensive processing and storage needs. These are available in a variety of locations:

- The Biomedical Informatics (BMI) at CCHMC provides HPC resources¹² including 100 CPU cores and 5Tb storage.
- The Ohio Supercomputer Center (OSC) has 4.3Pb storage and several compute resources including:
 - Pitzer Cluster: A 10,240-core Dell Intel Gold 6148 machine with a theoretical system peak performance 720 teraflops (CPU only)
 - Owens Cluster: A 23,392-core Dell Intel Xeon E5-2680 v4 machine with a theoretical system peak performance 750 teraflops (CPU only)
 - Ruby Cluster: A 4,800-core HP Intel Xeon machine
 - All OSC systems now support GPU Computing.

¹² <https://bmi.cchmc.org/resources/clusters/computational-cluster>

- Bridges, hosted by the Pittsburgh Supercomputer Center (PSC), has computational nodes that can supply 1.3018 Pf/s and 274 TiB RAM. The Bridges system also includes more than 6PB of node-local storage and 10PB of shared storage in the Pylon file system.
- The San Diego Supercomputer Center (SDSC) supports:
 - Comet supercomputer: A ~2 PFlop/s system featuring 1,944 nodes, each with two 12-core Intel Haswell processors, 128 GB memory and 320 GB of flash storage; and 4 large memory nodes, each with 4 Intel Haswell processors and 1.5 TB of memory;
 - Comet GPU: 72 GPU nodes: 36 NVIDIA K80 GPU nodes; 36 NVIDIA P100 GPU Nodes; each has 4 GPUs/node;
 - Data Resources: Over 7 PB of high-speed storage made available via Lustre parallel file systems, as either short term Performance Storage used for temporary files, or long term, non-purged Project Storage that persists for the life of the project. A Durable Storage resource provides a second copy of all data in Project Storage file system;

To keep up with present demand, the Liu lab needs local access to approximately 200 CPU cores, 1 GPU (Nvidia V100) node, and 100Tb storage.

2-5 Years

It is anticipated that the following upgrades will be undertaken in this time:

- The HPC in BMI or UC should meet or exceed 500 CPU cores and GPUs, 500TB-1PB data storage
- The next generation of sequencing machines (beyond Novaseq 6000 capabilities) will be available and used more widely
- Upgrades to the flow cytometry core facility will be required
- Upgrades to the genome editing core facility will be required

Beyond 5 Years

It is anticipated that the following upgrades will be undertaken in this time:

- The HPC in BMI or UC should meet or exceed 1000 CPU cores + GPUs, multiple PB to 1 EB of data
- Sequencing cost will reduce and resolution will increase
- Upgrades to the UC Flow Cytometry and Genome Editing Core Facilities will be mandatory as technology improves

3.3.D Process of Science

The process of science for the Liu lab is heavily based on sampling and sequencing, with analysis and result dissemination after. A basic workflow consists of:

- Sample preparation (wet lab)
- Sequencing, etc
- Data preprocessing as needed (smoothing)
- Data analysis using exiting codes
- Data storage (raw and analyzed pieces)

- Dissemination to collaborators

Present

The primary source of data used by this workflow are the outputs from genomics sequencing machines, including the Novaseq 6000 and HiSeq2500, that are located in the University of Cincinnati Gene Expression Core Facility or through external commercial vendors, including BGI and Novogene. The datasets output from these sources are approximately 2GB-100GB for each sample representing an individual genome. The gating factor is the length of time needed to pre-process the sample, and post process the resulting data set.

The analysis requires advanced computation and storage, running sequencing codes. The Liu lab currently uses local resources to store the zipped raw .fastq file outputted from the sequencing hardware, which may have been created locally or may be physically mailed from the external sequencing site. The first step of the analysis involves the use of the BWA/Bowtie package which reads the sample, and maps to reference genome files. The output of this process is a 'bam' file¹³. The second step of the analysis includes extracting the genotype, fragment length/coverage, DNA methylation, and other epigenetic information from the bam files.

The third step involves performing association testing with the phenotype and gene expression information to check the effect of genetic variation and epigenetic variations. The fourth step uses GPUs running machine learning tools (including deep learning) to study the fragmentation patterns in cell-free DNA and use it to decode the non-coding genetic variations. After these operations, the University of Cincinnati Flow Cytometry Core Facility is utilized to sort the cells into different subtypes and obtain the pure cell type for the downstream sequencing activities.

2-5 Years

It is anticipated that the Liu lab will sequence more and more. There are plans to develop large scale map-reduce computational framework to process the files.

Beyond 5 Years

It is anticipated that the Liu lab will continue trend of sequencing more samples at higher resolutions. Work will shift to different diseases vs. healthy conditions. It is expected that a large scale map-reduce computational framework will be developed to process the files.

3.3.E Remote Science Activities

The main driver of remote activity is the location of sequencing resources and processing facilities. Currently, some sequencing is done offsite, but as new devices are added to the local environment this will become less important. Computation is

¹³ <http://software.broadinstitute.org/software/igv/bam>

another driving factor. There are limited shared resources within the UC environment, implying that most computation must be done externally.

Present

The Liu lab will need to connect frequently with the following resources:

- The local BMI cluster in CCHMC Filesystem access makes transfer less involved than to external sites.
- Transfer the data to/from Ohio Supercomputer Center to transfer the data for the computation.
- Transfer the data to/from Pittsburgh Supercomputer Center
- Transfer the data to/from San Diego Supercomputer Center
- Transfer raw data from the public database resources, including BGI and Novogene, to help the data analysis

2-5 Years

The Liu lab will need to be able to store their data reliably over time, and have a stated preference for centralized storage at UC. Current storage at the various computational resources is often temporal, and this will not scale for research efforts.

Beyond 5 Years

Little is known about the changes beyond 5 years, it is expected that data set sizes will increase, but hardware could become more integrated with cloud solutions.

3.3.F Software Infrastructure

Software in the Liu lab is a mixture of proprietary tools that are associated with the specific instruments, community/open source tools for analysis, and internally developed scripts to assist with workflow. Almost all software is supported by members of the lab or local IT staff.

Present

Some of the software includes:

- CUDA for GPU
- HPC job schedule system (e.g. Slurm, LSF9 or PBS)
- Compilers & Programming Languages: c/c++, oracle Java 6, 7, 8, python, perl, R
- MATLAB

Data transfer software includes a mixture of tools including Globus, Aspera, lftp, and scp.

2-5 Years

The HPC/HTC stack of software is not expected to change significantly, but the proprietary software for tools may. It is also expected that routine data transfer

using high-performance and ubiquitous tools (e.g. Globus/GridFTP) will become more common.

Beyond 5 Years

There is little known about this time range, basic software needs are expected to remain stable or shift more towards service models.

3.3.G Network and Data Architecture

At the current time there are no special network or data requirements that cannot be met by the institutional network infrastructure.

3.3.H Cloud Services

Cloud service use is still in nascent stages. Cost is a factor, which also limits investment in converting codes and workflow tools.

Present

Both of Amazon AWS and Google Cloud resources have been used for prototyping. Costs remain high, so this is not a desirable resource when compared with the availability of the limited local and remote R&E resources.

2-5 Years

Costs are expected to drop, making the ability to burst into the cloud a more desirable feature for research use. Newer sequencing hardware may be more closely integrated with cloud resources.

Beyond 5 Years

With data volumes increasing, and cost to operate decreasing, it is expected that cloud will be more routine.

3.3.I Known Resource Constraints

Storage and computation are the largest factors in limiting research output at this time. The workflow to sequence, process, and store results will not change, but will increase in terms of volume and frequency.

Present

Persistent storage and backup capabilities (measured in the 100s of TBs range) is a requirement for the Liu lab. A near term need is the ability to share the results of research through an online portal/database. In addition to the base amount of storage, 10Gbps or greater network capacity to all aspects of the research workflow is required.

2-5 Years

Persistent storage and backup capabilities (measured in the 500TB - 1PB range) is an expected requirement for the Liu lab. Ability to use local CPU and GPU resources

(1000 CPU cores + at least 2 GPU nodes) with the distributed framework (e.g. Hadoop/SPARK) is highly desirable.

Beyond 5 Years

Persistent storage and backup capabilities (measured in the 10s of PB range) is an expected requirement for the Liu lab. Additional local CPU and GPU resources are expected. Peering relationships to facilitate access to cloud storage and computation is also expected.

3.3.J Parent & Affiliated Organizational Cooperation

The Liu lab uses extensive resources in the BMI cluster at CCHMC. These, in conjunction with other UC resources, are critical to the process of science.

The Ohio Supercomputer Center is providing a certain amount of startup support in the form of storage and cycle access.

The Liu lab has applied for time in XSEDE and has received allocations at the Pittsburgh Supercomputer Center.

3.3.K Outstanding Issues

Currently, the most challenging part of research in the Liu lab is data storage. Due to the multi-TB data sets that are generated monthly, there is not enough space to support immediate and long-term research needs. The cost of storage is prohibitive, and will impact the ability of the lab to be productive. A cost-effective solution for the UC community is desirable.

3.4 University of Cincinnati Corrections Institute (UCCI) Case Study

3.4.A Science Background

The University of Cincinnati College of Education, Criminal Justice, and Human Services (CECH) is home to a variety of academic, service, and research oriented centers that were created to serve as resources for students, staff, and community partners. These centers help facilitate partnerships between CECH and agencies within the Great Cincinnati-area and beyond. Research is statistical in nature, and involves analysis of records from law enforcement agencies. The CECH houses two of these centers discussed below - the UC Corrections Institute (UCCI) and the Institute of Crime Science (ICS).

UCCI has a mission to research, develop, disseminate, and implement evidence-based practices in corrections. This is an applied research program that includes both outcome and process evaluations. This effort impacts corrections agencies and organizations far and wide. From California to Maine, it has conducted and implemented research in all 50 states and spans international borders to Scotland, Singapore, and New Zealand. Serving both for profit and non-profit agencies, UCCI provides services for federal, state, local, and international governments as well as professional organizations to promote effective interventions for adult and juvenile offenders.

ICS combines the knowledge and skill of both academic researchers and criminal justice practitioners to solve real world problems. The ICS team includes world renowned experts in criminal justice research and law enforcement professionals who are subject matter experts. ICS team members' areas of expertise include: policing, violence reduction, violent street gangs, racial profiling, police legitimacy, social network analysis, co-offending networks, crime analysis, police staffing and program evaluation. ICS delivers evidence-based, empirically tested solutions, technical support and training to: national, regional, state, local and international law enforcement and criminal justice agencies.

3.4.B Collaborators

UCCI partners with agencies across the United States and its territories. Currently, UCCI has research partnerships with:

- Ohio Department of Youth Services
- Ohio Department of Rehabilitation and Correction
- Nevada Department of Public Safety
- Indiana Office of Court Services
- Vermont Department of Corrections
- Judiciary of Guam
- Ventura Co, California.

- Alvis, Inc (through a federal grant)
- Ohio's Bureau of Criminal Investigation

The ICS partners with agencies across the United States and Canada, including:

- The Cincinnati Police Department
- The Ohio Department of Public Safety
- The Fort Myers Florida Police Department
- The Hamilton County Heroin Coalition
- BrightView Treatment Centers
- Cordata Healthcare LLC

Each of these partnerships involves sharing of information and results. In some cases, data sets may simply be text records, and in other cases they can be video or audio files, for example, police body and dashboard camera footage.

3.4.C Instruments and Facilities

UCCI and ICS primarily uses computational and storage resources within the University of Cincinnati. Some are owned/operated locally, others are operated by affiliated groups such as IT@UC and the UC School of Medicine. In the case of the later, HIPAA compliant resources are a requirement for certain data sets.

Present

ICS primarily uses crime data from law enforcement entities and provides analytics using the same data. Once ICS has the data, it is moved into an in-house MySQL database, and stored on secure UC servers, owned by CECH and located in CECH. ICS has one dedicated server and has the ability to purchase additional CECH server space as needed.

2-5 Years

ICS will require HIPAA compliant server space for research related to opioid addiction. HIPAA compliant server space will need to be purchased from the Medical School who currently have HIPAA compliant servers located in the UC Datacenter or other entities on campus that have been certified.

Beyond 5 Years

ICS anticipates moving all of its databases, analytics, and applications to the cloud and using the web to access them for shared access and secure environments.

3.4.D Process of Science

UCCI and ICS data is accessed daily from remote locations. Access to each form of data implies a different access method:

- SFTP via agencies that support it, with storage into ICS created and housed databases.
- Proprietary dual authentication process for the remote data transfer to certain police departments

- External access to shared Microsoft Access and Excel resources
- Paper records transfer, which require entry into ICS created Access/Excel resources and eventual population into the ICS current analytics system
- Encrypted emails (limited file sizes)
- UC's secure BOX subscription

Present

Data is collected directly from police departments and other law enforcement entities via paper files as well as through online file transfers. When data is transferred electronically it is through encrypted email, SFTP, or through UC's secure Box. Data is then uploaded into MySQL and analyzed using SPSS or STAT.

UCCI partners with the School of IT for cloud based services to collect and store data agencies enter through a cloud based risk assessment system, developed by the School of IT. Data is entered directly into the system from agency staff.

ICS currently receives data daily from the Cincinnati Police Department. The data is transferred (using a dual authentication system) from a CPD server to a secure UC ICS server every morning at 2AM via an SFTP process developed by ICS. Once the data is transferred, it is extracted from the UC ICS server and downloaded into the ICS data visualization and analytics system that was developed in-house using MySQL and Java Script. ICS currently houses over seven million CPD data points.

ICS also receives data from seven other local police departments. Data is retrieved weekly, via a password protected direct access portal developed by PAMET,- the agencies' records management system vendor. Once the data is retrieved, ICS downloads it to the analytics and visualization system that was developed in-house using MySQL and Java Script.

2-5 Years

Police and corrections departments typically are not on the bleeding edge of technology. Many still use legacy versions of Microsoft operating systems and older forms of data organization and curation. Many still rely on paper records due to concerns about being able to manage and keep up with technology.

Enhancements in this time frame are not known, but are expected to become moderately automated. It is expected that as more records become digitized, there will be friction in dealing with file sharing, curation, and search. The HIPAA aspects of data will dictate technology changes and adoption patterns.

Beyond 5 Years

In the era of big data, there is a drive to have more automated ways to search and curate records. Not much is known about the trajectory, but computation and secure storage will continue to need to grow.

3.4.E Remote Science Activities

All data has a custodial home (the location where the data is generated and initially stored) in a different location than UCCI and ICS. As a result of this, most interactions with external entities involves using remote resources. Currently all processing and analysis are done at UCCI and ICS, but there are efforts to explore cloud services.

Present

Due to the support of legacy systems at partner institutions, there are no strong changes to the aforementioned methods of access and process. The only changes that are immediately expected are the addition of new sites and partners that will use similar retrieval methods.

2-5 Years

Changes on the collaborator side will force local changes. As more information becomes digitally native, challenges will crop up related to:

- Secure transmission of large data sets
- HIPAA compliant storage that scales with data growth
- HIPAA compliant processing that scales with data growth
- Support for legacy (e.g. Microsoft Windows, Access, Excel) systems on newer infrastructure

Beyond 5 Years

Not much is known beyond this time frame, but there are expectations of larger and more numerous data sets that can be mined for research.

3.4.F Software Infrastructure

Software is mostly developed in house for research purposes. Off the shelf tools are used for collaboration needs.

Present

CECH IT deploys, maintains, configures, and creates software packages using Microsoft System Center Configuration Manager (SCCM). The SCCM server is hosted by IT@UC.

CECH also uses:

- SFTP servers
- Microsoft file share, including home and shared drives on the CECH Cloud
- File Sender
- UC Box
- OneDrive
- Quest enterprise file management, for reporting and auditing

ICS uses MySQL, JavaScript, HTML5, D3.JS, Crossfilter.JS, and Bootstrap for its records management, analytics and visualization platforms.

2-5 Years

Secure transmission and storage of data will be more common as data sets grow and are digitized. It is expected that additional data management systems (e.g. map-reduce, elastic) will become more widely available for search and curation.

Beyond 5 Years

Not much is known beyond this time frame, but there are expectations of changes to the software stack as older systems are phased out

3.4.G Network and Data Architecture

At the current time there are no special network or data requirements that cannot be met by the institutional network infrastructure.

Present

IT@UC manages the backbone (switches, cabling etc.) network infrastructure. CECH IT supports from the wall (data jack) out.

All School of Criminal Justice (CJ) computers, including those utilized by the UCCI and ICS, are tied to the UC active directory domain. CJ staff have access to limited personal and shared storage space that is provided by CECH IT. The storage environment is a Microsoft storage cluster that is housed in the university datacenter. CJ has several virtual servers that are housed in the UC datacenter running on the CECH IT infrastructure. CJ also has access to Box and OneDrive for file storage.

2-5 Years

HIPAA compliant infrastructure will continue to be critical for some projects. UCCI and ICS may not have the ability to install and operate this infrastructure independently, thus will be looking to partner with IT@UC and the School of Medicine as needed.

Beyond 5 Years

Exponential growth is expected as more records become digitized.

3.4.H Cloud Services

UCCI and ICS partner with the School of IT for cloud based services. This is likely to increase as the number of sites increases.

Present

ICS and UCCI do not currently use cloud based services beyond secure file sharing.

2-5 Years

There is an anticipated need for cloud services as additional projects come on-line, and data needs to be migrated to the cloud to better run web-based applications.

Beyond 5 Years

Less is known for this period of time, but expectations are that cloud use will become important as partner agencies migrate.

3.4.I Known Resource Constraints

ICS has a current need for coders and money to pay them, for cloud based services and storage, for dedicated servers to house videos from law enforcement clients, and for the ability to certify one of the existing servers as HIPAA compliant.

3.4.J Parent & Affiliated Organizational Cooperation

ICS has a grant from OHIO 3rd Frontier to pay for limited opioid-based analytics and also a minor opioid evaluation grant from the Comprehensive Addiction and Recovery Act (CARA) through the Department of Health and Human Services.

3.4.K Outstanding Issues

Data storage for sensitive data (HIPAA) is a critical need for the current and future time scales.

3.5 Division of Statistics and Data Science at the University of Cincinnati Case Study

3.5.A Science Background

The Division of Statistics and Data Science (DSDS) consists of research faculty members, including Emily Kang, whose research focuses on developing statistical methodology and principled data analytical techniques to advance scientific discoveries in environmental, biomedical, and engineering sciences.

The faculty members are in particular interested in advanced statistical theories and scalable methods to better model and analyze hundreds of Gigabytes of remote sensing data from NASA, or with complex structures, such as data with varying spatio-temporal dependence, imaging data, or dynamic network data. The statistical methodological research is motivated by challenges in remote sensing, studies of climate change and related mitigation strategies, and biomedical research. It provides rigorous statistical tools to reveal, quantify, and validate scientific hypotheses in the presence of multiple sources of uncertainty.

Faculty members in DSDS collaborate widely, thus their research often requires downloading and storing data sets that are hundreds of Gigabytes from various external facilities. Once data is available, the software and the analysis results are also stored and shared with collaborators.

3.5.B Collaborators

Faculty member in DSDS regularly collaborate with the following external entities:

- Cincinnati Children's Hospital
- Jet Propulsion Laboratory
- The Pennsylvania State University
- Purdue University
- Case Western Reserve University
- The University of Alabama

Most, if not all, of the research requires downloading and storing hundreds of Gigabytes from various entities that include, but are not limited to:

- NASA and NOAA missions and observational data
- Anonymized biomedical data from collaborators at other institutes such as Vanderbilt University
- Data from the Alzheimer's Disease Neuroimaging Initiative (ADNI)

After the teams research work is completed, it is routine to return the result data sets to the same collection of collaborators.

3.5.B.1 Jet Propulsion Laboratory (JPL), Pasadena, CA

Data sets are shared in various ways between DSDS faculty members and their collaborators at JPL:

- JPL provides a large file transfer facility which is the most common way for JPL collaborators to share data with DSDS faculty.
- Some datasets are hosted by other parties, such as NOAA or NASA, and can be directly downloaded from website portals. Using this method, DSDS faculty and staff can directly download and store at local computational resources or the Ohio Supercomputer Center (OSC).
- Other Cloud sharing tools such as Box or Google drive are occasionally used when the researchers need to share data of small amounts of data, such as summaries of analyses, summaries of simulation results, drafts of technical reports, slides or manuscripts.

3.5.B.1 Cincinnati Children's Hospital Medical Center (CCHMC), Cincinnati, Ohio

Data sets are shared in various ways between DSDS faculty members and their collaborators at CCHMC:

- Some datasets are hosted by other parties (e.g. ADNI¹⁴) and can be directly downloaded from website portals. Using this method, DSDS faculty and staff can directly download and store at local computational resources or OSC.
- Other sharing tools such as USB drives, Box, or Google drive are occasionally used when the researchers need to share small amounts of data such as summaries of analyses, summaries of simulation results, drafts of technical reports, slides, or manuscripts.

3.5.C Instruments and Facilities

Data analysis and processing can be done using a mixture of resources at the University of Cincinnati or shared computational resources hosted elsewhere such as the resources at the Jet Propulsion Laboratory. This varies from researcher to researcher but is an acknowledged gap in resources given the amount of processing and storage that will be required into the future.

Present

Computation is currently done using the following resources:

- Local computation, purchased/hosted by DSDS faculty via grant funding. This is typically limited.
- Some DSDS faculty members have access to JPL clusters, while international students (even students on related projects) may not have access to JPL clusters due to restrictions on access to non-citizens.
- DSDS faculty members submit proposals to OSC and then use assigned resources for their teams to work on related research projects.

¹⁴ <http://adni.loni.usc.edu/>

Storage is another area of concern, as there is not a centralized storage mechanism for local users. This limits the size of data sets that can be processed for both local and remote use cases.

2-5 Years

With the growth of the size of data, it will be challenging to store or compute the data sets that are being imagined. By way of example, consider Sea Surface Temperature (SST) data from NASA. The data set is about 300MB daily (with expectation of growth as instruments increase their resolution). Space limitations have restricted the study to use data prior to the year 2002. Currently, due to space limitations, the research team have to focus on a small study region rather than investigating the SST changes at the global scale.

Beyond 5 Years

Data sets will continue to grow due to increases in resolution, amount captured, and availability of metrics. It is anticipated that cloud computing and storage will be critical, given the shortage of local resources to address the problem. There is not a clear idea of how data repositories will be linked to these resources at this time.

To continue the example of NASA, it is unclear if NASA will make storage/computation available to assist researchers. The current model is to download data and perform calculations locally, but questions have been raised about efficiency. Will researchers continue this method, or will computation/storage closer to the data facilitate a better workflow. What is the trade-off between data size vs. accuracy, model complexity vs. computing complexity?

3.5.D Process of Science

The workflow of DSDS faculty consists of:

1. Develop statistical methodology and models
2. Encode ideas into software using R, SAS, Matlab, Julia, Python, etc.
3. Identify relevant data sets
4. Acquire computational time, either local or remote, and download the data needed
 - a. Use of CPU and GPUs is common
 - b. Prototyping and longer-term runs are possible using resources on campus, once developed, as well as OSC.
5. Pre-process the data to eliminate noise, as needed
6. Run models, perform hundreds of simulations for validation and uncertainty quantification, and make changes as appropriate
7. Save the results and share them

Present

DSDS faculty need to design and run extensive simulation studies to validate the proposed new methods. These simulation experiments are usually run with

hundreds of replications in different settings so that researchers can investigate the robustness and properties of the proposed method, compare its performance with the state of the art, and perform uncertainty quantification. This step of simulation studies is an integral part of statistical methodology research as it provides empirical proof, explanation, and justification for utilizing the proposal statistical methods to analyze the scientific data and then address the scientific problems.

One of the key steps of the process discussed above is pre-processing. For example, satellite data is often noisy and can be spatially incomplete. It is necessary to process these datasets to smooth noise and fill in gaps in the data. When possible, using data from other missions or instruments) can provide a finer resolution or higher accuracy, and be more reliable in terms of uncertainty.

Researchers also need to carry out extensive simulation studies to validate the methods used and to compare the performance of the approach with other research results. The resulting methods are typically shared with collaborators as code and may be adopted or adapted for use at a collaborators institution.

2-5 Years

More storage (i.e., one month of Level-1 NASA OCO-2 data, approximately 600GB) and more powerful HPC facilities will be needed. The size of satellite data is constantly increasing. To enable DSDS faculty and their students to develop and apply their methods to these data, they will need to download and save these satellite data to HPC facility (their own personal computers become infeasible due to limited storage). To analyze or visualize these large datasets will require computing facilities with more nodes than currently available in order to shorten queuing time and to enable large-scale simulations.

Beyond 5 Years

Hopefully NASA and other agencies can provide a more flexible and powerful data portal that will allow users to access in cloud and do some analyses in cloud, rather than requiring researchers to download. In addition, it is expected that additional theoretical foundations and computational algorithms will be developed that will enable users to reduce data sizes without sacrificing the accuracy of the data analysis.

3.5.E Remote Science Activities

This team may continue to utilize remote data sets, computation, and storage. Without dedicated resources, it is expected to remain this way into the future.

Present

Researchers are primarily relying on external computation and storage for work, as there are not solutions provided by the department, college or university.

2-5 Years

Data sets will grow, and there is not a roadmap as to when HPC or HTC resources will be available at the University of Cincinnati. It is expected that DSDS faculty will invest time into exploring other external options where they can be found, including XSEDE, NSF-funded clouds, or commercial clouds. The lack of long-term storage options is likely to limit the ability to with with growing data sets.

Beyond 5 Years

The advancement of hardware used for analysis is not well known for this time period. It is expected that data sets will grow, and that tools will adapt to take on larger data volumes.

3.5.F Software Infrastructure

The software development environment is fluid and a critical part of the research process for DSDS faculty.

Present

DSDS faculty members develop methods and algorithms in R, Matlab, Python, and Julia at present time. Data is transferred via SFTP.

2-5 Years

The use of more advanced computational platforms such as HPC, HTC, or Cloud computing may encourage the use of other environments and tools. Data organizational methods and APIs are likely to evolve.

Beyond 5 Years

The advancement of software used for analysis is not well known for this time period. It is expected that data sets will grow, and that tools will adapt to take on larger data volumes.

3.5.G Network and Data Architecture

At the current time there are no special network or data requirements that cannot be met by the institutional network infrastructure.

3.5.H Cloud Services

DSDS has interest in exploring cloud services. Current use is limited to data sharing, typically through personal and institutional accounts such as Box.

Present

There are minimal explorations and use of the cloud beyond as a data sharing mechanism. The cost and time needed to migrate workflows to cloud computing environments is prohibitive.

2-5 Years

Computational and storage needs will grow, making cloud resources more attractive. Some data sets may become inherently tied to a cloud provider, such as those that NASA and NOAA provide.

Beyond 5 Years

Exponential data growth will outpace local and shared resources. Cloud operation is expected to be common.

3.5.I Known Resource Constraints

A succinct plan for institutional storage and computation resources is critical for DSDS faculty and staff. Without access to these resources, the ability to operate on large data sets (i.e., around 20 Gbytes per day for Level-1 NASA OCO-2 data, and 300MB per day for NASA SST data) is becoming severely limited.

3.5.J Parent & Affiliated Organizational Cooperation

UC has submitted proposals for NSF funding to add additional cores to the existing HPC Cluster. DSDS faculty members also apply for OSC individually for computing resources.

It will be good to know how faculty members should budget HPC resources in their proposals, which may require faculty members work and negotiate with an HPC facility, either OSC or the latest UC HPC clusters.

3.5.K Outstanding Issues

Due to need for extensive simulation and large-scale data analyses in research, DSDS faculty have requested access to additional HPC facilities. OSC is suitable for now when there is not also a need for more than 100GB storage space for data. A UC centralized HPC facility set-up would increase the flexibility of the research work. Such HPC resources not only are important to faculty's research but also should be included in training and education for the next generation of STEM workforce.

4. University of Cincinnati Network and Computational Environment

The University of Cincinnati network is operated by IT@UC.¹⁵ There are two network infrastructures maintained by campus:

- Enterprise Network (See Section [4.1 University of Cincinnati Enterprise Network Diagram](#))
- Science Network (See Section [4.2 University of Cincinnati Science Network Diagram](#))

This section also reviews the computational environment in Section [4.3 Computational Environment](#).

¹⁵ <https://www.uc.edu/about/ucit/about.html>

4.1 University of Cincinnati Enterprise Network Diagram

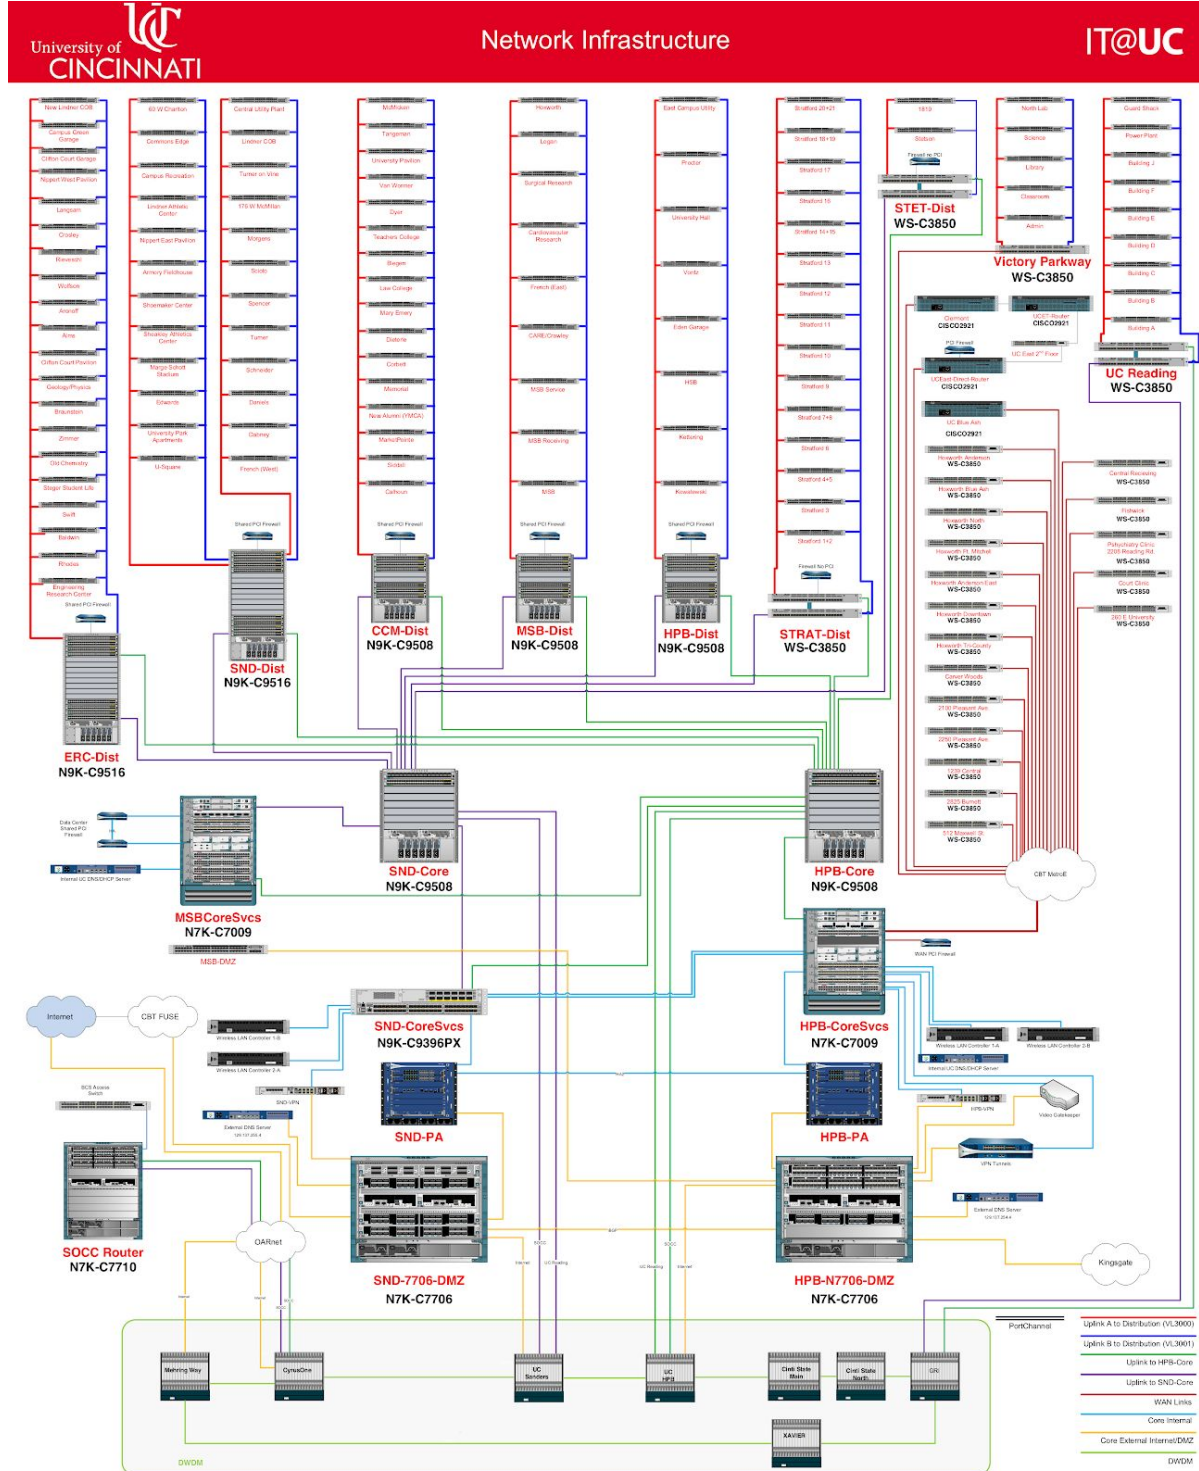


Figure 1 - Spring 2019, University of Cincinnati Enterprise Network

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.

External 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 has been constructed to prevent loss of service to any institution in the event of a fiber cut.

To improve research capabilities, IT@UC added a dedicated 100Gbps circuit to the CERF ring alongside the existing 10Gbps commodity circuits. The enhanced capability enables direct researcher access to a 100Gbps pipe, via the UCScienceNet (UCSN), UC's DMZ, connecting the university's main campus to OARnet's 100Gbps Internet2 backbone.

4.2 University of Cincinnati Science Network Diagram

UCScienceNet (UCSN), a 100Gbps Science DMZ modeled after ESnet's Science DMZ design, incorporates perfSONAR nodes 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 enables high-throughput transfers with minimal latency to ensure rapid delivery of large scientific data sets. The hardware employs bandwidth scalable from 40Gbps depending on research requirements and 100Gbps delivery from the aggregation layer outward to Internet 2 and other research and education networks .

UCSN, servicing eight research intensive locations within engineering, physics, criminal justice, and medicine, provides a friction-free network, creating a true Science DMZ to address the limitations of the existing enterprise network. IT@UC, in partnership with the UC Office of Research, provided funding to add additional endpoints to UCSN to expand the benefits of a high-speed networking to researchers in libraries and mathematics not connected during the initial deployment of UCSN.

This expansion required the deployment of Cisco Nexus 3000 switches deployed in strategic research areas. The Nexus 3000 will provide scalable 40Gbps back to the research core and 10Gbps to the high performance computing equipment, which is currently limited to 10Gbps or 100Mbps.

Expanding the friction-free campus networking architecture by eliminating campus and building level network infrastructure constraints, will enable formation of a research ecosystem encouraging diverse, multidisciplinary collaboration and partnerships to address complex Grand Challenge problems.

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. perfSONAR has been deployed to gather throughput statistics that are relevant to the use cases of researchers on UCSN.

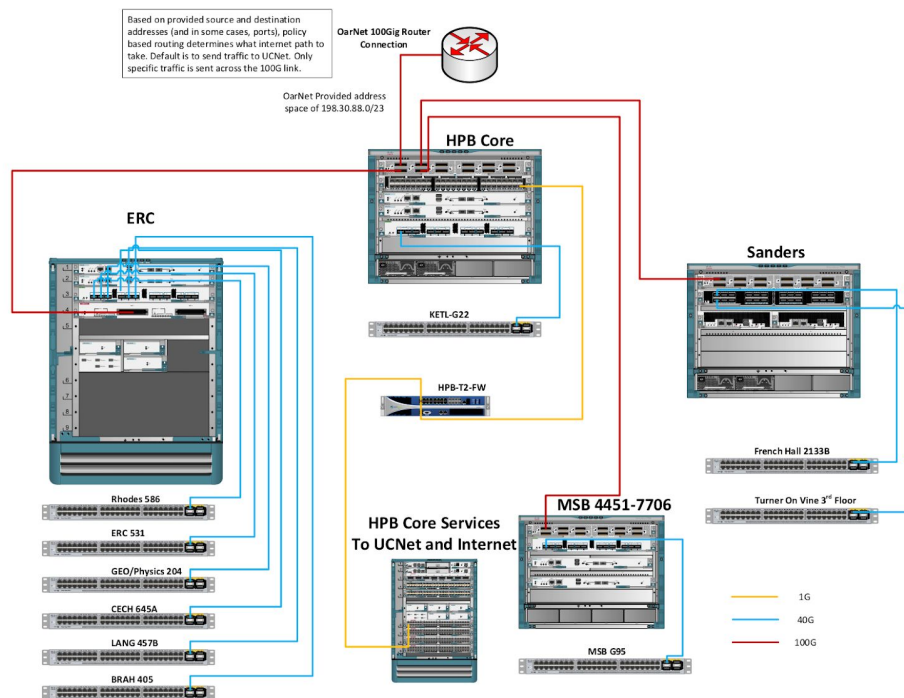


Figure 2 - Spring 2019, University of Cincinnati Science Network (UCSN)

4.3 Computational Environment

IT@UC's Research and Development department leads the IT Governance's Research and Development Topical Committee, which is charged with identifying, prioritizing and recommending computational resources for researchers. Membership includes computational researchers, undergraduate and graduate students, IT@UC cyberinfrastructure support personnel, and distributed research IT support personnel. Dr. Phil Taylor, Assistant Vice President, Research Infrastructure

and Development and Annette Ready, Associate Vice President, IT@UC, Innovation and Partnerships, are co-chairs of the committee.

In 2018, a partnership between the Office of Research, IT@UC and a set of the university colleges, departments, and researchers developed the Advanced Research Computing (ARC) initiative. ARC is a pilot trial of a central HPC cluster that will be available to all university researchers. The initial cluster was funded by the partnership, and an NSF MRI proposal was submitted in January 2019 to expand the cluster. A Faculty Advisory Committee is charged with developing a sustainability plan and business model utilizing the metrics from the 18-month pilot.

The ARC HPC Cluster configuration supports both CPU and GPU nodes that are connected with high-performance 100 Gbps Omnipath (OPA) interconnect. The 16 CPU nodes are utilizing Intel Xeon Gold 6148 2.4G, 20C/40T, 10.4GT/s, 27M Cache, Turbo, HT (150W) DDR4 2666 RAM 192GB per node, DDR4 2666. Overall this supports 50 teraFLOPS of peak CPU performance. The two GPU nodes are NVIDIA Tesla V100 32G Passive GPU with a peak performance of 224 teraFLOPS and a ZFS Storage Node supporting 96TB of raw storage, although initially configured to offer 43TB.

Research project storage is provided on an XAScaler 7990 EDR appliance with approximately 1.5 PB of raw storage. The parallel file system is integrated into the UC Central HPC switch fabric for maximum, non-blocking I/O performance and security. Performance is 15 GB/s sequential write, 20 GB/s sequential read. The system is scalable to multiple petabytes of capacity and can easily be expanded to meet additional storage demands over time.

5. Discussion Summary

On April 26, 2019, members of the EPOC team and staff from OARnet met with faculty and staff from the University of Cincinnati. This review was held in Cincinnati, Ohio, on the campus of the University of Cincinnati.

During the discussion, the following points (outside of clarifications to the Case Studies described in [3. University of Cincinnati Case Studies](#), and technology in [4. University of Cincinnati Network and Computational Environment](#)) were emphasized. The discussion is organized by subject matter area for readability.

5.1 High Energy Physics

A representative from the HEP program was not available during the in-person discussion. The review was read, and discussion focused on two key areas:

- Expectations regarding the schedule of LHC operations through the next 5+ years
- Data movement complications experienced by the research group

On point one, the publically shared schedule for LHC operations is known to be:

- 2015-2018: Run 2
- 2019-2020: Long Shutdown 2
- 2021-2023: Run 3
- 2024-mid 2026: Long Shutdown 3
- Beyond 2026: “High Luminosity” (HL-LHC)

This will impact all experiments (CMS, ATLAS, LHCb, etc.). While shutdown at this time, the collaborations are actively engaging in activities that are still data and network intensive:

- **Simulation** - This involves:
 - Creating simulated detector data sets (centrally managed by each detector group) at the various computational sites
 - Exchanging the data via transfers
 - Running analysis jobs on the simulation to improve the overall software operational profile
 - The simulated dataset sizes will gradually increase to match detector expected values.
- **Reprocessing** - Re-reading the tape output, which contains the complete raw data from Run 2, and sending to sites for analysis. This can be viewed in the same manner as a typical experimental run. The purpose of this is to re-analyze all the data looking for additional events of interest.

The discussion on data transfers focused on the use of Globus GridFTP. In January of 2018, Globus announced¹⁶ it was dropping long term support for the GridFTP

¹⁶ <https://www.globus.org/blog/support-open-source-globus-toolkit-ends-january-2018>

product, including both the underlying protocol and the set of related command line tools, and would only be supporting a non-compatible version going forward by licensed agreements. As a result of this, the Grid Community Forum¹⁷ created the open source Grid Community Toolkit (GCT) to support the free version of the GridFTP software.

Without having a member of Physics available to address the original concerns, there are two ways forward to address the changes in Globus GridFTP. University of Cincinnati can go forward with campus-wide adoption of the licensed version of Globus that would be usable by all colleges and departments. The Physics department could take advantage of this instead of acquiring their own licence. Alternatively, the Physics department could look into GCT to see if that solution would be more viable for the LHCb use case. . However, it is likely that LHCb is following the lead of other LHC experiments and moving to Rucio¹⁸, a new platform for data movement.

5.2 Department of Aerospace Engineering and Engineering Mechanics

Three use cases for Aerospace were discussed in detail (two of which were profiled via Case Studies, a third that was presented in person without a Case Study). GTSL has the most critical current needs with regards to networking, storage, and processing, and will require intervention by IT@UC immediately. Others are more future facing, but will benefit from the same infrastructure upgrades.

5.2.1 Fuzzy AI for Predictive Modeling

Work within this field is relatively new, and is still establishing many parts of the workflow. There are several opportunities to streamline and speed up the processes related to basic operations:

- **Data Movement** - Current methods mostly involve sharing of training data via non-sophisticated methods, for example, emailing spreadsheets or small sets of text-based data inputs. As more complex input is explored, for example audio or video files, the ability to move large data sets will become more critical to the process and require advanced tools.
- **Analysis** - Most algorithms used by this team can run on personal computing resources, scaling only at the core to processor level. No effort has been put into using non-local HPC or HTC resources. With the migration to larger and more complex input, it is anticipated that additional compute resources. These could be provided locally by UC or by commercial entities, such as Cloud services.
- **Storage** - Given the current data set sizes for spreadsheets and text input, storage is not an issue. Files are shared when they are needed between collaborators and there is no central method to store input or results. With a move to more sophisticated input centralized and scalable storage will be a factor for both internal and external collaborators.

¹⁷ <https://gridcf.org>

¹⁸ <https://rucio.cern.ch>

This project is part of a UC program that is looking at future technology and will receive some funding and support from the university to assist with technology requirements going forward. Support and upgrades may include additional network connectivity, and dedicated HPC resources.

5.2.2 Gas Turbine Simulation Lab (GTSL)

GTSL is an established facility with a number of partners federal agencies (NASA, NOAA) and industry (GE, Pratt-Whitney). The team has access to resources within the R&E community, including OSC and XSEDE, to drive much of the day to day technology requirements with regards to networking and computation.

Simulation is a core activity and product of this group. It is heavily leveraged as new algorithms and techniques are researched. For simulations to run seamlessly, there are several requirements:

- **Computing** - The simulations are highly parallelizable, thus most, if not all, of the work requires access to computation. Development of codes can be done locally, but testing can require many hundreds of hours on advanced computation resources that are often external to the facility. For example, when CFD codes are being tested to measure the placement of 'cooling holes' on a turbine fan (e.g. modeling to see if temperature can be controlled due to the physical design of the equipment), there can be a tremendous amount of individual data points modeled and measured. There can be as many as 27 million grid points per hole, and there may be 200 holes for each blade of a fan, with up to 64 fan blades. This resolution produces 100s of GBs of data per run, and changing a single input variable.
- **Storage** - Simulations by themselves are small, but there is a need to scale. All told, data is measured in GBs for an average dataset, and depends on these factors:
 - Each input data set must be saved
 - Each simulation output must be saved
 - Snapshots (moments in time during simulation run) are often saved when they can be. This helps restart in the event of a failure, as well as learn about algorithm run behavior.
 - Simulations are run hundreds of times, as input variables change.
- **Fast network access** - this depends on the type of collaboration:
 - Analysis Facilities: If the code and inputs are run on external computation, the results and snapshots will need to be transferred back to UC.
 - Research Collaborators: Sharing analysis and visualization products after research is critical.
- **Low latency between components** - During a simulation run, ideally results are visualized in real time, which is dependent on the speed in accessing the data. Some of the machines within the GTSL are a fast enough compute resource and have low enough latency for the data transfers to enable 'live'

viewing of the visualizations. This is easier if the computation is local to UC, and much harder as the computation location gets further away. At a minimum, this host should be considered for DMZ access on future networks.

Discussion and needs focused on these major themes:

- Upgrading network connectivity (in progress)
- Storage that will scale into the future that allows for local use, and integrates with local computation
- Ways to share results with collaborators
- Plans to increase local computation
- Ways to support visualization needs that may be remote

5.2.3 High Fidelity Computations and Models for Propulsion

This area of research was presented without a written Case Study. The overview of the research centered on the computational modeling that is being done, and the data challenges that exist. Highlights included:

- Heavy storage requirement related to snapshots. Due to the nature of the computation, there are individual checkpoints in the running of analysis that must be backed up to aid in visualization as well as recovery from errors. Each checkpoint is about 4-5GB in size, and there may be thousands of checkpoints taken during an execution. If this work is done remotely, the checkpoints must be transferred back to UC.
- Certain models are very CPU intensive. In some cases, a simple model can consume 250 million CPU hours, which is not readily available on UC systems. Use of external computation is a necessity and may occur at a number of locations (Georgia Tech, NASA, XSEDE, OSC, etc.).
- All data products are stored at UC long term, even those produced offsite. Data movement is an ongoing challenge.
- Resolutions of the simulations are increasing as the availability of CPU and storage increases. When storage or CPU is a limitation, the simulations are run at lower resolutions, which can be less appropriate for research approaches.

5.3 Bioinformatics

The medical and bioinformatics fields have well documented data growth patterns that are directly related to the improvements in instrumentation and the process of science:

- Tooling, in the form of sequencers and microscopes, has increased resolutions and speed to process, all while dropping in cost. Machinery that used to cost millions of dollars can now be purchased for thousands.
- The availability of tooling, previously only located at large specialized facilities, has increased to the point where most major facilities have several instruments available.
- The number of use cases has increased, and the time to process each sample has decreased, meaning that the number of associated output datasets has grown.

With these observations comes a simple set of facts:

- Data from these instruments must be stored.
- Data from these instruments must be processed.
- Data from these instruments must be shared.

CCHMC and UC Medical are at a critical point in the support of their researchers to provide enough processing and storage to facilitate the usage patterns of instruments now and to be able to grow into the coming years. Currently the output of a single genomics processing device can be 20-30GB, and hundreds of samples can be run during a single day across all research groups in the facility. Thousands of results are likely to be produced in a single year. Every one of these samples must:

- Be analyzed through advanced tools.
- Be able to have both the raw and processed results stored in a secure manner.
- Be shared with collaborators in a secure manner.

Local processing operated by BMI and UC cannot keep up with the current demands, and some research groups utilize resources at external locations (XSEDE centers, OSC, etc.). However, with remote processing comes the need for remote transmission of data sets, which relies on additional network connectivity.

Recommendations are for IT@UC to work with the Medical school to alleviate some of this pressure, namely in the form of centralizing storage and processing to help with the growing needs identified by the researchers.

5.4 University of Cincinnati Corrections Institute (UCCI)

Work within UCCI is gated on three primary roadblocks, all of which are not easy to offer an immediate solution:

- Data sources, primarily corrections and police departments, often use antiquated technology that are less effective than current approaches, but cannot be adapted.
- Data privacy is critical for all work that is being done.
- Data sets are increasing, both in terms of the breadth and complexity of data.

The first item is related to the collaborators they work with, which included various departments of corrections and police agencies around the world. Each has a different method for sharing and processing data. Some may be electronic, some may be physical. Of the electronic varieties, file formats and the databases in use are generally several years older than industry standards. Sharing methods are not sophisticated, and must be secure at all times, which adds overhead to data sharing as well.

Data privacy, given the nature of the work that is being done, is always a concern. In particular all systems and storage must be FISMA/HIPPA certified when dealing

with human subject data, for example medical records. For example, UCCI is working to store and process data related to ongoing work in opioid task forces for the state of Ohio that required a statewide collaboration space to access the sensitive data. The availability of certified technology to deal with this requirement is limited, and UCCI relies on infrastructure provided by other local sources, including the UC Medical School. Ideally, IT@UC could create and maintain this infrastructure for all departments that need it.

Lastly, data volumes are increasing in terms of the number of sources and the complexity of the data. More and more collaborating agencies are sharing records with UCCI. Given the complexity of the records that are shared in terms of format and content, the initial processing and storage steps can be very time consuming. The data itself is changing. Text-based records will always be present, and both audio and video are becoming more common. For example, UCCI will be the home to numerous dashboard and body camera videos from collaborating police departments. This data is typically uncompressed HD video and can span full shifts for a given officer. These records will need to be stored, cataloged, and searchable for a legally mandated number of days. During this time, UCCI staff may need to search for and share selected content when asked. The complexity of this task implies a need for significant processing, storage, networking, and sharing capabilities.

5.5 Division of Statistics and Data Science

The DSDS research is at a critical point with regards to technology adoption. Researchers are not able to make progress in processing data sets because they do not have access to enough local compute or storage resources

Figure 3 references the phenomenon known as the “long tail” of data: all of the data produced for a given field prior to a specific point (e.g. the “tail”) is equal to the amount of data being produced in a single period of time that is much shorter (e.g. the “head”). This impact has been documented in many fields including physics, astronomy, and genomics.



Figure 3 - Long Tail Diagram

For example, one research project is working with a set of NOAA-hosted weather and climate data. that spans several decades. . Due to size, only a subset set of the data can be operated on at a time. Figure 4 describes this workflow. This storage limitation also limits the ‘layers’ of data that can be explored. For example, a certain

simulation may function with many metrics (air temperature, water temperature, CO² measurements, etc), but data size restrictions may limit the input data set to fewer variables.

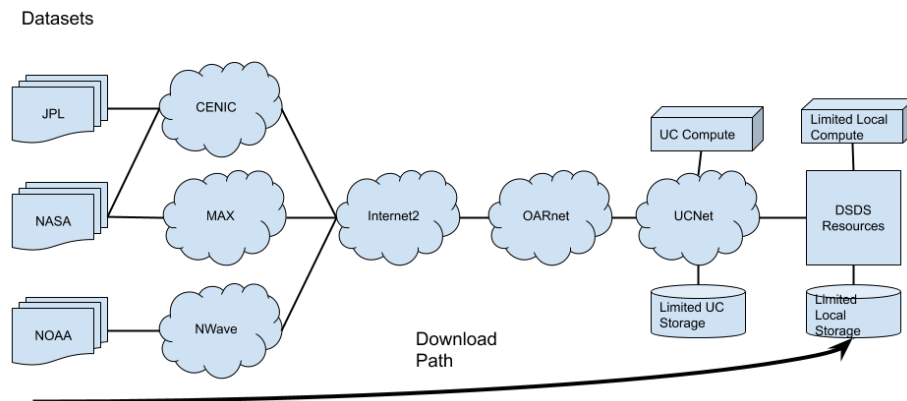


Figure 4 - DSDS Workflow Diagram

DSDS wants to prioritize the following research-critical requests:

- Having local storage to pull down larger and more complex data sets, estimated to be in the 100s of TBs to 10s of PBs in the coming years.
- Having access to on-site computing so they don't need to use external resource allocations to perform daily computational tasks.
- Understanding bottlenecks that exist when downloading data sets from external repositories, as well as better data transfer tools that can be employed to accelerate that process.

Further out investigations will involve:

- Exploring if the role of GPUs can accelerate research productivity, and finding available GPU resources.
- Acquiring computation time from external resources, as they are not always available at the current site.
- Exploring ways to utilize public cloud resources so that work can scale and burst as it becomes more cost effective.

5.6 University of Cincinnati Network and Computational Environment

The UC Science Network was designed with a specific set of use cases and access policies. Some of the findings during discussion included:

- Altering the peering arrangement with the regional networking provider, OARnet, which provides commercial and R&E connections. Current there are two peering points that are treated as different networks and land on different devices.
- Currently there is a physical separation of the networks that service enterprise connectivity and science connectivity. This separation has caused friction, in that:

- Resources onboarded for use on the research network can typically only reach locations with R&E connectivity (e.g. non-commercial resources), which results in challenges for certain use cases that span both worlds. As an example, consider a server that routinely exchanges data with XSEDE - the data transfer path would work fine. If the server ever needed to download a software update not available via R&E connectivity, it could not do so which puts the machine at risk for compromise through the use of older software packages.
- Resources not-onboarded for use of research network items are forced to access sites via the institutional firewall. This can be seen when trying to data “within” the campus, a user on the enterprise network must transfer data out of the enterprise network to the OARnet peering, and then back through the research network.
- Campus users were brought on specifically by IT when a need was identified
 - Only certain resources were permitted onto the network after review
 - Only certain resources could be accessed on each network
 - The lack of overlap is significant, and makes the research network less usable, thus there is much less of a benefit to users to migrate to use the science network infrastructure
 - The lack of use of the science network infrastructure has caused questions whether it is feasible for the long term
- There is a limited set of external locations that can be accessed via the research network - not the entire R&E routing take (e.g. white-list controlled). Thus when a user needs to access a site, they must make a formal request and wait for it to be allowed, which may take hours to days.
- The science network and enterprise network are operated by the same team of people. This is typically not a problem, but it means that enterprise-level operational approaches and behaviors are used for research, although running research networks often requires a different set of skills to execute properly.
- The dedicated ‘cyber engineer’ resource was lost due to a lapse in funding. This resource was traditionally the way that research use cases were identified and adapted.
- There are currently technology limitations that prevent high speed flows. Specifically, the use of smaller and less capable “data center” network hardware, that exists in the path of WAN-facing use cases that require more capable/larger capabilities (e.g. to support bulk data movement). The data center-type switches have lower latency, but do not allow for long-distance, high-bandwidth use cases to experience high performance due to a lack of buffering.
- There is no campus-level solution for dedicated data movement, either hardware or software. Some users have access to approaches in this space, but this is a very small number related to existing use cases, for example the LHCb work in the Physics department.

After discussing the research needs with users, and understanding both current and future needs, a set of requirements emerged that will be considered for future IT@UC services and upgrades:

- **Campus DMZ Refresh** - The IT@UC has funding to upgrade portions of the DMZ infrastructure and will be making a plan to do so. Aspects that will be considered include:
 - Creating a more open access policy for researchers that have an identified data needs, in additional departments to the current set
 - Adding more sites to the 'white list' of accessible endpoints, including known facilities for computation and collaboration. Ideally this would involve allowing most if not all of the R&E routing table.
 - Facilitating access to commercial sites and simplifying the BGP setup of the two networks.
 - Upgrading edge and core devices to support high-performance data movement needs by eliminating switching and routing bottlenecks.
- **Storage Support** - Procurement, deploying, and operation of a storage solution for the research community, as shown in Figure 5. This will include a more flexible allocation policy and enabling easier data sharing between groups. Storage will be more strongly integrated with other parts of the campus cyberinfrastructure, including both computation and data movement.




Figure 5 - Storage Pyramid

- **Data Movement Support** - Installation and operation of data movement hardware and software, for example Globus GridFTP or similar tools, that can be used by any researcher with a UC account. Integration with campus storage and compute solutions.
- **Computation Support** - Procurement, installation, and operation of campus HPC and HTC resources for research use. Availability of staff that can integrate workflows from research groups to utilize local computation instead of remote uses. Integration with campus storage and data movement solutions.

- **Research Engagement** - Reviving the CI engineering role to directly engage with researchers on an ongoing basis. Repeating this event with more groups to understand the needs and fix problems.
- **Exploring Sensitive Data Support** - Working with groups that require HIPPA enclaves (UCCI, UC Medical), and/or ITAR, to provide this service centrally. Will reduce cost and increase usability for a wider selection of campus. Note that the establishment of sensitive data infrastructure (particularly with ITAR), may imply forcing construction of new, separate, infrastructure that do not co-exist.




Figure 6 - Campus Network Diagram (High Level)

As shown in Figure 6, a new model for the network was discussed, which has several key features:

- Redundant connections to OARnet (100Gbps primary and 10Gbps backup)
- Blended connectivity via two campus core routers. This will prevent the ‘access’ issues that dominated the earlier design, where it was physically impossible to access parts of the global internet infrastructure.
- The Enterprise network is segregated from the research network, and can continue to be operated using current practices.
 - It is recommended that all BYOD and office machines (even though that may have a research use case) be installed on the Enterprise network.
 - Dedicated research machines should be considered for use on the Science network.
- The Science network will have two primary modes of operation:

- **Data Mobility:** The data transfer infrastructure should be directly connected to the main routing nodes to simplify the data path. These resources can be sealed to general access (e.g. no login) and only expose the data movement interface via a portal or set of Globus software, for example. They should have access to group storage for research.
- **General Access:** Other resources that require faster networking and have been pre-cleared and certified, can be considered for access to the Science network.
- Storage must be provided for campus users, and this storage must be integrated with instruments, data transfer hardware, and HPC hardware.
- Instruments (e.g. sequencers, microscopes, etc) are sensitive devices that must be protected. Separate their data paths and control paths (the later using firewalls and filters), are needed.
- Data paths, as shown in Figure 7, are now simplified

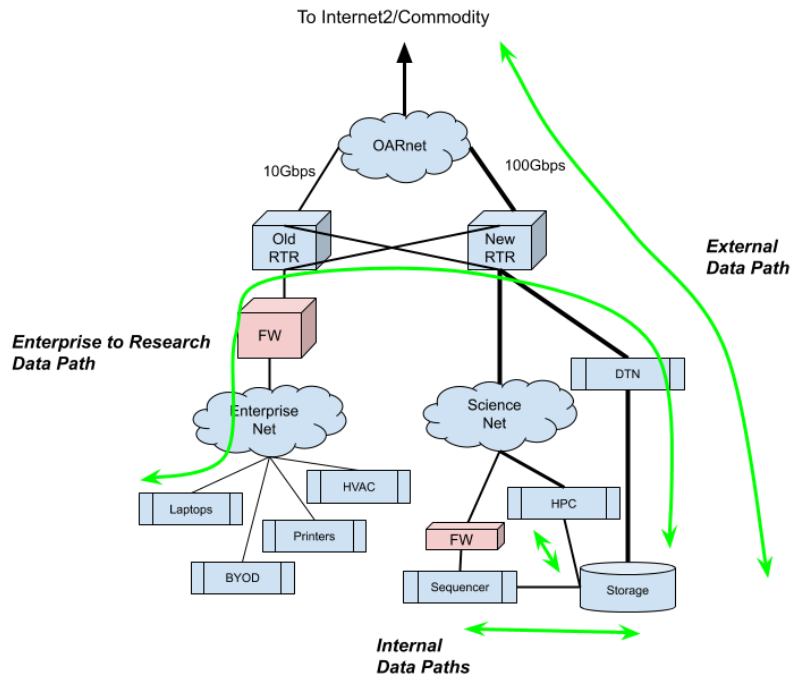


Figure 7 - Campus Network Diagram (Data Paths)

Following this general discussion, a more specific model has emerged, as shown in Figure 8.

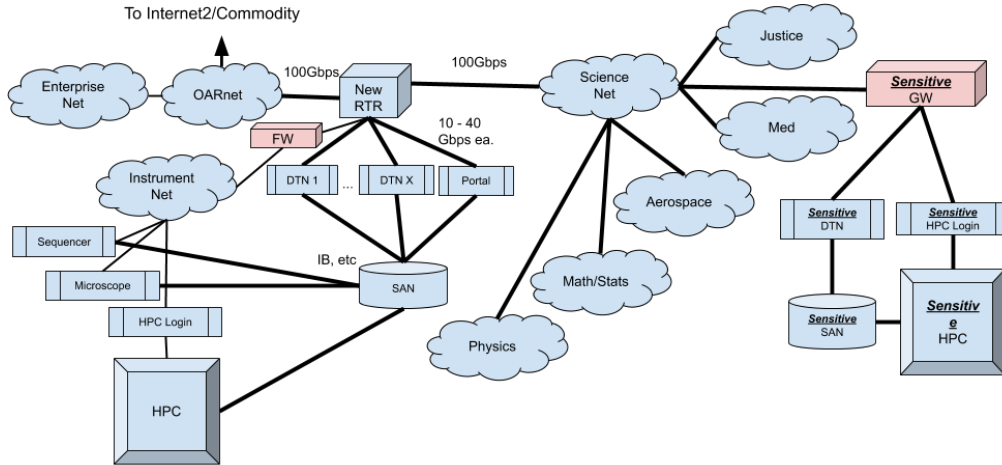


Figure 8 - Science DMZ Proposal (Logical)

There are several design highlights to consider in this expanded example:

- Instead of installing a single “large” DTN (e.g. 100Gbps), the installation of smaller, and more numerous, DTNs is more scalable. This allows different groups access to specific resources as demand grows. Portals can be added as well.
- All DTNs can be integrated into storage, and can also connect to HPC resources and instruments.
- The instrumentation control network can be protected by security that is more invasive, as long as the data path is not inhibited.
- Individual department deployments of DMZ resources can be installed and maintained, and still afford access to HPC, instruments, and storage. This access will only require an additional “hop” or two through fast infrastructure.
- An emerging “sensitive” use case can be constructed using separate infrastructure. Establishing FISMA/HIPAA/ITAR/EAR compliance requires special considerations. As such, this is typically different infrastructure and does not mingle with other research or enterprise use cases. Given the requirements of campus researchers, as seen in the use cases related to Medical and Criminal Justice, this investment would be quickly utilized and greatly beneficial.
- The network paths can be seen in Figure 9.




Figure 9 - Science DMZ Proposal (Paths)

Modifications to the existing network designs are subject to funding via proposals to the NSF, or University of Cincinnati capital investments.

6. Action Items

EPOC and OARnet recorded a set of action items from the University of Cincinnati Campus-Wide Deep Dive, including continuing the ongoing support and collaboration. These are a reflection of the Case Study reports, in person discussion, and other items specifically related to the support of scientific users.

1. **University of Cincinnati**, with the assistance of **OARnet** and **EPOC**, will work toward a new research network design pattern, and will attempt to provide a friction free network path to local and remote storage and compute resources.
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 identify and high-performance use cases. (e.g. Globus).
3. **University of Cincinnati** will deploy of 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.
4. **University of Cincinnati** will split off the functionality of operating the campus research network from that of the enterprise network. Having dedicated staff for the purpose of engaging with researchers on how to use network infrastructure.
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 working 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 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.

Appendix A – University of Cincinnati Cyberinfrastructure Plan

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 Cincinnati¹⁹.

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

- Modify and Enhance UCScienceNet (UCSN), a 100Gbps 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.
- 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.
- 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.
- Expand IT Resources & Services and Education for Researchers, collaborating with computational researchers to become an integral part of the research data management lifecycle in support of the University of Cincinnati 'Next Lives Here' strategy and 'Digital Futures' initiatives.

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.

Background

Supported by the university-wide IT governance structure, the IT@UC Strategic Plan establishes the core values of partnership, collaboration, and communication to drive the transformation, and focuses on three high-level strategies and priority investments for IT across the university:

¹⁹ As submitted for NSF Award #1541410, see also https://www.nsf.gov/awardsearch/showAward?AWD_ID=1541410

- Transformative eLearning - Drive adoption of 21st Century learning through the discovery, development, and deployment of a standardized eLearning ecosystem.
- Research and Knowledge Creation – Identify and enable a base of resources to support research across all disciplines, including high performance computing and data repositories, collaboration tools and education.
- Shared Services and Shared Architecture - Build and develop high-performance shared services and intentional interdependence between IT units to improve efficiency and cost-effectiveness.

The focus of the IT@UC Strategic Plan not only embraces but empowers the university's 'Next Lives Here' strategic direction, which makes a commitment to building the university's resource base and leveraging research.

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 the 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)

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.

This expansion will require the deployment of Cisco Nexus 3000 switches deployed in strategic research areas. The Nexus 3000 will provide scalable 40Gb back to the research core. It will also provide 10Gb to the high-performance computing equipment, which today is limited to 10/100Mb.

Expanding the friction-free campus networking architecture, and eliminating campus and building level network infrastructure constraints, will enable formation of a Research Ecosystem encouraging diverse, multidisciplinary collaboration and partnerships to address complex Grand Challenge problems.

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. perfSONAR is being used 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. The SOCC data center, provides real-time synchronization with data storage systems in the primary data center, replication of data backups, and both active-active and active-standby hardware for critical business continuity and disaster recovery scenarios.

Storage Capacity

Research Project storage is provided on an XAScaler 7990 EDR appliance with 1.488PB Raw storage. The parallel file system is integrated into the UC Central HPC switch fabric for maximum, non-blocking I/O performance and security. Performance is 15 GB/s sequential write, 20 GB/s sequential read. The system is scalable to multiple petabytes of capacity and can easily be expanded to meet increasing storage demands.

Compute Capacity

IT@UC's Research and Development department leads the IT Governance's Research and Development topical committee which is charged with identifying, prioritizing and recommending computational resources for researchers. Membership includes computational researchers, undergraduate and graduate students, IT@UC cyberinfrastructure support personnel, distributed research IT support personnel. Dr. Phil Taylor, Asst Vice President, Research Infrastructure and Development and Annette Ready, Assoc Vice President, IT@UC, Innovation and Partnerships, are co-chairs of the committee.

In 2018, a partnership between the Office of Research, IT@UC and university colleges/departments/researchers developed the Advanced Research Computing (ARC) initiative which is piloting a central HPC cluster available to all university researchers. The initial cluster was funded by the partnership and an NSF MRI proposal (with 30% cost share) was submitted in January 2019 to further expand the cluster. A Faculty Advisory Committee is charged with developing the sustainability plan and business model. This will utilize the metrics of the 18-month pilot to communicate the broad need for these resources at the University of Cincinnati

UC's Advanced Research Computing (ARC)

The Advanced Research Computing (ARC) initiative HPC Cluster configuration is equipped with 50 teraFLOPS of peak CPU performance and 2 NVIDIA Tesla V100 GPU nodes (224 teraFLOPS deep learning peak performance) connected with high-performance 100 Gbps Omnipath (OPA) interconnect.

CPU Nodes

16 - Intel Xeon Gold 6148 2.4G, 20C/40T, 10.4GT/s, 27M Cache, Turbo, HT (150W)
DDR4 2666 RAM 192GB per node, DDR4 2666

GPU Node

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

Omnipath HPC Networking infrastructure - Maximum Ominpath bandwidth
between nodes = 100Gbps

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. The 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 the data center. This IPv6 build out will be in coordination with the Business Application Services team.

Sustainability

IT@UC currently engages Cisco Smartnet for maintenance on all of the core, distribution, and optical equipment. Aruba maintenance is used for all 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 the 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 the 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.

Resources & Services

The IT@UC Research & Development office connects researchers with technical expertise, resources, training, and state-of-the-art IT services to support individual and multidisciplinary IT-enabled research projects. As an example, the Center for Simulations & Virtual Environments Research (UCSIM) is providing technical and hardware expertise, programming, and modeling support for virtual and augmented reality research collaborations with the Cincinnati Children's Hospital Medical Center TEAM VR Lab, the Air Force Research Lab Discovery Center, and the UC

Center for Cognition, Action, and Perception. Moving forward, the department will continue to grow its resource base and extend services to include expertise in data analytics, high performance computing, and data visualization. The department will continue to collaborate with local, regional and national experts in order to provide the level of research computing services necessary at a large, research intensive university.

Integration with State, National, and 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.