

Lawrence Berkeley National Laboratory

LBL Publications

Title

University of Wisconsin-Madison Campus-Wide Deep Dive

Permalink

<https://escholarship.org/uc/item/99c4v5xh>

Authors

Zurawski, Jason
Schopf, Jennifer
Addleman, Hans

Publication Date

2020-05-26

Peer reviewed



University of Wisconsin-Madison Campus-Wide Deep Dive

June 17-19, 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 Wisconsin-Madison Campus Deep Dive

Final Report

*University of Wisconsin-Madison
Madison, WI
June 17-19, 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-2001325¹.

¹<https://escholarship.org/uc/item/99c4v5xh>

Participants and Contributors

Hans Addleman, Indiana University
Paul Barford, Department of Computer Sciences
Steve Barnet, IceCube Neutrino Observatory
Michael Blodgett, ESnet
Brian Bockelman, Center for High Throughput Computing
Dan Bradley, Department of Physics
Dale Carder, ESnet
Jan Cheetham, Division of Information Technology
Patrick Christian, Division of Information Technology
Derek Cooper, Center for High Throughput Computing
Rob Kohlhep, College of Engineering
John Lalande, Space Science and Engineering Center
Miron Livny, Center for High Throughput Computing
Lauren Michael, Center for High Throughput Computing
Ajit Mohapatra, Department of Physics
Eric Montemayor, Department of Biochemistry
Dane Morgan, Computational Materials Group
Dirk Norman, Great Lakes Bioenergy Research Center, Wisconsin Energy Institute
Jerry Robaidek, Space Science and Engineering Center
Jeremy Sarauer, Division of Information Technology
Jeanne Skul, Division of Information Technology
Edgar Spalding, Department of Botany
Carl Vuosalo, Department of Physics
Elizabeth Wright, Department of Biochemistry
Jason Zurawski, ESnet

Report Editors

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

Contents

Disclaimer	2
Participants and Contributors	4
Report Editors	4
Contents	5
1 Executive Summary	10
2 Process Overview and Summary	12
2.1 Campus-Wide Deep Dive Background	12
2.2 Campus-Wide Deep Dive Structure	13
2.3 University of Wisconsin–Madison Campus-Wide Deep Dive Background	15
2.4 Organizations Involved	16
3 University of Wisconsin-Madison Case Studies	17
3.1 The University of Wisconsin-Madison Campus Case Study	18
3.1.1 Infrastructure Background	18
3.1.2 General Network Architecture	18
3.1.3 Network Core	18
3.1.4 Distribution	19
3.1.5 Access	19
3.1.6 Overlay Networks	20
3.1.7 Special Purpose Network Connections	20
3.1.8 Wide-area Network Connectivity	21
3.1.9 Research Network Support Tools	24
3.1.10 Planned Network and Research Support Enhancements	25
3.2 The Center for High Throughput Computing (CHTC) Case Study	26
3.2.1 Science Background	26
3.2.2 Collaborators	26
3.2.3 Instruments and Facilities	27
3.2.4 Process of Science	29
3.2.5 Remote Science Activities	29
3.2.6 Software Infrastructure	29
3.2.6.1 Data Archetypes	29
3.2.7 Network and Data Architecture	30
3.2.8 Cloud Services	30
3.2.9 Known Resource Constraints	30

3.2.10 Parent and Affiliated Organizational Cooperation	31
3.2.11 Outstanding Issues	31
3.3 Ahlquist Lab Virology Research	33
3.3.1 Science Background	33
3.3.2 Collaborators	33
3.3.3 Local and Remote Instruments and Facilities	34
3.3.4 Process of Science	34
3.3.5 Software Infrastructure	35
3.3.6 Network and Data Architecture	35
3.3.7 Cloud Services	35
3.3.8 Known Resource Constraints	35
3.4 Huisken Labs: Microscopy Case Study	37
3.4.1 Science Background	37
3.4.2 Collaborators	37
3.4.3 Local and Remote Instruments and Facilities	37
3.4.4 Process of Science	37
3.4.5 Software Infrastructure	38
3.4.6 Network and Data Architecture	38
3.4.7 Cloud Services	38
3.4.8 Known Resource Constraints	38
3.5 Tier-2 Computing Center for the Compact Muon Solenoid (CMS) Experiment at the Large Hadron Collider (LHC) Case Study	40
3.5.1 Science Background	40
3.5.2 Collaborators	40
3.5.3 Instruments and Facilities	41
3.5.4 Process of Science	41
3.5.5 Remote Science Activities	41
3.5.6 Software Infrastructure	42
3.5.7 Network and Data Architecture	42
3.5.8 Cloud Services	43
3.5.9 Known Resource Constraints	43
3.5.10 Parent and Affiliated Organizational Cooperation	44
3.5.11 Outstanding Issues	44
3.6 The Great Lakes Bioenergy Research Center (GLBRC) Case Study	45
3.6.1 Science Background	45
3.6.2 Collaborators	45
3.6.3 Local and Remote Instruments and Facilities	46

3.6.4 Process of Science	46
3.6.5 Software Infrastructure	47
3.6.6 Network and Data Architecture	47
3.6.7 Cloud Services	47
3.6.8 Known Resource Constraints	47
3.6.9 Parent and Affiliated Organizational Cooperation	48
3.6.10 Outstanding Issues	48
3.7 IceCube Neutrino Observatory Case Study	49
3.7.1 Science Background	49
3.7.2 Collaborators	49
3.6.3 Instruments and Facilities	50
3.7.4 Process of Science	50
3.7.5 Remote Science Activities	52
3.7.6 Software Infrastructure	52
3.7.7 Network and Data Architecture	52
3.7.8 Cloud Services	53
3.7.9 Known Resource Constraints	53
3.7.10 Outstanding Issues	53
3.8 UW-Madison Cryo-Electron Microscopy Research Center (CEMRC) Case Study	54
3.8.1 Science Background	54
3.8.2 Collaborators	55
3.8.3 Instruments and Facilities	55
3.8.4 Process of Science	55
3.8.5 Remote Science Activities	56
3.8.6 Software Infrastructure	56
3.8.7 Network and Data Architecture	57
3.8.8 Cloud Services	57
3.8.9 Known Resource Constraints	57
3.8.10 Outstanding Issues	57
3.9 Space Science and Engineering Center (SSEC) Case Study	58
3.9.1 Science Background	58
3.9.1.1 SDS	58
3.9.1.2 SIPS	59
3.9.1.3 S4	59
3.9.2 Collaborators	59
3.9.3 Instruments and Facilities	60

3.9.3.1 SDS	60
3.9.3.2 SIPS	60
3.9.3.3 S4	60
3.9.4 Process of Science	61
3.9.4.1 SDS	61
3.9.4.2 SIPS	61
3.9.4.3 S4	61
3.9.5 Remote Science Activities	61
3.9.6 Software Infrastructure	62
3.9.7 Network and Data Architecture	62
3.9.8 Cloud Services	63
3.9.9 Outstanding Issues	63
3.10 Plant Physiology and Computation-based Phenotyping Case Study	64
3.10.1 Science Background	64
3.10.2 Collaborators	64
3.10.3 Instruments and Facilities	64
3.10.4 Process of Science	64
3.10.5 Remote Science Activities	65
3.10.6 Software Infrastructure	65
3.10.7 Network and Data Architecture	65
3.10.8 Cloud Services	65
3.10.9 Known Resource Constraints	65
3.10.10 Parent and Affiliated Organizational Cooperation	65
3.10.11 Outstanding Issues	66
3.11 Computational Materials Case Study	67
3.11.1 Science Background	67
3.11.2 Collaborators	67
3.11.3 Instruments and Facilities	67
3.11.4 Process of Science	67
3.11.5 Remote Science Activities	68
3.11.6 Software Infrastructure	68
3.11.7 Network and Data Architecture	68
3.11.8 Cloud Services	68
3.11.9 Known Resource Constraints and Outstanding Issues	68
4 Discussion Summary	69
4.1 The Center for High Throughput Computing (CHTC), Ahlquist Lab Virology Research, and Huisken Lab Microscopy Research	70

4.2 Tier-2 Computing Center for the Compact Muon Solenoid (CMS) Experiment at the Large Hadron Collider (LHC)	71
4.3 The Great Lakes Bioenergy Research Center (GLBRC)	72
4.4 IceCube Neutrino Observatory	73
4.5 UW-Madison Cryo-Electron Microscopy Research Center (CEMRC)	74
4.6 Space Science and Engineering Center (SSEC)	75
5 Action Items	76
Appendix A - UW-Madison Campus Supernode Locations	77
Appendix B - UW-Madison Campus Backbone Network	78
Appendix C - UW-Madison Building Access Network (Logical)	79
Appendix D - Cyberinfrastructure Plan, University of Wisconsin-Madison, April 2018	81
OVERVIEW	81
GOVERNANCE	81
RESEARCH COMPUTING	81
NETWORK	82
Wireless	82
Network standards and capabilities	83
Software Defined Networking (SDN)	83
SECURITY	83
MIDDLEWARE	84
SOFTWARE	85
DATA	85
Storage	85
Data Centers	85
CYBERINFRASTRUCTURE STAFFING	85
EDUCATION AND TRAINING	86
Appendix E - UW-System Regional Networking Diagram	87

1 Executive Summary

In June 2019, staff members from the Engagement and Performance Operations Center (EPOC) and the University of Wisconsin–Madison Division of Information Technology (UW DoIT) met with researchers for the purpose of a Campus-Wide Deep Dive. The goal of this meeting was to help characterize the requirements for a number of research and educational activities, and to enable cyberinfrastructure support staff to better understand the needs of researchers. Material for this event includes both the written documentation from nine projects, and the campus technology organization, but also a writeup of the discussion that took place in person. The Case Studies highlighted the ongoing challenges that the University of Wisconsin–Madison will face in the coming years to support, encourage, and grow several emerging and established use cases.

The University of Wisconsin–Madison received an NSF award to help support upgrading the campus network in 2012, specifically to include a Science DMZ and monitoring equipment. However, there is a currently identified need to refresh equipment and identify and work more closely with current campus researchers. Updates to the state network, as well as the campus, are being planned and were discussed during this event.

As part of the overall review, the necessity of working more closely with research teams as they are still in the planning phase was discussed so that DoIT can better understand and adapt to changes in requirements as the research demands grow over time. Additional challenges with securing sensitive data, cybersecurity, and supporting collaborations were also discussed.

Action items from the meeting included:

- 1) UW DoIT will work to make the Science DMZ infrastructure more accessible for a greater number of use cases on and off campus.
- 2) UW DoIT will explore the use of sFlow/NetFlow data in more campus locations to understand traffic patterns.
- 3) UW DoIT will explore deployment of campus-wide data transfer nodes (DTNs) with Globus endpoints.
- 4) UW DoIT will continue to work with the Center for High Throughput Computing (CHTC) on ways to reduce friction for campus data movement.
- 5) UW DoIT will explore institutional storage options for campus users.
- 6) CHTC will continue to expand access to campus computing use cases such as those presented in the IceCube and Cryo-EM use cases.
- 7) UW DoIT will work with the Ahlquist and Huisken Labs to address data movement challenges to remote sites.
- 8) UW DoIT to work with the high energy physics CMS on network upgrades to support future large hadron collider (LHC) upgrades.

- 9) UW DoIT, the Great Lakes Bioenergy Research Center (GLBRC), and EPOC will work to address a wide area data movement problem to Michigan State University.
- 10) UW DoIT and GLBRC will continue to work on connectivity to remote field research locations.
- 11) UW DoIT will work with IceCube and National Energy Research Scientific Computing Center (NERSC) to address performance abnormalities with data transfers.
- 12) UW DoIT and Cryo-Electron Microscopy Research Center (CEMRC) will address architectural issues with network, compute, and storage as new instruments are installed and commissioned.

2 Process Overview and Summary

2.1 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 EPOC 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, the Deep Dive process offers 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 facilities². The EPOC team adapted this approach to work with individual science groups through a set of structured data-centric conversations and questionnaires.

² <https://fasterdata.es.net/science-dmz/science-and-network-requirements-review>

2.2 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 manage data resources (locally or remotely), 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.

2.3 University of Wisconsin–Madison Campus-Wide Deep Dive Background

In June 2019, EPOC and the University of Wisconsin–Madison organized a Campus-Wide Deep Dive to characterize the requirements for several departments and research projects on the campus. The representatives were asked to communicate and document their requirements in a case-study format (see [Section 3 University of Wisconsin-Madison Case Studies](#)). These include:

- [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#)
- [Section 3.2 The Center for High Throughput Computing \(CHTC\) Case Study](#)
- [Section 3.3 Ahlquist and Huisken Labs: Virology Research and Microscopy Case Study](#)
- [Section 3.4 Tier-2 Computing Center for the Compact Muon Solenoid \(CMS\) Experiment at the Large Hadron Collider \(LHC\) Case Study](#)
- [Section 3.5 The Great Lakes Bioenergy Research Center \(GLBRC\) Case Study](#)
- [Section 3.6 IceCube Neutrino Observatory Case Study](#)
- [Section 3.7 UW-Madison Cryo-Electron Microscopy Research Center \(CEMRC\) Case Study](#)
- [Section 3.8 Space Science and Engineering Center \(SSEC\) Case Study](#)
- [Section 3.9 Plant Physiology and Computation-based Phenotyping Case Study](#)
- [Section 3.10 Computational Materials Case Study](#)

The face-to-face meeting took place at the University of Wisconsin–Madison, in Madison WI, on June 17-19, 2019 (see discussion in [Section 4 Discussion Summary](#)). We document next steps in [Section 5 Action Items](#).

2.4 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 University of Wisconsin–Madison (UW–Madison) is a public research university in Madison, Wisconsin. Founded when Wisconsin achieved statehood in 1848, UW–Madison is the official state university of Wisconsin and the flagship campus of the University of Wisconsin System. UW–Madison is organized into 20 schools and colleges, which enrolled 30,361 undergraduate and 14,052 graduate students in 2018. Its comprehensive academic program offers 136 undergraduate majors, along with 148 master's degree programs and 120 doctoral programs.

3 University of Wisconsin-Madison Case Studies

There are nine scientific use cases, and one campus technology overview, provided by the University of Wisconsin-Madison. These are as follows:

- [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#)
- [Section 3.2 The Center for High Throughput Computing \(CHTC\) Case Study](#)
- [Section 3.3 Ahlquist and Huisken Labs: Virology Research and Microscopy Case Study](#)
- [Section 3.4 Tier-2 Computing Center for the Compact Muon Solenoid \(CMS\) Experiment at the Large Hadron Collider \(LHC\) Case Study](#)
- [Section 3.5 The Great Lakes Bioenergy Research Center \(GLBRC\) Case Study](#)
- [Section 3.6 IceCube Neutrino Observatory Case Study](#)
- [Section 3.7 UW-Madison Cryo-Electron Microscopy Research Center \(CEMRC\) Case Study](#)
- [Section 3.8 Space Science and Engineering Center \(SSEC\) Case Study](#)
- [Section 3.9 Plant Physiology and Computation-based Phenotyping Case Study](#)
- [Section 3.10 Computational Materials Case Study](#)

Each of these Case Studies provides a glance at research activities for the University, the use of experimental methods and devices, the reliance on technology, and the scope of collaborations. Estimates on data volumes, technology needs, and external drivers are discussed where relevant.

The University of Wisconsin-Madison is committed to supporting these use cases through technology advancements and is actively pursuing grant solicitations. The landscape of support will change rapidly in the coming years, and these use cases will take full advantage of campus improvements as they become available.

3.1 The University of Wisconsin-Madison Campus Case Study

Content in this section authored by Jan Cheetham, Patrick Christian, Jeremy Sarauer, and Jeanne Skul from the Division of Information Technology.

3.1.1 Infrastructure Background

The University of Wisconsin (UW) campus network spans over 180 buildings with wired and ubiquitous interior wireless network access, firewalls, and other telecommunication services such as video conferencing and voice-over-IP (VOIP). Switching and routing are implemented using Cisco switches and routers while Palo Alto firewalls and Aruba controller-based infrastructure is used for wi-fi (802.11n/ac) wireless networking. All networks are continuously monitored by a 24x7 Network Operations Center (NOC) support team and professionally managed by UW's central information technology group, DoIT Network Services, in collaboration with on-campus schools and colleges.

3.1.2 General Network Architecture

The network is architected in a traditional multi-layer, 3-tier hierarchical design of core, distribution, and access layers that is logically depicted in Figure 1.

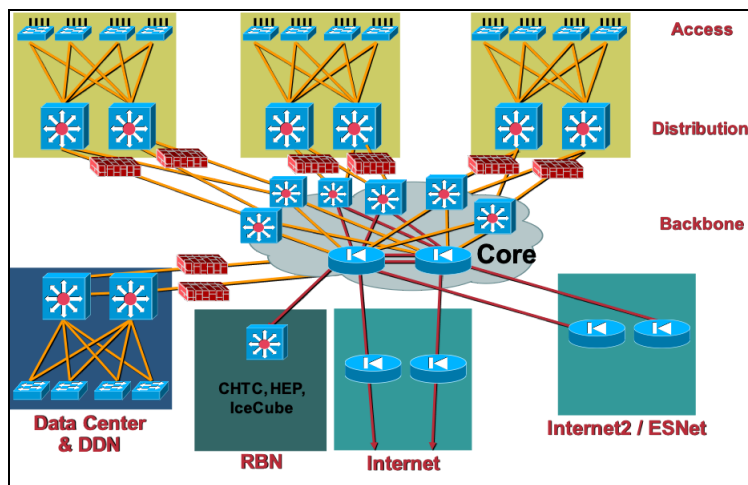


Figure 1: A logical diagram of the University of Wisconsin-Madison multi-layer 3-tier campus network. RBN is the Research Backbone Network, sometimes referred to as the Science DMZ, and the data center and distributed data network (DDN) is the network that interconnects the primary and backup data centers to the campus core network.

3.1.3 Network Core

The core of the network consists of 3 “supernode” locations (Computer Sciences, Animal Sciences, 432 North Murray (aka 432 East Campus Mall)) which geographically aggregate approximately 1/3 or about 60 buildings in each section, as shown in Appendix A.

Appendix B shows the physical implementation of the UW-Madison campus network - including connection speeds for the network. Each supernode area consists of a pair of (geographically separated) Cisco Nexus 7010 switches performing switching and routing functions which aggregate distribution node locations and connect Cisco ASR 9022 devices performing campus core and border routing functions. Also attached to the Nexus 7010 hardware are redundant (active/passive) virtual Palo Alto firewall (vsys) instances to enable independent, separately-managed firewalls with traffic separation from other groups.

3.1.4 Distribution

The distribution layer of the network aggregates multiple building entrance room (ER) or main distribution frame (MDF) nodes together in the campus network. Each distribution node aggregate is connected with a minimum of two active 10 Gigabits per second (Gbps) links to two geographically diverse Nexus 7010s (7k1 and 7k2 as depicted in Figure 2) in the core network, for a total of 20Gbps of network capacity from distribution to the campus core. The Nexus devices are configured in a virtual port channel (vPC) configuration creating a single logical node from two diverse nodes. Additional capacity between the distribution and campus core layers is added as network utilization requires. Packets (P) flow from distribution node A in Figure 2 to 7k1 or 7k2.

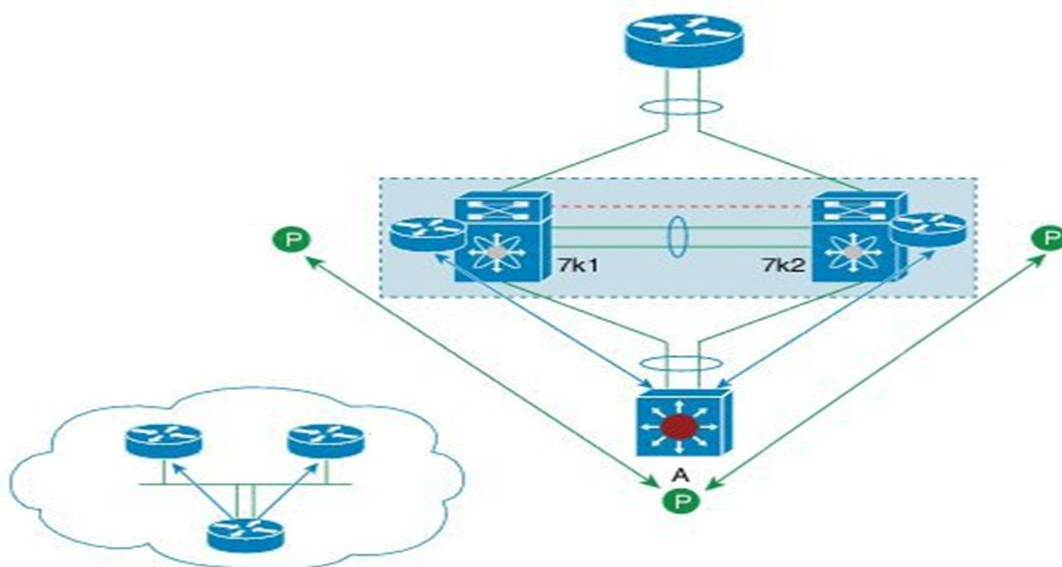


Figure 2: The implementation of the UW virtual port channel enables packets (P) to flow to a single logical node created from two diverse nodes.

3.1.5 Access

The access layer of the UW-Madison campus network, shown in Appendix C, is where most network connectivity is delivered to individually-connected terrestrial network users and wi-fi access points. One or more switch(es) in a switch stack

configuration located within a building telecommunication room(s) (TRs) or intermediate distribution frame(s) (IDFs) provide 10/100/1000 Megabits per second (Mbps) Ethernet connections to desktops and other wired devices. A small number of servers may also be connected at 1 or 10Gbps in a building's ER/MDF. In addition, this layer aggregates multiple in-building TRs/IDFs together via a separate switch stack and connects them to the distribution layer at minimum of 20Gbps (2x10Gbps). Additional 2x10Gbps of capacity from the ER/MDF uplink to the distribution layer is upgraded as network utilization requires – particularly for larger labs.

3.1.6 Overlay Networks

UW has several overlay networks built on top of the campus core network. These networks include:

- Research Backbone Network (RBN) (aka Science DMZ) – enables large, friction-free data transfers per ESnet Science DMZ³ concepts.
- Building Automation Network (BAN) – connects digital controls within buildings to a controller for management and monitoring.
- Payment Card Industry (PCI) - Payment Card Industry Data Security Standard (PCI DSS) compliant network.
- Data Center and Distributed Data Network (DDN) – network to interconnect primary/backup data center to campus core network as well as interconnect small data centers at select campus locations.

3.1.7 Special Purpose Network Connections

UW has several significant research projects that require special network connectivity due to unique network requirements. These projects include:

- **High Energy Physics (HEP-LHC) and IceCube**
 - Leveraging the identity management systems, the campus network, and data center and security monitoring operations, HEP and Ice Cube share a dedicated 100Gbps connection from their shared lab to the UW-Madison (Computer Science) Campus supernode⁴.
 - IceCube has moved its storage cluster from its current off-campus offices at Network222 to a new UW off-campus leased data center facility called OneNeck in Fitchburg, WI due to cooling and power constraints at Network222.. Computation resources are split between the UW-Madison campus and nodes at Network222. IceCube equipment at the OneNeck facility is connected by a fiber ring to the campus network at 100Gbps while Network222 is connected to the campus network via a fiber-ring with 8x10Gbps of capacity to its on-campus lab. The Ice Cube and HEP labs share a 100Gbps uplink to the campus core and to the world.
- **Center for High Throughput Computing (CHTC)**

³<https://fasterdata.es.net/science-dmz/>

⁴ <https://stats.net.wisc.edu/chmrl/rootindex.html>

- The CHTC manages a major, centralized HTCondor cluster and access to Open Science Grid (OSG) scientific computing resources on behalf of the UW campus. The main CHTC cluster is connected to the campus core (Nexus 7010 Computer Science area) through a 4x10Gbps connection to the campus core. A traditional multi-stage Clos architecture with 10Gbps server connections and 40Gbps leaf to spine connections links servers in a cluster together.
- **CloudLab**
 - CloudLab is an NSF funded project that supports a large-scale distributed infrastructure based at the University of Utah, Clemson University and the University of Wisconsin, on top of which researchers are able to construct many different types of clouds. At UW, 4x10Gbps connectivity to the Computer Science (CSSC) campus core Nexus 7010 infrastructure split between 2 geographically diverse locations on campus. Using vPC technology, the 4x10Gbps connections function as one large pipe to the campus core network.
- **Federal Information Systems Act (FISMA)**
 - UW operates a small FISMA-compliant data center with a completely isolated network environment to ensure FISMA data security standards compliance. The data center is connected to the campus core router for Internet access only via a 10Gbps link.

3.1.8 Wide-area Network Connectivity

In collaboration with the University of Minnesota, University of Iowa, and Iowa State University, UW built and operates the Broadband Optical Research Education and Science network⁵ (BOREAS-Net), shown in Figure 3. BOREAS-Net is a Regional Optical Network (RON) utilizing commercial telecommunications carrier grade optical infrastructure to light university-owned or controlled dark fiber strands to connect national exchange points in Chicago, Kansas City, and increasingly Minneapolis.

⁵ <https://www.boreas.net>

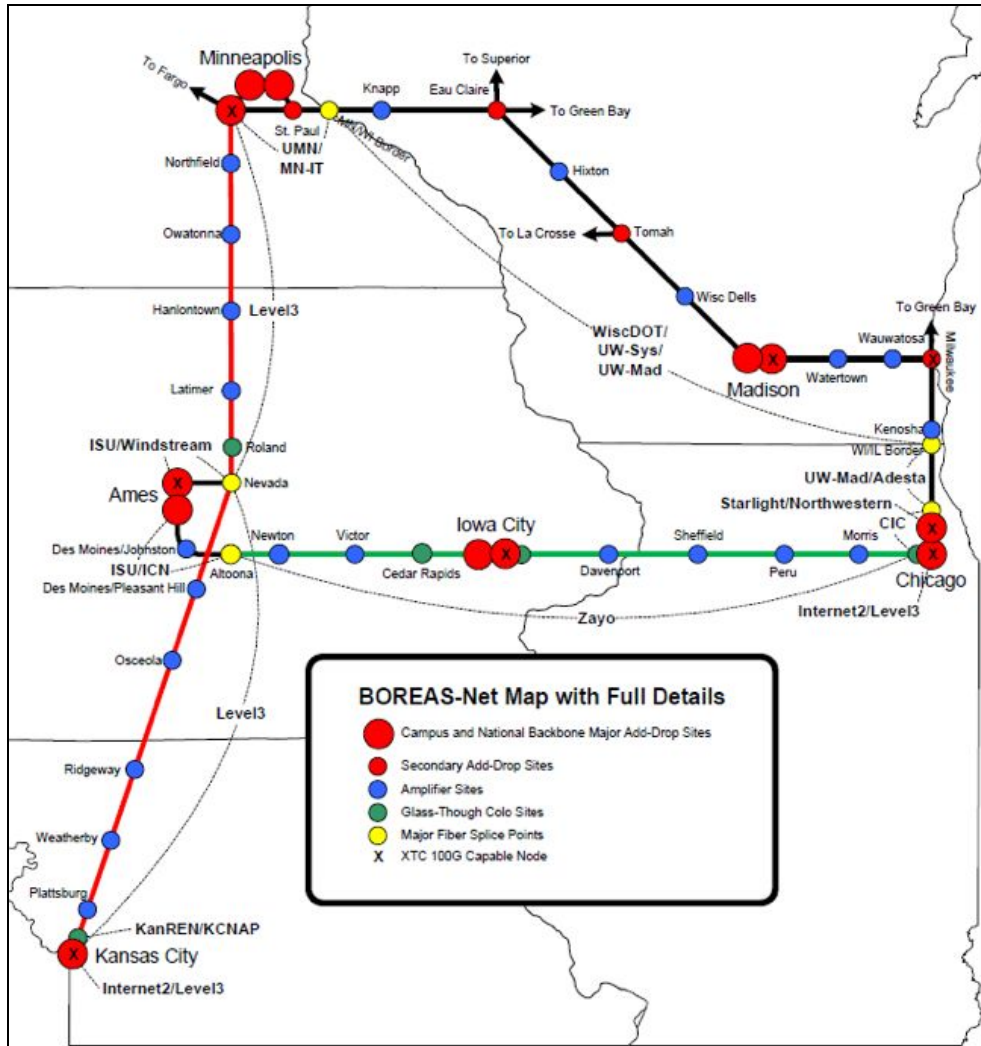


Figure 3 - The BOREAS-Net RON links UW, University of Minnesota, University of Iowa, and Iowa State University to exchange points in Chicago, Kansas City, and Minneapolis.

UW research network connections to scientific networks, such as Internet2 and ESnet, utilize a dedicated, optically-protected 100Gbps connection to the Big Ten Academic Alliance (BTAA) OmniPoP⁶ regional aggregation equipment in Chicago. Research connection traffic flows are bursty in nature, as shown in Figure 4, with occasional peaks to 100% though are currently about 55-65% utilization at the 95% of use.

⁶ <https://www.btaa.org/technology/omnipop/introduction>

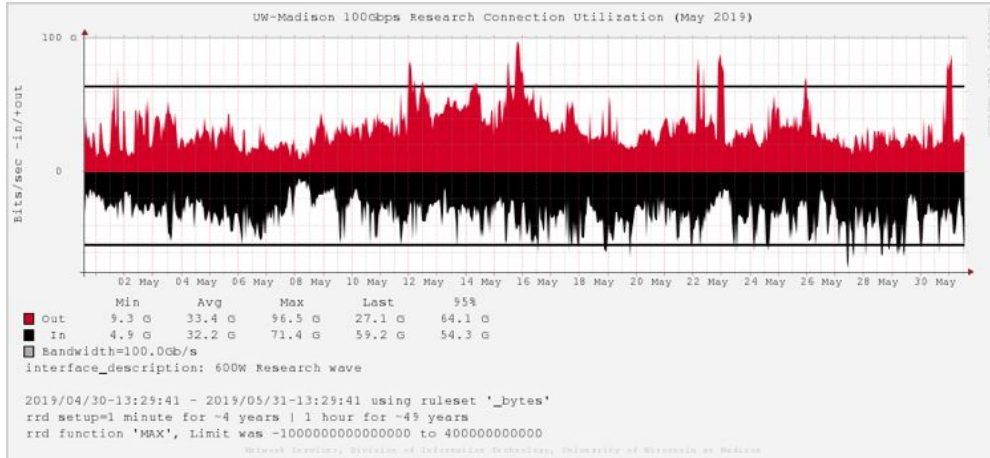


Figure 4 - UW 100Gbps research circuit utilization for May 2019.

BTAA OmniPoP provides UW with shared high-speed gateways to various research and education networks, including BTAA-shared 100Gbps links to both Internet2 (using Advanced Layer 2 (AL2S) and separately Advanced Layer 3 Service (AL3S)), a BTAA-shared 2x100Gbps connection to U.S unclassified federal labs (including Argonne National Laboratory (ANL) and Fermi National Accelerator Lab (FNAL)) via the Department of Energy’s Energy Science Network (ESnet), 100Gbps BTAA-shared connectivity to the LHC Open Network Environment (LHCONE), shared 10Gbps link to the StarLight exchange in Chicago, and other regional, national, and global research networks. UW is also connected to the Great Plains Network (GPN) in Kansas City at 2x10Gbps and uses this connection as a (small) tertiary connection to Internet2.

UW participates in the Northern Tier Network Consortium (NTNC) and operates portions of the network in North and South Dakota, as shown in Figure 5. BOREAS-Net is able to connect to Seattle via this network when necessary. Currently, NOAA’s N-Wave 10Gbps national backbone utilizes the Seattle-Chicago path.

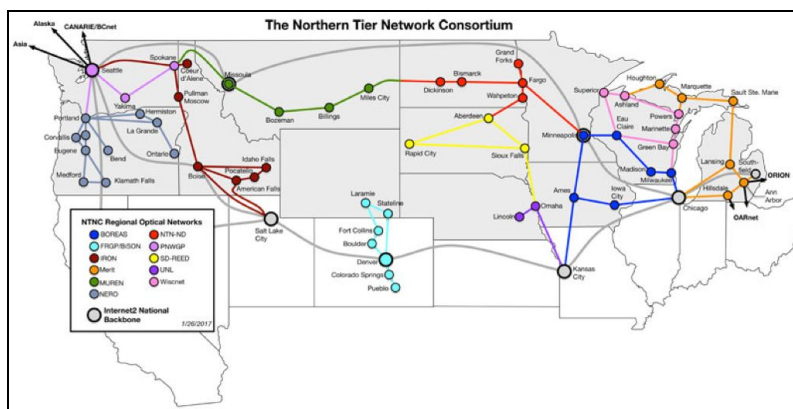


Figure 5: UW participates in the Northern Tier Network Consortium (NTNC) and operates portions of the network in North and South Dakota.

Off-campus groups, including the University Research Park in the greater Madison metropolitan area, are connected with 10Gbps service in a redundant ring configuration via dark fiber or wavelength services to the UW campus via the Metropolitan Unified Fiber Network (MUFN), shown in Appendix D. UW built and operates this network in collaboration with area education, government, non-profit, health and commercial partners.

Research flows occasionally utilize the “commodity” or general Internet. UW is locally, regionally, and nationally connected to the Internet by over 240Gbps of Internet access via the UW-System Network (see Appendix F). Commodity access is provided by two paid commercial providers, a 100Gbps circuit to Chicago using Telia-Carrier and two separate 10Gbps connections in Madison and Milwaukee to AT&T. Internet access is also provided by routed private peering networks, including Yahoo, Netflix, Facebook, Microsoft, Apple, Twitter, the Amsterdam Internet Exchange, Limelight, Hurricane Electric, Verizon Media and others. The UW campus is also connected to over 30Gbps of cached content delivery network services provided by Google, Netflix and Akamai.

There are various private 100 Gbps interconnections with Big10 universities and other national, regional, and local networks as well as Cloud providers, including Google, Amazon, and Azure. These include the recently provisioned 500Mbps Express Route connection via Internet2 to UW’s Azure virtual private cloud (VPC) and soon a 500Mbps Direct Connect path to UW Amazon Web Services (AWS) VPC instance. As necessary, connections may also be made to several telecommunications providers that are co-located in the data center, including ATT, CenturyLink, Verizon, TDS, MUFN, WiscNe,t and Spectrum Networks (formerly Charter and Time Warner Cable).

3.1.9 Research Network Support Tools

UW has perfSONAR nodes operating at several campus locations including:

- Campus commodity Internet border (10Gbps)
- Research backbone network (aka Science DMZ) (10Gbps)
- High-energy physics (HEP) gateway (1Gbps)
- WI Institutes of Discovery Center for High Throughput Computing node (CHTC cluster) (10GE)
- IceCube (off-campus location at 222 W. Washington) (10Gbps)
- UW-Madison Computer Science Department (1Gbps)

UW does not currently support any data transfer nodes (DTNs) or a Flash I/O Network Appliance (FIONA) node. There is also no current subscription for the Globus data transfer service, although several faculty have downloaded and used the Globus basic client to transfer data for their self-managed endpoints. Several

faculty have expressed interest in such a tool, but are not willing to pay the annual subscription fees themselves.

3.1.10 Planned Network and Research Support Enhancements

UW intends to provide the following network enhancements and capabilities in the following timeframes:

Present-2 years (current budget horizon)

- Upgrade Nexus 7010 core to Cisco Nexus 7700 or Cisco 9600 switches
- Add 2nd 100Gbps circuit to Chicago dedicated for UW research (ESnet, Internet2) (July 2020)
- Connect to Azure, AWS, and GCP cloud providers with layer 2 (Express Route, Direct Connect) links in addition to strong layer 3 routed capabilities via Internet2 (Azure done; AWS start in July'19; GCP expected 2020)
- Explore custom direct connect opportunities with one or two major cloud providers to improve current research data transfer rates to/from specific cloud providers

Next 2-5 years (current technology horizon)

- Investigate, design, and deploy next generation campus network, including enhanced network virtualization and network management technologies such as MPLS and eVPN, as well as connectivity to off-campus research facilities, including the College of Agriculture, several Life Sciences Research Stations, and the Division of Extension offices statewide
- Upgrade the BOREAS-Net RON to support 400Gbps or 800Gbps circuits, expected Summer 2023.
- Implement a shared 100Gbps connection to GPN in Kansas City for a redundant tertiary Internet2 access outside of Chicago.

Beyond 5 years (strategic planning)

- Connections to cloud connectivity are planned for multiple regions via Internet2 AL2S or perhaps stitching RONs together to connect to Ashburn, VA (Equinix facility) and via the Northern Tier Network to Seattle, WA (SIX exchange facility).

3.2 The Center for High Throughput Computing (CHTC) Case Study

Content in this section authored by Brian Bockelman, Derek Cooper, Miron Livny, and Lauren Michael from the Center for High Throughput Computing.

The Center for High Throughput Computing (CHTC) supports a variety of scalable computing resources and services for UW-affiliated researchers and their collaborators. High throughput computing (HTC) is the shared utilization of autonomous computational resources toward a common goal, where all the elements are optimized for maximizing computational throughput. CHTC provides over 400 million hours of computation for campus faculty on an annual basis. CHTC computation systems and personnel are funded by the National Science Foundation (NSF), the National Institutes of Health (NIH), the Department of Energy (DOE), the Morgridge Institute for Research, and various grants from the university.

3.2.1 Science Background

CHTC is not a single science but rather an organization that supports a broad set of sciences through development and promotion of advanced research computing on campus, nationally and internationally. CHTC has a broad mandate to support science though certain specialization areas where the approach to computation best fits specific scientific domain problems.

On campus, CHTC broadly uses the campus network to advance computing by transferring significant data flows across campus from digital data storage facilities on or off-campus to various computation resources primarily on campus for specific workflows and users. These include physical sciences groups such as the Wisconsin Energy Institute (WEI), the Wisconsin IceCube Particle Astrophysics Center (WIPAC) and High Energy Physics. It also includes many life science, social science, and humanities use cases.

Users generally submit jobs to a Condor submit node host, which then identifies worker nodes (usually on campus) that meet the specific requirements for the job, which is then scheduled to be performed on the worker nodes. In addition, approximately 10% of all UW computation hours (over 40 million compute hours) are obtained by scavenging unused computer cycles at over 60 other participating institutions across the United States using the Open Science Grid (OSG).

3.2.2 Collaborators

There are approximately 296 projects and at least 1583 users (many group accounts shared by multiple people) utilizing CHTC computation capabilities on campus to advance their respective science.

CHTC collaborates with research and science groups across the United States and internationally. CHTC works with very small to large institutions to enable all-scales

of research facilities. To meet this goal, a variety of tooling such as SCP and rsync are used.

UW-Madison collaborator groups include researchers in academic departments (agricultural applied economics, agronomy, anesthesiology, animal sciences, astronomy, atmospheric and oceanic sciences, bacteriology, biochemistry, biological systems engineering, biology, biomedical engineering, biostatistics and medical informatics, chemical and biological engineering, chemical engineering, chemistry, civil and environmental engineering, computer sciences, dairy science, economics, educational psychology, electrical and computer engineering, engineering physics, entomology, forestry, forest and wildlife ecology, genetics, geography, geoscience, history, horticulture, industrial engineering, life sciences communications, material science and engineering, mathematics, mechanical engineering, medical microbiology and immunology, medical physics, medicine, neuroscience, nuclear engineering, nutritional sciences, oncology, pathobiological sciences, pathology, pharmacy, physics, plant pathology, political science, primate center, psychiatry, psychology, radiology, School of Business (finance, marketing, real estate), soil sciences, space sciences and engineering, statistics, urology, and zoology as well as research centers, research cores, and other research services (biological magnetic resonance bank, bionates, computer-aided engineering, EarthCube, Engine Research Center, high-energy physics, Laboratory for Molecular and Computational Genomics, Laboratory for Optical and Computational Instrumentation, Nelson Institute, primate center, public affairs, small molecule screening facility, Waisman Center, WI Ice Cube astrophysics, Wisconsin Electric Machines and Power Electronics Consortium, Wisconsin Energy Institute).

A few specific large collaborations include:

- The high-luminosity Compact Muon Solenoid (CMS) LHC collaboration
- The Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP) collaboration
- The Wisconsin IceCube Particle Astrophysics Center (WIPAC)
- The Laser Interferometer Gravitational-Wave Observatory (LIGO) consortium

CHTC needs national and international network connectivity to other universities and federal labs with the ability to troubleshoot the network using monitoring and performance management tools.

3.2.3 Instruments and Facilities

The main data archive used for on-campus job submissions is 150TB of scratch storage (non-resilient) in the Computer Sciences building, although significant data clusters also exist at Chamberlin Hall (for HEP), the Space Science and Engineering Center (SSEC), the Department of Atmospheric and Oceanic Sciences, the Wisconsin Energy Institute (WEI), Ice Cube (WIPAC), the Wisconsin Institute of Discovery, Computer-aided Engineering (CAE), the Biomedical Research building

(BMRB), biostatistics, Van Vleck (mathematics), the Waisman Center, and several other campus facilities.

Researchers use network protocols to transfer data, including HTTP, SCP, and GridFTP. Typical jobs transfer 1-10GB while machine-learning jobs can range from 10GB – 500GB per job. CHTC doesn't support the same scale of data sets as a traditional supercomputer or high-performance computing applications, although specific job flows can be moderate to large, including jobs related to machine learning, which often require 150TB of storage in contrast to a HEP CMS job which uses only 30 to 40PB. Data is transferred from a data storage device (on or off-campus) to a CHTC submit node. The submit node transfers the data to a set of worker nodes, which then transfer it back to the submit node. The transformation, longevity, and general usage patterns of the scientific data vary based on the users. CHTC does not archive data for faculty members but rather provides a scratch space that is not backed up.

Each campus group upgrades their own cluster if and when needed. For example, HEP replaces approximately 25 to 40 machines per year with more powerful units. In addition, CHTC has received university support for more specialized resource needs, such as large memory computation nodes and more recently GPU nodes.

CHTC is also implementing a StashCache data federation node on campus and is collaborating across the U.S. and portions of Europe⁷. StashCache provides organizations with a method to distribute their data in a scalable manner to thousands of jobs without needing to pre-stage data at each site. The StashCache data federation is best suited for per-job data set sizes between 1 and 50 GB, with no more than 1 TB for a complete workflow. StashCache can be used to work with certain jobs that require a significant amount of repeated data reads from disk (e.g. 200 times) to perform multiple computations (e.g. LIGO project).

UW-Madison Division of Information Technology (DoIT) with support from the Vice-Chancellor of Research is also implementing a research storage facility for every faculty member. Each faculty member will be allocated five terabytes (5TB) of storage with more available for purchase at low rates. While this is not a CHTC initiative, it is anticipated that many faculty will use this allocation to store and retrieve data from CHTC computation jobs which will have impacts on the campus and wide area networks. CHTC is concerned about the performance of the high-capacity storage facility and the data transfers speeds that will be supported between the storage facility and the CHTC computation worker nodes.

Another initiative CHTC is undertaking is to move from a Gluster file system to an open-source file system called Ceph. Ceph uses Amazon Web Services (AWS) Simple Storage Service (S3) interfaces. Moving to Ceph will enable CHTC to better interface

⁷ <https://opensciencegrid.org/docs/data/stashcache/overview/>

to object storage systems at other national labs using S3 protocols as well as support better use of AWS' S3 storage services which researchers are demanding.

In addition, CHTC has acquired a significant GPU allocation at Argonne National Lab in the Chicago area. This was a demand that was not present three years ago, and work is ongoing to meet it.

Finally, CHTC continues to work towards providing integrated support in HTCondor for "burst workloads", where a researcher requires a short-term, large-scale computation resources from a cloud provider or other large national resource (e.g. TACC, SDSC, ANL, Oak Ridge National Lab, NERSC). An example of this need is more recent work with the Ice Cube project, where specific interesting astronomical observations require approximately 10,000 computer cores to quickly process the data and advance the science.

3.2.4 Process of Science

Standard access to CHTC resources are provided to all UW-Madison researchers, free of charge. Even external collaborators with an on-campus sponsor may be given access to resources. The group also offers hardware buy-in options for priority access to computing capacity on a case-by-case basis, though standard access is more than sufficient for the vast majority of CHTC users.

3.2.5 Remote Science Activities

CHTC works with approximately 80-100 higher education and national lab facilities across the United States as well as worldwide, including sites in Asia, Europe, and South America. Robust network connectivity to these locations is a requirement in all cases.

3.2.6 Software Infrastructure

CHTC's software infrastructure is HTCondor along with several other scientific domain specific "glide-in" packages to enable domains to seamlessly work with the main computation grid.

3.2.6.1 Data Archetypes

- Lots of data processing 1:1 transformation (more data intensive process)
- May have 1 to many back to 1 (data expands and then reduces back to original size)
- Simulation: configuration file used to simulate a process which then provides smaller outputs to the researcher
 - (now/future) training for machine learning: very large data input sets that are difficult to split into smaller worker loads so specific jobs lock up machine(s) for 1-2 days and then provide smaller output of approximately 1TB

- GPU use hasn't developed to a uniform (batch) factor with jobs queuing up yet partially due to limited (once/yr) and not centralized training (mostly done in labs yet) though expect this to occur within the next 2-3 years as computer science research on GPUs moves to more common uses in various domains (imaging, analysis, biology)
- CHTC is starting to see the emergence of pre-canned gateways (i.e. webpage, fills in parameters and then hits submit) with groups like the Great Lakes
- Bioenergy Research Center (GLBRC) starting to hit scale
- Most CHTC users use rsync, SCP, FTP to move data around
- Department of Energy (DoE) groups tend to use Globus while the National Institutes of Health (NIC) tend to use IBM's Aspera product to move data

3.2.7 Network and Data Architecture

Please see [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#)

3.2.8 Cloud Services

CHTC has existing integrations with Google, Azure, and AWS from a technical provider perspective, though there isn't significant utilization of this capability yet. Researchers are still finding cloud costs are higher than local resources though that cost model may shift at some point. The largest cloud services use case is a cloud bursting model, such as the Ice Cube group, which occasionally need to immediately and quickly evaluate interesting astronomical events and are willing to pay approximately \$1500/mo for intermittent use of cloud-based computation resources to meet their scientific needs.

Besides cloud bursting opportunities, several one-off activities use cloud services (e.g. MongoDB) where on-campus groups don't wish to use the campus computing infrastructure.

3.2.9 Known Resource Constraints

- CHTC has concerns about lower speed network connections for some buildings on campus that are still being connected at only 1Gbps
- WAN connectivity is cared for though on-campus networking
 - Lots of 1GE yet on campus. When CHTC begins supporting additional machine learning workflows, having 1GE connections may impact the overall performance
 - Making sure the high-capacity CHTC connectivity is kept in place at the Discovery Center (location of CHTC) (currently 4x10GE) vs upgrades to 40GE or 100GE
 - Campus has very distributed cyberinfrastructure and needs to maintain significant backbone links to various distributed nodes to ensure ample bandwidth exists
- It might be more efficient if the UW-Madison administration considered upgrading research computing infrastructure using an approach such as a fixed budget allocation as opposed to the current situation where one-time

funds are made available approximately every five years. It is possible another approach would enable access to additional computing resources at potentially lower cost.

- The DoIT 5TB baseline storage approach for faculty is a potential benefit to many research teams. There may be an emerging demand for a middle tier approach for end users needing 5-50TB that should also be considered. In general, experience has shown that when storage needs are over approximately 50TB or 100TB, typically the research teams will have sufficient resources to acquire the data storage in house.

3.2.10 Parent and Affiliated Organizational Cooperation

- Open Science Grid (OSG)
- Receive great support from local campus or regional network; from campus backbone up to national networks Condor is good
- Received CC* funds from NSF in 2012
- Need continuous upgrades to sustain computation resources

3.2.11 Outstanding Issues

- A previous incident showed how CHTC resource contention on campus could cause problems. During a routine window where CHTC was attempting to prepare a system for upgrade, a number of jobs were cleared off. During this time window, a heavy campus user (e.g. [Section 3.6 IceCube Neutrino Observatory Case Study](#)) started “flocking” jobs to the now available resources that were set for upgrade. This sudden influx of use caused both computational disruption for other users, as well as network congestion. Campus IT staff are working with CHTC on network issues, and CHTC has started working to address the phenomena with software upgrades.
- There is a need for new training modules for relatively new faculty to enable them to do certain computation activities on campus. For example:
 - “Hello world” training for faculty starting from scratch
 - How to render a portion of a complete picture into parts and then reintegrate back into the whole picture again.
 - Options for storing large data sets. E.g. New faculty woke up and have 1TB of data. They need guidance on where to store it, how to access research computing infrastructure; how to delete data from scratch storage, how to archive data
- At UW-Madison, billing for cloud resources is problematic and this matters to CHTC because they support elastic computing in the cloud through HTCondor. For example:
 - CHTC would like to be able to obtain cloud provider account and have campus funding string connected to account
 - Campus shows all budgeting and billing through an archaic process that doesn’t work well for researchers
 - Tools are just starting to be developed (i.e. Cloudwatch) to monitor usage

- It is currently a separate step to get added to bill monitoring software alerts, with only one person at DoIT responsible for taking care of user requests
- There is a perception among researchers that UW-Madison should waive overhead charges on cloud purchases similar to some other universities (i.e. UCSD).
- Groups like CHTC and WEI/GLBRC are big enough to do billing right
- In short, technology has progressed further/faster than campus tech

3.3 Ahlquist Lab Virology Research

Content in this section authored by Brian Bockelman, Derek Cooper, Miron Livny, and Lauren Michael from the Center for High Throughput Computing.

Paul Ahlquist is the Paul J. Kaesberg Professor of Oncology and Molecular Virology, a Professor in the Department of Plant Pathology, the Director of the John W. and Jeanne M. Rowe Center for Research in Virology, Associate Director of Basic Research for the UW Carbone Cancer Center, and an Investigator for the Howard Hughes Medical Institute.

3.3.1 Science Background

A primary focus of the Ahlquist lab is the study of novel, RNA-based pathways and virus-host interactions underlying replication, gene expression, and evolution by positive-strand RNA viruses, the largest class of viruses. Positive strand RNA viruses include many important human pathogens, such as hepatitis C virus, which chronically infects nearly 3% of the world population, and causes progressive liver damage and liver cancer, and the new SARS coronavirus. The Ahlquist lab is also studying selected replication processes of a reverse-transcribing virus, the hepatitis B virus, which is also a major human tumor virus. The studies integrate molecular genetics, genomics, biochemistry and cell biology to address fundamental questions in virus replication and virus-cell interactions.

Viruses are divided into six distinct classes differing in the type of nucleic acid in the virus particle and its replication pathways. Recently, research in the Ahlquist lab discovered multiple, detailed structural and functional parallels among the replication complexes of three of these six virus classes: positive-strand RNA viruses, retroviruses, and dsRNA viruses. These results imply a significant functional and evolutionary unification within virology, and have opened the door to many exciting new experimental questions and approaches that are now being pursued.

The Ahlquist lab has the first higher eukaryotic viruses that can direct genome replication, gene expression, and virion assembly in the genetically tractable yeast *Saccharomyces cerevisiae*. Using yeast genetics and genomics, a growing number of host genes required for viral RNA replication have been identified and research is ongoing to understand how these function with virus-encoded helicase- and polymerase-like replication factors and diverse cis-acting viral RNA signals to direct assembly and function of the membrane-bound viral genome replication complexes and other processes.

3.3.2 Collaborators

The Ahlquist Lab research is done using Titan Krios Cryo-EM machines located at the Howard Hughes Medical Institute (HHMI) locations, including the Janelia Research Campus in Ashburn, Virginia and the Pacific Northwest Center for Cryo-EM in Oregon.

3.3.3 Local and Remote Instruments and Facilities

The Ahlquist Lab use case involves researchers bringing hard drives to the remote facilities for use during their 72 hour booking of the scientific instruments. Typically this results in a data set size of 10-40TB. The data is then pre-processed by the remote facility and then shipped back to the Discovery Center at UW (where the CHTC is located, for processing. The stacks of images are built into 3d images and then reconstructed. Reconstruction can involve as many as 120,000 images per sample to permit looking for commonalities.

3.3.4 Process of Science

The process for the Ahlquist Lab is as follows:

1. Research teams travel to the Howard Hughes facilities, and bring their own storage capabilities with them.
2. An allocation of time on instruments will be used fully. During this time:
 - a. Samples are cycled through the instrument.
 - b. Initial calibration of sample data can be performed using local machines. Analysis work is not possible using local resources.
 - c. All raw and calibrated results are stored to local removable storage.
3. Research teams travel back to UW-Madison with removable media.
4. Processing is done primary using resources described in [Section 3.2 The Center for High Throughput Computing \(CHTC\) Case Study](#).
 - a. Data is downloaded from removable media to local HTC scratch storage resources.
 - b. The analysis process involves:
 - i. Pre-processing each image (10-100s of MBs to as much as 1GB)
 - ii. Multiple images are combined (stitched) when needed
 - iii. A manual step in which a researcher views each image and highlights areas of interest in order to train the software on what to look for. .
 - iv. An automated “Machine Learning” step, using the training data, to sift through the entire data set.
5. After analysis, curation is done by each researcher.
 - a. Typically, both the raw and processed data sets are preserved.
 - b. All results are migrated out of the CHTC facility after processing to departmental or personal storage resources.
 - c. There is no dedicated portal system for data sharing.

3.3.5 Software Infrastructure

The Ahlquist lab primarily uses the Matlab application for experiments, deployed via HTCondor⁸, along with a mixture of helper scripts written in a variety of scripting languages.

At present time, there is not a significant need for automated methods of data capture, processing, or curation.

3.3.6 Network and Data Architecture

For an overview of campus networking, please see [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#).

A high speed 20 Gb converged network ties together the various elements of technology within, and external to the laboratory. Each building network port has a one gigabit connection to a ten gigabit uplink on each floor. The data center has 10 Gb to 40 Gb uplinks for storage and servers. Intra-building fiber optics allow highly scalable networks between lab spaces and computational infrastructure while multiple fiber optics trunks among the building and diverse locations on campus provide fault tolerance as well as connect research collaborators. The University of Wisconsin provides 20GB and 40 Gb (CHTC) connectivity from Discovery to Campus, as well as to the external networks including ESnet, Internet2, and BOREAS-Net.

3.3.7 Cloud Services

Cloud services are used sporadically, mainly as a backup mechanism (e.g. Google Drive) for research data. Most researchers have only treated this as a “one way” relationship (e.g. push data in when done with active work), instead of making it an active part of the workflow.

Cloud computation is not currently used in this use case, and is not currently being explored.

3.3.8 Known Resource Constraints

Storage is a known problem, but is generally left to individual research users to address. This manifests in two unique ways:

- Research data storage is costly, but a new campus-wide storage initiative is now providing 5 TB per PI at no cost and additional storage at lower pricing than was previously available. Intra-campus data movement (e.g. to and from CHTC from other parts of campus) has historically been challenging due to performance limitations rooted in the security infrastructure. For example, a known problem with some of the firewall infrastructure had per-flow

⁸ <http://chtc.cs.wisc.edu/matlab-jobs.shtml>

limitations that resulted in the wide-scale use of portable storage instead of using networks to migrate data to locations around campus.

3.4 Huisken Labs: Microscopy Case Study

Content in this section authored by Brian Bockelman, Derek Cooper, Miron Livny, and Lauren Michael from the Center for High Throughput Computing.

3.4.1 Science Background

Jan Huisken created a technology called light sheet microscopy, which introduces minimal perturbation to the live specimen and captures sensitive biology in its truest functional state. His work focuses on the early development of zebrafish, a model organism widely used because its transparency allows direct observation of intact systems, and light-sheet imaging can produce striking images of cellular movements and beating hearts and real-time development of organs. The microscope set-up consists of an array of up to 12 cameras each capturing and streaming 800 Mbps of data (100 frames per second, 40 megapixels/frame) to an analysis server that is connected to the file server system. The file system server consists of 100 TB of SSD and 600 TB of hard drive with 4x40 Gb connections to the analysis server. The amount of useful data produced by this array over the course of a year can be as much as 6 PB.

3.4.2 Collaborators

The amount of data produced by this technique makes it challenging for remote observation of the data at other sites on campus as well as off-campus, due to bandwidth constraints. To make the technology accessible to other researchers, the Huisken lab has developed a portable, shareable light sheet microscope, called Flamingo, which was designed to meet the needs of researchers that can't afford their own light sheet microscope. Their approach shrinks a tabletop-sized technology down to the weight and dimensions of a suitcase.

3.4.3 Local and Remote Instruments and Facilities

The Huisken lab Flamingo project involves the sharing of instruments that are mailed to a lab anywhere in the world, configured remotely by Morgridge engineers, and then run experiments for one to three months. Data is stored on site, but if there are collaborative interests, the data may also be shared with the UW Discovery Center. UW collaborators also use the Flamingo microscopes to do experiments at the Discovery Center, and then ship the data to their departments. The cameras can generate as much as 850MBs per second, and it is common for experiments to generate 25TBs of data or more, depending on the frequency and duration of the sampling.

3.4.4 Process of Science

The process of science for the Huisken lab Flamingo project is as follows:

1. The instruments are shipped to the location that has requested the loan.
2. At the remote site, it is brought online and then controlled by UW-Madison staff

3. Researchers at the remote site prepare samples.
4. Generally, remote site researchers rely on UW-Madison staff to control the sampling/calibration process, but it is an option to control this processes locally as well.
5. Sampling and calibration is accomplished with limited local storage and computation resources to ensure things are working as expected
 - a. There must be local storage for results, but this can be plugable media.
 - b. There must be local computation for the basic analysis of calibration results.
6. When the local research is complete, the instrument is re-packed and sent back to UW-Madison

3.4.5 Software Infrastructure

The Huisken lab relies on a large number of analysis tools that are curated and written by local research groups for interacting with the instruments. These tools are typically deployed via HTCondor resources spread around the campus.

At present time, there is not a significant need for automated methods of data capture, processing, or curation.

3.4.6 Network and Data Architecture

Please see [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#)

3.4.7 Cloud Services

Cloud services are used sporadically, mainly as a backup mechanism for data generated at the main microscopy facility at the Wisconsin Institutes of Discovery and for researchers using the Flamingo instrument (e.g. Google Drive) for research data. Most researchers have only treated this as a “one way” relationship (e.g. push data in when done with active work), instead of making it an active part of the workflow.

Cloud computation is not currently used in this use case, or being explored.

3.4.8 Known Resource Constraints

Storage is a known problem, but is generally left to individual research users to address. This manifests in two unique ways:

- Research data storage is costly. A newly launched research data storage service provides up to 5 TB per PI at no cost. While this helps address some of the need for storage, maintaining storage for the large data volumes produced by light sheet microscopy generally requires additional investment by individual departments and units.
- Intra-campus data movement (e.g. to and from CHTC from other parts of campus) has historically been challenging due to performance abnormalities rooted in the security infrastructure. A known problem with some of the

firewall infrastructure had per-flow limitations that resulted in the wide-scale use of sneakernet techniques to migrate data to locations around campus. The Flamingo package represents on method of addressing these limitations.

3.5 Tier-2 Computing Center for the Compact Muon Solenoid (CMS) Experiment at the Large Hadron Collider (LHC) Case Study

Content in this section authored by Dan Bradley, Ajit Mohapatra, and Carl Vuosalo from the Department of Physics.

The Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland, exists to study the Higgs Boson, to search for its potential partners, to lead searches for Dark Matter, and to make extensive studies of Electroweak phenomena. Data describing collision events in the LHC are produced at CERN. The data streams from CERN to regional Tier-1 centers around the world where it is archived on tape for long-term storage. In the Americas, the Tier-1 center is the Fermi National Accelerator Laboratory (Fermilab), which provides 40% of the overall computational Tier-1 resources for CMS. From Fermilab, selected data is distributed to disks at the Wisconsin Tier-2 center and others for temporary storage while it is being processed and analyzed.

As a whole, the US provides 50% of the world-wide total Tier-2 computing resources. There are Tier-2 centers in the US at University of Wisconsin, MIT, the University of Florida, Purdue University, Vanderbilt University, the University of Nebraska Lincoln, the University of California San Diego, and Caltech University. The University of Wisconsin is a critical part of CMS computing, and provides about $\frac{1}{8}$ of the US resources.

3.5.1 Science Background

Analysis of the data occurs on computers at the Tier-2 centers. Most data read by analysis software is localized, but this is only loosely optimized. Analysis activity at Tier-2 and Tier-3 centers may result in streaming data directly between computing centers, which may bypass the local storage system. Previous generations of the analysis featured a rigid arrangement of data exchange (e.g. higher tiers to lower tiers), but usually opportunistic approaches are being used, which has resulted in more chaotic data flows and analysis patterns.

Simulated data (e.g. “monte carlo” data) is produced by computers at the Tier-2 centers and is either stored there for use in analysis or is transferred to the Tier-1 for long-term storage. Production of these simulations is a primary use case for Tier-2s, and most effort, outside of actual analysis, is spent on this activity.

3.5.2 Collaborators

The Wisconsin CMS Tier-2 center is embedded within the Open Science Grid. It is closely affiliated with the CMS Tier-1 center at Fermilab and the other US Tier-2 centers.

The Wisconsin CMS Tier-2 center supports approximately 250 physics analysts in addition to remote users at US Tier-3 sites and other facilities.

3.5.3 Instruments and Facilities

The LHC CMS experiment is the source of the data for the UW Tier-2 site, and is primarily provided by the CMS Tier-1 computing center at Fermilab in the form of archives on tape and associated disk cache and technologies for transferral of the data. The OSG and the CMS collaboration together package the primary open-source middleware software stack that the Tier-2 computing center relies upon to transfer, store, and access the data.

The UW networking team (DoIT) provides the LAN and WAN infrastructure and support for the 100Gbps connection to the campus backbone and wide area network. The campus backbone also enables the use of computers at the Center for High Throughput Computing (CHTC).

The Wisconsin CMS Tier-2 computing center currently consists of 13,000 compute cores and 8.2PB of disk. For performance and reliability, two copies of files are stored, so the usable disk space is 4.1 petabytes. 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. The average file size of the current 25M files is around 100MB. Around 250 users are served by the facility. Additional guest users from the OSG and CHTC make use of the facility when there is a lull in demand by CMS users.

3.5.4 Process of Science

Users of the Wisconsin CMS Tier-2 center run simulation and data analysis compute jobs on the HTCondor batch system. The compute jobs are submitted to the site remotely through the OSG middleware or using ssh access to the interactive job submission environment. The data sets hosted at the Wisconsin Tier-2 site are selected by the CMS collaboration, but users can also make requests for specific datasets.

In preparation for the onset of HL-LHC operation in 2026, improvements to CMSSW are planned to make use of GPUs, machine learning, and SIMD. Due to budget constraints, large-scale GPU deployment is not expected at the Wisconsin CMS T2 or other dedicated CMS sites, although a growing amount of GPU-enabled computers will be available at DOE facilities and commercial clouds. In the scenario where a significant amount of compute power is obtained from this type of external sites, the storage system at the Wisconsin CMS Tier-2 site will be under increased demand for access via the wide area network connection to those external facilities. In other words, it is likely that storage services will continue to be concentrated at CMS dedicated sites while an increasing fraction of the compute will be external.

3.5.5 Remote Science Activities

The LHC CMS experiment at CERN in Geneva, Switzerland, is the source of the raw data. The CMS Tier-1 computing center at Fermilab provides most of the data for the Tier-2 site in the form of archives on tape and associated disk cache and technologies for transferral of the data. The present pattern of data exchange between CMS facilities is expected to continue, but there will be increasing need, especially in HL-LHC operation starting in 2026, for large data flows between the Wisconsin Tier-2 site and additional compute sites such as DOE facilities and commercial clouds.

3.5.6 Software Infrastructure

Several software systems are currently used for data access and transfer. The CMS experiment has developed the PhEDEx system for managing data placement at computing centers (scheduling transfers and deletion). This will be replaced by Rucio in 2020.

Currently, data is written to the UW Hadoop storage system via GridFTP that run on most of the compute nodes, so there are roughly 300 GridFTP servers with 1Gbps network connections. Linux Virtual Server (LVS) is used to provide a single point of access for this server cluster⁹. In 2019, the project will transition from using GridFTP for writes to using xrootd, which is already used for read access. It is deployed in a similar way, with a large number of 1G servers collocated on computational machines.

Data is read using xrootd. The CMS software can make use of xrootd in different ways, including caching the data on the disk of the local computational machine. However, the primary mode of access is a streaming read with no disk-based cache.

RAW data recorded by the CMS detector is stored in binary formatted files by the online data storage system. The binary formatted files are then transformed at CERN into CMS ROOT-based event formatted files. Subsequent full reconstruction of the RAW detector data performed at the Tier-1 sites results in intermediate data formats with varying degrees of detail, size, and refinement which is directly used for physics analysis. The series of (re)processing tasks are performed using a collection of software tools designed and developed by the CMS collaboration, and is referred to as CMSSW.

3.5.7 Network and Data Architecture

The Wisconsin Tier-2 computers currently have 1G links to the UW LAN, which is composed of several layers. At the rack layer, there are Cisco Nexus 2200 Fabric Extenders (48G backplane). These are connected via 4x10G links to Cisco Nexus

⁹ <https://iopscience.iop.org/article/10.1088/1742-6596/1085/3/032004/pdf>

5000 switches (~1T backplane per computational resource). The two room switches are each connected via 8x10G links to the building Cisco Nexus 7000 switch. This is connected via 100G to the campus 100G backbone, which then connects at 100G to Chicago and the national research networks, including Internet2 and ESnet.

There is no facility-scale firewall in the network path, and most of the computers have publicly routable IP addresses. All firewalling is done internally by each computer.

perfSONAR is used to monitor the interconnectivity with other sites. While useful, the bandwidth measurements provided by perfSONAR are limited by the 10G connection of the bandwidth monitoring node. Ideally, the bandwidth monitoring would instead reflect the more relevant question of what is the aggregate bandwidth between the approximately 300 1G data transfer nodes and the data transfer nodes at other sites.

In the future, some existing 1G links will be updated to dual bonded (2G). New computers are planned to be purchased with 10G NICs.

3.5.8 Cloud Services

At present, cloud services are not used for computing, data storage or analysis, since it is beyond the current scope of the CMS Tier-2 program model and funding support. However, the CMS collaboration is actively pursuing the use of HPC and cloud computing resources to supplement the existing dedicated resources it owns in the form of computing centers around the world with a well defined Tier hierarchy. This is driven by the continuous increase of the computing/storage resources requirement foreseen for the High-Luminosity LHC era. As part of this effort, the Tier-1 center at FNAL has made significant progress in the development and operation of the HEPCloud framework that supports dynamic integration of heterogeneous HPC (NERSC, TACC, PSC) and Cloud (Amazon, Google etc.) resources elastically and as a transparent extension of the Tier-1 facility. In the future, implementation of a similar framework at the CMS Tier-2 centers in the US is highly likely in order to satisfy the evolving needs of the CMS experiment.

3.5.9 Known Resource Constraints

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

The wide area network connection is crucial for connecting the storage system to additional computing power from external sites such as DOE facilities and commercial clouds. The aggregate data read rate by jobs in the facility currently fluctuates from 20 to 10 Gbps. With 50x demand, this would be on the order of 1 to 5 Tbps, which, if it were directed to external sites, would overwhelm the current 100 Gbps WAN connection.

3.5.10 Parent and Affiliated Organizational Cooperation

The Wisconsin CMS T2 relies entirely on the campus network support group (DoIT) for the LAN and WAN hardware and configuration.

3.5.11 Outstanding Issues

Through the shutdown expected to take place in 2019-2020, the resources will not be idle. There are two activities that are ongoing:

- Simulation production
- Re-construction of all events from Run 2

In the case of simulation, it is not expected that the load will change significantly beyond production times. Data volumes will gradually increase to mimic the upgraded detector as the time gets closer to startup.

In the case of re-construction, all data will be read from CERN (off of tape) and distributed to the various centers around the world to rebuild the data sets. This can be viewed like a production use case, just constant for a period of weeks.

3.6 The Great Lakes Bioenergy Research Center (GLBRC) Case Study

Content in this section authored by Dirk Norman from the Great Lakes Bioenergy Research Center (GLBRC) / Wisconsin Energy Institute.

The Great Lakes Bioenergy Research Center (GLBRC) is a U.S. Department of Energy-funded Bioenergy Research Center led by the University of Wisconsin–Madison. With Michigan State University (MSU) and other partners, GLBRC is developing sustainable biofuels and bioproducts made from dedicated energy crops grown on marginal lands. The mission is to create biofuels and bioproducts that are economically viable and environmentally sustainable. The GLBRC is a part of the Wisconsin Energy Institute (WEI), and receives funding from the DOE Office of Science via the Biological and Environmental Research (BER) program office.

3.6.1 Science Background

The GLBRC works with a variety of datasets from images of field plots to genomic datasets. This data is created at the major institutions (UW and MSU) as well as DOE facilities such as the DOE Joint Genome Institute (JGI). As a result of such a wide set of research types, the data workflows are extremely varied. Some of the larger workflows involve transferring large images and large (up to 10s of TBs/day) genomic datasets between JGI, UW, and MSU.

One research goal is to better understand the fuel and chemical properties that result from biological sources via the study of genomics. Typical questions this facility assists to answer include:

- Evaluating where to grow crops, which may be different than where they are grown currently;
- Understanding the impact of climate and weather on crop growth; and
- Evaluating the impacts of feeding crops and when is the most effective times to do so.

3.6.2 Collaborators

GLBRC has several core collaborators:

- DOE Joint Genome Institute (Walnut Creek, CA)
 - Dozens of users creating hundreds of TB of data
 - All data transferred to UW data center
 - Connected via ESnet
- UW Biotech Center (Madison, WI)
 - Dozens of users creating hundreds of TB of data
 - All data transferred to UW data center
 - Connected via campus networking
- MSU Research Technology Support Facilities (East Lansing, MI)
 - Genomics, Proteomics, and Mass Spectrometry

- Dozens of users creating hundreds of TB of data. MSU maintains this data: They have collected 600TB of data over 10 years and expect another 600TB to be generated within the next 5 years
- Connected via regional peering in Chicago

Additionally, other collaborations (and data sharing arrangements) exist with:

- Joint BioEnergy Institute (JBEI) at Lawrence Berkeley National Laboratory
- The Center for Bioenergy Innovation at Oak Ridge National Laboratory and National Renewable Energy Laboratory
- The BioNanotechnology Laboratory (BNL) at the University of Illinois Urbana Champaign
- Environmental Molecular Sciences Laboratory at Pacific Northwest National Laboratory

Lastly, it is common to utilize photon sources, including the Advanced Light Source at the Lawrence Berkeley National Laboratory and the Advanced Photon Source at Argonne National Laboratory, in the course of research.

3.6.3 Local and Remote Instruments and Facilities

Each of the collaboration groups has a unique set of instruments.. Common instruments include genomics sequencers, microscopy, mass spectrometry, satellite and drone images, and a variety of sensor data.

Genomics is the current largest data producer, with a single sample producing data sets that range in size from 1-50GB, and tens or hundreds of data sets can be produced in a single day. Results from microscopy and mass spectrometry are smaller, and typically range in the 100s of MBs. Satellite and drone images can vary in size between MBs and GBs (depending on the satellite used and resolution of images), with entire data sets ranging from 1-10 GB. Lastly, sensor data is KB to MB in size, but can produce large volumes for heavily sensed environments, generally fields.

It is estimated that datasets currently fall into these categories for the time ranges requested:

- Present-2 years = 200TB/year of data created
- Next 2-5 years = 250TB/year
- Beyond 5 years = unknown

3.6.4 Process of Science

GLBRC frequently moves datasets between facilities and the two primary campuses at UW and MSU on a daily basis. These data sets are affiliated with the instrumentation mentioned above.

A typical use case (based on sensor data) is as follows:

1. Field deployed sensors will transfer data to a local base station. Sensors include:
 - a. Soil observations (chemistry, temperature, water, etc.)
 - b. Drone or fixed camera images
 - c. Weather (local stations and regional observations from NOAA)
 - d. Aerosol monitoring (gas levels, etc)
2. Base station transmits data back to GLBRC for storage.
3. Other data collection that is affiliated with this research (e.g. produced at JGI/Light Sources, etc) is also transmitted back to GLBRC for storage/analysis.
4. GLBRC utilizes dedicated HTC resources and HTCCondor for processing at both MSU and UW.
5. Results are stored locally. There is not currently a sharing mechanism (e.g. portal, etc) available.

Efforts are underway to explore flocking to institutional resources

3.6.5 Software Infrastructure

GLBRC uses Globus to transfer data between JGI (e.g. data storage and analysis) and the UW data center. Other data sync happens through use of older tools like rsync or object transfer via HTTP/multi-part.

3.6.6 Network and Data Architecture

GLBRC's primary data center is at UW Madison, within the Wisconsin Energy Institute (WEI) building. This building is connected to the UW backbone via 4x10Gb connections. Within the data center, all equipment is interconnected via 10Gb networking. GLBRC is not currently connected to the UW Science DMZ infrastructure, but could be in the future.

3.6.7 Cloud Services

There are no current plans to integrate the current workflow into a cloud environment for purely cost reasons. Pricing models, given the amount of data GLBRC uses, makes this prohibitive. Despite this, most of the workflow of GLBRC is being converted to utilize "cloud-enabled" technologies (e.g. Docker) that would facilitate a future migration if the costs allow.

3.6.8 Known Resource Constraints

Data transfers between UW and MSU frequently are less than 100Mbps, which can limit the ability for the GLBRC to collaborate. An investigation with UW DoIT is ongoing.

Connectivity to some research facilities has been historically poor, for example for field locations such as those located in Hickory Corners and Kalamazoo, MI. It is challenging to get fiber, let alone cellular connectivity to sensor networks deployed

in crop fields. UW IT is working with local Wireless Internet Service Providers (WISPs) to establish some connectivity.

3.6.9 Parent and Affiliated Organizational Cooperation

GLBRC receives excellent support from the UW networking and networking personnel and have worked with them to receive NSF campus infrastructure (e.g. CC* program) funding previously.

3.6.10 Outstanding Issues

As mentioned above, transfer of data between UW and MSU has been problematic in the past. Multi-gigabit transfer speeds between the two primary campuses would assist when creating data analysis workflows to move easily between research groups that are collaborating across campuses.

3.7 IceCube Neutrino Observatory Case Study

Content in this section authored by Steve Barnett from the IceCube Neutrino Observatory.

The IceCube Neutrino Observatory is designed to observe the cosmos from deep within the South Pole ice. Encompassing a cubic kilometer of ice, IceCube searches for nearly massless subatomic particles called neutrinos. These high-energy astronomical messengers provide information to probe the most violent astrophysical sources: events such as exploding stars, gamma-ray bursts, and cataclysmic phenomena involving black holes and neutron stars.

The Antarctic neutrino observatory, which also includes the surface array IceTop and the dense infill array DeepCore, was designed as a multipurpose experiment. IceCube collaborators address several big questions in physics, including the nature of dark matter and the properties of the neutrino itself. IceCube also observes cosmic rays that interact with the Earth's atmosphere, which have revealed fascinating structures that are not presently understood.

The National Science Foundation (NSF) provided the primary funding for the IceCube Neutrino Observatory, with assistance from partner funding agencies around the world. The University of Wisconsin–Madison is the lead institution, responsible for the maintenance and operations of the detector. Funding Agencies in each collaborating country support their scientific research efforts.

3.7.1 Science Background

IceCube is a neutrino detector built at the South Pole by instrumenting about a cubic kilometer of ice with 5,160 light sensors. Neutrinos of terrestrial, galactic, or extragalactic origin that interact with the ice, produce Cherenkov light that is used by IceCube to reconstruct their original energy and direction. Learning about the neutrino properties this way increases understanding of why the Universe looks the way it does today and how it will evolve in the future.

3.7.2 Collaborators

The IceCube collaboration comprises 50+ institutions supporting 400 faculty, graduate students, postdocs, and staff around the world. UW-Madison is the lead institution for the collaboration. The heaviest concentrations of collaborators are in the US and Northern Europe. The current list can be found online¹⁰. The international team is responsible for the scientific program, and many of the collaborators contributed to the design and construction of the detector.

Many collaborating sites offer access to local computing resources as in-kind contributions to the collaboration. Within that group are several sites that provide more significant resource contributions, including:

¹⁰ <https://icecube.wisc.edu/collaboration/institutions>

- National Energy Research Scientific Computing Center (NERSC) – Archive of raw detector data and a selected subset of derived data sets (500TB/yr)
- Deutsches Elektronen-Synchrotron (DESY) – Zuethen – Archive of filtered data set, level 2 data sets (250TB/year)
- University of Maryland – 128 GPU compute cluster

3.6.3 Instruments and Facilities

The primary instrument is the IceCube detector comprising over 5000 optical sensors deployed deep in the Antarctic ice at the geographic South Pole. These sensors instrument a cubic kilometer of ice which provides the detection medium for neutrinos and other particles passing through the instrumented volume of ice.

The instrument produces about 1TB/day of raw data. A subset of the data (10%) is transmitted north daily via satellite and stored in the data warehouse at UW-Madison. To perform the data reduction, a small compute cluster is maintained at the site in Antarctica to perform the online processing tasks of event reconstruction, triggering, final selection, archiving, and transmission to the north.

Yearly (typically in October), a shipment of disks from the station is sent back to the US. Once they arrive, they are sent to NERSC for extraction and processing. A campaign to compare a subset of data against the complete dataset (for investigation of interesting events) is then performed.

The northern hemisphere facility at the University of Wisconsin-Madison holds the primary data store and a computational facility to provide the processing to make the data science ready as well as handle more data intensive analyses. At a high level, the UW-Madison facility looks like this:

- 7,000 compute cores
- 10PB of storage (primarily based on the Lustre filesystem)
- 300 consumer grade GPUs
- A Science DMZ for distributing data to collaborating institutions as well as to the Open Science Grid and other international computing facilities.

In the coming years, NSF upgrades will facilitate installation of 7 new ‘strings’ of sensors into the ice. These new devices will be more advanced, but the increases in data production are expected to be less than an order of magnitude.

3.7.4 Process of Science

One of the primary goals for IceCube is to elucidate the mechanisms for production of high-energy cosmic rays by detecting high-energy neutrinos from astrophysical sources. IceCube performs physics analyses with neutrinos spanning six orders of magnitude in energy. Its contributions to fundamental scientific research are widely

recognized as an enormous success, garnering special recognition as the 2013 Breakthrough of the Year, according to the British journal *Physics World*.

IceCube has also obtained world-leading results in several neutrino physics related areas. The oscillation of atmospheric neutrinos has been measured in a previously unexplored energy range from 10 to 60 GeV. World-best limits have been set on the interaction cross section of dark matter particles with ordinary matter for a number of leading theoretical predictions. The limits on sterile neutrinos have been improved by one order of magnitude over accelerator searches. These are only a few examples of a long list of searches for physics beyond the Standard Model that are part of the IceCube research program.

Conversion of event rates into physical fluxes ultimately relies on knowledge of detector characteristics numerically evaluated by running Monte Carlo simulations that model fundamental particle physics, the interaction of particles with matter, transport of optical photons through the ice, and detector optics and electronics. Vast datasets containing simulations of background and signal must be produced and cataloged for use by the data analysts.

The main data processing center for IceCube is located at UW-Madison. Several other collaboration sites provide CPU and GPU clusters that are used as a federated distributed resource. The federated IceCube clusters, including the main data center at UW-Madison, currently deliver up to 500,000 GPU hours and 10,000,000 CPU hours per month.

As the detector collects more data, and measurement precision improves with higher statistics, an excellent understanding of experimental effects that are potential sources of systematic errors becomes essential. One of these effects is the properties of the Antarctic ice in which IceCube is embedded, and how light propagates through that ice. A good understanding of these experimental effects requires large amounts of simulation. Pursuing efficient access to a large amount of GPU computing power is therefore of great importance to ensure that future IceCube analysis reaches the maximum precision and the full scientific potential of the instrument is exploited. Because of this, IceCube is actively exploring possibilities to make efficient use of external computing facilities such as OSG HTC, as well as HPC systems or Supercomputers, especially those that provide substantial GPU capacity.

A typical muon event in IceCube creates over 10⁷ Cherenkov photons in the sensitive wavelength range, presenting a considerable computational challenge when more than a billion such muons need to be simulated to represent just a few days of detector data. However, we note that describing propagation of a large number of photons in a transparent medium is a computational problem of a highly parallel nature that is very well suited for GPUs. Although a single thread runs

slower than a typical modern computer CPU core, running thousands of them in parallel results in the much faster processing of photons on a GPU.

3.7.5 Remote Science Activities

The IceCube detector itself is in one of the most remote locations in the world (the geographic South Pole) and is network accessible exclusively by satellite. The NASA TDRSS (Tracking and Data Relay Satellite System) as well as low bandwidth access via Iridium systems are commonly used.

For analysis, data is distributed to grid sites around the world as well as to compute resources at collaborating institutions.

The Antarctic base station is limited in terms of capability (e.g. network connection, power, computation), but is expected to be upgraded marginally over the coming years as instrument data increases and machine efficiency improves. The facility has shared computational resources for all projects and that is expected to remain the same. In addition, cloud resources are being evaluated for the reconstruction of interesting events that may then trigger alerts for follow up observations from other instruments.

3.7.6 Software Infrastructure

Most software used is Free/Open Source Software (FOSS) or developed entirely in-house.

- Archive and transmission of observational experimental data is handled by an in-house application called JADE that is written primarily in Java.
- Data set transfer is handled predominantly by HTTP, SCP, gridFTP, and GlobusOnline. Rucio will replace Globus in the near future.
- Production data processing is handled by an in-house framework called IceProd. This package, written primarily in Python, handles the tracking of computational jobs, transfer of data (stage-in, intermediate data sets, stage-out of final results), and status of data sets. Currently, it is used most heavily for the production of simulation, but it will be used in processing of experimental data as well.
- Analysis and other physics software are homegrown and use a variety of languages, including C++, Python, and Java.

3.7.7 Network and Data Architecture

IceCube at UW is moving to a new datacenter, which will facilitate 2 x 100G connectivity to the UW Science DMZ. The old connection was 8 x 10G on the campus enterprise network. With this upgrade, it is expected that performance between the UW facility and NERSC will increase substantially.

Internal network capabilities (to HTC machines and storage) is funded by the NSF directly, but will be funded/managed by UW IT in the future. This will include

management of 10Gbps connected storage and computation machines from the current 1G connected infrastructure. This move is expected to address the data size increases that will occur in the 2023 event horizon.

3.7.8 Cloud Services

IceCube has explored a “bundled analysis” approach to cloud use. The process involves deploying specially constructed software packages to cloud resources for short-term use when there is a demand or cost drops enough to make things efficient to do so. This has been used to explore “interesting” events in the data when there is a time-based need, and lack of available computational resources locally or at NERSC.

3.7.9 Known Resource Constraints

The migration to the new facility eliminates many of the current problems with network connectivity, including issues during data migration to NERSC and across campus in the past.

3.7.10 Outstanding Issues

Commercial cloud infrastructure has not been widely used due to the prohibitive cost models. Currently IceCube pays for power, cooling, and networking. With the move to the new datacenter, UW campus will absorb these charges. By reducing these costs, it is expected that more resources can be put into on-premise storage and computation resources. It is expected that the NSF E-CAS project¹¹ could facilitate resources to facilitate a migration of some of this work toward the cloud in 2020 or beyond.

¹¹ https://www.nsf.gov/awardsearch/showAward?AWD_ID=1904444

3.8 UW-Madison Cryo-Electron Microscopy Research Center (CEMRC) Case Study

Content in this section authored by Eric Montemayor and Elizabeth Wright from the Department of Biochemistry.

The UW-Madison Cryo-Electron Microscopy Research Center (CEMRC) is a cross-campus initiative led by a coalition of partners including the Department of Biochemistry, the School of Medicine and Public Health, the Morgridge Institute for Research, the UW Carbone Cancer Center, the Office of the Vice Chancellor for Research and Graduate Education, and the College of Engineering Nanoscale Imaging and Analysis Center.

It is located in the Hector F. DeLuca Biochemical Sciences Complex on the University of Wisconsin-Madison campus. The CEMRC is dedicated to providing instrumentation, technical assistance, training, and access to Cryo-electron microscopy (cryo-EM) for the UW-Madison research community. When fully operational, the CEMRC will operate four Thermo Scientific cryo-microscopes utilizing Single Particle Analysis, Cryogenic Electron Tomography, and Micro-ED techniques. The microscopes are overseen by experienced staff who offer consultation and training in negative-stain and vitrified sample preparation, single particle analysis, tomography, data processing and additional computational support.

3.8.1 Science Background

Cryo-EM is the fastest developing and growing field in structural biology. This technology can be used to determine the structures of individual macromolecules to atomic resolution and the structures of larger complexes, viruses, and cells from nm-level to atomic resolution. The resulting data is used to understand mechanisms that underlie the function of biological systems and can lead to the development of new drugs or therapeutics.

CEMRC is designed to be shared. There are 2 primary usage models:

- Users physically visit the site from domestic or international locations
- Users send samples, but all data processing is done by local staff

Users are expected to be from the campus, as well as other universities, non-profits, governmental bodies, and commercial entities. At this time, there are no expected use cases of 'remote control', e.g. physically operating the instrument using robotics or network control.

To support these use cases, a dividing line regarding data stewardship is being established to set expectations: chain of custody is from the facility to the user directly. The facility will not be able to handle long-term storage of gathered data,

nor will it be able to provide extensive computational capabilities. Users are expected to have a plan in place for data migration after use of instrumentation.

In total, the instruments in this facility are capable of producing data sets that range in size from 4-16 TB per day depending on the resolution of the samples. Local storage can be leveraged temporarily, but will be purged after a set period of time.

Limited amounts of local processing exist to calibrate samples, but extensive use of CPU/GPU processing must be leveraged on external resources. It is typical to require 40hrs or more of processing time (which includes data movement for the larger data sets) to produce visualizations.

3.8.2 Collaborators

The full collaboration space is unknown. In 2018 there was an expectation of approximately 100 users/groups that could take advantage of the hardware. By late 2019 this space will grow to around 1000 as more instruments come online, and more researchers aim to use the technology. Researchers using these facilities may include:

- Academic research laboratories located at UW-Madison and globally. Global industrial research laboratories.
- US: Wisconsin, Ohio, Georgia, Illinois, Florida, California.
- Global: UK, Denmark, Switzerland, the Netherlands.

3.8.3 Instruments and Facilities

The cryo-EM resource, once operational in early 2020, will support four cryo-electron microscopes. Each of the instruments has several types of detectors that capture imaging data. The data collected from each can be from 1-4 TB per day per instrument. The current trajectory for hardware development is such that machines become obsolete before components break or wear out, so it is expected that the data volumes for new machines will grow as they are released or acquired.

Within the facility, a 1PB storage server and a small 10 node GPU/CPU cluster (~10 nodes) to run pre-processing procedures and initial processing pipelines are being added. Data can be stored temporarily for up to three months. Additional computation will be available on campus through other HTC resources (e.g. condor).

3.8.4 Process of Science

The instruments in the facility are used to collect imaging data for biological samples ranging from macromolecular complexes to intact cells. The resulting data is used to study fundamental aspects of biological and chemical systems and also for the design and development of drug and therapeutic targets.

A typical workflow is as follows:

1. Prepare samples for imaging

2. Perform imaging
 - a. Data is stored on a local storage server.
 - b. Pre-processing is done on local computation.
 - c. Re-imaging is done as required.
3. Advanced analysis takes place on either HTC or HPC resources, including a HTCondor cluster.. When the workload is heavy, other locations on campus can be used, as well as national computing resources and cloud resources.
4. The data is archived outside of CEMRC using campus resources. For visitors or remote users must transfer data off-site using software or removable media.

3.8.5 Remote Science Activities

There are potentially several remote approaches for this science use case:

- Local staff may export data to regional or national facilities for storage or computation.
- Other regional or national facilities may export reference data to CEMRC.
- Visitors to CEMRC can export their collected data to their home institutions.
- Local staff running samples for remote collaborators may exporting locally collected data to the home institution of the collaborator

These use cases will prove out more as the facility operates and expands. To address these needs, it is expected that different network profiles (e.g. controls network, data movement network, remote access network) will be available.

3.8.6 Software Infrastructure

The University of Wisconsin-Madison CryoEM facility uses both commercial and open source software packages for cryo-EM image processing. These include:

- REgularised Likelihood OptimisatioN, pronounce rely-on (***Relion***¹²) is a stand-alone computer program that employs an empirical Bayesian approach to refinement of (multiple) 3D reconstructions or 2D class averages in cryo-EM.
- ***EMAN/EMAN2***¹³: EMAN2 is a broadly based greyscale scientific image processing suite with a primary focus on processing data from transmission electron microscopes.
- ***cryoSPARC***¹⁴: CryoSPARC is the state-of-the-art platform used globally for obtaining 3D structural information from single particle cryo-EM data.
- ***IMOD***¹⁵: IMOD is a set of image processing, modeling, and display programs used for tomographic reconstruction and for 3D reconstruction of EM serial sections and optical sections.

¹² https://www3.mrc-lmb.cam.ac.uk/relion/index.php?title=Main_Page

¹³ <https://blake.bcm.edu/emanwiki/EMAN2>

¹⁴ <https://cryosparc.com>

¹⁵ <https://bio3d.colorado.edu/imod/>

- **Amira**¹⁶: Thermo Scientific Amira Software is a powerful, multifaceted 3D/4D+ platform for visualizing, manipulating, and understanding life science research data from many image modalities, including CT, MRI, 3D Microscopy, and other techniques.
- **ImageJ**¹⁷/**Fiji**¹⁸: ImageJ is a public domain Java image processing program. Fiji is an image processing package — a "batteries-included" distribution of ImageJ, bundling many plugins which facilitate scientific image analysis.

Data curation is a part of the workflow, and being adapted to explore the use of an institutional storage solution that utilizes Globus.

3.8.7 Network and Data Architecture

Please see [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#)

3.8.8 Cloud Services

Cloud resources (compute and storage) are known to the cryo-EM community, but to date these have not yet been used. Local users are expected to leverage institutional storage and compute resources primarily. It is possible that remote users could leverage cloud resources.

3.8.9 Known Resource Constraints

Cryo-EM technology is advancing rapidly and the data volumes produced by the instruments increase with each new iteration of hardware. At this time it is not yet known how the usage profile will change over time. It is expected that local use will dominate at first, with remote use cases (located on site or controlling remotely) increasing in the future.

Data storage and mobility are key concerns. It is expected that local storage will only be able to keep up with a short (e.g. < 3 months) time horizon and that users must migrate data off of resources immediately. The use of data transfer tools are being explored to facilitate this, but it is also expected that portable storage devices will be widely used.

3.8.10 Outstanding Issues

The facility is still making choices regarding expansion possibilities (e.g. instruments must be kept in a controlled environment, and thus there is a desire to pack more into space that is specifically designed). More instruments will imply a greater need for institutional computing availability, and network connectivity. When the facility accepts outside users (on site or remote), the number of users will increase.

¹⁶ <https://www.fei.com/software/amira/>

¹⁷ <https://imagej.nih.gov/ij/download.html>

¹⁸ <https://fiji.sc>

3.9 Space Science and Engineering Center (SSEC) Case Study

Content in this section authored by John Lalande and Jerry Robaidek from the Space Science and Engineering Center (SSEC)

The Space Science and Engineering Center (SSEC) is an internationally known research and development center at the University of Wisconsin-Madison. With a history of remote sensing innovation spanning more than 50 years, SSEC develops and utilizes space-, aircraft- and ground-based instrumentation to collect and analyze observations of the Earth's atmosphere, oceans, and land surface, as well as other planetary atmospheres to improve understanding of weather, climate, and atmospheric processes. SSEC is a world leader in developing the algorithms and designing the ground and archive systems necessary to process atmospheric data collected from geostationary and polar-orbiting platforms.

Housed within SSEC is the Cooperative Institute for Meteorological Satellite Studies (CIMSS), a world-renowned satellite meteorology research center, and also the SSEC Data Center, the world's largest online, geostationary weather satellite data archive. The SSEC Data Center provides high quality, geophysical data to researchers, but also to industries that are affected by weather – from agriculture to energy to aviation.

Research within the SSEC ranges from the study of new instrument technologies to data analysis, visualization, and product development. Numerous satellite data analysis algorithms developed by CIMSS and SSEC scientists are now used operationally by agencies such as the National Weather Service. SSEC scientists are committed to sharing their efforts, tools, and knowledge with the global research community.

3.9.1 Science Background

The SSEC is a diversified research and development center focusing on geophysical research and technology to enhance understanding of the atmosphere of Earth and other planets in the Solar system.

There are multiple projects and groups in SSEC that rely on high speed networking and large data sharing, including the SSEC Satellite Data Services (SDS), the NASA Visible Infrared Imaging Radiometer Suite (VIIRS), Atmosphere Science Investigator Processing System (SIPS), and the Supercomputer for Satellite Simulations and Data Assimilation Studies (S4).

3.9.1.1 SDS¹⁹

Led by SDS program manager Jerrold Robaidek, the SSEC Satellite Data Services ingests and distributes geostationary and low earth orbiting satellite data. This data is used for a wide variety of science projects at SSEC including the calibration and

¹⁹ <https://www.ssec.wisc.edu/datacenter/>

validation of new satellite instruments and also algorithm development for derived products to better understand cloud top heights, winds, and convective initiation.

3.9.1.2 SIPS²⁰

Science Investigator-led Processing Systems (SIPS) create NASA science data products using algorithms and software developed by the Suomi National Polar-orbiting Partnership (NPP) Science Team and deliver the products to a NASA Distributed Active Archive Centers (DAACs) for archive and distribution. NASA's Atmosphere SIPS contract was awarded to UW-SSEC PI Liam Gumley in 2014. Atmospher SIPS is one of five NASA projects to create a continuous Earth System Data Record (ESDR).

The Atmosphere SIPS processes data from the Suomi NPP satellite to produce VIIRS Level 2 cloud and aerosol products. The products generated by the Atmosphere SIPS are archived and distributed by the Level 1 and Atmospheric Archive and Distribution System (LAADS), a public data repository maintained by NASA and the Goddard Space Flight Center.

3.9.1.3 S4

The Supercomputer for Satellite Simulations and Data Assimilation Studies (S4) is a collaboration between the NOAA National Environmental Satellite, Data, and Information Service (NESDIS), NOAA Center for Satellite Applications and Research (STAR), and the UW-SSEC. Scott Nolin is the UW-SSEC PI.

S4 is an R&D resource open to UW-SSEC and NOAA scientists and their collaborators, and aims at maximizing the return on investment for projects related to satellite data assimilation. The S4 system supports two major activities:

1. Global regional scale satellite data assimilation experiments, including the assessment of their expected impact on forecast model accuracy, using current or future satellite sensors and allowing scientists to test new science/methodology, and
2. Conducting Observing System Simulation Experiments (OSSEs) for new sensors.

S4 serves multiple projects and researchers and data handling varies by project.

3.9.2 Collaborators

The SSEC has a large number of collaborators. S4 alone has approximately 100 outside collaborators. Notable collaborators include:

- NOAA (Maryland/DC)
- NASA (Maryland/Huntsville, AL)
- Universities (hundreds of Universities)

²⁰ <https://sips.ssec.wisc.edu>

- International organizations (EUMETSAT, ABoM, KKMA) Germany, Australia, South Korea, Japan
- For SDS, domestic and international commercial users

3.9.3 Instruments and Facilities

3.9.3.1 SDS

SDS uses satellite ingestor systems at UW-SSEC (including SSEC developed satellite ingestors), satellite rebroadcast, and terrestrial internet for input data. Linux servers and Lustre file systems located in the on-premise SSEC Data Center, are used for processing, storage, and distribution. The Linux servers typically connect using 10Gb Ethernet and Lustre file systems are connected via 40 or 56Gb infiniband.

SDS ingests more than 5 TBs/day and distributes more than 2 TBs per day. The SDS satellite data archive currently hold approximately 1.5PB of data and is expected to grow by 350 TB per year for the next 2-3 years, then at about 500 TB/year for the following 2-3 years.

3.9.3.2 SIPS

SIPS relies on Linux clusters and Ceph filesystems, including approximately 100 compute nodes and 50 Ceph storage hosts that provide 7.7PB of storage. Compute nodes are connected primarily with dual 1Gb Ethernet cards and the Ceph nodes support dual 10Gb Ethernet connections.

SIPS ingests over 600GB/day of satellite data and generates 1,250GB/day of products for distribution.

3.9.3.3 S4

The S4 system is a Linux cluster with Lustre storage that was updated in late 2018 to provide:

- 2,560 Intel Xeon Gold processors
- 3.2 PB of primary storage
- 1PB of scratch storage
- FDR Infiniband

The S4system also includes nodes from the previous system, which include:

- 1,600 Intel E5-2680v2 processors
- FDR-10 Infiniband

There is a likely upgrade within the next year to double the capacity of the system, similar to the 2018 update.

3.9.4 Process of Science

3.9.4.1 SDS

Data sets are ingested by SDS in multiple levels of processing including Level 0 (raw instrument data), Level 1 (calibrated and navigated data), Level-2 (derived products), and in multiple formats (NetCDF, GRIB, BUFR, etc.) Datasets are used in both near-real time and retrospective modes. Data are partitioned in subsets geographically, spatially, spectrally, and temporally.

The Datasets originate from multiple sources, including in situ instrument measurements, remotely sensed instruments, and computer model output. Remotely sensed observations will be used in a number of different ways. Raw data is sent to a ground station, calibrated, and navigated, and then relayed to a satellite for redistribution via direct broadcast. That data can also be received from various terrestrial routes. Some data come directly from international partners including the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), the Australian Bureau of Meteorology, the Korean Meteorological Administration, etc.; other data come from domestic partners including NOAA (National Oceanic and Atmospheric Administration), NASA (National Aeronautic and Space Administration), and NSF (National Science Foundation)-funded University programs.

3.9.4.2 SIPS

SIPS products are produced with a workflow system developed by the SIPS team that allows processing on the SSEC SIPS cluster, campus HTcondor resources, or Amazon cloud resources (experimental).

3.9.4.3 S4

S4 users have a particular focus on numerical weather prediction models such as the Finite Volume Cubed-Sphere dynamical core (FV3), Global Forecasting System (GFS), and Weather Research and Forecasting (WRF). The research for S4 is focused on improving these models via the assimilation of satellite data. Since it is research for multiple groups, and not operations-focused, there is not a single, formalized workflow as for the SIPS, but individual research groups develop their own workflows..

3.9.5 Remote Science Activities

Remote data both incoming and outgoing from SSEC is a core feature of SDS. While a large amount of data is ingested directly, for even that directly ingested data, network connections are key. The data is monitored for quality and, if needed, sourced from other sites via the internet to ensure that SDS provides and archives the highest quality data. It currently outperforms many official sources.

SIPS data products are transferred to NASA's Level-1 and Atmosphere Archive and Distribution System (LAADS) Distributed Active Archive Center (DAAC), part of the Terrestrial Information Systems Laboratory at NASA's Goddard Space Flight Center in Greenbelt, MD. The reliability and bandwidth of the connection to NASA GSFC is critical for the success of the SIPS project.

Most S4 users are remote collaborators that move input data and results between multiple institutions. Data connections to NOAA in Virginia are important, as they are the primary partner. S4 researchers are using the NOAA/NESDIS supercomputer, Theia, in Fairmont, West Virginia, as well as NASA and university resources in other locations. There has been an increase in collaborators via NOAA's partnership, and this is expected to continue since NOAA/NESDIS is focusing on growing their university partnerships.

3.9.6 Software Infrastructure

Data sets are shared using a broad suite of software, including:

- McIDAS Abstract Data Distribution Environment(ADDE)²¹
- Local Data Manager)(LDM²²/internet Data Distribution (IDD)²³
- RSYNC
- FTP
- SFTP, SCP
- Globus²⁴
- HTTP/HTTPS
- GRB fanout (GOES Rebroadcast data CADUs sent via TCP/IP stream)²⁵

Beyond this basic list, there is an array of software used for acquiring, analyzing, and visualizing data that is too wide to describe in this document. The SIPS project has the most specifically evolved in-house software, as a single purpose project. SDS has developed some software such as *McFetch*²⁶ specific for satellite data. Other examples of locally developed software packages include *CSPP*²⁷ and *CSPP-Geo*²⁸.

3.9.7 Network and Data Architecture

At SSEC we do use Globus and the Science DMZ. We expect to move more services such as the center's FTP server to the Science DMZ in the future.

²¹ <https://www.ssec.wisc.edu/mcidas/software>

²² <https://www.unidata.ucar.edu/software/ldm/>

²³ <https://www.unidata.ucar.edu/projects/idd/iddams.html>

²⁴ <https://www.globus.org>

²⁵ <https://www.goes-r.gov/users/grb.html>

²⁶ <https://mcfetch.ssec.wisc.edu/>

²⁷ <https://cimss.ssec.wisc.edu/cspp/>

²⁸ <http://cimss.ssec.wisc.edu/csppgeo/>

3.9.8 Cloud Services

SIPS has previously processed data in an Amazon cloud. This process identified network bandwidth limitations that recently have been improved by UW DoIT. SIPS may use Amazon resources for backlog or reprocessing in the future.

3.9.9 Outstanding Issues

SDS currently has a challenge with data access from Asia. There are often large latencies (e.g. 500ms or greater), which cause processing problems for ingesting data from some satellite data only available via internet from Asia, and for distributing to remote partners in a timely manner.

3.10 Plant Physiology and Computation-based Phenotyping Case Study

Content in this section authored by Edgar Spalding from the Department of Botany.

Edgar Spalding is a Professor of Botany at the University of Wisconsin-Madison. The goal of the Spalding lab is to better understand seedling growth and development. In particular, research examines light, gravity, temperature, and other variables that influence the vital processes that produce an independent seedling from a dormant embryo within a seed. The plant hormone auxin (indole-3-acetic acid) is a mobile coordinator of many of these processes. The project focuses on understanding the transport mechanism by using a rigorous thermodynamic framework and making electrophysiological measurements on isolated transporters and mutant plants lacking them.

Researchers in the Spalding lab employ methods from the fields of molecular biology, genetics, electrophysiology, and digital image analysis. The model species are primarily *Arabidopsis thaliana* and maize, the key to leveraging the genetic resources built up around these model systems is to increase the measurement throughput, so the lab is developing computational tools to automate measurements of phenotypes.

3.10.1 Science Background

Plant growth and development are measured by analyzing time lapse images, In addition to these measurements, the research involves developing image analysis algorithms and performing experiments. Data sources vary broadly, but include drones flying over corn fields to banks of microscopes.

3.10.2 Collaborators

Approximately 20 research groups working at universities across the U.S. use analysis tools developed by the Spalding lab . Increasingly, these are being served using CyVerse data storage and computing resources in the style of a scientific gateway or web service.

3.10.3 Instruments and Facilities

The primary instruments in current use include cameras, scanners, and hyperspectral imaging devices. These can be mounted to drones, put on posts for field work, or used in laboratory growth chambers or benchtop arrays. It is increasingly common to use small autonomous vehicles carrying lidar sensors in field work.

3.10.4 Process of Science

Phenotype information is extracted from the images using custom algorithms. Many pipelines require the experimenter to transfer the image data to CyVerse disks for

analysis. Others transfer data to the storage servers for processing on to [The Center for High Throughput Computing \(CHTC\)](#).

3.10.5 Remote Science Activities

CyVerse (remote) resources for storage and computing are widely used, along with the use cases to acquire data on plants located in indoor growth facilities and in fields at the off-campus Wisconsin Crop Innovation Center, located in Middleton, WI, approximately 7 miles from the UW-Madison campus.

3.10.6 Software Infrastructure

- Matlab is the primary software component used to develop the image analysis tools.
- HTCondor is used to manage the distributed high throughput computing of the images.

3.10.7 Network and Data Architecture

Primary networking is described in [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#).

Basic connectivity has been reliable to date, but access to experimental pipelines in CyVerse, hosted at University of Arizona has been problematic in the past, resulting in very low transfer speeds and transfer failures between UW-Madison and the Cyverse iRODS server, hosted at University of Arizona . An extensive troubleshooting exercise was needed to assist in working out this network trouble between UW and the University of Arizona, which did result in faster and more reliable end to end transfer performance between the Wisconsin and Arizona sites.

3.10.8 Cloud Services

Commercial cloud services are not used, and are not under current consideration.

3.10.9 Known Resource Constraints

Network connection to ag fields and off-campus research facilities. In particular, the connectivity to the West Madison Agricultural Research Station was limited to the speeds of a cellular network. Image data storage will become a problem as sizes and quantities increase.

Due to occasional network transmission issues to CyVerse, the use of overnight, off-hours transfers has become more common.

3.10.10 Parent and Affiliated Organizational Cooperation

The Center for High Throughput Computing is the primary provider of local support.

3.10.11 Outstanding Issues

The intensity of data acquisition will increase as new types of field sensors and autonomous ground and flight vehicles that transport them are deployed, as described above.. The agricultural research centers, including the Wisconsin Crop Innovation Center, do not have sufficient outdoor wireless network capabilities to enable the major changes in field-based research that is taking place. The fields themselves have no wireless capacity but the 5 year future includes 'smart fields' and an explosion of remote data that must be addressed.

Currently, other campuses with research in this field are in the process of preparing or using "smart fields," i.e., agricultural fields equipped with wired and wireless networking equipment, IoT sensor devices, and edge/cloud computing resources that support data collection, movement, and analysis needed for near-real time monitoring of crop plant physiology, soil and pest conditions, and other parameters that affect crop yield. . Currently, UW Madison does not have adequate network capabilities at most agricultural field stations to support this approach.

3.11 Computational Materials Case Study

Content in this section authored by Dane Morgan from the Computational Materials Group.

The Computational Materials Group (CMG) at UW-Madison uses atomic scale modeling to understand and design new materials. The group employs highly accurate *ab initio* (first-principles) techniques to study electronic structure and energetics of smaller systems, as well as interatomic potential modeling to study up to hundreds of millions of atoms. These core approaches are combined with a wide-range of other atomistic methods, including Monte Carlo, coarse graining, data mining, thermodynamics, and statistical physics. These tools allow deep and quantitative view of materials phenomena over an extensive range of time and length scales.

3.11.1 Science Background

This research is focused on atomistic simulations using molecular dynamics and *ab initio* quantum mechanical methods. There are also emerging use cases that use machine learning algorithms to fit deep learning models to images of materials. The use of electron microscopy is also common for characterizing materials, which generates large data sets that need to be shared with colleagues, for example, atom probe data sets, that can be several TBS in size, making forin data mobility challenges between UW-Madison and external collaborators.

3.11.2 Collaborators

The CMG group works with faculty and staff around the country and sometimes the world. Data is shared using cloud resources such as Box, Google Drive, or Dropbox, as well as email. Large image data sets are shared by the Materials Data Facility, hosted at the University of Chicago, using Globus GridFTP.

3.11.3 Instruments and Facilities

The research involves regular data movement to hosting sites like Figshare²⁹, Materials Data Facility³⁰, and National Institute of Standards and Technology³¹. Data set sizes range from 100s of MB to a few GB. Computation is done using local resources, as well as allocations on XSEDE for the molecular dynamics and quantum simulations.

3.11.4 Process of Science

A typical workflow consists of:

1. Pull data from web sites with materials databases, e.g., the Materials Data Facility.
2. Run calculations on computers at UW or national facility.

²⁹ <https://figshare.com>

³⁰ <https://materialsdatafacility.org>

³¹ <https://materialsdata.nist.gov/>

3. Analyze output using software on those computers. Share data and analysis with collaborators, sometimes off campus.
4. Share data/codes with the public as part of papers being published.

In the future, the team would like to be able to share codes, simulation tools, and fitted machine learning models in a form that can be easily used, ideally by a specific API.

3.11.5 Remote Science Activities

While not regular, there are often requirements to move data to/from characterization centers, such as the Advanced Photon Source at ANL or other similar light source facilities.

3.11.6 Software Infrastructure

The Materials Data Facility³² is used widely for this work, along with other tools such as Globus³³ Grid FTP for data resources.

3.11.7 Network and Data Architecture

The primary computational resources used are personal computer clusters connected by infiniband. Other aspects of the data architecture can be found in [Section 3.1 The University of Wisconsin-Madison Campus Case Study](#).

3.11.8 Cloud Services

Commercial cloud usage varies. There are limited use cases:

- There is an effort to explore the use of AWS, but this is experimental
- Google cloud CoLab has been, and will continue to be, used.
- Cloud storage (Box, Google Drive, Dropbox) for file sharing and backup (Backblaze) is common

R&E cloud usage is limited to Open Science Grid and XSEDE.

3.11.9 Known Resource Constraints and Outstanding Issues

We are struggling to find easy access to GPU computing. The campus has limited resources available currently (e.g. the “Euler cluster”), but CMG would benefit from stable and flexible support for standard GPUs.

An easy to use, and inexpensive, solution for data backup is also needed. An estimate of the current total data volume produced by the center is in the 10-20TB range yearly, with something that is scalable in terms of numbers of machines that can be integrated, along with additional storage.

³² <https://materialsdatafacility.org>

³³ <https://www.globus.org>

4 Discussion Summary

On June 17-19 2019, members of the EPOC team and staff from the University of Wisconsin–Madison Division of Information Technology (DoIT) met with campus researchers. This review was held in Madison, WI, to help characterize the requirements for a number of research and educational activities, and to enable cyberinfrastructure support staff to better understand the needs of the researchers they support.

As part of the overall review, the necessity of working more closely with research teams as they are still in the planning phase was discussed so that DoIT can better understand and adapt to changes in requirements as the research demands grow over time. Additional challenges with securing sensitive data, cybersecurity, and supporting collaborations were also discussed.

During the discussion, the following points (outside of clarifications to the Case Studies described in Section [3 University of Wisconsin-Madison Case Studies](#)) were emphasized:

- The Science DMZ infrastructure (constructed as a part of an NSF award to help support upgrading the campus network in 2012) is not easily accessible by all campus users. A physical configuration has to be made to enable to use the Science DMZ, which decreases its availability to the campus.
 - A desired outcome is to on-boarding additional research labs for the DMZ in an easier way.
 - The new datacenter infrastructure (e.g. OneNeck) will use the Science DMZ.
 - Several use cases profiled here (e.g. Icecube, SSEC) will be migrated onto the DMZ in the future.
- DoIT desires a more rich set of network monitoring data from the campus infrastructure. This can be accomplished with sFlow/NetFlow at all layers.
- DoIT will explore operating Data Transfer Node (DTN) infrastructure for all campus users in part to better understand how a Globus endpoint could operate in a campus environment.
- High Performance Computing (HPC) use cases and needs are not as well understood as that High Throughput Computing (HTC) use cases. Users that require HPC pursue allocations on NCSA resources, or deploy small clusters. DoIT wants to understand these use cases better to potentially consolidate services.

The following Case Studies were not discussed in person, but the text submitted by the researchers is also in the report for reference:

- [3.9 Plant Physiology and Computation-based Phenotyping Case Study](#)
- [3.10 Computational Materials Case Study](#)

4.1 The Center for High Throughput Computing (CHTC), Ahlquist Lab Virology Research, and Huisken Lab Microscopy Research

These three use cases were presented together, even though CHTC has broader impacts for campus at large. Representatives discussed the nature of the science, instruments, storage requirements, and common computational use cases and patterns.

The following represent top level findings from the discussion:

- CHTC provides a significant amount of processing for campus use, but sometimes it is not enough to meet all needs.
 - Emerging and changing use cases such as Cryo-EM, electron microscopy work, and IceCube, require access to resources at more frequent rates than can be delivered under the current framework.
 - Additional computation is added as funding allows.
 - Open Science Grid resources that are external to UW can be used to meet these needs at times.
- Storage is a significant problem for instrument-based science, and was a concern for the Ahlquist and Huisken Lab research projects. There is a need to have a campus-based solution to handle growing data volumes as opposed to each group building their own solution.
- Internal network performance (e.g. transmission of data within the campus) can sometimes be a challenge. This was attributed to crossing security zones that were not designed for high performance network use. Moving research groups that have significant intra-campus data movement needs to the campus Science DMZ should result in performance increases.
- Data transfers to external sites (e.g. Howard Hughes Medical Center) have also been problematic at times. Investigation has shown that the problems are almost always on the “remote” end due to a variety of problems (insufficient network capacity, security devices, use of non-efficient software to manage data transfer). UW DoIT makes every attempt to assist, but will require assistance from the remote side to improve the use case.

4.2 Tier-2 Computing Center for the Compact Muon Solenoid (CMS) Experiment at the Large Hadron Collider (LHC)

Use cases affiliated with LHC are often highly productive with regards to computational and networking needs. The Wisconsin CMS Tier2 Center fits this pattern, and is a heavy user of campus resources for scientific innovation. There are several points of discussion that came from this case study, many of which are focused on building for the future:

- Currently the Physics group supports their own storage and computation resources, but relies on UW DoIT for networking resources.
 - The network infrastructure consists of 1Gbps connections to individual compute resources and 10Gbps for outward facing connections.
 - There will be upgrades to make this 10Gbps for computation resources and 40Gbps/100Gbps for the outward facing connections.
- It is expected that data volumes will increase after the luminosity upgrades to the LHC (see case study for schedule).
 - The data volume increase will be gradual, in the form of simulation sizes increasing in preparation for production traffic.
 - Re-processing campaigns using old data will also take place as the time for production comes closer.
- Simulation work will continue (as it always has) through the LHC shutdown, and then ramp up significantly with the new data size metrics.
- The estimates presented in this case study with regards to data volumes are consistent with WLCG estimates and are expected to track closely to forecast.
- UW Campus migration from 100Gbps WAN connectivity to 400Gbps WAN connectivity is expected, with Physics is a primary use case to justify this need.

4.3 The Great Lakes Bioenergy Research Center (GLBRC)

GLBRC receives funding from several sources (NSF, DOE, and campus) and thus is able to make investments in computation, storage, and networking regularly.

Discussion on this case study found:

- A known performance issue between an affiliated site at Michigan State University (MSU) was discussed. EPOC has opened a ticket to work with GLBRC, UW DoIT, MERIT (regional network for Michigan), and the MSU campus to find a solution.
- Remote connectivity to field sites is a known problem without a great solution. UW DoIT is working to secure then use of cellular spectrum, Wireless Internet Service Providers (WISPs), microwave links, and fiber connectivity to fields that are located in Wisconsin and Michigan. The goal is to address data movement issues between field sensors, drones, and the GLBRC data stores on campus.
- GLBRC is not currently on the UW Science DMZ, but would like to be. Moving the GLBRC to the DMZ would also result in additional use of the CHTC resources for computation and enable easier data sharing to remote sites.
- GLBRC wants to explore migrating some of the workloads they do toward containers (e.g. Docker), which would facilitate future use of cloud resources (RandE or commercial). This would provide some agility for their more bursty workflows.
- Using Globus GridFTP for wide area transfers is likely to increase the performance of several GLBRC use cases.

4.4 IceCube Neutrino Observatory

IceCube receives a majority of funding from the NSF, but also receives some support from the UW campus. Discussion produced several findings:

- Data volumes for IceCube do not fluctuate heavily: the detector takes in a predictable amount of raw data per cycle.
 - Interesting events that are processed in semi-real time (through the use of satellite transmission) do not fluctuate in terms of size, but frequency over time may increase.
 - Upgrades to the detector in the coming years will not increase data by orders of magnitude.
- Performance between NERSC and UW/IceCube can be improved. This could be related to IceCube resources not natively being on the UW Science DMZ
 - The upgrade to a new data center, and availability of the Science DMZ connection, should result in immediate performance improvements
- IceCube, via the NSF collaborations with IRNC, is interested in exploring more performant networks to the South Pole.

4.5 UW-Madison Cryo-Electron Microscopy Research Center (CEMRC)

CEMRC is relatively new to UW, and thus is still planning deployment of instruments and infrastructure. Some of the findings from this discussion include:

- CEMRC will have a minimum of 4 new instruments located in a minimum of 2 different campus buildings. All of the instruments are known to generate significant data sets up to a TB in size.
 - Use frequency is unknown - but expected to grow as external users utilize the resources.
- CEMRC will function as a 'facility', where visiting researchers come to utilize the instruments.
 - A robust framework will be required to support the data movement needs of the remote users.
- CEMRC and UW DoIT will work closely to design a network architecture that supports both local and remote compute and storage resources that are integrated with the instruments.
 - Data movement will be a key part of this work - potentially through shared endpoints like Globus
- Data access is critical, and there are a number of options being discussed:
 - Standard methods (FTP, SSH, SCP, RSYNC)
 - BYOSS (Bring Your Own Storage) or removable media
 - Globus endpoints
 - Data portals
- CHTC will assist in providing some compute on campus. It is also expected that local resources (CPU and GPU based) will be available.
- CEMRC will need Science DMZ access for external data movement.

4.6 Space Science and Engineering Center (SSEC)

SSEC has known use cases and usage patterns, but will see growth in the coming years as instruments become more precise. Findings from discussions include:

- SSEC is not currently on the UW Science DMZ, but should be. Resources located in the new data center will facilitate this.
 - During outages (in particular the data center move outage), UW DoIT will be consulted on ways to keep production traffic going.
- SSEC could benefit from a more uniform data movement strategy that included portals and Globus tools.
- Commercial cloud computing is desirable, particularly since several federal partners (NOAA, NASA) are moving this direction.
- UW DoIT will continue to work with SSEC to upgrade pieces of the network architecture and to consult on planned improvements.
- SSEC has had data movement challenges in the past when transferring to international locations. UW DoIT has assisted, and will loop in EPOC as required.
- SSEC would like to move away from FTP/SFTP for some data transfer use cases, but external entities (such as NASA, NOAA) do not have a way to facilitate the use of tools like Globus.
- Data mobility into and out of the facility have not been reported in the past, but this may be due to the volumes expected, and the tools used. SSEC and UW DoIT are nonetheless expecting for data growth, and tool changes, that could expose friction for data movement tools. Future work will focus on instrumentation of the paths via the perfSONAR tools.

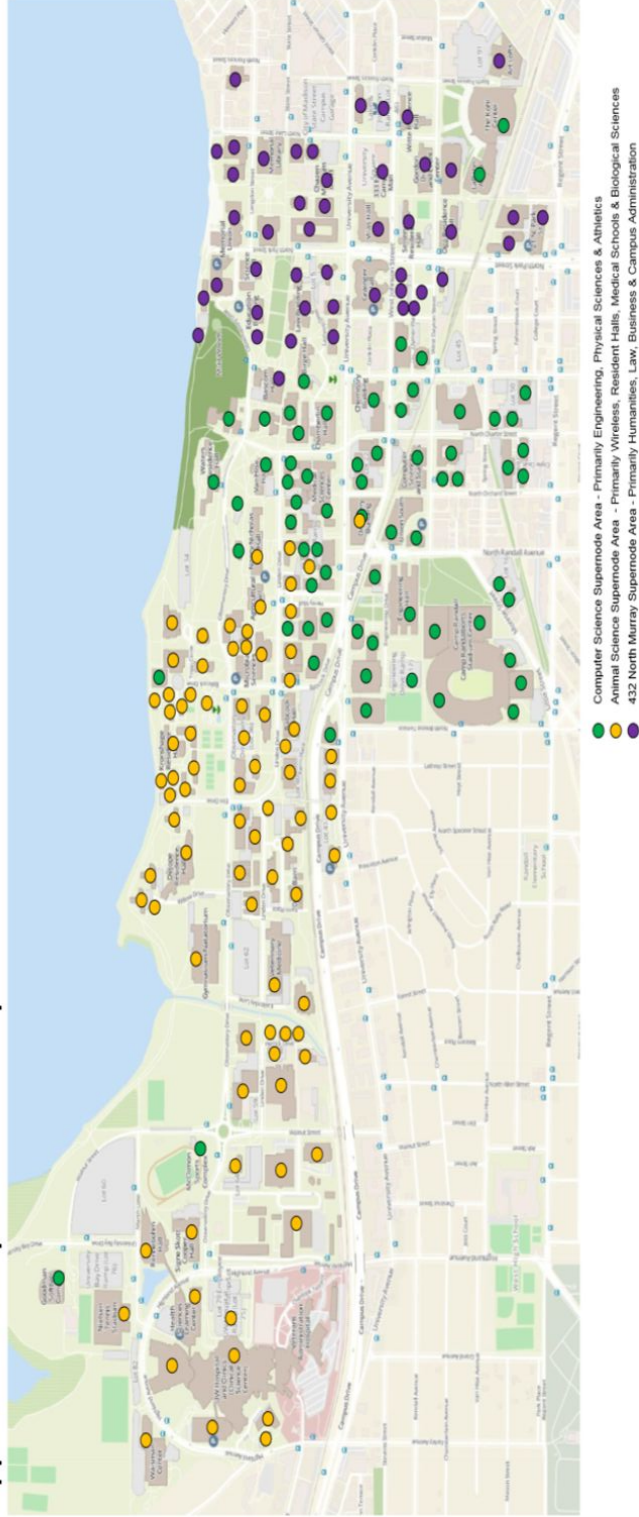
5 Action Items

EPOC and UW DoIT recorded a set of action items from the Deep Dive, continuing the ongoing support and collaboration. These are a reflection of the Case Study report and in person discussion.

1. UW DoIT will work to make the Science DMZ infrastructure more accessible for a greater number of use cases on and off campus (including at remote data centers). Use cases that can benefit include (but are not limited to):
 - a. GLBRC
 - b. IceCube
 - c. SSEC
 - d. CEMRC
2. UW DoIT will explore the use of sFlow/NetFlow data in more campus locations to understand traffic patterns.
3. UW DoIT will explore deployment of campus-wide DTNs with Globus endpoints as a service for the campus, and will work with groups such as GLBRC and SSEC to offload some data sets.
4. UW DoIT will continue to work with CHTC on ways to reduce friction for campus data movement.
5. UW DoIT will explore institution-wide storage options for campus users that goes beyond capabilities that exist today.
6. CHTC will continue to expand access to campus computing use cases such as those presented in the IceCube and Cryo-EM use cases.
7. UW DoIT will work with the Ahlquist and Huisken Labs to address data movement challenges to remote sites.
8. UW DoIT to work with CMS on network upgrades to support future LHC upgrades.
9. UW DoIT, GLBRC, and EPOC will work to address a wide area data movement problem to Michigan State University.
10. UW DoIT and GLBRC will continue to work on connectivity to remote field research.
11. UW DoIT will work with IceCube and NERSC to address performance abnormalities.
12. UW DoIT and CEMRC will address architectural issues including network, compute, and storage, as new instruments are installed and commissioned.

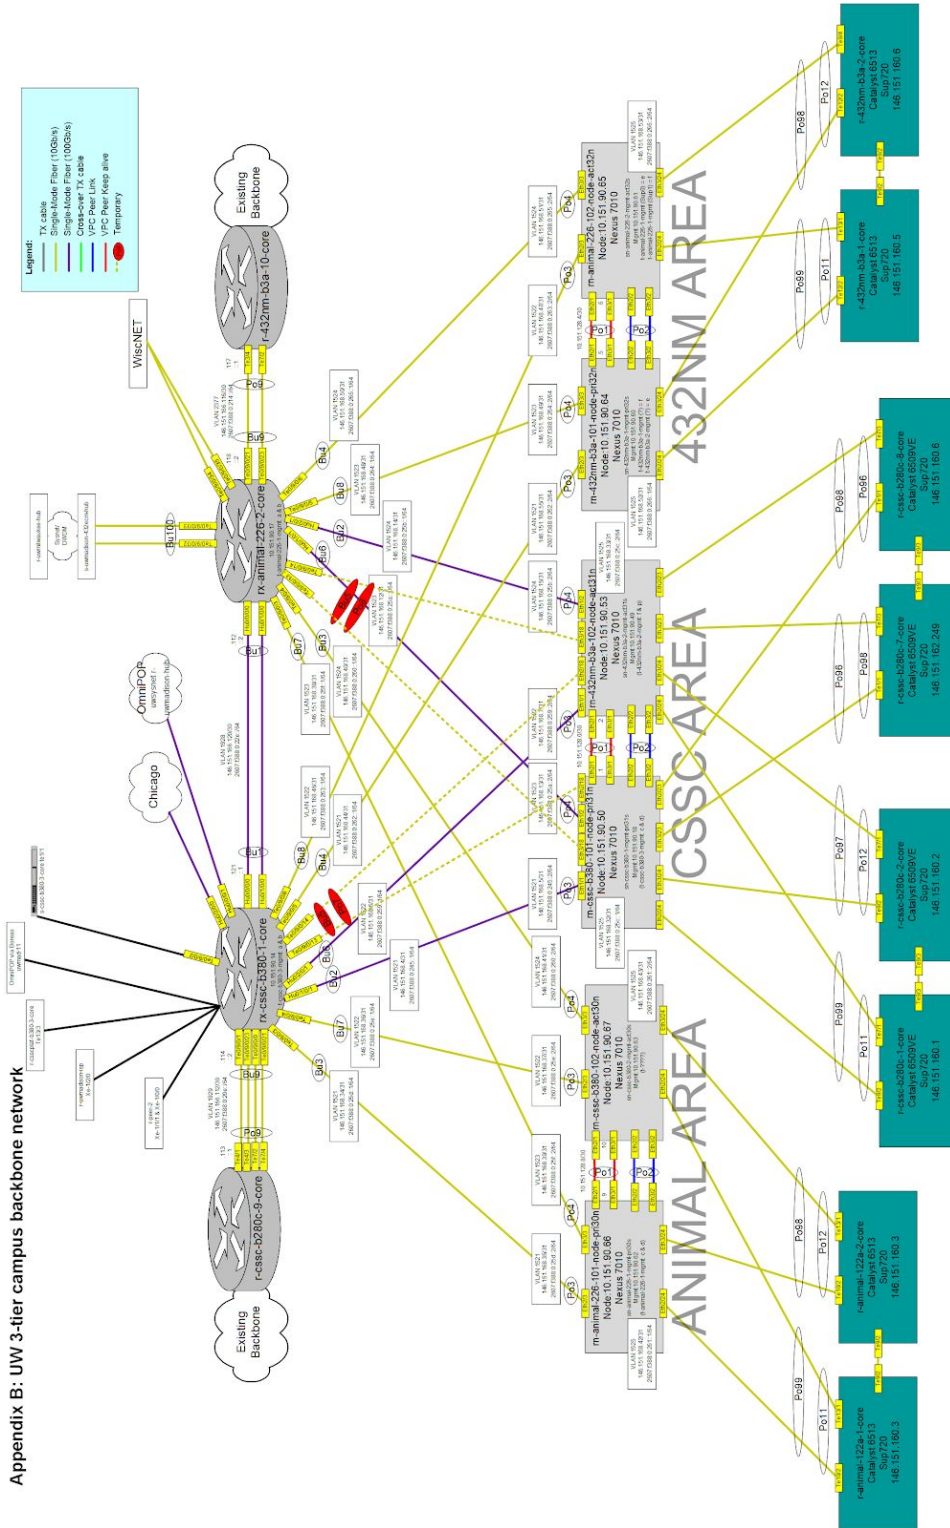
Appendix A - UW-Madison Campus Supernode Locations

Appendix A: Campus Network Supernode Areas

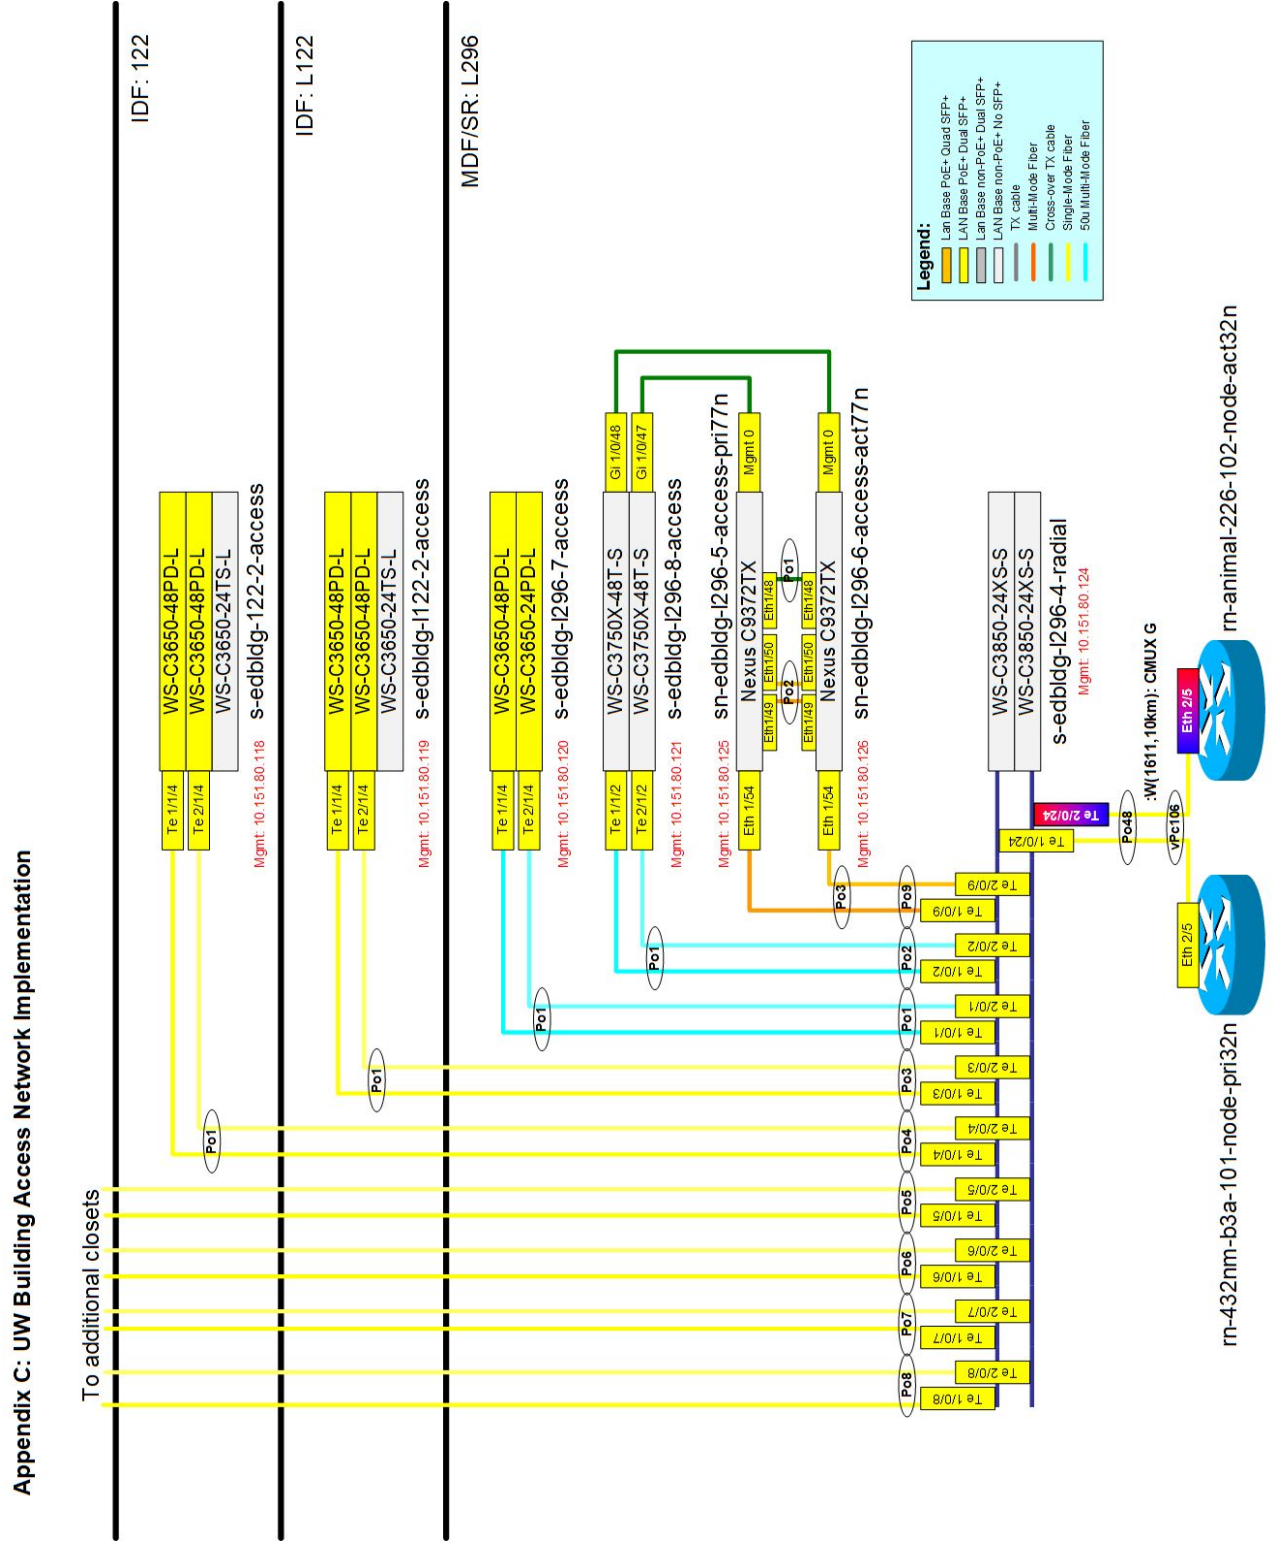


The UW campus network is divided into three Supernode areas, shown in yellow, purple or green. Each Supernode area consists of roughly 60 buildings and 1/3 of the campus.

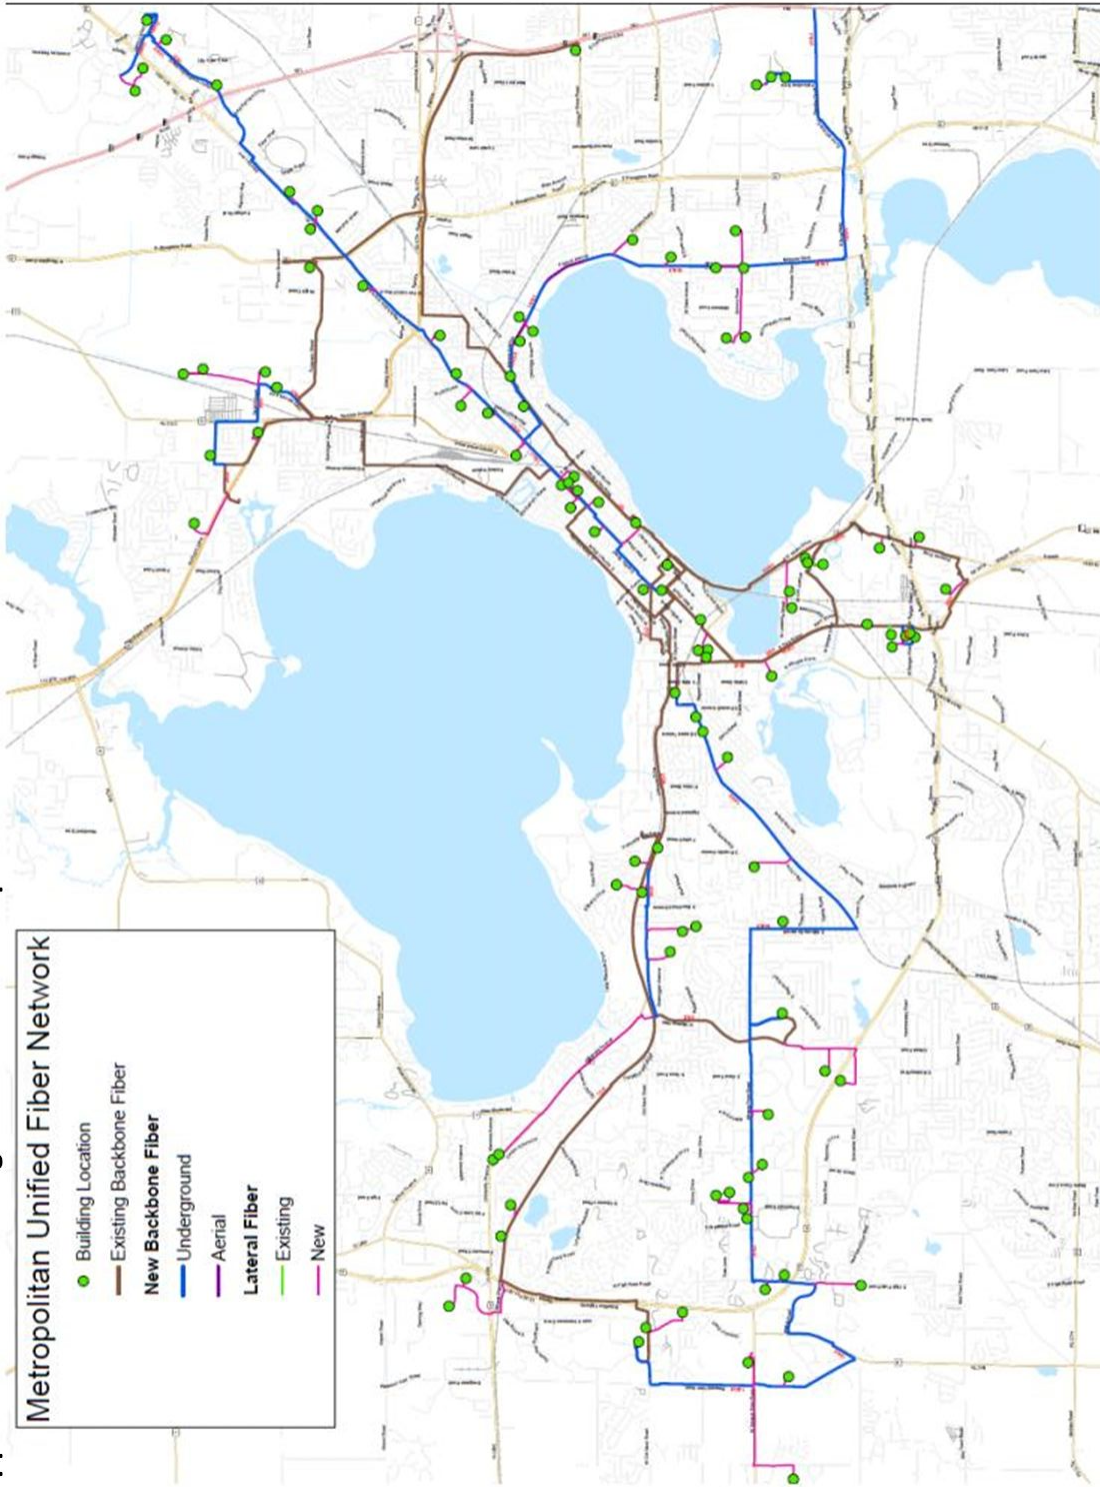
Appendix B - UW-Madison Campus Backbone Network



Appendix C - UW-Madison Building Access Network (Logical)



Appendix D - UW Building Access Network Implementation



Appendix D - Cyberinfrastructure Plan, University of Wisconsin-Madison, April 2018

OVERVIEW

UW-Madison's science and technology strategy is embodied in a core tenet of its mission: the "Wisconsin Idea," i.e., the creation of knowledge and innovations that benefit society. In accordance with this goal, research focuses on areas involved in global health, treatments for disease, new materials, technologies, and renewable energy sources, global food security, social and economic equality, and new understandings of matter and living systems. The Office the Vice Chancellor for Research and Graduate Education (VCRGE), the Wisconsin Alumni Research Foundation, and central administration provide support to researchers to help them develop new methods, build partnerships with the government and industry, and bring the products of their research to the relevant audiences and markets rapidly. Cyberinfrastructure (CI), a critical underpinning of all research activity, is provided through a combination of central, shared, and custom resources.

Research in the 21st century is data- and computation-intensive. UW-Madison's CI is designed to enable the computing power, network connectivity, storage, software applications, and security measures needed to support the world class research on campus and with our collaborators across the world. Our CI includes both physical hardware, technical resources, and human expertise to support the scientific work of researchers and enable more efficient approaches that maximize discovery and innovation. Research moves quickly and our vision for CI is one with the elasticity and modularity to build capacity where needed while also providing the flexibility to respond to new directions and approaches as they arise.

GOVERNANCE

UW-Madison's CI is based on close coordination between the offices of the VCRGE, the Vice Provost for Information Technology and CIO, the Center for High Throughput Computing (CHTC), and the Research Technology Advisory Group (RTAG). RTAG strategically aligns shared resources and campus investments to ensure that capabilities match the needs of a broad spectrum of research projects. Coordination and faculty governance is accomplished through monthly meetings of the RTAG, which is chaired by the CIO in the Office of the VCRGE.

The Vice Provost for Information Technology/CIO is responsible for networking, shared data centers, hosting of VMs, hosting of data, middleware, and campus-wide software licenses. As part of its responsibilities for core research services, the VCRGE coordinates investments in campus CI. The CHTC provides research operational and facilitator staffing and hosts the shared campus HTC and HPC compute resources. The Division of Information Technology (DoIT) operates the campus network, including the Science DMZ and manages a shared datacenter and cloud computing contracts.

RESEARCH COMPUTING

Advanced computational resources are vital to the work of UW-Madison researchers in all disciplines and to the discoveries needed to advance new energy alternatives, materials, and clinical therapies and to gain understanding of space, matter, and living processes. In many

cases, computational approaches have expanded the scale of the questions that can be asked and increased the need for more and more diverse computational resources. Machine learning, models that merge experimental data and simulations, high luminosity sensors, and cyber-physical sensor systems in agriculture and medicine are driving new requirements for edge, cloud, and on-site computational resources. Competitiveness of the UW-Madison in the research arena depends on a CI with diverse and widely accessible computing resources and professionals who support the infrastructure and help researchers use it effectively.

Shared computational resources are provided to UW-Madison's researchers and scholars through the CHTC, which procures, maintains and administers large-scale computing resources, providing over a million compute hours per day to researchers. Campus investments in 2016 added high memory capabilities, and additional accelerator resources are planned for 2018. The CHTC is the gateway to external computational resources on the Open Science Grid, which provided UW-Madison researchers with over 44 million compute hours in 2017. Departmental compute clusters, supported by extramural funds, provide compute and storage for researchers with specialized needs in Social Sciences, Mechanical Engineering, Physics, Space and Atmospheric Sciences, Biostatistics, and others.

Through the University's contracts with commercial cloud providers, researchers are able to take advantage of all Amazon Web Services and Microsoft Azure tools and agreements with Google Cloud are currently underway. An integration between HTCondor and AWS enables the CHTC to utilize spot pricing to scale up compute resources to meet the needs of OSG and campus researchers. The University is currently hiring a Director for Cloud Strategy to guide future integrations of research CI with the cloud.

NETWORK

The UW-Madison network consists of a redundant 100-Gb backbone and connections to all campus buildings that are being upgraded to at least 10-GB, with redundancy in each network node. An equitable funding model assures that network resources will be kept current into the future. UW-Madison has been fundamental to the establishment of the Broadband Optical Research Education and Science network (BOREAS) and UW System network. These regional optical networks interconnect UW System campuses together and to the CIC OmniPoP in Chicago, providing a high-speed gateway to various research networks, including Internet2's Advanced Layer 2 Service, ESNNet and other global research networks. An NSF Campus Cyberinfrastructure grant, awarded in 2012, enabled an upgrade of the campus wide area network to 100 Gb connectivity, as well as the frictionless architecture, security, and software defined networking capabilities that comprise UW-Madison's Science DMZ. PerfSonar nodes at the campus border enable monitoring and troubleshooting of latency and bandwidth issues encountered by network traffic traveling to and from campus. PerfSonar nodes deployed at points within departments provide metrics on connections between the campus backbone and local computing environments.

Wireless

The campus wireless network is ubiquitous in buildings but not yet available outside all buildings. Expansion of wireless connectivity in outdoor areas of campus is planned, in anticipation of tomorrow's Internet services and applications: a future likely to include thousands of smart devices, sensors, and applications utilizing cloud computing, with access to large data, and ubiquitous computing.

NSF funding has enabled installations that leverage the national GENI (Global Environment for Network Innovations) infrastructure and WiMAX and LTE technologies to provide connectivity in remote research sites and for high bandwidth applications such as mobile video streaming on campus. The WiNGS lab in the department of Computer Science explores mobile and wireless networking systems, including the Paradrop Wi-Fi routing system. The department of Electrical and Computer Engineering is home to research groups creating and testing new mmWave technologies and architectures for wireless communication and sensing.

The UW has active research in a number of data-intensive domains that rely on outdoor wireless communications, some on-campus and others at agricultural stations or urban areas outside of campus. These include precision agriculture, remote environmental sensing, smart energy grids, smart vehicles, and other areas of IoT research. The communication, computation, and quality of service needed to enable data collection, processing, and analysis in these activities, are creating needs for sustainable sensor networks and edge computing modalities that connect back to the campus network backbone.

Network standards and capabilities

The campus wired network is IPv6 enabled but the standard is not fully adopted across campus. The 802.11n IEEE WiFi standard for network throughput is fully deployed and is currently being upgraded to the 802.11ac Wave 2 standard. BPC 38 filtering to defeat IP spoofing is fully deployed across campus. In addition, campus provides a RPKI (Resource Public Key Infrastructure) service on top of the network.

Software Defined Networking (SDN)

SDN is being explored as an emerging network capability with applicability to scientific workflows that rely on moving data to computational resources, cloud computing, and high bandwidth wireless applications. NSF Campus Cyberinfrastructure funding has enabled experimentation with SDN as a means of circumventing network bottlenecks such as firewalls to allow faster and more efficient data movement across the computation network in the Science DMZ. In addition, SDN is being used as a method to allow applications and users to specify network connections in the NSF-funded CloudLab environment. SDN is also utilized in GENI architectures for gigabit wireless and mobile applications at the WiNGS lab.

SECURITY

The UW-Madison network is protected by active intrusion detection and mitigation systems and endpoint management tools. A UW-Madison risk management framework (RMF) aids researchers in determining the security risks to data in information systems and provides guidance on developing research systems with built-in security controls. The RMF is a key tool for researchers who must comply with federal regulations for protecting sensitive data, including HIPAA, FERPA, NIST 800-53 and NIST 800-171. In addition, a baseline departmental IT security policy provides guidelines for researchers and IT staff in academic departments for securing and monitoring workstations, servers, wired and wireless networks, and for building secure applications. UW-Madison is the home institution for the Software Assurance MarketPlace (SWAMP), a Homeland Security-funded project that develops software assurance tools available to all software developers.

UW-Madison Cybersecurity maintains a suite of tools to perform network security monitoring of the research and education network. At the most basic level, Netflow data is collected for

analysis and summarization of network traffic. For signature-based identification of threats, Suricata IDS and the Emerging Threats Pro threat intelligence feed from Proofpoint is used. Funding from a 2012 NSF CC NIE grant was used to integrate our Bro Network Security Monitor cluster to perform protocol level analysis and record metadata about activity on the Science DMZ. Data generated by these tools as well as logs generated by other systems such as authentication, VPN, firewalls, servers, etc., are forwarded to an enterprise log search and archive and Q Radar for further analysis, with plans to migrate to Elastic Stack as a logging platform in summer of 2018. A recent campus investment in systems for advanced monitoring, logging, and vulnerability scanning has expanded our capabilities and moved campus closer to establishment of a sustainable, formalized cyber security operations center.

MIDDLEWARE

Federated Web Access Control: The University has long been a member of InCommon, the US higher education and research identity federation. InCommon-supported connections allow University faculty, staff and students to access scores of service providers including research-related providers such as the NSF, NIH and the Indiana CTSI.

CI Logon <http://www.cilogon.org/> is a federated cyberinfrastructure access gateway by which UW faculty, staff and students can access significant national computing resources via a simple NetID login. As a member of InCommon, we are participating in the exploration of multi-factor authentication as a means to support higher levels of assurance. That in turn will provide a CI Logon path to more restricted, higher value national cyberinfrastructure resources.

DoIT Middleware and Enterprise Architects are also working with our Big Ten Academic Alliance colleagues to develop tools and procedures to ease the pain of integration that comes with the rapidly growing use of cloud-based services.

Global Research & Scholarship: The University participates in multiple national and international identity federations, including InCommon (www.incommon.org) and EduGAIN (www.edugain.org) to enable researchers to access data and resources across institutional boundaries. The University supports release of data to the Global Research & Scholarship Entity Category (<https://refeds.org/category/research-and-scholarship>), enabling our researchers to access resources from participating institutions around the world.

REFEDS Engagement: DoIT Middleware and Enterprise Architects are engaged in the Research and Education Federation Federations Group (REFEDS), which defines standards for interoperability and resource sharing across more than 40 higher education identity federations worldwide. REFEDS engagement enables the University to establish technology, infrastructure and policy that enables data and resource sharing with a global network of identity federations.

Federated Network Access: The University is connected to the U.S. eduroam www.eduroam.org wireless network access service and has been a production identity provider (IdP) since August 2011. This has allowed our researchers to easily access network resources at thousands of participating eduroam institutions around the world. Conversely, academic and other visitors to the University can access our networks with a simple login to their home institution.

Manifest Service: Web-based, delegated group and privilege management services are now available across campus for the purpose of controlling access to compute resources and data, as well as to foster communication and collaboration. This is a key support tool for academic, research and administration units across campus.

Internet2 TIER Initiative: UW-Madison is playing a prominent role within the group of institutions of higher education that have come together to make radical improvements in identity and access management (IAM) middleware for higher education and research. Modernized architectures, reference implementations and deployment guides for all aspects of IAM are being developed and packaged. One of the primary goals is to guarantee the long-term sustainability of key open source components of IAM by enlarging the community of people using and supporting them. Another is to increase local IT developer resources available for higher-value projects by freeing them from having to be devoted to building and maintaining home-grown IAM components.

SOFTWARE

The Division of Information Technology acquires bulk licenses for several popular research programs, including MatLab, Maple, SAS, and SigmaPlot, and makes them available to researchers at no cost through a campus software library. UW-Madison also has site licenses for several cloud-based collaboration and file sharing platforms with unlimited storage capacity, including Google Drive, Box, and LabArchives.

DATA

Storage

UW-Madison has traditionally operated numerous campus data centers and server rooms dedicated to maintaining local information systems. In 2012, an Administrative Excellence team conducted an assessment of security, power and backup capabilities and recommended aggregation of centers providing redundant services to improve efficiency, reduce costs and provide affordable, streamlined services in a Campus Computing Infrastructure (CCI). CCI is a campus sponsored and governed initiative delivering a portfolio of shared, scalable IT infrastructure services including virtual and physical server hosting, storage, and backup designed to meet the teaching, research, and service missions of UW-Madison. CCI is staffed by an engagement team that assists researchers in the selecting appropriate services, based on requirements. CCI currently administers over 750 virtual servers, 1.95 petabytes of production storage, and 1 petabyte of production tape for backup and archive. Cloud offerings have recently been added to the CCI portfolio and include Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) based on Amazon Web Services and Microsoft Azure.

Data Centers

The data requirements of researchers are rapidly outpacing the on-campus space, power, and cooling resources. This has driven an intensive exploration and planning effort which has produced options for off-campus facilities, including Tier 3 data center facilities at the OneNeck commercial provider located approximately 10 miles from campus.

CYBERINFRASTRUCTURE STAFFING

The CHTC is staffed by 22 technical personnel and two Research Computing Facilitators who consult with researchers on computational approaches, collaborate with a variety of additional

intra- and inter-campus entities providing research CI support and services. DoIT staffs 2-3 network engineers who participate in grant-funded projects, building knowledge inside central IT and providing expert technical support to projects such as CloudLab and IceCUBE. IT staff in system administrator, software engineer, and network engineer roles provide support to large research projects within departments on campus. The UW-Madison's newly launched Data Science Hub serves as a focal point for education and consulting on management, storage, and analysis of complex data sets. A data science facilitator position to consult with researchers on advanced data analysis methods is currently being filled. The RSAG group (Research System Administrator Group), serves as a community forum for exchange of methods and expertise in areas such as file storage systems, data transfer methods, cloud computing, between IT staff in the CHTC, local departments and research groups, and DoIT. Externally, UW-Madison is a partner institution in the multi-campus Advanced Research Cyberinfrastructure - Research and Education Facilitators project (ACI-REF) and is a member of the Campus Research Computing (CaRC) Infrastructures Consortium.

EDUCATION AND TRAINING

Numerous opportunities for researchers to learn new computational methods, develop sound coding practices, and hone their data management and code reproducibility skills exist on the UW-Madison campus. These include credit courses such as "Introduction to R for Scientific Research," (CS 368); workshops on data management, offered by the Library-based Research Data Services, individualized research computing training for individual research groups and departments and multi day Software and Data Carpentry workshops. UW-Madison's research computing facilitators are trained instructors for both Software and Data Carpentry and contribute to curriculum development for both programs. On the UW campus, the community of campus instructors has grown and Software and Data Carpentry workshops for 40 researchers each are now provided 3 times per year.

Appendix E - UW-System Regional Networking Diagram

