

Boston University Cyberinfrastructure Plan

I. BU Research Infrastructure

o THE MASSACHUSETTS GREEN HIGH-PERFORMANCE COMPUTING CENTER (MGHPCC)

The Massachusetts Green High Performance Computing Center (MGHPCC), established and operated by BU, Harvard, MIT, Northeastern, and the University of Massachusetts—five of the most research-intensive universities in Massachusetts—is dedicated to serving our growing research computing needs more collaboratively, efficiently, ecologically, and cost-effectively. BU high-performance computational resources are housed at the MGHPCC, taking advantage of an abundant source of clean renewable energy from the Holyoke Gas and Electric hydroelectric power plant on the Connecticut River. This project created a world-class high-performance computing center with an emphasis on green sustainable computing, offering unprecedented opportunities for collaboration among research, government, and industry in Massachusetts. MGHPCC is the first academic research computing center to achieve LEED Platinum certification, the highest level of green certification awarded by the Green Building Council.



MGHPCC comprises 33,000 SF of space optimized for high performance computing, a 19 MW power feed, and high-efficiency cooling to support up to 10 MW of load with provisions for expansion to 30 MW. The MGHPCC owns the 8.6 acre site, providing substantial room for expansion. Network infrastructure includes a fiber loop passing through Boston, the Northern Crossroads (NoX), and New York City, connecting BU to the MGHPCC at 40 Gb/s and providing Internet2 and other regional, national, and international networking connectivity at speeds up to 100 Gb/s from the MGHPCC. BU is also actively engaged in other national and regional consortia where we provide technical, strategic, and governance leadership.

o RESEARCH COMPUTING

Information Services & Technology (IS&T) Research Computing provides specialized hardware, software, training, and consulting for all areas of computational research at BU. Research Computing resources support a wide range of disciplines, from the physical sciences and engineering to emerging computational communities such as biostatistics, bioinformatics, genomics, neuroscience, microbiology, machine-learning, public and global health, economics, finance, and the social sciences. Resources are managed in close consultation with BU's Research Computing Governance Committee, our Center for Computational Science, and the Rafik B. Hariri Institute for Computing and Computational Science & Engineering.

Our primary computational facility is a Linux-based Shared Computing Cluster (SCC) with approximately 18,500 CPU cores and nearly 6.7 PB file-system. Seventy-eight SCC nodes are equipped with NVIDIA GPU accelerator cards, for a total of 302 GPUs. The SCC comprises heterogeneous subsets of nodes interconnected by GigE, 10GigE,

40GbE, QDR, FDR and EDR InfiniBand. The SCC uses Globus Online to expedite bulk data transfers with unimpeded 10GigE and in select cases up to 100GE of access to Internet2 and other research networks.

SCC resources are available on both a fair-share, no-cost basis to all faculty-sponsored research groups and via buy-in funded by researchers for priority use. Excess buy-in capacity is returned to the pool for general use.

SCC resources are managed by Research Computing at no charge. A dedicated service model provides fee-based management for other research computing systems that don't align with shared or buy-in models. Finally, a co-location model offers stand-alone rack space, power, cooling, and network access to proprietary resources.

Research Computing supports approximately 2,700 individual researchers working on more than 700 different research projects supported by more than \$673M in external funding across our Charles River, Fenway, and Medical Campuses.

Research Computing also provides both individualized consulting and group training. More than 60 classes each semester covering research computing basics, programming, mathematics, data analysis, and scientific visualization. Several hundred researchers attend these sessions each term.

- COMMUNICATIONS INFRASTRUCTURE

The BU campus network provides high-speed access for over 200,000 devices to all institutional information, communication, and computational facilities, as well as the Internet, regionally aggregated resources, and advanced networks such as Internet2. Tens of thousands of wired ports and wireless access points are interconnected via optical fiber and a robust hierarchy of high-speed routers and switches. The University has three Internet routable Class B IPv4 netblocks and a /32 IPv6 netblock space assignment.

Our network was designed and deployed with flexibility, reliability and scalability as core tenets. We adhere to best-practice distribution standards, from secure stacked network closets, patch panels, and cabling, to pervasive IP spoof checking via Reverse-Path Forwarding validation on all segments and institutional IP address verification at borders that comply with the Internet Engineering Task Force Best Core Practices under IETF BCP38. We also maintain a comprehensive database that maps the network to floor plans spanning hundreds of buildings.

Wired and encrypted 802.1x wireless network access is available throughout campus, including all residence halls and classrooms. The University is in the middle of a 4-year refresh of its wireless infrastructure to begin upgrading from current pervasive 802.11n to a combination of 802.11ac wave 2 and 802.11ax compliant radios. BU is a member of the eduroam federation supporting ease of cross institutional use of wireless resources supporting both visiting and travelling research faculty, students, and staff on its CRC and Fenway campuses with service being added to its Medical Campus in 2020. Redundant connections to both the Internet and Internet2 operate at 10-20 Gb/s with automatic failover.

II. Planning

Cyberinfrastructure planning at BU is an integral part of a comprehensive institutional IT strategy developed by stakeholders within an overarching governance framework. BU leadership established formal IT governance several years ago as described at bu.edu/tech/about/governance/.

Institution-wide IT strategy was formulated and published at bu.edu/tech/plan/ in 2014.

Note that, even prior to establishing formal governance and an IT strategic plan, the University had an extraordinarily rich and enduring history of commitment to shared infrastructure in support of research and education as highlighted in **Exhibit A** below.

The cyberinfrastructure goals highlighted below include and augment the 2015-20 BU Technology Plan Initiatives.

- COMMUNICATIONS INFRASTRUCTURE
 - Ensure rapid and secure provisioning of new services by deploying industry standard Identity and Access Management (IAM) infrastructure, including federated identity.
 - Deploy Eduroam and IPv6 transport for our clients and public facing Internet services. Towards that end, we have a /32 IPv6 allocation from ARIN that is selectively enabled to meet requirements for both research and production oriented applications. This followed an internal technical working group that formulated [1] IPv6 address allocation strategy, [2] a core protocol deployment plan to utilize IS-IS routing, [3] an assessment of network component IPv6 readiness, and [4] a plan for implementation of an IPv4 to IPv6 gateway. Eduroam was initially deployed in June 2014 to selective portions of the campus with an expectation that it will be pervasively available throughout the institution in the summer of 2019.
 - Provide unified, high fidelity communication and collaboration services for one-to-one, one-to-many, and many-to-many voice and video communication that integrate messaging, scheduling, shared workspaces, and social networks while employing social networking to enhance online research communities.
- RESEARCH COMPUTING
 - Support data lifecycle, including the capability to securely store, protect, share, publish, archive, and sunset research data in accordance with policy and data management.
 - Develop a dynamic, collaborative pool of human resources for the development and support of computational software for research.
 - Ensure that shared research computing services fully support researchers on both the BU Charles River, Medical, and Fenway campuses, enabling both optimal use of resources and maximum collaboration.
- REGIONAL RESOURCES
 - Build on the Massachusetts Green High Performance Computing Center (MGHPCC) to provide sustainable, shared research computing, and storage infrastructure to researchers.

- Join peer institutions and industry in developing safe and effective means to leverage cloud infrastructure for research, including high performance computing and big data.

Exhibit A: Campus and Regional Resources

Campus Resources

o COMMUNICATIONS INFRASTRUCTURE

BU developed one of the largest open-access computing and communication services in education in the early 70s, supporting thousands of subscribers. We were among the first institutions to establish an IP-based campus network in the early 80s and connect it to NSFNET in partnership with regional research universities.

BU continues to be a leader in regional networking. With more than thirty years of experience beginning with our establishing BITNET connectivity for New England in the early 80s and subsequently founding NEARnet, the New England component of NSFnet, with Harvard and MIT. The three institutions later founded the Northern Crossroads (NoX) in 1999 and BU was the first institution to connect to the vBNS, providing New England with shared access to the earliest vestige of Internet2. BU maintains a seat on the NoX Executive Committee as a founding member. We also participate and collaborate in the NorthEast Region and Education Network (NEREN), and Eastern Regional Network (ERN) activities.

o RESEARCH COMPUTING

BU is a pioneer in centrally supported computational research and the deployment of leading edge shared scientific computing resources. We were one of the first universities to install IBM 3090-based vector processing, Thinking Machines Corporation CM2 and CM5 parallel computers, SGI Origin2000, IBM Power4 p690, and IBM Blue Gene/L. In 1989, the University created the *Center for Computational Science* for institutional scientific computing and visualization support, providing focus to drive computational science and advanced research computing. That group subsequently evolved into *Research Computing*.

o INFORMATION SECURITY

Information Security is available to assist campus researchers with meeting security and compliance requirements for their research, whether those requirements are based on standards like NIST, regulations like HIPAA, or other data use agreements. Information Security staff can help complete and compile risk assessments, contingency plans, privacy impact assessment and other supporting documentation and has templates to help researchers complete System Security Plans (SSP) or can provide resources to help define Data Management Plans. Since these documents often rely in part on security controls provided by the institution, our processes can help relieve research teams from much of the documentation burden.

Regional Resources

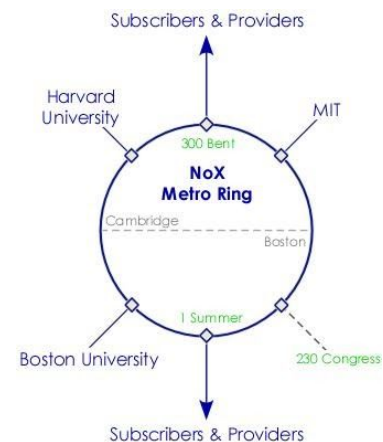
Boston University has a long history of collaboration and partnership with national computing centers and bridging local resources nationally. BU was a partner in the Metacenter Regional Affiliates program and with NCSA in the Partnerships for Advanced Computational Infrastructure. Today, BU participates in the XSEDE Champions program, facilitating the use of NSF resources by regional researchers. We are also working through XSEDE to bring other cyberinfrastructure campus bridging technologies to local and regional systems using Globus Online in a pilot project.

In collaboration with Harvard, BU runs the New England ATLAS Tier 2 Center, one of five such centers in the US and a participant in the Open Science Grid. ATLAS is one of two general purpose experiments running on the Large Hadron Collider at CERN involving over 3,000 physicists world-wide. The U.S. ATLAS Northeast Tier 2 (NET2) supports thousands of researchers with 10,000 hyperthreaded compute cores and approximately 10 PB of storage interconnected to the Internet2 and LHCone research networks at speeds of up to 100 Gb/s.

In 1998, BU joined forces once again with Harvard and MIT to found the Northern Crossroads (NoX), a regional collaborative established to share human, material, and intellectual resources to foster the development and delivery of advanced network services in New England. The NoX fosters aggregation and collaboration among participating institutions and our respective academic communities and consolidation of regional representation to the national and international community. NoX participants meet regularly to consider issues at hand and to hear from and talk with advanced network advocates about their commitment to R&E initiatives such as Internet2.



In 1999, the NoX established a robust service platform to support the goals of its participants, including sharing the cost of connections to Internet2 and provide a high-performance regional exchange for participants & commodity Internet service providers. In 2003, BU, Harvard, and MIT executed a twenty-year indefeasible right of use with Level3 for a 7.4 mile radius 144 optical fiber loop that runs throughout Boston and Cambridge providing access to major carrier facilities. Portions of this optical ring are allocated to each of the three partners and, in turn, to the NoX, greatly enhancing options for connectivity and providing essentially unlimited intercampus capacity. When subscribers connect to this exchange, local traffic remains within the exchange, improving performance and reducing demand on more congested and costly backbone networks. Since then the NoX has extended its Dense Wavelength Division Multiplexing (DWDM) and routing reach to NY and Virginia facilitating direct peering with major service providers (Google, Microsoft, and AWS), and has provided the Akamai Prolexic DDoS service offering to its membership.



In 2000, BU and the NoX joined forces with other R&E regional network service providers to form The Quilt where leaders throughout the regional networking community build on the intellectual capital and best practices of network service providers worldwide.

BU was an author of the 2000 proposal to extend Internet2 through the Sponsored Education Group Participant (SEGP) program. In 2005, BU was asked to participate in Internet2 strategic planning and has also served on various Internet2 governance and advisory councils, including the Network Planning and Policy Advisory Council (NPPAC) and Internet2's Governance and Nominations Committee (GNC). BU continues to represent the NoX in Internet2 Connector meetings and discussions.

BU is leading the Massachusetts Open Cloud (MOC) project, a collaboration between universities, industry and the Commonwealth to create a new public cloud, designed and implemented in Massachusetts. The MOC is the first cloud computing facility based on the "Open Cloud eXchange" (OCX)—a novel, transformative marketplace model for public cloud offerings. Partners in the collaboration include BU, Harvard, UMass Amherst, MIT, and Northeastern University, as well as the MGHPCC and Oak Ridge National Laboratory. The MOC is funded by a \$3M Commonwealth of Massachusetts MassTech Collaborative Matching Grant Award and more than \$16M of matching grants from industry participants. Efforts in the MOC have led the research computing groups at BU and Harvard to embark on a collaborative pilot project, the New England Research Cloud (NERC), whose aim is to deliver production quality cloud resources and services to their research communities. NERC aspires to extend this partnership with others in the MGHPCC to deliver cloud driven research computing services throughout the region.

The Commonwealth Computational Cloud for Data Driven Biology (C3DDB) is a computer system funded by the Massachusetts Life Sciences Center to support research that connects life science research with emerging innovative big data analytics. The C3DDB is located at the MGHPCC and operated jointly as a shared resource by MIT, BU, UMass, Northeastern, and Harvard.

BU is partnering with Harvard, MIT, UMass, MGHPCC, and Northeastern University in the NorthEast Storage Exchange (NESE) to design, build, manage, and deliver a multi-petabyte, multi-tier, shared storage exchange for science with a sustainable operating model. NESE leverages and extends existing regional infrastructure to support computationally and data intensive research, with an emphasis on lowering barriers to collaboration across institutions and disciplines.