

## Boston University Cyberinfrastructure Plan

### I. BU Research Infrastructure

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

BU recently partnered with government and regional institutions to establish the Massachusetts Green High Performance Computing Center (MGHPCC), operated by BU, Harvard, MIT, Northeastern, and the University of Massachusetts—five of the most research-intensive universities in Massachusetts—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 and the Northern Crossroads (NoX) and New York City. BU initially provisioned a pair of 10G connections from campus to the MGHPCC. BU is also actively engaged in other national and regional consortia where we provide technical, strategic, and governance leadership. BU recently returned to Internet2 strategic planning by participating in its 2014 Community Futures Forum.

#### 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 are typically used in computational science and engineering, simulation, modeling, bioinformatics, genomics, data analysis, and other disciplines that require high-performance computing, massive storage, or complex visualization. Resources are managed in close consultation with the Research Computing Governance Committee, the BU Center for Computational Science, and the Rafik B. Hariri Institute for Computing and Computational Science & Engineering.

Our primary computational facility is a Linux cluster with 3,500 Intel and AMD CPU cores, 103,996 NVIDIA GPU Cuda cores, more than 22 TB of main memory, and over 1 PB of disk, with nodes interconnected by QDR & FDR Infiniband. The cluster is composed of

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

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

Research Computing also provides consulting and training through more than 30 one to three hour training sessions each semester covering research computing basics, programming, mathematics, data analysis, and scientific visualization. Several hundred researchers attend these sessions each term.

#### o COMMUNICATIONS INFRASTRUCTURE

The BU campus network provides high-speed access for over 100,000 devices to all institutional information, communication, and computational facilities, along with the Internet, regionally aggregated resources, and advanced networks such as Internet2. On campus, tens of thousands of ports and wireless access points are interconnected via optical fiber and a robust hierarchy of high-speed routers and switches. The University has two full Class B IPv4 address space assignments.

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 to patch panels and cabling. We also maintain a comprehensive database that maps the network to floor plans spanning hundreds of buildings.

Encrypted 802.1x wired and wireless network access is available throughout campus, including all residence halls and classrooms. The University recently invested several million dollars to ensure the entire wireless network conforms to the most current 802.11n standards. Redundant connections to the Internet and Internet2 operate at 10G with automatic failover.

## II. Planning

At Boston University, cyberinfrastructure planning is an integral part of comprehensive institutional IT strategy developed by stakeholders within an overarching governance framework. BU leadership established formal IT governance several years ago and that framework is described at [bu.edu/tech/about/governance/](http://bu.edu/tech/about/governance/).

With that framework in place, institution-wide IT strategy was formulated and later completed and published in 2014. That strategy is outlined in the 2015-20 BU Technology Plan posted at [bu.edu/tech/plan/](http://bu.edu/tech/plan/).

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**.

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

## o 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.

Toward that end, we have a /32 IPv6 allocation from ARIN which we now announce on edge routers to Internet2 while an internal technical working group formulates [1] IPv6 address allocation strategy, [2] a core protocol deployment plan to transition from IS-IS to OSPFv3, [3] an assessment of network component IPv6 readiness, and [4] a plan for implementation of an IPv4 to IPv6 gateway. Eduroam was deployed in June 2014.

- 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.

## o 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 and Medical campuses, enabling both optimal use of resources and maximum collaboration.

## o 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: BU's History of Commitment to Shared Infrastructure in Support of R&E

### I. 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 has been a leader in regional networking for more than thirty years. We brought BITNET to New England in the early 80s, and we subsequently founded NEARnet, the New England component of NSFnet, with Harvard and MIT. The three institutions later founded the Northern Crossroads 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.

#### 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 and institutional scientific computing and visualization support, providing focus to drive computational science and advanced research computing.

### II. 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, with wider deployment planned later this year. Additional discussions regarding a Genesis II pilot are ongoing.

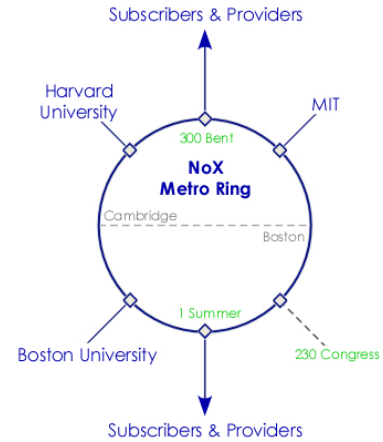
In collaboration with Harvard, BU runs the New England ATLAS Tier 2 Center, one of five such centers in the US and a participant in the Open Science Grid. ATLAS is one of two general purpose experiments running on the Large Hadron Collider at CERN and involves over 3,000 physicists world-wide. The NE ATLAS Tier 2 supports 1,200 researchers with 3,700 compute cores and more than 2 PB of storage.

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 at BU 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.



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).