

The New England Research Cloud - Concept Overview

May 14, 2020

Background

A decade ago the MGHPCC consortium was formed - a \$95 million, 90,000 square foot, 15 megawatt compute center in Holyoke, MA - with the space, power, and cooling capacity for more than 750 racks of computing equipment on a single shared floor. The resulting advantages of this facility and partnership are a simple operational model for cost efficiency and the ability to build out services that would not otherwise have been possible.

The vision of the New England Research Cloud (NERC) is to create a regional resource, rich with world-class services that include data storage, high-performance computing, and cloud computing. To this end, we are building a cohesive, sustainable partnership between research clusters and technology professionals with a common cloud framework that is tailored for data-driven discovery and will be available to many institutions in New England. This partnership will set NERC apart from today's public clouds and is critical to bridge the gap in skills that are needed for research and teaching professionals.

The NERC goals will enable the development of (1) a cost effective professionally operated on-premise production cloud service that includes various levels of services - Software-as-a-Service (SaaS); Platforms-as-a-Service (PaaS); and Infrastructure-as-a-Service (IaaS) - that includes emerging new technologies for hardware acceleration, (2) set standards of deployment and automation that will allow other institutions to easily deploy the full suite of services built within the NERC, and (3) expanded capabilities of the emerging workforce by connecting local universities as well as the innovation hubs and start-ups to our universities.

Goals: Immediate and Long-term

The pilot (build out) phase, which will run through the end of 2021, will focus on developing the technology to make it easy for researchers to take advantage of a suite of services (IaaS, PaaS, SaaS) that are not readily available today. More specifically, the goal of the pilot is to accomplish: (1) the creation of the building blocks needed for production cloud services, (2) begin collaboration with Systems Engineers from other institutions, (3) on-board select POC use cases from a variety of institutions in Massachusetts outside of MGHPCC consortium.

The longer term objectives will be centered around activities that will focus on (1) engaging with various OpenStack communities by sharing best practices and setting standards for deployments; (2) connecting regularly with the Mass Open Cloud (MOC) leadership to understand when new technologies they are developing with RedHat, Inc. - and as part of the new NSF funded Open Cloud Testbeds - might be ready for adoption into the production NERC environment; and (3) broadening the local deployment team of NERC to include our other partner universities within MGHPCC.

Through these efforts, NERC will lower the barrier of technology adoption to a much broader community by offering services and tools to researchers and students in an early stage of their learning; therefore, increasing their technology readiness. Within the MGHPCC consortium institutions alone, there are currently nearly 200,000 students with 20,000 active researchers using MGHPCC cluster computing resources. Within a fully realized vision of NERC, there could be hundreds of thousands of users across the region, including STEM-based classes in middle-school, high-school, and universities, all sharing common teaching platforms for data science and machine learning, as well as innovation hubs and start-ups connected to our universities.

Leadership

The NERC pilot will be developed and operated jointly by the BU Research Computing (Wayne Gilmore) and Harvard FAS Research Computing (Scott Yockel) groups, with shared leadership between their Directors (Tracy Schroeder and Anne Margulies). For more information, please visit <https://nerc.mghpcc.org/>.