

Northeast Tier 2...2019, 2020,...



1. Strategic Considerations
2. Tactics for 2019...2020...



Boston University  @BU_Tweets · 14m

President Brown today announced a proposed project to give data science a home on Comm Ave — a 17-floor building meant to resemble a stack of books that will be home to @BU_Computing, mathematics and statistics, and computer science.



BU to Build Data Sciences Center

Boston University announced its plans to build the BU Data Sciences Center. By bringing mathematics and statistics and computer science dep...

bu.edu

October 1, 2018

M.I.T. Plans College for Artificial Intelligence, Backed by \$1 Billion



The Massachusetts Institute of Technology is taking a particularly ambitious step in preparing students to develop, and consider the implications of, artificial intelligence. It is creating a new college, backed by a planned investment of \$1 billion.

Cody O'Loughlin for The New York Times

New York Times...*October 15, 2018*

Deals

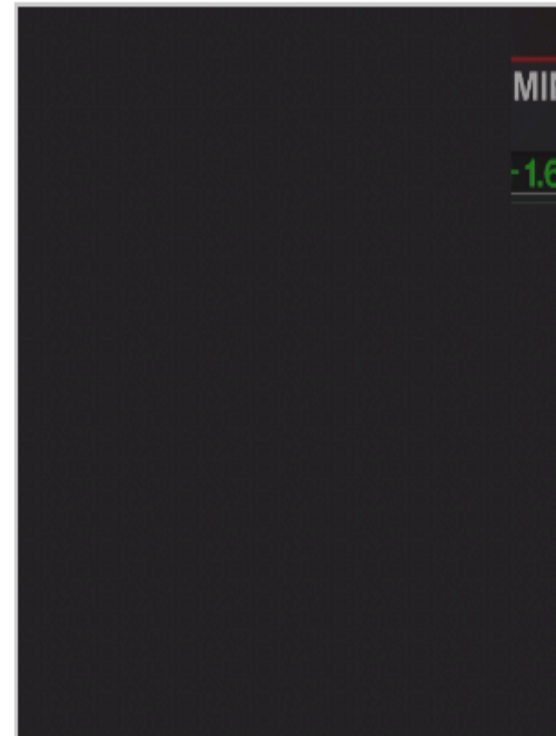
IBM to Acquire Linux Distributor Red Hat for \$33.4 Billion

October 28, 2018

By [Ed Hammond](#), [Kiel Porter](#), and [Alex Barinka](#)

October 28, 2018, 1:54 PM EDT Updated on October 28, 2018, 4:34 PM EDT

- ▶ Computer-services giant to pay \$190 a share for Red Hat
- ▶ Software maker will become a unit of IBM's Hybrid Cloud team



MGHPCC

Holyoke, Massachusetts

Near zero carbon footprint

>350,000 x86 cores now

Space/power/cooling for 785 racks

90,000 square feet of computer floor

Space for 2d building or pods

Consortium: BU, Harvard, MIT, NEU, UMASS,
Commonwealth of Massachusetts, DELL/EMC and Cisco



National compute & storage resources

San Diego Supercomputer Center:
- Comet: 46752 CPUs
- Comet GPU: 1728 processors
- Oasis (storage): 4000 TB



Pittsburgh Supercomputing Center
- Bridges: 21056 CPUs
- Bridges GPU: 1344 processors
- Bridges large memory: 160 CPUs
- Bridges Pylon (storage): 10000 TB

● Compute Resources
● Storage Resources

MGHPCCC



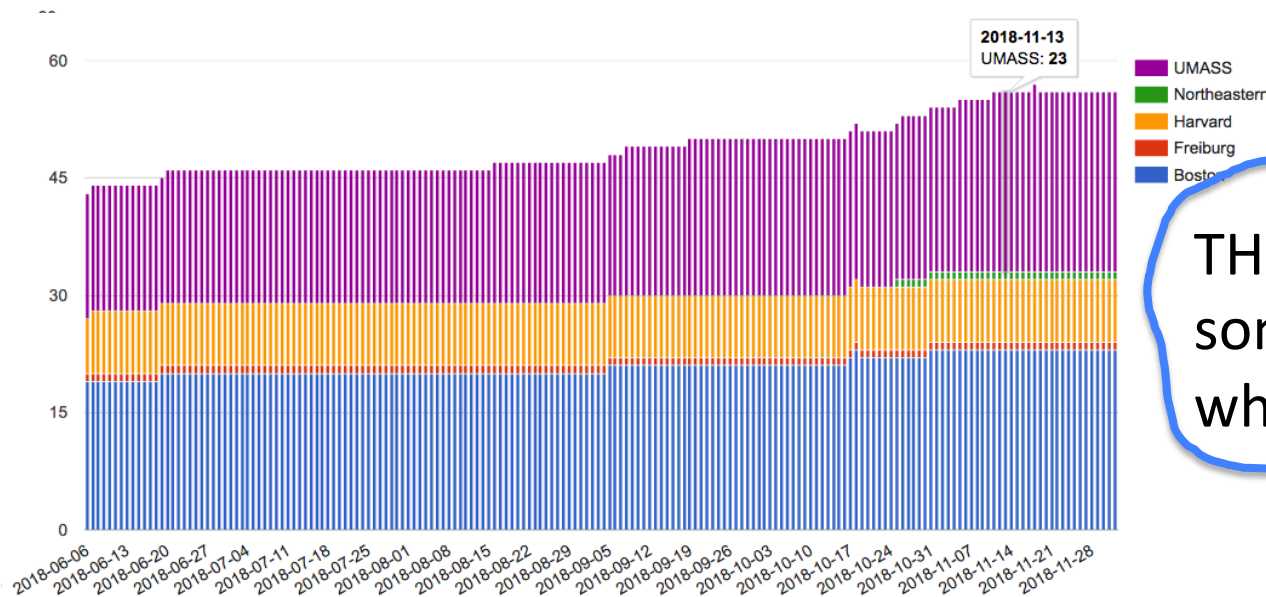
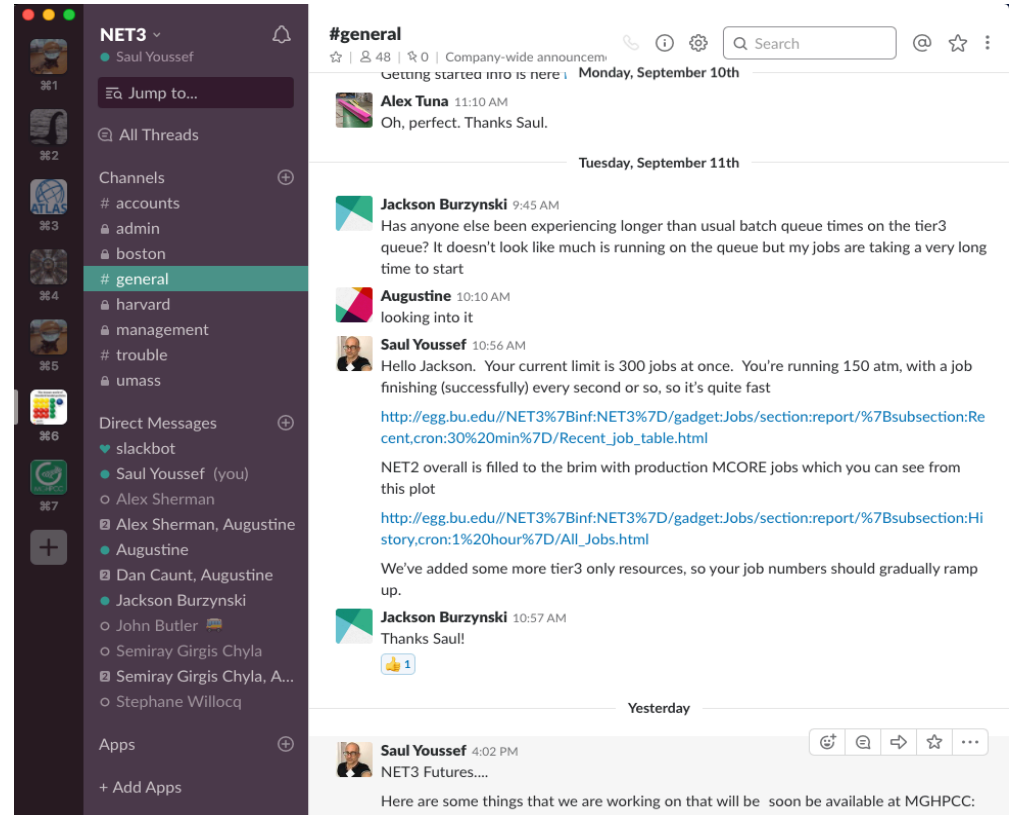
Texas Advanced Computing Center:
- Stampede2: 368280 CPUs
- Wrangler: 2304 CPUs
- Ranch (storage): 61440 TB
- Wrangler (storage): 10000 TB

350,000 cores
~50000 TB storage

Northeast Tier 3

Boston University
Harvard University
UMASS/Amherst


- Operated by Wayne's group as part of NET2
- 56 Users
- Gradually growing
- Almost as many UMASS users as BU users
- Essentially no extra work for Augustine
- Use of Slack is extremely helpful
- UMASS bought in with three nodes



THEME: If you're good at something, do it for the whole consortium.



Internet2, Internet1
ESNet, LHCONE,...



Up to
576
100G

CephFS/Posix

Globus

Gridftp

LHC Data

S3/Swift

Block storage



University buy-in, Project buy-in, NESE project funds, Inherited Equipment

MIT Supercloud

Education and Outreach

PILOT project including BU

THEME: If you're good at something, do it for the whole consortium.

- Home
- Requesting An Account
- Getting Started
- How To Use
- Jupyter Portal
- Online Courses
- Contact


Login using Touchstone

Systems and Software

This page lists information about the system and available [software](#), [languages](#), [compilers](#), [modules](#), etc. This is only a partial list, so if there is anything you are interested in that isn't listed here, please [contact us](#).

MGHPCC TX-E1 Specifications

Summary						
Number of Nodes		56				
Total CPU Cores		1348				
Total GPUs		10				
Distributed Storage		873 TB				
CPU-Only Nodes						
Processor	Nodes	Cores	RAM	Local Disk		
Intel Xeon	25	16	64 GB	16 TB		
Intel Xeon	7	2 x 14	256 GB	12 TB		
AMD Opteron*	20	2 x 16	192 GB	8 TB		
GPU Nodes (Intel Xeon CPU)						
GPU Type	GPU Interconnect	Nodes	GPUs	CPU Cores	RAM	Local Disk
Volta V100	NVLink	1	4	2 x 14	500 GB	2 TB
Volta V100	PCIe	3	2	2 x 14	500 GB	2 TB




Jeremy Kepner

Mathematics of Big Data

Spreadsheets, Databases, Matrices, and Graphs

Jeremy Kepner and Hayden Jananthan



Foreword by Charles E. Leiserson

MIT LINCOLN LABORATORY SERIES


* Coming soon!

2019...2020...Tactical

- Retire old HU nodes, absorb rest into BU pods
- Retire LSM in favor of rucio-mover
- RH7 upgrade => Singularity readiness
- New, small GPFS SSD system pool
- Networking upgrade including NET2<->NESE @200G
- Squid tune-up (re: John DeStephano, Nurcan)
- Harvester/single queue/SGE readiness
- Slate/k8 node acquisition & prep
- Creation of NESE, then MGHPCC-wide operations data co-op
- NESE ATLAS DDM endpoint creation & testing with gridFtp, xrootd,...(coordinate with Wei & Hiro)
- Technology evaluation for 2019..2020... NESE (learn from Xin re: tape carousels)

MIT "Supercloud" AI, Jupyter, Matlab, GPU-heavy

Operations
Data Co-op



Pooled worker nodes,
dynamically importable
into HPC or Cloud
environments

NESE Ceph



NESE Flash

NESE Carousel

Mass. Open Cloud, RH/IBM AI stacks, GPU-heavy

THEME: If you're good at something, do it for the whole consortium.

1st Deployment being cabled now. 12PB starter; 40+% buy-in already RedHat Partnership Going!

NESE: The Northeast Storage Exchange

Saul Youssef,¹ Scott Yockel,² Chris Hill,³ John Goodhue,⁴ Devesh Tiwari,⁵ and Mike Zink⁶
 Boston University,¹ Harvard University,² MIT,³ MGHPCCC,⁴ Northeastern University,⁵ University of Massachusetts⁶



Massachusetts Green High Performance Computing Center (MGHPCC)
 • 15 megawatts, 785 racks, approx. 2/3 full now
 • >350,000 x86 cores now
 • Redundant 100Gb fiber ring to national research networks
 • Secure, single use site
 • Room for second building and/or expansion via pods




GOALS

- To meet the storage needs of the data revolution for science, engineering, education and technology.
- To be the start of a national cyberinfrastructure in the Northeast U.S.

STRATEGY

- Take advantage of the basic economics of storage.
- Use our unique consortium and the MGHPCC facility as a starting point.
- Organize the project for long term growth, with institutional responsibilities, community building, long term planning, integration with education, and long term technology tracking.
- Use our unique environment for partnerships with universities, technology companies, biotechs, hospitals, institutes, and new data science centers.



Imagine shrinking HPC pods to the size of a building and connecting with MGHPCC



Context: The Consortium is comparable to the single Pacific Research Platform in a single building:

- 5 Universities, 9 campuses including Harvard and MIT
- 17,000 researchers
- More than 100,000 students
- Five new Data Science initiatives or institutes
- Looking to expand membership

PROJECT ORGANIZATION

- Management, planning, outreach, science coordination, sustainability: *All PIs*
- Operations: *Scott Yockel*, Harvard FAS Research Computing, BU, RedHat
- Networking: *Jefferson Burson*, Harvard University IT
- Technology planning: *Devesh Tiwari*, Northeastern University
- Ocean apps and iRods: *Christopher Hill*, MIT
- Block storage for clouds, CloudLab: *Mike Zink*, UMass/Amherst
- NET2, LHC applications, federated LHC storage: *Saul Youssef*, Boston University
- Collaboration with RedHat, Mass. Open Cloud: *Orran Krieger*, Boston University
- Authentication and Globus: *Jim Culbert*, MGHPCCC
- Education, Outreach, Open Storage Platform: *John Goodhue*, MGHPCCC

For Data Science, location matters!



DATA SCIENCE
 Every one of the five Consortium member universities has a new data science institute, initiative or program.

- Boston University Data Science Initiative
- Harvard Data Science Initiative
- MIT Institute for Data, Systems and Society
- Data Science at Northeastern University
- Center for Data Science, UMass/Amherst

- Every Consortium university will be able to create their own control room showing Northeast U.S. cyberinfrastructure operations including NESE. This will be great for
 - Community building
 - Student projects and training
 - Inspiring new generations of operations software based on Data Science and AI

Amazon Says 900 Jobs Will Be Added At New Boston Office



Partnership with RedHat


- Via BU/RedHat "Collaboratory"
- Design consultation...including with Sage Weil
- Installation help
- Embedded RedHat personnel
- Use of RedHat facilities in Boston
- Premium Software Subscriptions

RedHat
 Intel
 DELL/EMC
 Google
 NVIDIA
 Amazon
 Facebook
 Microsoft
 Wolfram Research
 Mellanox
 Cisco
 General Electric

Amazon to add 2,000 jobs in Boston



Internet2, Internet1
 ESNET, LHCONE, ...



CephFS/Posix Globus Gridftp LHC Data S3/Swift Block storage



University buy-in, Project buy-in, NESE project funds, Inherited Equipment

Model	Capacity	Performance	Notes
2 x 100Gb/s uplinks to 100Gb/s Fabric			
Basic Storage Unit			
1U	As dense as 60 bay J80D	12 x 10TB 7200 rpm spinning	
		Bluestore Ceph + CephFS	
		4 x 280GB Micron high endurance SSD	
		1 x 32GB NVMe	
		Ceph write-ahead log	
		2 x 100Gb/s	
		Matches drives I/O	
		1 x Single Socket Intel Xeon SkyLake 2.2 GHz (5.0 GHz Turbo), 35C/TDP with AVX-512 SMT0 coprocessor	

3.9 PB raw per rack when fitted
 Space for more buy-ins
 Start of UMass/Amherst buy-in



First Deployment

- Six racks at two sites
- 12 PB raw storage
- Small, flexible 120TB OSD
- Designed for all storage types: CephFS, Block storage, S3, Swift, Globus endpoint
- SSDs for Ceph Bluestore
- 100Gb/s networking

FIRST SCIENCE




- NESE storage for the U.S. ATLAS Northeast Tier 2 center (NET2, BU/ Harvard) is one of the first production applications. 6PB and growing.
- Federating NESE storage with OSIRIS/Also NSF-DIBBS and with tape storage at Brookhaven National Laboratory to produce a prototype national federation which can be scaled to the 1 EB needed in 2026 for the high luminosity LHC
- Ocean Modeling by Chris Hill's group at MIT uses data from ships, fixed stations, from space and from simulations.



- Jeffrey Lichtman's lab at Harvard produces ~10 PB per year of electron microscopy images of mouse brains with the goal of mapping neuronal circuits and understanding brain function
- Mark Friedl of BU needs to analyze 10s of PB of multi-spectral time-series satellite images of the Earth.



- Lars Hernquist's group at Harvard studies galaxy formation and cosmology with simulations which need 100s of TB for intermediate steps.
- A next-gen gene sequencing facility at Harvard needs to store and process tens of TB per day.



- Advanced microscopy at Harvard images mitochondria in single cells and produces approximately 1 PB of data per year.
- The CASA project at UMass/Amherst aims to revolutionize rapid response to hazardous weather. Data currently at DFW metacenter.

INTERVIEW: BU physicist on creating new, math-driven field of medicine



- Network physiology is a new field invented by Plamen Ivanov at BU where unique medically important insights can be extracted from multi-source human physiology data with fast enough access for clinical response.
- These are only some of the first expected NESE applications to be discussed at meetings like HPC Futures.

HPC Futures

Present, plans, and planning for computing at Boston area universities, institutes, hospitals and companies

Organization	Speakers	Programs/Conferences
Harvard	Scott Yockel, Jeffrey Lichtman, Mark Friedl, Lars Hernquist, Devesh Tiwari, Chris Hill	Harvard HPC, Harvard HPC Consortium
MIT	Chris Hill, Devesh Tiwari	MIT HPC, MIT HPC Consortium
UMass/Amherst	Mike Zink, Orran Krieger	UMass/Amherst HPC, UMass/Amherst HPC Consortium
Boston University	Scott Yockel, Jeffrey Lichtman, Mark Friedl, Lars Hernquist, Devesh Tiwari, Chris Hill	Boston University HPC, Boston University HPC Consortium
NET2 project	Scott Yockel, Jeffrey Lichtman, Mark Friedl, Lars Hernquist, Devesh Tiwari, Chris Hill	NET2 project

• We already have 40% buy-in

