


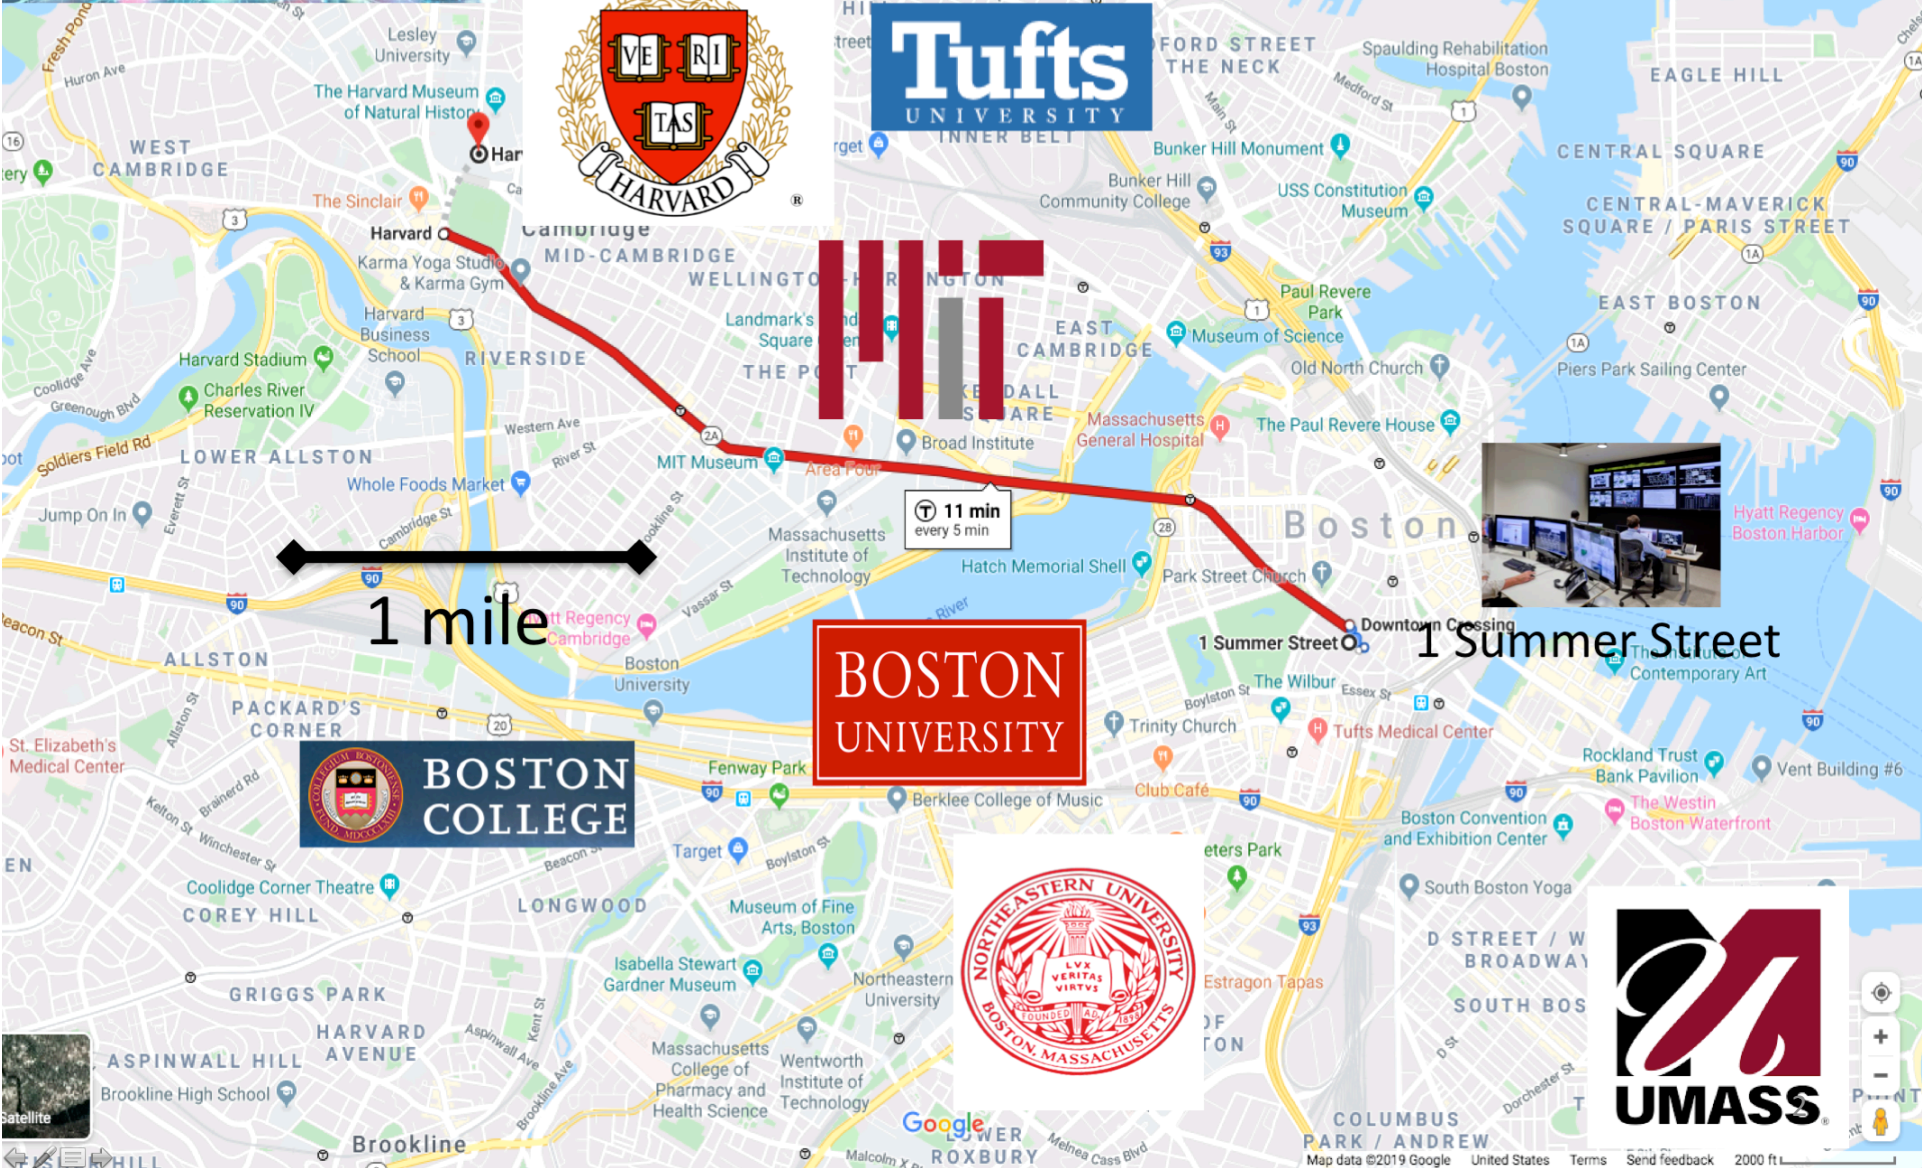
# NET2 Evolution

Facilities Meeting

October 30, 2019

Saul Youssef

← 100 miles



# MGHPCC

Boston University  
Harvard University

MIT

Northeastern University

University of Massachusetts

- 15 megawatt \$90M single purpose data center
- Near zero Carbon footprint
- Space power and cooling for 780 racks
- More than 300,000 x86 cores, millions of gpu cores
- 100Gb/s multi-fiber ring to internet2 and Esnet
- Three new top500 in the past few months
- Exascale storage plans via NESE project
- Located in Holyoke, MA
- Thousands of researchers, >100,000 student population



**FASRC CANNON**  
HARVARD'S LARGEST CLUSTER

- 100,000 CPU CORES  
3,000+ NODES
- 500 TB RAM  
40PB STORAGE  
2.5M CUDA CORES
- 29 MILLION JOBS/YR  
300 MILLION CPU HR/YR
- 3 DATA CENTERS @ 10K+ FT<sup>2</sup>  
BOSTON, CAMBRIDGE, & LEED PLATINUM GREEN DATA CENTER IN HOLYOKE, MA
- 500+ LAB GROUPS  
OVER 5500 USERS

CANNON: THE FASRC CLUSTER IS NAMED IN HONOR OF ANNIE JUMP CANNON, A PIONEER IN ASTRONOMY.

## 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



Boston University launches University-wide Computing and Data Science initiative, new building.

top500



MIT News  
ON CAMPUS AND AROUND THE WORLD

NEWS VIDEO SOCIAL FOLLOW MIT

Browse or Search

An \$11.6 million artificial intelligence computing cluster donated by IBM to MIT will come online this fall at the Massachusetts Green High Performance Computing Center (MGHPC) in Holyoke, Massachusetts.  
Photo: Helen HBM/HPCC

### IBM gives artificial intelligence computing at MIT a lift

Nearly \$12 million machine will let MIT researchers run more ambitious AI models.

Kim Martineau | MIT Quest for Intelligence  
August 26, 2019

IBM designed Summit, the fastest supercomputer on Earth, to run the calculation-intensive models that power modern artificial intelligence (AI). Now MIT is about to get a slice.

IBM pledged earlier this year to donate an \$11.6 million computer cluster to MIT modeled after the architecture of Summit, the supercomputer it built at Oak Ridge National Laboratory for the U.S. Department of Energy. The donated cluster is expected to come online this fall when the MIT Stephen A. Schwarzman College of Computing opens its doors, allowing researchers to run more elaborate AI models to tackle a range of problems, from developing a better hearing aid to designing a longer-lived lithium-ion battery.

RELATED
Satori
MIT Quest for Intelligence
MIT-IBM Watson AI Lab
Department of Brain and Cognitive Sciences

### Lincoln Laboratory's new AI supercomputer is the most powerful at a university

TX-GAIA is tailor-made for crunching through deep neural network operations.

SEPTEMBER 26, 2019 | Kylie Foy | Communications & Community Outreach Office



TX-GAIA is housed inside of a new EcoPOD, manufactured by Hewlett Packard Enterprise, at the site of the Lincoln Laboratory Supercomputing Center in Holyoke, Massachusetts. Photo: Glen Cooper.

The new TX-GAIA (Green AI Accelerator) computing system at Lincoln Laboratory's Supercomputing Center (LLSC) has been ranked as the most powerful artificial intelligence (AI) supercomputer at any university in the world. The ranking comes from TOP500, which publishes a list of the top supercomputers in various categories biannually. The system, which was built by Hewlett Packard Enterprise, combines traditional high-performance computing hardware — nearly 900 Intel processors — with hardware optimized for AI applications — 900 Nvidia GPU accelerators.

BU Rafik B. Hariri Institute for Computing and Computational Science & Engineering

CONNECT DIR

ABOUT CENTERS RESEARCH PEOPLE NEWS & EVENTS RESOURCES



## Collaboration Awarded an NSF Grant of \$5M to Create New Cloud Computing Testbed

Last month the National Science Foundation (NSF) awarded a significant grant to a team of researchers from Boston University, Northeastern University and UMass Amherst. The grant will support the development of a testbed for new cloud

# NESE: Northeast Storage Exchange

National Science Foundation CIF21 DIBBs award 1753840




Boston University

Harvard University

Massachusetts Institute of Technology

Northeastern University

University of Massachusetts



# PHASE I: HARDWARE SPECS



Single Socket Intel Xeon 4114 Skylake 2.2 Ghz (3.0 Ghz Turbo), 10C/20T with AVX-512 (AES Encryption offloading)

Bonded 2x10G SFP+ (NESE Net)  
Single 1Gbps OBM (Admin net)

Deployed at two sites with 100Gb ethernet.

12 x 10TB HDD



Data Drive

4 x 480GB SSD



3x RocksDB  
1x CephFS  
metadata OSD

1 x 64GB NVMe  
Optane



RocksDB-WAL

2x 100GB  
SATADOM



SW RAID1 Root  
Filesystem

**Starter deployment is 15 PB raw already  
with ~50% buy-in.**

#redhat #rhsummit



Respond effectively to data growth

The Dell EMC PowerEdge R740xd2 helps you plan for future growth with large internal storage and cost-efficient drive capacities. Deliver two-socket compute performance with flash and fast networking options to meet streaming demands. Simplify management of large data sets with automated administration and front-serviceable drives. The R740xd2 lets you keep your data safely on-premise with built-in security, even as you scale capacity.



<- Major buy of DELL r740xd2 with 14TB drives, 25G nics

6PB raw for NET2

3PB raw for Northeastern University

7PB more expected soon

# Northern Crossroads at MGHPCC



Demo saturating 2 gateways

gridftp.nese.mghpcc.org -> DNS round robin

2 x 100G  
dedicate  
for NET2

DELL S5048-ON

26 x 10G

ARISTA 100G  
 ARISTA 100G  
 NESE gateway 25G  
 NESE gateway 25G  
 NESE gateway 25G  
 NESE gateway 25G  
 NESE gateway 25G  
 NESE gateway 25G  
 4 more on order

DELL S5048-ON  
 DELL S5048-ON  
 OSD  
 OSD  
 .  
 .  
 .  
 OSD

...

nic.host	nic.name	nic.inet_addr	nic.RX_bytes_s	nic.TX_bytes_s	nic.RX_errors	nic.TX_errors
r1n1.nut.bu.edu	bond0	192.168.3.216	256.94 MB/s	763.09 MB/s	15985	0
r1n2.nut.bu.edu	bond0	192.168.3.217	255.57 MB/s	755.20 MB/s	0	0
r6n1.nut.bu.edu	p1p1	192.168.3.222	5.62 MB/s	0.08 MB/s	0	0
r6n2.nut.bu.edu	p1p1	192.168.3.223	6.12 MB/s	0.08 MB/s	1	0
r6n3.nut.bu.edu	p1p1	192.168.3.224	17.76 MB/s	0.10 MB/s	0	0
r7n1.nut.bu.edu	p1p1	192.168.3.225	168.67 MB/s	0.30 MB/s	0	0
r7n2.nut.bu.edu	p1p1	192.168.3.226	214.91 MB/s	0.30 MB/s	0	0
r7n3.nut.bu.edu	p1p1	192.168.3.227	521.25 MB/s	0.74 MB/s	0	0
r8n1.nut.bu.edu	bond0	192.168.3.228	2.00 MB/s	0.54 MB/s	0	0
r8n2.nut.bu.edu	bond0	192.168.3.229	1.75 MB/s	0.45 MB/s	0	0
r8n3.nut.bu.edu	bond0	192.168.3.209	236.70 MB/s	331.93 MB/s	0	0
r8n4.nut.bu.edu	bond0	192.168.3.210	192.14 MB/s	383.73 MB/s	107087822	0
r9n1.nut.bu.edu	bond0	192.168.3.212	227.13 MB/s	239.07 MB/s	0	0
r9n2.nut.bu.edu	bond0	192.168.3.213	359.11 MB/s	237.70 MB/s	0	0
r9n3.nut.bu.edu	bond0	192.168.3.214	273.79 MB/s	719.65 MB/s	0	0
r9n4.nut.bu.edu	bond0	192.168.3.215	261.62 MB/s	721.84 MB/s	403	0

5PB GPFS

8.2 PB raw -> 6.0 PB useable for NET2 with 8+3 EC<sup>7</sup>

1. Retire old Harvard worker nodes
2. Retire LSM software, switching to rucio-mover
3. 100G dedicated networking NET2-NESE
4. Docker containers with Gridftp & Wei's Adler32 callout
5. DNS round robin from gridftp.nese.mghpcc.org to gridftp01,02,...

## ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups Service object details Docs TWiki OLD

**Service**

Name: BU\_ATLAS\_Tier2\_SE\_173 Site.state: ACTIVE State: ACTIVE Operations: Show Changes log Edit

Type: SE Last modified: 2019-10-22 15:55 State updated: 2019-10-22 15:59

Site: BU\_ATLAS\_Tier2 Implementation: State comment: Object state auto changed to ACTIVE after the creation of DDMEndpoint (NESE\_DATADISK)

is\_new: True Description:

**Associated Protocols**

Name	Flavour	Endpoint	Basepath	State	Status	Activities	is_monitored	Description	Settings	Operations
BU_ATLAS_Tier2-SE-GRIDFTP-gridftp.nese.mghpcc.org	GRIDFTP	gsiftp://gridftp.nese.mghpcc.org:2811		ACTIVE		delete_wan/0, read_wan/0, third_party_copy/0, write_wan/0	False			

Add new Service protocol

Manage Service protocols for activity: READ WAN WRITE WAN DELETE WAN READ LAN WRITE LAN DELETE LAN Third party copy

**Service Resources**

Name	Basepath	Endpoint	DDMEndpoints	Extra settings	Operations
ATLASDATADISK	/atlas/datadisk/rucio		NESE_DATADISK   path=		

Add new Service resource

The plan is to use both NET2\_DATADISK (read/write)  
NESE\_DATADISK (read only)

...for our PanDA queues

We might give Globus endpoint space to our BU/Harvard/UMass Tier 3 peeps at some point.