

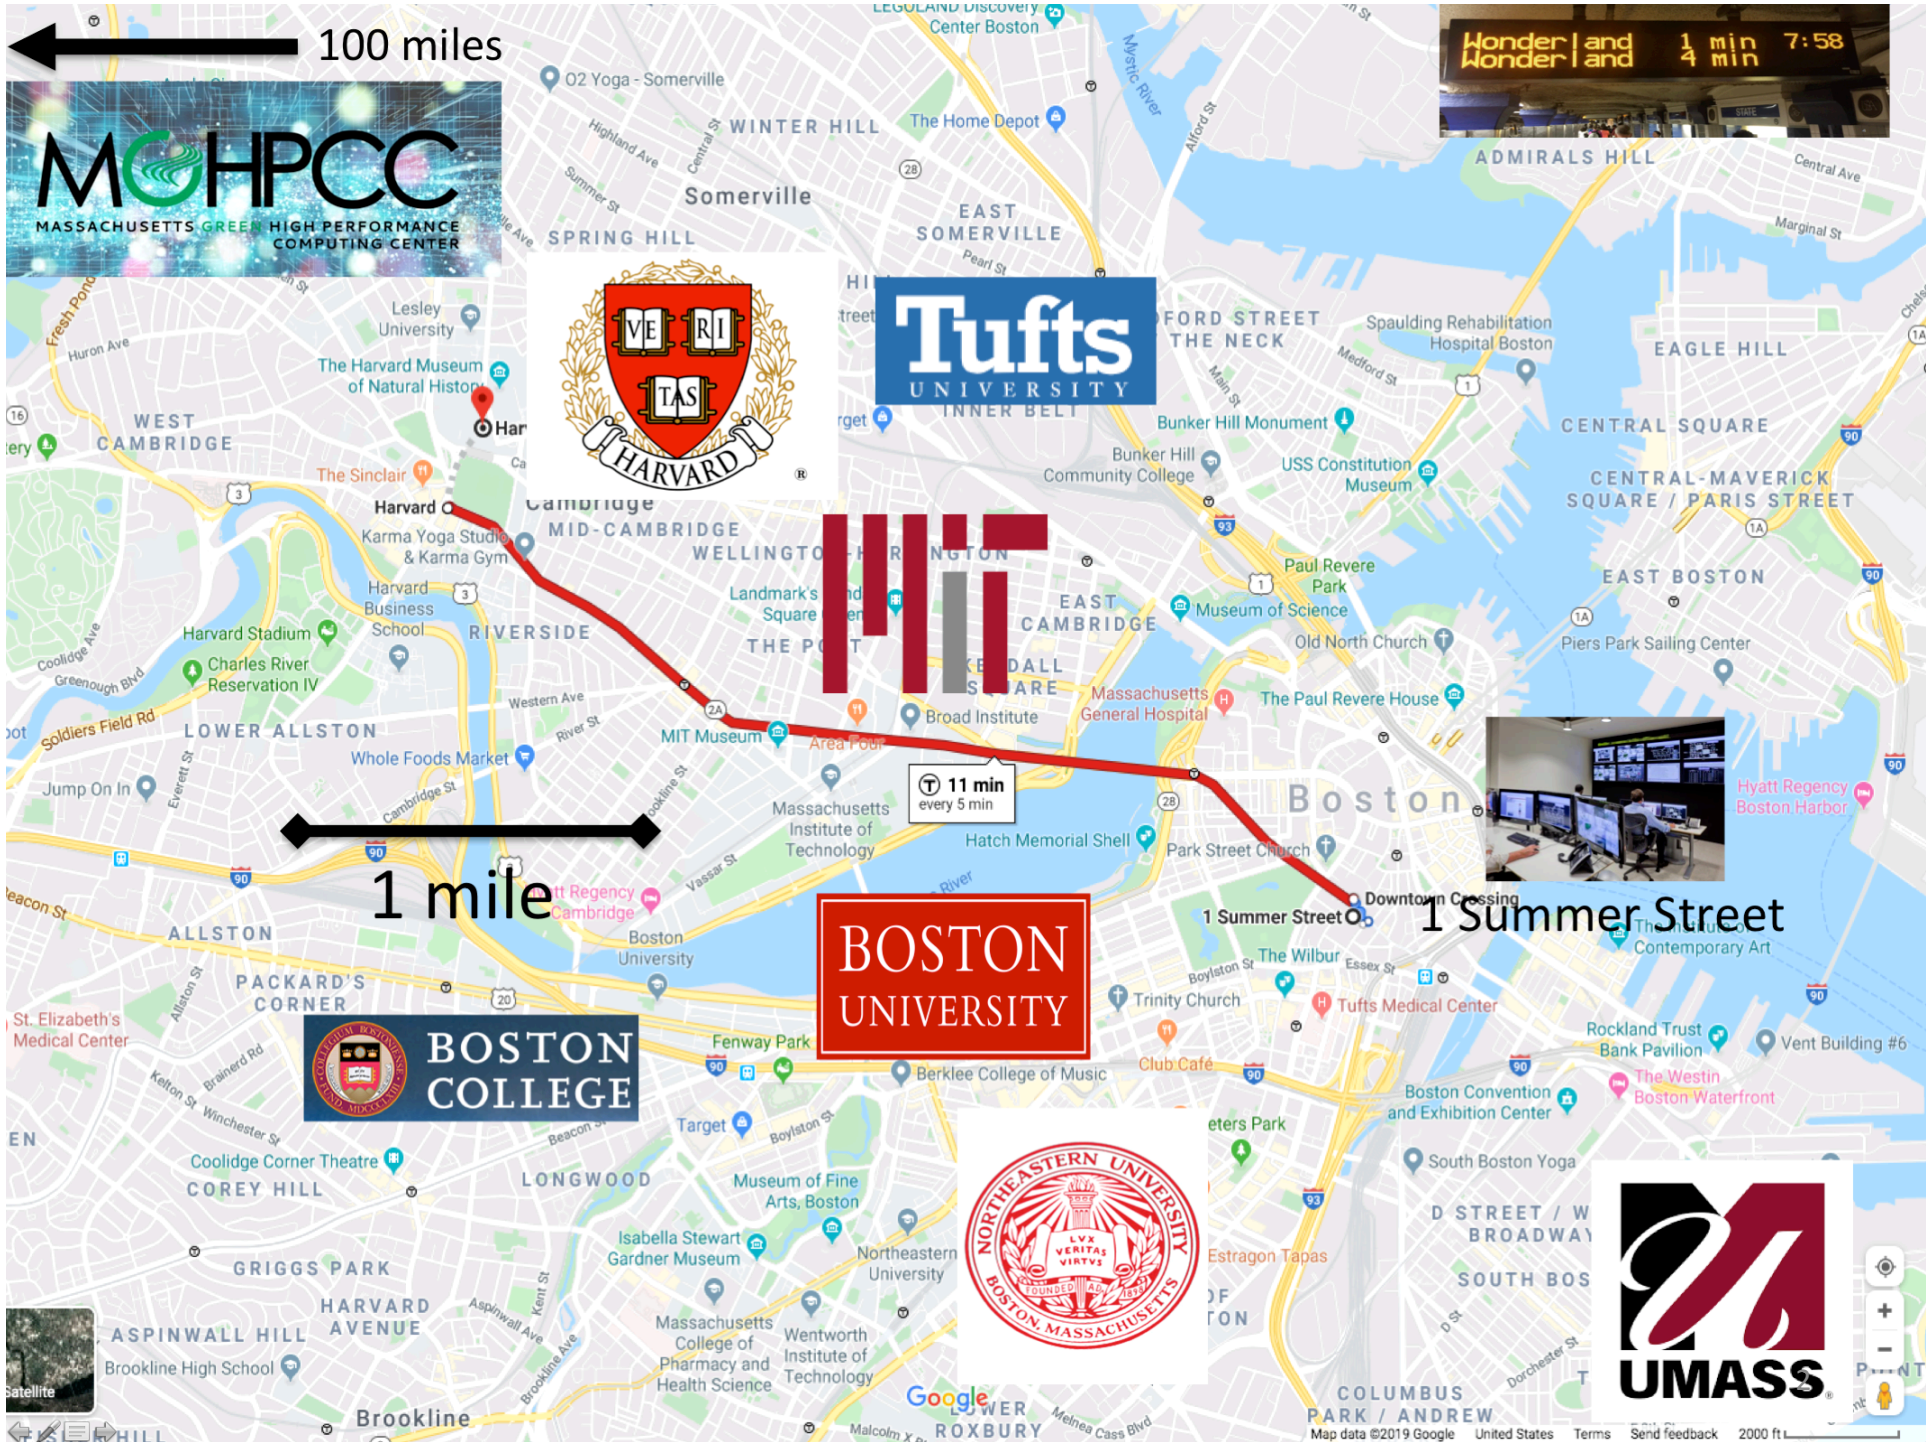
# Storage Evolution at NET2 with NESE

*Saul Youssef*

Boston University

2020-11-20

- Motivation
- Strategy
- NESE Exascale for HL-LHC









# 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 year
- Exascale storage via NESE project
- Located in Holyoke, MA
- Thousands of researchers, 200,000 student population








# DIBBS: Northeast Storage Exchange

PI: Saul Youssef, Co-PIs: Scott Yockel, Chris Hill, John Goodhue, Devesh Tiwari, Mike Zink  
Boston University, Harvard University, MIT, MGHPCC, Northeastern University, UMass, RedHat

Award #: 1753840



Large Hadron Collider, NSF, DoE

**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, & ILED 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, PIONEER IN ASTRONOMY

Harvard Research Computing



Harvard Data Science

**casa** Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere

Revolutionizing our ability to observe, understand, predict and respond to hazardous weather events

**Research Computing at Northeastern University**

Connecting the research community at Northeastern University with high performance computing solutions.

Northeastern University Research Computing  
CASA at UMass/Amherst is NSF funded

## New England Research Cloud (NERC)

**New Regional Stakeholders**

- Federal Funded Workforce Development: Northeast CyberTeam
- High Ed Consortia: Five College Consortium, Colleges of the Fenway, The Boston Consortium, ...
- Other Higher Ed: Boston College, Brandeis, Bentley, ...
- University incubator start-ups: Zapata Computing, ...
- K-12 support for data analytics learning.

**New Services**

- (NERC) Research Cloud Infrastructure & Software Services
- Leadership class HPC
- (NESE) Distributed Data & Archiving Services
- Collaboration / Facilitation & Support
- Innovation Hub / Industry Engagement
- Commercial Cloud & Services Broker

Proposal to Mass. Technology Council



Large Millimeter Telescope, NSF



Boston area research: 7 universities, 15 hospitals, dozens of major biotech companies, Intel, RedHat, IBM, Google, Amazon, Microsoft, Math Works, Wolfram, NVIDIA, many startups. NESE "posix cloud" deployment from 1 Summer Street.



NESE main deployment at  
MGHPCC, Holyoke, MA

# NESE



- Fast, cost-effective, sustainable shared storage.
- Long term commitment from Harvard operations team.
- 24 PB Ceph, expanding rapidly with 70% buy-in so far.
- Adding tape libraries for massive long term storage.

**CloudLab**

Flexible, scientific infrastructure for research on the future of cloud computing. Researchers use CloudLab to build their own clouds, experimenting with new architectures that will form the basis for the next generation of computing platforms.

Open Cloud Testbed, Massachusetts Open Cloud, NSF



Open Storage Network, NSF

**EAPS**  
Earth, Atmospheric and Planetary Sciences

Earth. Planets. Climate. Life.

The Department of Earth, Atmospheric and Planetary Sciences (EAPS) is the place at MIT where the turbulent oceans and atmosphere, the inaccessible depths of the inner Earth, distant planets, and the origins of life all come together under one intellectual roof.

EAPS IN THE NEWS  
UPCOMING EVENTS

MIT Earth, Atmospheric and Planetary Sciences



Institute for Applied Life Sciences

Massachusetts Life Sciences Center



ABOUT THE CEPH FOUNDATION

ceph foundation



RedHat/IBM

# The overall NESE strategy:

1. Strong science, engineering or educational need for a particular type of storage.
2. Economics favoring a shared facility.
3. Strong operations team that can take on a new facility.
4. Long term institutional commitment, both from an operations team and from the MGHPCC consortium.



# NESE Ceph



Figure 1. Image of the new Dell R740xd2, 2U rack systems with 24 drive density (12 front and 12 mid-plane all hot swappable). Other specifications include: 2 x Intel Xeon Gold 5218 2.3G, 16C processors; 12 x 32GB RDIMM 2666MT memory; 24 x 14TB 7.2K RPM SATA 6Gbps drives; Intel XXV710 Dual Port 10/25GbE network.

- ➔ Globus Endpoints
- ➔ CephFS -> NET2
- ➔ Block devices -> POSIX
- ➔ S3

- 24 PB Ceph
- 70% buy-in
- Rapidly expanding: 167 projects, 162 Pis, 112 organizations
- 6PB raw NET2, 8+3 EC, following Alastaire Dewhurst & RAL

**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 (MGHPCC) in Holyoke, Massachusetts.  
Photo: Helen HBM/GHPCC

### 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 (LISC) 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.

Big things are happening...

Exabyte scale shared storage



### NESE: Northeast Storage Exchange

National Science Foundation CIF21 DIBBs award 1753840



Boston University  
Harvard University  
Massachusetts Institute of Technology  
Northeastern University  
University of Massachusetts



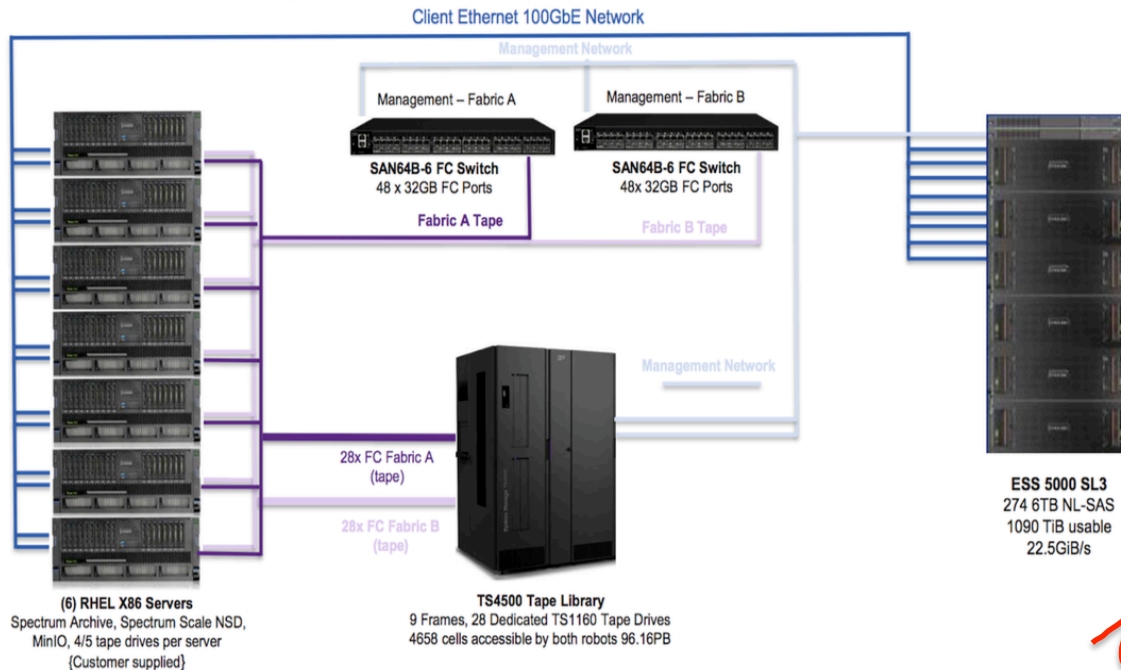


Organization	Hardware	Principal Contacts
MIT/CSAIL	Spectra Logic	Jonathan Proux and Greg Shomo
MIT/IS&T	Oracle	Nathan Thaler
BNL, LHC	Oracle, IBM	Xin Zhao and Tim Chou
SCC 2019	All vendors	Birds of a feather session
LANL	IBM and Spectra Logic	Julius Westerman
Fermilab	IBM	Bo Jayatilaka
Cambridge Comp.	All vendors	Jacob Farmer
NCSA	Spectra Logic	Jim Glasgow
LBNL	IBM	Nick Balthaser
TACC	Quantum	Frank Douma and Junseong Heo
MinIO	S3 software	CEO and CTO of MinIO Inc.

# NESE Exascale



## Spectrum Archive Architecture ESS



Nine frame IBM TS4500  
Library + ESS-5000 front end

133 PB capacity  
Max I/O: 11GB/s, scale up with drives.

Front end cache: 1 PB useable.


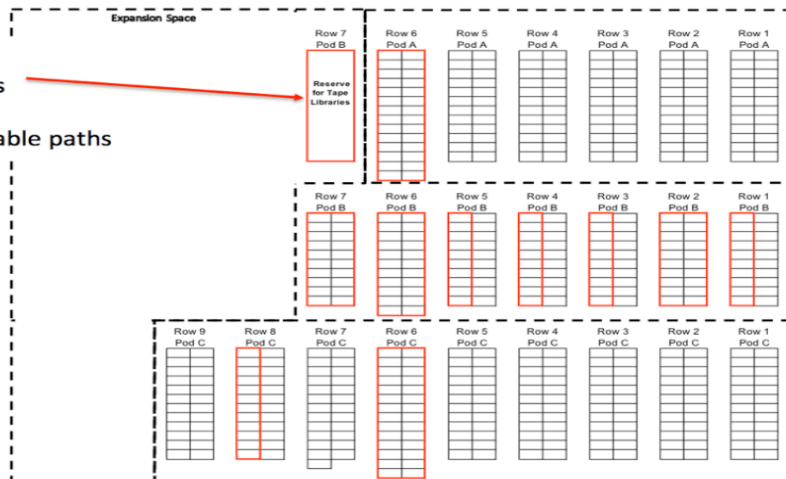
Protocols: POSIX and S3 via MinIO.

Software licenses and all maintenance  
paid for 5 years.

Extremely advantageous media buy-in  
price.

Immediate 50 PB buy-in from ATLAS.

Row 7, Pod A  
Space for four libraries  
Close to UPS power  
Away from "transit" cable paths

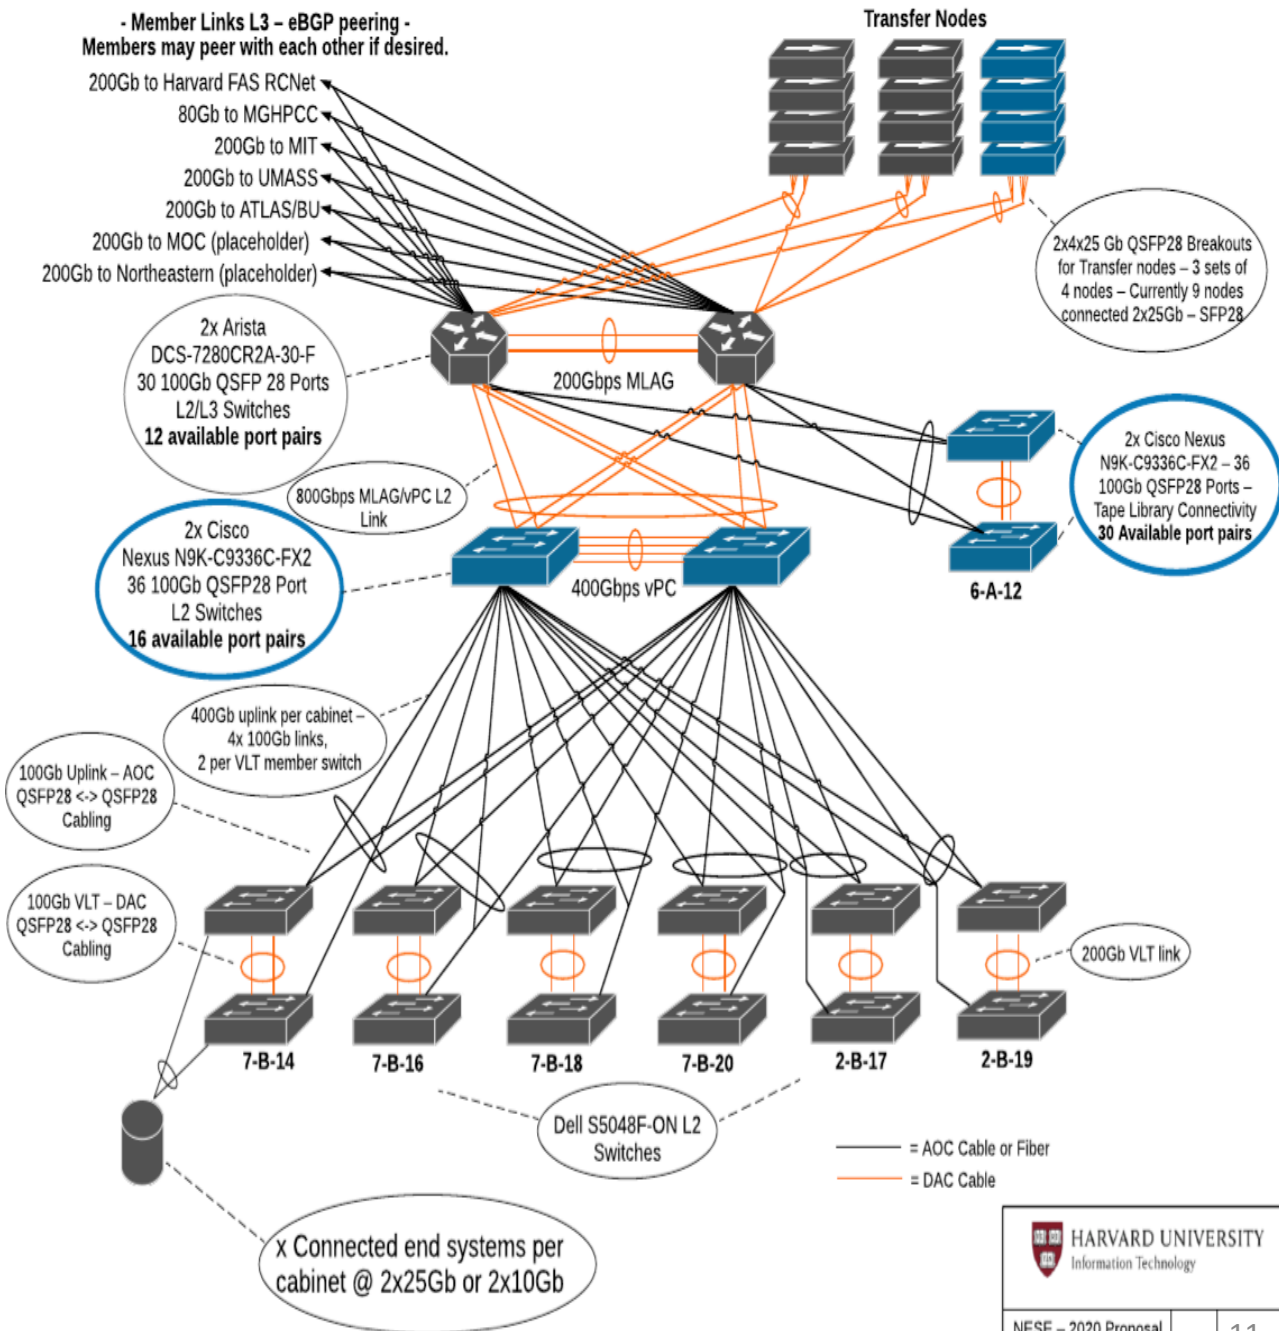




# NESE core network

## Expansion covers:

- 16 additional Ceph Racks
- 12 more projects/institutions
- Connects NESE Tape at 200 Gbps
- 6 NSDs POSIX/S3 via MinIO
- N x DTNs for HTTP-TPC



- NESE Exascale is happening
- 50 PB buy-in from ATLAS
- Working with the CMS/LHCb/Heavy Ion Tier 2 at MIT also
- We want to work with ADC, DDM, Carousel project, IRIS-HEP, WLCG...
- Ramping up to exascale by 2028 is quite feasible
- The strategy for NET2 is to migrate into much larger shared resources at MGHPCC

<http://nese.mghpcc.org>

<http://www.mghpcc.org>



