MGHPCC

MASSACHUSETTS GREEN HIGH PERFORMANCE COMPUTING CENTER

Massachusetts Green High Performance Computing Center

HPC UserForum

April 8, 2014

BOSTON UNIVERSITY · MIT · HARVARD · UMASS · Northeastern · CISCO · EMC² · COMMONWEALTH OF MASSACHUSETTS

The MGHPCC Data Center and Consortium

A partnership between 5 universities….


The Commonwealth, and industrial sponsors

Research Computing in Massachusetts

Medical imaging analysis

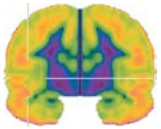

(BU, Harvard, Children's Hospital)

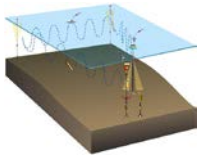

Ocean floor
(UMass, WHOI, MIT)

Understanding Evolution
(UMass)

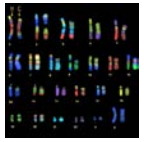

ATLAS/LHC

(BU, Harvard)

- **Hundreds of independent groups**
- **Diverse institutional cultures**
- **Frequent collaboration across organizations and disciplines**

Bicep2 / Gravity Waves
(Harvard)

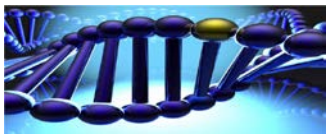

DNA samples -> Treatment Strategies
(UMass, MIT, Harvard Medical)

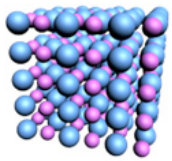

Computational Nanotechnology
(Northeastern, Harvard)

Origins of the MGHPCC

- Interest in institutional cooperation
  - Financial advantage
  - Opportunities that single institutions cannot address
  - Importance of university research to the Massachusetts economy
  - Meaningful impact on research competitiveness by state government

- Recognition that…
  - Computing is now fundamental to the university research mission
  - Incremental on-campus expansion is not a viable long term approach

Three Works in Progress

- Regional HPC Data Center
- University/Industry/Government Research Collaboration
- Regional economic development

MGHPCC Data Center

10 MW for compute / 15MW total

8 Acres and a 40MVA power feed for expansion

Green, low cost power

Communication Crossroads

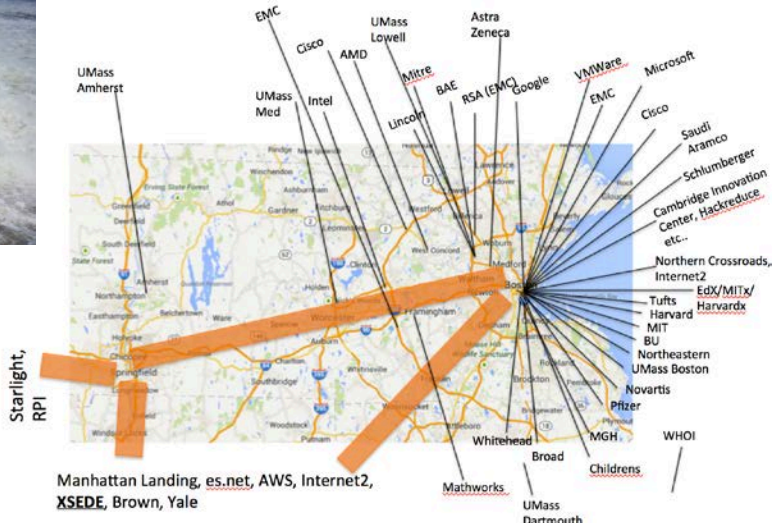

Minimizing Cost and Environmental Footprint

Minimize cooling load and chiller runtime

- Plate and frame heat exchangers 70% of the year
- Hot aisle containment
- No cooling in MEP areas

Minimize power distribution loss

- High Voltage distribution
- Eliminate a transformer tier

LEED Platinum Certification

- Dozens of details

Ownership and Operation of Computing Resources

Ownership		Management		Allocation
Research Project				
PI Coalition		Owner		Reservations
Department	X	Service Group	X	Job Scheduler
Campus/School		Outside Contractor		Virtual Machines
University-wide		etc		etc
University Coalition				

Four principal approaches

- Shared / Single-institution
  - University owned, fair-share usage
- Shared / Multi-institution
  - Jointly owned and operated by multiple institutions
  - Usage policy set by steering group
- Buy-in
  - Research groups contribute resources
  - Priority access to contributed resources; when-available access elsewhere
- Dedicated
  - Owned and used by a single group
  - Managed by the owner or a central service organization

Buy-in Owner Usage

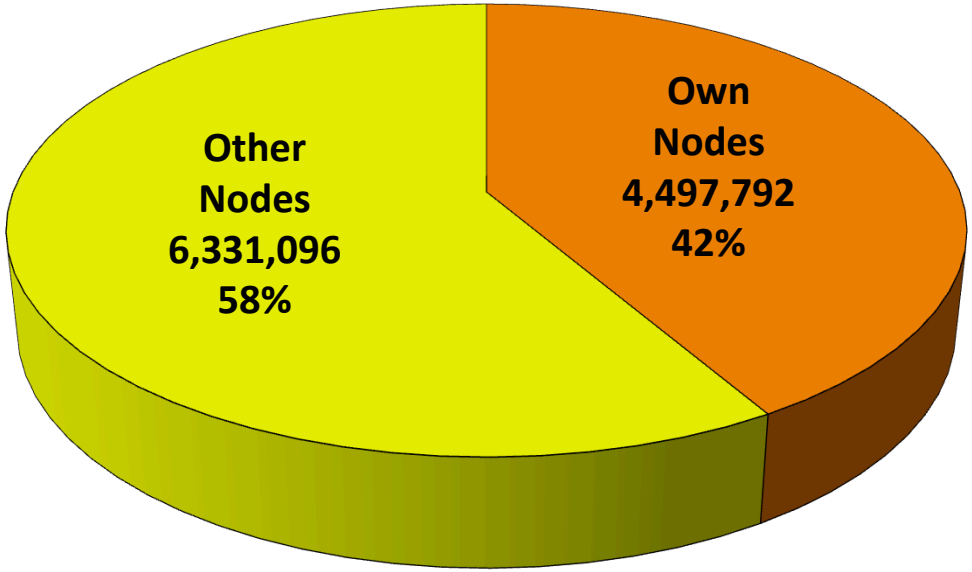

Buy-in Node Usage

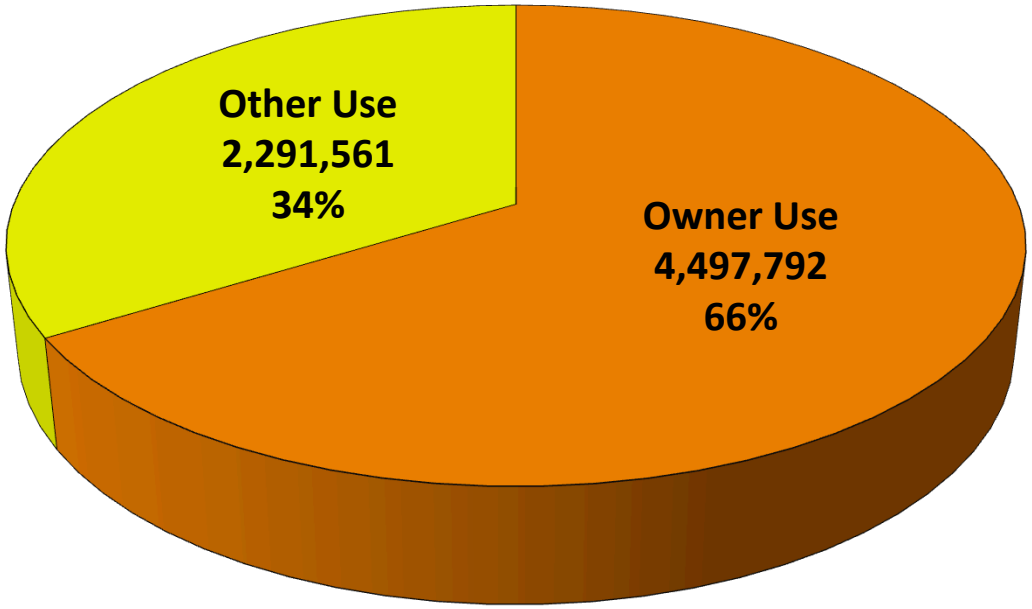

Diversity of Use

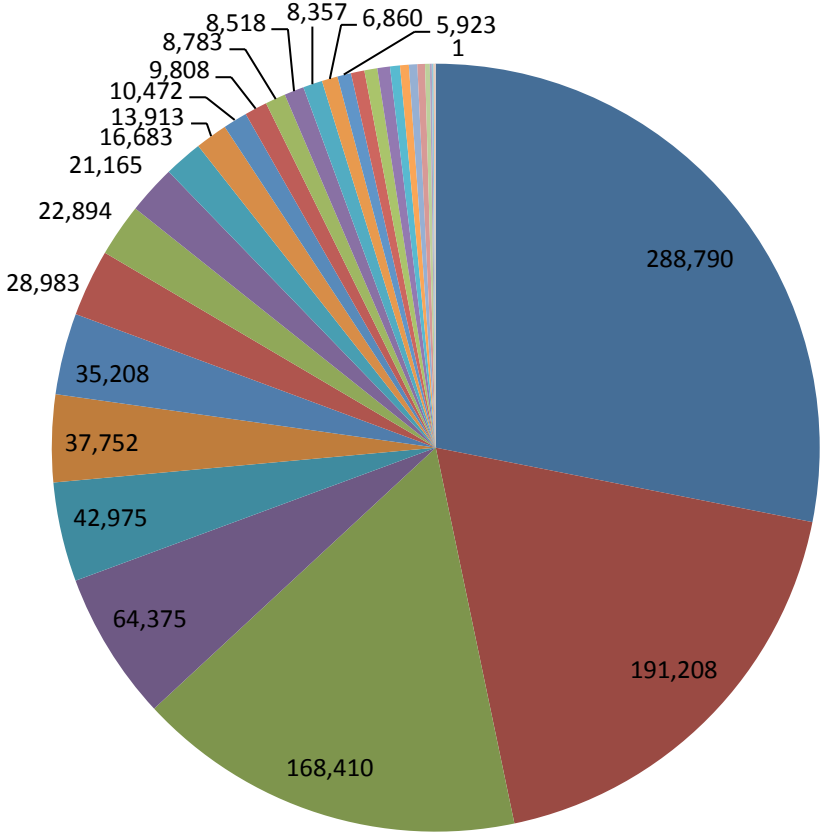

- Physiology & Biophysics
- Biomedical Engineering
- Chemistry
- Health Sciences
- Computer Science
- Earth & Environment
- Mechanical Engineering
- Astronomy
- Biomedical Genetics
- Training
- Information Systems
- Economics
- Mathematics & Statistics
- Biostatistics
- Physics
- Medicine
- Computational Neuroscience
- Finance
- Computational Biomedicine
- Biochemistry
- Biology
- Psychology
- Electrical & Computer Engineering
- Microbiology
- Neurology
- Administrative Sciences
- Bioinformatics
- Pharmacology

Security

- Facility: Enable compliance with common standards
  - HIPAA (Security Rule), NIST SP800-53 (PE controls), FERPA, MA CMR.17
  - protection, access control, managing sensitive information
  - Incident response, personnel

- Network: User-controlled interconnection
  - Every organization in the facility controls its own connectivity
  - “Neutral” space for shared systems
  - Meet Me switch for interconnections where needed

Model

A virtual building on each of 10 different campuses

Another virtual building for shared systems

Controlled, wire-speed connections between the virtual buildings

Governance

- Structure
  - Non-profit service corporation owned by five universities
  - Owns the facility and employs facility operations staff

- Oversight
  - Research mission and direction – Research Vice Presidents
  - Operational direction and goals – CIOs

- Day to day issues and opportunities
  - Research – Faculty working group
  - Operational – Business and IT working groups

Operational Collaboration

Operation and User Support

- Dedicated MGHPCC Staff focuses on facility operation
- User support comes from existing research computing groups
- Evolving Collaboration
  - Staying out of each others' way
  - Learning from each other
  - Unified networking plan
  - Coordinated vendor engagement
  - Acquisition and operation of jointly owned resources
  - Larger scale solutions to common problems

Research Collaboration

Seed Grant Program

Multi-university collaboration
Potential to lead to larger initiatives
14 projects over 2 years

engaging1::

Science via traditional HPC workflows and systems
Non-traditional workflows and technologies
Experiments with joint facilities integrated across universities

C3DDB

Dedicated to Life Sciences research
University and industry user community
Mix of conventional and forward-looking hardware

Closing Thoughts

- What we have done so far
  - **Strategic Infrastructure Investment**
    To accelerate research and education
  - **Collaboration Proof Point**:
    Pooling of resources across institutions can deliver powerful results
  - **Deliver Operationally**:
    >1 year of reliable, cost-efficient operation
- What lies ahead
  - **Make creative use of the resource**:
    Results that could not have occurred otherwise


MGHPCC
MASSACHUSETTS GREEN HIGH PERFORMANCE COMPUTING CENTER
BOSTON UNIVERSITY
UMASS
MIT
HARVARD UNIVERSITY
Northeastern
Bigelow Street
MGHPCC
P
Park