The Massachusetts Green High Performance Computing Center

High Performance Embedded Computing Workshop
September 22, 2011

John Goodhue

BOSTON UNIVERSITY · MIT · HARVARD UNIVERSITY · UMASS · Massachusetts It's all here. · EMC² where information lives · CISCO

MGHPCC

A partnership between 5 universities

With additional support from the commonwealth and industrial sponsors

Computing as “Third Leg” of Scientific Research

Run virtual experiments

Make sense of vast amounts of data

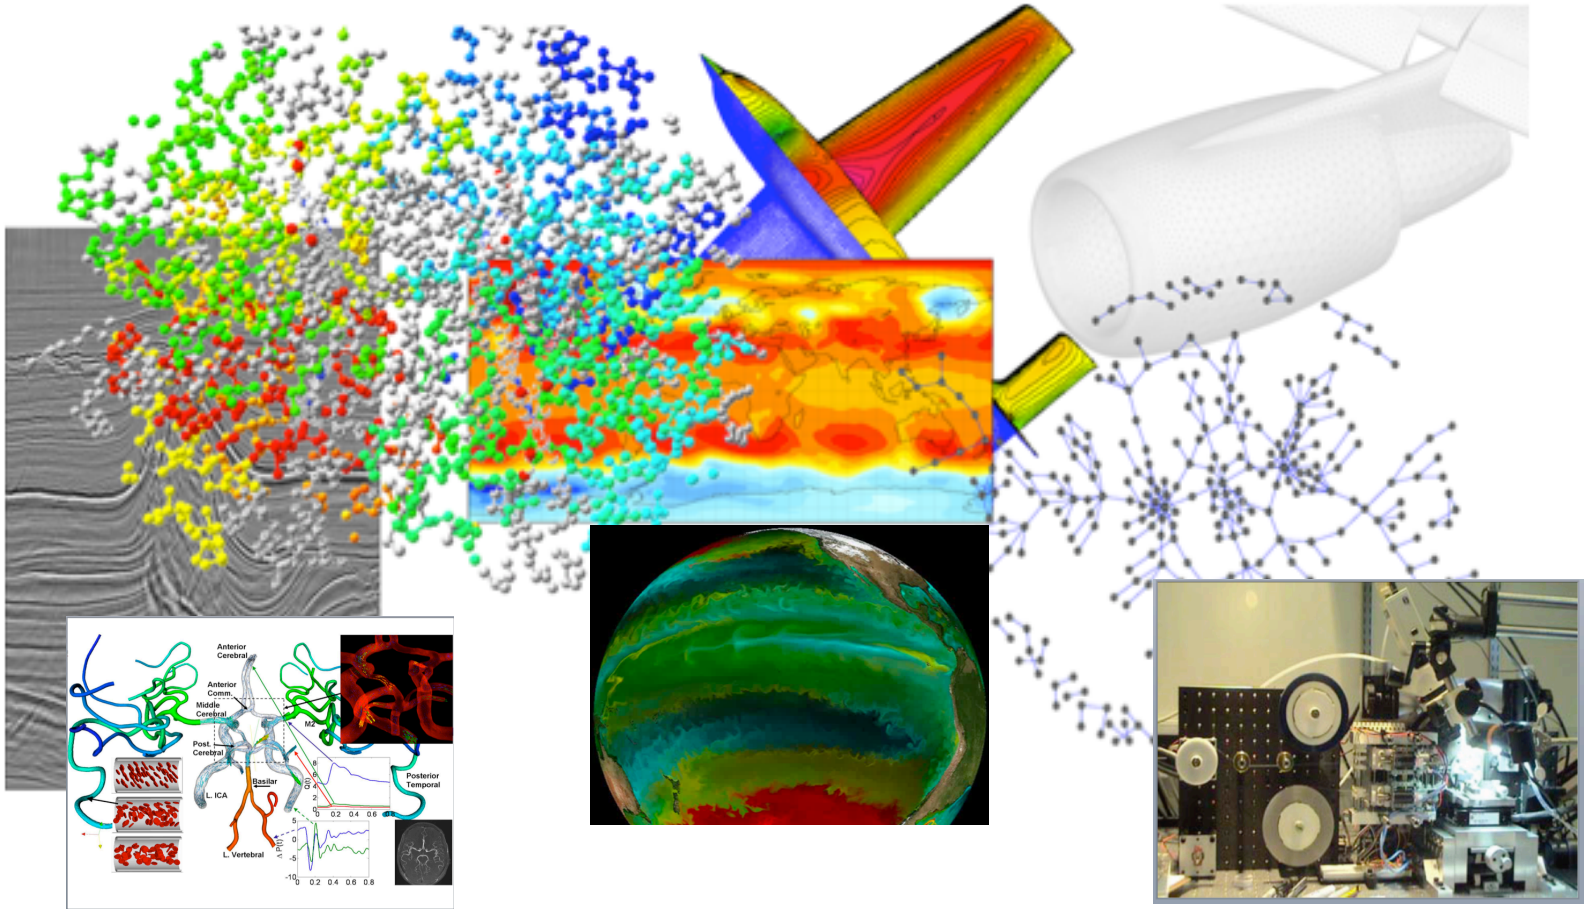

Time Line

Inception (Aug-Nov)
- Funding agreements
- Site purchase agreement
- Environmental testing
- Basis of Design
- Construction Manager Contract

Design and Planning (Nov-Aug) + Site Remediation (Feb-Oct)
- Building and IT design
- Business model and operating agreements
- Environmental remediation and building demolition
- Site acquisition

Construction and Move-in Planning (Aug-Nov)
- Implement construction plan
- Prepare for first wave of installations

We are Here

Move In (Q412)

Two Goals of the MGHPCC

Build and Operate a Data Center for Scientific Computing

Encourage Novel and Collaborative Computationally Intensive Research

MGHPCC Consortium

- Collaborative effort to strengthen R&D leadership
  - Most significant state/industry/university partnership in MA history
  - Focus on computationally intensive research
  - Leverage the data center as a research tool
  - Stimulate collaboration in other areas (e.g., cybersecurity)

- Research initiatives not achievable by any single institution
  - NSF STCs, ERCs, CISE Expeditions, DOE EFRCs, innovation-hub-like
  - An opportunity industry partners will want to participate in

- Catalyst for economic & workforce development
  - In the state, region and Holyoke (Innovation District)

Design of a Green Computing Facility

BOSTON UNIVERSITY

HARVARD UNIVERSITY

UMASS

Massachusetts It's all here.

EMC² where information lives

CISCO

The MGHPCC Approach to Green Design

- Pooling resources for higher efficiency
- Locate near clean/inexpensive energy and high bandwidth
- Improve an existing industrial site
- Optimize building and computer room layout
- Low loss power distribution
- Maximize use of outside air for cooling
- Pay attention to everything, including administrative space
- Slightly out of the box ideas

Pooling Resources for Higher Efficiency

- Reduce operating cost
- Improve and expand access to research computing infrastructure

Less of this...

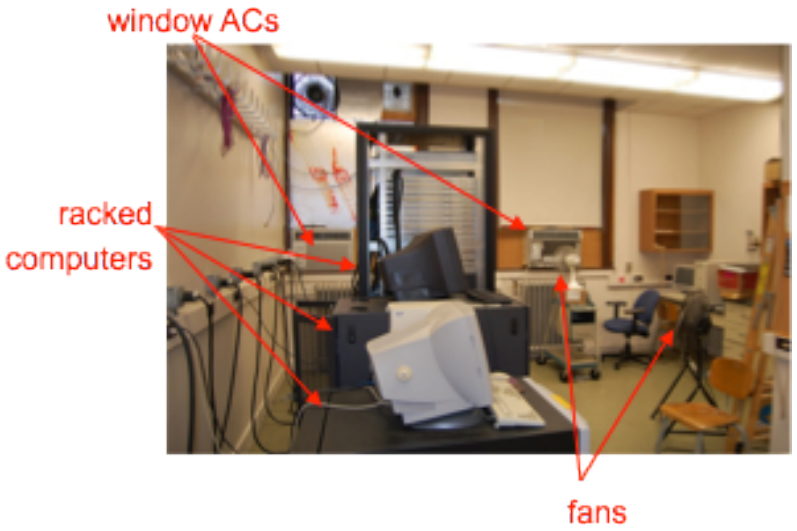

More of this...

Building a Community Cloud

Current Examples

Research Project
- BU Atlas

PI Coalition
- MIT Bates

Department
- MIT CSAIL

Campus/School
- UMass Medical
- Harvard FAS
- Harvard Medical
- BU Medical

University-wide
- Purdue
- Princeton
- Stanford

University Coalition
- MGHPCC facility
- MGHPCC Track II proposal
- MGHPCC MRI proposal

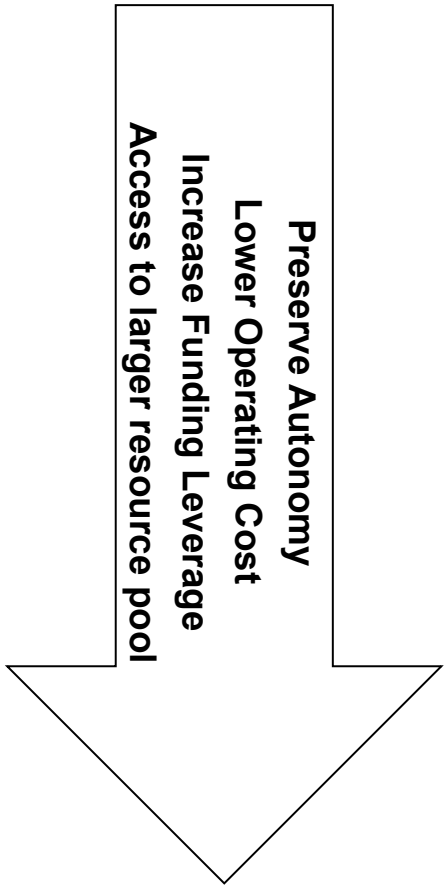

Goals

Preserve Autonomy / simplify operation
- Local administration and application support
- Faster install
- Space/power/cooling/security taken care of

Lower Operating Cost
- Leveraged capital investment
- Locate for low cost of power
- Design for cooling efficiency
- Amortize staff cost/expertise

Greater Funding Leverage
- Pooling of research grant funds
- Greater leverage for university subsidies
- Higher impact for major infrastructure grants

Access to larger resource pool
- Control access to what you own
- Opportunity to share idle resources

Ownership and Management 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				

The MGHPCC is prepared to support all of these usage models

Why Holyoke?

- Existing industrial site
  - Green, low cost power

- Fiber Crossroads

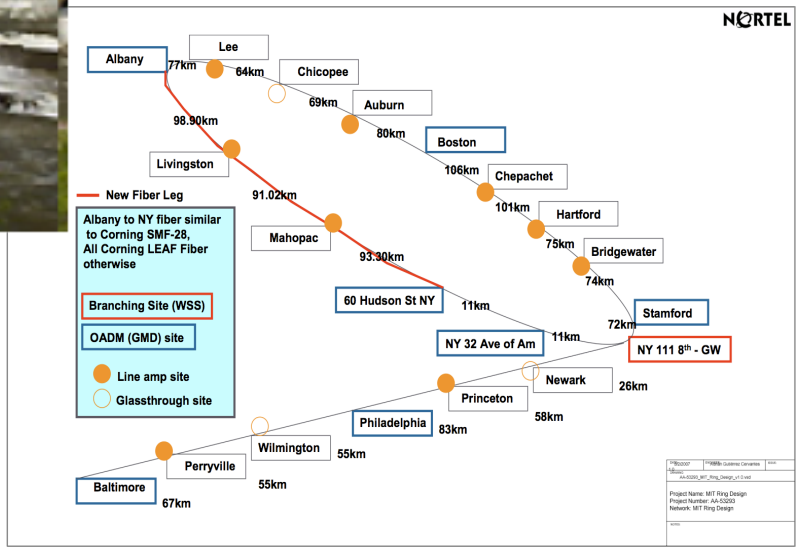

MGHPCC Site (8.6 acres total)

Computer Center
(90,300 Square Feet)

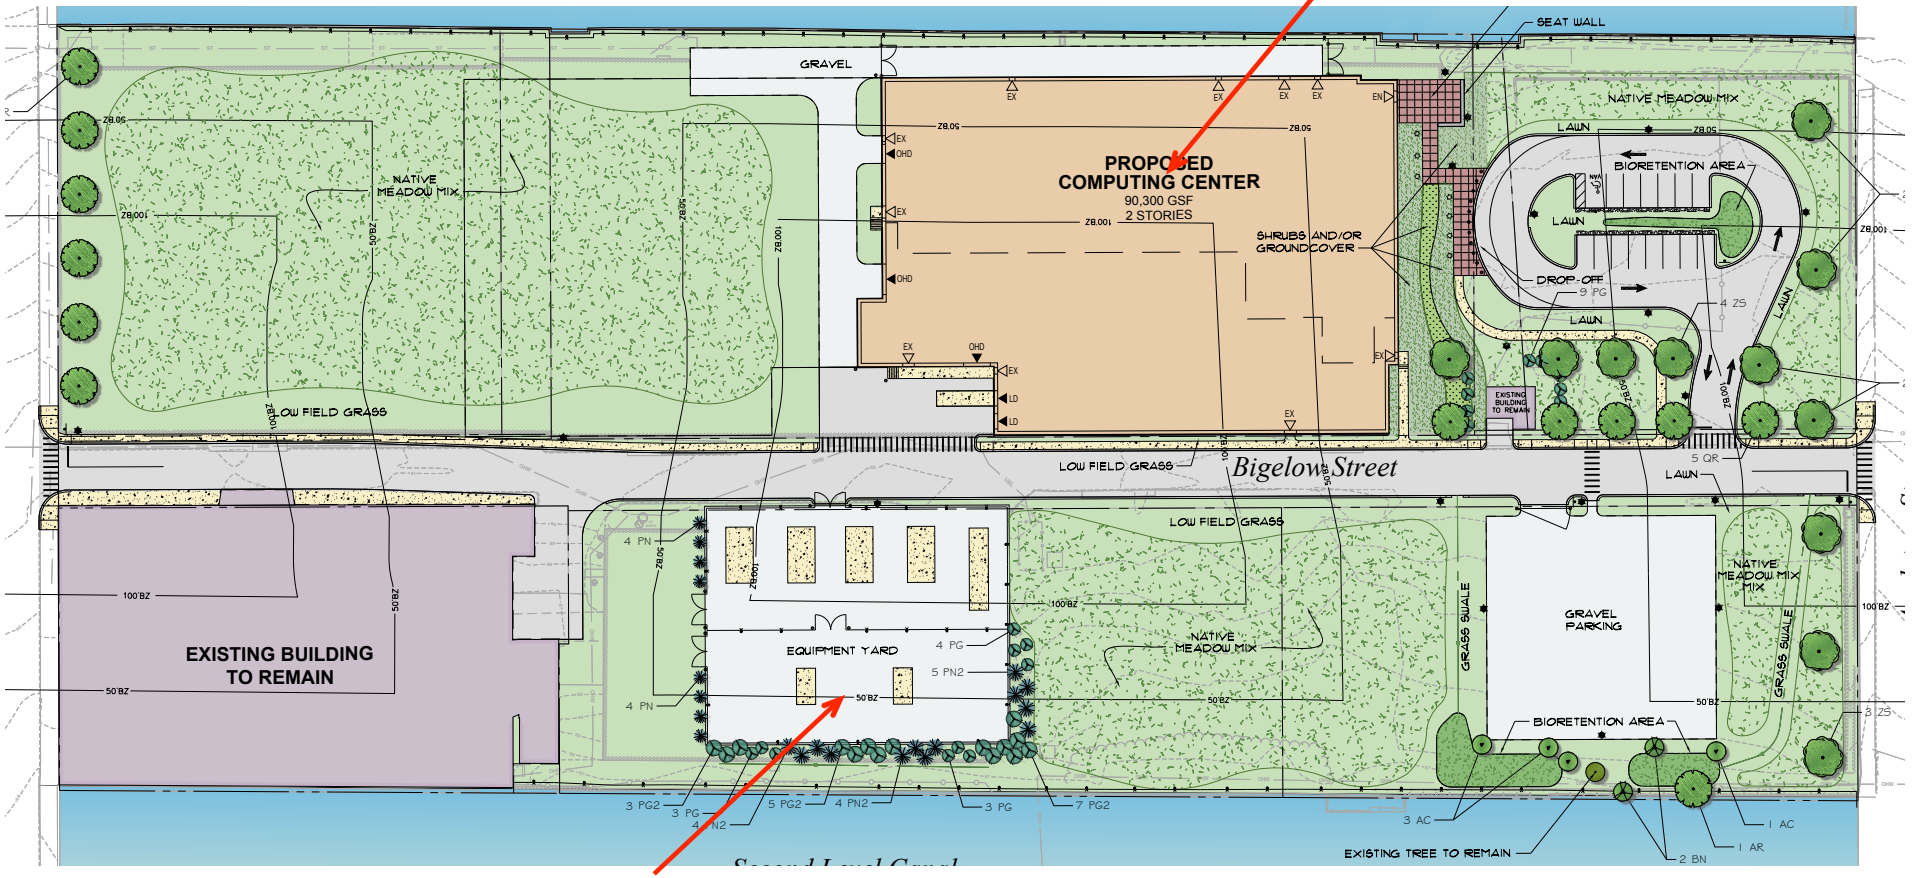

Transformer Yard
- **15MW Initial capacity**
- **15MW expansion**

MGHPCC Facility

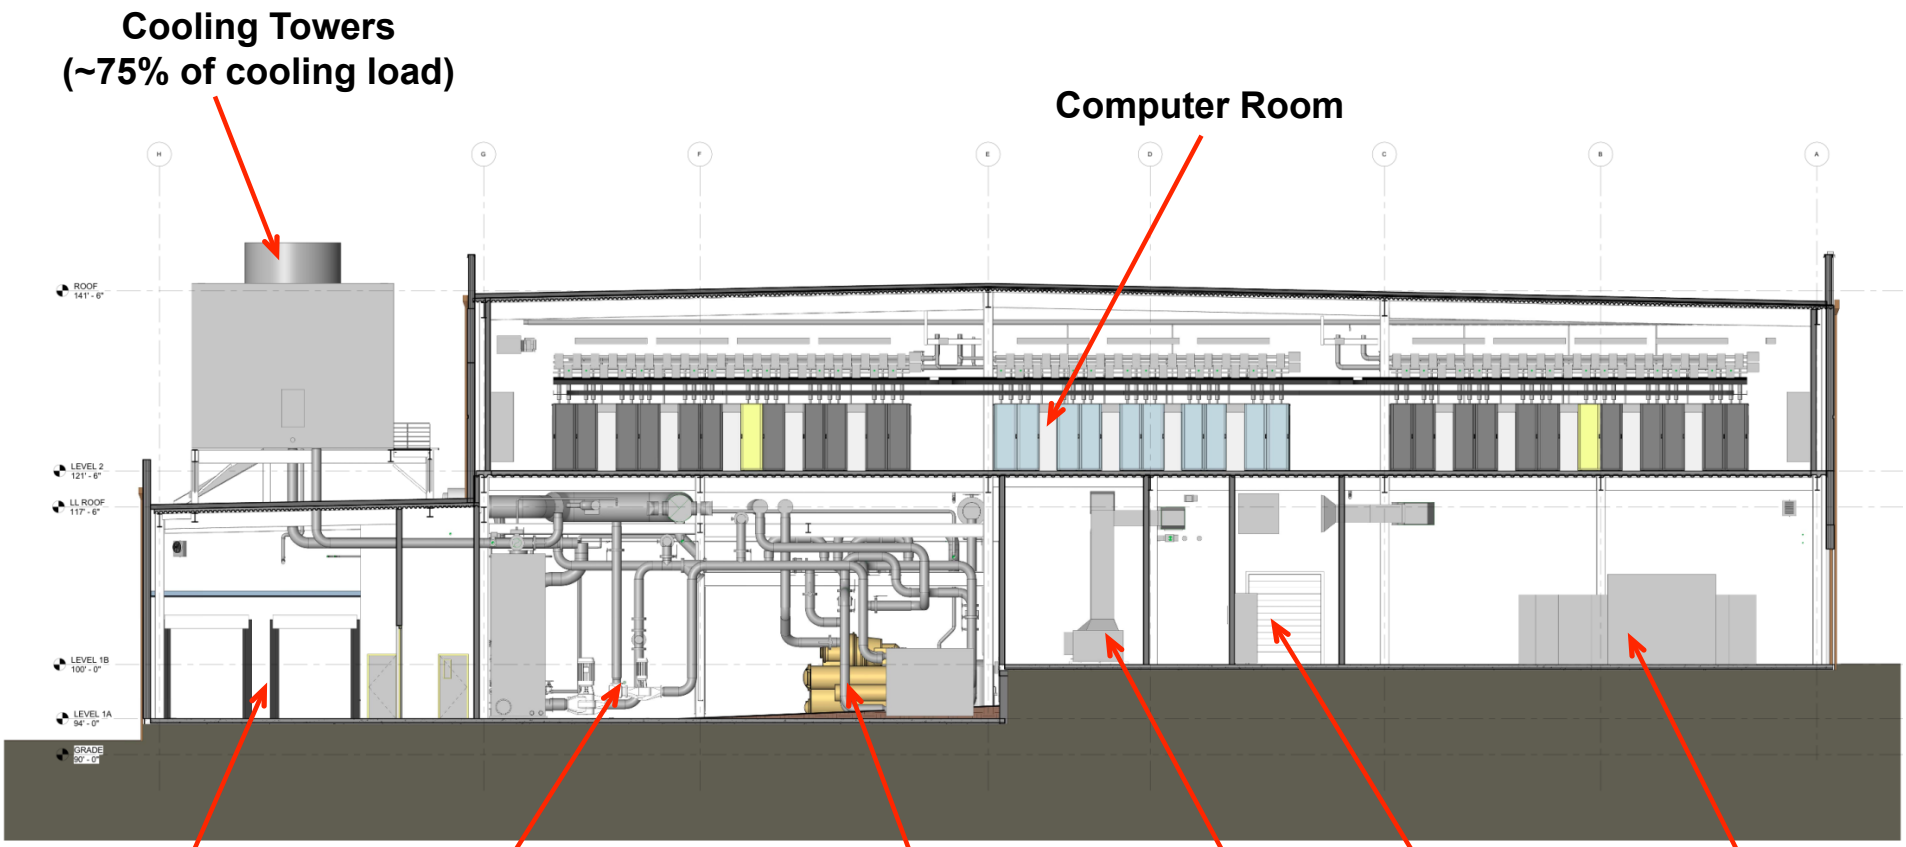

Computer Room

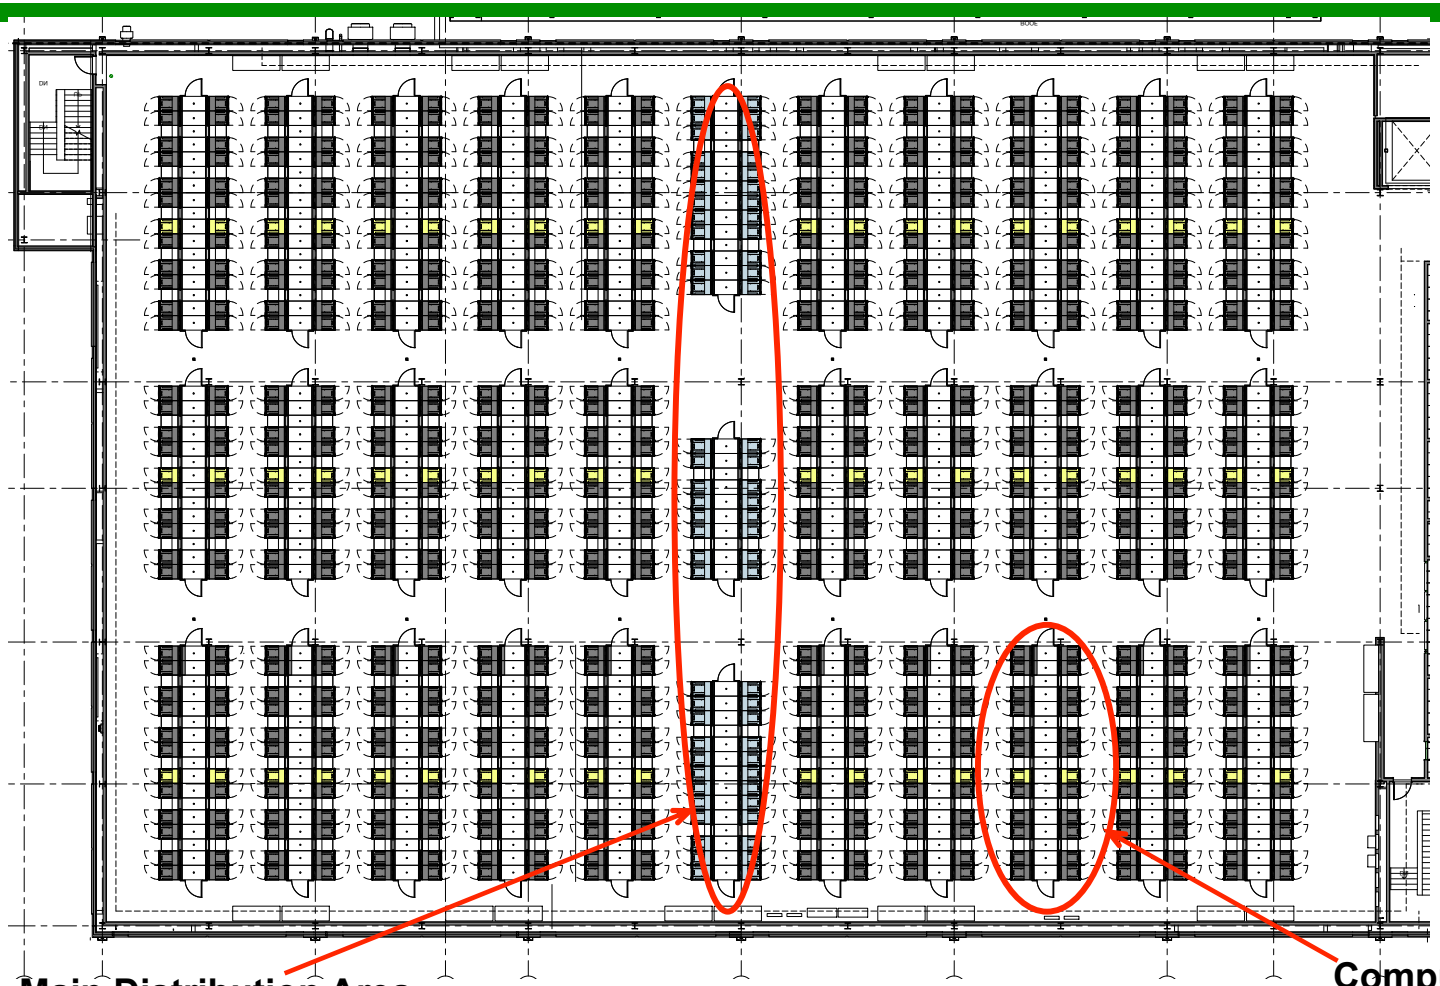

Space
- 33,500 sq ft
- 680 racks

Power and Cooling
- 10 MW for compute
- 5 MW for cooling
- Target PUE: 1.3

Network
- 10 Gb/s feeds to every campus

Main Distribution Area
- University Fiber Feeds
- University backbone switches
- Meet me switch

Compute Pod
- 20 or 24 Racks
- 14kW per rack average
- N x 10 GigE to the MDA

Low Loss Power Distribution

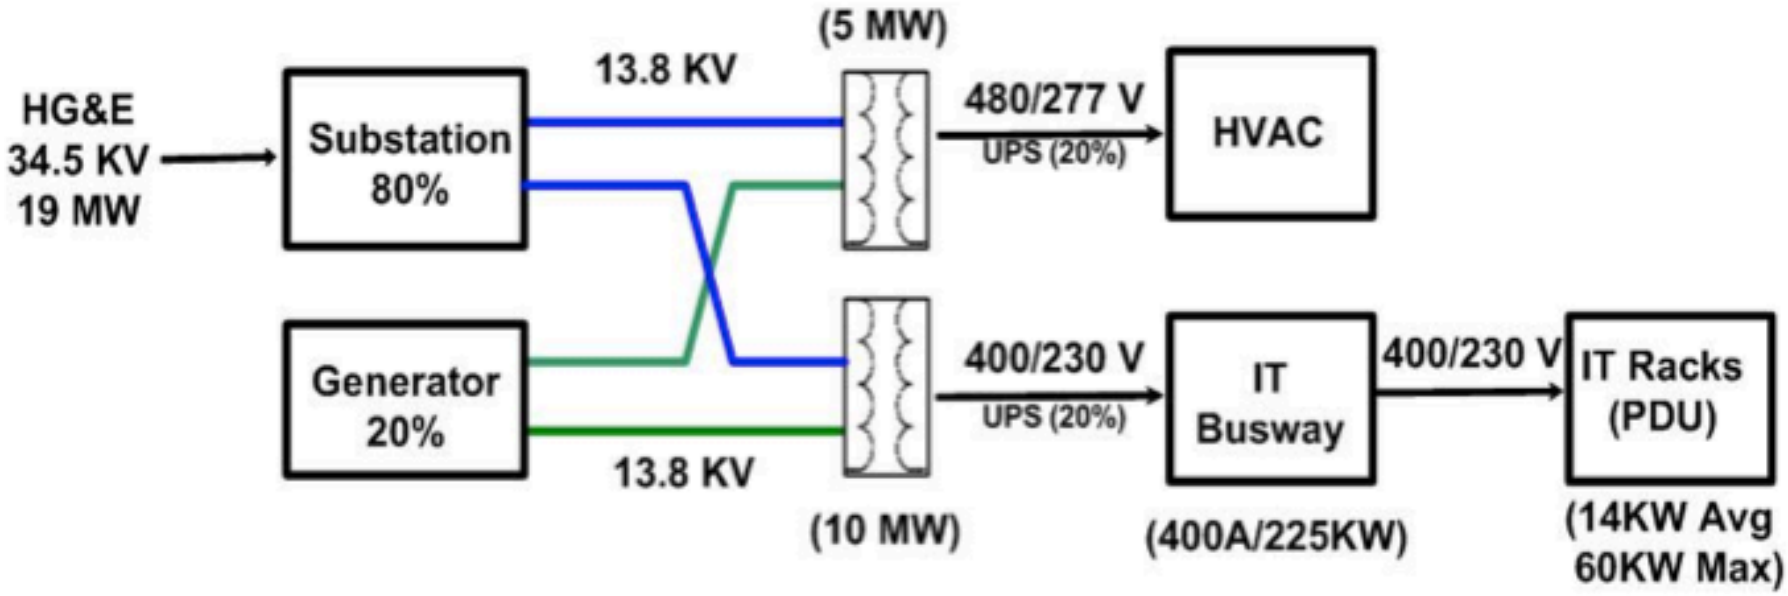

Pod Design

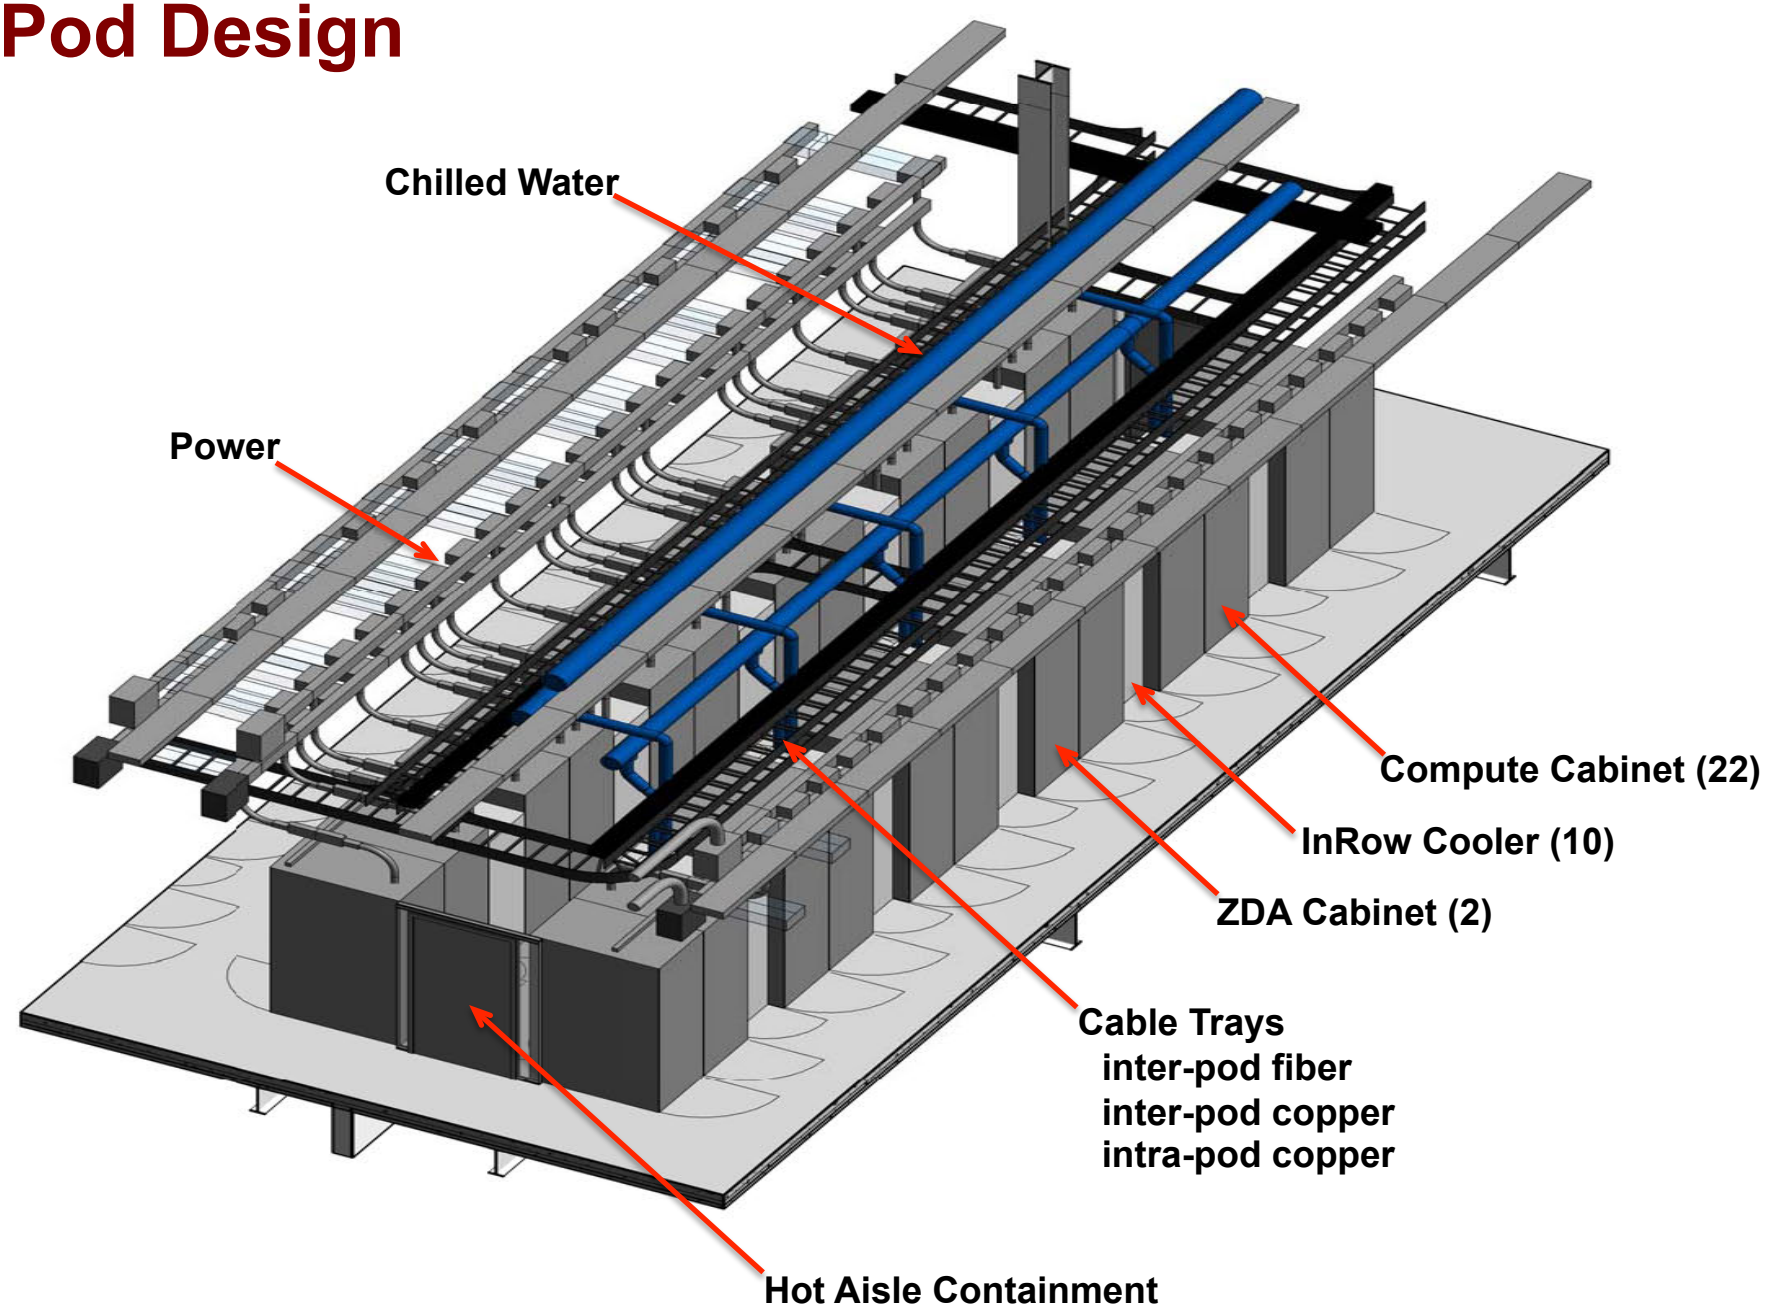
Chilled Water
Power
Compute Cabinet (22)
InRow Cooler (10)
ZDA Cabinet (2)
Cable Trays
inter-pod fiber
inter-pod copper
intra-pod copper
Hot Aisle Containment

Maximize Cooling Efficiency

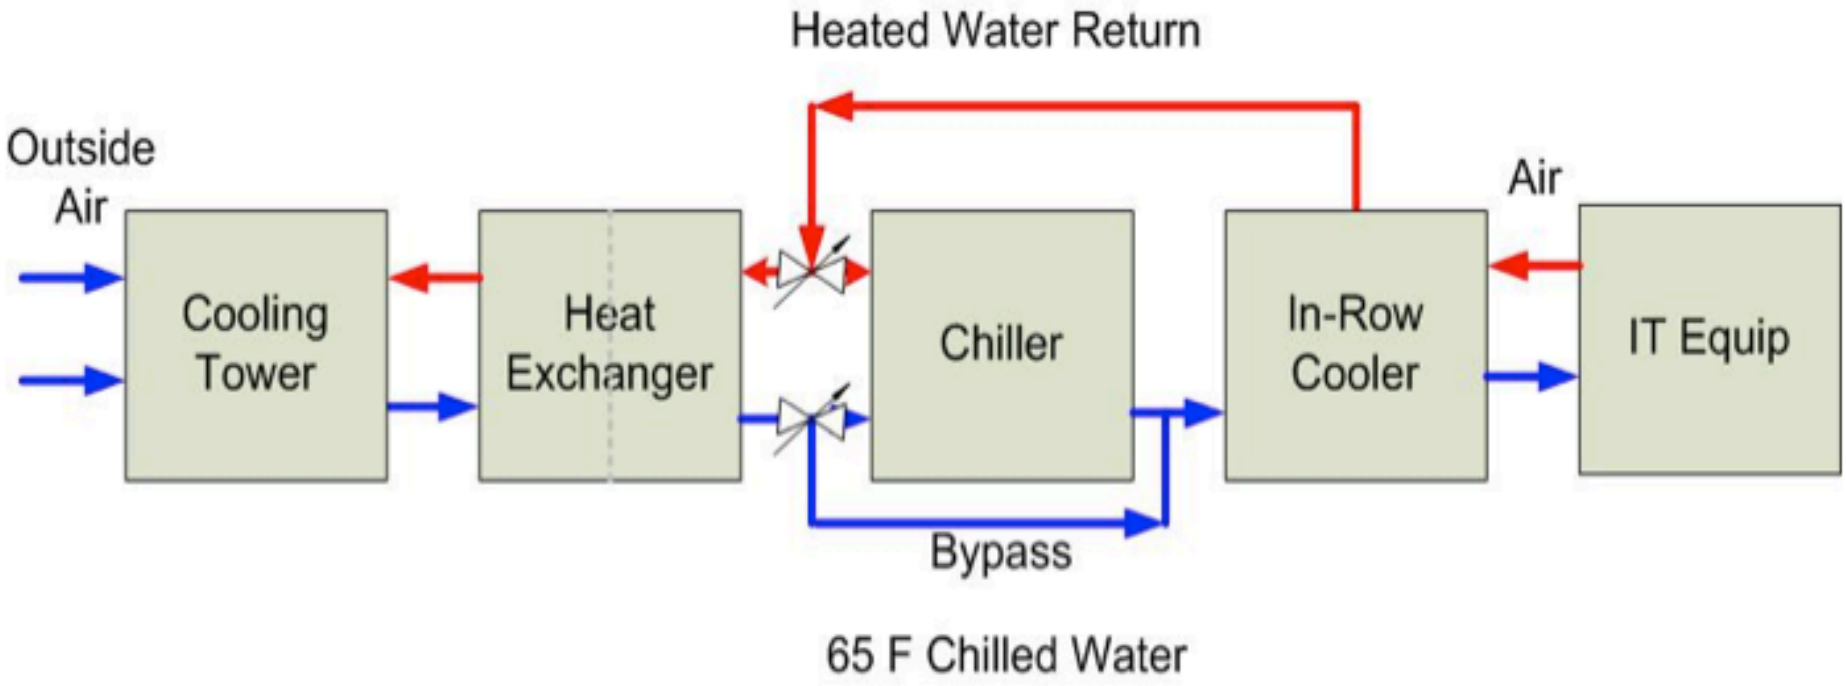

Cooling Modes of Operation

Economizer

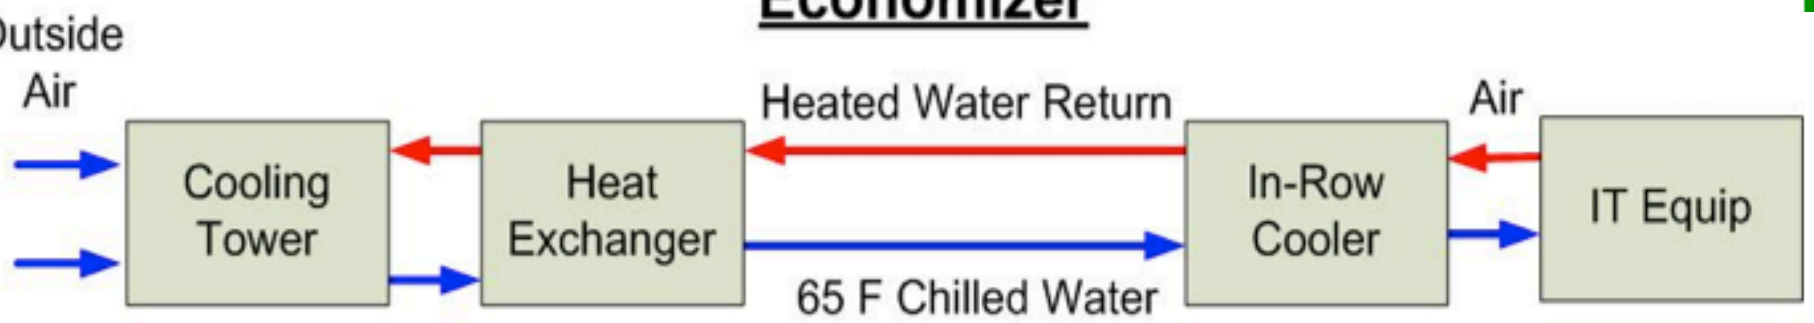

Precooling

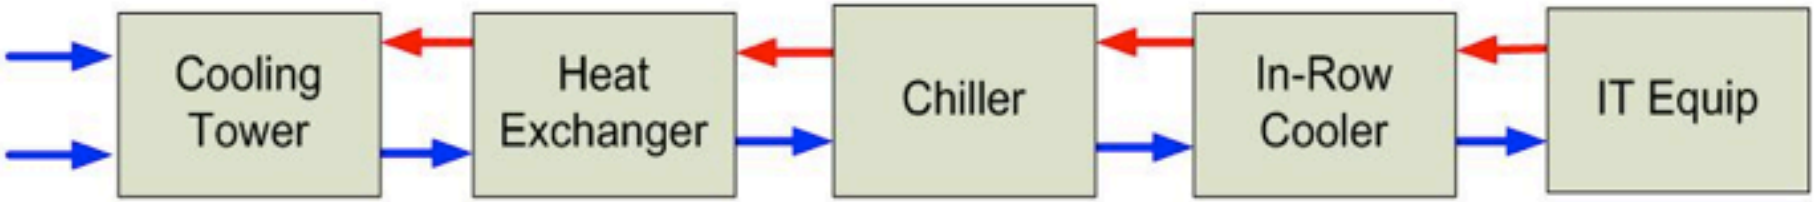

Chiller

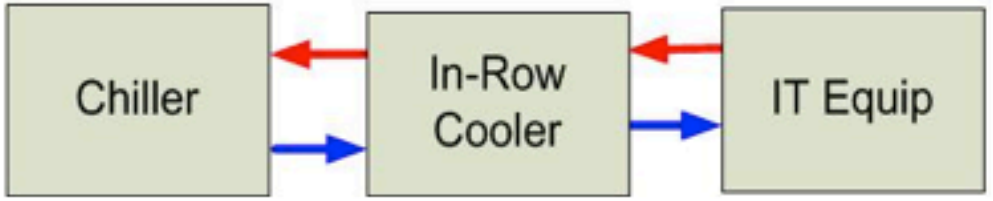

Administrative Space

- LEED certification
- Under-floor air distribution
- Low fan pressure in the air handling units
- Waterless/low flow bathroom fixtures
- Lighting layout

Slightly Out of the Box Ideas

- Canal water for cooling
- Energy storage for load shifting
- Reuse waste heat
- On site generation

Challenges Ahead

- Tuning facility operation for maximum efficiency
- Balancing flexibility and energy efficiency

Some Observations

- Pay attention to everything
- Observe the 80/20 rule
- Change is a constant concern
  - Must reconcile decade-scale plant lifetimes with Moore's law
- But there is much that stays relatively constant
  - Physics (Specific heat of air and water, conductivity of Copper, etc)
  - Cost, weight, and size of a fully loaded standard rack

Questions?

BOSTON UNIVERSITY

HARVARD UNIVERSITY

UMASS

Massachusetts It's all here.

EMC² where information lives

CISCO