Data Centers, Cloud and HPC Optimization

Daniel A. Reed
Vice President for Research and Economic Development
University Computational Science and Bioinformatics Chair
Computer Science, Electrical Engineering & Computer Engineering, and Medicine

THE UNIVERSITY OF IOWA

www.hpcdan.org

Disclaimers

Everything I am about to say is public information …

The opinions are mine (and I may disavow them later …)

Energy efficiency is a multivariate challenge

- Historical practice and conventional wisdom
- Culture and normative behavior
- Economics and social constraints
- *Science and technology (this is the easy part)*

Big change requires custom design and culture change

The data center energy efficiency problem is already (mostly) solved

- *Stop sweating PUE and focus on other things*

Orders of magnitude change outcomes

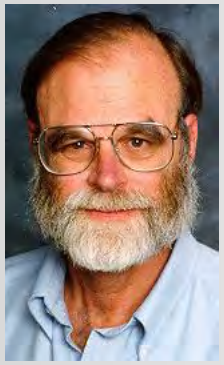

An insight from the late Jim Gray …

These are systemic problems

A computation task has four characteristic demands

Networking	Computation	Data access	Data storage
Delivering questions and answers	Transforming information to produce new information	Access to information needed by the computation	Long term storage of information

The ratios among these *and their* costs change the possible, technically, economically and politically

What's a cloud data center?

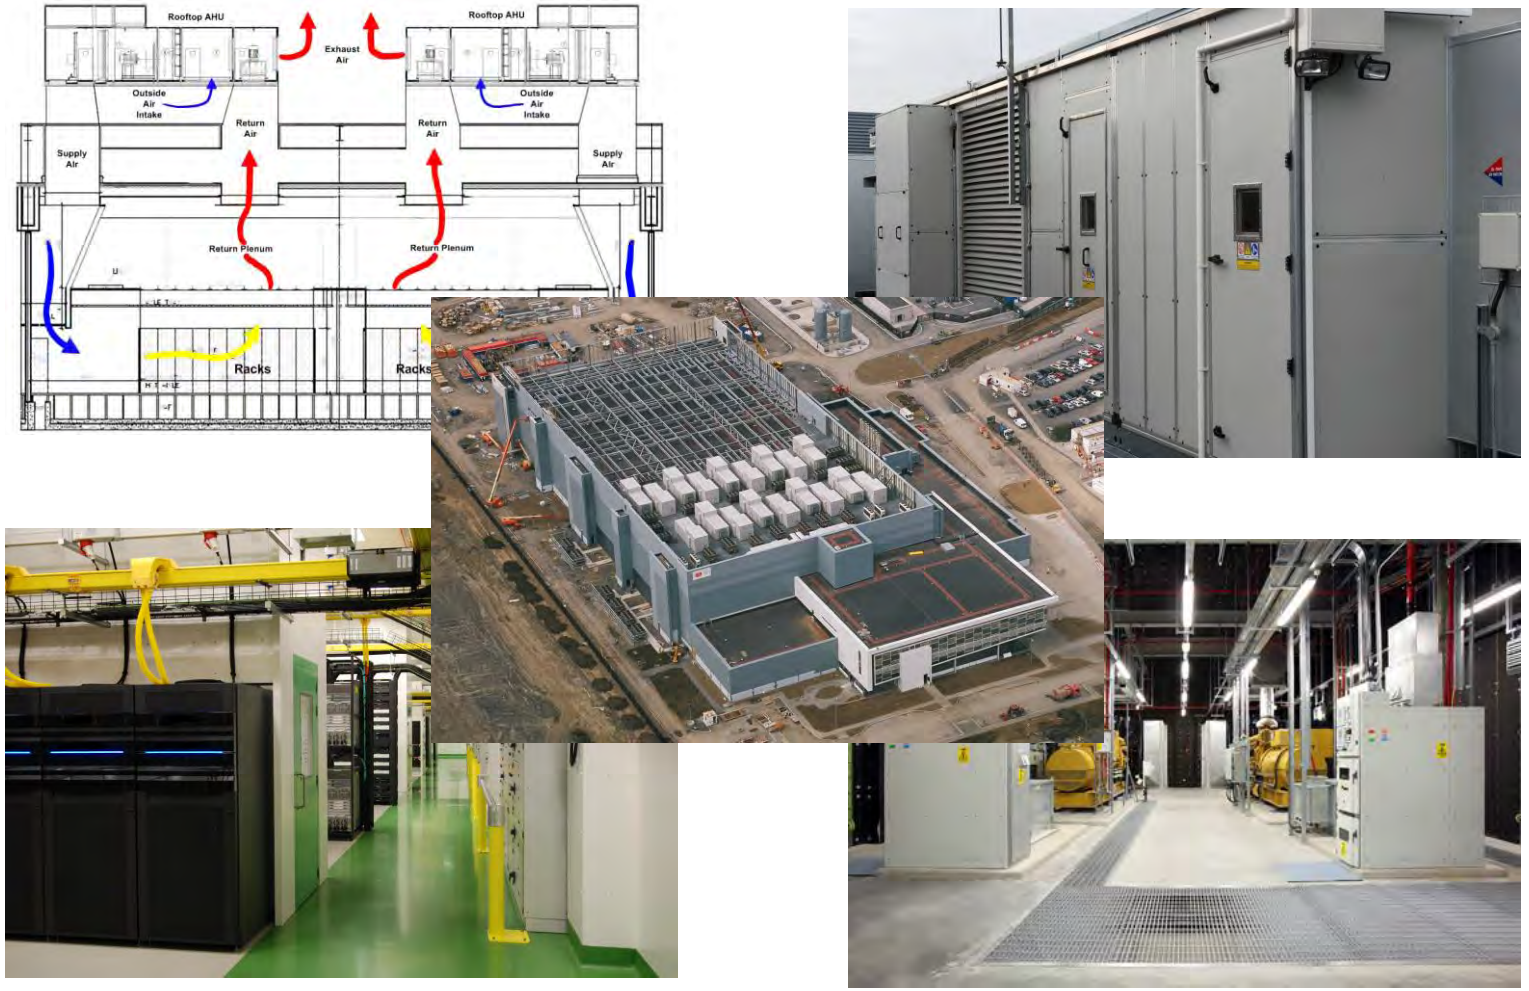

Microsoft Dublin: Gen2/Gen3 data center

Many optimization axes

- Power (availability/cost)
- Bandwidth
- Regulatory and tax structure
- User base
- Stability (geo and political)
- Time value of money

Open Compute vs. Custom

Commoditizing infrastructure

- Servers, storage and networks
- Data centers and facilities
- Management services

OPEN Compute Project

Specs & Designs News Community

Hacking Conventional Computing Infrastructure

We started a project at Facebook almost two years ago with a pretty big goal: to build one of the most efficient computing infrastructures at the lowest possible cost. We decided to honor our hacker roots and challenge convention by custom designing and building our software, servers and data centers from the ground up – and then share these technologies as they evolve.

The result is a data center full of vanity free servers which is 38% more efficient and 24% less expensive to build and run than other state-of-the-art data centers.

Why Open Hardware?

By releasing Open Compute Project technologies as open hardware, our goal is to develop servers and data centers following the model traditionally associated with open source software projects. That's where you come in.

Click here for talk and video archives from the May 2012 summit in San Antonio, TX

Join a Mailing List

NOVEMBER 15, 2013, 3:11 PM | 8 Comments

Amazon Bares Its Computers

By QUENTIN HARDY

Amazon Web Services

An Amazon Web Services data center. A company executive said its data storage system handles 1.5 million requests a second, and holds trillions of "objects," or individual stored items.

FACEBOOK
TWITTER
GOOGLE+
SAVE
E-MAIL

LAS VEGAS — However big and ambitious you think Amazon's plan to run the world's computing may be, you should probably think bigger.

In a startling talk Thursday evening, a vice president who oversees the internal engineering of Amazon's global computing system described how Amazon is building its own specialized computers, data storage systems, networking systems, even power substations and optical transmissions systems. In every case, he said, Amazon Web Services had developed ways

From OEMs to ODMs

- Lower margins and flexibility
- Focusing on differential value

THE UNIVERSITY OF IOWA

Microserver futures

Two competing ecosystems

- ARM and x86
- Different cultures and business models

SoCs and mass specialization

- Integrated networks
  - Electrical, then silicon photonics
- Accelerators
- Stacked DRAM (with capacity implications)
- 3-D memory stacks
- PCM and beyond

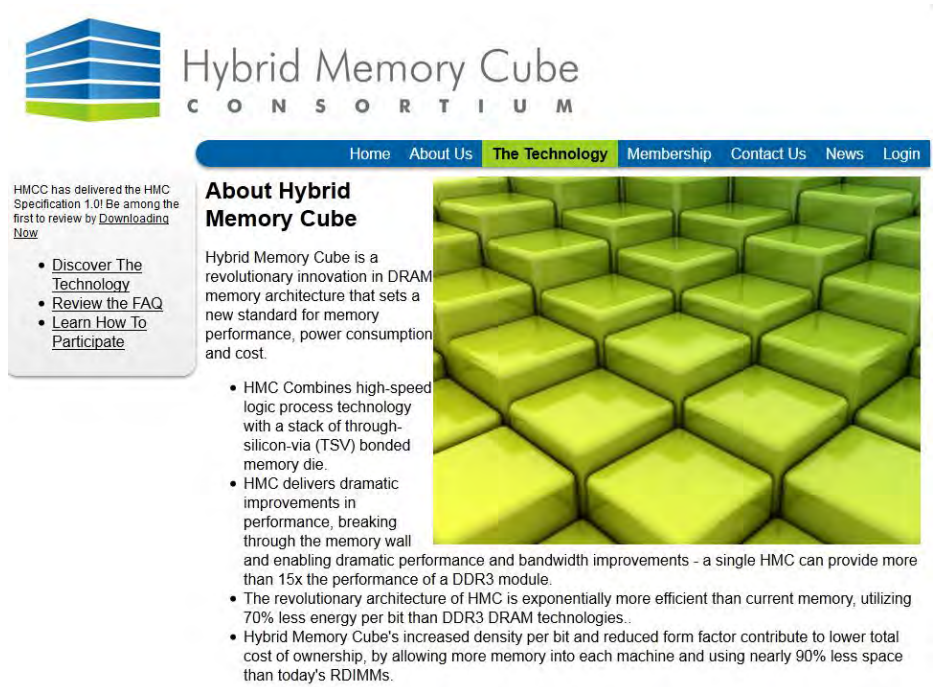

Driving trends

- Data movement energy costs
- Thermal dissipation and dark silicon
- Memory bandwidth constraints

Rethinking computing energy

Multiple energy sources
- Electrical grid, solar, wind, fuel cell, …

Multiple cost functions
- Energy pricing, carbon taxes, varying availability
- Hardware, data transfer bandwidth/latency …

Multivariate optimization and prediction
- Workload demand
  - Diurnal and seasonal
- Weather and seasonal models
- Auction-based energy pricing
- Infrastructure
  - UPS, optical fiber and computing

Scheduling subject to energy and reliability
- Cost, availability, resilience …

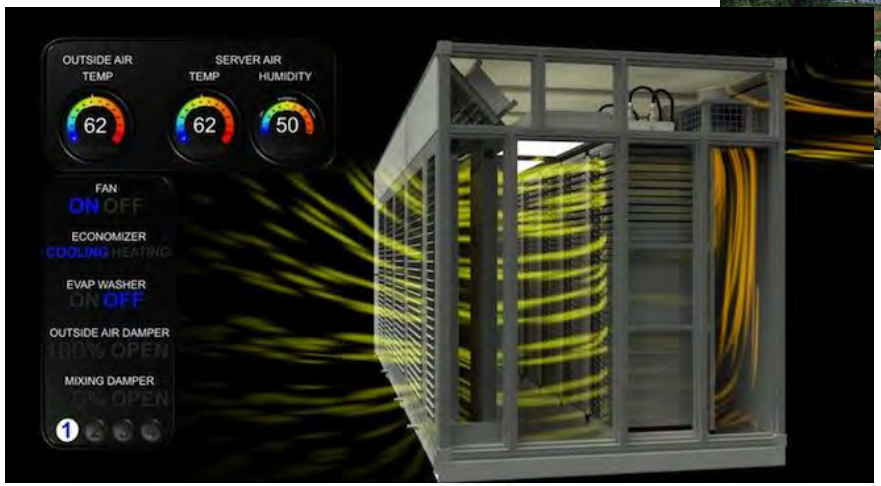

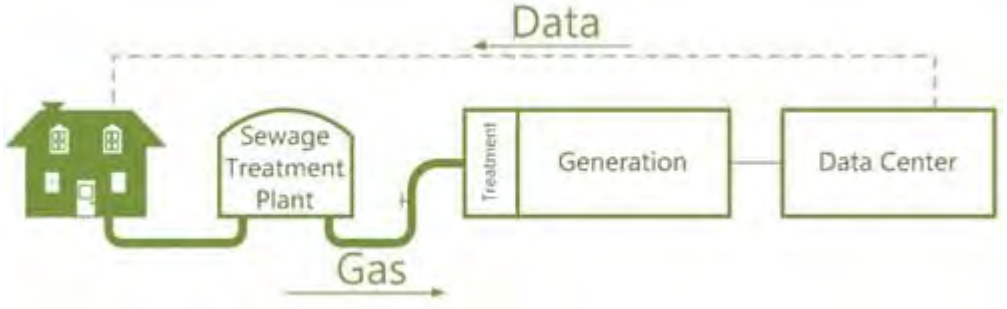

Microsoft Wyoming biogas prototype

THE UNIVERSITY OF IOWA

Big lessons

Supply chain optimization
- The advantage of scale

Specialized server design
- Workload specific optimization
- ODM, not OEM partnerships
- Functional accelerators

Network optimization
- Flatter networks
- Software virtualization and flow

Energy optimization
- Substations and generation
- Switchgear control

Systemic resilience
- Failure management, not avoidance

Looking forward: services appliances

National sovereignty will persist …

- … and a globalization privacy backlash is building

with deep implications

Transnational data flows

- Rising concerns around the world
- Business, sovereignty, protectionism and social issues

Implications

- "Zero touch" data appliances, operated by locals for locals
- Mega data centers for certain business and other social networks
- Local and regional policy frameworks

Culture and economics

Industry

- Capital is cheap (look at interest rates)
- Labor is expensive
- ROI drives behavior

Academia and government

- Capital is expensive
- Labor is cheap
- Other metrics drive success

To change the game, change the metrics …

- Infrastructure, personnel, social and political

Put another way, where you draw the bounding box shapes the answer

Discussion