

# Cloud Computing Is Driving Infrastructure Innovation

## Stanford Computer Forum

James Hamilton, 2011/4/12

VP & Distinguished Engineer, Amazon Web Services

email: [James@amazon.com](mailto:James@amazon.com)

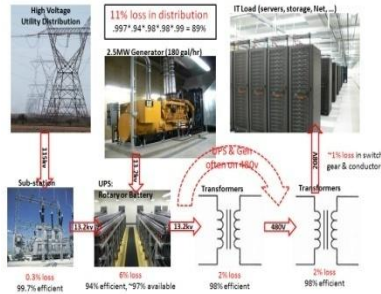
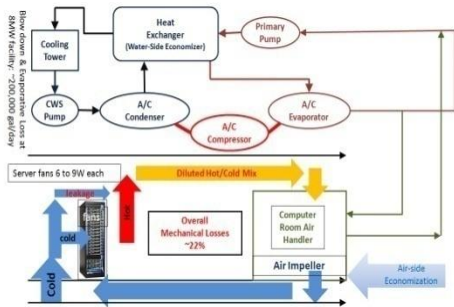

web: [mvdirona.com/jrh/work](http://mvdirona.com/jrh/work)

blog: [perspectives.mvdirona.com](http://perspectives.mvdirona.com)



# Agenda

- Quickening pace of infrastructure innovation
- Where does the money go?
- Power distribution infrastructure
- Mechanical systems
- Modular & Advanced Building Designs
- Sea change in networking



Talk does not necessarily represent positions of current or past employers



# Pace of Innovation

- Datacenter pace of innovation increasing
  - More innovation in last 5 years than previous 15
  - Driven by cloud service providers and very high-scale internet applications like search
  - Cost of infrastructure dominates service cost
  - Not just a cost center
- High focus on infrastructure innovation
  - Driving down cost
  - Increasing aggregate reliability
  - Reducing resource consumption footprint



# Perspective on Scaling



Every day, Amazon Web Services adds enough new capacity to support all of Amazon.com's global infrastructure through the company's first 5 years, when it was a \$2.76B annual revenue enterprise

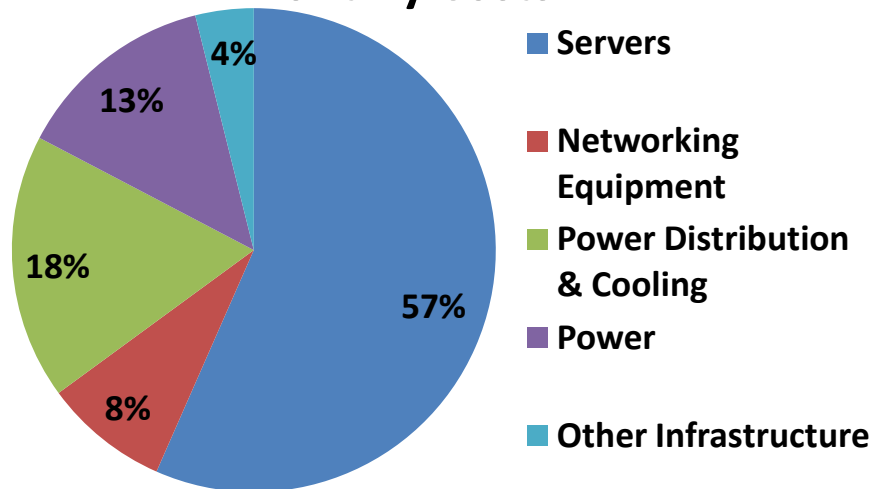
# Where Does the Money Go?

- **Assumptions:**

- Facility: ~\$88M for 8MW critical power
- Servers: 46,000 @ \$1.45k each
- Commercial Power: ~\$0.07/kWhr
- Power Usage Effectiveness: 1.45



## Monthly Costs



3yr server & 10 yr infrastructure amortization

- **Observations:**

- 31% costs functionally related to power (trending up while server costs down)
- Networking high at 8% of overall costs & 19% of total server cost (many pay more)

From: <http://perspectives.mvdirona.com/2010/09/18/OverallDataCenterCosts.aspx>

# Power Distribution

High Voltage  
Utility Distribution



115kv

**11% lost in distribution**  
 $.997 * .94 * .98 * .98 * .99 = 89\%$

IT Load (servers, storage, Net, ...)



Generators



13.2kv

UPS & Gen  
often on 480v

~1% loss in switch gear & conductors


208v

Sub-station




13.2kv


UPS:  
Rotary or Battery



13.2kv



Transformers



Transformers

480V

0.3% loss

99.7% efficient

6% loss

94% efficient, ~97% available

2% loss

98% efficient

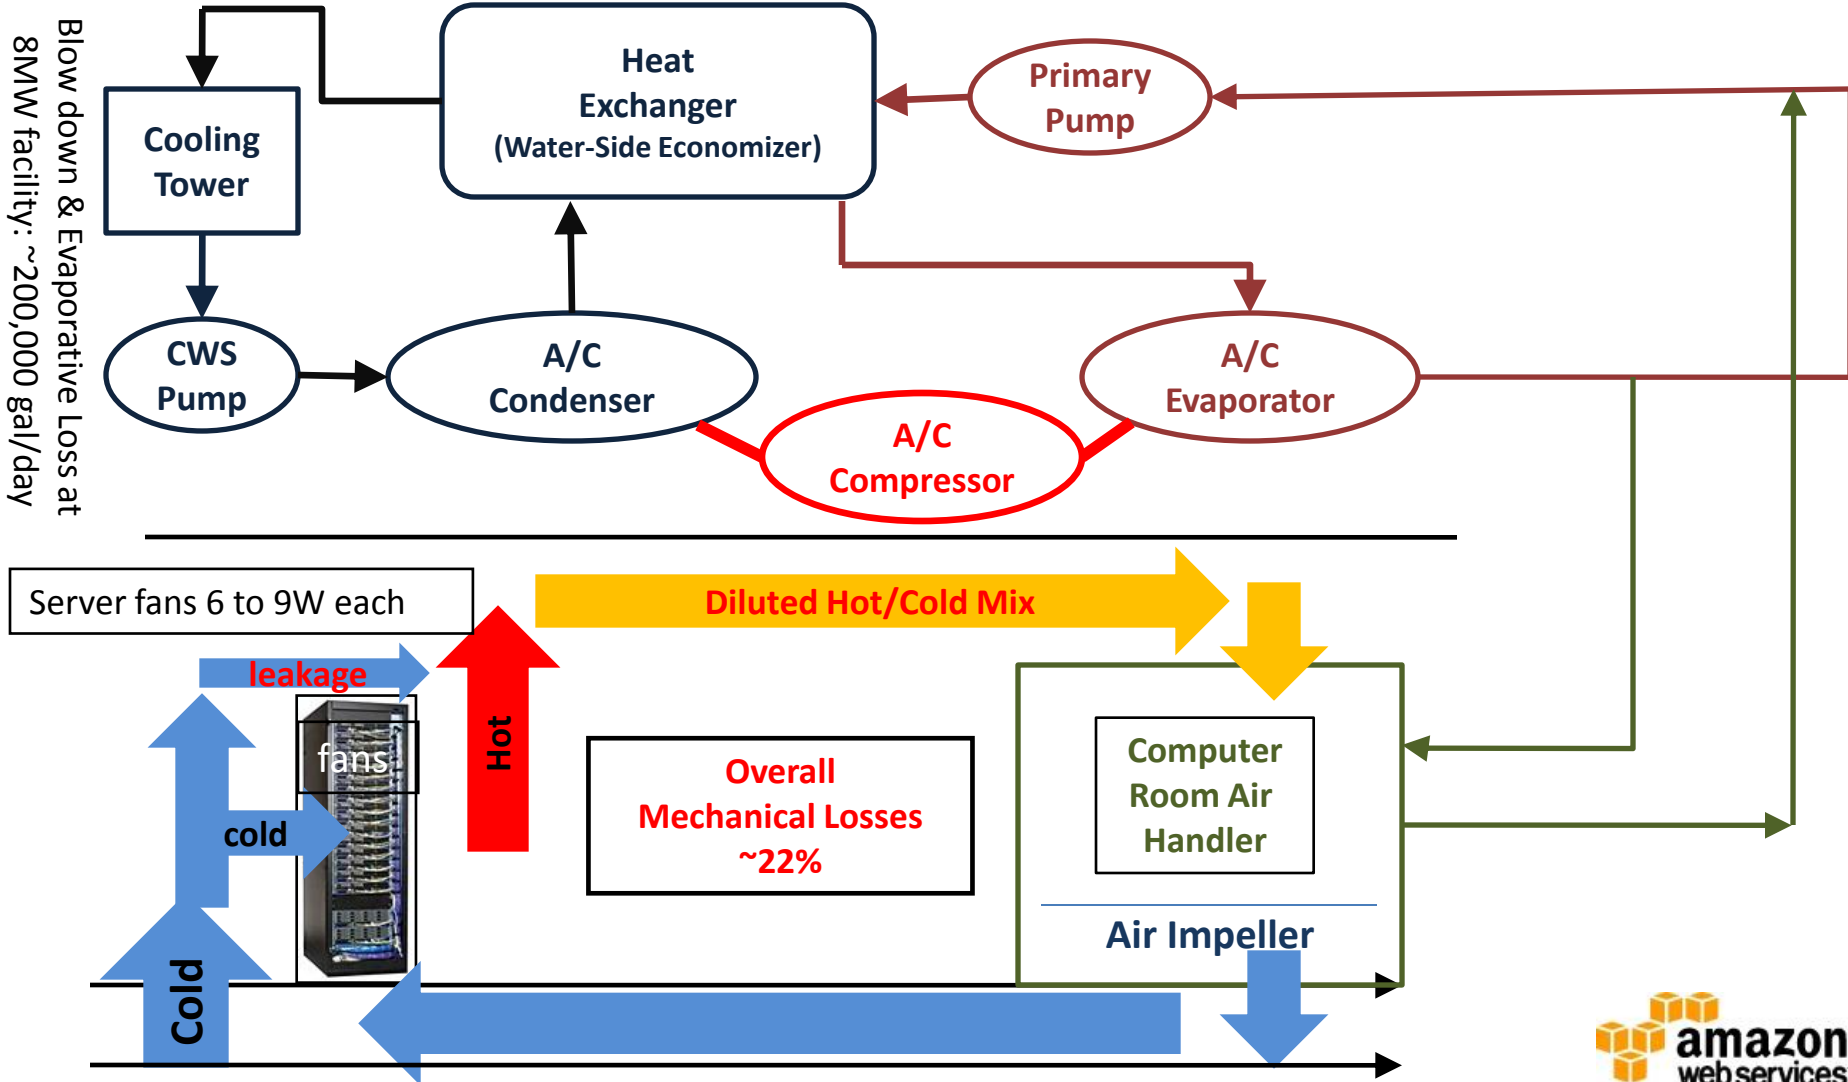
2% loss

98% efficient

Note: Two more levels of power conversion at server



# Mechanical Systems




# Hot Aisle Containment



WriteLine



Intel




Intel



# ASHRAE 2008 Allowable

## Psychrometric Chart

SI (metric) units  
Barometric Pressure 101.325 kPa (Sea level)  
based on data from  
Carrier Corporation Cat. No. 794-001, dated 1975



Most data center run in this range

ASHRAE Allowable Class 1


ASHRAE 2008 Recommended Class 1

64-81F


59F-90F

# Innovative Building Designs

- Evaporative cooling only
  - High pressure misting on right
  - Wet media design below
- Ductless full building cooling



Facebook Prineville above & below



# Modular and Pre-fab Designs



Microsoft ITPAC



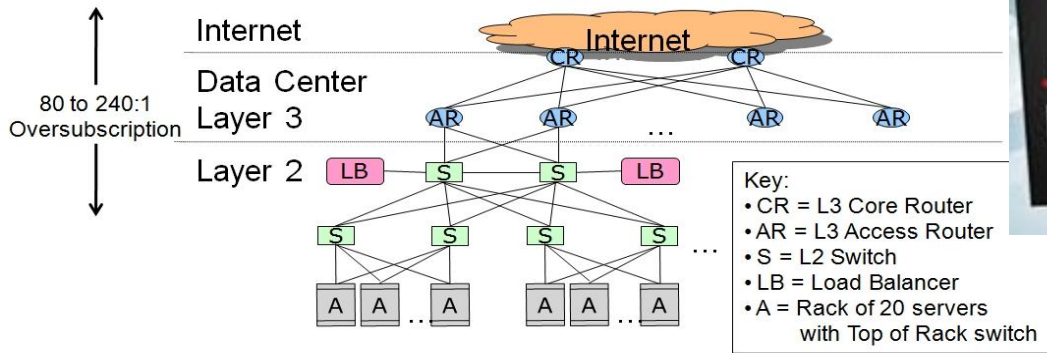

Amazon Perdix

- Fast & economic deployments
- Built in a factory setting
- Air-side economized
- ISO standard shipping containers also offered by Dell, HP, SGI, IBM, ...



# Sea Change in Networking

- Current networks over-subscribed
  - Forces workload placement restrictions
  - Goal: all points in datacenter equidistant
- Mainframe model goes commodity
  - Competition at each layer over vertical integ.
- Get onto networking on Moores Law path
  - ASIC port count growth at near constant cost
  - Competition: Broadcom, Marvell, Fulcrum,...



# Software Defined Networking

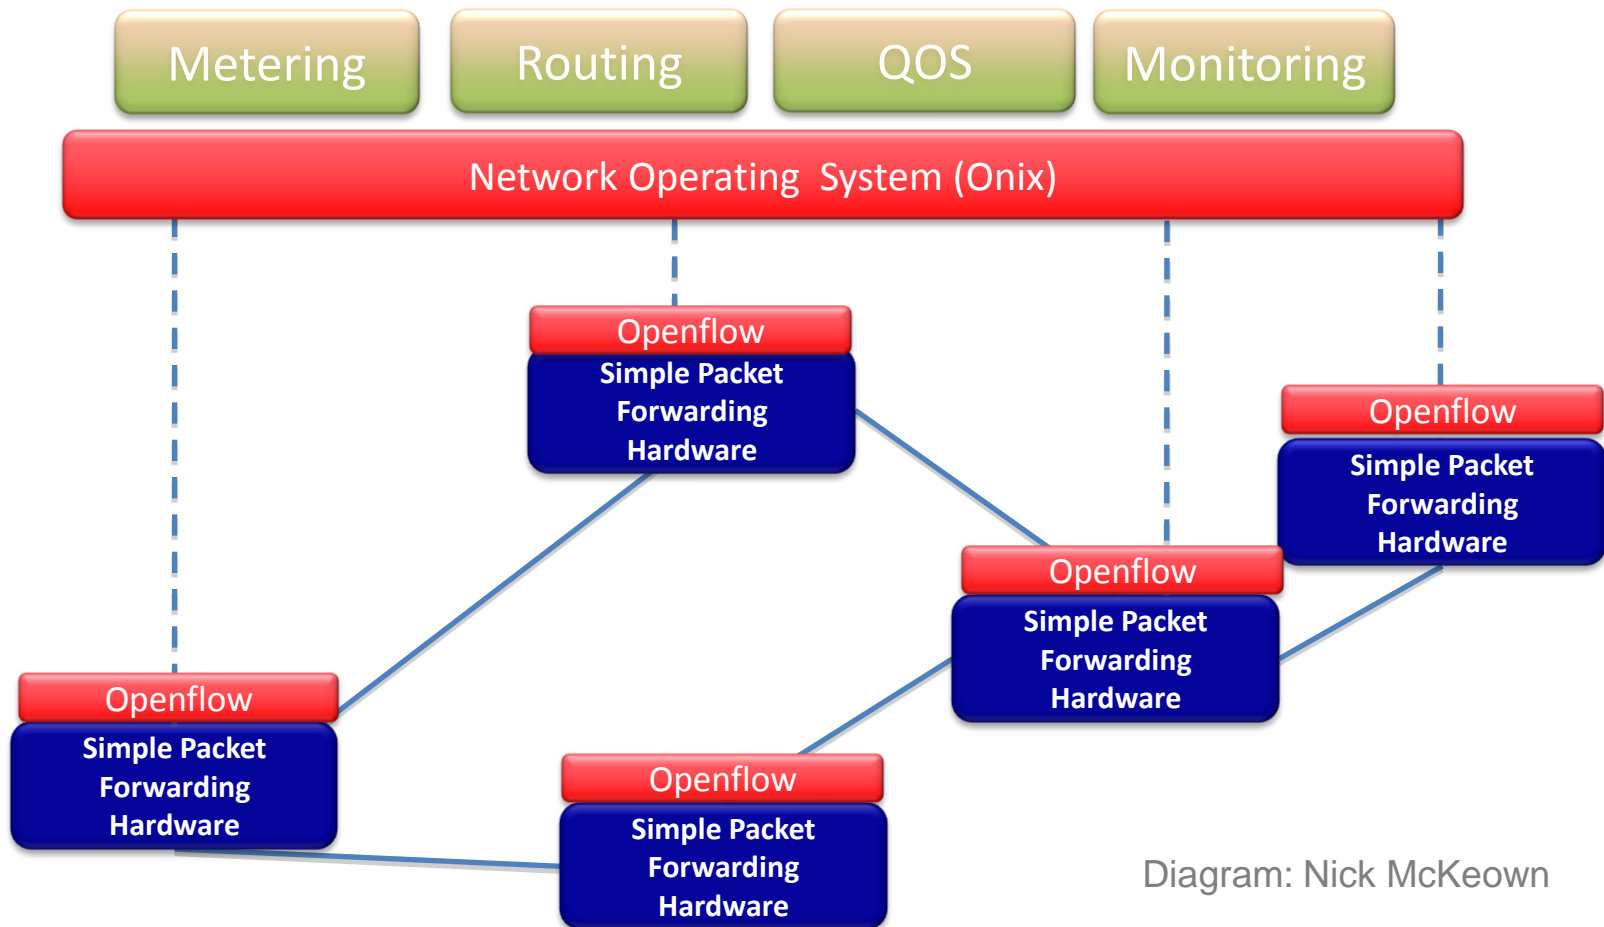


Diagram: Nick McKeown

- Distributed control plane with central control
  - Research examples: VL2, PortLand, & others
  - Onix/OpenFlow gaining industry support & traction quickly

# More Information

- **These Slides:**
  - [http://mvdirona.com/jrh/TalksAndPapers/JamesHamilton\\_StanfordComputerForum2011.pdf](http://mvdirona.com/jrh/TalksAndPapers/JamesHamilton_StanfordComputerForum2011.pdf)
- **Power and Total Power Usage Effectiveness**
  - <http://perspectives.mvdirona.com/2009/06/15/PUEAndTotalPowerUsageEfficiencyTPUE.aspx>
- **Berkeley Above the Clouds Paper**
  - <http://perspectives.mvdirona.com/2009/02/13/BerkeleyAboveTheClouds.aspx>
- **Data Center Costs**
  - <http://perspectives.mvdirona.com/2010/09/18/OverallDataCenterCosts.aspx>
- **Power Optimization**
  - [http://labs.google.com/papers/power\\_provisioning.pdf](http://labs.google.com/papers/power_provisioning.pdf)
- **Cooperative, Expendable, Microslice Servers**
  - <http://perspectives.mvdirona.com/2009/01/15/TheCaseForLowCostLowPowerServers.aspx>
- **Power Proportionality**
  - [http://www.barroso.org/publications/ieee\\_computer07.pdf](http://www.barroso.org/publications/ieee_computer07.pdf)
- **Email & Blog**
  - [James@amazon.com](mailto:James@amazon.com) & <http://perspectives.mvdirona.com>