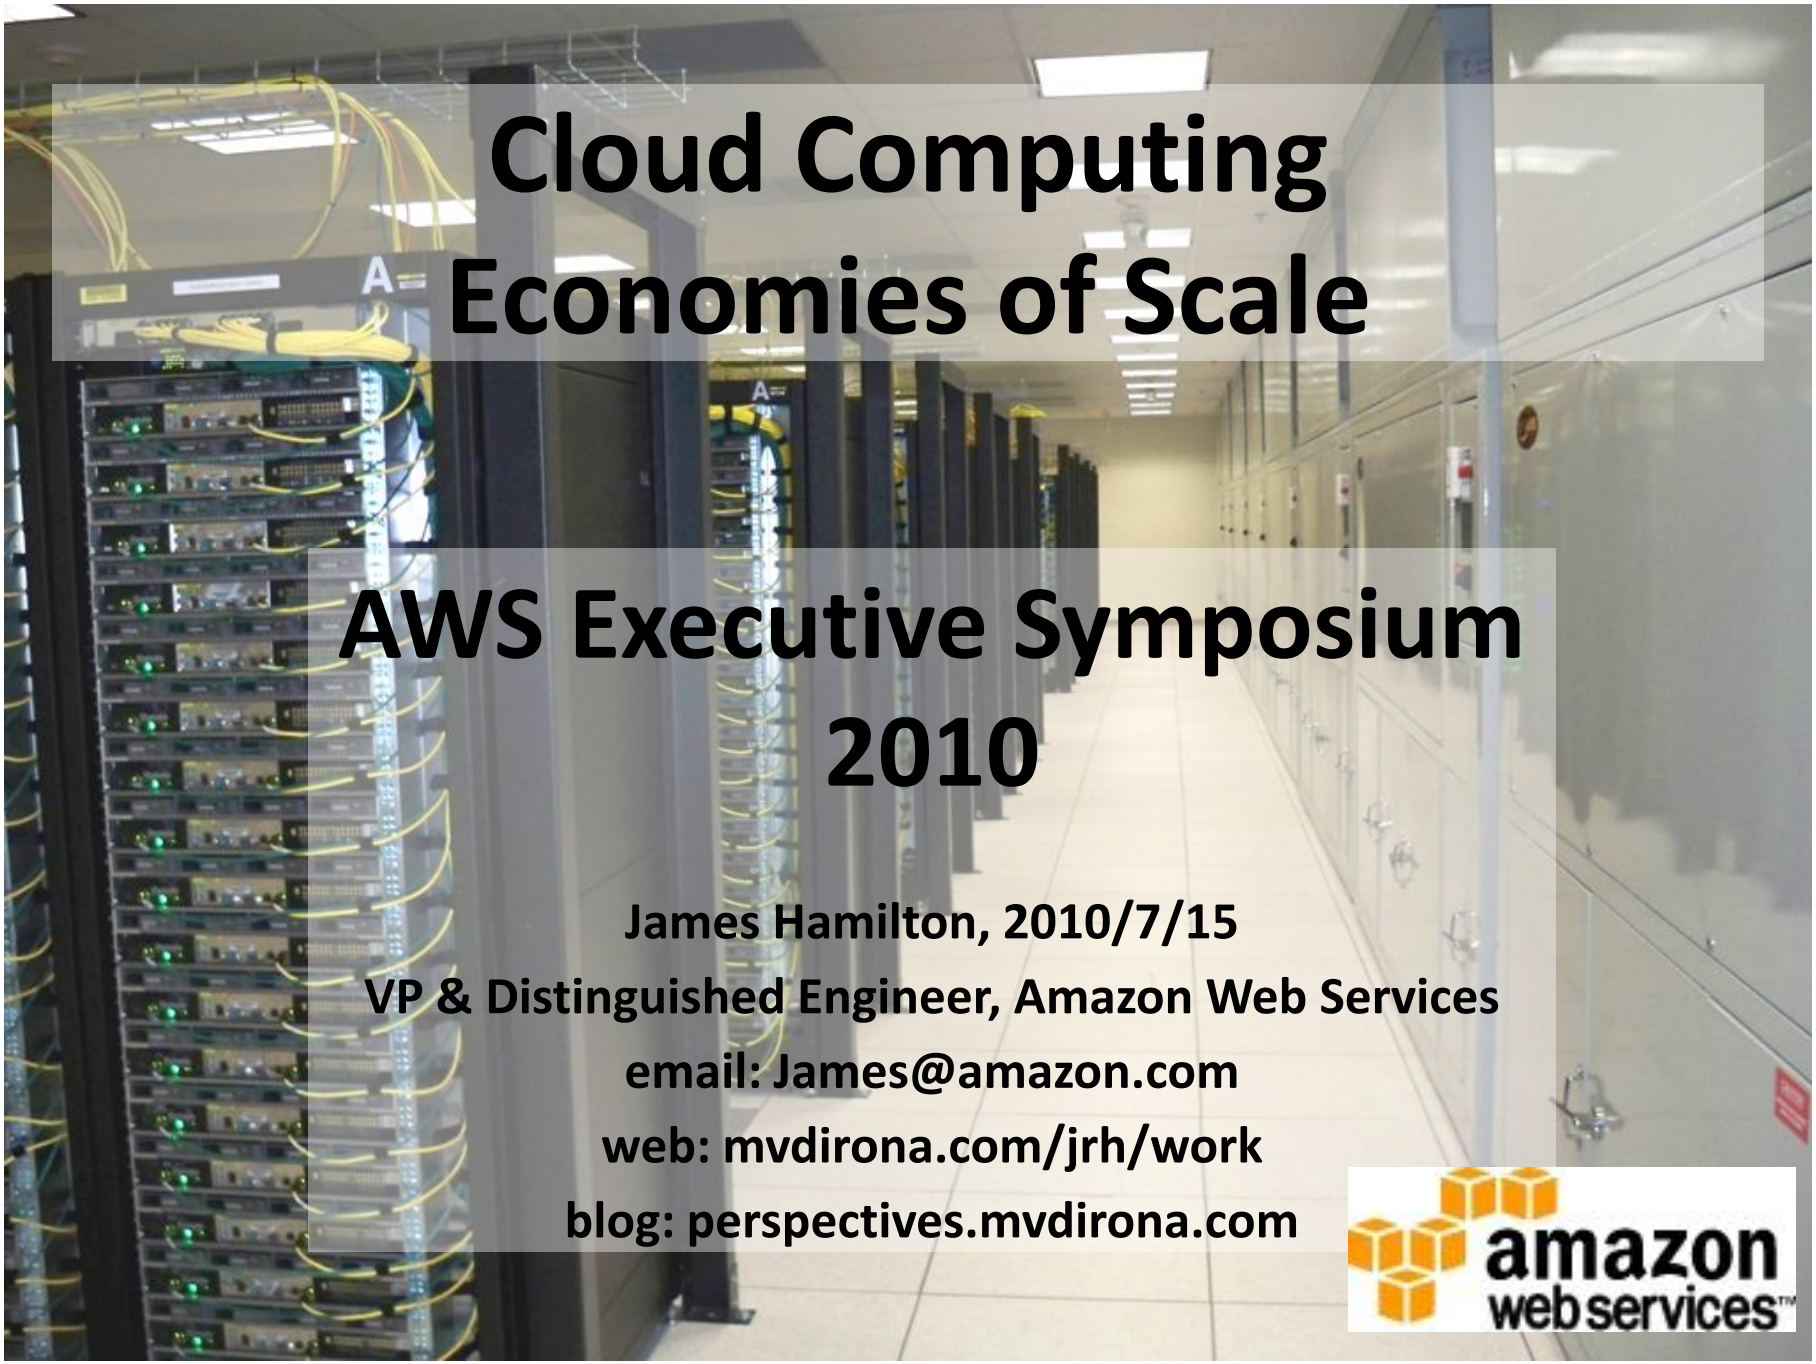

Cloud Computing Economies of Scale

AWS Executive Symposium 2010

James Hamilton, 2010/7/15
VP & Distinguished Engineer, Amazon Web Services
email: James@amazon.com
web: mvdirona.com/jrh/work
blog: perspectives.mvdirona.com

amazon web services™

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
- Power Distribution Efficiency
- Mechanical System Efficiency
- Server Design & Utilization
- Sea Change in Networking
- Cloud Computing Economics
  - Why utility computing makes sense economically

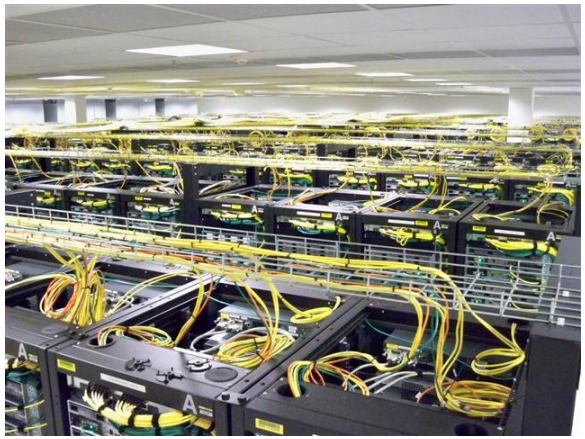

Economies of Scale

- 2006 comparison of very large service with mid-size: (~1000 servers):

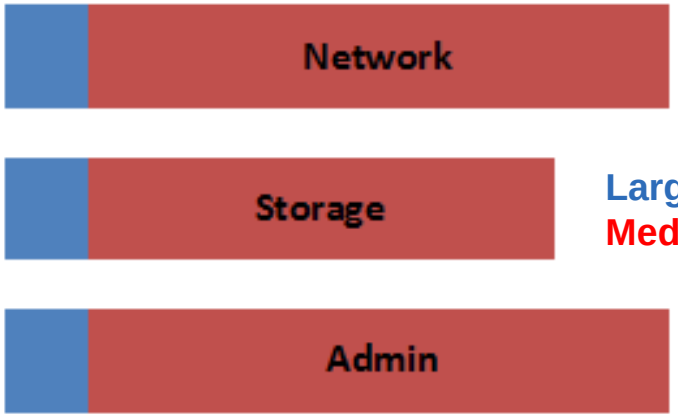

Large Service [$13/Mb/s/mth]: $0.04/GB
Medium [$95/Mb/s/mth]: $0.30/GB (7.1x)

Large Service: $4.6/GB/year (2x in 2 Datacenters)
Medium: $26.00/GB/year* (5.7x)

Large Service: Over 1.000 servers/admin
Enterprise: ~140 servers/admin (7.1x)

- Large block h/w purchases significantly more economic
  - Large weekly purchases offer significant savings
  - H/W Manufacturers willing & able to do custom designs at scale
- Automation & custom s/w investments amortize well at scale
- Summary: scale economics strongly in play

Where Does the Money Go?

- **Assumptions:**
  - Facility: ~$88M for 8MW facility
  - Servers: Roughly 46k @ $1.45k each
  - Server power draw at 30% load: 80%
  - Commercial Power: ~$0.07/kWhr
  - PUE: 1.5

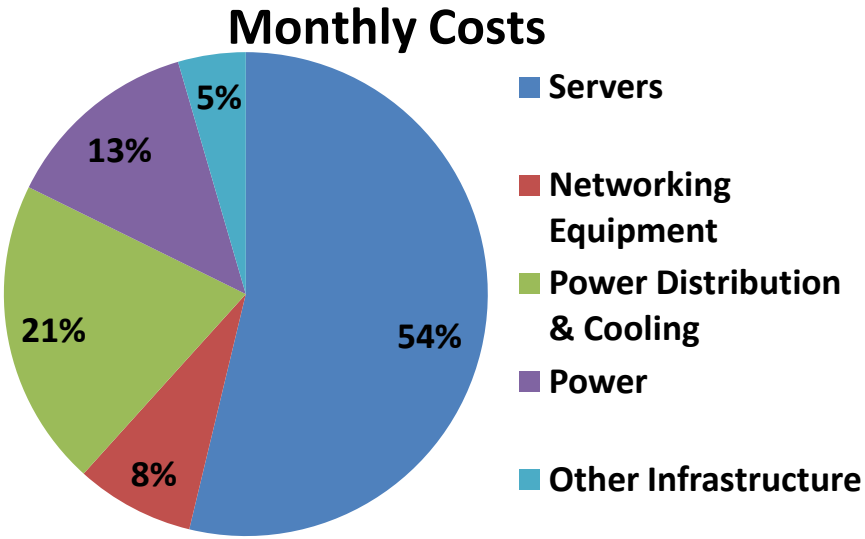

3yr server, 4yr net gear, & 10 yr infrastructure amortization

- **Observations:**
  - 34% costs functionally related to power (trending up while server costs down)
  - Networking high at 8% of costs & 19% of total server cost

Updated from: http://perspectives.mvdirona.com/2008/11/28/CostOfPowerInLargeScaleDataCenters.aspx

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
- Power Distribution Efficiency
- Mechanical System Efficiency
- Server Design & Utilization
- Sea Change in Networking
- Cloud Computing Economics
  - Why utility computing makes sense economically

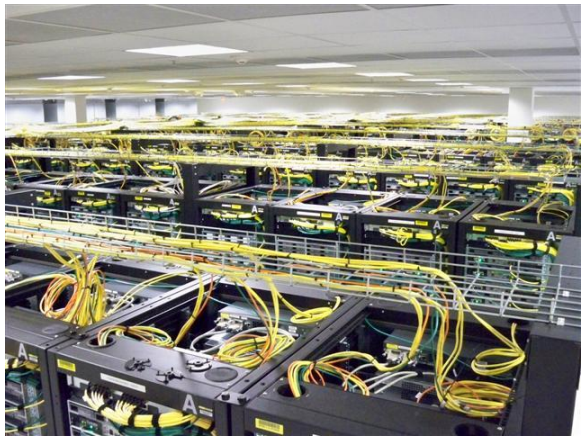

Power Distribution

High Voltage Utility Distribution

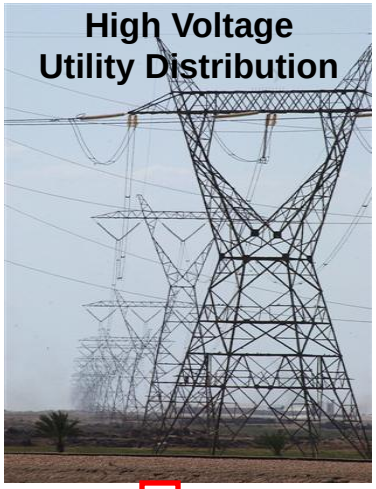

115kv

Sub-station

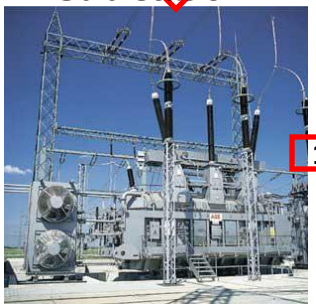

0.3% loss
99.7% efficient

13.2kv

11% loss in distribution
.997*.94*.98*.98*.99 = 89%

2.5MW Generator (180 gal/hr)

13.2kv

UPS:
Rotary or Battery

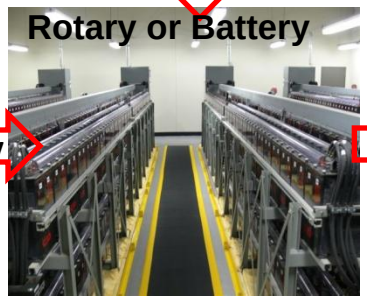

6% loss
94% efficient, ~97% available

13.2kv

UPS & Gen often on 480v

Transformers

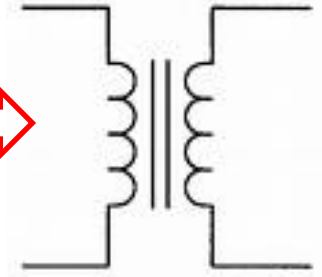

2% loss
98% efficient

480V

Transformers

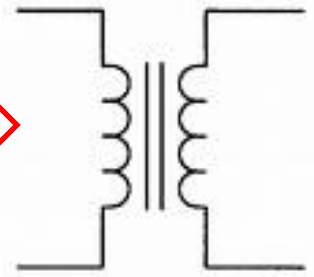

2% loss
98% efficient

208V

IT Load (servers, storage, Net, …)

~1% loss in switch gear & conductors

Conventional Mechanical Design

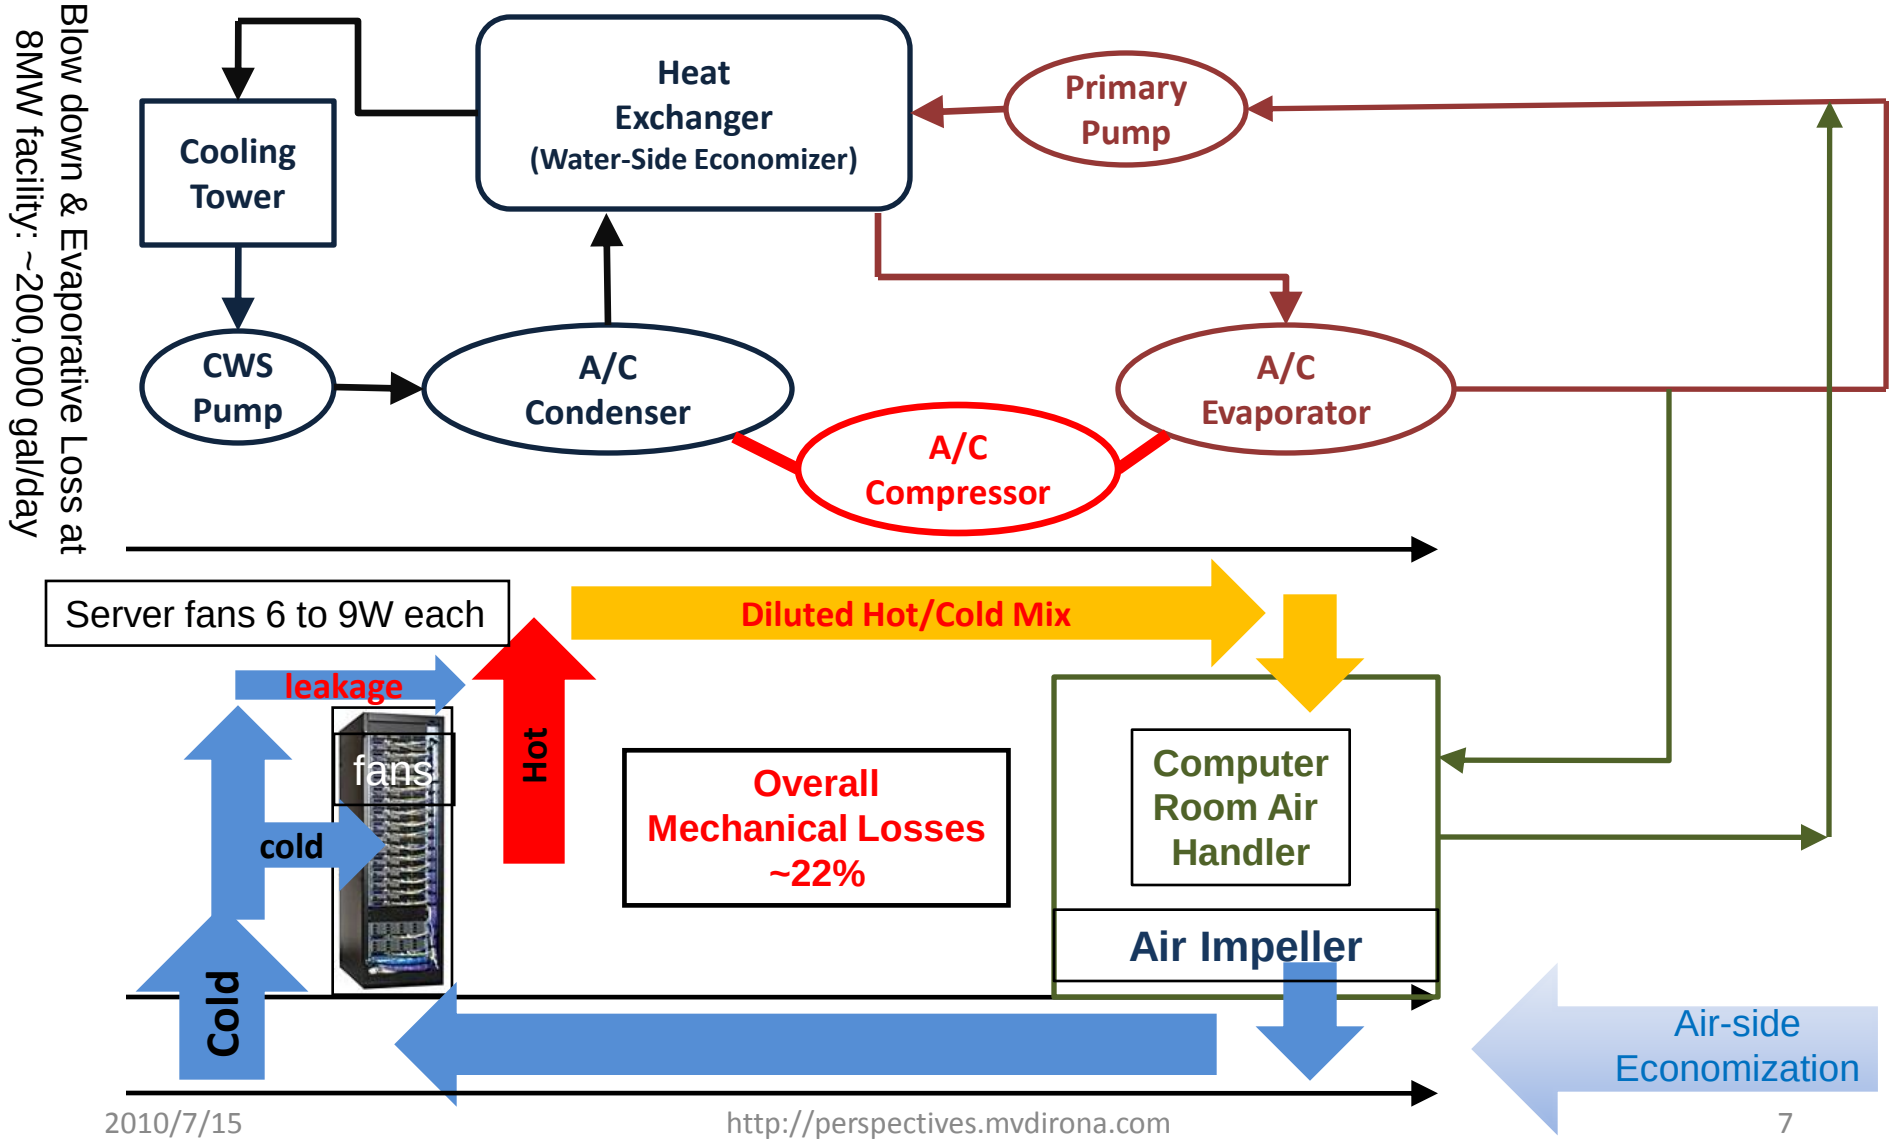

Hot Aisle/Cold Aisle Containment

WriteLine

Intel

Intel

Air-Side Economization & Evaporative Cooling

- Limiting factors to high temp operation
  - Higher fan power trade-off
  - More semiconductor leakage current
  - Possible negative failure rate impact
- Avoid direct expansion cooling entirely
  - Air side economization
  - Higher data center temperatures
  - Evaporative cooling
- Requires Filtration
  - Particulate & chemical pollution

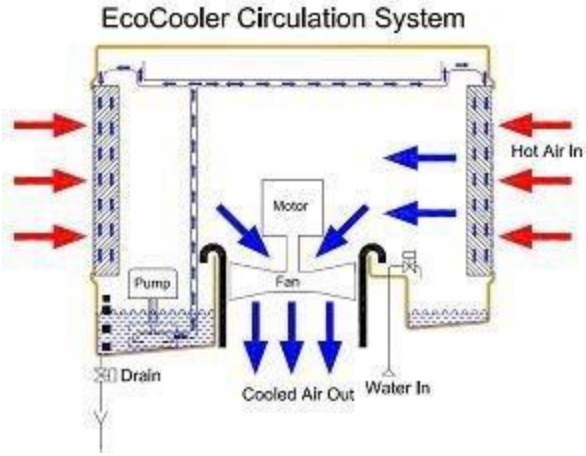

Server Innovation

- Shared Infrastructure Racks
  - Shared redundant PSUs & fans
  - e.g. Dell Fortuna & Rackable CloudRack
- Next Level: Multi-server on board
  - Intel Atom: SeaMicro
  - ARM: SmoothStone
- Very Low-Cost, Low-Power Servers
  - ARM, Atom, client & embedded CPUs
  - Cold storage (reduce CPU $ to GB)
  - Highly partitionable workloads: Web services, memcached
- Low utilization is still the elephant in room

Sea Change in Net Gear

- Current networks over-subscribed
  - Forces workload placement restrictions
  - Goal: all points in datacenter equidistant
- Mainframe model goes commodity
  - Competition at each layer rather than vertical integration
- OpenFlow: open S/W platform
  - Distributed control plane to central control
  - E.g. VL2, Portland, and others

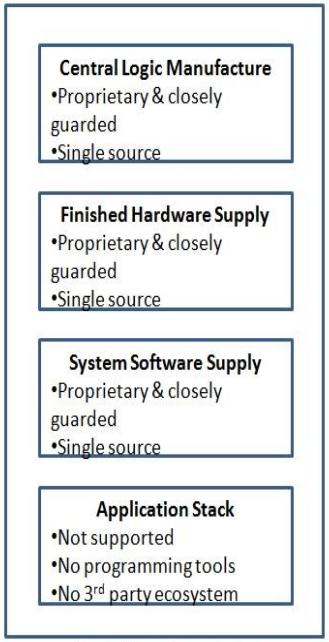

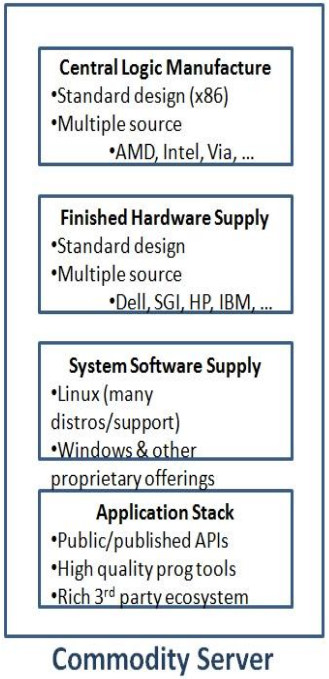

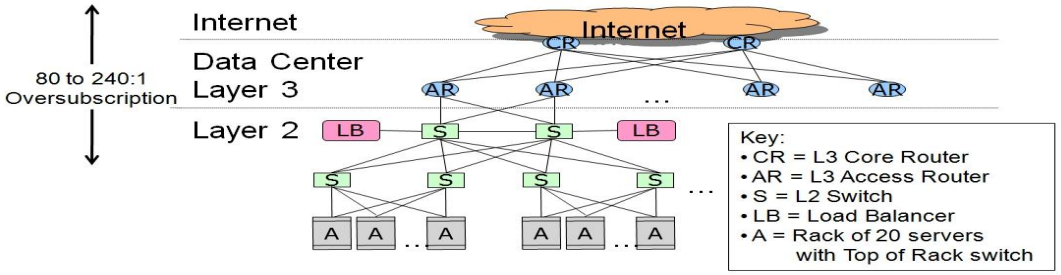

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
- Power Distribution Efficiency
- Mechanical System Efficiency
- Server Design & Utilization
- Sea Change in Networking
- Cloud Computing Economics
  - Why utility computing makes sense economically

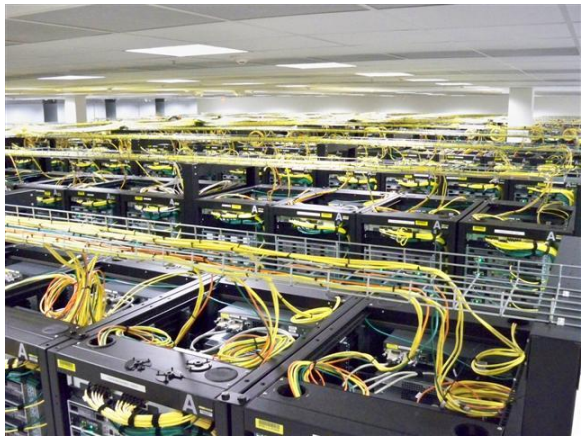

Infrastructure at Scale

- Datacenter design efficiency
  - Average datacenter efficiency low with PUE over 2.0 (Source: EPA)
    - Many with PUE well over 3.0
  - High scale cloud services in the 1.2 to 1.5 range
  - Lowers computing cost & better for environment
- Multiple datacenters
  - At scale multiple datacenters can be used
    - Close to customer
    - Cross datacenter data redundancy
    - Address international markets efficiently
- Avoid massive upfront data cost & years to fully utilize

AWS Approach

- Broad set of services:
  - Infrastructure Services
    - SimpleDB
    - Simple Storage Service
    - CloudFront
    - Simple Queue Service
    - Elastic MapReduce
    - Relational Database Service
    - Elastic Block Store
    - Premium Support
    - Virtual Private Cloud
  - Payments & Billing
    - Flexible Payment Services
    - DevPay
  - On Demand Workforce
    - Mechanical Turk
  - Alexa Web Services
    - Web Information Service
    - Top Sites
  - Merchant Services
    - Fulfillment Web Service
- "Open the hood" approach
  - Simple, layerable building block services
  - Component services are substitutable

Amazon Cycle of Innovation

- 15+ years of operational excellence
  - Managing secure, highly available, multi-datacenter infrastructure
- Experienced at low margin cycle of innovation:
  - Innovate
  - Listen to customers
  - Drive down costs & improve processes
  - Pass on value to customers
- AWS price reductions expected to continue

AWS Scale

- Looking at Quantcast Top 500k sites (11/2009)
  - 53% of cloud hosted sites are using AWS
  - 27% growth in AWS hosted site count Oct to Nov
  - More sites than all others combined
- High growth workloads: data Intensive computing, commercial HPC, analysis, & optimization
- Very high scale supports deep investment in automation, monitoring, operations, & funds faster innovation

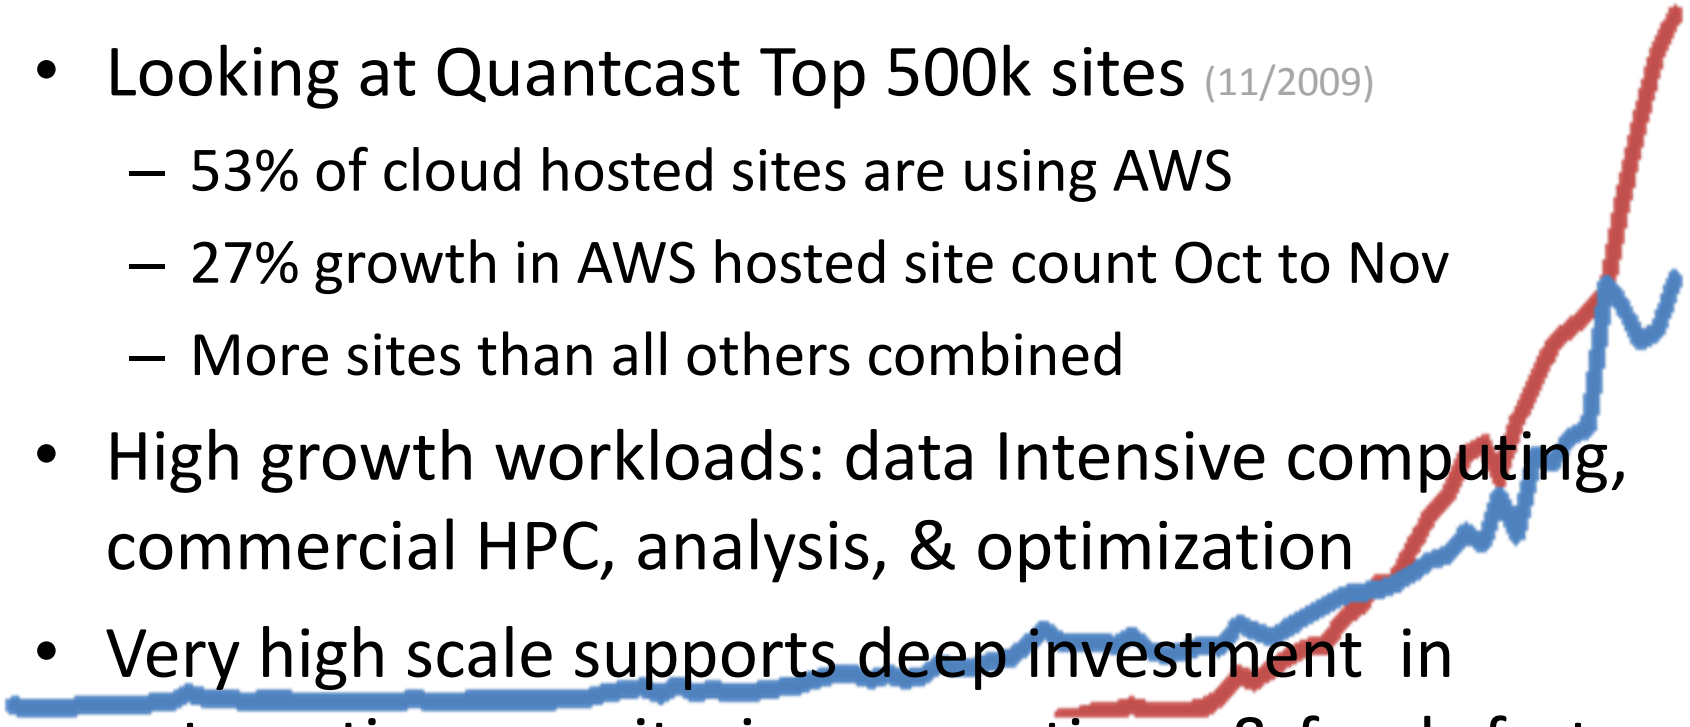

Sources:

- http://www.quantcast.com/top-sites-1
- http://www.jackofallclouds.com/2009/11/state-of-the-cloud-november-2009/

Utilization & Economics

- Server utilization problem
  - 30% utilization VERY good &10% to 20% common
    - Expensive & not good for environment
  - Solution: pool number of heterogeneous services
    - Single reserve capacity pool far more efficient
    - Non-correlated peaks & law of large numbers
- Pay as you go & pay as you grow model
  - Don’t block the business
  - Don’t over buy
  - Transfers capital expense to variable expense
  - Apply capital for business investments rather than infrastructure
- Charge back models drive good application owner behavior
  - Cost encourages prioritization of work by application developers
  - High scale needed to make a market for low priority work

Amazon Web Services Pace of Innovation

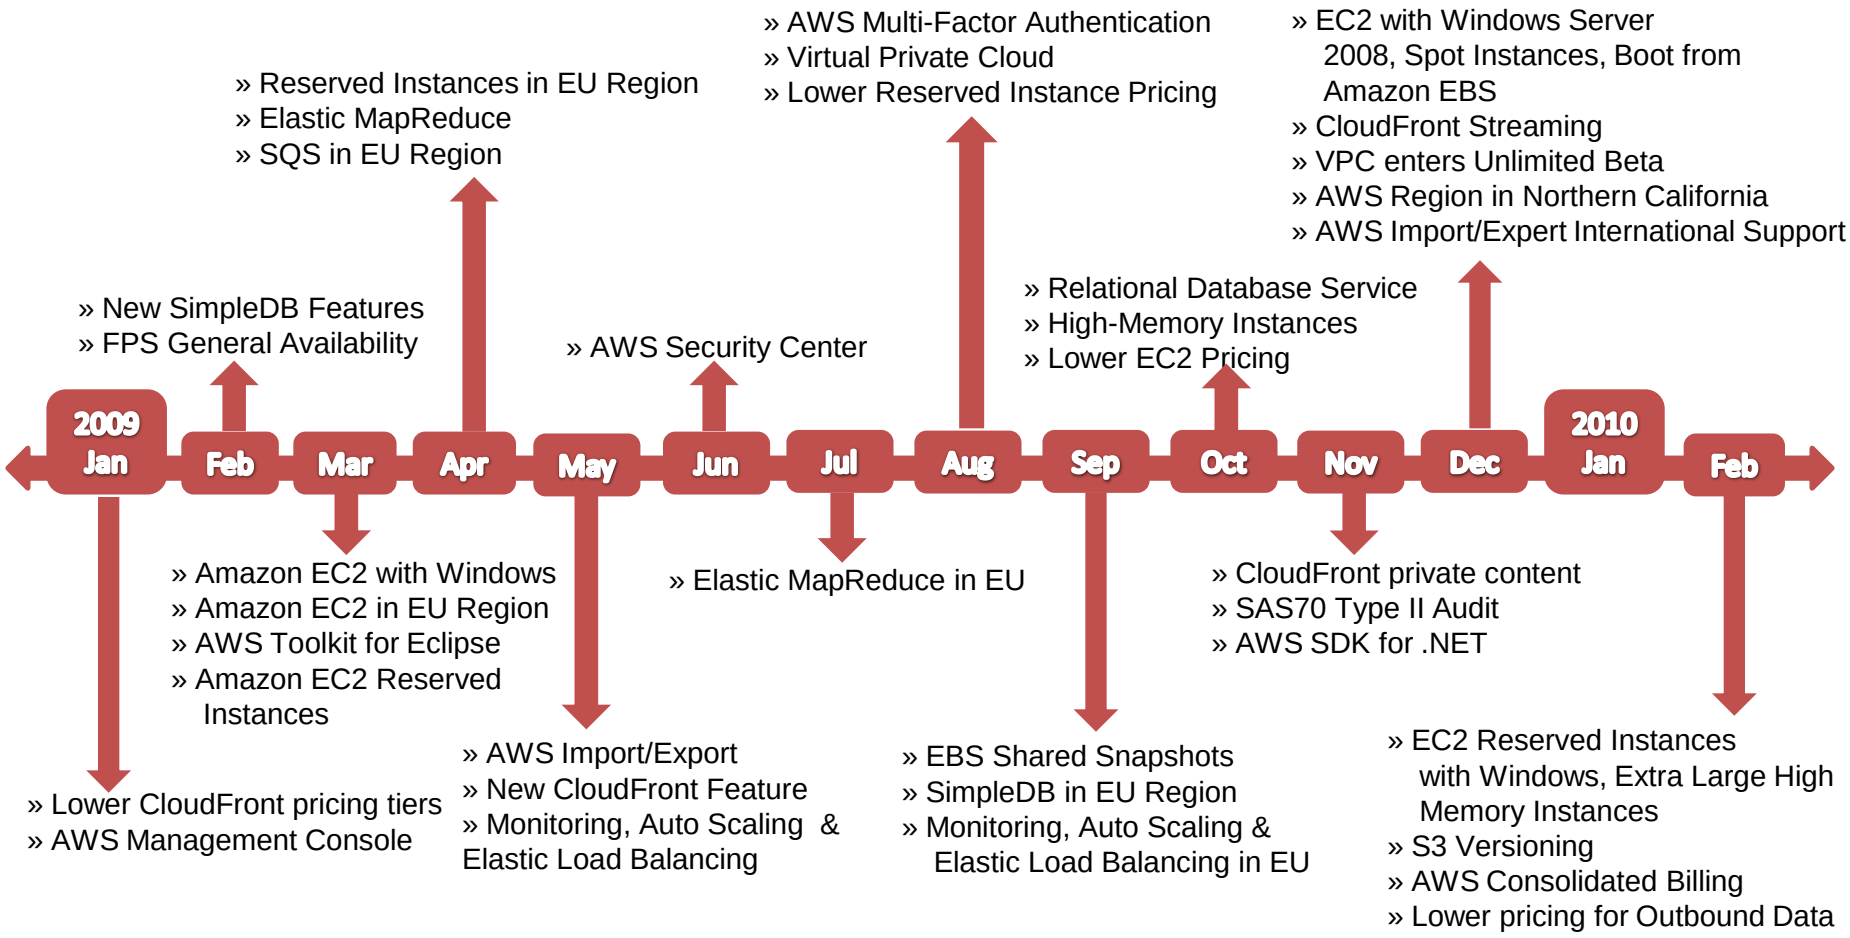

More Information

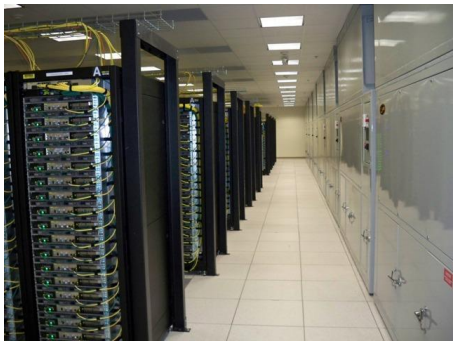

- **This Slide Deck:**
  - I will post slides to http://mvdirona.com/jrh/work later this week
- **Berkeley Above the Clouds**
  - http://perspectives.mvdirona.com/2009/02/13/BerkeleyAboveTheClouds.aspx
- **Degraded Operations Mode**
  - http://perspectives.mvdirona.com/2008/08/31/DegradedOperationsMode.aspx
- **Cost of Power**
  - http://perspectives.mvdirona.com/2008/11/28/CostOfPowerInLargeScaleDataCenters.aspx
  - http://perspectives.mvdirona.com/2008/12/06/AnnualFullyBurdenedCostOfPower.aspx
- **Power Optimization:**
  - 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
- **Resource Consumption Shaping:**
  - http://perspectives.mvdirona.com/2008/12/17/ResourceConsumptionShaping.aspx
- **Email**
  - James@amazon.com