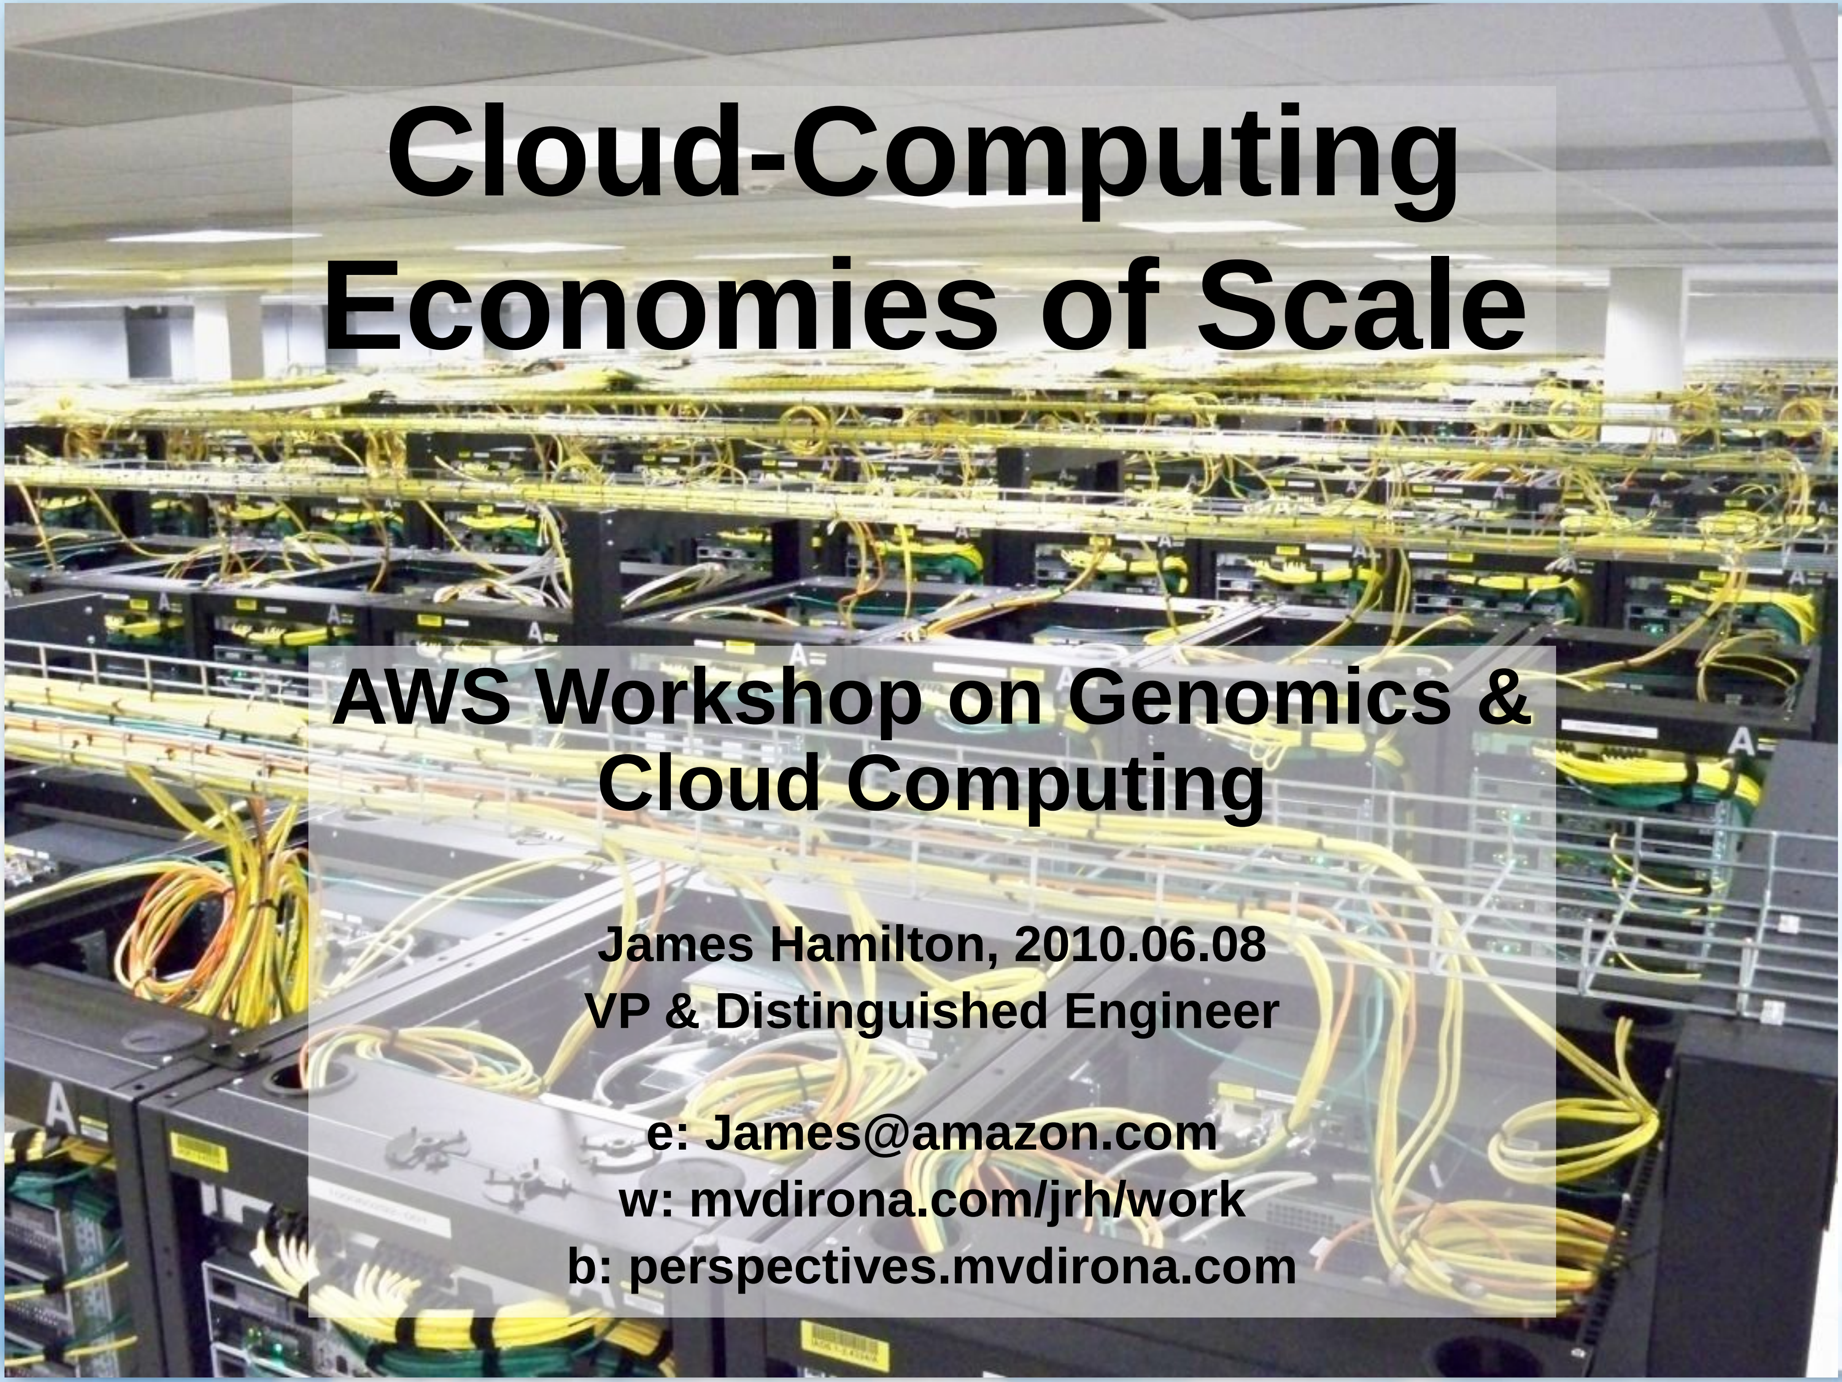

Cloud-Computing Economies of Scale

AWS Workshop on Genomics & Cloud Computing

James Hamilton, 2010.06.08
VP & Distinguished Engineer

e: James@amazon.com
w: mvdirona.com/jrh/work
b: perspectives.mvdirona.com

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
  - Where does the power go?
- Power Distribution Efficiency
- Mechanical System Efficiency
- Cloud Computing Economics
  - Why utility computing makes sense economically
- Summary

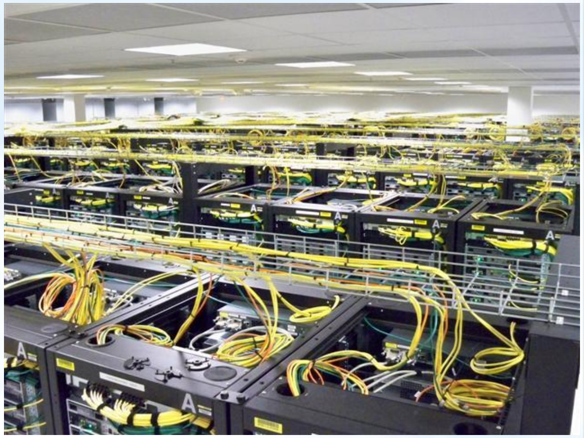

amazon webservices™

Economies of Scale

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

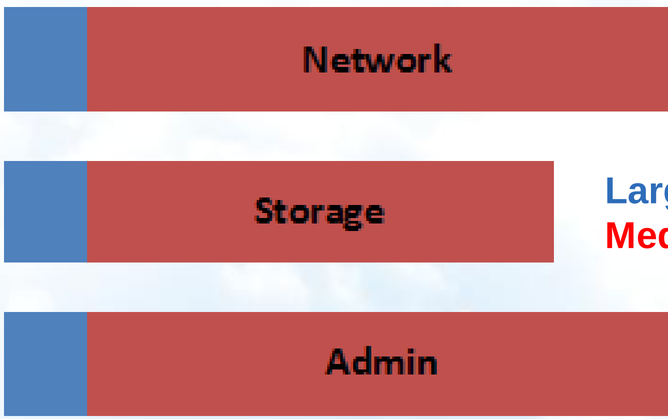

Large Service [$13/Mb/s/mth]: $0.04/GB
Medium [$95/Mb/s/mth]: $0.30/GB (7.1x)

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

Power & Related Costs [Will] Dominate

- **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

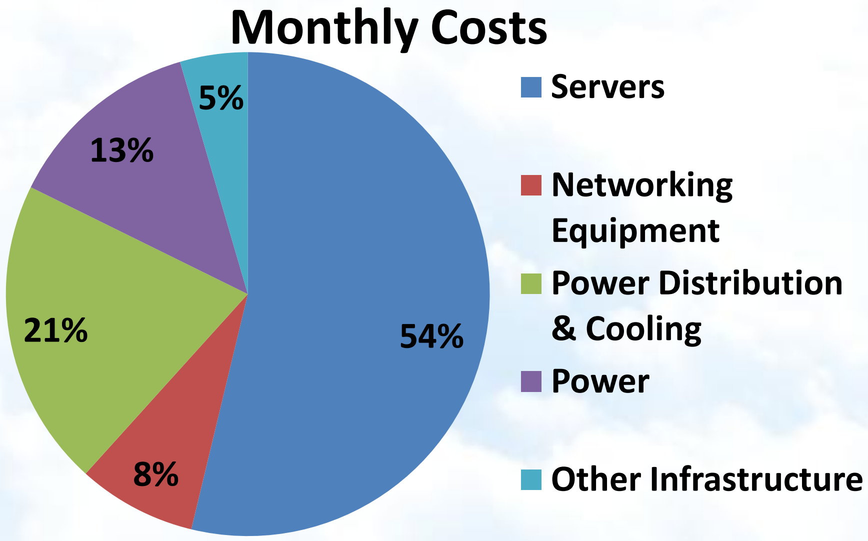

3yr server, 4yr net gear, & 10 yr infrastructure amortization

- **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

Where Does the Power Go?

- **Assuming a good data center with PUE ~1.5**
  - Each server watt loses ~0.5W to power distribution & cooling
  - IT load (servers & storage): 1/1.5 => 67%
  - Network gear <4% total power (5.8% of IT load)
- **Power losses are easier to track than cooling:**
  - Power transmission, conversion, & switching losses: 11%
    - Detailed power distribution losses on next slide
  - Cooling losses the remainder:100-(67+11) => 22%
- **Observations:**
  - Utilization & server efficiency improvements very leveraged
  - Networking gear very power inefficient individually but not big problem in aggregate
  - Cooling costs unreasonably high
  - PUE improving rapidly

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
  - Where does the power go?
- Power Distribution Efficiency
- Mechanical System Efficiency
- Cloud Computing Economics
  - Why utility computing makes sense economically
- Summary

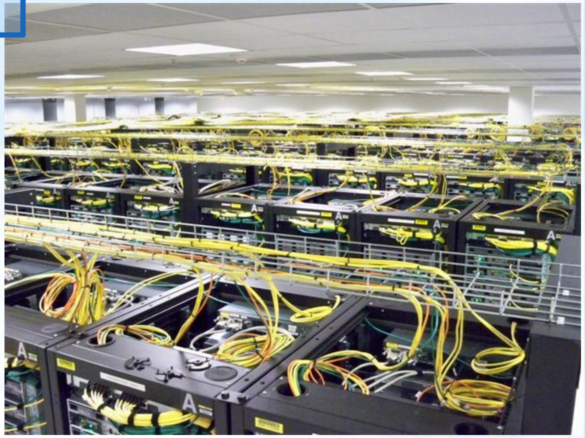

Power Distribution

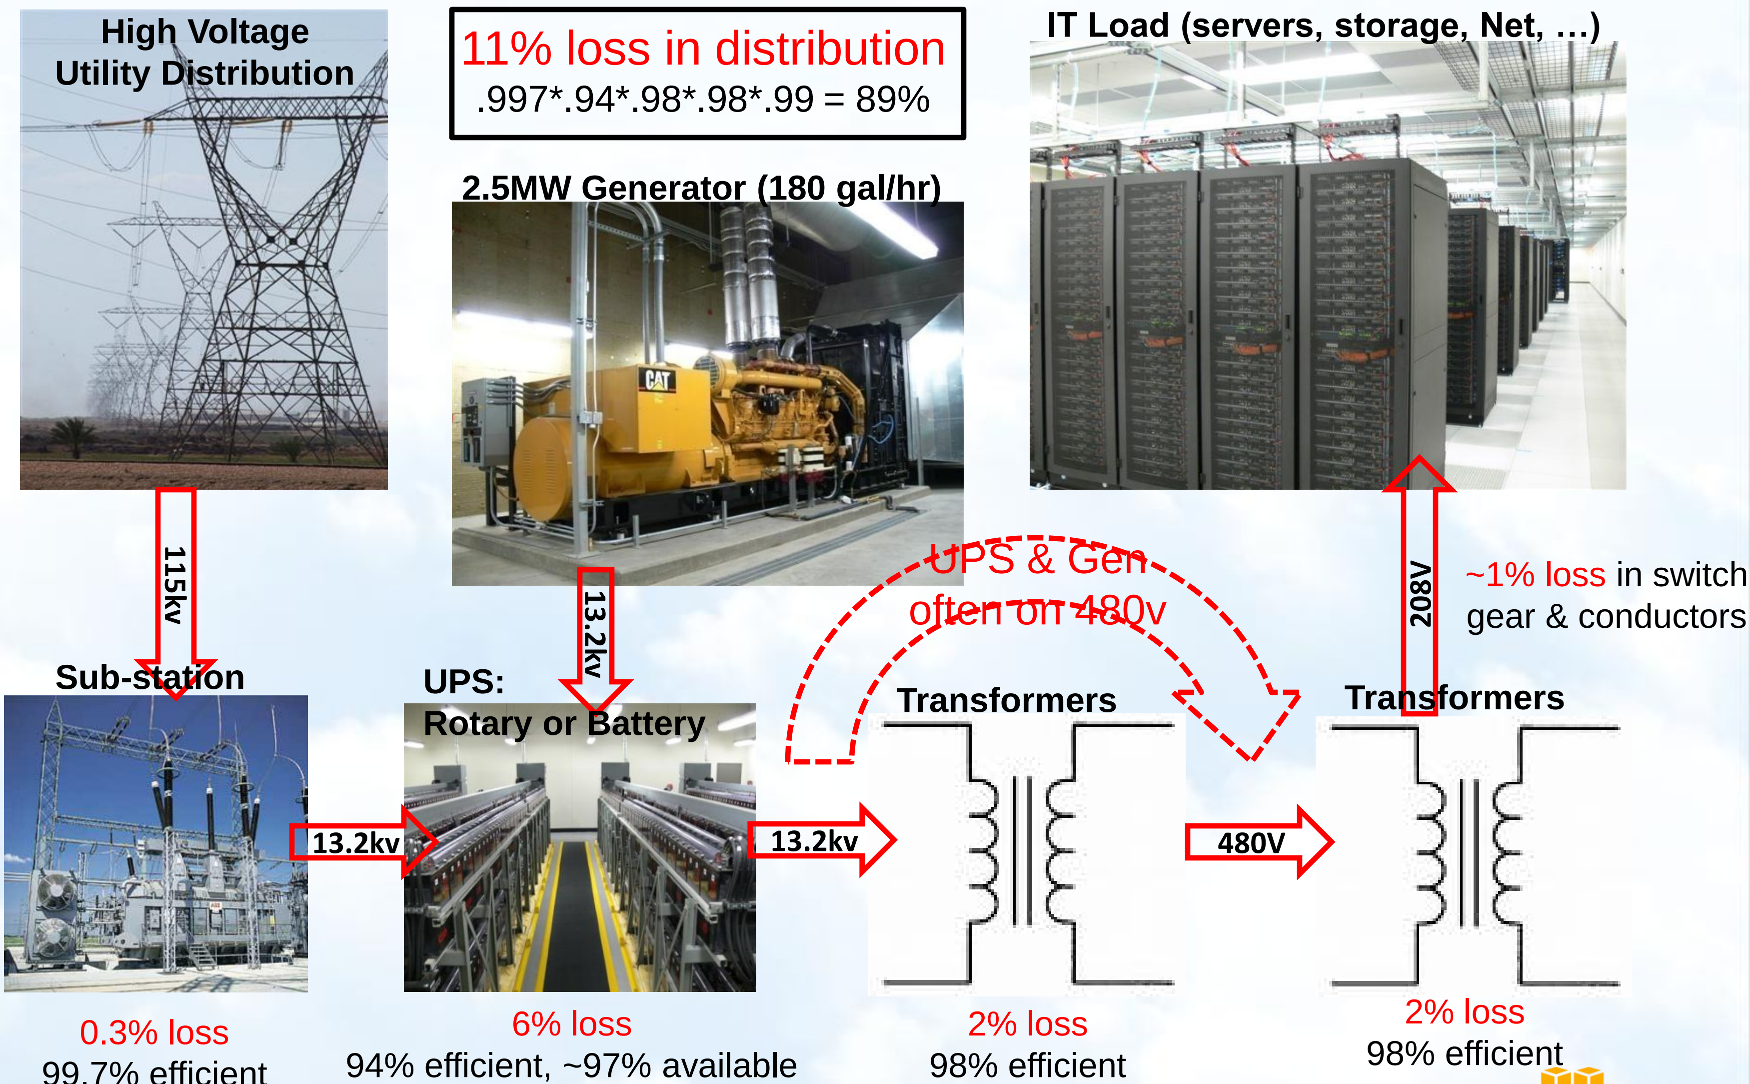

Power Distribution Efficiency Summary

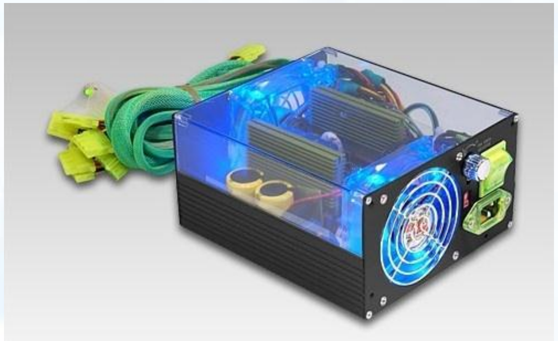

- Two additional conversions in server:
  1. Power Supply: often <80% at typical load
  2. On board step-down (VRM/VRD)
     - <80 Common & ~95% efficient available
- Rules to minimize power distribution losses:
  1. Oversell power (more potential load than provisioned power)
  2. Avoid conversions (fewer transformer steps & efficient UPS)
  3. Increase efficiency of conversions
  4. High voltage as close to load as possible
  5. Size VRMs & VRDs to load & use efficient parts
  6. DC distribution a fairly small potential gain

- But power distribution improvements bounded to 11%

amazon webservices™

Conventional Mechanical Design

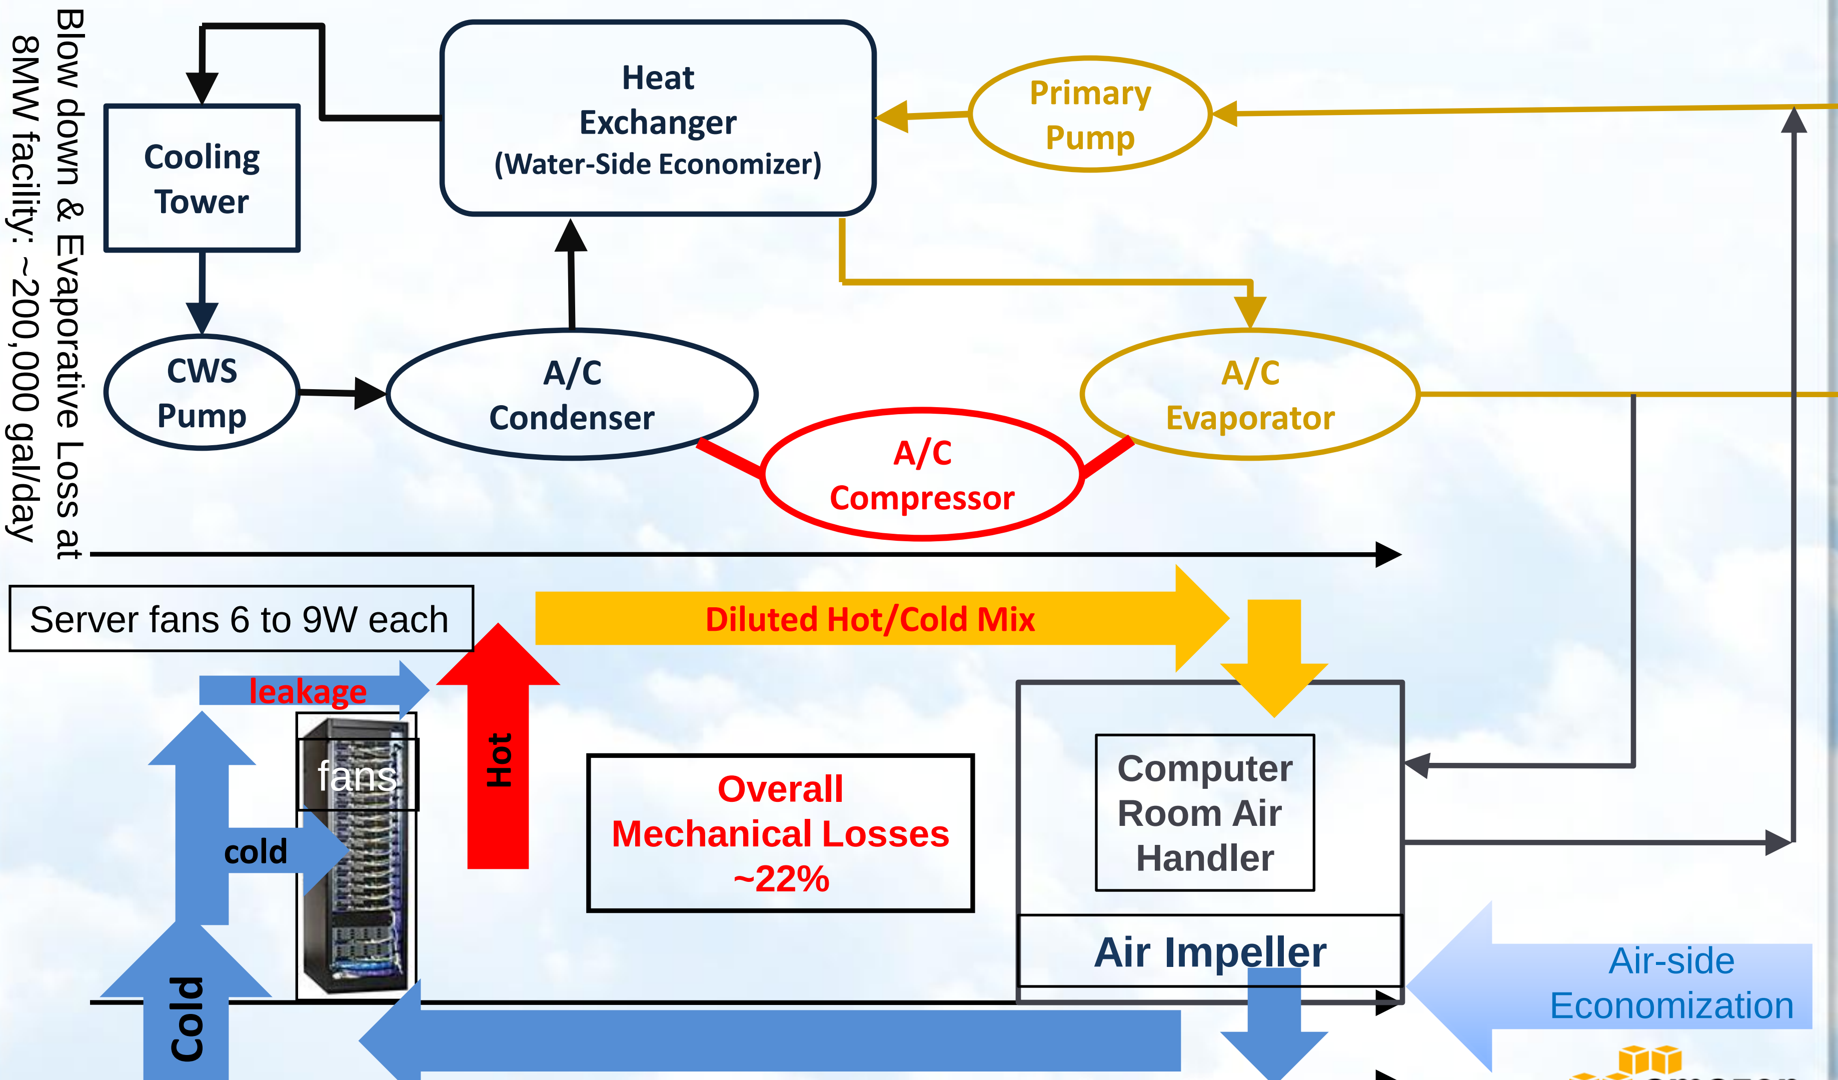

Mechanical Efficiency Summary

- Prioritized mechanical optimizations:
  1. Raise data center temperatures
  2. Tight airflow control, short paths & large impellers
  3. Cooling towers rather than A/C
  4. Air-side economization & evap cooling
     - outside air rather than A/C & towers

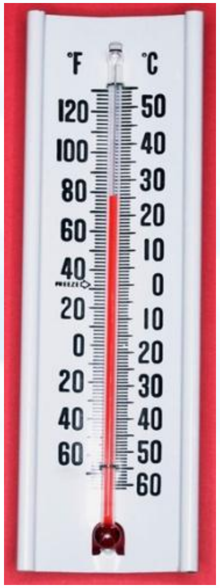

Agenda

- Follow the money in infrastructure
  - Infrastructure cost breakdown
  - Where does the power go?
- Power Distribution Efficiency
- Mechanical System Efficiency
- Cloud Computing Economics
  - Why utility computing makes sense economically
- Summary

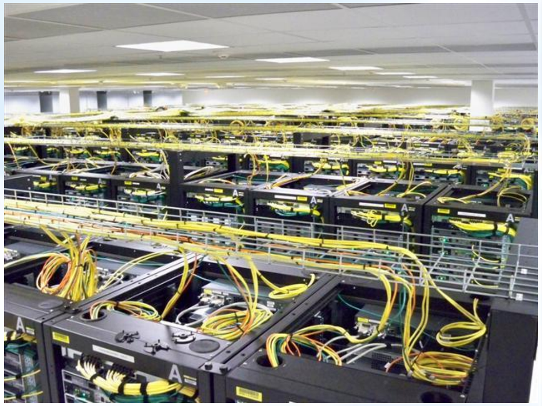

Investments at Scale

- Deep automation only affordable when amortized over large user base
  - Lack of automation drives both cost & fragility (human error)

- S/W investments at scale
  - Massive distributed systems investments such as Amazon Simple Storage Service & Elastic Block Store hard to justify without scale

- Special Skills with deep focus
  - Distributed systems engineers, power engineering, mechanical engineering, server h/w design, networking, supply chain, 24x7 operations staff, premium support,…

amazon webservices™

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

H/W Cost & Efficiency Optimization

- Service optimized hardware
  - Custom cloud-scale design teams:
    - Dell DCS, SGI (aka Rackable), ZT Systems, IBM, HP, …
- Purchasing power at volume

- Supply chain optimization
  - Shorter chain drives much higher server utilization
    - Predicting next week easier than 4 to 6 months out
  - Less overbuy & less capacity risk
- Networking transit costs rewards volume
- Cloud services unblocks new business & growth
  - Remove dependence on precise capacity plan

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

AWS Pace of Innovation

2009

May
» AWS Import/Export
» New CloudFront Feature
» Monitoring, Auto Scaling & Elastic Load Balancing

Jun
» AWS Security Center

Jul
» Amazon Elastic MapReduce in Europe

Aug
» AWS Multi-Factor Authentication
» Virtual Private Cloud
» Lower Reserved Instance Pricing

Sep
» EBS Shared Snapshots
» SimpleDB in EU Region
» Monitoring, Auto Scaling & Elastic Load Balancing in EU

Oct
» Amazon RDS
» High-Memory Instances
» Lower EC2 Pricing

Nov
» Amazon CloudFront Private Content
» SAS70 Type II Audit
» AWS SDK for .NET

Dec
» Amazon EC2 with Windows Server 2008, Spot Instances, Boot from Amazon EBS
» Amazon CloudFront Streaming
» Amazon VPC enters Unlimited Beta
» AWS Region in Northern California
» International Support for AWS Import/Export

2010

Jan

Feb
» Amazon EC2 Reserved Instances with Windows, Extra Large High Memory Instances
» Amazon S3 Versioning Feature
» Consolidated Billing for AWS
» Lower pricing for Outbound Data Transfer

Mar
» AWS SDK for Java
» Windows server 'Bring Your Own License' pilot program
» Amazon CloudFront: Singapore edge location, private content for streaming

Apr
» Combined AWS Data Transfer Pricing
» Amazon SNS
» Amazon Elastic MapReduce: custom cluster configuration option
» Amazon RDS: EU Region launch
» AWS Asia Pacific (Singapore) Region

May
» Amazon VPC: Europe launch
» Amazon RDS: Northern California Region, Multi-AZ Deployments, AWS Management Console support
» Amazon CloudFront: Access logs for streaming
» Amazon S3 Reduced Redundancy Storage

Jun
» Amazon CloudFront: HTTPS support, lower request pricing, NYC edge location

amazon webservices™

Summary

- Measure efficiency using work done/dollar & work done/joule
  - Server costs dominate all other DC infrastructure & admin at scale
  - 2/3 of total data center power is delivered to servers
  - Utilization poor: Servers are idle more than not
  - Conclusion: nearly ½ the provisioned power not doing useful work
- Considerable room for DC cooling improvements
- Cloud services drive:
  - Higher resource utilization
  - Innovation in power distribution & mechanical systems
  - Lower cost, higher reliability, & lower environmental impact

More Information

- **These Slides:**
  - http://mvdirona.com/jrh/TalksAndPapers/JamesHamilton_GenomicsCloud20100608.pdf
- **Power and Total Power Usage Effectiveness**
  - http://perspectives.mvdirona.com/2009/06/15/PUEAndTotalPowerUsageEfficiencyTPUE.aspx
- **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 & Blog**
  - James@amazon.com & http://perspectives.mvdirona.com

amazon web services™

18