OPEN
Compute Summit

Internet Scale Infrastructure Innovation

Open Compute Summit 2011

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

Agenda

- Quickening Pace Infrastructure Innovation
  - Influence of Cloud computing
- Power Distribution
- Cooling & Shell Designs

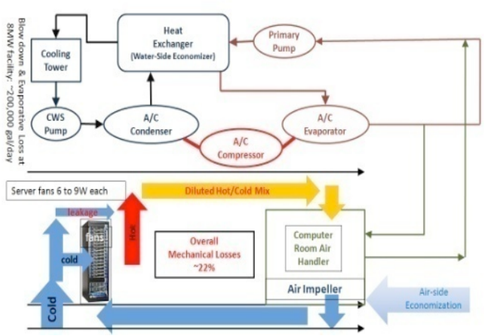

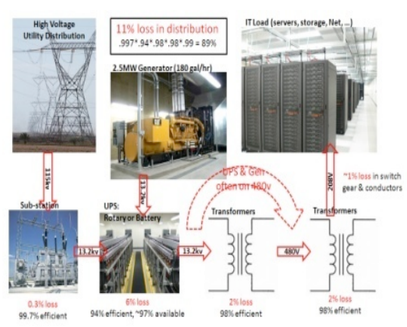

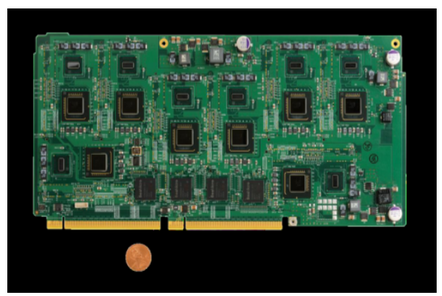

Talk does not necessarily represent positions of current or past employers

Quickening 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

Each day Amazon Web Services adds enough capacity to support all of Amazon.com's global infrastructure through the company's first 5 years, when it was a $2.76B enterprise 2011/10/

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

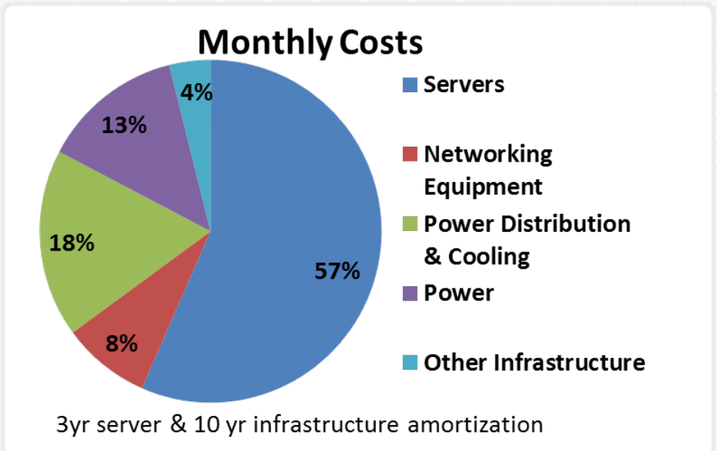

- Observations:
  - 31% costs functionally related to power (trending up while server costs down)
  - Networking high at 8% of overall costs & 19% of total server cost (many pay more)

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

Power Distribution

~11% lost in distribution - .997*.94*.98*.98*.99 = 89%

High Voltage Utility Distribution

115kv

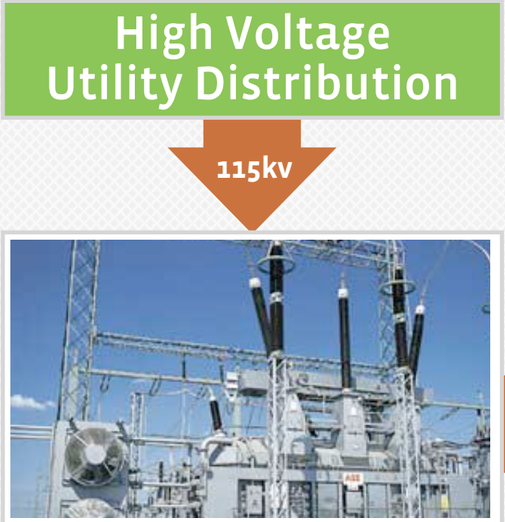

Sub-station

0.3% loss
99.7% efficient

13.2kv

Generators

13.2kv

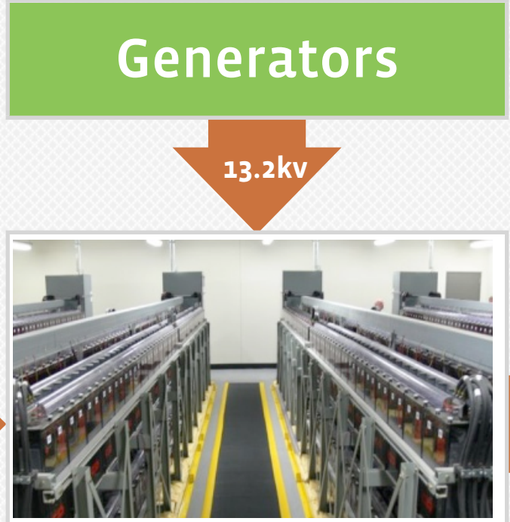

UPS: Rotary or Battery

6% loss
94% efficient, ~97% available

13.2kv

UPS & Gen often on 480V

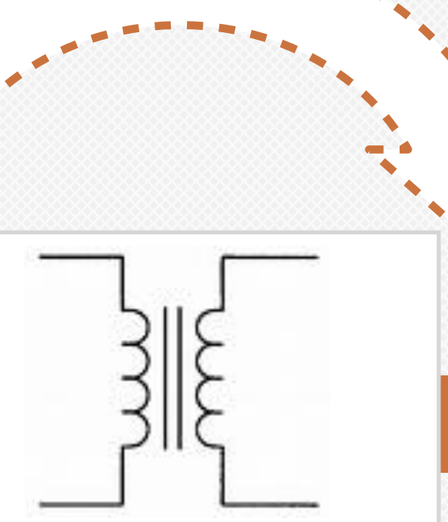

Transformers

2% loss
98% efficient

480V

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

208V

~1% loss in switch gear & conductors

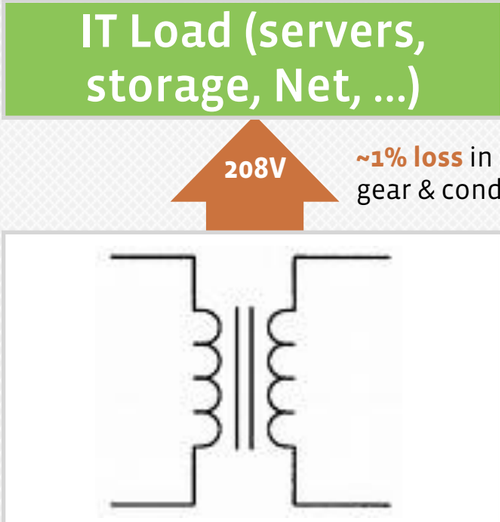

Transformers

2% loss
98% efficient

Note: Two more levels of power conversion at server

Power Distribution Efficiency Summary

- 2 more power conversions at servers
  5. Power Supply: often under 80% at typical load
  6. On board voltage regulators (VRMs or VRDs)

- Rules to minimize power distribution losses:
  - Oversell power (more load than provisioned power)
  - Avoid conversions (fewer & better)
  - Increase efficiency of conversions
  - High voltage as close to load as possible
  - Size voltage regulators to load & use efficient parts
  - High voltage direct current a small potential gain

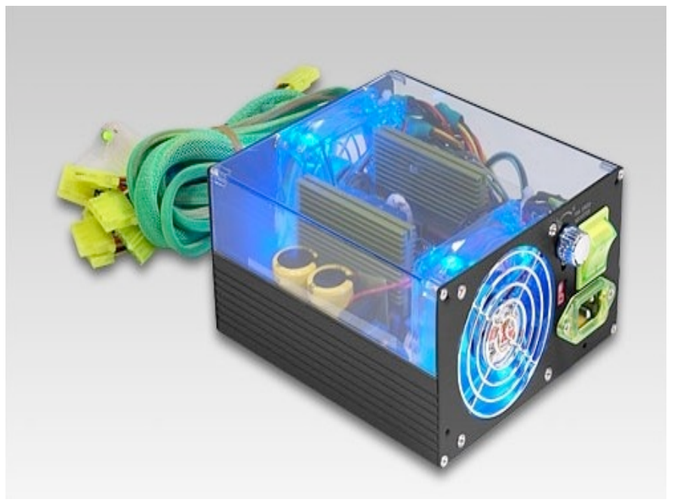

Mechanical Systems

20% of total power

Heat Exchanger (Water-Side Economizer)

Primary Pump

Cooling Tower

Blow down & Evaporative Loss at 8MW facility: ~200,000 gal/day

CWS Pump

A/C Condenser

A/C Compressor

A/C Evaporator

Server fans 6 to 9W each

Diluted Hot/Cold Mix

leakage

fans

Hot

cold

Cold

Overall Mechanical Losses ~22%

Computer Room Air Handler

Air Impeller

Hot Aisle Containment

Facebook Open Compute

WriteLine

Intel

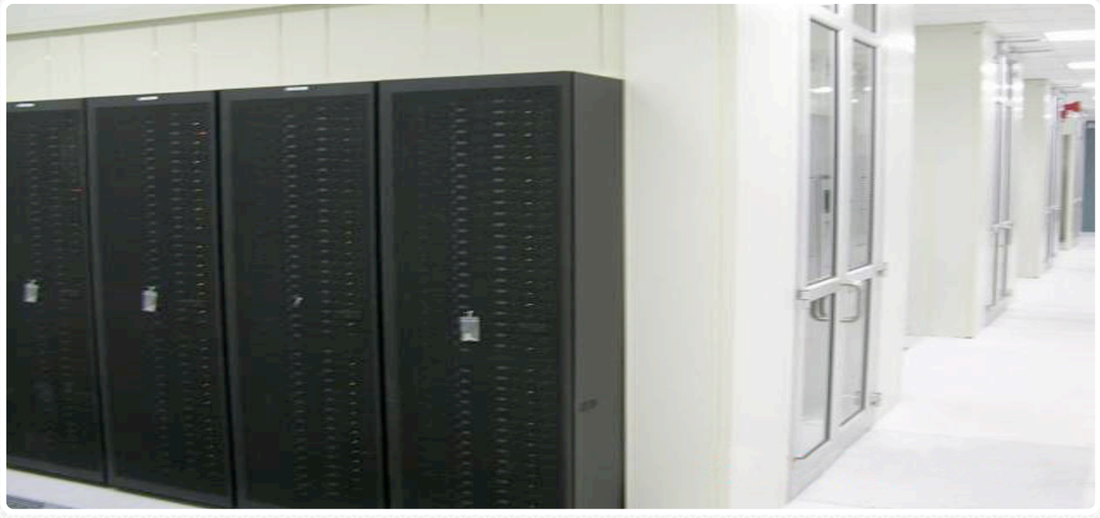

Intel

ASHRAE Recommendations

NEBS Telco Standard (~1970)

ASHRAE 2011 Allowable Class 1

ASHRAE 2011 Recommended Class 1

ASHRAE 2008 Allowable Class 1

ASHRAE 2008 Recommended Class 1

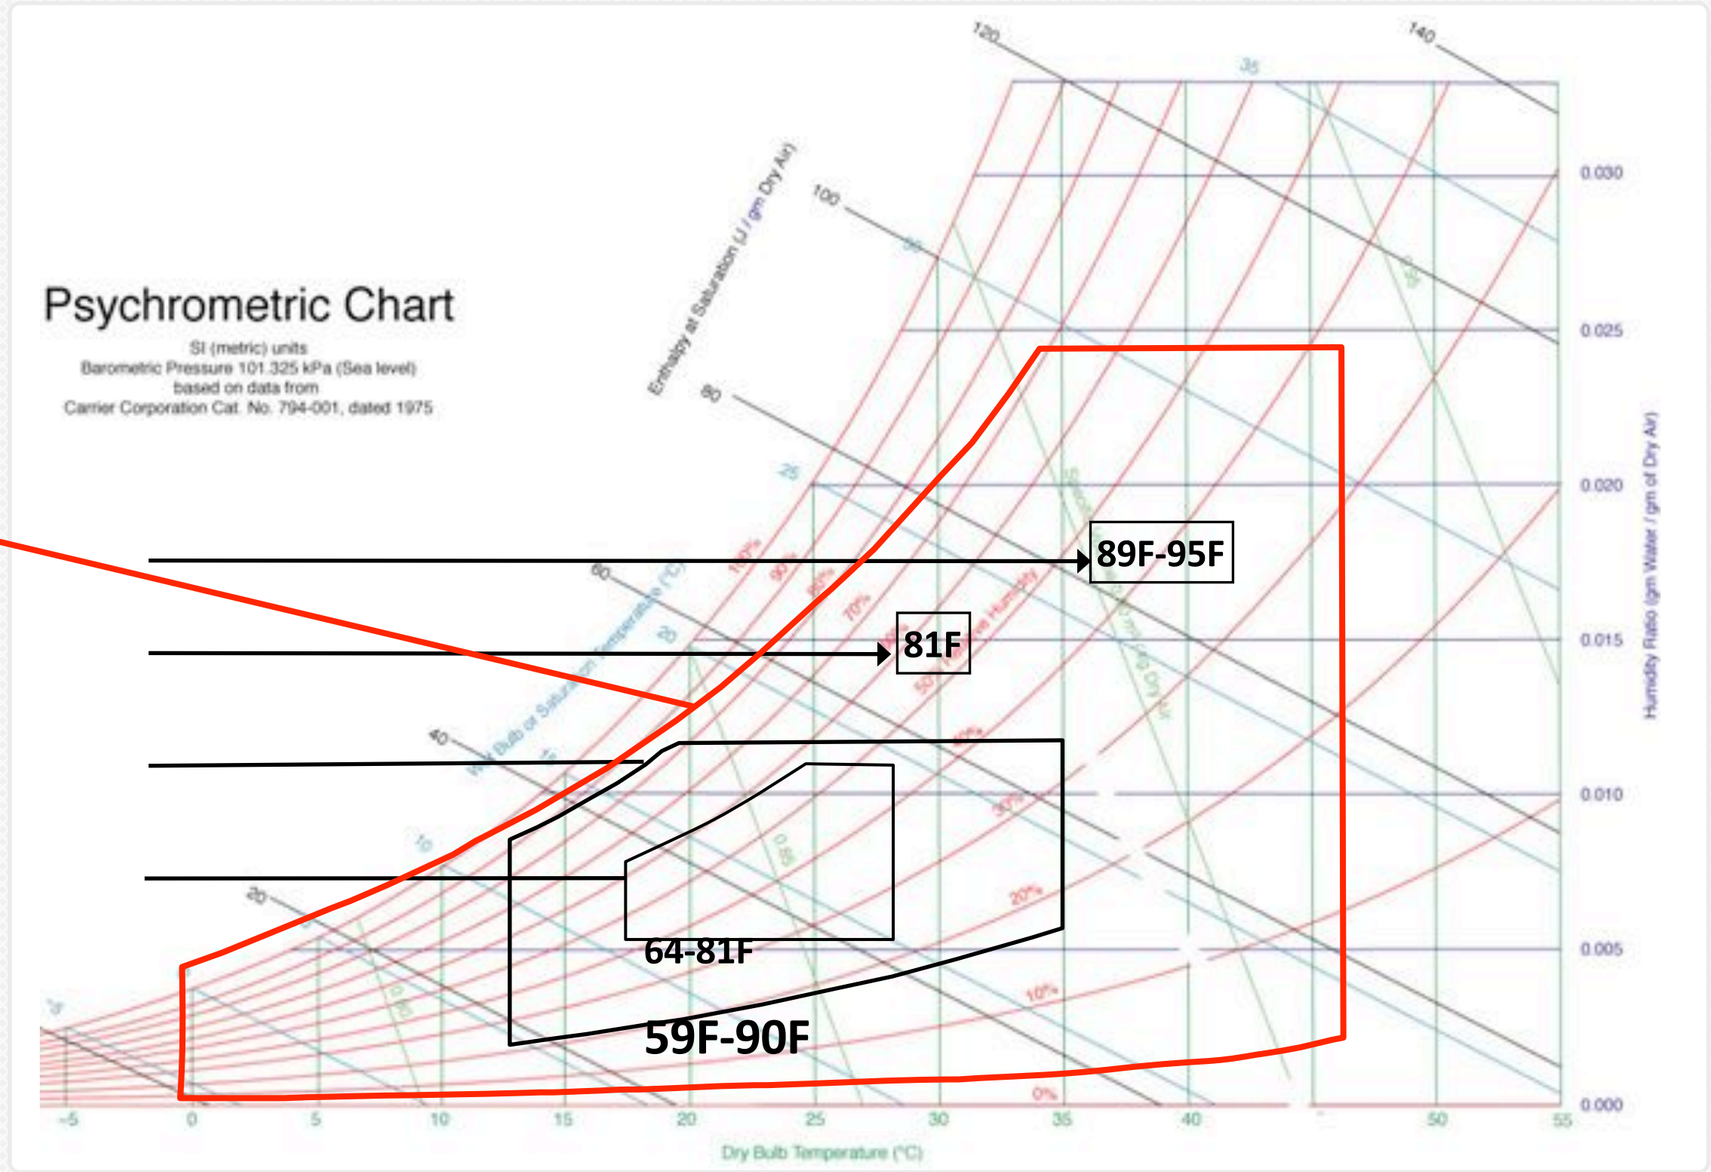

Most Datacenters Still Run Cold

Most datacenters run down in this range

ASHRAE 2008 Recommended Class 1

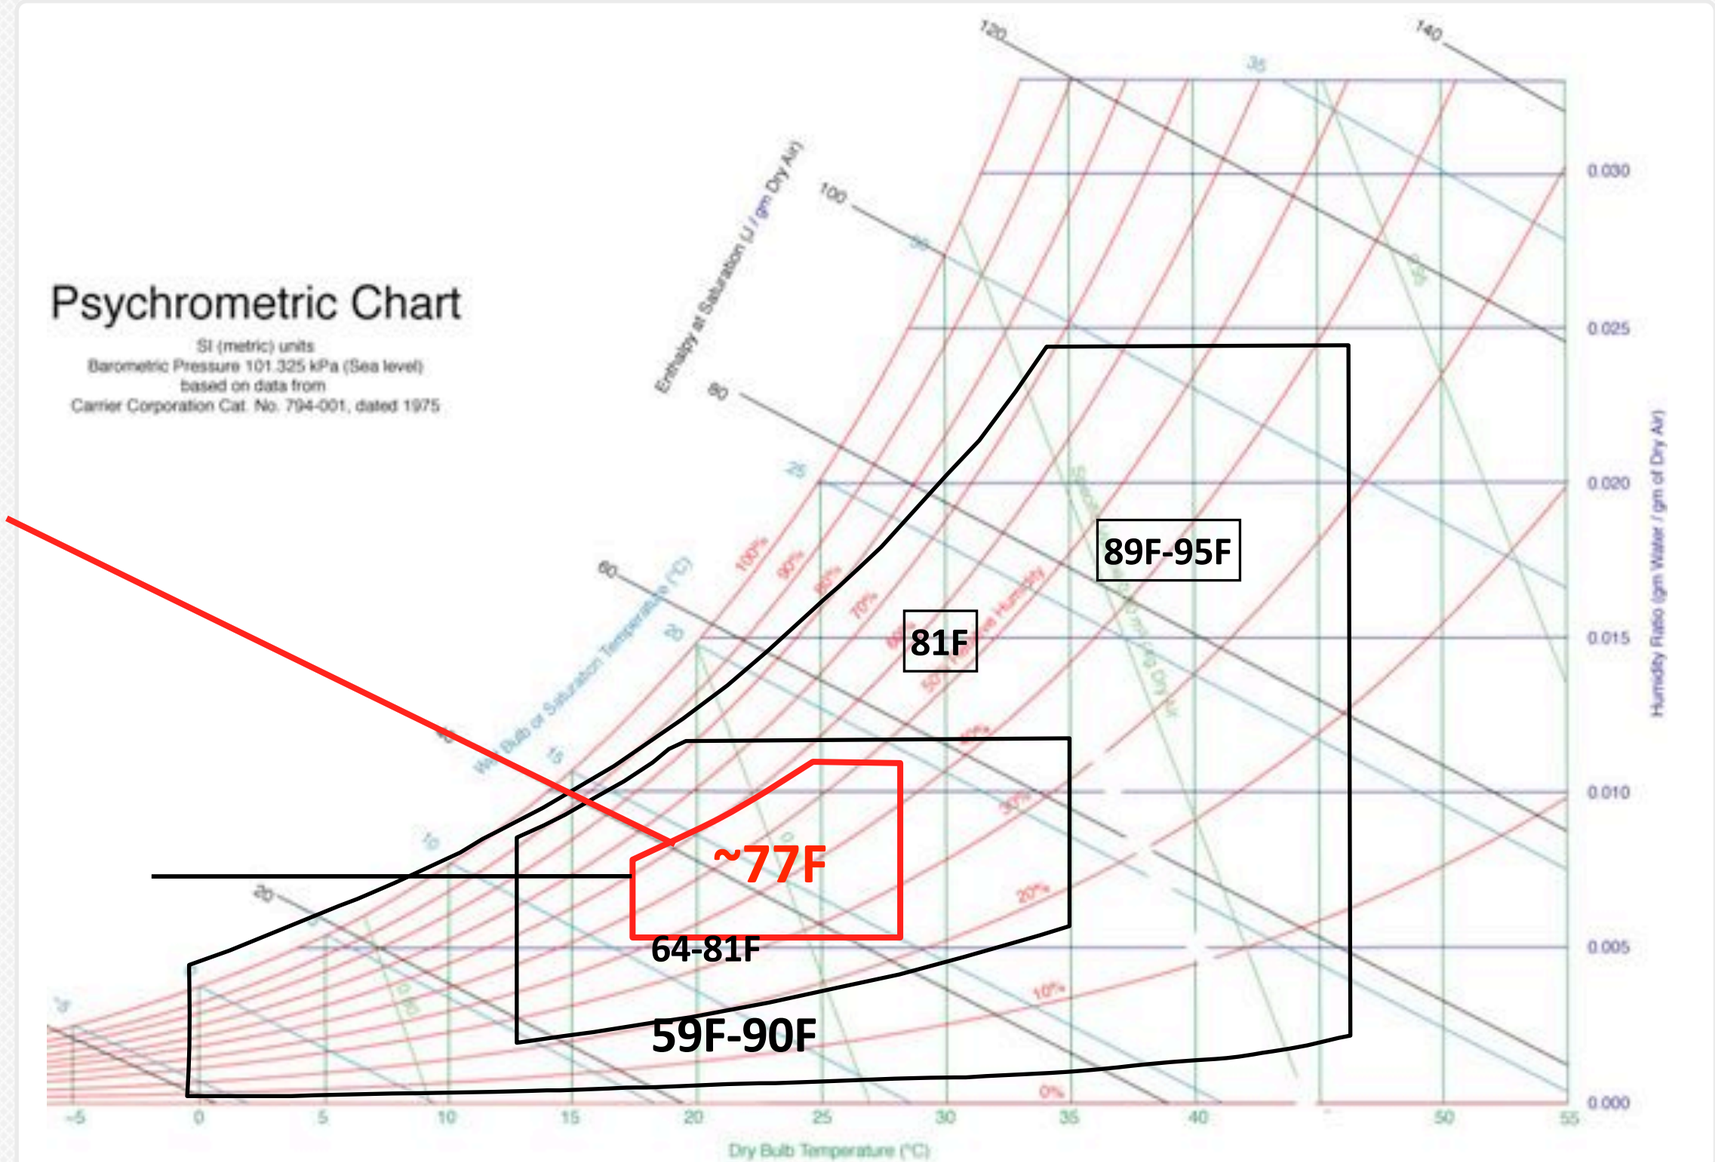

Avoiding Air Conditioning

- Component temps specs higher than historically hottest place on earth
  - Al Aziziyah, Libya: 136F/58C (1922)
- Just a mechanical engineering problem
  - More air or better mechanical designs
- Tradeoff: semi-conductor leakage & power to move more air vs cooling savings
- Currently available equipment temp limits:
  - 40C/104F: CloudRack C2 & most net gear
  - 35C/95F: Most of the server industry

Hard Drives: 7W- 25W
Temp Spec: 50C-60C

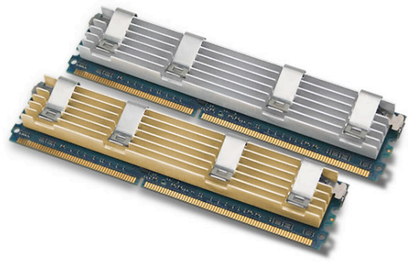

Memory: 3W - 20W
Temp Spec: 85C-105C

I/O: 5W - 25W
Temp Spec: 50C-60C

Processors/Chipset: 40W - 200W
Temp Spec: 60C-70C

Thanks to Ty Schmitt& Giovanni Coglitore

Innovative Shell Designs

- Evaporative cooling only
  - High pressure misting on right
  - Damp media design below
- Full building ductless cooling

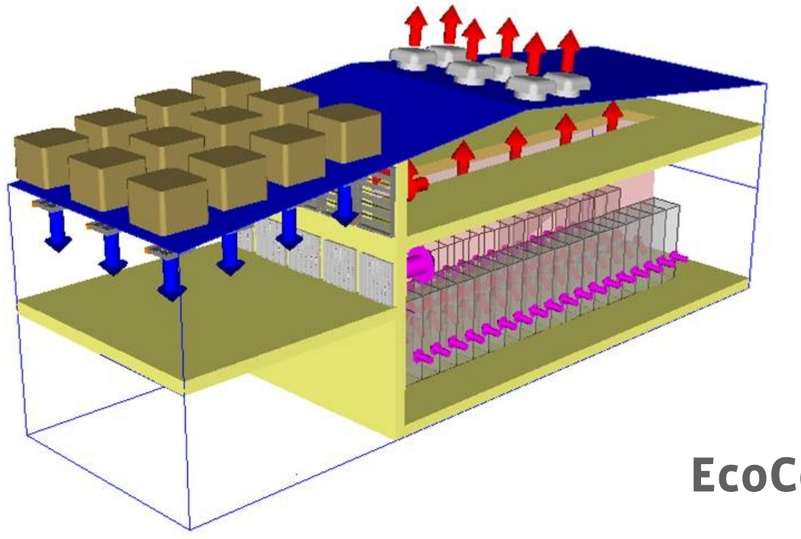

EcoCooling

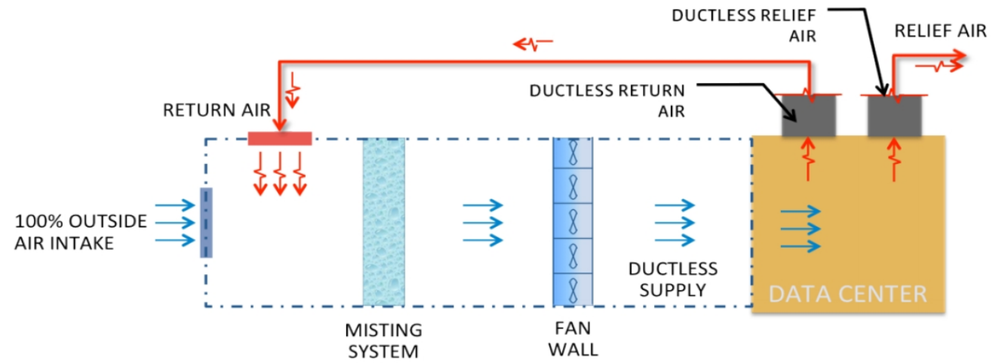

Facebook Prineville above & below

Modular and Pre-fab DC Designs

Microsoft ITPAC

- Fast & economic deployments
- Sub-1.2 PUE designs
- Air-side economized
  - In some cases no mechanical cooling
- ISO standard shipping containers offered by Dell, HP, SGI, IBM, …

Amazon Perdix

Questions?

- **Slides will be posted to:**
  - http://mvdirona.com/jrh/work
- **Perspectives Blog:**
  - http://perspectives.mvdirona.com/
- **Email:**
  - James@amazon.com