Pushing the limits of Datacenter Efficiency: Challenges and Opportunities

Kushagra Vaid
Partner Director, Hardware Engineering
Microsoft Online Services Division

Factors for Datacenter "Green-ness"

- Power delivery source (Coal, Hydro, Solar, …)
- Power delivery efficiency (PUE)
- Construction materials (Concrete, Steel …)
- Cooling techniques (Air, Water, Freon, …)
- Power consumption (Server, Network, …)
- Equipment disposal at decommissioning time

Factors for Datacenter "Green-ness"

- **Power delivery source (Coal, Hydro, Solar, …)**
- Power delivery efficiency (PUE)
- Construction materials (Concrete, Steel …)
- Cooling techniques (Air, Water, Freon, …)
- Power consumption (Server, Network, …)
- Equipment disposal at decommissioning time

Power delivery sources

- "Dirty" sources
  - Fossil fuels (Coal, Natural Gas), Nuclear
- Renewable sources
  - Solar, Wind, Hydro
  - Biofuel, Geothermal, Marine
- Tradeoffs
  - Energy costs and variability over time
  - Energy Provisioning costs and Infrastructure depreciation
  - Available options at Datacenter site location
  - Energy source reliability (max, avg, min)
  - Tax credits, PR

Factors for Datacenter “Green-ness”

- Power delivery source (Coal, Hydro, Solar, …)
- **Power delivery efficiency (PUE)**
- Construction materials (Concrete, Steel …)
- Cooling techniques (Air, Water, Freon, …)
- Power consumption (Server, Network, …)
- Equipment disposal at decommissioning time

Power delivery efficiency

- PUE = Total Facility Power / IT Equipment Power
  - Measure PUE at Server Inlet
- Components affecting efficiency: PDUs, UPS, Transformers, PSUs

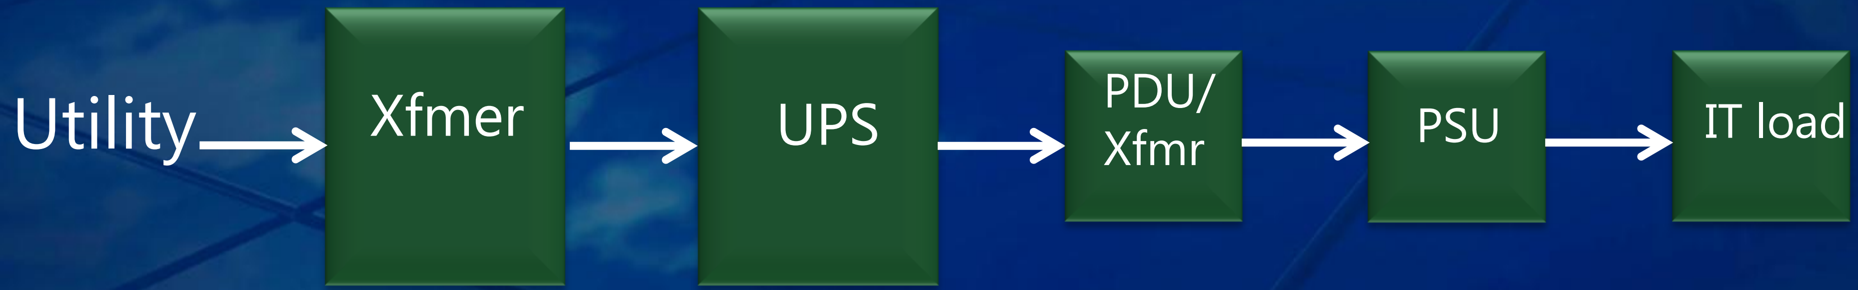

- Inefficient designs could have PUE ~2.0
  - i.e. for 1 watt available for IT load, 1 watt is wasted in overhead!
- Achieving high efficiencies
  - Minimize levels of power conversion
  - E.g. 480V 3Ø distribution, In-rack UPS, 95%+ Eff. PSUs
  - Facility PUE < 1.2 already achievable

Factors for Datacenter "Green-ness"

- Power delivery source (Coal, Hydro, Solar, …)
- Power delivery efficiency (PUE)
- **Construction materials (Concrete, Steel …)**
- **Cooling techniques (Air, Water, Freon, …)**
- Power consumption (Server, Network, …)
- Equipment disposal at decommissioning time

Traditional Datacenters (2007 era)

Airside economized Data Center (2011)

- Modular Datacener design
- Steel structure, concrete pad
- No mechanical cooling
- Ultra efficient water utilization
- Focus on renewable materials
- PUE < 1.2 achievable

Datacenter Operational Energy Use

Traditional

Source: EYP Mission Critical Facilities Inc., New York

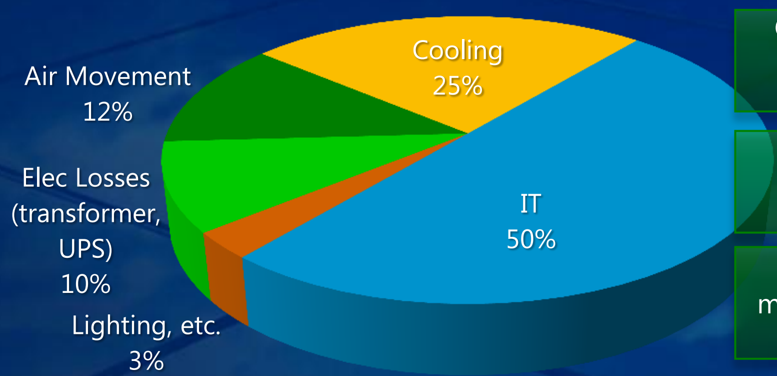

- Offline UPS technologies can drive Electrical losses substantially down
- Widening environment can remove chillers and drive Cooling to zero
- Virtualization, active power management increase IT return on investment

Modular

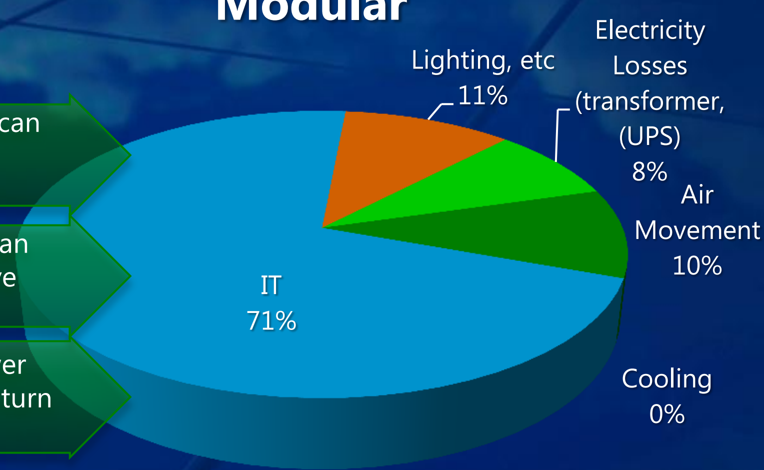

Traditional - PUE 2.0

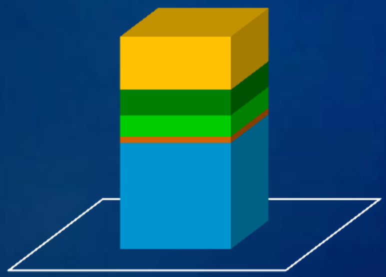

Optimization across Multiple functional Domains (Mechanical/Electrical & IT Utilization)

Modular - PUE 1.2

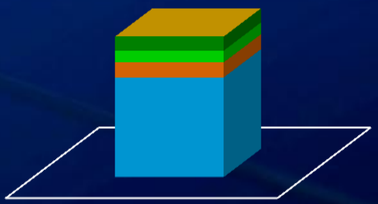

Factors for Datacenter "Green-ness"

- Power delivery source (Coal, Hydro, Solar, …)
- Power delivery efficiency (PUE)
- Construction materials (Concrete, Steel …)
- Cooling techniques (Air, Water, Freon, …)
- **Power consumption (Server, Network, …)**
- Equipment disposal at decommissioning time

Datacenter Energy usage breakdown

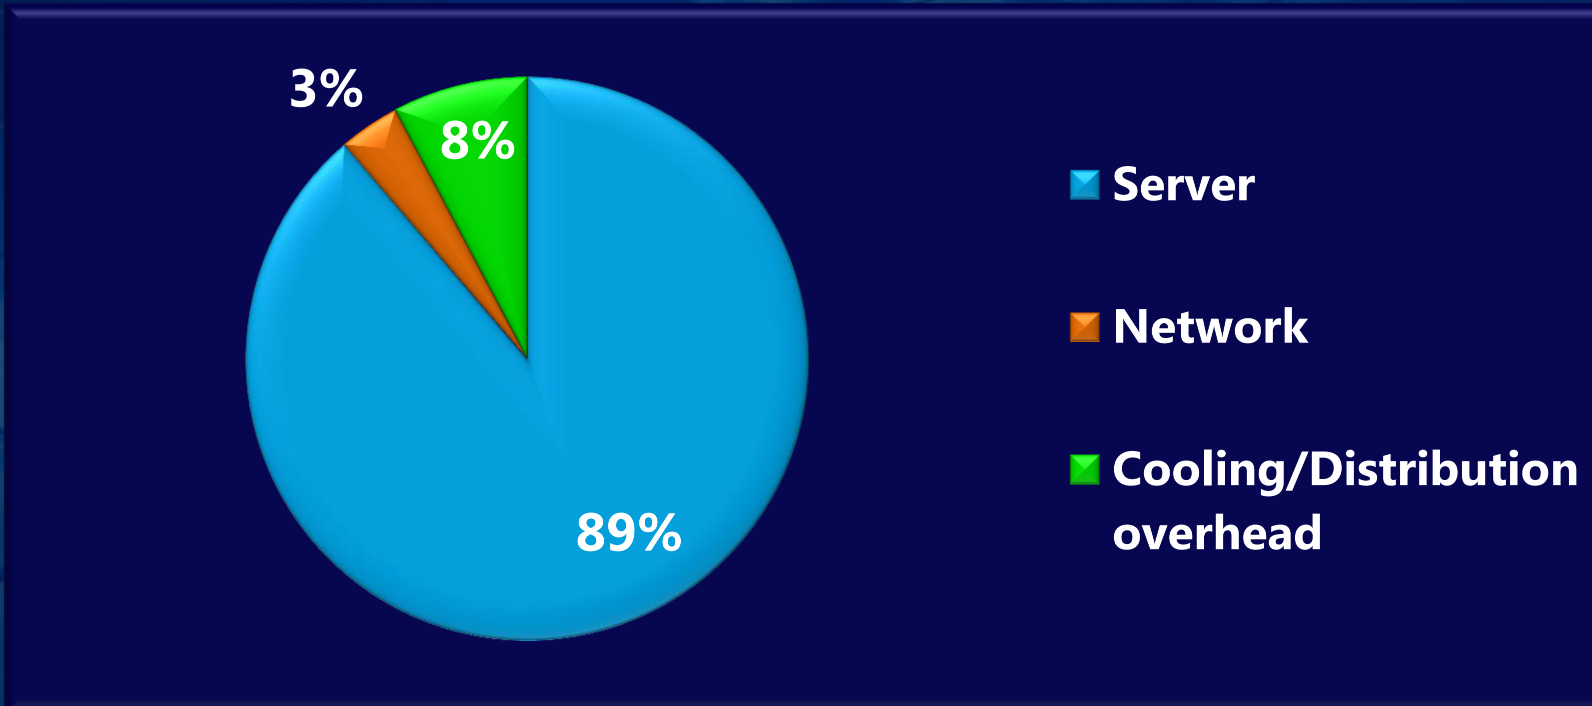

Power consumption reduction across all domains is necessary

Highest ROI for Server Power reduction

Reducing system power footprint

- Rightsizing systems design, E.g. LV CPU/Mem, SSDs for IOPS/W

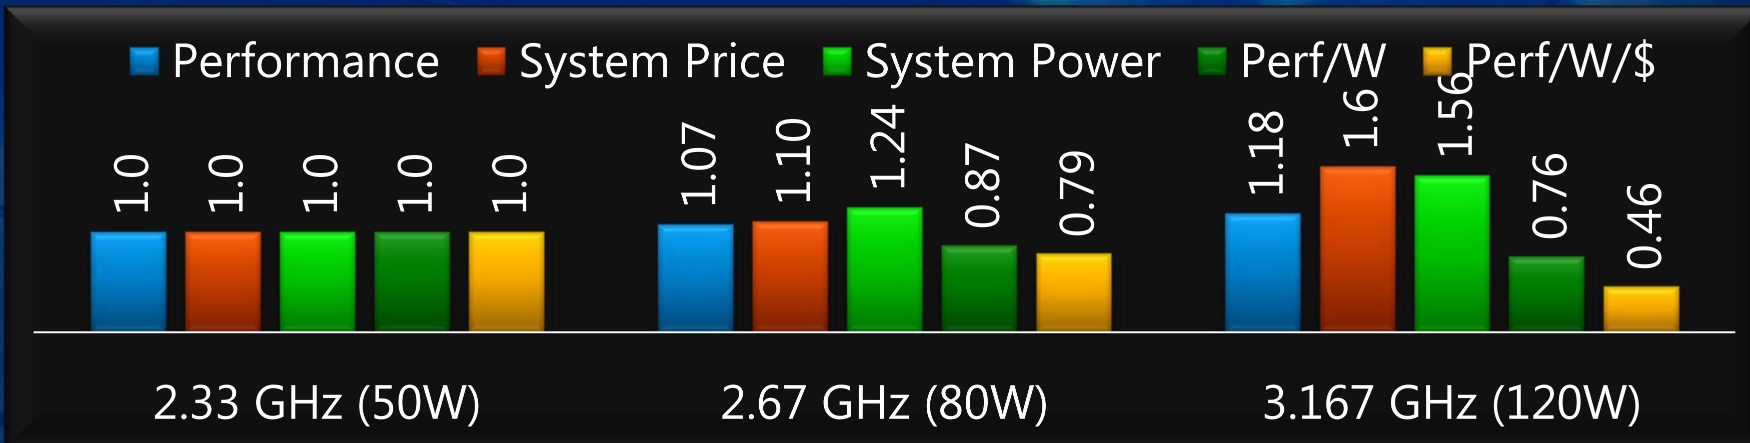

- Power-aware applications and system software
  - E.g. Core parking, Transaction load-balancing

- Silicon power reduction techniques
  - E.g. V-f scaling across independent domains, Aggressive clock gating

Summary so far

- Industry has made several advances on multiple aspects of datacenter efficiency for driving Green Computing
  - But work has mostly been done in independent silos
- Need to drive convergence and vertical integration for exploiting opportunities across functional silos
  - Challenging at industry level given diversity of solutions, lack of standardization, redundancy requirements etc.
- What role can the Large-scale Cloud Service providers play for spearheading the next wave of Green computing?
  - Open knowledge sharing for sharing best practices with industry
  - Solutions still inconsistent across providers, but industry may benefit as a whole from advancements

What next?

Pushing the limits – *The Disappearing Datacenter*

Brick-and-mortar
Chilled Water cooling

Concrete pad, Steel structure
Adiabatic cooling

No structure, open frame
Free Air-cooling

Possible only via HW/SW co-design

Datacenter Efficiency Opportunities

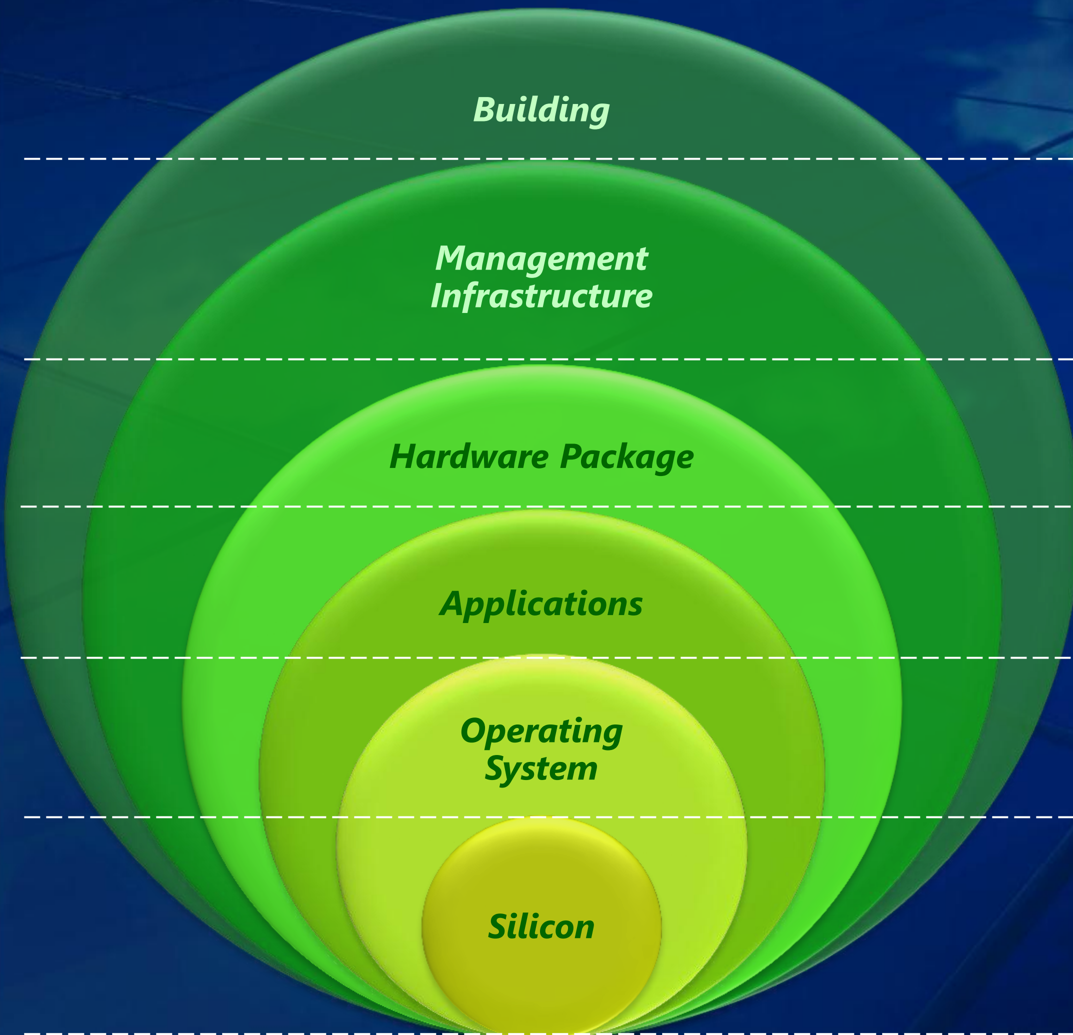

- Simplify Datacenter construction
- Improve Datacenter PUE
- **True Free-Air cooling**

- Centralized Power Management

- Rightsizing Server Components
- Efficient Power Systems

- **Energy Efficient Applications**
- **Self-healing and resiliency**

- Platform Power Management

- **Low-power system design**

Need to think holistically across the entire services stack

True Free-Air Cooling

- Use unconditioned outside air for cooling servers
- Benefits
  - Water-free operation → Environmental friendly
  - Reduced datacenter capex and opex costs
- High temperature operation has implications
  - Silicon Leakage Power
  - Component Reliability
    - High Temperature failures
    - Particulates/Contaminants
  - Heat removal and Airflow

Climate zones and Worst case temperature

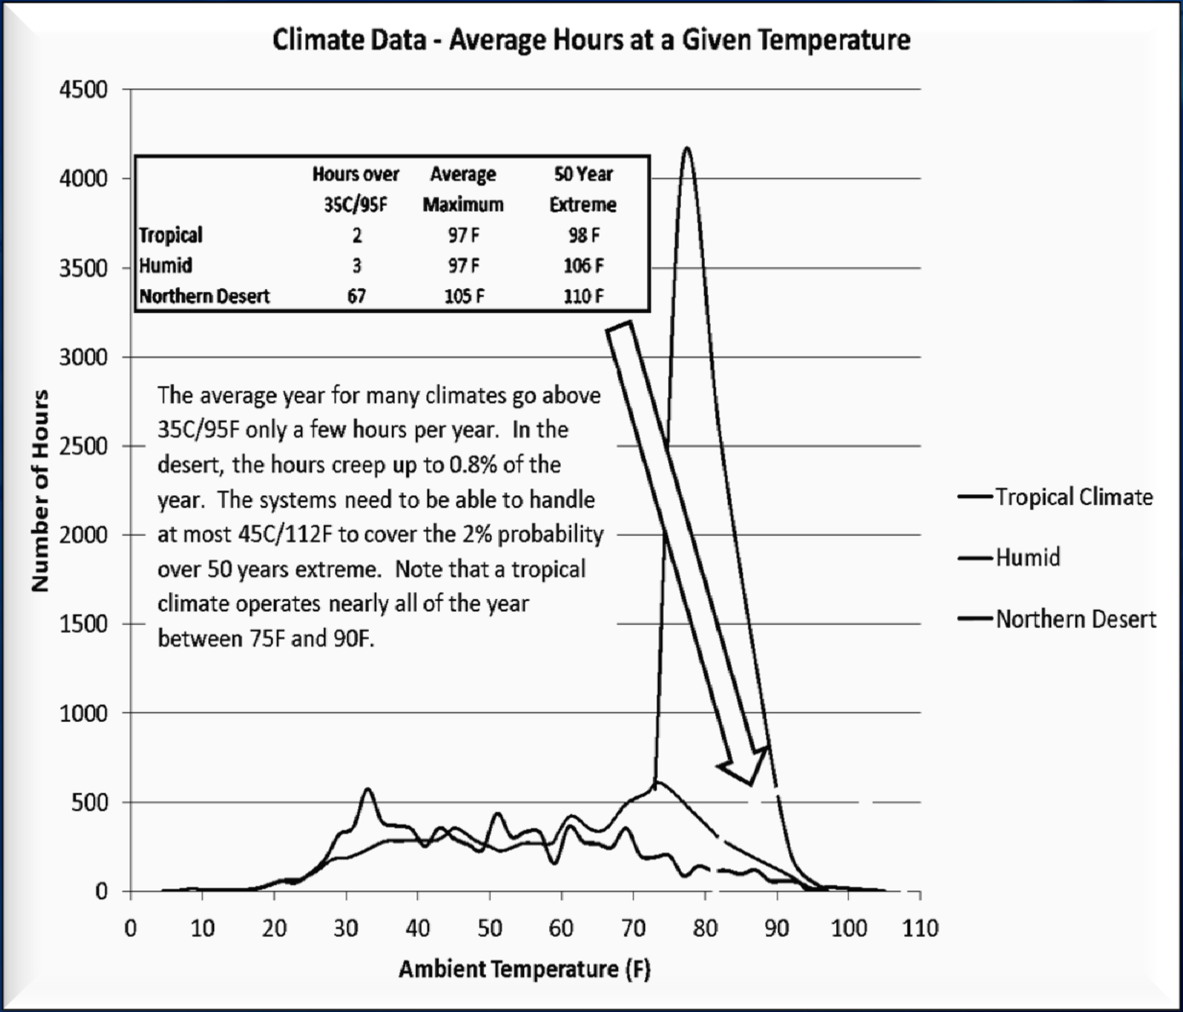

Chart shows typical temperature trends for three high temperature climates zones:

-- Tropical

-- Humid

-- Northern Desert

Very few hours of the year go over 35C/95F

Should adapt well to Server Rating of 10-35C/50-95F

System operation at High Temperatures

$$\mathit{Airflow\ to\ Cool\ a\ Server} \propto \frac{\mathit{Server\ Power}}{(\mathit{Exhaust\ Temp} - \mathit{Inlet\ Temp})}$$

More airflow is required with increase in temperature to compensate for …

- ❑ Increased heat removal
- ❑ Higher dynamic power (leakage)

One possible mitigation: Use System Power Capping

- ❑ Invoked above specific Inlet Temperature (e.g. 35C/95F)
- ❑ Reduces Server Power and hence Exhaust Temperature
- ❑ Slight performance impact under high workload conditions

What about reliability?

- HDDs most sensitive to temperature related failures
- Max rated temp 60C; but AFR increases 2x from 40C→60C!
- Mitigate via server design, e.g. HDD temp < Ambient+4C

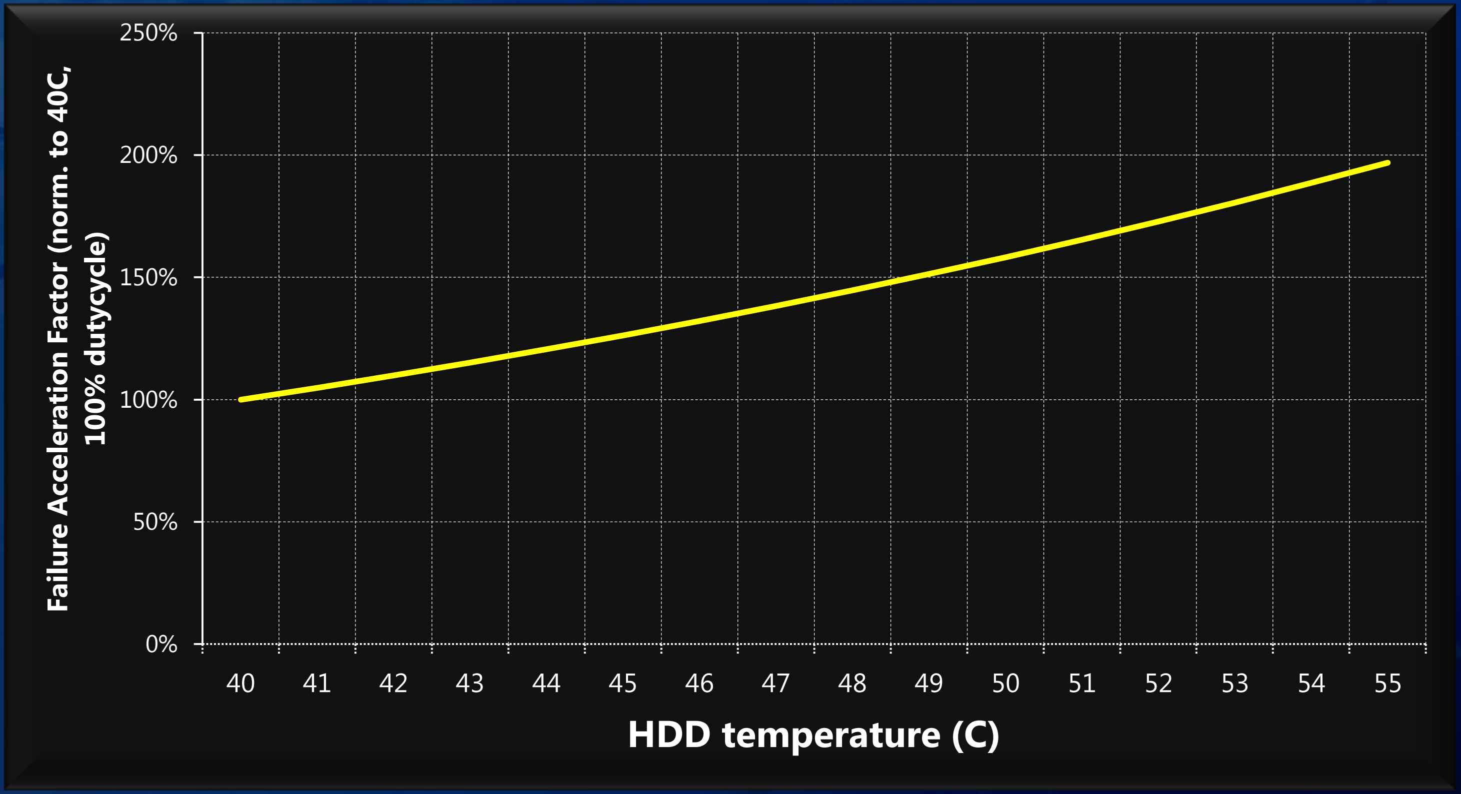

Architectural Innovations for system power reductions

- SoC designs – x64 and ARM

→

- NVRAM in memory hierarchy

Application manages Data locality

Memory (1st Cache) ↔ NVRAM (2nd Cache) ↔ HDD (App. Data)

- Distributed switching - Integrated on-chip fabric

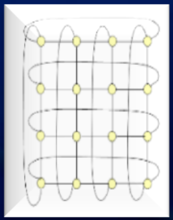

Torus/Mesh

- Packaging and cooling innovations

Self-Healing infrastructure

- 100% Availability doesn’t exist - Hardware will fail
- Building reliable cloud services involves planning for failures
  - Hardware failures (e.g. hard drives, power supplies)
  - Software bugs (e.g. BIOS, Drivers, App)
  - Human errors (e.g. Maintenance, Config changes)
  - Design for nuisance outages (maintenance) and utility power outages
- “Green” Benefits
  - Simplified Datacenters /IT gear → fewer materials and maintenance

Summary

- Progress has been made over past decade for Datacenter efficiency
  - …but mostly in independent functional silos
- Opportunities for further efficiency advances are in vertically integration the HW/SW ecosystem
  - Simple/Cheap Hardware & Reliable Software
  - Drive Open knowledge sharing for best practices with broader industry

Resources:

Datacenter infrastructure best practices:
http://www.globalfoundationservices.com/infrastructure/index.html

Microsoft Datacenter Blogs
http://blogs.technet.com/b/msdatacenters/