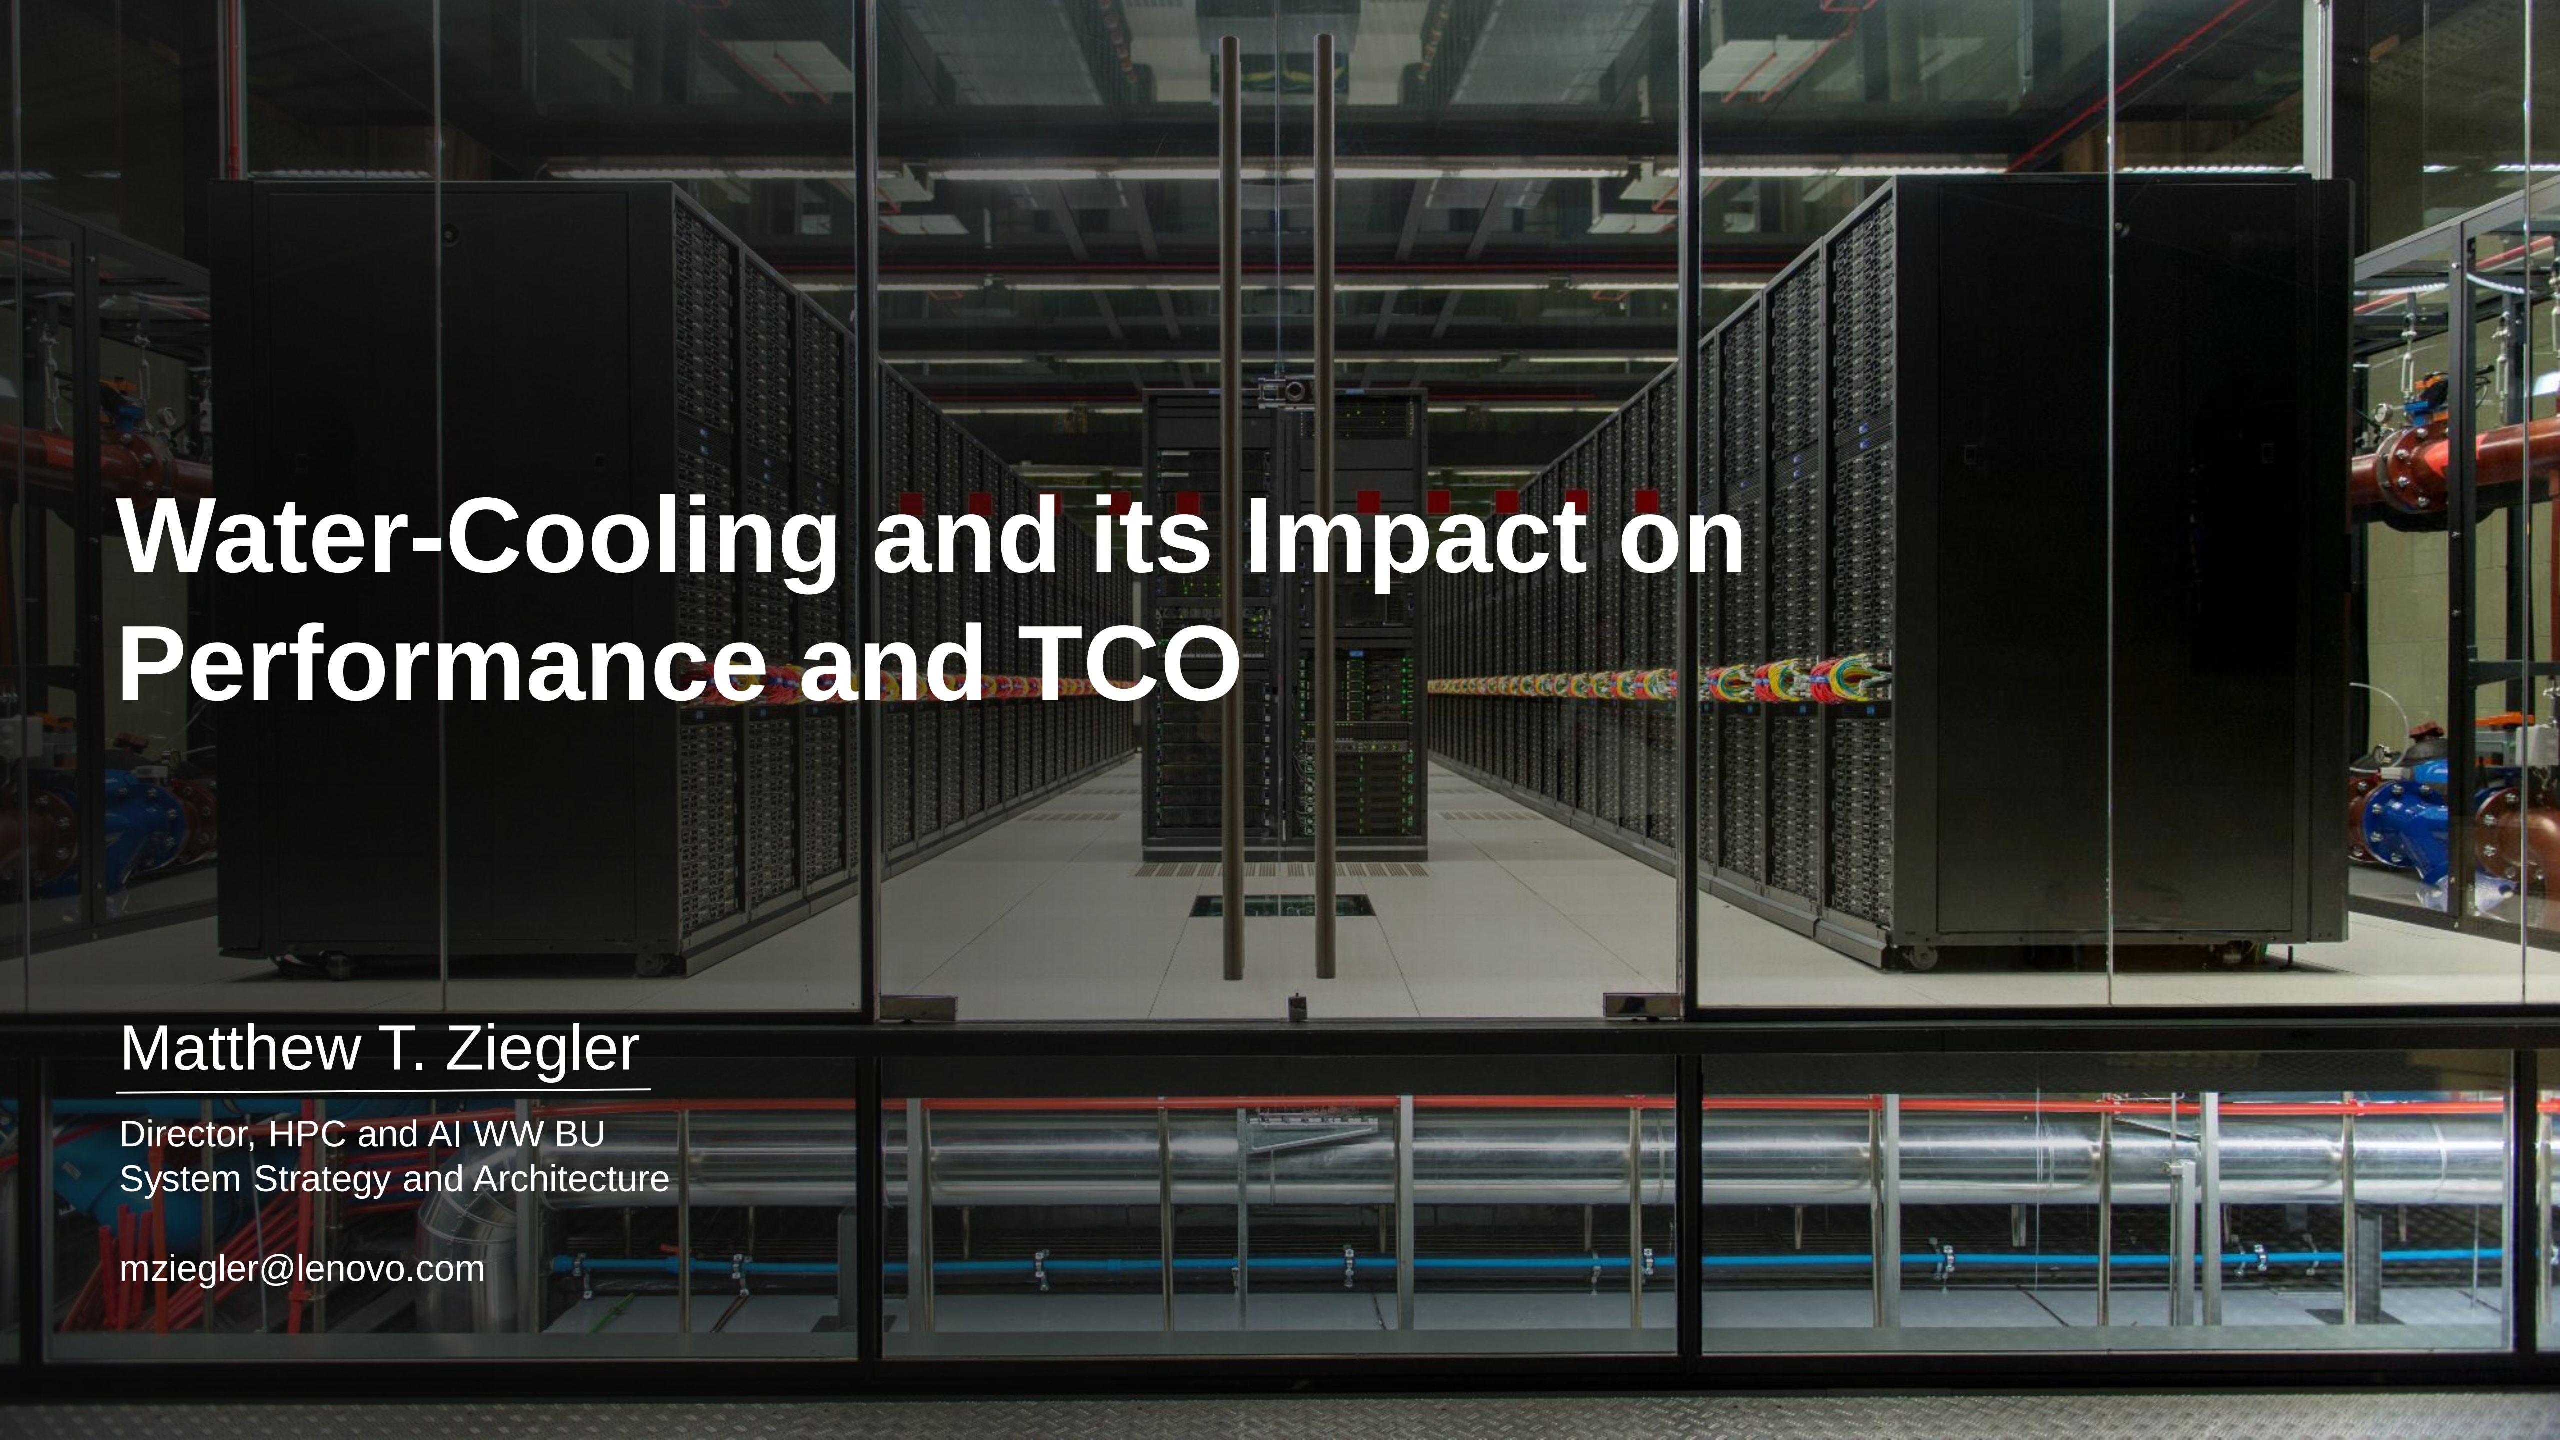

Water-Cooling and its Impact on Performance and TCO

Matthew T. Ziegler

Director, HPC and AI WW BU
System Strategy and Architecture

mziegler@lenovo.com

Agenda

- Why water cooling is becoming important?
- Server Power Trends
- Performance Impacts
- OPEX effects

Why do we have a problem?

- Higher TDP Processors
- Data Center Power/Space limits
- High Electricy Cost
- Performance is Power/Thermal capped
- Waste Heat Reuse

Lenovo

Power/Heat is changing the Datacenter Paradigm

Intel Xeon Server processor history

Release date	Code	Processor	core/chip	TDP(W)	Spec FP	Spec_fp Rate
2006/6/26	Woodcrest	Intel Xeon 5160	2	80	17.7	45.5
2007/11/12	Harpertown	Intel Xeon x5460	4	120	25.4	79.6
2009/3/30	Nehalem	Intel Xeon x5570	4	95	43.8	202
2010/3/16	Westmere-EP	Intel Xeon x5690	6	130	63.7	273
2012/5/1	SandyBridge	Intel Xeon E5-2690	8	135	94.8	507
2014/1/9	IvyBridge	Intel Xeon E5-2697v2	12	130	104	696
2014/9/9	Haswell	Intel Xeon E5-2699v3	18	145	116	949
2015/3/9	Bradwell	Intel Xeon E5-2699v4	22	145	128	1160
2017/7/11	Skylake	Intel Xeon Platinum 8180	28	205	155	1820

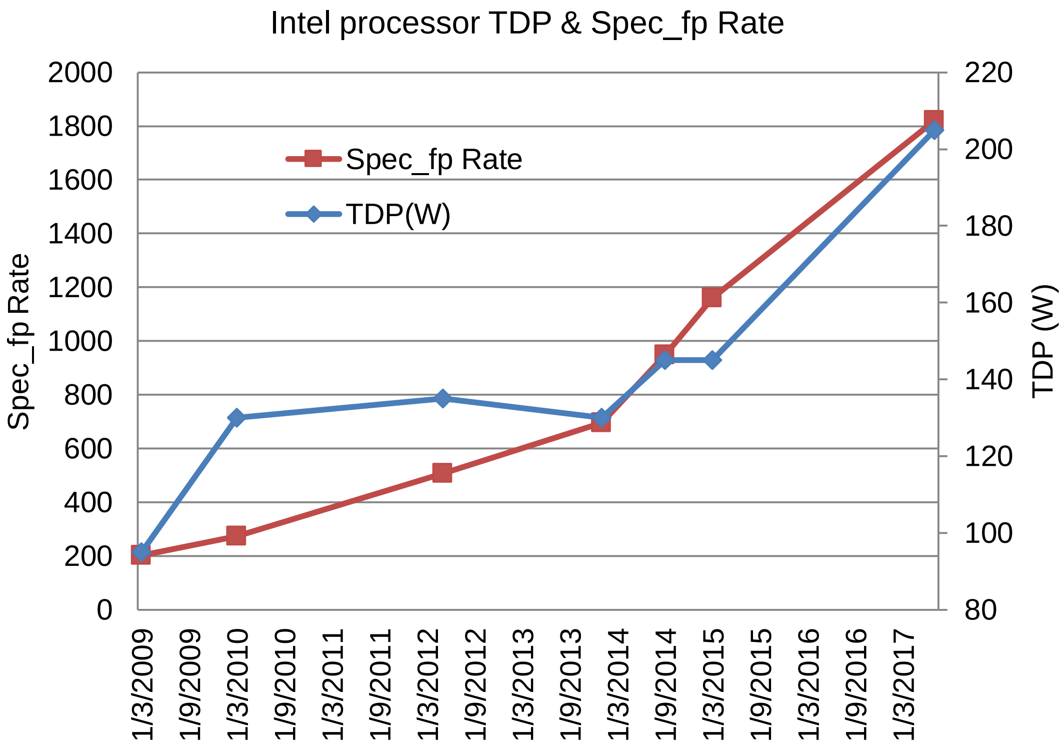

Processor performance trend

- Spec_fp rate with 2 processors/node has increased 40 times the past 11 years (2006 – 2017).
- The number of cores on the chip increase 14 times.
- After being flat, since 2014 TDP increases linearly with Spec_fp rate.
- Current maximum TDP is 205W. Knights Mill Xeon phi processor will be 305 W

To sustain increased performance servers will have to be less dense or use new cooling technology

2018 Lenovo. All Rights Reserved

Power Density Ever Increasing

Eli Lilly and Company
(#75, Nov 2006)
BladeCenter HS21 w/ Xeon 5160
2C 3.0GHz 80W

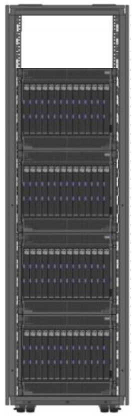

- **Rack: 56 Nodes, 224 Cores**
- **SPECfp2006 Rate: 2.548**
- **RackPower: ~20kW**

BSC – Mare Nostrum
(#16, Nov 2017)
Lenovo SD530 w/ Xeon 8160
24C 2.1GHz 150W

- **Rack: 72 Nodes, 3.456 Cores**
- **SPECfp2006 Rate: 110.160**
- **RackPower: ~33kW**

LRZ – SuperMUC-NG
(#?, Nov 2018)
Lenovo SD650 w/ Xeon 8174
24C 3.1GHz 240W mode

- **Rack: 72 Nodes, 3.456 Cores**
- **SPECfp2006 Rate: tbc**
- **RackPower: ~46kW**

How much heat can your DataCenter extract from a 19" rack?

2018 Lenovo. All Rights Reserved

The Future of Power and Performance

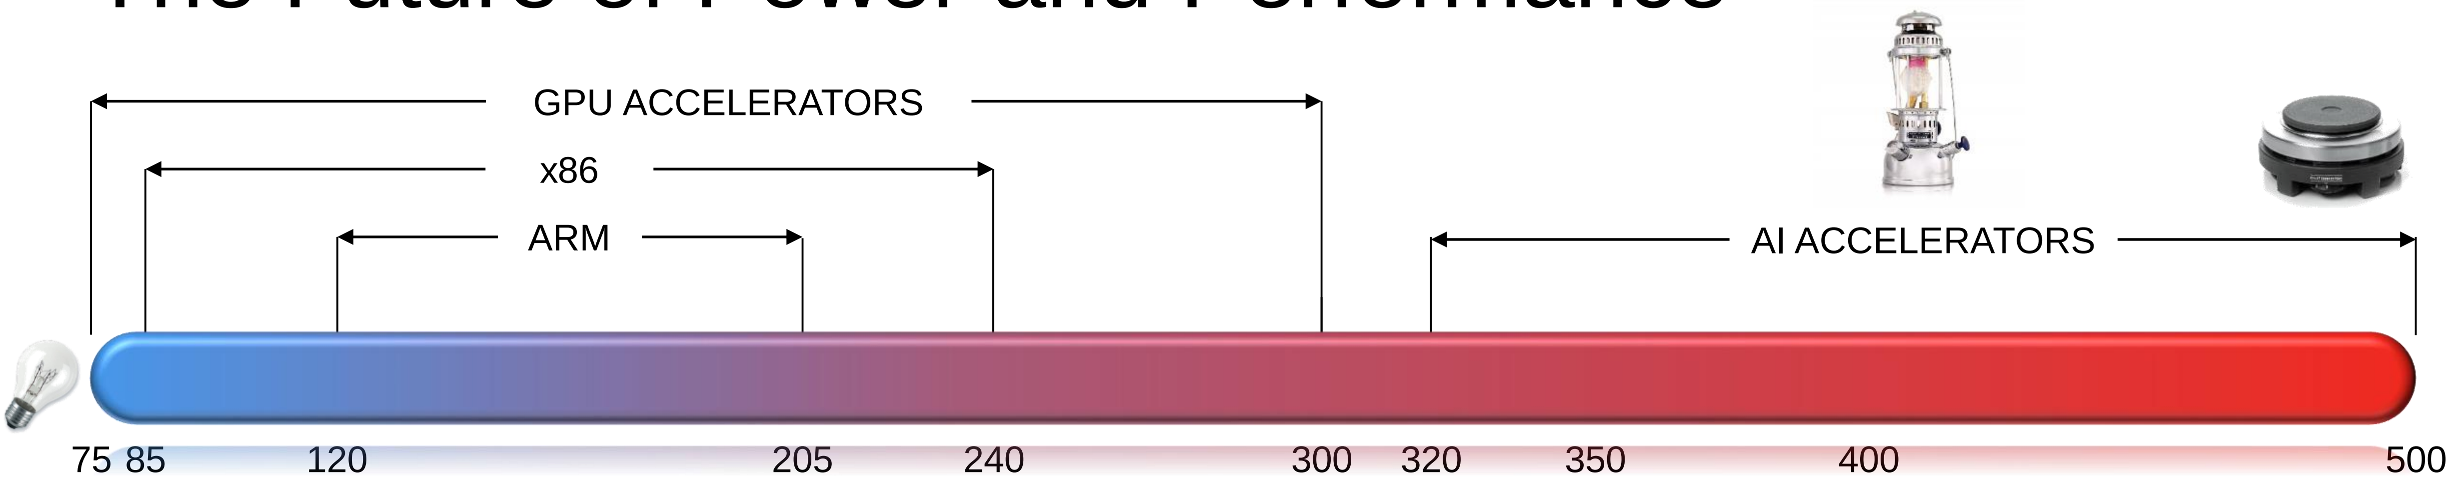

- Maintaining Moore’s Law + increased competition is resulting in higher processor power
- Increasing processor power, memory, NVMe adoption and I/O power growth will drive packaging and feature tradeoffs
- Rack power levels will challenge the data center – power delivery, heat handling, air flow delivery and floor loading
- **Smart thermal designs including water will become the norm**

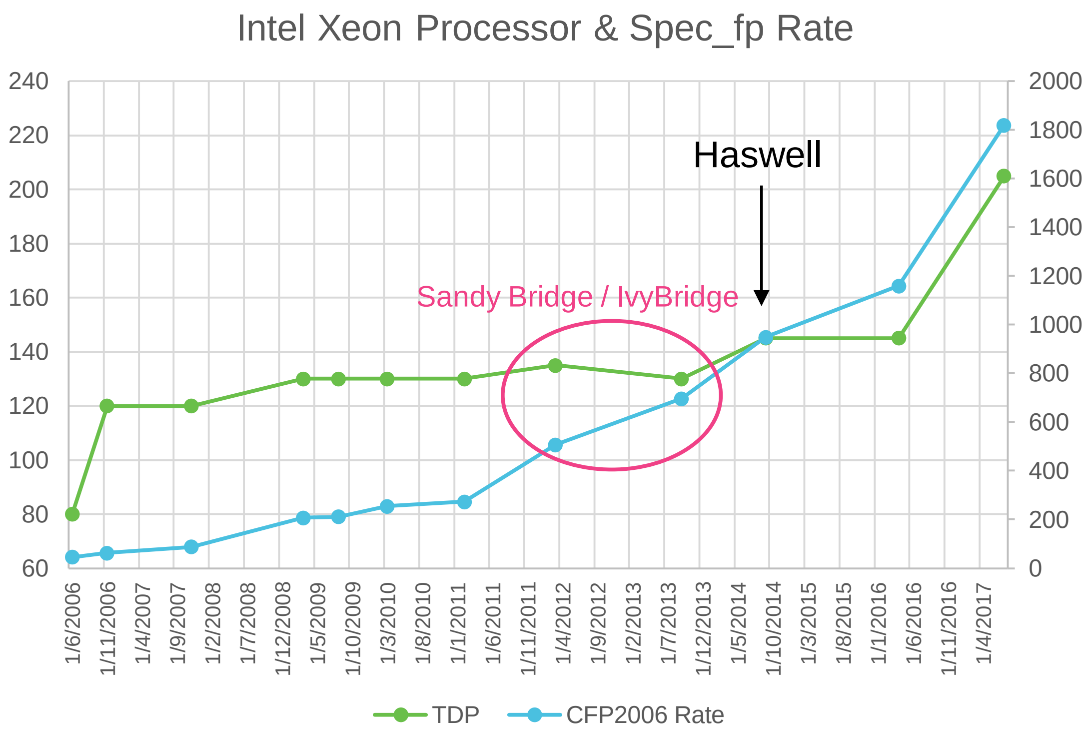

2018 Lenovo. All Rights Reserved

Silicon Roadmap Impacts on Air Cooling

- 100% of input power converted to heat: Exhaust Temp = Power/Airflow + Inlet Temp
- System airflow cannot keep up with silicon power increases: max 40-80 cfm per dense node
- Feature set tradeoffs will be required to fit in node thermal envelope:
  - Move I/O to front with Storage to reduce preheat issues?
  - Reduce superset of CPU and Memory power support: reduce # of DIMMs with high TPD CPUs?
  - Industry move to Direct Water Cooing?

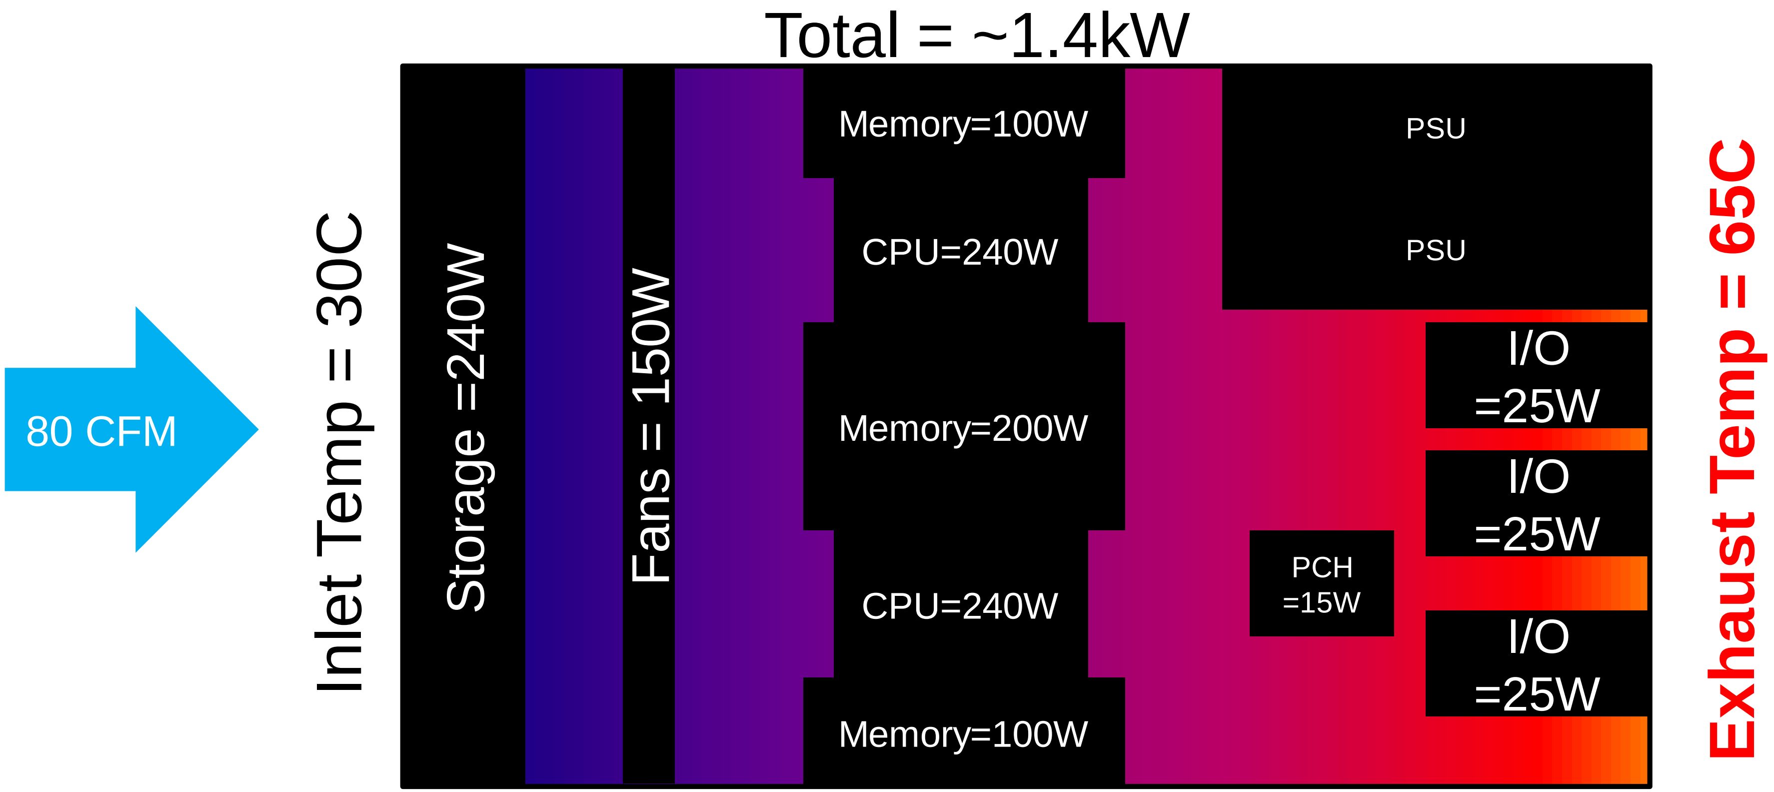

2018 Lenovo. All Rights Reserved

Server Power Trends – **ASHRAE* 2015-2020**

Market Requirements force IT manufacturers to maximize performance/volume creating high heat load/rack

Height	No. of Sockets	Heat Load / Chassis (watts)			Heat Load / 42U Rack			Increase 2010 to 2020
		2010	2015	2020	2010	2015	2020	
1U	1s	255	290	330	10,710	12,180	13,860	29%
	2s	600	735	870	25,200	30,870	36,540	45%
	4s	1,000	1,100	1,200	42,000	46,200	50,400	20%
2U	2s	750	1,100	1,250	15,750	23,100	26,250	67%
	4s	1,400	1,800	2,000	29,400	37,800	42,000	43%
4U	2s	2,300	3,100	3,300	23,000	31,000	33,000	43%
7U (Blade)	2s	5,500	6,500	7,500	33,000	39,000	45,000	36%
9U (Blade)	2s	6,500	8,000	9,500	26,000	32,000	38,000	6%
10U (Blade)	2s	8,000	9,000	10,500	32,000	36,000	42,000	31%

These rack heat loads will result in increased focus on improving data center ventilation solutions and localized liquid cooling solutions

*ASHRAE = American Society of Heating, Refrigerating, and Air-Conditioning Engineers. The group provides operating environment standards for datacenter operations.

2018 Lenovo. All Rights Reserved

Liquid Cooling Status

- Adoption of liquid cooling to date has been primarily driven by energy efficiency improvements and heat recovery
- Intel Purley Skylake 205W CPU TDPs has increased water cooling use to maintain density and for chassis reuse (e.g. Dell C6420 uses CoolIT for 205W TDP CPUs)
- High power CPU/GPU roadmaps will accelerate the adoption of liquid cooling
  - Component power is exceeding what can be cooled using forced convection at a node level
  - Rack level power is exceeding what can be cooled using forced convection at a data center level

2018 Lenovo. All Rights Reserved

Node Level Cooling Limits

- Component thermal requirements are exceeding what can be air cooled
  - Solutions with large heat sinks (≥ 2U) are possible, but with exponential fan power and acoustics

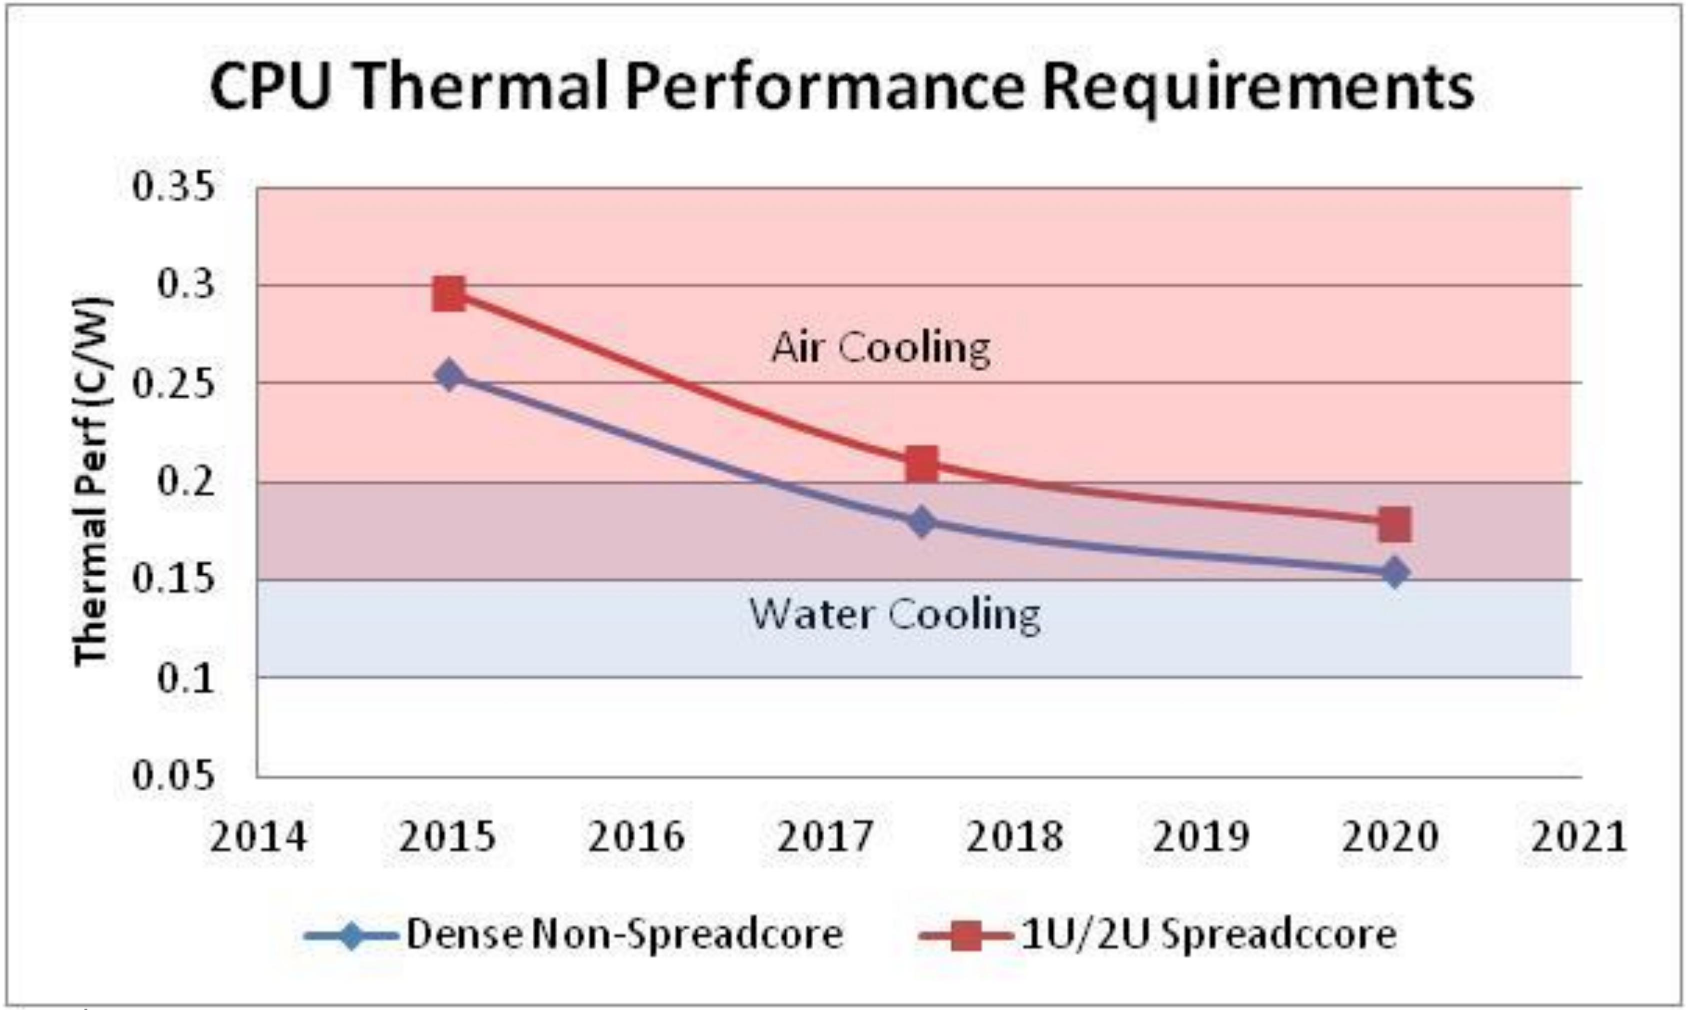

2018 Lenovo. All Rights Reserved

Rack Level Limits

- Node power density cannot be cooled at rack level = *Partial rack population or rack level power capping may be required*

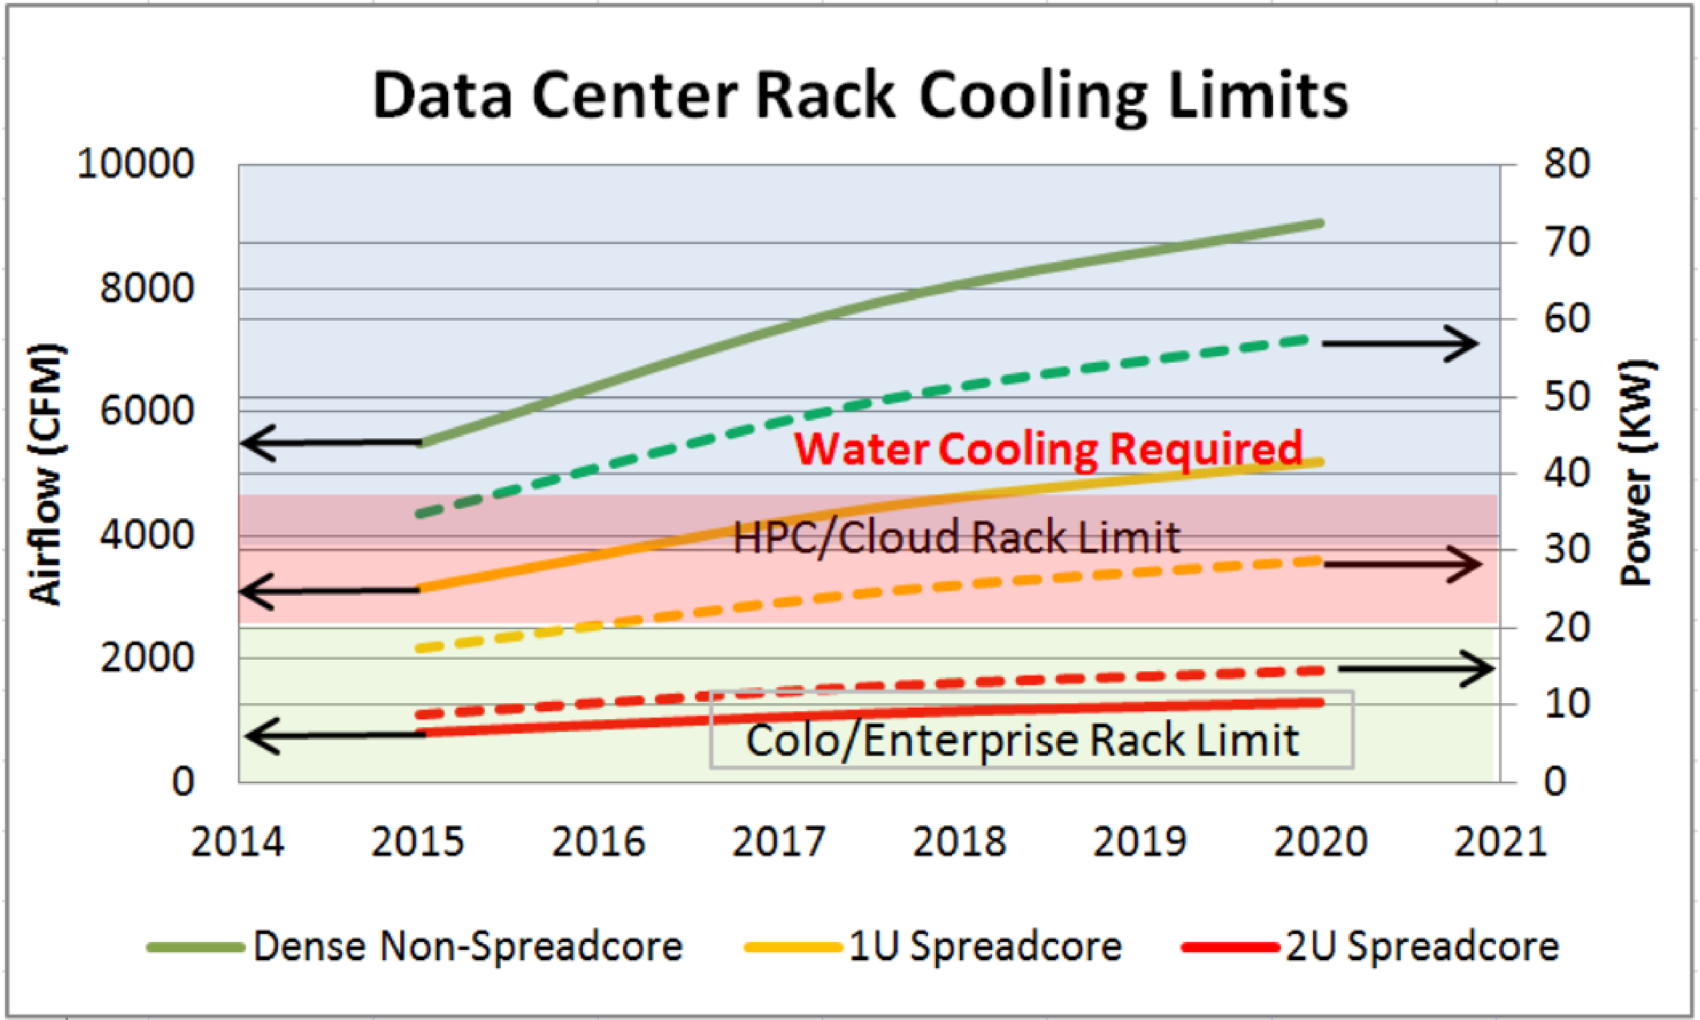

2018 Lenovo. All Rights Reserved

Rack Power by Segment through 2020

Co-location
350W per node
(2x120W TDP CPU, 12x16GB RDIMM, 2xSATA HDDs, 2x10GbE)

Enterprise
500W per node
(2x150W TDP CPU, 24x32GB RDIMM, 6x SATA HDDs, 2x10GbE)

Cloud
430W per node/2000W per chassis
(2x150W TDP CPU, 12x32GB RDIMM, 6x SATA HDDs, 2x10GbE)

HPC
550W per node/2500W per chassis
(2x205W TDP CPU, 12x32GB RDIMM, 2x SATA SDDs, 2x10GbE, 1xOPA)

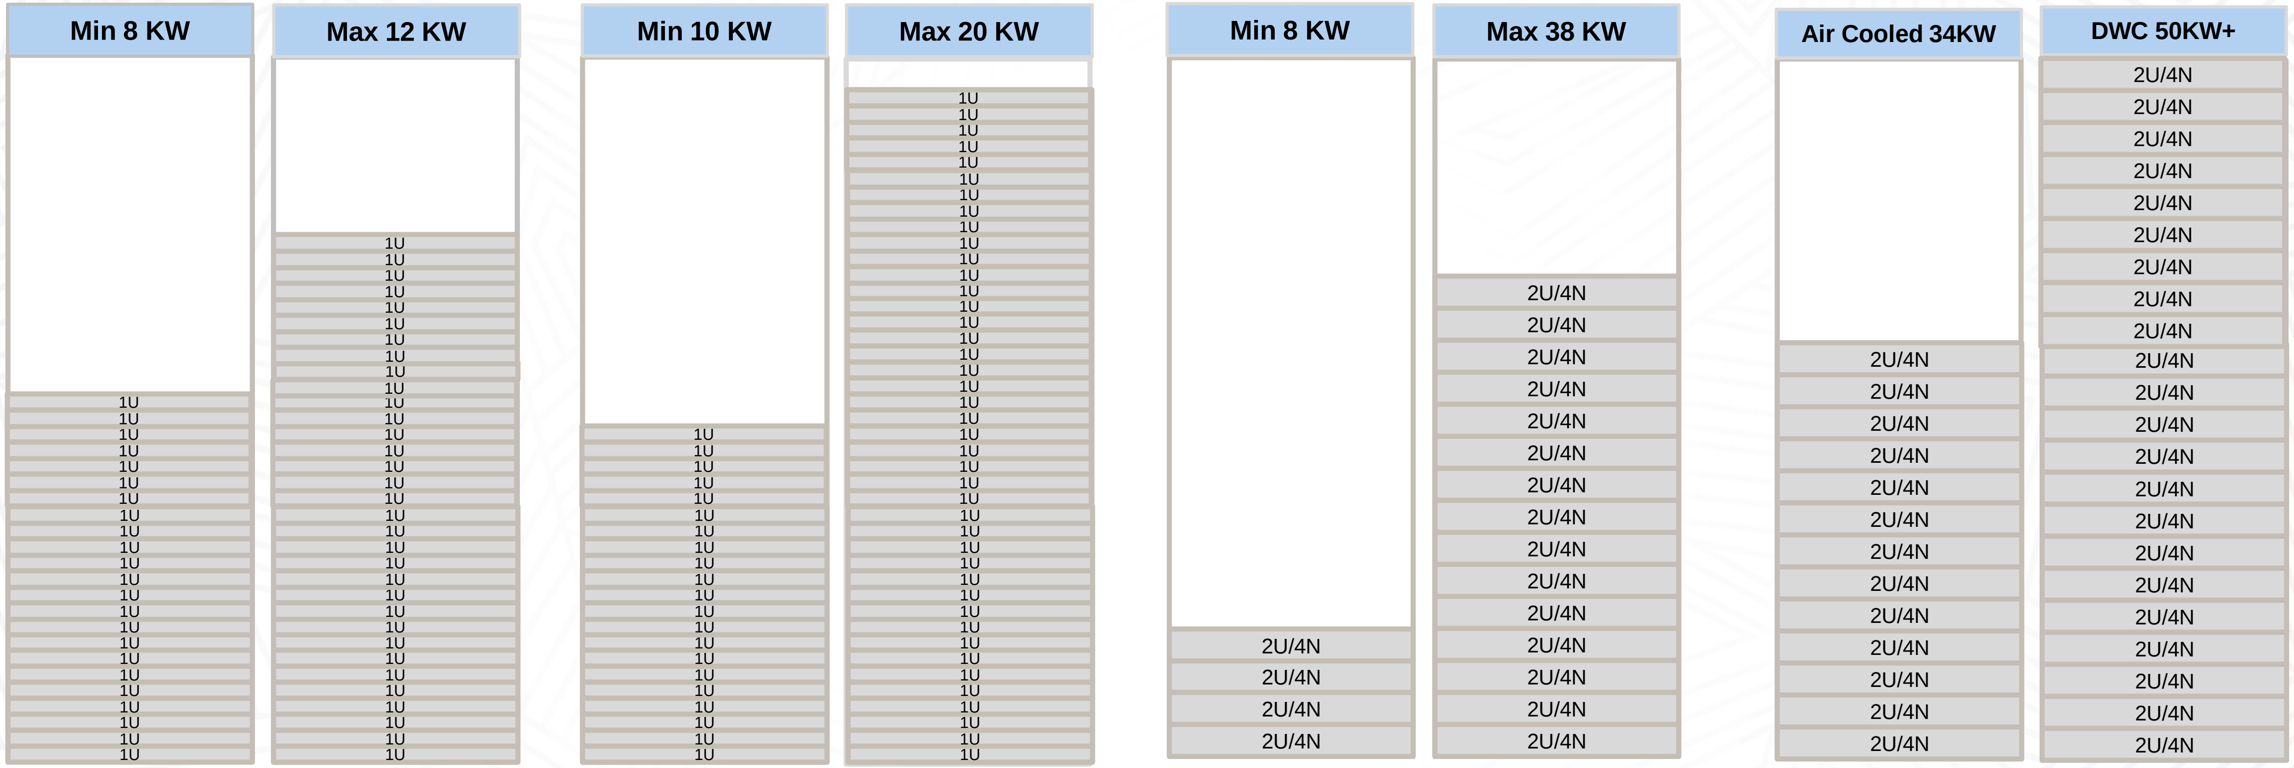

Power per node increasing due to:

- Step in CPU power to maintain Moore's law (Xeon → 235W, Xeon Phi → 400W) and increased competition (AMD Naples@180W, Nvidia GPU@300W)
- Increase in memory count (32 DIMMs per 2S) and adoption of NVMe for Storage and Memory

2018 Lenovo. All Rights Reserved

Cooling comparison

Air Cooled

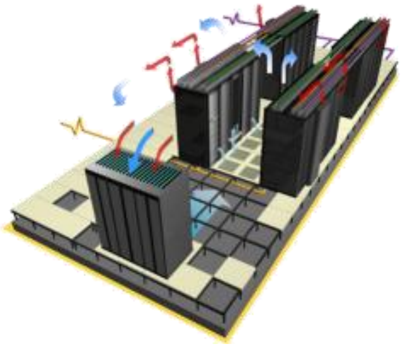

- Standard air flow with internal fans
- Fits in any datacenter
- Maximum flexibility
- Broadest choice of configurable options supported
- Supports Native Expansion nodes (Storage NeX, PCI NeX)

PUE ~2 – 1.5

ERE ~2 – 1.5

Choose for broadest choice of customizable options

Air Cooled with Rear Door Heat Exchangers

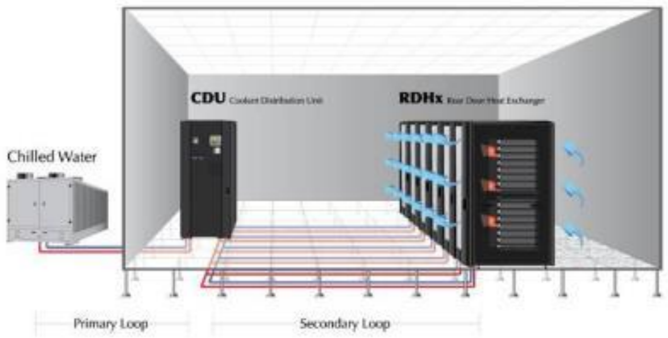

- Air cool, supplemented with RDHX door on rack
- Uses chilled water with economizer (18°C water)
- Enables extremely tight rack placement

PUE ~1.4 – 1.2

ERE ~ 1.4 – 1.2

Choose for balance between configuration flexibility and energy efficiency

Direct Water Cooled

- Direct water cooling with no internal fans
- Higher performance per watt
- Free cooling (45°C water)
- **Energy re-use**
- Densest footprint
- Ideal for geos with high electricity costs and new data centers
- Supports highest wattage processors

PUE ~ 1.1

ERE < < 1 with hot water

Choose for highest performance and energy efficiency

Lenovo

2018 Lenovo. All Rights Reserved

PUE

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

- **Power usage effectiveness** (**PUE**) is a measure of how efficiently a computer data center uses its power;
- Ideal value is 1.0
- **Does not** take into account how IT power can be optimised

ITUE

$$\text{ITUE} = \frac{(\text{IT power + VR + PSU + Fan})}{\text{IT Power}}$$

- **IT power effectiveness** (ITUE) measures how the node power can be optimised
- Ideal value is 1.0

ERE

$$\text{ERE} = \frac{\text{Total Facility Power} - \text{Treuse}}{\text{IT Equipment Power}}$$

- **Energy Reuse Effectiveness** measures how efficient a data center reuses the power dissipated by the computer
- ERE is the ratio of total amount of power used by a computer facility to the power delivered to computing equipment.
- An ideal ERE is 0.0. If no reuse, ERE = PUE

2018 Lenovo. All Rights Reserved

Value of Water Cooling Technology from Lenovo

- Lower processor power consumption (~ 6%)
- Higher TDP processor over air-cooled
- Higher density
- No fan per node (~ 4%)
- Constant Turbo Mode without power penalty (~7%)
- With DWC at 45°C, we assume free cooling all year long (~ 20%)
  - 90 % of heat goes to hot water leading to free cooling remaining goes to cold water

Total savings = ~35-40%

Lenovo

2018 Lenovo. All Rights Reserved

TCO: return on investment for DWC vs RDHx (*)

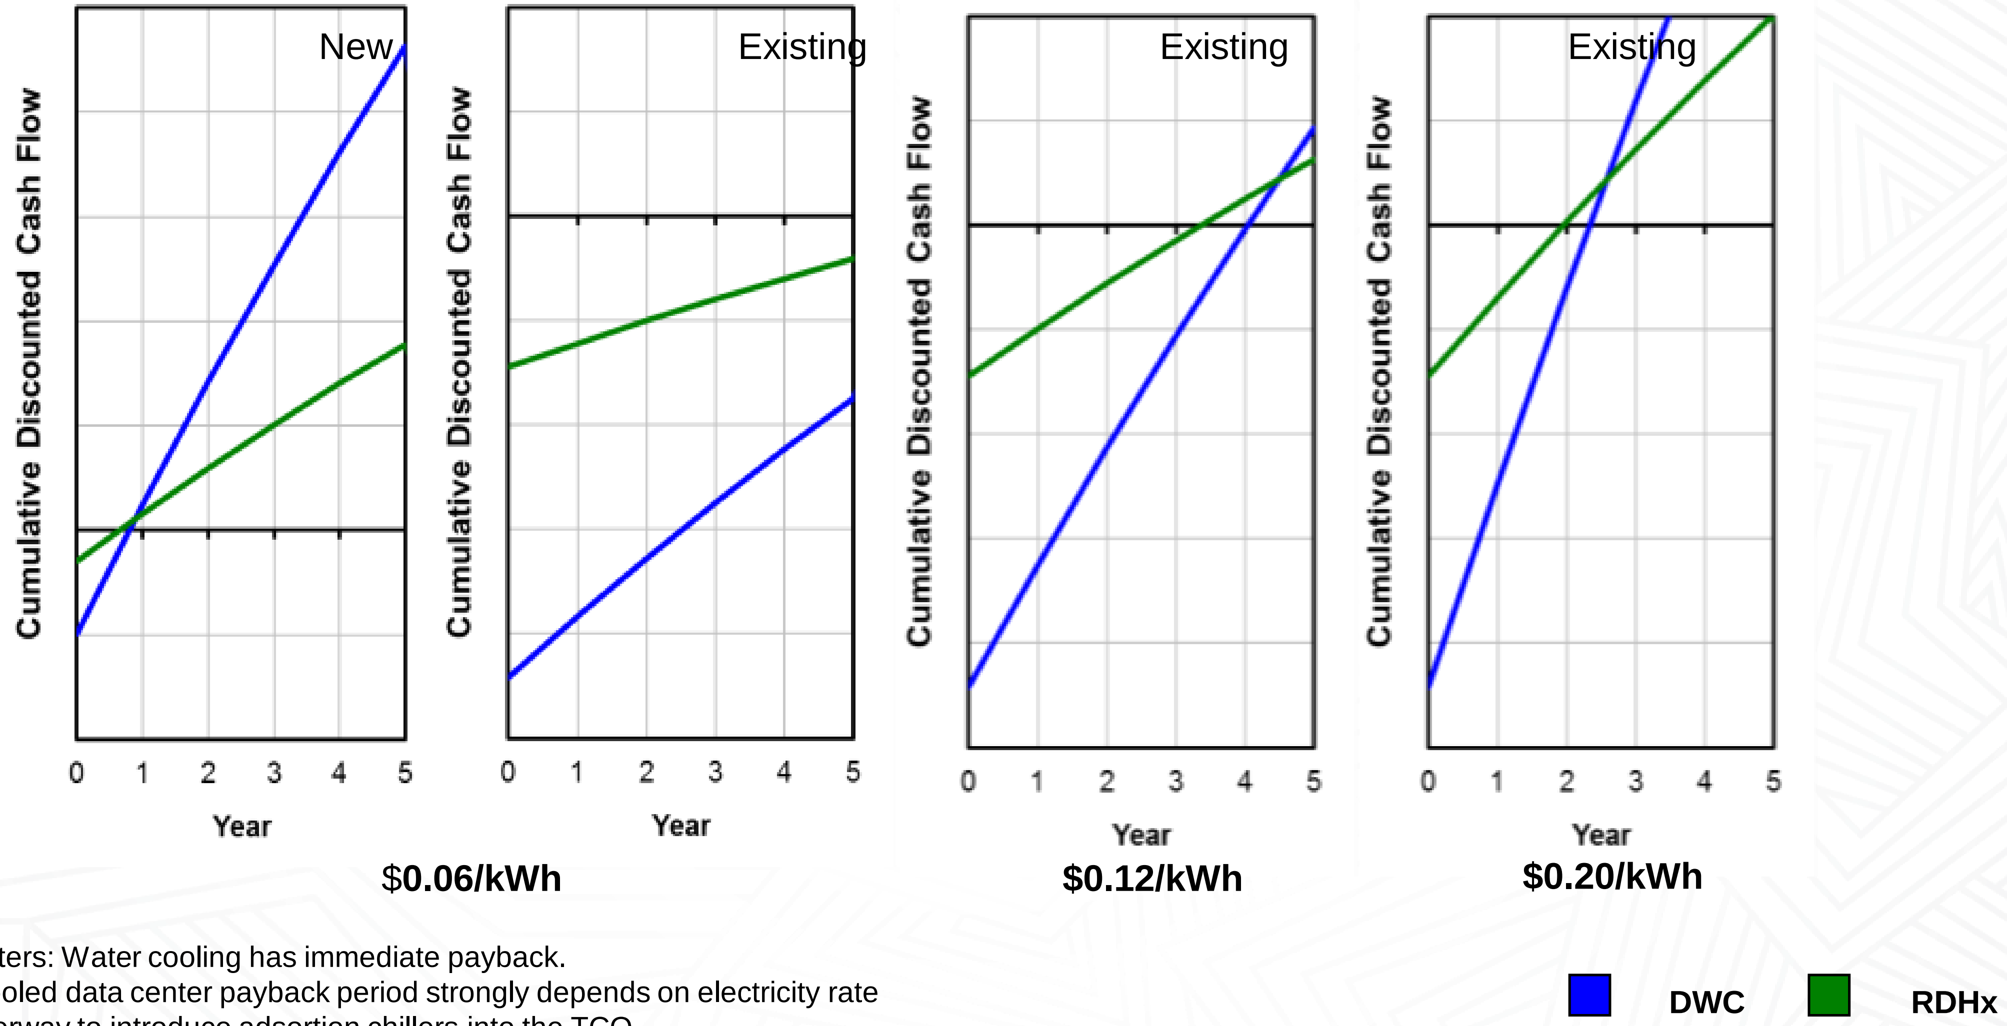

- New data centers: Water cooling has immediate payback.
- Existing air-cooled data center payback period strongly depends on electricity rate
- (*) : work underway to introduce adsortion chillers into the TCO

2018 Lenovo. All Rights Reserved

ThinkSystem Dense Optimized portfolio

Ready to adapt when you are, an ultra-dense hyperscale system for customers seeking the power and scalability to drive large, complex environments such as HPC.

More in less

- Innovative chassis enables greater density for Hyperconverged workloads; Designed for dense HPC architectures; Future proof with 3D XPoint

Ready to adapt

- Widest range of processors in a dense form factor; Max storage with 48TB of capacity; Stackable node design supports GPUs and specialized IO adapters

Modularity to transform

- Disaggregated IO design allows for multiple fabrics; Scalable management design simplifies infrastructure costs; Front and rear access for easy serviceability

ThinkSystem SD530/D2 Enclosure

Lenovo NeXtScale nx360 M5 WCT

Analytics	Big Data	High Performance Computing	Modeling & Simulation	Scientific

2018 Lenovo. All Rights Reserved

Lenovo

Summary

- Per Server power requirements are trending upwards making water cooling necessary in the future.
- Direct Water Cooling Technology from Lenovo can greatly reduced the overall OPEX cost burden
- Reduction in OPEX is dependent on the cost of energy per datacenter
  - Easier to recover in new datacenters that are building out with water cooling
  - Fluctuates greatly in existing hybrid datacenters
- Processing power in water-cooling servers is greater than air due to heat-transfer efficiency
  - Can run 100% of time in turbo mode

GROWTH

SOLUTIONS

SCALE

INNOVATION

APPLICATION

AWARENESS

2018 Lenovo. All Rights Reserved

thanks.