Smarter technology for all

Lenovo Neptune® Technologies and Liquid cooling Solution Update

Michael Chang | March 2023
Sr. Manager HPC & AI Enterprise Solution

Industry Trend & Cooling Challenges

Industry Trend – Chipset Power Increasing

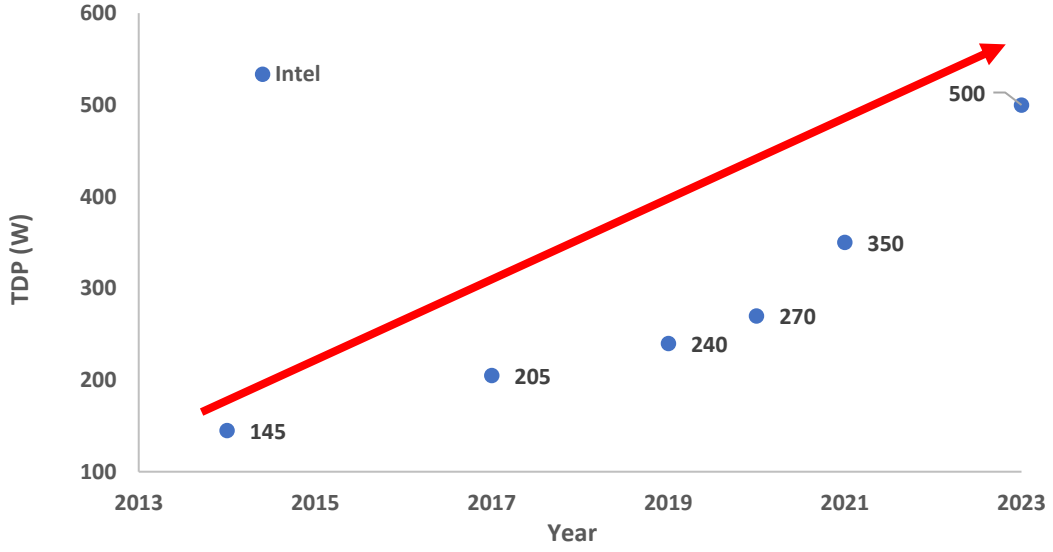

CPU

- Intel BirchStream AP CPU : 500W
- Intel BirchStream SP CPU : 350W

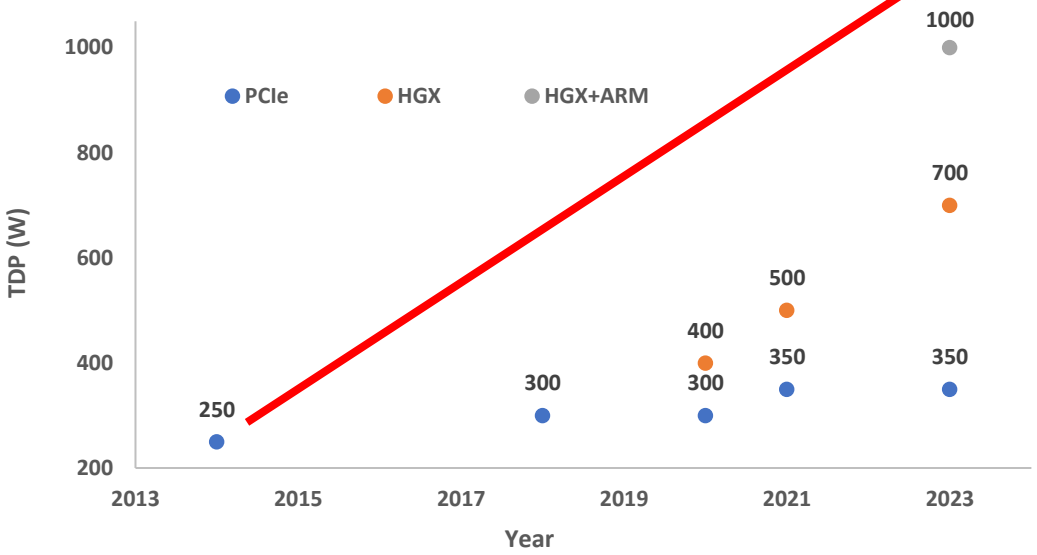

GPU

- Nvidia PCIe GPU card : 350W
- Nvidia HGX GPU : 700W
- Nvidia HGX GPU + ARM CPU: 1000W
- Intel GPU Max<=600W TDP

Component power increase rapidly.
Air cooling can not meet all requirement.

Traditional Cooling Approaches are Reaching Critical Limits

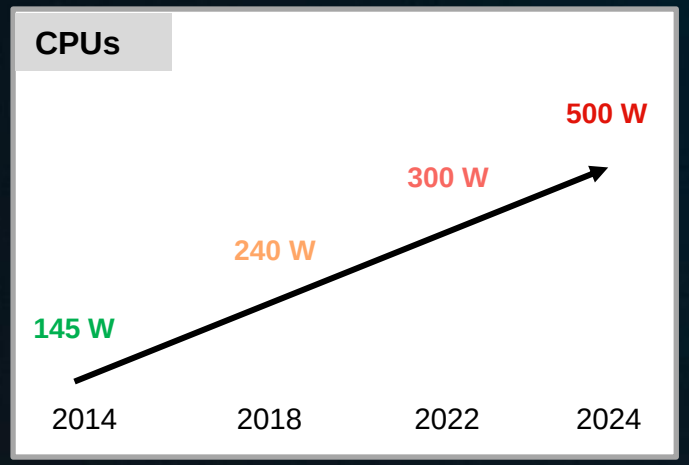

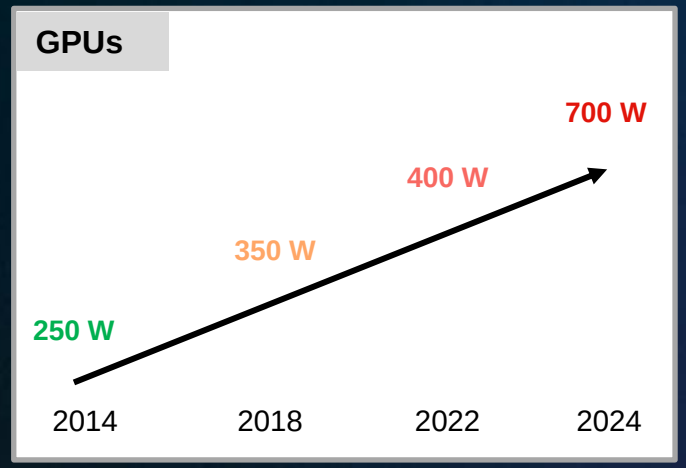

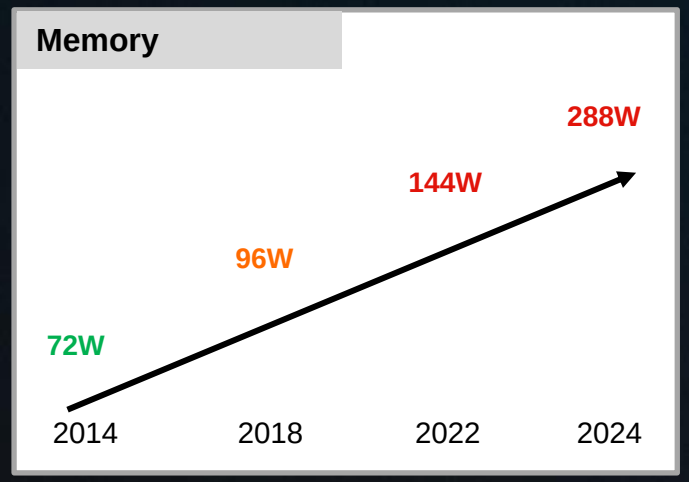

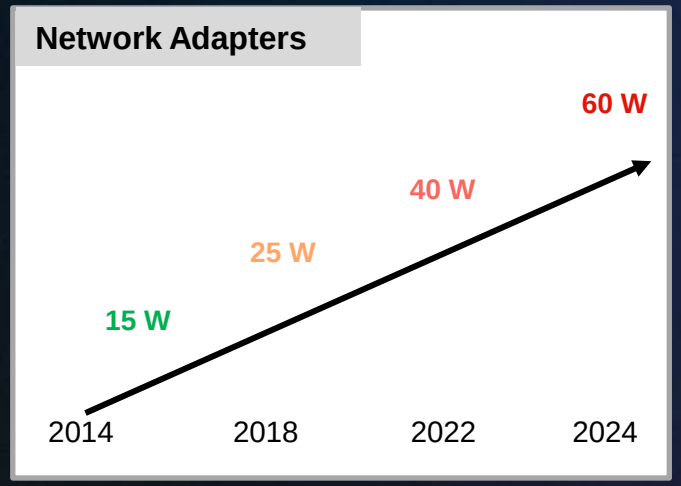

T-Case Temperatures Dropping

- The maximum temperature allowed at the processor Integrated Heat Spreader (IHS)
- The lower the temperature, the more challenges it is to maintain
- High power + low Tcase = lots of air to pull heat away quickly before reaching max Temp

Default TDP (W)	Tcase Max (ºC)
120/155	85
170	
180	81
200	
225	81
280	71
trend >300	trend <60

Traditional Cooling Approaches are Reaching Critical Limits

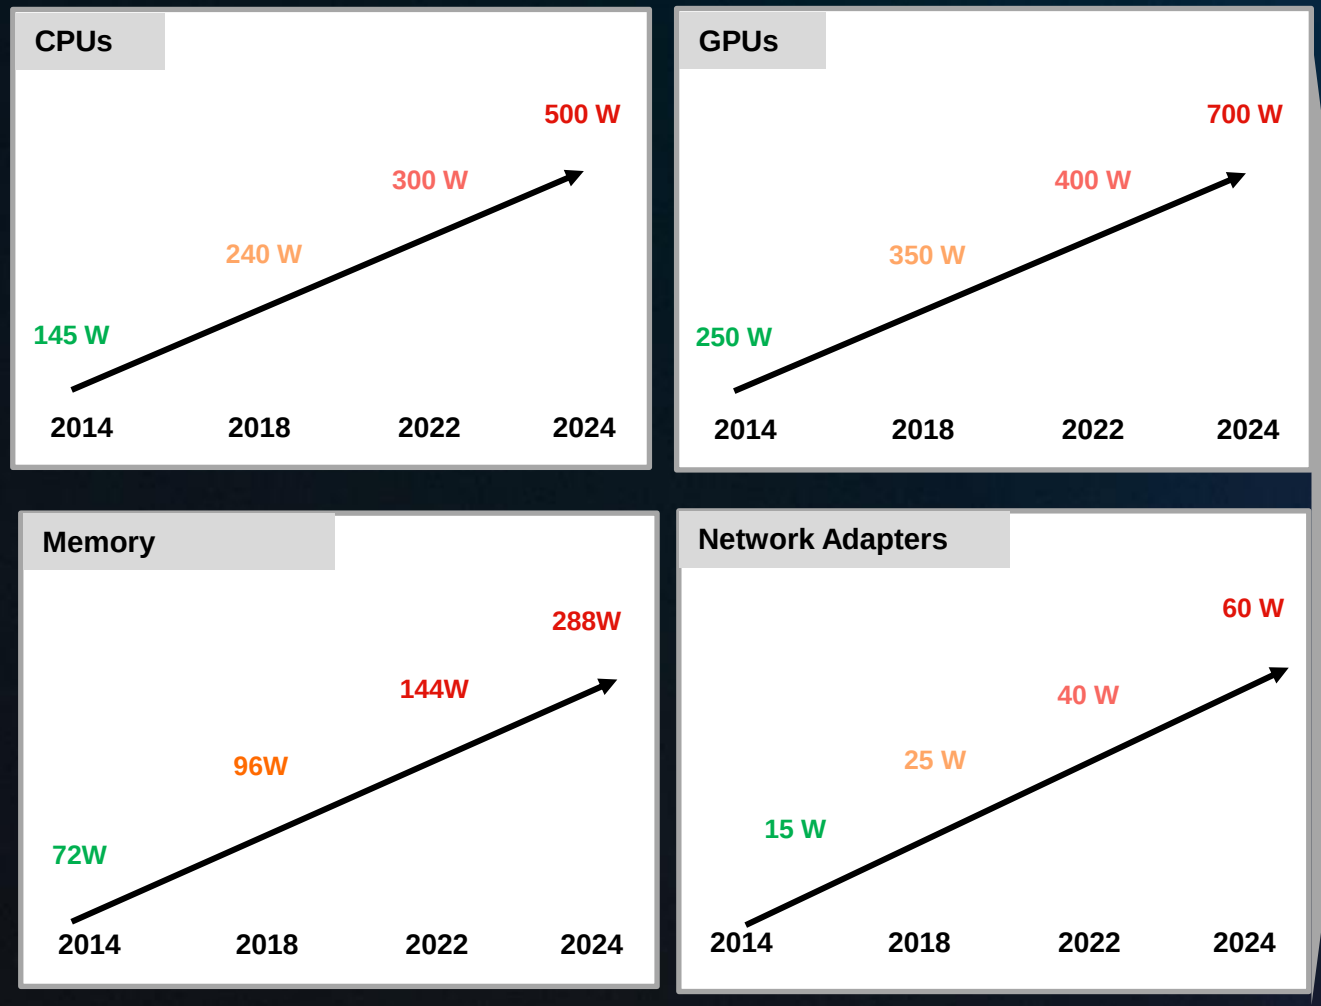

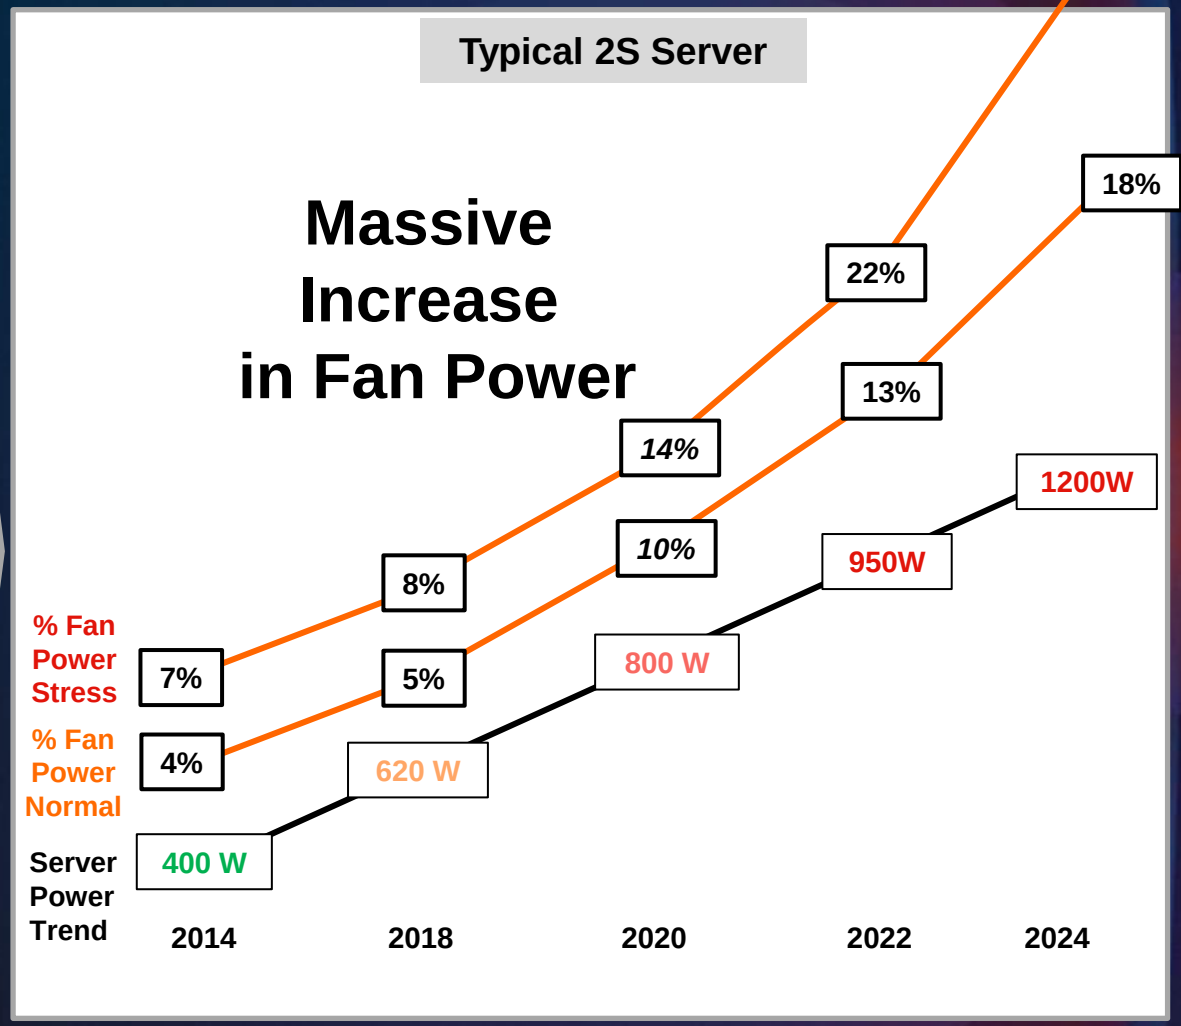

Traditional air cooling cannot keep up – we must rethink cooling

Cooling can consume massive amounts of water

- **+/- 1.8 liters per kWh (or 9,000 liters/yr)**
- **Emerging issue for many locations due to availability or cost**

More power = more data center infrastructure for power delivery and heat rejection

High power inside the server drives up component (Fan/VR/PSU) cost

High power reduces density, utilization, component life and reduces performance

Increasing power drives increasing fan speeds and higher acoustics (sound)

The benefits of reducing power consumption add up quickly

Lenovo Neptune® Technologies – Save Power

Neptune™

RACK WATER COOLING

Rear Door Heat Exchanger
(RDHX)

In-rack Cooling Distr. Unit
(CDU)

DIRECT WATER COOLING

Full System

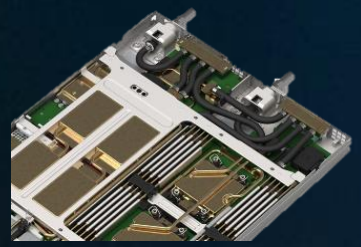

Core System

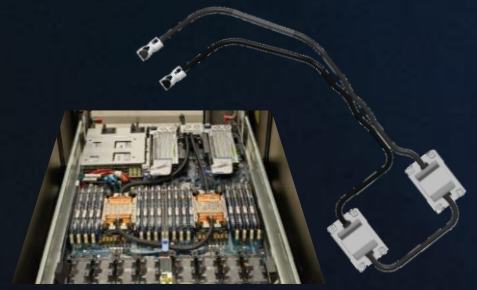

LIQUID ASSISTED COOLING

Thermal Transfer Module
(TTM)

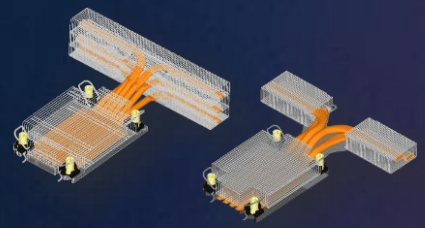

Liquid to Air Module
(L2A)

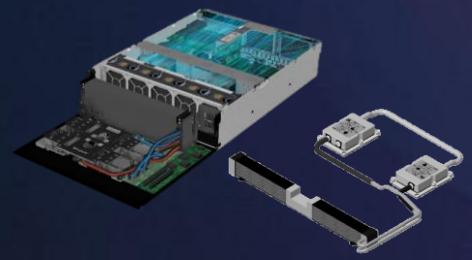

LIQUID IMMERSED COOLING

Self-Contained

Fully Submerged

submer

MIDAS IMMERSION COOLING

HPC&AI 2022/23 Server Portfolio with Intel Eagle stream

Ship Support	2/27 // 12/10	2/27 // 12/10	Mar 10th	May 26th
ThinkSystem	SR650 V3	SR630 V3	SD650 V3	SD650-I V3
				
Form Factor	2U	1U	6U12N	6U6N
CPU Sockets	2S Intel 350W / 2S	2S Intel 350W / 2S	2S Intel 350W	2S Intel 350W
DIMM Slots	32 / 24	32 / 24	16	16
2.5‘ Drive Slots	up to 40	up to 12	up to 4 (NVMe)	up to 2 (NVMe)
Acceleration	3x DW 300W or 6x SW 150W	2x SW 150W or 1x FHHL 150W	n/a	4x 600W OAM (Ponte Vecchio)
Cooling	Air (A4 ; A2 full perf.)	Air (A4 ; A2 full perf.)	Water (A4/W5 full perf.)	Water (A4/W5 full perf.)
Cabling	Rear	Rear	Front	Front
Typical HPC Workload	HPDA, AI Inference, Technical Computing	Technical Computing	Supercomputing	Exascale Supercomputing, AI Training

Blow Out Performance for Air Cooled Servers

The No Plumbing Liquid Cooling Advantage

Neptune™

Lenovo Thermal Transfer Modules (TTM)

1U TTM
For 1U & dense platforms

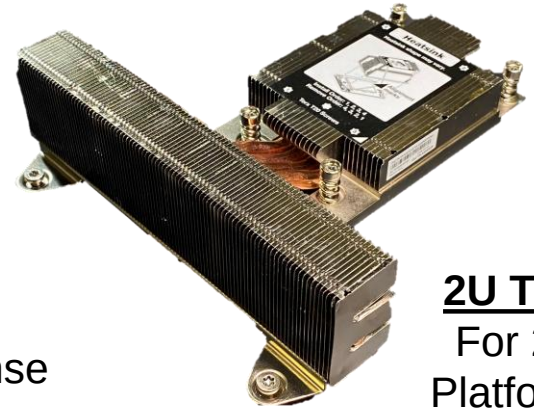

2U TTM
For 2U Platforms

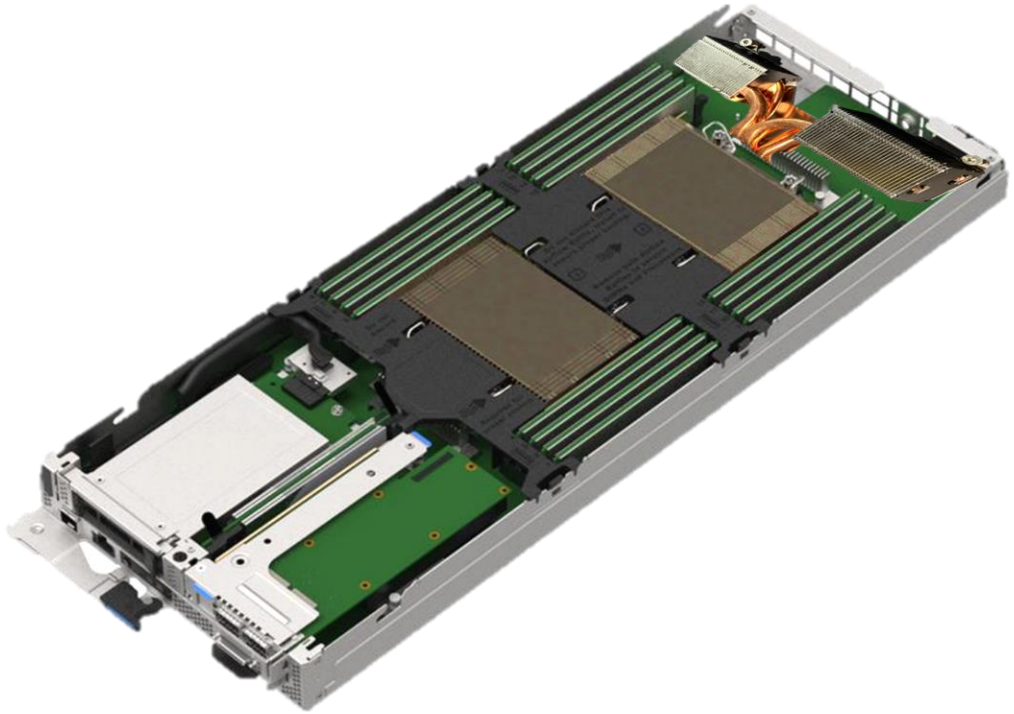

The power of liquid; Convenience of air

Runs quieter
Uses less power
Unlocks higher performance parts

Fits into any data center

The SR670 V2 with nvidia A (H)100 SXM + Lenovo Neptune

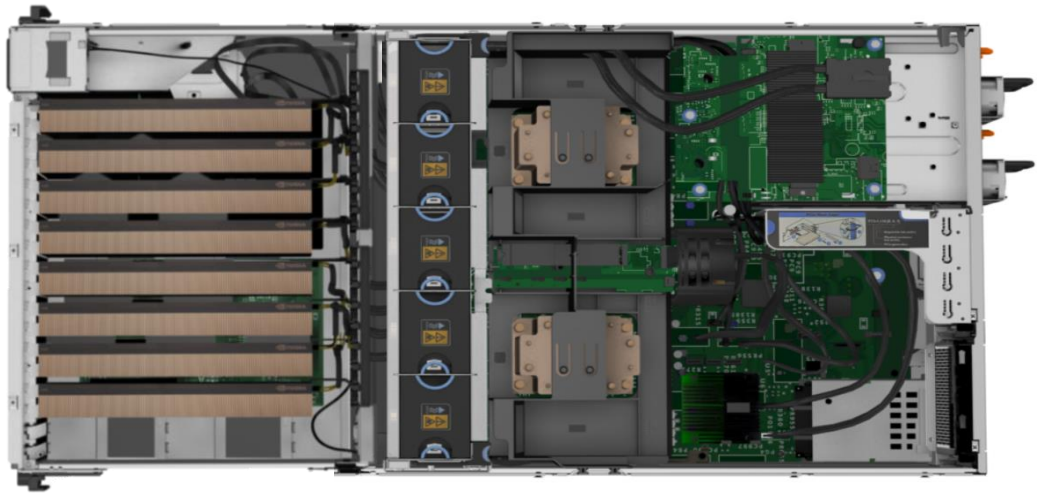

TRADITIONAL AIR COOLING

- Front GPU design, rear CPU elevated temp
- Super high airflow needed to move heat away
- High-speed fans
  - Loud
  - Massive power consumption

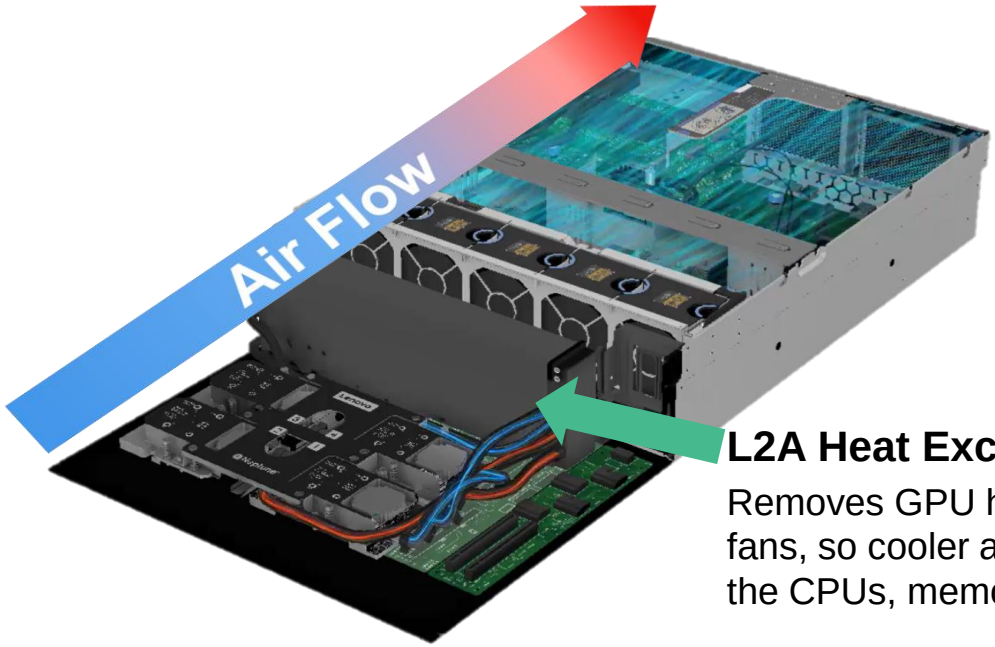

L2A Heat Exchanger

Removes GPU heat before the fans, so cooler air circulates thru the CPUs, memory & storage

LENOVO NEPTUNE L2A COOLING

- Liquid to Air Heat Exchanger (L2A)
- Fully sealed, no plumbing, no maintenance
- Circulates liquid efficiently over the GPUs
- Reduced fan speeds, lower power and quieter

Lenovo Neptune™

From Exascale to Everyscale™

SD650 V3
2 x Intel Xeon

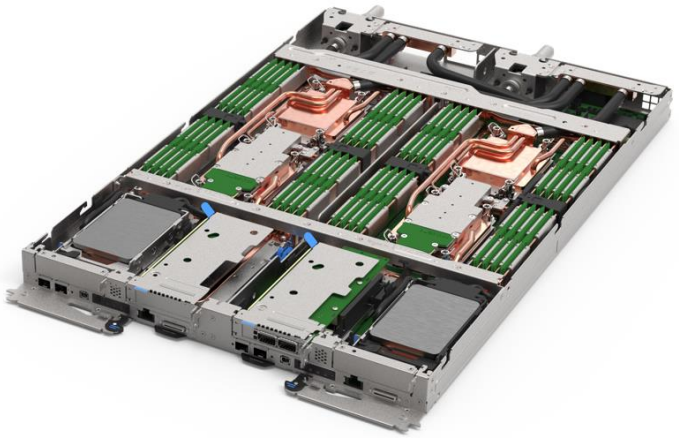

SD650-I V3
2 x Intel Xeon
+ 4 x Intel GPU

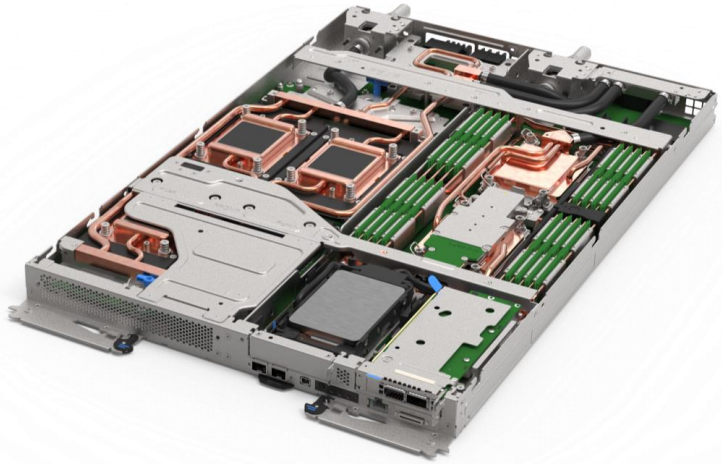

- Complete removal of all fans reduces moving components and drives down power consumption
- Liquid carries over 97% of all heat from the data center
- Allows use of unchilled/hot water for cooling (Over 110F inlet)
- High temperature effluent allows for possibility of recycling of heat
- Uses conditioned water for cooling – easy, safe, stable

#1 PROVIDER OF SUPERCOMPUTERS IN THE WORLD

NEARLY 1 IN 3 RUN ON LENOVO

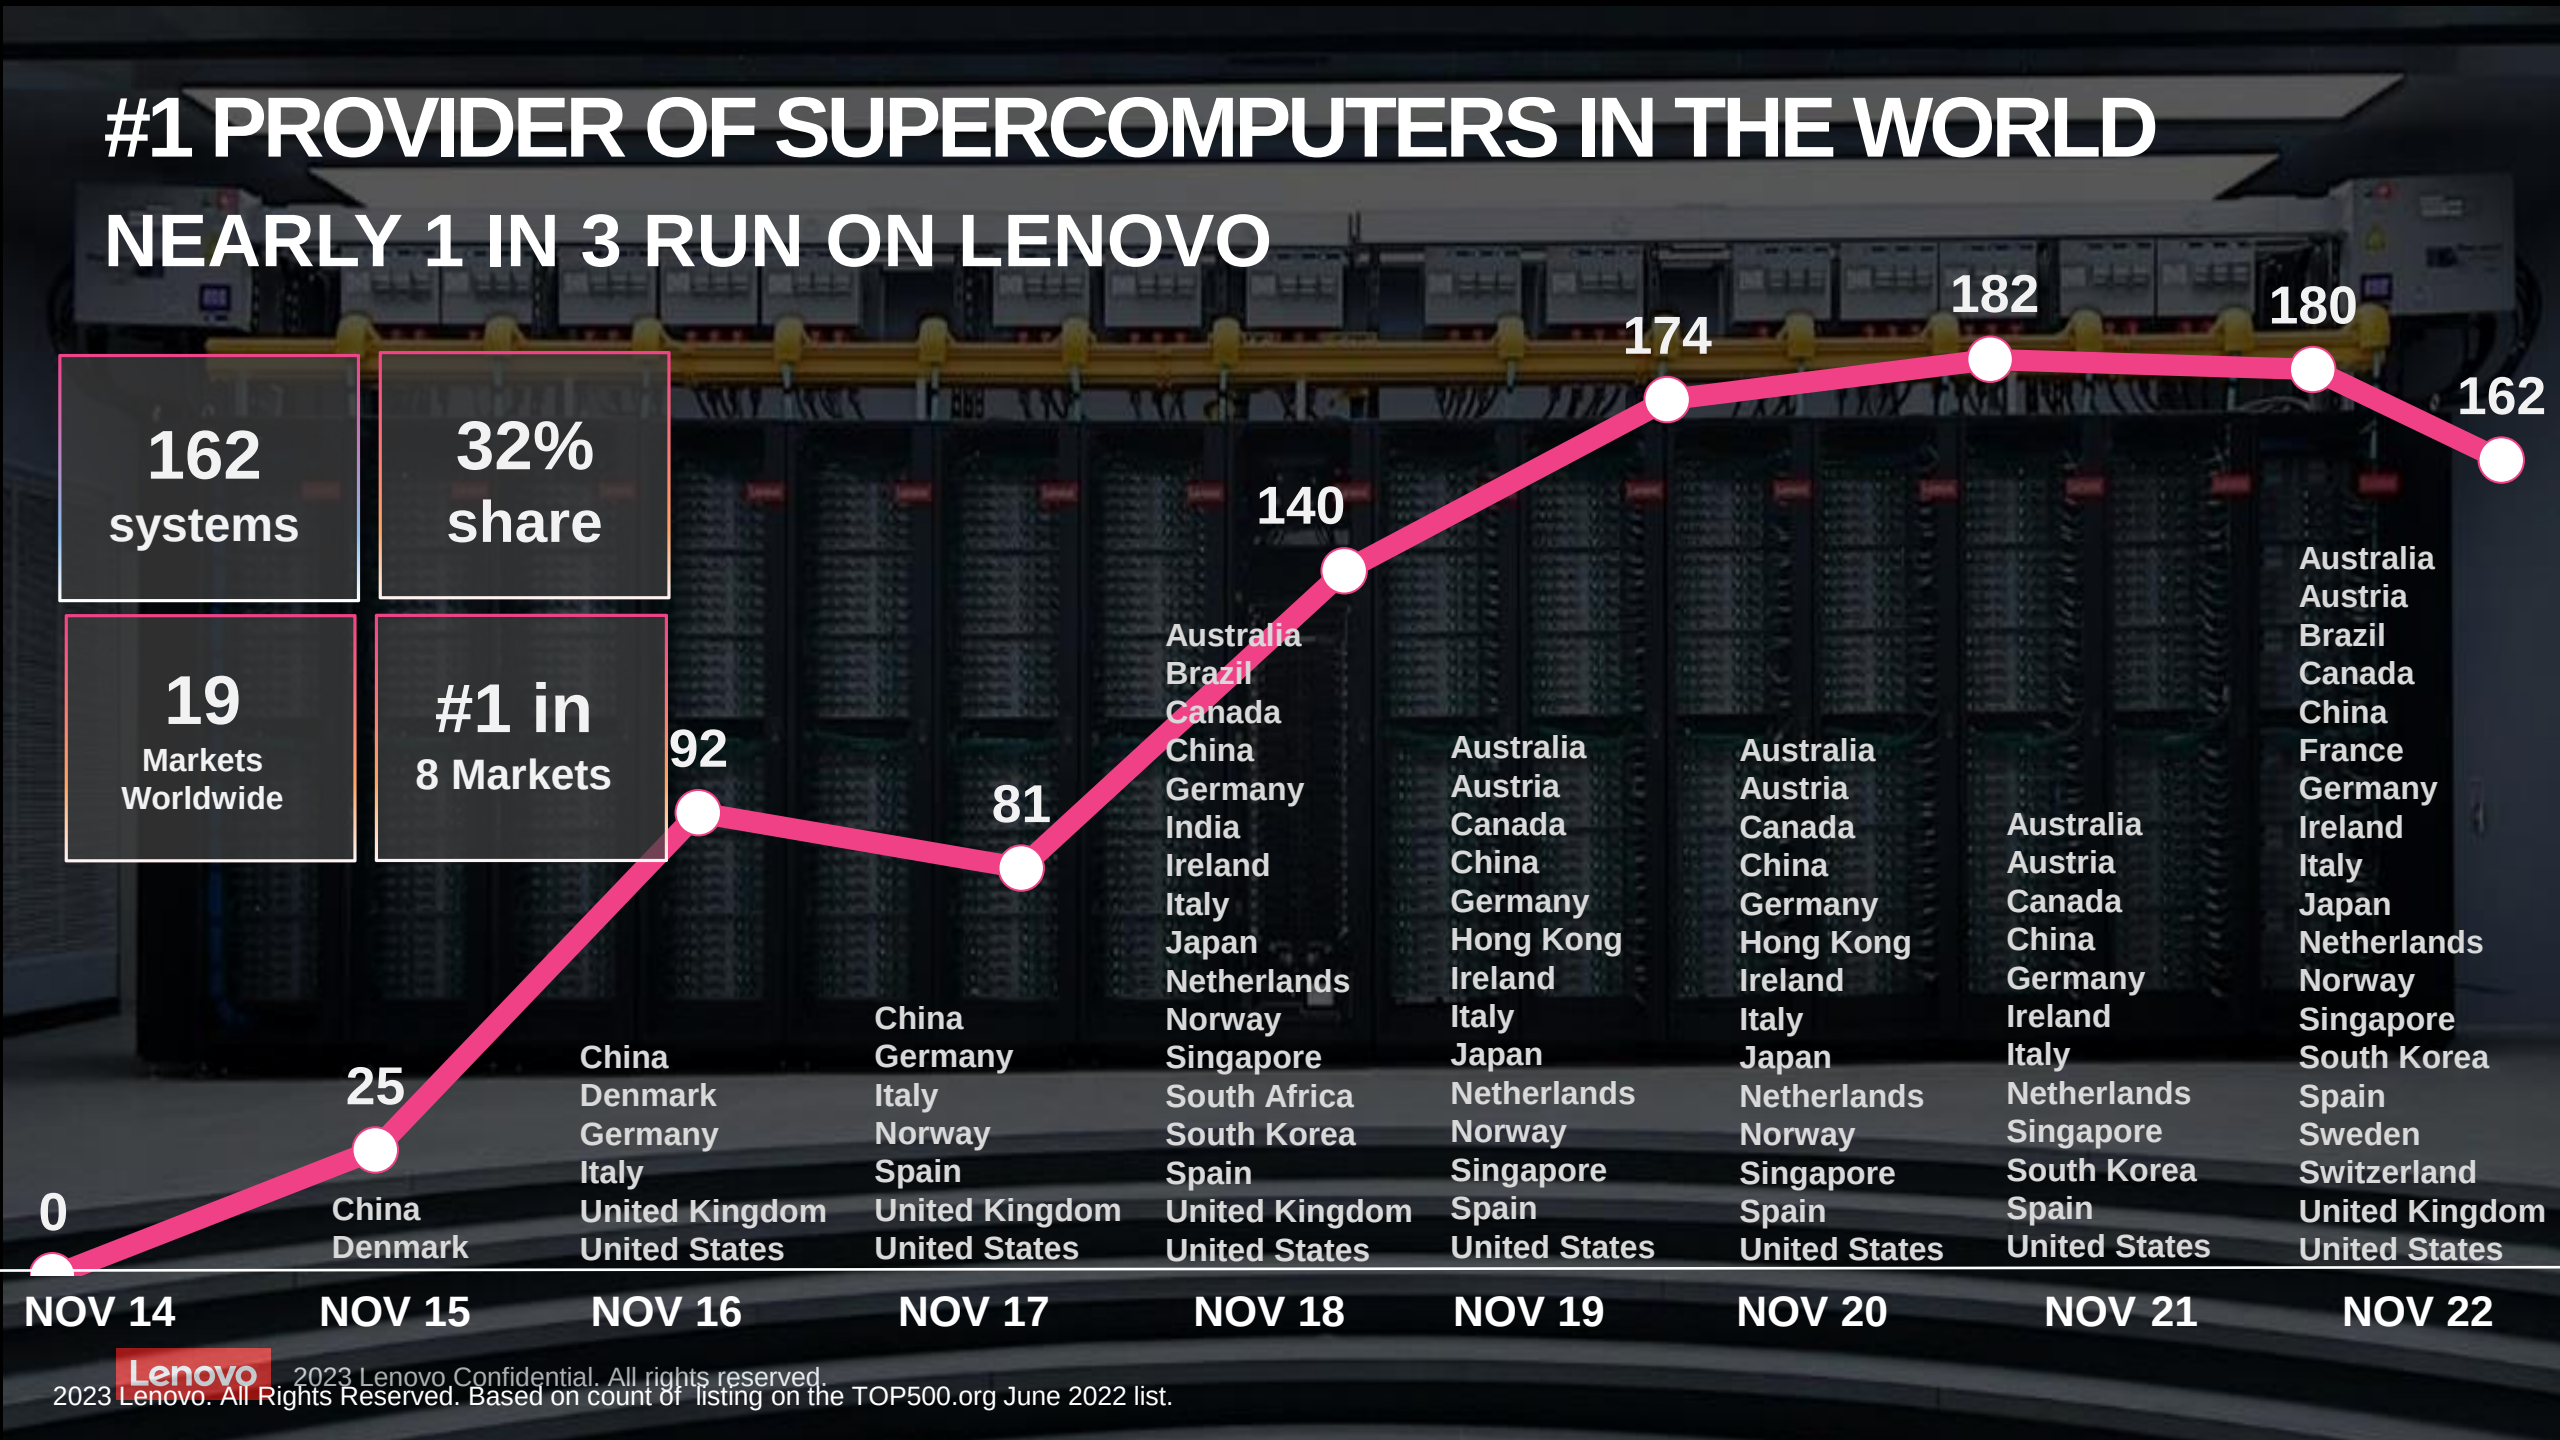

Lenovo DWC Solution

Data Center Cooling Type

1

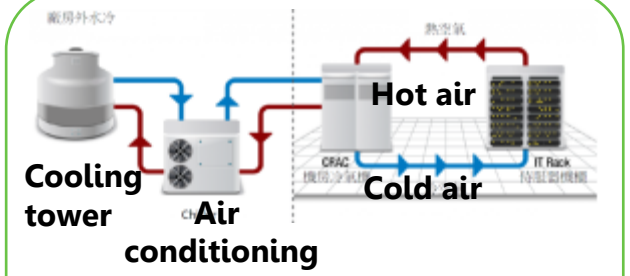

Air Cooling

- With server internal fans
- Rely on air conditioning
- Fit to any datacenter
- Max flexibility
- Max support config

100% heat to Air
PUE >1.5

2

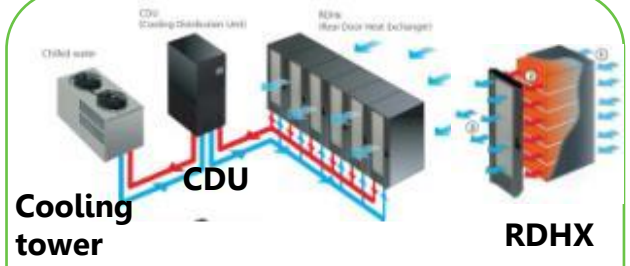

Air Cooling + Rear door heat exchanger

- With server internal fans
- Air cooling with RHDX to enhance cooling capability
- With Chilled water (18'C)
- Enable tight rack placement

100% heat to Air
PUE > 1.2

3

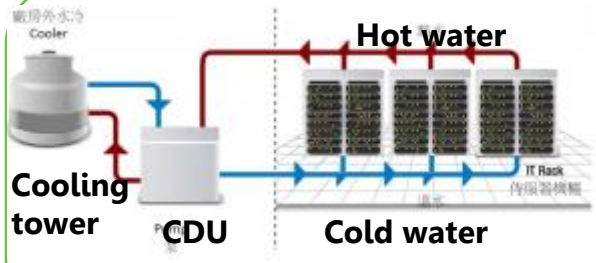

Direct Water Cooling

- No server internal fans
- Higher performance/per watt
- With warm water (45'C)
- >90% heat remove rate by liquid
- Support high power components

>90% heat to Liquid
PUE < 1.1

4

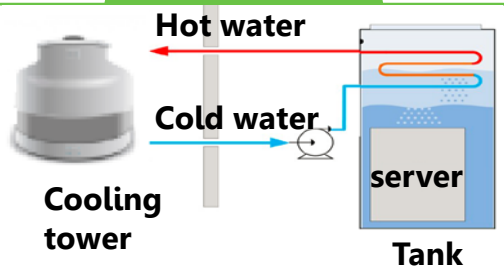

Immersion Cooling

- No server internal fans
- With 3M Novec liquid
- 100% heat-remove rate by liquid

100% heat to Liquid
PUE < 1.1

Data Center Cooling Type Comparison

* Lower is Better
** Higher is Better

Cooling type	Energy Usage PUE *	Rack Cooling Capacity **	Rack Floor Space *	Service Effort *	Reliability **	Economic (Cost)			Environment Impact		
						CAPEX *	OPEX *	Maintenance *	Heat Recycle **	Noise *	Pollution *
Air Cooling 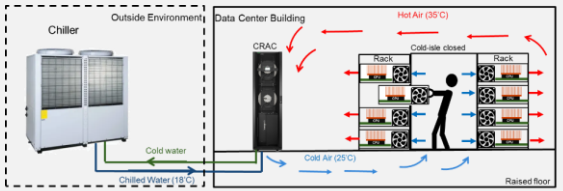	**>1.5**	**< 20kW**	**Typical**	**Low**	**High**	**Low**	**High**	**Low**	**N/a**	**High**	**Low**
Air Cooling + Rear Door Heat Exchanger 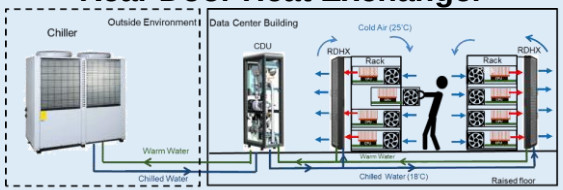	**>1.2**	**< 30kW**	**Typical**	**Low**	**High**	**Medium**	**Medium**	**Low**	**Low**	**High**	**Low**
Direct Water Cooling 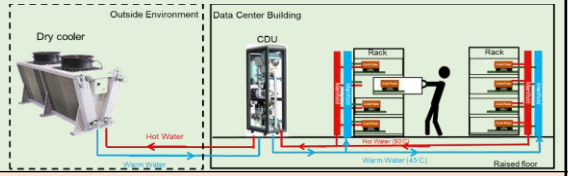	**<1.1**	**< 100kW**	**Typical**	**Medium**	**High**	**High**	**Low**	**Medium**	**High**	**Low**	**Low**
Immersion Cooling 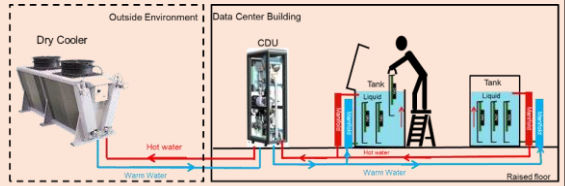	**<1.1**	**< 80kW**	**High**	**High**	**Low**	**Highest**	**Low**	**High**	**Medium**	**Low**	**Medium**

Lenovo Direct Water Cooling Product

- > 10 years experience
- Maximum heat remove by liquid
- Support various components
  - Intel CPU
  - Intel / Nvidia GPU

Liquid Cooling PSU

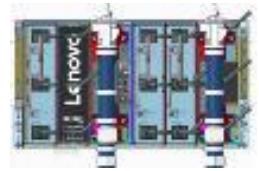

SD650 V2
2S Intel Whitley

SD650 V3
2S Intel EGS

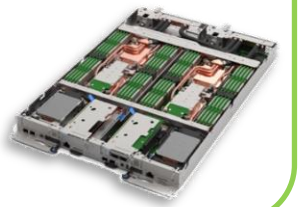

SD650-I V3
2S Intel EGS
4x Intel **Ponte Vecchio**

X3550 M3
1U2S blade

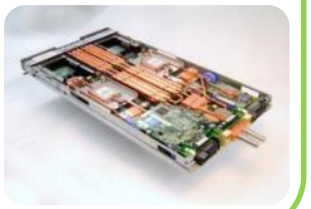

dx340M4
2S blade

Nx360 M5
2S Intel

SD650
2S Intel Purley

SD650-N V2
2S Intel Whitley
4x **Nvidia A100**

2009 **2012** **2015** **2017** **2020** **2022**

Lenovo ThinkSystem SD650 V3 / SD650-I V3

Feature	ThinkSystem SD650 V3 / SD650-I V3
Form factor	• 1U full wide double (CPU only) / single (GPU acc.) node tray in 6U6T Chassis (DW612 / DW612S) for 19inch rack cabinets
Processors	• NPI: 2x Intel “Sapphire Rapids” up to 350W TDP • R1: 2x Intel “Sapphire Rapids”/ “CPU Max” (HBM) up to 350W TDP • R2: 2x Intel “Emeral Rapids” up to 350W TDP
Memory	• 16x DDR5 4800 R/3DS DIMM ECC • Capacities: 16GB/32GB/64GB/128GB
Storage	• Up to 4x 7mm or 2x 15mm U.2/SATA/U.3/E3.S 1T (future proof) no drive/bp choice • 1x M.2 NVMe – liquid cooled, backplane-less • SW RAID only with Intel VROC
NIC	• 2x SFP28 25GbE LOM, NCSI (10Gb capable) • 1x RJ45 1GbE LOM, NCSI
PCIe	• 2x x16 PCIe Gen5 LP (each in place of half the storage) w/ CXL • Internal: 4x x16 Gen5 for Side Car expansion (all PCIe balanced)
Acceleration	• 4x Intel GPU Max<=600W TDP
Front Access	• All IO in front • Power LED Button, ID and System Health LED • KVM breakout connector, Pong
Rear Access	• 2x RJ45 on SMM for XCC/Daisy Chain • USB 2.0 dedicated to SMM • LED for ID, Error, Power and Heartbeat
Mgmt/TPM	• xClarity Controller (XCC), TPM 2.0
Power	• Up to 9x HS Air CFF v4 (2400W PT, 2600W TT) / Up to 3x HS DWC Power Supply (7200W) 80+ Titanium • N+1 redundancy (only air-cooled / without acceleration on DWC)
Cooling	• Up to 50°C warm water for component level cooling • Up to 100% cooling efficiency at 45°C inlet temperature* *depending on configuration/water inlet temperature

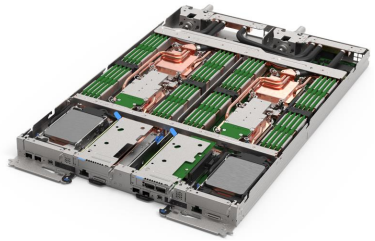

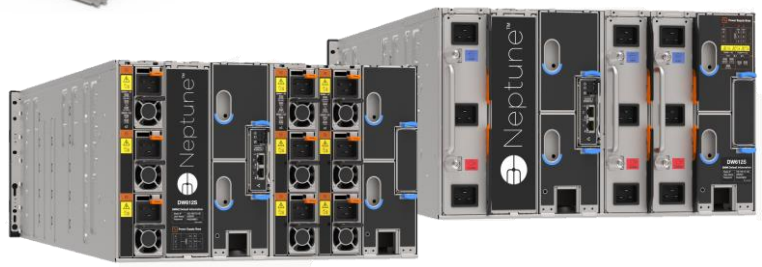

blue are changes over SD650 V2

Lenovo is the strategic partner of Intel on PVC x4 Subsystem

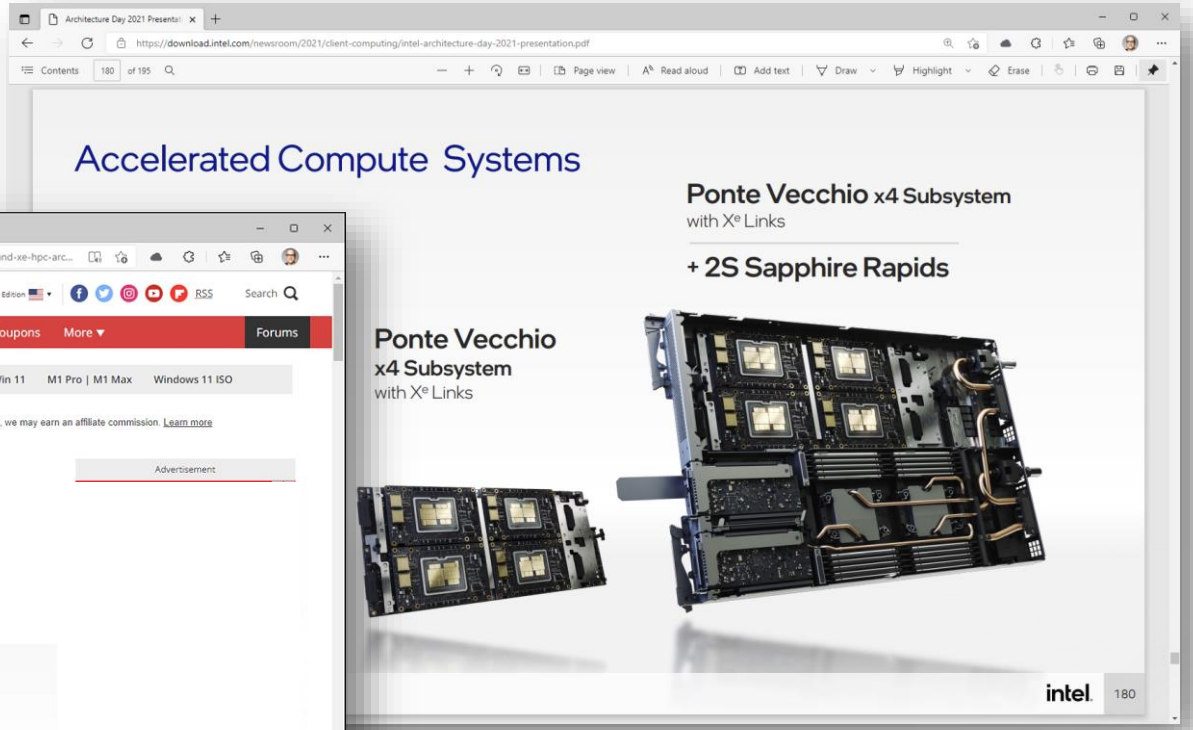

Board and system are designed to match perfectly.

The actual system shown at ISC'22 in the Lenovo NDA booth

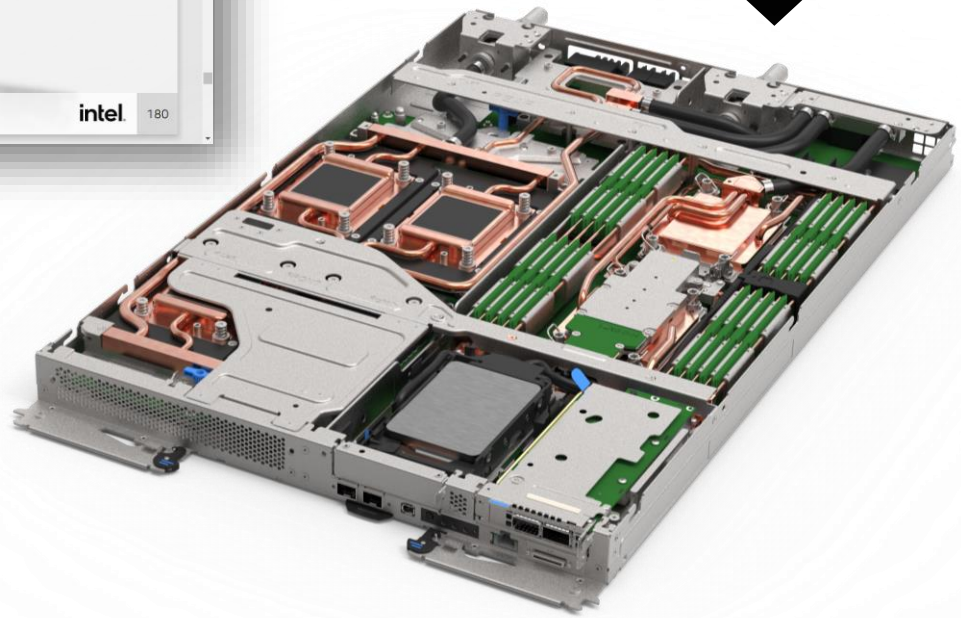

Lenovo is the contract developer for the base board.

Direct Water Cooling

- Cooling method.
  - Node level Cold plate contact with system components to transfer heat to CDU.
  - CDU exchange heat of primary side and secondary side.
  - Heat dissipate to outside environment through cooling tower in Primary side.

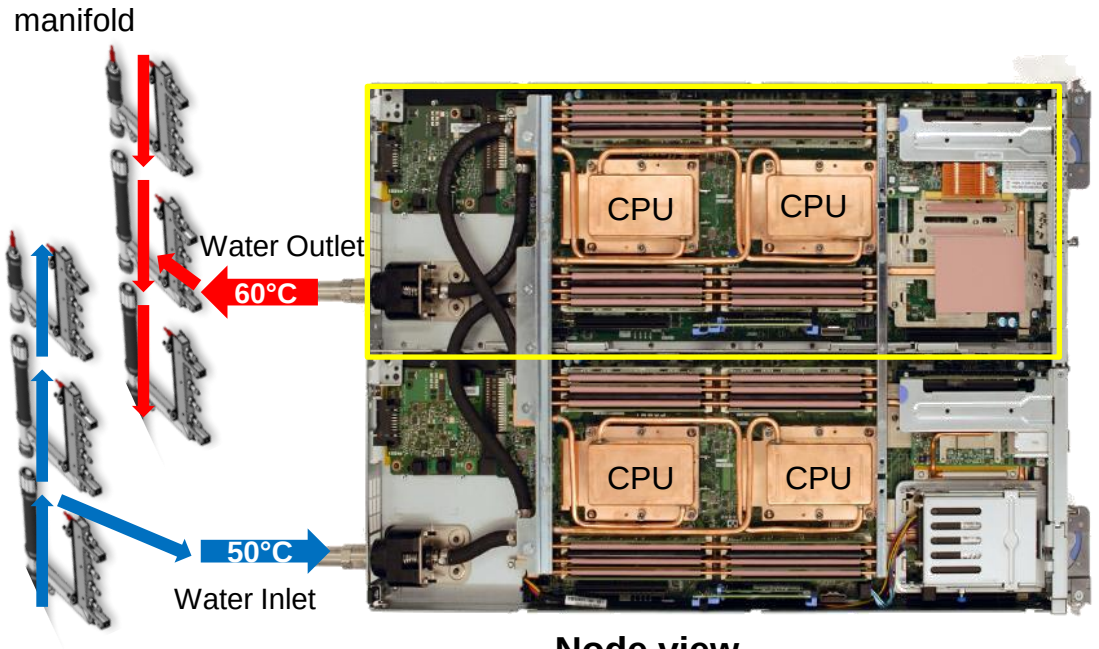

Node view

Manifold

Rack level

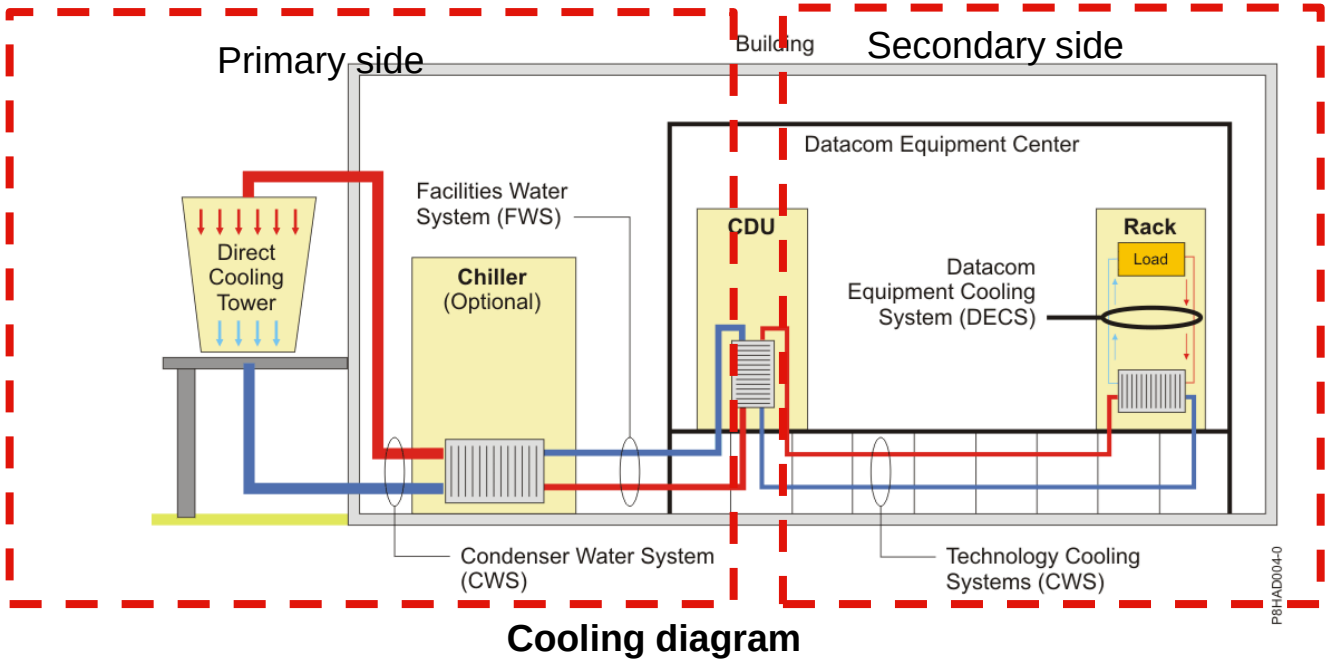

Cooling diagram

Lenovo ThinkSystem SD650 V3

72x Nodes Per Rack

6 x DW612 Chassis per rack

SD650 V3 Node

- 2 nodes per tray
- 2 CPUs + 2CPUs per tray

12x Nodes Per Chassis

6 x dual-node trays

DW612 Chassis

Lenovo Direct Water Cooling Solution Rack View

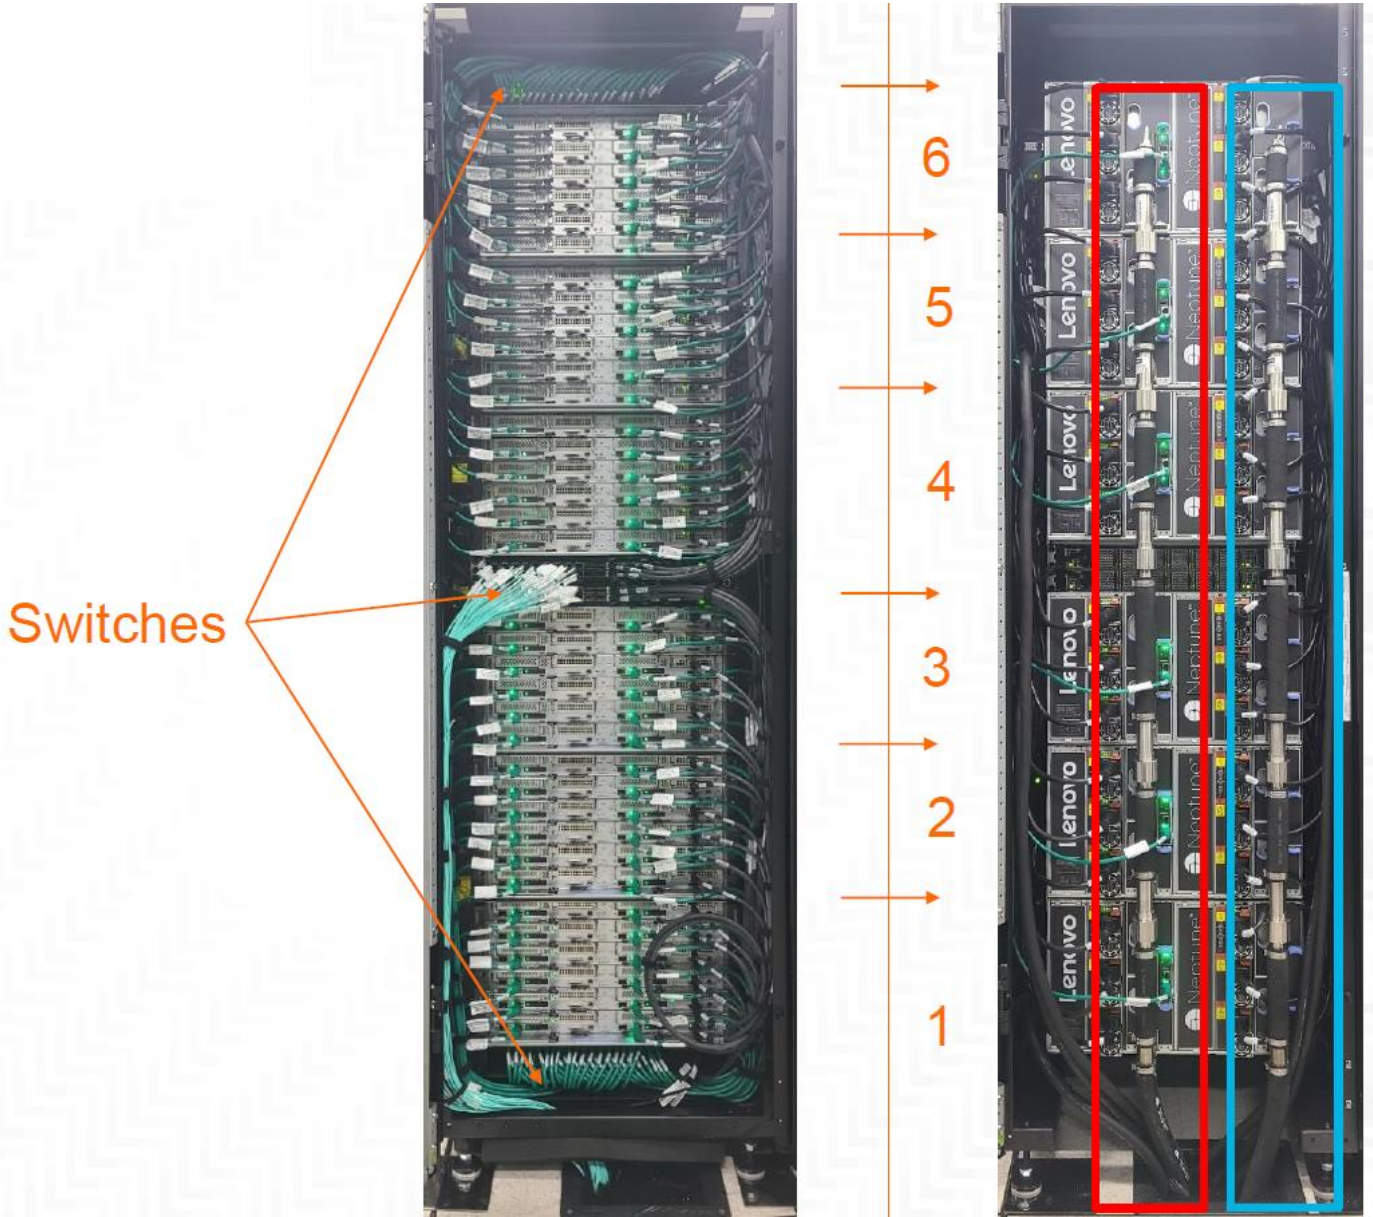

ESP Wins Rolling in Globally

Global Success

TUDA

LRZ
ZIB

BSC

Israel
PMO

Geely Auto

BP

SMN Argentina

NCI

Neptune™

Liquid cooling technology

Beyond carbon neutral to carbon negative

GOAL: Deliver computing that is carbon negative operationally
Requires energy efficiency + green power + energy re-use

No/Low Emission Power:
Now more affordable than ever before

Direct Liquid Cooling:
5th Generation ThinkSystem SD650 and SD650-N

Liquid to Liquid Exchanger:
Transfer heat from IT to facilities

Facility Reuse:
Hot Water, Heat, Pool

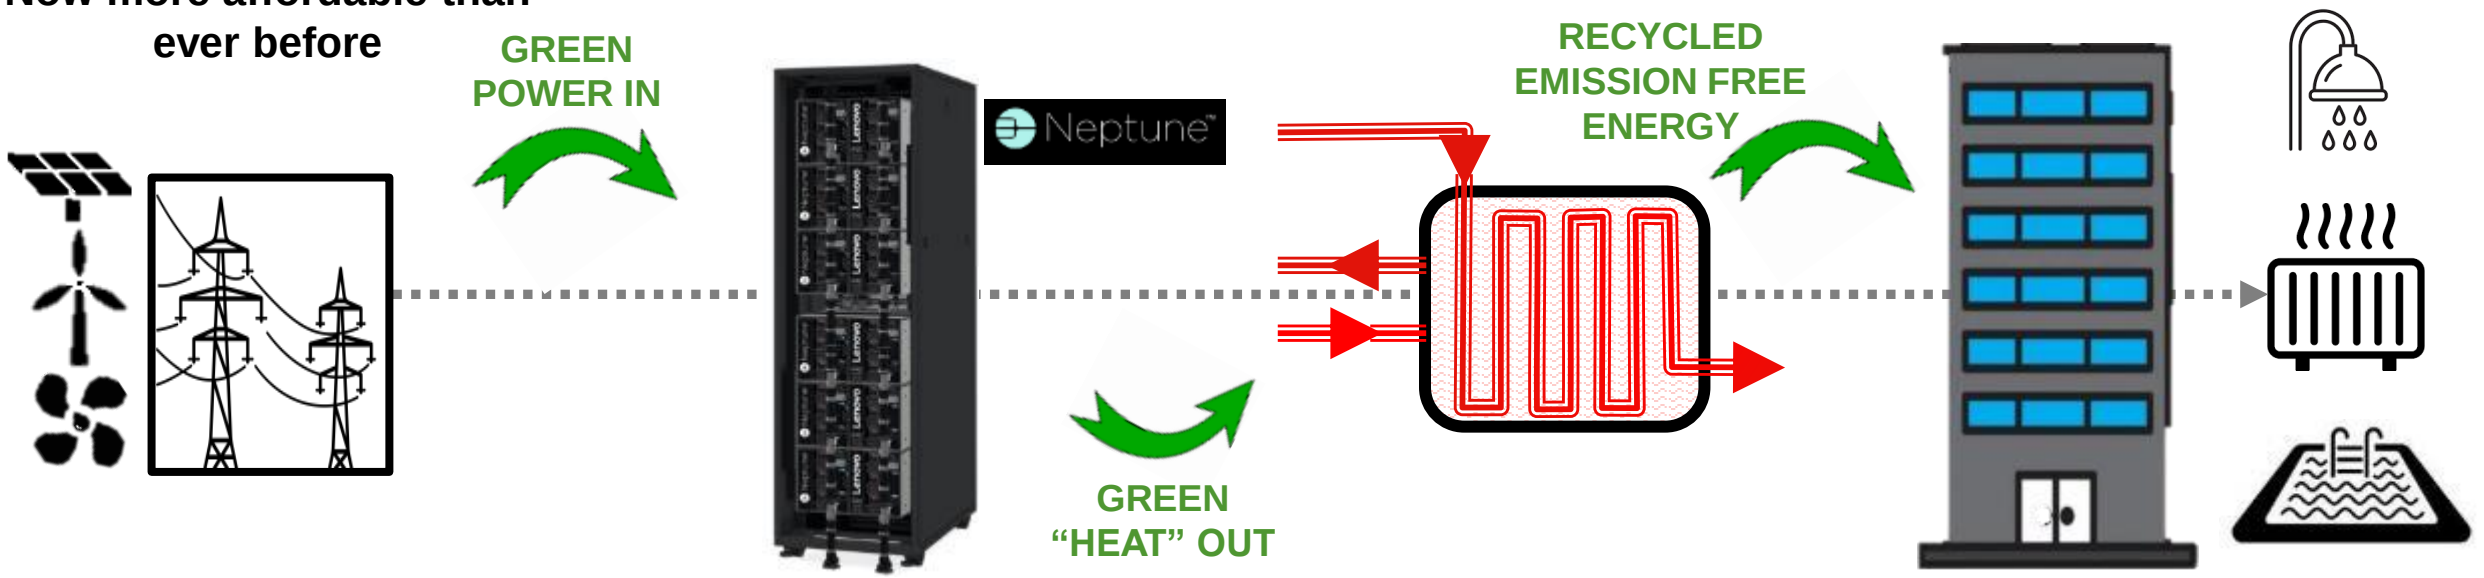

Smarter technology for all

Lenovo

thanks.