



ACCÆLSIUS

We Bring the Cool™

Everyone is talking about AI.

But is it really AI? **Hi performance GPU workloads** better description?

How are your organisation and your infrastructure being impacted?



AI is here, and it is  
here to stay.

It will transform our  
lives.

It will transform our  
businesses.

It will transform our IT  
infrastructures.

“ AI is one of the most important things humanity is working on. It is **more profound than electricity or fire.** ”

— Sundar Pichai, CEO  
Google



“ This is the **first time that a technology developed in Silicon Valley benefits the lives of everyday people so quickly and so tangibly.** ”

— Satya Nadella, CEO  
Microsoft



“ AI is on the fast track to **becoming ubiquitous** — at **home**, at **work**, and **every-where** in between. ”

60% of workers will use their own AI to perform their tasks.

— Predictions 2024: AI  
FORRESTER®

“ Generative AI could potentially raise annual labor productivity growth by around 1.5pp over a 10-year period, and eventually **raise global GDP by 7%.** ”

Gen AI raises the potential for a boom in labor productivity that significantly increases global growth.

— Joseph Briggs, Sr.  
Global Economist

Goldman  
Sachs



“ I am already seeing AI workloads leading to a broad proliferation of accelerated computing infrastructure. This will require investments in next-generation data center physical infrastructure to support new architectures with higher power and thermal management requirements. ”

— Lucas Beran,  
Research Director

DG





## Technology

Higher coding  
productivity



## Consumer

Conversion rate  
increase



## Biopharma

Research  
timeline  
reduction



## Financial Institutions

Higher fraud  
detection  
accuracy



## Entertainment

Higher quality of  
animated images



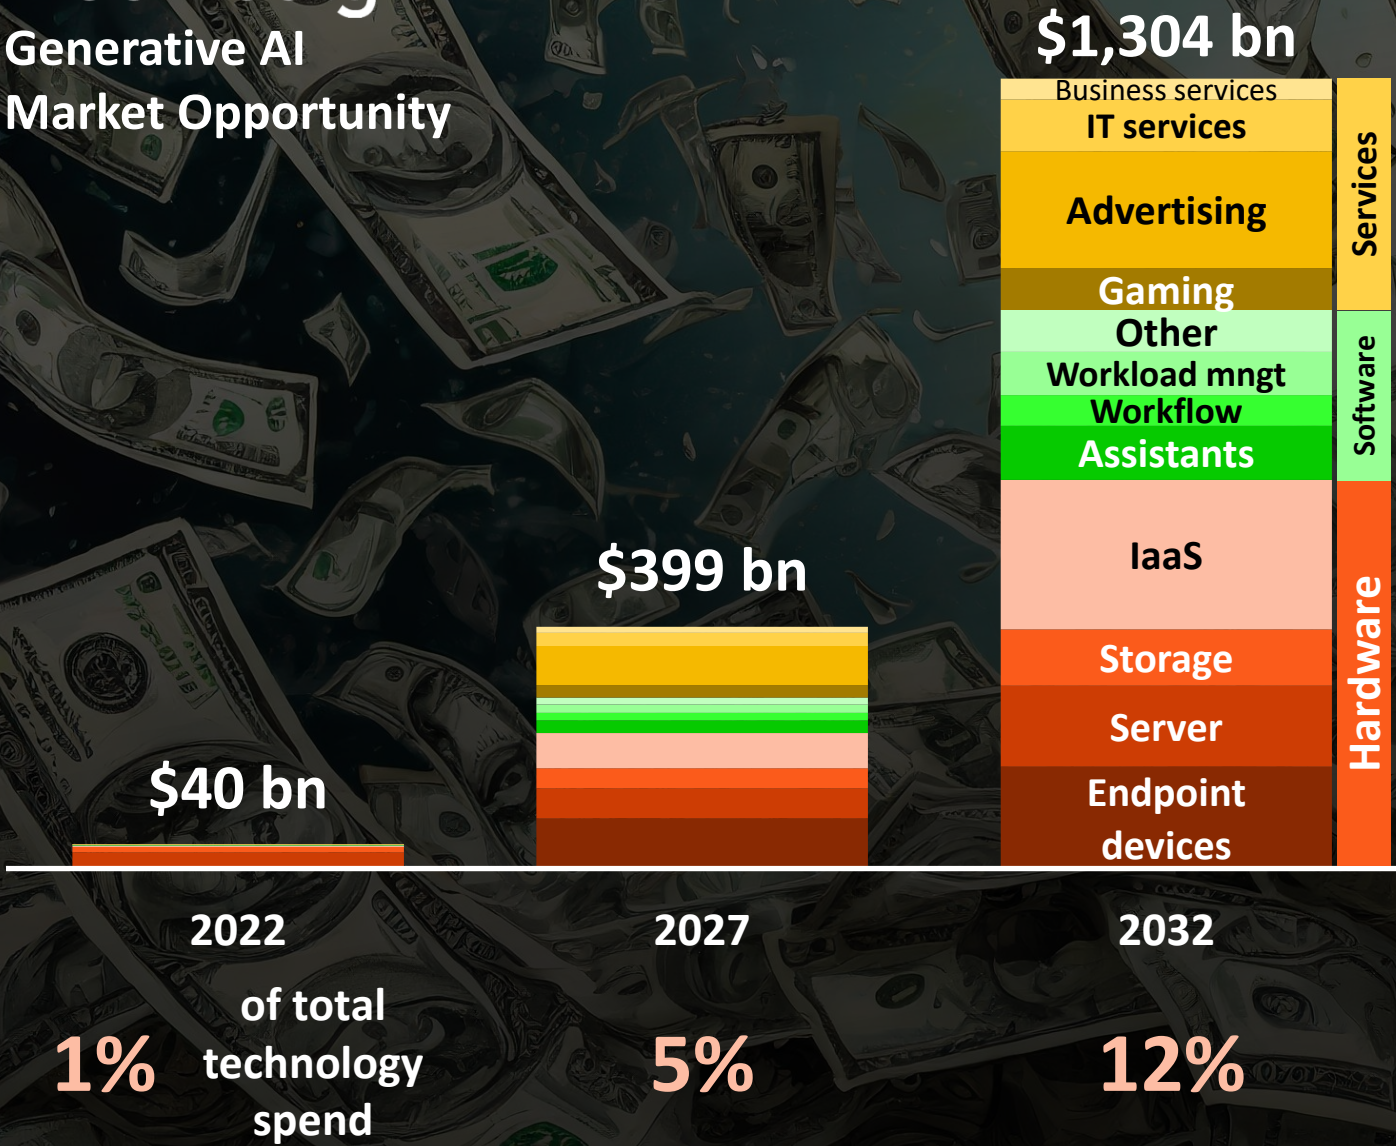
## Insurance

Customer service  
cost optimization

**Encouraging AI use-cases are bringing in profit to businesses across all industry verticals.**

# Bloomberg

Generative AI  
Market Opportunity



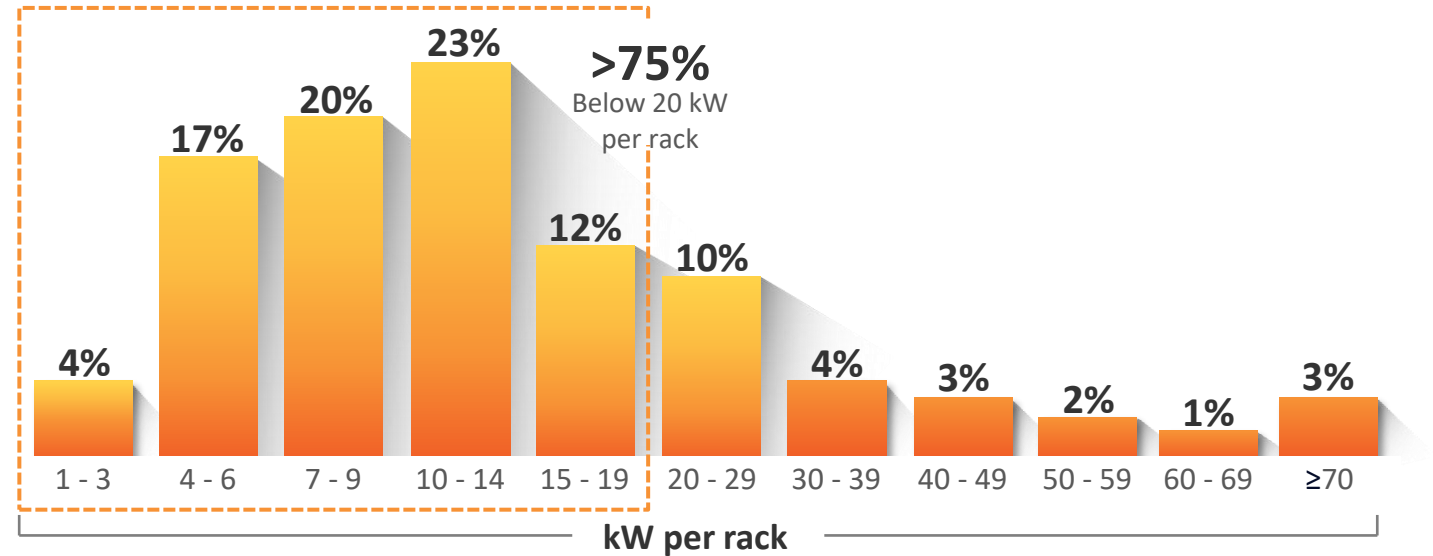
Gen AI is expected to be a **\$1.3-trillion business** by 2032.

Accelerated IT is  
changing the IT rack.

Currently most racks  
run loads up to 20kW,  
but 80% of IT  
managers agree  
**density is going up.**

70kW to 100kW  
projects live now!

Most racks are still in the low-density space.



But IT managers agree that rack density is on the rise.

Cloud / hosting / SaaS provider



Colocation / data center provider



Enterprise data center owner / operator



RAPIDLY  
INCREASED

SLOWLY  
INCREASED

REMAINED THE  
SAME

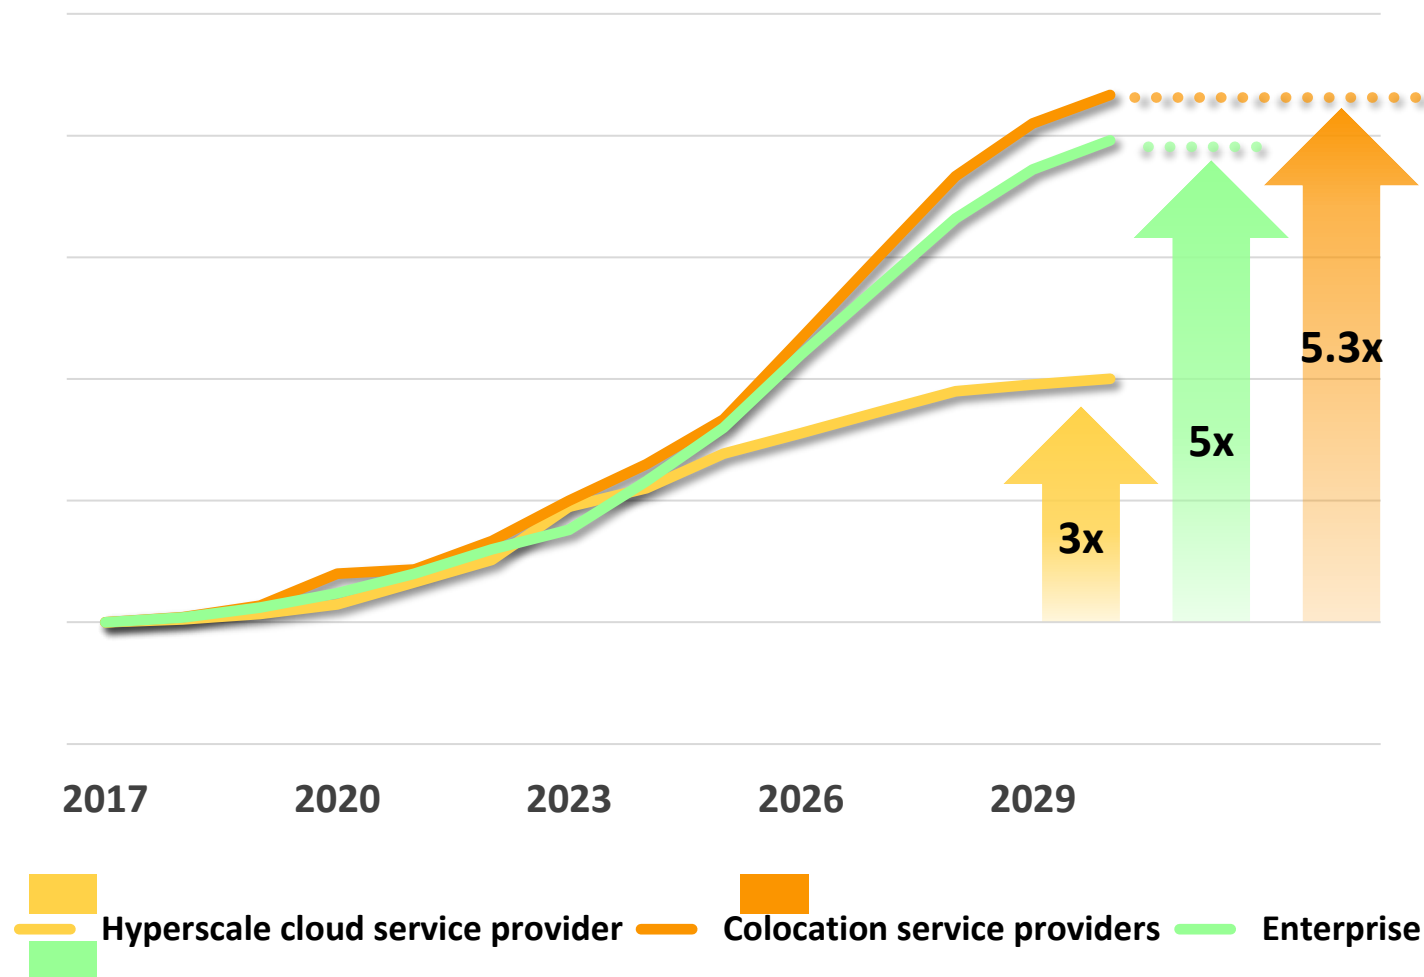
DECREASED

Projections point to doubling of rack densities every 12 months in coming years . In all environments, from enterprise to hyperscale.

We are already seeing this trend.

## Rack density growth is expected to take off.

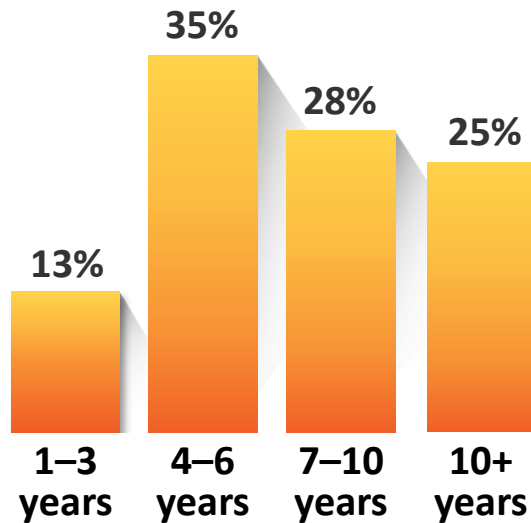
Avg kW per rack trend



Expectation that air cooling is **overtaken by liquid** as the dominant cooling technology.

*How long do you think air-cooling will be the dominant approach for data centers >1MW?*

Uptime  
Institute®



Source: Uptime Institute

## Liquid Cooling

In the next 2 years, we expect to see...

- Servers with **>2kW per U**.
- GPUs with over 600W, future **1.5kW** (400W/cm<sup>3</sup>).
- CPUs with over **300W**.
- Heat flux at the processors are **increasingly too high for air cooling**.

Overall trends:

- Trend towards **lower chip temperatures** and thus an increase in clock frequency leads to more computing power.
- Cooling **fluid temperature higher** than cooling air temperature.
- Up to **100% free cooling**.
- Better **waste heat utilisation**.
- **Higher density** in the data center.
- **Reduced space** requirement.

Tsunami or sea level increase?

Industry expects **liquid cooling to overtake air** in the next few years.

Accelsius is working side-by-side with Intel to develop power and cooling infrastructure and support its highest grade AI solutions.



Joint development of a product strategic for Intel's growth in the AI space.



Two different liquid cooling design options:

Refrigerant-to-air-cooled solutions handling one rack >40kW.

— or —

Refrigerant-to-liquid cooling system able to remove up to 160kW of heat load.



“ To support increasing thermal design power and heat flux for next-generation accelerators, Intel has worked with Accelsius and other ecosystem partners to enable an innovative cooling solution that will be critical in helping customers meet critical sustainability goals. ”

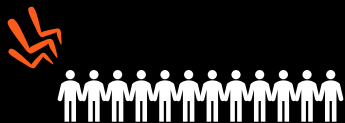
— Devdatta Kulkarni, Principal Engineer



Higher rack density brings on challenges and opportunities to IT infrastructures.

Are you already seeing these?

Skilled labor availability



Trade-offs between new builds and retrofits



Short turn-around times



Repeatable and quick-to-deploy designs



Costly IT gear



Airflow optimization



Profile of GPU loads



Higher capacity power systems



Power grid constraints



Higher amperage power distribution



Monetization of battery storage



PUE / CO<sub>2</sub> / sustainability targets



Wider and heavier duty racks



Compatibility of IT with cooling



Emerging cooling technology options



Expanding heat rejection capacity



Thermal chain



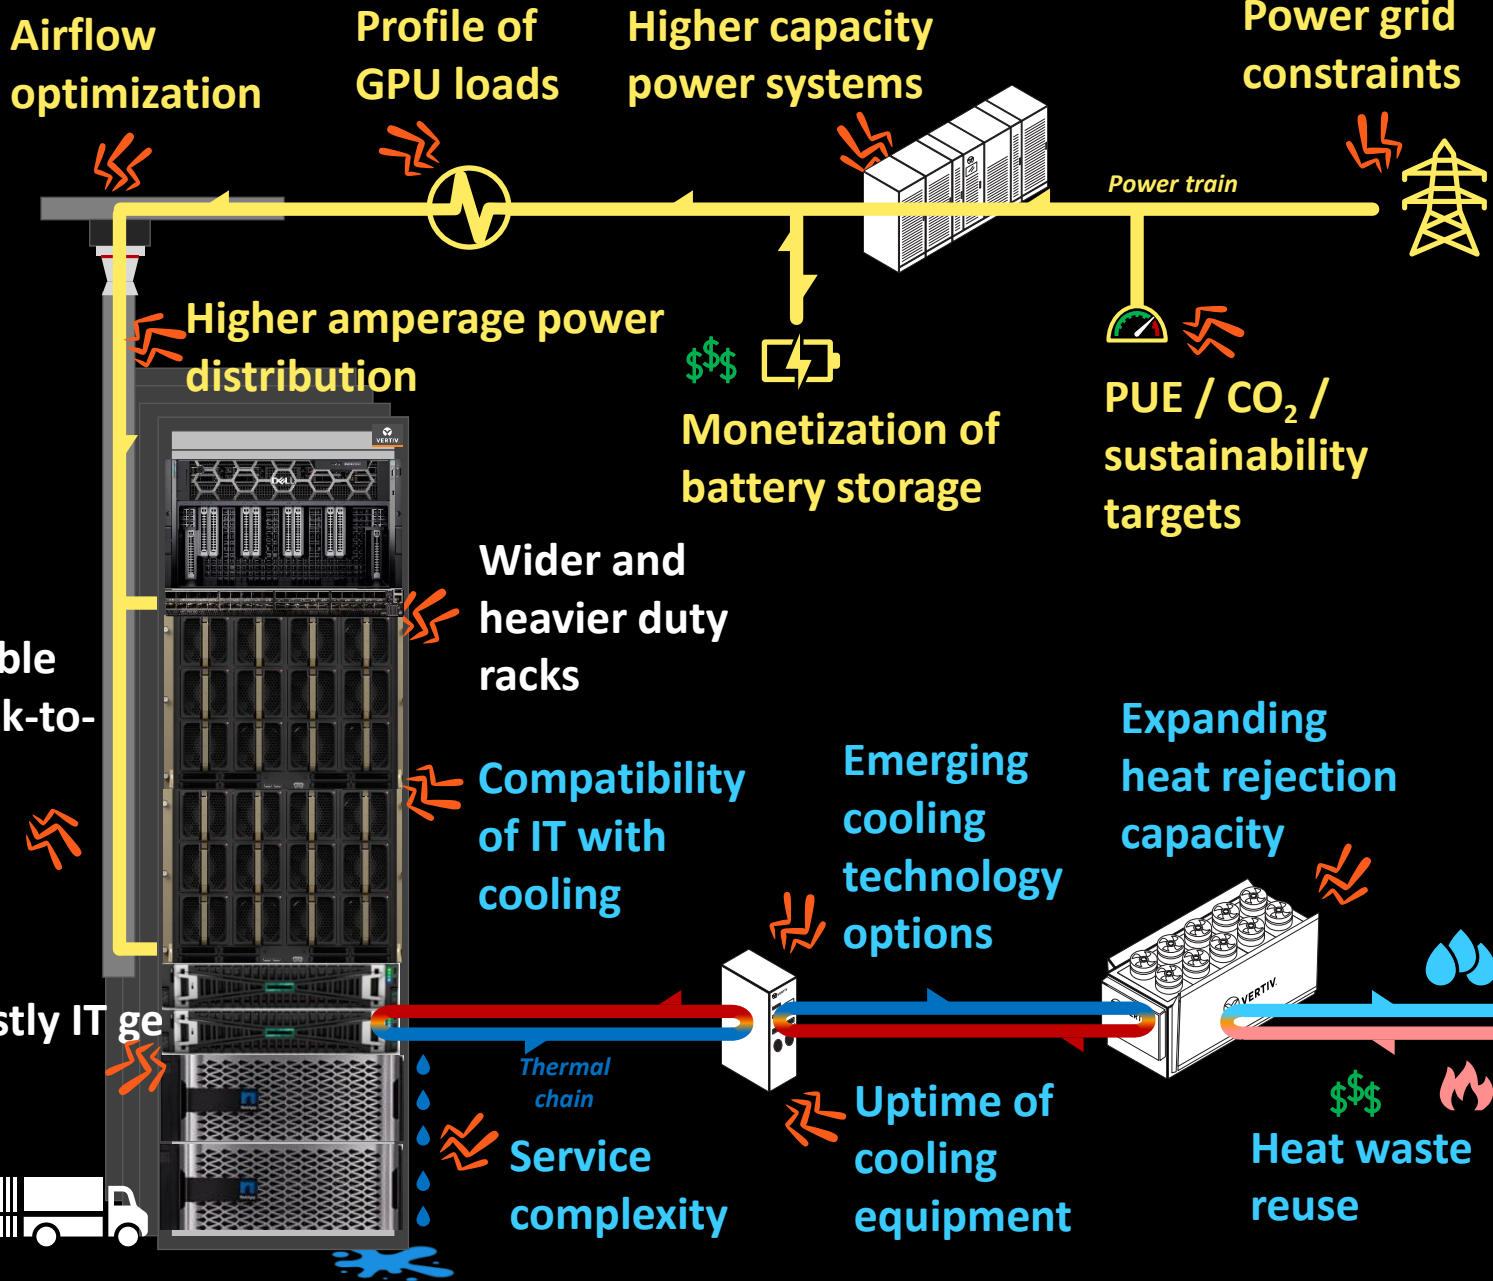
Service complexity



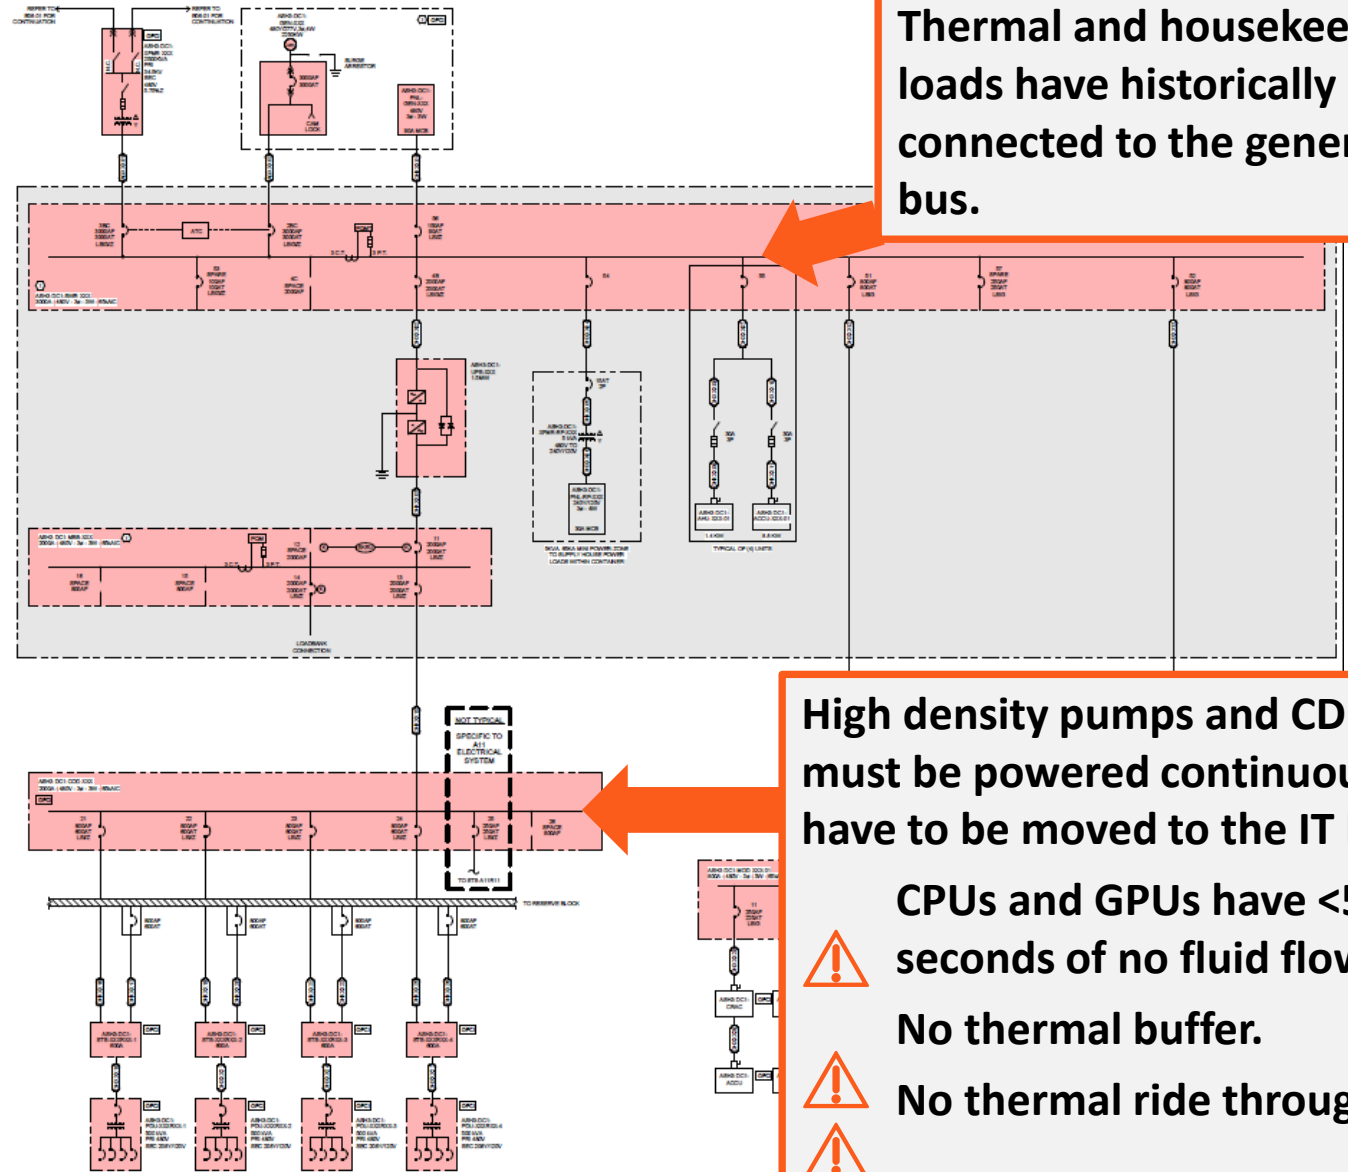
Uptime of cooling equipment



Heat waste reuse



Cooling equipment is added to the critical load as GPUs generate so much heat that thermal loads cannot go down.



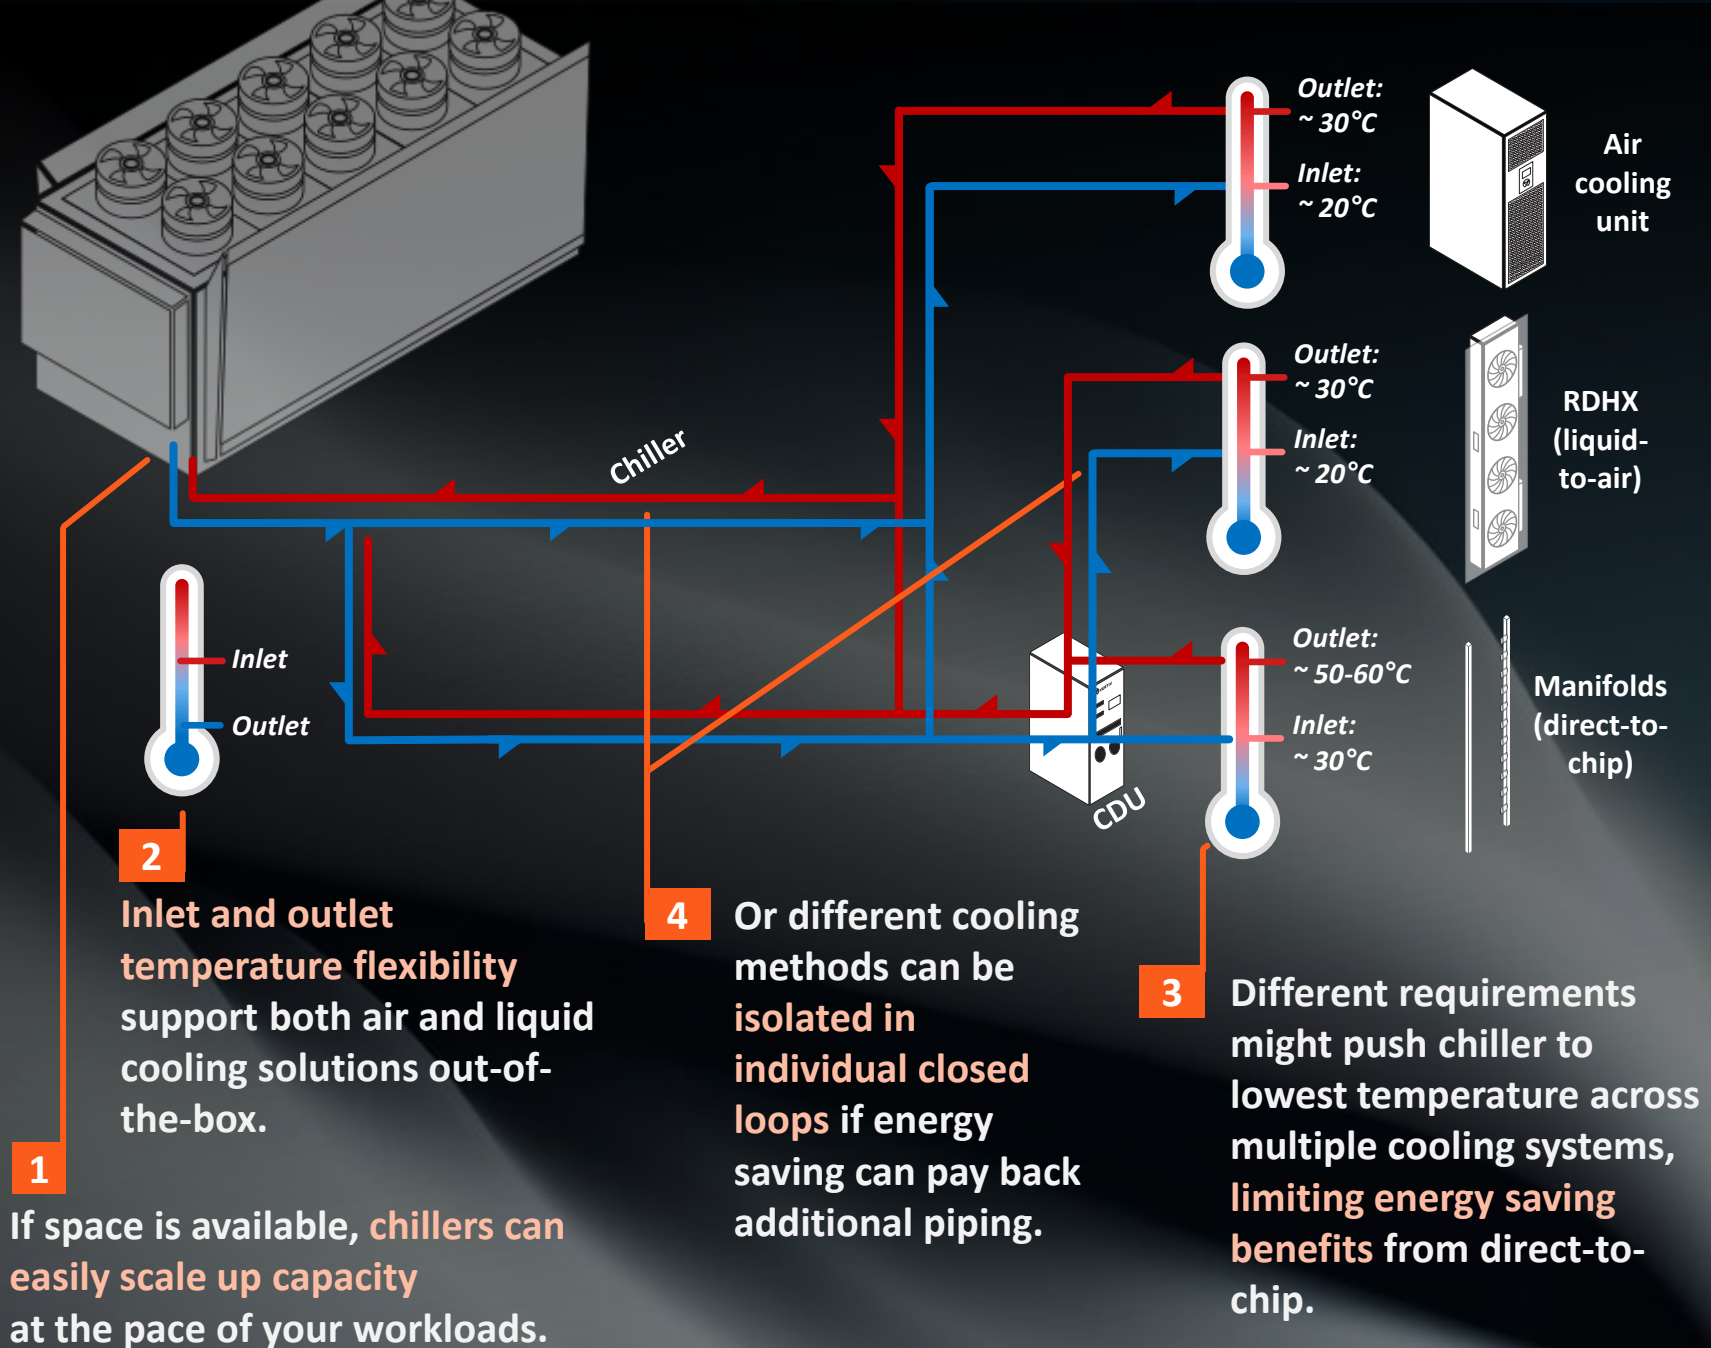
Thermal and housekeeping loads have historically been connected to the generator bus.

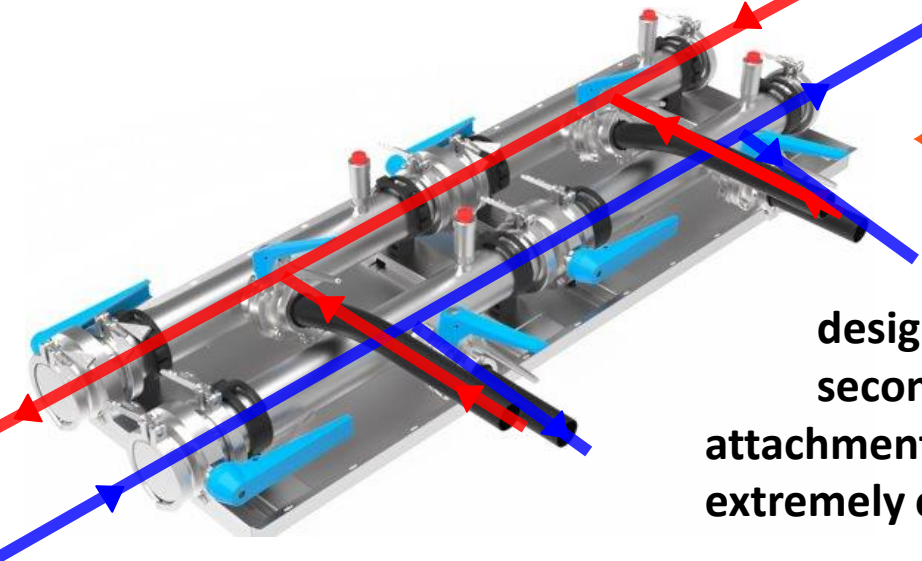
High density pumps and CDUs must be powered continuously so have to be moved to the IT bus.

- ⚠ CPUs and GPUs have <5 seconds of no fluid flow.
- ⚠ No thermal buffer.
- ⚠ No thermal ride through.

Primary fluid systems will need to adapt to new technologies.

Approach temperatures key to performance and efficiency.





**Cleanliness is paramount.**

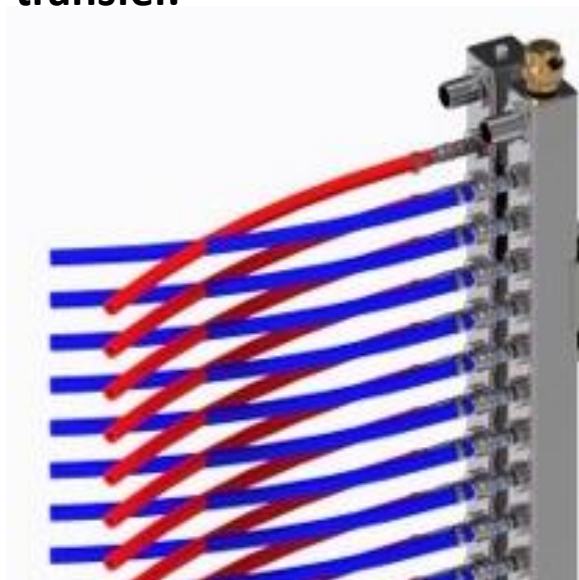
Unlike pumps and CDUs, current server designs do not offer multiple secondary fluid network attachment for redundancy, making them extremely critical.



Demanding **material quality** often food-grade or medical-grade to avoid corrosion and leaks.



Cold plates require **27µm filtration** to ensure optimal heat transfer.



Whether PG25 or DI water, fluid running in secondary networks require **specialised attention**:









- Flushing
- Regular testing
- Topping up
- Avoiding air in system

Secondary fluid networks are in direct contact with GPU cold plats and have become **as critical as power.**




# WE HAVE THE BEST THERMAL PERFORMANCE

At a Below Industry Average Price Point

	Air (1U)	1P Immersion	2P Immersion	1P Direct-to-Chip	NEUCOOL 2P Direct-to-Chip
Max TDP Cooled <u>with 30°C Facility Water</u> and Below Thermal Resistances (per socket)	165 watts	190 watts	415 watts	700 watts	1500+ watts 
Thermal Resistance (Watts/°C)	0.15	0.13	0.06	0.03	0.02 
Volume of Two-Phase Coolant	N/A	N/A	1,300 liters	N/A	30 liters 
Flow rate required to cool 500-watt CPU	N/A	N/A	N/A	2.2 LPM (assuming 4°C ΔT)	0.18 LPM (assuming 0.7 vapor quality) 
Total Heat Removed Per U with 30°C FWS	413 watts	475 watts	1038 watts	1400 watts	2000 watts 
Cost \$/Watt (CapEx)		\$3.21	\$4.23	\$1.53 	\$1.58 
Year 1 OpEx for 10MW @ 30°C	\$6.2M	\$2.5M	\$2.6M	\$2.4M	\$2.3M 

Sources: ASHRAE, The Green Grid, Analysis from Dell CTO Office, Accelsius Internal Tests, Chilldyne, The Gannet Group

 = Best in Class

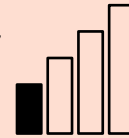
There is no “right” answer.

IT infrastructures will go through a **density journey** in the next few years.

Each density faces specific challenges and there are many solutions available.

**Low density**

*Up to 10kW*

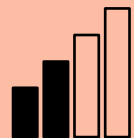


*Conventional compute workloads*

*Some accelerated compute, AI use case testing or inferencing workloads*

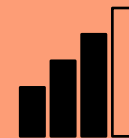
**Medium density**

*10 to 25kW*



**High density**

*25 to 80kW*

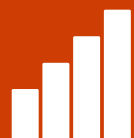


*Most accelerated compute, dedicated to both AI training and inferencing workloads*

*All accelerated compute, AI factory*

**Extreme density**

*> 80kW*



Existing cooling is not enough to meet needs of AI work-loads.

GPU density is outpacing traditional thermal deployment timelines.

