Evolution of Google's IT Liquid Cooling: From Prototype to 1GW Scale and Beyond

Jorge Padilla, Staff Product Design Engineer, Google
Madhu Iyengar, Principal Engineer, Google

August 24, 2025
Hot Chips Tutorial

Outline

- Motivation for liquid cooling
- Brief history of IT liquid cooling
- Google public announcements of liquid cooling
- Google liquid cooling system design
- Reaping the benefits of liquid cooling
- Design, deployment & operations for uptime
- Supply chain, shipping & handling
- Open standards
- Looking ahead
- Future challenges

AI Model Compute Scaling

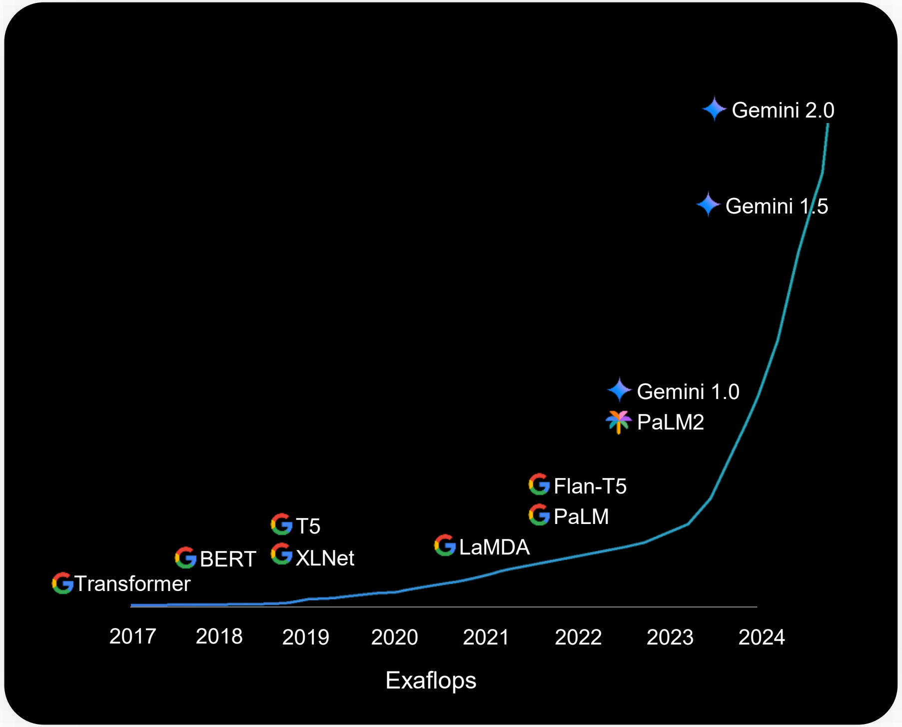

Recently released Gemini 2.5 !

Explosive growth in deployed ML capacity

Increasing demand for space and power which is driving the need for liquid cooling

A brief history of IT liquid cooling

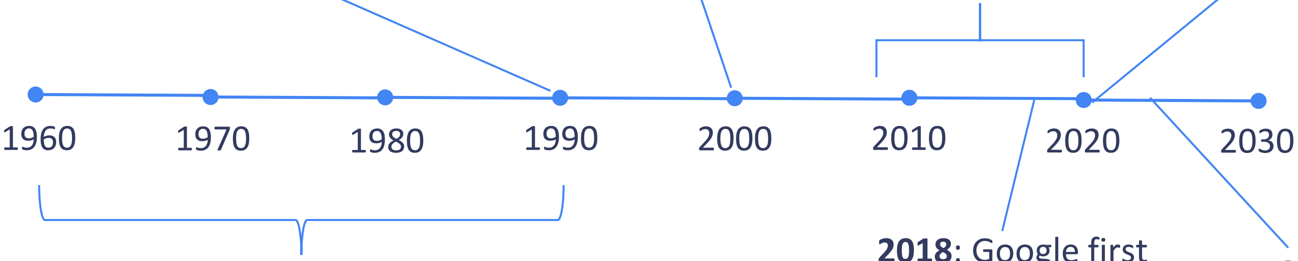

1960-1990: Liquid cooling IT in the industry

1990s: CMOS technologies led to significant reductions in chip powers such that air cooling was feasible and more economic

2000s: Chip power began to increase rapidly sparking renewed interest in liquid cooling R&D

2008-2020: Successful implementations of liquid cooling at different scales

2018: Google first public announcement of liquid cooled TPUs

2020s: AI boom and unprecedented chip powers leads to public announcements from hyperscalers adopting liquid cooling

2024: Google reports ~1 GW liquid cooled fleet capacity

Scale liquid cooling @ Google since 2018

- In 2018, Google publicly announced the introduction of liquid cooling in their data centers for machine learning TPUv3
- In 2020, Google announced the TPUv4 as their second liquid cooled ML platform.
- By 2024, Google's liquid cooled fleet had grown to ~1 GW total deployed fleet capacity

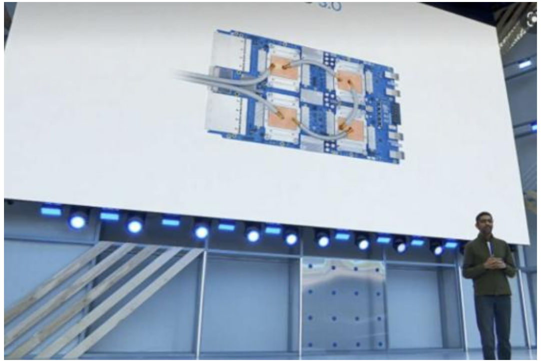

First announcement
2018 Google I/O developer conference
TPUv3

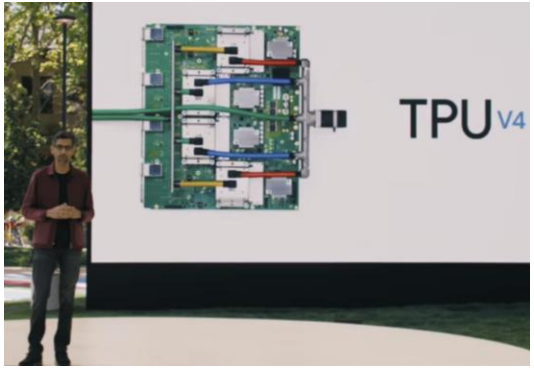

2020 Google I/O developer conference
TPUv4

2024 Google I/O developer conference
~1 GW of total deployed fleet capacity

2014-2024 Liquid Cooling R&D, Product Dev, Deployment

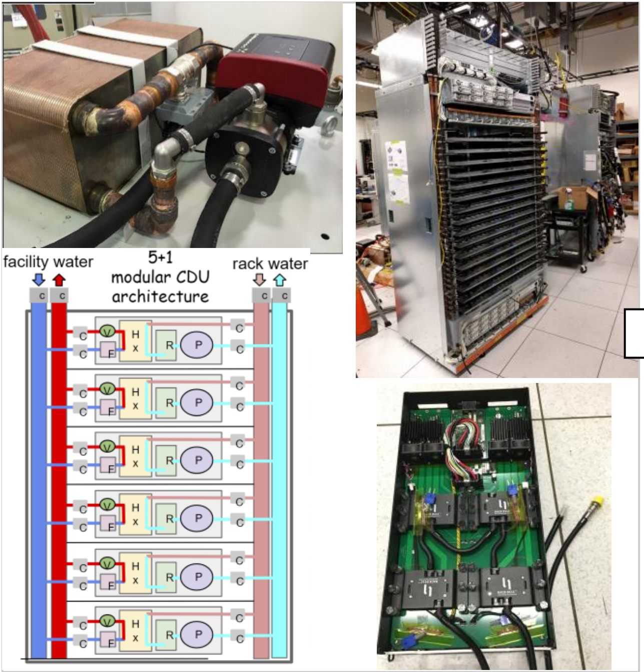

2014-2016

Ideation & experimentation Failure & iteration

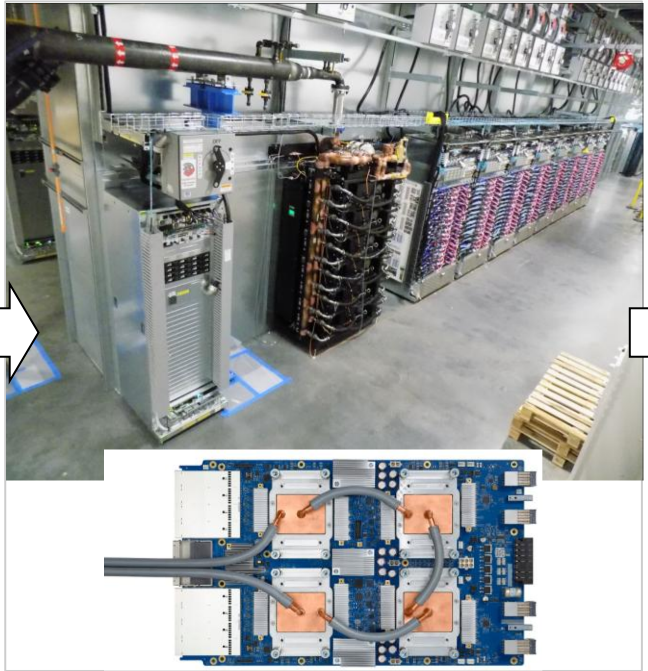

2017-2018

Market opportunity (TPUs) & cut in Productization & debugs

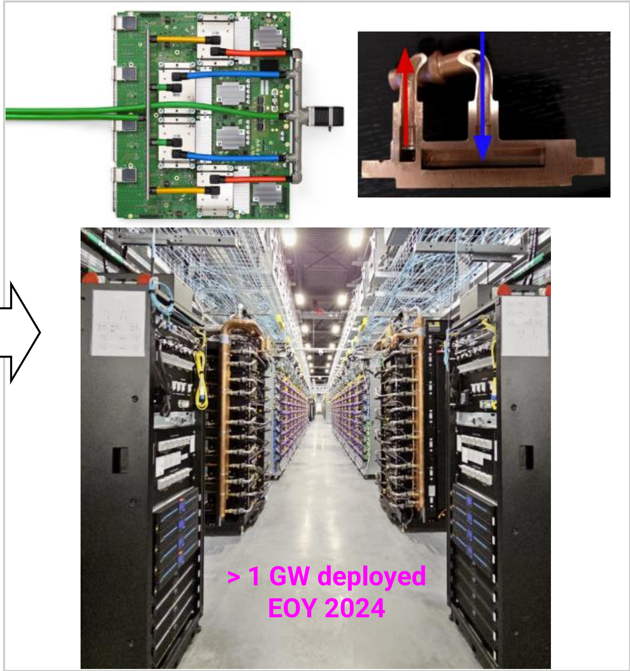

2019-2024

Design advancements & innovations
Scaling mfg/quality/deployment
Adoption across portfolio

How does liquid cooling at Google work?

- Coolant distribution Unit (CDU) pumps liquid (e.g. water) to/from IT racks via pipes
- Manifolds in IT racks distribute liquid to/from server, inside the server the liquid flows over chips
- Chip heat is transferred to liquid flowing through mechanically mounted cold plate copper fin passages
- Heat is ultimately rejected to data center water via heat exchangers in the CDU
- Uninterruptible Power Supply (UPS) rack provides backup power to CDU

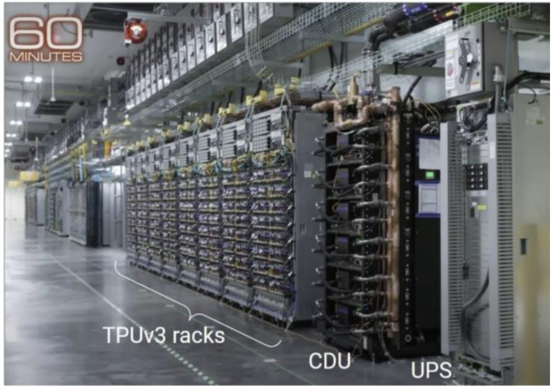

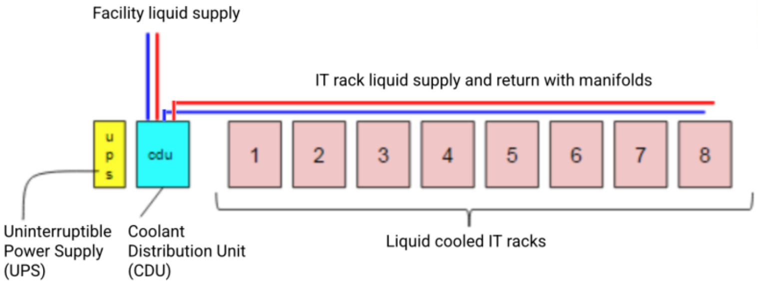

Facility liquid supply

IT rack liquid supply and return with manifolds

ups

cdu

1 2 3 4 5 6 7 8

Uninterruptible Power Supply (UPS)

Coolant Distribution Unit (CDU)

Liquid cooled IT racks

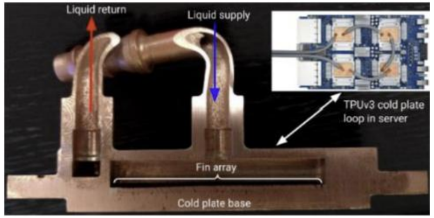

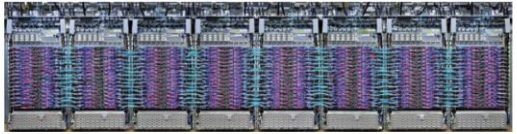

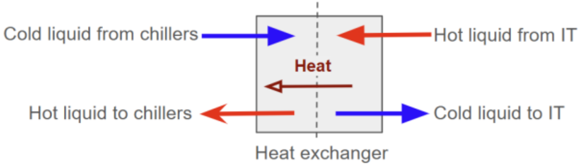

Cold liquid from chillers

Hot liquid from IT

Heat

Hot liquid to chillers

Cold liquid to IT

Heat exchanger

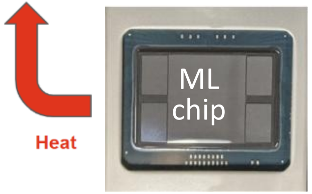

Heat

Google cold plate loop design

- Copper cold plate with microchannels
- Split flow configuration for lower thermal resistance (vs straight flow)
- TPUv3
  - Cold plates arranged in series → optimized to cool the “last chip” in the series
  - Passive
  - Lidded chip package
- TPUv4
  - Cold plate arranged in parallel → optimized for uniform flow
  - Active electro-mechanical cold plate loop → actively controlled valve for closed loop flow control
  - Bare die chip package → Cold plates with pedestal to reach down to contact bare die for better heat transfer from chip to cold plate

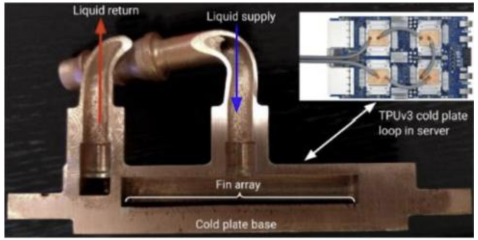

Split Flow Rca (center of copper lid) and Straight Through Flow Rca (center of copper lid)

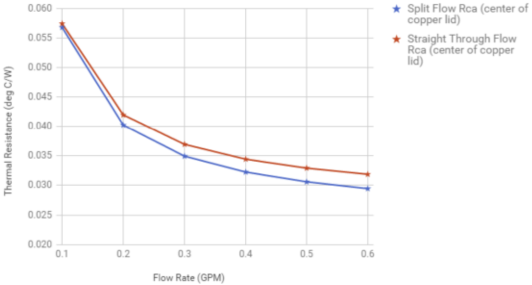

Split flow cold plate with microchannels

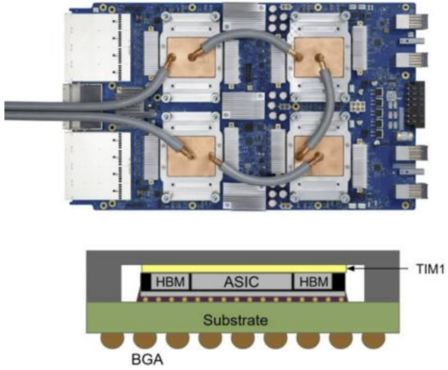

TPUv3 cold plate loop
Lidded chip package
Series configuration
Passive

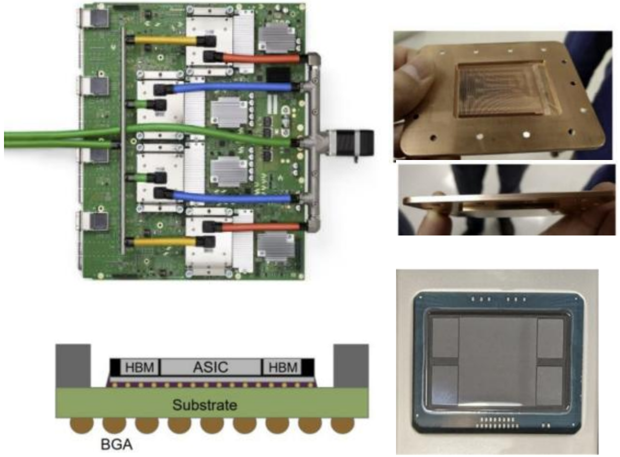

TPUv4 cold plate loop with pedestal
Bare die chip package
Parallel configuration
Actively controlled flow control valve

Google CDU design

- CDU rack consists of 6 modular Coolant Distribution Units (mCDUs)
- mCDU
    - Contains heat exchanger, pump, variable frequency drive, strainers & valve
    - Connect to CDU manifolds with flexible hoses and quick disconnect couplings
    - Quick disconnect couplings enable isolating each mCDU for service.
- Flexible hoses eliminate need for tight mechanical tolerances and facility deployment and service
- CDU rack contains expansion tank, air separator leak detection system, Programmable Logic Controller (PLC) and Power Distribution Unit (PDU)
- A valve on the facility side plumbing circuitry is controlled by the PLC to maintain a fixed approach temperature.
- The pump speed is controlled by the PLC to maintain a fixed pressure differential on the IT liquid side of the CDU.

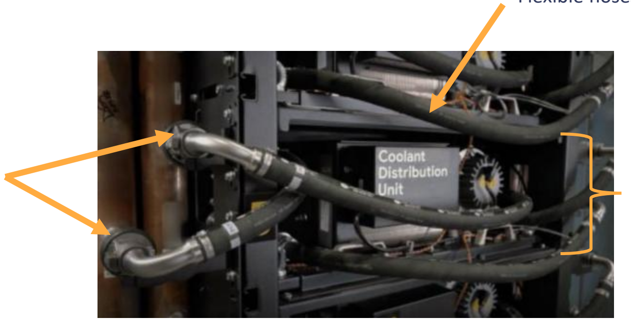

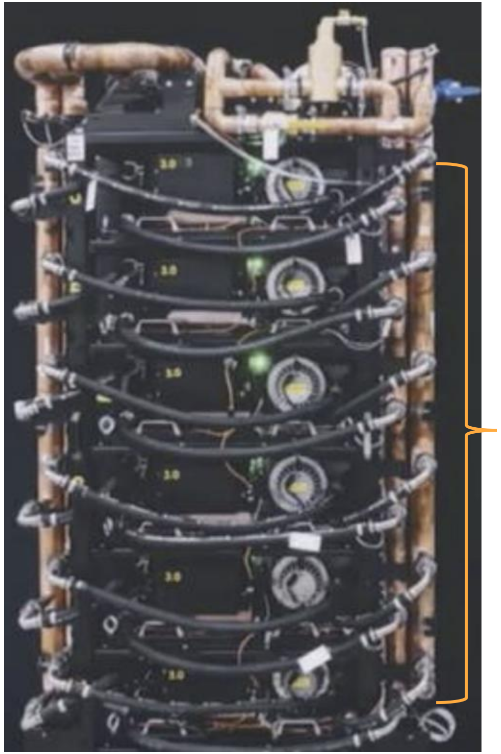

CDU rack with 6 mCDUs

Reaping the benefits of liquid cooling

- Liquid cooling enabled a **doubling of deployed chip density** and **quadrupling** the TPUv3 ML **supercomputer size** vs air cooled TPUv2
- **Thermal design power** (TDP) of the liquid cooled TPUv3 liquid cooling was **1.6x greater** than the air-cooled TPUv2
- TPUv3 **frequency was 30% greater** than TPUv2
- **Pump power consumption** per liquid cooled TPUv3 was **less than 5%** of the per chip **fan power** required for air cooling commensurate TDPs.

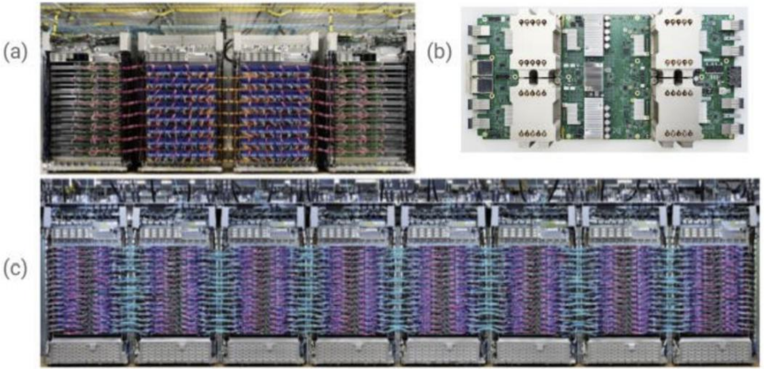

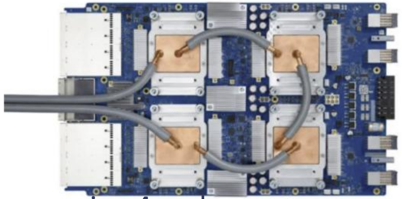

(a) Air cooled TPUv2 supercomputer consisting of 256 chips spanning 4 racks
(b) Top view of air cooled TPUv2 server
(c) Liquid cooled TPUv3 supercomputer consisting of 1024 chips spanning 8 racks

Design, deployment & operations for uptime

- **Google liquid cooled fleet uptime has consistently been 99.999% since 2020**
- **Design for uptime**
    - Of the 6 mCDUs, 1 is redundant such that the CDU capacity is based on N+1 mCDU redundancy.
    - Custom, certified hardware tools for deployment and operations
- **Deploy for uptime: installation & commissioning**
    - Leak testing
    - Filling and purging
    - Pump functional testing
    - Functional testing of telemetry, monitoring, alarms
- **Operations for uptime**
    - Preventative maintenance
    - Telemetry, monitoring & alarms
    - Corrosion and biogrowth detection
    - Leak detection and response with safety protocols
    - Emergency operating procedures
    - Maintaining charge pressure targets
    - Filtration

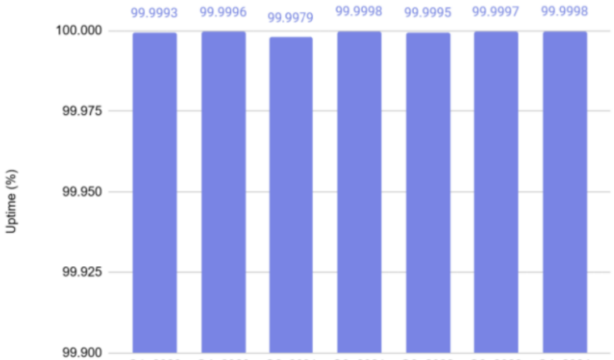

Liquid cooled fleet uptime of 99.999% consistently since 2020

“Feeder kit” custom tool used for deployment & operations

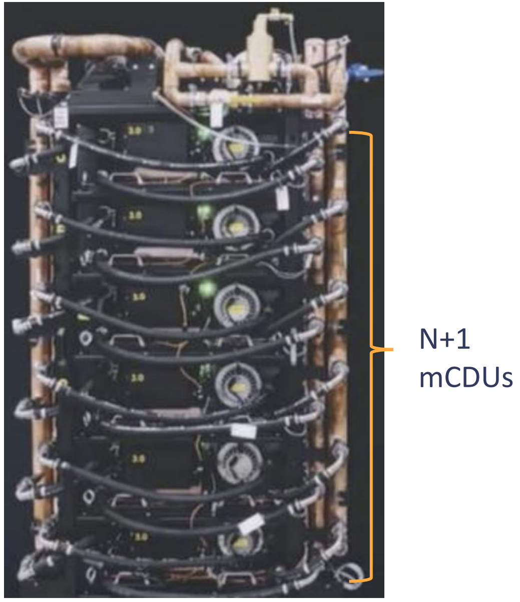

CDU with N+1 redundancy

Considerations for scale adoption

Supply chain, ecosystem

- Liquid cooled ITE is typically a high performance product
    - Requires the greatest standards of end-to-end thermal-hydraulic performance, quality, reliability & supply chain robustness
- Need multi-sourcing from healthy, stable & technically proven suppliers
    - Cold plates
    - Pumps
    - Coolant Distribution Units (CDUs)
        - Decouples facility cooling loop from ITE liquid cooling loop to deliver precisely controlled liquid temperature, pressure, flow and cleanliness to ITE
    - Heat exchangers, e.g. air/liquid & liquid/liquid
    - Hard and flexible tubing
    - Liquids, e.g. water, glycol, dielectric
- Interoperable, multi-source rack & tray level liquid interconnects
- Compliance certifications, e.g. UL, CE

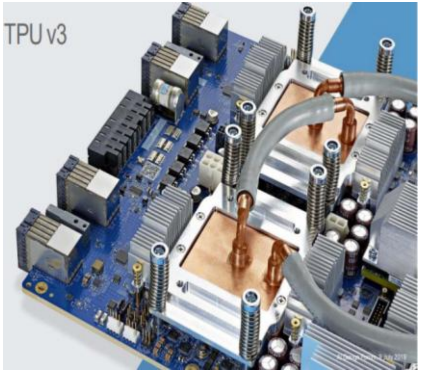

AI Design Forum at SEMICON West 2019, Cliff Young (Google Research)

Considerations for scale adoption

Testing, shipping & handling models: suppliers → system integrator → DC

- Need standardization of reliability & testing processes for leak & pressure joints
- Performance testing, e.g. thermal, hydraulic, mechanical, electrical
  - Component level @ supplier
  - Sub-system level @ system integrator
  - Full system level @ DC
- Shipping wet vs dry
  - Consider environmental conditions, e.g. temperature, pressure during shipment by air, land or water
  - Identify/develop tools for equipment filling, purging and drying processes

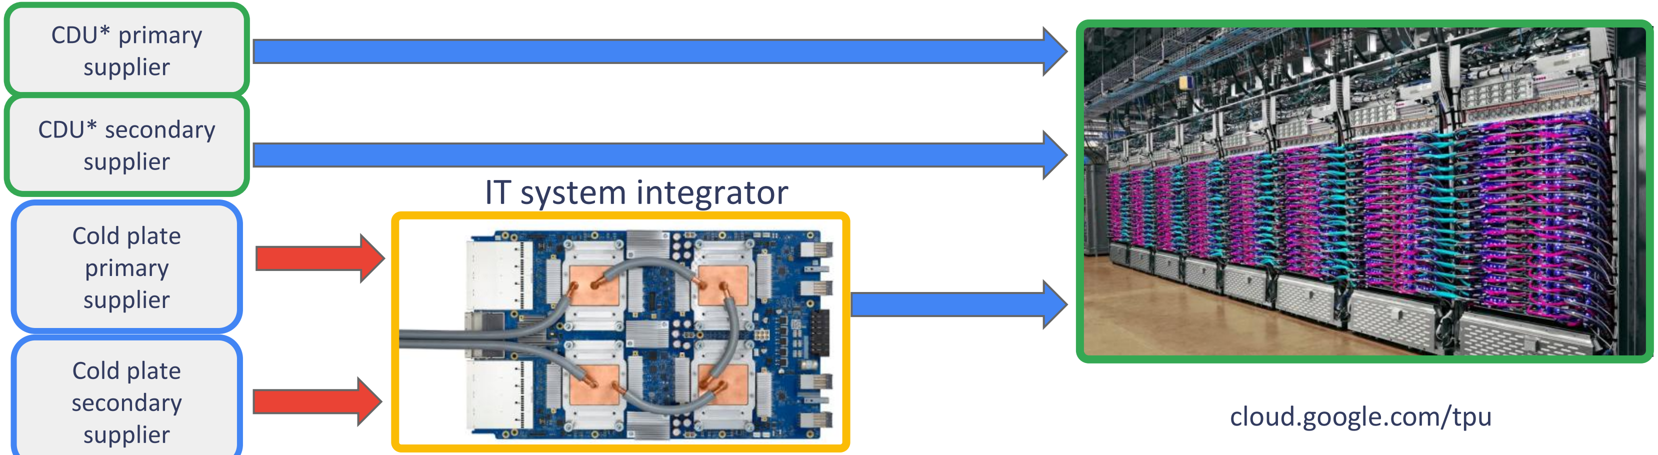

cloud.google.com/tpu

Google TPU v3,
cloud.google.com/tpu

* CDU: Coolant Distribution Unit
(pumping & heat exchange unit/system)

Considerations for scale adoption

Deployment model

- Need to develop a deployment model that enables efficient landing of high performance liquid cooled ITE
- DC building construction and facility readiness
- DC building & floor preparation & construction for landing of liquid cooled ITE
  - Piping
  - Power
  - Telemetry
- CDU installation and commissioning follows DC floor readiness
  - Leak testing
  - Filling and purging
  - Pump functional testing
  - Functional testing of telemetry, monitoring, alarms
- High performance ITE installation follows CDU commissioning

www.google.com/about/datacenters/gallery/#mayes-county-exterior-sunsets

www.google.com/about/datacenters/gallery/#douglas-county-pipes

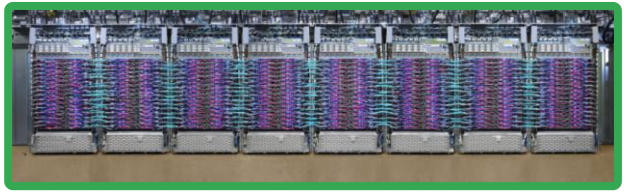

techcrunch.com/2019/05/07/googles-newest-cloud-tpu-pods-feature-over-1000-tpus/

Considerations for scale adoption

Ops and Serviceability

“Feeder kit” custom tool used for deployment & operations

- Ops & Serviceability
  - Preventative maintenance
  - Telemetry, monitoring & alarms
  - Corrosion and biogrowth detection
  - Leak detection and response with safety protocols
  - Emergency operating procedures
  - Maintaining charge pressure targets
  - Filtration
- Custom, certified DC liquid cooling tools & equipment are needed for deployment & operations
  - Tools for filling, purging, drying, testing, cleaning, etc.
  - May need not be available Off-the-Shelf → require development
- Engineering teams must design system with ops & serviceability in mind

Open liquid cooling initiatives and standards

- Scale adoption may be accelerated by leveraging and driving clarity on interface specifications.
- Some recommended interface specifications include:
    - Wetted materials (metallic & polymer) compatibility,
    - Fluid operating pressure
    - Temperature
    - pH ranges
    - Interoperable quick disconnects
    - Filtration
    - Reliability
    - Flammability
    - Power and electrical
    - Controls and monitoring
    - Serviceability and certifications
- Google co-authored the “Open Specification for a Liquid Cooled Server Rack” in 2018
- Google to share the Gen 5 CDU specification in 2025

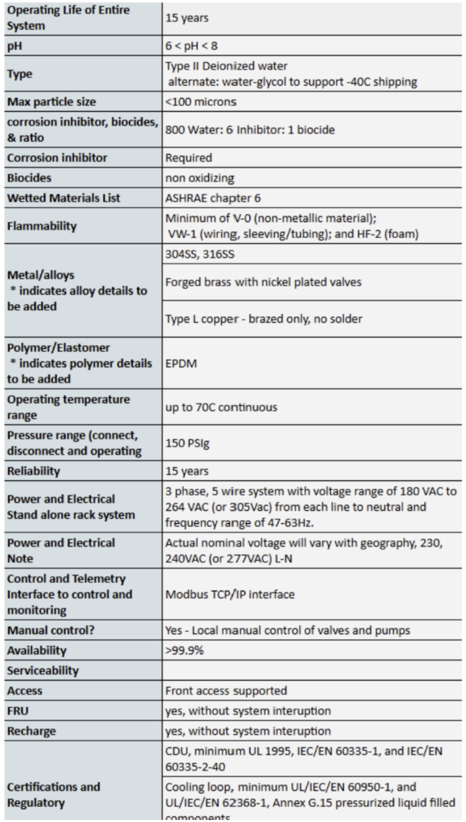

Operating Life of Entire System	15 years
pH	6 < pH < 8
Type	Type II Deionized water alternate: water-glycol to support -40C shipping
Max particle size	<100 microns
corrosion inhibitor, biocides, & ratio	800 Water: 6 Inhibitor: 1 biocide
Corrosion inhibitor	Required
Biocides	non oxidizing
Wetted Materials List	ASHRAE chapter 6
Flammability	Minimum of V-0 (non-metallic material); VW-1 (wiring, sleeving/tubing); and HF-2 (foam)
Metal/alloys * indicates alloy details to be added	304SS, 316SS
	Forged brass with nickel plated valves
	Type L copper - brazed only, no solder
Polymer/Elastomer * indicates polymer details to be added	EPDM
Operating temperature range	up to 70C continuous
Pressure range (connect, disconnect and operating	150 PSIg
Reliability	15 years
Power and Electrical Stand alone rack system	3 phase, 5 wire system with voltage range of 180 VAC to 264 VAC (or 305Vac) from each line to neutral and frequency range of 47-63Hz.
Power and Electrical Note	Actual nominal voltage will vary with geography, 230, 240VAC (or 277VAC) L-N
Control and Telemetry Interface to control and monitoring	Modbus TCP/IP interface
Manual control?	Yes - Local manual control of valves and pumps
Availability	>99.9%
Serviceability	
Access	Front access supported
FRU	yes, without system interuption
Recharge	yes, without system interuption
Certifications and Regulatory	CDU, minimum UL 1995, IEC/EN 60335-1, and IEC/EN 60335-2-40
	Cooling loop, minimum UL/IEC/EN 60950-1, and UL/IEC/EN 62368-1, Annex G.15 pressurized liquid filled components

Looking ahead…

Design principles

Design for fungibility for next gen AI

Flow, pressure, temperature across IT suppliers

Quality & reliability with wetted materials

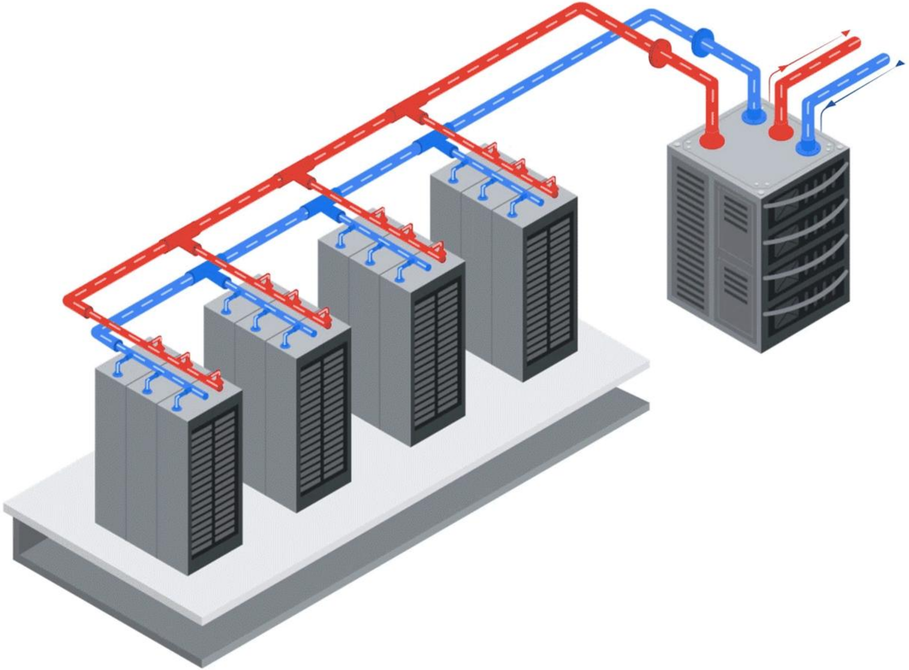

At scale insights

Google has > 1 GW of liquid cooling infra

~20 data centers with ~99.999 rack uptime

4th gen CDU embodies learnings, deploying now

Google’s 4th generation CDU

Enabling the industry

5th gen CDU…

Project Deschutes is in development

Get ready!

- 1 MW racks
- Project Deschutes CDU coming to OCP

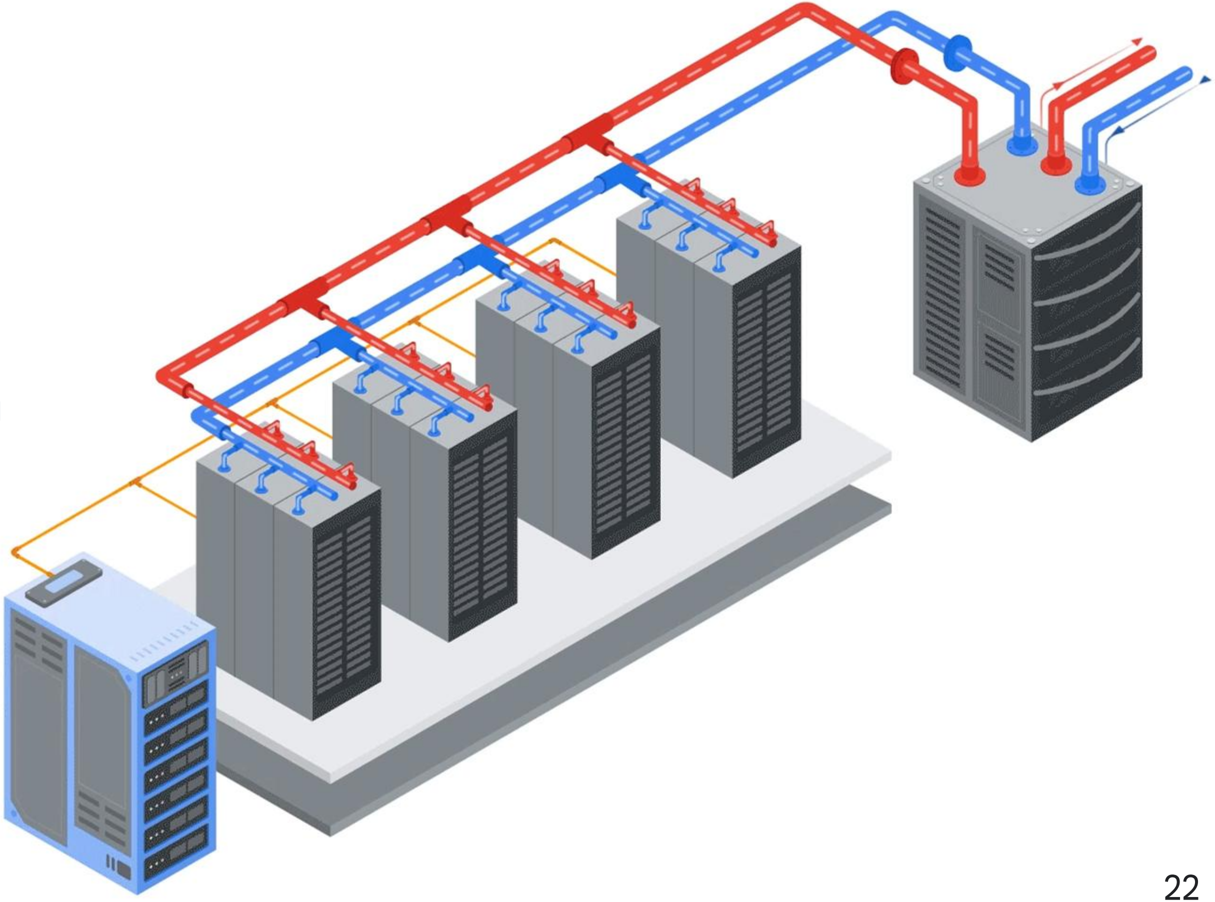

Challenges facing the industry

- **Performance increase require high power**: Rapidly increasing ML chip powers and hot spots will drive thermal-hydraulic requirements (greater liquid flow, pressure, and lower temperature coolant)
- **Evolving silicon tech**: New silicon structures have 3D geometries with tortuous heat paths
- **Supply**: AI demand is driving greater equipment quantities from an emerging supply chain
- **Speed**: AI demand is driving faster deployment times of large/complex liquid systems
- **Data Center**: Liquid cooling needs to be designed to be deployed and reliably operated in hyperscale facilities as well as 3rd party data centers. Sustainability concerns include construction, power, water.