


Hotspot-Aware Cooling for AI Data Centres: Precision Flow Optimization for Thermal Efficiency

By implementing a precision coolant flow system in **10%** of Microsoft's data centres (**30 sites**), we can unlock **\$2.4B** in total savings by **2050** while preventing over **6 million tons** of CO₂.



x

TKS

OPPORTUNITY

40% of Data Centre Energy Wasted on Cooling

AI accelerators generate intense, localized heat that traditional cooling systems struggle to manage. Cooling can consume up to **40%**¹ of a data centre's total energy. Standard cold plates distribute coolant evenly, which ignores heat concentration in specific areas. Even one hotspot can throttle the entire chip's performance. Additionally, narrow cooling channels increase pressure drop, requiring more energy to circulate coolant efficiently.

SOLUTION

Hotspot-Aware Cold Plate Designs using Computer Analysis

This system uses **topology optimization** and **fluid dynamics modeling** to design custom cold plates that **match** the chip's heat distribution. By analyzing a detailed 2D power map of the chip, the cooling layout is shaped to **direct more coolant to hotter regions** and less to cooler ones. Instead of using uniform channels, it generates **variable-density pathways**: narrow, tightly packed microchannels where cooling is most needed and wider, more open channels elsewhere.

IMPACT

13% Less Heat Rise, 55% Lower Pressure, \$2.4 Billion Saved by 2050

Reduces chip temperature rise by 13% and lowers coolant pressure drop by **55%**, leading to a **15–20%** reduction in total cooling energy use. This translates to **\$2–5 million** in annual savings **per site**. If adopted at just **10% of Microsoft data centres** (30), it could generate **\$2.4 billion** in cumulative savings by **2050**—enough to fund up to **five** new facilities. Over the same period, it would prevent more than **6 million tons** of CO₂ emissions, cut water use by up to **10.8 million** liters per site annually, and improve hardware longevity and rack-level performance.

Status Quo: Traditional cooling accounts for nearly **40%** of all data centre energy use

AI Chips produce uneven heat due to specific high power consuming regions (hotspots). However, uniform cooling (such as cold plates) spread all cooling evenly, failing to address the hotter localized regions.

While some regions are overcooled, these hotspots remain too hot, which slows down the entire chip which reduces performance and efficiency.

Additionally, traditional narrow straight channels result in high resistance levels, which thus requires more energy to pump coolant.

To compensate, chillers are used, which consume great quantities of water and electricity. Due to this inefficiency, cooling accounts for almost **40%**⁵ of all energy used in a data centre.

Traditional Cooling Wastes Energy and Limits Performance

40% of Data Centre Energy Wasted on Cooling



Cooling systems are one of the biggest energy drains in data centres — consuming up to **40%**⁶ of total energy use. At ~\$0.10/kWh, this equals **\$10.5 million** per data centre annually. With over 300 Microsoft data centres globally, this adds up to ~\$3 billion/year just to keep equipment cool.

Inefficient Cooling Systems Overcool Entire Chips



Most systems rely on uniform cold plates, which spread coolant evenly across the entire chip. These systems ignore hotspots, forcing coolant to cool low-heat areas unnecessarily. This inefficiency drives up power use and increases costs.

Hotspots Throttle Performance:

AI accelerators and high-performance servers produce intense, concentrated heat. Just one unchecked hotspot can cause the entire chip to throttle performance. Approximately **30%**⁷ of all data centre interruptions and outages are caused by server failures, many of which are heat-induced.



Our Recommendation...



Hotspot-Aware Cooling Tailors
Fluid Paths to Heat Maps

We propose an **optimized coolant routing system** that adapts to the chip's heat map, delivering more cooling to where it's needed most—**the hotspots**—and less to cooler regions.

In **hotspot** areas (where the chip gets very hot):

We use **narrower, densely** packed coolant channels

This increases **coolant contact** and **pulls away more heat** as fluid moves through narrow paths quicker

In **cooler** areas (that don't need much cooling):

We use **wider, spaced-out** channels

Reduce flow resistance (through lower pressure drop) and save pump energy, as well as efficiently **deliver coolant** to hotspots and **remove heated fluid** downstream

Why it works

By redirecting coolant based on each area's heat output, the system **avoids wasting energy** when overcooling the entire chip just to manage one hotspot.

How will these **optimized channels** be designed?

Use **topology optimization (TO)** to **custom-design** cold plate cooling geometry

By using **Computational Fluid Dynamics (CFD)** for heat transfer simulation, a computer finds the coolant layout that **most efficiently** targets chip hotspots.

↓ ↓ ↓
Adjusting variables like **density** and **diameter** of channels



This unique layout is then **manufactured** as a **cold plate** that can be fitted onto **existing chip racks** for improved cooling.

Technical Explanations

Goal

Minimize max junction temperature (temperature at the core of the chip, where all transistors and logical units are) **within tight pressure drop constraints** (prevent overuse of pump energy)

Case Study: Glacierware's Breakthrough in Cooling Efficiency.

Researchers from Corintis introduced a breakthrough cooling system called Glacierware, designed to improve data centre efficiency by targeting chip hotspots directly.

55% Lower Pressure Drop:

By precisely targeting hotspots, the system reduced coolant pressure requirements, cutting pumping costs.

Improved Hardware Longevity:

By reducing extreme heat fluctuations, server hardware was less prone to failure, extending its lifespan.

13% Lower Chip Temperatures:

The optimized coolant flow reduced chip temperature rise by ~13%, improving thermal stability.

3–6 MW Energy Savings:

Large data centres that adopted this solution saw a significant reduction in cooling power demand.

The Glacierware study shows that targeted coolant flow not only cuts cooling costs but also unlocks new performance potential, making it a powerful solution for sustainable, high-efficiency data centres.

Financial and Performance Benefits:

Save **\$2.4 billion by 2050**

Implementing Hotspot-aware Cooling

Hyperscale data centres consume **50–100 MW per site**, equivalent to the power usage of a small city. Cooling accounts for roughly **15 – 30 MW** of that energy, making it a critical factor in both cost savings and energy efficiency.

Performance and ROI benefits

Hotspot Aware Cooling **prevents thermal throttling** enabling **higher chip throughput**. It also allows for chips to run at full speed longer, meaning **more computations per server**, leading to **higher ROI per rack**.




Saving on Energy Usage

By directing cooling to only where needed, this system **eliminates up to 3–6 MW of unnecessary energy use**. This can reduce cooling total power consumption **by 10–20%**, translating to **\$4 million in annual savings per site**. Across 10% of Microsoft Data Centre's (30 sites), this would save **\$2.4 billion by 2050**.

[View Calculations](#)

Financial Overview – Unlocking Billions in Efficiency



Our precision-focused coolant flow system offers a sustainable and **cost-effective** solution for Microsoft's data centres, delivering impressive financial and environmental impact.

Annual Energy Savings: Each data centre can reduce cooling power use by **3–6 MW**, translating to **\$2–5 million** in annual energy savings per site.

Deployment Potential: Scaling the system to **30** data centres (**10%** of Microsoft's data centres).

Total Savings by 2050: With deployment starting in 2030 and sustained through 2050, this approach unlocks **~\$2.4 billion** in total savings.

[View Calculations](#)

How We Will Unlock \$2.4 Billion in Savings, Calculated

1. Energy Consumption per Site

- 3 MW Operation:
 $3 \text{ MW} \times 24 \times 365 = 26,280 \text{ MWh/year}$
- 6 MW Operation:
 $6 \text{ MW} \times 24 \times 365 = 52,560 \text{ MWh/year}$

2. Cost Savings per Site

Assuming an electricity cost of \$0.10/kWh (\$100/MWh):

- 3 MW reduction \$2.6M saved per year
- 6 MW reduction \$5.2M saved per year

Thus, each site can save **\$2.6M – \$5.2M** annually.

3. Global Impact (10% of Microsoft Data Centres = 30 Sites)

- Average savings per site:
 $(2.6\text{M} + 5.2\text{M}) / 2 \approx 4\text{M}$
- Total annual savings (30 sites):
 $4\text{M} \times 30 = 120\text{M/year}$
- Long-Term Savings (2030–2050, 20 Years)

$$120\text{M} \times 20 = 2.4\text{B}$$

$$\begin{aligned} &3\text{--}6 \text{ MW/hr} \\ &\times \\ &24 \text{ hours} \\ &\times \\ &365 \text{ days} \\ &= \\ &26,280 \text{ -- } 52,560 \\ &\text{MWh/year} \end{aligned}$$

Assumed \$4m
per data centre,
per year

Implemented in
10% of data centres
(30 centres)

Over time from
2030 to 2050
(20 years)

- Backed By:**
- Microsoft Infra & Energy Data
 - ASHRAE & CMU Cooling Benchmarks
 - Glacierware CFD & Pressure Drop Models

**\$2.4
Billion
by 2050**

[View Calculations](#)

By **2050**, an estimated **\$2.4B** will be saved as a result of the implementation of Hotspot Aware Cooling, which can then lead to the development of **5 new data centres.**

Cost Per Centre: **\$500M** (Land, Construction, Infrastructure, Materials)

Total Capacity Added: **150MW Power** (30MW/Centre)

Estimated Total Revenue: **\$750M** (\$150M/Centre)

Estimated Total Servers: **50,000** (10,000/Centre at 3KW/Server)

Assumptions: 500M/datacentre cost, 30MW powerload (12MW cooling), 150M/year revenue per centre from cloud services, and linear reinvestment of energy savings over time. Does not account for inflation, operational costs, or demand fluctuations. Savings are net of cooling efficiency gains; new centres use renewable energy.

[**View Calculations**](#)

Implementing Hotspot Aware Cooling:

Saves up to **10.8 million litres** of water per year per site, equivalent to the daily water consumption of **36,000 people**.

Lowers CO2 emissions by up to **21,000 tons per year** per site, equivalent to **removing 4,565 gas powered cars** from roads annually.

Reduces chip failure rates and increases hardware lifespan, which **lowers chip degradation rate by up to 30%**, resulting in lower replacement costs and reduced e-waste.



[View Calculations](#)

Our Simulation Demo

This simulation illustrates how localized cooling can dynamically respond to workload variations across different regions of a chip. Additionally, it visualizes thermal dynamics across different chiplets in an AMD-style CPU architecture. As thermal load increases in the performance-intensive zone, the system intelligently redirects cooling capacity to that area without expending unnecessary energy on lower-load regions.

By mimicking real-world heat maps, this demo shows how precision coolant flow enables us to target hotspots efficiently, reducing energy waste and improving overall thermal performance.

[View Full Simulation](#)




Two-step R&D Pathway for Seamless Integration into Existing Systems

STEP 1

Hotspot Aware Cold Plates

Hotspot-aware cold plates are designed to **retrofit** into existing chip systems, allowing Microsoft to deploy advanced cooling **without modifying racks or facility infrastructure**. This enables a **fast, low-disruption upgrade path** that improves thermal performance and reduces energy use immediately. These cold plates intelligently direct fluid to chip hotspots, lowering temperatures without the need for full hardware replacement.



STEP 2

Hotspot Aware Direct-to-chip cooling (D2C)

As Microsoft plans for future generations of AI hardware, the next phase focuses on co-developing **Direct-to-Chip (D2C)** cooling solutions. These systems **integrate microfluidic cooling directly into the chip package itself**, enabling even more efficient heat removal with minimal thermal resistance. D2C cooling will be built into next-gen silicon designs, offering deeper energy savings and unlocking higher compute density over the long term.

1

Phase 1:

R&D Testing for Feasibility & Create CFD-Backed Designs for Cold Plates

Research system design using simulations and infrastructure-specific inputs.

Key Actions:

- **Conduct CFD simulations** tailored to Microsoft's chip thermal profiles
- **Validate** topology optimization using real chip power maps
- **Define** pressure, flow, and **integration constraints** based on Microsoft's cooling architecture
- Begin **developing manufacturing masks** for cold plate prototypes
- **Coordinate** with component vendors to assess material compatibility and **supply chain feasibility**

Test prototype cold plates in controlled but production-similar environments.

Key Actions:

- Select **2–3 pilot data centres** (ideally with varied climates and rack densities)
- **Install cold plates** in a limited number of test servers within each site
- **Instrument** with thermal sensors and flow/pressure monitoring
- **Track** real-time performance vs. baseline (straight-fin designs)
- **Document** installation process, integration challenges, and feedback from operations teams

2

Phase 2:

Prototype Deployment in Pilot Racks

3

Phase 3:

Live AI Workload Testing & Refinement

Expand deployment and use production AI workloads to refine system performance.

Key Actions:

- **Deploy** cold plates across entire server pods running real AI training/inference jobs
- **Monitor** thermal throttling, PUE, and pump power consumption
- **Adjust** flow rates and geometry designs based on workload performance
- Begin **creating a repeatable integration playbook** for data centre teams
- **Analyze financials** based on measured ROI and energy savings

4

Phase 4:

Standardized Roll-out of Cold Plates

Standardize cold plate integration by retrofitting existing AI chips across Microsoft's AI-scale infrastructure.

Key Actions:

- Make cold plates **standard** in all future high-density server deployments
- Prioritize integration in **AI/HPC clusters, GPU pods, and inference zones**
- **Define retrofit criteria** for existing deployments based on cost and thermal ROI
- **Automate** coolant monitoring and pump controls via data centre management systems
- **Monitor long-term reliability** and maintenance cycles to build confidence for broader use

5

Phase 5.

Scaled Deployment of Optimized Cold Plates Across Global Sites

All out optimized cold plates across Microsoft's hyperscale footprint.

Key Actions:

- **Integrate** with global data centre design and procurement teams
- **Adapt designs** for region-specific cooling systems (e.g., evaporative vs. chiller outside of the U.S.)
- **Ensure compatibility** with diverse server layouts, rack densities, and power envelopes
- **Train regional operations teams** in cold plate handling and monitoring
- Use telemetry to **track** energy savings and PUE shifts by site

Build the foundational mechanical, fluidic, and operational design layer to support future DTC cooling deployment.

Key Actions:

- **Define** rack and board level hardware requirements for DTC-ready servers
- **Design** fluid routing architecture to support individual chip-level cooling loops
- **Evaluate** materials, connectors, and sealing methods for leak-proof microfluidic interfaces
- **Assess** serviceability and maintenance plans for DTC-equipped systems
- Work with thermal engineers and silic to **ensure DTC compatibility with future chips**

6

Phase 6.

Design Framework for Direct-to-Chip Integration

7

Phase 7.

Pilot Deployment of Embedded Microfluidic Cooling

Deploy and test DTC prototypes in controlled environments to validate cooling performance and reliability.

Key Actions:

- **Partner with leading chip manufacturers** (e.g., Intel, AMD, NVIDIA) for early **DTC prototype** access
- Build bench-level test rigs to **evaluate** thermal performance, flow control, and fluid reliability
- **Integrate optimized DTC cooling** into testbed servers in controlled lab environments
- **Monitor** thermal performance, leakage, pump stability, and service impact
- **Compare** against cold plate benchmarks to define break-even points and ROI

Establish DTC cooling as a co-designed industry standard across chips, boards, and data centre infrastructure.

Key Actions:

- **Formalize** DTC-ready mechanical, electrical, and fluid interface standards
- **Co-develop** future chip packages with **optimized DTC support built-in**
- **Publish Microsoft DTC Deployment Guidelines** for internal and partner adoption
- Work with **OEMs** to produce **turnkey** DTC-capable server SKUs
- Launch early deployments in next-gen AI accelerator clusters
- **Sustained execution** across data centre generations can unlock over **\$2.4B** in cumulative energy savings by **2050**, based on projected thermal efficiency and deployment scale

8

Phase 8.

Standardization & Development for DTC Adoption at Scale

Challenges: Complex Design and Performance Risks

Challenge 1: Complex Design and Integration

Custom cooling systems require specialized design, CFD modelling, and precision flow control, making installation more complex than standard cooling systems.

Possible Solution:

Develop modular cooling components that simplify integration into existing data centres to reduce downtime and complexity.

Challenge 2: Maintenance and Performance Risks

Hotspot-aware cooling systems need frequent calibration to ensure coolant flow stays optimized for changing workloads. Without this, performance may drop.

Possible Solution:

Create an automated monitoring system that tracks heat patterns in real-time and adjusts coolant flow accordingly, reducing the need for manual calibration.



APPENDIX

Sources

Calculations

Technical Explanations