# HeatSync Case Study

## Optimizing Cold Plate Design for Direct-to-Chip GPU Cooling

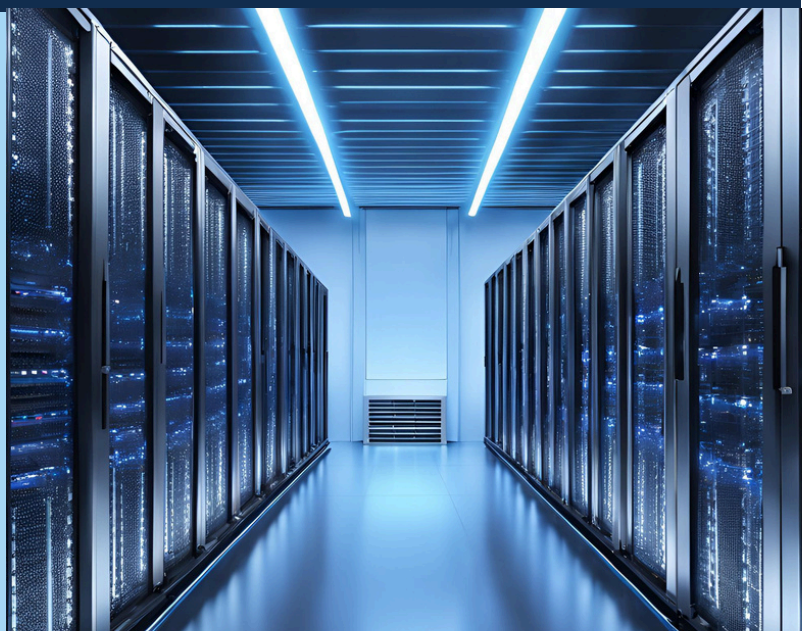

## OVERVIEW

A leading data center provider faced thermal management challenges with its **GPU-based servers,** where traditional air cooling proved insufficient. The objective was to develop a cold plate that could efficiently dissipate heat from GPUs, maintain low pressure drop, and ensure scalability for manufacturing. HeatSync applied a systematic design approach, integrating **CFD simulations**, **system-level analysis**, and rigorous **prototyping** to optimize cooling performance.

## CLIENT REQUIREMENTS

- **Efficient Heat Dissipation:** Design a dedicated cold plate for each of the four GPUs per server to effectively dissipate heat, ensuring optimal thermal performance and reliability.
- **Optimized Thermal Performance:** Maintain low GPU junction temperatures.
- **Hydraulic Efficiency:** Reduce pressure drop and improve coolant flow distribution across the system.
- **Manufacturability & Integration:** Ensure compatibility with existing form factors and refine production methods to enhance reliability.

## 1. INITIAL DESIGN

The design process considered form factor, material compatibility, and power load to establish a solid foundation. Copper and aluminum were evaluated for thermal performance, manufacturability, and cost-effectiveness.

The GPU junction temperature range was set per manufacturer specifications to ensure performance and longevity. These inputs shaped the cold plate's geometry, balancing thermal efficiency, hydraulic performance, and structural integrity. Preliminary calculations estimated dimensions and flow requirements, guiding detailed computational analysis.

## 2. CFD SIMULATION AND OPTIMIZATION

A high-fidelity CFD model was developed to refine the cold plate design, capturing thermal and hydraulic performance by monitoring GPU junction temperatures, temperature rise, and pressure drop. The model underwent rigorous validation, including grid independence verification, before being integrated into a Design of Experiments (DOE) framework for systematic microchannel and fin geometry optimization.

# Optimizing Cold Plate Design for Direct-to-Chip GPU Cooling

## 3. SYSTEM-LEVEL OPTIMIZATION

Once the cold plate design was refined, HeatSync focused on optimizing interconnections between the four cold plates per server to ensure maximum thermal efficiency. A system-level analysis was conducted to:

- Assess different coolant routing configurations for balanced flow distribution.
- Integrate both 1D and 3D simulations to accurately predict GPU junction temperatures under various operating conditions.
- Ensure hydraulic compatibility with the data center's cooling infrastructure.

## 4. PROTOTYPING AND TESTING

A prototype of the cold plate was manufactured and tested. Simulation-to-test correlation was within 6%, with a maximum junction temperature difference of 3.5°C.

## 5. RESULTS AND BENEFITS

- Reduced GPU temperatures by 20°C, significantly improving server reliability compared to air cooling.
- Achieved strong agreement between CFD predictions and experimental results.
- Optimized material selection and manufacturing methods for cost-effective production.
- Improved coolant flow distribution, minimizing pressure drop and enhancing system efficiency.

## 6. CONCLUSION

HeatSync successfully designed, optimized, and validated a cold plate for direct-to-chip GPU cooling, addressing thermal bottlenecks in high-performance data centers. By integrating CFD-driven design, system-level simulations, and rigorous testing, we delivered a high-efficiency cooling solution that enhances thermal performance, reliability, and manufacturability.

Contact HeatSync today to learn more about our advanced design, simulation, and testing solutions for direct-to-chip liquid cooling in data centers.

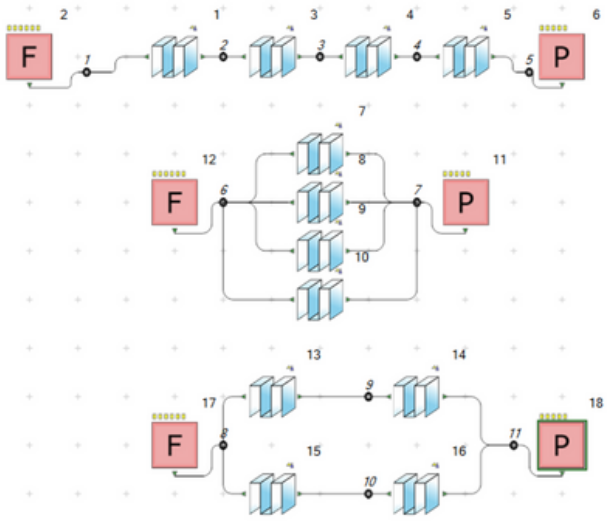


**HeatSync: Consortium of Thermal Management**

+1 (408) 650 2999 info@heat-sync.com

www.heat-sync.com Cupertino, CA 95014