Copper AM Cold Plate Optimization for High-Heat-Flux AI Data Centers

A liquid cooling OEM engineered a monolithic pure copper cold plate for a 900W-class AI accelerator using Cognitive Design, comparing CNC-machined and brazed multi-part construction against Green LPBF additive manufacturing. The optimized design delivers **27% higher thermal dissipation efficiency** at equivalent pumping power, with engineering lead time cut **from an estimated 10 weeks down to roughly 1.5 weeks**.

Engineering Challenge

Rising power density and shrinking packaging are pushing AI hardware toward a hard thermal ceiling. Thermal loads are climbing across the chip, the package, and the rack, and both conventional air cooling and standard liquid cooling are running out of headroom. Legacy multi-part cold plates, machined or skived then brazed to a cover plate, cannot resolve localized hot spots or scale channel geometry finely enough for next-generation accelerators.

1. **Heat Flux Mapping**: The chip's localized hot spots demand channel geometry tuned to its heat flux map, not a uniform design.
2. **Copper AM Manufacturability**: Pure copper LPBF requires a tight process window for wall thickness, powder evacuation, and laser parameters.
3. **Pressure Drop Budget**: Added heat exchange surface must still stay within a strict pumping power allowance.
4. **CFD Validation**: Complex lattice geometries are difficult to remesh manually in conventional CAE workflows.

The Solution: Cognitive Design Workflow

1. Design Space & Load Cases

Using Cognitive Design, the team imported the legacy cold plate geometry and defined Preservation Zones at the chip contact interface and the coolant manifold ports. Load cases included the accelerator's localized heat flux map and a maximum allowable pressure drop budget. Pure copper was carried forward as the material of record, targeting 98% relative density and 390 W/(m·K) as-built thermal conductivity, evaluated against a machined and brazed baseline, with the Green LPBF process window built in from the first iteration.

2. Generative Exploration

The Design Explorer generated **60+ variants in parallel** across gyroid and diamond TPMS families alongside jet impingement microchannel geometries, compared to the 2 to 3 concepts a conventional workflow allows. Variants were scored on **thermal resistance, pressure drop, and manufacturability**. The team selected a gyroid-based jet impingement hybrid for its balance of dissipation and pumping power.

3. Performance & Manufacturability (MDD & SDD)

Manufacturing-Driven Design constraints for pure copper LPBF, minimum wall thickness, powder evacuation paths, and laser parameter limits, were applied directly to the selected geometry. Simulation-Driven Design ran **meshless thermal-fluidic analysis in real time**, letting the team tune local wall thickness near hot spots down to 105 micron, without the meshing tax typical lattice geometries impose.

4. Validation / Final FEA

The team confirmed the design against the full load case: junction-to-coolant thermal resistance, hot spot temperature delta, and pressure drop against the allowable budget, confirming **27% higher dissipation efficiency at equivalent pumping power**. The optimized geometry exported directly as a simulation-ready volumetric mesh, enabling CFD validation in their standard CAE environment without manually remeshing the lattice.

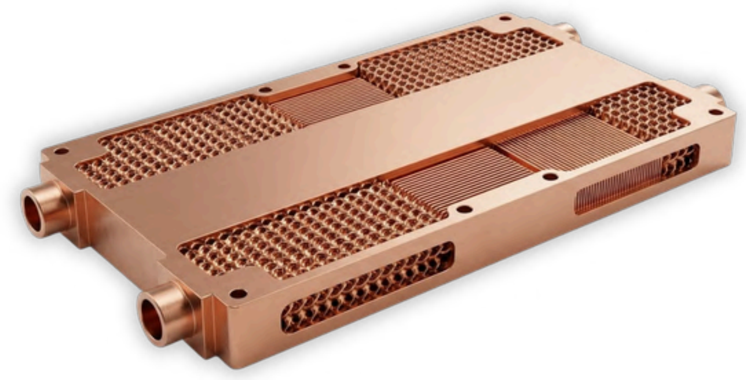

Optimized cooling plate with gyroid TPMS

Key Metrics Comparison

Metric	Legacy (Copper)	Optimized (Copper)	Improvement
Thermal Dissipation Efficiency (equivalent	Baseline	27% higher	27% improvement
Relative Density	100% (wrought reference)	98%	Near-bulk density
As-Built Thermal Conductivity	400 W/(m·K) (wrought reference)	390 W/(m·K)	Near-bulk performance
Minimum Wall/Fin Thickness	300 to 500 micron	105 micron	3 to 4x finer feature resolution