Pure copper, printed to the geometries AI chip cooling actually needs.

Redirecting MoldJet Dominant toward heat sinks and cold plates for semiconductor and data-center thermal management — where copper's conductivity matters more than its manufacturability problems.

Target: **hyperscale data centers** · **rack OEMs** · **chip packaging**

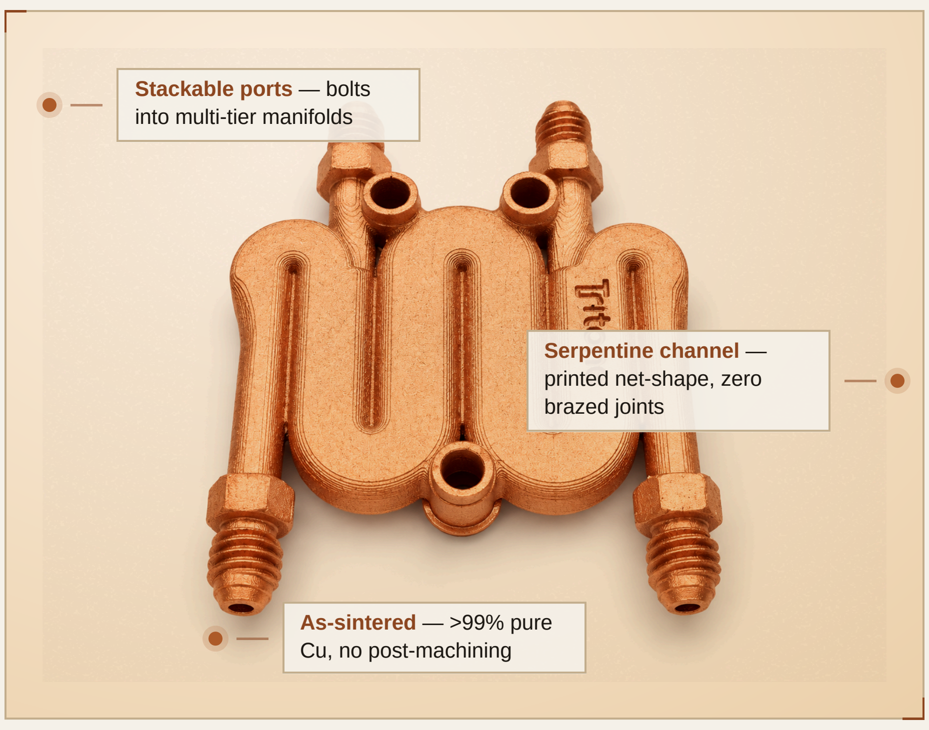

Direct-to-chip cold plate, stackable module · MoldJet Dominant — Fig. 1

$6.65B → $29.5B
Data-center liquid-cooling market, 2025–2033 (20% CAGR)

76%
of AI servers deployed by end-2026 expected to be liquid-cooled — Goldman Sachs

100–150 kW
per-rack thermal loads on today's AI accelerators — air cooling no longer viable

The opportunity

AI accelerators (NVIDIA Blackwell/Rubin, custom ASICs) now dissipate 1,500W+ per chip, pushing racks past the 40 kW ceiling where air cooling works. Direct-to-chip liquid cooling — cold plates and heat sinks in copper — is becoming the default architecture for hyperscale and HPC deployments. Major buyers are now qualifying next-gen cold-plate suppliers, and the qualification window is open now.

Why now
Rack thermal density has outrun what skived and brazed copper plates can economically deliver at volume.

Why copper
Highest practical thermal conductivity of any AM-compatible metal — but the hardest to print well.

Why MoldJet wins where others struggle

Processing a soft metal like copper is hard. Machining is limited, and flat ribbed plates made by skiving don't integrate cleanly into end-use components. In additive, laser powder-bed fusion fights copper's reflective nature — vendors compensate with exotic green/blue lasers, substitute Cu alloys, or accept degraded material properties. MoldJet's powder-free, mold-based process sidesteps this entirely: it never melts copper with a laser, so it never fights reflectivity in the first place.

The MoldJet approach
Printed net-shape, pure copper

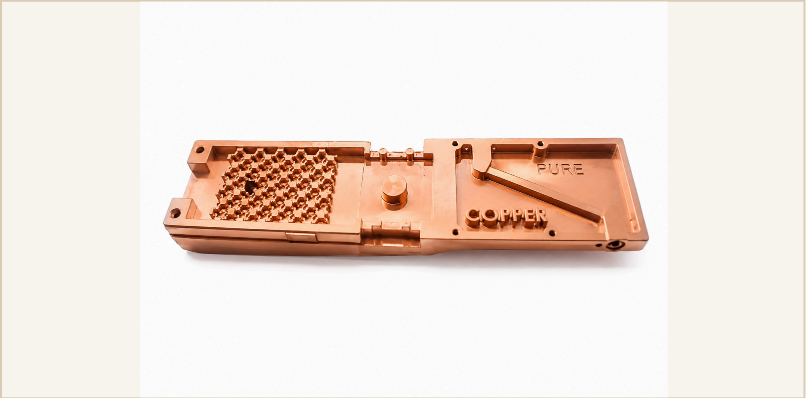

Internal lattices, tunnels and micro-channels printed directly — full density, no depowdering, no assembly.

VS

The machined incumbent
Skived, brazed, assembled

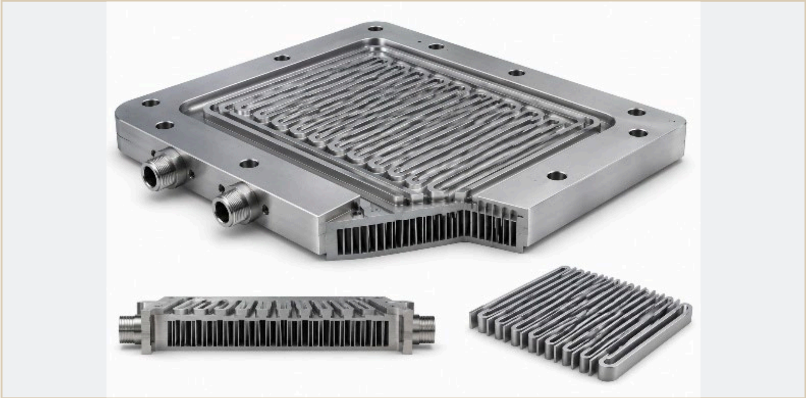

Today's standard: machined channels and skived fins, brazed from multiple pieces at real labor cost.

Tritone's benefits

Design freedom, no compromise
Micro-channels, internal tunnels and high-surface-area lattices that machining can't cut and laser-AM struggles to densify.

Cost and labor reduction
Infinite complexity at no extra cost, with assembly consolidation replacing multi-piece brazed builds.

Production-volume economics
MoldJet Dominant is built for scale manufacturing, matched to how rack OEMs actually buy.

No powder handling
No depowdering means clean inner-channel geometry — plus the safety and facility-cost advantage copper demands.

Proven precision and repeatability
Production-qualified tolerances and repeatable part-to-part performance across builds — not a lab curiosity.

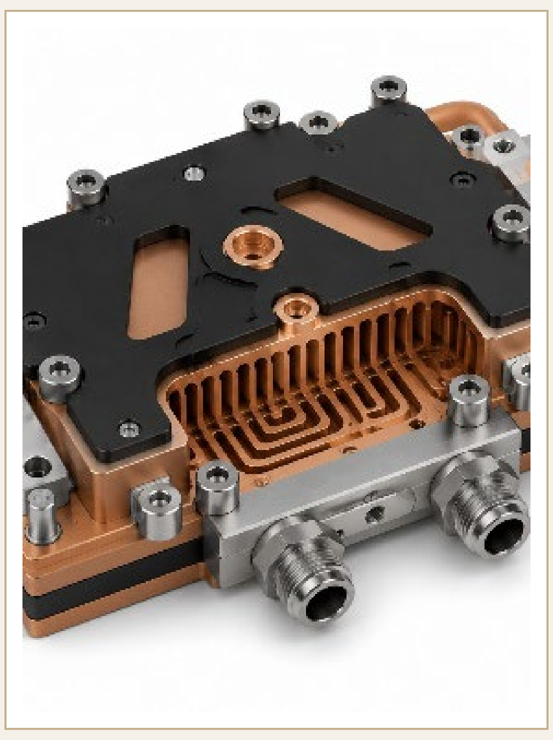

Market reference — the direct-to-chip cold-plate form factor the market has standardized today, machined and assembled from multiple pieces.

Target buyers

Hyperscale & colocation
Data-center operators qualifying next-generation cold-plate suppliers ahead of 2026 liquid-cooling rollouts.

Rack OEMs & integrators
Server/rack OEMs and cold-plate or thermal-module suppliers building direct-to-chip liquid-cooling assemblies.

Chip packaging & test
Semiconductor packaging and test houses building chip-level thermal solutions for GPU and custom-silicon programs.

Secondary verticals: EV power electronics · aerospace power systems — same pure-copper, complex-geometry case applies.