

What If the World had Fewer AI Data Centers?



New analysis from Silicon Valley Unicorn, Frote Systems, finds that better cooling could let AI data centers produce up to 40% more tokens output from the same power, potentially easing pressure on power grids, and water supplies from the communities hosting them.

San Jose, CALIFORNIA — October 12, 2026— The race to build AI data centers is colliding with a hard limit: there isn't enough power to go around. New analysis from Frote Systems, explained in their White Paper '*Cooling Is the Next Revolution*', argues that one of the biggest answers is hiding in plain sight. By optimizing the Thermal Stack and running AI chips cooler, the research finds data centers could produce over 40% more AI output from the same electricity. That could mean fewer new facilities are needed to meet the world's demand for AI.

Every AI answer, image or line of code is made of "tokens," the basic unit of intelligence AI systems produce and sell. AI data centers, increasingly called AI Factories, exist to make tokens. Their output is limited not only by chips or money, but increasingly by electricity. AI Factories face years-long waits for new utility connections, regulators are capping consumption, and communities are pushing back on the power and water these facilities draw.

Frote Systems' paper shows that heat is quietly eroding how much useful work each watt of power delivers. As an AI chip runs hotter, more of its power leaks away as waste heat instead of doing computation. That wasted power roughly doubles with every 10°C rise in chip temperature. In the paper's analysis, a 10°C cooler chip can produce about 15% more tokens per watt of power.

"The AI industry wants to build many more data centers. We think it should also be asking how much more it can get from the ones it already has," said Seshu Madhavapeddy, Founder and CEO, of Frote Systems. "A power-constrained world can't afford to waste a single watt of power. By treating cooling as a performance system rather than plumbing, the same building, the same chips and the same available power can deliver dramatically more tokens."



The paper traces heat's full journey through the Thermal Stack, from the silicon chip through its packaging, the thermal interface material that bond it to the coldplate, the coolant, and out to the cooling towers. Its key findings:

- **Over 40% more tokens output from the same power.** A well-engineered cooling system can raise an AI Factory's tokens per watt by over 40%. Because that extra output costs nothing more to produce, it flows straight to the bottom line.
- **Cooler chips last longer and make fewer mistakes.** Running chips at 65°C instead of 95°C can double GPU lifespan and reduce the risk of silent calculation errors.
- **Coldplate manufacturing has been a bottleneck.** Conventional coldplates are machined from solid copper, which limits their designs. Frore's **LiquidJet** coldplates are made with semiconductor manufacturing processes, enabling 3D ultra short-loop multi-stage structures customized to the specific power map of each GPU, that traditional machining methods cannot build. On its own, Frore Systems' LiquidJet coldplate can increase tokens per watt by 25%. The newest variant, **LiquidJet Diamond**, which integrates diamond to spread heat from a chip's extreme hotspots, can deliver a 35% increase.

The implications reach well beyond engineering. Applied across a Hyperscaler's building plans, a 40% gain could mean that seven new AI data centers deliver what would otherwise have taken ten. That is three fewer multibillion-dollar construction projects, three fewer grid connections in strained regions, and three fewer communities absorbing the power and water footprint of a hyperscale facility.



"Thermal Stack design has always been treated as a facilities line item," Madhavapeddy added. "But at the current scale of AI factory buildouts, it's a strategic boardroom decision. It shapes how many facilities a hyperscaler must build, how fast it can bring AI capacity online, and how it answers to the communities living around it."

Availability

Frore Systems White Paper 'Cooling Is the Next Revolution' is available to download free at [LINK]. An upcoming companion 12-part blog series, being shared over coming weeks explains the paper's findings in plain language, at [LINK].

About Frore Systems

Frore Systems is a pioneer in advanced thermal technologies that unleash performance across data centers and edge devices. The company's flagship solutions include **AirJet®**, the world's first solid-state active air-cooling chip used in consumer, industrial, and IoT markets delivering higher performance in ultra-compact, silent, light, dustproof and water-resistant edge devices; **LiquidJet®**, a 3D ultra short-loop multi-stage liquid cooling coldplate for data centers and **LiquidJet® Diamond** with an integrated diamond spreader, delivering higher GPU performance, AI token throughput, improved PUE and reduced TCO; and **LiquidJet® Nexus**, a lightweight integrated coldplate system that integrates multiple LiquidJets and eliminates all hoses, connectors and manifolds enabling ½U compute trays. Frore Systems' patented cooling technologies are integrated into products from major OEMs and system builders worldwide. Headquartered in Silicon Valley, with manufacturing operations in Taiwan, Frore Systems is redefining thermal architecture for the AI era. For more information, visit: www.froresystems.com.

Media Contact for CEO Interviews and more information:

Sue Ryan - VP Marketing, Frore Systems

sue@froresystems.com

Cell: +1 314 914 5008