



2025 EDITION

Your Partner for Liquid Cooling in AI Data Centers

Smarter Liquid Cooling, Vertically
Integrated with Flex

Table of Contents

- Introduction 3
- Industry Momentum 4
- Company Overview 5
- Technology Overview 6
- Rack-Level Performance 8
- Meeting AI Demands 10
- Global Supply Chain 11
- Lifecycle Support 12
- Conclusion 13



Introduction

The data center landscape is evolving

In 2025, the data center industry is at a pivotal moment. The rise of AI and machine learning has fundamentally changed compute infrastructure needs. Workloads powered by high-density GPUs and custom accelerators are pushing infrastructure to its physical and operational limits. Traditional air cooling, once sufficient for legacy systems, is now falling short as processor thermal design power (TDP) scales to unprecedented levels. According to Uptime Institute webinar titled “How Far Can We Go With Air?”, server TDP has increased by more than fivefold over the past 15 years.¹ Air cooling systems, which once managed 200–300W CPUs effectively, are now incapable of supporting the 700–2,700W chips dominating today’s AI ecosystems.²

“Air cooling systems, which once managed 200–300W CPUs effectively, are now incapable of supporting the 700–2,700W chips dominating today’s AI ecosystems.”

¹ Uptime Institute. Lawrence, Andy. Presentation “How Far Can We Go With Air?” March 25, 2025. Webinar.

² <https://flex.com/resources/the-future-of-data-centers-demands-advanced-cooling>

Industry Momentum

Liquid cooling market acceleration


According to Dell'Oro Group's January 2025 report, the Data Center Physical Infrastructure (DCPI) market is expected to reach \$61 billion in annual revenue by 2029, driven by a 14% compound annual growth rate (CAGR) from 2024 onward.³ These revised estimates reflect a dramatic surge in real-world demand, where shipments of AI-optimized hardware far exceeded projections. AI is no longer just a Tier 1 cloud challenge—demand is now expanding to Tier 2 providers, government-backed initiatives, and colocation environments.

As workloads intensify, Dell'Oro forecasts a steep rise in average rack power density from 15 kW to 60 kW, with leading deployments already reaching 120 kW per rack. Some forward-looking designs are preparing for 1 MW racks and beyond. This isn't theoretical: NVIDIA's latest reference architecture, unveiled at their annual NVIDIA GTC 2025 conference, supports 600 kW racks, while Google is actively planning 1 MW liquid-cooled designs.⁴

How are operators cooling their dense IT?

Even on the densest racks, perimeter cooling dominates

How do you currently cool your highest density cabinets? (n=572)



uptime
INTELLIGENCE

Dell'Oro identifies single-phase, direct-to-chip cold plates as the leading technology in liquid cooling—delivering the optimal balance of performance, simplicity, and serviceability. This aligns with findings from the Uptime Institute, which report that water-cooled cold plates are the most widely deployed liquid cooling solution among data center operators.⁵ These principles form the foundation of JetCool's liquid cooling portfolio.

³ <https://www.delloro.com/news/delloro-group-raises-market-forecast-for-data-center-liquid-cooling-and-rack-power-distribution/>


⁴ <https://www.datacenterdynamics.com/en/news/nvidias-rubin-ultra-nvl576-rack-expected-to-be-600kw-coming-second-half-of-2027/>

⁵ Uptime Institute. Lawrence, Andy. Presentation "How Far Can We Go With Air?" March 25, 2025. Webinar.


Company Overview

JetCool, A Flex Company

As a Flex Company, JetCool is bringing its industry-leading liquid cooling technology to a global scale. By integrating into Flex's extensive manufacturing ecosystem, JetCool now designs, builds, and delivers complete liquid-cooled rack and power solutions entirely in-house.



Solving power, heat, and scale challenges



■ Manufacturing services ■ Critical power products ■ Embedded power products ■ Cooling products

Technology Overview

SmartPlate™: Precision cooling for modern compute

JetCool's SmartPlate™ is an advanced cold plate that uses patented microconvective cooling® to cool the processor's hot spots directly—right where heat is generated. Unlike traditional approaches that rely on spreading heat across the entire surface, SmartPlate delivers superior thermal performance through a simpler, more efficient, and highly reliable design. It's precision cooling, redefined for next-gen compute.

SmartPlate cold plates are uniquely architected to perform best with warm coolant inlet temperatures, up to 60°C. This capability unlocks powerful operational benefits. First, it allows for year-round use of free cooling in most environments, eliminating the need for energy-intensive chillers and evaporative coolers.


JetCool's Microconvective Cooling Architecture



Second, it supports heat reuse systems that can redirect waste heat to nearby industrial, agricultural, or municipal applications. Uptime Institute has identified heat reuse as a key pillar of next-generation cooling strategy, yet very few cooling systems are optimized for such high inlet temperatures—SmartPlate is the exception.⁶

SmartPlate outperforms microchannel cold plates by 25%

Unlike microchannel solutions that often require 25-micron filtration and high-pressure pumps, SmartPlate operates at lower flow rates and pressures, reducing both CapEx and OpEx. It's available for a wide range of processors, including CPUs, GPUs, and AI accelerators from NVIDIA, AMD, and Intel. This allows operators to consolidate their thermal strategy under a single, scalable, and serviceable technology stack.



25%
Better Thermal
Performance


50°C
Coolant Inlet
Temperatures

⁶ Uptime Institute. Lawrence, Andy. Presentation "How Far Can We Go With Air?" March 25, 2025. Webinar.


Technology Overview

Proven performance for market-leading AI accelerators

In third-party testing, the JetCool SmartPlate for the NVIDIA GB300 demonstrated a thermal resistance of just 0.014 °C/W, including TIM, with a pressure drop of only 7.7 psi and 1.5 LPM/kW flow rate, validating its high-performance thermal design for next-generation AI accelerators.



By lowering thermal resistance, microconvective cooling technology minimizes the temperature difference between the coolant and GPU. This gives data centers two key options: increase coolant inlet temperatures to cut facility energy use or maximize power density to cool even higher-wattage processors. As data centers face the mounting challenge of balancing power, heat and scale, this technology affords them the luxury of a tenable, tailorable solution to best meet their escalating performance goals without compromising efficiency.



0.014 °C/W
Thermal Resistance

The industry-leading thermal performance further validates the capability of single-phase direct-to-chip cold plates using microjet impingement, supporting the highest power densities on the market today and providing headroom for even more demanding workloads.

Rack-Level Performance

SmartSense coolant distribution unit (CDU): designed for AI, ready for scale

Cooling can't keep up with compute—until now

A critical challenge identified in today's AI-driven data centers: cooling equipment isn't scaling as quickly as compute density. As GPUs become more powerful and rack power densities soar, operators are often caught between two difficult options—overbuilding cooling infrastructure, which risks stranding capacity and increasing cost, or underbuilding, which can result in thermal bottlenecks or premature obsolescence.

High-performance cooling in a compact 6U footprint

JetCool removes that trade-off with the SmartSense CDU, a high-performance, in-rack liquid-to-liquid coolant distribution unit (CDU) designed for AI computing. Despite its compact 6U form factor, SmartSense delivers the cooling capacity of much larger CDUs, supporting up to 300 kW per rack. This design allows operators to stack multiple CDUs within a single rack, reclaiming space and enabling higher compute density across the data hall.



SmartSense CDU

By supporting neighboring racks or entire rows—scaling to over 2 MW at the row level—SmartSense provides the thermal headroom needed to support even the most demanding workloads. Furthermore, as endorsed by ASHRAE, deploying smaller, dedicated CDUs optimizes CapEx and reduces risks associated with large cooling systems, ensuring a more responsive and efficient approach.⁷ It enables operators to maximize footprint efficiency while meeting both current and future cooling requirements.

PERFORMANCE METRICS	BOOST MODE*	REDUNDANT MODE
Cooling capacity	300 kW	160 kW
Approach temp. difference	7°C	4°C
Flow rate	400 LPM	280 LPM
Pressure head	22 psi	28 psi
Operating temperature range	10°C to 70°C	10°C to 70°C

*Boost Mode is defined as the SmartSense CDU operating with all pumps running.

7 ASHRAE TC 9.9, *The CDUs Critical Role of TCS and FWS Isolation in Cold Plate Deployments (March 2025)*, contributors: Tim Shedd et al., <https://tpc.ashrae.org/Documents?cmtKey=fd4a4ee6-96a3-4f61-8b85-43418dfa988d>

Rack-Level Performance

Modular, scalable, and future-ready

JetCool's SmartSense CDU is purpose-built to grow with customers as they begin scaling their AI infrastructure. Designed for flexibility, it creates a modular, end-to-end cooling ecosystem when paired with JetCool's SmartPlate™ cold plates adapting seamlessly as power and thermal requirements evolve.

With onboard telemetry, SmartSense continuously monitors pressure, temperature, and flow, offering predictive cooling insights and real-time responsiveness. This is particularly critical for generative AI workloads, which can produce sudden and unpredictable power spikes. Uptime warns in their webinar "How Far Can We Go With Air?" that these fluctuations can overwhelm traditional chiller systems—SmartSense mitigates that risk by dynamically adjusting coolant flow based on chip-level thermal data.




Built for reliability, and serviceability

To meet the demands of 24/7 mission-critical environments, SmartSense is built with redundant pumps, field-replaceable components, and remote diagnostics capabilities. It's backed by Flex-certified on-site servicing, ensuring reliability from the moment of deployment through the full lifecycle of the installation.

Together, JetCool's SmartPlate cold plates and SmartSense CDU provide a compact, scalable, and serviceable platform that enables operators to scale with confidence—without compromising on performance, density, or operational flexibility.


Streamline vertical integration for AI deployments



Meeting the Demands of AI Infrastructure

Techno-economic study of a 40 MW AI center in Virginia

In a case study conducted with a hyperscaler, JetCool modeled a 40 MW AI data center deployment using its SmartPlate™ cold plate technology. Compared to microchannel systems, this warm-liquid approach cuts annual cooling costs by 47%—from \$5.2 million to \$2.7 million—while maintaining high performance and, in most climates, eliminating the need for evaporative cooling.



A key driver of these savings is JetCool's ability to operate with warm coolant temperatures (ASHRAE 42°C and above), reducing both electricity and water consumption. The result: greater GPU utilization, lower operational costs, and improved sustainability for large-scale AI deployments.




Global Manufacturing and Supply Chain Localization

Resilient infrastructure starts with resilient manufacturing

Flex has built a globally distributed, regionally optimized supply chain to support faster delivery, supply continuity, and manufacturing agility. JetCool's liquid cooling solutions are manufactured at Flex sites in key regions including Penang, Malaysia; Guadalajara, Mexico; Tczew, Poland; and multiple locations across the United States. This regional footprint enables us to serve customers across North America, EMEA, and Asia-Pacific with shorter lead times and responsive, scalable production.

Flex's Global Support



This model also aligns with the growing demand from hyperscalers and government buyers for locally sourced and compliant infrastructure.

By manufacturing in-region and maintaining shared quality systems across all facilities, Flex ensures consistent product quality while reducing logistics costs, emissions, and exposure to global disruptions such as port congestion or export restrictions.

24x7
Support

30
Global Sites

14K
Field Service
Engineers

124+
Countries

Lifecycle support

True whole product solution

JetCool is more than a cooling technology provider—we are a long-term partner. With global service and support through Flex, our customers benefit from comprehensive lifecycle services that include deployment assistance, remote telemetry, maintenance, RMA logistics, and upgrade planning. With decades of experience supporting customers with liquid cooling around the world, Flex brings deep expertise in system design, deployment, and long-term operational management that will support SmartSense CDU customers.


Whether customers are deploying a rack or full-scale production, JetCool is trusted by hyperscalers, OEMs and enterprise customers to provide support from the first rack to full production scale.

The Flex platform

Throughout the value chain

Services to support the full product lifecycle

- **Manufacturing services**
- **Value-added fulfillment services**
- **Digital circular economy solutions**



flex®

Conclusion

The platform for the AI era

The transition to liquid cooling is no longer in the future—it's happening now. Driven by rapidly escalating power requirements, sustainability goals, and localized infrastructure mandates, operators need a trusted partner to migrate data centers efficiently and effectively.

JetCool's SmartPlate and SmartSense CDU are built for this new era—delivering unmatched thermal performance, rapid ROI, and simplified deployment. With active manufacturing in the U.S., Mexico, and Malaysia, and the ability to support everything from cold plate retrofits to full liquid-to-die deployments, JetCool offers the most resilient, efficient, and forward-compatible cooling platform on the market. Now is the time to uncomplicate liquid cooling—and scale with confidence.

Infrastructure Integration: Cold Plates, CDUs, Manifolds, Quick Disconnects



SmartPlate Cold Plates & SmartLid

Single-phase, DLC microconvective cold plates and cooling modules that target hot spots directly.



SmartSense CDU

In-rack, liquid-to-liquid 6U coolant distribution unit supporting 300 kW racks and 2 MW+ rows.



Rack Manifolds

Liquid rack manifolds configurable for standard 19" rack, ORV3 or custom racks with multiple form-factors, sizes, fitting and blind-mate connections.



Quick Disconnects

Diameter range 2 - 150mm, bidirectional self-seal, auto pressure relief with integration across node, cabinet, and manifold applications.



Contact Us



305 Foster Street, Suite 201
Littleton, MA 01460
+1 (978) 449 - 4600
user@example.com