INDUSTRY ANALYSIS · MARCH 2026

TWO-PHASE v SINGLE-PHASE.

Why two-phase direct-to-chip cooling reduces chiller load, simplifies the white space, and lowers operating cost for AI infrastructure

Lower Chiller Load	Simpler White Space	Lower Five-Year TCO
Fewer high-temperature chillers required	No separate technology cooling system	35% opex reduction, 12% TCO reduction

PUBLISHER
DARKNX LTD

FOCUS
Cooling Architecture

WEBSITE
darknx.com

OVERVIEW

Single-Phase vs Two-Phase: The Plant-Level Difference

This paper examines why two-phase direct-to-chip cooling is emerging as the more durable architecture for high-density AI environments. It focuses on lower chiller burden, simpler white-space architecture, and a lower five-year operating profile when compared with a single-phase design.

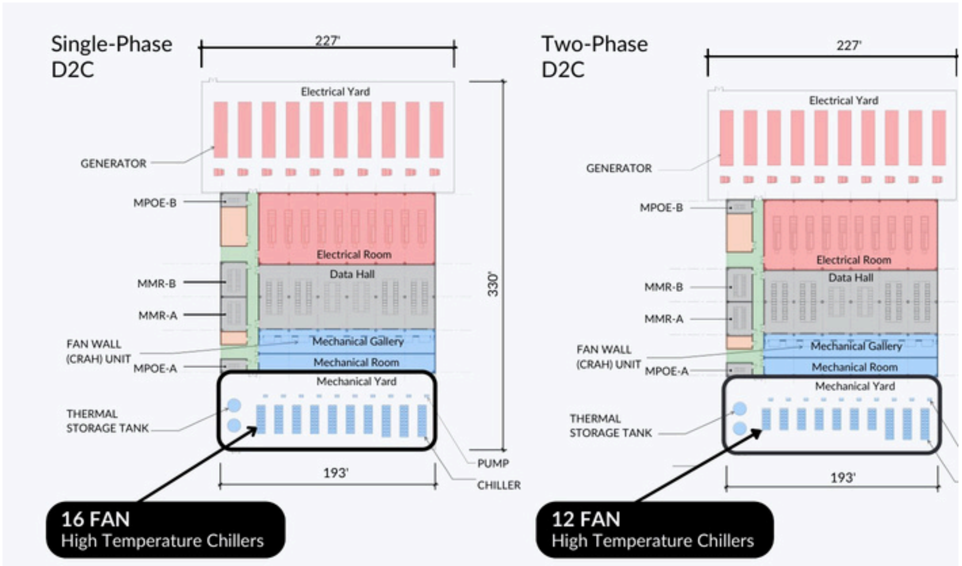

Jacobs reference design comparison: single-phase D2C requires 16 fan high-temperature chillers; two-phase D2C requires 12 — a reduction of four units and no separate technology cooling system.

16 → 12	-35%	-12%
High-temperature chillers reduced in Jacobs reference design	Reduction in annual operating cost (Austin reference case)	Reduction in five-year total cost of ownership

EXECUTIVE SUMMARY

The Architecture With Lower Long-Run Complexity

Dense AI racks have turned liquid cooling from a supporting system into a primary design decision. The relevant comparison is no longer only thermal feasibility at the chip — it is whether the selected architecture simplifies or multiplies the supporting plant over the life of the facility.

In the Jacobs 10 MW reference design, two-phase direct-to-chip cooling held chip-to-chiller capital cost essentially level with single-phase, while reducing annual operating cost by 35 percent and five-year total cost of ownership by 12 percent. The same comparison reduced the count of high-temperature chillers from 16 to 12 and removed the need for a separate technology cooling system.

A 2026 experimental study of retrofitted servers found CPU temperature fell from 93°C under air cooling to 56°C with two-phase hybrid cooling at 25°C inlet under full load, while the liquid loop captured up to 90 percent of the heat load at 45°C inlet air.

WHAT THE ARCHITECTURE NOW REWARDS	WHY IT MATTERS
Comparable capital entry point	Jacobs shows chip-to-chiller capital cost is effectively unchanged, making life-cycle performance the decision variable.
Lower chiller burden	Jacobs reduces high-temperature chillers from 16 to 12, lowering plant burden and ongoing energy consumption.
Integrated white-space loop	Removing a separate technology cooling system cuts interfaces, piping, and coordination points.
Lower five-year cost profile	When capital entry is comparable, a 12 percent lower five-year TCO becomes materially relevant.

WHY THIS MATTERS NOW

AI facilities are being designed around sustained rack power density, warmer water strategies, and faster hardware refresh cycles. In that environment, a cooling architecture that looks manageable at deployment can still become burdensome if it adds loops, chillers, or control dependencies as density rises.

THE COOLING CONSTRAINT HAS MOVED INTO THE WHITE SPACE

Why Two-Phase Changes the Baseline

For much of the last cycle, cooling decisions could be deferred until late in design because the penalty for extra loop complexity was manageable. Higher density has reduced that tolerance. Each additional loop, pump set, manifold path, or coordination point now carries more operational and commissioning weight.

In the Jacobs reference design, single-phase direct-to-chip cooling still depends on a separate technology cooling system and a higher number of high-temperature chillers. The two-phase configuration integrates manifolds into the MR250 and lowers the facility chiller count.

> AN INTERPRETATION OF COOLING CAPACITY
>
> In dense AI environments, cooling capacity means more than heat removal. It means what can be delivered, maintained, and expanded without adding avoidable thermal infrastructure around the rack.

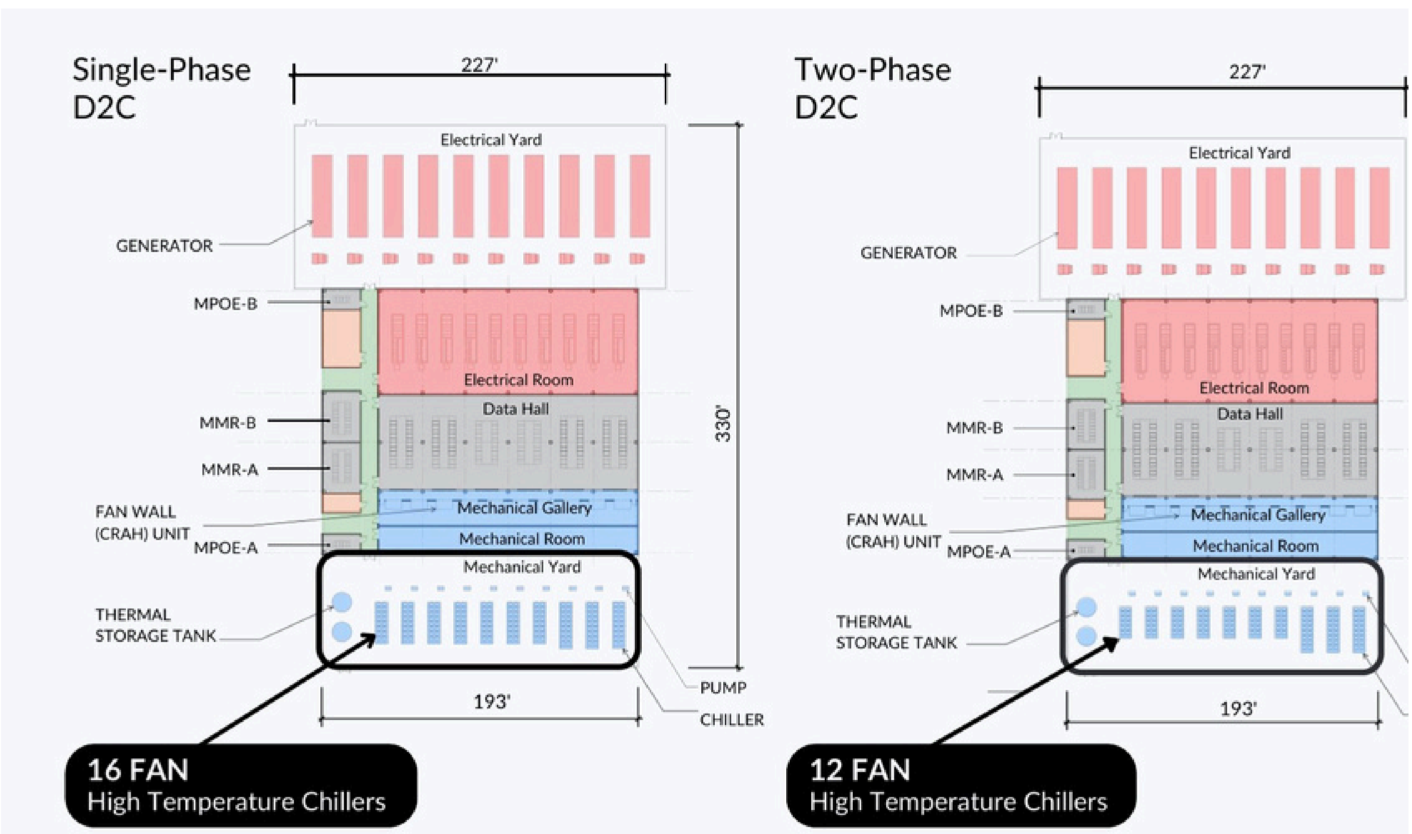

Jacobs reference design comparison showing lower high-temperature chiller count in the two-phase configuration and a simpler path through the white space.

SIMPLICITY IN THE WHITE SPACE

What Operators Need to See

The practical difference appears in how much thermal support must sit between the plant and the rack. Two-phase cooling consolidates more of that function inside the cooling architecture itself. That simplifies piping and controls in the Jacobs design, and experimental work published in 2026 further suggests that as inlet air temperature rises, the two-phase loop takes on a larger share of the heat load.

SINGLE-PHASE PATTERN	TWO-PHASE PATTERN	OPERATIONAL IMPLICATION
Separate technology cooling system	Integrated manifold architecture	Fewer thermal interfaces to build, monitor, and commission.
16 high-temp chillers (Jacobs)	12 high-temp chillers (Jacobs)	Lower plant burden and lower energy consumption.
More secondary piping and loop coordination	Cleaner white-space routing	Lower installation and maintenance complexity.
More air-side burden as temperatures rise	Higher liquid heat capture at warmer conditions	Reduces dependence on the air side as density rises.

WHAT DO OPERATORS NEED TO SEE?

Buyers and operators should evaluate cooling architecture as a system decision with measurable consequences for cost, visibility, and repeatability.

OPERATOR QUESTION	WHY IT NOW MATTERS
Has the design removed unnecessary loops?	Jacobs removes a separate TCS and cuts chiller count, reducing coordination points and failure exposure.
Can the plant support density without adding more mechanical layers?	Density is more durable when the supporting plant does not scale complexity at the same pace.
Will warmer operating conditions remain safe?	Peer-reviewed studies show lower CPU temperature, higher liquid heat capture, and stable operation at warmer inlet or water temperatures.
Does the architecture stay coherent as future hardware changes?	Thermal decisions now need to outlast one hardware generation.

LOWER OPERATING COST ACROSS THE COOLING LIFE CYCLE

A Five-Year Cost View

The Jacobs comparison makes the economic case unusually clear. Chip-to-chiller capital cost remains effectively comparable between the two designs, but annual operating cost falls materially in the two-phase configuration.

In the Austin reference case, annual operating cost declines from $1,043,799 in the single-phase design to $678,640 in the two-phase design, while five-year total cost of ownership declines from $15,604,571 to $13,774,735 — a 35 percent reduction in annual operating cost and a 12 percent reduction in five-year TCO.

	Single-Phase D2C	Two-Phase D2C	Savings (%)
CapEx*	$10,385,576	$10,381,534	0%
Annual OpEx	$1,043,799	$678,640	-35%
5-Year TCO**	$15,604,571	$13,774,735	-12%

*CapEx figures include chillers for the liquid cooling FWS, CDUs, TCS Loop, fluid, FWS extension, cold plates, manifolds, and QDs. It does not include FWS piping from chiller to white space, FWS pumps, and other minor supporting mechanical infrastructure.
**TCO analysis is based on a data center in Austin, TX. Detailed energy consumption for Austin, TX and other regions is available on request. Additionally, the TCO analysis includes maintenance and piping material costs from other third-party sources.

Jacobs reference case for a 10 MW data-centre design in Austin, Texas. Capital cost remains effectively comparable, while operating cost and five-year total cost of ownership improve in the two-phase configuration.

Those economics are consistent with broader technical findings. The cost delta is not being created by a higher capital hurdle — it is being created by **a simpler thermal system that can operate efficiently at warmer conditions over time.**

HOW DARKNX APPLIES TWO-PHASE COOLING

Infrastructure Architecture, Not Product Substitution

DARKNX approaches two-phase direct-to-chip cooling as an infrastructure architecture rather than a product substitution. The relevant question is not only whether two-phase performs better thermally, but where it simplifies the plant, protects density headroom, and reduces long-run operating burden.

Because DARKNX evaluates cooling architecture against density targets, plant loading, controls logic, and deployment sequence, published study data can be translated into site-level decisions rather than treated as isolated laboratory results.

CAPABILITY	HOW DARKNX APPLIES IT	WHY IT MATTERS COMMERCIALLY
Early cooling selection	Two-phase evaluated alongside power density, plant loading, and site phasing at concept stage.	Cooling decisions affect layout and delivery before procurement begins.
Integrated plant & white-space design	Mechanical, electrical, and controls choices are coordinated as one system.	Two-phase benefits are preserved when interfaces are not left to field reconciliation.
Repeatable module planning	Cooling architecture carried into modular or phased deployments without reopening the entire design.	Higher density capacity becomes easier to replicate.
Commissioning discipline	Validation includes manifold behaviour, control sequence, and operating envelope — not only hardware installation.	Performance claims become easier to trust in operation.
Evidence-led execution	External study data interpreted against real plant layouts, density targets, and commissioning sequence.	Published performance data is translated into delivery decisions rather than left as an abstract technical claim.

THE PRACTICAL PREMIUM

The premium belongs to cooling architectures that reduce complexity while supporting density. Two-phase becomes especially relevant when both reference-design and peer-reviewed studies point in the same direction: lower plant burden, lower operating energy, and stronger thermal performance at warmer inlet or water temperatures.

CONCLUSION

Two-Phase Is an Infrastructure Decision

Two-phase direct-to-chip cooling deserves attention because independent evidence now points in the same direction at different scales — lower chiller count, no separate technology cooling system, lower five-year cost, and stable operation at warmer water temperatures.

ABOUT THE PUBLISHER

At DARKNX, we develop and deliver high-density data-centre infrastructure for AI, HPC, and enterprise environments. We evaluate cooling architecture as part of the broader system that determines whether a site will energize cleanly, scale predictably, and absorb future density without avoidable redesign. Where two-phase cooling improves plant simplicity, commissioning discipline, or long-term expansion headroom, it becomes an infrastructure decision with direct commercial value.

EMAIL	PHONE	WEBSITE	OFFICES
info@darknx.com	+1 647 850 6315	darknx.com	Seattle, WA & Toronto, ON

REFERENCES

Accelsius LLC. Jacobs 10MW Data Center Reference Designs. V121725, 2026.

Ahmadi, V. E., Barestrand, H. A., Summers, J., and Yilmaz, C. Performance evaluation of two-phase direct-to-chip liquid cooling combined with air cooling for data centers. Case Studies in Thermal Engineering, 78, 107650, 2026.

Wang, L., Cheng, H., Yang, T., Yuan, W., and Ren, K. Modeling and Performance Analysis of a Pump-Driven Chip-Level Two-Phase Cooling System in Data Centers. Applied Sciences, 13(13), 7472, 2023.