



White paper

Transforming colocation facilities to meet demand for AI compute





Table of contents

Executive summary	3
The impact of AI on data center infrastructure	4
The AI-ready power system: Maintaining continuity amid changing requirements	5
Reducing demand on the grid	5
Meeting AI power needs	6
Delivering power to high-density racks	7
Removing heat from high-density pods: Optimizing efficiency and performance	8
Optimizing air and liquid cooling	8
Understanding liquid cooling infrastructure	9
Liquid cooling commissioning and management	10
Managing change with speed and efficiency	11
Using reference designs to streamline planning	11
The value of preconfiguration	12
Conclusion	12
References	13

Executive summary

Across every sector, business leaders recognize the potential of artificial intelligence (AI) to enhance operational efficiency, drive innovation, and elevate the customer experience. Management consulting company Oliver Wyman conducted a 2025 survey of CEOs from New York Stock Exchange-listed companies and found that 95% view AI as an opportunity for their business, and 83% are actively investing in AI across major use-case categories. ¹ Similarly, respondents to the 2025 IBM CEO Survey expect spending on AI to accelerate in the next two years, with 61% saying they are actively adopting AI agents and preparing plans to scale them across the enterprise.²

Despite executive enthusiasm, enterprise IT teams face significant hurdles in delivering on these aspirations. In particular, legacy infrastructure is often unable to support AI's unique power and cooling demands, impeding organizations' ability to scale their AI pilots. While cloud-based AI offers a simpler path for some use cases such as smart manufacturing, fraud detection and health records analysis have latency, privacy, or security requirements that make public cloud deployment impractical.³ As a result, many organizations will be best served by a hybrid AI strategy that combines public and private deployment models.

This creates an enormous opportunity for colocation and is already driving the next phase of expansion and growth. However, to capitalize on this opportunity, providers must embrace the challenge of delivering scalable, efficient, and reliable AI-ready infrastructure.


This white paper provides best practices for adapting existing infrastructure for AI support or developing new AI-ready facilities, enabling providers to reliably and profitably support the next era of computing.



¹ Kreacic, A., Romeo, J., Yentob, A., Lopez, J., and Spivak, V., 2025. Oliver Wyman Forum.
² IBM, 2025.
³ Thomas, C., Tayal, A., Stewart, D., Kearns-Manolatos, D., and Parveen, I., 2025. Deloitte.

The impact of AI on data center infrastructure

AI compute performs best when high-power graphics processing units (GPUs) are clustered in close proximity, often resulting in high-density racks. As AI adoption accelerates, rack power density is expected to rise even higher as more GPUs are packed into each rack and GPU power requirements escalate. Based on these trends, Vertiv projects the average multi-node density of an AI pod to grow from 40-to-100 kW to 100-to-250 kW in the coming years (Figure 1).



©2026 Vertiv Group Corp

Figure 1: The density of AI deployments is projected to double in the coming years.
Source: Vertiv

In addition, the power draw of AI systems is more variable than traditional IT loads due to the architectural differences between central processing units (CPUs) and GPUs. CPUs typically draw power at a relatively stable rate, while GPUs exhibit more variable power usage patterns.

The implications of these requirements extend from the utility grid to the chip and from the chip to heat reuse, fundamentally reshaping both power delivery and thermal management. Providers adapting their facilities for these demands will want to remain focused on the core factors critical to colocation success:

- **Reliability:** Customers who invest in AI technology will need high availability to generate a return on those investments. Meeting these expectations requires more adaptive power and cooling infrastructure that can maintain service level agreements (SLAs) while supporting denser and more variable loads.
- **Efficiency:** Cost-effective colocation depends on delivering reliable infrastructure with capital, energy, and operational efficiency that surpasses what enterprises can achieve on their own. High-efficiency cooling and power management are central to supporting the colocation value proposition.
- **Scalability:** Today, deploying modular infrastructure that can scale with demand represents only part of the solution to managing growth. With AI deployments driving increased power consumption across the industry, providers should consider strategies to supplement utility power to meet their growth objectives.

The AI-ready power system: Maintaining continuity amid changing requirements

Power is the backbone of data center availability, and AI requirements can stress legacy power systems in ways that increase the risk of downtime. Infrastructure providers like Vertiv work closely with chip manufacturers to adapt power system technologies to AI requirements.

Reducing demand on the grid

Growing demand for data center power, driven largely by AI, is already creating challenges for new data center development with longer delays securing power connections, unpredictable or limited energy supply, and higher interconnection costs in some areas.⁴ There are also signs that increasing demand is impacting grid reliability. According to Deloitte: “Some of the leading data center growth regions have experienced harmonic distortions, load relief warnings and near-miss incidents, and generation shutdowns.”⁵

Providers seeking to protect themselves from the risks associated with grid stress and instability are increasingly turning to “bring your own power” (BYOP) strategies.⁶ Battery energy storage systems (BESS), which typically leverage lithium-ion batteries sized to deliver from one to six hours of run time at full load, and grid support uninterruptible power supply (UPS) modes are being integrated into the data center power train. In addition to providing longer duration backup power than is available through the UPS system, these solutions use stored energy to reduce peak demand and stabilize power from the grid.

Working with a grid-interactive UPS, BESS enables greater flexibility and resilience by coordinating energy storage and delivery between the grid and the data center’s critical loads. In this scenario, in addition to its role in enabling seamless backup power, the UPS actively manages energy flows for peak shaving, load shifting, and dynamic grid support services. The UPS and BESS operate together, using real-time monitoring and intelligent controls to optimize charging, distribution, and backup functions, thereby maximizing availability, flattening demand curves, and reducing generator starts.

Solution profile: A close-coupled UPS-Switchgear package



Vertiv™ PowerNexus
©2026 Vertiv Group Corp.

A recent innovation in power system design highlights Vertiv™ PowerNexus, a close-coupled Vertiv™ Trinerger™ UPS-Switchgear integrated system. This solution reduces facility footprint, minimizes cable runs and electrical loss, and enables faster deployment and commissioning. The close coupling also supports real-time monitoring, minimizes points of failure, and allows for tighter coordination between protection, control, and emergency switching functions.

⁴ Judge, P., 2023. Data Center Dynamics.
⁵ Stansbury, M., Marchese, K., Hardin, K., and Amon, C., 2025. Deloitte.
⁶ Vertiv, 2024.


Meeting AI power needs

In addition to adding grid support capabilities, AI-ready UPS technologies can effectively and better support dynamic AI workloads, simplify scalability, and deliver improved operating efficiency.

Newer UPS platforms have been engineered and rigorously tested to manage the rapid load fluctuations generated by AI training and inference cycles. These systems deliver stable output voltage and current even when faced with abrupt step changes in GPU power demand, usually ranging from less than 30% to 100%, with peaks up to 150% of rated capacity within milliseconds. Through resilient design they deliver continuous power quality while supporting increased rack power densities.⁷ UPS systems can also actively manage and mitigate the effects of AI workloads with a set of advanced control features to either minimize the impact on the batteries or smooth the input power to reduce the stress on utility, generators and upstream infrastructure.

The patented Vertiv™ AI Load Simulator was developed to replicate the fast, variable power profiles of AI workloads, enabling the validation of UPS systems and related components under realistic operating conditions prior to deployment. By replicating GPU demand and cyclical load patterns, the simulator allows engineers to fine-tune control algorithms and settings to maintain stable output without unnecessary battery use, helping avoid overengineering and increasing confidence that power systems are AI-ready.⁸

Case study: EdgeConneX upgrades for AI-scale power demand



Background

EdgeConneX is a global data center provider that partners with service providers to deliver tailored solutions, enhancing access to content, cloud services, networks, and applications.

EdgeConneX partnered with Vertiv to address the unique power challenges posed by AI variable loads such as rapid load fluctuations and the need for seamless power transitions. Vertiv and EdgeConneX rigorously tested the performance of large UPS systems under the variable power demands of AI variable loads. The results demonstrated that Vertiv's UPS solutions provide reliable power continuity and energy efficiency, meeting the resilience and performance required for AI-driven data centers.

Company profile

A leading global data center provider specializing in efficient content delivery, cloud solutions, and application optimization across the world.

Industry: Data center and digital infrastructure

Region: Atlanta, USA

Challenge

AI workloads can generate highly variable and unpredictable power demands due to rapid shifts between compute-intensive and idle cycles, creating challenges such as:

- Managing sudden step changes in GPU power consumption without disruptions
- Providing continuous supply of power during transitions between utility and generator sources
- Accelerated battery aging caused by frequent recharging
- Maintaining compliance with strict IEEE 802.3 Class 1 power stability standards under all operating conditions

EdgeConneX chose a power solution that could reliably handle these challenges while maintaining operational continuity and efficiency in its AI-focused data centers.

Related links

[Evaluating the performance of Vertiv™ large UPS systems with AI variable loads](#)

[Vertiv's AI power load management solutions for data centers](#)

As AI adoption accelerated, EdgeConneX upgraded their power system and collaborated with Vertiv to validate the resilience of the Vertiv™ Trinerger™ UPS system under highly variable and sustained AI power loads. The Vertiv™ UPS demonstrated stable power output, seamless transitions between utility and generator, and overload capability during AI training cycles, enabling EdgeConneX to confidently support intensive AI workloads while maintaining operational efficiency and high uptime.



Scan the QR code to find the details of this case study.

Delivering power to high-density racks

At lower voltages, higher currents are required to support high-density racks, resulting in greater electrical losses, increased heat generation, and thicker cabling that raises installation complexity and labor costs. Ultimately, these considerations can create a cap on rack power density, limiting an operator's ability to meet customer requirements.

Many operators have already shifted from legacy distribution voltages to higher low-voltage distribution, such as 400/415V or direct 480V, to deliver more power with less current. As rack densities continue to climb, however, these new architectures will begin to experience similar issues as the legacy architectures they replaced. As a result, 800 VDC distribution is emerging as the next evolution in power distribution.⁹

Busways have become an important enabler of this evolution, providing a flexible, modular overhead system that efficiently handles high-capacity loads without the complexity and bulk of traditional cable runs. By supporting voltages above 480V with specialized open-track busways, which enable facilities to deliver greater amounts of power with lower currents, significantly reducing conductor size, installation costs, and transmission losses. Busways that feature hot-swappable tap-off boxes and modular design also streamline the additions and reconfigurations that are common in colocation facilities.

Solution profile: Designing for density



Vertiv™ SmartRun
©2026 Vertiv Group Corp.

Vertiv™ SmartRun is a pre-assembled, integrated plug-and-play infrastructure system that integrates high-density power distribution, liquid cooling, containment, and network support into one scalable overhead system. Designed for rapid deployment and maximum operating flexibility, Vertiv™ SmartRun allows power tap-off boxes and infrastructure elements to be repositioned or added without the need to shut down.



⁷ Raggi, T., 2025. Vertiv.

⁸ Remelli, A., 2025. Vertiv.

⁹ Vertiv, 2025.

Removing heat from high-density pods: Optimizing efficiency and performance

Comparing basic characteristics, air has relatively low thermal conductivity and a lower heat capacity than liquid. Air-cooling infrastructure has also reached its practical physical limits.¹⁰ From a business perspective, the economics of cooling IT equipment have rapidly shifted, particularly when considering liquid cooled applications. As energy costs represent one of the largest operational expenses (OpEx) in data center management, and with power usage effectiveness (PUE) ratios for air-cooled facilities typically hovering around 1.5-1.7, up to 30% of power is used toward non-computing functions.¹¹ ¹² Liquid cooling can help reduce PUE, resulting in direct operational savings.¹³ Whether as a transitional infrastructure solution or a complete shift, data center operators face increasingly complex challenges in implementing thermal management.

Optimizing air and liquid cooling

The dramatic power increases required by AI workloads directly translate into unprecedented heat loads that must be managed effectively to prevent thermal-induced shutdowns and CPU/GPU frequency throttling due to junction temperatures exceeding thermal design specifications.¹⁴ This creates a vicious cycle: higher power density requires more cooling, which in turn requires more power, which generates more heat.

Liquid cooling provides higher thermal transfer rates than air cooling, allowing systems to remove more heat with greater efficiency. Regardless of the method selected, liquid cooling solves this problematic cycle by bringing the cooling medium closer to the heat source, removing heat more effectively at the point of generation rather than trying to condition entire room environments.

However, in most high-density environments, liquid and air-cooling systems work together to provide effective thermal management. This is especially true in colocation data centers, where providers often support low-, medium-, and high-density zones and where RDHxs and direct-to-chip (DtC) liquid-cooling systems have become the technologies of choice.

Direct-to-chip liquid cooling targets high-power components—such as CPUs and GPUs—within each rack and is typically capable of removing over 70% of the total rack heat load. Meanwhile, the remaining estimated 30% of heat is generated by ancillary components such as memory modules, network interface cards, and power supplies.¹⁵ Complementary air cooling removes residual heat from the rack and maintains the ambient environment around it.

Solution profile: **Streamlining the transition to liquid cooling**



Vertiv™ CoolPhase Flex
©2026 Vertiv Group Corp.

The Vertiv™ CoolPhase Flex is a first-of-its-kind hybrid cooling solution that incorporates both air and liquid cooling modes in a compact, modular design, providing the ability to transition smoothly between air and liquid cooling as densities increase or operational needs change. This approach accelerates deployment, limits upfront retrofit requirements, and optimizes energy use by leveraging free cooling when possible.

With RDHxs, the air-cooling system must have the capacity to temporarily provide 100% of cooling to a high-density rack while the heat exchanger is being serviced. In addition, air-cooling systems will often be maintained or even upgraded to manage heat removal in lower-density zones, which may be experiencing their own relative density increase as IT systems supporting traditional workloads are refreshed or scaled.

Best practices for balancing air and liquid cooling capacity begin with defining the portion of total thermal load that will be managed by liquid systems and evaluating the ability of existing air-cooling infrastructure to effectively handle residual heat and support the ambient environment. Computational fluid dynamics (CFD) and digital simulations can then be used to model airflows, optimize piping layouts, and coordinate the operation of the air and liquid systems.

Modern data center complexity requires an engineering-driven approach that addresses the entire thermal chain—from chip-level heat management to facility-wide systems—through customized solutions that integrate cooling strategies seamlessly across all infrastructure layers.¹⁶ Liquid cooling can be implemented through a phased, modular approach that meets near-term customer demand while providing the flexibility to expand capacity in the future. However, a number of technical issues specific to fluid selection, distribution, and management should be addressed prior to the first phase of deployment.

A partner with deep experience in liquid and air cooling can provide services that include analyzing current and future capacity requirements for liquid and air cooling systems, positioning of liquid cooling loops and coolant distribution units, pipe routing, control system integration, monitoring points for fluid and leak detection, and potential compliance with local safety and environmental standards.

Understanding liquid cooling infrastructure

Effective liquid cooling infrastructure relies on strategic deployment of coolant distribution units (CDUs) and heat rejection capacity to efficiently manage thermal loads and protect customer equipment racks.

CDUs serve as the core functional element of a liquid cooling system, circulating coolant to specific racks or rows and isolating liquid cooling fluid loops from facility water systems for safety and temperature control. By isolating the fluid cooling loop and through integrated filtration systems, CDUs also prevent contaminants from reaching cold plates and IT components. Placing CDUs as close as possible to the racks they support is one recommendation to minimize fluid routing complexity, reduce risk from leaks, and improve response time to temperature changes. Another possible configuration is to place CDUs outside the whitespace in the mechanical gallery. These types of CDUs are often higher-capacity and can be combined together in a modular way to provide scalable cooling capacity for larger deployments while supporting the entire IT whitespace rather than just a single row of racks. These various deployment styles—in-rack, in-row, or gallery CDU layouts—enable flexible adaptation to the density and redundancy needs of each suite or containment zone.

Determining the appropriate liquid-to-liquid CDU configuration requires assessment of total heat loads, desired redundancy, available physical space, and routing logistics. Row-level or suite-level N+1 redundancy is often advised, depending on space availability. Centralized CDUs can simplify scalability while potentially reducing the total number of units and associated maintenance but require careful planning for pipe routing and fault domain isolation.

There is also a strong potential for providers to deploy liquid-to-air CDUs in colocation environments: they can deliver liquid to the chip and still utilize the existing infrastructure. So for a colo provider looking to transition to supporting liquid for a few customers or a few zones within the space, they can add those L2A CDUs and not have to completely redesign the surrounding spaces. Much like the L2L CDU deployment, this would also require an assessment if the existing infrastructure (e.g., perimeter air cooling) can support the increased heat load from the liquid-cooled racks.

¹⁰ ST Telemedia Global Data Centres, 2025.

¹¹ Kok, E., Rauer, J., et. Al., 2025. McKinsey & Company.

¹² Zhang, Q., Meng, Z., et. Al., 2021. Journal of System Architecture.

¹³ Vertiv, 2024.

¹⁴ Vertiv, 2025.

¹⁵ Theoretically, achieving full liquid cooling would require cold plates on every heat-generating component within the server. While technically feasible, as of this writing, it would be economically impractical as an approach that would add complexity and cost without proportional thermal benefit for lower-wattage components.

¹⁶ Vertiv, 2026.

Leak detection also requires careful attention to best practices. To maximize early identification and enable rapid containment, sensors should be placed at every critical connection and along the fluid network. Effective approaches include resistive cable sensors and spot sensors integrated with CDU monitoring and facility controls.

Chillers and other heat rejection-enabled deployments such as DX technologies with liquid cooling can work in tandem with CDUs to maintain precise supply temperatures across closed-loop cooling circuits. Chiller capacity should be matched to peak rack heat rejection, while also considering margin for potential future growth and the partial load operation common in colocation environments. Integration best practices call for clear separation of primary and secondary fluid loops, side-stream filtration to prevent contamination, pressure relief valve planning, and provision for routine coolant sampling and external flushing to verify system cleanliness and optimal performance.

Liquid cooling commissioning and management

Liquid-cooling systems introduce a range of requirements that differ substantially from traditional air-cooled environments, and providers will encounter new technical, operational, and maintenance challenges during initial deployment. Operators should be prepared for commissioning processes that are more rigorous and to adopt fluid-specific lifecycle maintenance practices.

For example, commissioning new liquid cooling systems involves thorough fluid quality verification, air removal from all pipe runs, validation of all sensors and alarms, and staged ramp-up with simulated heat loads where possible.

Best practices for maintaining liquid cooling systems include end-to-end lifecycle services that include ongoing coolant sampling and quality testing, routine leak detection inspections, and scheduled filter replacements to prevent particulate buildup. Operators should also be prepared to support regular system flushing, air purging, and cleaning protocols to maintain performance and extend equipment life. Experienced system and service providers like Vertiv have developed a comprehensive suite of solutions from end to end: planning, commissioning, deployment, maintenance, and the responsible removal of liquid cooling components and solutions.

Colovore adds liquid cooling



Colovore, a colocation provider focused on high-density applications, turned to Vertiv's rear door heat exchanger to enable customer deployments of fully packed racks up to 50 kW. The liquid cooling solution eliminated wasted rack space, reduced the need for additional cabinets, improved overall floor space efficiency, and supported advanced application profiles, including early-stage AI. By embracing liquid cooling, Colovore delivered enhanced operating and capital efficiency, positioning itself to meet the continuous growth in density and compute capability required by AI workloads.

Scan the QR code to find the details of this case study



Managing change with speed and efficiency

Scalability with speed will be instrumental in colocation providers' ability to capitalize on surging AI demand. As operators demonstrate their ability to reliably support AI workloads, the need to deploy additional capacity quickly will only intensify—making strategic approaches to scaling with speed critical to keeping pace with the market's AI adoption curve.

Using reference designs to streamline planning

By defining best practices for power, cooling, and system architecture, reference designs reduce complexity and integration risks, which is especially important when adapting to evolving AI requirements and high-density workloads.

Vertiv has collaborated with the leading chipmakers to develop a comprehensive portfolio of AI-focused reference designs spanning a range of retrofit and new build scenarios.

These designs are engineered to optimize rack densities, thermal management, and power delivery for current and next-generation GPU platforms. They currently encompass densities from 10 kW to over 140 kW per rack, with integration options for battery energy storage and unified control systems. With these proven templates as a guide, colocation providers can standardize infrastructure across sites, respond quickly to rapid technology shifts, and minimize deployment risks.

Modular, AI-ready colocation for Polar



Polar partnered with Vertiv to develop a flexible, modular data hall that can accommodate advanced liquid cooling technologies and support continuous operation for AI workloads. Vertiv delivered a prefabricated, AI-ready data center comprising factory-assembled and tested components, dual-busbar redundant power distribution, and N+1 electrical topology supporting high-density racks up to 110 kW. The project included grid-interactive, lithium-ion (Li-ion) battery-backed Vertiv™ Liebert® EXL S1 UPS and a concurrently maintainable chilled water network, plus Vertiv™ Liebert® AFC glycol-free chillers using low global warming potential (GWP) refrigerants.

Scan the QR code to find the details of this case study




The value of preconfiguration

Preconfigured, factory-integrated solutions deliver advantages that align with colocation providers' demand for speed, scalability, and support of advanced technologies. By testing and assembling modular infrastructure offsite—including cooling, power, containment, and monitoring—components arrive fully integrated and ready for rapid installation, cutting deployment times by up to 50% compared to traditional processes.¹⁷ This approach simplifies expansion as modular blocks can be added or upgraded in response to demand, allowing operators to scale capacity efficiently and adapt to evolving workload requirements.

Preconfigured systems also streamline support for hybrid liquid-air cooling, high-density power distribution, and dynamic workload management, all engineered and tested for compatibility and performance before delivery. Operators gain faster time-to-market (TTM), consistent quality, and more predictable costs.

Solution profile: Pre-configured AI-ready capacity



Vertiv™ MegaMod™ CoolChip
©2026 Vertiv Group Corp.

Vertiv™ MegaMod™ CoolChip is a high-density, pre-fabricated modular data center solution designed for the rapid deployment of AI and compute-intensive workloads. Modules integrate direct-to-chip liquid cooling for racks with capacities up to and exceeding 100 kW, with configurable air-to-liquid cooling ratios to match platform requirements. The solution includes high-efficiency rack power distribution, busway, and close-coupled UPS/switchgear for optimized power delivery and protection. With scalable capacities of multi-megawatts, Vertiv™ MegaMod™ CoolChip enables accelerated build times and reliable support for AI compute, all within a compact footprint.

Conclusion

The movement of AI workloads to colocation environments is already well underway. Providers who deploy AI-ready power and cooling technologies while leveraging standardized and modular deployment strategies will be well-positioned to capture this growing opportunity. With a focus on reliability, efficiency, and scalability, colocation providers can achieve sustained growth while meeting evolving customer requirements. For help assessing your needs or developing plans to support AI, [contact a Vertiv representative today](#).

References

1. IBM. (2025). *IBM study: CEOs double down on AI while navigating enterprise hurdles*. <https://newsroom.ibm.com/2025-05-06-ibm-study-ceos-double-down-on-ai-while-navigating-enterprise-hurdles>.
2. Judge, P. (2023). Data Center Dynamics. *How to build when the power isn't there*. <https://www.datacenterdynamics.com/en/analysis/how-to-build-when-the-power-isnt-there/>.
3. Kok, E., Rauer, J., Sachdeva, P., Pikul, P., Sutton, D., and Hasrouni, R., 2025. McKinsey & Company. *Scaling bigger, faster, cheaper data centers with smarter designs*. <https://www.mckinsey.com/industries/private-capital/our-insights/scaling-bigger-faster-cheaper-data-centers-with-smarter-designs>.
4. Kreacic, A., Romeo, J., Yentob, A., Lopez, J., and Spivak, V. (2025). Oliver Wyman Forum. *The CEO Agenda 2025: In black and white*. <https://www.oliverwymanforum.com/ceo-agenda/how-ceos-navigate-geopolitics-trade-technology-people.html>.
5. Raggi, T. (2025). Vertiv. *Understanding AI power demands: How standard Vertiv UPS systems support AI factory load dynamics*. <https://www.vertiv.com/en-emea/about/news-and-insights/articles/educational-articles/understanding-ai-power-demands-how-standard-vertiv-ups-systems-support-ai-factory-load-dynamics/>.
6. Remelli, A. (2025). Vertiv. *Vertiv AI Load Simulator: Innovation for precise power validation*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/blog-posts/vertiv-ai-load-simulator-innovation-for-precise-power-validation/>.
7. Stansbury, M., Marchese, K., Hardin, K., and Amon, C. (2025). Deloitte. *Can US infrastructure keep up with the AI economy?* <https://www.deloitte.com/us/en/insights/industry/power-and-utilities/data-center-infrastructure-artificial-intelligence.html>.
8. ST Telemedia Global Data Centres. (2025). *From air to liquid: The evolution of data centre cooling to support AI workloads*. <https://www.sttelemediagdc.com/sg-en/resources/air-liquid-evolution-data-centre-cooling-support-ai-workloads>.
9. Thomas, C., Tayal, A., Stewart, D., Kearns-Manolatos, D., and Parveen, I. (2025). Deloitte. *Is your organization's infrastructure ready for the new hybrid cloud?* <https://www.deloitte.com/us/en/insights/topics/digital-transformation/future-ready-ai-infrastructure.html>.
10. Vertiv. (2024). *Renewable energy for data centers: A path to grid independence*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/white-papers/renewable-energy-for-data-centers-a-path-to-grid-independence/>.
11. Vertiv. (2024). *How to implement liquid cooling at existing data centers*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/white-papers/how-to-implement-liquid-cooling-at-existing-data-centers/>.
12. Vertiv. (2025). *DCD AI Week: Engineering thermal resilience for next-gen workloads*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/educational-articles/dcd-ai-week-engineering-thermal-resilience-for-next-gen-workloads/>.
13. Vertiv. (2025). *800 VDC architectures will power AI expansion*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/blog-posts/the-vertiv-view-high-voltage-dc-will-power-ai-expansion/>.
14. Vertiv. (2025). *Vertiv 360AI: Accelerate your deployment*. <https://www.vertiv.com/4a4819/globalassets/documents/vertiv-360ai-br-en-emea-sl-71291-emea-rev2-082025-web.pdf>.
15. Vertiv. (2026). *The technology trends shaping the future of the data center*. <https://www.vertiv.com/en-us/about/news-and-insights/articles/frontiers/>.
16. Zhang, Q., Meng, Z., Hong, X., Zhan, Y., Liu, J., Dong, J., Bai, T., Niu, J., and Deen, M.J. (2021). *A survey on data center cooling systems: Technology, power consumption modeling and control strategy optimization*. *Journal of Systems Architecture*, 119, 102253. <https://doi.org/10.1016/j.sysarc.2021.102253>. <https://www.vertiv.com/en-us/about/news-and-insights/articles/educational-articles/dcd-ai-week-engineering-thermal-resilience-for-next-gen-workloads/>.

¹⁷ Vertiv, 2026.



Vertiv.com | **Vertiv Headquarters**, 505 N Cleveland Ave, Westerville, OH 43082, USA

© 2026 Vertiv Group Corp. All rights reserved. Vertiv™ and the Vertiv logo are trademarks or registered trademarks of Vertiv Group Corp. All other names and logos referred to are trade names, trademarks or registered trademarks of their respective owners. While every precaution has been taken to ensure accuracy and completeness here, Vertiv Group Corp. assumes no responsibility, and disclaims all liability, for damages resulting from use of this information or for any errors or omissions. Specifications, rebates and other promotional offers are subject to change at Vertiv's sole discretion upon notice.