VERTIV.

Frontiers

The technology trends
shaping the future
of the data center

2026

Contents

> The evolution of AI depends on highly integrated digital infrastructure deployed at unprecedented speed and scale. The data center industry must keep pushing the frontiers of innovation—driving advances from grid to chip, chip to heat reuse—while preparing for the next wave of breakthroughs.”

Scott Armul
Chief Product and Technology Officer, Vertiv

Executive summary

The data center industry is advancing into a new era defined by innovation, scale, and opportunity.

After two decades of steady evolution - when cloud computing reshaped location and scale but core infrastructure remained largely constant - the next wave of transformation is accelerating at unprecedented speed.

Driven by AI and accelerated compute, this new era is redefining how digital infrastructure is designed, deployed, and scaled. The pace of change is unmatched, creating new possibilities to push the frontiers of innovation.

Vertiv Frontiers offers a lens on the future - an exploration of macro forces and the technology trends reshaping digital infrastructure. It brings together the expertise of Vertiv specialists across power, thermal, IT systems, prefabricated modular infrastructure, advanced services, and AI infrastructure, reinforcing Vertiv's position as a leading voice guiding the future of critical digital infrastructure.

Macro forces driving data center innovation

The data center industry is being reshaped by powerful macro forces fueled by the rise of AI and accelerated compute. These forces are influencing every layer of digital infrastructure, spanning technologies, architectures, and industry segments.

At the center of this transformation is **extreme densification** - the defining macro force whose effects are felt across the entire data center and technology landscape. The additional macro forces align to this shift, representing broad changes that extend from the chip level to system-level integration and full facility design.

Extreme densification

Densification is a critical driver of modern data centers, accelerated by AI and HPC workloads pushing rack power well beyond 25 kW and often into triple digits. Systems that previously filled a data hall now fit into one rack-like unit or pod. This density increase creates additional complexity in power, cooling, and space.

Gigawatt scaling at speed

Data centers are now being built at unprecedented gigawatt scale. Factory-built, modular infrastructure blocks—combining power, cooling, and IT—are designed to scale from tens of megawatts to multi-gigawatt campuses. And all of this must be done at speed to meet the compressed timelines of AI-era capacity demands.

Data center as a unit of compute

The AI-era increasingly requires the data center to be built and operated as a single system. The 'unit of compute' is no longer just a chip—it's the entire system. Power, cooling, and compute must be highly integrated into one architecture – from rack to row to site. Component-level offerings are no longer strategic, and rising power density is driving the need for greater efficiency and tighter integration.

Silicon diversification

The chips powering AI are already diversifying to include custom silicon and other form-factors. Future data center infrastructure must be designed and optimized to support the full spectrum of chips and compute.

The technology trends defining market impact

In response to these macro forces, Vertiv has identified five key trends shaping the data center technology landscape. Each trend has been evaluated for its specific impact across critical technology areas. **Vertiv Frontiers** explores these trends in greater depth.

Trends

1	**Powering up for AI**	Hybrid AC and DC systems are pervasive but higher voltage DC is likely to become more prevalent as rack densities increase.
2	**Distributed AI**	AI will become increasingly critical to businesses. But how, and from where, AI services are delivered will depend on the specific requirements and conditions of the organization.
3	**Energy autonomy accelerates**	Operators are expanding on-site generation despite the grid being the preferred option for many. Investment in on-site is likely to continue until grid capacity expands and transforms.
4	**Digital twin-driven design and operations**	Data centers can be mapped and specified virtually using digital twins, and the IT and infrastructure will be integrated and deployed as units of compute. This approach will be key to achieving the gigawatt-scale buildouts required for future AI advancements.
5	**Adaptive, resilient liquid cooling**	Over time, liquid cooling systems could become analogous to circulatory systems with the ability to monitor and adapt. New forms of liquid cooling will emerge as the technology continues to mature.

The **technology trends** defining market impact

1 / Powering up for AI
2 / Distributed AI
3 / Energy autonomy accelerates
4 / Digital twin-driven design and operations
5 / Adaptive, resilient liquid cooling

These trends are already impacting advanced data center operators and will likely be adopted by the wider industry as the underlying technologies and drivers mature.

1 | Powering up for AI

Hybrid AC and DC systems are pervasive but higher voltage DC is likely to become more prevalent as rack densities increase.

Liquid cooling **(Trend 5)** has become closely linked to the growth in artificial intelligence (AI) workloads and accelerated compute. However, innovations in power management are as critical in driving the evolution of data center performance. One important area of power innovation is the move to higher voltage power distribution within the facility. Many existing data centers still rely on hybrid AC/DC power distribution from the grid to the IT racks, which includes three to four conversion stages and resulting inefficiencies. This approach could face challenges as densities increase largely driven by AI workloads.

AI compute platforms are rapidly pushing rack densities beyond 300 kW. In these high-performance environments, traditional 415 VAC or 480 VAC architectures face limitations in copper volume, thermal loss, and space efficiency. These limitations create potential constraints for achieving the necessary densities to support AI workloads. Shifting to higher voltage DC architectures enables significant reductions in current, conductors, and conversion stages while centralizing power conversion at the room level.

Innovations are already underway in higher voltage DC architecture, including recent announcements from NVIDIA, supported by Vertiv and others, around 800 VDC architecture. The 800 VDC technology enables more efficient power delivery.

Analysts, including Uptime Institute Intelligence, agree that AI, and specifically densification, will drive the adoption of higher voltage DC architectures within the data center. “Current data center power chains and electrical equipment are highly efficient and can support a wide range of capacities,” said Uptime Institute Intelligence senior research director Daniel Bizo.[1] “However, despite best efforts by facility designers and equipment manufacturers to enhance LV systems, they remain suboptimal for a high-density future.”

Looking ahead, it is likely AC power will be used more selectively, including at the grid interface—which may be relegated to back-up in some scenarios—with DC becoming an important approach for internal power distribution. The use of higher voltage DC architectures also aligns with the growth in energy storage within data centers and at the grid level. These issues are further discussed in **Trend 3: Energy autonomy accelerates,** where some data centers may evolve into DC microgrids that integrate on-site generation and energy storage.

Other future innovations in powering up for AI could also include a shift towards architectures leveraging more integrated power systems, including integration of medium voltage (MV), battery energy storage systems (BESS), uninterruptible power supply (UPS), and advanced power controls for managing the unique challenges of AI load profiles, the scale and density of training and inference facilities, and the evolving requirements of grid/utility compatibility and interaction.

[1] Bizo, D., 2024. Uptime Institute Intelligence.

> **“Higher voltage DC architectures are already being developed with deployments planned supporting increasing GPU rack and POD power densities. Facility operators and decision-makers can take immediate steps toward planning future-ready AI factories that will support these coming generations of AI servers."**
>
> **Kyle Keeper,**
> Senior Vice President, Power Management, Vertiv

Macro forces

- **Extreme densification:** Traditional power distribution is reaching its limits due to the densities of AI compute.
- **Data center as a unit of compute:** Deploying and managing power, thermal, and compute as a single system improves performance, resiliency, and efficiency.

Drivers

- Fewer conversion losses mean higher efficiency: Traditional AC power systems face significant energy losses due to multiple conversion steps.
- Builds on expertise developed for global telecom networks, integrated microgrids, and other applications.
- Integration with alternative energy: Alternative energy sources like solar panels or fuel cells natively produce DC power.

Challenges

- Higher voltage distribution may require stricter safety procedures and protocols.
- Some operators could face a lack of qualified staff for installation and maintenance.
- Higher upfront costs potentially until the technology matures further.

Related/example technologies

- Higher voltage UPS systems, high-efficiency DC busways, rack-level DC-DC converters, and DC-compatible backup systems.
- DC Microgrids, energy storage

Suggested actions

- Hyperscale and other early adopters of AI workloads and high-performance systems should actively investigate the benefits of higher voltage DC adoption – especially as part of on-site generation or microgrid deployment.
- Early access to reference designs and roadmaps allows customers to validate and deploy supporting infrastructure in parallel with emerging graphics processing unit (GPU) systems.
- Enterprise operators and smaller colocation providers should investigate and look for opportunities to trial the technology, potentially within high-density areas in existing facilities. Higher voltage DC architectures can be deployed into today’s traditional AC architectures.

2 | Distributed AI

AI will become increasingly critical to businesses. But how, and from where, AI services are delivered will depend on the specific requirements and conditions of the organization.

The billions of dollars invested into AI data centers to support large language models (LLMs) to date have been aimed at supporting widespread adoption of AI tools by consumers and businesses. Rather than being experimental technology used by a few advanced organizations or workers, AI is becoming a vital tool. The AI inference data center capacity required to support this shift could even eventually outweigh, by several times, current build outs to support large model training.[2]

However, exactly how business-critical AI is delivered is likely to differ between enterprises, business units, or even specific applications. Highly regulated industries, such as finance, defense, and healthcare, often need to maintain private or hybrid AI environments due to data residency, security, or latency requirements. According to technology analyst IDC, the financial services sector is expected to spend the most on AI solutions up to 2028 and will account for more than 20% of all AI spending.[3] Organizations may also choose to invest in AI-optimized private data centers to fine-tune proprietary models or process confidential data at scale. In such cases, organizations will likely adopt hybrid or federated architectures that combine local compute with cloud-based AI services.

According to technology analyst Gartner,[4] there will be a shift away from specialized models created by service providers to in-house domain-specific language models (DSLMs) created and owned by software providers and enterprise companies.

It is not clear yet how this AI model segmentation will translate into specific data center buildouts, but a proportion of DSLMs could require dedicated, owned AI infrastructure by businesses, including on-premise facilities and colocation. Some organizations will build out new on-premise capacity or look to retrofit or expand existing facilities, including through the use of modular technology.

Investment in AI infrastructure is also aligned to changes in enterprise-owned facilities and cloud-first strategies. However, just as with cloud services, costs, vendor lock-in, data portability, and other issues will shape decisions around where and how to source AI services. Availability of power is also an issue, with some enterprises opting to build-out existing data center footprints for AI to make best use of existing power capacity. The issue of data and AI sovereignty will also play a role in owned-AI infrastructure decisions.

> **“AI is becoming the cornerstone of enterprise value creation, shaping strategies and unlocking new possibilities. As data gravity and regulatory forces intensify, organizations will face pivotal choices about whether to build this foundation through trusted service providers or invest in their own next-generation data center infrastructure.”**
>
> **Martin Olsen**
> Vice President, Segment Strategy & Deployment, Vertiv

Macro forces

- **Silicon diversification:** Inference is likely to be built around the most appropriate chips for the application and workload.
- **Extreme densification:** Businesses will likely need to accommodate high-density zones within on-premise data centers.

Drivers

- Data sovereignty and privacy: Regulations may restrict sensitive data from leaving certain jurisdictions or enterprise-controlled environments, driving more on-premise investment.
- Latency and performance requirements: Applications demanding ultra low latency (e.g., real-time control, financial trading, edge analytics) may require local inference capabilities rather than remote cloud processing.
- Cost predictability: For predictable, high-volume workloads, owning infrastructure might still offer lower total cost of ownership (TCO).

Challenges

- Organizations exploring and deploying AI infrastructure on-premise will face challenges around specific technologies and skills related to extreme densification and associated power and cooling architectures.
- Rapid evolution of AI models and hardware means the infrastructure must be designed for multiple compute generations ahead.

Related/example technologies

- Advanced accelerators
- Prefabricated-modular data centers
- High-density, liquid-cooled AI servers

Suggested actions

- Organizations deploying AI infrastructure on-premise will require an ecosystem of skilled technology and service partners from design to operations.
- Deploying on-premises will require an assessment of whether existing facilities can be retrofitted to support AI infrastructure, or new capacity will be required, including modular buildouts.

[2] Noffsinger, J., Patel, M., and Sachdeva, P., 2025. McKinsey & Co.
[3] IDC, 2024.
[4] Gartner, 2025.

3 / Energy autonomy accelerates

Operators are expanding on-site generation despite the grid being the preferred option for many. Investment in on-site is likely to continue until grid capacity expands and transforms.

On-site energy generation has been essential for most standalone data centers for decades. However, persistent power availability challenges are making full 'energy autonomy' an almost unavoidable requirement for operators.

Data centers were facing an energy crunch even before the advent of AI. Large cloud operators more than doubled their electricity use from 2017 to 2021,[5] the year before the introduction of ChatGPT. In 2018, data centers consumed 1.9% of all electricity produced in the U.S.[6] Today, three years into the AI era, that number is 4.5%, and could be headed to 6% in 2026.[7]

In response, data centers are moving into self-generation, usually to bridge a gap, before new grid capacity can be brought online. Regional variation is also a factor with operators in some countries – especially those in developing regions – forced to invest in significant on-site generation due to grid instability.

Hybrid approaches to on-site power include microgrids, connecting to wind and solar generation, in combination with generators, UPS, and BESS.

On-site generation technologies being adopted by AI data centers include natural gas turbines and reciprocating engines to deliver scalable power and thermal energy for combined cooling, heat, and power (CCHP). These technologies can be deployed alongside power and cooling solutions and services, packaged as modular, pre-designed blocks, to shorten design cycles and standardize deployment.

Other technologies that could help meet power availability challenges include hydrogen fuel cells and nuclear power in the form of small modular reactors (SMRs). Geothermal could also be an option in some regions. In most cases, the choice of on-site power systems will be a mix and match exercise given the key considerations of cost, security, and availability.

> **"Data center operators may not aspire to become energy providers, but the rapid growth of AI and its massive power demands is driving those decisions. In many cases, that means embracing self-generation as a strategic bridge as grid capacity continues to evolve."**
>
> **Peter Panfil**
> Distinguished Engineer & Vice President, Technical Business Development, Vertiv

Macro forces

- **Data center as a unit of compute:** Managing the entire power train from grid to chip - as an integrated system - will help meet power availability and efficiency challenges.
- **Gigawatt scaling at speed:** On-site energy generation has become increasingly necessary to get power to the largest sites at speed.

Drivers

- Proliferation of AI and resulting increase in data center energy demands and costs.
- Limitations and instability of the power grid.
- Need to meet responsible business objectives.

Challenges

- Shortages in on-site generation equipment and skilled labor.
- Resistance/regulatory complications related to nuclear deployments.

Related/example technologies

- Gas turbines and related technology
- DC microgrids using solar power, batteries, and fuel cells
- Nuclear power: Small modular reactors (SMRs)

Suggested actions

- Self-powered "bring your own power and cooling" (BYOP&C) strategies can be applied across the spectrum of data centers from enterprise to hyperscale. They can also be tailored to specific needs, ranging from longer-duration batteries and BESS systems to low-impact natural gas generators and turbines, and even lower-impact hydrogen-based fuel cells. Start with a pilot that builds and complements existing power protection and backup infrastructures.

[5] EPRI, 2024.
[6] Shehabi, A., Newkirk, A., Smith, S., et. Al., 2024. Lawrence Berkeley National Laboratory.
[7] Stansbury, M., Marchese, K., Hardin, K., and Amon, C., 2025. Deloitte.

4 | Digital twin-driven design and operations

Data centers can be mapped and specified virtually using digital twins, and the IT and infrastructure will be integrated and deployed as units of compute. This approach will be key to achieving the gigawatt-scale buildouts required for future AI advancements.

Forms of digital twin technology have been used selectively in the data center for more than a decade, but usually to avoid hot spots or cooling inefficiencies, using computational fluid dynamics (CFD) and other approaches. It has been a tactical tool but rarely a strategic one.

AI-powered 3D models are now becoming transformative for engineering and manufacturing. Enhancing conventional CFD and computer-aided engineering (CAE) with technologies such as NVIDIA Omniverse™ libraries for real-time physically accurate simulation, supporting 3D design and digital twins, powered by NVIDIA accelerated computing, can enable real-time visualization, promising shorter development cycles and greater precision.

Digital twins can be used to create physics-based, photorealistic models of reference architecture. This enables real-time collaboration with architects, engineers, and operators, allowing rapid design iteration and validation before deployment. The digital twin simulates infrastructure performance, testing variables and failure scenarios virtually, which saves time and de-risks projects.

This capability extends into operations, providing closed-loop optimization and feed-forward controls to continuously enhance system performance, making it a dynamic tool for smarter infrastructure decisions across the lifecycle.

One of the biggest potential benefits of digital twin technology in the data center is that it enables design or deployment decisions to be simulated in a virtual space before being carried out in the physical world. This has the potential to reduce the cost or any negative impacts on resiliency or efficiency.

Digital twins also need to be kept current based on real-time data so integration with data center management software and digital services is also part of ongoing developments.

By combining a unit of compute design approach with digital twin technology, data center operators will be able to design, deploy, and operate facilities with increasing levels of efficiency and resilience. Data centers will no longer be manually designed buildings populated with IT, but digitally designed and operated units of compute.

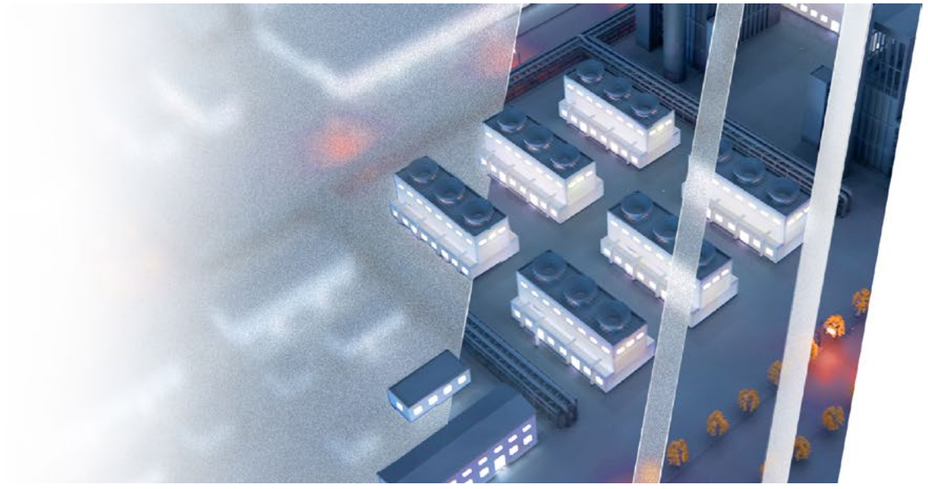

> **"Digital twin technology allows the entire infrastructure to be simulated, monitored, and optimized in real-time before and after construction. Infrastructure no longer operates independently of computing platforms; a data center must operate as a system that continuously adapts as computing loads change."**
>
> **Steve Blackwell**
> Vice President, Engineering, Vertiv

Macro forces

- **Data center as a unit of compute:** Digital twin technology can be used to manage the complexity of deploying highly integrated systems.
- **Silicon diversification:** Data center infrastructure will need to be designed and optimized around an increasing range of chips, which can be supported by digital twin simulations.

Drivers

- More accurate capacity planning by being able to model builds, implementing changes before they are undertaken in the real world.
- Real-time visualization, promising shorter development cycles and greater precision.

Challenges

- A well-developed instrumentation and data management strategy is required to make full use of digital twin technology.
- Data centers still designed and deployed along traditional product, technology lines, and IT/facilities divide can inhibit holistic approaches.

Related/example technologies

- NVIDIA Omniverse™
- Prefabricated modular technologies aligned with the unit of compute approach.

Suggested actions

- Organizations should investigate data center digital twin technology to understand its capabilities for improving both design and operation. The concept of the data center as the next unit of compute, while more abstract, should be used to identify opportunities to drive efficiencies via closer integration of power, thermal, and IT layers across the facility.

5 / Adaptive, resilient liquid cooling

Over time, liquid cooling technologies could become analogous to circulatory systems with the ability to monitor and adapt. New forms of liquid cooling will emerge as the technology continues to evolve.

Liquid cooling is recognized as a critical technology for heat removal for AI infrastructure. "Liquid cooling is moving from niche to necessary," according to Alex Cordovil, research director at industry analysts Dell'Oro.[8]

However, while a lot of attention has been focused on AI as a driver for liquid cooling adoption, AI can also be used to optimize the design and operation of liquid cooling systems. Liquid cooling has become mission-critical for a growing number of operators but AI could provide ways to further enhance its resilience and efficiency. For example, the emergence of adaptive capabilities will require innovations such as embedding sensors in cooling loops to monitor temperature, pressure, coolant quality, and flow rates. Specialized data center management software, and associated digital services, will also be important in providing predictive maintenance capabilities.

There could even be advances to some form of 'self-healing' capabilities in the future, but this will require additional advances in smart materials – potentially liquid metals[9] – to enable actual self-repair (but probably not for major faults or damage). AI could also manifest in the form of physical AI, such as autonomous robots for servicing and inspecting thermal chain equipment, especially in hard-to-access areas of the facility.

However, despite potential advancements in adaptive and automated servicing, installing liquid cooling can be a complex process. Data center teams will continue to require skilled service partners to address key issues, such as plumbing requirements, cooling distribution, custom design elements, risk mitigation strategies, and the solution architecture of the integrated systems within the liquid loop.

Looking ahead to how liquid cooling technology will evolve more broadly over the next five years and beyond, it could include approaches such as phase-change, different types of coolants, as well as the potential of full immersion.[10] As data center infrastructure becomes more complex, it demands greater customization, specialized capabilities, and an engineering-driven approach—not just to address heat at the chip level, but to manage the entire thermal chain. Solving localized heat challenges requires a deeper understanding of how cooling strategies interact across the system, facilitating seamless integration from the component to the facility level.

> **“Liquid cooling is evolving into a self-optimizing system—using AI for predictive maintenance to maintain peak performance, resiliency and efficiency through advanced management and intelligent control.”**
>
> **Nigel Gore**
> Vice President, High Density & Liquid Cooling, Vertiv

Macro forces

- **Silicon diversification:** Liquid cooling systems will need to be optimized for a wide range of chips, compute systems.
- **Extreme densification:** Rising rack densities are one of the key drivers for liquid cooling adoption versus traditional air cooling.

Drivers

- Adaptive capabilities: Advanced instrumentation and control technologies could drive efficiencies and reduce downtime through predictive and proactive management.
- Enhanced heat dissipation: Liquid cooling utilizes the higher thermal transfer capabilities of water or other fluids, which are orders of magnitude more effective than using air, to efficiently cool high-density racks.
- Reduced footprint: Liquid cooling also maximizes space utilization by enabling higher density within the same physical footprint.

Challenges

- Heat removal is becoming increasingly complex, raising the challenge of designing infrastructure that is flexible enough to integrate diverse cooling approaches at an advanced level.
- Traditional air cooling or hybrid approaches may be suitable for some enterprise AI-accelerated compute requirements (AI inferencing) near to mid-term.

Related/example technologies

- Smart embedded sensors for monitoring, fault detection.
- Advanced management and control software
- Smart materials for self-repair
- Autonomous robots for servicing and repair

Suggested actions

- Operators that have already deployed liquid should continue to monitor advances in new technologies, but also importantly, standards in areas such as fluids and rack/architectures. Operators — primarily enterprise operators — that have not deployed the technology should plan for flexible thermal infrastructure that can scale with silicon complexity, and address heat removal at the system level to optimize efficiency and TCO. Skilled partners will be required to implement reliable design, faster deployment, and long-term resiliency and performance.

[8] Dell'Oro Group, 2025.
[9] Yu. M., Zhang, C., Feng, J., et. Al., 2025. Materials Science and Engineering.
[10] Ebermann, H., 2024. Vertiv.

Summary

The impacts of AI are accelerating: rapid transformation due to increased densification, fast gigawatt-scale expansion, silicon diversification, and the emergence of data centers as computing units. The implications for data centers ripple across design, power, thermal management, IT systems, services, and software.

Vertiv anticipates organizations will face these challenges by adopting technology such as digital twins, on-site power generation, higher voltage DC distribution, advanced liquid cooling, and new form factors for AI deployment. In addition, Vertiv remains committed to enabling stakeholders to anticipate what's next—providing a lens into emerging technologies, guiding smarter investment decisions, and supporting the development of resilient, future-ready infrastructure.

Vertiv Frontiers aims to equip the data center community with the knowledge and perspectives to navigate a world defined by AI, advanced compute, and unprecedented innovation.

References

1. Bizo, D. (2024). Uptime Institute Intelligence. *AI to trigger radical overhaul of data center electrification.* https://intelligence.uptimeinstitute.com/

2. Dell'Oro Group. (2025). *Thermal management and cabinet PDU & busway lead growth; liquid cooling set to quintuple by 2029, AI build-out to propel data center physical infrastructure market to $63.1 billion by 2029, according to Dell'Oro Group.* https://www.delloro.com/news/ai-build-out-to-propel-data-center-physical-infrastructure-market-to-63-1-billion-by-2029/

3. Ebermann, H. (2024). Vertiv. *Immersion cooling systems: Advantages and deployment strategies for AI and HPC data centers.* https://www.vertiv.com/en-emea/about/news-and-insights/articles/blog-posts/advancing-data-center-performance-with-immersion-cooling/

4. EPRI. (2024). *Powering intelligence: Analyzing artificial intelligence and data center energy consumption.* https://www.epri.com/research/products/000000003002028905

5. Gartner. (2025). *Gartner forecasts worldwide IT spending to grow 9.8% in 2026, exceeding $6 trillion for the first time.* https://www.gartner.com/en/newsroom/press-releases/2025-10-22-gartner-forecasts-worldwide-it-spending-to-grow-9-point-8-percent-in-2026-exceeding-6-trillion-dollars-for-the-first-time

6. IDC. (2024). *Worldwide spending on artificial intelligence forecast to reach $632 billion in 2028, According to a new IDC spending guide.* https://my.idc.com/getdoc.jsp?containerId=prUS52530724

7. Noffsinger, J., Patel, M., and Sachdeva, P. (2025). McKinsey & Co. *The cost of compute: A $7 trillion race to scale data centers.* https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

8. Shehabi, A., Newkirk, A., Smith, S., Hubbard, A., Lei, N., Siddik, M., Holecek, B., Koomey, J., Masanet, E., and Sartor, D. (2024). Lawrence Berkeley National Laboratory. *2024 United States data center energy usage report.* https://doi.org/10.71468/P1WC7Q

9. 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

10. Yu. M., Zhang, C., Feng, J., Sun, Q., Yang, C., Zeng, X., Chu, P., and Tian, Y. (2025). Materials Science and Engineering. *Liquid metal alchemy: Unlocking self-healing gallium-based materials for next-generation electronics.* https://doi.org/10.1016/j.mser.2025.101073

Frontiers

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