E-book

Bring Your Own Power & Cooling (BYOP&C): An integrated approach to accelerate grid-to-chip data center deployment

Integrated power, cooling, and heat-reuse strategies for overcoming the AI-era infrastructure deficit

Table of contents

Executive summary

Unoptimized power and cooling infrastructure creates a hidden capital problem: stranded capacity buffers that consume investment without contributing to usable output. As AI workloads push rack densities well past 25kW (reaching the triple digits)[1] and interconnection queues stretch beyond five years, the cost of that fragmentation compounds: AI data centers face grid queues, volatile AI workloads, and cooling parasitics that stall usable IT capacity and delay time-to-revenue.

Vertiv recommends its Bring Your Own Power & Cooling (BYOP&C) model, which addresses the challenges by unifying on-site generation and co-optimizing liquid cooling, battery energy storage systems (BESS), and heat-reuse strategies into a converged grid-to-chip architecture. This integrated approach delivers speed to power in months instead of years, reclaims up to 5-15% of stranded capacity for compute, tightens PUE, and stabilizes spiky AI workloads. Vertiv serves as an orchestrator for the BYOP&C model, synchronizing power and cooling into a single architecture and maintaining resiliency through unified controls, predictive telemetry, and lifecycle services.

This white paper outlines the complete system architecture with a technically grounded framework for addressing power supply-demand imbalances, interconnection constraints, and thermal challenges with insights tailored to the operational, regulatory, and strategic priorities of the audience below:

- **Data center operators, energy managers, sustainability officers, and technology decision-makers** seeking on-site power strategies.
- **Grid infrastructure experts and policymakers** evaluating the benefits and challenges of distributed, behind-the-meter solutions.
- **Engineers and researchers** focused on advanced power and cooling design.
- **Investors and corporate leaders** exploring scalable, resilient infrastructure options.

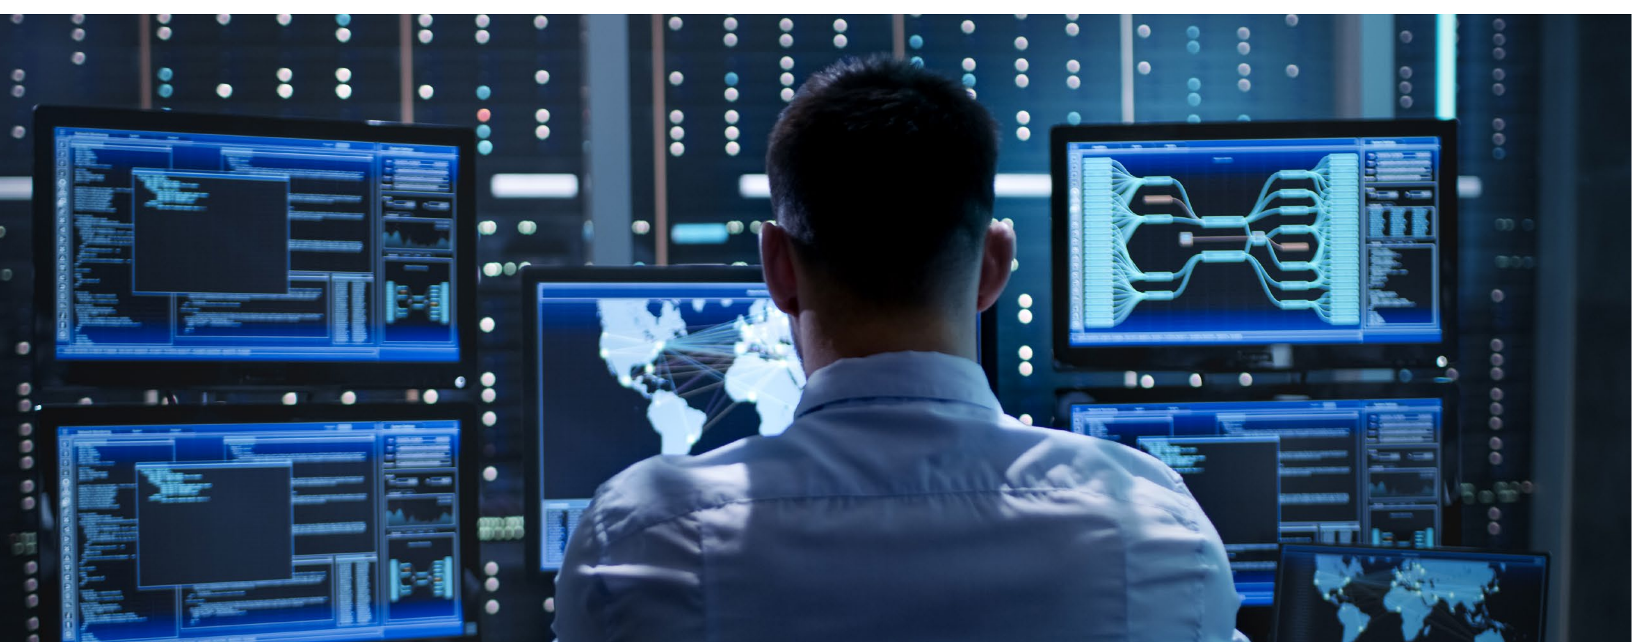

[1] Vertiv, 2026.

Introduction

Data center power demand is accelerating under the twin forces of cloud expansion and artificial intelligence (AI) workload growth. Cumulative new power capacity from 2023 to 2029 is expected to reach ~100 GW, with 13–20 GW added annually.[2]

AI-driven facilities alone could consume 945 terawatt hours (TWh) annually—more than Japan's electric consumption[3] —straining existing power infrastructure already constrained by regulatory and environmental limits on traditional diesel-based backup solutions.

Figure 1 underscores how this surge translates at the national level: U.S. data center power demand is projected to rise from 147 TWh in 2023 to more than 600 TWh by 2030, an increase of over fourfold in just seven years.[4] By the end of the decade, data centers could account for nearly 12% of total U.S. electricity use, placing them on par with major industrial sectors and amplifying the strain on existing infrastructure.

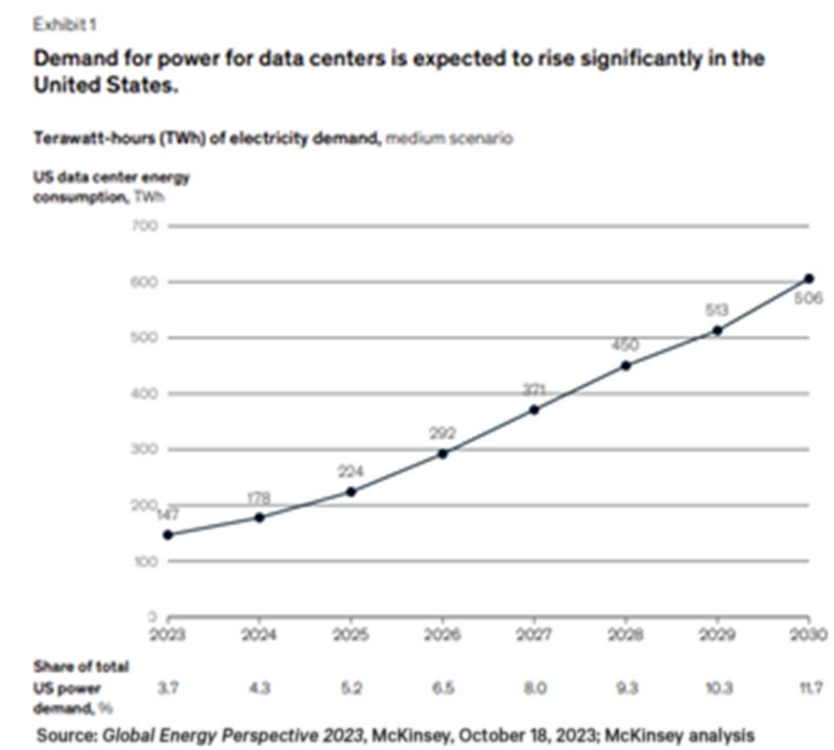

Figure 1. Power demand for data centers expected to rise significantly in the U.S.
Source: IEA, Goldman Sachs Global Investment Research, EIA

[2] Goldman Sachs Insights, 2025.
[3] International Energy Agency, 2025.
[4] Kok, E., Rauer, J., Sachdeva, P., et Al, 2025. McKinsey & Company.

Utility interconnection also remains a structural bottleneck for deployment. According to Lawrence Berkeley National Laboratory, median interconnection study durations for wind and solar projects now stretch two to three years, with interconnection request to commercial operations (IR-to-COD) timelines often exceeding five years.[5] Even battery and gas projects face mounting delays as queues grow.

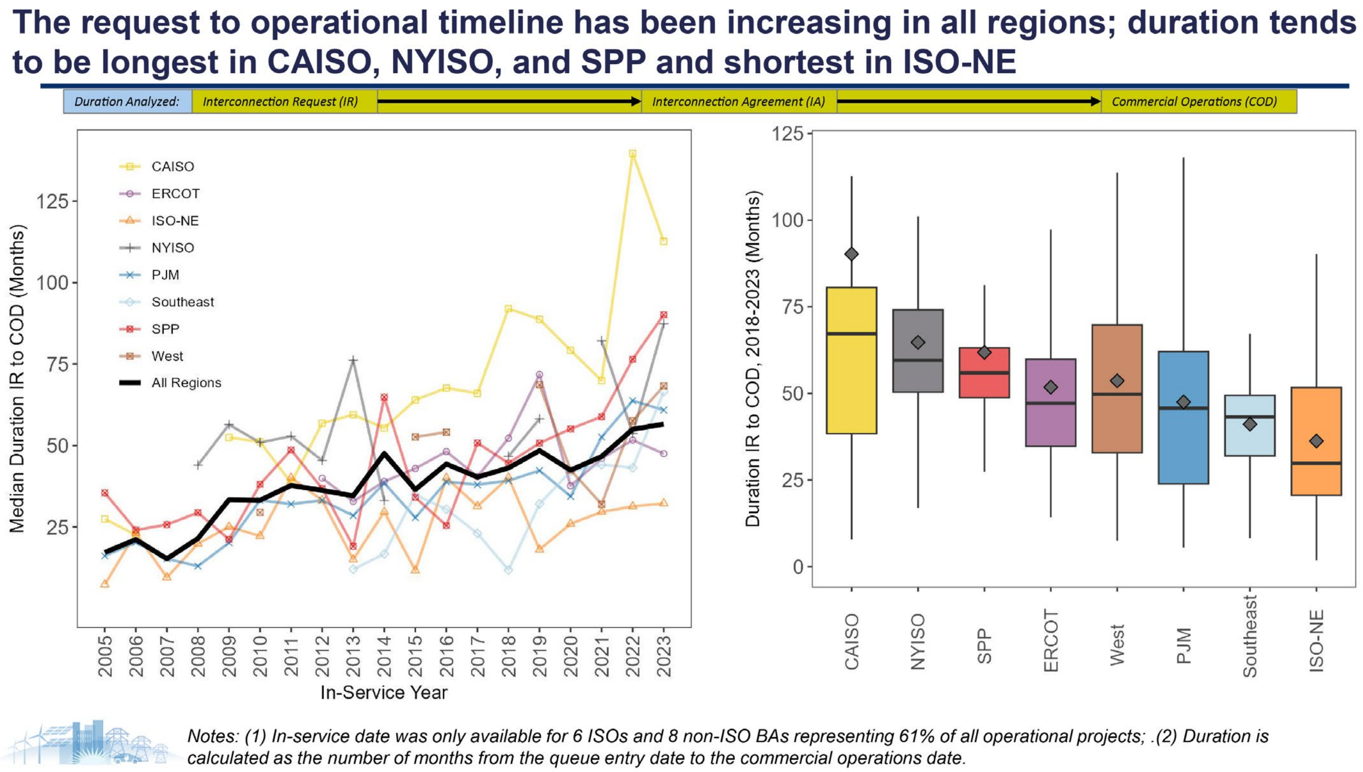

Figure 2. The median duration of IR-to-COD in months from 2005-2023

Source: Lawrence Berkeley National Laboratory

While AI-scale deployment is possible through grid-connected pathways, extended interconnection timelines can influence strategic planning. In scenarios where additional grid allotments are limited or study durations extend over multiple years, some developers explore remote or off-grid sites as complementary options. "Remote" refers to sites located away from congested transmission nodes, while "off-grid" describes locations powered mainly by on-site generation and storage.

Beyond power and site constraints, the AI era introduces a third pressure point: cooling. Rising rack densities make thermal management inseparable from power planning. The Uptime Institute's 2025 Cooling Systems Survey found that nearly 70% of operators have high rack density as the primary driver for liquid cooling adoption, citing that air cooling becomes inefficient beyond ~20 kW per rack.[6]

[5]Rand, J., Manderlink, N., Gorman, W., et Al, 2024. Lawrence Berkeley National Laboratory.

[6]Weinschenk, R., & Donnellan, D., 2025. Uptime Intelligence.

Together, these realities define the **three critical challenges of the AI era:**

1. Closing the power supply-demand gap created by surging data center workloads.
2. Overcoming grid and interconnection constraints tied to utility bottlenecks.
3. Managing thermal loads as rack densities exceed the limits of conventional cooling.

Vertiv's **Bring Your Own Power & Cooling (BYOP&C)** model addresses these challenges with an integrated grid-to-chip architecture that co-optimizes power and cooling for AI-scale deployment:

- **Power availability as the key constraint:** Bypasses multi-year interconnection queues with on-site generation.
- **Customer preference for integrated, end-to-end solutions:** Single integration partner for power generation, cooling, and controls across the full ecosystem.
- **Race against time to deploy AI-ready facilities:** Speed-to-power measured in months, not years.
- **Power usage effectiveness (PUE) as a critical parameter:** Co-optimized electrical and thermal design reduces conversion losses and tightens efficiency.
- **Efficiency goals:** Enabling lower emissions through waste heat recovery and reuse, higher operating temperatures, and modular scalability.

Bring Your Own Power & Cooling (BYOP&C): An integrated approach

The BYOP&C model rethinks data center deployment by unifying power generation and cooling. Instead of fragmented grid connections and siloed thermal systems, BYOP&C delivers a **grid-to-chip pathway,** converting fuel into reliable, conditioned power and cooling at the rack.

- **Integrated design:** On-site generation (e.g., natural gas, future nuclear, alternatives) is co-optimized with liquid cooling, electrical distribution, and UPS systems.
- **Deployment velocity:** By bypassing multi-year interconnection queues, facilities achieve speed-to-power in months.
- **Efficiency and resilience:** Coordinated electrical and thermal flows reduce conversion losses, stabilize inlet temperatures, and tighten PUE.
- **Scalability:** Modular blocks support AI-class densities and adapt to evolving energy sources and cooling technologies.

This integrated approach positions Vertiv as the enabler of **deployment-ready, grid-to-chip data centers** that closes the power supply–demand gap, overcomes grid constraints, and supports resilient infrastructure for the AI era.

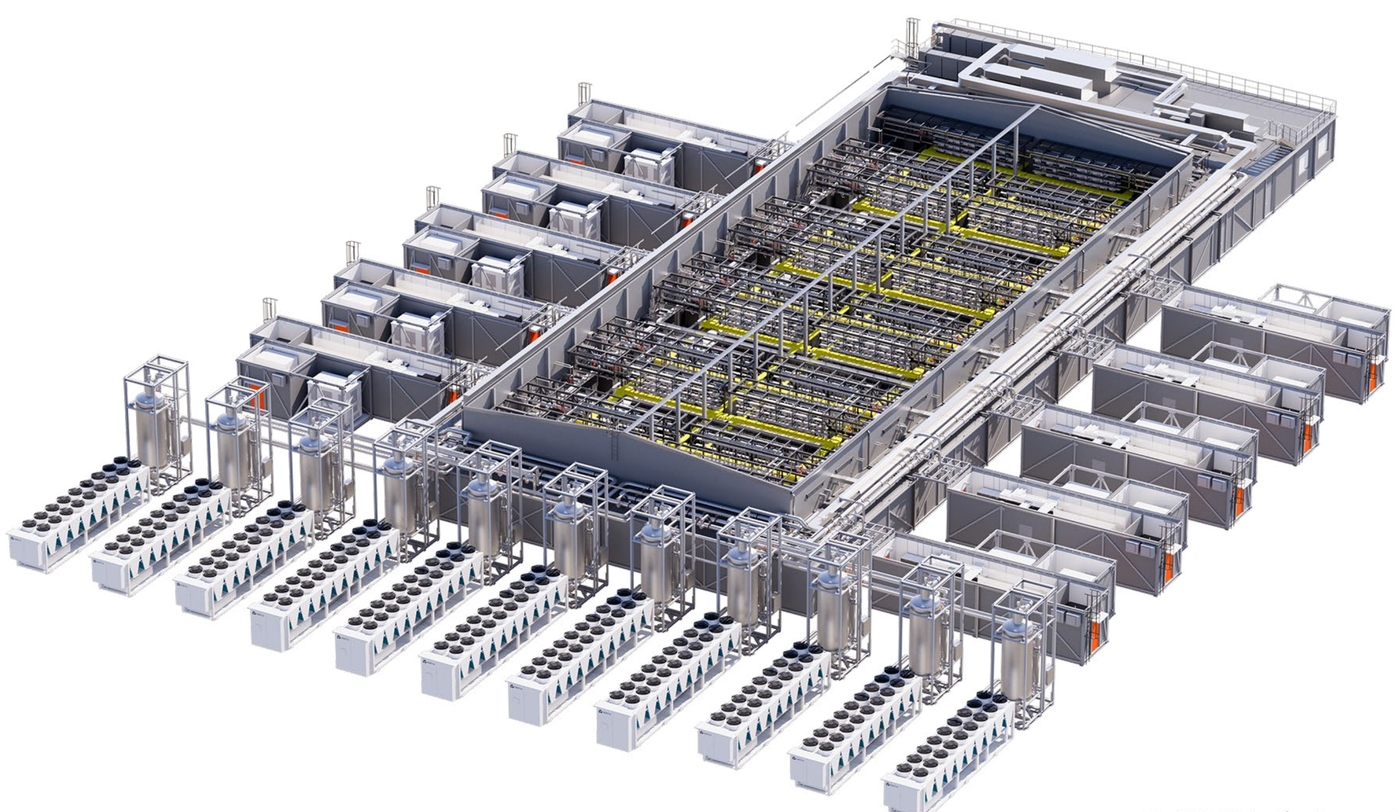

BYOP&C: Built on microgrids' features

On-site power generation is increasingly central to the data center strategy. Broadly, this solution falls under the **Bring Your Own Power (BYOP)** approach, where facilities deploy dedicated power generation assets to meet demand. Within BYOP, two primary configurations exist:

Microgrid (BYOP):

- Operate in parallel with the utility grid
- Provide grid services such as demand response or ancillary support
- Comply with grid codes and regulatory requirements
- Deployment timelines are often extended due to interconnection approvals

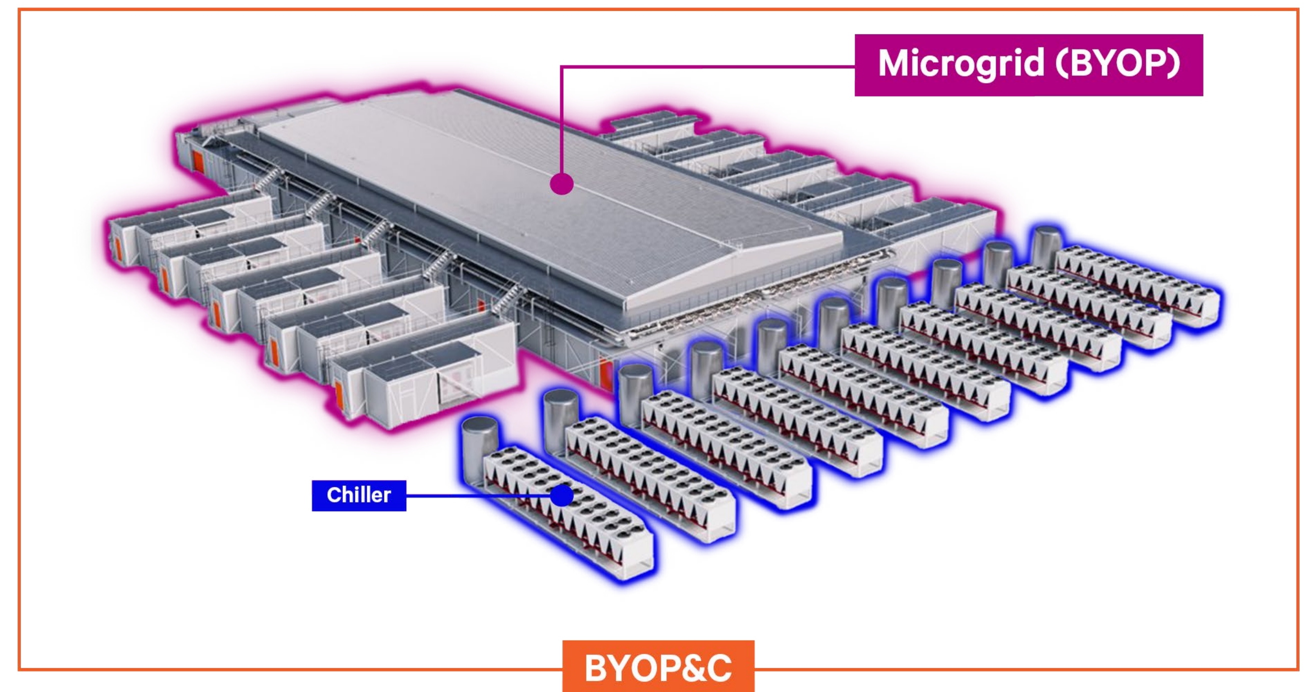

Figure 3. Expanding on the grid independence enabled by Bring Your Own Power (BYOP), Bring Your Own Power & Cooling (BYOP&C) integrates on-site generation with liquid cooling and heat-reuse strategies, accelerating deployment timelines, tightening PUE, and reducing stranded capacity across AI-scale facilities.

BYOP&C:

- Operates autonomously, independent of utility interconnection
- Enables faster deployment by reducing regulatory dependencies
- Integrates **power and cooling** into a single optimized solution
- Delivers improvements in both capital utilization and operating efficiency

While microgrids enhance reliability through distributed energy resources (DERs) and battery storage, BYOP&C extends the concept by embedding cooling into the architecture. This integration addresses rising power demand and thermal intensity, equipping facilities with power and cooling at scale.

Why BYOP&C now: From grid workaround to performance architecture

BYOP&C addresses grid queue concerns to facilitate speed-to-power in months, but its capabilities are beyond being seen as a temporary workaround for interconnection delays. The architecture is designed to support the realities of AI infrastructure scaling. Unoptimized designs force excess BESS and overhead cooling, eroding usable IT capacity. BYOP&C unlocks more capacity as organizations scale AI workloads.

Vertiv unlocks capacity in the BYOP&C model by serving as the orchestrator, integrating power, cooling, BESS, and controls into a single converged system. The performance gain comes from:

- **Heat capture and reuse:** Absorption cooling reduces parasitic electric load and reclaims up to 5-15% of stranded power for compute without additional grid allotments.
- **Coordinated controls:** Unified orchestration synchronizes generation, cooling, BESS, and IT demand.
- **Dynamic asset orchestration:** Predictive telemetry anticipates workload spikes, reducing reliance on oversized buffers.
- **Validated system integration:** Reference architectures and lifecycle service reduce risk and accelerate deployment.

BYOP&C is positioned not as a grid workaround but as a performance and capital efficiency architecture that scales with AI workloads — with Vertiv as the orchestrator, unlocking capacity and positioning organizations to benefit from a decisive advantage: transforming capital efficiency from survival engineering to optimization at scale.

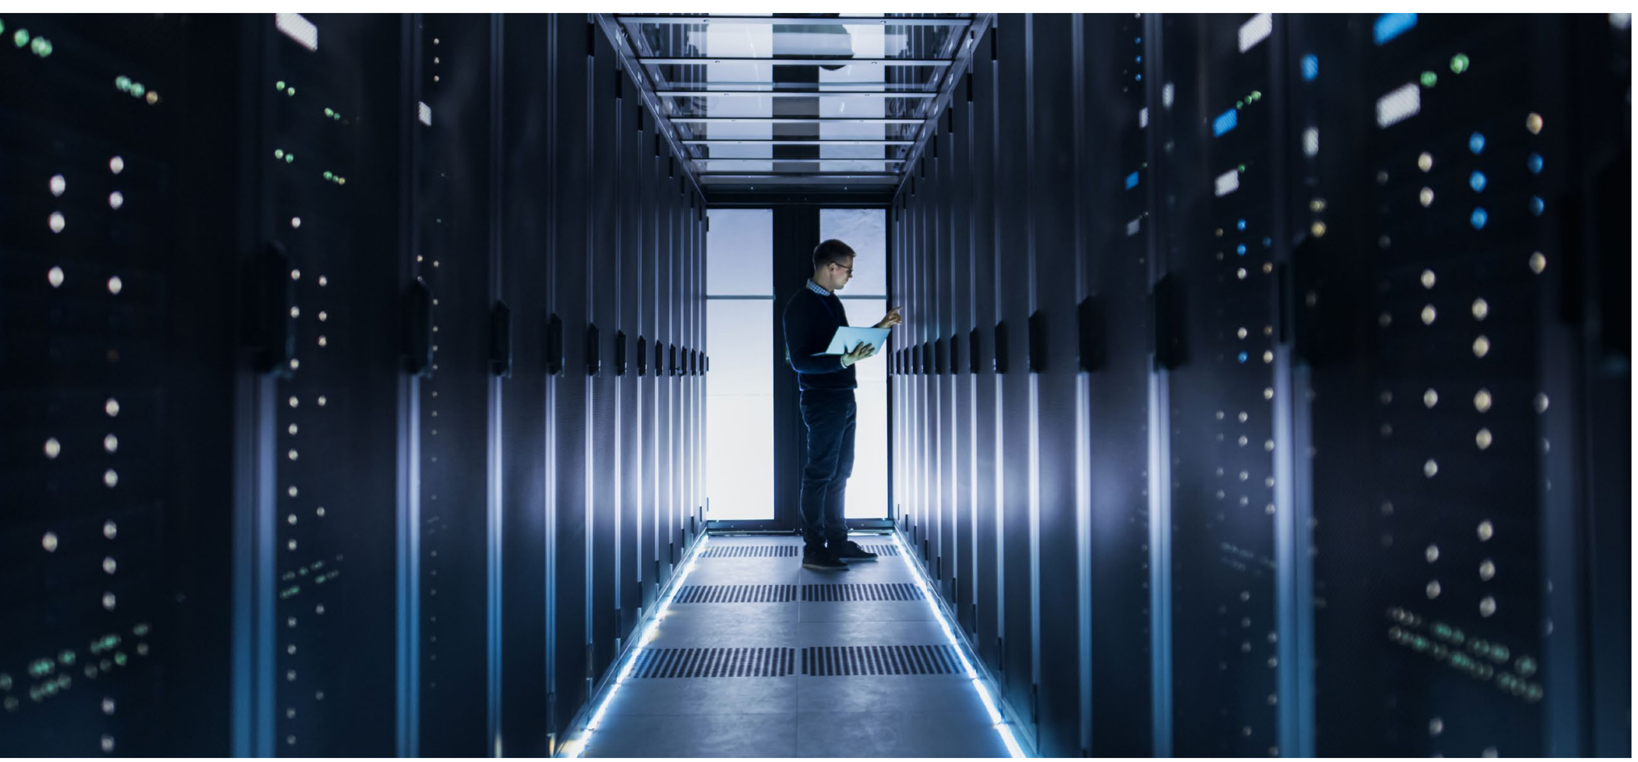

Understanding the BYOP&C ecosystem

There are many players within the BYOP&C ecosystem. The interdependencies between power and cooling, along with the reliability expectations of data center customers, necessitate significant collaboration across this ecosystem.

The BYOP&C ecosystem

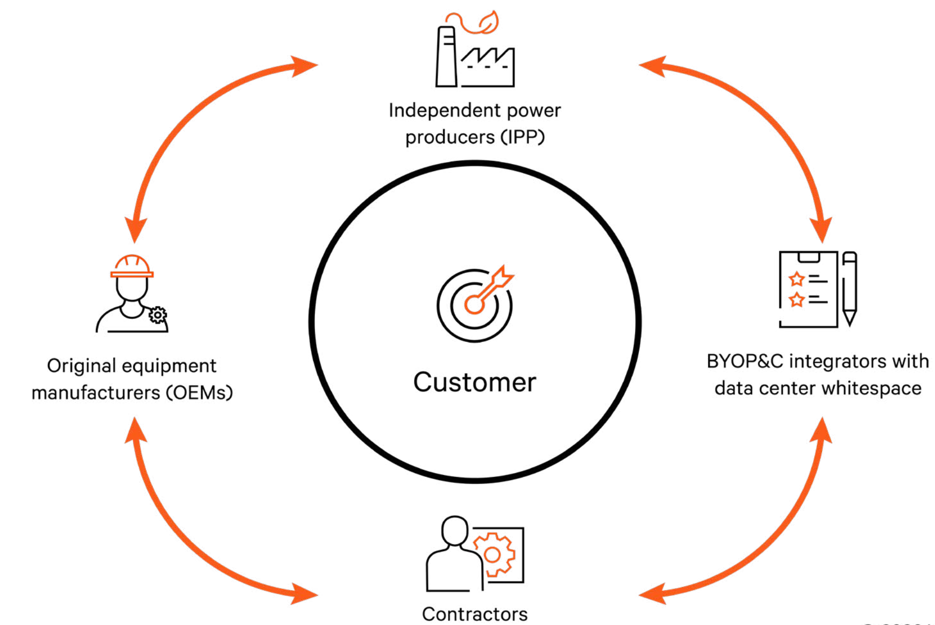

Figure 4. The key players with their respective roles in the BYOP&C ecosystem.

BYOP&C integrators with data center whitespace:

- Requirements analysis, including usage patterns
- Energy goal alignment
- Solution design customization based on the data center whitespace architecture
- Site layout optimization
- Performance simulation
- Technology integration roadmap development

Independent power producers (IPP):

- Power purchase agreement (PPA) execution
- Power plant ownership and operations
- Maintenance scheduling
- Original equipment manufacturer (OEM) service agreement management

Original equipment manufacturers (OEMs):

- Equipment manufacturing
- Capital expenditure considerations
- Lead time management
- Maintenance support provision

Contractors:

- Site construction management
- Permit acquisition
- Construction cost optimization
- Timeline management
- Efficiency maximization

Customers:

- Requirement specification
- Target setting
- PPA execution

The BYOP&C ecosystem requires coordination across integrators, independent power producers, OEMs, contractors, and customers. Each stakeholder plays a distinct role, and the value is realized when these contributions align into a unified, deployment-ready system.

To reduce engineering risk and accelerate deployment, the industry is moving toward **validated reference architectures—**modular, factory-integrated blocks that shorten design cycles from years to months. These architectures standardize the integration of turbines or reciprocating engines, battery energy storage systems (BESS), cooling, and controls into repeatable configurations.

Rather than treating each site as a one-off engineering project, BYOP&C leverages pre-engineered architectures that align with whitespace requirements, energy goals, and operational constraints. This approach enables rapid scaling, supports dynamic load responsiveness, and minimizes stranded capacity across diverse deployment scenarios.

The power layer: Power generation technology landscape

Choosing the right generation technology is central to BYOP&C deployment. AFCOM's 2025 State of the Data Center report notes that 62% of operators are exploring on-site generation and 19% have already implemented behind-the-meter solutions.[7]

The BYOP&C landscape will not be defined by a single technology, but by a mix of assets optimized for different roles. Natural gas engines and turbines provide reliable baseloads today. Larger steam turbines and industrial-scale reciprocating engines will expand near-term capacity. Nuclear microreactors are emerging as potential long-term complements, offering lower-carbon baseload potential.[8] Across all phases, BESS remains indispensable for transient stability, absorbing microsecond AI workload spikes that exceed turbine ramp rates.

This mix reflects the physics of deployment: some technologies excel at baseload, others at dynamic modulation, and others at stabilization. The winning strategy is **layered integration**, where each generation of technology asset builds upon the foundation of what came before.

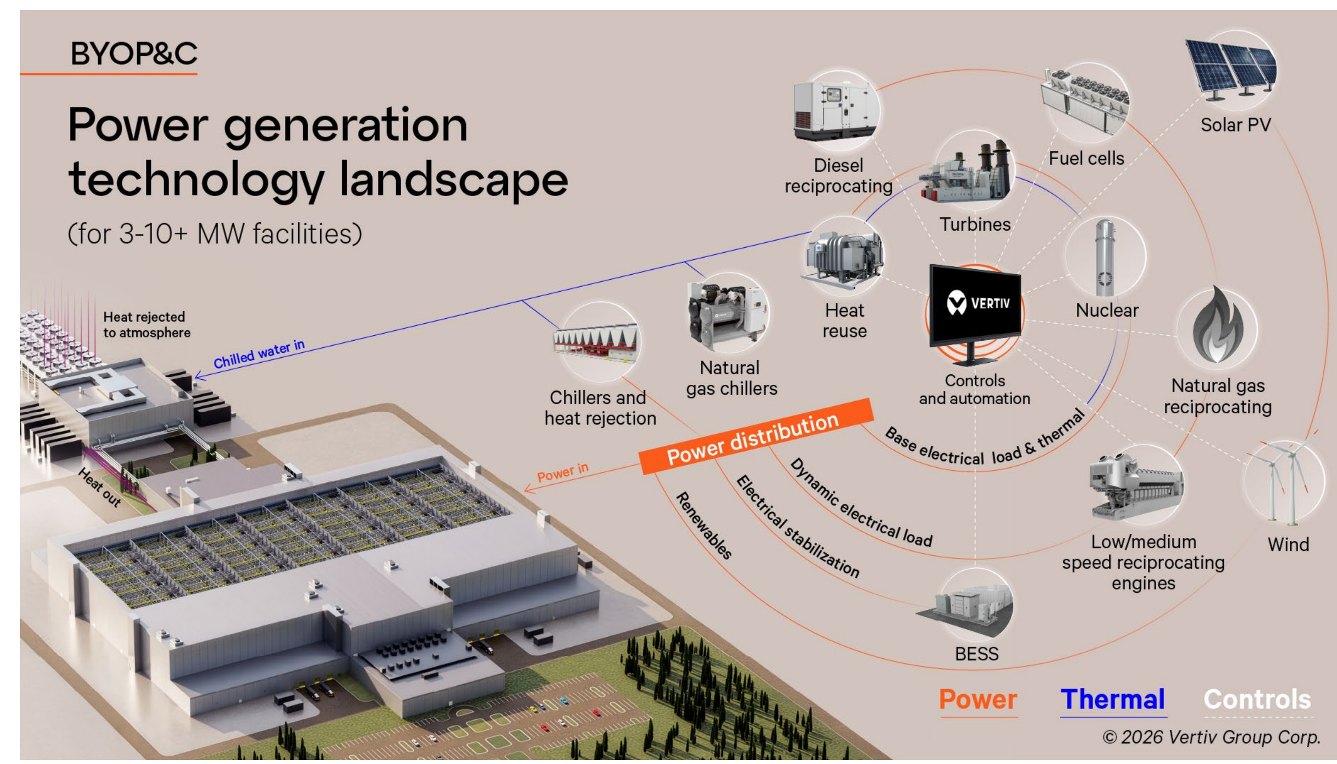

Figure 5: Power generation technology landscape for 3-10+ MW facilities with integrated power & cooling

[7] Kleyman, B., 2025. AFCOM.
[8] Vertiv, 2026

Application and timeline fit for large-scale solutions (3-10+ MW)

At 3-10 MW, BYOP&C is most immediately applicable to enterprise AI clusters, regional colocation facilities, and dedicated inference infrastructure. Grid connections remain an option at this scale, but interconnection delays routinely stretch deployment windows by years. Modular, factory-integrated configurations compress that timeline to months.

Large-scale deployments typically integrate:

- **Natural gas generators (gensets) with waste heat recovery (WHR):** Improve overall energy utilization by redirecting thermal output into cooling systems, reducing reliance on electrically driven chillers.
- **Battery energy storage systems (BESS):** Stabilize load, support grid interaction, and provide short-duration backup during demand spikes or transitions.
- **Future-ready assets:** Dual-fuel (hydrogen or H_2/natural gas or NG) gas turbines with WHR and BESS offer flexibility across energy sources and align with decarbonization pathways.

Battery energy storage systems (BESS) are not only short-duration backup assets, but also for critical stability components. In BYOP&C microgrids, BESS acts as a shock absorber, smoothing the micro-second spikes of AI workloads that exceed turbine ramp rates. This transient stability function sustains reliable operation under highly variable load conditions.

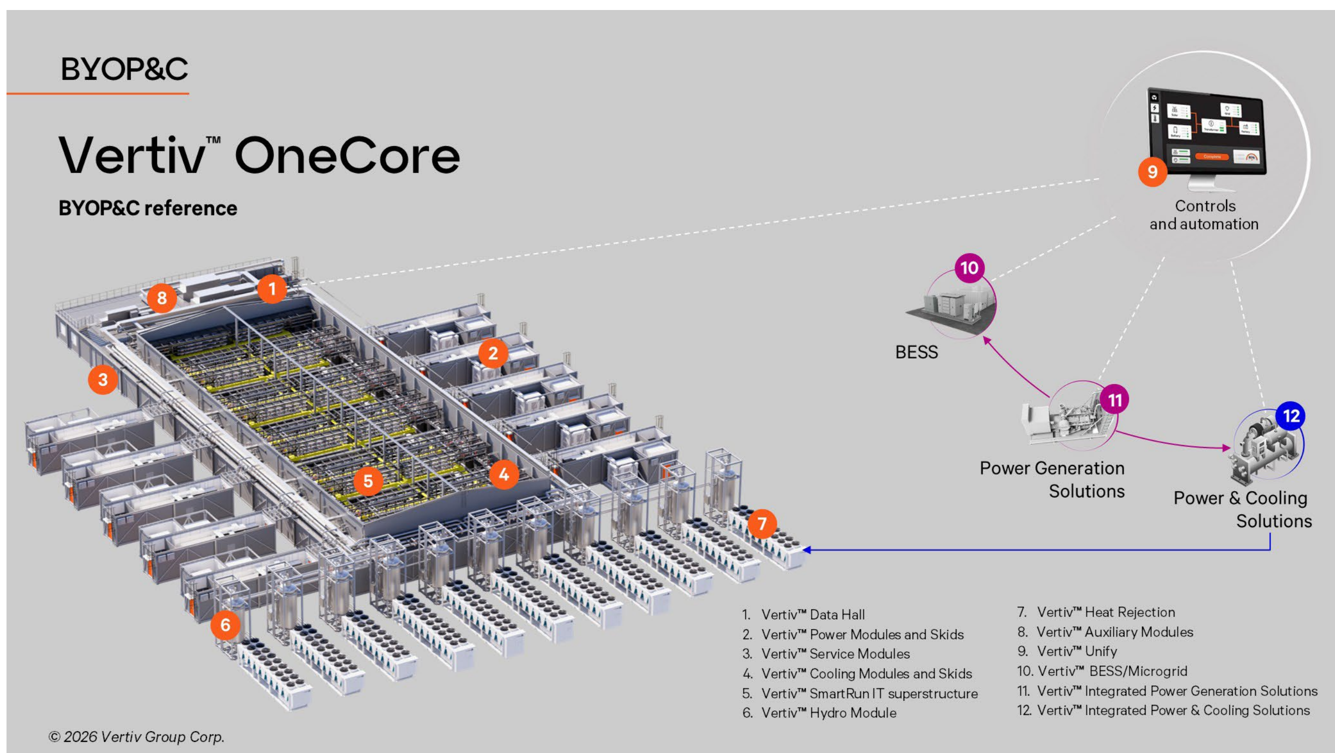

Figure 6: BYOP&C reference design with a Vertiv™ OneCore 12 MW hyperscale model

Readers can access further detail in the referenced white paper by scanning the QR code.

The schematic (Figure 6) illustrates a modular deployment approach for large-scale AI infrastructure integrating IT, power generation, energy storage, and cooling systems. The layout supports dynamic load responsiveness, thermal reuse, and phased scalability—key considerations for 100+ MW deployments and future-ready configurations.

AI workloads are highly variable, with power demand spiking during intensive training runs and tapering during idle or pre-processing phases. Large-scale solutions must therefore accommodate:

- Rapid ramp up and ramp down of demand without compromising stability.
- Integration with BESS to absorb short-term fluctuations and smooth transitions.
- Flexible generation assets capable of modulating output in real time.
- Waste heat recovery systems that remain effective across variable load conditions.

Vertiv has developed patented control solutions to address these AI load fluctuation challenges. These solutions use predictive telemetry and orchestration logic to smooth micro-second spikes, align generation and cooling assets, and reduce reliance on oversized buffers.

Timeline considerations

As AI workloads surge, operators must balance immediate deployment velocity with long-term infrastructure goals. The BYOP&C model supports a phased transition strategy, enabling facilities to deploy now while preparing for future energy landscapes.

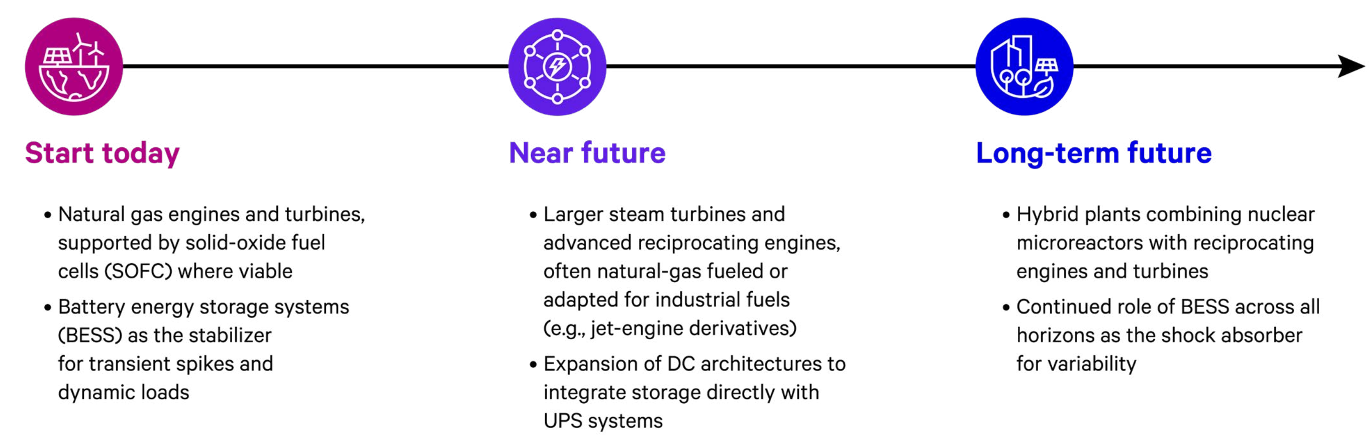

Figure 7: Timeline considerations for the phased transition strategy to the BYOP&C model.

Through the choice of power generation technologies, operators can address the dual challenges of supply-demand imbalance and deployment timelines. However, cooling integration remains a co-equal priority, making efficiency gains tangible and preparing facilities for the thermal realities of AI-class workloads. This interplay between power and cooling defines the transition point to the next consideration: why integrated solutions are the logical path forward.

The thermal layer: Why an integrated power and cooling solution is the way to go

Electric chillers consume material power, creating parasitic overhead that erodes usable IT capacity. In high-density AI data centers, separating power and cooling amplifies this burden[9], forcing duplicated infrastructure and excess BESS. Integrated approaches align generation, storage, and cooling within a unified architecture, reducing auxiliary load, improving responsiveness to dynamic workloads, and enabling waste heat reuse.

Thermal unlock: Reclaiming capacity through waste-heat-driven cooling

BYOP&C integrates heat reuse directly into the architecture. High-grade turbine exhaust drives absorption chillers, while lower-grade chip heat synergizes with liquid cooling loops. Together, they reduce reliance on electrically driven compressors and transform waste heat into usable cooling capacity. The benefits gained from this approach includes:

- **MW reclaimed for compute:** Heat reuse reclaims up to 15 of stranded power, converting it back into usable IT capacity
- **Reduced electric chiller load:** Less compressor demand lowers parasitic overhead and tightens PUE.
- **Lower operational expenses (OpEx):** Reduced auxiliary load and stabilized inlet temperatures extend equipment lifecycles.
- **Potential BESS right-sizing:** Smoother thermal response reduces dependence on oversized buffers.

BYOP&C supports both internal and external heat reuse. Inside the boundary, waste heat drives absorption chillers, reducing reliance on electric cooling and lowering demand. Outside, skids and pumps elevate thermal output for community or industrial use. Absorption cooling anchors internal reuse, while external export is viable when conditions permit.

[9] Zhang, L., Zhao, Z., and Chen, B., et al. 2025. DPI.

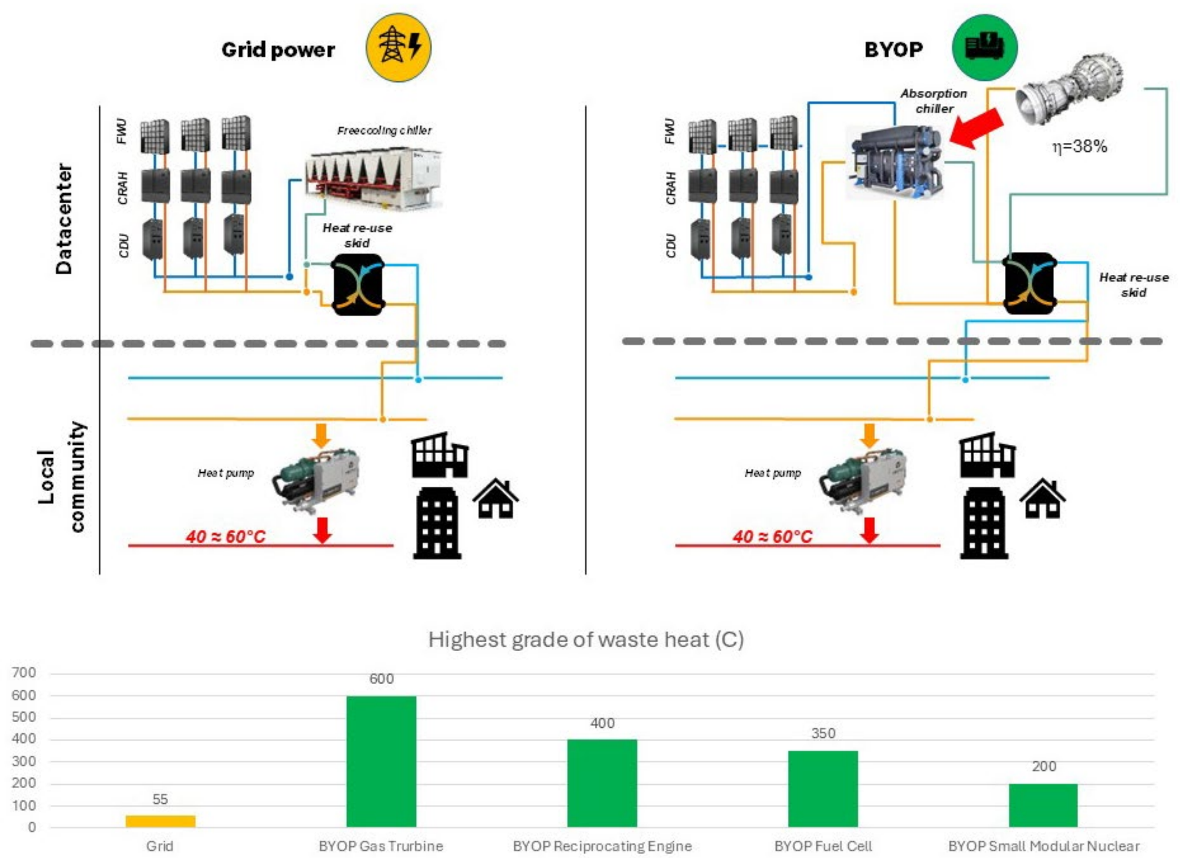

© 2026 Vertiv Group Corp.

Figure 8: Internal and external heat reuse for the BYOP&C model.

The outcome is clear: reclaimed capacity, reduced parasitic load, and lower operating cost. But these gains depend on specific boundary conditions that determine how effectively heat reuse can be deployed. Operators must evaluate the grade of available heat, system configuration choices, site factors, and deployment context for the architecture to deliver its full performance unlock.

- **Heat grade requirements:** Turbine heat is high-grade and ideal for absorption chillers; chip heat is lower-grade but aligns with liquid cooling strategies.
- **Configuration dependencies:** Single-loop vs. dual-loop designs balance retrofit flexibility with greenfield optimization.
- **Ambient/site factors:** External export is viable where district heating or industrial demand exists.
- **Deployment context:** Staged adoption[10] allows operators to start with internal absorption cooling and extend to external reuse when conditions permit.

BYOP&C embeds heat reuse as a performance lever, not an optional add-on. Properly integrated and orchestrated, it enables operators to approach nearly 1 GW scale for compute, transforming capital efficiency from survival engineering to optimization at scale.

[10] "Staged adoption" refers to a phased deployment approach in modular data center design: beginning with internal absorption cooling to manage immediate thermal loads, then expanding to external heat export (such as district heating) when infrastructure and market conditions permit. This incremental model reduces upfront complexity while enabling scalability over time.

The orchestration layer: Why controls determine whether BYOP&C works

The effectiveness of BYOP&C depends not only on hardware but on how electrical and thermal systems are coordinated. AI workloads are spiky and unpredictable, while mechanical generation assets respond slowly. Cooling systems and BESS must react in lockstep or risk stranded buffers and wasted capital. This mismatch (fast IT demand versus slower physical infrastructure) is the friction point that defines whether BYOP&C succeeds.

Vertiv unlocks capacity by serving as the orchestrator of the BYOP&C model. The control plane is the mechanism that converts assets into performance. Unified orchestration synchronizes generation, cooling, BESS, and IT demand, utilizing every watt and every thermal loop contributes directly to usable compute.

Feed-forward telemetry elevates this further. Instead of reacting to load changes, predictive signals from IT workloads inform proactive control, smoothing transients before they occur.

The outcome is improved stability, lower dependence on BESS for avoidable transients, tighter sizing of assets, and higher utilization across the system. This orchestration layer is where Vertiv unlocks the effectiveness of the BYOP&C model, serving as the glue that protects generation from load and load from generation, and positioning BYOP&C as a performance and capital efficiency architecture that scales with AI workloads.

The Vertiv layer: Why partner scale and lifecycle service matter

Building and operating BYOP&C at AI scale is not only a technical challenge. The larger risk is execution: whether systems can be delivered on time, integrated correctly, and supported over decades of operation. Sightline Climate's latest Data Center Outlook warns that up to 50% of the 2026 pipeline may not materialize, with up to 5 GW of the 16 GW announced capacity currently under construction. Around 11 GW has no visible construction progress despite being announced to come online in 2026.[11] Operators, insurers, and investors look for confidence that projects won't stall or fail.

Vertiv works within a validated ecosystem and combined service networks that extend from commissioning through lifecycle support. This approach reflects how integration and service coordination are handled across complex deployments.

Key elements include:

- **Supply chain coordination:** Established sourcing and logistics practices help keep projects aligned with schedules.
- **Field service footprint:** Global coverage with local presence supports rapid response and ongoing maintenance.
- **Commissioning and maintenance coordination:** Unified calendars and predictive service models prevent conflicts and extend asset lifecycles.
- **Validated interface ownership:** Clear responsibility for integration points across generation, cooling, BESS, and IT demand avoids gaps between vendors.
- **OEM and generation partner governance:** Structured alliances maintain interoperability and compliance across the ecosystem.

For operators, these practices support project continuity, investor confidence, and long-term adoption. BYOP&C is a performance architecture, provided that it is supported by partner scale and lifecycle service that make it practical to deploy and maintain at a global scale.

[11] Sightline Climate, Data Center Outlook, 2026.

Future outlook and recommendations

The evolution of data center power and cooling technologies is accelerating, but the trajectory remains uncertain. Operators must balance immediate deployment needs with long-term efficiency goals, recognizing that bridge solutions today must coexist with emerging technologies tomorrow.

Future assets must be provisioned as part of a mix, not a single technology. Some excel at baseload, others at dynamic modulation, and others at stabilization. Planning for compatibility across control systems, cooling architectures, and modular designs reduces retrofit costs and meets efficiency targets without compromising reliability.

The future of data center infrastructure will be defined by uncertainty in technology evolution and regulatory environments. Success will depend on adopting integrated, flexible designs that co-optimize power and cooling while providing end-to-end support across design, deployment, and operations. This approach enables operators to remain resilient while advancing toward reduced carbon operations.

Conclusion

BYOP&C is about **capital efficiency.** Treating power and thermal as one system eliminates stranded capacity buffers, making every investment contribute directly to usable compute. Moving from concept to deployment requires a clear sequence of actions that address both immediate risk and long-term scalability.

Operator action path:

1. **Quantify utility and interconnection delay risk:** Measure exposure to grid queue timelines and assess the impact on project schedules.
2. **Model IT load volatility and thermal profile:** Capture workload spikes, chip power density, and cooling demand to understand the real operating envelope.
3. **Simulate converged BYOP&C architecture options:** Include absorption cooling and heat reuse scenarios to evaluate parasitic load reduction and reclaimed capacity.
4. **Validate controls and failover logic:** Test orchestration across power, cooling, and BESS to confirm stability under spiky AI workloads.
5. **Select partners with established integration and global service support:** Working within a validated ecosystem helps align commissioning, lifecycle service, and interface ownership for smoother adoption.

Vertiv and its partners provide resources to help operators evaluate, design, and implement BYOP&C strategies that meet immediate demands and future objectives.

For further details on BYOP&C strategies and reference designs, visit the Vertiv BYOP&C resource center.

References

1. Goldman Sachs Insights. (2025). *AI to drive 165% increase in data center power demand by 2030.* https://www.goldmansachs.com/insights/articles/ai-to-drive-165-increase-in-data-center-power-demand-by-2030.

2. Graham, S. (2025). Vertiv. *Data Center Vision: How data center infrastructure will evolve to support AI and accelerated compute.* IDC Research. https://www.vertiv.com/49437d/globalassets/documents/reports/data-center-vision-how-data-center-infrastructure-will-evolve-to-support-ai-and-accelerated-compute_us53192425-ib.pdf.

3. International Energy Agency. (2025). *Energy and AI: International Energy Agency report.* https://www.iea.org/reports/energy-and-ai/executive-summary.

4. Kleyman, B. (2025). AFCOM. *State of the data center 2025: Executive summary.* https://afcom.com/news/692525/WHITEPAPER--State-of-the-Data-Center-2025-Executive-Summary.htm.

5. Kok, E., Rauer, J., Sachdeva, P., Pikul, P., Sutton, D., & 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.

6. Rand, J., Manderlink, N., Gorman, W., Wiser, R., Seel, J., Mulvaney Kemp, J., Jeong, S., & Kahrl, F. (2024). Lawrence Berkeley National Laboratory. *Queued up: 2024 edition – Characteristics of power plants seeking transmission interconnection as of the end of 2023.* https://emp.lbl.gov/sites/default/files/2024-04/Queued%20Up%202024%20Edition_R2.pdf.

7. Remelli, A. (2024). Vertiv. *Vertiv's AI power load management solutions for data centers.* hhttps://www.vertiv.com/en-us/about/news-and-events/articles/educational-articles/vertivs-ai-load-management-solutions-for-data-centers/

8. Vertiv. (2024). *Redefining efficiency: How and why data centers are embracing heat reuse.* https://www.vertiv.com/en-us/about/news-and-insights/articles/blog-posts/redefining-efficiency-how-and-why-data-centers-are-embracing-heat-reuse/

9. Vertiv. (2026). *Vertiv Frontiers 2026 Report: Navigating the future of AI-scale infrastructure.* https://www.vertiv.com/48d902/globalassets/content---assets-2025/documents/vertiv-frontiers-2026-report-en-gl-web.pdf.

10. Wang, Olivia (2026). *Data center outlook: half of 2026 pipeline may not materialize.* Sightline Climate. https://www.sightlineclimate.com/research/data-center-outlook.

11. Weinschenk, R., & Donnellan, D. (2025). Uptime Intelligence. *Uptime Institute cooling systems survey 2025: DLC adoption remains slow and steady.* https://intelligence.uptimeinstitute.com/sites/default/files/2025-07/UI%20Field%20181_Data%20center%20cooling.pdf.

12. Zhang, L., Zhao, Z., Chen, B., Zhao, M., & Chen, Y. (2025). DPI. *Zero carbon development in data centers using waste heat recovery technology: A systematic review.* https://www.mdpi.com/2071-1050/17/22/10101.

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

SL-80525 (R05/26)