eBook

FROM EQUIPMENT TO ECOSYSTEM

Embracing System-Level Simulation for Data Center Projects

The Short Version

- In the next five years we will create as much data center compute as was created over the past eight decades—a $6.7 trillion project to double data center capacity by 2030.
- The data center design and build sector is shifting toward mass production and product-based engineering, increasing the need for collaborative, simulation-based design platforms that are deeply integrated into the supply chain.
- Component suppliers who embrace these tools can more fully participate in the emerging data center ecosystem, offering ongoing services and long-term value to their customers while positioning themselves to better anticipate future developments in their fields.
- With an established track record in product-based engineering, Modelon is ready to help component manufacturers collaborate with the data center ecosystem, support their customers, and ensure that their R&D is pursuing profitable new technologies.

Table of Contents

3 How AI Is Rewriting the Rules of Data Center Infrastructure

4 How Data Center Productization Alters the Role of Component Suppliers

5 Using Collaborative System-Level Modeling to Improve Long-term Relationships

6 Using System Simulation to De-Risk Innovation

7 Design Faster. Operate Smarter. Compete at System Scale with Modelon Impact.

How AI Is Rewriting the Rules of Data Center Infrastructure

We're at the beginning of one of the largest infrastructure investment supercycles in human history. By 2030, it's estimated that we will have added 100 gigawatts (GW) of new data center capacity at a cost of roughly $6.7 trillion[1], effectively doubling global capacity.[2] Put another way, we're going to build as much data center capacity in the next five years as we did in the previous 80 years[3].

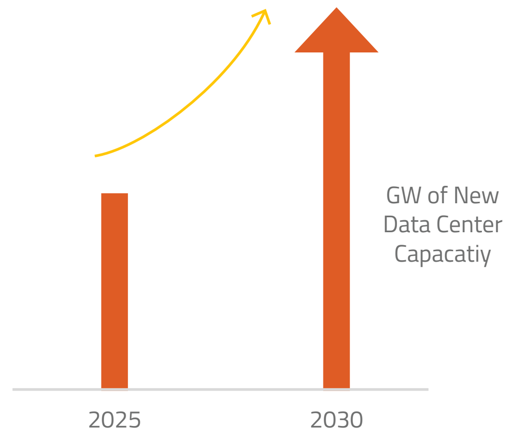

Until recently, the standard data center was an air-cooled 25 to 50 MW facility intended to host cloud services and SaaS for small and mid-sized enterprises. Now hyperscalers and other data center builders are prioritizing 100 to 500+ MW campuses meant to accommodate high-density liquid-cooled AI clusters [4].

The stakes are climbing higher with each data center evolution. If a failed chiller sends an older data center offline, they may lose half-a-million dollars in business over the course of a one-hour outage. If that same chiller fails in a new AI data center, they can lose a $3 million rack of GPUs in 60 seconds. These realities are fundamentally changing how data centers are designed, constructed, tested, commissioned, operated, and upgraded.

Designing and building data centers is rapidly evolving away from bespoke construction, and toward mass production. Throughout 2025 we saw a string of mergers and acquisitions that support this transition to product-based engineering emphasizing vertically integrated, full-stack solutions. This is further underscored by major players, like NVIDIA, who are shifting focus away from delivering standalone GPUs, and toward integrated rack-scale compute and the "AI Factory": prefabricated modules that arrive on-site tested, validated, and ready to plug in and power up.

Where do component suppliers fit into the data center ecosystem when data centers are less like construction projects and more like mass-produced products?

[1] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

[2] https://www.jll.com/en-us/insights/market-outlook/data-center-outlook

[3] https://blog.enconnex.com/data-center-history-and-evolution

[4]https://edgecore.com/2025-data-center-market-trends-report-building-for-whats-next/

How Data Center Productization Alters the Role of Component Suppliers

The ecosystem surrounding data centers design and construction is rapidly growing to resemble the automotive and aerospace industries—sectors that are similarly safety- and reliability-oriented. In these sectors, the entire supply chain is tightly integrated using simulation tools that support comprehensive physics-based system-level simulation, where an entire project is modeled end-to-end. This tight integration enables large-scale collaboration among designers, manufacturers, OEMs, and suppliers. Having a single shared model doesn't just streamline the initial design and creation of the product. That same system-level model often lives on as a digital twin, which can be fed data collected during real-world operations, then used to refine and optimize the product, as well as to guide future development of new components and systems.

Most parts of the data center ecosystem already internally embrace physics-based modeling in some way. Longtime proponents, like Meta, have used system simulation to reliably improve operations at their existing air-cooled data centers, using 32% less energy and 80% percent less water than the industry standard. [5]

But these simulation tools are not broadly integrated into the data center supply chain, making true collaboration a challenge. Component manufacturers remain siloed, and cannot enjoy the benefits of system-level insights, nor make more advanced contributions, assisting their customers in ongoing optimization and future planning.

By embracing system-level simulation on a collaborative platform now, component suppliers aren't just able to more quickly design and optimize solutions. They can also use these same comprehensive models as the basis for offering ongoing support services—like streamlined predictive maintenance—and to more confidently plan for what the next generation of technology should look like.

[5] https://engineering.fb.com/2022/09/14/data-center-engineering/data-centers-meta-thermal-simulation-optimization/

Using Collaborative System-Level Modeling to Improve Long-term Relationships

Component suppliers are increasingly focusing on digital engineering and after–sale services. The benefit to the supplier is clear: in businesses with a longer business cycle, like heavy-industrial manufacturing, it pays to have an additional revenue stream that is both higher margin and less volatile. But it's also good for their customers: predictive maintenance alone reduces unplanned downtime by 30% or more. Having the tools to foresee faults before they occur, rather than waiting for an equipment failure or performing maintenance on a set schedule regardless of real-world conditions, gives the customer long-term reliability while also furnishing the supplier steady revenue on a schedule they can anticipate and manage.

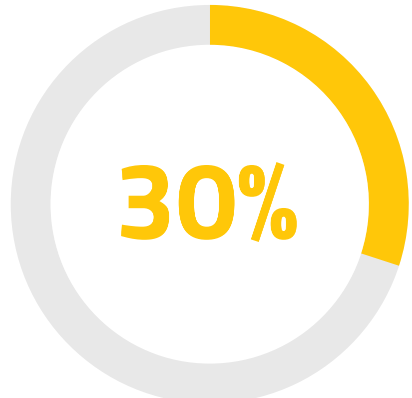

Unplanned downtime reduced by 30% or more with predictive maintenance.

Alfa Laval—a global heavy industry component supplier specializing in heat management, cooling, and fluid handling—readily acknowledges that growth in their "Connected Services," analytics, and predictive maintenance offerings is now a cornerstone of achieving their annual growth targets. In recent financial reports they've noted that their services are both higher-margin offerings, and have grown at a faster pace than sales, now accounting for more than 30% of their revenue [5.5]. That growth is driven by offerings like predictive maintenance and analytics services, which are possible because of the degree to which Alfa Laval has embraced system-level simulation. They maintain robust digital twins of the equipment and systems they have in the field, using live sensor data from that equipment to simulate equipment performance in real-time and accurately model wear.

Alfa Laval is also increasingly using their system-level simulation capabilities to run Energy Efficiency Simulations, allowing customers to see how upgrading specific equipment or changing control strategies can lower costs and fuel consumption while shrinking carbon footprints. Having accurate, flexible, mutli-fidelity digital twins is the key to being able to continue to grow these lucrative long-term support relationships.

Having accurate, flexible, mutli-fidelity digital twins is the key to being able to continue to grow these lucrative long-term support relationships.

[5.5] https://www.alfalaval.com/contentassets/3c19aad1e4a6431c9406e7e96247ecf5/wkr0006.pdf
https://mb.cision.com/Main/905/4142632/3418532.pdf

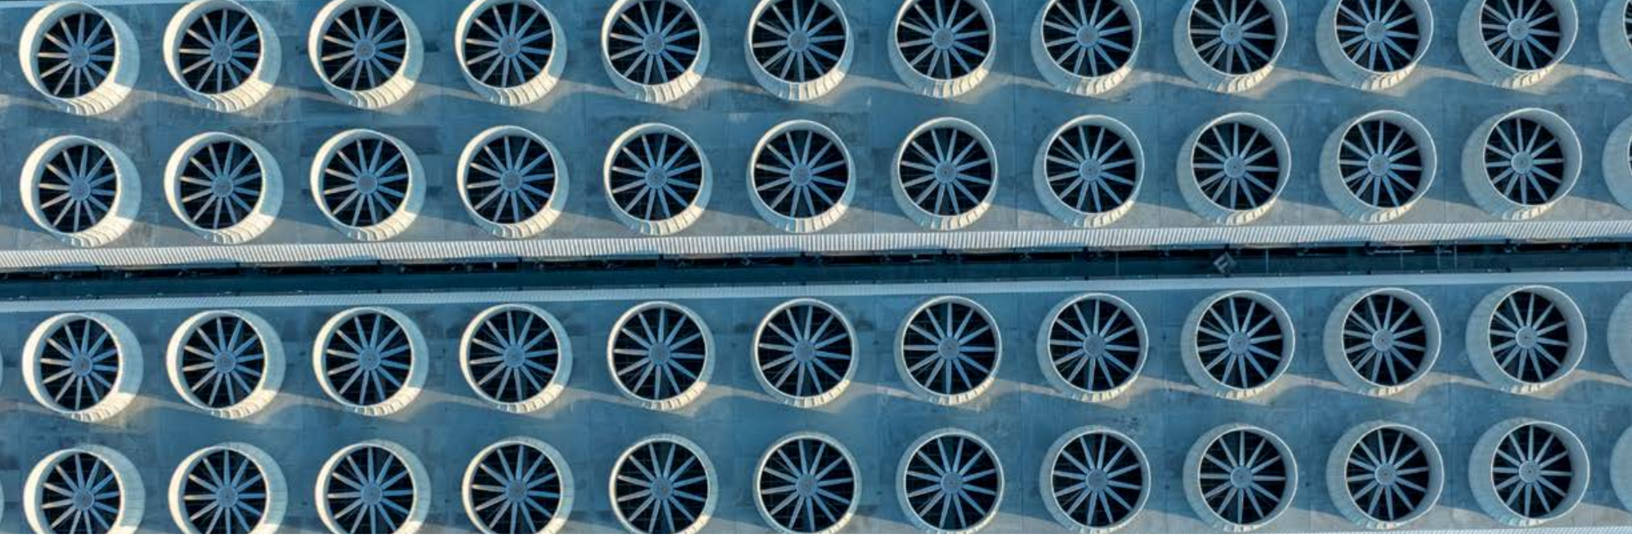

Using System Simulation to De-Risk Innovation

System simulation isn't just a design and engineering tool. It's also an invaluable resource for forward-looking component suppliers as they make strategic business decisions in determining which technology is the right one to develop, and when it is wisest to do so.

Imagine you are an established cooling technology OEM considering your next move. You know that traditional air cooling has reached its limits in data centers. Do you want to get ahead of the pack and roll out a two-phase cooling solution now? Or will single-phase liquid cooling remain viable long enough to justify developing that product instead?

It is a difficult question because there are so many trade-offs. Single-phase liquid cooling is simpler, with lower upfront costs and easier management and maintenance. Two-phase adds complexity (pressure control, fluid costs) but wins for heat transfer efficiency. It's ideal for the high chip power densities that are at the heart of AI/HPC data centers. When properly sized, configured, and optimized two-phase cooling can even potentially outperform simpler liquid-cooling solutions, in terms of overall PUE.

That said, moving prematurely carries tremendous risk: two-phase flow is a substantial change in cooling architecture, touching elements throughout the system, not just the components in direct contact with the cold plate and the chip. New two-phase cooling technology must manage boiling, condensation, pressure, and refrigerant distribution, with more complex control and safety concerns than with water loops.

- What are the implications for how you cool the refrigerant?
- Do you need a smaller or larger chiller?
- How must you change the control strategy for that chiller?
- Can you drop the chiller all together?
- How might that, in turn, change the design and size of the cooling tower?

With system simulation you can quickly explore the landscape of possibilities and get a sense of what avenues are worth pursuing. Then, during the design process, you can iterate through thousands of variations and see the ramifications of each design choice under real operating conditions. Once you narrow down your design candidates, you can insert them directly into the digital twin of an existing data center and show your partners and customers exactly what the wins will be and the future will look like.

Design Faster. Operate Smarter. Compete at System Scale with Modelon Impact.

Through our collaborative system simulation platform and support services, we're helping teams use physics-based modeling to streamline design, expand their after-sale and support services, and create digital twins that support both in-depth technical analysis and strategic business planning.

Modelon Impact is the AI-compatible, system-level simulation environment of choice for engineers, purpose built for complex systems, including liquid cooled infrastructures like data centers. It brings together cloud-native multi-physics modeling with vendor performance data, real-world operating conditions, your proprietary internal data and models, and industry-proven component libraries (like Modelon's Liquid Cooling Library and the Modelica Buildings Library).

Modelon Impact sits above existing component tools, coherently integrating thermal, hydraulic, and control domains to accurately capture the dynamic behavior of data centers and their cooling systems. It supports an iterative development workflow that enables smooth progression from conceptual design to control verification, validation, commissioning, optimization, and continuing upgrades and development. By relying on an open, extensible modeling system, Modelon Impact can represent real equipment and controls using industry validated models while integrating with existing engineering workflows that support cross-team collaboration.

Data center teams are already using these tools to evaluate control logic performance and improve energy efficiency and system controls resilience. In one 2024 study, physics-based simulations using the Modelica Buildings Library resulted in alternative cooling system control strategies that were 25% more efficient annually, and 18% more energy efficient during summer operation in Sacramento, CA.[6] Liquid cooled data centers are likewise using Modelon Impact to reduce energy use by 25%, create cooling solutions that are 300 times more water efficient, and cut development time in half while avoiding millions of dollars in risk[7].

The monumental demand for new high-performance compute has highlighted the need to collaborate throughout the data center ecosystem. It has further emphasized the importance of truly robust digital twins that model the actual physics and operational ground-truth of a complete system at the appropriate level of fidelity and complexity. Embracing these tools now is helping component manufacturers play a more meaningful role in using their expertise to optimize these systems and predict system performance as conditions change.

Take a Quick Tour of Modelon Impact

[6] https://www.tandfonline.com/doi/pdf/10.1080/23744731.2023.2276011

[7] updated against: https://modelon.com/industries/data-centers/

Move beyond components.
Compete at system scale.

As data centers shift toward product-based engineering and tightly integrated ecosystems, component manufacturers must go beyond standalone performance. Modelon enables teams to simulate, validate, and optimize components within the full data center system—unlocking faster innovation, stronger collaboration, and new service-based revenue streams. The result is not just better components, but a more strategic role in the evolving data center ecosystem.

- Strengthen OEM and hyperscaler partnerships with system-level performance insights
- Transition from product supplier to ecosystem partner with long-term value creation
- Design and validate components in the context of full data center systems, not silos

See Our Data Center Solutions