INTELLIGENCE UPDATE

AI-accelerated digital twins: lessons for operators

John O'Brien 28 Jul 2026

Many data center operators use management and control software — such as DCIM, building management systems (BMS) and supervisory control and data acquisition (SCADA) — to safeguard infrastructure resiliency and efficiency. Yet few today use engineering simulations in their day-to-day operations, let alone digital twins to manage the facility's full lifecycle. Historically, this was due to the cost and complexity of setting up and running computationally intense modeling. The business value was also unclear, as data centers underwent little change for long periods and often operated well below capacity (for related links on digital twins, see the end of this report).

However, the cost of engineering simulations has plummeted while data center complexity has increased. As IT racks serving AI take up a greater share of infrastructure alongside ongoing densification of IT, more sophisticated software-based methods for design, commissioning and change management are becoming necessary.

AI and densification present many technical challenges — from managing increasingly varied IT environments with different power and cooling needs to implementing new electrical equipment and hybrid or direct liquid cooling.

In this report, we explore some of the emerging applications being tested and deployed across AI factories and modern grid, power and cooling infrastructure.

Domain-specific applications are emerging

Digital twin simulations can effectively model design and change scenarios — facility-wide or within specific domains — helping operators manage growing complexity and its associated operational risks.

In the Uptime Intelligence report *Digital twins: the role of simulations*, we described how digital twins and digital simulation technologies are becoming interdependent, providing a framework for applications across the data center (see **Figure 1**).

Figure 1 The data center digital twin simulation framework

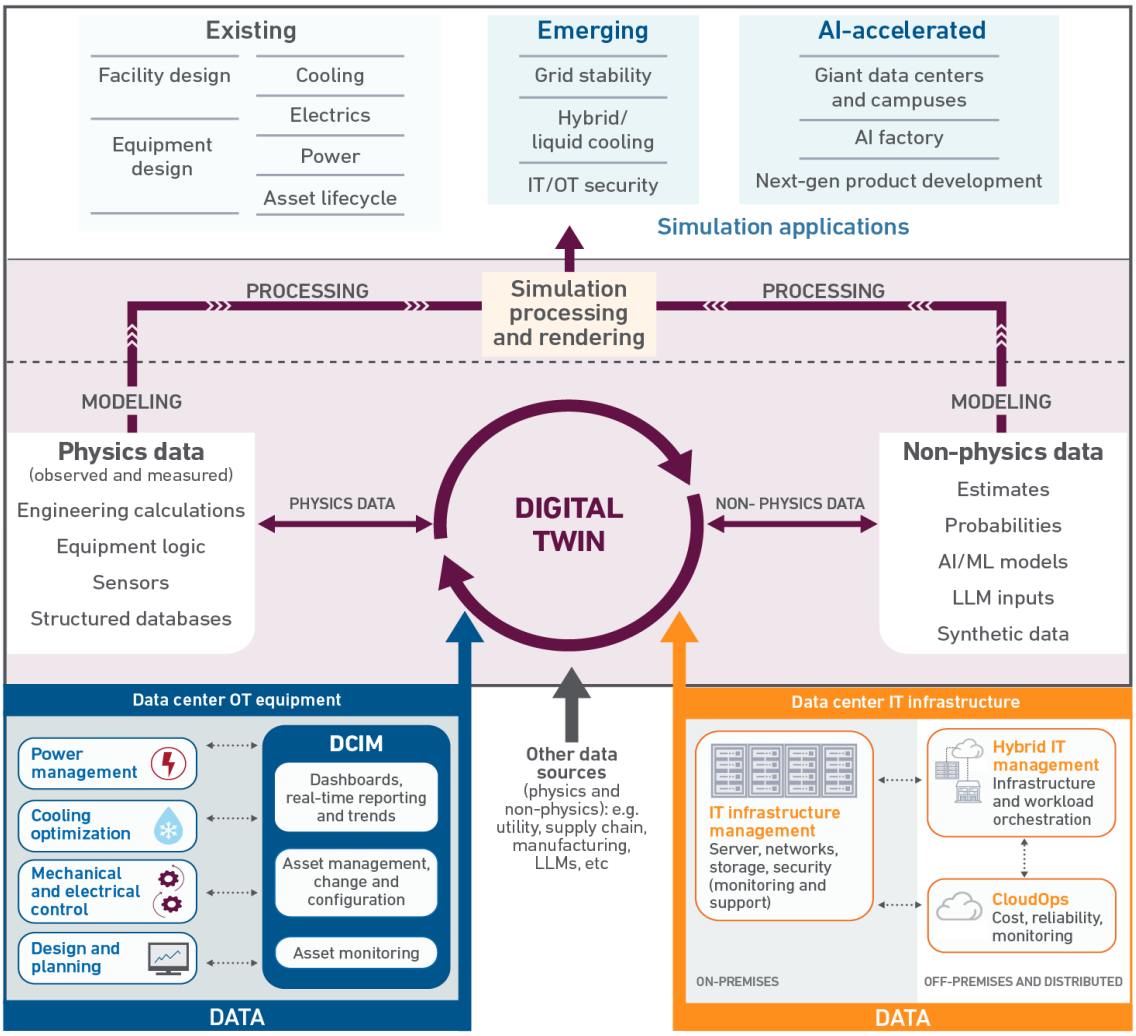

UPTIME INSTITUTE 2025

uptime INTELLIGENCE

AI-accelerated digital twins will use advances in computer processing GPUs and AI to render high-fidelity simulations. They have the potential to improve the understanding of complex system behaviors — provided they combine trusted physical data sources to perform AI modeling and inferences to support accurate predictions.

Overlaying lifelike interactivity within a virtual environment is likely to support collaboration between internal teams and external partners: particularly if this is adopted as a single pane of glass for designing, testing and iterating products and operations scenarios.

This approach could, for example, be transformative in identifying multisystem dependencies and risks, such as stressors on power, cooling and electrical systems caused by shifting workload patterns. Furthermore, modeling tools can reduce the time, cost and risk involved in delivering new data center capacity.

These objectives will resonate with many operators, especially those under pressure to improve lead times, costs and overall efficiency. At the same time, hype and vested interests make it difficult to identify which applications are genuinely worth considering beyond initial pilots.

Notably, Nvidia has taken a high-profile role in promoting digital twins for data centers with its Omniverse platform and Omniverse DSX Blueprint, closely linking them to its "AI factories"

concept — data centers purpose-built for dense IT systems serving AI workloads. However, Nvidia is by no means the first to offer such products. We have previously discussed alternatives, such as software as a service (SaaS) optimization tools that are less capital intensive.

Applications: some early examples

Several data center technology suppliers have partnered with Nvidia to develop digital twin solutions based on Omniverse in the past few years. A non-exhaustive list of vendors now contributing data directly into DSX Blueprint includes Cadence, Dassault Systèmes, Eaton, Jacobs, NScale, Phaidra, Procore Technologies, PTC, Schneider Electric (including its subsidiaries AVEVA and ETAP), Siemens, Switch, Trane Technologies and Vertiv.

Below are several promising applications where these digital twin solutions play a key role in design and/or product development. Power availability and reliability is currently the most active area of development. Any disturbance to power delivery will have an immediate impact on IT operations, while rapid fluctuations in large loads may adversely affect grid stability.

Power system resiliency

At an event in Milan, AVEVA (part of Schneider Electric) demonstrated the use of digital twins across multiple domains to investigate how the impact of GPU power profiles on both the grid and the data center resiliency can be mitigated. Schneider's ETAP digital twin models electrical performance from the grid to the chip, using dynamic simulations to analyze power flows and voltage stability to identify risks and predict potential faults.

In this use case, calculations are performed at very high precision to model real-world system behavior. For example, AVEVA's digital twin software analyzes GPU power data at half-cycle resolution to identify specific oscillation patterns under different AI workloads and power demands. The aim is to determine when and why power spikes occur as thousands of GPUs transition from one workload to another, since these rapid shifts in aggregate load can create power oscillations that affect grid stability. This modeling enables engineers to design and test power systems that can better mitigate these oscillations and reduce the risk of grid disturbances.

Medium-voltage power

ABB, Eaton, Siemens and Vertiv are all developing medium-voltage (typically 10 kV AC and above) power systems to support high-power-density AI infrastructure. This includes 800 V DC architectures, switchgear, UPS systems and battery storage. Integrated approaches are also emerging, including those developed by Eaton, Siemens Energy (low- to medium-voltage prefabricated skids), nVent (liquid cooling and electrical systems), Fluence (battery energy storage) for Nvidia Vera Rubin GPU deployments. The aim is to fully design, build and assemble flexible, modular medium-voltage infrastructure off-site before deployment on-site, simplifying installation time, expediting time-to-market and potentially lowering operational risks.

Grid-responsive power consumption

Startup energy management software supplier Emerald AI has developed a novel approach to AI

power consumption, based on feedback from the grid.

Emerald AI's software receives telemetry signals directly from the grid operator relating to grid stresses, weather events and other operating conditions. It can then respond by prioritising available IT power for the most critical AI workloads. The platform relies on on-site battery storage to shift some or all of the facility's power demand away from the grid to the battery system.

Emerald AI has conducted five real-world demonstrations with commercial data centers including Nvidia and Oracle Cloud in Virginia (US), National Grid (Nebius, UK), and Portland General Electric (US).

Trial with National Grid in the UK

The National Grid demonstration provides some useful insights into the business value impact.

Over a five-day period in December 2025, Emerald AI's demonstrator ran on a 130 kW cluster of 96 Nvidia Blackwell Ultra GPUs at Nebius' AI Factory in London, UK.

• National Grid Electricity Transmission (NGET) and Electric Power Research Institute (EPRI) issued 22 live dispatch events, including zero-notice emergency signals.

• The cluster autonomously reduced power consumption by 30% in under 40 seconds in response to an emergency scenario. It also followed load-reduction requests for up to 10 hours.

• Over the five-day period, it maintained compliance across more than 200 power targets while preserving service level agreements (SLAs) for critical workloads. However, no information was provided on the impact to non-critical SLAs.

GPU thermal optimization

A pilot led by cooling optimization specialist Phaidra claims a 75-80% reduction in GPU thermal spikes on one of CoreWeave's Nvidia DGX GB200 clusters. One of the key benefits is that improved thermal stability can help reduce failures rates while allowing operators to raise the liquid coolant temperature. This reduces the cooling energy required for heat rejection, which in turn can free up power for AI compute. Uptime Intelligence provides a more detailed analysis of this pilot and specific security considerations in this report: *Real-time telemetry requires modern, flexible cybersecurity*.

Research by other vendors echoes these findings. CFD supplier Cadence and engineering firm NV5 performed digital twin simulations using Nvidia DGX GB300 racks (approximately 150 kW nominal thermal design power) to demonstrate the benefits and viability of supplying coolant to the cold plates at 45°C (113°F) —the highest temperature formally supported by Nvidia. Similar simulations also indicate that limiting GPU power may also improve overall system efficiency, delivering up to 17% more AI performance per watt.

Other considerations

At this stage, data is collected from static repositories such as computer-aided design (CAD), building information management (BIM), product lifecycle management (PLM) and configuration management system (CMS) software.

However, trusted sources of real-world operational data would be valuable to help validate the models and simulations. Much of this physical data resides in operational systems such as DCIM, BMS or SCADA software, which monitor and control live equipment.

Integrating real-world operational data could be used to validate the model under simulated conditions. It could also support more robust predictive modeling of existing and future states.

The Uptime Intelligence View

AI factory digital twins and simulations are currently focused on high-end AI data center infrastructure — and this is supporting the investment in emerging applications. The vendors involved have a clear commercial interest in promoting their software and services for the design, construction and delivery of AI systems.

These innovations nonetheless, provide important indicators for those owning and operating existing infrastructures, and for those providing software into the data center.

Platforms such as Omniverse DSX are likely to remain uneconomic for most operators, and the repeatable blueprint is unlikely to suit facilities that operate a heterogeneous mix of systems.

Filling this gap could be the next big software opportunity, particularly for those able to integrate these forward-thinking concepts into their products.

Other related reports published by Uptime Institute include:
Digital twins: the role of simulations
Digital twins: reshaping AI infrastructure planning
Nvidia's vision: digital twins and automated facilities
Digital twins and DCIM: why data quality must come first

ABOUT THE AUTHOR

John O'Brien

29 Jul 2026

John is Uptime Institute's Senior Research Analyst for Digital Infrastructure Software, covering applications for IT-OT management, monitoring, control, and operational AI. Prior to Uptime, John was a research director at IT analyst firms IDC, GlobalData and Ovum/Informa.

jobrien@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. With over 4,000 awards issued in over 122 countries around the globe, and over 1,100 currently active projects in 80+ countries, Uptime has helped tens of thousands of companies optimize critical IT assets while managing costs, resources, and efficiency. For over 30 years, the company has established industry-leading benchmarks for data center performance, resilience, sustainability, and efficiency, which provide customers assurance that their digital infrastructure can perform across a wide array of operating conditions at a level consistent with their individual business needs. Uptime's Tier Standard is the IT industry's most trusted and adopted global standard for the design, construction, and operation of data centers.

Offerings include the organization's Tier Standard and Certifications, Management & Operations reviews and assessments including SCIRA-FSI financial sector risk assessment, the Sustainability Assessment, and a broad range of additional risk management, performance, availability, and related offerings. Uptime Education training programs have been successfully completed by over 100,000 data center professionals, such as the much-valued ATD (Accredited Tier Designer) and AOS (Accredited Operations Specialist). The Uptime Education curriculum has been expanded by the acquisition of CNet Training Ltd. In 2023.

Uptime Institute is headquartered in New York, NY, with offices in London, Sao Paulo, Dubai, Riyadh, and Singapore, and full-time Uptime professionals based in over thirty-four countries around the world.

For more information, visit www.uptimeinstitute.com