

INTELLIGENCE UPDATE

Digital twins: the role of simulations



John O'Brien

21 Jul 2025

Digital twin software: Part 2

Enterprises and operators are facing complex power, electrical and cooling challenges as they strive to meet the requirements of giant data centers, large-scale AI and high-density infrastructures. Much of the expansion taking place is on a scale not seen before and therefore requires new thinking and innovative solutions.

At Nvidia's GTC event in Paris in June 2025, data center digital twins (DC-DTs) and digital twin simulations (DTSs) were highlighted as potential solutions to address the challenges of designing and managing high-density IT. By modeling known "physical" systems through virtual simulations and scenarios, these technologies may be able to resolve complex engineering and infrastructure challenges. However, there are still numerous obstacles and risks to address.

While data centers have long utilized simulations such as computational fluid dynamics (CFD), these remain costly and time-consuming to render. As a result, many operators consider them impractical and economically unviable beyond the initial planning and design phases.

These limitations have also affected digital twins. In many cases, they have delivered only static data analytics or basic visualizations and dashboards, rather than presenting interactive, real-time insights suitable for live operational settings.

Technological advances are emerging that could position digital twins and DTSs as critical components of the data center management and control (DCM-C) software toolkit. In a few short years, high-performance GPUs and supporting AI software have enabled the development of interactive simulations that are both fast and visually impressive. While these technologies may benefit specific large-scale applications, it is unclear whether they offer tangible advantages for the majority of data center operators — many of whom are looking to adapt and optimize existing infrastructure.

This report examines some of the advancements in data center digital twin simulations and assesses their potential strengths, weaknesses and applications.

What is a digital twin simulation?

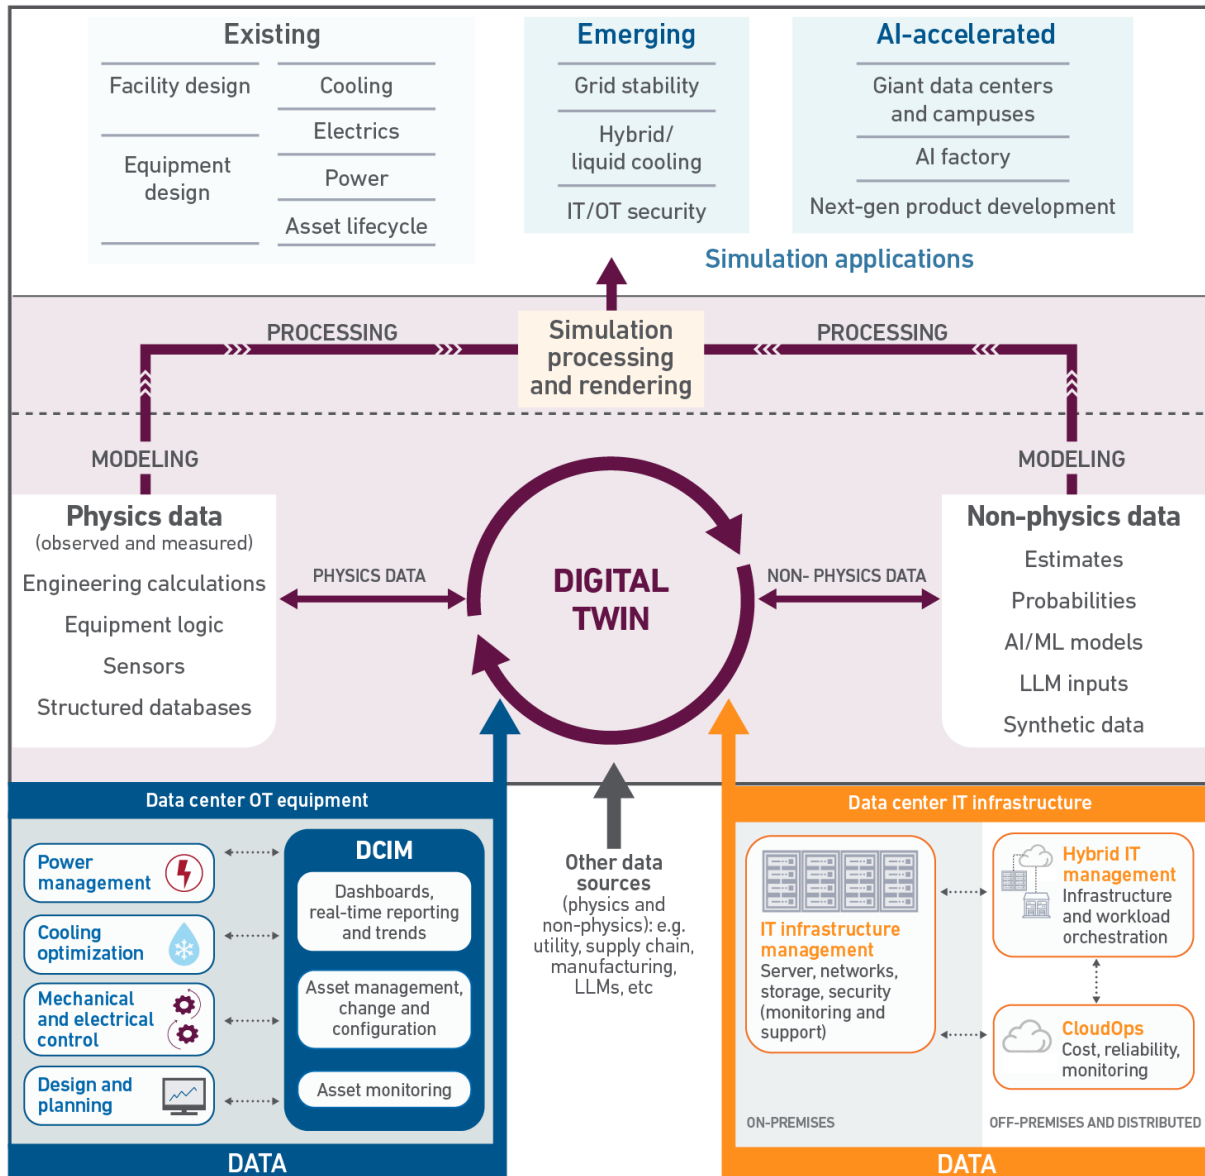
Digital twins and digital simulation technologies are becoming increasingly interdependent. A DC-DT is a software system that utilizes component libraries, software integrations and precision sensor data to create a digital replica of the physical assets within a data center (see [Digital twins: reshaping AI infrastructure planning](#)).

Digital twins employ simulations and models to test operating scenarios, make predictions and provide recommendations based on the virtual environment. They can help identify discrepancies between the virtual model and the physical facility — potentially revealing hidden faults.

DTSS combine digital twin technologies, data, visualizations (as discussed in Part 1) and interactive computer graphics such as 3D rendering and AR/VR environments.

Figure 1 provides a schematic of Uptime Intelligence's data center DTS framework. It has two core layers: the digital twin itself (located in the middle) and the simulation layer, which sits on top.

Figure 1 The data center digital twin simulation framework



UPTIME INSTITUTE 2025

uptime
INTELLIGENCE

Note, for a detailed analysis of the software categories shown at the bottom of this diagram, see [Data center management and control software: an overview](#).

Data from data center OT and IT systems can be supplied to the digital twin as follows:

- **OT** — Building management systems (BMS), SCADA and other programmable logic controllers (PLCs) can supply operational data such as power, electrical and water consumption, thresholds and settings, and fault and maintenance logs. Equipment sensors in the data halls can provide live information on environmental conditions — such as temperature, humidity and pressure — which can be valuable for interpreting operational data.
- **IT** — Data center infrastructure management (DCIM) software can provide IT asset and capacity management data, including information on IT servers, IT power, cabinets and system updates and changes. IT infrastructure management software may provide additional IT utilization, incident, fault and security logs. Hybrid IT management tools may provide useful third-party metrics related to service level agreements, key performance indicators and sustainability data from third-party

colocation and cloud providers.

How the framework operates

The digital twin acts as the data integration and data modeling engine. It connects to DCM-C software, equipment sensors and other external data sources (see bottom of **Figure 1**). Data is ingested via APIs and OT system protocols.

The simulation layer serves as the processing and rendering engine, which translates the digital twin model into an interactive visualization.

In high-performance computing — and increasingly in next-generation industrial twin applications — simulations are rendered using GPUs, which are much better suited to this task than CPUs.

Physics data is the foundation

Physics data and physics laws are foundational for DTS modeling and predictions. Physical data can be directly observed and measured — for example, the dimensions of a data center or the temperature of a cooling system. CFD, the most well-known DTS technology used in data centers, is a branch of fluid dynamics. It uses the Navier-Stokes equations to model the flow of gases and fluids through physical systems to determine the pressure (Pascals, Pa), density (kg/m^3), velocity (m/s) and viscosity (Pascal-seconds, Pa·s).

CFD can simulate how physical obstacles, the location of IT equipment and the operation of cooling fans might affect HVAC system performance. For over two decades, it has been used in data centers to help design cooling systems, optimize layouts and identify potential inefficiencies or points of failure.

CFD-based simulations can take many hours to process using traditional CPU-based IT infrastructure — but more suitable hardware is now available. For example, supplier Cadence revealed in Paris that a data center CFD cooling simulation running on Intel Xeon Gold 5120 CPUs takes 8 hours and 15 minutes to process. In contrast, the company demonstrated that a more advanced version of the same simulation could be rendered in under 10 minutes on Nvidia GB200 GPUs — 48 times quicker. This difference highlights why CFD has previously been impractical in live environments, and why GPU-based architectures may now offer a solution.

Prominent emerging applications

It is still early days, but advanced digital twin simulations are being applied to some of the biggest challenges that face operators when designing and building high-density data centers. Exploration is taking place in three critical areas:

- Understanding the requirements and risks around hybrid air- and liquid-cooled systems.

- Identifying and addressing power supply issues from both the grid and on-site systems when faced with an increase in demand.
- Analyzing and resolving electrical demand challenges associated with high-performance AI compute.

Liquid and hybrid cooling

AI facilities still typically require as much as 30% air cooling, even when the chips and components are liquid-cooled. Heat from surrounding circuitry can transfer to cabinets, requiring conventional rear-door heat exchangers. At the same time, many data centers continue to operate traditional workloads alongside AI, requiring a hybrid cooling approach that allocates the appropriate cooling method to the workload and IT infrastructure. CFD software, such as Cadence's Flow Network, models fluid flow through data center pipe networks to simulate the performance of CDUs, heat exchangers and liquid-cooled IT scenarios. It can also help explore potential issues, such as the failure of direct liquid cooling systems, to help determine the required size of the buffer tank in the event of a failure.

Power and electrical systems

It is now widely acknowledged that AI training clusters exhibit significant power fluctuations. Predicting the behavior of a specific compute cluster spanning several racks — and potentially hundreds or even thousands of GPUs — during a training run is a challenge.

At the same time, AI infrastructure power requirements are increasing rapidly, from around 125 kW per rack today to as much as 600 kW per rack by 2028, with some projections even suggesting 1 MW per rack system in the near future.

To address some of the unknown risks associated with these AI power demands, ETAP is developing a closed-loop simulation of GPU power behavior — from chip to grid — using Nvidia's Omniverse platform. This includes modeling on-site power generation (batteries and generators), UPS and active harmonic filters for AC and DC power conversions. The simulation also evaluates the resulting impact on GPU loads, power quality and stability.

The Uptime Intelligence View

For most organizations, deploying advanced digital twin simulations will remain impractical for the foreseeable future. This may change if the cost and complexity of deploying GPUs for rendering can be significantly reduced.

For many operators, physics simulations using CPUs are likely more feasible for use in design, planning and retrofitting applications. Lighter optimization software tools may be more practical in operational settings. As it stands today, next-generation digital twins and simulations will likely appeal only to the largest enterprises and operators, particularly those planning to build out multi-gigawatt campuses.

Other related reports published by Uptime Institute include:

[*Data center management and control software: an overview*](#)

[*Data center management software: optimizing the IT*](#)

[*Digital twins: reshaping AI infrastructure planning*](#)

[*Hold the line: liquid cooling's division of labor*](#)

ABOUT THE AUTHOR



John O'Brien

John is Uptime Institute's Senior Research Analyst for Cloud and Software Automation. As a technology industry analyst for over two decades, John has been analyzing the impact of cloud migration, modernization and optimization for the past decade. John covers hybrid and multi-cloud infrastructure, sustainability, and emerging AIOps, DataOps and FinOps practices.

[**jobrien@uptimeinstitute.com**](mailto:jobrien@uptimeinstitute.com)

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. Its Tier Standard is the IT industry's most trusted and adopted global standard for the proper design, construction, and operation of data centers – the backbone of the digital economy. For over 25 years, the company has served as the standard for data center reliability, sustainability, and efficiency, providing customers assurance that their digital infrastructure can perform at a level that is consistent with their business needs across a wide array of operating conditions.

With its data center Tier Standard & Certifications, Management & Operations reviews, broad range of related risk and performance assessments, and accredited educational curriculum completed by over 10,000 data center professionals, Uptime Institute has helped thousands of companies, in over 100 countries to optimize critical IT assets while managing costs, resources, and efficiency.