Scaling AI-Enabled Digital Twins Across the Data Center Lifecycle with OpenUSD

White Paper 302

Version 1

Data Center Research & Strategy

by Kevin Brown & Alison Matte

Executive summary

This paper explains why Schneider Electric is prioritizing engagement with open digital twin frameworks, specifically OpenUSD, and why we endorse a USD-based, AI-enabled digital twin platform approach exemplified by NVIDIA Omniverse. It describes how this architectural direction will build on many existing digital twin and infrastructure management solutions to enable a new class of AI-powered digital twins capable of supporting higher-confidence planning, design, construction, and operational decisions across the full data center infrastructure lifecycle.

RATE THIS PAPER ★★★★★

Key takeaways

1. Most data center digital twin and infrastructure management tools deliver value within narrow scopes: design, simulation, build, visualization, monitoring, alarming, and localized optimization. Their impact is constrained by proprietary data models, limited interoperability, low or inconsistent model fidelity, and fragmentation across systems and lifecycle phases. As a result, they function well as point solutions but do not scale into a single, persistent, facility-level digital twin.

2. NVIDIA Omniverse addresses these limitations through a different architectural approach. Built on OpenUSD and accelerated computing, it provides a shared digital foundation that allows models, data, and design intent to persist across planning, design, construction, and operations. Rather than replacing existing tools, it connects and extends them within a common scene that supports system-level reasoning, high-fidelity visualization, and AI-driven insights.

3. Fully unified, lifecycle-scale data center digital twins are technically feasible today but remain early in adoption. Maturity will depend on organizations aligning tools, data pipelines, and operating models around open, platform-based architectures. For this reason, companies such as NVIDIA, Microsoft, Intel, ETAP, AVEVA, and Schneider Electric (among many more) are actively supporting OpenUSD and AI-enabled platforms built on it.

4. For owners and operators, the value lies in faster, higher-confidence decision-making across the full site lifecycle. A shared system-level context reduces manual interpretation, surfaces risks/constraints earlier, and enables safer, more predictable change. All this will ultimately improve time-to-power, time-to-revenue, and total cost of ownership as infrastructure complexity increases.

Introduction

NVIDIA defines digital twins as “virtual representations of products, processes, and facilities used to design, simulate, and operate their physical counterparts.” In data centers, this capability provides a continuously accessible, risk-free environment for understanding system behavior and evaluating change before acting in the physical world.

Digital twin solutions for data centers are not new. However, existing offerings range from static 3D visualization models to physics-based simulations synchronized with live operational data. They vary significantly in scope, fidelity, synchronization, and intelligence.

Most solutions in use today struggle to keep pace with the scale, density, and throughput of modern data centers and AI factories. Fragmented models, limited interoperability, and constrained use of accelerated computing and AI prevent many technically feasible use cases from being realized.

NVIDIA developed Omniverse to address these limitations. It is a USD-based, AI-enabled digital twin platform designed to support persistent, system-level digital representations across the full data center lifecycle.

Note, while the capabilities and use cases described in this paper are technically feasible, with several already demonstrated in pilot and early deployments, the adoption of fully unified, lifecycle-scale data center digital twins is still in its early stages and expected to mature as organizations integrate data, tools, and operating models around these platforms.

To help encourage this development, this paper explains why Schneider Electric is prioritizing engagement with open, extensible digital twin frameworks, specifically OpenUSD, and why we endorse AI-enabled digital twin platforms built on an OpenUSD framework, as exemplified by NVIDIA Omniverse.

Terms used in this paper

To avoid ambiguity, the following terms are used consistently throughout this paper and are defined once here.

Scene

A scene is a structured digital representation of assets, systems, and their relationships assembled to describe a physical or logical environment.

In the context of a data center, a scene captures:

- What exists (assets and logical constructs)
- How it is connected (topology and dependencies)
- What it means (semantics and intent)
- How it behaves (operating states and constraints)

A data center scene may represent an asset or subsystem (e.g., UPS, electrical distribution), a facility, or an entire site and serves as the shared context for simulation, analysis, and AI.

Universal Scene Description (USD / OpenUSD)

USD is an open framework for describing and composing complex digital scenes. It defines how assets are identified, structured, related, and assembled into a shared representation that can evolve over time.

OpenUSD is the open, governed specification stewarded by the Alliance for OpenUSD, ensuring USD can be implemented consistently across vendors, tools, and industries.

In the context of data center digital twins, OpenUSD provides a common scene foundation that allows electrical, mechanical, IT, and facility systems authored by different tools and vendors to coexist and remain interoperable across lifecycle phases.

Semantics

Semantics refer to the shared, machine-readable meaning of what an asset is, how it behaves, and how it relates to other assets.
In a digital twin, semantics ensure that a UPS, breaker, cooling unit, or logical construct (such as a redundancy zone or operating mode) means the same thing across tools, simulations, AI models, and lifecycle stages.

Preserving semantics is what allows models to be reused, combined, and reasoned over without reinterpretation or manual translation.

USD representation vs. SimReady representation

A USD representation is the minimum OpenUSD-based description required for an asset or system to exist within a shared scene. It defines identity, structure, spatial context, and relationships, but alone it does not enable simulation behavior.

A **SimReady representation** builds on a USD representation by adding validated semantics, physical properties, and behavioral information required for physics-based simulation and AI-driven analysis.

While not every asset must be SimReady initially, most of the differentiated value of USD-based, AI-enabled digital twin platforms emerges when assets are enriched with simulation-ready (SimReady) representations. These representations allow assets to participate directly in physics-based simulation and AI-driven analysis, enabling higher-confidence decisions, automation, and system-level insight.

Note, developing a SimReady representation of a complex infrastructure asset requires substantial engineering rigor. Beyond geometry and connectivity, assets must include validated physical properties, control behavior, protection logic, and dynamic characteristics appropriate to the simulation domain. The required depth of modeling depends on the use case and is comparable to traditional engineering-grade simulation environments.

Limitations of traditional digital twin approaches

With traditional digital twin and infrastructure management tools, asset identity, semantics, and design intent are difficult to preserve across tools and over time, leading to fragmented models, semantic drift, inconsistent fidelity, and limited reuse of insights. These constraints reduce trust in simulation outputs, restrict system-level reasoning, and confine many solutions to advisory or exploratory use cases rather than predictive, prescriptive, or automation-ready applications. These limitations become increasingly acute at the scale, density, and complexity of modern data centers and AI factories.

Key limitations include:

- **Narrow and fragmented scope:** Models are built around isolated systems, tools, or project phases, preventing lifecycle continuity and coordinated cross-domain reasoning.
- **Proprietary architectures:** Closed data models and formats limit interoperability, preserve silos, and hinder composition into a persistent facility-level twin.
- **Limited simulation fidelity* and trust:** Disconnected and recreated models lack consistent behavioral detail, constraining reliable prediction and validation.
- **Incomplete system representation:** Even when facility-wide, models often omit logical constructs (e.g., redundancy zones, operating modes, workload groupings), limiting true system-level analysis and scenario evaluation.

* Note: Higher fidelity does not automatically equal higher value. Fidelity only matters when it directly supports a decision, automation, or risk reduction. Particularly in the early planning and design phases (site capacity planning, analyzing land constraints, etc.) high fidelity models of rack-level airflows or breaker settings, for example, adds no value in that case.

How a new architecture closes the gaps

These limitations are architectural, not merely functional. Addressing them requires a different approach to how digital twins are structured, integrated, and sustained across the lifecycle. This approach was impractical at industrial scale until only recently.

NVIDIA list[1] advances in four areas that now make this possible:

[1] https://www.NVIDIA.com/en-us/glossary/digital-twin/

- Open scene frameworks (OpenUSD): Enable assets and systems from many tools and vendors to be composed into a single, persistent digital representation, or scene.
- Accelerated computing: Makes large-scale visualization and simulation feasible in near real time.
- Physics-based simulation and graphics: Move digital twins beyond visualization into behaviorally accurate models.
- AI: Enables faster model creation, scenario exploration, and reasoning over complex system interactions.

These technologies only deliver sustained value when applied on top of a shared, governed scene foundation (see Figure 1). Without it, acceleration and AI simply speed up isolated models rather than enabling lifecycle continuity.

OpenUSD provides this foundation. It defines how assets are represented, composed, and maintained so that identity, structure, and meaning persist as models evolve. This allows electrical, mechanical, IT, and facility systems authored independently to coexist and interact within a single scene without being rebuilt at each transition.

When this foundation exists, AI and accelerated compute act as force multipliers rather than point optimizations.

Figure 1

Architectural overview of an OpenUSD-based, AI-enabled digital twin integrating infrastructure assets, engineering tools, and operational systems.

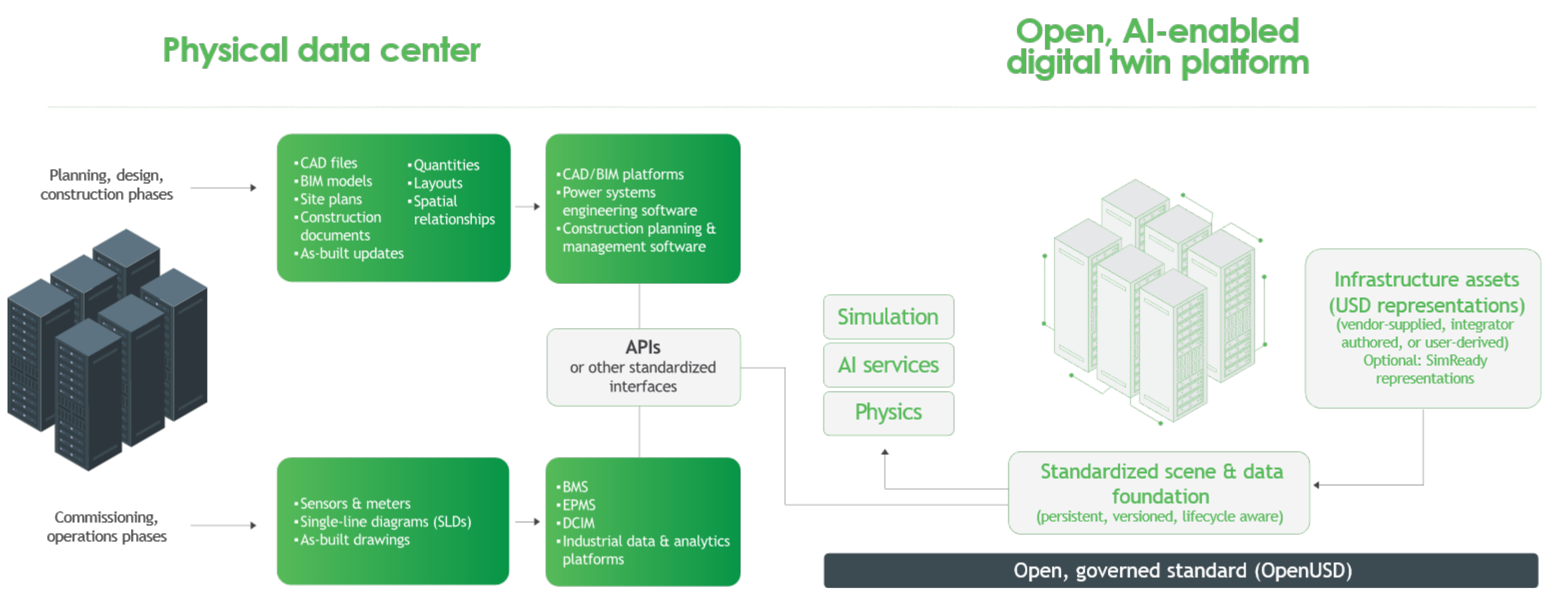

What Omniverse can enable

NVIDIA Omniverse operationalizes this architecture. It is not a replacement for DCIM, BMS, EPMS, BIM, CAD, Industrial intel platforms (e.g., AVEVA CONNECT), or power system modeling tools (e.g., ETAP). Instead, it acts as a platform layer that connects them into a shared OpenUSD-based scene.

In practical terms, Omniverse:

- Maintains a common digital representation of assets and systems using the OpenUSD framework.
- Integrates data from design and operational tools via APIs or other standardized interfaces.
- Supports GPU-accelerated simulation, rendering, and AI workflows.
- Enables multi-user collaboration on the same evolving digital twin.

Domain-specific tools (e.g., DCIM, BMS, EPMS) remain authoritative systems of record. Their value increases when their outputs are synchronized into a shared scene where system-level context, simulation, and AI can be applied consistently.

This enables a shift from phase-specific, disposable models to a single digital twin that persists and evolves from planning through operations. Design intent, assumptions, and constraints remain embedded rather than being lost between handoffs.

Practical use cases enabled by this architecture

Examples that are difficult or impractical without a shared OpenUSD-based foundation include:

1. **Continuous, multi-domain co-simulation: evaluating impact of increasing rack power density.**

   Simultaneously observe in one place impacts on transformer heating, breaker loading, coolant loop and airflow dynamics, and SLA risk without having to simulate and reconcile multiple models.

2. **Design intent carried into operations: flagging violations of tier or availability assumptions even if no alarms are triggered.**

   Early assumptions (e.g., density limits, redundancy intent, control philosophies, risk tolerances) remain embedded in the scene from one phase to another. Operational decisions are evaluated against original design intent, not just the current state. Without USD, these assumptions and intent often live in spreadsheets or people's heads once the site is launched.

3. **AI-assisted scenario generation and exploration: automatically discover non-obvious scenarios that could push a facility to instability.**

   Basically, humans don't have to enumerate all the "what-ifs". GenAI could propose stress/fault scenarios and failure chains. Reinforcement learning could explore many permutations using accelerated compute to uncover areas of risk that you would not have been aware of otherwise.

4. **Synthetic data generation for rare or unsafe conditions: training AI models to recognize early signs of cascading electrical failures that have never occurred on site.**

   USD-based platforms facilitate using AIs to train on things you don't want to happen. Synthetic data is needed since real world data would presumably be unavailable. USD more easily enables synthetic data creation that matches real geometry, physics, and controls. This physically accurate synthetic data generated from the same scene (as reality) could be used for planning and operations. It enables AI training for fault detection, anomaly recognition, and eventually autonomous controls.

5. **Human-AI collaboration inside the digital twin: with engineers, operators, and AI sharing the same context, you could use natural language queries and AI copilots to operate directly on the scene.**

USD-based platforms could provide the necessary data (real-time, structured semantics, spatial understanding) to enable this. Using screenshots and siloed dashboards of traditional tools alone likely would not provide enough meaning for this to work. With an AI copilot you could imagine, for example, an operator asking something like "Show me the safest way to add 20 kW of load per rack in this pod while maintaining N+1 and minimizing chilled-water flow increases." …and getting a visually grounded, simulation-backed answer.

6. Brownfield retrofits and expansions

   For existing facilities, the ability to more easily layer new assets and systems onto an existing digital twin helps owners evaluate retrofit options while applying and preserving historical context and constraints.

In all cases, value comes from reasoning across systems and time, not visualization alone.

What this means for data center and AI factory owners

As infrastructure scale and complexity increase, traditional tools and manual workflows struggle to keep pace. USD-based, AI-enabled digital twin platforms offer a more scalable operating model, one that enables system-level reasoning at machine speed.

For planning, this means faster exploration and earlier constraint detection. For construction and commissioning, simulation-first validation—including support for electronic site acceptance testing (e-SAT) and pre-commissioning verification before energization. For operations, continuous insight into how current conditions compare to original design intent, enabling real-time performance optimization and improved energy efficiency, including higher tokens per watt.

These outcomes depend on AI operating on a coherent architectural foundation. Platforms such as NVIDIA Omniverse combine OpenUSD-based scene representation with accelerated simulation and AI services, enabling lifecycle continuity at scale.

This is why Schneider Electric is engaging with OpenUSD and endorsing USD-based, AI-enabled digital twin platforms as the best path forward.

For data center and AI factory owners, here is a summary of the opportunity that these new platforms will bring…

- Improved visibility and shared understanding across stakeholders, enabling better, faster coordination and alignment.
- Higher-quality, more confident decisions with reduced human error, driven by consistent, system-level context.
- Faster analysis and identification of constraints and risks, before they impact cost or schedule.
- Safer, faster, and more predictable change, from design modifications through ongoing operations and maintenance.

Together, these benefits translate into faster time to power and revenue, as well as lower total cost of ownership compared to traditional, siloed approaches. As infrastructure complexity and sprawl continues to increase, platforms built on open standards and designed for AI-driven insight will be increasingly essential to operating data centers at the pace, scale, and reliability modern workloads demand.

Next steps

For large data center and AI factory owners seeking to apply the concepts described in this paper, the following steps can help translate architectural principles into practical progress towards the full promise of USD-based, AI enabled platforms:

- Assess your current digital twin and infrastructure management landscape. Evaluate how your existing tools (e.g., DCIM, BMS, EPMS, power system modeling, BIM) are used today, where models are fragmented across lifecycle phases, and where context or fidelity is lost during transitions from planning to operations. Consider industrial intelligence platforms like AVEVA CONNECT and digital twin platform NVIDIA Omniverse as a path to addressing those limitations.

- Prioritize platforms that are OpenUSD based and architected for lifecycle continuity. When evaluating future digital twin platforms, look beyond individual features and assess whether the solution is built on open, governed standards (such as OpenUSD) that preserve asset identity, semantics, and intent across tools, vendors, and lifecycle phases. Even if adoption begins with narrow, use-case-specific digital twin solutions, this foundation ensures continuity of data and a credible upgrade path as organizational needs and digital maturity evolve.

- Explore how best-in-class domain tools are beginning to integrate into platform-scale digital twins. Rather than replacing proven tools, understand how they can contribute to and benefit from a shared digital twin. For example, examine how ETAP used NVIDIA Omniverse to develop the world's first digital twin to simulate AI factory power requirements.

- Learn how industrial data and analytics platforms will extend digital twin value. Review how AVEVA integrates its industrial intelligence, data management, and analytics capabilities with Omniverse to support real-time data contextualization, simulation, and AI-driven insight across complex facilities. These integrations illustrate how operational data and enterprise analytics can enrich, not replace, the existing digital twin.

- Engage with partners that support open ecosystems and long-term flexibility. Choose technology partners and infrastructure vendors that actively support open standards and are investing in validated, simulation-ready ("SimReady") asset models with documented physical and behavioral characteristics, and who participate in ecosystem collaborations that reduce lock-in and preserve optionality as data center designs and operational requirements evolve.

About the authors

Kevin Brown is Senior Vice President and Chief Marketing Officer Secure Power, Data Centers & Global Services at Schneider Electric. With decades of experience in critical infrastructure and energy technologies, Kevin focuses on advancing resilient, efficient, and sustainable digital infrastructure. He is a frequent industry speaker and contributor on the future of data centers, edge computing, and the evolving demands of AI-driven workloads.

Alison Matte is the Director of Sustainability & Innovation for Schneider Electric's Data Center Segment. She leads strategic research on sustainable, resilient, and high-performance data center architectures, with a focus on how emerging power, cooling, and digitalization technologies shape the future of large-scale compute and AI infrastructure. Alison's work translates technical analysis into practical guidance for operators, developers, and policymakers, helping the industry reduce environmental impact while meeting rapidly growing capacity and efficiency demands.

Acknowledgements

Special thanks to Patrick Donovan for his significant contribution, and to Brice Martinot-Lagarde, Joseph Reele, Tanuj Khandelwal, Marc Garner, and Francisco Ramirez Fajardo for their expertise and technical review of this white paper.

RATE THIS PAPER ★★★★★

Life Is On Schneider Electric

Resources

Browse all white papers and executive reports
Data Center Research & Strategy Papers

Browse all TradeOff Tools™
Data Center Research & Strategy Tools

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@se.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm