Composable Bindings: Simplified System Integration

Embracing a modern integration approach for the Digital Engineering workflow

NVIDIA Microsoft

v1.0 | November 2025

Executive Summary

Organizations seeking to embed intelligent systems into industrial operations face a fundamental integration challenge. No single software platform or technology stack comprehensively addresses the diverse requirements of AI-driven operations—from predictive maintenance and energy optimization to adaptive manufacturing and real-time simulation. Operational AI – intelligent systems that unlock efficiency, resilience, and sustainability at scale require more than advanced algorithms. Success depends on an integrated system-of-systems architecture that enables seamless data exchange across disparate platforms, legacy systems, and modern AI tools.

But integration remains one of the most complex and resource-intensive barriers to AI adoption. Traditional approaches require extensive custom development, create maintenance burdens, and struggle to scale as organizations add new capabilities.

This whitepaper introduces **Composable Bindings**, a modern design pattern for simplified integration between software solutions, developed through the collaboration of Microsoft and NVIDIA - two leaders in cloud platforms and AI infrastructure. The approach reduces integration complexity, accelerates deployment timelines, and enables organizations to build flexible AI ecosystems without having to build costly point-to-point integrations over again.

Contributors

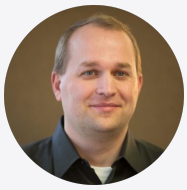

Christoph Berlin
VP of Engineering - *Microsoft*

Aaron Luk
Director of Product Management, Omniverse - *NVIDIA*

Contents

Digital Engineering as a Core Strategy

Digital Engineering represents a shift from manual, document-centric practices to model- and data-centric workflows that span the entire lifecycle of engineered systems and the products they create. While industries have their own interpretations of workflows, the industry-agnostic V-Model outlines a methodology of system requirements cascading down the left side into progressively detailed designs, and corresponding verification and validation activities rise on the right side. The concept of Digital Engineering Loop builds on top of the V-model and creates a digital thread through the entire process, enabling continuous, bidirectional flows of information across domains, disciplines, and organizations.

Enabling a Digital Engineering strategy necessitates a principled approach of interoperability across all software stacks. It is crucial to ensure that models, data, and processes can be exchanged, reused, and trusted across diverse toolchains and multi-vendor ecosystems. This entails not only technical integration but also governance frameworks that align the intent, context, and accountability of partners. By weaving these connections, digital engineering establishes a digital thread that links design, simulation, implementation, and operations.

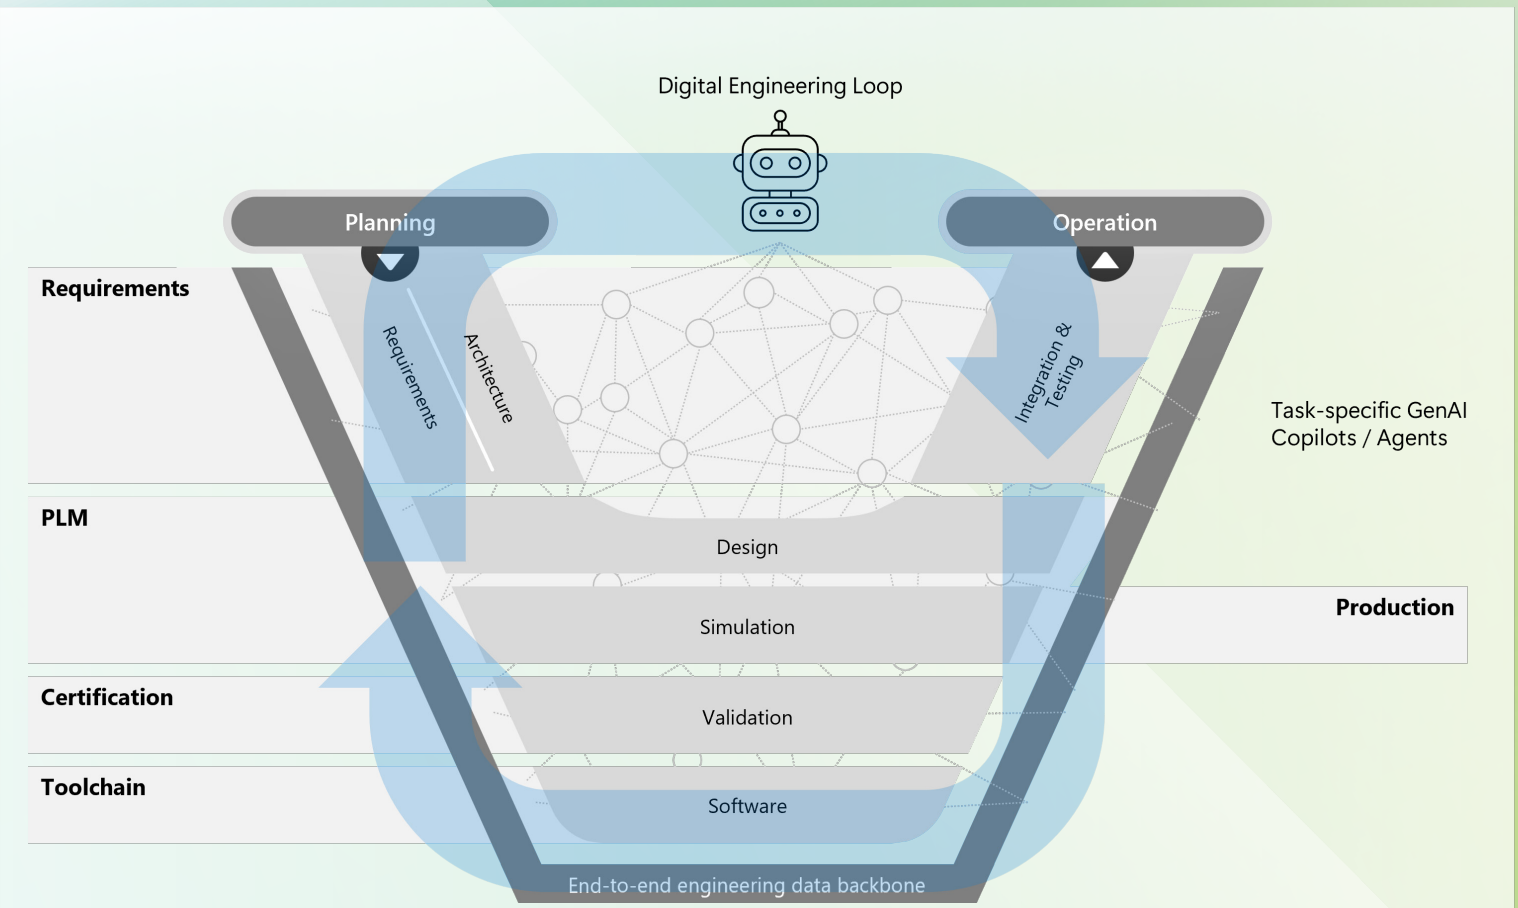

As systems become more cyber-physical, distributed, and software-defined, organizations are increasingly complementing the V-model with alternative frameworks such as the double V-model, the W-model, and lifecycle-spanning digital twin architectures. These models better capture the iterative, federated, and collaborative nature of modern engineering, where success depends less on isolated optimization and more on the ability to co-create within an open, extensible ecosystem.

Visualization a Key Requirement for Shift-Left Methodology in Digital Engineering

Simulation and visualization play one of the most pivotal roles in digital engineering. They enable organizations to adopt a shift-left strategy, where design, integration, and validation activities move earlier in the lifecycle, closer to the point where requirements are first defined. Traditionally, flaws or integration gaps surface late, during prototyping or commissioning, when changes are costly and delays are most damaging. By leveraging visualization twins, companies can validate assumptions, test performance, and explore alternatives in the digital world long before physical assets are manufactured or operationalized. This delivers measurable business value: it reduces capital expenditure on late-stage rework, accelerates time to market by minimizing downstream surprises, and strengthens strategic agility by allowing leaders to make earlier, data-driven investment decisions with confidence.

Visualization provides the engine room of digital engineering, where high-fidelity models of mechanical, electrical, software, and control systems as well as operational tasks, such as shift management or scheduling, are being tested and integrated into shared environments. These environments allow engineers to validate system behavior under realistic conditions, ensuring that the “real deal” has already been proven digitally before deployment. When connected through an intelligent digital thread, visualization not only helps validates designs but also creates a closed feedback loop: operational data from the field can be fed back into upstream models, continuously refining accuracy and enabling predictive maintenance, optimization, and lifecycle extension.

Crucially, visualization must also support multi-vendor collaboration by allowing partners to contribute domain-specific models into a federated environment without losing intellectual property control. This federated approach ensures that complex, cyber-physical systems can be validated as a whole, even when no single organization owns every component. As industries move toward distributed, software-defined automation, visualization becomes the proving ground where interoperability, safety, and performance are validated before the physical system is ever built, making shift-left not just a technical practice, but a strategic imperative.

Composable Bindings to Supercharge the Digital Engineering Loop

While Digital Engineering is essential to modern operational efficiencies, integrating all those heterogenous systems with each other is overly complex due to the high degree of integration required across disparate systems. For example, when organizations attempt to connect visualization environments with Product Lifecycle Management (PLM) systems, they frequently encounter significant costs and manual effort because the tools do not natively interoperate. Systems must be integrated, synchronized, and maintained across heterogeneous platforms, creating brittle workflows that depend on custom scripts, point-to-point integrations, or manual re-entry of information. This lack of seamless connectivity not only slows down engineering teams but also undermines the very promise of digital engineering by introducing errors, duplications, and delays. With integration costs being unreasonably high, the value of digital engineering strategies gets undermined as integration cost outweighs savings.

To address these concerns, a more modern approach to scalable integration is required, something we named Composable Bindings. Composable Bindings describe an integration strategy where systems are connected in ways that are flexible, context-aware, and resilient to change. Unlike rigid, point-to-point integrations that lock systems together and create fragility, Composable Bindings allow connections to adjust as requirements, tools, or data models evolve. They don't require every system to speak the same language natively; instead, they build upon an ecosystem of common data models that scales via conceptual data mapping and open standards that let each system contribute its strengths without being forced into a brittle one-to-one relationship.

Conceptual Approach to Composable Bindings

Industrial organizations collect massive volumes of machine asset telemetry, sensor readings, performance metrics, and event logs, streamed continuously from equipment in the facility, such as a data center or a manufacturing plant. This telemetry is ingested into a unified data lake, where it is stored at scale, often in raw or semi-structured formats. The data lake excels at aggregation and business intelligence, but by itself it does not provide the visualization needed for the digital engineering. On the other side of the architecture sits a **visualization engine,** designed to transform data into 3D scenes and more broadly speaking to build the industrial metaverse. Participants of the Digital Engineering loop see visualizations of the telemetry events in the context of the full scene.

Quite often challenges arise when customers want to bring these two worlds together. The telemetry in the data lake is rich but heterogeneous, different assets may produce data in varying formats, with inconsistent KPI, and at different cadences. The visualization engine, meanwhile, expects data to be

delivered in a structured, contextualized form it can render meaningfully. Without a scalable integration strategy, bridging the gap requires brittle, point-to-point pipelines: custom scripts to reshape data, manual mappings to align bespoke schemas, and ongoing maintenance to keep the connections alive as assets, schemas, or visualization requirements evolve. This creates friction, slows down innovation, and undermines the promise of real-time insight.

Challenges in current integration strategies:

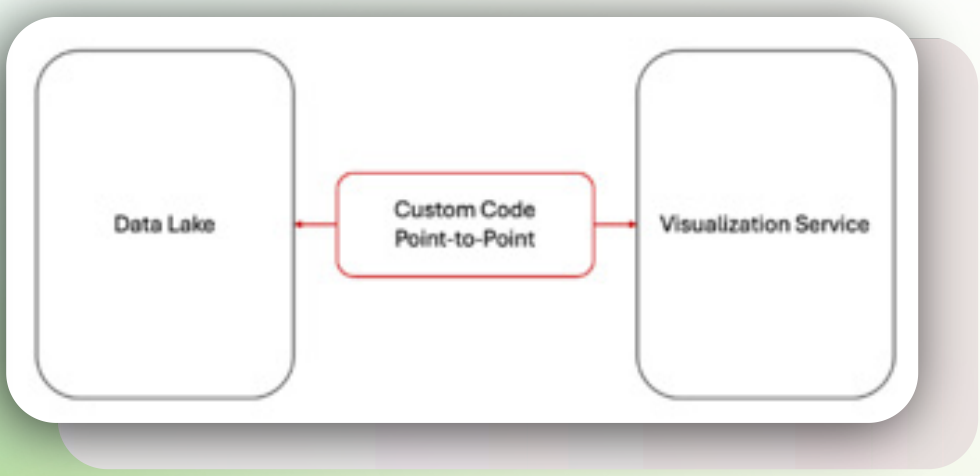

This is precisely where Composable Bindings come into play. Instead of forcing rigid one-to-one integrations, composable bindings act as a flexible integration layer between systems, such as the data lake and visualization engine. They can ingest telemetry in its native form, enrich it with metadata, and harmonize it into a semantic model that preserves context. At the same time, they can adapt that harmonized data into the structures and APIs the visualization engine requires, aggregating, filtering, or reshaping on demand. The result is a resilient, future-proof connection: the data lake and visualization engine can evolve independently, while the composable binding ensures they remain interoperable.

Composable Bindings simplify the integration by tethering solutions through minimal dependencies:

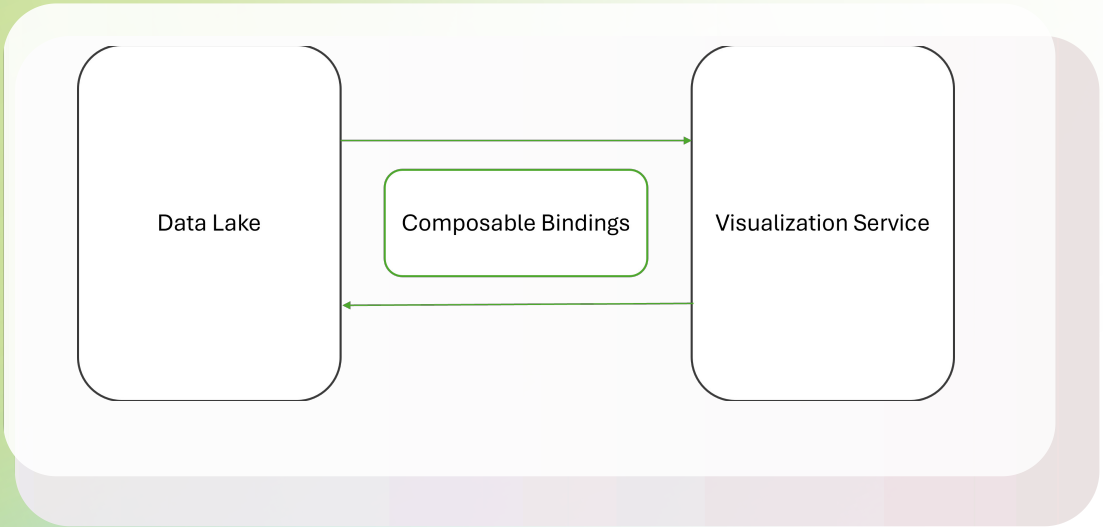

Putting Composable Bindings into Practice

Let's look at how these principles manifest in practice. We have constructed a reference architecture and implemented a proof of concept to serve next-generation data center design and layout. For these gigaLet's look at how these principles manifest in practice. We have constructed a reference architecture and implemented a proof of concept which simulates a digital engineering loop to optimize next-generation data center design and layout. For these gigawatt data centers, digital engineering and operational AI are paramount to maximizing energy efficiency and token output as AI Factories, for example by adjusting the cooling system in response to an analysis of real-time observability of temperature events on compute racks. Event processing, visualization and simulation at scale are key for the digital engineering loop to optimize the most comprehensive range of real-world configurations and scenarios.

Real-world facilities such as these AI Factories encompass many data sources (including differential pressure, liquid flow rate, and leak detection status in addition to the temperature events mentioned above), representing numerous characteristics and observability for the thousands of operational equipment that need to be visualized and simulated. Vast amounts of operational events (tens of thousands per second) representing real-world updates must travel through the digital twin, aggregating observability metrics in context for analysis and optimization via agentic AI.

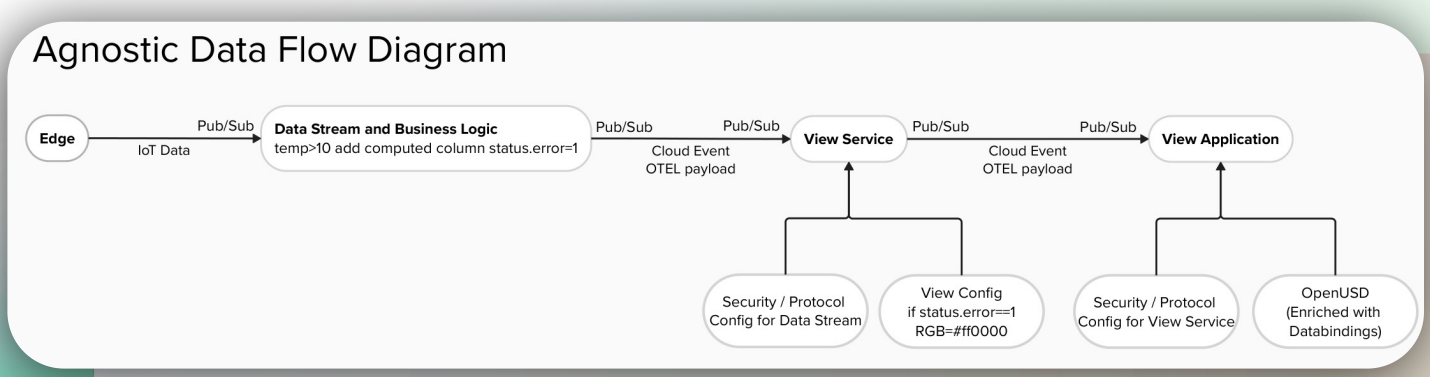

Our proof of concept, illustrated in the diagram above, showcases a modular, scalable data flow architecture for compute rack monitoring, beginning with each rack streaming telemetry to the data lake, representing systems such as business intelligence systems or any other software solutions.

On the Data Lake side:

- Telemetry data, such as OPC UA, Open Telemetry (OTEL) or any other message schema is ingested into a central processing layer (e.g. Eventstream) & stored in an Eventhouse for further processing.
- From there, business logic such as temperature thresholds can be applied to enrich the event with analytic fields such as status indicators.
- The processed data is then packaged into an OTEL schema , representing a common data model for the event (containing fields such as object identifier, temperature value, and status), ensuring interoperability and future extensibility across different transport protocols.
- The enriched OTEL messages are published to event topics via CloudEvents, supporting multiple consumers, including a View Service which can also configure another level of business logic to interpret the observability event and status into 3D visual indicators, and enabling robust, scalable distribution.

On the visualization side, an application subscribes to the relevant event topics from the View Service, interprets the incoming data, and updates the user interface in real time via Universal Scene Description (OpenUSD), e.g., changing rack colors based on temperature event status.

The proof of concept utilizes Composable Bindings to realize the full data flow from the edge through Microsoft Fabric into NVIDIA Omniverse, while the reference architecture is applicable to bind objects and ontologies in data lakes to their 3D manifestations in the visualization and simulation at scale.

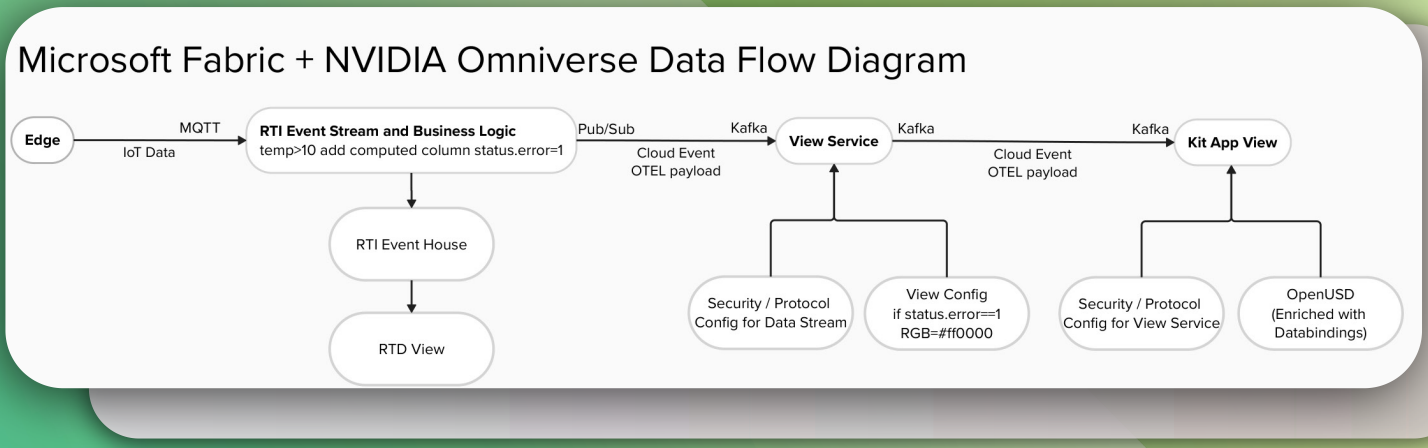

This approach demonstrates a clear separation of concerns, allowing end users to focus on business logic (e.g., the temperature thresholds and status visualization) and their own data structure (e.g., temperature events from the Building Management System, delivered via MQTT), while ISVs can leverage standardized, scalable infrastructure for integration, visualization, and simulation.

Landscape for the Content Journey and Standards Development

Advancing digital engineering and Physical AI relies on a robust content journey and standardization strategy. This approach centers on model- and data-centric methodologies, enabling continuous information exchange and seamless integration across a wide array of tools and systems. A key objective is to empower equipment manufacturers to provide virtual representations of their physical products, which are then used to populate digital twins and simulation environments.

OpenUSD serves as a vital entry point and aggregator for 3D visualization and simulation. By aggregating 3D information such as geometry and materials along with domain metadata from multiple sources, OpenUSD ensures clarity and interoperability for both human users and AI-driven processes. In parallel, OTEL standardizes the collection and transmission of telemetry data, supporting consistent monitoring and observability across distributed systems. CloudEvents complements these efforts by providing a unified payload structure for event-driven data exchange, facilitating seamless integration between platforms and tools.

The standardization process encompasses mapping existing data, developing new schemas, and establishing best practices for asset structuring. The Alliance for OpenUSD (AOUSD) formalizes these into normative standards for 3D interoperability and composability. OTEL ensures operational data models are standardized and easily integrated, and CloudEvents acts as a common envelope for event-driven communication. Together, these open standards enable the creation of modular and scalable architectures for digital twin and simulation workflows.

A major challenge in this journey is adapting to rapidly evolving technology while maintaining consistent standards. The process requires iterative prototyping, community engagement, and refinement of common data models and operational behaviors into widely adopted standards. OpenUSD, OTEL, and CloudEvents all support this iterative, community-driven approach, allowing frameworks to evolve with industry needs and ensuring compatibility as new methods and technologies emerge.

To address these challenges, a structured approach to standardization is essential. This includes developing modular data flow architectures that leverage industry standards for data transport and observability and event data, and CloudEvents provides a unified event messaging paradigm for

interoperability. Segmenting content requirements across distinct simulation capabilities key to a modular, prioritized approach to fully represent robots and equipment in a digital twin. Active participation in standardization—through learning, certification, tool adoption, and knowledge sharing—is crucial for bridging gaps and driving innovation across industries.

Next Steps: Activating the Digital Thread

The immediate next steps focus on enhancing the reference architecture for further configurability, scalability, and performance. Key priorities include migrating data binding logic from hard-coded application code to declarative configurations, enabling more flexible and maintainable integration for diverse use cases. Further separation of concerns includes partitioning the specifics for the transport protocols from the topical specifics for the data bindings.

We are also stress testing the system by increasing data update frequency and scaling the number of racks, ensuring the architecture can handle realistic industrial workloads. Performance metrics and tuning are assessed in collaboration with domain experts to align with real-world operational requirements.

Plans are underway to open-source the reference design, fostering broader adoption and community contributions. This will allow partners and ISVs to extend the baseline for their own scenarios, accelerating the adoption of standardized, scalable digital twin solutions across the industry.

Get started with OpenUSD and with many available OpenUSD resources, including the self-paced Learn OpenUSD, Digital Twins, and Robotics training curriculum. Learn more about OpenTelemetry and CloudEvents and consider how your own existing point to point integrations can adopt these and other open standards to more effectively scale your solutions in an agentic world.