Digital Twins and the data backbone

Enabling the next generation of intelligent systems

SurrealDB

Table of Contents

Executive summary

Digital Twins and the data backbone

The shift to system-level intelligence

Digital twins have evolved from niche aerospace simulation tools into the foundational control systems of modern enterprises. As organisations integrate artificial intelligence and expand these virtual replicas across global supply chains, smart cities, and healthcare networks, digital twins are no longer just passive operational mirrors. They are becoming active, AI-driven decision environments that allow leaders to simulate, test, and optimise strategies before executing them in the physical world.

The architectural bottleneck

The primary barrier to scaling digital twins is no longer simulation software or AI capabilities - it is data architecture. Building an enterprise-grade digital twin requires continuous real-time data ingestion, complex relationship mapping, and temporal state tracking. Attempting to support these interconnected requirements with traditional, single-model architectures results in "architectural sprawl"- a fragile patchwork of separate graph, time-series, document, and relational databases stitched together with complex middleware.

The multi-model foundation

This whitepaper explores why the next generation of intelligent systems demands a unified data backbone. By consolidating document storage, native graph relationships, and time-series telemetry into a single multi-model engine, platforms like SurrealDB 3.0 eliminate integration overhead and latency. We outline how transitioning to a unified data architecture provides the strategic foundation necessary to build scalable, real-time, and AI-ready digital twins.

Key takeaways at a glance:

From assets to ecosystems: how digital twins are shifting from monitoring single machines to simulating entire enterprise ecosystems and Digital Twins of the Organisation (DTOs).

The hidden cost of polyglot persistence: why stitching together specialised databases introduces synchronisation risks, governance challenges, and operational friction.

The multi-model advantage: how natively unifying graph, document, and temporal data enables real-time responsiveness, scenario branching, and safe AI experimentation without crossing system boundaries.

Actionable technical implementation: concrete examples of how to model, populate, and query complex industrial assets and their dependencies within a single, coherent data layer.

Enabling the next generation of intelligent systems

A digital twin is a dynamic, data-driven, virtual representation of a system or physical object that uses real-time data to accurately reflect its real-world counterpart's behaviour, performance and conditions. Unlike static models, it continuously ingests real-time data, reflects current state, simulates future scenarios, and increasingly incorporates artificial intelligence to optimise decision-making.

Digital twins are rapidly evolving from specialised engineering tools into foundational infrastructure for modern enterprises. Originally developed to simulate high-risk physical systems in aerospace and manufacturing, digital twins are now being deployed across supply chains, cities, healthcare systems, defence operations, energy networks and financial ecosystems.

Several converging forces are accelerating their adoption: the proliferation of IoT sensors and connected devices; advances in real-time data processing; GPU-driven simulation platforms; generative AI and autonomous agents, and the need for resilience in increasingly complex systems.

As organisations face growing operational volatility - from geopolitical disruption to climate risk and supply chain instability - the ability to simulate, test and optimise before acting has become a strategic advantage.

In a survey of 660 C-suite executives and senior leaders across 11 industries, 62% maintained that there had been immense value from the technology. 8 out of 10 organisations found that digital twins helped to reduce their carbon emissions, and 50% found their ROI led to between 21-40% revenue growth.[1]

However, the success of digital twins is not determined solely by simulation software or AI capabilities. It depends fundamentally on data architecture. Digital twins are not simply visual replicas; they are living data systems that must model relationships, track state over time, process high-velocity events, and scale across distributed environments.

"In aerospace, resilience and precision depend on reliable, transparent data - it's the foundation that enables AI, digital twins, and the next generation of intelligent operations".

-Taylor Brown, COO and Co-founder, Fivetran

Digital twins are becoming the control systems of modern enterprises. The organisations that invest in the right foundations today will be best positioned to operate intelligently in an increasingly complex world.

[1] *Hexagon Digital Twin Report 2024*

The evolution of the digital twin

From aerospace simulation to enterprise infrastructure

The concept of the digital twin can be traced back to aerospace engineering, where mission-critical systems required precise modelling and continuous monitoring. During NASA's Apollo programme, engineers maintained ground-based replicas of spacecraft systems to simulate anomalies and test corrective strategies in real time. While the term "digital twin" was not yet widely used, the principle was established: a live, virtual counterpart of a physical system could reduce risk and enhance decision-making.

In the decades that followed, industrial sectors such as aerospace, automotive and advanced manufacturing developed increasingly sophisticated simulation models. These early systems were primarily design-focused. They supported prototyping, performance testing and stress analysis before production.

The next major shift occurred with the rise of Industry 4.0 and the expansion of IoT networks. Sensors embedded in machines, factories and infrastructure began streaming operational data continuously. This enabled digital representations to evolve from static design models into dynamic operational mirrors. For the first time, digital models could reflect real-world state in real time.
Today, digital twins are entering a third phase of evolution.

The convergence of AI, simulation and real-time data

Several technological developments are transforming digital twins into intelligent systems:

- **high-performance computing and GPU-based simulation platforms** enable realistic modelling of entire factories, cities and energy grids
- **cloud and edge computing architectures** allow distributed ingestion and processing of massive data volumes
- **artificial intelligence and machine learning** enable predictive and prescriptive insights rather than simple monitoring
- **generative AI** introduces the ability to simulate alternative scenarios, test policy changes, or model system adaptations dynamically.

Digital twins are increasingly being described not merely as mirrors of reality but as decision environments - platforms in which organisations can explore the consequences of actions before implementing them.

In parallel, governments and research institutions are advancing national digital twin programmes, particularly in infrastructure and urban planning. Enterprises are expanding digital twin initiatives beyond individual machines to encompass supply chains, logistics networks, and even workforce planning.

The evolution is clear
Digital twins are moving from asset-level optimisation to system-level intelligence.

Current industry applications of digital twins

Digital twins are already delivering measurable value across multiple sectors. While implementations vary in scope and sophistication, a common pattern emerges: organisations use digital twins to reduce uncertainty, optimise performance, and improve resilience.

The following are key sectors where digital twins are becoming operationally significant.

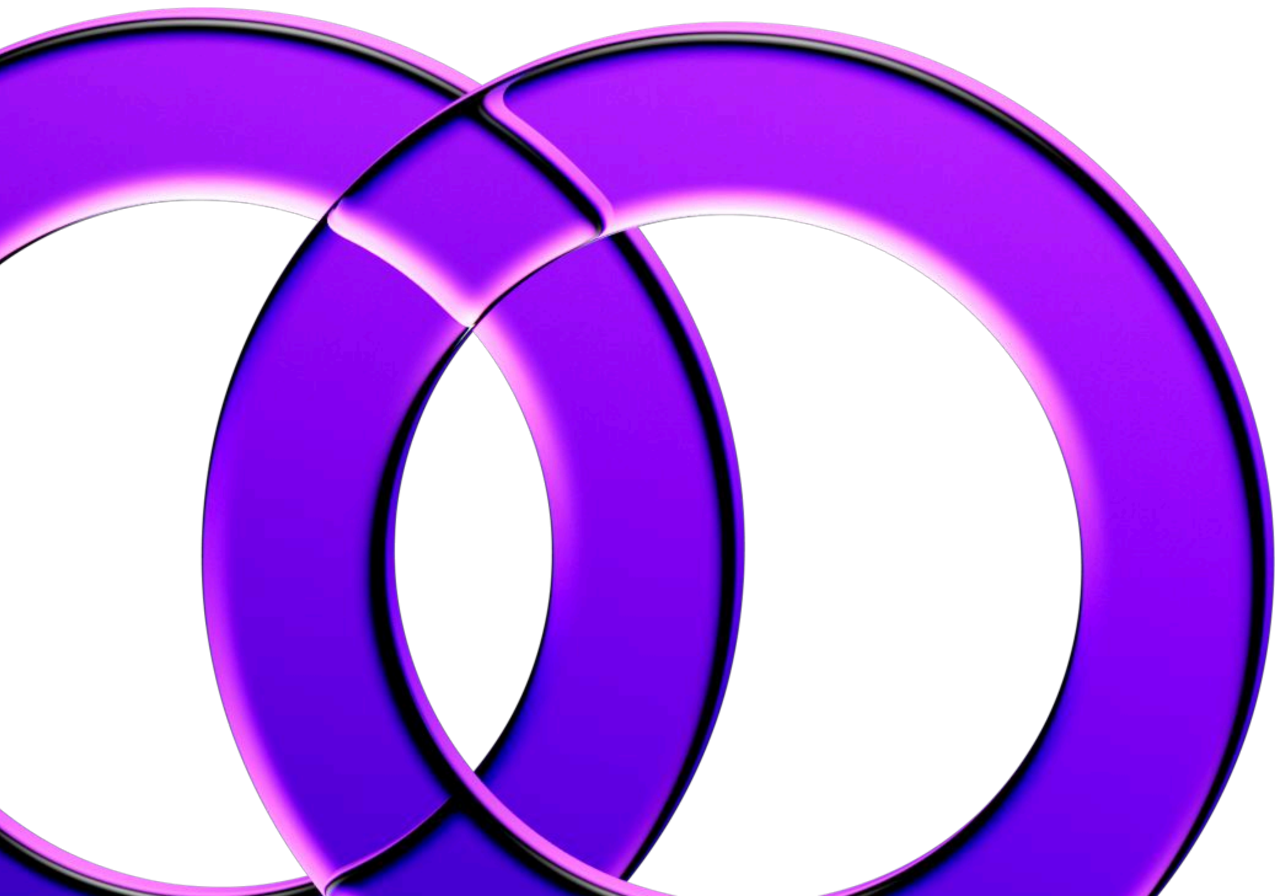

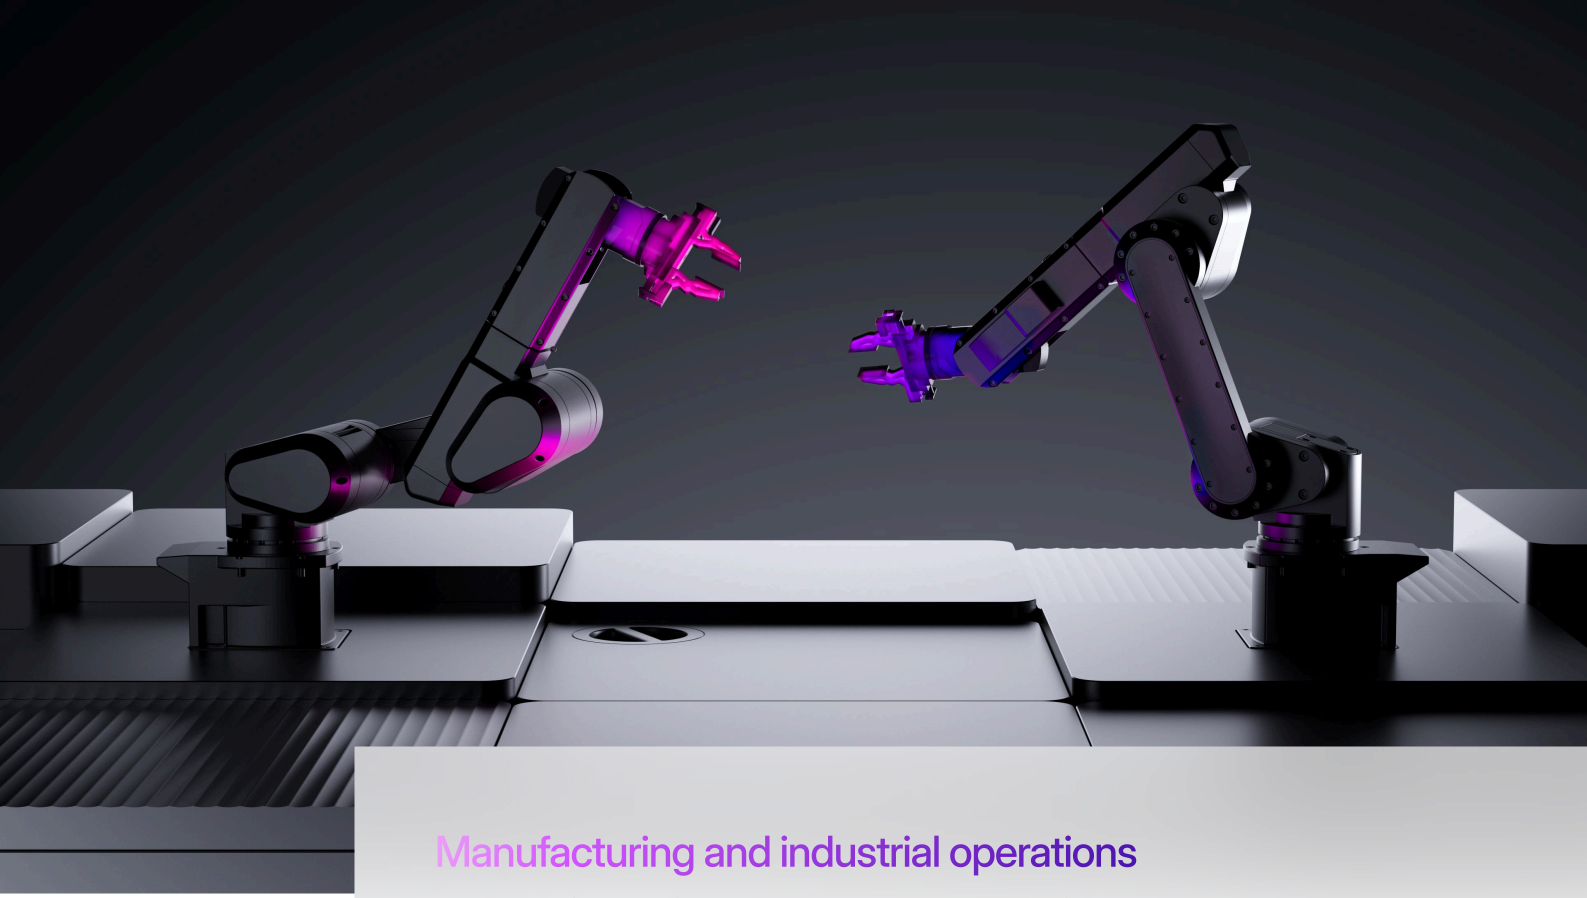

Manufacturing and industrial operations

Manufacturing remains one of the most mature domains for digital twin deployment.

The challenge

Modern factories involve complex coordination between machines, robotics, supply inputs and human operators. Downtime is costly, quality variation is expensive, and process inefficiencies compound across production lines.

The digital twin approach

Manufacturers create digital replicas of production lines, machinery and workflows. These systems ingest sensor data continuously, enabling: predictive maintenance, throughput optimisation, quality control modelling, robotics coordination simulation and energy usage optimisation. Advanced platforms now integrate physics-based simulation with AI-driven forecasting, enabling factories to test production changes virtually before implementing them physically.

Business outcomes

→ **Reduced unplanned downtime**

→ **Improved asset lifespan**

→ **Increased production efficiency**

→ **Lower operational risk**

Digital twins in manufacturing are evolving from maintenance tools into comprehensive operational control systems.

Supply chains and logistics

Global supply chains are increasingly volatile, interconnected and exposed to external shocks.

The challenge

Disruptions - from geopolitical tensions to extreme weather - can cascade rapidly across suppliers, ports, transport networks and distribution centres. Traditional planning tools often lack the dynamic modelling capabilities required to anticipate these effects.

The digital twin approach

An end-to-end digital twin of a supply chain models suppliers, transportation routes, inventory levels, demand signals and risk variables as an integrated system.

This enables organisations to: simulate supplier disruption scenarios; model inventory rebalancing strategies; optimise routing under constraint conditions; and stress-test demand fluctuations.

Business outcomes

→ **Reduced stockouts and overstock**

→ **Increased resilience**

→ **Improved response time to disruptions**

→ **Data-driven strategic planning**

Digital supply chain twins are increasingly integrated with AI systems capable of recommending mitigation strategies.

Smart cities and infrastructure

Urban environments represent some of the most complex systems managed by public authorities.

The challenge

Cities must coordinate transportation, utilities, emergency services, energy systems and environmental monitoring - all under fiscal and regulatory constraints.

The digital twin approach

City-scale digital twins combine infrastructure data, traffic systems, energy consumption patterns, environmental metrics and demographic data into integrated simulation environments.

Applications include: traffic flow optimisation; flood and climate impact modelling; energy grid balancing; infrastructure lifecycle planning; emergency response simulation.

National digital twin initiatives are emerging to connect built environment data across public and private stakeholders.

Business and societal outcomes

→ **Reduced congestion and emissions**

→ **Improved infrastructure planning**

→ **Enhanced disaster preparedness**

→ **More efficient public spending**

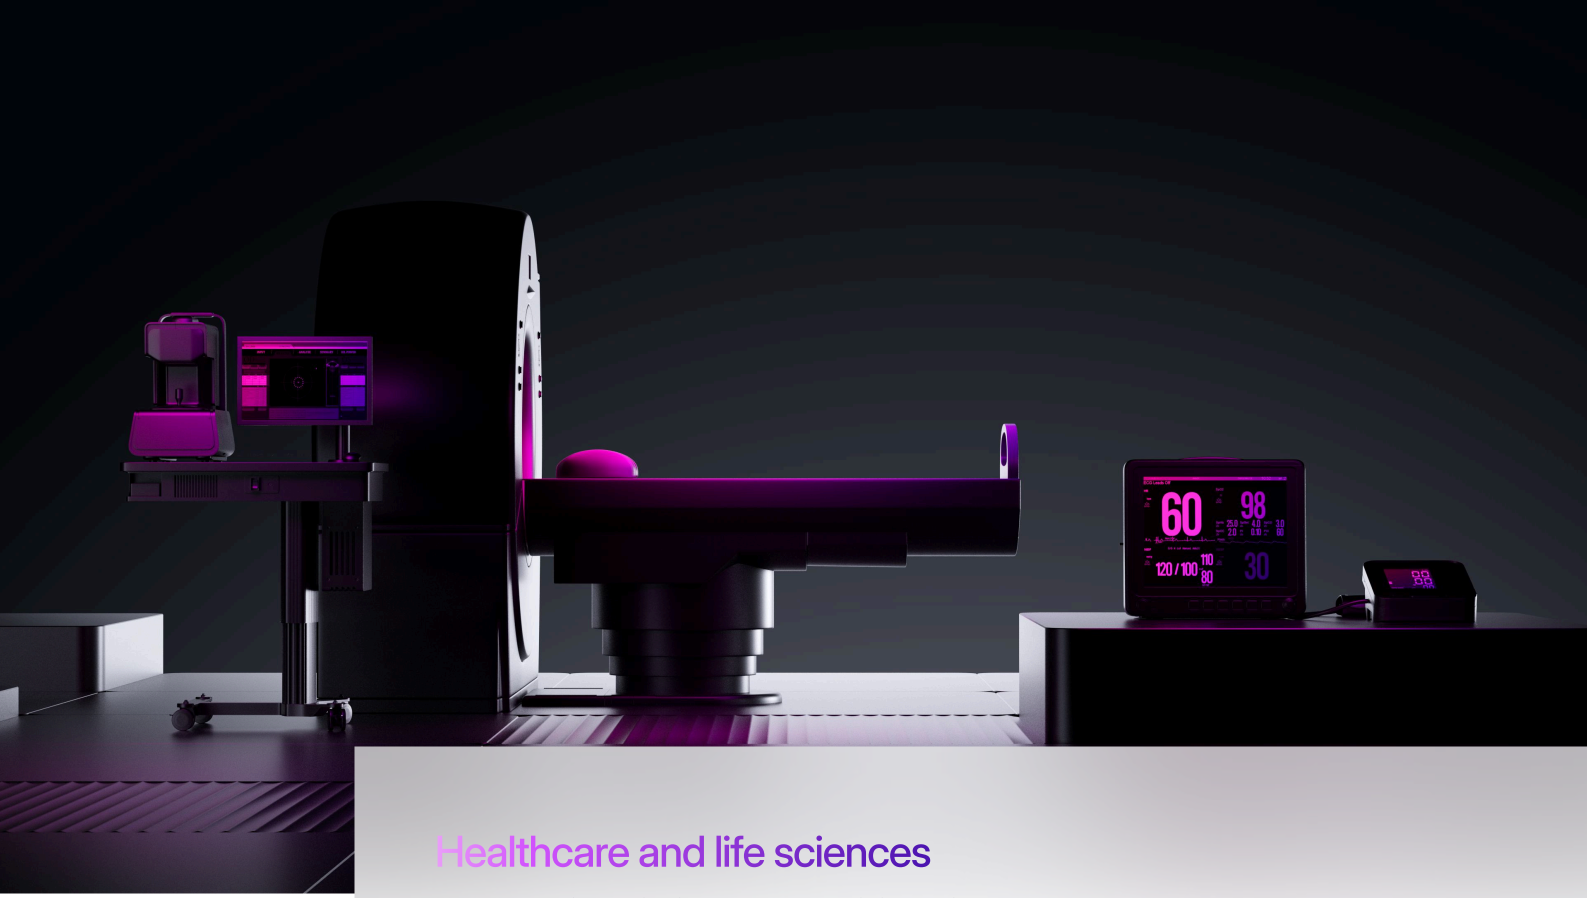

Healthcare and life sciences

Healthcare is beginning to explore digital twins at both system and individual levels.

The challenge

Medical systems involve complex interactions between patient data, treatment pathways, biological variability and resource constraints.

The digital twin approach

Digital twins are being explored for: patient-specific treatment modelling; hospital operations optimisation; drug discovery simulation; synthetic data generation for research; epidemiological modelling.

In research contexts, digital twins allow hypotheses to be tested virtually before clinical deployment, reducing cost and accelerating discovery.

Business and societal outcomes

- → **More personalised treatment strategies**
- → **Reduced clinical risk**
- → **Faster therapeutic development**
- → **Improved operational efficiency**

Although still emerging, healthcare digital twins represent a transformative frontier.

Energy, utilities and environmental systems

Energy systems are undergoing rapid transformation driven by decarbonisation and distributed generation.

The challenge

Grid operators must balance supply and demand across renewable, conventional and storage assets while managing volatility.

The digital twin approach

Energy digital twins model: generation assets; transmission networks; storage capacity; consumption patterns; and weather variables

Simulation enables operators to predict load imbalances, optimise distribution and model infrastructure upgrades.

Business and societal outcomes

→ **Improved grid stability**
→ **Lower energy waste**
→ **Better renewable integration**
→ **Reduced carbon intensity**

A data-driven pattern across industries

Despite sector-specific applications, digital twin deployments share several characteristics:

- they model **complex, interconnected systems**
- they depend on **continuous real-time data ingestion**
- they require modelling of **relationships and dependencies**
- they track **state changes over time**
- they increasingly integrate **AI for prediction and optimisation**

These shared attributes point toward a central insight: digital twins are fundamentally data architectures as much as they are simulation platforms.

"The oil & gas industry is actively deploying twins in areas like maintenance planning, inspection, and process safety, including sim ops. What I really see a call for now is scenario evaluation and being able to guide actions in real time - intervention intelligence, where you want to know not only why something is happening, but what you can do about it. As digital twin and AI technologies intersect, it will get easier to provide these types of insights."

- Marty Gonzalez, Innovation & Technology Principal, BP.

Emerging and future use cases

Digital twins are evolving beyond asset-level optimisation into system-level intelligence. While current deployments focus on improving operational efficiency, the next generation of digital twins will increasingly serve as decision platforms, policy simulators, and AI-enabled control environments.

> *"I'm excited about Gen AI and AI in general, because the pace of development is much faster there – that means we are able to bring some of those improvements into the technology used in the development of digital twins."*
>
> *-Tarun Dhall, Compressions Business, Siemens Energy*

Several emerging applications illustrate how digital twins are expanding in scope and ambition.

The digital twin of the enterprise

One of the most significant emerging concepts is the **Digital Twin of the Organisation (DTO)** - a dynamic model of enterprise processes, resources, workflows, and dependencies. Unlike traditional enterprise resource planning (ERP) systems, which record transactions, a DTO simulates how changes ripple through the organisation.

A mature DTO may model: organisational structures and reporting lines; workforce capabilities and capacity; supply dependencies and vendor risk; IT systems and infrastructure interdependencies; and financial flows and performance drivers.

This allows leadership teams to test structural changes - such as reorganisations, outsourcing decisions, automation strategies or market expansions - before implementation.

For example:

what is the impact of shifting production between regions?

how would a hiring spree or freeze affect delivery timelines?

what are the cascading risks if a key vendor fails?

In volatile economic conditions, such scenario modelling becomes a strategic tool rather than an analytical exercise.

AI agents operating within digital twins

Digital twins are increasingly becoming controlled environments in which AI agents can operate safely. As generative AI and autonomous decision systems mature, enterprises face a critical governance question: how can AI systems be tested and validated before affecting real-world operations?

Digital twins provide a structured, simulated environment in which: AI systems can test operational decisions; policies can be stress-tested under varied scenarios; resource allocation algorithms can be optimised; and risk exposure can be evaluated. Observe outcomes and refine strategies, and only then make a decision on how much real-world execution if at all should be given to an AI system.

In this sense, digital twins may become a prerequisite for trustworthy AI in complex enterprise settings.

To illustrate how this interaction might work in practice, consider an industrial facility operating hundreds of interconnected machines. Within the digital twin, each machine is represented as an entity linked to other components, suppliers, maintenance schedules and operational constraints. Sensor data continuously updates the twin with information about temperature, vibration, utilisation and performance.

An AI agent analysing this environment might begin by traversing the network of relationships within the twin to identify components that are approaching failure thresholds. Rather than recommending maintenance in isolation, the agent can analyse broader system dependencies: which machines support critical production lines, which spare parts are available, and how maintenance windows would affect operational output.

Before any action is taken in the real world, the agent tests its proposed strategy inside the digital twin environment.

Within this sandboxed scenario, the agent might simulate:

- **delaying maintenance by several hours to avoid peak production time**
- **scheduling replacement parts through alternate suppliers**
- **rebalancing production workloads across other machines.**

Each scenario can be evaluated against multiple objectives, including operational efficiency, safety constraints and financial cost. The results of these simulations can then be compared with historical system behaviour and organisational policies. Only after this analysis is complete would the system recommend an action to human operators or, in highly controlled environments, trigger automated execution.

Governance and trust

Before any action is taken in the real world, the agent tests its proposed strategy inside the digital twin environment.

Within this sandboxed scenario, the agent might simulate:
- delaying maintenance by several hours to avoid peak production time
- scheduling replacement parts through alternate suppliers
- rebalancing production workloads across other machines.

Each scenario can be evaluated against multiple objectives, including operational efficiency, safety constraints and financial cost. The results of these simulations can then be compared with historical system behaviour and organisational policies. Only after this analysis is complete would the system recommend an action to human operators or, in highly controlled environments, trigger automated execution.

Governance and trust

This approach provides an important governance mechanism for organisations deploying AI systems. Rather than allowing AI agents to operate directly on live infrastructure, the digital twin becomes an intermediate decision layer. It enables organisations to observe how AI systems reason about complex environments, evaluate their proposed actions, and measure the potential consequences before any real-world intervention occurs. In effect, the digital twin becomes a testing and validation environment for operational AI.

As enterprises deploy increasingly autonomous systems, this capability may become critical for maintaining safety, transparency and regulatory compliance. Digital twins allow organisations not only to optimise operations but also to govern how intelligent systems interact with the real world.

Climate, ESG and resilience modelling

As regulatory and stakeholder pressure increases, organisations must assess environmental and climate-related risks with greater precision.

Emerging digital twin applications include: carbon footprint modelling across supply chains; infrastructure stress testing under extreme weather scenarios; energy optimisation strategies; water usage and resource efficiency modelling; and regulatory impact simulation. By integrating environmental data streams with operational systems, digital twins can move sustainability from reporting to optimisation.

For infrastructure-heavy industries - energy, transportation, real estate - climate scenario simulation may become essential to strategic planning and insurance risk management.

Financial system and market simulation

Financial institutions and large enterprises are beginning to explore digital twins of economic and financial systems. These models may simulate: liquidity flows; counterparty exposure; credit risk propagation; market stress scenarios; portfolio rebalancing strategies. While financial modelling has long existed, the integration of real-time transactional data, behavioural signals, and AI-driven simulation marks a new phase.

Digital twins in finance may evolve from static stress tests into continuously updating system-level simulations that inform capital allocation and regulatory compliance.

Human-centric and healthcare twins

Healthcare digital twins are advancing from system-level hospital modelling to patient-specific applications. Emerging research explores: personalised treatment modelling; disease progression simulation; surgical outcome prediction; drug interaction modelling; and synthetic data generation for clinical research. While regulatory and ethical considerations remain substantial, digital twins offer the possibility of testing treatment pathways virtually before applying them in practice.

In parallel, workforce modelling in healthcare systems may help optimise staffing levels, resource allocation, and emergency preparedness.

Infrastructure autonomy and self-optimising systems

The long-term trajectory of digital twins points toward partially autonomous infrastructure. In energy grids, manufacturing plants and logistics hubs, digital twins may: continuously optimise routing and load balancing; detect anomalies and self-correct; adjust configurations dynamically based on demand; and coordinate distributed assets across geographies. As automation increases, digital twins shift from descriptive and predictive systems to prescriptive and adaptive ones.

This evolution depends not only on AI capability but on reliable real-time data integration and system-wide state awareness.

Digital twins as system-of-systems platforms

Perhaps the most transformative development is the emergence of system-of-systems digital twins. Rather than modelling isolated assets, organisations are beginning to connect: manufacturing twins; supply chain twins; financial twins; workforce twins; and environmental twins. When integrated, these form a comprehensive simulation layer for the enterprise. In this model, decision-making becomes less reactive and more experimental. Strategies can be tested in a controlled digital environment before implementation in reality. This represents a structural shift in how organisations manage complexity.

A data-centred strategic inflection point

Across these emerging applications, several themes are consistent: the scope of digital twins is expanding from individual assets to interconnected ecosystems; AI is becoming embedded within twin environments; real-time, high-volume data ingestion is a prerequisite; relationships between entities - not just attributes of individual assets - are central to value creation; and temporal modelling (how systems evolve over time) is increasingly important.

As digital twins grow in scale and ambition, they demand more than visual simulation engines or analytics dashboards. They require robust, flexible, and unified data architectures capable of representing complex, evolving systems.

The architectural attributes required for digital twins

As digital twins expand from isolated simulations to enterprise-wide intelligence platforms, their success increasingly depends on architectural foundations rather than visual modelling tools alone. A digital twin is not simply a dashboard, a 3D rendering, or a predictive model. At scale, it becomes a continuously evolving system of record and simulation - representing entities, relationships, state transitions, constraints, and behaviours across time. This complexity places specific and often underappreciated demands on the underlying data architecture.

The following attributes are consistently required in effective, enterprise-grade digital twin systems.

Continuous real-time data ingestion

Digital twins derive value from their ability to reflect the current state of the systems they represent. This requires ingesting and processing: sensor data from physical assets; transactional business events; environmental and contextual data streams, and external signals such as market conditions or weather patterns.

In many environments, data arrives at high velocity and in heterogeneous formats. The architecture must support streaming ingestion while maintaining data integrity and consistency. Without reliable real-time integration, a digital twin becomes a historical model rather than a live operational mirror.

Modelling complex relationships

Real-world systems are defined not merely by individual components but by the relationships between them.

A manufacturing robot depends on upstream supply; a supplier connects to multiple distribution hubs; a patient interacts with clinicians, medications and environmental factors; and an energy grid balances generation assets across geographic regions. Capturing these interdependencies requires a data model that can represent: many-to-many relationships; hierarchical structures; dynamic dependencies; network topologies, and conditional interactions.

As digital twins scale from asset-level representations to system-level models, relationship modelling becomes central rather than peripheral.

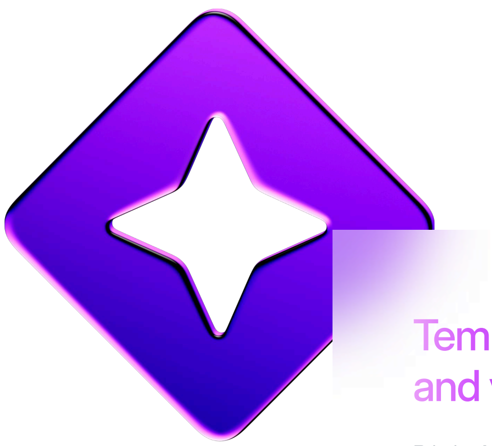

Temporal state and versioning

Digital twins are inherently time-based systems. They must answer questions such as: what is the current state? What was the state at a given point in time? How did the system evolve? What would happen if we revert to or branch from a previous configuration?

This requires robust handling of: event histories; state transitions; audit trails; versioned queries of past state; and scenario branching.

In advanced use cases - particularly those involving simulation or AI-driven optimisation - the ability to “rewind,” replay or fork system state becomes essential.

Multi-modal data representation

Digital twins do not operate on a single data type. They commonly combine: structured transactional data; semi-structured or document-based records; graph-like relationship networks; time-series sensor streams; geospatial information; and unstructured contextual inputs. Attempting to manage these in isolated systems can introduce latency, integration complexity and operational overhead.

An effective digital twin architecture must accommodate multiple data modalities while maintaining coherence and performance.

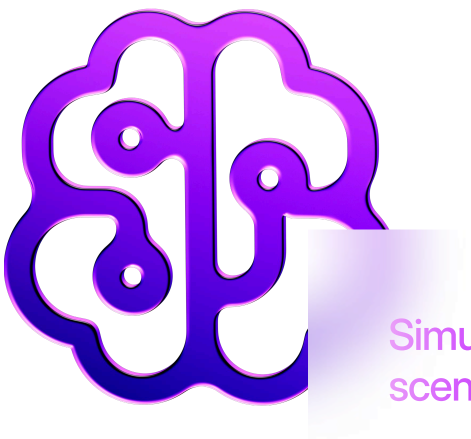

Simulation and scenario management

Beyond representing the current state, digital twins enable experimentation. This introduces additional architectural demands: parallel scenario environments; sandboxed simulation states; controlled execution environments; and comparison between simulated and live outcomes. Supporting scenario management requires the ability to isolate state changes, branch data structures, and evaluate hypothetical outcomes without corrupting production data.

This capability becomes increasingly important as AI systems operate within digital twins to test strategies before real-world execution.

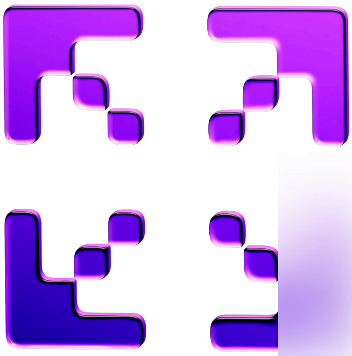

Scalability and distribution

Enterprise digital twins often span: multiple facilities; global supply networks; distributed energy assets, and national infrastructure systems. Data must be processed close to where it is generated while remaining synchronised across distributed environments.

Architectures must therefore support: horizontal scalability; geographic distribution; resilience under load; and fault tolerance. As digital twins extend across organisational and geographic boundaries, distributed design is no longer optional.

Security, governance and access control

Because digital twins frequently represent mission-critical systems - infrastructure, defence assets, financial systems, healthcare data - security is foundational.
Requirements often include: fine-grained access control; role-based permissions; data segregation; encryption at rest and in transit; and comprehensive auditability. In multi-stakeholder environments, governance frameworks must ensure that sensitive data remains protected while enabling collaboration.

Interoperability and integration

Few digital twins operate in isolation. They must integrate with: ERP systems; CRM platforms; manufacturing execution systems; IoT frameworks; AI and analytics platforms; and visualisation engines. Architectural flexibility and open interfaces are critical to prevent fragmentation, and an overly rigid data layer can quickly become a bottleneck as digital twin initiatives expand.

A structural insight

Across all industries and emerging use cases, a pattern becomes clear: a digital twin is not just a model of something - it is a living representation of a complex, interconnected, time-evolving system.

That representation must: capture relationships; track state changes; process real-time events; support simulation branching; scale across distributed environments; and enable secure collaboration.

Traditional single-model database approaches often struggle to support all of these requirements without introducing architectural sprawl - multiple databases for different data types, complex integration layers, and synchronisation challenges.[2]

As digital twins grow in scope, the choice of data backbone becomes a strategic decision rather than a technical afterthought.

[2] *openai.com/index/scaling-postgresql/*

Design principles for digital twin data platforms

As digital twins evolve from experimental projects to enterprise infrastructure, the underlying data platform must support a set of architectural principles that reflect the complexity of real-world systems.

1. Relationship-centric modelling

Digital twins represent networks of dependencies between assets, people, processes and environments. Platforms should support modelling relationships directly, allowing systems to trace dependencies and analyse cascading effects across interconnected components.

2. Multi-modal data support

Digital twins combine multiple data forms - structured records, semi-structured documents, relationship graphs and time-series telemetry. Effective platforms must support these data paradigms without forcing organisations to fragment their architecture across specialised systems.

3. Temporal awareness

Because digital twins represent systems evolving over time, platforms must support efficient handling of historical state, event streams and scenario branching. Understanding how systems change is as important as understanding their current state.

4. Real-time responsiveness

Operational twins depend on continuous data ingestion from sensors, enterprise systems and external data feeds. Platforms must be able to ingest and query real-time updates while maintaining transactional consistency.

5. Architectural simplicity

As digital twins expand across organisational boundaries, maintaining a simple and coherent data architecture becomes critical. Reducing integration layers and synchronisation overhead improves reliability and scalability.

These principles highlight an important reality:

digital twins are not defined by simulation tools alone. Their effectiveness depends on whether the underlying data platform can represent complex systems in a unified and flexible way. A new generation of databases has begun to emerge with these requirements in mind.

The limits of single-model architectures

Organisations implementing digital twins today typically adopt one of several architectural approaches. Some rely on specialised graph databases to model relationships across assets and dependencies. Others combine time-series platforms with relational databases to manage sensor telemetry alongside operational records. Cloud providers have also introduced managed digital twin frameworks that integrate simulation tools with IoT services.

Each of these approaches can support particular aspects of a digital twin system. However, as implementations expand beyond isolated use cases into enterprise-wide platforms, organisations often encounter increasing architectural complexity. Systems designed around a single data paradigm may require additional components to represent relationships, handle streaming data, or maintain transactional state across interconnected systems.

Consider a manufacturing twin:

machines and sensors generate time-series data

production runs generate transactional events

components depend on suppliers in a complex network graph

configuration files and asset descriptions may be semi-structured

simulated scenarios require temporary state branches.

If each of these elements is handled by a separate specialised system, the architecture quickly becomes fragmented. Data must be synchronised across platforms. Queries become cross-system operations. Latency increases. Governance becomes more complex. The digital twin becomes a coordination challenge rather than an intelligence platform.

At scale, a digital twin is:

a graph of relationships (dependencies, hierarchies, networks)

a document of state (attributes, configurations, metadata)

a stream of events (sensor updates, transactions, changes)

a temporal system (tracking how entities evolve over time)

a transactional platform (ensuring consistency and reliability)

This hybrid nature does not align neatly with single-model database design. Traditional enterprise systems were built around specific data paradigms. Relational databases excel at structured transactional records - orders, invoices, inventory tables. Document stores offer flexibility for semi-structured data. Graph databases specialise in relationship modelling. Time-series databases optimise sensor ingestion. Digital twins, however, require all of these capabilities simultaneously.

One of the hidden risks in digital twin programmes is architectural sprawl: the gradual accumulation of specialised tools stitched together through middleware and custom integrations. While polyglot persistence can be effective in narrow contexts, at enterprise scale it can introduce: operational complexity; data synchronisation risk; higher infrastructure costs; increased latency; fragmented governance; and slower innovation cycles.

Why multi-model databases are particularly suited

As digital twins mature from experimental pilots to enterprise infrastructure, organisations are discovering that the greatest constraint is often not simulation software, nor AI capability, but the structure of the underlying data systems.

> *Data issues are currently the top two challenges facing organisations: in particular, data quality, and data integration and interoperability.*[3]

Digital twins are fundamentally data environments. They model entities, relationships, events and states - continuously and at scale. When the data backbone is fragmented or overly rigid, the twin becomes difficult to evolve, integrate or scale. This is where architectural choices begin to matter strategically.

Multi-model architectures - which support graph, document and relational paradigms within a unified system - are increasingly aligned with the structural needs of digital twins. A unified multi-model data platform reduces these integration burdens by allowing diverse data structures to coexist within a coherent system. This simplification is not merely technical; it has organisational impact.

It enables:

Faster experimentation

Easier scaling across departments

Clearer governance frameworks

Lower operational overhead.

In environments where digital twins span multiple business units, simplification becomes a competitive advantage.

[3] *Hexagon: The Digital Twin Industry Report 2024*

Supporting AI and simulation

As AI systems increasingly operate within digital twins, they require access to highly connected relationship data; efficient querying of current and historical states; transactional reliability; and flexible schema evolution. Multi-model systems can allow AI agents to traverse relationships, analyse documents, evaluate temporal changes and execute controlled updates without crossing system boundaries. This becomes especially important when simulation environments branch from live systems and must later reconcile results.

The ability to manage these workflows within a unified data backbone reduces friction and improves reliability.

Strategic data foundations

Digital twin initiatives often begin as isolated projects - a factory optimisation programme, a supply chain resilience model, a city infrastructure pilot. Over time, however, organisations seek to connect these initiatives into system-of-systems platforms. At that stage, architectural constraints become visible. A data backbone that can represent relationships natively, evolve schema safely, support real-time ingestion, and operate across distributed environments becomes foundational rather than optional. Multi-model databases are emerging as strong candidates for this role.

Platforms designed from the outset to unify relational, document and graph capabilities within a single engine can align closely with the hybrid requirements of digital twins.

Among the new generation of databases, systems such as SurrealDB 3.0 have been designed with this unification principle at their core - enabling structured queries, native relationship modelling, and flexible data representation within a single, consistent framework. As digital twins become more ambitious, the alignment between architectural requirements and database capabilities becomes increasingly strategic.

> *"We stand at the edge of an informational revolution - one where the fusion of physical and virtual worlds will redefine how we live, work, and thrive.... By embracing innovation, fostering collaboration, and investing in secure data sharing infrastructure, we can unlock extraordinary value across every sector"*
>
> *-Mark Enzer, in TechUK Digital Twinning of Everything and Everyone*

> *"Data is becoming the new currency - and Digital Twinning will accelerate this shift. Just as social platforms monetised personal data, Digital Twins will do the same across industries, with real-time, high-value datasets driving innovation, simulation, and decision-making."*
>
> *-Tech UK - Business Model Proposition's 2030-2050*

Technical deep dive

The preceding sections have explored digital twins from a strategic and architectural perspective. This section provides a deeper technical examination of how digital twin systems can be modelled and supported in practice. For technical teams evaluating implementation pathways, the following considerations are central.

Modelling a digital twin as a graph of entities

At its core, a digital twin can be expressed as a graph:

- **nodes** represent entities - machines, components, suppliers, employees, patients, energy assets
- **edges** represent relationships - dependencies, ownership, connectivity, causation
- **attributes** capture state - temperature, configuration, capacity, financial value
- **events** record change - sensor updates, transactions, interventions.

Graph modelling enables traversal across relationships to:

- **identify upstream dependencies of a failed component**
- **trace supply chain exposure to a regional disruption**
- **map cascading impact across a grid.**

Native relationship support avoids expensive join operations and improves clarity of system modelling.

Temporal modelling and state evolution

Digital twins must capture how entities change over time.

Common technical patterns include:

- **event sourcing (storing immutable event logs)**
- **state snapshots with versioning**
- **hybrid event/state models.**

The system must support queries such as:

- "what was the configuration at 14:32 yesterday?"
- "simulate the system as it existed prior to the last policy change."
- "branch from this state to test an alternative scenario."

Efficient temporal modelling requires first-class support for time-indexed data and historical querying.

Real-time ingestion and event processing

In high-velocity environments (industrial IoT, smart grids, logistics networks), ingestion pipelines must:

- **accept high-frequency updates**
- **validate and normalise data**
- **maintain transactional consistency and avoid bottlenecks**

The persistence layer must sustain continuous writes while supporting concurrent analytical queries.

This dual workload - transactional plus analytical - is characteristic of digital twin environments.

Scenario branching and simulation isolation

Simulation requires the ability to:

- **fork system state**
- **apply hypothetical changes**
- **compare outcomes**
- **merge or discard branches.**

Architecturally, this may involve:

- **namespace and database isolation**
- **versioned data sets**
- **snapshot duplication strategies**
- **controlled write permissions.**

Database systems that allow flexible schema and controlled state evolution can simplify this process.

Distributed and edge architectures

Large-scale digital twins often combine:

- **edge data processing**
- **regional aggregation nodes**
- **centralised analytical environments.**

The data backbone must support:

- **synchronisation across nodes**
- **conflict resolution**
- **secure replication**
- **resilience under network disruption.**

This distributed design is particularly relevant in energy grids, defence applications and global supply chains.

Why unified multi-model engines matter technically

From a technical standpoint, unified multi-model systems can provide:

- **native graph relationships alongside document storage**
- **structured query capabilities without sacrificing flexibility**
- **strong typing and validation**
- **transactional guarantees**
- **reduced need for external synchronisation layers.**

For digital twin systems, this means:

- **relationship traversal without separate graph stores**
- **event ingestion without specialised time-series engines**
- **document flexibility without losing query power**
- **simplified AI data access patterns.**

SurrealDB 3.0 exemplifies this class of unified architecture, enabling developers to model entities, relationships and evolving states within a single engine and query language.

For organisations building digital twin infrastructure, this architectural coherence can reduce complexity while enhancing expressive power.

Technical perspective

Digital twins are not merely simulations; they are continuously evolving representations of complex systems. Supporting them effectively requires more than visual modelling tools or AI overlays. It requires a data backbone capable of: representing interconnected systems; tracking state across time; processing high-velocity events; enabling safe experimentation; and scaling across distributed environments.

As enterprises move toward intelligent, simulation-driven operations, the database ceases to be background infrastructure and becomes a strategic enabler.

How might a digital twin be represented in a unified data system?

Even a relatively simple industrial twin typically needs to capture several things at once, sometimes split across several systems:

- **the asset itself**
- **its current operating state**
- **its relationships to other components**
- **incoming sensor readings**
- **maintenance or operational events over time.**

Many current implementations combine specialised technologies - for example, graph databases for dependency modelling, time-series platforms for telemetry, and relational systems for transactional data. While effective in certain contexts, such polyglot architectures introduce integration overhead that unified data platforms aim to reduce. In a unified multi-model environment, however, they can be expressed within a single data layer.

The following illustrative example shows how a turbine, its sensor readings, and its upstream dependencies might be modelled.

Example

A simple industrial asset twin

The following illustrative example shows how a turbine, its sensor readings, and its upstream dependencies might be modelled.

```
-- Core asset table
DEFINE TABLE asset SCHEMAFULL;

DEFINE FIELD name       ON asset TYPE string;
DEFINE FIELD kind       ON asset TYPE string;
DEFINE FIELD status     ON asset TYPE "ordered" | "operational" | "decommissioned";
DEFINE FIELD location   ON asset TYPE string;
DEFINE FIELD updated_at ON asset TYPE datetime;

-- Sensor reading table
DEFINE TABLE sensor_reading SCHEMAFULL;

DEFINE FIELD asset_id    ON sensor_reading TYPE record<asset>;
DEFINE FIELD metric      ON sensor_reading TYPE string;
DEFINE FIELD value       ON sensor_reading TYPE number;
DEFINE FIELD recorded_at ON sensor_reading TYPE datetime;

-- Relationship table for dependencies between assets
DEFINE TABLE depends_on TYPE RELATION IN asset OUT asset;
```

This schema allows the system to represent both the state of an asset and the network of relationships around it.

A specific twin might then be populated like this:

```
CREATE asset:turbine_17 CONTENT {
    name: "Primary Turbine 17",
    kind: "turbine",
    status: "operational",
    location: "Plant A",
    updated_at: time::now()
};

CREATE asset:cooling_loop_3 CONTENT {
    name: "Cooling Loop 3",
    kind: "cooling_system",
    status: "operational",
    location: "Plant A",
    updated_at: time::now()
};

RELATE asset:turbine_17->depends_on->asset:cooling_loop_3
    CONTENT { critical: true };

CREATE sensor_reading CONTENT {
    asset_id: asset:turbine_17,
    metric: "temperature",
    value: 83.4,
    recorded_at: time::now()
};
```

Example query

Retrieving the twin context

A useful property of a digital twin is that it should be possible to query both the asset and its surrounding context in one place. For example:

```
SELECT
    id,
    name,
    kind,
    status,
    location,
    →depends_on→asset.{ id, name, status } AS upstream_dependencies
FROM asset:turbine_17;
```

This kind of query illustrates an important principle: the twin is not just a row in a table. It is a connected object with relationships, context and evolving state.

A second query might retrieve the most recent operating signals:

```
SELECT *
FROM sensor_reading
WHERE asset_id = asset:turbine_17
ORDER BY recorded_at DESC
LIMIT 5;
```

Together, these queries show how a digital twin can combine, within a single model:

current state

historical events

dependency relationships

live operational signals

This example is deliberately simple, but the principle scales.

The same pattern can be extended to represent:

production lines and their dependencies

supply chain nodes and disruptions

hospital equipment and patient pathways

energy assets and grid relationships

city infrastructure and transport networks

What matters is not only that the data can be stored, but that it can be traversed, updated, queried and extended without moving constantly between disconnected systems.

Digital Twin architecture

from physical systems to intelligent decision support

A useful property of a digital twin is that it should be possible to query both the asset and its surrounding context in one place. For example:

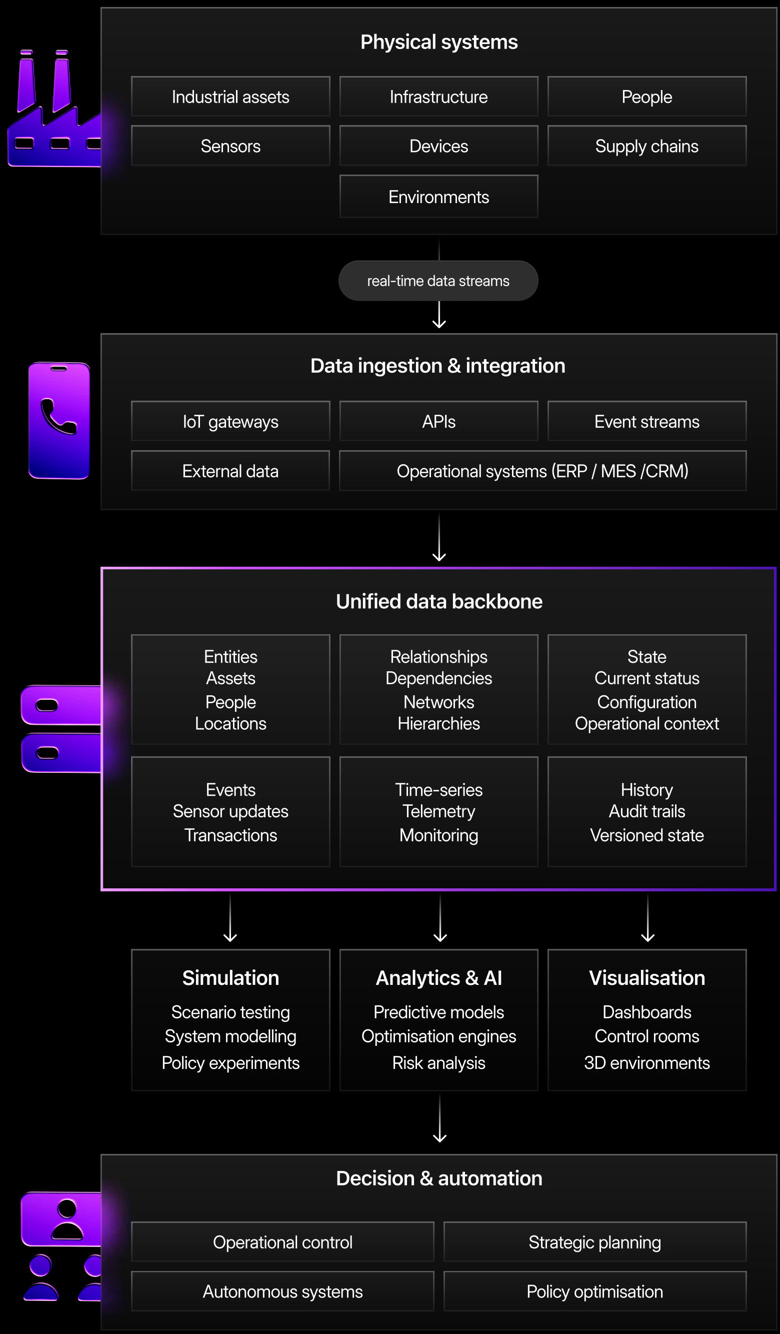

A digital twin architecture connects real-world systems with simulation, analytics and operational decision-making. At its centre is a unified data backbone capable of modelling entities, relationships, events and time-evolving state.

Enabling digital twins with a unified data backbone

Digital twin initiatives frequently begin with simulation software or AI platforms. Yet as implementations scale, organisations often discover that the most significant long-term constraint lies deeper - in the structure and flexibility of the underlying data layer.

A digital twin is, ultimately, a living data system. It must model entities and relationships, track state changes across time, process real-time events, support scenario branching, and maintain strong governance controls - often across distributed environments.

Meeting these requirements with fragmented or overly specialised data systems can introduce unnecessary complexity. Synchronisation overhead, latency, integration maintenance and governance challenges can dilute the intended agility of the twin itself. This is where unified, multi-model data platforms become particularly relevant.

Architectural alignment

SurrealDB 3.0 has been designed around the principle of unifying data paradigms within a single, coherent engine. Rather than requiring separate systems for relational records, document storage, or relationship modelling, it enables these representations to coexist natively.

This architectural alignment is particularly relevant for digital twin systems, which inherently combine: structured transactional data; rich relationship networks; semi-structured configuration and metadata; time-evolving state; event-driven updates. By supporting these data modalities within a unified framework, SurrealDB can simplify the data backbone that underpins digital twins.

Native relationship modelling

Because digital twins often depend on modelling complex interdependencies - supply chains, infrastructure networks, asset hierarchies - the ability to represent relationships as first-class constructs is critical. SurrealDB provides native relationship capabilities alongside structured querying, allowing interconnected systems to be expressed clearly without relying on external graph databases or complex join logic. For system-of-systems twins, where cascading dependencies must be traced efficiently, this capability can become strategically important.

Evolving state and flexibility

Digital twins are not static. They expand in scope, incorporate new data sources, and evolve as organisational needs change. SurrealDB’s flexible data modelling approach allows schemas to adapt without sacrificing query precision or integrity. This flexibility can support iterative digital twin development, where pilots mature into enterprise-wide platforms. As simulation requirements grow - particularly where scenario branching and temporal state tracking are involved - adaptable data structures reduce friction and development overhead.

Simplifying the data stack

Many digital twin architectures evolve into multi-database environments: one system for transactions, another for graph modelling, another for document storage, and additional layers for streaming and synchronisation. While such architectures can function, they increase operational complexity. A unified engine can reduce integration surfaces, streamline governance, and simplify deployment - particularly in distributed or edge environments.

For organisations seeking to build digital twins as long-term infrastructure rather than isolated projects, reducing architectural sprawl can materially improve sustainability and scalability.

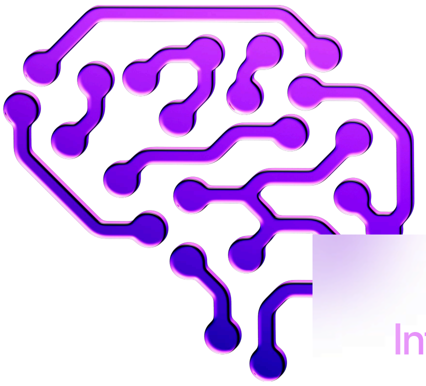

Infrastructure for intelligent systems

As digital twins increasingly serve as environments for AI-driven optimisation and autonomous decision support, the demands placed on the data layer intensify. Systems must support: high-throughput ingestion; relationship traversal; temporal queries; secure access controls; and distributed deployment.

SurrealDB 3.0 is engineered to address these dimensions within a single platform, offering an approach aligned with the emerging needs of intelligent, simulation-driven enterprises. Rather than treating the database as passive storage, digital twin architectures benefit from viewing the data backbone as an active enabler of system intelligence.

Building the intelligent infrastructure of the future

Digital twins are moving from experimental innovation to strategic infrastructure.

> *The global digital twin market size was valued at USD 24.48 billion in 2025 and is projected to grow from USD 33.97 billion in 2026 to USD 384.79 billion by 2034, exhibiting a CAGR of 35.40% during the forecast period.* [4]

What began as simulation tools for aerospace and manufacturing has evolved into a broader capability: the ability to model, monitor and optimise complex systems before acting in the physical world. Across manufacturing, supply chains, cities, healthcare, energy and finance, digital twins are enabling organisations to reduce uncertainty, improve resilience and make better-informed decisions.

> *"In an era of intensifying competition, skills shortages, and unstable demand, the need for modernised production facilities has never been greater. To stay ahead, companies must prioritise the integration of Digital Twins into their processes."*
>
> *-Philippe Delannoy, Dassault Systèmes*

At the same time, their scope is expanding. Digital twins are no longer limited to individual assets. They increasingly represent interconnected ecosystems - organisations, networks, infrastructure systems and environments. As artificial intelligence becomes embedded within these environments, digital twins are becoming platforms for experimentation, validation and controlled autonomy.

This evolution brings a structural shift. A digital twin is not merely a visual model or analytics dashboard. It is a continuously evolving representation of relationships, state changes, events and dependencies across time. Its effectiveness depends on the strength and flexibility of its data foundation.

Organisations embarking on digital twin initiatives face a strategic choice: whether to build on fragmented, specialised systems stitched together over time, or to establish a coherent data backbone capable of supporting multi-modal, interconnected and evolving systems from the outset.

As digital twins mature into enterprise-wide intelligence platforms, the database ceases to be background infrastructure and becomes a central enabler of capability. The organisations that recognise this - and invest accordingly - will be best positioned to operate with foresight rather than reaction, simulation rather than assumption, and intelligence rather than instinct. Digital twins are not simply a technology trend. They are a new layer of operational thinking - one that connects data, simulation and decision-making into a unified system.

Building that layer thoughtfully will define the next generation of intelligent enterprises.

[4] *www.fortunebusinessinsights.com/digital-twin-market-106246*

March 2026

Written by Tobie Morgan Hitchcock
surrealdb.com