McKinsey Digital

Digital twins: The foundation of the enterprise metaverse

Companies can leverage digital twins in a way that delivers significant value today—while building the engine for the enterprise metaverse of tomorrow.

This article is a collaborative effort by Joshan Abraham, Guilherme Cruz, Sebastian Cubela, Tomás Lajous, Kayvaun Rowshankish, Sanchit Tiwari, and Rodney Zemmel, representing views from QuantumBlack, AI by McKinsey and McKinsey's Digital and Technology, Media, and Telecommunications Practices.

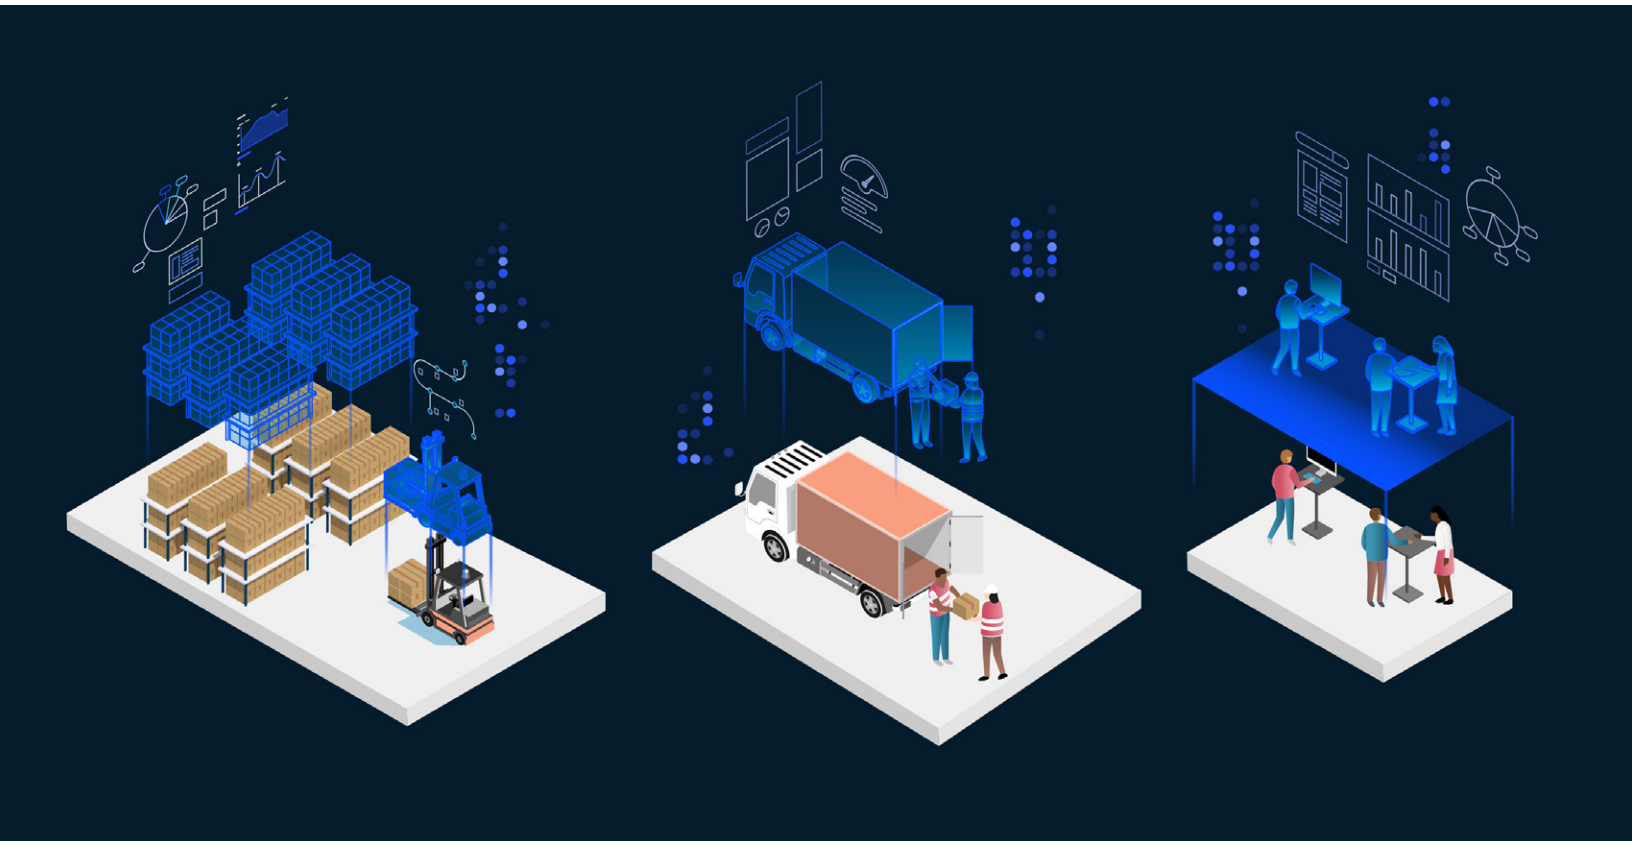

Illustration by Jing Zhang

October 2022

Organizations are set to look dramatically different in coming decades. We envision a world where the lines between physical and digital environments blur. Every asset, process, or person within and related to an enterprise will be replicated virtually—and connected. As a result, nearly every aspect of work can take place solely digitally or, at the least, before it does so physically. Immersive experiences, enabled by augmented reality (AR) and virtual reality (VR), will allow employees to gain real-world product design experience and training from their desks as they manipulate 3-D digital replicas of equipment. Mass simulations and AI technologies will use data streams from across and beyond the enterprise to help senior executives predict what's next with tremendous precision and prescribe the best course of action in even the most turbulent of times.

This is the enterprise metaverse—a digital and often immersive environment that replicates and connects every aspect of an organization to optimize experiences and decision making. The journey toward the enterprise metaverse, while still aspirational, has already begun with the development of the engines that will power it: digital twins.

Much of the discourse today has been about a commercial metaverse where consumers and their avatars can interact with brands and buy physical and digital products. However, we believe that the impact of the metaverse could be as high in the enterprise as in the consumer space—if not higher. The metaverse could enable optimized decision making from the C-suite to the front lines, customized and immersive employee and customer experiences enabled by AR and VR, autonomous AI use cases (such as proactively self-healing equipment) that aren't possible today, and entirely new product development mechanisms and processes.

Although it will take time for this vision to fully materialize, leading companies around the world are already building its foundations: digital twins. Ultimately, the enterprise metaverse will be powered by dozens of interconnected digital twins that replicate everything from physical assets (like products and office buildings) to people (such as customers and employees) to core business processes and often interact with the physical environment without human intervention (see sidebar "Imagine this future in the enterprise metaverse").

Sufficient technology and know-how exist to build digital twins today, and doing so creates significant value now rather than taking years. One telecom and technology player, for instance, reduced its capital and operating expenses by 10 percent thanks to a digital twin of its network assets. The twin can optimize capital spending, regulate usage patterns, identify failure points, and automatically initiate digital interventions based on unique network insights. Over the next decade, the company estimates that its digital twins will deliver billions of dollars in cumulative financial impact as they enable additional AI use cases and increase the amount of data-driven decision making across the organization.

In this article, we outline how companies at any level of digital maturity can begin the journey toward the enterprise metaverse by building out their foundation of value-driving digital twins today. It's important to note that the enterprise metaverse is not the goal in and of itself. We believe that organizations are best served by focusing on iterating, evolving, and adding twins only if they could have a real impact. In this way, many firms could ultimately arrive at their version of the enterprise metaverse.

From one digital twin to the enterprise metaverse

Companies can begin the journey by starting with just one digital twin that has a data product at its core, evolving it over time to provide increasingly powerful predictive capabilities. They can then move on to interconnecting multiple digital twins to unlock even more use cases and, finally, layer on the additional technologies required to transform this network of digital twins into the enterprise metaverse.

What is a digital twin?

There are many different definitions of digital twins, so it's important to start by answering this question. We see a digital twin as a virtual representation

Imagine this future in the enterprise metaverse

— A digital version of your end-to-end supply chain, from raw materials to delivery, continuously replicates in real time.

— It's linked to supplier information so it provides an early warning of disruption in a vendor's production capabilities for a key component in your flagship product. Managers receive a real-time report of existing inventory buffers, alternative suppliers, and comparable parts. After agreeing to find a new supplier, managers can simulate their vendor transition plans and choose the company that minimizes the impact of the changeover.

— Upon selecting the new source, vendor-onboarding and purchasing-order processes are automatically initiated. Now that a component from the new supplier has been chosen, the R&D organization receives a 3-D replica of it, and its impact on customers and existing processes is automatically simulated.

— Next, your virtual factory simulates any resulting production disruptions and gives leaders recommendations to ensure that production quality remains high by optimizing workforce and shipping schedules during the changeover.

— Your virtual retail store proactively sends store managers recommendations for updating store layouts and product mixes to fill any temporary gaps on shelves and for training employees so they're ready to answer customers' questions about any product changes.

— The result: the time the item is out of stock falls from months to days, financial costs are near zero, employees experience minimal disruptions, and customer satisfaction increases.

of a physical asset, person, or process. The twin comprises data collected from multiple sources, a layer of behavioral insights derived from the data, and visualizations. But a simple 3-D visualization or stand-alone simulation would not be considered a digital twin. Multiple AI use cases, "what if" simulations, and additional visualizations can be built on top of it.

As an example, a digital twin could provide a 360-degree view of customers, including all the details that a company's business units and systems collect about them—for example, online and in-store purchasing behavior, demographic information, payment methods, and interactions with customer service. It would also generate insights derived from the data, such as the average length of a customer service call. AI use cases leveraging the twin could include customer churn propensity models or a basket of the next products a customer would be likely to buy.

Alternatively, the twin might replicate the operation of real-world assets or processes (such as an entire factory production line or critical pieces of machinery) and generate information on average equipment downtimes or the average time for completing a product assembly. AI use cases it supports could include predictive maintenance and process automation and optimization.

Digital twins speed the time to market (and value) of many applications and use cases because development teams don't have to spend time cleaning and restructuring raw data every time they build an application. As a result, we often find that digital twins reduce the time needed to deploy new AI-driven capabilities by up to 60 percent and capital expenditures and operating expenditures by up to 15 percent. They can improve commercial efficiency by about 10 percent.

CEOs are increasingly recognizing the importance and power of digital twins and even mention them with growing frequency on earnings calls.[1] Our research indicates that 70 percent of C-suite technology executives at large enterprises are exploring and investing in digital twins. This interest, combined with rapidly advancing supportive technologies, is driving market estimates for digital-twin investments of more than $48 billion by 2026—a 58 percent compound annual growth rate.[2] Already, we're seeing some advanced implementations (see sidebar "Examples of recent digital twins").

The beginning: One digital twin

Which twin an organization builds first is determined by its priority value drivers and potential for reuse, balanced by business support and feasibility factors, such as the availability, quality, and accessibility of data. These are some examples:

A pharmaceutical company prioritized helping providers to improve patient outcomes. It started with a patient twin that enabled providers to deliver the right drug-safety content in the channel the patients would prefer at a given moment.

An automotive OEM started with parts twins to help it refine the most important features of the parts. The twins helped increase these parts' contribution margins by 5 to 10 percent.

The telecom and technology player mentioned earlier started with network twins, focusing initially on optimizing capital spending. To deploy newer use cases, such as predictive maintenance, this company has since evolved its twins to incorporate real-time performance and behavioral data.

Organizations typically start by building a data product, which is the core of a digital twin. A data product delivers a high-quality, ready-to-use data set that people across an organization can easily access and apply to different business challenges. Initially, a data product provides a few layers of data structured together to emulate an entity from the physical world and improve the performance of a critical business function. For example, to help

[1] Pratik Akhairamka, "What CEOs talked about in Q4/2021: Disrupted supply chains, vaccine mandates, and the metaverse," IOT Analytics, January 19, 2022.
[2] "Digital twin market worth $48.2 billion by 2026," MarketsandMarkets, June 2022.

Examples of recent digital twins

— *Emirates Team New Zealand* A digital twin of sailing environments, boats, and crew members enables the team to test boat designs without physically building them. The digital twin enables the sailing team of this reigning America's Cup champion to evaluate thousands of hydrofoil design concepts (instead of just hundreds) in the quest for a winning design.

— *Anheuser-Busch InBev* A brewing and supply chain digital twin enables brewers to adjust inputs (for instance, "add more hops to mixer #3") based on active conditions and can automatically compensate for production bottlenecks when, for example, vats are full. It also gives the company's production engineers remote assistance and augmented-reality capabilities for live troubleshooting on how to fix pump leaks and other common issues.

— *SoFi Stadium (in California)* To help optimize stadium management and operations, a digital twin aggregates multiple data sources, including information about the stadium's structure and real-time footfall data.

— *Space Force* This branch of the US Armed Forces is creating a digital twin of space, including replicas of extraterrestrial bodies and satellites (both US and foreign), as part of its commitment to be a digital-first organization.

— *SpaceX* A digital twin of the Dragon capsule enables operators to monitor and adjust trajectories, loads, and propulsion systems. The goal is to maximize reliability and safety during transport.

optimize employee scheduling, an organization might create a data product of its bench of retail employees, with information on their roles, availabilities, tenures, and store locations. A data product can often be created and deliver value in just a few months. It will continue to generate value by serving as a single source of truth, continually enhanced over time, that can serve as the basis for future use cases.

Once the data product is supporting a few use cases, its capabilities can be rapidly expanded. Teams can, for example, enrich it with layers of behavioral data, such as propensity models and likely interaction responses, by feeding data generated by the use of the data product back into it. Visualizations can also be added to enable simulations of different scenarios. These enhancements unlock more (and more powerful) predictive and prescriptive use cases and turn your data product into a digital twin.

An employee digital twin, for example, could help a company develop an AI-driven coach that provides real-time nudges to improve the performance and productivity of employees. It might also help managers identify the right combination of interventions to keep employee satisfaction high.

One chemical manufacturer prioritized building a factory process digital twin to optimize the production of enzymes for use by biofuel, agricultural, and consumer packaged goods companies in everyday products. The company's initial process twin structured data across a few key process parameters—such as average pH, airflows, and dissolved oxygen levels—to help engineers optimize the set points and schedules of enzyme production processes.

To further build out the twin, the company added even more data. For example, it ran agile two-week development sprints to incorporate

If customers don't enjoy the checkout process, the enterprise metaverse offers an opportunity to reinvent the experience in a digital context to achieve better outcomes.

sensor data for delivering live feeds of production status. It could then deploy automated tools and use cases to identify the root cause of any reductions in throughput. In one year, the digital twin enabled the company to increase production as much as adding an entire new production line would have done.

Compounding value by interconnecting digital twins

When companies interconnect two or more digital twins, they can simulate the complex relationships among different entities and generate richer behavioral insights for even more sophisticated use cases and greater value.

An organization might, for example, connect its employee twin with digital twins of its customers, retail stores, inventory, sales, and customer process flows. In this way, it could achieve outcomes such as these:

- simulating the end-to-end impact of business and market changes on retail stores
- creating a true omnichannel experience that provides for seamless pause-and-resume customer journeys across channels through improved data tracking
- optimizing store layouts by responding to shifting customer preferences
- assessing different compensation and staffing models by sales, employee performance, and the characteristics of local stores
- enhancing next-best-product recommendations with in-store data, such as browsing patterns and inventory levels

When these twins interconnect with each other, they become smarter as the dependencies and correlations get incorporated into the modeling.

Transforming the digital-twin network into the enterprise metaverse

As organizations begin to connect these multiple twins of different business domains, functions, and operational ecosystems, the enterprise metaverse could start to take shape. A retailer, for instance, could connect the digital twin of its retail store to digital twins of its warehouses, supply chain, call center, and more until every part of the organization was replicated, sharing insights, simulating scenarios, and enabling automation and AI use cases.

On top of the digital-twin foundation, companies could then build a layer that stitches together all the digital applications and analytics they've created on the back of the twins. APIs will integrate the system.

Finally, companies can add a unified consumption layer to give employees and customers integrated,

immersive experiences that leverage augmented and virtual reality.[3]

As implementations mature, leaders will want to shift from simply replicating what exists today to digitally reengineering entire processes and experiences, saving precious time and resources in the end. Say, for example, that customers don't enjoy the checkout experience or that engineers struggle with existing product design processes. The enterprise metaverse offers an opportunity to reinvent these experiences and processes in a digital context to achieve even better outcomes rather than replicate the existing experiences and processes.

How to build the first twin

Building and scaling a digital twin requires a three-step approach: creating a blueprint, building the initial digital twin, and then boosting its capabilities.

Create a blueprint

Aligning stakeholders on a clear vision of digital twins is a crucial first step. A blueprint should define the types of twins the organization will pursue, the order for building them to maximize value, the way their capabilities will evolve, and their ownership and governance structures (see sidebar "Potential priority twins, by industry"). Without all this, we've seen companies build disparate single-use twins with limited engagement by the business and no way to attribute value from use cases back to the twin. The answers to eight key questions can help leaders create this blueprint.

1. What is the overall vision for digital twins?
2. Which digital twins give our enterprise the greatest leverage and offer the best opportunities for reuse?
3. What is the total value at stake?
4. What are the highest-value, most feasible use cases we should deliver first, and what is the process to attribute value to digital twins?
5. What data layers and attributes do we need to collect?
6. How will we source and model the data?
7. What models will be built on top of the data? What is the end-state architecture?
8. How will the project team work together with business users to deliver the digital twin and use cases it supports?

Build the base digital twin

With the blueprint in place, the project team then builds the basic digital twin over three to six

[3] Although AR and VR have become synonymous with the commercial metaverse, in some cases they won't be essential in the enterprise metaverse. For example, AR and VR technologies will be necessary to enable employees to "enter" digital experiences such as maintenance staff physically interacting with virtual replicas of manufacturing equipment to gain repair experience right from the desk. However, AR and VR may not be needed for replicating, say, end-to-end customer service processes.

Potential priority twins, by industry

Manufacturing: Factory process, supply chain, equipment

Healthcare: Hospital, patients, nursing workforce

Telcos: Network, customers, call center

Retail: Supply chain, store, call center

Insurance: Customers, insured assets, claims process

Oil and gas: Refinery, equipment, workforce

Utilities: Powerplants, grid network, B2B customer equipment

Transportation: Planes, ships, trucks; logistics network

months. The build phase begins with assembling the core data product. To do so, data teams engineer structured and unstructured data to ensure their quality and usability. This in turn enables the development of visualizations and allows data science professionals to build out one or two initial use cases that generate additional data and insights—and create an early digital twin.

Organizations don't need perfect data or a state-of-the art technology platform to get started. We've seen companies with varying levels of data and platform maturity successfully build digital twins. There are, however, several keys to building a successful digital twin. They include the following:

Build a strong, balanced data model. Base the data model on the most demanding use cases and aggregate the data at the lowest levels of granularity. Strive for the appropriate level of accuracy in the data model. Higher accuracy comes with higher marginal costs. Lower accuracy will not have an impact.

Establish governance and benefits attribution from the start. Set up the right data governance and ownership model for using the digital twin, implement the process for attributing its impact, and manage the economics of all use cases built on top of the twin.

Ensure cross-team engagement. Dedicate people from business, IT, and data science teams, and ensure that technical and business teams create solutions and visualizations together. Assign a strong product manager who can lead the build team in an agile manner and persuade it and other stakeholders to accept an iterative approach.

Boost capabilities

Once the digital twin's initial use cases are up and running, it's time to expand its capabilities by adding more data layers and analytics to support new use cases. At this stage, companies often advance their twins from simply representing assets, people, or processes to providing simulations and prescriptions through the use of AI and advanced modeling techniques.

Digital twins have already become critical business tools for leading companies. However, the technology is accessible for any organization, no matter their level of digital sophistication. As a result, we expect digital twins will soon become key tools for optimizing processes and decision making in every industry. Executives are not only investing in digital twins today but also regarding the enterprise metaverse as a matter of "how soon" rather than "if." In the near term, these efforts can provide tremendous value, enabling companies to curate data—just once—for hundreds of use cases that deliver deep insights on complex business issues and optimize outcomes in real time. In the longer term, these investments lay the groundwork for the enterprise metaverse that will transform how every organization in every industry operates and unlock the full promise of data and AI.

Joshan Abraham is an associate partner in McKinsey's New York office, where **Guilherme Cruz** is a partner and **Tomás Lajous, Kayvaun Rowshankish,** and **Rodney Zemmel** are senior partners; **Sebastian Cubela** is an associate partner in the Miami office; and **Sanchit Tiwari** is a principal data scientist in the Chicago office.

The authors wish to thank Ignacio Ferrero and Tanay Jalan for their contributions to this article.

Designed by McKinsey Global Publishing