

WHITE PAPER

The Transformative Impact of the Industrial Metaverse

2025

By Tilman Buchner, Martin Kleinhans, Daniel Küpper, and Jonathan Brown



Contents

- 03** Introduction
- 04** Opportunities in Four Domains
- 05** Opportunity 1: Improving 3D Interoperability
In Engineering Processes
- 07** Opportunity 2: Improving the Accuracy and
Speed of Simulations
- 11** Opportunity 3: Enabling Robotic Systems to Perform
Human-like Tasks in Dynamic Environments
- 14** Opportunity 4: Leveraging Factory Data in
Operations with LLMs and Factory Ontologies
- 15** Starting the Journey to Success
- 16** Conclusion
- 17** About the Authors

Introduction

Digital technology has been transforming manufacturing for decades, yet significant challenges persist throughout the factory life cycle. From product and factory design to automation and operations, producers often struggle to achieve seamless integration and optimization.

Recent technological innovations, however, have the potential to reshape the world of manufacturing through a new concept: the Industrial Metaverse. In essence, the Industrial Metaverse is a persistent platform that serves as a digital replica of an organization and its operational environment. It integrates machine, process, product, and human data in a bi-directional flow between the physical and virtual worlds.

Crucially, the Industrial Metaverse is not a single technology but a concept that captures the convergence of a wide range of IT and operational technologies (OT). These include cloud and edge computing, the Industrial Internet of Things (IIoT), predictive and generative AI, digital twins, augmented and virtual reality, and open standards. Accelerated computing is a key enabler of the complex ecosystem forming around these technologies. In the Industrial Metaverse's most advanced stage, Embodied AI integrates AI with robotics to create physical intelligence—making it possible to automate handling tasks that are too versatile and complex for traditional robotics.

In this paper, we explore the vast potential of the Industrial Metaverse and how its adoption can redefine the future of manufacturing. By embracing this paradigm, factories of the future will no longer be constrained by physical prototyping or incremental improvements. Instead, they will be “born in simulation”—designed, tested, and optimized in virtual environments—before operating in reality.

Opportunities in Four Domains

In the Industrial Metaverse, advanced digital models serve as the ultimate source of truth, accurately mirroring the real world and continually updating in real time to reflect changes. These digital models enable data continuity across the factory lifecycle, providing a foundation for scaling digital and AI use cases effectively. The urgency for such a transformation is clear: a recent [BCG study](#) revealed that only 26% of companies have succeeded in scaling AI initiatives beyond pilots and prototypes.


The Industrial Metaverse offers the tools needed to unlock the full potential of digital transformation in manufacturing across four critical domains: engineering, simulation, automation and robotics, and operations.

As shown in Exhibit 1, our analyses find that each of these areas stands to benefit significantly:

- 30% improvement in planning efficiency through enhanced engineering processes
- 40% reduction in quality costs enabled by synthetic data generation
- 25% decrease in conversion costs via cognitive robotics
- 21% productivity boost in factory operations driven by GenAI

EXHIBIT 1

The Industrial Metaverse Offers Four Key Opportunities in Manufacturing



Sources: BMW's "Neue Klasse" electric vehicle production facility in Debrecen, Hungary (Engineering); BCG Automotive Tier-1 client (Simulation); BCG study (Automation & Robotics, Operations); BCG (images).

COPQ = cost of poor quality

Opportunity 1: Improving 3D Interoperability in Engineering Processes

The design and commissioning of factories requires seamless collaboration among diverse disciplines—mechanical design, electrical engineering, software development, and simulation. Hundreds of experts from different fields, often spread across company boundaries, must coordinate their efforts. Yet fragmented workflows and incompatible IT systems frequently hinder progress.

This lack of integration means that engineers struggle to anticipate the downstream effects of their design decisions. Realizing the vision of “Design for X”—whether for assembly, automation, repair, or recycling—requires a bi-directional solution that bridges these gaps.

This challenge is well-known, but existing solutions (such as the Jupiter Tessellation file format) have fallen short of expectations owing to the problem’s technical complexity. Although these approaches simplify data exchange across heterogeneous toolchains, they fail to address the problem comprehensively and have seen limited adoption. As a result, there is no effective interchange service to notify engineers of design changes or their downstream implications, nor to enable feedback from downstream processes to inform design decisions.

A truly holistic solution must address two fundamental challenges. First, a unified hierarchical model must organize and link elements such as geometric information, simulation data, and factory automation information, accounting for their interdependencies. Second, a mechanism is needed to interpret changes in the model, enabling different applications with diverse data formats to stay synchronized.

A parallel can be drawn with the emergence of HTML, the standard that transformed the internet into a cohesive and interoperable platform. Until recently, manufacturing lacked such a standard. The introduction of the Universal Scene Description (USD) format now offers a game-changing solution. (See “USD Evolves from Animation to Manufacturing.”)

USD Evolves from Animation to Manufacturing

The USD format was originally conceived by Pixar Animation Studios in 2012 to address the growing complexity of digital scenes in animated films. Designed to streamline collaboration and manage intricate scenes across various applications and studios, USD proved indispensable for projects like Toy Story 4. This blockbuster, with its millions of models, textures, lights, and colors, would not have been possible without USD’s robust framework.

In August 2023, Pixar took a significant step by announcing the Alliance for OpenUSD and making USD open source. In addition to well-known software vendors such as Autodesk and Adobe, the open-source alliance includes Foxconn, NVIDIA, and Siemens, introducing USD to the manufacturing world. As a result, the format has expanded beyond its original focus on geometry and textures to include material properties, part numbers, and IoT data—key elements for industrial applications.

USD is much more than a file format. It is a comprehensive, high-performance interchange service and ecosystem for describing, composing, simulating, and collaboratively constructing 3D scenes in near real time.

Leveraging the USD standard, an open ecosystem that fosters adoption and interoperability formed around the comprehensive NVIDIA Omniverse platform. Industrial software vendors (ISVs)—such as Siemens, Autodesk, Visual Components, and Ansys—have contributed by building connectors that link Omniverse to their proprietary software portfolios. These connectors enable engineers and developers to extend interoperability across diverse content tools, data types, workflows, and platforms.

The increased transparency provided by the adoption of such a platform can significantly improve planning efficiency. For instance, BMW reported efficiency gains of 30% within their processes.

In a nutshell, the combination of the USD standard, the NVIDIA Omniverse platform, and its open ecosystem of ISVs serves as a critical enabler for seamless interoperability across organizational boundaries. (See “Use Case Example: Bridging Design and Automation with an Integrated Toolchain.”)

Use Case Example: Bridging Design and Automation with an Integrated Toolchain


Exhibit 2 illustrates how the integration of Siemens, NVIDIA, and AWS technologies creates a seamless toolchain that transforms manufacturing workflows from product design to factory automation, and vice versa. For example, if engineers adjust the radius of a component during the product design phase, they can immediately understand the downstream implications of this change on factory automation processes, such as robotic handling tasks, in real time. In this setup, a collision detection service uses simulation to identify potential conflicts, such as collisions between robotic arms and other objects.

This process is facilitated by tools like Siemens NX, which generates a USD file containing essential product updates about geometry, kinematic models, and material properties. The NVIDIA Omniverse platform enables collaborative, real-time work by distributing updates to a scene across connected applications. This allows the implications of changes in one system (for example, design software) to be understood in another (such as automation software), and vice versa. Shifting industrial software applications and the Omniverse platform to the cloud (for example, AWS) significantly enhances real-time collaboration within this solution.

The true potential of this approach is realized by scaling it across an entire organization and, if necessary, its suppliers—enabling unprecedented levels of efficiency and coordination throughout the manufacturing process.

EXHIBIT 2

The Integration of NVIDIA, Siemens, and AWS Technologies Creates a Seamless Toolchain



Source: BCG

Opportunity 2: Improving the Accuracy and Speed of Simulations

The domain of simulation encompasses a wide array of tools that support various engineering disciplines. These include finite element methods, computational fluid dynamics, and plant and manufacturing process simulations. The Industrial Metaverse promises to revolutionize this domain through two major advances: closing the simulation-to-reality gap and enabling new computational methods that accelerate simulation.

Closing the Simulation-to-Reality Gap

A transformative application of simulation technology is in the field of vision systems. Simulations can generate synthetic images of potential quality defects, which are then used to train vision systems before deployment. However, this approach has been hindered by the persistent simulation-to-reality gap. (See “Understanding the Gap.”)

Understanding the Gap

The gap between simulated environments and the real world can be attributed to two main factors: appearance and content.

Appearance. Synthetic (computer-generated) images often differ from real-world images at the pixel level. These discrepancies are caused by limitations in rendering technologies, which struggle to accurately replicate the complexity of objects, materials, and their physical properties. For example, in a synthetic rendering of a car’s instrument panel, the textures and reflections may appear less realistic compared to an actual photograph. This lack of realism can diminish the effectiveness of simulations in applications requiring high visual precision.

Content. Synthetic environments often lack the diversity and contextual richness found in real-world scenarios. This is due to limited variability in objects, scenes, and behaviors within simulations. For instance, a simulated cityscape may exclude spontaneous dynamics, such as random pedestrian movements or changing weather conditions. Such omissions hinder the ability of AI systems trained on synthetic data to generalize effectively to real-world conditions, especially in use cases like urban navigation or scenario planning.

The potential opportunity to close the simulation-to-reality gap is widely recognized. A recent [BCG study](#) revealed that 40% of respondents identified continuous optimization of AI models in simulation as a critical long-term priority over the next five to ten years.

Recent enhancements in realistic light reflections and textures, paired with the capability to easily modify and randomize every aspect of a scene, have significantly improved the accuracy of synthetic training datasets in simulations. (See “Use Case Example: Accelerating Industrial Inspection with Synthetic Data Generation.”)

Use Case Example: Accelerating Industrial Inspection with Synthetic Data Generation


Machine vision systems in industrial inspection often struggle with two critical issues: missing actual defects (false negatives) and generating pseudo-errors (false positives). These challenges stem from systems' reliance on rigid rule-based approaches, which are ill-suited to handle the variety of defects and fluctuating environmental conditions often present in manufacturing. Creating perfect rule sets for every scenario is nearly impossible. These issues ultimately lead to manual inspections and higher quality costs. A recent BCG client engagement highlighted that up to 40% of quality costs could be saved by eliminating pseudo-errors.

AI-based models could offer a solution. However, training such systems requires vast quantities of annotated image data, which is costly and time-consuming to generate and unavailable at the start of production.

Closing the simulation-to-reality gap solves this problem and enables full automation. Photorealistic images of components or products can be generated in simulations based on existing computer-aided design (CAD) models from engineering and automatically annotated, as shown in Exhibit 3. This synthetic training data can then replace real images for training highly accurate AI models.

EXHIBIT 3

Photorealistic Images Can Be Generated in Simulations




Source: BCG

Exhibit 4 highlights how synthetic data generation can be fully automated through scripting. The process begins by importing a 3D design model into applications developed with the NVIDIA Omniverse platform. The Omniverse platform software then facilitates the generation of defects or image variations. With the help of scripts, parameters such as the width, depth, or distribution of scratches on a surface can be precisely defined.

Once configured, the system automates the generation of thousands of annotated images—up to 2,000 in this example—with the push of a button. These synthetic datasets can then be used to train machine vision systems for quality control.

EXHIBIT 4

Fully Automating Synthetic Data Generation Through Scripting of Defect Simulations



Source: BCG (images)

GUI = graphical user interface

Using New Computational Methods to Boost Speed

Recent breakthroughs in accelerated computing are driving significant advancements in simulation capabilities. Traditional simulation tools rely on numerical methods to solve complex mathematical equations, such as those in fluid mechanics. Although these methods are effective, they are computationally intensive and often require days to complete.


Emerging technologies, such as Physics-Informed Neural Networks (PINNs), are addressing this challenge. PINNs leverage AI to emulate the physical behavior of systems, offering significant speed improvements without sacrificing accuracy. This breakthrough makes simulations economically viable for a broader range of applications, including renewable energy, weather forecasting, and health care. (See Exhibit 5.) For instance, Siemens Gamesa utilized PINNs to accelerate the wake optimization of wind farms by a factor of 4,000. AI frameworks such as NVIDIA PhysicsNeMo are enabling such advancements.

EXHIBIT 5


Emulating the Physical Behavior of Systems Dramatically Boosts Simulation Speeds

Training neural networks using both data and physical equations

Problem: Harmonic oscillator




Neutral network



Training step: 970

— Exact solution
— Neutral network prediction
● Training data

Physics-informed neutral network



Training step: 14550

— Exact solution
— Neutral network prediction
● Training data
● Physics loss training locations

Physics-informed neural networks are broadly applicable

Renewable Energy

Siemens Gamesa: 4,000x speed-up of wind turbine wake optimization




Climate Change

FourCastNet: 45,000x speed-up of extreme weather prediction




Health Care

Achieve high-fidelity results for blood flow in inter-cranial aneurysm



Industrial High-Performance Computing

NETL: 10,000x faster build of high-fidelity surrogate models



Source: NVIDIA; BCG

Crucially, AI will not replace simulation in the future. Instead, these technologies will increasingly work together, complementing each other's strengths. For instance, simulations can generate synthetic training data to build and refine AI models (simulation-assisted AI) and AI models can accelerate particularly complex calculations (AI-assisted simulation). This powerful combination paves the way for achieving the vision of factories and products born in simulation and operating in reality, unlocking significant cost savings.

Opportunity 3: Enabling Robotic Systems to Perform Human-like Tasks in Dynamic Environments


Robotics is a critical component for the future of manufacturing. In a recent [BCG study](#), more than 90% of respondents identified advanced robotics and production automation as core elements of their future factory vision.

Despite this strong emphasis, many handling operations remain too versatile and complex for traditional industry robotics. (See Exhibit 6.) This limitation arises because robots lack an awareness of their surroundings and rely on rigid, rule-based programming. For every task a robot is expected to perform, its control logic must define specific rules. However, because the real world is highly dynamic, it is impossible to predefine a set of rules covering all potential scenarios a robot may encounter.

EXHIBIT 6

Many Operations Are Too Versatile for Traditional Automation

Robotic systems lack a situational understanding and causal model of the real world



Sources: Siemens; BCG

To understand the complexity of the problem, consider what happens when a robot is asked to open a bottle. While current machine vision systems can identify the bottle in an image or video feed, completing the task requires far more than object recognition. The robot must:

- Understand what the object is, including its parts (for example, the cap and neck).
- Know how to interact with it, anticipating how the bottle will react based on its material properties and physical behavior.
- Plan a sequence of actions (for example, gripping, twisting, and removing the cap).
- Dynamically adjust its workflow to account for unexpected changes or errors.

To perform these tasks, robots require “physical intelligence”—the ability to perceive, understand, and interact with the physical world. This advanced capability can be provided by Embodied AI, which integrates AI with physical entities, such as robots, allowing them to learn from and respond to their environments. Groundbreaking innovations in Embodied AI can be grouped into three critical areas: visual perception, workflow planning, and object manipulation.

Visual Perception

Visual perception is a prerequisite for enabling robots to understand how to interact with objects. This understanding goes beyond simple object recognition to include identifying an object’s orientation in three-dimensional space.

Effectively building this capability today requires a combination of advanced technologies, including photorealistic rendering, generative pre-trained transformer (GPT) models, and simulation environments. Together, these technologies create Large Vision Models (LVMs), which are pre-trained on diverse, often synthetically generated datasets. LVMs can handle tasks with minimal or no additional training (that is, few-shot or zero-shot learning), a critical feature for enabling robots to work with previously unseen objects or in unfamiliar environments.

This innovation is particularly transformative for applications such as machine tending in sheet metal processing. The traditional vision systems utilized to recognize components require costly retraining—up to \$10 per component type—whenever the geometry changes. This makes flexible automation economically unfeasible for custom or small-batch production. By leveraging one-shot learning, new LVMs can reduce adaptation costs to nearly zero, making flexible automation viable.

Although vision perception is critical for enabling robots to build a comprehensive scene understanding, this capability alone does not provide the situational understanding required for truly dynamic robotic performance.

Workflow Planning

“We need machines that understand the world; [machines] that can remember things, that have intuition, have common sense—things that can reason and plan to the same level as humans,” asserts Yann LeCun, Meta’s chief AI scientist. Although no current AI systems meet these criteria, various startups and tech companies are working on building systems that would possess such an understanding of how the world works and the ability to reason about the consequences of actions.

These companies are developing “world models” that aim to provide a basic representation of a given environment, such as a warehouse. Robotic systems can use such representations to achieve specific goals. For example, a system tasked with unloading a shipment could use a world model to develop a sequence of actions to accomplish that goal: using the forklift to unload, cutting off the banderole, and opening the packages. The system would not mimic patterns it has seen but apply a deeper understanding to reason through the steps required for unloading.

Table: Comparing the Capabilities of Large Vision Models and World Models

Feature	Large Vision Model	World Model
Focus	Processing visual data	Modeling the dynamics of an environment
Data Input	Images, videos (often static)	Multisensory, time-dependent
Applications	Computer vision (understanding images)	Reinforcement learning, robotics, simulation
Objective	Understanding visual information	Predicting behavior and planning
Temporal Dynamics	Mostly static (image by image)	Considers temporal dynamics

The most prominent approaches in this area are Fei-Fei Li’s World Labs, OpenAI’s video generator Sora, and Google DeepMind’s Genie. However, these approaches face significant challenges, including the issue of hallucination, common in LLMs. Training data for a world model must strike a balance between being broad enough to cover diverse scenarios and specific enough to allow the AI to deeply understand the nuances of those scenarios.

Once consistent and reliable environmental models can be generated, AI systems will be better equipped to bridge the gap to the real world. This advancement will enable robots to achieve a level of situational awareness essential for executing complex workflows and achieving dynamic goals.

Until this point is reached, however, several innovative approaches are significantly improving robots' planning capabilities and enabling programming to become increasingly implicit rather than explicitly defined. Notable advancements include the following:

- Automatic Domain Randomization (ADR), pioneered by OpenAI, is a simulation-based technique that exposes robots to automatically generated, increasingly complex variations of their environment. This approach improves a robot's ability to generalize to real-world conditions. In manufacturing, ADR could transform robotics by enabling systems to handle a broader range of unpredictable tasks, such as assembling diverse components with precision or adapting to dynamic assembly line conditions.
- Google X's Robotic Transformer Model (RT-1) leverages large-scale datasets and transformers to train general-purpose robots, emphasizing real-world applicability across diverse robotic tasks with end-to-end learning.
- Meta's work on concept learning via Video Joint Embedding Predictive Architecture (VJEPA) explores how robots can infer abstract concepts and generalize effectively from minimal data, enabling creative problem-solving in unfamiliar scenarios.

While ADR focuses on domain variability for robust task performance, RT-1 prioritizes scalability and real-world adaptability, and VJEPA delves into higher-level abstraction and conceptual understanding. Together, these innovations highlight a convergence of approaches that could revolutionize robotics in manufacturing, merging adaptability, scalability, and conceptual reasoning.

Object Manipulation

Enhancing a robotic system's ability to manipulate objects in the real world is essential for expanding its functionality. This capability spans simple pick-and-place operations to highly complex, contact-rich assembly processes that require advanced technologies such as soft-body simulation. Robotics still faces significant constraints in contact-rich assembly and soft-body simulation, where precise interaction with complex, deformable, or irregular materials is critical. Contact-rich assembly, such as in electronics manufacturing or assembling intricate machinery, requires robots to manage fine tolerances and dynamic forces. For example, inserting fragile components like a pin into a delicate connector demands advanced force control and accurate sensing, which remain challenging.

Similarly, soft-body simulation—essential for tasks like handling food, fabric, or human tissue in medical robotics—is challenged by computational complexity and the need to accurately model soft material deformation.

Current advancements, such as tactile sensors, real-time physics engines, and machine learning algorithms, have enabled robots to simulate and adapt to the behaviors of soft materials. Although these technologies promoted progress in tasks like folding clothes or performing minimally invasive surgeries, achieving human-level dexterity and reliability remains an unmet goal.

The near future, however, holds promise. PINNs (discussed above) and hybrid models (combining real-time simulations with learned priors) have the potential to transform fields such as flexible manufacturing and delicate object handling.

Impact: Simplifying Implementation, Expanding Possibilities

As Embodied AI matures, it will enable smarter, faster, and more adaptable factories. By introducing physical intelligence, automation systems will shift to implicit programming, allowing robots to learn from and collaborate with humans. Imagine a robot that seamlessly transitions from assembling intricate electronics to inspecting and sorting parts—all without reprogramming.

This innovation not only broadens the scope of handling tasks that robots can perform but also significantly reduces the time and effort required to engineer robotic solutions. Implementation times could be cut by 50%, and automation engineering hours reduced by 70%. These efficiencies will democratize access to robotics, making them a viable option for smaller manufacturers. The resulting ability to automate tasks that were previously manual is a key driver of the 25% reduction in conversion costs enabled by the Industrial Metaverse.

Opportunity 4: Leveraging Factory Data in Operations with LLMs and Factory Ontologies

Although Industry 4.0 has revolutionized the manufacturing sector, the full potential of data and analytics remains unrealized. The introduction of LLMs, such as ChatGPT, has reignited optimism about accessing and leveraging factory data through GenAI. These models excel at understanding user intent, as well as analyzing unstructured and semi-structured data, identifying patterns, and extracting insights.

However, the fragmented and siloed nature of data within a manufacturing environment—spanning machinery, production lines, logistics, and enterprise systems—poses a significant challenge. Without a structured framework to map relationships between data objects and processes, LLMs struggle to generate meaningful factory insights. Additionally, they face difficulties in grasping hierarchical and specific relationships inherent to complex industrial data if they lack proper context or predefined structures.

To harness the full potential of LLMs in manufacturing, companies should invest in building factory ontologies. A factory ontology systematically organizes and defines the relationships between all available factory data. Currently, much of this knowledge remains uncoded, residing within the experiential knowledge of employees who understand how and where information is stored and the interconnections between datasets. Without a formal structure, complex, multi-dataset analysis is difficult, if not impossible.

A promising solution involves combining an LLM with an industrial Unified Namespace (UNS)—a centralized, structured real-time data broker that integrates all factory data with contextual relationships. By providing a semantic and contextual framework, the UNS enables the LLM to interpret and reason about the relationships and processes effectively. This integration contributes to an estimated 21% productivity boost in factory operations driven by GenAI.


A UNS also solves many of the inherent limitations of the classical point-to-point connections typically used in the automation pyramid. (See Exhibit 7.) It simplifies scalability by centralizing data exchange, enhances contextualization with structured information, and ensures real-time availability for faster decision-making. Additionally, it promotes interoperability through open standards, streamlines maintenance with a decoupled architecture, and enables advanced analytics by unifying data for IIoT and AI applications. These benefits make UNS essential for modern, scalable, and intelligent manufacturing systems.

Combining LLMs with a factory ontology enables advanced reasoning on factory data by bridging human intent with machine-readable data structures. LLMs excel at interpreting user intents expressed in natural language, transforming vague or complex queries into actionable instructions. Meanwhile, the factory ontology provides a structured mapping of data relationships, such as machine hierarchies, process flows, and material dependencies, offering the necessary context for precise data interpretation and analysis. Together, this synergy allows operators and systems to query factory data intuitively, extract meaningful insights, and automate decision-making processes with a deep understanding of interconnected factory systems. This approach enhances efficiency, reduces errors, and paves the way for smarter, more adaptable manufacturing operations.


EXHIBIT 7

A UNS Increases Interoperability on the Shopfloor by Standardizing Communication and Data Exchange Between Applications

From: Hierarchical and point-to-point connectivity



To: Event-based, hub-and-spoke topology



Source: BCG

ERP=enterprise resource planning; MES=manufacturing execution system; PLC=programmable logic controller; SCADA = supervisory control and data acquisition

Starting the Journey to Success

Embracing the transformation to the Industrial Metaverse requires a clear understanding of its concepts and the enabling technologies, capabilities, and processes. Given the scope and complexity, companies should take the following structured steps:

- 1 Set an ambition and vision.** Define your organization's ambition level and vision for adopting the Industrial Metaverse. Clearly outline what success looks like and establish measurable objectives.
- 2 Assess the current state.** Analyze the status quo to identify existing value leakages and inefficiencies. Pinpoint the most impactful use cases to address through the Industrial Metaverse.
- 3 Design the strategy.** Develop a target picture, strategy, and road map for implementation. This begins with evaluating and prioritizing the most beneficial use cases to ensure maximum impact. Additionally, companies need to build an ecosystem of partners, as no single vendor or organization can deliver a comprehensive solution. Addressing existing gaps is also critical, requiring a focus on both people- and technology-related measures. This includes launching capability-building programs, initiating upskilling efforts, and making strategic make-or-buy decisions to bridge any resource or expertise gaps.
- 4 Implement pilots.** Test new processes, platforms, and applications through pilot projects. Focus on achieving tangible results while iterating and improving as quickly as possible.
- 5 Scale proven solutions.** Roll out successful use cases across the organization. Expand the transformation by launching additional pilot projects to address new opportunities.

By following this step-by-step approach, companies can navigate the complexity of adopting the Industrial Metaverse and build a scalable foundation for success.

Conclusion

The Industrial Metaverse represents a transformative leap for manufacturing, blending advanced technologies like digital twins, GenAI, and robotics with the power of immersive, interconnected environments. By enabling factories to be “born in simulation and work in reality,” this paradigm unlocks unprecedented opportunities for efficiency, adaptability, and innovation. Success will hinge on strategic planning, developing the right ecosystem of partners, and investments in both technology and talent. Front-running organizations can redefine the future of manufacturing—creating smarter, more resilient factories that not only address today’s challenges but also shape industrial operations for decades to come.

About the Authors



Tilman Buchner
Partner & Director
Munich
Buchner.Tilman@bcg.com



Martin Kleinhans
Project Leader
Munich
Kleinhans.Martin@bcg.com



Daniel Küpper
Managing Director & Senior Partner
Cologne
Kuepper.Daniel@bcg.com



Jonathan Brown
Managing Director & Partner
Hamburg
Brown.Jonathan@bcg.com

For Further Contact

If you would like to discuss this report, please contact the authors.



Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://www.bcg.com). Follow Boston Consulting Group on [LinkedIn](#), [Facebook](#), and [X \(formerly Twitter\)](#).

