DGS
EMBRACE INNOVATION

From Generative AI to Physical AI: When Artificial Intelligence Interacts with the Real World

EMBRACE INNOVATION

Contents

After AI that generates comes AI that operates 04

What is Physical AI? 05

Opportunities and Responsibilities 06

Physical AI and Security 07

World Models: The Evolutionary Horizon 08

From language to experience 09

When Intelligence enters operations 12

Physical AI is an organizational capability 15

What can be done today 17

Beyond the Digital Twin 19

From sim-to-real to sim-to-real under governance 22

The trust layer of Physical AI 25

The four enabling capabilities of Physical AI 27

What changes for CIOs, COOs, CISOs, and Innovation Leaders 30

Key questions for the board 32

Operational Roadmap 34

A look ahead to World Models 36

The next generation of AI will be situated, governed, and operational 38

DGS
EMBRACE INNOVATION

Francesco Di Bianco

Offering Leader Physical AI di DGS

"The next evolution of artificial intelligence depends on the ability to integrate AI, physical systems, and operational knowledge into ecosystems capable of understanding the real world before interacting with it."

Elena Dalle Cort

Solution Consultant di DGS

"Digital twins, OT systems, sensor technologies, and AI are converging toward a new operational paradigm: an intelligence capable of interpreting context, anticipating the consequences of actions, and supporting decision-making in the physical world."

Foreword

Physical AI marks the shift from systems that interpret information to systems capable of understanding operational context and interacting with the real world. Value emerges when data, processes, digital twins, OT systems, and industrial knowledge converge within an integrated ecosystem, transforming distributed information into reliable decisions and context-aware actions. The foundations of this evolution are being built today through integration, governance, observability, and control.

After AI that generates comes AI that operates

Over the past few years, artificial intelligence has transformed the way organizations create content, process information, access knowledge, and support decision-making processes. The most recent phase of this evolution has been driven by Generative AI and Large Language Models (LLMs), which have made interactions with data, documents, language, code, and enterprise knowledge more immediate and accessible.

Today, a new technological shift is emerging.

> **Artificial intelligence** is no longer limited to generating text, images, insights, or instructions. It is progressively moving beyond the purely digital domain and entering physical processes, **industrial environments**, **infrastructure**, **logistics**, **maintenance**, **safety**, **security**, and **operations**.

EMBRACE INNOVATION

What is Physical AI?

This is the realm of Physical AI: an evolution of artificial intelligence in which software systems, sensors, actuators, robots, digital twins, intelligent agents, Operational Technology (OT) platforms, and simulated environments work together to create systems capable of perceiving their context, interpreting it, anticipating its evolution, and supporting decisions or actions in the real world.

This evolution should not be viewed merely as an advancement in robotics. Physical AI represents a new phase of enterprise AI: the transition from intelligence that interprets information to intelligence that operates within complex, physical, and dynamic environments.

For organizations, the challenge extends beyond the adoption of new technologies. It involves the ability to integrate AI, data, infrastructure, industrial systems, security, governance, and operating models into a single, reliable architecture.

Physical AI combines software systems, sensors, actuators, robots, digital twins, and OT platforms to understand and operate in the real world.

Opportunities and Responsibilities

Physical AI opens up significant opportunities: more advanced predictive maintenance, intelligent quality control, autonomous inspections, operator assistance, production process optimization, adaptive infrastructure management, and greater operational resilience.

This evolution should not be viewed merely as an advancement in robotics. Physical AI represents a new phase of enterprise AI: the transition from intelligence that interprets information to intelligence that operates within complex, physical, and dynamic environments.

> For organizations, the challenge extends beyond the adoption of new technologies. It involves the ability to integrate AI, data, infrastructure, industrial systems, **security**, **governance**, and **operating models** into a single, **reliable architecture**.

OPPORTUNITIES

1. Predictive maintenance
2. Quality control
3. Operator support
4. Operational resilience

RESPONSIBILITIES

1. People
2. Assets
3. Facilities
4. Critical infrastructure

EMBRACE INNOVATION

Physical AI and Security

With Physical AI, security is no longer limited to protecting data, applications, and digital infrastructure—it extends directly into the physical world. Intelligent models, autonomous agents, edge systems, and Operational Technology (OT) platforms can influence the behavior of machines, industrial facilities, and operational processes.

In this context, cybersecurity must ensure data integrity, model protection, identity management, edge security, and the traceability of actions performed by intelligent systems.

The objective is to ensure that every decision and every action remain verifiable, attributable, and governable, rather than simply preventing unauthorized access or technological compromise.

It is precisely this ability to maintain control over data, models, agents, and behaviors that links security to the principles of digital sovereignty.

World Models: The Evolutionary Horizon

World Models represent the most advanced evolutionary horizon of this trajectory: models capable of building a dynamic representation of the world, simulating future scenarios, understanding causal relationships, and predicting the consequences of actions. However, they are not a mandatory starting point.

The value of Physical AI can already be realized today through mature technologies and progressive paths of experimentation, simulation, governance, and operational adoption.

> **The challenge for organizations is to start now: to build the technological, organizational, and trust foundations that will enable AI to enter operations in a sustainable, secure, and scalable way.**

From language to experience: why AI must learn from the real world

The first wave of generative artificial intelligence found its primary domain in language. Language models demonstrated the ability to synthesize information, create content, generate code, support analysis, interact with users, and make knowledge more accessible.

This transformation has profoundly changed the relationship between people, data, and digital systems.

> It has brought AI closer to the day-to-day reality of organizations and accelerated the adoption of intelligent tools across a wide range of business workflows.

Yet this form of intelligence remains, by its very nature, largely disembodied.

Language models learn from the symbolic record of human experience: text, images, code, conversations, documentation, and both structured and unstructured data. They can describe a machine, interpret a process, recommend an action, or generate a procedure.

However, they do not interact directly with physical objects, operational constraints, environmental conditions, real-world failures, or material consequences.

> Physical AI emerges precisely from this transition: bringing artificial intelligence into a more direct relationship with the real world.

This is not simply about enhancing a model's ability to recognize patterns. It is about building systems capable of perceiving their environment, understanding context, evaluating scenarios, anticipating consequences, and supporting decisions that have an operational impact.

Generative AI

Understands data
Generates outputs
Digital domain

Physical AI

Understands context
Supports actions
Physical domain

This evolution is particularly relevant in sectors where value is generated not only through information management, but through the ability to improve physical processes, including manufacturing, utilities, energy, logistics, infrastructure, transportation, healthcare, regulated environments, and mission-critical operations.

In these contexts, artificial intelligence cannot be viewed as a software layer separated from the rest of the organization. It must be integrated with assets, facilities, field data, OT systems, control platforms, operational procedures, safety standards, and organizational responsibilities.

This is where Physical AI meets the domain of Industrial AI: the application of artificial intelligence to operational, production, and infrastructure processes, with the goal of improving efficiency, resilience, quality, adaptability, and decision-making capabilities.

The distinction is not marginal. In the information domain, AI supports understanding. In the physical domain, AI must contribute to action. And every action demands a higher level of reliability, control, and accountability.

KEY INDUSTRIES

1. **Manufacturing**
2. **Utilities**
3. **Energy**
4. **Transportation**
5. **Logistics**
6. **Healthcare**

EMBRACE INNOVATION

When Intelligence enters operations: robots, sensors, agents, and physical systems become part of the same cognitive stack

Physical AI is often associated with advanced robotics, autonomous systems, or humanoid machines. While understandable, this perspective is only part of the picture.

The real transformation is not about the robot itself. It is about the ability to build an ecosystem in which artificial intelligence, sensors, actuators, OT systems, software agents, digital twins, simulation platforms, edge computing, and cloud technologies work together to connect perception, context, decision-making, and action.

From this perspective, Physical AI should be understood as a new enterprise architecture.

Physical AI reaches its full potential when it becomes an integral part of the industrial ecosystem rather than an isolated technology initiative.

To make accurate decisions and operate effectively in the physical world, intelligent systems must have access to a comprehensive view of the operational context, integrating data and processes across the entire value chain—from design and Product Lifecycle Management (PLM) to production planning, from ERP and MES systems to industrial assets, sensors, and Operational Technology (OT) infrastructure.

INDUSTRIAL ECOSYSTEM

PLM → ERP → MES → OT → AI

AI GENERATED

In this context, Industrial AI represents the integration layer that connects intelligent models, industrial processes, and enterprise systems, transforming fragmented data into shared operational knowledge.

> Physical AI creates value when supported by end-to-end integration across data, processes, and enterprise systems, bringing together design, planning, production, and operations within a single ecosystem capable of turning data into actionable operational intelligence.

Only through this unified view can intelligent systems be provided with the context required to understand their environment, make consistent decisions, and act effectively in the physical world.

For this reason, Physical AI cannot be approached solely as an artificial intelligence or robotics initiative. It requires the convergence of multidisciplinary capabilities—including generative and agentic AI, robotics, industrial automation, system integration, data engineering, enterprise platforms, and cybersecurity—to connect data, processes, and technologies across the entire industrial lifecycle.

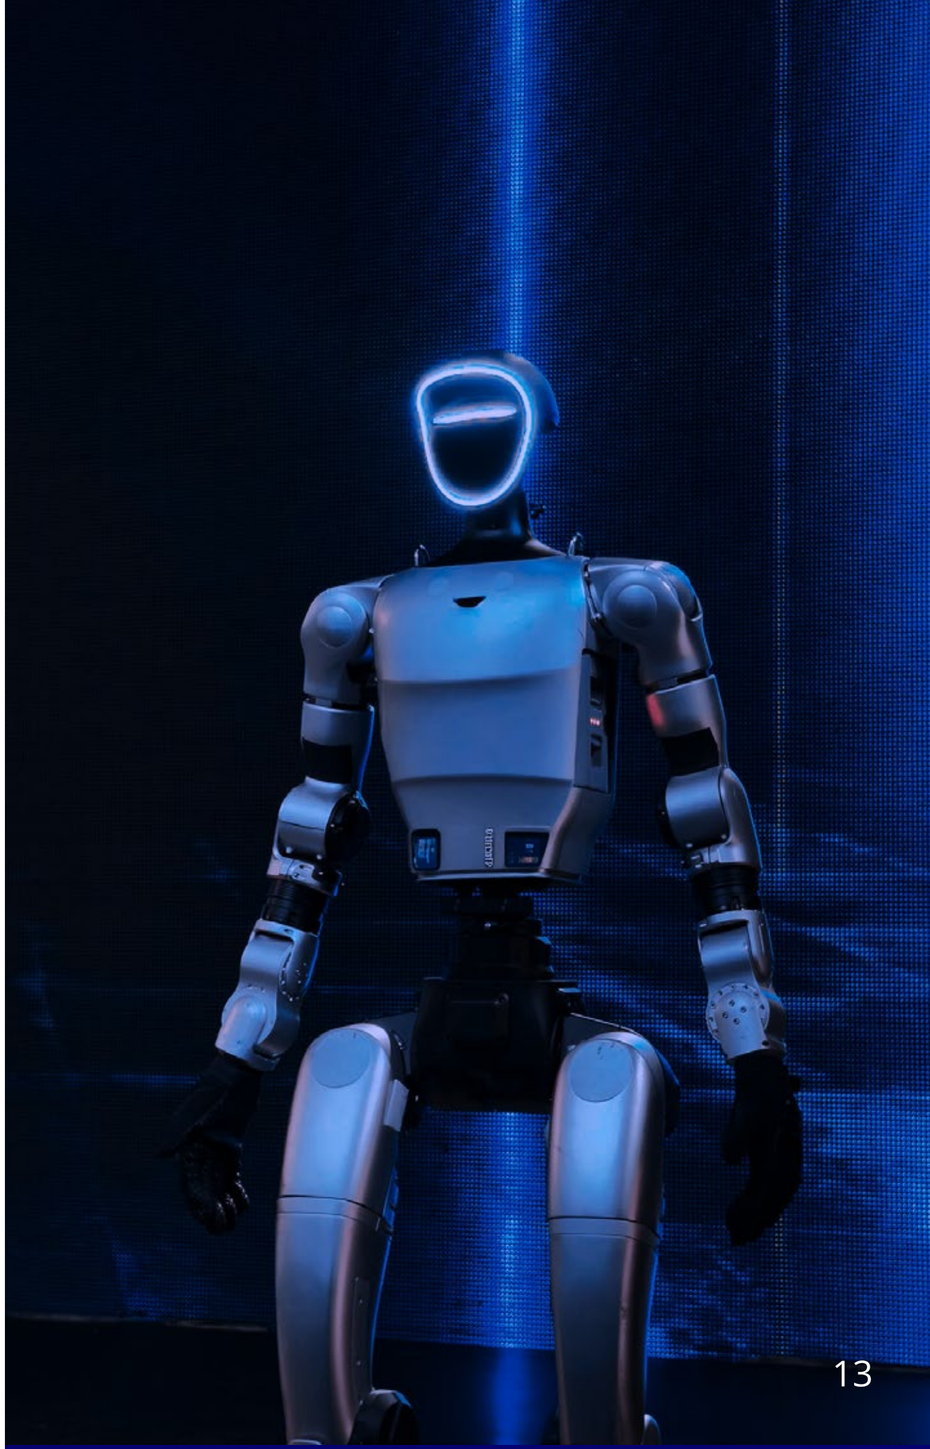

It is precisely this ability to integrate the full production ecosystem that enables the development of reliable, governed, and scalable cyber-physical systems, transforming Physical AI from a technological experiment into a tangible driver of business competitiveness.

Within this architecture, multiple components coexist:

1. **AI and multimodal models**
2. **Computer vision systems**
3. **Sensors and physical data sources**
4. **IoT and Industrial IoT devices**
5. **OT and industrial systems**
6. **Digital twins and simulation environments**
7. **Edge and cloud platforms**
8. **Robotics and automation technologies**
9. **Software agents**
10. **Monitoring, control, and security systems**

EMBRACE INNOVATION

Physical AI is an organizational capability, not just a technology

Physical AI is not simply the next stage in the evolution of robotics or artificial intelligence. It is the convergence of disciplines that have traditionally operated separately—including AI, robotics, industrial automation, cybersecurity, OT systems, simulation, and data engineering—which must function as a single, coherent system.

The true critical success factor lies not in any individual algorithm, but in the ability to orchestrate these capabilities across the entire lifecycle of an intelligent physical system.

AI is no longer an isolated software component. It is part of a complex ecosystem that encompasses:

1. **AI**
2. **robotics**
3. **OT**
4. **networking**
5. **edge computing**
6. **cybersecurity**
7. **digital twin**
8. **safety**
9. **human factors**

> This is a **fundamental point**: **Physical AI** cannot be approached as a stand-alone technology innovation project. **It requires a systemic capability**.

> Organizations seeking to bring AI into physical processes will need to adopt an integrated approach, in which traditionally separate disciplines converge throughout the entire system lifecycle—from design and simulation to integration, testing, deployment, monitoring, maintenance, security, and governance.

This also implies an organizational transformation

> Physical AI does not belong to a single business function. It is not solely the responsibility of IT, operations, innovation teams, or cybersecurity departments. It is a cross-functional domain that requires collaboration among professionals with diverse expertise and shared objectives.

This is why the concept of AI Systems Engineering is becoming increasingly important. It is no longer enough to develop intelligent models; organizations must design intelligent systems.

Designing an intelligent system means considering not only model performance, but also data quality, infrastructure reliability, connection security, integration with existing processes, behavior under abnormal conditions, the possibility of human intervention, decision traceability, and long-term operational sustainability.

DGS
EMBRACE INNOVATION

What can be done today: Physical AI is not just a future promise

One of the most common risks when discussing Physical AI is to immediately focus on the most advanced scenarios: humanoid robots, general-purpose autonomy, systems capable of learning from the real world with human-like flexibility, and fully mature World Models seamlessly integrated into industrial processes.

While this perspective is important, it should not overshadow a key point for organizations: many of the building blocks of Physical AI are already available today and can deliver tangible business value.

> Physical AI does not require organizations to wait for the full maturity of World Models to become useful.

Value can already be created through the integration of existing and increasingly mature technologies, including digital twins, computer vision, simulation systems, industrial sensors, advanced automation, predictive models, control platforms, intelligent agents, and edge solutions.

EMBRACE INNOVATION

The range of applications is broad

MANUFACTURING

Physical AI can support quality control, production line optimization, defect identification, predictive maintenance, and operator assistance

LOGISTICS

it can contribute to intelligent flow management, material handling, space optimization, and coordination between automated systems and human operators

UTILITIES INFRASTRUCTURE

it can enable monitoring, anomaly detection, automated inspections, distributed asset management, and improvements in operational continuity

ENERGY

it can support the management of complex facilities, failure prevention, and process optimization in highly critical environments

HEALTHCARE

it can enhance operational environments, internal logistics, intelligent devices, and highly controlled processes

Today, the real value does not necessarily lie in envisioning fully autonomous systems. It lies in developing progressive, measurable, and well-governed use cases.

Physical AI can begin with specific processes and clearly defined objectives: reducing errors, increasing efficiency, improving safety, anticipating anomalies, supporting operational decisions, minimizing downtime, enhancing quality, and strengthening the resilience of complex environments.

From this perspective, the topic should be approached without excessive futurism. **Physical AI is not just a vision of what may come**. It is a journey that organizations can begin today, particularly when supported by controlled experimentation, co-creation, testing, and a gradual approach to operational readiness.

Beyond the Digital Twin: from descriptive models to contextual understanding

The digital twin has been one of the key enablers of industrial transformation. It has allowed organizations to represent assets, processes, and environments in digital form, enabling monitoring, simulation, predictive analytics, and optimization.

Through a digital twin, an organization can observe the behavior of a system, collect data, simulate scenarios, and improve decision-making.

This capability has already represented a significant leap forward compared to more traditional approaches to asset and operations management.

Physical AI, however, introduces a further step forward.

A digital twin describes, replicates, and monitors a system.

Physical AI seeks to leverage that representation to better understand the operational context and support decisions that produce tangible effects in the real world.

> This shift is not only technological.
> It is cognitive.

It is no longer enough to know what is happening. Organizations must be able to understand what could happen if a variable changes, if a machine operates under abnormal conditions, if a robot alters its trajectory, if an agent recommends a specific action, if an operator interacts with an intelligent system, or if a process is optimized in real time.

This is where causality becomes critical.

An industrial system must do more than classify an object, recognize an image, or detect a deviation from a standard. It must understand the consequences of actions. It must be able to evaluate alternatives, estimate outcomes, recognize constraints, anticipate deviations, and support decision-making in dynamic conditions.

The distinction is substantial.

Recognizing an anomaly means identifying a deviation.
Understanding the context means interpreting why that anomaly occurs, what consequences it may generate, which actions are available, and what risks they entail. **This is the trajectory along which World Models are emerging.**

RECOGNIZE

Identify a deviation.

UNDERSTAND

Interpret causes, impacts, and consequences.

> Understanding the context means interpreting why the anomaly occurs.

EMBRACE INNOVATION

World Models can be viewed as the next evolution of simulation and representation models: systems capable of building an internal representation of their environment, continuously updating it through sensory data, and using it to simulate future scenarios.

Looking ahead, World Models may represent for Physical AI what Large Language Models have represented for Generative AI:

> a cognitive infrastructure capable of generalizing, simulating, and enabling new forms of operational intelligence.

From sim-to-real to sim-to-real under governance: from simulation to real-world action must be controlled, observable, and secure

Many technology innovation initiatives demonstrate promising results in simulated environments, prototypes, or controlled settings. The real leap occurs when these systems are deployed in real-world environments.

In the case of Physical AI, this transition is particularly delicate.

A system may function correctly in a laboratory. It may perform well in simulation. It may demonstrate advanced capabilities in a protected environment. But once it enters operations, it must contend with far more complex variables: people, machines, critical assets, environmental conditions, regulatory constraints, legacy systems, OT networks, security requirements, business procedures, and operational responsibilities.

The question, therefore, is not simply: Does it work? The question becomes:

> Can we trust it when it acts in a real-world environment?

EMBRACE INNOVATION

> This is where the concept of Sim-to-Real Under Governance emerges.

The transition from simulation to reality cannot be managed as a simple technical handoff. It must be **controlled**, **observable**, and **secure**.

This means establishing validation, testing, monitoring, and intervention mechanisms throughout the entire lifecycle of the system.

1. **Validation**
2. **Testing**
3. **Monitoring**
4. **Supervision**
5. **Alert**
6. **Observability**

It means defining which behaviors are acceptable, which thresholds must never be exceeded, which data must be tracked, which events should trigger an alert, which actions require human oversight, and under what conditions the system's autonomy must be restricted or halted.

From this perspective,

> **simulation** is no longer merely a design tool. **It becomes a governance capability**.

It enables organizations to test scenarios, anticipate risks, assess performance, validate behaviors, stress-test edge conditions, and reduce uncertainty before deployment.

This issue is particularly critical because Physical AI operates in environments where failure is not merely a software error. It can become a matter of safety, operational continuity, accountability, compliance, or reputation.

> For this reason, organizations will need to develop a new discipline for adoption: moving beyond proof-of-concept initiatives toward **progressive pathways** of **readiness**, **validation**, and **control**.

The transition to real-world environments will require dedicated testing environments, evaluation criteria, governance frameworks, reliability metrics, observability systems, and response capabilities.

Today, for many advanced Physical AI scenarios, the regulatory and certification landscape is still evolving.

However, the growing adoption of intelligent systems designed to interact with physical assets, infrastructure, and people is accelerating the need for more structured standards, validation methodologies, and certification frameworks.

> In this context, the ability to design solutions according to the principles of security, observability, traceability, and control already represents a critical factor in organizational readiness.

Organizations that begin today to develop the necessary skills, validation capabilities, and governance models will be better prepared to adopt the future certification and compliance frameworks that will accompany the widespread deployment of Physical AI.

READINESS FRAMEWORK

1. **Testing**
2. **Governance**
3. **Certification**
4. **Compliance**
5. **Observability**
6. **Risk Management**

EMBRACE INNOVATION

The trust layer of Physical AI: governance, security, and observability become prerequisites for scaling AI in the physical world

As AI moves into the physical world, the nature of risk changes as well.

> In the **information domain**, an error may result in an inaccurate response, an incorrect piece of content, a recommendation that requires verification, or a decision that needs to be reconsidered.

> In the **physical domain**, however, the consequences may affect assets, people, facilities, infrastructure, operational continuity, and safety.

Physical AI introduces new areas of exposure:

1. **manipulable sensors**
2. **compromised actuators**
3. **vulnerable OT systems**
4. **opaque multimodal models**
5. **autonomous agents**
6. **unpredictable interactions between software and the physical environment**
7. **context-driven decisions that may be difficult to interpret**
8. **field data that is not always complete, accurate, or reliable**

For this reason, cybersecurity must be embedded into the architecture from the outset rather than added as an afterthought.

The trust layer of Physical AI encompasses policies, control mechanisms, monitoring, validation, data protection, OT security, operational accountability, auditability, and incident response capabilities.

At the same time, growing attention is being given to requirements related to technological sovereignty and value chain control.

Particularly in the most critical industrial and infrastructure environments, organizations increasingly favor approaches that provide greater transparency, stronger data control, and greater independence in the strategic components of their technology ecosystem. Although this area is still evolving, the trend is expected to play an increasingly important role in shaping the architectural and governance choices of future Physical AI systems.

> Trust is becoming an infrastructure layer in its own right.

For many organizations, the ability to scale Physical AI will depend less on the availability of advanced algorithms and more on their ability to establish security, governance, and accountability mechanisms suited to emerging cyber-physical environments.

> This is a critical shift: Physical AI cannot be framed solely as a way to automate processes or improve efficiency.

It must be understood as the capability to embed intelligence into physical processes while maintaining control, reliability, and accountability.
The competitive advantage will not belong solely to those who adopt these technologies first, but to those who can make them reliable, governable, and sustainable over time.

A Physical AI system must be observable, understandable, stoppable, correctable, and secure. This is the fundamental condition that allows AI to move beyond digital workflows and into real-world processes without increasing the organization's overall risk profile.

The four enabling capabilities of Physical AI: understand, simulate, protect, act

To move Physical AI from experimentation to operations, organizations need a simple framework that connects strategic vision with operational execution.

One possible approach is built around four enabling capabilities:

> understand, simulate, protect, and act.

These are not strictly sequential phases, but four interconnected dimensions that must coexist throughout the lifecycle of an intelligent physical system.

EMBRACE INNOVATION

UNDERSTAND

> Understanding means mapping the operational environment in which the system will operate.

It involves identifying assets, processes, data, sensors, physical constraints, OT systems, human interactions, operating conditions, and risk factors.

This is the first step in ensuring that Physical AI is not treated as an isolated technology project.
Before introducing intelligent systems into the real world, organizations must understand where they will operate, which data they will rely on, which processes they will interact with, what dependencies exist, and what potential consequences their actions may have.

Contextual understanding is also a prerequisite for building trust. A system may be highly advanced from an algorithmic perspective, but if it is not integrated into the operational environment, it risks delivering limited value or introducing new vulnerabilities.

SIMULATE

> Simulation means creating scenarios, testing behaviors, evaluating outcomes, and reducing risk before deployment in production environments.

Simulation makes it possible to observe how a system may behave under different conditions, including rare, anomalous, or difficult-to-test situations that cannot easily be replicated in the real world.

In Physical AI, simulation is not only a design tool. It is a mechanism for building trust, measuring performance, identifying anomalies, validating responses, and preparing for deployment into operational environments.

The more a system is expected to interact with people, critical assets, or complex infrastructure, the more simulation becomes an essential governance capability.

EMBRACE INNOVATION

PROTECT

Protection means securing models, data, OT systems, agents, physical interfaces, decision-making processes, and communication channels.

In Physical AI, protection and governance are not secondary considerations. They are prerequisites for scale.

Security must encompass the entire ecosystem: from data quality and sensor protection to OT network resilience and model robustness, from access management to the ability to detect anomalous behavior.

Without security, observability, and control, AI in the physical world remains a fragile experiment.

ACT

Acting means bringing AI into operations through mechanisms for supervision, continuous monitoring, human oversight, and response management.

The objective is not to automatically replace the role of the organization with autonomous systems. Rather, it is to enhance the organization's ability to operate in a smarter, more adaptive, and more secure way.

The degree of action must be proportional to the maturity of the system, the level of risk within the environment, and the criticality of the process involved.

In some cases, Physical AI will provide recommendations. In others, it may automate specific tasks. In more advanced scenarios, it may orchestrate complex systems. In every case, however, the levels of autonomy, oversight, accountability, and intervention must be clearly defined.

This framework helps make the journey less abstract and more operational. It demonstrates that Physical AI is not merely a technological frontier, but a transformation that must be built methodically and with purpose.

EMBRACE INNOVATION

What changes for CIOs, COOs, CISOs, and Innovation Leaders

Physical AI is a cross-functional issue that spans operations, technology, and risk management. It does not belong to a single stakeholder within the organization.

For this reason, one of the most common mistakes would be to position it exclusively within a specific function:

1. innovation
2. IT
3. operations
4. industrial automation
5. cybersecurity

In reality, Physical AI cuts across all of these domains.

For CIOs, it represents an integration challenge. Data, platforms, legacy systems, cloud and edge infrastructures, applications, and enterprise architectures must interact seamlessly with physical systems, operational environments, and distributed data sources.

For COOs, it represents a powerful lever for process transformation. Maintenance, quality, productivity, operational safety, business continuity, and resilience can all be reimagined through systems capable of understanding context and supporting action.

For CISOs, it introduces a new cyber-physical attack surface. The challenge is no longer limited to protecting data, applications, or digital infrastructure, but extends to systems that perceive, make decisions, and interact with the physical world.

For innovation leaders, it represents a high-potential yet high-responsibility area of experimentation. The challenge lies in building tangible use cases while avoiding isolated, demonstrative initiatives that are difficult to scale.

Physical AI therefore requires a shared governance model. It cannot be managed solely as an innovation project, an IT initiative, or an evolution of industrial automation.

It is a cross-disciplinary challenge in which operational value, risk, security, compliance, adoption sustainability, and scalability must be evaluated together.

This cross-functional nature makes Physical AI particularly relevant at the board level.

Bringing AI into the physical world means making decisions that go far beyond efficiency gains. It requires addressing broader questions of **accountability**, **security**, **resilience**, and the **future operating model** of the enterprise.

DGS
EMBRACE INNOVATION

Key questions for the board: what needs to be kept under control

Before bringing AI into the physical world, organizations need to understand what must be controlled.

For any enterprise, approaching Physical AI should not begin with technology. **It should begin with the right questions**.

Before adopting solutions designed to interact with physical environments, boards and executive teams should establish a clear understanding of objectives, scope, risks, and governance requirements.

Several questions become essential:

> Which processes could become more adaptive through Physical AI?
>
> Which physical or operational activities could be supported by intelligent systems?
>
> Which decisions can be delegated, supported, or automated?
>
> Which sources of physical data are available, and what is their level of quality?
>
> Which sensors, OT systems, or digital twins are already in place across the organization?
>
> Which critical assets or processes could be affected?
>
> Which regulatory, operational, or safety constraints must be considered?
>
> How will system behavior be validated before go-live?

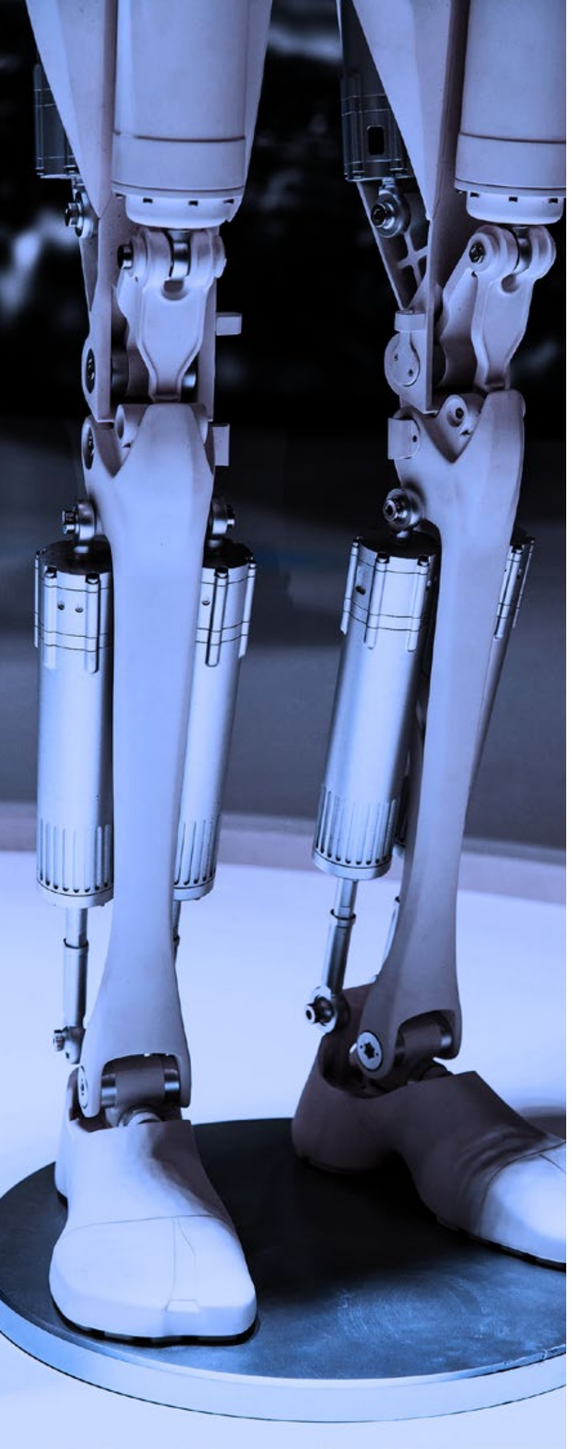

DGS
EMBRACE INNOVATION

How will the system be monitored after deployment?

Who is accountable if an AI agent behaves unexpectedly?

Which controls must be in place to stop, correct, or constrain system behavior?

How will sensors, models, data, and actuators be protected?

Which skills and capabilities are required to govern the solution throughout its lifecycle?

Which use cases can deliver tangible value today?

Which pilot initiatives can help prepare the organization for the future evolution toward World Models?

These questions help shift the conversation from a technology adoption mindset to one of **organizational readiness.**
They reinforce the idea that Physical AI is not simply a technology initiative, but a transformation that requires **risk awareness, governance capabilities, and a long-term strategic vision.**

EMBRACE INNOVATION

Operational Roadmap: from experimentation to readiness – preparing for enterprise Physical AI

Physical AI requires a progressive journey, not a **leap into the unknown**.

Organizations can start with focused, measurable, and high-potential initiatives, gradually building the capabilities, infrastructure, and governance models needed for long-term success. An operational roadmap can be structured around five key steps.

1 | Identify high-potential processes

The first step is to identify the processes where Physical AI can create tangible business value.
These are often processes that are repetitive, complex, risk-sensitive, data-intensive, or characterized by operational variability.
The goal is not to select the most impressive use case, but the one that is most relevant to the organization—where improvements in the ability to perceive, predict, or act can generate measurable impact.

2 | Map data, assets, and physical systems

Physical AI depends on the quality of its understanding of the operational environment.
Organizations need access to physical data, sensors, OT systems, digital twins, process information, historical data, operational knowledge, and integration with existing platforms.
Many organizations already possess some of these elements, but often in a fragmented manner.
The initial assessment should determine which data is available, which systems need to be integrated, which information is missing, and which constraints must be taken into account.

DGS
EMBRACE INNOVATION

3 Assess the maturity of digital twins and simulation environments

Organizations that already have digital models, simulation platforms, or advanced monitoring capabilities can accelerate their journey toward Physical AI.
Where these capabilities do not yet exist, they should be developed progressively.
Simulation becomes a space for learning, validation, and risk reduction. It enables organizations to test behaviors before deployment, evaluate scenarios, and prepare for the transition into real-world environments.

4 Establish governance, security, and validation criteria

Before scaling, organizations must define how intelligent systems will be tested, monitored, controlled, and secured.
This includes accountability models, policies, cybersecurity requirements, observability capabilities, reliability criteria, intervention mechanisms, and auditability.
Governance must accompany the initiative from the very beginning. It cannot be added later, once the system is ready for production deployment.

5 Launch pilot projects and co-creation initiatives

Experimentation should be practical, measurable, and value-driven.
Innovation labs, test environments, co-creation initiatives, and controlled pilot programs enable organizations to validate solutions, gather evidence, refine assumptions, and progressively build readiness.
The approach should remain pragmatic: start with clearly defined use cases, establish measurable success criteria, evaluate outcomes, and define the conditions required for scalability.

> This roadmap serves a dual purpose. On the one hand, it enables organizations to generate value today. On the other, it helps prepare them for the future evolution of Physical AI, including the more advanced developments associated with World Models.

EMBRACE INNOVATION

A look ahead to World Models: the evolutionary horizon of Physical AI

World Models represent one of the most promising directions in the evolution of Physical AI.

> If Large Language Models have made linguistic intelligence scalable, World Models have the potential to make operational intelligence scalable.

the ability to build dynamic representations of the world, simulate future scenarios, evaluate alternatives, and anticipate the consequences of actions.

Over time, this evolution could make Physical AI systems more adaptive, better able to generalize beyond training data, and more effective in managing complex environments.

From Language Models to World Models

TODAY – LLMs	TOMORROW – WORLD MODELS
Understand language	Understand context
Generate content	Simulate scenarios
Process knowledge	Evaluate alternatives
Support decision-making	Predict consequences

EMBRACE INNOVATION

A World Model seeks to build an internal representation of its environment, continuously update it based on incoming data, and use it to predict what may happen in response to a given action.

This capability is particularly relevant for Physical AI because systems operating in the real world must deal with variability, uncertainty, and consequences.

> The real evolution lies not simply in recognizing the world but in understanding its underlying dynamics.

Physical AI is the foundation on which organizations can begin building today. World Models represent one of the developments that may significantly expand its potential in the years ahead.

WHAT MAKES WORLD MODELS DIFFERENT?

1. **Dynamic representations of the environment**
2. **Continuous updates through data**
3. **Simulation of future scenarios**
4. **Prediction of consequences**
5. **Management of variability and uncertainty**

EMBRACE INNOVATION

The next generation of AI will be situated, governed, and operational

The next phase of artificial intelligence will not be defined solely by its ability to generate content, automate cognitive tasks, or produce insights.

It will be defined by its ability to understand the real world, interact with physical environments, support operational decisions, and act reliably within complex processes.

> Physical AI represents this transition: the evolution from AI that interprets information to AI that becomes part of operations.

For organizations, the real challenge is to start building today the foundations required to bring AI into physical processes: data, systems, simulation capabilities, skills, governance frameworks, cybersecurity measures, and control mechanisms.

UNDERSTAND

Understand the context

GOVERN

Ensure trust and safety

ACT

Generate value in operations

This is not about chasing robotics as the latest technology trend, nor about waiting for the most advanced technologies to reach full maturity.

Physical AI creates value when it is embedded within business processes, validated in realistic environments, protected through appropriate security measures, and governed throughout its entire lifecycle.

World Models point to a possible future evolution of this trajectory. Achieving enterprise readiness for Physical AI, however, is already a priority today.

The transition will be organizational, operational, and trust-driven—not merely technological.

Because bringing AI into the real world means building systems capable not only of acting, but of acting in ways that are **understandable**, **secure**, and **accountable**.

Physical AI is not the distant future of artificial intelligence. It marks the beginning of a new phase—one in which AI becomes **situated**, **governed**, and **operational**.

DGS
EMBRACE INNOVATION

CONTACT US

@dgs_spa

DGS SpA

www.dgsspa.com