Digital Twin Consortium Publishes Industrial AI Agent Manifesto, Led by XMPro

Industry-first governance framework defines ten laws for deploying trustworthy AI agents in safety-critical industrial operations

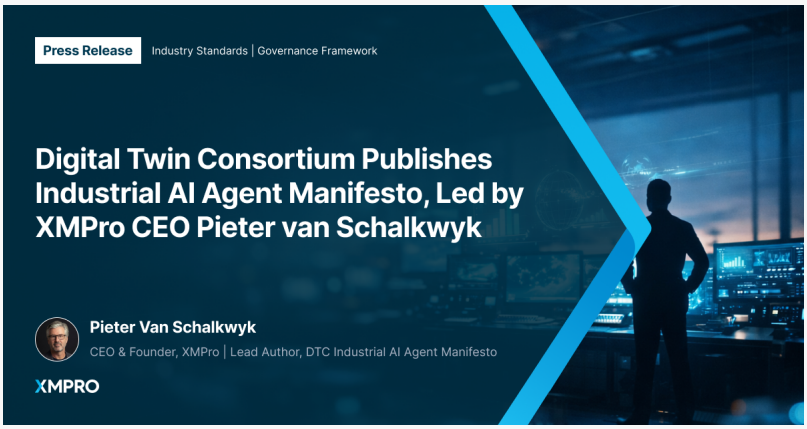

Digital Twin Consortium Publishes Industrial AI Agent Manifesto, Led by XMPro CEO Pieter van Schalkwyk

DALLAS, TX, UNITED STATES, March 2, 2026 /EINPresswire.com/ -- The Digital Twin Consortium (DTC) has published the Industrial AI Agent Manifesto — the industry's first comprehensive governance framework for deploying trustworthy AI agents in safety-critical industrial environments. XMPro CEO and founder Pieter van Schalkwyk served as lead author and chair of the DTC Composability Framework Working Group that developed the framework.

Now available on the Digital Twin Consortium's website, the manifesto establishes ten governance laws that define what industrial AI agents must do to operate safely and securely in production environments spanning healthcare, manufacturing, energy, mining, aerospace, and building operations. "Autonomous AI agents are no longer optional in industrial operations — but deploying them without governance infrastructure is a risk no responsible operator should accept," said van Schalkwyk.

> Autonomous AI agents are no longer optional in industrial operations — but deploying them without governance infrastructure is a risk no responsible operator should accept."
>
> *Pieter Van Schalkwyk - XMPro CEO*

"This framework defines the engineering requirements that make industrial autonomy durable: safety that is structurally guaranteed, not probabilistically predicted."

The Governance Gap

The manifesto responds to an urgent industry challenge. As AI agents move into safety-critical

operations — influencing patient care, process control, structural integrity, and grid stability — most organizations cannot demonstrate what their agents decided, why they decided it, or whether unsafe outcomes were structurally prevented.

Research firm Gartner has predicted that 40 percent of agentic AI deployments will fail by 2028 — not because the underlying AI lacks capability, but because the governance infrastructure required to operate these systems safely at scale does not yet exist.

At the same time, experienced industrial operators are retiring faster than replacements can be trained. The manifesto identifies properly governed AI agents as one of the few tools capable of preserving institutional knowledge and supporting thinner workforces in an era of increasing operational complexity.

The Ten Laws of Trustworthy Autonomous Operations

The manifesto defines ten governance laws derived from decades of safety-critical systems development and formalized through real-world experience with autonomous AI agents in industrial environments:

Deterministic Validation and Execution — Predictable, reproducible behavior in safety-critical decisions

Physics-Aware and Process-Aware Intelligence — Agents must respect physical constraints and encode process models, not just statistical patterns

Symbolic Primacy with Sub-Symbolic Intelligence — Symbolic reasoning architecturally bounds AI behavior, creating inherent trustworthiness

Separation of Control with Standardized Interoperability — Agent cognition architecturally separated from action execution

Emergency Stop, Human Override, and Graceful Degradation — Non-negotiable capabilities for safe operation

Interoperability with Operational Systems — Integration mediated through semantic models, not direct protocol interfaces

Auditability and Transparency — Complete decision provenance for regulators, operators, and stakeholders

Progressive Autonomy with Safety Boundaries — Graduated autonomy levels mapped to human roles, approvals, and safety criticality

Multi-Agent Safety Orchestration — Coordination of specialized capabilities with clear safety hierarchies

Safe and Secure Continuous Learning — Learned improvements deployed only through controlled processes maintaining safety guarantees

A detailed companion technical brief including conformance criteria for vendor evaluation, a deployment maturity assessment framework, and domain-specific implementation guidance is available on the DTC GitHub repository.

From Bounded to Bonded Autonomy

"Organizations that adopt governance frameworks like the Industrial AI Agent Manifesto will move faster toward autonomous operations, not slower," said van Schalkwyk. "When you can demonstrate to regulators, boards, and operators that your AI agents are constrained by architecture — not just training data — you earn the trust that unlocks real operational autonomy."

The manifesto introduces the concept of a progression from bounded autonomy — where agents operate within architecturally constrained safety boundaries — to bonded autonomy, where sustained compliance under independent verification enables certified, auditable, and insured autonomous operations.

Digital twins play a central role in the framework, providing proven infrastructure for state capture and replay, domain constraint enforcement, validation architecture, audit trails, and fleet-scale policy management — capabilities already deployed in production across DTC member organizations.

Industry Collaboration and Next Steps

The manifesto was developed by the DTC Composability Framework Working Group, chaired by van Schalkwyk, with contributions from Sean Whiteley of Axomem and editorial oversight from Dan Isaacs, CTO of the Digital Twin Consortium, and Will Thompson of the DTC.

DTC members including NIST, MITRE, TUV SUD, and leading academic and research institutions are already collaborating on the framework's next phases, which include domain-specific implementation guidance, standards development in partnership with ISO/IEC JTC 1/SC 41 and IEEE, verification frameworks for compliance testing, and reference architectures for production implementation.

Organizations interested in exploring collaboration opportunities can contact the DTC at agent_governance@engage.digitaltwinconsortium.org.

About XMPro

XMPro is pioneering agentic operations for industrial enterprises. The XMPro platform enables organizations to progress from real-time visibility through AI-assisted decision-making to autonomous operations — coordinating multiple AI agents to detect, decide, coordinate, and execute across complex industrial environments. With production deployments at elite global operators across mining, oil and gas, and process industries, XMPro delivers progressive decision intelligence on a single, governed platform spanning the full autonomy continuum. For more information, visit www.xmpro.com.

About the Digital Twin Consortium

The Digital Twin Consortium, a community of the Enterprise Data Management Association (EDMA), enables organizations to move from digital twin concepts to real-world practice. With over 200 members across industry, academia, government, and technology providers, the DTC develops standard requirements, reference architectures, and implementation guidance for digital twin adoption. For more information, visit www.digitaltwinconsortium.org.

Wouter Beneke - Marketing Lead
XMPro
email us here
Visit us on social media:
LinkedIn
Facebook
YouTube
X

This press release can be viewed online at: https://www.einpresswire.com/article/896719115

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.