




Closing the AI Deployment Gap

A Manifesto for Physical AI, Built to be
Deployed, Trusted, and Sustained.

An Ecosystem Initiative Founded by Lynx



An Open Invitation

The next phase of AI will not be defined by what is possible in the lab, but by what can be deployed, trusted, and sustained at the edge in the real world. This is an open invitation to silicon, platform, security, framework, integration, and operational leaders to help define the next phase of AI at the edge.

The Challenge

AI That Works in the Lab is Failing in the Field

For a decade, progress in AI has been defined by the lab: bigger models, better benchmarks, faster training. That era is now giving way to something more demanding, the move of AI out of IT and into OT, into the physical systems the real world runs on. This is the domain of Physical AI.

And here a hard truth emerges. Most AI, despite its promise, does not cross the boundary from development into edge deployment. It performs on benchmarks and in prototypes, then breaks down inside physical systems that must meet strict requirements for timing, safety, reliability, SWaP, security, and longevity.

This challenge defines the AI Deployment Gap: the widening disconnect between advances in AI capability and the ability to operationalize those capabilities in real-world, mission-critical environments. Models may succeed in development, but deployment demands something more: systems that can be trusted, verified, secured, and sustained over years of operation.



Four Gaps Must Be Closed

Four areas collectively create the deployment gap:

The Execution Gap

AI is probabilistic; mission systems must be deterministic. Execution must be bound, predictable, and verifiable.

The Integration Gap

AI must coexist with operating systems, middleware, legacy code, and heterogeneous, mixed-criticality hardware.

The Governance Gap

It is no longer enough to know what a system decides. We must know how, when, and under what conditions.

The Lifecycle Gap

Real systems live for years or decades. AI within them must be upgradeable, certifiable, and defensible over time.

The Opportunity

What Becomes Possible When We Close the Gap

The gap is real, but so is the prize. Closing it unlocks a whole class of systems the industry has struggled to field, and an opening for those who help define how it is done.

Closing the AI Deployment Gap Enables:

- **Trusted Autonomous Systems:** That act safely without a human in the loop.
- **Deployable Physical AI:** Moved out of the lab and into real machines.
- **Scalable Edge Intelligence:** Running under real-world compute and power constraints.
- **Certifiable AI-Enabled Platforms:** That meet the bar regulated industries' demand.
- **A Faster Path from Innovation to Operations:** Closing the distance between demo and deployment.

Why Lynx is Calling for Others to Join This Initiative


For decades, Lynx has helped organizations deploy software into systems where determinism, security, resilience, and lifecycle management are non-negotiable. Through MOSA.ic.AI® and our work across mission-critical environments, we believe the next challenge is enabling AI to meet those same standards, and that no single company can meet it alone.

Closing the AI Deployment Gap

We are learning to deploy artificial intelligence into the physical systems the real world depends on, by doing it, and by helping others do it. Through this work, we have come to value:

- **Deployment** over demonstration
- **Systems** over models
- **Predictable execution** over peak performance
- **Observable accountability** over opaque capability
- **Sustained Resilience** over momentary reliability

Each of these priorities reflects a simple reality: in operational environments, success is measured not by what a system can do under ideal conditions, but by how consistently, transparently, and reliably it performs when it matters most. At the edge, in the real world, there is no abstraction layer that can hide failure.



What We Stand For

Principles for Physical AI

If AI is to operate in the real world, it must be built on foundations that reflect the realities of that world. These are not enhancements. They are prerequisites, and the commitments a signatory makes.

- **Determinism by Design:** Execution must be predictable, even when outcomes are probabilistic. Latency is bound; timing is engineered, not hoped for.
- **Observability by Default:** System behavior is visible at every layer, at runtime, not reconstructed after the fact.
- **Integration Across the Stack:** AI operates seamlessly across heterogeneous silicon, operating systems, middleware, and legacy components, never in isolation.
- **Governance at Runtime:** Constraints and policies are enforced as the system operates, enabling traceability, control, and certification.
- **Security Without a Patch:** Systems defend themselves throughout deployment, resilient even when a fix is not yet available.
- **Engineered for the Lifecycle:** Built to be sustained, upgraded, and re-certified across the years, often decades, that real systems must endure.
- **Deployment as a First-Class Concern:** Deployment is designed in from the first decision, not bolted on at the end.
- **Ecosystem Effort:** No single vendor can close the gap. Silicon, platform, security, framework, integrator, and operator must align around a common goal.

Moving the Conversation Forward

The AI Deployment Gap Initiative creates a forum for organizations across the technology ecosystem to exchange perspectives on what it takes to deploy trustworthy AI in real-world, mission-critical systems. By bringing together expertise across silicon, platforms, security, software, integration, and operations, the initiative can help connect conversations that often take place independently across the industry.

What We Hope to Advance

- A stronger shared understanding of real-world AI deployment challenges
- Practical insights informed by experience across the technology stack
- Greater visibility into approaches being explored across mission-critical industries
- Continued exchange of ideas, use cases, and lessons learned
- Broader awareness of what is required to move AI from development into trusted operation

The Ecosystem Taking Shape

Founding Signatories

The AI Deployment Gap Initiative is launching with organizations that bring experience and perspectives from across software, security, engineering, and mission-critical systems. Each founding signatory who has endorsed the principles shares an interest in advancing deployable, accountable, secure, and sustainable AI. Their participation reflects a shared recognition that closing the deployment gap requires perspectives from across the technology ecosystem. Together, these organizations form the initiative's initial foundation and help bring greater visibility to the challenges surrounding real-world AI deployment.



Continue Exploring the Initiative

Visit the AI Deployment Gap Initiative landing page to learn more about the organizations supporting the effort, explore perspectives, resources, and discussions focused on advancing trusted AI deployment.

www.lynx.com/closing-the-ai-deployment-gap

© 2026 Copyright Lynx

The information herein is subject to change at any time after the date of publication. Lynx does not guarantee the accuracy of the information herein beyond the date of publication. All third-party company and product names mentioned, and marks and logos used, are trademarks and/or registered trademarks of their respective owners. Lynx trademarks are the property of Lynx.