

LYNX MOSA.ic.AI®

A Unified Execution Platform for Deterministic,
Certifiable AI in Mission-Critical Systems

LYNX MOSA.ic.AI is a unified CPU and GPU software platform that provides a deterministic, certifiable execution environment for deploying AI and advanced workloads in mission-critical edge systems.

As AI adoption accelerates across aerospace, defense, autonomy, and industrial systems, organizations face growing integration complexity, certification burden, and non-deterministic behavior. MOSA.ic.AI addresses this challenge by transforming AI from an experimental capability into a deployable, operational system function.

From AI Innovation to Operational Deployment


Modern AI development ecosystems are optimized for rapid prototyping, model training, and high-performance execution. They are not designed to meet the requirements of safety-critical systems. Rather than replacing existing AI ecosystems, MOSA.ic.AI provides the execution discipline needed to operationalize them and turn AI from a prototype capability into a deployable, certifiable system function.

- **From Training to Deployment:** Enables AI models trained in frameworks such as PyTorch and TensorFlow to execute deterministically in real-world systems.
- **Execution, Not Experimentation:** Focuses on runtime inferencing, where timing, scheduling and system behavior must be guaranteed.
- **Deterministic AI Behavior:** Ensures bounded execution time, fixed memory allocation, and predictable system response.



One Execution Platform

- **Unified Architecture:** Combines CPU mission computing and GPU acceleration into a single, cohesive platform.
- **Governed Execution Model:** Applies consistent isolation, scheduling and lifecycle control across compute domains.
- **Heterogeneous Compute Support:** Operates across ARM and Intel x86-64 architectures.



Two Configurations

- **MOSA.ic.AI Graphics Edition:** A unified, deterministic platform for mixed-criticality systems, enabling secure CPU/GPU coordination, visualization, and real-time processing within a certifiable architecture.
- **MOSA.ic.AI Compute Edition:** An advanced configuration that incorporates ComputeCore, enabling deterministic GPU-accelerated AI workloads aligned with mainstream neural network frameworks. This configuration enhances AI performance, supports hardware-accelerated inference, and provides a path toward certification for AI-enabled systems.

Reduced Integration Complexity

- **Eliminates CPU/GPU Silos:** One platform, not multiple integrated solutions.
- **Simplifies System Architecture:** Consistent execution and safety model across all workloads.
- **Supports Lifecycle Evolution:** Update components without redesigning the system.
- **Future-Proof Design:** Leverage open standards, mitigate obsolescence, avoid costly rework and recertification cycles.

Reduced Integration Complexity

- **Eliminates CPU/GPU Silos:** One platform, not multiple integrated solutions.
- **Simplifies System Architecture:** Consistent execution and safety model across all workloads.
- **Supports Lifecycle Evolution:** Update components without redesigning the system.
- **Future-Proof Design:** Leverage open standards, mitigate obsolescence, avoid costly rework and recertification cycles.

Deterministic Execution for AI and Real-Time Systems

- **AI Inference as a Deterministic System Function:** Controls scheduling, memory access, I/O timing, and partitioning to ensure predictable behavior.
- **Static Execution Model:** Removes runtime unpredictability through fixed configuration.
- **Bounded Performance Guarantees:** Ensures deterministic system behavior and worst-case executing timing.
- **Mixed-Criticality Support:** AI and safety-critical workloads coexist with controlled and bounded interference.

Assured Isolation and System Integrity

- **Separation Kernel Foundation:** Built on LynxSecure®, with hardware-enforced isolation.
- **Non-Bypassable Isolation:** Prevents cross-domain interference and protects system integrity.
- **Controlled Communication:** Secure data exchange between domains.
- **Concurrent OS Support:** LynxElement®, LynxOS-178, third-party RTOS, Linux, Windows, and bare-metal applications.

Primary Use Cases

- **Fault-Tolerant, Cyber-Resilient Platforms:** Resists cyber attacks and contains faults through hardware-enforced isolation and a trusted computing base.
- **Avionics Sensor Fusion:** Real-time perception and data fusion with guaranteed timing.
- **Autonomous Flight and Navigation:** Deterministic AI for UAV/UAM and mission systems.
- **Pilot Assist Systems:** AI-driven decision support within certified flight environments.
- **Defense ISR and Targeting Systems:** Predictable AI behavior under contested conditions.
- **Industrial Robotics and Automation:** Safe, AI-enabled automation with real-time control.
- **Edge AI Systems:** Deployment in disconnected, degraded or low-latency environments.

Why LYNX MOSA.ic.AI

The industry needs a way to safely operationalize AI. MOSA.ic.AI delivers:

- A deterministic execution platform for AI at the edge.
- A unified CPU + GPU architecture governed under a single model.
- A clear path to deploy AI safely in regulated environments.
- Reduced integration burden, certification risk, and lifecycle cost.
- A foundation for mixed-criticality, AI-driven mission systems.



© 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.

edge@lynx.com

www.lynx.com