LYNX

Lynx Certifiable Compute Portfolio on Intel® Architecture

Mixed-Criticality CPU and GPU Platforms

Executive Overview

Modern aerospace and defense systems demand a difficult balance: high performance compute, advanced graphics, AI acceleration, system consolidation, all while maintaining rigorous safety standards without increasing certification risks. Lynx delivers a complete Intel platform solution for safety- and mission-critical edge capable systems by offering fully integrated CPU, GPU and supporting full stack ecosystem solution:

LYNX MOSA.ic.AI®

- Solution platform enforcing deterministic execution of partitioned CPU & GPU workloads enabling constrained Graphics & AI use cases for the embedded edge
- Mixed integrity / heterogeneous system unification
- DO-178C DAL A certifiable multicore runtime encompassing: LynxSecure™, LynxOS-178®, LynxElement™, CoreSuite™ 2.0, and ComputeCore™

LynxSecure™

- Type 1 separation kernel hypervisor providing the foundation for system level isolation & security
- Fully virtualized platform
- Designed in a least privileged manner

LynxOS-178® & LynxElement™

- High-assurance guest RTOS operating environments

Certification & Resilience

- Adherent to mandated safety & security standards (FAA, EASA, ATC, etc.) and key objectives (CAST-32a and ARP-4754a)

Engineered for Reliability

- Long lifecycle, supply assurance, extended maintenance and engineering services

CoreSuite™ 2.0

- HMI / Graphics framework based on a Vulkan SC foundation layer enabling Khronos® conformant open standards. In addition to monitoring mechanisms for resilient edge deployment

ComputeCore™

- AI capable BLAS, FFT, & NN building blocks providing a path to deterministic model execution and autonomous system use cases (ML & CV)

Professional Services

- Baseline platform tailoring & application-level accommodations to meet unique integrated system deployment needs

Temperature Screening & Storage

- Extended temperature range
- Dynamic temperature range device resilience

Engineered for Reliability

- Long lifecycle, supply assurance, extended maintenance and engineering services

Intel & Lynx Enabled Technology

- GPU task context switching & time pre-emption
- Watch dog monitoring
- Thermal throttling control
- Ring buffer automation / self-propelling
- Resource director technology (RDT)
- Cache allocation technology (CAT)
- DisplayPort multi-stream (MST)
- AEP & FMEA safety assessment
- Multi-engine (Graphics / Compute | Codec | BLT)
- In-band ECC
- Time-sensitive networking (TSN)

By pairing Intel's industry-leading performance with Lynx's safety-critical software architecture, Lynx delivers a unified execution platform where CPU and GPU workloads, including AI inference, operate under a governed and bounded model. This approach transforms AI from a development artifact into a deployable, certifiable system function.

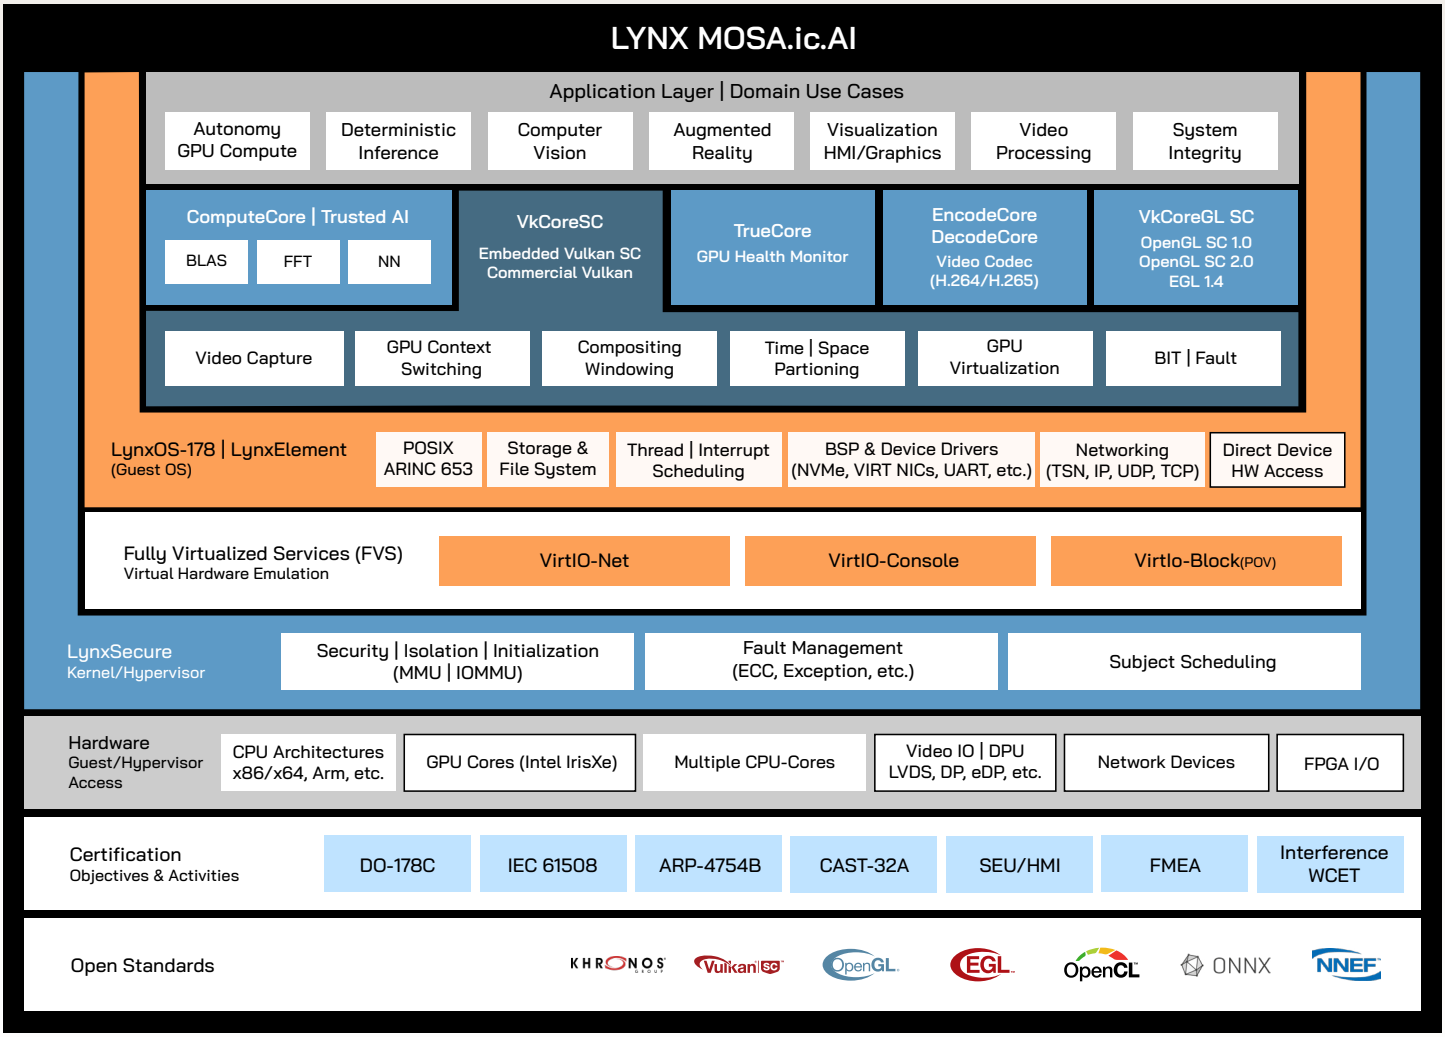

Why LYNX MOSA.ic.AI & Intel

- Foundation for mixed-Integrity multicore CPU & GPU system consolidation
- Secure partitioning of heterogeneous workloads & system unification
- Path to DO-178C DAL A deployment & standards compliance
- DMA engine interference mitigation (WCET) & hazard assessment
- Safety-critical graphics (OpenGL SC 1.0, OpenGL SC 2.0, Vulkan SC)
- Resilient AI inference as a governed runtime function
- Deterministic GPU compute (BLAS, FFT, & Neural Network Primitives)
- Controlled I/O and DMA behavior plus networking protocols
- Long lifecycle BSP and platform sustainment
- Unified CPU and GPU architecture under a governed model
- Reduced integration and lifecycle risk

What Lynx Enables on Intel

- Mixed-integrity system consolidation on multicore CPUs
- DO-178C DAL A deployments
- Multicore / multi-engine interference mitigation (WCET)
- Safety-critical graphics (OpenGL SC 1.0 / 2.0, Vulkan SC)
- Resilient AI inference as a governed runtime function
- Deterministic GPU compute (BLAS, FFT, and Neural Network Primitives)
- Controlled I/O and DMA behavior plus networking protocols
- Secure partitioning of heterogeneous workloads
- Long lifecycle BSP and platform sustainment

Why LYNX MOSA.ic.AI on Intel

- Deterministic execution for AI and heterogeneous workloads
- Unified CPU and GPU architecture under a governed model
- Clear path toward certification of AI-enabled systems
- Reduced integration and lifecycle risk
- Foundation for mixed-integrity mission systems

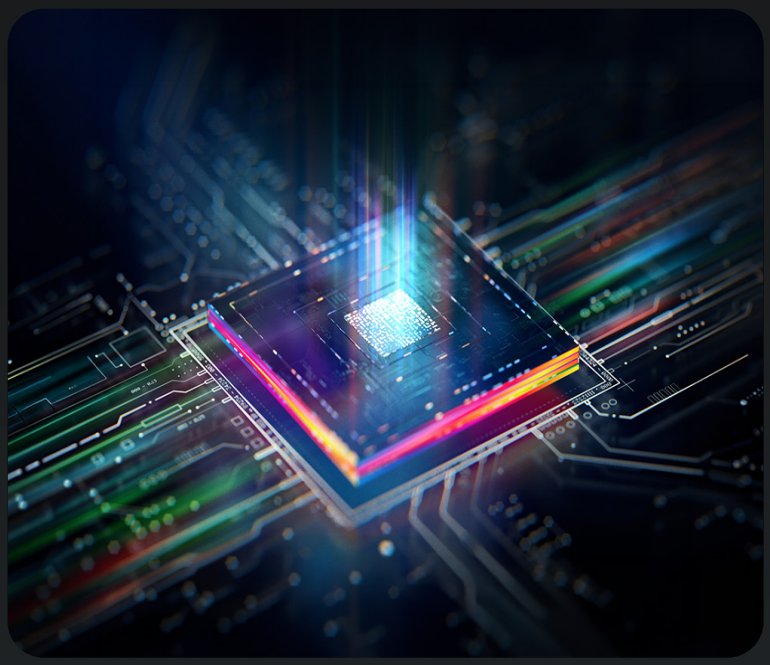

Unified CPU + GPU Execution Architecture

MOSA.ic.AI unifies CPU mission computing and GPU acceleration into a single governed platform.

- **Unified Architecture:** Eliminates CPU/GPU silos through a cohesive execution model
- **Governed Execution Model:** Applies consistent isolation, scheduling, and lifecycle control across compute domains
- **Heterogeneous Compute Support:** Optimized for Intel x86 architectures with integrated GPU acceleration
- **Mixed-Integrity Enablement:** AI, graphics, and safety workloads coexist with controlled and bounded interference

This architecture reduces integration complexity while enabling predictable system behavior.

Reduced Integration Complexity

MOSA.ic.AI reduces system complexity by applying a unified execution model across CPU and GPU domains.

- Eliminates CPU / GPU silos
- Provides a consisten safety execution framework
- Supports lifecycle evolution without full system redesign
- Reduces recertification scope through controlled paritioning

This approach lowers integration burden while improving long-term system maintainability.

From AI Development to Operational Deployment

Modern AI ecosystems are optimized for training and experimentation, not for safety-critical deployment. MOSA.ic.AI introduces the execution discipline required to operationalize AI on Intel platforms.

- **Training & Validation to Deployment:** Models developed in frameworks such as PyTorch and TensorFlow execute deterministically within partitioned runtime environments
- **Execution, Not Experimentation:** Inference operates under controlled scheduling, memory, and I/O constraints
- **Deterministic Behavior:** Bounded execution time, fixed resource allocation, and predictable system response

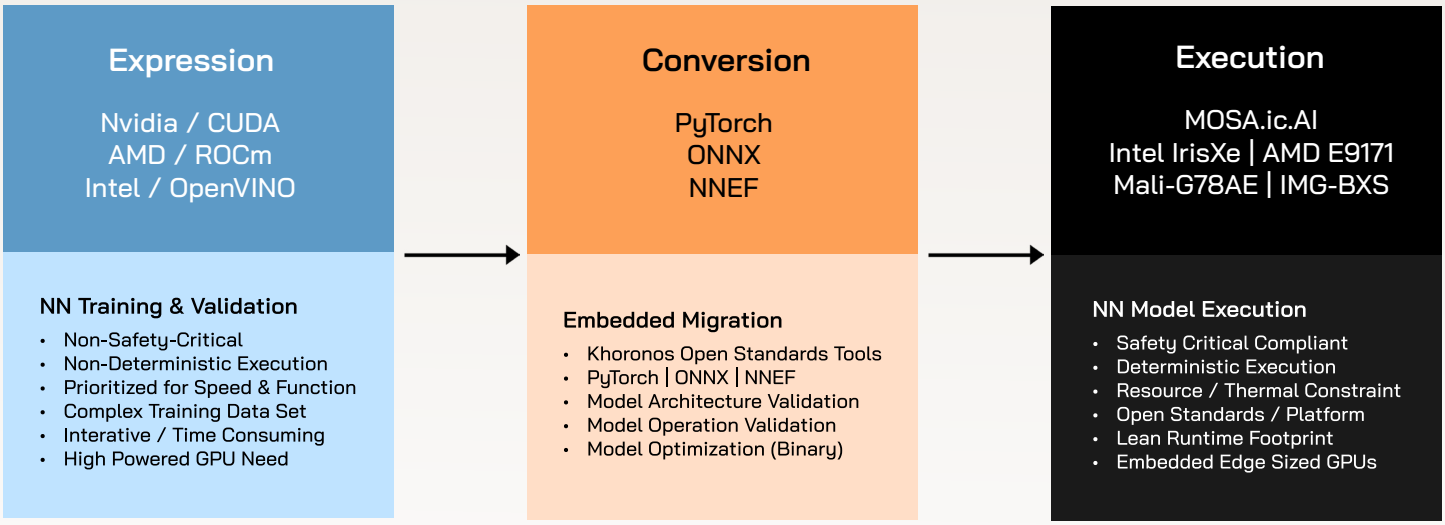

This model treats AI inference as a controlled system function aligned with certification expectations, rather than a dynamic software process.

LYNX MOSA.ic.AI Configurations

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. Customers will gain better optimization between CPU and GPU safety domains, reduced integration burden, and a cleaner architectural path for heterogeneous compute systems within a safety-oriented framework.

Graphics Edition

MOSA.ic.AI brings together the existing capabilities of MOSA.ic (CPU safety) and CoreSuite (GPU safety) into a unified system-level solution. Customers will gain better optimization between CPU and GPU safety domains, reduced integration burden, and a cleaner architectural path for heterogeneous compute systems that can support AI workloads within a safety-oriented framework.

Compute Edition

MOSA.ic.AI is available with a more advanced configuration incorporating ComputeCore, a CoreSuite extension specifically focused on supporting AI hardware accelerated workloads in safety-relevant systems. ComputeCore enables deterministic GPU accelerated AI capabilities aligned with mainstream neural network frameworks with a path to certification. This version enhances AI workload support, provides optimizations designed for accelerated compute, and strengthens the foundation for safety-conscious AI deployment.

ComputeCore: Deterministic AI Edition

ComputeCore provides a structured and deterministic execution environment for GPU-accelerated AI workloads. It bridges mainstream AI frameworks and safety-critical runtime environments by constraining how models execute on GPU hardware.

Key Capabilities

- Deterministic execution of neural network inference workloads
- AI capable BLAS, FFT, & NN building blocks for ML & CV system use cases
- Alignment with common AI frameworks while enforcing static runtime behavior
- Hardware-accelerated compute with bounded execution characteristics
- Integration with partitioned CPU/GPU scheduling under LynxSecure

ComputeCore enables AI workloads to operate as predictable system functions, supporting timing analysis and certification-aligned assurance objectives.

LynxSecure

LynxSecure is a separation kernel type 1 hypervisor that provides robust isolation for mixed integrity workloads.

Features

- Hardware-assisted virtualization (VT-x)
- DMA isolation (VT-d) and SR-IOV device partitioning
- Strong time and space separation
- Reduced Trusted Computing Base (TCB)
- VirtIO support

Supported Guest Operating Environments

- LynxOS-178 (DO-178C certifiable RTOS)
- LynxElement (unikernel-based runtime)
- Linux (mission / non-safety partitions)
- Third-party RTOS (e.g., Zephyr®, commercial RTOS)

Within MOSA.ic.AI, LynxSecure enforces non-bypassable isolation and controlled communication between CPU and GPU domains, preserving system integrity under mixed-integrity workloads.

Making Multicore Intel Practical for Certifiation

Intel multicore processors provide substantial performance advantages, but shared resources introduce interference and nondeterministic behavior that complicate certification. Achieving certification requires more than hardware virtualization. It requires a controlled execution model governing compute, memory, and I/O behavior across workloads.

Lynx addresses this by combining Intel hardware capabilities with a separation kernel architecture and a deterministic execution model aligned with MOSA.ic.AI. This approach enables multicore performance to be used within predictable and bounded conditions.

Hardware-Assisted Isolation

- VT-x for virtualization and privileged control
- VT-d for DMA remapping and device isolation
- SR-IOV for high-performance device partitioning

These capabilities provide the foundation for isolating workloads and managing shared resources.

Deterministic Execution Controls

- Time, and space partitioning across CPU and GPU domains
- Controlled scheduling of safety, mission, and AI inference workloads
- Isolation of shared resources to limit cross-domain interference
- Fixed system configuration to reduce runtime variability function

LYNX

Benefits

- Clear separation between safety-critical and mission workloads
- Controlled and auditable resource access
- Reduced reliance on software-only mitigation techniques
- Smaller and more defensible certification boundary
- Support for technology insertion and software reuse

Cache and Shared Resource Control

LynxSecure utilizes Intel Resource Director Technology (RDT) and Cache Allocation Technology (CAT) to enforce deterministic use of shared resources.

- Allocate defined LLC regions to workloads
- Reduce cache contention
- Improve execution predictability
- Support worst-case execution time (WCET) analysis

These mechanisms are essential for maintaining bounded interference in multicore systems running mixed-integrity and AI workloads.

I/O Determinism and Device Management

Deterministic execution extends beyond compute into I/O and device behavior.

Intel virtualization capabilities allow Lynx to:

- Preserve isolation during high-performance data movement
- Safely share devices across partitions
- Control DMA behavior
- Manage interrupt routing predictably

This ensures that data movement and external interactions remain consistent with system-level timing and certification requirements.

CoreSuite 2.0 on Intel GPUs

CoreSuite extends deterministic execution into the GPU domain, enabling graphics and compute workloads to operate under the same governed model as CPU workloads.

Capabilities

- Safety-critical Khronos conformant standards: OpenGL SC 1.0, OpenGL SC 2.0, & Vulkan SC
- Video encode/decode (H.264/H.265)
- Runtime integrity monitoring and fault reporting
- GPU task context switching & time pre-emption
- Resource virtualization and mixed integrity
- GPU compute primitives supporting AI and signal processing

RTOS Ecosystem Support

CoreSuite 2.0's OS abstraction layer enables integration with:

- Linux
- VxWorks® HVP
- Deos™
- Integrity-178®

This allows GPU acceleration to be incorporated into existing system architectures without compromising certification objectives.

Inventory Assurance Services and Long-Term Supply

Lynx complements software determinism with hardware lifecycle assurance.

For select platforms, including Intel Tiger Lake, Lynx provides:

- DTR-screened SoCs
- Production yield optimization support
- Inventory management and secure storage
- Long-term availability programs

These services reduce supply chain risk and support sustained deployment.

Intel Hardware Ecosystem Support

Lynx supports a broad range of Intel platforms, including:

- Intel Tiger Lake (active certification)
- Intel Raptor Lake
- Skylake, Apollo Lake, Bay Trail

Deployments span leading aerospace and defense hardware providers, including Curtiss-Wright, Mercury Systems, Congatec, and Supermicro.

Lynx works closely with these partners to enable deterministic operation, develop BSPs, and support certification-aligned integration.

Platform Enablement and BSP Engineering Services

Accelerating Intel-Based Deployments

Lynx provides comprehensive platform enablement services to accelerate adoption of Intel architectures in aerospace and defense systems.

Safety-Aligned BSP Strategies

For programs requiring certification, Lynx supports safety-aligned BSP approaches consistent with DO-178 objectives. Lynx collaborates closely with customers, silicon vendors and hardware partners to ensure platform software:

- Aligns with certification plans and controls certification scope
- Supports deterministic system behavior and minimizes integration risk

These deployments demonstrate architectural maturity, repeatability, and reduced risk for avionics programs.

Engage with Lynx

Lynx enables confident adoption of Intel multicore CPU and GPU technologies for safety- and mission-critical systems.

To learn how Lynx can accelerate your Intel platform deployment, contact us.

edge@lynx.com
www.lynx.com