OCUDU ECOSYSTEM FOUNDATION

WNDRVR

Resilient Edge Operations — AI-RAN & Live Workload Continuity

Wind River Systems | OCUDU Ecosystem Foundation | MWC 2026

OCUDU ECOSYSTEM FOUNDATION

Executive Summary

Wind River and the OCUDU Ecosystem Foundation's OCUDU Technical Project community jointly demonstrated resilient, AI-native edge operations — delivering deterministic low-latency RAN processing, automated fault detection, and live workload migration without service interruption. The result is a validated, open-source foundation for mission-critical distributed networks that reduces anomaly detection from hours to minutes, cuts TCO by up to 30%, and enables sovereign, carrier-grade AI-RAN deployments at scale.

At a Glance

Partners	**Wind River Systems · OCUDU Ecosystem Foundation**
Platform	Wind River Cloud Platform (StarlingX-based)
Use Case	AI-RAN & Open RAN — Resilient Edge / Mission-Critical
Demonstrated At	MWC 2026
TCO Reduction	Up to ~30% vs. legacy RAN
Anomaly Detection	Hours > Minutes
Data Scale	70+ TB of network data per week
Reliability	Carrier-grade five-nines class availability

The Challenge: Keeping Mission-Critical Networks Running at the Edge

Distributed networks and mobile platforms operate in some of the most demanding conditions in modern computing. Low-latency, deterministic processing and continuous availability must be maintained across sovereign, resource-constrained edge environments — often with intermittent connectivity, heterogeneous hardware, and no tolerance for downtime.

Before this collaboration, operators faced four compounding challenges:

- Ensuring low-latency, local inference and RAN control to meet real-time requirements while minimising dependence on fragile backhaul links.
- Providing deterministic platform behaviour and lifecycle management so critical workloads remain available during upgrades, failures, or mobility events.
- Automating fault detection and remediation — reducing detection and recovery time — to support resilient, autonomous operations.
- Delivering secure, supply-chain-aware, and sovereign deployment options across multi-vendor ecosystems without sacrificing performance.

What Was Holding the Industry Back

Four structural barriers blocked progress before this work began:

- Fragmented multi-vendor stacks with limited edge compute and power resources, and inconsistent standards that made interoperable Open RAN/ AI-RAN solutions hard to replicate at scale.
- No reliable low-latency local inference or deterministic platform behaviour at the edge, forcing real-time RAN control to depend on fragile backhaul connections.
- Immature operational tooling and lifecycle management for distributed sites, meaning upgrades, migrations, and fault remediation risked service disruption and required extensive manual intervention.
- Unresolved security, governance, and sovereignty concerns — combined with costly, unproven ROI for AI at the edge — prevented broad adoption in mission-critical environments.

The Solution: A Unified, Deterministic Edge AI Platform

Wind River delivered an integrated software stack and operational framework that brings deterministic real-time compute, secure cloud-native edge infrastructure, Edge AI/AI-RAN capabilities, and workload lifecycle tooling into a validated, interoperable package for distributed and mobile deployments.

Edge AI / AI-RAN	Cloud-Native Edge Platform	Lifecycle & Mobility
Local inference runtimes + AI-RAN xApp/rApp integrations — closed-loop RAN optimisation and real-time decisioning at the network edge.	Wind River Cloud Platform (StarlingX-based) — a deterministic, scalable, and security-hardened Kubernetes compute infrastructure optimized for distributed edge environments, with carrier-grade high availability for VNFs, CNFs, and AI workloads.	Orchestrated live VM/container migration + zero-touch provisioning — workload continuity during upgrades, failures, and mobility events.

Together, these elements solved the low-latency and availability constraints, automated fault detection and remediation, enabled continuous mission services during change events, and provided a validated, secure path for deploying AI-enabled Open RAN and private 5G at the edge.

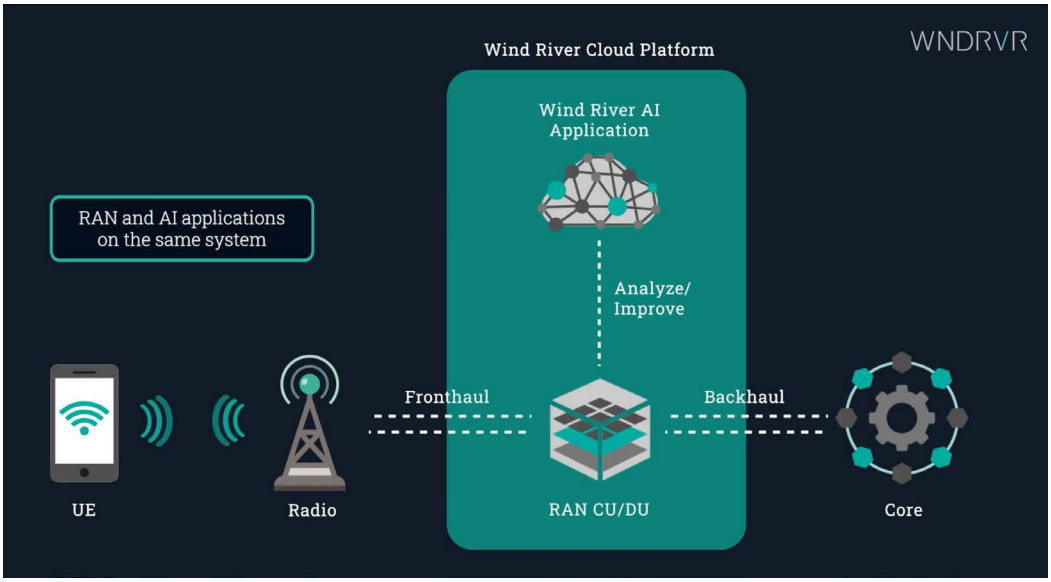

OCUDU
ECOSYSTEM FOUNDATION

OCUDU's Critical Contribution

The OSS project provided the open, carrier-grade reference RAN stack and ecosystem framework that underpinned the entire demonstration — enabling Wind River to integrate AI-RAN and Open RAN components on a validated CU/DU baseline. Specifically, OCUDU:

- Provided the open RAN reference architecture that accelerated interoperability testing with partner hardware and runtimes.
- Enabled Wind River to validate AI-RAN and vRAN workloads running deterministically on the same validated platform, rather than on separate, fragile stacks.
- Served as the carrier-grade baseline that demonstrated real workload portability and multi-vendor interoperability in a live production environment at MWC 2026.

The Role of Open Source: From Community to Carrier-Grade

Open source provided the reference infrastructure and components that enabled rapid integration, reproducibility, and multi-vendor interoperability throughout the demonstration:

- StarlingX / Kubernetes formed the cloud-native edge foundation, enabling scalable, community-driven infrastructure that both Wind River and the OCUDU community could build upon and extend.
- Yocto / Debian provided the hardened, supply-chain-verified OS layer — combining open-source reproducibility with carrier-grade security and compliance.
- The OCUDU open source project provided the AI-RAN network function (an open RAN CU and DU)

Wind River built on these community-driven components while adding hardened, supply-chain-verified artifacts and lifecycle tooling to meet carrier-grade and sovereignty requirements, creating an approach where open standards and production-grade reliability coexist.

Results and ROI: What the Numbers Say

The MWC 2026 demonstration validated the solution under real-world conditions, producing measurable outcomes across performance, economics, and operations:

Metric	Result	What It Means
Anomaly Detection Speed	Hours > Minutes	AI-RAN telemetry correlation slashes fault detection time
TCO Reduction	Up to ~30%	Consolidating Open RAN + AI-RAN workloads on unified platform vs. legacy RAN
Network Data Scale	70+ TB/week ingested	Enables long-term behavioural learning and predictive optimisation
Reliability	Five-nines class (99.999%)	Carrier-grade HA via live upgrades and deterministic lifecycle tooling
Deployment Speed	Weeks > Minutes	Automated lifecycle management scales to thousands of distributed sites
Latency	Real-time edge inference	Closed-loop RAN tuning improves responsiveness vs. centralised processing
Vendor Lock-in	Eliminated	Open RAN + open-source foundations enable multi-vendor hardware freedom
Workload Continuity	Zero service interruption	Live VM/container migration during maintenance and mobility events

Beyond the metrics, the broader business impact is transformational: operators can reduce mean-time-to-repair, unlock new revenue streams through monetisable edge services (including V2X), and transition from reactive network management to autonomous, AI-driven operations.

What Makes This Approach Different

Traditional edge deployments rely on separate, fragile stacks for RAN and AI workloads — requiring manual tuning, site-specific integration, and accepting service disruptions during change events. This collaboration breaks that model:

- **Unified deterministic platform:** AI-RAN and vRAN workloads run on the same validated, hardened infrastructure — eliminating the operational complexity of managing separate stacks.
- **Live workload migration:** Deterministic VM and container migration preserves service continuity during upgrades, maintenance, and mobility events — a capability that alternative approaches cannot match without service interruption.
- **Automated fault detection:** AI-RAN telemetry correlation reduces anomaly detection from hours to minutes, enabling autonomous remediation that dramatically lowers MTTR.
- **Sovereignty by design:** Supply-chain-aware builds, hardened images, and role-based management ensure that data and operational control remain with the operator — critical for government and mission-critical deployments.
- **Vendor agnostic** CU/DU software implementation from the OCUDU open source project that can be easily adapted to running on any Edge platform.

OCUDU
ECOSYSTEM FOUNDATION

Next Steps: From Demonstration to Production at Scale

The MWC 2026 demonstration marks a beginning, not an endpoint. The next phase of this collaboration focuses on three priorities:

- Validating demo artifacts in production environments — operators are invited to participate in an instrumented OCUDU trial with Wind River to measure real-world KPIs including MTTR, detection time, service continuity during migration, and TCO.
- Scaling AI-RAN and edge AI workloads from pilot demonstrations to full commercial deployments across thousands of distributed sites, leveraging automated lifecycle management and validated reference stacks.
- Deepening ecosystem contribution, with Wind River and the OCUDU community continuing to contribute upstream to open source projects and standards bodies — ensuring the blueprint remains open, reproducible, and extensible for all operators.

GET INVOLVED

Join a short, instrumented OCUDU trial with Wind River to validate the demo artifacts in your environment and measure KPIs including MTTR, detection time, service continuity during migration, and TCO. Explore the technical blueprints and validation reports from the Wind River and OCUDU community. Contact Wind River telco experts to schedule a customised deep-dive session — and join the movement toward resilient, autonomous, and interoperable distributed networking for 5G and 6G.

OCUDU EF's Getting Started Guide | **ocudu.org** | **windriver.com/industries/telecommunications**