

Enhancing Defense Capabilities with EdgeRunner AI

Learn how domain-specific AI adapters optimize decision-making for defense and intelligence operators while improving data sovereignty and operational resilience.

Solution Benefits

- **Data privacy and security.**
Keep AI on-device, with technology that the military trusts.
- **AI that works offline.**
Use AI anywhere, anytime, independent of internet connectivity.
- **Personalized AI.**
Custom AI adapters are tailored to your role and mission.
- **Scalability and flexibility.**
Deploy AI across any Intel® architecture-powered device ecosystem with ease.

Executive summary


In the rapidly evolving landscape of defense operations, AI is crucial for making swift, informed decisions in high-stakes environments—on the battlefield, during training maneuvers, and at headquarters. However, traditional cloud-based AI solutions pose significant challenges: they require constant internet connectivity, introduce data privacy risks, incur unpredictable per-use costs, and lack domain-specific accuracy. Intel and EdgeRunner AI joined forces to create an air-gapped, private AI solution that solves these challenges. The solution doesn't require internet connectivity; instead, it delivers on-device AI agents, customized for specific Military Occupational Specialties (MOS), to AI PCs powered by Intel® Core™ Ultra processors. Now, military personnel can use AI without needing Wi-Fi or the internet, with confidence that data remains highly secure.

Challenge: Cloud-based AI often poses difficulties for defense organizations

AI is changing the way defense organizations operate. Approximately 62% of surveyed defense organizations worldwide have already integrated AI into their operations.¹ The benefits of AI in defense are clear, as illustrated by these examples²:

- The use of AI in military logistics has increased efficiency by approximately 35%.²
- AI-powered surveillance systems have reduced ground patrol costs by 40%.²
- AI-assisted target identification accuracy in military applications exceeds 92%.²

Traditionally, AI requires connectivity—however warfighters and field personnel often operate without reliable access to the cloud. Even when connectivity is available, concerns about data privacy and spiraling costs are always present. What's more, standard AI solutions, even those that use advanced retrieval-augmented generation (RAG) and large language models (LLMs), have limited knowledge of and customization for military missions and projects.



AI can accelerate tasks like logistics planning, operations reporting, and readiness assessments. However, defense organizations need secure, mission-specific AI that runs entirely on-device—helping to ensure data sovereignty, operational resilience, and cost predictability.

Solution value: Customized, air-gapped AI

EdgeRunner AI has partnered with Intel to create a Generative AI (GenAI) solution that helps defense organizations to solve AI challenges. The solution is:

- **Personalized.** The solution provides customized, role-specific AI assistants, accelerating accurate decision making. These personalized adapters are trained for different Military Occupational Specialties (MOS), with MOS-specific data and information. Users can simply click on their MOS from a drop-down menu for instant access to customized data, LLMs, and agents.
- **Agentic.** AI agents can take actions and complete tasks on behalf of the user, streamlining workflows and increasing productivity. Capabilities include summarization, question-and-answer interactions, translation, transcription, code generation, image recognition, text-to-speech, and speech-to-text. EdgeRunner AI can quickly build custom adapters for specific customers as requested.
- **Air-gapped.** The solution runs without requiring access to the cloud or internet, enhancing data privacy and security. However, the solution can connect to the internet for real-time information or collaboration applications, if desired. EdgeRunner releases an updated version of the application on a quarterly basis, which includes the latest LLMs and performance improvements.
- **Performant and efficient.** Because the agents, RAG, and LLMs run locally, latency is near zero.³ EdgeRunner uses the Intel Distribution of OpenVINO™ toolkit to optimize popular open-source LLMs, such as Llama or Mistral, to run efficiently on the Intel Core Ultra processor's eXtensible processing unit (XPU).
- **Cost-effective.** Customers incur zero inferencing costs because inference is completely on-device, eliminating the need for cloud hosting costs. This helps keep costs under control and increases pricing predictability.
- **Easy to use.** Customers can install the solution in as little as 15 minutes⁴, with no machine learning expertise. The user interface is intuitive.



AI PCs powered by Intel

An AI PC is a client/edge device that uses AI technologies to elevate productivity, security, and more. AI PCs can take many form factors, including laptops, companion compute edge devices, embedded form factors, and even desktops and workstations.

AI PCs powered by Intel leverage XPU architecture. This is a comprehensive computing platform that integrates multiple types of processing units to handle a wide variety of workloads. This architecture includes CPUs, GPUs, NPUs, and other specialized accelerators. The idea is to provide a flexible and efficient computing environment that can dynamically allocate tasks to the most suitable processing unit. This helps optimize performance, power efficiency, and cost for diverse applications, ranging from AI to general-purpose computing.

- **CPU:** a high-performance Intel® central processor for fast response and low-latency, lightweight AI tasks
- **GPU:** a built-in Intel® graphics processing unit for high throughput and performance parallelism
- **NPU:** a built-in, power-efficient Intel® neural processing unit for sustained AI workloads and AI offload





Figure 1. Solution architecture for EdgeRunner AI.

Figure 1 illustrates the EdgeRunner AI solution stack. Note that the OpenVINO optimization is done by EdgeRunner AI prior to customer download.

Example use cases

EdgeRunner AI, powered by Intel® technology, can enhance real-time analysis, decision-making, operational efficiency, and security, empowering defense personnel to operate more effectively in both combat and intelligence-gathering environments. Here are three common use cases:

- **Logistics planning in contested environments.** When you need information instantly to keep troops safe, the last thing you want to see is “No Internet” on your screen. The EdgeRunner AI solution runs effectively in denied, degraded, intermittent, and limited bandwidth (DDIL) environments—giving you the data you need when you need it.
- **Maintenance and operations planning at the tactical edge.** Defense organizations deploy a vast array of equipment, such as planes, drones, trucks, tanks, autonomous land vehicles, and more. EdgeRunner AI’s MOS-specific adapters provide assurance and personalized responses to military-specific questions.
- **Combat medicine in the field.** As medics work to save lives, keeping their location secret is crucial. Because the EdgeRunner AI solution does not send data externally, it increases operational security by limiting transmission of signals.



Conclusion

Unique in the AI industry, EdgeRunner AI provides a domain-specific, state-of-the-art LLM-based AI assistant that runs completely locally on-device. LLM optimization for Intel Core Ultra processors delivers powerful and performant AI, while the air-gapped solution helps increase protection for sensitive military data. With EdgeRunner AI, defense organizations can operate at the speed of intelligence—securely, locally, and without compromise.

Visit [EdgeRunnerAI.com](https://edgerunnerai.com) to learn more and access a free trial.

Learn more

- [EdgeRunner AI](#)
- [Intel® Core™ Ultra processors](#)
- [Intel Distribution of OpenVINO™ toolkit](#)



¹ Gitnux, April 2025, "[AI In The Military Industry Statistics](#)."

² Gitnux, April 2025, "[AI In The Military Industry Statistics](#)."

³ Source: [EdgeRunner website](#), accessed June 25, 2025.

⁴ Source: Based on EdgeRunner AI's customer interactions.

Performance varies by use, configuration, and other factors. Learn more at intel.com/PerformanceIndex.

Performance results are based on testing by Intel and may not reflect all publicly available security updates. See configuration disclosures for details. No product or component can be absolutely secure.

AI features may require software purchase, subscription or enablement by a software or platform provider, or may have specific configuration or compatibility requirements. Data latency, cost, and privacy advantages refer to non-cloud-based AI apps. Learn more at intel.com/AIPC

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

Other names and brands may be claimed as the property of others. © Intel Corporation

0825/CP/CAT/PDF 366190-001EN