Solution Brief

Healthcare

# Building the Next Generation of AI-Ready Patient Monitoring

## For medical OEM architects and product managers: Kontron embedded platforms on Intel® Core™ processor and Intel® Core™ Ultra processor, characterized with Intel® Health & Life Sciences (HLS) AI Suite

Patient monitoring is moving from single-purpose display-and-alarm devices to **software-defined edge systems** that fuse real-time waveforms, video, wearable inputs, alarm context, and operational telemetry. Increasingly, that analysis runs as concurrent AI pipelines close to the patient rather than in the cloud.

For OEM architects and product managers, the hard part is making platform decisions today that won't limit the roadmap tomorrow. Kontron, an Industry Solutions Builders Partner, uses the Intel® Core™ CPU and Intel® Core™ Ultra CPU for a product line of boards and chassis. These run today's monitoring workloads, and feature integrated GPU and NPU acceleration that leaves room for tomorrow's AI-enabled workloads either on the device or across a connected fleet of monitoring endpoints.

**Intel® Health & Life Sciences (HLS) AI Suite** adds a repeatable, healthcare-specific way to **demonstrate and measure** representative multimodal monitoring workloads on that hardware, so design teams can quantify performance early in the design cycle.

### Key Takeaways

- **Compact, AI-ready compute**: Intel Core processor / Intel Core Ultra CPU processor with integrated GPU and NPU for heterogeneous execution.
- **Fit for multimodal monitoring**: vitals simulation, rPPG, AI-ECG, and vision workloads running concurrently.
- **Secure and manageable at scale**: Intel vPro®, Intel® AMT, and TPM (Intel® PTT) help protect the platform and lower BOM cost.
- **Performance you can measure**: HLS AI Suite reports CPU, IGPU, and NPU utilization and headroom on Kontron hardware.
- **Headroom for what comes next**: a standards-based interoperability stack and an ECG foundation model already run alongside the visible pipelines.

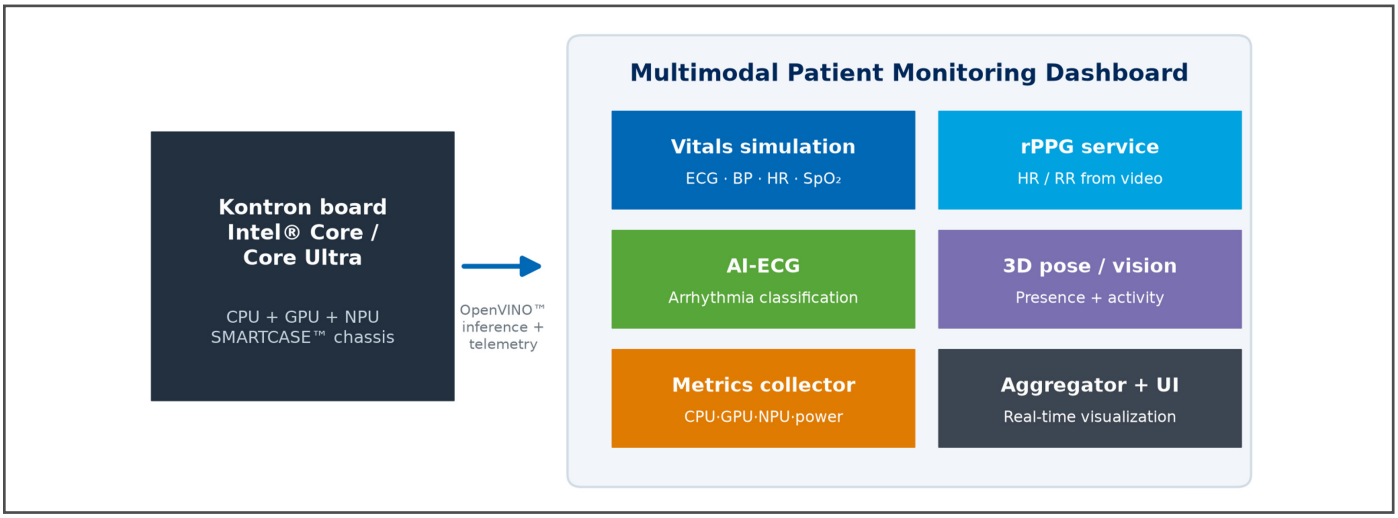



**Figure 1.** Intel HLS AI Suite running representative multimodal patient monitoring workloads on a Kontron Intel Core processor / Intel Core Ultra processor board in a SMARTCASE™ chassis.

## Patient Monitoring Is a Multimodal Edge AI Workload

Modern monitoring already fuses multiple sensor streams in real time, applies complex signal processing, and operates within tight power and always-on reliability envelopes. The market is shifting from isolated, threshold-based monitoring toward **continuous, AI-enabled monitoring**, pushed by distributed care, remote monitoring, hospital-at-hospital, constrained clinical resources, and the need to reduce false alarms. The global patient monitoring market is projected to grow from **$54.3B in 2025 to $118.2B by 2033 (10.5% CAGR)**[i].

The engineering implication: next-generation platforms must support more sensing, more AI, and more continuous operation, running several pipelines *concurrently* while holding BOM cost, thermals, and validation complexity in check. Kontron's Intel-based platforms and the HLS AI Suite are designed to inform that class of decision.

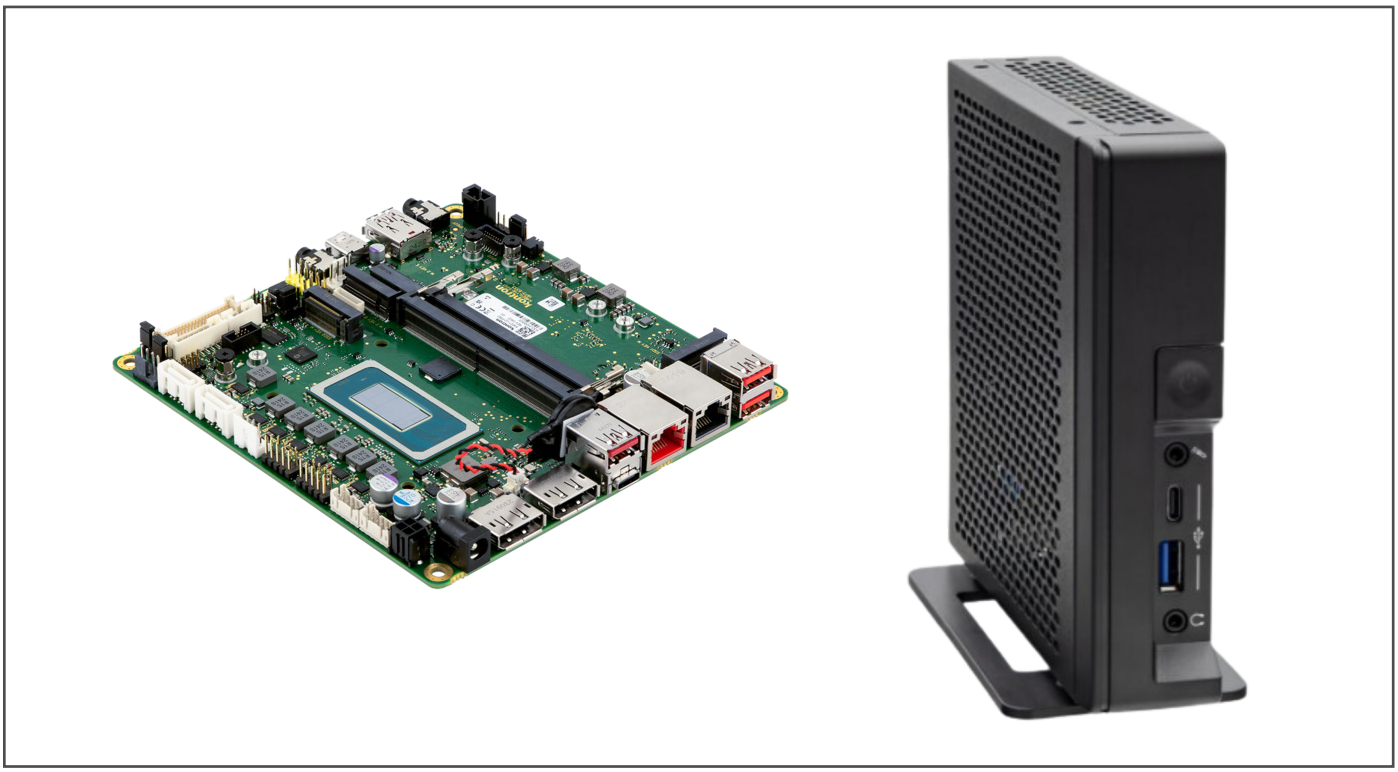


**Figure 2.** The performance of the Intel®-based Kontron K4021-U mSTX is available in a circuit board or in a SMARTCASE chassis.

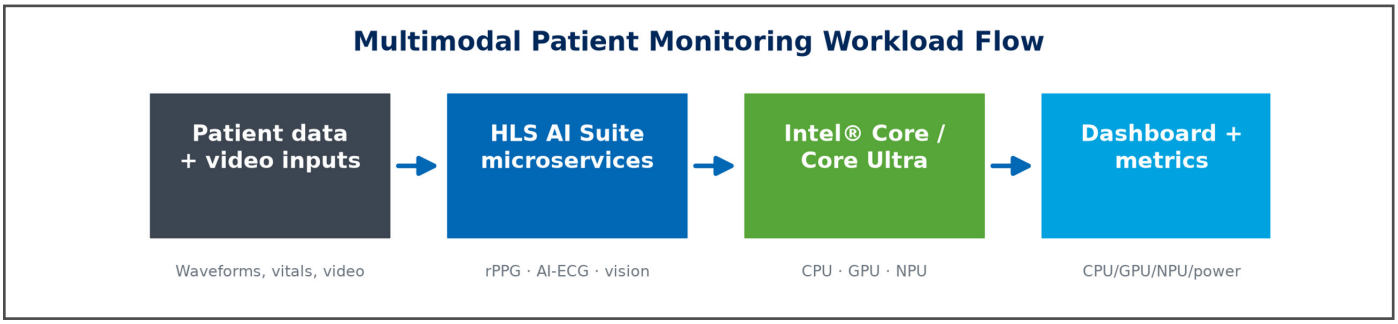



**Figure 3.** Representative multimodal patient monitoring workload flow evaluated with Intel HLS AI Suite on Kontron hardware: waveform data and video in, heterogeneous CPU/IGPU/NPU acceleration, and real-time performance metrics out.

| Use-Case Fit | Performance Visibility | Deployment Relevance |
|---|---|---|
| The suite models representative monitoring workflows: multiparameter vitals simulation, rPPG from video, AI-ECG arrhythmia classification, and vision-based presence/activity tracking. | The metrics collector records CPU, IGPU, NPU, memory, and power utilization while workloads run, showing where acceleration lands and how much headroom remains. | The microservice architecture shows how monitoring applications can be deployed, updated, and scaled across a fleet of edge systems. |

## How the HLS AI Suite Maps to Kontron Hardware

Each HLS AI Suite microservice exercises the Kontron platform differently, so teams can see how concurrent monitoring pipelines land across the CPU, integrated GPU, and NPU.

| HLS AI Suite Component | Role in Patient Monitoring Evaluation | Intel Engine |
|---|---|---|
| **Patient Monitoring Assets Service** | Prepares AI assets, reference models, and OpenVINO™ IR artifacts for downstream services; typically runs once at startup. | **CPU** |
| **rPPG Service** | Uses video preprocessing and the MTTS-CAN model to estimate heart and respiratory rate from video. | **IGPU / NPU (CPU fallback)** |
| **3D Pose / Vision Service** | Detects body key points from monocular RGB cameras for contactless presence and activity monitoring. | **IGPU / NPU** |
| **AI-ECG Arrhythmia Service** | Runs lightweight deep-learning classification on ECG signals end-to-end, without cloud dependency. | **CPU / NPU** |
| **OpenICE Interoperability Layer** | Open-source reference implementation of the ICE interoperability standard. Supplies simulated devices and waveform sources that stand in for a full patient monitor stack. | **CPU** |
| **HuBERT-ECG Foundation Model** | Transformer ECG foundation model pre-trained on 9.1 million 12-lead recordings; represents the next tier of cardiac inference above per-task classifiers. | **iGPU** |
| **Patient Monitoring Aggregator** | Central gRPC endpoint for waveform and numeric vitals; maintains state and time-series histories. | **CPU** |
| **Metrics Collector** | Collects CPU, GPU, NPU, memory, and power telemetry for benchmarking and platform sizing. | **System** |
| **UI Service** | Web dashboard visualizing current and historical vitals alongside optional system metrics. | **CPU / IGPU** |

## Two Workloads Running Out of Frame

The dashboard shows rPPG, pose estimation, and ECG analysis. Two further components carry a large share of the compute and are the reason the sample application works as a platform sizing exercise: an interoperability stack that behaves like a real monitor, and an ECG foundation model on the integrated GPU.

### OpenICE

*Medical Device Plug-and-Play Program, Massachusetts General Hospital*

- Open-source reference implementation of the Integrated Clinical Environment standard, AAMI 2700-1 (previously published as ASTM F2761-09).
- Built on RTI Connext, an implementation of the OMG Data Distribution Service (DDS) standard, with software adapters for patient monitors, ventilators, and infusion pumps.
- Ships with device and patient simulations, so the sample application drives a realistic multi-device environment with no physical bedside hardware attached.
- Runs continuously on the CPU as the device-integration and aggregation tier, the part of a production monitor most often underestimated when sizing silicon.

### HuBERT-ECG

*Self-supervised ECG foundation model, published 2024*

- Pre-trained on 9.1 million 12-lead ECGs drawn from four countries and spanning 164 cardiovascular conditions.
- Transformer encoder behind a convolutional feature extractor, released in three sizes (30.5M, 93.1M, and 188.6M parameters) so the model can be matched to available compute.
- Reported AUROC ranges from 84% in low-data fine-tuning to 99% at scale, and above 90% for 140 conditions when trained on all 164 simultaneously.
- Fine-tuned by its authors across 21 independent datasets and more than 1,600 diagnostic and prognostic targets, including single-lead settings, indicating how broadly one pre-trained model can be re-targeted.
- Executes on the integrated GPU through OpenVINO, which is what makes a transformer of this class practical inside a fanless bedside enclosure.

## What This Means for Platform Selection

Patient monitoring is early in its AI adoption curve, and the compute floor is moving. A monitor that renders waveforms and evaluates threshold alarms has long been served by entry-class silicon. Adding standards-based device interoperability, and then transformer inference on physiological signals, changes the class of part required. The table below sketches that progression as a planning aid for design conversations, not as a benchmark result.

| Monitoring Capability Tier | Representative Workload | Compute Typically Engaged |
|---|---|---|
| **Conventional bedside monitor** | Waveform acquisition and display, numeric vitals, threshold alarms, network push to a central station. | **Intel Atom® SoC or Intel® Processor N-series** |
| **Interoperable monitor** | The above, plus DDS-based multi-device integration, ICE-style aggregation, and FHIR export to hospital systems. | **Intel Core processor** |
| **AI-enabled monitor** | The above, plus concurrent foundation-model ECG inference, camera-based rPPG, and pose estimation. | **Intel Core Ultra (CPU, iGPU, and NPU)** |

The goal of the sample application is to make the third row demonstrable on a system an OEM can actually build, and to show what capacity is left once it is running.

## Kontron on Intel Core processor and Intel Core Ultra processor: Product Fit

Kontron's embedded lineup gives OEM design teams a compact, industrial starting point. The Kontron K4021-U mSTX combined with the SMARTCASE chassis family is a great building block and meets Intel's high performance and compatibility standards for edge AI.

### K4021-U mSTX Platform Highlights

- Intel® Core™ Ultra 7 155U processor or Intel Core™ Ultra 5 125U processor options (Intel® Core™ Ultra Series 1 SoC).
- DDR5-5600 memory, two SO-DIMM sockets, up to 96 GB.
- Intel® Graphics driven by Intel® Xe architecture (up to 8 Xe-cores); up to 4 × 4K or 2 × 8K displays.
- Integrated NPU for dedicated AI acceleration, enabling heterogeneous CPU/IGPU/NPU execution.
- Rich I/O: M.2 storage/WLAN, GPIO, dual COM, USB 3.2 Gen2, Thunderbolt™ 4 / USB4®, 2.5 GbE with TSN.
- Designed and approved for 24/7 operation across a 0–60 °C temperature range.

### SMARTCASE Chassis Highlights

- SMARTCASE S502 is designed for the K4021-U mSTX motherboard.
- Ultra-compact 0.87-liter housing (36 × 147 × 165 mm).
- Horizontal or vertical operation with VESA and wall-mount options.
- Front and rear USB access; optional WLAN antennas, COM cable, DC-in holder, and chassis-intrusion kit.
- A compact system building block for patient-room endpoints, carts, wall-mounted nodes, and bedside gateways.

### Security, Manageability, and BOM Cost

Patient monitoring platforms handle sensitive data and are deployed at fleet scale, so security and remote manageability cannot be afterthoughts. Selecting the Intel Core processor and the Intel Core Ultra processor with **Intel vPro** builds those capabilities into the platform itself, which keeps the **bill of materials (BOM)** lower than adding discrete components and extends manageability across an entire installed base.

| Capability | What it Enables For Patient Monitoring Platforms |
|---|---|
| **Intel vPro** | Hardware-based, below-the-OS security controls that help protect data, detect threats, and support recovery from attacks. |
| **Intel Active Management Technology (AMT)** | Out-of-band remote management, including power control, remote KVM, and remediation. Fleets of monitoring endpoints can be maintained and recovered without an on-site visit. |
| **TPM via Intel Platform Trust Technology (PTT)** | Firmware-based TPM 2.0 for measured boot, key storage, and platform integrity, providing a hardware root of trust without a discrete TPM component on the BOM. |
| **Memory & disk protection** | Platform features including memory encryption and verified/immutable boot help protect data at the point of care against physical and bus-level attacks. |

## What This Provides, and What It Does Not

Neither Intel nor Kontron builds medical devices or makes clinical claims. What this collaboration provides is narrower and more practical: **representative reference workloads, characterized platforms, and a repeatable measurement method**. Medical OEMs still design, validate, and certify their own products. This work shortens the platform-selection stage of that effort.

Kontron's Intel Core processor and Intel Core Ultra processor lineup, paired with the SMARTCASE chassis family, combines a compact industrial form factor with the security and manageability of Intel vPro and the healthcare-specific performance visibility of the HLS AI Suite.

## Engage Early: Pilots and Proofs of Concept

| 1. Evaluate | 2. Validate | 3. Engage |
|---|---|---|
| Run Intel HLS AI Suite on a Kontron Core processor/ Intel Core Ultra processor platform to exercise representative multimodal monitoring workloads. | Use the metrics collector to characterize CPU, IGPU, NPU, memory, and power utilization for your target design point. | Bring Kontron and Intel into pilots and PoCs early in the design cycle to de-risk architecture and roadmap decisions. |

## References


- Intel, "Preparing for the Wave of AI-Enabled Patient Monitoring," Intel Community Blog, accessed 2026. https://community.intel.com/t5/Blogs/Products-and-Solutions/Healthcare/Preparing-for-the-Wave-of-AI-Enabled-Patient-Monitoring-Health/post/1739927
- Intel Industry Solution Builders, "Health and Life Science AI Suite," accessed 2026. https://builders.intel.com/intel-technologies/software/edge-ai-suites/healthcare-life-sciences-ai-suite
- Open Edge Platform GitHub, HLS AI Suite — multimodal patient monitoring "How It Works." https://github.com/open-edge-platform/edge-ai-suites/blob/main/health-and-life-sciences-ai-suite/multi_modal_patient_monitoring/docs/user-guide/how-it-works.md
- Kontron, K4021-U mSTX product information. https://www.kontron.com/en/products/k4021-u-mstx/p186072
- Kontron / DirectIndustry, SMARTCASE™ S502 chassis information. https://www.directindustry.com/prod/kontron/product-5073-2802808.html
- M. Plourde, D. Arney, J. M. Goldman, " OpenICE medical device interoperability platform overview and requirement analysis," peer-reviewed overview. https://pubmed.ncbi.nlm.nih.gov/28734113/
- MD PnP Program, "OpenICE version 2.0" release announcement (aarch64, RTI Connext 7.3, FHIR updates). https://github.com/mdpnp/mdpnp/releases/tag/v2.0.0
- E. Coppola, M. Savardi, M. Massussi, M. Adamo, M. Metra, A. Signoroni, "HuBERT-ECG as a self-supervised foundation model for broad and scalable cardiac applications," medRxiv. https://doi.org/10.1101/2024.11.14.24317328
- HuBERT-ECG source code and model weights. https://github.com/Edoar-do/HuBERT-ECG and https://huggingface.co/Edoardo-BS/hubert-ecg-base
- Market data: Grand View Research, global patient monitoring market estimate (as cited by Intel)


[i] https://www.grandviewresearch.com/industry-analysis/patient-monitoring-devices-market

intel ai

## Notices & Disclaimers


This document and the Intel HLS AI Suite are intended for development, evaluation, and benchmarking. They are not intended for production, clinical, or diagnostic use. Neither Intel nor Kontron manufactures medical devices or makes any medical, clinical, safety, or regulatory claims.

Performance varies by use, configuration, and other factors. Learn more at intel.com/PerformanceIndex. Performance results are based on testing as of the dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details.

No product or component can be absolutely secure. Your costs and results may vary. Intel technologies may require enabled hardware, software, or service activation.

Intel, the Intel logo, Intel Core, Intel vPro, and OpenVINO are trademarks of Intel Corporation or its subsidiaries. Kontron and SMARTCASE are trademarks of Kontron. Other names and brands may be claimed as the property of others.


1126/AL/DJA/PDF Please Recycle (if printed) 370450-001US