

ASX RELEASE

25 June 2026

ASX: NVU

EMASS Achieves AI Function Fusion with Voice Recognition and Real-Time Keyword Detection

First demonstration on ECS-DoT of keyword detection (92% accuracy) and voice recognition fused on a single always-on chip at sub-milliwatt power, entirely on-device.

Highlights

- **First demonstration by EMASS of AI function fusion on ECS-DoT:** keyword detection and voice recognition running together on a single always-on co-processor at sub-milliwatt power, with no host-processor involvement during sensing.
- **Sub-milliwatt always-on operation:** both applications run always-on at under half a milliwatt of average power (400–500 μ W), on silicon independently benchmarked at up to 20x the energy efficiency of leading edge-AI chips¹.
- **Accurate and instant:** the keyword-detection model is a compact 8-bit network demonstrated at 92% accuracy with approximately 4 ms inference latency, fast enough to feel instantaneous to the user.
- **Always-on co-processor, host stays asleep:** ECS-DoT performs continuous listening and wakes the device's main processor only on a recognised command or a verified user, complementing rather than replacing the host CPU.
- **Security and privacy at the edge:** keyword detection and voice recognition run entirely on-device, so audio never leaves the chip, materially reducing the privacy surface of an always-on voice interface.
- **Interoperable with industry standard hardware:** implemented on the ECS-DoT evaluation board with a standard PCM digital MEMS microphone, with no specialised hardware required.
- **Applicable across multiple high-value markets:** next-generation computing devices including tablets and PCs, alongside consumer, professional and industrial always-on endpoints, positioning ECS-DoT as a always-on intelligence layer for the next generation of devices.

Nanoveu Limited (ASX: NVU, OTCQB: NNVUF) ("Nanoveu" or the "Company"), a technology company specialising in advanced semiconductor, visualisation and materials sciences, is pleased to announce EMASS has had a breakthrough demonstration in real-time keyword detection and voice recognition fused as two AI functions, running entirely on its ECS-DoT edge AI co-processor. Both capabilities run at sub-milliwatt average power and around 4 ms latency directly on ECS-DoT, bringing accurate, low-latency voice intelligence to computing devices such as tablets and PCs while the host processor remains asleep until a command or an authorised user is detected.

Natural Language Communication Is the New Way to Communicate with Devices and AI

Voice is becoming the primary way we control computing devices.




Figure 1: Showcasing the shift from buttons and touch to voice and natural language as the emerging primary control interface for computing devices

Voice is becoming a primary control surface for computing devices. Manufacturers increasingly seek those experiences to run on-device for responsiveness, privacy and reliability, rather than depending on the cloud, with Voice AI Agents market projected to reach **USD 47.5 billion** by 2034².

Delivering always-on voice on a compact, power-constrained device has traditionally been limited by power and thermal budgets. Keeping a microphone pipeline and a capable processor continuously active to listen for commands can increase power consumption and thermal load. Conversely, duty-cycling those processors to save power can introduce latency and increase the risk of missed activations.

Delivering always-on, instantly responsive voice processing has traditionally involved a three-way trade-off between power, latency and privacy:

- **Power:** keeping microphones and a capable processor continuously active delivers responsiveness, but drains the battery and adds thermal load.
- **Latency:** duty-cycled and wake-on-demand schemes preserve power, but introduce delay and depend on repeatedly waking higher-power processors.
- **Privacy:** an always-listening microphone pipeline that processes audio off the sensor creates an always-on privacy surface.

Sub-Milliwatt AI Function Fusion: Recognition and Keyword Detection All in One Chip

EMASS has demonstrated two complementary voice capabilities running entirely on the ECS-DoT edge AI co-processor. The first is keyword detection: recognising spoken commands and wake words in real time. The second is voice recognition: confirming whether the person speaking is the enrolled, authorised user. Together, they let a device respond to voice instantly and personalise or gate that response to the right person, all on-chip. This is the first time ECS-DoT has been demonstrated running two AI functions simultaneously. Rather than choosing between recognition or keyword detection, the chip performs both at once, on-device, with audio never leaving the device.

- **Low power:** each application ran always-on at under half a milliwatt of average power (400–500 μ W), enabling continuous operation without materially affecting battery life.
- **Real-time response:** keyword detection and voice recognition were demonstrated at the co-processor, without first waking a higher-power processor.

- **Privacy:** audio does not leave the device, materially reducing the privacy surface of an always-on voice interface.




Figure 2: Diagram depicting ECS-DoT running keyword detection and voice recognition simultaneously on a single co-processor at under 0.5 mW average power, with audio never leaving the chip

Both capabilities were implemented on the ECS-DoT evaluation board using a standard PCM digital MEMS microphone, with no additional specialised hardware required. The keyword-detection model is a compact 8-bit network that detects the most likely spoken and was tested over 1,500 times to measure how reliably it identifies the correct spoken keyword.

92% Top-1 and 97% Top-3 Accuracy

When the model makes a prediction, it produces a ranked list of likely words rather than a single answer. Two industry-standard measures describe how often it gets things right. The first, known as top-1 accuracy, counts a result as correct only when the model's single best guess exactly matches the spoken word. This is the strictest measure and the one the machine learning community treats as the benchmark. **On this measure the model reached 92%.** The second, top-3 accuracy, counts a result as correct when the right word appears among the model's three best guesses, and **on this measure the model reached 97%.**

The voice-recognition model produces a per-user voice signature matched against an enrolled template to distinguish the authorised user from others. The sub-milliwatt result was achieved through aggressive model compression and tuning tailored to the ECS-DoT architecture, combined with the chip's on-chip energy management.

These are two AI models running together on a single low-power co-processor, which is itself a demonstration of ECS-DoT's multi-model design. The same engine that runs these voice models is built to run image and other sensor models, so the capability demonstrated here forms the foundation for broader on-device, multimodal intelligence rather than a single-function voice block.

Why Our Solution Outperforms Traditional Voice Architectures

Comparison of Energy Between ECS-DoT and Conventional Solutions




Figure 3: Chart depicting the energy comparison between a conventional always-on microphone and host CPU solution versus the ECS-DoT co-processor at sub-milliwatt baseline power

Existing common approaches to voice activation tend to fail at least one leg of the trade-off:

- Always-on microphone with host processing is responsive but power hungry in comparison: the microphone and a capable processor draw current continuously, the device processes audio off the sensor, and the main processor stays awake to do it.
- Duty-cycled and wake-on-demand systems save power but add lag in the trade off: they preserve power but introduce delay and depend on repeatedly waking higher-power processors, so the device is not continuously ready to respond.

ECS-DoT resolves this by performing the always-on sensing itself, at sub-milliwatt average power, and waking the host processor only when a recognised command or a verified user calls for it. The host CPU is complemented, not replaced; it sleeps until ECS-DoT has something worth waking it for. Because keyword detection and voice recognition run on-device, audio never needs to leave the device.

PROPERTY	ALWAYS-ON MIC + HOST CPU	DUTY-CYCLED / WAKE-ON-DEMAND	ECS-DOT CO-PROCESSOR
Power continuous	No	Partial	<0.5 mW demonstrated
Latency always-ready	Yes, but costly	Wake delay	~4 ms at co-processor
Privacy audio on-device	Audio off-sensor	Partial	Never leaves chip
Speaker verification on-device	Host only	No	Yes, per-user signature

Figure 4: Table showing property comparison across voice architectures, contrasting always-on microphone with host CPU, duty-cycled systems, and keyword detection with voice recognition on ECS-DoT

Beyond Voice: A Multimodal, Multi-Sensor Co-Processor

This milestone is voice, but the capability underneath it is general. The same engine that lets ECS-DoT extract meaning from an audio stream at sub-milliwatt power applies across other sensing modalities. ECS-DoT is multi-sensor and multimodal by design, built to run vision and motion models alongside audio within a single ultra-low-power, always-on budget. On a computing device such as a tablet or PC, one low-power co-processor can deliver layered, sensor-fused intelligence

Modality	Capability
Voice	Instant, personalised, hands-free control through keyword detection and voice recognition, demonstrated at sub-milliwatt average power.
Vision	On-device image classification and presence or user recognition; answering who is at the device and whether anyone else is present, without images leaving the device.
Other Sensors	Motion and additional sensor streams fused with audio and vision to drive smarter always-on behaviour.

Because all of this runs on ECS-DoT rather than the host processor, these capabilities can be layered onto a device without the power and thermal cost of keeping the main processor awake, and at the low latency an always-on experience requires. Today that is voice; the same platform is designed to scale to image and beyond, positioning ECS-DoT as an always-on, multimodal sensing co-processor for computing devices.

Next Steps

EMASS will continue to advance the always-on voice solution and progress it toward commercial deployment:

- **Technical development:** Expanding the keyword vocabulary, hardening keyword detection and voice recognition performance across users and operating conditions, and progressing toward production-ready reference designs for computing devices.
- **Commercial engagement:** Active discussions with computing-device manufacturers, processors and platform integrators on deploying on-device voice modules.
- **Design-ins:** Pursuing design-in engagements to integrate ECS-DoT-based voice processors into next-generation tablets, PCs and always-on consumer and industrial endpoints.

Mark Goranson, Nanoveu's CEO of Semiconductor Technologies, commented: *"Today we are demonstrating two AI voice models, keyword detection and voice recognition, running together on ECS-DoT at sub-milliwatt average power, which is significant in its own right. But the bigger point is what it proves about the platform: ECS-DoT is multi-sensor and multimodal, so the same always-on engine is built to run image and other sensor models on these devices too. That provides optionality for device makers to add layered, low-power intelligence across voice, vision and beyond, while the main processor stays asleep."*

References

1. Based on MLPerf Tiny benchmark testing refer announcement 22 October 2024 comparing ECS-DOT against publicly available benchmark results from peer edge-AI semiconductor companies.
2. Markets.us Voice AI Agents Market, April 2025

This announcement has been authorised for release by the Board of Directors.

-ENDS-

Nanoveu Media

Alfred Chong, Nanoveu MD and CEO

P: +65 6557 0155

E: info@nanoveu.com

About Nanoveu Limited

Further details on the Company can be found at <https://nanoveu.com/>.

EMASS is a pioneering technology company specialising in the design and development of advanced systems-on-chip (SoC) solutions. These SoCs enable ultra-low-power, AI-driven processing for smart devices, IoT applications, and 3D content transformation. With its industry-leading technology, EMASS will enhance Nanoveu's portfolio, empowering a wide range of industries with efficient, scalable AI capabilities, further positioning Nanoveu as a key player in the rapidly growing 3D content, AI and edge computing markets.

EyeFly3D™ is a comprehensive platform solution for delivering glasses-free 3D experiences across a range of devices and industries. At its core, EyeFly3D™ combines advanced screen technology, sophisticated software for content processing, and now, with the integration of EMASS's ultra-low-power SoC, powerful hardware.

Nanoshield™ is a self-disinfecting film that uses a patented polymer of embedded Cuprous nanoparticles to provide antiviral and antimicrobial protection for a range of applications, from mobile covers to industrial surfaces. Applications include Nanoshield™ Marine, which prevents the growth of aquatic organisms on submerged surfaces like ship hulls, and Nanoshield™ Solar, designed to prevent surface debris on solar panels, thereby maintaining optimal power output.

Forward Looking Statements This announcement contains 'forward-looking information' that is based on the Company's expectations, estimates and projections as of the date on which the statements were made. This forward-looking information includes, among other things, statements with respect to the Company's business strategy, plans, development, objectives, performance, outlook, growth, cash flow, projections, targets and expectations and related expenses. Generally, this forward-looking information can be identified by the use of forward-looking terminology such as 'outlook', 'ambition', 'anticipate', 'project', 'target', 'potential', 'likely', 'believe', 'estimate', 'expect', 'intend', 'may', 'mission', 'would', 'could', 'should', 'scheduled', 'will', 'plan', 'forecast', 'evolve' and similar expressions. Persons reading this announcement are cautioned that such statements are only predictions, and that the Company's actual future results or performance may be materially different. Forward-looking information is subject to known and unknown risks, uncertainties and other factors that may cause the Company's actual results, level of activity, performance, or achievements to be materially different from those expressed or implied by such forward looking information.