

ASX RELEASE

16 June 2026

ASX: NVU

EMASS Achieves Microphone-Free Keyword Spotting at Under 1 Milliwatt on ECS-DoT

Real-time keyword spotting from bone-conduction vibration on the ECS-DoT edge AI SoC, at under 1 milliwatt and with no microphone advancing the hearables reference design announced in March 2026¹.

Highlights

- **Novel sensing architecture using bone conduction:** EMASS demonstrates always-on, real-time keyword spotting on the ECS-DoT edge AI SoC from jawbone vibration, capturing speech through the body rather than the air at a sub-milliwatt power envelope.
- **Sub-1mW always-on operation:** the complete keyword-spotting application runs always-on at under 1 milliwatt of computational power, on silicon independently benchmarked at up to 20x the energy efficiency of leading edge-AI chips.
- **Overcomes the key trade-off in wearable voice design:** noise immunity, low power and real-time response delivered simultaneously, where traditional architectures have to weigh two of the three.
- **Built-in noise immunity:** ambient noise is excluded at the source with no noise-cancellation processing required, and because ambient audio is never captured, user privacy is inherently protected.
- **Demonstrated on market-leading hardware:** ECS-DoT is compatible with the STMicroelectronics IMU, with no specialised hardware required.
- **Hearables market forecast to grow from USD 62.22 billion in 2026 to USD 107.1 billion by 2031²,** with voice emerging as the primary control interface for next-generation devices
- **Applicable across multiple high-value markets:** AR glasses, VR headsets, professional and industrial communications, and defence and critical-infrastructure systems.
- **Commercial engagement:** developing production-ready reference designs across the hearables and AR ecosystem, with active customer engagement ongoing.

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 a breakthrough in always-on edge AI sensing. EMASS has demonstrated fully functional, real-time keyword spotting on its ECS-DoT edge AI SoC using bone-conduction vibration sensing at under 1 milliwatt of computational power, addressing the two most persistent constraints in wearable voice interface design: power consumption and ambient noise. This advances the bone-conduction hearables reference design EMASS announced in March 2026: rather than using motion to wake an audio subsystem, ECS-DoT now performs keyword spotting directly on the vibration signal, removing the microphone from the detection path.

¹ Refer to ASX Announcement dated 5 March 2026

² Data from Mordor Intelligence – Hearables Market Size & Share Analysis, 14 January 2026

The technique captures speech through the body rather than the air, sensing the micro-vibrations that propagate through the user's jawbone as they speak. Because the signal is taken from vibration rather than air, ambient noise is largely excluded from the signal path, and ECS-DoT's sub-milliwatt inference budget allows the system to remain always on without materially affecting battery life. The result has direct application across next-generation hearables, AR glasses, and professional, industrial and defence communication systems.

Hearables: Market Sizing and Key Problems

Hearables market size




Figure 1: Showcasing Projected Market Growth of the Hearables Market²

The global hearables market is forecast to grow from USD 62.22 billion in 2026 to USD 107.1 billion by 2031, driven by demand for smarter, more capable devices in ever-smaller form factors. Central to this next generation of devices is a shift in how users interact with them: voice is becoming the primary control interface. As hearables evolve from audio accessories into AI-assistant endpoints, users expect to activate and command their devices through speech alone, without buttons, taps or gestures.

Delivering this has forced manufacturers into a three-way trade-off between noise immunity, power and latency:

- **Background noise immunity:** Microphones capture everything in the air. Making voice control work in wind, crowds and machinery noise requires continuous noise-cancellation and signal processing, which consumes significant power and still degrades in the loudest environments.
- **Power:** Keeping microphones and their processing pipeline continuously active delivers responsiveness, but drains the battery and creates an always-listening privacy surface.
- **Latency:** Motion-triggered and duty-cycled wake systems preserve power, but introduce delay and depend on repeatedly activating higher-power processors for voice analysis.

Designers can optimise for any two of these properties; achieving all three simultaneously has remained out of reach. The constraint is fundamental to the silicon: battery life and thermal limits remain the primary technical hurdles in hearable design, and placing an AI chip next to the sensor has traditionally meant a processor with power consumption far too high for compact, battery-constrained products such as earbuds and headsets.

Our Solution: Noise-Immune, Always-On Voice Sensing at Under 1 Milliwatt

EMASS has demonstrated fully functional, real-time keyword spotting on its ECS-DoT edge AI SoC, capturing speech not through the air but by sensing the micro-vibrations that propagate through the user's jawbone as they speak, and doing so at under 1mW of computational power. Today, such sensors are found primarily in premium earbuds and communication headsets, where they improve voice isolation and call clarity. EMASS is pioneering a new use for this sensing modality: detecting voice commands and wake words directly from bone vibration, at the sensor level, without waking the host processor.




Figure 2: Showcasing bone conduction with ECS-DoT resolving the noise, power and latency trade-off.

This architecture delivers three properties that have historically been difficult to combine in a wearable voice interface:

- **Noise immunity:** Speech is captured through the jawbone rather than the air, so ambient noise is largely excluded from the signal path. Wind, crowds and machinery noise do not enter the processing pipeline, and no noise-cancellation processing is required. This noise immunity comes from the sensing method itself rather than from additional processing.
- **Low power:** The complete keyword-spotting application ran in real time at under 1 mW of computational power, enabling always-on operation without materially affecting battery life.
- **Real-time response:** Keyword detection was demonstrated at the sensor level, without first waking a higher-power processor for voice analysis.

The same architecture also delivers a privacy benefit at no additional cost. Because ambient audio never enters the system, nearby conversations cannot be inadvertently captured, materially reducing the privacy surface of always-on voice interfaces.

The demonstration was implemented using a standard STMicroelectronics inertial measurement unit (IMU) configured to detect jawbone vibrations, paired with ECS-DoT, with no additional specialised hardware required. The sub-milliwatt result was achieved through aggressive model compression and AI model scaling tailored to the chip's architecture, combined with ECS-DoT's advanced on-chip energy management, and confirms the platform's ability to perform edge AI inference across vibration-based sensing modalities.

Why Our Solution Outperforms Traditional Voice Architectures

Conventional



Power is spent at **every stage** fighting noise that was captured unnecessarily.

EMASS



Ambient noise **never enters the signal path**; the host sleeps until a keyword is detected.

*Figure 3: Flowchart of conventional microphone signal path versus EMASS
 bone-conduction sensing with ECS-DoT*

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

- **Always-on microphones are responsive but power hungry:** The mic draws current the whole time it is listening, and the noise cancellation needed to work in real conditions draws more on top. It is also recording everything around it, which is a privacy concern. In louder environments, more power is consumed processing out ambient noise than the microphone captured in the first place.
- **Motion-triggered and duty-cycled wake systems save power but add lag:** It preserves power, but introduce delay and depend on repeatedly activating higher-power processors for voice analysis, so the device is not continuously ready to respond.

Bone-conduction sensing with ECS-DoT avoids this trade-off rather than balancing within it. In microphone-based architectures, noise immunity comes at a power cost, because more processing is required to extract the voice from surrounding noise. With bone conduction, the ambient noise is not captured in the first place, so neither the processing cost of removing it nor the privacy exposure of recording it arises.

ECS-DoT is what makes this deployable. Bone-conduction sensors have existed for years, but pairing them with always-on AI inference required power budgets no compact device could support. ECS-DoT's sub-milliwatt inference closes that gap, allowing the device to remain always on while staying power-efficient and private.

For next-generation hearables, this supports voice-first interaction in demanding environments such as factories, busy streets, public transport and live events, while preserving battery life in power-constrained devices.

PROPERTY	ALWAYS-ON MICROPHONE	MOTION-TRIGGERED WAKE	BONE CONDUCTION + ECS-DoT
○ Noise immunity	Processing-dependent	× None	✓ Structural
⊞ Low power	× No	● Partial	✓ <1 mW demonstrated
⚡ Low latency	✓ Yes	× Wake delay	✓ Real-time at sensor
🔒 Privacy	× Always listening	● Partial	✓ Audio never captured

Figure 4: Table showing property comparison across wearable architectures

AR, VR and Defence: Applicable to Markets Beyond Hearables

The sensing architecture demonstrated here extends well beyond earbuds. At its core, ECS-DoT is extracting meaningful signal from mechanical vibration, in real time, at under 1 milliwatt of computational power, and that capability applies across the full spectrum of voice-first wearable and edge devices:

- **AR glasses** with discreet, hands-free control designed to work even when the user speaks quietly.
- **VR headsets** where always-on voice input enables natural interaction without controller dependence or battery penalty.
- **Professional and industrial communication gear** where noise immunity, privacy and battery life are mission-critical.
- **Defence and critical infrastructure communication systems** where reliability and discretion are non-negotiable.

Beyond keyword spotting, EMAS believes bone-conduction sensing combined with ultra-low-power edge AI could enable future applications such as voice authentication, user-intent detection, context-aware wearable AI, and other advanced human-machine interfaces.

Next Steps

EMAS will continue to advance the bone-conduction sensing solution and progress it toward commercial deployment:

- **Technical development:** Expanding the keyword vocabulary, hardening detection performance across users and operating conditions, and progressing toward production-ready reference designs.
- **Commercial engagement:** Active discussions with device manufacturers and module makers across the hearables and AR ecosystem.
- **Design-ins:** Pursuing design-in engagements to integrate ECS-DoT-based vibration sensing into next-generation products.

Mark Goranson, Nanoveu's CEO of Semiconductor Technologies, commented: *"This demonstration matters as much for what it confirms about ECS-DoT as a platform as for what it enables in wearables. Recognising speech directly from bone-conduction vibration at under one milliwatt, with no microphone in the signal path, shows the chip can extract a meaningful signal from mechanical vibration in real time, within a power budget that compact, battery-powered devices can support. We*

see that core capability as applicable well beyond hearables, to other products that need to interpret vibration on tight power and thermal budgets, and it is central to how we are positioning ECS-DoT."

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.