eMemory 1Q26 Earnings Call Transcript

May 8th, 2026, 16:00-17:00 Taiwan Time

OPENING REMARKS

Dr. Charles Hsu, Chairman

Good afternoon, everyone. In the first quarter of this year, eMemory continued to deliver record-high operating performance. In addition to sustained growth in 3nm licensing for AI server CPU applications, we also successfully expanded into new areas such as high-speed interface applications.

Meanwhile, as AI infrastructure rapidly evolves, industry demand for hardware security is undergoing a fundamental upgrade. The Open Compute Project (OCP)'s Foundation Chiplet System Architecture (FCSA), introduced in 2025, explicitly requires every chip within a chiplet system to incorporate a Hardware Root of Trust. This signifies that every chip in future AI systems will require built-in hardware security mechanisms to ensure system integrity and data security.

Furthermore, the rapid advancement of Agentic AI, together with increasing government requirements worldwide for Post-Quantum Cryptography (PQC), is driving what we believe to be the most significant hardware security upgrade cycle in the past two decades. In response to this trend, we are actively evolving from an IP provider into a system-level security solution provider. Our internally developed Hardware Security Module (HSM) product is expected to be officially licensed for customer mass production deployment upon completion of U.S. National Institute of Standards and Technology (NIST) validation. We believe eMemory will be one of the key beneficiaries of this global hardware security upgrade trend.

At the same time, the core technological foundation of eMemory remains Logic NVM. We have long focused on enabling non-volatile memory technologies through standard logic processes. In addition to being embedded into SoCs in IP form, these technologies also have the potential to evolve into standalone memory products.

Recently, global Flash memory supply has remained constrained. eMemory's key advantage lies in our ability to directly leverage mature logic process capacity to produce non-volatile memory, avoiding the capacity bottlenecks associated with conventional dedicated Flash manufacturing lines. We are currently working closely with multiple foundries to advance related technologies and commercialization efforts.

This development not only carries strong strategic value for foundries and end-system customers, but also has the potential to become the next major growth driver for eMemory's revenue and profitability, following OTP and PUF security solutions.

Next, we invite our CFO, Mr. Joseph Hsia, to present the operating results for the first quarter of 2026.

FINANCIAL RESULTS

Joseph Hsia, Financial Officer

Q1 2026 Financial Results

Good afternoon, everyone. Thank you for joining us. Allow me to walk you through our financial results for the first quarter of 2026.

Q1 revenue was NT$ 1.09 Billion (One billion and ninety-four million NT dollars), increasing 4.4% quarter-over-quarter and 20% year-over-year.

Operating expenses were NT$ 432 Million (Four hundred and thirty-two million NT dollars), up 4.6% sequentially and up 10.8% year-over-year.

As a result, our operating income reached NT$ 662 mil (Six hundred and sixty-two million NT dollars), representing an increase of 4.2% quarter-over-quarter and an increase of 26.8% year-over-year.

Operating margin was 60.5%, relatively flat sequentially, down 0.1 percentage point but increased by 3.2 percentage points year-over-year. Net income for the quarter was NT$ 596 mil (Five hundred and ninety-six million NT dollars), experienced an increase of 5.9% sequentially and an increase of 29.1% year-over-year.

EPS for this quarter was NT$ 7.98.

Revenue across Different Streams

Next, let's move on to revenue mix by licensing and royalty.

Licensing in the first quarter accounted for 34.8% of the total revenue, up 9.9% sequentially and up 58.6% year-over-year. On U.S. dollar basis, licensing grew 7.8% quarter-over-quarter and 64.4% year-over-year.

Royalties on the other hand contributed 65.2% of the total revenue, increasing 1.6% sequentially and increasing 6.2% year-over-year. On U.S. dollar basis, there was a 1.2% decrease quarter-over-quarter but a 11% increase year-over-year.

Overall, **Revenue** increased by 4.4% quarter-over-quarter and 20% year-over-year. On U.S. dollar basis, the growth was stronger at 1.8% quarter-over-quarter and 25.2% year-over-year.

Revenue by Technology

With that, I will break down revenue contribution by specific IPs.

NeoBit accounted for 20% of total revenue in the first quarter. Its licensing revenue decreased 28.4% sequentially and decreased 13.7% year-over-year, while royalties increased 3.1% sequentially but down 4.2% year-over-year.

NeoFuse accounted for 59.8% of the total revenue in the first quarter. Its licensing revenue increased 39.8% sequentially and increased 20.7% year-over-year. In terms of royalty revenue, NeoFuse royalties increased by 0.5% sequentially and increased by 7.4% year-over-year.

For PUF-Based Security IPs, it contributed 12.2% of the total revenue in the first quarter. Its licensing revenue increased 21.5% sequentially and increased 606.9% year-over-year, while its royalties increased significantly, accounting for about 1% of total royalties.

MTP technology accounted for 8% of total revenue in the first quarter. Its licensing revenue decreased 5.5% sequentially but increased 37.9% year-over-year. Royalty from MTP was down 3.1% sequentially but increased by 45% year-over-year.

Royalty Revenue by Wafer Size

Now, let's look at royalties for 8-inch and 12-inch wafers.

8-inch wafers accounted for 33.7% of royalties, down 5.5% sequentially and down 15.9% year-over-year. On U.S. dollar basis, this represents a sequential decrease of 8.2% and a year-over-year decrease of 12%.

12-inch wafers contributed 66.3% of royalties, up 5.6% sequentially and up 22.5% year-over-year. On U.S. dollar basis, this represents a sequential increase of 2.8% and a year-over-year increase of 28%.

A total of 139 licensed product tape-outs were completed in the first quarter. Further details will be provided in the management report, which will be released shortly after this earnings call.

Next, I would like to invite our President Michael Ho to share more about our future outlook.

FUTURE OUTLOOK

Michael Ho, President

Good afternoon, everyone. In the following section, I will address our future outlook.

Licensing Revenue is driven by an increase in the number of licenses from both foundries and fabless customers, as well as higher average selling prices (ASP) per license. Demand remains strong for advanced-node technologies, AI-related applications, security-related solutions, and next-generation Flash technologies.

Royalty Revenue continued to accelerate, driven by upgrades of existing products, higher ASPs from new advanced-node applications, expanding PUF royalty contributions, and a growing mix of higher-royalty-rate MTP products.

New IP Technologies

1. Continued development of next-generation OTP and PUF-based hardware security solutions targeting sub-2nm process nodes.
2. NeoFlash is expanding into both embedded and standalone 1T Flash applications, leveraging logic-process-based architectures for improved scalability and cost efficiency.

Business Development Platforms

1. **Chiplet Security Platform**

   The company is collaborating with ecosystem partners to develop end-to-end Chiplet security solutions, covering key technologies including supply chain security, Chiplet authentication, secure communication, and hardware root of trust. These solutions are designed to address the increasing security requirements driven by AI and heterogeneous integration architectures.

2. **Data Center Caliptra Platform**

   Targeting data center and AI server applications, we are developing a comprehensive IP platform that combines PUFrt (PUF-based Root of Trust) with a Security Subsystem architecture. This platform enables customers to accelerate the adoption and mass production of hardware security solutions compliant with the Caliptra framework and future data center security standards.

3. **Root of Trust / CPU Platform Collaboration**

   Our hardware security technologies have been successfully integrated into the latest-generation AI/AGI CPU platforms of major CPU vendors, providing chip-level Root of Trust, Secure Boot, device identity, and secure key protection capabilities. These technologies enhance trustworthiness and security across AI systems from chip to system level.

4. **HSM Edge Server Platform**

   Our PUF-based HSM Edge Server is positioned as a key infrastructure component for future Security as a Service (SECaaS) platforms. The platform supports critical applications including secure OTA updates, privacy protection, device identity management, and Post-Quantum Cryptography (PQC) migration for automotive, medical, industrial control, Edge AI, and other high-security systems, helping customers build sustainable next-generation security architectures.

This concludes my comments. I will now hand over to Felix, Head of Digital Marketing.

FEATURED TOPIC

Dr. Felix Hsu, Head of DM

(Page 13: Hardware Security for AI Agents)

Good afternoon, everyone. Thank you for being here.

The AI industry has spent the last few years building agents that can act on the world — but not enough time building the foundation that makes those actions trustworthy.

That serious need for security is what we're going to talk about today.

(Page 14: When AI Agents Reach the Real World, Security Must Be Absolute)

An AI agent is not a chatbot. A chatbot answers. An agent **acts**. It runs in a loop — it gathers context, it reasons, it decides, it executes in the real world, and it records the result. Then it does it again. And again.

There are three key things we need to know about this loop

First — it's autonomous. The agent decides for itself. There is no human pressing "approve" between steps.

Second — it runs at machine speed. We're talking over a thousand decisions per minute. No compliance officer, no security team, no human reviewer can keep up with that in real time. It is mathematically impossible.

And third, this is an important point, different steps carry vastly different stakes. When the agent is reasoning internally, a mistake is recoverable. But the moment that loop touches the external world — money, contracts, identity, infrastructure — the stakes become absolute. There is no undo button on a wire transfer. There is no rollback on a signed contract.

(Page 15: Integrity Risks are Mostly Solved, Authorization Ones are Not)

At every single step of that agent loop, there are two distinct categories of risk.

The first is **integrity risk** — whether the agent is thinking with the right information. Whether it's being fed false data. Whether its reasoning engine is the one we deployed, or has been silently swapped. Whether its memory has been tampered with.

The second is **authorization risk** — whether the agent is allowed to do what it's about to do. Whether the action was approved by a human. Whether the agent is even who it claims to be.

These are fundamentally different problems. And here's what the market has missed: most of the security industry today only addresses one of them — usually integrity,

through software signing and model verification. The authorization side — the matter of who has the right to act — is largely being handled by software policies that an attacker, or even a misaligned agent, can route around.

You cannot have governance without solving both. A perfectly authentic model executing an unauthorized transaction is still a catastrophe. An authorized agent running tampered logic is still a catastrophe.

Both columns have to be risk-free. Today, in production AI deployments, they are not.

(Page 16: Real Governance Can’t Be Achieved Without Hardware Root)

Every major regulatory framework now being written — the EU AI Act, the emerging U.S. governance frameworks, the financial sector guidelines — they all converge on the same four requirements.

Human-in-the-loop approval for high-stakes decisions. Tamper-proof audit logs. Protection of models from theft. And zero-trust authentication on every interaction.

Every AI vendor in the world will tell you they offer these things. And every one of those promises, in current implementations, has the same fatal weakness: **it lives in software.** In the case **when PUF is absent**, having software is not sufficient, as shown in the right column, particularly, **software signatures can be forged. Software audit trails can be rewritten by anyone with sufficient privilege. Software keys can be extracted. Software identity can be impersonated.**

What this means in practice is that today's AI governance is, to use a phrase we keep hearing from CISOs, "governance theater." It looks compliant on paper. It does not survive contact with a determined attacker.

The only way to make these promises **enforceable** rather than aspirational is to follow the middle column, knowing what governance requires, and anchoring the promises in something that **we know the exact identity of**, **with signatures that cannot be copied or extracted, keys that cannot be stolen, which essentially is describing hardware**. In other words, it has to be a property of the silicon itself — unique to each chip, generated at the moment of need, and never stored anywhere it can be stolen, all arrows pointing to PUF.

(Page 17: PUF is the Hardware-Anchored Trust for Agentic AI)

We provide the foundation layer. The hardware root of trust. And on top of that foundation, three properties become possible — properties that no software-only stack can deliver.

Identity. Every AI agent — every device, every chiplet, every endpoint — gets an unforgeable identity that is intrinsic to its silicon. It cannot be cloned, because the physical randomness it derives from cannot be replicated, even by the foundry that made the chip.

Integrity. Models, prompts, and agent logic stay sealed to the hardware that's authorized to run them. Steal the model file, and it's a useless blob of encrypted bytes. There is no key file to steal alongside it, because the key is never written down.

Authority. Human-in-the-loop approval, audit logs, policy enforcement — all the things regulators are mandating — become cryptographically enforceable rather than procedurally promised. The audit log is no longer "what the system says happened." It is provable, in court if necessary, by anyone with the public verification keys.

The AI governance platforms, the agent orchestration tools, the compliance dashboards — they are all valuable, and they are all building real businesses. In fact, eMemory is positioned at the hardware layer that everything else depends on.

In summary, by moving security from the software layer to the silicon itself, eMemory enables agentic AI to operate with the level of autonomy and authority required for high-stakes industries like finance, healthcare, and autonomous infrastructure.

That concludes my talk, thank you for your time. Next, we will enter the Q&A session.

CLOSING REMARKS

Dr. Charles Hsu, Chairman

For more information about our PUF-based security IPs and technology, we encourage you to visit our PUFsecurity website at https://www.pufsecurity.com/ and check out our articles and other materials.

Thank you once again for your patience and support for eMemory. We will continue to work hard on technology and IP innovation and PUF-based hardware security solutions for our customers and bring higher returns for our shareholders. Thank you!