

Penguin Solutions 3rd Quarter Fiscal 2026 Earnings Call July 7, 2026

Thank you, operator. Good afternoon and thank you for joining us on today's earnings conference call and webcast to discuss Penguin Solutions' third quarter fiscal 2026 results. On the call today are Kash Shaikh, Chief Executive Officer, and Nate Olmstead, Chief Financial Officer.

You can find the accompanying slide presentation and press release for this call on the Investor Relations section of our website. We encourage you to go to the site throughout the quarter for the most current information on the company. I would also like to remind everyone to read the note on the Use of Forward-Looking Statements that is included in the press release and the earnings call presentation. Please note that during this conference call, the company will make projections and forward-looking statements, including - but not limited to - statements relating to; statements about the company's growth trajectory, financial outlook, and preliminary expectations for future fiscal periods; business, plans and strategy; product development and innovation; market demand and shifts; supply chain conditions and cost assumptions; leadership transitions; strategic agreements; and existing and potential collaborations.

Forward-looking statements are based on current beliefs and assumptions, are not guarantees of future performance, and are subject to risks and uncertainties, including, without limitation, the risks and uncertainties reflected in the press release and the earnings call presentation filed today as well as in the company's most recent annual and quarterly reports. The forward-looking statements are representative only as of the date they are made and, except as required by applicable law, we assume no responsibility to publicly update or revise any forward-looking statements. We will also discuss both GAAP and non-GAAP financial measures. Non-GAAP measures should not be considered in isolation from, as a substitute for, or superior to our GAAP results. We encourage you to consider all measures when analyzing our performance. A reconciliation of the GAAP to non-GAAP measures is included in today's press release and accompanying slide presentation.

And with that, let me turn the call over to Kash Shaikh, CEO. Kash?

Good afternoon and thank you for joining our third quarter fiscal 2026 earnings call.

Penguin Solutions delivered an exceptional quarter, with record results that reflect disciplined execution and strong customer traction for our AI Factory Platform strategy. At the company level, we delivered record net sales and significantly higher-than-anticipated EPS. Building on our excellent third quarter, we are again raising our full year outlook for both net sales and EPS. All of our business units performed well.

Growth was outstanding in our AI-driven businesses. In Q3, our AI-driven businesses represented 74% of total company net sales and grew 104% year over year. These businesses consist of Integrated Memory and non-hyperscale AI Infrastructure solutions.

In these businesses, AI-driven demand continued to outpace net sales growth, contributing to a growing backlog. This gives us confidence that the AI opportunity is expanding, as enterprises increasingly adopt agentic AI workloads powered by inference at scale.

Our results this quarter demonstrate that Penguin Solutions is evolving into a leading AI Factory Platform company. We believe we are still in the early stages of a significant long-term profitable growth opportunity.

Before discussing our performance in more detail, I want to address our finance leadership transition. As announced on June 1st, Nate Olmstead will step down as Chief Financial Officer on July 8 to pursue an opportunity in a different industry.

Nate has been an important partner during a meaningful period of transformation for Penguin Solutions. His leadership helped strengthen our financial and operational foundation, and I want to thank him for his contributions and wish him well.

I am also pleased that Aaron Johnson, our Vice President of Finance and Accounting, will serve as Interim CFO effective July 9th. Aaron brings more than 16 years of public-company experience in the technology sector and a deep understanding of our business. We have also initiated a search for a permanent CFO. This transition does not change our operating priorities, financial discipline, or focus on execution.

Today, I'll cover four key topics before reviewing our third quarter performance:

- First**, the market environment as AI moves into production-scale inference and agentic AI;
- Second**, how our AI Factory Platform is positioned to address this opportunity;
- Third**, the customer and partner momentum that validates our strategy; and
- Fourth**, how we are executing across product innovation, operations, and go-to-market.

Since our April earnings call, the demand environment has strengthened as AI continues to transition from early prompt-and-response experimentation to production-scale inference and agentic AI. A simple way to think about the shift is that early AI answered questions, while agentic AI performs work.

In the early phase of AI adoption, many workloads were transactional: a user asked a question, AI generated an answer, and the session ended. Agentic AI is different. Agents are persistent, context-rich, and task-oriented. They can operate continuously across workflows, applications, and data sources, moving AI from an adviser to an operator.

For example, in software development, each engineer can now have multiple agents that can write, test, review, document, and monitor code. In business operations, agents can monitor customer activity, analyze supply chain risks, prepare follow-ups, and update workflows.

As inference workloads scale and agentic AI deployments accelerate, infrastructure requirements for production AI environments continue to expand across the full AI data center technology stack. Demand is increasing not only for GPUs and accelerator-attached high-bandwidth memory, or HBM, but also for general-purpose compute, including CPUs, as well as memory, storage, and networking.

Every GPU deployment depends on a surrounding layer of general-purpose compute and memory to feed data, coordinate workflows, manage context, and connect agents to enterprise applications. We believe this trend is strengthening demand for our memory solutions.

These persistent, context-rich workloads require more memory capacity, faster access to context, and better orchestration. As AI moves to inference at scale, the industry is increasingly recognizing that memory, not compute alone, is becoming one of the primary bottlenecks for large-context AI inference performance.

We designed our MemoryAI™ appliances specifically to address this bottleneck and improve both inference performance and token economics.

We believe that for enterprises, sovereign AI initiatives, and neocloud customers, AI factories are becoming essential for optimizing time-to-first-token, performance, and token economics for inference.

These developments reinforce the strategic premise behind Penguin Solutions: production-scale AI infrastructure requires a full-stack AI Factory Platform approach. Our platform connects AI infrastructure, memory, operations software, and operational execution to help customers deploy and operate at scale.

Our second topic today is how Penguin is becoming a leading AI Factory Platform company at the intersection of memory and AI infrastructure.

Our platform combines five core elements with our partner ecosystem:

- 1) ClusterWareAI™, operating system software for AI factories
- 2) MemoryAI™ and Integrated Memory solutions
- 3) Advanced Computing systems under our ComputeAI™ brand,
- 4) OriginAI™ factory architectures, and
- 5) End-to-end design, build, deploy and managed services.

We believe our Platform differentiation comes from combining these proprietary products and services into an integrated offering.

Customers increasingly need more than just AI hardware procurement. They need architecture that performs in production at scale, an accelerated deployment path, and an operating model that helps them generate profitable revenue or deliver operational efficiencies.

In production AI, time to production deployment is directly tied to time to revenue, and our platform is designed to compress both. Customers need the ROI and superior token economics that our Penguin AI Factory Platform can provide.

Turning to our third topic: customer and partner momentum.

With our land-and-expand go-to-market strategy, we continue to add new customer logos and deepen engagements with existing customers.

In Q3'26, we added four new AI Infrastructure customer logos. New logo acquisitions are important in part because they often lead to repeat business. For example, across the 4 trailing quarters from Q3'25 to Q2'26, we added 13 new AI Infrastructure logos, and 7 of those customers have already increased their business with us. New customer engagements can involve longer sales cycles, but they also support deeper customer relationships, repeat business, and more durable long-term growth.

Deepgram is a strong example of the power of our partner model and technical differentiation. Penguin designed and deployed an optimized inference environment built with Dell infrastructure and NVIDIA technology to support enterprise voice AI workloads. In Q3, we expanded our engagement with Deepgram to support additional production capacity as its business scales, and Deepgram also acquired our ClusterWareAI™ product and additional services.

In Q3, we also expanded our previously disclosed engagement with a tier-one financial institution. This expansion includes our Compute Express Link or CXL-powered MemoryAI™ KV Cache Server, ClusterWareAI™ software, and services, with Dell providing AI compute.

Deployments such as Haein in South Korea, which provides local neocloud GPU-as-a-service, continue to demonstrate our ability to support AI infrastructure requirements across a range of customer environments. So does our recent Q3 customer win with a leading quantitative trading firm.

Another important proof point is Spectra, a new sovereign deployment. This supercomputer system was built and deployed through our collaboration with Sandia National Laboratories, and NextSilicon. For Penguin, Spectra demonstrates our ability to design and deploy complex systems in a demanding national security environment.

In our memory business, across the trailing 4 quarters from Q3'25 to Q2'26, we added 16 new logos, and 5 of those customers subsequently increased their business with us.

We also saw continued expansion with a generative AI customer that is purchasing our CXL memory expansion cards to support inference workload solutions.

These relationships validate our “land-and-expand” strategy. Once a customer experiences the value, performance, and ROI of our AI Factory Platform, we have an opportunity to deepen and scale our relationships.

Finally, our partner ecosystem continued to recognize Penguin’s capabilities. We were recently named an NVIDIA AI Factory Specialized Partner, a recognition of our expertise in designing, building, and deploying, and managing full-stack, NVIDIA-based AI factory infrastructure. We were also recognized this quarter as the 2026 Dell Technologies Global Alliances Americas AI Partner of the Year.

Finally let me share some insights into our product, operations and go-to-market execution.

In our product portfolio, we continue to invest in the areas where we see potential for durable differentiation, including our MemoryAI™ products, our ClusterWareAI™ operating system software, and our OriginAI™ reference architectures.

MemoryAI™ is focused on addressing the memory bottleneck that emerges as AI inference workloads become contextually larger, more concurrent, and more latency sensitive.

Our CXL-based MemoryAI™ KV Cache Server is designed to keep large-context inference closer to the workload, allowing large language models to respond faster without needing to re-process entire datasets for every prompt.

MemoryAI™ KV Cache can dramatically improve AI Factory efficiency by enabling superior token economics for inference and agentic AI workloads, delivering:

- 1) Up to 2X higher inference performance;
- 2) Up to 8X lower time-to-first-token latency; and
- 3) Expanded memory capacity beyond GPU HBM by leveraging CXL memory that is approximately 4X–5X more cost-effective than GPU HBM.

The Tier 1 financial services customer we added in Q2 purchased additional MemoryAI™ KV Cache servers in Q3 for their on-prem AI Factory, which is initially focused on inference and agentic AI for code generation using open-weight LLMs. We see AI-assisted code generation as a common use case for inference-focused, on-premise AI factories, where the right architecture can deliver superior token economics.

Our early investment in Celestial AI, a pioneer in Photonic Fabric technology, reflects our long-standing focus on memory architecture innovation. We continue to advance our MemoryAI™ Photonic Memory Appliance, or PMA, through our continued partnership with the Celestial AI team, now part of Marvell. This appliance is designed to extend memory capacity and bandwidth for large-context AI inference environments.

Together, our CXL memory expansion cards, CXL-based MemoryAI™ KV Cache Server appliance and MemoryAI™ PMA appliance reinforce our innovation agenda and long-term vision for AI memory infrastructure.

ClusterWareAI™ addresses another critical need: managing and orchestrating complex AI infrastructure across heterogeneous environments.

ClusterWareAI™ is a hardware vendor-agnostic operating system for AI infrastructure in the data centers. It acts as a unified control plane across GPUs, CPUs, memory, and networking, allowing the infrastructure to operate as a single cluster while reducing manual configuration and supporting faster time to production.

We recently introduced a new AI Factory Operations Agent, which provides data center administrators with a conversational, natural-language interface for operational insights. This is the first in a planned family of agents designed to simplify and automate cluster operations and increase administrator productivity with a human-in-the-loop approach.

OriginAI™ brings these technologies into validated reference architectures that can reduce deployment complexity and risk. Faster deployment can mean faster time to revenue for our customers. We recently introduced OriginAI™ inference solutions.

Together, these solutions are designed to leverage Penguin Solutions' 4 billion hours of GPU runtime experience. They also draw on more than 30 years of expertise delivering advanced memory and AI infrastructure solutions.

As our deployments continue to scale across enterprise, sovereign AI, and neocloud customers, our focus is on increasing repeatability. This includes architectures, software, services, and operations. Over time, this can strengthen both customer outcomes and platform economics.

Within Penguin, we are applying agentic AI across our own operations with a focus on measurable business outcomes that support our profitable growth plans. This includes AI-assisted software development, often referred to as "vibe coding," and AI-powered productivity tools designed to accelerate our product cycles, improve time to market, and make our supply chain more scalable and responsive.

We are also managing supply chain, working capital, and growth investments with discipline as we scale to meet AI-driven demand.

Turning to our third quarter performance, we delivered record quarterly net sales of \$479 million, up 48% year over year and 40% sequentially. Non-GAAP operating income was a third-quarter record at \$64 million, up 67% from the year-ago quarter, demonstrating the strong operating leverage in our model as we scale.

AI-driven demand continued to drive growth in Integrated Memory and in the non-hyperscale AI infrastructure businesses within Advanced Computing. Together, these two businesses represented 74% of total company net sales and grew 104% year over year.

Advanced Computing net sales totaled \$138 million, representing 29% of the total company, and growing 4% year over year. Demand and customer engagement of our AI Infrastructure business within Advanced Computing remains very strong across enterprise, sovereign AI, and neocloud customers. Our non-hyperscaler AI Infrastructure business is scaling rapidly, with third quarter net sales up 81% year over year.

In Q3, this non-hyperscale AI Infrastructure business represented 58% of total Advanced Computing net sales versus 33% in the third quarter of last year.

The shift to inference at scale with agentic AI, aligns directly with Penguin's AI Factory Platform strategy, and the differentiated role that we play at the intersection of memory and compute infrastructure.

Our MemoryAI™ appliance continues to gain traction. It generated both revenue and new bookings this quarter. The pipeline also continues to strengthen across enterprise, sovereign AI, and neocloud customers.

Memory business net sales were outstanding at \$275 million, up more than 111% year over year supported by both higher volume and pricing. Our strong results reflect a fundamental shift: agentic AI is driving sustained, structural demand for memory, reinforcing our view that this AI-driven demand is more durable than a traditional cyclical memory upturn.

We are anchoring our memory business on the data center market, where demand is driven by agentic AI. We are growing our engagements across networking, hyperscale infrastructure, and enterprise computing customers for their data center products.

We exited the third quarter with a very strong backlog. This reflects robust data center-focused AI demand. This demand continues to outpace net sales growth giving us a strong backlog entering the 4th quarter.

Just as important, we are broadening that growth across a wider set of memory customers. This is creating a more durable customer base. This momentum reinforces the expanding reach of our data center-focused memory solutions. It also highlights the strength of the opportunity ahead.

Our CXL memory expansion cards continue to gain traction, generating both revenue and new bookings this quarter, while our customer pipeline continues to strengthen.

Our **LED** business also performed well in the quarter, with Q3 net sales totaling \$66 million, up 7% year over year.

Our strategic priorities remain centered on memory and AI Infrastructure within Advanced Computing. At the same time, LED continues to be managed with discipline. It generates positive cash flow and continues to execute against its innovation roadmap.

On June 1, we announced an improved financial outlook. At that time, we expected full-year fiscal 2026 net sales and diluted EPS to be at the high end of our previously issued ranges. Today, based on our third quarter results and

supported by strong AI-driven demand, we are further increasing our fiscal 2026 outlook for both net sales and diluted EPS. Nate will provide more details in a moment.

Looking ahead, we intend to continue balancing growth investments with operational discipline. Our priorities are clear.

- **First**, invest in differentiated software, AI memory, and compute infrastructure solutions.
- **Second**, execute with discipline and speed.
- **Third**, add new logos and deepen existing customer relationships.
- **Fourth**, grow our partner ecosystem and diversify our customer base.

Together, these actions are designed to support more consistent and predictable profitable growth.

In closing, AI is entering a production phase. This phase is reshaping both the architecture and economics of data centers. As agentic AI makes inference more persistent and demanding, several capabilities are becoming more essential. These include memory, orchestration, full-stack integration, and operational discipline.

We believe we are strategically positioned at the intersection of memory and AI infrastructure to lead the next phase of AI: inference at scale powering agentic AI workloads.

We are still in the early innings of a very large growth opportunity.

Thank you to our employees, customers, and partners. Our teams are energized by the opportunity ahead, and we appreciate their continued focus on execution and customer delivery.

With that, I will turn the call over to Nate for a closer look at our financial results and outlook.

Thanks, Kash. I will focus my remarks on our non-GAAP results, which are reconciled to GAAP in our earnings release tables and in the investor materials available on our website.

With that, let me now turn to our third quarter performance. In the quarter, both net sales and profits were significantly higher than expected, driven primarily by accelerating AI-driven demand for our memory products and continued adoption of AI infrastructure solutions, which drove a 40% sequential increase in our overall net sales and 42% sequential increase in our non-GAAP operating income. For Q3 FY26, total Penguin Solutions net sales were a record \$479 million, up 48% year over year. Non-GAAP gross margin came in at 28.1%, above our expectations, and down 3.6 percentage points versus Q3 last year. Non-GAAP operating margin was 13.4%, up 1.5 percentage points versus last year, and non-GAAP diluted earnings per share were \$0.84, up 79% year over year and up 62% versus last quarter.

In the third quarter of fiscal 2026, our overall product net sales were \$414 million, representing 87% of total company net sales, and growing 60% versus the prior year. Services net sales totaled \$65 million, or 13% of total net sales, and were down 1% versus the prior year.

Net sales by business segment were as follows:

- In Advanced Computing: Q3 net sales were \$138 million, representing 29% of total company net sales, and grew 4% year over year and 19% sequentially. This sales increase reflects stronger AI Infrastructure sales to enterprise customers, which more than offset reduced sales to hyperscale customers. Within Advanced Computing, our non-hyperscale AI Infrastructure business continues to scale rapidly, with net sales up 81% year over year in the quarter. In addition to accelerating growth in this part of our business, we continue to make good progress on diversifying our net sales to new customer categories. In Q3, the non-hyperscale AI Infrastructure business represented 58% of total Advanced Computing net sales versus 33% in the third quarter of last year. We continue to see strong demand from enterprise, neocloud and sovereign AI customers, and expect these customer categories to represent an increasing share of our Advanced Computing net sales over time.
- In Integrated Memory: Q3 net sales were \$275 million, representing 57% of total company net sales, up 111% year over year and 60% sequentially; and
- In Optimized LED: Q3 net sales were \$66 million, representing 14% of total company net sales, and were up 7% versus the same quarter last year.

Non-GAAP gross margin for Penguin Solutions in the third quarter was 28.1 percent, down 3.6 percentage points year over year and down 3.1 percentage points sequentially, primarily attributable to the ongoing wind down of our Penguin Edge business and a shift in the overall mix of sales across business units. These factors were partially offset by AI-driven demand, which supported favorable pricing in our Integrated Memory business during the quarter.

Non-GAAP operating expenses for the third quarter were \$70 million, up 9% year over year and up 14% versus last quarter. The increase in operating expenses sequentially is due to normal seasonality, increased investments in R&D, including for our ClusterwareAI™ software and MemoryAI™ solutions, and higher variable compensation as a result of our strong net sales and profit performance.

Q3 non-GAAP operating income was \$64 million, a third-quarter record, up 67% year over year and up 42% sequentially. Operating margins were up 1.5 percentage points versus the prior year and 0.2 points sequentially, driven by strong operating leverage.

Non-GAAP diluted earnings per share for the third quarter were \$0.84, up 79% versus Q3 last year and up 62% versus the prior quarter.

Adjusted EBITDA for the third quarter was \$68 million, up 51% year over year, and up 34% versus the prior quarter.

Turning to the balance sheet; for working capital, our net accounts receivable totaled \$704 million, compared to \$293 million a year ago, with the increase primarily driven by significantly higher memory sales volumes and prices. Days sales outstanding remained healthy at 53 days, up from 47 days a year ago and up from 50 days last quarter.

Inventory totaled \$498 million at the end of the third quarter, up from \$184 million a year ago, reflecting increased memory costs, growth in our memory and AI infrastructure businesses, and strategic purchases to maximize supply for anticipated future memory demand. Days of inventory were 42 days, up from 36 days a year ago and down from 51 days last quarter, primarily due to the timing of receipts and shipments.

Accounts Payable were \$736 million at the end of the quarter, up from \$272 million a year ago, due primarily to higher memory costs, growth in our memory and AI infrastructure businesses, and the timing of purchases and payments. Days payable outstanding was 62 days compared to 53 days last year and 63 days last quarter. The year over year and quarter over quarter movements were due to the timing of purchases and payments.

Despite the significant growth in our net working capital, our cash conversion cycle remained within our typical range at 33 days, an improvement of 5 days compared to Q2 and up 3 days versus last year due to the timing of purchases and payments.

Consistent with past practice, days sales outstanding, days payables outstanding, and inventory days are calculated on a gross sales and gross cost of goods sold basis, which were \$1.22 billion and \$1.09 billion, respectively, in the third quarter.

As a reminder, the difference between gross and net sales is primarily related to our memory business's logistics services, which are accounted for on an agent basis, meaning that we only recognize the net profit on logistics services as net sales.

Cash, cash equivalents and short-term investments totaled \$440 million at the end of the third quarter, down \$295 million from Q3 last year and down \$49 million sequentially. The year over year decrease was primarily due to proceeds from the issuance of preferred shares in Q2 of last year offset by debt repayments for our term loan in Q4 of last year. Sequentially, the cash decrease was primarily due to investments in working capital to fund our growth and partially offset by approximately \$40 million received from proceeds from the disposition of our 19% equity investment in Zilia Technologies.

Third quarter cash flows used for operating activities totaled \$75 million, compared to \$97 million provided by operating activities in the prior year quarter. The decrease in cash flow in the quarter versus last year was due primarily to investments in net working capital to support the significant growth in our memory and AI infrastructure businesses.

For those of you tracking capital expenditures and depreciation, capital expenditures were \$3 million in the quarter, and depreciation was \$5 million for the quarter.

Wrapping up our cash flow activities, we spent \$9 million to repurchase approximately 466,000 shares in the third quarter under our stock repurchase program. As of May 29, 2026, an aggregate of \$56 million remained available for the repurchase of our common stock under our current authorizations.

And now turning to our outlook.

Given our solid performance over the first nine months and an improved Q4 outlook for our memory business, we are raising our full-company net sales and non-GAAP diluted EPS outlook for the year, which at the midpoint now calls for 22% net sales growth and \$2.60 of non-GAAP diluted EPS, up from our previous outlook of 12% net sales growth and \$2.15 of non-GAAP diluted EPS. As a reminder, our full year outlook assumes that we will continue to diversify our customer sales mix and does not include any Advanced Computing AI hardware sales to hyperscale customers. And also consistent with our assumptions from last quarter, our FY26 financial outlook reflects the ongoing wind down of our high margin Penguin Edge business. We expect sales from this business to essentially cease by the end of fiscal 2026.

The combined effect of these two assumptions reduces our FY26 growth outlook by approximately 14 percentage points at the total company level year over year, and approximately 30 percentage points within Advanced Computing.

With that said, our full year net sales outlook reflects the following full year growth ranges by segment:

- for **Advanced Computing**, we are improving our full-year outlook compared with our previous outlook and now expect net sales to decline 15% to 20% year over year. As it has previously, this outlook reflects the Penguin Edge and hyperscale hardware sales impacts mentioned earlier.
- for **Memory**, we are increasing our full-year outlook, and we now expect net sales to grow between 90 and 95% year over year, driven by continued AI-related demand and a favorable memory market environment.
- and for **LED**, we are also improving our full year outlook and now expect net sales to decline approximately 5% year over year.

Our non-GAAP gross margin outlook for the full year is now 28.5% plus or minus 0.5 percentage points. We adjusted our gross margin outlook up by half a percentage point to account for favorable memory pricing, which helped our Q3 gross margins. Our Q4 outlook assumes less pricing favorability than we experienced in Q3, and we therefore expect some downward pressure on gross margins as we exit the year.

Our full year expectation for total non-GAAP operating expenses has increased to \$260 million, plus or minus \$5 million.

For the full year FY26, we now expect a non-GAAP diluted share count of approximately 56 million shares, up from our prior outlook, primarily reflecting the expected dilutive impact from our convertible debt as a result of higher share prices. As a result of this dilutive impact, we expect our non-GAAP diluted share count in Q4 FY26 to be approximately 62 million shares.

Our non-GAAP full year diluted earnings per share is now expected to be approximately \$2.60, plus or minus 5 cents.

Our forecasted FY26 non-GAAP tax rate is now 20%, down from 22%, reflecting increased pretax income in jurisdictions with lower tax rates. Note, our Q3 non-GAAP tax rate was 17.3%, which reflects the year-to-date true up of this new, lower full year tax rate.

While we expect to use this normalized non-GAAP tax rate throughout FY26 and beyond, the long-term non-GAAP tax rate may be subject to changes for a variety of reasons, including the rapidly evolving global and US tax environment, significant changes in our geographic earnings mix or changes to our strategy or business operations.

Given the strength of current AI-driven demand trends, we want to give some preliminary color on our initial expectations for net sales and non-GAAP EPS growth in fiscal 2027. Based on current customer signals and our expectations about ongoing AI-driven demand, our preliminary fiscal 2027 view contemplates both total company net sales growth and non-GAAP EPS growth of approximately 30% from the midpoint of our full-year FY26 outlook. The EPS outlook factors in the higher share count beginning in Q4'26 that I mentioned earlier. This is an initial planning view, and we expect to provide a full FY27 outlook on our next earnings call.

Our outlook for fiscal year 2026 and our preliminary view of FY27 are based on the current environment and our current assumptions, including, among other things, assumptions relating to customer demand, the global macroeconomic environment, ongoing supply chain constraints, and supply costs, especially as they relate to our Advanced Computing and Integrated Memory businesses. This includes extended lead-times for certain components that are incorporated into our overall solutions, impacting how quickly we can ramp existing and new customer projects and fulfill customer orders. Our outlook also contemplates the industry-wide higher costs for memory, which may slow customer demand for our products and solutions and may lower our gross margins in our Advanced Computing and Memory businesses.

We believe the combination of accelerating AI-related demand, an expanding enterprise, neocloud and sovereign AI customer base, and our disciplined operating model positions Penguin well for continued profitable growth.

Please refer to the non-GAAP financial information section and the reconciliation of GAAP to non-GAAP measures tables in our earnings release and the investor materials on our website for further details.

With that Operator, we are ready for Q&A.