Comments from CEO Lip-Bu Tan and CFO Dave Zinsner

Intel's chief executive officer and chief financial officer offer comments after the company released its fourth-quarter 2025 earnings.

The following are the prepared remarks provided on Intel's fourth-quarter 2025 earnings conference call held at 2 p.m. PST on January 22, 2026. These remarks include forward-looking statements that are based on the environment as seen by the company as of the time of the call and, as such, are subject to various risks and uncertainties. They also contain references to non-GAAP financial measures that the company believes provide useful information to investors. Refer to the company's earnings release for the fourth-quarter 2025, most recent annual report on Form 10-K and other filings with the SEC for more information on the risk factors that could cause actual results to differ materially from the company's expectations and additional information on non-GAAP financial measures, including reconciliations where appropriate to the corresponding GAAP financial measures.

January 22,2026 – Lip-Bu Tan, chief executive officer of Intel (bio):

2025 was a year of solid progress. Over the last 10 months, we have established the foundation for the new Intel – a more focused and execution-driven company. We simplified our organization and greatly reduced bureaucracy to improve efficiency and accelerate decision-making. We also recruited new leaders from the outside, and empowered key leaders from within. We strengthened our balance sheet, forged strong new partnerships, and deepened relationships with existing, as well as new customers.

I am encouraged by my conversations with our customers and partners around the world. I'm hearing a clear and consistent message: they see the progress we're making, and they want Intel at the table as they navigate their own transformations.

The opportunity in front of us is meaningful and significant. The era of artificial intelligence is driving unprecedented demand for semiconductors across the entire compute landscape from AI-accelerated and traditional datacenters, into the network and enterprise domains, all the way out to client and edge devices.

Rapid deployment of AI workloads across this diverse environment will require heterogenous silicon solutions leveraging CPUs, embedded-NPUs, discrete and integrated GPUs, ASICs, and XPUs. In addition, we will need to see innovations in the software stack along with new emerging technologies like photonics, memory interfaces, interconnect and quantum to name just a few.

The breadth of our IP and know-how across silicon design, system-level integration, wafer manufacturing, and advanced packaging uniquely positions us to capitalize on these AI-driven trends and capture sustainable profitable growth.

This will not happen overnight, and our execution needs to continue to improve. But we will stay humble as we address the work ahead, and we will never be satisfied.

Our Q4 was another positive step forward. Revenue, gross margin and EPS were all above our guidance. We delivered these results despite supply constraints which meaningfully limited our ability to capture all of the strength in our underlying markets. We are working aggressively to address this and better support our customers' needs going forward.

Looking ahead for 2026, we will continue to position Intel to capture the significant growth opportunity AI presents across all our businesses. We will do this by strengthening our client franchise, advancing our datacenter, AI accelerator and ASICs strategies, and continuing to build a trusted US Foundry.

Let me start with our core X86 franchise, which remains the most widely deployed compute architecture in the world. The deployment of AI is only amplifying the importance of X86 – from orchestration and control planes to inference, edge workloads and agentic AI.

In our client computing group, we strengthened our position in both consumer and enterprise notebooks with our Core Ultra Series 3 line-up – formerly known as Panther Lake and built on our most advanced Intel 18A manufacturing process.

We committed to delivering our first Series 3 SKU by end of 2025, and we exceeded that commitment by delivering our first three SKUs. While we still have work to do, I am encouraged by the steady progress on Intel 18A yields, and Naga and his team remain laser focused on additional improvements as they ramp Series 3 into the high volume needed to meet strong customer demand.

Our client momentum was on full display at CES earlier this month, where we formally launched Series 3 with our OEM partners. Powering over 200 notebook designs, Series 3 will be the most broadly adopted and globally available AI PC platform we have ever delivered.

Along with our next generation Nova Lake, coming at the end of 2026, we now have a client roadmap that combines best-in-class performance with cost optimized solutions, giving me confidence that we are on a path to fortify market share and profitability, in both notebooks and desktops, over the next several years.

In addition, the PC is becoming an important part of the AI infrastructure. The surge in AI workloads is driving massive demand for data centers, but cloud capacity alone cannot meet the scale of inference needed especially in a power constrained environment.

This is accelerating the push towards hybrid-AI – splitting workloads between cloud and client – which offers clear advantages in performance, cost and control. We are working closely with ecosystem partners to seamlessly enable hybrid-AI and we are encouraged by the opportunity to grow the installed base and accelerate the refresh rates over time.

Let me now turn to DCAI. To support our AI objectives, I believe that our traditional server and accelerator roadmaps must advance together. To reinforce this alignment, I centralized our Data Center and AI businesses under Kevork, ensuring tight coordination across CPUs, GPUs and platform strategy.

Demand for traditional servers continues to be very strong and we are focused on ramping available capacity to support the meaningful uptick we are seeing – including partnering with key customers to support their needs beyond 2026.

The continuing proliferation and diversification of AI workloads is placing significant capacity constraints on traditional and new hardware infrastructure, reinforcing the growing and essential role CPUs play in the AI era. This is, and will continue to benefit the ongoing ramp of Granite Rapids as well as our mainstream products, Sapphire and Emerald Rapids.

We have also made decisive changes to simplify our server roadmap, focusing resources on the 16-channel Diamond Rapids and areas to accelerate the introduction of Coral Rapids where we can. With Coral Rapids, we will also re-introduce multi-threading back into our datacenter roadmap. We also continue to work closely with Nvidia to build a custom XEON fully integrated with their NVLink technology to bring best-in-class x86 performance to AI host nodes.

Over the last several quarters, we have been developing a broader AI and accelerator strategy that we plan to refine in the coming months. This will include innovative options to integrate our x86 CPUs with fixed-function and programmable accelerator IP. Our focus is on the emerging wave of AI workloads – reasoning models, agentic and physical AI, and inference at scale – where we believe Intel can truly disrupt and differentiate.

Our long-term ambition is clear: To rebuild Intel as the compute platform of choice for the next era of AI driven computing, grounded in world class engineering, an accelerated roadmap, and a renewed culture of execution.

We're also building momentum in ASICs as customers seek purpose-built silicon for AI, networking, and cloud workloads. Our combination of design services, IP building blocks, and manufacturing capabilities positions Intel well to solve specialized problems at scale. This is not a new area for us, although it's one that I'm committing significantly more focus, resources and investment dollars, including leveraging my own experience at Cadence Design supporting and growing this market.

Finally, we remain focused on the long-term objective of building a world-class wafer and advanced packaging foundry anchored in trust, consistency, and execution. As I have said before, building a foundry business will take time and considerable effort and resources. While still early in our journey, we have hit some early important milestones worth highlighting.

We are now shipping our first products built on Intel 18A – the most advanced semiconductor process developed and manufactured on U.S. soil. As stated earlier, yields continue to improve steadily as we work to ramp the supply needed to meet strong customer demand.

In addition, Intel 18-AP continues to progress well, and we are engaging with internal and external customers on this node, delivering our 1.0 PDK at the end of last year.

Intel 14A development remains on track and we have taken meaningful steps to simplify our process flow and improve our rate of performance and yield improvement. We are developing a comprehensive IP Portfolio on Intel 14A, and we continue to improve our design enablement approach. Importantly, our PDKs are now viewed by customers as industry standard.

Engagements with potential external customers on Intel 14A are active, and we believe customers will begin to make firm supplier decisions starting in the second half of this year, and extending into the first half of 2027.

We also have the opportunity to provide strong differentiation in advanced packaging, particularly with EMIB and EMIB-T. We are focusing on improving quality and yield to support customers' desire for ramps beginning in the second half of 2026.

In closing, as I reflect on 2025, I am proud of the resilience and commitment our team has demonstrated. We exit the year with a stronger foundation and a clearer roadmap for 2026 and beyond.

The opportunity ahead is meaningful and significant as AI-driven computing expands all the markets we serve. But I am also mindful of the challenges ahead of us and transparent about areas that we are doing well and areas we need to improve.

In the short-term, I am disappointed that we are not able to fully meet the demand in our markets. My team and I are working tirelessly to drive efficiency and more output from our fabs, and while yields are in-line with our internal plans, they are still below where I want them to be. Accelerating yield improvement will be an important lever in 2026 as we look to better support our customers.

As I said earlier, we are on a multi-year journey – it will take time and resolve, but my team and I are committed to rebuilding this iconic American company and increasing the long-term value for our shareholders. I would like to thank my team for their hard work over the course of the last 10 months.

I look forward to updating you on our progress as we continue this journey together, including hosting an investor day in the second half of this year at our headquarters in Santa Clara.

Dave Zinsner, chief financial officer of Intel (bio):

We remain encouraged by the fundamental drivers of demand across our core markets. Fourth quarter revenue was $13.7 billion, at the high end of the range we provided in October. We experienced strong growth across all our businesses benefiting from the AI infrastructure build-out – with AI PC, traditional server and networking revenue all up double digits sequentially and year over year.

Q4 marks the fifth consecutive quarter of revenue above our guidance even as we navigate industry-wide supply constraints for our key products. Non-GAAP gross margin came in at 37.9 percent, approximately 140 basis points ahead of guidance on higher revenue and lower inventory reserves partially offset by increased mix of outsourced client products and the early ramp of Intel 18A to support the launch of Core Ultra Series 3 code named Panther Lake.

We delivered fourth quarter non-GAAP earnings per share of $0.15 cents versus our guidance of $0.08 cents driven by higher revenue, stronger gross margins and continued spending discipline.

Q4 operating cash flow was $4.3 billion with gross capex of $4 billion in the quarter and positive adjusted free cash flow of $2.2 billion. Nvidia's $5 billion investment closed in Q4 as expected.

For the full year, revenue was $52.9 billion down slightly year over year due to constraints across our own manufacturing network and with external suppliers which limited growth, especially in the second half. Full year non-GAAP gross margin was 36.7 percent up 70 basis points on reduced period charges.

Full year non-GAAP EPS was $0.42 cents up $0.55 cents year over year on lower period charges and improved operating leverage. Specifically, non-GAAP OpEx of $16.5 billion was down 15% vs 2024 as we executed actions to reduce complexity and bureaucracy in the business and drive improved execution.

For the full year, we generated $9.7 billion in cash from operations and made $17.7 billion of gross capital investments with capital offsets of approximately $6.5 billion. Although Adjusted Free Cash flow was minus $1.6 billion in 2025, we produced $3.1 billion in the second half as cash from operations more than doubled half on half. We exit 2025 with $37.4 billion of cash and short-term investments bolstered by further monetization of Mobileye, the completion of our stake sale of

Altera to Silverlake, accelerated funding from the US Government, and investments by the Softbank Group and Nvidia. In addition, we repaid $3.7 billion of debt.

Looking back, 2025 marked an important year of progress against our key priorities, even as we know we have more work ahead. Internally, we re-organized and right sized the team to become customer centric and engineering focused while shoring up our balance sheet to give us more flexibility to pursue our goals. We've navigated a market that has shifted from tariff driven uncertainty in the first half to an intense AI driven demand environment constrained by supply in the second half. 2025 demonstrated the staying power of the x86 ecosystem across client and datacenter, and the importance of our manufacturing assets as we launched Core Ultra Series 3 on Intel 18A – the most advanced process fully developed and manufactured in the United States. Both create a firm foundation on which to build the new Intel.

Moving to segment results.

Intel Products' Q4 revenue was $12.9 billion, up 2 percent sequentially. CCG revenue was down 4 percent quarter over quarter even as AI PC units grew 16%, and DCAI was up 15% reflecting strong demand for traditional server compute. These results reflect our efforts to balance our constrained supply with strong datacenter demand while maintaining support for our client OEM partners. Where possible we' re prioritizing our internal wafer supply to datacenter and leveraging an increased mix of externally sourced wafers in client.

CCG revenue was $8.2 billion and in line with our expectations. We estimate the client consumption TAM was greater than 290 million units in 2025 marking two straight years of growth off the post-Covid bottom in 2023 and the fastest TAM growth since 2021.

Within the quarter, CCG launched 3 SKUs of Series 3, ahead of our expectations of one. Performance reviews have been extremely favorable with up to 27 hours of battery-life, a 70% gen-on-gen improvement in graphics, and performance on industry standard benchmarks that is 50 to 100% better than peers.

DCAI revenue was $4.7 billion, up 15 percent sequentially, above expectations and the fastest sequential growth this decade. Revenue would have been meaningfully higher if we had more supply.

While the market continues to benefit from more power efficient CPUs stimulating a refresh cycle, all indicators point to the growing and essential role CPUs will play within hyperscale and enterprise AI datacenters as inference driven AI usage expands. The world is shifting from human prompted requests to persistent and recursive commands driven by computer to computer interactions. The CPU's central function coordinating this traffic will drive not only traditional server refresh, but new demand that grows the installed base.

In addition, due to the networking demand for the AI infrastructure buildout, our custom ASIC business grew more than 50% in 2025, 26% sequentially, and reached an annualized revenue run rate greater than $1 billion in Q4. This strength provides our ASIC team a solid base to pursue a $100 billion TAM opportunity.

Operating profit for Intel Products was $3.5 billion, 27 percent of revenue, and down approximately $200 million quarter-over-quarter, on an increased mix of outsourced products and seasonally higher operating expenses.

Intel Foundry delivered revenue of $4.5 billion, up 6.4 percent sequentially, on increased EUV wafer mix. EUV wafer revenue grew from less than one percent of wafer outs in 2023 to greater than 10 percent in 2025. External Foundry revenue was $222 million in the quarter driven by projects with the US government and the de-consolidation of Altera. Intel Foundry operating loss in Q4 was $2.5 billion, and $188 million worse quarter-over-quarter, driven by the early ramp of Intel 18A.

Within the quarter, Intel Foundry met key 18A and 14A milestones. With the official launch of Core Ultra Series 3, Intel Foundry is the only semiconductor manufacturer in the world shipping gate all around transistors with backside power for revenue. These advanced wafers are rolling off our production lines in Oregon and Arizona; here in the United States. Finally, our continued progress on Intel 14A demonstrates our commitment to research and develop the world's most important technology on US soil.

Turning to All Other.

Revenue came in at $574 million and was down 42% sequentially due to the Q3'25 deconsolidation of Altera. The primary components of All Other in Q4 were Mobileye and IMS. Collectively the category delivered an operating loss of $8 million. I'm pleased with the early momentum at Altera as an independent company with a new leadership team. Their industry-leading programmable fabric, developer productivity-driven software tools, and a large installed base positions them well to drive long-term value creation.

Now turning to guidance. During the second half of 2025 we supported strong demand for our products with intra-quarter wafer production and inventory on hand. As we enter 2026, our buffer inventory is depleted, and the mix shift in wafers towards servers, which began in Q3, will not come out of fab until late Q1'26. As a result, and as we stated last quarter, our internal supply constraints are most acute in Q1.

In light of these dynamics, we are forecasting a Q1 revenue range of $11.7 to $12.7 billion. The midpoint of $12.2 billion reflects a lower end of seasonal Q1. Within Intel Products, we forecast a more pronounced revenue decline in CCG than in DCAI as we continue to prioritize internal supply to our server end markets. We expect Intel Foundry revenue up double digits quarter over quarter helped by continued mix shift to EUV wafers and Intel 18A pricing.

At the midpoint of $12.2 billion, we forecast a gross margin of approximately 34.5 percent with a tax rate of 11 percent and break-even EPS, all on a non-GAAP basis. Gross margin is down sequentially due to lower revenue, increased 18A volumes, and product mix.

Let me take a few moments to provide some color for your full year 2026 model.

First, from a revenue perspective, we expect our factory network to improve available supply beginning in Q2 and for each of the remaining quarters in 2026. Within the server market, customer feedback and our own market intelligence point to the likelihood of a strong year of growth for DCAI. Finally, client CPU inventory is lean and there is excitement for Series 3. In contrast, over the last several months, industry-wide supply for key components like DRAM, NAND and substrates has come under increasing pressure due to intense demand to support the rapid expansion of AI infrastructure. Rising component pricing is a dynamic we continue to watch closely, especially relative to the client market, and could limit our revenue opportunity this year.

For OpEx, we target 2026 operating expenses of $16 billion. We expect noncontrolling interest, or NCI, to net to approximately $325 million in Q1 and be approximately $1.2 billion for the year, on a

GAAP basis. NCI is expected to grow meaningfully again in fiscal 2027. Our share count is forecast to be 5.1 billion shares in Q1 and grow in line with our stock-based compensation going forward.

As we think about our capital expenditures for 2026, we're working to balance our ability to drive capital efficiencies with our need to respond to the demand signals we're receiving. Previously, we said CapEx would be down but are now planning for a range of flat to down slightly, and for expenditures to be more weighted to the first half. As a reminder, CapEx in 2026 would be to support demand in 2027 and beyond. We expect to generate positive adjusted free cash flow for the full year and we are planning to retire all $2.5 billion of maturities as they come due this year.

I'll wrap up by saying that Q4 was another solid quarter to mark our 5th consecutive quarter of overdelivering to our guide. We exit 2025 increasingly confident in the long-term sustainability of the end markets we serve. We believe our improved balance sheet, thanks in part to the trust of our strategic partners, combined with the strong talent we have will enable us to meaningfully participate in the next wave of computing as the industry pushes for returns on their AI investments. I look forward to providing you, our shareholders, an update on what this future means to you at our Analyst Day later this year.

Closing – Lip-Bu Tan, chief executive officer:

Thank you again for joining us today. As we move forward, we remain focused on disciplined execution and deep collaboration with our customers to seize the meaningful opportunities created by the AI era. While we have a lot more to do, we are confident in the foundation we've built and the progress underway. We look forward to providing another update in April.

About Intel

Intel (Nasdaq: INTC) designs and manufactures advanced semiconductors that connect and power the modern world. Every day, our engineers create new technologies that enhance and shape the future of computing to enable new possibilities for every customer we serve. Learn more at intel.com.