

Hewlett Packard Enterprise

Fourth Quarter Fiscal 2024 Earnings Conference Call

Thursday, December 05, 2024, 5:00 PM
Eastern

CORPORATE PARTICIPANTS

Antonio Neri - *President, Chief Executive Officer*

Marie Myers - *Chief Financial Officer*

Paul Glaser - *Head of Investor Relations*

PRESENTATION

Operator

Good afternoon and welcome to the Fourth Quarter Fiscal 2024 Hewlett Packard Enterprise Earnings Conference Call. At this time, all participants will be in a listen-only mode. We will be facilitating a question and answer session towards the end of the conference. Should you need assistance during the call, please signal a conference specialist by pressing the "*" key followed by "0." As a reminder, this conference is being recorded for replay purposes.

I would now like to turn the presentation over to your host for today's call, Paul Glaser, Head of Investor Relations. Please go ahead, sir.

Paul Glaser

Good afternoon. I'm Paul Glaser, Head of Investor Relations for Hewlett Packard Enterprise. I would like to welcome you to our fiscal 2024 fourth quarter earnings conference call, with Antonio Neri, HPE's President and Chief Executive Officer and Marie Myers, HPE's Chief Financial Officer.

Before handing the call to Antonio, let me remind you that this call is being webcast. A replay of the webcast will be available shortly after the call concludes. We have posted the press release and the slide presentation accompanying the release on our HPE investor relations web page. Elements of the financial information referenced on this call are forward-looking and are based on our best view of the world and our businesses as we see them today. HPE assumes no obligation and does not intend to update any such forward-looking statements.

We also note that the financial information discussed on this call reflects estimates based on information available at this time and could differ materially from the amounts ultimately reported in HPE's quarterly report on Form 10-K for the fiscal year ended October 31st, 2024. For more detailed information, please see the disclaimers on the earnings materials relating to forward-looking statements that involve risks, uncertainties, and assumptions. Please refer to HPE's filings with the SEC for a discussion of these risks.

For financial information we have expressed on a non-GAAP basis, we have provided reconciliations to the comparable GAAP information on our website. Please refer to the tables and slide presentation accompanying today's earnings release on our website for details. Throughout this conference call, all revenue growth rates, unless noted otherwise, are presented on a year-over-year basis and adjusted to exclude the impact of currency. Finally, Antonio and Marie will reference our earnings presentation in their prepared comments.

With that, let me turn it over to Antonio.

Antonio Neri

Thank you, Paul, and thank you all for joining us today. Before I share my comments on our quarterly and yearly performance, I would like to provide an update on the Juniper Networks acquisition.

We have received approvals from key geographies around the world, including the European Union, UK, India, South Korea and Australia, which have given unconditional clearance and see the competitive benefits of the transaction. In the United States, we are engaged with the Department of Justice, and we expect the review will continue into the new calendar year. We anticipate the deal will close in the early part of 2025, within the previously stated timeframe.

Both HPE and Juniper continue to believe the transaction will enable us to provide a complete portfolio of modern, secure networking solutions that offer essential foundations for both Hybrid Cloud and AI. The combined companies' enhanced ability to compete and innovate will fundamentally improve the networking industry. This transaction will also strengthen US National Security interests by advancing HPE's position as a strong US Innovator among global technology companies.

Turning to results, we had an exceptional fourth quarter to cap off a strong year. In actual dollar terms, we generate a record quarterly revenue of \$8.5 billion, up 15% year-over-year, exceeding \$8 billion for the first time in our history. Our results reflect the continued adoption of HPE GreenLake and the acceleration of our revenue growth in AI.

Q4 and fiscal year 2024 results exceeded our expectations for revenue, dilute and net earnings per share and free cash flow. We draw sequential revenue growth across each of our segments, a continuation of the positive momentum we saw build throughout the year across our business.

Q4 constituted our third consecutive quarter of improved year-over-year growth rates and revenue, non-GAAP operating profit, and non-GAAP diluted net earnings per share. Contributing significantly to our performance was record server revenue of \$4.7 billion, up 31% year-over-year, a third consecutive quarter of double-digit growth.

Sequentially, AI systems revenue was up 16% to \$1.5 billion, exiting fiscal year 2024 with \$6.7 billion in cumulative orders since Q1 of fiscal year 2023. After the close of Q4, I am pleased to share that our AI systems demand remains strong. We have received new orders bringing AI systems backlog to over \$3.5 billion.

Our traditional server business also continues to reflect the improvement market dynamics, as we draw our third consecutive quarter of year-over-year double-digit orders growth.

Hybrid cloud revenue grew 18% year-over-year and 21% sequentially, resulting in revenue of \$1.6 billion, with increasing revenue for HPE Private Cloud and continued growth in HPE Storage.

We are seeing customers like Lupin Limited, a large India pharmaceutical company, repatriate workloads from the public cloud to enhance their data security amid regulatory and compliance requirements.

Customer adoption of our HPE Alletra Storage MP solutions continues to rise at an accelerated pace. Since launch, we have sold approximately 3,000 systems. Miami-Dade County, the seventh most populous county in the United States, implemented HPE Alletra Storage MP to consolidate an existing storage footprint to lower TCO and improve performance and sustainability.

Our Intelligent Edge business achieved its third consecutive quarter of orders growth, reflecting our expectations for an improving demand environment. We are particularly pleased with the double-digit year-over-year orders growth that we saw in data center networking, an important growth market for our HPE Aruba Networking business.

HPE financial services financing volumes of \$2.1 billion rose to an all-time high this quarter, driven by strong demand for HPE GreenLake and AI Financing. It is clear customers need more investment capacity whether to deploy AI or to accelerate the adoption of Hybrid Cloud.

Our record-breaking revenue performance in Q4 resulted in record profitability with non-GAAP operating profit of \$938 million. That translated to non-GAAP operating margin of more than 11%, a 110 basis points increase sequentially.

As a result of our exceptional Q4 performance, fiscal year 2024 revenue growth topped 3% year-over-year, exceeding \$30 billion as we added more than 9,000 new customers across the portfolio. Combined with our discipline operating expense management, which accelerated in the back-half of the year, we generated a record-breaking \$2.3 billion in free cash flow, exceeding our guidance of at least \$1.9 billion for the year.

We closed the year on a strong note, exceeding our full year commitments for revenue, diluted net earnings per share, and free cash flow. We are very pleased with our results and excited about the opportunities ahead for HPE.

It is clear our strategy is working. We have a tremendous set of assets, including extensive solutions from edge to cloud, enabling us to meet our customer needs across AI, networking, and hybrid cloud. Providing a full stack of product and services clearly differentiates us versus our competition and accelerates value for our shareholders.

In Q4, we continue our investment in innovation, which further enhances our product portfolio and positions us to capture additional market share. AI is clearly transformative for our customers, and HPE is positioned to lead in the next wave of innovation. At our recent AI day, we announced the industry-first, 100% fanless direct liquid cooling architecture, a critical capability for customers to leverage the next generation of direct liquid cooled GPUs and CPUs, while facilitating peak performance as they manage increased energy densities.

This industry-first architecture is unique to HPE and results from decades of experience in liquid cooling technologies. Two weeks ago, at SC24, we announced our newest 400 gigabits HPE Slingshot direct liquid cool networking fabric of products, engineered to support both generative AI cluster workloads and exascale supercomputing.

These innovations are enhancing the performance of our AI systems and supercomputing product lines, differentiating us from the competition. At SC24, we also announced that HPE has delivered the fastest verified supercomputer in the world to the US Department of Energy's Lawrence Livermore National Laboratory. El Capitan is not only the world's fastest supercomputer at more than 1.7 exaflop, it is also among the most energy efficient in the world.

HPE now has deployed seven of the top 10 fastest supercomputers in the world, including numbers one, two, and three on the list. We also continue to innovate to help accelerate enterprise adoption of AI. Since launching HPE Private Cloud AI, we have expanded our collaboration with Deloitte to help businesses of all sizes to deploy AI solutions tailored for their industry-specific use cases such as C-suite AI and many others.

In addition, we introduced the Unleash AI Partner Program to grow our ecosystem and expand customers' ability to address their AI use cases with HPE Private Cloud AI. While we are still in early days, interest is strong with hundreds of customers in the pipeline, with many compelling proof-of-concept underway.

We closed our first deals during Q4, including RWE, a leading company in the field of renewable energy, which selected HPE Private Cloud AI to unlock insights from weather data. Customers are also asking us to help them simplify their VMware private clouds and optimize their virtualization costs.

At HPE Discover Barcelona two weeks ago, we launched HPE VM Essentials, which enables customers to manage their virtualization states across HPE VMware and many others. We have built on our recent acquisitions of Morpheus Data and OpsRamp, giving customers operational flexibility and allowing them to no longer be locked into a single virtualization vendor.

With HPE VM Essentials, they can achieve up to five times lower TCO. As a part of our HPE VM Essentials, we also provide an enterprise-grade open-source KVM virtualization solution. This allows customers to dramatically reduce costs, while increasing their flexibility to innovate across a full suite of hybrid cloud operations capabilities in an on-premises multi-stack and a multi-cloud public environment.

In storage, we continue to invest in our multi-protocol HPE Alletra Storage MP platform as we transition our storage portfolio to an AI-driven, cloud-native, and disaggregated architecture. Our new HPE Alletra Storage MP X 10000 is a revolutionary new high-performance object storage designed for AI exabyte scale and optimized for high-speed unstructured data lakes with rapid restore for backup and recovery.

Our AWS S3 compatible object storage interface and deduplication, supports up to 20 times data reduction and streamlines integration with any available backup solution in the market. With HPE Alletra Storage MP X 10000, customers can achieve up to six times faster object storage performance compared to competitor solutions in the market.

Finally, our Innovation Intelligent Edge are helping us to advance HPE Aruba Networking Central with AIOps copilot capabilities built into our HPE GreenLake cloud. During the quarter, we announced two new solutions to enhance the networking operators experience and security for our customers.

The first is the integration with HPE OpsRamp to monitor third-party devices. We also introduced a behavioral analytical-based network detection and response capability, where we are leveraging network and telemetry to train AI models to monitor customer IoT devices. We have an important place in the market, as a leading technology innovator. I am proud that HPE team members around the globe bring our innovations to customers every day to help them modernize their IT infrastructure and transform for the better.

I am very pleased with our exceptional fourth quarter results and the resulting outperformance in several of our full year commitments to our shareholders. Our strategy is well aligned to the megatrends reshaping networking, hybrid cloud, and AI, and our unique product portfolio is well-positioned to capture new customers and accelerate value for our shareholders. We expect the pending acquisition of Juniper Networks to further enhance our portfolio providing customers with complete edge-to-cloud solutions.

I am extremely excited about the significant opportunity we have in the coming fiscal year to drive increased value for our shareholders. I am optimistic about what we can achieve and look forward to the year ahead.

Now let me turn the call over to Marie, who will provide more details about our Q4 results and guide for the first quarter. Marie.

Marie Myers

Thank you, Antonio and good afternoon. We ended fiscal 2024 on solid footing and exceeded what we said we would do. We achieved record quarterly revenue, delivered solid profitability, and generated free cash flow above our guidance. We executed well in a dynamic environment, growing revenue sequentially in each segment as investments in AI are leading to infrastructure refreshes.

In Q4, we also maintained disciplined cost management and improved non-GAAP operating profit. A few highlights from our 2024 results. We grew AI systems revenue more than 150% to \$4.1 billion and met our server non-GAAP operating margin target. We continued our product transition within storage and ended the year with Alletra MP accounting for a meaningful portion of total storage orders. Our transition to more software defined storage will drive a greater mix of higher-margin recurring revenue over the long-term.

Lastly, in Intelligent Edge, we rightsized our cost structure as we navigated a digestion period. We grew fiscal 2024 revenue 3% to \$30.1 billion exceeding the high-end of the outlook range we provided in September.

Non-GAAP diluted net EPS of \$1.99 and free cash flow of \$2.3 billion both came in above our revised outlook ranges. In total we returned \$826 million to shareholders via dividends and share repurchases. Overall, our results reflect focused execution and improving demand as customers implement AI strategies.

Let me dive into the details of the quarter. We delivered record quarterly revenue and our second highest non-GAAP diluted net EPS, both of which came in above the high end of guidance. Total revenue grew 15% year-over-year and 9% quarter-over-quarter to a record \$8.5 billion. Growth in the quarter was led by strong conversion of AI systems backlog, refreshes in traditional compute, better than expected performance within Hybrid Cloud, and a continued recovery in networking.

We are pleased to report non-GAAP diluted net EPS of \$0.58 above the high-end of guidance despite dilution from the issuance of our convertible preferred securities in September which was not factored into our outlook. ARR grew 48% year-over-year to \$1.9 billion, nearly doubling since Q1 2023 primarily driven by HPE GreenLake. We added approximately 2,000 new HPE GreenLake customers during the quarter ending the year with approximately 39,000 unique customers.

I am happy to report continued demand for our differentiated AI system offerings, which resulted in strong double-digit sequential revenue growth to \$1.5 billion. This record contribution from AI systems revenue, combined with lower mix of Intelligent Edge, drove gross margin down 390 basis points year-over-year and 90 basis points quarter-over-quarter to 30.9%.

Non-GAAP operating expenses decreased 9% year-over-year, as reported and were down \$5 million quarter-over-quarter, as we focus on streamlining our cost structure and are closely managing discretionary expenses. Non-GAAP operating margin was 11.1%, up 140 basis points year-over-year and up 110 basis points sequentially. Strong profitability rolled through to free cash flow, which exceeded our expectations and totaled \$1.5 billion. Fiscal 2024 free cash flow was \$2.3 billion our highest ever for the full year.

GAAP diluted net EPS of \$0.99 and non-GAAP diluted net EPS of \$0.58 were both above guidance, benefiting from better than expected OI&E and expense management. GAAP EPS benefited primarily from a lower than previously expected tax expense on a aggregated H3C dispositions. Non-GAAP diluted net EPS excludes the gain, we recognized on the partial sale of our H3C investment at \$162 million in net costs, primarily from stock-based compensation expense, acquisition and other related charges and amortization of intangibles.

Now let's turn to the segment results. Our server business was a key driver of record quarterly revenue and grew double-digits year-over-year for the third consecutive quarter. Server revenue achieved an all-time high of \$4.7 billion, up 31% year-over-year and up 9% quarter-over-quarter with sequential growth in AI systems and traditional servers.

The traditional compute business continued its momentum during the quarter and grew sequentially for the fourth consecutive quarter, driven by ongoing refreshes to our Gen11 server products, which carry higher AUPs. Adoption has been strongest in North America and Europe, as customers are investing in new workloads. In the fourth quarter, Gen11 accounted for more than two-thirds of our core compute revenue and contributed to AUP growth.

In AI systems, we continue to grow the business to new highs. AI systems orders during the quarter were in-line with our expectations at approximately \$1.2 billion. However, we had an order de-book in Q4, leaving our net orders for the quarter at approximately \$500 million. Subsequent to the end of the quarter, we have received orders that bring our current backlog to over \$3.5 billion. As we have mentioned before, AI systems orders can be lumpy and this is an example of that. Fourth quarter server operating margin was 11.6%, up 150 basis points year-over-year and up 80 basis points quarter-over-quarter.

Moving to Hybrid Cloud, we grew revenue 18% year-over-year and 21% quarter-over-quarter to \$1.6 billion, materially exceeding our guidance of a slight revenue increase. Revenue growth was led by Private Cloud and the continued ramp of Alletra MP. During the quarter, we received our first orders for Private Cloud AI as we target a growing pipeline that includes corporations across verticals and regions.

Interest is strong in manufacturing, education, and financial services. We have also deployed Private Cloud AI internally and are moving use cases from the Public Cloud due to better cost performance and compliance.

In storage, we are balancing investments in owned IP products against market trends. Alletra MP remains ahead of our expectations and is the fastest ramping storage product in our company's history.

Keep in mind, these sales carry a higher portion of deferred software and services revenue, which takes longer to translate to the P&L but benefits margins long-term. Hybrid Cloud operating margin was 7.7%, up 390 basis points year-over-year and up 260 basis points sequentially, predominantly due to better OPEX controls.

Now, on to Intelligent Edge. We believe the business remains at a path of recovery and customers have largely digested excess inventories. Total Intelligent Edge revenue was \$1.1 billion, down 20% year-over-year and essentially flat quarter-over-quarter. Sequentially, we saw growth in services, WLAN products and software, partially offset by declines in switching and campus.

Importantly, we are seeing more large deals in our pipeline, giving us confidence that demand is improving. Keep in mind, we were still benefiting from the drawdown of networking backlog in Q1 of fiscal 2024. Operating margin of 24.4% was down 270 basis points year-over-year but up 200 basis points quarter-over-quarter due to lower OPEX in line with our expectations.

Lastly, Financial Services. Our Financial Services business generated \$893 million of revenue, up 2% year-over-year and up 1% quarter-over-quarter. Financing volumes increased 41% year-over-year to a new all-time high of \$2.1 billion. Our Q4 loss ratio remained steady near 0.5% and return on equity totaled 17%. Operating margin was 9.2%, up 120 basis points year-over-year and up 20 basis points quarter-over-quarter.

Moving to cash flow and capital allocation. We generated strong operating cash flow of \$2 billion and free cash flow of \$1.5 billion in the quarter, both of which exceeded our expectations due to higher revenue and increased collections. For the full year, free cash flow was a record \$2.3 billion, above our guidance of \$1.9 billion. Q4 cash conversion cycle was a negative 12 days, down eight days from Q4 '23 and down 16 days from last quarter.

Inventory ended the year at \$7.8 billion, up 2% quarter-over-quarter due to the nature of our AI systems business. We are focused on reducing inventory during fiscal 2025 as we convert AI systems backlog and grow our storage business. During Q4, we returned \$169 million via dividends and \$50 million via share repurchases to common shareholders, respectively. We returned a total of \$826 million during the full year.

Moving to our outlook for fiscal Q1 2025. As Antonio said, we expect to close the Juniper transaction in early 2025 at which time we will provide combined company guidance for fiscal 2025.

However, let me provide some high-level thoughts, as to how we are thinking about the upcoming year. Overall customer conversations indicate higher IT spending in 2025 with multiple tailwinds that should contribute to revenue growth. We expect continued recovery in traditional compute and growing adoption of AI systems by enterprises and sovereigns, although we expect orders to remain competitive and lumpy.

In Hybrid Cloud, we are beginning to see customers accelerate digital transformation projects in order to execute on AI strategies. And in networking, we expect demand will modestly recover throughout the year. For the first quarter we expect year-over-year revenue growth to be in the mid-teens.

On a sequential basis, this is in-line with normal seasonality. We expect server revenue to be down quarter-over-quarter. We expect flat to modest growth sequentially in traditional compute and lower AI systems contribution following a very strong quarter. Server operating margin will be closer to 10% to 11%, as customers navigate the transition to NextGen GPUs.

In Hybrid Cloud, we expect a sequential decline in the first quarter due to stronger than seasonal growth in Q4, with operating margin in the mid-single-digit range. And in networking, we are still managing prolonged sales cycles and expect the business to be around historical seasonality of flattish sequential growth. We expect Intelligent Edge operating margin to be in the low 20% range due to mix.

Gross margin will benefit from a more favorable revenue mix, and we expect a modest sequential decrease in OPEX. OI&E is expected to positively contribute to the bottom-line with higher net interest income more than offsetting \$0.02 of headwinds associated with the cost of our Juniper Financing.

Overall, we estimate Q1 GAAP diluted net EPS to be between \$0.31 and \$0.36 and non-GAAP diluted net EPS to be between \$0.47 and \$0.52 based on 1.4 billion diluted weighted average shares outstanding, including \$76 million of dilution from our convertible preferred securities.

For free cash flow, we typically consume cash in the first quarter. We intend to procure components needed to meet current and future AI systems demand, while continuing to manage inventory. For Q1, we intend to buy back shares of common stock at a similar pace to Q4. We remain committed in the long-term to our balanced capital allocation framework, our dividend, and our investment-grade rating.

In summary, we delivered strong Q4 results and ended fiscal 2025, well-positioned to drive profitable growth and generate free cash flow that we can reinvest into the business while also returning capital to shareholders.

With that, I'll turn it over for Q&A.

QUESTION AND ANSWER

Operator

Thank you. We will now begin the question and answer session. To ask a question, you may press "*" then "1" on your touchtone phone. If you are using a speakerphone, please pick-up your handset before pressing the keys. To withdraw your question, please press "*" then "2". We also request that you ask only one question. We will now pause for just a moment to assemble our roster.

And your first question today will come from Wamsi Mohan with Bank of America. Please go ahead.

Wamsi Mohan

Hi, yes. Thank you so much and congrats on the strong quarter and guide, especially given the dilution from the preferred stock. As we look at your AI business, can you maybe give us some sense of how your AI pipeline is shaping up, given that the backlog levels kind of dipped in the quarter? And I know you made some comments about how it's backed up intra-quarter. But can you give us some sense of how the pipeline looks and what the composition of that could be? Thank you.

Antonio Neri

Yes. Thanks, Wamsi. This is Antonio. And good afternoon. The pipeline is a multiple of the current backlog which we shared as of a day ago, which is basically more than \$3.5 billion. And inside our pipeline, obviously remember there is four key segments of the market. Obviously, the biggest driver continues to be what I call segment one or model builders, hyperscalers as well as now the tier 2, tier 3 that they are in the business to provide GPU capacity on demand on a per hour basis. So they continue to be the two biggest driver. But what I'm really pleased, and Marie and I have made those comments in our remarks, is the fact that we see now the enterprise AI pipeline continue to grow steadily.

There are hundreds of customers on that pipeline. The size of the deal will vary obviously because it's not the same size of the deal of the other two segments. But there are some very interesting large-size deals in the pipeline, and they are doing proof-of-concept. But the positive news is that, remember that we made the announcement of the offer that co-engineered solution with NVIDIA in June, and we made it available only on September 09.

So really, we had seven weeks to be able to collect orders. We already collected orders in enterprise AI and we already closed deals in enterprise AI. And two weeks ago, we made additional enhancement to that offer, particularly with a partner ecosystem, including Deloitte. So very pleased. This is going to be a growth driver in 2025, as we go forward. And then on the first two segment obviously, the transition, I would say, to the latest GPUs that will happen over the course of the next few quarters.

Wamsi Mohan

Very good. Thank you.

Paul Glaser

Thank you, Wamsi. Next question please.

Operator

And your next question today will come from Meta Marshall with Morgan Stanley. Please go ahead.

Meta Marshall

Great, thanks and congrats on the quarter. Maybe diving into gross margins. Can you just talk about some of the puts and takes on maybe particularly server margin this quarter and just kind of the weight of some of the XD servers versus El Capitan recognition or just some of the dynamics that play on gross margins would be helpful? Thank you.

Marie Myers

Hi, Meta, good afternoon. It's Marie. And first of all, I'll just add up by saying, as you know, we typically don't get into gross margin at the segment level, but let me give you some color around how we think about our op-margins. First of all, I'd say look, from a service segment perspective, really pleased with the performance in the quarter. We were at 11.6%. And if you think about it, we actually had our highest ever AI server revenue of \$1.5 billion. So I think you can see just the performance there.

In terms of the color and how to think about it, El Capitan actually was included in the quarter, so we did actually get the final acceptances. A couple of points to note. If you look at the mix inside server, we see traditional compute particularly Gen11, we're about two-thirds of the mix so far to-date, and that particularly has higher configurations. And as a result, it is driving richer AUPs and obviously higher structural margins. And also, frankly, we've been really successful about passing over commodity costs. And if you look at both the GPU servers, it is competitive but you can see the revenue scale is helping us out there. And then finally, I think Meta we've been really prudent with our cost controls and all of that's playing off in terms of our OPEX discipline. So that's how we're thinking about server margins.

Antonio Neri

But in addition to what Marie said Meta, is this was the third consecutive quarter of double-digit year-over-year orders growth in the traditional server. And when I think about the opportunity ahead of us, there are just for HPE, 0.5 million units in the installed base that must be refreshed

for several reasons. Number one, they are doing very little work and consuming a lot of energy. Number two is because they have reaching or they are reaching the depreciation cycle. And when enterprise customers think about deployment of AI, they need to free up space, power and cooling to absorb this new system. So this is a big opportunity for us, both from a units perspective and obviously Gen11 comes with a different configuration, as Marie said, and then the services attach that comes with it.

Paul Glaser

Thank you, Meta. Next question, please.

Operator

And your next question today will come from Tim Long with Barclays. Please go ahead.

Tim Long

Thank you. I was hoping I could ask a two-parter on the AI service. First, can you just clarify that de-book comment that came in? It sounds like \$700 million de-book of backlog. Was that a competitive loss or someone had double ordered or something else that was going on there? And then the second part, following on the gross margin? Can you talk a little bit about when you're looking at AI business, what are the metrics you are looking at? Are you...do you have thresholds for gross margin or cash flow or something else that prohibits you from going at certain deals where maybe ODMs or lower gross margin OEMs would go? Thank you.

Antonio Neri

Sure, and thank you for the question. As we said before, orders during the quarter were in-line with our expectations. In fact, we've booked \$1.2 billion of new orders in Q4. However, we de-booked a large order, which was \$700 million during the same quarter, because we had concern with a specific customer. And we proud ourselves of a very strong control of this environment, and we continue to be very vigilant about how we engage and making sure we manage the risk. And we felt that order carry a lot of risk, and therefore we felt it was very prudent to de-book it and focus on other areas of the business. And that's why as we keep going throughout the next few weeks, up to yesterday, we ended up with a backlog now, which is more over the \$3.5 billion. So that's what happened. It has nothing to do with competitive. It has nothing to do with everything else. We had a concern with a specific customer, and we decided it was prudent to de-book it.

Marie Myers

Yes. And Tim, just to add on how we think about margin structure and AI systems. I think we've said this in the past, but look we have a framework that we use for our pricing, and it's absolutely our intention to remain disciplined both on cost and price as we pursue profitable growth. I think that has been our mantra. But obviously then, we're out there balancing those gross margin pressures with cost control, strategic buys, and you've seen that in terms of our pricing discipline. And I would add...and I think Antonio made this comment earlier that we do see over time, improving profitability as the market moves more to, I think, the two sectors that Antonio referred to, both in terms of sovereign and enterprise where we see services attach, et cetera. So just use that as a way to think about the longer-term sort of margin structure as well.

Paul Glaser

Thank you, Tim. Next question, please.

Operator

Your next question today will come from Amit Daryanani with Evercore. Please go ahead.

Amit Daryanani

Good afternoon. Thanks for taking my question. I guess, if I just look at the corporate level, you've had a couple of quarters with gross margins have been down a fair bit year-over-year. I think it was down like 390 basis points this quarter? Could you just talk about like, a), what's driving this gross margin degradation for you folks? But it clearly looks like you're able to offset all of this with some really good OPEX control. As we just talk about what's driving the gross margin decline? And then how low can you take OPEX at this point, given your expectations for revenue growth in the next fiscal year? Thank you.

Marie Myers

Yes, no worries, Amit. Good afternoon. It's Marie. So why don't I take that one? On gross margin, so for Q4, as you rightly said, it was 30.9% so we're down both on a year-on-year and Q-on-Q basis. If you step back and look at the overall revenue mix, perhaps one of the largest drivers is really just the lower contribution of the Intelligent Edge revenue. And in fact, in Q4, Amit, it was actually the lowest for the entire year. And obviously, we had a higher contribution inside of servers of AI systems revenue and product mix. So that obviously impacted gross margin. But what I would say is despite that headwind, we actually grew gross profit dollars both year-on-year and Q-on-Q.

Now going forward, I think I said in my prepared remarks that we do actually expect to see gross margin will be up modestly in terms of the improved business mix and pre-dominantly, that's driven by the Intelligent Edge mix. But from a longer-term perspective, just two things to keep in mind. First of all obviously, when we closed Juniper that is going to be a big driver, and has a very positive impact on both gross and operating. We expect that 50% of operating profit going forward will come from Juniper.

And secondly as I mentioned earlier, I think when I answered Tim's question, we do expect, particularly in the AI space, those higher margin opportunities as we move more into seeing the adoption of sovereign and enterprise. On the cost side of the house, really pleased with the work we've done so far to-date. We've been really focused on discretionary costs throughout '24. We're moving more into areas of productivity. Amit, as we look at '25, we are actually using GenAI inside the company, using it in finance to help drive productivity. So we see opportunities like that throughout the company. I would just say at a company as large as HPE, there's plenty of opportunity, Amit.

Paul Glaser

Thank you, Amit. Next question please.

Operator

And your next question today will come from Simon Leopold with Raymond James. Please go ahead.

Simon Leopold

Thank you for taking the questions. I think in your commentary on the outlook, you indicated that traditional servers would be flattish to even slightly up sequentially in the January quarter, which I think defies seasonality. So I just want to verify if I got that correct. Maybe if you could unpack what's happening in that part of the business. Thank you.

Marie Myers

Yes. So this is Marie here, Simon. So yes, you did get that correct in terms of servers. And I think I did mention in my prepared remarks that we did see a lower AI mix in terms of server systems in Q1. And I think as we've said in the past, this business is lumpy so you just bear that in mind as you think about it.

In terms of traditional servers and the outlook on the business really nothing's changed. We continue to see the momentum in that business. I think Antonio described it very well earlier when he mentioned two-thirds of the mix is already at Gen11 in traditional servers, and we see that continuing. So don't expect that to change fundamentally. But just remember, it's a pretty competitive environment out there, both...for both traditional servers and AI systems right now in the market. So that's the way I'd sort of think about it. So that's, why we felt it was prudent, given the current backlog we've got to guide 10% to 11% right now.

Antonio Neri

Also, Simon obviously, in addition to the traditional yearly seasonality, now we have gone back to the normal seasonality we used to have a few years ago when just for the COVID and the recovery that happened after that. On the AI side, one-third...more than one-third of our backlog is Blackwell-related. And so, you need to consider that the trajectory of that more than \$3.5 billion being Blackwell-related. You need time to convert that into revenue. And that's why Marie mentioned of the lumpiness of revenue conversion sometime or the orders intake. But notwithstanding, we are very pleased with the revenue we had in Q4. We are confident about the continued growth in 2025, as we transition the GPUs to this new type of silicon.

Paul Glaser

Thank you, Simon. Next question please.

Operator

Your next question today will come from Aaron Rakers with Wells Fargo. Please go ahead.

Aaron Rakers

Yes. Thanks for taking the question. Antonio, kind of building off of that last comment. I'm curious as we think about the Blackwell platform cycle ahead of us, and your comment on the backlog build that you're seeing. Can you just update us where...I think a lot of that initial demand seems to be centered around the GB200 NVL, that racks scale system. Where does HPE stand on that platform shipping those system configurations? When do you think that those would ship in volume for HPE? Thank you.

Antonio Neri

Yes. I think you're right, Aaron. I think again, in the segmentation, I'd like to operate within the segment one and two, call it, service provider, hyperscalers, and model builders. They tend to lead with the latest technology, and that transition is going to happen over '25. And so obviously, we are working with the customers and NVIDIA as they become available, and we will be shipping those in volume in 2025. And so, when you think about calendar year right, I will say, in the first half of 2025 is when you start seeing growth in that platform.

In the sovereign space, I don't think that's the case. I think it's a mix of multiple configurations, because they are different type of workloads. And then in enterprise I argue, it will be an N minus 1 or 2, because the enterprise is less about Blackwell or Grace Blackwell. It's more about the simplicity of the deployment and the time to value. And that's why when you look at today, our HPE Private Cloud AI is all now based on H100. And so...and that's plenty for what they

do, whether it's dragging a model or fine-tuning a model or doing inferencing. So that's why you can generalize things here, but it is fair to say people in the business to build and train models will lead with the latest technology and they will transition in the first half of 2025. And then the enterprise will stay one or two generations behind because obviously, they'll need all that computing power.

Paul Glaser

Thank you, Aaron. Next question please.

Operator

And your next question today will come from Toni Sacconaghi with Bernstein. Please go ahead.

Toni Sacconaghi

Hi yes. Thank you. I'm not sure if you said it but perhaps you can just confirm, did your AI pipeline grow sequentially? And then on the Hybrid Cloud side, it was a huge jump sequentially, dramatically higher than your guidance as well and the biggest sequential increase we've ever seen. Was there any unusually large deal in there? I know you closed Morpheus also, but I don't think that was really material. Can you add any color about whether there was anything unusual? Did you pull forward deals and hence you think next quarter is going to be seasonally weaker than usual? But again, there was an enormous differentiation relative to normal history in your own guidance. Thank you.

Antonio Neri

Yes, Marie, you want to take that?

Marie Myers

Yes. I want to just address Hybrid Cloud, Toni. Look, first of all, I'd start out and say, look, we're really pleased with the performance of Hybrid Cloud in the quarter. And as you mentioned, revenue was up almost 18% year-on-year and 22% quarter-on-quarter, so really strong growth. It was really driven by both Private Cloud and frankly just the ongoing adoption of HPE Alletra MP Storage. So I think both of those were just great drivers of contribution. I think you see in the deck...on the investor deck, we had 39,000 customers on GreenLake. And our MP Storage, I think, is one of the fastest adoptions ever in the history of the company. So I think you're starting to see those benefits pass-through.

And in fact, AAR was up 48% as well. So overall, really good on the top-line, and frankly, pleased with the trajectory from a margin perspective as well, as you rightly said. We're continuing to scale the business, manage our cost structure all of that's paying off. And the Q1 guide actually does continue some of the momentum on the margin. We expect that to be relatively sequentially flattish. And given the strong Q4, we do expect to have a modest sort of decrease in revenue which is more in-line with the sequential on Hybrid Cloud. So overall I think very pleased and want that trajectory to continue, Toni.

I'll turn it over to Antonio now.

Antonio Neri

Just to add a couple of things there, Toni, and before I answer the AI question. So first of all, we now see the result of many, many quarters of hard work in Hybrid Cloud. Remember, Hybrid Cloud is led by the GreenLake platform plus storage plus all the ancillary services. And Alletra MP Storage today is the fastest-ever follow ramp in the history of the company, already on a \$1 billion annualized basis. And remember, that platform actually could be higher. But

remember, a portion of that revenue gets deferred because it's a subscription of the software to the infrastructure. And that's why I'm so excited with the focus of the platform, including the latest introduction with the Storage MP X 10000 because that addresses the AI needs at scale, particularly exabyte scale for unstructured data.

The second thing is I'm really proud, and by the way, on Alletra MP, we sold already 3,000 new systems. On the GreenLake obviously, the platform has 39,000 customers. Throughout the year, we added more than 9,000 new logos to HPE. And now with Morpheus Data and OpsRamp, Zerto on the ransomware protection plus the GreenOps, the FinOps, people realize that GreenLake is a true hybrid by design and is core how we deliver infrastructure and the software that comes with it. So that's what's driving the growth in the Hybrid Cloud space. And more to come there because also Aruba has a huge component of subscription through the GreenLake platform.

And then on the AI, yes, the pipeline grew quarter-over-quarter. And as I said, it's a multiple of the current backlog at more than \$3.5 billion, but also the composition of the pipeline keeps shifting. And my...what I'm really pleased is that the enterprise side has grown very rapidly on the back of the offer that now it's available for ordering for customers.

Paul Glaser

Thank you, Toni. Next question please.

Operator

And your next question today will come from Samik Chatterjee with JP Morgan. Please go ahead.

Samik Chatterjee

Hi, thanks for taking my questions. If I can change gears here a bit on Intelligent Edge. You talked about the strong order patterns. But when I look at the revenue performance in the quarter, I know you were up modestly, but it did seem like overall from at least relative to some other peers that have reported, the sequential growth was a bit more limited relative to sort of quarter-over-quarter increases here. So I mean, is there...how are you thinking about sort of the next year playing out? And when you look at order trends, are they coming in as expected or are you seeing a bit more weakness than what you would have expected maybe 90 days ago? And within that portfolio, how you are thinking about where...as you said, enterprises are going to spend more relative to their IT budgets? What are you hearing from them in terms of priorities within that portfolio? Thank you.

Antonio Neri

Absolutely. And thank you for the question. So as I said in my remarks, we had the third consecutive order growth in Intelligent Edge. But you don't see it fully translating into revenue, which looks like more flattish. In fact, that's the case for even Q1 is because the growth areas that accelerated are very much subscription based.

So this is SD-WAN with SSE, called the SASE framework, had a very nice quarter. The second is the subscription services to HPE GreenLake Aruba Central, which is very important in the way you manage that infrastructure. But we are now seeing is also customers placing orders for future quarters related to Wi-Fi 7. And as we look at that, what I'm really excited is that when you upgrade your access points at some point, you also need to upgrade your campus switching.

And the reason why it is necessary because the access points are powered over Ethernet from the switching port. But today, I will say, the one area has not yet fully returned to what we want to see is the campus switching itself. And then last but not least, data center networking using more intelligent features that our CNX platform brings to table has had double-digit growth again. And so that's why I'm excited about the combination with Juniper Networks because bring complementarity to that portfolio.

And when I look at the AI, I see three unique opportunities. Number one, the interconnectivity within data center to data center. That requires basically a big router. And obviously Juniper really has a leading platform with a 100-gig. Once you are inside the data center, you need a high performance fabric. You need to lay the pipes inside the data center to be able to connect rack-to-rack and row-to-row. And then inside the rack itself, you need a high-performance connectivity.

This is why the HPE Slingshot 400-gig that we just announced two weeks ago has played a massive role, especially as customers will consider other type of silicon in that environment. And you need an open architecture and that's where the Slingshot brings on. But I believe the market is on recovery, three data points already, three quarters of growth. And this is what I believe in the context what Marie just said early on, on the question of margins and contribution of each of the businesses. I argued networking is the biggest upside HPE will have in '25 and '26.

Paul Glaser

Thank you, Samik. Next question, please.

Operator

And your next question today will come from David Vogt with UBS. Please go ahead.

David Vogt

Great, guys. Thank you for the [technical difficulty]. Antonio, can you dig in a little bit into the de-platforming of the orders? You mentioned risk, but can you get into specifics in terms of what was the definition of risk? And how do you think about AI orders going forward? I know it's lumpy but just trying to get a sense for what was particularly risky about that orders. And how should we think about that going forward? Thank you.

Antonio Neri

I mean, listen, we look at the customers from a variety of angles. Obviously, the ability to verify end user, ability to obviously fulfill the payment of the orders and risk associated with allocated components and the like. And we felt that particular order had a lot of risk and we decided not to fulfill it because of our stringent controls. And my view is that this was the right thing to do.

Marie Myers

Yes. And I'd just add, look, HPE has a very strong control environment. As we've said before, orders and backlog are subject to the booking push-outs and cancellations prior to RevRec [ph]. So, we maintain those tight controls.

Paul Glaser

Okay, thank you. Next question please.

Operator

Your next question today will come from Matt Niknam with Deutsche Bank. Please go ahead.

Matt Niknam

Hi, thanks so much for taking the question. You sounded fairly constructive around IT spend, particularly heading into '25. So, I'm just wondering what's changed if anything, over the last three months mainly in terms of macro backdrop. And I guess, more specifically, have you seen any pickup or change in behavior post-election? Thanks.

Antonio Neri

Well, first of all Matt, welcome. Welcome to the community. Pleased to have you here. So, the macro environment versus a year ago has improved, no question whatsoever. Obviously they are unique areas where the customers are paying more attention. Focus on AI clearly is very high on the agenda. Focus on implementing a more flexible hybrid model because AI at the core is a hybrid workload and a data-intensive workload. Definitely, that's another big area.

And both of those spaces, HPE has the right to play and win with our unique focus and innovation. I think obviously, cyber continues to be very important. But then also by geography, maybe varies a little bit. But as I said earlier, in server, we are on the cycle for modernization and improve efficiency of footprint cost per workload. I talked about the TCO per workload and as well as the power consumption. A lot of focus on sustainability. This was very obvious to me, Matt, when I was two weeks ago in Barcelona. We had the biggest conference ever in Europe with more than 4,500 customers all very keen to understand our solutions in the context of sustainability.

And then one of the areas, again, I will reinforce what I said earlier networking needs to catch up to the demands of AI. And that's why we believe that the combination of HPE and Juniper is so important to drive that next wave of innovation. But networking clearly is going to be a topic for 2025 and beyond. So that's what we see. When I speak to customers, which I do more than 50% of the time the AI is pretty much focused on the enterprise side. It is very much focused about business productivity and time to value. And even in our case right, we have many used cases, which many are already deployed in production, which allows us to do some of the things that Marie talked early on in efficiency and cost reduction. But we're still early and that's the excitement, which means the opportunities are yet to come at much bigger scale. And if I want to wrap my commentary about why we are delivering the performance we have delivered in Q4, I will use only one word, innovation.

Our strategy is an innovation-driven strategy that makes us more relevant, that eventually customers buy more from us, that increases value for shareholders. So for me, it's a testament that the strategy is working. And then as I look at these mega trends and the customer needs, we are uniquely positioned even with the pending Juniper acquisition to address those needs.

Paul Glaser

Okay. Thank you. Operator, one more question, please.

Operator

Of course. And your final question today will come from Asiya Merchant with Citigroup. Please go ahead.

Asiya Merchant

Great. Thank you for squeezing me in. The enterprise services on the AI side of things, it looks like you continue to book services there on your orders. Just if you can double click on that, relative to where you were last quarter, what kind of services? Is that a growing percentage of

your orders revenues? And where is HPE winning as you look at those services pipeline as well? Thank you.

Antonio Neri

Thank you for the question. Our services orders related to AI went up quarter-over-quarter. So...and the type of services customers are looking for is you can use two different ways to think about it, the day zero, day one, and day two. The day zero is where you help customers define the used cases, define the return on investment and the like. The day one is where you design and implement, and the day two is to run it. And HPE is uniquely positioned in all these three days, if you will. And therefore, we see quite a bit of engagement upfront. But also the partnerships with companies like Deloitte, and I mentioned that are so important to reach customers, to bring more a business process-oriented approach is essential. They also focus on additional AI practices. And we have many other conversations, by the way. You will hear more about that in the next few weeks, few announcements there.

And then obviously, on design and implementing and running, HPE has all the capabilities that we need. But on the designing and implementing, this unique approach with the Private Cloud AI has proven to be the right value proposition because many of these enterprises don't have the time, don't have the expertise, and they want a partner that can help them do that. And then on the runtime, many of these instances would not land in their own data center, will land into what I call co-locations. And HPE already has three disposition enterprise stacks with our Private Cloud AI, so that they can turn it on when they are ready to deploy those used cases. And that's why we should expect continued improvement on the services side, which remains the source for gross margin as we go forward.

Paul Glaser

Very good. That was our last question. Antonio, if you have...

CONCLUSION

Antonio Neri

Well, we are right on time, and I want to thank everyone for participating today. Again, very pleased with our results in Q4, capping off a very strong year. You can see the momentum of revenue, operating profit dollars. We did what we said. I mean, that's fundamentally the bottom-line. And we had a record number of metrics there. But I'm excited about 2025. I think we have an amazing set of assets that will continue to drive profitable growth and drive shareholder value. And depending Juniper acquisition, as I said in my remarks, we are working through the process. We are confident this will close in the early part of 2025. And we are working very, very collaborative with the DOJ. So, nothing gives me pause that would not happen. So, if I don't speak to you, happy holidays to you and your family, and we'll reconvene in early 2025.

Operator

The conference has now concluded. Thank you for attending today's presentation. You may now disconnect.