Hewlett Packard Enterprise

First Quarter Fiscal 2024 Earnings Conference Call

Thursday, February 29, 2024, 5:30 PM ET

CORPORATE PARTICIPANTS

Antonio Neri - *President, Chief Executive Officer*

Marie Myers - *Chief Financial Officer*

Shannon Cross - *Senior Vice President, Chief Strategy Officer, Investor Relations*

Presentation

Operator
Good afternoon, and welcome to the First Quarter Fiscal 2024 Hewlett Packard Enterprise Earnings Conference Call. My name is Gary, and I'll be your conference moderator for today's call. At this time, all participants will be in 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, Ms. Shannon Cross, Senior Vice President and Chief Strategy Officer and Investor Relations. Please proceed.

Shannon Cross
Good afternoon. I'd like to welcome you to our fiscal 2024 first quarter earnings conference call with Antonio Neri, HPE's President and Chief Executive Officer and Marie Myers, HPE's Chief Financial Officer.

After 25 years on Wall Street and over 20 years covering HPE, I'm very excited to join HPE as Chief Strategy Officer. I look forward to working with Jeff Kvaal and the rest of the IR team, and I look forward to seeing many of you in the months ahead.

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-Q for the fiscal quarter ended January 31st, 2024.

For a 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 otherwise noted, 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 and their prepared comments.

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.

With that, let me turn it over to Antonio.

Antonio Neri
Thank you, Shannon. Good afternoon and thank you for joining us today. In the first quarter, we are proud to have outpaced our profitability expectations while advancing our long-term strategy. We also continue to scale our recurring revenue, achieving the second highest year-over-year growth rate since we started tracking ARR in late 2019. This is a promising indicator for our ongoing portfolio shift to higher margin revenues. But overall, Q1 revenue performance did not meet our expectations.

During this call, I will address three key points. First, I will touch on our revenue, which was lower than expected in large part because network and demand softened industry-wide and because the timing of several large GPU acceptances shifted. Additionally, we did not have the GPU supply we wanted, curtailing our revenue upside. Second, I will address how we are streamlining our reporting segments, accelerating a new specialized sales model, managing our spending, and reinforcing execution discipline. Third and most importantly, I will discuss the progress we're making in executing our long-term strategy, which we remain confident in.

Let me first address revenue in the quarter. Similar to peers in the market, we saw campus networking product demand weakened and the decline later in the quarter was greater than expected. This was a large headwind relative to our expectations. Customers are taking longer to digest prior orders than we had anticipated which partially offset the benefit of our backlog entering the quarter. And Europe and Asia were areas of relative softness.

We expect weakness in the networking market to persist, which is likely to impact revenue through fiscal year 2024. That said, we anticipate some improvement late in fiscal year 2024 as inventory clears and we ramp into the purchasing season for state and local education customers in the United States.

AI server demand remains very strong, evidenced by our growing cumulative order book. However, GPU availability remains tight, and our delivery timing has also been affected by the increasing length of time, customers require to set up the data center space, power and cooling requirements needed to run these systems. As a result, overall AI server orders conversion was below our expectations. However, the AI contribution to ARR is increasing.

We continue to prioritize profitability. I am pleased that we have delivered non-GAAP gross margin of 36.2%, which is up 200 basis points year-over-year and above 600 basis points from fiscal year 2018. This performance helped our Q1 non-GAAP diluted net earnings per share grow to $0.48, which was above the midpoint of our guidance range despite lower-than-expected revenue, illustrating the positive impact of our pivot to higher growth, higher margin revenue.

We know that the current environment will require continued discipline in how we execute. We are accelerating new specialized sales motions to maximize opportunities and improve order linearity across our portfolio. And further cost management will remain an important competency for us in fiscal year 2024.

We also found opportunities to streamline our reporting segments. We have now combined the Compute and HPC and AI segments into a single Server segment that integrates general-purpose computing, high-performance computing, supercomputing, and AI systems. This will enable us to maximize the opportunities across the entire AI life cycle, from training to tuning to inferencing and execute with agility. And as previously discussed, we have simplified our hybrid cloud strategy by putting all-related products, software, and services into one business unit.

Our new hybrid cloud segment will further accelerate customer adoption of the HPE GreenLake hybrid cloud platform.

Turning to our strategy, while we are experiencing cyclicality in submarkets, I am more confident than ever in our long-term strategy that is aligned to key market mega trends. In edge, over the last several quarters, we have gained share in the campus networking markets, and our strategic investments have paid off. Most recently, we have seen strong growth in SASE, an offering bolstered by our acquisition of Silver Peak in 2020 and Axis Security in 2023. Our sales pipeline for our private 5G offering is also growing rapidly, following our acquisition of Athonet in 2023.

In hybrid cloud, HPE GreenLake continues to resonate in the marketplace and was the primary driver of the highest Q1 year-over-year rise in ARR over the four plus years we have been reporting it. Our ARR grew 41% year-over-year to more than $1.4 billion in Q1, and we continue to expect ARR growth of 35% to 45% as we look ahead.

We're also capitalizing on cross-selling opportunities when customers come to us for our AI solutions and realize we can meet their storage needs as well. In AI, we are capturing the explosion in demand for AI systems. Our cumulative accelerator processing unit orders rose to $4 billion in the quarter, driven by demand across HPE Cray EX and XD solutions, as well as HPE ProLiant Gen11 AI-optimized servers.

HPE orders represent nearly 25% of our total server orders since the first quarter of fiscal 2023. Our pipeline is large and growing across the entire AI life cycle from training to tuning to inferencing. We are starting to see AI demand pull-through for other solutions, including storage.

We expect our server and hybrid cloud segments to grow sequentially through the fiscal year. Server revenue stands to benefit from AI system demand, improving GPU supply, and our continued mix shift to HPE ProLiant Gen11. Hybrid cloud will benefit from continued HPE GreenLake storage demand and the rising productivity of our specialized sales force.

Our customers continue to validate our value proposition. As one example, we are building for Eni, one of the world's largest energy providers, a new HPE Cray EX supercomputer that will reach more than half an exaflop performance. The system will be one of the most powerful in the world for enterprise use and will accelerate AI-driven scientific discovery to advance efforts in energy.

We also have been awarded the deal for Poland's most powerful supercomputing system located at the Academic Computer Centre CYFRONET of the AGH University of Science and Technology. Based on the HPE Cray EX supercomputer and HPE Slingshot interconnect fabric with NVIDIA Grace Hopper GPUs, the system will be used to support modeling simulation and AI driven scientific research needs, including training and tuning of large language models. We're also seeing growth in AI inferencing. For example, Coles Supermarkets, a leading Australian retailer, implemented an HPE ProLiant Gen11 AI inferencing solution in Q1, which helps with video imaging to reduce store stock lost to theft and incorrect scanning at the checkout.

In the quarter, we expanded our strategic collaboration with NVIDIA targeting the enterprise segment of the market. We introduced a preconfigured solution for enterprise customers to fine tune AI large language models with the private data to accelerate inferencing. Our HPE machine learning development environment software and unique HPE supercomputer and IP are critical parts of the solution, alongside NVIDIA AI Enterprise software. We have built a strong and growing sales pipeline for this new offering.

We also continue to build out our HPE GreenLake hybrid cloud platform services. Earlier this quarter, we announced an expanded HPE GreenLake for File Storage that is designed for generative AI. This solution is highly differentiated through a high-performance file system solution specifically designed for AI applications. We believe it better positions us to take market share in storage by addressing the previously underserved segment of the file market.

Innovation like this continues to attract customers to the HPE GreenLake platform, which connects 3.8 million network devices and supports more than 31,000 customer organizations, up approximately 8% from last quarter. One new HPE GreenLake customer is the US Navy Fleet Numerical Meteorology and Oceanography Center, which produces critical models of weather and ocean conditions for the US and coalition forces worldwide. They turned to HPE GreenLake to improve predictability, accuracy and speed of their modeling while reducing costs.

This week, I attended Mobile World Congress where AI and private 5G were the key topics among Telco's and service providers alike. With our pending acquisition of Juniper Networks, HPE's portfolio will expand to better serve these unique customers from the edge of the network to core 5G to cloud.

Customers were very interested in our integration of Athonet private 5G capabilities into our Intelligent Edge portfolio, as well as, in our cloud Open RAN and vRAN solutions. They were also very eager to explore the massive new market opportunity AI inferencing presents at the edge of the network. That is one of the reasons why we're so excited about our pending Juniper Networks acquisition.

Combining our complementary portfolios will supercharge HPE's edge-to-cloud strategy, accelerating our entire portfolio with AI-enabled innovation. When our proposed acquisition closes, we will create a new networking innovator with our comprehensive portfolio for customers and partners. The transaction is expected to double the size of our networking business, which will be the core foundation of HPE, covering the anticipated $180 billion market opportunity with our combined IP.

From a financial perspective, this transaction is also compelling for our shareholders. In the first year post-close, we expect accretion to non-GAAP earnings per share and in the long term, higher non-GAAP gross and operating margins. We are working to secure regulatory approvals in several jurisdictions. We are hopeful that regulators will recognize that this acquisition is centered around driving further innovation for our customers. We continue to expect that the transaction will close later this calendar year or in early calendar 2025.

In summary, Q1 2024 was a mixed quarter for HPE. We achieved strong profitability and drove new record ARR growth with overall revenue short-term expectations given the softening networking market, GPU deal timing and to some extent, GPU availability.

We are focused on execution as we navigate the fluctuations in demand we see in certain areas of the market. Marie will take you through our adjusted guidance, which reflects our latest thinking about the year ahead.

This quarter is a moment in time and does not at all dampen our confidence in the future ahead of us. We have taken the right actions to maximize value for our shareholders. The work we are doing now, combined with our technological edge and our strategy that has never been more relevant, will position us to convert on the long-term opportunities in front of us across edge, hybrid cloud, and AI.

Before we transition, I'm delighted to welcome Marie Myers as our new CFO. Having worked with her at HP before the separation, it is a pleasure to partner with her again. I admire her passion for and skill at fueling innovation and performance. I am confident that Marie is a great fit for this role and expect she will help drive the next phase of growth and shareholder return for HPE.

I will now turn the call over to her for details about our segments and our outlook. Marie.

Marie Myers
Thank you, Antonio. I'm pleased to be with you today on my first earnings call as HPE's CFO. I've long admired HPE's impressive transformation, and there has never been a more exciting time to be part of this company.

We have a growing addressable market, a proven strategy and a differentiated portfolio that is levered to long-term market trends around networking, hybrid cloud and AI. I believe we have a significant opportunity ahead of us. I'm excited to partner with Antonio and the rest of the outstanding HPE team to capitalize on this opportunity and drive value for our shareholders. And as Antonio mentioned, we have much to be proud of.

Financial highlights in the quarter included record gross margins and expense discipline, which helped lift non-GAAP EPS to the high end of our guidance range. Demand for our traditional server and storage products has stabilized. Demand for our HPE GreenLake offerings was evident in a healthy ARR growth and demand for our AI systems remains robust.

However, demand in Intelligent Edge did soften due to customer digestion of strong product shipments in fiscal year '23, which is lasting longer than we initially anticipated and is the primary reason Q1 revenue came in below our expectations.

GPU availability and deal timing also contributed. We are taking swift action to address these headwinds by curtailing costs and driving efficiencies across the business.

With that, let's take a closer look at the details of the quarter. Revenue fell 14% year-over-year in constant currency to $6.8 billion. Please recall, we had significant backlog consumption in Q1 '23, particularly in traditional servers and storage. Backlog has now largely normalized across our business with the exception of our APU products.

We have strong momentum in HPE GreenLake. ARR exceeded $1.4 billion in Q1. Storage and networking software and services are the fastest growth elements of ARR. Our software and services mix rose 400 basis points year-over-year to 69%. ARR is the best indicator of our model transformation to our as-a-service offerings. This growth validates what our customers are telling us, that HPE GreenLake is a differentiated value proposition in the market.

Our Q1 non-GAAP gross margin was 36.2%, it rose 200 basis points year-over-year driven by a mix shift to Intelligent Edge and favorable input cost management. We are pleased that our non-GAAP gross margin is up 600 basis points from fiscal year '18 to a company record. This illustrates the ongoing pivot to higher growth, higher margin revenue across our portfolio.

Given the current networking market dynamics, we are taking this moment to streamline and simplify our operations. We are focused on controlling the things we can control. Prudent cost management and disciplined execution are important regardless of the macro environment and are even more critical at times like these.

The benefits of our focus are already evident in our Q1 non-GAAP operating expenses, which declined 4.7% year-over-year at 9.5% sequentially to $1.7 billion. The strong Q1 non-GAAP gross margins and OPEX discipline led to Q1 non-GAAP operating margins of 11.5%, which is down only 30 basis points year-over-year despite less revenue scale. Favorable timing on some corporate expenses was a tailwind to Q1 operating expense and will appear in Q2.

GAAP EPS of $0.29 and non-GAAP EPS of $0.48 exceeded the midpoint of our guidance range. Our diluted share cap was approximately 1.3 billion, and non-GAAP diluted net earnings per share exclude $251 million in costs primarily from stock-based compensation expense, amortization of intangibles and noncash loss on investments.

As we manage the business with focus and discipline, we will also invest to capitalize on the sizable opportunities associated with moving through the interrelated inflection points in networking, hybrid cloud and AI all at the same time. HPE is evolving into a more simple, more agile company that is even better positioned to pursue our growth opportunities and to evolve our mix of products, services and software, and to drive structural higher profitability.

Turning to our segment results. In addition to the new segments we discussed at SAM, we also created a new Server segment that combines our prior compute and HPC and AI segment under one streamlined segment that offers solutions for our customers' training, tuning, and inferencing AI needs across their AI life cycle.

Server revenues were $3.4 billion in the quarter, which was down 23% year-over-year. This new segment had a difficult year-over-year compare as, in Q1 '23, the business made significant progress against its backlog and benefited from HPE Cray EX revenue for Frontier. We are capturing the robust growth in AI demand.

Our cumulative APU orders since Q1 '23, which include APU attached products and services in our HPE Cray EX, XD and ProLiant systems, rose approximately $500 million sequentially to now $4 billion.

Our pipeline is robust and GPU supply, while still constrained, is improving. Our APU product revenue increased sequentially to well over $400 million. We are now converting our APU orders into revenues, and yet very strong demand means our APU backlog is over $3 billion. Our pipeline is well above that.

AI revenue in the quarter included our first revenues from AI cloud offering and from our HPE GreenLake win with a large hyper-scale customer. Our traditional high-performance computing and supercomputing revenue fell seasonally, sequentially following a strong Q4. Revenue from our traditional server business increased sequentially. We expect this trend to continue given a structural mix shift to higher AUP Gen11 and rising input costs. Gen11 servers nearly doubled sequentially to 30% of the mix in Q1.

Including HPE, Cray XD takes the mix of next-generation servers to 44%. We are encouraged that our Gen11 pipeline is starting to include AI inferencing activity and enterprise applications. We expect AI inferencing to gather momentum in fiscal year '24. As a reminder, our Gen11 services come with an attached subscription to our Compute Ops Management software, which lifts our margin structure.

Our Q1 operating margin was 11.4%. While up sequentially, the margin was down 430 basis points year-over-year given declining revenues.

Intelligent Edge revenues were $1.2 billion or up 2% year-over-year. Demand for our campus switching and Wi-Fi products eased materially, particularly in Europe and Asia, and was the largest contributor to our Q1 revenue gap. We continue to see mid-single-digit growth benefits from our existing backlog but expect these to normalize going forward as we are now approaching our typical range.

Our total channel inventory is within the normal range overall, inclusive of a pocket in the SMB business. We also continue to make progress in our new TAMs of data center networking, private 5G and SASE.

The Intelligent Edge portfolio of subscription revenue grew well above 50% as we are benefiting from growing attach from our strong fiscal year '23 revenue growth. We are pleased with our 29.4% operating margin, which was up 1,000 basis points year-over-year.

The new Hybrid Cloud segment includes our storage businesses and now the HPE GreenLake portion of our Server business. Revenues of $1.2 billion were down 10% year-over-year. Our traditional storage business was down year-over-year on difficult compares, given backlog consumption in Q1 '23.

Total Alletra subscription revenue grew over 100% year-over-year and is an illustration of our long-term transition to an as-a-service model across our businesses. We are starting to see AI server demand pull through interest in our file storage portfolio. We are also already seeing some cross-selling benefits of integrating the majority of our HPE GreenLake offering into a single business unit.

Our operating margin was 3.8%, which was down 200 basis points year-over-year. Lower revenues and a high mix of third-party product impacted margins in the traditional storage business.

Our HPE Financial Services revenue was down 2% year-over-year and financing volume was $1.4 billion. Our operating margin of 8.5% was up 130 basis points year-over-year. We are successfully passing through interest rate increases, and our asset management margins are returning to normal. Our Q1 loss ratio remained steady at 0.5%.

Turning now to cash flow and capital allocation. We generated $64 million in cash flow from operations and consumed $482 million in free cash flow this quarter. HPE typically consumes cash in the first half of the year and generates cash in the second half.

Our first quarter free cash flow benefited from some prepayments associated with pending HPE GreenLake deals and HPE Cray XD shipments. Our cash conversion cycle was seven days, which is a reduction of eight days from Q1 '23. Our days payable and days of inventory were both higher, given the lower revenue and strong demand for APUs. We returned $172 million in capital to shareholders in Q1, primarily through our dividend.

Before I discuss our outlook, let me first recap the key drivers that factor into our expectations for Q2 and the full year. For Server, we expect improving GPU supply to drive sequential revenue increases through fiscal year '24. Given improving supply and the timing of installations, we expect segment revenue and therefore, corporate revenue to be heavily weighted to the second half of the year. We expect the blended margin of the new segment to be flattish for the remainder of the fiscal year.

For Intelligent Edge, we expect the market to remain soft throughout the year. Our cost reduction efforts will take time to show their benefits, which will result in margin pressure in Q2.

Later in our fiscal year, we expect our normal channel inventory position and seasonal strength in the education market to be a modest revenue tailwind. We expect the full year margin to be in the mid-20% range. For Hybrid Cloud, we expect sequential increases through the year as our traditional storage business improves and HPE GreenLake momentum continues. We expect meaningful progress through the year.

With that context, let me turn to our outlook. For Q2, we expect revenues in the range of $6.6 billion to $7 billion or up slightly sequentially at the midpoint. We expect GAAP diluted net EPS to be between $0.20 and $0.25, and our non-GAAP diluted net EPS between $0.36 and $0.41. In part, given the higher prepayments in Q1, we expect negative free cash flow in Q2.

Let me remind you that we have two drivers of potential lumpiness in our business. One is the timing and customer acceptances of large Cray wins, including Cray XD wins, which should ramp beginning in Q2. The second is the start of certain large HPE GreenLake deals. For some of these large deals, our segments recognize the related hardware on installation. In our consolidated results, we eliminate the segment revenue and recognize it over time. Both large deals and higher eliminations are indicators of our confidence in the second half of the fiscal year.

For fiscal year '24, we are revising our outlook, primarily given the current networking market headwinds. Let me remind you that we are excluding the H3C earnings and gain on sale from our non-GAAP results beginning in fiscal year '24. We now expect revised ranges for constant currency revenue and non-GAAP operating profit growth of 0% to 2%, GAAP diluted net EPS of $1.81 to $1.91 and non-GAAP diluted net EPS of $1.82 to $1.92. The mix shift from Intelligent Edge to Server should also weigh on our gross margins.

We now expect the full year non-GAAP gross margin to be slightly down from our prior full year expectation of 35%. We expect the impact of the cost actions we have initiated to materialize in the second half and leave fiscal year '24 OPEX to be flat to down from fiscal '23 OPEX. We expect our operating margin to be flattish year-over-year. We now expect OI&E to be $200 million to $250 million headwind versus our prior $300 million expectation, given better Q1 free cash flow and more favorable interest rate assumptions. We expect the effect on currency to be immaterial.

We expect free cash flow to be at least $1.9 billion in fiscal year '24. We expect significantly stronger free cash flow in the second half of the year led by improvements in inventories as AI service shipments ramp. We reiterate our commitment to our dividend, which was raised 8%

from fiscal year '23 to fiscal year '24; debt repayment to maintain an investment-grade credit rating; and in the long term, returning capital to shareholders through share repurchases.

To conclude, we executed well in Q1 amidst a challenging market backdrop. We are pleased with our margin and EPS results while understanding our slower networking product demand and GPU availability and timing impacted our revenue performance.

We have taken prompt action to further reduce our costs and continue to manage our expenses prudently while we advance our long-term strategy. AI server demand is strong. Demand has stabilized for our traditional server and storage products, and our HPE GreenLake momentum is robust. We will continue to invest in IT inflections, networking, hybrid cloud and AI, to drive our pivot to higher growth, higher margin revenue.

I look forward to engaging with you in the months ahead, as does our new Chief Strategy Officer, Shannon Cross, who has joined us following a distinguished career as a Wall Street analyst and now has oversight of our IR function.

We will appreciate your input and questions along the way, and we can get started with that right now. I'll open it up now for your questions about the quarter.

QUESTION AND ANSWER

Operator
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 pickup your handset before pressing the keys. To withdraw your question, please press "*" then "2." We also request that you only ask one question.

The first question today is from Meta Marshall with Morgan Stanley. Please go ahead.

Meta Marshall
Great. Thanks. Maybe on the GPU available…or GPU delays that you're seeing in far as acceptance. Just what are you seeing in terms of…you identified that power and some of these things were conditions for what the delays were, but just what are you seeing in terms of how long those delays are going to take and how long the delays and the acceptances are going to be? Thank you.

Antonio Neri
Well, thank you, Meta. Good afternoon. So as I said in my prepared remarks, we had a couple of deals that slipped from Q1 into future quarters because customers are taking a little bit longer to prepare the data center space, getting the power and the cooling ready. And obviously, those deals will come as we complete those installations.

And then on the GPU side, obviously, we continue to experience a tight environment, although we are seeing some improvements. We have already a lot of GPUs that we already built, but the customers will take time to accept those systems. But the reality is that we need more supply against the backlog that we announced today, which was $3 billion at the end of the Q1. So that's what we see today. And as we go-forward, we expect that improvement to happen, and that's why we are confident on the conversion of the GPU orders into revenue as we go along, not just because of the GPU availability but also the acceptances.

Shannon Cross
Great. Thank you very much Meta. Gary can we have the next question.

Operator
And the next question is from Amit Daryanani with Evercore. Please go ahead.

Amit Daryanani
Thanks for taking my questions. I guess, Antonio, maybe just talk about…if I look at the revenue shortfall in Jan quarter, how much of that do you think is because your customers are pushing out their delivery schedules because there isn't enough power versus you just didn't have enough GPUs? If you were to think about those two buckets. And then as you think about the full year guide, perhaps I didn't appreciate this. But can you just talk about what are you expecting for the networking segment, Intelligent Edge to do in the 0% to 2% guide right now?

Marie Myers
Hi, Amit. Good afternoon. Nice to hear from you. So I will take…I'll answer the first part of your question around the revenue. So in the first quarter with respect to what drove the $300 million on the revenue, was mostly actually networking. We did have one deal that moved out. I think Antonio mentioned that in his prepared remarks. And then in terms of just how we're thinking about the second half of the guide, we are expecting a very strong second half, and that's predominantly driven actually by AI systems revenue. We are expecting networking to be slightly more favorable in the back half, but we really see the trough of networking in Q2, Amit.

Shannon Cross
Great. Gary. Thank you, Amit. Can we have the next question?

Operator
And the next question is from Simon Leopold with Raymond James. Please go ahead.

Simon Leopold
Thanks for taking the questions. I wanted to see if we could drill down a little bit in terms of understanding what's changed in the Intelligent Edge versus 90 days ago. And two elements are crossing my mind. One is, really around the pending Juniper deal, whether that's influencing customers to maybe hold off purchases because of the uncertainties that might be affecting their decision-making as to what happens after you're combined. And the other part is just wondering, if there's inventory that's been sitting in the channel, why didn't you know about it or why didn't you see it. Just trying to get an understanding of sort of what you've learned over these last 90 days. Thank you.

Antonio Neri
Thank you, Simon. So first of all, we saw an acceleration of the demand softness in the back…at the end of the Q1, really in January, whether it's now people coming back, but obviously, the reality is we saw that as a headwind to our revenue in Q1. I have to say, we do not have a channel inventory problem. Actually, we are in great position, particularly from the enterprise customers and the enterprise product, we do not have that issue at this point in time. What we do see is customers are taking longer with the product we already shipped to them to install it and eventually go through the next cycle. And that's why we said with Marie, we're going to start seeing a slight improvement on the back half with Q2 being the trough. And part of the back half also is the traditional buying season in the United States with state and local education. The pipeline is very good. We have not lost one single deal that I can point to,

neither because of the slowdown or customers deferring, not because of the announcement of the acquisition of Juniper.

Marie Myers
And maybe, Simon, just to add to Antonio's comments, the only place where we saw a slightly elevated pocket of inventory was in SMB, which is a pretty small part of our business.

Antonio Neri
Yes.

Shannon Cross
Thank you, Simon. Gary, can we have the next question?

Operator
And the next question is from Toni Sacconaghi with Bernstein. Please go ahead.

Toni Sacconaghi
Yes, thank you. Sorry, I just have one clarification and a question. Marie, on the [Indiscernible] that the backlog dropdown has contributed, I think, mid-single-digits to the Intelligent Edge or something more broad than that? So can you just comment or clarify exactly what the backlog contribution was, and I suppose there's none going forward. And then just on my question…sound pretty excited about sequential growth over the course of the year, both in servers and storage. Maybe you can just elaborate on why you see that. Is servers…do you see sequential growth in traditional servers, non-accelerated [indiscernible]. Thank you.

Marie Myers
Hi, Toni. Yes, good afternoon. So maybe I'll take the first part of the question, and I'll turn to Antonio for the second part. So, look, in terms of the backlog, we really don't disclose the backlog on Edge. But I think what we've said is that we've seen our backlog sort of revert back to normalized levels with the exception, obviously, of our APU or AI system. So that's how we're thinking about the backlog. I think in terms of just some context and commentary, you've seen in the industry that the market has definitely softened. And I think as Antonio said earlier, we saw that late in the quarter. So that's how we characterized the networking demand. And I would say that we do expect the trough in Q2 and to be slightly more favorable in the back half. So that's how we're thinking about the networking market playing out for the year.

And I think I'll turn to Antonio to comment to servers.

Antonio Neri
Yes, Toni. Thank you for the question. So, on the server and server side, obviously, we see signs of stabilization now, has been now a couple of quarters with sequential order improvement. But the reality is that, as we said in our opening remarks, we continue to see the mix shift in the traditional servers, as you call it, to Gen11. By the end of the year, we should be approximately 60% of the way there. Obviously, those services come with a different set of structural configurations and pricing, which obviously is higher. At the same time, we're going to see cost inflation. We believe that's going to be the case, and therefore, we have to eventually pass those as well. But the number of units have been very stable or slightly improving. And that's an important indicator because ultimately, that also drives our attach rate of our operational services, which in the quarter was very, very good. So that's why we are confident that sequential improvement from here on.

And then on the storage side, obviously, AI is going to be a pull-through demand for us. We introduced a new offer now specifically for File. And remember that HP Alletra continues to grow from here on. And a portion of that HPE Alletra revenue is also in the ARR because remember that software now is completely disaggregated from the solution itself, which means you have the CAPEX portion of the revenue recognized in quarter and the subscription part of the software amortized over the period of the contract. So that's why ARR is growing because of the subscription in networking, which was up significantly, the storage and obviously, AI now also contributing to the ARR as well.

Shannon Cross
Thank you, Toni. Gary, can we have the next question?

Operator
The next question is from Aaron Rakers with Wells Fargo. Please go ahead.

Aaron Rakers
Yes, thanks for taking the question. I wanted to ask about the server market. Maybe two parts. I guess when we look at some of your peers, I mean it seems to be that the lead times have improved on some of the GPUs, particularly the H100. So, I'd be curious of kind of like, can you talk a little bit about what you've seen on lead times there in terms of your ability to deliver on some of this backlog? How that's changed over the course of this last quarter? And then any thoughts on traditional server recovery? How do we think about the pace of that embedded in your expectations looking through this year?

Antonio Neri
Yes, sure. I think I covered the larger part of that question with Toni's question about sequential improvement in the traditional server, which obviously, is still very CPU centric. On a combined Server, right, now 25% of the total volume is APUs, which obviously GPU is the biggest portion of it. So, we expect that sequential improvement driven by recovery in demand in units and then obviously, the shift to Gen11, which is important in this transition.

On the GPU lead times, they have come down but it's still elevated. We're talking about 20 plus weeks at least lead time. And so…and then it's going to be a combination of multiple type of GPUs, right, because there is still demand for the prior generation to H100. Obviously, the majority of the demand today is on H100. Going forward, we're going to have the Grace Hopper H200 and others, right, including MI 300x and the like. And the difference for us is that because we have a unique network in interconnect fabric, we can support all of them. So that's an important differentiation that I think everyone needs to remind because while a lot of volume today is NVIDIA and then on the supercomputing, which is also an AI business, by the way, we support all three of them. And so that, for us, gives us the optionality to convert the orders that we have and future orders we see in the pipeline with a little bit more flexibility, I will say.

Shannon Cross
Thank you, Aaron. Gary, can we have the next question?

Operator
The next question is from Wamsi Mohan with Bank of America. Please go ahead.

Wamsi Mohan
Yes. Thank you so much. You said some of your demand in AI systems is coming in, we had GreenLake. Can you help us understand the linkage between your view of AI revenue and ARR growth?

Antonio Neri
Yes, Wamsi. I can start, and then Marie, feel free to add. I mean, the fact of the matter is that when you look at that $4 billion cumulative orders a significant portion is going to go through the HPE GreenLake. If you recall last year, I announced that a hyperscaler placed an order with us, and that order is going through the GreenLake platform. And so, that's why you see a breakdown over time of the AI GPU orders going through the ARR, which is fine, you know, ultimately, they give us the ability to attach other services, which is important to remember here, because remember, when it goes to the HPE GreenLake, in many cases, we are actually running those systems for the customer. It's not just shipping the system to the customer. We actually put it in a location, whereas our data center footprint with our cooling and power, and then we attach our services, which are the run time plus other things we do. And why it's important also the growth in ARR, because that drives margin expansion and accretion over time. So that's what's going on in addition to the fact that, obviously, now we crossed 31,000 customers on HPE GreenLake platform. To put it in context, Wamsi, that's almost 3,000 customers in one quarter. I mean, 8% up quarter-over-quarter, 3,000 customers. And everything we do from the software perspective, it now is a subscription whether you sell HPE ProLiant Gen11, let's say, an AI-optimized server, the software to connect the server actually runs to GreenLake. Obviously, Alletra software runs through the GreenLake. A lot of the Aruba software, including Aruba Central, a subscription and now you have AI as well.

Marie Myers
And maybe Wamsi, I'll just put a couple of numbers around the APU or the AI system orders that we saw, too. So we ended the quarter actually with $3 billion in backlog, so we really nearly tripled actually year-on-year. And in terms of just the link back to ARR, we shipped around $400 million in revenue, but we had incremental revenue actually that went into ARR. And that sort of underscores the growth that we saw in ARR and expect to see going forward as well, Wamsi.

Shannon Cross
Thank you, Wamsi. Gary, can we have the next question?

Operator
The next question is from Samik Chatterjee with JP Morgan. Please go ahead.

Samik Chatterjee
Hi, thanks for taking my question. I guess, Antonio, you referenced the increase in demand on AI that you see related to inferencing workloads on the enterprise side. Can you maybe talk a bit in terms of how these deployments are looking different to what you've been doing on the AI training side and maybe with some hyperscalers? And given the lead times, is this demand going to more materialize in relation to revenue more in fiscal '25 or is that just a fair estimate about, given the lead times? Thank you.

Antonio Neri
Yes. No, thank you. That's an excellent question. Obviously, I spoke about the AI life cycle, training, tuning, and inferencing. Obviously, the training side has been more focused on the hyperscalers, so tier 2, tier 3 type of providers or companies that are funded well to build these

large language models. But when you look at enterprise, most of the enterprises are going to take a model and fine-tune it to give a context to the model with their data. And that, it can happen in multiple locations, right? It can happen in the data centers or potentially in [indiscernible] or in some cases in the public cloud, but we see more focus on where they can control the data in a secure environment.

And then the inferencing side, it can happen in a data center or in a public cloud once they are all trained, but also at the edge. In fact, we showcase a lot of the inferencing cases at the edge in mobile congress at the edge of the network. Think about use cases like the Coles Supermarket, right? So there is a lot of data to the video footage captured in the stores that video footage needs to be inferenced right there at that given moment with zero latency in order to deliver the outcome. And there are others manufacturing and the like. And in fact, one of the use cases we saw also for inferencing is a large bank that now are doing some fine-tuning and inferencing to do risk management and other things.

So we, I will say we are kind of getting into it. I will say the growth will happen in the second half in '25. Definitely, the lead times will play a role. But I'm very encouraged about the momentum we see and the opportunity we have also with the combination of Juniper, because most of these inferencing requires the network connectivity to deliver it. And that to me is one of the thesis why we went ahead with that acquisition.

Shannon Cross
Thank you, Samik. Gary, can we have the next question?

Operator
The next question is from Tim Long with Barclays. Please go ahead.

Tim Long
Thank you. Just on the…another one on the AI server side. Can you talk a little bit about how you guys are thinking about profitability for the businesses as we get more accelerated compute in your servers? And if you can also kind of break that down between maybe the Cray businesses and the standard compute? Is there going to be more of a margin gap in those two businesses when looking at more traditional going to accelerated? Thank you.

Antonio Neri
Sure, I can start. I will say, listen, I think if you look at our Server segment that we just published, we delivered very strong performance. I mean, we are in the target range we committed well back of 11% to 13%. And that, the fact that we're bringing together give us the flexibility, opportunity to maximize the blended margin here as we go forward. To give a reference, when you sell an EX system, generally it's a liquid cool system that tends to gravitate to the supercomputing side or large AI clusters of thousands of GPUs. But an EX-system supports today up to 80,000 GPUs in one system, and that's because of our interconnect fabric HPE Slingshot. In fact, some of those systems have 80,000 GPUs and maybe 40,000 CPUs in one [indiscernible] system. But then you have other customers are many 2,000 or 4,000 GPUs. And depending on which location they pick, they need liquid cooling, we deploy those.

Now, generally speaking, the XD-platform…the Cray XD platform is the one that has the density and is more cool oriented and ability to mix much different configurations, and that's where the vast majority of the action is today in AI. And ProLiant Gen11 actually is more used for inferencing or in some areas of fine-tuning as well. So we have the flexibility to meet all those demands with our unique IP. And on top of that, we actually lay our machine learning development environment. In fact, there are customers that come to us just for the MLDE environment. Later on, we pull the server.

Now on AUPs, I will tell you that when you sell an XD depending [indiscernible] 20 times the value of a traditional server with CPUs. And in EX, it can be up to 35 times. And so, as we go forward, the ability to optimize margin through the configs and attach the services, whether it's our data center services plus the software and the operational services allows us to really drive the best outcome for our shareholders.

Shannon Cross
Thank you, Tim. Gary, we'll take one final question.

Operator
And that final question will come from Lou Miscioscia from Daiwa Capital Markets. Please go ahead.

Lou Miscioscia
Hey, thank you for taking my call. Antonio, I guess the question I have is, since you're talking a lot about data centers. I'm wondering what's going to happen is that the vast majority of x86 applications going to start to shift over to really be accelerated with GPUs due to the concern of more of those coming to end? And what I'm asking is not really inference and it's not training. These are just normal applications, sort of like the same way architectures shifted from IBM mainframes or PA years-and-years ago the x86 eventually to that and cloud. Do you think that, that's going to shift over to running on GPUs?

Antonio Neri
Well, thank you for the question. I think we need to understand there are two worlds that will coexist. There is the cloud-native world. Think about the cloud-native world where you have thousands and thousands of applications running on thousands and thousands of servers and they share everything. That architecture will exist for a long, long time because it's cost efficient. And the realization is that those applications were designed for that type of environment, for the traditional monolithic approach to more a cloud-centric approach.

And then you have these AI applications where you may have one application…only one, running on thousands and thousands of servers, which have accelerated compute. And it's a little bit far-fetched to say everything is going to move there. I argue that you will have inferencing solutions that a CPU will be just fine. Think about your phone, right? The phone will have, at some point, the ability to manage a large language model, let's say, 20 billion or 30 billion parameters, or the PC, maybe in the 80 billion to 100 billion parameters, but when you go up higher than that, obviously, you need potentially a server at the edge. And then eight…what I'll call eight way GPU will be the right way to go. So I argue there will be a mix in the transition here for a long period of time, not everything will go to a GPU and also depends how these large language models and all the AI applications get constructed.

Now you asked us…you made another interesting point, which I wanted to make sure the…all of you remember. We, as a company, have two now public instances of AI, power with renewal energies where we are supporting some of these customers, including a hyperscaler, and going forward, enterprise customers because they don't have the space and the cooling and the understanding how to run the system of scale. That's a unique differentiation Hewlett Packard Enterprise have in addition to build systems and ship them. And I think that's an opportunity for us because that will drive stickiness to our HPE GreenLake platform, which obviously will drive recurring revenues, but better attach software and services down the road. And Juniper will play a huge role in that environment.

Shannon Cross
Thank you, Lou. Let me now turn it back to Antonio for concluding remarks.

CONCLUSION

Antonio Neri
Well, thank you, Shannon. And thank you, everyone. I know you have been covering multiple calls today. I know it's late on the East Coast, but I will leave you with a few comments. Number one, we have the right strategy and the right team at the right time. This quarter, obviously, it was a little bit mixed because of the revenue. But remember, a lot of revenue also went through the ARR, so we need to understand that implication going forward. I'm very confident about the future. And the moves we have made and continue to make, including the Juniper acquisition, will allow us to participate in this inflection point with a unique IP. Everybody, obviously, is focused about this AI momentum and the server side, but you need more than servers.

AI will drive the need for more ports that means you need more networking bandwidth, that's for sure. Also, let's not forget, we need to do this responsibly. One of the things I'm really proud about our company is the commitment to social responsibility. Doing all of this, addressing the sustainability and the ethical challenges and the responsibility around AI. But listen, just we came out two weeks ago where HPE was ranked number one in the Just [ph] Capital, ranking something that we are proud of it, and I know shareholders value all of that. We have to take some actions here. We are really focused on the strong execution and discipline, something we have shown now for six years plus. And that's why I'm confident in the adjusted guidance we provided with Marie. And as we get into '25, obviously with the pending acquisition, I feel HPE will be even in a stronger position as we get through 2024. So, thank you for your time and hope to connect with you soon.

Operator
Ladies and gentlemen, this concludes our call for today. Thank you. You may now disconnect.