

# Hewlett Packard Enterprise

## Second Quarter Fiscal 2024 Earnings Conference Call

Tuesday, June 04, 2024, 5:00 PM Eastern

### **CORPORATE PARTICIPANTS**

**Antonio Neri** - *President, Chief Executive Officer*

**Marie Myers** - *Chief Financial Officer*

**Jeff Kvaal** - *Head of Investor Relations*

## **PRESENTATION**

### **Operator**

Good afternoon and welcome to the Second 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 toward 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, Jeff Kvaal, Head of Investor Relations. Please proceed.

### **Jeff Kvaal**

Good afternoon. Welcome to our second quarter fiscal 2024 earnings conference call with Antonio Neri, HPE's President and CEO, and Marie Myers, HPE's CFO.

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 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 the 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 April 30<sup>th</sup>, 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 SEC filings 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 presentations accompanying today's earnings release on our website for details. Throughout the call, all revenue growth rates are presented on a year-over-year basis and adjusted to exclude the impact of currency, unless otherwise noted.

Finally, Antonio and Marie will reference our earnings presentation in their prepared remarks. With that, Antonio.

### **Antonio Neri**

Good afternoon and thank you for joining us today. HPE delivered a very solid performance in the second quarter, with revenue and non-GAAP diluted net EPS exceeding our outlook range – driven by AI systems revenue more than doubling from our first quarter. I am very optimistic about where we are headed.

AI demand continues to accelerate, with cumulative AI systems orders reaching \$4.6 billion dollars this quarter. We have a robust pipeline in this business, though large AI orders can

cause fluctuations during the quarter.

We anticipate continued revenue growth, driven by increased AI systems demand, continued adoption of HPE GreenLake, and ongoing improvement in the traditional infrastructure market, including servers, storage, and networking. Due to our confidence in the second half of fiscal year 2024, we are raising our full-year revenue and non-GAAP EPS guidance and reiterating free cash flow guidance. Marie will provide more specifics in her remarks.

While we focus on translating strong AI customer demand to revenue growth, we continue to drive cost discipline to operate more efficiently and to preserve the ability to make targeted investments, which will sustain our growth into the future. We are being prudent with our spending, and reduced operating expenses in the first half, as compared to the prior-year period. We are also driving business process simplification across the company – including through digitization and automation with AI.

Demand for HPE's AI systems is accelerating at a faster pace, and our solid execution enabled us to more than double our AI systems revenue sequentially to over \$900 million dollars, helped by supply chain conversion through improved GPU availability. Our lead time to deliver NVIDIA H100 solutions is now between 6 and 12 weeks, depending on order size and complexity. We expect this will provide a lift to our revenues in the second half of the year.

Enterprise customer interest in AI is rapidly growing, and our sellers are seeing a higher level of engagement. Enterprise orders now comprise more than 15 percent of our cumulative AI systems orders, with the number of enterprise AI customers nearly tripling year-over-year. As these engagements continue to progress from the exploration and discovery phase, we anticipate additional acceleration in enterprise AI systems orders through the end of the fiscal year.

We are winning in AI for several reasons.

Our leadership in AI at scale -- and across the entire AI lifecycle – from training to fine tuning to inferencing, continues to attract new customers. These customers range from new generative AI model builders to broader service providers, to sovereign states, to traditional enterprise customers, to large hyperscalers.

For example, in partnership with Microsoft, we are extending the Azure AI platform to HPE infrastructure, which provides Microsoft with additional capacity to serve even more customers, including OpenAI.

HPE has decades of experience in the design, manufacturing, and management of air- and liquid-cooled systems, including the data center infrastructure to reliably deliver the highest levels of computing performance. Customers appreciate our AI-at-scale expertise and intellectual property, as well as our unique liquid-cooling manufacturing and services capabilities.

As accelerated computing silicon innovation advances, higher power densities demand direct

liquid cooling technologies. Building direct liquid cooling AI systems is complex and requires manufacturing expertise and infrastructure, including power, cooling, and water.

With more than 300 HPE patents in direct liquid cooling, proven expertise, and significant manufacturing capacity for these kinds of systems, HPE is well positioned to help customers meet the power demands for current and future accelerated computing silicon designs.

Our leadership in AI was once again validated in May, when the latest Top500 list of the world's most powerful supercomputers was released. HPE now has four of the top ten world's fastest supercomputers, all of which are direct liquid cooled. Two of these systems are exascale supercomputers, with Frontier still the world's fastest and Aurora now breaking the exascale barrier. I am also proud that we have built 7 of the world's top 10 energy-efficient systems, according to the latest Green500 list. This experience makes us an attractive partner to sovereign states and governments pursuing AI strategies.

HPE benefits from a strong ecosystem of AI partners, including NVIDIA. We introduced co-engineered enterprise solutions with NVIDIA last year to streamline the model development process – as well as to enable enterprises to fine-tune large language models with their private data to accelerate inferencing. I am excited that NVIDIA CEO Jensen Huang will join me at HPE Discover Las Vegas in just two weeks. Together, we will unveil new, exciting, and differentiated innovations that will simplify and accelerate enterprise AI adoption and deployment.

Enterprise customers are already responding to our unique AI portfolio. For example, CUBOX, a facial and image recognition company in Korea, is developing generative AI models using HPE AI systems to enhance identity verification at locations like the Incheon International Airport in South Korea.

JT Group, based in Japan, operates a pharmaceutical business and is planning to use our HPE AI systems to support AI model training and simulations to accelerate drug discovery.

And we just announced that we will power Scaleway's AI cloud service offering using our HPE AI systems. The new service will make powerful computing accessible to companies to support their various AI workloads and use cases.

As we capitalize on the AI growth opportunity, we also see indications of market recovery in the traditional and cloud infrastructure markets.

Orders for traditional servers grew sequentially and year-over-year, driven by enterprise, public sector, and SMB customers in North America and Europe. We are seeing no indication of cannibalization from accelerated computing demand. And revenue grew sequentially as customers transition to higher-AUP HPE ProLiant Gen 11 servers.

Differentiated, customer-centric innovation positions us well to capture this market recovery.

For example, our leading HPE GreenLake hybrid cloud is attracting new customers. In the second quarter of fiscal 2024, the number of customer organizations using HPE GreenLake

increased sequentially by almost 9 percent to 34,000. And our as-a-service lifetime Total Contract Value grew to more than \$15 billion dollars in Q2 with our annualized revenue run rate, or ARR, growing 39 percent year-over-year.

Demand is increasing for our HPE GreenLake on-premises private cloud solutions.

The Defense Information Systems Agency, a combat support agency of the United States Department of Defense, selected HPE to develop a distributed, hybrid multi-cloud platform prototype on HPE GreenLake as part of an effort to simplify the organization's management of disparate IT infrastructure and resources across public and private clouds.

In storage, we accelerated the transition of our portfolio in Q2 to meet the needs of hybrid cloud and AI.

Several weeks ago, HPE introduced significant new functionality to our HPE Alletra storage offerings. We rounded out our block offering by extending the hybrid capabilities of HPE Alletra block storage to AWS, doubling its capacity to address more customers, and enhancing its automation capabilities with generative AI. Earlier this quarter, we introduced new HPE GreenLake for File Storage capabilities with options specifically targeting the unstructured data demands of AI.

We have also added significant specialist sales capacity in recent months. While it takes time to activate new sellers and bring them to full productivity in a market with a long sales cycle, we expect increased order-to-revenue conversion in the future.

The enhancements to our winning portfolio, complemented by a more focused sales force, position HPE to strengthen the already-robust customer adoption of HPE Alletra. More than 1,000 new HPE Alletra MP systems have been deployed to date, which is the fastest product ramp in the history of our company.

In networking, the market remained in transition during the quarter, as customers continued to work through their current inventory.

As demand in this segment gradually returns, we believe our broad portfolio positions us well. We expect modest sequential networking demand improvement, driven by the state and local purchasing season in the United States.

We announced significant new innovations during the quarter to align with HPE's broader AI strategy. These solutions include generative AI capabilities to improve AIOps, and Wi-Fi 7 access points that capture edge data for AI inferencing. In addition, last month we launched new security and AI observability tools to help fight AI cyber risks. And just yesterday, we expanded the most complete private 5G and Wi-Fi portfolio in the market with the launch of HPE Aruba Networking Enterprise Private 5G. All of this will be delivered through our HPE GreenLake cloud platform.

Customers are responding to our networking innovation. In the last few months, customers ranging from Houston Airports to Bethesda, Health Group to Mercedes-Benz Stadium, to the University of Maryland have turned to HPE to enhance the experience for their visitors, residents, and employees.

As we look to our future in networking, we continue to be very enthusiastic about the proposed acquisition of Juniper Networks. We are currently in the regulatory process for this transaction and expect to close by the end of 2024 or early 2025.

As I mentioned earlier, we continue to invest in innovation while we drive operational and cost discipline to continue to improve our cost structure. We are focused on reducing complexity in our business processes, as well as implementing automation and AI across the company to enhance customer service, R&D, and team member productivity.

I also want to note that we recently announced we have restructured the sale of our stake in H3C. Since the transaction is large and complex, the required regulatory approvals were taking longer than previously anticipated. So, we agreed to restructure the sale with UNIS. The updated agreement provides HPE with the opportunity to sell a significant portion of the shares in the coming months. The new payment structure has no impact on the pending Juniper Networks transaction, as the structure of the deal financing does not rely on any of the H3C proceeds.

Finally, you likely saw that we announced in May that we have agreed to divest our Communications Technology Group, to enhance our strategic focus in high-growth, high-margin parts of the market, including the service provider and enterprise markets.

In closing, I want to reiterate that I am proud of a very solid performance in Q2. It shows the alignment of our strategy and innovation to major market opportunities. We have greater optimism about the second half of the year, leading us to raise our full-year revenue and non-GAAP EPS guidance.

AI is creating growing demand across our portfolio, and we see significant opportunities across customer and business segments. Our competitive advantages – from deep expertise in standing up AI systems, to our differentiated HPE GreenLake cloud, to our networking and storage offerings -- position us well.

We have an excellent team, and I am confident in our ability to continue executing with discipline to take advantage of incredible opportunities presented by this era of innovation.

I am looking forward to sharing with you our latest breakthrough innovation and partnerships across AI, hybrid cloud, and networking later this month at HPE Discover Las Vegas. We are excited to be the first corporate keynote at Sphere, and I hope to see many of you there.

I'd now like to hand it over to Marie, who will talk about the details of our segments and our outlook. Marie, over to you.

**Marie Myers**

Thank you, Antonio, and good afternoon everyone. It's a pleasure to be here with you all after my first full quarter as HPE's CFO. Over the past three months, I have become even more excited about our opportunities across AI, hybrid cloud, and networking. We remain in the very early days of AI, yet it is already driving strong interest, pipeline, orders, and revenue across our portfolio, from servers to storage to services to financing. Our AI system revenue inflected positively in Q2. We are winning deals in the AI market now and are well positioned for additional demand from enterprises and sovereigns in the future. Our differentiation includes decades of liquid cooling expertise, which we expect to become even more in demand with future iterations of chips, including NVIDIA's. In short, we see AI as a long-term driver of our financial results and as a pillar of our strategy to pursue higher-growth, higher-margin revenues.

We are very pleased that we have exceeded our expectations in Q2 across key metrics. We exceeded the mid-point of our revenue guidance by \$400 million dollars, non-GAAP diluted net EPS was above the high end of our range, and free cash flow exceeded \$600 million dollars. Improving enterprise demand for traditional servers on top of the expected sharp ramp in AI servers drove the outperformance. Our AI orders are healthy, Intelligent Edge is set to grow sequentially beyond Q2 as expected, and AI emerged as a driver of our healthy HPE GreenLake momentum. We are seeing rapid growth in AI system revenue. Overall, I am very pleased with our performance in Q2 and am excited about our continued progress through fiscal 2024.

Let's take a closer look at the details of the quarter. Revenue grew 4 percent year-over-year and 7 percent quarter-over-quarter in constant currency to \$7.2 billion dollars. This exceeded the midpoint of our prior guidance by approximately \$400 million dollars. We have strong momentum in HPE GreenLake. The number of customers that have adopted HPE GreenLake rose 9 percent sequentially. ARR grew 39 percent year-over-year to above \$1.5 billion dollars in Q2. Storage and networking are typically the fastest growth elements of ARR and both retain robust growth rates. This quarter, AI was the fastest growth component of ARR.

Our software and services mix rose approximately 200 basis points year-over-year to 67 percent. 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 key differentiator. We expect HPE GreenLake's value proposition to key customers, including enterprises and sovereigns, to sharpen with the advent of AI.

Our Q2 non-GAAP gross margin was 33.1 percent, which was down 310 basis points sequentially and year-over-year, driven by a mix shift from our higher-margin Intelligent Edge revenue to Server revenue, plus an unfavorable mix within Hybrid Cloud. Our Q2 non-GAAP operating expenses fell 1.6 percent year-over-year despite our revenue growth of 4 percent.

Our opex discipline partially offset lower non-GAAP gross margins and held the non-GAAP operating margin decline to 200 basis points sequentially and year-over-year to 9.5%. The opex discipline plus higher revenue drove GAAP diluted net EPS of \$0.24 and non-GAAP diluted net EPS of \$0.42. The latter exceeded the high end of the guidance range on strong revenue and cost discipline. Our non-GAAP diluted net earnings per share excludes \$247 million dollars in net costs, primarily from stock-based compensation expense, amortization of intangibles and acquisition, and other related charges.

We are managing the business with focus and discipline and evolving into a simpler, more agile company. We are also investing to capitalize on growth from the interrelated inflection points in AI, hybrid cloud, and networking, and to drive structural higher profitability over time.

Let's turn to our segment results. Server revenues were \$3.9 billion dollars in the quarter. This was up 16% sequentially and up 18% year-over-year. Strength in both AI systems and traditional servers drove the healthy revenue growth.

Our cumulative AI system product and services orders since Q1 2023 rose approximately \$600 million dollars sequentially to \$4.6 billion dollars. I am very pleased our AI system product revenue more than doubled sequentially to over \$900 million dollars. This strong revenue growth allowed us to make progress against our backlog which is now \$3.1 billion. Given the growing importance of our services business, we have updated our AI disclosures for this quarter to include services. Services is a small portion of our AI systems metrics at present, though we expect it to become more meaningful over time. Our differentiation -- with liquid cooling, software, HPE GreenLake, and increasingly services -- is resonating in the market. We have seen a threefold increase in our enterprise AI customer base in the past year.

Revenue from our traditional server business increased sequentially. We expect this trend to continue. Demand is improving as enterprises digest prior purchases and gain more comfort with the macro outlook. The structural mix shift to higher AUP Gen 11 servers is ahead of our expectations and we are able to pass through rising input costs. We are encouraged that our Gen 11 pipeline is starting to include AI inferencing activity in enterprise applications and we see more evidence of adoption in the enterprise in Q2.

Our Q2 operating margin was 11.0 percent. This was down 40 basis points sequentially and was in line with the expectations we laid out last quarter for operating margins near the low end of our long-term 11-13% range. While pricing remains aggressive in the server market, particularly in AI systems, we remain disciplined in cost and price as we pursue profitable growth.

Hybrid Cloud revenues of \$1.3 billion dollars were up 1 percent sequentially and down 9 percent year-over-year. We are already seeing some cross-selling benefits of integrating the majority of our HPE GreenLake offering into a single business unit. I mentioned the 39% growth in ARR this quarter.

Our traditional storage business was down year-over-year. The business is managing two long-term transitions at once. We've talked about our migration to the more software-intensive Alletra platform. This is reducing current period revenue growth though locking in future recurring revenue. Storage ARR growth of over 50 percent year-over-year offers early confidence in the migration. The second transition is from block storage to file storage driven by AI. While early, this is also on the right trajectory. Our new file offerings plus the sales force investment Antonio mentioned tripled our pipeline of file storage deals sequentially in Q2.

Our operating margin was 0.8 percent, which was down 300 basis points sequentially and 110 basis points year-over-year. Reduced revenue scale and an unfavorable mix of third-party product in traditional storage was the largest driver of the sequential change.



Intelligent Edge revenues were \$1.1 billion dollars. Revenues fell 9 percent sequentially and 19 percent year-over-year. Backlog consumption created difficult compares with both prior periods. Our backlog is now at normal levels. The demand environment remains soft and large enterprises have yet to return to the market in force. However, we do see some green shoots that give us confidence networking will transition to modest sequential growth beyond Q2 as we had expected. Our channel inventory remains within the normal range. Wi-Fi has grown sequentially for two consecutive quarters. Growth remains strong in software and services. Attach rates and renewals for Aruba Central, SASE, and our AIOps software remain strong. The Intelligent Edge portfolio of subscription revenue grew above 50 percent year over year.

The segment operating margin of 21.8 percent was down 760 basis points sequentially and 290 basis points year-over-year. As expected, the lower revenues, reduced mix of switching business, and less revenue from backlog were the primary drivers. As we indicated last quarter, we have reset our opex plan for the year to account for lower revenue and expect the Intelligent Edge operating margin to be back in the mid-20 percent range by Q4.

Our HPE Financial Services revenue was up 1 percent year-over-year and financing volume was \$1.7 billion dollars. Our operating margin of 9.3 percent was up 80 basis points sequentially and 40 basis points year-over-year. Our Q2 loss ratio remained steady below 0.5 percent.

These results are what we have come to expect from this high-quality, predictable business. However, underneath these steady results, FS is already adapting to drive AI growth across the business. Year to date, nearly half a billion dollars of our \$3 billion in financing volume went to AI wins with both cloud and enterprise customers. This illustrates our prior point that AI is driving demand to every one of our businesses.

Turning now to cash flow and capital allocation. We generated \$1.1 billion dollars in cash flow from operations and \$610 million dollars in free cash flow this quarter. HPE typically consumes significant amounts of cash in the first half of the year and then generates cash in the second half. We are ahead of traditional free cash flow patterns thus far in fiscal 2024 given higher than expected net income in Q2, prepayments for AI systems, and timing of working capital payments.

Our cash conversion cycle was negative four days, which is a reduction of 28 days from Q2 2023. Our days of inventory and days payable were both higher to support our expected growth in AI system revenue in the second half. We returned \$214 million dollars in capital to shareholders in Q2, including \$169 million in dividends and \$45 million in share repurchases. Our year-to-date capital return is \$386 million.

Let's turn now to our forward view. We expect a materially stronger second half led by AI systems, traditional servers and storage, networking, and HPE GreenLake. Let me recap the key drivers that factor into our expectations for Q3 and the full year.

For Server: We expect improving GPU supply for AI systems and improving demand for traditional servers to drive sequential revenue increases through fiscal year 2024. While the rising AI systems mix is a gross margin headwind, we are balancing this with higher-margin

services revenue, improving scale and cost discipline. We expect the segment operating margin to be approximately 11% for the fiscal year.

For Hybrid Cloud: We expect slight sequential revenue increases through the year. HPE GreenLake growth should continue and traditional storage should improve slightly. We expect operating margin to improve modestly to the mid single digit range through the year as HPE GreenLake deals mature, new products ramp, and our sales force optimization gathers momentum.

For Intelligent Edge: We anticipate slight sequential growth in Q3 and Q4 driven primarily by seasonal education spending rather than improving markets. We continue to expect our cost reduction efforts to materialize in the second half and our full year operating margin to be in the mid-20 percent range.

With that context, let me turn to our outlook. For Q3, we expect revenues in the range of \$7.4 to \$7.8 billion dollars. We expect GAAP diluted net EPS to be between \$0.29 cents and \$0.34 cents and non-GAAP diluted net EPS between \$0.43 cents and \$0.48 cents.

For fiscal year 2024. We now expect constant currency revenue growth of 1 to 3 percent, which is up from our prior 0 to 2 percent range. We reiterate our non-GAAP operating profit growth guidance of 0 to 2 percent. We are reducing our GAAP diluted net EPS guidance by \$0.20 cents to \$1.61 cents to \$1.71 cents to incorporate the recent updates to our H3C proceeds. We are raising our non-GAAP diluted net EPS guidance up \$0.03 cents to \$1.85 cents to \$1.95 cents. This incremental \$0.03 cents, or \$0.06 cents annualized, reflects the contribution from the retained portion of our H3C stake. We are also increasingly comfortable with the high end of the non-GAAP diluted net EPS range given our OI&E and operational improvement. We are excluding from our non-GAAP results the gain on sale from our H3C and CTG divestments.

This year's mix shift from networking to AI systems should weigh on our gross margins. We expect the fiscal 2024 non-GAAP gross margin to be below our full year expectation of 35% from our analyst day. To balance the mix shift, we are driving further simplicity and efficiency across the business. We are accelerating our generative AI capabilities, such as implementing HPE-specific large language models in chatbots for our sales and service representatives. As I mentioned last quarter, prudent cost management, simplified processes, and disciplined execution across cycles are key tenets of our long-term journey towards higher margins.

These cost actions will be evident in financial results in the second half of fiscal 2024. We now expect fiscal year 2024 opex to be down modestly from fiscal 2023 opex. Our prior view was flat to down. This includes a sequential increase in Q3 for marketing before a sequentially lower Q4, which will serve us well heading into fiscal 2025. We expect our fiscal 2024 operating margin to be flattish year over year. We now expect OI&E to be less of a headwind this year. We anticipate a \$150 million dollar headwind vs our prior expectation of a \$200 to \$250 million dollar headwind given a one-time benefit in Q2 and the retained portion of our H3C stake. We expect the effect of currency to be immaterial.

Our strong first half free cash flow increases our confidence we will deliver at least \$1.9 billion dollars in fiscal year 2024. We expect significantly stronger free cash flow in the second half of the year led by higher earnings given our ramp in AI systems. This does not include the \$2.1

billion dollars that we expect to receive from Unisplendour this fiscal year as a result of our recently restructured agreements to sell our stake in H3C. We expect working capital to be neutral to free cash flow as we expect declines in inventory to balance declines in accounts payable.

We remain committed in the long-term to our balanced capital allocation framework, including our target of returning 65-75% of free cash flow to shareholders. In the near-term, we expect to continue share repurchases at a pace in line with Q2 as we prudently manage our balance sheet ahead of the anticipated receipt of the H3C proceeds and the Juniper transaction closing. The proposed Juniper deal remains on track to close in late 2024 or early 2025 as planned. We remain committed to our dividend and to our investment grade rating.

To conclude: Our solid Q2 results illustrate how comprehensively AI is affecting our portfolio. We are capturing profitable growth opportunities in the AI market. We are excited for Discover and look forward to seeing many of you at our IR Summit.

I'll open it up now for your questions.

## 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 pick up 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 is from Amit Daryanani with Evercore. Please go ahead.

### Amit Daryanani

Thanks a lot, and congrats on a nice print. Yes, I guess my question is really on the AI system side. You folks obviously had some really good revenue performance here. I think business pretty much doubled sequentially. I'm hoping you could spend some time talking about how should we think about margins on the AI side versus maybe segment averages because your margins appear to be holding up a lot better versus what they have done at Dell or Super Micro, for example. So, I'd love to just understand how do we think about AI system margins versus corporate averages? And if you just touch on what do you think differentiates your AI solutions versus companies like Super Micro, Dell, that'll be super helpful. Thank you.

### Marie Myers

Hi Amit, good afternoon. It's Marie, and happy to take the question. And look, obviously, we shipped \$900 million of AI revenue in the quarter. And I think as you saw, our margin rate on our server segment from an operating profit perspective was 11%. So, look, honestly as a CFO, very pleased that's what we guided and that's what we achieved, and that's what we've guided for the second half of the year.

In terms of what's the puts and takes and the headwinds and tailwinds, let me give you some color around it. So, from a headwind perspective, obviously we're dealing with a tougher inflationary commodity market, particularly even in the second half. And we also see, as you saw in the quarter, a greater mix of AI servers actually in our overall server mix. And as you

know and you've heard, it's a pretty competitive market out there I might add, too. But clearly, we're managing it very well. And I'd say what's benefiting us is a couple of really important tailwinds. One, is what we are seeing in terms of the Gen11 transition. That part of the business brings a higher AUP. And then also frankly, Amit it's the work we're doing around costs. So as a result, I think this really illustrates that we are very disciplined around both price and cost and frankly, it's part of our strategy to drive profitable growth. So pleased with the results, and that's how we are guiding the back half of the year, Amit.

### **Antonio Neri**

So, Amit, on the differentiation, I will summarize this on four key elements. One is our ability to deliver and run systems at scale, so AI system at scale, that's a unique expertise and we have decades of experience.

Number two is our infrastructure cooling intellectual property. We actually have all the IP necessary to cool systems in three different ways for them other. Our manufacturing footprint, which is very unique, we have one of the largest water-cooled manufacturing footprints in the world with two very important locations in the US and in Europe, which are close to customers.

And then last but not least is services. What I think people are coming to realize that running the system of scale requires unique services capabilities. And that's why, with Marie, we started showing you what the services pull-through is, which is also over time, a lever to improve the gross margin in this business. And we cover all aspects from day zero which is consulting, to day one, which is advisory and professional services design and then...and build, and then day two, which is a running part with our...in our operational services side and deep expertise when it comes down to this system of scale, including direct liquid cooling. All those four key elements are a big differentiation for us.

### **Jeff Kvaal**

Thank you, Amit. We move to the next question.

### **Operator**

And the next question is from Aaron Rakers with Wells Fargo. Please go ahead.

### **Aaron Rakers**

Yes. Thanks for taking my question. I guess sticking on the AI topic, if I could first ask, when you referenced the AI enterprise customers starting to show up, and I think the comment on the conference call was...it's now north of 15% of your AI orders. Can you give a little bit more context of that? What has that been over the last couple of quarters? I'm just trying to think about the trajectory of that.

And then Antonio, on the liquid cooling side, as we and investors think about Blackwell product cycle from NVIDIA. I'm curious of...can you be a little bit more specific of exactly where, from a technology perspective, you differentiate at liquid cooling? Is there something unique that HPE does within the 300 patents that you would want to highlight for us, as sustainably differentiated. Thank you.

### **Antonio Neri**

Yes. Thanks, Aaron. I think it's fair to say that the enterprise demand has started to pick in Q1 and really accelerate in Q2. Early on, if you go back to the Q1 2023 all the way to, call it, calendar 2023, a lot of demand came from the model builders, the service providers, and

hyperscalers and you see the reference to our partner, Microsoft. But as I think about that, I think three segments of the market. I think about DaaS segment, which is driving a lot of demand for a lot of GPUs and accelerated compute in general. The second segment is sovereign clouds and that starts to pick up now. And that's, think about tens of customers around the globe because obviously, our countries.

And then there is enterprise customers', which will be thousands of customers over time. So today, we are very pleased with the momentum in enterprise AI. A lot of the conversations are happening. I was in Europe the last two weeks and I was in London, I was in Paris, I was in Madrid. All of them want to adopt AI. And I think they are attracted to HPE because of the trust, the quality, and the ability to deliver a simplified experience to our HPE GreenLake platform.

As for the differentiation, HPE has three different ways to cool systems. So, one is the traditional way, which is called the liquid-to-air cooler, think about that, basically running water supply in chilled locations where basically cools the air around the systems. Everybody has done that for a long time.

The second is what most of the industry is doing today, which is what I call 70% direct liquid cooling or hybrid liquid cooling. Those companies still use fans to cool aspects of the systems. Some of our competitors talk about direct liquid cooling, but that's exactly what they're doing, and they are doing only a hybrid direct liquid cooling.

And HPE has...and by the way, in that environment, we have 10 systems already in market today that we are shipping and configuring for customers. And then we have what I call 100% direct liquid cooling. And this is a unique differentiation HPE has because we have been doing 100% direct liquid cooling for a long time. And today, there are six systems in deployment, and three of them are for generative AI. And as we go to the next generation of the silicon and you talk about Blackwell, when you go to the B200, that will require 100% direct liquid cooling.

And that's a unique opportunity for us because you need not only the IP and the capabilities to cool the infrastructure but also the manufacturing side. And that's why I said early on to the previous call that HPE has one of the largest water-cooled manufacturing capabilities. So that is what really differentiate us. And even on the hybrid way to cool the system, HPE is not using off-the-shelf solution. We use what we call an ARC, which stands for adaptive rack cooling system, it's a unique design. And the good news is that we can actually colocate in the same data center and in the same aisle, direct liquid cooling and air-cooled systems, which is unique because customers don't want to retrofit data centers. So that's the opportunity.

**Jeff Kvaal**

Aaron, thank you very much, Gary...

**Operator**

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

**Mary**

Hi, this is Mary on for Meta. I just had a question for you on Intelligent Edge. Where are you in the Intelligent Edge inventory digestion and when do you expect to emerge from it?

**Marie Myers**

Yes. Hi, good afternoon. And in terms of where we're at Intelligent Edge, I think as I commented in my prepared remarks, we saw Q2 as being the trough period, and we've been

transitioning through that throughout the quarter. And if you look at the market and where it's at right now, I'd say that our channel inventory right now is in really good shape. And we did mention in the guide that we do expect a modest sequential improvement in networking in the back-half of the year.

### **Antonio Neri**

Yes. And I will say in addition to what Marie said, what I'm really excited is that we see interesting areas of growth happening now. If you go back to Mobile World Congress and you see even announcements like we made yesterday, the enterprise private 5G is picking up significant momentum. Of my almost 40 meetings I had in Barcelona, more than half were about enterprise private 5G. And, so yesterday, we announced the most complete enterprise private 5G on the back of the Athonet acquisition.

### **Jeff Kvaal**

Thanks very much Mary, Gary...

### **Operator**

The next question is from Toni Sacconaghi with Bernstein. Please go ahead.

### **Toni Sacconaghi**

Yes, good afternoon and thank you for taking my question. I just wanted to follow-up on the guidance. You talked about enthusiasm for the second half, but you beat revenues this quarter relative to your expectations by \$400 million and by guiding up an additional percent, you're actually only guiding up the full year by \$300 million. So, I'm wondering, are you just being conservative, given the commentary around enthusiasm and forces at work in the second half or how do we reconcile that discrepancy? And then also just on AI servers for the second half, I think you talked about 6 weeks to 12 weeks lead times. So, if you have \$3 billion in backlog and lead times for 6 weeks to 12 weeks, why can't you deliver \$3 billion in AI systems like next quarter or certainly in the second half? Thank you.

### **Marie Myers**

Yes. Hi Toni, good afternoon. So, this is Marie. I'm going to take the first question just on the guide. So let me just clarify the guidance in terms of how we put it together. So, we raised the guide from...to 1.85 to 1.95, and that actually was the pass-through on that 19% stake in H3C. So, we actually put \$0.03 related to the 19% stake. What I did point to though, Toni, is I pointed to the higher end of the range, so that's really what's giving us confidence based on the increase that we made on revenue. So, you're seeing that higher top-line and then also the confidence I got around just the cost discipline. So you've seen that just in the last couple of quarters where we've had a really strong scrutiny and focus around OPEX. And plus, we did have some favorability in OI&E this quarter, that albeit, that was actually just a one-timer. So, I just want to make that point of clarification. So, overall, Toni, keeping the guide at 1.85 to 1.95, but really pointing to the higher end of the guide in terms of just the confidence that you articulated. So, I'll turn it over to Antonio to cover the second question.

### **Antonio Neri**

Yes. Toni, I think there is an opportunity to potentially exceed that. I think the limiting factor is not the supply, to be honest with you, it's the availability of data center space. I made this comment in Q1, if you recall, data center space and power and cooling. And so, some...we are working with the customers to time everything correctly, 6 to 12 weeks, think about it, maybe less than a quarter, but then you have to go and install it. And there is a nice percentage of our deals in generative AI, which are all actually GreenLake. And so, while we can recognize the

revenue upfront, we are deferring all the services piece of it. So, it really is going to come down to the timing of the data center and the power and cooling. And if that all aligns correctly, then we may have an opportunity to do better. But we felt prudent at this point in time to keep it the way it is and raising by 1%.

**Jeff Kvaal**

Toni, thanks very much, Gary...

**Operator**

The next question is from Samik Chatterjee with JPMorgan. Please go ahead.

**Joe Cardoso**

Hi, thanks for the question. This is Joe Cardoso on for Samik. So maybe just following-up on the AI question. You're seeing a sequential decline in AI backlog this quarter, admittedly off of a very robust revenue quarter. But maybe you can just discuss the pipeline you're seeing. And sitting here, like what are some of the puts and takes that investors should be considering when seeing the sequential decline in backlog despite what appears to be a strong underlying demand environment overall? Thanks.

**Antonio Neri**

Sure. Thank you. Yes, our backlog was slightly down quarter-over-quarter, also was based on the demand we took in and also the fact that we converted more. I think you have to realize some of these deals, particularly larger deals, take time and are little bit lumpy. Some of them have to go through the financing side. And so, we feel good about where we are today. But in terms of the pipeline, think about multiples of the current backlog, multiples of the current backlog. So, as we go through the next weeks and quarters here, we feel very confident in our ability to capture that AI for all the reasons I described before, and also the ability to close these deals as some of them may include, by the way, the need to provide data center space as well because we are not just building the system. And these are all generative AI systems, by the way, all of them that we also need to run it for customers.

**Jeff Kvaal**

Thanks very much Joe, Gary...

**Operator**

The next question is from Simon Leopold with Raymond James. Please go ahead.

**Victor Chiu**

This is Victor Chiu in for Simon. The doubling the AI system's revenue is a pretty sharp inflection this quarter. Can you just tell us how much of the increase was driven by allocation improvement versus ability to ship orders? And kind of just help us understand the dynamic there a little bit.

**Antonio Neri**

No, I don't think it has to do anything with location. I read all this thing about location. We have a fantastic partnership with NVIDIA, it's about lead times in the general H100, which obviously NVIDIA made significant improvements on their own. And then obviously, as we start ramping our manufacturing processes, we actually become better and better at the revenue conversion. And I would say, I want to thank my team because my team did a fantastic job and we feel good about as we go forward. So, it's a combination of multiple things, better lead times on supply. We feel pretty good about the ability to deliver systems 6 to 12 weeks to Toni's question, and

some of them are straightforward, it's just shipments and some are with our HPE GreenLake offering wrapped around.

**Jeff Kvaal**

Thanks very much Victor, Gary...

**Operator**

The next question is from Ruplu Bhattacharya with Bank of America. Please go ahead.

**Ruplu Bhattacharya**

Hi, thanks for taking my questions. It's Ruplu filling-in for Wamsi today. Can you talk about how much of the GreenLake revenue and ARR growth came from AI? And then Antonio, do you think that having GreenLake is helping you sell AI systems? And I also wanted to clarify, Antonio, you talked about sovereign AI. Can you talk about what innings you're in? And are there specific requirements of sovereign AI where you think HPE is well positioned to satisfy?

**Marie Myers**

Hi Ruplu, we look forward to seeing you tomorrow. So, in terms of your question around just ARR and AI, actually, it was the fastest growth element of ARR in Q2, followed by storage and networking. So, I think, as I said in my prepared remarks, we are starting to see AI to sort of move through our entire portfolio. So pleased with the progress that we had this quarter, and I'll turn it over to Antonio to add some more context.

**Antonio Neri**

Yes. On the second part of the question, I think we are early at this point in time. I think there is more to be seen. And I want to make sure I captured exactly what you said. Can you repeat the last part because you were a little bit breaking.

**Ruplu Bhattacharya**

Yes, I just wanted to ask about sovereign AI. What innings are you in, and are there specific requirements that you can satisfy?

**Antonio Neri**

Yes. So, on sovereign AI, I said it's early, early on. I think there is a lot of engagements right now happening at the country level. The good news, we have very good reach across the board. And it's a combination of both. There is a combination of providing what I call generative AI locations where customers, enterprise customers can get access to, to a sovereign cloud that the government may be helping funding at the same time.

And the other one is what I call supercomputing power, right, itself. And so, the two are very well aligned to the sovereign AI opportunities. And that's why I'm excited because HPE already provides to a lot of the sovereign governments supercomputing. So now we can extend into generative AI. If you think about the example of the UK Bristol, it's a generative AI system that the UK is funding as a part of the Bristol University, that the venture is going to be open to start-ups and enterprises in the UK, to either train models or do other type of research. And that's what we see.

**Jeff Kvaal**

Okay. Thanks very much Ruplu. Gary, we have two more questions, please.

**Operator**



And the next question is from Ananda Baruah with Loop Capital. Please go ahead.

**Ananda Baruah**

Hi, good afternoon guys. Really appreciate taking the question. Just a quick clarification and a quick question, Antonio. Did you say that in the last 12 months, AI growth was driven by service providers? And was it mobile, you said? And then the question is, what's a good way to think about cloud, your gen AI cloud scale opportunity, hyperscalers and Tier 2s going forward? Thanks a lot.

**Antonio Neri**

Yes. Now, thank you. I said that the segment where we have seen, obviously the vast majority of action and demand is the model builders. Those are the big companies that build large and small language models. Obviously, you saw our comments about our partner with Microsoft and the extension of their capacity to OpenAI. That's an example of a model builder. But also, there are other service providers. And in fact, in my remarks, I mentioned Scaleway, which is a French service provider that provides the capacity for the local French model builders. In fact, there is a lot...there is a very vibrant ecosystem in France about building AI models that will use that capacity to train the models. That's what I referred to at this point in time. And that's what we see and that's what we've seen. So, we are very interested in not just the model builders but Tier 2, Tier 3, which also is going to be a big driver. But understanding that enterprise, ultimately what the action is going to happen in terms of fine-tuning and deploying these AI models over time.

**Jeff Kvaal**

Okay. Thank you very much Ananda. And last question, Gary.

**Operator**

And the last question is from Lou Miscioscia with Daiwa. Please go ahead.

**Lou Miscioscia**

Hi, thanks for getting me in. And so, my question is really about GPU or accelerated diversification. Obviously, AMD and Intel and others are starting to come out. So, as you think about calendar 2024, when do you think that your system will be for this? And what do you think demand will be and then maybe continue that into 2025?

**Antonio Neri**

Yes. I mean, listen, I'm very pragmatic about these things. Today, in generative AI, the market leader is NVIDIA. And that's where we have aligned our strategy, that's where we have aligned our offerings. And as I made my remarks earlier, we have 10 systems already in the mixed cooling environment and six systems or six offers also in the direct liquid cooling environment. So, we are aligning with NVIDIA today, and that's why you are going to see at Discover, Jensen coming on stage with me to talk about what we're doing together. Now when you go into the sovereign space where there may be some components of supercomputing down the road, obviously, they designed their own systems with our help, and that will be a mixed environment. But just in the last month or so, we opened a new system in the Los Alamos laboratory, which was actually an NVIDIA system with HPE and is direct liquid cooled, so we are clear about that.

But then there are also systems that will come in 2025 that may have different type of accelerators. Over time, we are going to be time to market. But right now, we have aligned to NVIDIA and that's what we're doing. And that's why I think it will be a great opportunity for you to join us at HPE Discover because you can see everything we talked today on the floor. Every

system I just referred to in my slides, every comment I made in my answers to the questions, you can come and see it. These are systems and IP that we are shipping today and you're going to see our time to market with this silicon in addition to all the services in HPE GreenLake.

So, it's going to be an amazing opportunity also because we are doing the keynote at Sphere, we're going to have probably inside the Sphere more than 17,000 customers and partners joining us. So, I know we don't have a lot of...the normal time for questions, but I will say thank you for joining today. As a recap, the quarter was very solid. Our AI system revenue more than doubled to \$900 million, allowing us to obviously exceed our revenue and non-GAAP earnings per share guidance. Because the demand in AI is strong, we have a multiple of...in the backlog, in our pipeline and some aspects of the traditional infrastructure market is recovering, plus our disciplined execution on pricing and in cost, we are raising both the revenue and EPS guidance.

And on the revenue, there could be something more obviously, but it's depending on some of the timing. And then obviously, on the EPS, we are comfortable with the higher end of the range, as Marie said. So again, we are very pleased. We feel the second half is going to set up really well. And then obviously, we are looking forward to seeing many of you at HPE Discover in just less than two weeks. Thank you for your time today.

## **CONCLUSION**

### **Operator**

Ladies and gentlemen, this concludes our call for today. Thank you.