

i4D EVENT SERVICES, INC.

HPE Securities Analyst Meeting 2023

Thursday

October 19, 2023

**Ubiquis
61 Broadway – Suite 1400 – New York, NY 10006
Phone: 212-346-6666 ♦ Fax: 888-412-3655**

HPE Securities Analyst Meeting 2023

[START RECORDING]

[Music]

[Audio starts at 0:28:15]

MODERATOR: Please welcome Vice President of Investor Relations, Jeff Kvaal.

[Applause]

MR. JEFF KVAAL: Good afternoon. Welcome to the 2023 HPE Securities Analyst Meeting. Thanks for joining us in our return to New York. We haven't been here since before COVID. So 4 or 5 years now. It's nice to see so many familiar faces in the audience. I hope for you New Yorkers, this counts as an in-office day... and thank you to the many folks that are tuning in virtually.

The agenda, we have Antonio first. He'll give us a look at how we're pivoting to a higher growth, higher margin, recurring revenues, also positioning ourselves for sustainable and profitable growth.

Through his talk, Antonio will introduce the heads of three of our new segments. So we've got Phil Mottram, from Intelligent Edge, Fidelma Russo in our Edge to Cloud and also Justin Hotard from the HPC and AI. They'll give us deeper dives into each of their own businesses.

We'll then take a break until 4:00 o'clock. Jeremy will come on after. During the break, please give the execs a little bit of space. We'll have plenty of time for Q&A at the end. When Jeremy's presentation wraps up we will have approximately an hour or so for Q&A.

Okay. So with that, I guess, let me mention that once the webcast does conclude that's when we'll be issuing our press release and putting all of our presentation materials on our website. That is at [HPE.com/investor/SAM2023](https://www.hpe.com/investor/SAM2023). Also we'll put a replay up on our website. We'll keep that up there for about a year.

Lastly, let me get to the disclosures. The event may include forward looking statements involving risks, estimates and assumptions. HPE assumes no obligation to update those

statements. And with that, it's now my pleasure to welcome Antonio Neri, HPE's President and CEO.

[Applause]

ANTONIO NERI, President and Chief Executive Officer

MR. ANTONIO NERI: All right. Thank you. Well good afternoon and thank you for joining us today. It is exciting to be back in the New York Exchange. I had the honor this morning to ring, again, the opening bell. It's always fun to be here. Welcome to those of you joining us virtually via the webcast.

Together with my leadership team today, I'm excited for the opportunity to demonstrate to you how much additional value our strategy can create for our shareholders in the near and long term. I'm very proud of the progress we have made in the last five years. The world around us has changed immensely. And so have we.

Today as I prepare to share with you the opportunity we see for HPE in the next three years, I will also reflect on the last three years and what I have demanded of HPE and enterprises around the world. They require vision, agility and resilience from every enterprise around the globe as well as we were all confronted with unimaginable challenges.

In that time period HPE has transformed to become 'the' Edge to Cloud company. We have anticipated what customers will need as they recover from the pandemic which we call the New World Pandemic.

We have offered unparalleled innovation. And we have executed extremely well. Since the beginning of Fiscal Year 2021, HPE has successfully deployed our strategic pivot to higher margin and growth areas of the market which has created strong value for shareholders.

We have delivered 13.4% revenue growth since then compared to the last trailing 12 months with particular exceptional growth in reoccurring revenue. Since we began our strategic as a service pivot in 2019 and started tracking annualized revenue run rate or ARR, we have tripled it.

Last quarter ARR rose to a record of 48%, exceeding a long-term target growth rate of between 35% and 45% which we shared last year at SAM. HPE has grown our annual non-GAAP diluted net earnings per share by 43.5% between Fiscal Year

'20 and the trailing 12 months of the most recent quarter. In Fiscal Year '20, it was \$1.54 and it is now \$2.21.

Our non-GAAP gross margin growth has also been impressive. We have added 310 basis points to our non-GAAP gross margin since Fiscal 2021 and in the most recent trailing 12 months period at 34.8%. Our non-GAAP operating margins improved by 280 basis points in the same period.

We have made this profitability gains as a result of diversifying our business to grow the software and services rich parts of our portfolio while maintaining discipline in pricing particular in our Compute business segment. And we have introduced profitable subscription based cloud services with HP GreenLake which enhances the stickiness of our solutions with our customers.

Greater profitability has also enabled us to significantly increase free cash flow generation in the last three years. From the start of Fiscal Year 2021 through the most recent trailing 12 months, we have more than tripled free cash flow generation. And we are on track to deliver within the range we guided during Fiscal Year '23 Q3 earnings.

Over approximately the last three years, we have strategically focused our portfolio on areas of high growth and high margin. We generated more than 121% total shareholder return as compared to a 37% return for the S&P 500 over the same period.

Importantly, this return represents a CAGR [phonetic] of 31% compared to the 11% CAGR for the S&P 500.

As we produce higher profit and greater cash flow we deliver on our commitment to provide direct capital into our investors. Through our disciplined capital allocation framework we return about \$11 billion to shareholders through dividends and share purchases from the start of Fiscal Year 2018 through our most recent quarter. Obviously that included the spinoff proceeds.

But this capital return benefit represents 122% of our free cash flow over that period. This strong, consistent result would have been admirable in any period but I'm particularly proud of them coming off such turbulent three years. Frankly I believe the market should give us more credit for the

accomplishments because they're a testament to the relevance of our strategy, the differentiation of our portfolio, and the strength of our execution.

We delivered when it mattered the most. And as we continue to execute against these strategies in Fiscal Year '24 and over the next three years, we will continue to accelerate value for our shareholders.

Today I will discuss HPE's compelling market opportunity for sustainable, profitable growth. And at the conclusion of our time together I hope you walk away with a different understanding of a few things.

First how HPE's bet on Intelligent Edge is paying off with a sizeable proportion of our revenue growth segment operating profit coming from this business segment. How we plan to sustain our momentum through time expansion and market share gains.

Second how our Hybrid Cloud business is positioned to become a sustainable center for growth through the scale of HPE GreenLake with investments to make us the number one leading Hybrid Cloud provider.

Third how AI has expanded over our market opportunity to drive what we expect will be profitable growth in the next fiscal year adding to our overall Compute in HPC and AI businesses.

And finally how we continue to make shifts to those areas of our portfolio combined with operational discipline will drive profitable revenue growth for our shareholders.

HPE's strategy is aligned to the significant market trends we see today around Edge, Hybrid Cloud and AI all which will create profitable market expansion opportunities that will help us fuel our growth.

Customers continue to validate our strategy, turning to us to power critical business transformations. And even in this micro economic environment, we continue to see them prioritizing data first, digital transformation initiatives. And those initiatives increasingly include AI which is invigorating today's IT spending.

We anticipate the overall total addressable market across our portfolio will rise by nearly \$100 billion by the end of

Fiscal Year 2022 to more than \$340 billion by Fiscal Year 2026. We will eliminate some of the crossovers because obviously you have some elements of the portfolio that can be added in different segments.

And we intend to capture this growing TAM [phonetic] across the megatrends of Edge, Hybrid Cloud and AI with a mix of HPE's business segments. Our approach will include innovation in key markets where we already have a very strong position as well as customer acquisition, expansion in high margin adjacent markets.

Let's start at the Edge. At the Edge we are expanding our innovation in security, private 5G, and data center networking.

In Hybrid Cloud, we continue to expand in Private Cloud, the segment of the market we anticipate growing at approximately 10% CAGR. We can also capture a larger portion of the hybrid cloud market with our entry into the growing AI-powered IT operations management segment through our recent acquisition of OpsRamp.

And in AI we are well positioned to capture the sizeable growth with our full stack AI native architecture. That includes AI infrastructure growing at a 23% CAGR and AI platform software growing at a 32% CAGR.

Ultimately we anticipate that the AI market where we play growing by roughly 2.4 times to almost \$150 billion by 2026. Clearly there is a great opportunity for HPE in those growth segments.

But let's detail each one today you can emerge with a better appreciation for our differentiated innovation and the strength in each category and our plans to further capitalize on the opportunity ahead of us.

Our Intelligent Edge business is a prime example of what HPE can do to maximize future growth with the right investment at the right time. We have invested nearly \$6 billion in our HPE Aruba networking business organically and inorganically since I became CEO.

It is on track to be a more than \$5 billion annual business in Fiscal Year '23 and beyond. Poised to generate sustained revenue growth and the highest profitability of any of our

business segments. It is a critical part of our business both from a strategic and financial standpoint.

Our portfolio delivers what we call security first, networking from Edge to Cloud to access scale which is very important point that we're going to talk later about that. In Fiscal Year '24, we will build on our successes in our HPE Aruba networking, the core business of campus and branch with a goal of continuing to take share in the enterprise midmarket and NSNB [phonetic] customer segments.

We are expanding our offers in growing markets like security, private 5G and data center networking which we have entered through organic innovation and acquisitions this year. We intend to increase our already high gross margin in this business. By Fiscal Year 2026, we expect our total addressable market at the Edge and networking to grow to about \$94 billion.

Customers will continue to expect a unified networking experience that is secure, automated, and high performing as well as flexible subscription and consumption-based solutions. Our value proposition is absolutely suited to address these needs.

This adds up to the significant opportunity for HPE and for our shareholders. Phil Mottram will now provide more details about our leadership position and priorities in this very important segment. So I would like to welcome Phil.

Phil, please.

[Applause]

**PHIL MOTTRAM, Executive Vice President General Manager,
Intelligent Edge**

MR. PHIL MOTTRAM: Thank you Antonio. So my name is Phil Mottram. And I've been running the Intelligence Edge business now for a little over two years. It's been an incredible journey for me and my team. As you look around this room today and you admire the great decoration and you're scanning the ceiling occasionally, I also want you to pay attention to the Aruba access points that have been installed in the alcoves in the corners. So NYSE made great technology choices and we're proud to be the wireless provider here at the New York Stock Exchange. So nice one, NYSE, if that's a phrase we can use.

If you've been following HPE for the last few years, you've undoubtedly heard Antonio refer to us as the Edge to Cloud Company. The first part is all about the connectivity. And that's where our Intelligent Edge business comes into play.

In the last two years we've seen significant growth in this segment. By the end of '23, we expect the annual revenue of our Intelligent Edge business to increase by nearly \$2 billion from FY '21 levels through consistent growth over that timeframe. The third quarter of FY '23 marked the 11th consecutive quarter of year-over-year growth and this segment represented 20% of the company's revenues.

In addition to the topline revenue growth, we've also seen an increase in profits since 2021. Through intentional efforts towards operational efficiency as well as portfolio management and a focus on margin-rich products, we've increased our profitability. In Q3, we marked a record level of operating profit margin and we accounted for 49% of the company's total segment operating profits.

So we've done this by capitalizing on the high demand in the market for our product and solutions, strategically adding more product areas to expand our TAM [phonetic] and focusing on our delivery mechanisms. We're giving the customers the option for consumption or subscription purchasing through our network service offerings and through HPE GreenLake.

With Aruba Central on HPE GreenLake, we're one of the only companies giving customers the ability to manage all of their network solutions from one single platform. As we go forward we will continue to follow the network in order to identify selling and upselling opportunities from the access points in the ceiling, the campus switch, to the SD-1 network.

Selling these incremental products to our existing customers will continue to drive revenue in core business segments. We'll also be following that work even further into the data center, integrating it with security and adding private 5G and delivering it all through the same unified cloud native experience.

Additionally, we're weaving sustainability into everything that we do. This is something we feel is vitally important and we've heard loud and clear that it's a priority for our customers as well. To give you an example of this, our asset lifecycle management program within HPE Financial Services

reclaims and refurbishes used assets on our customers. And over the last three years, we've given more than 8.2 million devices a second life ultimately reducing mounds of material that's directed towards landfills.

Bringing all of this together with confidence that we have the right strategy in place to drive sustainable, profitable growth and deliver value for our shareholders.

As we look ahead we believe that trends in the marketplace around hybrid work, security, and data center requirements will create new opportunities for us. And while there are some headwinds related to the overall networking market growth projections, we expect to continue to sustain our Intelligent Edge business segment performance through TAM expansion, continued market share gain in key segments.

We continue to innovate and enhance our offering and receive validation from leading industry analysts in this regard. In the Edge networking space, we're targeting TAM of about \$94 billion by 2026. And that's one and a half times the TAM that we pursued in 2022.

And within that I see four distinct areas of growth for our business. The first area is security. We believe that the networking market is converging with the security market under a framework called SASE, or Secure Access Service Edge. And we're not alone in that belief.

Gartner believes that 60% of enterprise customers will buy SASE from a single vendor within three years. SASE has five key elements. And we already have two of them with networking and firewalls. The acquisition of Axix Security earlier this year allowed us to add the missing three elements of SASE: namely ZTNA, SWG, and CASB.

Within HPE networking, HPE Aruba networks, sorry, we've always been great at protecting things on the network. Adding Axis Security now gives us the ability to be able to protect people. So we can protect people and things across multiple network devices all from one location on HPE GreenLake.

We have a compelling and highly differentiated offer and we're poised to gain share in this high growth market that will represent a \$3 billion TAM by 2026.

The second area of growth for our business is in private 5G. So the rollout of 5G has advanced mobile networks around the world. And in many countries governments now assign 5G spectrum to be used by enterprises privately which is what is referred to as private 5G.

We believe that private 5G and Wi-Fi will exist side by side and enabling new customer use cases where outdoor coverage and latency is important such as in ports, defense applications, and other scenarios like mining operations and sporting events.

Indeed, we recently provided private 5G and Wi-Fi 6E connectivity for the Ryder Cup in Italy. So we provide all ways of connectivity to 250,000 spectators whilst also delivering secure, private network for the operations staff across the hundreds of acres that made up the venue. So hats off to us for that.

So private 5G represents a significant opportunity for HPE to enable a unified customer experience with security and policy management. Additionally this represents a significant opportunity to HPE to capture market share in this space at healthy gross margins.

The third area of growth for our business is in data center networking which is an area that will grow to a \$19 billion TAM by 2026. Today our HPE Aruba networking CX switch platform is already sold into many data centers around the world.

We've spent 12 months enhancing the capability of our switching products and have a 2-year roadmap that is highly differentiated through our HPE Aruba networking CX operating system which will enable enterprise customers to deploy cloud native networking in their data centers or Co-lo facilities.

We expect to continue to capture market share from existing customers, especially those who want to extend the management ease of Aruba central into their data center.

The fourth area of growth for our business is in NaasS, or Network as a Service. The NaaS market is expected to a \$3 billion opportunity by 2026 which is good news. But the better news is that we already have a clear lead in this market. With our customers we can take them on a journey to NaaS as opposed to some of the more recent startups in this

space which require customers to rip and replace all of their network technology in order to deploy NaaS.

Our intelligence services capability, or the managed services piece of NaaS, offers a 98% reduction of network events through automation and utilization of AI, ML-driven analytics. NaaS also represents an important way in which we're delivering on our commitments to sustainability because NaaS creates a circular economy by extending the life of assets for customers.

For these reasons and more, we continue to see demand from customers for NaaS solutions. Today we have a growing number of customers across different verticals including retail, hospitality, higher education, and manufacturing. And some of the customers who are already consuming our NaaS services, the Home Depot, CarMax, and KPMG.

And now in addition to large and global customers, we're now scaling it into the commercial and midmarket segments through and with our channel partners.

So those are the four areas that we expect to see significant growth in our business. Aggressively pursuing these growth opportunities while continuing to capture revenue in our core segments is how we will drive sustained business momentum. And our ability to do that can all be put down to one key point: customers buy HPE Aruba networking because we are customer centric.

Aruba Central which is our HPE GreenLake, cloud controlled play is the industry's first AI-powered, cloud native architecture designed to connect, detect and automate the Edge. This gives customers a simple but powerful management tool to drive efficiencies in our business.

And the three core layers of our portfolio each deliver specific benefits to customers. First, our unified infrastructure helps eliminate operational silos and streamline operations. We do this by bringing together wired to wireless and SD-WAN technologies across Campus, Branch, Remote Worker and Data Center locations into a single cloud native management and orchestration tool. And that's Aruba Central.

Second, our Edge to Cloud security simplifies network security through a combination of built-in, zero-trust

[phonetic] features focused on users, devices and IoT and assess the integration framework that combines One Edge functions and integration with third party cloud-delivered security services.

And third, our AI and automation software enables our customers to resolve issues quickly, ideally, before they impact the business while also helping IT departments operate more efficiently.

We have a robust portfolio of products and solutions that deliver what our customers want and need to harness the power of the Edge. And while we've seen fantastic growth in recent years we are not standing still. We have the right strategy in place and we will continue to accelerate shareholder value by sustaining growth in core segments, expanding our overall TAM, and gaining market share in key, high growth, margin rich areas.

Back to you Antonio.

[Applause]

MR. NERI: HPE is a pioneer in hybrid cloud which represents our largest market opportunity. Recognizing the strategic importance of this market and our differentiated value proposition led by HPE GreenLake, last month we announced we will create a new Hybrid Cloud business unit at the start of Fiscal Year 2024. Establishing this new segment reflects the maturity of our offering and sets us apart from our peers.

Our Hybrid Cloud business unit will bring together into a single Hybrid Cloud, segment of Storage and Compute as-a-service. All the offers related to that include our HPE GreenLake Private Cloud and software solutions. Our updated operating model will enable greater focus and efficiency, faster execution, and a superior cloud native experience for our customers and our partners.

Our simplified operating model also incorporates and self-transformation to enhance the execution between business units and our go to market function. Through these changes we will engage customers more directly on the offerings they want the most.

HPE GreenLake is at the heart of our hybrid cloud strategy and offering its growth truly remarkable. At the end of Q3, we support 27,000 unique customer logos and 3.4 million

connected devices. Today more than 1,100 partners transact HPE GreenLake Edge and Hybrid Cloud offerings showcasing HPE GreenLake's incredible cloud strength and scale.

HPE GreenLake also plays a key role in expanding our gross margin with software and services comprising and increasing portion of our ARR mix, now nearly at 70% at the end of Q3.

Fidelma Russo, our Chief Technology Officer, has been leading our HPE GreenLake cloud platform development and will take on the responsibility of accelerating our Hybrid Cloud services growth opportunity when she becomes the general manager of the new Hybrid Cloud business segment on November 1st.

I would like to invite her to talk about how it will enhance our leadership position in this growing market. Fidelma?

[Applause]

FIDELMA RUSSO, Executive Vice President Chief Technology Officer

MS. FIDELMA RUSSO: Thanks Antonio. Good afternoon everyone. I'm honored to lead our future Hybrid Cloud business segment which I believe will be a large and enduring growth engine for HPE and will create long-term shareholder value.

Enterprises across the world are moving from a Public Cloud first approach to a hybrid approach. And they've come to the realization that as their data grows and becomes more distributed in nature it is neither scalable or nor economical to rely solely on the Public Cloud. The emergence of generative AI has made this reality even more clear. And this creates a tremendous opportunity for us.

Our new Hybrid Cloud business segment accelerates our participation in several large and attractive markets across Storage, private cloud, and infrastructure software. Combined, we are looking at a \$160 billion-plus market that is highly profitable. And as enterprises continue to embrace hybrid, this market will see robust growth for the future.

To capitalize on the market opportunity, our strategy is to take share in Storage, scale private cloud with the momentum we have with HPE GreenLake and expand into infrastructure software.

Ultimately this strategy starts and ends with our customer. We offer the most differentiated customer value proposition

among our competitors including cloud providers. We have a growing portfolio of market-leading offerings and we are seeing strong customer adoption. As Antonio said, we are a pioneer and leader in Hybrid Cloud with HPE GreenLake. And the reason we've been successful, and the reason we've been able to differentiate versus the Public Cloud, is because when we show up to our customers, our mission is to solve three of their biggest challenges. Transform their business through the power of data, modernize their IT infrastructure with a true cloud experience, and dramatically simplify the operations of their multi-generational and multi-cloud IT estates.

Let me give you some more color on how we address each of these customer challenges, and how different aspects of our portfolio come together to drive an integrated value proposition through the power of our GreenLake platform. Everything starts with data, which is the most valuable asset companies have, and it is at the center of any digital transformation. Its gravity and distributed nature are what drive enterprises to embrace the hybrid cloud. And this is why Storage is foundational to our hybrid cloud strategy. With our HPE Alletra line of Storage products, we offer customers a value proposition none of our competitors can match. We have engineered a truly modern, scalable platform with unified support for block and file that is capable of meeting the most demanding requirements for AI. It has been engineered for a true hybrid cloud experience. And it's made possible with an architecture that's built from the ground up to be cloud native and software defined. We continue to deliver industry leading uptime and resilience, enabled by our market leading AI ops capabilities. And since its launch, HPE Alletra has seen the fastest customer adoption of any Storage platform in HPE history.

In addition to our product leadership, we are also increasing investments in our go to market, both in our direct sales force and in the channel. And with these investments, we are poised to take share from our competitors. But not only are customers trusting us with our data, they are also trusting HPE to build, and in many cases, operate their hybrid cloud infrastructure on top of their data, through HPE GreenLake. Our entire strategy around HPE GreenLake is built on three foundational beliefs, which differentiate us from the Public Cloud. First, customers should not have to choose between

the agility of the Public Cloud and the performance and control of their private infrastructure. Second, customers want choice and freedom from lock-in. And third, they operate multi-generational, multi-vendor IT estates that need to coexist with the Public Cloud.

Over the last five years, we have built a multibillion-dollar HPE GreenLake franchise, and it is the envy of our competitors. And we are continuing to double down on that success. This year, we added a host of new cloud native offerings and capabilities, including hybrid multi-cloud orchestration for virtual machines, containers, and bare metal, a full suite of private cloud offerings that enable customers to self-manage or choose a fully managed experience, and a portfolio of world class AI infrastructure delivered as a service. The portfolio clearly resonates with customers, and it is allowing us to extend our market leadership. Now, we understand that the adoption of hybrid introduces several operational challenges for our customers. And through a curated strategy of organic investments and acquisitions, we have built a compelling set of SaaS offerings aimed at helping customers simplify the data management and protection of their hybrid multi-cloud environment while reducing the risk of Public Cloud lock-in.

So we have focused on three critical customer needs, AI powered monitoring and observability for day two operations and beyond through our acquisition of OpsRamp, unified data access through our HPE, Ezmeral data and analytics suite, and that helps customers move and transform their data for use in AI and other applications. And then data lifecycle management and protection through our suite of offerings, including Zerto Disaster Recovery. All of our software is built for hybrid to give our customers the flexibility they need to simplify their hybrid cloud operations at any location. Our offers support multi-vendor, multi cloud environments, and enable native integrations for both HPE GreenLake and major Public Clouds. Our strategy with our HPE GreenLake SaaS offerings is to drive aggressive penetration across our customer base. All of our Hybrid Cloud offerings, Storage, private cloud, and SaaS, as well as HPE Aruba Networking, and HPE Compute offerings, are natively delivered through our HPE GreenLake cloud platform. And we will deliver our AI offers natively on the platform over the next couple of years.

The platform enables our customers to have a consistent cloud-based management experience across all our offerings. It provides a set of essential services, including consumption analytics, and a sustainability dashboard that helps customers understand and reduce their carbon footprint. And as more customers adopt our cloud platform, they are reaping the operational benefits of a truly unified cloud-based management model, which reduces upfront CapEx and ongoing OpEx running costs. And the more HPE offerings our customers consume through HPE GreenLake, the bigger the operational benefit they gain, and the higher the value we have generated for our shareholders.

So to bring our differentiation to life, I'd like to share a couple of examples of how our customers are leveraging HPE's hybrid cloud portfolio to solve their business challenges. For those of you who have a great time watching football, you will be glad to know the Dallas Cowboys were looking for a solution to allow their players and coaches to watch and analyze video footage across different locations. And our high-performance Alletra Storage solution with multi-sized replication perfectly suited their needs. Danfoss loved the manage for you aspect of our private cloud solution. And it allows them to focus on delivering business innovation globally while reducing their carbon footprint. So there's no question that HPE GreenLake is winning in the market, and hybrid cloud is the driving force behind it. We are seeing tremendous adoption with our customers and partners, with more than 27,000 unique customers using the HPE GreenLake platform. We are creating significant shareholder value because we are expanding our infrastructure offerings into new and higher growth markets. Our recurring revenues are growing at more than a 35% CAGR, which is faster than the Public Cloud. And we are delivering greater profitability through our software and services rich portfolio.

In closing, HPE has a winning Hybrid Cloud strategy. We have a highly differentiated value proposition that is grounded in a deep understanding of our customers evolving needs with respect to data, infrastructure modernization, and hybrid cloud operations. The differentiation comes from being a pioneer, and as a service with HPE GreenLake, our unrivaled HPE edge to cloud portfolio, and it's further enhanced by years of curated M&A, organic investment in our SaaS portfolio and our HPE GreenLake platform. Value proposition

is clear really resonating with our customers, and our continued success will fuel sustained growth and profit expansion for our shareholders. Now back to Antonio. Thank you.

MR. NERI: Well, thank you for that. I'm excited to have Fidelma now lead this great opportunity we have in hybrid cloud. Just as we have a very differentiated hybrid cloud value proposition, we also uniquely positioned with a compelling High Performance Computing and AI offering, as this market dramatically expands. The markets extreme acceleration in the last year and anticipated growth over the next several years reflects the enterprise realization that they must embrace AI or they will lose their competitiveness and get left behind. As organizations lean into AI, they are discovering a few things. First, the data intensity of the AI workloads require a hybrid by design solution, rather than a cloud only approach. They need a solution across the entire AI lifecycle, from training, to tuning, to inferencing. And given the energy and datacenter services required for large scale AI workloads, sustainability must be built into the technology from the start.

HPE addresses these needs with an end-to-end portfolio designed for the full spectrum of use cases, spanning large scale AI model development, training and inferencing, as well as unique liquid cooling data center services expertise. Customers are attracted to HP's market leading supercomputing capabilities, differentiated network and interconnect IP, AI specific software, and services expertise. This capabilities position us favorably in a rapidly growing, total addressable market, comprising HPC and supercomputing AI infrastructure and AI platform software. Our pursuit of an outsized share of this market opportunity will deliver real value to shareholders, especially when—with our path to AI driven profitability charted in the next fiscal year.

Our investments in our unique silicon and software are expected to drive a larger profitable revenue base. I would like to welcome Justin Hotard, who leads our HPC and AI segment to discuss how we'll plan to capitalize on the boom in AI opportunity. Justin.

JUSTIN HOTARD - Executive Vice President, General Manager, High Performance Computing, AI & Labs

MR. JUSTIN HOTARD: Thank you, Antonio. AI is driving the next wave of investment in innovation, resulting in value creation in the IT market and the global economy. It will have a transformative impact, similar to what we saw with web, mobile and cloud. And we are seeing demand shift dramatically as our customers realize the potential of AI to deliver business transformation. Today, I want to make sure I cover three points with you. First, the AI market is growing rapidly and HPE is well positioned to profitably capture that market expansion. Second, we have a right to play and win in AI, and we're already winning in the market today. You already are seeing this in our sales pipeline, in our orders, and our revenue. Third, we deliver solutions for the AI lifecycle across training, tuning, and inferencing that will drive higher growth and margin expansion.

AI adoption is fueling significant market growth. We're targeting three areas of the AI market, Supercomputing, the AI Infrastructure, and AI Software platform. As you can see, our total addressable market for these segments will grow at nearly 24% CAGR to \$146 billion by 2026. By focusing on these areas, we are empowering customers to transform their businesses to dramatically enhance productivity, accelerate innovation, and create new revenue models. Let me cover each in depth. First, supercomputing. HPE is a market leader in super-in delivering the world's leading supercomputing and high-performance computing solutions. We hold the number one position of performance share for the world's 500 fastest supercomputers and deliver the majority of the world's top 10 most efficient supercomputers. Our leadership positions us to continue to capture share in the markets forecasted double digit growth. We will continue to invest in our supercomputing technology and leverage our global service delivery footprint.

The second is the AI Infrastructure market. By 2026, this market is expected to be seven times larger than our core supercomputing market and reach \$86 billion. A significant amount of this growth will be driven by AI model training, which is one of the most computationally intense workloads of our time. Customers require powerful infrastructure capable of the scale and performance that supercomputing technologies deliver. For this reason, key features of supercomputing

technologies, including management software, networking, and liquid cooling, will become requirements for AI infrastructure. These are technologies where HPE has unique intellectual property that positions us to deliver differentiated solutions to customers.

Third, we have an opportunity above the infrastructure layer in the AI software platform market. This market is expected to grow to \$49 billion by 2026. And we plan to capture share through continued investment and innovation. We will extend our current software suite and invest in new offerings to specifically target market opportunities across the entire AI lifecycle. As this portion of our portfolio grows, we anticipate healthy margins that will be accretive to our business. We believe our strategy to focus on these three segments across the AI lifecycle of training, tuning, and inferencing, will lead to increased market share in two ways. First, through our core supercomputing business and trusted global brand will create new opportunities and expand our share of wallet within our existing customer base. And second, by further building on these core offerings with purpose-built AI solutions and HPE's global presence, we expect to attract new customers to win even more market share. In fact, it's the combination of the market opportunities across supercomputing, AI infrastructure, and the AI software platform that differentiates HPE from the competition.

Now, let me transition and cover how the growing AI market will be driven by Compute and data intensive workloads, and we'll need an architecture that HPE is uniquely positioned to provide. This AI native architecture is different from the cloud architecture. Traditional cloud architectures are optimized to run multiple workloads on a single server. They were not designed with AI in mind. There are four key attributes of an AI native architecture. First, the architecture should be designed with computing and infrastructure that is at the scale of supercomputing to take full advantage of computing capacity. This is essential to train, retrain, and tune AI models, using large quantities of data with efficiency, speed, and accuracy to accelerate time to value. Second, ensuring a truly hybrid design provides the flexibility to train, tune and deploy AI models in any environment. Customers need the flexibility to integrate data that exists in the Public Cloud, the private cloud, and

the edge. For example, enterprises need to deploy inferencing where they can deliver real time insights. This will be necessary for large language models, as enterprise use cases scale, just as it already is for Compute vision applications across autonomous driving and medical imaging.

Third, an open hardware and software ecosystem deliver strong advantages. Having an open AI ecosystem expands training and inference market solutions to accelerate adoption for customers and support ongoing innovation in the broader ecosystem. And fourth, these architectures must be sustainable by design. By 2028, it's estimated that AI workloads will grow at a 35% CAGR and require about 20 gigawatts of power within data centers. Customers will need an architecture that satisfies this demand at a new level of energy efficiency to minimize the impact of their carbon footprint. HPE is uniquely positioned to deliver this new AI native architecture through our technology differentiation and expertise.

For many years, we've made strategic investments in AI, and led AI focused research in Hewlett Packard labs. As a part of this, we have a powerful blend of strong technology, intellectual property, expertise, and talent to deliver an AI native architecture. This positions us to be a long-term market leader and to capture market share in the nearly \$150 billion AI TAM I previously highlighted. Our multiyear investment plan targets three distinct areas to enhance our differentiated position in the market that will set us apart from the competition. These include software, high performance networking, and supercomputing infrastructure. We expect these investments to significantly expand our HPC and AI business segment operating margins by 2026.

First, our machine learning development platform is proven to train quickly, efficiently, and at scale to create reliable and accurate models that can make valuable predictions and reduce business risk. With our current software and the investments we're making across the AI lifecycle, we will deliver a robust platform of open-source software. This software stack also integrates with the HPE Ezmeral Data Fabric that Fidelma covered in her overview. By integrating the two platforms, we're making it easy for customers to manage their entire data lifecycle.

Second, we deliver ethernet-based, high-performance networking with HPE Slingshot, which is purpose built and proven for significant AI scaling. This technology currently powers Frontier, the world's fastest supercomputer, with nearly 40,000 GPUs, at the United States Department of Energy's Oak Ridge National Laboratory. HPE Slingshot will also power the upcoming Aurora exascale supercomputer, which features more than 60,000 GPUs at Argonne National Laboratory. HPE Slingshot connects all GPUs to operate as one single large supercomputer. This makes it possible to train trillion parameter AI models in one single instance. HPE Slingshot is built upon the ethernet standard, and it already supports a variety of GPUs and accelerators, including those from Nvidia, AMD, and Intel. This creates market options for customers' training needs as the market continues to scale.

And finally, to support the reliability and resiliency of AI systems required for large scale training, we offer powerful and integrated supercomputing solutions to address these needs. These include our expertise in system integration and our supercomputing software. As a part of our supercomputing infrastructure, we have extensive experience in R&D that focuses on sophisticated liquid cooling solutions. Liquid cooling is essential to deliver sustainable datacenter infrastructure for AI. Our solutions can drive up to a 20% performance improvement per kilowatt over air cooled solutions and consume 15% less power. Further, we believe next generation accelerators will require liquid cooling in every system to meet power and thermal demands. Our intellectual property in liquid cooling infrastructure and our advanced manufacturing capabilities in this area position us well to capture this demand.

However, technology alone is not sufficient. Customers training AI models don't have the time to develop the expertise to operate their AI infrastructure and platform. This is an opportunity that positions us to offer unique innovations like a virtual private cloud for AI with HPE GreenLake for Large Language Models that we announced in June. This is a turnkey cloud service that integrates our AI infrastructure and AI software platform for model training, where we see significant global demand. The service is also designed to lower carbon footprint.

As I covered, our intellectual property in AI is proven and it positions us well to accelerate innovation in the market. This is why customers from various verticals are turning to us. I'd like to point out a couple of examples that illustrate how our AI focused solutions are making a transformative impact in the market today. We have a strong partnership with the US Department of Energy and its national laboratories to co-design and co-develop powerful supercomputers, as we have with exascale, to accelerate national initiatives in AI driven science and innovation. Argonne National Laboratory, for example, is creating a series of generative AI models at the exascale level that will be trained on HPE supercomputers.

In the pharmaceutical industry, Recursion Pharmaceuticals, a leading tech bio company, uses advancements in AI to accelerate and industrialize the discovery of new drugs. Recursion uses the HPE machine learning development environment to manage its large-scale AI training jobs on its AI supercomputer. This software significantly speeds up model training across more than 25 petabytes of biological and chemical data and improves team collaboration. Customers like Taigo Cloud and Crusoe Energy partner with us to leverage our expertise and global supercomputing leadership to deliver full stack solutions that integrate our industry leading infrastructure, open-source software, and on-site services.

And our supercomputers and AI software are also behind the work of foundation model builders that create prebuilt models used by broader enterprises. For example, Aleph Alpha has built a powerful Large Language Model in five languages using HPE supercomputers and machine learning development platform. This large language model is already used today as an advanced AI digital assistant across banks, automotive manufacturers, legal firms, to accelerate business outcomes. They turned to HPE to scale their training capabilities in a virtual cloud, versus relying on deploying and managing their own supercomputer and software. As our launch customer for HPE GreenLake for Large Language Models, Aleph Alpha will extend their customer reach globally and accelerate their growth. As this business scales, the combination of top line growth and higher margin offerings within our portfolio will help us achieve double digit operating margins.

AI is a fast-growing market that will be fundamental to transforming businesses and accelerating economic growth. We believe we are well positioned to profitably capture the market opportunity and win in AI to drive higher growth and margin expansion. With our trusted expertise, differentiated IP, and long-term sustained market leadership, we can capture significant value from the AI market. This is an exciting opportunity for our company, and we look forward to fueling the value creation from AI for our customers, and in turn, for our shareholders. And now, let me turn it back over to Antonio.

MR. NERI: All right. Well, thank you, Justin. Each of these areas have—we just presented, provide our customers and our shareholders with tremendous opportunity. And the engine that powers our ability to capture that opportunity is Compute, which produces cash flow to invest in our business and deliver direct capital return to our shareholders. As you all know, there is a cyclical nature to the Compute business. Over the last several years digital transformation drove increased investment to modernize infrastructure. Now customers are focused on digesting those investments. We will be very intentional about how we execute in Compute during this cycle to maintain our scale and industry leading profitability. We are focusing on capturing every unit while maintaining balance in our operating margin performance. We are capturing opportunities in a steady Compute market, from cloud repatriation from edge and IoT workloads, the amount from the telecom and 5G sector, and service providers, and from the emerging needs from AI inferencing solutions. Justin talked about that.

We expect continued demand next year for service with now GPUs and other Compute accelerator types. Combined with our shift to HPE ProLiant Gen 11 servers, which deliver significant greater performance compared to the previous generation, we anticipate a tailwind in Compute average unit price in fiscal year 2024. We measure our progress in part to tracking the servers we sell with accelerated processing units, or APUs. That is because customer can use a variety of Compute accelerators beyond just GPUs to support AI workloads. We have a complete Compute portfolio to address the entire AI lifecycle across training, tuning, and inferencing. We also continue to see strong interest in our HPE services, which is additive to both revenue and margin.

Our world class HPE services team helped us to deliver great customer experience, and customers can gain maximum benefit from HPE GreenLake, from our software designing and building their infrastructure, and running their hybrid IT estates with one exceptional end to end customer experience.

Just as Compute serve as an engine to power accelerated growth, HPE Financial Services continue to be a critical competitive advantage for HPE. This business creates smarter IT lifecycle solutions for our customers and partners through offerings that combine insights, financial expertise, and deep-rooted focus on sustainable IT. It is also strategically important for us as we increase our other service business to HPE GreenLake. Customers can transition to HPE GreenLake using our asset Lifecycle Management services to assess-to access efficient technology and cloud consumption models, creating value for our business.

Going forward, we see even higher demand from our customers as they put more emphasis on finding ways to accelerate their sustainability goals through our services and the - - economy solutions. For our investors, HPE Financial Services offerings and its best-in-class return on equity, provide a great source of profit that expand our earnings. In addition to driving impressive organic innovation across our portfolio, we continue to be opportunistic in making the right acquisitions. So far this year, we have acquired five businesses to accelerate our strategy and enhance our capabilities. When we assess opportunities to make organic and inorganic investments, we focus on the potential to drive higher level of recurring revenue and profitability. We maintain particular interest in investment that help us innovate and grow profitably at the edge, in Hybrid Cloud, and in AI. We will continue to be opportunistic in making beneficial and creative acquisitions, while following our disciplined return base framework and ensuring integration success.

Jeremy will address more specific on the long-term sustainable value creation and free cash flow generation objectives we pursue on behalf of our shareholders. When I reflect on HPE's strong business momentum, I'm particularly proud that a great deal of that strength come from solution to help customer advance their important environmental, social, and governance objectives. In fact, in fiscal year 2022 alone, we generated about \$1.3 billion in net revenue

that we can directly attribute to sustainability engagements with customers. In addition, our portfolio enhances opportunity for our customers to achieve their environmental goals. I am very proud that HPE is one of the only two global IT companies to have a net zero target of 2040 or sooner approved by the science-based target initiative.

Last year, we reduced emission directly within our controlled by more than 1/5 from the 2020 baseline and reduced our overall carbon intensity by 2% year over year. We also believe a diverse and engaged workforce fuels innovation and performance. Since 2017, our employee engagement score has risen 20 percentage points to 83%. And while we have work to do, and so the entire IT industry, I am proud of our progress in diversifying our workforce. Our US workforce is 32% ethnically diverse, and more than 26% of executive positions worldwide are held by women at HPE. We also have a very diverse board of directors. Half of our independent directors are female. And we are innovating with a strong ethical compass. We win the right way.

In the last year we have put in place even more comprehensive governance around AI to enable us to size the opportunity this technology possesses in the right way. We have the right strategy aligned with the key market megatrends. We have the right team with strong focus on delivering business outcomes for our customers. And while the world is navigating uncertainty, we are confident HPE will continue to accelerate value for our shareholders. To speak about shareholder value creation, Jeremy Cox will take the stage after a short break. As you know, Jeremy is serving as our interim Chief Financial Officer while we continue our internal and external search for a permanent CFO. And we will be back onstage as soon as the market close. Thank you for your time and attention.

[Break 01:35:49 - 01:55:56]

JEREMY COX, Senior Vice President and Interim Chief Financial Officer

MALE VOICE 15: Please welcome Senior Vice President and interim Chief Financial Officer, Jeremy Cox.

MR. COX: All right, well, good afternoon and welcome back. I hope you guys enjoyed your break. You know, as an 18-year veteran of HPE, it's a privilege to be serving as the interim CFO. And I'm very happy to be with all of you guys today. Now that our leaders have presented our vision, strategy, and differentiation, let me translate that into our financial thesis and our commitments. We'll demonstrate how pivoting our mix to higher growth, higher margin, and recurring revenue is accelerating value creation for our shareholders. My main goal today is to help you understand how our strategy will deliver additional value through sustainable profitable growth and increased capital returns. We're pleased with our progress throughout FY23. We've managed through some ongoing macroeconomic challenges, yet our business is performing well. We're reiterating our Q4 guidance for revenue of \$7.2 to \$7.5 billion and our FY23 guidance for revenue growth of four to six percent in constant currency. We expect currency to be approximately a 300 basis point headwind for the year.

We're also reiterating our Q4 and FY23 non-GAAP EPS guidance. We continue to expect Q4 non-GAAP diluted net EPS of \$0.48 to \$0.52 cents, and FY23 non-GAAP diluted net EPS of \$2.11 to \$2.15. We also reiterate our free cash flow guidance of \$1.9 to \$2.1 billion. Given our business performance and the scale of opportunity in front of us, we have chosen to make some targeted investments in Q4, to accelerate our pivot and Intelligent Edge, Hybrid Cloud, and HPC and AI, which we've offset with higher OI&E than we originally expected. We've also incurred some additional GAAP expense within our cost optimization plan, including certain real estate charges, which we don't expect to be meaningful in future periods. Consequently, our full-year GAAP and non-GAAP operating profit will be slightly lower than our guidance, with non-GAAP operating profit growth to be approximately four percent, versus our prior guidance of six to seven percent. FY23 GAAP EPS is now expected to be \$1.35 to \$1.39. We are accelerating these investments, and yet remaining within our long-term financial framework, where we continue to expect non-GAAP operating profit to grow faster than revenue over our outlook period.

So our strategy at HPE is definitely accelerating value creation for our shareholders. We are improving our growth and margin profile by shifting our mix towards higher growth and higher margin segments. We're adding as-a-service

software and service revenue across all of our segments, and our capital allocation strategy balances investments to drive further long-term revenue and growth in our free cashflow and consistently returns capital to our shareholders.

So here's how you can expect to see this strong mix come through in our financials in the next three years. We expect our growth businesses, our Intelligent Edge, Hybrid Cloud and HPC and AI segments will contribute increasingly more revenue to HPE, and the combination should exceed 50 percent of our total segment revenue by FY26. Already for FY23, they should be more than 45 percent of our total segment revenue. That's an increase from approximately 40 percent in FY22, and our operating profit trajectory is even more dramatic. We expect a mix of total segment operating profit from these segments to reach over 60 percent in FY26, and that continues our recent trend. We expect to see our growth businesses at approximately 45 percent of total segment operating profit in FY23, driven largely by improvements in the Intelligent Edge, compared to under 30 percent in FY22.

HPE Green Lake is a key element of our portfolio mix shift, with healthy customer demand. Continued ARR growth means our as-a-service products will represent more than a significant percentage of overall revenue, and at richer gross margins. We are rolling forward our long term ARR target CAGR by one year through FY26, and thus reiterating our commitment to 35 to 45 percent CAGR growth. This puts ARR on track to more than double and reach approximately 10 percent of our revenue by FY26.

In our new segment structure, we'll continue to disclose ARR in the same way. ARR growth is set to drive further gross margin expansion for the company. Gross margins for our as-a-service offerings are already meaningfully higher than the same products and services sold through a transactional capex model, and as we build momentum in standard offerings in SaaS, our mix of high margin software and services within our ARR has risen to 61 percent in FY 22 to 68 percent in the most recent quarter, and should reach the mid to upper 70 percent range in FY26.

So let me step back and discuss how that rising mix of software and service revenue appears in an HPE Green Lake deal. This Compute and Storage HPE customer opted to purchase HPE Green Lake in lieu of a traditional capex deal. The deal

had a normal profile of revenue growth, margins and term length. Customers in a transactional model typically overprovision hardware capacity and purchase at a lower amount of revenue from software and services attached. We have normalized the software and services revenue in this deal example at \$100. This HPE Green Light customer was happy with the flexibility to purchase less capacity to start, though with more than twice as much software and services revenue. These additional services delivered an improved experience in operational savings for the customer. Over time, rising usage led the customer to increase its spending with HPE, including three times as much software and services revenue at a similar hardware level. The Green Lake model meant we received more high margin revenue over the deal life cycle, and it created a happy customer who's likely to stick with us, as our annual customer churn is only three percent. Along with higher software and services contribution, the rising usage also pulls through additional hardware, and we get another bite at the apple when the customer decides to replace their hardware. Our HPEFS management cycle business enables us to repurpose the fleet of depreciated assets into other customers. There, this lowers their cost and thus increases our margins. So the bottom line, HPE Green Lake drives a better value for our shareholders through higher revenue, margins, and customer retention for HPE.

And our end markets and product sets are evolving. Our new segment structure, effective November 1st, will improve our portfolio alignment and accelerate our go-to-market motion. We'll report out Q1'24 results under the new structure, and prior to Q1 earnings, you can be assured we plan to file restated historical segment financials for the new reporting structure. While there are several product lines and shifts between existing segments for various optimization purposes, most of the changes are not overly material to the financials of the existing segments. The more material change is establishing the Hybrid Cloud segment, which Fidelma spoke to earlier. It combines the existing Storage segment and the as-a-service Compute business, and the software business previously reported in our Corporate Investment and Other segment.

So now let me take you into a deep dive of our current segments. You've heard we expect megatrends in Edge, Hybrid Cloud, and AI to drive our TAM to rise higher than \$340

billion by 2026. In the coming slides, I'll show the financial goals of each of our segments. We'll also provide a table of TAMs by existing segments in tables for your reference.

The Edge and Networking TAM is growing at a 10 percent CAGR to \$94 billion by 2026. The portion addressable by our Intelligent Edge segment is growing in the mid-single digits, and our business is performing well above this. We deliver 53 percent year-over-year constant currency revenue growth in Q3'23 and expanded our operating margin by more than 1300 basis points to 29.7 percent. Last year, at SAM, we talked about the mid 20 percent range by FY25. Well, we're ahead of schedule on operating margin. However, we're keeping this mid 20 operating margin for the segment for the outlook period. Improving supply is now allowing us to make progress against our order book and deliver on deals we won in prior periods. We're confident that the combination of shared gains and the expanded TAM that Phil mentioned will allow us to continue to outgrow our end markets. We expect revenue growth to be in the low double digits from FY22 to FY26, a figure that's front-end loaded, given our growth this year. And I also want to highlight our as-a-service business in the Intelligent Edge, such as Aruba Central and Security, because it's now both sizable and growing rapidly. As growth here continues, we will see more of a tradeoff between current period revenue and future period revenue. Such strong growth in FY23 sets a high bar for comparison in FY24. Even so, we expect FY24 revenue to be slightly up relative to FY23. We're forecasting a continued tailwind to revenue in the first part of FY24 from our significant order book. Our HPE Green Lake and Storage portfolios are well positioned to capture the TAM growth we expect to see in the Hybrid Cloud, where we expect to see the TAM to grow at a 7 percent CAGR to \$164 billion in 2026. We estimate the TAM for Storage alone will grow at a 3 percent CAGR to reach \$66 billion in 2026. We believe we'll grow at market rates from FY22 to FY26. We haven't included significant AI contribution in our growth outlook, but we do believe AI will pull through Storage demand.

Our revenue driver in Storage will be ongoing strength of our HPE Alletra product line which has grown in triple digits pace over the last five quarters to numbers that are no longer small.

Also, as-a-service is now the fastest growing portion of the business, and like Intelligent Edge, trades some current period revenue for future period revenue. Our SaaS business growth is one reason we expect operating margins to improve to the mid-teens range in the outlook period for this segment, which will be evident in the operating margin structure of the new hybrid cloud segment over the same period.

We have a significant opportunity in HPC and AI with AI, as we're seeing great momentum in the overall HPC and AI segment. We expect the AI TAM to grow at a 24 percent CAGR to \$146 billion by 2026. We estimate the HPC and supercomputing TAM, which this segment addresses directly, will grow at a 22 percent CAGR to \$78 billion in 2026. So, we talked about AI activity after Q1. We talked about wins after Q2 and order book after Q3, when HPC and AI order book rose to more than \$3 billion. We recognized a modest amount of AI revenue in Q3, and we expect a similar amount in Q4. However, given the specific interest in AI demand, we will begin to report a new metric. Total HPE orders booked that include accelerated processing units, or APUs. APU includes GPU based orders, within both the HPC and AI and Compute segments. However HPC and AI represents the large majority. Year to date, total APU, HPE APU orders are over \$3 billion. This metric captures all APU orders across a range of suppliers. We expect the strength of our order book and the pipeline opportunities we see in front of us to drive near and long-term growth. We forecast solid double digit revenue growth in the HPC and AI segment between FY22 and FY26. For modeling purposes, you can assume for now, this CAGR will be front end loaded. We said at last year's SAM that we expect this business to carry a double-digit operating margin over time. While we've been investing and will continue to invest in AI, as we can see in our - - demand growth, our long-term expectation for this margin structure has not changed. We continue to expect approximately 10 percent operating margin by FY26. Operating margin improvement is a focus for this business, as Justin noted. When we expect to benefit from an increasing investment in software and continued learnings for the supply, from the supply chain challenges to improve supercomputing execution, gross margins, and capital intensity.

While its early days, our commercial AI wins, on average, have come at margins above typical, above levels typically commanded for relevant product lines. And you know this business benefits from scale, which AI demand is absolutely supercharging. Our Compute business is a steady value contributor at HPE, and it continues to be a critical component of our overall company scale. We estimate the Compute TAM at \$63 billion in 2026, excluding the Tier 1 market where we have chosen to limit our presence. We expect the market to grow at four percent through 2026, but to be mostly flat when excluding China.

We're preparing for growth in the inference market, and are expecting to capture that upside through this segment. Earlier this year, we induced, introduced several AI optimized and GPU intensive HPE ProLiant Gen 11 servers. However, we're not embedding significantly above normal GPU growth for inferencing in our revenue outlook. If or more likely when such growth does materialize, our product portfolio between HPC and AI and Compute is well positioned to capture the entire life cycle of opportunities from tuning or training to tuning to inferencing. We remain focused on capturing profitable unit share with stable revenues while managing our long term and market leading operating margin target of 11 to 13 percent. This includes an assumption that component costs will rise in FY24. Let me also reiterate that our Compute segment comprises a portion of our total, or of HPE's total server revenue. We also recognize server revenue in our HPC and AI segment. We expect HPE's combined server revenue to deliver mid-single digit growth in FY24 to approximately \$16 billion.

HPEFS remains a steady and strategic driver of healthy returns in equity across economic cycles, and it facilitates our as-a-service pivot in multiple ways. This segment is particularly critical as a support for our HPE Green Lake business because it creates investment capacity for our customers. It supports our as-a-service pivot, with its best in class lifecycle asset management business. The fleet management concept is becoming way more powerful, as customers' mindsets are shifting from paying for specific equipment to paying for capacity. We expect the business to sustain a mid-teens return on equity with a mid-single digit revenue growth, CAGR and an upper single digit operating margin through FY26. This

revenue growth is slightly above the long-term trend, given the rising interest rate environment.

Okay, let's translate that into a multi-year financial outlook. Our three-year revenue CAGR is expected to be two to four percent in constant currency. This outlook incorporates the current macroeconomic picture, and it affects risk and considers a revenue base, revenue based year that benefitted from an order book consumption, which means we expect demand growth to be above revenue growth in FY24. We aren't incorporating significant further growth in AI demand beyond what we have current line of sight to. We expect non-GAAP diluted net EPS CAGR of three to five percent. Our structural non-GAAP effective tax rate is 15 percent, from 14 percent in FY23, largely due to, largely due to removing H3C earnings from our pre-tax non-GAAP earnings. Excluding H3C and the tax rate change, we forecast our non-GAAP diluted net EPS CAGR to be seven to nine percent, and our philosophy remains that we will grow our operating profits faster than revenue. We expect our mix-shift, including as-a-service revenue and the rising profitability on that revenue, to be strong contributors to our operating profit growth over this time. Our mix-shift to higher growth revenue has lifted our visibility into free cash flows. As a result, we're increasing the amount of free cash flow we intend to return to shareholders to approximately 65 to 75 percent between FY24 and FY26. Above our recent historical target return of between 50 and 60 percent. This includes a combination of targeted dividend increases and larger share repurchases. But we'll be thoughtful around the target on both directions, given we'll continue to follow our returns-based framework.

And for cash flow, we expect our FY 24, free cash flow to be between \$1.9 and \$2.1 billion. That guidance is flat year over year with our current FY23 guidance. However, it's important to note, FY24 free cash flow will rise 10 percent year over year when adjusting for approximately \$200 million in cash flow we receive from H3C in FY 23. This year-over-year improvement on an adjusted basis will be driven by lower cost, lower cash transformation costs. We also expect improving working capital, although the surge in capital intensive AI slightly constrains us for now. Importantly, we expect free cash flow growth to significantly exceed net income growth for the next three years. Our conversation of non-GAAP net earnings to free cash flow is steadily

improving, as you can see on our slide. We are on track for approximately 70 percent in FY23, a sustained improvement from FY21, or FY22, FY21, and FY20. And we expect to reach approximately 90 percent by FY26.

Let's now turn to our FY24 outlook. We expect revenue growth of two to four percent in constant currently, and FX to be a modest 50 to 100 basis point headwind. We're expecting AI demand to be a key driver to above trend HPC and AI segment growth, and Intelligent Edge revenue to be slightly up year over year. Non-GAAP gross margins is forecasted to remain in the 35 percent range. Continued contributions from the Intelligent Edge, Storage and HPE Green Lake are to balance the pressure on Compute, gross margins after and above trend FY 23. We are prioritizing investments in higher growth opportunities we've talked about today, balanced with cost discipline in lower growth businesses. For modeling purposes, you can expect OI&E to be approximately negative 300 million. We expect our structural non-GAAP tax rate to be approximately, to be 15 percent, so GAAP diluted net EPS is expected to be between \$1.83 and \$2.03, and non-GAAP diluted net EPS is to be between \$1.82 and \$2.02. Finally, we plan to generate FY24 free cash flow of between \$1.9 and \$2.1 billion. This does place us below our initial plan of more than \$6.5 billion between FY22 and FY24. The biggest changes to our assumptions are, one, we no longer are assuming cash flow from H3C, and two, robust AI growth suggests supply chain and working capital won't normalize as planned.

We expect free cash flow to follow our typical seasonal pattern of negative in the first half of the year, and significantly positive in the second half of the year. And, we intend to increase our dividend by eight percent in FY24, given our rising confidence in free cash flow. We expect our FY24 corporate performance to be weighed to the back half of the year, while we are expecting the first half to benefit from normalizing order book in the Intelligent Edge business, our HPC and AI revenue growth will likely be governed by GPU availability well into FY24.

Our non-GAAP diluted net EPS is expected to grow by four percent from a normalized level. The bridge between the mid-points of our FY23 and FY24 non-GAAP diluted net EPS includes a few non-operational headwinds I'll walk you through.

The H3C put exercise creates a 17 cent non-GAAP EPS headwind. We are not planning to accrue dividends from H3C in FY24, nor will we report H3C earnings in our non-GAAP income in FY24. However, we're required to continue to include GAAP income in GAAP income, our proportionate amount of H3C earnings generated in FY24, while we still own the interest. We continue to expect to receive H3C cash proceeds in the first half of calendar year 24. Other OI&E is likely to be a ten cent headwind. We benefitted in FY23 from FX hedging, and other gains, and interest rates are higher. As I mentioned, our structural non-GAAP tax rate is forecasted to be 15 percent in FY24, and this equates to a two cent headwind in EPS, and again, is largely attributable to removing H3C earnings from our pre-tax, non-GAAP income.

We also have some operational headwinds in the coming year, including normalizing margin structure in Compute segment, following order book consumption during FY23 and planned investments in our growth businesses. The combination should reach approximately 34 cents at the mid-point of our guidance range. However, we intend to more than offset these headwinds with operational improvements in revenue growth, specifically in our higher growth, higher margin businesses, and ongoing efficiency gains. Operating margin target gains should be particularly evident in our HPC and AI and Storage segments. And as you know, we follow a discipline returns based framework to ensure we maximize shareholder value. Our rigorous investment evaluation process balances investments for growth with capital returns to shareholders. Our top priorities are investing to capture high return opportunities while remaining committed to dividends, opportunistic and material share repurchases, and retaining our investment grade credit rating. Our acquisitions, going back to Aruba Networking, and more recently to OpsRamp, Athonet and Axis Security have followed our disciplined ROI based framework.

We are confident in our ability to grow organic revenue, EPS and free cash flow. With the premise of our M&A to be, to accelerate that growth with acquisitions that are value accretive. We are expecting to receive a significant amount of cash when our H3C deal closes, and we'll update you with our plans for that cash at that time. However, we do not intend to hold excess cash over the long time, though some reserve is prudent during uncertain economic times, or during a technology inflection, such as the AI.

So let me close with a recap of our financial messages. We see health growth in the IT industry, and we have products across our portfolio that leave us very well positioned to capture that growth. As a result, our expectation is for sustainable two to four percent revenue growth and constant currency, and we're pivoting our portfolio to higher growth, higher margin revenue, led by our efforts in the edge, HPE Green Lake, HPE Alletra Storage, and AI. We intend to expand operating profit faster than revenue over the outlook period, which translates to a three to five percent non-GAAP net diluted EPS growth, or seven to nine percent normalized for our H3C divestiture and tax rate change. Finally, we expect free cash flow growth in excess of net income growth. This gives us line of sight to increase our target percentage of free cash flow returned to our shareholders to 65 to 75 percent. It also gives us a target, it also allows us to target an eight percent increase in our dividend in FY24.

I've thoroughly enjoyed working more closely with all of our business unit leaders in the last couple of months as we finalize our FY24 plan, and I have great confidence that our thoughtful strategy will provide very attractive investment returns. We look forward to your questions and feedback in a bit, but first, let me turn it back over to Antonio to close us out.

[applause]

MR. ANTONIO NERI: All right, thank you, Jeremy. Well, thank you, Jeremy. We believe we have presented a clear case for the compelling value in HPE will deliver to our shareholders. With a total addressable market opportunity growing at about 1.4 times, I am confident in our ability to capture that opportunity in a way that delivers full investors for several key reasons. First, while some microeconomic challenges remain, I'm optimistic about the math. HPE is more relevant than ever because we have innovated the categories that are most important to our customers now and in the future. Second, we have made very sound investment choices over the last five years to pivot our portfolio to a diverse set of businesses with greater profitability potential and they are paying off. As we unlock greater growth from these markets, our investors are poised to share in higher returns. Finally, we have an experienced and personal leadership team who execute with vision, a sense of urgency, and a commitment to culture. Our company recognizes the need now more than ever

to continue investing our workforce to attract and retain the necessary talent to execute on the growth strategy we communicated today.

As we left SAM, the SAM stage last year, I shared with you that HPE does not wait for the next big thing to happen. We accelerated what comes next for our customers, our company, and our shareholders. Our foresight and spirit of innovation has always served us well and positioned us to realize the full benefit of what is on the horizon for us and for our industry. We expect HPE to capitalize on the opportunity where this goes here today, at the Edge, Hybrid Cloud and the exciting opportunity that AI presents. We believe strongly in the long-term profitability potential of HPE and what we offer as an investment opportunity. I would like now to take your questions. We will give the crew a little bit of time here to reset the stage, and then we will invite the presenters back on stage so we can invite your questions, and then also, the audience on the webcast can submit question on the bottom of the screen. So...

[background noise]

Q&A

MR. NERI: I'm going to sit here.

JEFF KVAAL: In the middle.

MR. NERI: I'll sit here. And then Jeremy. Justin there. Great.

MR. JEFF KVAAL: All right, thank you all for your presentations. I wanted to let you all know that in addition to the executives that we have on the stage today, we have a few other members of the executive committee that may help us out. Gerri Gold is over here. She runs our HPE FS business. Neil McDonald in Compute has an important wedding anniversary today, so representing Compute is Krista Satterthwaite who is our Compute SVP and General Manager. We have some mics floating around, so if you would like to, yep, okay, great, you guys know the drill.

MR. NERI: As expected.

JEFF KVAAL: Super. Let's go with Wamsi first right here in the middle.

MR. WAMSI MOHAN: Thank you for the presentation, Wamsi Mohan, Bank of America. I guess I want to kick off around the assumption of Compute because you kind of broke the Compute TAM into non-AI and, you know, the AI opportunity for inferencing. You're seeing a lot of your competitors talk about AI servers for, within sort of the standard category for training, as well, so kind of wondering why you're not categorizing any of those kind of servers within, for training specifically, within Compute, and do you see the entire AI opportunity necessarily in the HPC AI TAM only?

MR. NERI: So, first of all, we don't see the AI opportunity just on HPC and AI. That's why Jeremy made that comment in his remarks. You have to think about the server category, okay? The server category in our company has two distinct segments, the Compute segment, which is the traditional enterprise and what I call the cloud centric kind of infrastructure, whether it's tier one, two, or three. We see that also in other verticals like telecommunications, and obviously more at the edge, as well. And then you have what I call the HPC and supercomputing side, which includes the traditional HPC business, the supercomputer, and now what Justin talked about AI, infrastructure. Much of the AI infrastructure for training opportunity is captured in his side, but there are aspects of training which may be captured in computing, as well. But most of the inference in the end is in the traditional Compute because you may deploy a small class or you may deploy one server with eight GPUs, and that's why Jeremy was very good in saying when you think about the entire server category as a, as a whole. We expect that business to grow mid-single digits, approximately \$60 billion, I think you quoted, and, but the important message you need to take away is that for AI, we cover entire lifecycle, and particularly on the training tuning, because we have large set of capabilities across supercomputing as one aspect, AI infrastructure, which is gonna grow seven times, you know, supercomputing, and then the AI software platform, which is a huge differentiation to be able to deploy this system at scale. Fact of the matter is that we're gonna give you an order for what we call APUs, which includes GPUs and other accelerators, which includes both Compute and HPC and AI.

MR. JEFF KVAAL: Okay, I'm gonna go with Aaron in the front here. Thank you.

AARON: Yeah, thanks for taking a question. So, first of all, I just wanted clarification. The \$3 billion number you talked about for AI, I think it was new orders. There was numbers in the past you've thrown around, pipeline at 3, orders at \$1.6. I think it started at \$800 million. I just want to clarify, what's the comp of that number previously? And then on the AI narrative, I'm going to get a little bit technical here, but the, the slingshot is a core competency of the company, so I'm curious of where we stand, as we think about these large AI models, you know, infrastructure deployments, there's a lot of inertia around InfiniBand versus ethernet, and there's this migration from 400 to 800 gig. I'm curious to where you stand on slingshot from a technology perspective to keep up with that, that architecture.

MR. NERI: So maybe Jeremy, if you can talk about the orders and then you can talk about the slingshot.

MR. COX: Sure, so we have talked historically about wins, AI wins at \$1.6 billion. We've also talked about how that contributed to our order book in HPC and AI that was around 3 billion at the end of Q3. I think the better compare of what we're saying now is these are, this is representative for what we have orders booked in the year. So, that obviously will feed into the order book, but that is, there may be other factors that feed into the order book. So these are GPU based orders that are booked within the year, and that's over 3 billion and we see a path for that to continue to grow. I would say also, on top of that, that the 1.6 billion was more capturing what we were seeing in the near-term demand around the H100 Nvidia piece. I would say that a meaningful majority of the over \$3 billion order number is made up of that category.

MR. NERI: So think about it this way, it is 3 billion orders which are in our books, and we will take time to bill, ship, all that time, and that's what we, we were very clear in our calendar that we want to give you the order, right? Because sometimes between the win and the order, there is a time delay. You win it, but then you have to win it, it takes time and all that. So \$3 billion is the order for the balance of the year, so far.

AARON: Year to date?

MR. NERI: Year to date. Okay, slingshot.

MR. COX: Slingshot, so first of all, we get caught up a little bit in mixing interconnects, network interconnects that are used for east, west traffic and in a cluster versus 400 and 800 gig at a data center switching side. If you think about interconnects in a cluster, there's two, two players in the market today, InfiniBand, from Nvidia and Slingshot. Those are the two things providing high performance interconnect for east west traffic in clusters. If you think a little bit about what WAMSI was asking earlier on box, there's not a lot of people buying one box of GPUs today, okay? They're buying them in 512 GPU building blocks, and so that whole thing has to be connected by a network, right? They're, and that, and that network is InfiniBand as the default for the reference architecture, the H100 reference architecture. However, we have demand outside of that in our supercomputing business for Slingshot, for interconnecting both H100s, but also grace hopper, as well as MI300, and so we're, and we obviously are running that on the Frontier Supercomputer and Argon supercomputer, which is powered by Intel GPUs, so we've got broad demand for that interconnect. That's different than the data center network that sits on top, and if you think about why that's different, it's that at the data center level, I've got to handle a whole bunch of different kinds of traffic, and so the speed matters, cause I've got to manage the inefficiency. The interconnect is all about tuning and optimization for latency and loss because I've got to run one parallel workload. I've got to run one large training run, and the supercomputer world, I'm running one large simulation for weather prediction or something like that, and that's why they're very different. To give you a technical metric, we did a little bit of analysis on Slingshot a couple, maybe a year or two ago, and we found that Slingshot at 200 gig was 92 percent effective when running Rocky, RDMA over converged ethernet, for those of you that enjoy the technical side of this, versus 50 percent for ethernet. So, that means that, you know performance, you know, packets, data, efficiency was far higher because we're running one very specific workload. We're trying to address GPUs or CPUs across distributed memory. That's why east, west traffic and the interconnect for the clusters is different than the, the top network. Does that, I know I went deep on that. I'm in tech, so I want to make sure I--

MR. NERI: And the other thing on that, part of the investments we are making, right, so we talk about investing in a

supercomputing IP, obviously that supercomputing IP a lot is what he talked about, which is the networking piece, which is slingshot, and the software to manage that, the, the latency. That will follow, you know, 400 and 800, right? But that's because of the ethernet, but the architecture is so different, right? Cause you're connecting every node to every node, and if you're, Frontier is 9,408 nodes, times 9,408 nodes, cause you have to look at it as one entity, and that's where our advantage is, and that's why we can use it elsewhere.

MR. JEFF KVAAL: Okay, why don't we go to Toni right next to you, and we'll come right down the line. Can I just remind you to introduce yourself on the mic, please.

MR. TONI SACCONAGHI: Toni Sacconaghi from Bernstein. Sorry, I just want to follow up on that last question. You're saying you have 3 billion in orders here today for accelerated Compute. It sounds like the vast majority of that is in HPC and AI.

MR. COX: Correct.

MR. SACCONAGHI: You said your HPC and AI backlog is 3 billion. It's been 2 billion or more every quarter for the last 12 quarters, so is the 3 billion in orders like truly incremental? Like how, how do we think about that? I don't think you addressed the specific question, what was that exact comparable metric last year through this date? Because I'm still struggling to jive the huge order number with the backlog that you're attributing to AI.

MR. COX: Yeah.

MR. SACCONAGHI: And then second, related to that, can you comment on the wait times for delivery on your AI servers, how concentrated your customer base, and what kinds of customers have generated the orders.

MR. NERI: Yeah, why don't we start with the first one and then back to you because the second is interesting and it ties to the customer stories we shared.

MR. COX: Right, when we thought about this metric, Toni, we wanted to make sure we were being thoughtful, not about just the near term demand like in H100, which again was more the focus of our prior discussion of the \$1.6 billion wins. We wanted to capture the full life cycle and, and you can see

that again within the lower inferencing area, within the Compute area, all the way up through the supercomputing space, and then obviously in the high performance space, where you're seeing the H100 activity happen more, and so you know, the prior period compared to that, I would say you can think about the Computer portion of that has not significantly increased on a year over year basis, less than 10 percent of that total, of that total order base. Supercomputing at this point a little bit more than that, but not, again, not a significantly high percentage. So the vast majority right now is falling in that HPC, that HPC area that's attaching with the H100s, and this is the metric that we're gonna continue to provide to you guys. I appreciate that some of the order book dynamics, that had a multi-year component to it. As you know, some of these supercomputing deals take multi-years to resolve, so we thought this order metric would be a better way for you guys to digest the activity that's currently happening.

MR. NERI: And in the prior periods, Toni, there's also the fact that you may have a frontier or Aurora in some of those that just wipes out significant backlog in one go and the you build against, but this is what he's referring, that 2023 the mix of the, the bookings is, has changed quite dramatically.

Okay, so want to talk about the customer? What type of customers' demands we're seeing?

MR. JUSTIN HOTARD: Yeah, so I touched on some of the segments, you know, and I would say the way to think about it is, if you think about some of those cloud service providers, they tend to be buying in these 512 GPU chunks that I touched on earlier, so I, in terms of concentration, I would think of them kind of buying large clusters in that stage for, for training. Now, they may buy multiples, but that's a good building block to think about. Then, the A, the foundation model builders that are, that are buying directly, I would say, again, that's probably the largest size they're building, so I, they may be buying something that's half or a full, a full pod of GPUs, I think 256 or 512, and then we have customers that are buying, you know, subsets of that, and I think about, I think about that as like traditional enterprise customers, existing HPC and AI customers that are adding capabilities. What we haven't seen yet is material demand from national AI initiatives and national AI centers, and I think we'll start to see some of that come in. That may

be more of a mix of some of our traditional super computing systems, so they may buy Cray X systems. They may buy Cray XDs or a combination, and a combination of that would be, if you look at the win we announced with Kowst [phonetic], they bought a small grace hopper cluster, which is basically for AI training. The reason that's so important and so differentiated is that there's a lot of customers looking at Gen AI, not just for, you know, broad LLM applications, and I touched on this with Argon. There are customers looking at training scientific models, doing expert models with deep technical knowledge. Those customers realize that they may need something that's a little different than just the standard building block today, and that's why they're coming to us, and I think you'll, you know, hopefully you'll see us announce some more, some more news on that shortly, but there's, you know, we see these, each of these markets as very different today.

MR. NERI: And you have to think about the - - right? A lot of the momentum that we have heard about in the market is about this large language model. That's one way to think about it, but the other ones are traditional AI models like Computer vision, mot simulation, molecular docking, climate. Climate is pretty significant, as well, and diffusion, text, picture to text. We have all this, and one of the things we have, which is a key initiative, we have been supporting those models for a long time. Now they become more pervasive, right.

MR. SIMON LEOPOLD: Thank you. Simon Leopold with Raymond James. Two, if I might. First one, I want to see if we can get a better understanding on the impact on the financial statements of rising and high interest rates, how that factors in to the business, and the second one is probably for Phil is, what are your market assumptions when you talk about slight growth in the, in the edge business, what are you assuming is going on in the world around HP and what's your, basically, assumption on your share position, basically.

MR. NERI: So you want to talk about the interest in general?

MR. COX: Yeah, I would think about that, Simon, in two parts. As you likely know, the large majority of our company debt is attributable to our FS business, and there, that business is effectively able to capture the high interest rates and pass

that on to customers as its attracting on the margin there, and so that business is actually driving higher revenue performance in the HPE EFS business that I mentioned earlier, as a result of that, that direction.

On the core business, we mentioned a 10 cent headwind to EPS. That is definitely a component of that headwind are the higher interest rates as we're rolling over bonds and we're having to incur higher interest rates on those bonds as we rolled.

MR. PHIL MOTTRAM: And then, on the market share assumptions that we are making, so we'll assume that market is negative next year, so the market drops next year, and we offset that through continued market share gains, which we have been gaining market share the last couple years, and then also the expansion of the TAM into the new areas that I talked about in my presentation, namely security, data center networks, private 5G and NAS, so that's the assumption that we are making.

MR. JEFF KVAAL: Asiya [phonetic], right, right up front. Thanks, Kate.

MS. Asiya: Thanks, Asiya from Citi Research. Just if I can on the Storage, which I understand now is in the hybrid cloud segment, I think, the underlying assumptions would be as you shift more to Alletra and to your own software IP, the margins on that would be higher than the mid-teens that you've been reporting. I think in the revised model, that, that's kind of what you're reiterating, that the hybrid segment would still be in the mid-teens, so maybe you can walk us through, are there some underlying investments still that are going on in Storage? And then you talked about market share gains, as well, specifically as it relates to Storage. Maybe you can kind of lay out where you see strong opportunity to gain share there, and you know, who would be the traditional players? Is it the legacy players that you're going after? Is it some of the more newer software offerings there?

MS. FIDELMA RUSSO: You want me to start?

MR. NERI: Go ahead, yep.

MS. RUSSO: Okay, so in terms of Storage investment, we've invested a lot over the last number of years in the new

architecture. We continue to invest on the, on the protocol build-out on the architecture, so we're really poised in the block space to take share, and our investment in the go to market side of it is now complimenting the engineering investment we've put in, so there's continued investment in new protocols, and also the investment in the go to market side. In terms of taking share, you know, we've got about a 10 percent market share within Storage. We've got a lot of runway to go, and we're very confident about the cloud enabled and the platform based attractiveness of the architecture with our customers, and we see us not against traditional players, but also against, when customers are thinking about should I put my data in the Public Cloud or should I put my data on prem, and so especially as AI is, you know, more top of mind on storage buyers, as well, we are seeing more and more conversations about how do I expand my storage footprint on prem.

MR. NERI: And part of that investment that Fidelma talked about, protocol, is the file part of this because that ties directly to the IP's.

MR. JEFF KVAAL: Can we go over to Meta over here on the side in the front?

MS. META MARSHALL: Thanks, Meta Marshall, Morgan Stanley. In a, numerous of you guys talked about kind of customers exploring different paths on AI. I guess I just wanted to get a sense of where do you think customers are on the journey of even kind of trying to discover what their approach is going to be, and when do you see as the timeline of some of those decisions. And then maybe second question, you guys talked about a double digit operating margin on HPC, just what are kind of the critical components from getting where you are today to that double digit. Thanks.

MR. JUSTIN HOTARD: Want me to take those two?

MR. NERI: Yeah, go ahead.

MR. JUSTIN HOTARD: Meta, I'll answer the double digit one. It's pretty first. Operating margin is pretty simple. It's scale in the business, which gives us leverage, and then it's just the, it's the mix shift to more of our IP rich offerings, and that includes the penetration of our software stack, obviously continued penetration of slingshot as a part of that, the growth and demand we see, and then the complete

solution we talked about where we're integrating, you know, our services and our supercomputing IP and that software stack in HPE Green Lake for LLMs, so that's the driver of margins.

On your first question, the journey, look, I think where we are right now is, here are a number of customers spending, you know, making massive investments around foundation models. A lot of it is around LLMs for, you know, for broad commercial applications, and that has not, that has not broadly deployed into massive investments in the enterprise yet. It, it's largely, I think it's largely the early days. We think the enterprise build out is going to be massive. We also think that the initial models, which are largely general purpose language models, across many languages around the world are going to be, are going to be the beginning of a much broader and more robust build out of models and training, and that will be things like technical use cases, you think about scientific models, you think about, obviously financial and trading is one we'll probably continue to build, just given the dynamics and the economics, but that will go on and on and on, and that's what, that's what we think we're at the beginning of something very big. It's also why we're being very conservative in terms of how we're thinking about the business in our forecasting, because we're really looking at the demand that we see today. But we, we think this is going to be the beginning of a long investment, and you know, enterprises, most enterprises I talk to and Krista can comment as well, are very productivity focused today. As I touched on, I think productivity is the start. Over time, I think we're going to see people look at new revenue models, and we'll see that part of the business emerge, as well.

MR. NERI: Yep, that's why I want to make the comment that a lot of the new customers are about developing these foundation models. Most of the enterprises are going to leverage these foundation and tune these models with their own data, but they are right now more focused on productivity like we are in our own company, right? How we leverage that, whether that is in the coding base with our R&D team, whether it is in our operations, in our services team with bots and, and the like. But I think the biggest focus for them is how they maintain control of the intellectual property, the data, in a way that allows them to deploy these models with confidence and the

accuracy, because that's the other important piece he talked about. It has to be accurate, right? So.

MR. JEFF KVAAL: Okay, why don't we go with David over here, two more over here, livestream, and then we'll come back up to the front and then you folks.

MR. DAVID VOGT: Thank you, David Vogt from UBS. Two, if I may, one for Justin, one for Phil. As inferencing becomes a bigger part of the AI story going forward, how is Slingshot sort of competitively positioned against the legacy ethernet vendors, as we move in that direction. And then for Phil, you talked about four different categories driving growth in edge and Aruba, you talked about networking convergence with security, also data center switching. Those are two pretty crowded markets, so can you maybe touch on what you bring to the table, besides sort of the operating system, which you touched on briefly in terms of how you're positioning yourself vis a vis the competitors in the marketplace already pursuing a similar strategy. Thanks.

MR. HOTARD: Yeah, so I think if you look at it, again, as you think about Slingshot, think of it as east west traffic within a cluster, so as inferencing workloads scale and there's more demand, there probably are parts of inferencing that will, will have, you know, large capacity, right? Think about content distribution networks as a parallel in the, in the internet. If you think about, the other part of inferencing, which is why we're focused on the entire life cycle, getting down to the edge, these are, this could be down into your mobile device, or into a, obviously a car as we see with autonomous driving, so there's going to be a lot of places that inferencing gets used. We do think there's going to be a demand for, you know a cloud, a cloud service or a large cluster of inferencing solutions as part of the, the broader market build out. Phil, I'll let you answer the other one.

MR. MOTTRAM: Yeah, so I think your question was crowded markets in both - - switching and security, and then how do we believe that we'll be successful in those parts of the market. So, on the compass switching element, I mean, our strategy there is to follow the networks, so obviously the access point connects into the compass network and we find customers that use the platform, Aruba Central, you know, once they're comfortable with that, it's a logical extension

to go into compass switching, and we find it, you know, there's a big player, obviously, in the compass switching space that we come up against quite regularly and customers like our, customers see and the fact that we have one platform linking all of the products together, whereas some of our competitors have different platforms for different products, so that's on the compass switching side of it. On the security side of it, it is, I think that's a market that's going to go through a number of changes, isn't it? Cause on the security side of it, you've got the three elements being ZTNA, SWG, and CASB [phonetic], and most enterprise customers right now don't buy those three from one player. They're often buying them from three different players, and then on the network side, you know, you've got network firewall and the SD WAN side of it, so I agree that there's a lot of players in the market, but there aren't many players with the full range of five capabilities. So I actually think we've got a really, really good option there, albeit the market is going to have to go through a number of steps to consolidate.

MR. JEFF KVAAL: Sidney, please, yeah.

MR. SIDNEY HO: Great, thanks. Sidney Holt with Deutsch bank. Just want to follow up with that question. You talk about Intelligent Edge growing about low double digits over a three year period. Which of these opportunities that you highlighted are more near term, call it one to two years, versus longer term? Do you have all the products and services already in your portfolio, and in the near term, are you expecting the operating margin, that this is to go back to that mid-20 percent already in fiscal 24? Thanks.

MR. MOTTRAM: I mean, just on the, on the capability side of it, in the products, I mean, on the data center side of it, we have been investing in the products in that space in the last 12 months and have got a pretty well defined roadmap for the next couple of years, and we do low hundreds of millions already in the data center space. It's kind of difficult, actually, for us to, customers buy switches. We don't always know whether they're going into a data center or a compass location, but for the most part, we know that some of the products are already going into data centers, so we've got a good starting point there, and the sales teams are pretty accustomed to selling in data centers, so data center is a very real opportunity for us today. And we've got a good

growth there, and we've been adding more resources, so I feel pretty good about that.

Private 5G is an earlier space. I mean, the company that we acquired is pretty low in revenue terms, but there is a lot of interest, because what private 5G is allowing companies to do is, companies with very outdoor locations like, I don't know, oil refineries and those sorts of places, they haven't really been able to benefit from network technology before because they struggled with covering large outdoor spaces, and private 5G enables that, so we've seen, since we made the acquisition, a lot of interest from many, many different vertical sectors, albeit the starting point from a revenue perspective is relatively low, and the same is true on security. The company that we bought, we're very, very pleased with the technology and the people, and when customers go through proof of concepts of the technology, every proof of concept that we have had with the customers since we've bought the company, they've always proceeded to contract, so we're very happy with the technology, but again, it's relatively low revenues at this point, so we've got a way to scale.

And then with regards to the margins and OP, I mean, on the security side, they would be better than the margins that we currently see on the private 5G would be there or thereabouts maybe slightly below. Data center would be typical to the margins that we, again, would currently see, and then on the opex side of it, you know, this, this market, we should benefit, on the top line potentially, things get more aggressive next year, cause the market contracts and therefore, potentially you have to do a bit more discounting, discuss, when we feel like we've got a pretty strong proposition, but then when you go to the opex line, during the supply chain crisis, we had to fly everything around the world by air, which is more expensive. We used to have to pay expedite fees to supply us to get equipment moving, and as the supply chain is now opened, we're now, the opportunity to move equipment around by ship and not pay expedite fees, so this put some takes in the overall P&L, but I think the statements that Jeremy made earlier are very doable.

MR. COX: I would just add to that, I would, I would say, yeah, it's the dynamic of similar to Compute, we had high-priced backlog that was, as it was being worked down with lower logistic costs, we saw above-trend results. And so that's

going to start to normalize a bit, and we're continuing to make investment in this business. And so while I expect us to be in FY24, maybe on the higher end of that mid-20s range, over the three-year outlook, we would expect to be in that range.

MR. JEFF KVAAL: Why don't we take one from the webcast? It's been staring at us here—

[Crosstalk]

MR. ANTONIO NERI: [Interposing] Yeah, there's one here actually.

MR. KVAAL: Yeah.

MR. NERI: You want to read it?

MR. KVAAL: I'll paraphrase it, really. Actually I can read it. It's well, so it's \$200 million in H3C cash flow you mentioned for fiscal '23. How does that compare to prior years? And is that a full-year number?

MR. COX: Sure. So maybe a bit of the confusion here is, is that we have to differentiate between the earnings that are incorporate for EPS purposes and then the cash that we ultimately include in our free cash flow, which are really from dividends, which are paid in arrears. And so that may be the disconnect for the person that asked this. And so the adjustment we've made to the \$200 million on a year-over-year basis is reflecting the cash received in the dividends. The adjustment we made to EPS was reflective to our amount of earnings that benefitted from the addition of those H3C earnings into our non-GAAP EPS.

MR. NERI: But it has been fairly...

MR. COX: Yeah, over time it would match. We do have some dynamics with the disposition where under rules in China, certain dividends aren't allowed to be paid after a certain period of time, and so that will be an impact in '24 as well.

MR. KVAAL: And is that number essentially flat-ish; has it been growing? How can they look back a year or two?

MR. COX: I think that's the relative adjustment for '23.

MR. KVAAL: Why don't we go with Samik [phonetic] up here in the front? Thanks, Kate [phonetic].

SAMIK: Hi, Samik from J.P. Morgan. Had one for Justin and one for Jeremy. Justin, when you look at the different parts of the AI stack that you're offering, software, networking, and then the Compute solutions, like how you are thinking about the go-to-market, in terms of do you want to sort of remain the full-solution provider? Or is it better sort of adoption than quicker adoption to be chased if you segregate that thesis and sort of use your differentiation in each? And for Jeremy, you talked about H3C proceeds - -. You updated the investors on that. When you think about excess cash, can you just define that for us, how you think about what sort of cash on the balance sheet that you need and what are your priorities once you get that cash in excess of what you want to, what are the priorities?

MR. JUSTIN HOTARD: I think the way to think about this simply is in supercomputing, we're a full provider. And that's necessary because it's almost impossible to deliver a supercomputer without delivering the whole system. It's really, really challenging. There are some people that have tried it, but it's hard. In the AI infrastructure space, we're selling, we obviously aspire to be a full provider in that space.

But we're also selling parts of our stack where it makes sense. So some cases we're obviously selling H100 clusters. I mentioned Recursion Pharmaceuticals. They're a software customer only today. They bought our machine-learning development environment. And then you've got a customer like Aleph Alpha that initially bought infrastructure from us, then bought software, and is now our launch customer for HPE GreenLake for large-language models.

MR. COX: So on the cash-flow side from H3C, look, I think we're constantly balancing between our desire to return more capital to our shareholders, but also making sure we have the right amount of capital to make the strategic investments to grow our business. And so we assess that balance through our ROI-based framework that we talked about. I think as the H3C proceeds come in, we have now directionally suggested to you guys hey, we have a better confidence in our long-term outlook on free cash. That's driving our expectation of increasing the amount of share we purchase in dividends as a total percentage of cash flow back to our shareholders to the 65 to 75% range. And so I think those things will be

factored in when we get to the cash position when H3C is delivered in kind of the mid-calendar '24 period.

MR. NERI: So again, I think it's important when you step back and look at what we guided for capital return in '24 and the direction for the next three years, is all the operational side, right. So a better line of sight, the pivot is happening. Obviously we have seen the gross margin increase in our portfolio. And that gave us the confidence to increase from the 50 to 60% to now 65 to 75%. And in that, increase the dividend by 8% at the time when the dividend comes to be paid.

And ultimately, when we receive the proceeds sometime in '24, we're going use the same capital framework. But the confidence is not the \$3 billion, \$3.5 billion after tax; it will be less obviously. But it's about our ability to execute the plan that we just gave. And then when the cash comes in, we'll assess the best return. But we don't expect to hold excess cash in our balance sheet.

MR. KVAAL: Want to go with Mike and then we'll go with Lou.

MR. NERI: Yeah.

MR. MICHAEL NG: Thank you, Mike Ng from Goldman Sachs. I just have two. The first one is just on Intelligent Edge. I was just wondering if you'd talk a little bit about the curve over this mid-term outlook, obviously very front-end-weighted. Is that just a function of some of the backlog production that we'll see over the next two years? I'm just trying to reconcile that with some of the growth initiatives that you have that you mentioned before and why doesn't that help to accelerate that growth in fiscal '25 and fiscal '26? And then just a quick follow-up on H3C.

MR. NERI: Yeah, go ahead.

MR. MOTTRAM: Yeah, you're exactly right. I mean obviously this year, we have benefitted by a significant backlog retirement, and we have some of that next year, albeit at a lower rate. But we do benefit from this portfolio expansion. And - - next year, we're lapping very significant growth obviously for this year, so probably next year is going to be more subdued, as Jeremy highlighted. But then as you go out year two, year three after that, we would expect to build way more momentum in the new areas that I highlighted today, so data

centers, security, and private 5G. I think some of those could be quite meaningful in the sort of medium term.

MR. NERI: Yeah, but I think it's important on this point. In the last two years, we added \$2 billion to this business. That doesn't go away. Let's be clear. We added \$2 billion. We are building a new platform, if you will. And we're going to grow from that new baseline the numbers that Jeremy explained. In that, there is obviously continued market share in the campus branch, which includes wireless and switching and wide data network. Security will be a driver of that, obviously, with SD-WAN.

But then there is the new areas that takes longer to mature because you need to either fully integrate in the platform, customer-use cases will come to life, go-to-market is ramping. So that's why we gave you the guidance that we gave. But let's not walk back on the fact that we just added \$2 billion in that space as the new baseline.

MR. NG: Great, that's very clear. And then my next question is just around some of the comments around Compute ASPs benefitting from accelerated servers in fiscal '24.

MR. NERI: Yeah.

MR. NG: I guess how meaningful of a benefit is that? Are you seeing demand for accelerated servers right now? And then if you could provide some additional comments around the GPU availability point, that'd be very helpful, thank you.

MR. NERI: Yeah, I mean so the AUP, so average unit price, when you have some sort of accelerator - in particular GPU, it's significantly higher than the traditional CPU. The question is the margin side of the house, right. So that's why our constant approach with the operations team Mark [phonetic] runs and Krista Satterthwaite who is here representing Compute on pricing discipline, is very, very critical. So as the shift continue to happen, and I said we see evidence of that, but it's still very small relative to the size of the CPU base in traditional compute, call it single server. The structural side of that will grow.

And also remember with Gen 11, we also have another structural change, which obviously we can add more memory, more Storage, and more everything associated with that. And also there is a component of subscription as part of

GreenLake to lifecycle that server. That's why we said on an aggregate, the average unit price in '24 plus the rise of cost on the traditional components, which it's just a function of demand and supply, after the capacity kind of shortages, is going to happen. So the question is how the transition happen is what we built in our guidance is what we have line of sight, but in the context of the full lifecycle.

MR. COX: I would maybe just add one point to that. On a year-over-year basis though, we'll still see some pressure on AUPs. While we're seeing some of these tailwinds that Antonio described, we're coming off a significantly inflated AUP position in FY23. So part of the compute story will be some of that compression in AUP, even though we'll see tailwinds start carrying that out of-

MR. NERI: [Interposing] In the second half.

MR. HOTARD: Just on lead time, I think what I'd say is, I think Toni [phonetic] this, I may not have answered it earlier, is we continue to see demand well on excess of supply, and we don't see that changing right now. So that continues to be the view we have going into '24. I think that's why Jeremy touched on it being governed by availability.

MR. COX: Yeah.

MR. NERI: And our mix of demand is, can be quite significant, right, so if you have supercomputing with 40,000 GPUs versus a customer that has 512.

MR. HOTARD: Yeah.

MR. NERI: Right. Yes?

MR. KVAAL: Lou.

MR. LOUIS MISCIOSCIA: Thank you. Lou Miscioscia here, Daiwa Capital Markets. So asking a little bit of a different question, ARM recently went public, and you've got companies like Nvidia, Apple, AWS using your chips. Just wondering is there any opportunity for HPE to include any ARM products that could lower your costs, or alternatively, do you see as possibility, any risks in the area?

MR. NERI: To include ARM products?

MR. MISCIOSCIA: ARM products, yeah.

MR. NERI: We already do. Well we already have a great partnership with a company called Ampere. We already integrated that into Gen-11 platforms. We already have interest from customers that have unique workloads. They actually were not typically AI workloads. We had other type of workloads. Maybe the core count and the power and all that was perfectly suited for that workload. But it's still early. And remember that part of the challenge I will say, but I think is an opportunity, is the entire ecosystem of the application compilers and software have to come all together. I think we're at a different mature level than we were maybe three, two, four years ago.

MR. HOTARD: Yeah.

MR. NERI: But we already have in the supercomputer, I mean you can talk about that.

MR. HOTARD: Yeah, I would add two things. When Nvidia announced Grace Hopper, they announced the win that we had with them with CSCS in Switzerland, or Grace Hopper, so there's a solution with Cray X server [phonetic] on the supercomputing side that is ARM-based. The other thing is, in May we announced a system that the University of Bristol in the UK, that is a Grace-Grace superchip system, which is ARM-based. So as Antonio said, it's very much about the tool chain. But we're across, whether it's the Compute portfolio or in HPC in AI, we see demand, we're supporting the demand strategically, and working closely with partners like Ampere and Nvidia.

MR. NERI: We have that capability already in our products to support.

MR. KVAAL: I think we have one sort of back center. Yes, please.

MR. SID NAG: Hi, this is Sid Nag from Gartner Research. I had a quick question in terms of your longer-term strategy. It seems like you're at a crossroads as a corporation. So on one hand, you're competing with the incumbents, and you're competing with the hyper-scalers, right. So in other words, do you look at the hyper-scalers based on the coopt-and-compete strategy and continue to compete with the incumbents? And where do you see the growth coming?

MR. NERI: I mean we talk about the world being hybrid. And in that experience, you have to include the public aspect of that experience. And hence why GreenLake as a platform is

our cloud platform, which is truly designed with hybrid principle, we actually support the three largest cloud, 'cause customers have data workloads in those Public Cloud. And our goal is to give them the best cloud-native experience whether on the Public Cloud or on-prem. And at Discover, we made several announcements with Amazon Web Services. We have a partnership with Azure for different things. And so we have to continue to include that in our thinking. And in other aspects, we are going to compete, right, when it comes down to delivering specific outcomes by delivering an on-prem or a co-lo Edge infrastructure that have the same principles. But ultimately, give to them an Edge-to-cloud hybrid experience that matters.

I think the growth, I mean private cloud, as I said, right, is growing at 10%. That's a great opportunity for us, honestly. That's why we enter the market with a series of offerings that I think are better suited for what the customers are looking for, with choice, because we support all the run times available in the market. So it's just a balance, right. But ultimately it's what the customer is looking for and how we address those needs in the best possible way. If ultimately selling services through our hybrid solutions, through data protection services that they back up in a Public Cloud, that's fine. It can be either. But ultimately, we are driving to a cloud-native experience that delivers the best economics with the best experience that's truly hybrid.

MR. KVAAL: Did I see a couple over here? No, I guess not.

MR. NERI: I got one here.

MR. KVAAL: Yeah, why don't we go to the web first, and the question is you noted 18 cents in investments in 2024. Are these product or go-to-market investment, which businesses?

MR. NERI: Yeah, it's a combination of both. So obviously in Phil's business, right, we continue to make investments in the product side. Phil talked about the continued progression of our data center products, right, with the Aruba OS-CX, which gives the same experience if you're in the data center in the campus branch, but obviously have different features and functionality, but is all managed through the same platform. Also, it has integration that is done with 5G and - - and Axis security. But also, it's

invested as well in the go-to-market, continue to expand that coverage and specialization. All that investment in go-to-market are about specialization, much deeper specialization in Security, in Hybrid Cloud, obviously growing the coverage, Fidelma talk about the Storage.

In HPC and AI, solution architecting is important; pre-sales is important. And then on the R&D side, obviously there is a sizeable investment that we are driving in AI is silicon. When you are in the silicon, you asked the question, it takes multi-year investment in silicon. It is not like a quarter and you're done, right. So it takes a continuous investment and the software, because the software we are developing, the machine-learning development environment is what we call the pass layer to make sure the data automation and data pipeline for AI takes place, because one of the biggest pain points customer have is that how they bring together the data to be trained on that model, right. And so those are the things we are investing in as well. So there's a combination, but it's ultimately Edge, hybrid cloud, and AI, and obviously Fidelma has other investment in GreenLake as well. But also we've made some investment in Compute, silicon root of trust and other software elements that are very important, which we can monetize to our subscription-base models.

MR. KVAAL: All right, we've got a follow-up from Wamsi [phonetic]. We have time for one or two more; we'll get to you Matt for sure.

WAMSI: Thanks for taking the follow-up. I was wondering if you could share any color on the \$3 billion APU orders in terms of just average and median order size, so we can understand if that's, how much of that is largely HPC-driven? I mean are these kind of orders of \$100 million, or are we talking about \$250,000? What's, call it kind of the distribution look like?

MR. NERI: I mean it's obviously large-size deals.

MR. HOTARD: More sizeable than...

MR. NERI: More sizeable.

MR. HOTARD: Yeah, I think I covered the, I think Tony asked the question about the customer profile. I think you could sort of think about the mix in that construct, the types of customers. And obviously the large, the more super-pod size,

512-GPU-sized clusters people are buying, the more disproportionate they shape that mix.

MR. NERI: Yeah. But think about a higher- - number, yeah.

MR. KVAAL: We've got, yeah, we've got another...

MR. NERI: Oh, here, that side, we haven't been on that side.

MR. MATT Sheerin: Yeah, thank you, Matt Sheerin [phonetic] from Stifel. A question on your guidance for fiscal '24, you talked about a 42-cent tailwind in terms of operational improvements, and it's going to offset some of those headwinds-

MR. NERI: [Interposing] Yeah.

MR. SHEERIN: You talked about. Could you go into that? That's a pretty big number. Are we talking about restructuring, will there be charges, headcount reduction, things like that?

MR. NERI: Yeah, well it's almost like the number we had last year, if I remember correctly. It was at 45 cents I think last year or 40-plus cents in operational improvement. But it's actually the combination of the mix shift, right, which is driving higher gross margin dollars as the mix of the business continue to shift. Also, there is ongoing improvements in our operations effectiveness, which obviously we drive as a part of our ongoing improvement. We do that all now within our domain. We build those investments in our typical plans, not the typical way that was done many years ago. And so also we continue to be way more efficient from the back-end operations with our COO. But we have done a great job in driving efficiency in our real estate, in our supply chain with digitization, and all the things you expect.

But I think the vast majority of this is the continued mix-shift in our portfolio that drives higher gross profit. And also the go-to-market, the go-to-market's pretty interesting evolution, right, because since I became CEO, I have taken deliberate action with our chief sales officer to simplify the structure, eliminate roles, drive simplicity in our sales compensation, more focus, productivity at the sales-rep level. Those are all bundled together in that 42 cents.

MR. SHEERIN: Thank you.

MR. KVAAL: Sure, we've got Aaron upfront.

MR. AARON RAKERS: Yeah, Aaron Rakers again, Wells Fargo. Thanks for the follow-up. I just want to, I keep going back to this AI stuff, so I apologize. I think in the past in the commentary, you've talked about a lot of that being under the GreenLake umbrella. And so as we all try and think about this path and the pace of monetization of this large order book, how do we think about that in the context of GreenLake, right? Wouldn't it come out and actually rev/rec over a period of time under the contractual GreenLake model? And do you take the capital intensity related to that as well? You take the upfront capex?

MR. COX: So the GreenLake model, we have the benefit of our HPEFS business that ultimately incurs that leverage cost at the end of the day. And that provides the capacity for us to grow that portion of the business. From a revenue-recognition perspective, it depends at the end of the day on the term length in the contract with the customer. The shorter-term periods, you would expect to see while we recognize revenue upfront in the segments, we would be eliminating that and recognizing it over time for the company. Whereas you have longer-term arrangements with the customer, those would require upfront recognition, both for the segment and for the company. But overall, and to your broader question of how much of that GreenLake business was within that \$3-billion-plus base, there's a fair amount. At this point, I would say the majority is capex, transactional capex. But there's a good chunk of that GreenLake business embedded in that number.

MR. NERI: Yeah. We are extremely aware of the capital intensity, so we're driving different commercial terms in the way they pay us so that we can ensure that there is ongoing, sometimes prepayment of that.

MR. HOTARD: Yeah.

MR. NERI: But ultimately, as I think about the type of customers, right, so I think you're going to have a mix of different type of transactions in GreenLake. You're going to have those who say Antonio, Justin, I need a virtual private cloud with this size for me for the next three years because I'm going to continue to build a model, training, retraining, and is not a multi-tenant solution. It is in a public instance

that leverages our data center services and power efficiency and expertise that we have to run the system of scale. But they are dedicated to that, right. So that's I think is going to be...

MR. HOTARD: Yeah.

MR. NERI: Quite the vast majority of it. And then there will be those enterprises that may want to tune their data for a period of time, tune their model with a period of time - - for a month or two and then they said okay, I don't need that anymore. But still from the data/intellectual property perspective, they want to make sure that's contained, right. They don't want to put data in a multi-tenant-type of environment; that's very obvious.

MR. HOTARD: I mean maybe just two things. Given the demand that we've been talking about, most customers when they're coming to us even to buy a cluster on HPE GreenLake, it's more about getting the operating expertise and getting it up and running faster. They want control of the cluster. They're not interested in using it for a period of time. They're willing to commit because of the limited supply.

The second thing that Antonio touched on, this the multi-tenancy, and this is different than the Public Cloud. And I touched on this a little bit in my comments, but Public Clouds, your running multiple applications on one server, right, in most cases. Or you maybe scale it out to a few servers. It's kind of like what we do on our iPhones or laptops. In a multi-tenant environment in HPC and in AI, it's exactly the same, you're allocating a large portion, if not the entire cluster, to one job.

And the application is running in parallel over the whole thing. And that actually is where we have a lot of experience, because we built system-management software that does this for supercomputing customers. And it's exactly, it's very parallel to what we do in that space. It's almost exactly the same. And so as Antonio touched on, as the market matures and we see demand for that tuning workload, we think that's going to be another opportunity. But we haven't, we have the core technology for it, but we haven't seen that demand emerge yet.

MR. NERI: Yeah.

MR. KVAAL: We have time for one more? All right, we've got one right up front. Samik, please.

SAMIK: Hi, Samik, J.P. Morgan. Justin, going back to AI again, liquid cooling, a lot of discussion nowadays on that. How do you see that as a differentiator relative to where the comparers are on that front? But it also seems like that ecosystem to provide liquid cooling is not really completely developed yet. When you think about partnerships or even M&A, are you looking to enable, what are you doing to enable that ecosystem so that you can sort of build on that differentiation?

MR. HOTARD: Yeah, there's multiple elements to the ecosystem. This is a place, Cray has a history of working in liquid cooling for decades. So we've actually got lots of experience in implementing it. When you think about the liquid-cooling ecosystem, you need to think about it across a set of dimensions. The first is the, let's start with the data center. Data center has to be liquid-cooling ready, and so that's a place where there's a bunch of capital buildout that's happening, that it's an important part of the investment, 'cause you actually need to plumb the data center for liquid cooling.

Second for us in terms of IP is how you design the systems and architecture and how you manufacture them. We have a factory in Chippewa Falls. We've set up a factory in the Czech Republic to manufacture this. To my knowledge, they're the two largest-scale factories in the world for building liquid-cooling systems. And then you have to design the systems for liquid cooling, which is quite different than just simply air cooling, because you're not actually, you're dealing with heat dissipation across every component.

So think about, you think about fans and you're like well I can just replace the fans. You can't just replace the fans, 'cause you've now got to make sure you're liquid-cooling the CPU, the accelerator, the memory, the drives. If you're going to put a full liquid-cooling system in and have it be fully, run neutral, you actually have to liquid-cool the fabric and the switches, right, as well. So that's what we do in a supercomputer, and that's the kind of capability and the IP we're bringing more broadly.

The third part of it is then you need people that actually know how to deploy these things and manage the services and implement them. And that may sound simple, but it turns out if you pull out the wrong plug when you're servicing one of these things. You can have a pretty significant water leak.

And the other part of servicing it is you actually have to maintain the water. The water running through these things, if you think about it, is getting heated and cooled. It's getting subjected to different thermal dynamics. Sometimes it's getting subjected to light. And so we have to manage that, so we have to chemically treat the water. And we have expertise in all these areas, from manufacturing to system design to the service and deployment.

And then of course there's the ecosystem behind that of components and parts. But I think when you look at it, the three things that we do today, we're well ahead of the market, and we understand how to make this repeatable and scalable. And I see it even in how fast we can deploy a supercomputer today that's liquid-cooled versus what we were able to do just a couple years ago when I...

MR. NERI: The bottom line, we are unique in that space. This would make us attractive because customers who need to deploy at the scale that needs the services, and our HPE services business is a significant point of differentiation. You can go buy a commodity and then have someone else build it for you. We deploy the entire solution for the customer, from beginning to the end. And that IP is very unique for us, and that's where it requires an AI-native architecture. And I think that's something that we're going to leverage as we go forward. So I think we are—

MR. KVAAL: [Interposing] Yeah, why don't we hold it there.

MR. NERI: Yeah, so we appreciate your time and investment, so we are happy to be back at the New York Stock Exchange. We want to thank the New York Stock Exchange for the hospitality. I think we have walked you through our strategy, very detailed. We have given you a lot of information; I know it takes time. But one of the feedback you had given to us is the fact that you want to understand how we compete in the market. I think we have laid the case very clearly in each of the business segments, what is our differentiation. So I hope you got greater insight on the value HPE has captured through our

strategic pivot, which obviously you see in the numbers, through higher-growth and higher-margin areas.

And as the market expands in these areas where we have made investments or we're going to double-down the investment, I think it's making us not just relevant, but also will create greater value for our shareholders. I believe HPE is an excellent opportunity to invest. I think our guidance is logical, considering some of the things we have to deal with H3C or not. But we are very confident in what we're doing. And the reason why I'm confident is because customers are coming to us. That's the bottom line. When you spend the time, I spend, as I do, more than 50% with customers, it gives me the confidence that we can go and execute what we show you and accelerate that value for our shareholders. So thank you. I know there is a little bit of mingling here. Thank you for the audience...

MR. KVAAL: Little mingling.

MR. NERI: On the webcast-

MR. KVAAL: [Interposing] Yeah.

MR. NERI: --Thank you for staying with us.

MR. KVAAL: Not too much though.

MR. NERI: Thank you.

MR. KVAAL: Thank you.

[END RECORDING]