NVIDIA Corporation (NASDAQ:NVDA) Q2 2023 Results Conference Call August 24, 2022 5:00 PM ET

Company Participants

Simona Jankowski - Head of IR

Jensen Huang - President and CEO

Colette Kress - EVP and CFO

Conference Call Participants

C.J. Muse - Evercore ISI

Vivek Arya - Bank of America Securities

Matt Ramsay - Cowen

Toshiya Hari - Goldman Sachs

Aaron Rakers - Wells Fargo

Atif Malik - Citi

Joseph Moore - Morgan Stanley

Stacy Rasgon - Bernstein

Srini Pajjuri - SMBC Nikko

Timothy Arcuri - UBS

Harlan Sur - J.P. Morgan

Ross Seymore - Deutsche Bank

Operator

Good afternoon. My name is Regina, and I will be your conference operator today. At this time, I would like to welcome everyone to NVIDIA's Second Quarter Earnings Call. All lines have been placed on mute to prevent any background noise. After the speakers' remarks, there will be question-and-answer session. [Operator Instructions] [Technical Difficulty]

Simona Jankowski

Conference call for the second quarter of fiscal 2023. With me today from NVIDIA are Jensen Huang, President and Chief Executive Officer; and Colette Kress, Executive Vice President and Chief Financial Officer. I'd like to remind you that our call is being webcast live on NVIDIA's Investor Relations website. The webcast will be available for replay until the conference call to discuss our financial results for the third quarter of fiscal 2023. The content of today's call is NVIDIA's property. It can't be reproduced or transcribed without our prior written consent.

During this call, we may make forward-looking statements based on current expectations. These are subject to a number of significant risks and uncertainties, and our actual results may differ materially. For a discussion of factors that could affect our future financial results and business, please refer to the disclosure in today's earnings release. Our most recent Forms 10-K and 10-Q and the reports that we may file on Form 8-K with the Securities and Exchange Commission. All our statements are made as of today, August 24, 2022, based on information currently available to us. Except as required by law, we assume no obligation to update any such statements.

During this call, we will discuss non-GAAP financial measures. You can find a reconciliation of these non-GAAP financial measures to GAAP financial measures in our CFO commentary, which is posted on our website.

With that, let me turn the call over to Colette.

Colette Kress

Thanks, Simona.

This was a challenging quarter. Total revenue of $6.7 billion was down 19% sequentially and up 3% year-on-year, below the $8.1 billion outlook we provided on our last earnings call. As we indicated in our pre-announcement press release on August 8, we experienced a shortfall to our expectations driven primarily by weaker Gaming revenue. Today, we will share with you more details on our Q2 results and Q3 outlook.

Starting with Gaming. Revenue of $2.04 billion was down 44% sequentially and down 33% year-on-year, reflecting challenging market conditions. As discussed in May, we expected a sequential decline in Gaming revenue due to softness in Europe related to the war in Ukraine and COVID lockdowns in China. The decline in Gaming GPU revenue was sharper than anticipated, driven by both lower units and lower ASPs. Macroeconomic headwinds across the world drove a sudden slowdown in consumer

demand. We implemented programs with our Gaming channel partners to adjust pricing in the channel and to price position current high-end desktop GPUs as we prepare for a new architecture launch.

As noted last quarter, we had expected cryptocurrency mining to make a diminishing contribution to gaming demand. We are unable to accurately quantify the extent to which reduced crypto mining contributed to the decline in gaming demand. While Gaming navigates significant short-term macroeconomic challenges, we believe the long-term fundamentals in Gaming remain strong.

NVIDIA RTX has redefined computer graphics and is now supported by almost 300 games and applications. NVIDIA's GeForce GPUs are the most coveted brand by gamers, representing 15 of the top 15 most popular GPUs on Steam. Gaming has emerged from the pandemic an even more popular form of entertainment and social connectivity. Estimated GeForce sell-through is up over 70% since before the pandemic, and peak concurrent users on Steam are also up more than 70% over the same time period. GeForce NOW registered members now exceed 20 million. This quarter, we added 80 more titles, including the hugely popular Genshin Impact, bringing our total to over 1,350.

Moving to Professional Visualization. Revenue of $496 million was down 20% sequentially and down 4% from a year ago. A sequential increase in mobile revenue was more than offset by lower desktop revenue, particularly at the high end. As macroeconomic headwinds intensified, enterprise demand slowed and OEMs worked to reduce inventory. We expect these trends to persist in Q3. While ProViz is undergoing a near-term adjustment after doubling last year, we believe we have expanded the market opportunity over the last couple of years with AI and Omniverse workloads. We believe hybrid work is here to stay, and with it, the need for collaborative 3D design enabled by professional graphic workstations, both at home and in the office as well as in the cloud.

In June, we announced a partnership with Siemens to enable the industrial levers and AI-powered digital twins, connecting Siemens Xcelerator platform to NVIDIA Omniverse. This connection opens Siemens to the vast ecosystem of NVIDIA Omniverse and NVIDIA to Siemens ecosystem of the world's largest industries. Earlier this month at SIGGRAPH, the premier computer graphics conference, we announced advancements to several foundational technologies of the metaverse, defined as the 3D version of the Internet.

First, NVIDIA Omniverse Avatar Cloud Engine will enable businesses to create and deploy assistants and avatars, transforming interactions across a range of industries. We also unveiled 11 new Omniverse connectors, bringing the total number of connectors to the Omniverse USD ecosystem to 112. And finally, we released SDKs for the new field of neural graphics, which intertwine AI and graphics to help automate the creation of virtual world.

Moving to Automotive. Revenue of $220 million increased 59% sequentially and 45% from the year-ago quarter. Strong growth was driven by auto AI solutions, which include AI cockpit and self-driving revenue, with particular strength in self-driving as new energy vehicle design wins ramp into volume.

We believe Q2 was an inflection point for our automotive revenue as NVIDIA Orin has great momentum. During the quarter, we announced rollout plans of new vehicles from OEM partners, NIO, Li Auto, JIDU and Human Horizons as well as Pony.ai's line of self-driving trucks and robotaxis, all built on NVIDIA DRIVE. Looking forward, we expect our $11 billion automotive design win pipeline to translate to continued growth.

Moving to Data Center. Revenue of $3.81 billion grew 1% sequentially and 61% year-on-year. Although a record, this was somewhat short of our expectations as we were impacted by supply chain disruptions. Revenue from hyperscale customers nearly doubled year-on-year. Sequentially, sales to North America hyperscale and cloud computing customers increased but were more than offset by lower sales to China hyperscale customers affected by domestic economic conditions. Vertical industries grew both, sequentially and year-on-year. Key workloads driving growth include natural language processing, recommender systems, autonomous vehicle fleet, data processing and training and cloud graphics.

Let me share a couple of customer examples. Pinterest transitioned to 100x larger recommender models by moving its inference from CPUs to NVIDIA GPUs. Its ability to deploy a higher-quality model at high throughput and low latency resulted in a 16% increase in engagement, a critical metric for the company, which has over 400 million users and 300 billion images.

And Tesla recently upgraded its supercomputer to use over 7,000 A100 GPUs for autopilot training. From a product perspective, networking led growth this quarter with strong demand from our high-speed Ethernet adapters and design win momentum toward next-generation adopters, including the ConnectX-6 and ConnectX-7. We also see growing interest from cloud service providers for our new Spectrum-4 400 gigabit per second Ethernet networking platform.

Additionally, we are ramping into the upcoming launches of our next generation platforms. The Hopper architecture flagship H100 data center GPU is in production. Grace is our first CPU. Top computer makers, including Dell, HPE, Inspur, Lenovo and Supermicro are adopting the new NVIDIA Grace CPU Superchip and Grace Hopper Superchip to build the next generation of supers. 72% of the systems on the latest top 500 list of the world's fastest supercomputers are powered by NVIDIA, including 31 of 39 new systems. NVIDIA's [Indiscernible] Selene supercomputer ranks at number 8 in the top 500 and is the world's fastest enterprise supercomputer. Moreover, 22 of the top 30 systems on the Green 500 list of the most energy-efficient supercomputers are powered by NVIDIA.

Significant advances in software technologies are key to our platform performance. In the past two years, our A100-based platform has delivered 6x more performance as measured by the MLPerf industry benchmark, largely through new software technologies and optimizations. Last month, we announced an update to the NeMo Megatron framework that can speed up the training of large language models by up to 30%, improving a multi-hundred million dollar AI infrastructure by 30% translates to significant value for customers.

LLM are one of the most important neural networks today, ranging in size from tens of billions to over 1 trillion parameters. Learning from text, they can be used for real-time content generation, tech summarization, customer service chat box and question answering for conversational AI interfaces. Currently, these capabilities are available to early access customers to run on NVIDIA DGX SuperPOD and NVIDIA DGX Foundry as well as in Microsoft Azure cloud with other platforms available soon. We are working with the industry leaders in large language models, a very active and exciting space of AI.

Moving to the rest of the P&L. GAAP gross margin was 43.5%, and non-GAAP gross margin was 45.9%. Gross margin includes $1.22 billion in charges for inventory and related reserves based on revised expectations of future demand and $122 million for warranty reserves. These charges incurred in the quarter reflect purchase commitments that we made during the time of severe component shortages and our current expectation of ongoing macro uncertainty. We believe our long-term gross margin profile is intact.

GAAP operating expenses were up 36% from a year ago and down 32% sequentially as Q1 included a $1.35 billion acquisition termination charge related to the Arm transaction. Non-GAAP operating expenses were up 38% from a year ago and up 9% sequentially. These increases were driven primarily by employee growth costs, as well

as increases in salaries to support our employees during this high inflationary environment, and engineering development of new products coming to market.

We have slowed operating expense growth, balancing investments for long-term revenue growth while managing near-term profitability. Our full year non-GAAP OpEx is expected to grow over 30%. During the first half of fiscal 2023, we returned $5.5 billion to shareholders in the form of share repurchases and cash dividends. We plan to continue share repurchases. We have nearly $12 billion remaining under our authorization through December of 2023.

Let me turn to the outlook for the third quarter of fiscal '23. We expect Gaming and ProViz revenue to decline sequentially as OEMs and channel partners reduce inventory levels to align with current levels of demand and prepare for our new product generation. We expect that decline to be partially offset by sequential growth in data center and automotive. Revenue is expected to be $5.9 billion plus or minus 2%. GAAP and non-GAAP gross margins are expected to be 62.4% and 65%, respectively, plus or minus 50 basis points. GAAP operating expenses are expected to be approximately $2.59 billion. Non-GAAP operating expenses are expected to be approximately $1.82 billion. GAAP and non-GAAP other income and expenses are expected to be an expense of approximately $10 million, excluding gains and losses on nonaffiliated investments. GAAP and non-GAAP tax rates are expected to be 9.5%, plus or minus 1%, excluding any discrete items. Capital expenditures are expected to be approximately $550 million to $600 million.

Further financial details are included in the CFO commentary and other information available on our IR website.

In closing, let me highlight upcoming events for the financial community. We will be attending the Jefferies Conference in Chicago on August 30th and the Goldman Sachs Conference in San Francisco on September 12th. And we will be holding a financial analyst Q&A with management following Jensen's GTC keynote on September 20th. Our earnings call to discuss the results of our third quarter of fiscal 2023 is scheduled for Wednesday, November 16th.

We will now open the call for questions. Operator, can you assist? Would you please poll for questions?

Question-and-Answer Session

Operator

[Operator Instructions] Your first question comes from the line of C.J. Muse with Evercore ISI.

C.J. Muse

I think the question we all have is what is normalized revenues for gaming for you guys? Obviously, this is a challenge to you as well. But curious how you're thinking about it today. Is the fiscal '20 recovery post the first half '19 correction an appropriate framework, or was that inflated by crypto as well? And I guess, as part of that, how do we think about the cascading in of the new product cycle? And is there potential for future reserves needed to be taken if gaming does not meet your new updated outlook? Thanks so much.

Colette Kress

Well, let me start first there and see if I can assist in terms of how to think about after we get through our completion just now of Q2 and what we have provided here for guidance for Q3. Across those two quarters, the Q2 of '23, the Q3 of '23, we have likely undershipped gaming to our end demand significantly. We expect that sell-through or essentially our end demand for those combined two quarters of Q2 and Q3 to be approximately $5 billion.

Now, on top of this, keep in mind that we do have gaming growth drivers to consider for the future. These can include our new gaming product introductions that are around the corner as well as new segments of the market that we plan to reach with our gaming technology to just name a couple. I'll turn it over to Jensen to talk a little bit more about that.

Now, regarding any further types of write-downs on this perspective, we did a thorough assessment with this quarter, not only just looking at what we needed for this quarter, but what we need for the long term. Keep in mind, our inventory provisions and write-downs that we took into account had to reflect some of the purchasing that we did a supply back more than a year ago when we were still in extreme supply shortages in almost all of our products. And so, this was an opportunity for us to resize given the macroeconomic conditions, what we needed in terms of supply. So, our expectations were higher, and we took this opportunity to write them down to what our current expectations are.

I'll turn it over to Jensen to see if he wants to add more.

Jensen Huang

Yes. Thanks, Colette. C.J., our sell-through is off the highs in the beginning of the year, but it's still very solid. In fact, sell-through is -- has increased 70% since pre-COVID, pre-pandemic. And so, it's very clear that gaming is -- the fundamentals of gaming are strong, and this medium is really doing well. Not to mention the gaming platforms are being used -- our gaming PCs are being used for influencers, people sharing content, creating content, V bloggers, VTuber, there's all kinds of new ways of engaging and spending time with video games.

Our strategy is to reduce the sell-in -- reduce the sell-in this quarter, next quarter to let channel inventory correct. Obviously, we're off the highs, and the macro condition turned sharply worse. And so, our first strategy is to reduce sell-in in the next couple of quarters to correct channel inventory. We've also instituted programs to price position our current products to prepare for next-generation products.

Ampere is the most popular GPU we've ever created. It is in the top 15 most popular gaming GPUs on Steam. And it remains the best GPUs in the world, and it will be very successful for some time. However, we do have exciting new next-generation coming and it's going to be layered on top of that. And so, we've taken -- we've done two things. We've reduced sell-in to let channel inventory correct and we've implemented programs with our partners to price position the products in the channel in preparation for our next generation.

All of this we anticipate were working towards a path to being in a good shape going into next year. Okay? So, that's what our game plan is.

Operator

Your next question will come from the line of Vivek Arya with Bank of America Securities.

Vivek Arya

Actually, I just wanted to clarify, Jensen. So, should we assume that a gaming sell-in will kind of stay at these levels into your Q1 or Q2, or depending on new product launch, it might recover? So just wanted to make sure our baseline assumption is set there.

And then, my question is actually similar on the data center. Sales are pretty strong right now, but there is a concern that data center CapEx could be the next shoe to drop in this rolling correction in semiconductors. I'm curious, what's your sense of utilization of your data center shipments? And what is the risk that there could be a correction in the data

center given some of the macro caution expressed by some of the hyperscaler and enterprise customers?

Jensen Huang

Thanks, Vivek. The sell-through -- the sell-through, as I mentioned earlier, of GeForce is solid. The end market gaming demand is solid. It's off the highs, which was really high recently in the beginning of the year. And so, we have -- and because we were building for such a vibrant market, we found ourselves with excess inventory. And so, our strategy is to sell well below -- sell in well below the current sell-through levels in the marketplace to give the channel an opportunity to correct. We'll do that for a couple of quarters or so. We believe that by the end of the year, we'll be in a good shape going into next year. And so, I hope that answers your question. But, the important thing is our sell-in rate is far below what is happening in the market for sell-throughs. The sell-through is solid, has increased 70% since pre-COVID. And so, the gaming market is really quite vibrant.

On the second question, on data center end markets, we hear fairly broadly that GPU supply is in shortage in the cloud. We hear quite broadly that demand for GPU rentals far exceeds current supply. And it's fairly sensible to us, partly because the number of use cases for GPUs in the cloud has grown quite a bit. If you look at one particular segment in just managing -- collecting data and managing the data of the AV fleet and using that data to train AI models, using that data to reconstruct HD maps, the usage of GPUs in the cloud for just that one application has grown a lot. And furthermore, there's the deep learning-based recommender systems has demonstrated such significant effectiveness. And it helps internet service providers to enhance engagement, enhance click-through rate. And so that -- so this particular form of recommender systems is going to really drive a fair amount of data processing and machine learning in the cloud.

And then, of course, over the last several years, a very important model has emerged called transformers. You and I've spoken about this model several times in the past. And it's been found that this transformer model, this large language -- this language model, which when scaled up in size, exhibits really spectacular and effective capabilities for -- to be used to learn skills with either few shots or almost no shot, meaning it could learn skills, it could perform skills that it has never learned because the knowledge was somehow encoded from the large amount of data that it had learned from.

And so, this large language model area of innovation is used in, of course, conversational chat, Q&A summarization, text generation, image generation. But very importantly, it's being used in life sciences for understanding chemistry. We've done

some very important work in this area ourselves called MegaMolBART, understanding proteins, understanding DNA to learn the language of these large -- very, very large, spatially as well as temporally or sequentially types of data.

And so, the impact of this area is really quite worth staying close to. It's called large language models. I think Stanford did a paper that called it the Foundation models that could be used for training all kinds of other types of AIs.

And so, we're seeing a great deal of demand for GPUs in the cloud. We were challenged this quarter with a fair amount of supply chain challenges because as you know, we don't just sell the GPU chip, but these systems are really complex with a large number of chips in the system components that we offer like HGX.

And so kitting -- all of the components that have to come together for us to be able to deliver the final component. And then furthermore, these data centers sit idle until the last piece comes together. And the last piece includes very complicated switches and very complicated NICs and networkings and cables. And so these -- building these high-performance computing data centers at very large scale for the world's cloud is not particularly easy. And so the supply chain challenges have been somewhat disruptive. But the demand is there.

And on top of that, we're ramping into Hopper, which is really a fantastic generation.

Operator

Your next question will come from the line of Matt Ramsay with Cowen.

Matt Ramsay

I wanted to follow up on the data center question that Vivek just asked from a couple of angles. I guess, Colette, the first angle being, in the release, you guys talked about pulling high $200 million of revenue into the July quarter from October, but also with supply chain challenges, maybe some deliveries that were meant for July got pushed back into October. So if you could talk a little bit more about those dynamics and just clarify for us that the October sequentially up data center guidance is actually clean of any pull-ins.

And Jensen, the second part, the moving pieces, networking stronger in data center, U.S. hyperscale stronger, China hyperscale weaker. If you could kind of walk us through the trends that you're seeing into the October and January quarters. And in those sort of

those breakouts and when clarify for us when you think H-100 will really start to drive revenue.

Colette Kress

Great. Thanks for the question. And it's kind of a little bit of an add-on to some of the statements that Jensen was discussing regarding our supply chain and what we're seeing today. Our supply chain during the quarter really was quite difficult, was quite challenging to work through. Our platforms, including HDX, networking chips, cables, switches, were very important to the customers. It's not just about us selling the GPUs. So even though customers orders components themselves, they're looking from us what we may refer to as kits, kits that go with those GPUs for them to stand up their data centers.

We also experienced supply disruptions internally with our logistics and our component availability. Some of our supply arrived very late in the quarter. We had very little time from a logistics and availability to get those things out. Customers were impacted as well by availability of key third-party other components that we weren't offering, which were slowing down some of their deployments. So what we did in our Q2 orders that couldn't be delivered in Q3, given that some of these supply constraints existed, and we had Q3 demand where we did have supply in Q2. So we worked with customers to optimize that supply and demand, and that's what we've disclosed to you.

Jensen Huang

Let me answer the questions about the North American and the China hyperscalers. The Chinese hyperscalers and the Chinese Internet companies really, really slowed down infrastructure investment this year, particularly starting in -- they've been rather slow in building out and really accelerate -- well really slowed down in Q2. This slowdown can't last forever. And the number of new technologies in software, the number of people who are using clouds and the number of cloud services is continuing to grow. And so I fully expect investment to return. They're a very important market for us, a very large market for us. And the fact that North American hyperscalers doubled year-over-year our revenues at North American hyperscalers, and that was offset by declines in China said something about the slowdown in China. And so I don't think that's going to last forever. I think it's going to return.

With respect to Hopper, we're in full production now. And we're racing to get Hopper 2, all of the CSPs are dying to get them. And it goes with our HGX, which is multiple Hoppers on a system tray, it's really a supercomputer in a motherboard, if you will. And it

goes along with it networking gear and switch gear. And so there's the enormous amounts of resources apply from all of the CSPs around the world and ourselves to get Hopper. We expect to ship substantial Hoppers in Q4.

Operator

Your next question will come from the line of Toshiya Hari with Goldman Sachs.

Toshiya Hari

I had a question regarding the $1.22 billion inventory charge, maybe for Colette, on this one specifically. In the CFO commentary, I think you stated that the inventory charge is related to weaker demand in both data center and gaming. I think the gaming side is pretty clear based on your comments so far. Curious what's changed on the data center side? Is it mostly the Chinese hyperscalers that Jensen just spoke to? Or is there something else going on in terms of how you're thinking about demand in data center over the next couple of quarters? And related to this, curious if the delay in Sapphire Rapids at Intel is having an impact on your business in the near to medium term?

Colette Kress

Thanks for the question. So our inventory charges, as we commented, we're taking a thorough look of not only this last quarter as well as the quarter that we're guiding, but looking over the long term of what we need for demand and then what we had in terms of supply. Remember, we had purchased this very early on in the year as we needed to, to drive the commitment of the supply that we already have. And so what is happening here for the data center, we had great, high expectations. We still have very strong, solid growth projection for data center as well. We're going to take this opportunity for some of the prior architecture pieces to write down those given what we see as just a change in terms of our expectations going forward.

So you are correct. There are also pieces in there for gaming. We have written down some silicon and chips as the macroeconomic conditions and you get ready for our future product launches take into account, but there's also components, services and capacity in some of the other drivers that are incorporated in those write-downs.

Jensen Huang

Our Hopper supports previous generation CPUs. But I guess, next-generation GPUs, CPUs, Sapphire Rapids and Genoa after that as well as Graviton. And so -- so we certify

and test across all of the CPUs because the cloud service providers demand it. And they intend to deploy NVIDIA accelerators, NVIDIA Hoppers across a large number of CPUs.

There is no question that the delay is disruptive and a lot of engineers have to scramble. It would have been a lot easier if next-generation CPUs were to have executed more perfectly. However, Hopper goes into an environment with CSPs where they connect our PCI Express connectors to old generation, current generation CPUs as well. And so nobody likes the delay. The next-generation CPUs will trigger a refresh of infrastructure and new servers. And so I'm super excited about them. However, we're going to be able to go to market plentifying with Hopper supporting existing infrastructure.

Operator

Your next question will come from the line of Aaron Rakers with Wells Fargo.

Aaron Rakers

Yes. And a lot of it's kind of been talked a little bit about, but I just -- Colette, I want to think about the numbers a little bit more. I guess, based on the guidance commentary, where I've come up thinking about would be like a 30-plus percent sequential decline in gaming and Professional Visualization and maybe kind of low to mid-single-digit growth in Data Center and Auto. So I guess, is that the right context?

And with that and with Hopper coming down the pipeline, we're going to have a lot of questions around kind of the deceleration of year-over-year growth in Data Center, I believe. Do you think that Hopper, as that comes fully available, it sounds like in fiscal 4Q, that you actually see Data Center growth reaccelerate as that product cycle materializes?

Colette Kress

Yes. Thanks for the question. And that's a pretty good understanding of our guidance. And we do expect, yes, gaming to decrease, not in the dollar amount that it decreased between Q1 and Q2. So that may be of our 2 areas of the decline, our gaming and ProViz, that may be about 3/4 of it associated with that gaming. And then Professional Visualization would probably be about 1/4 of the 2 areas that will decline.

Our Data Center yes, we do expect it to grow. It may grow about what we just saw between Q1 and Q2. We'll continue to look at it. There may be some more opportunity there. And automotive, very similar to our thoughts at the very beginning of the quarter, we are expecting continued growth through each of the quarters of this fiscal year. We

felt that Q2 was an inflection point. So we'll continue to grow into Q3 and hopefully Q4 going forward.

I'll turn it over to Jensen to see his thoughts in terms of Hopper, what Hopper brings to us in Q4 and expectations.

Jensen Huang

The first thing I'd say, Aaron, is that we are selling in or we're selling far below the market demand, far -- excuse me, far below the market sell-through. And the reason for that is to allow the inventory the channel inventory, the OEM inventories to correct. And this allows us to prepare for our next generation. And our next generation has Hopper for compute, but we also have the next generation for computer graphics that will be coming to market.

Hopper is a giant new generation because it is designed to perform this new type of AI model called Transformers. It has an engine inside it called Transformer engine with numerical formats and pipelines that allows us to do a spectacular job on Transformer-type of models, which includes large language models, but it also includes computer vision models that are now able to be processed with this new type of AI model called Transformers. And so I fully expect Hopper 2 to be the next springboard for future growth. And -- and the importance of this new model, Transformer, can't possibly be understated and can't be overstated. This is the impact of this model across robotics, computer vision, languages, biology, chemistry, drug design is just really quite spectacular. And I'm sure that you've been hearing about this new breakthrough in AI, and Hopper was designed for this.

Operator

Your next question will come from the line of Atif Malik with Citi.

Atif Malik

Colette, can you talk about networking versus compute? Were they both supply constrained in the July quarter? And are they both sequentially growing in the October quarter?

Colette Kress

So within our Q2 results, we have been continuing to improve our supply for networking. We have a lot of important products that the CSP needs, many of our customers' needs, and we have been working to really improve that supply. And we were able to set very

strong growth in terms of networking, both sequentially and year-over-year. And as we move into the next quarter, we're going to have to see which is going to be growth larger. We're just going to have to take a lot once we finish that quarter. But our supply for compute is here. But as we've discussed, sometimes it's important that they have many of our other components that we provide in networking at the same time that we are providing the GPUs. So sometimes those are very important for us to deliver together. So we always have to keep that in mind. So it's not always supply constrained, but there are certain parts of it that are.

Operator

Your next question will come from the line of Joseph Moore with Morgan Stanley.

Joseph Moore

Great. I wanted to ask why did the supply constraints hit you guys so hard this quarter? I mean you've done such a great job in the last couple of years outgrowing really everybody in a very challenging supply chain environment. It hasn't tripped you up at all. And it seems like now it's kind of hitting you fairly hard at a time when, in other cases, it's kind of easing. So I'm just kind of curious, what is it about the timing and how long does it sort of take? Or is that a Hopper issue? Is it related to other components? As you talked about, how long does it take to clear those issues up?

Colette Kress

Let me start and see if Jensen wants to add on to it. Our execution has absolutely been phenomenal. When you think about the challenges of we're almost putting together a full data center for our customers and getting it shipped out. So we're no different in the same way that the CSPs are challenging. We're setting up their data centers as we're such an integral part of that. And so networking has been short of supply. These are the same supply issues that some of our CSPs are having. So our supply arrived a little bit late in the quarter for some of our key products that we needed to get out. And putting that together caused some disruption in our logistics and distribution. We were pleased in terms of reaching the leverage of networking that we did, but we did have some challenges this quarter.

Operator

Your next question will come from the line of Stacy Rasgon with Bernstein.

Stacy Rasgon

I wanted to go back to that data center inventory charge. You listed data center first when you talked primarily related to data center and gaming. Can you give us a feeling for how much of that charge was data center versus gaming?

And then to follow up on that, it did sound like to me that it was mostly Ampere and not Hopper change in expectations. Can you verify that or clarify it? And just talk about what's happened with your expectations for Hopper? Have they gone up or down? Or has there been any change at all relevant to that inventory charge?

Colette Kress

Yes, Stacy. Regarding our inventory charges that we had, when you think about what we have in supply, whether it be chips, components, whether it be memory. Remember, a lot of these things can be used interchangeably across the 2. Additionally, the value of our inventory for data center is much different than the value of what we have for gaming from an overall cost perspective. So we're creating capacity opportunity, putting together all of those systems in terms of data center. It is prior architectures. Absolutely, this is not a question regarding anything of our future products coming to market. Nothing on the inventory provision has to do with that. So we took this as looking at the macroeconomic conditions, as we've discussed. Our expectations, our plans were higher. They're still quite solid that we see in demand both for gaming as well as solid for data center, and that will continue. But we did have to just take a rightsizing of that note.

Jensen Huang

Hopper was designed for transformers. The new transformers was going to be important. Nobody could have predicted the profound importance of large language models. Large language models, excitement, innovation, ideas, companies, start-ups, industries, all exceeding everyone's expectations. I don't think anybody could have predicted the impact of Transformers as it scaled up to these giant sizes. There's a fair amount of literature now written about language models that were smallish in the old days, in the beginning several years ago, 3 years ago. And the ones that are in the hundreds of billions and moving towards probably several trillion parameters, the effectiveness of the AI is really quite spectacular. And to have AI that was never trained on a particular skill and yet within 1 shot or 1 shot of trying or even no shots, are able to perform that skill is beyond anybody's expectations, I would think. And so I think the -- the success of Hopper is -- reflects the amount of work and pent-up demand for large training systems that Hopper is going to go into. If that's an indicator, I think Hopper is going to be a spectacular success.

Operator

Your next question will come from the line of Srini Pajjuri with SMBC Nikko.

Srini Pajjuri

I have a question on gross margins, Colette. The 65% non-GAAP number you're guiding to, does that include or assume any additional write-offs on the inventory front?

And then just to expand on that, your mix is probably a tailwind to gross margins given that gaming is down significantly and data center is up a little bit. And I'm just wondering, are there any other offsets? Because I would have thought gross margins could actually be better than your longer-term model because of the mixed tailwinds?

Colette Kress

Yes. Thanks for the question. So our gross margins outside of the inventory charges in Q2 as well as going into Q3 is really about our sales mix that we have and probably also to understand that our sales mix in the next quarter for GPUs is not in the high end. And so that has impacted our gross margin as we move into Q3.

You are correct. We do expect that data center will assist in our gross margins but we also have growth plans in auto. Auto is below our company average, and so that will tend to offset some of those upper bound things that we will see in terms of data center. From time to time, there's always a small amount of scraps that we will have in our gross margin estimates. So nothing material is planned. But there is small scraps that may occur from quarter-to-quarter that are included in our gross margins.

Operator

Your next question will come from the line of Timothy Arcuri with UBS.

Tim Arcuri

Colette, I had a clarification and then a question. So my clarification is whether October gross margins are benefiting at all from the sale of previously written down inventory. And then my question is whether you can give us the enterprise cloud split in data center because it sounds like the mix shifted more towards enterprise in July. And I think investors might want to see that as risky in the face of enterprise clearly slowing. So I'm wondering if you can give us that.

Colette Kress

Okay. No. No, there is nothing in our Q3 regarding those inventory provisions that we took in terms of earning that back, in terms of our Q3. Regarding our split between our hyperscalers and data centers and what we refer to as our vertical industries. They always tend to be about the same, 50% for one, 50% for the others. They're still in about that range. We had discussed that our China hyperscales did not drive growth in terms of sequentially here. And so that did influence in terms of the hyperscales, but still we are approximately in that 50%, 50%.

Operator

Your next question will come from the line of Harlan Sur with J.P. Morgan.

Harlan Sur

On the weakness in ProViz, this is an enterprise-focused business, right? Granted, it's a somewhat narrow vertical market, but it does sort of play into the market concerns that consumer is weak now, enterprise is the next shoe to drop. So is the decline in ProViz attributed to enterprise spending weakness? Or is there some other dynamic? And why or why not is this not a leading indicator for your enterprise and vertical industry segments within your data center business?

Colette Kress

So our Pro Visualization business, just in such a short time, they were in short supply and really trying to feed so much of the industry's desire to both refresh and attracted to the new RTX workstations that were there. We were fueling both mobile as well as desktop for that market and growing quite nicely. Remember, all of last year, doubled the size of ProViz in 1 year. So right now, the OEMs are concentrating on their levels of inventory. We want to make sure that they can get through that inventory. But it -- keep in mind, very similar to our discussion on gaming, there is still solid demand. There is still solid demand. We just have to correct some of the inventory, but we still see both the opportunities that we've created for the market for these to be long-standing.

And your second question, remind me again?

Jensen Huang

Whether the broader enterprise market is -- the verticals are going to be affected by that? I would say, first of all, we don't know. Second of all, unlike like our Workstation business, our ProViz business, there's no installed base. Most of the ProViz sales tend to be tend to be upgrades or replacements from something that has -- our installed base

of 3 or 4 or 5 years that people -- whatever upgrade cycle they happen to have. And so in the case of ProViz, the companies that are buying are ProViz, our ProViz systems likely already have systems that they've been using. And so if they were to tighten up ProViz for whatever reason, the people could continue to use what they have.

In the case of our AI business, there's no real installed base. These are all brand-new things that people are growing into. And the productivity benefits or the cost savings benefits of using autonomous systems is fairly profound. And it's not so much that the demand isn't out there. Everybody would like to be more productive. Everybody would like to save more money. Everybody would like to move faster. It's just that AI understanding and AI's use is still spreading. And so we're delighted by the rate of growth and the rate of adoption of enterprise. My sense is that our AI business and our Viz business have very different characteristics for that reason.

But what Colette said earlier is about our ProViz businesses last quarter is absolutely true, which is OEMs realizing that the end market is slowing and taking the opportunity to correct their inventory.

Operator

Your final question will come from the line of Ross Seymore with Deutsche Bank.

Ross Seymore

Let me ask a question. I just had a longer-term question about once your gaming business normalizes, with the absence of crypto in a general sense and with the merge coming, et cetera, how do you view the pricing environment? I know you guys really weren't raising prices like we saw in the MSR key premiums in the aftermarket. But generally, your mix was quite rich over the last year or 2. You're going to have lovely coming in that will obviously help the mix sequentially versus the last couple of quarters. But how do you think about it normalizing? Is that $2.5 billion per quarter sell-through rate comes, is the ASP mix across your stack about the same? Or does it go down because of the absence of that crypto tightening dynamic?

Jensen Huang

I would say that without crypto dynamic, the mix would go down. However, the overall trend long term, the ASP is drifting up. And the way to think about that is a game console, when -- my first game console was $99. Lately, game consoles are selling for about $599. And the reason for that is because it's more useful than ever. You use your gaming console for your greatest form of entertainment, and you use it for a very, very

long time. And GeForce essentially is a game console inside your PC. And we've always believed that the ASP of GeForce should drift towards the average selling price of a game console. And so it should be something along the lines of $500 or so roughly at this time.

We also have GeForce in the cloud. And because GeForce in the cloud is hosting many gamers simultaneously, it tends to want to be a much more powerful GeForce. And so our cloud gaming GeForce tends to be -- our cloud gaming graphics tends to be a much higher end.

And so -- and then, of course, there's the design aspect of it. Most designers and most creators are able to use GeForce these days. And they use their PC to create content, and much of that content goes into video games and/or they're using video games to create their artistic content. And so the GeForce is not just for gaming for them. The GeForce is essentially their creative work station as well. And so there are several dynamics that are causing the ASP of GeForce to go up, and we've been seeing this trend for several years now.

Operator

Thank you. I will now turn the call back over to Jensen for closing remarks.

Jensen Huang

Thanks, everyone. We're navigating our supply chain transitions in a challenging macro environment. In Gaming, our partners and ecosystem are responding to a sudden slowdown in consumer demand and correcting channel inventory. Still, the fundamentals of gaming are strong. We'll get through this over the next few months and go into next year with our new architecture. I look forward to telling you more about it at GTC next month.

In Data Center, AI where computers are helping us write software that was impossible before is driving a computing revolution and transforming every industry. NVIDIA's leadership in full stack data center scale, accelerated computing has made us the ideal partner for companies racing to leverage the power of AI. Even with the current macroeconomic headwinds, demand for our data center products have never been stronger.

The next wave of computing is coming. With AI and 3D graphics advances, developers will extend the Internet with virtual world overlays that connect to the physical world. This next evolution of the Internet is called metaverse. We created Omniverse to

connect the digital and physical world and be an open platform for creating and operating metaverse applications. The immediate applications for Omniverse span product design, manufacturing and operations. Omniverse is off to a great start.

Our automotive revenue is inflecting, and we expect it to be our next $1 billion business. Autonomous driving is one of the biggest challenges AI can solve, and computing opportunity for us spans the data center to the car. Autonomous driving will transform the auto industry into a tech industry. Automotive is one of the first to transform into a software-defined tech industry that all industries will be. We're building NVIDIA AI and NVIDIA Omniverse to be the engines for the world's enterprise to become software-defined, AI-powered technology companies.

I look forward to next month's GTC conference, where we will share new advances of RTX reinventing 3D graphics and gaming. AI's continuing breakthroughs and building the metaverse, the next evolution of the internet. So join us. We look forward to updating you on our progress next quarter. Thank you.

Operator

This concludes today's conference call. You may now disconnect.