

12-Aug-2025

CoreWeave, Inc. (CRWV)

Q2 2025 Earnings Call

CORPORATE PARTICIPANTS

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

OTHER PARTICIPANTS

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Kash Rangan

Analyst, Goldman Sachs & Co. LLC

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Raimo Lenschow

Analyst, Barclays Capital, Inc.

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Michael Turrin

Analyst, Wells Fargo Securities LLC

Gregg Moskowitz

Analyst, Mizuho Securities USA LLC

Tyler Radke

Analyst, Citigroup Global Markets, Inc.

Brad Sills

Analyst, BofA Securities, Inc.

Mike Cikos

Analyst, Needham & Co. LLC

MANAGEMENT DISCUSSION SECTION

Operator: Thank you for standing by. My name is Tina, and I will be your conference operator today. At this time, I would like to welcome everyone to the CoreWeave Second Quarter 2025 Earnings Call. All lines have been placed on mute to prevent any background noise. After the speakers' remarks, there will be a question-and-answer session. [Operator Instructions]

It is now my pleasure to turn the call over to Deborah Crawford, Vice President of Investor Relations. You may begin.

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

Thank you. Good afternoon and welcome to CoreWeave second quarter 2025 earnings conference call. Joining me today to discuss our results are Mike Intrator, CEO; and Nitin Agrawal, CFO.

Before we get started, I would like to take this opportunity to remind you that our remarks today will include forward-looking statements. Actual results may differ materially from those contemplated by these forward-looking statements. Factors that could cause these results to differ materially are set forth in today's earnings press release and in our quarterly report on Form 10-Q filed with the SEC.

Any forward-looking statements that we make on this call are based on assumptions as of today, and we undertake no obligation to update these statements as a result of new information or future events.

During this call, we will present both GAAP and certain non-GAAP financial measures. A reconciliation of GAAP to non-GAAP measures is included in today's earnings press release. The earnings press release and an accompanying investor presentation are available on our website at investors.coreweave.com. A replay of this call will also be available on our Investor Relations website.

And now, I'd like to turn the call over to Mike.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

Thanks, Deborah, and good afternoon, everyone. CoreWeave had a standout second quarter as we continue our hyper growth journey against the backdrop of unprecedented demand for our AI cloud services.

Adoption is expanding rapidly with the enterprise increasingly viewing AI as a strategic imperative and CoreWeave as the force multiplier that enables adoption, innovation and growth for training, as well as inference workloads.

As a result, revenue grew better than expected 207% year-over-year to \$1.2 billion for the second quarter, with adjusted operating income of \$200 million. This marks the first quarter in which we reached both \$1 billion in revenue and \$200 million of adjusted operating income. Scaling our capacity and services remains key ingredient for our success in this structurally undersupplied market.

To that end, we ended the quarter with nearly 470 megawatts of active power, and we increased total contracted power approximately 600 megawatts to 2.2 gigawatts. We are aggressively expanding our footprint on the back of intensifying demand signals from our customers, ensuring that we maintain a durable, multiyear runway for growth.

We are now on track to deliver over 900 megawatts of active power before the end of the year. We ended the second quarter with \$30.1 billion in contracted backlog, up \$4 billion from Q1, and doubling year-to-date. This includes not only the \$4 billion expansion with OpenAI we previously discussed, but new customer wins ranging from large enterprise to AI start-up.

Importantly, we've also signed expansion contracts with both of our hyperscale customers in the past eight weeks. Our pipeline remains robust, growing and increasingly diverse, driven by a full range of customers from media and entertainment, to healthcare, to finance, to industrials and everything in between.

The proliferation of AI capabilities into new use cases and industries is driving increased demand for our specialized cloud infrastructure and services. For instance, while its early stages in the first half of 2025, we saw more than a 4x increase in our VFX cloud service product, Conductor, and entered a multiyear contract for NVIDIA's GB200 NVL72 system with Moon Valley, an AI video generation startup that lets filmmakers craft professional-grade clips with granular cinematic control.

We are seeing an increased adoption in the financial services sector as we expand our relationship in proprietary trading like Jane Street, and are adding Mega cap bank clients like Morgan Stanley and Goldman Sachs. We are also seeing significant growth from healthcare and life science verticals, and are proud of our partnership with customers like Hippocratic AI, who builds safe and secure AI agents to enable better healthcare outcomes.

In short, AI applications are beginning to permeate all areas of the economy, both through start-ups and enterprise, and demand for our cloud AI services is aggressively growing. Our cloud portfolio is critical to CoreWeave's ability to meet this growing demand. Our focus on delivering the industry's most performant, purpose-built AI cloud infrastructure makes us the platform of choice for both training and inference across incumbent AI labs and new entrants alike.

We're helping these customers redefine how data is consumed and utilized globally as their critical innovation partner, and we are being rewarded for our efforts as they shift additional spend to our platform. We continue to execute and invest aggressively in our platform up and down the stack to deliver the bleeding edge AI cloud services, performance and reliability that our customers require to power their AI innovations.

For example, during the second quarter, we delivered NVIDIA's GB200 NVL72 and HGX B200 at scale deployments fully integrated into CoreWeave's Mission Control for reliability and performance management. Mission Control continues to be the cornerstone of CoreWeave's ability to scale at breakneck speed, building a fully automated and rigorous process for cluster lifecycle management with unmatched visibility for our customers.

In addition to CAIOS, we began our private preview of an innovative archive tier object storage product with automatic tiering and industry-leading economics, and with a simplified cost structure that makes optimizing storage costs for start-ups and enterprises seamless. As a result, customers are shifting petabytes of their core storage to CoreWeave in the form of multiyear contracts.

We are providing support for additional third-party storage systems tightly integrated into CoreWeave's technology stack with large scale production deployments of VAST, WEKA, IBM Spectrum Scale, DDN and Pure Storage.

With Weights & Bias (sic) [Weights & Biases] (00:07:56), we deliver an integrated full-stack observability feature, giving researchers immediate feedback to diagnose the factors impacting performance and reliability of their AI workloads from the data center through network fabrics and storage, GPUs and up to their machine learning code.

We launched the CoreWeave and Weights & Bias (sic) [Weights & Biases] (00:08:20) Inference Service, utilizing our incredibly reliable compute platform to power a research-friendly API for state-of-the-art AI models including OpenAI's new open source model, Meta's Llama 4, DeepSeek, Kimi K2 and QnA3.

This new product allows customers to easily bring AI inference into production on their applications with tight integration into our Weave product, ensuring visibility into the service, quality and safety. We continued our investment in SUNK, Slurm on Kubernetes, which is used by many of the largest AI labs and enterprises in the world, providing improved identity federation, research segmentation and scale.

We began introducing flexible capacity products to help our customers better manage their end customer demand. In addition to our on-demand and reserved inference offering, our spot product is in customer preview, and will be introduced additional capacity product over the second half of the year.

We also saw significant growth in our backbone and networking service as one of our largest AI lab customers leveraged our networking backbone to connect its multi-cloud inference infrastructure. Our product development roadmap is robust, and we are excited to announce new cloud services and capabilities over the remainder of the year, that will further accelerate growth within the AI ecosystem and empower customers to meet their evolving business needs.

We have entered new parts of the capital markets and accessed new pools of capital, driving our cost of capital lower. We priced both of our inaugural and second high yield bond offerings in the past three months. These transactions were upsized due to strong demand and were priced at lower interest rates.

More recently, we closed on a landmark secured GPU financing with many of the world's leading banks, a novel financing structure that CoreWeave has pioneered. As evidenced by these transactions, our access to the capital markets not only remains robust, but is deepening. We are grateful for this support in our mission, and expect to continue to access less expensive capital sources as we continue to execute.

We will continue to verticalize our platform and enhance our control, efficiency and differentiation fueled by our investment both up the stack as you saw with our acquisition of Weights & Biases last quarter, and down the stack as highlighted by our proposed acquisition of Core Scientific last month.

Our ability to scale state-of-the-art infrastructure will further be bolstered by the more than \$6 billion data center investment we've announced in Lancaster, Pennsylvania, as well as a large data center project in Kenilworth, New Jersey, that we are co-developing via a joint venture with Blue Owl. These new sites are perfect examples of our broader data center strategy, which allow us to provide a mix of both large-scale training and low latency inference compute across the country.

Now, I'd like to come back to our proposed acquisition of Core Scientific. We believe the combination will accelerate value creation for shareholders of both companies. Both the CoreWeave and Core Scientific management teams and boards have evaluated this transaction extensively and concluded this is the best for both companies and their shareholders.

The rationale behind the deal is quite simple and powerful. Verticalization creates tremendous operational and financial efficiencies that will strengthen our ability to serve our customers at scale. Owning the infrastructure will allow CoreWeave to scale faster and more efficiently. The integration of Core Scientific meaningfully advances our capacity to operate one of the largest and most sophisticated AI cloud platforms in the world.

Upon closing, CoreWeave would own approximately 1.3 gigawatts of gross power capacity across Core Scientific's national data center footprint, with an incremental 1 gigawatt or more available for future expansion. This scale enhances our flexibility to take on new projects and meet accelerated customer demand.

In addition, the acquisition would drive the immediate elimination of more than \$10 billion in future lease liability overhead, as well as a more streamlined and efficient operating model. As a result, we anticipate \$500 million in fully ramped annual run rate cost savings by the end of 2027, benefiting both the Core Scientific and CoreWeave shareholders directly.

Vertical integration will allow us to finance infrastructure more efficiently, furthering one of our key objectives in lowering our cost of capital and enabling us to grow in a more capital efficient manner. We and Core Scientific look forward to discussing the transaction with you in the months ahead. Our respective teams are already engaged in pre-integration planning to ensure we're ready to hit the ground running.

To that end, we are executing with pace and purpose amidst the market in which the supply/demand imbalance is only deepening, as new enterprise adopters increasingly compete with large AI labs for limited capacity and services. We're building on our leadership across all key success criteria from power, access to AI cloud service performance to revenue and backlog growth. And we will keep getting stronger as we verticalize our data center infrastructure and cloud services. I am excited about the momentum that we are building upon. I want to thank our customers, teams and business and financial partners for making it possible.

Now here's Nitin.

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Thanks, Mike; and good afternoon, everyone. Our strong second quarter results highlight the unprecedented demand environment we are seeing and our continued execution to rapidly scale our AI cloud platform to meet that customer demand. Our growth continues to be capacity-constrained with demand outstripping supply.

Since our Q1 call, we have signed expansion contract with both our hyperscaler customers. We also closed our acquisition of Weights & Biases and announced our proposed Core Scientific acquisition. In addition, we successfully raised \$6.4 billion in the capital market through two high yield offerings and a delayed draw term loan, all of which have opened access to new capital pools at an increasingly lower cost of capital.

Turning now to Q2 results. Q2 revenue was \$1.2 billion, growing 207% year-over-year, driven by strong customer demand. Revenue backlog was \$30.1 billion, up 86% year-over-year and doubled year-to-date. While our revenue backlog is expected to scale rapidly over time, growth rates will fluctuate from quarter-to-quarter given the nature of our large committed contract business model, timing and size of new contract signings and revenue recognition.

Operating expenses in the second quarter were \$1.2 billion, including stock-based compensation expense of \$145 million. We continue to ramp our investments in data center and server infrastructure to meet growing

customer demand, which contributed to the increase in our cost of revenue and technology and infrastructure spend in Q2. In addition, the increase in sales and marketing was largely driven by marketing spend to accelerate new customer acquisition and raise awareness of our differentiated capabilities. The increase in G&A was largely driven by professional services.

Adjusted operating income for Q2 was \$200 million compared to \$85 million in Q2 2024. Our Q2 2025 adjusted operating income margin was 16%. Net loss for the quarter was \$291 million compared to a \$323 million net loss in Q2 of 2024. Interest expense for Q2 was \$267 million compared to \$67 million in Q2 of 2024 due to increased debt to support our infrastructure scaling, partly offset by lower cost of capital. Adjusted net loss for Q2 was \$131 million compared to a \$5 million adjusted net loss in Q2 of 2024. The adjusted net loss was impacted by increases in interest expense due to scaling of our infrastructure, partly offset by growth in adjusted operating income.

Adjusted EBITDA for Q2 was \$753 million compared to \$250 million in Q2 of 2024, scaling more than 3x year-over-year, and our adjusted EBITDA margin was 62%, roughly in line with Q2 of last year.

Turning to capital expenditures. CapEx in Q2 totaled \$2.9 billion, which is up over \$1 billion quarter-over-quarter as we scaled rapidly to meet our accelerating customer demand. We are executing at a massive scale and the demand continues to outpace supply. As a reminder, CapEx consists primarily of investments in property and technology equipment and is calculated as a change in gross PP&E minus the change in construction in progress. Construction in progress represents infrastructure not yet in service, so it's not yet revenue-generating. In addition, the timing of data center capacity coming online and generations of GPUs being placed into service could drive significant variation quarter-to-quarter, an example of which you will see in our Q4 CapEx ramp.

Now, let's turn to our balance sheet and strong liquidity position. We have designed our capital structure to enable rapid scaling. As of June 30, we had \$2.1 billion in cash, cash equivalents and restricted cash. Other than payments on OEM vendor financing and self-amortizing debt through committed contract payments, we have no debt maturities until 2028. As Mike mentioned, we continue to see strong success in the capital markets.

Growing at a rapid pace and executing at scale requires a unique and sophisticated approach to securing the funding required. CoreWeave continues to be not only the leading AI technology partner, but also the leading innovator in financing the infrastructure required by the world's most advanced AI labs and enterprises. Since the beginning of 2024, we have secured over \$25 billion of debt and equity to fund the build-out and scale the leading AI cloud platform. In May, we launched and closed our first unsecured high yield offering of \$2 billion, which was upsized by \$500 million due to strong demand. In July, we reentered the market and raised an additional \$1.75 billion, also oversubscribed at a lower interest rate.

More recently, we closed our third delayed draw term loan facility. This \$2.6 billion facility completes the financing for the \$11.9 billion OpenAI contract we announced in March. Notably, the transaction was completed at a cost of capital of SOFR plus 400, a 900 basis point decrease from the non-investment grade portion of our prior facility DDTL-2 and was the first one to be fully underwritten by top tier banks. Together, these financings highlight our ability to drive a sustained reduction in our cost of capital and the increasing depth of access we have to the capital markets, both of which were stated goals during our IPO.

Turning to tax. Again in Q2, we recorded an income tax provision despite a net loss due to impacts from non-deductible items and the valuation allowance on net deferred tax assets. Our tax rate might fluctuate significantly in the future due to similar factors.

Now turning to guidance for Q3 and for full year 2025. As Mike mentioned, we are seeing an acceleration of customer demand and our pipeline remains robust and increasingly diversified. We are still operating in a structurally supply-constrained environment, where demand far outstrips supply for our products and services. Our operations and engineering teams are working relentlessly to deploy more capacity faster for our customers.

With the strong demand backdrop, we expect Q3 revenue in the range of \$1.26 billion to \$1.30 billion. In addition, we anticipate Q3 adjusted operating income between \$160 million to \$190 million as we are quickly ramping our capacity to meet customer demand. As we have discussed earlier, as we deploy scaled capacity and bring large chunks of capacity online, we incur some costs prior to revenue generation. The scale of our deployments relative to our base implies that these costs ahead of revenue have a short-term impact on our margins.

We expect our Q3 interest expense to be in the range of \$350 million to \$390 million, impacted by increased debt to support our demand-led CapEx growth, partly offset by increasingly lower cost of capital. We expect our CapEx for the third quarter to be \$2.9 billion and \$3.4 billion. In addition, like last quarter, we expect stock-based compensation to remain slightly elevated throughout the year for grants issued in connection with the IPO and incremental hiring to support our growth.

Moving to full year guidance. For the second quarter in a row, we are raising our full year revenue guidance. For 2025, we now expect revenue in the range of \$5.15 billion to \$5.35 billion, a \$250 million increase from our prior guidance of \$4.9 billion to \$5.1 billion, driven by continued strong customer demand. We expect adjusted operating income in the range of \$800 million to \$830 million, unchanged from our prior guidance as we remain cost disciplined while rapidly scaling our deployments at an unprecedented rate to end the year with over 900 megawatts of active power. We expect CapEx in the range of \$20 billion to \$23 billion, unchanged from our prior guidance in the backdrop of continued strong customer demand. A significant portion of our full year CapEx will fall in Q4 due to the timing of go-live dates of our infrastructure.

We had an outstanding first half of the year and our outlook remains strong. We are entering the second half of the year in an excellent position with strong execution in delivering at scale for our customers, as well as execution in the capital markets and a robust backlog coupled with a very healthy demand pipeline. As we move into the second half, we'll continue investing to meet the needs of our growing customer base, while reinforcing our leadership in this transformational market.

Thank you to our investors and analysts for your support and engagement. We look forward to updating you on our progress in the quarters to come.

With that, we will move to Q&A.

QUESTION AND ANSWER SECTION

Operator: [Operator Instructions] Our first question comes from the line of Keith Weiss with Morgan Stanley. Please go ahead.

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Excellent. [indiscernible] (00:25:50-00:26:08)

Q

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Keith, we can't hear you. Operator?

A

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

Perhaps we can go to the next question and then we can come back to Keith.

A

Operator: Yes. Our next question comes from the line of Kash Rangan with Goldman Sachs. Please go ahead.

Kash Rangan

Analyst, Goldman Sachs & Co. LLC

Thank you very much. Congrats on a really spectacular finish to second quarter. Wondering if you could talk about the renewal of the hyperscaler contracts. I think one of the two is a particularly larger one. And I'm curious to see if this means that you have greater confidence that they will renew, not just expand the interim motion, but how is it more likely that they renew the big contract that they first signed with you in 2024?

Q

And one thing for Nitin, how do you look at the tweaks that you can put into the business to achieve even better return on assets as the company continues to lower its cost of capital? When you look at the core deployment model, maybe that has something to do with how quickly you can translate the bookings into revenue. Clearly, you saw upside in the quarter on that front. But what are the things that the company has uncovered to continue to give you conviction that the company can earn a higher and higher rate of return on your capital going forward as it translates into those revenue line item. Thank you so much.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

Thank you, Kash, and I appreciate the comments on the second quarter. We're really excited about how we've closed out the first half of the year. When we think about contracts with our hyperscaler clients, for that matter really with any of our clients, we generally don't focus on the concept of renewals. We focus on the concept of expansion. And the reason that we focus on the concept of expansion is because, generally speaking, the clients are purchasing hardware that is appropriately state-of-the-art for their use case. And as new hardware comes out, as new hardware architectures are released, they tend to come back in and purchase the same top tier infrastructure in their next renewal. And so, we're excited about renewals when we get them with our hyperscale clients, we're excited about the renewals that we get with any of our clients across the board.

A

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

A

Thanks, Mike. And Kash, to second part of your question, there are few things that we are already executing on. You've seen us kind of acquire Weights & Biases, which is our attempt to go upstack and deliver more value-added services for our customers. You've seen us, with our proposed acquisition, of going down the stack and verticalizing and continuing to get cost savings. We've talked about the anticipated fully ramped \$500 million by the end of 2025 (sic) [2027] in terms of savings. In addition to that, we also talked about how we continue to scale rapidly for our customers and continue to reduce the time from when we start deployment to when customers go online. In addition, we remain cost conscious and disciplined across every vector in our business as we continue to scale this business at an unprecedented rate. All of those factors are working great for us and we continue to deliver great results for the company.

Kash Rangan

Analyst, Goldman Sachs & Co. LLC

Q

Thank you very much.

Operator: Our next question comes from the line of Keith Weiss with Morgan Stanley.

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Q

All right. Can you guys hear me now?

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

A

Yes.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Got you.

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Q

Excellent. All right. Sorry about that technical difficulty. Congratulations on a fantastic Q2. You guys are really putting a lot of emphasis on the growth side of the equation there. So, great to see. I want to ask one question on the kind of demand side of the equation, one on the supply side of the equation. On the demand side of the equation, I think lot of investors have the impression that CoreWeave is handling a lot of training revenues, but a lot of the stuff that, Mike, you were talking to in terms of what you guys are doing with the software, as well as the customers coming in the door speak to doing more inference, more applications being built on the platform. So, can you talk to us a little bit about sort of the mix of business that you're seeing and also the fungibility of the platform, the ability to handle both the pre-training and the inference workloads over time?

And then on the supply side of the equation, you talked about being supply-constrained. Can you give us some sense of where the most acute supply challenges are? Is it on the chip level? Is it on the power level? Like, where do you guys expect to see those constraints in the near term and how much of that can you guys work against? Like, is there any fungibility of where you can more quickly or where would it take longer to solve those supply constraints?

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Sure. Thank you for the question. Let me start off with some comments around our infrastructure and the way in which we see our clients are consuming the compute that we're able to provide. So, when we build our infrastructure, we really build our infrastructure to be fungible, to be able to be moved back and forth seamlessly between training and inference, right. Like, our intention is to build AI infrastructure, not training infrastructure, not inference infrastructure. It's really infrastructure that allows our clients to be able to support the workloads that they need to be able to drive to be successful.

We have seen a massive increase in our workloads that are being used for inference and we're able to monitor that by the profile that the power is being consumed within the data centers. So when you have big training runs that come on and off, that's a step function of power consumption, either up or down, as opposed to when you are using your compute for inference, which is much more incremental in its nature. In addition to that, the infrastructure that we're building has been increasingly been used for a chain of reasoning, which is driving a substantial amount of consumption on the inference level. And that's very exciting for us. As I always say, inference is the monetization of artificial intelligence and we are extremely excited to see that use case expanding within our infrastructure.

On the second question in terms of the supply side, at the end of the day, right now, it's the powered shells that are the choke point that is causing the struggle to get enough infrastructure online for the demand signals that we are seeing, not just within our company. It's the massive demand signals that you're seeing across the industry. And at the end of the day, what we are looking at, and I think what you're hearing across the board is that this is a structurally supply-constrained market. It is a market that is really working hard to try and balance. And there are fundamental components at the powered shell, at the power in terms of the electrons moving through the grid, at the supply chains that exist within the GPUs, the supply chains that exist within the mid-voltage transformers, there's a lot of different pieces that are constrained. But ultimately the piece that is the most significant challenged right now is accessing powered shells that are capable of delivering the scale of infrastructure that our clients are requiring.

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Q

That's super helpful. Thank you so much.

Operator: Our next question comes from the line of Mark Murphy with JPMorgan. Please go ahead.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

Thank you so much. Congratulations on a robust RPO backlog figures and declining cost of capital. It's a great combination. Mike, we've heard commentary that many governments actually around the world want to build their own version of Stargate and that they've begun to reach out. Can you comment on any developments with respect to some of the sovereign governments that want to build modern AI data centers? And what do you think might determine whether they're comfortable using a US-based provider such as CoreWeave? And then I have a quick follow-up.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Yeah. So, it's a very broad question. And you're going to have different jurisdictions, different sovereigns that are going to react differently to that question. What we have seen is that many of the sovereigns are really looking for best-in-class technical solutions to allow them to build the infrastructure that will allow their aspirations within artificial intelligence to be as successful as possible. We have a tremendous number of sovereigns that are beginning and discussing and talking through how to go about doing this, what technology to use, what software stack to use, where it should be placed, right up and down the line. And we are very confident that we will continue to expand our footprint within the sovereign cloud universe.

There are other jurisdictions that are going to be less welcoming to tech coming out of the US. And that's just the nature of the way the world is going to unfold for this. We've had some success in Canada. We're really excited about our partnership with Cohere up there. We think they're doing a wonderful job and that infrastructures is really well-positioned to be successful. We've done a really good job expanding infrastructure across Europe. We feel like we'll be able to reach clients across the European theater, and we look for our clients to lead us into new jurisdictions where they will become the anchor tenants that will allow us to expand the build that we do and the software delivery systems that we create in order to let them become as successful as they would like to be within the AI infrastructure component of the market.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

Okay. Okay. Understood. And as a follow-up, I believe you said CoreWeave signed expansion contracts with both hyperscaler customers in the past few weeks, and just since it's August 12, could you clarify – did you mean that those expansions are already reflected in the Q2 backlog figures? In other words, you're saying that you did incremental business in the month of June or did you mean that those expansions were signed in July and August?

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

A

Thanks, Mark, for your question. One of those contracts was signed in Q2 and is reflected in the Q2 revenue backlog number. The other one was signed in Q3 and will be reflected in our Q3 revenue backlog number.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

And is there any sense of scale on those, Nitin, or whether it's core GPU services versus an extension into Weights & Biases or are you unable to give that kind of detail yet?

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

A

It does include services elements of our portfolio and we'll give a wholesome update in Q3 earnings on the revenue backlog at the end of Q3.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

These contracts are for GPU compute.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

Excellent. Great to hear. Thank you so much.

Operator: Your next question comes from the line of Raimo Lenschow with Barclays. Please go ahead.

Raimo Lenschow

Analyst, Barclays Capital, Inc.

Q

Obviously, we have like this big debate out about this imbalance of demand and supply, and you talked about it a little bit. Like, from listening to you, it sounds more structural, i.e., that is kind of out longer. Can you talk a little bit about that? Because we obviously have like Microsoft who was like maybe we are in balance soon, but then they pushed it up by another six months. Listening to you sounds a little bit longer. Kind of what are the data points for you on that one? And then the other follow-up I had was, as you do more inference, how important is latency and hence location of data centers? That's a number the data is coming up a lot with us. Thank you. Congrats.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Sure. So, we have been unwavering in our assessment of the structural supply constraint that exist in this market. I think that there are other entities that have repositioned, restated and rethought how they are going to deliver infrastructure and when they are going to deliver infrastructure. But we have never wavered from our belief that the market is structurally supply-constrained, and that is based on our discussions and relationships with the largest, most important consumers of this infrastructure in the world. And so, I can't speak to how other organizations are thinking about it. I can only speak to it from what our position is based on our relationships with the buyers that come in looking for the specific solution that we provide. And that is that this market has significant structural supply constraints.

As far as the latency goes, I would encourage you to think about latency through a lens of use case, right. And so, when you are in a chain of reasoning, query, latency is not particularly important. The compute is going to be more impactful than the latency or the latency and the relative distance to the query. If you're in a different type of workload, latency becomes more important. Our approach has been since the early days to try and place our infrastructure as close to the population centers as we can in order to have the optionality associated with a low latency solution. Having said that, as we move through this cycle of developing artificial intelligence, as we see new models coming out, as we see chain of reasoning gain more traction, there's definitely going to be significant demand for latency insensitive workloads that will be able to live in more remote regions.

Raimo Lenschow

Analyst, Barclays Capital, Inc.

Q

Very clear. Thank you. Congrats.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Thank you.

Operator: Our next question comes from the line of Brad Zelnick with Deutsche Bank. Please go ahead.

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Q

Great. Thanks so much. And I'll echo my congrats as well. Mike, my question follows Keith's and Raimo's about inference. How should we think about the economics of inferencing versus training? And then I have a follow-up as well. Thanks.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Yeah. So, look, for our business model, the inference consumption and the training consumption, the economics are identical. The overwhelming majority – and we spoke about this in our last earnings call. The overwhelming majority of our infrastructure has been sold in long-term structured contracts in order to be able to deliver compute to our clients that need to consume it for training and for inference over time. And so, we don't see a real fluctuation in the economics associated with inference or training.

Having said that, I think that it stands to reason to think that when a new model is released and there is a rush to explore the new model, to use the new model, to kind of drive new queries into it, you will see a spike in demand within a given AI lab that may cause there to be a spike in the short-term pricing associated with inference. And we see those. But as we said before, the on-demand component of compute is a very small percentage of our overall workloads. And we are observing inference cases on older generations of hardwares, the A100s, the H100s. They're still being recontracted out. They're being bought on term in order to serve the inference loads that people continue to have, continue to see and need compute to be able to serve.

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Q

Thanks. And that's actually relates to my follow-up question because in your prepared remarks, you talked about flexible capacity products coming online in the back half and a spot product in customer preview. Can you just expand on what that looks like, maybe different than what you're already offering? What GPU generations will be available and how we might think about pricing versus reserved or the take-or-pay style contracts that you more typically do?

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Yeah. So, look, we're going to continue to build up our on-demand and spot pricing offering, right. It's going to take time. The biggest challenge that we have is that every time we're able to build capacity, it is immediately consumed by one of our existing or a new client that wants to expand their exposure to additional compute to be able to serve their models. And so, that has been a continual challenge for us. I guess, it's a good problem to have, but it's a problem for us. But that product, we're working diligently to be able to expand that capacity so that we're able to provide more of a spot product. A big part of that, just so you understand, is so that we can identify new users of compute, identify new companies that are coming into existence, identify new use cases that need compute so that we can build services that are appropriate for them, that allow them to build their businesses and sell their product into the market. And so, we really want to be able to do that and we want to be able to have that offering, but it is challenging in a market that is so demand-constrained.

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Q

Makes a lot of sense. Thank you.

Operator: Our next question comes from the line of Michael Turrin with Wells Fargo. Please go ahead.

Michael Turrin

Analyst, Wells Fargo Securities LLC

Q

Hey great. Thanks very much. Appreciate you taking the question. You mentioned we'll see some variability on the backlog number \$30 billion, nearly 2x where you were a year ago, but also fairly consistent with where you were last quarter when you add in the OpenAI expansion.

So, I think it would just be useful as we're all getting CoreWeave, just if you could help us calibrate a bit more on what to expect from that metric going forward. How often – is it the case that you can find a customer scale to move the needle sequentially there? And where does that \$30 billion sit relative to the opportunities you still see in front of you? Thanks very much.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Thank you. So, first of all, it's important to understand that the demand for compute that we're seeing is from our largest, most important clients, is expanding in scale, in magnitude, this is a planetary rebuild of the infrastructure that they require in order to be able to deliver their products to the market. And so, when we're looking at our pipeline, and we are looking at the contracts that are in that pipeline that we are working on, they are extremely significant, they will move the needle.

Having said that, these contracts are heavily negotiated, and they do take a significant amount of time in order to move through the cycle to make sure that everything is done correctly, so that we can successfully deliver the product and quality that our clients require.

And so, we think that you're going to continue to see step functions in compute as these large clients take large blocks of compute over long periods of time from CoreWeave. They like our product, they like the way we deliver compute, they like the performance of the compute, and they will continue to buy from us as long as we can continue to build and deliver this infrastructure.

Michael Turrin

Analyst, Wells Fargo Securities LLC

Q

Thank you.

Operator: Our next question comes from the line of Gregg Moskowitz with Mizuho. Please go ahead.

Gregg Moskowitz

Analyst, Mizuho Securities USA LLC

Q

Great. And I'll add my congratulations. In the Q2, can you give us a sense of how successfully you were able to repurpose older GPU clusters that had come off contract? Any changes today vis-à-vis how this was trending around the start-of-the-year?

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Yeah. So, thank you for the question. It's a great question. So, what we are seeing is, we are seeing the infrastructure that is being delivered off of these contracts being re-contracted out for additional term in order to be able to continue to deliver that compute largely for inference.

And so, we're talking about the H100s, we're talking about the A100s, we're talking about delivery of this compute into contracts that are anywhere between one and three years in extension after the initial contract is over. And so, we're pretty excited about that. We've also seen things and this came up in the last call where the OpenAI contract was contracted out for five years with two additional one-year extensions, which also provides a significant amount of transparency into how people view the run-out of compute as it becomes an older generation.

Gregg Moskowitz

Analyst, Mizuho Securities USA LLC

Very helpful. Thank you.

Q

Operator: Our next question comes from the line of Tyler Radke with Citi. Please go ahead.

Tyler Radke

Analyst, Citigroup Global Markets, Inc.

Yeah. Thank you for taking the question. Two for me. So, one question just on timing. You talked about the big CapEx ramp in Q4. Obviously, the revenue guide also implies a pretty big step-up in Q4. Can you just help us understand the timing aspect there, particularly with CapEx a little bit lighter than we expected in Q2? Is this simply kind of delays related to Blackwell or is it specific to kind of contracts that you're expecting to ramp in Q4?

And then, the second question, just on the cost side, Nitin, you did highlight some increased costs. Obviously, you took up the revenue guide, but left operating income unchanged for the full year. So, could you just elaborate like what are the specific costs that you have to incur ahead of these contracts? Like, what is kind of coming in a bit higher than you expected? Just that would help us out in thinking through the mechanics. Thank you.

Q

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

Yes. So, I'll take the first one and then Nitin can follow up on the cost side. So, when we're looking at our build-out in ramp, it goes through a series of necessary steps, right? And so, the way that this is going to work is we are going to build from where we are right now, an additional 400 plus megawatts of power into our online and delivered compute and power. That is followed by the CapEx spend when the power is available, which is then followed by the revenue.

And so, we are very comfortable with the ramp that we are seeing in front of us in order to deliver the 900 megawatts plus of power as we go through Q4. It is going to be back-loaded, as Nitin said, we knew that it was going to be back-loaded as we came in and we're watching the build out and scaling of that infrastructure very systematically as we continue to move through the year.

A

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Thanks, Mike. And the other piece I would add to that is that we've been operationally preparing for this ramp up in executing and delivering this power by the end of the year. So, we are ready to kind of go through that exercise at this moment.

When we think about the costs in particular, we do incur costs especially associated with data center leases and expenses coming online, as we deploy this infrastructure and get it ready for our customers before we start generating revenue on that infrastructure.

A

That does create a timing mismatch, especially when you're adding capacity at the unprecedented scale we are adding, which is what you see reflected in our margin profile for the short duration as these customer contracts ramp up and the infrastructure associated with them is delivered to these customers.

Tyler Radke

Analyst, Citigroup Global Markets, Inc.

Thank you.

Q

Operator: And your next question comes from the line of Brad Sills with Bank of America. Please go ahead.

Brad Sills

Analyst, BofA Securities, Inc.

Wonderful. Thank you so much. And I'll echo the congratulations on a real solid Q2 here. I wanted to ask about the different segments of the market here. You think about that in the Big Three, the big enterprise AI companies or hyperscalers, if you will, and then you have the AI labs and then the enterprises themselves.

Q

As you've been embarking on this more product-led, software-led sale with Weights & Biases and the investment in Kubernetes, [ph] SUNK (00:53:22) for Kubernetes, for example. Are you seeing more of the pipeline, more of the business coming in weighted towards that next wave of AI lab and then eventually enterprises themselves? Any commentary on just the end segments? Thank you.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

Yeah. We're seeing incredibly broad-based demand for the compute and it is coming from the massive labs, right? Like that's clear. But what is less probably recognized in the market is that you're getting real green shoots in the sectorial growth of different parts of this market.

A

And like, we try to make a little bit of a reference in the initial statement, right, like we saw, within our VFX, we saw companies like Moonvalley, right? That's incredible, right? Like it's a new area of real growth from a new lab that's building products for a different part of the market and that's really exciting.

The financial players are really great to see. That's an enterprise a type of client, but it also represents a different use for the actual computing infrastructure, uncorrelated revenue source for us that we're really excited about. The big banks are starting to really show up and they are massive consumers of compute.

And so, the productization within the enterprise companies like Hippocratic AI, these are really representative of different parts of the economy starting to adopt AI, using it to deliver services to different parts of the economy. And we think that's tremendously exciting.

Keep in mind that you've got a [ph] scale prop (00:55:17), right? And that when you have a company like OpenAI or an entity like OpenAI consuming compute, they're just doing it at an order of magnitude that these other companies have not achieved yet. And so, we're excited to see the green shoots we think it's fantastic. We love that it is broadening the consumption of compute.

But we are also well aware that for the time being, these really large consumers of compute will dominate the client component of our pipeline. One of the things that we're incredibly excited about is how Weights & Biases

has like impacted our pipeline, right. Weights & Biases brings in 1,600 new clients. They bring in clients like British Telecom.

It's just – it's fantastic to see us forming relationships with these enterprise clients that are experimenting, are learning, are integrating AI into what they do. And it gives us an opportunity to position ourselves to become a supplier to them of the software and the hardware that they're going to require to be successful.

Brad Sills

Analyst, BofA Securities, Inc.

Q

Super exciting. Thank you, Mike. And then one more, if I may, please. I know cost of debt is a big focus for you and congratulations on the last two capital raises, bringing that down.

As equity investors were not in those conversations with creditors, would love to get your kind of sense from those conversations. What are the puts and takes that are driving that cost of debt down? Thank you so much.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

It's hard to overstate how excited I am about the progress that we have been able to make within the capital markets, within the debt markets. In a very fundamental way, what we have been doing is, we have been bringing to bear the largest part of the capital markets, the debt markets, to the problem of building and scaling infrastructure and it is an absolute necessity that that part of the market, those pools of capital, are able to come to bear, because of the size, scale and cost of what needs to be done to build compute at a planetary scale.

And so, it's just the incredible progress, and I'm incredibly proud of the team here that has been able to deliver the quality of infrastructure, that lenders can understand, get their arms around and underwrite. And it's been three deals that have really gone through this step function. There were two deals within the high yield space and then there was the DDTL, the Delayed Draw Term Loan, and that came in at SOFR plus 400 basis points.

That entirely changes the economics that are embedded in the contracts that we are delivering to the market. And it is a step function of massive importance. When Nitin is able to say, hey, we were able to drop our non-investment grade borrowing costs by 900 basis points, that's a seismic level shift in the cost of capital.

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

A

Great. Thank you. Operator, we have...

Brad Sills

Analyst, BofA Securities, Inc.

Q

Wonderful. Thank you, Mike.

Deborah Crawford

Vice President-Investor Relations, CoreWeave, Inc.

A

...sorry. We have time for one last question.

Operator: Thank you. And your final question comes from the line of Mike Cikos with Needham & Company. Please go ahead.

Mike Cikos

Analyst, Needham & Co. LLC

Q

Great. Thanks for taking the questions and squeezing me in here, guys. A bit of a two-parter. I just wanted to come back to the Weights & Biases. Great to hear on the 1,600 clients and how that's impacting the pipeline.

Imagine this goes hand-in-hand with some of the increased OpEx investment we're seeing. But can you talk about where we are on the sales front, as far as getting in front of those customers? You obviously had the fully connected developer event, but how is that tracking?

And then, the second piece I wanted to ask you about for on-demand and spot availability, again, I think, Mike, you had said it's a high quality problem for sure that as soon as it comes off contract, it's going back on. Would you need on demand and spot availability to increase as another avenue for new logo acquisition to tie to inference with these newer customers or not necessarily? Thank you.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

A

Yeah. So, look, the integration with Weights & Biases between the two companies has been fantastic. It's been one of the things that is so exciting to the legacy CoreWeave employees and the legacy Weights & Biases employees. I think that putting two organizations into a room that are so incredibly focused on the clients leads to incredible outcomes.

And there have been three real products that have been developed by the combination of the Weights & Bias (sic) [Weights & Biases] (01:00:17) and the CoreWeave teams, right? And so, we've integrated Weights & Bias (sic) [Weights & Biases] (01:00:22) into the Mission Control integration, which gives Weights & Bias (sic) [Weights & Biases] (01:00:27) clients historically the ability to now access the incredible observability in the Mission Control product to improve the performance of their use of the AI infrastructure.

We've built a Weights & Bias (sic) [Weights & Biases] (01:00:44) Inference product, which once again allows them an incredible amount of control over how they're using compute, what they're using compute, how it's impacting from really the data center all the way up and that level of transparency, clarity is just an incredible differentiator for the services that we provide that differentiate our compute from other providers and that's been a huge step function.

And then, the final product is actually the Weave product that has been pushed out by Weights & Biases and CoreWeave, which will allow folks to really make massive strides in optimizing all the way from their GPU through the code in their model to be able to drive performance. And so, that's super exciting for us, and we do think that it will lead to increased traction. We have always been an organization that has leaned on the concept of land and expand.

We get a client to come in and try our infrastructure. The incredible, the performance of our infrastructure leads to a deeper relationship, leads to larger contracts, leads to a broadening of how they use us to drive their business. And that's a tried and true.

When I go back to the Moonvalley and the VFX side of the house, it's the same concept when you acquired Conductor. It's how you go about introducing yourself to this new area of the market and build upon it and expand upon it and bring in new clients. And so, that acquisition is really moving along exactly how we had hoped it would.

With regards to your question around on-demand and spot, like I tried to say this earlier, I'll go about it a different way is that like, different use cases approach compute in different ways. They need different things. The profile of compute, the tooling that they're going to require, all those things are incredibly important.

And by having on demand, what you're doing is, you're creating a portion of our infrastructure that new players can use, so that they can build new product, they can open up new markets. And we really want to be able to continue to expand our footprint there because we think it's important.

We've been very successful with it in the early days of our growth, and we think that it's an important part of what we will need to do to continue to provide such incredible productization of our compute over time.

Mike Cikos

Analyst, Needham & Co. LLC

Q

Thank you very much.

Michael Intrator

Co-Founder, Chairman, President & Chief Executive Officer, CoreWeave, Inc.

All right. So, thank you for coming through the Q2 earnings call with us. We appreciate the questions and your interest. Our standout second quarter results reflected continued execution across every dimension of the business.

We're scaling rapidly to meet unprecedented demand for our purpose-built AI cloud platform, which continues to lead the industry in both performance and scale. We remain confident that 2025 will be a landmark year for CoreWeave. Our momentum is real, our strategy is working, and we are just getting started.

Thank you again for joining us today. I look forward to speaking to you at the next quarterly earnings call. Thank you.

Disclaimer

The information herein is based on sources we believe to be reliable but is not guaranteed by us and does not purport to be a complete or error-free statement or summary of the available data. As such, we do not warrant, endorse or guarantee the completeness, accuracy, integrity, or timeliness of the information. You must evaluate, and bear all risks associated with, the use of any information provided hereunder, including any reliance on the accuracy, completeness, safety or usefulness of such information. This information is not intended to be used as the primary basis of investment decisions. It should not be construed as advice designed to meet the particular investment needs of any investor. This report is published solely for information purposes, and is not to be construed as financial or other advice or as an offer to sell or the solicitation of an offer to buy any security in any state where such an offer or solicitation would be illegal. Any information expressed herein on this date is subject to change without notice. Any opinions or assertions contained in this information do not represent the opinions or beliefs of FactSet CallStreet, LLC. FactSet CallStreet, LLC, or one or more of its employees, including the writer of this report, may have a position in any of the securities discussed herein.

THE INFORMATION PROVIDED TO YOU HEREUNDER IS PROVIDED "AS IS," AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, FactSet CallStreet, LLC AND ITS LICENSORS, BUSINESS ASSOCIATES AND SUPPLIERS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE SAME, EXPRESS, IMPLIED AND STATUTORY, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NEITHER FACTSET CALLSTREET, LLC NOR ITS OFFICERS, MEMBERS, DIRECTORS, PARTNERS, AFFILIATES, BUSINESS ASSOCIATES, LICENSORS OR SUPPLIERS WILL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING WITHOUT LIMITATION DAMAGES FOR LOST PROFITS OR REVENUES, GOODWILL, WORK STOPPAGE, SECURITY BREACHES, VIRUSES, COMPUTER FAILURE OR MALFUNCTION, USE, DATA OR OTHER INTANGIBLE LOSSES OR COMMERCIAL DAMAGES, EVEN IF ANY OF SUCH PARTIES IS ADVISED OF THE POSSIBILITY OF SUCH LOSSES, ARISING UNDER OR IN CONNECTION WITH THE INFORMATION PROVIDED HEREIN OR ANY OTHER SUBJECT MATTER HEREOF.

The contents and appearance of this report are Copyrighted FactSet CallStreet, LLC 2025 CallStreet and FactSet CallStreet, LLC are trademarks and service marks of FactSet CallStreet, LLC. All other trademarks mentioned are trademarks of their respective companies. All rights reserved.