

10-Nov-2025

CoreWeave, Inc. (CRWV)

Q3 2025 Earnings Call

CORPORATE PARTICIPANTS

Michael Intrator

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

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

OTHER PARTICIPANTS

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Kash Rangan

Analyst, Goldman Sachs & Co. LLC

Tyler Radke

Analyst, Citigroup Global Markets, Inc.

Michael Turrin

Analyst, Wells Fargo Securities LLC

Brent Thill

Analyst, Jefferies LLC

Raimo Lenschow

Analyst, Barclays Capital, Inc.

Amit Daryanani

Analyst, Evercore Group LLC

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Brad Sills

Analyst, BofA Securities, Inc.

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 Third 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] Thank you.

It is now my pleasure to turn the call over to CoreWeave.

Unverified Participant

Thank you. Good afternoon and welcome to CoreWeave's third 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 to be 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.

Good afternoon, everyone, and thank you for joining us. CoreWeave once again delivered an exceptional quarter showcasing the accelerating momentum underlying our business as AI adoption proliferates globally across industries. We continue to operate in a highly supply-constrained environment where the demand for CoreWeave best-in-class AI cloud platform far exceeds available capacity. This insatiable customer demand is a clear signal that the world's leading companies trust CoreWeave to power their most critical AI workloads.

In Q3, we beat expectations, delivering revenue of \$1.4 billion, up 134% year-over-year. We added over \$25 billion in revenue backlog in the third quarter alone, bringing us to over \$55 billion in revenue backlog to end Q3, almost double Q2 and approaching [ph] 4 times (00:03:06) year to date. Further, CoreWeave has reached \$50 billion in RPO faster than any cloud in history. These results demonstrate the deep confidence customers have in CoreWeave, the company they trust as their essential cloud for artificial intelligence.

We continue to scale aggressively even as the industry remains capacity constrained. We expanded our active power footprint by 120 megawatts sequentially to approximately 590 megawatts, while growing our contracted power capacity over 600 megawatts to 2.9 gigawatts. This leaves us well-positioned for future growth, with more than 1 gigawatt of contracted capacity available to be sold to customers that we expect to largely come online within the next 12 to 24 months.

In Q3, we executed large-scale compute contracts with many of our largest customers, including Meta and OpenAI. Each represents a meaningful expansion of existing relationships and a diversification away from any single customer.

We also grew our relationship with a leading hyperscaler, marking the sixth contract with this customer to date. In fact, 9 of our 10 largest customers have now executed multiple agreements with us, the only exception being a new customer we onboarded in Q3.

CoreWeave is the force multiplier that empowers pioneers to accelerate breakthroughs in AI innovation. These are the world's most sophisticated AI organizations, and once they experience the performance, flexibility, and reliability of CoreWeave Cloud, they consistently expand with us. That is the strongest validation we could ask for.

Our exceptional growth illustrates just how quickly AI adoption is progressing beyond the frontier AI labs and hyperscalers. Broader global demand and our recent large wins are driving diversification of our revenue base. For example, the number of customers that exceeded \$100 million of revenue over the last 12 months tripled year-over-year. AI-native and enterprises across sectors are embracing CoreWeave to transform operations and unlock new sources of innovation, productivity, and growth.

At the forefront of foundation model development, poolside selected CoreWeave to power its mission to build artificial general intelligence and enable the deployment of agents across enterprises, while Periodic Labs is using

CoreWeave to push the boundaries of scientific discovery and computational research. At the application layer, we added AI-native customers like Jasper, who chose CoreWeave as their cloud partner as they transform the digital marketing landscape.

We are also seeing incredible momentum within enterprises. CrowdStrike chose CoreWeave to advance the development of AI agents for cybersecurity, while Rakuten is using our platform to transform their visual language models, helping to achieve greater transparency, reproducibility, and speed in their AI workloads. We also saw further expansion with a wide range of enterprise customers, including a leading software design platform and a large telco operator in the US.

Our reach now extends into the public sector, a market with unique performance and security requirements. We recently launched CoreWeave Federal to bring our cloud services to the US government agencies and the defense industrial base. Already, NASA is leveraging our services to advance scientific exploration at its Jet Propulsion Lab. We are honored to help strengthen America's AI infrastructure, enabling agencies to accelerate innovation and address critical missions and our national interests.

These recent wins underscore that we are enterprise-ready. With our customer base broadening across verticals and geographies, we are excited to welcome Jon Jones as our first Chief Revenue Officer. Jon joins us from AWS, where he served as Global Head of Startups and Venture Capital. Jon is a strong addition to our team and will play an important role scaling our global revenue organization and driving expansion through this next phase of growth.

Next, as I move to discuss our growing data center footprint, I want to briefly touch on our previously proposed acquisition of Core Scientific, which was terminated in October. While the deal made sense strategically for both companies, the valuation required by their shareholders was simply not a price that was appropriate for CoreWeave, particularly because the outcome of the transaction in no way adversely impacts our ability to achieve our growth ambitions in the coming years. Instead, we will continue to work closely with Core Scientific on the approximately 590 megawatts of capacity we have already leased.

Our disciplined approach to expanding our capacity footprint ensures we are meeting the surging global demand for CoreWeave's Cloud Services. As I mentioned, we grew our contracted power capacity to 2.9 gigawatts this quarter as we diversified across size, geography, and developers enhancing resilience and flexibility across our portfolio. As of Q3, no single data center provider represents more than approximately 20% of our contracted power portfolio.

In the past quarter, we added eight new data centers across the US, strengthening our domestic coverage, with additional expansions underway across Europe, including a major new presence in Scotland, which is being developed in partnership with the UK government. And as we announced over the course of the summer, we have embarked on self-build projects to further accelerate our footprint and provide us greater operational control.

While we are experiencing relentless demand for our platform, data center developers across the industry are also enduring unprecedented pressure across supply chains. In our case, we are affected by temporary delays related to a third-party data center developer who is behind schedule. This impacts fourth quarter expectations, which Nitin will discuss shortly.

Having said that, the customer affected by the current delays has agreed to adjust the delivery schedule and extend the expiration date. As a result, we maintain the total value of the original contract and the customer preserves their capacity for the full duration of the initial agreement, demonstrating the confidence they have in

our ability to provide the most performance solutions in market. We are incredibly proud of our technical accomplishments and our customers continue to tell us that CoreWeave is the absolute best place to run AI workloads.

In the third quarter, we continued to deliver many of the initial scaled deployments of the GB200s, while once again being first to market, this time with the GB300s, further highlighting our incredible track record of operational excellence.

CoreWeave's industry leadership is unmatched. We are the only cloud provider to submit MLPerf inference results for GB300s, setting the benchmark for real-world AI performance. And just last week, SemiAnalysis once again recognized our dominance, awarding CoreWeave its highest possible distinction, its Platinum ClusterMAX ranking, for the second time ahead of more than 200 providers, including the hyperscalers and emerging neoclouds. No other cloud has achieved this once. CoreWeave has done it twice, underscoring yet again that CoreWeave stands alone at the forefront of the AI cloud.

Demand for AI cloud technology remains robust across generations of GPUs. For example, in Q3, we saw our first 10,000-plus H100 contract approaching expiration. Two quarters in advance, the customer proactively re-contracted for the infrastructure at a price within 5% of the original agreement. This is a powerful indicator of customer satisfaction as well as the long-term utility and differentiated value of the GPUs run on CoreWeave's platform.

CoreWeave is the world's first AI cloud at hyperscale, comprising compute, storage, networking, and software purpose-built for AI workloads. Our growing cloud portfolio is underpinned by an expanding suite of software and services that help our customers build, train, and deploy new products faster.

In addition to Mission Control, our proprietary orchestration solution which is critical to autonomously operate our AI cloud at the bleeding-edge, we recently launched CoreWeave AI Object Storage, a fully managed storage service that eliminates any friction of moving data between regions, clouds, and tiers with zero egress or transaction fees.

CoreWeave's AI Object Storage delivers the highest amount of throughput of AI workloads, while cutting the customer's cost by more than 75%. We have already seen tremendous interest in this offering, adding a number of initial customers, including frontier AI labs like Mistral.

Across our entire storage platform, we have seen rapid customer adoption eclipsing \$100 million in ARR in Q3. Combined with our unique global network backbone purpose-built for AI, this positions CoreWeave as the hub for customers and their key AI workloads, enabling consistent best-in-class performance and seamless user experiences when utilizing CoreWeave Cloud or a secondary provider. [ph] We've (00:14:57) supplemented these capabilities with further expansion of our observability and security suites to ensure that CoreWeave is best positioned to handle all of our customers' critical workloads regardless of the use case or geography.

Our role over the last few years has been to support the pioneers who are developing and improving AI. Now, we are expanding our role to help put AI to work, from the tools that developers require to build AI to the solutions that the physical world requires to adopt AI. We've used M&A as a key tool to accelerate this journey, including the recently announced acquisitions of OpenPipe, Marimo, and Monolith.

With OpenPipe, we quickly integrated their solutions into our broader fine-tuning product suite and introduced the first publicly available serverless reinforcement learning tool. With Marimo, we are expanding CoreWeave's

exposure to and impact within the open-source community, starting with entry-level exploration and prototype. Both OpenPipe and Marimo fit seamlessly with the capabilities of Weights & Biases where we are rapidly growing the developer base reliant on CoreWeave's holistic platform. With Monolith, we are expanding these capabilities into the physical world to unlock the monetization of AI today, initially focusing on industrial use cases with an established enterprise customer base and mature workloads, including leading auto OEMs like Nissan and Stellantis.

Through the rapid and successful launch of new products and services, we are expanding our addressable market and growing with our customers. We are fundamentally evolving the capabilities of CoreWeave, which is creating beachheads and expansion opportunities into new markets, all in the service of further supporting the rapid growth of AI and enabling AI builders and innovators to get to market faster and more reliably, and drive ROI.

Our engagements are getting more sophisticated, as evidenced by our partnership with CrowdStrike, which will unlock and accelerate partner-driven growth. Our new storage product and partnership with VAST Data is another example of accelerating both our product portfolio and partner go-to-market motions, and allows us to compete in new markets where we previously had limited or no offerings. This facilitates customer-driven platform adoption and product-led growth, creating tailwinds for our business.

As I close, I want to emphasize what truly sets CoreWeave apart. We are the essential cloud for AI, combining unmatched technical and operational excellence with a rapidly diversifying customer base. We deliver the most performant infrastructure, the fastest time to market, and the most advanced capabilities in the industry. The world's leading AI innovators choose CoreWeave because we enable them to move faster, scale smarter, and achieve outcomes that simply are not possible anywhere else.

Our momentum has never been stronger and the opportunities ahead continue to expand. Powered by exceptional products, an extraordinary team and unrivaled execution, CoreWeave is ready to enter the next phase of growth as a full-stack AI service provider and hyperscale. The future runs on CoreWeave and we are just getting started.

With that, here's Nitin.

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Thanks, Mike, and good afternoon, everyone. Our impressive third quarter results reinforced the relentless demand for CoreWeave and our focused execution in building the essential cloud for AI. As Mike shared, we continue to execute within a highly supply-constrained environment which we expect to persist for an extended period of time. Our continued focus on delivering the most performant solution in the market and investing up and down the stack is spurring growth and diversification across our customer base from new enterprises and AI natives to expansion with existing customers.

Now turning to Q3 results, Q3 revenue was \$1.4 billion, up 134% year-over-year, driven by robust customer demand and strong execution. Revenue backlog for the quarter ended at \$55.6 billion, almost doubling in the third quarter alone. Demand remains robust for not just the Blackwell platform, but across our GPU portfolio. In the third quarter, we signed a number of deals for older generations of GPUs, adding new customers and recontracting existing capacity.

The breadth of demand for CoreWeave's cloud services has enabled us to reduce our customer concentration significantly. Today, no single customer represents more than approximately 35% of our revenue backlog, down from approximately 50% last quarter and even more meaningfully from approximately 85% to begin the year. Additionally, as of Q3, more than 60% of our revenue backlog is tied to investment-grade customers. This is what successful execution against our stated goal of platform and customer diversification looks like.

Operating expenses in the third quarter were \$1.3 billion, including stock-based compensation expense of \$144 million. We continue to ramp our investments in data center and server infrastructure to execute against our growing revenue backlog, which contributed to the increase in our cost of revenue and technology and infrastructure spend in Q3. In addition, the increase in sales and marketing was driven by investments in marketing and scaling our go-to-market organization to capture the rapid growth of AI opportunities across enterprises and AI natives. The increase in G&A was driven by professional services and head count.

Adjusted operating income for Q3 was \$217 million compared to \$125 million in Q3 of 2024. Our Q3 adjusted operating margin was 16%. Adjusted operating income was better than expected due to higher revenue, lower costs due to timing of data center deliveries from our third-party partners, and improved fleet efficiencies.

Net loss for the third quarter was \$110 million compared to a \$360 million net loss in Q3 of 2024. Interest expense for Q3 was \$311 million compared to \$104 million in Q3 of 2024 due to increased debt to support the scaling of our infrastructure, partly offset by the benefit from better interest rates on our debt as we make further progress in lowering our cost of capital.

Adjusted net loss for Q3 was \$41 million compared to approximately breakeven in Q3 of 2024, while adjusted EBITDA for Q3 was \$838 million compared to \$379 million in Q3 of 2024, increasing more than 2x year-over-year. Our adjusted EBITDA margin was 61%.

Turning to capital expenditures, CapEx in Q3 totaled \$1.9 billion, lower than anticipated due to the delays Mike mentioned related to deliveries from a third-party data center provider. The meaningful growth in construction in progress to \$6.9 billion, an increase of \$2.8 billion quarter-over-quarter, is a direct result. As a reminder, construction in progress represents infrastructure not yet in service and is excluded from CapEx until it is deployed.

Now let's turn to our balance sheet and strong liquidity position. As of September 30, we had \$3 billion in cash, cash equivalents, restricted cash and marketable securities. Growing rapidly and operating at scale demands a strategic approach to securing capital. CoreWeave has established itself as the leading AI cloud and the leading innovator in financing the infrastructure required to power the world's most advanced workloads for enterprises and AI labs.

We continue to make great progress in strengthening our capital structure and lowering our cost of capital. In Q3, we amended the DDTL 2.0 Facility by increasing its remaining drawable capacity by over \$400 million to create a new \$3 billion tranche at SOFR plus 4.25%, which is significantly below the original cost of the facility. As we discussed previously, we also closed DDTL 3.0 in the third quarter priced at SOFR plus 400 basis points, which represents a 900-basis-point decrease from the non-investment-grade portion of our prior facility.

Going forward, we expect to continue to be able to finance at lower spreads as our capital providers increasingly appreciate our best-in-class execution as well as the durable cash flow and visibility that underpin our take-or-pay customer contracts. Further, we raised \$1.75 billion in senior notes in July, extending our exposure to the high yield market at a cost 25 basis points lower than our inaugural offering in May. Year to date, CoreWeave has

successfully secured \$14 billion in debt and equity transactions to support our execution on our rapidly growing backlog and efficient scaling for long-term growth. Other than payments related to OEM vendor financing and self-amortizing debt through committed contract payments, we have no debt maturities until 2028.

Turning to tax, in Q3, we recorded a non-cash tax benefit, primarily due to the impact of One Big Beautiful Bill. While the size of the impact to Q3 was one time in nature due to a year-to-date catch-up, we expect the change in law to enable CoreWeave to realize cash tax savings in future periods.

Now turning to guidance, as mentioned, the delays in powered shell delivery associated with the data center provider will have an impact on our fourth quarter results. These delays are temporary, and as Mike noted, the affected customer has agreed to adjust the delivery schedule to preserve their capacity for the full duration and the total value of the original agreement. With that backdrop, we now expect 2025 revenue in the range of \$5.05 billion to \$5.15 billion. In addition, we anticipate 2025 adjusted operating income between \$690 million to \$720 million and expect to end the year with over 850 megawatts of active power.

In Q4, we will be bringing online some of the largest-scale deployments in our company's history. This will have a near-term impact on adjusted operating margin due to the timing difference between when data center costs are first incurred and when we start recognizing revenue. We expect 2025 interest expense in the range of \$1.21 billion to \$1.25 billion, driven by increased debt to support our demand-led CapEx growth, partly offset by an increasingly lower cost of capital.

Moving to CapEx, we now expect 2025 CapEx in the range of \$12 billion to \$14 billion. We expect this reduction in CapEx from our prior guidance will be mostly reflected by a corresponding increase in construction in progress due to the buildup of infrastructure waiting to be deployed following the delivery of powered shell capacity. As such, the vast majority of the remaining CapEx we have previously anticipated to land in Q4 will now be recognized in Q1.

In addition, given the significant growth in our backlog and continued insatiable demand for our cloud services, we expect CapEx in 2026 to be well in excess of double that of 2025. These investments in our infrastructure platform will strengthen our competitive moats and support our continued hypergrowth.

In closing, we delivered a record third quarter and remain more confident than ever in the long-term trajectory of our business. Over the course of this year, we've made tremendous progress, accelerating our revenue backlog growth that now exceeds \$55 billion while diversifying our customer base, executing strategic partnerships and acquisitions to strengthen and broaden our platform, accessing new capital pools that meaningfully reduce our cost of capital, and scaling both our capacity and organization at an unprecedented pace. This progress enables us to seize the opportunities in front of us today and create a strong foundation for years to come.

Our addressable market continues to expand not only as AI adoption proliferate across industries and use cases, but also through deliberate business decisions we've made to broaden our product portfolio and capture greater wallet share across the industry. CoreWeave is reaching escape velocity, scaling more rapidly and efficiently and solidifying our leadership as the essential cloud for AI.

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 move to Q&A.

QUESTION AND ANSWER SECTION

Operator: [Operator Instructions] We respectfully request that you limit questions to one. And our first question comes from the line of Mark Murphy with JPMorgan. Please go ahead.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

Thank you, Michael. Yeah, very discussion we have across the AI landscape, we hear that bookings are booming and, obviously, that applies to CoreWeave, but the bottlenecks around power and manpower are just becoming so severe. Can you speak to that situation relating to the third-party provider? Specifically, is it a shortage of power or manpower? Is it something outside of that with GPUs or memory or storage? And then have you spoken to your other third-party providers to get a sense of their own trending relative to schedule and whether they think they can hold on their – or deliver on their commitments into early next year?

Michael Intrator

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

A

Let me kind of take that question apart a few different ways, right? So, first of all, you're correct, it is very frustrating for our clients, it's very frustrating for us because of the kind of systemic challenges that exist within the supply chains that are necessary to deliver a – the global infrastructure that's required for artificial intelligence. Having said that, we have taken a number of steps along the way here to really drive home our ability to manage that environment, which is going to be challenging, into the future.

We've really spent a lot of time diversifying our data center providers. We have created a significant portion of the company dedicated to being able to facilitate and assist with the operational component of delivering infrastructure. We've set up our own self-build efforts, including Kenilworth and Lancaster, Pennsylvania. So you see us kind of really spreading out and ensuring that we're doing everything that is possible to limit the damage associated with or the delays associated with delivering this infrastructure, which is just overwhelming the supply chains.

Now, when you have a diversified portfolio of paths to infrastructure, the relative impact of each delay becomes small. You'd just be able to draw on different data centers as you're getting it delivered. And so we really look at this as this is a significant block of infrastructure that's come on late, but the fact that the ultimate end customer that's going to be consuming this infrastructure has shifted the contract back to allow us to be able to deliver the full contract value in spite of the delays really speaks to the value that the customers get out of our infrastructure.

So, you're going to be hearing this theme repeated again and again as you talk to not just CoreWeave but you talk across the space. And it is a real challenge at the powered shell level, it's not a challenge for power, right? There's plenty of power right now and we believe that there will be ample power for the next couple of years, but really where the challenge is, is the powered shell.

Mark R. Murphy

Analyst, JPMorgan Securities LLC

Q

And so, Michael, does this not relate to Core Scientific in any way or is this totally removed from that situation that you've gone through?

Michael Intrator

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

A

So I'm not going to speak to any specific one of our data center providers. We're working with all of our data center providers to do everything we can to facilitate the ultimate delivery of the infrastructure that they're going to deliver to us. We've had some incredible success getting infrastructure delivered to us as you can continue to see us scaling. You saw us hit 500 – approximately 590 megawatts, we're up 120 megawatts since the last call, so you are seeing a significant amount of success as we continue to scale delivery.

But I don't think it really matters who the individual data center provider is. This is a systemic problem that the industry is going to have to deal with for the foreseeable future. The important part here is that – or the important part from my seat is that the infrastructure which is undergoing a delay is not going to impact our backlog and our ability to extract the full value from the contracts that we're going to deliver on.

Operator: Our next question is from the line of Keith Weiss with Morgan Stanley. Please go ahead.

Keith Weiss

Analyst, Morgan Stanley & Co. LLC

Q

Excellent. Thank you, guys, for taking the question and congratulations on another super impressive quarter in terms of building out that backlog. You're right, just we've never seen this in terms of any cloud provider being able to build out that quickly. Mike, I wanted to ask you a question that's been asked of us a lot that we're hearing a lot on CNBC, and it's really about sort of the risk of overcapacity. But I think it's more narrow than that, that people are worried about overcapacity from – or of what's being contracted by AI labs out there.

The question I want to ask you though is how we should think about your guys' infrastructure and the infrastructure that you build and how fungible that infrastructure really is. When you're building out for a particular customer, those data centers, is that usable for any customer, is it usable for inference and training, or do you really build to suit to a certain customer that would lock you in and give you kind of less degrees of freedom, if you will, if one customer is doing better or worse?

Michael Intrator

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

A

Yeah, Keith, that's an excellent question. It's actually something that we've spent a lot of time thinking about here as we kind of proceed with our relationships with all our customers. And so in short, the infrastructure is fungible. It would be able to be transferred from one client to another. The infrastructure is built to the most demanding specs so it's able to be used for training, it's able to be used for inference. We really have thought a lot about making sure that we maintain as much optionality, as much flexibility within our infrastructure build as possible.

And I want to highlight for everyone that a lot of that flexibility, a lot of that fungibility really does tie back to the incredible software suite that we provide that allows for such effective use of the infrastructure, right? Like when SemiAnalysis did their annual kind of review of the alternatives out there, there's a reason that CoreWeave has come back time and time again as singular as the best solution for this type of infrastructure that exists in the world, and that includes the hyperscalers, the neoclouds, and everyone else that's trying to deliver this infrastructure.

We just do a great job and we believe that there's a lot of value that we are protecting by providing such a robust software suite to be able to deliver infrastructure.

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

Kash Rangan

Analyst, Goldman Sachs & Co. LLC

Q

Hi. Thank you very much and impressive backlog growth. Two things that I wanted to just touch upon. One is, Mike, I think you've talked about how you're going to be diversifying your contractors in the data center side. Maybe you could give us a honest to goodness update on how far are we away from potentially reaching a point where any disruptions that have nothing to do with your business should not affect your revenue outlook, how far away are we from that point.

And secondly, when you look at the developments, I mean, nobody expected a – maybe some did, but a \$250 billion contract for OpenAI with Microsoft, nobody expected a \$300 billion contract OpenAI with Oracle. All of sudden, the – certainly CoreWeave has got a unique value proposition, being able to stand up GPU clusters very quickly, very effectively at the speed of though almost. But in a landscape where we're talking hundreds of millions of dollars being awarded to the hyperscaler giants, what gives you the uniqueness three to four years from now when things have sort of settled into a supply equals demand, or what – when we look back at CoreWeave, what will be the shining value proposition that keeps you in the game at that point? Thank you so much and that's it from me.

Michael Intrator

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

A

Yeah, thank you. Let me break that question into two pieces, right? The first question you asked is about diversification and when does it start – when does it stop, kind of causing dislocation in in our numbers as we're delivering them quarter-to-quarter. And what I would like to focus you on here is as the individual builds become smaller relative to the size of the entire portfolio of data centers that we are running, the impact of being a couple of weeks late will become less and less meaningful in the general accounting of what's going on, right?

So when you're delivering 590 megawatts of power and you have a step function of 200 megawatts or 300 megawatts, it's a material percentage that's going to be delivered over the next quarter, right? As we become larger and larger and start to build out the full 2.9 gigawatts of power that we have, having a data center that's 100 megawatts delayed a week or two is not going to have a material impact.

And as Nitin said, we expect the overwhelming majority of that 2.9 gigawatts of power to be brought into service over the next 12 to 24 months. And so that'll give you a really good idea of how the curve begins to become more smooth as we get larger and the relative impact of each data center becomes smaller. And so that's the first part on the scaling side.

The second part is the question you're asking has been asked of us since we started this business. Why is CoreWeave going to be able to deliver GPUs faster? Why are we going to be able to deliver the GPUs that NVIDIA uses to run its MLPerf? Why are we going to be able to create software that is going to define the space? And with each quarter, you see us extending the lead with which we have because of the customization of our cloud to the use case that is required. And once again, you saw us in the in the SemiAnalysis, like we're singular in this, we are out there building our product offering, we're building or buying additional capacity to further decommoitize the compute that we're delivering. And so a company that's built singularly to deliver this type of compute will be effective on a go-forward basis.

Operator: Our next question is from the line of Amit Daryanani. Please go ahead.

Michael Intrator

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

Amit, are you on mute?

A

Operator: Amit, your line is open.

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Operator, let's go to the next question and we'll come back to Amit. [indiscernible] (00:44:09). Operator, can we go to the next question and we'll come back to Amit?

A

Operator: Our next question is from Tyler...

[indiscernible] (00:44:23)

Operator: ...Radke with Citi. Please go ahead.

Tyler Radke

Analyst, Citigroup Global Markets, Inc.

Hey, hopefully, you can hear me okay. Thanks for taking the question. So double-clicking on some of the delays that you called out in the quarter, can you just help us understand the implications on 2026? I know, Nitin, you provided some high level commentary on CapEx but, I mean, just given the visibility you have, particularly on the 24-month component of RPO, how should we be thinking about sort of the revenue implications of the shift? Is this a delay that you think kind of gets fully resolved into Q1? And should we see sort of a step-up in growth rate next year relative to this year? Just any color on that would be helpful.

Q

Michael Intrator

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

Yeah, so I'll start then I'll hand it over to Nitin. I think it's important to understand that the ramp that we are seeing is at a – is associated with the infrastructure from a single provider and we are parallel pathing with other providers for other contracts. And so you're going to see a short-term impact associated with this delivery, and then what you're going to see is the – our ability to accelerate through the year back to schedule. So the overwhelming majority of the delay that you're seeing should be taken care of within Q1 of next year. Nitin?

A

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Yeah, Tyler, that is correct. The vast majority of the CapEx pushout that we experience in Q4 will be done in Q1. And as you can imagine, we're going to ramp the capacity through the course of Q1 for this. As Mike earlier mentioned, the impact on the total revenue associated with the customer is not impacted here because we've been able to adjust the delivery dates associated with the customer so that the customer keeps the full capacity as well as the contract value associated with it.

A

We will share more details around the 2026 build and our revenue plan associated in the next earnings, but as we highlighted in this quarter, given the strong customer demand that you see, even – that is demonstrated in our

revenue backlog growth as well as the continued customer demand we see, we expect 2026 CapEx to be well more than double of that of 2025.

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

Michael Turrin

Analyst, Wells Fargo Securities LLC

Q

Hey, thanks very much. I appreciate you taking the question. I want to just try to tie some of the commentary together, because the bookings growth clearly stands out and there are a lot of questions just around the sequencing. And so, it sounds like what you're saying is the supply chain impacts you're seeing are more single customer specific. What I'm trying to get a better sense of is does this at all impact the cadence at which you're able to sign on new customers or is this more tied to post-ramp signing and one more specific customer environment?

And just as a small follow-up, does the NVIDIA deal specifically show up in the backlog metric? It might be useful to hear you expand on what that deal opens up, given it's a bit of a different structure there as well. Thanks very much.

Michael Intrator

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

A

Sure. So, the – there is no impact on our ability to bring on more clients. I think it's important to understand that we're parallelizing the build of infrastructure. There was a problem at one data center that's impacting us, but there are 32 (sic) [41] (00:48:23) data centers in our portfolio, all of them are progressing to one extent or another. And so, that is – each one of those is independent. And as Nitin spoke earlier, we have 2.9 gigawatts' worth of contracted power that will come on in the next 12 to 24 months. And we are going to be looking to fill that with clients that are going to be using that, which will have a substantial impact on our revenue on a go-forward basis. This one data center will catch up and then we will move forward from there.

Want to talk about the NVIDIA deal?

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

A

Yeah. Michael, on the NVIDIA deal perspective, we're really excited about this deal. This contract allows for the capacity contracted and reserved for NVIDIA to be interrupted and resold to different customers. So, the nature of this contract allows us to offer our services profitably to a wide range of smaller customers, such as high-growth AI labs that prefer shorter and lower upfront commitments, while eliminating any utilization risks for capacity from our site. So, we're really excited about this.

Given the flexibility in the contract to interrupt and to resell capacity accounting rules require us that we exclude the amount we expect to be resold to other customers from RPO. To be clear, if not resold, this capacity will remain committed to NVIDIA and will be recognized as revenue. So, you'll see this NVIDIA contract in our revenue backlog, but not in our RPO to a large extent.

Michael Intrator

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

A

So, just to follow up with that for a moment there, we – as Nitin said, we're extremely excited about this because what this contract is going to allow us to do is to provide infrastructure to emerging companies, startups,

companies that are struggling to get access to the computing infrastructure that they require to be able to build their business. And so, the interruptibility here is an incredibly powerful tool for the resiliency and opportunities for new companies to become part of CoreWeave's broader offering. And we're really excited about this. We think it's a great structure. It is a deal with NVIDIA. They fully underwrite the economics because we will sell the compute to them.

And I want to be clear that this really does represent an incredibly disciplined way of financing the compute in order to be able to reach parts of the market that we have been unable to reach, or anyone for that matter, has been unable to reach up to this point.

Operator: Our next question comes from the line of Brent Thill with Jefferies. Please go ahead.

Brent Thill

Analyst, Jefferies LLC

Nitin, I just wanted to be clear, you cut CapEx by 40% for the year. And just to be clear, this is from one customer, correct? This is – you're not assuming other delays across the board, correct?

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

That is correct. So, this is associated with a single provider, data center provider partner and the delays associated with that. And as we talked in our prepared remarks, most of it, [ph] a large (00:51:53) majority of it is going to be recognized in Q1. And in Q4, you're going to see a major impact on build-up of construction in progress associated with the build-up related to it.

Brent Thill

Analyst, Jefferies LLC

Okay. Terrific.

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

Raimo Lenschow

Analyst, Barclays Capital, Inc.

Perfect. Thank you. As we think about the CapEx next year, Nitin and Mike, can you speak as well about the sources of funding a little bit because what we've seen from a lot of the other players is that leasing is coming up a lot more. You talked about CapEx, which is kind of what you need to do. How do you think about that path for you going forward between the different ways of kind of funding the business, which might give you even more flexibility? Thank you.

Michael Intrator

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

Thank you. So, look, we've driven innovation on the technology side and we've driven innovation on the financing side, right? The way that I look at this is that we will look at the full suite of potential ways of financing and expanding our footprint. And then we will choose whatever is the most cost-effective way of increasing our scale and serving our clients.

And so, if leasing is the path, that's the path we'll go. But we've seen a lot of different structures, we've created a lot of different structures that have given us access to capital over the past three years. And we believe that we're going to explore the full suite of those as we look forward. We don't sign customers without knowing where the financing is going to come from. We go deal by deal and we make sure that we have the physical data center spoken for, we have the power spoken for, we have the GPU spoken for, and we have the financing spoken for in order to ensure that we are able to successfully deliver compute to them.

Operator: Our next question comes from the line of Amit Daryanani with Evercore. Please go ahead.

Amit Daryanani

Analyst, Evercore Group LLC

Yeah, hopefully this works better.

Q

Michael Intrator

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

We got you.

A

Amit Daryanani

Analyst, Evercore Group LLC

All right. Perfect. Mike, I was hoping if you could just talk about as you shift from third-party data center providers to perhaps do more of your own self-build, how does that impact your CapEx and time to market for power as you go forward? Would love to just understand how do you think that optimal mix looks like and what the CapEx requirements could be as you perhaps go more towards self-build versus third-party data center providers. Thank you.

Q

Michael Intrator

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

Yeah. So, I want to be clear, we're not saying that we're going to go self-build and not use third-party data center providers. What we are saying is that self-build is a component of the way that you go about de-risking delivery across the broader portfolio. And so, we're going to go ahead and we're going to continue to work with our partners who provide data center capacity that allow us to co-locate at their facilities, that builds facilities for us. All of that is going to continue to be true.

A

We need that capacity in order to be able to continue to move and operate at the speed and scale that we are. We just look at self-build as an additional piece of the puzzle. It puts us closer to the physical infrastructure. It embeds us deeper into the supply chain around the world so that we have firsthand information. We just think that you need to be on both sides of this fence in order to be as effective as you can be de-risking what is a complicated supply chain environment.

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

Brad Zelnick

Analyst, Deutsche Bank Securities, Inc.

Great. Thank you so much for taking the question. Mike, with 2.9 gigawatts in committed power and over 1 gigawatt yet to be contracted out to customers, meanwhile, we continue to see a number of other large deals get

Q

announced industry-wide. How do you think about and how might you frame for us the pacing on contracting out the remaining capacity given the demand is insatiable out there?

Michael Intrator

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

A

Yeah. So, look, thanks for the question here. The fact that there are other deals getting contracted out there is incredible validation for the supply and demand environment that we have been describing for years now, right? There is no entity that has the capacity to be able to deliver infrastructure globally in order to meet the demand that's being driven by the largest technology companies in the world, by the largest AI labs in the world, by government, by enterprise. All of these things are coming to bear. And so, your – the fact that there are other deals going to other players is part and parcel for the fact that we, like the hyperscalers, like the AI labs, like the data centers are being overwhelmed by demand. It is just reinforcing and validating the theme that we've been talking about.

We think that at the end of the day, the product that we deliver, right, which is a full stack, everything from the hardware all the way through the software, is the most valuable representation of this infrastructure that can be delivered to the market. And we continue to think that that will drive a significant amount of demand for our infrastructure. As far as the remaining capacity goes, we're being very thoughtful about continuing to drive diversification across our cloud. We're continuing to think about different applications that are going to be meaningful contributors to the way the world will work in the future. And we are allocating that infrastructure to those parties as quickly as we can in order to ensure that they are successfully able to launch their products, their enterprises.

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

A

And, Brad, couple of things to kind of keep in mind here as we kind of talked about in our prepared remarks. Today, approximately – no customer represents greater than approximately 35% of our revenue backlog, which is meaningfully down from where we began the year at 85%, and 60% of our revenue backlog is with investment-grade customers. So, vectors that we are very thoughtful around as we take care of the capacity that we have available to be sold.

Operator: We have time for one final question. Our final question comes from the line of Brad Sills with Bank of America. Please go ahead.

Brad Sills

Analyst, BofA Securities, Inc.

Q

Oh, great. Thank you so much. I did want to ask a question around this concept of the [ph] power shell as (00:59:19) the bottleneck here, Mike. Is there any IP that CoreWeave has that you contribute to the build-out of these data centers? Any learnings from this delay that you might be able to apply to other contracts? I'm just trying to get a sense for how much is in your control here to kind of solve for this bottleneck issue that you're experiencing in this one contract itself. Thank you.

Michael Intrator

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

A

Yeah. What I would say is, Brad, I don't think that I would say that our learning has come from this one delay. We've been operating in a systemically supply-constrained market globally now for three years. We understand how difficult it is, and with each additional wave of demand, the market gets tighter and tighter.

So, when you ask what are we doing to position ourselves on a go-forward basis, what I would really encourage you to think about is the fact that we've built out an entire organization within CoreWeave that is capable of helping us build and deliver additional capacity on the self-build side. That's where you embed yourself into the supply chains, you understand where the power is, how it's being contracted. You understand what it takes to build the powered shells because you're doing it yourself, in addition to the fact that you're using other third-party providers. Those are the type of relationships that will enable us to be as successful as possible in what is going to be a challenging environment for quite a while.

Operator: Thank you. And that concludes our question-and-answer session for today. I would now like to turn the conference over to Michael Intrator for closing remarks.

Michael Intrator

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

Thank you all for joining us today. As we wrap up, I want to emphasize how proud we are of the strong foundation we've built this year and the incredible momentum driving our business forward. Our team's exceptional execution to build the essential AI cloud has positioned CoreWeave to capture a significant and expanding market opportunity. We appreciate your support and engagement and we look forward to updating you on progress next quarter. Thank you. Have a good night.

Operator: This does conclude today's conference call. You may now disconnect.

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.