

26-Feb-2026

# CoreWeave, Inc. (CRWV)

Q4 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

### Josh Baer

*Analyst, Morgan Stanley & Co. LLC*

### Amit Daryanani

*Analyst, Evercore ISI*

### Mark R. Murphy

*Analyst, JPMorgan Securities LLC*

### Brent Thill

*Analyst, Jefferies LLC*

### Gabriela Borges

*Analyst, Goldman Sachs & Co. LLC*

### Ben Reitzes

*Analyst, Melius Research LLC*

### Brad Zelnick

*Analyst, Deutsche Bank Securities, Inc.*

### Michael Turrin

*Analyst, Wells Fargo Securities LLC*

---

## MANAGEMENT DISCUSSION SECTION

**Operator:** Hello, and thank you for standing by. My name is Tiffany and I will be your conference operator today. At this time, I would like to welcome everyone to the CoreWeave Fourth Quarter and Fiscal Year 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]

I would now like to turn the call over to CoreWeave. Please go ahead.

---

### Unverified Participant

Thank you. Good afternoon, and welcome to CoreWeave's fourth quarter and fiscal year 2025 earnings conference call. Joining the call 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 Annual Report on Form 10-K 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](https://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. 2025 was a defining year for CoreWeave. We generated more than \$5.1 billion of revenue, up 168% year-over-year; grew our contracted revenue backlog to \$66.8 billion, an increase of \$11.2 billion sequentially and more than \$50 billion year-over-year; reached more than 850 megawatts of active power as of December 31; and added approximately twice as many new reserved instance customers in Q4 versus any quarter in our history. We delivered these results while quickly resolving the data center delays we discussed last quarter, delivering the impacted deployments ahead of our Q3 earnings call expectations.

CoreWeave is the fastest cloud in history to reach \$5 billion in annual revenue. We remain in the early stages of the most transformative infrastructure build-out in history, and CoreWeave is at the forefront, building and operating some of the largest purpose-built AI clusters for the world's most demanding workloads. Strip away the complexity, and four fundamentals define where we stand: one, a demand environment that remains relentless, driving rapid adoption from an increasingly diversified set of hyperscalers, AI-native and enterprise customers; two, expanding opportunities for new margin-accretive avenues to monetize CoreWeave Cloud, unlocked by the evolution of our platform beyond GPUs and the recent expansion of a partnership with NVIDIA; three, a rapidly growing data center footprint underpinned by unmatched execution and a strategic approach to capacity expansion; and four, a disciplined financial model deliberately designed to invest ahead of revenue to fulfill contracted demand, backed by \$66.8 billion in revenue backlog and providing strong visibility into durable cash flows, attractive returns, and the ability to drive down our cost of capital. We will speak to each of those today.

Regarding demand, the signals we are seeing across hyperscalers, AI natives and enterprise customers are only intensifying as AI workloads get more complex, models scale faster and adoption continues to proliferate. The breadth of this demand has translated to deepening engineering relationships with our largest customers and material progress on diversification.

CoreWeave is supporting the next generation of AI pioneers. As I mentioned in Q4, we added approximately twice as many new reserved instance customers as any prior quarter in our history, including AI native and enterprise companies like Cognition, Cursor, Mercado Libre, Midjourney and Runway. We also expanded our relationships with some of our largest partners, including both of our existing hyperscale cloud customers. Demand accelerated from each of these customer types, while pricing remained stable throughout 2025, trends that we have seen continue as we have started 2026.

In total for the year, we grew the number of customers committed to spending at least \$1 million on CoreWeave Cloud by nearly 150%. These are not onetime infrastructure deployments. They represent sophisticated multi-product opportunities, the early chapters of enduring platform relationships and a growth engine that compounds as AI becomes more deeply embedded in how these companies operate.

We are also seeing significant increase in demand for prior generations of GPU architectures, where supply also remains constrained. Average H100 pricing in Q4 was within 10% of where it started the year, while average

A100 pricing increased in 2025. From our customers, we understand the demand for this infrastructure is largely for inference use cases, which are proliferating rapidly. We are signing this infrastructure into new reserved instance contracts ahead of when it becomes available, firmly putting to bed concerns about demand for older generation SKUs.

With largely all of our new 2026 capacity allocated, we continue to work diligently to expand our footprint to meet the overwhelming needs of existing and prospective customers for both near and long term. These trends reinforce our conviction in the durability of demand and the longevity of this technology, while running on CoreWeave Cloud. In light of the insatiable demand environment and the persistent signals we are seeing from customers, we are accelerating our road map with the objective of adding more than 5 gigawatts of additional data center capacity beyond our already contracted footprint by 2030.

Moving on to new avenues to monetize CoreWeave Cloud. Our platform is evolving as we unlock margin-accretive avenues for growth through new products and services, as well as offering our proprietary cloud stack outside of our data centers to the broader NVIDIA ecosystem. AI natives and enterprise customers are not just consuming our core GPU infrastructure. They are engaging with our unified platform at significantly higher rates across CPU, storage, software and development tools. The opportunity to add additional value to these customers as their AI workloads mature is substantial and represents meaningful upside over time.

This is already showing up across our platform. For example, approximately 80% of CoreWeave Cloud customers paying at least \$1 million per year have adopted one or more of our storage products. Additionally, we are seeing strong cross-selling momentum with Weights & Biases, as we added hundreds of millions of CoreWeave Cloud TCW from Weights & Biases customers in the second half of the year. We have also accelerated the development of CoreWeave's proprietary cloud stack, reference architecture and related software solutions, including SUNK and Mission Control, which orchestrate every layer of our purpose-built cloud and increasingly define the CoreWeave customer experience.

In January, we announced NVIDIA intends to test and validate our platform, including our software and reference architectures, to work towards including those offerings within NVIDIA's reference architecture for cloud, enterprise and sovereign customers. Already, we are seeing select customers license SUNK as their default research cluster management platform across their multi-cloud footprint.

We expect the broader distribution of our proprietary cloud stack to become a growing source of higher margin revenue over time. The ability to monetize our platform, both inside and now beyond our own data centers through third-party licensing agreements, substantially expands our addressable market. This represents tangible, long-term upside potential that is not reflected in the 2026 guidance, as we are providing it today.

On to execution. We ended the year with more than 850 megawatts of active power, adding approximately 260 megawatts in the fourth quarter alone across 43 active data centers, up from 32 at the start of the year. We contracted close to 2 gigawatts of additional power in 2025, ending the year with more than 3.1 gigawatts of contracted capacity, virtually all of which we expect to come online by the end of 2027. Our contracted but not yet active capacity represents latent revenue potential that we will monetize as built and delivered.

We will continue to strategically source land, power, and data center shell infrastructure. Particularly looking beyond 2026, we see robust opportunities in the current market for CoreWeave to grow our contracted power capacity, and we'll also selectively leverage our expanded collaboration with NVIDIA to accelerate our road map further to better meet demand.

Operating at this scale and pace is inherently complex. When disruptions surface, we move decisively through disciplined coordination across teams and partners. We quickly cleared the delays discussed in our third quarter earnings call and, in total, we have now delivered more than 50,000 Grace Blackwells to the impacted customer, deploying servers on a rolling basis and delivering them within weeks of receiving access to the requisite data center infrastructure.

We are delivering at this breakneck speed across several different sites, handing over tens of thousands of GPUs to different customers simultaneously. We believe CoreWeave is the only cloud platform that can move at this pace while providing the industry-leading performance and reliability that drives customers' trust and allows us to capture additional wallet share. The feedback and the results we are seeing from our closest customers is inspiring. Grace Blackwell running at scale on CoreWeave Cloud is revolutionary.

In Q4, we became the first cloud platform to reach NVIDIA's Exemplar Cloud status for GB200, while remaining SemiAnalysis' sole Platinum ranked AI cloud. We expect to remain at the forefront of execution and innovation across the AI cloud stack as we become one of the first to bring NVIDIA's new Rubin GPU platform to market in the second half of 2026, while expanding our product portfolio to include NVIDIA's Vera CPU and Bluefield storage.

The integration of these newer technologies into our proprietary cloud platform will help power new capabilities, including agentic workflows for our customers. The pace of our execution also explains why our capital expenditures for Q4 came in above guidance. Our teams were able to bring infrastructure into service ahead of our expectations, which we view as a high-quality acceleration of revenue capacity for 2026.

To put our current scale into perspective, according to third-party estimates, CoreWeave today is larger than the 15 largest neoclouds across North America and Europe combined. Bringing more than 260 megawatts online in a single quarter requires simultaneously orchestrating hardware, networking, storage and purpose-built software across more than 100,000 GPUs and millions of interconnected system components, all in near-perfect unison. This is among the most operationally complex undertaking in the technology industry. It is also what CoreWeave does better than anyone.

And finally, turning to our financial and business model. In 2026, we expect our CapEx will be at least \$30 billion, more than 2x the CapEx in 2025. I want to frame that number in clear terms. This is a reflection of the extraordinary amount of contracted demand in front of us. Our revenue backlog has grown to \$66.8 billion, and the vast majority of our intended capital deployment is to directly support this long-dated contracted demand, where we have direct visibility into our long-term margins, underpinned by durable cash flow. The dividends of these investments will compound, as you will hear from Nitin, as he provides some commentary around our targets for 2027 and beyond, in addition to our 2026 guidance.

This backlog will continue to be primarily financed with the asset-level delayed draw term loans that we introduced to this market. We expect to continue to reduce our weighted average cost of capital along the way while unlocking broader industry participation in the facilities. There is a diverse and growing demand to participate in CoreWeave's capital market journey. We have cultivated a phenomenal financing vehicle for our business that enables us to scale at the pace of artificial intelligence, while staying on our targeted path to reach investment grade.

Before I turn it over to Nitin, let me leave you with a few final thoughts. We have \$66.8 billion of contracted revenue backlog, with every contract for new capacity expected to begin generating revenue by year-end 2026. We are delivering cloud infrastructure and converting it to revenue today. We are building and operating some of

the largest purpose-built AI clusters for the world's most demanding workloads at a pace and quality second to none.

The demand driving the build-out is relentless, diversified and growing, with customers engaging across our broadening product suite. Our contracted backlog gives us and you clear visibility into the trajectory ahead. The expansion of our collaboration with NVIDIA positions CoreWeave's platform as the natural destination for cloud, enterprise, and sovereign customers seeking optimal AI infrastructure performance inside and now beyond our own data centers.

Our ability to see into the future of AI innovation and build towards its requirements is unmatched. We will continue to invest and grow this incredible market advantage with disciplined and contracted cash flows, as we deploy capital strategically to expand capacity, deepen our product suite, and develop the AI cloud that our customers demand. Our priority remains clear: deliver the most performant, reliable and efficient AI platform for our customers at global scale. The market is accelerating, and CoreWeave is primed to be both the beneficiary and the enabler of the AI revolution.

With that, I'll turn it over to Nitin.

---

## Nitin Agrawal

*Chief Financial Officer, CoreWeave, Inc.*

Thanks, Mike, and good afternoon, everyone. Throughout 2025, we executed with discipline against the strategy we laid out for the year, beginning with our IPO. We significantly diversified our customer base, more than doubled our contracted and active power capacity, and strengthened our balance sheet by unlocking new funding sources, while lowering our weighted average cost of capital. We also broadened our product portfolio, both organically and inorganically, successfully completing four strategic acquisitions to pull forward our road map. The current pace of the market and scale of demand for CoreWeave Cloud has created a clear opportunity, and in 2026, we are investing deliberately to meet it, accelerating our plans to further extend our leadership position.

Turning now to Q4 results. Revenue was \$1.6 billion in Q4, up 110% year-over-year, driven by robust customer demand and exceptional execution. Full-year revenue was approximately \$5.1 billion, up 168% year-over-year. Demand for CoreWeave Cloud continues to intensify with revenue backlog for the quarter ended at \$66.8 billion, up more than 4x this year alone.

As Mike noted, we made significant progress diversifying our customer base across hyperscalers, AI natives and enterprises, a stated goal at our IPO last year. Moreover, the customers are committing their foundational AI workloads to CoreWeave for longer periods of time, resulting in the average weighted contract length increasing from roughly four years to roughly five years.

We have \$66.8 billion of revenue backlog, with every contract for our new capacity expected to begin generating revenue by year-end. We are delivering on our commitments today. These commitments are being made to current and past GPU generations as a part of broader customer road maps, with active conversations already underway for future SKUs.

Operating expenses in the fourth quarter were \$1.7 billion, including a stock-based compensation expense of \$157 million. We were able to deploy our data center and server infrastructure faster than expected, while bringing online more capacity this quarter than any in our history. This drove the corresponding increase in our cost of revenue and technology and infrastructure spend. In addition, the increase in sales and marketing was driven by investments in scaling our go-to-market organization to capture the rapid growth of the AI opportunity.

The increase in G&A was driven by professional services related to M&A and financing activities, public company costs and additional head count to support our growth.

Adjusted EBITDA for Q4 was \$898 million, compared to \$486 million in Q4 of 2024, increasing nearly 2x year-over-year. Our adjusted EBITDA margin was 57%. Adjusted operating income for Q4 was \$88 million, compared to \$121 million in Q4 of 2024. Our Q4 adjusted operating margin was 6%. Adjusted operating income was lower than expected as a result of deploying infrastructure ahead of our expectations.

Net loss for the fourth quarter was \$452 million, compared to a \$51 million net loss for Q4 of 2024. Interest expense for Q4 was \$388 million, compared to \$149 million in Q4 of 2024 due to increased debt to support the scaling of our infrastructure. Adjusted net loss for Q4 was \$284 million, compared to \$36 million in Q4 of 2024.

Turning to capital expenditures. CapEx in Q4 totaled \$8.2 billion and \$14.9 billion for the full year, higher than anticipated due to our team's ability to put infrastructure in service ahead of our expectations. As we previewed last quarter, the meaningful growth in construction in progress in Q4 to \$9.4 billion, an increase of \$2.5 billion quarter-over-quarter, reflects the significant scale of infrastructure we are on track to deliver in the near term.

As a reminder, construction in progress represents infrastructure not yet in service and not yet being depreciated. As these assets come into service, they will drive incremental revenue and corresponding depreciation. Our financing structure is designed to match this deployment model. The large majority of our term debt is structured as delayed draw facilities, meaning capital is only drawn as the data centers are operationalized.

While global supply chains remain complex amid persistent supply-demand imbalances, we have consistently navigated these challenges through operational discipline and strategic sourcing. Our track record of bringing infrastructure online at scale gives us confidence in our ability to adapt and continue accelerating capacity deployments in 2026 and beyond.

Now let's turn to our balance sheet and strong liquidity position. As of December 31, we had \$4.2 billion in cash, cash equivalents, restricted cash and marketable securities. We continue to make significant progress in strengthening our capital structure and lowering our weighted average cost of capital. In Q4, we raised approximately \$2.6 billion via our inaugural convertible senior notes offering, where investor demand dramatically exceeded the offering size, leading to its upsize.

We also expanded our revolving credit facility in the quarter to \$2.5 billion to manage liquidity and support our various growth initiatives. In total, in 2025, we secured more than \$18 billion of debt and equity, working with more than 200 investment partners and financial institutions, reflecting the depth and diversity of capital committed to CoreWeave's growth.

As Mike discussed, in January, we announced the expansion of our commercial relationship with NVIDIA, which was accompanied by a \$2 billion investment in CoreWeave, in support of our platform, team, and shared vision for the AI infrastructure at scale. Our efforts in Q4 and over the past year to optimize our financial structure and lower our weighted average cost of capital is evidenced by the 300 basis points decline in our weighted average interest rate during the year, and represents a total reduction of nearly 600 basis points since 2023. To put that in perspective, the 300 basis points improvement represents nearly \$700 million in annualized interest savings based on our Q4 debt balance.

Going forward, we expect to continue to be able to reduce our weighted average cost of capital, as capital providers and rating agencies increasingly appreciate our best-in-class execution, as well as the durability of and

visibility into the cash flows that underpin our take-or-pay customer contracts. We have no debt maturities until 2029, other than self-amortizing contract-backed debt and OEM vendor financing.

Turning to tax, we recorded a noncash tax benefit in Q4, driven primarily by the impact of One Big Beautiful Bill Act. Our tax rate might fluctuate significantly in the future due to similar factors.

Demand continues to intensify and diversify across all customer categories. We are accelerating investments deliberately to capture the contracted demand, and the long duration of those commitments give us clear cash flow visibility to deliver best-in-class cloud margins as the deployed capacity matures.

We expect 2026 CapEx of \$30 billion to \$35 billion, which is more than double our 2025 investment. Substantially, all of it is tied to our already signed customer contracts that we intend to bring online this year, as we expect to double our active power capacity to more than 1.7 gigawatts by year-end.

As I have described previously, when new capacity comes into service, data center lease costs, including power and depreciation expense commence, while customer revenue ramps over subsequent months. In 2026, this effect is amplified by the scale of our deployment program. We will be bringing online roughly double the capacity of 2025, which means a corresponding increase in depreciation running ahead of associated revenue recognition.

For full-year 2026, we expect revenue of \$12 billion to \$13 billion, representing approximately 140% growth year-over-year at the midpoint. We expect adjusted operating income of \$900 million to \$1.1 billion. We anticipate margins will ramp sequentially from low-single digits in Q1, expanding in each of Q2 and Q3, and returning to low-double digit levels by Q4 as deployed capacity matures and revenue scales against the existing cost base.

Our 2026 margin progression is a result of deliberate investments we are making to meet the insatiable demand for our platform. As our business and growth normalize, we remain confident in our ability to achieve 25% to 30% margins over the long term. Our mature revenue contracts generate contribution margins in the mid-20s, which, combined with the ramp-up of margin-accretive products and services we continue to unlock, gives us conviction in our ability to achieve this target.

Our 2026 guidance excludes any potential meaningful revenue or margin benefits from the further monetization of CoreWeave's proprietary cloud stack to other NVIDIA cloud, enterprise, or sovereign customers, which we do expect to begin in 2026 and to become more meaningful in the coming years. This represents tangible long-term potential upside.

For Q1 specifically, we expect revenue in the range of \$1.9 billion to \$2 billion. We expect Q1 adjusted operating income between \$0 and \$40 million. Q1 represents the trough in our annual margin trajectory as we expect our CapEx deployment to be \$6 billion to \$7 billion of infrastructure in the quarter, as we continue to bring online significant further capacity beyond the approximately 260 megawatts we added in Q4. Our Q1 interest expense is expected to be in the range of \$510 million to \$590 million.

The long-term nature of our contracted revenue backlog provides us with visibility well beyond 2026. As we continue on our hyper-growth trajectory, we expect to exit 2026 with annualized run rate revenue of \$17 billion to \$19 billion, which we expect to grow to more than \$30 billion of annualized run rate revenue as we exit 2027. We are not building towards this trajectory speculatively. Contracted customer demand, deep strategic partnerships, active infrastructure deployment, industry-leading capabilities, and a thoughtful approach to capital markets give us the confidence to put these numbers forward.

We delivered a strong fourth quarter and full year, capping a transformative 2025. We grew our contracted revenue backlog to \$66.8 billion, while meaningfully diversifying our customer base, secured more than \$18 billion in debt and equity capital at progressively lower costs, and strengthened our platform through new product, services, and strategic acquisitions. We entered 2026 with 850 megawatts of active power across 43 data centers, on track to exceed 1.7 gigawatts by the year-end, with every contract for our new capacity expected to begin generating revenue this year.

Our 2026 investment program is fully supported by contracted demand. And as we noted, our guidance excludes the potential upside from licensing CoreWeave's proprietary cloud stack, which we expect to begin contributing in 2026 and scale in the years ahead.

Thank you. We look forward to your questions.

## QUESTION AND ANSWER SECTION

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

**Josh Baer**

*Analyst, Morgan Stanley & Co. LLC*

Q

Hi. This is Josh Baer on for Keith. And thank you for the question. Congrats on a good quarter. You came in nicely ahead on CapEx, and it's great to hear the delivery delays resolved quicker than expected. Trying to align that with seeing the active power, which is more in line and revenue guidance in the range, which is well below like the typical level of upside. So I was hoping you could unpack some of those dynamics. If you're moving faster, why didn't that show up in the active power and the revenue? Maybe it did. Thanks.

**Nitin Agrawal**

*Chief Financial Officer, CoreWeave, Inc.*

A

Thanks, Josh, for your question. As we deploy capacity, a lot of that capacity came online towards the end of the quarter, and you're going to start seeing the monetization of it in 2026. We continue to build our capacity at a rapid pace. As we talked in our prepared remarks, we will continue to deploy that capacity for 2026 throughout the year as well, including Q1, and that's the impact that you're seeing. Relative to the Q1 number, like, we are basically providing the guidance for the first time for 2026 and Q1 at this moment.

**Josh Baer**

*Analyst, Morgan Stanley & Co. LLC*

Q

Okay. Thank you. And really great to see the list of enterprise customers. I was hoping you could unpack what that type of contract and deal looks like from those enterprise customers. We have a great sense for what a hyperscaler mega contract looks like. But any chance you could run through size, duration, prepayment, pricing associated with those enterprise customers? Thanks.

**Michael Intrator**

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

A

So we don't speak to individual contracts. But what you're seeing is, in an environment where there is so much intense competition for the product we deliver, which is the most cutting-edge computing infrastructure delivered

through CoreWeave Cloud, the contracts largely look very similar to the hyperscale contracts in tenure. And we work with each of the individual enterprise clients as we're putting together an appropriate structure for their business model and for their clients.

**Josh Baer**

*Analyst, Morgan Stanley & Co. LLC*

Great. Thank you.

Q

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

**Amit Daryanani**

*Analyst, Evercore ISI*

Yes. Thanks for taking my question. I guess, my question is really around the cost of financing, especially given the \$30 billion-plus kind of CapEx number we have for the year. I'm just wondering, as you continue scaling capacity, can you sort of quantify where you estimate your blended cost of capital is. How has that really evolved over the last 12 months? And when you negotiate with these data center operators, how do they assess your credit profile? Is it really tied to your customer contracts and who they are, or is it something else? And then just on the financing side, does the NVIDIA credit support, the guarantor framework, help translate into a measurable step-down in your borrowing cost, you think, in 2026?

Q

**Michael Intrator**

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

Yeah. Thank you for the question or questions. So, look, we've made incredible progress at the company. As the company matures as a business, as we have more extensive track record of operating this infrastructure, working with the clients, delivering infrastructure, you've seen our cost of capital drop 300 basis points in the last 12 months. You've seen it drop 600 basis points over the last two years. We expect that that will continue. It is a trend that's being driven by our business increasingly performing well with these ETL structures.

A

When you're talking about the data centers, we added close to 2 gigawatts worth of infrastructure in 2025. And just to give you scale and perspective, at the end of 2024, we had 1.3 gigawatts in our portfolio. So you've seen a material increase in our capacity to access, build, and drive data center contracts. And we're excited about that. It's an important stepping stone for us as we continue to kind of drive the business. And once again, the ability to enter into contracts with that scale of data center capacity is once again a reflection of the business maturing, the creditworthiness and scale of the business increasing.

As far as the data center operators go, I can tell you what I think, right? And what I think is that data center operators are very interested in working with CoreWeave. They're looking for a diversified portfolio of tenants in their data centers. They're looking to go ahead and get exposure to a company like CoreWeave that represents so much of the AI infrastructure that's going to be ultimately delivered. And so they kind of look at us as a pure-play way of really getting access to the scaling of artificial intelligence, and they want that exposure.

As far as our relationship with NVIDIA in terms of accelerating our ability to get access to data centers, I think the perspective that you should take here is that, obviously, working with an investment-grade counterparty as the offtake will have an impact on the cost of capital or the cost that is associated with the data center. Obviously, working with NVIDIA, which we do selectively, but certainly not exclusively when we're building out our data center portfolio, will have a positive impact on the costs associated with our data center footprint.

**Amit Daryanani**

*Analyst, Evercore ISI*

Q

Perfect. Thank you very much for all the insights.

**Operator:** Your 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. And congratulations on just very, very strong bookings. Mike, some of the AI models have demonstrated a pretty gigantic leap forward in the last couple of months. And the one that's in the headlines is Claude Code. But I don't think we have seen models yet that were fully deeply trained on some of these gigantic Blackwell or GB200 or NVL data centers, really the stuff that CoreWeave has pioneered and mastered.

I'm curious what you're hearing in the marketplace just in terms of how those Blackwell-based models are coming along. If we end up seeing [ph] GBT6 (00:40:49) or any of the other ones in the next three to six months. Do you think it's going to feel like a huge step forward in their capabilities or is it looking more like a steady evolution on the Blackwell systems?

**Michael Intrator**

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

A

Look, the Blackwell systems are amazing, right? They represent the next step function in computing power that allows these data scientists, these companies that are driving the models to be able to build and scale infrastructure in a way that they just haven't been able to historically. And my expectation is, and certainly every indication from the model companies, is that the rate of increasing performance from these models, we're just getting going.

Now, it is early in the deployment of Grace Blackwell, right? Like, there are not that many clusters that exist at the size and scale that we talked about, we have already delivered. As those clusters come online within our portfolio, within the global portfolio, I think it stands to reason, and you will see step functions in performance that are associated with this new technology. We're really excited about it. And I think our customers are extraordinarily excited about it because they understand what they're going to be able to do with this technology that is so incredibly performing, both from a training perspective, but also from an inference perspective when it becomes available for them.

**Mark R. Murphy**

*Analyst, JPMorgan Securities LLC*

Q

And Mike, thank you for that. And just by extension, because you just said the word inference, how are you weighing the merits of focusing on the NVIDIA reference architecture? It's obviously very powerful for the massive training runs and some work beyond that. Just the other side would be any inclination to work with custom ASICs that they do legitimately seem to offer better inferencing price performance. And then I'm – and then, obviously, NVIDIA's acquisition of Grok maybe kind of – I don't know if you think that that sort of resets the playing field in a way that seeing NVIDIA reference architecture might kind of reign supreme even for inferencing. I'm just wondering how you sort of project that forward.

**Michael Intrator**

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

A

Yeah. So, look, whenever we have these calls, whenever I'm asked about this, I kind of speak to the way that we've gone about building our business, which is we are client-led. Our clients are coming to us and they are telling us that the infrastructure that they need in order to drive their business. And I want to be clear that when they say infrastructure, it's not a training infrastructure, it's not inference infrastructure. It's AI infrastructure, right? And they're coming to us specifically because we're able to deliver such an incredibly performing [indiscernible] (00:43:59) of the NVIDIA technology.

They know we're great at it. That's why they come to us. Are they looking for other technologies from other providers? That stands to reason. But what I believe is that – or what I know is that we are unable to catch up with the demand signals that are coming in for the product that we deliver. And so we are going to focus on continuing to drive the solution that we have that is so performant and that has overwhelmed our ability and the market's ability to deliver infrastructure for the past three years.

**Mark R. Murphy**

*Analyst, JPMorgan Securities LLC*

Q

Thank you very much.

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

**Michael Intrator**

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

A

Brent, we can't hear you.

**Brent Thill**

*Analyst, Jefferies LLC*

Q

Hey. Good afternoon. Nitin, I had a quick question just on the guide. And just from a perspective, I know when you look at the revenue guide, you were in line, op income a little lower and your CapEx was way higher. I guess, it just – it kind of illustrates even the guide you gave us all that some of the metrics can really vary.

I'm just curious just in terms of how you're thinking about the guide going forward. Are some of the variables out that you've taken out from maybe what you saw in Q4? Are those variables taken out? Or has your guidance changed a little bit? Again, I know this is incredibly difficult to make an estimation, but some of the numbers were effectively kind of outside the range of what you initially gave us.

**Nitin Agrawal**

*Chief Financial Officer, CoreWeave, Inc.*

A

Thanks, Brent, for your question. So from a guide perspective, let me break it down by a few variables here. We talk about the CapEx numbers that is fundamentally in service of our contracted customer backlog, which we disclosed this quarter to be at \$66.8 billion. And that's what is driving the investment in our platform. And when you think about the revenue ramp, we talked about that as well that almost all of our – most of our – all of our contracts that we are – our backlog would start generating revenue in this year. So that's the ramp that you are seeing. We delivered 850 megawatts of power in this fiscal year. And for the year, for the 2026, we expect to be at 1.7 gigawatts of power.

When you think about margin, when all of this comes together in margins, our margin progression effectively is a result of these deliberate investments that we are making to meet the insatiable demand that we have in our platform. We talked about Q4. Q4 alone, we brought 30% of our total active power base, which naturally creates

some near-term margin compression as capacity costs ramp ahead of full revenue maturity and recognition. As I mentioned in my remarks, Q1 represents the trough of what we would see in margins. And then from there on, as we scale into the capacity deployed, we will expand margins quarterly from there, returning to low-double digits by Q4.

We also talked a little bit about long-term trajectory of this business. Our strategy and our management philosophy has continued to be to invest in terms of customer demand with contracted backlog, which is what still continues to be the case. Over the long term, how it manifests itself in our business as its growth normalizes, we remain confident in our ability to achieve 25% to 30% margins. The factors that give us confidence in that domain, if we look at our mature fully ramped contracts and that portfolio, that generates contribution margins in the mid-20s.

We continue to ramp up our margin-accretive products and services in our product portfolio. For instance, we had announced in Q3 that our storage revenue on our platform eclipsed \$100 million in ARR. Today, we also discussed how attach rates for storage are now at 80% in our large customer base. While not included in our 2026 guidance, we see tangible long-term upside potential from further monetization of CoreWeave's proprietary cloud stack to other NVIDIA cloud, enterprise, and sovereign customers. So what you're seeing in 2026 is a reflection of the acceleration of the growth in the back of our existing backlog, which continues to grow with customer demand.

---

**Michael Intrator**

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

A

Yeah. So I just – I wanted to add a couple of things here, right? And so our margins reflect the cost of building tomorrow's revenues, right? That's what we're doing, right? As Nitin said, the fundamental margins at a stabilized facility are in the mid-20s, right? And as we continue to build our infrastructure, as there is more infrastructure online, that will come to bear. The variance that you're seeing is a function of how much infrastructure we are bringing on versus the installed capacity.

We brought on 260 megawatts worth of power in Q4. It's fully a third of our installed capacity. So the variance that you're going to see there is higher. As we continue to grow and scale our company, as the incremental data center capacity that we bring on becomes relatively smaller, you will see less variance from us, right?

This is what the acceleration looks like, making the decision to go ahead and invest to pull in tomorrow's revenue, to be able to serve our clients, is a fundamental strategic decision that the company made. And we're doing this extremely responsibly because we're not doing this, we're going to build it and they're going to come. We're doing this leaning into the backlog of contracts that we have already sold.

---

**Brent Thill**

*Analyst, Jefferies LLC*

Q

Great. Thanks for the color.

---

**Operator:** Your next question comes from the line of Gabriela Borges with Goldman Sachs. Please go ahead.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Hey. Good afternoon. Thanks for taking my question. Nitin, I wanted to ask you about the diversity of customers that you have on your platform. I'm curious if you can share with us your observations on how the unit economics

or how the attractiveness of how these customers are using the CoreWeave platform is different between types of customers. So, a little bit of a broad question. I know that your pricing model is based on dollars per GPU and then the length of the committed contract, but curious if you could share your observations on customer behavior across the different cohorts. Thank you.

---

**Nitin Agrawal**

*Chief Financial Officer, CoreWeave, Inc.*

A

Yeah. So I think a lot of this depends upon the variables that we've talked about in terms of how we structure our contracts. The term length, the amount of upfront payment, the generation and the demand for that capacity at that moment, all dictate into it. Fundamentally, as Mike described, like, our contracts look mostly similar across our customer base with the exception, of course, the volume element that we look at things when we are talking about larger customers versus smaller customers. But across the board in our customer profile, we look to generate similar economics for the infrastructure that we are generating as with the market dynamics go on.

Mike talked about how, for Hoppers, we are continuing to see incremental demand and demand of re-contracting those Hoppers at about 10% of the original ASPs when they were first contracted, to A100s where the ASPs are actually increasing as we write newer contracts. Those dynamics are broader market dynamics, but across our customer base, the economics kind of look relatively similar.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Appreciate the color.

---

**Michael Intrator**

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

A

Yeah. The only thing I would add is that one of the things that I do think is very exciting is that many of the enterprise customers, many of the smaller AI native customers, they're coming on to our infrastructure, and they have the ability to use the H100s and the A100s. Their ability to build product and to be able to serve inference with that infrastructure is a really wonderful sign of the depth and resiliency of the bid that is looking for compute. It's new use cases, people doing things that we've never seen before. It's really exciting.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Thank you for the thoughts.

---

**Operator:** Your next question comes from the line of Ben Reitzes with Melius Research. Please go ahead.

---

**Ben Reitzes**

*Analyst, Melius Research LLC*

Q

Hey, guys. Thanks a lot. I wanted to ask the other side of – I believe it was the Brent Thill question. One of the things going on in the market is there's all this CapEx and not enough margin or cash flow necessarily in the near term. And so I think that's what he was getting at is that you guys were talking about your margins being 20% to 25% over the long term, but – and then your confidence to get there.

What is your confidence in the rate of the CapEx growth? Like, we understand you're spending now to get the revenue. But what is – does that growth rate moderate more than we're thinking as those margins go up? Do you

need to keep spending at this kind of upside versus the Street on the CapEx side in order to hit those numbers? Just a little bit more color on the CapEx side that balances Brent's question on the margin would be really helpful in the rate of that trajectory. Thanks a lot.

---

**Michael Intrator**

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

A

Yeah. Thanks, Ben. It's a good question, right? And I think it's really important that we deconstruct that question into at least two pieces, right? The first piece is that our business is built on a success-based model where clients come to us and buy long-term contracts to get access to the infrastructure to build their business, right? And so when a hyperscaler comes to us and buys infrastructure for five years, they are going to be purchasing that infrastructure from us at a fixed price for five years.

And that is a very stable way to go about building our revenue and our margins. And that's where Nitin gets the confidence around the margins from is because we know not only what we're going to make now, but we know what we're going to make as those contracts move through the five-year cycle that they have been put under a take-or-pay contract. And so that is a very stable way to go about building a business. It's a very stable way to go about getting access to the capital markets, so that we're able to finance the builds, right? And that's a fundamental building block of how CoreWeave has gone about building its business really since the beginning.

The second question kind of, I think, embedded in that is what is our confidence interval that there will be a continuation of demand for computing infrastructure. And this is a great opportunity for me to talk a little bit about what we're seeing in the demand stack, because the demand stack is actually fascinating, right? Not only are we seeing the proliferation of demand across the economy, going from where it was initially really housed within the hyperscaler clouds and the foundation models. You're now seeing it kind of explode into the enterprise. You're seeing it move into sovereign. You're seeing all these new participants beginning to come in and securing the infrastructure that they need.

You're also seeing a really fascinating component where it's moving from just the GPU, which was the initial wave of demand that we used to launch our company, but really now starting to move out into storage, into CPUs, into – and that is really a function of the portfolio of clients that are using our infrastructure really extending into the application layer.

And so, as far as the demand goes, we've got our fingers on the pulse in a way that very few companies in the world have. We are getting information fed in from across the entire economy as people are trying to get access to the infrastructure that they require. So we're very confident in our contractual position relative to the portfolio of clients that we have. And we're very excited about the portfolio of different types of clients and we're very excited about the different types of compute that they're using. And so it's really threefold of information, threefold of confidence around the drive towards accessing compute via CoreWeave.

---

**Nitin Agrawal**

*Chief Financial Officer, CoreWeave, Inc.*

A

And a couple of data points that I'll just add is we discussed this in our prepared remarks. For our 2026 CapEx, substantially all of it is tied to already signed customer contracts that we intend to bring online this year, and we expect to double up deployed power capacity. So, that's one data point.

The second data point that I want to make sure that you get is from an EBITDA margin perspective, I know you talked about where is the cash associated, the EBITDA margin perspective, we generated 57% margin this quarter. When I think about the long-term contracted customer contracts that we have and those scaled customer

contracts, we talked about they are in the mid-20s from a contribution margin perspective. When you look at the EBITDA margins for those, they are in the 70 percentage zone. So that is something that gives us confidence in terms of when our contracts scale, they generate a lot of cash for us.

**Ben Reitzes**

*Analyst, Melius Research LLC*

Okay. Thanks, guys.

Q

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

**Brad Zelnick**

*Analyst, Deutsche Bank Securities, Inc.*

Great. Thanks so much for taking the question. And congrats on an absolutely amazing year. Just, guys, as we look to Rubin contracts ahead, are you seeing demand for similar five-year duration? And is there anything different about how these deals are priced, the amount of prepayment that you'd expect or anything else? And how should we think about the ROIC on these deals versus prior-generation contracts and the economics that you've outlined in your original S-1? Thank you.

Q

**Michael Intrator**

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

Yeah. So, thanks. It hasn't even been a year yet, although it does feel like it. I think we've got 11 months under our belt. So, look, in the written statements that we made, we talked a little bit about the fact that one of the trends that we're seeing is the extension of the contracts from an average of four years up to an average of five years. And obviously, that is great for us. It's stabilizing for our business. It gives us a lot of confidence around building and scaling the infrastructure.

A

And so that's sort of the first thing I'd highlight to you is that the clients that are coming in and using this infrastructure, they're gaining more confidence in the longevity of the usefulness of the infrastructure, which once again gets back to many of the important components that are required for us to build our business. The pricing on the infrastructure, we kind of take a – we think about that from a margin perspective. And as the infrastructure changes in price, we are altering our pricing of how we deliver it to clients in order to kind of target the type of margins that Nitin spoke to. And so, like, those are some of the trends that you're seeing.

Now as far as prepay goes, prepayment is always a lever which we can use with clients to change the economics around a contract. One of the things that's very exciting for us is, as we continue to drive down our cost of capital, our dependency on prepayments is reduced.

**Brad Zelnick**

*Analyst, Deutsche Bank Securities, Inc.*

All right. Thank you.

Q

**Operator:** Your final 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. I appreciate you taking the question. This is for Nitin or Mike. Can you just speak to what gives you confidence in the \$30 billion run rate by 2027? And how much of that is already booked versus business your team needs to go get? And maybe as a second part, just if you could speak to any change in demand you're seeing? Mike, you touched on some of the segments of the market, but just any change in demand you're seeing across those segments and how you prioritize across those as well? Thanks very much.

---

**Michael Intrator**

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

A

I'm going to do it in reverse order, if that's okay. One of the things...

---

**Michael Turrin**

*Analyst, Wells Fargo Securities LLC*

Q

Totally okay.

---

**Michael Intrator**

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

A

Yeah. So one of the things that we, as a business, are very interested in is making sure that we have a diversified perspective on what the compute is being used for. And so we're really out there working with everyone that consumes compute in every way that we can, because we feel like that gives us the best view on where the demand is going to come from.

And I've said this before, one of the things that's going to happen here is as compute becomes more available, you're going to see businesses that don't even exist yet, ideas that don't even exist yet have an opportunity to come into existence. Those will be new clients of ours. And we want to be able to pick them up right at the beginning as they go ahead and build these new incredible businesses.

Moving on to your question around the \$30 billion run rate, like, what we are doing is we are taking the contracted power that we have and we are projecting out when the existing contracts that have already been sold, and like I said, we are virtually sold out in 2026 of all of our capacity, and then continuing to add contracts that will be allocated once they come online in 2027. And we have vast and sustained interest from our clients to get more capacity to bring on more compute.

And these are some of the largest, most creditworthy companies in the world. These are some of the most important AI labs in the world. These are the people that are building the AI future. And they are trying to secure infrastructure through CoreWeave, and it's really exciting. And when you move through that exercise, we have a lot of confidence in the \$30 billion number that we've put out there.

---

**Michael Turrin**

*Analyst, Wells Fargo Securities LLC*

Q

Thank you.

---

**Michael Intrator**

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

All right. So, like I said, we've got 11 months in here, and I appreciate all of you working with us as we've built this company. So as we wrap up here, I want to thank the CoreWeave team and our partners. None of these accomplishments would have been possible without you. I'm incredibly proud and humbled of the execution across the organization from product velocity and innovation to operational excellence and financial rigor. The

focus and intensity across our organization is what enables us to continue our hyper-growth trajectory and the size of the incredible opportunity that lies ahead.

Thank you, all, for joining us today. We appreciate your support, and we look forward to updating you in the future. Thank you.

**Operator:** Ladies and gentlemen, this concludes today's call. Thank you all for joining. 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 2026 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.