Corrected Transcript

11-Aug-2026

CoreWeave, Inc. (CRWV)

Q2 2026 Earnings Call

CORPORATE PARTICIPANTS

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

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

OTHER PARTICIPANTS

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Brad Zelnick
Analyst, Deutsche Bank Securities, Inc.

Irvin Liu
Analyst, Evercore ISI

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Brett Knoblauch
Analyst, Cantor Fitzgerald & Co.

MANAGEMENT DISCUSSION SECTION

Operator: Hello everyone. Thank you for joining us and welcome to CoreWeave's Second Quarter 2026 Earnings Call. After today's prepared remarks, we will host a question-and-answer session. [Operator Instructions]

I will now hand the conference over to CoreWeave. Please go ahead.

Unverified Participant

Thank you. Good afternoon and welcome to CoreWeave's second quarter 2026 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 because of factors that 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. Q2 was an exceptional quarter for CoreWeave. We outperformed our plan across the board, with the operating leverage we have been building, beginning to show up clearly in our results. Extraordinary execution across the organization drove record financial performance, rapid capacity growth, broadening customer demand and continued platform innovation.

We generated record revenue of $2.6 billion, up 112% year-over-year, increased revenue backlog to $104 billion, while driving rapidly expanding enterprise adoption. This figure does not include the over $25 billion of net new customer commitments added in the early weeks of Q3. We continue to execute on our power strategy, reaching 1.5 gigawatts of active power, adding nearly 500 megawatts more than any quarter in our history and more than tripling year-over-year. We remain firmly on track to reach at least 8 gigawatts by 2030.

We grew adjusted operating income to $128 million, with margins expanding meaningfully as our scale increasingly translates into operating leverage. And we continued to broaden our technology stack, delivering seven new AI platform capabilities and achieving multiple industry first milestones that enable customers to build, deploy and operate AI faster and at greater scale. Our incredible progress is a testament to the entire organization, and the business is only getting stronger.

In Q2, the customer contracts we signed came with contribution margins, we expect to be 5 to 10 percentage points above those added in recent quarters. This is more than a collection of milestones. It is evidence that the AI market is developing in line with the convictions on which we built CoreWeave. We believe that the AI era is here and will ultimately touch every part of the global economy. That every organization is being transformed, creating an opportunity to reinvent established markets and create entirely new ones. That the future will be led by the pioneers who seize that opportunity.

Both AI native companies at the frontier and changemakers inside established enterprises. That learning and iterating at lightspeed are now tablestakes for AI leadership. And we believe those pioneers need a new kind of platform to unleash AI's potential at scale. These beliefs are the operating assumptions that drive our strategy. They guide our product roadmap, our capital allocation, our partnerships, and ultimately how we serve our customers.

Today, I want to discuss how our vision is translating into four areas. One, broadening demand and customer adoption. Two, the continuous AI development tools we provide on our platform. Three, the performance and economics enabled by our AI native architecture. And four, the power and supply chain foundation that will support years of growth. AI is transforming every organization.

The debate around future demand for AI cloud infrastructure will likely continue, but what we know with certainty is informed by our customer's actions. Demand continues to intensify as the market broadens across sectors, geographies, workloads and generations of GPU architecture. AI is no longer confined to frontier model labs. It is becoming embedded in software, industrial systems, financial markets, enterprise workflows and national security missions. We see that breadth in our backlog, in the new commitments we have signed and in the utilization and pricing environment across our platform.

Pricing and margins for our Blackwell and Vera Rubin SKUs are setting new highs while pricing for prior generation SKUs is at or above where it was years ago. Our near-term capacity remains effectively sold out. That is translating into signed commitments on increasingly favorable terms from a broadening set of customers and is positioning CoreWeave to gain market share for years to come.

Every organization is being challenged to rethink what is possible. AI is not simply accelerating existing processes. It is enabling enterprises to redesign core functions, create new offerings and enter markets that did not exist before. Caterpillar is a powerful example. Together, we will deploy NVIDIA's Vera Rubin platform to support Caterpillar's physical AI training and inference at industrial scale.

Using CoreWeave's AI cloud infrastructure as its data factory, Caterpillar will train specialized models that enhance the intelligence and productivity of autonomous construction equipment. Life sciences is emerging as another important growth vertical for CoreWeave as organizations tackling some of the world's most complex scientific challenges increasingly turn to our platform. We recently welcomed Isomorphic Labs as a new customer and are excited to support them in their mission to solve all disease.

Financial services remains a major growth area for us as well. We recently added Flow Traders and IMC to our growing roster of systematic trading firms. These customers are using CoreWeave to develop and deploy the next generation of AI models for quantitative trading. They select our platform because of our ability to orchestrate high performance workloads with the speed, reliability and efficiency those applications require.

And in the public sector, our collaboration with Leidos marks an important step in the growth of CoreWeave Federal. Together, we are working to accelerate the delivery of secure AI capabilities for defense, national security and intelligence missions.

The future of AI is being built by a new class of innovators, some are AI native companies operating at the frontier. Others are changemakers inside established enterprises who are willing to challenge the status quo. CoreWeave serves both. Companies such as Decart are using our platform to develop Oasis 3, the first API accessible world model for physical AI. IBM is using CoreWeave to experiment securely with reinforcement learning, agent tool use and model evaluation. And through Monolith, our specialist field engineers, work directly alongside customers such as Nissan and ZF to accelerate the development of enterprise ready AI applications.

Demand now extends beyond our infrastructure as well. Through CoreWeave Omni, we are seeing significant interest from sovereign enterprise and cloud customers alike. In the past few weeks we signed our first deal, which will begin to scale in 2027. These examples differ by industry and use case, but the pattern is consistent. AI is moving from experimentation into core operations, and the organizations that act decisively are creating an advantage. Deployment is no longer the finish line.

As AI moves into production, the way applications are built is changing, and the leaders will be those who learn and iterate the fastest. For the last several years, many organizations treated a model like a deliverable. Train it, deploy it, and move on. Enterprises no longer operate that way. Training, inference, evaluation and improvement now form a single continuous loop. Models and agents in production generate real world data. That data informs evaluation, driving new experiments which improve the model or application before being redeployed into production. The loop repeats and capability compounds over time. That shift fundamentally changes both the demand curve and the economics of AI.

Compute is no longer a one time requirement, concentrated at the beginning of a model's life. It becomes an ongoing requirement that grows with every application in production and every cycle of improvement. Our AI native platform was built for this. It spans cutting edge cloud infrastructure, a rapidly growing managed inference business, leading developer tooling and agent solutions, and a best in class orchestration and observability layer powered by mission control. Together, these capabilities give customers one integrated environment.

Customers deploy applications through CoreWeave Inference using our models or the ones they have customized with our serverless capabilities. Monitor performance with Weights & Biases, evaluate applications in production, experiment with new models, refine their performance through serverless reinforcement learning or sandboxes, and validate every change against quality, performance and cost before returning it to production.

CoreWeave ARIA, our AI research and iteration agent, accelerates that process further by analyzing thousands of evaluation runs, surfacing insights in minutes, and recommending the next experiment. That allows customers to compress the time between an idea and experiment and a production improvement at lightspeed.

In Q2, we introduced seven new AI platform capabilities and achieved multiple industry firsts. These innovations were built alongside our customers and partners to solve real production challenges. That is why we are seeing such strong adoption. And just this week, we surpassed 1 billion model training runs tracked on our platform. Behind that number are millions of experiments, thousands of research breakthroughs, and a growing community of engineers, researchers and organizations building the next generation of AI.

Our AI development services, which carry higher margins, are also being adopted by a broader set of customers than our core cloud. That is proving to be a natural customer expansion path because the developers building AI applications today are the AI cloud infrastructure customers of tomorrow. By serving them early, we are establishing relationships that naturally expand as AI workloads scale.

We have seen an explosion of growth in our managed inference platform in the few months since its launch, with growth constrained only by our near-term capacity. Across serverless offerings and dedicated deployments, CoreWeave is monetizing tokens, while giving customers flexibility in how they consume our platform. Companies such as Grammarly and [ph] You.com (14:09) are moving from experimentation to real production traffic, running AI coding agents, serving their own fine tuned models, and deploying open weight models at scale.

Customers shouldn't have to trade speed for cost and on CoreWeave they don't. They choose our platform for the combination of total cost of ownership, quality, breadth of service and performance that is reflected in our consistent leadership across cost per token and speed to first token on leaderboards like artificial analysis for open source models, including Kimi K2.6, K2.7 Code, GLM 5.2 and MiniMax M3. And that leadership is converting directly into revenue.

In the past few months, since its launch, booked ARR for our managed inference platform has grown from $1 million to more than $100 million. We expect to exit 2026 with at least $250 million of managed inference ARR. Pioneers need a different kind of platform. The continuous AI life cycle cannot be supported by simply adding GPUs to a general purpose cloud. It requires a new approach from power cooling and rack design through networking, orchestration, observability, developer tools and managed services. That is why CoreWeave is purpose built for AI.

Our platform is singular in its depth, breadth and technical capability. In Q2, we became the first cloud provider to bring up and validate NVIDIA's Vera Rubin NVL72, leveraging our innovations in software defined liquid cooling and rack management to extend our track record of being first to market. We also set new MLPerf records for training and inference with open source models, running on the NVIDIA's Grace Blackwell platform and in our tests achieved the lowest cost per token for inference.

However, performance alone is not enough. Customers need enterprise grade security and observability, reliability and compelling economics. According to Signal65, CoreWeave delivers total cost of ownership, estimated to be up to 47% lower than the average hyperscaler. Customers also require our platform, which integrates these capabilities with a broader portfolio of storage, CPU and networking services across a distributed footprint of data centers globally. In July, Gartner named CoreWeave a visionary in its 2026 Magic Quadrant for cloud AI infrastructure. We believe that recognition provides additional independent validation of our approach to building the AI cloud.

These achievements are not isolated technical milestones, they translate directly into faster deployment, higher utilization, better application performance, and lower cost for customers. Pairing product depth with best in class performance, quality and market leading TCO is a winning formula for our customers and for CoreWeave.

CoreWeave is the foundation for AI at scale. This market requires a foundation at a magnitude unlike anything that came before. That means securing power, sites, cooling, hardware, storage, networking and supply chain inputs well ahead of need. And operating them as one integrated system.

As I shared at the top of the call, we ended Q2 with 1.5 gigawatts of active power, adding close to 500 megawatts in the quarter alone. To put that in perspective, we added more power in Q2 than any single neo cloud operates in total today, according to third party estimates.

Critically, our scale is working in our favor and the math gets better from here. Each new deployment is landing against a much larger installed base than it was even one quarter ago. As that base grows, each new build becomes a smaller part of the whole, while contracted revenue from existing deployments remains in place. This is how we are transforming scale into operating leverage. It is why margins expanded in Q2 and why we expect them to continue expanding sequentially during Q3 and Q4.

We are also securing the ingredients required to sustain growth over a multiyear horizon. Contracted power grew to 3.7 gigawatts in Q2. Since quarter end, we have added roughly 500 megawatts, bringing contracted power to 4.2 gigawatts as of today. These figures exclude more than 1.5 gigawatts of further potential power from powered land we have accumulated, options we have to expand at existing sites and LOIs we have executed. Our first several self-builds are already well underway, including our first site expected to come online later this year.

Powered land forms the foundation for deeper vertical integration, giving us greater operational control and supporting enhanced long term margins. We are also expanding globally and have contracted more than 1 gigawatt of power outside the United States, including recently entering the APAC region with 360 megawatts in Indonesia. That will begin coming online in approximately 18 months. We expect international markets to become a major driver of growth as we meet customers where they and their end users operate.

All in, we have excellent visibility to our target of at least 8 gigawatts by 2030. We expect demand to meaningfully exceed supply for years. In that environment, access to power is only part of the equation. Just as important is each necessary component required to deliver the AI cloud at scale, building on our close partnerships with NVIDIA and our OEM and ODM partners, our recent long term agreement with Solidigm is one illustration of how we are de-risking access to the critical inputs needed to serve our customers.

Our investments in technology, capacity, vertical integration, supply chain and global expansion all flow from the same vision, AI is increasingly pervasive. The pioneers who move fastest will lead, and they will require a platform capable of supporting continuous learning and deployment at unprecedented scale.

Before I turn to Nitin, I want to reiterate that CoreWeave enters the second half of the year with more momentum than at any point in our history. AI is reshaping every industry. The pioneers that are building need more than compute, and that is why they come to CoreWeave for an AI cloud designed for the full AI lifecycle, serving any workload from frontier training to rapidly scaling inference.

Demand continues to exceed supply across sectors, geographies and generations of infrastructure. We are delivering at extraordinary scale to a diverse set of customers with improving operating leverage and visibility into the power and critical components required to sustain growth for years. The opportunity ahead is generational. CoreWeave is the essential cloud for AI. Our conviction in our strategy has never been stronger and our execution continues to reinforce it.

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

Nitin Agrawal

Chief Financial Officer, CoreWeave, Inc.

Thanks, Mike, and good afternoon, everyone. Q2 was an exceptional quarter for CoreWeave, marked by intense customer demand, significant ramp of our active capacity and continued execution against our product and financing roadmaps. Perhaps most importantly, Q2 marked the quarter in which we saw margins inflect, expanding sequentially as we had discussed over the past several quarters.

Before diving into results, I wanted to spend a few moments touching upon how demand dynamics are evolving in the current environment, as well as its implications on cash flows and the value of our rapidly growing infrastructure footprint. Demand for CoreWeave Cloud remains exceptionally strong across the entirety of our customer base, with demand from multiple customers for each GPU we bring online.

We are being disciplined in how we allocate our scarce cloud capacity. We are prioritizing opportunities that are strategically important, adding new customers while deepening long term existing relationships and delivering attractive returns that, as Mike noted, are expanding further.

The scale of our AI products and services, beyond GPUs, also continues to ramp significantly as customers consolidate spend with us. These margin accretive businesses, including storage, CPU, networking and software, already exceed $400 million of ARR as of Q2. We expect they will continue to expand rapidly.

Simply put, customer spending on CoreWeave has gone up as customers recognize the increased value we deliver. And this operating margin improvement came before our July pricing changes, which included an approximately 25% increase across SKUs in response to the current demand environment and the increasing ROI, our customers are observing from their investments in the CoreWeave platform as they shift to inference. We're also passing through component price increases.

In terms of how this translate to cash flows, as we previously discussed, a typical five year contract carries strong and still expanding unit economics across [ph] its term. (25:13) But those economics do not arrive evenly. The cost, primarily in the form of CapEx, is front loaded, requiring a combination of debt, customer prepayments and other corporate level capital to finance its build out. Once the cluster is delivered, contracted revenue ramps becoming predictable and highly cash flow generative. This is all considered in our underwriting of expected margins before a contract is signed.

The deployment delivers attractive returns, fully repaying asset level debt used to fund the CapEx while generating significant additional free cash flow. So when an initial contract ends, the cluster no longer has any leverage and we are free to recontract that cloud infrastructure or offer it to the market.

We will have generated an attractive return even before the prospect of further monetizing the cloud infrastructure. Every resale or renewal is incremental on top of the returns already earned within the initial term. What we are seeing today is that the upside of re-contracting is real as we remain largely sold out of prior generations of NVIDIA GPUs in addition to the current SKUs.

So as our earlier generation fleets roll off their original contracts, they offer the potential to deliver strong returns in the subsequent years. We are seeing this across our [ph] Ampere (26:43) and Hopper fleet. As an example, we recently signed an A100 contract that extends into 2029 at an attractive price. As a reminder, this SKU was introduced in 2020.

Clusters of prior generations of architecture offer installed energized production grade compute already running at scale. They come with a proven ROI for customers. In a market where new capacity is supply constrained and costs are rising, AI cloud infrastructure in production is a scarce, valuable asset. While we have built a business whose economics do not rely on recontracting after initial customer term, increasingly we are seeing longer utilization at higher prices, offering the potential for significant further upside. With these tailwinds at our back, we are more confident than ever in the long term ROI of our product and capacity investments, enabled by our industry leading AI cloud services.

Now turning to Q2 results. Revenue was $2.6 billion in Q2, up 112% year-over-year and 24% sequentially, driven by continued strong execution and customer demand for CoreWeave's AI cloud platform.

Revenue backlog ended the quarter at $104 billion, up 246% year-over-year. As Mike noted, this does not include the over $25 billion of net new customer commitments we added early in Q3. Of the existing backlog, more than 50% is attached to a contract where customer delivery has commenced. We expect this figure to reach more than two thirds of our Q2 backlog by the end of this year.

Operating expenses in the second quarter were $2.6 billion, including a stock based compensation expense of $165 million. The increase in our operating expenses was a direct result of scaling our active power, while converting backlog into revenue. This drove the corresponding increases in our cost of revenue and technology and infrastructure spend.

In addition, the increase in sales and marketing was driven by increased investment in our go to market organization as we further diversify our customer base and expand into new products and markets.

G&A increased driven by personnel cost to support our growth, while continuing to moderate versus revenue growth. Adjusted EBITDA for Q2 was $1.5 billion compared to $753 million in Q2 of 2025, doubling year-over-year. And adjusted EBITDA margin was 59%. Adjusted operating income for Q2 was $128 million compared to $200 million in Q2 of 2025 and up from $21 million last quarter well above the high end of our guidance as operating leverage comes into a business with scale. Adjusted operating margin was 5%. Margins expanded as we scaled despite continuing to incur significant ramp costs.

Net loss for Q2 was $626 million, compared to a net loss of $290 million in Q2 of 2025. Interest expense for Q2 was $640 million compared to $267 million in Q2 of 2025, driven by increased debt to support the continued scaling of our infrastructure and delivery of our contracted customer commitments.

We recorded an income tax provision despite a net loss due to valuation allowance on net deferred tax assets. As noted last quarter, absent significant discrete items or a change in circumstances, our tax rate should remain broadly consistent over 2026. Adjusted net loss for Q2 was $567 million, compared to a net loss of $130 million in Q2 of 2025.

Turning to capital expenditures. CapEx in Q2 totaled $9.4 billion, slightly above the high end of our guided range. Higher CapEx in the quarter reflects customer deliveries accelerating. Construction in progress, CIP, increased to $11.9 billion from $9.6 billion quarter-over-quarter, signaling the significant amount of PP&E we expect to deploy early in Q3 based on the large amount of power we received very late in Q2.

In fact, in June, we brought on more than 300 megawatts of active power, which makes June itself larger than any full quarter in our history. As Mike noted, the global supply chain remains complex. We continue to navigate these challenges with operational discipline and leveraging our partner relationships, including new ones, like Solidigm, to strategically source required inputs.

Turning to our balance sheet and strong liquidity position. As of June 30th, we had more than $6.9 billion in cash, cash equivalents, restricted cash and marketable securities. In Q2, we made significant progress in strengthening our balance sheet and expanding the depth and breadth of our access to capital, raising approximately $18 billion across a combination of debt, convertibles and equity.

These transactions included several firsts, like our inaugural Eurobond, as well as our first ever delayed draw term loan backed by HPC infrastructure issued in the public markets. Our most recent financing, our second publicly syndicated term loan, marked another significant milestone as the first to include shorter duration customer contracts. The transaction, which priced during one of the most dislocated weeks for credit this year, was met with meaningful interest. Despite the environment at the time, we chose to complete the transaction at its full size. Our spreads have retraced since. Perhaps most importantly, the transaction demonstrated the credit markets growing conviction in long term value of NVIDIA infrastructure running on CoreWeave Cloud. This financing is significant as it unlocks our ability to serve critical part of the enterprise market at scale, while also allowing us to accelerate the ramp of our managed inference platform and grow our exposure to shorter dated contracts that typically come at a higher ASP and margins.

These transactions attracted broad and deep investor participation, highlighting the significant interest we continue to see in supporting CoreWeave's journey. The combination of these transactions brings us to over $32 billion of debt and equity capital secured to date. Over the past year, we have reduced our weighted average cost of debt by almost 300 basis points, representing approximately $1.1 billion of annualized interest saving based on our end of Q2 debt load.

Turning to guidance. As a result of continued strong execution, we now expect to end the year with more than 1.85 gigawatts of active power, up from our previous guidance of more than 1.7 gigawatts. In terms of how this flows through the second half, we expect Q3 revenue to be in the range of $3.45 billion to $3.6 billion. We expect Q3 adjusted operating income of $200 million to $260 million as margins continue to sequentially expand, reaching low teens in Q4.

Q3 interest expense is expected to be in the range of $860 million to $940 million, reflecting the growth in our debt balance to finance our accelerating deployments. We expect CapEx to be $11.5 billion to $13.5 billion based on the significant amount of new capacity we will be delivering to customers.

Moving on to full year, our disciplined execution and the momentum we are seeing across our customer base gives us confidence in raising our full year 2026 revenue guidance to $12.4 billion to $13.2 billion, and adjusted operating income to $960 million to $1.15 billion. As a result of our increased expectations around capacity to be delivered to customers this year, as well as some of our significant recent wins, we now expect 2026 CapEx in the range of $35 billion to $39 billion. Finally, we're also raising our expected end of year annualized run rate revenue to $18.5 billion to $19.5 billion.

The long term nature and attractive margins underpinning our contracted revenue backlog continue to provide us with clear visibility and we are confident in the targets we have put forward.

In closing, Q2 demonstrated the strength of the demand environment for CoreWeave's full technology stack and the discipline of our operating model. We strategically expanded our customer base to support the next wave of enterprise AI adoption at increasingly attractive margins. Customers are expanding their spend with CoreWeave to leverage the full depth of our AI native platform. We remain on track for our sequential margin expansion through the balance of the year and we have made significant additional progress on our capital structure, reducing our weighted average cost of capital, while securing the financing required to support our long term growth plan.

We look forward to seeing many of you at our annual developer conference, fully connected, in September, where you will hear from our leadership and customers alike in how our platform is accelerating AI in production. Thank you.

With that, we will open up for questions.

QUESTION AND ANSWER SECTION

Operator: We will now begin the question-and-answer session. Please limit yourself to one question and one follow-up. [Operator Instructions] Your first question comes from the line of Samik Chatterjee with JPMorgan. Your line is open. Please go ahead.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Q

Hey, guys. Thanks for taking the questions and congrats on a strong overall [ph] front year. (38:24) Maybe just a couple of topics. One, you did mention the renewal opportunity with shorter term contracts as some of the older contracts come off expiration to leverage sort of the pricing that we're seeing in the market. Can you just help us think through, as you engage in some of that discussion with customers, what you're finding in terms of typically customer intent in terms of contract period and how much of your installed base of equipment is maybe up for renewal over the next few years? If you can give us a sense of how to think about the magnitude of that opportunity? And I have a follow up. Thank you.

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

A

Thank you for the question. And excited to spend a little bit of time with you talking about what was a truly outstanding quarter for the company, across our infrastructure, across our software, across our solutions, across our sales and contracts with new clients and existing clients. Yeah. One of the most exciting components of what we are beginning to understand and what we believe the market is providing real insight into right now is that the older generations of infrastructure continue to have significant value for use cases within many of the consumers of AI. And we've talked about this literally for years now that the most bleeding edge solutions that are coming out of NVIDIA that we build into our cloud and deliver to our most demanding customers, that's really important for some of the most cutting edge use cases.

But within that environment, within that ecosystem, there are an enormous number of other use cases that can make use of older, more, later dated SKUs. And the fact that we have been able to go ahead and sell a GPU who's architecture was from 2020, in a contract that was fully priced out to 2029, really provides some insight into what the future is going to look like as this infrastructure comes off contract.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

Yeah. In terms of the capacity that's coming up for renewal, Samik, that is a very limited part of our fleet and the ASPs of the older generation remain higher or at levels that we've seen about a year ago. The second part of the piece that is very interesting in our business is as these fleets come off maturity, it allows us to have a great product in terms of managed inference to sell for our customers, which, as Mike noted, is a very fast evolving nature of our business, part of our business, which we expect to continue to grow rapidly and expect to have about $250 million of ARR by the end of the year.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Q

Okay. And for my follow-up, thanks for all the details there. But in terms of the follow-up, can you talk about the supply chain a bit. You're obviously navigating it pretty well to bring capacity online. But in terms of the agreement that you have now with Solidigm, for example, how are you looking at sort of the need to maybe do something more broad based across supply chain in terms of longer term agreements to assure yourselves [ph] support (41:57) supply as well, so that you can continue to sort of execute on the powered, on the capacity that you want to bring online? Thank you.

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

A

Yeah, it's great question. So look end of the day my job is to ensure that this company has the capacity to deliver the product that our clients require. And in order to do that, we need to aggressively manage a complicated supply chain. And that supply chain includes everything from land, power and [ph] shell (42:29) through GPUs and networking through memory, all of which is being challenged by the growth and expansion of artificial intelligence.

In order to do that we have built over the last several years really long standing, deep relationships with our ODMs, our OEMs, NVIDIA, the companies that supply us with memory, all of them. And what we've done is we've thought about what is necessary to ensure that we have access to the infrastructure and the components and the capital that we need in order to deliver our products at a acceptable price and quality to our clients. And it's one of the things that's just embedded in the DNA of CoreWeave. That's what we do. It's part of what we do every single day is nurture these relationships and ensure that we have access to everything that we need in order to deliver the product, which is NVIDIA infrastructure delivered through our cloud.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

One thing, Samik, to note here is the increase in the value of output of the CoreWeave Cloud has outpaced the value of the input increases that we are currently experiencing in the supply chain. And as a result of it, what you're seeing is margins expand as Mike noted in his comments around, the typical contribution margins that we saw last quarter were 5 to 10 percentage points higher than what we've observed in the recent quarters.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Q

Great. Thank you. Thanks for taking my questions.

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

Brad Zelnick
Analyst, Deutsche Bank Securities, Inc.

Q

Great. Thanks very much and congrats on the strong execution. My first question, I wanted to ask about your managed inference offering, which is off to a really strong start. What are your initial learnings and what are the factors that inform your thinking on allocating capacity to it going forward versus your traditional take or pay deals? And I have a follow-up to that as well. Thanks.

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

A

Yeah, thank you for the compliment. It really was a great quarter for us. We're very excited about it. Look, when we think about our offering, we really think about it holistically. And we have made enormous strides through the last several years to focus on building scale through these long-term take or pay contracts.

As we've hit hyperscale, we understand that we are going to need to broaden our offering to provide the products that our clients needs, that to deliver products that have higher margins, to provide the software solutions, to provide access to CPUs, all of the things that are necessary for our clients to be successful. And when we think about the lessons that we've learned as we've gone through this unbelievable and unique scaling of our managed inference product, which went from $1 million to $100 million inside of a single quarter, we really think about the fact that that is an incredible opportunity for us to offer the most bleeding edge compute that we have, but also a wonderful way for us to access and use contract -- GPUs that are coming off contract in a way to extract maximum value for the company over time.

So look, the market is very deep. We think that we have an embedded advantage because of our control over the silicon, and we think that we're going to be very successful in that market over time.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

And Brad, one thing to note here is we announced yesterday around our DDTL 5.5 closing, and that shows that the capital markets are extremely interested and supportive of CoreWeave's product in terms of underwriting shorter duration contracts, which is definitely a tailwind as we look at these markets to support our customer needs.

Brad Zelnick
Analyst, Deutsche Bank Securities, Inc.

Q

Thank you, Nitin and thank you, Mike. That actually leads to my next question. So on the 5% to 10% better margin that you're seeing on recent deals that you're signing, can you help unpack the drivers? How much is a function of shorter duration deals versus strong competitive differentiation or other factors? And what are you seeing more broadly just out there in the market as it relates to pricing? Thank you.

Michael Intrator A

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

Sure. It's a combination of a lot of things, and it's difficult to deconstruct it. The infrastructure that we deliver through the CoreWeave Cloud is more valuable to our customers than any other solution that they can encounter. The quality of the platform, the reliability of the infrastructure, the security, the TCO, all of those things contribute to customers coming back to us again and again and expanding their footprint within our cloud and infrastructure.

And so there is a piece of it which is they understand how much more valuable a given piece of infrastructure is delivered through us. The second piece of it is many of our clients are monetizing their products, and so they are more aggressive about coming in and willing to pay us higher margins because they need access to the compute that will allow them to be successful. This is a phenomenon that's occurring across the infrastructure space, but it's particularly occurring within our ecosystem. And it's very exciting to see as the premium product that we deliver is priced in a premium fashion by the consumers of this compute.

Brad Zelnick Q

Analyst, Deutsche Bank Securities, Inc.

Thank you.

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

Irvin Liu Q

Analyst, Evercore ISI

Hi, thank you. This is Irvin Liu on for Amit. I had one and a follow-up. So my first question is, it sounds like there's upward pressure to pricing across multiple vectors, including the higher value you provide to your customers, the pass through of higher component costs and the re-contracting opportunity coming up. So with that in mind, should we still think of kind of say $18 billion to $19 billion in ARR as kind of the exit target for 2027?

Nitin Agrawal A

Chief Financial Officer, CoreWeave, Inc.

Yeah. So we increased the exit ARR number that we provided in guidance to you folks right now at $18.5 billion to $19.5 billion for 2026. So that is baked in our guidance that we provided to you.

Irvin Liu Q

Analyst, Evercore ISI

Okay. Got it, got it. Thank you for that. And then for my follow-up, I think the regulatory backdrop for data centers appears to be increasingly difficult, there have been reports of local opposition to data centers. With this in mind, can you talk about your confidence level in deploying more than 3 gigawatts of active power by the end of next year, and kind of your roadmap to 8 gigawatts by the end of the decade?

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

Irvin, your question is very timely and very important for the entire AI space and the entire data center space. I guess, I'll start with we believe that the certain communities have moved forward with moratoriums and we feel like moratoriums are, they're not going to impact the demand for this infrastructure. They are going to impact where this infrastructure gets built. And so our approach to how you engage with the stakeholders is that you have to be extremely collaborative with the communities that ultimately host the infrastructure, and that's based on transparency. You have to work with the local governments, you've got to work with the utilities, you've got to work with the policymakers in order to allow yourself to ensure that what you're building fits into the communities that you're entering. Ultimately, at the end of the day, it is in our interest and it is in their interest for us to be good neighbors of their community.

A lot of that comes down to making sure that you're paying for grid upgrades, so that it doesn't fall or impact the rate base. You create an enormous number of construction jobs. There are long term jobs that are left, that survive within the data centers. The data centers that are being built contribute to the tax base. All of these things are incredibly important to how you enter into a community and how you engage that community as you're building this infrastructure that is so necessary in order to be able to provide America's AI leadership.

And so, when we talk through the numbers with you guys, we are basing our progress on where we are today and what we have guided here, none of those numbers will be impacted by the regulatory pushback as of today. We are comfortable with it. We continue to expand. We continue to engage, our data centers, are best in class, and we expect to be held to that as we continue to build our infrastructure across the globe.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.
A

And Irvin, just to give you some numbers in perspective here, if you look at our gigawatts contracted today, they're already at 4.2 gigawatts contracted. In addition, we have about 1.5 gigawatts of powered land, options to execute LOIs that we have executed. That puts you close to about 6 gigawatts already in terms of how we think about power and it's middle of 2026. So we remain well on track to execute against our stated goal of greater than 8 gigawatts of active power by end of 2030.

Irvin Liu
Analyst, Evercore ISI
Q

Got it. Thank you for the color.

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

Raimo Lenschow
Analyst, Barclays Capital, Inc.
Q

Thank you. Congrats from me as well. I just, I wanted to talk a little bit about the growing importance of inference for you guys and how does your fleet need to evolve? Because inference need [ph] to (53:23) a lot more CPU, a lot more storage. Can you do that in the existing data centers? Do they need to evolve? Can you speak to that as well to make sure we have the capacity there going forward?

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

A

Yeah. It's a great question, it's a question we've been talking about now for several quarters. We believe that when you're building infrastructure, you don't build infrastructure for training and you don't build infrastructure for inference. You build AI infrastructure. And when you build AI infrastructure, you need to ensure that you have all of the components to be able to serve the full AI loop, everything from training through inference as it cycles back and forth, as it moves through the iterations that are required in order to serve your clients and those companies that are consuming this. And so really the infrastructure that we built will move seamlessly into the ability to serve inference over time.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Q

Okay, perfect. And then one follow-up is like obviously with the news from Meta, yesterday, a lot of questions that we faced today was around doing AI in the edge, et cetra and then, all these concerns came up again. Like, can you talk about like how you see the market evolving between edge, smaller clouds, neoclouds and [ph] hyperscale? (54:47) Thank you.

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

A

Yeah, one of the things about CoreWeave that should never be underestimated is we sit at the epicenter of an incredible amount of information flow from across the entire industry, right. The hyperscalers use us, the labs use us. You've seen it enterprise begin to scale within our platform. The information flow that's coming back and feeding us the clues to how the world is going to look in the future has been incredibly powerful for us in terms of how we position ourselves and our compute to serve our clients.

Look, at the end of the day, we believe that there are workloads that are going to be served from the edge, and there are workloads that do not require the same level of latency, protection and we have built our cloud to be able to serve both of those constituents effectively. And we will continue to build in that fashion. We will be informed by our clients continuously whether they need a little bit more of edge, they need a little bit more of scale. That is not as latency sensitive. All of those things are being fed to us on a continual basis.

And so, yes, we do see workloads on the edge and yes, we do see workloads that don't require to be on the edge. And we are very, very comfortable that the scale of our infrastructure and the ability to move it back and forth is going to provide a competitive advantage for CoreWeave over time.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Q

Okay. Thank you.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

Raimo, to your point in terms of increased competition. Even with this increased competition, we are seeing demand, pricing and margin all expanding, which is a signal for the growth in the CoreWeave kind of product as well as our growth overall in an already massive TAM that exists.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Q

Yeah, exactly. Yeah. Okay. Makes sense. Thank you.

Operator: Your next question comes from the line of Michael Turrin with Wells Fargo Securities. Your line is open. Please go ahead.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

Hey, great. Thanks very much. Appreciate you taking the question. I realize there's likely some rounding here, but you added an impressive 500 megawatts of active power in the quarter. The revenue, if we're looking at the sequential adds, is fairly consistent with last quarter. We can hear all the commentary around the uplift that's coming. So maybe help us think through the linearity of capacity added and when that 300 megawatts added in June started to hit more of a steady state in terms of model contribution and also would be useful as the second part to hear any early market signals you're gathering on Vera Rubin monetization and what the uplift there could look like versus prior generations? Thank You.

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

Yeah. So as you mentioned, right, you saw Q3, we added about -- sorry, Q2 we added about 500 megawatts of power, 300 megawatts of that alone was in the month of June, which is higher than any amount of power that we've added in any [ph] historic (58:04) prior quarter for CoreWeave. So definitely that power was back end loaded in terms of Q2, which you will start seeing kind of come through in Q3 and Q4 in our business.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

[ph] Thank you. (58:20)

Nitin Agrawal
Chief Financial Officer, CoreWeave, Inc.

A

And then in terms of...

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

A

Yeah. So let me maybe I'll take a moment to speak to Vera Rubin. Vera Rubin is a generation that is seeing the margin expansion right from the start and so it's really exciting for us, the demand for the Vera Rubin platform is enormous and the pricing power that CoreWeave has been able to garner with its CoreWeave Cloud, delivering that infrastructure really bodes well. When we were talking about that, 5% to 10% margin step function that we're seeing, a lot of that is coming in in the Vera Rubin SKU. We're excited about where that's going to lead. We think that it's going to be a very, very successful SKU for CoreWeave and CoreWeave's clients.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

Thanks very much.

Operator: Your next question comes from the line of Brett Knoblauch with Cantor Fitzgerald. Your line is open. Please go ahead.

Brett Knoblauch
Analyst, Cantor Fitzgerald & Co.

Q

Hi, guys. Thank you for taking my question and congrats on the very strong quarter. Mike, I guess just kind of based on the prepared remarks, it looks like the price environment has never really been better for older generation and obviously newer generations, [ph] if you see here. (59:45) As you look at the GPU fleet that's maybe rolling off contracts, can you talk about the cadence of how you guys look to either recontract that or kind of put it on spot or in your inference products and how far in advance of the roll off of those contracts would you look to kind of make that decision?

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

A

Yeah, so look, it's a good question. It's one we're working through. Keep in mind that the environment for inference is incredibly dynamic and it is scaling so fast as we kind of struggle to keep up with the build out of new infrastructure, the flexibility that we are given because we have infrastructure coming off line, allows us to continue to scale the inference product as we're continuing to explore exactly how big, how extensive is the managed inference opportunity for us.

Some of the infrastructure that comes off line we go ahead and we place back into a term contract if we think the economics warrant putting it in. And the economics include both the term that we're able to garner, as we think about the long term stability of the company and the long term obligations that we need to support as we continue to build and scale the company. But we also do recognize that in the short term, there is an opportunity to sell on shorter term contracts and extract additional margin on this infrastructure as the world tries to catch up with what is a systemic disequilibrium that has really existed for several years now and will continue to exist for the foreseeable future.

Brett Knoblauch
Analyst, Cantor Fitzgerald & Co.

Q

I think it kind of leads to my next question. I guess, off the back of DDTL 5.5, where you were able to kind of get funding for shorter duration contracts combined with kind of this data center [indiscernible] (01:01:49) political atmosphere that's kind of taking off. It feels like you guys should be quite well positioned given you are the most scaled to realize the most price benefits. How does the success of DDTL 5.5 change the way you view on kind of target durations on a go forward basis? Is that an avenue you want to use more to maybe extract more margin in shorter duration contracts given the useful life is there?

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

A

Yeah. I mean, you're exactly right. The execution of the DDTL 5.5 puts CoreWeave in a position where we get to populate the curve in terms of what we think is the most profitable configuration for term leasing. And so we want to sell our compute on long term contracts. We also want to sell it on shorter term contracts to extract additional margin. And we have been really, really aggressive about doing that. We were the first ones to bring the DDTL 5.5 to market in order to be able to really plug into those.

Just one more really important part of the short term contracts that I think it's important for everyone to understand. When you're thinking about enterprise, enterprise tends to want to enter into contracts that are not five years in length. They tend to think in shorter cycles than that. And so by enabling the financing market to support the contracts in DDTL 5.5, we're able to go ahead and diversify our term, so that we're able to support additional contracts. It opens up a whole new markets for us. These clients want to buy compute for two years or three years, and that's not a market that was easily accessible to us until we were able to bring DDTL 5.5 to market. And now that market will accelerate meaningfully as we're able to offer compute to our customers on a timeframe that they are able to consume it, buy it and enter into contracts to purchase it from us.

Brett Knoblauch
Analyst, Cantor Fitzgerald & Co.

Q

Awesome. Well said it. Thank you, Mike.

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

A

Thanks.

Operator: This concludes our question-and-answer session. I will now turn the call back to Mike Intrator for closing remarks.

Michael Intrator

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

Before we sign off, I want to thank our team, customers and partners for their trust, hard work and commitment to CoreWeave. None of these achievements would have been possible without you. I'm incredibly proud of what we have accomplished together and for what comes next as we build the essential cloud for AI. Thank you all for joining today. We appreciate your support and we look forward to updating you on our progress in the quarters to come.

Operator: This concludes today's call. Thank you for attending. You may now disconnect.