CoreWeave Inc - Q3 2026 Earnings Call *(Proofread Transcript)*

Call Participants

EXECUTIVES

Operator

Mike Intrator
Chief Executive Officer

Nitin Agrawal
Chief Financial Officer

Presentation

Operator

Hello, everyone. Thank you for joining us and welcome to CoreWeave's Third 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.

Unidentified Executive

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

Mike Intrator, Chief Executive Officer

Good afternoon, everyone, and thank you for joining us. CoreWeave delivered an exceptional third quarter. Revenue is growing rapidly. Margins are rising faster than planned. Customer demand remains extremely strong, and our newest contracts are among the most profitable we have ever underwritten. Our original fleet of cloud infrastructure is recontracting at premiums, producing cash flow in excess of what we assumed when we made those early investments. And our funding is anchored by secured lending, customer prepayments, and operating cash flows. Together, this momentum leaves our business resilient and well positioned to navigate the current environment.

Now, the specifics. Revenue reached $3.7 billion, up 168% year-over-year and above the high end of our guidance. We added more than $1 billion quarter-over-quarter. Adjusted operating income reached $379 million, well above our $200 million to $260 million guidance, with margins moving into the low double digits a quarter earlier than planned. Revenue backlog grew to approximately $133 billion. And we grew active power to more than 1.7 gigawatts, doubling our active footprint since the start of the year.

Last quarter, we talked about the operating leverage that would emerge as our infrastructure reached scale, and we diversified our business across short and long-term contracts. This quarter, you can see that progress in our results and the guidance we've provided. We also talked about the conviction behind CoreWeave, that AI would become part of how organizations across the economy work, and that the people building that future would need a different kind of cloud. That conviction is becoming more tangible every quarter.

Our customers are using AI to advance biological research, support patient care, develop autonomous systems, and help people learn. Our AI cloud platform gives them the ability to pursue those ambitions and bring their useful applications into everyday life. Powering these promising ideas requires the tools to build and deploy AI on reliable compute at attractive economics. Our job is to provide that foundation and earn compelling returns on our capital.

Given the recent volatility in the capital markets, I want to explain how these fit together. The value we deliver to customers earns their long-term commitments to CoreWeave. Those commitments underpin our financing approach, and disciplined investment allows us to keep expanding the platform our customers depend on while compounding shareholder value. Underneath it all is one central principle. At every decision, we look for the optimal economic choice, how we fund, what we charge, and how fast we build.

Building an AI cloud at hyperscale requires substantial capital and disciplined execution. We understood that at CoreWeave's inception and designed our business model and financing strategy around the long-term customer contracts that capital would serve. This strategy allows us to evaluate every new source of capital, choosing to pursue the opportunities that optimize returns depending on market conditions. At the end of Q3, our weighted average cost of debt across our entire business was 7.9%, down 250 basis points year-over-year.

Looking ahead, we will not use the unsecured debt market to fund our business at current yields. Nor do we need to. Consistent with our historical approach, we expect to finance the majority of the remaining CapEx needed to deliver our backlog through delayed draw term loans matched to specific customer contracts. Through 2027, more than 75% of committed CapEx is tied to investment-grade and enterprise customers, with the remainder tied to the world's leading AI labs and AI-native companies.

The pricing of these facilities reflects the underlying customer commitments and our proven ability to deliver. They can finance anywhere from 70% to 100% of the associated CapEx. And since our IPO have priced somewhere from 225 to 550 basis points above SOFR with high-quality investment-grade customers at the tight end of that range and AI labs towards the wide end. Today, that translates to less than 10% at the high end based on current rates. These facilities allow us to draw capital as the equipment is purchased and delivered. The cash flow from those contracts pays interest and repays the principal over the contract term while generating additional cash flow for CoreWeave.

As planned, we have already paid down more than $4 billion of debt through this mechanism. This includes paying down nearly half of our first two delayed draw term loans, which are our most expensive facilities. For the facilities funding long-term commitments, by the time the customer contract ends, the debt is repaid, and the infrastructure is owned free and clear by CoreWeave.

That paid-off infrastructure is becoming a meaningful source of capital in its own right. More of it comes off contract every year and generates revenue without new investment. The equipment is bought, the data centers are built. That lowers our capital intensity as we grow.

We also have several other material sources of capital, including prepayments. We take an economic approach to prepayments, weighing the impact to pricing against the cost of financing. In the past, that has led us to the unsecured debt markets. In the current environment, prepayments are more attractive. This year, we have collected prepayments on 70% of our deals, with customers funding up to 100% of the contract value upfront.

In the next three months alone, we will collect approximately $9 billion in prepayments and deposits from signed customer contracts. This includes more than $3 billion that relate to recontracting existing cloud infrastructure, which we sold at meaningful premiums to their original price. We also expect EBITDA less interest, which excludes the benefit of prepayments, to more than double in 2027, increasing our ability to internally fund capital needs. Further, OEM financing supports additional deployments, particularly as an efficient means to fund capacity serving short-term opportunities. Taken together, we expect these sources to fund the significant majority of our capital needs through 2027, with only a limited amount expected to be raised at the parent company level.

Our spending discipline matters just as much as our access to capital. Beyond our contracted commitments, our success-based CapEx model allows us to adjust the pace of our investment based on customer demand for CoreWeave Cloud and supply constraints, including regulatory impacts and the capital markets. That is a risk

mitigant. We commit the large majority of our capital only once the contracts and infrastructure are in place. Nitin will address this in more detail shortly.

Now let me connect that financing model to the economics you are seeing in our results. As more contracted infrastructure enters service, revenue grows against investments we have already made, and each new deployment becomes a smaller share of a much larger business. That is the operating leverage we described last quarter. Adjusted operating margin was approximately 1% in Q1, 5% in Q2 and exceeded 10% in Q3.

The economics of new business are strengthening as well. In July, we increased pricing across SKUs by approximately 25%, followed by another 10% increase later in the quarter. Pricing has continued to move higher in Q4, more than offsetting cost increases. We are also allocating more capacity to short-term opportunities, where the pricing justifies it, selling capacity at more than $40 million per megawatt annually, while pricing for new infrastructure on four and five-year contracts is moving to $20 million to $25 million per megawatt.

Long-term commitments provide the visibility that supports our buildout. Short-term contracts let us respond to a strong pricing environment and earn premium returns. We weigh that trade-off contract by contract, assessing the full economics of an investment, the revenue it can earn, the cost of operating it, and the cost of financing it. Strong demand gives us choices in how we best create durable value.

We see the same strength in our prior generation fleet. In Q3, pricing for Amperes and Hoppers increased 30% to 50% quarter-over-quarter, and at prices above their original contracts. Different workloads require different combinations of performance and cost. Earlier generations continue to offer compelling economics, particularly when they are part of a reliable scaled cloud. We are signing multi-year commitments extending beyond the GPU's depreciable lives, driving margins even higher in their later years as price and operating costs remain fixed but depreciation dissipates.

More than 130,000 GPUs come off contract in 2027, as do another approximately 170,000 GPUs in 2028. We already have a robust pipeline of customers competing for this capacity. At current pricing levels, 2027 renewals would generate approaching $4 billion of incremental annualized revenue, with no additional CapEx required. This grows to nearly $10 billion of annualized incremental revenue when including 2028 renewals.

Furthermore, if data center construction across the sector were to slow, installed compute becomes scarcer, and the knock-on effect is that GPU pricing may stay higher for longer. In that scenario, being first matters. Our installed base comes off contract fully paid years ahead of most others in the market.

We're well positioned across a range of market conditions. If the industry builds quickly, we grow with it. If supply is delayed, our installed paid-off fleet becomes more valuable. We do not build our plan on that scenario. The plan rests on contracted demand, success-based CapEx, and the contracts we are already signing.

Last quarter, I described how training, inference, evaluation, and improvement are becoming a continuous AI development loop. An application in production generates information about where it works and where it falls short. Developers use that information to improve it, test the changes, and put a better application back into use. That process requires sophisticated software that helps people learn from each experiment and highly dependable cloud infrastructure throughout the life of the application. This is what CoreWeave does best. We deliver a unified platform that makes hundreds of thousands of GPUs work together as a high-performance, full-stack cloud integrating networking, cooling, orchestration, and observability. This is why customers choose CoreWeave for their most important workloads.

In September, we introduced CoreWeave Forge to bring our AI development capabilities together in one environment. We are making the products and services we developed alongside leading AI labs accessible to enterprise customers so they too can build leveraging their own expertise and data.

In Q3, 70% of the new cloud customers we added were already Forge customers. Wayve is one example of how our relationship expands as the work becomes more demanding. They began using our software and development tools to experiment and assess model performance, and now use CoreWeave Cloud to train their family of world models for autonomous driving.

The Chan Zuckerberg Initiative is using CoreWeave for computational biology research. Hippocratic AI is training and deploying agents to support patient communication and care navigation. These are efforts to understand disease more deeply and help people get the information and support they need through their care. In engineering, our teams work with organizations including Nissan and Aston Martin Aramco Formula 1 to shorten testing cycles and reduce the data required to reach target accuracy. The broader promise of physical AI is to help engineers develop better systems and bring useful capabilities into the real world sooner.

MasterClass uses Forge to evaluate conversations with its teaching agents and improve the experiences they run on CoreWeave Cloud. The opportunity is to make learning more responsive to the individual, with help available when a person needs it. These customers are doing the scientific, clinical, engineering, and educational work that combines human discovery with AI. CoreWeave gives them a platform on which they can pursue it at scale and speed. That is why both the platform and the infrastructure of our AI cloud matters.

The distance between a promising idea and a useful application depends in part on whether people can run the experiments, evaluate the results, and afford to serve the application reliably. More than half of our new customers this year came from industries outside AI labs and technology, and 10% of Fortune 500 companies are now CoreWeave customers. The number of organizations committing more than $100 million to CoreWeave Cloud increased more than 2.5x year-over-year. As customers find more valuable uses for AI, they need more from the platform. That broadens our opportunities and deepens our relationships.

We have 10x as much qualified demand as available capacity through the end of 2027, and demand for 2028 is already several times larger than our available capacity, even before we have brought it formally to market. We now have more than 1.7 gigawatts of active power. We continue to expect to reach more than 1.85 gigawatts of active power by year-end. We also do not see material risks to our 2027 deliveries based on ongoing construction at fully permitted sites.

This year, we made international expansion a strategic priority. Among our 4.6 gigawatts of contracted power, we now have more than 1.3 gigawatts outside of the US, spanning Canada, Europe, and Asia, built where our customers asked us to be present. In Q3, we entered new markets like Indonesia and India while growing our footprint in Denmark and Spain. Our global AI backbone connects our infrastructure across all regions, giving customers consistent performance, software, and reliability regardless of where their compute infrastructure is located. Our focus remains on delivering the capacity customers have committed to use and earning strong returns on every deployment.

Before I turn it over to Nitin, let me return to where I started. We built CoreWeave to be an enduring business. We are growing profitably on an operating basis and at scale. Our financing strategy will continue to link capital to committed customer contracts underpinned by attractive returns and drawn from a broad and diverse set of sources. Tying these together allows us to calibrate our pace of investment to customer demand, the regulatory landscape, and the capital markets.

We are delivering capacity ahead of plan, and the value of that capacity is increasing with our original fleet recontracting at a premium. We continue to work to earn the trust of the communities where we invest and the customers we serve. And our cloud platform combines leading price performance and reliability with the innovation, software, and development tools enterprises depend on to accelerate their AI strategies.

We have reached an inflection point with our business becoming increasingly durable while the fundamentals accelerate. We are confident in our ability to deliver for the customers, shareholders, and the communities we are partnering with for the long term.

With that, here's Nitin.

Nitin Agrawal, Chief Financial Officer

Thanks, Mike, and good afternoon, everyone. Mike shared the momentum in our business across revenue growth, margin growth, demand, and rising contribution margins, renewals on our original infrastructure, and our capital plan. I will put the numbers behind them.

Q3 demonstrated that our model is delivering the operating leverage we said it would as we scale. We grew active power to more than 1.7 gigawatts of critical IT load, increased revenue by more than $1 billion sequentially, and expanded adjusted operating margins into the double digits, 5 percentage points higher than Q2.

Three things drove the quarter. First, we continue to convert contracted power to active power and revenue-generating capacity at a rapid pace. Second, demand and pricing are strengthening, driven by the returns customers are realizing from their AI spend with CoreWeave. Our expanded operating footprint gives us the flexibility to capture more of that upside through premium-priced and short-term opportunities. Third, platform adoption continues to broaden and deepen with customers across industries using a wider range of our products and services.

Our AI products and services continue to scale as customers consolidate more of their AI spending with CoreWeave. Customers are moving from experimentation to production, and the returns they measure are driving our demand. Our managed inference platform launched earlier this year has already exceeded $250 million of

booked ARR. Growth of this business is constrained primarily by the capacity allocated to it. Taken together, these dynamics are increasing the returns on our product and capacity investments and reinforcing our confidence in the long-term economics of the business.

We recognize we are operating in a dynamic macro environment. We are not a growth-at-all-costs business. But today's environment is very profitable for companies that have demand for their differentiated technology and that can execute. Our success-based CapEx model ties investment to customer commitments. Each contract we enter is underwritten to cover its associated expenses and produce a profit.

In each underwriting, we account for the GPUs, the data centers, the power, and the corporate overhead behind the contract. That has not changed. So to reaffirm what we have said previously, we expect the contracted cash flow from our existing customer commitments to be sufficient to service and repay the debt we have outstanding today, as well as any additional debt we expect to raise to deliver these contracts. For the avoidance of doubt, that includes both asset-level secured debt and unsecured debt sitting at the parent level. And that is before counting renewals.

Now turning to our third quarter results. Revenue was $3.7 billion, up 168% year-over-year and above the high end of our guidance, reflecting strong execution and a favorable demand and pricing environment for CoreWeave's AI cloud platform and services. Revenue backlog ended the quarter at approximately $133 billion, up 138% year-over-year, reaching another record high. Of this backlog, more than 50% is attached to a contract where we expect customer delivery to commence by the end of this year.

Operating expenses were $3.5 billion, including stock-based compensation expense of $160 million. The increase primarily reflects the cost of scaling our active power and converting backlog into revenue, driving corresponding increases in our cost of revenue and technology and infrastructure spend. Sales and marketing expense continues to support our go-to-market strategy and diversification of our customer base and multi-product adoption. G&A increased primarily due to investments required to support our scale while continuing to moderate versus revenue growth.

Adjusted operating income was $379 million and adjusted operating margin was 10%, expanding sequentially from 5% in the second quarter and 1% in the first quarter of 2026. We achieved low-double-digit margins one quarter ahead of the plan we stated to start the year. Our margin outperformance reflects strong execution, a favorable demand and pricing environment, and the operating leverage that comes with scale.

We continue to see attractive opportunities across both long and short-term contracts. The scale and flexibility of our platform allow us to allocate available capacity strategically to higher-margin, short-term contracts when those opportunities arise. This portfolio approach contributed meaningfully to our strong Q3 results.

Adjusted EBITDA for Q3 was $2.3 billion compared to $838 million in Q3 of 2025, up 169% year-over-year and 49% sequentially. Our adjusted EBITDA margin was 62%, the highest it has been in the last five quarters. This strong performance translated into $2 billion of operating cash flow in Q3. As we scale, operating cash flow will continue to grow and become an increasingly important source of funding for our growth.

Net loss for Q3 was $709 million compared to a net loss of $110 million in Q3 of 2025. Interest expense for Q3 was $846 million compared to $310 million in Q3 of 2025. Adjusted operating income at $379 million nearly tripled from the second quarter, while interest expense rose about a third. That is the path to earnings.

The same driver runs through adjusted net loss, $529 million against $41 million a year ago. 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.

Turning to capital expenditures, CapEx totaled $14.5 billion in Q3 above the high end of our guided range, reflecting strong execution that allowed us to pull forward certain deliveries from Q4. That same execution drove construction in progress, CIP, down to $9.3 billion from $11.9 billion last quarter as we deployed significant infrastructure against power that became available late in Q2. Year-to-date, we've brought online more than 850 megawatts of power bringing our total live footprint to more than 1.7 gigawatts, doubling our active power footprint from year-end 2025 in just three quarters. This pace of deployment reflects the strength and repeatability of our execution model, powered by the full-stack solution we developed to accelerate the speed at which we ramp new capacity.

Turning to our balance sheet and strong liquidity position. As of September 30, we had $6.5 billion in cash, cash equivalents, restricted cash, and marketable securities. In addition, we have nearly $5 billion of undrawn capacity under committed DDTL facilities and significant customer prepayments and deposits arriving over the next three

months. Since our IPO last year, we have demonstrated continued access to the capital markets. This is a direct result of our consistent track record of execution as we convert our $133 billion of revenue backlog into cash flow.

We expect approximately $9 billion of customer prepayments and deposits in the next three months alone, and for that number to continue to grow in 2027. As Mike noted, more than $3 billion of these prepayments are related to existing cloud infrastructure where no additional CapEx is needed. We expect EBITDA less interest, which excludes the benefit from growing prepayments, to more than double in 2027.

This increase contributes to our growing ability to internally fund our capital needs. We expect asset-level debt to cover the significant majority of our remaining capital needs. We have two fundamentally different types of debt sitting on our balance sheet. Debt that sits at the asset level, directly financing our customer contracts backed by committed cash flow and cloud infrastructure. And debt that sits at the parent company that is serviced by parent-level cash flow. Over the last year, we used parent-level unsecured debt to help fund our growth as these instruments priced at levels that optimized our overall economics.

As Mike noted, at current yields, we do not intend to issue additional parent-level unsecured debt through 2027. By contrast, the asset-level secured debt market has remained stable, with our loans continuing to trade around par value. Lenders have direct visibility to the cash flow we are generating as we deliver on our committed customer contracts. This market remains robust, and its depth for CoreWeave has grown since our IPO.

In fact, we expect to launch our next DDTL in the coming weeks. This will be an investment-grade facility with multiple rating agencies indicating a high investment-grade rating. Based on our early discussions, we expect the facility to price at a spread that implies a 7% to 7.5% yield based on current SOFR while financing more than 90% of the associated CapEx. More than 75% of our remaining committed CapEx for 2027 not yet attached to an existing DDTL is tied to these types of investment-grade and enterprise customers. For the remainder of our backlog, we expect to leverage similarly structured DDTLs as customer prepayments and OEM financings serve as the other primary sources of capital to fund our growth.

As we look beyond our committed customer backlog, we see other avenues we can also utilize. CoreWeave Omni provides an additional lever to grow with significantly lower capital intensity. By extending the CoreWeave platform to third-party infrastructure, Omni allows us to monetize incremental customer demand without funding the purchase of the underlying assets ourselves. We have already signed our first customer and have received significant interest from a number of others who want to run their workloads on CoreWeave while maintaining direct ownership of their infrastructure. This is another tool in our toolkit, allowing us to balance our capital needs with customer demand based on the prevailing environment.

Turning to guidance. Our strong third quarter performance and continued execution give us confidence to raise our full-year 2026 revenue guidance to $13 billion to $13.3 billion and adjusted operating income guidance to $1.275 billion to $1.375 billion, which represents a 26% increase midpoint to midpoint. We continue to expect more than 1.85 gigawatts of active power at year-end. We expect CapEx of $37 billion to $39 billion, raising the low end of our prior range of $35 billion to $39 billion. This is as a result of our expectation to deploy infrastructure and deliver it to customers sooner than planned.

For the fourth quarter, we expect interest expense of $1.05 billion to $1.15 billion. Based on the capacity we expect to deploy and the contracted revenue embedded in our backlog, we are raising our ARR guidance, and we expect to exit 2026 with annualized run rate revenue of $20 billion to $21 billion.

Looking ahead, we expect adjusted net loss to reduce meaningfully in 2027 and to become adjusted net income positive in the first half of 2028 and for the full year 2028. Looking at 2027 specifically, we retain flexibility in how quickly we execute our growth plan. We will deploy capital only where we have robust demand, access to financing on attractive terms, and expected returns that meet our thresholds. As market conditions evolve, we can calibrate our future capital commitments accordingly. This does not impact our ability to deliver on our existing customer commitments or to meaningfully exceed our ARR target for 2027. This means we have levers to control our CapEx.

Should we choose to focus only on our committed backlog, we would expect CapEx for 2027 to remain in line with our 2026 CapEx investment while revenue and EBITDA grow significantly. To deliver this capacity, we already have nearly $5 billion available to draw under already raised DDTLs, and more than $9 billion of the required infrastructure is already on our balance sheet, sitting in construction in progress, awaiting placement into service. Our spend is further supported by the approximately $9 billion of prepayments and deposits we expect to collect in the next three months from signed customer contracts, growth in cash flows from EBITDA, and additional DDTL and asset-backed facilities we expect to raise, including the near-term transaction I mentioned earlier.

Accounting for these sources, we expect to raise less than $5 billion in incremental parent-level capital through 2027. In short, we feel confident in our ability to fund our business for growth even in the current macroeconomic conditions.

Our performance this quarter is the result of deliberate choices we made years ago about how to build and fund this business with attractive returns. We built differentiated technology purpose-built for AI. We developed the operating capabilities to bring highly performant cloud infrastructure online at speed and scale. And we designed a success-based CapEx model that ties investment to contracted demand with a diverse set of capital sources that ensure CoreWeave can finance its growth with discipline and at reasonable cost.

We take our role in the AI buildout extremely seriously. This is a capital-intensive business, and we are stewards of that capital. We deploy it against committed customer contracts at targeted contribution margins, profitably and at scale. We have the ability to accelerate or moderate that pace based on market conditions. That discipline will continue to guide how we grow. The opportunity ahead is extraordinary. AI is transforming how the world builds, discovers, and creates, and CoreWeave was built to power that transformation.

With that, we will open up for questions.

Question And Answer

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 Gabriela Borges with Goldman Sachs. Gabriela, your line is open. Please go ahead.

Gabriela Borges, Analyst, Goldman Sachs

Hi, good afternoon, and thanks for all the incremental detail that you shared here. Mike, I want to start with a high-level industry question, which is all of the supply-demand data points that you shared sound very attractive. One of the concerns that I hear from investors is there is a wall of new supply coming online in the industry in the next one to two years, and we don't have perfect visibility into the demand algorithm. So we could have a debate around when supply-demand normalizes, whether it's next year or five years from now. My question for you is, walk us through the implications for the CoreWeave model if and when the industry gets into a period of normalized or oversupply. Thank you.

Mike Intrator, Chief Executive Officer

Thank you for the question, and thank you for everyone who's come to the quarterly call. It was an extraordinary quarter for us, and we're absolutely excited to share the results. With regards to your question, the concept that there is infrastructure coming to market faster than demand has grown has been a prevailing view that has existed now for several years. And what we have seen and what I think you're seeing across the board from other providers of infrastructure, from other consumers of infrastructure, is that demand for compute is growing faster than the capacity to deliver that infrastructure. And that is being echoed by the expansion of the margins that we are seeing on each additional contract. And so yes, you're correct, there is a significant amount of infrastructure that is coming to market, but the proliferation of this technology, the diffusion of this technology across the economy is outpacing the ability of that infrastructure to be able to deliver the tokens, to be able to deliver the infrastructure and the clusters to be able to train and deliver infrastructure.

We are largely sold out through 2027. There is overwhelming demand for that, and we are also sold out through 2028. We're not sold out through 2028, but any infrastructure that we have placed on the market has been sold, and so we have significant demand that we could deliver that infrastructure into.

Nitin Agrawal, Chief Financial Officer

Yes. A couple of data on points, Gabriela, on that piece is, as we see the demand across the industry, what we are seeing, that demand is transforming into pricing power. Pricing for long-dated customer contracts is going up to $20 million to $25 million per megawatt and for four and five-year contracts. And for shorter duration contracts, it's up to â€“ as high as up to $40 million per megawatt. So we're seeing that demand translate into pricing power and our ability to extract margin for our premium product infrastructure and cloud products.

Gabriela Borges, Analyst, Goldman Sachs

Yes, that's all good. Yes, please.

Mike Intrator, Chief Executive Officer

The final point there is we could sell all of the capacity that we have to current buyers that are in the queue. We are choosing not to. We're choosing to be patient. We're choosing to bring the infrastructure to market it is, further along, because we believe that that will result in better economic output for the outstanding infrastructure that we have.

Gabriela Borges, Analyst, Goldman Sachs

That all makes sense. And then my follow-up for you is, I really like this 20 to 25 updated megawatt data point. My question is on the contracts that you've already signed at committed megawatt monetization prices. How do you manage the risk between when you lock in that number on the revenue monetization side versus the cost side of the equation, which can sometimes be ephemeral and/or move up in the period of time between when you contract the megawatts with the customer versus when you actually bring it online? How do you manage that disconnect in timing?

Nitin Agrawal, Chief Financial Officer

Yes. Thank you so much for the question. And you're right, the way we manage that is, our CapEx is fundamentally success-based. As we sign customer contracts, we also place the corresponding POs and lock in the pricing for those as we move ahead. The additional piece is when we are signing customer contracts, we've already identified the data centers in which those customer contracts will land. We very well understand the costing of those data centers and the cost associated with operating that infrastructure. So as we sign these customer contracts, we are basically making sure that the economics stay relatively locked through the time when we make the deployments for those happen.

Gabriela Borges, Analyst, Goldman Sachs

Thank you for the detail. I appreciate the updates.

Operator

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

Brad Zelnick, Analyst, Deutsche Bank

Great. Thank you so much, and a lot of great disclosure and a lot of comforting details. I wanted to ask, how, if at all, did some of the recent concerns in credit markets across the industry impact your approach to contracting new capacity beyond 2027? Perhaps, as we think around time to power, customer types, duration, and when you line up the funding.

Mike Intrator, Chief Executive Officer

Yes. Thank you for the question. And it's a great opportunity for us to clarify that we haven't brought 2028 to market because we are waiting for that to season further. With regards to how we think about power and capacity on an ongoing basis, we have a number of levers that we use to ensure that we are staying in sync with the demand signals that we have. And you can see that in the pipeline of power that we're bringing to market, the timing of that power relative to where we're providing guidance on what demand looks like. And so we're not getting out too far on the curve in terms of where we're locking in power because we're trying to coordinate our roadmap for capacity with the roadmap of the clients.

One of the things that you'll hear us speak about again and again is our focus on maintaining a broad, diversified client base, which allows us to have incredibly accurate information flow in terms of who's going to need compute when, how much compute they need across multiple industries, whether it's the AI labs, whether it's robotics, whether it's the S&P 500 clients. It's a really, really robust information flow that's coming across our desk as we think about how we're going to enter into power contracts on a go-forward basis.

Brad Zelnick, Analyst, Deutsche Bank

Thank you, that's helpful. And Mike, if I could just ask a follow-up for you, with some of the less established AI clouds going through various growing pains out there, does this potentially present opportunity for you ahead, either in terms of customer choice, consolidation, or pricing power that you have? Thanks again.

Mike Intrator, Chief Executive Officer

Yes. So I think that's also a great question. One of the things that we have done is we have gotten to scale, and with that, we have positioned ourselves very well to be able to take advantage of this market that's going through a degree of volatility. At the end of the day, an AI cloud is a combination of a software layer that allows people to effectively manage their infrastructure, as well as an infrastructure build in and of itself. And what we believe is that we have done a wonderful job, probably the best in the market, delivering the software layer that's required.

And as far as the infrastructure goes, that is a scale game. That is a game where you're going to need to get to material scale in order to be an effective player. We've known that since the beginning. We have built our business around that concept in terms of matching our long-term contracts with a diversified group of buyers as we sit here and try to continue to drive to scale. The end result is our software is the best solution and the hardware is available at scale. Both of those give us an opportunity to be very selective about which clients we want to work with, which clients we want to pull onto our infrastructure and partner with as they begin their AI journey or as they continue to scale it.

Brad Zelnick, Analyst, Deutsche Bank

Thank you very much.

Operator

Your next question comes from the line of Samik Chatterjee with J. P. Morgan. Samik, your line is open. Please go ahead.

Samik Chatterjee, Analyst, J.P. Morgan

Great. Thank you. Thanks for taking my questions. Michael, listen, strong margins in the quarter. You're guiding to also significant acceleration or improvement in the margins into the next quarter as well. I was just wondering if you can flesh out the role that the short-term contracts are playing in this margin upside that you're seeing.

You've historically talked about also bringing capacity earlier in the quarter, that tends to drive margin upside. But how should we think about maybe execution versus the contribution from shorter-term contracts in terms of the upside on margin, both in the third quarter and going into the fourth quarter guide? And I have a follow-up.

Nitin Agrawal, Chief Financial Officer

Thanks so much for that question. Our margin performance this quarter was a reflection of our strong execution, and both a favorable demand and pricing environment. As well as our ability to sell premium products to our end customers. All three of those vectors contributed to our overperformance this quarter. And we continue to expect those vectors to continue in our favor through the course of the near term.

Mike Intrator, Chief Executive Officer

One of the things that we think about a lot here is we really like the stability associated with the long-term contracts. We really like the margin additions that come with the short-term contracts. And we really like the ability to deliver compute through our services like Managed Inference.

And all three of those are growing very rapidly. And all three of those provide an incredible opportunity for us to build and scale our business. We think that our ability to provide additional volume into the short-term business, additional volume into our Managed Inference product, is a function of the scale of the business that we have built on the back of the long-term contracts. And so you'll continue to see us work through the process of how we're going to allocate this valuable infrastructure across the three main vectors of opportunity that are in front of us.

Samik Chatterjee, Analyst, J.P. Morgan

And Michael, a lot of concerns recently following the elections about what the regulatory environment or regulatory pressures in the US would mean. Any way to help us with how you're thinking about things, particularly as you look at your long-term plans for 2030, sort of 8 gigawatts, and what role maybe your international expansion plays to-- given this current macro backdrop that you have? Thank you.

Mike Intrator, Chief Executive Officer

So, to begin with, we do not anticipate any impact from the seismic regulatory shift that's occurred. We believe that based on the permitting status as well as the construction status of the data centers that our backlog is going to go

into, we are very comfortable with our pipeline. And so that's a high-level comment. We believe that what has occurred is a statement around how you build. We want to build responsibly. It's been a hallmark of how we've entered communities early, how we've worked with the communities that host our data centers, how we go about building and interfacing and being good neighbors. And so we think that that is an opportunity for us to continue to scale our business within the new political reality that exists.

And as you correctly identified, CoreWeave made it a priority to begin to scale our footprint globally. We have added an additional 1.3 gigawatts outside of the US. You saw us enter Indonesia. You saw us begin our voyage in India, in Spain, in Denmark, in Canada. All of these locations are an opportunity for us to build a diversified portfolio of infrastructure that is required by our clients, but also provides us the ability to build infrastructure in communities that are excited about us coming and expanding our footprint with them.

Samik Chatterjee, Analyst, J.P. Morgan

Great. Thank you. Thanks for taking my question.

Operator

Your next question comes from the line of Karl Keirstead with UBS. Karl, your line is open. Please go ahead.

Karl Keirstead, Analyst, UBS

Okay. Thanks very much, Mike. You made a comment about a desire to diversify your client base. I think as everybody on the call hears the frontier labs signal a desire to lease and stand up their own compute going forward, I think we'd all love a little bit of color on how CoreWeave is going to react to that and the pace at which you would like CoreWeave to allocate incremental capacity to non-lab, specifically enterprise customers. So perhaps you could articulate that planned mix shift and tie it maybe to the growth in the backlog that you saw in the third quarter and offer a little color on what the customer composition of that backlog growth looked like. Thanks so much.

Mike Intrator, Chief Executive Officer

Yeah, one of the things that you hear me talk about almost regardless of what part of the business we're talking-- what we're focused on is the concept of diversifying our customer concentration, our data center concentration, so forth and so on. Look, it's pretty incredible that CoreWeave gets to stand in front of you guys and say, hey, 10% of the Fortune 500 are CoreWeave customers. 10% of our backlog is from financial services. There is real diversification there. More than half of the new customers that we added this year are not within the AI space.

It's really been a focus of ours as we've understood right from the beginning that in order to build a cloud, you have to be able to open up and address the long tail of enterprise clients that are going to need this type of infrastructure, that are going to need the type of software that we can provide through Forge as they begin to build out the tools that they need in order to be able to build and grow their business. And so we see that we're going to have long-term relationships with the AI labs. We're going to have long-term relationships with the clouds, the hyperscale clouds. We're going to have long-term relationships with the social media companies, but we've also made it a priority across the company to really begin to build out and diversify those types of companies that are running on our infrastructure.

Nitin Agrawal, Chief Financial Officer

And Karl, just to add, Forge is a great vehicle for us to make entry into the enterprise customer base. And we are already seeing great results come out of that. In Q3, 70% of our new cloud customers that we added were already Forge customers. Forge is how enterprise customers begin to deploy infrastructure. And that's a great opportunity for us to make headways into the enterprise market.

Karl Keirstead, Analyst, UBS

Okay. Thank you both. I'll leave it at one question, give some room to others. Thank you.

Operator

Your next question comes from the line of Ben Reitzes with Melius Research. Ben, your line is open. Please go ahead.

Ben Reitzes, Analyst, Melius Research

Hi, Mike and team, thank you. Nice to be speaking with you. Hey, Mike, you gave a lot of detail. I do know that you guys reached out to Texas, the new governor-elect there, and was just wondering if you could clarify how that went and also what your exposure is there. I know you mentioned that you don't see any delays. I was wondering what gave you that conviction. Maybe after that, reach out, just a little more specific and what your exposure is to that state. I have a quick follow-up.

Mike Intrator, Chief Executive Officer

Sure. So within Texas specifically, the overwhelming majority of our power is either online or fully permitted and under construction. And that is where our statement around, that we do not anticipate impact to the data centers that will ultimately host the backlog that we've brought to market.

Having said that, we also believe that it's important to establish relationships with new governors that come in or new politicians, new regulators that are going to be in charge of or in oversight of their jurisdictions, just like we had with politicians and governors and other regulators that sit over infrastructure historically. And so, my letter to her was to reach out and say, hey, we are nonpartisan. We are looking to continue to build our infrastructure. We are a good operator, and we are ready to work with you and the communities within which host our facilities, and we're excited about that.

Ben Reitzes, Analyst, Melius Research

Okay. Thanks, Mike. And then just, I may have missed it, but just a little more color on whether you've received your Vera Rubins and ramping them, and maybe even the Vera CPU racks, and how important that was for the raise guidance. In the near term and even through your ARR next year? Thanks.

Mike Intrator, Chief Executive Officer

Yes, so I'll start it out here. CoreWeave was very proud to be the first company, first cloud to bring the Vera Rubin online and to deliver it to Cognition, which was the first client in the world to have Vera Rubin in production. That was really exciting. It was really exciting for Cognition. It was really exciting, I think for the space at large. It's the introduction of a new incredibly powerful architecture that will be one of the defining components of the next leg of artificial intelligence.

We're also very excited about deploying the Vera CPU independently. We think that it's got a tremendous amount of opportunity to serve parts of the AI infrastructure space that is really exciting and growing incredibly rapidly. And so we're there. It's a big part of what we're going to be building over the next several years. I think it's a big part of what everyone is going to be building over the next several years. And it was our privilege to be the first ones to actually bring it online and deliver it into production.

Nitin Agrawal, Chief Financial Officer

And Ben, to answer your second question around the contribution around what drove the performance this quarter as well as our raised guidance for next quarter, and as I mentioned, three things drove our quarter. First, we continued to convert contracted power to active power and revenue-generating capacity at a very rapid pace. Second, the demand and the pricing environments are strengthening, and that is helping us realize better returns on our investments. And third is that our platform adoption is continuing and broadening across the domain with our customers across different industries using a wider range of our products and services. All three vectors contributed to this quarter as well as expected to contribute further as we raised our guidance for Q4 as well.

Ben Reitzes, Analyst, Melius Research

Thanks, Nitin.

Operator

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

Amit Daryanani, Analyst, Evercore ISI

Perfect. Thanks a lot, and thanks for all the details as well. I have two hopefully quick ones. The first is, you folks cited $40 million of annual revenue per megawatt on short-term contracts versus $20 million to $25 million on the longer term. On a comparable hardware utilization basis, how do the returns differ after financing costs? And then Mike, maybe you can just talk about what portion of capacity are you comfortable having on shorter-term contracts?

Nitin Agrawal, Chief Financial Officer

Yes. So as Mike mentioned, look, our business is built on the three fundamental elements of the predictable cash flows for long-term committed customer contracts, the premium products in terms of the higher-stack services like managed inference, as well as these short-dated customer contracts. When you look at our business, we manage that effectively in terms of looking at each opportunity, understanding the returns it will generate, the cost of financing that infrastructure, and the demand environment for each one of those vectors. That is how we internally allocate capacity.

When we think about the returns, of course, when you're looking for shorter-duration contracts, they pay back CapEx in a much faster duration. Somewhere in the tune of 6 to 18 months. And when you look at the longer-dated customer contracts, they pay back CapEx somewhere in the domain of hardware CapEx, somewhere in the domain of one to two years at the current pricing levels that we are seeing for long-dated customer contracts.

Mike Intrator, Chief Executive Officer

The other piece of it that is actually really interesting is that, with the initial contracts coming to a conclusion, like we said earlier, 130,000 GPUs are coming off contract in 2027, and then an additional 170,000 are coming off contract in 2028. Those GPUs are going to make exceptionally great opportunities for us to be able to monetize into short-term contracts. I'd also like to highlight all of that infrastructure has been fully paid off at the end of their initial contract. That was the construction of the delayed draw loans. And so they are fully amortized and gives us the opportunity to really have a meaningful impact on cash flow.

Amit Daryanani, Analyst, Evercore ISI

Absolutely. And maybe just on that topic, I think you said initially it's $4 billion of incremental revenues from renewed contracts in '27. $10 billion in '28. Is that just the total revenue contribution or is that the uplift you expect? And is this thing you've already renewed or you expect to do? And just any color on that, because those numbers seem very powerful, and you certainly have a lot of EBITDA margins with those, I think.

Nitin Agrawal, Chief Financial Officer

Yes. Those are outside our revenue backlog that you see because that is for contracted capacity today. As we recontract this capacity, that is the available opportunity at the current pricing levels for the ARR to be generated against that infrastructure. Which, as you highlighted, is deployed infrastructure for which we do not have incremental CapEx.

Mike Intrator, Chief Executive Officer

For clarity, I want to make sure that it's understood that in our $9 billion prepayment deposit, there is in excess of $3 billion that was prepayment for existing infrastructure that has been fully paid off. And so that's $3 billion that comes into the company that doesn't require any debt service whatsoever. And it's for infrastructure that has already run through a full cycle within the DDTLs in the initial contracts that makes it so powerful.

Amit Daryanani, Analyst, Evercore ISI

Perfect. Thank you very much.

Operator

This concludes the Q&A session. I will now turn the call back to Mike Intrator for closing remarks.

Mike Intrator, Chief Executive Officer

Before we sign off, I want to thank our team, customers, and partners for their trust in CoreWeave and their commitment to what we are building together. None of these achievements would have been possible without you. I am incredibly proud of what we have accomplished 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.