11 - 12 - 2026

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

CORPORATE SPEAKERS:

Chance Moreland

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

Corvex, Inc.; Chief Financial Officer
Jay Crystal
Corvex, Inc.; Co-Founder & Co-Chief Executive Officer
Seth Demsey
Corvex, Inc.; Co-Founder & Co-Chief Executive Officer

PARTICIPANTS:

Bill Papanastasiou
Chardan Capital; Analyst
Kingsley Crane
Canaccord Genuity; Analyst
Fedor Shabalin
B. Riley Securities; Analyst

PRESENTATION:

Operator^ Hello, everyone. Thank you for joining us, and welcome to the Corvex Third Quarter 2026 Earnings Conference Call.

After today's prepared remarks, we will host a question-and-answer session. (Operator Instructions).

I will now hand the conference over to Chance Moreland, Chief Financial Officer. Chance, please go ahead.

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

Chance Moreland^ Good morning, and welcome to Corvex's Third Quarter 2026 Earnings Conference Call. I'm Chance Moreland, Chief Financial Officer.

Joining me today are Jay Crystal, Co-Founder and Co-Chief Executive Officer; and Seth Demsey, Co-Founder and Co-Chief Executive Officer.

Before we begin, I remind you that our remarks today include forward-looking statements. Actual results can differ materially from these statements. The factors that could cause this difference are in today's earnings press release and in our quarterly report on Form 10-Q for the quarter ended September 30, 2026, which we expect to file with the SEC later today. Our forward-looking statements use assumptions as of today. We do not undertake to update them.

On this call, we discuss both GAAP and non-GAAP financial measures. We use Adjusted EBITDA, and we refer to contracted annualized recurring revenue on live compute, which is an operating metric rather than a financial measure. A reconciliation of Adjusted EBITDA to net loss is in today's press release and in our Form 10-Q. Adjusted EBITDA is not a substitute for GAAP results. Other companies calculate similar measures in different ways.

The press release and a replay of this call are on our investor relations site at investors.corvex.ai. Two framing points.

First, Corvex completed its merger with Corvex Legacy Holdings on March 19, 2026, so prior year periods contain only the legacy healthcare business, and year-over-year comparisons are not meaningful. The third quarter is our first quarter with a clean sequential comparison to a full quarter of the AI cloud computing business, and that is the comparison I will use.

Second, new to this quarter, we will open the call to questions from our analysts after the prepared remarks. We expect to continue this format for subsequent earnings calls.

I will now turn the call over to Jay.

Jay Crystal^ Thank you, Chance. Good morning and thank you for joining us.

90 days ago, I told you what we do, why it works, and how you should measure us. On our second earnings call as Corvex, I'd like to do those same three things again, this time focusing on our execution and growth. I'll start with the three numbers I asked you to measure us on, then walk through what we said we would do and what we did, why customers are choosing us, and what we are focused on next.

Then, Seth will take you deeper into the technology, including Corvex Token Factory, and Chance will take you through the numbers.

In August, I asked you to measure us on three things: contracted power, contracted annualized recurring revenue on live compute, and the capacity we bring into service. Here is where each one stands.

In terms of contracted power, all 7.7 megawatts of critical IT capacity in the plan we laid out in August is now under contract. 1.5 megawatts is live, and the remaining 6.2 megawatts across Midwest-1 and the Mid-Atlantic 1 expansion are on track to be ready for service in the fourth quarter of 2026.

In August, we also brought in Garrett Brams from Cushman & Wakefield to lead data center sourcing and development. His job is to find power on the scale and timelines that match customer demand, and he is off to a terrific start. Contracted annualized revenue on live compute stands at 21.6 million as of September 30 and 21.7 million as of today, in line with the approximately22 million shared on our August 14 call.

And a quick reminder, live ARR looks backwards. The growth in this business starts with our power funnel. We match power with customer demand, finance the project, and then deliver it. In terms of capacity brought into service, 128 NVIDIA Blackwell GPUs went live during the quarter at our Mid-Atlantic data center. These GPUs, delivered in late July, were the final tranche of the multi-year offtake on the NVIDIA Blackwell cluster we announced on August 4. We will report these numbers on the same basis every quarter.

Turning now to what we said we would do last quarter. On August 31, we laid out a specific plan. We said we would grow critical IT capacity from approximately 1.5 megawatts to approximately 8 megawatts by the end of 2026, with a right of first refusal on an additional 12.5 megawatts in our Midwest-1 site that could take total Corvex capacity past 20 megawatts in 2027.

We said both expansions sit inside existing enterprise data centers with live utility power. We also highlighted that we were looking for two kinds of customers for that capacity. One was a growing AI-native company with qualified interest in at least three years of offtake, and our goal was to use the deployment as an opportunity to establish what could become a very scalable relationship.

The other was a counterparty with the credit quality and the appetite to take not just the first phase of our Midwest-1 data center, but the entire data center. We also noted that Corvex Token Factory would move out of [closed alpha] in the third quarter, and we said we would keep underwriting every deployment before we sign it and fund growth in ways that limit dilution.

Now let's turn to what we did since our last call.

We've made demonstrable progress securing offtake with the targeted customer profiles and terms we described a quarter ago. For Mid-Atlantic 1, we are in final discussions with a rapidly growing AI native, one we can grow alongside to take over 1,000 NVIDIA Blackwells for three years on terms that clear our hurdle rate.

For Midwest-1, we are encouraged with how offtake discussions are shaping up relative to the three tests we set out for this site -- credit quality, contracting the full data center in an upside case, and the growth potential of the overall partnership. We look forward to announcing news when these agreements are complete.

On capacity, Midwest-1 is on track to be ready for service in the fourth quarter of 2026. I want to be precise about this term, because we'll use it every quarter. Ready for service means the facility is powered, cooled, and otherwise ready for us to begin operating the computing cluster. It is not the date when the customer compute goes live and revenue begins, which follows hardware delivery, installation, and commissioning after the site is ready for service.

Securing live power, certain data center electrical and cooling equipment, and permits are the long poles for most of this industry. In Midwest-1, they're now largely behind us. From the start of construction to anticipated ready for service, we expect to deliver in under two quarters. The Mid-Atlantic expansion is also substantially complete, with final equipment installation and commissioning remaining. This is on track to be ready for service in the fourth quarter as well.

On Token Factory, we said the third quarter, and we launched it publicly on that timeframe. Developers and enterprises can register at no cost, get an API key, and run open-weight models with zero data retention by default. I'm excited for Seth to tell you more in a moment about why it matters and why we're excited to provide open-weight inference on our platform.

Now let's turn to why customers choose Corvex. They select us on trust in three things: our technical and operational expertise, our ability to deliver power on the scale and timeframe they require, and our ability to structure and finance complex projects.

On the first, we're engineers at every level of the stack, with decades of experience building and operating large-scale distributed systems. That depth is how we are able to enable customers to succeed, even when it relates to issues outside of our platform. For example, we had a customer with a workload that ran 30% slower than expected. We took a look for them, profiled and optimized that workload, and two days later, we returned it to them running six times faster.

Regarding power, we're fortunate to work with exceptional partners, and we're developing great talent internally. Collectively, our team and partners are enabling Corvex to secure attractive data centers, and we're on track to rapidly deliver that capacity to the market. Our power pipeline continues to grow.

On economics, nothing has changed in terms of how we think about them, so I will be brief. Our AI factory revenue is contracted, fixed-term, and take-or-pay. Customers pay for reserved capacity regardless of utilization. Every project clears a return threshold before we sign. Power cost, hardware cost, financing cost, contract term, and residual value all go into the model, and so does the counterparty credit profile.

Every project is financed on its own, customer prepayments first, asset-level debt secured by the contract and equipment second, and equity only where it earns its place. I said in August we would rather report a smaller contracted number than a larger one that cannot earn back its cost of capital. I remain very encouraged with our ability to develop deals with attractive customers, economics, and financing terms.

Now let me give you the three priorities for the next 90 days.

One, bring Midwest-1 and Mid-Atlantic expansion to ready for service on schedule while completing offtake for these sites. Two, convert more power we have access to into contracted capacity and keep growing our footprint. Three, grow demand for the Corvex Token Factory through partners and channels.

With that, let me hand it to Seth to take you inside the technology. Seth?

Seth Demsey^ Thank you, Jay. And good morning, everyone.

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

I want to spend a few minutes on three things: what Corvex Token Factory is and why we built it, what we're doing on security, and what the engineering team delivered this quarter.

Let me start with Token Factory. We launched Corvex Token Factory on September 30th. The simplest way to describe it is, it's a way to buy intelligence by the token instead of by the GPU.

A developer or an enterprise sends a request to the standard API, we serve it with an open-weight hosted model on our infrastructure, and they pay for the tokens that they use. No cluster to provision, no hardware to manage, and by default, we retain nothing, not the prompt, not the response. We built it because of a shift in how we see AI being consumed in production.

Frontier closed-weight models are remarkable, but they're expensive, and not every use case needs a frontier model. For a large and rapidly growing share of real production work, open-weight models now deliver results that are capable of powering real-world workflows and actions. In fact, we run our own engineering this way.

A small share of requests go to frontier models for planning and the hardest problems, and the large majority of routine requests run on open-weight models powered through Corvex Token Factory. Every workload that makes that shift needs somewhere to run. This is the demand we're building for.

And for Corvex specifically, it changes the economics of our infrastructure. Capacity sold by the token can earn more per GPU hour than the same GPU under a bare metal contract when there is demand. Pricing follows from the model. We run open-weight models on our own infrastructure and price at a material discount to frontier labs, while still pricing at a significant margin above the cost of serving, rather than charging a premium for excessive access to a closed model.

And switching is deliberately easy. Token Factory is compatible with both the OpenAI and Anthropic API conventions, so a developer changes an endpoint, a key, and a model name, and their existing code and tools, Claude Code, Codex, Aider, keep working. It's still early.

We're weeks past launch, and the first customers are telling us why it works for them. One, the chief executive of an engineering firm told us speeds and availability made Token Factory reliable enough to become an execution layer in their workflows, and that knowing their code is not used for training was non-negotiable.

Our focus over the next two quarters is demand, onboarding partners and channels first, then direct enterprise sales, and continuing to ship against the roadmap. We're not commenting on Token Factory revenue yet, and we expect a meaningful contribution to begin in 2027.

Turning now to security and our Assured AI solution. The dirty secret of AI is, at best, data is encrypted at rest and in transit and then decrypted before it is processed inside the CPU, the GPU, and stored in memory. Model weights, fine-tuning data, and prompts are all exposed in those moments. Every provider asks you to trust their policies and their peoples.

We're building Assured AI so that customers do not have to. Assured AI uses trusted execution environments to prove cryptographically that no one, including Corvex, has access to a protected model or the data it processed. In March, we were among the first companies to complete a verified production

deployment of confidential computing on NVIDIA HGX B200 systems with CPU and GPU attestation through Intel Trust Authority, NVIDIA NRAS, and encrypted NVLink traffic between GPUs.

We have been running protected workloads for certain customers for more than a year. We are integrated in Assured AI into Token Factory so that lower-cost inference and verifiable privacy come together in a single platform. Today's standard Token Factory endpoints run with zero data retention, and the confidential computing layer is what makes the operator access not just out of policy, but impossible.

The customers who care most, regulated enterprises, financial firms, and public sector organizations, are also the customers with the strongest and most durable relationships once you are in. We expect to have proof points to share over the coming quarters.

Turning now to the three items the engineering team delivered this quarter. First, on Blackwell, we completed delivery of a multi-year cluster announced on August 4th, including the liquid-cooled HGX B200 capacity. That cluster has been fully in production since July 29th.

Second, on Midwest-1, our real estate team under Garrett Brams is converting an existing enterprise data center into a liquid-cooled, AI-ready facility, on track to be ready for service in the fourth quarter of 2026.

Third, on our platform software, the next version of our cloud management software, which we said in August had moved from planning into execution, remains in execution and on track, with the first end-to-end stack now being deployed in our lab.

It will serve as the basis of all of our deployments moving forward and the base on which we continue to develop our Assured AI security offerings for both cloud and inference customers.

Finally, in the third quarter, we received notices of allowance for two patents covering Corvex's technology, bringing our portfolio to two allowed patents and 10 pending applications.

With that, I turn it over to Chance. Chance?

Chance Moreland^ Thank you, Seth.

As I said at the top, the useful comparison this quarter is sequential, so I will compare the third quarter to the second quarter throughout.

Total revenue for the third quarter was $5 million, compared with $3.8 million in the second quarter. All of it came from AI platform and services.

For the nine months, total revenue was 9.3 million. The sequential increase reflects a full quarter of the capacity that came online during the second quarter and partial quarter contribution from the final tranche of the Blackwell cluster, which came online in late July. That timing is also why third-quarter revenue is below one quarter of the 22 million contracted annualized revenue figure we gave you in August. A portion of that capacity was live for only a part of the period. Revenue was in line with what we indicated in August.

Contracted annualized recurring revenue, or contracted ARR, on live compute was 21.6 million as of September 30, and is 21.7 million as of today. The definition is unchanged. The annualized value of fixed contractual fees on capacity that has been delivered, accepted by the customer, and is generating revenue as of the stated date. It excludes contracted capacity that is not yet live. It is a point-in-time operating metric, not a forecast, and not a GAAP measure, and it will not equal annualized GAAP revenue in any quarter where capacity comes online mid-period. Going forward, we will give you both the quarter-end figure and the as-of-today figure.

Turning now to expenses. Cost of revenue for AI platform and services, excluding depreciation and amortization, was 2.5 million, compared with 2.1 million in the second quarter. Depreciation and amortization was 2.9 million, compared with 2.7 million in the prior quarter.

Technology and infrastructure expense was 1.3 million, compared with 1.4 million. Sales and marketing was 0.8 million, compared with 0.7 million. And general and administrative was 14.8 million, compared with 12.1 million. As I explained last quarter, G&A is dominated by non-cash stock-based compensation from the replacement equity awards issued at the merger. Total stock-based compensation in the quarter was 11.3 million, of which 9.3 million sat in G&A.

The remaining expense from those awards, approximately 143.6 million as of the merger, is recognized through 2030, so you should continue to expect a large non-cash and reasonably predictable charge each quarter. The awards themselves are detailed in Note 12 of the 10-Q, and I would point you to the share count rather than the income statement to understand their economic effect.

Loss from operations was 17.4 million. Interest expense was 0.3 million, which is new this quarter and reflects the equipment financing I will describe in a moment, including amortization of financing costs. Net loss for the third quarter was 17.5 million compared with 12.8 million in the second quarter. Net loss attributable to common stockholders was 17.5 million.

Adjusted EBITDA was negative 3.2 million for the quarter, flat with the second quarter. The definition is unchanged, and the reconciliation is in the press release and will be in the 10-Q.

Now turning to the balance sheet and cash flow. We ended the quarter with $42 million of cash and cash equivalent, compared with $21.7 million at June 30, plus $0.7 million of restricted cash. The increase reflects the approximately $31 million of net proceeds from our private placement we closed in early September, and $6.9 million of net proceeds from equipment financing, offset by $9.8 million of capital expenditures and $7 million of operating cash use.

Property and equipment net was 39.1 million compared with 31.4 million. We ended the quarter with $6.9 million of debt compared with none at June 30. This is the equipment financing for the Blackwell expansion we announced on August 4th, held at the project level, secured by the equipment and the related customer contract, and is non-recourse to Corvex Inc.

Finance lease liabilities were 8.5 million, and operating lease liabilities were 4.9 million. Deferred revenue was 3.2 million, 1.9 million current and 1.3 million non-current, compared with 3.7 million at June 30. The decrease reflects recognition of previously received customer prepayments as the capacity under those agreements generates revenue.

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

Cash used in operating activities was $7 million in the quarter compared with $5.3 million in the second quarter. This includes approximately $1.8 million of payments to wind down our legacy connected devices business, approximately $2.5 million in rent prepayments, and $1.2 million of prepaid software. This was offset by a $2.8 million vendor refund collected in July. Capital expenditures were 9.8 million, including capitalized software of 0.2 million.

Turning now to our capital plan. I said in August that scaling this business will require additional capital and that our funding requirements depend on the pace and structure of expansion. The principle is one we described in August. Each project is underwritten to clear its cost of capital before we sign, and we fund it as far as possible with customer prepayments and asset-level debt before we consider equity.

Turning now to our capital structure. As of November 10, 2026, we had approximately 39.1 million shares of common stock outstanding and approximately 21,800 shares of Series D preferred stock outstanding, convertible into approximately 21.8 million common shares, approximately 60.9 million shares on an as-converted basis. That compares with approximately 56.6 million as converted at July 8, and the change reflects the September private placement of approximately 4.3 million as-converted shares at $7.75 per share.

On the at-the-market facility we announced on October 14, we have not used it to date. We put it in place to give us fast, opportunistic access to equity in measured amounts, primarily to fund the equity portion of specific deployments with attractive returns rather than general operating needs. We will be disciplined about when and how much we use it.

On internal controls, the material weakness disclosed in our 2025 Form 10-K remains unremediated as of this quarter's filing. In August, I told you our plan, engage an advisory firm, complete risk assessment and scoping during the third quarter, then move to control design and documentation, then testing.

We completed the risk assessment and scoping in the third quarter. We are now designing and documenting controls across entity level, business process, and IT general controls, starting with order to cash and building those controls into our new procure-to-pay and ERP systems as they go live.

We expect to begin testing operating effectiveness in the first half of 2027. Remediation will not be complete until the controls have operated for a sufficient period and have been tested. We will keep updating you each quarter until this is closed.

With that, I will hand it back to Jay.

Jay Crystal^ Thanks, Chance.

Let me close the way I closed in August, because the point has not changed while the evidence has. The bottleneck in AI infrastructure is not demand, and it's not chips. It's the ability to put a finished, energized data center in the same place and time as hardware, software, and a creditworthy customer, and to deliver this union reliably and on the customer's timeframe.

In August, that was the description of the job. Today, it's a description of the quarter. 6.1 megawatts of critical IT capacity scheduled to come into service, advanced discussions with partners to take that capacity, a product launched on schedule, and funding structured around prepayments and asset-level debt before equity.

What is still in front of us has not changed either. Our capacity needs to grow by multiples. Our revenue is still concentrated in a small number of customers.

Token Factory is launched, and we're focused on scaling this product. Our controls work is underway but not finished. None of that is hidden, and none of it is news to us. We remain excited, though, to tackle these challenges and increase shareholder value by doing so.

With that, Operator, we're ready to take questions.

QUESTION & ANSWER:

Operator^ We will now begin the question-and-answer session. (Operator Instructions).

Your first question comes from the line of Bill Papanastasiou with Chardan. Your line is open. Please go ahead.

Bill Papanastasiou^ Hey, good morning. Thank you for taking my questions.

For the first one, I was just hoping you could share color on the financing environment and whether you have recognized any notable changes from the prior quarter.

Furthermore, I'm not sure -- I didn't catch it. Are you able to share what the cost of capital was for the USD.AI loan, and do you anticipate tapping USD.AI going forward for future equipment financing? Thank you.

Chance Moreland^ Sure, Bill. Yes, it's Chance. Thanks for the question.

On the first question about funding environment in general, I would say we remain in active discussions with our financing partners and are really encouraged about the continued trends that we share. I would say the interest from their side remains as robust as we saw earlier in the summer. We're all looking at the interest rate environment and adapting as broader trends roll through the industry. But confidence in the space, interest in the asset class, and confidence in Corvex, frankly, remains as strong as it has been in a long time.

As it relates to the USD.AI loan that we announced in August, yes, I can announce it was a 10% interest rate on that facility. And the USD team, they have been great partners, and yes, we hope to continue to do business with them as we go forward.

Bill Papanastasiou^ Great. And just as a follow-up there, with respect to the economics of some of the offtakes that you're signing for GPU capacity, are you seeing any improvement in terms of prepayment or the appetite to increase prepayments, and just general GPU pricing? Hope to get some more color there. Thanks.

Jay Crystal^ Sure, Bill. This is Jay.

In terms of prepayments, yes, we see customers competing more aggressively for capacity, in part as data center capacity becomes more scarce, and that is an area where certain aspects of the financing

market are showing up. It's raising the tenant credit quality that is required to get into sites. But Corvex has thus far successfully been able to navigate around that and secure power in certain cases where others have not been able to.

So, with respect to customer prepayments, yes, we're seeing the percentages push up relative to where we've seen them in the past. And then with respect to pricing, SemiAnalysis continues to be a relatively good guide for industry benchmarks. As capacity becomes more scarce and demand density builds for clusters as they come online, that enables us to position the cluster to effectively achieve or exceed our hurdle rates as we bring these forward. So, we remain encouraged about our ability to bring financed transactions that exceed our hurdle rate to market.

Bill Papanastasiou^ Thanks. And then just the last one for me. Jay, you just mentioned that data center capacity is harder to come by. I was hoping you could share a little bit more on a boots-on-the-ground perspective. Are you seeing more competition emerge as a result of regulatory headwinds impacting players that are going after large-scale capacity and need to essentially wind it down to subscale or smaller-scale data center deployments? I'm curious to hear your perspective there. Thanks.

Jay Crystal^ Yes. It's always been a competitive market, but some of the recent events, especially with regard to Texas, have certainly led to increases in demand showing up in other places throughout the country for U.S.-focused offtake. Through the lens of Corvex, we've been ahead of these trends in terms of where we have chosen to grow and have secured rights to expand, for example, at our Midwest-1 site.

We are encouraged about our ability to navigate some of these because there are really sort of two trends that are occurring. While there is an increase in demand pressure, at the same time, especially within the past few weeks, landlords are looking for improvements and enhancements in tenant credit quality in terms of the offtake that we bring.

As a company, we've always been very focused on tenant credit quality and offtake through the partnerships that we form and the customers that we choose to bring. And so that is proving to be an asset for us as we navigate this environment, both in terms of financing for clusters, and then the financing that enables for landlords that choose to partner with us. So, through that lens, we're not discouraged by the environment, and in certain ways, it's becoming a competitive advantage for us.

But, yes, to answer your question, we are seeing increased levels of demand. It just needs to be filtered through the lens of tenant credit quality to compete and win that incremental power.

Bill Papanastasiou^ Great. Thank you for answering my questions, and best of luck in Q4.

Jay Crystal^ Yes. Appreciate it, Bill.

Operator^ Your next question comes from the line of Kingsley Crane with Canaccord Genuity. Your line is open. Please go ahead.

Kingsley Crane^ Great. Thanks for taking my questions.

Really encouraged to hear about the release of Token Factory. Just curious your thoughts on how you expect that offering could evolve over the next year, and just maybe what percentage of total

consumption could be occupied by something like Token Factory versus a GPU hour or a take-or-pay arrangement, if you look maybe a year or two out. Thanks.

Jay Crystal^ Yes. Thanks for the question, Kingsley.

In terms of Token Factory and how we see it evolving over the next 12 months, especially in terms of, yes, the percentage of capacity we might fill with it, for the clusters and larger footprints that we're bringing online, those are generally destined for long-term committed take-or-pay relationships with third-party customers.

The way we're planning to scale Token Factory, we've reserved a portion of our H200s and existing GPU fleet for that we're using to seed the effort, and then we've secured capacity on a third-party asset-light basis that enables us to grow cost-effectively, and also match -- better match the duration of the compute liabilities that we would bring on with the duration of the demand that we see through Token Factory.

So, we're effectively trying to grow, matching assets and liabilities in a controlled fashion, and thus far, we've been able to match capacity with demand as we grow that platform. Over time, the expectation as we think about a year or two out would be as we have stable base workloads, contracted guaranteed throughput, and other forms of forward-looking visibility, that would give us the confidence to begin to bring some of that capacity in-house.

And then there may be selective opportunities to use compute as we do these larger deployments to seed a portion of that compute toward the Token Factory, though for the larger-scale GPU clusters, that would be a minority percentage of the cluster that gets allocated to the Token Factory.

Kingsley Crane^ Okay. That's helpful. And that sort of flows into the second question of, just curious your thoughts on trends in spot pricing or committed pricing for different tiers of GPUs, whether that's H100 or Blackwell or Vera Rubin, and then just how you feel your business model can perhaps benefit from increases in spot. Maybe that's more related to Token Factory if you're planning on signing more committed arrangements for the majority of capacity. Thanks.

Jay Crystal^ Yes. Yes, great question. We see what you see in terms of increases in spot prices, but we've also seen increases in committed financing. To date, just to ensure that we lock in the certainty of the offtake with the liabilities that come along with deploying these clusters at scale, we've elected to allocate that capacity toward longer-term offtake in order to de-risk the outcomes for investors while still hitting very attractive rates of returns that exceed our hurdle rate.

In terms of spot, we've experimented with it in the past, and while it generated attractive yields, the market fluctuated over time with the portion of capacity that we allocated to spot. In other cases, a customer who may have been using the spot capacity for training moves off, and then we have to replace that customer.

So, it led to a lot of energy being dedicated toward maintaining a high utilization rate on the spot capacity in order to, when we think about effectively the effective revenue per token, the effective revenue per GPU hour filtered through the lens of utilization, at the time, it was too much effort relative to growing the business in a more scalable fashion through long-term committed offtake.

But again, just to circle back to your question, it's something we certainly keep an eye on. We may choose to experiment with in the future. And we are seeing upward movement in prices for both committed and spot, and where that could flow through to our business as longer-term offtake, even six-month and one-year deals from some of our legacy servers come back up for renewal.

There are opportunities to increase average pricing there. And then for the new deals that we're bringing to market, we're quite pleased with the economics that we're seeing come through.

Kingsley Crane^ Okay. Really helpful. Appreciate the time. Thank you.

Jay Crystal^ Of course. Thanks.

Operator^ Your next question comes from the line of Fedor Shabalin with B. Riley. Your line is open. Please go ahead.

Fedor Shabalin^ Thank you very much, Operator, and good morning, everyone.

My first question is on unit economics and the potential security premium. Specifically, how do these unit economics compare between a plain bare metal GPU lease and the full Corvex deployment with, call it, storage managed services and Assured AI attached? Maybe if you can frame it as the revenue per GPU hour per megawatt and incremental gross margin as a premium for the security. Thank you.

Jay Crystal^ Yes. Well, thank you for the question. Yes, great question.

As we think about, I guess the relative pricing differences between bare metal, managed Kubernetes, and then Assured AI, and some of the more advanced offerings, there is a premium that comes along with it. So, in terms of our upcoming deployments, part of the challenge, I guess, is when you filter that premium through the power footprint, the rack density of the servers that we're choosing to deploy. So, for instance, with our upcoming NVIDIA Blackwell cluster, we're quite encouraged in terms of where that agreement is going to shake out in terms of revenue per watt or per megawatt.

As we look to other offerings, even ones that would take our managed Kubernetes offering, where we do get a price premium on a per GPU per hour basis, and our Assured AI offerings, when you filter that through the power efficiency of the compute, even if we're seeing an increased premium in the GPU hourly rate that we contract for, if the density is different on the machines, that can impact how that flows through and is measured in terms of revenue per watt.

So, yes, we do see stratification and the opportunities to definitely distinguish pricing between bare metal, managed Kubernetes, and more of the software-driven offerings, all the way up to the revenue per watt that we anticipate being able to achieve over time on Token Factory. But we also filter it through the lens of the computing efficiency, the power efficiency of the machines that we're operating it on.

Fedor Shabalin^ Thank you very much for this. And my follow-up is on CapEx to build the capacity [already] secured. Pardon me if I missed it, but of the roughly 6.3 megawatts coming online this quarter, how much of the GPU networking and fit-out CapEx is already committed versus still uncommitted and in agreements? If you can just size the remaining spend for each site, that would be very helpful. Thank you.

Corvex (Q3 2026 Earnings)

Corvex Third Quarter 2026 Earnings Conference Call

Jay Crystal^ Yes, great question. A lot of, in terms of when we make commitments, is dictated by the timing of customer offtake. So, there may be certain items that we commit to in advance just due to lead times. For instance, we've publicly talked about committing to power in advance, in order to deliver the facilities on time to meet the needs of customers that we had high confidence would be coming through the funnel.

And in other cases, there are similar long lead items that we procure in advance. But the overwhelming majority of capital is ultimately committed and packaged and brought to market, simultaneous with when we memorialize the offtake agreements.

Fedor Shabalin^ Thanks for that. Thanks for answering my question and continue best of luck. Thank you.

Jay Crystal^ Thank you.

Operator^ There are no further questions at this time. This concludes today's call. Thank you for attending. You may now disconnect.