

S&P Global

Market Intelligence

QumulusAI, Inc. NasdaqGM:QMLS

Earnings Call

Tuesday, August 25, 2026 10:00 PM GMT

CALL PARTICIPANTS	2
PRESENTATION	3
QUESTION AND ANSWER	8

Call Participants

EXECUTIVES

Michael Maniscalco

Chairman & CEO

Scott Carroll Krosnowski

CFO & Secretary

Zachary Graeve

Head of Investor Relations

ANALYSTS

Allen Robert Klee

*Maxim Group LLC, Research
Division*

Brett Anthony Knoblauch

*Cantor Fitzgerald & Co., Research
Division*

Michael John Donovan

*Compass Point Research &
Trading, LLC, Research Division*

Vasili Papanastasiou

*Chardan Capital Markets, LLC,
Research Division*

Presentation

Operator

Hello, everyone. Thank you for joining us, and welcome to the QumulusAI Second Quarter 2026 Earnings Call. [Operator Instructions] Today's call is being recorded. I will now hand the conference over to Zachary Graeve, Head of Investor Relations. Zachary, please go ahead.

Zachary Graeve

Head of Investor Relations

Good afternoon, and welcome to QumulusAI's Second Quarter 2026 Earnings Call. I'm Zachary Graeve, Head of Investor Relations. And with me today, I have Mike Maniscalco, our Chief Executive Officer; and Scott Krosnowski, our Chief Financial Officer. This being our first earnings call as a public company, I'd like to spend a moment on a few items before we begin.

Our earnings release crossed the wire this afternoon, and that release, along with supplemental materials are posted on our Investor Relations website at investors.qumulusai.com. This call is being webcast live, and a replay will be available on that site. I'd like to remind everyone that certain statements made on this call may constitute forward-looking statements.

These statements include those about our strategy, capacity and deployment plans, customer agreements and future performance. Please be advised that these forward-looking statements are covered under the safe harbor provisions of the Private Securities Litigation Reform Act of 1995 and the QumulusAI avails itself of the protections of the safe harbor for these statements.

Forward-looking statements are based on management's current expectations and assumptions as of today and are subject to risks and uncertainties that could cause actual results to differ materially from those stated or implied. Those risks are described in our Risk Factors section of the registration statement on Form S-1 as amended, file no. 333-292514 filed with the Securities and Exchange Commission. Those factors may be updated in our subsequent filings, so please read these reports and future filings that QumulusAI will make with the SEC.

QumulusAI disclaims any obligation to update or revise any forward-looking statement, except as required by law. Please note that we may refer to non-GAAP measures in today's call. Any non-GAAP measure will be paired with the most directly comparable GAAP measure and reconciliations are included in today's release. Now with that, I'll turn the call over to Mike.

Michael Maniscalco

Chairman & CEO

Thanks, Zachary, and thanks to all of you for joining us. This is our first call as a public company. So I'll start the way I did on the NASDAQ stage in July with an expression of gratitude to our customers, employees, shareholders, Board members, and partners.

A number of them committed to this company well before there was much evidence to go on. We're thankful for that and excited about the future, which is the reason the rest of this call will focus on executing on the opportunity ahead. Let me tell you how I got here because it explains a lot about what this company is.

My background is software, computer science by training, technology entrepreneur by trade and much of that has been up the stack. So when I stepped out of the AI infrastructure business in January of last year, I went back to what I knew. I started building things, leveraging the power of the emerging Agentic coding tools, and I watched those tools get very good, very fast.

What I experienced change my mind. If software was going to become this easy to build, then software wasn't the constraint anymore, whatever sat underneath it was. And what sits underneath the AI tools are

the models and the compute to power them. The heavy AI infrastructure, power, data centers and GPUs were the enabler of every good thing I was watching happen.

So less than 8 months after I left the infrastructure space, I came back and I came back through Atlanta. I was in town visiting Georgia Tech, my alma mater where I met the QumulusAI team. QumulusAI's own story runs on a parallel track. Two companies, both founded in 2019, one built data center assets and power, the other built blockchain managed services.

In December of 2022, they merged into what is now a QumulusAI. And in April of last year, we acquired the Cloud minders, which brought GP as a service as a foundation, 2 halves of an answer put together. Through the early buildup of our Bitcoin business, we also ended up holding power in our own right.

We have long-term land leases on 2 sites, 1 in Oklahoma and 1 in Texas, totaling roughly 39 megawatts available today. That's in addition to our existing HPC data center lease footprint. We have 8 megawatts of compute completely sold today with the final GPU deployments actively in progress.

Last week, we announced a colocation agreement in Metro Atlanta. Our home market for up to initial 3.75 megawatts. That agreement also carries a right of first offer for up to 7 megawatts of additional capacity which the provider may pursue at the same site and could bring the Atlanta location to 10.75 total megawatts.

Now looking back, the guidance we issued on July 14, 2026, contemplated 18 megawatts of HPC capacity by year-end. The 8 we had plus 10 more to develop. Atlanta is the first 3.75 of that 10. If you add the 39 megawatts under lease, that's roughly 57 megawatts in our combined footprint going into 2027. That's real opportunity ahead of us.

Now to the quarter. On July 16, we began trading on the NASDAQ Global Market under the ticker QMLS. We did it as a direct listing. Our registration statement was declared effective on July 14, and we rang the opening bell at the NASDAQ MarketSite on July 29. On July 17, the day after we began trading, we were approved as an NVIDIA Cloud Partner, 2 things in the same week and both are about access.

One, to the capital markets, the other to the supply chain that the capital buys from. However, neither one is the goal. They're both tools for the same job, turning signed demand into deployed GPUs. One funds it, the other sources it. And we do that faster than most companies in this market. And for our customers, speed is everything.

It's also why we've built our model around hyper speed rather than hyperscale. This was the quarter our model started proving itself. Our GPU fleet went from 952 to 3,088 a 224% increase from Q1. It's worth highlighting that most of this was compute, we had already sold coming online and revenue more than doubled year-over-year.

Additionally, gross margin expanded to 66.6% from 55.1%. And Scott will talk about that later in today's call. But those 3 numbers tell one story, AI compute, we had already sold came online. And whenever compute comes online, that's when the economics of this business show up.

On the commercial side, we signed 21 new direct customer contracts totaling \$169.7 million in the second quarter. That brings our total signed contract value to \$282.5 million across 40 contracts with a weighted average term of 2.2 years, shorter than what you'll see from the largest players, but that's deliberately so. I'll come back to why later.

Everything we sold this quarter was contracted before the hardware was energized meaning customers are committing and paying ahead of delivery. Again, Scott will talk about that more a little later in this call. But I also want to say something about that \$282.5 million because it's the real number I'd watch. It isn't revenue. It's signed multiyear demand sitting in front of a company that recognized \$6.7 million in the quarter.

The distance between those 2 figures is the entire operating challenge of the business, and it's our job to close it. The only way you close it is by putting compute on the floor at hyperspeed. The customer base

has changed, too. Direct customer relationships now account for more than 96% of our recurring revenue up from less than 10% a year ago.

We've completed the transition away from dependence on a single marketplace to direct multiyear agreements with customers themselves with marketplaces used for what they're generally good at, tailored user experiences and filling short-term gaps between term customers as they roll on and off, and customer contracts are getting bigger. If you look at our recent announcements, the trend is clear. What we're signing now is materially larger than what we signed 6 months ago. As we secure more power, we expect that trend to continue.

Now I want to take a step back for a minute and talk about why we're built the way we are, because there's a lot of noise in this category. We are not a powered shell developer building sites to lease to the next hyperscaler. We are not brokering powered land to data center developments, and we're not brokering GPU clusters. We deploy and operate GPU clusters for our customers at scale, and we made a deliberate choice about how the traditional way to build AI infrastructure goes like this.

And I think this is important to understand. You do an initial data center design and seek data center development approval for about a year. Then you engineer it and then you finance it. You spend 24 to 36 months in construction. You roll in the GPUs, you spend a few more months commissioning. And 3 to 4 years later, you've delivered a token generating cluster. Now, look at the pace that AI is moving, then look at NVIDIA's chip road map. Neither of those paces line up with a 4-year build. So we don't lead with 4-year builds. We look for pockets of power. Call it, 2 to 50 megawatts, where the power and Shell already exists or can be ready for service soon, and we deploy quickly.

Most of that is colocation, and it means we bring capacity online in months, not years. Now that's hyperspeed. There is a trade in that, and I'll name it. It's gigawatt scale campuses. Gigawatt scale monolithic builds are important for AI, but they are for everybody, and we're not chasing those at the moment.

We are currently focused on customers who need capacity today, not those planning for 2030. That market is large and is underserved. What we get in exchange is speed. And in this market, speed is a core differentiator but speed alone does not win. There are 5 things customers actually buy. We call it our FACTS framework.

First, flexibility. Whether it's around location, vendors or architectures such as bare metal versus Kubernetes or Rocky versus InfiniBand. We provide customers that flexibility. Second is access. As I mentioned, our customers struggle to find capacity. So what we do, we provide access and capacity that's available when a customer needs it. And we do that with a fair and clear price that customers can grow into. That's our cost structure.

Next is trust. We're building long-term partnerships through reliability and SLAs or isolation and data privacy, we're here to support our customers' needs. Lastly, it's speed and all 5 matter. But one of them decides most deals right now. Customers just aren't opening with price. They're opening with when can I be running?

What we see is an industry that's compute-constrained top to bottom. From the largest providers down to the start-ups that can't get enough compute to keep pace with their customer demand. When supply can't keep pace with customer demand, speed stops being a feature, it becomes the decision.

That's the near-term reason we go after small pockets of power. It is the fastest way to get a customer running. There's a second thing that comes from building this way. because we typically deploy into pockets of existing power rather than financing four-year builds, we don't need to lock in every megawatt into longer-term 5-year contracts.

Most of our contract value is multiyear, but part of the book comes up for renewal each year, and GPU prices have been rising, not falling, which means in a market like this, we'd rather have some capacity repricing than all of it locked rates we set years ago and that's deliberate. But we're also making a longer-term bet on that same footprint.

Today, distribution is a means to an end. Over time, we think it becomes the end itself. As workloads shift from training to inference, they get more particular about where they run. For example, budget starts to matter more. Latency begins to matter. Proximity to the end users starts to matter. So does the security and compliance posture of a specific location.

So when customers start choosing in those terms, and network of sites in different places stops being an operational complexity and becomes a differentiated product. Now we'll see how quickly that develops. In the meantime, the same footprint is doing exactly what we need it to do today. Delivering supply at the highest speed. That's why we built this company inference first.

Tokens are the output. And getting tokens to customers faster than anybody else in our business is winning today. Getting them to the right place is the business we're building towards tomorrow. Now let me put it together. In the second quarter alone, we more than tripled our GPU fleet.

We signed 21 new contracts, we gained access to the capital markets and a place inside NVIDIA's partner ecosystem. Our go-to-market motion is working. We activate capacity, sign demand and deliver the supply. Those are the 3 dials and all 3 move this quarter. Thank you.

Now I'll turn it over to Scott to walk you through our financial reports in more detail. Over to you, Scott.

Scott Carroll Krosnowski
CFO & Secretary

Thanks, Mike. Good afternoon, everyone. I'll take revenue first, then walk down the income statement, then finish with cash on the balance sheet. First, revenue. Revenue was \$6.7 million in the second quarter, up 118% compared to \$3.1 million from the same period last year.

The driver is compute power, which grew to \$5.6 million from \$1.3 million a year ago, up 328% or \$4.3 million as GPUs activated primarily on term compute contracts representing 84% of total revenue in the second quarter compared to 43% from the same period last year. Our legacy lines moved as expected, together mining hosting and cryptocurrency mining were about 16% of revenue this quarter down from roughly 39% in the first quarter and 57% a year ago.

We expect this downward trend to continue as we activate more GPUs. Now margin. Gross profit was \$4.5 million compared to \$1.7 million a year ago. Gross margin expanded to 66.6% from 55.1%. Two things drove the improvement, and both are clear sequentially and margin moved 29 percentage points from 37.5% in the first quarter.

GPU activations grew faster than the associated HPC co-location costs, and we recognized approximately \$0.4 million in curtailment credits at our Oklahoma site. Credits for returning power to the grid during peak demand periods that reduce power costs.

Down to the operating line. Operating loss for the quarter was \$7.7 million compared to \$2.2 million a year ago. the loss widened, and I want to be specific about why because the reason is favorable. Depreciation and amortization were \$6.9 million in the quarter against \$1.1 million a year ago and up \$4.2 million quarter-over-quarter.

That sequential step is the direct consequence of the HPC assets Mike described coming online and driving top line and margin growth. Adjusted EBITDA, which excludes interest, taxes, depreciation and noncash items, was a loss of \$0.8 million compared to a loss of \$0.3 million a year ago.

Sequentially, the loss narrowed from \$2.8 million in the first quarter. That's a non-GAAP measure. The most directly comparable GAAP measure is net loss and a full reconciliation is in today's release. The year-over-year comparison reflects the public company and personnel investments we have made. The sequential improvement reflects revenue beginning to scale against them.

General and administrative expenses were \$4.3 million, up \$1.9 million from a year ago, driven by increased headcount costs associated with becoming a public company and scaling up infrastructure operations.

Sales and marketing expenses were \$0.9 million, up from \$0.4 million a year ago, driven by increased headcount and costs associated with promoting and marketing our HPC services. And below the bottom line, Net loss was \$22.8 million compared to net income of \$12.1 million a year ago. The year-over-year comparison is significantly affected by noncash accounting in both periods.

The prior year result included a \$14.5 million noncash gain related to the remeasurement of our investment in the cloud minders upon acquiring the remaining interest. In the current quarter, the largest single item is a \$19.2 million noncash loss associated with the issuance of convertible notes, which did not represent a cash outflow.

The charge was partly offset by a \$6.2 million in noncash fair value gains on the notes and related option. Net interest expense was \$2 million, primarily related to loan draws and lease financing.

Turning to cash. Operating cash flow for the first 6 months was positive at \$22.3 million against \$0.8 million used a year ago. The driver is customer prepayments. Deferred revenue rose \$30.5 million, on our major multiyear contracts, we see prepayments in the range of 10% to 35%, and that money in some instances, supports the deployment it's attached to.

Investing activities used \$36.3 million primarily purchases and deposits of HPC and related equipment. Financing activities provided \$42.2 million. Cash and restricted cash ended the period of \$39.9 million, up from \$11.7 million at year-end. That includes \$19.9 million that have been restricted pending our public listing, which became available following the direct listing on July 16.

On the balance sheet, total assets were \$215 million, up from \$91.7 million on December 31, driven by finance right-of-use assets, property and equipment and deposits on power equipment.

On the liability side, we carry \$55.5 million of convertible notes payable, a \$38.7 million additional convertible notes option liability, \$18.9 million of protocol loans and finance lease liabilities of \$45.8 million. One note on how the term compute contract value shows up in our statements. Mike gave you a total signed contract value of \$282.5 million, in our financial statements, you'll see a related but narrower figure, remaining performance obligations of \$173.1 million as of June 30th. That's the accounting measure, direct customer compute contracts only and only what was signed as of quarter end.

Two last points. First on unit economics. Our most recent Blackwell contracts are generating between \$18 million and \$20 million of annualized revenue per megawatt across our installed base, the blended figure is closer to \$16 million per megawatt. The gap is pricing power. Our deployments are currently focused on the blackwell chips. The price per GP hour has been firming across the market and customers will pay for GP using capacity that come online quickly.

All 3 of these factors are working in our favor and the revenue we earn per megawatt is widening as we scale. Moderator, we are ready to take questions.

Question and Answer

Operator

[Operator Instructions] Your first question comes from the line of Michael Donovan with Compass Point.

Michael John Donovan

Compass Point Research & Trading, LLC, Research Division

Congrats on the progress. So we could discuss the AI XP opportunity. How are economic structure across Qumulus and site partners, particularly around CapEx contributions and revenue sharing. And can we get a sense of timing for bringing on initial sites of the 25 sites that you indicate in the presentation?

Michael Maniscalco

Chairman & CEO

Yes. Thanks, Michael. I appreciate it. So I want to keep those somewhat separate because I think that if you reflect back on the statements I made a few minutes ago, where we're seeing a lot of success in deploying chips quickly is into third-party colocation capacity that's available RFS this year.

So a lot of the effort that our land powered sale team and colo team is doing is out looking for colo capacity that's undergoing, say, renovations or just coming online that is in the single-digit megawatts that we can move into quickly. And that go-to-market motion has been working really well for us. So that's where you're seeing a lot of growth.

That's where you're seeing our new colo announcements recently because that's the fastest path to get chips online. So that being said, the AI XP initiative and those sites that we announced are still in progress. That project is making good progress. They've been posting pretty good updates online.

The collective groups are responsible for bringing the land, the power of the shell to us. And our responsibility at that point is to once connectivity and power and cooling is ready, we roll in the racks and bring the GPUs online. So that's the role we play in that arrangement. The Sites are energizing in phases right now.

And really, just to be a little more specific, the first site in Wichita is really the one we're focused on today. The additional sites are still a little further out in our pipeline. That's helpful. I appreciate it.

Operator

Your next question comes from the line of Brett Knoblauch with AI infrastructure.

Brett Anthony Knoblauch

Cantor Fitzgerald & Co., Research Division

As we look at the remainder of the year, I'm just curious on the power ramp. I think previously, you talked about getting up to 16 or 18 megawatts. So I guess -- is that still the target? And then the \$16 million to \$18 million or \$16 million revenue per megawatt figure of the kind of existing base. At what point would you expect that to start to show up or all those contracts to be kind of up and running where we see that on the 8-megawatt figure that you guys disclosed?

Michael Maniscalco

Chairman & CEO

Yes. Thanks I'll answer the first part and Scott, if you want to round the thing out upfront, done, please feel free. But on that, so if we look back at the 18 megawatts that we guided to on July 14, we had put out a plan to hit an additional 10 megawatts on top of the 8 megawatts that we already had, so right?

So 18 megawatts total with 8 already leased, and that is actually fully sold, and some of that's generating revenue now. We expect all of that 8 megawatts to generate revenue this year or to begin generating

revenue this year. But that additional capacity, we just announced the 3.75 that we've added to that capacity.

So we're continuing to push really hard on that and still see a fair amount of smaller scale capacity out in the market that's available RFS this year. And to give you a little more specifics there, our preference is to find, call it, 12 megawatts that's available RFS this year but can expand beyond that. So if it's a 3.75, that has the capability to expand to 10, that's a much better fit for our longer-term needs.

However, we're also seeing smaller pockets of, say, stranded 1 to 2 megawatts that are still out there that haven't been prioritized so we still see some good options out to hit the '18 that we were looking back at July.

Brett Anthony Knoblauch

Cantor Fitzgerald & Co., Research Division

Maybe just on the CapEx front, kind of all the contracts that you've signed so far, I guess, how much maybe GPU purchases or CapEx or GPU-related CapEx is left for you guys? And then I'll hop back in queue.

Michael Maniscalco

Chairman & CEO

Okay. And I think you asked about the revenue per mile as well. So Scott, I'll take the revenue per, you can jump in on the CapEx piece. But we are seeing, I think, the latest economics are coming into the \$18 million to \$20 million per megawatt on the revenue side of the equation. And then Scott, do you want to comment on the CapEx side of it?

Scott Carroll Krosnowski

CFO & Secretary

Sure. Yes. The -- all of the CapEx that is scheduled to go into the 8 megawatts has been ordered and financed. So -- and as Mike alluded to -- Mike mentioned, we expect all 8 megawatts to be active and revenue producing by the end of the year.

Operator

Your next question comes from the line of Bill Papanastasiou with Chardan. Please go ahead.

Vasili Papanastasiou

Chardan Capital Markets, LLC, Research Division

Congrats on the progress and the inaugural earnings call as a public company. First one for me. So as mentioned in the print that Qumulus has signed a total of 21 new direct contracts in the quarter.

Maybe you could just spend some time speaking about what customer profiles you're targeting as the company scales. Are you looking for a particular enterprise customers? And what goes into that decision?

Michael Maniscalco

Chairman & CEO

Yes. Thanks, Bill. And we're excited to get this first one in front of us and behind us. And so thank you for that. Yes, so the new contracts we've announced that is made up of a total of 8 customers. We have announced a couple of those customers where the customers have been willing to -- but generally speaking, what we're seeing as far as those types of customers, it's a combination of things.

Some of it is some of our marketplace partners, which we've been working with for the past couple of years. Some of it is with the -- and the way I explain this is they're the AI inference GPU model as a service, generative AI as a service type companies that are seeing really strong customer growth.

They've usually gone out and raised significant capital from brand name venture Capitalist to for argument sake I'll say they've gone on raise to, say, \$100-plus million from a16z or Sequoia or the likes of those

categories. And they're in a rapid growth phase with a lot of customer and they need to compute to satisfy I would say that's where a lot of the near-term growth is coming from, we call those the AI natives.

And then in our pipeline, we do have a lot of interest and continue to look at the enterprise segments. However, and enterprise segments tend to make decisions and move a little less quickly than say, the AI natives. But we are working with a lot of those verticals.

I'm happy to talk about some of the verticals. I think we could call that out in our investor slides from July, and I don't think that's changed too much, but we continue to focus on the enterprise verticals as well.

I would say, as we are proving ourselves in the market as we are bringing more supply and that word is getting out to market. And as we are increasing our organizational maturity level, the quality of our customers, we also see increasing. So that's another positive that all highlight.

Vasili Papanastasiou

Chardan Capital Markets, LLC, Research Division

Great. Appreciate that color. And then hoping you'd be able to provide an update on the GPU financing environment. Have you seen any changes in economics could become more favorable now that you're a public company and you're seeing this shift in your customer profile base?

Michael Maniscalco

Chairman & CEO

Yes. I mean, I would -- the way I've explained this pretty consistently over the last couple of months, is it seems like the capital market side of the equation is moving quickly to try to keep pace with how quickly the AI space is moving. So we are -- we've got our historical ways of funding.

Those have been working out really well. As a reminder, there has been a combination of equity and debt. A lot of that -- a portion of that debt has come from USDI in the Permian Labs agreement. Some of that has come from traditional lease agreements with companies like Tech Finance.

So that's been working really well, but we're also seeing new announcements from our peers and from our partners about other creative financing opportunities and options. And really, from our standpoint, it's about aligning the right capital to the right asset, whether it's a data center power agreement or GPU or other, and then the right customer. and then understanding what the right cost of capital is. And yes, I would say there are more options now than ever that we're seeing out there, and that's encouraging.

Vasili Papanastasiou

Chardan Capital Markets, LLC, Research Division

And then last one for me. apologies if I missed this, but can we just touch on the prior guidance figure of \$300 million exiting in the year. Did you mention that, that was still in play? And can you speak to the demand funnel from July 14, when that guidance figure was published? Are you seeing the demand funnel kind of grow at the same faster or slower clip than was previously expected?

Michael Maniscalco

Chairman & CEO

Yes. So we did put out the guidance on July 14 with a \$300 million ARR target. We are not reaffirming any of the guidance today, but what I can tell you is we're not currently demand constrained.

We're not able to satisfy all the customer requests that our sales team is and marketing are starting up and bringing in and our partners are bringing in. So really the largest constraint to our growth right now is the land powered shell in the near term. And so that's why we've got a big focus on our land powered shell team on going out and finding that additional power.

Specifically in that July guidance was 18 additional megawatts or 18 total megawatts at the end of the year on top of the 8 we already had. So a lot of focus is on that, both in the end of the year, but also into '27 even '28. I feel really good about the land powered shell team and the team in itself. I think that we've done a great job at -- having capital markets and the Land Power shell teams focus on those efforts, while

the operations team is heads down focusing on execution and delivering the contracts and bringing in the revenue on the contracts we've announced. So that's working really, really well. I'd just like to think that's a good point to highlight.

Operator

Your next question comes from the line of Allen Klee with Maxim Group LLC. Please go ahead.

Allen Robert Klee

Maxim Group LLC, Research Division

Could you talk a little about the competitive environment of companies you were going up against for GPU compute?

Michael Maniscalco

Chairman & CEO

Yes. Thanks, Allen. Absolutely. I talked a lot about this in a lot of our historical materials and also, I think, in our recording that the differentiation is really about the what we are not, right? There are a lot of neoclouds pursuing the big, big opportunities, hundreds of megawatts gigawatts right now.

Those things are signing out years out. And to the biggest companies in the world, that's really important to the largest frontier model training companies, they need to lock down that capacity for years to come to train and host the largest model. So that is an important part.

I can't emphasize that enough. Those gigawatt scale campuses in those longer-term projects are extremely important to the future of AI and to continue the momentum. However, there's just a lot of near-term demand of customers who are calling and saying, "Hey, we would love compute yesterday, if you could get it to us."

That is a theme we hear from our sales team. We hear it from our customers. And we hear it from the market. It's not just us. If you look -- listen to our peers calls and you listen to hyperscaler calls. I mean everybody is saying we are restricted -- growth constrained by supply. I think it's the exact words that some of them are using right now.

So our differentiated approach right now is really -- it's a combination of things, and I refer to it as a right to play, right? The right to play in this space is, do you have access to capital? Do you have the land power shell to get the chips into? Do you have the partnerships in the supply chain to procure those chips and highlight the NVIDIA cloud partner relationship.

That's a key piece of it, but also our OEM relationships with the likes of the Super Micro, Lenovo, Dell, et cetera, right? And then it's can you bring those online and operate those effectively and do you have the team to do it. And this is my -- one of my favorite things to say as a former public company CTO is it feels great to have a CTO I can really trust and lean on.

So you put those pieces together, and I think we check the boxes, right? So then it's -- what is the additional differentiation that we're seeing and what's making us win these contracts, and it really is to me, speed and fair costs.

We can get it to our customers quickly because we're not out looking for 150 megawatts at a time, and our customers are open to a few megawatts of compute at a time, if not less. And we can move that a lot faster. So I think to summarize it, it's that sub-50 megawatt thesis, where we believe that we can deploy and find pockets of power shell, roll the GPUs, get them online a lot faster if we are focused on smaller to more moderate pockets of compute rather than large deployments.

Allen Robert Klee

Maxim Group LLC, Research Division

One other question. Could you just remind me of when you're trying to finance GPUs for the debt component, are those -- how are -- how much does it typically amortize and is there any changes in rates with the current environment or that you're seeing now versus maybe 6 months ago?

Michael Maniscalco

Chairman & CEO

Scott, I might defer that to you. I think you can probably speak the clearest in claims right now.

Scott Carroll Krosnowski

CFO & Secretary

Yes. So most of our financing arrangements are 3 years. We do have a few that are longer, but most of those are financed over 3 years, which lines up typically with some of our recent longer-term compute contracts. In terms of -- our partners, we have some great partners that we've already done some financing relationships with, and we expect to continue to grow with them.

So as we get bigger, as we deploy more, we expect, obviously, the cost of capital to go down. But that's really, really what we're focused on is finding the lowest cost of capital in each case. Thank you.

Operator

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

Copyright © 2026 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages. S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not act as a fiduciary or an investment advisor except where registered as such. S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its Web sites, www.standardandpoors.com (free of charge), and www.ratingsdirect.com and www.globalcreditportal.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.

© 2026 S&P Global Market Intelligence.