
**UNITED STATES
SECURITIES AND EXCHANGE COMMISSION
Washington, D.C. 20549**

Form 8-K

CURRENT REPORT

Pursuant to Section 13 or 15(d) of the Securities Exchange Act of 1934

Date of Report (Date of earliest event reported): May 08, 2026

IREN LIMITED

(Exact name of registrant as specified in its charter)

Commission File Number: 001-41072

Australia
(State or other jurisdiction of
incorporation)

Not Applicable
(IRS Employer
Identification No.)

Level 5, 55 Market Street, Sydney, NSW 2000 Australia
(Address of principal executive offices, including zip code)

+61 2 7906 8301
(Registrant's telephone number, including area code)

(Former name or former address, if changed since last report)

Check the appropriate box below if the Form 8-K filing is intended to simultaneously satisfy the filing obligation of the registrant under any of the following provisions:

- Written communications pursuant to Rule 425 under the Securities Act (17 CFR 230.425)
- Soliciting material pursuant to Rule 14a-12 under the Exchange Act (17 CFR 240.14a-12)
- Pre-commencement communications pursuant to Rule 14d-2(b) under the Exchange Act (17 CFR 240.14d-2(b))
- Pre-commencement communications pursuant to Rule 13c-4(c) under the Exchange Act (17 CFR 240.13c-4(c))

Securities Registered Pursuant to Section 12(b) of the Act:

Title of each class	Trading symbol(s)	Name of each exchange on which registered
Ordinary shares, no par value	IREN	The Nasdaq Stock Market LLC

Indicate by check mark whether the registrant is an emerging growth company as defined in Rule 405 of the Securities Act of 1933 (§230.405 of this chapter) or Rule 12b-2 of the Securities Exchange Act of 1934 (§240.12b-2 of this chapter).

Emerging growth company

If an emerging growth company, indicate by check mark if the registrant has elected not to use the extended transition period for complying with any new or revised financial accounting standards provided pursuant to Section 13(a) of the Exchange Act.

Item 2.02 **Results of Operations and Financial Condition.**

On May 7, 2026, IREN Limited (the “Company”) conducted a conference call to discuss its financial results for the three months ended March 31, 2026. A copy of the transcript of the conference call is furnished herewith as Exhibit 99.1 and is incorporated herein by reference.

The information contained in this Item 2.02 and Exhibit 99.1 attached hereto shall not be deemed to be “filed” for purposes of Section 18 of the Securities Exchange Act of 1934, as amended (the “Exchange Act”), or otherwise subject to the liabilities of that section, nor shall it be deemed incorporated by reference in any filing under the Securities Act of 1933, as amended, or the Exchange Act, except as shall be expressly set forth by specific reference in such filing.

Item 9.01 **Financial Statements and Exhibits.****Item 9.01(d)** **Exhibits**

Exhibit No.	Description
99.1	May 7, 2026 conference call transcript.
104	Cover Page Interactive Data File (embedded within the Inline XBRL document)

SIGNATURES

Pursuant to the requirements of the Securities Exchange Act of 1934, the registrant has duly caused this report to be signed on its behalf by the undersigned hereunto duly authorized.

IREN LIMITED

By: /s/ Daniel Roberts
Daniel Roberts
Co-Chief Executive Officer and Director

Date: 8 May 2026



IREN Q3 FY26 Results Transcript

Presenters:

- Daniel Roberts, Co-Founder & Co-CEO
- Anthony Lewis, Chief Financial Officer
- Kent Draper, Chief Commercial Officer
- Michael Power, Vice President, Investor Relations

Operator: Good day and thank you for standing by. Welcome to the IREN Q3 FY26 results presentation. At this time, all participants are in a listen-only mode. After the speakers' presentation, there'll be a question and answer session. To ask a question during the session, you need to press star one, one on your telephone keypad.

You will then hear an automated message advising your hand is raised. To withdraw your question, please press star one, one again. Please be advised that today's conference is being recorded. I would now like to hand the conference over to our first speaker today, Mike Power, Vice President of Investor Relations. Please go ahead.

Mike Power (Vice President, Investor Relations): Thank you, Operator. Good afternoon and welcome to IREN's Q3 FY 2026 results presentation, and thank you for your patience as we get assembled. I'm Mike Power, VP of Investor Relations, and with me on the call today are Daniel Roberts, Co-Founder and Co-CEO, Anthony Lewis, CFO, and Kent Draper, Chief Commercial Officer.

Before we begin, please note that this call is being webcast live with an accompanying presentation. For those dialed in by phone, you can elect to ask a question through the moderator after our prepared remarks.

I would like to remind everyone that certain statements made during this call may constitute forward-looking statements. Those statements are based on current expectations and assumptions and are subject to risks and uncertainties that could cause actual results to differ materially.

Please refer to slide 2 of the accompanying presentation and SEC filings for more information. During today's call, we will also refer to certain non-GAAP financial measures. A reminder: a reconciliation to the most directly comparable GAAP measures is included at the end of the presentation.

So with that, I will turn the call over to Dan Roberts.

Daniel Roberts (Co-Founder & Co-CEO): Thanks, Mike, and thank you, everyone, for joining us today. Eight years ago, when Will and I founded this business, we spent a lot of time thinking about what the digital future actually meant for the physical world.



We talked about films like *The Matrix* and *Ready Player One*, not as science fiction, but as a signal. Worlds where digital adoption was total, instantaneous, and infinite. The insight we kept coming back to was this: digital adoption curves can go from zero to one overnight, but the real world doesn't scale that way.

Power, infrastructure, land, data centers – these take years to permit, finance and build. The bigger the demand, the harder delivery becomes. That gap between exponential digital growth and the physical world's ability to service that structural disconnect is exactly what we set out to solve.

That scarcity is now defining where AI infrastructure gets built and who can build it. Eight years later, that thesis is playing out exactly. This quarter, we demonstrated what disciplined execution against it looks like at a global scale.

In AI infrastructure, secured power is only valuable if it can be converted into customer-ready compute. That conversion is hard. It requires site control, grid connection work, permitting, design, procurement, construction, GPU installation, networking, commissioning, financing, and customer delivery, all coming together on tight timelines.

IREN's strength is bringing those pieces together. We have experienced site teams, standardized designs, and repeatable construction processes that allow us to build across multiple sites in parallel.

As we scale, each phase builds on the prior phase. The template becomes more repeatable, the procurement and construction process becomes more efficient, and the site teams carry that experience forward.

That is where IREN has built its moat, and why real assets and real capabilities are harder to replicate than they might appear. That execution capability is showing up in the numbers. More capacity, more revenue, stronger funding certainty, and in the partnerships we're announcing today.

This was a significant quarter and a significant week. Let me run through the highlights.

On capacity, we increased secured power to five gigawatts, added new sites in Europe and APAC, energized Sweetwater 1 on schedule, and have Horizon 1 GPU commissioning now underway for Microsoft.

On customers, all of our operational capacity is fully contracted. We are not chasing demand. We are racing to build supply fast enough to meet it. In this market, the moment compute comes online, it goes to work.

That is the nature of the structural imbalance of AI infrastructure between supply and demand, and it is why time-to-compute is the most important metric we track. We increased ARR under contract to \$3.1 billion, remain on track to hit \$3.7 billion exiting calendar 2026, and this week signed a \$3.4 billion five-year AI Cloud contract with NVIDIA.

The first step in a broader strategic partnership I will come to in a moment. On capital, we had \$2.6 billion of cash at 30 April, and we continue to progress GPU, data center, and corporate level financing initiatives to support the next phase of build-out.





But the headline today is the NVIDIA partnership, and it deserves a little more than a bullet point. Let me explain what this partnership actually means.

We are working with NVIDIA to support deployment of up to five gigawatts of NVIDIA DSX-aligned AI infrastructure across our global data center platform, alongside DGX environments and the DSX AI factory reference architecture.

The \$2.1 billion NVIDIA investment is structured to reflect that. Their rights to invest only vest as NVIDIA GPU infrastructure is deployed across IREN campuses and only fully vest upon deployment of 600,000 GPUs.

NVIDIA's capital is directly tied to execution. That's not a passive financial investment. NVIDIA is a partner who wins as we deliver. The \$3.4 billion AI Cloud contract announced today, supporting NVIDIA's own internal workloads, is the first step in that partnership.

Eight years ago, Will and I set out to build the infrastructure the digital world would need. Today, the world's leading AI infrastructure company has chosen IREN as the partner to help build it.

This next slide shows exactly how we do it against this. So here is our plan. In 2026, we are targeting 480 megawatts of AI Cloud capacity, 150,000 GPUs, and 3.7 billion of ARR by year-end.

That is the near-term plan and the clearest bridge from capacity to revenue. In 2027, we are scaling to 1,210 megawatts, with an additional 730 megawatts currently under construction across British Columbia and Texas, including Childress and the initial phase at Sweetwater 1.

The construction flywheel we are running in 2026 carries directly into this next phase. Beyond 2027, we are building against a five gigawatt global power portfolio. North America, our new European platform in Spain, and an APAC pipeline anchored by large-scale Australian opportunities.

The sequence of delivery matters because it dictates time-to-compute, and time-to-compute is what drives revenue. Each phase supports the next, that's how the platform compounds.

One more thing before we move on. This week we welcomed Mirantis into the IREN family. Six hundred and fifty engineers, operators and customer support professionals who have spent more than a decade running cloud infrastructure for over 1,500 enterprise customers globally. To Alex and the whole Mirantis team, welcome. I'll come back to what this means for our delivery capability later on.

But let me start with 2026, where construction and customer demand are coming together most visibly. The 2026 expansion is focused on delivering 480 megawatts of AI cloud capacity across Childress, Prince George, and Mackenzie.

This is where the roadmap translates into near-term deployments, customer handoffs, and ARR conversion. We'll start with the largest and most complex 2026 work stream, the 300 megawatt Horizon 1 to 4 liquid-cooled deployment at Childress, where NVIDIA GB300 NVL72 installations are now underway.





Horizon 1 is scheduled for Microsoft handoff in Q3, and Horizons 2 to 4 remain on track for delivery by the end of this year. This is a major execution milestone. It demonstrates our ability to design, build, fit out and commission large-scale next-generation liquid cooled infrastructure for a hyperscale customer on an accelerated schedule.

We have around 3,000 workers on-site right now. That level of activity reflects both the urgency of AI infrastructure demand, and also the depth of our execution capability on the ground.

Importantly, the model is repeatable. Horizon 1 establishes the build template. Each subsequent phase benefits from the same design, supply chain, construction sequencing, and site team. That is how we drive faster deployment phases every time.

Alongside the liquid-cooled build, we're also converting existing air-cooled capacity into AI Cloud deployments across British Columbia and Childress. In British Columbia and Childress, we are progressing 180 megawatts of air-cooled AI Cloud capacity by leveraging existing infrastructure.

At Prince George, all air-cooled GPUs have now been delivered and are either operating or undergoing commissioning across the 50 megawatt site. At Mackenzie, 80 megawatts of data center capacity has been prepared for GPU installations commencing in the second half of 2026.

Finally, at Childress, data center retrofits are underway across an initial 50 megawatts ahead of GPU deliveries in the second half of this year. This is a capital efficient part of the roadmap. We are taking existing sites and converting them toward higher value AI Cloud workloads.

It works because we already have the operational teams, infrastructure, and site control in place. Air-cooled capacity can come online faster than liquid-cooled. In a market where time-to-compute is everything, that speed is a commercial advantage, and we are using it.

We're already seeing this dynamic play out commercially with capacity continuing to be contracted ahead of commissioning as customers prioritize speed to market. We now have \$3.1 billion of ARR under contract, including approximately \$700 million of ARR associated with a \$3.4 billion five-year contract for Blackwell GPUs to be deployed across 60 megawatts of air-cooled capacity at Childress for NVIDIA.

Against the full 2026 expansion, we are targeting \$3.7 billion of ARR by year-end, across 150,000 GPUs. The remaining uncontracted capacity represents approximately 50,000 air-cooled GPUs scheduled for delivery in phases through the second half of this year. Demand for that capacity is robust. Our focus is on using our time-to-compute advantage to secure the right customer mix.

With the 2026 plan on track, let me turn to what comes next. The 2027 expansion, where the platform scales to 1,210 megawatts. So the 2027 plan is about demonstrating that what we are building in 2026 is not a one-off. It is a repeatable, scalable model that should accelerate over time.





Here's what that looks like in practice. In British Columbia, Canal Flats is another example of converting existing infrastructure into AI Cloud capacity. We plan to retrofit all 30 megawatts of existing air-cooled capacity to support AI workloads. Capital efficient, fast to execute, and consistent with the same model we are running at Prince George and Mackenzie.

In parallel, Childress continues to be the largest single contributor to the 2027 setup. With both new liquid-cooled capacity and additional air-cooled retrofits adding a total of 400 megawatts of gross capacity.

At Childress, the 2027 plan includes 100 megawatts of additional liquid-cooled IT load for Horizons 5 and 6, as well as retrofitting an additional 250 megawatts of existing air-cooled capacity. Of that 250 megawatts, approximately 60 megawatts will be deployed to support the NVIDIA AI Cloud contract.

The combination of new liquid-cooled data centers and air-cooled retrofits gives us real flexibility. We can support next-generation high-density deployments while continuing to use existing infrastructure where it is the right technical and economic fit. That flexibility is part of what makes Childress such a productive campus.

In parallel, Sweetwater becomes the next major Texas campus in the 2027 plan. At Sweetwater 1, the high-voltage substation has been energized on schedule and construction is now underway for the initial 200 megawatts IT load phase of liquid-cooled data centers.

Energizing the substation is an important milestone. It moves Sweetwater from development into execution and establishes the electrical foundation for the broader site build-out.

Sweetwater 1 is being designed for next generation chip architectures, including the NVIDIA Vera Rubin. Like Childress, we are deliberately sequencing the build so that the first phase creates the backbone for faster subsequent phases.

The first 200 megawatts is not just the first 200 megawatts, it is the foundation for a much larger site. The commercial pipeline for our 2027 capacity is anchored on the same principle that is driving everything we are building.

Our vertical integration is a genuine advantage for customers because we control more of the critical path than anyone else in this market. Power, land, data center construction, the pieces that cause delays for others are the pieces we own and control. Businesses want certainty that capacity will be available when promised.

The phased 2027 build-out plan gives us a concrete basis for those conversations, and we are having them. We are in the process of negotiating large-scale AI cloud deployments across our 2027 capacity today.

Demand is not the constraint, however. It is highly unlikely to be the constraint. The priority is delivering capacity on schedule and converting our time-to-compute advantage into durable, long-term customer relationships.





We do expect the customer mix to evolve over time. Hyperscalers, AI natives, enterprises and on-demand use cases, but we do not need to force that outcome. The platform will attract the right customers as it continues to scale.

Beyond 2027, the same execution model extends into a much larger five gigawatt global platform. We now have five gigawatts of secured power. To put that in context, that is not a pipeline number or an aspiration. That is secured power and it represents one of the largest portfolios assembled for AI infrastructure anywhere in the world.

The question now is how we build against it. The answer is a phased global platform across North America, Europe and APAC, with additional development opportunities beyond that. Let me walk you through each region.

We'll start with North America, which remains the largest component of the long-term platform. In North America, the next major phase is driven by Sweetwater and Kiowa, our flagship gigawatt-scale campuses in Texas and Oklahoma, where data center capacity is expected to commence ramping across 2027 and 2028.

We also have multiple development projects advancing through the connection processes, including Batch Zero candidates in Texas, which represent some of the most strategic, valuable grid connection opportunities in the country.

The North American pipeline has a natural progression of scale. Childress demonstrates the operating model today, Sweetwater expands it across an even larger campus, and Kiowa provides the path to another hyperscale tier opportunity as power ramps from 2028.

Every campus builds on the last. That's the compounding effect of having secured the right land and power positions early. At the same time, we're expanding the platform into Europe through Spain.

Today, we announced the acquisition of Nostrum Group, and with it our entry into Europe. The transaction adds 490 megawatts of secured power in Spain, a gigawatt scale development pipeline, and a team of more than 50 people across development, engineering, construction, and operations.

But what it really adds is a platform and the right people to build it. I want to acknowledge Gabriel Nebreda and the Nostrum team. Gabriel spent nearly two decades in European energy, at EDP Renewables managing gigawatts of operating assets across multiple European markets, and most recently as CEO of EDP Solar.

He understands European power infrastructure as well as anyone, and we're excited to have him leading IREN's European platform. Spain is the right place to start. Supportive AI policy, abundant renewables, lower build costs, and strong connectivity into broader European demand.





Europe is a market where power availability and grid timelines are increasingly shaping where customers can actually deploy. Spain gives us a credible, scalable answer to that question. This is not just a power acquisition, it's the establishment of IREN's European platform.

From Europe, we move to the other side of the world, and an opportunity that matches the scale of everything we've just described. Australia is obviously not a new idea for us. We've been progressing large-scale Australian projects towards secured grid access for some time.

We think the opportunity here is as significant as anywhere in our portfolio. This is why: Asia-Pacific is home to roughly 4.8 billion people, around 60% of the world's population. That includes some of the fastest growing AI demand markets on Earth. Indonesia, Singapore, Japan, Korea.

The infrastructure requirement to service that demand is enormous and it is largely unmet. Australia is uniquely positioned to serve it. Abundant renewables, a trusted jurisdiction, strong rule of law, and as the submarine connectivity map shows, direct fiber links into major demand centers across the region. It is the natural anchor point for AI infrastructure service in APAC.

We are already seeing hyperscalers and frontier labs make significant commitments to Australian operations and we intend to be a major part of that story. Beyond Australia, we continue to progress global development opportunities that extend IREN's runway further still.

The platform we are building is designed to create scale into demand wherever it develops, and the pipeline gives us the flexibility to do exactly that. That is the global platform, secured power across North America, Europe, and the development pipeline extending into APAC and beyond.

But securing power and building data centers is only part of the equation. The other part is what happens when the compute goes live, how it is deployed, managed, and supported for customers at scale. That is where I'd like to spend a moment on Mirantis.

This week, we welcomed Mirantis into the IREN family. I want to take a moment to acknowledge that. Six hundred and fifty people joined IREN this week. Engineers, operators, customer support professionals, a team that has spent more than a decade building and running cloud infrastructure for over 1,500 enterprise customers globally.

That track record speaks for itself. What they bring is specific. Their k0rdent AI platform manages AI infrastructure across bare metal, virtual machines, and Kubernetes environments. Exactly the complexity our customers are dealing with as deployments scale.

They are also a founding ISV partner of the NVIDIA AI Cloud Ready initiative, which means they are already deeply embedded in the same ecosystem we are building into. As we scale, delivery is not just about bringing GPUs online. It is about what happens after.





Provisioning, monitoring, supporting customers through increasingly complex environments. Mirantis strengthens all of that. We are already seeing it, and they will play a central role in supporting our NVIDIA AI Cloud contract.

To Alex and the whole Mirantis team, a big welcome. We're super excited to have you. So what you've heard today is a Company that has secured power at scale, is contracting revenue at scale, and is now building delivery capability at global scale. Anthony will now walk you through how we are funding that.

Anthony Lewis (Chief Financial Officer): Thanks, Dan. The capital strategy is designed to support the phased build-out of capacity Dan discussed, while maintaining flexibility and capital discipline. As of April 30, we had \$2.6 billion in cash and cash equivalents.

We expect this, together with operating cash flows, GPU financing, and additional financing initiatives to support our near-term Capex program, which includes delivery of the Microsoft contract and deployment of air-cooled capacity across Mackenzie and Childress.

For GPU Capex, we are leveraging secure debt and customer prepayments. As we have noted previously, approximately 95% of Microsoft GPU-related Capex is expected to be funded through prepayments and GPU financing. We have work streams underway for additional GPU financing to support upcoming deployments.

On the data center side, expect our financing approach to evolve as projects move from development to construction and contracting and ultimately to stabilized operations. Early stage development can be supported by balance sheet capacity and corporate level sources.

As projects reach construction and customer contracting milestones, asset and project level financing can be introduced. As assets are stabilized, refinancing and capital recycling can help support future builds. And I note we will continue to maintain a disciplined balance of debt and equity as the platform continues to scale.

I will now turn to the financial results, which continue to reflect the transition underway from Bitcoin mining to AI Cloud.

Revenue was \$144.8 million for the March quarter, compared to \$184.7 million in the prior quarter. Within that, Bitcoin mining revenue was \$111.2 million, down from \$167.4 million, driven by a lower average Bitcoin price and the ongoing decommissioning of mining hardware ahead of GPU installations.

This was partially offset by continued growth in AI cloud services revenue, which increased to \$33.6 million compared to \$17.3 million in the prior quarter. Cost of revenues decreased by \$25.9 million, primarily due to electricity costs from reduced Bitcoin mining capacity.

Net loss for the quarter was \$247.8 million, impacted by non-cash impairments of \$140.4 million, primarily related to the decommissioning of mining hardware, as well as \$23.7 million of unrealized losses related to cap calls associated with our convertible notes.





As we continue to transition our remaining Bitcoin mining operations towards AI Cloud, we expect to incur additional non-cash impairments associated with decommissioning mining hardware. These outcomes reflect the strategic reallocation of infrastructure toward AI Cloud growth, which we believe is the higher-value, long-term opportunity.

Adjusted EBITDA was \$59.5 million compared to \$75.3 million in the prior quarter, primarily on account of the revenue and cost of revenue items noted above. So as noted, the quarter reflects the ongoing transition from Bitcoin mining to growing AI Cloud.

As Dan noted earlier, we continue to target \$3.7 billion in ARR by the end of calendar 2026. We expect that ramp to be back-end weighted with Microsoft revenue and revenue from the additional 50,000 GPUs procured during the quarter, expected to begin ramping in Q3 2026. I will now turn back to Dan for closing remarks.

Daniel Roberts (Co-Founder & Co-CEO): Thanks, Anthony. Eight years ago, Will and I asked a simple question. What does the world need to build the right digital future? The answer was power, land, data centers and compute, and the ability to bring them all together at scale, faster than anyone else. Today, that thesis is playing out and we're just getting started. With that, we will open the call for Q&A.

Operator: Thank you. As a reminder, to ask a question, please press star one, one on your telephone keypad and wait for your name to be announced. To withdraw your question, please press star one, one again. Please stand by as we compile the Q&A roster. Just a moment for our first question, please. First, we have Mike Ng from Goldman Sachs. Please go ahead.

Mike Ng (Goldman Sachs, Analyst): Good afternoon, and thank you for the questions, and congratulations on all the progress. I just had two questions, if I could. First, on the five-year NVIDIA AI Cloud contract, I was just wondering if you could talk a little bit about how many GPUs are being supported by the 60 megawatts and the cost per GPU.

Then second, for Sweetwater and Oklahoma, I think you mentioned the data center capacity is coming in, in '27 and '28. I was just wondering if you could talk a little bit about at what point do those sites become marketable or maybe they already are and what milestones do you typically need to hit to increase the likelihood of a tenant being willing to take that out. Thank you very much.

Daniel Roberts (Co-Founder & Co-CEO): No problems.

Kent Draper (Chief Commercial Officer): Happy to take that one, Dan. So with respect to your first question, we haven't disclosed the specific amount of GPUs but, as we mentioned on the call, approximately 60 megawatts of air-cooled Blackwells and we think that the contract value that we're getting, and obviously the relationship that we continue to build with NVIDIA, is very beneficial coming out of that contract. Importantly, this is a managed services deployment and so it shows our ability to be able to service different segments of the market as we move forward.



With respect to your second question, as Dan mentioned earlier, we are still seeing extremely strong levels of demand within the industry, certainly outstripping supply and what we continue to see as we move forward is that capacity becomes increasingly scarce further out than people were expecting. So if we rewind even a number of months ago, '27 people were thinking that there was a relatively decent amount of capacity available. We're already seeing that capacity available in '27 is extremely scarce and that is continuing to push into '28 now as well. So for us, there is certainly the ability to market those sites for '27 and '28 online dates.

As Dan mentioned earlier, we're working through the type of customers that we bring into the mix and making sure that we are structuring the contracts in the right way to enable a flywheel at our end. But certainly the demand signals are very strong.

Daniel Roberts (Co-Founder & Co-CEO): Maybe just to add to that quickly, Kent, I think to directly answer the question, there's nothing stopping us contracting that capacity today. It just gets easier the closer you get, so the focus is on time to compute. The demand we know is there and all it does is make the conversations and the negotiations that we are having live time for a lot of that capacity much easier when you've got a defined construction and delivery plan, rather than trying to make things up on the fly in parallel with a full-form agreement.

Mike Ng (Goldman Sachs, Analyst): Thank you very much, appreciate the thoughts.

Operator: Thank you. Next, we have Paul Golding from Macquarie. Please go ahead.

Paul Golding (Macquarie, Analyst): Thanks so much and congrats on all the progress and the new relationships coming in-house. I wanted to ask about air-cooled GPUs in general. So it sounds like with the 60-megawatt deployment at Childress for NVIDIA, that will be an air-cooled deployment along with the rest of the uncontracted capacity that you're deploying across British Columbia and Texas.

Air-cooled is going to represent a meaningful part of the strategy; I just wanted to ask how you see efficiencies as well as hardware performance looking so far based on the deployments that you've planned for and how we can look at that from a financial perspective as well as we think about the model and the air-cooled opportunity. Thanks so much.

Kent Draper (Chief Commercial Officer): So in terms of efficiency and performance, I mean what we're deploying across the air-cooled portfolio is the latest generation of NVIDIA air-cooled GPUs, being Blackwells. So they perform extremely well, there is very high demand for those across all Blackwell GPU types and certainly continue to see customers finding a very good degree of performance versus cost efficiency from those units over time.

Sorry, Paul, I didn't quite understand the second part of your question in relation to how that converts into revenue over time?





Paul Golding (Macquarie, Analyst): That's right, Kent, so just wondering with retrofitting and repurposing of Bitcoin mining infrastructure for these air-cooled deployments, how that seems to be working out, maybe from a margin perspective relative to some of the liquid-cooled deployments that you're doing around the Horizon projects, just given the simpler cooling opportunity there.

Kent Draper (Chief Commercial Officer): Yes, from an operational margin perspective, it is slightly more efficient than the liquid-cooled deployments, but where we get the real benefit is, as Dan mentioned earlier, it's very capital efficient because we're taking existing air-cooled data centers that require relatively little Capex to retrofit them, compared to brand-new build, liquid-cooled facilities. So that is the major difference in terms of the two. At an operating margin level, yes, air cooled is probably slightly higher, but immaterial.

Paul Golding (Macquarie, Analyst): Thanks. If I could just sneak one more in around Europe and the Nostrum acquisition, as we think about the roadmap there, are you looking to use a similar form factor to what you've used either at Horizon or with air-cooled facilities, or is there a bespoke form factor you plan to leverage from that platform as you do the European rollout? Thanks.

Kent Draper (Chief Commercial Officer): Yes, so one of the things that attracted us to the Nostrum opportunity and we've been looking at Europe for a while, is that they did have significant land holdings that came as part of that and access to a large amount of secured power. So that gives us quite a large degree of flexibility as we build out that platform over time as to the form factor that we use.

Typically in Europe you do tend to see slightly more condensed buildouts, but we do have the ability there to utilize our typical modular design that we use across North America, which obviously may well bring construction advantages with it. So that was one of the key elements that we saw in terms of the platform that they have and the projects they've developed.

Paul Golding (Macquarie, Analyst): Thanks so much, Kent.

Operator: Thank you. Next, we have Brett Knoblauch from Cantor Fitzgerald. Please go ahead.

Brett Knoblauch (Cantor Fitzgerald, Analyst): Perfect, thanks, guys. Congrats on the, I guess, multiple acquisitions over the last week and the NVIDIA partnership and deal. I wanted to touch on Mirantis a bit because I thought that was important to the long-term story. Can you maybe just elaborate how that fits into your go-to-market motion, how it might accelerate your go-to-market motion when it comes to landing these enterprise deals, which is also what it seems like the NVIDIA partnership wants you to do as well?

Kent Draper (Chief Commercial Officer): Yes, happy to take that one initially, Dan and then you can add. So it brings with it a number of elements that we think are significantly attractive to our business. The ability to deploy quickly, the ability to service enterprise customers that may require a higher level of software over and above bare metal.





They also, as a large company that has very big internal engineering resources, bring very good capability on the software development side. That can flow through to the business, not only in terms of the software stack, but also the operations of these large clusters more generally. Further to that, again, having serviced customers for decades, they have an extremely well-built-out customer support function internally. So all of those elements are things that attracted us to the Mirantis team and are able to add to the existing skill set and customer service support that we've already built up internally.

Brett Knoblauch (Cantor Fitzgerald, Analyst): Awesome. Then if I could maybe just do a follow-up, just double clicking on the capacity ramp for '27, am I right in thinking that of the 730 megawatts, 450 will come from the remaining Childress capacity and, I guess, the 280 would be coming from Sweetwater?

Kent Draper (Chief Commercial Officer): That's correct.

Brett Knoblauch (Cantor Fitzgerald, Analyst): Thank you, guys.

Operator: Thank you. Next, we have Nick Giles from B. Riley Securities. Please go ahead.

Nick Giles (B. Riley Securities, Analyst): Yes, thank you operator. Hi, everyone. Guys, congrats on all the developments here. I know the IREN team has a lot of experience in developing infrastructure in Australia, but maybe less under the IREN platform. So I was curious if you could walk us through some of the key differences, specifically in power procurement, maybe commercial strategy, so on and so forth. Thanks.

Daniel Roberts (Co-Founder & Co-CEO): Sure. Look, in some ways, Australia is very similar to other markets and the operation of the electricity market in Australia, managed by AEMO, is very similar to what we see in Texas as ERCOT. There are markets in Australia which resemble Texas in other ways: lots of land, good transmission line capacity, good fiber connectivity and abundant renewables which isn't located close to other demand centers, similar to what we see in West Texas. So there are a lot of parallels.

The reality is, Texas is just an easier place to do business and we've been able to accelerate faster there, but it hasn't stopped us continuing to incubate projects down in Australia, and we're getting far closer to those projects becoming a bit more of the reality. I think the demand environment and the ability to service APAC and the demand constraints that we're seeing and hearing in our conversations with hyperscalers means that Australia looks like a fantastic frontier for us, and we'll look to accelerate that in parallel with North America and Europe.

Operator: Thank you. Just one moment for our next question, please. Next, we have Michael Donovan from Compass Point. Please go ahead.

Michael Donovan (Compass Point, Analyst): Hi guys, thanks for taking my question. Congrats on the progress. How should we think about regional customer mix as the platform expands? Are certain markets globally better suited for enterprise and sovereign AI customers versus hyperscalers and does that change the expected contract structure or margin profile?



Daniel Roberts (Co-Founder & Co-CEO): Look, it's going to evolve and there's a lot of unknowns around this, but if you break it down, hyperscale contract can mean two things. They can mean hyperscale is using capacity for their own purposes in terms of training and servicing workloads, such as their own AI models, or it can mean they're just acting as intermediaries to aggregate capacity for end customers that we're talking to directly. So obviously in the case of the latter, whether you're dealing with a hyperscaler or going directly to the end customer, the end demand is the same.

Then you've got different types of workloads, so inference and training. Inference is a little more latency sensitive. Training, you can probably afford a bit more latency. Indicatively, we've had conversations around training models in Australia. Yes, the USA to Australia is a long geographic distance, but it's actually not that far over fiber, particularly where you're talking about training models and given where inference sits today as well, we're all using ChatGPT or Claude, the response times are still, I guess, adjusting to the level of demand and the supply to service it.

So look, it will evolve over time and our objective is to build out an expansive ecosystem of end customers. The partnership with NVIDIA is designed around that. The Mirantis integration into our business is designed to help facilitate that over time in addition to all the near-term operational capabilities that it brings out, so the goal is very much to build out that diversified customer base over time across all of those markets.

Michael Donovan (Compass Point, Analyst): Appreciate that. A follow-up, if I may, can you help bridge the 490 megawatts in Spain from secured power to time to first token? What has to happen before construction begins?

Kent Draper (Chief Commercial Officer): So that is secured power and the sites across the portfolio there are secured as well. So, from here, it's a matter of working through final design permitting, which is already well advanced at a number of those sites and then ultimately construction of those facilities. But one of the elements that we found very attractive was the near-term security of power, so that is power that is available on a timeline that we think is going to tie in very well to general European demand. We are already seeing a number of direct requests from existing and new customers for European capacity.

Michael Donovan (Compass Point, Analyst): Great, thank you.

Operator: Thank you. Next we have John Todaro from Needham & Company. Please go ahead.

Austin Ortiz (Needham & Company, Analyst): Hi, this is Austin Ortiz on the line for John Todaro. Maybe just a quick question on how do you intend to finance the build-out for the recently announced NVIDIA deal? It seems to be around five gigawatts, so just any color on that would be helpful. Thank you.

Anthony Lewis (Chief Financial Officer): I can take that. Yes, so the Capex involved for the retrofitting of the air-cooled data centers in Childress is pretty modest in the scheme of things. In terms of the Capex for GPU, obviously we've got a range of financing sources available to us. That obviously includes initiatives at the





corporate level, but we can also look to finance GPU acquisitions in various ways in the debt capital markets, through debt capital as well. So we'll be looking at all those initiatives.

Austin Ortiz (Needham & Company, Analyst): Understood.

Daniel Roberts (Co-Founder & Co-CEO): In terms of the five gigawatts more broadly, maybe just to address that and the plan, so that's obviously a lot of capital today, but the reality is you don't need all that capital day one. There's an S-curve of construction that takes time; it takes years to deliver this. This is the whole point around time to compute. It's not just a case of getting power and land; it's assembling multi-thousand construction teams and actually delivering it. The funding for that just is progressive over time.

So as we've seen, as we continue to deliver, we continue to drive revenue, we can reinvest that revenue in Capex and it continues to unlock more and more financing sources over time. Part of the partnership with NVIDIA we've announced, they've got the ability to invest in IREN as we commission GPUs, but equally, there's other support mechanisms being discussed to the extent that we need them. But the reality is capital markets are open, they've been very supportive of our plan and we anticipate that continuing. The moment that that changes, there's a whole world of capital out there in terms of other options, whether you're creative around private markets or otherwise. When you look at the GPU financing, which is the lion's share of that Capex, the Microsoft contract is a great template. We finance 95% of that Capex at an average interest rate of about 3% through prepayments and GPU financing. So the capital is out there as long as you sign good contracts and you show that you can execute and operate this capacity.

Operator: Thank you. Just a moment, please. Next we have Joseph Vafi from Canaccord Genuity. Please go ahead.

Joseph Vafi (Canaccord Genuity, Analyst): Thanks guys. Good morning, good afternoon. My congratulations here as well on the great progress. Just a couple of thoughts or just some of your thoughts here, just to gauge demand out there. I know you threw out \$3.1 billion contracted going into \$3.7 billion contracted in ARR here, exiting the year. Your confidence in that uncontracted capacity and signing contracts, how is the demand out there for that, say, extra \$0.5 billion of ARR and what kind of clients you may be looking to bring on board there? Then I'll have a quick follow-up.

Kent Draper (Chief Commercial Officer): Yeah, I think – sorry, go ahead, Dan.

Daniel Roberts (Co-Founder & Co-CEO): Sorry, Kent. Look, again, we're trying to reiterate this as much as we can and I'm very happy for someone to point it out, but there are no idle GPUs. The prospect of then there being GPUs sitting there unused, given how structurally constrained this market is, let alone the near term, but in the medium term, it's not the focus. We are having a lot of customer conversations, but all of our operational capacity is fully contracted.



We're contracting substantial portions of capacity before it even arrives and we're in discussions with a variety of customers all the way from hyperscale clients down to AI native labs for all of that 2026 and 2027 capacity. So when a signature is put on paper, it just flows naturally. Our conviction is around the demand/supply and you cannot tap into that unless you bring the capacity online. This is the thing; a customer contract doesn't deliver revenue. Having compute online delivers revenue and that has been the focus.

Joseph Vafi (Canaccord Genuity, Analyst): Got it. Thanks for that, Dan.

Kent Draper (Chief Commercial Officer): I was going to say many of the same things. The one addition I would add is particularly for our air-cooled capacity where we are adding substantial amounts across second half of '26 and into the early part of '27, there is very significant demand on those timelines. That is the most constrained portion of the market and that is directly what is leading into the dynamic that Dan discussed, where there just are not idle GPUs that are not being used in this market. Everything on shorter-term timelines is extremely attractive to counterparties.

Joseph Vafi (Canaccord Genuity, Analyst): Got it. That's great color, Dan, Kent. Just then on your strategy and philosophy around customers and diversification there. If you are in the catbird seat here relative to fulfilling demand from multiple parties, how are you looking at broadening, deepening, diversifying that customer mix over time? Thank you very much.

Daniel Roberts (Co-Founder & Co-CEO): It's something that we're looking closely at, Joe. There is no set formula as to the proportional splits between different types of customers. There are benefits in having hyperscale clients in terms of financeability, contractual certainty, but there are also consequences in terms of price because you're not servicing the end customer in many of those instances. But the ability to service the end customer has been something we've focused on since day one.

All of our early deployments have been very focused on non-hyperscale customers and getting as close to AI natives and enterprise as we can. So the Mirantis acquisition certainly helps that. I'm not going to sit here and say we're going 100% hyperscale, we're going 100% AI native end market. The reality is that that blend will just emerge organically over time. This, again, is part of the close working relationship we've got with NVIDIA. We've spent a lot of the last fortnight in their San Jose office working through how we service all types of customers all the way from the trillion-dollar hyperscalers through to the emerging AI scale-ups where a lot of this innovation and development is taking place.

It's funny, I was speaking to someone the other day, you don't need a sales team in this market, particularly when you've got NVIDIA. They see the whole ecosystem, the introductions, the referrals, putting us in touch with anyone that needs capacity, it's just happening so organically, so quickly, live time, that it'll just play out a good way. But I think a combination of hyperscale, a hot combination of other is absolutely the goal.

Joseph Vafi (Canaccord Genuity, Analyst): Got it. Congrats. Very exciting times. Thanks, Dan.





Operator: Thank you. Our last question comes from Ben Sommers from BTIG. Please go ahead.

Ben Sommers (BTIG, Analyst): Hey yes, good afternoon and thank you for taking my questions. So, I was curious a little bit on older generation GPUs. I know you've talked in the past as you've seen the useful life of older generations for like H100s extend out further than maybe people originally thought. So I'm curious what you're seeing on the demand profile there and just what potentially type of workloads are going on to those older generation GPUs.

Kent Draper (Chief Commercial Officer): Yes, the comments that we made about no idle GPUs, that applies to all GPUs, not just latest generation. So older generations, A100s, H100s, H200s are all effectively fully utilized across the industry, so the demand picture continues to be strong. In some instances, you're actually seeing pricing for older generation units climbing significantly and there's a number of observable pricing points out there in the market where you can see that happening.

Yes, the type of demand may shift over time. You may have older generations being used more for inference, but also those older generations are equally suitable for certain types of training. So we just see strong demand across the board, both on the inference and the training side and that continues to drive demand and elongated life cycles for those older generations of equipment.

Ben Sommers (BTIG, Analyst): Great, thank you. Then just on potentially future conversations that you're having for potential contracts down the line, is there any talk of pre-payment structures similar to that of Microsoft? Just curious what you're hearing in the market on that end.

Kent Draper (Chief Commercial Officer): Yes, it certainly plays a role in a number of those conversations and we are still seeing pre-payments being on the table in a large number of instances. Now, it obviously factors in as part of the overall equation, so it's not the single factor that you're looking at. Everything has to go together with a combination of term length, prepayment, creditworthiness, price, but prepayments are certainly very much on the table in the current environment.

Ben Sommers (BTIG, Analyst): Great, thank you for taking my questions.

Operator: Thank you. I see no further questions at this time. I'll now pass to Dan for closing remarks.

Daniel Roberts (Co-Founder & Co-CEO): Thanks, operator. Thanks everyone for joining us today. We remain focused on execution, delivering the 2026 plan, advancing the 2027 build-out and positioning our now global platform for the opportunity beyond that. We look forward to updating you as we deliver. Thanks, everyone.





Forward-Looking Statements

This transcript contains “forward-looking statements” within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, that involve substantial risks and uncertainties. Forward-looking statements generally relate to future events or IREN’s future financial or operating performance. Forward-looking statements include information concerning possible or assumed future results of operations, including descriptions of our business plan and strategies, revenue targets and trends we expect to affect our business. These statements often include words such as “anticipate,” “believe,” “may,” “can,” “should,” “could,” “might,” “plan,” “possible,” “project,” “strive,” “budget,” “forecast,” “expect,” “intend,” “target,” “will,” “estimate,” “predict,” “potential,” “continue,” “scheduled”. Forward-looking statements may also be made, verbally or in writing, by members of our Board or management team in connection with this news release.

These forward-looking statements are based on management’s current expectations and beliefs. These statements are neither promises nor guarantees, but involve and are subject to known and unknown risks, uncertainties and other important factors that may cause IREN’s actual results, performance or achievements to differ materially from any future results performance or achievements expressed or implied by the forward-looking statements, including IREN’s ability to successfully execute on its growth strategies and operating plans, achieve its targeted annualized AI Cloud revenue, continue to develop its existing data center sites, design and deploy direct-to-chip liquid cooling systems, and diversify and expand into the market for high performance computing solutions (including the market for cloud services and potential colocation services, along with other important factors discussed under the caption “Risk Factors” in IREN’s Annual Report on Form 10-K, filed with Securities and Exchange Commission (the “SEC”) on August 28, 2025 and our other filings with the SEC. These and other important factors could cause actual results to differ materially from those indicated by the forward-looking statements made in this press release. Any forward-looking statement included in this press release speaks only as of the date of such statement. Except as required by law, IREN disclaims any obligation to update or revise, or to publicly announce any update or revision to, any of the forward-looking statements, whether as a result of new information, future events or otherwise.



