PUBLIC

INDUSTRY ANALYSIS

Chips Depreciate, Power Does Not

Bitcoin Miners Pivoting to AI: Who Carries the Depreciation?

PUBLISHER
Bitplanet Research Lab

DATE
2026 · 07 · 28

EXECUTIVE SUMMARY

The two models divide along a single line: who books the depreciation

WHAT HAPPENED

GPUs, the core of AI compute, fall a generation behind within 18 to 24 months, and rental rates for the prior-generation H100 fell 60 to 70 percent, from a late-2024 peak of roughly $8/hour to about $2.5–2.9/hour by mid-2026[1]. Within the same shift from mining to AI, Hut 8 secured **$16.8B** in contracted value by leasing power and buildings alone under 15-year triple-net terms, under which the tenant bears taxes, insurance and maintenance[3], while IREN bought GPUs outright and grew its owned book value from $76M to **$551M** in six months[6].

WHY IT MATTERS

The genuinely scarce and long-lived asset is not the chip but grid-connected power and the data center building. The leasing model that hands depreciating chips to the tenant (Hut 8) and the operating model that absorbs them directly (IREN) differ entirely in revenue character, debt structure and earnings volatility, and which one is superior comes down to what an investor prioritizes.

WHAT TO WATCH

How long the power shortage persists is what separates the two models. Morgan Stanley and other major institutions expect the supply-demand imbalance to run at least through 2027–2028[11]. The pace at which secured power and lease contracts convert into operating revenue is the test: of Hut 8's 9,085MW development platform, 710MW is energized[3], and industry-wide roughly 25 percent of leased capacity has been delivered[10].

$16.8B
Hut 8 contracted value ($42.8B with renewal options)
[3]

−65%
H100 rental rate decline (late-2024 peak to mid-2026)
[1]

$551M
IREN owned GPU book value (Dec 31, 2025; $76M six months prior)
[6]

2.6×
IREN quarterly D&A growth (YoY) ($47.3M→$121.2M, +22% QoQ)
[8]

710MW
Hut 8 energized capacity (of a 9,085MW platform)
[3]

CONTENTS

01 THE QUESTION

Is the chip really a melting asset?

An uncomfortable question sits at the center of the AI data center boom. GPUs, the core of AI compute, fall a generation behind within 18 to 24 months and lose value sharply. Rental rates for the prior-generation H100 fell 60 to 70 percent, from a late-2024 peak of roughly $8/hour to about

$2.5–2.9/hour by mid-2026, across multiple GPU price trackers[1]. Is this investment of tens of billions of dollars in chips a house built on fast-melting ice?

The suspicion is not idle. Satya Nadella, chief executive of Microsoft and among the largest buyers of AI chips, reportedly said on a technology podcast that he wants to avoid being locked into several years of depreciation on a single chip generation, and that the biggest competitor to a new NVIDIA chip is the previous generation (press account of the podcast, secondary citation)[2].

Yet in the same conversation Nadella reportedly said the opposite as well. What actually constrains the company is not a shortage of chips but the absence of power and warm shell, meaning data center space already fitted with power and equipment so that chips can simply be installed. Chips sit in warehouses while no energized data center is available to house them[2].

That is where this report begins. Chips do depreciate. That much is settled. But it is not the central question. The genuinely scarce and long-lived asset is not the chip but grid-connected power and the data center built on it. Chips are swapped out every two to three years, while power contracts and buildings run ten or fifteen. What an investor should weigh, then, is not whether chips depreciate but **who holds those depreciating chips on the balance sheet**. That question splits companies that describe themselves identically, as moving from mining to AI, into two entirely different kinds.

02 TWO PATHS

Two paths: colocation leasing and cloud operation

The market usually debates whether these companies are still miners or already AI businesses. That distinction is not enough. What actually separates them is who buys the chips and who carries their depreciation. On that basis the paths diverge in two (Table 1). Measured by individual contract size, Hut 8's Beacon Point ($9.8B) is the largest, followed by IREN's Microsoft contract ($9.7B) (Figure 1). A 15-year colocation lease value and a 5-year cloud contract value differ in term and revenue character, however, so comparing them on size alone has limits.

Colocation leasing. The company builds only the power and the shell, and the tenant brings in the chips. The company collects rent. It is, in effect, a landlord. The tenant supplies its own furniture, the chips, and pays operating costs and electricity as well, under what is known as a triple-net (NNN) lease. Whether the chips go obsolete is the tenant's risk. Hut 8, Cipher and TeraWulf belong to this model.

Cloud operation. The company buys and runs the chips itself and sells the resulting compute capacity, measured in GPU hours. Here it is not a landlord but the owner of an internet cafe. High

utilization raises earnings, but when the machines go obsolete the owner takes that loss. The depreciation stays on the company's books. IREN falls into this category.

The landlord and internet cafe comparison has its limits, though. Stability does not make the landlord superior. In exchange for absorbing depreciation, the direct operator holds what the lessor cannot: the operating know-how to run GPU clusters efficiently, the software that manages data centers and allocates resources, direct customer relationships and the operating data they generate, and the upside of capturing the entire compute margin rather than a fixed rent. A landlord cannot earn beyond the contracted rent, while an operator keeps the excess when demand surges. In short, the leasing model is built to reduce risk, and the operating model to hold the technology stack and the margin in full. Which is the better structure is taken up later.

The difference between the two models shows up clearly in the accounts. The claim that power does not depreciate holds cleanly, in accounting terms, only for the leasing model. The landlord sells long-lived assets and never takes hold of chips that go technically obsolete at speed. The direct operator carries that obsolescence risk instead. In return it takes the entire compute margin those assets earn, rather than a rent. Which model is superior is not settled by accounting alone.

	Colocation leasing	Cloud operation
Representative firms	Hut 8, Cipher, TeraWulf	IREN
Chip (GPU) ownership	Held by the tenant	Owned and debt-funded by the company
Depreciation risk	Transferred to the tenant	Retained on the company's books
Revenue form	15-year triple-net rent (take-or-pay)	GPU hours × utilization
Debt matching	Long-dated debt against long-lived assets (16.5 years)	Debt against chips with short technical lives
Analogy	Landlord	Internet cafe owner

Table 1. Core distinction between the two models. Source: Bitplanet Research.

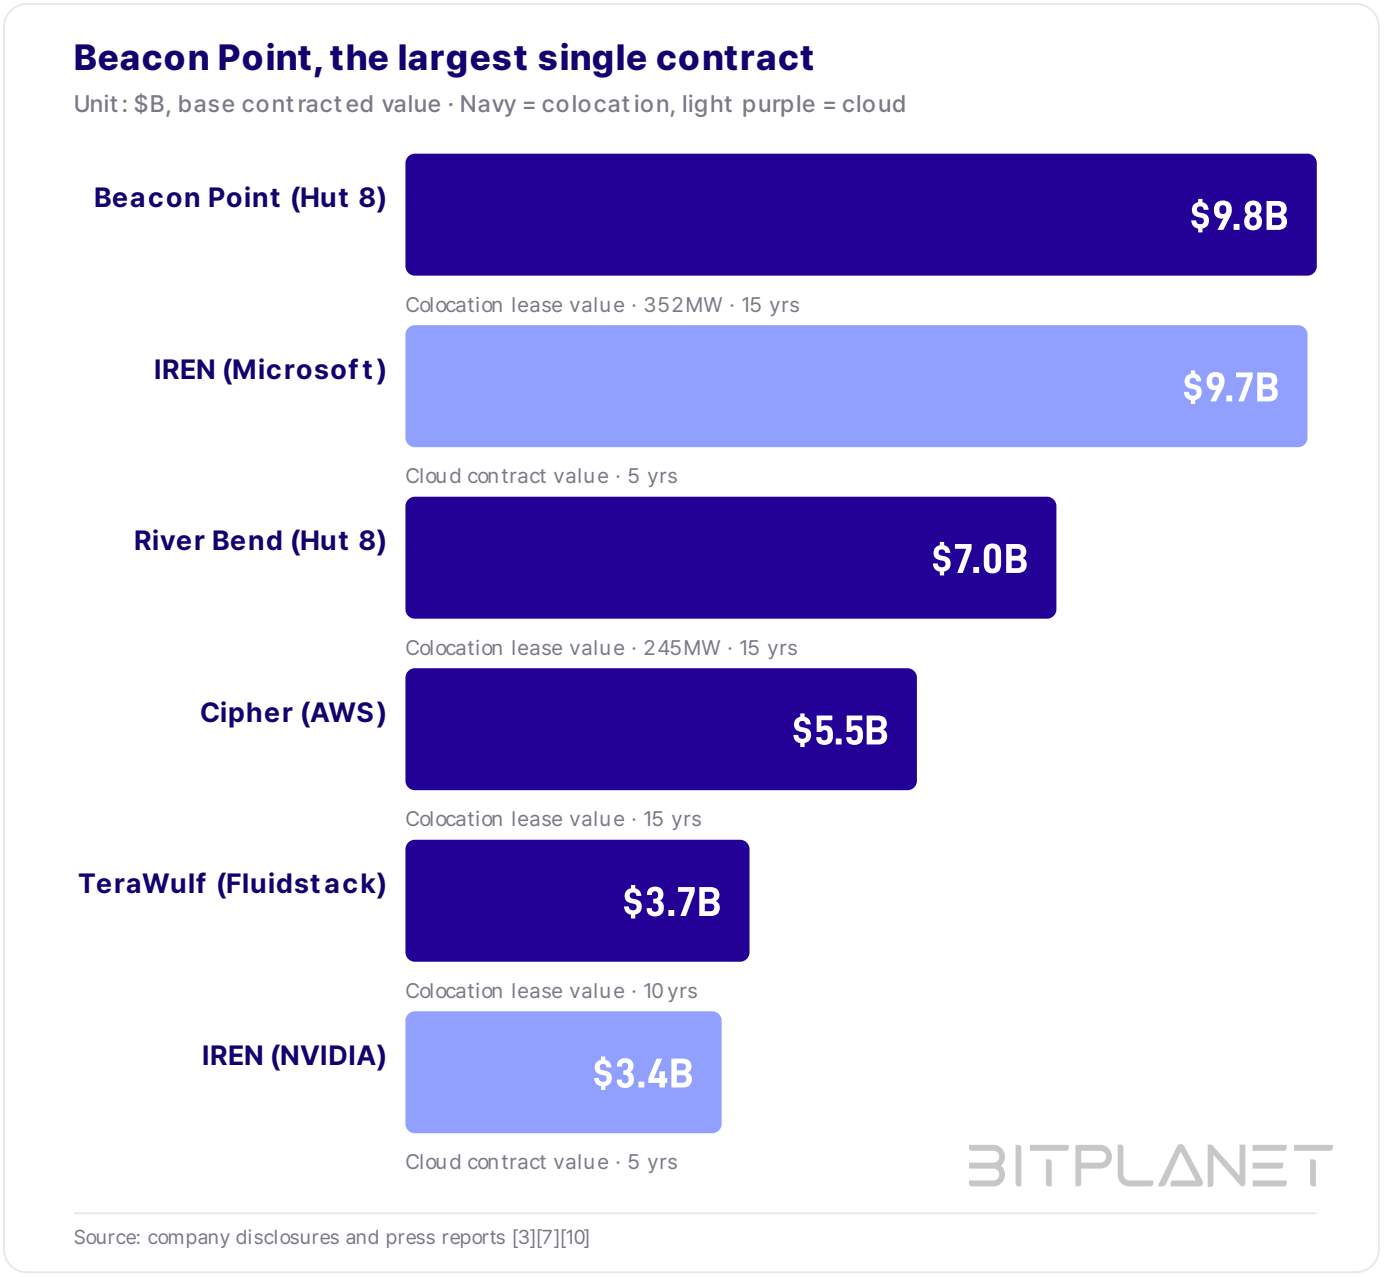

03 MODEL A

Colocation leasing: passing depreciation to the tenant

Hut 8 is the textbook case for the leasing model.

A triple-net lease is one in which the tenant bears all three net items: taxes, insurance and maintenance. In an ordinary lease the landlord covers these, while under triple-net terms the tenant takes them on entirely. The landlord effectively rents out land, building and power and collects a fixed rent on schedule. In a data center the tenant also brings in the most expensive and fastest-ageing equipment of all, the chips. Nowhere on Hut 8's balance sheet is there a depreciating chip.

Hut 8 has signed leases at two sites to date. Beacon Point in Texas (352MW) and River Bend in Louisiana (245MW) are both 15-year triple-net leases, and both tenants sit at the upper end of investment grade. Base contracted lease value across the two totals **$16.8B**, on 597MW of IT capacity[3]. Reading that figure requires looking at delivery and energization alongside it. Beacon Point has commercialized its first phase, while River Bend is under construction with initial delivery scheduled for the second quarter of 2027[3]. The $16.8B is contractually fixed, but most of it is not yet operating.

If tenants exercise every renewal option embedded in the contracts, total contract value could rise to **$42.8B** (Figure 2)[3]. That figure is not committed revenue but potential value conditional on renewal, and must be read separately from the base contracted value.

Tenant credit reinforces the model as well. The Beacon Point tenant is an unnamed technology company that Hut 8 describes only as an investment-grade firm with a market capitalization above $1T. River Bend is leased by the cloud start-up Fluidstack, behind which stands an investment-grade company reported in the press to be Google, which has backstopped the leased capacity[4]. The credit standing behind the rent is therefore big-tech grade.

From a $16.8B base to $42.8B including renewal options

Unit: $B (contracted lease value) · As of May 6, 2026 · River Bend 245MW, Beacon Point 352MW

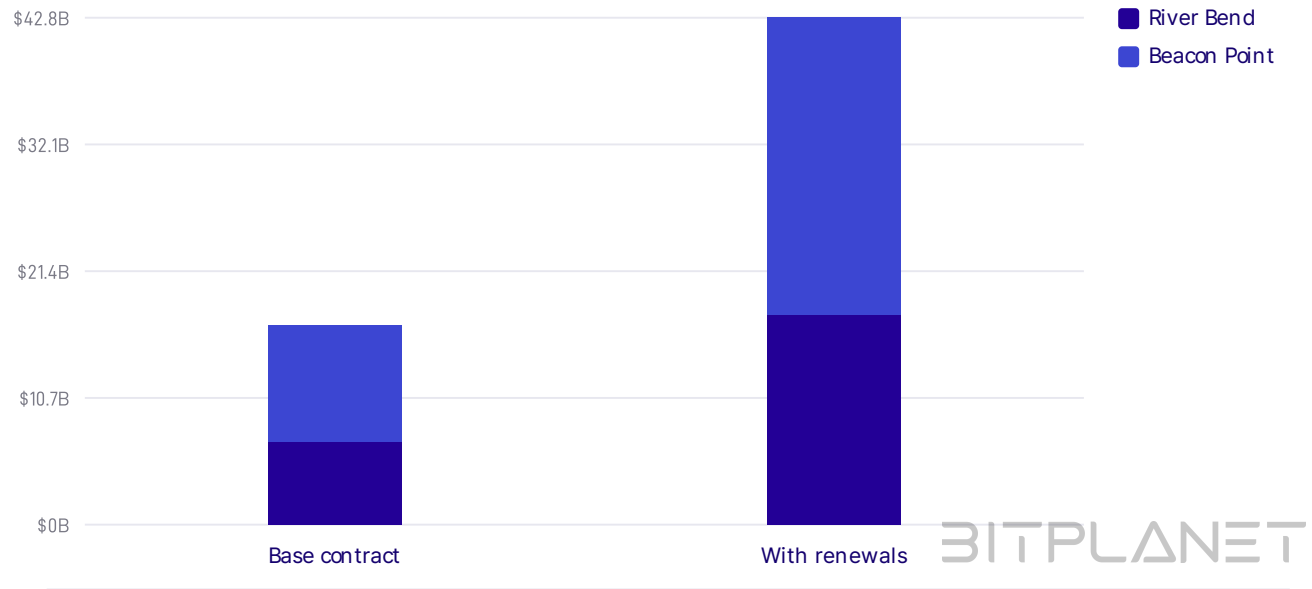

Source: Hut 8 Q1 2026 results materials [3]

The funding structure reveals the core of this model. Hut 8 financed River Bend construction with **$3.25B of fully amortizing, investment-grade senior secured notes maturing in 16.5 years**. The structure is non-recourse to the parent and non-dilutive. The maturity is what deserves attention. Debt with a 16.5-year term was matched against a 15-year lease and against long-lived building and power assets. Long-dated debt was set against assets that last, not against chips that lose value quickly. The reason this was possible is simple. **Hut 8 did not buy the chips.**

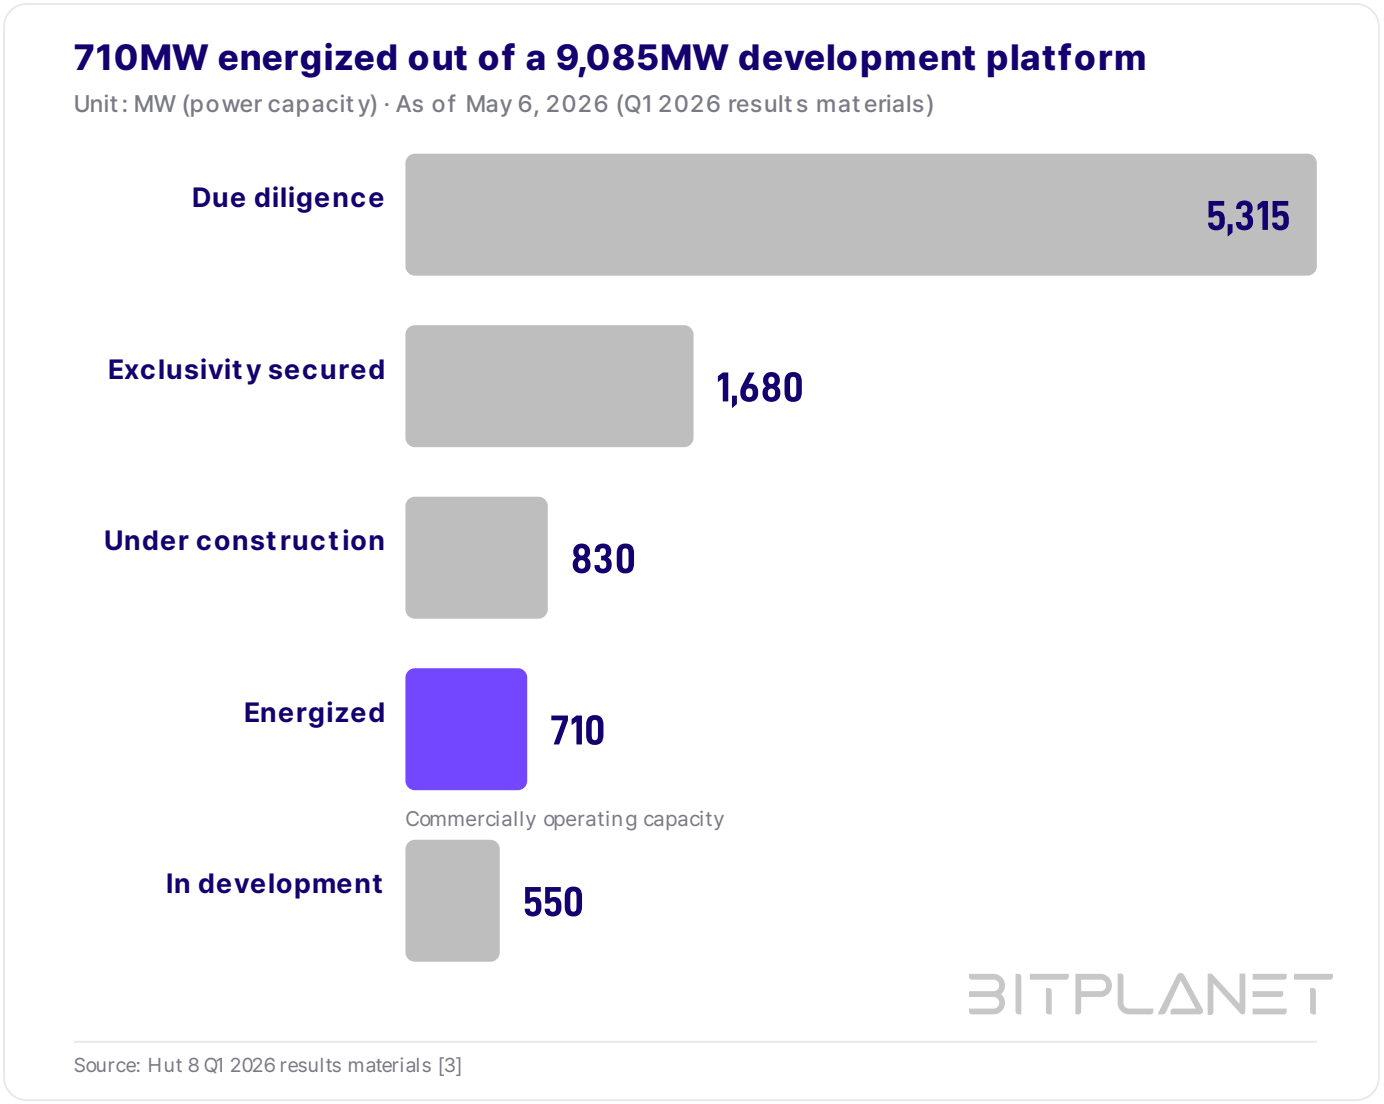

The model has its shadows, and there are two of them.

First, most of it is still a promise. Hut 8's development platform exceeds 9GW, but only 710MW is energized and commercially operating. The remaining 8.4GW sits in due diligence, exclusivity, development and construction. Securing power is one thing; energizing it, signing a tenant and collecting rent is another.

Second, bitcoin is a variable. Hut 8 runs mining and bitcoin holdings alongside the leasing business. The company posted a net loss of $253.1M in the first quarter of 2026, and the mark-to-market loss on its bitcoin holdings exceeded that net loss[5]. Even a leasing model that looks clean sees its income statement move when the bitcoin price does. Section 5 returns to this point.

04 MODEL B

Cloud operation: absorbing depreciation directly

IREN took the opposite path.

IREN buys GPUs outright and rents out the resulting compute under the name AI Cloud. Its customers are Microsoft and NVIDIA. Like an internet cafe owner, it can earn more than a rent when utilization is high. In exchange it takes the loss directly when the machines go obsolete.

IREN's GPU assets on a gross basis grew from **$76.0M on June 30, 2025 to $550.6M on December 31, 2025**[6]. Quarterly depreciation and amortization reached **$121.2M** in the third quarter of fiscal 2026, ended March 31, 2026, about 2.6 times the $47.3M of the prior-year quarter and roughly 22 percent above the $99.2M of the preceding quarter[8]. The multiple looks dramatic year over year, but the sequential comparison of plus 22 percent is the better read on the current trend. This D&A covers all property and equipment, including mining rigs, data centers and buildings, not GPUs alone, so the growth in GPU book value cannot be treated as the sole cause of the entire increase (Figure 4).

GPU book value up from $76M to $551M in six months

Unit: $M (GPU hardware, gross book value before accumulated depreciation) · As of June 30 and December 31, 2025

Source: IREN Form 10-Q, Note 9 [6]

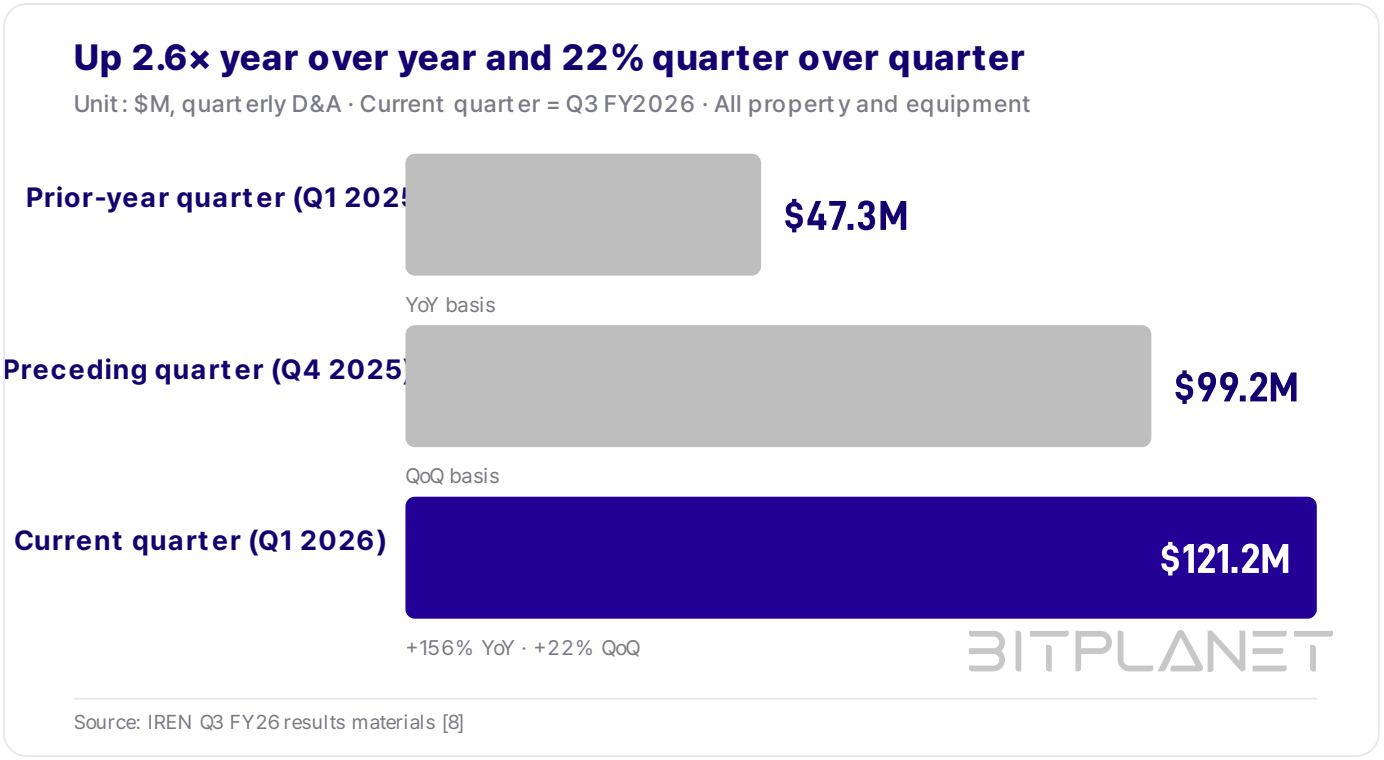

What is notable is that IREN's large impairments come not from GPUs but from older mining rigs. At the end of 2025 IREN impaired its S21 Pro and T21 bitcoin miners, citing early replacement in the course of the AI transition, and shortened their useful lives to the middle and second half of 2026[6]. No large impairment has yet come from GPUs, but IREN has already demonstrated how shortening an asset's accounting life worsens earnings.

Why, then, do IREN's earnings swing so widely? The nine-month net loss of $18.6M in fiscal 2026 looks small not because operations improved but because of a **net gain of $533.9M on financial instruments**. That gain arises from remeasuring, at market prices each quarter, the derivatives attached to five series of convertible notes maturing between 2029 and 2033. Those derivatives were contracted in advance to limit share count growth from the convertibles, and their carrying value moves with the share price. As the share price fell well below the conversion price, the carrying value of those contracts moved sharply, and that movement drove the reported result. It is a number that comes from the capital structure, far removed from operations.

The customer base also deserves a closer look. IREN's two pillars are Microsoft and NVIDIA, and the NVIDIA arrangement is structurally unusual. NVIDIA is at once (1) a customer that ordered $3.4B of cloud capacity, (2) an investor holding warrants to purchase up to $2.1B of IREN stock, and (3) the supplier of the chips going into that cloud[7]. It sells the chips, invests in the company and then rents the chips back, an arrangement in which the customer, investor and supplier roles overlap. Under that structure, contracted revenue cannot be read on the same basis as purely external demand. The Microsoft contract is different in character, being purely external demand.

IREN ARR, much of it pre-operational or assumption-based

Unit: $B, annualized run-rate revenue · IREN Q3 FY26 · Only Microsoft is operating

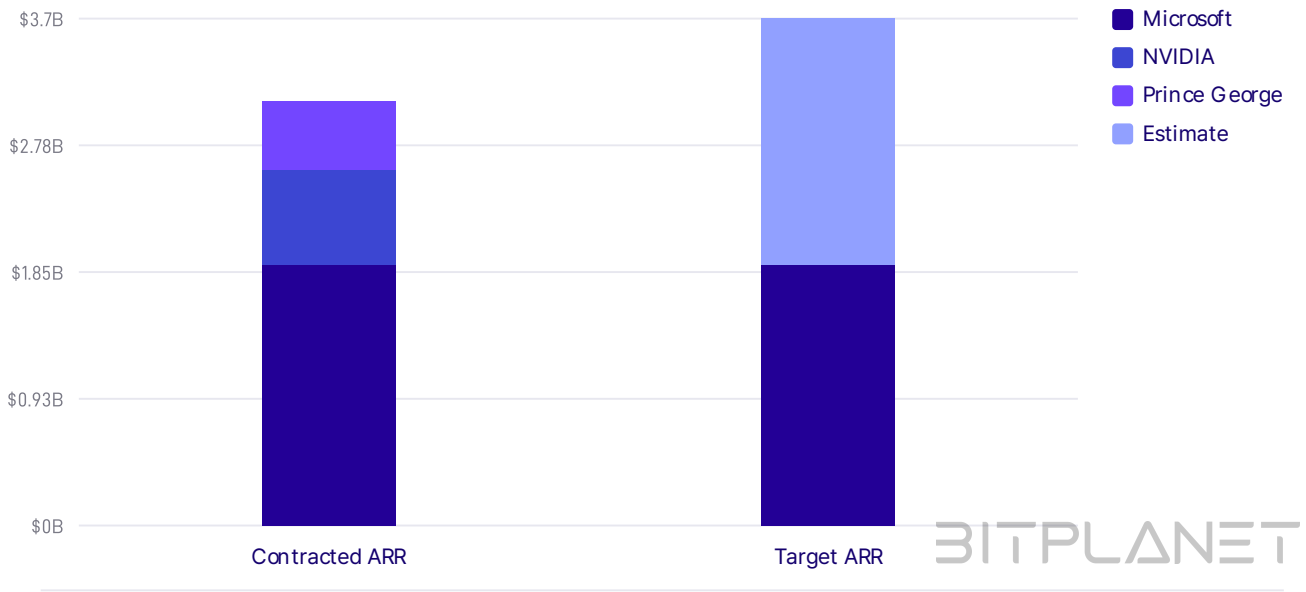

Source: IREN Q3 FY26 materials [8]

The annualized run-rate revenue (ARR) the company presents also breaks down unevenly. Contracted ARR of $3.1B consists of Microsoft ($1.9B, operating), NVIDIA ($0.7B, ramping in early 2027 and therefore contributing close to nothing today) and Prince George (about $0.5B, still pipeline). Of the $3.7B target ARR, the $1.8B outside Microsoft rests on an internal assumption of deploying roughly 74,000 GPUs and is an estimate rather than a signed contract (Figure 5)[8].

Balance: IREN is also stepping toward the leasing model

IREN is not moving in one direction only. The company has said it will diversify its revenue base beyond AI Cloud into colocation data center development. It is shifting weight toward the leasing model that passes depreciation to the tenant.

05 HEAD TO HEAD

Which model holds the stronger hand?

Set the two models side by side (Table 2).

	Colocation leasing	Cloud operation
Strengths	No depreciation on the books · 15-year rent from big-tech-grade credit · Long-dated debt matched to long-lived assets	Large upside · Captures the entire compute margin rather than a rent
Weaknesses	Capped upside · Only 710MW energized of 9GW (execution risk) · Bitcoin holdings swing earnings	Holds obsolescing chips and 6% debt directly · Effectively two customers · Overlapping roles in the NVIDIA arrangement (customer, investor, supplier)
In short	Holds only long-lived assets and passes depreciation on	Absorbs depreciation and reaches for the upside

Table 2. Strengths and weaknesses by model. Source: Bitplanet Research.

The strengths of the leasing model are clear. No depreciating chip sits on the books, rent runs steadily for 15 years from a high-credit tenant, and long-dated debt can be matched to that rental stream. Its weaknesses come from the same structure. Upside is capped at the rent, execution risk remains in carrying a 9GW pipeline through to energization and lease-up, and in Hut 8's case bitcoin holdings amplify earnings volatility.

The strength of the direct operating model is that its upside is uncapped. It takes the entire compute margin rather than a rent, so earnings scale further when GPU economics hold. Its weaknesses sit on the other side. The company itself carries chips that go technically obsolete along with the 6 percent debt raised to buy them[9], has effectively only two customers in Microsoft and NVIDIA, and in one of those the customer, investor and supplier roles overlap, so that contract must be distinguished from purely external demand.

Reliance on a small number of customers, however, must be applied to the leasing model just as evenly. Hut 8's $16.8B of contracted value is likewise concentrated in two tenants, the unnamed technology company at Beacon Point and Fluidstack at River Bend, and each is a 15-year contract with a single counterparty. If having two customers counts as a risk for IREN, then Hut 8's dependence on two contracts and two tenants, and the fact that one of those tenants relies on a backstop for its credit, deserve the same weight. Both models share concentration in a few large counterparties and renewal risk at contract maturity.

How long will the current advantage last? The primary reason the colocation model enjoys deep demand and high margins today is that power is scarce. That scarcity will not unwind quickly. Transmission expansion takes years in permitting alone, and lead times for high-voltage transformers have stretched to three to five years. Morgan Stanley and other major institutions expect the supply-demand imbalance to run at least through 2027–2028[11]. North American colocation vacancy is at effectively record lows, around 1 percent, and more than 90 percent of

capacity under construction is already pre-leased[14]. Hut 8's advantage in securing power early therefore has grounds to persist for at least several more years.

INTERPRETATION What comes next is the harder question. This looks less like a structural moat than a premium created by a temporary imbalance. Once the grid catches up in the late 2020s, bring-your-own-power spreads, or hyperscaler capital spending normalizes, the scarcity of power plus shell as a product diminishes and rental bargaining power may converge toward ordinary real estate economics. Colocation demand is also concentrated in single-tenant leases with a few hyperscalers, which carries renewal cliff and counterparty concentration risk[12]. Hut 8's two contracts are not exempt, and with only 710MW energized of 9GW (Figure 3), converting the pipeline into rent requires demand to hold through the delivery dates, including River Bend in the second quarter of 2027. The test is not whether power has been secured but whether it is turned into operating revenue[10].

On the other side, it is worth weighing whether the demand growth IREN expects in return for absorbing depreciation will actually materialize. Demand for operators that own GPUs outright and rent compute capacity, the so-called neoclouds, is growing quickly. CoreWeave, the leader among them, reported a backlog of roughly $100B on its first-quarter 2026 results released on May 7, 2026, has passed 1GW of energized power and targets 8GW by 2030[13]. Backlog and target figures like these can shift materially each quarter, so the reporting date matters. The industry as a whole is short of power rather than short of demand: as of March 2026, compute leases signed between hyperscalers and neoclouds over the preceding six months exceeded $100B, and North American colocation vacancy stood at effectively record lows[14]. As the center of gravity in AI workloads shifts from training to inference, the higher-margin compute service layer served by operators that hold chips, software and customers together is itself expanding. IREN's bet is an attempt to take that layer directly, where margin and demand are structurally moving. That holds only if utilization supports it and the risks of customer concentration in Microsoft and NVIDIA and of the overlapping-role arrangement are cleared.

To summarize, the power bottleneck does make Hut 8's leasing model the more profitable one with deeper demand at present. Whether that advantage would still surpass IREN once the bottleneck eases is harder to assert. If the premium shifts from holding power to delivering compute as power becomes plentiful, today's ranking could reverse. Which model prevails depends heavily on a timing variable: how long, and how severely, power stays scarce.

One trait runs through both models. Neither company's GAAP net loss comes from operations; both come from marks and impairments (Figure 6). Hut 8's net loss of $253.1M in the first quarter of 2026, ended March 31, stemmed largely from a $295.7M mark-to-market loss on bitcoin[5]. IREN's net loss of $247.8M in the third quarter of fiscal 2026, ended on the same date, was driven by $140.4M of impairment on mining and data center equipment and $121.2M of depreciation[8]. Reading the bottom line alone as evidence that both businesses are failing would be wrong.

Distinguishing the periods matters as well. The $247.8M above is a **single-quarter** net loss for IREN and rests on a different period from the nine-month cumulative net loss of $18.6M cited earlier. IREN's quarterly results ran from a gain of $384.6M to losses of $155.4M and $247.8M, summing to roughly the negative $18.6M nine-month figure[8]. The common item behind the large first-quarter gain and the two subsequent losses is the mark on financial instruments. Over nine months that mark was a net gain of $533.9M, large enough to narrow the cumulative loss to $18.6M, while in the third quarter alone the same item was a $23.7M unrealized loss that widened the quarterly deficit. The same item points in opposite directions depending on how the period is cut. Both the nine-month gain that improved IREN's bottom line and the third-quarter loss that widened it come from the capital structure rather than from operating performance.

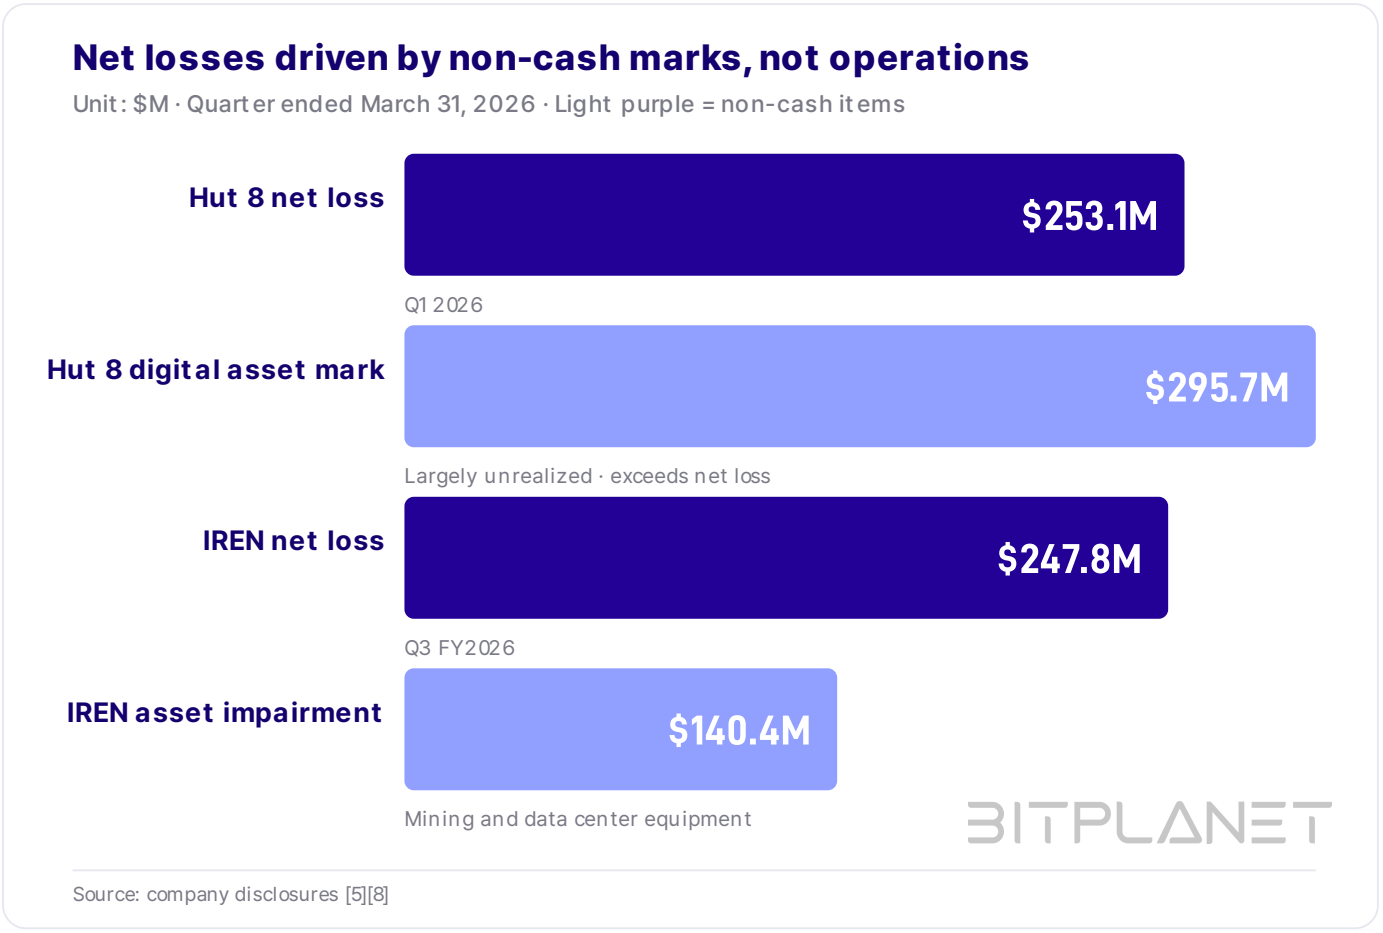

The distinction between the two models is what matters in the end. The genuinely long-lived assets are power and buildings, while the asset exposed to rapid technical obsolescence is the chip. Even so, declaring either side the outright winner is difficult, because the answer turns on what an investor wants. **For those who prefer a stable evolution, using large capital to secure infrastructure in power and buildings and passing depreciation risk to creditworthy tenants under long contracts, colocation leasing (Hut 8) is closer to the answer; for those who prefer a technology-led evolution, holding the chips directly and capturing both the compute margin upside and the full technology stack, cloud operation (IREN) is.** That IREN, which holds the chips itself, has begun diversifying into colocation shows that the boundary between the two models is being adjusted rather than fixed.

A diagnosis from the US asset manager VanEck summarizes the point. Valuing miners as AI infrastructure comes down to two things: energized power in megawatts, and tenant credit. Companies holding leases command more than 10x their energized power, while those with only a pipeline receive 2–6x (Table 3, Figure 7). Industry-wide, only about 25 percent of leased AI capacity has actually been delivered, and a near-term funding gap of $50B remains[10]. The next premium comes not from winning contracts but from energizing the power already secured.

Key contract comparison

Company	Model	Contract	Capacity	Term	Tenant (credit)
Hut 8	Colocation	Beacon Point (TX) $9.8B	352 MW	15 yrs	Unnamed big tech (investment grade)
Hut 8	Colocation	River Bend (LA) $7.0B	245 MW	15 yrs	Fluidstack (Google-backed)
Cipher	Colocation	AWS $5.5B	300 MW	15 yrs	AWS (investment grade)
TeraWulf	Colocation	Fluidstack $3.7B	200+ MW	10 yrs	Fluidstack + Google backstop
IREN	Cloud	Microsoft ~$9.7B	200 MW	5 yrs	Microsoft (investment grade)
IREN	Cloud	NVIDIA $3.4B	~60 MW	5 yrs	NVIDIA (customer, investor, supplier)

Table 3. Major contracts by model, size, term, capacity and tenant credit. Source: company disclosures and VanEck[10].

Critical IT load larger in colocation than in cloud contracts

Unit: MW, critical IT load · Navy = colocation, light purple = cloud

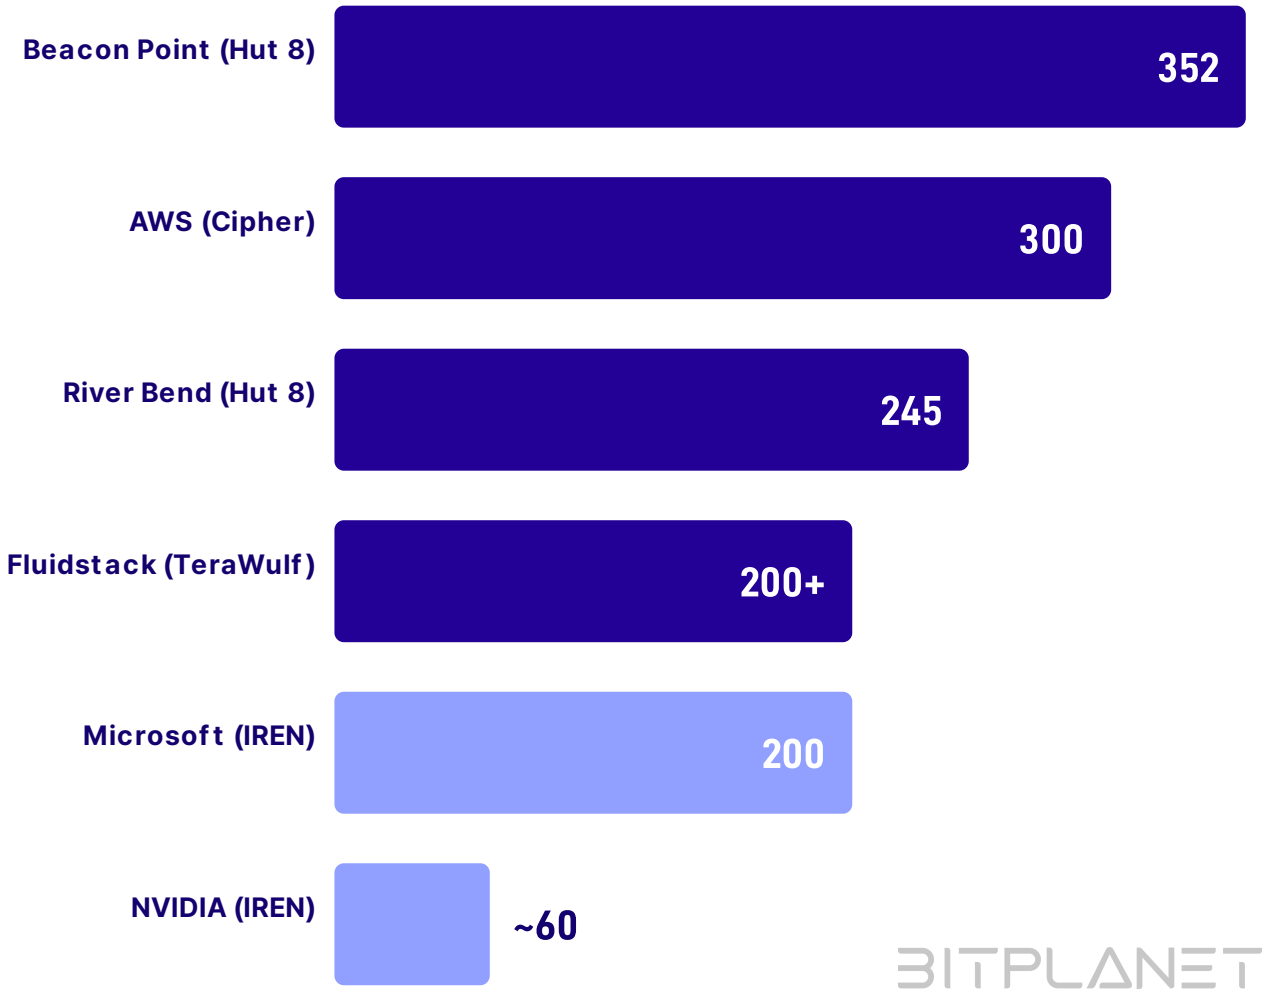

Source: company disclosures and SEC filings, VanEck [3][7][10]

06 KOREA

Power is the bottleneck in Korea too, but the lessor carries the depreciation

Power decides outcomes in Korea's data center business as well. A contracted power application is the process by which a data center asks KEPCO in advance for grid capacity. Applied volume rose from 60MW before 2020 to 3,091MW in 2023, a rise of more than 50x in three years. Over the same period the number of applications went from 2 to 47[15]. Supply, by contrast, is blocked. The power grid impact assessment is a scheme that reviews in advance whether the grid can absorb new demand. On a cumulative basis through March 2026, first-stage technical reviews in the Seoul metropolitan area covered 522 applications and 33,592MW. Of those, 279, or 53.4 percent, were ruled unable to be supplied. Of the 243 that cleared the first stage, 24 advanced to full review and

10, totaling 1,010MW, received final approval. That is an approval rate of 1.9 percent by application count and 3 percent by capacity[16] (Figure 8).

Scarcity pushed prices up and pulled contracts forward. Metropolitan colocation rents rose from about KRW 140,000 per kW per month in 2019 to about KRW 250,000 in 2025, a gain of more than 70 percent over six years. Leased and pre-leased share of supply reached 99.7 percent in 2024 and 99.4 percent in 2025. Even capacity not yet completed has been taken early, at 16 percent for 2027 and 23.5 percent for 2028. Global cloud providers account for about 170MW, or 54 percent, of leasing demand. Domestic firms including Naver, Kakao, Coupang and NHN take 110MW, or 34 percent, and together the two groups make up 88 percent[16]. The pattern matches the pre-leasing concentration seen in North America.

Open the contracts, however, and the direction differs. Korean colocation leases are frequently fully fitted, with the lessor supplying the power and cooling equipment alongside the building. Under that structure, equipment depreciation and reinvestment obligations arising from technology change remain with the lessor. That is the reverse of triple-net terms, under which the tenant bears capital expenditure and maintenance (secondary citation; primary report not obtained)[16]. The timescales diverge as well. A building serves 30 to 40 years, while a GPU generation turns over in two to three. Power draw per rack is also rising from 15 to 25kW to 40 to 100kW for AI servers. Equipment that is current at completion may be unable to serve demand within five to ten years[16].

Policy is changing, but the power problem remains. The Special Act on the Promotion of the AI Data Center Industry (Act No. 21759) was promulgated on June 9, 2026 and takes effect on March 10, 2027. It includes exceptions to the power grid impact assessment for non-metropolitan data centers and consolidated permitting, among other measures. Exempted capacity and applicable criteria, however, were left to the enforcement decree and are not yet set[17]. A proposed exception allowing direct liquefied natural gas (LNG) power purchase agreements was dropped from the final bill[15].

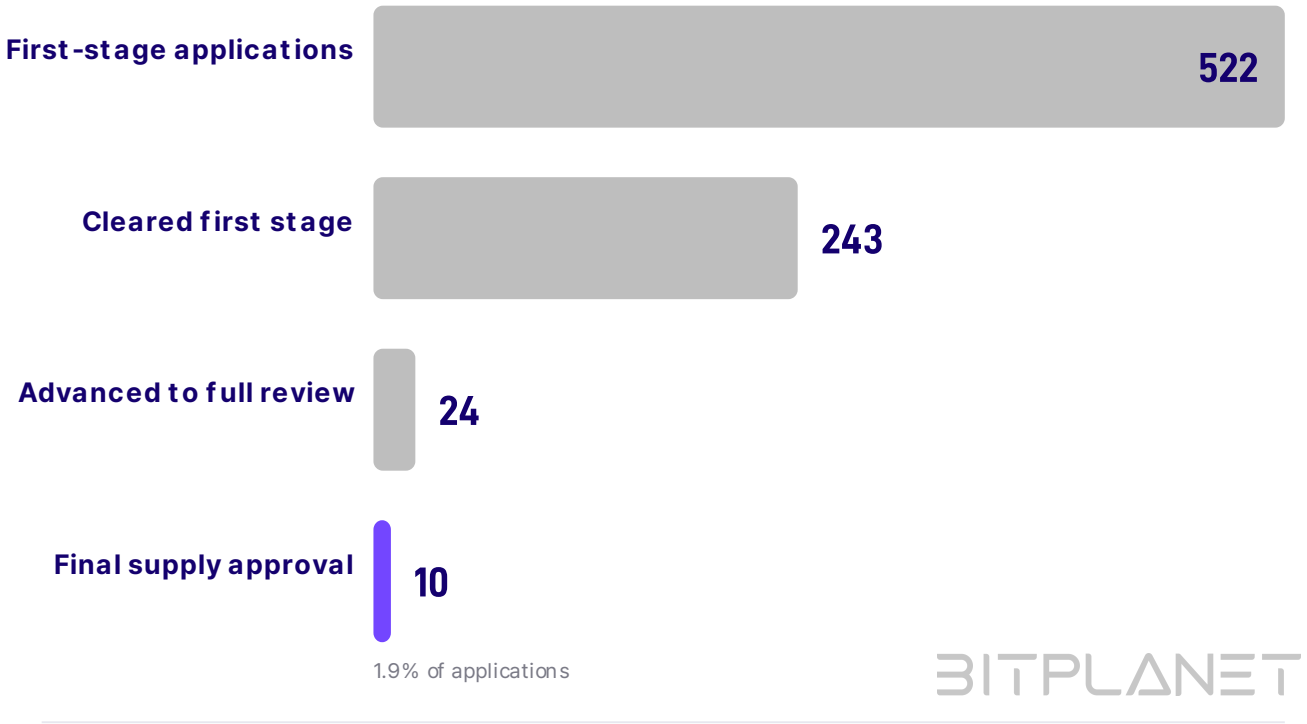

Source: CBRE Korea materials [16]

INTERPRETATION Korea has virtually no listed company that has moved from bitcoin mining to AI. What this section establishes is therefore not a transition path for miners. It extends only to the point that power constraints govern where data centers sit and how their contracts are written. What deserves attention is that the side carrying the depreciation is the reverse of North America. The Korean lessor holds power as a scarce asset and absorbs equipment depreciation at the same time. Applying this report's two-model distinction to Korea requires checking the contract structure first. Which domestic operator has the better strategy is outside this report's scope. The issuer conducts digital asset and AI infrastructure business in Korea, so this section overlaps with the issuer's own operating environment.

C CONCLUSION

Conclusion: who is holding the chip?

Chips melt and power does not. The opening question, whether chips depreciate, was therefore the wrong one. Of course they do. What matters is **who is holding the chip** when it goes obsolete. The path taken by those who, like landlords, hold only power and shell and pass depreciation to the tenant (Hut 8, Cipher) diverges

from the path of those who, like the internet cafe owner, absorb the chips and the debt directly (IREN), even within the same shift toward AI.

Nadella's remark, though, is the starting point of the question rather than this report's answer. The power bottleneck he emphasized, the chip depreciation he guarded against, and ultimately who absorbs that depreciation are what separated the two models. Which is right depends on how long power stays scarce. While power is dear, the leasing model that passed depreciation on (Hut 8) leads; the moment that scarcity eases, the premium moves from those who hold power to those who run compute. There is no single answer to the question, only a matter of what to prioritize.

Key conclusion

The genuinely long-lived assets are power and buildings, not chips. Whether to pass depreciation to the tenant and take steady rental income (Hut 8) or to absorb it and reach for compute margin upside (IREN) is not a question with one fixed answer, and the outcome depends heavily on a timing variable: how long power stays scarce.

R REFERENCES

References

#	Title · URL	Publisher	Date
[1]	Analysis of GPU useful life and H100 rental rate declines. The 60–75% fall from peak is cross-checked by multiple GPU price trackers (Silicon Data H100 Rental Index, CloudZero, Spheron and others) https://www.stanleylaman.com/signals-and-noise/gpus-how-long-do-they-really-last	Stanley Laman Group and others	2025-11
[2]	Press account of Nadella remarks on avoiding depreciation and the shortage of warm shell (press report of podcast remarks; secondary citation, original timestamp not directly verified) https://techcrunch.com/2025/11/03/altman-and-nadella-need-more-power-for-ai-but-theyre-not-sure-how-much/	TechCrunch	2025-11-03
[3]	Hut 8 Q1 2026 results materials and 8-K: contracted value $16.8B (597MW IT), $42.8B with renewals, 9,085MW platform (710MW energized), Beacon Point commercialized, River Bend delivery Q2 2027, $3.25B senior secured notes Primary https://www.sec.gov/Archives/edgar/data/0001964789/000110465926055894/hut-20260506xex99d1.htm	Hut 8 · SEC	2026-05-06

#	Title · URL	Publisher	Date
[4]	Hut 8 IR materials (River Bend 1GW ROFO expansion note) and VanEck sector note (June 2026) **Primary** https://www.sec.gov/Archives/edgar/data/0001964789/000110465926019393/hut-20260225xex99d1.htm	Hut 8 · SEC	2026-02-25
[5]	Hut 8 Q1 2026 results materials: net loss $253.1M (including roughly $295.7M of bitcoin mark-to-market loss), revenue $71.0M **Primary** https://www.sec.gov/Archives/edgar/data/0001964789/000110465926055894/hut-20260506xex99d1.htm	Hut 8 · SEC	2026-05-06
[6]	IREN Form 10-Q (as of December 31, 2025), Note 9 property and equipment: GPU hardware gross book value $76.0M (June 30, 2025) → $550.6M (December 31, 2025), shortened miner useful lives and impairment **Primary** https://www.sec.gov/Archives/edgar/data/0001878848/000187884826000015/iren-20251231.htm	IREN · SEC	2026-02
[7]	IREN Secures $3.4bn AI Cloud Contract with NVIDIA (five-year managed cloud plus warrants over 30M shares at $70, up to $2.1B) **Primary** https://iren.gcs-web.com/news-releases/news-release-details/iren-secures-34bn-ai-cloud-contract-nvidia	IREN	2026-05-07
[8]	IREN Q3 FY26 results materials (quarter ended March 31, 2026): quarterly net loss −$247.8M, impairment $140.4M, quarterly D&A $121.2M (prior year $47.3M, preceding quarter $99.2M), quarterly net result path +$384.6M → −$155.4M → −$247.8M, nine-month net loss −$18.6M (net mark gain +$533.9M, Q3 mark loss −$23.7M), contracted ARR $3.1B, of the $3.7B target ARR $1.8B assumes roughly 74k GPUs deployed **Primary** https://www.sec.gov/Archives/edgar/data/0001878848/000187884826000025/irenreportsq3fy26results.htm	IREN · SEC	2026
[9]	IREN Closes $3.65bn Investment-Grade GPU Financing (blended 6.00%, secured on GPUs and Microsoft cash flows) **Primary** https://www.globenewswire.com/news-release/2026/06/01/3304211/0/en/IREN-Closes-3-65bn-Investment-Grade-GPU-Financing.html	IREN · GlobeNewswire	2026-06-01
[10]	A Framework for Valuing Bitcoin Miners as AI Infrastructure (about 25% of leased capacity delivered, near-term gap of roughly $50B, leases at 10×+ versus pipeline at 2–6×) **Primary** https://www.vaneck.com/us/en/blogs/digital-assets/matthew-sigel-a-framework-for-valuing-bitcoin-miners-as-ai-infrastructure/	VanEck (Sigel, MacMaster)	2026-06
[11]	Energy Markets Race to Solve the AI Power Bottleneck (power constraints into 2027–2028, rise of bring-your-own-power) **Primary** https://www.morganstanley.com/insights/articles/powering-ai-energy-market-outlook-2026	Morgan Stanley	2026
[12]	The AI Bubble: Hidden Risks and Opportunities (reliance on single hyperscaler leases, renewal cliff and duration risk) **Primary** https://www.man.com/insights/the-ai-bubble	Man Group	2026

#	Title · URL	Publisher	Date
[13]	CoreWeave Q1 2026 results 8-K: backlog of roughly $100B, more than 1GW of energized power, 8GW target by 2030 Primary https://www.sec.gov/Archives/edgar/data/0001769628/000176962826000220/coreweave1q26earningspress.htm	CoreWeave · SEC	2026-05-07
[14]	AI Data Center Build Advances at Full Speed (compute leases above $100B, record-low colocation vacancy, more than 90% of capacity under construction pre-leased) Primary https://about.bnef.com/insights/commodities/ai-data-center-build-advances-at-full-speed-five-things-to-know/	BloombergNEF	2026-03
[15]	Korean data center contracted power applications (60MW before 2020 → 3,091MW in 2023; 2 → 47 applications) and how the LNG direct power purchase exception was dropped from the AI Data Center Act. Press report citing Ministry of Climate, Energy and Environment data (secondary citation) https://www.mt.co.kr/tech/2026/05/23/2026052210211399740	Money Today (머니투데이)	2026-05-23
[16]	Korea Data Centre Investment report: metropolitan power grid impact assessment approval rate of 1.9%, colocation rents up more than 70% over six years, pre-leasing share, depreciation attribution under fully fitted leases, physical obsolescence. Accessed via press coverage as the primary report was not obtained (secondary citation) https://www.kharn.kr/news/article.html?no=31267	CBRE Korea · KHARN (콜드체인뉴스)	2026-07-12
[17]	Special Act on the Promotion of the AI Data Center Industry (Act No. 21759), promulgated June 9, 2026, effective March 10, 2027. Basis for non-metropolitan power grid impact assessment exceptions, consolidated permitting and special district designation Primary https://www.law.go.kr	Korean Law Information Center (국가법령정보센터)	2026-06-09

A APPENDIX

Data limitations and formulas for our own estimates

① Limitations

- **Beacon Point tenant identity:** Hut 8 describes the tenant only as an investment-grade technology company with a market capitalization above $1T, so the identity is not disclosed[3].
- **River Bend backstop provider:** The company backstopping Fluidstack's leased capacity is reported in the press to be Google, but this is a secondary citation and no confirmation from the parties has been obtained[4].

- **Nature of IREN's target ARR:** Of the $3.7B target ARR, $1.8B is an estimate resting on an internal assumption of roughly 74,000 GPUs deployed rather than a signed contract. Within the $3.1B contracted ARR, the NVIDIA portion of $0.7B ramps in early 2027 and contributes close to nothing today[8].
- **Pipeline execution risk:** Of Hut 8's 9,085MW development platform, 710MW is commercially operating. Whether secured power converts into lease revenue depends on demand holding through the delivery dates, including River Bend in the second quarter of 2027[3][10].
- **Asymmetric sample:** The colocation model is examined through three cases in Hut 8, Cipher and TeraWulf, while the direct operating model rests effectively on IREN alone. Statements about how the direct operating model behaves are therefore shaped heavily by IREN's characteristics, and generalizing them across the model has limits. CoreWeave and others are cited only as industry context.
- **Source of the analytical frame:** The axis of valuing miners as AI infrastructure by energized power in megawatts and tenant credit is borrowed from VanEck[10]. This report adds the question of who books the depreciation to distinguish the two models.
- **Secondary nature of the backstop:** The company said to support Fluidstack's credit at River Bend, Google, rests on press reporting as secondary information, and confirmation from the parties has not been obtained[4].

② Formulas for our own estimates

Not applicable (this issue). No figure in this report is labeled as our own estimate. The multiples and changes cited in the text, including roughly 2.6 times year over year, plus 22 percent quarter over quarter and the 60 to 70 percent decline, along with the check of quarterly results summed against the nine-month cumulative figure, are all derived by dividing or adding disclosed figures from the cited sources, without separate price or exchange rate assumptions.

PUBLISHED BY Bitplanet Research Lab WRITTEN BY Taewon Kim REVIEWED BY Sooyoung Kim

Conflict of interest disclosure The issuer, Bitplanet, conducts digital asset and AI infrastructure business and holds a view of power and compute infrastructure as core assets. The central argument of this report, that the genuinely long-lived assets are power and buildings rather than chips, may align with the issuer's own business narrative, and readers should weigh that in forming their judgment. The issuer also holds bitcoin, so the earnings volatility from bitcoin marks identified at Hut 8 in this report may apply similarly to the issuer. As of the date of writing, the issuer holds no equity stake in and has no contractual relationship with the companies discussed in this report, including Hut 8 and IREN.