


Neoclouds



Source: iStock/aleksandarvelasevic

Neoclouds are fast-growing companies that specialise in leasing clusters of AI chips to firms needing computing power at short notice. Positioned at the very centre of the AI boom, they have emerged as critical enablers of innovation. As traditional cloud hyperscalers face capacity constraints and struggle to keep up with surging demand, neoclouds are seeing rapid growth. However, their data centre expansion is often financed through debt, making them highly controversial investments at the heart of the “AI bubble” debate.

Jakub Dubaniewicz

Senior Equity Analyst

jakub.dubaniewicz@syzgroup.com

Introduction

At the start of the year, investors seeking AI exposure were limited to mega caps such as Nvidia, Microsoft, or Alphabet. As the boom accelerates, however, new AI infrastructure specialists – neoclouds - have joined the market. CoreWeave listed in March 2025, Nebius shares started earlier in September 2024 rose to prominence in summer this year. Several listed crypto miners (Iren, Hive, Riot) are pivoting towards AI datacentre services too.

Neoclouds are reshaping both the technology and investment landscape. Hyperscalers still dominate, but their legacy architectures are not always ideal for GPU-heavy AI workloads. Neoclouds, built specifically for accelerated computing, can deploy capacity faster and with greater flexibility.

For investors, neoclouds offer a pure-play on the AI data centre cycle. Their growth is impressive, but capex requirements are immense, forcing sustained capital raising. That makes them simultaneously attractive, controversial, and volatile.

In recent quarters, neoclouds have reported triple-digit revenue growth, sold-out capacity, and huge backlogs tied to multi-year

hyperscaler contracts. Yet the debate remains heated: the same growth is funded by heavy leverage, widening losses, and dependence on rapidly evolving technology (i.e. what’s the useful life of an AI chip?). Share prices have surged in last 12 months, but swings are dramatic, reflecting a sector still defining its long-term economics.

The following Q&As explain what neoclouds do, why they matter, and the key issues investors should consider as this young sector evolves.

Q&A

Q: What exactly are neoclouds?

A: They are cloud providers built specifically for high-performance AI computing. Rather than offering broad cloud services, neoclouds focus on delivering large clusters of the latest GPUs for training and deploying advanced AI models. Their narrow specialisation allows speed, efficiency, and rapid scaling.

Q: Who are the leading players?

A: Table below lists the largest and most notable neocloud companies:

Company	Market cap / Valuation	Key contracts & partnerships	Notes
CoreWeave NASDAQ: CRWV	\$45bn	Meta (\$14.2bn), OpenAI (\$22.4bn),	Largest neocloud with \$55.6bn backlog (as of Q3 2025). Revenue expected to rise from \$4bn in 2025 to \$9.4 bn in 2026. Nvidia owns 7% stake, sell chips to Coreweave with priority access, and provides the revenue backstop guarantee (buys back unsold capacity).
Nebius NASDAQ: NBIS	\$24bn	Microsoft (\$19bn over 5 years), Meta (\$3 bn)	Spun out of Yandex, now a major GPU supplier with strong hyperscaler ties.
IREN NASDAQ: IREN	\$15bn	Microsoft (\$9.7bn over 5 years)	Former crypto miner. Recently signed a \$9.7bn deal with Microsoft. Now guiding to \$3.4bn revenue run-rate by end-2026.
Hive Digital TSXV: HIVE	\$1bn	None disclosed	no big-ticket contracts yet but building capacity expected to support \$1.3–4.5bn in annual revenue.
Lambda Labs Private	Valued at \$2.5bn (last funding round Feb 2025)	Microsoft (largest client), Nvidia partnership	GPU-focused provider with Nvidia key technology partner and Microsoft the top client (deal value undisclosed). Reportedly pursuing next private funding round at \$4–5bn.
Crusoe Private	Valued at \$10 bn (last funding round Oct 2024)	OpenAI / Stargate ecosystem	Not a pure neocloud but key player using stranded energy to power GPU datacentres.

Q: How big is the neocloud market?

A: Still tiny relative to the hyperscalers – less than 2% of global cloud revenue – but expanding at extraordinary speed. Growth often re-rates overnight when multi-year deals are announced. CoreWeave’s contracts with Meta and OpenAI, Nebius and Iren both signed long-term agreements with Microsoft. show how crucial these firms have become in relieving the industry’s GPU bottlenecks.

Q: Why are neoclouds growing so fast?

A: Demand for cutting-edge GPUs far exceeds what hyperscalers can deliver. AI workloads require specialised, dense, power-intensive infrastructure that traditional cloud environments were not designed around. Neoclouds, built “AI-first”, can deploy new GPU clusters in days. That ability to fill supply gaps has driven explosive revenue growth: CoreWeave +124% YoY, Nebius +117%, IREN +304% in the latest quarter.

Q: How neoclouds differ from traditional cloud providers?

A: Cloud hyperscalers were built around CPU-based, general-purpose computing. These systems were never designed for the intense, parallel workloads AI demands.

Neoclouds flip that model: they offer rapid, “bare-metal” access to the latest graphic processing units (GPUs), simple pricing, and near-instant scalability. Their infrastructure is streamlined for speed and cost efficiency, without the layers of legacy software or complex billing structures that slow down hyperscalers. For AI developers, that means faster model training, greater flexibility, and fewer headaches.

Q: What are neoclouds particularly good at?

A: In short, neoclouds excel at three critical things:

1. **Securing access to the newest GPUs**, often through priority agreements. While hyperscalers try to design own custom chips to reduce reliance on Nvidia, neoclouds remain fully GPU-centric, making them reliable, high-volume customers and preferred partners.
2. **Efficiency**: their streamlined operations and advanced cooling systems allow dense GPU packing and high utilisation.
3. **Financial innovation**: neoclouds have developed financial structures – typically combining asset-backed financing and revenue-backed credit – that allow rapid scaling in a capital-intensive business.

Q: Who uses neoclouds?

A: AI-heavy organisations that need instant GPU access: start-ups, research labs, gaming firms, and large enterprises running major AI projects. They benefit from rapid scaling, straightforward pricing, GPU availability, and potentially lower costs.

Even hyperscalers such as Microsoft or Meta rely on neoclouds to cover short-term capacity gaps. Their long-term commitments highlight the strategic role neoclouds play in the AI ecosystem.

Q: How large is their capital expenditure?

A: Very large. Neocloud capex should reach around \$60 bn in 2025 (JP Morgan’s estimate), up more than 120% YoY. While still far below the combined \$300bn+ capex of Tier 1 hyperscalers, it exceeds what those same hyperscalers were spending annually only a few years ago. The scale underlines both the opportunity and the risk.

Q: Why are crypto miners pivoting to neocloud business?

A: Because their existing facilities – GPU-heavy infrastructure with abundant power – are nearly ideal for AI workloads. As crypto mining economics worsened due to halving events, rising electricity and operational costs, repurposing for AI became a natural evolution rather than a reinvention. CoreWeave, IREN, and Hive exemplify this shift.

Q: How are neoclouds funding their colossal investment needs?

A: Through a mix of equity and significant secured debt, often backed by GPU fleets and long-term customer contracts. CoreWeave has leaned heavily on secured debt, leveraging its GPU fleet and multi-billion-dollar client contracts as collateral. Nebius has supplemented its financing with both senior debt and substantial equity offerings. Usually, capital raising tends to be tied to massive commercial wins.

Q: Should investors be concerned about neoclouds’ rising debt?

A: Yes, but context matters.

Large amount of expensive debt makes neoclouds vulnerable to any dip in utilisation or a slowdown in chip availability. Take CoreWeave as an example: its \$15bn debt load looks steep against its 2025 EBITDA estimate (\$2.5bn), but less so against its 2026 forecast (\$6.5bn). Its borrowing costs are high, with unsecured notes priced around 9%, compared with single-digit rates for established hyperscalers such as Microsoft or Oracle.

Take-or-pay agreements with credit-worthy anchor clients (Microsoft, Meta, OpenAI) provide predictable revenue streams. This helps unlock private credit and JV financing. The model allows rapid growth, but it also raises questions about leverage, interest costs, and dilution. The big debate is whether future margins and utilisation can support such capital intensity, especially as AI chips depreciate quickly. Management commentary has been reassuring, but scepticism will persist until the model proves itself over several cycles.

Q: Is Oracle turning into the largest neocloud, or the smallest hyperscaler?

A: Both descriptions fit. While Oracle runs a traditional cloud platform, its OCI unit has embraced the neocloud model, investing heavily in Nvidia GPU clusters and selling raw GPU hours to major clients, including OpenAI.

Oracle’s advantage is its balance sheet: Oracle can pursue the same strategy as neoclouds but with far cheaper funding, thanks to its highly profitable and cash generative enterprise software division. It is, in effect, applying the neocloud playbook from the financial position of an established software giant.

Oracle’s recent share price reflects this shift. The stock rallied sharply after announcing a large, multi-year revenue backlog tied to OpenAI, only to sell off as investors questioned cloud margins, dependence on OpenAI, the heavy need for debt financing, and the near-term negative FCF outlook. In effect, Oracle’s shares behaved more like those of a neocloud start-up than a mature big tech company.

Welcome to Syzerland®

For further information

Banque Syz SA

Quai des Bergues 1
CH-1201 Geneva
T. +41 58 799 10 00
syzgroup.com

Jakub Dubaniewicz

Senior Equity Analyst
jakub.dubaniewicz@syzgroup.com

This marketing document has been issued by Bank Syz Ltd. It is not intended for distribution to, publication, provision or use by individuals or legal entities that are citizens of or reside in a state, country or jurisdiction in which applicable laws and regulations prohibit its distribution, publication, provision or use. It is not directed to any person or entity to whom it would be illegal to send such marketing material.

This document is intended for informational purposes only and should not be construed as an offer, solicitation or recommendation for the subscription, purchase, sale or safekeeping of any security or financial instrument or for the engagement in any other transaction, as the provision of any investment advice or service, or as a contractual document. Nothing in this document constitutes an investment, legal, tax or accounting advice or a representation that any investment or strategy is suitable or appropriate for an investor's particular and individual circumstances, nor does it constitute a personalized investment advice for any investor.

This document reflects the information, opinions and comments of Bank Syz Ltd. as of the date of its publication, which are subject to change without notice. The opinions and comments of the authors in this document reflect their current views and may not coincide with those of other Syz Group entities or third parties, which may have reached different conclusions. The market valuations, terms and calculations contained herein are estimates only. The information provided comes from sources deemed reliable, but Bank Syz Ltd. does not guarantee its completeness, accuracy, reliability and actuality. Past performance gives no indication of nor guarantees current or future results. Bank Syz Ltd. accepts no liability for any loss arising from the use of this document.