


Estimated price: between USD \$30,000 and \$50,000 each.

The H100 chip sells between USD \$30,000 and \$50,000 each. This is the bedrock for all superscalers that purchase from Nvidia, and they are usually sold as part of a system that can cost between USD \$2.8m to \$3.4m. However, the first quarter of 2026 saw larger incoming orders for the newer chips H200 (upwards of USD \$40,000) and Blackwell B200 (starting at USD \$50,000).

Inside the H100 GPU casing:



The AI Financial Ecosystem: Fragile or Foundational?

Posted on 10 September 2026 by Johndn7

Artificial intelligence may prove to be the defining technological development of our generation. The productivity improvements, scientific discoveries and new industries it could create are difficult to quantify because the technology remains at such an early stage of adoption. Yet the extraordinary expectations surrounding AI have already produced something that can be measured: one of the largest investment cycles in financial history.

Hundreds of billions of dollars are flowing into semiconductors, data centres, cloud infrastructure and AI companies. NVIDIA estimates that four trillion US dollars could ultimately be invested (combined effort) in AI infrastructure before the end of the decade (conservatively speaking), while the five largest hyperscalers are expected to spend extraordinary amounts expanding computing capacity. The important question for investors is increasingly not simply how much money is being spent, but where that money is coming from and where it subsequently goes. The largest technology companies (for example Microsoft, Amazon, Alphabet) are investing in AI developers by way of cloud computing purchases, AI developers purchase NVIDIA

processors, and NVIDIA itself is investing billions into companies that ultimately require its hardware for example Microsoft, Amazon, Alphabet). One company's expenditure is increasingly becoming another company's revenue, which can then help justify further investment back into the same financial ecosystem.

Against this extraordinary concentration of capital, precious metals provide an interesting comparison. At the time of writing gold trades at approximately AUD \$6,123 per ounce, silver AUD \$93.77 per ounce and platinum AUD \$2,620 per ounce. These markets operate very differently from the increasingly interconnected AI investment system.

What is circular trading?

Circular trading traditionally describes securities or other assets being repeatedly bought and sold between a group of parties, potentially creating the appearance of greater market activity, liquidity or demand than genuinely exists. In its more problematic forms, coordinated circular trading can become a form of market manipulation because essentially the same economic interest is repeatedly transferred without genuine independent demand. Moreover, every bubble offers historical examples of its use to make money for a small group of people.

That is **not** an accurate description of what the major AI companies are doing. Real capital is being invested, real processors are being manufactured, real data centres are being constructed and real computing services are being consumed. The developing AI ecosystem is better described as circular financing, reciprocal investment and commercial interdependence. The concern is not that transactions are fictitious, but that the investors, customers and suppliers increasingly overlap, limiting the stakeholders.

Consider the relationship between Microsoft, NVIDIA and Anthropic. In November 2025, NVIDIA committed to invest up to USD \$10 billion in Anthropic while Microsoft committed another USD \$5 billion. Anthropic simultaneously committed to purchase USD \$30 billion of Microsoft Azure computing capacity, initially using up to one gigawatt of capacity powered by NVIDIA systems. Microsoft therefore invests in Anthropic, Anthropic spends money with Microsoft, and Microsoft purchases the infrastructure required to provide that computing capacity, including NVIDIA technology (who invested in Anthropic).

OpenAI provides an even larger example. In February 2026, OpenAI announced USD \$110 billion in new investment, including USD \$50 billion from Amazon, USD \$30 billion from NVIDIA and USD \$30 billion from SoftBank. OpenAI simultaneously expanded its relationship with Amazon and secured next-generation NVIDIA computing infrastructure. By the end of March, the funding round had grown to USD \$122 billion in committed capital. Amazon's relationship is particularly illustrative because it is simultaneously an investor and infrastructure provider, while OpenAI has committed to consume substantial Amazon Web Services (AWS) computing capacity. These are legitimate strategic partnerships with obvious commercial logic. Nevertheless, they create a web of relationships in which the distinction between investor, supplier and customer is becoming increasingly blurred.


Does circular AI investment obscure true demand?

This does not mean AI demand is artificial. OpenAI has reported enormous consumer adoption, while cloud providers continue to report expanding AI workloads. NVIDIA's latest results similarly indicate demand extending beyond the largest hyperscalers into AI laboratories, enterprises and sovereign customers. There is therefore considerable evidence that genuine demand exists.

The more difficult question is how much of today's extraordinary AI demand is ultimately being driven by end users, and how much is being supported by capital circulating between the companies building the AI ecosystem. If a technology company invests billions into an AI developer, which subsequently uses part of that capital to purchase computing capacity from the investor or hardware produced by another investor, genuine transactions and revenue are generated; however, some of that demand is being supported by capital originating within the same relatively concentrated ecosystem. This distinction becomes increasingly important as valuations depend upon continued revenue growth. Hyperscaler capital expenditure creates revenue for semiconductor manufacturers such as NVIDIA. NVIDIA's rapidly rising revenue reinforces expectations of extraordinary AI demand, supporting valuations across all semiconductor and AI-related companies. Those valuations can improve access to capital, which finances additional infrastructure and investment, creating further orders for processors and cloud capacity. The cycle can therefore become self-reinforcing.

There is nothing inherently unsustainable about this if end-user demand ultimately grows sufficiently to justify the current infrastructure buildout. Railways, electricity networks and the internet all experienced periods when infrastructure investment ran ahead of immediate demand, yet the infrastructure subsequently supported enormous economic development. The vulnerability appears if investment continues expanding faster than the revenues ultimately available from consumers outside the investment loop. A second vulnerability exists in the fact that, unlike railways, electricity networks and the internet where infrastructure can last multiple decades, the life span of the infrastructure for an AI buildout has a lifespan of about two to five years maximum. That is the lifespan of the extraordinary amount being invested into the AI space currently.

Cash and Equivalents to Total Assets Across AI Hyperscalers



How important is NVIDIA to the AI financial ecosystem?

Few companies illustrate the scale of the AI boom better than NVIDIA. Once primarily associated with graphics processors, NVIDIA has developed into the dominant supplier of the accelerated computing infrastructure underpinning modern AI. Its processors, networking equipment, systems and software now form



part of the infrastructure used by major cloud providers, AI laboratories, enterprises and governments. The company's financial growth has been extraordinary. For the three months ending 26 July 2026, NVIDIA reported revenue of USD \$96.2 billion, an increase of 106% from the corresponding period a year earlier. Data Centre revenue alone reached USD \$89 billion, increasing 117%. NVIDIA expects approximately USD \$108 billion in revenue for the following quarter. Its GAAP gross margin reached 75%, while quarterly net income was almost USD \$59.7 billion.

NVIDIA also reported that the cloud industry's backlog now exceeds USD \$2 trillion, while capital expenditure among the five largest hyperscalers is expected to approach USD \$800 billion in 2026 and USD \$1.3 trillion in 2027. NVIDIA is therefore not merely selling processors into the AI boom. It occupies a critical position through which an enormous proportion of the industry's infrastructure expenditure passes. Yet NVIDIA is increasingly doing something else: **helping to finance the ecosystem that buys its technology**. Its investments in companies including OpenAI and Anthropic mean NVIDIA is not simply benefiting from rising AI expenditure but is also providing capital to businesses whose expansion requires enormous amounts of computing infrastructure. NVIDIA's latest financial disclosures also demonstrate that its investment activities have become substantial alongside its operating business.

This creates an unusual position. NVIDIA can be a hardware supplier to a cloud provider, an investor in an AI company using that cloud provider and a beneficiary when increased AI demand causes the cloud provider to purchase additional NVIDIA infrastructure. Each transaction has a legitimate commercial purpose, but collectively they demonstrate why NVIDIA has become integral to the web of relationships supporting the AI capital cycle.


NVIDIA's H100 Flagship GPU



Estimated price: between USD \$30,000 and \$50,000 each.

The H100 chip sells between USD \$30,000 and \$50,000 each. This is the bedrock for all superscalers that purchase from Nvidia, and they are usually sold as part of a system that can cost between USD \$2.8m to \$3.4m. However, the first quarter of 2026 saw larger incoming orders for the newer chips H200 (upwards of USD \$40,000) and Blackwell B200 (starting at USD \$50,000).

Inside the H100 GPU casing:



What happens if one part of the AI system weakens?


This is where interdependence becomes vulnerability. The Magnificent Seven and other leading AI companies are often discussed as separate investment opportunities, yet their earnings are increasingly connected. One company's capital expenditure becomes another company's revenue, while one company's investment can finance another company's expenditure. Things can become intense when you consider that, if revenue is soft for a particular quarter, pre-selling can change the figures on a cash basis even though goods and services have not been deployed.

Imagine that a major AI developer discovers that customer revenue is growing considerably more slowly than anticipated. It may respond by reducing planned computing expenditure. The cloud provider expecting that demand could then reduce future data-centre expansion. Lower infrastructure spending means fewer orders for NVIDIA processors, networking equipment and associated components. NVIDIA's expected revenue growth could subsequently be revised lower, affecting its valuation and potentially sentiment towards semiconductor companies more broadly. If investors interpret that slowdown as evidence that expected AI returns

have been overestimated, valuations across cloud providers, data-centre operators and AI-related companies could also adjust. The circularity that reinforces growth can therefore operate in reverse.

This matters beyond technology because these companies now occupy an extraordinary position within global equity markets. As discussed in our [previous analysis](#), more than half of the increase in S&P 500 market capitalisation between 2021 and 2025 came from the Magnificent Seven, while concentration among the largest US companies has reached historically elevated levels. A serious repricing of AI expectations would consequently not remain confined to a handful of technology stocks. It could affect major sharemarket indices, retirement portfolios, international investors and confidence across other risk assets.

The Circulatory Flow of Capital Within AI-Orientated Companies



Examples of recent deals:

- ① In November 2025, NVIDIA invested USD \$10 billion and Microsoft USD \$5 billion into Anthropic. Anthropic simultaneously committed to purchase USD \$30 billion of Microsoft Azure computing capacity, initially using up to one gigawatt of capacity powered by NVIDIA systems.
- ② In February 2026, OpenAI announced USD \$110 billion in new investment, including USD \$50 billion from Amazon, USD \$30 billion from NVIDIA and USD \$30 billion from SoftBank. While Amazon committed USD \$50 billion to OpenAI, OpenAI simultaneously expanded its existing agreement with Amazon Web Services by another USD \$100 billion over eight years and secured next-generation NVIDIA computing infrastructure.

● Investment ● Infrastructure ● Services



Can AI change the world without every investment succeeding?

Perhaps this is the most important distinction for investors. The debate should not be reduced to whether AI is a bubble or whether AI will change the world. Both questions can have very different answers. A transformative technology does not guarantee that every company developing it will justify its valuation, nor does a market correction invalidate the technology itself. The vulnerability in today's AI ecosystem comes from concentration and interdependence. NVIDIA's extraordinary growth depends partly upon hyperscalers continuing their unprecedented infrastructure spending. Hyperscalers increasingly depend upon AI developers and customers using that infrastructure. AI developers depend upon continuing access to enormous amounts of capital and compute. Investors, meanwhile, have already assigned substantial valuations on the expectation that the resulting revenue will eventually justify the expenditure, but will it? Pioneering companies do the heavy lifting, but that does not necessarily equate to success.

For precious metals investors, the contrast is unusually clear. Gold and silver do not need a hyperscaler to expand capital expenditure, another company to raise billions of dollars or a customer to purchase enough computing capacity to validate an earlier investment. Physical bullion has no management team directing capital into another company, no earnings forecast dependent upon a strategic partner and no interconnected corporate balance sheet through which financial stress can migrate. That does not make precious metals immune to market volatility, nor does it diminish the extraordinary potential of artificial intelligence. It simply places them at the opposite end of the structural spectrum. The AI financial ecosystem is becoming more centralised, concentrated and mutually dependent as enormous amounts of capital circulate through a relatively small group of companies. Physical precious metals remain decentralised assets whose existence and ownership are independent of that corporate network.

If AI fulfils its promise, the companies building it may create extraordinary wealth. If expectations falter, however, the same web of relationships that has accelerated the boom could equally transmit disappointment with remarkable efficiency. The greatest strength of the AI investment cycle may be how effectively its participants reinforce one another. Its greatest vulnerability may prove to be exactly the same thing.