




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Data Center GPU Pricing 2026: The Full AI Pricing Index

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-07-20 · data center gpu pricing · nvidia h100 price · gpu cloud pricing · h100 vs h200 vs b200 · gpu rental cost per hour · ai infrastructure cost · gb200 nvl72 pricing · gpu shortage 2026 · cloud gpu comparison

Original article: <https://intuitionlabs.ai/articles/data-center-gpu-pricing-2026>



In this article

1. [Data Center GPU Pricing 2026: The Full AI Pricing Index](#)
 2. [Executive Summary](#)
 3. [Introduction and Background](#)
 4. [The Cloud GPU Rental Market in 2026: Pricing by Provider and GPU Model](#)
 5. [Hardware Acquisition Costs: Buying H100, H200, B200, and GB200 Systems Outright](#)
 6. [Commitment Structures: On-Demand, Reserved, and Spot Pricing Economics](#)
 7. [Data Analysis and Evidence](#)
 8. [Case Studies and Real-World Examples](#)
 9. [xAI's Colossus Cluster: Speed as a Pricing and Deployment Strategy](#)
 10. [CoreWeave's OpenAI and Meta Contracts: Pricing Frontier-Scale Capacity](#)
 11. [Stargate: A \\$500 Billion Commitment to Physical AI Infrastructure](#)
 12. [China's Export-Restricted Chip Black Market: Pricing Under Sanctions](#)
 13. [Implications and Future Directions](#)
 14. [Conclusion](#)
 15. [Frequently Asked Questions \(FAQs\)](#)
-

Data Center GPU Pricing 2026: The Full AI Pricing Index

Executive Summary

Data center GPU pricing in 2026 is defined by a wide spread rather than a single number. An NVIDIA **H100** 80GB GPU rents for as little as \$1.38 to \$1.49 per GPU-hour on marketplace and boutique clouds and as much as \$11.68 to \$12.29 per GPU-hour on hyperscaler on-demand instances, a spread across materially different purchasing options. Marketplace and interruptible capacity, reserved commitments, and managed multi-GPU cloud instances are not functionally equivalent or directly comparable as a single on-demand rate range. Google Cloud's own pricing page lists the 8-GPU **a3-highgpu-8g** H100 instance at \$88.49 per hour on-demand, or roughly \$11.06 per GPU, while Amazon's EC2 P5 page confirms up to 8 H100 or H200 GPUs per instance in clusters scaling to 20,000 GPUs ([cloud.google.com](#)) ([aws.amazon.com](#)). The newer **H200**, with 141GB of HBM3e memory versus the H100's 80GB, has about 76% more memory capacity; its hourly price must be compared within the same provider, region, instance configuration, and billing model rather than expressed as a single market-wide premium. NVIDIA's Blackwell-generation **B200** runs \$2.12 to \$6.04 per hour depending on provider tier and contract length, and a full **GB200 NVL72** rack (72 Blackwell GPUs) prices out at \$10.50 to \$27 per GPU-hour, or \$756 to \$1,944 per hour for the whole rack ([gputaas.com](#)) ([spheron.network](#)).

Purchasing hardware outright follows the same wide bands. A new H100 costs \$25,000 to \$40,000 per card depending on PCIe versus SXM5 form factor, an 8-GPU DGX H100 system runs \$300,000 to \$460,000, and an 8-GPU H200 DGX system is quoted at roughly \$400,000 to \$500,000 ([cloudzero.com](#)) ([intuitionlabs.ai](#)).

Used H100 cards that sold for \$40,000 in late 2023 now trade for \$6,000 to \$22,000 on secondary markets, an 85 percent decline that reflects both Blackwell's arrival and a supply base that has broadened from a handful of hyperscalers to more than two dozen specialist clouds (cloudzero.com). Break-even between renting and owning a single H100 lands around 18 to 24 months of continuous, high-utilization operation once power, cooling, and networking are counted, which is why most [organizations below hyperscaler scale](#) still rent (cloudzero.com).

The pricing spread is widening, not narrowing, because of a genuine supply squeeze. Contract pricing for H100 and H200 GPUs climbed roughly 40 percent between October 2025 and March 2026 as HBM3e memory costs passed through from Samsung and SK Hynix collided with surging demand, and industry trackers describe finding eight-GPU nodes of Hopper-class capacity on short notice as "no longer routine" (gpuaas.com). Lead times for Blackwell PRO GPUs stretched to three to seven months by early 2026, and the Financial Times, cited by Reuters, reported that NVIDIA's export-restricted chips more than doubled in price on China's black market in the first half of 2026 (electropages.com) (reuters.com). Demand-side pressure is equally stark: NVIDIA reported record Data Center revenue of \$75.2 billion for its fiscal first quarter of 2027 (the three months ended April 26, 2026), up 92 percent year over year, while the four largest hyperscalers, Amazon, Google, Meta, and Microsoft, plan a combined [\\$630 billion in 2026 capital expenditures](#), a 62 percent jump from 2025's \$388 billion (nvidianews.nvidia.com) (datacenter richness.substack.com). Building the data centers that house this compute now costs \$20 million or more per megawatt for AI-optimized facilities, versus \$7 to \$12 million per megawatt for conventional builds, with fully built-out hyperscale campuses modeled at \$45 to \$55 billion per gigawatt (gigaenergy.com). For buyers, the practical takeaway is that provider selection and contract structure, not the GPU model alone, now determine the bill: reserved multi-year commitments cut 25 to 50 percent off on-demand rates, and specialist "neoclouds" such as CoreWeave, Lambda, and Nebius consistently undercut AWS, Azure, and Google Cloud on-demand pricing by roughly 40 to 70 percent for identical hardware.

Introduction and Background

Twenty-four months ago, a single NVIDIA H100 was a scarce commodity that commanded whatever a hyperscaler chose to charge. By mid-2026, the same chip is sold through more than two dozen distinct channels, from spot marketplaces to reserved hyperscaler capacity blocks to secondary-market resale, at prices that can differ by a factor of eight or more for the identical part (cloudzero.com). This report is a data-driven index of what data center **graphics processing units (GPUs)**, the specialized processors that perform the parallel matrix math underlying artificial intelligence (AI) model training and inference, actually cost across the market as of July 2026. It draws on official cloud provider pricing pages, NVIDIA's own financial filings, independent pricing trackers that publish verified vendor-by-vendor rate tables, and reporting on the largest AI infrastructure contracts signed to date.

The market this report covers spans three NVIDIA GPU generations still in active commercial use: the **Hopper-architecture H100** (80GB HBM3 memory, released 2022) and its memory-upgraded sibling the **H200** (141GB HBM3e), and the **Blackwell-architecture B200** (192GB HBM3e) alongside its rack-scale sibling the **GB200 NVL72**, which packages 72 B200 GPUs and 36 Grace CPUs into a single liquid-cooled cabinet (gpuaas.com) (spheron.network). Buyers access this hardware through four distinct channels that this report

treats separately throughout: **hyperscaler clouds** (AWS, Microsoft Azure, Google Cloud, Oracle Cloud Infrastructure), **specialist or "neocloud" providers** (CoreWeave, Lambda, Nebius, Crusoe), **long-tail and marketplace providers** (Vast.ai, Runpod, ThunderCompute, Hyperstack), and **direct hardware purchase** for **organizations building owned infrastructure**. Each channel prices the same underlying silicon differently because each bundles a different mix of networking, service-level agreements, compliance certifications, and sales overhead into the hourly rate (mercatus-ai.com).

NVIDIA does not publish an official price list for its data center chips, which is precisely why third-party price discovery matters and why this report leans on verified, dated vendor quotes rather than manufacturer suggested pricing (intuitionlabs.ai). Pricing is also unusually volatile in 2026: on-demand H100 rates fell sharply through 2025 as AWS cut prices by roughly 44 percent that June and competition from specialist clouds intensified, then reversed course as a fresh supply squeeze, driven by high-bandwidth memory (HBM) shortages rather than GPU die availability, pushed contract pricing back up through the first half of 2026 (intuitionlabs.ai) (electropages.com). Every figure in this report is anchored to a specific measurement date, generally between March and July 2026, because a rate quoted without a date is close to meaningless in a market moving this fast.

This report is written for technology and operations leaders evaluating AI infrastructure budgets, including teams in regulated, data-intensive fields such as life sciences, where GPU-backed workloads now span genomic sequence modeling, molecular simulation, and clinical natural-language processing pipelines that were not economically feasible on general-purpose compute even three years ago. IntuitionLabs, a life-sciences and AI consultancy, tracks GPU pricing as part of its own infrastructure cost research and has independently published H100 rental comparisons and hardware acquisition cost guides that corroborate the figures assembled here, though it is not itself a GPU cloud or hardware vendor and does not appear as a pricing option in the comparisons that follow (intuitionlabs.ai).

The Cloud GPU Rental Market in 2026: Pricing by Provider and GPU Model

Cloud GPU pricing in 2026 sorts cleanly into four provider tiers, and the gap between the cheapest and most expensive tier is the single largest lever available to a buyer, larger than the choice between GPU generations. Hyperscalers (AWS, Azure, Google Cloud, Oracle) charge the most because their on-demand rate bundles a broad managed-services ecosystem, global compliance certifications, and enterprise support that many AI workloads do not strictly require, while specialist "neoclouds" that run leaner software stacks and price aggressively for reserved-capacity deals routinely undercut hyperscaler on-demand rates by 40 to 70 percent for the same chip (gpucompare.com).

Table 1 collects published H100 rate snapshots from several billing models. Its per-GPU-hour figures are a convenience normalization, not a like-for-like on-demand index: marketplace and interruptible offers, reserved commitments, and managed multi-GPU instances differ in availability, service scope, and included resources.

Table 1: NVIDIA H100 published cloud-rate snapshots by provider, normalized to \$/GPU-hour (not a like-for-like on-demand index)

Provider	Tier / billing basis	Published \$/GPU-hour rate	Notes
Vast.ai	Marketplace	\$1.49	Interruptible marketplace floor (gpucloudcost.com)
ThunderCompute	Long-tail	\$1.38	Virtual GPU model, shared infrastructure (deploybase.ai)
Runpod	Long-tail	\$1.99 to \$2.89	PCIe Community Cloud to Secure Cloud SXM (intuitionlabs.ai)
Together AI	Specialist, reserved	\$3.09	Reserved cluster, 91 to 180 day term (gpucloudcost.com)
Nebius	Specialist	\$3.85	HGX H100 80GB, 3.2 Tbps InfiniBand (gpuccompare.com)
Lambda	Specialist	\$3.99	8x H100 SXM5 node, on-demand (gpuccompare.com)
CoreWeave	Specialist	\$4.76	NVIDIA HGX H100 published hourly price (CoreWeave Classic Pricing)
AWS EC2 (p5)	Hyperscaler	\$6.88	On-demand, post mid-2025 price cut (gpucloudcost.com)
Oracle Cloud (OCI)	Hyperscaler	\$10.00	Bare-metal 8x H100 node (gpuccompare.com)
Google Cloud (a3-highgpu-8g)	Hyperscaler	\$11.06	On-demand list, official pricing page (cloud.google.com)
Microsoft Azure (ND H100 v5)	Hyperscaler	\$12.29	On-demand, highest tracked list rate (gpuccompare.com)

The apparent spread in Table 1 should not be interpreted as a single market-wide premium for identical service. For example, Google Cloud states that its accelerator-optimized VM price includes attached GPUs, predefined vCPUs, memory, and bundled Local SSD storage where applicable; it is therefore not equivalent to a bare or interruptible marketplace GPU rate. Compare prices only after matching region, GPU form factor and count, billing model, interruption terms, and included compute, storage, networking, and support. ([Google Cloud's accelerator-optimized VM pricing](#))

Answering "**nvidia h100 gpu price 2026**" directly: as of mid-2026, expect \$1.38 to \$3.50 per GPU-hour on marketplace and long-tail clouds, \$3.85 to \$6.16 per GPU-hour on specialist neoclouds, and \$6.88 to \$12.29 per GPU-hour on hyperscaler on-demand instances, an overall range documented consistently by DeployBase (\$1.38 to \$11.68 across 28+ providers, March 2026), GPUCloudCost.com (\$1.49 to \$12.29 across 17 providers, June 2026), and CloudZero (\$1.38 to \$8.00+, with a market median of \$2.29 to \$3.12), with AWS's own instance page confirming the P5 family as the delivery vehicle at hyperscaler scale (deploybase.ai) (cloudzero.com) (aws.amazon.com).

For "**h200 vs b200 pricing**," the two chips are different products and should be compared using matched workload and contract assumptions. NVIDIA specifies the H200 at 141GB of HBM3e memory and 4.8 TB/s of memory bandwidth; compared with an 80GB H100, that is about 76% more capacity. NVIDIA characterizes the capacity as nearly double the H100's and the bandwidth as 1.4 times higher. A universal hourly-price premium cannot be inferred across providers because published prices vary by region, instance bundle, availability, and

commitment type. ([NVIDIA H200 GPU](#)) Google Cloud's own pricing page confirms this tier directly: the a3-ultragpu-8g instance, built on 8x H200, lists at \$84.81 per hour on-demand, roughly \$10.60 per GPU, while the Blackwell-based a4-highgpu-8g (8x B200) is not sold as a standard on-demand SKU but carries a published spot rate of \$39.63 per hour, roughly \$4.95 per GPU ([cloud.google.com](#)) ([cloud.google.com](#)). The B200 itself is an entirely different chip: NVIDIA's Blackwell architecture packs 208 billion transistors in a dual-die design, 192GB of HBM3e at 8.0 TB/s, and native FP4 precision support that raises inference throughput roughly 2.5x to 4.5x over Hopper ([gpuaas.com](#)). B200 on-demand pricing runs \$2.12 spot to \$4.99 to \$6.04 per GPU-hour at specialist providers, with 36-month reserved contracts as low as \$2.25 per GPU-hour, while AWS's Blackwell instance list rate reaches roughly \$14.24 per GPU-hour ([gpuaas.com](#)). Because B200 supply is still constrained relative to Hopper, most teams in 2026 are advised to buy or reserve H100 or H200 capacity today and evaluate Blackwell on a two-year upgrade cycle rather than wait ([mercatus-ai.com](#)).

At the top of the price ladder sits the **GB200 NVL72**, NVIDIA's rack-scale Grace Blackwell system that is not rented as individual GPUs but as Superchip or rack-node allocations. CoreWeave lists GB200 NVL72 on-demand access at roughly \$10.50 per GPU-hour through 4-GPU instances, Oracle Cloud prices bare-metal access around \$16 per GPU-hour, and Azure's ND GB200 v6, generally available since March 2025, runs approximately \$27 per GPU-hour ([spheron.network](#)) ([spheron.network](#)). Scaled to a full 72-GPU rack, that range implies \$756 to \$1,944 per hour, 1.7x to 4.5x the cost of the same 72 GPUs deployed as nine separate 8xB200 nodes, a premium that only pays back when a workload's inter-GPU communication overhead is large enough that the NVL72's 130 TB/s all-to-all NVLink fabric meaningfully cuts training or serving time ([spheron.network](#)).

Legacy Ampere-generation **A100** GPUs, NVIDIA's pre-Hopper flagship, remain in active commercial use for cost-sensitive fine-tuning and inference that does not require FP8 acceleration. A100 80GB on-demand pricing runs \$1.50 to \$3.00 per GPU-hour on hyperscalers, some of which have phased the chip out entirely, \$1.10 to \$1.50 on specialist clouds, and \$0.80 to \$1.10 on long-tail providers ([mercatus-ai.com](#)). Cloud A100 rates of roughly \$1.29 to \$2.50 per hour remain cheaper than H100 rates on paper, but the H100 delivers three to five times better throughput on transformer workloads, so an A100 job that takes 100 hours can still cost more than an H100 job that takes 30, even at the lower hourly rate, which is why cost per completed run, not cost per hour, should drive the comparison ([cloudzero.com](#)).

Hardware Acquisition Costs: Buying H100, H200, B200, and GB200 Systems Outright

For organizations planning owned infrastructure rather than cloud rental, the acquisition-cost picture answers "cost to build an ai data center" at both the chip and facility level. *Table 2* summarizes verified purchase pricing for individual GPUs, complete 8-GPU systems, and rack-scale Blackwell hardware as of Q1 to Q2 2026.

Table 2: Data center GPU and system purchase costs, plus GB200 NVL72 cloud-equivalent rate (verified Q1 to Q2 2026)

Item	Price range	Source detail
H100 PCIe 80GB, new	\$25,000 to \$30,000	Standard PCIe slot, no NVLink (deploybase.ai)
H100 SXM5 80GB, new	\$35,000 to \$40,000	NVLink 4.0, 900 GB/s interconnect (cloudzero.com)
H100, used or secondary market	\$6,000 to \$22,000	85% below the 2023 peak of \$40,000 (cloudzero.com)
DGX H100, 8-GPU system	\$300,000 to \$460,000	Full system: CPUs, networking, storage (deploybase.ai)
DGX H200, 8-GPU system	\$400,000 to \$500,000	Quoted range per market analysis (intuitionlabs.ai)
GB200 NVL72 rack (72 GPUs)	\$756 to \$1,944/hour cloud-equivalent	No single retail price published; sold as rack-node cloud allocations (spheron.network)

The 85 percent collapse in used H100 pricing deserves emphasis: cards that traded for \$40,000 in late 2023 changed hands for as little as \$6,000 by mid-2026, a decline driven by Blackwell's arrival making Hopper-class hardware progressively less economical for inference workloads even as demand for AI compute overall keeps climbing ([cloudzero.com](#)). At the same time, new H100 street pricing has not collapsed nearly as far, because HBM3 memory and packaging costs remain a hard floor under manufacturing cost regardless of secondary-market discounting. Viewed across the full product lineup, new-card pricing runs roughly \$10,000 to \$15,000 for an A100, \$25,000 to \$40,000 for an H100, and \$30,000 to \$50,000 for the newest B200 cards, a curve that tracks each generation's memory bandwidth and precision-format gains more closely than it tracks raw compute alone ([cloudzero.com](#)).

Buy-versus-rent math depends heavily on utilization. At a representative \$25,000 purchase price and a \$3.00 per hour rental benchmark, the raw break-even point is roughly 8,333 hours, about 347 days of continuous use, but that calculation ignores power, cooling, and rack infrastructure, which shift the realistic break-even to 18 months or more of near-100 percent utilization ([cloudzero.com](#)). A single SXM5 GPU pulls up to 700 watts, meaning a full 8-GPU node draws roughly 5.6 kilowatts under load, translating to real annual electricity costs in addition to \$10,000 to \$100,000 in incremental cooling and power infrastructure per rack ([deploybase.ai](#)). Purchasing tends to make financial sense only when GPUs run 24/7 for 18 or more months, the buyer already has (or budgets \$400,000-plus for) data center infrastructure, and in-house staff can manage firmware, drivers, and liquid cooling, a profile that describes hyperscalers and large AI labs far more than mid-market engineering teams ([cloudzero.com](#)).

Beyond the chips themselves, the facility that houses them is now the larger line item for anyone building at scale. Conventional data centers cost \$7 to \$12 million per megawatt of capacity, but AI-optimized facilities, which must accommodate GPU-dense racks, liquid cooling, and high-voltage power distribution, run \$20 million or more per megawatt, and hyperscale AI campuses at gigawatt scale are modeled at \$45 to \$55 billion per gigawatt for a fully built-out ecosystem ([gigaenergy.com](#)) ([gigaenergy.com](#)). That construction budget is separate from the IT equipment inside it: at 100 megawatts of capacity, servers, networking gear, and GPUs alone can run \$2.5 to \$4 billion, and transformer lead times from legacy manufacturers can stretch past 50 weeks, a supply constraint that increasingly rivals GPU allocation as the binding factor on how quickly new AI capacity can come online ([gigaenergy.com](#)) ([gigaenergy.com](#)). Independent trackers segment likely buyers by scale: a small team running one to four GPUs for inference and fine-tuning should budget roughly \$500 to

\$5,000 per month, a mid-sized cluster of 8 to 64 H100 or A100 GPUs runs \$15,000 to \$250,000 per month, and a large cluster of 256 or more H100, H200, or B200 GPUs runs \$1 million to \$40 million or more per month (gpucloudcost.com) (gpucloudcost.com).

Commitment Structures: On-Demand, Reserved, and Spot Pricing Economics

Beyond provider choice, the contract structure a buyer selects moves the effective hourly rate by 30 to 80 percent, which is why two organizations renting the identical H100 SXM chip from the identical provider can pay materially different bills. Three structures dominate the 2026 market, and understanding each is central to answering "[cloud gpu rental pricing comparison](#)" and "[ai gpu cost per hour](#)" in practical terms rather than headline list-price terms.

On-demand pricing carries no commitment and the highest per-hour rate; it suits unpredictable workloads, unproven providers, or short evaluation windows. **Reserved** capacity, typically committed for one to three years, cuts cost 25 to 50 percent versus on-demand at the same provider, with one-year terms averaging 25 to 35 percent savings and three-year terms averaging 35 to 50 percent, rising past 50 percent for custom five-year-plus commitments (mercatus-ai.com). **Spot** or preemptible instances undercut on-demand by 30 to 80 percent in exchange for the risk of short-notice reclamation, typically a two-minute eviction warning on AWS, which makes them suitable for checkpointed training runs and batch workloads but poor fits for production inference with uptime requirements (deploybase.ai). Google Cloud's own preemptible A3 instances illustrate the discount structure directly, published at roughly 60 to 91 percent off on-demand pricing depending on availability (deploybase.ai).

Table 3 reports AWS's published effective hourly rates for EC2 Capacity Blocks for ML. These are reservation rates, not AWS on-demand prices, and cannot be used as a cross-provider comparison.

Table 3: AWS EC2 Capacity Blocks for ML, P5 H100 effective hourly rates by region, \$/GPU-hour (checked October 10, 2026)

Instance type	Region group	Effective hourly rate per instance	Effective hourly rate per H100
p5.48xlarge (8× H100)	US East (N. Virginia, Ohio), US West (Oregon and N. California), US East (Atlanta) Local Zone, Asia Pacific (Tokyo, Jakarta, Mumbai), Australia (Sydney), Europe (London, Stockholm), and South America (São Paulo)	\$47.757	\$5.970
p5.4xlarge (1× H100)	US East (N. Virginia, Ohio), US West (Oregon), Asia Pacific (Tokyo, Mumbai), Australia (Sydney), Europe (London), and South America (São Paulo)	\$5.970	\$5.970

(Source: [AWS EC2 Capacity Blocks for ML Pricing](#), checked October 10, 2026. AWS now lists a single P5 rate across all of these regions; when this report was first published in July 2026 the same page listed \$41.528 per p5.48xlarge instance (\$5.191 per H100) in the US regions and \$37.76 (\$4.720 per H100) in the other listed regions. AWS states that these rates are next scheduled to be updated in January 2027. Rates are subject to change.)

Table 3 is useful only for comparing AWS Capacity Blocks across the listed regions and P5 configurations. It does not establish on-demand, spot, or multi-year reserved prices for AWS or any other provider. A complete cost comparison should obtain dated quotes for the same region, instance configuration, contract term, storage, networking, and support requirements; additional services can materially change the total invoice.

This structure also governs "**gpu prices for llm training**." A full fine-tune of a 7-billion-parameter Mistral model on a single H100 SXM at \$2.69 per hour for 20 hours costs \$53.80 in raw compute, while the same task using parameter-efficient LoRA fine-tuning on a cheaper A100 at \$1.39 per hour for the same duration costs \$27.80, roughly half, for broadly comparable output quality ([deploybase.ai](#)). At the median H100 cloud rate, a single reserved GPU running continuously costs roughly \$1,964 per month, \$23,567 annually, and \$70,700 over three years, a baseline that scales linearly (before volume discounts) to the multi-hundred-GPU clusters that frontier model training runs require ([deploybase.ai](#)). Right-sizing GPU form factor also matters more than most teams expect: an inference workload running on SXM at \$2.69 per hour that does not actually need NVLink multi-GPU interconnect can move to PCIe at \$1.99 per hour and save 26 percent with no measured performance loss, since the SXM premium only earns its keep when multi-GPU communication bandwidth is actually the bottleneck ([deploybase.ai](#)).

Data Analysis and Evidence

The quantitative backdrop to every price quoted above is a demand curve that is still steepening in 2026, not flattening, which is the central fact behind "**ai infrastructure cost trends 2026**." NVIDIA's own financial results are the clearest signal: for its fiscal first quarter of 2027, the three months ended April 26, 2026, the company reported record total revenue of \$81.6 billion, up 85 percent year over year and 20 percent sequentially, with GAAP gross margin of 74.9 percent ([nvidianews.nvidia.com](#)). Data Center revenue specifically, the segment that includes H100, H200, B200, and GB200 sales, hit a record \$75.2 billion, up 92 percent year over year, and within that, Data Center compute revenue alone reached \$60.4 billion (up 77 percent year over year) while Data Center networking revenue reached \$14.8 billion (up 199 percent year over year), evidence that buyers are increasingly paying as much for the interconnect fabric around GPUs as for the chips themselves ([nvidianews.nvidia.com](#)). NVIDIA guided its following quarter to \$91.0 billion in revenue, plus or minus 2 percent, explicitly excluding any Data Center compute revenue from China given ongoing export restrictions ([nvidianews.nvidia.com](#)). Update (October 2026): for its second quarter of fiscal 2027, the three months ended July 26, 2026, NVIDIA reported revenue of \$96.2 billion, up 106 percent year over year, and Data Center revenue of \$89.0 billion, up 117 percent, and guided its third quarter to \$108.0 billion, plus or minus 2 percent, again assuming no Data Center compute revenue from China ([nvidianews.nvidia.com](#)).



On the buyer side, the four largest hyperscalers disclosed a combined \$630 billion in planned 2026 capital expenditures during their respective early-2026 earnings calls, a 62 percent increase over the record \$388 billion the same four companies spent in 2025 (datacenter richness.substack.com). Individually: Amazon guided to \$200 billion for 2026, up from \$125 billion in 2025; Google guided to \$175 to \$185 billion, up from \$91 billion; Meta guided to \$115 to \$135 billion, up from \$72 billion; and Microsoft guided to \$110 to \$120 billion, up from \$90 billion (datacenter richness.substack.com) (datacenter richness.substack.com). Most of that capital converts directly into GPU purchases and the data center shells that house them, which is the mechanical link between hyperscaler earnings calls and the cloud rental prices in Table 1. AWS's own P5 documentation illustrates the resulting scale directly, describing EC2 UltraClusters that interconnect up to 20,000 H100 or H200 GPUs on a petabit-scale nonblocking network capable of up to 20 exaflops of aggregate compute, with Anthropic among the early customers stating publicly that it expected the P5 generation "to deliver substantial price-performance benefits" over the prior GPU generation at that scale (aws.amazon.com) (aws.amazon.com).

Supply-side scarcity compounds the demand story. Industry trackers reported H100 and H200 contract pricing climbing roughly 40 percent between October 2025 and March 2026, driven primarily by HBM3e memory cost pass-throughs from Samsung and SK Hynix rather than by GPU die shortages, with roughly half of tracked specialist providers reporting no Hopper-class capacity coming off contract at all during that window (gpuaas.com). This is the substance behind "**gpu shortage pricing impact**": lead times for Blackwell PRO-class GPUs stretched to three to seven months by early 2026, RTX 6000 Ada 48GB cards rose to \$7,400 to \$7,800, and RTX PRO 6000 96GB cards climbed to \$9,450 to \$9,800, all reflecting a memory-constrained rather than silicon-constrained bottleneck (electropages.com). The scarcity is broad-based rather than

confined to the newest chips: sourcing data cited by trade press shows NVIDIA's H200 NVL PCIe and H100 NVL PCIe among the most requested parts even as buyers simultaneously widen their search to older Ada-generation and legacy cards to find any available inventory ([electropages.com](https://www.electropages.com)). Analysts distinguish this cycle explicitly from the 2021-2022 GPU shortage, which was driven by speculative cryptocurrency mining demand that eventually collapsed; the 2026 tightness is described as structural, tied to persistent AI infrastructure investment cycles rather than a speculative bubble that self-corrects ([electropages.com](https://www.electropages.com)).

Export controls add a further pricing distortion. The Financial Times, in reporting cited by Reuters, found that NVIDIA's export-restricted AI chips more than doubled in price on China's black market over the first half of 2026, based on interviews with multiple Chinese chip traders, illustrating that geopolitical restrictions on legitimate sales channels create their own parallel pricing structure entirely disconnected from the on-demand and reserved rates in this report's core tables ([reuters.com](https://www.reuters.com)).

Case Studies and Real-World Examples

xAI's Colossus Cluster: Speed as a Pricing and Deployment Strategy

Elon Musk's xAI built the clearest illustration of what pricing looks like when speed, not unit cost, is the binding constraint. Quoted 18 to 24 months for conventional data center construction, xAI instead repurposed a former Electrolux appliance factory in Memphis, Tennessee, and deployed 100,000 NVIDIA H100 GPUs there in 122 days, then doubled the cluster to 200,000 GPUs in a further 92 days ([introl.com](https://www.introl.com)) ([introl.com](https://www.introl.com)). By December 2025, the cluster comprised 150,000 H100, 50,000 H200, and 30,000 GB200 GPUs, described by the deployment's infrastructure partner as the largest fully operational, single-coherent AI training cluster in the world ([introl.com](https://www.introl.com)). Powering that density required 35 gas turbines capable of 420 megawatts alongside 208 Tesla Megapack battery units and a 150-megawatt utility substation built in 97 days versus the roughly 2.5 years such projects normally require, with the facility drawing approximately 250 megawatts by late 2025 and xAI separately committing \$80 million to a wastewater recycling facility to support cooling at that scale ([introl.com](https://www.introl.com)). Networking relied on NVIDIA's Spectrum-X Ethernet fabric rather than the InfiniBand standard used by most large training clusters, achieving 95 percent data throughput against the roughly 60 percent typical of standard Ethernet at this scale, evidence that Ethernet-based fabrics can now support frontier-scale training economically ([introl.com](https://www.introl.com)).

CoreWeave's OpenAI and Meta Contracts: Pricing Frontier-Scale Capacity

CoreWeave, a specialist GPU cloud provider, illustrates how frontier AI labs now lock in capacity years in advance through multibillion-dollar bilateral contracts rather than spot-market rental. CoreWeave's relationship with OpenAI began with an \$11.9 billion, five-year agreement in March 2025, expanded by \$4 billion in May 2025, then expanded again by \$6.5 billion in September 2025, bringing the cumulative contract value to \$22.4 billion ([cnbc.com](https://www.cnbc.com)) ([cnbc.com](https://www.cnbc.com)) ([cnbc.com](https://www.cnbc.com)). Separately, in April 2026, CoreWeave and Meta Platforms announced an expanded, long-term agreement to provide AI cloud capacity through December 2032 for

approximately \$21 billion, with the dedicated capacity to include some of the earliest commercial deployments of NVIDIA's next-generation Vera Rubin platform, a deal CoreWeave itself described as "a clear signal of the industry's" accelerating demand for high-performance AI infrastructure capacity (coreweave.com) (coreweave.com). By the time of its 2025 public listing prospectus, CoreWeave operated 32 data centers powered by more than 250,000 NVIDIA GPUs across the United States and parts of Europe, and separately committed \$6 billion to a 100-megawatt data center in Lancaster, Pennsylvania (cnbc.com). These contracts illustrate that at the frontier-lab scale, GPU access is priced and negotiated as multi-year committed capacity rather than the hourly rates that dominate Tables 1 through 3, but the underlying economics, reserved capacity beating on-demand, are the same principle scaled up by roughly six orders of magnitude.

Stargate: A \$500 Billion Commitment to Physical AI Infrastructure

OpenAI, Oracle, and SoftBank's Stargate initiative demonstrates how GPU demand cascades into gigawatt-scale physical construction commitments. First announced in January 2025 at \$500 billion and 10 gigawatts of planned capacity, the partnership announced five additional U.S. data center sites in October 2025 that brought total planned capacity to nearly 7 gigawatts and over \$400 billion in committed investment over three years, putting the project ahead of schedule to secure its full \$500 billion commitment by the end of 2025 (openai.com) (openai.com). A separate July 2025 agreement between OpenAI and Oracle to develop up to 4.5 additional gigawatts of Stargate capacity was itself valued at more than \$300 billion over five years (openai.com). At the flagship Abilene, Texas campus, Oracle began delivering the first NVIDIA GB200 racks in June 2025 and had already started early training and inference workloads on that capacity by the time of the announcement, a concrete example of GB200-generation hardware moving from announcement to production inside roughly a year (openai.com). Separately, NVIDIA itself announced a \$100 billion investment in OpenAI to support the buildout, and Microsoft committed \$4 billion to a second Wisconsin data center in the same period, illustrating how capital now flows in multiple directions between chip suppliers, cloud providers, and AI labs simultaneously (cnbc.com).

China's Export-Restricted Chip Black Market: Pricing Under Sanctions

The fourth case illustrates what GPU pricing looks like when the legal supply channel is closed by policy rather than by capacity. With U.S. export controls barring the sale of NVIDIA's most advanced AI chips to China, a parallel black market has formed, and the Financial Times, in reporting relayed by Reuters in June 2026, found that prices for these restricted chips had more than doubled over the preceding six months based on interviews with multiple Chinese chip traders (reuters.com). This case matters for the broader pricing index because it demonstrates that the supply-and-demand dynamics documented in Tables 1 through 3, driven by memory shortages, hyperscaler capex, and reserved-versus-spot economics, operate even more acutely once legitimate market access is restricted, and it explains why NVIDIA's own Q2 fiscal 2027 revenue guidance explicitly assumed zero Data Center compute revenue from China rather than model a highly uncertain, sanctioned market (nvidianews.nvidia.com).

Implications and Future Directions

Three structural forces will keep shaping data center GPU pricing through the remainder of 2026 and into 2027. First, memory, not GPU die fabrication, has become the binding supply constraint: HBM3e cost pass-throughs from Samsung and SK Hynix were the primary driver of the roughly 40 percent H100 and H200 contract price increase between October 2025 and March 2026, and unless HBM output expands materially, price pressure on Hopper and Blackwell-class GPUs alike is likely to persist even as GPU die supply itself improves (gputaas.com). Second, the buy-versus-rent calculus continues to favor rental for all but the largest, most utilization-disciplined organizations, since the reserved, long-tail rate of \$1.30 to \$1.80 per GPU-hour now sits within striking distance of owned-cluster economics without the operational burden of managing firmware, cooling, and depreciation directly (mercatus-ai.com). Third, physical infrastructure, power procurement, transformer lead times exceeding 50 weeks, and construction costs of \$20 million or more per megawatt, is increasingly the true bottleneck on how fast new GPU capacity can be deployed, a shift visible in xAI's turn to gas turbines and Tesla Megapacks and in Stargate's framing of its commitment in gigawatts rather than GPU counts (gigaenergy.com) (openai.com).

For buyers outside the hyperscaler tier, including mid-market AI teams, research labs, and specialized enterprise deployments in fields such as life sciences that increasingly run genomic foundation models, molecular dynamics simulations, and large-scale clinical text pipelines on rented GPU capacity, three practical shifts follow from this data. Reserved and long-tail capacity now beats hyperscaler on-demand pricing by wide enough margins, 40 to 70 percent in many comparisons, that provider selection alone can swing an annual GPU budget by hundreds of thousands of dollars at modest cluster sizes (mercatus-ai.com). Workload-to-hardware matching is the second lever: a team serving a 70-billion-parameter model in FP16 needs roughly 140GB of VRAM, which a single H200 covers at \$3.72 to \$6.00 per hour where an equivalent H100 deployment would need two GPUs billed at double the rate, and choosing PCIe over SXM for single-GPU inference that does not need NVLink saves a further 20 to 26 percent with no measurable performance cost (gputaas.com) (deploybase.ai). Consultancies advising on AI infrastructure procurement, including firms serving regulated life-sciences organizations, are increasingly treating GPU cost modeling as a first-class part of technology strategy rather than an afterthought delegated entirely to engineering, given that the same workload can cost 8 to 9 times more or less depending purely on which of the four provider tiers is selected; one such firm's own independent review of the H100 rental market, conducted separately from the trackers cited elsewhere in this report, found rates spanning \$1.49 to \$6.98 per hour across more than fifteen providers, a range consistent with every other source in this index despite being compiled independently (intuitionlabs.ai).

Looking forward, the Blackwell Ultra generation (B300) and NVIDIA's announced Vera Rubin platform are already entering early commercial deployment. At least one specialist provider prices B300 access at \$8.55 per hour on-demand or \$3.67 per hour spot, and CoreWeave and Meta's April 2026 agreement specifically calls out Vera Rubin as part of the contracted capacity, suggesting the pricing ladder documented in this report will gain at least one more rung, with correspondingly higher price ceilings, before 2026 is over (spheron.network) (coreweave.com). Update (October 2026): NVIDIA reported in August 2026 that the Vera Rubin platform is ramping into full production, with racks running at partners including CoreWeave, Google Cloud, Microsoft Azure, Oracle Cloud Infrastructure and Nebius, and AWS lists its 8-GPU Blackwell Ultra instance, p6-b300.48xlarge, at \$142.416 per hour on-demand in US East (N. Virginia), roughly \$17.80 per GPU-hour (nvidianews.nvidia.com) (aws.amazon.com). Whether that new generation compresses Hopper and

early Blackwell pricing further, the way B200 arrival already discounted secondary-market H100s by 85 percent, or simply adds a new premium tier on top of an already-expensive market, will depend largely on whether HBM memory supply, the binding constraint identified throughout this report, expands fast enough to keep pace with hyperscaler capital spending that shows no sign of decelerating in 2026.

Conclusion

Data center GPU pricing depends on GPU generation, provider, region, exact instance SKU, included resources, and billing model. The H200 has 141GB of HBM3e memory versus the H100's 80GB, about 76% more capacity. NVIDIA describes the H200 as having nearly double the H100's memory capacity and 1.4 times its memory bandwidth. ([NVIDIA H200 GPU](#))

The evidence in this report supports comparing published offers only within matched conditions. A marketplace or interruptible GPU rate, a reserved-capacity rate, and a managed multi-GPU VM list price are different transactions with different availability, support, and included resources. Google Cloud, for example, states that accelerator-optimized VM prices include attached GPUs, predefined vCPUs, memory, and bundled Local SSD storage where applicable. ([Google Cloud's accelerator-optimized VM pricing](#))

For infrastructure budgeting, first define the workload, region, GPU configuration, uptime requirement, and contract term. Then compare offers with the same billing model and included services, and evaluate cost per completed workload rather than treating a normalized per-GPU-hour figure as a universal market price. Ownership-versus-rental economics likewise require organization-specific utilization, power, cooling, financing, depreciation, and operating-cost assumptions; this report does not establish one general break-even point.

Frequently Asked Questions (FAQs)

What does an NVIDIA H100 cost per hour in 2026? There is no single comparable H100 hourly price. Verify a current quote for the required region and configuration, then compare offers with the same billing model. In particular, do not compare an interruptible marketplace offer or a reserved commitment with an on-demand managed VM as if they were equivalent; managed GPU VM prices can include GPUs, predefined vCPUs, memory, and local storage. ([Google Cloud's accelerator-optimized VM pricing](#))

How much does it cost to rent versus buy a GPU cluster for LLM training? A single H100 at the median cloud rate of \$2.69 per hour costs roughly \$1,964 per month or \$23,567 per year to rent continuously; buying the same GPU outright costs \$25,000 to \$40,000 plus power and cooling, with realistic break-even at 18 months or more of near-full utilization ([deploybase.ai](#)) ([cloudzero.com](#)).

Is the H200 worth the price premium over the H100? It depends on the workload and on a matched provider quote. The H200 has 141GB of HBM3e memory and 4.8 TB/s of memory bandwidth; NVIDIA describes this as nearly double the H100's memory capacity and 1.4 times its memory bandwidth. Compare total cost and performance using the same region, instance configuration, availability terms, and commitment type rather than assuming a universal hourly premium. ([NVIDIA H200 GPU](#))

What does an NVIDIA A100 cost compared to newer GPUs? A100 80GB on-demand pricing runs \$0.80 to \$3.00 per GPU-hour depending on provider tier, cheaper than H100 or H200, but the H100 delivers three to five times better throughput on transformer workloads, so the A100 remains attractive mainly for fine-tuning and inference tasks that do not need FP8 acceleration ([mercatus-ai.com](https://www.mercatus-ai.com)) ([cloudzero.com](https://www.cloudzero.com)).

Why is there such a large gap between the cheapest and most expensive GPU cloud providers?

Hardware amortization costs are nearly identical across provider tiers, but hyperscalers layer on far higher sales, support, and margin overhead, roughly \$1.00 to \$2.00 per GPU-hour in margin alone versus \$0.20 to \$0.50 at long-tail providers, which is why identical silicon can differ in price by 8 to 9 times ([mercatus-ai.com](https://www.mercatus-ai.com)).

How much does it cost to build an AI data center in 2026? AI-optimized facilities run \$20 million or more per megawatt of capacity, versus \$7 to \$12 million per megawatt for conventional data centers, with hyperscale campuses at gigawatt scale modeled at \$45 to \$55 billion per gigawatt, plus \$2.5 to \$4 billion in servers, networking, and GPUs per 100 megawatts of capacity ([gigaenergy.com](https://www.gigaenergy.com)) ([gigaenergy.com](https://www.gigaenergy.com)).

Is the current GPU shortage likely to ease in 2026? Available evidence suggests no near-term easing. The shortage is driven by structural HBM memory constraints tied to sustained AI infrastructure investment rather than a speculative demand spike, which analysts distinguish explicitly from the 2021-2022 cryptocurrency-driven shortage that eventually collapsed on its own ([electropages.com](https://www.electropages.com)).

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://cloud.google.com/products/compute/pricing/accelerator-optimized>
2. <https://aws.amazon.com/ec2/instance-types/p5/>
3. <https://gpuaas.com/blog/h100-vs-h200-vs-b200-which-gpu-to-rent-2026>
4. <https://www.spheron.network/blog/nvidia-gb200-nvl72-guide/>
5. <https://www.cloudzero.com/blog/h100-gpu-cost/>
6. <https://intuitionlabs.ai/pdfs/nvidia-ai-gpu-prices-h100-27k-40k-h200-315k-8-gpu-cost-guide.pdf>
7. <https://intuitionlabs.ai/articles/biotech-ai-implementation-cost>
8. <https://www.electropages.com/blog/2026/03/fusion-worldwide-gpu-shortage-and-price-increases-2026>
9. <https://www.reuters.com/world/asia-pacific/nvidias-banned-ai-chips-double-price-chinas-black-market-ft-reports-2026-06-24/>
10. <https://intuitionlabs.ai/articles/biotech-ai-budget-benchmarks-report>
11. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027>
12. <https://datacenterrichness.substack.com/p/hyperscalers-plan-630-billion-in>
13. <https://www.gigaenergy.com/blog/how-much-does-it-cost-to-build-an-ai-data-center>
14. <https://intuitionlabs.ai/articles/kimi-k3-life-sciences-regulated-data>
15. <https://www.mercatus-ai.com/blog/cloud-gpu-pricing>
16. <https://intuitionlabs.ai/articles/h100-rental-prices-cloud-comparison>
17. <https://www.gpucostcompare.com/>
18. <https://gpucloudcost.com/>

19. <https://deploybase.ai/articles/nvidia-h100-price>
20. <https://www.coreweave.com/pricing/classic>
21. <https://www.nvidia.com/en-us/data-center/h200/>
22. <https://aws.amazon.com/ec2/capacityblocks/pricing/>
23. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-second-quarter-fiscal-2027>
24. <https://introl.com/blog/xai-memphis-colossus-100000-gpu-supercomputer-infrastructure>
25. <https://www.cnbc.com/amp/2025/09/25/coreweave-openai-6point5-billion-deal.html>
26. <https://www.coreweave.com/news/coreweave-and-meta-announce-21-billion-expanded-ai-infrastructure-agreement>
27. <https://openai.com/index/five-new-stargate-sites/>
28. <https://aws.amazon.com/ec2/pricing/on-demand/>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.