

Venture Compute in AI: Compute-for-Equity Programs and Partnerships (2024-2025)

Introduction – The Rise of “Venture Compute”

As the AI boom drives **massive demand for GPU computing**, startups often struggle to afford or even access the hardware needed to train large models. In response, a new model of support – sometimes dubbed “*venture compute*” – has emerged. **Data center operators and investors are offering cutting-edge compute resources (e.g. GPU cloud credits or dedicated clusters) to startups in exchange for equity stakes or strategic partnerships.** This approach helps early-stage AI companies get critical infrastructure **without upfront capital**, while providers and backers secure a place on the startup’s cap table or a future customer. The result is a growing landscape of **compute-based incubators, accelerators, and VC alliances** aimed at democratizing access to AI infrastructure.

VC-Led GPU Clusters – Compute for Equity

Leading venture capital firms have started **stockpiling GPUs and building private supercomputers** to give their portfolio startups a competitive edge. A prime example is **Andreessen Horowitz’s “Oxygen” program**, a **massive Nvidia H100 cluster** (reportedly 20,000+ GPUs ¹) that **portfolio companies can use at cost**. In exchange, **a16z receives equity stakes or rights** in those startups, effectively trading discounted GPU time for ownership ². The goal is to let new AI companies train models without paying cloud market rates, leveling the playing field against tech giants ³ ⁴. As a16z partner Anjney Midha explained, during the H100 supply crunch “startups were being deprioritized by large clouds... We offer GPUs when they need it, without a long-term commitment... Instead, they give a16z a stake in their business” ³ ².

Importantly, a16z is *not alone*. Former GitHub CEO **Nat Friedman and investor Daniel Gross** built the **“Andromeda” cluster with 4,000 GPUs** in mid-2023 to support their AI startup investments ⁵ ⁶. Their portfolio companies get access to this private supercomputer **at below-market rates** as part of the deal ⁶. **Y Combinator** has also arranged GPU access for its accelerator companies – through cloud partners rather than owning hardware. **YC startups can run training workloads on a dedicated GPU cluster provided via partnerships (most recently with Google Cloud)** ⁷. This gives each YC batch member a shot at model training runs using **Nvidia GPUs (or Google TPUs) at no cost**, courtesy of the YC–Google alliance.

Several VC firms have forged **cloud partnerships** instead of building clusters. For instance, **Sequoia Capital teamed up with Google Cloud** in 2024 to offer **up to \$500,000 in free cloud credits** for AI startups in Sequoia’s portfolio ⁸. (Google already offered \$350k credits to qualified AI startups generally, but this deal ups the ante for Sequoia companies ⁹.) Sequoia’s rationale: early-stage teams with only a few million in funding can see much of it drained by cloud bills, so subsidizing compute protects their runway ¹⁰. In return, Google gains an influx of promising AI customers on its platform. Similarly, **Index Ventures launched a program in late 2023 with Oracle Cloud** to grant its startups **free access to an Oracle-managed GPU cluster** ¹¹. And **Microsoft’s venture arm (M12)** has reserved **thousands of Azure GPU instances** specifically for startups in its portfolio and partner funds (including YC), ensuring those companies can tap Azure’s AI infrastructure as they grow ¹¹. Even newer funds like

Conviction (Sarah Guo's) quietly operate smaller GPU pools for their startups ¹¹. In all these cases, the **VC provides (or arranges) compute as a value-add on top of capital**, strengthening their investment by accelerating the startup's R&D. Venture firms see it as a competitive edge to *"win deals"* with top AI founders: indeed, **a16z's massive GPU purchase was in part to attract coveted AI startups by meeting their infrastructure needs** ¹².

Cloud & Data Center Providers Offering Startup Programs

Beyond VCs, many **GPU cloud providers and data center companies** have inaugurated programs to support AI startups – sometimes for equity, other times simply to build goodwill and future customers. Below is a landscape of notable initiatives:

Provider / Program	Infrastructure & Scale	Program Terms	Partners / Notes
CoreWeave Accelerator (USA)	Specialized cloud with 10+ <i>NVIDIA GPU SKUs</i> (A100, A5000, H100 etc.) ¹³ ; large-scale GPU clusters on demand.	Cloud credits + discounts on CoreWeave's GPU cloud ¹⁴ . No equity taken – instead, startups get intro's to VCs and join CoreWeave's client community ¹⁵ .	<i>Strategic:</i> CoreWeave connects startups with top VCs and ensures they can scale on CoreWeave's infrastructure. ¹⁶ Program launched 2022.
Lambda (Lambda Labs) (USA)	GPU Cloud & hardware provider. Offers on-demand and reserved instances with latest GPUs (currently NVIDIA H100 , and plans for H200 and Blackwell) ¹⁷ . Also sells GPU servers and workstations.	No formal equity program announced (Lambda focuses on paid services). However, provides <i>volume discounts, professional support</i> , and even "1-Click" multi-GPU cluster orchestration for startups needing to scale quickly.	<i>Notes:</i> Lambda's cloud is used by many AI startups for its bare-metal performance and custom ML tooling. Emerging programs may come as competition heats up.
Voltage Park (USA)	24,000 NVIDIA H100 GPUs acquired for its cloud ¹⁸ ; state-of-the-art Dell HGX H100 servers, InfiniBand networking (3.2 Tb/s) ¹⁹ . Data centers in Texas, Virginia, Washington ²⁰ .	Ultra-low pricing: GPUs as cheap as \\$1.89/hour (for H100 tier) ²¹ . <i>No equity required</i> – backed by a non-profit fund to subsidize costs ²² . Offers on-demand or long-term leases; scale from single GPU to thousand-GPU clusters in 15 minutes ²³ .	<i>Philanthropic model:</i> Funded by billionaire Jed McCaleb's Navigation Fund with \ \$500M to make AI compute affordable ²² . Already leasing capacity to notable startups like Imbue, Character.ai, and Atomic AI .

Provider / Program	Infrastructure & Scale	Program Terms	Partners / Notes
Scaleway Startup Program (EU)	"AI Supercomputer" cluster with 1,016 NVIDIA H100 GPUs in Paris ²⁴ ; fully containerized with NVIDIA AI Enterprise software on sovereign EU cloud.	€9,000 in cloud credits over 6 months, plus technical guidance and support ²⁵ . No equity taken. Startups also get priority access to H100 instances and compliance-ready (GDPR) infrastructure ²⁶ .	<i>Alliances:</i> Works with NVIDIA Inception to reach AI startups ²⁷ . NVIDIA and Scaleway's 2023 partnership offers H100 access + LLM tooling to European startups ²⁸ ²⁴ . Emphasizes EU data sovereignty (important for European AI ventures).
Oracle for Startups (Global)	Oracle Cloud (OCI) with both NVIDIA A100/H100 and AMD MI300X GPU instances ²⁹ . OCI's newest clusters scale to <i>16,384 GPUs (MI300X)</i> and even larger upcoming Nvidia Blackwell superclusters (planned 131k GPUs) ³⁰ .	Free cloud credits (amount varies by startup stage) and 70% discount on OCI services for two years ³¹ . No equity required. Startups also get go-to-market support and Oracle customer network access.	<i>VC Partnerships:</i> Oracle has provided dedicated GPU clusters to VC-backed startups (e.g. an Oracle-managed cluster free for Index Ventures companies) ¹¹ . Oracle's aggressive GPU investments (including 30,000+ MI300X deal and planned Blackwell capacity) aim to attract AI startups to OCI ³⁰ .
Northern Data "AI Accelerator" (EU)	Northern Data GenAI Cloud in Frankfurt running NVIDIA HGX H100 pods (InfiniBand-connected), powered by 100% renewable energy ³² . European Tier-IV green data centers.	Complimentary H100 compute for selected early-stage startups. ³² Program cohort winners get free access to high-performance GPU clusters + mentorship from industry experts. No direct equity taken (ND may cultivate future clients).	<i>Initiative:</i> Launched 2024 to <i>"democratize access to AI"</i> ³³ . Sponsored by NVIDIA and Gigabyte for hardware; showcased startup successes (e.g. Voxist, GroundedAI) at its AI Accelerator Summit ³⁴ .

Provider / Program	Infrastructure & Scale	Program Terms	Partners / Notes
Reliance Jio & India AI Mission (India)	Planned hyperscale AI cloud in India via Jio–NVIDIA partnership. Targeting NVIDIA Blackwell GPUs (next-gen) with 2,000MW data center capacity – “ <i>over an order of magnitude more powerful than India’s fastest supercomputer</i> ” ³⁵ . Green-energy powered data center in Jamnagar.	GPU-as-a-Service at ultra-competitive rates – Jio aims to offer the world’s lowest pricing for AI compute to Indian startups and researchers ³⁶ . Part of a government-backed mission; likely no equity, but strategic national initiative.	<i>Strategic:</i> Jio secured priority supply of Nvidia’s Blackwell chips ³⁶ . By 2025, once live, this could give Indian AI startups affordable access to thousands of top-tier GPUs domestically. Jio’s model echoes its telecom disruption (low-cost data) now applied to AI ³⁷ ³⁸ .

Table: Key “Compute-for-Equity” and Startup Compute Programs (2024) ¹⁴ ²⁴ ¹¹

Infrastructure Models – GPUs, Orchestration, and Data Locality

Across these programs, the **infrastructure is cutting-edge and tailored for AI**. NVIDIA’s GPUs dominate: the **A100 and H100 Tensor Core GPUs** (40GB/80GB memory) became the industry standard for training large models in 2022-24. Most providers have built around H100 clusters, often in the thousands of GPUs ¹⁸ ²⁴. For example, CoreWeave and Lambda both offer on-demand H100 instances, and **Voltage Park’s entire 24k-GPU farm is H100-based** with NVLink and InfiniBand for multi-node training ²⁰ ¹⁹. **Next-gen hardware** is also on the horizon: Lambda is already advertising upcoming **NVIDIA Blackwell** GPUs (the expected post-H100 generation) on its roadmap ¹⁷, and Oracle announced availability of **Blackwell-powered “zettascale” superclusters** scaling to 131k GPUs in the near future ³⁰. In India, Jio’s partnership explicitly centers on **securing Blackwell GPUs for its new AI cloud** ³⁶.

Notably, **AMD’s MI300 series accelerators** are emerging as an alternative in some programs, addressing supply constraints and diversifying options. Oracle Cloud was the first major cloud to deploy the **AMD Instinct MI300X (192GB) GPUs** in 2024, with clusters up to 16,384 MI300X GPUs in one RDMA fabric ²⁹ ³⁹. This gives startups on OCI an opportunity to train LLMs on high-memory GPUs (able to fit ~66B-parameter models on one MI300X) ³⁹. Other clouds like Vultr have also begun offering MI300 instances ⁴⁰. While NVIDIA still has the software ecosystem edge, these programs are infrastructure-agnostic if it means more capacity for hungry startups.

Orchestration and deployment are key considerations. Many providers emphasize **bare-metal performance with containerized flexibility**. For example, Oracle’s GPU instances are bare metal (no hypervisor overhead) with fast interconnects ⁴¹, and Scaleway’s H100 cluster comes pre-integrated with NVIDIA’s AI Enterprise suite for container orchestration ²⁸ ²⁷. **Kubernetes-based platforms** are often used under the hood – CoreWeave’s “Mission Control” and Lambda’s one-click clusters suggest K8s or similar schedulers managing multi-GPU workloads. Startups can typically deploy via **standard Docker containers or APIs** and scale out to dozens or hundreds of GPUs on demand. **InfiniBand networking** is commonly featured for distributed training: Voltage Park uses NVIDIA Quantum-2 IB at 3,200 Gbps ¹⁹, Northern Data links H100 nodes with IB, and Oracle’s clusters leverage NVIDIA’s ConnectX-7 NICs for RDMA ³⁰. This ensures near-linear scaling and low latency for jobs like large language model training that span many GPU nodes.

Data locality and compliance have also become selling points. With concerns around data privacy and regional regulations, some programs cater to “**sovereign cloud**” requirements. For instance, Scaleway operates in France under EU data laws, allowing European startups to train on H100s while keeping data in-region – crucial for GDPR and industries like healthcare ²⁶. Similarly, Oracle highlights its various global regions (including EU and Asia) and even offers cloud deployments *within customers’ own data centers* for sensitive workloads. In India, the forthcoming Jio/NVIDIA infrastructure aims to keep AI data and computation on Indian soil, aligning with the country’s digital sovereignty goals. **Green energy** is another aspect of infrastructure: Northern Data touts that its GenAI cloud is powered by 100% renewable energy ³², and Jio is incorporating green power in its new data center to cut costs and carbon footprint ³⁷. These factors can be important for both compliance and cost (sustainable energy often means less volatile power costs for energy-hungry GPU farms).

Program Structures – Pricing, Equity, and Deal Terms

The **business models** behind these programs vary in structure, balancing the cost of expensive hardware with the upside of startups’ success:

- **Credits and Discounts:** Most cloud providers offer **free credits to startups up front**, often time-limited. For example, Scaleway’s program gives €9k over 6 months ²⁵, and **Oracle for Startups provides an initial credit pool (often \$10k+ scaling up by stage) and then a steep 70% discount on further usage for two years** ³¹. Google’s AI startup program (through Google for Startups Cloud) similarly offers up to \$350k in credits for eligible companies ⁹ – and even more (up to \$500k) for those in the special Sequoia partnership ⁸. These credits essentially act as **compute funding** – they reduce burn rate dramatically, which can be as valuable as a cash investment for an AI company.
- **Equity for Compute Deals:** Some programs explicitly **take equity or warrants** in exchange for providing subsidized compute. The clearest example is a16z’s Oxygen: instead of charging normal rates or giving pure grants, a16z negotiates a stake in the startup commensurate with the value of the GPU access ². While exact percentages aren’t public, such stakes are typically in the single digits, akin to an accelerator. This can be structured via a **SAFE (Simple Agreement for Future Equity)** or a side letter in the VC deal – effectively treating, say, \$X of GPU time as \$X investment. The benefit to startups is they get critical compute *now* and the cost is diluted into their next equity round. As TechCrunch noted, Oxygen can prevent founders from having to raise an overpriced round just to pay cloud bills ⁴² – instead, they get GPUs from their investor and “fundraise” less, preserving equity for when valuation is fairer ⁴². Smaller VC clusters (Friedman/Gross’s Andromeda, Conviction, etc.) likely operate on a similar “in-kind investment” logic, though they may simply bundle it with their standard equity check.
- **No-Equity Grants:** On the other end, some initiatives are essentially **promotional or mission-driven** and do not ask for equity at all. CoreWeave’s and Scaleway’s accelerators fall here – they give credits and support in hopes that the startups eventually become paying customers on their platform. **Northern Data’s free H100 program** likewise does not take equity, nor does it require purchase commitments – it’s positioned as advancing AI innovation (and of course building relationships with promising companies). **Voltage Park’s model** is unique in that it’s **non-profit-backed**; it offers *ultra-low pricing* (\$1.89/GPU-hour for H100) with no strings, effectively a philanthropic subsidy to the entire AI startup ecosystem ²¹ ²². Similarly, government or industry-backed programs (like India’s IndiaAI mission via Jio) aim to provide *at-cost or below-cost computing* to domestic startups as an investment in the nation’s tech growth, rather than for

equity. In these scenarios, the “return on investment” is indirect – future cloud revenue, ecosystem growth, or societal benefit, instead of a slice of the startup’s stock.

- **Hybrid and Milestone-based Models:** Some partnerships combine elements. For instance, a VC firm might invest cash *and* throw in cloud credits from a partner cloud (common in deals where the VC has an arrangement with a provider). There are also concepts of **milestone-based compute allocations** – while not widely publicized, it’s conceivable a deal could stipulate that a startup gets, say, an extra \$100k in cloud credits unlocked when they achieve a certain technical milestone or next funding round. This mirrors how incubators might release funds on milestones. Although concrete examples of *compute-for-milestones* aren’t public, **the stage-gated approach is used in some accelerators for cash and could apply to credits as well** (e.g. Microsoft for Startups has levels of Azure credits as a startup progresses) ⁴³ ⁴⁴ .
- **Pricing Structures:** For startups that do pay (or after credits are exhausted), **pricing per GPU-hour** is a critical factor. Specialized GPU clouds often undercut major providers: e.g. **Voltage Park’s \$1.89/hour H100 rate** undercuts equivalent H100 instances on AWS (which can be \$2.75–\$3+/hour) – this is possible because of subsidized hardware costs ²¹ . CoreWeave also advertises lower-than-big-cloud pricing for GPUs (with spot instances or reserved deals offering further savings). Many providers offer **tiered pricing**: on-demand hourly vs. reserved usage (12+ month contracts) at lower rates ⁴⁵ . Startups with consistent needs can sign longer leases (Voltage Park and others offer monthly leasing of full GPU servers, which can be easier to budget for). Some platforms (e.g. Lambda, Paperspace) also have **spot market or spare capacity pricing** which can drop costs if the timing is flexible. The net trend is that **startups in these programs are paying significantly less per GPU-hour than the market norm**, whether via credits or discounts or special rates. This “compute leverage” is exactly the point – *providing compute cheaply in the early days so the startup can focus on innovation instead of cloud bills*.

Strategic Alliances Driving These Models

A remarkable feature of the 2024–2025 landscape is the **tight collaboration between infrastructure providers and venture investors** – a symbiosis fueling the venture compute trend:

- **Cloud-VC Partnerships:** Major cloud companies actively seek out VC alliances to onboard promising AI startups. Google’s partnership with Sequoia (and earlier Y Combinator) to give free credits is one example ⁸ ⁴⁶ . Oracle’s collaboration with Index Ventures is another ¹¹ , as is its work with Microsoft’s M12 and others to host their portfolio compute ¹¹ . These partnerships are win-win: the cloud gets high-growth customers locked in early, and the VC’s companies get resources that would normally be out-of-reach. Startups often receive *concierge support* in these arrangements too – e.g. Google assigns solution engineers to help Sequoia startups optimize costs ⁴⁷ . Such deep integrations blur the line between investor and vendor in a constructive way.
- **Hardware Vendor Programs:** **NVIDIA’s Inception** program is a global startup support program that has grown to tens of thousands of members. While Inception itself doesn’t hand out equity or run a cloud, it **connects startups with GPU resources and discounts**. For instance, **NVIDIA partnered with Scaleway to funnel Inception startups to Scaleway’s H100 credit program in Europe** ²⁷ . Inception members get benefits like early access to NVIDIA hardware, discount vouchers (for NVIDIA hardware or cloud credits on partners), and technical training ²⁷ . This effectively complements venture compute programs – a startup might get cloud credits from a provider plus belong to NVIDIA Inception for additional support. NVIDIA also works closely with

cloud providers (AWS, Azure, Oracle, etc.) to ensure their latest GPUs (A100, H100, now Blackwell) are available to startups – sometimes **seeding programs with loaner hardware or special pricing** for those use cases.

- **Startup Funds by AI Giants:** The **OpenAI Startup Fund** (a \$100M fund launched 2021) exemplifies another alliance type: backed by OpenAI and its partners (including Microsoft), it invests in startups aligned with AI advancements. While not openly advertised, it's understood that companies in OpenAI's portfolio (e.g. Descript, Mem, etc.) **gain privileged access to OpenAI's technology and Azure's AI infrastructure** as part of that relationship. Microsoft, having invested billions in OpenAI itself, has a strategic interest in routing AI startups to Azure – and indeed **Microsoft's expanded Azure credits program for AI startups** (via Microsoft for Startups) now offers *up to \$150k in Azure credits plus \$2,500 in OpenAI API credits* for qualifying companies ⁴⁸ ⁴⁹. This shows how cloud vendors are aligning with AI research leaders to attract the next generation of innovators.
- **Venture + Nonprofit Initiatives:** An interesting twist is the involvement of nonprofits and governments. We saw how **Jed McCaleb's Navigation Fund bankrolled Voltage Park with \$1B to offer cheap AI compute** ²² – essentially a philanthropic venture compute model to *“advance AI capabilities and democratize access”*. Governments are also stepping in: the **IndiaAI compute mission** aggregates resources from multiple data center firms (Jio, CtrlS, E2E Networks, etc.) to create a national GPU pool at low cost for startups ⁵⁰ ⁵¹. Such alliances often tie into broader strategic goals (e.g. ensuring a country's startups aren't left behind in the AI race due to lack of infrastructure).

In summary, 2024's venture compute landscape is characterized by **tight-knit alliances across the stack – VCs, cloud providers, hardware makers, and even nonprofits/governments are coordinating to fuel AI startups with compute power**. This cooperative approach is a response to the unique challenges of this AI wave: unlike past software startups that needed only laptops and cloud VMs, today's cutting-edge AI ventures require *access to supercomputing-class hardware from day one*. The industry has recognized that and is collectively lowering the barriers through innovative partnerships.

2024–2025 Trends and Outlook

The concept of trading compute for equity or providing “compute capital” is rapidly maturing. In 2024 we saw **major deals and expansions** that solidify this model:

- **Explosion of GPU Capacity for Startups:** Venture-backed GPU clusters went from experimental to mainstream. By late 2024, a16z had arguably the largest GPU cluster of any private entity ¹, and multiple other funds launched similar (if smaller) efforts ⁶ ¹¹. This arms race among investors is likely to continue into 2025, especially as new GPU generations (H200, Blackwell, MI300X, etc.) come online. **The result: top AI founders now consider access to compute as part of an investor's pitch.** It's not just “how much money can you invest?” but “can you also get me the GPUs or credits I need to train a billion-parameter model next month?”
- **Startups Scaling Faster (and Cheaper):** Early evidence shows startups in these programs can achieve milestones that would have been impossible otherwise. For example, one of the first CoreWeave Accelerator startups, Descriptor.AI, noted that *“the CoreWeave program has been invaluable – it allows us to deploy the GPU-accelerated workloads we need to grow our business”* ⁵². Many similar anecdotes are emerging of companies that trained larger models or iterated faster thanks to free or subsidized compute. By alleviating the **need to spend millions on cloud bills**

in seed stage, these programs let startups direct funds toward hiring talent and product development. This could yield **better AI products on shorter timelines**, a crucial competitive factor in a hot market.

- **Ecosystem Fragmentation and Choice:** A positive trend for startups is the **diversification of compute suppliers**. In 2020, if you needed serious AI compute you essentially went to AWS, Azure, or GCP and paid list price (with maybe a startup credit promo). In 2025, a 10-person startup has a menu of options: specialized clouds like CoreWeave or Lambda with potentially better pricing/performance, non-profit or government-backed pools like Voltage Park or Jio that offer low rates, and big-cloud programs with hefty credits. This **competition is driving innovation in pricing models** (spot markets, reserved deals, credits-for-equity) and in technology (e.g. Oracle pushing multi-vendor GPU clusters, NVIDIA enabling cloud-native supercomputers for partners). We're also seeing startups mix-and-match – using credits on one cloud for part of work, borrowing time on a VC's cluster for another, etc., enabled by more portable ML workflows (Docker containers, PyTorch Lightning, Ray, etc. that can run anywhere).
- **Challenges and Limitations:** Despite the promising model, venture compute programs do have limits. Often the **free credits or subsidized access are capped** – startups may still face a steep cloud bill as they grow usage beyond the program's grant. There's also a question of **sustainability**: GPU prices are high, and not every provider has a billionaire benefactor or \$7.5B debt financing like CoreWeave ⁵³. As demand keeps booming, these programs might become more selective (choosing startups with the most promise, or tying compute to equity more tightly). Additionally, some startups worry about **lock-in or dependency** – e.g. if you build heavily on a free cluster and that program ends, migrating workloads can be non-trivial. Providers are trying to mitigate this by using standard platforms (so workloads are portable) and by extending discounts gradually rather than a sudden cutoff.
- **The Democratization Mission:** A clear trend in messaging is *“democratizing AI compute.”* Whether it's CoreWeave's founding mission or Voltage Park's nonprofit charter, or Jio's ambition to make GPU power as accessible as cheap data, the industry narrative is that **compute should not be a monopoly of tech giants**. Jensen Huang of NVIDIA himself remarked that there's unprecedented demand from startups – and programs like Inception and these cloud partnerships are meant to **empower the next OpenAI or DeepMind that might currently be two founders in a garage**. The **AI divide** (startups vs. Big Tech in access to compute) is something these 2024 initiatives actively aim to bridge.

Going into late 2025 and beyond, we can expect **further innovation in venture compute**. Some foresee *marketplaces* where unused GPU capacity from one company can be lent to another for equity or tokens. (In fact, NVIDIA has hinted at software to create cloud GPU marketplaces ⁵⁴.) Others predict large AI labs might spin out their infrastructure as a service to startups – for instance, if OpenAI built excess capacity for itself on Azure, it could allocate some to its Startup Fund companies as a bundled perk. **Traditional accelerators and VCs may all need a compute strategy**: just as having a network of advisors became standard, having a go-to GPU supply might become a must-have for serious AI investors.

In summary, **“compute-for-equity” has moved from a novel idea to a practical instrument in launching AI startups**. The landscape now spans dedicated GPU clouds trading credits for future stakes, venture firms buying hardware instead of just writing checks, cloud providers co-investing via generous credits, and even nations treating AI compute as infrastructure for innovation. This convergence of infrastructure and venture capital is enabling startups to push the state of the art in AI –

turning GPU horsepower into the new venture capital, and in doing so, accelerating the next generation of AI breakthroughs.

Sources: The information above is drawn from a range of 2022–2025 reports and announcements, including TechCrunch and Forbes coverage of VC GPU clusters ³ ² ⁶ , official program pages and blogs from CoreWeave, Scaleway, Oracle, etc. ¹⁴ ²⁴ , and news on partnerships and infrastructure deployments from NVIDIA, Reuters, and others ³⁶ . These illustrate the rapidly evolving market of venture compute and its real-world implementations.

¹ ² ³ ⁴ ⁷ ⁴² Andreessen Horowitz is renting a bunch of Nvidia GPUs to its portfolio companies | TechCrunch

<https://techcrunch.com/2024/10/23/andreessen-horowitz-helps-founders-meet-compute-needs-with-oxygen-private-gpu-cluster/>

⁵ ⁶ ¹¹ Los inversores en IA están cortejando a las empresas emergentes con equipos informáticos masivos

<https://forbes.com.mx/los-inversores-en-ia-estan-cortejando-a-las-empresas-emergentes-con-grupos-informaticos-masivos/>

⁸ ⁹ ¹⁰ ¹² ⁴⁶ ⁴⁷ Google (GOOGL.US) partners with Sequoia Capital on cloud

<https://www.ainvest.com/news/google-googl-us-partners-with-sequoia-capital-on-cloud-2410101075f284a874752f3c/>

¹³ ¹⁴ ¹⁵ ¹⁶ ⁵² coreweave.com

<https://www.coreweave.com/blog/coreweave-launches-startup-accelerator-program-democratizing-access-to-gpu-compute-in-the-cloud>

¹⁷ Lambda | GPU Compute for AI

<https://lambda.ai/>

¹⁸ ²⁰ ²¹ ²² Voltage Park plans to lease \$500M worth of Nvidia H100s • The Register

https://www.theregister.com/2023/10/30/voltage_park_gpu_farm/

¹⁹ ²³ ⁴⁵ Voltage Park, Next Generation GPU Cloud Infrastructure for AI

<https://www.voltagepark.com/>

²⁴ ²⁶ ²⁷ ²⁸ NVIDIA and Scaleway Speed Development for European Startups and Enterprises | NVIDIA Blog

<https://blogs.nvidia.com/blog/scaleway-european-startups-inception/>

²⁵ Startup Program - | Scaleway

<https://www.scaleway.com/en/startup-program-new-form/>

²⁹ ³⁰ ³⁹ ⁴¹ AMD GPUs join Nvidia's in Oracle Cloud Infrastructure | Fierce Electronics

<https://www.fierceelectronics.com/ai/amd-gpus-join-nvidias-oracle-cloud-infrastructure>

³¹ [PDF] Oracle for Startups VC Connect

<https://www.oracle.com/a/ocom/docs/vc-connect-brochure.pdf>

³² ³³ ³⁴ AI Accelerator - Home

<https://northerndata.de/ai-accelerator>

³⁵ Nvidia partners with Reliance Jio for 2,000MW of AI data centers ...

<https://www.sdxcentral.com/news/nvidia-partners-with-reliance-jio-for-2000mw-of-ai-data-centers-tata-for-ai-cloud-in-india/>

³⁶ ³⁷ ³⁸ Jio Platforms: Jio brings data playbook to build out accessible and affordable AI - The Economic Times

<https://economictimes.indiatimes.com/tech/artificial-intelligence/jio-brings-data-playbook-to-build-out-accessible-and-affordable-ai/articleshow/116773997.cms?from=mdr>

- 40 **AMD MI300 Gains Momentum with New Vultr & Oracle Cloud Wins**
<https://nand-research.com/amd-mi300-gains-momentum-with-new-vultr-oracle-cloud-wins/>
- 43 **Microsoft for Startups Founders Hub**
<https://foundershub.startups.microsoft.com/>
- 44 **Microsoft Startups \$150k Funding- everything you need to know**
https://www.reddit.com/r/AZURE/comments/1e2fiz9/microsoft_startups_150k_funding_everything_you/
- 48 **How to Get \$2500 in OpenAI Credits - API**
<https://community.openai.com/t/how-to-get-2500-in-openai-credits/881589>
- 49 **Just received USD 2,500 OpenAI Credits Join Microsoft for Startups...**
https://www.linkedin.com/posts/gateremark_just-received-usd-2500-openai-credits-activity-7213585227489636352-7SR_
- 50 **Reliance Jio, Other GPU L1 Bidders In IndiaAI Mission To ... - Inc42**
<https://inc42.com/buzz/reliance-jio-other-gpu-l1-bidders-in-indiaai-mission-to-get-workloads-soon/>
- 51 **India's Reliance partners with Nvidia to build large language model**
<https://techcrunch.com/2023/09/08/reliance-nvidia-india/>
- 53 **CoreWeave raises \$2.3 billion in debt collateralized by Nvidia chips**
<https://www.reuters.com/technology/coreweave-raises-23-billion-debt-collateralized-by-nvidia-chips-2023-08-03/>
- 54 **Nvidia software aims to create marketplace for AI computing power**
<https://www.reuters.com/business/media-telecom/nvidia-software-aims-create-marketplace-ai-computing-power-2025-05-19/>