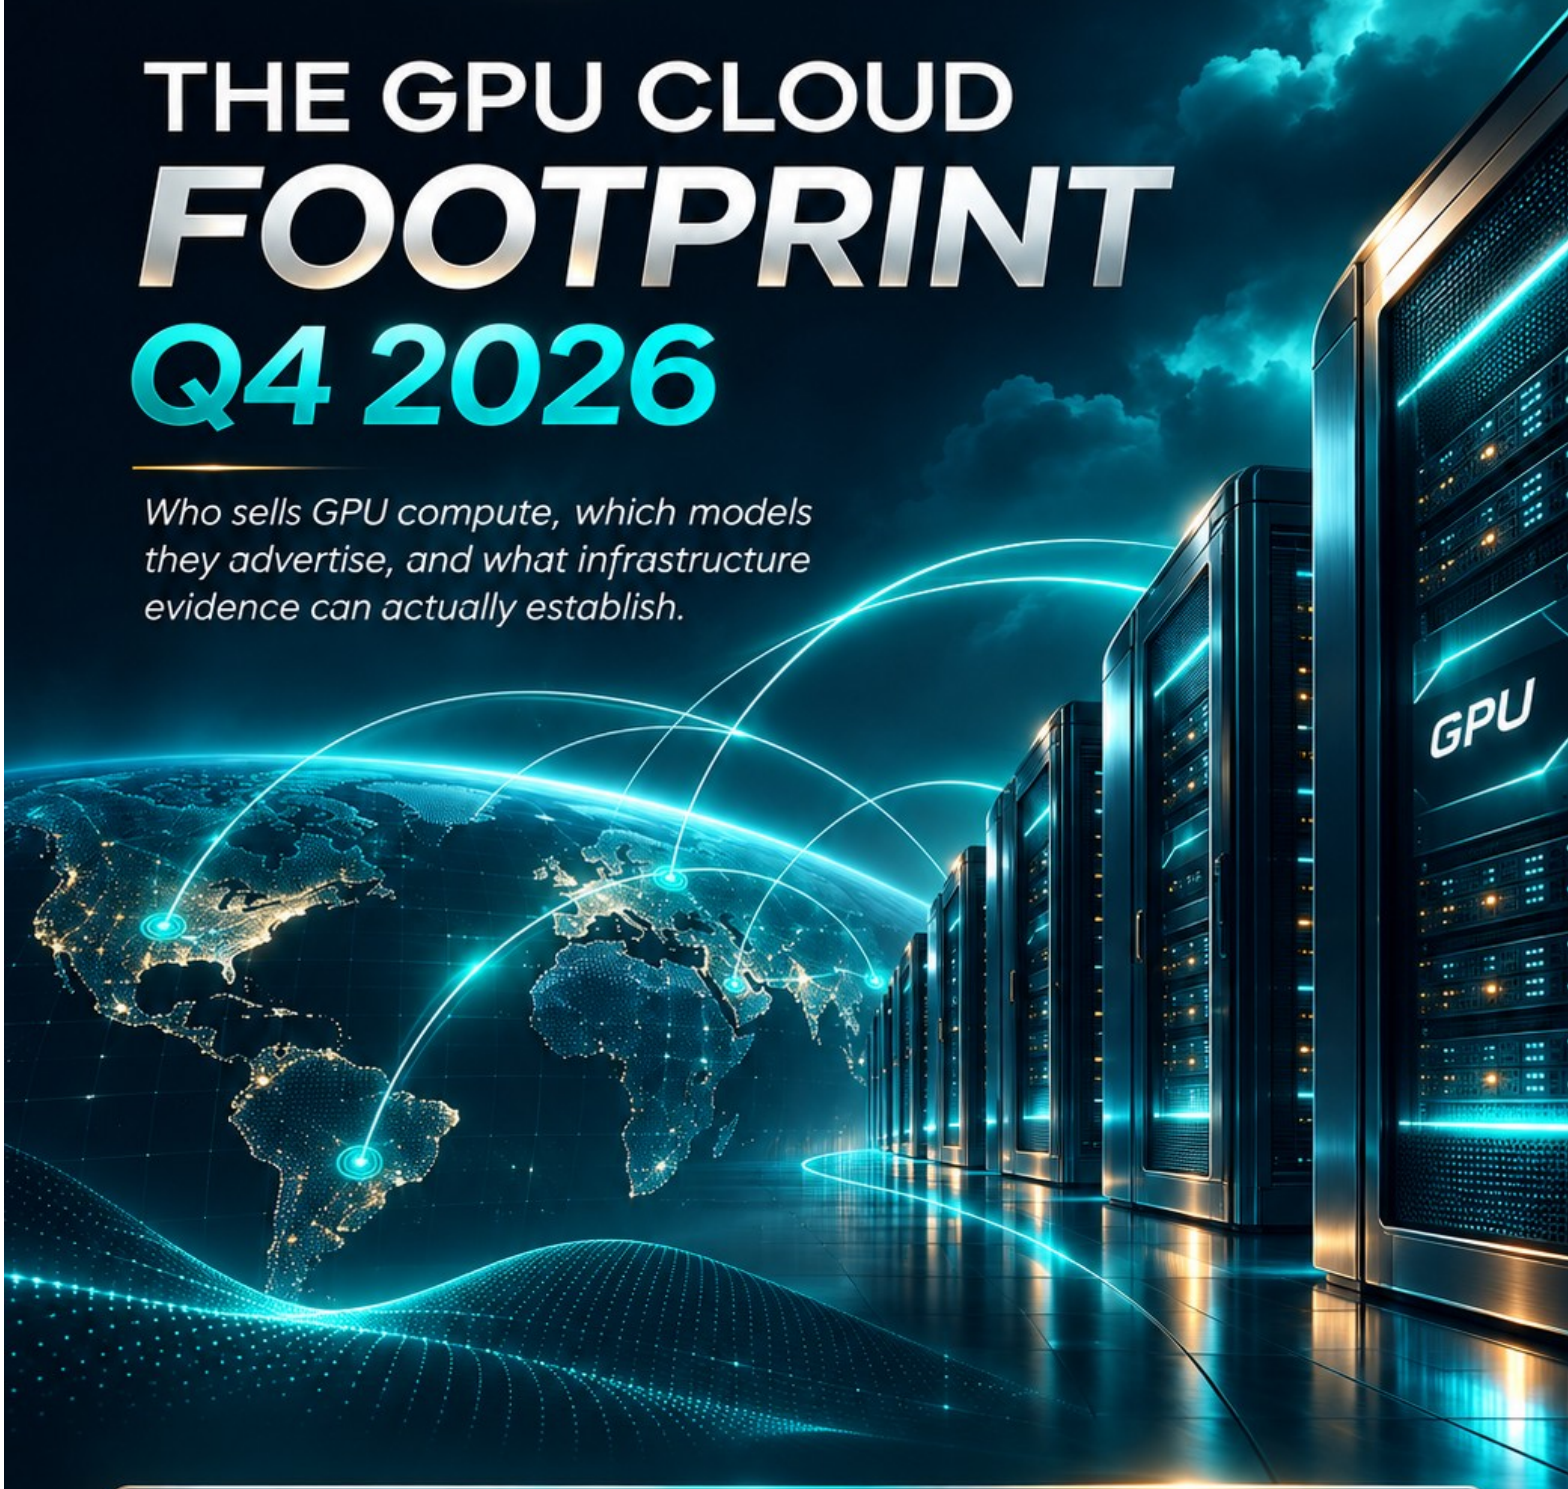
VIABANDWIDTH

RESEARCH / INFRASTRUCTURE INTELLIGENCE

THE GPU CLOUD FOOTPRINT Q4 2026

Who sells GPU compute, which models they advertise, and what infrastructure evidence can actually establish.



339	243	222	126
CONFIRMED DOMAINS	DIRECT-CLASSIFIED	NAMED-GPU DOMAINS	DOMAINS WITH PREFIXES

OCTOBER 2026 EVIDENCE REFRESH | CHECKPOINT 162

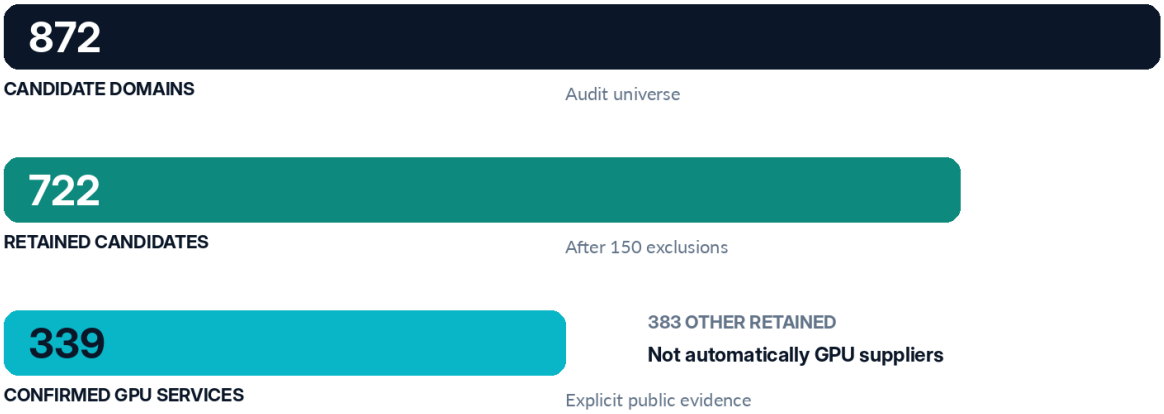
01 / EXECUTIVE READ

The footprint behind the GPU market

A market intelligence report for GPU cloud operators, infrastructure researchers and technical buyers. October 2026 evidence, reconciled through November 5.

The market behind the GPU listing

The GPU market is wider than the familiar cloud brands. ViaBandwidth screened 872 candidate domains and found explicit GPU-service evidence for 339 under its frozen October rules. A November review then identified 387 provisional supplier groups, including hyperscalers absent from the earlier screen. Those figures describe different populations, not market growth. The commercial question is how much of each supplier's offering and operating footprint can be substantiated.



What this means for service providers

A GPU model on a product page makes a service visible. Model-by-region availability, capacity-sourcing disclosure and clear operating responsibility make it easier to assess. This report separates published GPU offers from network presence, service geography and verified infrastructure, helping operators see where the market is transparent and where it is not.

How to read the evidence

A provider domain, a corporate group, a marketplace and a GPU fleet owner are different entities. Listed GPU models do not establish live inventory. Network records do not establish deployment sites. July and October populations cannot support a growth rate.

Source throughout: ViaBandwidth checkpoint 162, October 25–28, 2026. Publication-safe aggregates supplied in the accompanying research package.

02 / THE EVIDENCE POPULATION

Where the 872 candidates went

The difficult part of sizing the GPU market is separating genuine service offers from adjacent businesses, duplicate brands and unverified claims.

The October candidate screen retained 722 domains, including a frozen 339-domain explicit-service cohort. Subsequent November investigation restored candidate providers and documented duplicate brands, exclusions and unresolved sites. Therefore 339 must be labelled an October evidence cohort, not the current size of the GPU supplier market. See the updated supplier-group analysis later in this paper.

The largest categories outside the strict confirmed set

122 Insufficient evidence, access or configuration

121 Outside bounded GPU scope / general compute

48 Colocation, managed or sales-led infrastructure

24 Alias / canonical family / duplication treatment

16 Hardware sales, procurement and integration

13 Software / orchestration / inference without raw GPU tariff

A retained candidate is not automatically a GPU provider

Other cases include GPU-adjacent services (13), lifecycle limitations (10), advertised GPU services outside the strict label (7), broker/marketplace referrals (6) and research-only access (3). These groups reconcile to the 383 other retained domains; they cannot be added indiscriminately to the 339.

> This October distribution reconciles its own original screening rules. The November 5 registry uses a more granular disposition process and different inclusion rules; its updated counts must not be added to or subtracted from these bars as if they were growth.

03 / PROVIDER COMPOSITION

Who was in the October evidence cohort?

Where a supplier sits in the sales chain shapes what customers can verify about capacity, location and accountability.



Direct	71.7% of confirmed domains	243
Reseller	11.8% of confirmed domains	40
Marketplace	4.4% of confirmed domains	15
Unclassified	12.1% of confirmed domains	41

The graphic describes the frozen October cohort of 339 confirmed GPU-service domains, not November's provisional group-level supplier population. Its 243 direct, 40 reseller, 15 marketplace and 41 unclassified roles came from rule-based directory classifications. They do not establish physical GPU ownership.

What sort of businesses are these?

Traditional cloud 103 domains	**Neocloud** 62 domains	**Sovereign / regional** 44 domains
Residual / combined 34 domains	**Marketplace sector** 26 domains	**Managed ML ops** 24 domains
Dedicated hosting 19 domains	**Distributed** 16 domains	**Colo + GPU** 11 domains

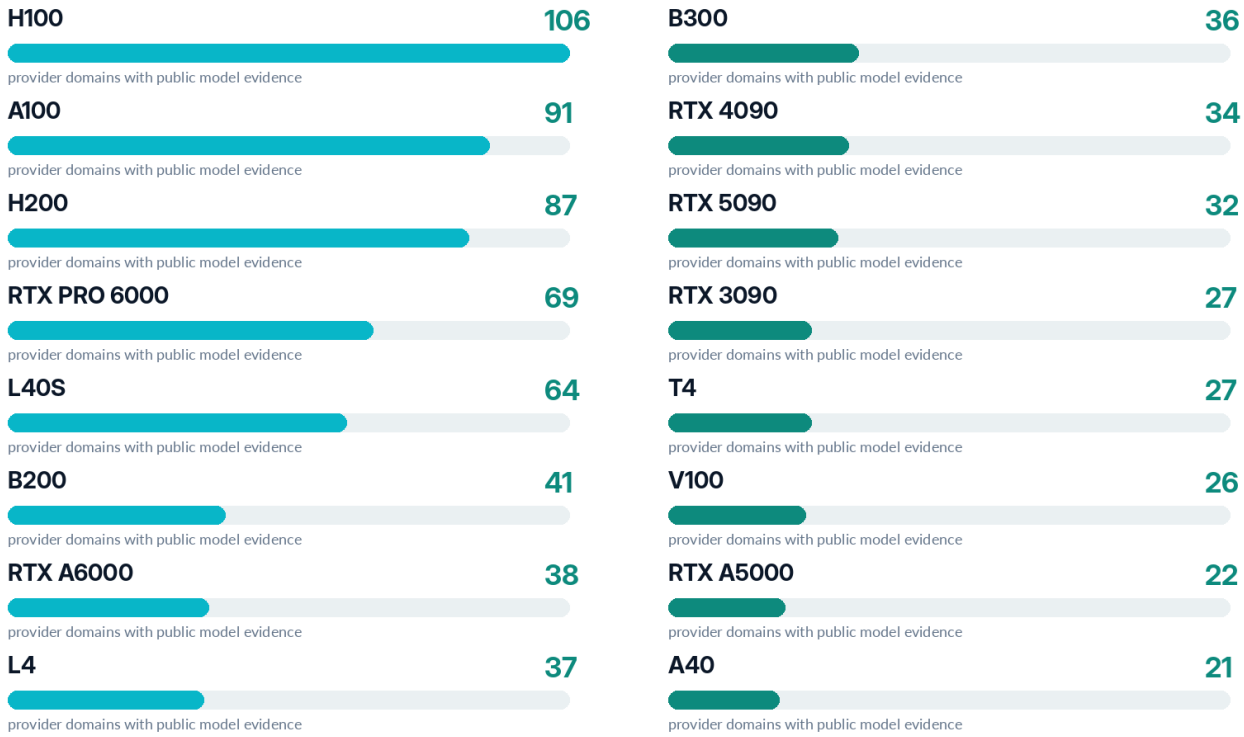
All 339 domains reconcile. The 34-domain residual combines missing or unmapped sector labels with a low-frequency mapped sector under the disclosure rule. Sector marketplace (26) is a business-positioning category; operating-type marketplace (15) is a distinct supply-intermediation classification.

04 / ACCELERATORS

Which GPUs appear in public evidence?

Published model lists describe the menu, not the size or location of a live GPU fleet.

Within the frozen 339-domain October cohort, 222 (65.5%) have a bounded named-GPU model record; the 117 without named-model evidence are unknown, not model-free. Models can appear across multiple providers, and availability at a particular data center is not established by an aggregate model label.



Leading models within 339 confirmed domains, October 2026. Counts are conservative minimums based on bounded public model evidence; models overlap across domains. Only 222 domains have named-model evidence. Values do not measure installed inventory, stock, or available fleet capacity.

Pricing evidence is a separate layer

The scoped price ledger contains 377 H100 offering rows from 90 domains, 394 A100 rows from 85 domains, and 318 H200 rows from 81 domains. Offer rows are not equivalent GPU-hour prices, and a priced configuration does not guarantee stock.

> For context only, July reported 339 H100-listing operators and 294 A100-listing operators among 588 directory records. Those counts used different evidence rules and cannot be compared with October minima as a change in availability.

05 / ADVERTISED MODELS BY OPERATING TYPE

The same accelerator, different route to market

The same accelerator can be sold by its operator, a reseller or a marketplace. The route to market changes what can be checked.

GPU MODEL	DIRECT	RESELLER	MARKETPLACE	UNCLASSIFIED
H100	73	20	6	7
A100	66	15	5	5
H200	61	13	6	7
L40S	40	11	6	7
B200	26	10	<5	<5
B300	25	6	<5	<5

Model evidence by directory operating type. Direct here means a classified business role, not audited GPU ownership. Suppressed cells are omitted, not zero; totals count provider domains, not chips or rentable stock.

Generation presence is overlapping

64
FRONTIER
B100, B200, B300, GB and newer AMD

186
ESTABLISHED
H100, H200, L4, L40S and others

129
PRIOR GENERATION
A100, V100, T4 and others

Minimum confirmed domains with at least one named accelerator from each generation grouping. A domain can appear in multiple groups; these totals should not be summed.

The 339-domain evidence boundary matters: the 4,511 scoped public-price rows are a different unit and cannot be converted into a universal count of available GPUs.

06 / NETWORK INTELLIGENCE

Network indicators in the frozen October cohort

Network identifiers reveal something about an organization's connectivity. They do not tell us which buildings host its GPUs or who owns the accelerators.

Among 243 direct-classified domains, 103 have nonzero retained routed-prefix counts. Across all 339 confirmed domains, 126 have nonzero prefix counts and 133 retain at least one ASN identifier. The stored "network_verified" flag equals prefix_count > 0; it is not an additional independent verification method.

103 / 243
DIRECT-CLASSIFIED
42.4% with nonzero prefix count

126 / 339
ALL CLASSES
37.2% with nonzero prefix count

133 / 339
ASN EVIDENCE
39.2% with retained ASN ID(s)

RETAINED ROUTED PREFIX COUNT / CONFIRMED DOMAINS

0	1–9	10–99	100+
213 domains	33 domains	51 domains	42 domains

"Network-verified" equals prefix_count > 0 in this snapshot; it is not a second independent test.

The 126 nonzero-prefix results and 126 "network-verified" results are identical by definition, not two independent confirmations. Prefix bands use merged upper ranges for disclosure. Routing scale is not GPU inventory.

What network evidence can and cannot show

ASN identifiers and observed prefixes are network indicators. They may relate to an associated network entity and do not confirm legal ownership of a server, installation of a GPU, its deployed country, or live ordering capacity. Zero recorded prefixes does not prove resale.

07 / GEOGRAPHIC EVIDENCE

Provider-claimed geography: useful but incomplete

Service-region claims matter to technical buyers, but they are not proof of GPUs installed at a particular building. The November evidence applies stricter sourcing rules.

GEOGRAPHY: PARTIAL PROVIDER-REPORTED EVIDENCE

October cohort with carried-forward claims and limited new official-page evidence.

96 / 339	24	243 / 339
DOMAINS WITH CLAIMS	COUNTRIES REPRESENTED	GEOGRAPHY UNRESOLVED
28.3% of the confirmed cohort	Only among covered domains	No qualifying country evidence

Most frequently claimed countries within the partial sample:

United States 27	Germany 26	India 25	Netherlands 25	United Kingdom 20

Source limits: 90 baseline claims; +6 domains with new evidence; 0 independently supported GPU locations.

What the earlier State report established

The July edition reported operators advertising deployments in 48 countries, with the United States (42), Germany (32), India (32), the Netherlands (29) and the United Kingdom (28) leading its published country table. Those figures remain historical July observations.

What local directory evidence adds

The October review found legacy provider-claimed geography for 96 of 339 domains across 24 named countries. Most were older directory claims without retained quotations. The November master review separately accepted GPU-service country evidence for 89 company IDs in 41 countries (149 company-country pairs), drawing on source-linked product and location records. These are not the same population or acceptance standard.

> Neither geography series is a physical facility census or independent audit of live GPU installations. The new location amendments add 27 source-linked records, some duplicating company-country pairs; they must not simply be added to the 149 accepted pairs. The following research update details valid comparisons and examples.

08 / Q3 AND Q4: THE CORRECT COMPARISON

Two snapshots. Two definitions. No invented trend.

The published July cohort cannot be reconstructed record-for-record.

JULY Q3 PUBLICATION

588

Indexed GPU operator records

Directory listing / model-tag basis

NO DELTA

OCTOBER Q4 CHECKPOINT

339

Confirmed provider domains

Strict public-evidence basis

What Codex recovered

The July publication reported 588 indexed operator records, but the exact frozen rows could not be recovered. July directory backups expose 589 GPU-tagged rows, a later curated API service filter yielded 187, and the October research cohort contained 339 evidence-confirmed domains. These are distinct selection rules rather than comparable market counts. Source filters and row-level membership must be reconciled before any growth calculation.

Why counts cannot be subtracted

An indexed tag, a verified commercial GPU offer, a corporate group and an independently operated GPU site are separate concepts. No numerical percentage change between July and the November group registry is supported; missing membership and differences in category definitions preclude a matched-cohort comparison.

What this evidence can support going forward

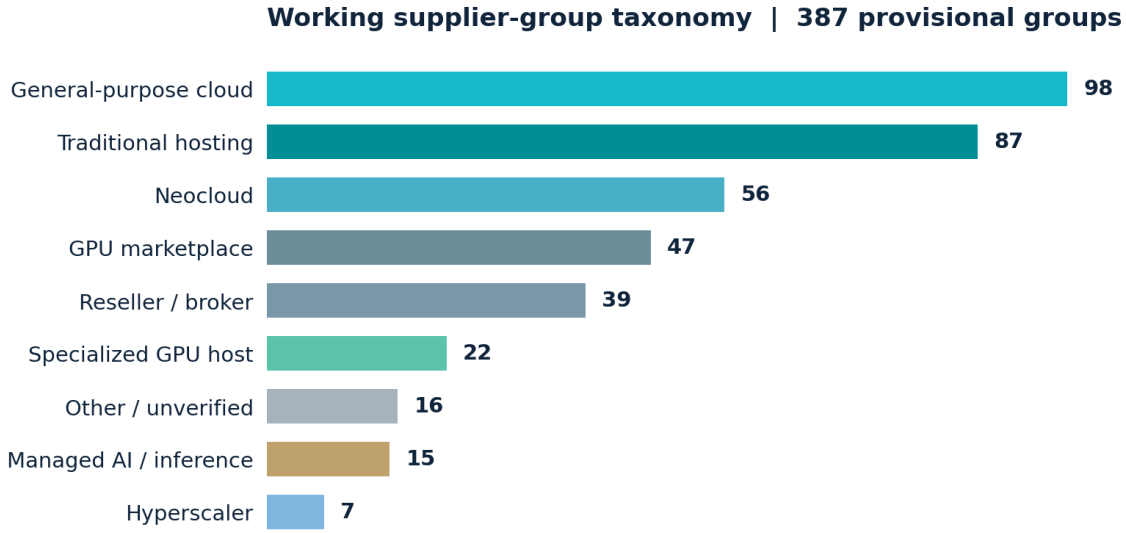
The October 339-domain cohort remains a reproducible baseline for model and network diagnostics. The November 387-group working registry broadens coverage but requires legal/brand deduplication, service-scope checks and a repeatable taxonomy before it can become the headline supplier census.

09 / UPDATED PROVIDER REGISTRY

A market of sellers, not one fleet

The November reconciliation expands the picture of who participates in GPU service delivery, while keeping corporate ownership and operating roles distinct.

872 original domains / 114 targeted checks completed / 387 provisional supplier groups



Source: provisional_387_group_taxonomy.csv; includes marketplaces, brokers, managed platforms and seven hyperscalers. Many labels remain inherited, not freshly audited. Chart categories sum to 387.

Why this is not “supplier growth”

The earlier 339 count described domain-level evidence inside a frozen October cohort. The 387 figure is a November group-level working reconciliation that includes restored providers and four absent hyperscalers. Changing the denominator is a research improvement, not evidence that 48 new suppliers entered the market.

The ownership question remains unanswered

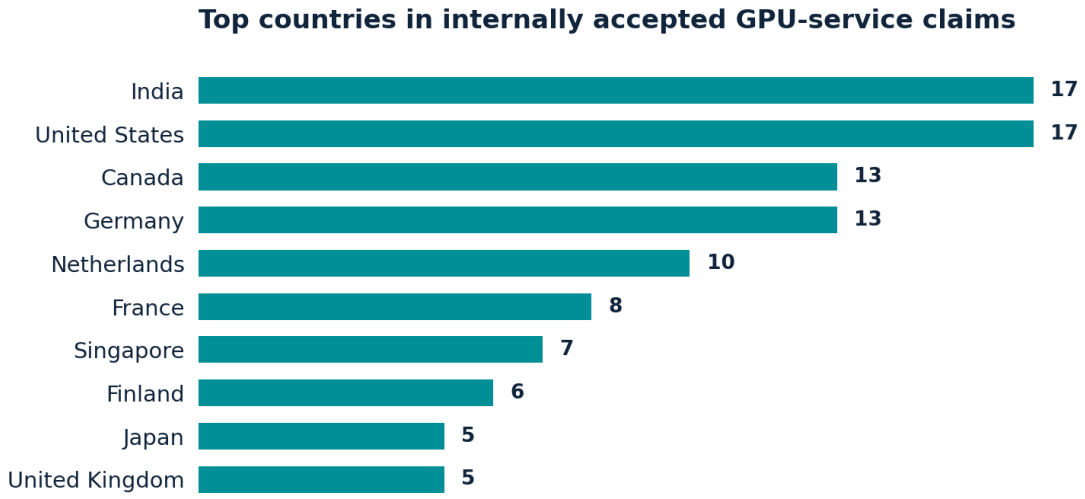
The 387 groups are not necessarily independent legal entities or owners of GPU stock. Brand relationships (including Hut 8/Highrise and NVIDIA/Brev), 16 inaccessible websites and inherited service labels remain as explicit review exceptions. No headline global supplier census is approved.

10 / COUNTRY-LEVEL GPU CLAIMS

Where providers say GPUs are served

Service regions help customers plan deployments. They do not, on their own, disclose the physical facility or who owns the GPU hardware.

89 company IDs / 41 named countries / 149 accepted company-country pairs



Source: master reconciliation view v_mr_accepted_service_countries, 5 Nov. Country count is a presence of accepted advertised service evidence, NOT installed capacity, unique data centers or availability. No Q3-to-Q4 country growth inference.

What the newer evidence adds

A separate 27-row November location-amendment ledger documents site-selectable offers and named deployments, including Beyond in Poznan (B200 GPUaaS), BIT in the Netherlands, LayerStack in Hong Kong, GTS in Romania and fromVOLT in Amsterdam. Multiple site listings for one company must not be counted as separate providers.

Example first-party sources: beyond.pl/en/services/gpu-as-a-service-gpuaas/ | bit.nl/en/services/gpu-hosting/ | gts.ro/en/gpu-as-a-service | fromvolt.com/compute . Details and stock may change.

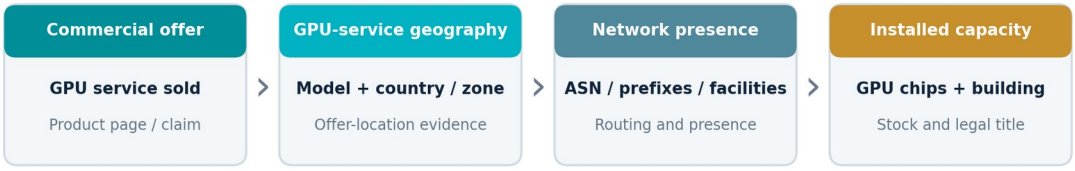
What remains unknown

The master geography ledger is supplier-advertised service coverage, not a GPU facility census. A separate 13-provider infrastructure sample now supplies 44 model-to-region/zone links across provider-reported operational rows. Cloud regions, PeeringDB sites, announced projects and live GPU stock remain distinct evidence categories.

11 / INFRASTRUCTURE REALITY CHECK

The four different kinds of footprint

A GPU cloud may involve a service seller, a hardware operator, a network operator and a building owner. Those roles need not belong to the same company.



Each layer answers a different question. Network presence cannot establish a physical GPU installation.

FROZEN OCTOBER NETWORK COHORT

339 provider network summaries | 133 domains with ASN IDs | 126 with retained prefix counts

The master archive also retains 4,833 facility observations (network presence). But its detailed prefix-observation and reverse-DNS tables are empty. A facility observation does not identify a GPU cluster; the retained “network verified” flag is simply derived from whether a prefix count is nonzero.

Why this matters for procurement

A provider’s H100 or B200 listing answers which products it advertises. It does not answer whether the model is in stock, in which region it runs, which party operates the hardware, whether dedicated GPU capacity is included in the contract, or whether traffic routes through an independent upstream. Those questions require offer-level and site-level evidence.

Evidence threshold for the next release

A targeted 13-provider validation supplies model-to-region evidence, but does not change the 339-domain October cohort or establish who owns installed accelerators. Keep the 387-group taxonomy provisional until legal-brand joins and inherited labels are checked. IP-level analysis still needs the separate local prefix/rDNS export.

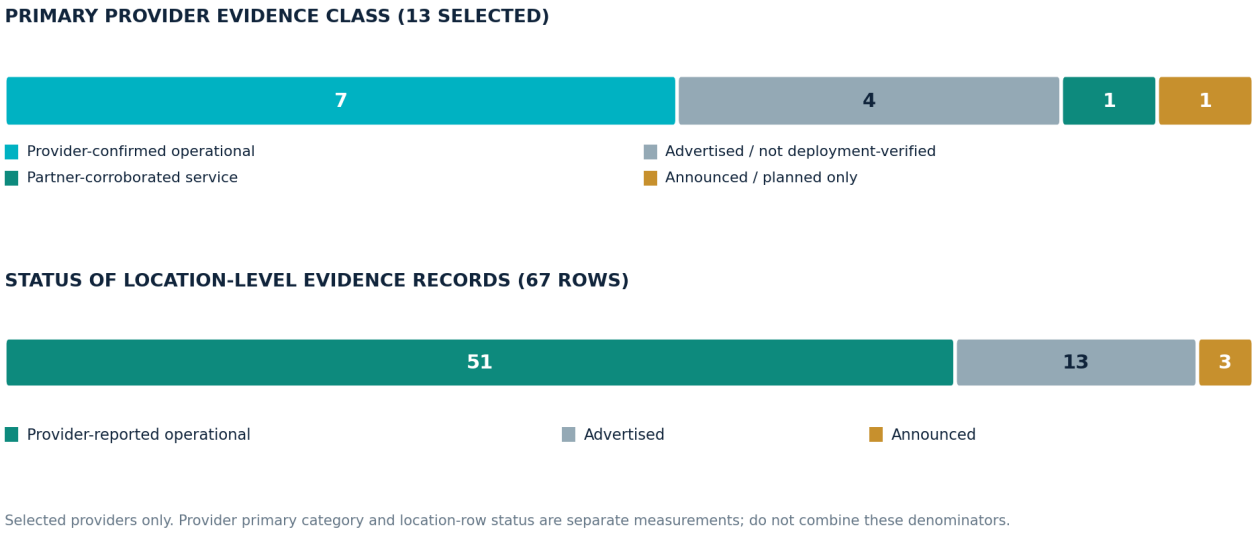
Reconciliation inputs: ViaBandwidth_Master_Provider_Reconciliation.zip; ViaBandwidth_Provider_Verification_2026-11-05.zip; ViaBandwidth_Footprint_Q4_Final_Research_Corrected.zip. No production database changes were made.

12 / TARGETED GPU INFRASTRUCTURE CHECK

Where the model meets the map

A selected 13-provider investigation traces published models to operational service regions and tests how much physical infrastructure information those claims reveal.

13 providers / 67 location records / 44 named-model operational links / 16 operational countries



Selected providers only. Provider primary category and location-row status are separate measurements; do not combine these denominators.

A critical denominator distinction

Nine providers have at least one operational service row, but only eight carry an operational primary provider-tier label. Civo also has operational LON1 coverage alongside a separate planned Hertfordshire project, so primary provider tiers and deployment-row statuses must not be conflated.

What the evidence actually reveals

IBM and DigitalOcean publish model-to-cloud-zone or data-center matrices, but not building-level proof of GPU ownership. Exoscale has additional partner corroboration for its service operator and GPU-zone matrix, not independent physical inspection. AlphaVPS names Telepoint and Hetzner as site operators for GPU nodes, separating GPU service operation from facility ownership.

Announced is not operational

Silicon Network's proposed clusters and Civo's planned Hertfordshire Vera Rubin project are excluded from operational counts. ColoCrossing's advertised GPU-capable services and ten cities lack model-specific city verification. All 275 retained network joins remain network-only observations, not GPU deployments.

Evidence basis: provider classification (13 rows); deployment ledger (67 rows); network join (275 rows). Examples: cloud.ibm.com/docs/vpc?topic=vpc-profiles; docs.digitalocean.com/products/droplets/details/gpu-availability/; ocre.vshn.ch/ai/; alphavps.com/solutions/ai-gpu; colocrossing.com/server/ai-dedicated-servers/. Reviewed 5 Nov 2026.

13 / FINDINGS AND METHODOLOGY

Research conclusions

A GPU listing makes a provider discoverable. Verifiable operating details make its footprint meaningful.

1. Supplier role matters.

A commercial GPU seller may be a direct cloud, traditional host, neocloud, broker, marketplace or managed inference platform. The November ledger provisionally allocates 387 groups across these roles, but many categories are inherited and company/legal-entity identities remain unverified. These are not 387 GPU-fleet owners.

2. Model labels need a denominator.

At least 106 confirmed domains have bounded H100 evidence, among 222 with any named model. No inference about available GPU stock is warranted for the 117 domains without named-model evidence.

3. Routing evidence is a derived diagnostic.

The 103 direct-classified domains with a retained nonzero prefix count are a subset of 126 across the full cohort. The “network_verified” flag is defined by this same count; the indicator is not a second independent test.

4. Geography is partly recovered, not market-complete.

The master registry accepted 149 advertised service-country pairs across 89 company IDs and 41 countries. A separate purposively selected 13-provider check yielded 51 operational-labelled rows in 16 countries, 44 with an explicit GPU model to region or zone. These observations are not a census, do not identify 51 physical facilities, and cannot support a global country ranking.

5. Build the next quarter on a frozen rule set.

For the next quarter, keep company identities, GPU offerings, model-region availability, physical facility evidence and network presence as separate measures. Preserve source dates and quotations so any change can be traced to a comparable population.

Scope and citation

October research freeze: 28 October 2026. November 5 inputs: master reconciliation SQLite, 114-case verification, service-country ledger, and a separate selected 13-provider infrastructure review (67 deployment records, 275 network-only joins, 14 unresolved research questions). The targeted sample is not a market estimate. The Exoscale partner source corroborates service operation; no physical ownership census was performed.

Suggested citation: ViaBandwidth, The GPU Cloud Footprint: Q4 2026, October checkpoint 162 with November 5 registry and targeted infrastructure evidence; research edition (5 November 2026).

viabandwidth.com | research@viabandwidth.com