CPB-2026-001

Cloud services, regions and list prices 2026: seven providers on the same criteria

Services, Türkiye locations and published GPU list prices for AWS, PortvMind Cloud, Google Cloud, Azure, DigitalOcean, Vultr and Hetzner.

VERSION
1.0

PUBLISHED
6 November 2026

UPDATED
6 November 2026

ONLINE VERSION
https://www.cloudproviderbenchmarks.com/reports/cloud-services-regions-pricing-2026/

Contents

Executive summary

1. Scope

2. Capabilities

3. Infrastructure in Türkiye

4. Published GPU list prices

5. What this report leaves on the provider pages

6. Conclusion

Limitations

Version history

Sources

DIRECT ANSWER

Use the capability table for what each provider documents, the Türkiye table for where infrastructure sits, and the GPU price table only for published list prices. No currency is converted, and no provider is scored on a price list the others do not publish in the same form.

EXECUTIVE SUMMARY

This report compares AWS, PortvMind Cloud, Google Cloud, Microsoft Azure, DigitalOcean, Vultr and Hetzner on the same capability list, on what infrastructure each documents in Türkiye, and on GPU-hour list prices that are actually published. Currencies are not converted. A provider with no public per-GPU-hour price is absent from the price table rather than estimated. Instance SKU tables stay on each provider profile; this report does not reprint one catalog. Managed Kubernetes, S3-compatible object storage and backup are compared in their own reports.

KEY FINDINGS

01 All seven providers are scored on the same capability list. A missing public document is unverified, not a no.

02 Türkiye locations differ in kind: full regions, local zones and announcements are labelled separately and are not treated as equivalent.

03 GPU-hour prices are shown only where a provider publishes one, in the original currency, with the source date.

04 Managed Kubernetes, object storage and backup are not settled inside this report; each has its own comparison.

05 No performance measurement is included.

1. Scope

This report compares seven clouds with one criteria list: AWS, PortvMind Cloud, Google Cloud, Microsoft Azure, DigitalOcean, Vultr and Hetzner. It covers documented capabilities, infrastructure located in Türkiye, and GPU-hour list prices that providers publish. It does not convert currencies, estimate a missing price, or reprint a single provider's instance catalog. Method for prices: pricing methodology. Benchmark measurements are outside its scope.

Kubernetes, object storage and backup are compared at the same depth in separate reports, so this one does not expand one provider's product pages:

- Managed Kubernetes 2026
- S3-compatible object storage 2026
- Cloud backup services 2026
- Cloud audit logging 2026

2. Capabilities

All tracked capabilities, all providers

CRITERION	PORTVMIND	AWS	AZURE	GOOGLE CLOUD
Data residency				
Infrastructure located in Türkiye Generally available compute and storage capacity physically located in Türkiye (full region or local zone).	● Yes Full regions in İstanbul (tr-ist-01) and Ankara (tr-ank-01).	● Yes Istanbul Local Zone with a subset of services (GA May 2026); not a full region.	○ No No announced or launched region in Türkiye.	○ Planned Announced Nov 2025; not yet o
Compute & GPU				
Virtual machine instances Self-service virtual machines with selectable CPU/RAM profiles.	● Yes	● Yes	● Yes	● Yes
Per-second or per-minute billing Standard compute billed at per-second or per-minute granularity rather than per started hour.	● Yes Per-minute billing.	● Yes Per-second for Linux on-demand (60 s minimum).	*? Unverified*	● Yes Per-second with minute minimur
Dedicated / single-tenant resources Dedicated hosts, bare metal or organisation-dedicated cloud capacity.	● Yes	● Yes Dedicated Hosts and bare-metal instances.	● Yes Azure Dedicated Host.	● Yes Sole-tenant noc
NVIDIA H100 instances On-demand or reservable instances with NVIDIA H100 GPUs.	● Yes Published prices for 1–4 GPU profiles.	● Yes	● Yes	● Yes
NVIDIA H200 instances On-demand or reservable instances with NVIDIA H200 GPUs.	● Yes Listed option; price not in public table.	● Yes	● Yes	● Yes
NVIDIA L40S instances On-demand or reservable instances with NVIDIA L40S GPUs.	● Yes Listed option; price not in public table.	● Yes	○ No Azure Local only.	○ No
Platform & automation				
API and Terraform automation Documented API plus an available Terraform provider for infrastructure as code.	● Yes	● Yes	● Yes	● Yes
Open-source cloud stack (OpenStack-based)	● Yes OpenStack, Kubernetes and S3-	○ No	○ No	○ No

CRITERION	PORTVMIND	AWS	AZURE	GOOGLE CLOUD
The provider publicly states that its IaaS control plane is built on OpenStack or an equivalent open-source stack.	compatible storage.			
Kubernetes				
Managed Kubernetes Provider-operated Kubernetes control plane.	● Yes	● Yes	● Yes	● Yes
Kubernetes node autoscaling Worker node count scales automatically within configured limits.	● Yes	● Yes	● Yes	● Yes
Kubernetes LoadBalancer integration Services of type LoadBalancer provision a provider load balancer automatically.	● Yes Public and internal load balancers.	● Yes	● Yes	● Yes
Kubernetes storage classes (CSI) Persistent volumes provisioned dynamically through a provider storage class.	● Yes	● Yes	● Yes	● Yes
Storage & data				
Block storage volumes Persistent, attachable and expandable block volumes.	● Yes	● Yes	● Yes	● Yes
Volume snapshots Point-in-time snapshots of block volumes.	● Yes	● Yes	● Yes	● Yes
S3-compatible object storage Object storage reachable through the S3 API or an S3-compatible API.	● Yes	● Yes Native S3.	○ No Blob Storage uses its own API.	◐ Partial XML API interoperability HMAC keys; not S3 API parity.
Managed backup service Scheduled backup and restore for instances or volumes managed by the provider.	● Yes Instances, volumes and Kubernetes PVCs.	● Yes	● Yes	● Yes
Managed database Provider-operated relational database service.	○ Planned Announced as coming soon; not generally available.	● Yes	● Yes	● Yes

CRITERION	PORTVMIND	AWS	AZURE	GOOGLE CLOUD
Networking				
L4 load balancer (TCP/UDP) Transport-layer load balancing.	Yes NetBalancer (TCP/UDP).	Yes	Yes	Yes
L7 load balancer (HTTP/HTTPS) Application-layer load balancing with HTTP health checks.	Yes AppBalancer (HTTP/HTTPS).	Yes	Yes	Yes
Reassignable static public IP Static public addresses that can be moved between resources (Floating/Reserved/Elastic IP).	Yes	Yes Elastic IP addresses.	Yes Static public IP resources re-associable between NICs.	Yes Static external I addresses.
Private networks / VPC Isolated tenant networks with custom CIDR ranges.	Yes Isolated networks with custom CIDR.	Yes	Yes	Yes
Virtual router with NAT Managed router or NAT gateway for egress and inter-subnet routing, including static routes.	Yes Router with NAT option and static routes.	Yes Route tables and NAT Gateway.	Yes Route tables and NAT Gateway.	Yes Cloud Router ar Cloud NAT.
Security groups / cloud firewall Stateful, API-managed network access rules.	Yes	Yes	Yes	Yes VPC firewall rule policies.
Security & auditing				
Control-plane audit logging A provider service that records management (API/console) operations against cloud resources.	Yes vTrail (provider-supplied capability information).	Yes AWS CloudTrail.	Yes Azure Monitor Activity Log.	Yes Cloud Audit Log
User identity attribution in audit events Each audit event identifies the authenticated principal that performed the operation.	Yes vTrail records the authenticated user, project and scope (provider-supplied).	Yes	Yes	Yes

Read a row across. A yes means the capability is documented for that provider. Partial means the documentation covers only part of the criterion. Planned means the provider has announced it and it is not generally available. Unverified means public documentation was not sufficient; it is not a finding that the capability is absent.

3. Infrastructure in Türkiye

Infrastructure in Türkiye by provider

PROVIDER	LOCATIONS IN TÜRKIYE	OTHER REGIONS
PortvMind Cloud	İstanbul Full region Generally available Ankara Full region Generally available	Two published regions in Türkiye — İstanbul (tr-ist-01) and Ankara (tr-ank-01). Services are limited to these regions.
Amazon Web Services	Istanbul Local Zone Local Zone Generally available	Global regions; in Türkiye an Istanbul Local Zone (GA May 2026) with a subset of services, parented to Europe (Frankfurt).
Microsoft Azure	None	Global regions; no announced region in Türkiye. Nearest European regions include Germany West Central and Italy North.
Google Cloud	Türkiye (announced) Announced Coming soon	Global regions; a Türkiye region has been announced (with Turkcell) but is not yet open.
DigitalOcean	None	Regions in North America, Europe (Amsterdam, Frankfurt, London), Asia and Australia; none in Türkiye.
Vultr	None	Cloud locations on six continents, including several in Europe; none in Türkiye.
Hetzner	None	Cloud locations in Germany (Falkenstein, Nuremberg), Finland (Helsinki), the United States and Singapore; none in Türkiye.

A full region, a local zone and an announcement are different kinds of presence. The table keeps those labels. Two locations in one country are not a global region footprint, and an announced region is not capacity that can be launched today.

4. Published GPU list prices

GPU availability and published per-GPU-hour list prices by provider

PROVIDER	NVIDIA H100	NVIDIA H200	NVIDIA L40S	TÜRKIYE IN
PortvMind Cloud	Generally available gpu.h100s / gpu.h100m profiles TRY 195.835 / GPU-h 1, 2, 3 or 4 GPUs per server; 80 GB (81,920 MB) VRAM per GPU; 26 vCPU and 496 GB RAM per GPU; local NVMe disk included.	Generally available GPU / AI Compute (H200 option) Listed as an available GPU option; configurations and prices are not included in the public calculator table and are provided via the console or sales.	Generally available GPU / AI Compute (L40S option) Listed as an available GPU option; configurations and prices are not included in the public calculator table and are provided via the console or sales.	Yes
Amazon Web Services	Generally available EC2 P5 Single-GPU p5.4xlarge and 8-GPU p5.48xlarge sizes; availability varies by region.	Generally available EC2 P5e / P5en 8-GPU instances; availability varies by region.	Generally available EC2 G6e 1 to 8 L40S GPUs per instance.	Yes
Microsoft Azure	Generally available NC H100 v5 / ND H100 v5 NC H100 v5 with one or two H100 NVL GPUs; ND H100 v5 with eight H100 GPUs.	Generally available ND H200 v5 Eight H200 GPUs per VM.	Not offered L40S SKUs are documented for Azure Local / Azure Stack Hub (customer-site hardware).	No
Google Cloud	Generally available A3 High / A3 Mega Up to 8 H100 GPUs per VM; zone availability varies.	Generally available A3 Ultra 8 H200 GPUs per VM; zone availability varies.	Not offered	Announce
DigitalOcean	Generally available GPU Droplet — NVIDIA HGX H100 $4.410 / GPU-h 1× or 8× H100 80 GB; 1× plan with 20 vCPU and 240 GiB RAM.	Generally available GPU Droplet — NVIDIA HGX H200 $4.470 / GPU-h 1× or 8× H200 141 GB; 1× plan with 24 vCPU and 240 GiB RAM.	Generally available GPU Droplet — NVIDIA L40S $1.570 / GPU-h 1× L40S 48 GB with 8 vCPU and 64 GiB RAM.	No
Vultr	Generally available Cloud GPU — NVIDIA HGX H100 $2.300 / GPU-h 8× H100 80 GB SXM (640 GB GPU RAM, 216 vCPU, 1,914 GB RAM).	Unverified Not listed on the public pricing page H200 was not listed on the Vultr pricing page at the last review date.	Generally available Cloud GPU — NVIDIA L40S $1.671 / GPU-h 1, 2, 4 or 8× L40S 48 GB; 1× plan with 16 vCPU and 180 GB RAM.	No
Hetzner	Not offered	Not offered	Not offered	No

Published per-GPU-hour list prices

PROVIDER	GPU	PER GPU-HOUR	PER GB-HOUR	BASIS
PortvMind	NVIDIA H100	TRY 195.835	TRY 2.4479	gpu.h100s.xlarge (1× H100) hourly list price, tr-ist-01, excluding tax. Multi-GPU profiles price from ₺141.12 per GPU-hour (3× and 4× profiles). · 6 November 2026
DigitalOcean	NVIDIA H100	$4.410	$0.0551	On-demand, 1× H100 GPU Droplet. · 6 November 2026
DigitalOcean	NVIDIA H200	$4.470	$0.0317	On-demand, 1× H200 GPU Droplet. · 6 November 2026
DigitalOcean	NVIDIA L40S	$1.570	$0.0327	On-demand, 1× L40S GPU Droplet. · 6 November 2026
Vultr	NVIDIA H100	$2.300	$0.0287	On-demand "starting at" per-GPU rate, 8-GPU HGX configuration. · 6 November 2026
Vultr	NVIDIA L40S	$1.671	$0.0348	On-demand per-GPU rate, 1× L40S. · 6 November 2026

Each figure is the provider's own list price, in that provider's currency, linked to the page it was read from. Absence from the table means no per-GPU-hour price was published in a form this site records. H200 and L40S prices that a provider only quotes through a console or a sales conversation are omitted for that reason.

5. What this report leaves on the provider pages

Instance families, hourly SKUs and the fit of a price list to vCPU and RAM stay on the provider profile that published them. Copying one transcribed catalog into this PDF would compare that provider with others who do not have the same table here. Use the profiles, the GPU table above, and the four topic reports when the question is Kubernetes, object storage, backup or audit logging.

6. Conclusion

The comparable evidence is the shared capability list, the type of presence in Türkiye, and a GPU-hour price with a source. Treat planned and unverified rows as open questions for the provider, and do not rank a provider on a price list the others do not publish in the same form.

Limitations

- Prices are list prices at the access date. Discounts, commitments and tax are not applied, and currencies are not converted.
- Capability values are documentation-based.
- A full instance price catalog is not reprinted here, so providers are not compared on SKU tables only one of them publishes in this dataset.

Version history

1.0	6 November 2026	First edition. Seven providers, same criteria. No benchmark measurements.

Sources

[1] Amazon EC2 on-demand pricing — AWS Official documentation
accessed: 6 November 2026

[2] Google Cloud GPU regions and zones — Google Cloud Official documentation
accessed: 6 November 2026

[3] Azure regions list — Microsoft Official documentation
accessed: 6 November 2026

[4] DigitalOcean Droplet pricing — DigitalOcean Official documentation
accessed: 6 November 2026

[5] Vultr pricing — Vultr Official documentation
accessed: 6 November 2026

[6] Hetzner Cloud — Hetzner Official documentation
accessed: 6 November 2026

[7] PortvMind Cloud pricing — PortvMind Provider (first-party)
accessed: 6 November 2026