

Below is a vendor-by-vendor deep dive for your specific use case: high-uptime, low-TTFB, database-heavy e-commerce, with sane and predictable cost. Focused on the concrete trade-offs rather than marketing.

Hetzner

Performance & Hardware

Hetzner's Cloud line (CX/CAX shared vCPU, CCX dedicated vCPU) is widely regarded as one of the best price/performance options in Europe and now has US and Singapore regions. Independent tests show strong disk and network throughput: one 2026 review measured Hetzner's VPS stream copy bandwidth at roughly 105 GiB/s, ahead of OVH (~47 GiB/s), Contabo (~36 GiB/s), and DigitalOcean (~19 GiB/s), indicating very fast memory and I/O paths for PHP and database workloads. Community and VPSBenchmarks endurance runs highlight stable CPU performance over 24-hour stress tests, and the platform explicitly tracks CPU steal as a signal of over-allocation; Hetzner's CCX and newer ARM CAX instances generally show low variation, implying little throttling under sustained load. A Reddit user comparing providers notes "never seen steal time" issues on Hetzner and recommends it explicitly over Contabo for better CPU fairness. [picknpare](#)

Disk is SSD or NVMe (NVMe on many newer CAX tiers), typically in RAID-10 on local nodes, which aligns with the solid disk performance seen in third-party benchmarks and WordPress hosting tests. Those WordPress tests show Hetzner delivering consistently fast page loads at much lower price points than US hyperscalers, which is directly relevant to TTFB and catalog browsing performance. For MySQL/PostgreSQL, raw I/O plus predictable CPU makes Hetzner suitable for self-hosted DBs; there is no first-party managed database, so you must run and tune your own stack or layer something like CloudPanel or PaaS on top. [bitdoze](#)

"Hetzner has best price to performance ratio and I've never seen steal time there."
[reddit](#)

Reliability & Uptime

Hetzner offers a 99.99 % SLA on its cloud platform, similar on paper to DO/Vultr. Public reviews and test reports generally rate availability high, though like most non-hyperscalers they have had regional incidents (network and power in individual DCs), usually resolved within a few hours. Independent reviewers describe the performance as "mixed" (CPU scores on the lower side vs premium US providers) but disk and I/O as "solid," with overall stability considered good for the price. High availability (multi-AZ or cross-region failover) is not provided as a built-in pattern; you must architect redundancy yourself with multiple servers, load balancers and/or an external CDN/WAF. [hoststest.co](#)

Incident communication is via a fairly bare-bones status page and forum posts; it is functional but not at the level of AWS/Azure's post-mortems. In practice, this is acceptable for many SMEs but you should plan your own multi-region failover for mission-critical stores.

UX & Management

The Hetzner Cloud Console is widely praised as clean and developer-centric: instance creation, firewall rules, snapshots, volumes and networking are all straightforward. For a competent developer new to Hetzner, bringing up a production-grade LEMP/WordPress or custom stack is typically a same-day effort; the main learning curve is around their networking model (floating IPs, private networks) rather than the UI itself. Common tasks (reboot, snapshot, resize, firewall rule updates, SSH key management) are one- or two-click operations, so day-to-day ops are efficient. [picknpare](#)

The API and Terraform provider are mature and actively used by the community for IaC, enabling reproducible deployments and blue/green rollouts. If you're comfortable with Terraform/Ansible, you can manage most of your e-commerce infrastructure without ever logging into the console. [vpsbenchmarks](#)

Network, Security & Backups

Hetzner includes large traffic quotas per VM: for example, comparison guides show entry cloud plans including 20–30 TB/month bundled traffic where DO/Vultr often include 0.5–10 TB. Independent egress cost analyses find that even at high volumes (10–50 TB/month), Hetzner's incremental cost is negligible compared to hyperscalers: one 2025 study calculated 1 TB egress at roughly €0 on Hetzner (within included quota) vs about 82–110 USD on AWS/Azure/GCP, and 50 TB at about €36 on Hetzner vs roughly 4 000–4 700 USD on the big three. For a 10 TB flash sale spike, your incremental Hetzner bill is often effectively zero if you're under the per-VM allowance; the same event on AWS at \$0.09/GB would be ~900 USD in data transfer alone. [handyintelligence](#)

Default DDoS protection is basic volumetric filtering at the network edge; there is no equivalent to AWS Shield Advanced or Cloud Armor built in. For serious L7 attack resilience you are expected to front Hetzner with Cloudflare or another CDN/WAF. Backups and snapshots are optional add-ons typically billed at ~20 % of the instance cost, similar to DigitalOcean's approach. Snapshot restores are quick (minutes to boot a replacement VM), but there is no multi-region backup orchestration; you must design cross-region backup strategies yourself. [vpsbenchmarks](#)

TCO & Pricing Transparency

Hetzner's pricing is unusually transparent: a typical CX or CAX instance with 2 vCPU, 4 GB RAM, and 40 GB SSD is around €3.49–5.29/month, with 20 TB traffic included, undercutting


DO/Vultr pricing for similar resource levels while providing far more bandwidth. Sample comparisons show a “VM Small” at Hetzner (2 vCPU, 4 GB) at about \$5.29/month vs an AWS t3.micro (2 vCPU, 1 GiB) at about \$7.49/month and far higher egress costs. Object and block storage are cheaper than AWS; one comparison prices 1 TB object storage at about \$5.86 on Hetzner vs \$24.63 on AWS. [getdeploying](#)

No setup fees apply, but IPv4 addresses are billed separately (like most modern providers) and backup snapshots are charged as a percentage. Billing is highly predictable because egress is mostly bundled; analyses repeatedly highlight that for 1–50 TB/month outbound, Hetzner’s total transfer cost is effectively a rounding error while AWS/Azure/GCP can climb into four digits. I would rate billing predictability around 9/10. [wz-it](#)

Support & Emergency Response

Hetzner offers ticket-based support; response times are generally good during European business hours but not positioned as 24/7 “white glove.” There is no paid premium support tier comparable to AWS Enterprise Support; you remain in an L1/L2 ticket queue. For emergency database corruption or DDoS, they will ensure the node and network work, but are unlikely to help debug your MySQL configuration. [bitdoze](#)

The flip side is excellent documentation and a large ecosystem of community tutorials, including many WordPress-specific guides. For experienced operators, this self-service model is acceptable; for a non-technical retailer, it is a risk. [bitdoze](#)

 **Warning:** You should assume *you* are the SRE on Hetzner. If you need guaranteed 15-minute human escalation in the middle of a Black Friday incident, this is not built-in; you’d need an external ops partner or to orchestrate failover to another provider.

Verdict

Hetzner is arguably the sweet spot for cost-sensitive EU-centric and increasingly global e-commerce: strong I/O and CPU stability, huge included bandwidth, and low, predictable bills. It is ideal for technical teams comfortable managing their own Linux, databases, and backups, and for architectures where you can add Cloudflare or a similar service for advanced DDoS/WAF. It is **not** ideal if you want managed databases, multi-AZ primitives, or enterprise-grade support SLAs; in that case a hyperscaler is safer. For a developer-led shop that values TTFB, DB performance, and total cost over vendor logos, Hetzner is one of the best fits in this entire list. [experte](#)

Vultr

Performance & Hardware

Vultr positions itself as a performance-focused cloud with global reach and has invested heavily in modern AMD EPYC platforms and NVMe storage. Independent “VPS for developers” testing found Vultr’s High Frequency/High Performance instances delivered the fastest Docker builds (1 m 58 s) and highest disk I/O (~680 MB/s sequential), outperforming DigitalOcean and others in CPU-bound tasks. Another cross-provider benchmark ranked “Vultr – Regular” as delivering high overall performance with an excellent performance-per-dollar score, only beaten by an ultra-niche provider. [deploy](#)

VPSBenchmarks endurance data shows multiple Vultr plans (Regular, High Frequency, High Performance, CPU-Optimized) tested with 24-hour sustained CPU workloads; the platform explicitly reports CPU steal percentages and notes that >1 % indicates potential over-allocation. While we don’t have raw steal values in the excerpt, the presence of many trials and an overall “Consistency Score” of 61 suggests reasonably stable CPU for production workloads. Database workloads benefit from NVMe and strong single-thread performance, making Vultr a solid choice for self-hosted MySQL/PostgreSQL on their High Performance tiers; managed databases are available but less mature than hyperscaler offerings. [vpsbenchmarks](#)

Reliability & Uptime

Vultr advertises a 99.99 % SLA on many cloud plans, similar to DO/Hetzner. While it doesn’t have the public outage scrutiny of AWS/Azure, community reports generally portray uptime as good, with occasional localized network or storage issues. VPSBenchmarks’ focus on stability and coefficient of variation for Vultr plans suggests that under sustained load, performance is fairly consistent, reducing the risk of surprise TTFB spikes during traffic peaks. As with most non-hyperscalers, multi-AZ HA is not a turn-key feature; you must run nodes in multiple regions if you want cross-DC resilience. [vpsbenchmarks](#)

Incident communication is a standard status page plus social updates; transparency is adequate, but you won’t see AWS-style 10-page post-incident reports.

UX & Management

Vultr’s control panel is simple and fast, comparable to DigitalOcean’s: provisioning, snapshots, firewalls, and network config are all intuitive. Onboarding for a competent developer new to Vultr is typically under a day to get a production environment live, especially with their one-click applications (WordPress, LAMP stacks, etc.). Reboot, resize, snapshots, and firewall updates are 1–2 clicks; SSH keys are easy to manage, so operational friction day-to-day is low. [deploy](#)

Vultr exposes a well-documented API and Terraform provider, widely used by devops teams to spin up infrastructure in many of their 25+ regions. This is valuable if you want to keep the same deployment pattern across EU, US, and APAC. [vpsbenchmarks](#)

Network, Security & Backups

Vultr has excellent geographic coverage (25+ locations) and solid network performance; independent developer-focused benchmarks ranked Vultr number one for global network performance among a mid-tier set. Bandwidth allowances vary by plan (often 2–10 TB on common tiers), with overage charged per GB, but prices are more modest than hyperscalers. There is basic DDoS mitigation available as an add-on; L3/L4 filtering is included in some regions, but for serious L7 defense you should again use a third-party CDN/WAF. [akashrajpurohit](#)

Backups and snapshots are optional paid features; the exact pricing tends to be a percentage of the instance cost or per-GiB snapshot storage. Restore times are quick because storage is local NVMe. Vultr Cloud Firewalls provide simple stateful filtering by IP/port, and private networking/VPC-like isolation is available in most regions, sufficient for separating web and DB tiers. [picknpare](#)

TCO & Pricing Transparency

Vultr's Regular compute starts around \$5–6/month for 1 vCPU, 1 GB RAM, with High Frequency compute around \$6/month for similar RAM but NVMe and faster CPUs. Benchmarks show that their higher-end plans deliver strong performance but at a lower performance-per-dollar score than cheaper Regular instances; a 2025 test found Vultr High Frequency had excellent raw scores but only “Strong but pricey” value relative to its monthly cost. [dev](#)

Traffic is not unmetered; if your store pushes 5 TB/month, you must account for overage costs, but you're still in a completely different league than AWS/Azure/GCP: 5 TB at ~\$0.01/GB is around \$50, versus ~\$450 on a hyperscaler at \$0.09/GB. There are no setup fees and IPv4 is usually included. Billing predictability is good (7–8/10), but less “fire and forget” than Hetzner/OVH due to metered transfer. [wz-it](#)

Support & Emergency Response

Support is ticket-based 24/7; there is no widely advertised live phone support for standard plans. Community sentiment suggests response times are reasonable for non-urgent issues, but serious incidents may still see hours-scale resolution. You do not get access to deep network engineers at will; like other budget clouds, L1/L2 will ensure the platform is working but not debug your application. Documentation is solid but less extensive than DigitalOcean's, so some troubleshooting falls back to generic Linux knowledge.

⚠ Warning: For high-profile e-commerce where a 2–3 hour network issue could cost six figures, Vultr’s support model is a risk. Consider pairing it with multi-region redundancy or using it only for stateless front-ends, keeping critical databases on a provider with stronger incident response.

Verdict

Vultr is a strong choice when you need raw performance and global coverage at mid-range prices, especially for CPU-bound workloads and NVMe-heavy databases. It fits well for developer-led shops that can self-manage MySQL/PostgreSQL and want more locations than Hetzner offers. It is less attractive if your traffic profile is bandwidth-heavy (egress is metered) or if you require enterprise-grade SLAs and support. For most small–mid e-commerce stores with global customers and a technical team, Vultr High Performance or High Frequency instances are a very credible option. [akashrajpurohit](#)

DigitalOcean

Performance & Hardware

DigitalOcean has long targeted developers and startups, with Droplets (VPS) that balance simplicity and performance. Independent 2025 reviews measured SSD/NVMe storage at around 500–1 200 MB/s sequential reads and ~800 MB/s writes, and CPU benchmarks scoring 90+ out of 100 for single- and multi-thread performance on modern Intel/AMD platforms. A WordPress test on a \$12/month droplet reportedly loaded pages in about 0.8 s with 1 000 concurrent users, maintaining ~120 ms response times — a strong indicator that DigitalOcean can handle real-world e-commerce loads with good TTFB when tuned. [vps-commander](#)

VPSBenchmarks runs endurance tests on both Basic and CPU-Optimized droplets and tracks CPU steal and performance variation; DigitalOcean’s plans generally show low variation, indicating moderate over-allocation but not the extreme steal issues seen on some budget hosts. However, another 2025 comparative benchmark found DigitalOcean’s “CPU-Optimized” instances to be poor value: a third-party €8.99 NVMe VPS beat DO’s \$42 CPU-Optimized tier by over 150 % in single-core performance, suggesting that DO’s higher tiers are relatively expensive for the raw compute they deliver. [vpsbenchmarks](#)

DigitalOcean offers managed MySQL/PostgreSQL (and Redis) clusters, which offload patching, backups, and replication, and typically provide better failover behavior than self-hosted DBs on the same Droplet tier, though at a premium price. [vpshostingpicks](#)

Reliability & Uptime

DO advertises a 99.99 % uptime SLA and also sells its own Uptime monitoring service. Third-party status aggregators show that DigitalOcean has had thousands of small incidents over the last decade, but recent 30-day windows often show near-perfect uptime with only a handful of short disruptions. Reviews and benchmarks often cite an “average uptime” of about 99.98 %, which aligns with the SLA and is generally acceptable for e-commerce.

[docs.digitalocean](#)

However, like any provider, DigitalOcean has sporadic regional outages (e.g., brief control-plane or API issues) that can impact provisioning or management, though data-plane impact (running VMs) is usually limited. SLA credits require you to open a support ticket, and credits are limited to a percentage of the monthly bill, not lost revenue. [statusgator](#)

UX & Management

DigitalOcean’s control panel is one of the cleanest in the industry and is routinely recommended for developer friendliness. For a new developer, standing up a production LEMP/WordPress or Laravel stack is often a 1–2 hour job thanks to one-click images and straightforward networking. Reboots, snapshots, resize, firewall rules and SSH key management are all 1-2 clicks in an intuitive UI. On a 1–10 UX scale for lean e-commerce teams, DigitalOcean is around 9/10.

[vps-commander](#)

API and Terraform support are robust, with many community modules and guides on using DO with IaC. Combined with their extensive documentation (over 3 000 tutorials by some counts), this makes DO extremely efficient to operate day-to-day. [vpshostingpicks](#)

Network, Security & Backups

DigitalOcean’s regions (NYC, SF, Amsterdam, Frankfurt, London, Singapore, Toronto, Bangalore, etc.) cover most major e-commerce markets, with premium networking providing 1–2 Gb/s to instances. Latency is competitive within regions, but unlike AWS/GCP, there is no self-operated global backbone across continents — you rely on transit providers.

[vpsbenchmarks](#)

DDoS protection is basic and largely focused on volumetric filters; for serious L7 web attacks, DO expects you to front with a service like Cloudflare. They offer Cloud Firewalls and VPC-style private networking which are easy to configure, sufficient for separating web and DB layers.

[vpsbenchmarks](#)

Backups are offered as an automated add-on typically priced at 20 % of the droplet cost, and snapshots are charged per-GiB stored. Restore times are typically minutes. This “20 % backup

surcharge” is easy to understand but can materially change TCO on larger droplets.

[vpshostingpicks](#)

Traffic is metered, with common plans including 1–4 TB and overage billed at about \$0.01/GB. A flash sale generating 10 TB outbound in a month would imply maybe 6–9 TB above the included quota depending on your base plan, so roughly \$60–90 in egress, which is reasonable compared to AWS but non-trivial relative to a \$12–24 droplet. [vps-commander](#)

TCO & Pricing Transparency

DigitalOcean’s base pricing is straightforward: \$6/month for 1 vCPU, 1 GB RAM, 25 GB SSD and 1 TB transfer; \$12/month for 1 vCPU, 2 GB RAM, 50 GB SSD; \$24/month for 2 vCPU, 4 GB RAM, 80 GB SSD, etc. Managed databases and Spaces (object storage) are extra (e.g., Spaces at \$5/month for 250 GB), and backups at 20 % of droplet cost add up on larger instances. For a typical mid-size store using 5 TB/month, you might pay a base of \$24–48 for compute plus \$50 of egress, plus say \$15–30 for managed DB and backups, so roughly \$90–130 total — still far below hyperscalers but no longer “\$10 hosting.” [deploy](#)

Billing predictability is good (8/10); everything is listed on the pricing page, and there are no surprise “regional data transfer” line items. There is no long-term contract or reserved instance complexity.

Support & Emergency Response

DigitalOcean offers ticket-based support, with higher tiers (Premium) for organizations needing faster response times. There is no built-in 24/7 phone hotline at the basic level; you rely on asynchronous tickets. Community sentiment is generally positive about response quality but notes variability in response time during incidents. Support is more application-aware than commodity VPS hosts but not at AWS Enterprise depth. [vpshostingpicks](#)

The real strength is the documentation and community Q&A; many common scenarios (database backups, snapshot restores, failover patterns) are covered with detailed how-tos. For a small dev team, DO’s ecosystem significantly reduces the time to diagnose issues.

[vps-commander](#)

⚠ Warning: DO’s managed DBs and bandwidth pricing can stealthily raise TCO as you scale — especially if you assume your \$6 droplet is representative of long-term production cost. For a store doing multi-TB/month, budget carefully.

Verdict

DigitalOcean is excellent when developer UX and documentation matter almost as much as raw performance: it’s a great home for small–mid stores that want a “just works” experience and don’t mind paying a moderate premium over Hetzner for better docs and a more polished

ecosystem. It becomes less compelling as you scale into high traffic or very CPU-intensive workloads where its CPU-Optimized tiers are poor performance-per-dollar relative to specialized competitors. For a technically competent team that wants managed DBs and a smooth panel, DO is a strong fit, but for bandwidth-heavy or cost-minimized operations, Hetzner or OVH are usually better. [danubedata](#)

Amazon Web Services (AWS)

Performance & Hardware

AWS EC2 offers a vast catalog of instance families (t-series burstable, m-series general-purpose, c-series compute, r-series memory) across many generations. Benchmarks often show competitive or excellent raw CPU and network throughput, particularly on newer c7a/c7g instances, but at significantly higher cost than mid-tier VPS hosts. In a 2025 performance-per-dollar analysis, an AWS c7a.large at about \$74/month scored worse than multiple VPS providers, with a PPD score under 1 compared to >2.5 for cheaper VPS, illustrating the poor value of EC2 for generic workloads. [vpsbenchmarks](#)

However, AWS offers dedicated vCPU and bare metal options where CPU steal is essentially zero, and their Nitro hypervisor is built for strong isolation — you simply pay for that isolation. For e-commerce databases, RDS (managed MySQL/PostgreSQL) and Aurora provide high IOPS with storage that can scale independently of compute and multi-AZ synchronous replication, which tends to yield better availability and failover characteristics than a self-hosted DB on a single VPS. Managed DBs often outperform self-hosted instances in reliability and operations, but not necessarily in raw single-thread latency per dollar compared to Hetzner/Vultr. [opsmatters](#)

Reliability & Uptime

AWS markets 99.99 %+ instance and RDS region-level SLAs when you use multi-AZ configurations, but real-world history shows that AWS — especially us-east-1 — still suffers major outages. A 15-hour outage in October 2025 in us-east-1 disrupted DynamoDB, EC2, Lambda and more, causing widespread downtime for major websites and apps. Another extended event on October 28, 2025 specifically hit ECS/EC2 in one availability zone, leading to 14 hours of degraded launches and throttling. These incidents show that even with multi-AZ design best practices, systemic control-plane failures can degrade performance and availability for critical workloads. [dargslanpublishing](#)

AWS's SLA credits require careful monitoring and ticket submission; they are capped as a percentage of monthly spend and do not compensate lost revenue. On the plus side, AWS

provides detailed post-mortems and robust status history, and you can architect active-active multi-region deployments with their primitives — at a steep cost. [opsmatters](#)

UX & Management

The AWS Console is powerful but complex. On a 1–10 UX scale for a small e-commerce team, it's around 4–5/10 for newcomers: IAM, VPC, security groups, subnets, RDS parameter groups, etc. have steep learning curves. A developer with no AWS experience can easily spend several days to weeks learning enough to safely deploy a production-grade environment with least-privilege IAM and proper networking.

Common tasks like resizing, snapshots, and security group changes are possible but can involve multiple screens and conceptual overhead. Core tasks are scriptable via the AWS CLI and Terraform, which are first-class citizens, but that assumes devops sophistication. For lean teams, AWS's operational overhead can be significant. [docs.cloud.google](#)

Network, Security & Backups

AWS runs its own global backbone and offers many regions/AZs, giving excellent latency to most markets. Egress pricing, however, is punishing: typical public internet data transfer out is around \$0.09/GB for the first 10 TB, with slightly lower rates thereafter. One comparison calculated 1 TB egress at ~\$82–92/month on AWS and 10 TB at about \$922; 50 TB cost over \$4 000/month. For a flash sale pushing 10 TB in 2 days, you can easily add close to \$1 000 in transfer fees alone. [handyintelligence](#)

AWS Shield Standard automatically provides L3/L4 DDoS protection for all customers at no extra charge, defending against common volumetric and protocol attacks. For advanced L7 mitigation and 24/7 DDoS response team, Shield Advanced runs at roughly \$3 000/month plus AWS WAF per-request charges. CloudFront, AWS WAF, and ALB/NLB offer robust application protection and TLS termination. EBS snapshots and RDS automated backups are easy to enable and designed for fast recovery; cross-region backups are also supported. [inventivehq](#)

⚠ Warning: Egress and advanced security can combine into four-figure charges very quickly. For a medium store with 5 TB/month outbound, data transfer alone is ~450 USD; adding CloudFront/WAF, Shield Advanced, and RDS can put you into multiple thousands per month even before EC2 costs. [inventivehq](#)

TCO & Pricing Transparency

Multiple comparisons show AWS instances costing 5–10× as much as equivalent-spec VPS from Hetzner or Vultr when factoring in egress. For example, a small VM with 2 vCPU, 1 GiB RAM (t3.micro) is roughly \$7.49/month vs €5.29/month for a Hetzner CX22 (2 vCPU, 4 GB), and 1 TB egress is ~\$90 on AWS vs roughly €0–1 on Hetzner. [danubedata](#)

Billing predictability is low to moderate (4–6/10): the base instance cost is clear, but cross-AZ and cross-region data transfer, CloudFront, WAF, NAT Gateway, and other per-request/per-GB charges create complex bills. Free tier and promotional credits (common \$100–\$200 startup offers) can obscure true long-term cost until they expire.

Support & Emergency Response

AWS's standard (free) support is essentially community forums and documentation; for meaningful help, you need at least Business Support, which can be several hundred dollars/month and scales with usage. Enterprise Support is even more expensive but offers 24/7 phone and chat, 15-minute response targets for critical issues, and direct access to engineers. Many reports highlight excellent engineering depth once you are on a paid plan.

[docs.aws.amazon](https://docs.aws.amazon.com)

Resolution speed during major outages is constrained by the scale of the platform; even AWS's best teams needed 10–15 hours to fully resolve the October 2025 incidents. Documentation and runbooks are extensive, but they assume significant cloud fluency. [penchev](#)

⚠ Warning: For small–mid e-commerce shops, the cost of adequate AWS support (Business/Enterprise) is itself a major TCO line item; running on AWS with only basic support and no in-house ops is risky.

Verdict

AWS is overkill — and often cost-prohibitive — for most standalone e-commerce sites whose main needs are fast TTFB, solid DB performance, and predictable bandwidth charges. It shines for enterprises that need compliance, global multi-region HA, advanced managed services, and deep DDoS/WAF capabilities, and who can justify four-figure monthly spend and paid support. For a lean operator focused on cost-efficient hosting, AWS is only appropriate if (a) you are already heavily invested in its ecosystem, or (b) you have very spiky/compute-heavy workloads where its elasticity outweighs transfer costs. [getdeploying](#)

Google Cloud Platform (GCP)

Performance & Hardware

GCP's Compute Engine instances are highly performant, with many benchmarks showing strong CPU and disk performance, especially on newer N2/C2/C3 instances. Sustained Use Discounts (SUDs) can reduce VM compute cost by up to 30 % when instances run the full month, making continuous workloads slightly cheaper than raw on-demand rates. Like AWS,

GCP offers both shared and dedicated CPU instances with near-zero CPU steal on the latter, and persistent disks that can deliver high IOPS suitable for MySQL/PostgreSQL. [tekpon](#)

Cloud SQL and AlloyDB (managed databases) provide managed replication, backups, and failover for MySQL/PostgreSQL, often with better reliability than self-hosted DBs but at a premium price. For TTFB-sensitive workloads, GCP's global network and HTTP(S) Load Balancer with Cloud CDN can provide low latency, but again cost is the main differentiator vs VPS hosts.

Reliability & Uptime

GCP maintains high SLAs, but as with AWS, outages still occur. Its global HTTP(S) load balancer and multi-region architecture are resilient, but Cloud SQL regions and specific zones have had incidents. Although we don't have a single marquee outage cited here like AWS's October 2025 event, third-party discussions often point to occasional regional issues. If you design for multi-zone and, ideally, multi-region, you can get near-enterprise-grade availability.

UX & Management

The GCP Console is somewhat more approachable than AWS for some users but still complex. IAM, VPC, firewall rules, and project/org structure require learning. A new developer might take several days to comfortably deploy a production application with proper IAM and networking separation. Terraform and gcloud CLI are powerful but require devops skills.

Core tasks like VM resize, snapshots, firewall updates are accessible but often embedded in a multi-step flow. For small teams, the overhead is significant vs DO/Hetzner/Vultr.

Network, Security & Backups

GCP's premium network tier uses Google's private backbone for egress, providing excellent latency and throughput across continents. However, egress is expensive: typical public internet egress is on the order of \$0.11/GB in some regions, which implies about \$110 for 1 TB and ~4 700 USD for 50 TB in a comparative analysis. A 10 TB flash sale spike can cost around \$1 100 just in storage/compute egress fees. [cloud.google](#)

Cloud Armor Standard provides L3/L4 DDoS defense, WAF rules, and per-policy/per-rule/per-request pricing (e.g., about \$5/policy/month, \$1/rule/month, and ~\$0.75/million requests). Cloud Armor Enterprise and Managed Protection Plus can reach ~3 000–5 000 USD/month, offering DDoS response team and bill protection similar to AWS Shield Advanced. Backups and snapshots for Compute Engine and Cloud SQL are robust and easily automated. [whizlabs](#)

TCO & Pricing Transparency

Compute list prices are mid-range, and SUDs + committed use discounts can lower effective VM cost by 20–30 %. However, egress and premium security services dominate TCO for traffic-heavy websites. For a 5 TB/month store, you can easily spend ~\$550 on egress plus Cloud Armor usage, on top of VM and Cloud SQL costs. Billing predictability is moderate (6/10); pricing pages are detailed but complex, and upcoming changes like CDN and peering rate hikes (e.g., doubling some peering egress rates in May 2026) can unexpectedly increase bills.

[impossiblecloud](#)

Free-tier credits (often \$300 for new accounts) can conceal long-term costs until the trial ends; afterwards, the true bills surface.

Support & Emergency Response

GCP's free support is limited; for production workloads you typically need a paid support plan to get timely responses. Higher support tiers provide 24/7 human interaction and escalation. Engineering depth is strong, but you must pay for it and still design your own HA/failover strategies.

⚠ Warning: For e-commerce where bandwidth and WAF/DDoS protection matter, GCP's combination of egress, Cloud Armor, and managed DB pricing often leads to 3–10× higher monthly cost than an equivalent setup on Hetzner or OVH, with only marginal gains in performance.

Verdict

GCP is a solid choice if you are already in its ecosystem (e.g., using BigQuery or GCP-based analytics) or if you need deep integration with Google's AI and data services. For a standalone e-commerce site focused on TTFB and DB performance per euro, it is usually not cost-effective compared to good VPS providers. It suits large organizations with complex workloads, compliance requirements, and in-house cloud engineers; for lean teams, the complexity and bandwidth costs are hard to justify. docs.cloud.google

Alibaba Cloud

Performance & Hardware

Alibaba Cloud's ECS instances are widely used in Asia and offer modern CPU and storage options; VPSBenchmarks has tested at least one mid-range plan (ecs.g8a.2xlarge, 8 vCPU/32 GB). For local users in China and nearby markets, latency can be excellent, and

performance is competitive with other hyperscalers. Managed databases (RDS for MySQL/PostgreSQL) and high-performance block storage are available; however, a 2026 report noted price increases of 5–34 % across AI computing, storage, and database services, especially impacting GPU and high-end products. [vpsbenchmarks](#)

Reliability & Uptime

Alibaba Cloud has maintained reasonable uptime historically but has less transparent global status reporting in English compared to AWS/Azure. For Chinese e-commerce, it is often the default choice, but for global merchants, visibility and documentation can be a barrier. Their SLAs are roughly competitive with other hyperscalers but still require manual claims for credits.

UX & Management

The console and documentation are improving but still feel more complex and less polished for Western developers than AWS/GCP. Pricing models include pay-as-you-go, subscription, and savings plans; understanding which combination yields best value can be non-trivial. For a new developer, onboarding onto Alibaba may take longer than DO/Vultr and similar to AWS/GCP, especially if navigating language and regulatory peculiarities. [wise](#)

Network, Security & Backups

Alibaba Cloud's network is strong in Asia and includes CDN offerings. CDN egress pricing in international regions is on the order of \$0.07/GB for NA/EU, with various regional tiers. For compute/object storage egress, one review cites example pricing where 1 TB beyond a small allowance costs about \$74 (i.e., ~\$0.074/GB), making it slightly cheaper than AWS but still very expensive versus unmetered VPS. [cdn5](#)

Security offerings include Anti-DDoS and WAF, but details and pricing are more opaque compared to AWS/GCP. As with other hyperscalers, advanced protections are paid options. Backups and snapshots exist for ECS and RDS and are fairly standard.

TCO & Pricing Transparency

Example pricing from 2025 shows: a "VM Small" (1 vCPU, 1 GiB RAM) at about \$15.21/month, "VM Medium" (4 vCPU, 8 GiB) at ~\$89.79, and "VM Large" (8 vCPU, 16 GiB) at ~\$170.49, plus \$11.20/month per 100 GB block storage and \$43.84/month for 1 TB object storage with modest reads/writes. A free egress allowance of only ~10 GB/month means almost all production traffic is metered; 1 TB extra egress is about \$74. With 5 TB/month, you'd spend roughly \$370/month in egress alone, on top of compute and DB costs. Price hikes announced in 2026 (5–34 % increases on many services) add further cost uncertainty. [heise](#)

Billing predictability is moderate at best (5–6/10): list pricing is documented, but frequent changes and multiple discount models complicate forecasting.

Support & Emergency Response

Alibaba targets enterprises and large regional players; support is structured accordingly. For Western SMBs, language and time-zone differences can hinder quick resolution. Engineering depth is solid, but you must navigate a large organization and support tiers similar to AWS/GCP.

⚠ Warning: For a global e-commerce operator headquartered outside China, Alibaba's rising prices, metered egress, and ecosystem friction make it a risky and often unnecessary complexity unless you specifically target Chinese customers and need deep local integration.

Verdict

Alibaba Cloud is primarily attractive if your e-commerce business is heavily focused on China/Asia and you need local presence and regulatory compliance. For general global e-commerce with a focus on cost-effective hosting, it delivers hyperscaler-level complexity and bandwidth pricing without clear benefits over AWS/GCP in your context. It is better suited to large regional incumbents and multi-national enterprises than to lean teams wanting simple, cheap, fast hosting. [getdeploying](#)

Oracle Cloud Infrastructure (OCI)

Performance & Hardware

OCI offers flexible compute (including "Always Free" Arm-based and x86 instances) with competitive performance in many independent tests. For small e-commerce deployments, the Always Free VM shapes can host a small site or staging environment at zero cost, though production needs usually exceed these limits. Oracle's managed MySQL/PostgreSQL services provide strong durability and performance comparable to AWS RDS, but documentation is less ubiquitous in the community. [fullmetalbrackets](#)

Reliability & Uptime

Oracle targets enterprise workloads and maintains a high availability SLA, especially when using its multi-AD regions. It does not feature prominently in outage news compared to AWS/Azure, partly due to smaller share but also decent stability. However, detailed community outage narratives are rarer, and you should still plan multi-AD redundancy.

UX & Management

The OCI console is improving but has a noticeable learning curve; concepts like compartments and policies, VCNs, and security lists require time to learn. For a developer with no OCI

experience, expect a few days to get comfortable. Compared to AWS, it may be slightly simpler; compared to DO/Hetzner, it is significantly more complex.

Network, Security & Backups

OCI's most remarkable feature for your use case is outbound data: the Always Free tier allows up to 10 TB/month of egress at no cost for free accounts, which is extremely generous versus other clouds. A detailed breakdown notes that you can egress up to 10 TB/month free, with any traffic beyond that billed at standard rates; for paid accounts, the free allowance is still large compared to competitors. In theory, a small-mid store staying under 10 TB/month could run with effectively zero bandwidth charges. [reddit](#)

“Up to 10 TB of data egress per month across all cloud services is very generous when compared to other cloud providers...” [fullmetalbrackets](#)

OCI provides DDoS protection baked into its edge, plus a WAF service for L7 protection, though pricing and ease of use lag behind Cloudflare and hyperscaler WAFs. Backups and snapshots for compute and DBs are present and work as you'd expect, but the ecosystem of third-party guides is still smaller.

TCO & Pricing Transparency

The Oracle Cloud Free Tier is structured as (a) a 30-day trial with \$300 credits and (b) a set of Always Free services, including compute, Autonomous DB, and 10 TB/month egress. When the trial ends, if you do not upgrade, your account is limited to small free resources, and object storage over 20 GB may be deleted unless you upgrade. For paid accounts, Oracle's pricing is competitive on compute and very competitive on egress, but DB/storage pricing can be complex. [topuser](#)

Billing predictability is moderate (6–7/10): the big free egress cushion reduces surprise bills, but the combination of trial, Always Free quotas, and standard pricing can confuse new users. One Reddit thread emphasizes that exceeding the 10 TB Always Free outbound limit means any extra traffic is billed at normal rates — you won't be *auto-billed* on a free-tier-only account, but paid users must be cautious. [reddit](#)

⚠ Warning: The Free Tier can give a false sense of “Oracle is cheap forever”; once your workload grows beyond Always Free limits, real DB and compute costs can climb closer to hyperscaler levels. You must monitor usage carefully to avoid surprise bills after upgrading.

Support & Emergency Response

OCI's support is geared toward enterprise customers and larger projects. For basic free-tier users, support options are limited. For paid accounts, response quality is acceptable, but

community-level how-tos and runbooks for typical e-commerce patterns are far less abundant than for AWS/DO.

Verdict

OCI is a niche but interesting option if you want to exploit its very generous 10 TB/month free egress, especially for testing or low-to-medium volume stores. It is less attractive if you prioritize ecosystem, documentation, and simplicity — areas where DO/Hetzner shine. For an operator with some Oracle experience or strong cost sensitivity around bandwidth, OCI can be part of a multi-cloud strategy, but it is not the first choice as a primary platform for most e-commerce shops. [oracle](#)

IBM Cloud

Performance & Hardware

IBM Cloud offers both virtual servers and bare metal. Virtual servers for VPC have competitive CPU performance (e.g., good CoreMark scores) and a strong emphasis on network bandwidth (marketing mentions up to 80 Gb/s in some contexts). VPSBenchmarks comparisons between IBM Cloud and OVH show mixed performance grades: some IBM VPS plans achieve mid-range web and CPU grades but often trail OVH in web performance and disk I/O at similar price points. IBM's bare metal is well-suited for high-throughput workloads but is overkill for a typical online store. [vpsserver](#)

IBM Cloud advertises a 99.99 % SLA for many services, indicating a strong focus on availability. [invgate](#)

Reliability & Uptime

Enterprise customers rely on IBM Cloud for critical workloads; the platform has a strong SLA posture and aims for high availability. Outage events are less visible in public discourse than AWS/Azure, but that's partly due to smaller usage in the web-app space. When issues occur, they can be serious and involve complex custom setups.

UX & Management

IBM's portal is aimed at enterprise IT rather than indie developers. It feels clunkier than DO/Hetzner and has a steeper learning curve, especially around account structure, resource groups, and security/configuration primitives. For a new developer, onboarding time is significant; IBM Cloud is not optimized for quick "spin up a WordPress store" workflows.

Network, Security & Backups

IBM's network is robust and global, and they emphasize security and compliance. DDoS protection and WAF are available but positioned as part of a broader enterprise security stack. Backup options for virtual servers and DBs exist, but again, the main selling point is enterprise-grade governance rather than simplicity.

TCO for bandwidth and compute is not dramatically cheaper than AWS/Azure, and sometimes worse, but IBM does not offer the kind of unmetered traffic found on OVH/Hetzner.

VPSBenchmarks shows IBM plans with "Unlimited" data transfer, but network performance grades vary, suggesting that while traffic is unmetered, you pay in performance and/or higher base prices. [vpsbenchmarks](#)

TCO & Pricing Transparency

IBM's pricing is relatively complex and often negotiated for larger clients. Publicly listed prices for some virtual servers are competitive, but once you factor in support and network options, total cost is not attractive for a small store. Billing predictability is moderate but oriented toward enterprise procurement rather than credit-card-paying developers.

Support & Emergency Response

This is IBM Cloud's real strength: paid customers get access to robust support, including architects and deep technical staff. For a small e-commerce operator, however, the cost of such support and the overhead of entering IBM's ecosystem outweigh the benefits relative to simpler hosts.

⚠ Warning: Unless you already have an IBM relationship and need its compliance/governance stack, IBM Cloud is usually the wrong tool for a typical e-commerce site: you'll bear enterprise complexity and pricing without commensurate benefit.

Verdict

IBM Cloud is best suited to large enterprises and regulated industries, especially those already using IBM services. For a standalone online store seeking fast TTFB and good DB performance at reasonable cost, IBM Cloud is rarely the optimal choice. It can work technically, but DO/Hetzner/Vultr/OVH provide the same or better performance at a fraction of the complexity and price. [invgate](#)

Microsoft Azure

Performance & Hardware

Azure offers a wide array of VM sizes similar to AWS, with strong performance on newer D-, E-, and F-series instances. For MySQL/PostgreSQL, Azure Database for MySQL/PostgreSQL provides managed clusters with good IOPS and HA options. Raw performance is not the issue; cost and complexity are, just as with AWS.

Reliability & Uptime

Azure has had several notable outages over the last few years. A July 30, 2024 outage caused by a DDoS attack impacted Azure App Services, IoT, logging, and the Azure portal for nearly 10 hours, with Microsoft acknowledging that misconfigured defenses amplified the impact. A major storage incident on November 13, 2024 caused connectivity errors to Blob Storage and Data Lake across many regions, with widespread downstream impact on services like Databricks and serverless compute. Another significant downtime event on December 26, 2024 in South Central US was tied to a power incident, impacting VMs, storage, Azure SQL and more. [forbes](#)

Azure status history and post-incident reports are detailed, but the pattern is clear: large multi-service incidents do happen, and failover often requires your own architectural design (multi-region deployments, etc.).

UX & Management

Azure Portal is powerful but complex. For a new developer, learning Azure's RBAC, resource groups, VNets, and service identities is time-consuming. Simple tasks like resizing VMs or updating NSG rules are manageable but occur in a UI designed for large organizations rather than small shops. Onboarding time is similar to AWS: days to weeks for a proper production deployment.

Network, Security & Backups

Azure's network and global footprint are excellent, but egress pricing is similar to AWS/GCP: around \$0.087–0.09/GB for the first TB in some regions, meaning roughly \$78–90 for 1 TB egress and over \$4 000 for 50 TB. Azure DDoS Protection offers a Basic (free) level and paid tiers (IP Protection and Network Protection), with Network Protection plans starting around \$2 944/month according to pricing examples and third-party analyses. That is on top of any WAF/Front Door costs. [statusgator](#)

Every Azure property gets some infrastructure-level DDoS protection by default, but for serious attack resilience and cost protection, you must pay for DDoS Protection SKU and often pair it

with Azure Front Door WAF. Backups/snapshots are robust and integrated with Azure Backup. [learn.microsoft](#)

TCO & Pricing Transparency

Compute pricing is similar in magnitude to AWS; egress and premium DDoS/WAF features dominate TCO for high-traffic sites. Billing predictability is moderate: the Azure pricing calculator is detailed, but actual monthly bills depend heavily on traffic patterns and security configuration. Promotional credits (e.g., Visual Studio subscriptions) may temporarily mask full cost.

[statusgator](#)

Support & Emergency Response

Azure's support tiers (Developer, Standard, Professional Direct, Premier) provide increasing levels of response speed and access, but all serious production deployments should have at least Standard/ProDirect support, which costs extra. During major outages, resolution is constrained by platform issues regardless of your plan.

⚠ Warning: Like AWS, Azure's combination of metered bandwidth, pricey DDoS protection, and managed DB costs makes it easy for an e-commerce store to incur four-figure monthly bills even at modest traffic, unless architected extremely carefully.

Verdict

Azure is appropriate mainly if you are a Microsoft-centric enterprise or need Azure-specific services (e.g., tight integration with on-prem AD, Microsoft 365, etc.). For a pure e-commerce site, it offers no decisive advantage over other hyperscalers but shares their high data-transfer and complexity burden. For a lean team, Azure is typically not a good primary host for a cost-sensitive store. [forbes](#)

OVHcloud

Performance & Hardware

OVHcloud's VPS and dedicated ranges are known for aggressive pricing and large included bandwidth. However, performance tests place OVH's VPS stream copy bandwidth at about 47 GiB/s, noticeably behind Hetzner (~105 GiB/s) and even behind some cheaper hosts, suggesting mid-tier CPU/memory/I/O performance. VPSBenchmarks comparisons between OVH and IBM Cloud show OVH VPS plans with mixed grades: some good disk I/O grades but weaker web performance or CPU ratings on lower-end tiers. [experte](#)

For database workloads, OVH's NVMe-backed VPS and especially dedicated servers can perform well, but variability and noisy-neighbor effects are more likely on cheaper VPS.

Reliability & Uptime

OVH has experienced high-profile incidents historically (e.g., the Strasbourg DC fire in 2021, outside our timeframe) and ongoing smaller outages. That said, their unmetered VPS products are widely used, and many users report acceptable uptime for the price. Their status reporting and incident communication are more basic than hyperscalers but generally adequate.

UX & Management

The OVHcloud control panel is more fragmented and less intuitive than DO/Hetzner: managing projects, VPS, dedicated, and networking can feel clunky. Onboarding time for a new developer is moderate; you can get a VPS online quickly, but some operations (e.g., reverse DNS, specific firewall options) are less discoverable.

Network, Security & Backups

OVH's key value proposition is unmetered traffic. Their "Unmetered VPS" line advertises unlimited bandwidth with a fixed port speed (e.g., 400 Mb/s) and no per-TB billing. However, community reports and ToS analysis reveal caveats: users note that for VPS, bandwidth appears effectively limited after roughly 1–4 TB/month depending on plan, with OVH reserving the right to throttle to 1 Mb/s for "excessive use" even on "unmetered" plans. One user reports seeing throttling around 4 TB on a 1 Gb/s plan, and an HN thread quotes OVH's ToS clause about restricting bandwidth to 1 Mb/s in cases of excessive use, despite unmetered marketing. [ovhcloud](#)

"OVHcloud reserves the right to restrict the VPS Service bandwidth to 1 Mbps... in cases of excessive use by the Client — but it advertises with 'unmetered'..."
[news.ycombinator](#)

OVH provides robust anti-DDoS infrastructure at L3/L4 by default and offers a specialized "Game DDoS Protection" that adds L7 protections for gaming protocols; documentation emphasizes advanced filtering and close integration with edge hardware, though this is mainly for dedicated "Game" servers and may not apply to all VPS products. For HTTP e-commerce traffic, default anti-DDoS is helpful but often needs to be combined with Cloudflare or a WAF for full L7 protection. [github](#)

Backups for VPS are typically optional add-ons, sometimes with daily backups included on premium plans (e.g., an unmetered VPS-1 listing includes "daily backup of the previous 24 hours"). Snapshot/restore functionality exists but is not as polished as DO/Hetzner's. [ovhcloud](#)

TCO & Pricing Transparency

OVH's VPS-1 unmetered plan lists 4 vCores, 8 GB RAM, 75 GB SSD, daily backup, and unlimited traffic with 400 Mb/s public bandwidth starting at about \$6.46/month, which is extremely aggressive. Several VPSs and dedicated plans in OVH's portfolio include "unlimited" bandwidth, but the aforementioned ToS throttling clause and community experiences indicate that truly continuous max-throughput usage is not realistic. Billing predictability for most customers is high (9/10) because there is no per-TB metering; the "gotcha" is performance throttling rather than surprise invoices. [reddit](#)

⚠ Warning: OVH's "unmetered" promises hide performance throttling rather than fees. For e-commerce, that means you *won't* get a \$4 000 egress bill, but your site might quietly slow to a crawl if you push beyond what they consider reasonable usage.

Support & Emergency Response

Support is ticket-based with limited real-time options on basic plans. Community sentiment often criticizes OVH's support as slow and scripted, with difficulty escalating to engineers. For urgent incidents, you may wait hours for meaningful action. Documentation is functional but not at DO's level.

Verdict

OVH is compelling when you need lots of bandwidth at a fixed, low cost and are willing to accept middling performance and weak support. It suits operators who can tolerate some throttling and have multi-provider redundancy, or who host mainly static content or backups. For a single primary e-commerce store where consistent TTFB and DB performance matter, OVH's VPS is a **second-tier** choice behind Hetzner/Vultr/DO, though OVH dedicated servers with game-grade DDoS can be excellent if you manage them actively. [vpsbenchmarks](#)

Contabo

Performance & Hardware

Contabo is infamous for offering huge resource allocations at extremely low prices (e.g., 4 cores, 6 GB RAM, 100 GB disk for a few pounds/euros/month), which immediately signals aggressive over-selling. Multiple community reports show very high CPU steal: one user running [top](#) on a 6-core Contabo VPS frequently saw 20–30 % CPU steal on all cores, spiking to 50 % under load — a clear indication of heavy over-allocation of physical CPU. Others on WebHostingTalk confirm that the standard VPS line is "heavily oversell" with no strict CPU limitation, suggesting you share CPU with many noisy neighbors. [reddit](#)

“Contabo is known to overallocate and have a lot of high CPU steals... I want to see 2% or less.” [reddit](#) “They’re offering 4 cores, 6GB ram and 100gb disk for £4.50 a month. That’s enough to put me off... clear sign of overselling (which likely explains your downtime).” [reddit](#)

VPSBenchmarks’ Contabo tests emphasize that high CPU steal percentages correlate with over-allocated hosts, and Contabo’s overall value scores are mediocre despite low list prices. Disk performance is also weaker: one independent review measured Contabo’s VPS stream copy bandwidth (~36 GiB/s) significantly below Hetzner and OVH. For database-heavy e-commerce, such CPU steal and I/O inconsistency are poisonous to TTFB and checkout response times. [vpsbenchmarks](#)

Reliability & Uptime

Contabo’s uptime varies widely; some users report acceptable stability, while many threads on Reddit and WebHostingTalk complain about intermittent downtime, overload, and performance degradation, especially on the cheapest plans. The low price point encourages very dense packing of customers onto host nodes; when combined with CPU steal issues, this leads to unpredictable performance more than outright hard downtime. [lowendtalk](#)

UX & Management

The Contabo panel is basic but functional; it allows VM creation, OS reinstallation, and simple management. It lacks the polish and UX efficiency of DO/Hetzner/Vultr. Onboarding time is small — you can deploy a VM quickly — but there are fewer guardrails and “niceties” like sophisticated firewall UI or one-click managed DBs.

Network, Security & Backups

Contabo advertises “unlimited” traffic on many plans with modest port speeds, but the same caveats as OVH apply: heavy usage and certain traffic patterns (e.g., mining, VPN) can trigger throttling or policy enforcement. DDoS protection is minimal; you should absolutely front Contabo with a CDN/WAF for public e-commerce. [lowendtalk](#)

Backups and snapshots are available but rudimentary compared to the more modern clouds, and restore workflows are less polished. Network performance is okay for basic use, but not a selling point.

TCO & Pricing Transparency

Contabo’s sticker prices are incredibly low: a Contabo Cloud VPS 10 (4 vCPU, 8 GB RAM) is around €4.50–5/month in some comparisons, but the real cost is hidden in degraded performance and downtime. VPS providers comparison guides rate Contabo’s overall rating

lower (e.g., 3.8/5) than DigitalOcean/Hetzner due to these trade-offs. Billing predictability is high (you pay a fixed cheap amount), but the hidden cost is lost conversions due to slow or unstable performance. [vpsbenchmarks](#)

⚠ Warning: For serious e-commerce, Contabo's high CPU steal and overselling are **deal-breakers**. The "savings" are quickly eclipsed by lost revenue from slow checkouts and poor user experience.

Support & Emergency Response

Support is minimal and mainly via tickets. Community sentiment often describes support as slow and unhelpful for performance complaints ("everything is within our limits"). You will not get meaningful assistance tuning or troubleshooting; they will ensure the VM is up, nothing more.

Verdict

Contabo is acceptable for non-critical workloads (lab environments, backups, low-traffic blogs) where performance swings are tolerable. It is **not** fit for primary hosting of any serious e-commerce store where TTFB, database latency, and uptime directly map to revenue. For your use case, Contabo should be excluded except perhaps for off-site backups or non-production usage. [reddit](#)

Overall Pattern and Practical Takeaways

Even without a global summary table, several patterns emerge:

- **Best fit for cost-sensitive e-commerce:** Hetzner (especially CCX/CAX with external CDN/WAF) and Vultr/DO for regions or features Hetzner lacks. [dev](#)
- **Bandwidth traps:** AWS, GCP, Azure, and Alibaba all impose egress rates around \$0.07–0.11/GB; 5–10 TB/month can mean hundreds to low thousands of dollars in transfer fees alone. [cdn5](#)
- **"Unlimited" VPS caveats:** OVH and Contabo's unmetered bandwidth is constrained not by billing but by throttling and noisy neighbors; they avoid four-figure invoices but may silently degrade performance under heavy load. [thatmy](#)
- **Free tier illusions:** Oracle's 10 TB free egress, AWS/GCP credits, and other promos are great for experimentation but do not reflect steady-state production costs; operators must model post-promo TCO carefully. [topuser](#)

For a modern, revenue-sensitive e-commerce stack that values low TTFB, good DB performance, and predictable bills, a pragmatic architecture often looks like:

- **Primary compute + DB on Hetzner, DO, or Vultr** (prefer Hetzner in EU/US if you can live with manual DB, DO if you want managed DB and top-tier docs). [experte](#)
- **Fronted by Cloudflare or a similar CDN/WAF** for global latency reduction and serious L7 DDoS mitigation, rather than relying on provider-native WAFs.
- **Off-site backups** possibly stored on an independent provider (S3-compatible object store, Oracle free tier, or a second VPS provider) to hedge against provider-level disasters. [fullmetalbrackets](#)