

Virtual Machine Hosting: Choosing the Right Cloud for Your Workload

AWS, Azure, Google Cloud, and the alternatives — a practical guide for enterprise IT buyers.

Dave Hoon
Chief Technology Officer, Norseman Defense Technologies · February 2026

The decision to move workloads to virtual machine hosting used to be simple: pick a cloud provider, pick an instance size, and go. **That era is over.** Today's enterprise IT leaders face a market with five credible hyperscalers, dozens of specialized providers, and a new generation of bare-metal-as-a-service options that blur the line between cloud and dedicated infrastructure. *The wrong choice doesn't just cost money — it creates technical debt that takes years to unwind.*

At Norseman, we've helped organizations across federal, commercial, and healthcare markets navigate infrastructure decisions where the stakes are high and the marketing is louder than the reality. This is a practical, opinionated guide to the commercial VM hosting landscape — built for enterprise IT buyers who need to make the right call the first time.

The Big Three: AWS, Azure, and Google Cloud

The hyperscaler market is effectively a three-party system, with each provider having carved out distinct advantages that make them genuinely better choices for specific workloads — not just marketing differentiation.

- **Amazon Web Services (AWS)** — The market leader with the broadest service catalog and the deepest ecosystem. AWS EC2 remains the gold standard for VM flexibility, offering more instance families than any competitor: compute-optimized (C-series), memory-optimized (R and X-series), storage-optimized (I and D-series), GPU instances (P and G-series), and the increasingly popular Graviton ARM-based instances that deliver **40% better price-performance than comparable x86 instances** for the right workloads. AWS wins on breadth, ecosystem maturity, and the sheer number of certified partner solutions available. *If your team doesn't have a strong reason to go elsewhere, AWS is the defensible default.*
- **Microsoft Azure** — The enterprise integration play. If your organization runs Microsoft 365, Active Directory, SQL Server, or any significant Windows workload, Azure's hybrid licensing benefits — particularly the **Azure Hybrid Benefit** — can reduce VM costs by up to 85% compared to pay-as-you-go pricing. Azure's tight integration with Entra ID (formerly Azure AD), Microsoft Defender, and the broader Microsoft security stack makes it the natural home for organizations already deep in the Microsoft ecosystem. Azure also leads in hybrid cloud scenarios through **Azure Arc**, which extends Azure management to on-premises and multi-cloud environments. *For Windows-heavy shops, Azure's total cost of ownership often beats AWS despite higher nominal instance pricing.*
- **Google Cloud Platform (GCP)** — The engineering-first choice. GCP's Compute Engine consistently benchmarks competitively on raw performance, and Google's networking infrastructure — built on the same backbone that powers Search and YouTube — delivers measurably lower latency for globally distributed workloads. GCP's **Sole-Tenant Nodes** provide dedicated physical hosts for compliance-sensitive workloads, while **Confidential VMs** offer hardware-level memory encryption for sensitive data processing. Google's pricing model, including sustained use discounts that apply automatically without committed use agreements, makes it attractive for variable workloads. *GCP is the right answer for data-intensive, globally distributed, or Kubernetes-native workloads — and it's underused by enterprises that default to AWS or Azure without evaluating it.*

Oracle Cloud Infrastructure: The Sleeper Pick

Oracle Cloud Infrastructure (OCI) is consistently underestimated by enterprise IT teams conditioned to think of Oracle as a database company. **That's a mistake that's costing organizations real money.** OCI's VM pricing is aggressive — often **50–70% below equivalent AWS instances** — and its bare-metal instances offer consistent, predictable performance that cloud-native VMs can't match for latency-sensitive workloads.

The calculus is particularly compelling for Oracle database workloads. Organizations running Oracle Database Enterprise Edition on-premises are eligible for **Bring Your Own License (BYOL)** on OCI, which eliminates the licensing cost entirely for cloud instances. Combined with OCI's lower compute pricing, the TCO advantage over running Oracle workloads on AWS or Azure can be *substantial* — often 60–80% lower for equivalent database performance.

OCI's **Always Free tier** is also genuinely useful for development and testing: two AMD-based VMs with 1 OCPU and 1 GB RAM, plus 200 GB of block storage, at no cost indefinitely. For small workloads and proof-of-concept environments, it's worth evaluating seriously before paying for AWS or Azure dev instances.

Specialized Providers: Where the Hyperscalers Fall Short

The hyperscalers are optimized for scale and breadth. For specific workload patterns, specialized providers consistently outperform on price, performance, or simplicity.

- **Hetzner** — A German cloud provider with data centers in Europe and the US that delivers bare-metal and virtual machine pricing that makes hyperscaler costs look absurd for straightforward compute workloads. A dedicated server with 32 cores, 128 GB RAM, and 2× 1.92 TB NVMe SSDs runs around $150/month. The equivalent workload on AWS would cost **five to ten times more.** Hetzner is the right answer for compute-intensive workloads that don't require hyperscaler-specific services and can tolerate a less mature managed services ecosystem. *European data residency requirements are a natural fit.*
- **Vultr and Linode (Akamai Cloud)** — Developer-friendly cloud providers with predictable flat-rate pricing and global data center coverage. Both are significantly cheaper than AWS, Azure, or GCP for general-purpose VMs, with simpler interfaces and faster provisioning for teams that don't need the complexity of hyperscaler consoles. Linode's acquisition by Akamai added enterprise CDN and DDoS protection capabilities that make it a credible option for web-facing workloads where edge performance matters.
- **CoreWeave** — The GPU compute specialist. For AI inference, machine learning training, and high-performance computing workloads, CoreWeave's H100, A100, and H200 GPU availability and pricing consistently outperform AWS, Azure, and GCP. As demand for GPU compute exploded with generative AI adoption, the hyperscalers have faced chronic GPU shortages. CoreWeave built its infrastructure specifically for GPU-dense workloads and can deliver capacity — and pricing — that the hyperscalers often can't match. *For organizations building AI pipelines, CoreWeave deserves serious evaluation alongside the big three.*

Cost Management: The Hidden Variable

Hyperscaler pricing is intentionally complex. **The published on-demand rate is almost never what sophisticated buyers pay.** Every major provider offers significant discounts through committed use:

- **AWS Reserved Instances and Savings Plans** — 1 or 3-year commitments deliver 30–72% discounts versus on-demand pricing. Compute Savings Plans offer flexibility across instance families and regions. *Any production workload running continuously should be on a savings plan.*
- **Azure Reserved VM Instances** — Similar 1 or 3-year commitments, with the added benefit of combining with Hybrid Benefit for Windows Server and SQL Server to maximize savings.
- **GCP Committed Use Discounts** — 1 or 3-year commitments for compute resources, plus automatic sustained use discounts for VMs running more than 25% of the month.
- **Spot and Preemptible Instances** — All three hyperscalers offer deeply discounted (60–90% off) instances that can be reclaimed with short notice. For fault-tolerant, batch, or stateless workloads, spot instances are one of the highest-leverage cost optimizations available.

FinOps discipline — treating cloud spend as a first-class engineering concern — consistently delivers 20–40% cost reductions in organizations that implement it systematically. Tools like **AWS Cost Explorer, Azure Cost Management, and GCP Recommender** surface rightsizing opportunities that are often left on the table in organizations without dedicated cloud financial governance.

Self-Hosted and Hybrid: The On-Premises Tier

The public cloud conversation dominates industry media, but for a significant portion of enterprise workloads — particularly in regulated industries, federal environments, and organizations with substantial existing infrastructure

investment — **the right answer isn't a hyperscaler at all.** It's a self-hosted virtualization platform running on owned or leased hardware, managed with the same operational discipline as any cloud environment.

This conversation has never been more relevant. **Broadcom's 2023 acquisition of VMware triggered immediate, dramatic licensing restructuring** that left thousands of organizations facing cost increases of 3–5× overnight. VMware's perpetual licensing model — the foundation of enterprise virtualization for two decades — was replaced with mandatory subscription bundles that forced customers to pay for capabilities they didn't need. The result has been the largest disruption in enterprise VM hosting in a generation, and a genuine re-evaluation of on-premises virtualization platforms that were previously overlooked.

Nutanix: The VMware Migration Destination

Nutanix has emerged as the primary beneficiary of VMware customer defection, and for good reason. **Nutanix AHV (Acropolis Hypervisor)** is a mature, enterprise-grade hypervisor that ships as part of the Nutanix Cloud Infrastructure stack — included in the base license at no additional cost. For organizations paying VMware per-socket or per-core licensing, the elimination of a separate hypervisor licensing cost alone can justify the migration.

Nutanix's architecture is hyperconverged infrastructure (HCI) — compute, storage, and networking managed as a single software-defined system across commodity or purpose-built hardware nodes. The operational model is deliberately cloud-like: **Prism Central**, Nutanix's management plane, provides a single console for VM lifecycle management, performance monitoring, capacity planning, and policy enforcement across all clusters — whether on-premises, in colocation facilities, or extended to public cloud through **Nutanix Cloud Clusters (NC2)** on AWS or Azure.

- **VMware migration tooling** — Nutanix Move automates VM migration from VMware vSphere environments with minimal downtime, handling disk format conversion, network mapping, and cutover orchestration. Organizations that have spent years building VMware infrastructure can migrate without rebuilding from scratch.
- **Licensing simplicity** — Nutanix's per-node subscription model is predictable and scales linearly. There are no per-VM charges, no per-core multipliers, and no surprise licensing audits. What you pay is what you get.
- **Data resilience built in** — Nutanix's distributed storage fabric provides data protection, compression, deduplication, and erasure coding natively, without requiring a separate SAN or NAS investment. *For organizations running VMware on traditional three-tier infrastructure (compute + SAN + network), consolidating onto Nutanix HCI typically reduces hardware footprint and operational complexity simultaneously.*
- **Federal and compliance credentials** — Nutanix maintains FedRAMP authorization for cloud deployments and supports STIG-hardened configurations for DoD environments. For federal customers evaluating on-premises VM platforms, Nutanix's compliance posture is well-established.

HPE Morpheus: Multi-Cloud Management and Private Cloud Orchestration

HPE Morpheus occupies a distinct position in the VM hosting landscape — it's not a hypervisor, it's a **cloud management platform (CMP) and infrastructure orchestration layer** that sits above the hypervisor and unifies management across on-premises infrastructure, private clouds, and public clouds through a single interface and API.

Morpheus supports VMware vSphere, Nutanix AHV, KVM, Hyper-V, and OpenStack as on-premises compute targets, alongside AWS, Azure, GCP, and OCI as public cloud targets. **The value proposition is operational unification** — a

single self-service catalog, a single cost management view, a single automation framework, and a single policy engine regardless of where the workload runs.

- **Private cloud without the complexity** — Morpheus provides OpenStack-level functionality — tenant isolation, self-service provisioning, quota management, chargeback — without the operational overhead of running OpenStack. For organizations that want a genuine private cloud experience on their own hardware, Morpheus delivers it on top of the hypervisor they already own.
- **Day-2 operations at scale** — Morpheus includes integrated lifecycle automation: provisioning workflows, configuration management (Ansible, Puppet, Chef integration), backup orchestration, patch management, and decommissioning. VM sprawl — the silent cost driver in most virtualized environments — is addressable through Morpheus's policy engine and reclamation workflows.
- **Cost visibility across environments** — Morpheus ingests cost data from public cloud providers and calculates on-premises costs using configurable rate cards, delivering a unified FinOps view across hybrid environments. *For IT leaders trying to make honest comparisons between running workloads on-premises versus in public cloud, Morpheus provides the data to make that call with confidence.*
- **HPE ecosystem integration** — For organizations running HPE ProLiant, Synergy, or Alletra infrastructure, Morpheus integration with HPE OneView and HPE GreenLake creates a fully managed private cloud experience with infrastructure-as-code provisioning down to the hardware layer.

The VMware Displacement Opportunity

The Broadcom-VMware situation deserves direct attention because it's reshaping the self-hosted VM market in real time. **Organizations that renewed VMware contracts before the acquisition or are now facing renewal under Broadcom's new licensing terms are in a fundamentally different economic position than they were two years ago.**

The typical VMware-to-alternative migration evaluation follows a predictable pattern: initial sticker shock at the new Broadcom pricing, a request for alternatives, a Nutanix proof-of-concept, and ultimately a decision point between absorbing the VMware cost increase, migrating to Nutanix on-premises, or accelerating a cloud migration that was already on the roadmap. Each path has a legitimate case depending on workload characteristics, team capabilities, and capital vs. operating expense preferences.

What the market has learned is that **the decision is rarely binary between VMware and public cloud.** Nutanix as the hyperconverged platform, Morpheus as the management layer, and selective use of public cloud for burst capacity and cloud-native services represents a hybrid architecture that delivers the operational benefits of cloud without the wholesale transfer of workloads — and costs — to a hyperscaler.

Updated Decision Framework

With self-hosted and hybrid platforms included, the workload-first framework becomes more complete:

- **Microsoft-centric workloads** (Windows Server, SQL Server, Active Directory) → Azure with Hybrid Benefit, or on-premises Nutanix/Hyper-V if data sovereignty or existing hardware investment applies
- **Oracle database workloads** → OCI with BYOL for cloud, or on-premises with Nutanix for maximum performance and licensing control

- **Globally distributed, latency-sensitive applications** → GCP for cloud-native; Nutanix NC2 for hybrid with consistent management
- **GPU-intensive AI/ML workloads** → CoreWeave for cloud; on-premises GPU servers managed through Morpheus for sensitive or air-gapped environments
- **VMware environments facing Broadcom renewal** → Evaluate Nutanix AHV migration; use Morpheus to bridge the transition period and unify management post-migration
- **Organizations wanting private cloud without public cloud cost exposure** → Nutanix HCI plus HPE Morpheus for self-service provisioning, chargeback, and lifecycle automation
- **Federal, classified, or air-gapped workloads** → On-premises Nutanix (FedRAMP authorized, STIG-hardened) with Morpheus for orchestration; public cloud only where authorization exists
- **Straightforward compute without managed service dependencies** → Hetzner, Vultr, or Linode for dramatic cost reduction in non-sensitive environments

The Multi-Cloud Reality

Most enterprise organizations end up with workloads spread across multiple providers — not by design, but by accumulation. Business units make independent decisions, acquisitions bring inherited infrastructure, and SaaS vendors run on their own preferred clouds. **Multi-cloud is the reality. Multi-cloud strategy is the choice.**

The organizations that manage multi-cloud environments successfully treat it as a deliberate architectural decision with intentional governance — standardized networking through transit architectures, unified identity through federation, consistent security posture through cloud security posture management (CSPM) tools, and FinOps processes that span providers. The organizations that struggle treat each cloud as an isolated silo and pay the complexity tax indefinitely.

The right cloud for your virtual machine workloads isn't the one with the best marketing, the largest customer logo reel, or the most aggressive sales team. It's the one that aligns with your workload characteristics, your licensing position, your team's existing skills, and your total cost of ownership over a realistic planning horizon. Take the time to build that analysis — the decision compounds for years.

Norseman's infrastructure practice helps commercial and government organizations evaluate, architect, and optimize cloud environments across all major providers. If you're navigating a cloud platform decision or looking to reduce your current cloud spend, contact us to start a conversation.

Published at norseman.com/insights/commercial-vm-hosting-comparison. To talk with Norseman Defense Technologies, call 410.579.8600 or email sales@norseman.com.