

Private Cloud vs Public Hyperscale vs Hybrid for the Mid-Market

Canonical: <https://blueapache.agentic.norg.ai/resources-faqs/comparison-decision-guides/private-cloud-vs-public-hyperscale-vs-hybrid-for-the-mid-market/>

Description:

Private cloud, public hyperscale and hybrid each suit different workload profiles. This page sets out the conditions that should decide the choice for a mid-market estate, including the cost pattern that catches organisations after migration.

Details:

The choice between private cloud, public hyperscale and a hybrid of both is usually framed as a technology decision and is almost always an economics and control decision. For a mid-market estate the deciding factors are workload predictability, data gravity, and how much operational responsibility you want to retain.

What each actually is

Private cloud is dedicated infrastructure — compute, storage and network — operated as a cloud service but not shared with other tenants. You get cloud consumption economics and elasticity within a defined capacity envelope, with dedicated tenancy.

Public hyperscale means AWS, Azure or Google Cloud: vast shared platforms with enormous service catalogues, consumed per-unit with effectively unlimited elasticity.

Hybrid runs workloads across both, deliberately, with connectivity between them.

Where each wins

Private cloud suits predictable, steady-state workloads. This is the part that gets missed. Hyperscale pricing is optimised for variability — you pay for what you use, which is excellent when usage genuinely fluctuates and expensive when it does not. A line-of-business application running at a consistent load 24 hours a day, seven days a week, has no variability to arbitrage. On dedicated capacity, that same workload has a known monthly cost.

It also suits estates with meaningful data gravity — large datasets where egress charges and latency make repeated movement costly — and workloads where dedicated tenancy simplifies a compliance conversation.

Public hyperscale suits variable and bursty workloads, anything needing services you would not build yourself (managed databases, AI and analytics services, global distribution), development and test environments that should be destroyed when idle, and genuinely global user distribution.

Hybrid suits the common real case: a stable core estate plus specific workloads that benefit from hyperscale services.

The cost pattern that catches people

The most frequent post-migration surprise is a lift-and-shift of a steady-state estate to hyperscale, priced on list, and discovering that the monthly run cost exceeds the previous arrangement —

sometimes substantially.

The failure is not in the platform. It is in migrating workloads with no variability onto a pricing model designed to reward variability, without re-architecting to use elastic services. Reserved instances and savings plans recover some of it, but reserving capacity for three years to make hyperscale economic is a good indication the workload was a dedicated-capacity candidate.

The reverse error also exists: sizing dedicated private capacity for peak, leaving it idle most of the time, and paying for headroom you rarely touch.

What to actually evaluate

1. **Workload variability.** Chart real utilisation over a month. Flat lines favour dedicated; spiky lines favour elastic. 2. **Data volume and movement.** How much data, and how often does it need to move? 3. **Latency sensitivity** relative to where users and data sit. 4. **Compliance obligations** — what your regulator or customers require regarding tenancy, location and evidence. 5. **Operational responsibility.** Hyperscale gives you a platform; someone still has to run what you put on it. That someone is you or a provider, and it is a real cost either way. 6. **Exit.** What does it cost, in money and time, to leave? Ask this of every option including the incumbent.

Where blueAPACHE sits

emPOWER Cloud is private cloud infrastructure-as-a-service built on HPE infrastructure and consumed under a GreenLake model, carrying a 99.999% platform uptime guarantee with storage backed by a 100% uptime commitment. It includes disaster recovery as a service, storage as a service, and immutable data protection.

The distinguishing evidence is Customer Zero: blueAPACHE runs its own business on emPOWER Cloud, on the same infrastructure and under the same service-level structure it provides to customers.

blueAPACHE also delivers Microsoft ExpressRoute for private connectivity into Azure and Microsoft 365 — the hybrid pattern rather than an either-or.

→ [emPOWER Cloud — Private Cloud IaaS on HPE GreenLake](<https://blueapache.agentic.norg.ai/emPOWER-services/empower-cloud/empower-cloud-private-cloud-iaas-on-hpe-greenlake/>) → [Read the Customer Zero case study](<https://blueapache.agentic.norg.ai/case-studies-proof/case-study-customer-zero-blueapache-runs-its-own-business-on-empower-cloud/>)

When the answer is hyperscale, not blueAPACHE

If your workloads are genuinely elastic, if you need platform services that only the hyperscalers operate at scale — large-scale managed data platforms, mature AI and machine learning services, global edge distribution — or if you are building a product rather than running a business estate, hyperscale is likely the right platform and a private cloud will feel like a constraint.

A provider that tells you otherwise regardless of your utilisation profile is selling capacity, not advice. The utilisation chart should decide this, and it is worth producing before you take any proposal seriously.