

HPE GreenLake: G-Cloud Pricing Template (Example)

Overview of HPE GreenLake Services

HPE GreenLake is a hybrid cloud, pay-per-use platform that delivers infrastructure and cloud services on-premises while providing a public-cloud-like consumption experience. It provides managed on-premises infrastructure, including compute, storage, networking, and associated services. Pricing is consumption-based with a minimum commit and variable usage above that commit.

How the GreenLake Pricing Model Works

- Baseline Capacity (Minimum Commit): Minimum committed capacity billed monthly.
- Buffer Capacity: Additional onsite buffer (typically 20–30%) available for burst usage.
- Usage-Based Billing: Pay only for consumption above the baseline.
- Management & Support: Hardware management, updates, and support included.

Example G-Cloud Pricing Table

Pricing Component	Unit of Measure	Minimum Commit	Buffer Capacity	Description	Price (POA)
vCPU / Compute Instance Usage	Per vCPU per Month	Customer-defined baseline	20–30%	Metered compute consumption	POA
VM Instance Consumption	Per VM per Month	Minimum VM baseline	Auto-scale	Usage based on VM provisioning	POA
RAM Consumption	Per GB RAM per Month	Defined baseline	Burst buffer	Memory consumption billed monthly	POA
Block Storage	Per GB per Month	Baseline GB/TB	20–30%	High-performance block storage	POA
File/Object Storage	Per GB per Month	Customer baseline	Buffer available	Scalable storage services	POA
Network Throughput	Per Gbps per Month	Baseline bandwidth	On-demand burst	Managed networking service	POA