



GLENCORE

[Glencore](#) is one of the world's largest global diversified natural resource companies and a major producer and marketer of more than 60 commodities. With approximately 150,000 employees and contractors across 30+ countries, the company's operations span mining, processing, refining, and distribution across 50 global locations, serving the automotive, steel, power generation, battery manufacturing, and oil sectors.

Industry

Natural Resources / Mining & Metals

Products

- VMware® Cloud Foundation®

Glencore Controls Memory Costs with VCF NVMe Memory Tiering

When server memory costs surged, Glencore needed a fundamentally different approach to infrastructure economics. The global natural resource company used NVMe memory tiering capability, a standard feature in VMware Cloud Foundation (VCF), to decouple memory capacity from DRAM cost. By pairing high-performance DRAM with high-capacity NVMe drives in a single unified memory pool, the Glencore architecture team substantially reduced per-server provisioning costs versus pure DRAM expansion. The result is the ability to support disaster recovery capacity and application memory requirements with a financial model that scales.

About Glencore

Glencore is one of the world's largest diversified natural resource companies and a major producer and marketer of more than 60 commodities. With approximately 150,000 employees and contractors operating across more than 30 countries, Glencore's industrial and marketing activities span mining, processing, refining, and distribution, serving customers in the automotive, steel, power generation, battery manufacturing, and oil sectors. Headquartered in Baar, Switzerland, the company manages a geographically dispersed IT environment.

Managing memory costs across a global enterprise at scale

Running IT infrastructure for one of the world's largest commodity enterprises is not a simple proposition. Glencore operates roughly 70 vCenters across approximately 50 locations, a footprint that spans continents, time zones, and operational environments. At the center of that architecture are the infrastructure engineers responsible for defining the standards, blueprints, and reference architectures that the Glencore IT organization follows worldwide.



Glencore's infrastructure engineers have long been responsible for defining the group standards, blueprints, and reference architectures used across the global IT organization. With the introduction of VCF in production, the team gained a more effective mechanism for translating these established standards into consistent platform configurations worldwide. VCF fleet management capabilities help reinforce governance while simplifying implementation across the global infrastructure estate.

Before that transformation could fully take hold, the team faced a problem that was escalating fast, driven by the soaring cost of server memory. DRAM prices have climbed sharply in recent years, and Glencore experienced this increase with stark clarity. A server procurement quote more than tripled in price over the original price in a matter of months. That single data point crystallized what the team already understood: relying solely on DRAM to meet enterprise memory requirements was becoming financially untenable.

A further complication existed at the Glencore headquarters. The team operates a metro cluster configuration with workloads split across two sites designed to run all production at either location independently. To support this resilience model, no more than 50 percent of memory on any host can be in active use at any given time. This constraint effectively cuts usable DRAM in half, a significant operational factor when memory costs are already running high.

Glencore was already realizing substantial operational value from VCF. NVMe memory tiering now extends that advantage. The DRAM cost savings it delivered provided a concrete, quantifiable proof point that reinforced the platform's business case, adding measurable cost avoidance

on top of the capabilities the team had already come to rely on.

The team even found small areas where the nature of application memory use could be positively impacted by VCF memory tiering. For example, Glencore has a global developer Virtual Desktop Infrastructure (VDI) system. Standard VDI builds require 48 to 96 gigabytes of RAM per virtual desktop; actual utilization typically runs between 10 and 15 percent. A messaging service requires Glencore to reserve a minimum of 128 gigabytes, yet active memory consumption is minuscule compared to the stated requirement. Using NVMe memory tiering the team could accommodate these application memory minimums without a correspondingly large outlay for DRAM memory.

NVMe memory tiering: Building intelligence into the platform

VCF addresses the memory cost challenge through NVMe memory tiering, a capability that combines fast DRAM (Tier 0) with high-performance NVMe solid-state drives (Tier 1) to present a single contiguous memory pool to virtual machines. The system intelligently manages memory placement. Active, frequently accessed pages remain in DRAM for maximum performance, while cold or dormant pages migrate to the NVMe tier automatically.

The economics are compelling. A 1.6-terabyte NVMe drive carries roughly the same price as 64 gigabytes of DRAM. The default Glencore configuration doubles available memory by matching DRAM capacity with an equivalent NVMe allocation, with 512 gigabytes of DRAM paired with 512 gigabytes of NVMe yielding one terabyte of total memory. The cost differential between DRAM and NVMe memory is significant, and at Glencore's scale, that differential compounds across dozens of sites and hundreds of servers.

Glencore's analysis quantifies the difference. Expanding a server by 1.6 TB using a single 1.6TB NVMe device reduces memory costs substantially and therefore overall server costs when compared to provisioning a similar amount of DRAM memory. VCF memory tiering gives Glencore a cost-effective way to right-size DRAM purchases, especially where workloads rarely exceed 80% active memory, which is common in the metro environment where half of DRAM is unassigned for resiliency purposes.



The Glencore team first encountered NVMe memory tiering on vSphere 8, where setup required a degree of manual configuration. The guides were clear and the process manageable, but it required deliberate effort. The transition to the latest version of VCF removed that friction entirely. Memory tiering is now a single cluster configuration setting. What once required careful, step-by-step intervention became a routine operational task.

The team's response to these capabilities was decisive: every new Cisco UCS server deployed across Glencore infrastructure now includes an NVMe card designated for memory tiering. What began as an evaluation of a specific feature has become a foundational design principle.

The deployment path itself reflected the team's disciplined approach to production risk. Glencore ran an extensive technical evaluation across multiple environments, testing the platform's limits with stretched configurations that placed management and workload domains across geographically separated sites. These included Brisbane and Sydney in Australia and locations in Baar, Switzerland. The team worked through the full VCF feature set, including automation integration and identity broker deployments. When early access to VCF 9.1 became available, Glencore made a deliberate decision to retire the test environments and rebuild on VCF. The team arrived at production readiness already deeply familiar with the platform.

Today, 50% of Glencore Linux developer VDI systems use VCF NVMe memory tiering in production.

More memory, lower costs, and no developer complaints

The impact of NVMe memory tiering at the Glencore Baar headquarters is already measurable, and the trajectory points toward significant expansion.

Today, 50 percent of Glencore Linux developer VDI systems use VCF NVMe memory tiering in production. Across the full workload cluster in Baar, approximately 5 percent of all workloads are using the capability, a figure that reflects an early-stage rollout, not a ceiling. Every new server added to the environment ships with an NVMe card, extending the tiering footprint with each hardware refresh cycle.

The metro cluster configuration in Baar, which had previously constrained usable memory to 50 percent of installed DRAM, benefits directly from tiering. NVMe memory effectively doubles the memory available to workloads within the same physical host footprint, turning a resilience-driven constraint into a manageable operating condition rather than a persistent capacity bottleneck.

Looking ahead, the operational objective is to fully utilize installed DRAM under normal conditions and treat NVMe tiering as the overflow buffer that activates only in cases of high server utilization or the event of a site disaster, eliminating the standing 50 percent reservation and recovering that capacity for active production workloads. Glencore is also looking to extend beyond VDI to other workloads such as business applications and databases.

Broadcom, through its portfolio of VMware products, has given the Glencore infrastructure team what they needed to make a compelling business case, enforce a consistent global standard, and build infrastructure that scales without proportionally scaling cost. The architecture is set. The standard is in place. And with every new server that ships with an NVMe card, the savings compound.