


The storage foundation for Neoclouds.

Whether you rent GPUs by the hour, are launching a GPU cloud, or scaling an established one, VDURA delivers hyperscaler-grade storage on commodity hardware you can site, secure, and scale in any region. Flash feeds the GPUs, WD UltraSMR HDDs hold petabyte capacity, all in one global namespace.



60%

Lower acquisition cost vs all-flash

Up to 2.7 TB/s

Throughput per rack to saturate GPUs

<1 day

Software-defined deploy, any colocation

Why VDURA × WD for Neoclouds



01 MULTI-TENANT BY DESIGN

Every renter isolated. No noisy neighbors.

Software-defined storage runs on commodity servers you procure and site in any country or colocation. Per-tenant AES-256 encryption, KMIP keys, VIP, and VLAN isolation keep workloads — and tenants — cleanly separated.

AES-256 • KMIP • VLAN ISOLATION • NO LOCK-IN

03 NEOCLOUD TIME-TO-MARKET

GPU storefronts live in a day.

New nodes light up as fast as you rack them. One flash node fans out to 6–7 GPUs. Per-tenant VIP, namespaces, and billing APIs make multi-tenant clouds turnkey.

MULTI-TENANT • TERRAFORM / CSI • 60% LOWER COST

02 PROVEN HPC HERITAGE

25 years feeding the most demanding compute.

Founded by Garth Gibson, co-inventor of RAID and architect of the pNFS standard. The parallel file system behind NASA, Boeing, and Airbus.

1,000+ DEPLOYMENTS • 50+ COUNTRIES

04 MIXED-FLEET ECONOMICS

Flash for speed. Disk for scale.

NVMe flash and WD UltraSMR live in one global namespace and data plane. No second object store, no external movers. Ride independent cost curves for flash and disk.

26 PBe/RACK • 2×+ PERF/WATT • ONE NAMESPACE

ONE PLATFORM ACROSS THE ENTIRE AI PIPELINE



1,500+

Nodes per cluster, linear scaling

6 nines+

Availability in production

Build your Neocloud on a proven foundation.

[Book a Neocloud workshop →](#)

vdura.com/industries/neocloud