DATASHEET

VDURA Data Platform V5000

NVMe Performance and Petabyte-Scale Capacity in One Parallel File System

75 GB/s	2.7 TB/s	26 PBe	100/0 to 2/98
PER FLASH NODE	PER RACK, ALL-FLASH	PER RACK, MIXED FLEET	FLASH / HDD RATIO

V5000 Key Features

- DirectFlow high-performance parallel Client
- RDMA CPU-bypass data paths
- S3 protocol support
- Native Kubernetes CSI plug-in
- VeLO: Velocity Layered Operations
- Elastic Metadata Engine, up to 20x
- VPODs: Virtualized Protected Object Devices
- Multi-Level Erasure Coding
- Context-Aware Tiering
- SMR HDD optimization
- Multi-tenancy with per-tenant VLANs and keys
- REST API and infrastructure-as-code providers
- Prometheus metrics, Grafana dashboards
- End-to-end AES-256 encryption, KMIP
- Data compression in each Storage Node
- Non-disruptive expansion and upgrades
- Ethernet and InfiniBand
- VDURACare Premier: 10-year drive replacement

The VDURA® Data Platform V5000™ Certified Platform is engineered for AI and HPC pipelines that demand relentless GPU feed rates. Built on industry-standard servers, V5000 pairs NVMe flash with optional HDD capacity expansion in a Mixed Fleet architecture: flash sized for performance, HDD for capacity, and a cost-balanced growth path from a pilot program to production on the exabyte scale.

V5000 runs VDURA Data Platform V12, VDURA's flash-tuned parallel file system, streaming multiple terabytes per second per rack from a single global namespace. V12 delivers 20% or more additional throughput, 20x metadata acceleration and a 20% or greater reduction in cost per TB at the same performance level versus V11, with a minimum downtime in-place upgrade. Working with the DirectFlow® Client, VDURA offers parallel redundant data paths that scale linearly, safeguard data with enterprise-class durability and keep day-to-day management simple.

V5000 Director Nodes

Director Nodes anchor the VDURA control plane, running the flash-optimized VeLO™ (Velocity Layered Operations) metadata engine that coordinates billions of file operations and keeps AI pipelines saturated. Each 1U server combines high-core CPUs, plentiful DRAM and dual 200 or 400 Gb/s links, and the V12 Elastic Metadata Engine scales metadata performance dynamically across nodes. Toolless service and redundant power maximize uptime.

V5000 Flash Nodes

Flash Nodes form the performance-driven foundation of the data plane. Each 1U node holds 12 NVMe SSDs from 3.84 TB to 122.88 TB and delivers up to 75 GB/s. Configure as all-NVMe flash for maximum performance, or pair with HDD capacity expansion for economical bulk storage. RDMA moves data between GPU servers and storage with no CPU in the path.



Figure 1. V5000 Flash Node.

V5000 Capacity Expansion Enclosures

Two 4U, 19-inch capacity expansion enclosure options are available, fully populated with 16, 24, 30 or 32 TB HDDs: one with 78 drives for one-meter deep racks and one with 108 drives for 1.4-meter racks. V12 SMR HDD optimization unlocks 25 to 30% more usable capacity per rack. Both include two Titanium-level redundant power supplies and connect over SAS 4 (24G).



Figure 2. V5000 capacity expansion enclosure.

V5000 Performance Specifications	
Throughput	Up to 75 GB/s per Flash Node Up to 2.7 TB/s per 42U rack, all-flash; 200 GB/s with capacity expansion
IOPS	Up to 1.2M per Flash Node; 3.6M for a minimum filesystem Up to 45M per 42U rack, all-flash; 9M with capacity expansion
Capacity*	Up to 17.7 PBe per 42U rack, all-flash; 26 PBe with capacity expansion
Data Durability	Up to 8 nines, increasing with scale
Data Availability	Six nines and above in production environments
Protocols	DirectFlow Client (POSIX), S3, CSI; RDMA
Security	End-to-end AES-256, KMIP key management, per-tenant keys, RBAC

* Effective capacity after data compression

42U Rack Configuration Options (not including switches)	
1U Director Node (x3 minimum)	
VCH-5000-D1N	Director Node
1U Flash Node (x7)	
VCH-5000-F1N	Flash-only, 12 x NVMe SSD (3.84 to 122.88 TB)
VCH-5000-F1N	Flash with HDD capacity expansion support
4U JBOD, 78 or 108 Drive (x7)	
VCH-5000-J108-3456	108 x 32 TB HDDs
VCH-5000-J108-3240	108 x 30 TB HDDs
VCH-5000-J108-2592	108 x 24 TB HDDs
VCH-5000-J108-1728	108 x 16 TB HDDs
VCH-5000-J78-2496	78 x 32 TB HDDs
VCH-5000-J78-2340	78 x 30 TB HDDs
VCH-5000-J78-1872	78 x 24 TB HDDs
VCH-5000-J78-1248	78 x 16 TB HDDs

VeLO metadata engine

The distributed, flash-optimized metadata engine coordinates file operations across the realm. The V12 Elastic Metadata Engine scales metadata performance across Director Nodes as the system grows.

VPODs and erasure coding

VPODs™ (Virtualized Protected Object Devices) isolate failures to one virtual device, not the server. Per-file client-side erasure coding (Dual Parity, Single Parity or Striped Mirror) combines with optional Multi-Level Erasure Coding, so durability increases with scale.

Context-Aware Tiering

Data lands on flash and automatically tiers to HDD based on size, access pattern and temperature, keeping the flash tier lean and hot. Active datasets, checkpoints and model files stay on NVMe; completed work drains to high-density HDD without administrator intervention.

Multi-tenancy and automation

Tenants are isolated with per-tenant VLANs, per-tenant AES-256 keys (BYOK), and per-tenant telemetry with hourly metering. The REST API handles tenant lifecycle, quota and billing automation: three calls create a tenant, assign capacity and provision a volume that is mountable in seconds, with consumption data granular enough to invoice from.

Surprising simplicity

Every node is managed as one realm from a single GUI, CLI or API. Prometheus-compatible metrics feed Grafana dashboards. Data compression runs in each Storage Node with zero client-side impact.

Ultimate configurability

Configure from 100% flash to 98% HDD in one namespace. Scale out with additional Flash Nodes for throughput and IOPS, or additional capacity expansion enclosures for bulk storage, without re-architecture. V12 upgrades in place with zero downtime.

V5000 Technical Specifications	VCH-5000-D1N	VCH-5000-F1N	VCH-5000-J78-*	VCH-5000-J108-*
Description	Director Node server, 1U	Flash Node server, 1U	78-drive JBOD, 4U	108-drive JBOD, 4U
Performance	N/A	Up to 1.2M IOPS Up to 75 GB/s	Up to 13.2 GB/s	Up to 17.8 GB/s
Media	12 x NVMe SSD slots	12 x NVMe SSD, 3.84 to 30.72 TB	78 x HDD, 16 to 32 TB	108 x HDD, 16 to 32 TB
Connectivity	2 x NDR/NDR200 or 400/200/100 GbE, RoCE v2	4 x 200 Gb or 2 x 400 Gb, InfiniBand or Ethernet, RoCE v2	SAS 4 (24G) to Storage Node	SAS 4 (24G) to Storage Node
Dimensions W x D x H (without bezel)	17.5 x 33 x 1.75 in. 45 x 84 x 5 cm	17.5 x 33 x 1.75 in. 45 x 84 x 5 cm	17 x 32 x 7 in. 43 x 81 x 18 cm	17 x 41 x 7 in. 43 x 104 x 18 cm
Weight with Drives	35 lbs. / 15.9 kg	38.8 lbs. / 17.6 kg	207.2 lbs. / 94.5 kg	263.2 lbs. / 119.4 kg
AC Power Input	200 to 240 V, 50/60 Hz, C14 inlet (all models)			
Power Consumption (typical / max at plug)	636 W / 981 W	742 W / 1158 W	839 W / 1155 W	1081 W / 1448 W
Thermal Output (typical / max)	2647 / 4097 BTU	2532 / 3951 BTU	2862 / 3940 BTU	3688 / 4940 BTU
Power Supplies	1+1 Titanium	1+1 Titanium	1+1 Titanium	1+1 Titanium

About VDURA

VDURA is the high performance data storage infrastructure for AI factories and Neoclouds. Powered by HYDRA, its software-defined, mixed-fleet architecture keeps every GPU fed, every tenant isolated, and every cluster online, delivering hyperscale-class durability and economics with up to 8 nines of durability and a heritage of more than 1,000 production deployments. Visit vdura.com for more information.

Worldwide Office
1-888-726-2727
info@VDURA.com

Headquarters
Pittsburgh, PA, USA

Offices & R&D
Milpitas, CA and Niwot, CO, USA

EMEA
Oxford, UK
emeainfo@VDURA.com

APAC
Sydney, Australia
apacinfo@VDURA.com

China
Shanghai, China
chinainfo@VDURA.com

© 2026 VDURA, Inc. All rights reserved. VDURA, the VDURA logo, VeLO, VPOD, VDURAbility, V5000, DirectFlow, VDURACare Premier and Context-Aware Tiering are trademarks or registered trademarks of VDURA, Inc. in the U.S. and/or other countries. All other trademarks, registered trademarks, trade names, company names and service marks are the respective properties of their holders. Specifications subject to change without notice.
v1.5 | October 2026