



ENTERPRISE PRODUCT BROCHURE

NexaVM nHPS

Hyper-Performance Block Storage

A high-performance distributed block storage system — software-defined and cloud-native, delivering 1M random IOPS per volume with sub-millisecond latency to fully unleash next-generation hardware.

No.1

SPC-1™ Leaderboard — first distributed storage to break the centralized-array monopoly

27.2M

IOPS delivered in SPC-1 benchmark with record-breaking latency

0.144_{ms}

Average latency — 0.217 ms peak, under half of peer products

4,096


Nodes per system — zero service interruption during scale-out

PRODUCT OVERVIEW

Software-Defined. Cloud-Native. Built for Speed.

NexaVM nHPS (NexaVM Hyper-IO Storage) is a high-performance distributed block storage system built on NexaVM's proprietary architecture. Software-defined and cloud-native, it enables agile deployment and simplified management — easily reaching **1M random IOPS** and **sub-millisecond concurrent latency per volume** to fully unleash next-generation hardware.

nHPS serves high-performance databases, large-scale high-performance virtualization, containers, and HPDA storage — empowering core business growth across finance, telecom, energy, media, healthcare, and education industries.



nHPS software architecture — from application workloads down to commodity X86/ARM servers

ALL-FLASH & HYBRID FLASH IN ONE

One Platform, Every Workload

No separate storage procurement for workloads with different requirements — nHPS software-defined storage freely combines storage media to meet performance and capacity needs on demand. All enterprise workloads run on a single storage platform, avoiding the resource waste and operational complexity of siloed deployments.

PRODUCT ADVANTAGES

Ultimate Performance

SPC-1™ NO.1

27,201,325

IOPS · peak latency 0.217 ms · average latency 0.144 ms

Built on a fully in-house high-speed storage engine, nHPS is the **first-ever distributed storage to break the traditional centralized-array monopoly** on the SPC-1 leaderboard, taking the No.1 spot. Key resource usage is only **one-third of peer products** on the SPC-1 list, with IOPS up to **29.51% higher** and latency **nearly halved**.

The nHPS storage engine leverages flash, storage-class persistent memory, and NVMe over RDMA with an in-house minimalist high-performance IO architecture — shortening the IO path, boosting concurrent throughput, reducing thread switching, and eliminating locks to maximize the potential of new storage media and accelerate end-to-end data transfer.


LINEARLY SCALABLE ARCHITECTURE

Scale from 3 to 1,024 Nodes — No Ceiling, No Downtime

nHPS adopts a fully distributed architecture that scales out to linearly increase total capacity and performance. Its sharded, evenly-spread no-bottleneck design and automatic load balancing place data and metadata evenly across servers — **no component in the data I/O path or O&M framework limits scaling**.

With **zero service interruption**, a cluster scales from 3 commodity rack servers to 1,024 — completely removing the scale ceiling of traditional SAN controller clusters. IOPS grows nearly linearly as nodes are added, with barely any latency fluctuation.

Aggregate IOPS Scales Linearly with Nodes



Latency remains stable at ~0.35–0.37 ms across all scale-out stages — barely any fluctuation.



Multi-Tier Reliability, Always-On Business

- Elastic erasure coding and multi-replica redundancy, with placement policies at drive, server, and rack levels
- Intelligent sub-health handling correlates underlying and business metrics to precisely handle drive, network, and node sub-health
- Built-in data consistency verification detects and corrects silent data corruption end to end
- Active-Active / Active-Standby stretched clusters support 1+1, 2+1, and 2+2 topologies



Key Enterprise Storage Features

- ROW snapshots, consistency-group snapshots, volume cloning, online expansion, thin provisioning
- iSCSI multipathing; CHAP mutual authentication, in-flight encryption, fine-grained RBAC
- Metro stretched clusters and cross-cluster asynchronous replication for DR solutions
- Graphical O&M console, drive locate LED / roaming, whole-node replacement, scheduled health inspection

FULL CLOUD ECOSYSTEM COMPATIBILITY

Native Interfaces for the Modern Cloud Stack

nHPS provides fully functional iSCSI and NVMe drivers for mainstream compute virtualization platforms, and fully supports the vhost interface for deep integration with virtualization and cloud platform partners — optimizing front-end and back-end IO performance for more competitive converged compute-storage products and solutions.

10x

Read IOPS vs iSCSI over NVMe-oF & RDMA, at sub-millisecond latency

5x

Write IOPS vs iSCSI — strongly supporting the cloudification of core businesses

3

Ecosystem interfaces: VAAI (VMware), CSI (Kubernetes), vhost (cloud platforms)

TYPICAL USE CASES

One Storage Platform for Every Critical Workload


Virtualization

Database

Container

AI

HPDA

PRODUCT SPECIFICATIONS

Interface Types	iSCSI, NVMe-oF, Vhost	Max Capacity per Volume	512 TB
Supported Nodes per System	3 – 4,096	iSCSI MPIO	Up to 6 paths
Max Storage Pools per Cluster	64	Max Snapshots per System	65,535
Max Volumes per Cluster	640,000	Max Snapshots per Volume	65,535
Max Drives per Storage Pool	50,000	Max Clones per Volume	128
Max Drives per Server Node	Unlimited	Network Types	RoCEv2, TCP
Max Capacity per Server Node	Unlimited	Drive Types	NVMe SSD, SATA SSD, HDD

Ready to unleash your core business?

Contact NexaVM for a proof of concept tailored to your workload.

www.nexavm.com