Lab Report

High-Density HDD ZFS Performance Evaluation on AIC SB407-VA server

Author: Igor Misko, Senior Product Marketing Engineer,
Business Development Storage Products,
Toshiba Electronics Europe GmbH

Introduction

High density HDD platforms continue to play a critical role in environments where vast sequential datasets must be moved, analysed, or archived efficiently. As organisations build peta-byte scale data lakes, long retention backup repositories, and virtualisation infrastructures with large numbers of concurrently active workloads, the Toshiba Electronics Europe (TEE) HDD Laboratory team wanted to put the following under test: how much performance can be extracted from HDD only ZFS systems when all external bottlenecks are removed?

To answer this, the team engineered a directly attached, non-networked testbed: a 60 drive ZFS configuration based on **Toshiba's MG11ACA24TE 24 TB HDDs**, organised as 6 × RAIDZ2 virtual devices (vdevs) – 8 data + 2 parity – deployed in an AIC SB407-VA platform and tuned to avoid external throughput bottlenecks. By excluding network interfaces and SSD based caching layers, the system exposes the **true, native behaviour of RAIDZ2** under sustained load, providing insight into what users can realistically expect when designing large ZFS pools for bulk storage scenarios.

Picture 1: AIC SB407-VA

The findings of this report are particularly relevant for architects of **data intensive workflows**, including:

- **Data lake and analytics pipelines**, where multi terabyte sequential reads and writes dominate,
- **Backup and archival repositories**, which require predictable, sustained throughput for ingest and restore operations,
- **Suitable for VM/container back ends** when most I/O is large-block sequential or cached; for many active, small-random workloads, add NVMe (special vdev/L2ARC) or move hot sets to SSD/NVMe tiers. **Enterprises evaluating all HDD ZFS deployments** for sequential and streaming workloads as a cost efficient alternative to hybrid or flash accelerated systems.

Helping users and customers to understand the practical limits of HDD only ZFS performance and on top providing a data driven basis for planning cluster layouts, virtual devices (vdev) topologies, and capacity efficient architectures is the scope of this lab report.

Configuration

Physical Platform

Chassis: AIC SB407-VA - 4U high-density storage chassis, 60 bays

Deployment model: Network-less (no network interface (NIC) data path during tests; no external fabric)

Host Hardware

CPU: 1 × Intel® Xeon® Platinum 8468

Memory: 4 × DDR5 4800 MHz 16 GB (Micron MTC10F1084S-1RC56BD1 QLFF) → 64 GB total

HBA: Broadcom 9500-16i, IT mode (no RAID offload)

Storage Media

Drives: 60 × Toshiba MG11ACA24TE (24 TB) enterprise HDDs

Caching devices: None (no SLOG, no L2ARC)

Software

Platform: TrueNAS SCALE 25.10.1

Filesystem: ZFS (ZVOL used for Virtual Machine (VM) block)

ZFS Pool and ZVOL

Pool: AIC60-pool

Topology: 6 × RAIDZ2 vdevs, each 10 disks (8 data + 2 parity) → 60 total

ZVOL: 750 TB, volblocksize=128 K, sync=disabled, exposed to VM as a block device

Virtualization

Hypervisor: KVM/QEMU (TrueNAS SCALE)

Guest: Windows Server 2022 Standard

vCPU: 1 vCPU (8 cores)

Guest RAM: 12 GB

Storage path: VirtIO-SCSI / VirtIO-Block to 750 TB ZVOL

Guest Operating System Storage Configuration

The ZVOL device presented to the Windows Server 2022 virtual machine was initialised as a single logical volume and formatted using ReFS with a 64 KB allocation unit (cluster) size. This configuration aligns the guest file system with the underlying ZFS volblocksize=128 K, reducing write amplification and ensuring efficient large-block sequential I/O. No additional Windows-side caching or tiering mechanisms were enabled.

Networkless Test Methodology

All benchmarks were executed locally on the host via the VM; no network adapters are used for data I/O. This eliminates NIC bandwidth and TCP/IP overhead. Results therefore reflect the intrinsic limits of the HDD array + RAIDZ2 + ZVOL + VirtIO stack.

Benchmarking Procedure

Tools & Workloads

Tool: FIO

Patterns: Large sequential read/write

Scripts:

```
fio --filename=test100 --size=100G --direct=1 --rw=read --bs=128k --iodepth=128 --time_based
--runtime=240 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=24
--norandommap --randrepeat=0

fio --filename=test100 --size=100G --direct=1 --rw=write --bs=128k --iodepth=64 --time_based
--runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=12
--norandommap --randrepeat=0
```

Note: ZVOL bypasses ARC for data; ARC remains metadata-oriented.

Controls

No L2ARC or SLOG configured.

Tests run long enough to avoid cache warm-start effects.

FIO Results

Sequential Read (128 K, QD≈1024, 16 threads): 16.0 GiB/s (17.2 GB/s), ~131k IOPS, avg latency 7.8 ms, p99 8.8–9.5 ms (p99.9 ~24 ms)

Sequential Write (128 K, QD64, 12 threads): 12.0 GiB/s (12.9 GB/s), ~98.5k IOPS, avg latency 7.7 ms, p99 51 ms

Interpretation

The configuration delivers approximately 85–90 per cent of the theoretical maximum for this vdev width and drive class. The dominant limiting factors are the RAIDZ2 parity overhead and the aggregate bandwidth and IOPS capability of the HDDs. In this operating regime, CPU resources, system memory, and virtualisation overheads do not act as performance bottlenecks. Such behaviour makes the system particularly well suited to workloads that process or stream very large files in sustained sequential operations, such as multi terabyte backup ingestion or bulk scientific dataset transfers.

Use Cases

The system's performance profile makes it suitable for a range of data intensive workloads, particularly those that benefit from sustained sequential throughput or broad parallelism across multiple vdevs:

Large-Scale Media or Scientific Data Lake (Sequential Throughput Showcase)

High-bandwidth ingest and egress of large files, e.g., media mastering, satellite imagery, simulation outputs.

Backup Repositories for Veeam/Commvault

Capacity-efficient, predictable large-block throughput for periodic backups and synthetic fulls.

High-Concurrency VM or Container Storage (Random I/O Showcase)

Wide striping across six vdevs supports many concurrent streams (latency not comparable NVMe-class devices).

Database Archival or Analytics Warehouse

Optimised for sequential scans, ETL, and reporting over cold and warm datasets.

Note: Use cases described in this report assume workloads that either do not require strict synchronous write guarantees or employ application level data protection mechanisms.

CPU Configuration and Expansion Capabilities

The system operates with a single installed Intel® Xeon® Platinum 8468 processor. Even with only one CPU populated, the platform provides substantial compute headroom for basic local analytics, which is relevant in this network-less deployment where data processing must occur directly on the storage host.

The underlying motherboard is dual-socket-capable, allowing installation of a second Xeon CPU. A second processor can be dedicated to local analytics or data processing tasks without impacting storage operations, preserving throughput characteristics of the ZFS stack.

TOSHIBA

The platform also exposes multiple PCIe 5.0 x16 slots, enabling straightforward installation of one or more GPU accelerators. Graphics cards can be easily installed, and the system can be upgraded for heavy AI/ML workloads (e.g., feature extraction, inference, vector search, or model fine-tuning) while the storage host continues to serve bulk data.

In addition, the chassis provides six NVMe bays, which enable a hybrid storage configuration for AI and analytics workflows:

- NVMe devices can be assigned to SLOG (for safe synchronous writes), metadata-special vdevs, or L2ARC to accelerate metadata-heavy and small-I/O tasks.
- NVMe can also host high-IOPS scratch space or dataset staging for AI pipelines, while the 60-disk HDD pool supplies capacity and sequential bandwidth for bulk data.
- This separation allows GPUs to stay fed with low-latency working sets from NVMe, while HDDs handle large, sequential reads/writes efficiently.

Operational note: When expanding for AI workloads, ensure adequate power, cooling, and airflow for GPUs and NVMe drives, validate PCIe lane mapping (to avoid slot bifurcation conflicts with the HBA and NICs if/when added), and confirm IOMMU/virtualisation settings if GPUs will be passed through to VMs or containers.

Conclusion

A fully network less, HDD only ZFS system, built from 60 × 24 TB Toshiba MG11ACA24TE drives in a 6 × RAIDZ2 layout and exposed as a 750 TB ZVOL, can achieve sustained throughput of ~15.3 GiB/s for reads and ~12 GiB/s for writes. These results represent roughly 85–90% of the theoretical limit for this vdev width and drive class. By eliminating NICs, SSD caching tiers, and external bottlenecks, the testbed reveals the native sequential I/O capabilities of RAIDZ2 on high capacity HDDs. The architecture delivers predictable, cost- and capacity efficient performance suitable for large data lake ingest, multi terabyte backup operations, scientific workloads, and other environments requiring reliable bulk sequential throughput.

Note of thanks to our partners

This lab report is the result of dedicated collaboration.

"I would like to thank all our partners for their valuable support on this project.

AIC provided the SB407 VA high density storage platform used as the foundation of our 60 drive ZFS testbed. Broadcom supported us with the Host Bus Adapter HBA 9500 16i in IT mode, enabling direct, bottleneck free access to the full HDD array. Together with our Toshiba MG11ACA24TE 24 TB hard disk drives, this setup enabled a controlled, network-less evaluation of pure HDD only ZFS performance, highlighting the sustained throughput achievable at scale without SSD acceleration or network bottlenecks."

Igor Misko
Senior Product Marketing Engineer
Business Development Storage Products
Toshiba Electronics Europe GmbH

Interested in performing a test in our HDD Innovation Lab yourself or planning a visit?
We'd be happy to support you.

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

info@toshiba-storage.com
toshiba-storage.com

Issued 02/2026