



SOLUTION BRIEF

# XINAS: HIGH-PERFORMANCE NFS STORAGE FOR AI & HPC ON SUPERMICRO AS-1116CS-TN SERVERS (SOLUTION VALIDATION)



*Supermicro AS -1116CS-TN*

## Executive Summary

Modern data-intensive workloads (AI/ML training, HPC simulation, genomics, and high-resolution media pipelines) demand shared storage that delivers extreme throughput at low latency and remains stable even in the face of drive failures. In these environments, “traditional NAS” becomes the limiting factor, and NFS is often considered a bottleneck rather than a performance enabler.

xiNAS changes that assumption. xiNAS is Xinnor’s high-performance NFS server solution designed for AI, HPC, and other throughput-hungry environments.

xiNAS is made of three key components:

- xiRAID Classic engine, to provide data protection with exceptionally low overhead
- XFS layer with optimal filesystem geometry matching the RAID layout
- NFS layer tuned for high-performance access over RoCE and or InfiniBand.

In this document, we present the validation of xiNAS on a Supermicro NVMe server, demonstrating performance and resilience across multi-client and multi-server scenarios, including degraded and rebuild states.

## xiNAS Explained

xiNAS is a high-performance NFS solution built on Supermicro NVMe servers and powered by Xinnor's xiRAID engine, delivering throughput that reaches the practical limits of modern storage hardware. Explicitly designed for AI and HPC workloads, xiNAS keeps DGX/HGX systems continuously busy with optimized data delivery.




Figure 1 - xiNAS Solution Architecture

The solution delivers 5-10x<sup>1</sup> higher performance than standard NFS implementations through a carefully optimized stack that includes tuned filesystem settings to avoid write amplification, optimized NFS server and client configurations, and smart network settings. This architecture enables more than 95% GPU<sup>2</sup> utilization in AI training environments while remaining cost-effective for midrange installations.

Key advantages of the xiNAS approach include:

- No proprietary client software required - works with standard NFS protocols
- Simple deployment and maintenance - leverages familiar NFS infrastructure
- Industry-leading performance - using standard NFS semantics with RDMA transport
- Linear scalability - aggregate bandwidth grows with server count in scale-out deployments

At the core of xiNAS is xiRAID, Xinnor's patented software RAID engine designed for NVMe environments. xiRAID delivers superior performance in both normal and degraded modes, with minimal CPU overhead for checksum calculations and no dedicated

hardware requirements. The engine supports flexible RAID configurations (levels 0, 1, 5, 6, 7.3, 10, 50, 60, 70, and N+M) across any drive type and capacity.

## Goals and Validation Scope

The validation was guided by the following practical objectives:

- Establish a hardware baseline to understand the maximum capability of the NVMe media.
- Quantify the efficiency of software data protection (RAID) relative to raw hardware.
- Measure sequential throughput (GB/s) for large-file and streaming workloads and small block transactional performance.
- Demonstrate multi-client scalability and multi-node aggregation.
- Assess business continuity by measuring performance impact during drive failure and rebuild.

## Architecture Topology

### Hardware Building Blocks




Figure 2 - Server Components Layout

### NAS Servers (System Under Test): 2x Supermicro AS-1116CS-TN

- CPU: 1x AMD EPYC 9455 (48 cores)
- Memory: 384GB DDR5
- NICs: 2x NVIDIA BlueField-3 (dual 200GbE) per node
- Data media: 12x Micron 3.84TB Gen5 NVMe per node

- Boot: Micron 7450 480GB M.2

#### **Load Generator Clients:** 2× Supermicro ASG-1115S-NE316R

- CPU: 1× AMD EPYC 9535 (64 cores)
- Memory: 768GB DDR5
- NICs: NVIDIA BlueField-3 (dual 200GbE) + ConnectX-7 (400GbE)

#### **Network Fabric**

- High-speed switching (SN-5600), configured for RoCE v2, enabling RDMA transport for NFS.

#### **Software Stack**

- OS: Ubuntu 24.04 (server and clients)
- Kernel: 6.8.0-85-generic
- Storage engine: Xinnor xiRAID Classic v4.3.0
- Load generation: fio v3.36
- NFS transport: NFS over RDMA (RoCE)

#### **Configuration**

Each 3.84TB NVMe drive is split into two namespaces, enabling a “two-array strategy” that isolates bulk data I/O from metadata/journaling activity:

- Capacity/data array: xiRAID RAID 6 (10+2) built on the large namespaces (~3.80TB per drive)
  - Stripe size: 64 KiB
  - Block size: 4096 bytes
  - Purpose: high-efficiency, high-throughput sequential I/O with dual-drive fault tolerance
- Journal/Metadata accelerator: xiRAID RAID 10 (6+6) built on the small namespaces (~38GB per drive)
  - Purpose: extremely low-latency write target for filesystem log activity

#### **Performance Highlights**

##### **Backend Efficiency (xiRAID Local Baseline)**

On local block-device testing (bypassing filesystem and NFS), the platform demonstrated near-maximum utilization of the NVMe subsystem:

- Sequential read: 176 GB/s
- Sequential write: 61.3 GB/s (full-stripe) and 47.1 GB/s (1 MiB block size)
- Efficiency vs theoretical: 97–100%, with low RAID compute overhead (RAID calculation CPU averages remained only a few percent in healthy mode)

These results indicate that data protection is delivered with minimal performance cost, preserving CPU headroom for networking demands.

## Single-Server xiNAS (NFS, Multi-Client)

When adding XFS and NFS over RDMA on top of xiRAID, with one NAS server, we measured these performances:

- Sequential read: 74.5 GB/s
- Sequential write: 39.5 GB/s

This level of throughput aligns with “all-flash array-class” performance using standard NFS semantics, enabled by NVMe speed plus RDMA transport efficiency.

## Two-Server Scale-Out (NFS, Multi-Client, Aggregated Throughput)

Scaling to a two-node xiNAS deployment (and scaling the clients accordingly) produced high aggregate bandwidth:

- Sequential read (2 servers): 117 GB/s (network-limited)
- Sequential write (2 servers): 79.6 GB/s

Write throughput demonstrated near-linear scaling as nodes were added, confirming readiness for larger clustered deployments where aggregate bandwidth grows with server count.

## Small-Block / Transactional Performance (IOPS)

For metadata-heavy and small-block workloads (4K random I/O, multi-client):

- Random read: 990k IOPS (~265 µs latency)
- Random write: 587k IOPS (~430 µs latency)

This indicates the architecture is not only strong at streaming bandwidth, but also capable of servicing high-operation-rate workloads often seen in AI pipelines (many small files), build farms, and mixed analytics environments.

## Resilience and Business Continuity (Degraded and Rebuild)

A key requirement for production storage is maintaining service levels during drive failures and recovery. In multi-client, multi-server scenarios:

**Healthy baseline:** 117 GB/s sequential read, 79.6 GB/s sequential write (as reported in previous paragraph)

### Degraded (one failed device/namespace):

- Read: 107 GB/s (~8.5% drop vs healthy)
- Write: 81 GB/s (comparable to healthy baseline)

### Active rebuild:

- Read: 102 GB/s (drop <13% vs healthy)
- Write: 75 GB/s (remains close to baseline given rebuild workload contention)

The data shows that read-intensive workloads (common in AI training and fine-tuning) experience minimal disruption during a drive failure and rebuild, while write performance is barely affected.

## Conclusion

These tests, jointly conducted by Xinnor and Supermicro, validate that xiNAS on Supermicro AS-1116CS-TN NVMe servers can deliver enterprise-class, scale-out NAS performance via NFS over RDMA while maintaining strong resiliency.

Key outcomes demonstrated in validation:

- Extreme throughput with a single node and strong scale-out aggregation across two nodes (up to 117 GB/s read and 79.6 GB/s write).
- High storage efficiency with drive failure protection, reaching 97–100% of theoretical NVMe capability in backend tests.
- Operational stability during failure and rebuild, with read throughput dropping <13% during rebuild and write throughput staying near the baseline.
- Excellent performance also for small files and random operations.

These results position the solution as a high-performance shared storage foundation for AI, HPC, and other data-intensive environments requiring both speed and predictable behavior under real-world fault conditions.

## References

1. <https://www.truenas.com/blog/truenas-electric-eel-performance/>
2. [https://msstconference.org/MSST-history/2025/Presentations/Day%201%20\(September%202022nd\)/MSST\\_XINNOR.pdf](https://msstconference.org/MSST-history/2025/Presentations/Day%201%20(September%202022nd)/MSST_XINNOR.pdf)

## SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements. Visit [www.supermicro.com](http://www.supermicro.com)

## XINNOR

Xinnor is an Israeli-based software development company that specializes in creating innovative data storage solutions. Their main product is xiRAID, a software RAID engine that delivers exceptional performance. Visit [www.xinnor.io](http://www.xinnor.io)

## AMD

AMD is the high-performance and adaptive computing leader, powering the products and services that help solve the world's most important challenges. Our technologies advance the future of the data center, embedded, gaming and PC markets.

Founded in 1969 as a Silicon Valley start-up, the AMD journey began with dozens of employees who were passionate about creating leading-edge semiconductor products. AMD has grown into a global company setting the standard for modern computing, with many important industry firsts and major technological achievements along the way. Visit [www.amd.com](http://www.amd.com)