


Xinnor's xiSTORE Running on Celestica's SC4200 (Athena G2) Provides Unmatched Performance for New AI Workloads

XINNOR



Celestica™

INTRODUCTION

The advent of Artificial Intelligence requires very fast storage to keep the CPU busy.

Deep learning models, the workhorses of AI, devour data during training, constantly analyzing massive datasets to learn and improve. GPUs, the specialized processors powering many AI tasks, excel at crunching these numbers, but they get sluggish and inefficient if data isn't delivered fast enough. GPUs are expensive resources, and maximizing their utilization is the key. Slow storage bottlenecks create idle time, where the GPU waits for data instead of working. Faster storage ensures a smooth flow of information, keeping the GPU constantly "fed" and churning out results. Every second a GPU waits for data translates to longer training times, impacting development schedules and resource costs.

To address the requirement of AI workload, Xinnor introduced xiSTORE.

ABOUT xiSTORE

xiSTORE is a software Defined Storage (SDS) solution for the HPC and AI markets, based on xiRAID, the fastest and most reliable RAID engine, combined with Lustre FS clustered file system and commodity hardware, to provide efficient, flexible and scalable storage infrastructure.

Based on building blocks architecture xiSTORE offers highest flexibility and performance for both NVMe and HDD storage setups with industry fastest drive rebuild time. xiSTORE supports different RAID levels including RAID5, 6, 7.3, N+M, nested and declustered. To offer even more reliability xiSTORE applies silent data corruption protection. Moreover, xiSTORE offers Virtual Machine Management and distributed file systems support.

When performance matters, xiSTORE can be implemented entirely on NVMe SSD.

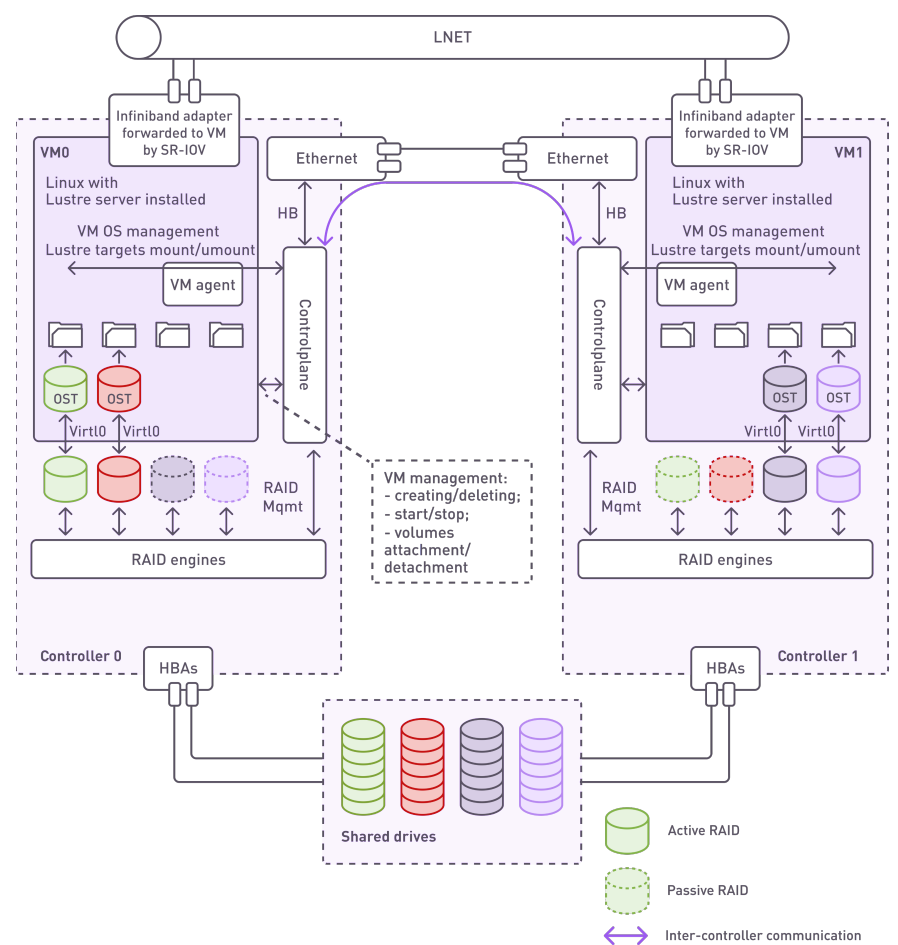


Figure 1. xiSTORE Architecture

ABOUT THE SC4200 PLATFORM

SC4200 (Athena G2) is a next-generation 2U rackmount NVMe platform featuring 24 PCIe NVMe dual-port SSDs and redundant computing nodes powered by Intel® Xeon® Scalable processors.

Harnessing the power of dual Intel® Xeon® Scalable Processors, the system ensures exceptional performance during peak demands, while its all-flash architecture provides swift responsiveness tailored to diverse business applications and environments. Redundant data access to all hot-swappable NVMe modules and power supplies, coupled with optional battery backup units, maximizes uptime, reinforcing system reliability and data integrity even in challenging scenarios.



Figure 2. First testing configuration

Its dual controller architecture with shared NVMe storage makes the SC4200 platform the ideal companion to xiSTORE.

TEST CONFIGURATION AND RESULTS

Tests were performed on the following hardware configuration:

- **CPU:** 2x Intel(R) Xeon(R) Gold 6336Y CPU @ 2.40GHz;
- **Memory:** 32 GB x 16 = 512 GB, Samsung M393A4K40DB3-CWE;
- **Drives:** 22x Kioxia CM6-R KCM61RUL3T84 dual port NVMe;
- **RAID Configuration:** 20 NVMe drives used for 4 RAID 6 (8d+2p) (2 namespaces per NVMe to split each NVMe to 2 block devices, to avoid PCIe x2 limits), 2 NVMe drives for RAID 0;
- **Lustre configuration:** 4 Virtual machines each with 1 OSS/ 1 OST (MDS with 2 NVMe RAID 0 is located on one of the virtual machines);
- **OS:** Rocky 8.7, Lustre 2.15.2;
- **20 Lustre clients:**
 - o Oracle linux 8.8;
 - o Lustre 2.15.3.
- **Performance benchmarking tool:** IO

Test configuration setup is presented on the picture:

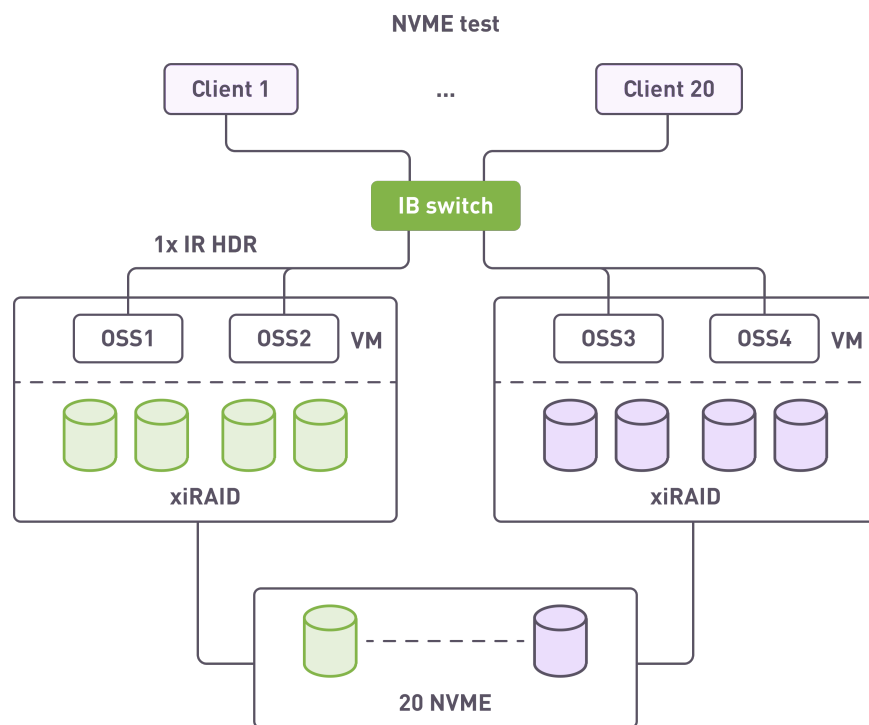


Figure 3. Test Configuration

IOR Test parameters are:

```
ior -F -w -b 8g -t 1m -e
```

After testing the following results were achieved:

| | Random Read | Sequential Read | Sequential Write |
|-------------|--------------|-----------------|------------------|
| Performance | Over 3M IOPS | Over 82 GBps | Over 63 GBps |

CONCLUSION

Through extensive testing, it has been conclusively demonstrated that the combined solution comprising xiSTORE and Celestica's SC4200 platform exhibits remarkable performance capabilities. The synergy between SC4200's robust configuration and xiSTORE's advanced functionality not only delivers high performance but also provides exceptional reliability. The performance demonstrated by Xinnor, combined with the no-single-point-of-failure architecture are an optimal solution to provide fast storage to run complex AI models and optimize the usage of GPUs.

The logo for Xinnor, featuring a stylized 'X' symbol followed by the word 'INNOR' in a bold, uppercase, sans-serif font.

Learn more about us at **xinnor.io**
request@xinnor.io
+972 43 740 203

All Rights Reserved © 2024 Xinnor