Case Study: Enhancing High-Performance Computing at FAU with xiRAID and Lustre

About FAU

Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) is one of Germany's leading research universities, renowned for its excellence in a wide range of fields, including artificial intelligence, machine learning, atomistic simulation, and high-performance computing. In order to support these research activities, FAU hosts the Erlangen National High Performance Computing Center (NHR@FAU), one of nine national HPC centers in Germany. NHR@FAU operates "Alex", one of their two GPU clusters, which is integral to various HPC, ML, and AI applications. Often, these require the handling of millions of small files.

The Alex cluster is designed for advanced research in areas such as molecular dynamics simulations in the field of mRNA vaccine research, studies on the mode of action of enzymes in DNA repair, and applications in machine learning, such as gesture recognition. Alex is equipped with 656 Nvidia GPUs (A100 and A40 Tensor Core models) and AMD Epyc CPUs, providing immense computational power. The A100 nodes are interconnected with a high-speed HDR InfiniBand network, making it one of the most powerful and energy-efficient systems globally, ranking 16th on the Green500 list as of November 2023.

Challenge

In 2023, NHR@FAU procured seven servers, each with 24 NVMe SSDs to build a storage cluster based on the CephFS system, with two-way replication. However, during testing, CephFS underperformed, particularly in terms of write throughput and metadata operations, and CephFS could not cope with the mixed interconnect (HDR InfiniBand for the A100 nodes but only 25 GbE for the A40 nodes). The cluster's high-performance requirements for AI and ML workloads necessitated a more efficient and scalable solution, especially in terms of read/write performance, metadata handling, and capacity optimization. Indeed, machine learning involves reading millions of small files to train the models, which requires high I/O rates. With its double replication architecture, CephFS could not handle this kind of I/O efficiently. For this reason, in 2024 NHR@FAU decided to look at a file system that could better utilize the cluster's NVMe storage and meet the demands of HPC and AI workloads.

Solution

Given the limitation shown by CephFS, MEGWARE, a leading system integrator in Europe, recommended a solution based on Xinnor's xiRAID with Lustre, a parallel distributed file system designed for large-scale cluster computing. Thanks to the unique data path combined with the optimal usage of AVX technology on modern CPU, xiRAID efficiently handles both high sequential and random workloads in both read and write operations, while using very limited system resources.

To validate the benefits of this solution, MEGWARE helped NHR@FAU in deploying xiRAID + Lustre on half of the available storage cluster (four servers) and comparing its performance with the CephFS implementation on the remaining three servers.

Test Architecture

The performance was compared between the following configurations:

xiRAID + Lustre Configuration:

- 4 NVMe servers

- 1x NVIDIA ConnectX-6 100Gbit/s Ethernet
    - NVIDIA A40 nodes connected via 25Gbit/s Ethernet
    - theoretical single NVMe server limit via Ethernet: 12.5GB/s
- 1x NVIDIA ConnectX-6 200Gbit/s HDR InfiniBand
    - used for NVIDIA A100 nodes
    - theoretical single NVMe server limit via IB: 25GB/s
- xiRAID RAID6 (8+2) for Data and RAID1 (1+1) for Metadata
- Lustre 2.15.5

CephFS Configuration:

- 3 NVMe servers

- 1x NVIDIA ConnectX-6 100Gbit/s Ethernet
    - used for NVIDIA A40 and A100 nodes (each connected via 25Gbit/s Ethernet)
    - theoretical single NVMe server limit via Ethernet: 12.5GB/s

- 3-way replication

Testing involved using eight clients to simulate workloads on both the xiRAID + Lustre and CephFS configurations. The setup was not an entirely fair comparison as the solution based on xiRAID had a theoretical network limit of 100GB/s while the CephFS configuration had a theoretical limit of 25GB/s (8x A100 clients each at Gbit/s).

Test Results

The performance results clearly demonstrated the superiority of the xiRAID + Lustre configuration over CephFS:
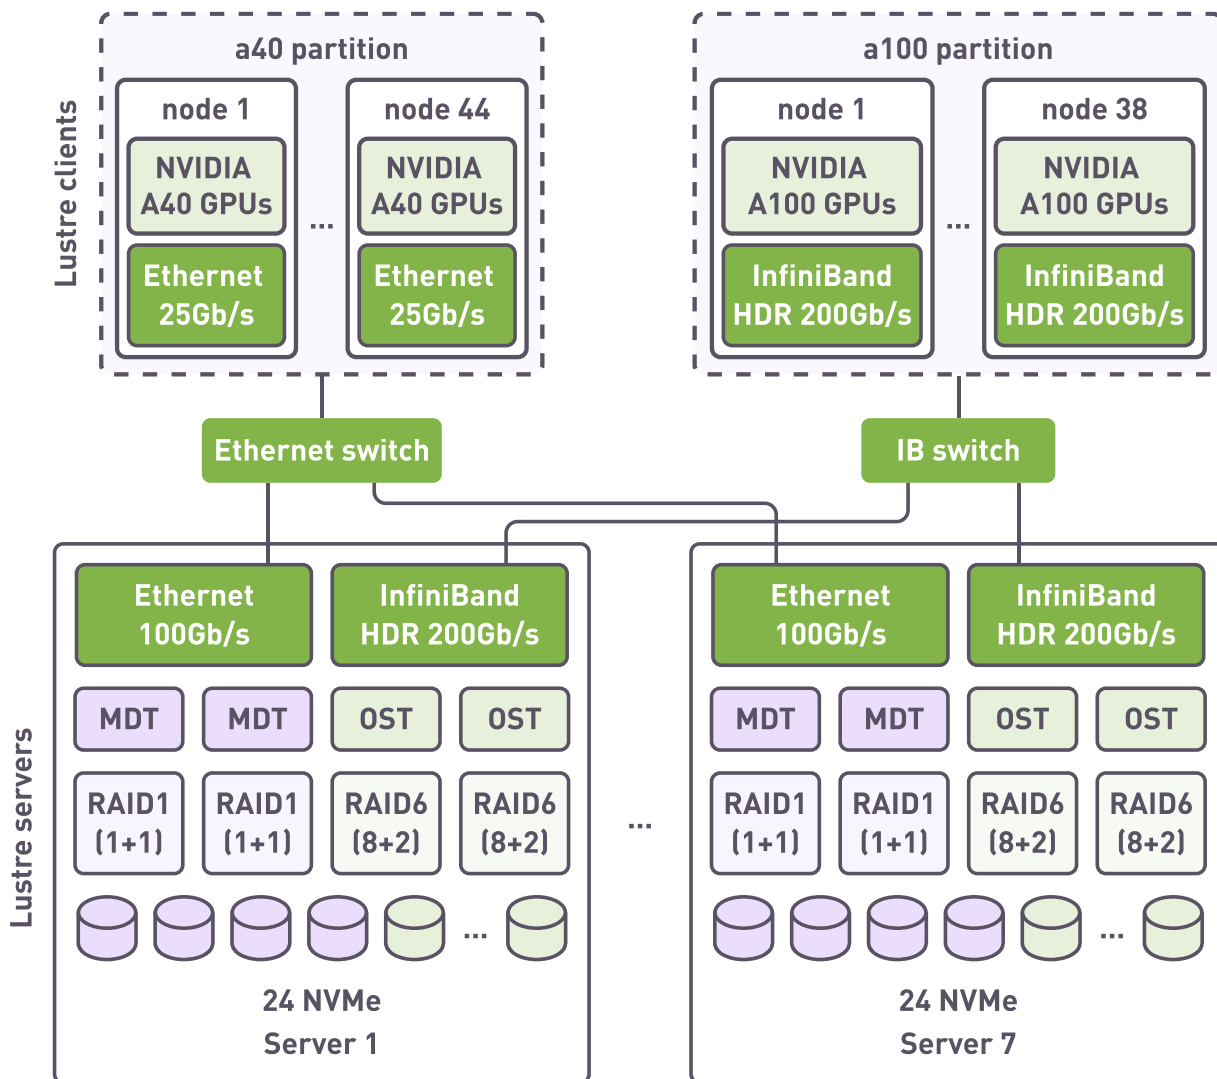
xiRAID + Lustre (4 nodes):

- Write Throughput: **67 GB/s**
- Read Throughput: **90.6 GB/s**

CephFS (3 nodes):

- Write Throughput: **11.3 GB/s**
- Read Throughput: **23.4 GB/s**

The results showed that xiRAID + Lustre achieved 3–4 times higher throughput compared to CephFS, which cannot be justified with the differences in the test setup.

ℵINNOR   MEGWARE

Implementation

To optimize capacity and performance, after validating the benefit of the solution based on xiRAID, the entire cluster was reconfigured to use xiRAID+Lustre with 14 MDTs (1+1) in RAID1 and 14 OSTs in RAID6 (eight drives for data and two for parity), resulting in a usable capacity of approximately 775TB, that is about 30% more than previously available from CephFS because of its two-way replication. To be fair, this Lustre implementation was not protected from the failure of one server node, while CephFS could cope with this situation thanks to the double replication. Nonetheless, high-availability was not a requirement for this installation, as it is used for scratch data.

In case high-availability is needed, Xinnor supports it on xiRAID and Lustre using Storage Bridge Bay servers (servers hosting in the same box two independent compute nodes that share the same drives) or Ebof (Ethernet JBOF) connected to multiple diskless servers.

The configuration implemented at NHR@FAU maximized performance and capacity by leveraging the NVMe storage and the high-speed InfiniBand network infrastructure available on the Alex GPU cluster.

Conclusion

By transitioning to xiRAID + Lustre, NHR@FAU significantly enhanced the performance and scalability of its Alex GPU cluster. The new system provided a substantial boost in read and write throughput, better metadata performance, and nearly double the usable storage capacity compared to CephFS. This transition showcases the power of xiRAID + Lustre solution for high-performance storage solutions in research-intensive environments like NHR@FAU.

ℵINNOR   MEGWARE