



University in Singapore Case Study:

BeeGFS Cluster Storage Solution
powered by xiRAID

XINNOR

 ON DEMAND
systems

Overview

When it comes to innovation, Singapore has always been leading in Asia and the world. Nowadays, with the advent of Artificial Intelligence, it makes no exception. Its universities are leading in AI research projects across diverse fields. For example, in healthcare, researchers are using deep learning for medical imaging analysis to enhance disease diagnosis and predictive analytics to optimize patient treatment plans. In natural language processing, they are developing multilingual models and advanced conversational agents for customer service applications.

To support these and more innovative research, one of the leading universities in Singapore implemented overtime multiple GPUs and Compute systems, including:

- 4 GPU Nodes with Nvidia A40
- 3 GPU Nodes with Nvidia A100
- 22 Compute Nodes
- 1 Compute High Memory Node

To keep these systems busy with data the university tasked On Demand System (ODS), a leading HPC system integrator in Singapore, to find the optimal storage solution.

Requirements

The main requirements from the university were:

1. Fast access to data: Facilitate swift data retrieval for GPU-accelerated applications and high-performance computing tasks.
2. Data protection: Ensure data integrity and availability.
3. Ease of deployment and management: Implement a solution that simplifies ongoing maintenance efforts.
4. Cost-optimization: The solution had to meet stringent university budget

5. Scalability: The solution should be provisioned to handle future growth in term of clients and performance requirements

Deployment overview

To meet all the requirement, On Demand System (ODS) used On Demand PFS and architected a solution based on 2 server nodes based on NVMe drives with BeeGFS Parallel File System and protected by Xinnor's xiRAID:

Infrastructure Setup:

On Demand PFS - 2-Node BeeGFS Cluster each with:

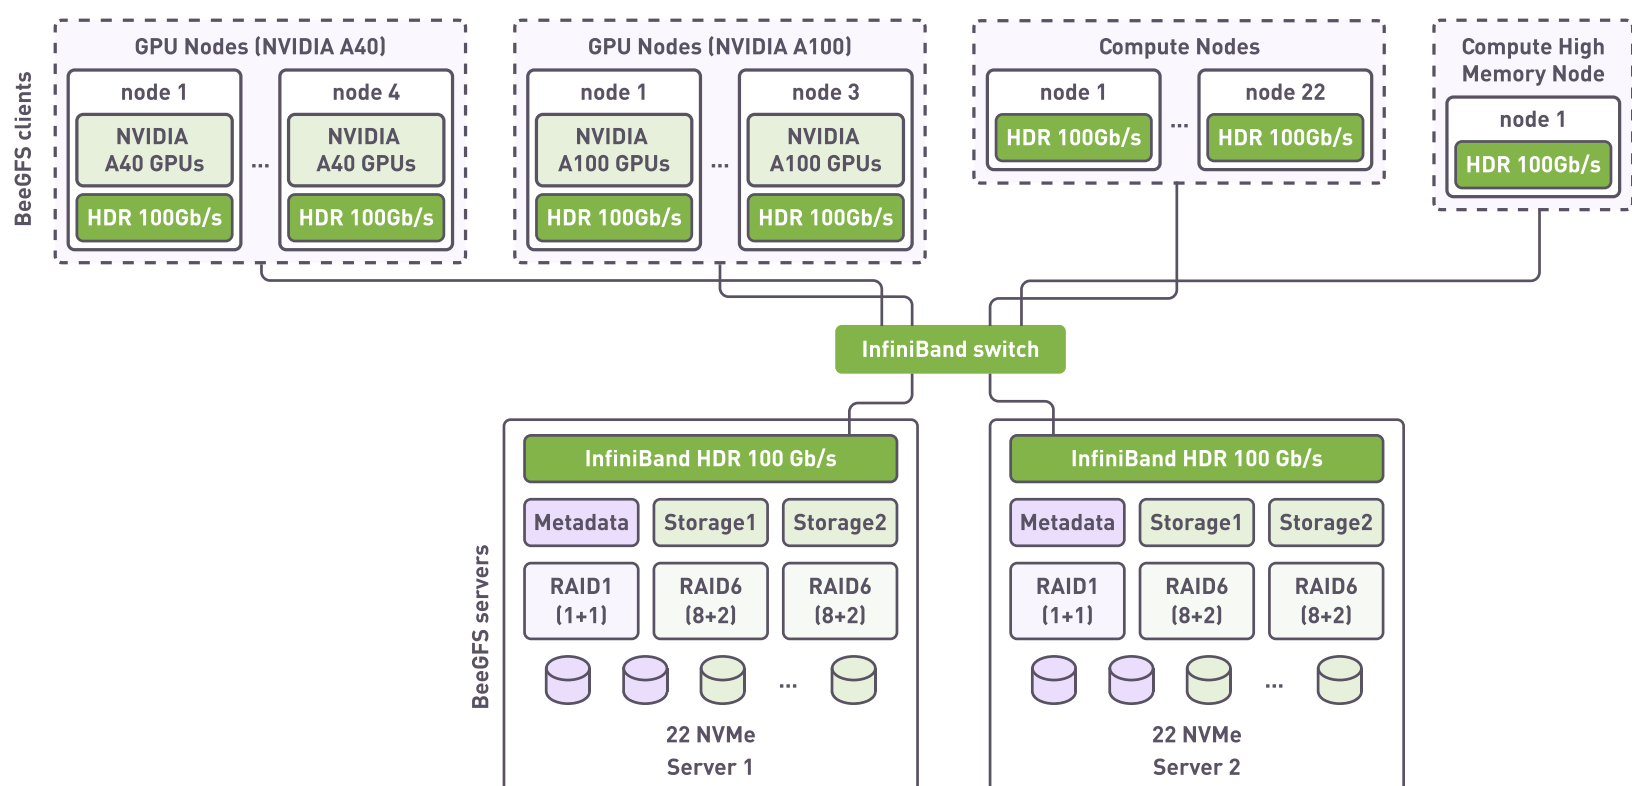
- 2x AMD EPYC™ 7763 (Milan:64C,280W,2.45G)
- 16x 32GB DDR4 3200
- 2x 480GB Enterprise SATA SSD
- 24x 15.36 TB Gen4 NVMe SSD
- Single Port HDR100Gbs Adapter

To protect the data in case of one or multiple drive failure and deliver maximum performance to the file system, On Demand System (ODS) selected to use Xinnor's xiRAID Classic.

xiRAID Classic is an advanced software RAID engine designed to handle the high level of parallelism of NVMe SSD. It leverages optimal utilization of Advance Vector Extension (AVX) technology of modern CPUs to quickly run the parity checksums and combines it with its proprietary lockless data path, the ability to dynamically allocate the stripe to the available CPU core, to evenly distribute the load across all the available CPU cores, minimizing overall CPU utilization and maximizing the performance.

xiRAID was deployed in each Server as follows:

- 2x drives in RAID1 for MetaData
- 10x drives in RAID 6 (8+2) for OST1
- 10x drives in RAID 6 (8+2) for OST2
- 2x drives as Global Spares



RAW capacity: 737TB

Net user capacity: 492TB

This configuration was selected to maximize support high IOPS and throughput by providing more than sufficient capacity to handle the metadata and almost 500TB of storage to the users. The architecture is designed to survive multiple drive failures and by having spare drives pre-installed in the servers it avoids downtime for maintenance in case one or more drives need to be replaced.

xiRAID Benefits

xiRAID is engineered to optimize data protection and access speeds within storage environments. In this deployment, xiRAID played a critical role in:

1. **Data Protection:** xiRAID employs advanced algorithms to provide redundancy and fault tolerance, ensuring that data remains safe even in the event of multiple drive failures. The implementation in RAID 6 allows to withstand the concurrent failure of 2 drives in each OST, giving peace of mind in term of data accessibility.
2. **Performance Acceleration:** the cluster accelerated by xiRAID was able to saturate the 2x100Gb Infiniband ports, with measured sequential read performance of 24.7GB/s. This speed allows for faster data retrieval rates, particularly beneficial for GPU-accelerated compute tasks.
3. **Ease of deployment:** xiRAID Classic exposes a block device within the linux kernel, for easy integration with the file system. Mounting BeeGFS over xiRAID block device was straightforward requiring minimal manual intervention, which accelerated overall deployment time.
4. **Scalability:** in case in the future more performance is required, it will be sufficient to add more InfiniBand cards or increase their bandwidth. Indeed, the read performance measured within the cluster exceeds 215GB/s.
5. **Cost optimization:** xiRAID doesn't require any hardware. There's no need to install a x16 PCIe card dedicated to run the RAID, so no need to waste the PCIe slot and related PCIe lanes that can be better reserved for future network expansion.

Conclusion

The deployment of an On Demand PFS storage solution, accelerated and protected by xiRAID, at this leading university in Singapore exemplifies the successful integration of advanced storage technology in an academic environment. The solution not only met the immediate goals of providing fast data access and enhanced reliability, but it is also ready to support future research expansions. xiRAID played a pivotal role in ensuring ease of deployment, robust data protection, and impressive performance, thereby empowering researchers and engineers at the university to push the boundaries of scientific inquiry and technology development.

XINNOR

Learn more about us at **xinnor.io**

request@xinnor.io

+972 43 740 203

All Rights Reserved © 2024 Xinnor