

Data Archive challenges for sPhenix 2025

Thursday, 20 March 2025 16:20 (22 minutes)

With the Run2025 for sPhenix, it comes the higher data throughput and data volume requirements.

The sustained data throughput required for sPhenix2025 is 20GB/sec. Once started in mid-April, this sustained data stream will be steadily constant with no breaks through December. The projected data volume is 200PB.

In order to meet these data throughput and volume requirement, we must rebuild our data storage archive systems...

1. Data movers.

a. Replace data movers with new PCIe-4 architecture and FC HBA's. Increase the number of data movers to nine servers from four.

2. Ethernet connections.

a. Replace NIC adapters with dual 100GbE PCIe-4.

3. Disk arrays.

a. With the NAND options too expensive, we decide to stay with HDD spindles

b. Increase disk arrays to nine units from three.

c. Increase each disk array multipath connections to four channels from two.

4. OS I/O related parameter tunings

5. Upgrade HPSS software to resolve occasional hanging processes on movers

6. Purchase additional two 9-frame IBM tape libraries

a. To lower the costs, we decided to go for LTO9 technology, instead of enterprise tape technologies.

b. Two 9-frame tape libraries are needed to meet the data volume requirements

c. To make the most usage of all tape drives, we decide to evenly strip injected data across four tape libraries with 100 LTO9 drives (25 drive in each library).

7. Benchmark testing

a. We ran benchmark testing on existing 3 disk arrays with 4 movers and 56 tape drives

b. With ongoing injection of 8.5GB/sec, the concurrent migration to tape is 11GB/sec.

c. With ongoing injection of 12GB/sec, the concurrent migration to tape is 9.9GB/sec.

d. The CPU usage on each mover is at 80%.

e. This benchmarking numbers give us confidence that our new configurations will comfortably meet the 20GB/sec sustained throughput.

Conclusion:

With the new designs and fine-tunings, we have identified a solution to the sPhenix2025 data requirements.

Presenter: CHOU, Tim (Brookhaven National Laboratory)

Session Classification: Data Management & Big Data

Track Classification: Track 6: Data Management & Big Data