AI & Machine Learning

Before the age of AI, typical datacenter read/write ratios were 4:1. With AI, it's 5000:1*. AI leaders are moving away from HDDs to QLC SATA SSDs. The Micron® 5210 delivers.

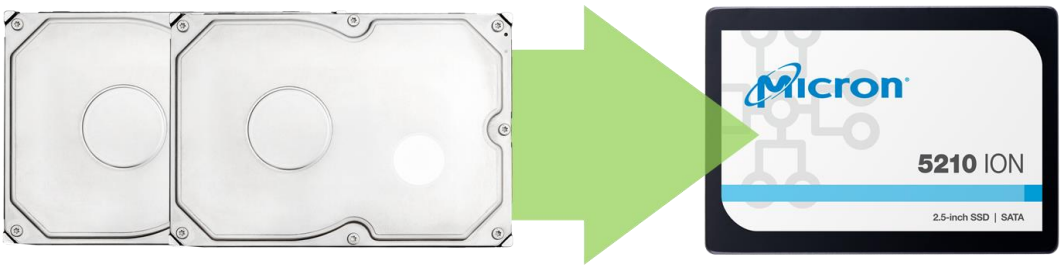

Micron 5210 QLC SSD vs. Legacy HDDs

TensorFlow Dataset Size	HDD	5210	Improvement
23GB	10 min	1 min	9 mins
230GB	90 min	11 min	1.3 hours
1150GB	7.5 hours	1 hour	6.5 hours
2300GB	15 hours	2 hours	13 hours

Colfax Research test report data: https://colfaxresearch.com/micron-qlc-ssd/

5210 Advantage 13 Hours Saved!

How much is your time worth?

Typical AI & ML Workload

*Source: EnterpriseStorageforum.com: "Data Storage, AI, and IO Patterns http://www.enterprisestorageforum.com/print/storage-technology/data-storage-ai-and-io-patterns.html

Storage access pattern: random reads & writes
Storage IO size: 8K to 1GB (or more)
Read/write ratio: 99% read / 1% write *
How the workload works:

- 50TB+ datasets moved from storage to AI servers (training)
- Once training completes, refined datasets fed back to storage
- Storage can be major bottleneck: dataset sizes dwarf system memory
- Cost of idle GPUs (waiting on storage) often dwarfs QLC SSD cost

Ready to learn more? Read Colfax's in-depth ML research