

HDD vs SSD



Ken Smith

Background

2009 Started PSI ingest with HDD storage

2013 Switched PSI ingest to SSD & gained 10x ingestion speed

2015 Tested SSDs in RAID0. ATLAS ingestion begins.

2019 Outgrew single SSD - now operating all servers in RAID0

2020 January - tested ingestion of Lasair data with SSDs but only marginal improvements

Direct I/O Tests (Bonnie++) - HDD vs SSD

	random seeks (ms)	sequential output (ms)	sequential input (ms)
lasair-dev-db HDD	52	1225	1215
psdb3 HDD	117	326	515
lasair-dev-db NVME-SSD	3	40	38
psdb2 PCI-SSD	7	126	11
psdb3 SATA-SSD	4	227	3

Random seeks SSD is 17 times faster
sequential input & output is 30 times faster.

Database Speed Test (part I - Ingestion)

Bypasses the Kafka/Avro packet parsing

Just reads 55,331,759 Gaia DR2 records from a CSV file and forces them into the database

2 databases - one stored on HDD, the other symlinked to SSD

Multithreaded (multiprocessing) so uses max cores available

lasair-dev-db = 14 cores (64GB memory)

Disk type	Elapsed ingest time (minutes)	Ingestion speed (rows/sec)
HDD	168	5,489
SSD	165	5,589

Surprisingly little improvement when writing to SSD

All these writes are sequential

Implies the MySQL processes are the bottleneck, not the disk type

Database Speed Test (part 2 - Random Cone Searching)

Randomly shuffle the input data before running the cone searches.

RESET CACHE command in MySQL DOES NOT WORK!

10,000, 100,000 and 1 million cone searches of real objects from the 55 million gaia source table

Testing not only the lookup, but of the database ability to return the results back to us

(i.e. no blank results)

Search for the nearest object so for a cone search of 10,000 random objects, I should get 10,000 objects back from the database.

Each search done 5 times.

Number of Cone Searches	HDD (s)	SSD (s)
10,000	44	7
100,000	365 *first run only	68
1,000,000	1682 *first run only	676 *first run only

Conclusions & Further work

We can get up to 6x improvement in speed by switching to SSD - but this makes deployment harder

We know that the database can deal with 5,000 inserts / sec - SSD or HDD

Resolve speed issues in parsing Kafka