QLC Considerations in the Age of AI

Dave Verburg & Alex Scrabeck

IBM

Challenge of Data Explosion

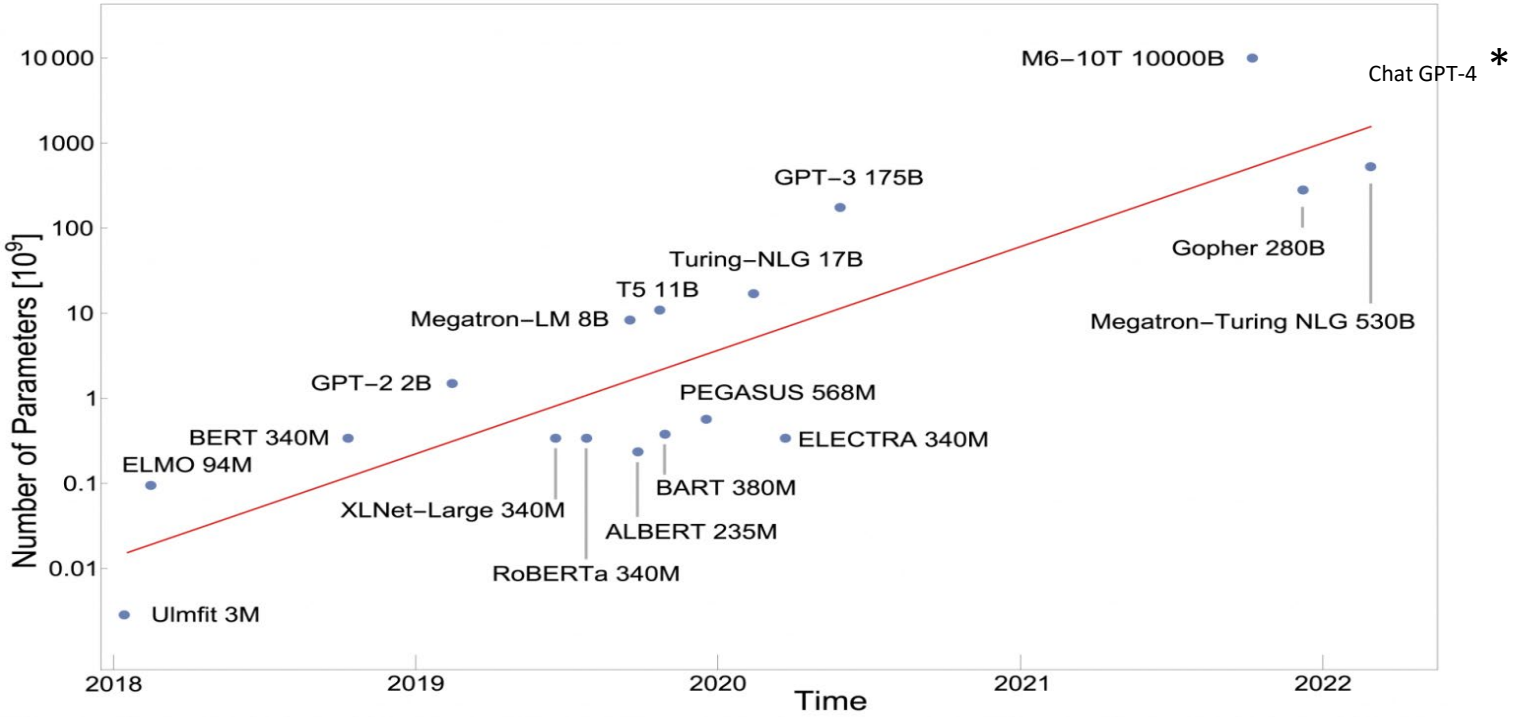

Source: Contact Software Blog

AI Model Parameters growing exponentially

- Real time AI workloads require ingestion of large amount of data
- CPU performance not growing as fast as demand
- Data movement is a key concern

Leading to …

- Move compute closer to the source
- Processing on the edge
- SSDs assigned to specific purposes

QLC Considerations

- **Performance**
  - SSD Clear winner
- **Cost**
  - QLC better price than TLC; still lagging HDD
  - What about TCO?
- **Endurance**
  - Is HDD the winner here?
- **Future**
  - Our prediction

Storage Futures in the age of AI

- **QLC SSD gaining traction**
  - Higher Capacity at lower costs (and as a result lower power per TB)
    - Samsung, Solidigm shipping 61 TB SSD, Samsung “envisions” 120 TB class SSD
    - SK Hynix is working on a 300 TB solid-state drive*
    - IBM FCM currently 38.4 TB usable, 115.2 TB effective with compression*
    - Max capacity of HDD ~30TB, with larger form factor
  - QLC NAND has 2^4 voltage states vs 2^3 for TLC NAND
    - TLC faster, especially for writes, but for AI data lakes, need mostly reads
    - TLC endurance better (for writes), but most QLC better rated endurance than HDD
    - For AI training, need to feed the beast (quickly); training involves a lot of reads
      - QLC SSD much better performance than HDD

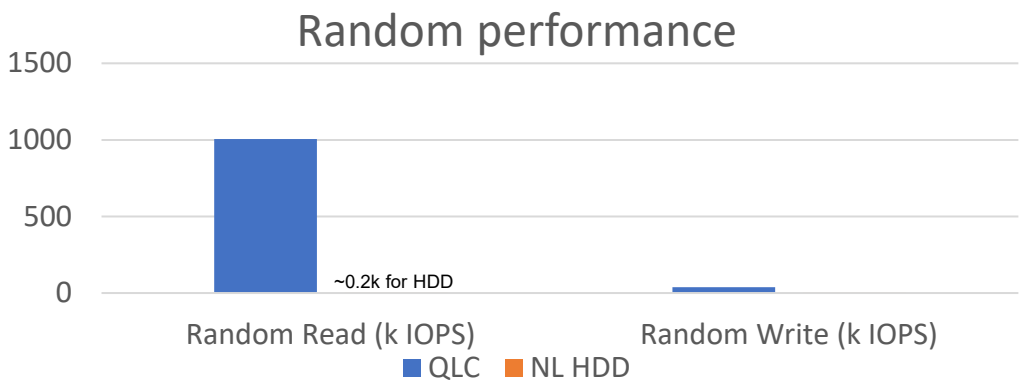

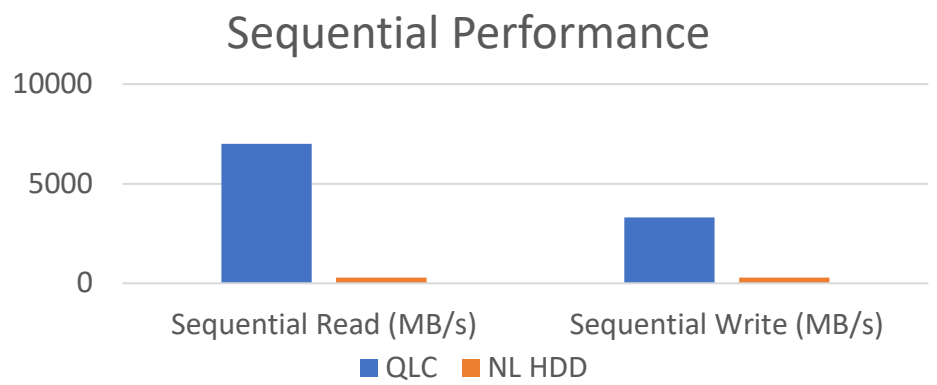

*Sources: anandtech.com, tomshardware.com, www.ibm.com, supplier specs

Cost Comparison

- Market data shows QLC SSDs narrowing the gap between TLC SSDs and HDDs.
- HDDs remain the most cost-effective in $/TB.
- The cost gap shrinks when considering total cost of ownership (TCO), as shown on the next page

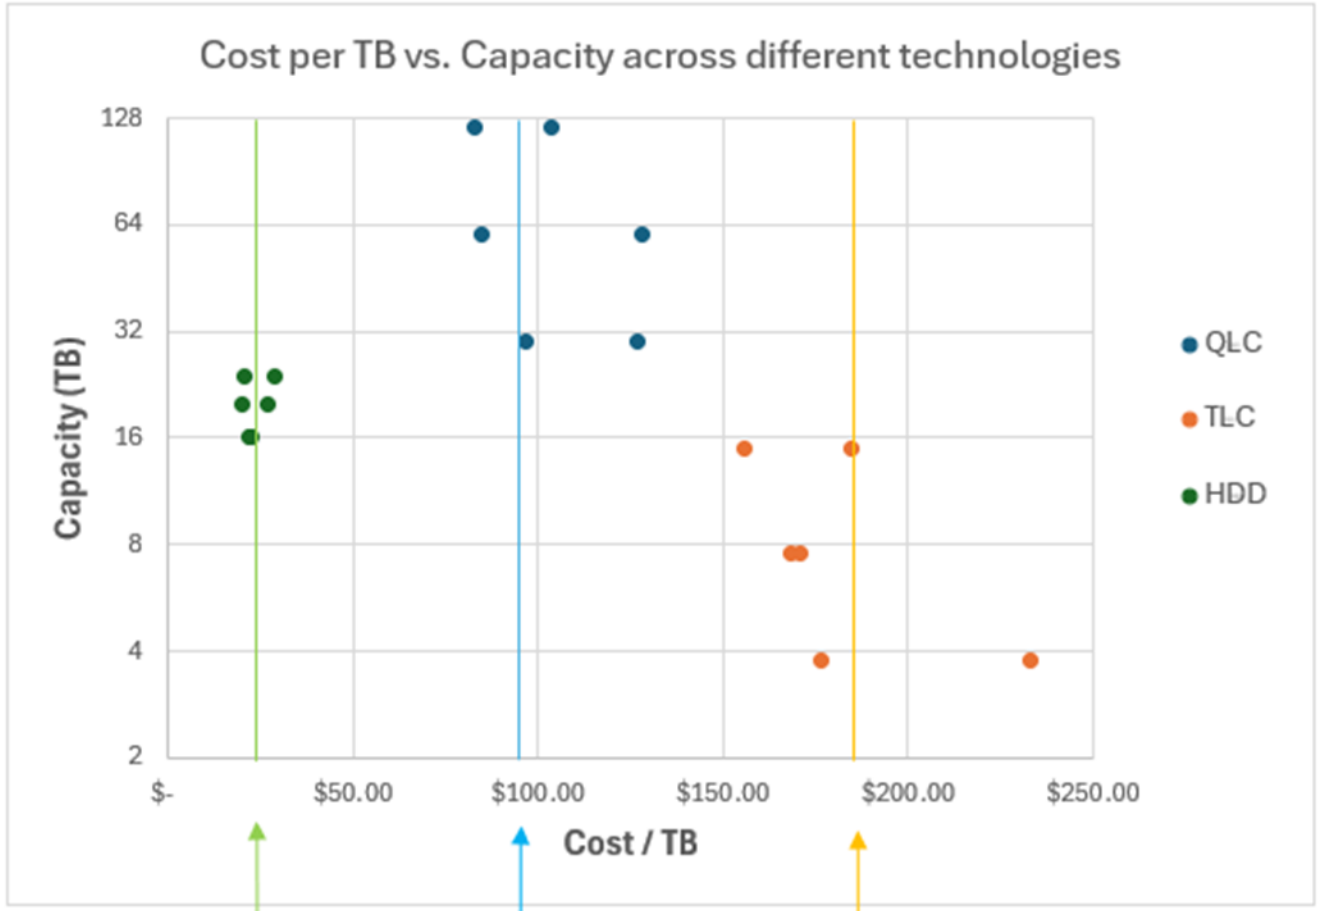

Average Cost/TB HDD	$	23.53
Average Cost/TB QLC	$	103.73
Average Cost/TB TLC	$	181.54

Will QLC have a cheaper TCO than HDD?

- Assumes more replication needed for HDD, less utilization
- QLC prices need to fall in order to get to cost cross-over for raw TB
- QLC wins with operational expenses, HDD wins with acquisition cost

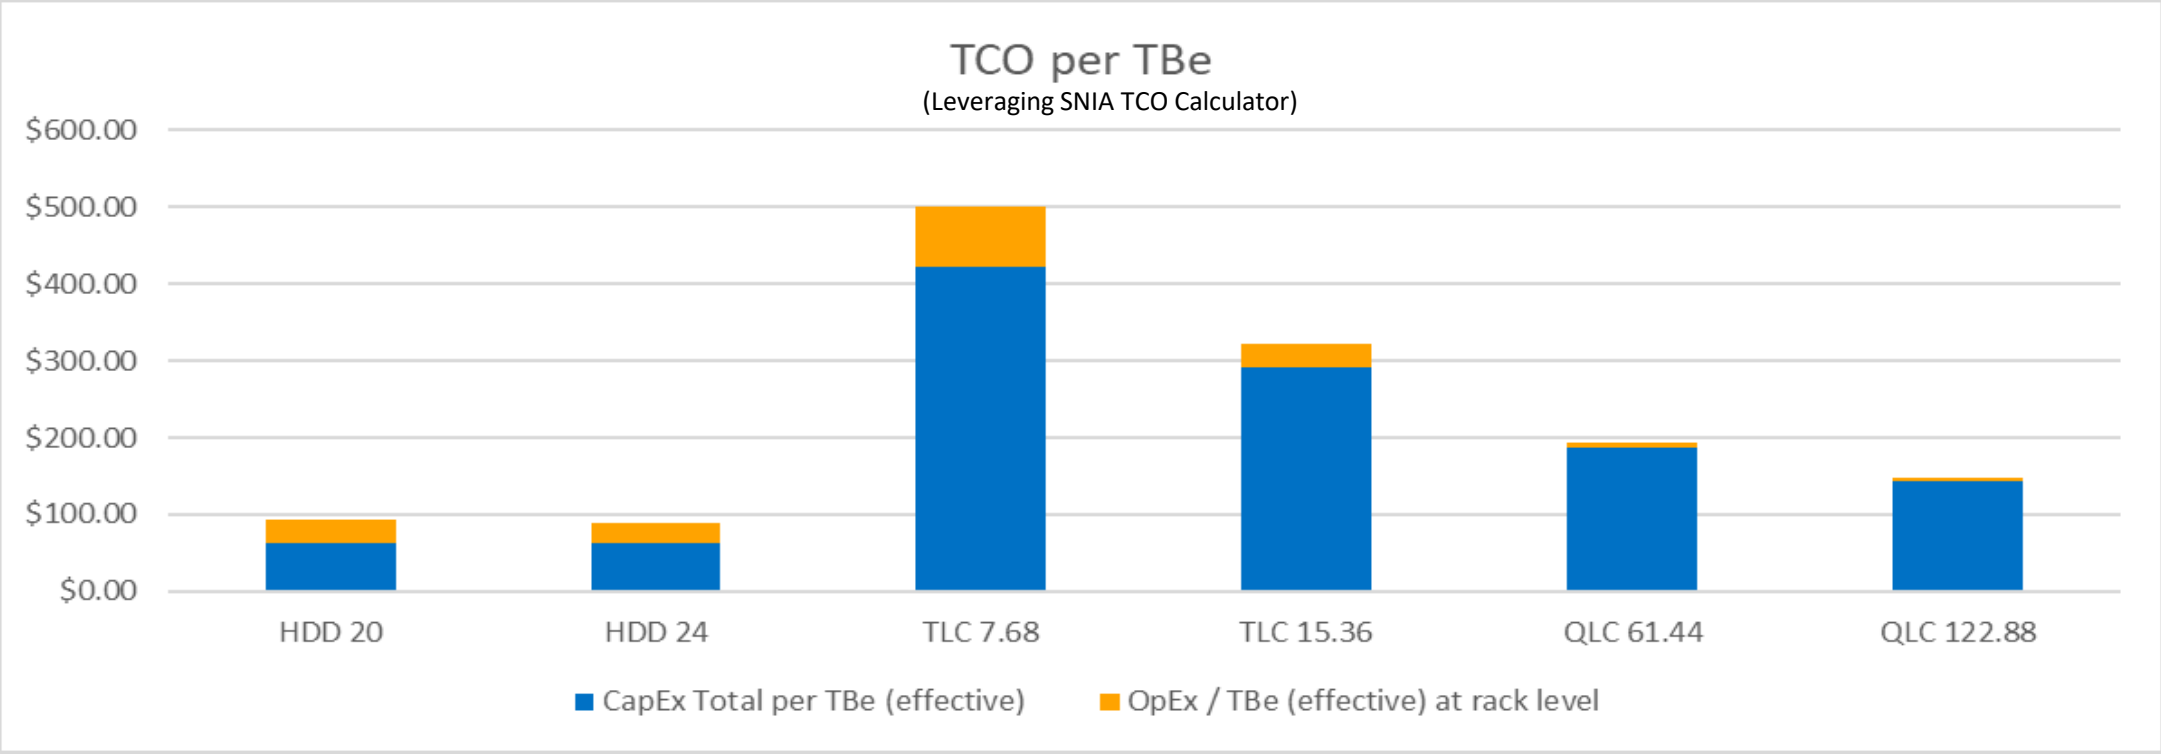

Endurance

- Common thinking … HDD have better endurance than QLC SSD
  - According to Google AI:
    - Higher Write Endurance: HDDs can generally handle more write cycles before experiencing issues compared to QLC SSDs ….
- Actual data
  - NL Limit currently 550 TB per year
    - Equates to 0.05 DWPD if will still be the limit when reach 30 TB HDD
    - Both Read and Write limitations; Perhaps 0.02 DWPD limit for Writes
    - 1/10th the endurance of typical QLC SSD, ~0.2 DWPD
      - Lower performance will likely lead to less writes for HDD

Sustainability Calculations

- Key factors
  - CO_2 used for creation
    - SSD has higher “embodied CO2” on creation
    - CO2e creation per TB is improving!
    - Energy taken from SNIA calculator
  - CO_2 used for operation
    - SSD is much lower; expect to improve
    - Could improve with more green energy
  - Cooling
    - SSDs run reliably at higher temperatures

CO2e creation for HDD vs SSD, 5-year life

Storage	Energy (KWh)	OPEX CO_2e (Kg)	CAPEX CO_2e (Kg)	Total CO_2e (Kg)
HDD/TB ~2021	183.9	79.6	20	99.6
SSD/TB ~2021	56.9	24.6	160	184
HDD/TB current	46.1	17.1	1.1+2 (+ drawers)	20.2
SSD/TB current	19.8	7.3	24.0	31.3

Source: https://futurumgroup.com/insights/are-ssds-really-more-sustainable-than-hdds/, Seagate website

Storage Futures in the age of AI

- **QLC SSD gaining traction**
  - SSD reliability better than HDD
    - HDDs are also more vulnerable to vibration/heat
    - Need more data for QLC … more voltage levels, higher bit error rate
  - Better performance, better Op Ex sustainability
  - Retention may be a concern
- **Near Line HDD will still have a place**
  - Crossover isn't happening as soon as predicted
    - 2021 prediction 2026; 2023 prediction 2029
  - Tailor-fit application with mixed media type architecture
  - Key building blocks in AI data cycle – content archival for future training*
  - Lower cost of acquisition, lower Cap Ex Sustainability (Scope 3 CO_2)
  - Will be difficult to meet all industry needs with SSD
    - SK Hynix new DRAM fab in Korea will cost $3.86 billion, take 1.5+ years*
  - Expect SSD % of total EB to take off in a few years
  - A hybrid approach may be good for cost sensitive applications

*Sources: tomshardware.com, www.ibm.com, blocksandfiles.com, Reddit (Darkhoarder)