

Tape Cold Storage at Meta

Qingzhi Peng

In collaboration with
James Munninghoff
Jon Holdman
Madhavan Ravi
Shankar Yanamandram
and Meta Tape Archive Storage Team

June 24, 2022

@ 12th Annual Fuji Global IT Executive Summit

OUTLINE

- What do we care about in tape solution?
- Meta archival and tape
- Tape TCO modelling
- Summary

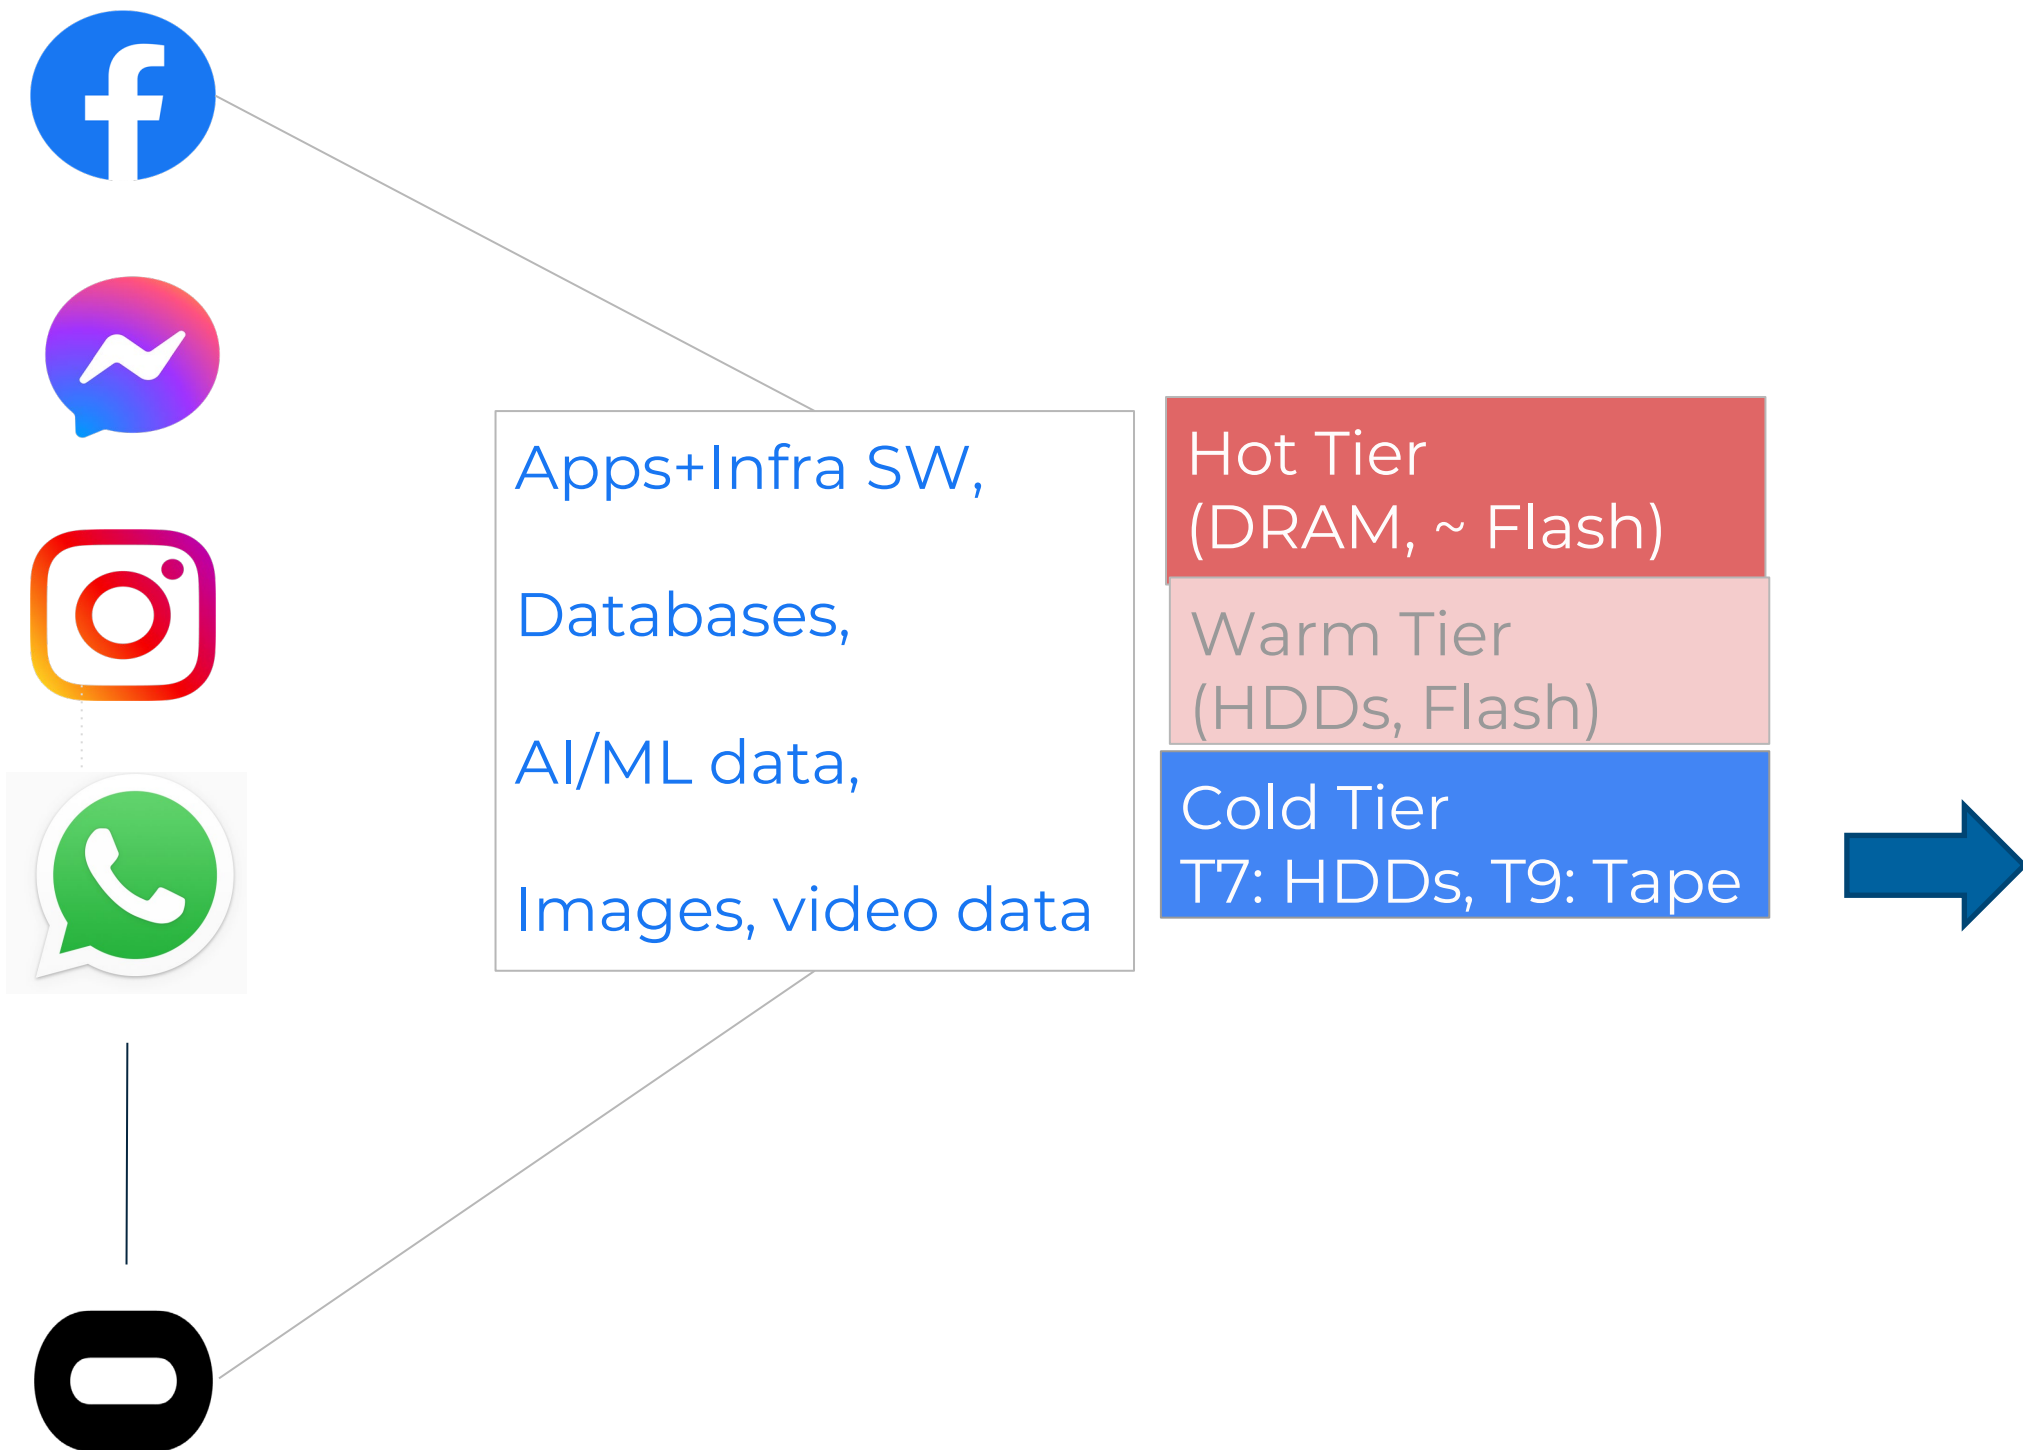
Tape in Hyperscaler

- Tape storage in hyperscale is relatively new and still evolving
 - We are learning as we develop the tape solution and increase in scale
- Need collective effort between hyperscaler customers and supply partners to shape the tape use case to its strength
 - Example 1: Olympic athletes 100 meter dasher and marathon runner
Both are great athletes, but they will not win a medal if not in the right event
 - Example 2: HDD vs SSD

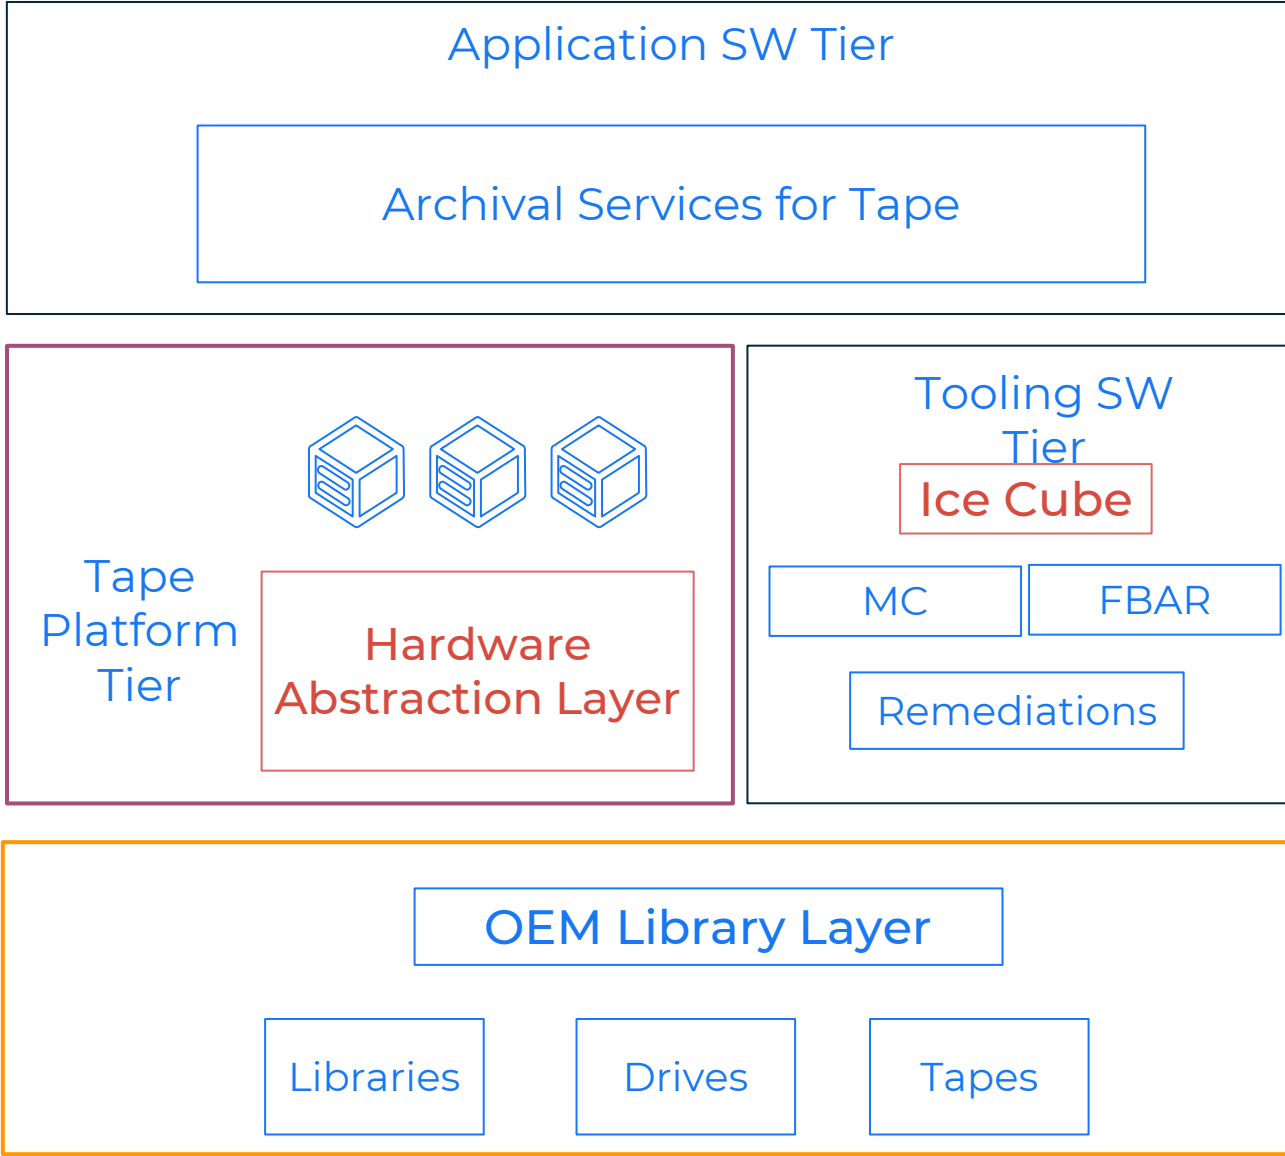
What Do We Care About Tape Solution?

- Normalized TCO for the storage solution
 - Technology growth: drive and media recording density and capacity growth
 - Costs of data IOs
- Serviceability and availability of the tape system
- Library fill time - days required to write storage capacity of library
 - Rewrite Rate (Touch Rate): how fast the full library capacity can be written
- Power - power draw per unit of logical storage capacity
- Density - Storage capacity per unit floor space
 - Capacity per library rack
 - Size per rack
- Supply and ecosystems

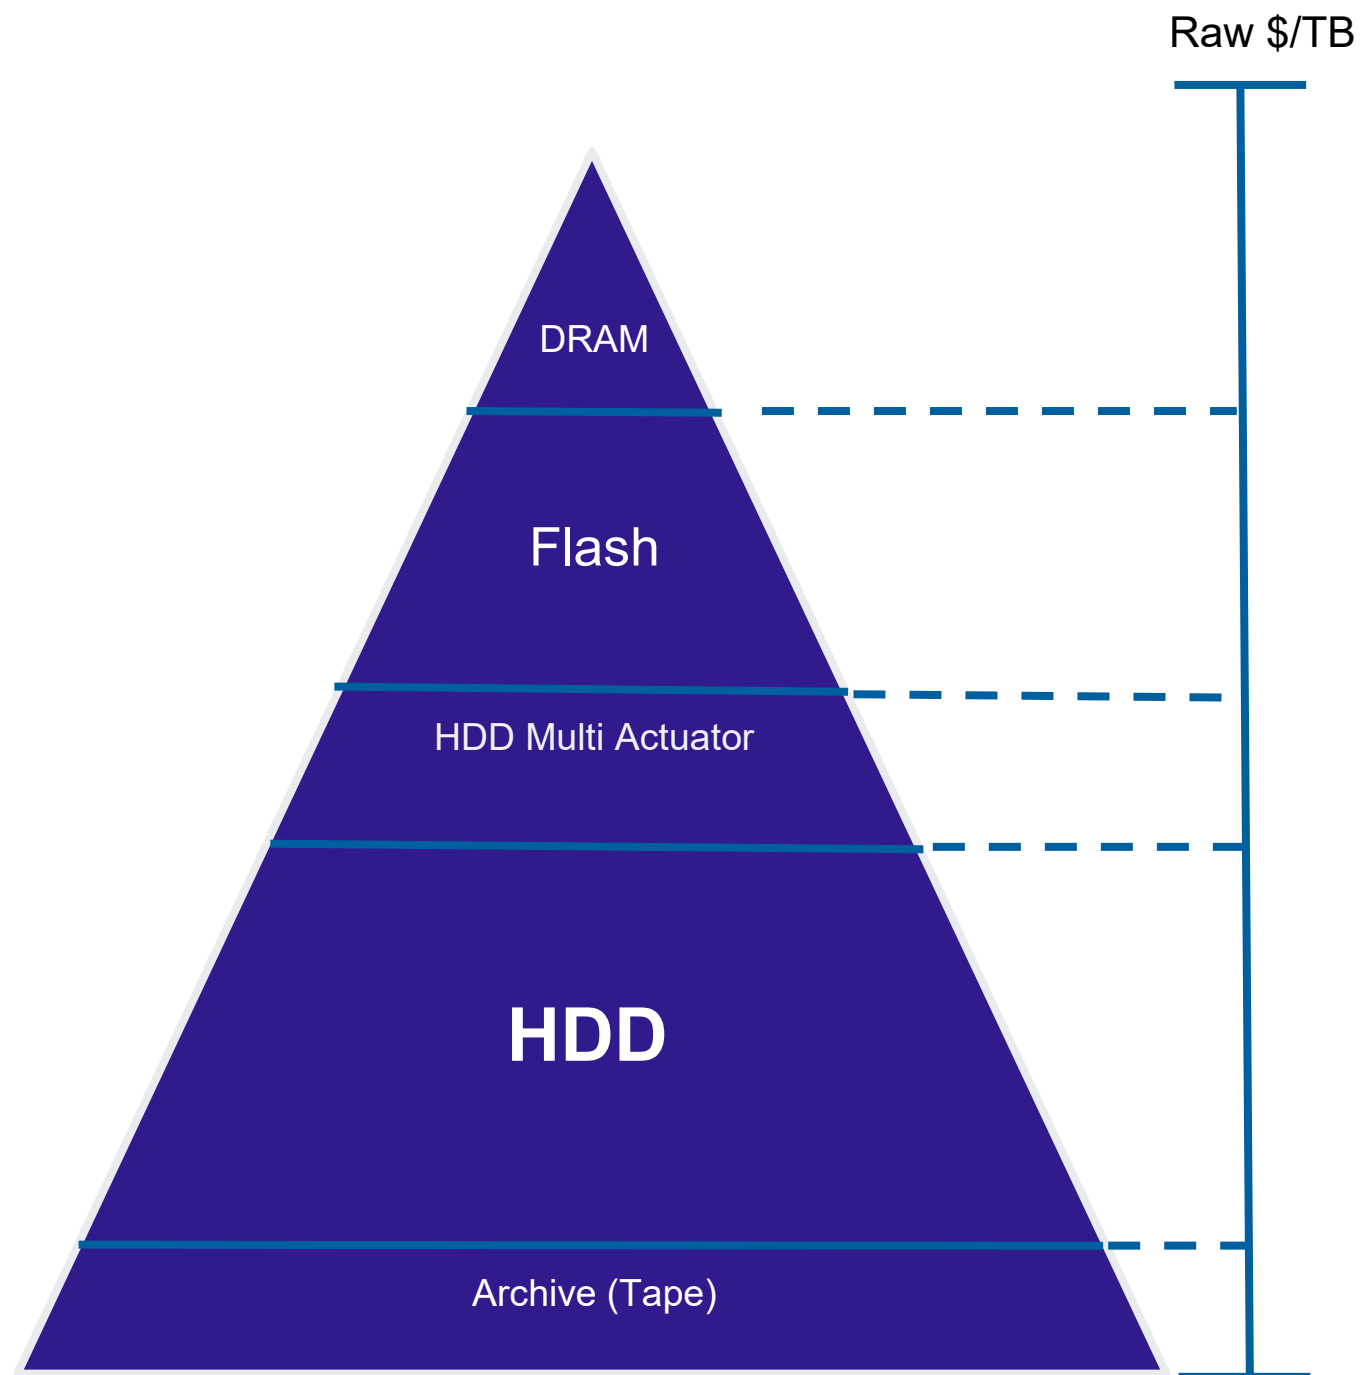
Tape Workload and System



FB Tape Archival Stack



Storage Cost Pyramid



- Tape is placed (and assumed) **STRAIGHT** at the lowest cost tier in the storage pyramid.
- Understand and demystify some of the tape TCO to better utilize and optimize tape
 - Components and drivers for the tape TCO

Tape Library System TCO Analysis Model

$$\text{Total Cost} = \text{CapEx}_{\text{system}} + \text{OpEx}$$

$$\text{CapEx}_{\text{system}} = n \cdot P_{\text{cart}} + C_0 + \frac{n}{r} \cdot P_{\text{drive}} + \frac{n}{r} \cdot (DR) \cdot C_1$$

$$\text{TCO}_{\text{CapEx}} = \text{CapEx \$ per physical TB} = \frac{\text{CapEx}_{\text{system}}}{\text{Total Capacity}} = \frac{\text{CapEx}_{\text{system}}}{(n \cdot S)}$$

$$r = \frac{(DR)}{\beta \cdot (TR) \cdot S} \quad \beta = 0.0317 \frac{\text{MB/s}}{\text{TB} \cdot \text{Yr}}$$

$$\text{TCO}_{\text{CapEx}} = \frac{P_{\text{cart}}}{S} + \frac{C_0}{(n \cdot S)} + (TR) \cdot \beta \cdot \left(\frac{P_{\text{drive}}}{(DR)} + C_1 \right)$$

$$\text{TCO}_{\text{OpEx}} = \frac{E_0}{(n \cdot S)} + (TR) \cdot \beta \cdot \left(\frac{E_{\text{drive}}}{(DR)} + E_1 \right)$$

$$\text{TCO} = \frac{P_{\text{cart}}}{S} + \frac{C_0 + E_0}{(n \cdot S)} + (TR) \cdot \beta \cdot \left(\frac{P_{\text{drive}} + E_{\text{drive}}}{(DR)} + (C_1 + E_1) \right)$$

Cartridge
\$/TB

Fixed TCO
from frame

Data IO costs

C_0 : Fixed cost/overhead including frames and racks

n : # of tape cartridges

S : Size of each tape cartridge (in TB)

P_{cart} : Cost of each tape cartridge

P_{drive} : Cost of each tape drive

r : Library cartridge to drive ratio

DR : Drive IO rate

TR : Library touch rate (# of rewrites per year)

C_1 : Normalized (to total drive IO) cost of data staging, network, and head nodes

E_0 : operation cost (e.g. power and facility) for frames and racks

E_1 : Normalized (to total drive IO) operation cost of data staging, network, and head nodes



- Encoding ratio (Physical → Logical)
- Hardware life cycle



Logical Byte TCO over Lifecycle

Tape System TCO Components

- Fixed frame/overhead costs and expenses
 - Higher capacity will help reduce the % of frame/overhead
- Cartridge costs
 - Relatively simple: cartridge price / capacity
 - Often viewed as the cost of tape storage but it is only the partial cost
- Data movement (IO) costs
 - Scaled by the IO throughput rate
 - Driven by the costs and expenses of drive, network, and data staging

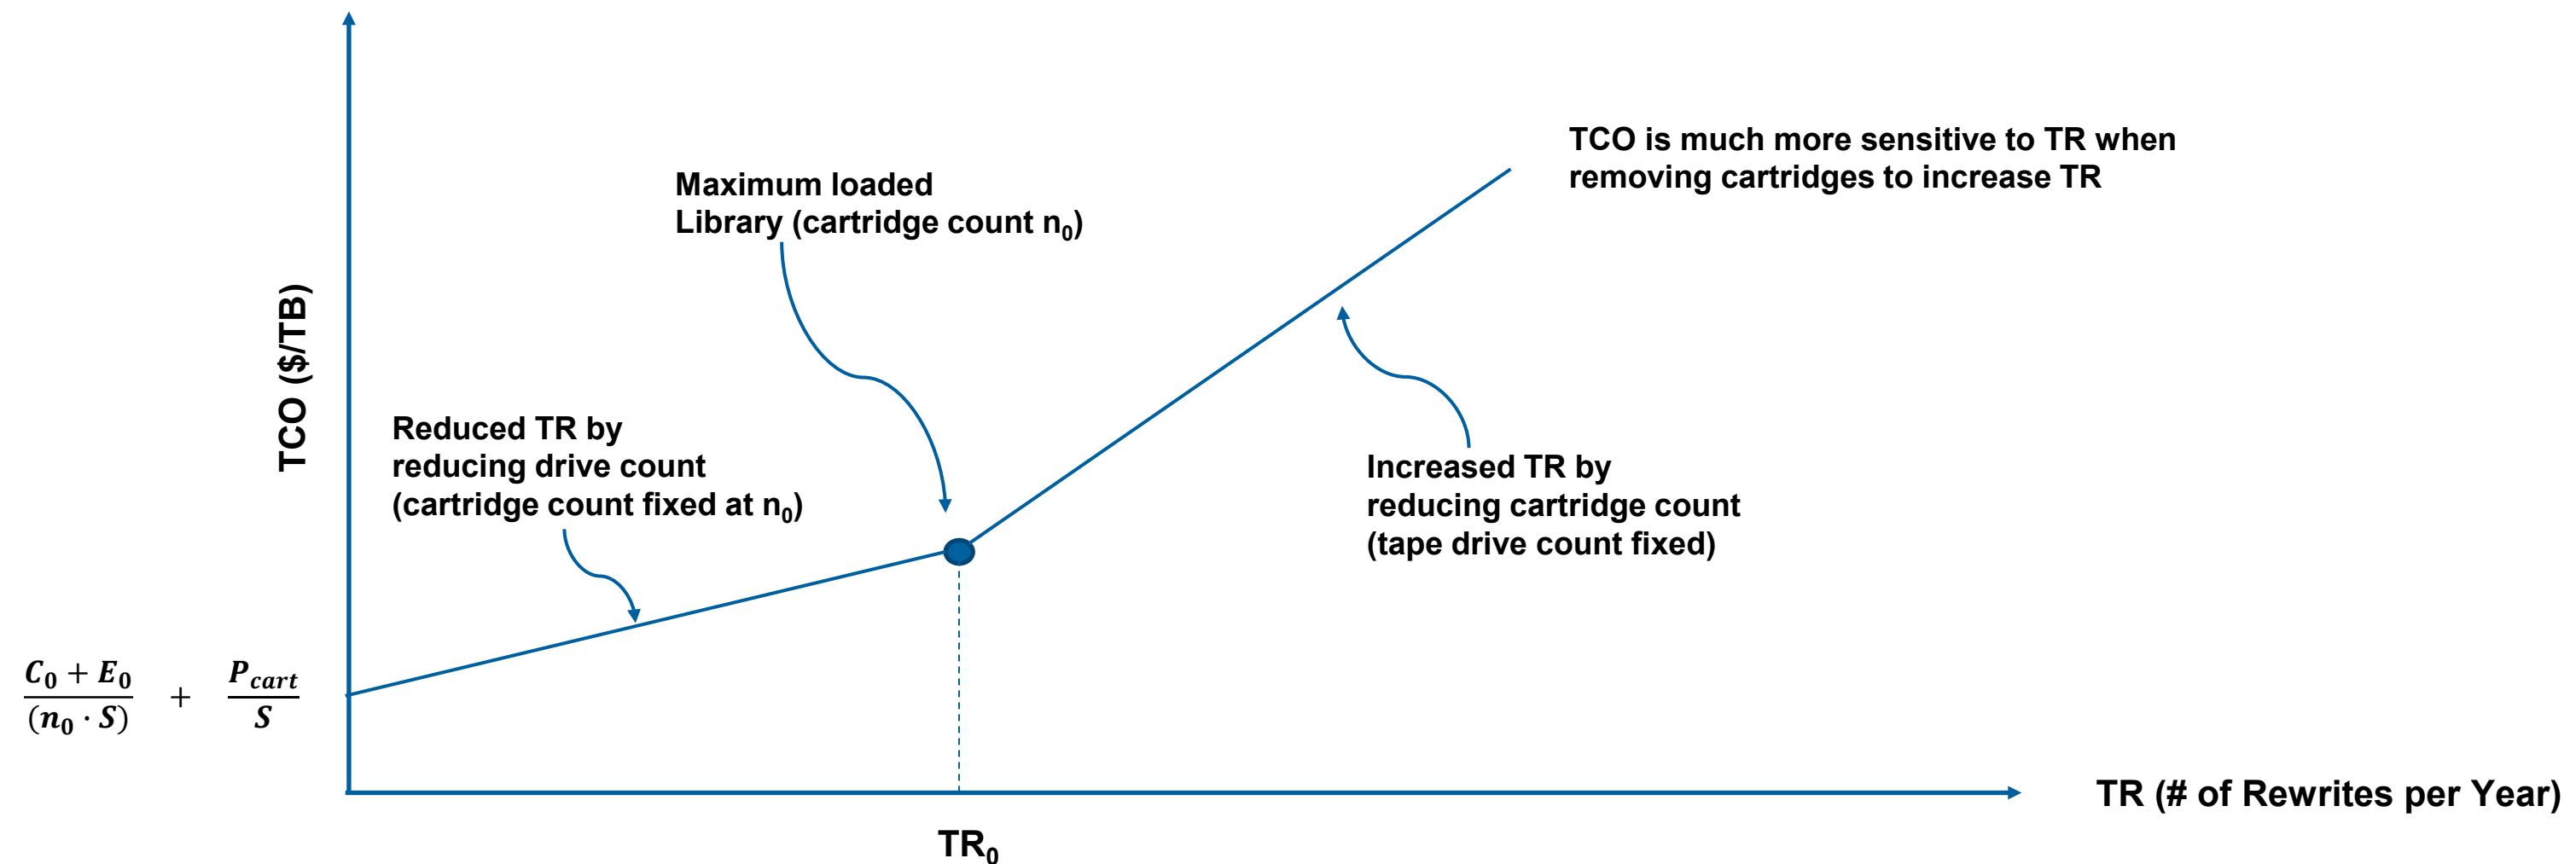
A Closer Look at the IO Impacts on TCO

Case 1: Adjusting TR by drive count while fixing cartridge count

$$TCO = \frac{P_{cart}}{S} + \frac{C_0 + E_0}{(n_0 \cdot S)} + (TR) \cdot \beta \cdot \left(\frac{P_{drive} + E_{drive}}{(DR)} + (C_1 + E_1) \right)$$

Case 2 : Increasing TR by lowering library cartridge count while fixing the total drive count

$$TCO = \frac{P_{cart}}{S} + TR \cdot \left[\frac{C_0 + E_0}{(n_0 \cdot S) \cdot TR_0} + \beta \cdot \left(\frac{P_{drive} + E_{drive}}{(DR)} + (C_1 + E_1) \right) \right]$$



TCO SUMMARY

- TCO depends on system configuration and sensitive to IO (# of rewrites per year)
- From TCO perspective, tape systems is more preferred toward at lower IO region
- If IO needs to increase beyond the fully a loaded tape library system by reducing cartridge count, a higher cost of IO will incur.