Track 1: Data and Metadata Organisation, Management and Access

Challenging the economy of tape storage with the disk-based one

Sang Un Ahn¹, Latchezar Betev², Eric Bonfillou², Heejune Han¹, Jeongheon Kim¹, Seung Hee Lee¹, Bernd Panzer-Steindel², Andreas-Joachim Peters², Heejun Yoon¹ ¹KISTI, Daejeon, South Korea ²CERN, Geneva, Switzerland





CHEP2023 @ Norfolk







Motivation

- Trying to compare purchasing and maintaining cost (a.k.a TCO) of two different types of custodial storages: tape and CDS (disk-based)
 - Purchasing cost includes hardware, software, licenses, installation and related training
 - Maintaining cost includes service contract (warranty), upgrade, power consumption, paid maintenance
 - Human resources for operation and maintenance
 - Cost for others such as floor space, cooling and outages not taken into account
 - Tape and CDS share the same floor and cooling system at KISTI, no economical impact on operation with outages while experiment point of view it could be significant

Disclaimer

- Supply chains, price tags, business models of manufacturers are different from regions and countries
 - Comparing to other East Asia countries, i.e. China and Japan, the market size of South Korea is rather small and there aren't local competitors to foreign manufacturers in server and storage domains
- Some predictions are based only on estimates that have not actually been procured n.b. prices could happen to be varied at the point of bidding

Architecture

Tape (2012~2021)



CDS (2021~)





SAS 2x12Gbps

Introduction to CDS

- a.k.a RAIN configuration
- I/O and commercial HSM software
- Better data accessibility: any data available at any time
- Avoiding vendor lock-in due to monopoly in tape market

CHEP2019 Proceeding [Link] vCHEP2021 Proceeding [Link]

 A disk based storage designed to store and preserve data from the ALICE experiment by accommodating EOS with its erasure code implementation,

Simpler architecture: removing additional disk buffers in front of tape library for

Provided to the ALICE experiment for commissioning at the early of 2021 and in production since November 2021, replacing the tape library completely



EOS Components on CDS







Usable space = 12 PB (out of 18 PB raw capacy)



- Thanks to spare FSTs.
 - Data are still accessible if 6 FSTs are offline
 - Data can be written if 2 FSTs are offline
 - One node (= 2 FSTs) can be turned off for maintenance at any time



• Data loss rate in a year is $\approx 8.6 \times 10^{-5}$ %, if 5 disks fail simultaneously in 5 independent FSTs considering 1.17% of AFR in practice cf. vendor published AFR is 0.35% (AFR = Annualized Failure Rate)





Cost Analysis

Tape Acquisition Cost (2012-2017)

- 1 PB Tape library (175,380,000 KRW = 153,270 USD)
 - IBM TS3500, TS1140 drives (x8) & 3592 4 TB media (x260, 1PB)
 - Frames (x5) capable to host up to 3 PB
 - Dual robotics
- TSM license (117,960,000 KRW = 103,088 USD)
- GPFS licenses (38,110,000 KRW = 33,305 USD)
- SAN switches (6,550,000 KRW = 5,724 USD)
- Additional Tape Cartridges
 - (2014) 560 TB (34,000,000 KRW = 29,713 USD)
 - (2016) 1,600 TB (96,200,000 KRW = 84,071 USD)
- Disk buffer (784,600,000 KRW = 685,683 USD)
 - Hitachi USP 600 TB (procured in 2010)
 - Hitachi VSP 600 TB (procured in 2015 to replace Hitachi USP)

Acquisition Costs Tape Libarary Cartridges TSM GPFS SAN SW Disk buffer (600TB) Sub Total

Total 3,200 TB + 600 TB Disk buffer 1,094,856 USD to 2017

2012	2013	2014	2015	2016	2017	TCO up
175,380,000						175
		34,000,000		96,200,000		130
117,960,000						117
38,110,000						38
6,550,000						e
425,600,000			359,000,000			784
338,000,000	0	34,000,000	359,000,000	96,200,000	0	1,252





Tape Maintenance Cost (2012-2017)

	2012	2013	2014	2015	2016	2017	TCO up to 2017
Acquisition Costs							
Tape Libarary	175,380,000						175,380,000
Cartridges			34,000,000		96,200,000		130,200,000
TSM	117,960,000						117,960,000
GPFS	38,110,000						38,110,000
SAN SW	6,550,000						6,550,000
Disk buffer (600TB)	425,600,000			359,000,000			784,600,000
Sub Total	338,000,000	0	34,000,000	359,000,000	96,200,000	0	1,252,800,000
Maintenance							
Tape drives replacement (x3)							0
Paid maintenance (per year)							0
GPFS Upgrade							0
Sub Total	0	0	0	0	0	0	0
Grand Total	338,000,000	0	34,000,000	359,000,000	96,200,000	0	1,252,800,000
VAT incl.	371,800,000	0	37,400,000	394,900,000	105,820,000	0	1,378,080,000
HR	70,000,000	70,000,000	70,000,000	70,000,000	70,000,000	70,000,000	420,000,000
Terabytes (usable)	1,040		560		1,600		3,200
KRW/TB							561,900
USD/TB	1 USD =	₩1,144.26					491.06

- 5 Year Warranty + Extended SLA due to penalty
- 1 FTE dedicated to Tape operation up to 2017
- No Power consumption data
- TCO up to 2017
 - 561,900 KRW / TB (= **491.06** USD / TB)





Tape Acquisition & Maintenance Cost (2018 quote) No real procurement performed for tape in 2018

- Average estimates for Tape library with over 7 PB capacity from 3 vendors (Spectra, IBM & Oracle) : 687,500,000 KRW = 600,824 USD
 - Including frames, robotics, softwares, etc.
- Additional estimates for IBM TS3500 Maintenance (226,300,000 KRW = 197,771 USD)
 - Including TS1140 tape drives replacement, GPFS upgrade



TCO for Tape (2012-2018) 281,692.37 KRW / TB = 246.18 USD / TB

	2012	2013	2014	2015	2016	2017	TCO up to 2017	2018 (quote)	Total (2017 Quote
Acquisition Costs									
Tape Libarary	175,380,000						175,380,000	687,500,000	86
Cartridges			34,000,000		96,200,000		130,200,000		13
TSM	117,960,000						117,960,000		11
GPFS	38,110,000						38,110,000		3
SAN SW	6,550,000						6,550,000		
Disk buffer (600TB)	425,600,000			359,000,000			784,600,000		78
Sub Total	338,000,000	0	34,000,000	359,000,000	96,200,000	0	1,252,800,000	687,500,000	1,94
Maintenance									
Tape drives replace (x3)							0	41,310,000	4
Paid maintenance							0	39,072,000	3
GPFS Upgrade							0	145,920,000	14
Sub Total	0	0	0	0	0	0	0	226,302,000	22
Grand Total	338,000,000	0	34,000,000	359,000,000	96,200,000	0	1,252,800,000	913,802,000	2,16
VAT incl.	371,800,000	0	37,400,000	394,900,000	105,820,000	0	1,378,080,000	1,005,182,200	2,38
HR	70,000,000	70,000,000	70,000,000	70,000,000	70,000,000	70,000,000	420,000,000	70,000,000	49
Terabytes (usable)	1,040		560		1,600		3,200	7,000	
KRW/TB							561,900		2
USD/TB	1 USD =	₩1,144.26					491.06		

- High-end enterprise-class disk storage used as disk buffer
- Open-source tape HSM (like CERN CTA) can help reduce acquisition & maintenance costs, but more FTE could be required
- Electricity not included (we didn't measure)

Tape IBM TS3500 TCO Model (KRW)



CDS Acquisition Cost (2019-2022)

- 18 PB (909,570,700 KRW = 794,898 USD)
 - 40Gbps 2-port NICs, 2x 12G 2-port SAS HBAs)
 - 18x DELL/EMC ME484 Enclosures (84x 12TB NL-SAS disks)
 - Including Racks (x4), 40Gbps Ethernet Switches (x2)
- EOS Erasure Coding layout applied
 - 4 Parity mode (qrain) + 2 Spares
 - 12 PB of usable space

	2019	2020	2021	2022	TCO up to 202
Acquisition Costs					
Servers	39,910,800				39,910,8
Racks	22,000,000				22,000,0
Enclosures	781,659,900				781,659,9
Ethernet SW (40Gbps)	66,000,000				66,000,0
Sub Total (VAT incl.)	909,570,700	0		0	909,570,7



• 10x DELL PowerEdge R730 (Intel Xeon E5-2637 v4 @ 3.5GHz, 8c, 2x SAS SSD 480GB, 2x SATA SSD 960GB, 2x

Total 12,000 TB (Usable) 794,898 USD

1 USD = 1,144.26 KRW (2012-2023 avg.)



CDS Maintenance Cost (2019-2022)

	2019	2020	2021	2022	TCO up to 2022
Acquisition Costs					
Servers	39,910,800				39,910,800
Racks	22,000,000				22,000,000
Enclosures	781,659,900				781,659,900
Ethernet SW (40Gbps)	66,000,000				66,000,000
Sub Total (VAT incl.)	909,570,700	0		0	909,570,700
Maintenance					
Disk replacement					0
Paid maintenance (6%)					0
Electricity			15,800,000	16,300,000	32,100,000
Sub Total (VAT incl., 6%))	0	0	0	0	0
Grand Total (6%)	909,570,700	0	0	0	909,570,700
Terabytes (usable)	12,096				
KRW/TB (6%)					75,196
USD/TB (6%)					65.7

1 USD = 1,144.26 KRW (2012-2023 avg.)

- 5 Year Warranty
 - 1.9% of AFR, about 30 disks replaced in a year
- Electricity (32,100,000 KRW = 28,053 USD)
 - Power consumption measured since
 Jan 2021
- No additional FTE compared to Tape
- TCO up to 2022
 - 75,196 KRW / TB
 - 65.7 USD / TB

)53 since

CDS Expansion & Maintenance Cost (2023-) Based on 2022 Quotes & Prediction

- 4 PB of usable space (720,000,000 KRW = 629,227 USD)
 - Current pledges will fulfill the requirement of ALICE T1 shares in 2025
 - Due to COVID-19 & global supply chain inconsistency, dramatic increase in disk and enclosure prices
- 86,630 KRW (75.71 USD) / TB (2019) ⇒ 148,748 KRW (129.99 USD) / TB (2023.01, ①71%) Paid Maintenance contracts required after 2024
 - Maintenance rate in negotiation : 4% (Governmental guide), 6-8% (Conventional rate), or up to 20% (Vendor demand)
- Electrical energy increases as the amount of stored data grows (see the next slide)

CDS Power Consumption (2021-2022)





```
Average kilowatt per month
```

Dec

TCO for CDS (2019-2025) 109,992 KRW / TB = 96 USD / TB (with 6% of maintenance rate)

	2019	2020	2021	2022	TCO up to 2022	2023	2024 (Prediction based on quotes)	2025 (Prediction based on quotes)	Total (2022 +)
Acquisition Costs									
Servers	39,910,800				39,910,800				
Racks	22,000,000				22,000,000				
Enclosures	781,659,900				781,659,900				7
Ethernet SW (40Gbps)	66,000,000				66,000,000				
Sub Total (VAT incl.)	909,570,700	0		0	909,570,700	0	720,000,000	0	1,6
Maintenance									
Disk replacement					0				
Paid maintenance (4%)					0		36,382,828	36,382,828	
Paid maintenance (6%)					0		54,574,242	54,574,242	1
Paid maintenance (20%)					0		181,914,140	181,914,140	3
Electricity			15,800,000	16,300,000	32,100,000	16,815,823	17,347,969	17,896,955	
Sub Total (VAT incl., 4%))							53,730,797	54,279,783	1
Sub Total (VAT incl., 6%))	0	0	0	0	0	0	71,922,211	72,471,197	1
Sub Total (VAT incl., 20%)							199,262,109	199,811,095	3
Grand Total (4%)							773,730,797	54,279,783	1,7
Grand Total (6%)	909,570,700	0	0	0	909,570,700	0	791,922,211	72,471,197	1,7
Grand Total (20%)						0	919,262,109	199,811,095	2,0
Terabytes (usable)	12,096						4,032		
KRW/TB (4%)									
KRW/TB (6%)					75,196				
KRW/TB (20%)									
USD/TB (4%)									
USD/TB (6%)					65.7				
USD/TB (20%)	1 USD =	₩1,144							

• Conservatively, taking the worst scenario of paid maintenance (20%), TCO for CDS will be 109.9 USD / TB

• Disk replacement cost included in current warranty and paid maintenance

• 3% of linear power consumption increase is assumed (based on the measurement in 2021-2022)



CDS DELL ME484 TCO Model (KRW)



Conclusions

		TCO (KRW/TB)	TCO (USD/TB)
Tape (2012-2018)	Procured	281,692.37	246.18
	Lowest estimate	193,034	168
CDS (2019-2025)	4%	107,736	94
(per Maintenance rate including annual 3% increase	6%	109,992	96
of electricity)	20%	125,784	110

- Remarks
 - Lowest estimate for tape: with cheaper disk buffer (JB0 library (0.45M USD for 7 PB in 2018)
 - Current high price tag of disks/enclosures and increase much worse for upcoming years

1 USD = 1,144.26 KRW (2012-2023 avg.)

• Lowest estimate for tape: with cheaper disk buffer (JBOD type, 0.6PB for 0.1M USD) and the lowest estimate for tape

• Current high price tag of disks/enclosures and increase of power consumption & electricity fares would make CDS TCO



