Backblaze

How Long Do Hard Drives and SSDs Live and What Can They Tell Us Along the Way

Andy Klein | andy@backblaze.com

Agenda

Andy Klein

7 years - Drive Stats Guy
25 years - Marketing
6 years - Sys Admin
8 years – Developer

- The world spins on hard drives
- The data we've collected
- Hard drive failure rates
- Fun facts I'll bet you didn't know about hard drives
- Can you predict drive failure?
- Hard drives versus SSDs

Hard Disk Drive (HDD)

AKA: Hard Drive, Disk Drive

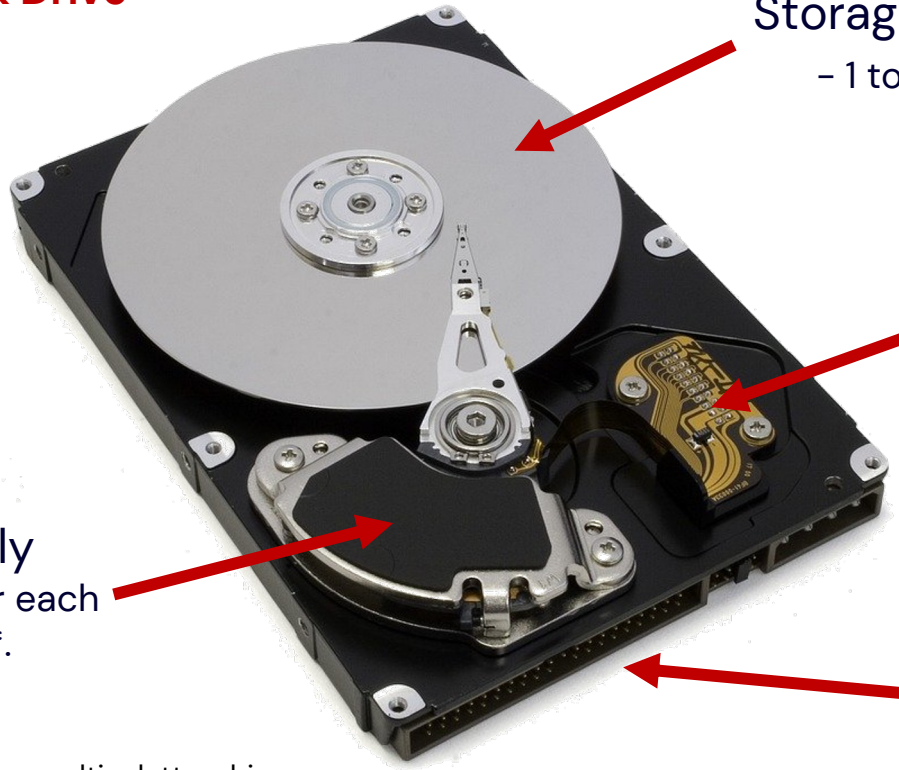

Storage (Disk Platter)
- 1 to 9 platters

Controller

Read/Write Assembly
- 1 read/write head for each side of each platter*.

Connectors

* - Except top and bottom in some multi-platter drives.

Worldwide Byte Shipments by Storage Media Type

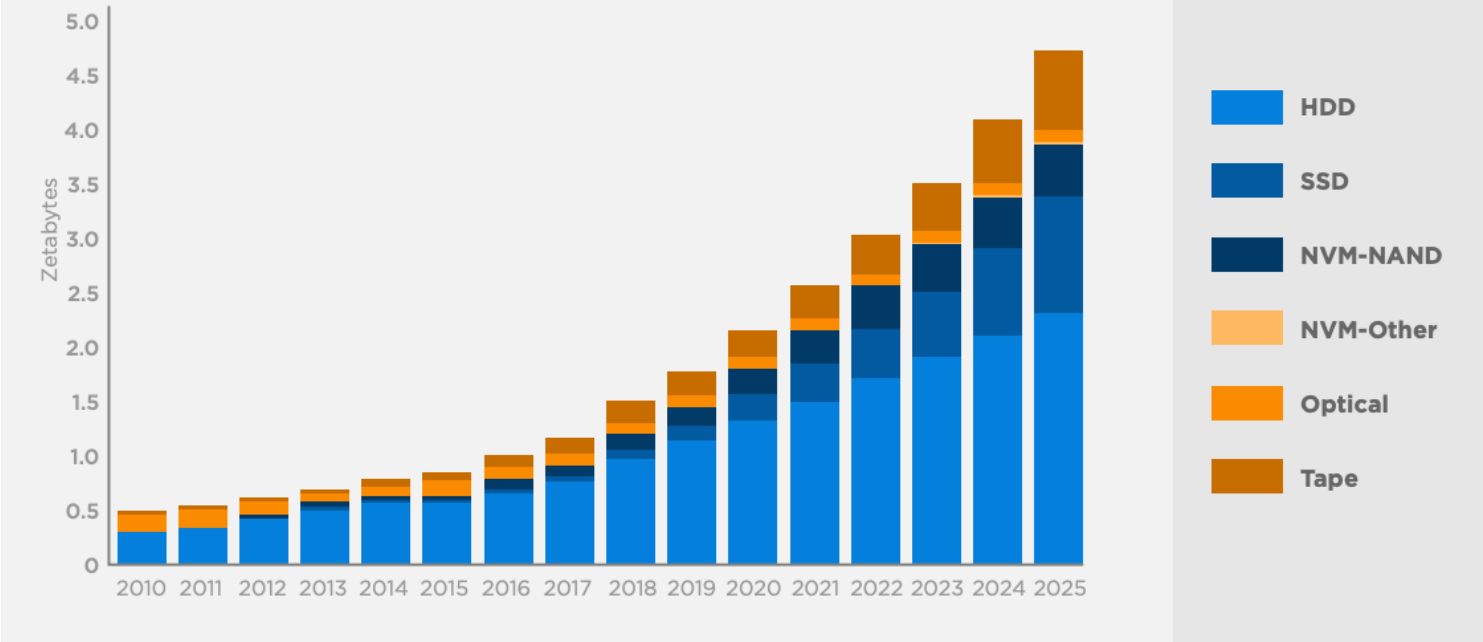

Source: Data Age 2025, sponsored by Seagate with data from IDC Global DataSphere, Nov 2018

Report: https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

Search for: “idc seagate data age 2025”

Lifetime Drive Stats

As of 6/30/2022

215,424
Active HDD data drives

2.4 Exabytes
Active HDD storage

293,611,067
Drive Days

11,219
Drive Failures

1.39%
Lifetime Annualized Failure Rate

The Drive Stats Data

- Collected and stored since April 2013
  - Use smartmontools package to collect data
    - https://www.smartmontools.org/
  - All drives in operation each day to create a CSV file for each day
  - About 1,400 files, 360 million records, 15GB of ZIP data, 100GB of raw data
- The data is open source

Backblaze

www.backblaze.com/drivestats

SNIA IOTTA Repository

iotta.snia.org/traces/reliability

Drive Data Collected Each Day

date	serial_number	model	capacity_bytes	failure	Smart_1_ normalized	Smart_1_ raw
12/7/21	Z305B2QN	ST4000DM000	4000787030016	0	98	2766
12/7/21	PL1331LAHG1S4H	HGST HMS5C4040ALE640	4000787030016	0	100	0
12/7/21	ZACH007	ST8000NM0055	8001563222016	1	81	139015
12/7/21	ZA130TTW	ST8000DM002	8001563222016	0	83	100901
12/7/21	ZA18CEBF	ST8000NM0055	8001563222016	0	81	140551
12/7/21	PL2331LAH3WYAJ	HGST HMS5C4040BLE640	4000787030016	0	100	0

More SMART stats ›››

SMART Stats:

- Collected using Smartmontools.
- There are 255 pairs of values per drive.

SMART Stat Attributes:

- Smart_1: Read Error Rate
- Smart_5: Reallocated Sector Count
- Smart_9: Power On Hours

en.wikipedia.org/wiki/S.M.A.R.T

.

More Drives ⌄⌄⌄⌄

Drive Day: The data collected for one drive for one day.

What's a Drive Failure

Reactive Failure

- The drive will not spin up or connect to the OS.
- The drive will not sync or stay synced in a storage array.

Proactive Failure

- Triggered by SMART stats, FSCK, etc.
- Reviewed by Backblaze before action is taken

Data Center:	Sac0
Pod:	pod-000-1113-01
Drive:	drive_0057
Tasks:	Replace Data Drive
Action:	Proactive
Reason:	High Offline Uncorrectable (SMART)
Brand:	HGST
Model:	HGST HUH721212ALN604
Serial:	8AJK007BH
Size:	12TB Drive
Notes:	5 Reallocated_Sector_Ct - 82 197 Current_Pending_Sector - 276 198 Offline_Uncorrectable - 266 199 UDMA_CRC_Error_Count - 0 9 Power_On_Hours - 23422 Found ATA error that is 2 hours old - CONSIDER REPLACING THIS DRIVE

Computing Annualized Failure Rate (AFR)

1. Define AFR cohort and period:
   a. Cohort = Model (All models active as of 12/31/2021)
   b. Period = 2021
2. Obtain Drive Days and Drive Failures for the cohort and period.
   a. Drive Days = 65,929,573
   b. Drive Failures = 1,820
   c. Drive Count = 202,759
3. Apply Formula: AFR = (Drive Failures / (Drive Days / 365)) * 100

AFR = (1820 / (65929573 / 365)) * 100 = **1.01%**

This method accounts for drives with different drive days within the period.

*Psst: That's why using Drive Count can give a different answer, e.g. 5,000 drives ((1820 / 202759) * 100) =* ***0.89%*** *AFR*

Hard Drive Failure Reports

- Quarterly Reports
  - All drives models in operation at the end of the quarter.
  - HDD drives
  - Separate report for SSDs
- Results for
  - Most recent quarter
  - Annual
  - Lifetime
- Find reports at
  - www.backblaze.com/blog
  - Search for drive stats

AFR by Manufacturer Quarter by Quarter

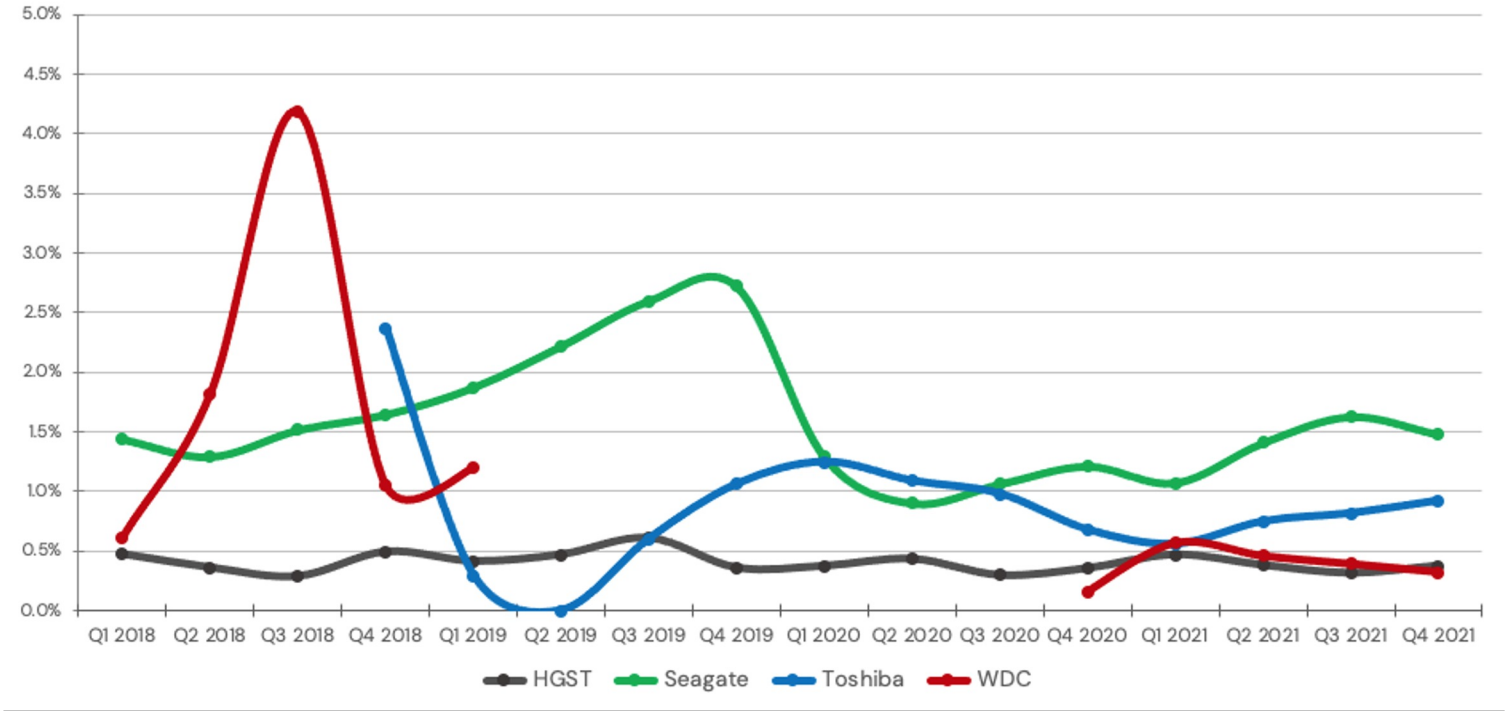

Lifetime Gold Medal Winners for Q2 2022

MFG	Model	Drive Size	AFR	Confidence Interval	
				Low	High
WDC	WUH721816ALE6L0 *(1)*	16TB	0.13%	0.00%	0.50%
Toshiba	MG08ACATEY	16TB	0.58%	0.40%	1.00%
WDC	WUH721414ALE6L4	14TB	0.29%	0.20%	0.40%
HGST	HUH721212ALE600	12TB	0.33%	0.20%	0.50%
HGST	HUH721212ALE604	12TB	0.46%	0.40%	0.60%
Seagate	ST10000NM0086	10TB	1.48%	1.20%	1.80%
HGST	HUH728080ALE600	8TB	0.60%	0.40%	0.90%
Seagate	ST6000DX000	6TB	0.87%	0.70%	1.10%
HGST	HMS5C4040BLE640 *(2)*	4TB	0.40%	0.40%	0.40%

(1) – Not available for retail sale in the US and Canada
(2) – Available as rehab drives only

Fun facts I'll bet you didn't know about hard drives

- Hard drives and the bathtub curve
- Temperature and failure
- Temperature and drive size
- Turn it off or leave it on
- How long to hard drive last?

Failure Rate Over Time: The Bathtub Curve

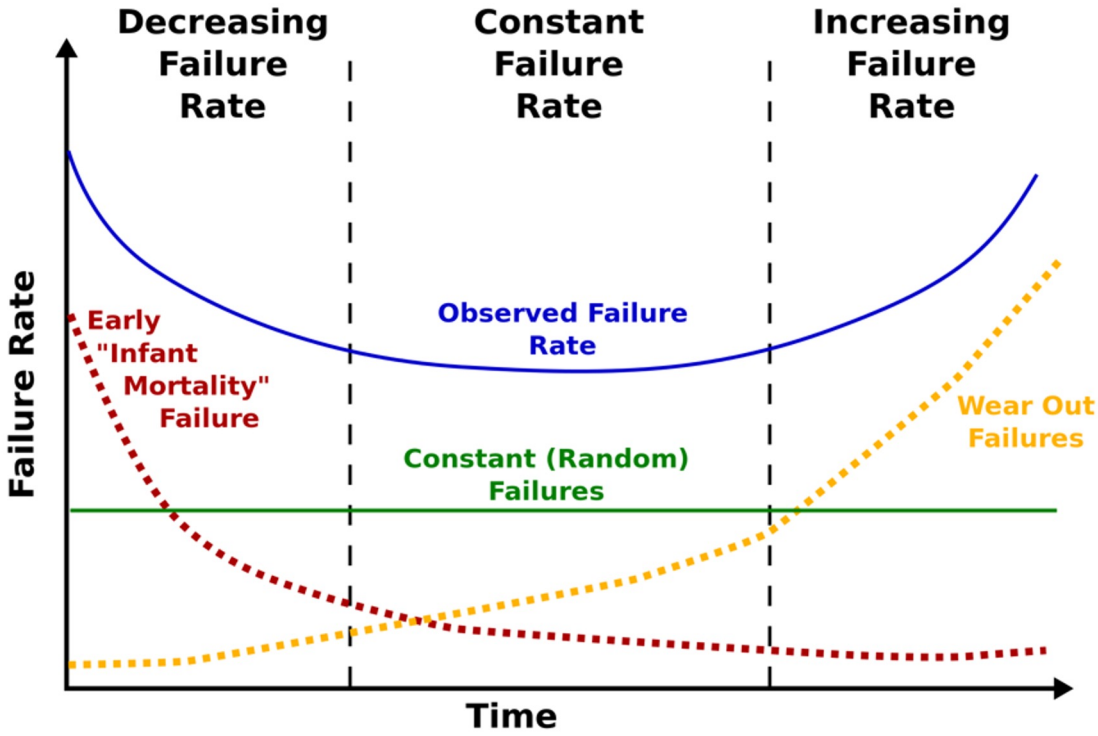

Drive Failure Over Time: 2013

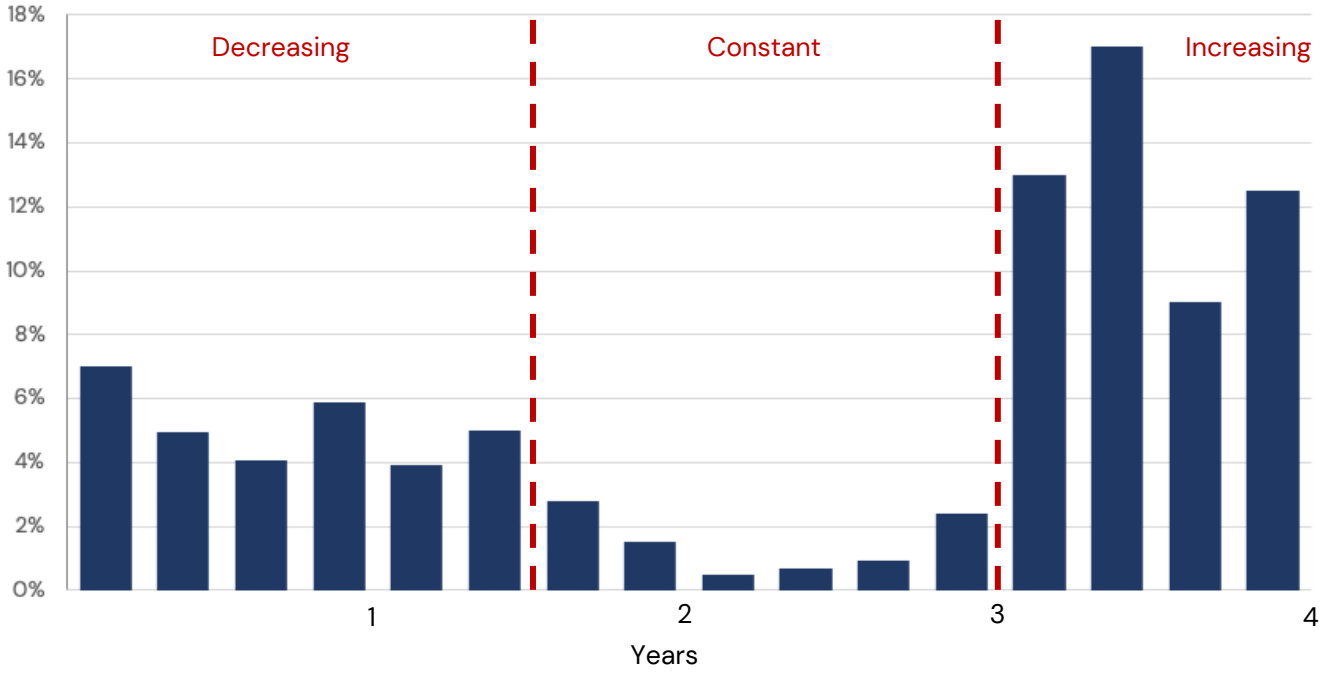

2021: Drive Failure Over Time

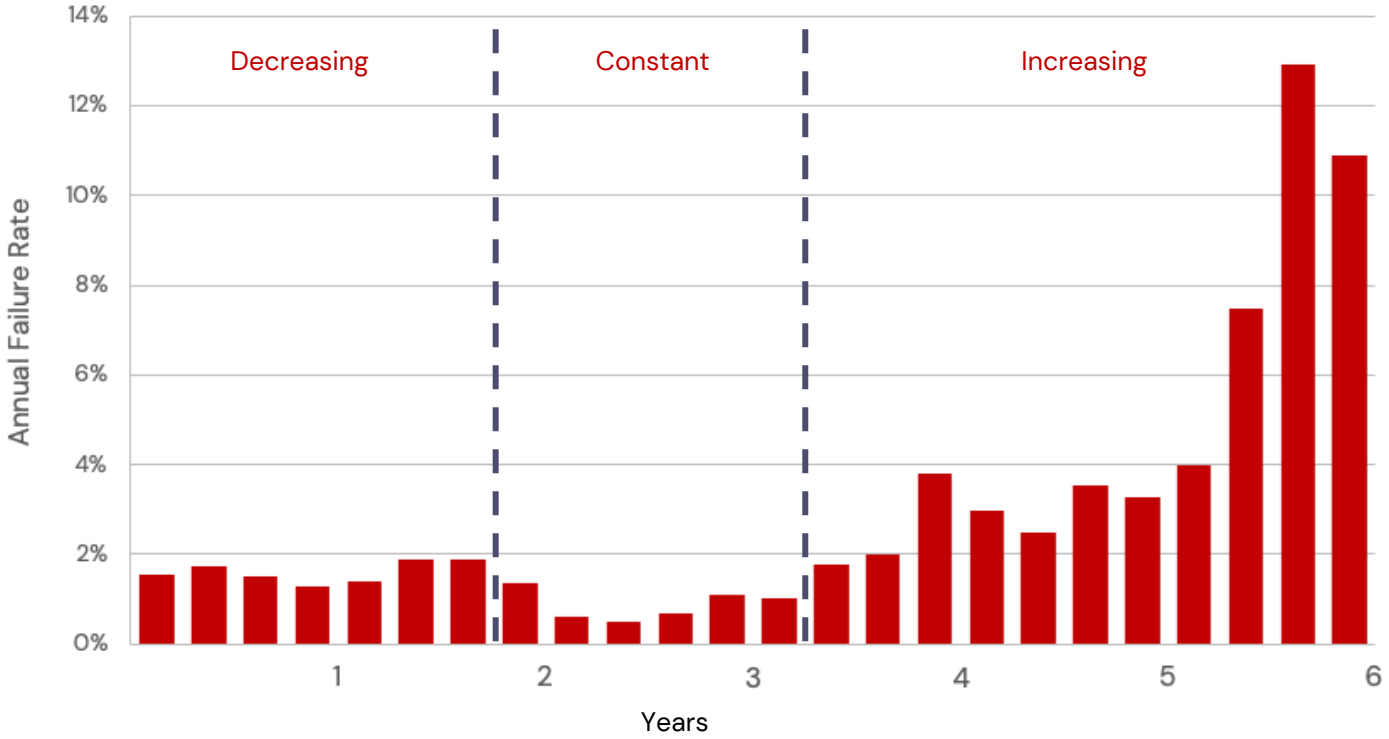

Drive Failure Over Time: 2013 versus 2021

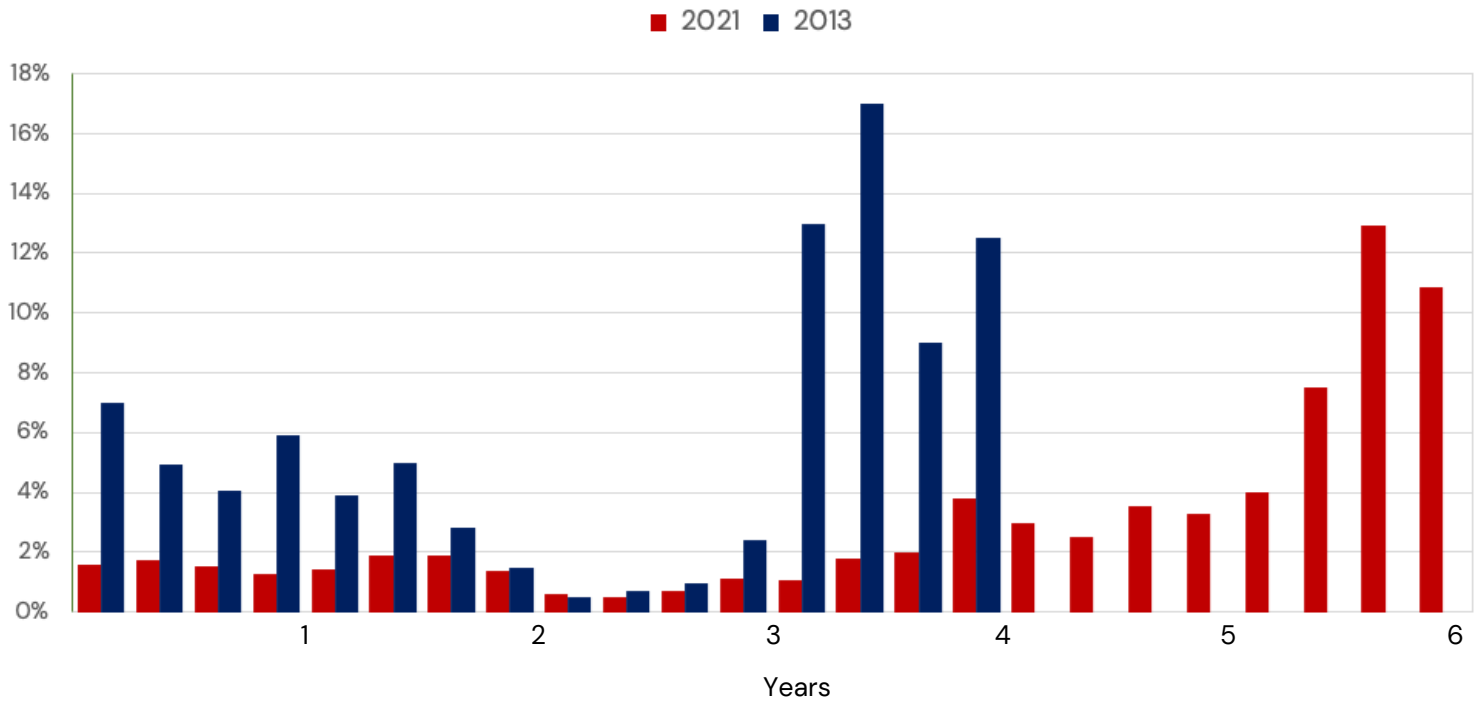

The 2021 New Normal

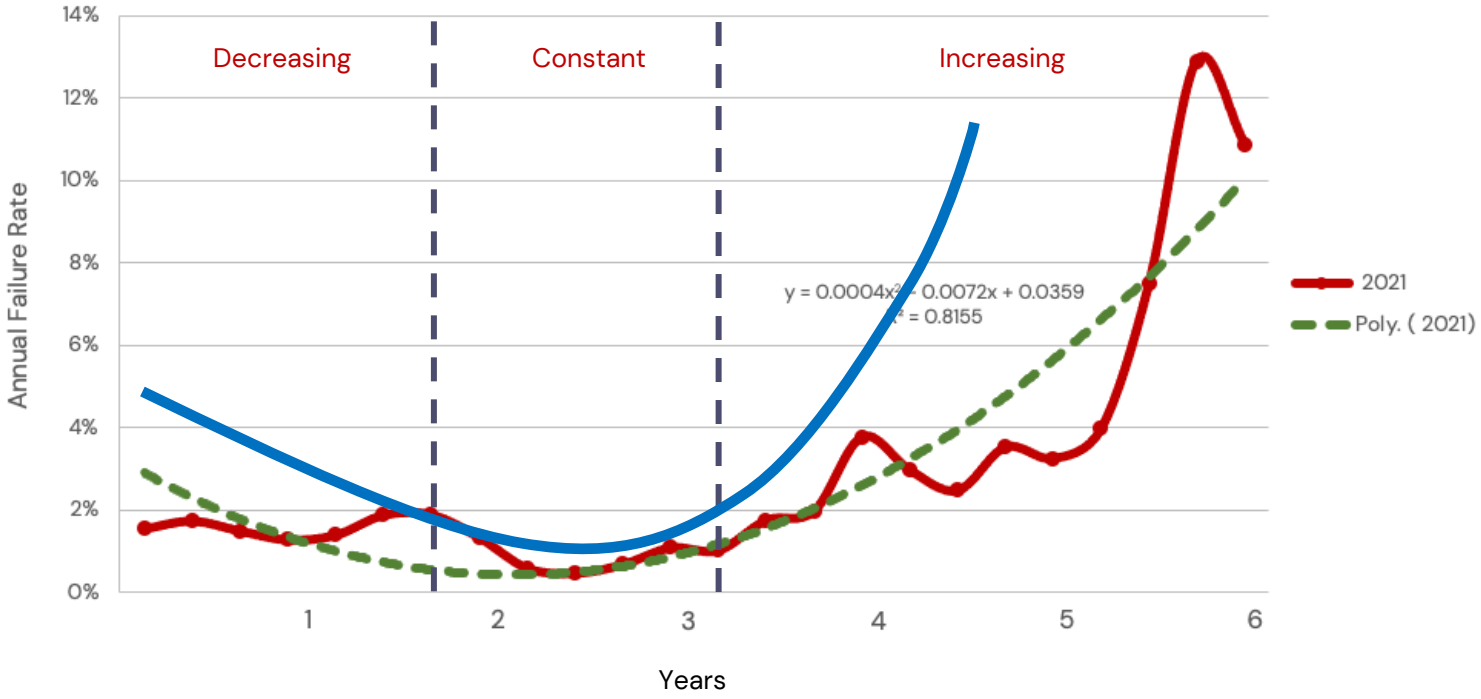

Temperature of Operational vs. Failed Drives (SMART 194)

Operational (Good) Drives

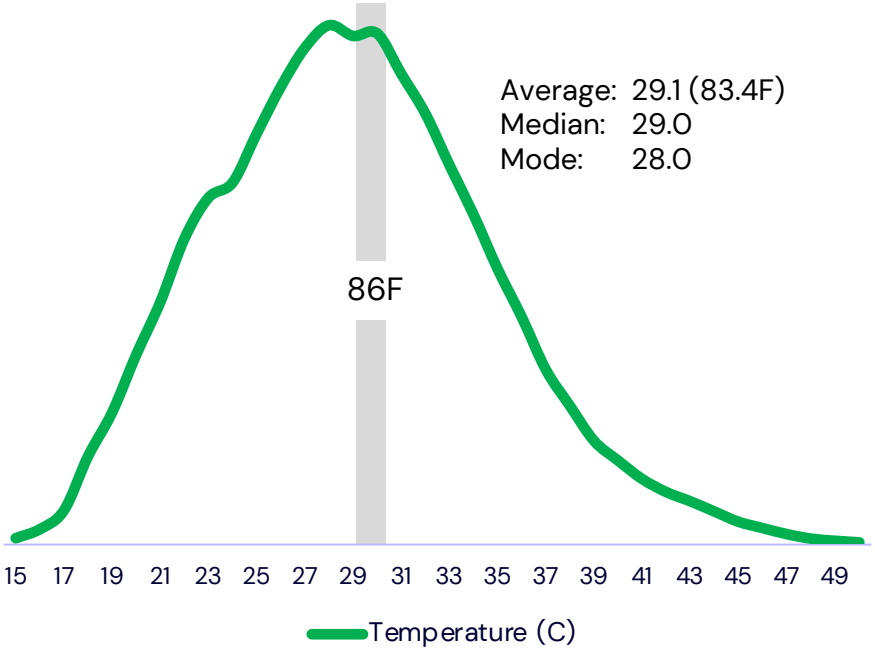

Failed Drives

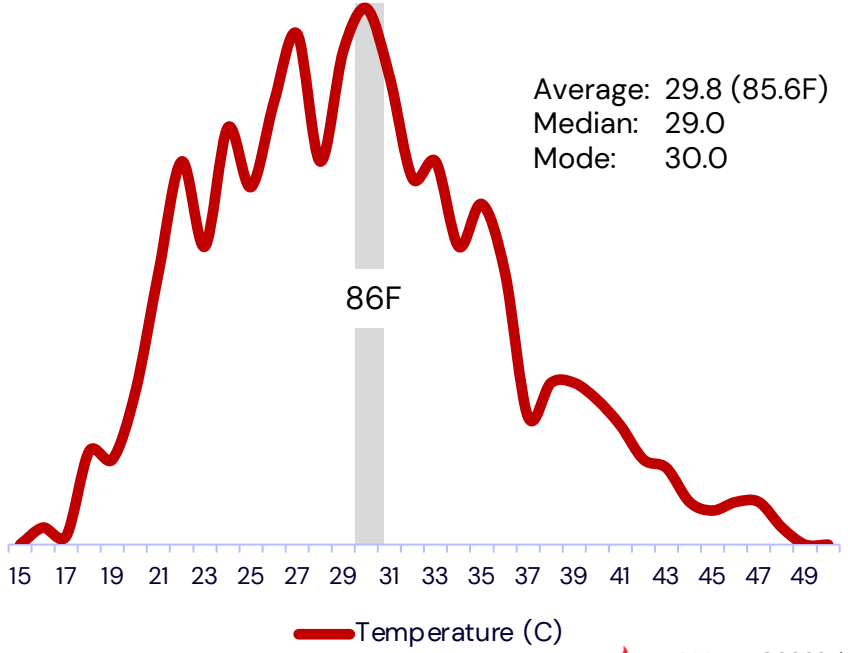

Temperature versus Drive Size

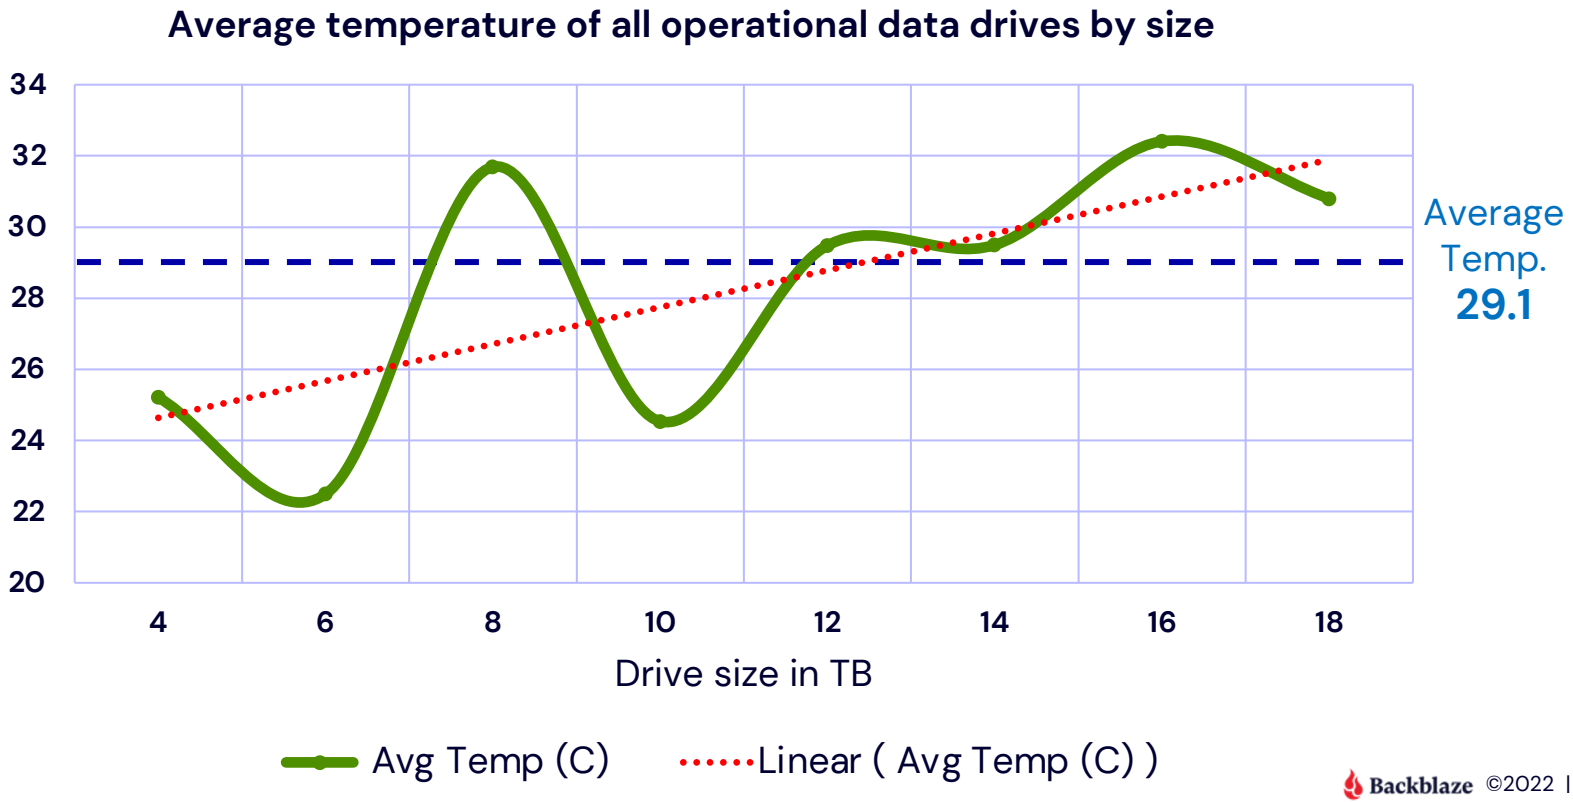

Power Cycing (Turn it off or leave it running)

Let’s compare the average number of power cycles

Good Drives
17.46

Failed Drives
21.24

Power Cycles

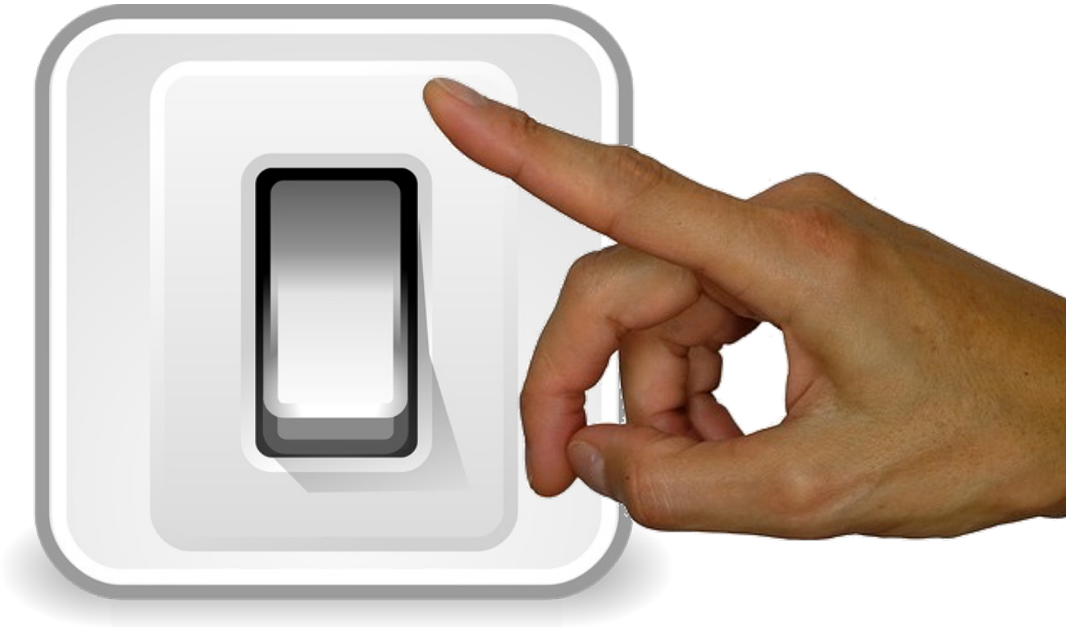

Hard Drive Life Expectancy (Survival Curves)

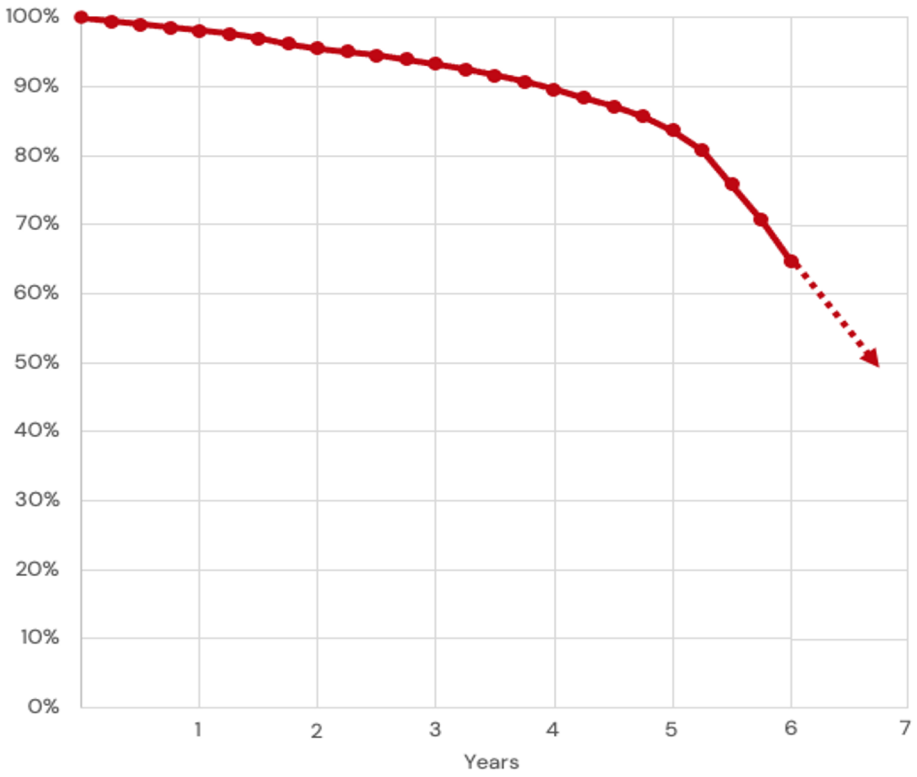

Kaplan-Meier Projection

Year	Survival Rate
1	98%
2	95%
3	92%
4	90%
5	84%
6	64%
6.75	50%

Hard Drive Life Expectancy – Select 8TB Drives

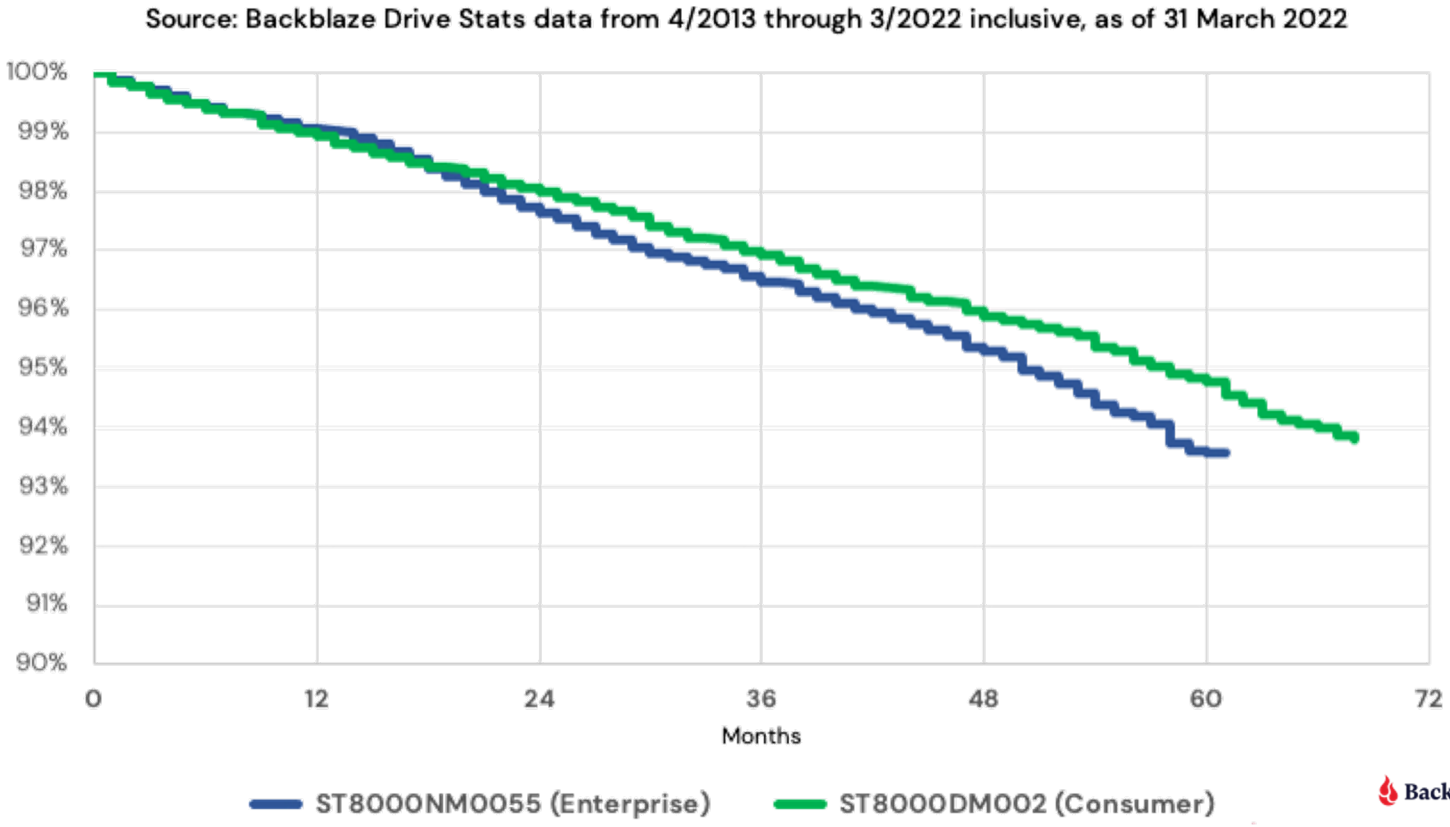

Hard Drive Life Expectancy – Select 12TB Drives

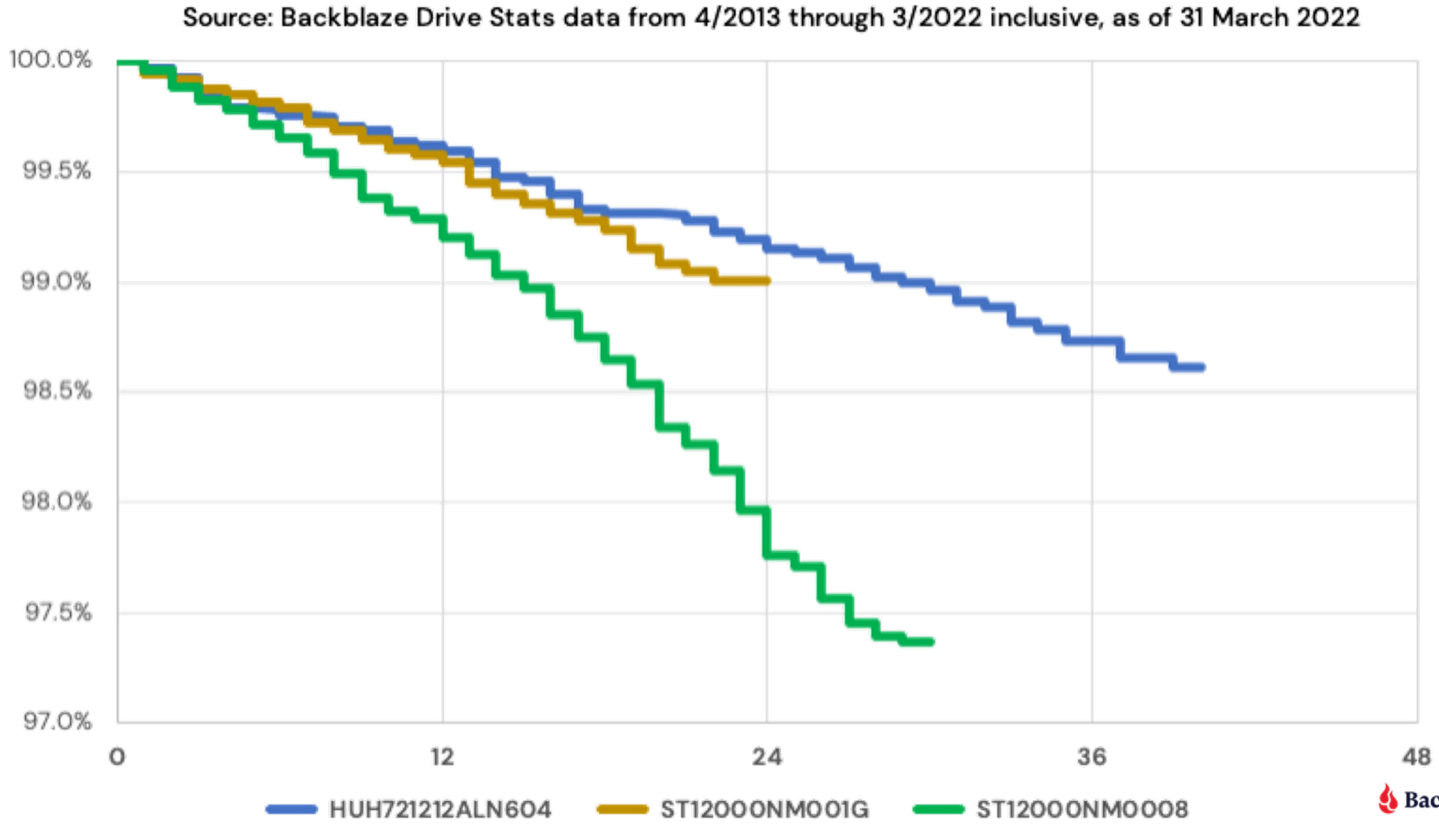

Hard Drive Life Expectancy – Select 14TB Drives

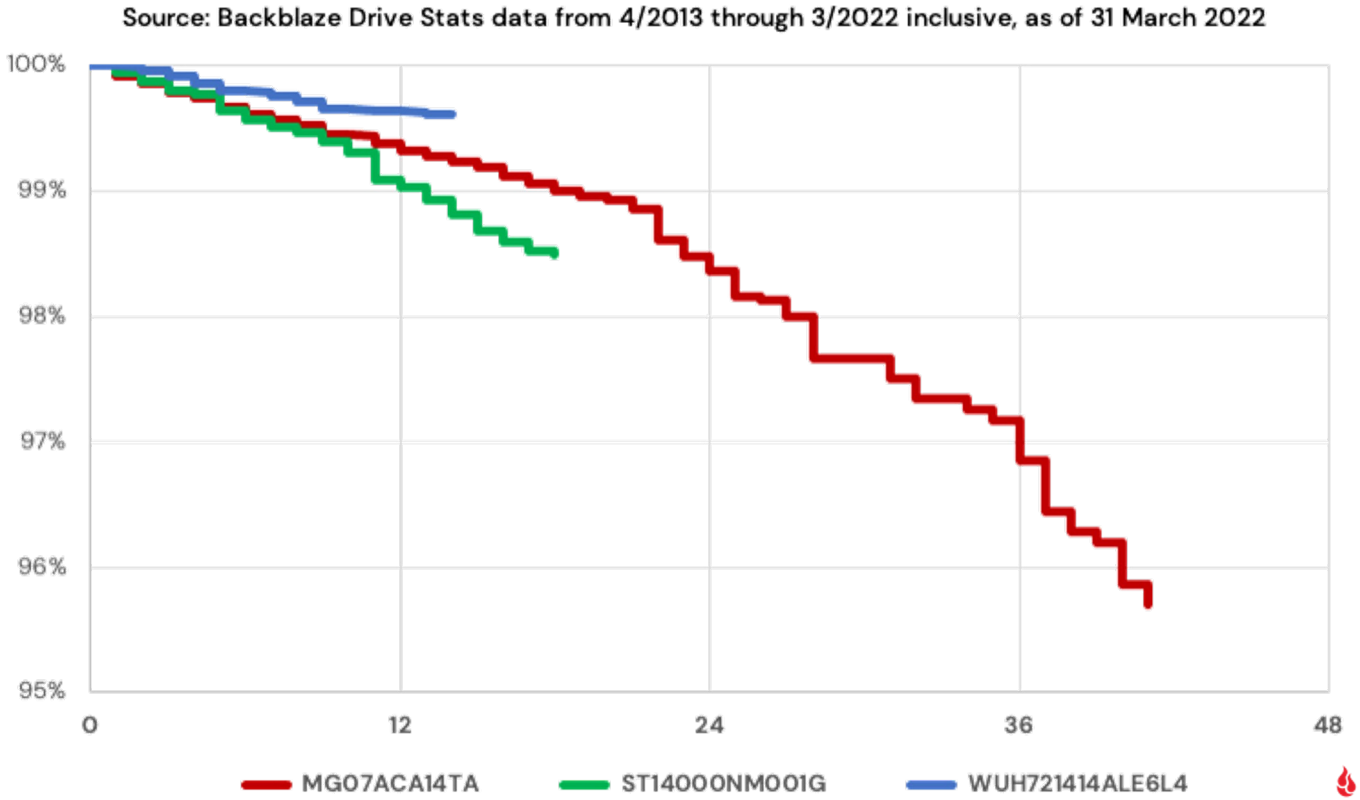

Actual vs. Projected Life Expectancy – Toshiba 14TB Drive

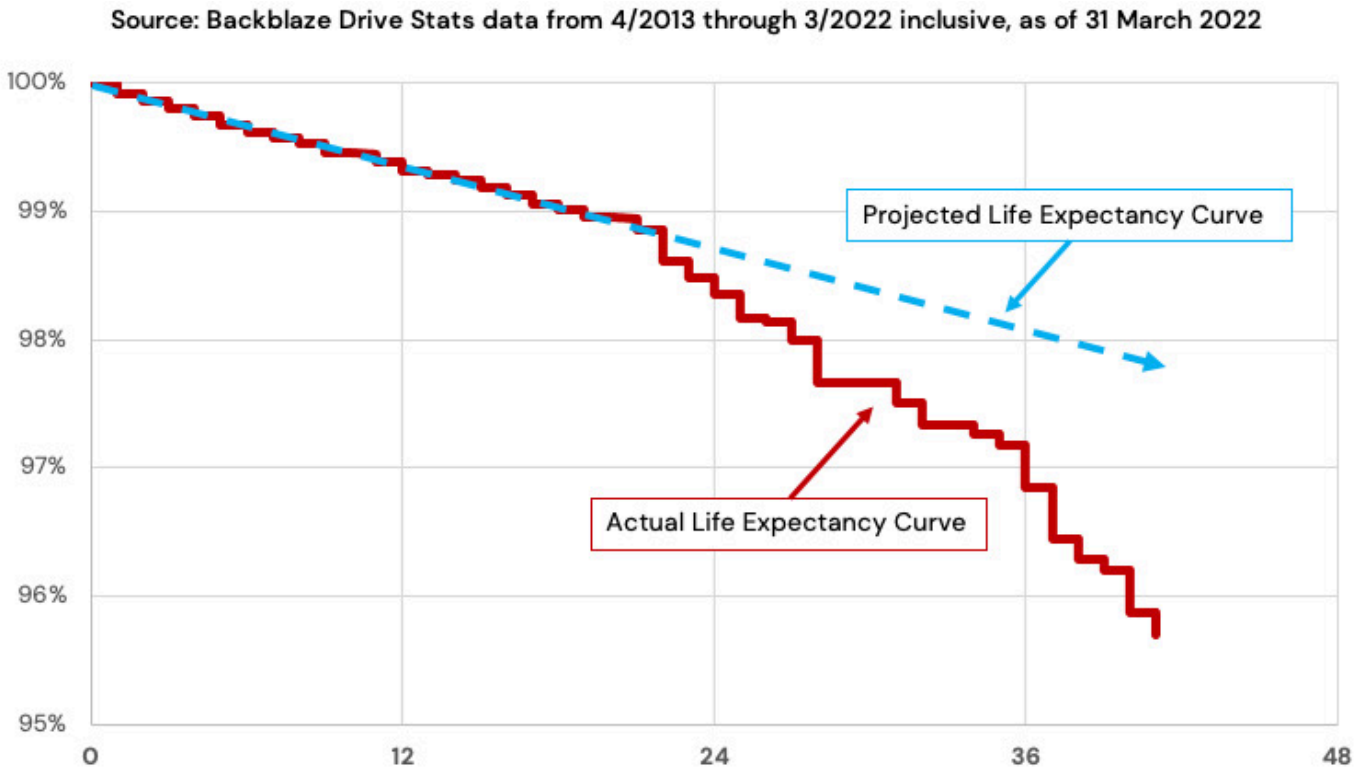

Predicting Hard Drive Failure

2016

Paper: Predicting Disk Replacement towards Reliable Data Centers

Authors: Botezatu, Mirela & Giurgiu, Ioana & Bogojeska, Jasmina & Wiesmann, Dorothea. (2016).

Location: *https://dl.acm.org/doi/10.1145/2939672.2939699*

...several others...

2021

Paper: Interpretable Predictive Maintenance for Hard Drives

Authors: Maxime Amram, Jack Dunn, Jeremy J. Toledano, Ying Daisy Zhuo

Location: *https://www.sciencedirect.com/science/article/pii/S2666827021000219*

Using Machine Learning to Predict Hard Drive Failure

Backblaze Drive Stats Data
+ Optimal Decision Trees
+ Survival Curves

= Predictions on long-term & short-term drive health

Drive Stats Data: https://www.backblaze.com/b2/hard-drive-test-data.html

Optimal Decision Trees

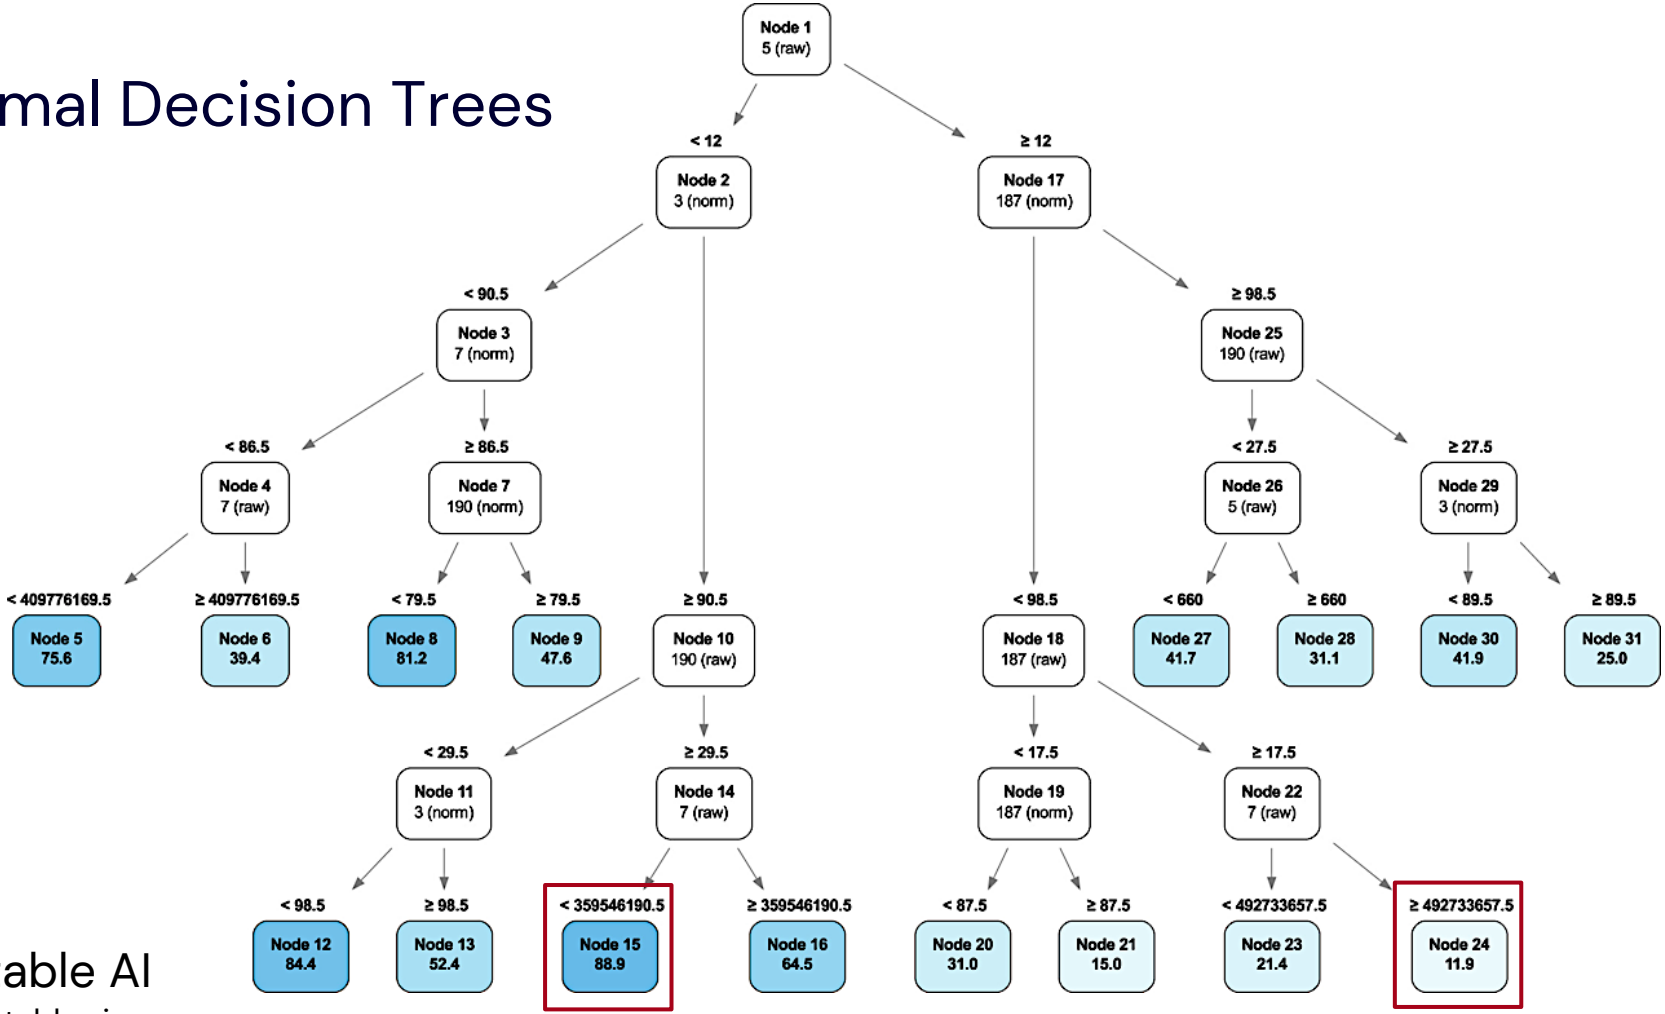

Predicting Short-term Drive Health

Example survival curves for selected cohorts

Healthy Drives

Node 15

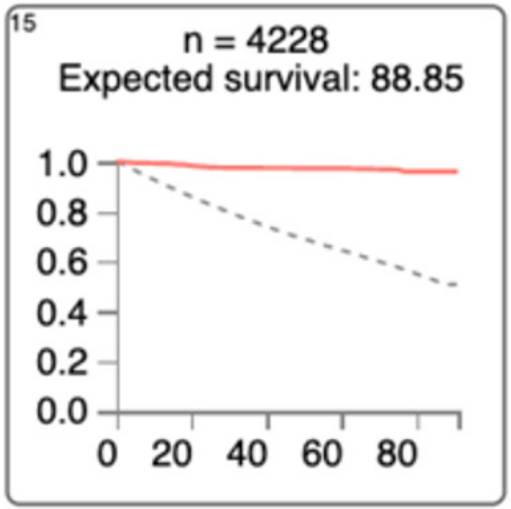

5 (raw) < 12

3 (norm) < 90.5

10 (raw) >= 29.5

7 (raw) >= 359546190.5

Predicting Short-term Drive Health

Example survival curves for selected cohorts

Unhealthy Drives

Node 24

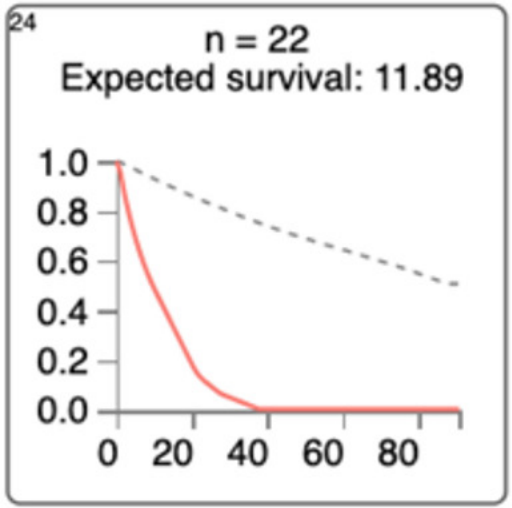

5 (raw) >= 12

187 (norm) < 98.5

187 (raw) >= 17.5

7 (raw) >= 492733657.5

SSD (Solid State Drive)

Psst: Saying SSD Drive is redundant, but saying SS Drive is probably worse…

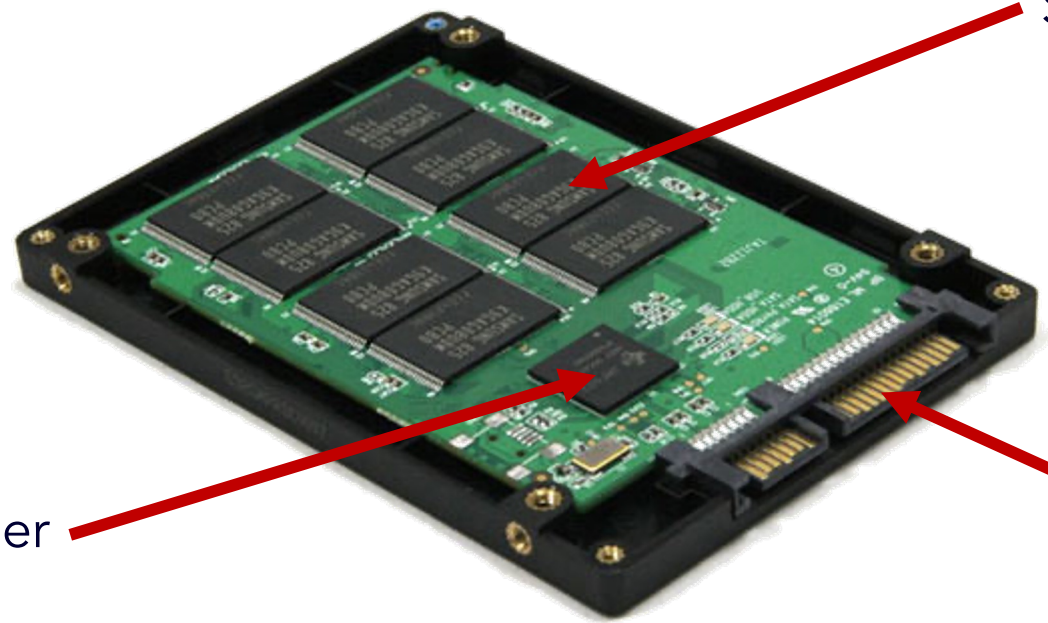

SSD versus HDD (Hard Disk Drive)

The Tale of the Tape

- Speed – SSD
- Electricity – SSD
- Cost – HDD
- **Reliability – ???**

SSD versus HDD Reliability

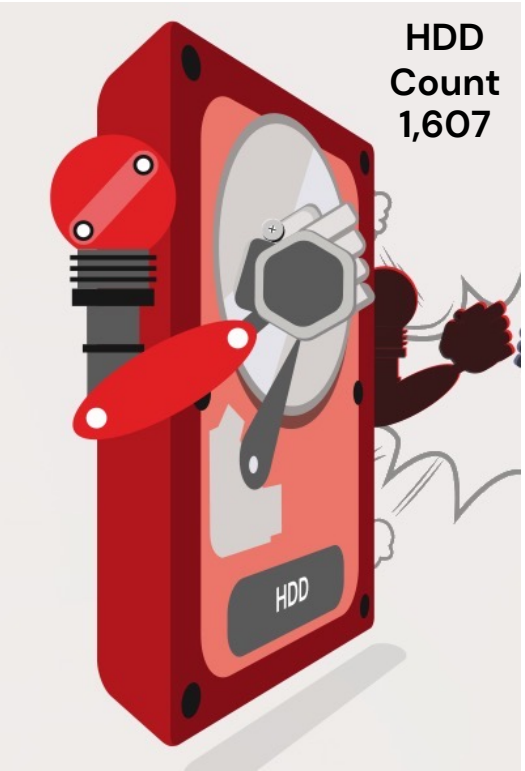

Boot Drives

- Boot the server
- Regular daily activity to read/write/delete log files for system access and diagnostics.

SSD versus HDD Reliability

HDD Annualized Failure Rate (Lifetime)

6.76%

SSD versus HDD Reliability

HDD
Annualized
Failure Rate
(Lifetime)

6.76%

SSD
Annualized
Failure Rate
(Lifetime)

1.22%

SSD versus HDD Reliability

HDD
Annualized
Failure Rate
(Lifetime)

6.76%

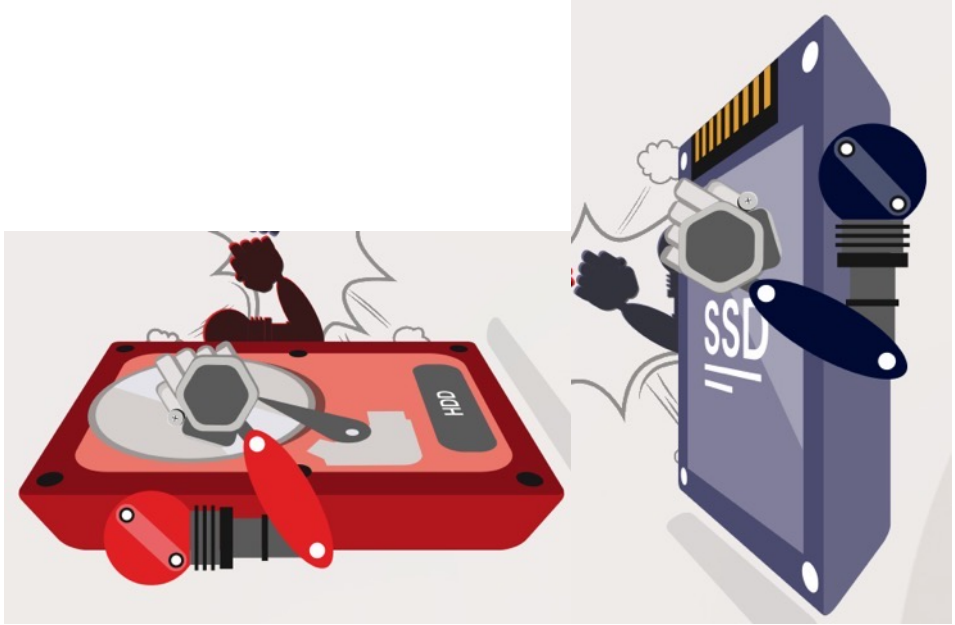

SSD
Annualized
Failure Rate
(Lifetime)

1.22%

SSD versus HDD Reliability

HDD
Annualized
Failure Rate
(Lifetime)

6.76%

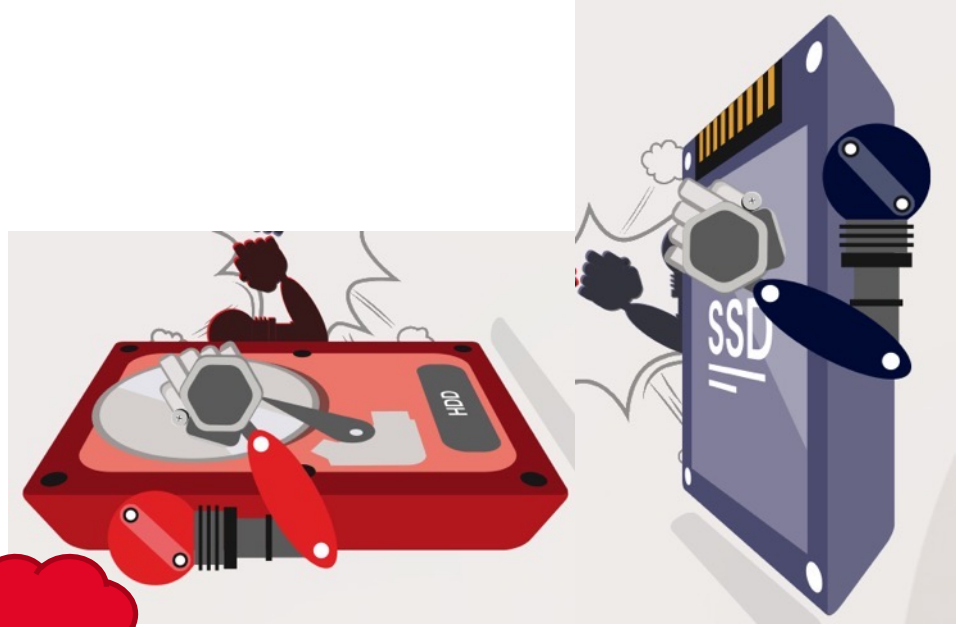

SSD
Annualized
Failure Rate
(Lifetime)

1.22%

SSD and HDD Lifetime AFR

Reporting Period: 04/2013 through 12/2021 inclusive

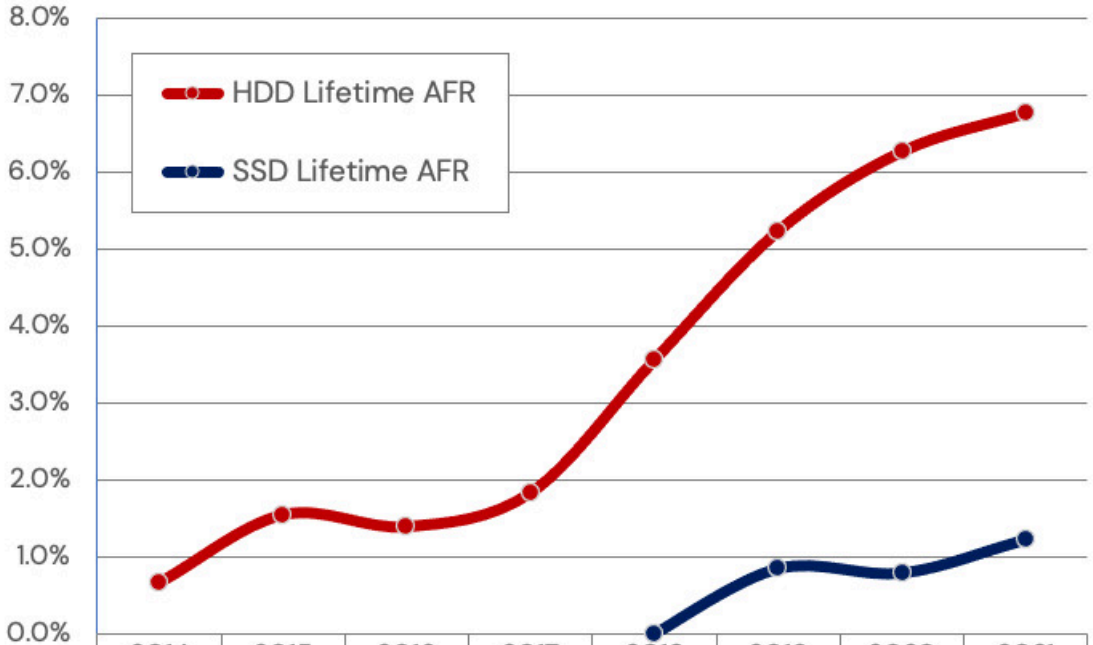

	2014	2015	2016	2017	2018	2019	2020	2021
HDD Lifetime AFR	0.66%	1.54%	1.38%	1.83%	3.55%	5.23%	6.26%	6.76%
SSD Lifetime AFR					0.00%	0.84%	0.79%	1.22%

SSD and HDD Lifetime AFR

Reporting Period: 04/2013 through 12/2021 inclusive

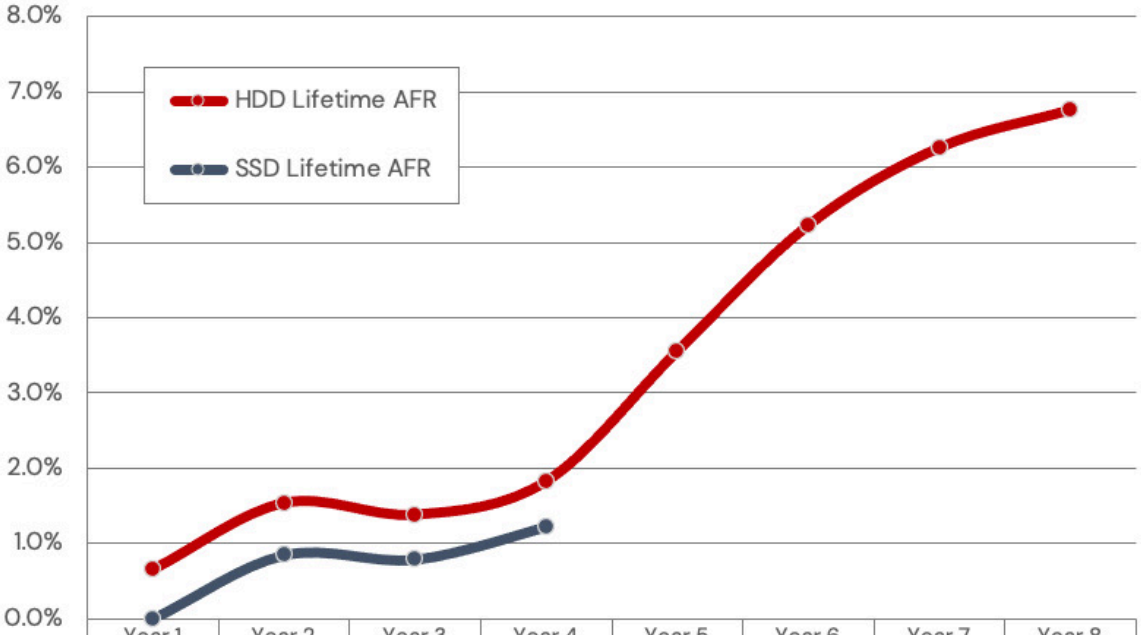

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8
HDD Lifetime AFR	0.66%	1.54%	1.38%	1.83%	3.55%	5.23%	6.26%	6.76%
SSD Lifetime AFR	0.00%	0.84%	0.79%	1.22%				

SSD and HDD Lifetime AFR

Reporting Period: 04/2013 through 12/2021 inclusive

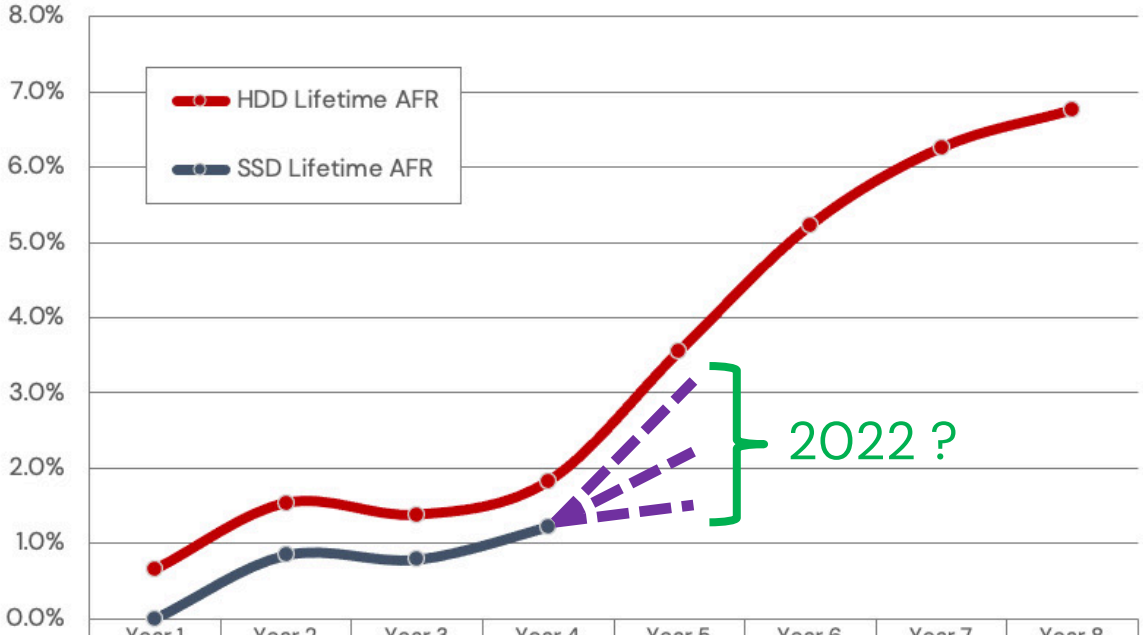

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8
HDD Lifetime AFR	0.66%	1.54%	1.38%	1.83%	3.55%	5.23%	6.26%	6.76%
SSD Lifetime AFR	0.00%	0.84%	0.79%	1.22%				

Backblaze Annual SDD Failure Rates for 2019, 2020, and 2021

Reporting periods are for each calendar year inclusive

			Annual 2019			Annual 2020			Annual 2021		
MFG	Model	Size (GB)	Drive Days	Drive Failures	AFR	Drive Days	Drive Failures	AFR	Drive Days	Drive Failures	AFR
Crucial	CT250MX500SSD1	500							1,689	2	43.22%
Dell	DELLBOSS VD	500	12,560		0.00%	31,566		0.00%	63,710		0.00%
Micron	MTFDDAV240TCB	240				1,473		0.00%	33,478	7	7.63%
Micron	MTFDDAV240TDU	240				128		0.00%			
Seagate	SSDSCKKB480G8R	480				4		0.00%			
Seagate	ZA250CM10003	250				47,161		0.00%	276,281	8	1.06%
Seagate	ZA2000CM10002	2,000	1,166		0.00%	1,460		0.00%	1,267	1	28.81%
Seagate	ZA250CM10002	300	7,100		0.00%	154,144	5	1.18%	204,287	2	0.36%
Seagate	ZA500CM10002	500	6,346	1	5.75%	6,583		0.00%	6,515		0.00%
Seagate	SSD	300	15,350		0.00%	39,407	1	0.93%	39,147	1	0.93%
			42,522	**1**	**0.86%**	**281,926**	**6**	**0.78%**	**626,374**	**21**	**1.22%**

Backblaze Cumulative SSD Annualized Failure Rate by Drive Model

Selected Models: Cumulative from the quarter when the first drive was in placed operation

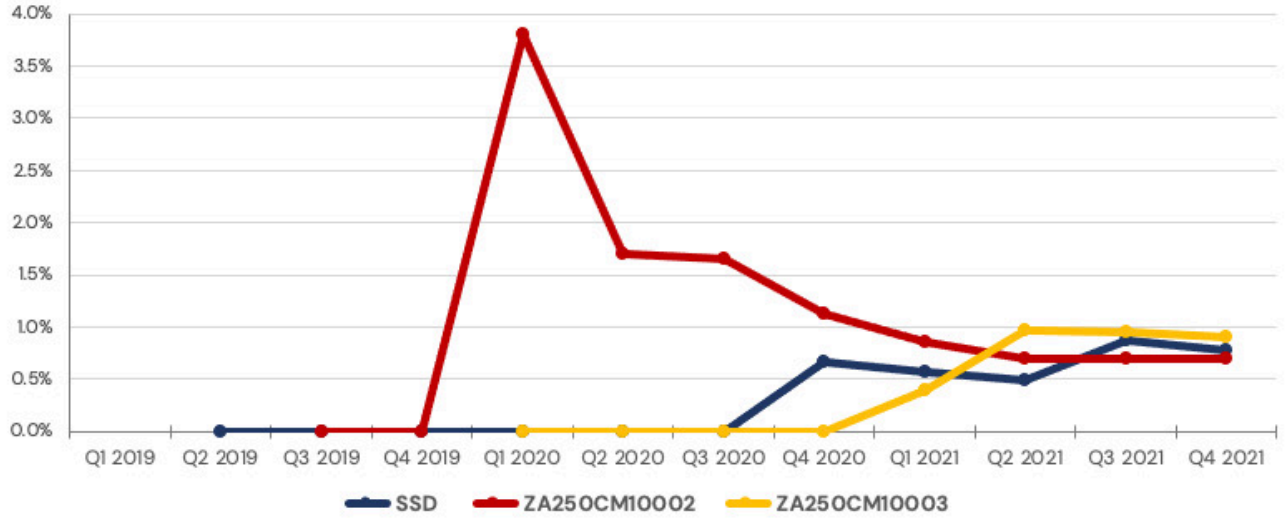

Seagate Model	Q1 2019	Q2 2019	Q3 2019	Q4 2019	Q1 2020	Q2 2020	Q3 2020	Q4 2020	Q1 2021	Q2 2021	Q3 2021	Q4 2021
SSD		0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.7%	0.6%	0.5%	0.9%	0.8%
ZA250CM10002			0.0%	0.0%	3.8%	1.7%	1.7%	1.1%	0.9%	0.7%	0.7%	0.7%
ZA250CM10003					0.0%	0.0%	0.0%	0.0%	0.4%	1.0%	0.9%	0.9%

Drive Count for Selected Models

Seagate Model	Q1 2019	Q2 2019	Q3 2019	Q4 2019	Q1 2020	Q2 2020	Q3 2020	Q4 2020	Q1 2021	Q2 2021	Q3 2021	Q4 2021
SSD		37	96	107	108	108	108	108	108	108	107	107
ZA250CM10002			7	157	347	436	550	558	559	556	560	562
ZA250CM10003					1	112	219	406	589	746	959	1090

Andy's Rules for Hard Drives (and Mostly for SSDs too)

- **Back them up!**
- Buy the type you need
  - Consumer, Surveillance/Video, NAS, Data Center, Enterprise
  - Watch out for SMR drives
- They will eventually fail
  - Consumer drives – 3+ years
  - Enterprise drives – 5+ years
- Keep them comfortable
  - Room temperature, not near heaters, sunlight, A/C
  - They don't like electrical shocks – think surge protector
  - They don't like vibration (i.e., don't drop, throw, or fling)

Summary

- The world spins on hard drives
- The data we've collected
- Hard drive failure rates
- Fun facts I'll bet you didn't know about hard drives
- Can you predict drive failure?
- Hard drives versus SSDs

Backblaze

Questions

Backblaze

Thank You

Andy Klein | Backblaze