Backblaze

What Ten Years of Drive Stats Data Can Tell Us

Andy Klein | andy@backblaze.com

Agenda

Andy Klein
Drive Stats Guy
Product Marketing/Management
Sys Admin/Network Admin
Developer

- Since we last met at MSST
- The data we collect
- Let’s talk about drive failure
- The bathtub curve lives on
- Let’s talk about drive survival
- Using Machine Learning to predict drive failure
- If we have time: Bonus content

Drive Stats Since the Last Time at MSST

	For active hard drive models as of	
	Q1 2017	Q1 2023
Hard Drive Count	95,054	237,278
Hard Drive Storage Capacity	351 Petabytes	2.7 Exabytes
Drive Days	62 Million	355 Million
Drive Failures	5,674	13,660
Lifetime Annualized Failure Rate	2.27%	1.40%

Data Collection

- Data has been collected and stored since April 2013
  - Use smartmontools package to collect data
    - https://www.smartmontools.org/
- All data and boot drives in operation each day
- Data for each day is stored in a CSV file
  - About 3,700 files
  - 380 million records
  - 15GB of ZIP data
  - 110GB of raw data

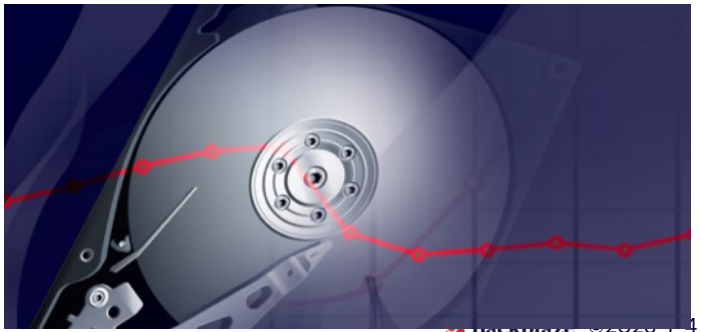

The Data is Open Source

backblaze.com/drivestats
or
backblaze.com/b2/hard-drive-test-data.html

iotta.snia.org/traces/reliability

Drive Data Collected Each Day

date	serial_number	model	capacity_bytes	failure	Smart_1_ normalized	Smart_1_ raw
12/31/22	2Z305B2QN	ST4000DM000	4000787030016	0	98	2766
12/31/22	**PK1331PAJ2V6WS**	**HGST HDS724040ALE640**	**4000787030016**	**0**	**100**	**0**
12/31/22	ZACH007	ST8000NM0055	8001563222016	1	81	139015
12/31/22	ZA130TTW	ST8000DM002	8001563222016	0	83	100901
12/31/22	ZA18CEBF	ST8000NM0055	8001563222016	0	81	140551
12/31/22	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 ATA, FSCK, SMART stats, Shard integrity
- Reviewed by Backblaze before removal is taken

Proactive Failure Death Spiral

Apr 25 01:04:05 Device: /dev/sdax [SAT], ATA error count increased from 38 to 39
Apr 25 05:34:04 Device: /dev/sdax [SAT], ATA error count increased from 39 to 41
Apr 25 07:04:05 Device: /dev/sdax [SAT], 64 Currently unreadable (pending) sectors (changed +8)
Apr 25 07:04:05 Device: /dev/sdax [SAT], 64 Offline uncorrectable sectors (changed +8)
Apr 25 07:34:05 Device: /dev/sdax [SAT], ATA error count increased from 41 to 42
Apr 25 09:34:06 Device: /dev/sdax [SAT], ATA error count increased from 42 to 43
Apr 25 10:04:05 Device: /dev/sdax [SAT], ATA error count increased from 43 to 46
Apr 25 12:04:12 Device: /dev/sdax [SAT], 88 Currently unreadable (pending) sectors (changed +24)
Apr 25 12:04:12 Device: /dev/sdax [SAT], 88 Offline uncorrectable sectors (changed +24)
Apr 25 12:04:12 Device: /dev/sdax [SAT], ATA error count increased from 46 to 55
Apr 25 12:34:06 Device: /dev/sdax [SAT], ATA error count increased from 55 to 58
Apr 26 06:34:05 Device: /dev/sdax [SAT], ATA error count increased from 58 to 67
Apr 26 07:04:06 Device: /dev/sdax [SAT], ATA error count increased from 67 to 68
Apr 29 10:04:04 Device: /dev/sdax [SAT], ATA error count increased from 68 to 69
Apr 29 11:04:05 Device: /dev/sdax [SAT], ATA error count increased from 69 to 71
Apr 29 12:34:12 Device: /dev/sdax [SAT], 168 Currently unreadable (pending) sectors (changed +96)
Apr 29 12:34:12 Device: /dev/sdax [SAT], 168 Offline uncorrectable sectors (changed +96)
Apr 29 12:34:13 Device: /dev/sdax [SAT], ATA error count increased from 71 to 90
Apr 29 13:04:05 Device: /dev/sdax [SAT], 216 Currently unreadable (pending) sectors (changed +48)
Apr 29 13:04:05 Device: /dev/sdax [SAT], 216 Offline uncorrectable sectors (changed +48)
Apr 29 13:04:05 Device: /dev/sdax [SAT], ATA error count increased from 90 to 120
Apr 29 13:17:05 kernel: [13594389.932700] sd 16:3:0:0: [sdax] Stopping disk

Proactive Failure Using SMART Attributes

SMART attributes and thresholds are vendor, and most times, model specific

ID#	ATTRIBUTE_NAME	FLAG	VALUE	WORST	THRESH	TYPE	UPDATED		CURRENT_RAW_VALUE
1	Raw_Read_Error_Rate	0x000f	049	049	050	Pre-fail	Always	-	FAILING_NOW 13818560
3	Spin_Up_Time	0x0003	091	091	000	Pre-fail	Always	-	0
4	Start_Stop_Count	0x0032	100	100	020	Old_age	Always	-	9
5	Reallocated_Sector_Ct	0x0033	100	100	010	Pre-fail	Always	-	0
7	Seek_Error_Rate	0x000f	087	060	030	Pre-fail	Always	-	547583970
9	Power_On_Hours	0x0032	022	022	000	Old_age	Always	-	68553
10	Spin_Retry_Count	0x0013	100	100	097	Pre-fail	Always	-	0
12	Power_Cycle_Count	0x0032	100	100	020	Old_age	Always	-	9
183	Runtime_Bad_Block	0x0032	100	100	000	Old_age	Always	-	0
184	End-to-End_Error	0x0032	100	100	099	Old_age	Always	-	0
187	Reported_Uncorrect	0x0032	100	100	000	Old_age	Always	-	0
188	Command_Timeout	0x0032	100	100	000	Old_age	Always	-	0
189	High_Fly_Writes	0x003a	100	100	000	Old_age	Always	-	0
191	G-Sense_Error_Rate	0x0032	100	100	000	Old_age	Always	-	0
192	Power-Off_Retract_Count	0x0032	100	100	000	Old_age	Always	-	0
194	Temperature_Celsius	0x0022	032	042	055	Old_age	Always	-	32
197	Current_Pending_Sector	0x0012	100	100	000	Old_age	Always	-	0
198	Offline_Uncorrectable	0x0010	100	100	000	Old_age	Offline	-	0
199	UDMA_CRC_Error_Count	0x003e	200	200	000	Old_age	Always	-	0

We Don’t Use Drive Count

Date	Drive Count
1/1/2022	202,759
3/31/2022	207,478
6/30/2022	215,011
9/30/2022	226,309
12/31/2022	230,921

For 2022 we had 2,963 failed drives. What drive count should we use? Let’s try 12/31/2022.

= (2,963 / 230,921) * 100 = **1.28%**

Why is it incorrect?

- Using Drive Count assumes you have the same number of drives over the entire period of measurement.
- Our environment is dynamic with drives entering and leaving almost daily.

Computing Annualized Failure Rate (AFR)

AFR = (Drive Failures / (Drive Days / 365)) * 100

1. Define AFR cohort and period
   a. Cohort = Model (All models active as of 12/31/2022)
   b. Period = 2022
2. Obtain Drive Days and Drive Failures for the cohort and period
   a. Drive Days = 78,760,434
   b. Drive Failures = 2,963
3. Apply Formula

AFR = (2,963 / (78,760,434 / 365)) * 100 = **1.37%**

Putting the Data To Work

- Quarterly Reports Since 2016 Containing:
  - Quarterly (data for the quarter specified)
  - Annual (data for the year specified)
  - Lifetime (data back to April 2013)
- Other Topics
  - Semi-annual SSD Drive Stats
  - Drive failure over time
  - Does drive temperature matter
  - Predicting Drive Failure with SMART Stats
  - Many more

Backblaze Hard Drive Stats for 2022

Backblaze Hard Drive Stats for Q3 2022

Backblaze Hard Drive Stats for Q2 2022

Backblaze Hard Drive Stats for Q1 2022

The Helium Factor and Hard Drive Failure Rates

HDD vs SSD: What Does the Future for Storage Hold?

500 Petabytes and Counting

What is HAMR...

Yes, Backblaze Just Ordered 100 Petabytes of Hard Drives

Hard Drive Cost Per Gigabyte

What SMART Stats Tell Us About Hard Drives

HGST 8TB Drives – Helium Makes Them Fly

A Look at Backblaze's Toshiba Hard Drives

Reliability Data Set For 41,000 Hard Drives Now Open Source

16TB Drive Models with a Lifetime AFR < 1% as of Q1 2023

MFG	Model	Drive Count	Drive Days	Drive Failures	AFR	Confidence Interval
WDC	WUH721816ALE6L0	2,699	1,444,994	6	0.15%	0.30%
WDC	WUH721816ALE6L4	14,098	2,442,657	19	0.28%	0.30%
Toshiba	MG08ACA16TEY	5,288	2,439,952	36	0.54%	0.30%
Toshiba	MG08ACA16TA	5,194	1,090,335	19	0.64%	0.40%
Seagate	ST16000NM001G	21,910	9,584,565	222	0.85%	0.30%

14TB Drive Models with a Lifetime AFR < 1% as of Q1 2023

MFG	Model	Drive Count	Drive Days	Drive Failures	AFR	Confidence Interval
WDC	WUH721414ALE6L4	8,412	7,004,125	55	0.29%	0.20%
Toshiba	MG07ACA14TA	38,156	33,737,232	854	0.92%	0.10%

Honorable Mention

MFG	Model	Drive Count	Drive Days	Drive Failures	AFR	Confidence Interval
Seagate	ST14000NM001G	10,760	8,383,052	260	1.13%	0.60%

12TB Drive Models with a Lifetime AFR < 1% as of Q1 2023

MFG	Model	Drive Count	Drive Days	Drive Failures	AFR	Confidence Interval
HGST	HUH721212ALE600	2,606	3,274,145	27	0.31%	0.20%
HGST	HUH721212ALE604	13,151	9,549,226	130	0.47%	0.20%
HGST	HUH721212ALN604	10,690	15,719,532	290	0.53%	0.10%
Seagate	ST12000NM001G	12,805	10,718,697	226	0.76%	0.10%

Other Drive Models with a Lifetime AFR < 1% as of Q1 2023

MFG	Model	Size	Drive Count	Drive Days	Drive Failures	AFR	Confidence Interval
HGST	HMS5C4040BLE640	4TB	12,728	34,617,142	391	0.41%	0.10%
HGST	HMS5C4040ALE640	4TB	3,722	15,582,325	223	0.52%	0.10%
HGST	ST6000DX000	6TB	884	3,945,301	96	0.89%	0.30%
Seagate	HUH728080ALE600	8TB	1,116	2,071,099	47	0.83%	0.40%

Is the Bathtub Curve…

…Leaking?

Failure Rate Over Time: The Bathtub Curve

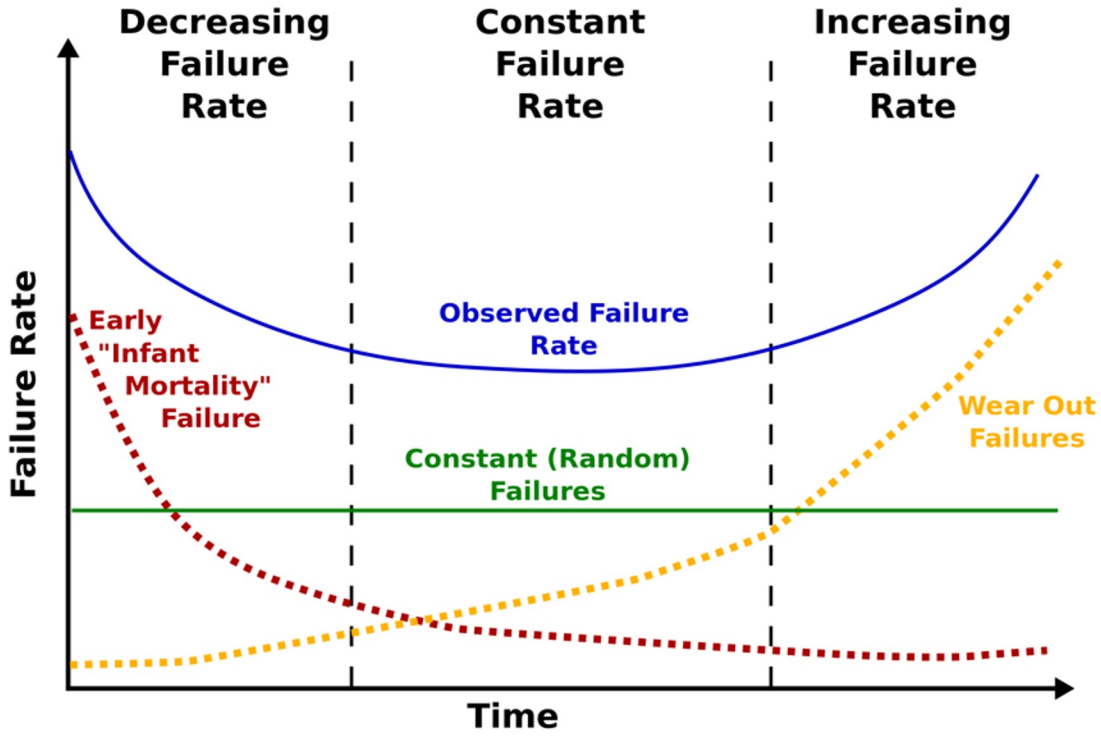

Drive Failure Over Time: 2016

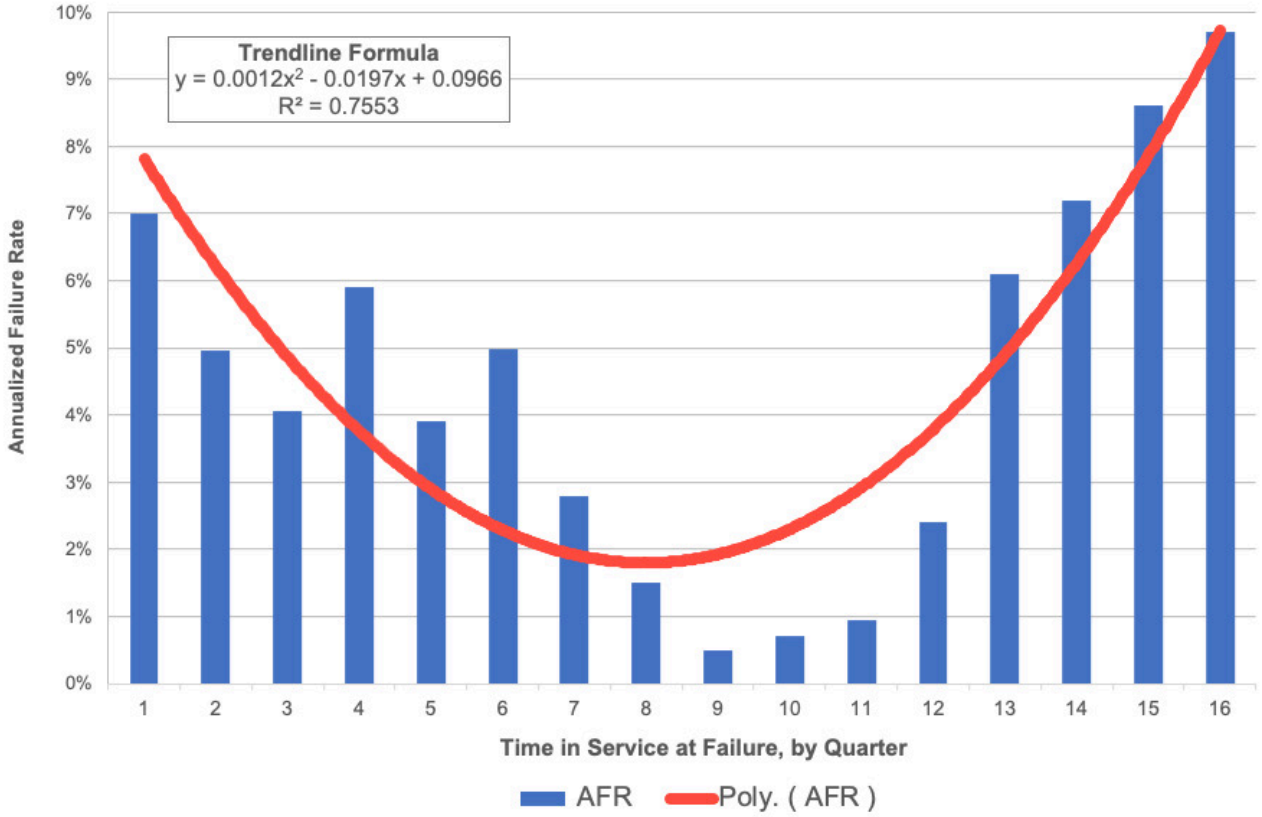

Drive Failure Over Time: 2016

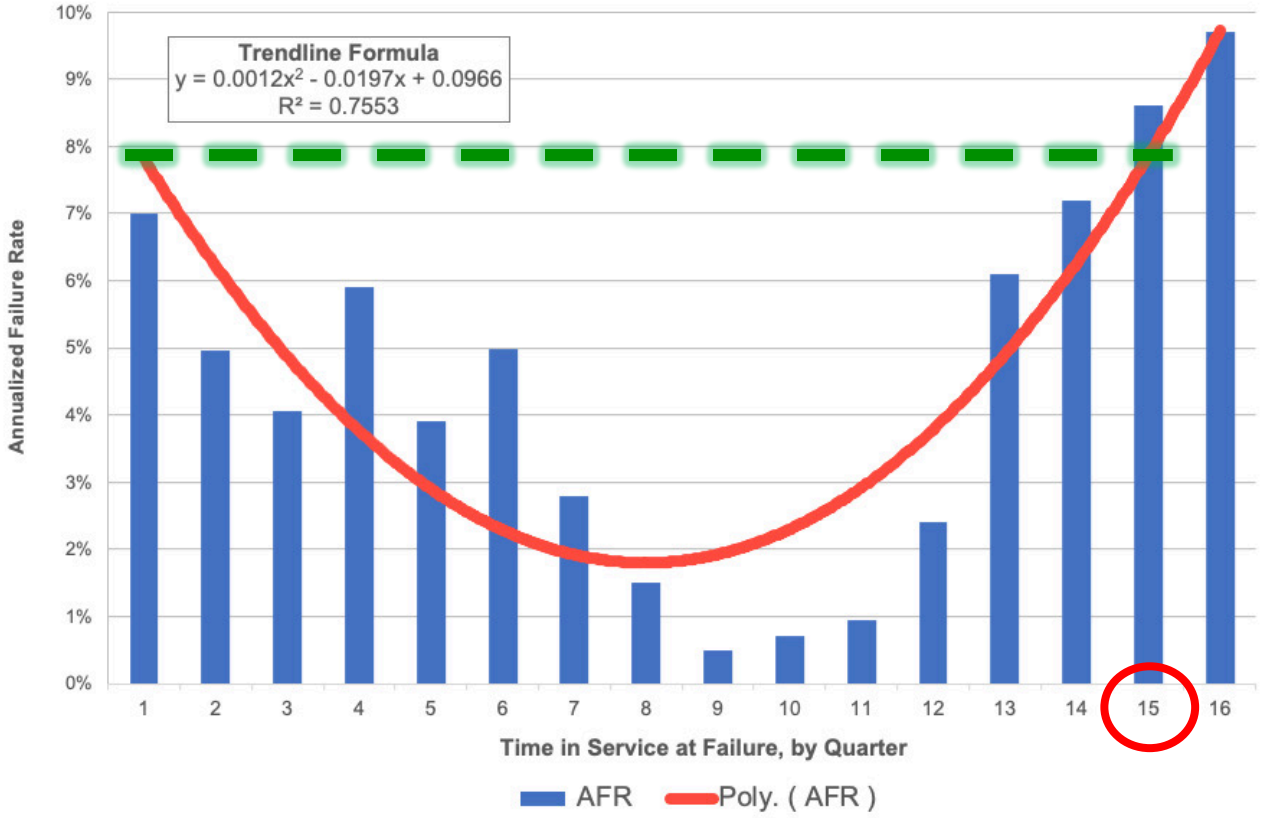

Drive Failure Over Time: 2022

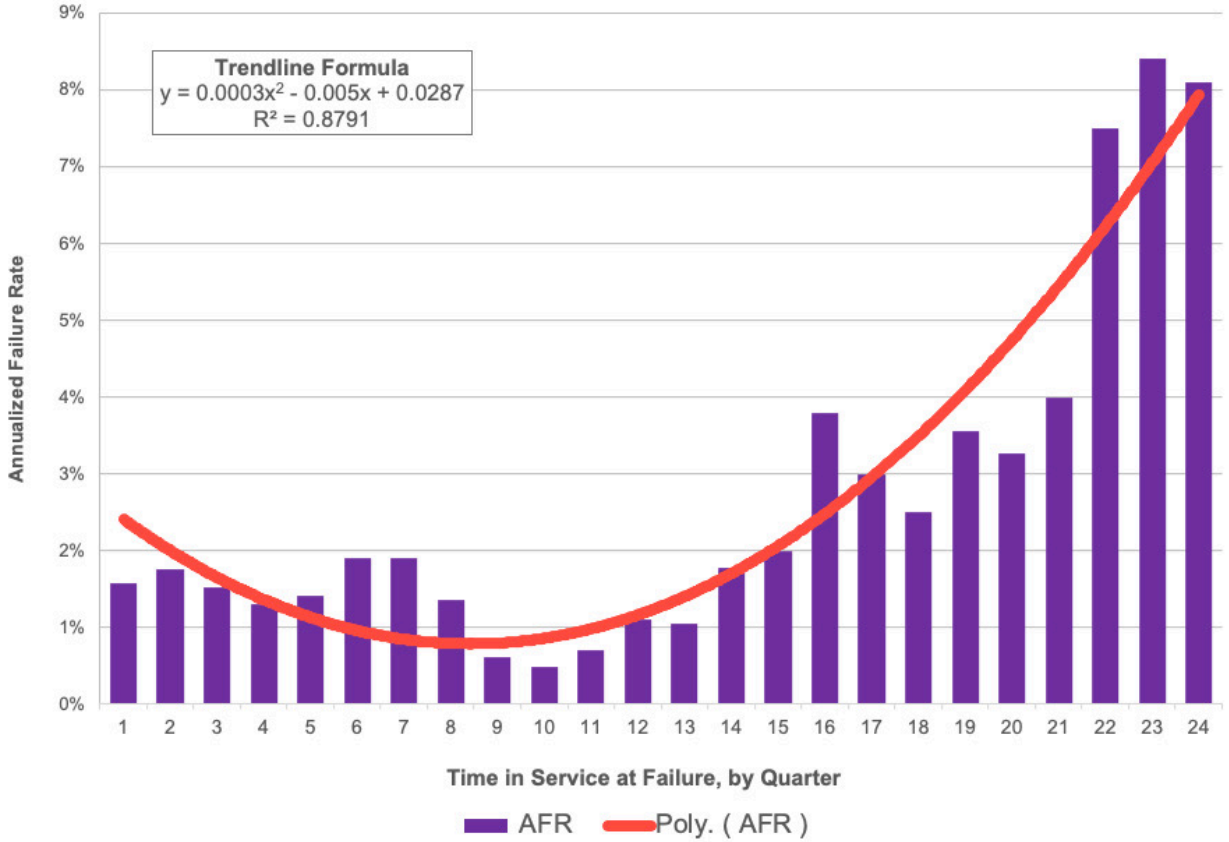

Drive Failure Over Time: 2022

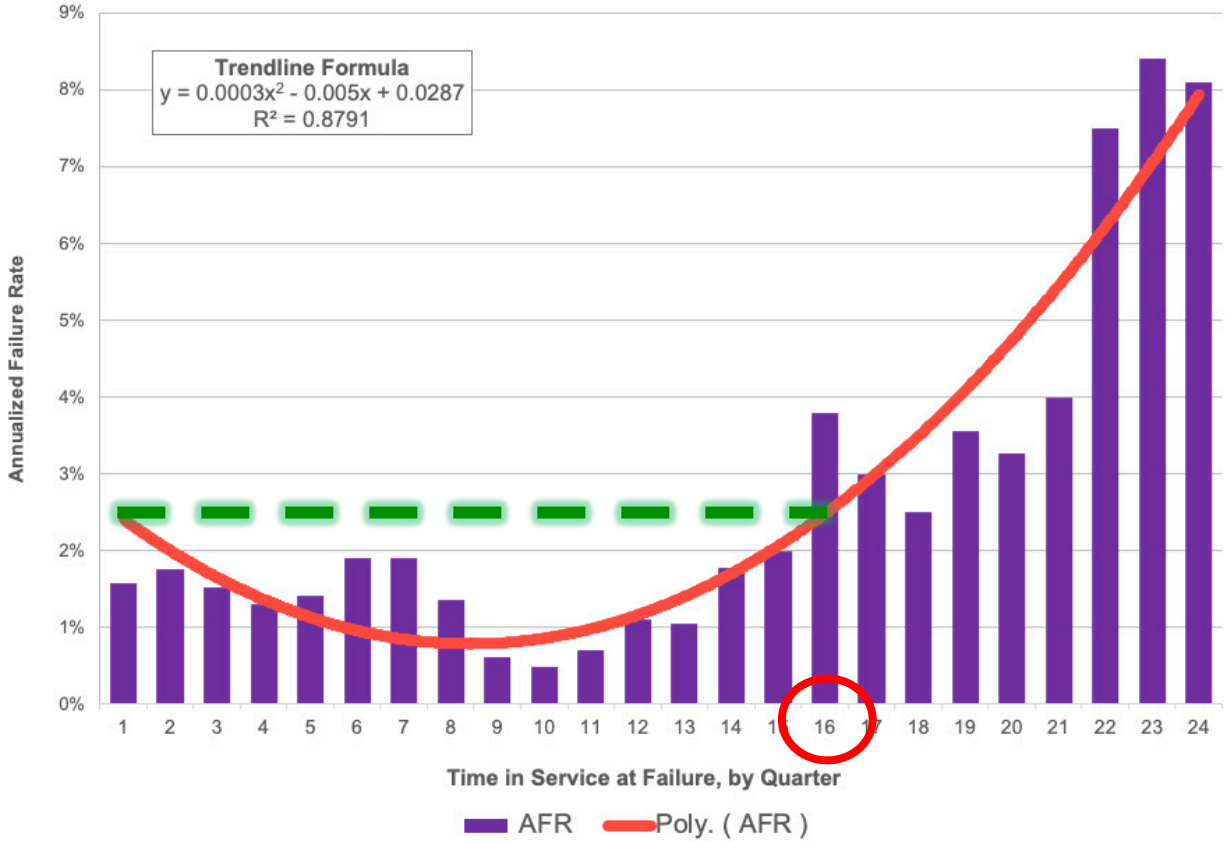

What is the Life Expectancy of a Hard Drive?

Using our drive stats data in combination with Kaplan-Meier curves has been done previously in various forms by others including Ross Lazarus, Simon Erni, and Tom Baldwin. We thank them for their collective efforts.

t	d	n	1-d/n	S(t)		t	S(t)	N
0		38664		1		0	1	1
1	31	38586	0.9991966	0.9991966		1	1	1
2	21	38532	0.999455	0.99865204		1	0.9991966	2
3	27	38489	0.9992985	0.99795148		2	0.9991966	2
4	20	38436	0.99947965	0.9974322		2	0.99865204	3
5	29	38376	0.99924432	0.99667846		3	0.99865204	3
6	20	38201	0.99947645	0.99615666		3	0.99795148	4
7	19	37770	0.99949696	0.99565554		4	0.99795148	4
8	12	33426	0.999641	0.9952981		4	0.9974322	5
9	22	32241	0.99931764	0.99461895		5	0.9974322	5
10	5	31303	0.99984027	0.99446008		5	0.99667846	6
11	19	30043	0.99936757	0.99383116		6	0.99667846	6
12	18	27477	0.99934491	0.9931801		6	0.99615666	7
13	10	26223	0.99961866	0.99280136		7	0.99615666	7
14	10	24220	0.99958712	0.99239145		7	0.99565554	8

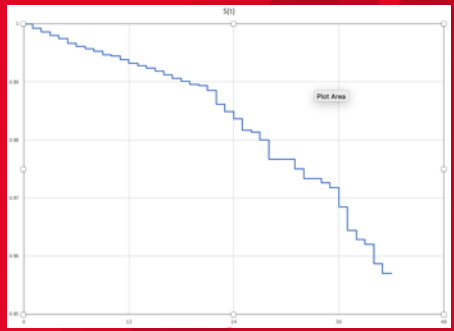

Hard Drive Life Expectancy (Survival Curves)

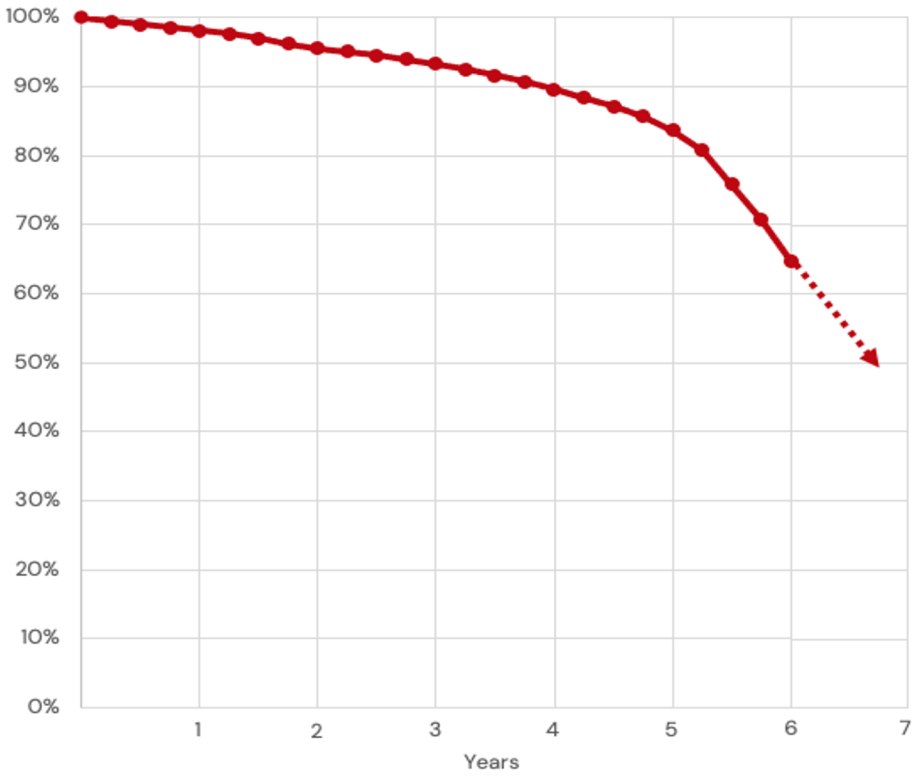

Kaplan-Meier Projection

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

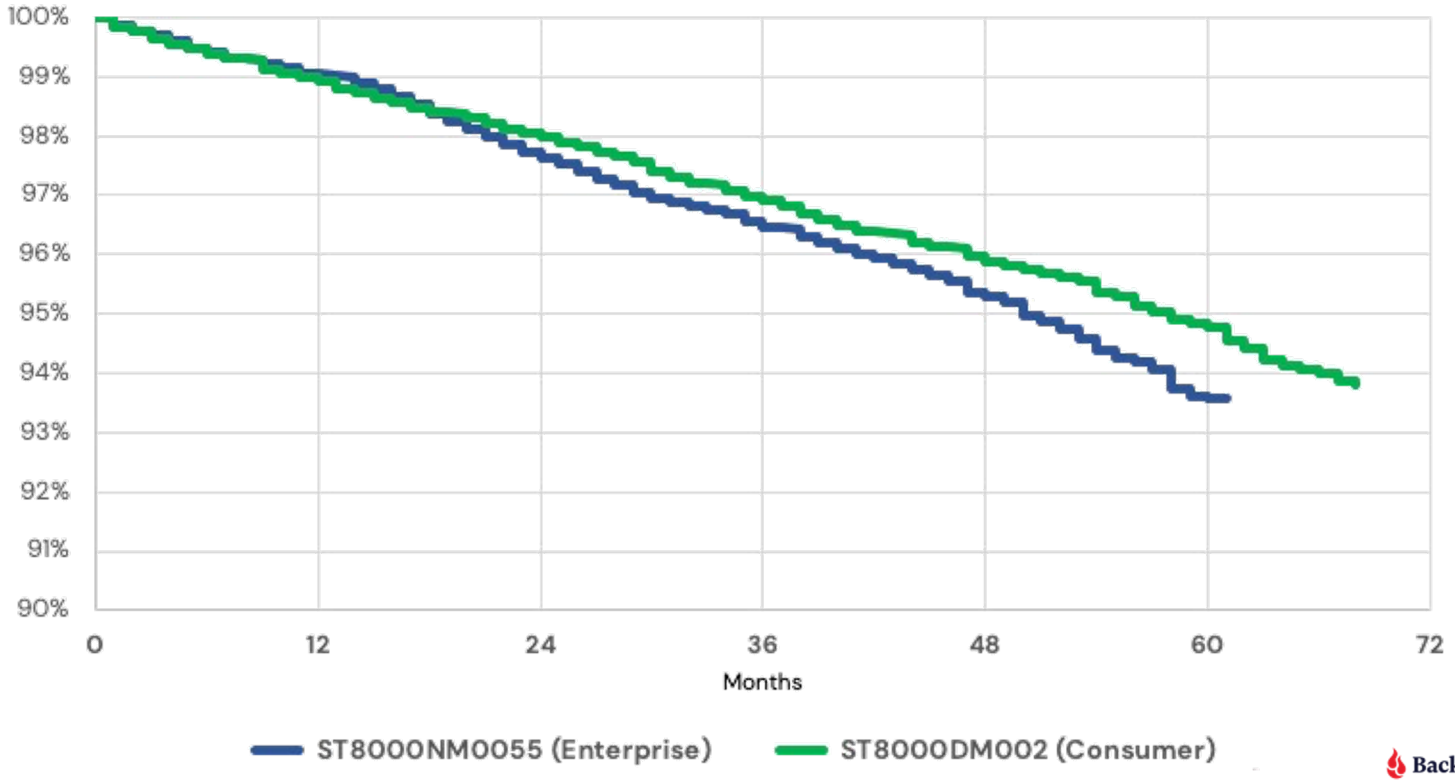

Hard Drive Life Expectancy – Select 12TB Drives

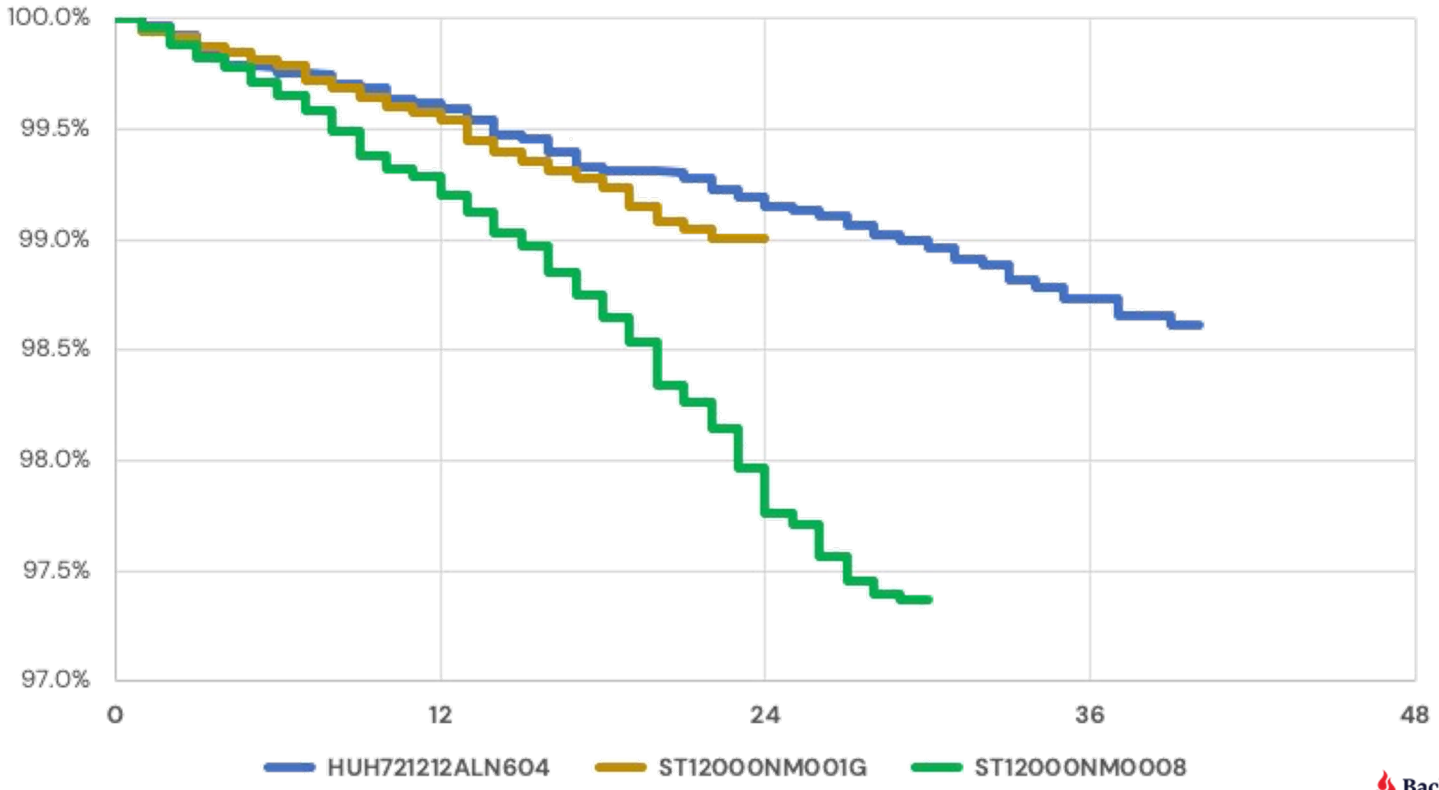

Hard Drive Life Expectancy – Select 14TB Drives

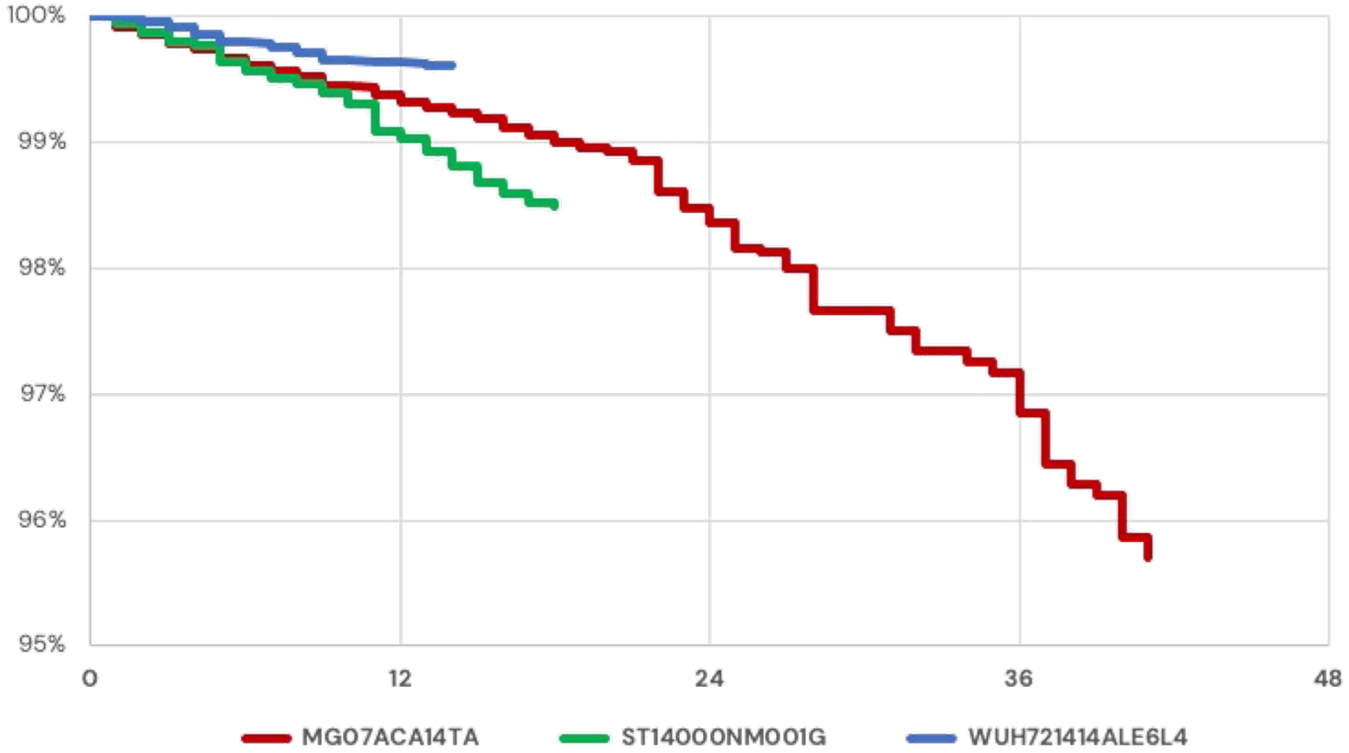

Predicting Hard Drive Failure

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

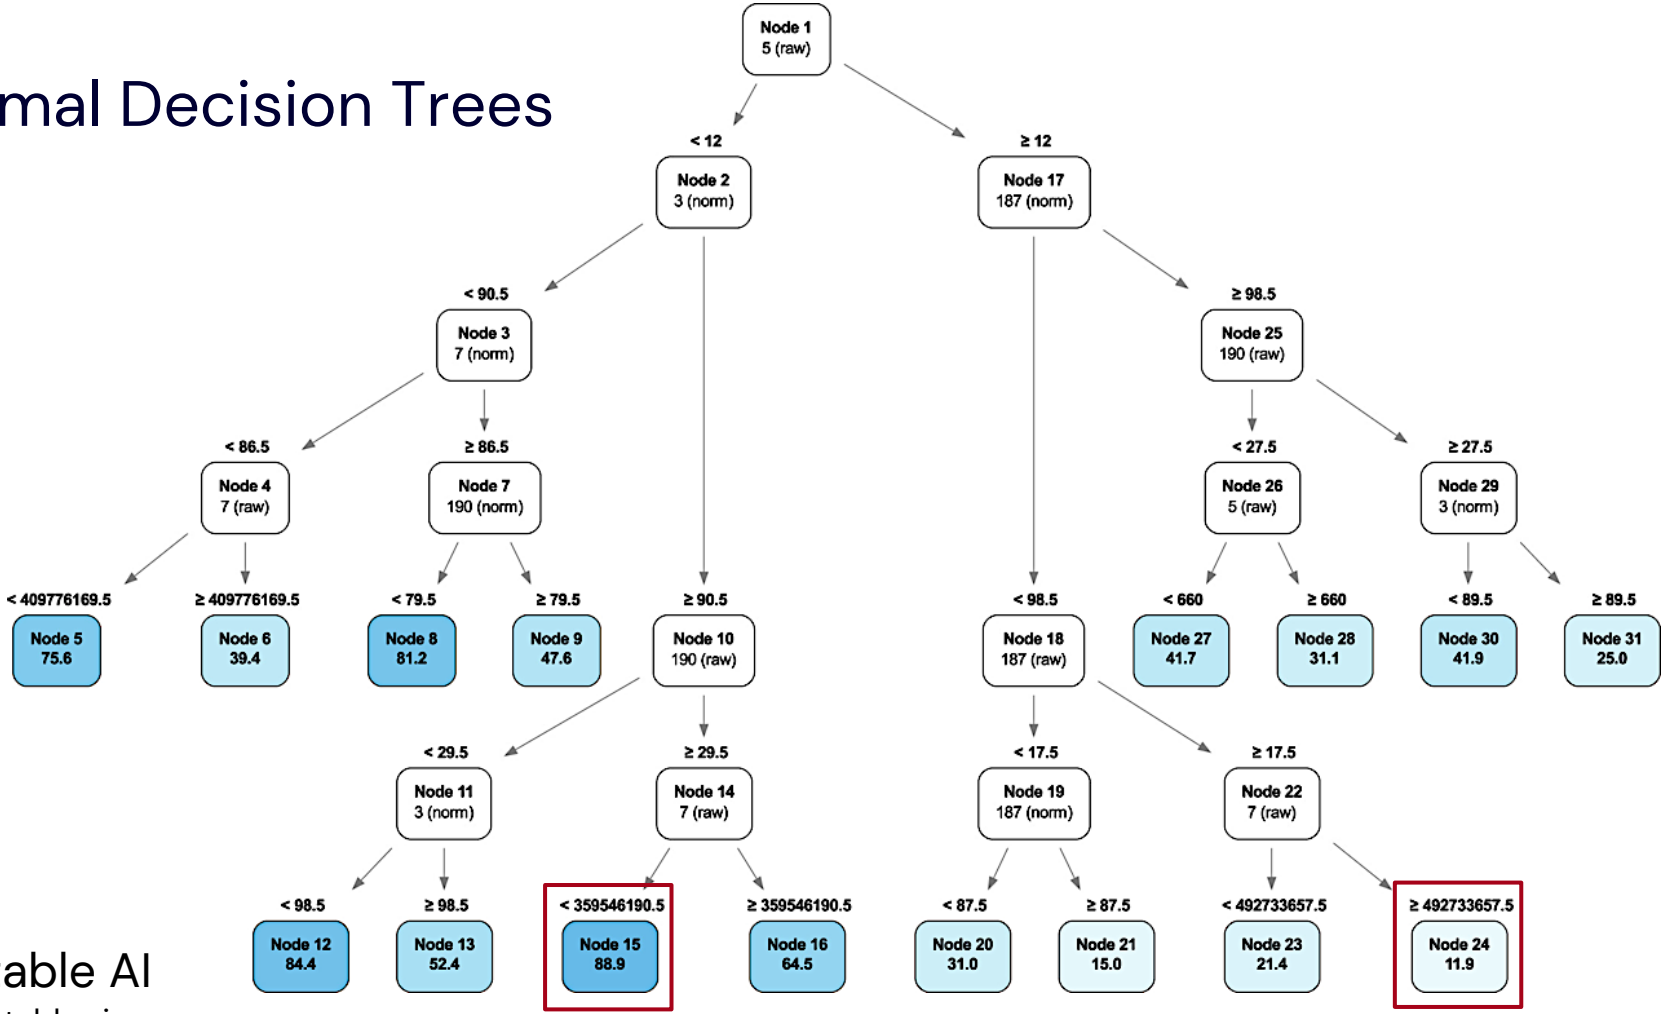

Predicting Short-term Drive Health

Example survival curves for selected cohorts

Healthy Drives

Node 15

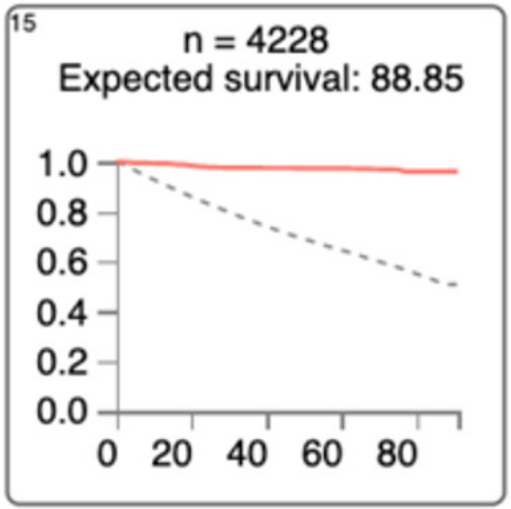

5 (raw) < 12

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

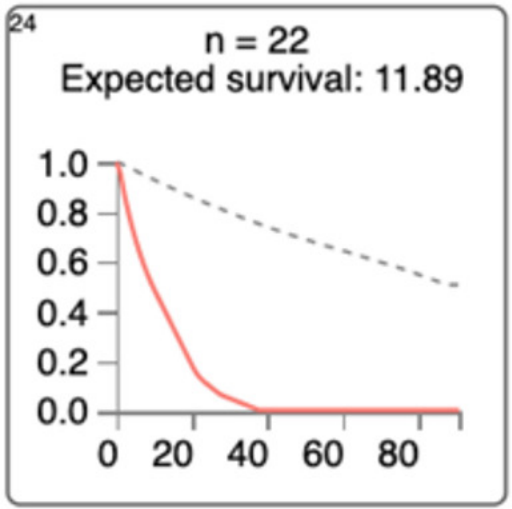

5 (raw) >= 12

187 (norm) < 98.5

187 (raw) >= 17.5

7 (raw) >= 492733657.5

Summary

- Since we last met at MSST
- The data we collect
- Let’s talk about drive failure
- The bathtub curve lives on
- Let’s talk about drive survival
- Using Machine Learning to predict drive failure
- If we have time: Bonus content

Average Temperature versus Drive Size

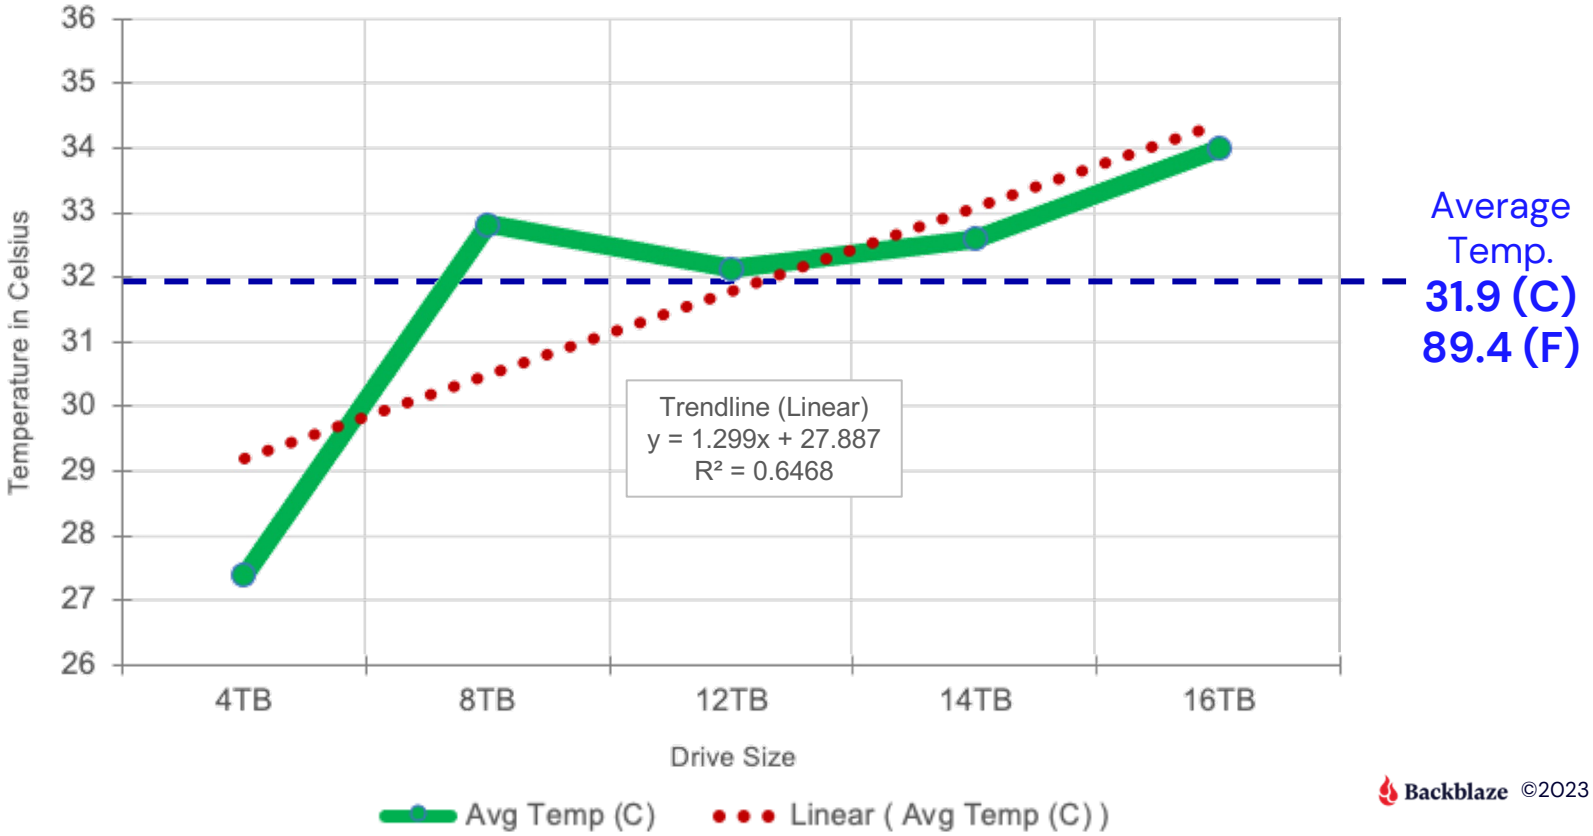

AFR for SSD and HDD Boot Drives

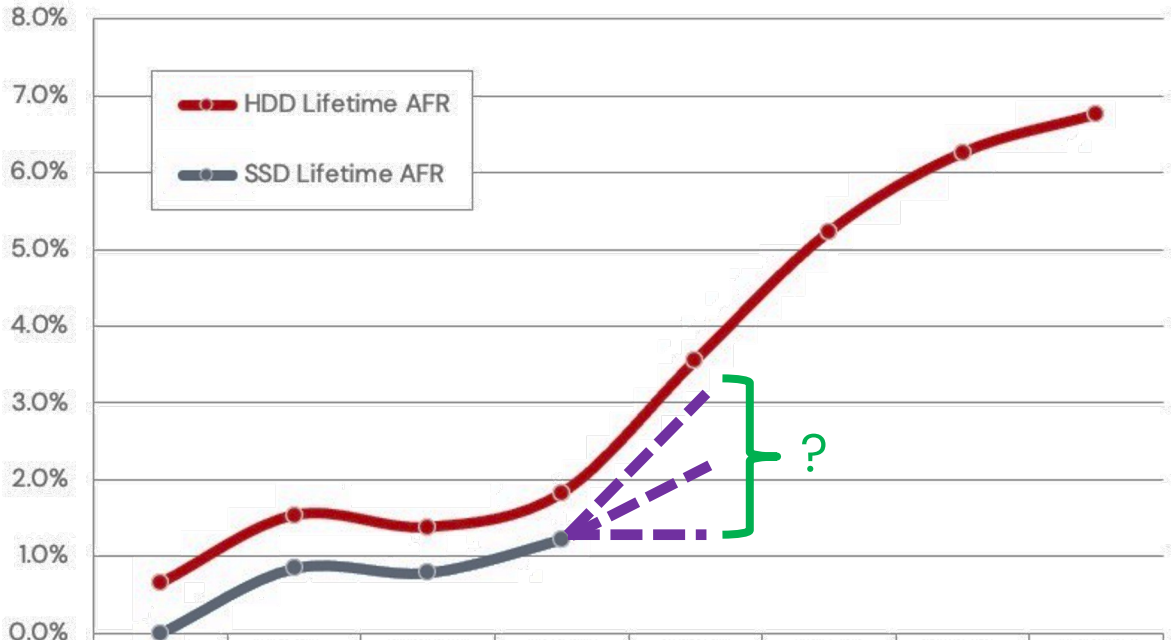

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8
HDD Lifetime AFR	0.66%	1.54%	1.38%	1.83%	3.55%	5.23%	6.26%	6.76%
SSD Lifetime AFR	0.00%	0.84%	0.79%	1.22%				

AFR for SSD and HDD Boot Drives

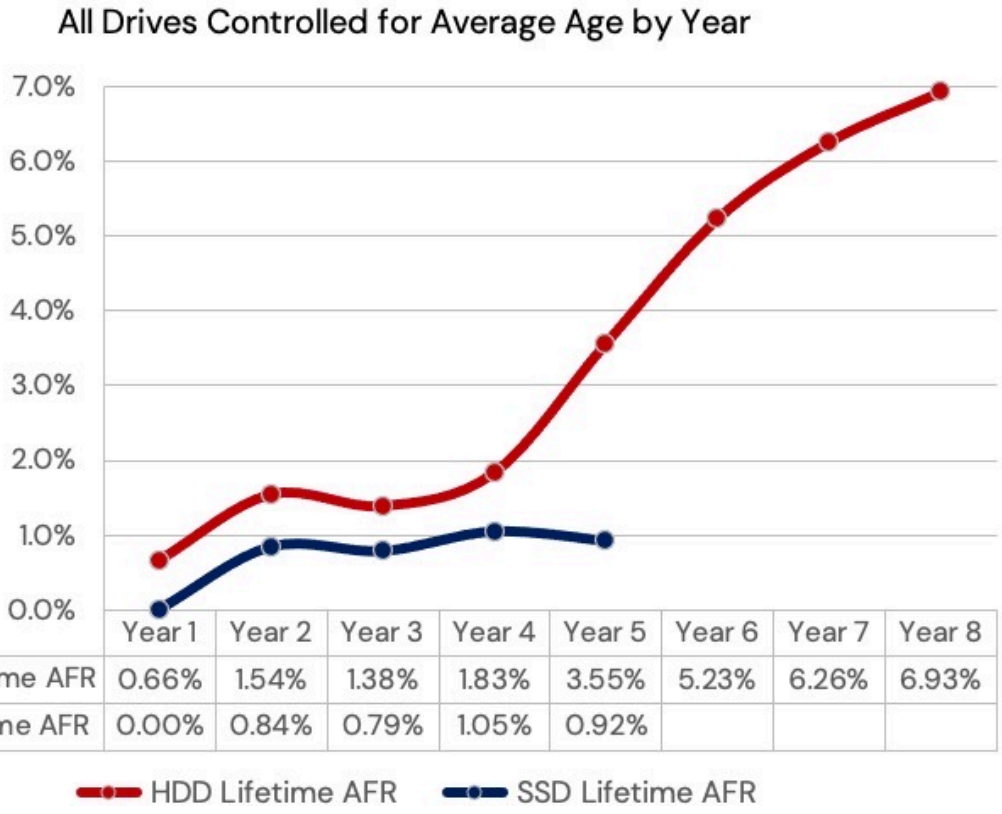

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8
HDD Lifetime AFR	0.66%	1.54%	1.38%	1.83%	3.55%	5.23%	6.26%	6.93%
SSD Lifetime AFR	0.00%	0.84%	0.79%	1.05%	0.92%			

Backblaze

Thank You

Andy Klein | Backblaze