

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

Virtual Conference  
September 28-29, 2021

# Facts, Figures and Insights from 250,000 Hard Drives

Andrew Klein, Storage Cloud Evangelist, Backblaze

# Overview

- Where do these 250,000 drives live and work
- What is a drive failure?
- Drive failure and...
  - Power cycling
  - Time
  - Temperature
- Predicting drive failure

**1.8 Exabytes**

Available data storage

**4 data centers**

California (2), Arizona & Holland

**178,166**

Active HDD data drives

**260,461**

Total HDD data drives

**246,100,245**

Lifetime drive days



# Where the drives work

Speaker  
Photo Will  
Be Placed  
Here

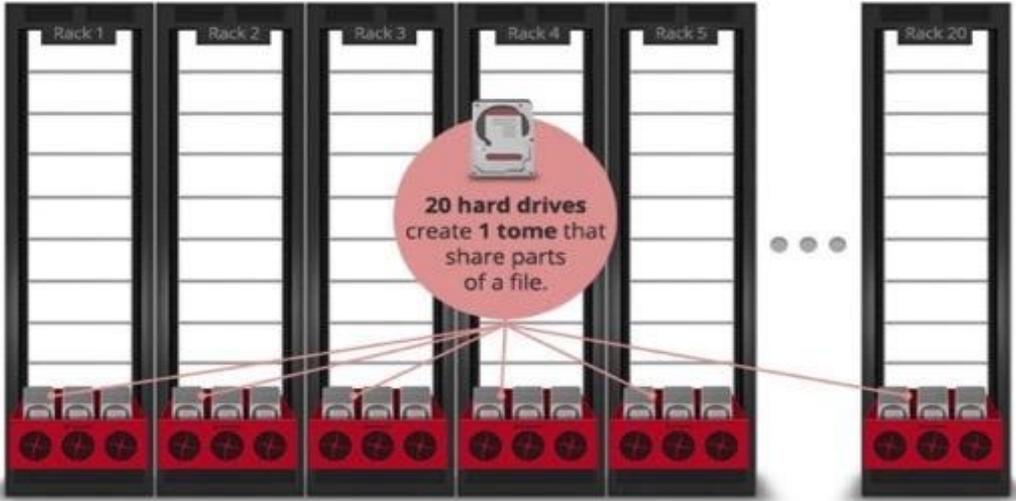
Storage Server



60 drives

\* 20 =

Backblaze Vault



20 drives is 1 tome \* 60 tomes in a 1 vault  
= 1,200 drives in 1 vault

# Collecting Drive Data

- Use smartmontools package to collect S.M.A.R.T. data
  - <https://www.smartmontools.org>
- Collect data once a day from each drive
- Started keeping data in April 2013
- Open Sourced the data in February 2015
  - <https://www.backblaze.com/b2/hard-drive-test-data.html>

Speaker  
Photo Will  
Be Placed  
Here

# Drive Data Collected Each Day

date	serial_number	model	capacity_bytes	failure	Smart_1_normalized	Smart_1_raw
6/7/21	Z305B2QN	ST4000DM000	4000787030016	0	98	2766
6/7/21	PL1331LAHG1S4H	HGST HMS5C4040ALE640	4000787030016	0	100	0
6/7/21	ZACH007	ST8000NM0055	8001563222016	1	81	139015
6/7/21	ZA130TTW	ST8000DM002	8001563222016	0	83	100901
6/7/21	ZA18CEBF	ST8000NM0055	8001563222016	0	81	140551
6/7/21	PL2331LAH3WYAJ	HGST HMS5C4040BLE640	4000787030016	0	100	0

More stats



Speaker Photo Will Be Placed Here

**Smart Stats:** There are 255 pairs of values per drive. Examples:

- Smart\_1: Read Error Rate
- Smart\_5: Reallocated Sector Count
- Smart\_9: Power On Hours

[en.wikipedia.org/wiki/S.M.A.R.T.](https://en.wikipedia.org/wiki/S.M.A.R.T.)

More drives



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

# What is Drive Failure?

- The drive will not spin up or connect to the OS.
- The drive will not sync, or stay synced in a storage array.
- The statistics we track show persistent values above our thresholds.



Speaker  
Photo Will  
Be Placed  
Here

# Two Types of Failures

## Reactive Failure

- We react to the failure

## Proactive Failure

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

## In Either Case

- Removed drives are quarantined until they are no longer needed.
  - Cloning, fault analysis, etc.

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

Speaker  
Photo Will  
Be Placed  
Here

# Annualized Failure Rate

1. Define AFR cohort and period:
  - a. Cohort = Model ABC123
  - b. Period = Q2 2021
2. Obtain Drive Days and Drive Failures for the cohort and period.
  - a. Drive Days = 409,124
  - b. Drive Failures = 17
  - c. Drive Count = 5,000
3. Apply Formula:  $AFR = (\text{Drive Failures} / (\text{Drive Days} / 365)) * 100$

$$AFR = (17 / (409,124 / 365)) * 100 = 1.52\%$$

Speaker  
Photo Will  
Be Placed  
Here

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

# Drive Failure and ...

Power cycling, time, temperature...

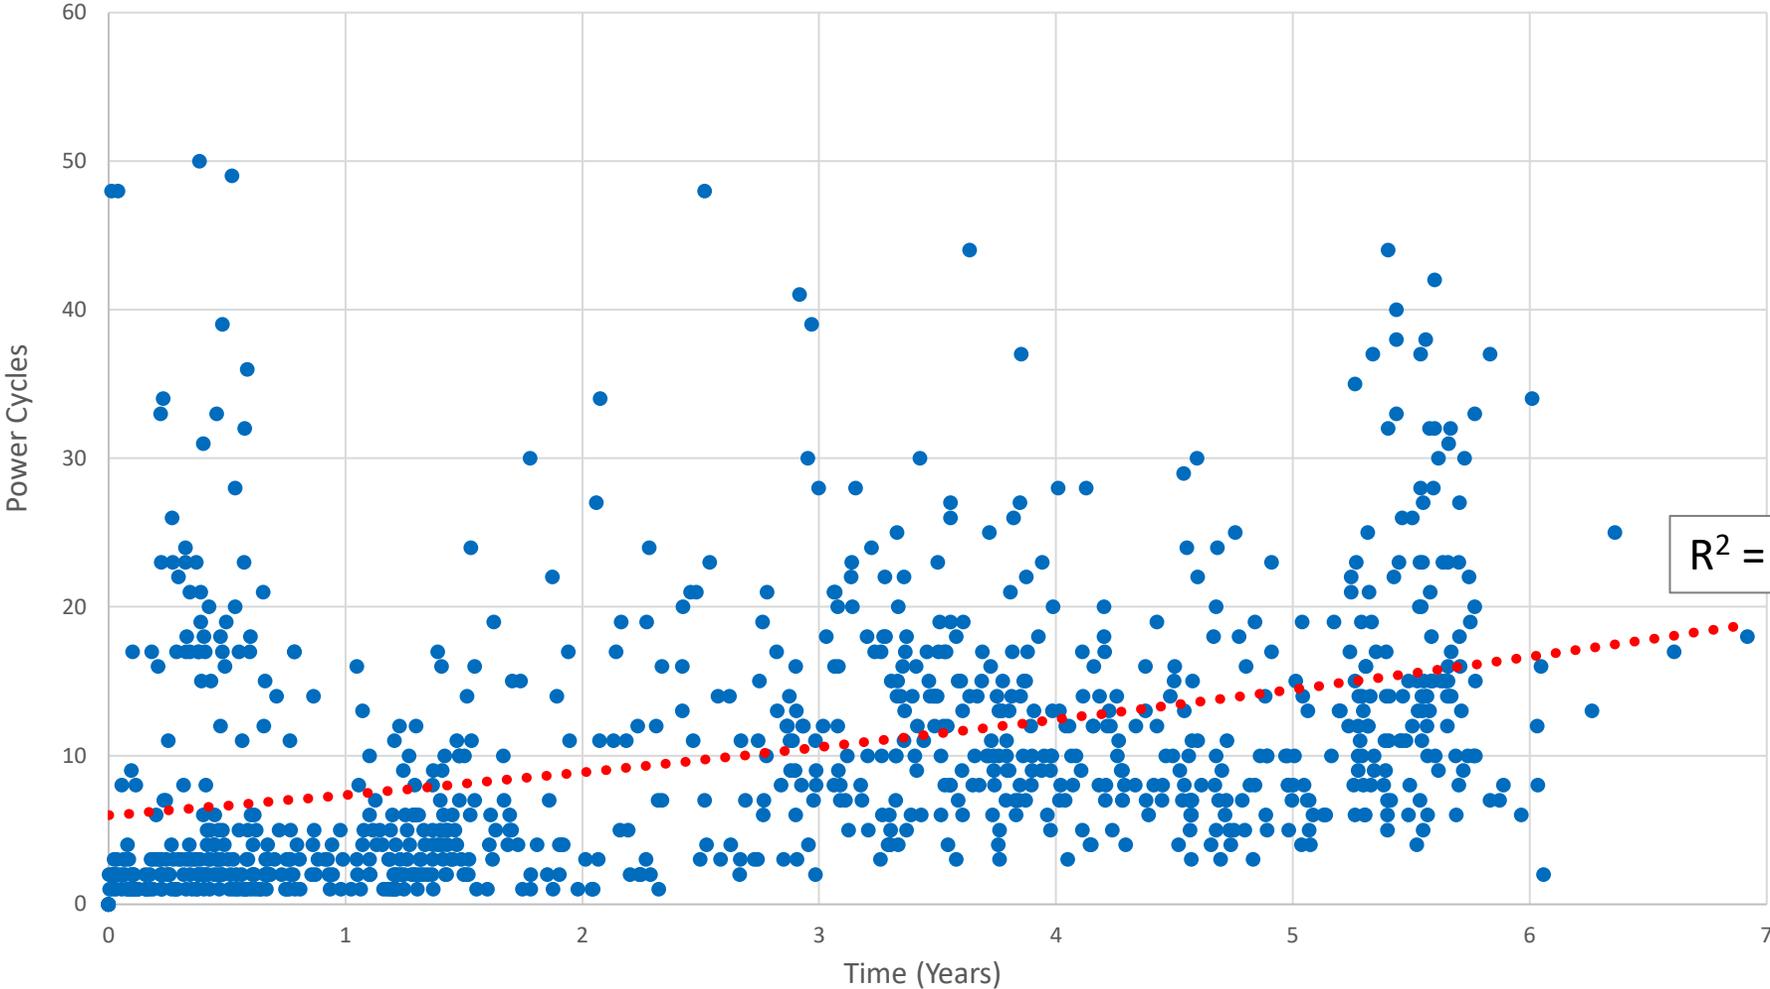
# Power Cycling

## Average Number of Power Cycles for Each Drive

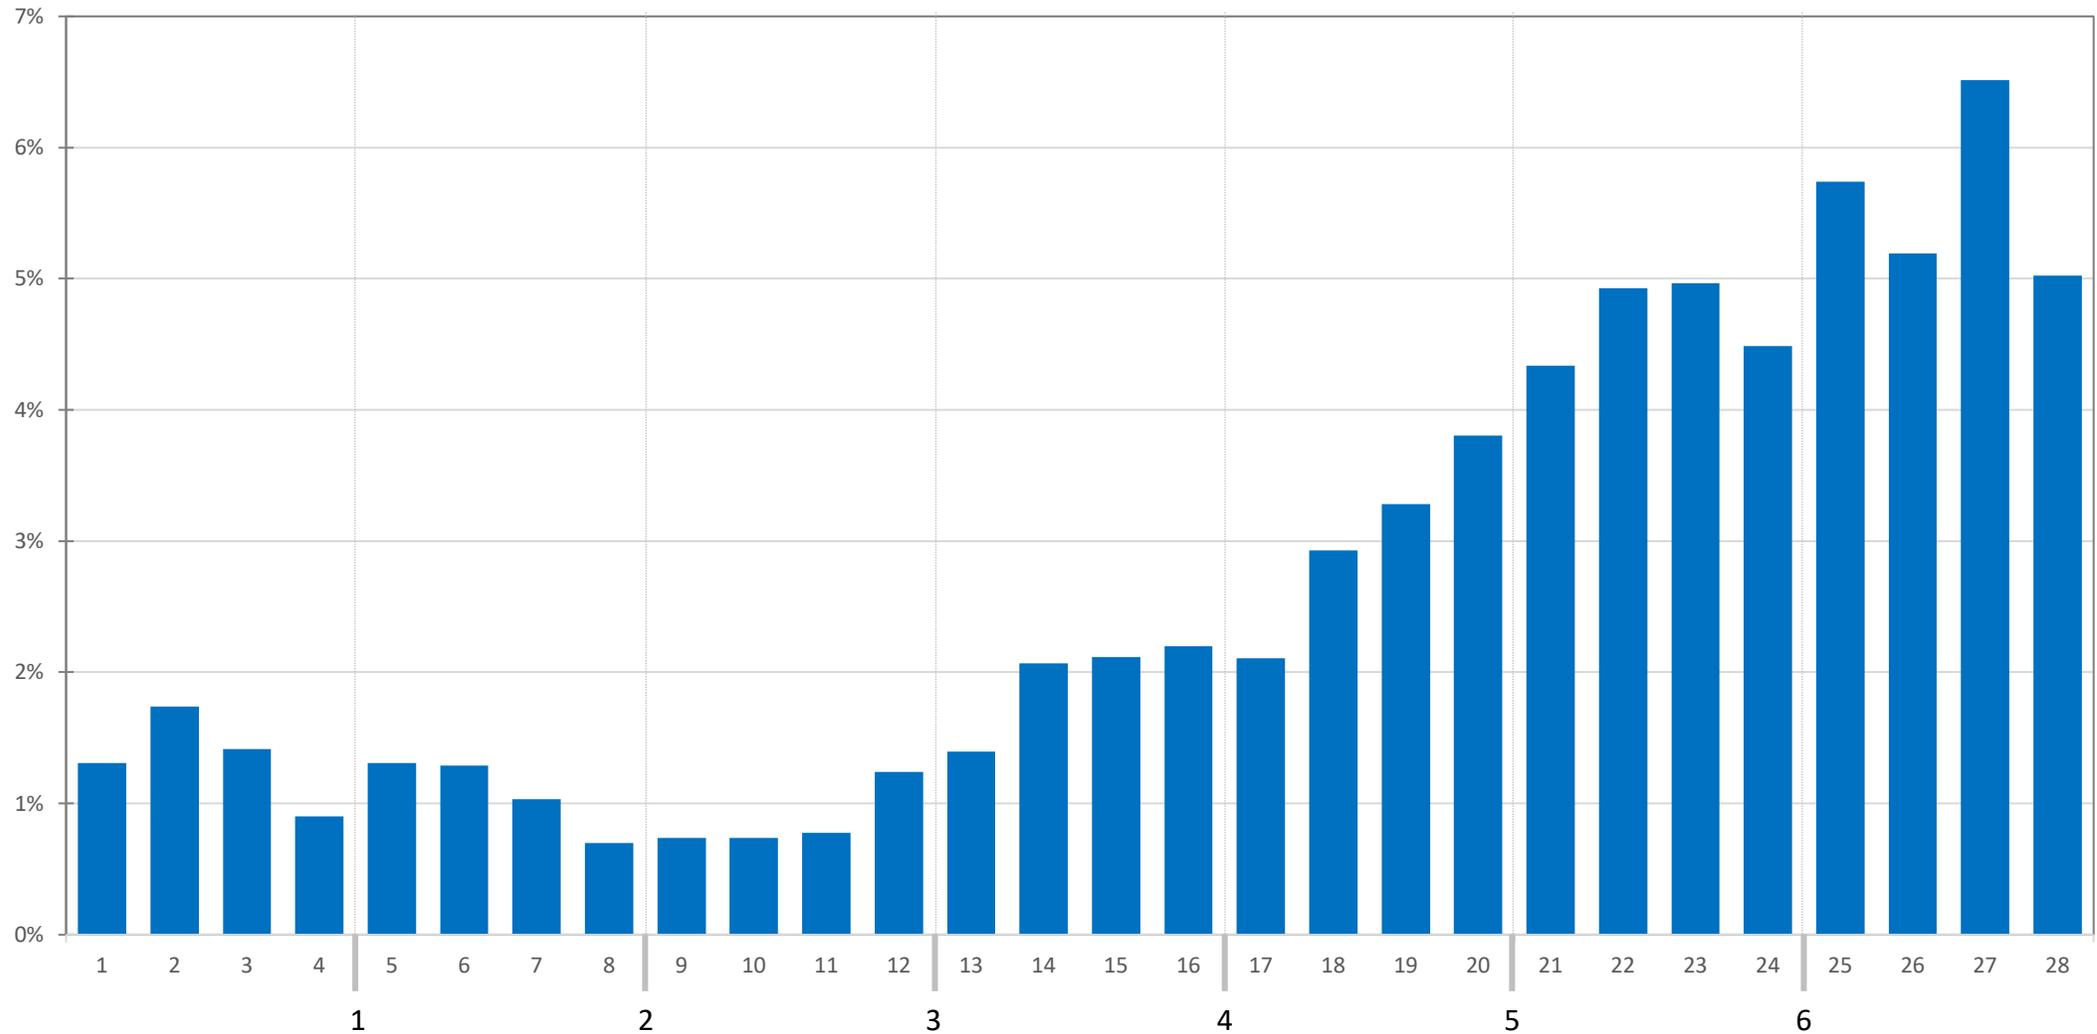


# Power Cycling

Number of Power Cycles a Failed Drive Experienced over Time

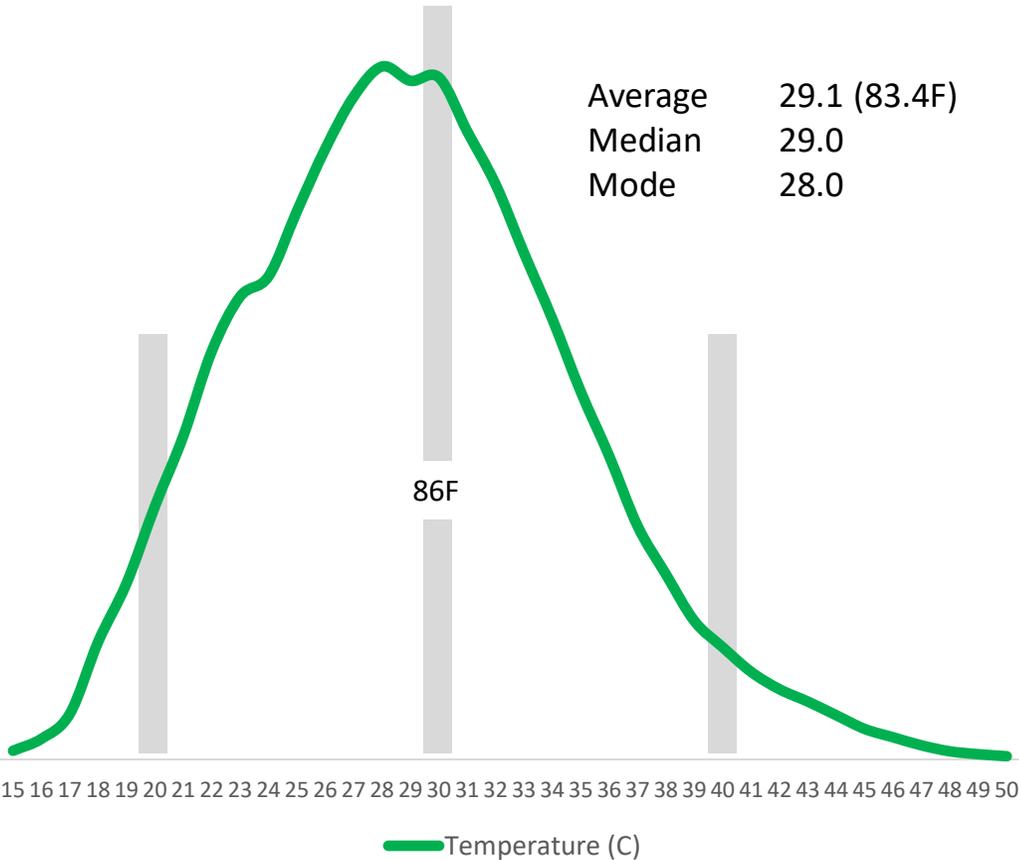


# Annualized Failure Rate Over Time (Quarters)

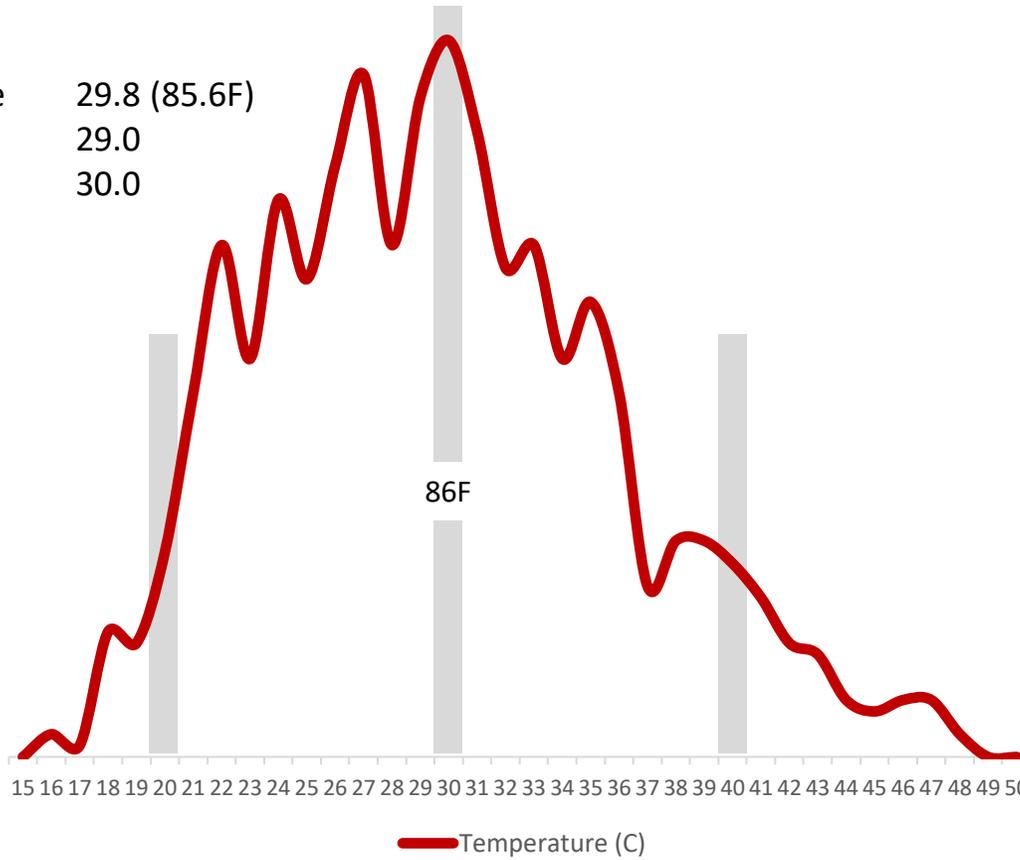


# Temperature of Operational versus Failed Drives

Operational (Good) Drives



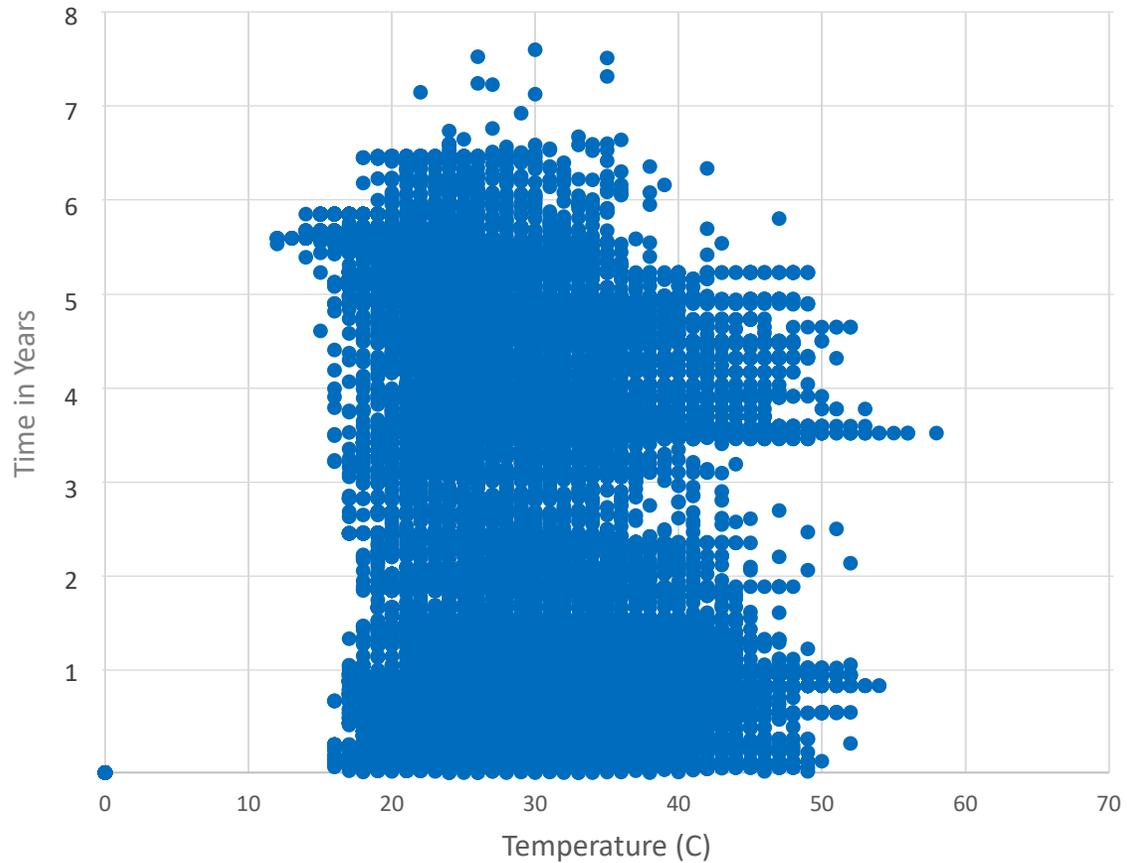
Failed Drives



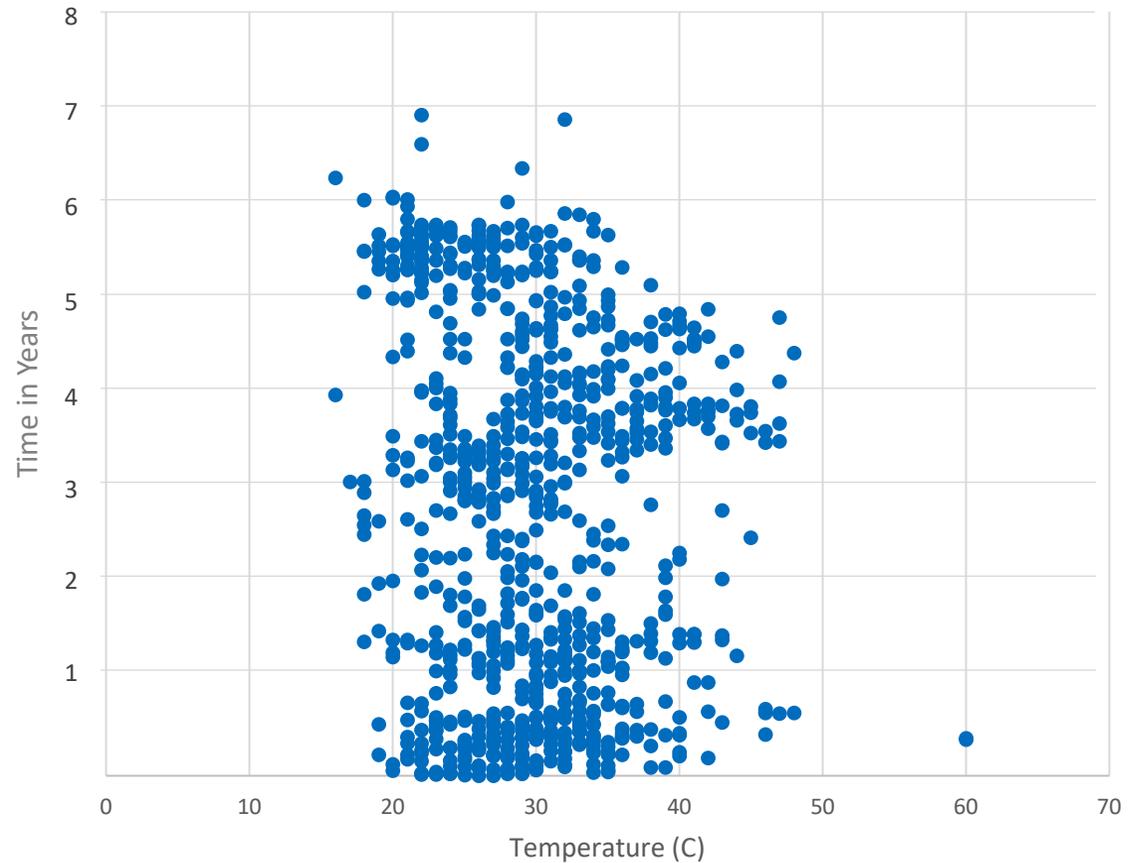
AFR for observation period: 1.04%

# Relationship of Drive Temperature to Drive Age

Operational (Good) Drives

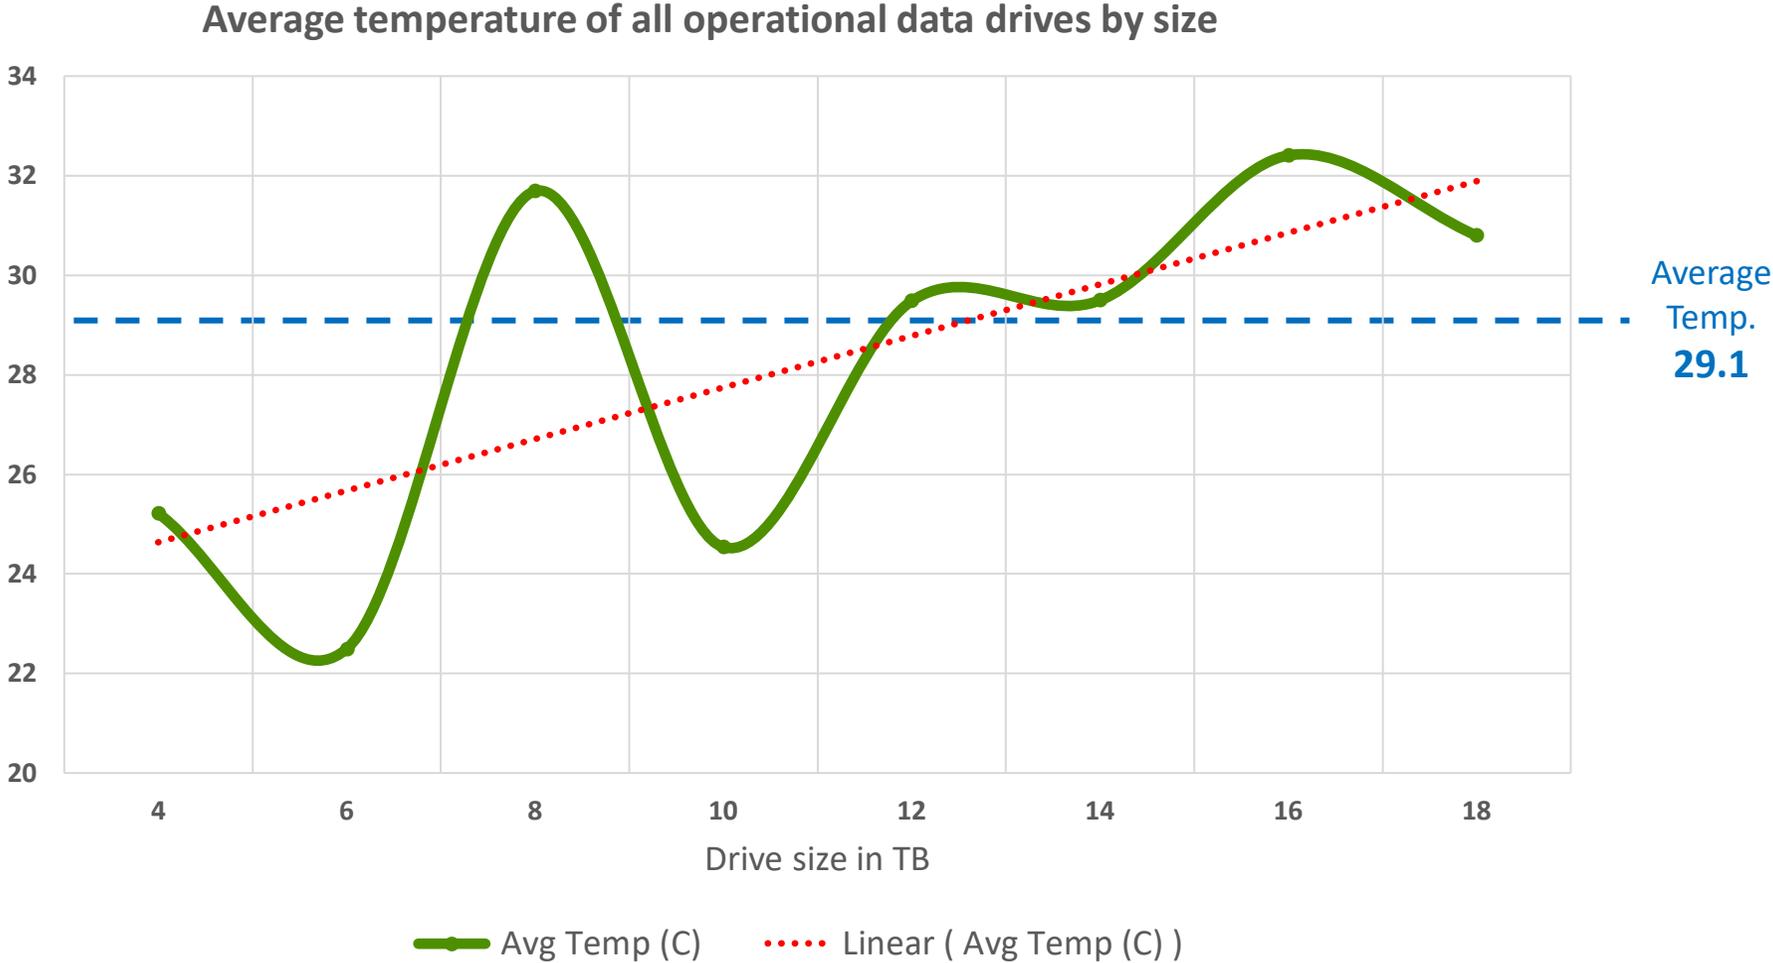


Failed Drives



AFR for observation period: 1.04%

# Temperature versus Drive Size



# Predicting Drive Failure

Yesterday and Today

# Back in 2017

Speaker  
Photo Will  
Be Placed  
Here

## SMART Attributes as Failure Detection

Percentage of drives with SMART attribute RAW value > 0

Drive Status	SMART 5 Reallocated Sectors Count	SMART 187 Reported Uncorrectable Errors	SMART 188 Command Timeout	SMART 197 Current Pending Sector Count	SMART 198 Uncorrectable Sector Count
Operational	1.1%	0.5%	4.8%	0.7%	0.3%
Failed	42.2%	43.5%	44.8%	43.1%	33.0%

1) Failed drives as of one day prior to being marked as failure



2017 Storage Developer Conference. © 2017 Backblaze Inc. All Rights Reserved.



22

# 2021 versus 2017

Speaker Photo Will Be Placed Here

2017

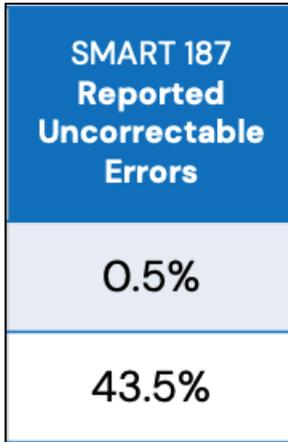
Drive Status	SMART 5 Reallocated Sectors Count	SMART 187 Reported Uncorrectable Errors	SMART 188 Command Timeout	SMART 197 Current Pending Sector Count	SMART 198 Uncorrectable Sector Count
Operational	1.1%	0.5%	4.8%	0.7%	0.3%
Failed	42.2%	43.5%	44.8%	43.1%	33.0%

2021

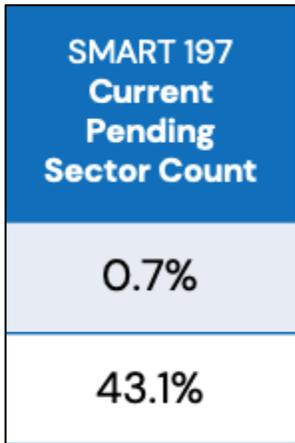
Drive Status	SMART 5 Reallocated Sectors Count	SMART 187 Reported Uncorrectable Errors	SMART 188 Command Timeout	SMART 197 Current Pending Sector Count	SMART 198 Uncorrectable Sector Count
Operational	2.7%	23.5%	0.0%	99.8%	0.5%
Failed	47.9%	57.1%	10.9%	43.1%	26.1%



# Voted off the Island



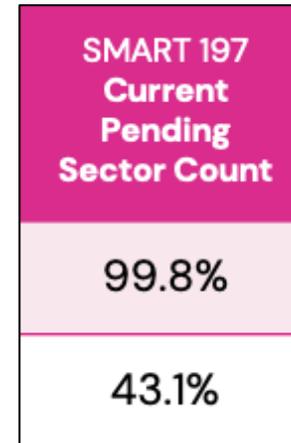
2017



- Only reported by Seagate 4TB drives
- Does not decline/reset



2021



- Most larger drives report, but value is the same as SMART 1
- Non-sensible values for SMART 197

Speaker  
Photo Will  
Be Placed  
Here

# 2021 versus 2017

Speaker Photo Will Be Placed Here

2017

Drive Status	SMART 5 Reallocated Sectors Count	SMART 187 Reported Uncorrectable Errors	SMART 188 Command Timeout	SMART 197 Current Pending Sector Count	SMART 198 Uncorrectable Sector Count
Operational	1.1%	0.5%	4.8%	0.7%	0.3%
Failed	42.2%	43.5%	44.8%	43.1%	33.0%

2021

Drive Status	SMART 5 Reallocated Sectors Count	SMART 187 Reported Uncorrectable Errors	SMART 188 Command Timeout	SMART 197 Current Pending Sector Count	SMART 198 Uncorrectable Sector Count
Operational	2.7%	23.5%	0.0%	99.8%	0.5%
Failed	47.9%	57.1%	10.9%	43.1%	26.1%



# Predicting Drive Failure with SMART Stats

## ■ 2016

### ■ Predicting Disk Replacement towards Reliable Data Centers.

- Botezatu, Mirela & Giurgiu, Ioana & Bogojeska, Jasmina & Wiesmann, Dorothea. (2016).
- <https://dl.acm.org/doi/10.1145/2939672.2939699>

## ■ 2021

### ■ Interpretable Predictive Maintenance for Hard Drives

- Maxime Amram, Jack Dunn, Jeremy J. Toledano, Ying Daisy Zhuo
- <https://www.sciencedirect.com/science/article/pii/S2666827021000219>
- Optimized Decision Trees

# Summary

- Where do these 250,000 drives live and work
- What is a drive failure?
- Drive failure and...
  - Power cycling
  - Time
  - Temperature
- Predicting drive failure



Thank You

Please take a moment to rate this session.

Your feedback is important to us.