What Can One Billion Hours of Spinning Hard Drives Tell Us?

What	Storage Developer Conference
When	September 2016
Who	Gleb Budman, CEO
	@GlebBudman

Overview

- Our environment
- How we diagnose “sick” drives
- Drive Reliability Findings
  1. Reliability over time
  2. Enterprise vs. consumer drives
  3. Effects of temperature
  4. Effects of power cycling
- Review and Questions

Introducing Backblaze

B2 Cloud Storage

The lowest cost cloud storage on the planet: $0.005/GB a month.
Try it and get the first 10 GB free on us.

Email

Password

Sign Up for B2

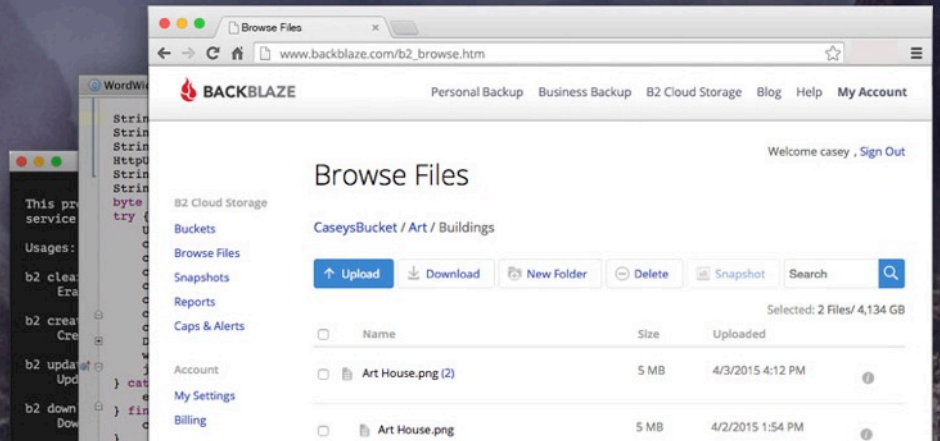

Our Environment

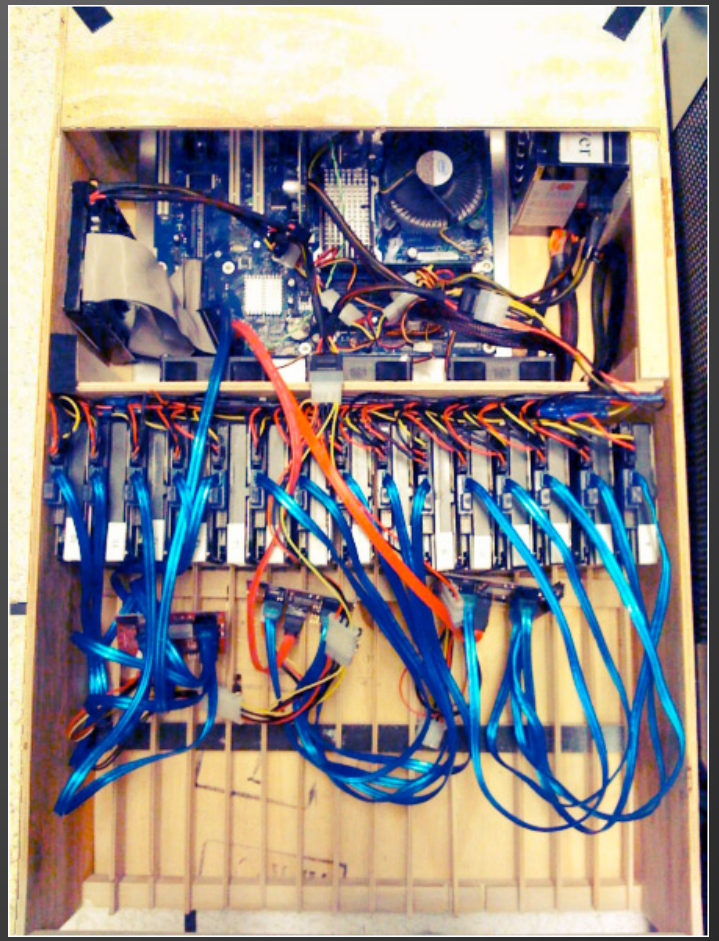

Once upon a time…

…Storage Pod 1.0

Speed bumps…

…Success too

Backblaze Vaults

Storage Pod v6.0


250
PETABYTES

How do we diagnose “sick” drives?

Defining a Drive Failure

- The drive will not spin up or connect to the OS.
- The drive will not sync, or stay synced, in a RAID Array (see note below).
- The SMART Stats we use show values above our thresholds.

- Note: Stand-alone Storage Pods use RAID-6. Backblaze Vaults use our own open-sourced implementation of Reed-Solomon erasure coding instead. Both techniques have a concept of a drive not syncing or staying synced with the other member drives in its group.

SMART attributes we use for failure detection

Attribute	Description	Reported by
SMART 5	Reallocated Sectors Count	All
SMART 187	Reported Uncorrectable Errors	Seagate
SMART 188	Command Timeout	Seagate
SMART 197	Current Pending Sector Count	All
SMART 198	Uncorrectable Sector Count	Seagate

Detection rates of SMART attributes in aggregate

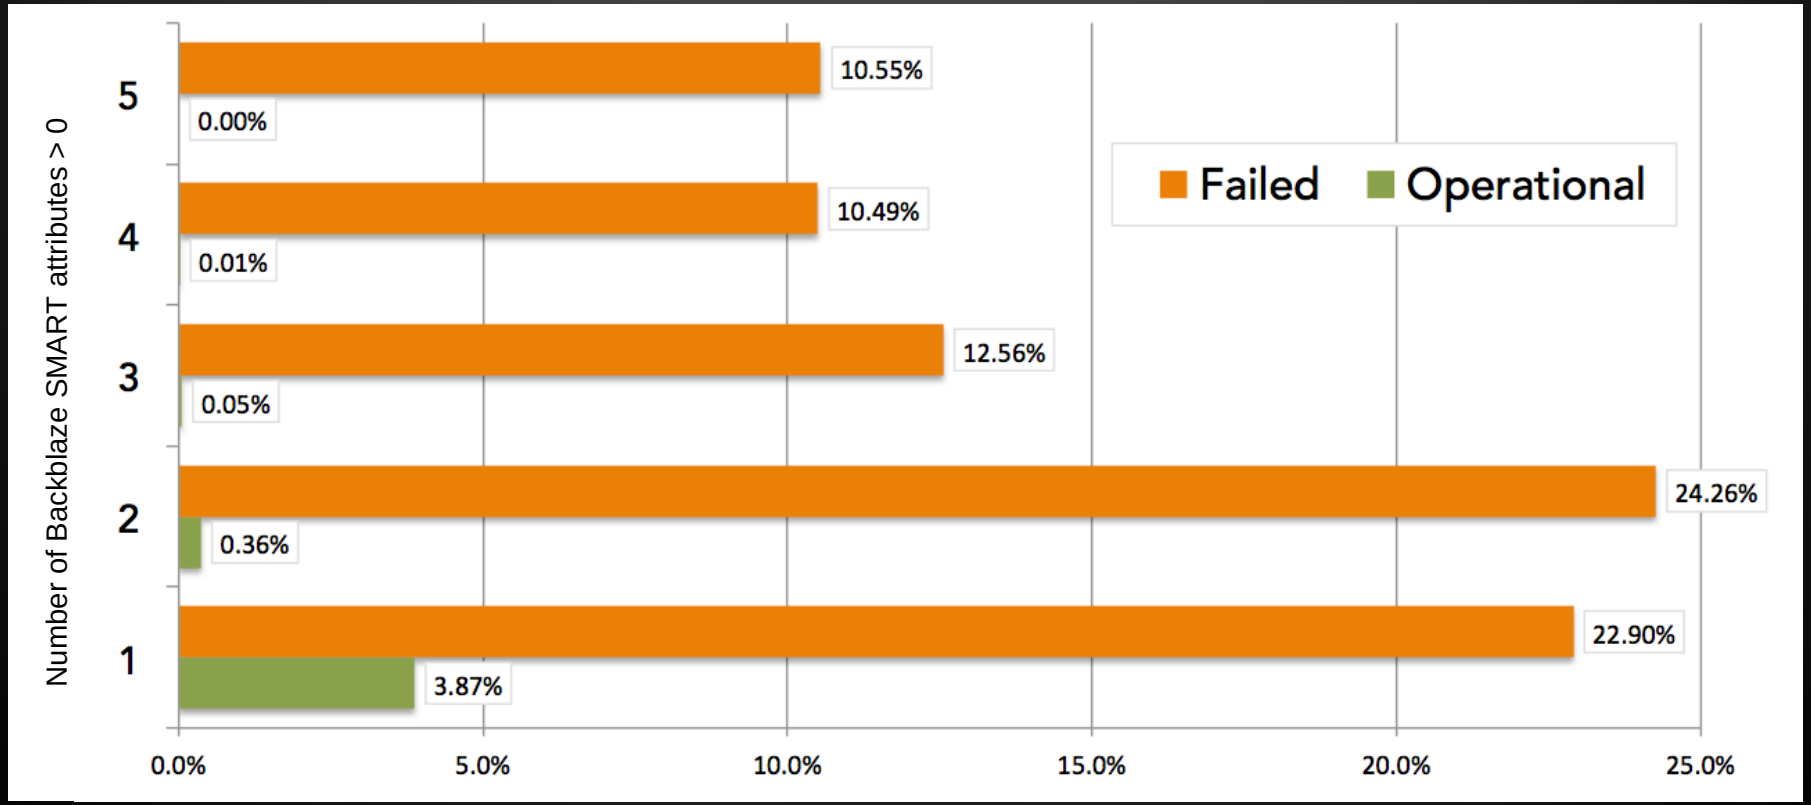

SMART attributes for failure detection

Does it matter? You decide.

4.2%

Percentage of
operational drives
with 1 or more our
5 attributes > 0

76.7%

Percentage of
failed drives
with 1 or more our
5 attributes > 0

Detection rates of Operational vs Failed drives

Percentage of drives with SMART attribute RAW value > 0

Drive Status	SMART 5	SMART 187	SMART 188	SMART 197	SMART 198
Operational	1.1%	0.5%	4.8%	0.7%	0.3%
Failed	42.2%	43.5%	44.8%	43.1%	33.0%

1) Operational drives as of July 31, 2016

2) Failed drives as of one day prior to failure

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

Correlation of stats used for failure detection

	SMART 5	SMART 187	SMART 188	SMART 197	SMART 198
SMART 5	1	0.034	0.026	0.064	0.043
SMART 187	0.034	1	0.007	0.025	0.033
SMART 188	0.026	0.007	1	0.000	0.006
SMART 197	0.064	0.025	0.000	1	0.808
SMART 198	0.043	0.033	0.006	0.808	1

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

What other SMART stats have we looked at to see if they can help predict drive failure?

SMART 189 – High Fly Writes

- Detects when a recording head is flying outside its normal range of operation.
- Adds to counter each occurrence.
- Seagate only in our dataset.

Failed Drives - 47.0%

Operational Drives - 16.4%

Outcome: Needs more study, its more about frequency and distribution of occurrence than quantity.

SMART 10 – Spin-Retry Count

- The count of retries of each spin start attempt after the initial failure.
- Only found in some HGST and Toshiba drives.

Failed Drives - 1.48%

Operational Drives - 0.07%

Outcome: In our case we don't power cycle drives much, so we don't see this error. But, could be a sign of impending failure.

So, what is the failure rate of drives?

Drive Stats

- Drives in operation: 68,877 as of July 31, 2016
- Number of failures: 4,794
- Operational Drive Hours: > 1Billion
- 6.9% of our drives have failed, BUT…

Annualized Failure Rate

AFR is just 3.89%

AFR = **(100*drive-failures)/(drive-days/365)**

Drive Failure Rates Over Time

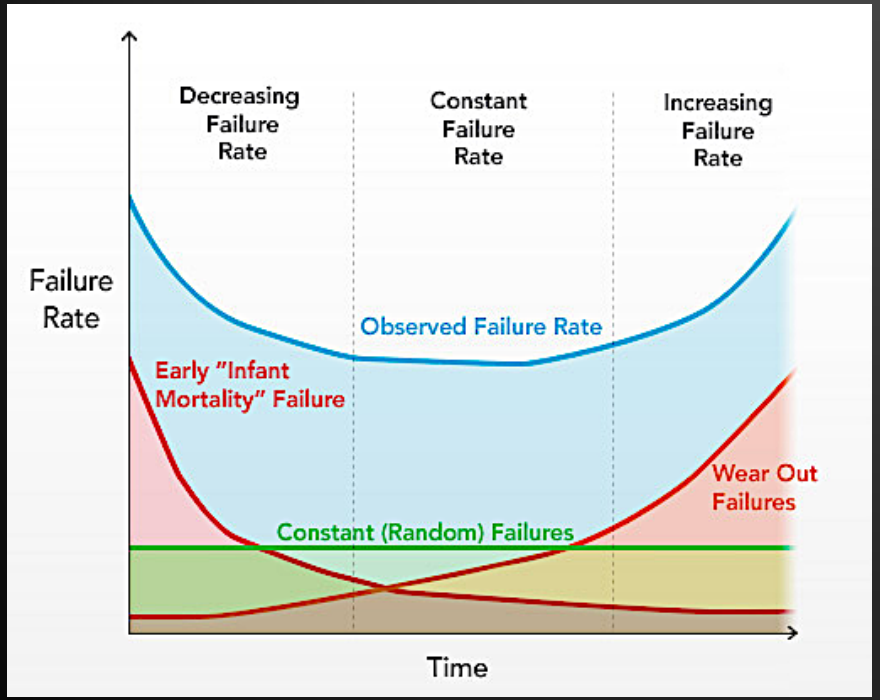

The bathtub curve

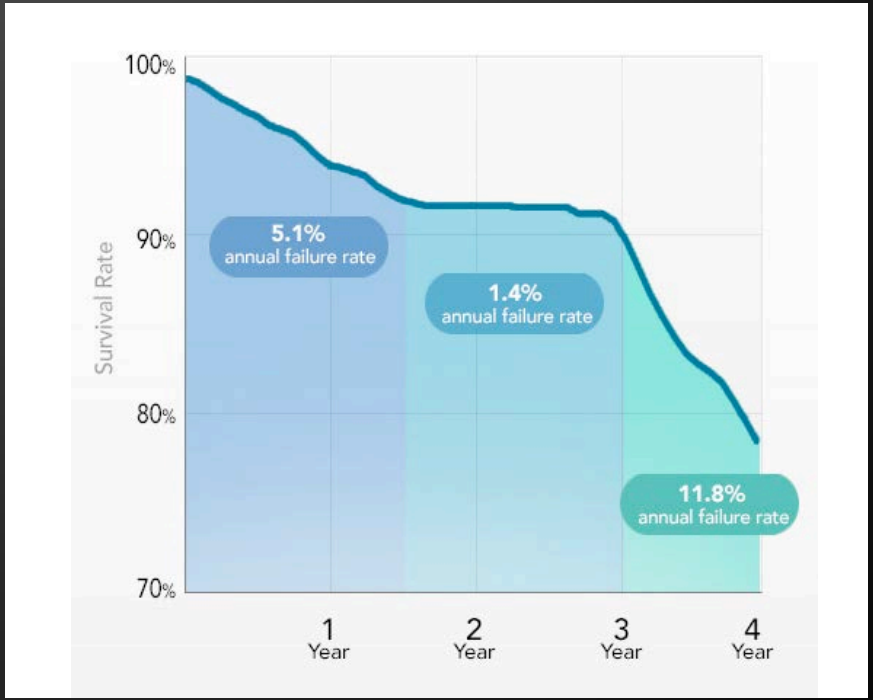

The curve in reality

Are enterprise drives more reliable?

Enterprise versus Consumer drives at Backblaze

	Enterprise Drives	Consumer Drives
Drive-Years of Service	368	14,719
Number of Failures	17	613
Annualized Failure Rate	4.6%	4.1%

At Backblaze the *failure rate of our enterprise drives is actually higher than that of our consumer drives!*

Drives fail when they're
too cold,
too hot, or
just right?

Operational Drives by Temperature

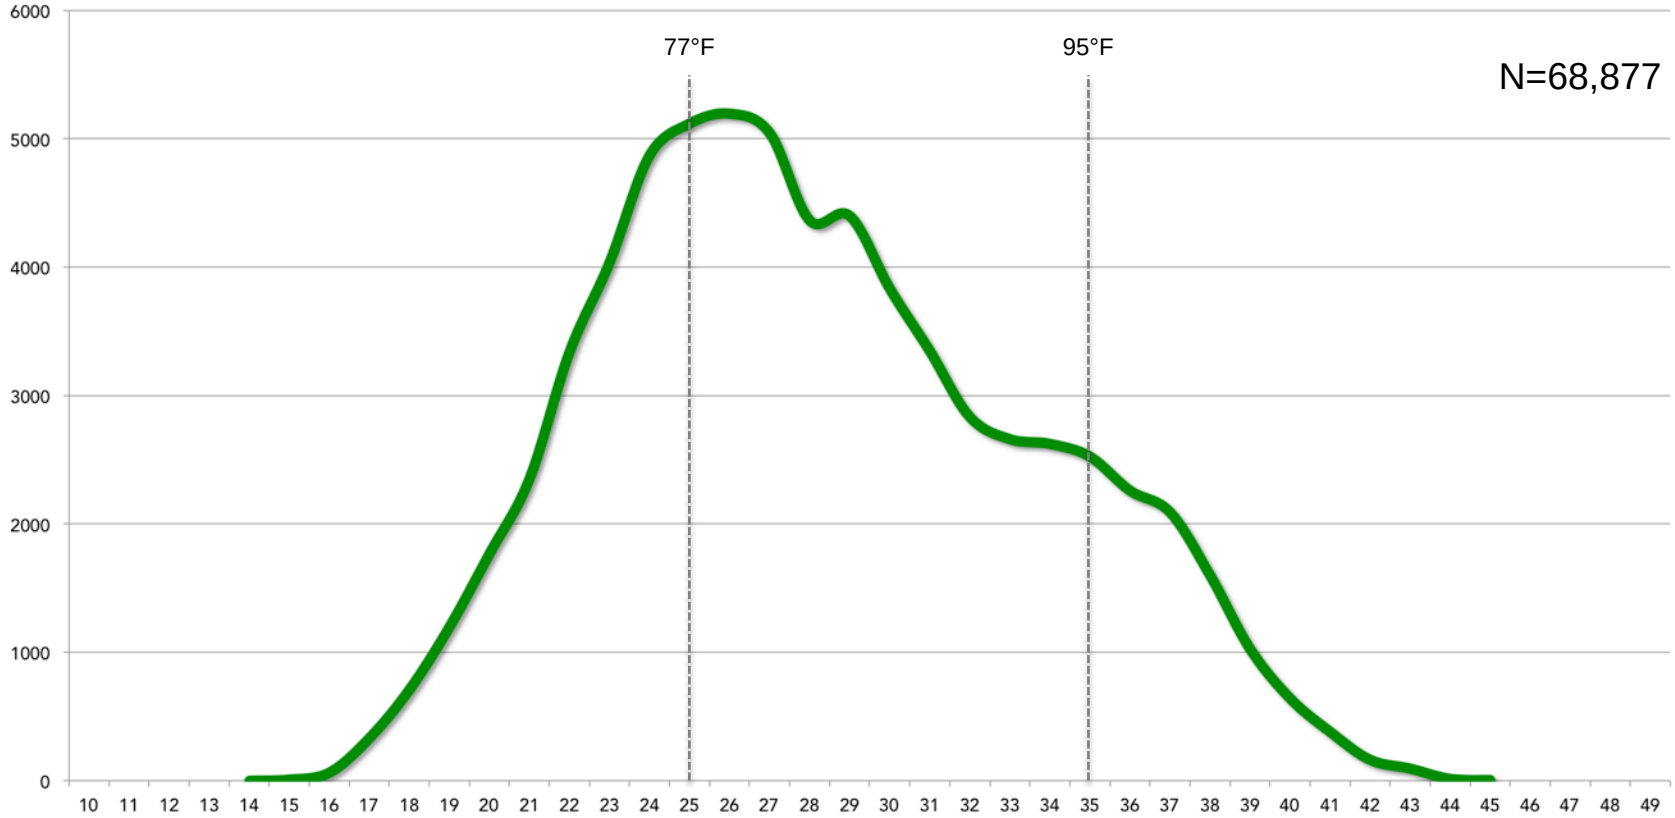

Failed Drives versus Operational Drives by Temperature

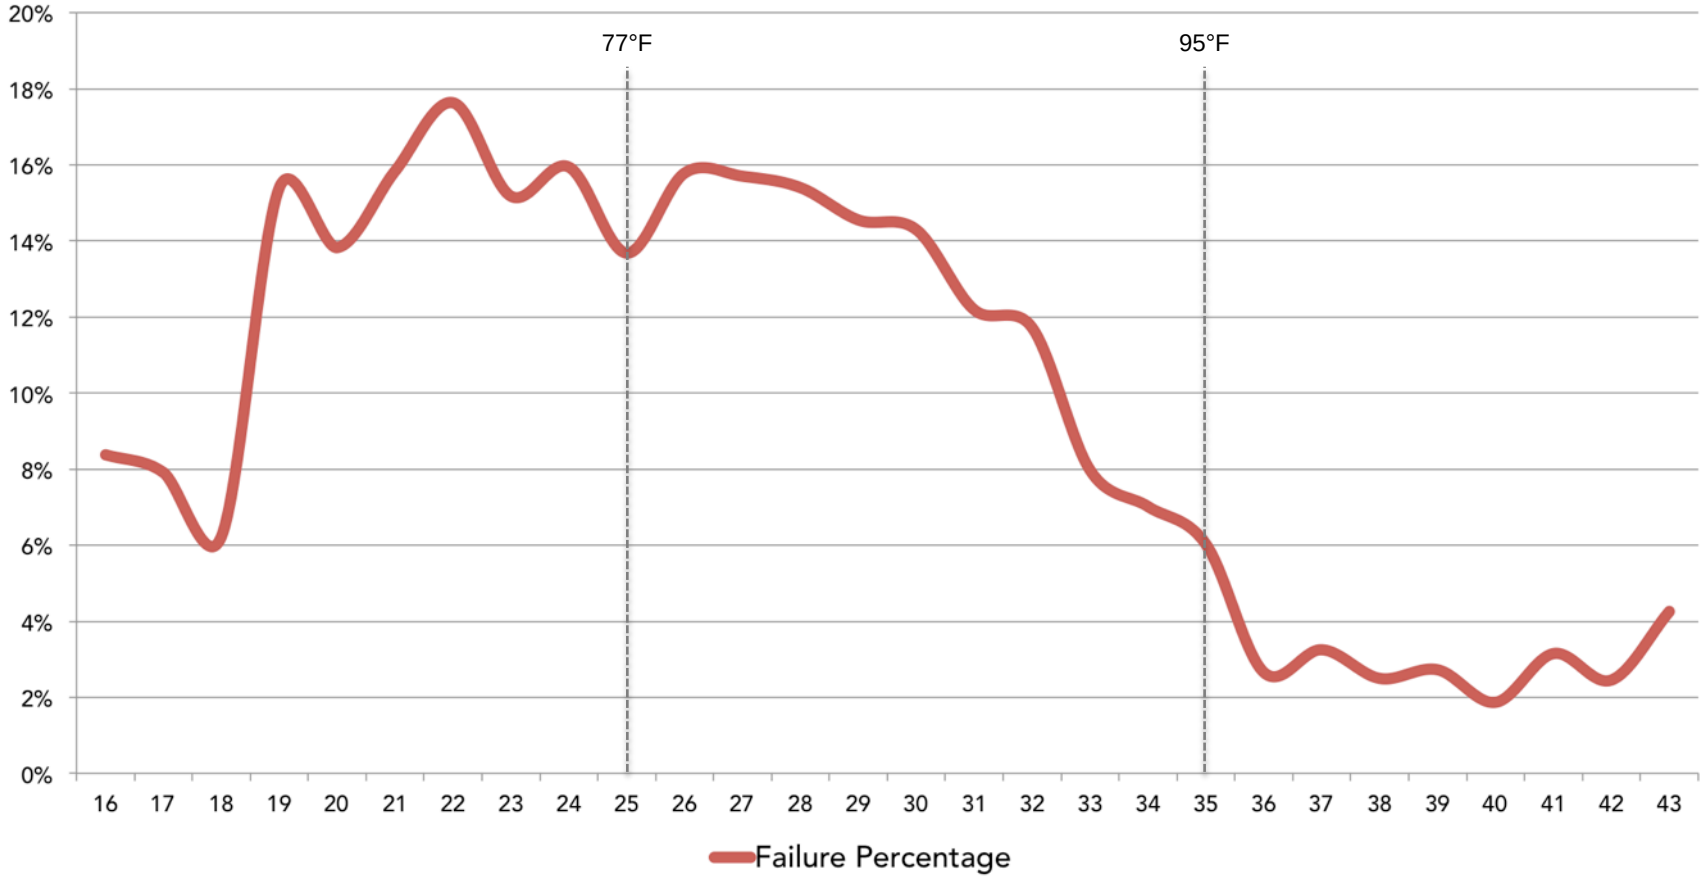

Failed Drives versus Operational Drives by Temperature and MFG

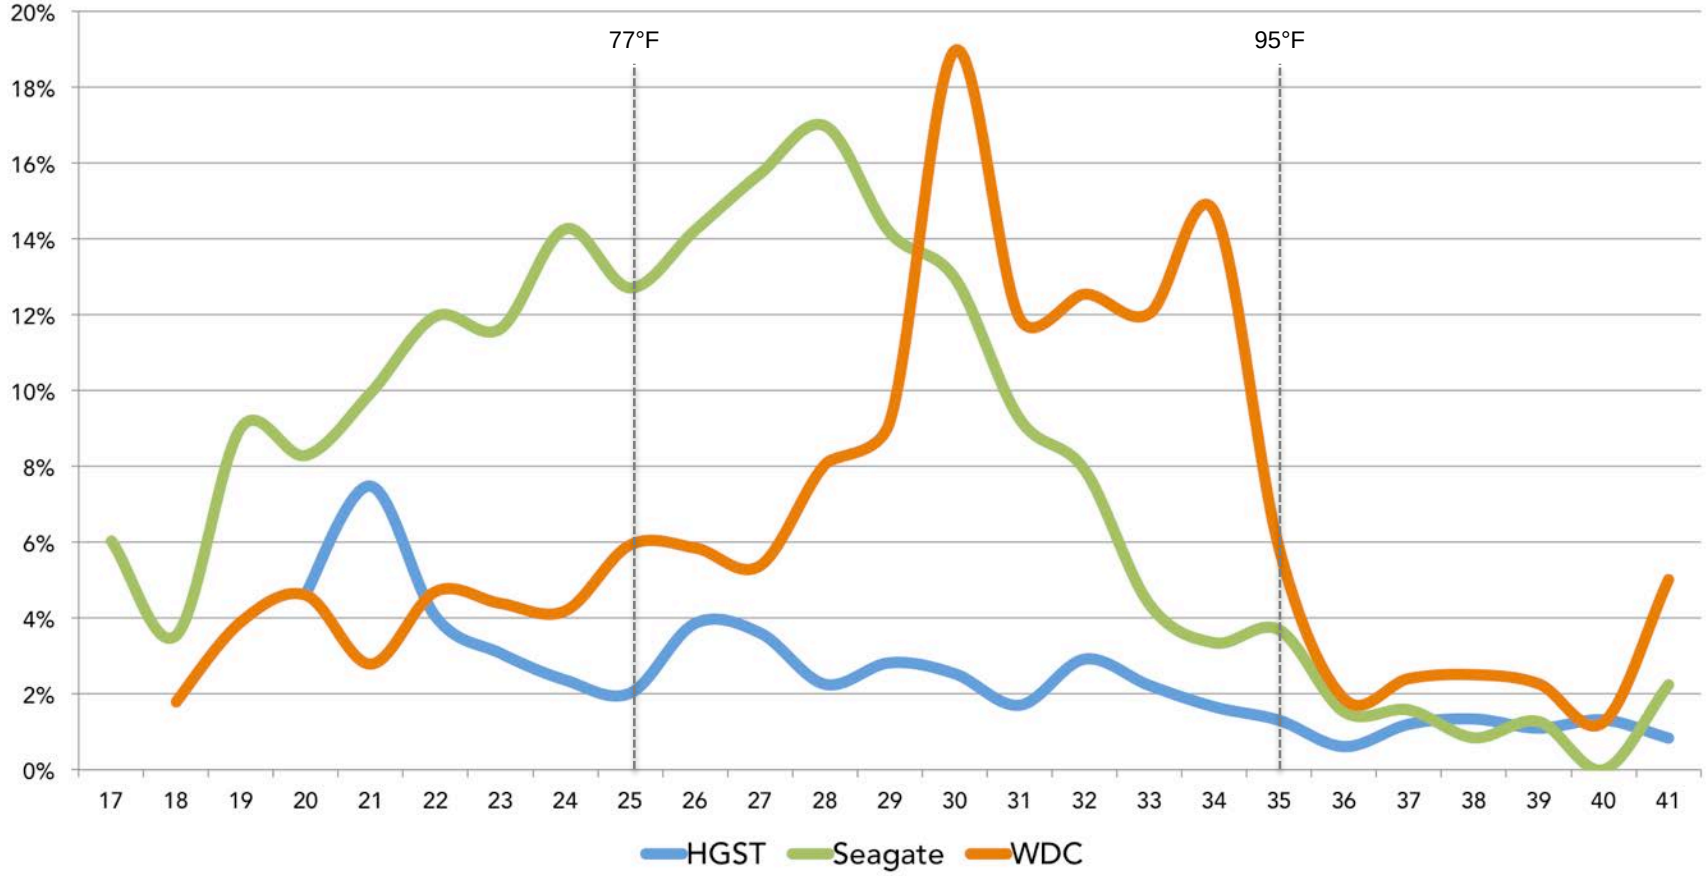

SMART vs. Temperature

Does temperature correlate to our detection stats?

	SMART 5	SMART 187	SMART 188	SMART 197	SMART 198
SMART 194 Temperature	(0.036)	(0.015)	0.005	(0.019)	(0.008)

Is power cycling drives bad?

SMART 12 – Power Cycle Count

- The count of full hard disk power on/off cycles.
- We only power cycle drives when there’s a problem with the Storage Pod they are in.

The average number of times the Failed Drives were power cycled

27.7

The average number of times the Operational Drives were power cycled

10.2

Review

- Our environment
- How we diagnose “sick” drives
- Drive Reliability Findings
  1. Reliability over time
  2. Enterprise vs. consumer drives
  3. Effects of temperature
  4. Effects of power cycling
- Review and Questions

Backblaze.com/Storage-Pod

@GlebBudman