


Behind the Curtain of Backblaze Hard Drive Stats



What

MSST 2017

Where


Santa Clara, California – May 2017

Who

Andrew Klein – Backblaze

Overview

- ◆ History
- ◆ Drive Stats
 - ◆ Reliability over time
 - ◆ Enterprise vs. consumer drives
- ◆ SMART Stats
 - ◆ Is predicting drive failure possible?
 - ◆ Other attributes to consider



History

- Incorporated – 4/20/2007
- V1.0 launched – 9/16/2008
- Storage Pod 1.0 – 9/1/2009
- 10 PB stored – 12/19/2010
- 1st Drive stats post – 11/12/2013
- HD Data published – 2/4/2015
- 300 PB stored – 11/12/2016



Speed bumps...

...Success too





300

PETABYTES

82,516 Drives

8 TB – 12,365

6 TB – 2,334

5 TB – 45

4 TB – 61,347

3 TB – 6,425

Hard Drive Failure Rates at Backblaze



The Basics

- ◆ Use smartmontools package to collect data
 - ◆ <https://www.smartmontools.org/>
- ◆ Collect data once a day from each drive
- ◆ Started keeping data in April 2013
- ◆ All drives in DC
 - ◆ Customer data drives
 - ◆ Boot drives
 - ◆ Management systems, etc.

Daily SMART Stats Data

date	serial number	model	capacity bytes	failure	smart 1 normalized	smart 1 raw
3/31/17	MJ0351YNG9Z0XA	Hitachi HDS5C3030ALA630	3000592982016	0	100	0
3/31/17	MJ0351YNG9WJSA	Hitachi HDS5C3030ALA630	3000592982016	0	100	0
3/31/17	PL1321LAG34XWH	Hitachi HDS5C4040ALE630	4000787030016	0	100	0
3/31/17	MJ0351YNGABYAA	Hitachi HDS5C3030ALA630	3000592982016	0	100	0
3/31/17	PL2331LAHDBJPJ	HGST HMS5C4040BLE640	4000787030016	0	100	0
3/31/17	PL1331LAHG1S4H	HGST HMS5C4040ALE640	4000787030016	0	100	0
3/31/17	PL2331LAGN2YTJ	HGST HMS5C4040BLE640	4000787030016	0	100	0



More attributes

More hard drives



Lifetime Drive Stats

- ◆ Period: April 2013 – March 31, 2017
- ◆ Drives used: 95,054
- ◆ Number of failures: 5,674
- ◆ Drive days (all drives): 62,147,829
- ◆ At first glance
 - ◆ 5.97% of our drives have failed, BUT...

www.backblaze.com/b2/hard-drive-test-data.html

Annualized Failure Rate

AFR is just 3.33%

$$\text{AFR} = (100 * \text{drive-failures}) / (\text{drive-days} / 365)$$

Hard Drive Annualized Failure Rates for Q1 2017

Reporting period 1/1/2017 - 3/31/2017 inclusive

MFG	Model	Drive Size	Drive Count	Drive Days	Drive Failures	Annualized Failure Rate
Toshiba	DT01ACA300	3TB	46	3,956	-	0.00%
HGST	HDS5C3030ALA630	3TB	4,380	383,788	10	0.95%
HGST	HDS723030ALA640	3TB	974	83,918	4	1.74%
WDC	WD30EFRX	3TB	1,025	94,973	3	1.15%
HGST	HDS5C4040ALE630	4TB	2,624	232,760	1	0.16%
HGST	HMS5C4040ALE640	4TB	8,482	644,282	6	0.34%
HGST	HMS5C4040BLE640	4TB	15,339	959,967	37	1.41%
Toshiba	MD04ABA400V	4TB	146	12,551	1	2.91%
Seagate	ST4000DM000	4TB	34,540	2,981,251	267	3.27%
Seagate	ST4000DX000	4TB	170	15,261	15	35.88%
WDC	WD40EFRX	4TB	46	3,956	-	0.00%
Toshiba	MD04ABA500V	5TB	45	3,870	-	0.00%
Seagate	ST6000DX000	6TB	1,891	162,506	3	0.67%
WDC	WD60EFRX	6TB	443	38,271	3	2.86%
HGST	HUH728080ALE600	8TB	45	3,870	-	0.00%
Seagate	ST8000DM002	8TB	9,861	815,494	23	1.03%
Seagate	ST8000NM0055	8TB	2,459	37,559	2	1.94%
Totals			82,516	6,478,233	375	2.11%

Hard Drive Annualized Failure Rates

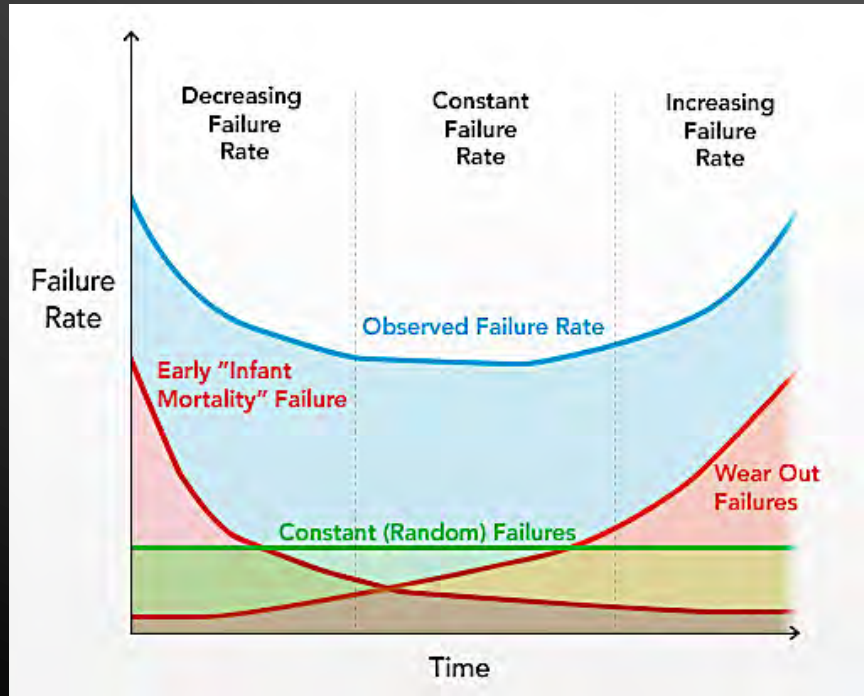
Reporting period April 2013 - March 31, 2017

MFG	Model	Size	Drive Count	Annualized Failure Rate	Confidence Interval	
					Low	High
HGST	HDS5C3030ALA630	3TB	4,380	0.84%	0.7%	1.0%
HGST	HDS723030ALA640	3TB	974	1.96%	1.5%	2.5%
Toshiba	DT01ACA300	3TB	46	3.89%	1.6%	8.0%
WDC	WD30EFRX	3TB	1,025	5.63%	4.8%	6.5%
HGST	HDS5C4040ALE630	4TB	2,624	0.88%	0.7%	1.1%
HGST	HMS5C4040ALE640	4TB	8,482	0.64%	0.5%	0.8%
HGST	HMS5C4040BLE640	4TB	15,339	0.68%	0.5%	0.8%
Seagate	ST4000DM000	4TB	34,540	3.00%	2.9%	3.1%
Seagate	ST4000DX000	4TB	170	7.51%	5.6%	9.8%
Toshiba	MD04ABA400V	4TB	146	1.50%	0.4%	3.8%
WDC	WD40EFRX	4TB	46	2.28%	0.5%	6.7%
Toshiba	MD04ABA500V	5TB	45	2.34%	0.3%	8.4%
Seagate	ST6000DX000	6TB	1,891	1.30%	0.9%	1.7%
WDC	WD60EFRX	6TB	443	5.59%	4.2%	7.2%
HGST	HUH728080ALE600	8TB	45	2.10%	0.3%	7.6%
Seagate	ST8000DM002	8TB	9,861	1.60%	1.2%	2.0%
Seagate	ST8000NM0055	8TB	2,459	2.38%	0.2%	7.4%
Totals			82,516	2.07%	2.0%	2.1%

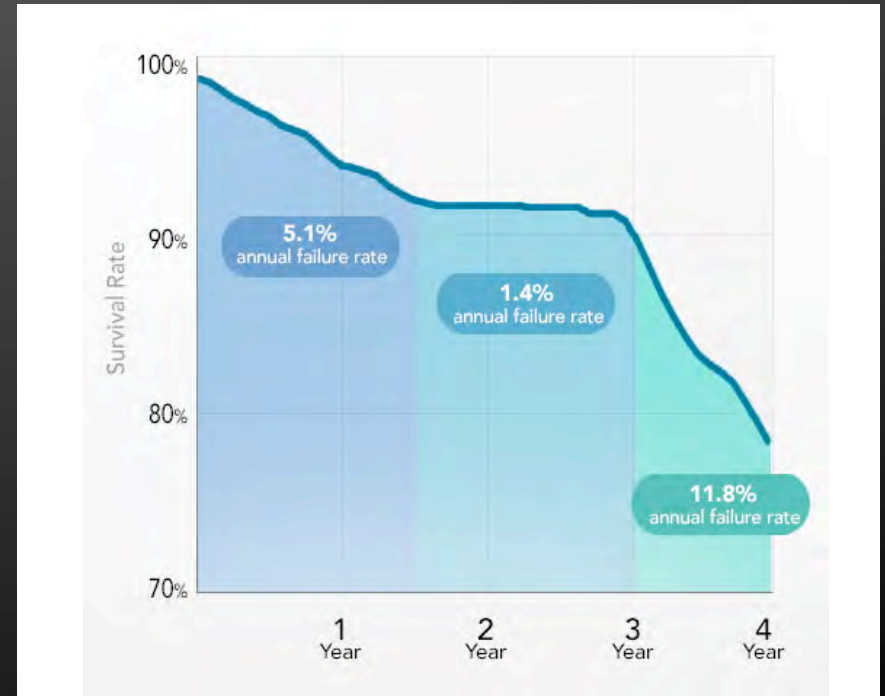
Defining a Drive Failure

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

Drive Failure Rates Over Time

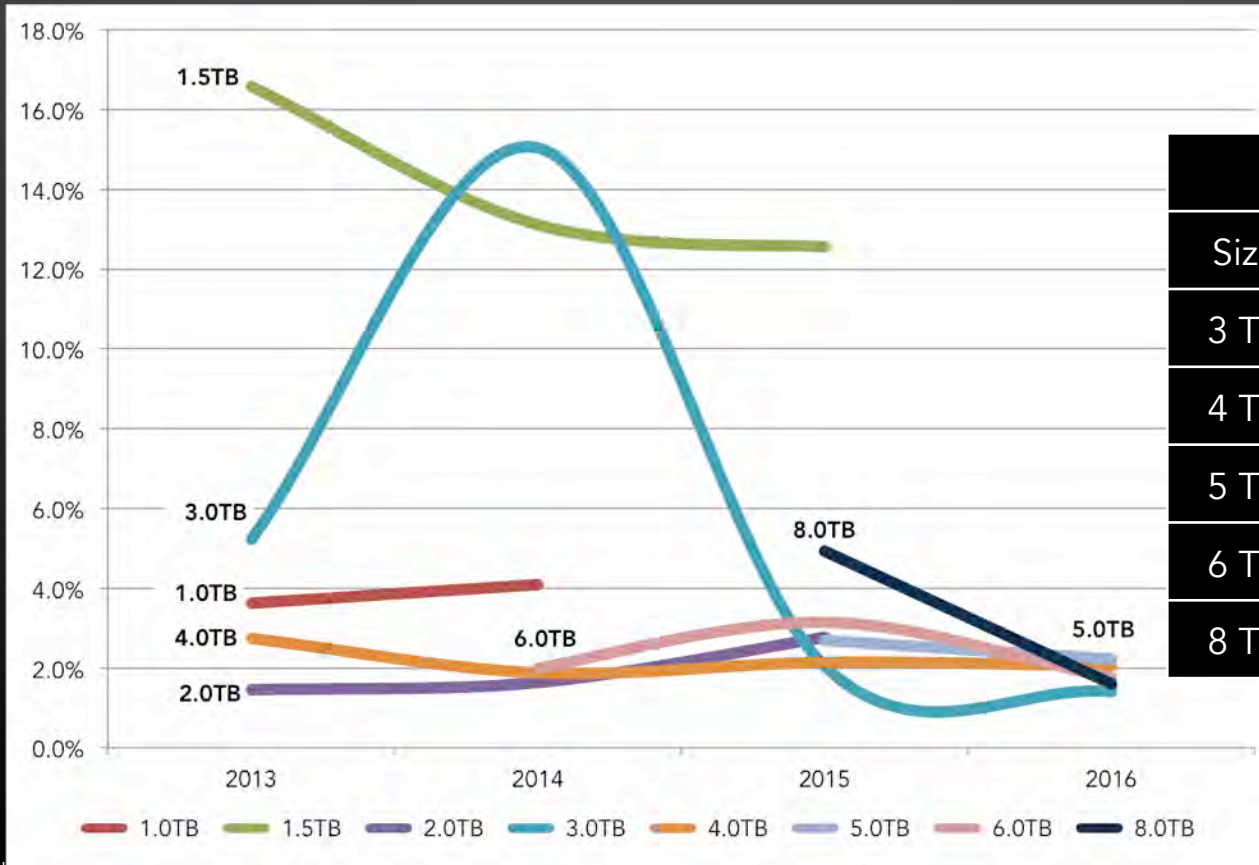


The bathtub curve



The curve in reality

Hard Drive Failure Rates Cumulative by Year



As of Q1 2017		
Size	Q1	+/- vs Q4
3 TB	1.72%	0.32%
4 TB	2.17%	0.12%
5 TB	2.30%	0.08%
6 TB	2.24%	0.48%
8 TB	1.60%	0.00%

Consumer drives
in a DC?

Yes.



Enterprise versus Consumer drives – Then...

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

Enterprise versus Consumer drives – Now...

	Enterprise Drives Seagate 8 TB ST8000NM0055	Consumer Drives Seagate 8 TB ST8000DM002
Drives in Service	2,459	9,861
Drive-Years of Service	84	4,320
Number of Failures	2	69
Annualized Failure Rate	2.38%	1.60%

Disk Properties that Matter to us

- ◆ Cost per GB – right now \$0.025 - \$0.03 / GB
- ◆ Power – The lower the better
- ◆ Fits our usage
- ◆ Failure Rates
- ◆ Warranty
- ◆ Speed

Can you diagnose
a "sick" drive?



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 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) 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

SMART attributes for failure detection

Does it matter? You decide.

4.2%

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

76.7%

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

Rules – Seagate Drives

Rule	Outcome	Confidence
If SMART_197_raw < 2 and SMART_188_raw > 0 and SMART_1_normalized \in [0,117]	Healthy	100%
If SMART_197_raw >= 2	Replace	100%
If SMART_197_raw < 2 and SMART_188_raw > 0 and SMART_1_normalized > 117	Replace	80%
If SMART_197_raw < 2 and SMART_188_raw = 0 and SMART_187_normalized < 100 and SMART_240_raw < 14780 billion	Replace	97%

Rules extracted
from a decision
tree model trained
on the Seagate
dataset

Predicting Disk Replacement towards Reliable Data Centers. IBM Research, Zurich Switzerland, [Mirela Botezatu et. al.]
<http://www.kdd.org/kdd2016/subtopic/view/predicting-disk-replacement-towards-reliable-data-centers>.

Rules – Hitachi/HGST Drives

Rules extracted from a decision tree model trained on the Hitachi/HGST dataset

Hitachi/HGST	If SMART_197_raw > 1 and SMART_3_raw > 626	Replace	100%
Hitachi/HGST	If SMART_197_raw > 5 and SMART_3_raw < 626 and SMART_5_raw > 17	Replace	92%
Hitachi/HGST	If SMART_197_raw > 1 and SMART_3_raw < 626 and SMART_5_raw < 17	Replace	100%
Hitachi/HGST	If SMART_197_raw < 1 and SMART_5_raw < 7200 and SMART_3_raw > 629 and SMART_1_raw \in [0,109]	Healthy	97%

Predicting Disk Replacement towards Reliable Data Centers. IBM Research, Zurich Switzerland, Mirela Botezatu et al. <http://www.kdd.org/kdd2016/subtopic/view/predicting-disk-replacement-towards-reliable-data-centers>.

**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.

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

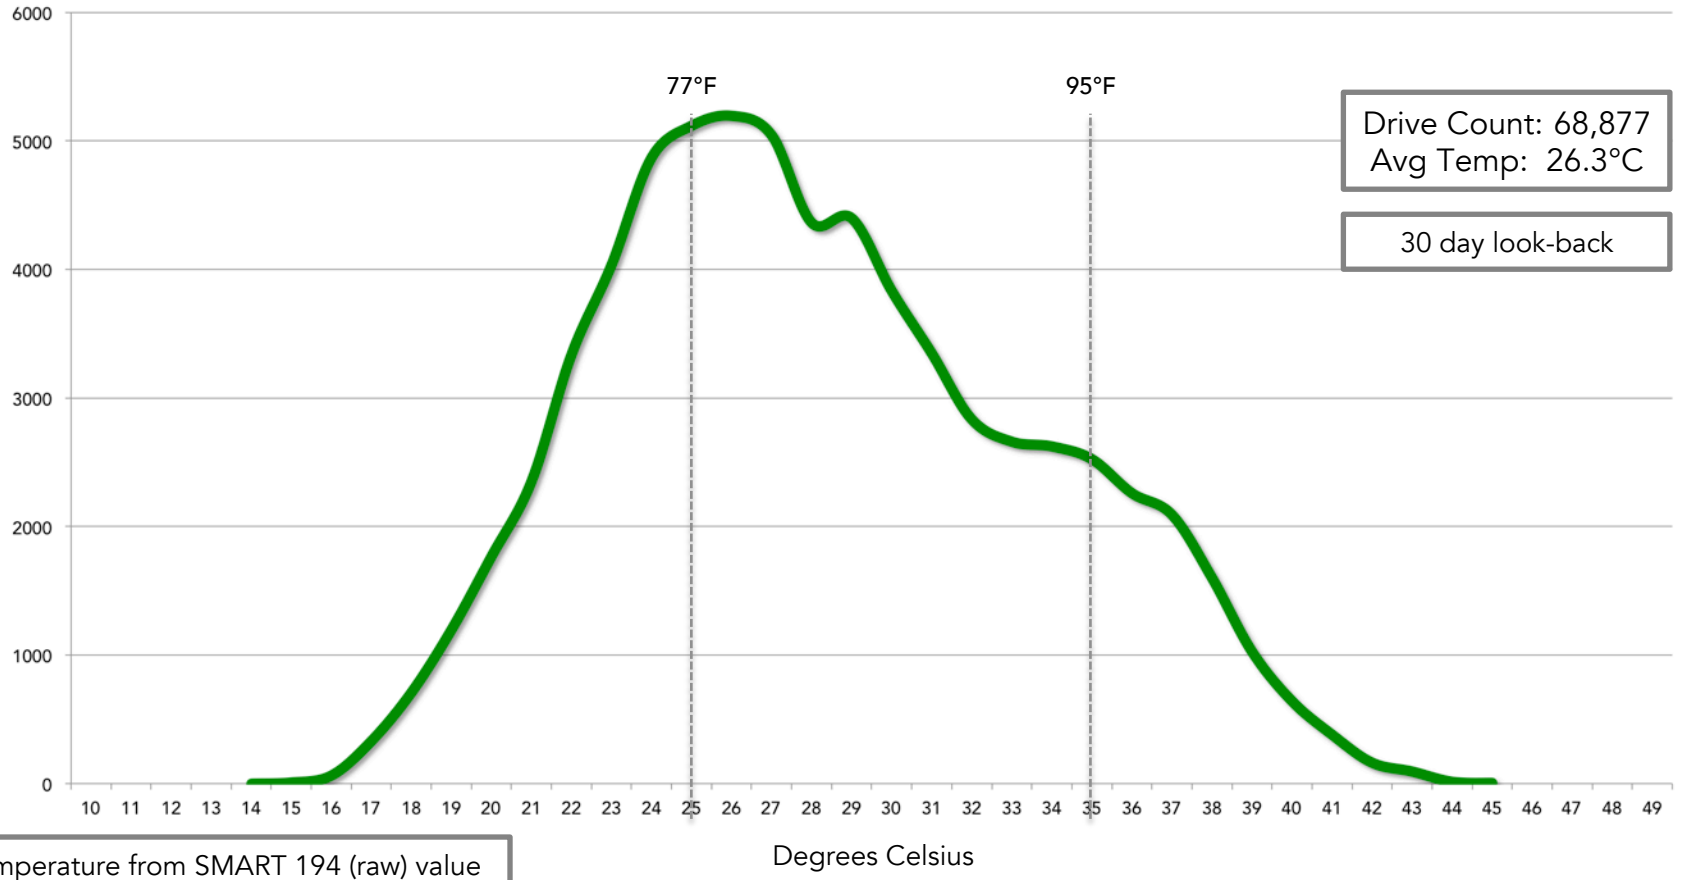
Drives fail when
they're too cold,
too hot, or
just right?



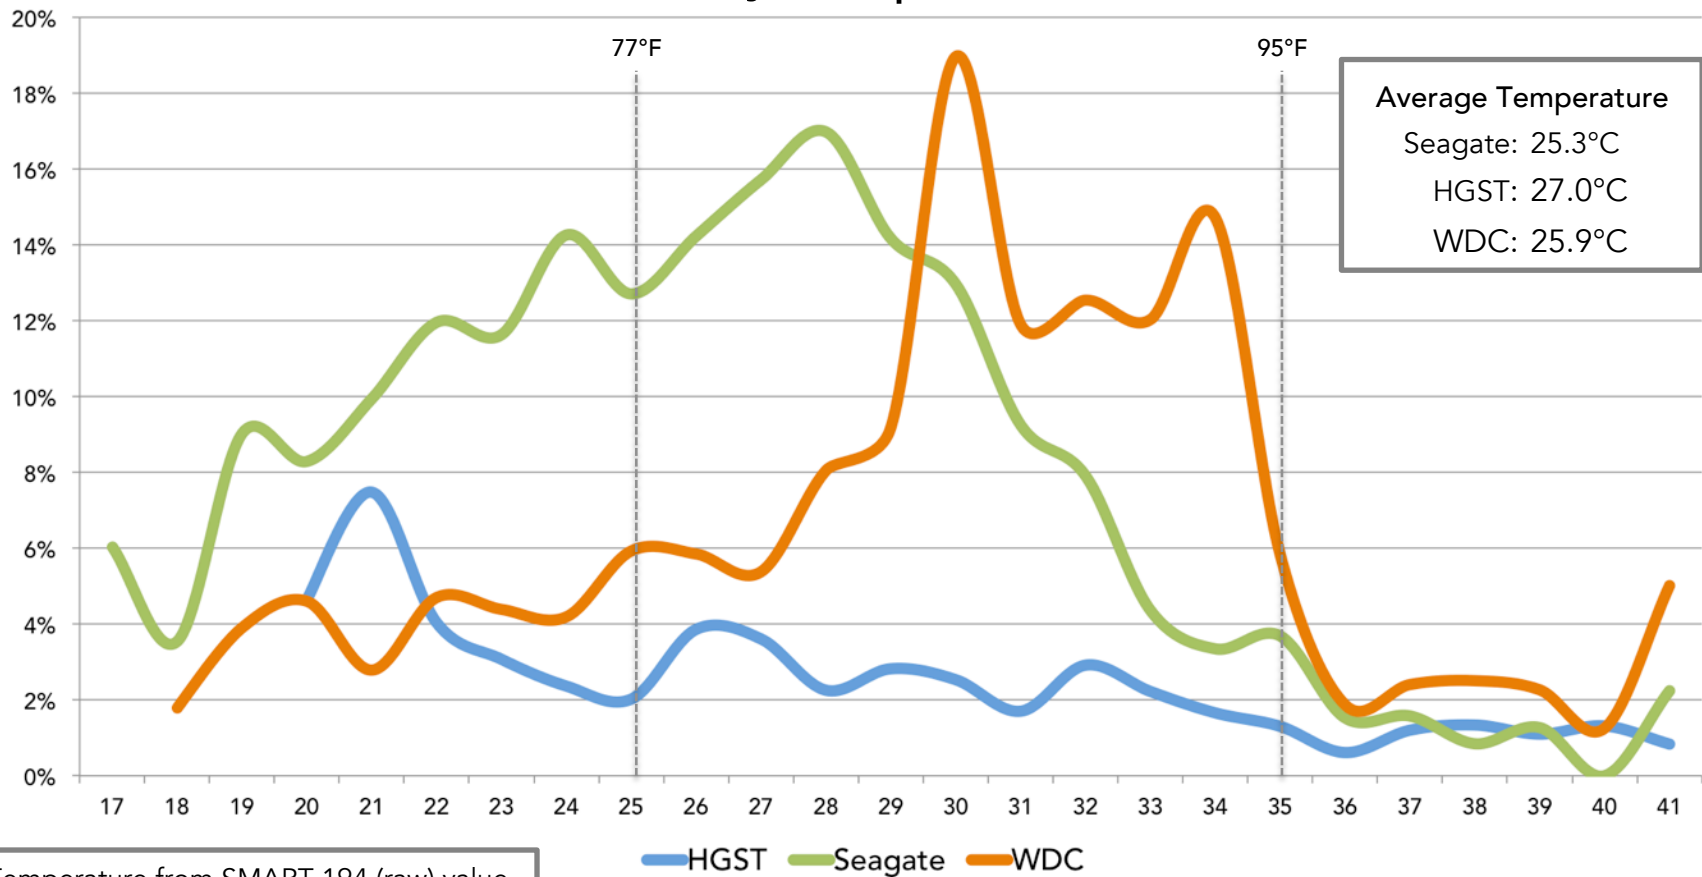
Temperature: Operational vs Failed Drives

	Operational	Failed
Drive Count	68,877	4,794
Average Temperature	26.3°C / 79.3°F	25.6°C / 78.0°F
Temperature Range	14°C - 45°C 57.2°F - 113.0°F	13°C - 46°C 55.4°F - 114.8°F
Median	27.0°C / 80.6°F	25.0°C / 77.0°F
Mode	27.2°C / 80.9°F	24.0°C / 75.2°F

Operational Drives by Temperature



Failed Drives by Temperature and MFG



Temperature from SMART 194 (raw) value

— HGST — Seagate — WDC

Review

- ◆ History
- ◆ Drive Stats
 - ◆ Reliability over time
 - ◆ Enterprise vs. consumer drives
- ◆ SMART Stats
 - ◆ Is predicting drive failure possible?
 - ◆ Other attributes to consider



Questions

More: backblaze.com/b2/hard-drive-test-data.html