

# Impact of Temperature on Hard Disk Drive Reliability in Large Datacenters

Sriram Sankar  
Microsoft Corporation  
Redmond, WA  
srsankar@microsoft.com

Mark Shaw  
Microsoft Corporation  
Redmond, WA  
mshaw@microsoft.com

Kushagra Vaid  
Microsoft Corporation  
Redmond, WA  
kvaid@microsoft.com

**Abstract**—When datacenters are pushed to their limits of operational efficiency, reducing failure rates becomes critical for maintaining high levels of healthy server operation. In this experience report, we present a dense storage case study from a large population of servers housing tens of thousands of disk drives. Previous studies have presented divergent results concerning correlation between temperature and hard disk drive failures. In our paper, we specifically establish correlation between temperatures and failures observed at different location granularities: a) inside drive locations in a server chassis, b) across server locations in a rack and c) across multiple racks in a datacenter. We also establish that temperature exhibits a stronger correlation to failures compared to the correlation of disk utilization with drive failures. Thus, we show that temperature-aware server and datacenter design plays a pivotal role in datacenter reliability. Following our case study, we present a reliability model for estimating hard disk drive failures correlated with the datacenter operating temperature. We use a physical Arrhenius model with empirically derived coefficients for our model. We show an application of the model for selecting the datacenter inlet temperature setpoint for two different server storage configurations. Finally, with the help of a datacenter cost discussion, we highlight the need to incorporate reliability-aware datacenter design for increased efficiency in large scale datacenters.

**Keywords**- Datacenter, Reliability, Temperature, Hard Disk Drive, Correlation

**Submission Category**- Practical Experience Report

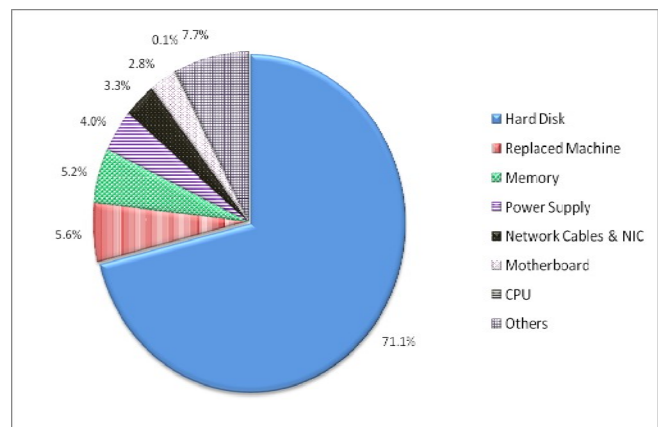
## I. INTRODUCTION

As large enterprises move to modular datacenters [8] and efficient cooling practices become prevalent [5], we move closer to the limits of cost efficiency achievable in the power domain. The next order of datacenter efficiency is to make the *entire computation capacity* provisioned in the datacenter available to the application deployment at all times. Hence increasing reliability is of primary importance in large scale datacenters.

Server components are composed of electrical and mechanical parts, and hence they are prone to failures. Frequent failures reduce infrastructure availability and increase the cost of datacenter operations. In addition, the server design in itself could be a major catalyst for most of the server component failures. For instance, we found that a

particular drive location in a dense storage configuration under a fairly constant workload was continuously exposed to high temperature conditions, even under nominal inlet temperature to the server. We found a higher number of drives in this location failing more often, thereby showing strong correlation to operating conditions. Understanding the reason behind such failures enabled us to address the design issues, thereby increasing the availability of machines for the particular online service. Increase in availability of machines increases service availability for the application. In turn, it decreases the cost associated with application scaling and delays the need to build new datacenters. Hence, having more servers that are readily available affects the financials of a large enterprise.

Server component failures have indeed been recognized as important and prior works have studied individual component reliability, such as for hard disks [15, 12] and memory [16]. Figure 1 presents actual data on the different kinds of failure types observed over a period of two years from our typical large-scale datacenters housing more than 100,000 servers. We see clearly that hard disk drives account for 71% of the known failures, making it the most dominant failing part. This is in part due to the mechanical moving parts of the disk drives and also due to the extensive use of SATA drives in this deployment. SATA disk drives are known for failing more often than SAS drives, but cost less for storage capacity/\$ [10]. Memory failures constitute about 5.2% of the total failures, including configuration errors. The actual percentage of memory that was replaced to correct memory failures was 4%. This is close to numbers reported



**Figure 1.** Breakdown of Hardware Component Errors in a Large Datacenter (2 years failure data)

by Schroeder et al [16]. Some errors like Network cable errors and NIC issues are also worth noting. It shows that hardware and network configuration issues do result in service issues on the datacenter floor. The ‘others’ category includes storage/power backplane issues, wrong wiring, and other issues that do not fit into the major component buckets. Given that hard disk drives are the most significant failing component and recent previous studies established no conclusive relationship between temperature and hard disk drive failures [12], we set out to evaluate whether temperature relatable to datacenter and server design has stronger correlation to failure.

In particular, our paper makes the following contributions:

1. We present a *temperature case study* from a large population of servers, where we observed chassis and datacenter design leading to increased number of server failures. We show correlation between temperature and observed failures inside the server chassis, across servers in a rack and across multiple racks in the datacenter.
2. We present a *reliability model* (based on Arrhenius model) for hard disk drives that can estimate expected annualized failure rate (AFR) based on different inlet operating temperature.
3. We show how our reliability model can be used to select *efficient datacenter operating temperature* setpoints for different server configurations. We present a comparative cost analysis for a reliability-aware datacenter design.

The rest of the paper is structured as follows: Section 2 presents details on the datacenter infrastructure while Section 3 discusses the major failure determinants for hard disk drives. Section 4 presents a case study from a population containing tens of thousands of disk drives and establishes correlation between temperature and observed failures over a period of 1 year. Section 5 presents our reliability model and elaborates on how the model could be used to determine temperature operating setpoint in large scale datacenters. Section 6 discusses a comparative cost analysis between increased failures vs power reduction. Section 7 presents related work in this field. Section 8 summarizes the lessons learnt and Section 9 concludes the paper.

## II. EXPERIMENT INFRASTRUCTURE

### A. Hardware Infrastructure

All the servers in this study are identical with dual CPUs and an additional storage enclosure containing up to 40 SATA drives in a RAID 1+0 configuration. Figure 2 shows the layout of the disk drive columns across a server chassis. In our server chassis, we are able to fit 5 disk drives (3.5” SATA HDD) across the length of the server. The traditional datacenter racks have a cold aisle from which cold air is pulled across the server and exhausted out in the hot aisle (HDD1->HDD5). Hence the air gets preheated by the time it

reaches the interior hard disk drives and leads to higher temperatures for those hard disk drives. We abstract out the specifics of the actual configuration details (including number of disk drives, logical layout) to preserve proprietary configurations; however the statistical study was done over the entire sample. All the facilities that are discussed in this study are raised-floor datacenters and have cooling equipment to control temperature inside the datacenter.

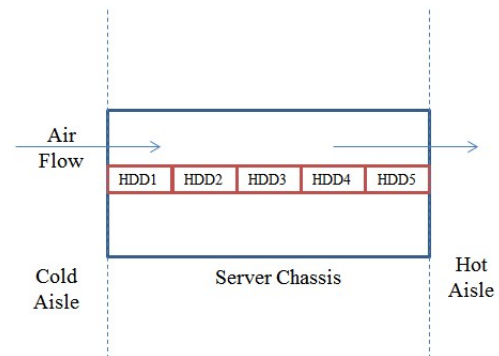


Figure 2. Hard Drive Columns within Server Chassis

### B. Management Infrastructure

Large scale enterprises have a management infrastructure in place that monitors the different kind of failures that occur in their server deployments at scale, similar to Autopilot [11], Microsoft System Center, or the Cluster-level infrastructure [9] and System Health Infrastructure [12] used at Google. The management infrastructure records and reports software and hardware errors encountered in operation. We use a similar infrastructure to monitor and classify the server failures in our largest server deployments.

## III. HARD DISK DRIVE FAILURE DETERMINANTS

### A. Age of the Disk Drive

Several previous studies have established different failure rates with respect to the age of the disk drive population [12]. A typical failure curve across age resembles a Weibull bathtub shaped curve with a large number of infant mortality, stable mid-life curve and steady increase in failures again at older age. In our study, most of the disk drives are of similar age since all the servers were deployed around similar timeframe when the datacenter became operational. Hence the age factor does not become a major determinant for our study. This is extremely beneficial since this helps isolate the impact of other factors on failure rates in datacenters.


### B. Vibration and SMART monitors

There could be significant vibration due to dense storage; however modern hard disk drives balance internal vibration through vibration compensation techniques in the servo mechanism of the hard disk drives [7]. We currently do not have metrics that expose the level of induced vibration, and measuring the impact of vibration is one of

our projects that are currently underway. We do collect several SMART information from the disk drive population, including Reallocated Sector count, Seek errors, Spin up time, ECC errors, Temperature etc. Though we see SMART counters being indicative of some failures, a predictive methodology is hard to obtain. For one of our large populations, such a methodology would have been able to account for less than 20% of all disk failures. We do not present the details here in interest of space. Previous conclusions made by Pinheiro et al [12] also suggest that SMART counters do not provide a confident way of predicting hard disk drive failures.

### C. Utilization vs Temperature

The remaining two significant failure determinants are disk utilization and temperature. We need to isolate the impact of these two metrics that are location dependent. One of the primary factors that can cause more wear on the hard disk drive is the disk utilization (we use utilization as a proxy for workload duty cycle), which denotes the amount of activity on the hard disk drive. According to the volume and data layout, certain disks might be more stressed than other disks (for instance, a data volume in SQL might have higher level of activity than a Backup volume). We conducted a preliminary investigation to determine which of these two metrics is highly correlated to hard disk drive failures.



**Figure 3.** Temperature shows better correlation to Failures than Sectors Read/min and Sectors Written/min


Figure 3 presents the results of the analysis on a total of 10000 hard disk drives spread across two clusters. We correlated the ‘sectors read/ minute’ and ‘sectors write/ minute’ experienced by the disk drive in a particular location as seen by the controller over its entire lifetime, to the failures observed in that location over a year. On the other hand, we also correlated the temperature observed in those disk locations to the number of failures. We plot the resulting coefficient of correlation in Figure 3. As can be seen from the figure, the read and write activity on the disk drives correlate minimally with the failures. However, drive temperature inside the chassis shows stronger correlation to

disk failures in the particular location within the chassis (R value for temperature is above the critical R value line, at  $df=30$  for a two-tailed test at level of significance = 0.01). Hence for the remainder of the paper, we concentrate on disk drive temperature and do an in-depth temperature measurement and correlation analysis across disk drive locations inside chassis, location of a server within a rack and locations of racks in a datacenter.

## IV. TEMPERATURE CASE STUDY

We present a case study where specific datacenter design parameters and a dense storage chassis design resulted in higher number of disk failures, under high operating temperature. The case study was conducted in a raised-floor datacenter, containing tens of thousands of hard disk drives in a dense storage server and failure data was collected for a period of 1 year.

The result of our study is surprising since earlier studies [12] establish that disk drive failures do not increase with increase in temperature in the field. Figure 4 shows the actual HDD temperature in increments of one degree and the corresponding AFR for our entire population. We see clearly that with increase in HDD temperature, the AFR rate increases. There are some datapoints at the end of the spectrum that have smaller number of samples and hence a higher skew. For the major part of the distribution (shown by



**Figure 4.** Distribution of average temperatures and failure rates (Increase in temperature results in increase in AFR)

PDF columns), we see that AFR steadily increases as HDD temperature increases. Interestingly, we found that the certain disk locations in the heavy storage enclosure were exposed to high temperature for a longer duration even under nominal inlet operation temperatures. We also observed a significant delta T between the inlet temperature at different locations in the datacenter. In the next section, we present our analysis and observations categorized by location

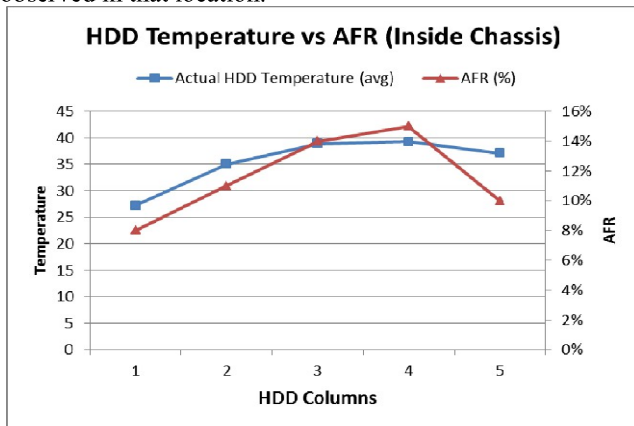
granularity. We divide our correlation analysis into three distinct temperature impact zones:

- Drive locations inside the server chassis
- Server locations within a rack
- Multiple rack locations across the datacenter

There are different factors that come into play for each of these temperature zones. We shall discuss each in more detail in the following sections.

#### A. Correlation inside the Server Chassis

Server design is an important factor in determining the availability of machines in a datacenter. Depending on the placement of the hard disk drives, there could be significant variation in drive temperature. This is especially true in the case of dense storage, since cold air flows from the front of the storage enclosure to the back. Given that the workload running on the disk drives are similar (no significant duty cycle variations), we can establish the correlation if there are more failures for drives which experienced higher operating temperatures. We present the layout of a dense storage device in Figure 2 that was used in our case study. There are five hard disk drives columns where HDDs are arranged one behind the other from the front of the enclosure to the back. Hence the air gets preheated by the time it reaches the interior hard disk drives and leads to high temperatures for those drives. This results in an increase in number of failures observed in that location.



**Figure 5.** Temperature and Failure correlation for Hard Disk Column location within Server Chassis; Correlation Coefficient,  $R = 0.79$  for entire population.


Figure 5 shows the average temperature observed in each hard disk drive column (1 through 5) across all the machines under this study. Note that the temperatures increase from 27 C in the front-most hard disk drive (HDD1) to 39 C in the fourth hard disk drive column (HDD4). This is just the average temperature measurement, and there were hard disk drives that were at temperatures greater than 45 C in hotter parts of the datacenter as shown in Figure 4. The last drive (HDD5) closer to the hot aisle has a reduced temperature due to heat dissipation at the outlet. The corresponding total failures observed across the entire server population over a

period of 1 year are denoted by the AFR line. Note that we present Annual Failure Percent (which is a measured population based value and should not be considered as the Annualized Failure Rate, which is a calculated metric that manufacturers provide) for our population that is on continuous mode of operation throughout the year (For a discussion on different annual failure rates, please see [2]). Out of the hard disk drives that were in the front-most part of the server chassis (HDD1), only 8% failed, whereas, for the fourth hard disk drive (HDD4) around 15% of the total disks failed. This is almost twice the number of failures compared to the front of the chassis. This result shows a strong correlation between temperatures observed through the SMART logs collected at the machines and the observed failures reported in this datacenter. In fact, the correlation coefficient measured across the entire population for (average temperature for drive locations inside the chassis, number of failures) pair is  $R = 0.79$ , which is significantly high. Our experience with this dataset does point out that lower temperature locations do have lower failures, and as system designers it is a strong motivation for reducing temperature impact inside a chassis design. Fan speed and airflow management helps reduce such temperature impact.

**Observation 1:** *There is a significant correlation ( $r = 0.79$ ) between actual hard drive temperature inside a server chassis design and the number of drive failures. Hence chassis design should incorporate temperature reduction optimizations.*

#### B. Correlation across Servers in a rack

A datacenter rack consists of multiple server chassis arranged on top of each other. The cool air comes through vents closer to the bottom of the rack and rises upwards. It is pulled across the server as it rises up and that direction is horizontal (as shown in Figure 2). However as it moves up through the vertical direction, there is an increase in air temperature due to heat dissipation. There are also other mechanical impacts such as the differences in air pressure (cfm) at different server locations within a rack. In this



**Figure 6.** Temperature and failures correlation within a rack ordered by server number;  $R = 0.91$

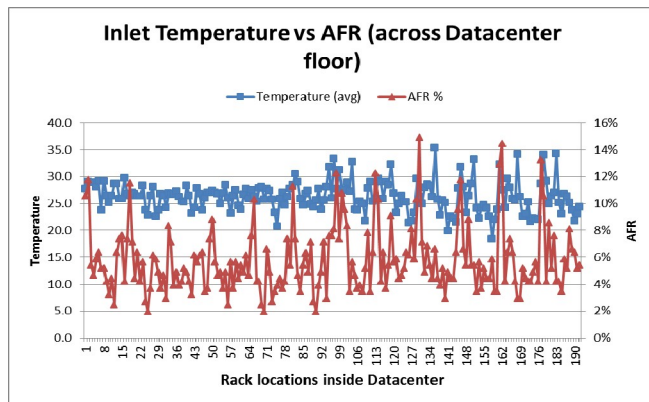
section we explore if the server location and inlet temperature observed at each location correlates with the number of disk failures observed at that server location.

From Figure 6, we see that for the cooler servers (Location 9, 10, 11, 12) that are on the bottom of the rack, the number of failures is lesser (10%) as compared to hotter servers (Location 2) at 13% failure rate. This shows a strong correlation between server locations inside a rack and the number of failures. This again reiterates our observation that temperature and air movement across a rack are significant determinants for server failures. The correlation coefficient computed for (inlet temperature for server location within rack, number of failures) pair is  $R = 0.91$ .

**Observation 2:** *There is a significant correlation ( $R = 0.91$ ) between the inlet temperatures observed with respect to the position of the server in the rack and number of failures for that server. The higher the average inlet temperature at a server location within a rack, the higher the number of failures.*

### C. Correlation across multiple rack locations

Having seen that drive bay location and server location temperatures are indeed major determinants for number of failures observed in that location, we also determine whether the temperatures observed across rack locations inside the datacenters are correlated to the number of failures observed.



**Figure 7.** Temperature and Failure correlation across rack locations inside datacenter:  $R = 0.30$  for entire population.

Figure 7 presents the temperature observed at the particular rack location (averaged across the servers in the rack). Every cluster in the datacenter has two columns of nine racks each. Each column has an inlet cold aisle and a corresponding hot aisle. Every rack has 12 servers.

One important observation from Figure 7 is that we would expect the Temperature line to be fairly horizontal at a fixed datacenter setpoint temperature. However this is not the case and there is significant variation in temperatures

across the datacenter floor. This is possible due to a variety of reasons including inefficient hot aisle/cold aisle containment, other networking or server gear venting hot air into the cold aisle and hot air recirculation around the edges. There are other significant patterns observable from Figure 7, especially that the rises in temperature are accompanied by rises in failures, however we note that there are several places in the figure where this is not the case. However, the correlation coefficient for the entire set of data (temperature at datacenter location, failures at that location) is  $R = 0.30$ . There is indeed a positive correlation and is statistically significant (critical value of  $R$  at  $df=120$  is 0.232 for a two-tailed test at level of significance = 0.01). Also, it is clear that the lower temperature racks have lower failures and hence the motivation to be temperature-aware in datacenter and server design is still valid.

**Observation 3:** *There can be varying degrees of deviation from the Datacenter setpoint temperature in different parts of the datacenter floor. Hence hot and cold aisle containment solutions are needed for higher efficiency in traditional datacenters.*

## V. RELIABILITY MODEL AND APPLICATION

In order to quantify the impact of different datacenter inlet temperatures experienced by the servers, we needed to come up with a model for measuring the reliability of the hard disk drives that are the primary failure components in the system. We used a physical Arrhenius model and estimated the activation energy based on the failures from the field. Earlier studies have estimated duty cycle has a negative effect on AFR (higher duty cycles have higher accelerated failures) [1]. We factor in the effect of duty cycle in the proportional multiplier for Arrhenius model in the next section. Using this model, we estimate the AFR (Annualized Failure Rate) and consider that to be a baseline for comparison between different datacenter inlet temperature decisions.

### A. Arrhenius model for AF

The failure rate due to elevated temperature is governed by the Arrhenius equation [1]. The Arrhenius acceleration factor (AF) can be expressed as

$$AF = A * e^{\left(\frac{E_a}{kT}\right)}$$

Where,

$A$ , is a proportional multiplier

$E_a$ , is the activation energy determined empirically

$K$ , is the Boltzmann's constant that relates energy at the particle level with temperature observed at macro level,

$T$ , is the absolute temperature observed at elevated temperature points respectively.

Acceleration Factor (AF) can also be expressed as the ratio between the time it took to fail under normal temperature versus the elevated temperature. Rewriting above equation,

$$\ln(t_2/t_1) = \left(\frac{E_a}{k}\right) * \left(1/T_2 - 1/T_1\right)$$

Where,  $t_2$  is the time for failure with elevated temperature and  $t_1$  is the time to failure with normal temperature.

Activation energy  $E_a$  [6], can be calculated from the above equation. We know empirically from Section IV-A that we had almost twice the number of failures with 12 C increase in temperature. Substituting the values in the equation, we get  $E_a = 0.464$  eV. We estimate the proportional multiplier (A) for the Arrhenius Acceleration Factor equation to be 1.25 based on workload duty cycle expectations. This multiplier is calculated as a function of the duty cycle expected and the duty cycle rated by the manufacturer (similar to [1]). We base our calculations on the worst case duty cycle for the workload (100%). We use the above empirically calculated value to compose the Arrhenius model for estimating Acceleration Factor at different temperatures.

HDD Temp	Acc Factor (AF)	AFR	AFR relative to 40 C
40 C	1.257	2.75	100%
41 C	1.328	2.91	106%
42 C	1.402	3.07	112%
43 C	1.480	3.24	118%
44 C	1.562	3.42	124%
45 C	1.648	3.61	131%
46 C	1.737	3.80	138%
47 C	1.831	4.01	146%
48 C	1.930	4.23	153%
49 C	2.033	4.45	162%
50 C	2.141	4.69	170%
51 C	2.254	4.94	179%
52 C	2.372	5.19	189%
53 C	2.495	5.46	198%
54 C	2.625	5.75	209%
55 C	2.759	6.04	219%

**Table 1.** HDD temperature and corresponding AFR (40C is baseline)

Table 1 shows the increase in Acceleration Factor and the corresponding impact on reliability (AFR). We use the 40C row as the baseline temperature and AFR value since it is derived from typical HDD manufacturer data sheets (eg: [13]). We see that operating the hard disk drive at 55C increases the AFR by almost twice when compared to the AFR quoted by manufacturers at 40C. The table provides a handy reference sheet for expected failures when the hard disk drive experiences a particular temperature. Given a chassis design, it is straightforward to compute the delta T

observed by the hard disk drives at different location inside the chassis. We computed the delta T from SMART logs and when running a constant workload at specific temperatures. We can then use Table 1 to estimate the failure rate for the particular chassis design, given the corresponding datacenter inlet temperature. Thus, this provides a methodology for selecting datacenter setpoint based on expected reliability.

### B. Application to Datacenter Setpoint Selection

This section discusses the application of Table 1 in selecting the datacenter setpoint temperature. The setpoint temperature determines the chilled water temperature. The lower the setpoint temperature required, higher the energy required by the chiller units to bring down the temperature. Hence, fixing an optimal setpoint temperature by a data-driven reliability-aware approach would lead to energy conservation and better efficiency at the datacenter.

Table 2 presents two server chassis design. One design contains the HDDs in the front, and therefore is exposed to the cold aisle. The delta T between the temperatures experienced by the front HDDs and the datacenter setpoint temperature is minimal (1 C). The other server design consists of the inner HDDs, which has HDDs arranged one behind the other. We present only the case of the worst HDD in the design. Because of preheating, the delta T in cold temperatures is 20C. However as the air gets hotter, the chassis fans will be sped up to prevent the HDDs from overheating with a delta T of 10C. For temperatures in-between, the delta T will be assumed to be linear. Hence at inlet of 50C, the hottest drive experiences a temperature of 60C. We assume that for temperatures below 40C there is no AFR increase and we keep that as baseline AFR and compute the relative AFR from that datapoint.

Inlet Temp		HDD's in Front, ΔT 1°C		Buried HDDs Design, ΔT 20°C cold de-rated to ΔT 10°C hot	
		HDD Case Temp	Relative AFR	HDD Case Temp	Relative AFR
10 C	50 F	11 C	100%	30 C	100%
15 C	59 F	16 C	100%	34 C	100%
20 C	68 F	21 C	100%	38 C	100%
25 C	77 F	26 C	100%	41 C	106%
30 C	86 F	31 C	100%	45 C	131%
35 C	95 F	36 C	100%	49 C	153%
40 C	104 F	41 C	106%	53 C	189%
45 C	113 F	46 C	138%	56 C	231%
50 C	122 F	51 C	179%	60 C	281%

**Table 2.** Choosing Datacenter Setpoint for a) HDDs in Front, b) Buried HDDs

As we can observe from the table, a front facing hard disk drive design experiences fewer failure events at 50C inlet temperature. However, the buried HDD design has significant increase in the relative AFR of the disk drives. Hence we need to make the decision about housing the second design in a datacenter more carefully. If the threshold for disk failures can be fixed, (say at 1.05X the advertised AFR rates, a 5% increase over baseline), then we need to adjust the datacenter setpoint inlet temperature for a datacenter having the second design at 25C. However, if all

our servers had the first design, then the setpoint temperature could be 40C. The 15C delta between these two setpoints is a significant temperature delta to operate a datacenter. A 15C difference in setpoint temperature is close to 150KW difference on the datacenter floor. Hence it is useful to have such a methodology in place for setting datacenter setpoint temperature.

**Observation 4: Datacenter setpoint temperature should be selected in a reliability-aware manner to avoid potential increases in server failures due to temperature impacts.**

## VI. COST DISCUSSION (FAILURE VS POWER USAGE)

We discuss the cost opportunity due to a reliability-aware design in this section. We take both types of server designs presented in Table 2 - (a) *HDDs in Front* and (b) *Buried HDD* design. We then evaluate two scenarios. Scenario 1 describes the design decision to have a conservative setpoint of 25 C for both the designs. Scenario 2 describes the design decision to fix an aggressive 40 C setpoint temperature for both.

Design	Scenario 1 Conservative 25C setpoint	Scenario 2 Aggressive 40C setpoint
HDD in Front	⊖ Lost opportunity for power reduction	⊕ No AFR increase; optimal power usage
Buried HDDs	⊕ No AFR increase, correct setpoint	⊖ AFR increase by 89% over current AFR

**Scenario 1 (Opportunity Cost):** In Scenario 1 (inlet temperature setpoint of 25C), it is clear from Table 2 that both the designs will not have AFR increase. However we are not utilizing the opportunity to operate the datacenter at increased temperature for *HDDs in Front* design which is still under acceptable AFR limits at 40C. The differential of 15C accounts for 150KW in power consumption by the cooling units in the datacenter. 150KW can power an additional 300 servers consuming 500W each (one disk consumes ~10W). Hence from the datacenter operational perspective this scenario is inefficient for the HDDs in Front design.

**Scenario 2 (Cost of increased failures):** In Scenario 2 (inlet temperature setpoint of 40C), we would make the datacenter operation efficient for the *HDDs in Front* design. However for the *Buried HDD* we would see additional 89% failures from the current failure rate. If we fit these numbers to our case study population and if we assume original AFR to be 5% and that all drives were at the elevated temperature due to an inlet temperature of 40C, then we would see thousands of additional drives failing out of the population. Including drive replacement, service and warranty cost on the server, this is indeed a significant concern for datacenter

operations. The cost of a failure is estimated to be 15% of the server cost [18]. When we include the server downtime due to increased failures, this design decision could indeed prove costly. Hence choosing operating temperature for either design should be a reliability-aware decision.

In summary, our work motivates the need to have a data-backed decision on selecting operating temperatures inside a datacenter. Even under relaxed datacenter environmental specifications, there could be significant economic loss due to increased failures if we decide to be over-aggressive in energy efficiency and operate at high temperatures with less cooling infrastructure. On the other side, we might lose the potential advantage of operating at higher temperatures for certain server designs, if we do not understand server reliability under high temperatures. Hence, a deep understanding of server design and a calculated approach is necessary to choose optimal operating conditions for increased efficiency in datacenter.

## VII. RELATED WORK

Server component failures and reliability are yet to be understood completely. There has been several previous research in this field, but have generated conflicting results, especially in relation to subjects like the impact of temperature on disk drive failures. With respect to large scale installations, Gray et al [4] observed failure rates ranging from 3.3-6% in two large web properties at Microsoft. Schwartz et al [17] report failure rates of 2-6% in the drive population at the Internet Archive. Elerath et al [3] report that end-user failure rates can be as much as ten times higher than what the drive manufacturer might expect in their study on server class disk drives. Schroeder et al [15] find that in the field, annual disk replacement rates typically exceed 1%, with 2-4% common and up to 13% observed on some systems. The authors also present interesting per-component failure percentages for three different type of systems that they considered. They also report a significant overestimation of mean time to failure by manufacturers. Schroeder et al [14] in their study of failures in petascale computers review sources of failure information for compute clusters and storage systems, and project corresponding failure rates.

One of the most closely related works is by Pinheiro et al [Pinheiro2007], which identified correlation between disk errors and SMART attributes from a large population of serial and parallel ATA drives. This paper also concluded that temperature and activity levels had less correlation to disk failures and was a surprising result when compared to previous studies [19, 1]. Yang et al [19] establishes that a 25 C delta in temperature derates the MTTF by a factor of 2 in their study on Quantum hard disk drives. Cole et al [1] from Seagate, present thermal de-rating models showing that MTTF could degrade by close to 50% when going from operating temperatures of 30C to 42C. Our results agree with the observations made by Cole. Also interestingly,

Vishwanath et al [18] report no correlation between failures and location of servers within a rack. We find in our case study that temperature does have a strong correlation to failures (within chassis, racks and across racks). We propose that temperature impacts for datacenter scale environments should be factored in knowing the server configuration and datacenter inlet temperature range. We present a summary of the observations made during the course of this study in the next section.

## VIII. EXPERIENCE SUMMARY

This section summarizes a few important lessons we learned during the course of our analysis that led to better datacenter and server chassis designs. One major lesson was that the deviation between inlet temperature observed at the server and the datacenter setpoint temperature is quite significant. Setting the datacenter operational temperature lower does not necessarily reduce the temperature experienced by servers in different parts of the datacenter. Hence hot aisle and/or cold aisle containment solutions are necessary for traditional datacenters to prevent significant temperature delta across the datacenter floor.

Another important factor that affected server component failure was the server chassis design relative to cooling solutions. Since hard disk drives were arranged horizontally behind one another in the path of air flow, the air gets preheated before it reaches the deeper hard disk drives. Hence though the server inlet temperature could have been under specifications, hard disk drives that were internal were continuously operating at higher temperatures. The failure rates reflected this phenomenon in Section IV-A. These lessons motivated us to construct reliability models that can estimate expected failures for different datacenter and chassis designs. We incorporate this into our design cycles and optimize for efficient datacenter operation for large scale online services. Note that the temperature impact can be countered by different cooling mechanisms including fan speeds and heat sinks.

## IX. CONCLUSION

Our paper shows significant correlation between temperature and observed failures across drive locations inside a chassis, server locations in a rack and rack locations inside a datacenter. Our study also highlights that temperature-aware datacenter design is critical to reduce server failures. We present a model that can be used with relative ease to estimate hard disk drive failures due to elevated temperature. Using the model, we present a reliability-aware methodology for setting datacenter operating temperature. We propose that datacenter and server chassis designers use these reliability specific metrics in the design phase to reduce the number of server failures in a large scale datacenter.

## REFERENCES

- [1] Gerry Cole. Estimating drive reliability in desktop computers and consumer electronics systems. SeagateTechnology Paper TP-338.1, November 2000
- [2] J. G. Elerath. AFR: problems of definition, calculation and measurement in a commercial environment. In Proc. of the Annual Reliability and Maintainability Symposium, 2000.
- [3] Jon G. Elerath and Sandeep Shah. "Server class disk drives: How reliable are they?" In Proceedings of the Annual Symposium on Reliability and Maintainability, pages 151 – 156, January 2004.
- [4] Jim Gray and Catherine van Ingen. Empirical measurements of disk failure rates and error rates. Technical Report MSR-TR-2005-166, December 2005.
- [5] S. Greenberg, E. Mills, W. Tschudi, P. Rumsey, and B. Myatt. Best practices for data centers: Lessons learned from benchmarking 22 data centers. ACEEE Summer Study on Energy Efficiency in Buildings, 2006.
- [6] D.J. Groebel, A. Mettas, F.B. Sun; "Determination and interpretation of activation energy using accelerated-test data," Reliability and Maintainability Symposium, 2001. Proceedings. Annual , vol., no., pp.58-63, 2001
- [7] G. Guo and J. Zhang. Feedforward Control for Reducing Disk-Flutter-Induced Track Misregistration. IEEE Transactions on Magnetics, 39(4):2103–2108, July 2003
- [8] J. Hamilton. "An Architecture for Modular Data Centers". In CIDR 2007.
- [9] Urs Hoelzle , Luiz Andre Barroso, The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Morgan and Claypool Publishers, 2009
- [10] HP White Paper, "Assessing and Comparing Serial Attached SCSI and Serial ATA Hard Disk Drives and SAS interface", October 2003
- [11] Michael Isard, Autopilot: Automatic Data Center Management, in Operating Systems Review, vol. 41, no. 2, pp. 60-67, April 2007
- [12] E. Pinheiro, W. D. Weber, and L. A. Barroso. Failure trends in a large disk drive population. In Proc. of the FAST '07 Conference on File and Storage Technologies, 2007
- [13] Seagate Constellation ES drive datasheet.
- [14] Bianca Schroeder, Garth Gibson. "A large scale study of failures in high-performance-computing systems." International Symposium on Dependable Systems and Networks (DSN 2006).
- [15] Bianca Schroeder , Garth A. Gibson, "Disk failures in the real world: what does an MTTF of 1,000,000 hours mean to you?", Proceedings of the 5th USENIX conference on File and Storage Technologies, February 13-16, 2007, San Jose, CA
- [16] Bianca Schroeder , Eduardo Pinheiro , Wolf-Dietrich Weber, DRAM errors in the wild: a large-scale field study, Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems, June 15-19, 2009, Seattle, WA, USA
- [17] Thomas Schwartz, Mary Baker, Steven Bassi, Bruce Baumgart, Wayne Flagg, Catherine van Ingen, Kobus Joste, Mark Manasse, and Mehul Shah. Disk failure investigations at the internet archive. 14th NASA Goddard, 23rd IEEE Conference on Mass Storage Systems and Technologies, May 2006.
- [18] K. V. Vishwanath and N. Nagappan. Characterizing cloud computing hardware reliability. In Proceedings of the 1st ACM Symposium on Cloud Computing (Indianapolis, Indiana, USA, June 10 - 11, 2010). SoCC '10
- [19] Jimmy Yang and Feng-Bin Sun. A comprehensive review of hard-disk drive reliability. In Proceedings of the Annual Symposium on Reliability and Maintainability, pages 403 - 409, January 1999