

# Disk Drive Workload Captured in Logs Collected During the Field Return Incoming Test

Alma Riska

Erik Riedel

**Seagate**

We turn on ideas



# Motivation

- Understand the workload of an entire family of drives
  - Salient characteristics across a set of applications that are targeted by a drive family
- Measurement on single systems (either operational or experimental) may be difficult to generalize over the entire usage scenarios for a drive family
- Complement the analysis obtained from detailed measurements in system with coarse grain information obtained over a larger set of drives

# Data Source

- Drives returned back to Seagate
  - Enterprise-level drives
  - Cheetah 10K RPM
  - Cheetah 15K RPM
- Upon return drives go through the Field Return Incoming Test
  - Extracts information from drives
    - It may not be possible if the drive is failed
  - Information consists of cumulative attributes
- Attributes analyzed in this paper
  - Power-on hours
  - Spin-ups
  - Bytes READ
  - Bytes WRITTEN
  - SEEKs

# Available drives per family

- Drives are grouped based on the family (10K or 15K)
  - Although used in enterprise systems the set of applications and reliability and performance requirements are expected to be different for the two families
- Drives are also distinguished by their lifetime
  - Drives less than a month are expected to have experience only integration activity
  - Drives more than a month are expected to have experienced both integration and normal in-field operation
    - More information about how drives are used in the field from the applications they support

Drive Family	Total drives	Less than 1 month	More than 1 month
Cheetah 10K	197,013	43,999	153,014
Cheetah 15K	108,649	19,557	89,092

# Advantages of the data set

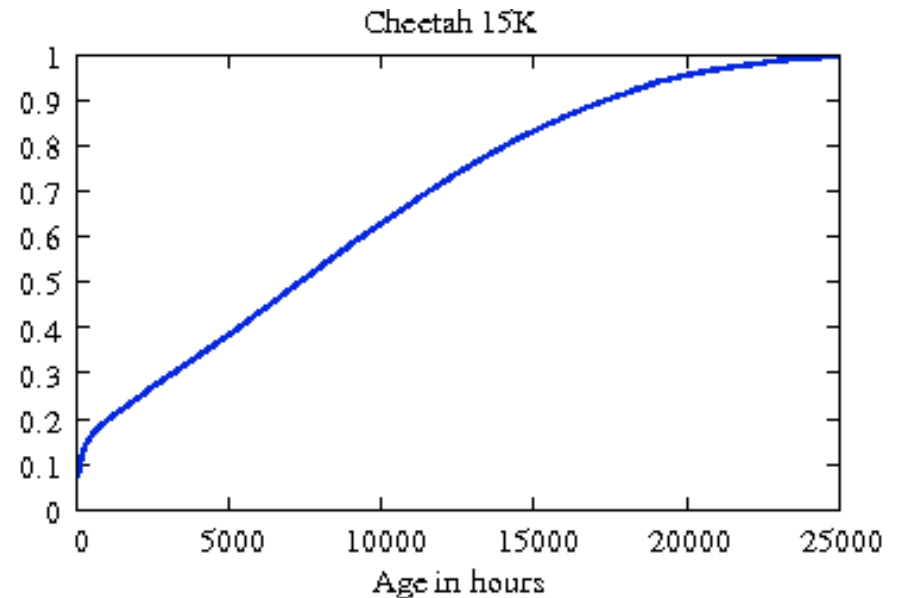
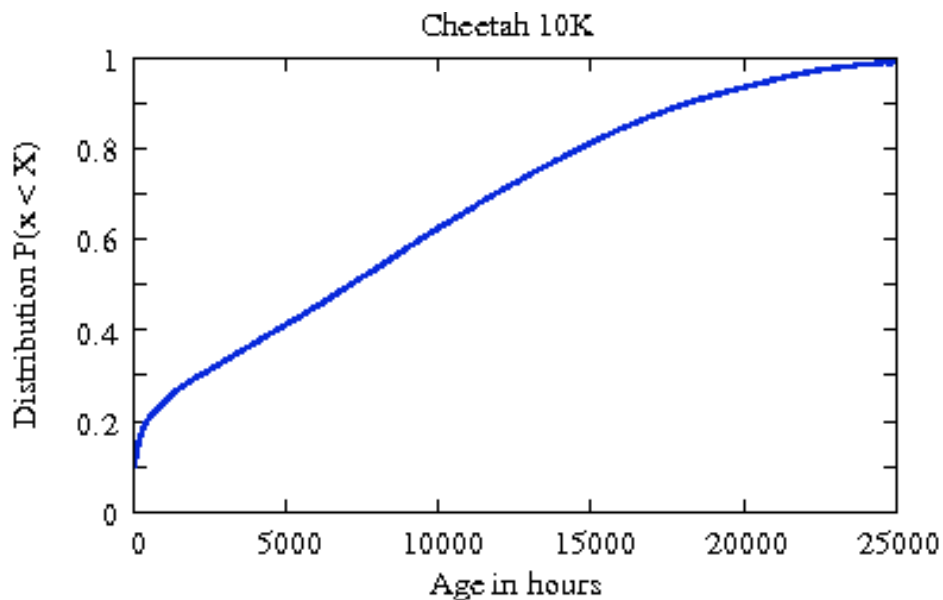
- A large set of available drives
  - Thousands of drives
- Drives in the field integrated in the entire spectrum of storage systems
  - Each type of storage system is represented accordingly to its popularity in the set
- We expected to capture correctly the dynamics across the entire spectrum of applications supported by a given drive family

# Drawbacks of the data set

- Granularity of the information in the data set
  - One cumulative number per attribute
  - Only average behavior
  - Dynamics of workload behavior over time is not available here
- The data set comes from drives returned because of issues
  - The selection of the set is not random and may be biased
  - Drives with issues may experience heavier load than other drives
  - We have evaluated here only drives whose data is still accessible (severity of the problem is low)

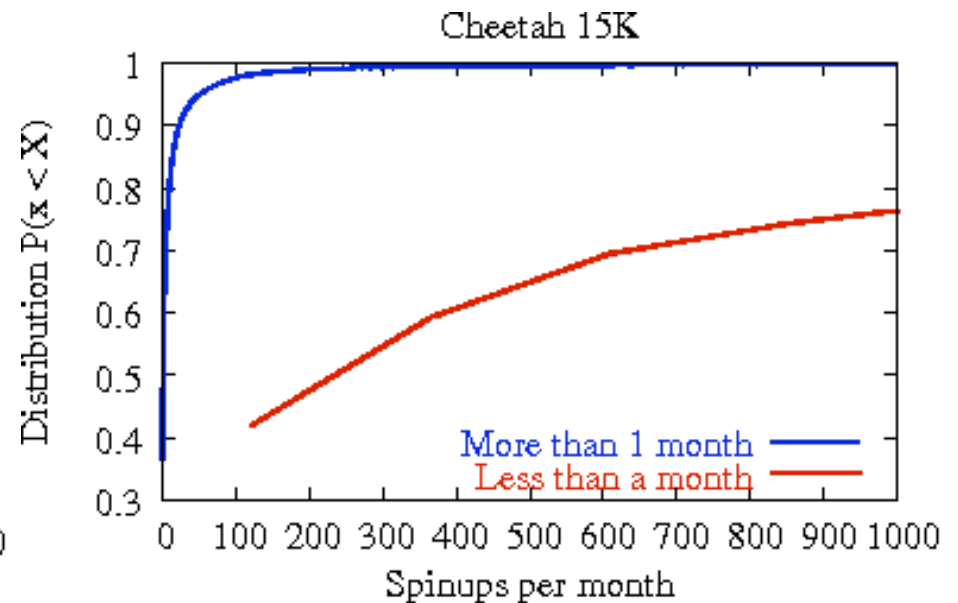
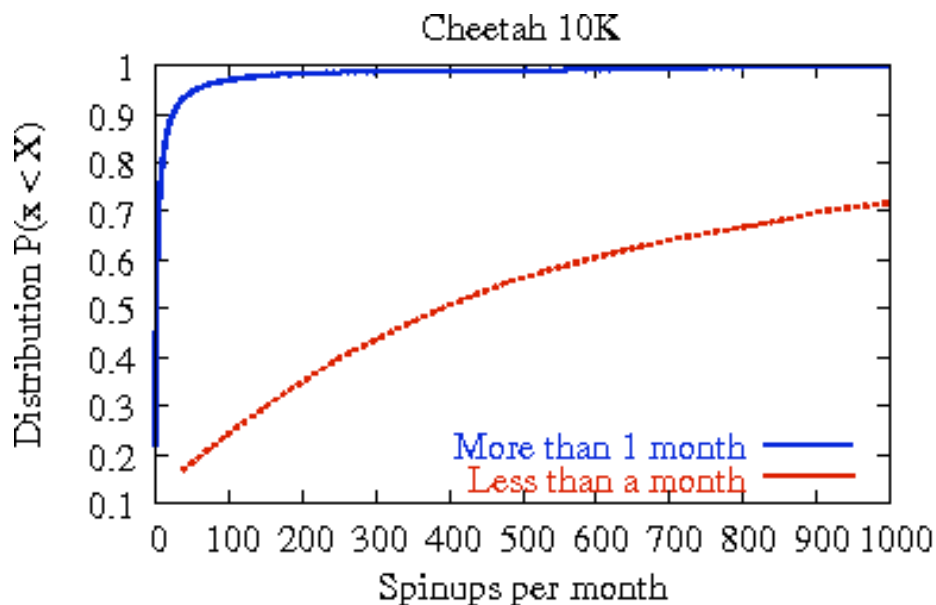
# Power-On hours of the drives

- In our set there is a wide range of drives
  - From drives with a few months of life to drives that have been in the field for about three years
    - The drive families we analyze have been in the field for at most 3 years
  - 50% of drives more than a year in the field



# Powering up/down the drives

- Distribution of the number of Spin-ups per month is significantly different in drives more than a month old and less than a month old
- During integration drives go through more intense operation
- During normal operation drives almost never get shutdown





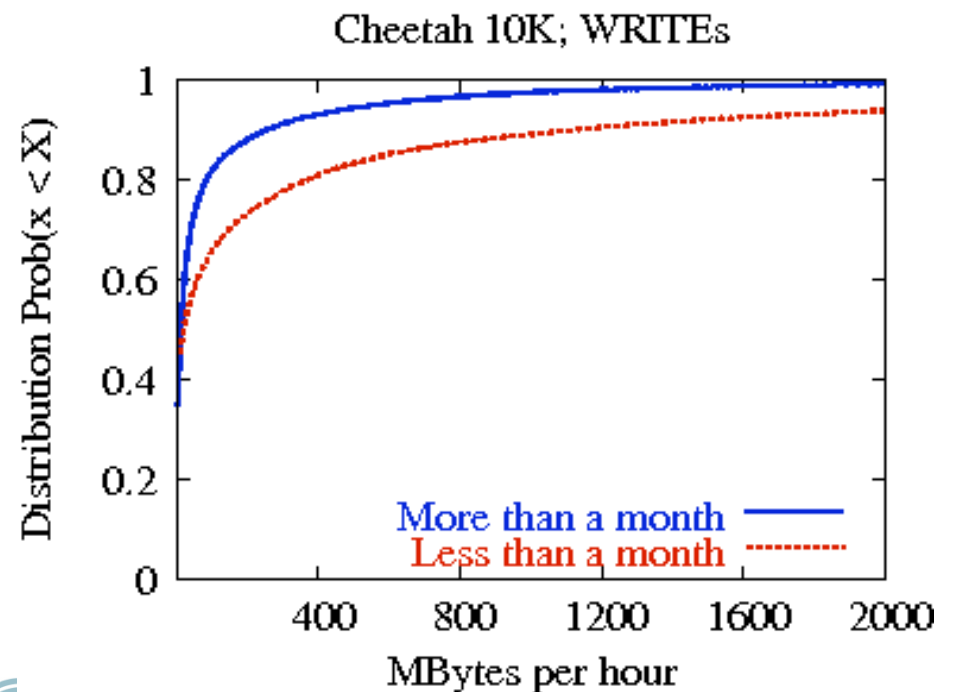
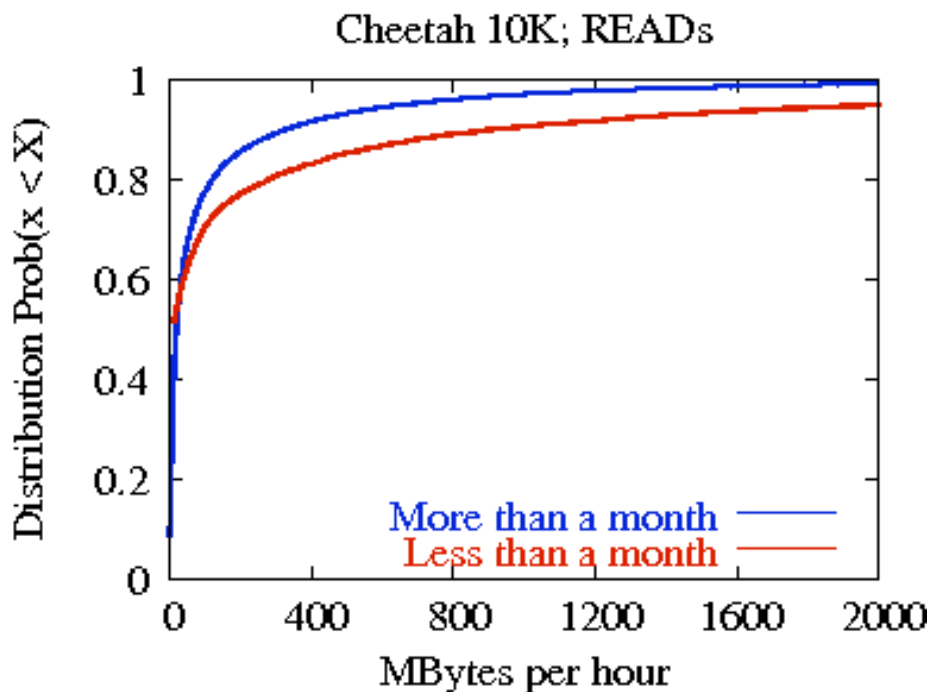
# Mbytes READ and WRITTEN per hour

Less than one month old drives				
Drive Family	Mean READ	CV READ	Mean WRITE	CV WRITE
Cheetah 10K	348 MB	7.85	394 MB	5.83
Cheetah 15K	250 MB	3.58	436 MB	2.33

More than one month old drives				
Drive Family	Mean READ	CV READ	Mean WRITE	CV WRITE
Cheetah 10K	140 MB	2.94	127 MB	4.91
Cheetah 15K	191 MB	2.98	197 MB	3.34

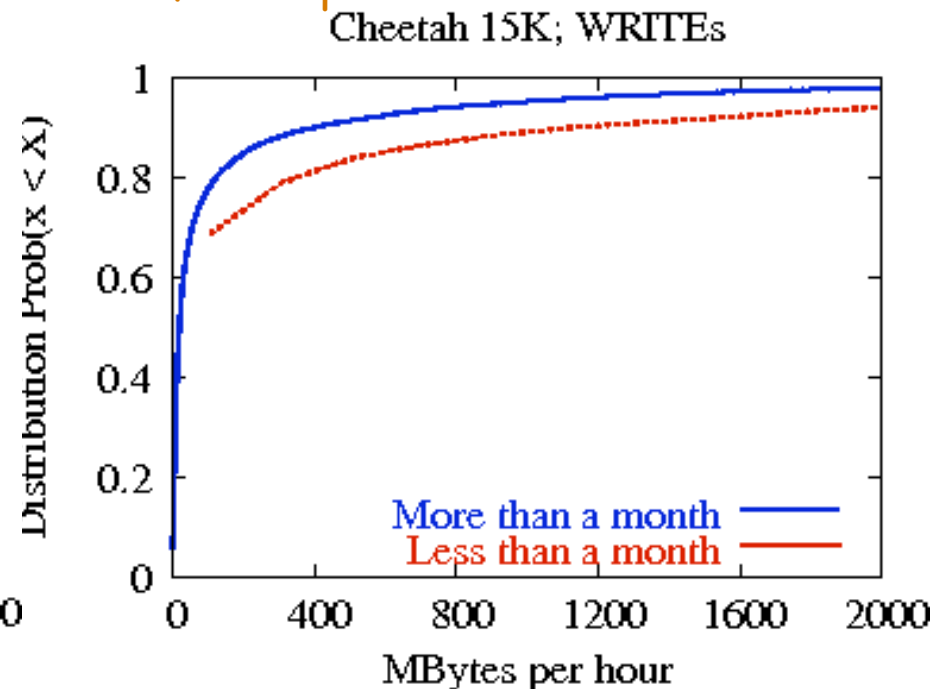

# Bytes READ/WRITTEN - 10K RPM

- Distribution of Mbytes READ/WRITTEN per hour
- Variability much higher during integration
- More stability during normal operation
- Median at 15-20 MB per hour
  - Drives lightly utilized, 3% more than 1GB per hour throughout lifetime
- Except integration, READs/WRITEs behave similarly



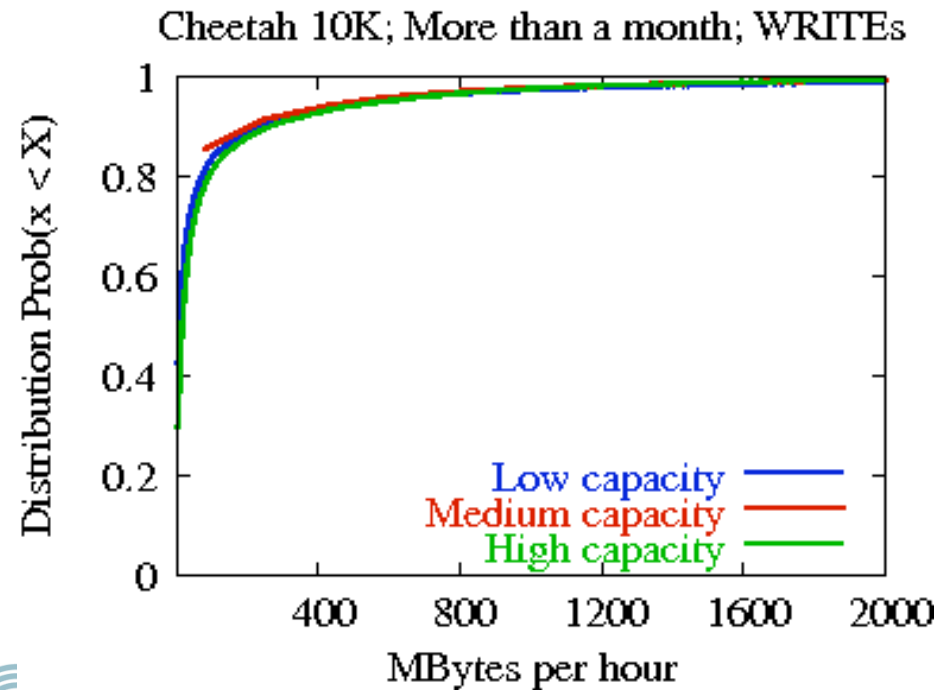
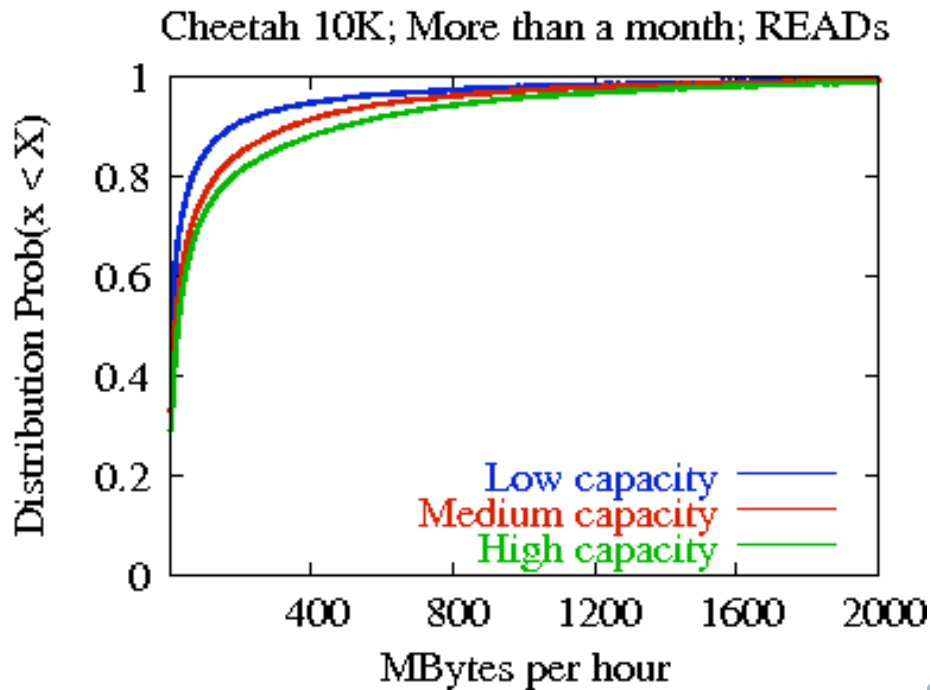
# Bytes READ/WRITTEN - 15K RPM

- Distribution of Mbytes READ/WRITTEN per hour
- Less difference between two age-based categories than in 10K RPM drives
- READs/WRITEs behave similarly
- Slightly more work than 10K RPM drives
  - 5% of drives transfer more than 1GB of data per hour



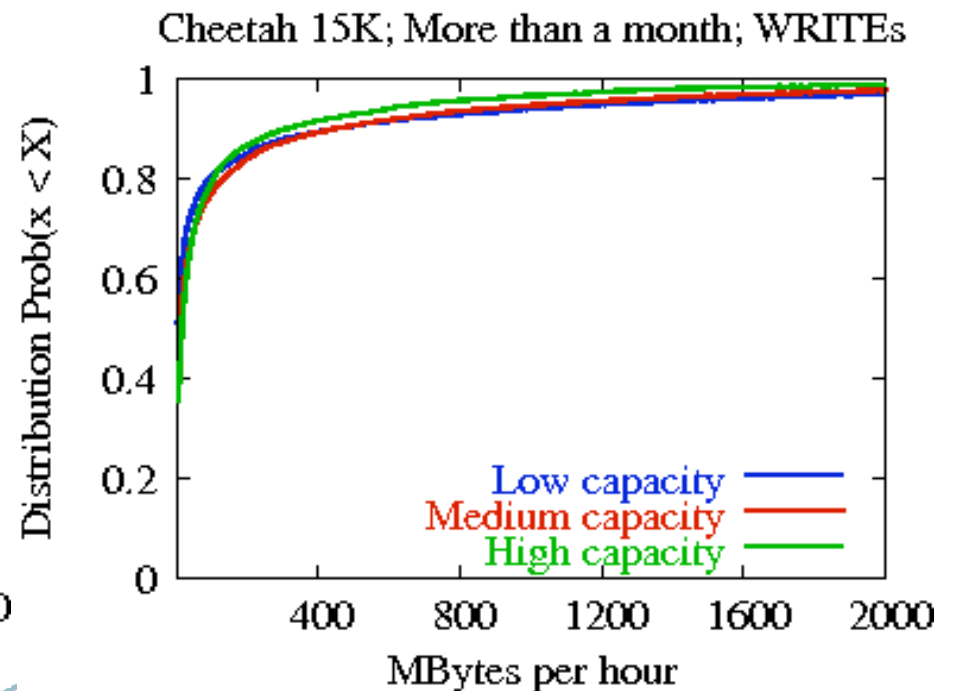
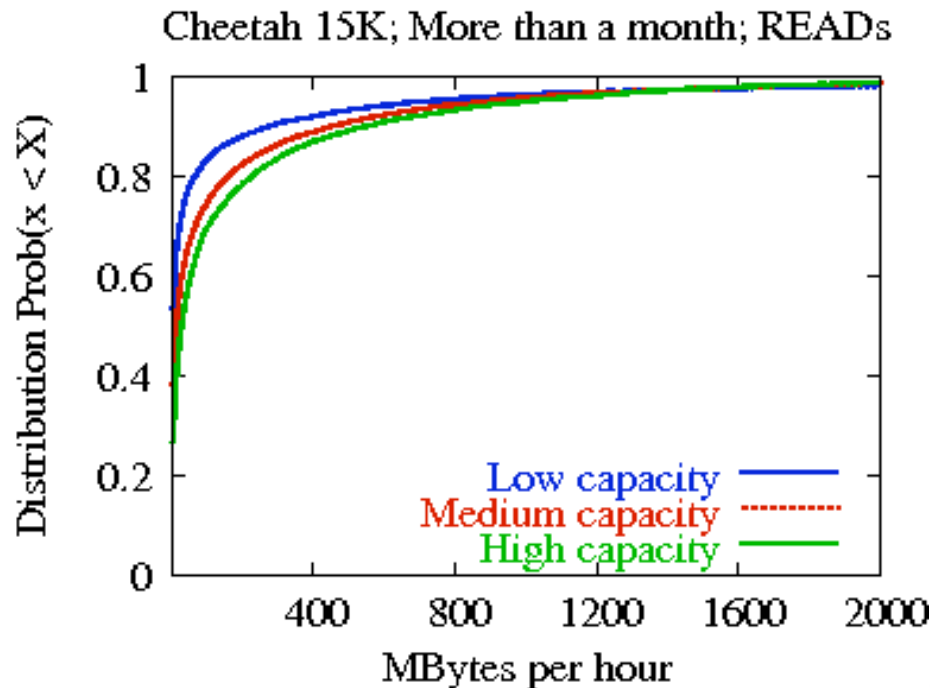
# Bytes READ/WRITTEN by capacity - 10K RPM

- Each drive family has three capacities and we further categorize our drives based on their capacity
- There is almost no difference in the amount of bytes written
- There is difference in the amount of bytes read



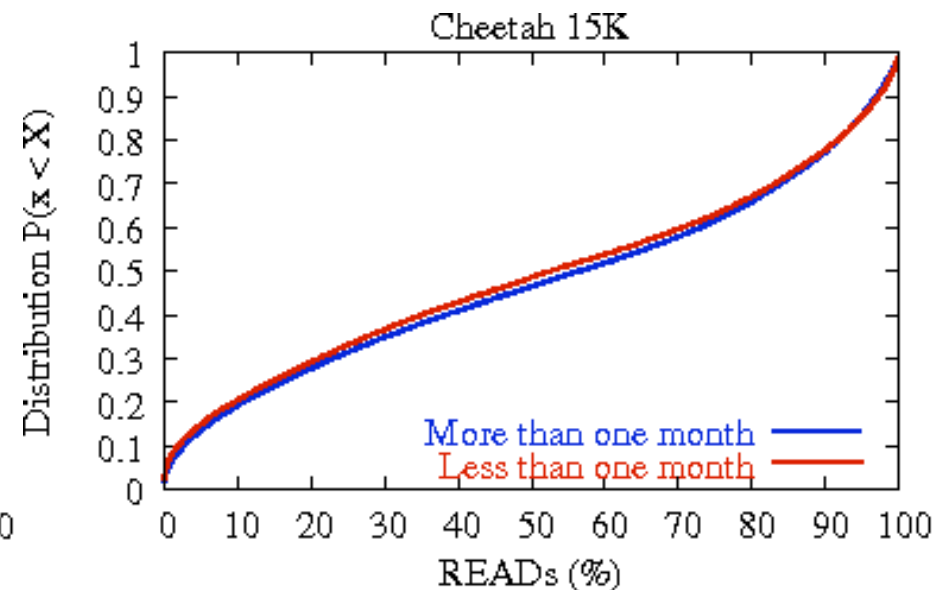
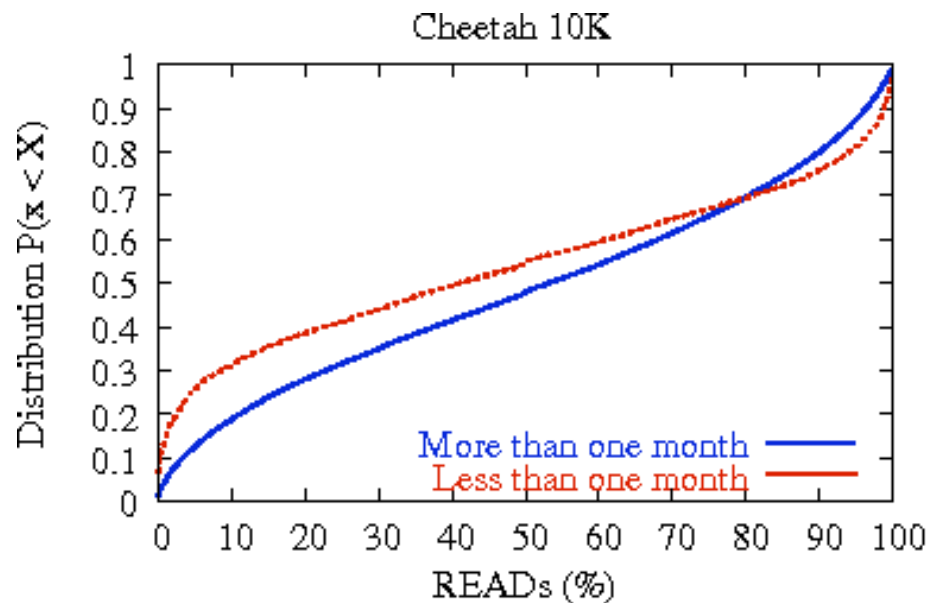
# Bytes READ/WRITTEN by capacity - 15K RPM

- Each drive family has three capacities and we further categorize our drives based on their capacity
- There is almost no difference in the amount of bytes written
- There is difference in the amount of bytes read



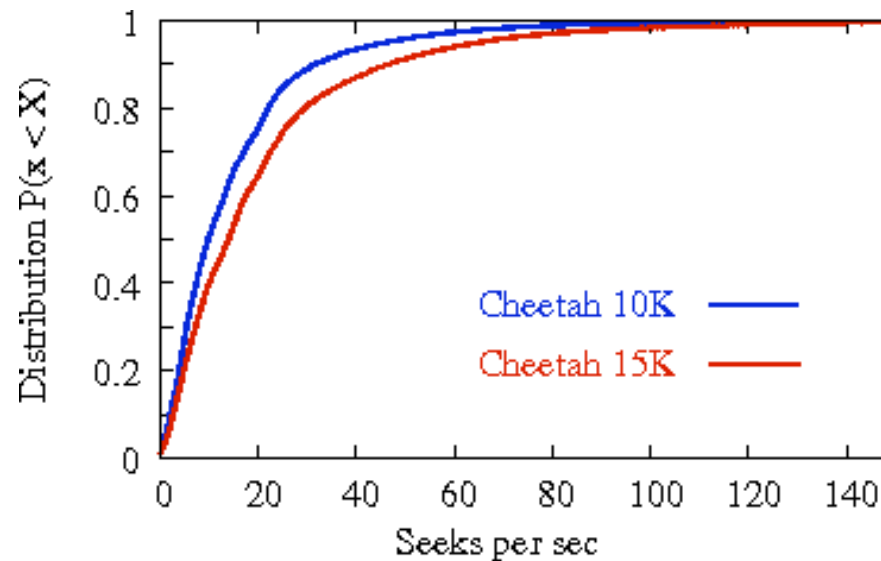
# READ/WRITE ratio

- We estimate the portion of bytes READs in the total bytes transferred
- In 10K RPM drives there is more writing during integration which means that their data is stored in the beginning and accessed afterwards
  - This is not the case for 15K RPM drives
- More drives seem to WRITE more than they READ
- This is very significant when it comes to the special handling of the WRITE traffic throughout the IO path



# Distribution of Seeks

- Among the available information is the cumulative number of seeks
- Estimate the average number of seeks per second
- 15K RMP complete more seeks than 10K RPM drives in average
  - 30 seeks per second vs. 40 seeks per seconds
- A seek takes in average a few (2-3) ms
  - 5-10% of drives with about 70-100 seeks per second which represents the portion of drives experiencing consistent high utilization throughout their life



# Conclusions

- Information extracted from the returned drives is used to characterize workloads processed by an entire drive family
- During the Field Return Incoming Test, the number of spin-ups, amount of bytes READ and WRITTEN, and the amount of SEEKS over the entire lifetime of a drive were analyzed
- Workload variability across drives is high in particular during the early life (integration) of the drive
- Drives are generally underutilized with about 10% of them experiencing utilizations that would affect the availability of idleness available to complete activities that enhance reliability, performance, and power consumption
- The majority of drives WRITE more than they READ