SDC 19

September 23-26, 2019
Santa Clara, CA

Data Science of QoS

Dave Holmstrom
SSD System Architecture
Micron Technology, Inc.

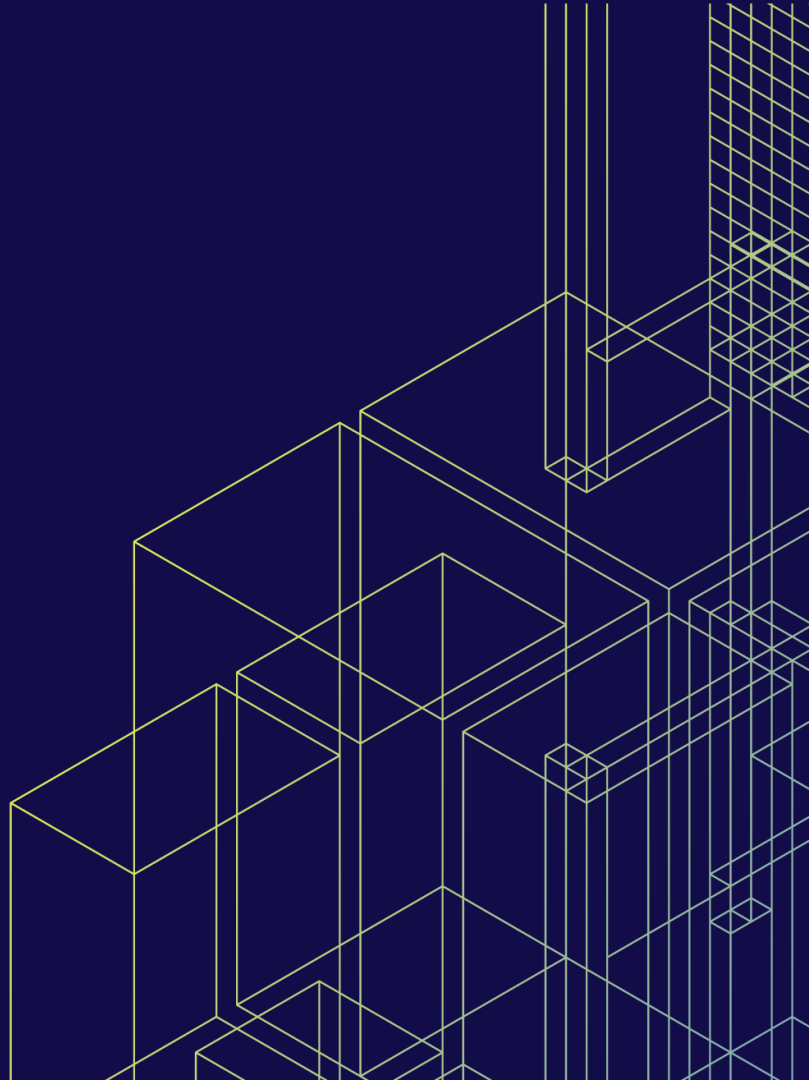

Disclaimer

- This is about data
- *Not workloads*
- *Not SSDs / HDDs*
- *Not success criteria*

QoS

Quality of Service

- The term Quality of Service (QoS) is a measure of latency (quicker is better) at a specific consistency level
- Example: A pizza company says they can deliver in 30 minutes or less. If out of 100 deliveries, one of them is later than 30 minutes, the QoS for the pizza delivery is 99/100, or 99% below 30 minutes

QoS Traditional View

- In this case, 1,000,000 IOs are measured, and a very large portion (999,000) are complete under 44.3ms
- There are 1,000 I/Os are outliers taking longer than 44.13ms.
- QoS is 99.9% (999,000/1,000,000) = 44.13ms. This is often referred to as “Three nines at 44 milliseconds”

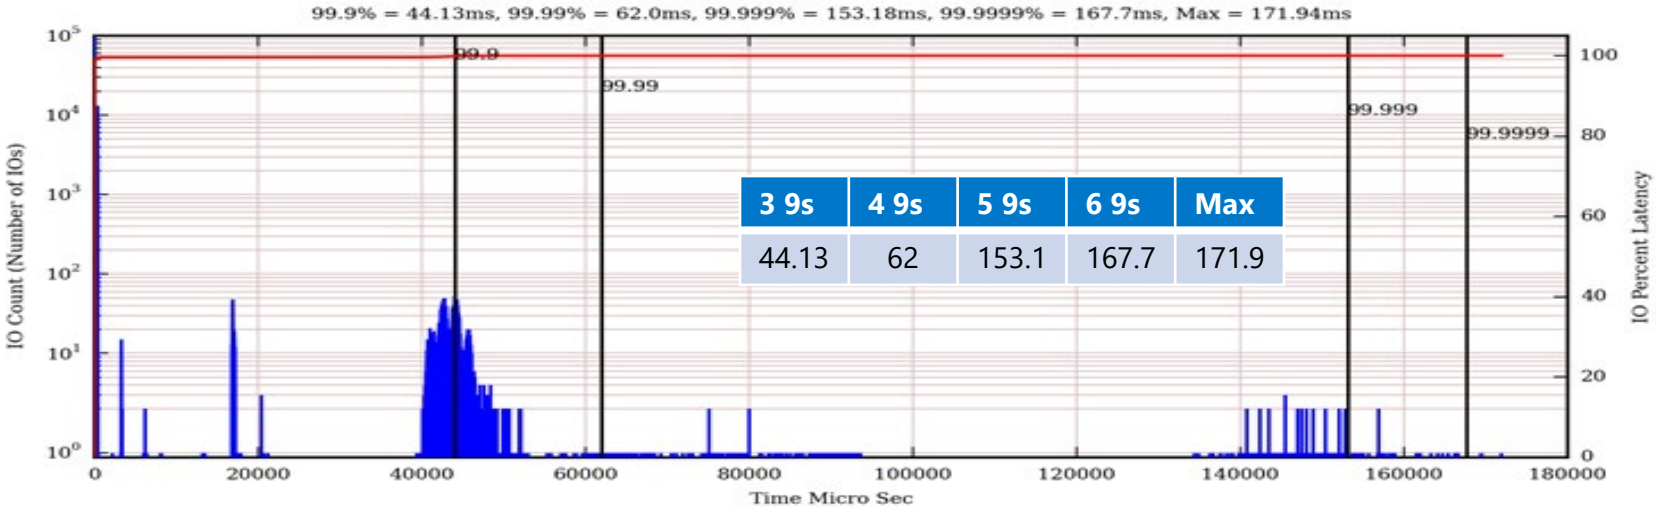

3 9s	4 9s	5 9s	6 9s	Max
44.13	62	153.1	167.7	171.9

Improving QoS Measure

Ensure Consistency

- Device at steady state for the current workload and test platform
- IOPS level should be stable across the measurement

In this example, the first interval shows higher performance and will shift QoS

Percentiles should be based on good data

- Good data requires
  - Repeatability
  - Reliable data up to the level of percentile of interest
  - Understanding the measured workload
- Reliable data cannot stand on any one single datapoint. Individual datapoints can be unreliable and can represent random events.

Sample Size MATTERS

- Sufficient samples are needed for mathematical significance in a QoS measurement
- According to the Central Limit Theorem, if a large uniform distribution of sample data is not a “normal distribution”, statistically, the minimum number of samples should be well beyond “just enough”
- For storage latencies, at least **100x more samples** beyond the base requirement is needed to ensure the measurement is accurate and repeatable (i.e. 1,000,000 points are needed to yield a reliable 4 9s)

Sample size effect on resolution

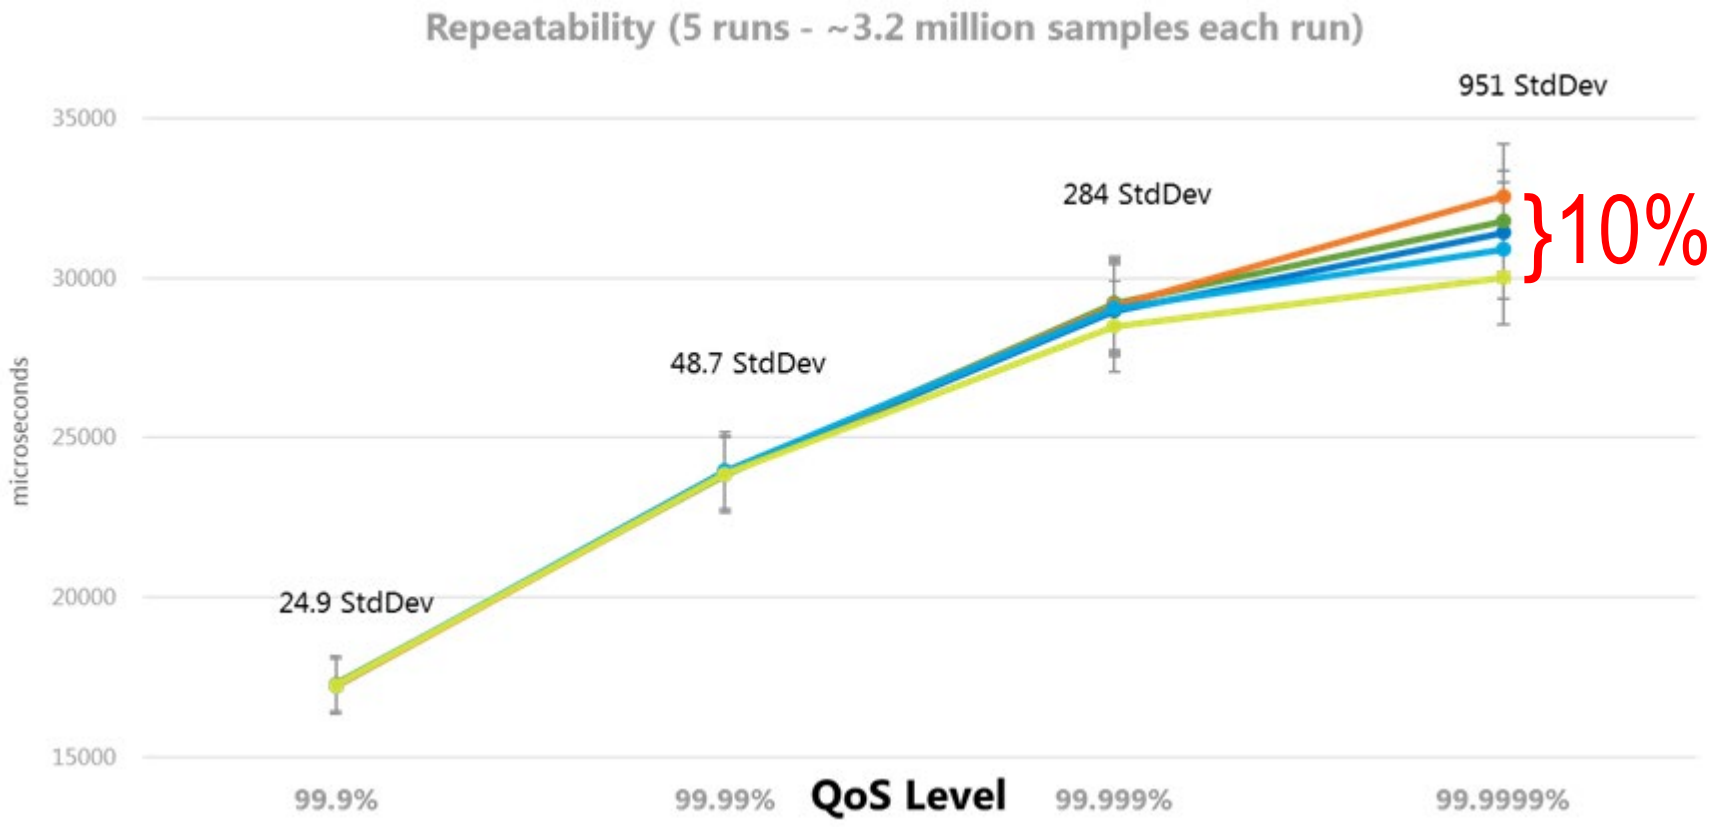

 D. Holmstrom 9/26/2019.

Calculating percentiles

- Default percentile calculations can return values that were not measured
- Use the next ***higher*** latency when no latency is within a calculated percentile

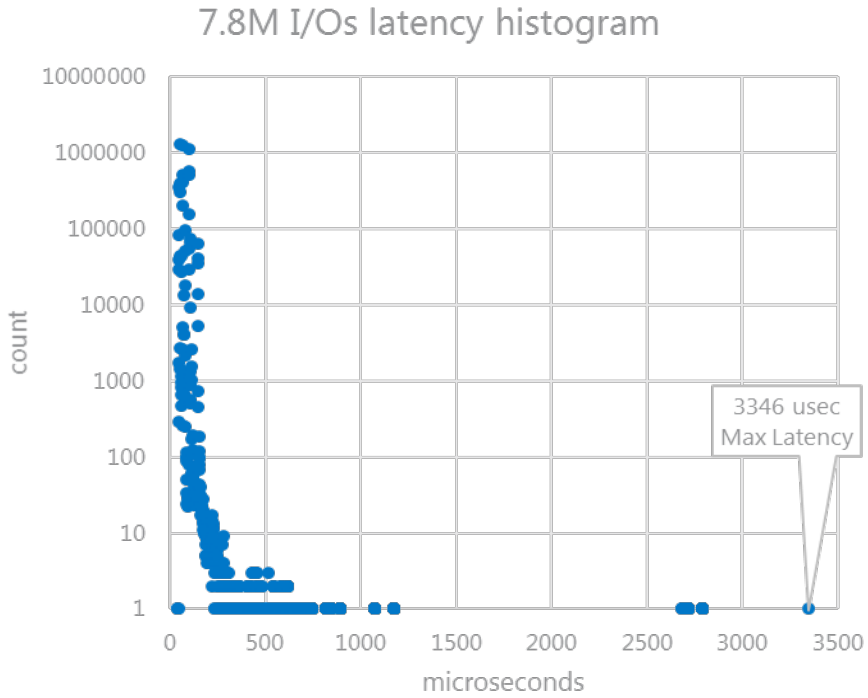

Enough data to calculate
Up to 6 9s are measured values

QOS levels (usec) 7.8M samples

QOS usec	3 9s	4 9s	5 9s	6 9s	7 9s	8 9s	9 9s	Max
Next Higher	148	180	455	2686	3346	3346	3346	3346
Linear (default)	148	180	453	2683	2907	3302	3341	3346

Not enough data to calculate
These are estimated values
causing misleading data

A better view of QoS

Waterfall View: QoS is more than a number

- Traditional views of histograms and tables don't tell the whole story
- Implementing a 1-CDF waterfall view illustrates continuous QoS levels
- 1-CDF clarifies driving factors for levels of 9s
- This view, when combined with histograms, to gives insights into the innerworkings of a device

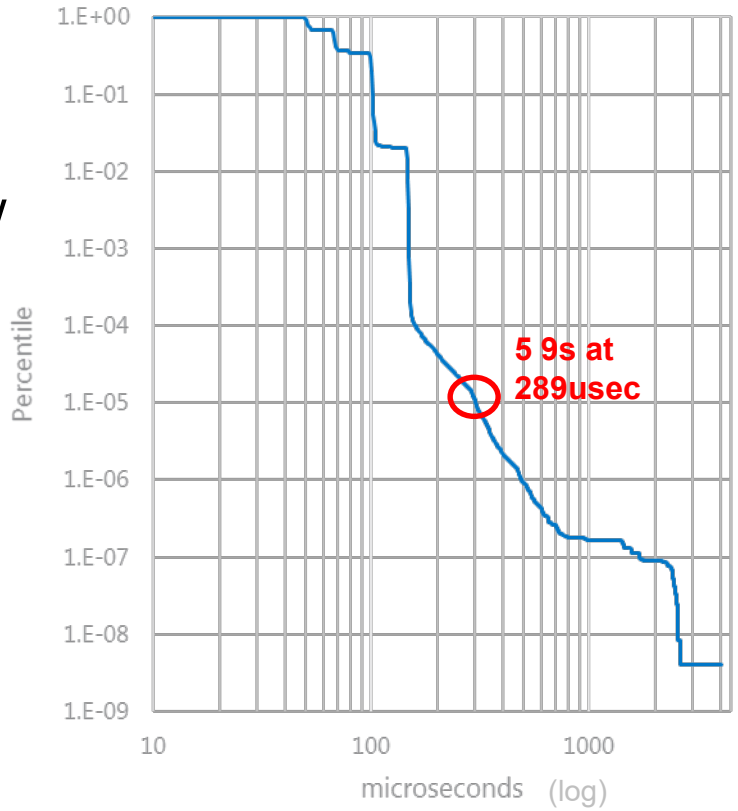

Combined view

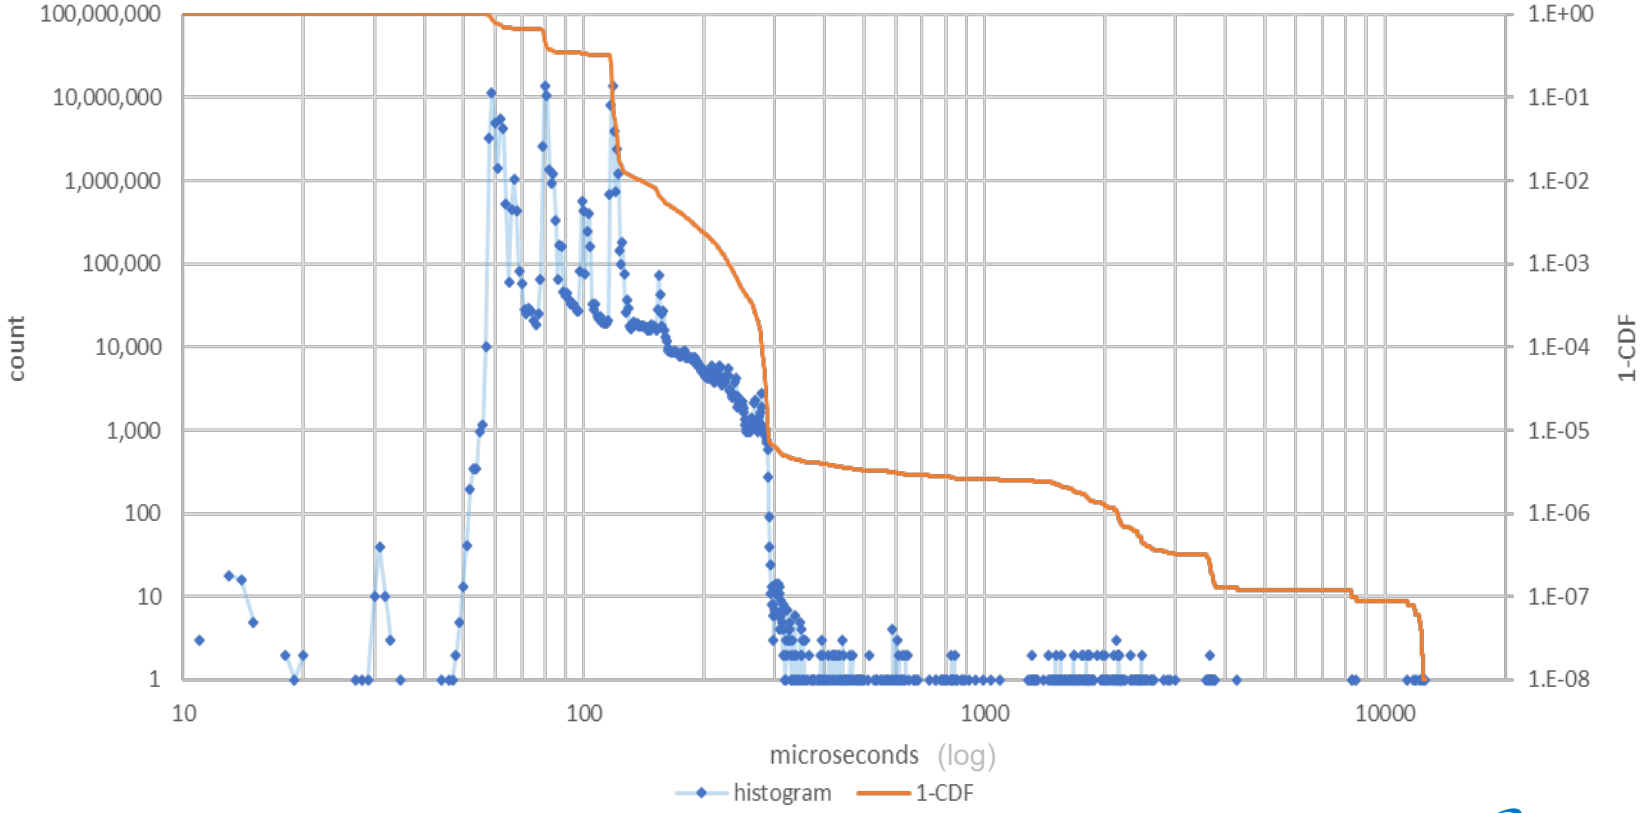

Multiple runs reveal variations

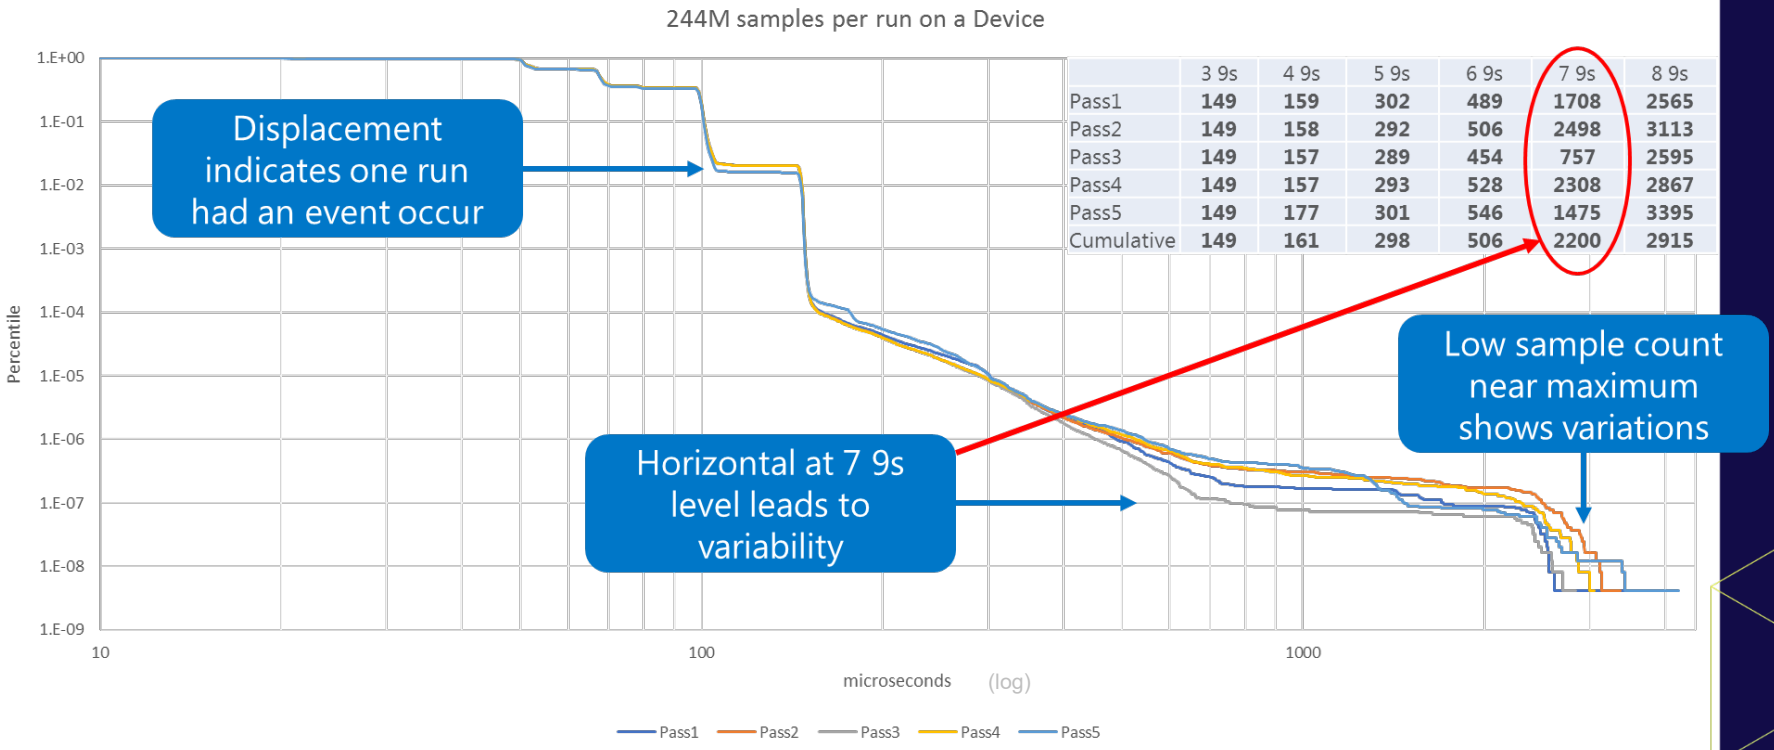

	3 9s	4 9s	5 9s	6 9s	7 9s	8 9s
Pass1	149	159	302	489	1708	2565
Pass2	149	158	292	506	2498	3113
Pass3	149	157	289	454	757	2595
Pass4	149	157	293	528	2308	2867
Pass5	149	177	301	546	1475	3395
Cumulative	149	161	298	506	2200	2915

Run Time QoS

QoS in real time

- Calculating QoS on an interval returns point in time of service levels
- The service levels can be aggregated across the test interval
- Output indicates variability and expose periodic long latency events

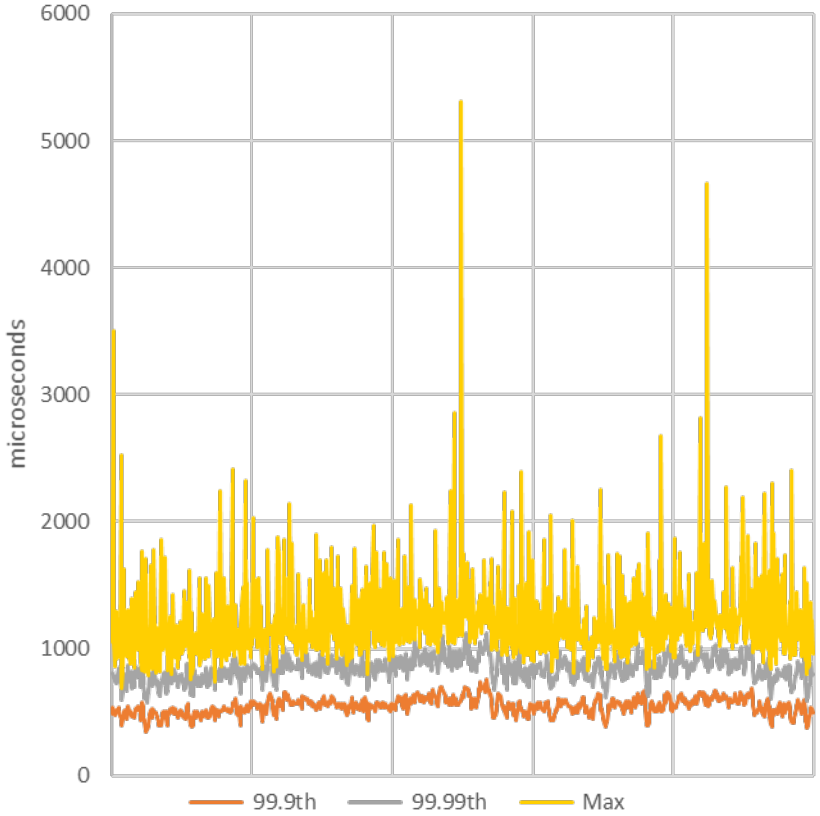

Previously unseen behavior emerges

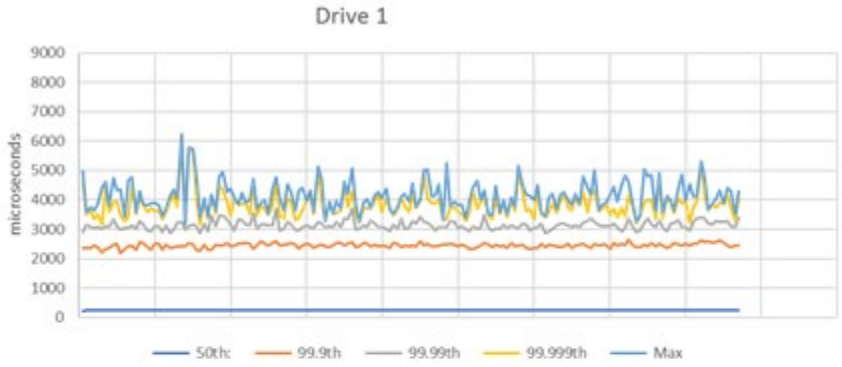

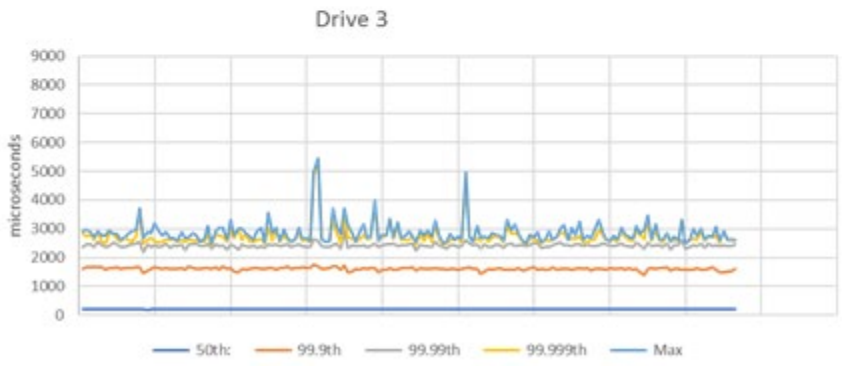

A New View of QoS

Average	14.6116	Median	11	StdDev	34.83609	Min	0	Max	5309	Sample Co	99999972
99.9:	556	99.99:	855	99.999:	1124	99.9999:	1615	99.99999	2388		

Write

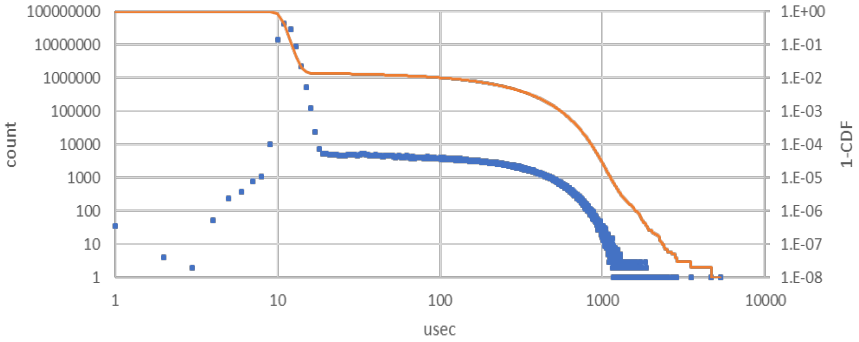

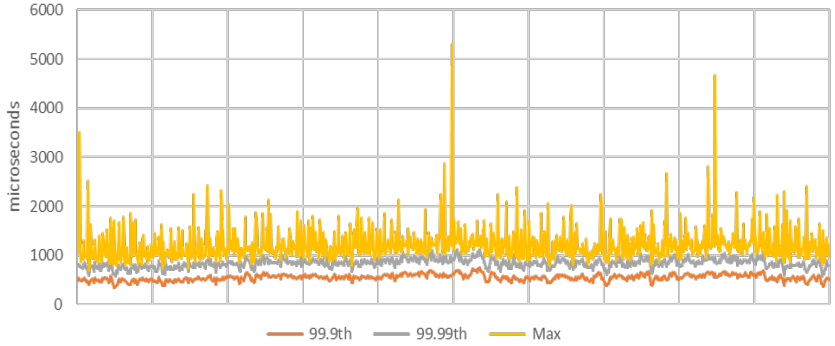

Average	87.40794	Median	80	StdDev	26.18199	Min	1	Max	12585	Sample Co	100000000
99.9:	230	99.99:	279	99.999:	289	99.9999:	2151	99.99999	8484	99.99999	12585

Read

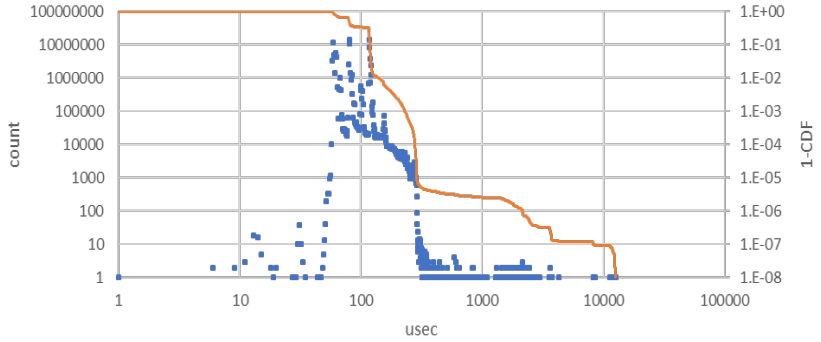

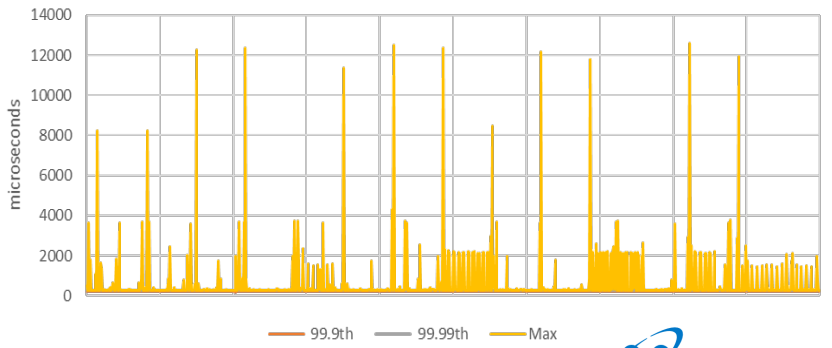

Questions?