Flash Memory Summit

Debugging High Latency Bursts in Data Center Workloads

Presented by:

Rob Dobson, Director Product Marketing

Mike Carlin, Software Architect

TELEDYNE LECROY
Everywhereyoulook™

35 YEARS OF NAND FLASH MEMORY INNOVATION

About the Presenters

Rob Dobson

About

- Director of Product Marketing for Teledyne LeCroy, OakGate SSD Test Solutions, Workload*Intelligence*™ Analytics and Epic Chamber Products
- Prior to Teledyne LeCroy, held management and executive roles with Hewlett-Packard, Intel, Vision Solutions (acquired by Clearlake Capital Group) and Ethentica
- Co-author of three separate U.S. Patents for technologies created in association with "cryptographic authentication for electronic commerce."

Mike Carlin

About

- Lead Software Architect for Teledyne LeCroy's WorkloadIntelligence™ Analytics Software Portfolio
- Oversees the software and systems engineering projects and programs for the Analytics group
- Before Teledyne LeCroy, had stints with Hewlett Packard Enterprise and Clear Capital
- Bachelor of Science in Computer Science for California State University-Sacramento

Agenda

- Background
- SSDs and the Real World
- Challenges
- Analytics and Benefits
- Real World Use Case
  - Problem
  - Configuration and Statistics
  - Investigation
  - Observations and Analysis
  - Recommendations
  - Summary
- SSD Comparison with Real-World Workload
- Questions

Background

- Two years ago, we were approached by several current customers indicating a need for analytics to help them address real-world workloads
- Our Data Center customers (and partners) indicated they needed help understanding how their applications interacted with their storage systems
- Additionally, our Enterprise SSD customers pointed out they had no way of capturing or testing real-world workloads
- In response, the company developed a portfolio of patented applications called WorkloadIntelligence™ DataAgent, Analytics and Replay

SSDs and the Real-World

- Flash-based SSD technology behaves differently than mechanically-based legacy HDDs
- Synthetic workloads never give an accurate reproduction of real-world production workloads
- No method for Data Centers to share and replay production I/O workloads with SSD suppliers to help improve the storage devices in their environment and help shorten their qualification cycle

Challenges in the Data Center

Flash Memory Summit

1. Inability to Capture Production Workload Problems When They Occur
2. Lack of a Deep Analytical Tools to Identify the Issue
3. Flash-based SSDs that Behave Differently than Mechanically-based Legacy HDDs
4. No Method to Share & Replay Production Workloads with SSD Suppliers
5. SSD Technology and Feature Sets that Move Quickly

Workload*Intelligence*™ Software Portfolio

Innovative, **patented** software to capture, analyze and replay real-world production workloads to **help…**

- Data Centers identify, characterize and solve performance problems
- SSD Suppliers improve the performance of Enterprise SSDs in Data Center environments
- Application Developers maximize performance on custom hardware and/or cloud platforms

Workload*Intelligence*™ Analytics

What Workload*Intelligence*™ Enables

Flash Memory Summit

Key Features

Capture and Graph Millions of Individual Data Points	Compare & Contrast Multiple Runs for Different SSDs	Sync & Compare Analyzer Data, Power Data and/or any Time-based Data

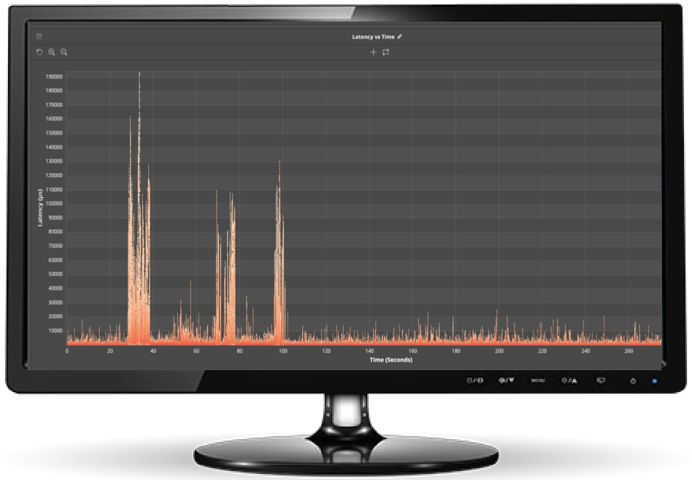

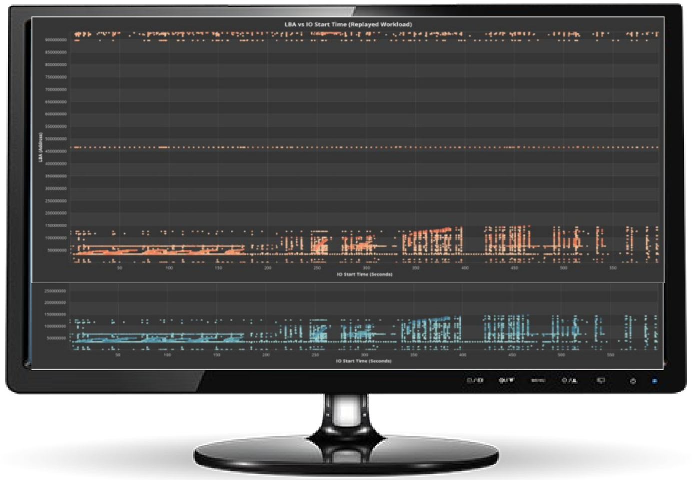

Get deeper insight into past or future events	Understand what variation may exist between SSDs	Discern correlation of different time-based events

Key Benefits

Three Products, One Solution

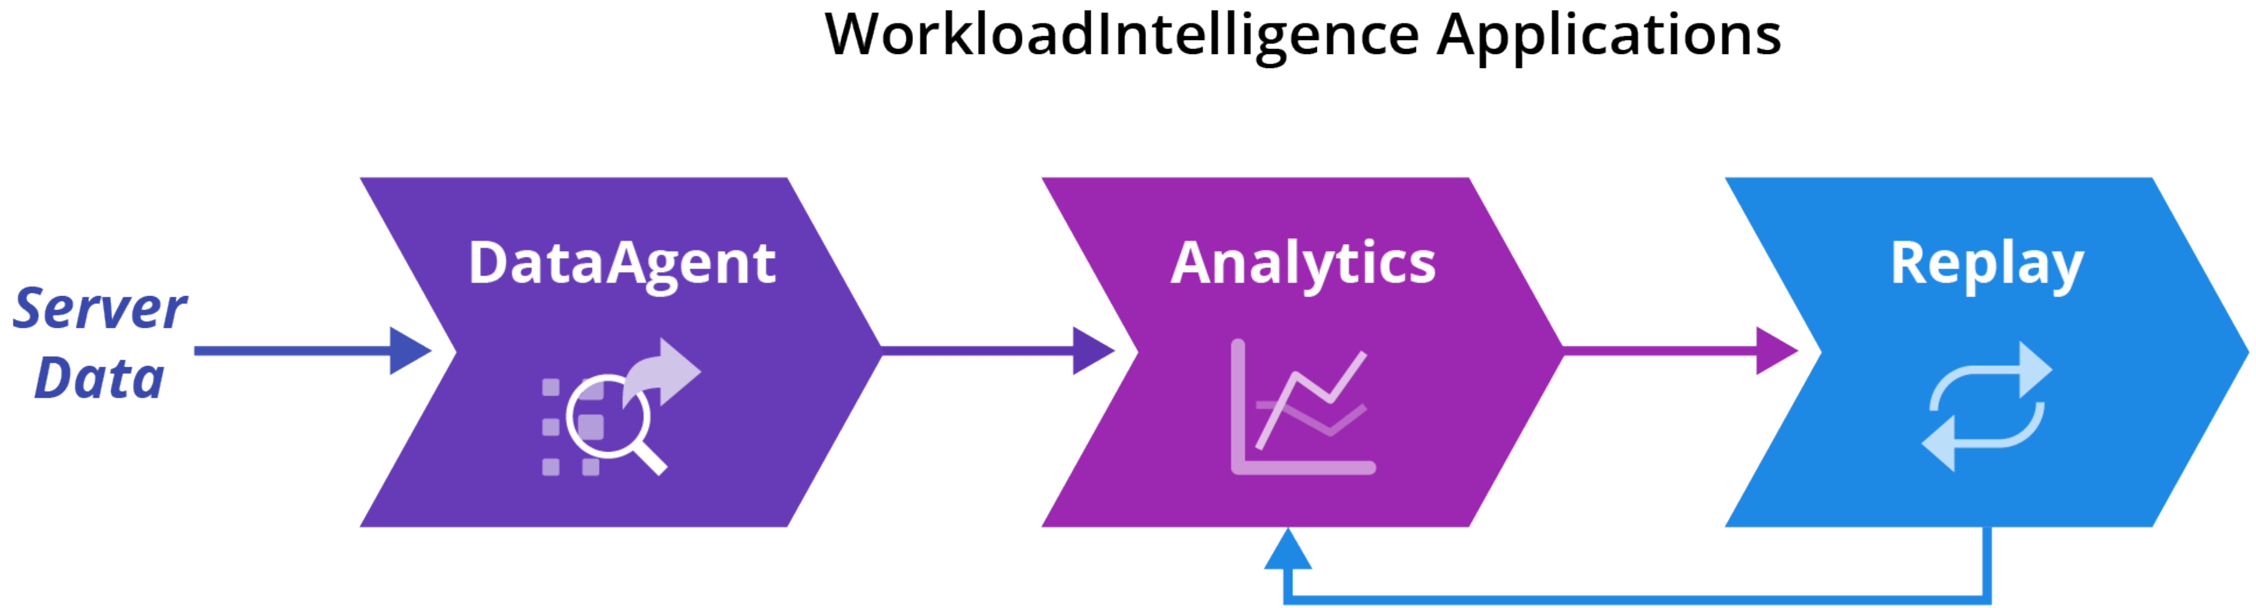

- Collect
- Filter
- Export

- Filter
- Visualize
- Analyze

- Exercise
- Stress Test
- Compare

Workload*Intelligence*^{TM} Workflow

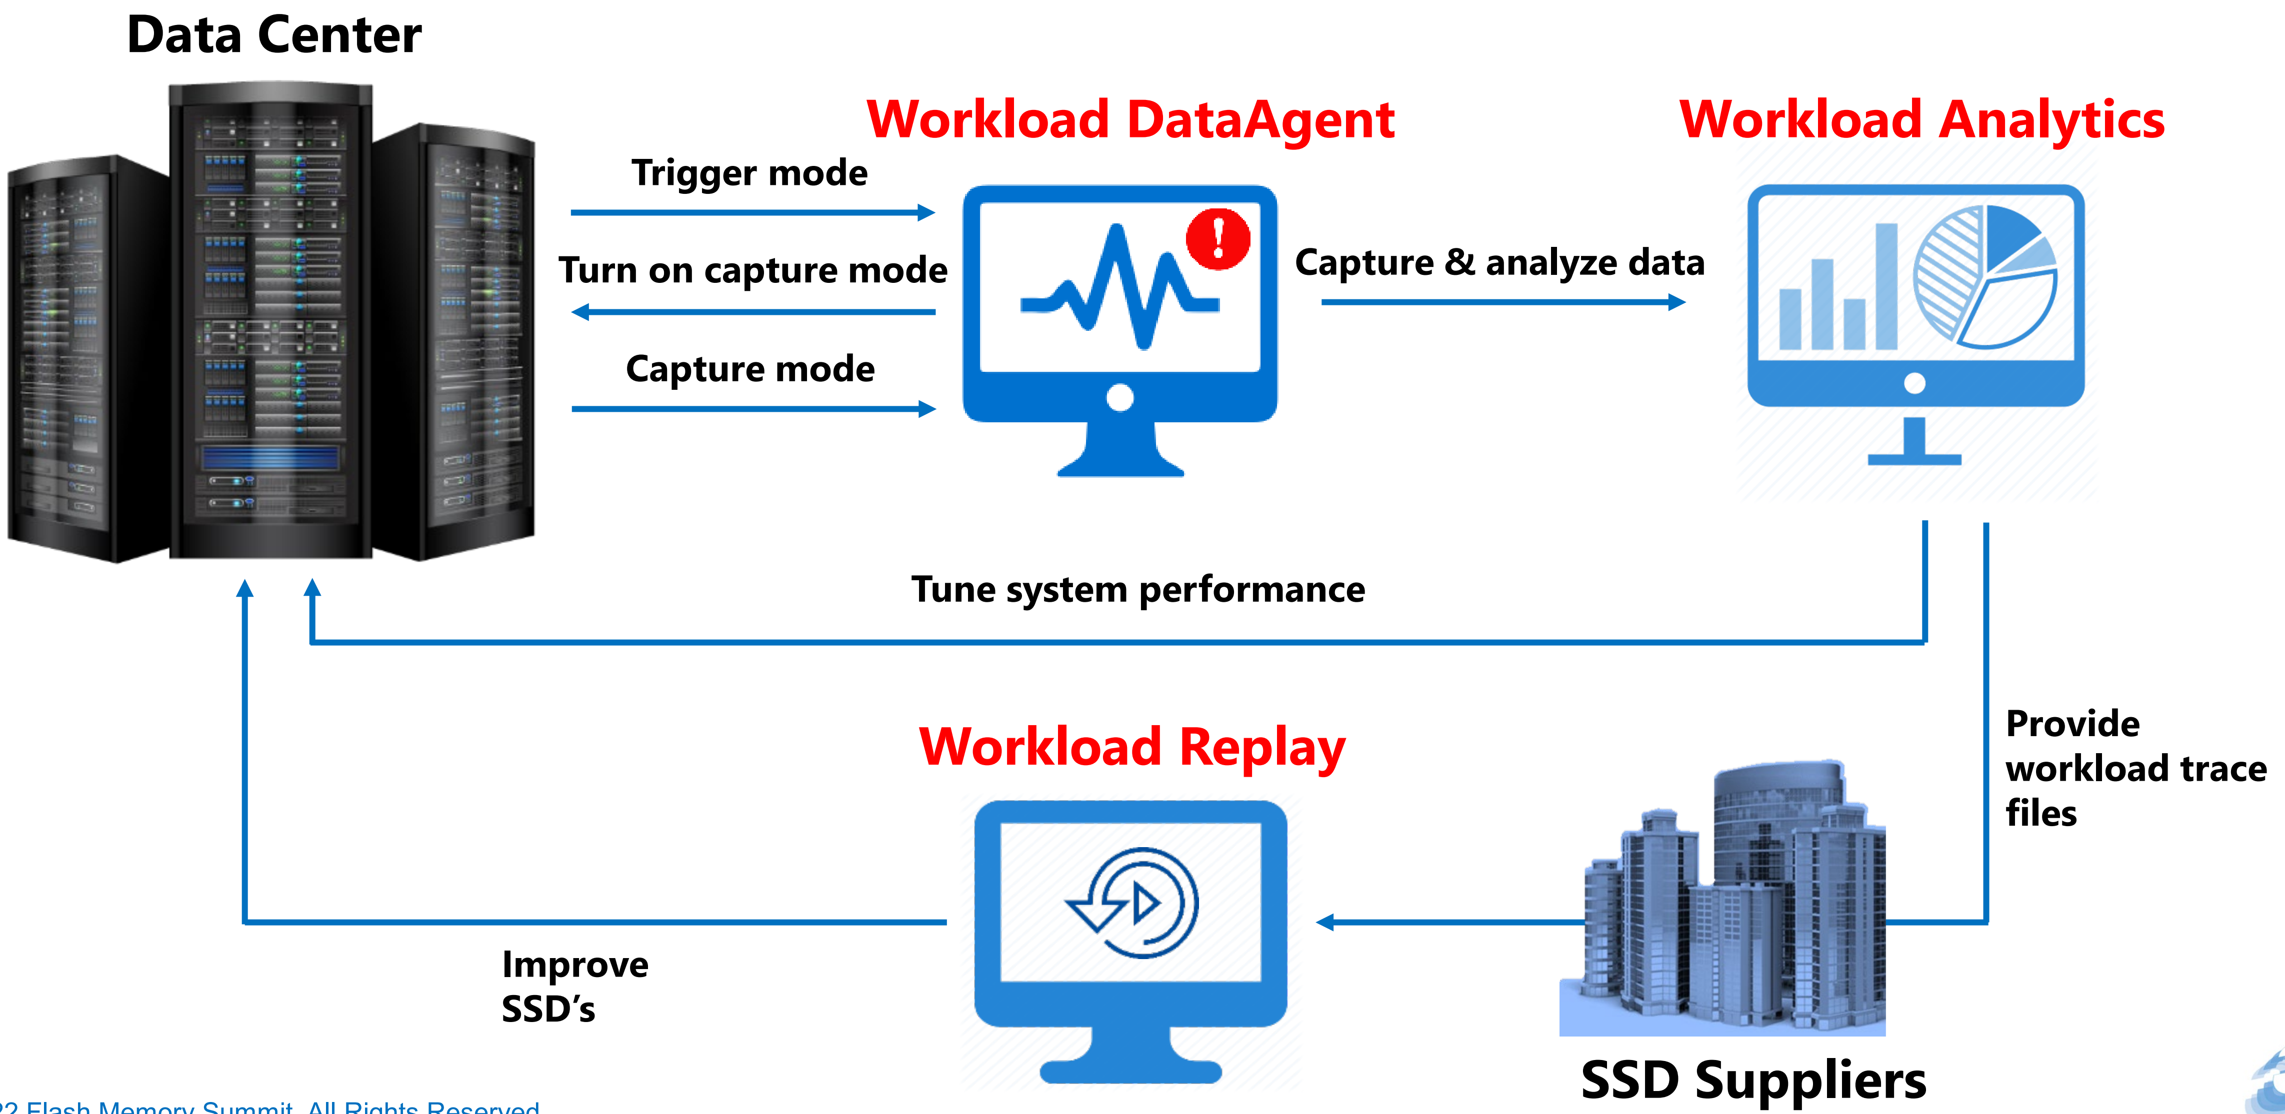

Investigation

The Problem

- Hyperscale Data Center customer was observing unanticipated high-latency bursts with their workloads running on production server(s)
- Customer was unable to characterize and resolve the issue

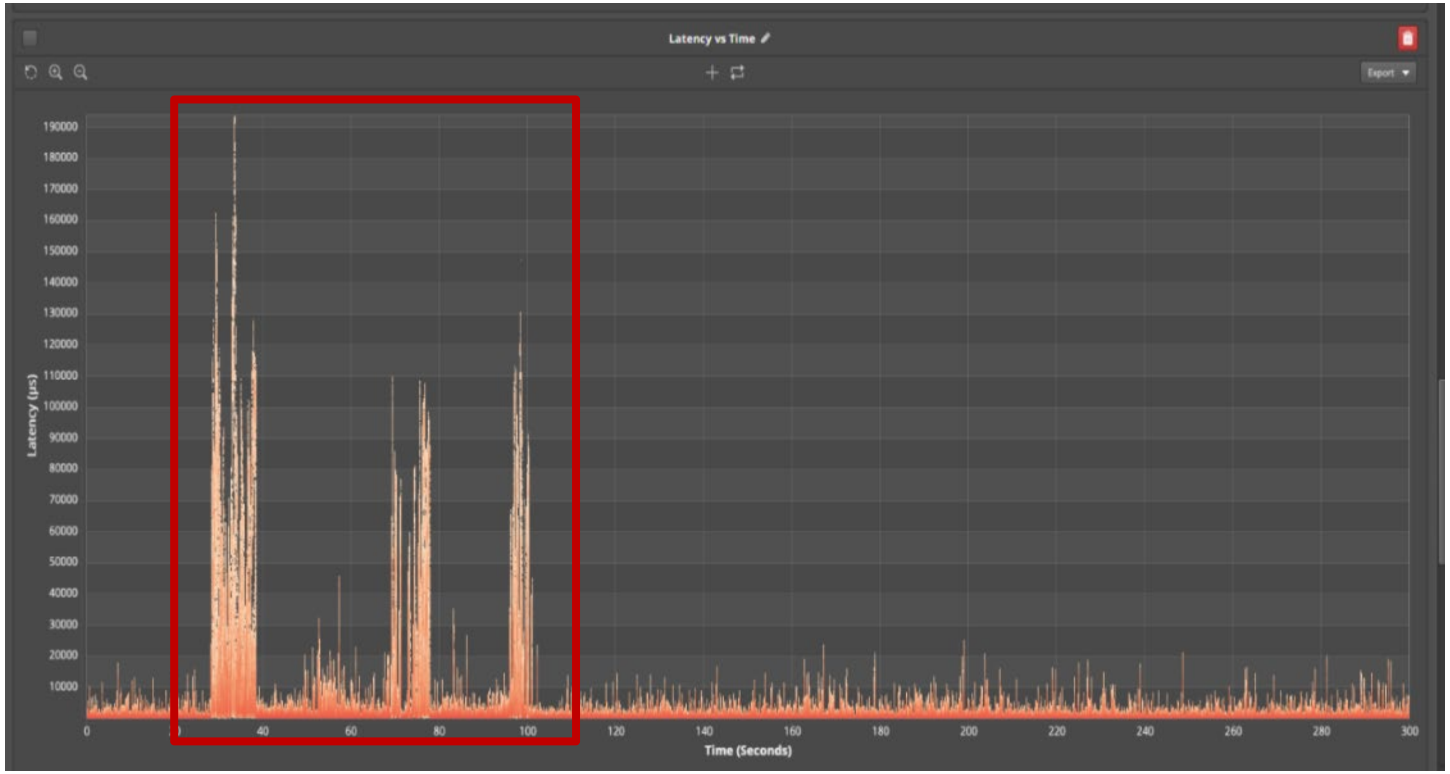

Configuration and Statistics

Customer	Hyperscale Data Center Customer
System	Single 56-core CPU System
Trace Length	5 minute Linux Block Trace
I/O Profile	Read I/O Intensive
Processes	700+ processes
Transactions	157 Million

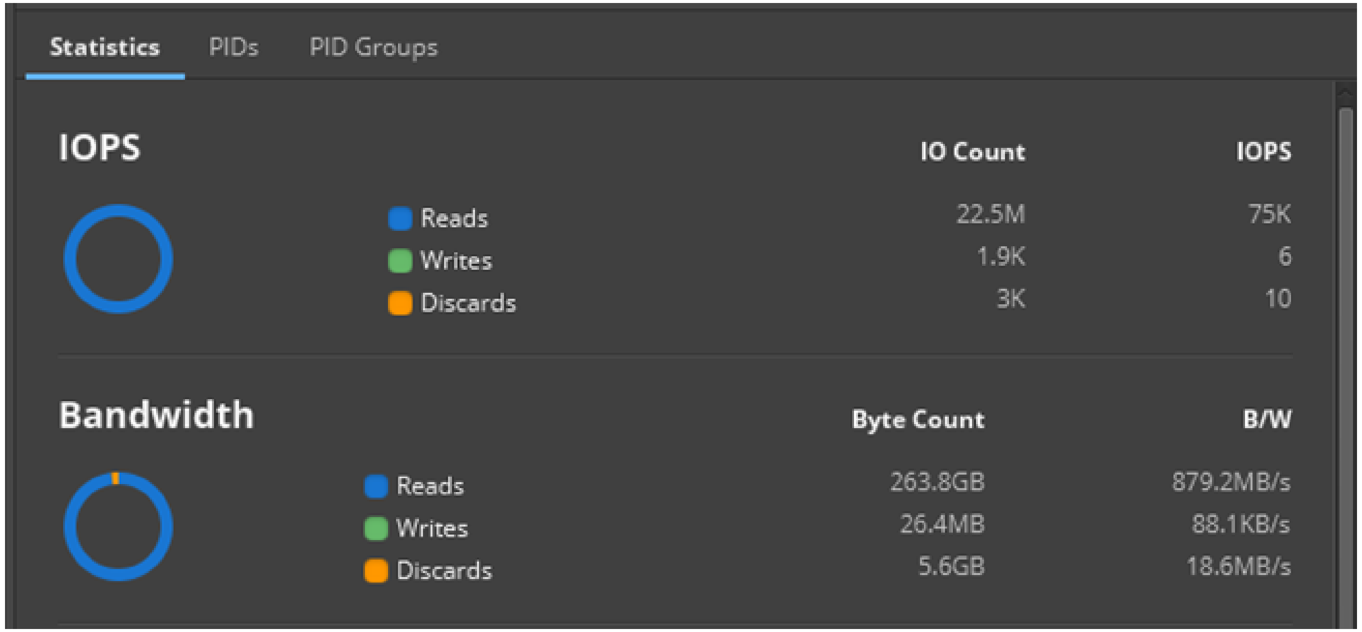

Investigation Scope

1. Imported the hyperscale Data Center workload into the Analytics application
2. Identified and isolated the I/O latency anomalies
3. Created a customized filtered “view” of the workload
4. Zoomed-in on the segment of the filtered workload view
5. Synchronized the zoom time-stamps for all performance graphs
6. Compared and examined the synchronized graphs
7. Correlated the process IDs with key events
8. Shared the results with our customer

Observations and Analysis

Queue Depth and Process Changes

Highlights three separate side-by-side graphs: QDepth vs Time, PID Count QDepth vs Time and the original Latency vs Time, all synchronized to the latency burst timestamp

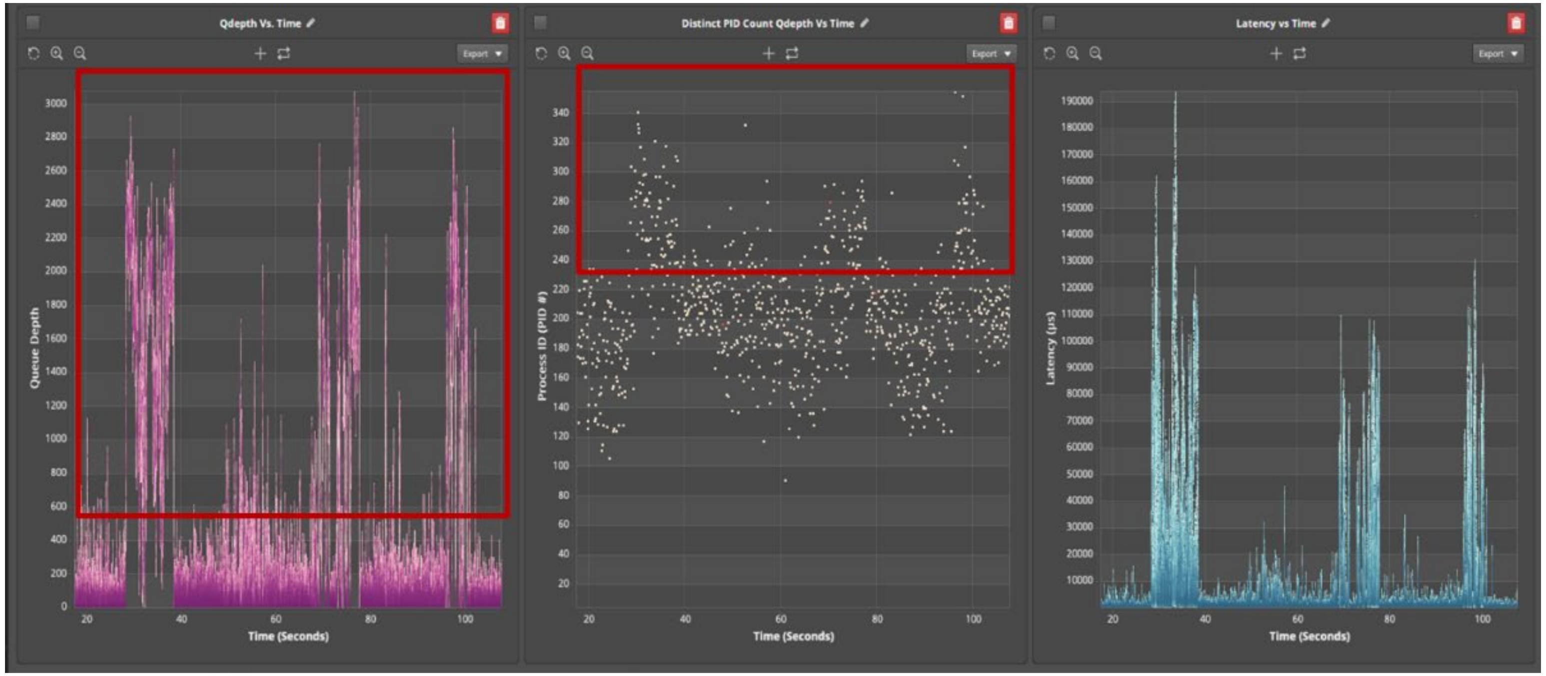

Key observations include a significant increase in queue depth and a 4x increase in the processes triggered at the time of the latency event

Large Trim Operation

Highlights Read Latency vs Time

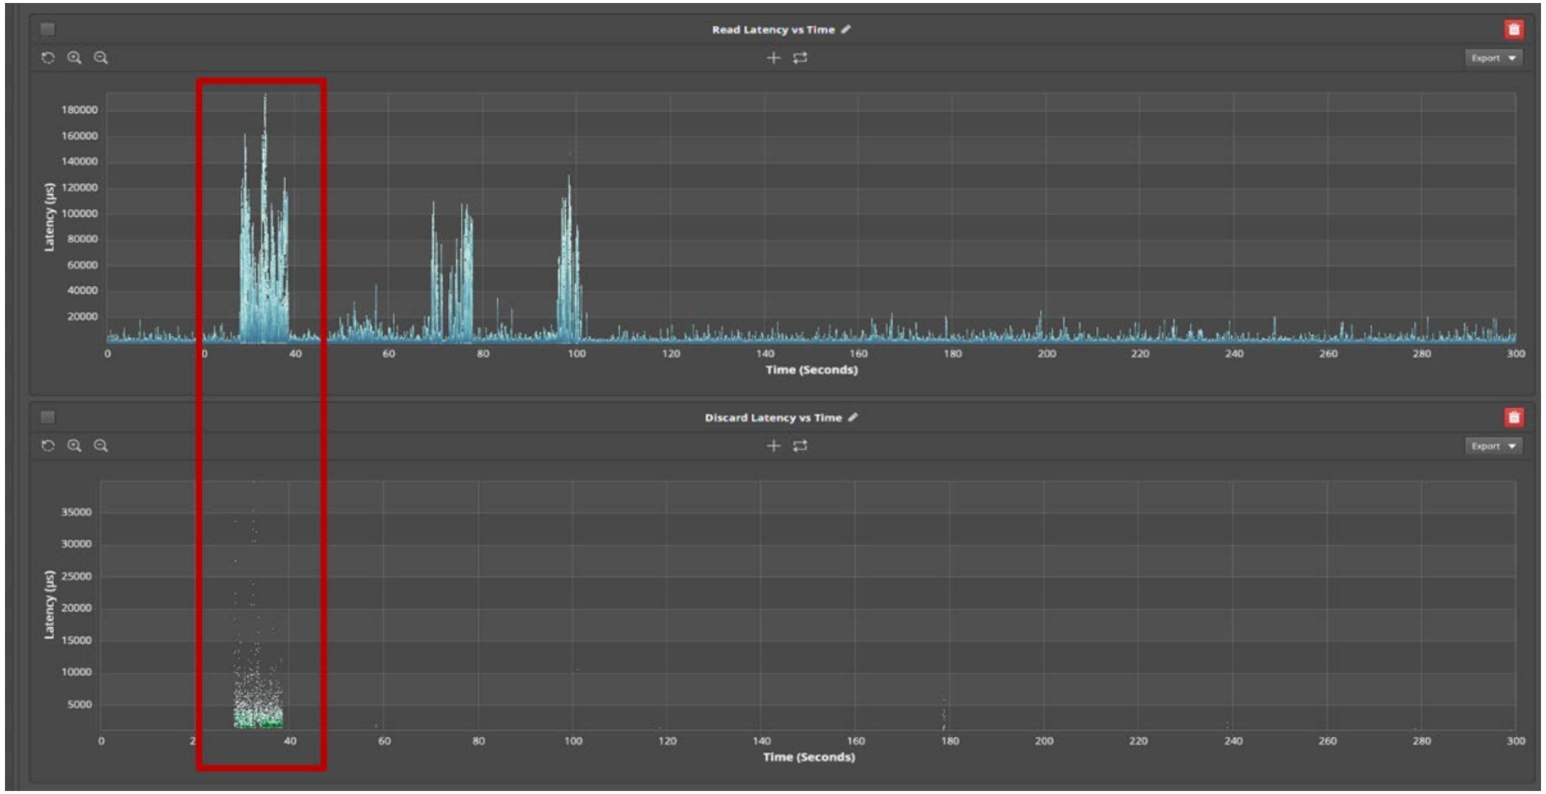

The key observation was the large set of trim (discard) operations occurring at the time of the latency event.

Decrease in IOPs

Highlights IOPs vs Time

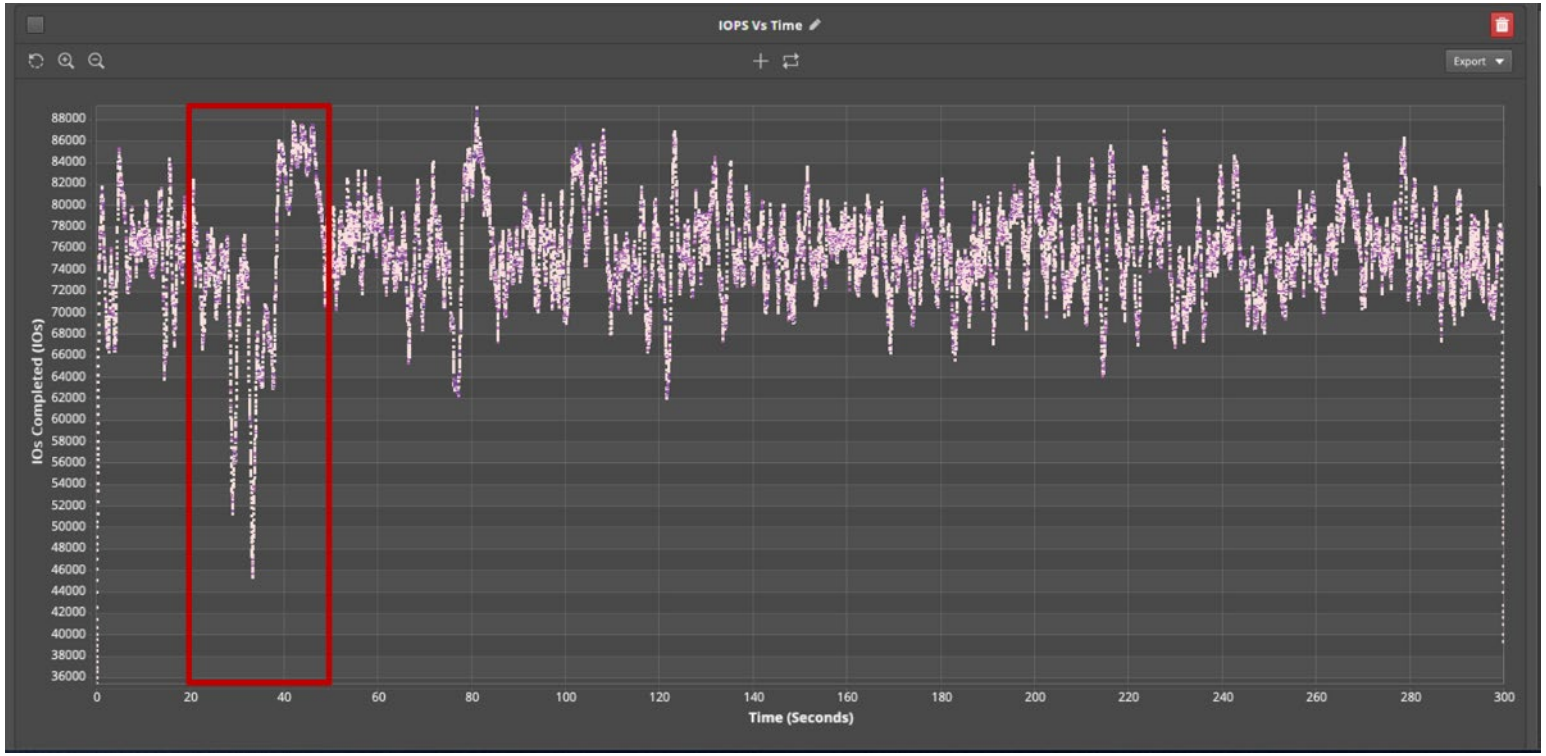

The key observation is a large decrease in I/O operations at the time of the latency event.

Inefficient Use of CPU Cores

Highlights Latency vs CPU

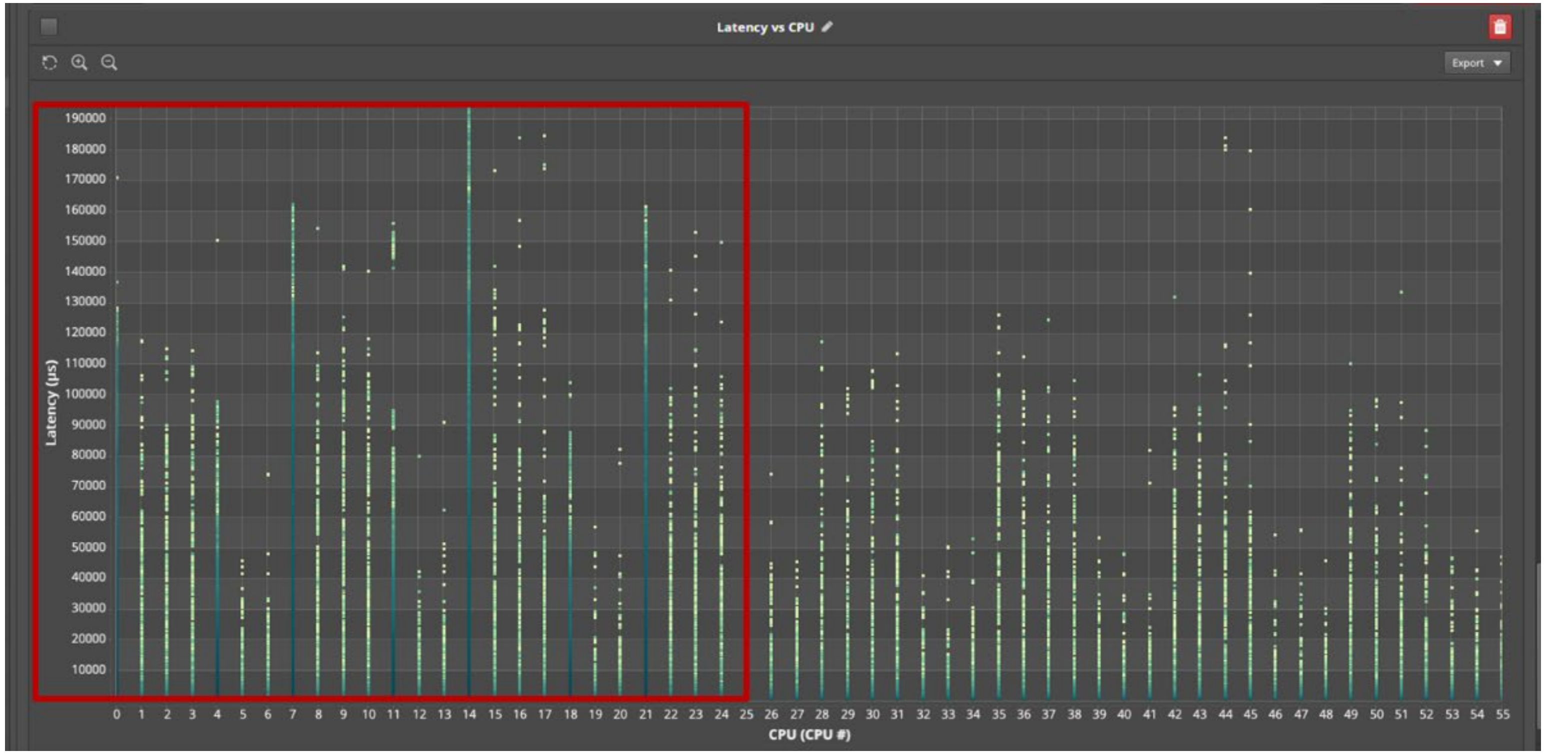

Key observations are an inefficient use of the 56 cores available to the application. That is, approximately eight cores were fully loaded during the latency burst and only a single core serviced the trim in the same time period

Inefficient Use of CPU Cores

Highlights I/O Size vs CPU

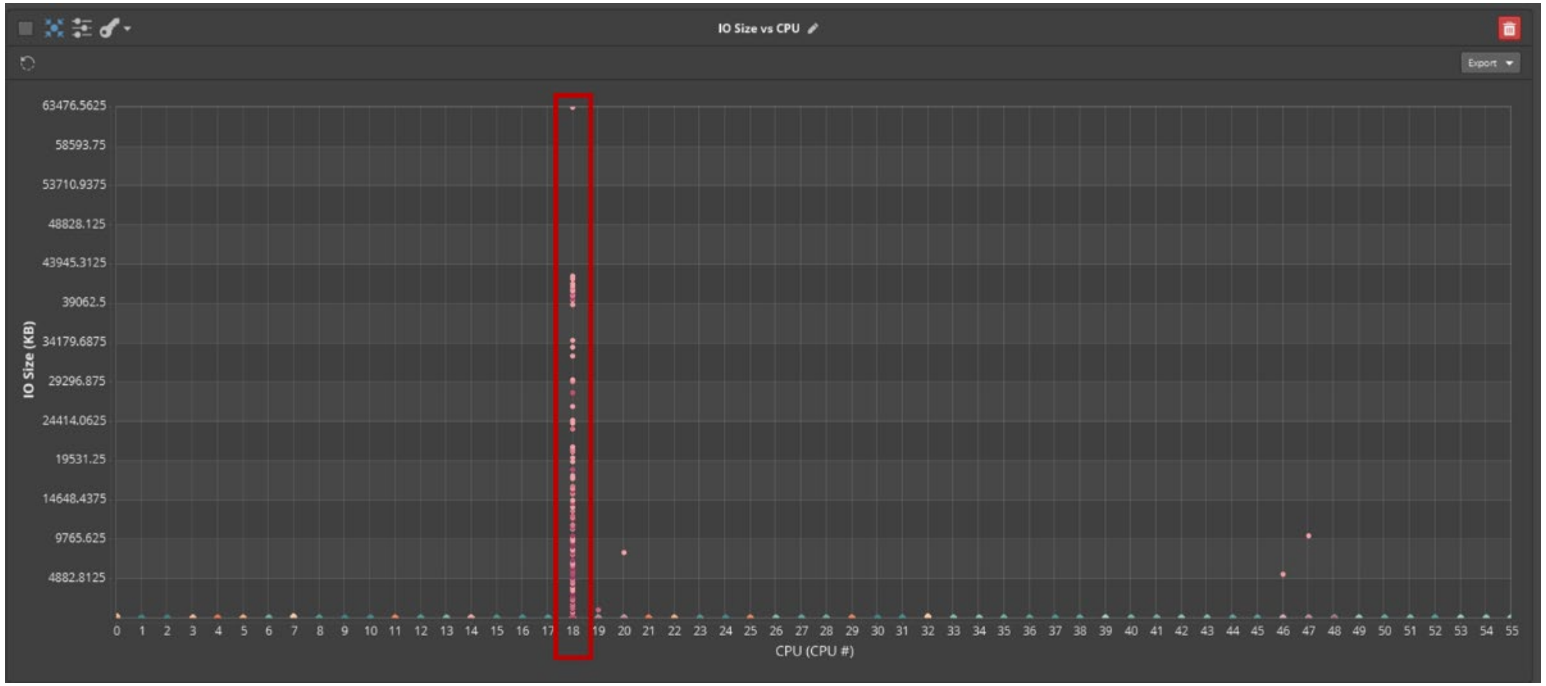

The key observation is that only a single core serviced the trim (discard) in the same time period.

Rogue Process

Flash Memory Summit

Highlights Process ID vs Time

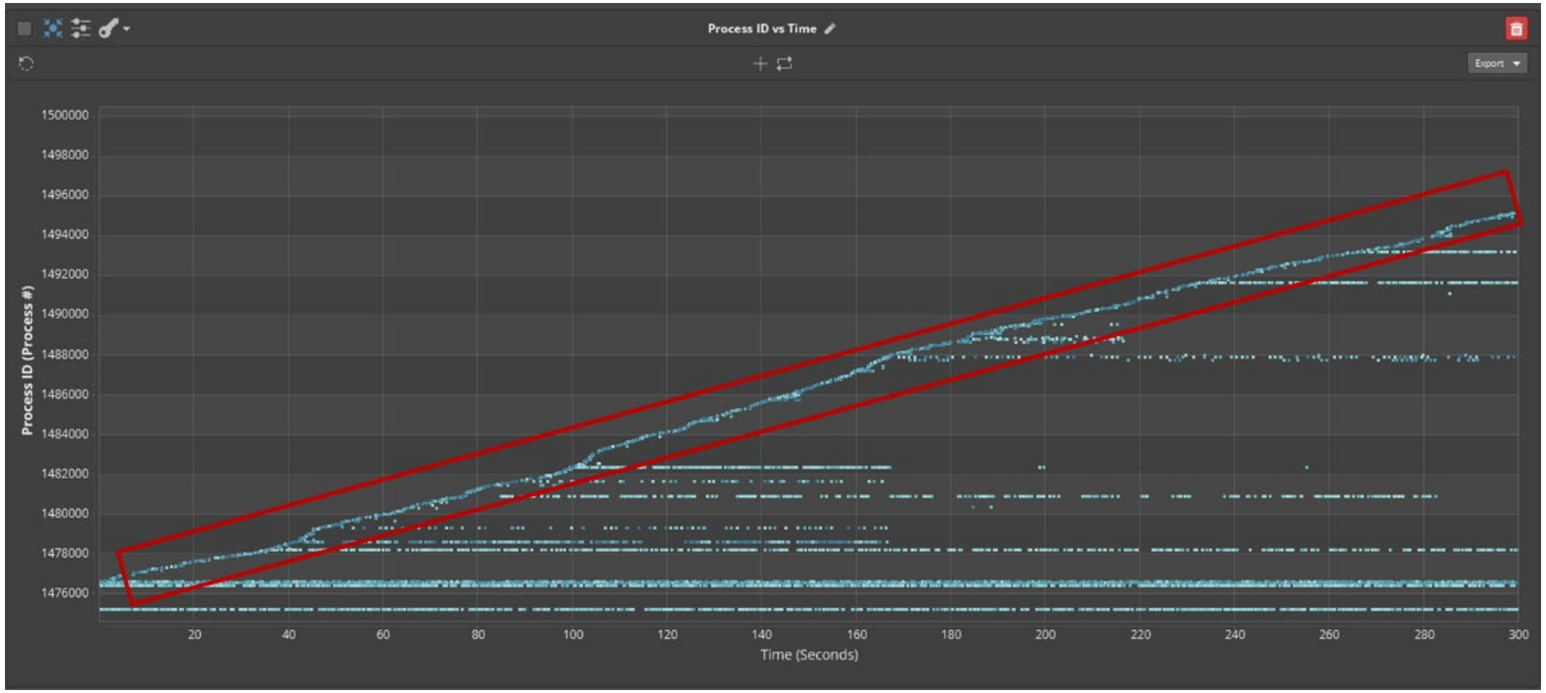

This additional observation was not specifically related to the latency bursts but highlights a “rogue” or unoptimized process being continually respawned.

YEARS OF NAND FLASH MEMORY INNOVATION

Key Recommendations to Customer

1. Review the Linux OS Kworker process settings that initiated the multi-second Trim
2. Examine the additional processes that were kicked off by host at the time of the latency event
3. Investigate the rogue or unoptimized process to determine whether it’s actually a rogue process or a process that could be doing more before being recycled
4. Review the CPU core allocation with the application team to understand whether the cores could have been allocated more efficiently
5. Tune the Application and Linux block layer to achieve better workload performance

Further Analysis – Workload Comparison on other SSDs

Workload Comparison on Other SSDs

- For this analysis, we utilize three (3) latency metrics to characterize how SSDs handle the same workload:

  1. **Baseline latency:** the average latency when the drive is operating normally and not experiencing high bursts of latency
  2. **Peak latency:** the max latency exhibited by the drive during the workload execution
  3. **Peak latency recovery:** the amount of time it takes for the drive to return to its baseline latency if the drive experiences a burst of high latency. If the drive never experienced a burst of high latency, this metric is not applicable

Solid State Drives Used During Analysis

- Table 1 highlights the **Consumer SSD** capacities and capabilities

SSD	Protocol	Capacity (GB)	Seq. Read MB/s	Seq. Write MB/s	Rand. Read IOPS	Rand Write IOPS	Form Factor
Drive A	NVMe	480	2,500	2,000	550,000	550,000	AIC
Drive B	NVMe	512	3,000	1,600	240,000	240,000	M.2
Drive C	NVMe	512	3,500	2,300	370,000	550,000	M.2
Drive D	NVMe	500	3,400	2,500	410,000	330,000	M.2
Drive E	NVMe	500	5,000	2,500	430,000	630,000	M.2
Drive F	NVMe	500	3,430	2,600	NA	NA	M.2

Table 1

- Table 2 highlights the **Enterprise SSD** capacities and capabilities

SSD	Protocol	Capacity (GB)	Seq. Read MB/s	Seq. Write MB/s	Rand. Read IOPS	Rand Write IOPS	Form Factor
Drive G	NVMe	1,000	2,800	1,100	465,000	70,000	U.2
Drive H	NVMe	960	3,000	1,200	400,000	38,000	M.2
Drive I	NVMe	960	2,000	1,200	245,000	67,000	M.2

Table 2

Replay Technology

- **Workload*Intelligence*™ Replay software** ensures the integrity of the Workload I/Os during replay on dissimilar drives
- The software provides **three different modes** to faithfully reproduce and to ensure integrity of the workload given it might have been produced on a different SSD drive:

1. **Skip IO:** which tells the replay application to skip I/Os that have LBAs that exceed the maximum LBA
2. **Wrap LBA:** which tells the replay application to wrap the exceeded LBA values starting from 0 again to complete the total number of LBAs in the original trace file
3. **Scale LBA:** which tells the replay application to scale the LBAs to fit the LBA range

Replay Reproduction

IO Size vs Time (Original)

IO Size vs Time (Replay)

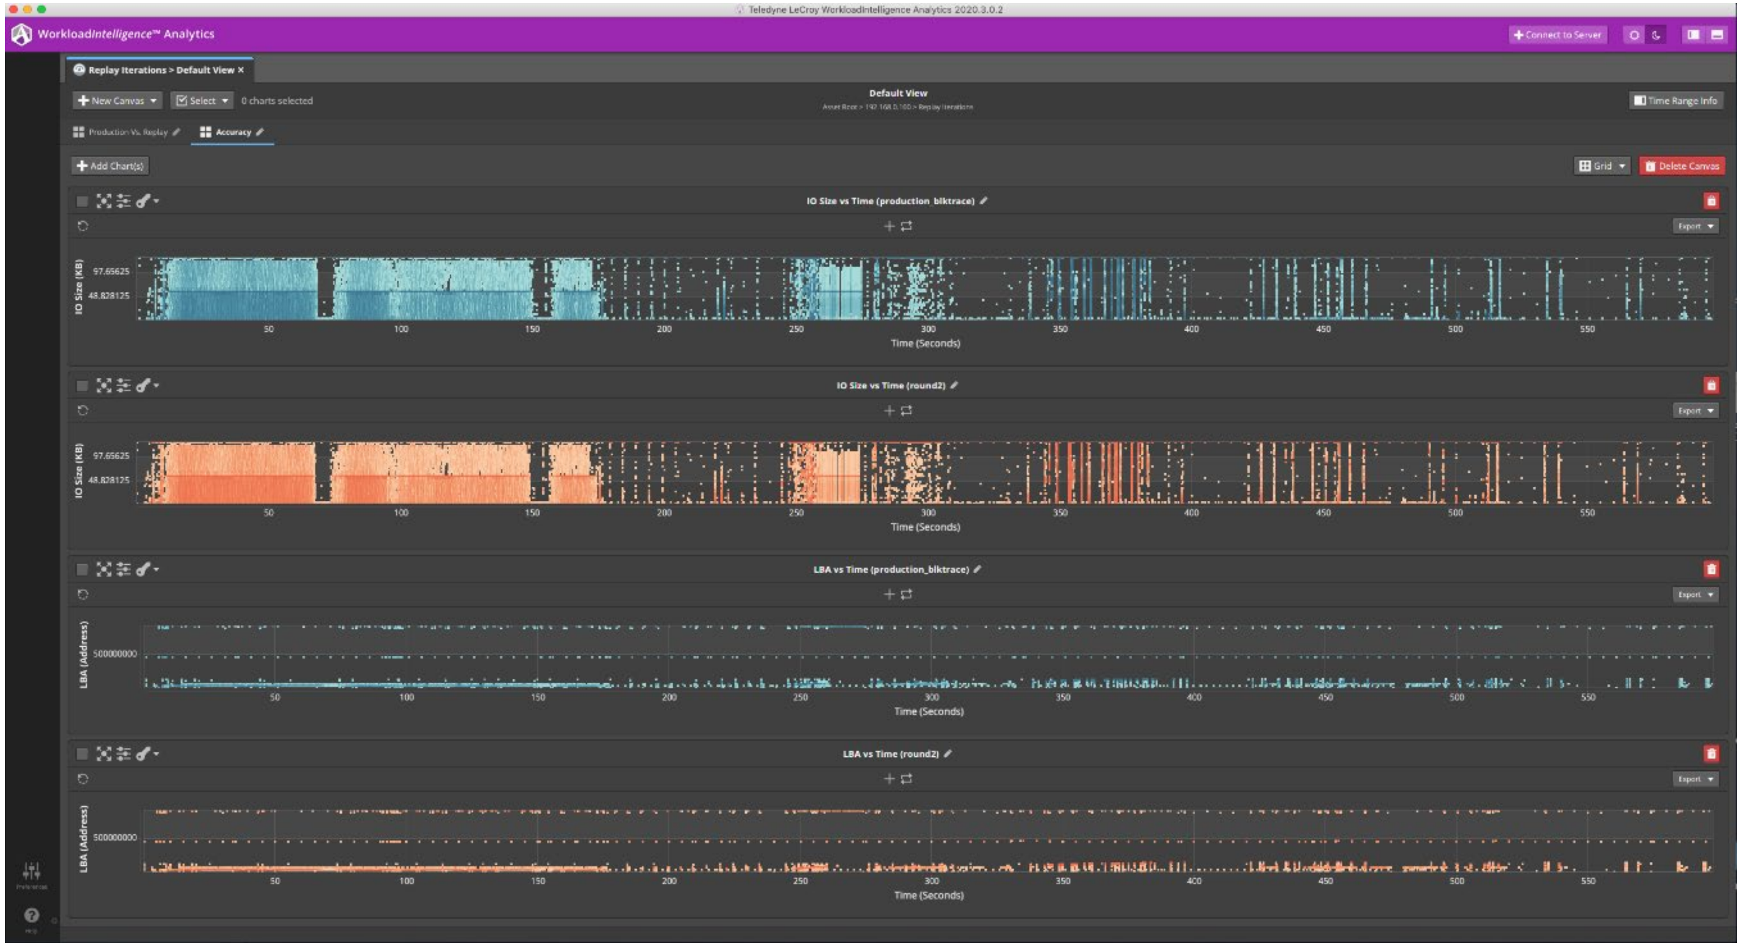

LBA vs Time (Original)

LBA vs Time (Replay)

Results – A Few Examples

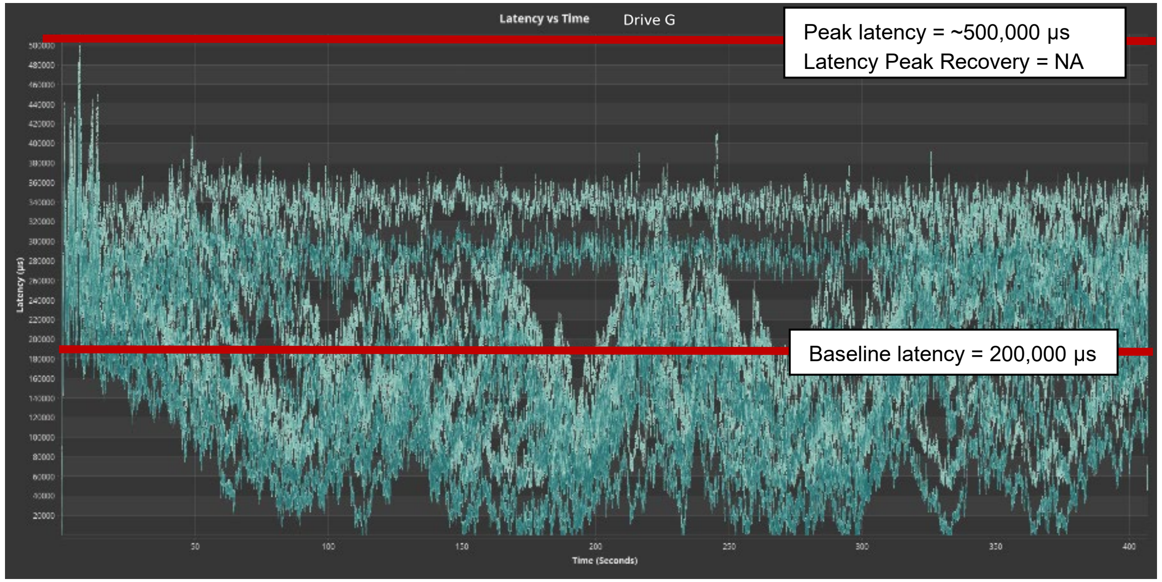

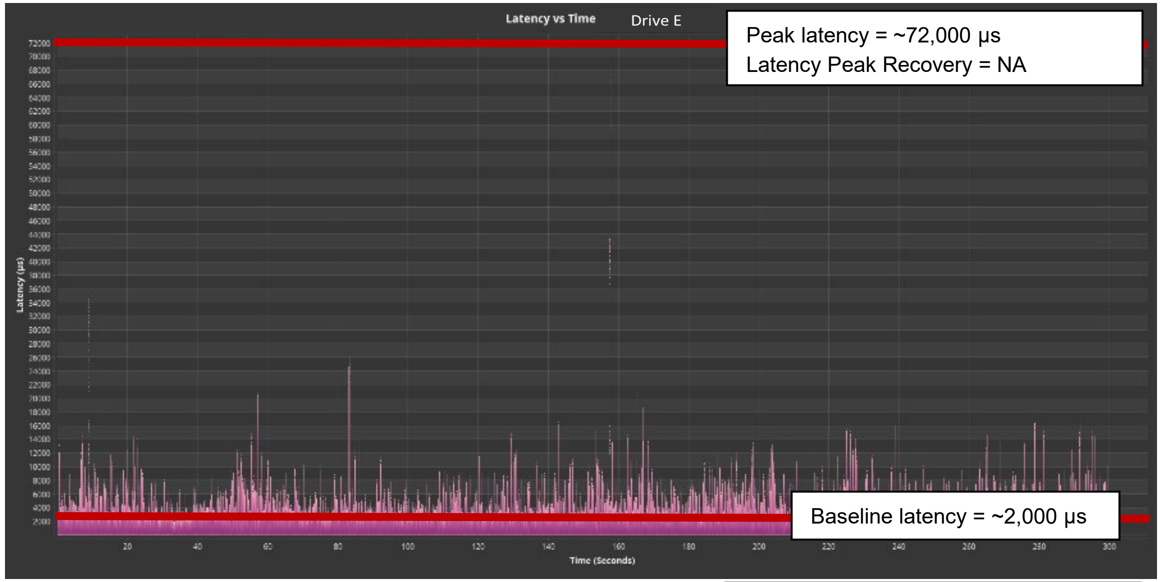

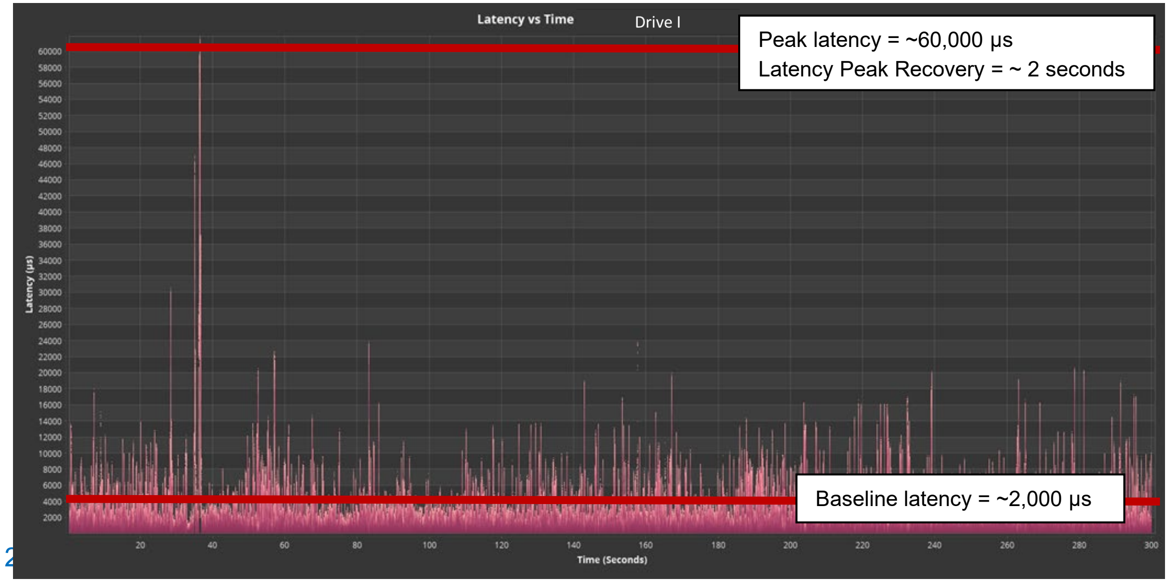

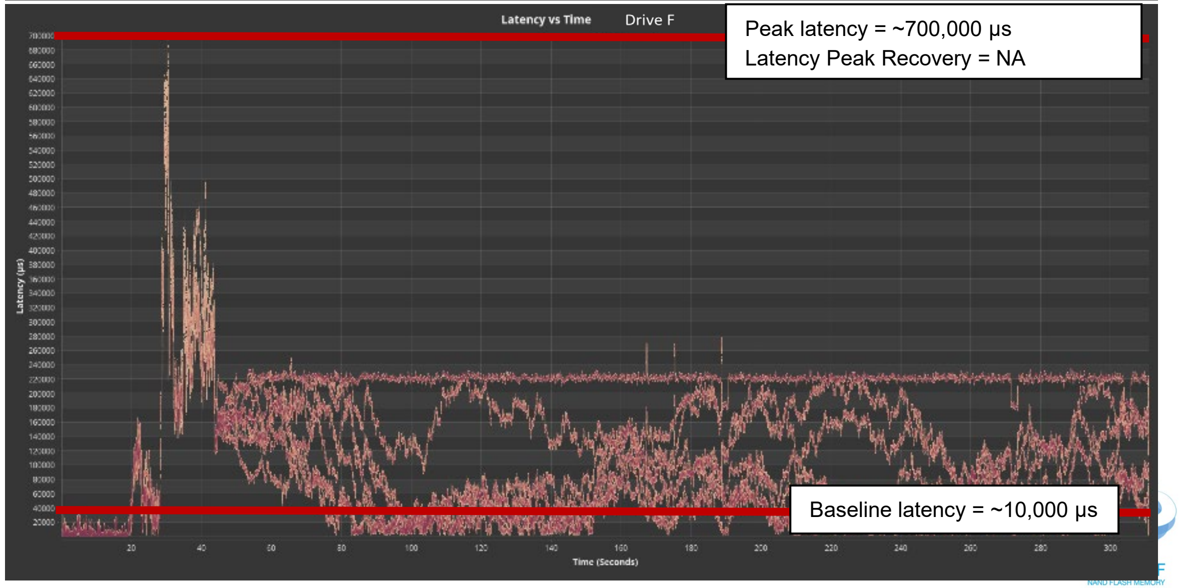

Complete Results

Consumer SSD results (Table 1)

- Results were quite mixed
- Drive A and Drive D performed admirably, much better than the original drive, and would be better candidates to utilize for this data center workload
- All of the drives, except for Drive B and Drive F, were able to recover quickly from the peak latency occurrence

Enterprise SSD results (Table 2)

- Results were also mixed
- Drive H and I performed well, again much better than the original drive, and would be better candidates to utilize for this data center workload
- Both Drive H and Drive I were able to recover quickly from the peak latency occurrence
- Drive G had unanticipated results and more analysis is needed to understand why it did not perform well with this workload.

SSD	Protocol	Baseline Latency (µs)	Peak Latency (µs)	Latency Peak Recovery (sec)
Original	NVMe	3,000	190,000	10 seconds
Drive A	NVMe	400	8,000	NA
Drive B	NVMe	4,000	320,000	15 seconds
Drive C	NVMe	2,500	72,000	NA
Drive D	NVMe	1,200	13,000	NA
Drive E	NVMe	2,000	72,000	NA
Drive F	NVMe	NA	700,000	NA

Table 1

SSD	Protocol	Baseline Latency (µs)	Peak Latency (µs)	Latency Peak Recovery (sec)
Original	NVMe	3,000	190,000	10 seconds
Drive G	NVMe	200,000	500,000	NA
Drive G (4K)	NVMe	40,000	400,000	20 seconds
Drive H	NVMe	900	10,000	NA
Drive I	NVMe	2,000	60,000	2 seconds

Table 2

Summary

- **WorkloadIntelligence has been used successfully to:**
  - Find system and application latencies affecting performance
  - Detect database latencies affecting overall application performance
  - Characterize current and future workload performance
  - Identify rogue or unoptimized processes
  - Ensure applications were utilizing and allocating CPU cores efficiently
  - Tune CEPH storage solutions especially when utilizing large, clustered configurations
  - Tune the Application and Linux/Windows block layer to achieve better workload performance

Visit us @ Booth 919

For a Demonstration of WorkloadIntelligence™ in Action!