Flash Memory Summit

OakGate TECHNOLOGY

Optimizing SSD Performance Efficiently with Realistic Workloads

Presented by:

Bob Weisickle, CTO and Co-Founder

OakGate Technology

Santa Clara, CA
August 2019

Data Center Industry Challenges

Flash Memory Summit

OakGate TECHNOLOGY

Lack of adequate analytic tools to identify and assess the challenges with their production data center workloads

Capture

Need to effectively capture and analyze performance data while maintaining privacy and security

Analyze

Once the data has been captured, how does an organization get value from this data? How do they find the signal in the "noise?"

Visualize

Optimize

Once the signal has been detected, how does an organization take action to optimize performance?

Share

How to share results with SSD partners to help them pro-actively improve the product?

The Case for Workload Analytics

Hard Disk Drives

Solid State Drives

SSD provides

- Improved application performance
- Greater storage utilization
- Reduced operating expenses
- Lowered total cost of ownership

Identifying Bottleneck?

However…

- Replacing slow spinning disks with faster flash storage doesn't remove the bottleneck. **It just moves it somewhere else**
- Additionally, it may **shift over time** as processors and memory advance and applications and firmware are updated
- And the introduction of NVMe will **likely move it again**

Shared Value Proposition

SSD and Storage System OEMs get:

- Production workloads that can be replayed to test and validate their drives
- Real-world performance data that can help optimize the drives/systems

Datacenters and Cloud SPs get:

- Drives validated and tested with production workloads
- Real-world performance data to make intelligent business decisions on purchases and configurations

Both get a deeper insight of storage system behavior and performance

Workload Characteristics

- **I/O Mix:** is the workload read heavy, write heavy or balanced?
- **I/O Type:** does the workload write or read data sequentially or randomly?
- **Data/Metadata Mix:** does the workload read or manipulate metadata more so than actual data?
- **Block or File Size Distribution:** does the workload write in small or large blocks?
- **Host Resource Usage:** CPU, Process IDs, Memory
- **Data Efficiency:** does the workload have highly redundant or compressible data so that functions like deduplication and compression work effectively?
- Is the workload prone to **specific hot spots?**
- How do all of the above characteristics **change over the relevant time period?**

Flash Memory Summit

Workload Analytics Workflow

OakGate TECHNOLOGY

1. Import data center production workload trace(s) →
2. Create customized filtered "view" of the workload →
3. Zoom-in on segment of the filtered workload view →
4. Synchronize zoom time-stamps for all performance graphs →
5. Compare and examine synchronized graphs →
6. Run deep analytics for additional detailed graphs →
7. Correlate process IDs with key events →
8. Share results with data center teams and SSD suppliers

Use Case

High Latency Bursts within Hyperscale Data Center Workload

Background

- Hyperscale Data Center Customer
- Single 56-core CPU system (28 cores with Hyper-Threading enabled)
- 5 minute workload (Linux “Block Trace” – 5Gb file)
- 157 million transactions
- 700+ individual process IDs
- Read-intensive application workload

Latency vs Time

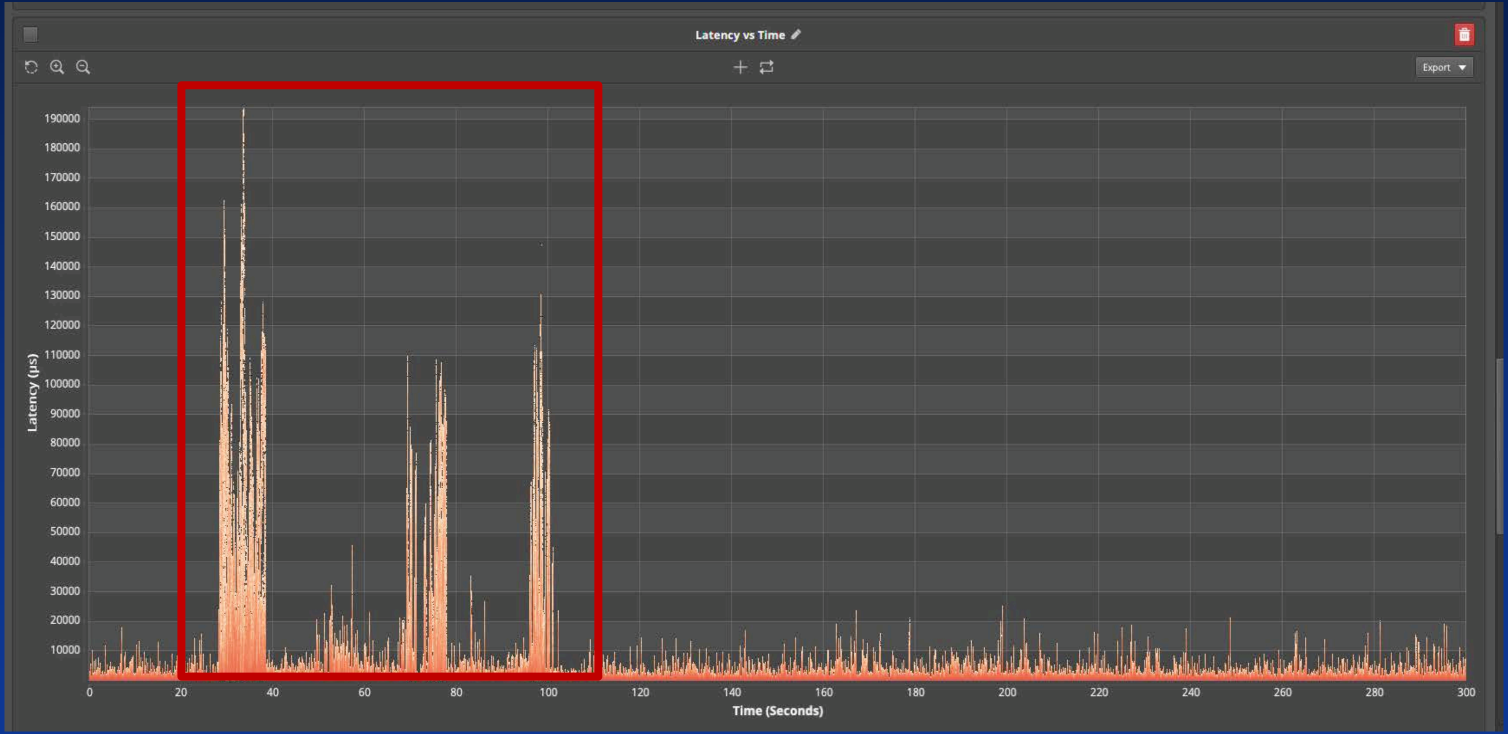

Latency vs Time Graph

QDepth and Process ID Count

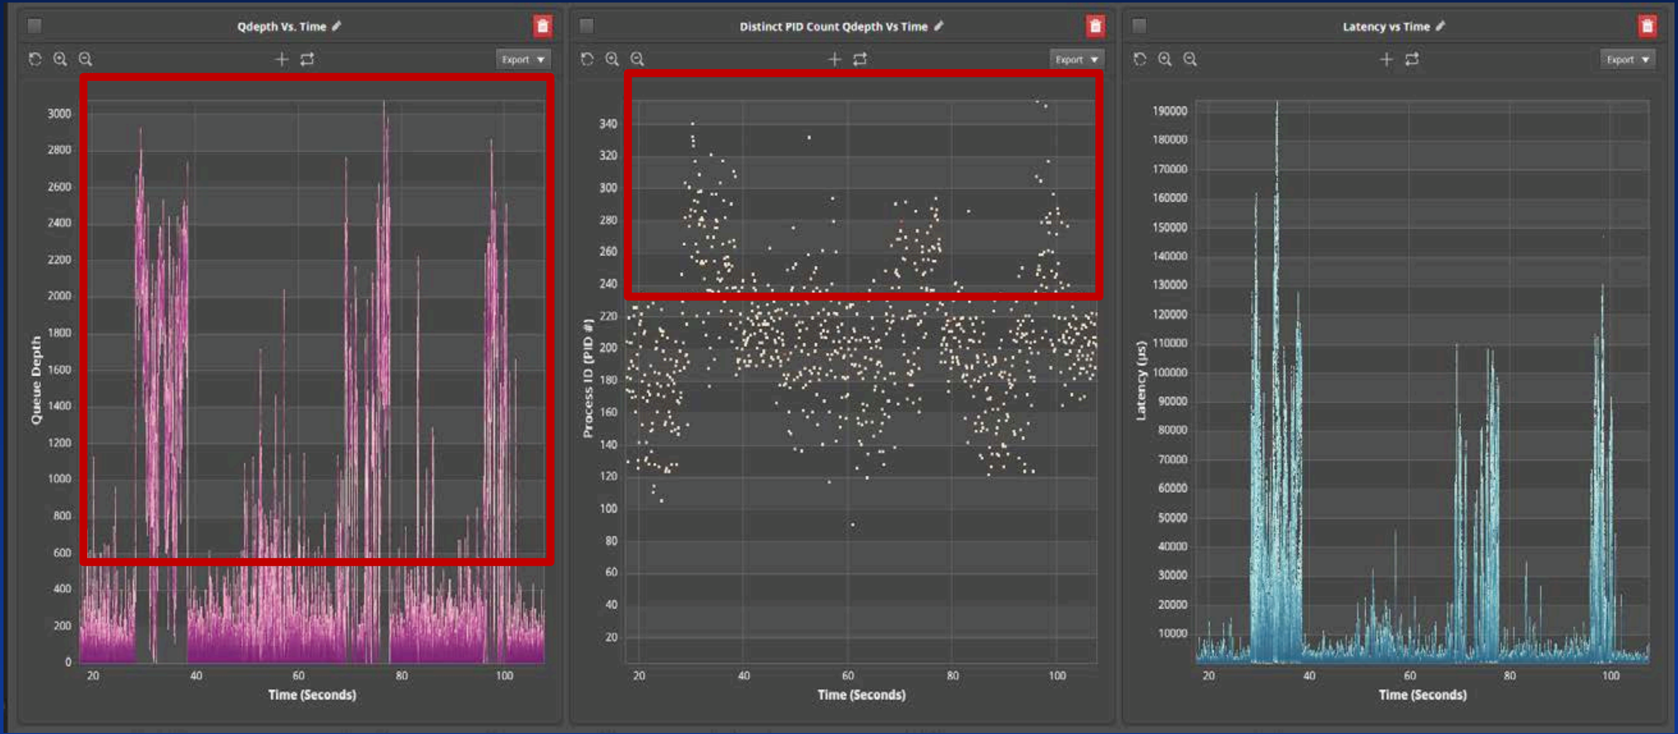

Queue Depth and Process ID Count correlated with latency time stamp

Discards vs Time

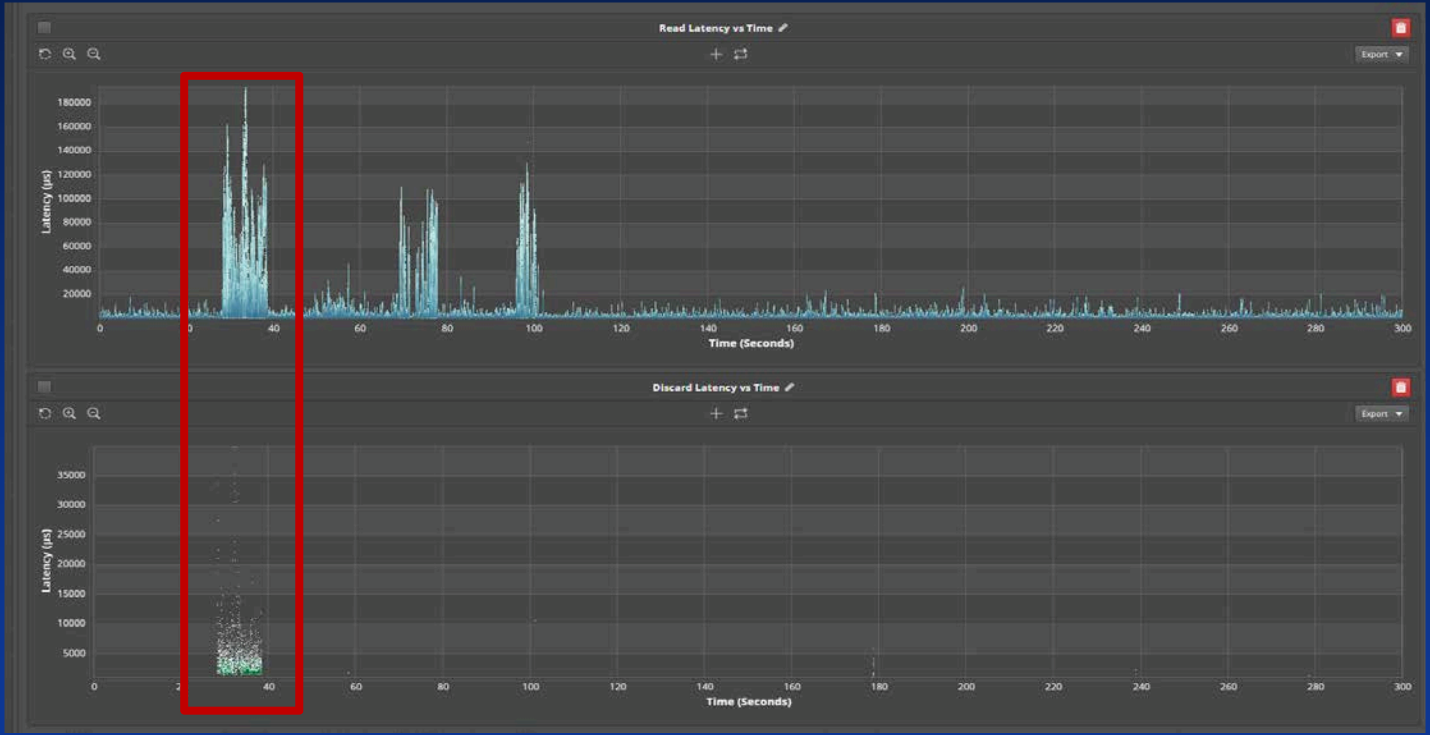

Discards vs Time correlated with Latency

IOPs vs Time

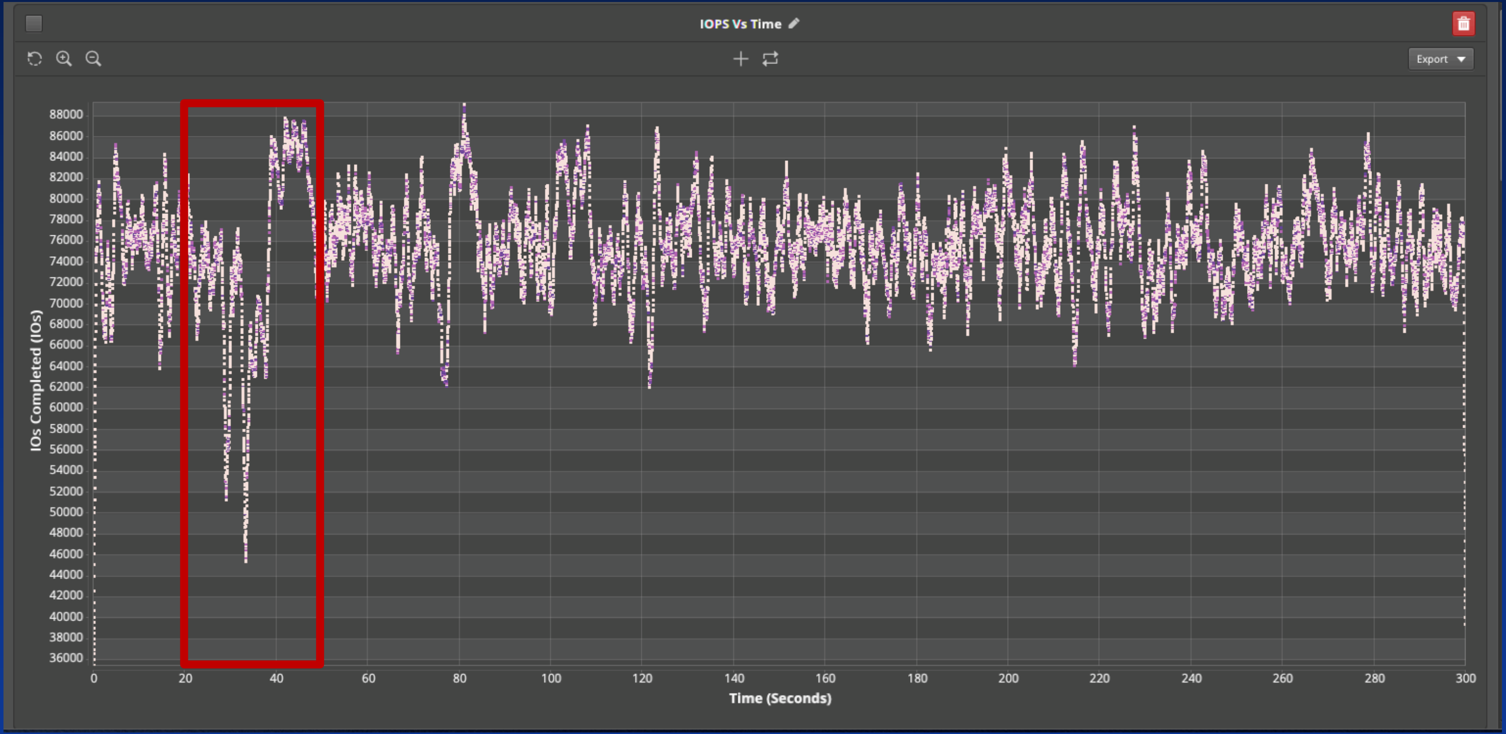

IOPs vs Time Graph

CPU vs Latency

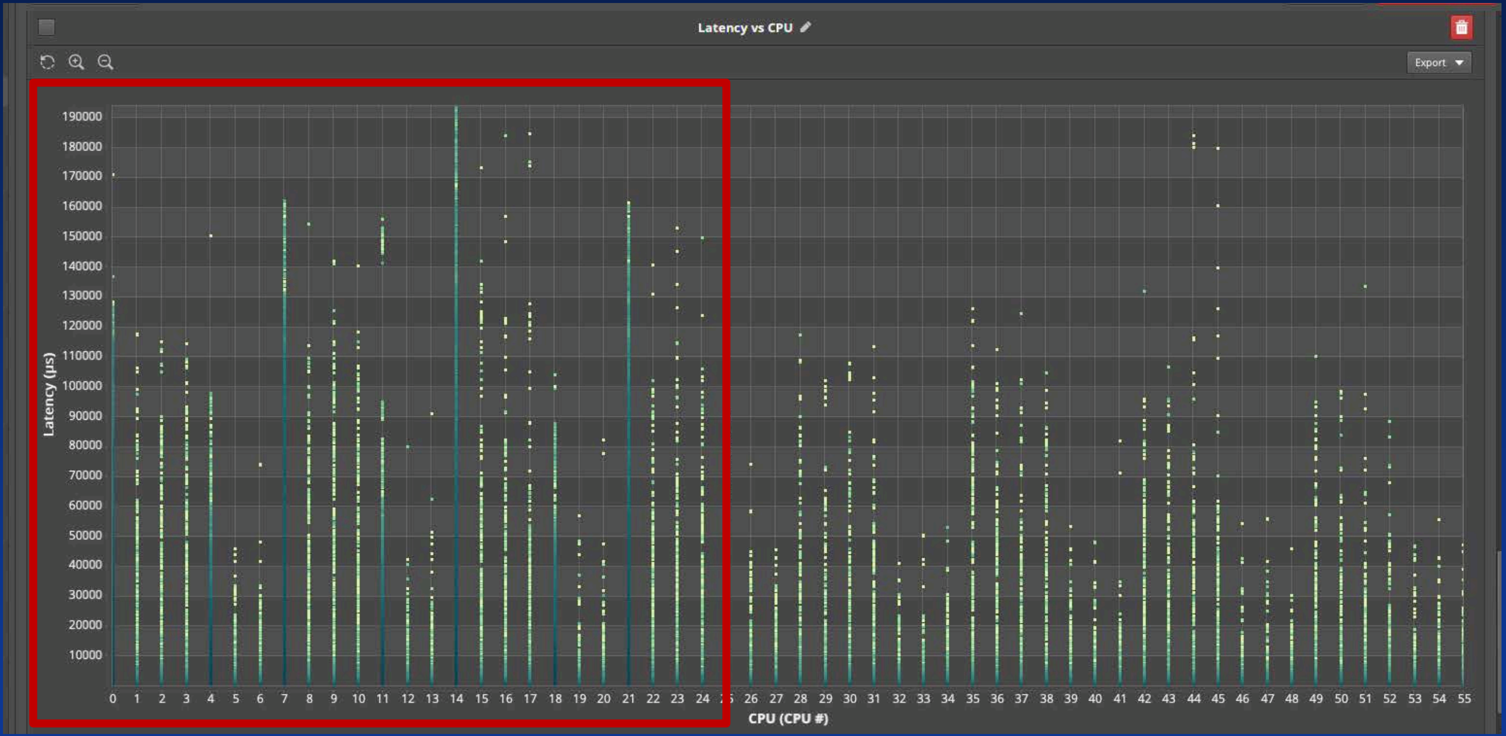

CPU vs Latency

Positive Customer Reactions

- Provided deeper insight into what is actually happening at the storage layer than we could get previously
- Provided “Feedback” mechanism to help the application team improve overall performance
- Workload Replay allowed comparison of workload behavior and performance across various drive vendors
- Provided insights to tune Linux block layer to achieve better workload performance

Wrap Up

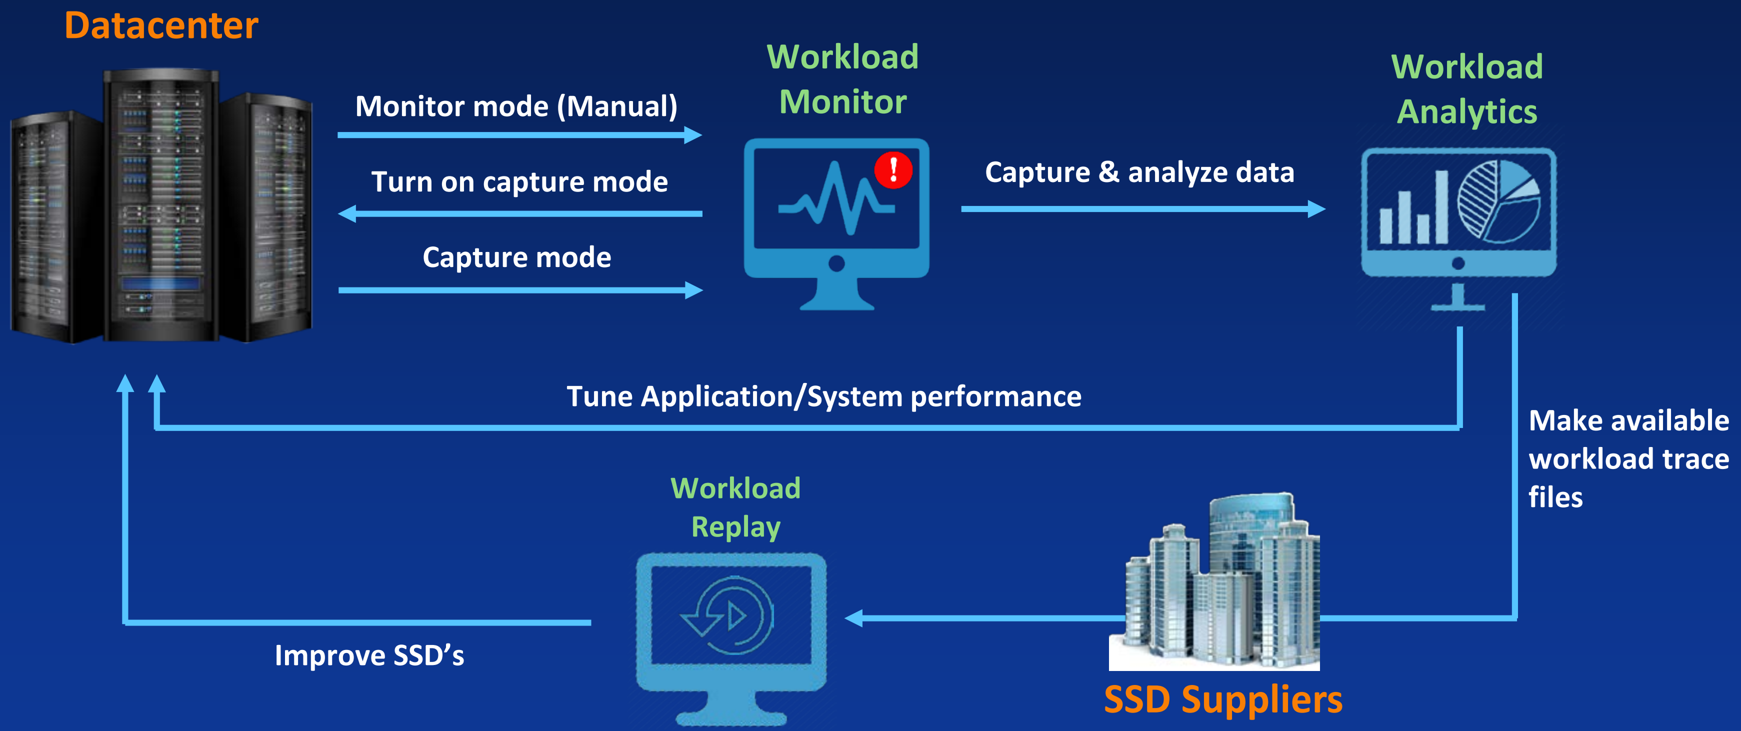
Datacenter
Monitor mode (Manual)
Turn on capture mode
Capture mode
Workload Monitor
Capture & analyze data
Workload Analytics
Tune Application/System performance
Make available workload trace files
Workload Replay
Improve SSD’s
SSD Suppliers

Audience Q&A