

September 23-26, 2019 Santa Clara, CA

Real World Data Center Workloads

& Advanced Storage Architectures

Eden Kim Calypso Systems, Inc.



Datacenter Real World Workloads

& Advanced Storage Architectures



The "Holy Grail"

SD (2)

- Real World Workloads have long been called the "Holy Grail" of Computing & Storage
- Advanced Software and Storage Design, Performance and Validation Depends on your Real World Workload
- Real World Workloads are very different from Synthetic Lab workloads and change depending on where the workload is captured

Real World Workloads are Important for Enabling Emerging Technologies



What Are Real World Workloads?

Constantly Changing IO Streams and QDs



IOs are Affected at Each Layer of Abstraction



Real World Workloads are Dynamic & Unique:

SD[®]

Σ	Cumulativ	e Workload	×			
🗹 F	RND 64K R	18.5%	842,361			
v :	SEQ 0.5K W	17.0%	775,127			
v F	RND 8K R	10.0%	456,175			
v :	SEQ 8K R	8.4%	382,972			
v F	RND 4K W	4.0%	182,251			
v :	SEQ 64K W	3.7%	169,571			
v :	SEQ 64K R	3.4%	155,798			
🔽 F	RND 4K R	2.92%	132,777			
I	RND 8K W	2.78%	126,550			
	SEQ 4K R	2.20%	100,309			
Tota Sel	Total IOs of 5,086 streams: 4,551,062 Selected 9 streams: 3,223,582 (71%) E					

IOPS

Dynamically Changing IO Streams & QDs



Comparing

Real World Workloads

Comparing Real World Workloads

1. Retail Web Portal: Drive0/Drive1



9 IO Streams = 71%	Ave QD
66:34 RW Mix	Median
Retail Store Events	Max QE

2. GPS Nav Portal: Drive C



67:33 RW MIX SEQ 128b R = 70% IOs

Max QD =

368

3. GPS Nav Portal: Drive 0



SD[®]

19

306

OD =



Analysis of Real World Workloads

Free Capture & Analysis Tools at TestMyWorkload.com

TestMyWorkload Log in | Log in Up Download Demo How To Use What It Means FAQ Support Download new White Paper by Calypso Systems: Real World Storage Workload: Retail Web Portal SQL Workload & Comparison Testing to Datacenter Storage Join Our Community to Help **Understand Your** . . . **IO Profiler: Prompt** Real World Workloads! **IO**profiler by CALYPSO Use the IOProfiler to capture, view and analyze your real world storage workloads on your laptop, desktop or server Define profiling parameters: 5 min ‡ Duration: 10 us 💠 Temporal resolution: Spatial resolution: 1 % \$ Capture level: ✓ block device file system mixed Download Start Register or sign in and download the IOProfiler applet for your Operating System (OS) to get your workload results. Start the capture by selecting the desired time and resolution or use the default settings. Once the capture is complete, you will be prompted to click a button to upload the capture to your account for view and analysis. Try several captures doing the same and different activities and see how it affects your workloads (reads/writes).



www.TestMyWorkload.com

SD (9)

- Official SNIA SSSI Real World Workload Capture Tools and Reference Workloads
- Free IO Capture Tools for Windows, Linux, Mac
- Reference SSS TWG and User Workload Captures
- White Papers, Tutorials, FAQ

IO Stream Map, Workload Description, PIDs

SD



IO Stream Map, LBA Range Hits, Metrics

1. IOs: IO Stream Threshold & ART/MRT



2. MB/s: Amount Transferred; Ave/Max QD



3. IO Stream & LBA Range: IO Sequentiality

SD (9



Workload Data Analytics: View by IOs or MBs; IO Streams & Threshold; Response Times; QDs; IO Bursts & Sequentiality

Workload Descriptions, PIDs, IO Streams, Metrics

4. Workload Descriptions



5. Process IDs

Workload Streams (by frequency) RND 64K R SEQ 0.5K W RND 8K R SEQ 8K R RND 4K W SEQ 64K R + 10PS 2 am Data Back-up Early M reve 74-Hr Work or Σ: \$ Clear/Select all Sort by name salservr.exe 79.9% 3.456.063 System 13.5% 584.528 2.3% 99,524 svchost.exe inetinfo.exe 2.1% 92,756 0.9% 40,446 w3wp.exe Isass.exe 0.5% 20 742 xcopy.exe 0.2% 10,708 8,707 rundll32.exe 0.2% 0.1% 5.600 iexplore.exe LogonUI.exe 0.1% 2,960 WerFault.exe 0.0% 1,397 18:00:10 20:00:00 22:00:00 2 Earl 0.0% 467 csrss.exe 451 Far.exe 0.0% 22 4% 774 341 SEO BK W 3 SEO 0.5K W RND 8K P 12.7% 438,71 SEQ 4K W winlogon.ex 0.0% 437 SEO SK P 11 034 300 505 6.2% 1,940 SEQ 64K R 4.4% 151,226 SEQ 0.5K W 4.19 SOLCMD EXE 0.0% 246 1,883 SRND 8K W 3.4% 118,791 S RND 1K W 3.1% 0.0% 213 dllhost.exe 1.00% 1,772 SEQ 512K R 1.73% 59,917 SRND 64K W 3.1% 1.60% 55,312 RND 2K W 1.428 RND 32K R 1.50% 52.001 RND 2.5K W 0.80% Selected 36 of 36 processes 0.72% 1.283 RND 16K R 1.49% 51.41 RND 8K W Apply Selection Cancel 178.057 Total IOs of 395 3 456 063 Total IOs of 14

6. IO Streams & Metrics

SD[©]



Workload Data Analytics:

Cumulative, Event, Time Point, Range Workloads; Process IDs (PIDs); IO Stream Combinations & Metrics

TRIMs, IO Burst, IO Sequentiality, Compression, Dedupe SD®

7. TRIMs



9. Compression & Duplication



Workload Data Analytics: TRIMs, TRIM MB/s, TRIM ART/MRT; IO Bursts & Sequentiality; Compression & Duplication

8. IO Burst & Sequentiality

Real Time Plot Playback - WAF, QoS, TPBW

SD (9



Workload Data Analytics: Write Amplification Factor (WAF); 5 9s QoS Response Times; Total PB Written (TPBW)



NVMe U.2 SSD

Optimization Using Real World Workloads

NVMe U.2 SSD: Replay Test of Retail Web Portal Workload

Replay Test: Retail Web Portal - U.2-H v U.2-M v U.2-S





Real World Replay: U.2-H > U.2-M > U.2-S

SD©

Individual Streams: RND 4K R/W vs SEQ 0.5K W

Synthetic Benchmark: RND 4K R/W: U.2-S > U.2-H

Real World Replay: SEQ 0.5K W: U.2-H > U.2-S

Retail Web Portal 9 IO Stream Workload

A0	Retail Web Portal – 73:27 RW			
	RND 6	4K R	26.1%	
	SEQ 0	.5K W	24.0%	
	RND	8K R	14.2%	
	SEQ 8	8K R	11.9%	
	RND 4	4K W	5.7%	
	SEQ 6	64K W	5.3%	
	SEQ 4	4K R	4.8%	
	RND	4K R	4.1%	
	RND	8K W	3.9%	



SD @



Persistent Memory

Optimization Using Real World Workloads

PM Optimization – PM Aware App: DRAM v DCPMM



SD[®]

PM R/W Optimization – R Dominant to PM, W Dominant to DRAM

All Data Plo	t	1						
Optane-FsDaxMmap2M 107 GB Test R17.1-8512. ipmctl - KB 2Drive 9 Stream 6s (28 min)Auto Replay NativeQD by MAX QD - Profile #5 Web Portal 2,000 Store Retail Chain 24 hr - Drive0 Drive1 Cumulative Workload (9 streams) - Step=6sec (WIN- C255D607N82_20161602_0000.xml)								
Time = 00:42	● IOPS = ● Avg RT = ● 5 9s = ● Max RT = ● CPU Sys = ● IO Wait = ● RW Mix = ● Total RW Mix = ● Fills = ● OIO ● WAF = TPBW/1Y 701,534 3,318 MB/s 0.0055 ms 0.20 ms 0.85 ms 0.194 % 0 % 5.30% R : 94.7% W 27.6% R : 72.4% W 160 = 7 3.19 = 292							
60 000	100 R 100 R <td< td=""><td>100% Read</td></td<>	100% Read						
14 (MB/s)								
20 000 Throughpu 20 000		- 0% F 0% \						
10 000		100%						
0	00:00 00:05 00:10 00:15 00:20 00:25 00:30 00:35 00:40 Compare Plots Playback Controls Close	- Write						

SD[®]

PM WAF Optimization – WAF by IO Streams; Endurance TPBW/Year





September 23-26, 2019 Santa Clara, CA

Questions?

Thank You!

Eden Kim, CEO Calypso Systems, Inc.

Chair - SNIA Solid State Storage Technical Working Group

edenkim@calypsotesters.com

www.TestMyWorkload.com

Please take a moment to rate this session.

Your feedback matters to us.