How To Be a Part of the Real-World Workload Revolution

Capture and Analyze Your Own Application Workload

Introduction by Tom Coughlin and Jim Handy
Authors: Eden Kim, Calypso Systems
Jim Fister, The Decision Place

Introduction

Real-world digital workloads often behave very differently from what might be expected. The equipment used in a computing system may function differently than anticipated. Unknown quirks in complicated software and operations running alongside the workload may be doing more or less than the user initially supposed. To truly understand what is happening, the right approach is to test and monitor the systems' behaviors as real code is executed. By using measured data designers, vendors and service personnel can pinpoint the actual limits and bottlenecks that a particular workload is experiencing. The SNIA Real-World Workload Capture program allows users to measure their own actual workload and compare it to other workloads doing similar operations and to optimize systems and configurations for the greatest performance and efficiency. In addition, capturing and sharing these anonymous workload traces will help the general community to anticipate and avoid common system problems and bottlenecks. We can all benefit by sharing traces of our digital workloads through the SNIA SSSI Real-World Workload Capture program.

Tom Coughlin, Coughlin Associates

Jim Handy, Objective Analysis

Table of Contents

About the SNIA

The Storage Networking Industry Association (SNIA) is a not–for–profit global organization, made up of member companies spanning the global storage market. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For more information, visit http://www.snia.org.

About the Solid State Storage Initiative

The SNIA Solid State Storage Initiative (SSSI) fosters the growth and success of the market for solid state storage and persistent memory in both enterprise and client environments. Members of the SSSI work together to promote the development of technical standards and tools, educate IT communities about solid state storage and persistent memory, perform market outreach that highlights the virtues of solid state storage and persistent memory, and collaborate with other industry associations on solid state storage and persistent memory technical work. SSSI member companies represent a variety of segments in the IT industry (see http://www.snia.org/forums/sssi/about/members).

For more information on SNIA's Solid State Storage activities, visit www.snia.org/forums/sssi and get involved in the conversation at http://twitter.com/SNIASolidState.

Executive Summary

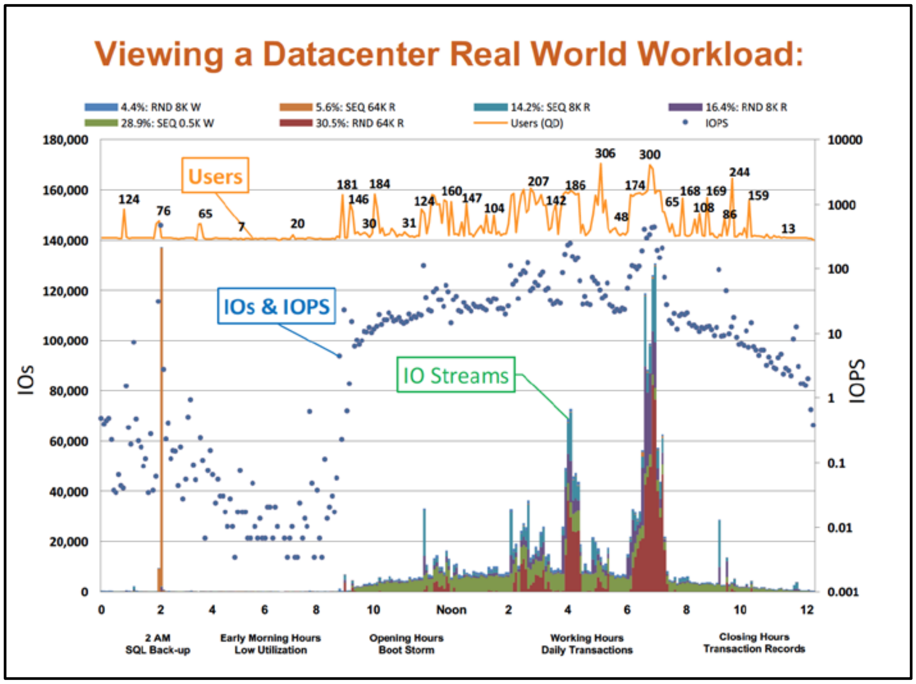

Analysis of real-world workloads has often been called the ultimate goal of computer storage solutions. This is because the workload function is often the most important factor in determining the performance of deployed systems. There is an opportunity to be a part of this revolution by capturing and sharing workloads from actual use.

By participating in the SNIA Solid State Storage Initiative (SSSI) Real World Workload Capture program, you contribute to the advancement of the industry as well as find what your workload looks like to aid in your storage purchasing decisions.

Real-world workloads can be generally defined as the combination and sequence of I/O Streams that occur when going from one point to another in the software/hardware stack. Real-world storage workloads are created when applications send and receive I/O streams to and from storage.

New I/O trace and capture tools, like the free tools available at www.TestMyWorkload.com, now allow any user to easily capture their real-world workload. This paper introduces and explains what real-world workloads are, how you can capture and share them, what they mean, and how a program like this benefits the industry overall. The audience includes IT, professionals, application developers, storage designers, manufacturers, and the intellectually curious user and technology geek.

Characteristics of a Real-World I/O Workload

Most system or software architects will talk about a particular workload in terms of its read/write balance, for instance, a “read-dominant system.” This is typically a measure of the overall mix of the workload as a cumulative measure, essentially measuring all the reads and writes over time and finding a basic percentage. However, at critical times the balance could vary widely from the general result.

A simple example would be video streaming. When looked at from the application layer, it intuitively appears to be sequential reads of large blocks. While this can be true at the file system, the underlying block I/O level can be very different, consisting of mostly random 4K and 8K blocks going to the disk. While real time video frames at the file system may be spooled and viewed from DRAM, the I/O Stream traffic to storage is usually a mix of random 4K and 8K read/write blocks which match the page size of the solid-state drive, or SSD.

Especially in modern solid-state drives, it’s likely that the file is distributed in available page sizes rather than in a single, contiguous block, further complicating the read and write profiles. Further, if there’s a need to transcode video streaming from a server to multiple clients, the actual results might be even more complex.

Real-world workloads are all unique, even those that do similar functions. They are comprised of a combination of I/O streams, different queue depths, and multiple process IDs in the system. A measure of any workload will find both application and system processes behaving in similar ways, though that behavior is likely to change over time with multiple runs.

Unlike synthetic workloads that are comprised of a fixed and constant workload intended to measure storage performance outside the range of normal operation, real-world examples are comprised of a constantly changing combination and sequence I/O Streams and queue depths. In addition to having many I/O Streams, these Streams are comprised of an assortment of transfer sizes from small byte size accesses to larger transfers of non-traditional sizes.

Real-world workload composition often does not match commonly understood definitions or assumptions. For example, a commonly accepted industry definition for a synthetic SQL workload is a random 8K 65:35 read/write mix. However, let’s look at a real-world read dominant (72:28 R/W mix) SQL workload for a retail web portal seen below. First, note that the mix is 72:28 R/W, not 65:35 R/W. Second, there are only 10% random 8K read and 2.78% random 8K write. Third, note that there were 5,086 unique IO Streams over 24 hours, 71% of which are comprised of the most frequently occurring nine I/O Streams.

Of these nine I/O Streams, the dominant four streams are random 64K read (18.5%), sequential 0.5K write (17%), random 8K read (10%), and sequential 8K read (8.4%) for a total of 63.9% of the total I/Os.

How To Be A Part of the Real-World Workload Revolution

Looking at the I/O stream map, we can see the sequence and combination of different I/O streams as they relate to different activities. The 2AM back-up load is comprised of 45.1% sequential 64K write and 41.6% sequential 64K read - reading from Drive0 and writing to Drive1. During the morning boot process, the I/O activity is dominated by sequential 0.5K write. Finally, the daily transactions are comprised of a collection of I/O streams reflecting the dominant I/O streams over a 24-hour period.

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

A 2 am Back-up		
☑ SEQ 64K W	45.1%	156,395
☑ SEQ 64K R	41.6%	144,397
☐ RND 4K W	1.83%	6,362
☐ SEQ 2M R	1.40%	4,851
☐ RND 4K R	0.95%	3,294
☐ RND 32K R	0.75%	2,595
☐ RND 24K R	0.56%	1,940
☐ SEQ 60K R	0.54%	1,883
☐ SEQ 40K R	0.51%	1,772
☐ RND 8K R	0.44%	1,519
Total IOs of 298 streams: 346,825		
Selected 2 streams: 300,792 (87%) E		

Other system activity will affect the I/O characteristics. If there is significant user or application-level activity, the queue depth of the I/O stream will likely increase, affecting the overall behavior of the individual application. Systems today run a variety of background processes which will increase overall I/O activity. Finally, the configuration of the application will also change the workload characteristics. Running on a virtual machine, for instance, will make for a different pattern than running on bare metal. As such, it's useful to take consecutive captures over time as the workload content changes over time and evolves with the system use case and platform environment.

Obviously, the large variety of real-world workloads, and the variability of any given one, complicates both the development and proliferation of I/O devices in the marketplace. While modern solid-state disks cover a variety of use cases, the variability of both the hardware and the applications that use them creates a need to better understand the environment before building a solution, especially one for scalable deployment across a large enterprise.

Opportunity: Provide Workloads to Improve the Marketplace

As noted, the Storage Networking Industry Association (SNIA) is a global organization dedicated to developing standards and education programs to advance storage and information technology. The Solid State Storage Initiative (SSSI) in SNIA supports the acceptance and growth of Solid State Storage (SSS) and Persistent Memory (PM) in the marketplace. Our member companies, such as Calypso Systems, educate vendor and user communities, and support SNIA technical work. By working together, the SNIA SSSI participants benefit the industry.

The SNIA is an impartial arbiter of the industry that focuses on advancing the proliferation of solid-state technologies in their many forms. As a SNIA and SSSI member, Calypso Systems is an SSSI-certified performance test lab providing third-party testing using advanced methodologies and test systems. SNIA utilizes Calypso RTP/CTS as the test system of record for capturing I/O workloads.

By providing workloads, participating companies can help expand the information used to develop and test solid-state storage products and drivers to better serve the industry overall. While this benefits the industry overall, it most benefits the participants by creating product better suited to their application characteristics. Additionally, the participants have the ability to get information on their own workloads immediately, which enables better development and deployment options. SNIA and SSSI are calling for workloads at all stages of development.

Capturing and Analyzing Workloads

There are three components to utilizing real-world workloads:

- Capture: SSSI and Calypso offer free I/O capture applets available for download for Linux, Windows, and Mac. The applets and data repository can be located at TestMyWorkload.com. After loading the application, open and start the capture after defining the capture duration/length, time resolution, and other factors.
- Analyze: All data is collected as binary information in tables. No personal data or actual data is collected. The I/O capture process averages statistics on each observed I/O and saves the data locally or at the TestMyWorkload site. It is then processed into an I/O stream map for immediate or later analysis.
- Test: The captured workloads can be converted into sophisticated test scripts that can be used on a variety of hardware. Users can test any logical storage combinations connected directly or indirectly to the server containing the capture data. In this way, the workloads are immediately useful as a test bed for performance.

How To Be A Part of the Real-World Workload Revolution

Typically, a capture process will have almost no overhead, between 1-4% of CPU usage. As said above, no actual data is captured, but representative block data is tracked and stored. Therefore, there are no concerns for user or application data. Captures can be taken from any logical storage recognized by the target OS, at either file system or block level. Additional types of captures can be tuned using customized tools available commercially from Calypso.

The results of the test can be visualized in different ways. Simple I/O stream maps detail the actual read and write data, providing instant understanding of the workload characteristics. The maps can be parsed to provide significant information at any point of the capture, allowing analysis of key areas that most impact performance.

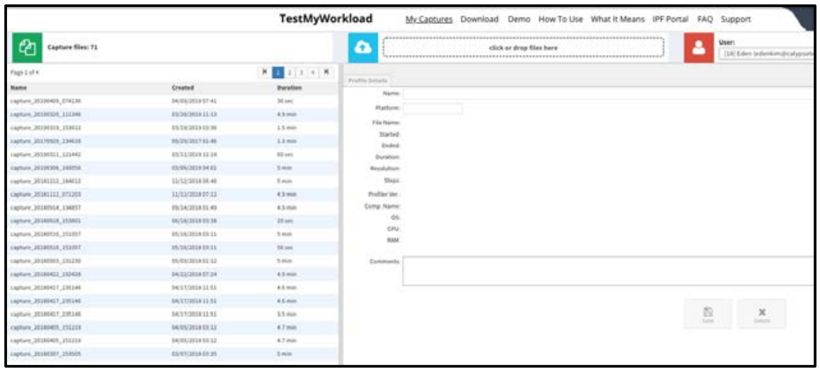

Additional visual options allow the workload to be displayed showing spatial locality of each I/O by logical block address (LBA) range, random or sequential accesses, and the like. This would enable the user to examine storage tiering strategies, view traffic on the various actual devices, and other useful analyses. The workload can also be filtered by process ID, time, event, and a variety of other views.

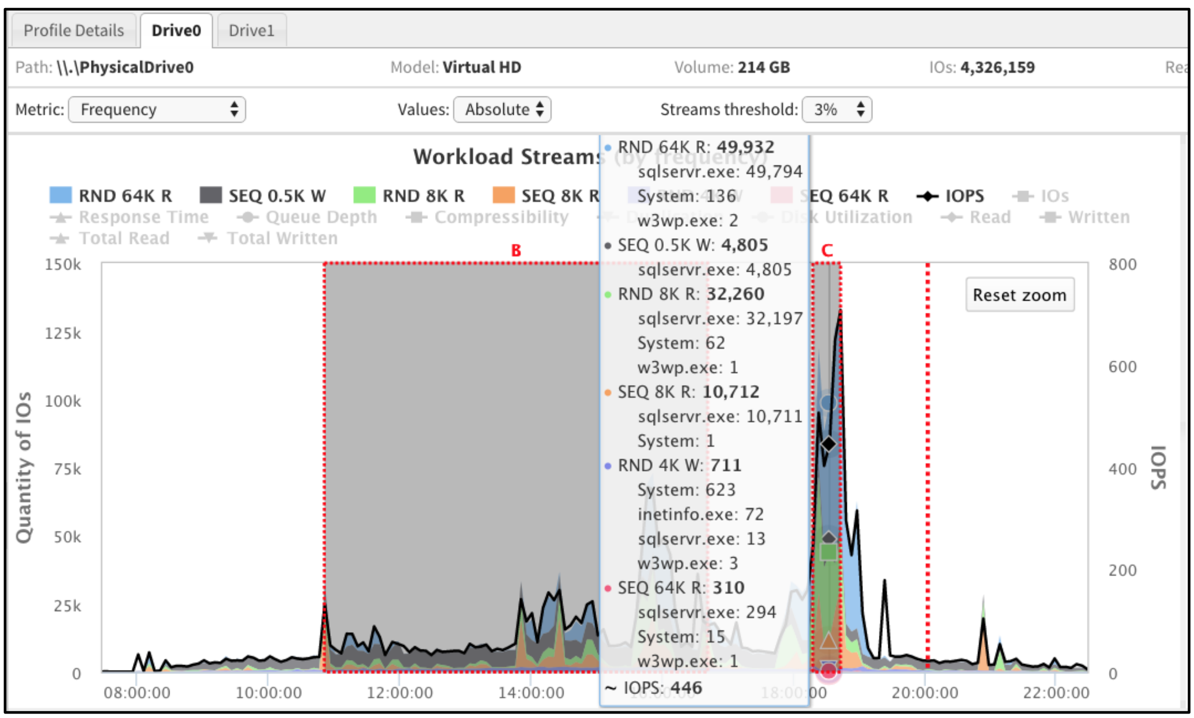

How To Be A Part of the Real-World Workload Revolution

Test scripts can be created from the parsed I/O stream map. Once captured, the replay can select the entire workload, or can focus on key areas such as individual streams critical to the application. Using the commercial tool, automatic scripts can be generated that provide even more detailed test methods that enable more workload and performance testing.

There are a variety of benefits that result from building an anonymized set of workloads.
From an industry hardware perspective, this provides a better guideline for the development of new technologies. Having the data to test methodologies in the design phase makes a better product.

Software developers can use workload testing to better understand the characteristics of the application. In many cases, the structure of a program can dictate the I/O activity, and it's possible that better knowledge during development creates better-behaved programs in the real world.

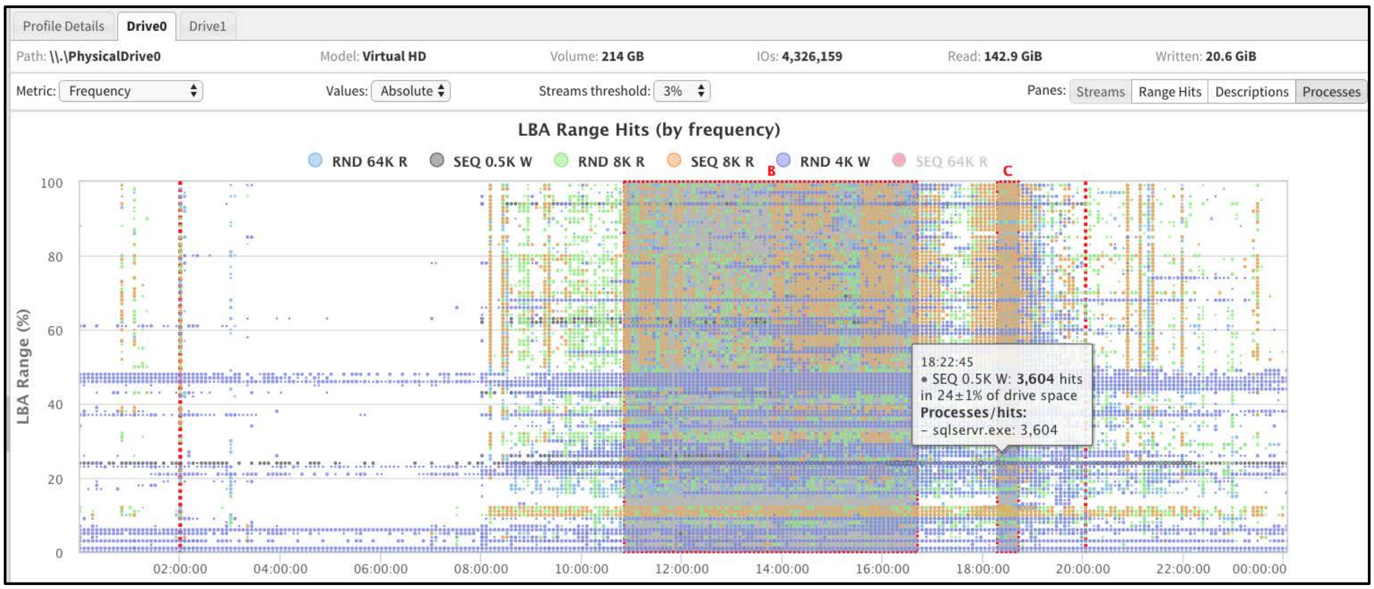

IT professionals and consultants can utilize the increased knowledge of their environment to make more informed decisions on the makeup of new systems and the workload distribution on existing architecture. Especially where large-scale deployments are concerned, the configuration of local and remote storage can make a significant difference in the capability, and costs, of the long-term solution.

Conclusion: How to Benefit

It's very likely that utilizing advanced tools such as workload I/O monitoring and analysis can improve localized results for any user. However, the combined results of many workloads can provide input to manufacturers and software developers to build better products. As such, the need for everyone to act is critical to the health of the industry.

For IT organizations and solutions providers, a simple workload test provided to TestMyWorkload.com can help improve the knowledge base used to build the solutions themselves. Further, using the simulations from software providers can give a better idea of the right configuration in the R&D phase, ensuring that better performance and cost targets can be achieved.

Software developers have an opportunity to provide baseline workloads for use by IT and hardware providers, and they can also utilize other variants of their own software to better determine real-world conditions. The opportunity to find corner cases that allow a company to fix race conditions ahead and to test new I/O hardware more easily provide significant benefit to the company.

For hardware manufacturers, the opportunity is to build a greater expanse of tests that better characterize the needs of solid-state memory and storage. Particularly, driver development can be adjusted to be more responsive to actual system demands instead of synthetic ones.

For more information and to view/contribute to the repository and download tools, go to TestMyWorkload.com.

About the Authors

Eden Kim is CEO of Calypso Systems, Inc., a leading test and measurement company specializing in solid state and persistent memory performance test. Eden is Chair of the SNIA Solid State Storage Technical Working Group, which authored the PTS 2.0.1 for SSD and RWSW PTS 1.0.7 for Data Center Storage. Eden is also Chair of the SNIA SSSI TechDev Committee and member of the SSSI Governing Board. Mr. Kim has authored various white papers which can be found on the SSSI knowledge page at www.snia.org/forums/sssi/knowledge/whitepapers.

Jim Fister is Principal of The Decision Place, a company focused on driving new technology and business strategies to the marketplace. He worked for Intel for 26 years, and has over 30 years of computing industry experience. Jim currently acts as Director of Technology and Application Enabling for Persistent Memory and Solid State Drives for SNIA. When not kicking dirt clods all over trails in Central Oregon, Jim also chairs a non-profit focused on STEM and CTE opportunity for youth.

Jim Handy of Objective Analysis has over 35 years in the electronics industry including 20 years as a leading semiconductor and SSD industry analyst. He has worked at leading semiconductor suppliers including Intel, National Semiconductor, and Infineon. A frequent presenter at trade shows, Mr. Handy is known for his technical depth, accurate forecasts, widespread industry presence and volume of publication. He has written hundreds of market reports, articles for trade journals, and white papers, and is frequently interviewed and quoted in the electronics trade press and other media. He posts blogs at www.TheMemoryGuy.com and www.TheSSDguy.com.

Tom Coughlin, President, Coughlin Associates is a digital storage analyst and business and technology consultant. He has over 37 years in the data storage industry with engineering and management positions at several companies. Coughlin Associates consults, publishes books and market and technology reports (including The Media and Entertainment Storage Report), and puts on digital storage-oriented events. He is a regular contributor for forbes.com and M&E organization websites. He is an IEEE Fellow, President of IEEE-USA, and is active with SNIA and SMPTE. For more information on Tom Coughlin and his publications and activities go to www.tomcoughlin.com.

SNIA SSSI | SOLID STATE STORAGE

Storage Networking Industry Association
4360 Arrows West Drive • Colorado Springs, CO 80907 • Phone: 719-694-1380 • Fax: 719-694-1389 • www.snia.org