How IOs Traverse the I/O Software/Hardware Stack

And Its Impact on Performance and Testing

SSSI White Papers - Expanded from "*Understanding SSD Performance Using the SNIA SSS PTS*"
www.snia.org/forums/sssi/knowledge/education

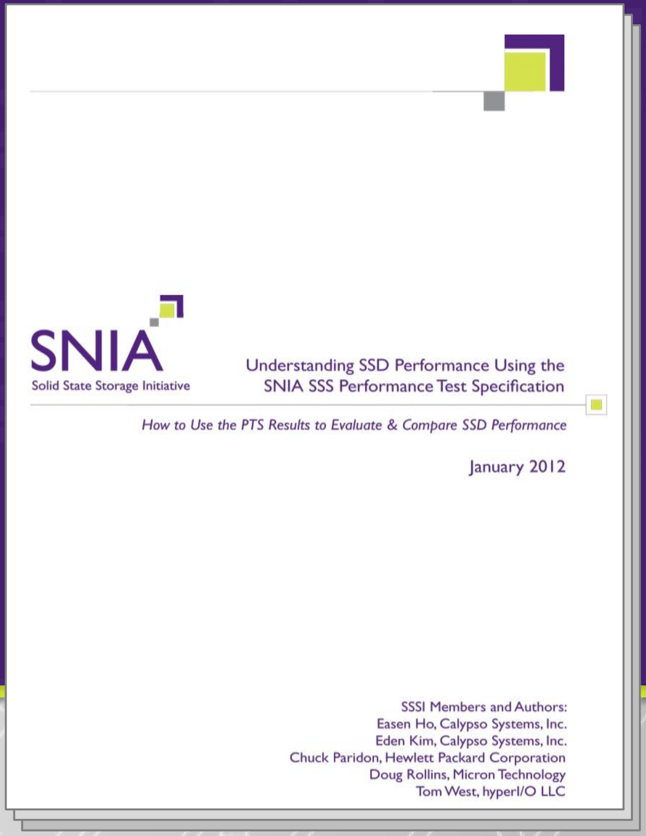
SNIA
Solid State Storage Initiative

Understanding SSD Performance Using the SNIA SSS Performance Test Specification

How to Use the PTS Results to Evaluate & Compare SSD Performance

January 2012

SSSI Members and Authors:
Easen Ho, Calypso Systems, Inc.
Eden Kim, Calypso Systems, Inc.
Chuck Paridon, Hewlett Packard Corporation
Doug Rollins, Micron Technology
Tom West, hyperI/O LLC

SNIA
Solid State Storage Initiative

About the Presenters

Eden Kim is Chair of the SNIA Solid State Storage Technical Working Group and a member of the SNIA Solid State Storage Initiative Governing Board.

Mr. Kim was recognized in 2010 as the SNIA Outstanding Contributor for his work with the Solid State Storage Initiative and SSS Technical Working Group. Mr. Kim has been Chair of the SSS TWG since 2009 and has ushered the PTS through to publication.

Mr. Kim is CEO of Calypso Systems, Inc. which is the developer of the Calypso RTP / CTS SSD test platform. Calypso provides SSD Test and Measurement equipment and services to the solid state storage industry.

Mr. Kim previously founded hard disk drive test companies Media Measurements, Inc., Swan Instruments, Inc. and acquired Scotts Valley Instruments. Mr. Kim received his BA/JD from the University of CA.

Tom West has been a co-chair of the SNIA I/O Traces, Tools, and Analysis (IOTTA) Technical Working Group (TWG) since 2005 and led the establishment of the SNIA IOTTA Repository. He is also the author of the Filesystem Performance Profile within the SNIA Storage Management Initiative Specification (SMI-S) and actively participates in the SSS TWG, particularly in I/O trace analysis related to the PTS.

Tom is a named inventor in 2 U.S. patents. His storage background includes 15 years at Storage Technology Corp., where he held advisory engineering positions in Systems Engineering and led the design/development of IBM-compatible enterprise mainframe disk storage subsystems.

Tom is the President of hyperI/O LLC, an Independent Software Vendor that designs and develops innovative software tools for both measuring and monitoring disk and file I/O operation performance; the firm has also been engaged in SSD firmware consulting.

Download this deck at www.snia.org/forums/sssi/pts

1. **Brief Background Information about the PTS**
2. The IO Software/Hardware Stack
3. Examples of How IOs Traverse the IO SW/HW Stack
4. Impact of Traversing the IO SW/HW Stack
5. Conclusion

SNIA PTS-C & PTS-E Specifications: Standardizing SSD Performance Test

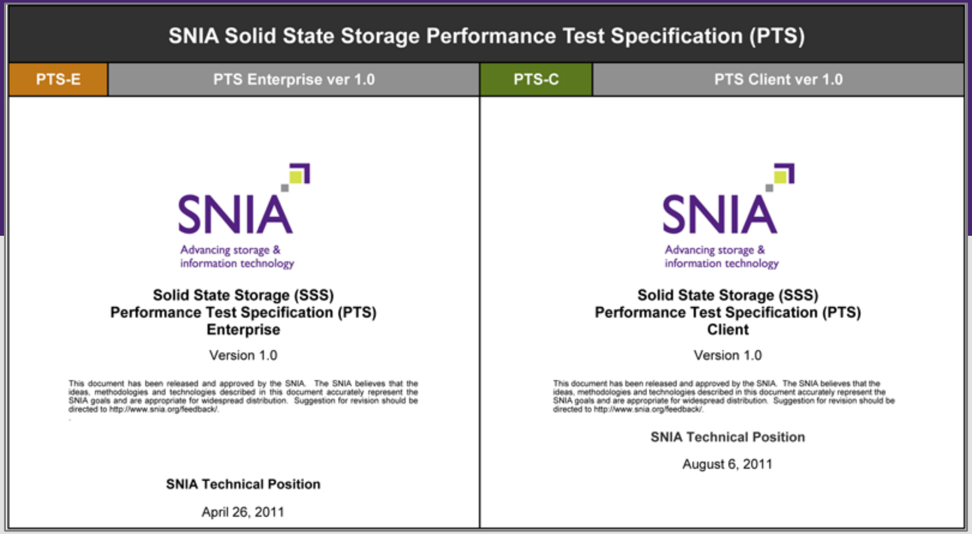

SNIA SSSI Solid State ***Performance Test Spec*** link:

www.snia.org/tech_activities/standards/curr_standards/pts

Understanding SSD Performance Project link:

www.snia.org/forums/sssi/pts

Understanding SSD Performance ***White Paper & Powerpoint*** link:

www.snia.org/forums/sssi/knowledge/education

Understanding SSD Performance ***Webcast*** link:

www.brighttalk.com/webcast/663/40549

PTS Provides a Standardized Methodology to Compare SSD Performance

Download this deck at www.snia.org/forums/sssi/pts

1. Brief Background Information about the PTS
2. **The IO Software/Hardware Stack**
3. Examples of How IOs Traverse the IO SW/HW Stack
4. Impact of Traversing the IO SW/HW Stack
5. Conclusion

What is the IO Software/Hardware Stack?

- Disk I/O operations are handled by a particular stack of I/O-related components (the IO SW/HW stack)
  - ✓ Operating-system (OS) software
  - ✓ Hardware components
- IO operations issued by applications traverse through the various components within the IO SW/HW stack
- A variety of factors and variables can impact:
  - ✓ The use of these components
  - ✓ The interactions amongst these components

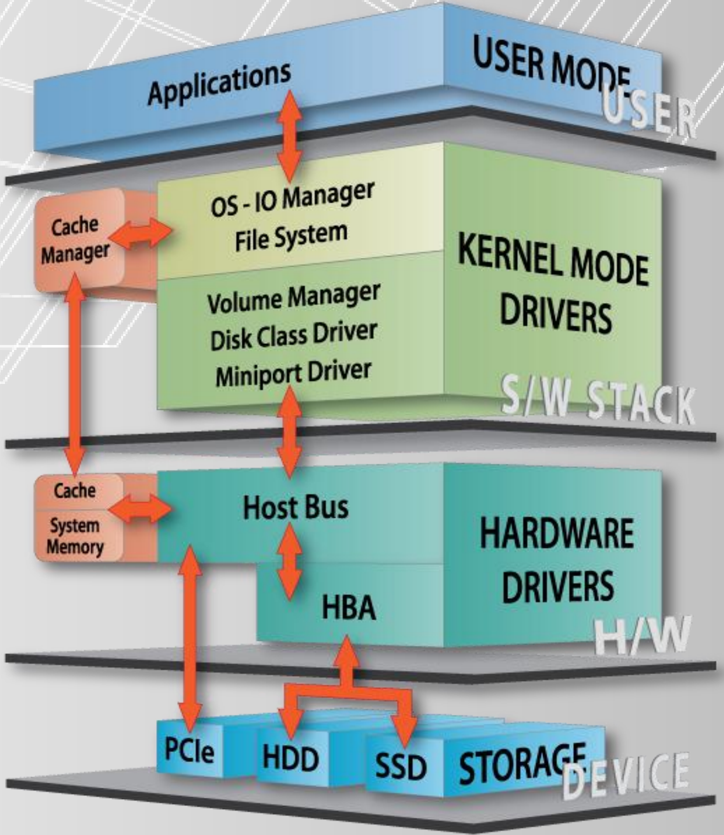

IO SW/HW Stack: OS Software and Hardware Components

Why Is Understanding the IO Stack Important?

- How IOs traverse the IO SW/HW stack has a direct impact upon performance (also data integrity, cost, etc.)
- Better understanding of the IO stack can be beneficial when:
  - ✓ Resolving performance problems
  - ✓ Managing storage Quality-of-Service (QoS)
  - ✓ Configuring overall system options and usage
  - ✓ Evaluating storage device options, etc.
  - ✓ **Considering IO performance testing approaches and results**

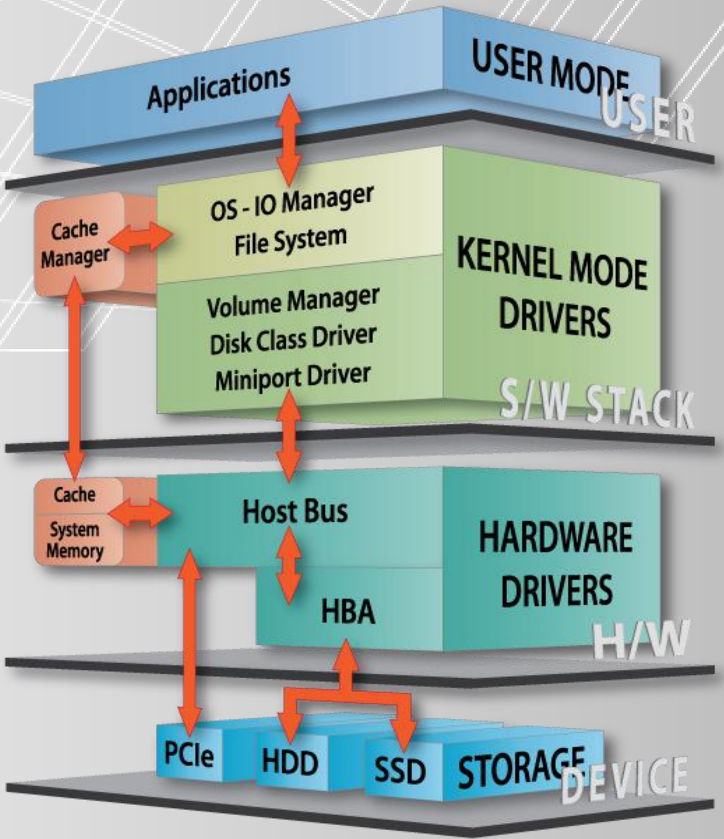

IO Stack – Directly Impacts IO Performance and Testing

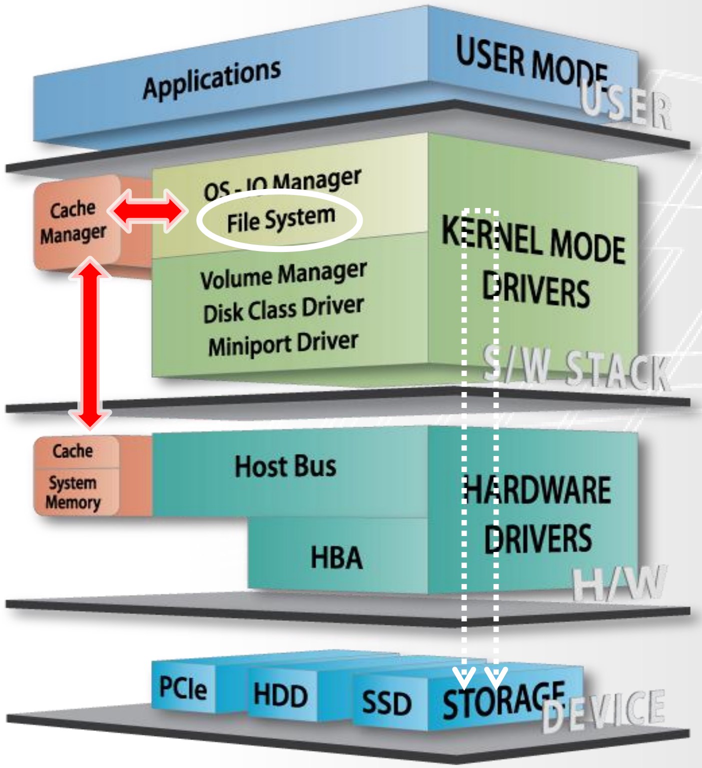

File System Test

- Specific File IO operations are issued by using the File System
- The IOs can be subject to **system file caching** (depending upon the particular test tool/program used)
- **Location of file** upon storage device is dependent upon allocation by the File System
- IOs can also be subject to **file fragmentation**
- Original IO can be different at the Device level due to the factors noted above – can lose 1-to-1 correspondence

File System Testing is subject to a variety of IO Stack factors

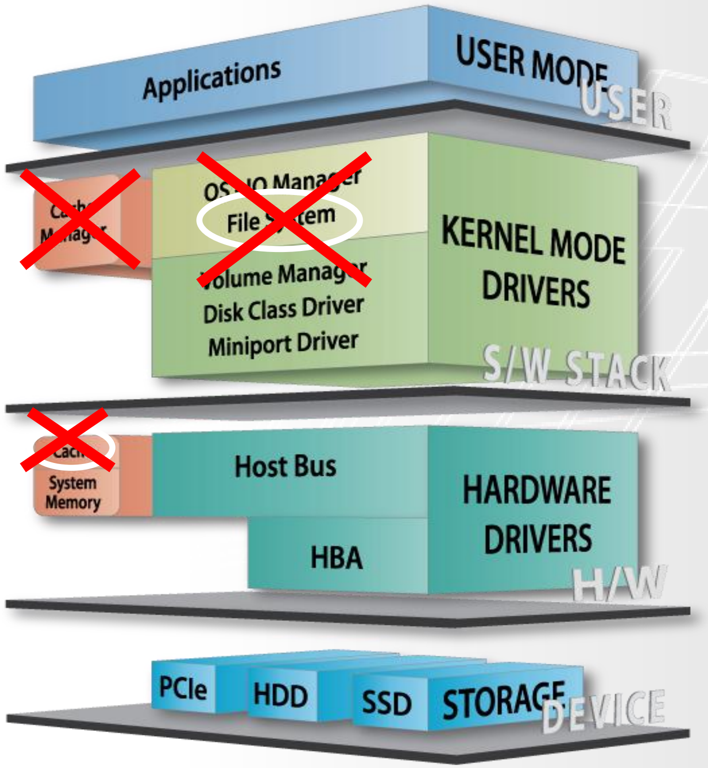

PTS Synthetic Device Level Test

- Applies a known, repeatable test stimulus that targets Block IO Devices directly (does **not** go through the File System)
- Uses Specified Test Workloads (Access Patterns, Data Pattern)
- Specifies LBAs allowed to be used (ActiveRange and ActiveRange Amount)
- Prescribes the specific Test Methodology to be used

Focuses on **innate** Storage Device performance capabilities

Download this deck at www.snia.org/forums/sssi/pts

1. Brief Background Information about the PTS
2. The IO Software/Hardware Stack
3. **Examples of How IOs Traverse the IO SW/HW Stack**
4. Impact of Traversing the IO SW/HW Stack
5. Conclusion

Cached IOs: Reads and Writes

1. Application issues file IO operation
2. File System passes to Cache Manager
3. Cache Manager Data Transfer
4. File IO Operation Completion

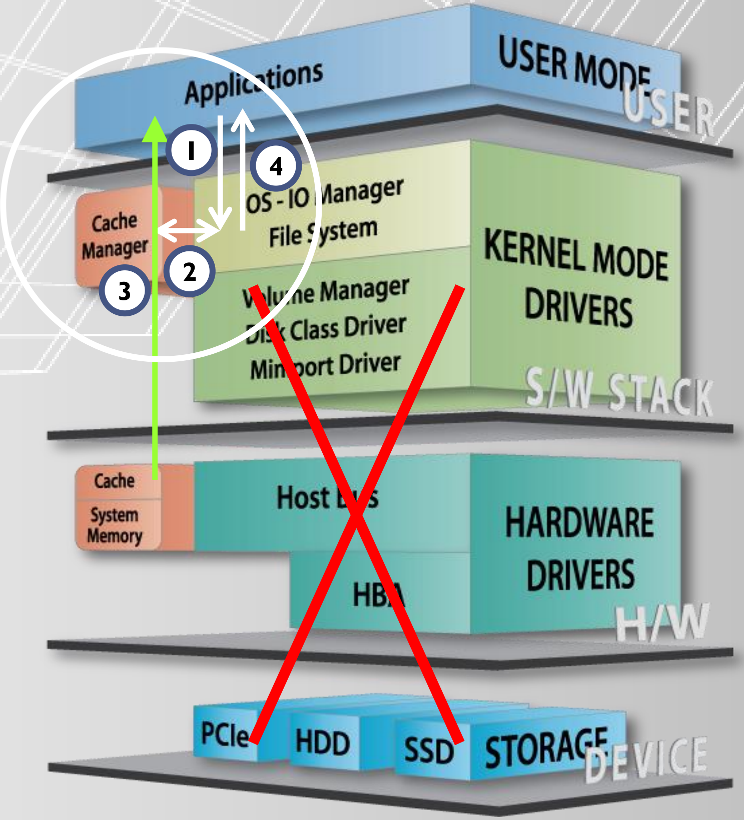

Not every application IO results in a Device IO

Cached IOs: Read Ahead

1. Cache Manager "Prefetch" processing
2. Read IO operation traverses the stack
3. Data Transferred to the System File Cache

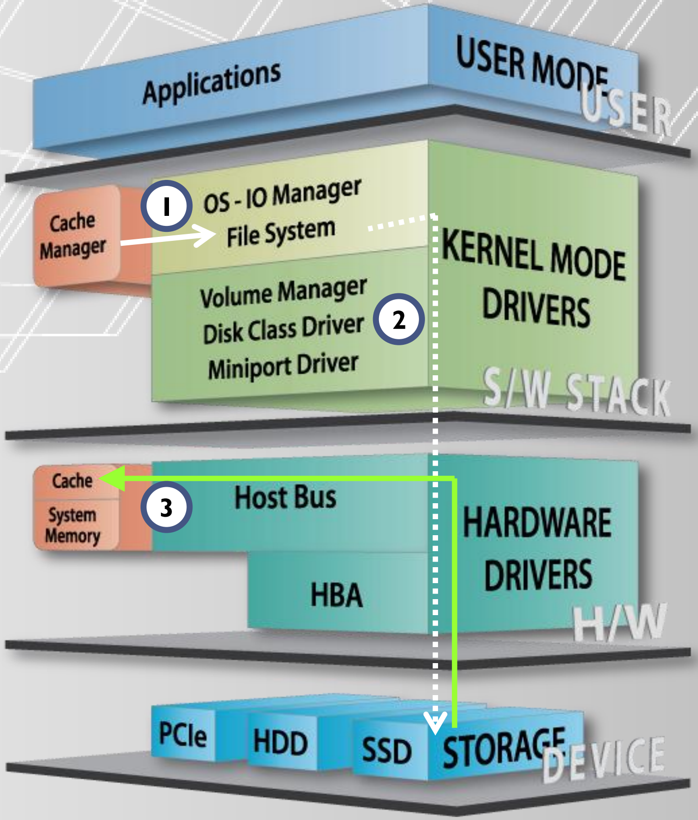

Are separate from – and can be concurrent with – application IOs

Cached IOs: Write Behind

1. Cache Manager "Lazy-Write" processing
2. Write IO operation traverses the stack
3. Modified Data Transferred to the Storage

***Note** that until it is written to storage, the modified data in the system file cache is subject to loss of data in the event of a power loss.*

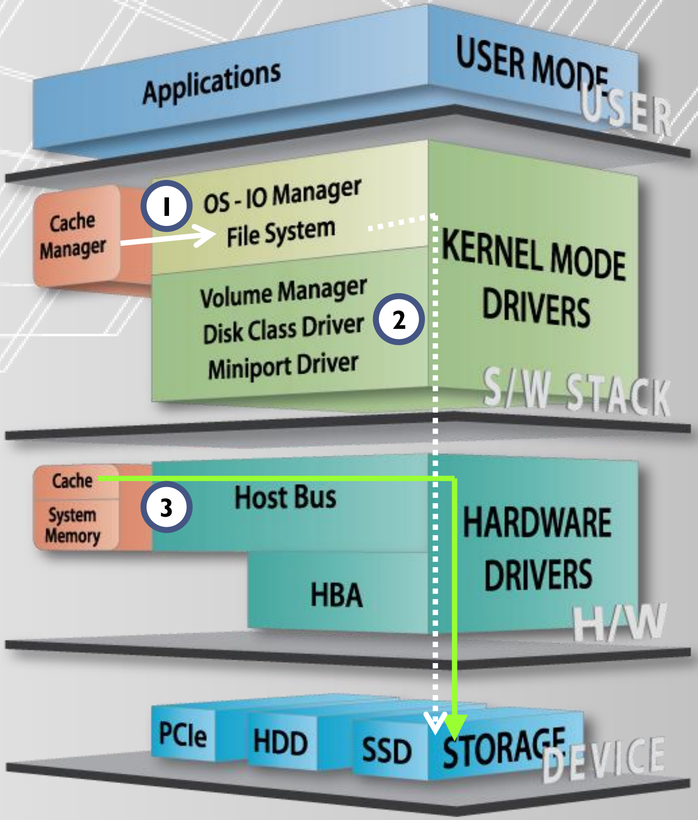

Also separate from – and can be concurrent with – application IOs

SNIA Solid State Storage Initiative

Fragmented File IOs

1. Application issues file read or write IO
2. File System processes the File IO
3. Storage receives **multiple** IO commands

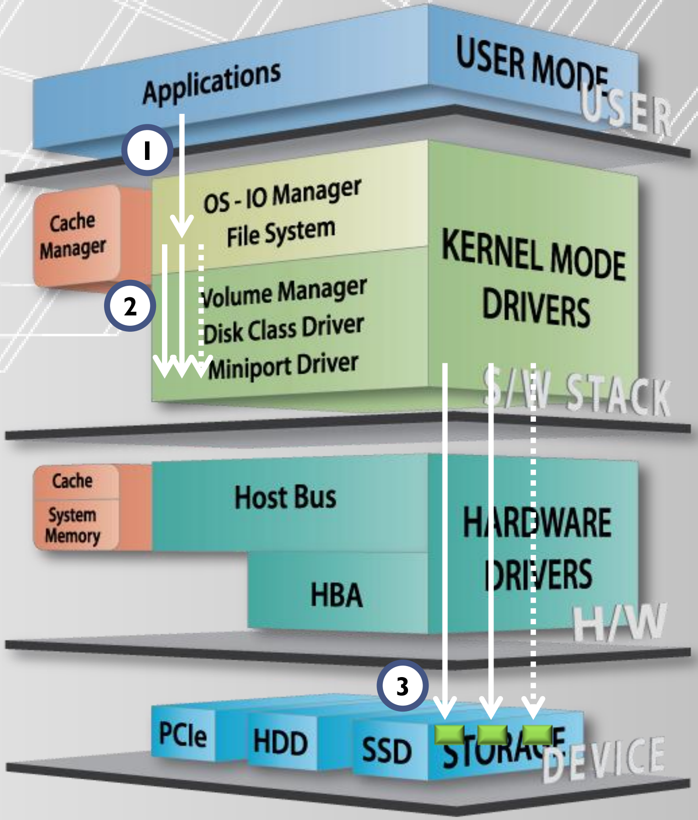

Application IO count is *NOT* equal to the Device IO count

Split IOs due to Large Data Transfers

1. Application issues file read or write IO

2. File System processes the File IO

3. <u>Multiple</u> Device IO commands due to the large data transfer size (e.g., 256 KiB)

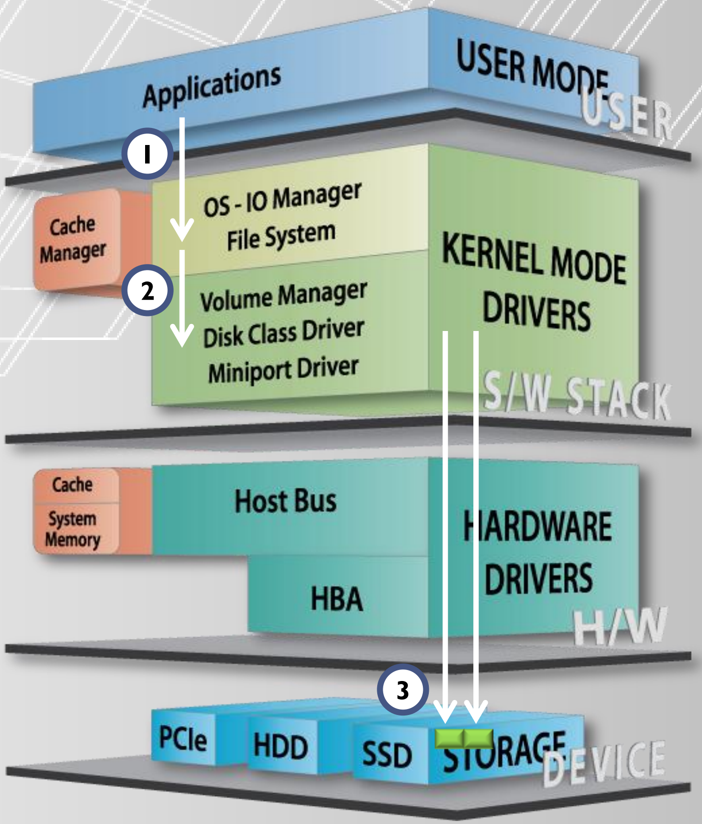

Application IO count is *NOT* equal to the Device IO count

Concurrent Sequential Access IOs

1. Application Yellow issues read IO to its **file**
2. Application Red issues read IO to its **file**
3. File System processes the File IOs
4. Two Device IO commands, each to a separate non-adjacent location upon the storage device

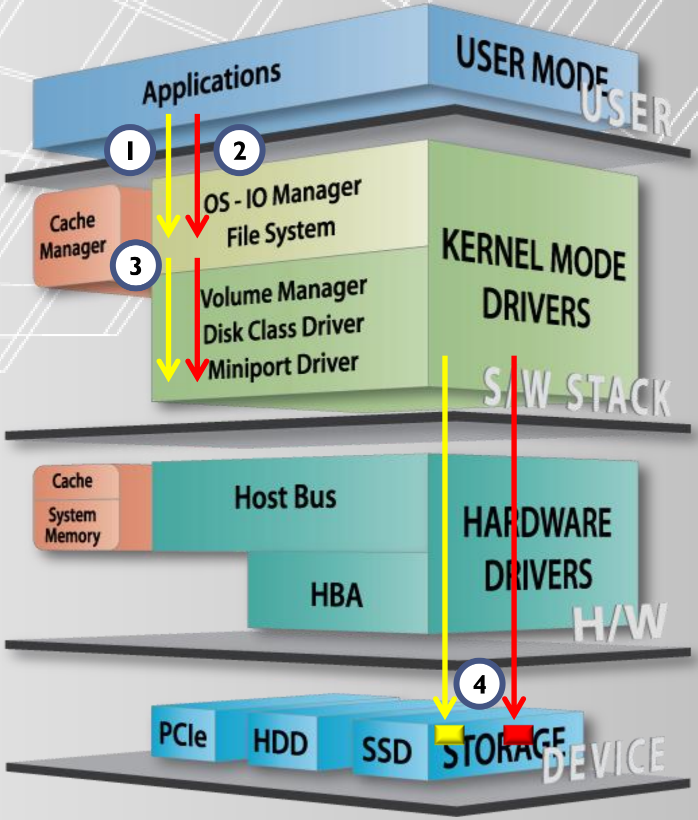

The second Device IO is a RANDOM IO from the perspective of the Storage Device

Concurrent Sequential Access IOs (continued)

5. Second sequential **file** IO by App Yellow
6. Second sequential **file** IO by App Red
7. File System processes the File IOs
8. Two Device IO commands, each to a separate non-adjacent location upon the storage device

Various concurrent IO operations, "fragmented file" IO operations, timing considerations and other factors can have an impact upon the actual access patterns.

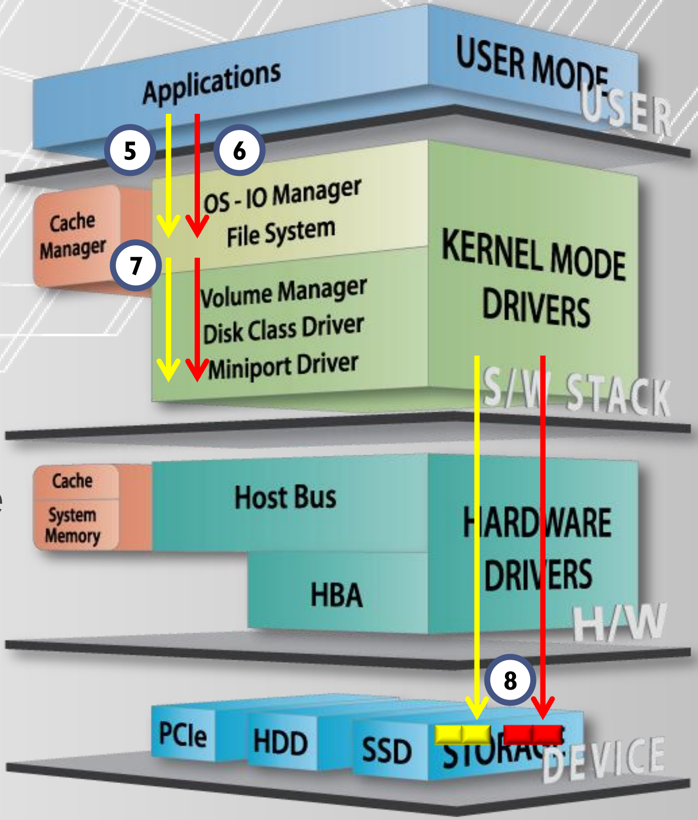

RND and SEQ accesses are RELATIVE to where IOs are performed/observed within the IO Stack

File System “MetaFile” IOs

1. Application OPENS file
2. File System processes OPEN request, which might require metafile data transfer IO
3. Application issues file write IO, then closes file
4. File System updates metadata within metafile, which (eventually) requires data transfer IO to the storage device

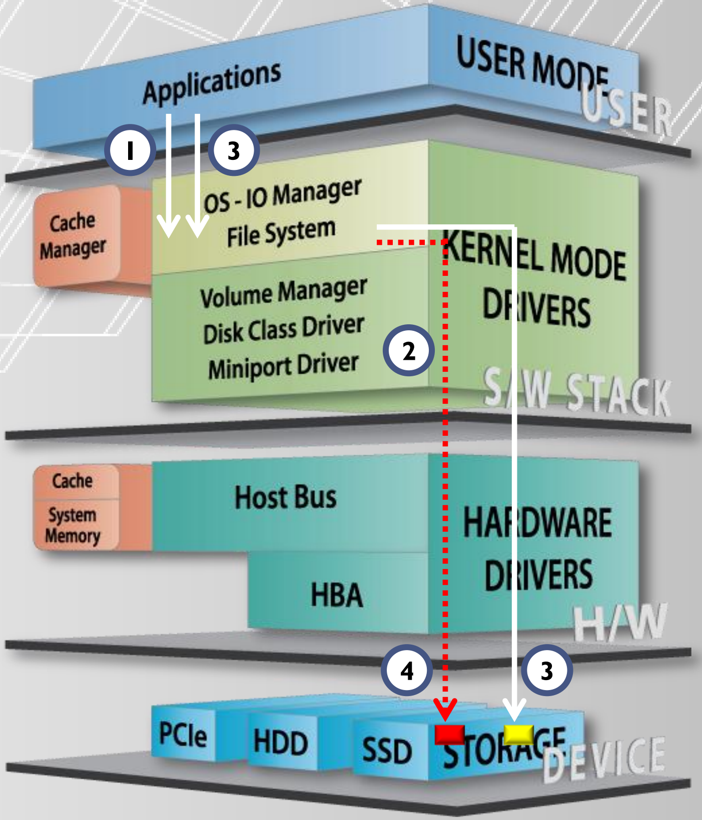

Can introduce RND accesses when intermingled with application IOs

Download this deck at www.snia.org/forums/sssi/pts

1. Brief Background Information about the PTS
2. The IO Software/Hardware Stack
3. Examples of How IOs Traverse the IO SW/HW Stack
4. **Impact of Traversing the IO SW/HW Stack**
5. Conclusion

IOs are Impacted by Traversing the IO Stack

The various paths that an IO operation might take through the IO SW/HW stack can have an impact upon (as examples):

- ✓ IO operation response times
- ✓ Storage device usage and utilization
- ✓ CPU and overall system utilization

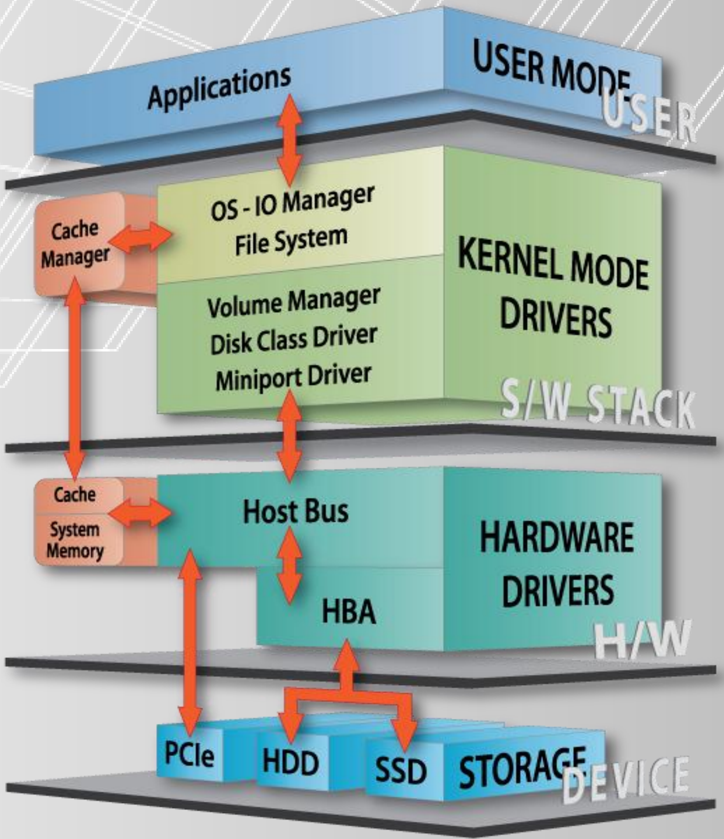

Traversing the IO Stack can have a **significant** impact upon various performance factors

Several Key Impact Considerations

- Application IO activity does not necessarily equate to storage device IO activity
  - ✓ IO operation performance metrics can be different
- Consider various viewpoints of IO operation performance:
  - ✓ An individual IO SW/HW stack component perspective
  - ✓ An overall system perspective
  - ✓ The particular application/workload perspective

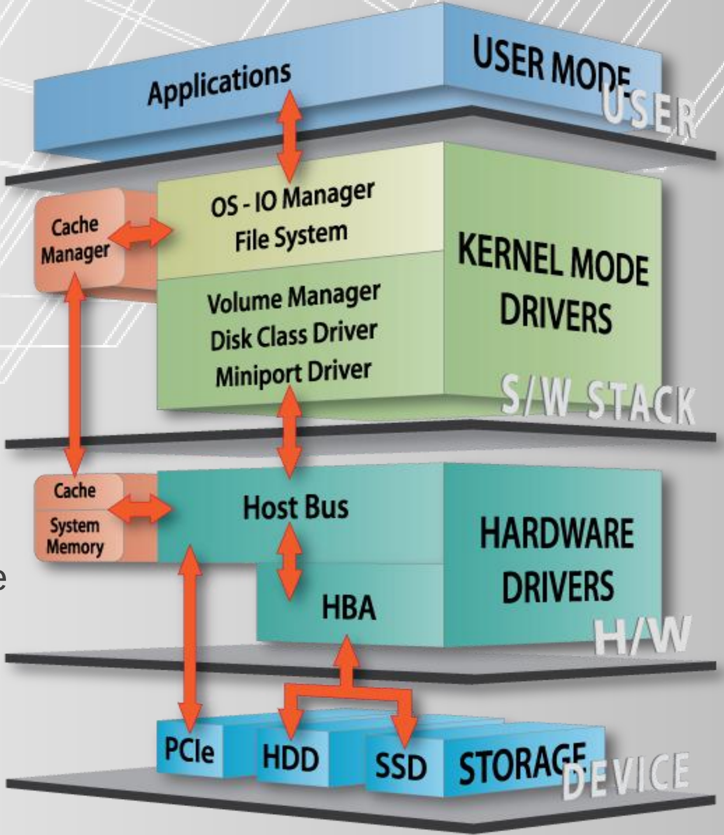

Perspective: Observed Performance depends upon where you are standing in the stack

Impact upon IO Performance Testing

What and ***how*** of tools used to test IO performance:

- ✓ Which particular IO SW/HW stack components are actually being used and exercised?
- ✓ How does the tool actually operate?
- ✓ How are the tool's reported IO performance measurements and "scores" actually calculated?
- ✓ How well documented are the items above?

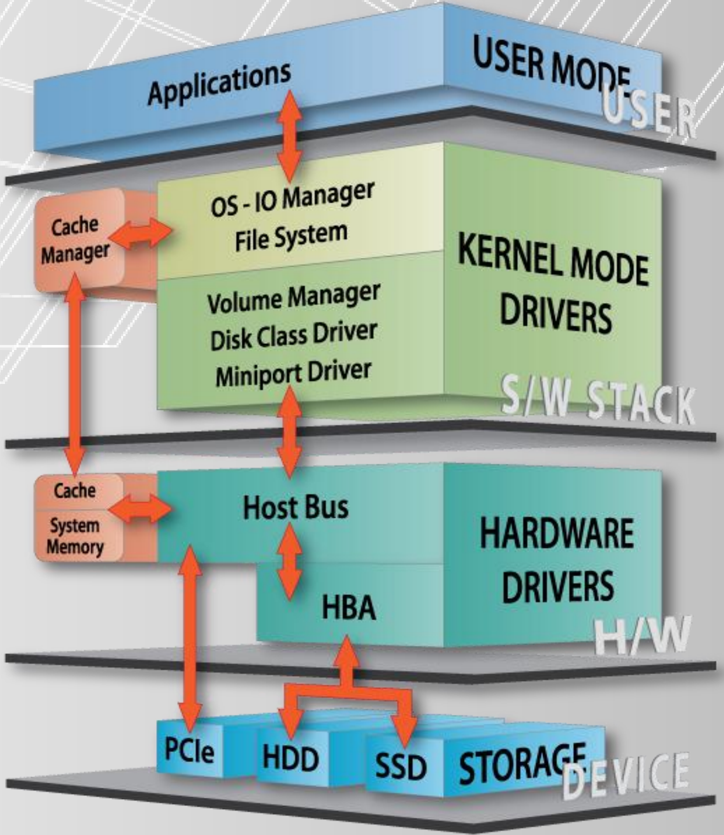

Must view IO performance test results within the context of the IO stack

SNIA SSS PTS and IO Performance Testing

- Currently specifies a synthetic **device level** test
- Prescribes the exact methods and parameters required to be used to generate the IO activity
- Prescribes the conditions under which the test measurements are to be made
- Prescribes the measurements to be reported along with the reporting format

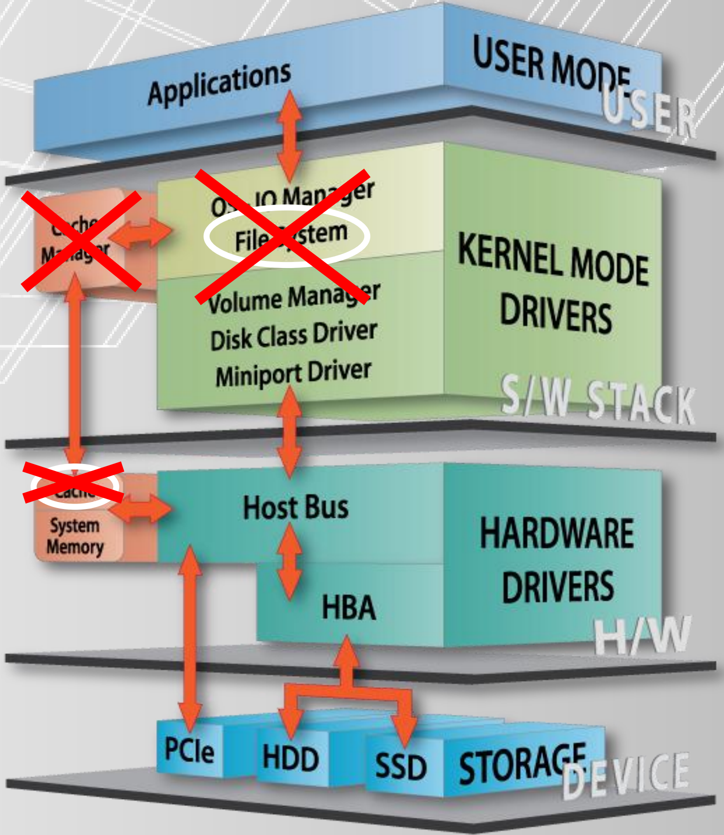

Standardized testing of **innate** SSD performance capabilities

Download this deck at www.snia.org/forums/sssi/pts

1. Brief Background Information about the PTS
2. The IO Software/Hardware Stack
3. Examples of How IOs Traverse the IO SW/HW Stack
4. Impact of Traversing the IO SW/HW Stack
5. **Conclusion**

Conclusion

How IOs traverse the IO SW/HW Stack

- Can involve a variety of different paths, variables, and factors
- Can have a direct impact upon performance
- Is an important consideration for IO performance testing as well as overall application/ workload IO performance

IO Performance Testing Tools and Performance Measurement

- Understand what your test tool actually tests, how it operates, and exactly what/how it measures and reports
- PTS specifies device level testing to focus on **innate** SSD performance capabilities

PTS enables fair, accurate, and repeatable SSD Performance Comparison Testing

SSD Client & Enterprise PTS Comparison Data

Available at

www.snia.org/forums/sssi/pts

SNIA Solid State Storage Initiative