



Flash Memory Summit

Open Source Software for Managing Performance Analysis of Flash Storage

Shing Lee

Technical Director

ADATA Technology Co., Ltd.

shing_lee@adata.com

www.adata.com



Flash Memory Summit

Adata Technology



May, 2001
Founded

2014 No.2
DRAM
module
Manufacturer

2016 No.4
Global SSD
Market
Share



Flash Memory Summit

Milestones of Adata SSD



ADATA®



2018

Joint Development of Industry Standard
Announced NGSFF for Enterprise & Data Center

Ranked No.4
in WW market share
in 2016

SSD ranked No.1
share in Japan in 2012

Start

SSD Business &
Development in 2007

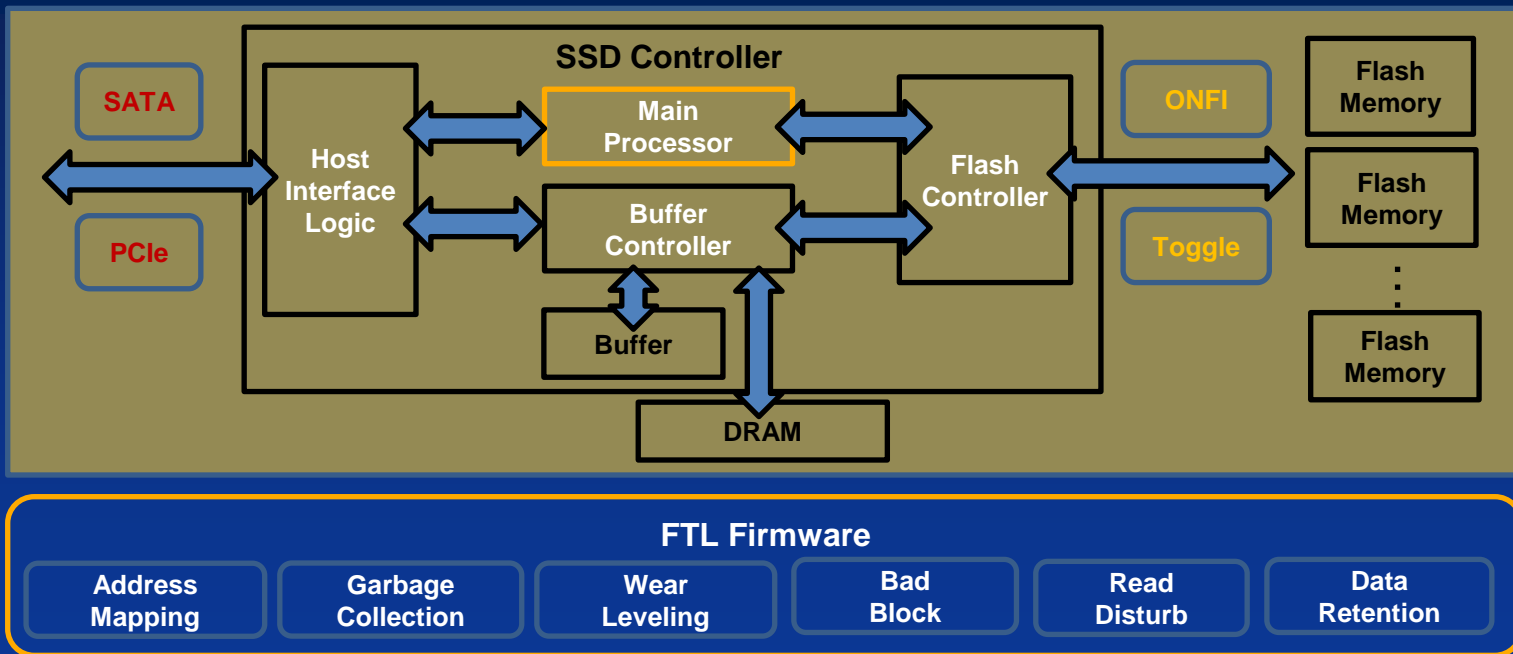
Founded
in 2001





SSD Architecture

The SSD is mainly composed of a main control chip, a flash memory chip, and a firmware algorithm. The performance is very different and affected by the different of main control, the type of the flash memory, and the number of channels, etc.

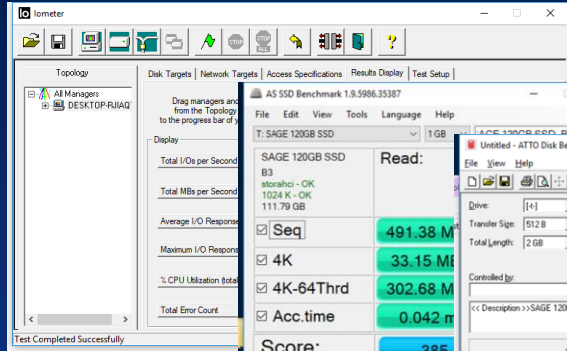




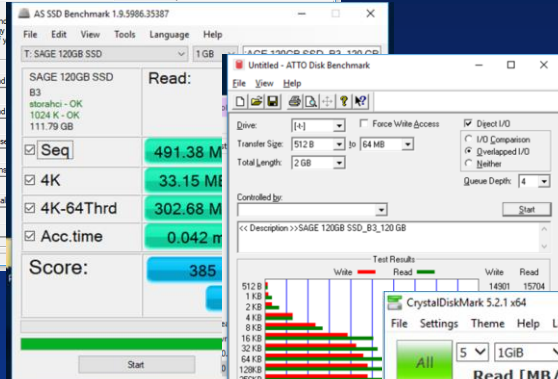
A Variety of SSD Testing Tools

Performance Evaluation Tests

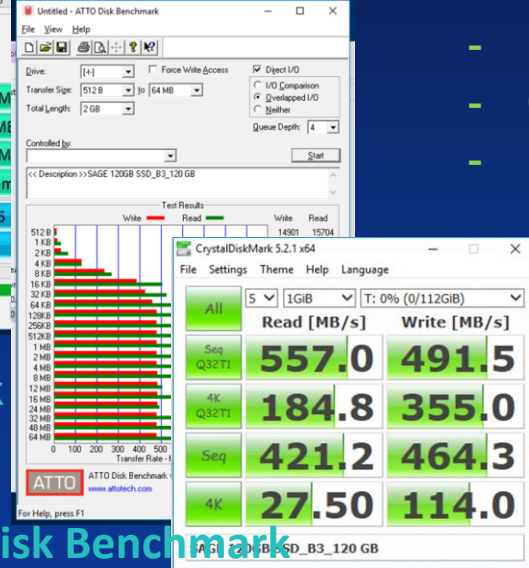
- Sequential Read/Write
- Random Read/Write
- Response time
- Queue Depths
- Read/Write ratio



IOMeter

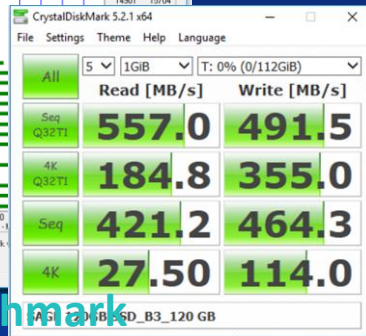


AS SSD Benchmark



ATTO Disk Benchmark

Crystal Disk Mark



The performance is affected by the transmission data structure, test tool settings, and system settings. The software tools sometimes do not have real performance!!



Workload by applications

With different workload, SSD has different performance.

Application	Payload Size	QD	Sequential	Random	Read %	Write %
Web servers	4K	4	25	75	95	5
	8K	4	25	75	95	5
	64K	4	25	75	95	5
Exchange Email	4K	4	0	100	67	33
	64K	4	0	100	67	33
Media Streaming	64K	64	100	0	98	2
File Servers	8K	4	25	75	90	10
Database OLTP	8K	4	0	100	70	30
OS paging	64K	64	100	0	90	10
Video on Demand Ran	512K	64	0	100	100	0
Video on Demand Seq	512K	4	100	0	100	0
Web server logs	64K	64	0	100	100	0
Decision Support system	8K	4	100	0	0	100
SQL logs	64K	64	100	0	0	100
DBMS	8K	4	33	67	100	0
Archive	2M	4	5	95	55	45
Medical Imaging	1M	64	95	5	5	95



Flash Memory Summit

Background Issues

1. How do we make a wise decision with a pool of testing software?
2. The internal methods are unknown to the users as black-box testing is applied.
3. Does the result of the testing software agree with the user experience?



Flash Memory Summit

Our Solution

SNIA Performance Test Specification




Open Source Software Package



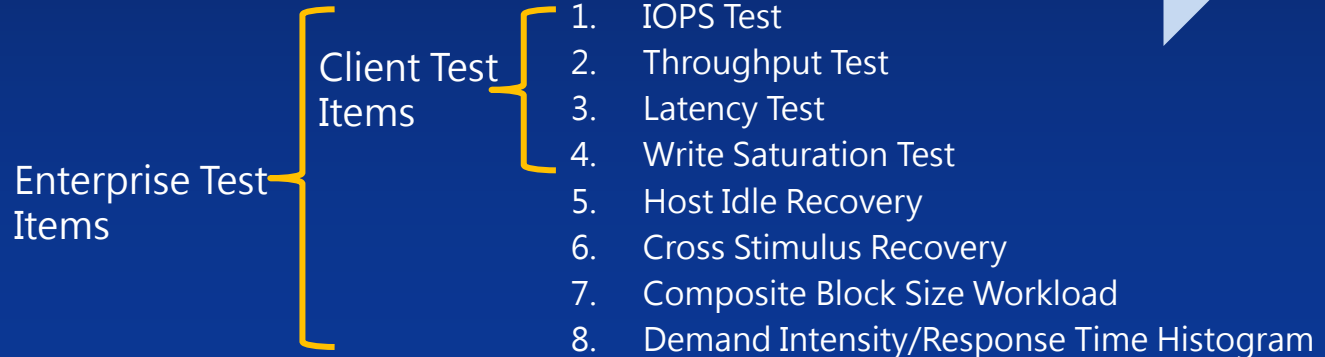
Solid State Storage (SSS) Performance Test Specification (PTS)



PTS, developed by SNIA is a method to **compare the performance specifications of SSDs**. These specifications define a set of device level tests and methodologies which enable comparative testing for Enterprise and Client systems.




www.sina.org





SSD Performance Phases



Peak ⇨ **Decline** ⇨ **Steady**



Open Source Baseline: cloudharmony/block-storage

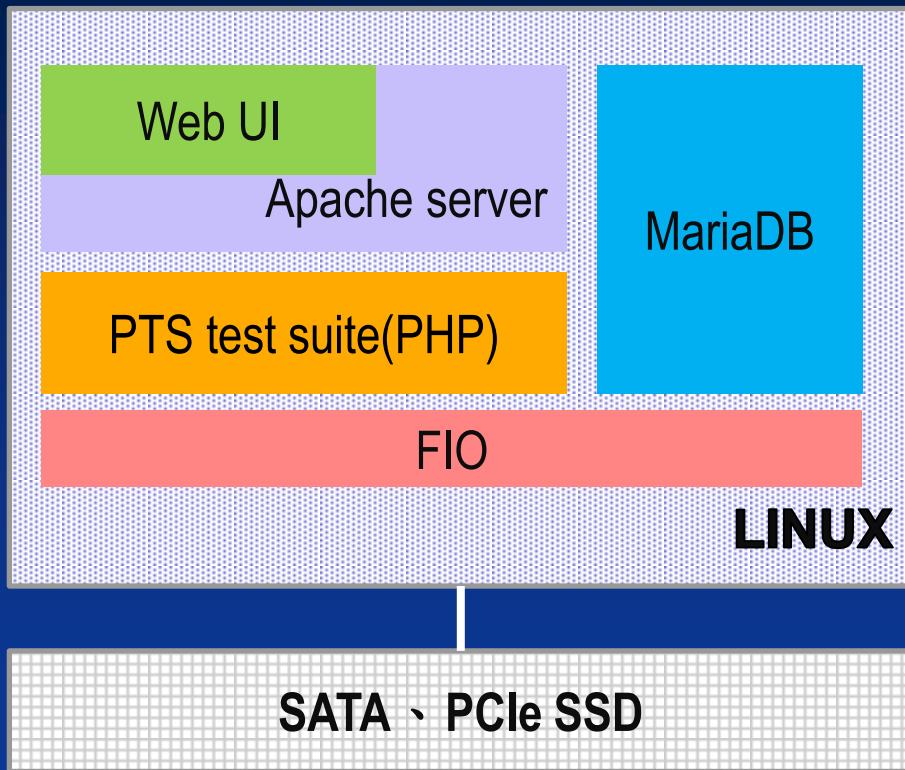
```
root@ADATA: /home/root/pts/test_report
root@ADATA: /home/root/pts/test_report# ./run.sh --threads=4 --oio_per_thread=16 --target=/dev/sdb --test=iops \
> --test=throughput --ss_max_rounds=25 --verbose

02/08/2018 01:57:32 UTC 0s      run:43      Starting block storage tests [iops, throughput] using
targets [/dev/sdb] and 86400s timeout
02/08/2018 01:57:32 UTC 0s      run:50      Starting IOPS block storage test
02/08/2018 01:57:32 UTC 0s      BlockStorageTest:1688  Attempting to purge non-rotational target /dev/sdb wit
h --nosecureerase=1; --notrim=0; --nozerofill=0
02/08/2018 01:57:32 UTC 0s      BlockStorageTest:1701  ATA secure erase not be attempted for /dev/sdb because
--nosecureerase argument was specified (or implied due to lack of --secureerase pswd argument)
02/08/2018 01:57:32 UTC 0s      BlockStorageTest:1707  Attempting TRIM for volume /dev/sdb using command blkd
iscard /dev/sdb >/dev/null 2>&1; echo $?
02/08/2018 01:57:42 UTC 10s     BlockStorageTest:1711  TRIM successful for target /dev/sdb
02/08/2018 01:57:42 UTC 10s     BlockStorageTest:1748  Target /dev/sdb purged successfully using trim
02/08/2018 01:57:42 UTC 10s     run:59      Target purge successful - continuing testing
PHP Notice: Undefined index: precondition time in /home/root/pts/lib/BlockStorageTest.php on line 1932
02/08/2018 01:57:42 UTC 10s     BlockStorageTest:1932  Attempting workload independent preconditioning (2X 12
8k sequential writes on entire device). This may take a while..
PHP Notice: Undefined index: precondition time in /home/root/pts/lib/BlockStorageTest.php on line 1936
02/08/2018 01:57:42 UTC 10s     BlockStorageTest:1940  Attempting workload independent precondition pass 1 of
2
02/08/2018 01:57:42 UTC 10s     BlockStorageTest:344   Starting fio using command: fio --name=global --direct
=1 --ioengine=libaio --refill_buffers --scramble_buffers=1 --blocksize=128k --rw=write --numjobs=1 --iodepth=16 --
group_reporting --output-format=json --size=100% --name=wipc-1-1 --filename=/dev/sdb
```

**Command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers --
scramble_buffers=1 --blocksize=128k --rw=write --numjobs=1 --iodepth=16 --
group_reporting --output-format=json --size=100% --name=wipc-1-1 --filename=/dev/sdb**

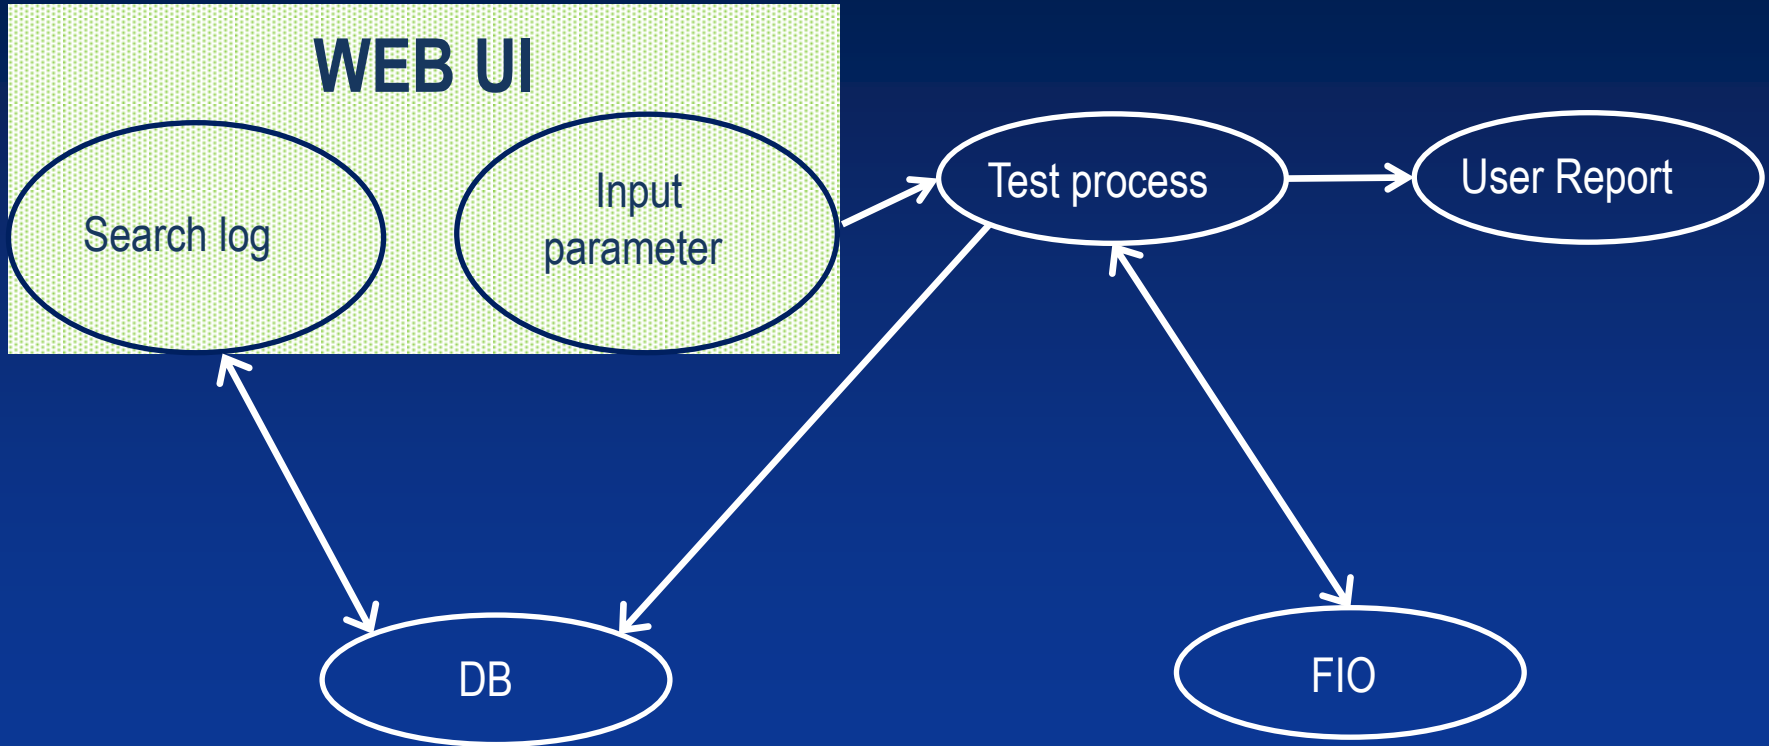


Software Architecture





Data / Command Flow





User Interface Work Flow

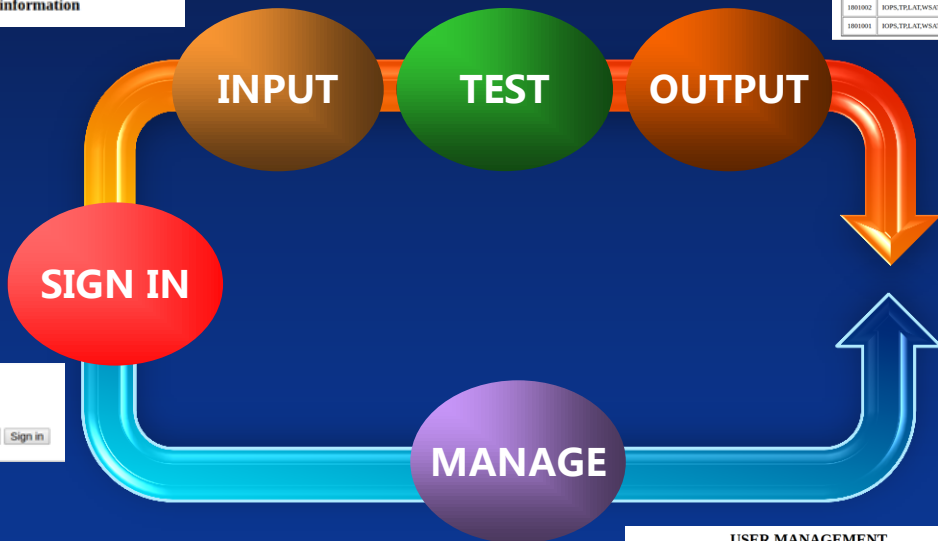
Choose your action

Target : ADATA SP920SS(dev/sdb)

Confirm information

```
Start testing...
02/06/2018 07:39:17 UTC 0s run:43
02/06/2018 07:39:17 UTC 0s run:59
02/06/2018 07:39:17 UTC 0s BlockStorageTest:1975
02/06/2018 07:39:17 UTC 0s BlockStorageTest:1976
02/06/2018 07:39:17 UTC 0s BlockStorageTest:1983
02/06/2018 07:39:17 UTC 0s BlockStorageTest:2035
02/06/2018 07:39:17 UTC 0s run:59
02/06/2018 07:39:17 UTC 0s BlockStorageTest:2219
02/06/2018 07:39:17 UTC 0s BlockStorageTest:2227
02/06/2018 07:39:17 UTC 0s BlockStorageTest:244
02/06/2018 16:04:39 UTC 38322s BlockStorageTest:358
02/06/2018 16:04:39 UTC 38322s BlockStorageTest:2231
02/06/2018 16:04:39 UTC 38322s BlockStorageTest:2227
02/06/2018 16:04:39 UTC 38322s BlockStorageTest:344
02/06/2018 20:07:19 UTC 44882s BlockStorageTest:358
02/06/2018 20:07:19 UTC 44882s BlockStorageTest:2231
```

SN	ITEM	REPORT	INFO
1002003	DBTH	Download	LINK
1002002	DBTH	N/A	LINK
1002001	DBTH	N/A	LINK
1001003	IOPLTFLATWSA1JRH	Download	LINK
1001002	IOPLTFLATWSA1JRH_XS1ECW1DBTH	N/A	LINK
1001001	IOPLTFLATWSA1JRH_XS1ECW1DBTH	N/A	LINK



SIGN IN


ID

USER MANAGEMENT

Total 1 data [Add a new user](#)

ID	NAME	PERMISSION	LAST LOGIN	STATUS	EDIT
1000370	洪志明	Read Write/Delete Administrator	2018-02-06 16:28:27	Enable	Update Delete



**SSS PTS TEST**
Love Life Dream

Test Log

Choose your action

Target : ADATA SP920SS(/dev/sdb)

Confirm information

Vendor
Model No.
Capacity
S/N
Firmware version
Interface
NAND Type

Specification type: Enterprise Client

Test item : iops throughput latency wsat hir xsr cbw dirth

Purge mode: No purge ATA secure erase NVMe format namespace TRIM zero fill

Max test rounds:

TOIO: [iops](#) [throughput](#) [wsat](#) [hir](#) [xsr](#) [by test operator choice](#)

Thread Count: Queue Depth:

WSAT test optional workload: WSAT test optional time period: Hr

DIRTH test parameter [dirth test read write percentage and block size by test operator choice](#)

RW: Read % Block Size:

Show test log: ON OFF



Test



SSS PTS TEST

Test Report

Start testing...


```

02/06/2018 07:39:17 UTC 0s      run:43      Starting block storage tests [dirth] using targets [/dev/sdb] and 86400s timeout
02/06/2018 07:39:17 UTC 0s      run:50      Starting DIRTH block storage test
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:1975 Attempting to purge non-rotational target /dev/sdb with --nosecureerase=0; --notrim=0; --no
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:1978 Attempting ATA secure erase for target /dev/sdb
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:1983 ATA secure erase successful for target /dev/sdb
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:2035 Target /dev/sdb purged successfully using secureerase
02/06/2018 07:39:17 UTC 0s      run:59      Target purge successful - continuing testing
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:2219 Attempting workload independent preconditioning (2X 128k sequential writes on entire device
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:2227 Attempting workload independent preconditioning pass 1 of 2
02/06/2018 07:39:17 UTC 0s      BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 16:04:39 UTC 30322s    BlockStorageTest:358 fio execution successful for step wipc with 64 IOPS (8.05 MB/s). There are now 1 results fo
02/06/2018 16:04:39 UTC 30322s    BlockStorageTest:2231 Workload independent precondition pass 1 of 2 successful
02/06/2018 16:04:39 UTC 30322s    BlockStorageTest:2227 Attempting workload independent preconditioning pass 2 of 2
02/06/2018 16:04:39 UTC 30322s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:07:19 UTC 44882s    BlockStorageTest:358 fio execution successful for step wipc with 134 IOPS (16.77 MB/s). There are now 2 results
02/06/2018 20:07:19 UTC 44882s    BlockStorageTest:2231 Workload independent precondition pass 2 of 2 successful
02/06/2018 20:07:19 UTC 44882s    run:69      Workload independent preconditioning successful - continuing testing
02/06/2018 20:07:19 UTC 44882s    BlockStorageTestDirth:866 Starting workload dependent preconditioning for DIRTH R/W WDPC (1 of 3)
02/06/2018 20:07:19 UTC 44882s    BlockStorageTestDirth:591 Initiating workload dependent preconditioning and steady state for DIRTH test
02/06/2018 20:07:19 UTC 44882s    BlockStorageTestDirth:599 Starting 60 sec DIRTH rand write preconditioning round 1 of 25, test 1 of 30 [name=x1-0_10
02/06/2018 20:07:19 UTC 44882s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:08:19 UTC 44942s    BlockStorageTest:358 fio execution successful for step wdpc with 935 IOPS (7.31 MB/s). There are now 1 results f
02/06/2018 20:08:19 UTC 44942s    BlockStorageTestDirth:607 Test x1-0_100-rand-n01 was successful
02/06/2018 20:08:19 UTC 44942s    BlockStorageTestDirth:617 Added IOPS metric 935 from preconditioning round 1 of 25 for DIRTH steady state verificati
02/06/2018 20:08:19 UTC 44942s    BlockStorageTestDirth:599 Starting 60 sec DIRTH rand write preconditioning round 1 of 25, test 2 of 30 [name=x1-0_10
02/06/2018 20:08:19 UTC 44942s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:09:20 UTC 45003s    BlockStorageTest:358 fio execution successful for step wdpc with 636 IOPS (4.97 MB/s). There are now 2 results f
02/06/2018 20:09:20 UTC 45003s    BlockStorageTestDirth:607 Test x1-0_100-rand-n02 was successful
02/06/2018 20:09:20 UTC 45003s    BlockStorageTestDirth:599 Starting 60 sec DIRTH rand write preconditioning round 1 of 25, test 3 of 30 [name=x1-0_10
02/06/2018 20:09:20 UTC 45003s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:10:22 UTC 45065s    BlockStorageTest:358 fio execution successful for step wdpc with 769 IOPS (6.01 MB/s). There are now 3 results f
02/06/2018 20:10:22 UTC 45065s    BlockStorageTestDirth:607 Test x1-0_100-rand-n03 was successful
02/06/2018 20:10:22 UTC 45065s    BlockStorageTestDirth:599 Starting 60 sec DIRTH rand write preconditioning round 1 of 25, test 4 of 30 [name=x1-0_10
02/06/2018 20:10:22 UTC 45065s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:11:24 UTC 45127s    BlockStorageTest:358 fio execution successful for step wdpc with 601 IOPS (4.7 MB/s). There are now 4 results fo
02/06/2018 20:11:24 UTC 45127s    BlockStorageTestDirth:607 Test x1-0_100-rand-n04 was successful
02/06/2018 20:11:24 UTC 45127s    BlockStorageTestDirth:599 Starting 60 sec DIRTH rand write preconditioning round 1 of 25, test 5 of 30 [name=x1-0_10
02/06/2018 20:11:24 UTC 45127s    BlockStorageTest:344 Starting fio using command: fio --name=global --direct=1 --ioengine=libaio --refill_buffers
02/06/2018 20:12:26 UTC 45189s    BlockStorageTest:358 fio execution successful for step wdpc with 682 IOPS (5.34 MB/s). There are now 5 results fo

```




Report

 **SSS PTS TEST**

Test Report

TEST LOG

Total 6 datas

SN	ITEM	REPORT	INFO
1802003	DIRTH	Download	LINK
1802002	DIRTH	N/A	LINK
1802001	DIRTH	N/A	LINK
1801003	IOPS,TP,LAT,WSAT,HIR	Download	LINK
1801002	IOPS,TP,LAT,WSAT,HIR,XSR,ECW,DIRTH	N/A	LINK
1801001	IOPS,TP,LAT,WSAT,HIR,XSR,ECW,DIRTH	N/A	LINK

©2018 ADATA Technology Co., Ltd. All rights reserved.



SN	1802003
START TIME	2018-02-06 15:39:17
FINISH TIME	2018-02-07 15:18:15
PTS TYPE	enterprise
ITEM	DIRTH
OPERATOR	洪志明
SSD INFO	<ul style="list-style-type: none">• Vendor:ADATA• Model:SP920SS• SN:1452TN000123AQ1• FW ver:1.08• Interface:SATA 3.0, 5.0 Gb• NAND type:MLC
COMMAND	/home/alan/project/block-storage
PDF REPORT	Download
HTML DATA	Download
TEST LOG	Download



Management



If testing is not normal completed, delete the check file.

USER MANAGEMENT

Total 1 datas [Add a new user](#)

ID	NAME	PERMISSION	LAST LOGIN	STATUS	EDIT
10003789	洪志明	Read WriteDelete Administrator	2018-02-06 16:28:27	Enable	Update Delete

©2018 ADATA Technology Co., Ltd. All rights reserved.

ADD USER


ID	<input type="text"/>
NAME	<input type="text"/>
SATAUS	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
PERMISSION	<input type="checkbox"/> Read <input type="checkbox"/> Write <input type="checkbox"/> Delete <input type="checkbox"/> Adn
<input type="button" value="Add"/> <input type="button" value="Reset"/>	

Update Member Data

ID	<input type="text" value="10003789"/>
NAME	<input type="text" value="洪志明"/>
STATUS	<input checked="" type="radio"/> 開啟 <input type="radio"/> 關閉
PERMISSIONS	<input checked="" type="checkbox"/> Read <input checked="" type="checkbox"/> Write <input checked="" type="checkbox"/> Delete <input checked="" type="checkbox"/> Administrator
<input type="button" value="UPDATE"/> <input type="button" value="RESET"/>	

DELETED

NAME	PERMISSION	LAST LOGIN	STATUS
洪志明	Read Write Delete Administrator	2018-02-06 16:57:50	Enable






PTS Report Header

Test Run Date Start: 05/11/2018 08:00 UTC				Test Run Date Finish: 05/14/2018 07:15 UTC				
IOPS Test - Report Page								
SNIA SSS TWG		Solid State Storage Performance Test Spec		IOPS - Block Size x RW Mix Matrix			Rev.	PTS-E 2.0
Vendor: Memblaze		SSD Model: ADATA SR2000CP					Tested By:	ADATA
Test Platform		Device Under Test		Setup Parameters		Test Parameters		
Ref Test Platform	Block Storage	Mfgr	Memblaze	Data Pattern	RND	Data Pattern	RND	
Motherboard	Micro-Star International Co., Ltd X299 SLI PLUS (MS-7A93)	Model No.	SR2000CP	AR	100%	AR & Amount	100%	
CPU	Intel Core(TM) i5-7640X 4.00GHz	S/N	FL175100258	AR Segments	N/A	Test Stimulus 1	IOPS Loop	
Memory	8 GB	Firmware ver	001008K0	Pre Condition 1	SEQ 128K W	RW Mix	Outer Loop	
Operating System	Ubuntu 16.04.4 LTS	Capacity	3726.02 GiB	TOIO - TC/QD	TC 1/QD 16	Block Sizes	Inner Loop	
Test SW	fio	Interface	PCIe 3.0 X8	Duration	2X Device Capacity	TOIO - TC/QD	TC 2/QD 16	
Test SW ver	2.2.10	NAND Type	TLC	Pre Condition 2	IOPS Loop	Steady State	3 - 7	
Test SN.	1805001	PCIe NVM		TOIO - TC/QD	TC 2/QD 16	Test Stimulus 2	N/A	
HBA	N/A	Purge Method	FORMAT namespace	SS Rouds	3 - 7	TOIO - TC/QD	N/A	
PCIe	3.0 x16	Write Cache	Disable	Notes		Steady State	N/A	
IOPS Steady State Convergence Plot - All Block Sizes								



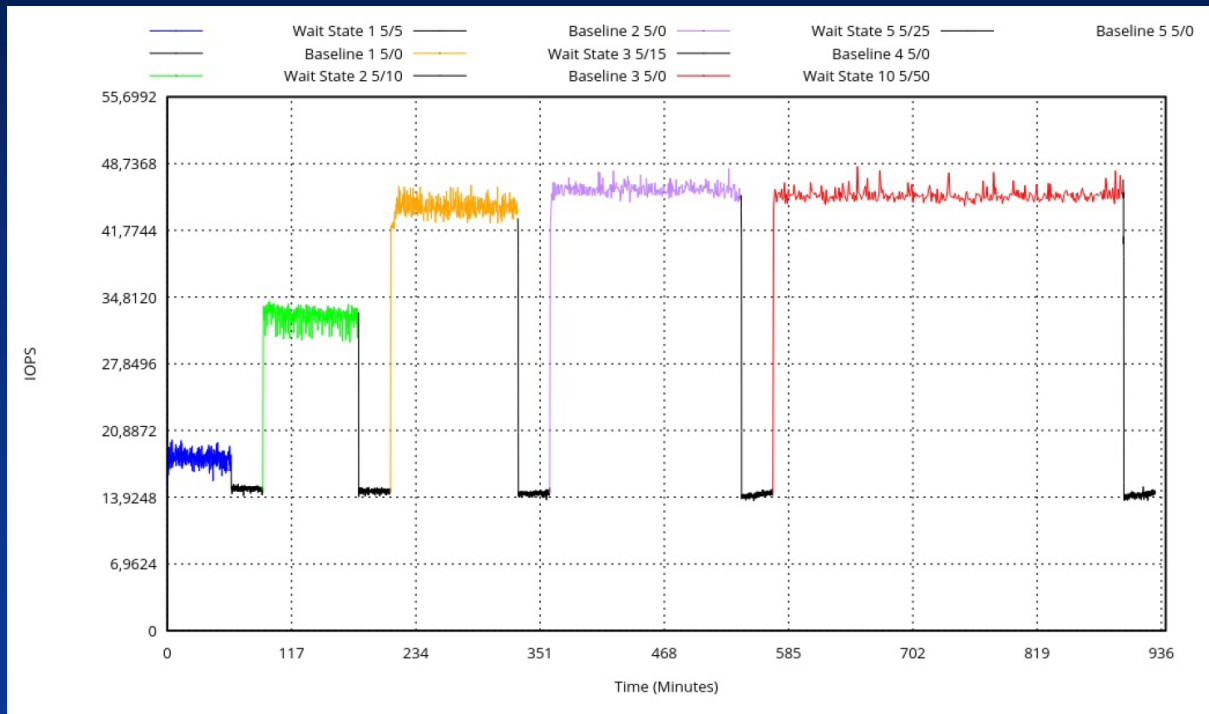


IOPS Test : All RW Mix & BS – 2D Plot



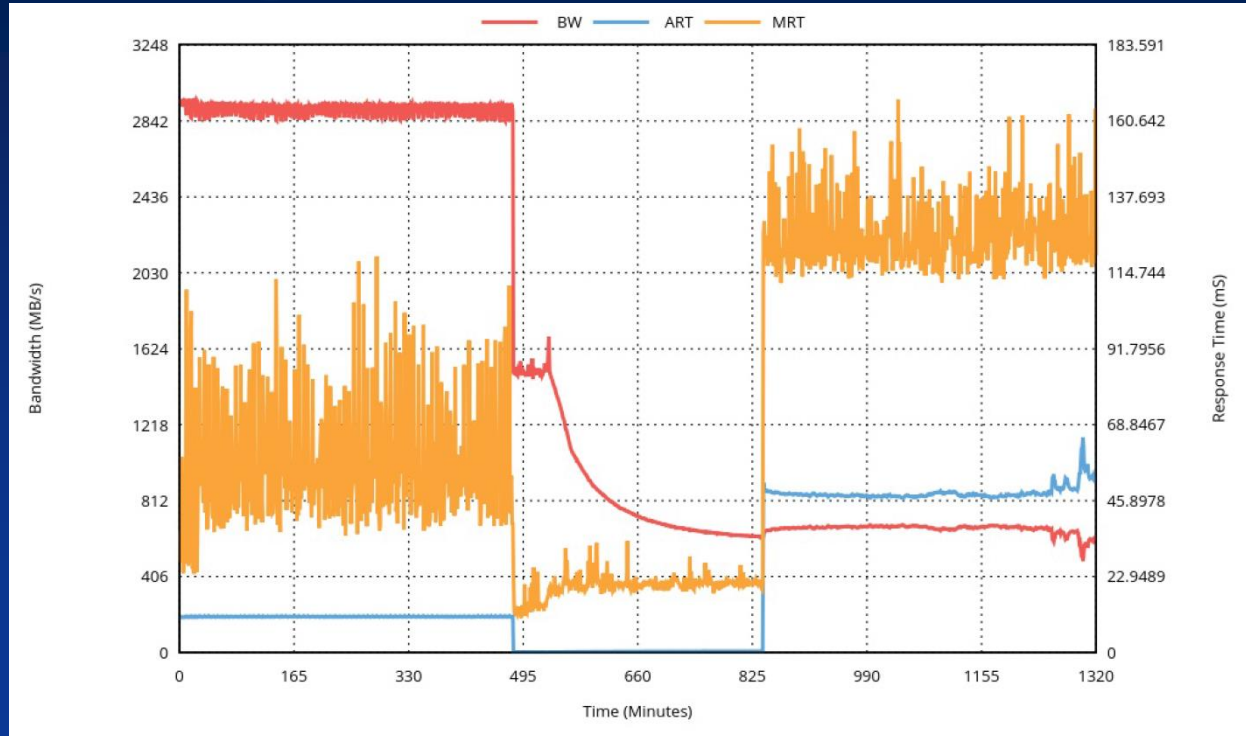


Host Idle Recovery Test : IOPS vs. Time – All Wait States






Cross Stimulus Recovery Test : Throughput, ART & MRT vs. Time





Composite Block Size Workload Test: CBWT Demand Variation Plot





Flash Memory Summit

Open Source Availability

The work has been made open-source on Github and is readily available to all performance testers.

Commit Type	Description	Time Ago
Alan-ADATA update to PTS 2.0		Latest commit 2d8bcf2 3 days ago
database	update to PTS 2.0	3 days ago
test	update to PTS 2.0	3 days ago
web	update to PTS 2.0	3 days ago
.gitignore	first commit and change to ADATA's version	2 months ago
INSTALL	update to PTS 2.0	3 days ago
LICENSE	first commit and change to ADATA's version	2 months ago
README.md	update to PTS 2.0	3 days ago

<https://github.com/Alan-ADATA/SSS-PTS-TEST>



Conclusion

- Full compliance with SNIA PTS specifications
- Auto-generation of complete test reports
- Graphical interface instead of command line interface
- Excellence in database management
- Available as an open resource on GitHub



Flash Memory Summit

Thank You!

**Questions & Follow-up:
shing_lee@adata.com
www.adata.com**