

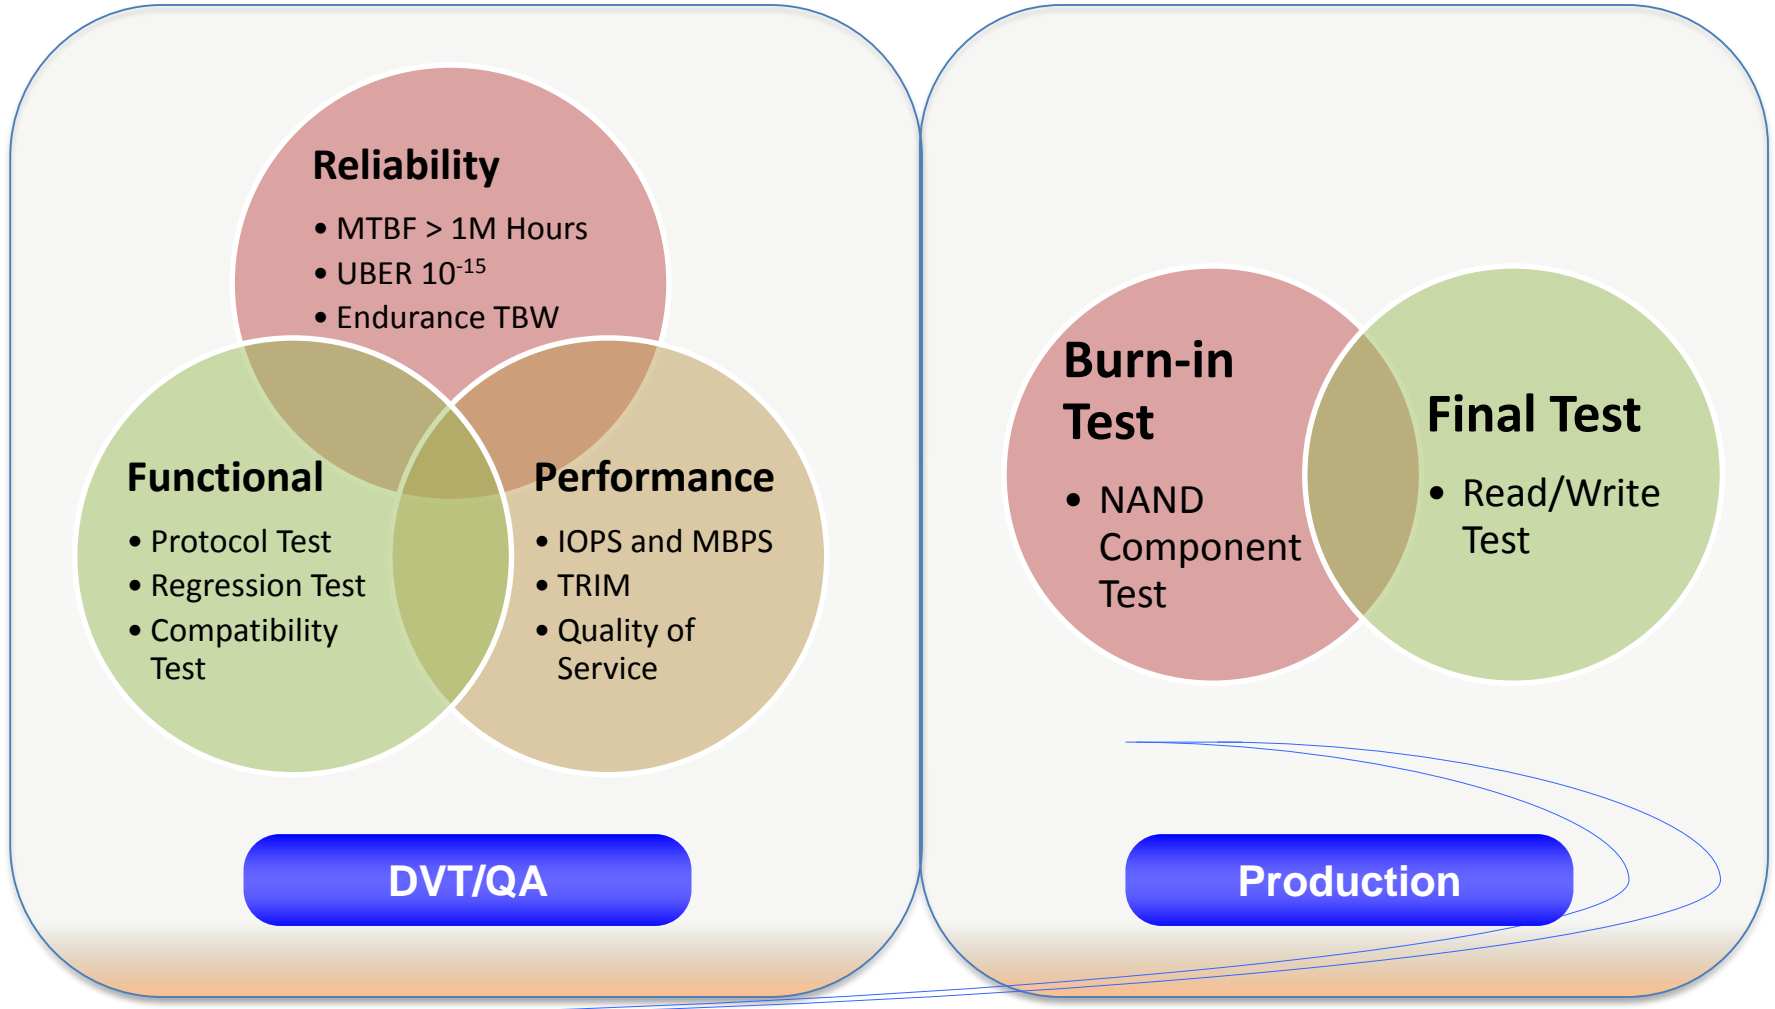


# Creating Effective SSD Test Suites

Joseph Chen, ULINK Technology, Inc.

August 12<sup>th</sup>, 2013

# SSD Testing Overview



# SSD Testing Matrix

	EVT/DVT	RDT	Production
Functional	Protocol Test Electrical/Power Test EMI/Agency/Logo Tes		Final Test
Reliability	Data Compare Test POR/Power Cycle Error Handling Test	MTBF Test Endurance Test/TBW UBER Test	
Performance	IOPS MBPS Test TRIM Performance Test Quality of Service		
Compatibility	System Install/Boot Test Application SW Test System Operation Test		
Burn-in			NAND Component Scan

# Protocol Test Summary

DUT	Sum	PCT1	PCT2	PCT3	PCT4	PCT5	PCT6	PCT7	PCT8	PCT9	PCT10	PCT11	PCT12	PCT13	PCT14	PCT15	PCT16	PCT17	PCT18	PCT19	PCT20
	Errors	SEN	IDF	SFS	MAN	WCF	RWB	ATA	DID	DCS	SMS	SCT	PWM	PWS	IPM	SST	SSP	PHY	DSM	NCQ	SEC
H1	10																		NS	5	
H2	10																		NS		
H3	10									1		3					1		NS	5	
H4	10									1		3					1		NS	5	
H5	10							3				3							NS	1	
H6	3																		NS		
H7	3																		NS		
H8																			NS		
S1	55						6	3	NS	NS		NS	4	12	NS	NS	NS	NS	6	5	19
S2	22							1			3	3							3	1	11
S3									1		2	NS								NS	
S4	BRK							BRK	NT												
S5	7											2									5
S6												2									5
S7	18	NS						3			3				NS				7	5	NS
S8	18	NS						3			3				NS				7	5	NS
S9										NS	NS										5
S10	5									NS	NS										5
S11	7									NS	NS								2		5
S12										NS	NS								2		5
S13	2									NS	NS										2
S14	2									NS	NS										2
S15	18							1				3									7
S16	18							1				3									7
S17	18							1				3									7
S18								1					1						1	2	
		0/26	0/26	0/26	0/26	0/26	1/26	10/26	1/18	2/18	4/25	9/23	3/25	1/25	0/22	0/24	2/24	0/24	7/25	17/24	5/23

Test	Abb	Name
PTC1	SEN	SecurityEraseNormal
PTC2	IDF	IdfyInfo_SATA
PTC3	SFS	SetFeature_SATA
PTC4	MAN	MandatoryCmds
PTC5	WCF	WrCacheFlushTime
PTC6	RWB	RdWrBoundaryCk
PTC7	ATA	ATACmds
PTC8	DID	DCO_IdfyInfo
PTC9	DCS	DCO_SATA
PTC10	SMS	SmartSet
PTC11	SCT	SCT
PTC12	PWM	PwrMgt
PTC13	PWS	PowerState
PTC14	IPM	IPM_Cmplt
PTC15	SST	SSPState
PTC16	SSP	SSP
PTC17	PHY	PhyEvtCnt
PTC18	DSM	ATA_DSM
PTC19	NCQ	NCQ
PTC20	SEC	SecuritySet

Legend	NT	Not Tested
	BRK	Broken
	NS	Not Supported

- Total Fail Rate: 14% (64/480)
- Top Fail Tests: NCQ/SCT/ATA/SEC
- Average Fail Items: 2.4 (64/26)
- Average HDD Fail Count: 4.8 (39/8)
- Average SSD Fail Count: 11.1 (201/18)
- Broken: 1 SSD

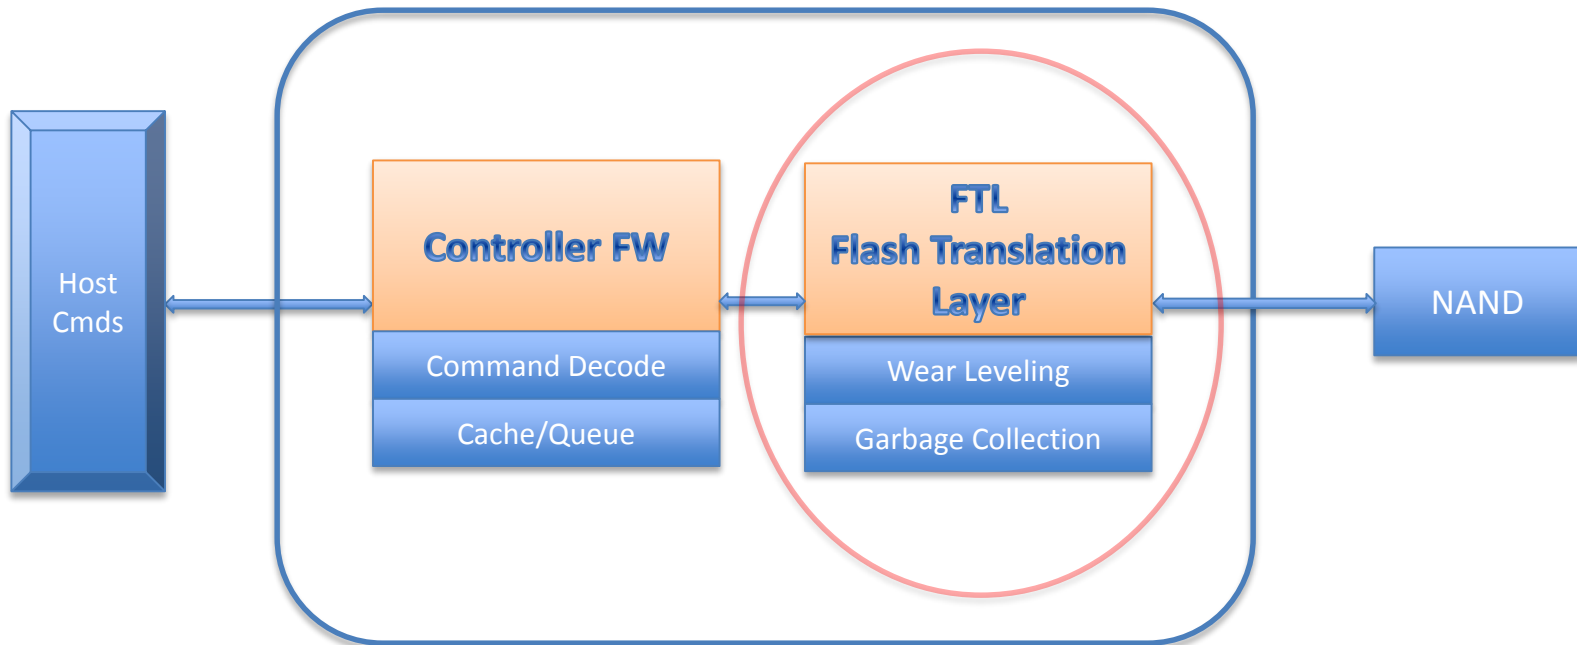
	HDD	SSD
Good	4	0
Bad	4	17
Ugly	0	1

# SSD Reliability Test

Should SSDs and HDDs be tested the same way?

Test	Spec	HDD	SSD
MTBF	1M Hours	Mechanical - Spindle Head/Media	NA
Load/Unload	60K	Mechanical - Spindle Head Stiction	NA
Power Cycle	50K	Spindle Motor Ramp/Latch/Park	NA
Endurance	TBW	NA	NAND
Power Interrupt	5K	NA	FTL

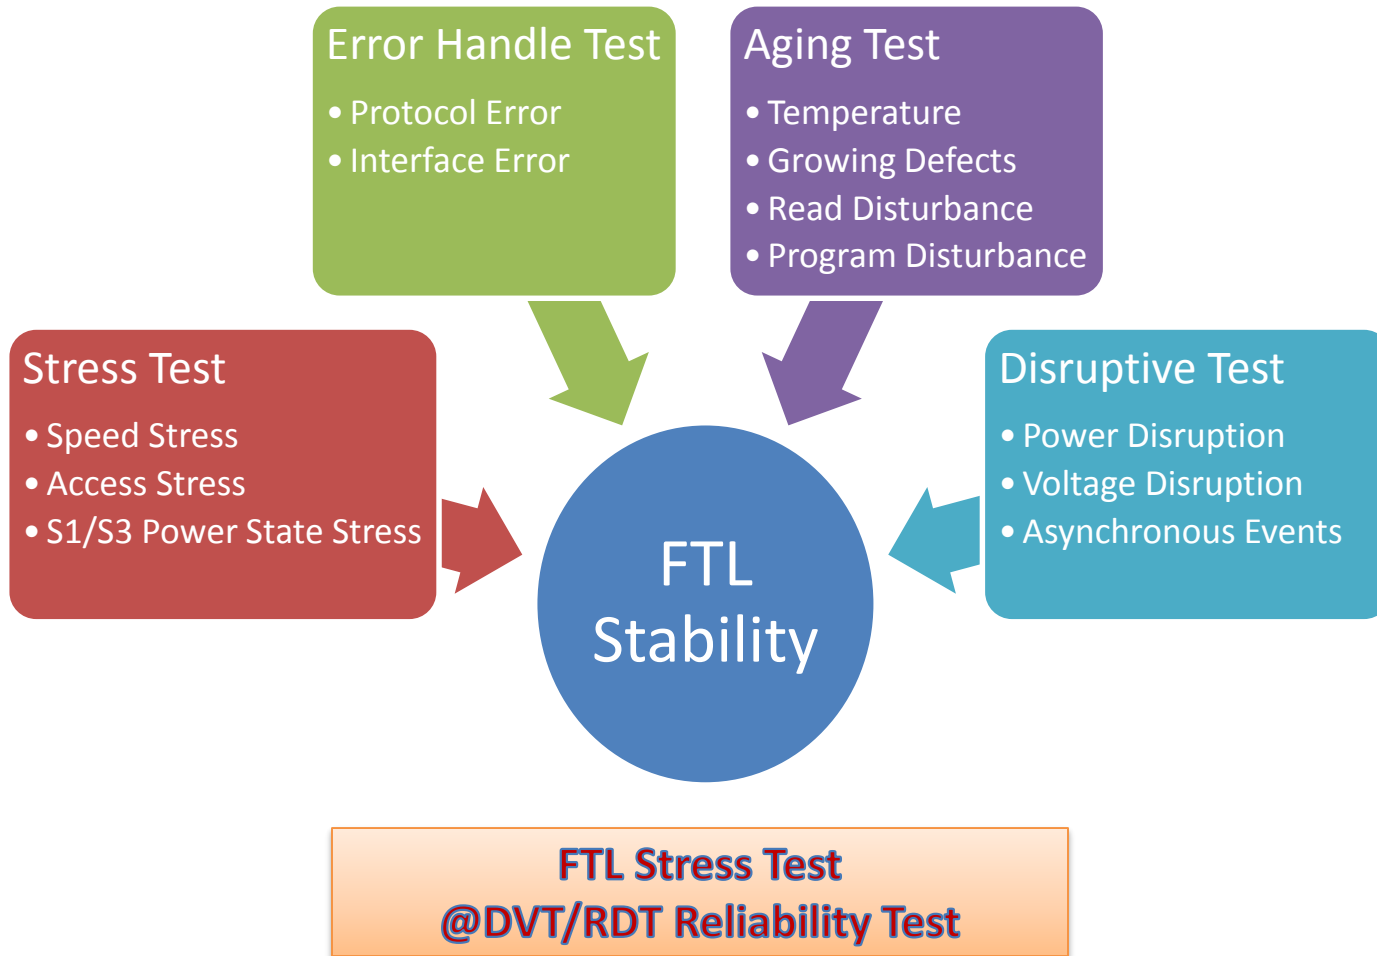
# FTL in SSD FW Block Diagram



# FTL LBA Mapping

HDD Zone Map	SSD FTL Map
<p><b>Fixed Zone Map</b> Map on cylinder/head/sector Map created on factory Burn-in</p>	<p><b>Variable Mapping</b> High degree of variance Complicate meta data management</p>
<p><b>Static Mapping</b> Map does not change on host command</p>	<p><b>Dynamic Mapping</b> Map changed on host command</p>
<p><b>No Update</b> Map does not change on Idle</p>	<p><b>Continuous Update Mapping</b> Map changed on background tasks (Wear Leveling/Garbage Collection)</p>
<p><b>Slow Mapping</b> Mechanical seek/latency (~200 IOPS, 5ms)</p>	<p><b>Fast/Ultra Fast Mapping</b> High IOPS (~50K, 20us) 0.4% background mapping</p>

# How to Test FTL?





# Power Interruption Test (Unintended Shutdown Test)

Prepare  
with  
known  
data

Power off  
device

Power on  
the  
device

Verify  
device  
ready  
timing

Verify  
read and  
write  
function

Compare  
data

# Power Interrupt Test Summary

	Capacity	Dead Device	Power Up Ready Timeout	Read/Write Error	Good Data Compare Error (Shorn Write)	Prep Data Compare Error (Flying Write)
SSD1	128 GB	N	0	0	81	0
SSD2	400 GB	N	0	1	0	0
SSD3	128 GB	N	0	0	108	58
SSD4	128 GB	Y	BRK	BRK	BRK	BRK
SSD5	128 GB	N	0	0	0	0
SSD6	120 GB	N	10	0	0	0
SSD7	480 GB	N	0	0	0	0
SSD8	480 GB	N	0	0	0	0

Test for 100 Hours / 100 Loops



# Power Interruption Test Log

Loop: 1, POR= 0.234 Sec, NumWrDmaCmd= 106, GoodDataMiscmp= 0/100, PwOffDataMiscmp(m): 6/006 [mmmmmm]  
Prep Data Miscompare Count=0

Loop: 2, POR= 0.250 Sec, NumWrDmaCmd= 103, GoodDataMiscmp= 0/100, PwOffDataMiscmp(m): 3/003 [mmm]  
Prep Data Miscompare Count=0

Loop: 3, POR= 0.249 Sec, NumWrDmaCmd= 107, GoodDataMiscmp= 0/100, PwOffDataMiscmp(m): 7/007 [mmmmmmm]  
Prep Data Miscompare Count=0

Loop: 4, POR= 0.234 Sec, NumWrDmaCmd= 103, GoodDataMiscmp= 0/100, PwOffDataMiscmp(m): 3/003 [mmm]  
Prep Data Miscompare Count=0

Loop: 5, POR= 0.234 Sec, NumWrDmaCmd= 108, ERROR: Buffer Miscompare

READ Buffer: LBA= 122547520 (0x74DED40) Offset: 0x8

00000 40 ED 4D 07 00 00 00 00 42 A0 59 12 42 A0 59 12 @.M.....B.Y.B.Y.

00010 42 A0 59 12 42 A0 59 12 42 A0 59 12 42 A0 59 12 B.Y.B.Y.B.Y.B.Y.

00020 42 A0 59 12 42 A0 59 12 42 A0 59 12 42 A0 59 12 B.Y.B.Y.B.Y.B.Y.

00030 42 A0 59 12 42 A0 59 12 42 A0 59 12 42 A0 59 12 B.Y.B.Y.B.Y.B.Y.

WRITE Buffer: LBA= 23 (0x17) Offset: 0x8

00000 40 ED 4D 07 00 00 00 00 4D 1B 60 F5 4D 1B 60 F5 @.M.....M.`.M.`.

00010 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 M.`.M.`.M.`.M.`.

00020 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 M.`.M.`.M.`.M.`.

00030 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 4D 1B 60 F5 M.`.M.`.M.`.M.`.

**MISCOMPARE ERROR:** Good Data Miscompare Error: Command#: 83, Error LBA: 074DED40h, Code:1003, Pointer: 8h

# SSD Endurance Spec Variations

Name	Descriptions	Workload Specified
<b>DWPD</b> (Drive Write Per Day)	<ul style="list-style-type: none"> <li>Ten full drive write per day, 5 year warranty (Seagate)</li> <li>Ten drive write per day for 5 years (Intel/Seagate)</li> </ul>	No
<b>GBPD</b> (Giga Byte Per Day)	<ul style="list-style-type: none"> <li>Minimum of three years of useful life under typical client 3 years workloads with up to 20 GB of host writes per day (Intel)</li> <li>20GB/day of host writes for 5 years under typical client workloads (OCZ)</li> </ul>	Typical Client Workload
<b>TBW</b> (Tera Byte Write)	<ul style="list-style-type: none"> <li>TBW: Total bytes written (Micron/Kingston/WD)</li> <li>TBW: Tera byte write (Sandisk)</li> </ul>	No
<b>PBW</b> (Peta Byte Write)	<ul style="list-style-type: none"> <li>Lifetime endurance (8KB): Up to 14 PB (Intel)</li> <li>4 kilobytes (KB) write endurance of up to 1.5 petabytes (PB) with 20 percent over-provisioning (Intel)</li> </ul>	Yes, 4KB/8KB
<b>JEDEC 218/219</b>	<ul style="list-style-type: none"> <li>JEDEC TBW: JEDEC 218/219 TBW rating with client/enterprise workload</li> </ul>	Yes

Note: Information quoted from the public specifications of the list companies for the purpose of illustration

# Endurance Test Workload

□ Endurance TBW definition issue:

- ❖ Not a real world workload
- ❖ TBW is based on **WAF**, WAF is based **Workload**

$$TBW = \frac{(SSD\ Capacity * NAND\ Cycles) * (1 + OP)}{2 * WAF[Wkld]}$$

*TBW = Tera Byte Write (Endurance Rating)*

$$OP = Overprovision = \frac{Physical\ Capacity}{Logical\ Capacity} - 1$$

*WAF = Write Amplification Factor*

*2 = Guard Band for wear leveling effect*

□ Solution: Use JEDEC Workload for Endurance Tests

# Single Port Endurance Test

	IOPS (example)	Time/ JEDEC 218 Client	TBW/ JEDEC 218 Client	TBW / Day	Days / 72 TBW
<b>1 Port</b>	<b>12K</b>	<b>3 Hr</b>	<b>0.8</b>	<b>6.4</b>	<b>11.5</b>
<b>2 Ports</b>	<b>6K</b>	<b>6 Hr</b>	<b>0.8</b>	<b>3.2</b>	<b>22.5</b>
<b>4 Ports</b>	<b>3K</b>	<b>12 Hr</b>	<b>0.8</b>	<b>1.6</b>	<b>45.0</b>

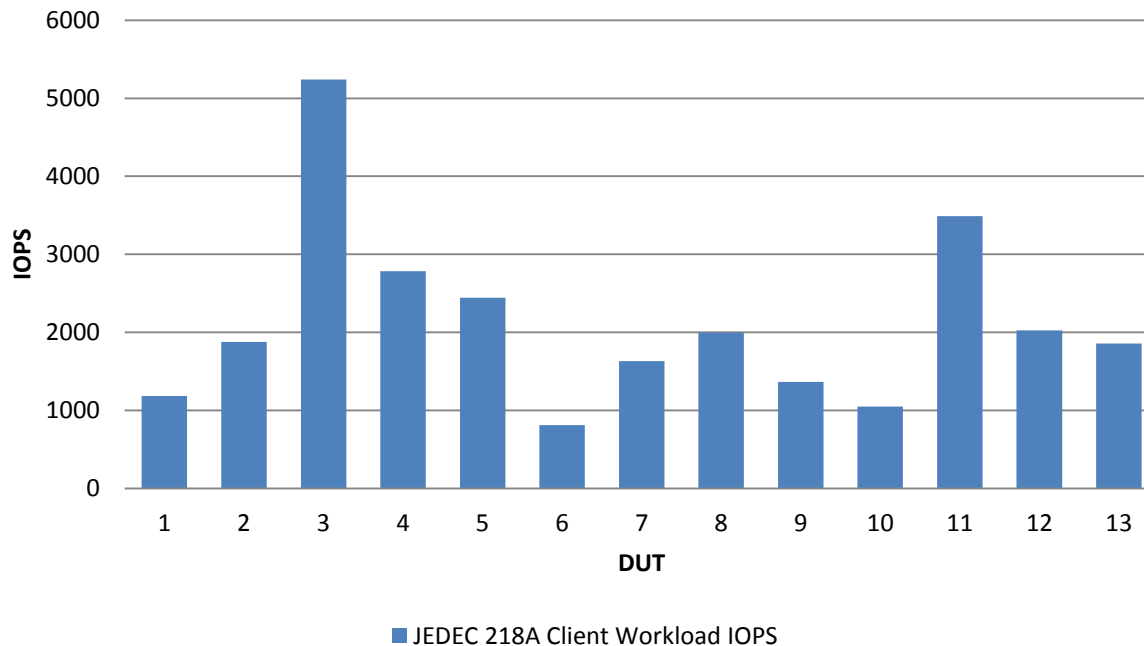
Single port test system significantly reduce test time

# Four Corner Performance Test?

- ❑ Four Corner: 4K Random R/W and Sequential R/W
- ❑ SNIA Standardized SSD Performance
  
- ❑ Issues – “Unreal” test condition
  - ❖ Not a real world workload
  - ❖ No TRIM performance tested
  
- ❑ Solution – Use JEDEC 218 Client Workload
  - ❖ Real world workload
  - ❖ TRIM commands supported

# Performance Test Summary

JEDEC Wkld Performance




DUT	Capacity	IOPS
S1	128 GB	1186
S2	256 GB	1879
S3	100 GB	5241
S4	160 GB	2783
S5	80 GB	2442
S6	80 GB	813
S7	120 GB	1632
S8	128 GB	1997
S9	250 GB	1363
S10	250 GB	1050
S11	480 GB	3489
S12	240 GB	2023
S13	120 GB	1856

Tested with JEDEC Client Workload with 38M Commands Issued



# Trim Performance Comparison

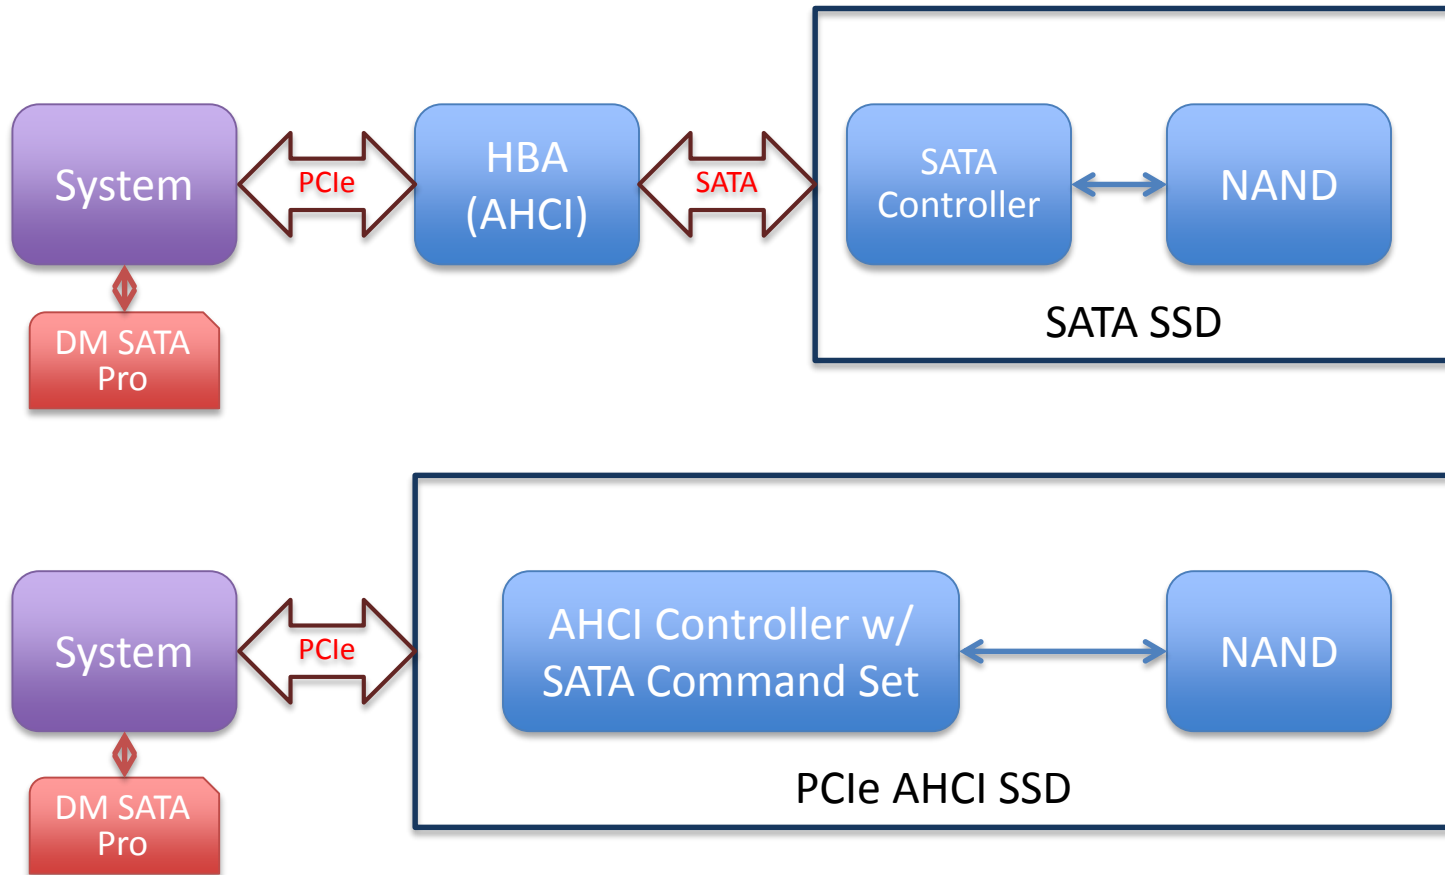
Trim Performance Comparison



DUT	GB	IOPS no Trim	IOPS Trim	% Diff
S1	128 GB	1186	772	65.09%
S2	256 GB	1879	667	35.50%
S3	100 GB	5241	4615	88.06%
S4	160 GB	2783	3368	121.02%
S5	80 GB	2442	2944	120.56%
S6	80 GB	813	1559	191.76%
S7	120 GB	1632	2430	148.90%
S8	128 GB	1997	2542	127.29%
S9	250 GB	1363	1185	86.94%
S10	250 GB	1050	1132	107.81%
S11	480 GB	3489	2490	71.37%
S12	240 GB	2023	1753	86.65%
S13	120 GB	1856	1644	88.58%

Tested with JEDEC Client Workload with 38M Commands Issued

# PCIe AHCI SSD Testing



DriveMaster SATA Pro is compatible with both SATA SSD and PCIe AHCI SSD

# PCIe AHCI Test Tool

## AHCI Compliance Test Software

- Provide utility to walk through every command and protocol.

## AHCI Hardware and Protocol Debug Tool

- Tool help debug and find the problem.

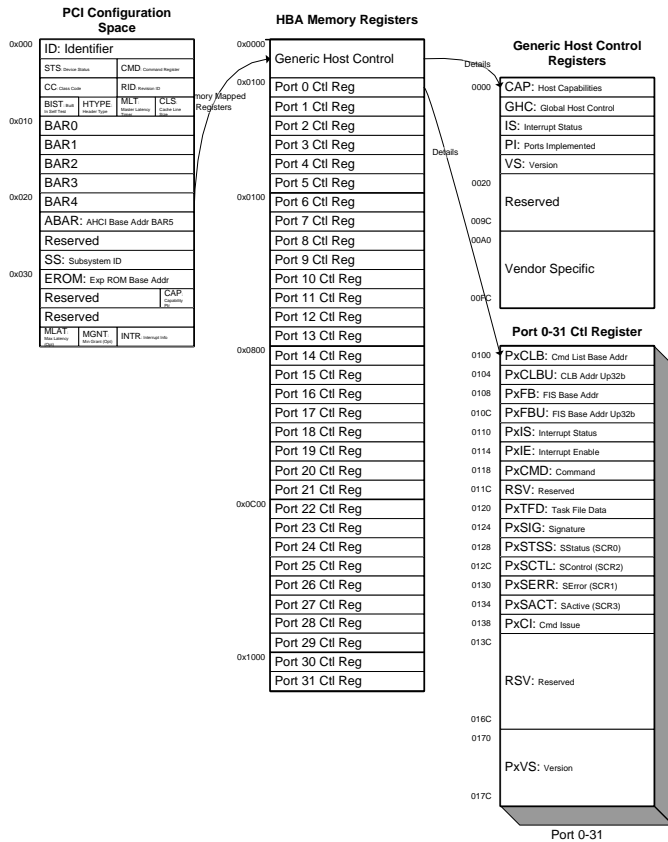
## AHCI Traffic Generator for Device Interface Test

- Issuing commands with different combinations of commands, corner cases, and error conditions. Decipher all of the complicated traffic and review the result.

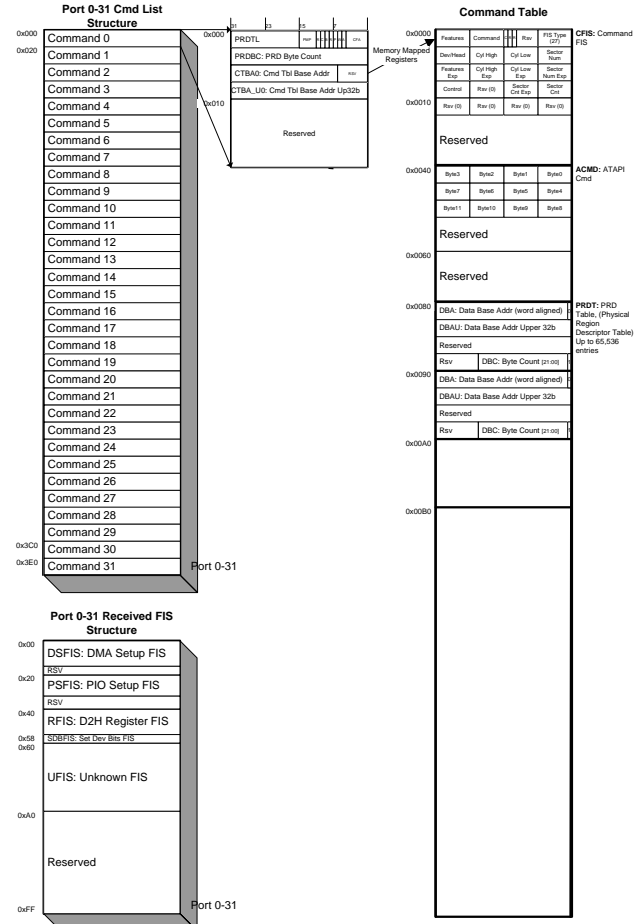
DriveMaster provides consistent, repeatable, automated, and comprehensive PCIe AHCI test solution.

# PCIe AHCI Register Structure

## AHCI HBA Registers




## System Memory

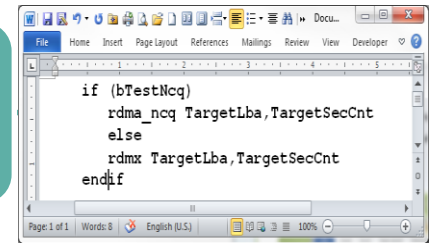


# PCIe AHCI Testing Support


Interactive AHCI Register Window



Manual Command Activation



Ball-of-Wax AHCI Validation



# Summary

## ❑ Creating Effective SSD Test Suites

### ❖ Functional

- Compliance Test (SATA-IO)
- Protocol Test

### ❖ Reliability

- Regression Test for FTL
- JEDEC Workload Endurance Test

### ❖ Performance

- Benchmark JEDEC Real-Life Workload IOPS
- Benchmark TRIM Performance



## ❑ ULINK SSD Testing Suites fulfil the need!

Thank  
You

**THANK YOU!**

ULINK Technology Inc.  
3120 De La Cruz Blvd., #117  
Santa Clara, CA 95054

Tel: 408.446.8455  
Fax: 775.796.8472  
Email: [contact@ulinktech.com](mailto:contact@ulinktech.com)

ULINK Technology, Inc.  
3120 De La Cruz Blvd, Ste 117  
Santa Clara, CA 95054  
1-408-446-8455

[www.ulinktech.com](http://www.ulinktech.com)

Contact (at) [ulinktech.com](mailto:ulinktech.com)