Testing Dual-Port NVMe SSD

Sneha Nadig, Product Application Engineer
Advantest America, Inc.

Contents

- What is Dual-Port NVMe?
- Dual-Port Functionality in NVMe
- Form Factors for Dual-Port NVMe
- Advantages of Dual-Port NVMe
- NVMe Dual-Port vs. SAS Dual-Port
- Testing Dual-Port NVMe

What is Dual-Port NVMe?

- Traditionally single x4 device is split into two x2 devices
  - Port A and Port B
- Methods to access the device:
  - Either Port A or Port B
  - Both Ports simultaneously
- Supported form factors:
  - U.2 (SFF-8639)

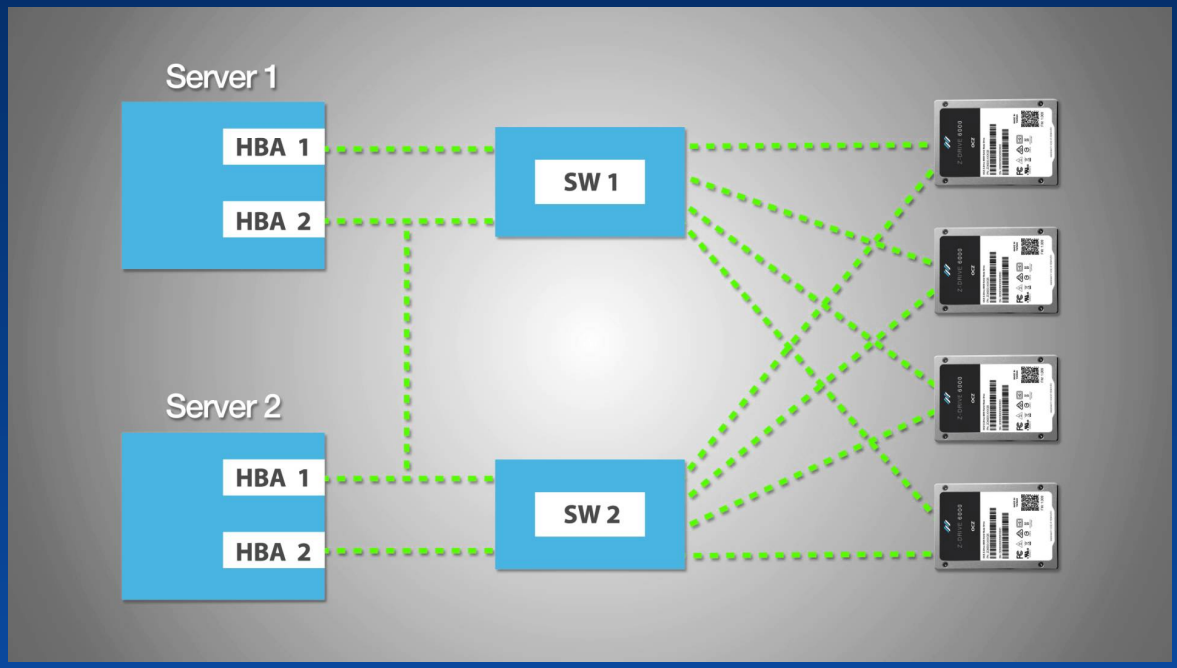

Image Source: OCZ Storage Solutions

Enterprise PCIe Pin Out

Drive	Usage	Signal Description	Name	Mating	Pin #
		Ground	GND	2nd	S1
input	SAS+SATA	SAS/SATA/SATAe 0 Tx+	S0T+ (A+)	3rd	S2
input	SAS+SATA	SAS/SATA/SATAe 0 Tx -	S0T- (A-)	3rd	S3
		Ground	GND	2nd	S4
output	SAS+SATA	SAS/SATA/SATAe 0 Rcv -	S0R- (B-)	3rd	S5
output	SAS+SATA	SAS/SATA/SATAe 0 Rcv +	S0R+ (B+)	3rd	S6
		Ground	GND	2nd	S7
input	Dual Port	ePCIe RefClk + (port B)	RefClk1+	3rd	E1
input	Dual Port	ePCIe RefClk – (port B)	RefClk1-	3rd	E2
input	ePCIe opt	3.3V for SM bus	3.3Vaux	3rd	E3
input	Dual Port	ePCIe Reset (port B)	ePERst1#	3rd	E4
input	ePCIe	ePCIe Reset (port A)	ePERst0#	3rd	E5
		Reserved	RSVD	3rd	E6
input	SATAe +SAS4	Reserved(WAKE#/OBFF), SASAct2	RSVD(Wake#) /SASAct2	3rd	P1
Bi-Dir	SATAe	SATAe Client /SAS reset	sPCIeRst/SAS	3rd	P2
input	SATAe	Reserved (DevSLP#)	RSVD(DevSLP#)	2nd	P3
output	SATAe + ePCIe	Interface Detect (Was GND-precharge)	IfDet#	1st	P4
	all	Ground	GND	2nd	P5
	all			2nd	P6
NC	SAS+SATA	Precharge SATA, SATAe, SAS only	5 V	2nd	P7
NC	SAS+SATA			3rd	P8
NC	SAS+SATA			3rd	P9
	all	Presence (Drive type)	PRSNT#	2nd	P10
Bi-Dir	all	Activity(output)/Spinup	Activity	3rd	P11
	all	Hot Plug Ground	GND	1st	P12
input	all	Precharge All – 12V Only power for ePCIe SSD	12 V	2nd	P13
input	all			3rd	P14
input	all			3rd	P15

ePCIe → Enterprise PCIe (separate from SATA/SAS)

SATAe → SATA Express (Client PCIe- muxed on SATA/SAS signals)

SAS4 → SAS x4

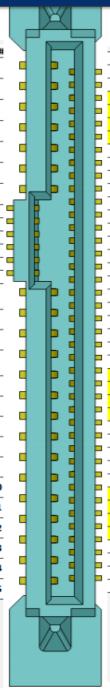

Pin #	Mating	Name	Signal Description	Usage	Drive
E7	3rd	RefClk0+	ePCIe Primary RefClk +	ePCIe	input
E8	3rd	RefClk0-	ePCIe Primary RefClk -	ePCIe	input
E9	2nd	GND	Ground		
E10	3rd	PETp0	ePCIe 0 Transmit +	ePCIe	input
E11	3rd	PETn0	ePCIe 0 Transmit -	ePCIe	input
E12	2nd	GND	Ground		
E13	3rd	PERn0	ePCIe 0 Receive -	ePCIe	output
E14	3rd	PERp0	ePCIe 0 Receive +	ePCIe	output
E15	2nd	GND	Ground		
E16	3rd	RSVD	Reserved		
S8	2nd	GND	Ground		
S9	3rd	S1T+	SAS/SATAe 1 Transmit +	SAS+SATAe	input
S10	3rd	S1T-	SAS/SATAe 1 Transmit -	SAS+SATAe	input
S11	2nd	GND	Ground		
S12	3rd	S1R-	SAS/SATAe 1 Receive -	SAS+SATAe	output
S13	3rd	S1R+	SAS/SATAe 1 Receive +	SAS+SATAe	output
S14	2nd	GND	Ground		
E17	3rd	RSVD	Reserved		
E18	2nd	GND	Ground		
E19	3rd	PETp1/S2T+	ePCIe 1 /SAS 2 Transmit +	ePCIe+SAS4	input
E20	3rd	PETn1/S2T-	ePCIe 1 /SAS 2 Transmit -	ePCIe+SAS4	input
E21	2nd	GND	Ground		
E22	3rd	PERn1/S2R-	ePCIe 1 /SAS 2 Receive -	ePCIe+SAS4	output
E23	3rd	PERp1/S2R+	ePCIe 1 /SAS 2 Receive +	ePCIe+SAS4	output
E24	2nd	GND	Ground		
E25	3rd	PETp2/S3T+	ePCIe2 / SAS 3 Transmit +	ePCIe+SAS4	input
E26	3rd	PETn2/S3T-	ePCIe2 / SAS 3 Transmit -	ePCIe+SAS4	input
E27	2nd	GND	Ground		
E28	3rd	PERn2/S3R-	ePCIe 2 / SAS 3 Receive -	ePCIe+SAS4	output
E29	3rd	PERp2/S3R+	ePCIe 2 / SAS 3 Receive +	ePCIe+SAS4	output
E30	2nd	GND	Ground		
E31	3rd	PETp3	ePCIe 3 Transmit +	ePCIe	input
E32	3rd	PETn3	ePCIe 3 Transmit -	ePCIe	input
E33	2nd	GND	Ground		
E34	3rd	PERn3	ePCIe 3 Receive -	ePCIe	output
E35	3rd	PERp3	ePCIe 3 Receive +	ePCIe	output
E36	2nd	GND	Ground		
E37	3rd	SMClk	SM-Bus Clock	PCIe opt	Bi-Dir
E38	3rd	SMDat	SM-Bus Data	PCIe opt	Bi-Dir
E39	3rd	DualPortEn#	ePCIe 2x2 Select	Dual Port	input

Dual-Port Functionality in NVMe

- Dual-Port functionality can be enabled with DUALPORTEN# (E39)
- If DUALPORTEN# is left open, dual-port functionality is not enabled and device will be configured in x4
- If DUALPORTEN# is grounded, dual-port functionality is enabled and device is configured in dual x2

Form Factors for Dual-Port NVMe SSD

- Currently supported Form Factors:
  - U.2 (SFF 8639)
  - AIC
- New Form Factors:
  - EDSFF
  - NF1

Advantages of Dual-Port NVMe

- Dual Port provides failover capabilities
  - Ability to connect to two hosts simultaneously

NVMe Dual-Port vs. SAS Dual-Port

NVMe	SAS
No improvement in performance	Significant improvement in performance for read test (Read BW (simultaneous dual-port)) = 2x (non-dual port mode)
Provides failover capabilities	Provides failover capabilities
Overall performance(NVMe PCIe Gen 3 x4): • Read BW: ~3300 MB/s • Read IOPS: ~800K	Overall performance (SAS-3) dual port mode: • Read BW: ~1900 MB/s • Read IOPS: ~270K Source: Toshiba
Form Factor: • U.2 (SFF 8639) • AIC	Form Factor: • U.2 (SFF 8639)

Power On & Link Status of Dual-Port (1/2)

- Link Status display with lspci

```
$ lspci -s 0d:00.0 -vvv | grep -A6 LnkCap
                LnkCap: Port #0, Speed 8GT/s, Width x2, ASPM L0s, Exit Latency L0s <4us, L1 unlimited
                        ClockPM- Surprise- LLActRep- BwNot- ASPMOptComp+
                LnkCtl: ASPM Disabled; RCB 64 bytes Disabled- CommClk+
                        ExtSynch- ClockPM- AutWidDis- BWInt- AutBWInt-
                LnkSta: Speed 8GT/s, Width x2, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
                DevCap2: Completion Timeout: Range AB, TimeoutDis+, LTR-, OBFF Not Supported
                AtomicOpsCap: 32bit- 64bit- 128bitCAS-
$ lspci -s 0e:00.0 -vvv | grep -A6 LnkCap
                LnkCap: Port #1, Speed 8GT/s, Width x2, ASPM L0s, Exit Latency L0s <4us, L1 unlimited
                        ClockPM- Surprise- LLActRep- BwNot- ASPMOptComp+
                LnkCtl: ASPM Disabled; RCB 64 bytes Disabled- CommClk+
                        ExtSynch- ClockPM- AutWidDis- BWInt- AutBWInt-
                LnkSta: Speed 8GT/s, Width x2, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
                DevCap2: Completion Timeout: Range AB, TimeoutDis+, LTR-, OBFF Not Supported
                AtomicOpsCap: 32bit- 64bit- 128bitCAS-
^
```

Power On & Link Status of Dual-Port (2/2)

- Link Status display with Advantest MPT3000 platform

```
*************************************
PCIe Port a Link Status:
LinkUp L0          = true
LTSSM              = 0x10, LTSSM_L0
LinkWidth          = x2
LinkSpeed          = 8.0G
ActiveLanes        = 00000011 (0x03)
ValidLanes         = 00000011 (0x03)
LinkUpCount        = 1
LinkRetrainCnt     = 3
*************************************

*************************************
PCIe Port b Link Status:
LinkUp L0          = true
LTSSM              = 0x10, LTSSM_L0
LinkWidth          = x2
LinkSpeed          = 8.0G
ActiveLanes        = 00000011 (0x03)
ValidLanes         = 00000011 (0x03)
LinkUpCount        = 1
LinkRetrainCnt     = 3
*************************************

Power up PCIe levels!
Power up was successful
```

Testing Dual-Port NVMe (Production Test Flow – 1/4)

- User can access one port at a particular point of time
  - Either Port A or Port B
- Complete LBA address range can be accessed from both ports

Testing Dual-Port NVMe (Production Test Flow – 2/4)

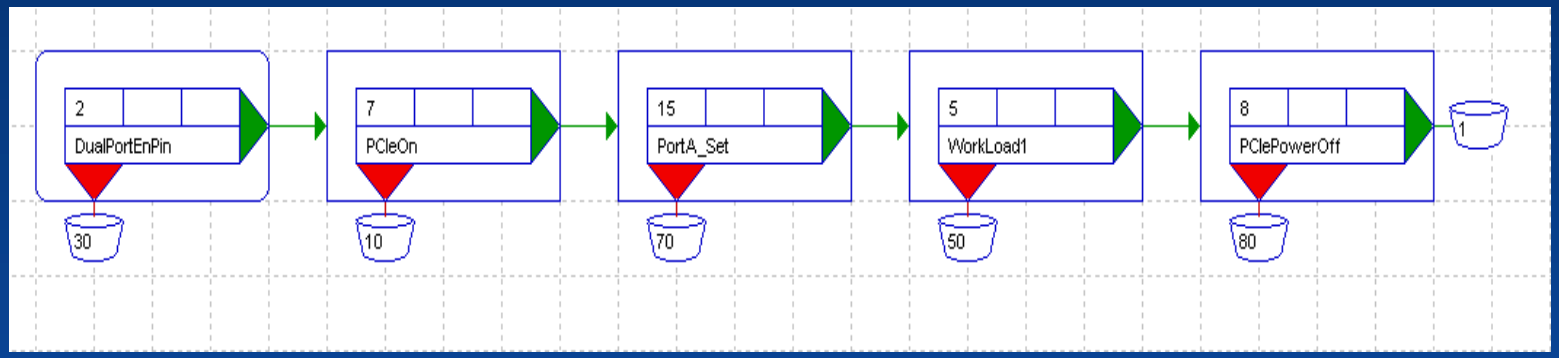

Testing Dual-Port NVMe (Production Test Flow – 3/4)

- Example Test Case:
  - Write to a particular LBA range using Port A
  - Read from the same LBA address range using Port B

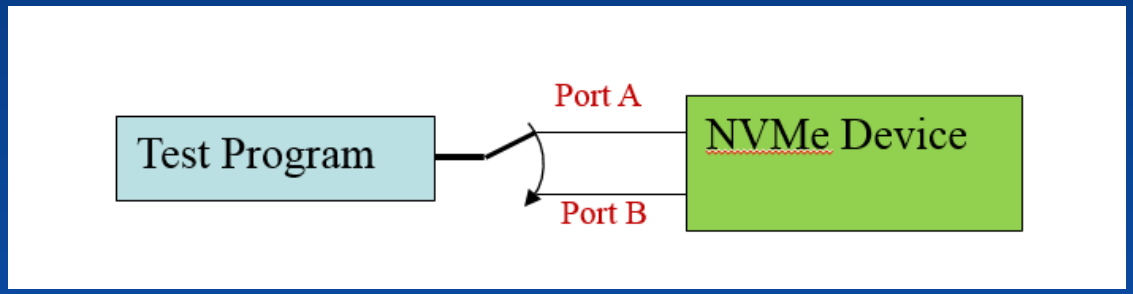

Testing Dual-Port NVMe (Production Test Flow– 4/4)

- Checks if both Port A and Port B is functional
- Checks if same link rate can be negotiated on both Ports
- Part of Production Test Flow

Testing Dual-Port NVMe (QA Test Flow – 1/4)

- User can access both ports simultaneously.
- *Simultaneous Read:* Both ports can access the same LBA address range
- *Simultaneous Write:* LBA address range should be exclusive to the ports.

Testing Dual-Port NVMe (QA Test Flow – 2/4)

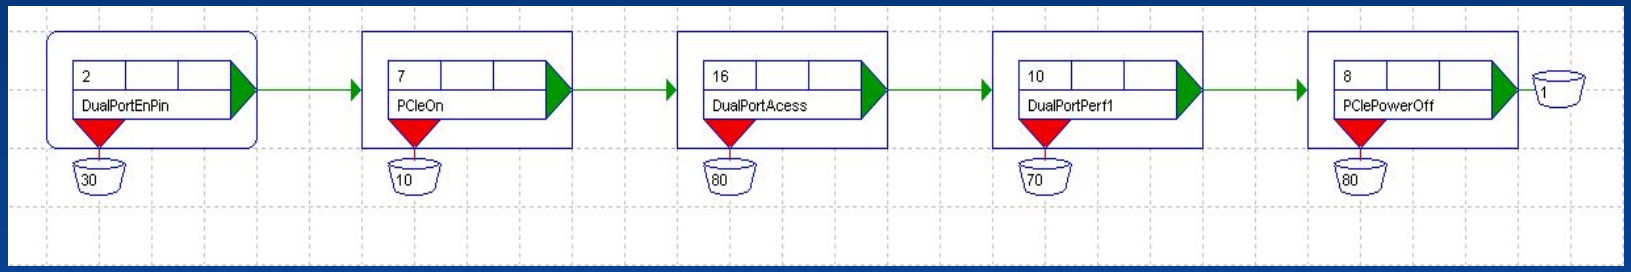

Testing Dual-Port NVMe (QA Test Flow – 3/4)

- Example of Test Case (Simultaneous Write)
  - Write to exclusive LBA address ranges using Port A and Port B simultaneously.

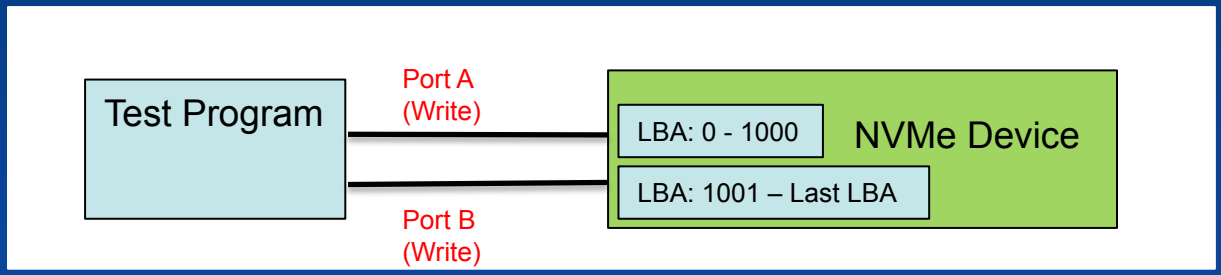

Testing Dual-Port NVMe (QA Test Flow– 3/4)

- Example Test Case (Simultaneous Read)
  - Read from the same LBA address range using both Port A and Port B

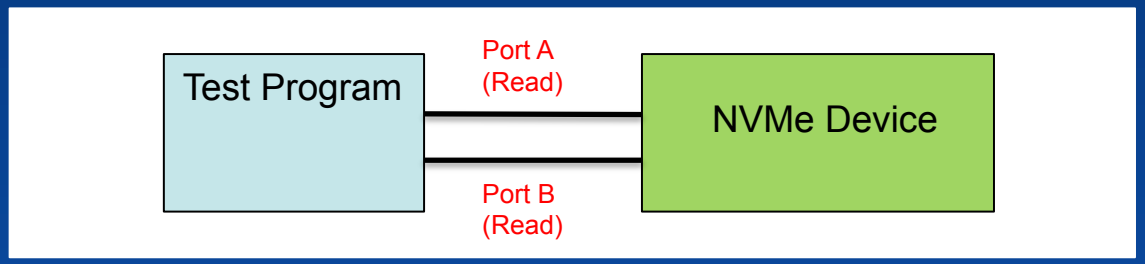

Testing Dual-Port NVMe (QA Test Flow– 4/4)

- Checks if both ports can be accessed at the same time
- Checks if there is no performance degradation with simultaneous dual-port access
- Part of QA Test Flow

Conclusion

- Dual-port devices ensure redundancy and provides failover capability
- Provides ability to connect to two hosts simultaneously
- Used in Enterprise SSD market
- Current form factor is U.2

Flash Memory Summit

More Info: NVMe Dual-Port Testing

- Please visit us at Booth # 606 for to know more about NVMe Dual-Port Testing

Thanks