



Flash Memory Summit

Challenges in Zoned Namespace (ZNS) Testing

Kean-Yau Chaw

Advantest

ZNS: An Endurance Architecture – Part 1

August 2, 2022

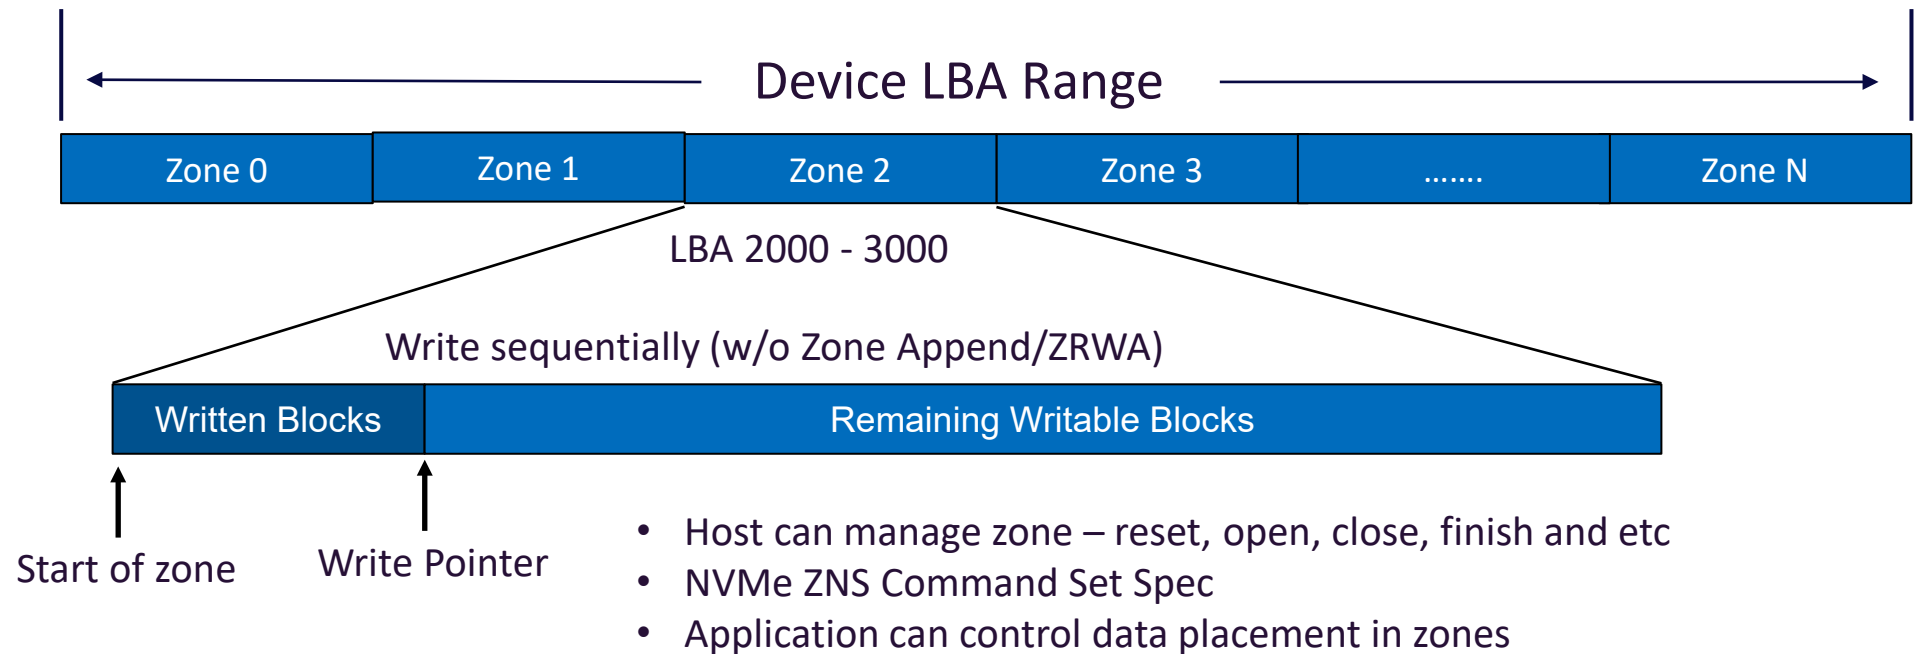
- ZNS vs Conventional Namespace
- Benefits of ZNS
- 5 ZNS Test Challenges
- ZNS Devices & Market Segment
- Summary

ZNS vs Conventional Namespace

1. Conventional Namespace



2. Zoned Namespace (ZNS)



ZNS is supported in Linux Kernel 5.9, and NVMe CLI tool v1.12

Summary of ZNS vs Conventional Namespace

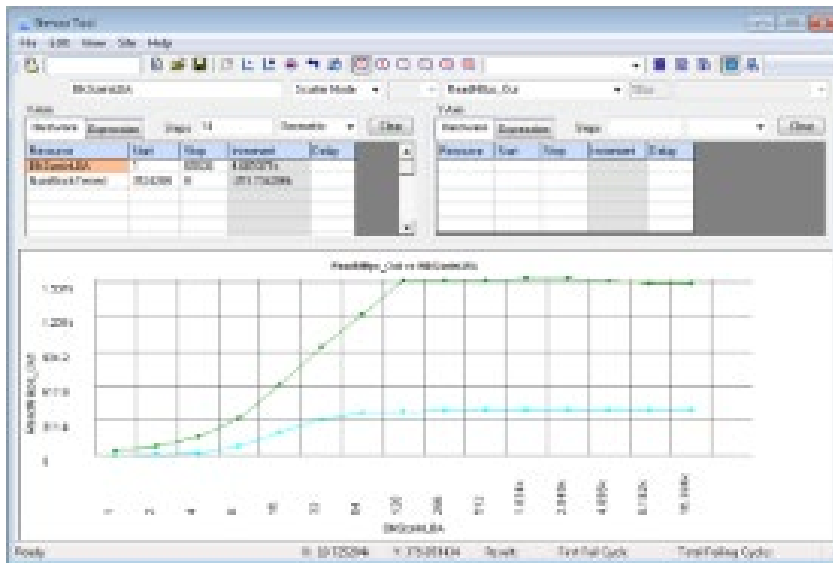
No	Conventional Namespace	Zoned Namespace (ZNS)
1	SSD reset NAND Flash before write (Garbage collection)	Host reset NAND Flash before write (Reset zone)
2	Write/Read using LBA	Write/Read using LBA that falls into a zone assignment (Exp: LBA 100 in zone 2)
3	Write in any LBA	Write sequentially based on write pointer within the zone (If ZRWA/Zone Append is not supported)
4	NA	Zone can be managed by open, close, finish, reset and etc.
5	NA	Application can write data into specified zone

Benefits of ZNS

No	Benefit
1	Extend SSD Lifespan (WAF Reduction)
2	Improve Performance & QoS
3	Lower Cost (Less Over-Provisioning, DRAM)

ZNS Test Challenge 1

- Achieve maximum bandwidth in ZNS device with high parallelism



ZNS Test Challenge 2

- Lack of debug capability feature: User doesn't know zone access sequence

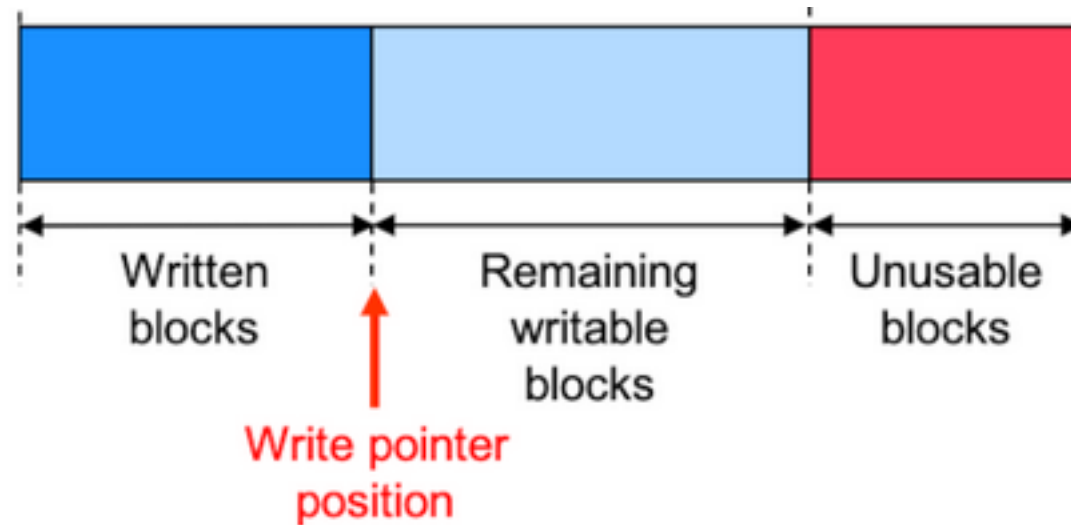
Zone pref result accessed zone by seq#:

seq#	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
0	1031	3339	1253	2881	2656	2244	1187	100	2348	2362	2328	3066	1031	1253	2328	965	1051	1195	554	390

```
Display explicitly opened zones:
znsi: nr_zones=          zone_size_lbas=          zone_capacity_lbas=          max_open_zones=
 0 SLBA:          WP:          Cap:          State: exp_opened  Type: seqwrite_req  Attrs: 0
 1 SLBA:          WP:          Cap:          State: exp_opened  Type: seqwrite_req  Attrs: 0
 2 SLBA:          WP:          Cap:          State: exp_opened  Type: seqwrite_req  Attrs: 0
 3 SLBA:          WP:          Cap:          State: exp_opened  Type: seqwrite_req  Attrs: 0
 4 SLBA:          WP:          Cap:          State: exp_opened  Type: seqwrite_req  Attrs: 0
Total zones in this state = 5
```

ZNS Test Challenge 3

- Host needs to retain write pointer in open zone to get maximum bandwidth (Query from device with ZNS command is slow)



ZNS Test Challenge 4

- Host needs to execute actions (reset zone, open zone, finish zone) automatically for bandwidth test

Zone
management
send

Reset
zone

LAST 1024 Command LIST:

(DWs are in hex)			DW1	DW2	DW3	DW4	DW5	DW6	DW7	DW8	DW9	DW10	DW11	DW12	DW13
86:	ios	20200701_15:17:39.203480318	400079	2	0	0	0	0	0	0	0	0	0	0	4
87:	ioc	20200701_15:17:39.203523282	0	0	10002	10040									
88:	ios	20200701_15:17:39.204213275	410001	2	0	0	0	0 a2ff1000	3 884b3000	4	0	0	1f	0	
89:	ioc	20200701_15:17:39.204317218	0	0	10003	10041									
90:	ios	20200701_15:17:39.204373790	420001	2	0	0	0	0 a2fd1000	3 884b3000	4	20	0	1f	0	
91:	ioc	20200701_15:17:39.204474062	0	0	10004	10042									
92:	ios	20200701_15:17:39.204489028	430001	2	0	0	0	0 a2fb1000	3 884b3000	4	40	0	1f	0	
93:	ioc	20200701_15:17:39.204579343	0	0	10005	10043									
1015	[1023]	20200701_15:56:44.409739008				205744 ns WRITE				Ox3143420 + Ox20	Ox8:0000 Good	0 [0: 0]			
1016	[0]	20200701_15:56:44.409920768				1816 Finish Zone				Ox3143440 + Ox20	Ox4d:0000 Good	0 [0: 0]			
1017	[1]	20200701_15:56:44.410101760				181024 ns WRITE				Ox3143460 + Ox20	Ox44:0000 Good	0 [0: 0]			
1018	[2]	20200701_15:56:44.410264064				3391504 ns ZFINI				Ox3180000 + Ox20	Ox6:0000 Good	0 [0: 0]			
1019	[3]	20200701_15:56:44.410307840				206016 ns WRITE				Ox3143480 + Ox20	Ox7f:0000 Good	0 [0: 0]			
1020	[4]	20200701_15:56:44.410449664				141728 ns WRITE				Ox31434a0 + Ox20	Ox20:0000 Good	0 [0: 0]			
1021	[5]	20200701_15:56:44.410614784				164960 ns WRITE				Ox31434c0 + Ox20	Ox73:0000 Good	0 [0: 0]			
1022	[6]	20200701_15:56:44.410740992				126192 ns WRITE				Ox31434e0 + Ox20	Ox3e:0000 Good	0 [0: 0]			
1023	[7]	20200701_15:56:44.415158016				4417024 ns ZFINI				Ox3100000 + Ox20	Ox8c:0000 Good	0 [0: 0]			

ZNS Test Challenge 5

- Host needs to access multiple zones concurrently

TestLabel	TestDlog	Start Zone	End Zone	BlkSize InLBA	NumBlock Tested	ReadPercentage	Random Address Percentage
100RW	DLOG_UNCHANGED	0	-1	_128	0	0	100
Number OfZones	ZoneOptions	NumberOf Concurrent Zones	AlignLBA_Opt	NCQ_Opt	Random Seed_Opt	InitialData Value_Opt	DataCompareOpt
20	RSTBEFWRT_FNSHFULLWRTZONE	12	8	12	0	-1	DATA_COMPARE_DISABLE
DataGen Type_Opt		BufferData_Opt	Entropy Percent	Loops_Opt	Timeout_opt	NumIOPerZRWA Zone_Opt	ZRWARandom Percent
USER_BUFFER_DATA		BufPat_RANDOM	0	1	0	3	0

ZNS Devices & Market Segment

- Data center / Cloud
 - U.2 Form Factor
- Hyperscaler
 - EDSFF Form Factor



Image from [Solid State Drive Form Factors | SNIA](#)

- ZNS Test Challenges
 - Functionality test in compliance of ZNS spec
 - Multiple form factors
 - New debug tool
 - High parallelism with maximum bandwidth
- ZNS device and system are in market. However, conventional namespace device and system are still the dominant.

Questions?

