FDP: Flexible Data Placement Methods for the Benefit of Conventional Applications

Luca Bert, Micron Technology, Inc.

Problem statement

- Future NAND technologies (QLC, PLC,…) will not be able to guarantee same endurance and performance margins as older nodes
- Data layout and NAND management can help bridging the gap and FDP is a front runner in this approach
- FDP requires Host SW to be FDP-aware and existing solutions will need to be ported to the new protocol to benefit of it
- Can conventional solutions, “non-FDP-aware”, benefit of the same, at least in some configurations?

Conventional SSD data layout

- Conventional applications have no understanding of NAND data layout
- Multiple app data get inter-mixed on NAND
- This has negative impact on:
  - Endurance (due to higher Write Amplification)
  - Performance (due to data scattering)

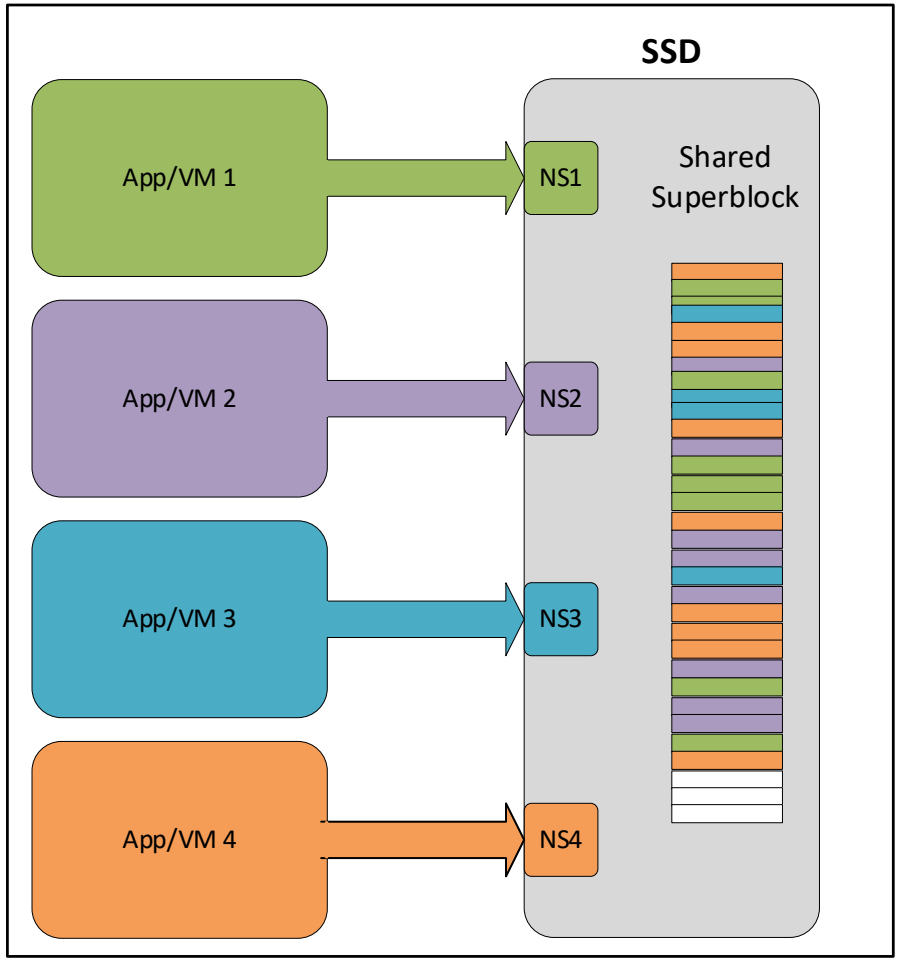

FDP opportunities for non-FDP aware applications

- It is conceivable that in multi-tenant or virtualized environments SSD could leverage FDP in backend
- Does not require App/VM changes and allows most (but not all) FDP advantages
- SSD may implicitly tag NS access to specific regions (RUH tagging)
- Data separation in dedicated region will improve performance stability and reduce Write Amplification

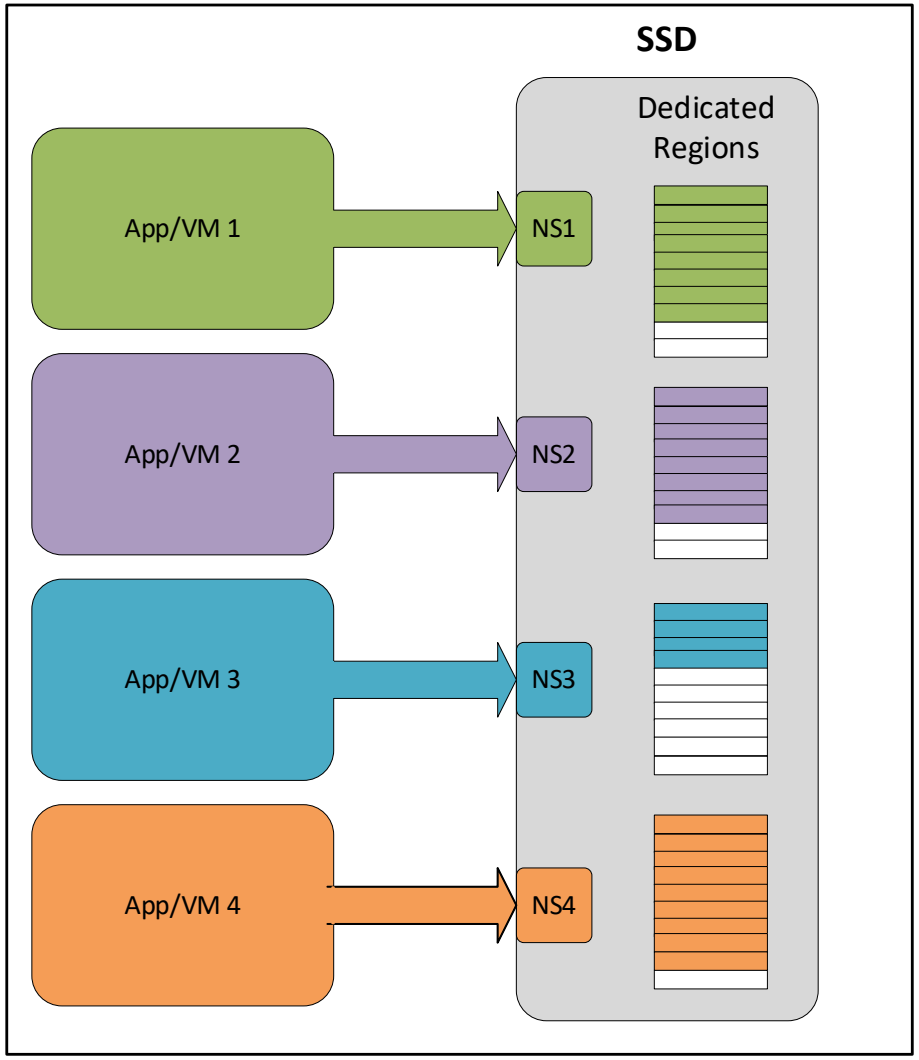

Simple simulation case

- Use reference benchmarks to test Write Amplification improvements in Conventional and FDP-aware configurations
- Focus on Persistently Isolated FDP
- Use FIO, JESD219A.01 and variant, Real-Life benchmarks (Aerospike)

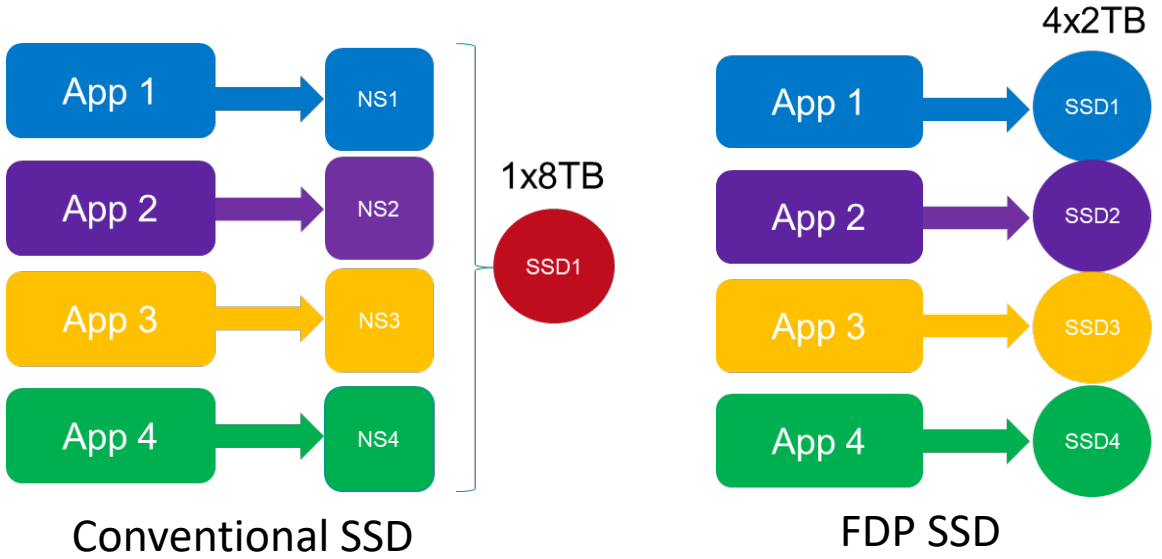

Simple simulation case

- FIO Apps profiles:
  - App 1: Seq_write_4k_64QD
  - App 2: Seq_write_16k_32QD
  - App 3: Seq_write_64k_16QD
  - App 4: Seq_write_256k_8QD
- Run with JEDEC JESD219A.01
- Run with Aerospike
  - 4DB 2TB each
  - 32K, 64K, 128K and 256K write sizes
  - Benchmark: YCSB WL A (50% R, 50%W, Uniform distribution)

	Write Amplification		
	Conv SSD	FDP SSD	%
FIO	1.63	1.02	-37%
JESD219A.01 (512B)	3.95	3.8	-4%
JESD219A.01 (4KB)	3.73	3.48	-7%
Aerospike	1.84	1.09	-41%

Observations

- All benchmarks benefit of FDP-aware configurations
- Real Life benchmarks are 37%-41% but given WAF baseline is 1, improvement is much higher
- Real Life benchmarks (Aerospike and, in lesser degree, FIO) benefit much more than purely synthetic (JESD)
  - We cannot rely on benchmarks who do not follow realistic behaviors
  - There is an opportunity to improve existing WAF characterization methods

	Write Amplification		
	Conv SSD	FDP SSD	%
FIO	1.63	1.02	-37%
JESD219A.01 (512B)	3.95	3.8	-4%
JESD219A.01 (4KB)	3.73	3.48	-7%
Aerospike	1.84	1.09	-41%

Questions?