Exploring Performance Paradigm of HMB NVMe SSDs

Presenter(s):

Pradeep S R,
Associate Technical Director,
SSIR-Samsung Electronics

Ranjith Thottathil,
Associate Staff Engineer,
SSIR-Samsung Electronics

Disclaimer

This presentation and/or accompanying oral statements by Samsung representatives collectively, the "Presentation" is intended to provide information concerning the SSD and memory industry and Samsung Electronics Co., Ltd. and certain affiliates (collectively, "Samsung").While Samsung strives to provide information that is accurate and up-to-date, this Presentation may nonetheless contain inaccuracies or omissions. As a consequence, Samsung does not in any way guarantee the accuracy or completeness of the information provided in this Presentation.

This Presentation may include forward-looking statements, including, but not limited to, statements about any matter that is not a historical fact; statements regarding Samsung's intentions, beliefs or current expectations concerning, among other things, market prospects, technological developments, growth, strategies, and the industry in which Samsung operates; and statements regarding products or features that are still in development. By their nature, forward-looking statements involve risks and uncertainties, because they relate to events and depend on circumstances that may or may not occur in the future. Samsung cautions you that forward looking statements are not guarantees of future performance and that the actual developments of Samsung, the market, or industry in which Samsung operates may differ materially from those made or suggested by the forward-looking statements in this Presentation. In addition, even if such forward-looking statements are shown to be accurate, those developments may not be indicative of developments in future periods.

Agenda

- HMB Overview
- Performance Parameters Overview
- HMB vs DRAM Comparison
- HMB Tuning Comparison
- Conclusion

HMB Overview

What & Why of HMB?

- Host Memory Buffer (HMB)

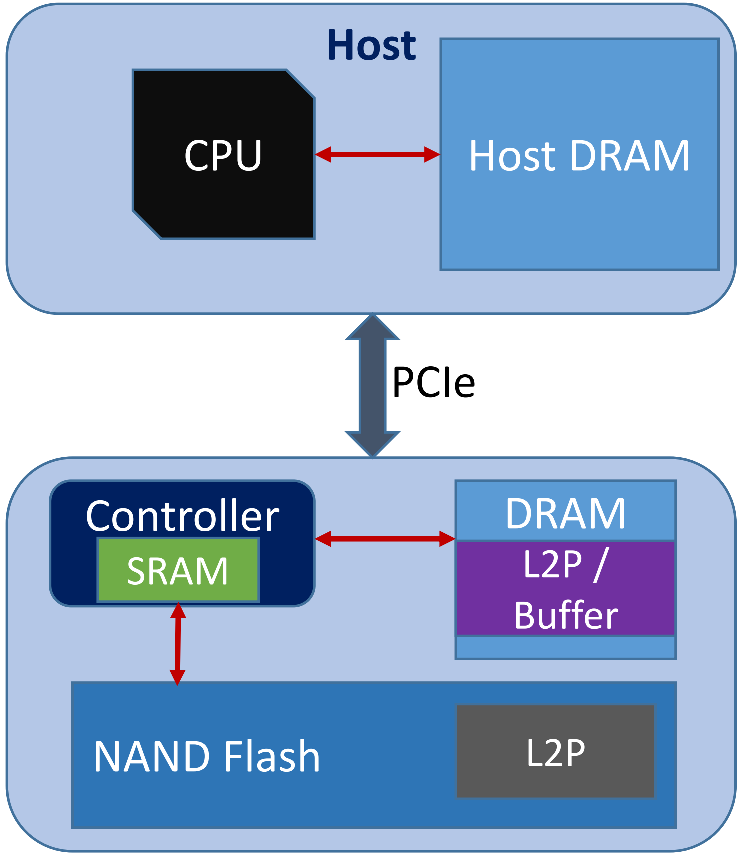

Conventional SSD w/ DRAM

HMB Motivation:

- Cost Reduction
- Reduced Size
- Market Adoption

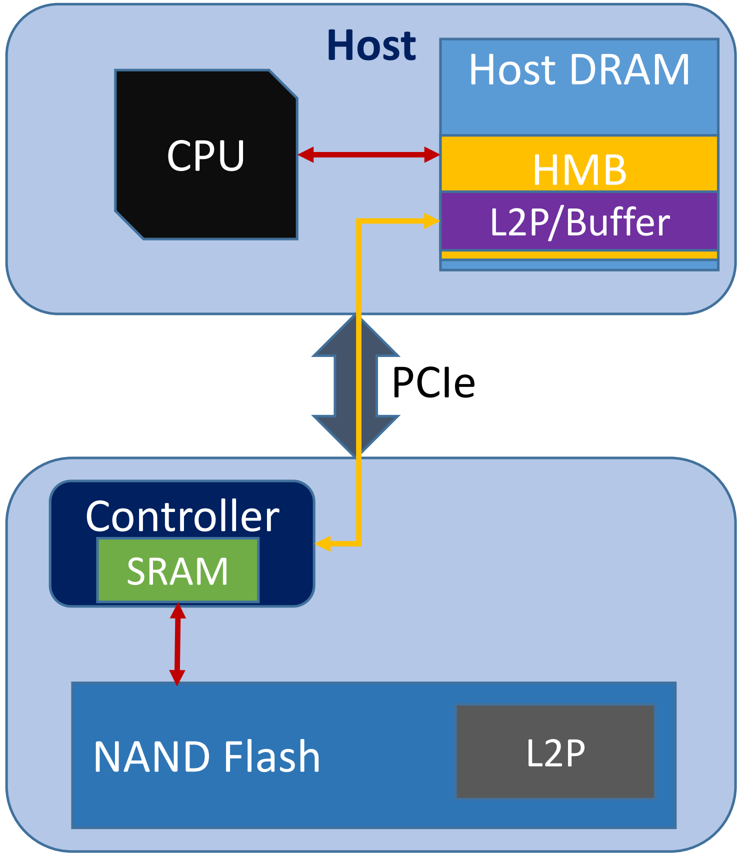

Value SSD w/ HMB

NVMe HMB Feature

- OS & Device Support
  - Windows 10 OS & Above
  - *Linux 5.12,Ubuntu 20.12 & Above
  - NVMe 1.2 & above compliance devices
- Device Requirements
  - HMB Preferred Size (HMPRE)
  - HMB Minimum Size (HMMIN)
- Host configuration of HMB
  - Enable HMB
  - HMB Size
  - Descriptor Structure with each entry having Buffer address & chunk size
  - Memory Return

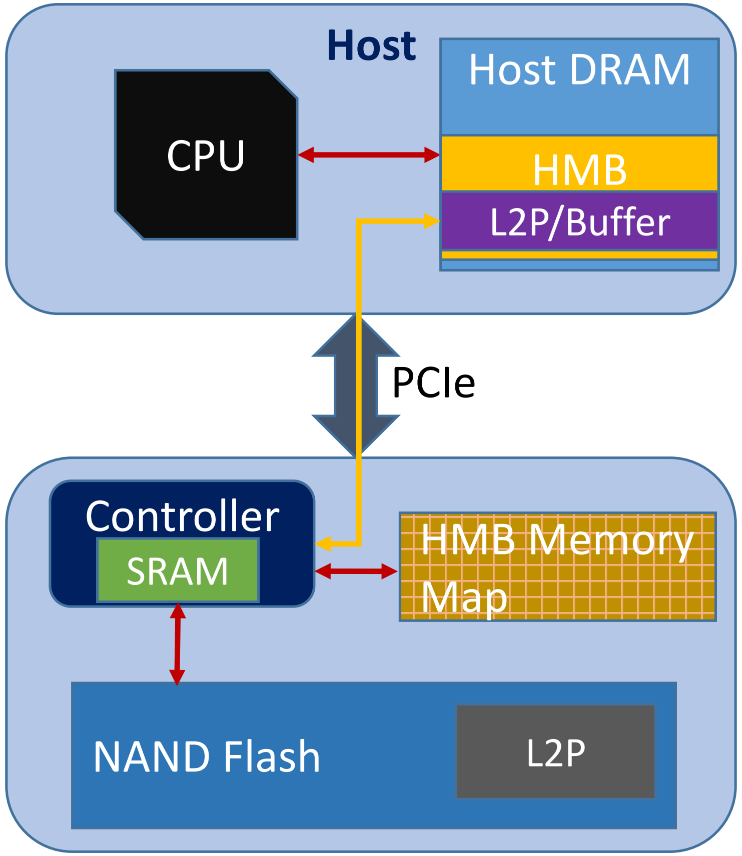

*Reference : https://www.phoronix.com/review/samsung-980-linux

Performance Parameters Overview

Key Factors

Device Factor(s)	Host Factor(s)
Number of Queues	IO Chunk size (Page alignment)
Queue Depth	Workload Types
DRAM Size	HMB Size
IO Chunk size (Transfer limits)	Number of Threads

Experiment set up

Sl No	Item	Configuration
1	NVMe SSD “A” with DRAM	DRAM Size: 1GB
2	NVMe SSD “B” with HMB	HMBPRE Size: 64MB HMMIN Size: 16MB **Note:** Both “A” & “B” are of same Density & NAND Type, Controllers are different.
3	Device Driver	Custom Test Driver* **Note:** Custom driver used in order to override the Queue creation at initialization
4	**Focus Area(s) of Experiment**	▪ Number of Queue ▪ HMB Size allocation ▪ Drive States
5	Tool	IO Meter
6	Workloads	Standard Sequential & Random workloads
7	Host Hardware	Windows 10 OS, 32GB DRAM, Intel i7 ASROCK Z690 Taichi Razor Edition

***Disclaimer(s):**

- Custom Test Driver is in-house & not built for performance, thus numbers can vary across inbox driver.
- This Experiment purpose is to find behavior patterns Only
- Performance may vary depending on various factors & each supplier’s firmware policy

HMB vs DRAM Comparison

Workloads: Sequential Write (128K-T1-QD32)

- DRAM device leads in all Drive states & gap widens multifold at aging
- HMB device at 1Q configuration shows lowest performance in Aging

Normal: No Pre-Condition

Sustain: Turbo area Pre-Conditioned

Aging : Full Drive Pre-conditioned

Max-Q : Max Queues Supported by Device

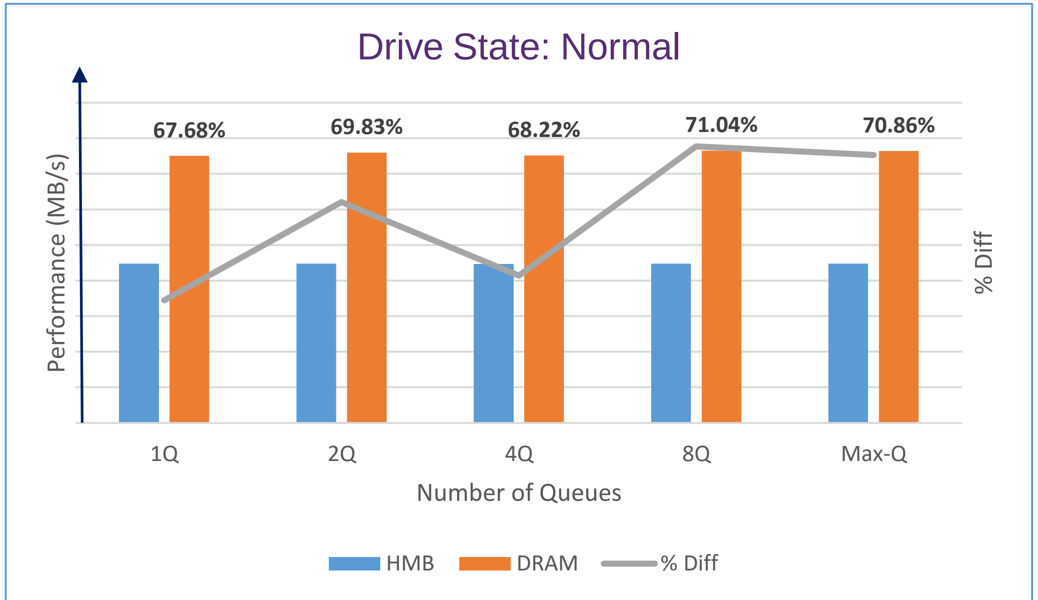

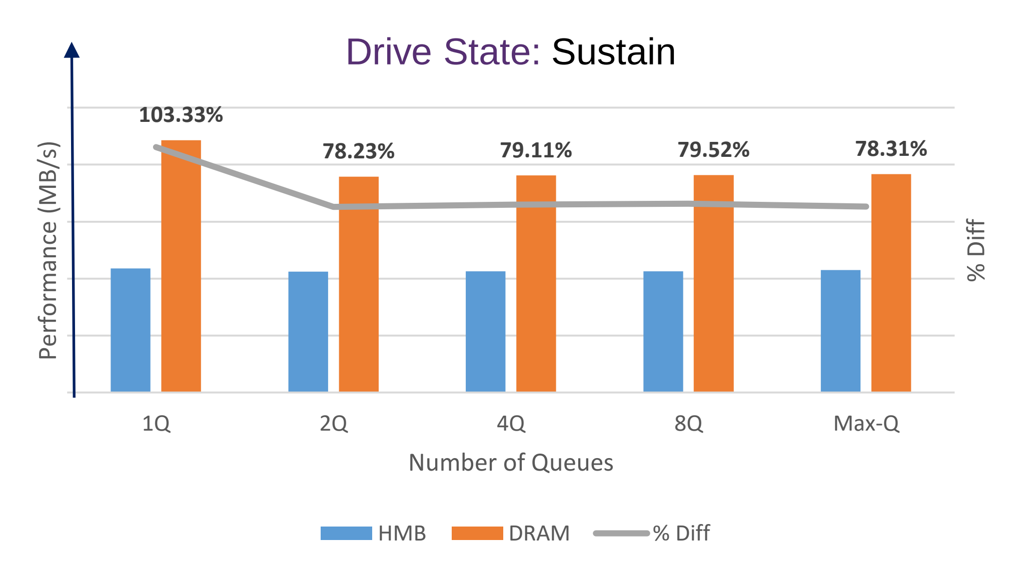

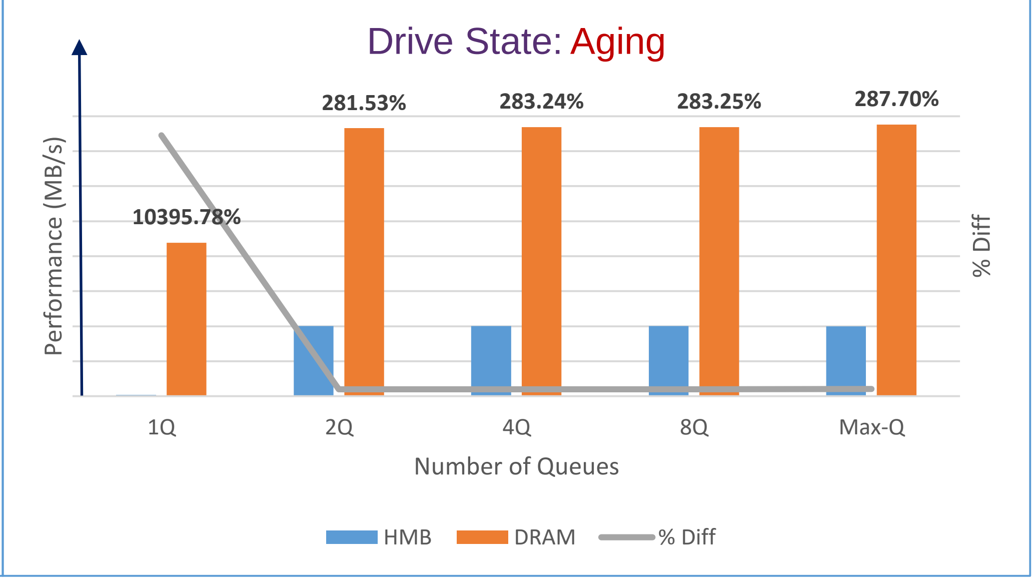

Disclaimer: Results are based on Custom test driver only

Workloads: Sequential Read (128K-T1-QD32)

- DRAM device leads in all Drive states & gap widens multifold at aging
- 1Q configuration has better results for HMB device than in sequential write case in aging

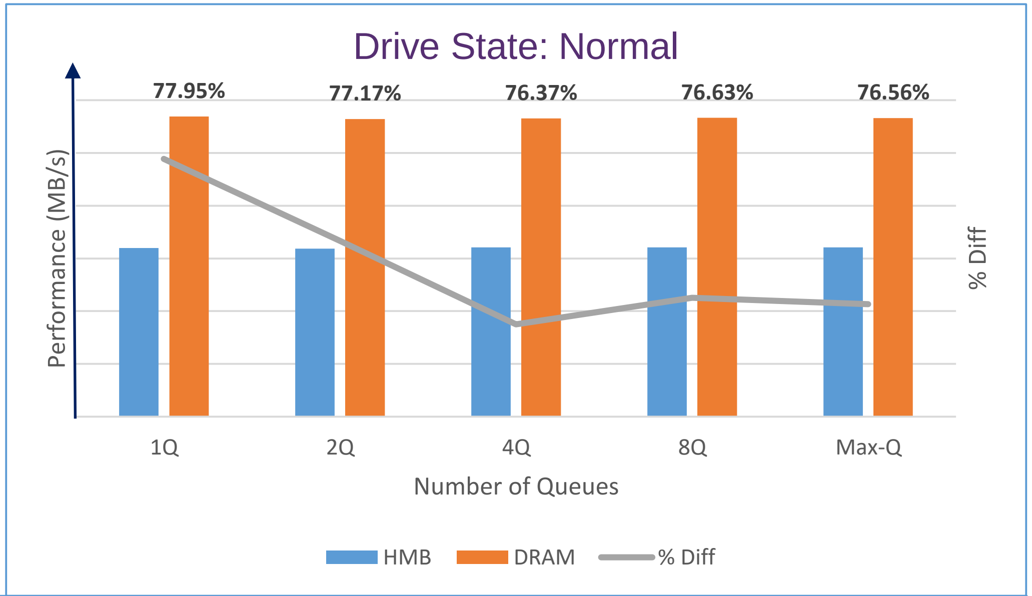

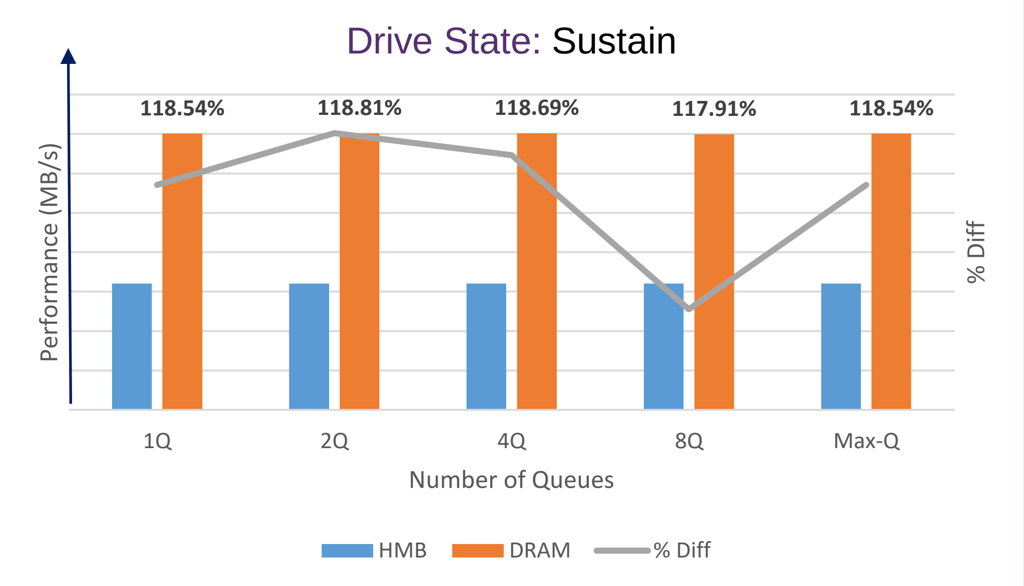

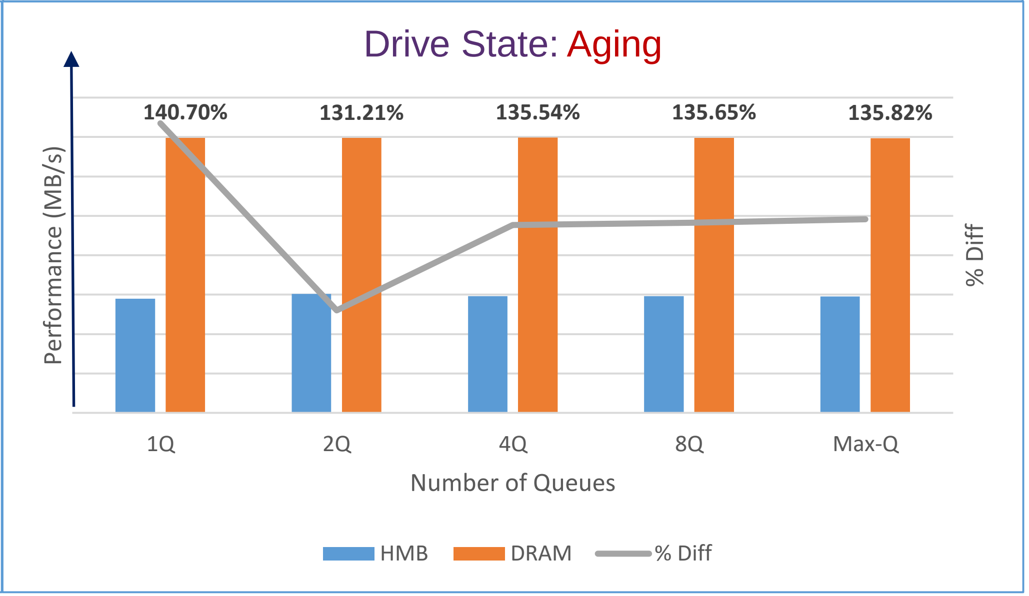

Disclaimer: Results are based on Custom test driver only

Workloads: Random Read (4K-T1-QD32)

- HMB device performed better than DRAM device in all Drive states & Queue Configurations
- HMB is more Consistent at Aging state a well.
- In Aging with increase in number of Q's, DRAM is able to perform better

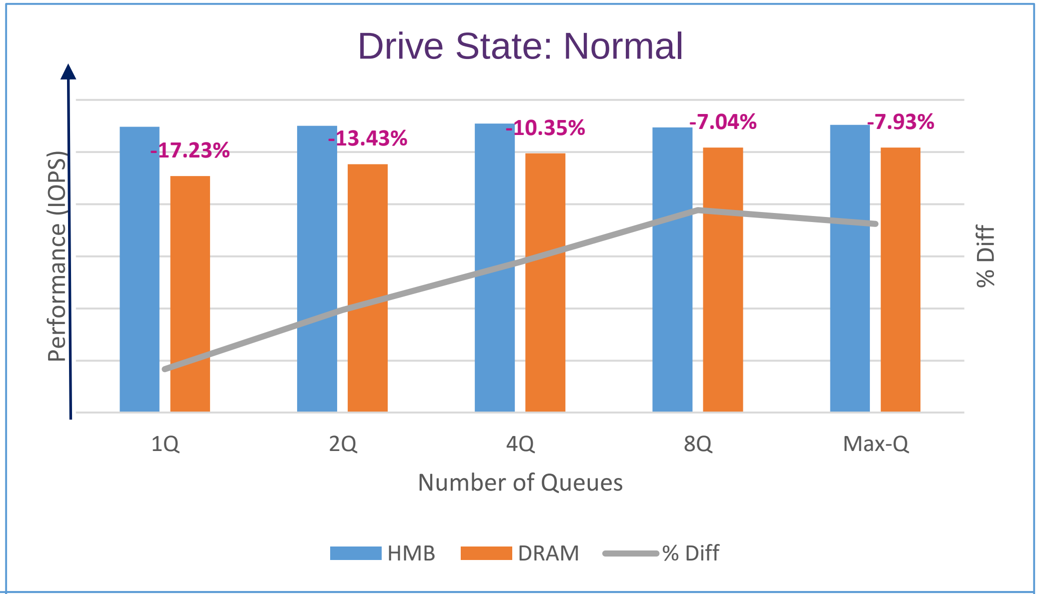

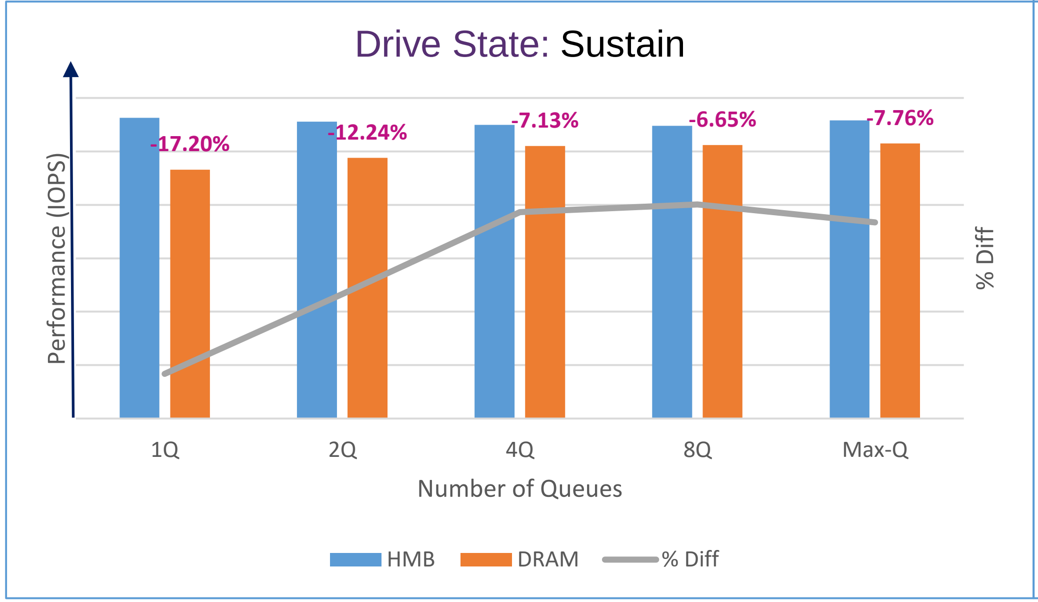

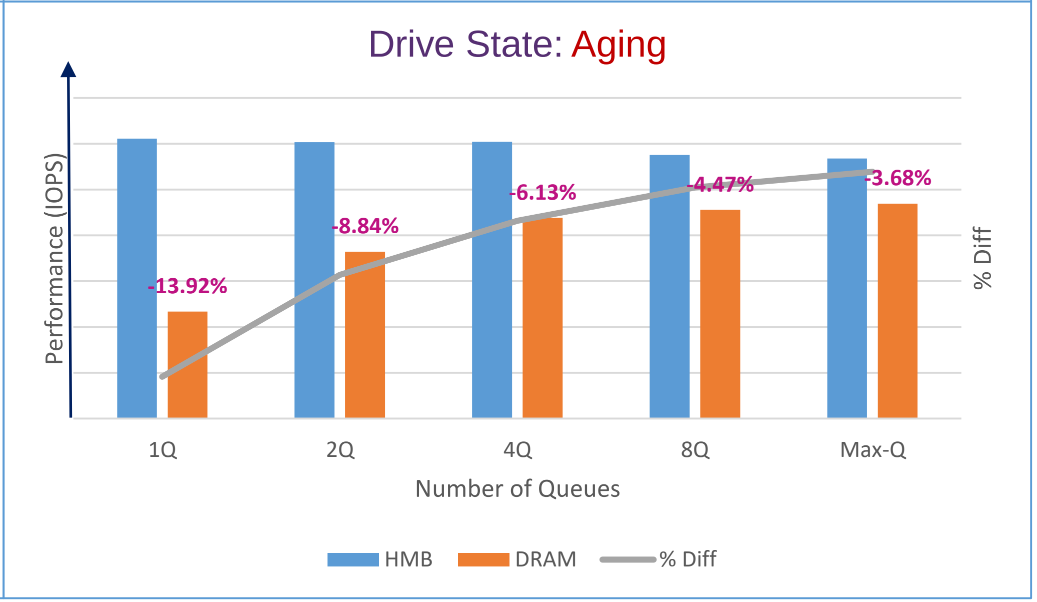

Disclaimer: Results are based on Custom test driver only

Workloads: Random Read (4K-T16-QD32)

- HMB device performed better than DRAM device in all Drive states & Queue Configurations
- HMB is more Consistent at Aging state a well.
- In Aging with increase in number of Q's. DRAM is able to perform better

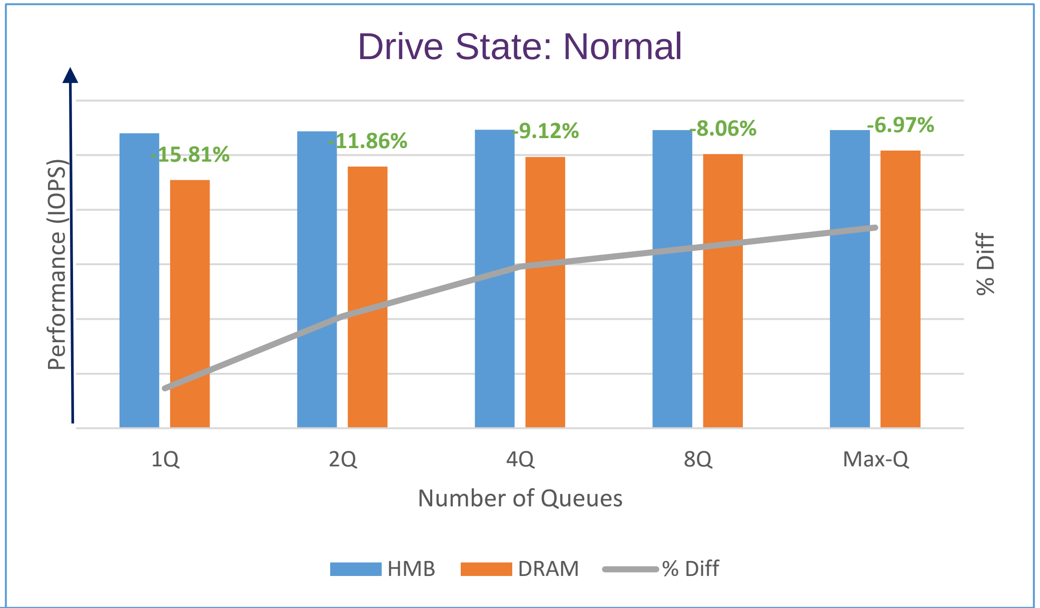

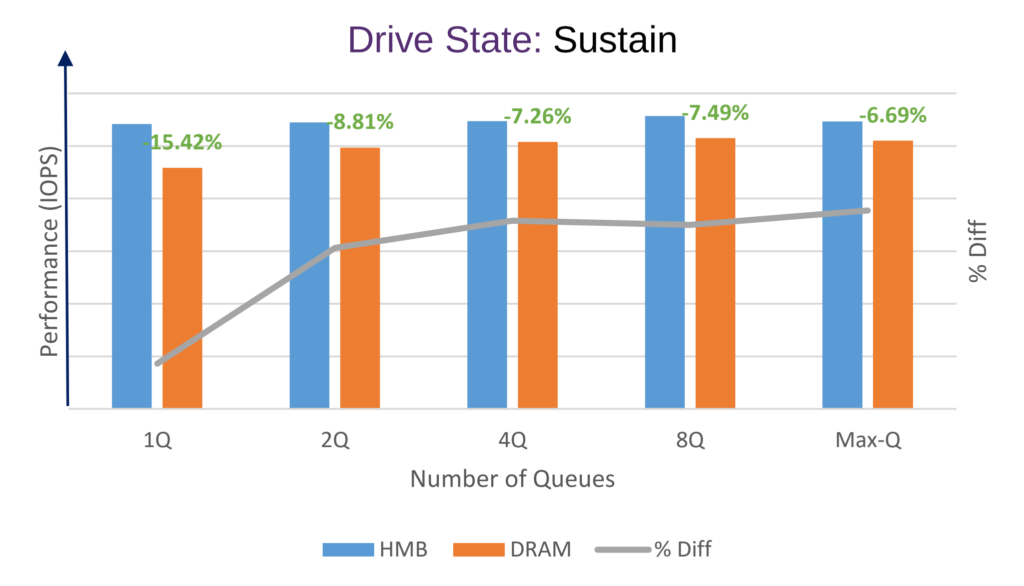

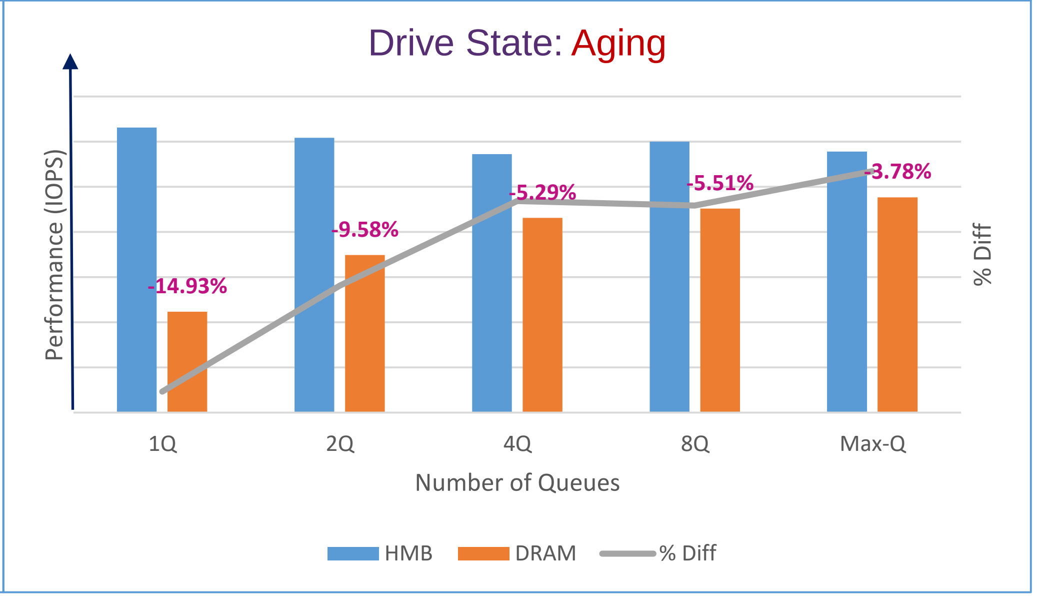

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T1-QD32)

- In Normal & sustain States both device perform on par across majority of Queue states
- In Normal & Sustain state Performance increased as number of Q's increased
- In Aging DRAM device leads by large margin

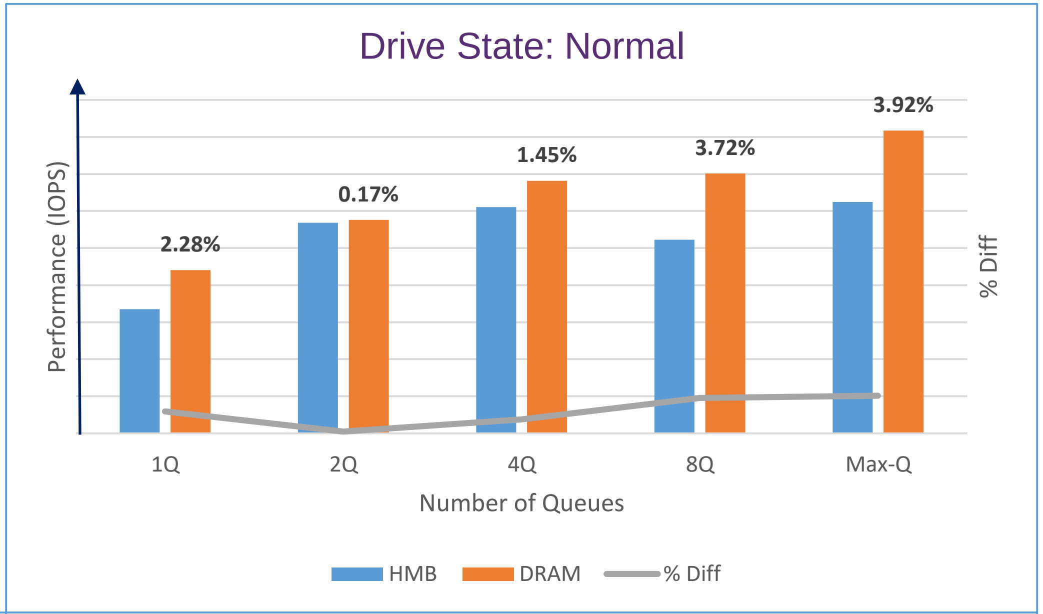

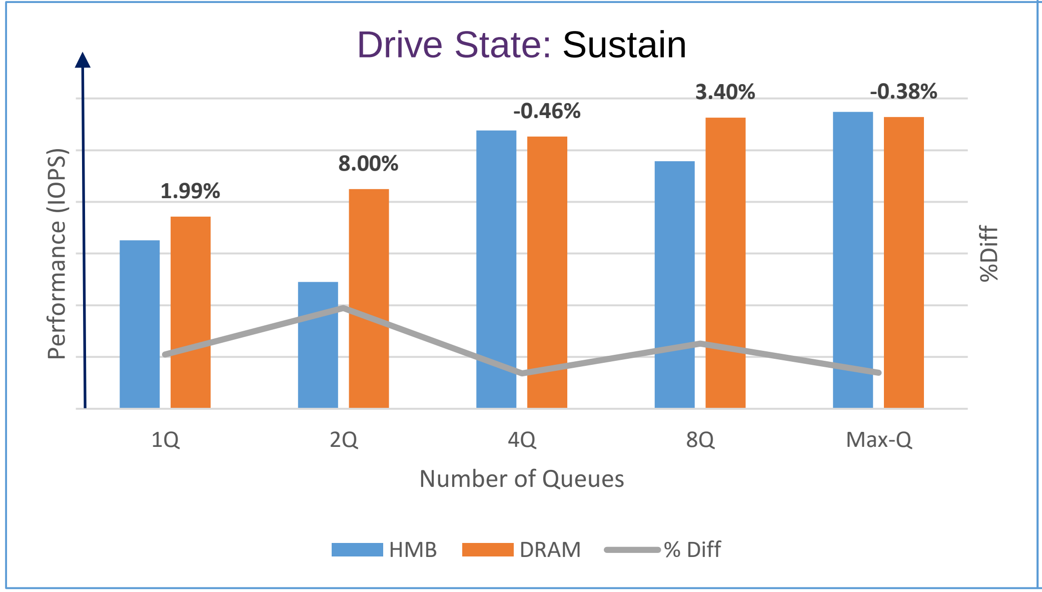

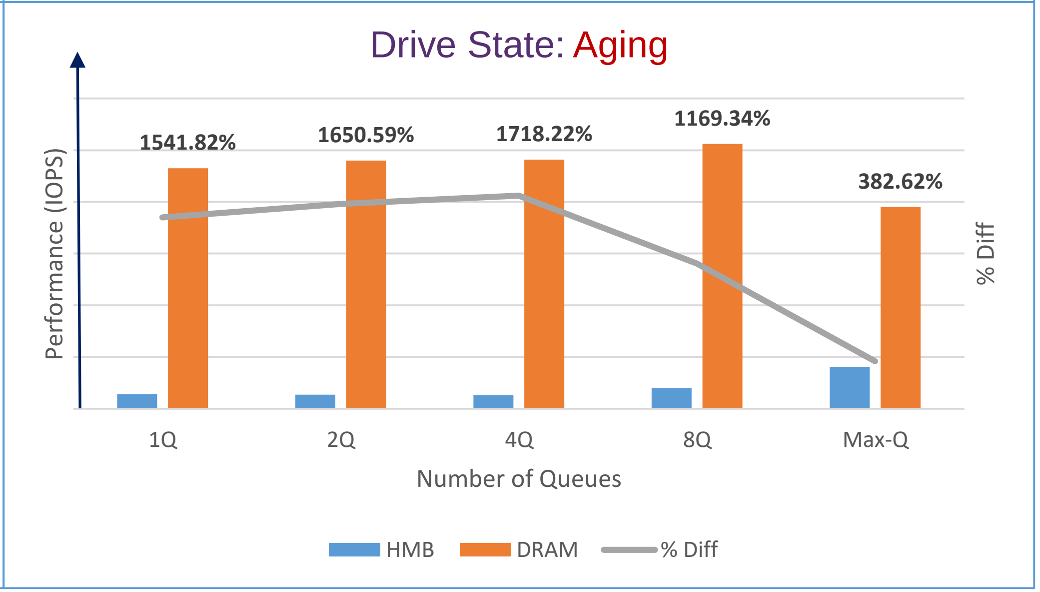

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T16-QD32)

- In Normal & sustain States both device perform on par across majority of Queue states
- In Normal & Sustain state Performance increased as number of Q's increased
- In Aging DRAM device leads by large margin

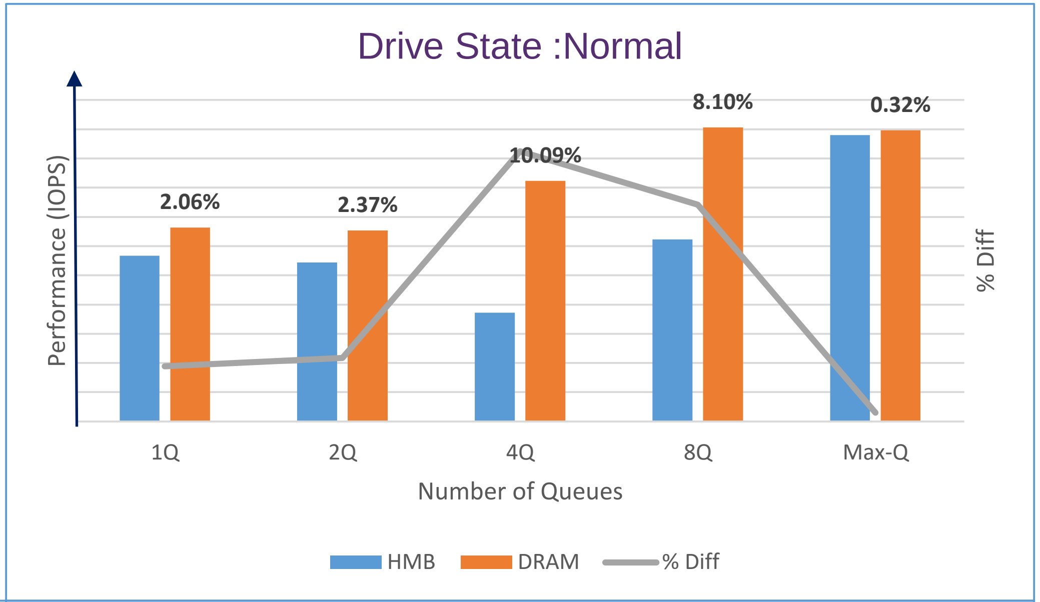

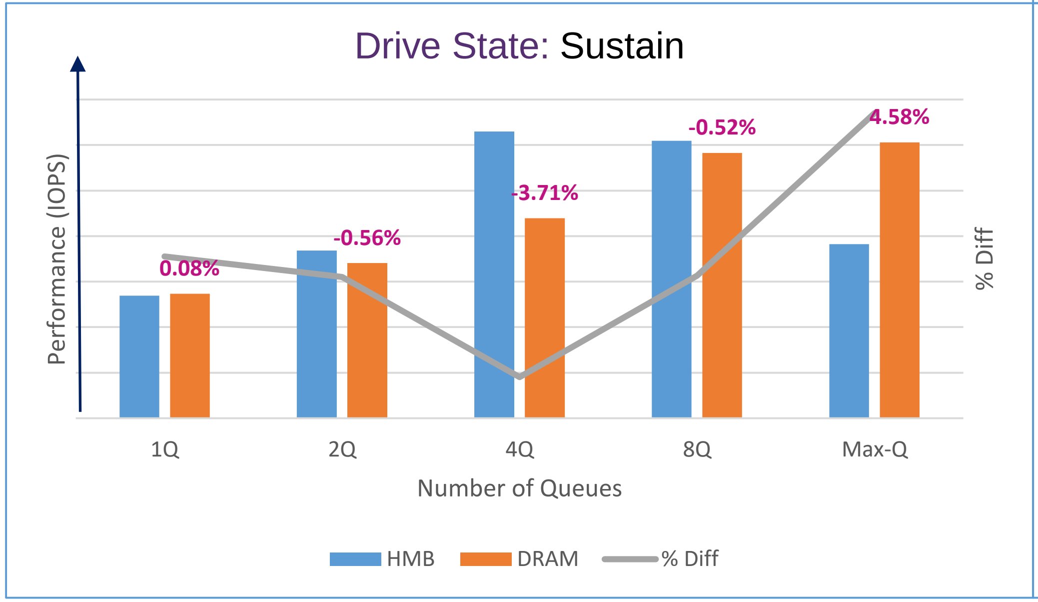

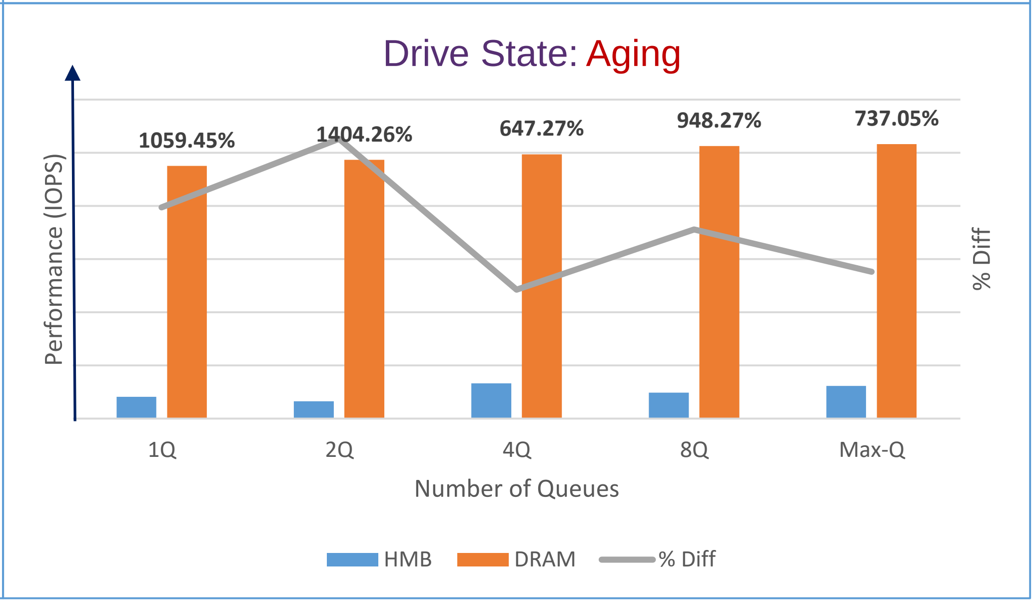

Disclaimer: Results are based on Custom test driver only

Multi-Density Comparison

Workloads: Random Read (4K-T1/16-QD32)

- For low density, the DRAM is on little higher side than HMB in Aging state across T1 & T16 workloads

Density: High

Density: Low

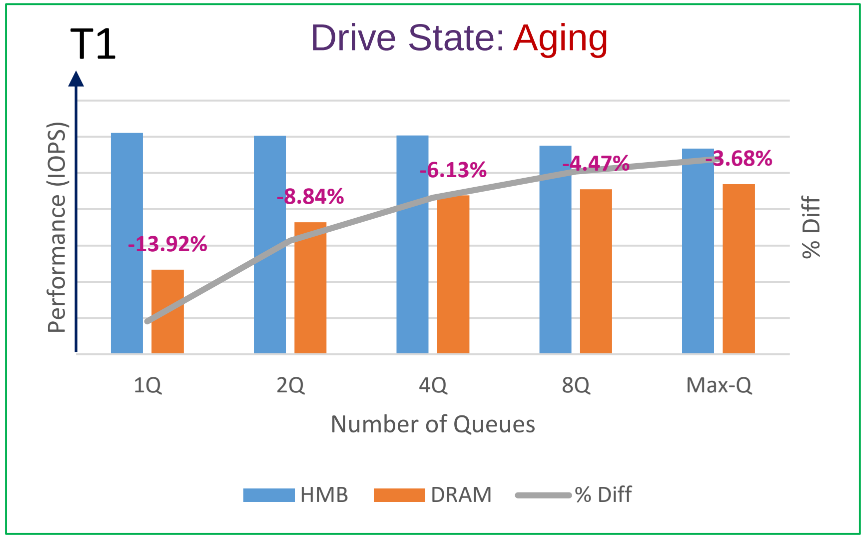

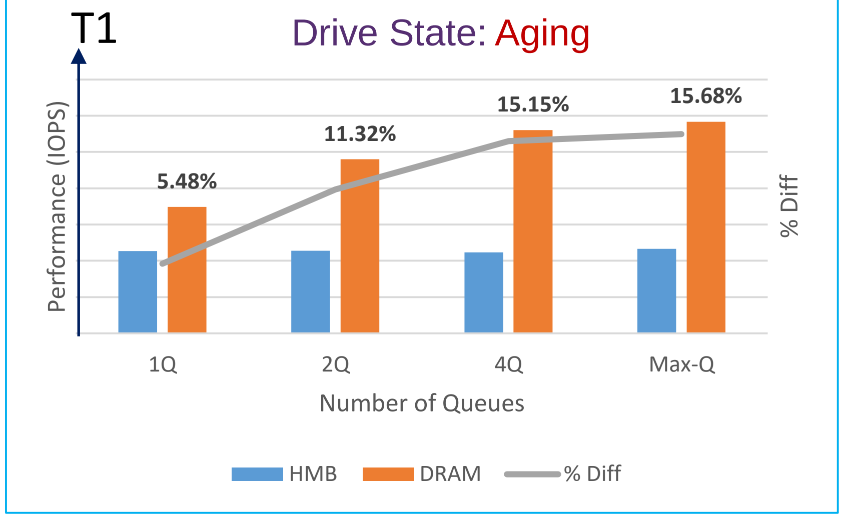

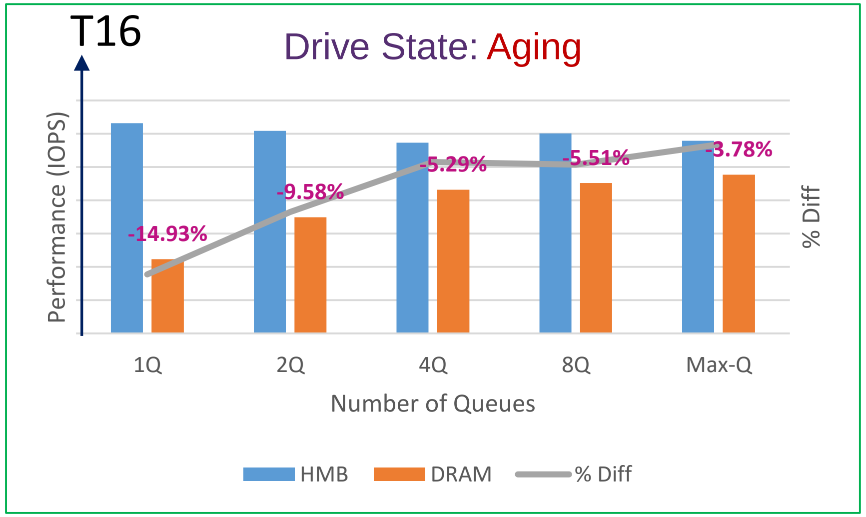

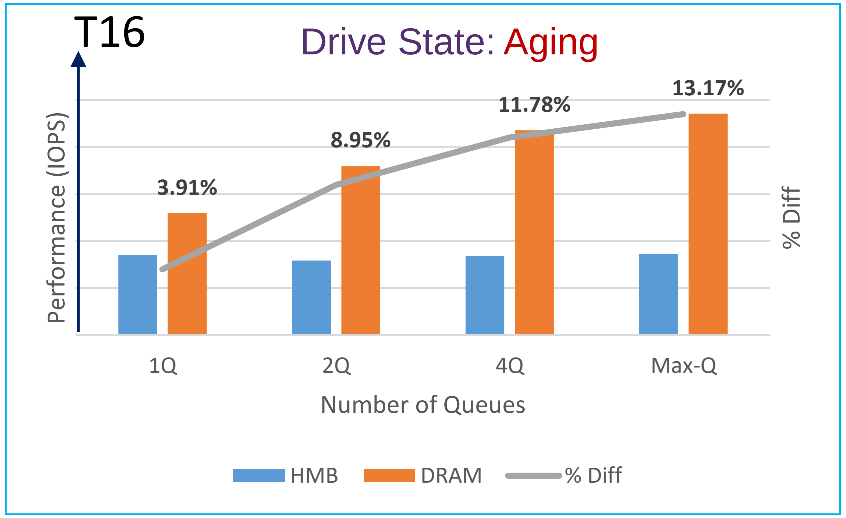

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T1-QD32)

- Low density shows better DRAM performance at low Q configurations in Normal & Sustain states.

- Low density shows uniform performance at all Q configurations in sustain state.*

***NOTE**: T16 has similar observation in Sustain state only.

Density: High

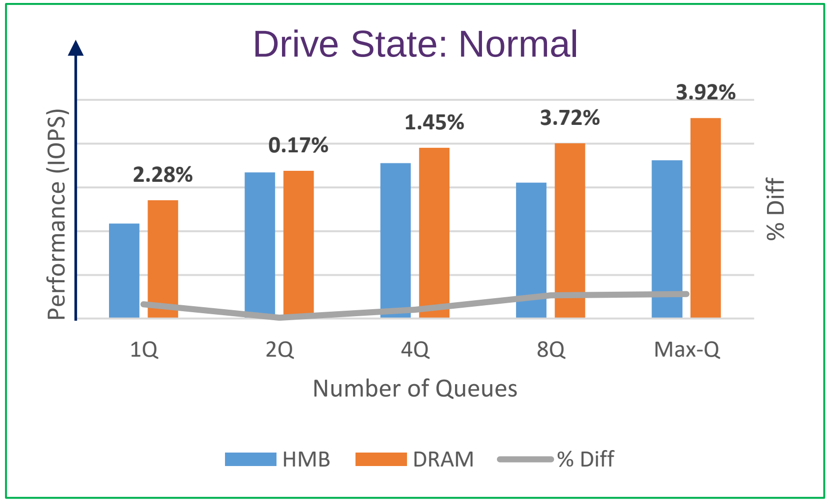

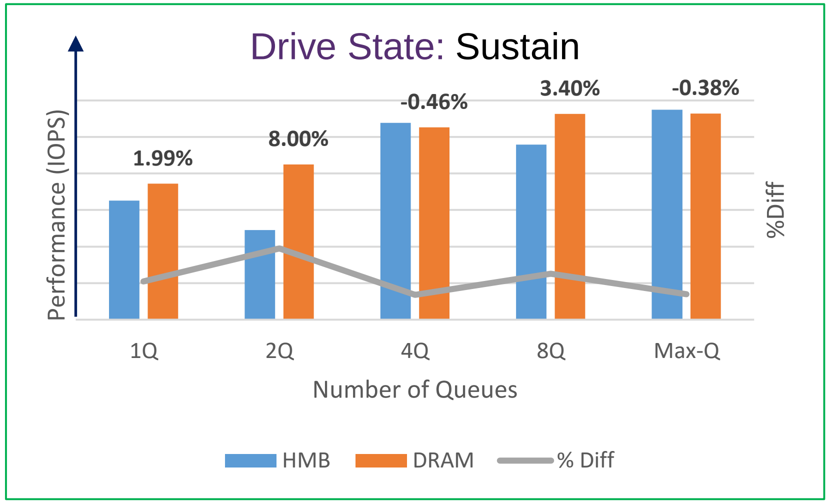

Density: Low

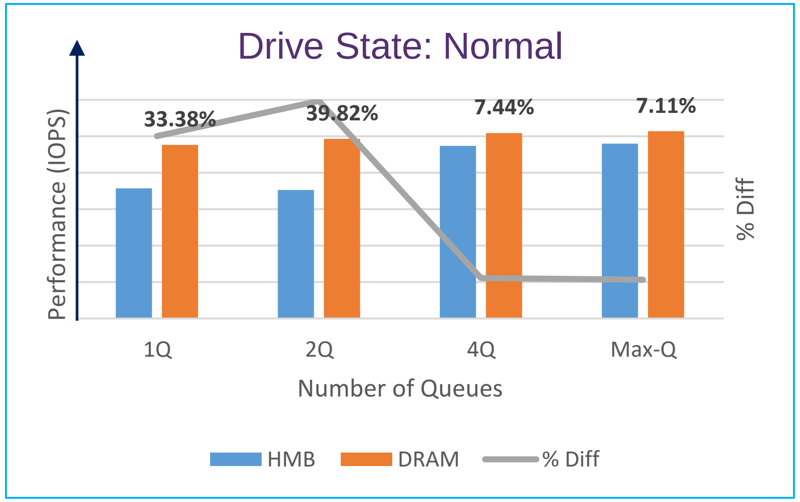

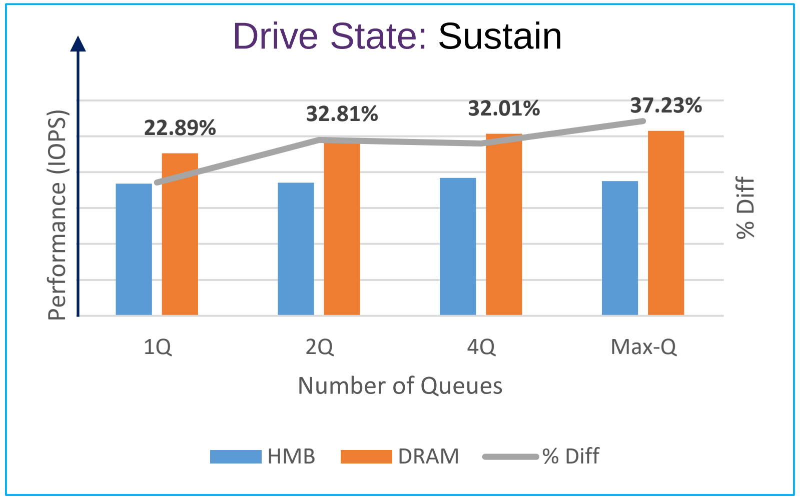

Disclaimer: Results are based on Custom test driver only

HMB Tuning Comparison

Max HMB : Allocation same as Preferred HMB size (HMPRE)
Half HMB : Half the preferred HMB size allocated
Min HMB : Allocation same as Minimum HMB size (HMMIN)

Workloads: Sequential Write (128K-T1-QD32)

- HMB Size variation didn't impact performance across all states & Q-Configuration.
- Aging state performance drastically reduced across all Q-Configuration & badly hit in 1Q-Configuration.

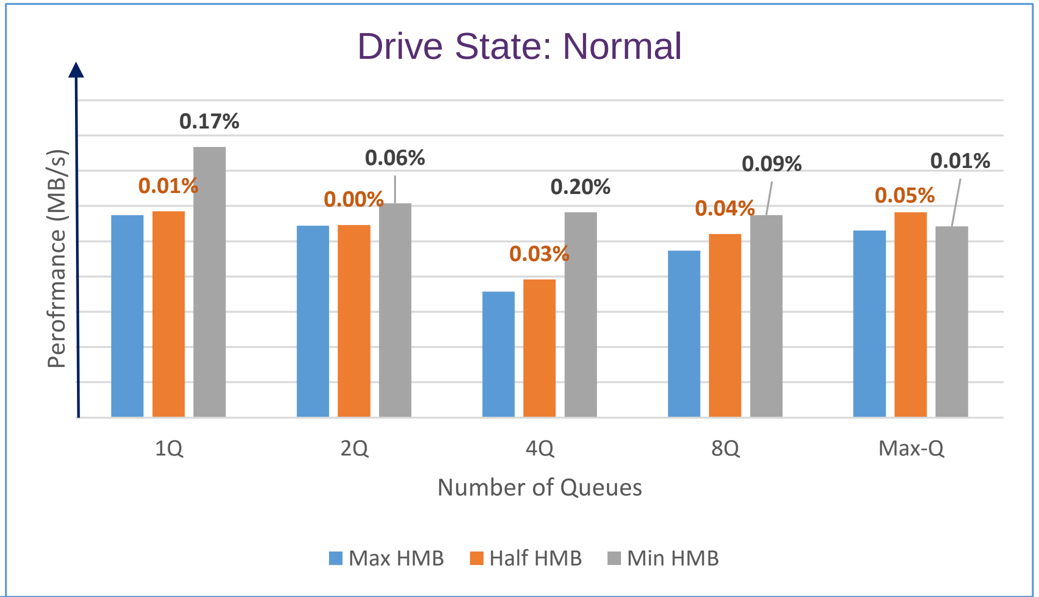

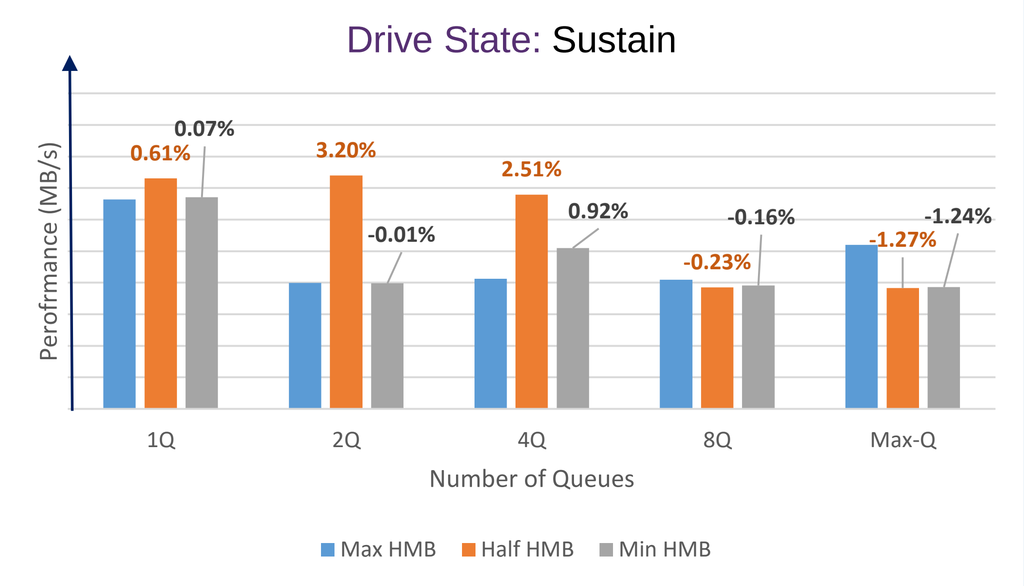

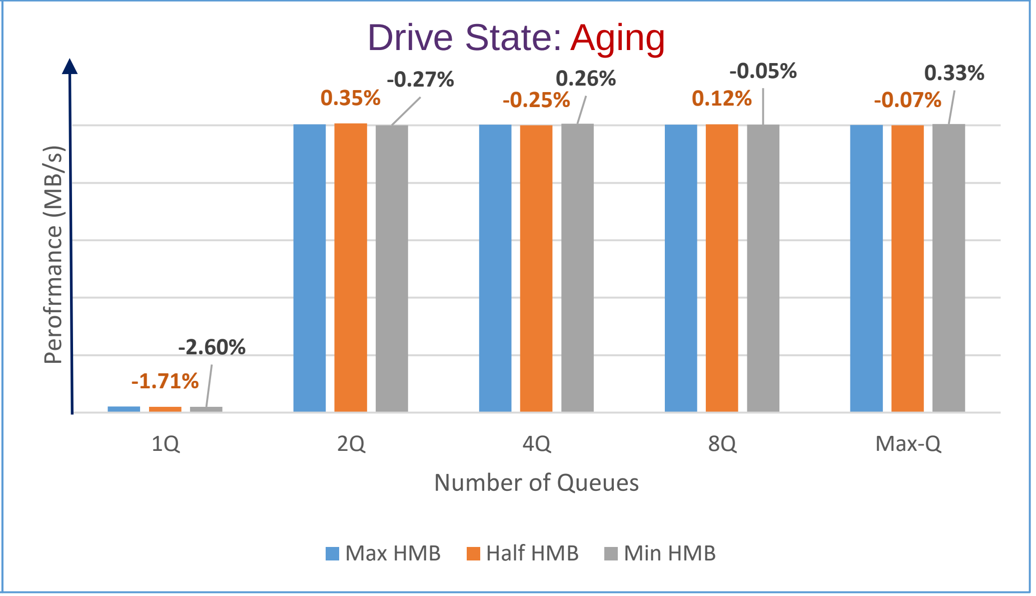

Disclaimer: Results are based on Custom test driver only

Workloads: Sequential Read (128K-T1-QD32)

- HMB Size variation didn't impact performance across all states & Q-Configuration.

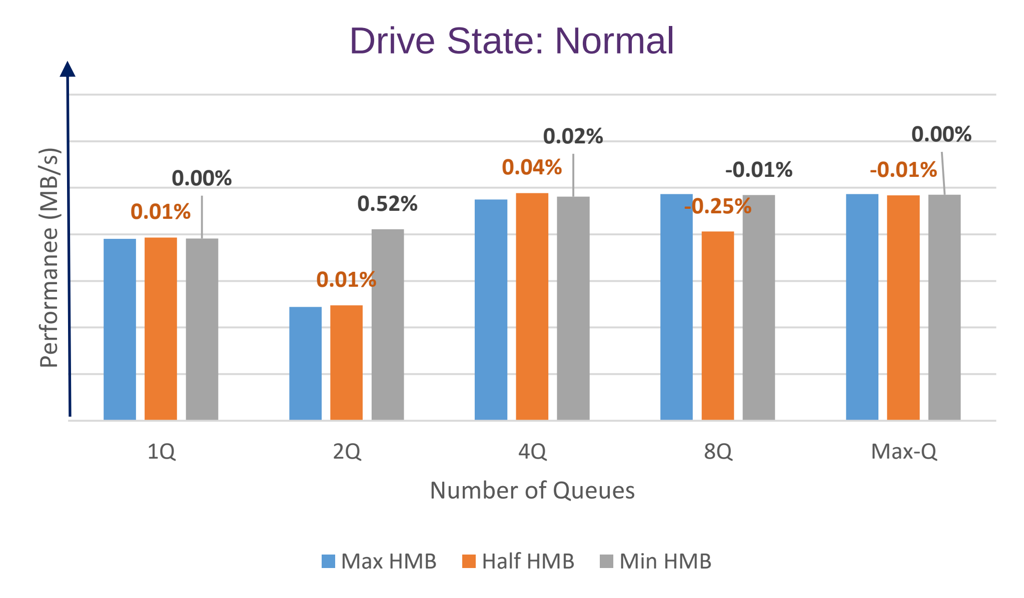

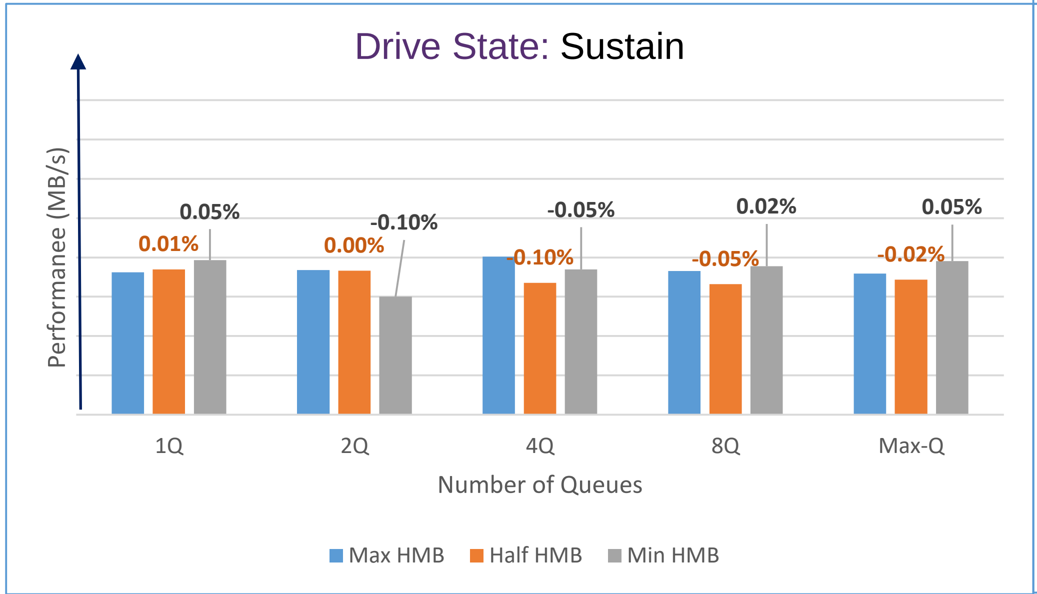

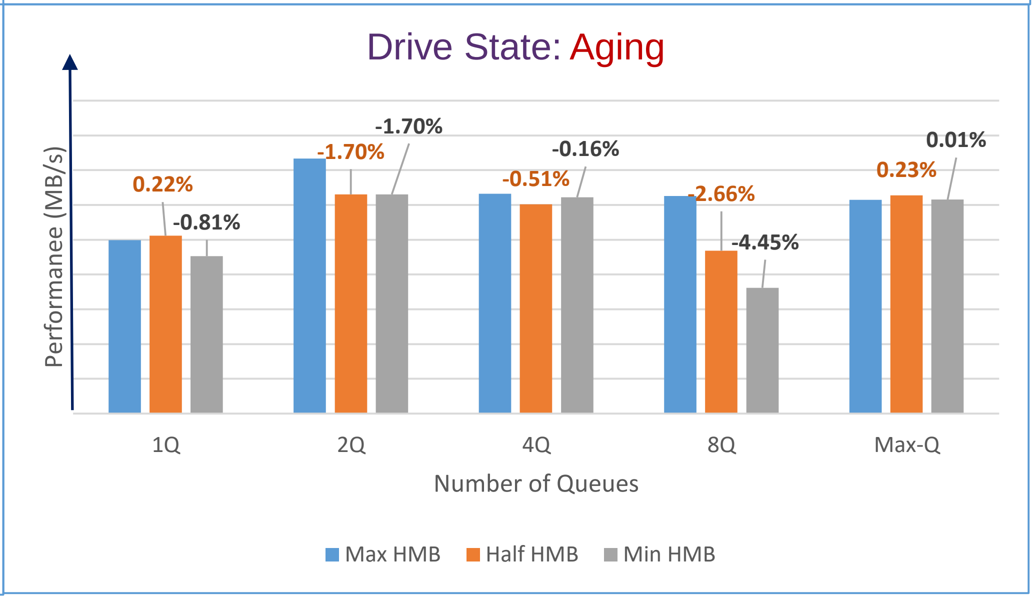

Disclaimer: Results are based on Custom test driver only

Workloads: Random Read (4K-T1-QD32)

- Performance better for Max HMB allocation in majority cases.
- In Aging Half HMB size allocation is more consistent across all Queue configurations

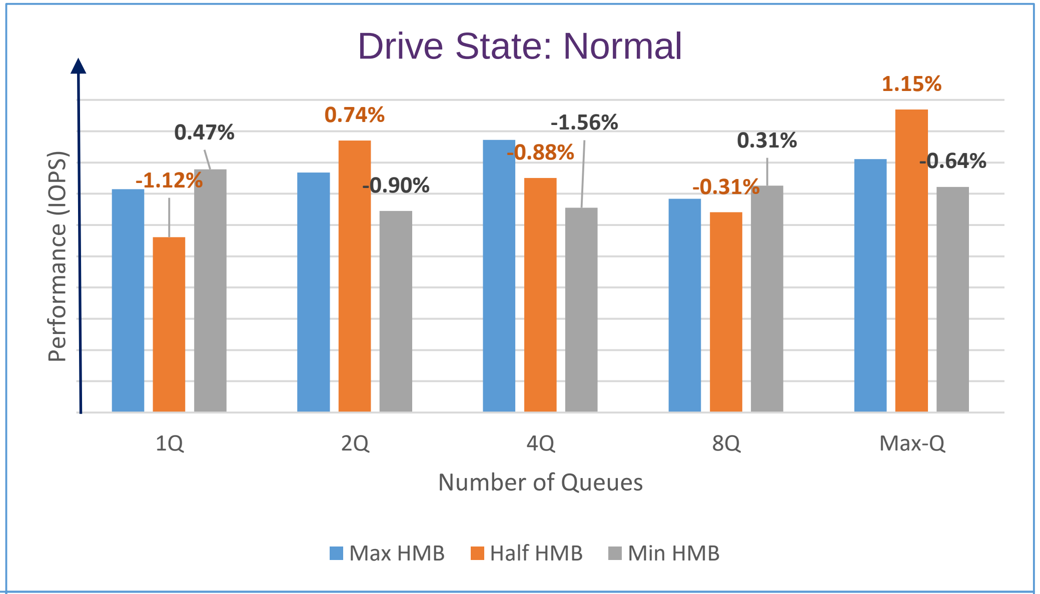

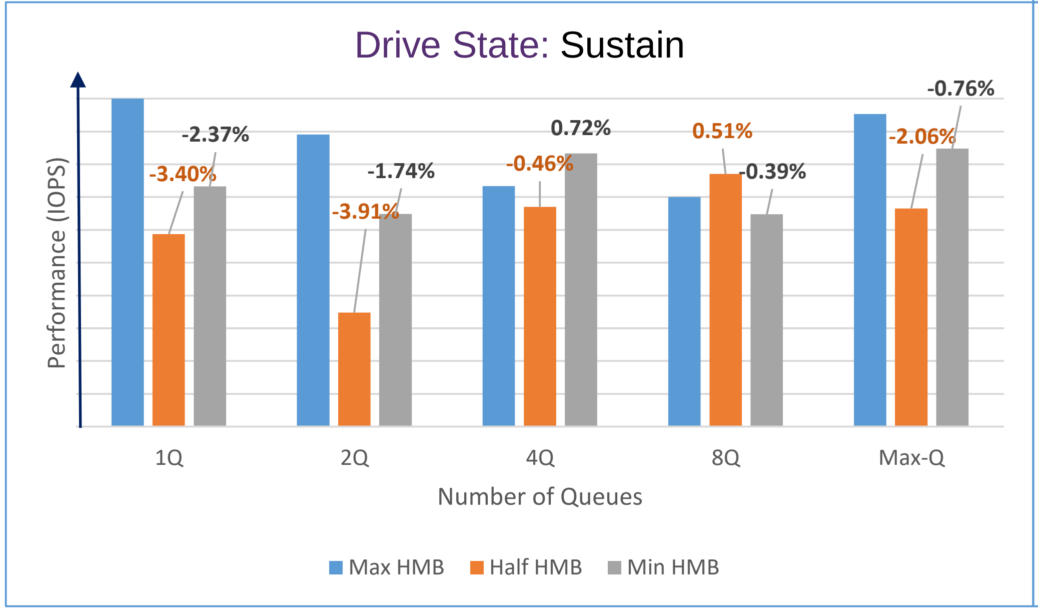

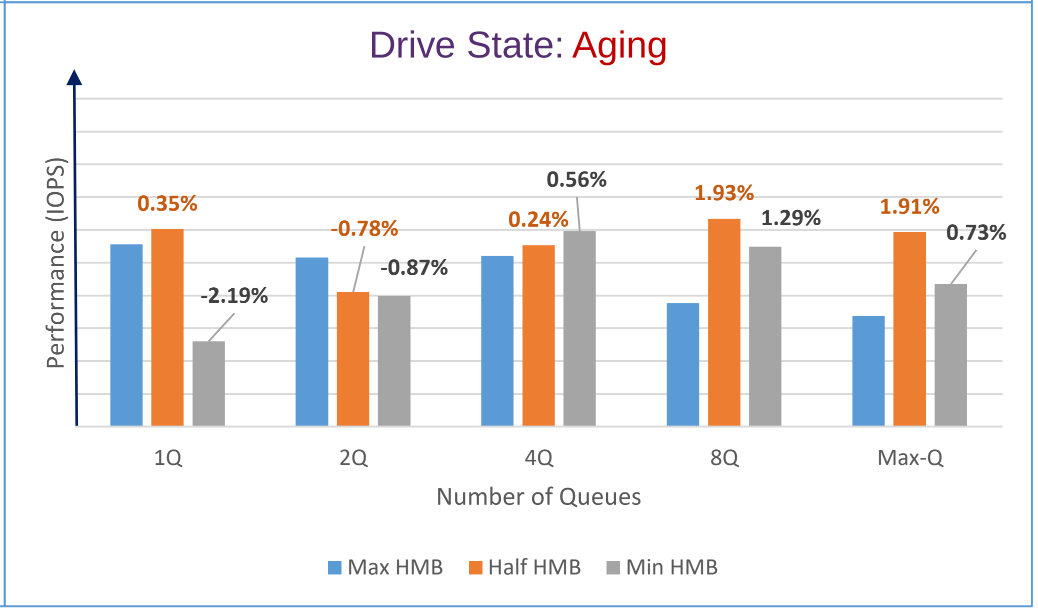

Disclaimer: Results are based on Custom test driver only

Workloads: Random Read (4K-T16-QD32)

- HMB allocation Half the preferred size performed better in Normal & Aging states .
- Half HMB Size allocation worked relatively poor in sustain state across all Queue configurations

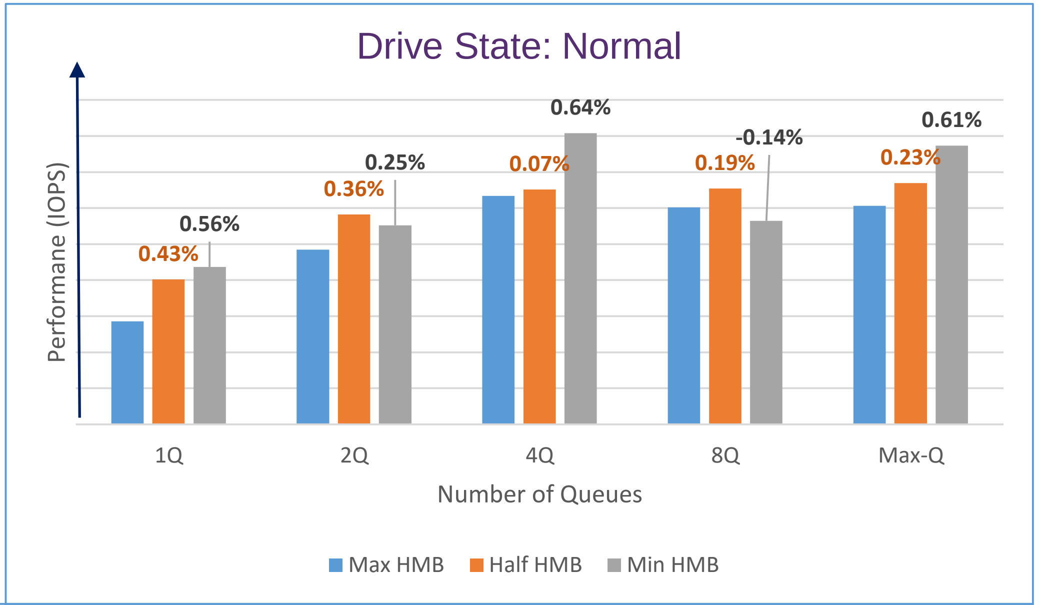

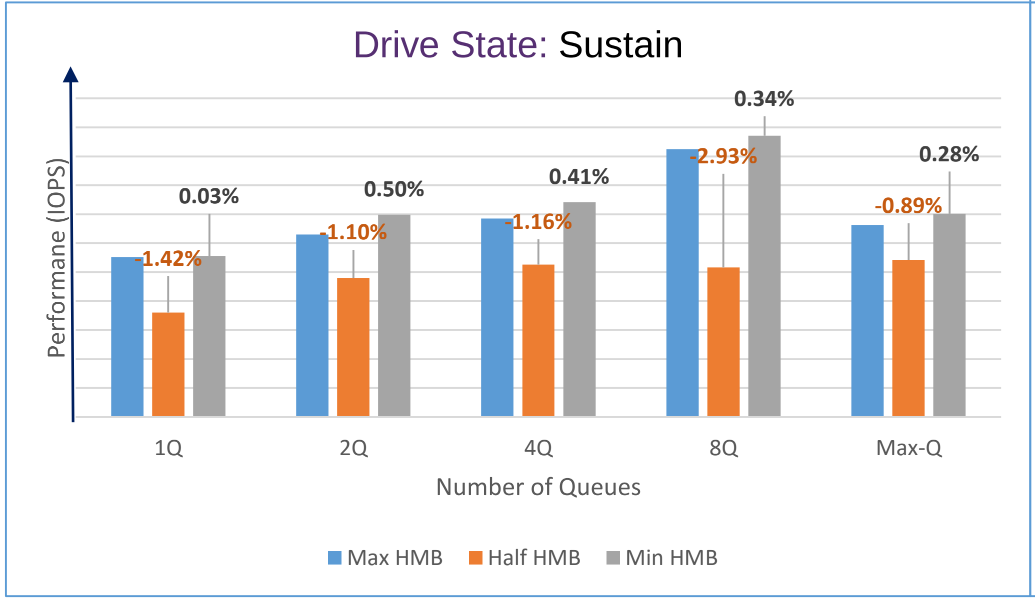

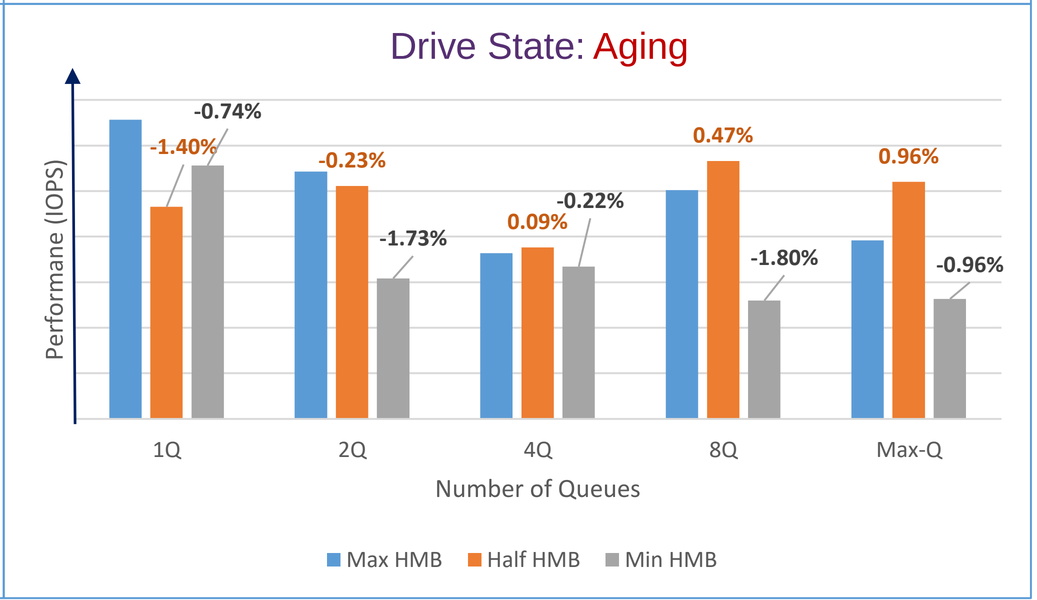

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T1-QD32)

- Normal & Sustain state performance has no major impact due to HMB Size variation
- Min HMB allocation shows Aging performance decreased with increase in Queues configured.
- Max HMB allocation shows Aging performance increased with increase in Queues configured.

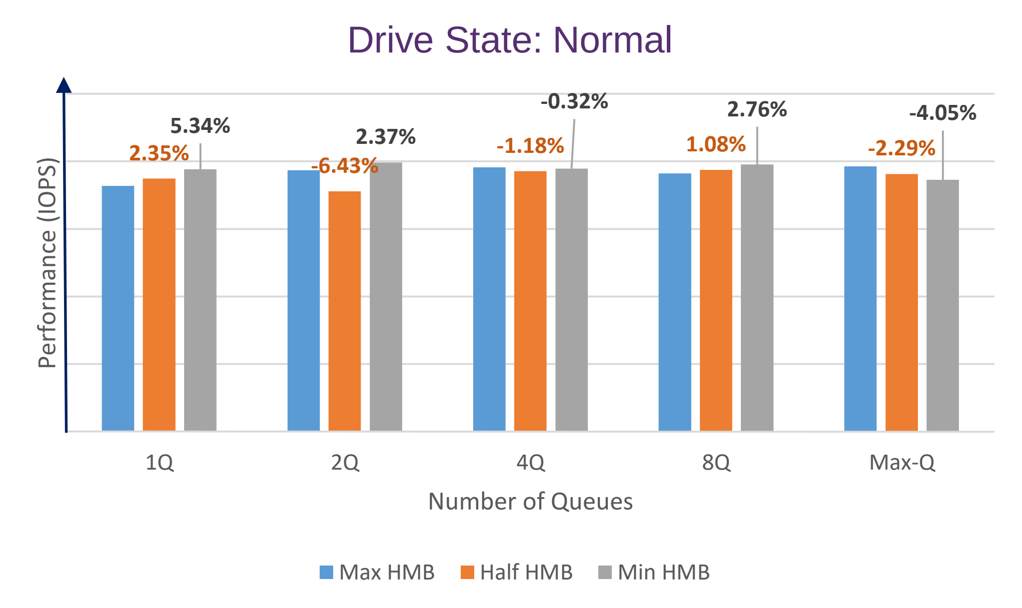

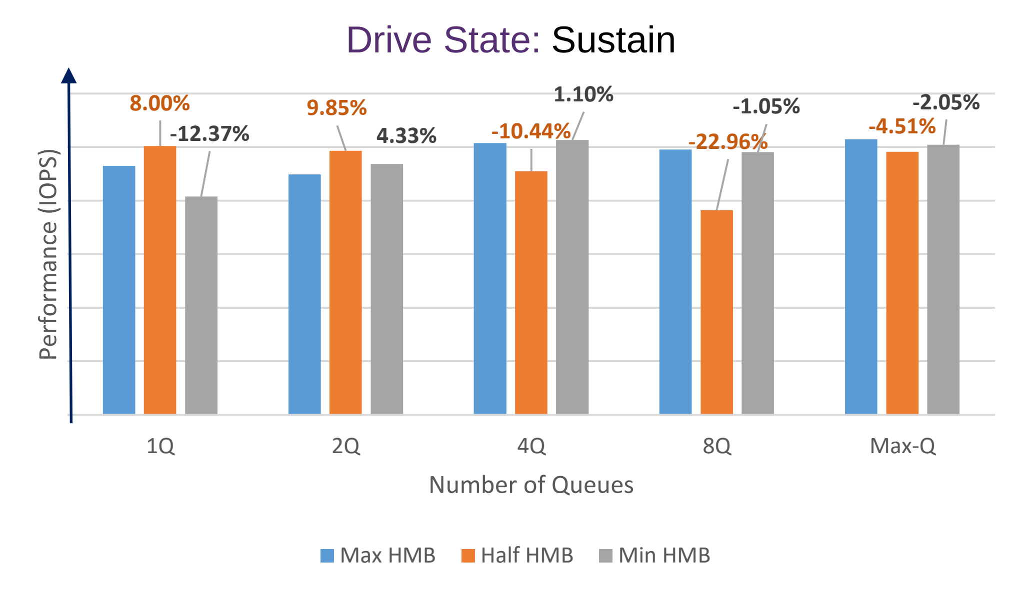

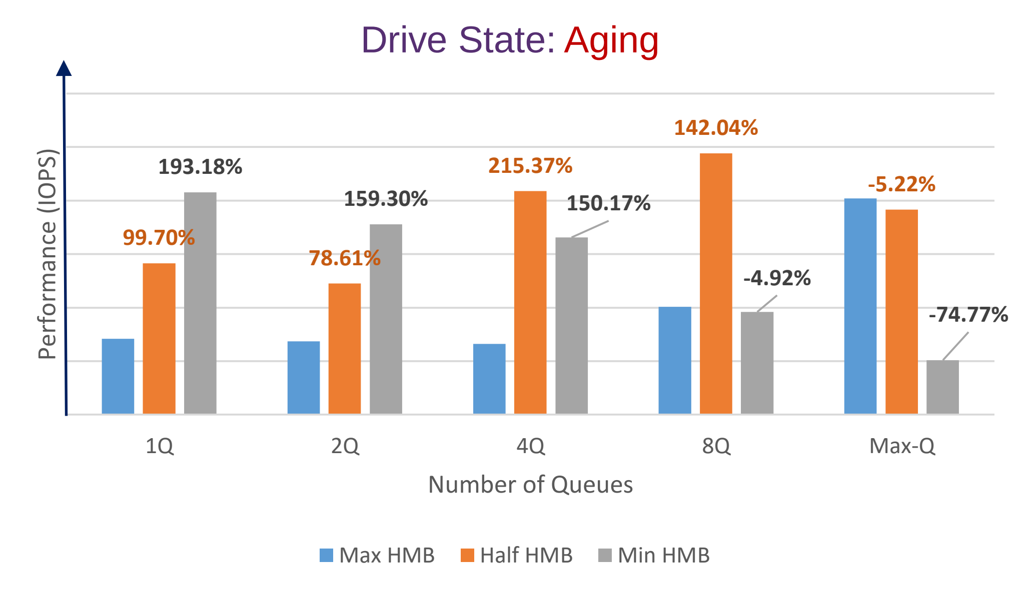

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T16-QD32)

- Min HMB size allocation shows better results in majority of configurations & drive states
- Half HMB size allocation shows low performance in majority of cases in sustain & aging states

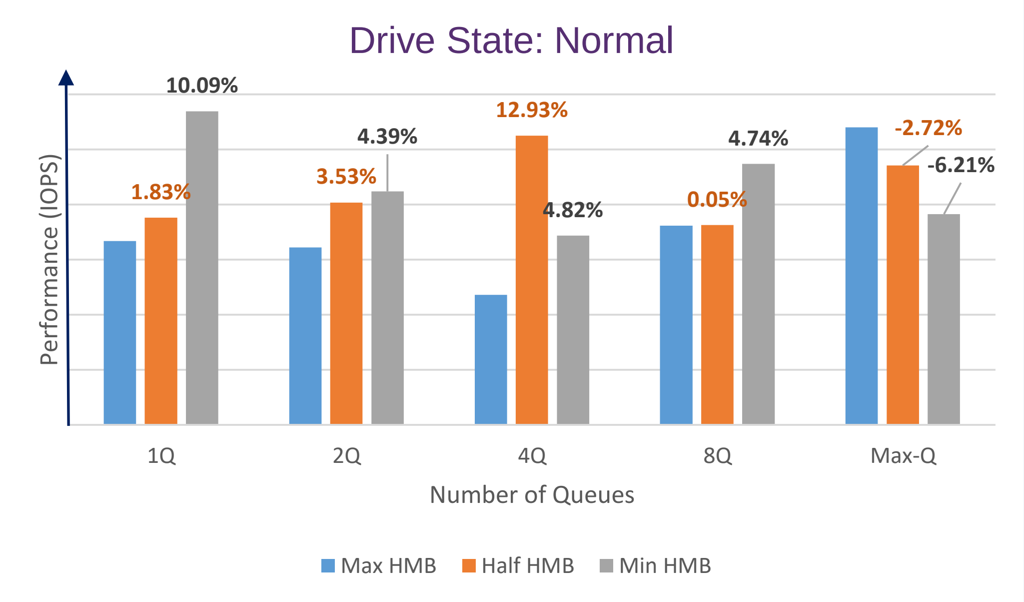

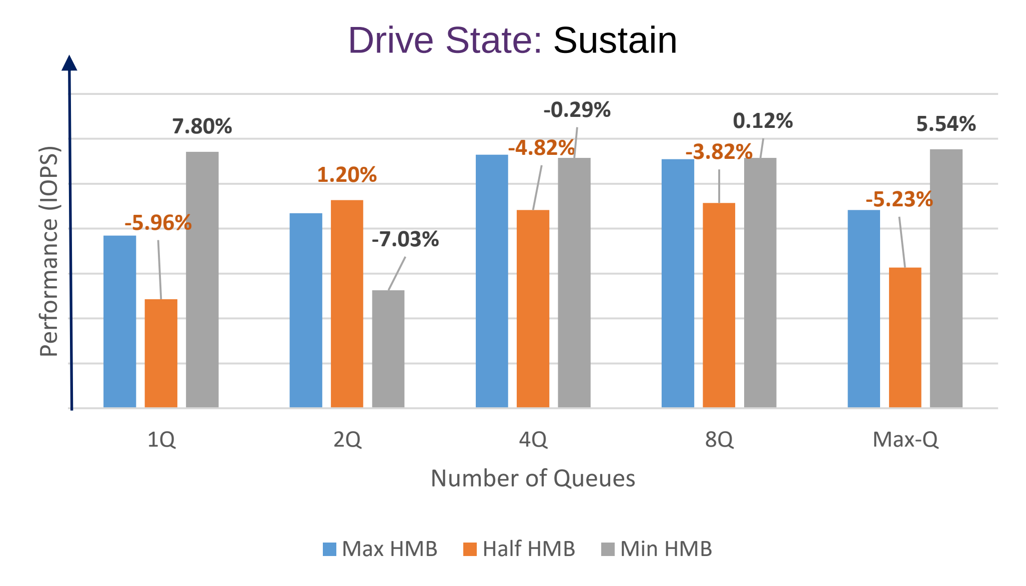

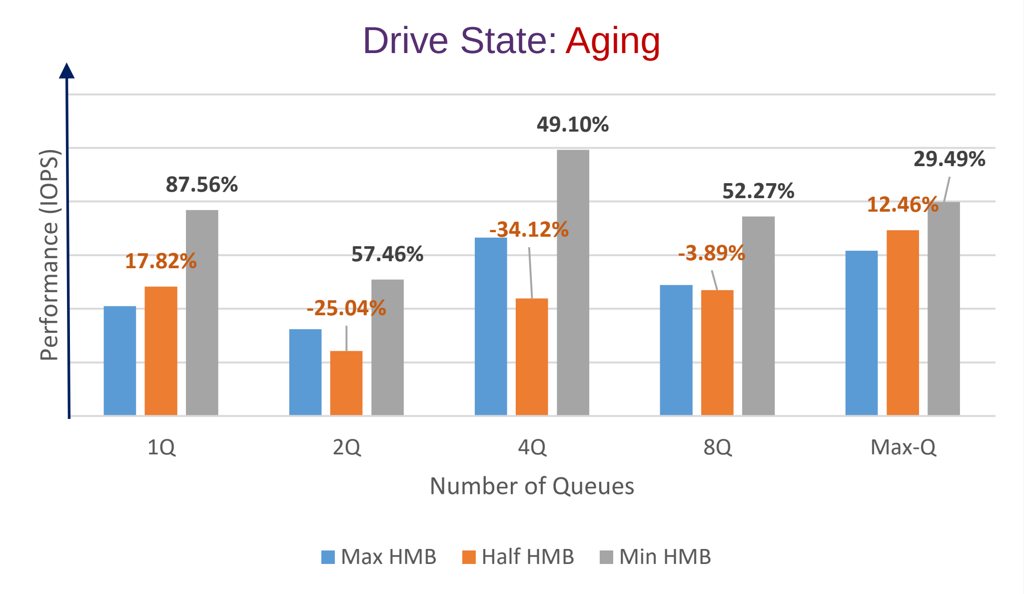

Disclaimer: Results are based on Custom test driver only

Multi-Density Comparison

Workloads: Random Read (4K-T16-QD32)

- Low density show not much difference across HMB sizes in sustain state

Density: High

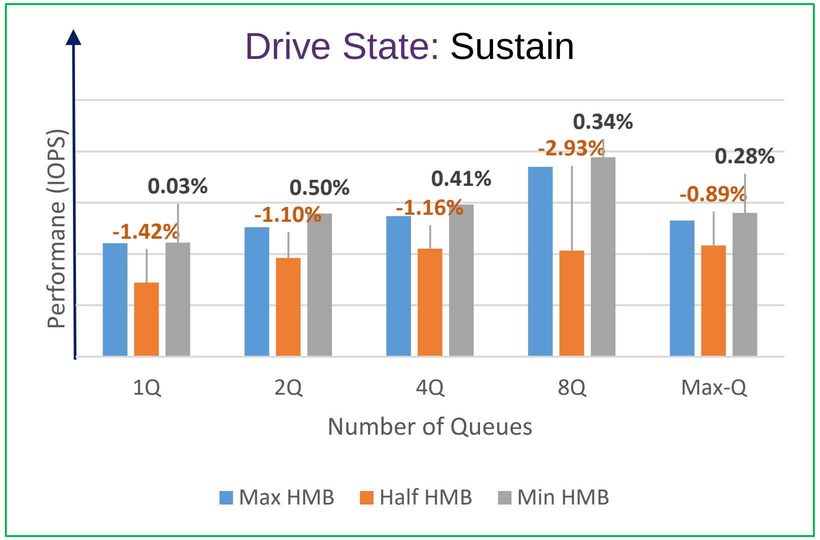

Density: Low

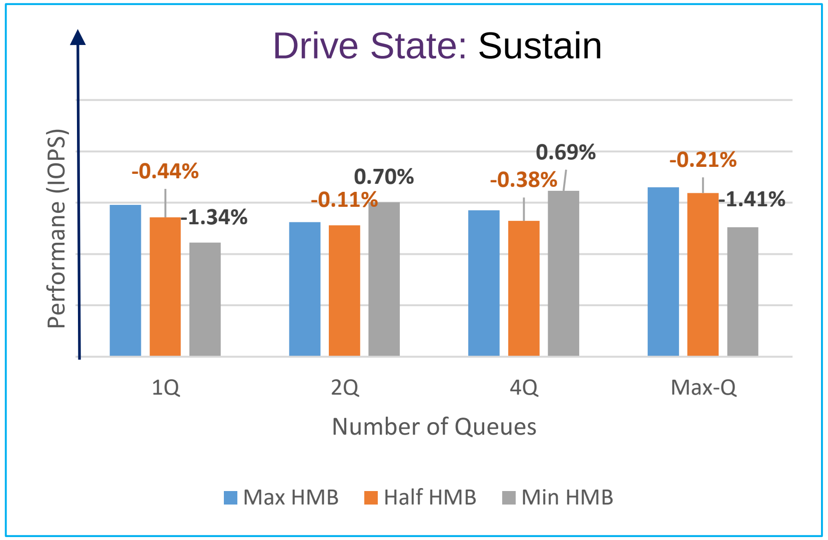

Disclaimer: Results are based on Custom test driver only

Workloads: Random Write (4K-T1-QD32)

- For low density, Half HMB Size has better values for low Q configuration in normal state
- For low density, Min HMB Size has been more consistent at aging across all Q configurations

Density: High

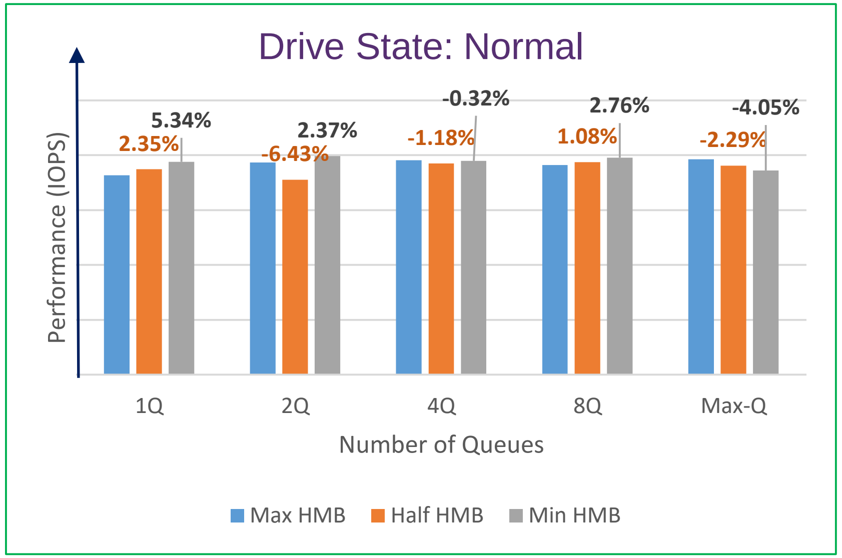

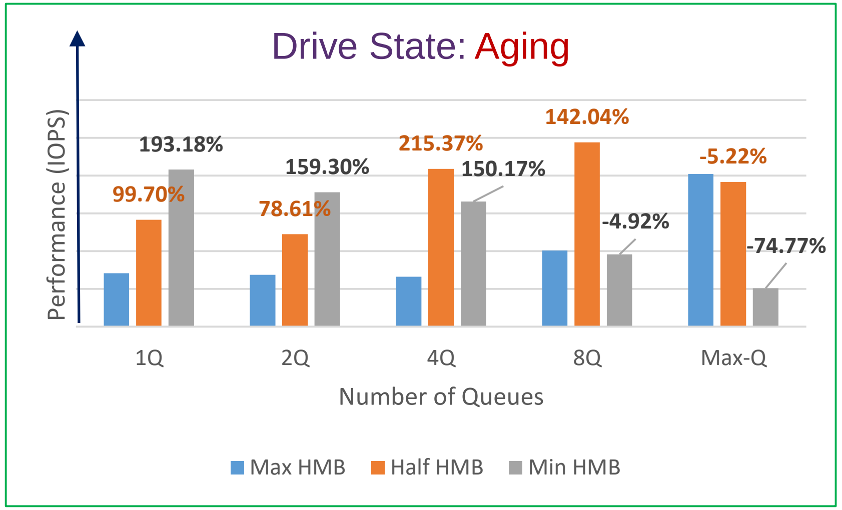

Density: Low

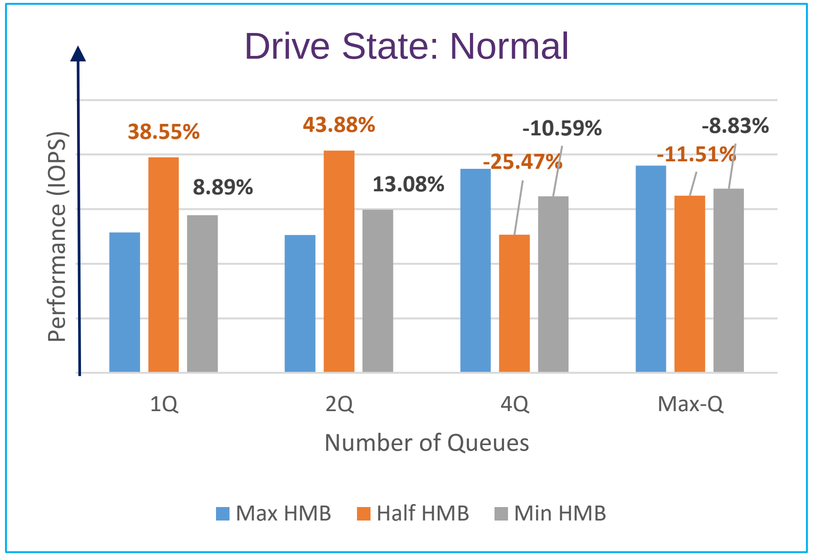

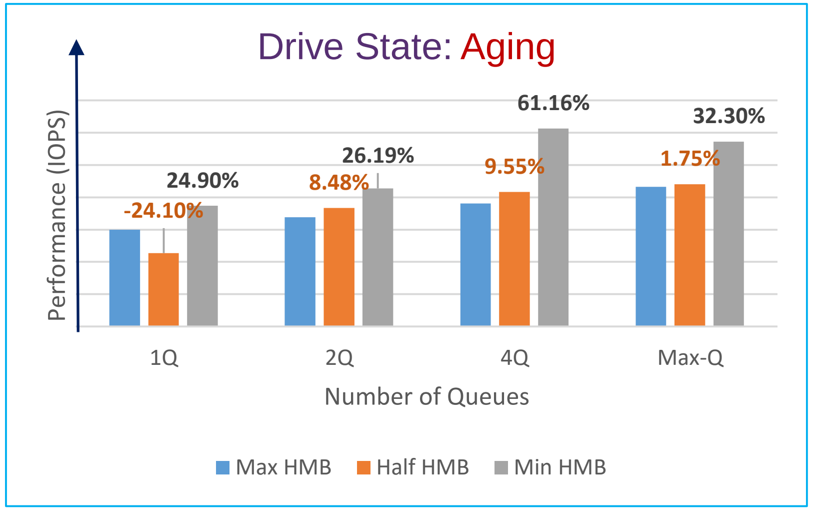

Disclaimer: Results are based on Custom test driver only

Conclusion

Conclusion

- Optimal Queue choice across different HMB sizes & drive states: 4Q configuration
- Max & Half Preferred Size Allocation for HMB worked better in most of workloads & Drive states across densities
- Max-Q configuration worked best with lowest gaps possible between DRAM & HMB across drive states

Work Loads	HMB Suitability	Recommendation
Random Read	✓	Suits better for Larger densities
Random Writes	$\cong$✓	Competitive enough till Aging
Sequential Read	✗	
Sequential Writes	✗	

Thank You