SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Exploring Performance Paradigm of HMB NVMe SSD's


Pradeep SR

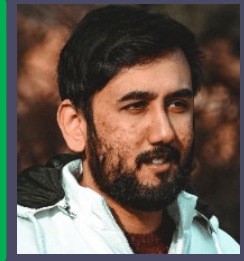
Ranjith T

Presented by: Pradeep SR, Ranjith T
Samsung Electronics

Disclaimer


Pradeep SR Ranjith T

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


Pradeep SR


Ranjith T

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

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

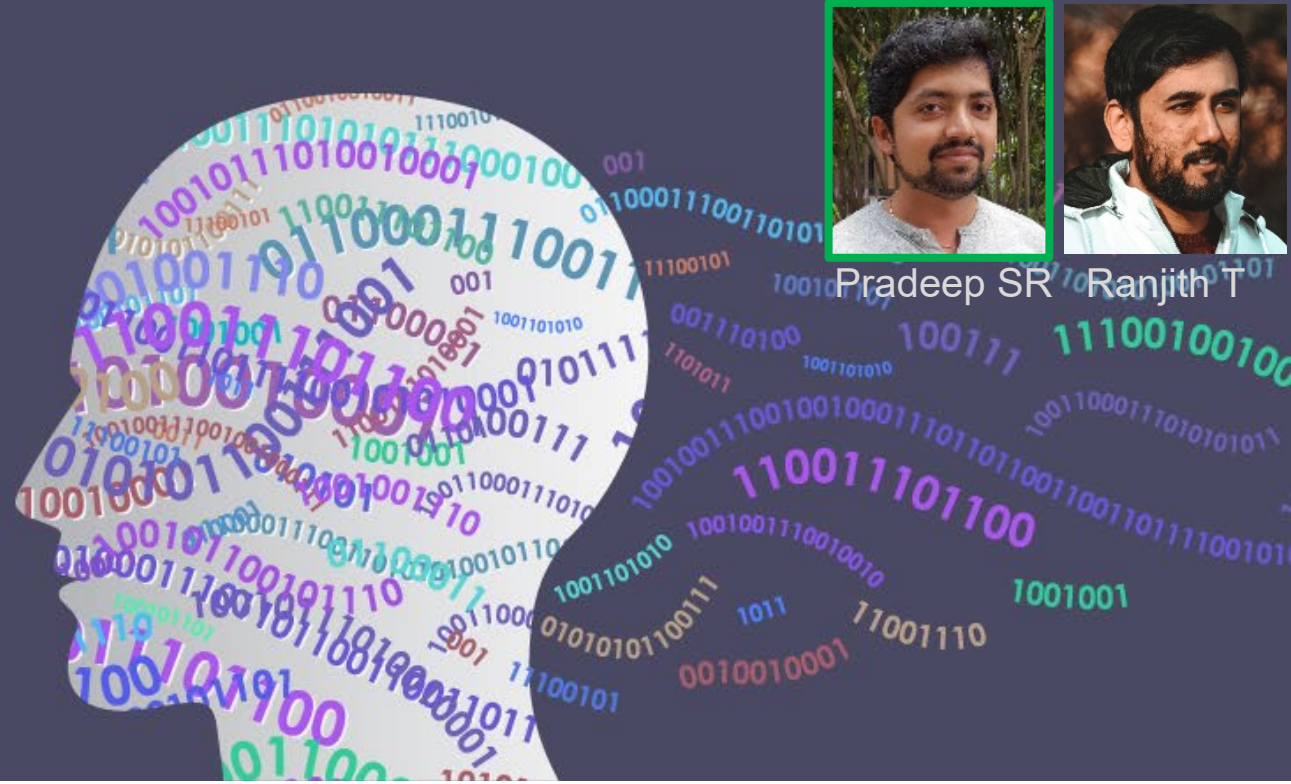

Pradeep SR Ranjith T

HMB Overview

What & Why of HMB?

Pradeep SR Ranjith T

- Host Memory Buffer (HMB)

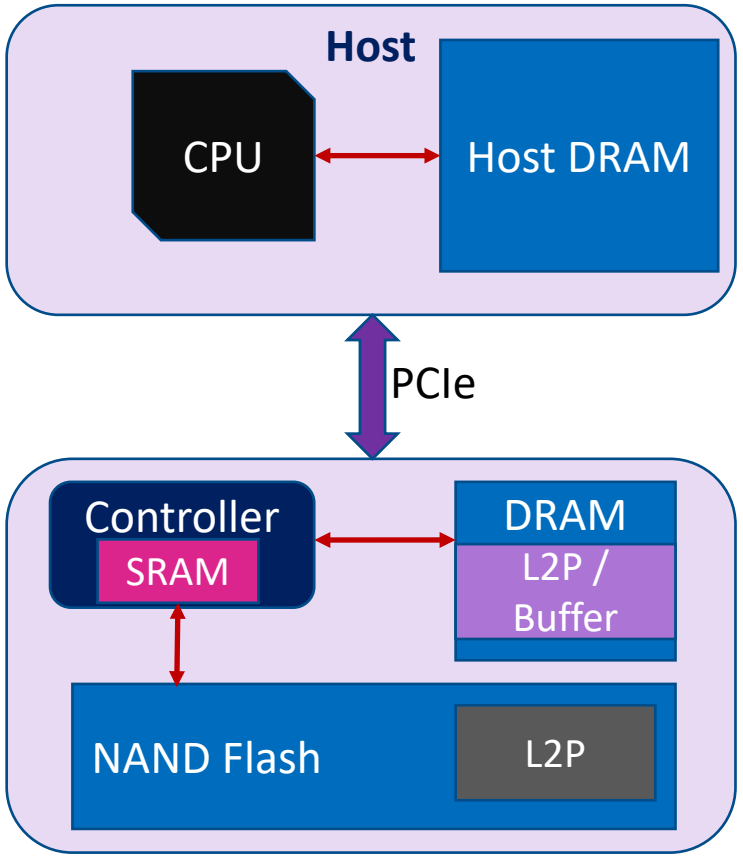

Conventional SSD w/ DRAM

HMB Motivation:

- Cost Reduction
- Reduced Size
- Market Adoption

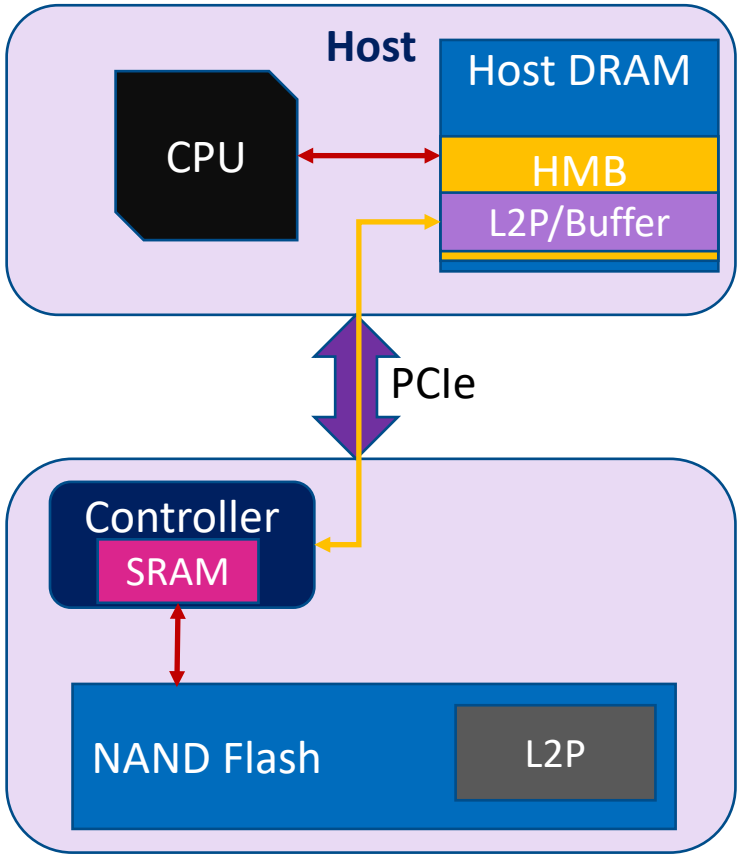

Value SSD w/ HMB

NVMe HMB Feature


Pradeep SR


Ranjith T

- 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
- OS & Device Support
  - Windows 10 OS & Above
  - *Linux 5.12,Ubuntu 20.12 & Above
  - NVMe 1.2 & above compliance devices

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

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

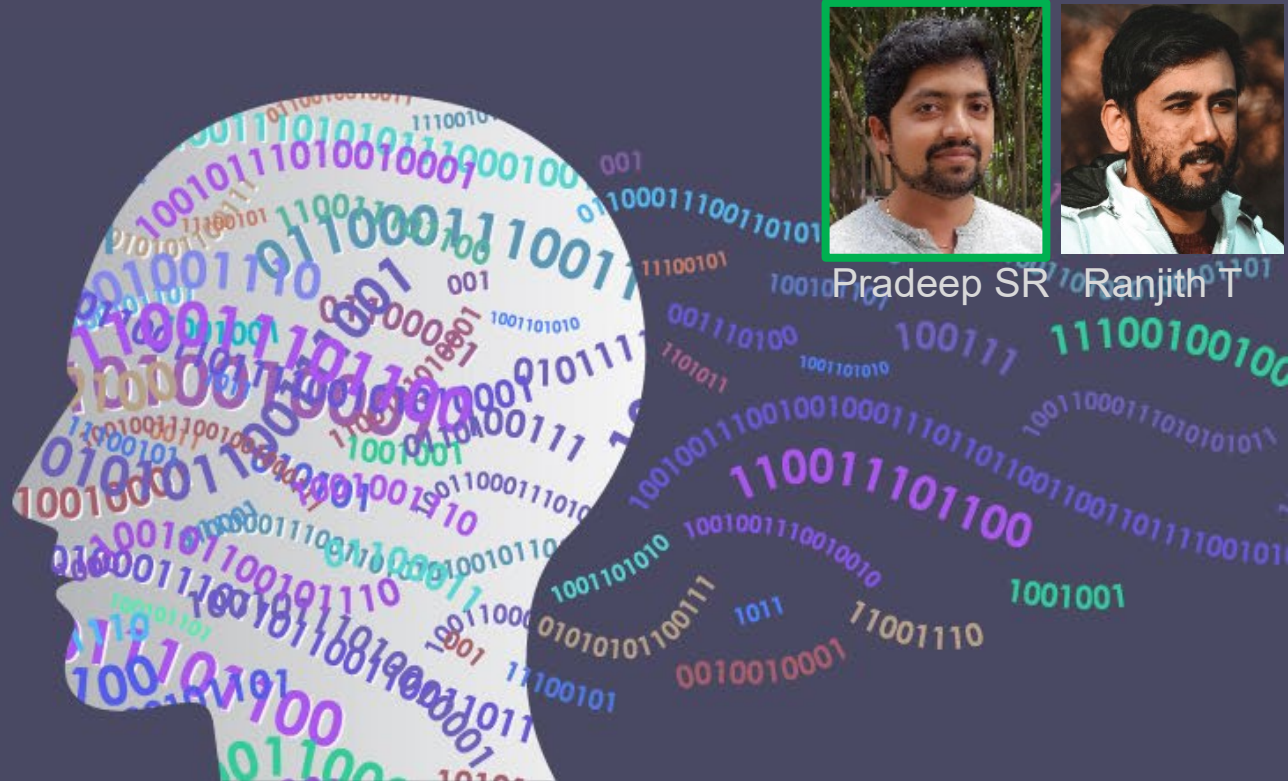

Pradeep SR Ranjith T

Performance Parameters Overview

Key Factors


Pradeep SR


Ranjith T

- Number of Queues
- Queue Depth
- IO Chunk size
- Workload types
- DRAM/ HMB Size
- Number of Threads

Experiment set up


Pradeep SR


Ranjith T

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 OS & DRAM Size	Windows 10 OS, 32GB DRAM

***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

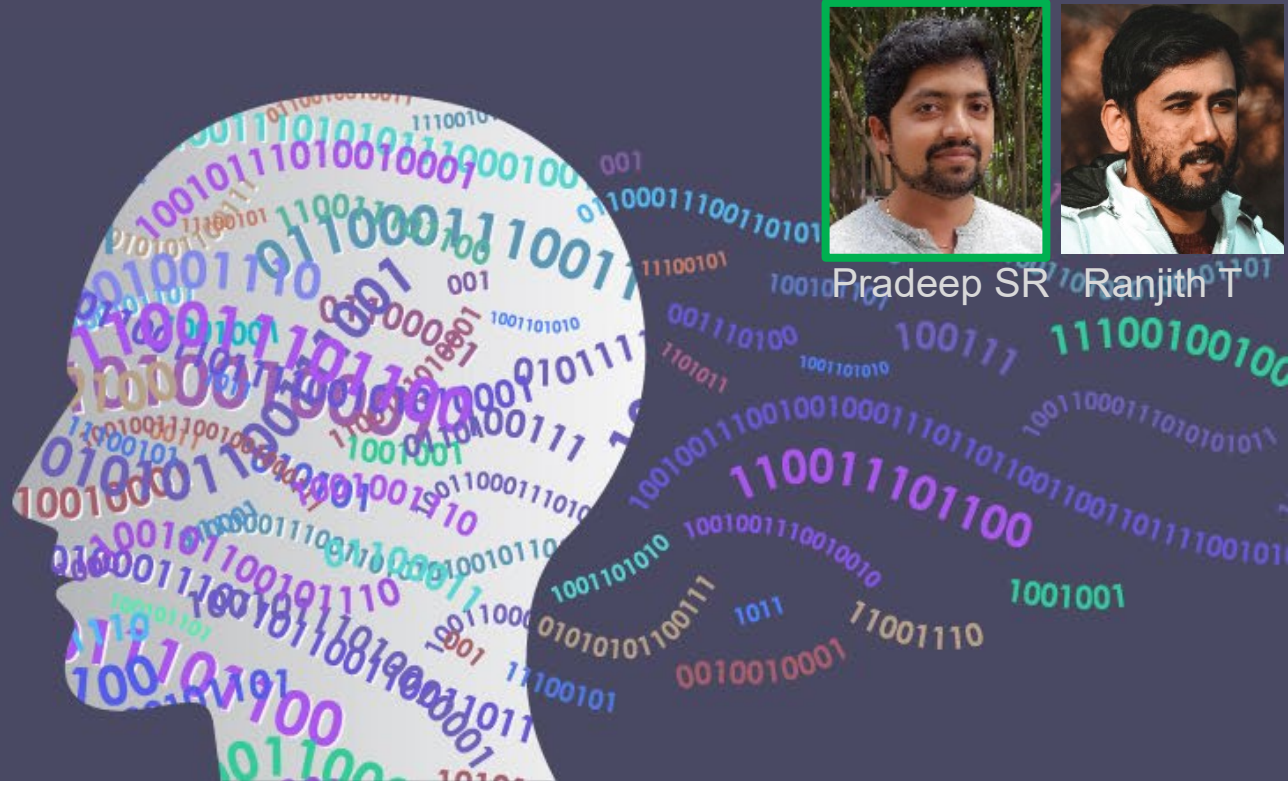

Pradeep SR

Ranjith T

HMB vs DRAM Comparison

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


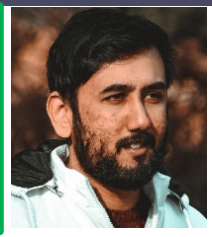

Pradeep SR Ranjith T

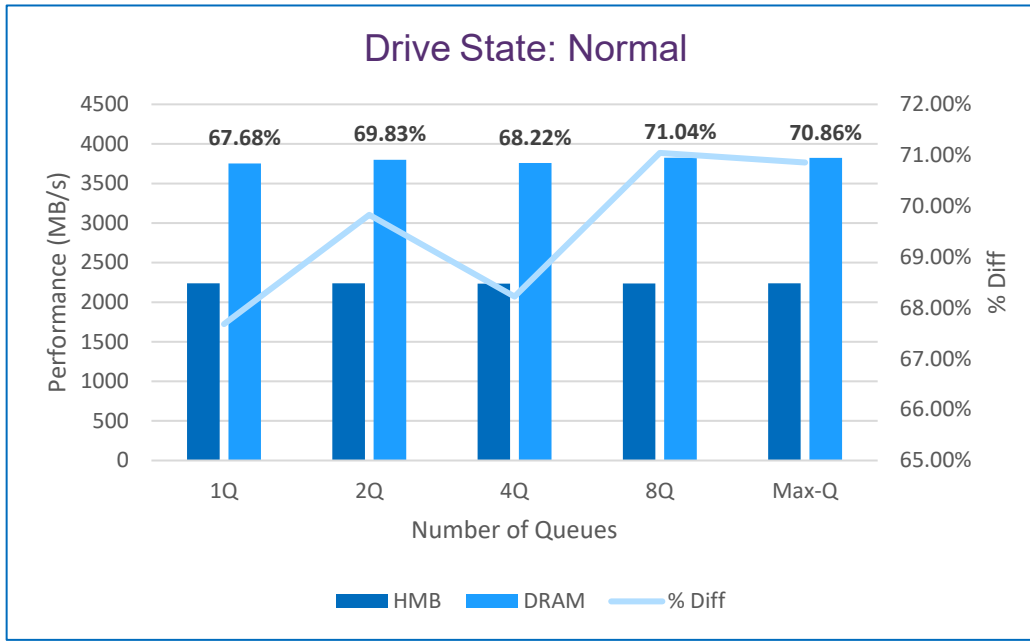

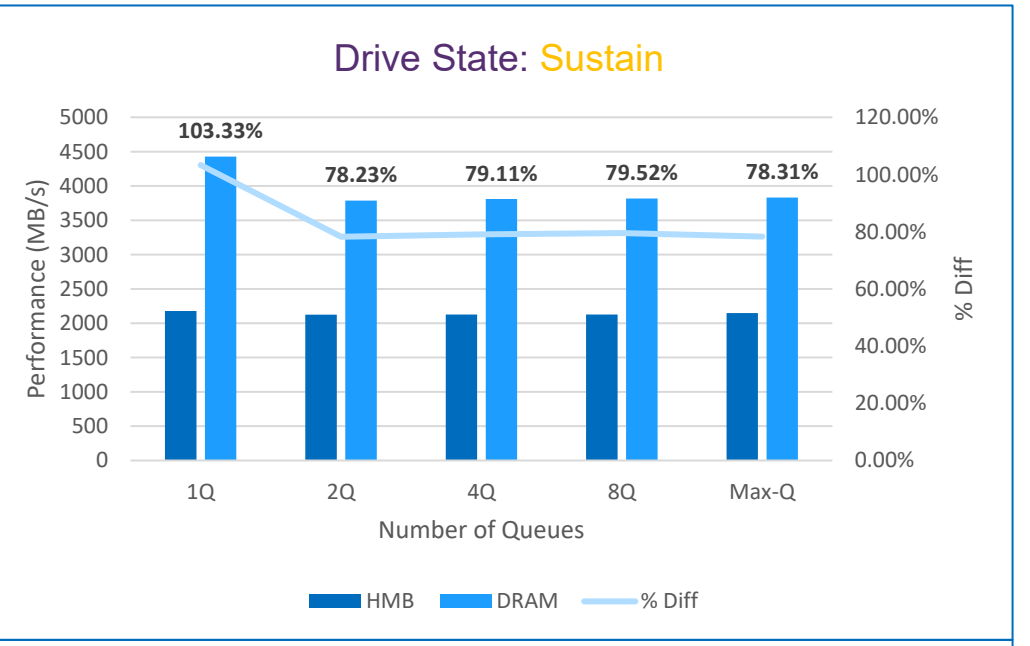

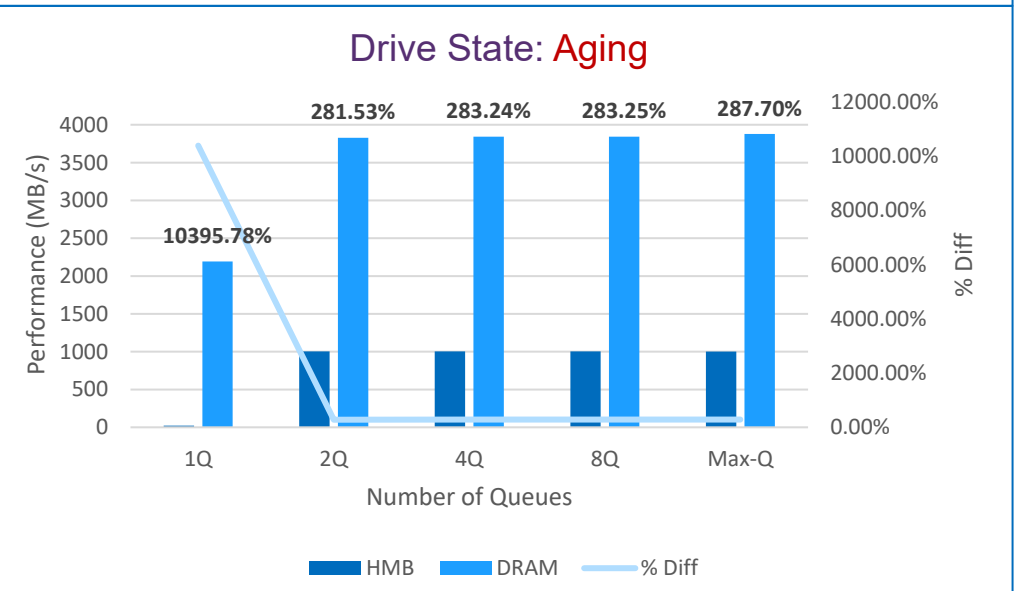

- Key Observations:
  - DRAM device leads in all Drive states & gap widens multifold at aging
  - HMB device at 1Q configuration shows lowest performance in Aging

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)


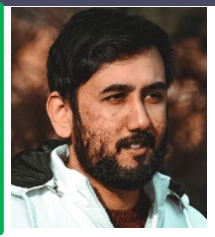

Pradeep SR Ranjith T

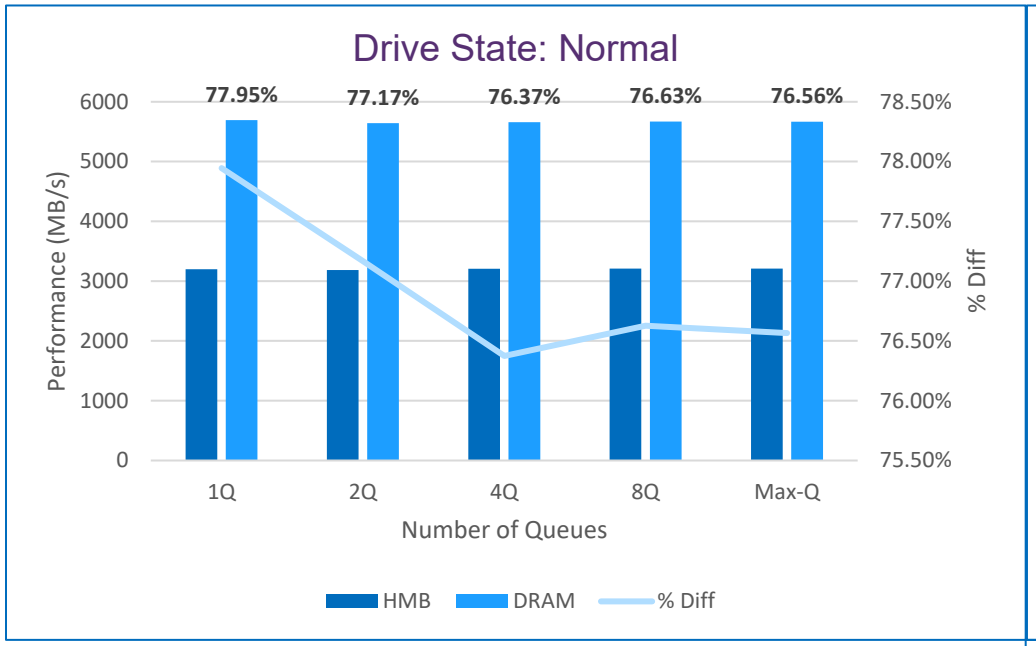

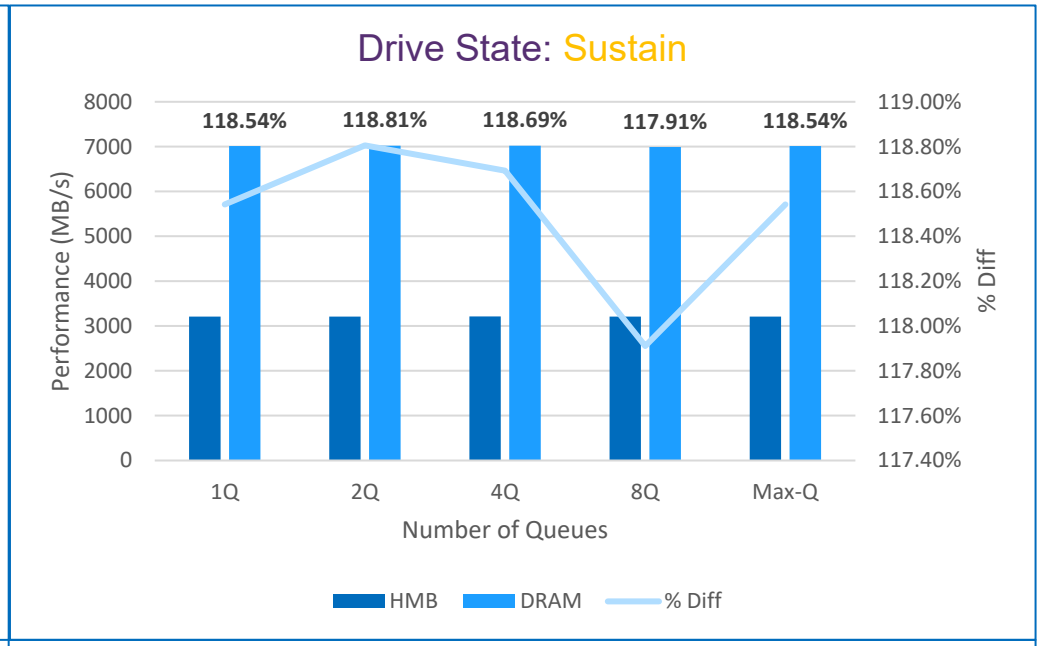

- Key Observations:
  - DRAM device leads in all Drive states & gap widens multifold at aging
  - 1Q configuration has better results for HMB device than in sequential write case in aging

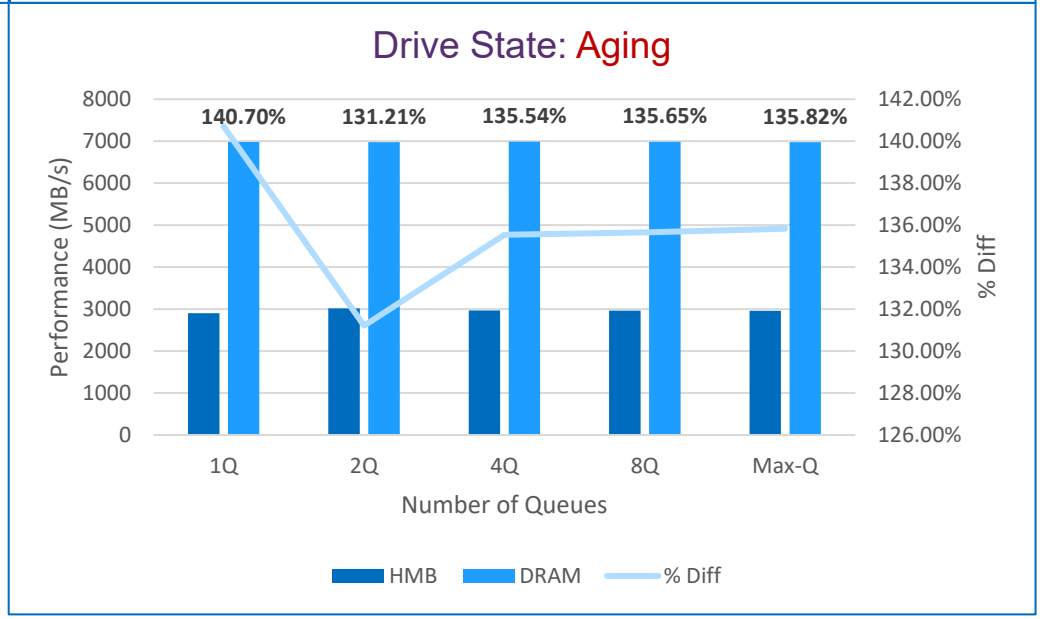

Disclaimer: Results are based on Custom test driver only

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


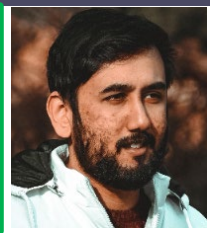

Pradeep SR Ranjith T

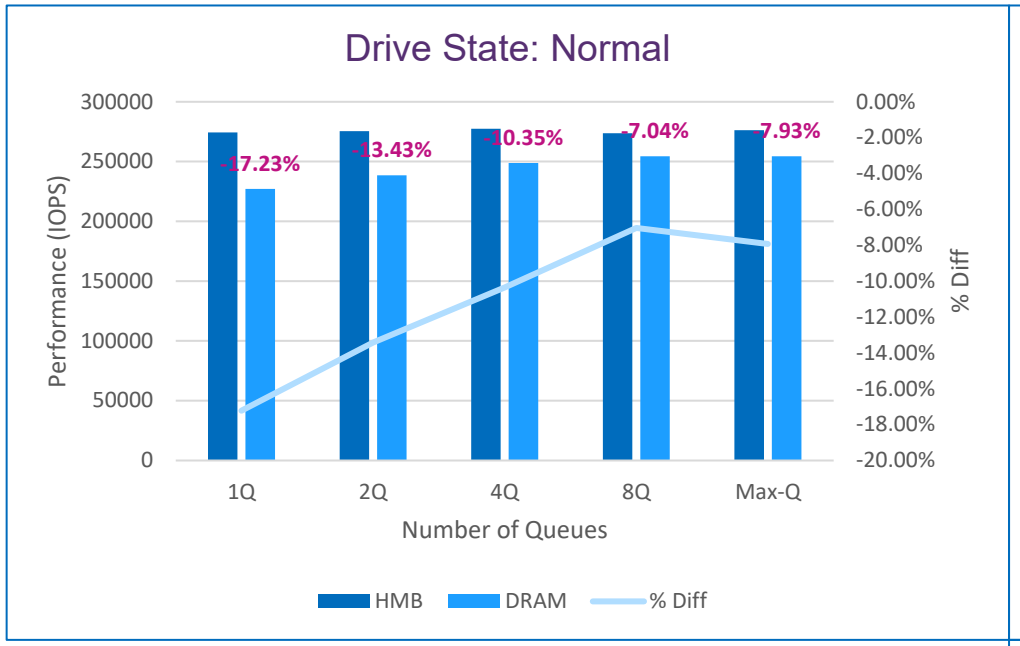

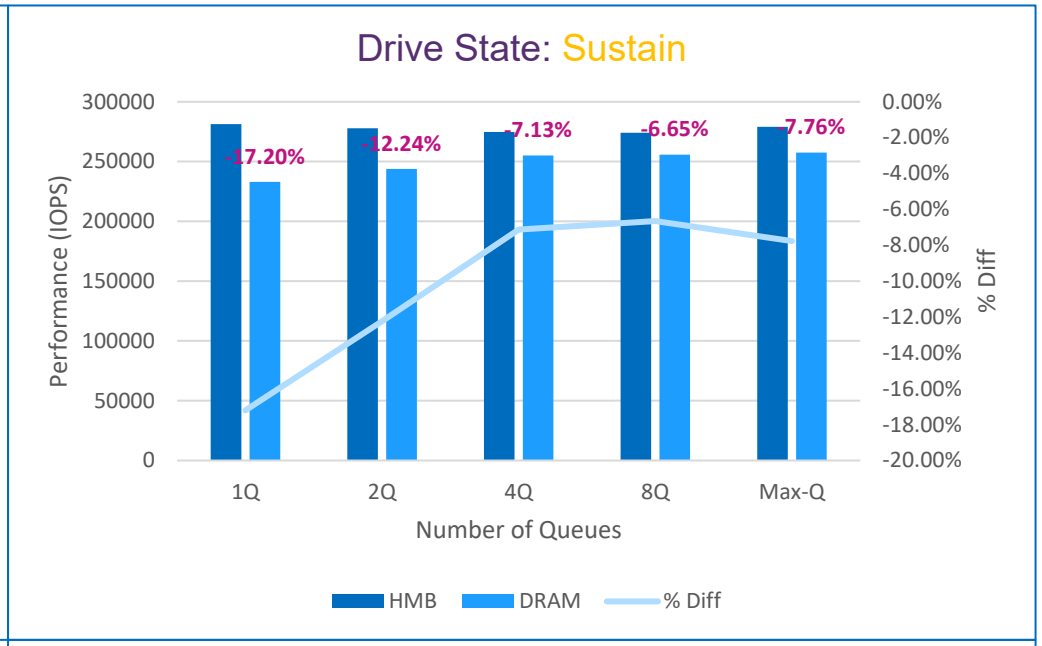

- Key Observations:
  - HMB device performed better than DRAM device in all Drive states & Queue Configurations
  - HMB is more Consistent at Aging state a well.
  - In Aging with increase in number of Q's, DRAM is able to perform better

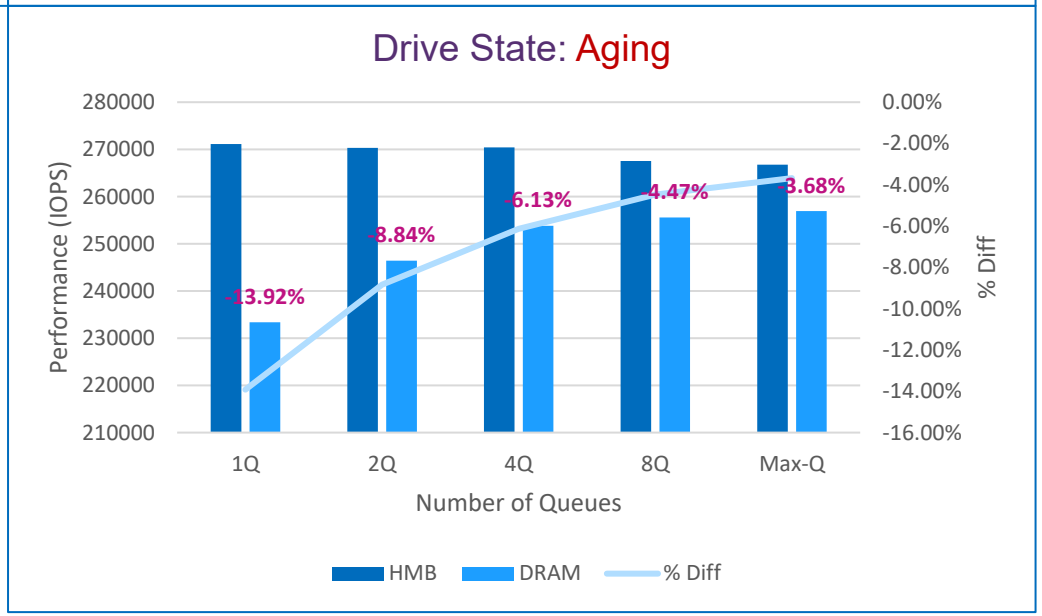

Disclaimer: Results are based on Custom test driver only

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


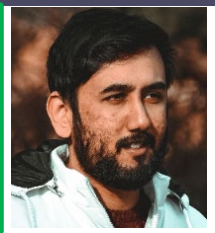

Pradeep SR Ranjith T

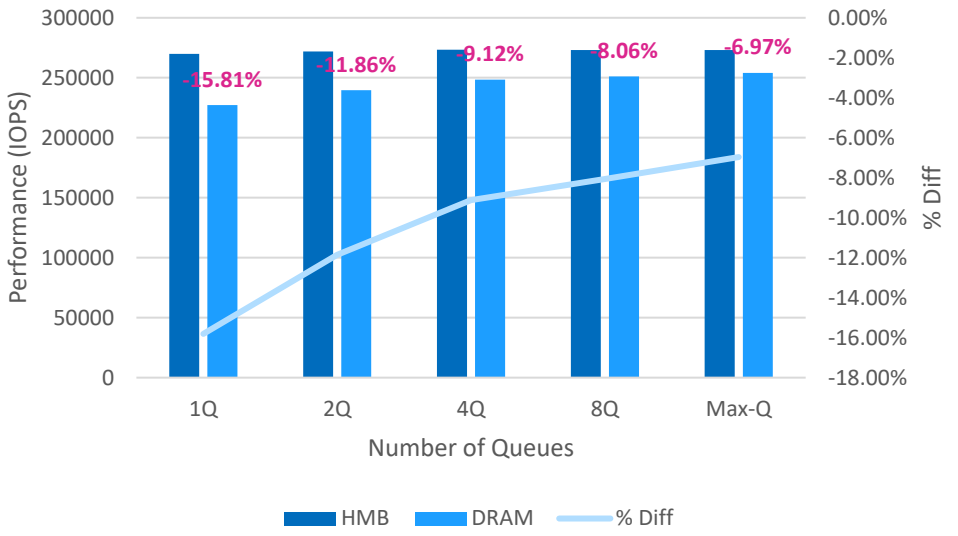

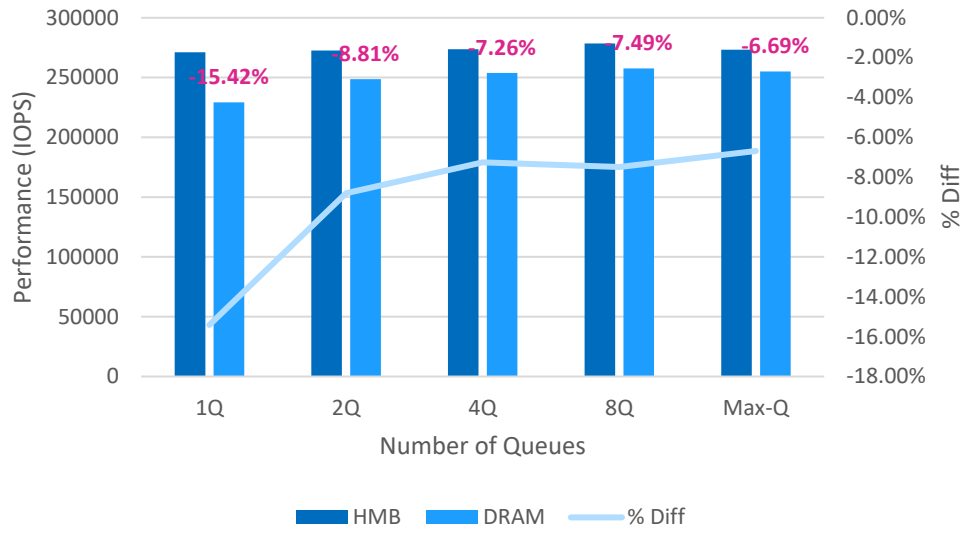

- Key Observations:
  - HMB device performed better than DRAM device in all Drive states & Queue Configurations
  - HMB is more Consistent at Aging state a well.
  - In Aging with increase in number of Q's. DRAM is able to perform better

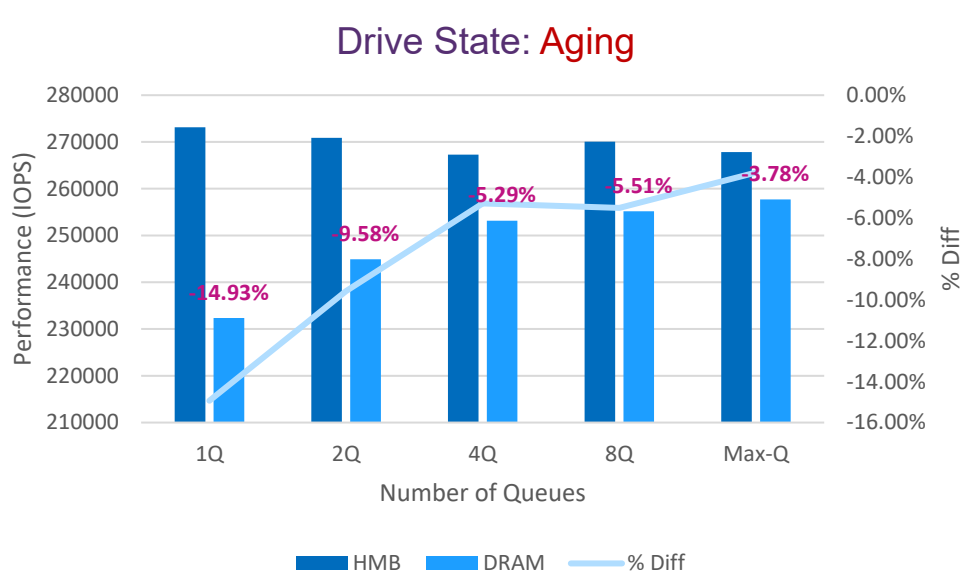

Disclaimer: Results are based on Custom test driver only

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


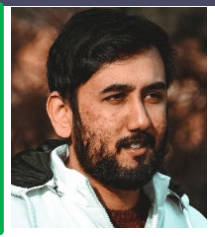

Pradeep SR Ranjith T

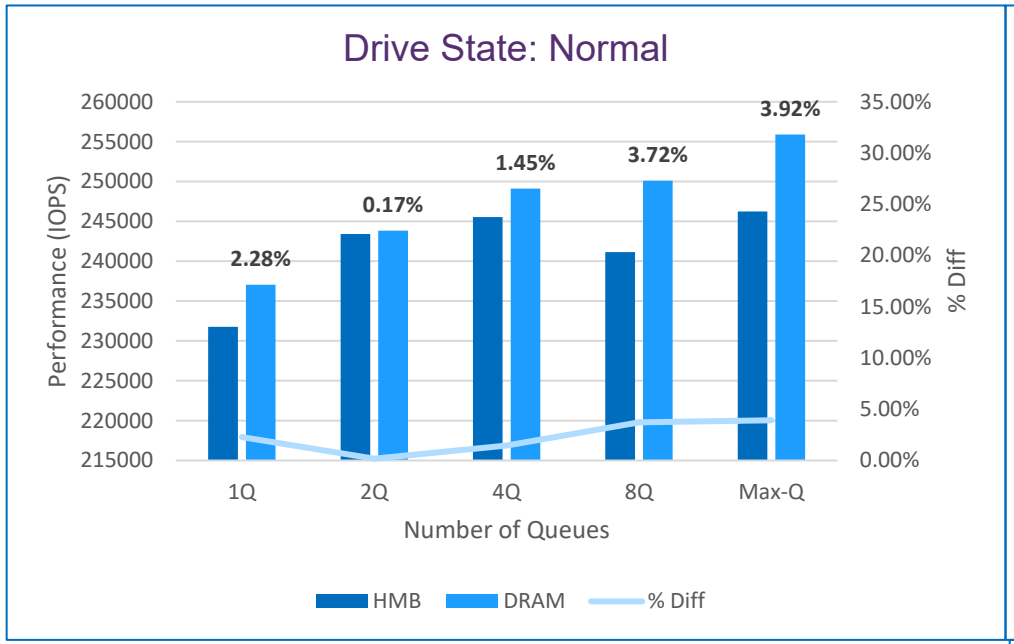

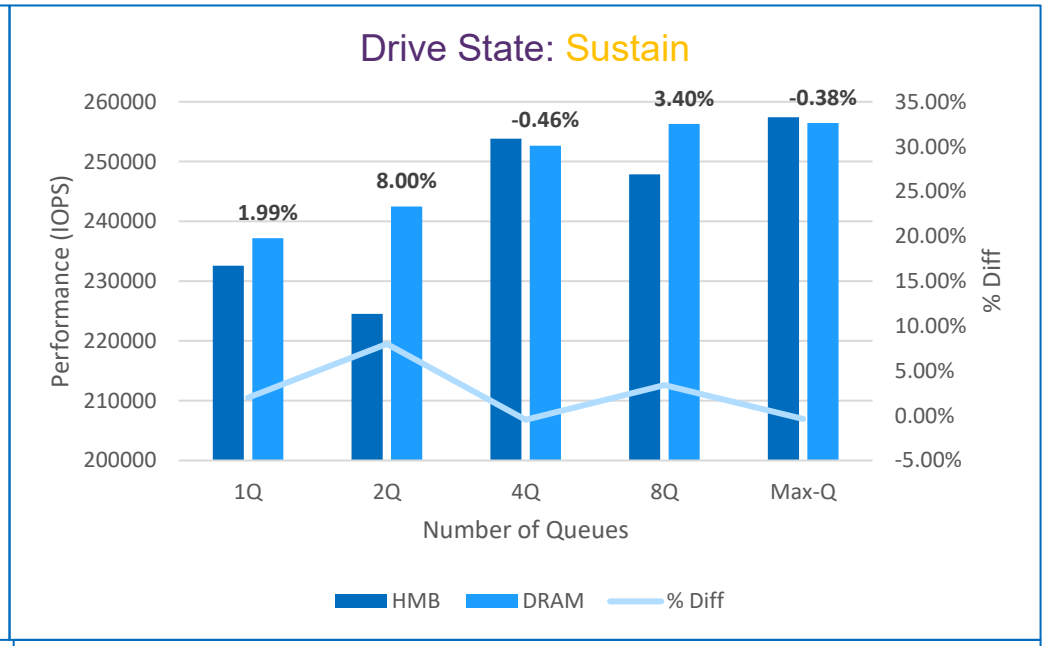

- Key Observations:
  - In Normal & sustain States both device perform on par across majority of Queue states
  - In Normal & Sustain state Performance increased as number of Q's increased
  - In Aging DRAM device leads by large margin

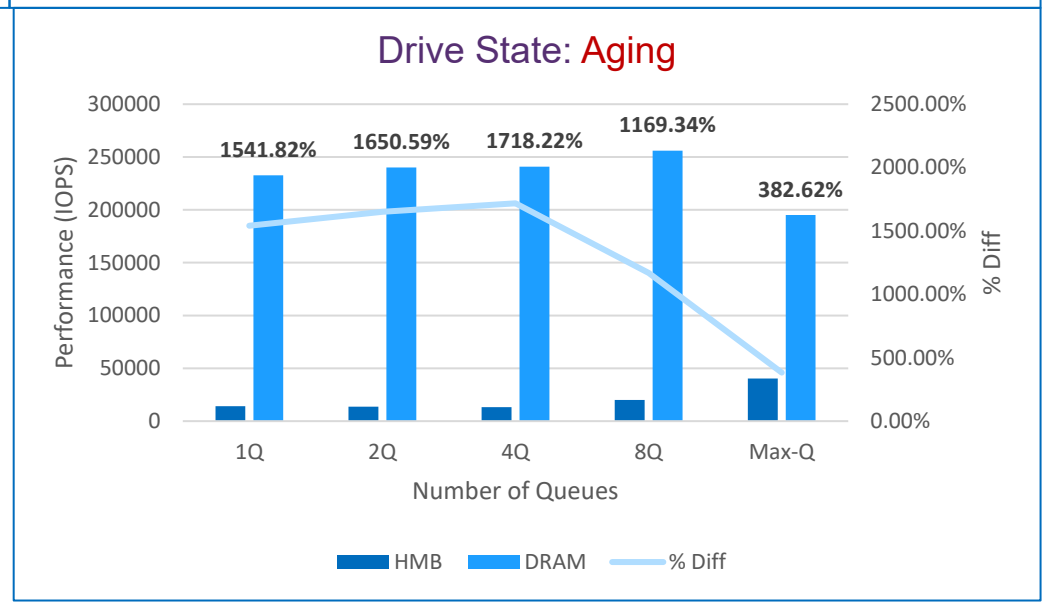

Disclaimer: Results are based on Custom test driver only

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


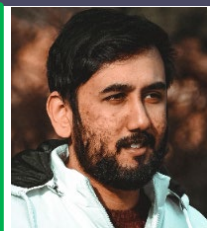

Pradeep SR Ranjith T

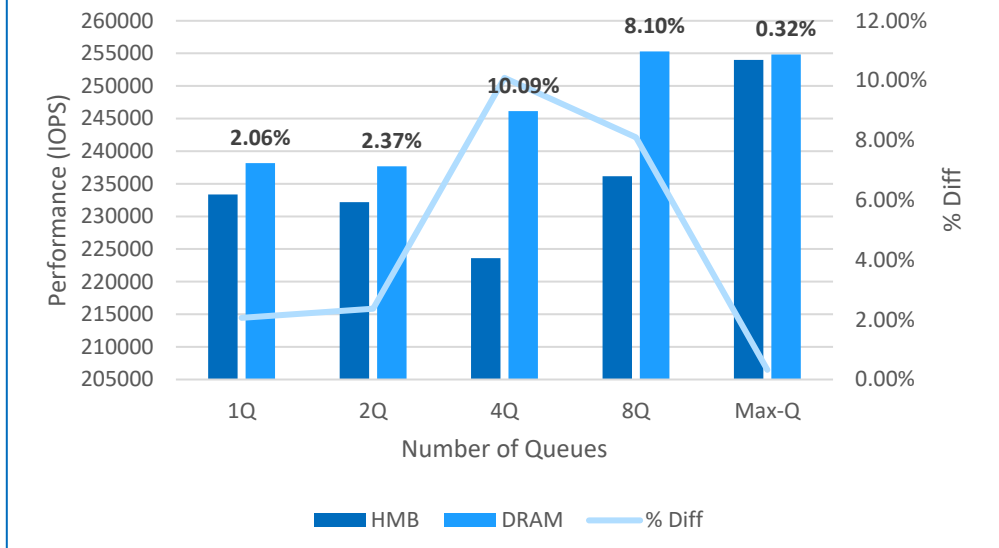

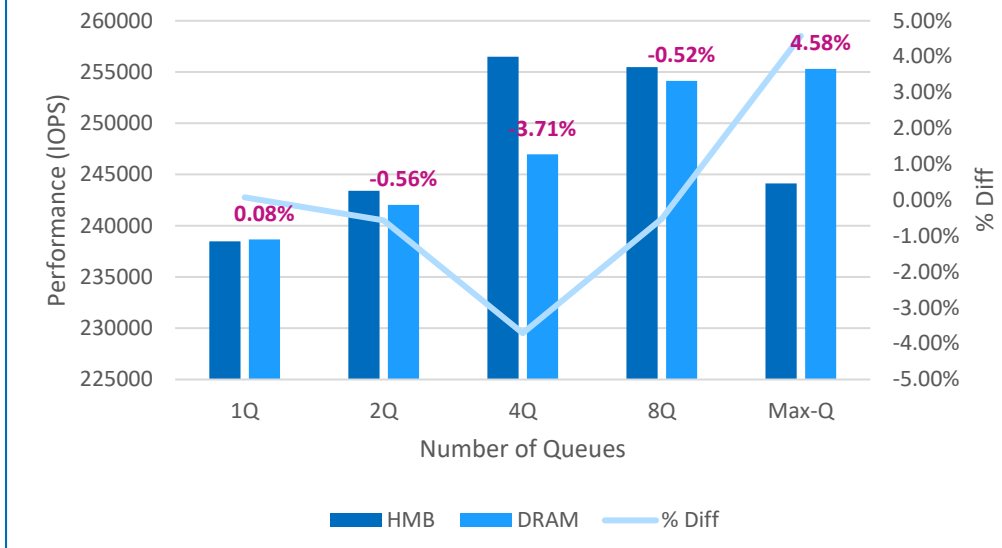

- Key Observations:
  - In Normal & sustain States both device perform on par across majority of Queue states
  - In Normal & Sustain state Performance increased as number of Q's increased
  - In Aging DRAM device leads by large margin

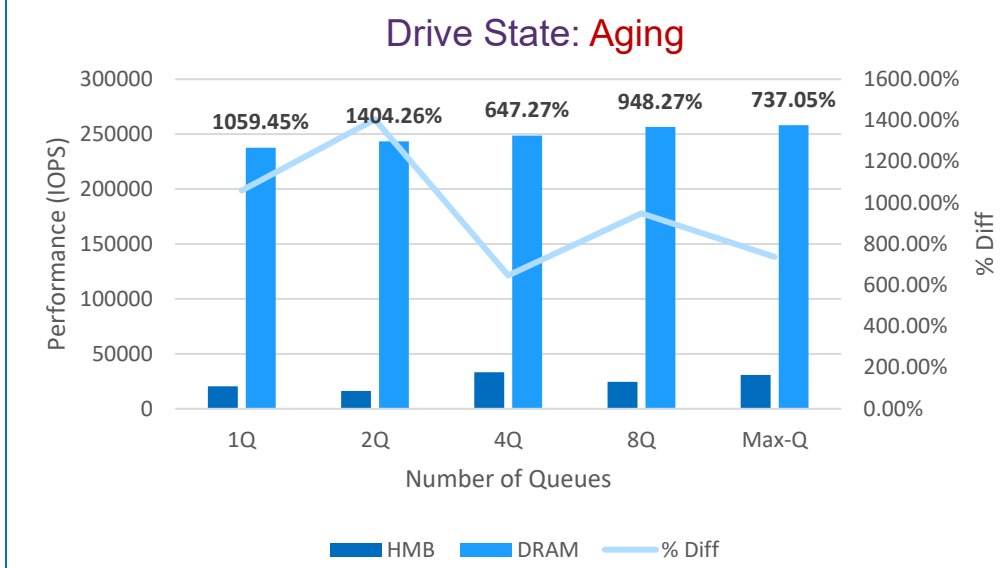

Disclaimer: Results are based on Custom test driver only

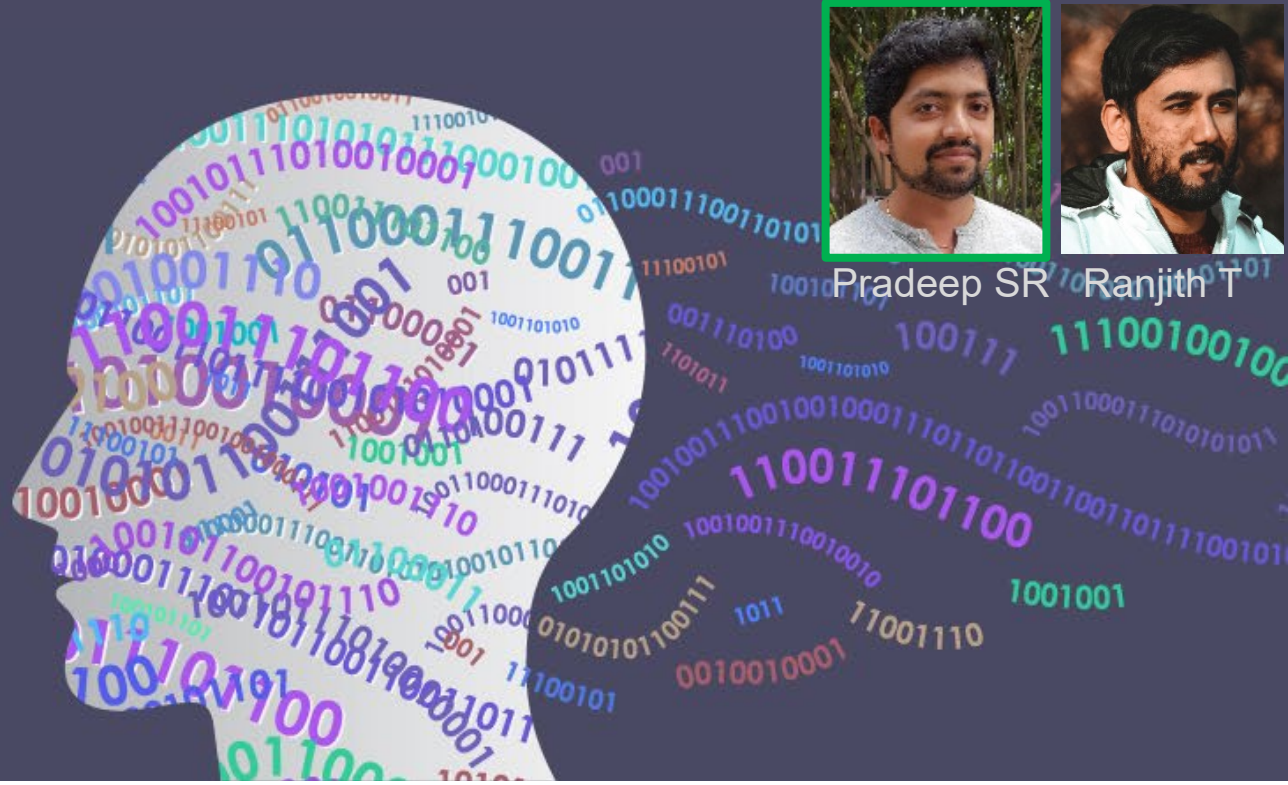

Pradeep SR Ranjith T

HMB Tuning Comparison

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


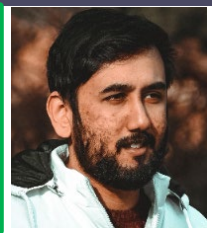

Pradeep SR Ranjith T

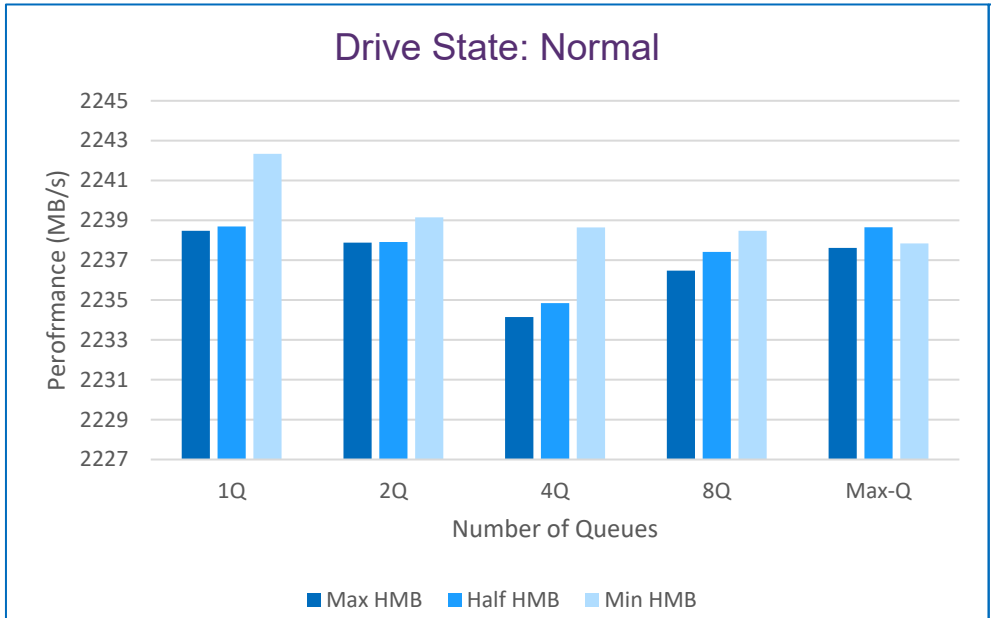

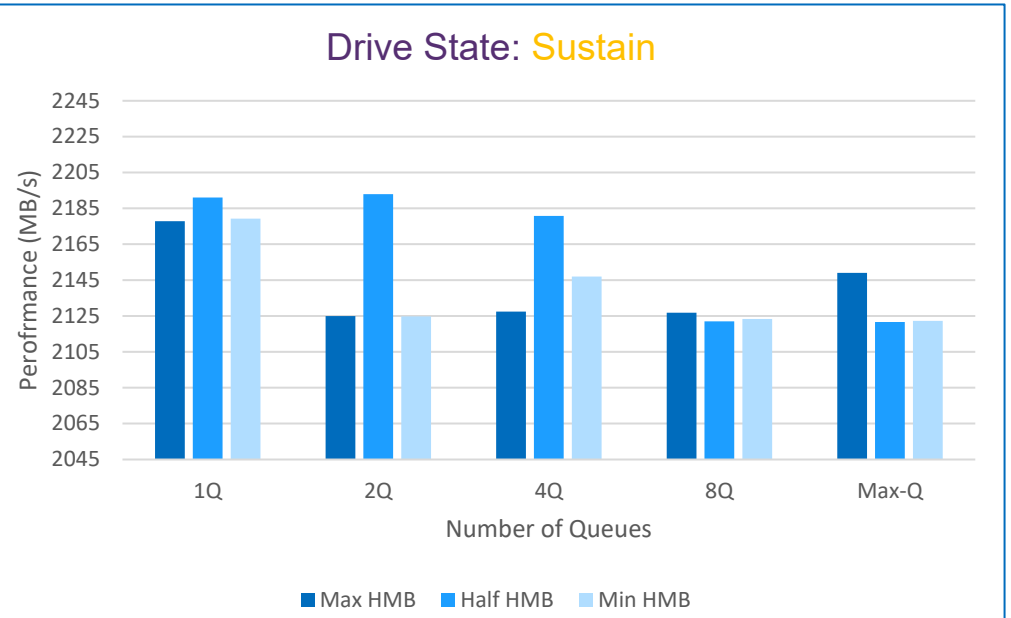

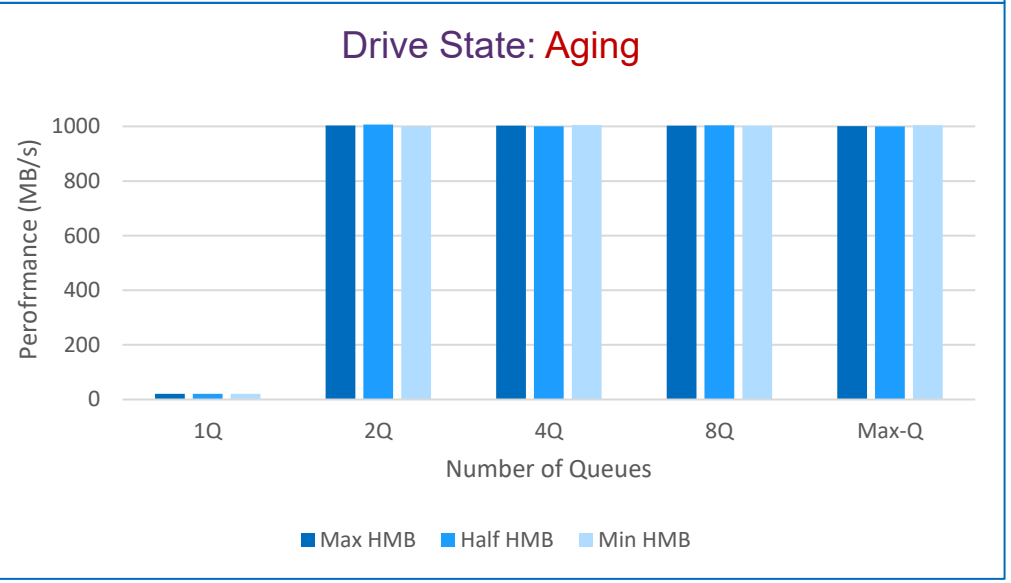

- Key Observations:
  - HMB Size variation didn't impact performance across all states & Q-Configuration.
  - Aging state performance drastically reduced across all Q-Configuration & badly hit in 1Q-Configuration.

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)

Disclaimer: Results are based on Custom test driver only

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


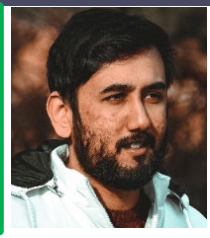

Pradeep SR Ranjith T

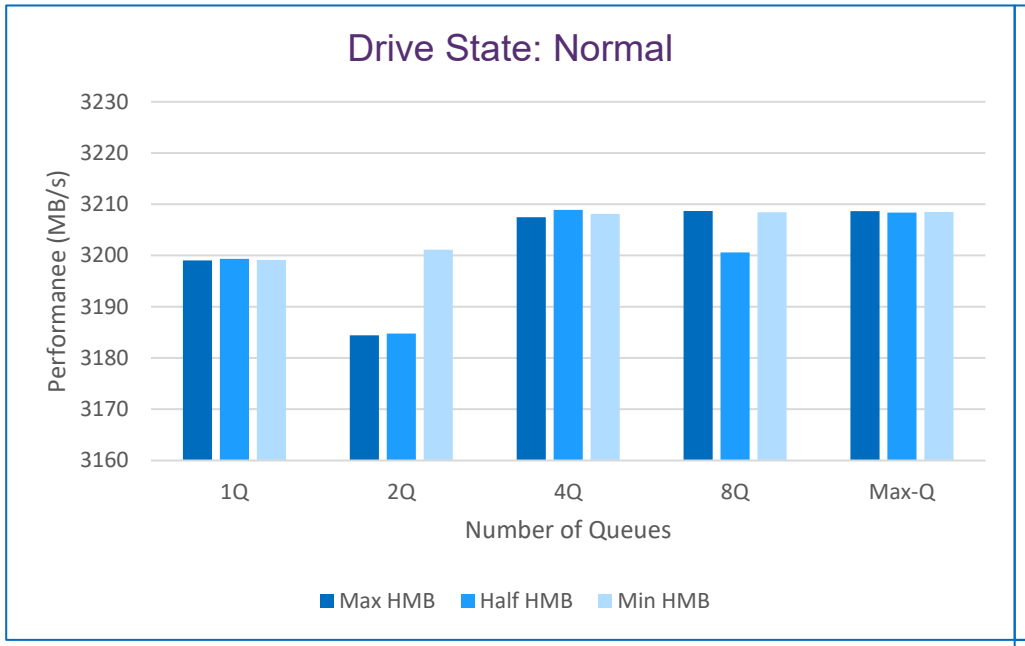

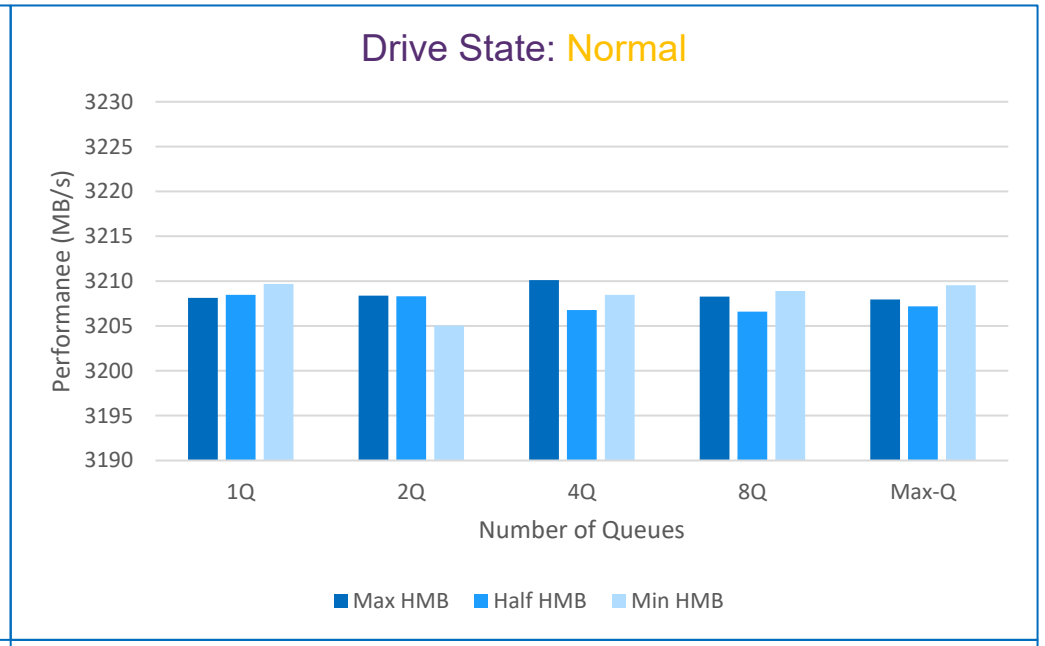

- Key Observations:
  - HMB allocation changes didn't impact in Normal & Sustain states.
  - In Aging 8Q configuration shows decrease in performance is proportional to decrease in HMB Allocation

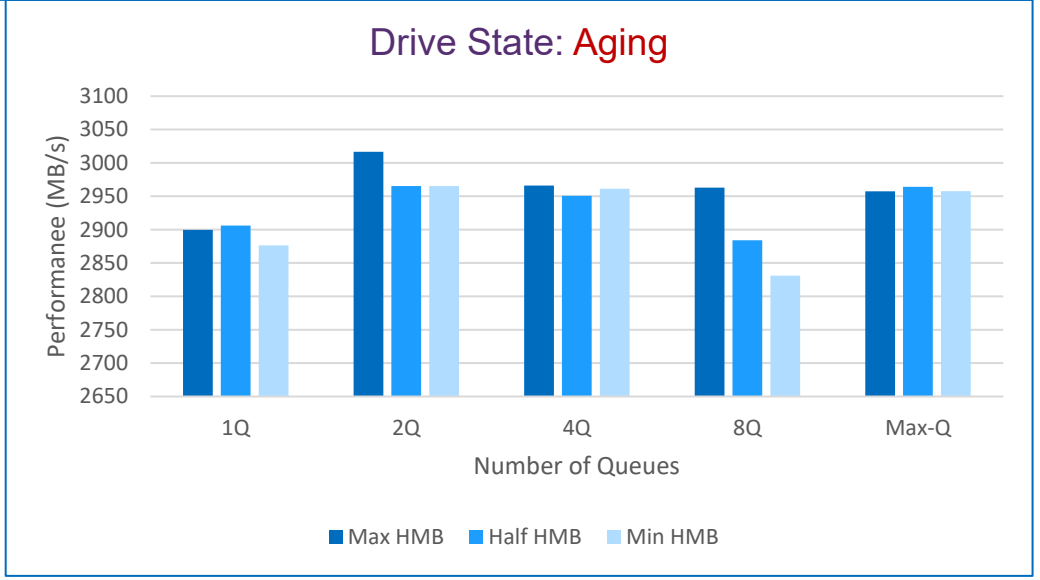

Disclaimer: Results are based on Custom test driver only

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


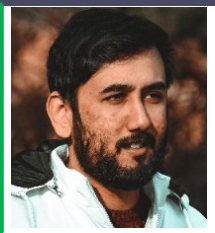

Pradeep SR Ranjith T

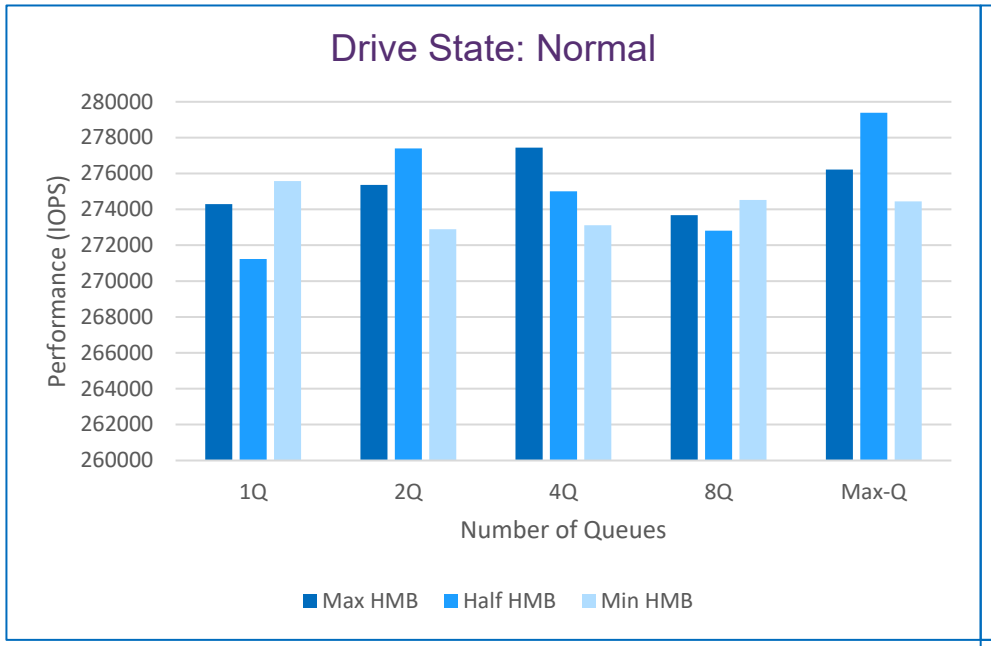

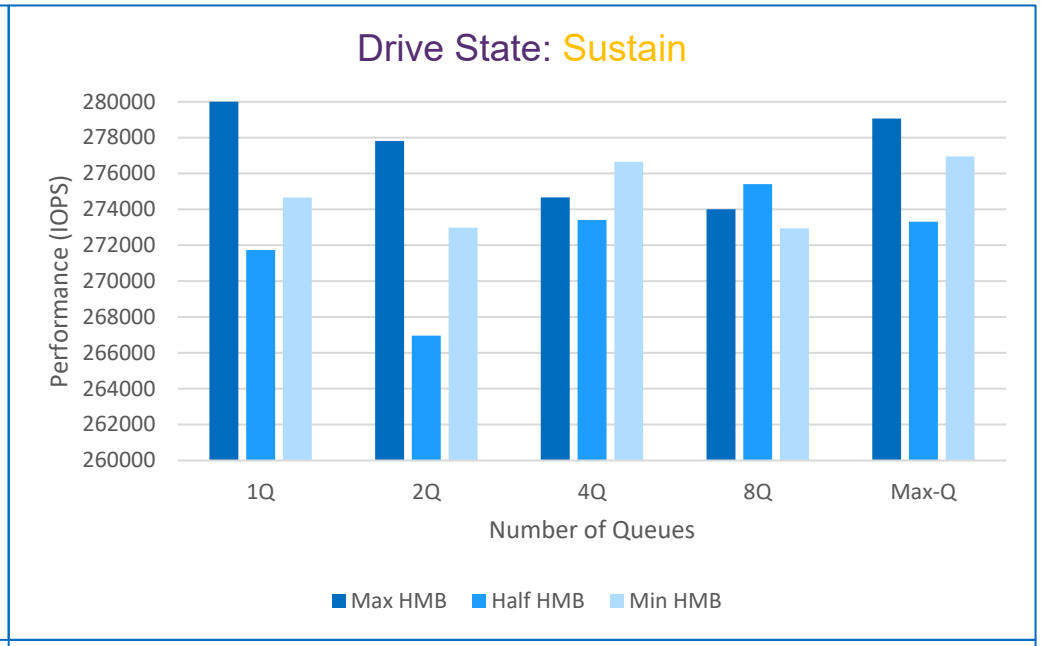

- Key Observations:
  - Performance better for Max HMB allocation in majority cases.
  - In Aging Half HMB size allocation is more consistent across all Queue configurations

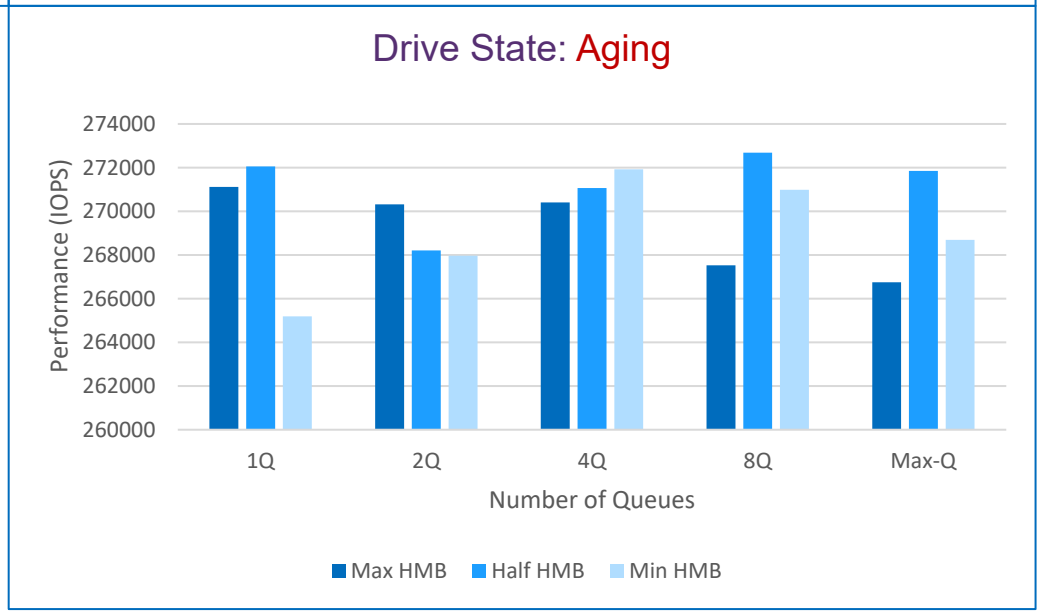

Disclaimer: Results are based on Custom test driver only

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


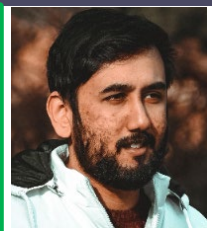

Pradeep SR Ranjith T

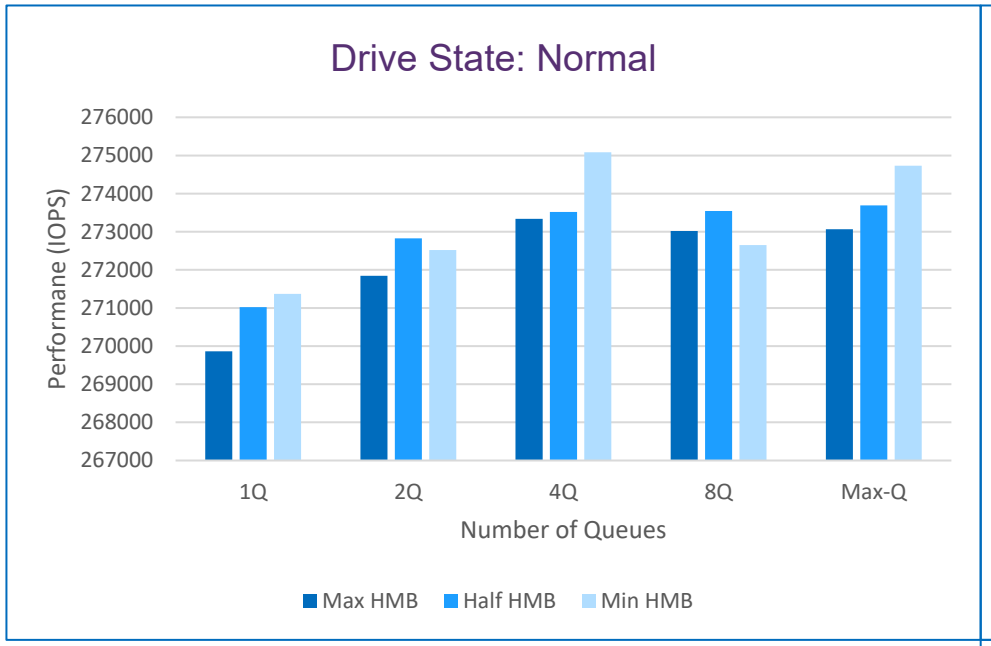

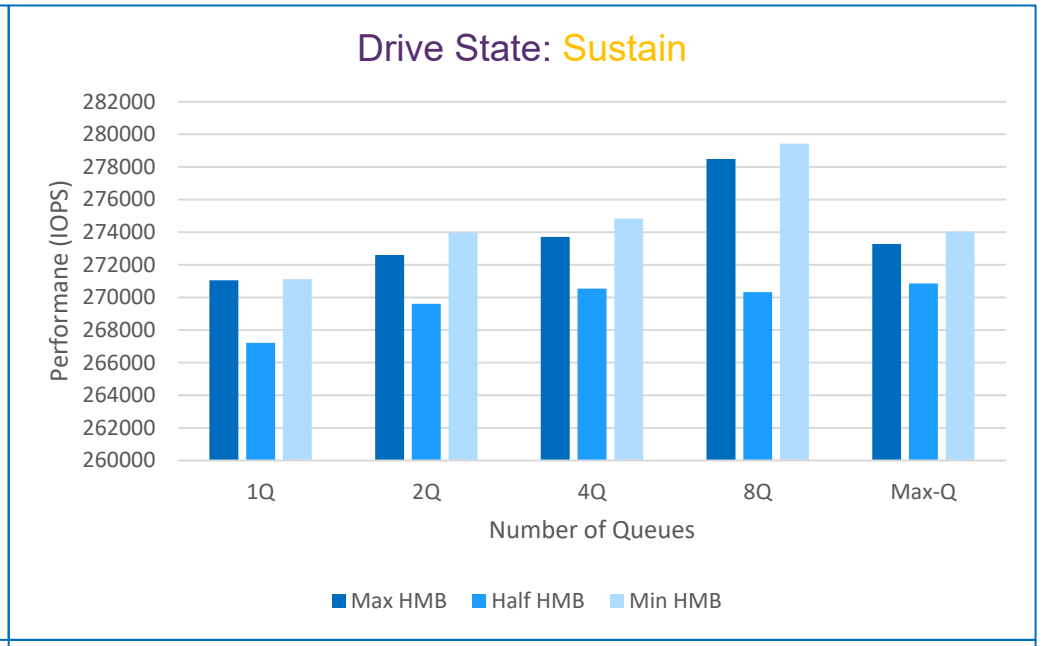

- Key Observations:
  - HMB allocation Half the preferred size performed better in Normal & Aging states.
  - Half HMB Size allocation worked relatively poor in sustain state across all Queue configurations
  - Min HMB allocation shows Aging performance decreased with increase in Queues configured.

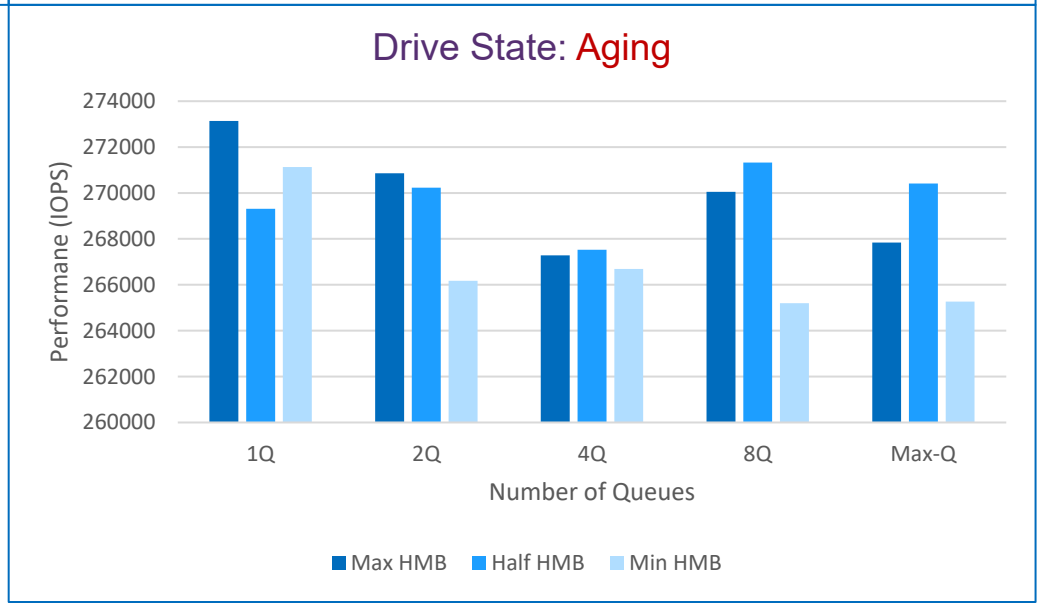

Disclaimer: Results are based on Custom test driver only

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


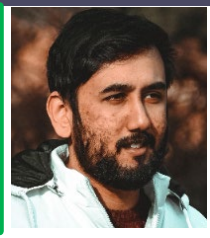

Pradeep SR Ranjith T

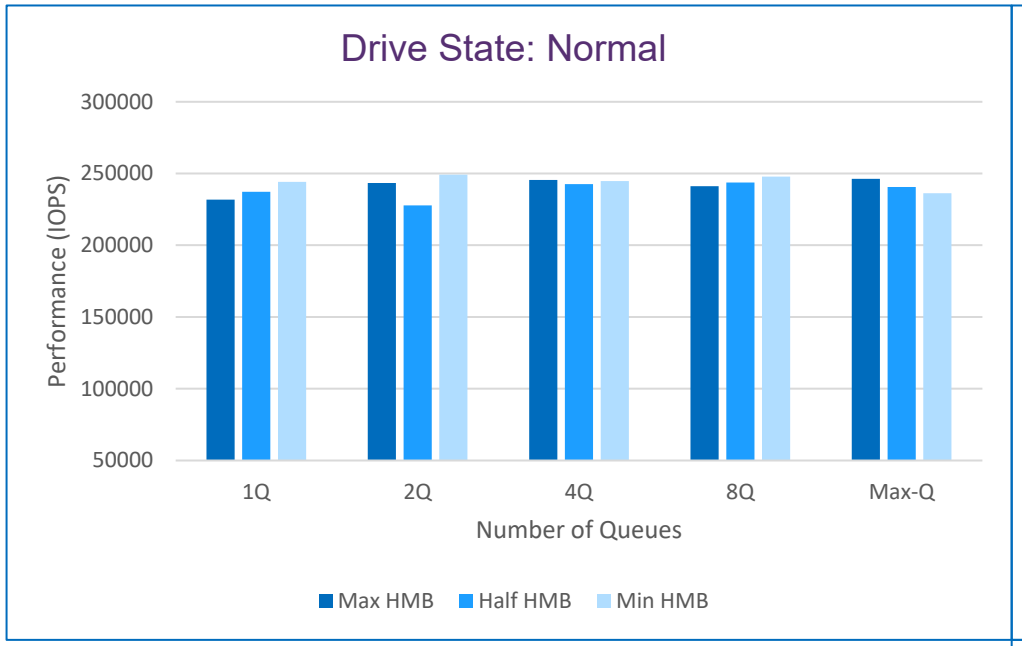

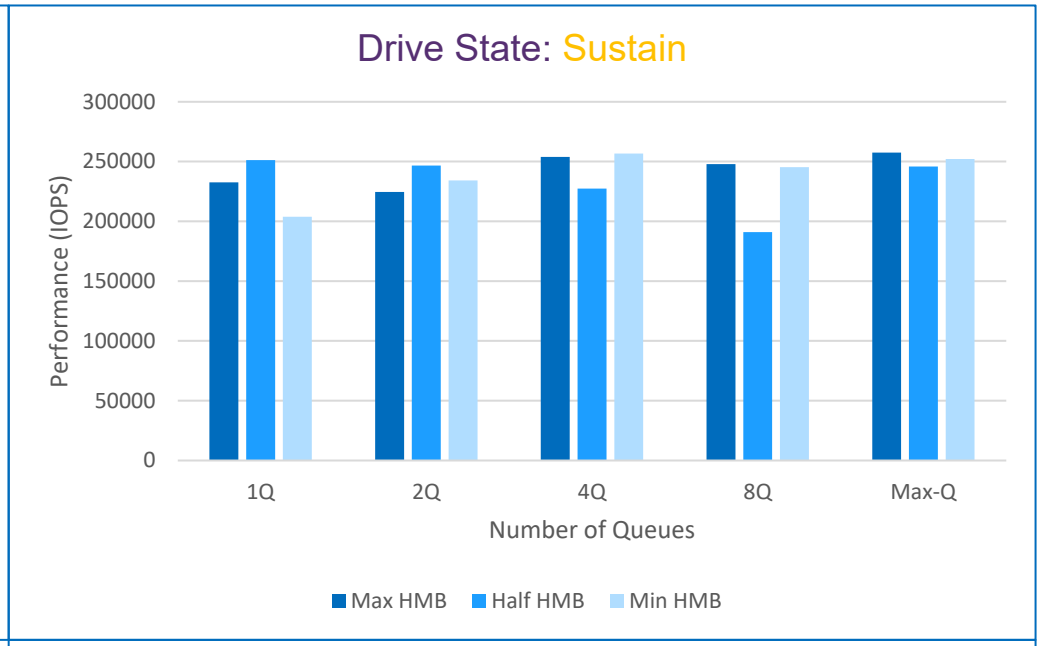

- Key Observations:
  - Normal & Sustain state performance has no impact due to HMB Size variation
  - Min HMB allocation shows Aging performance decreased with increase in Queues configured.
  - Max HMB allocation shows Aging performance increased with increase in Queues configured.

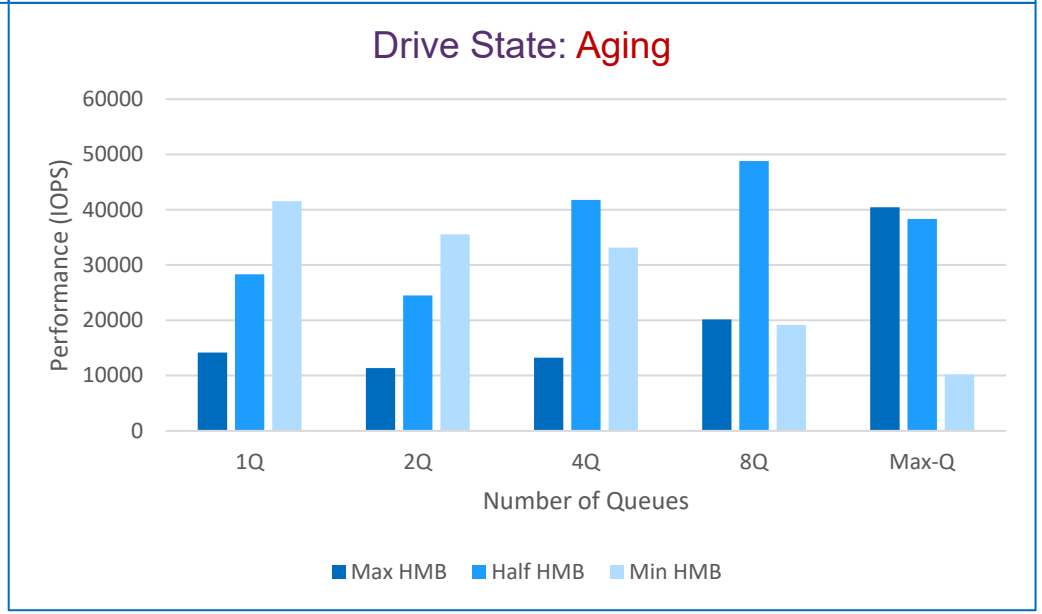

Disclaimer: Results are based on Custom test driver only

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


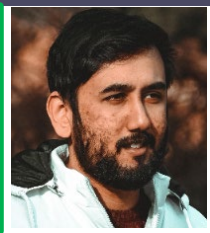

Pradeep SR Ranjith T

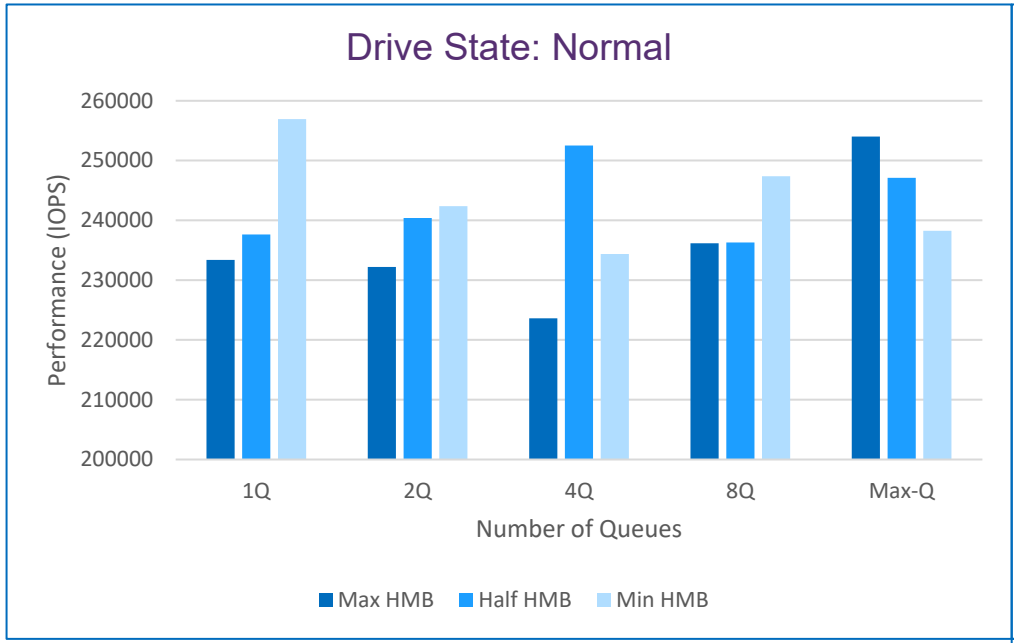

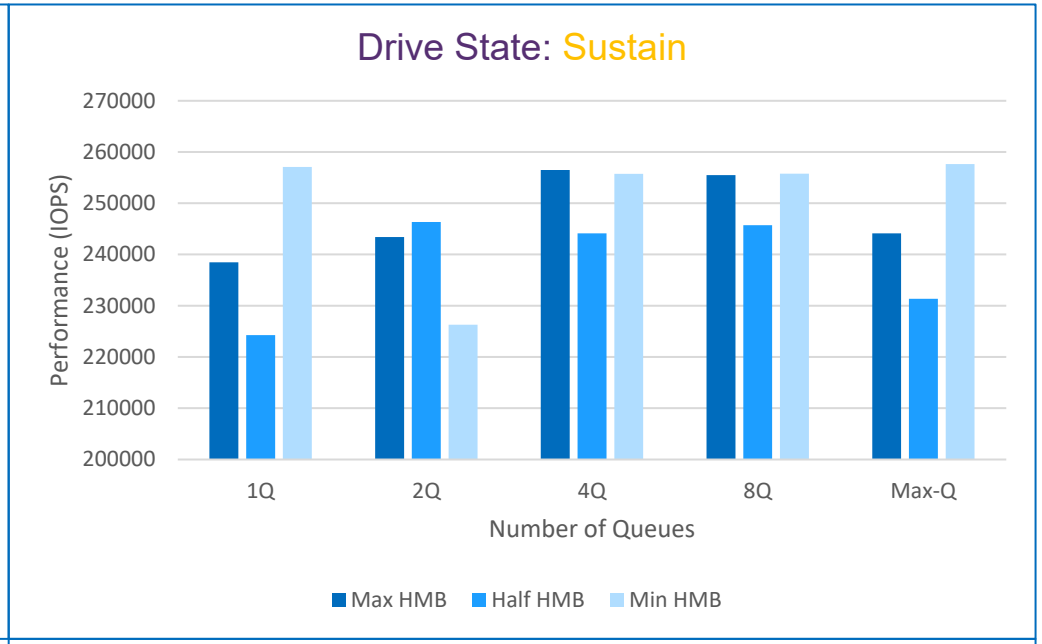

- Key Observations:
  - Min HMB size allocation shows better results in majority of configurations & drive states
  - Half HMB size allocation shows low performance in majority of cases in sustain & aging states

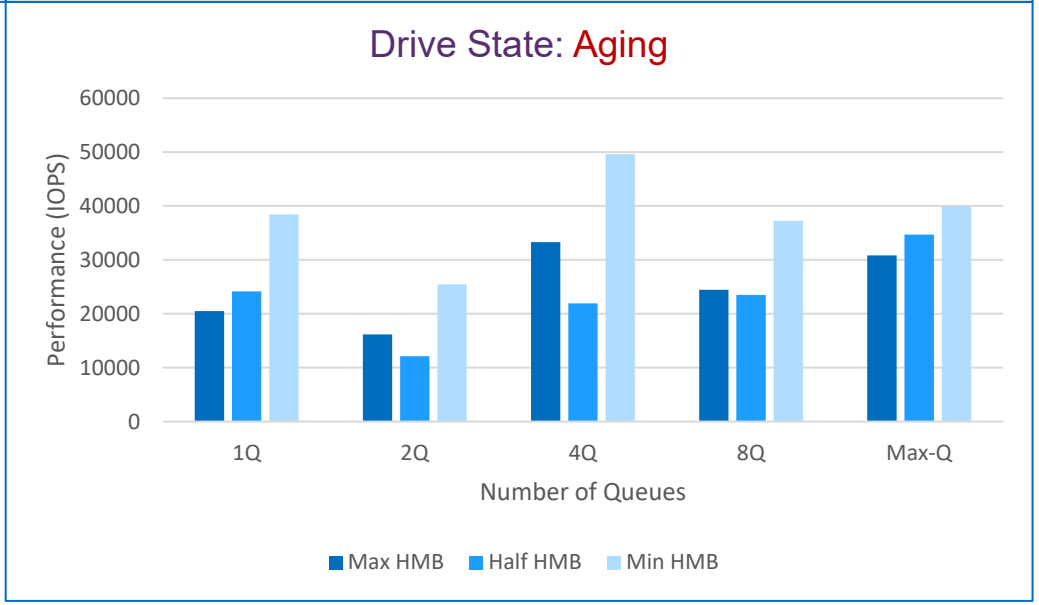

Disclaimer: Results are based on Custom test driver only

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

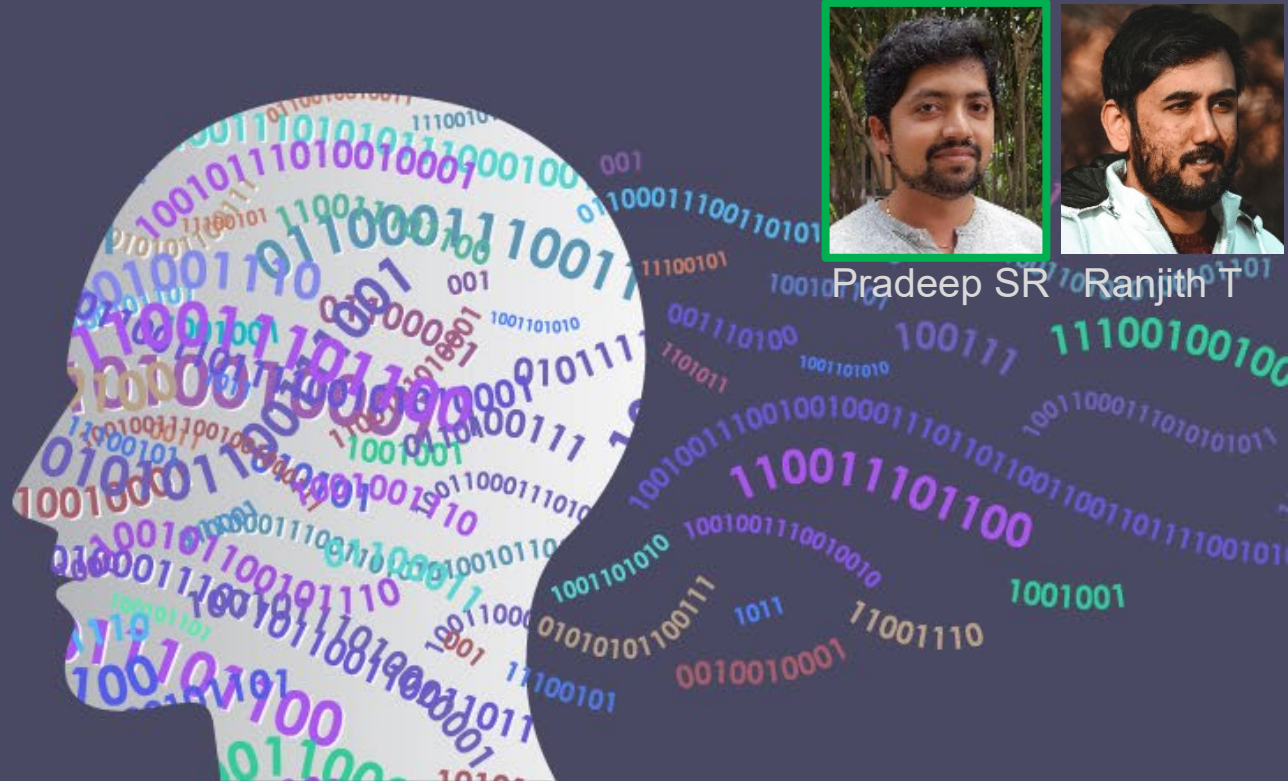

Pradeep SR Ranjith T

Conclusion

Conclusion


Pradeep SR Ranjith T

- On Sequential Workloads DRAM leads much ahead of HMB
- On Random Read workloads HMB worked better to DRAM
- On Random Writes HMB is competitive enough, but lose sheen as drive ages
- HMB Worked best with Max-Q configuration bridging gap with DRAM
- HMB Size Variation didn’t impact Sequential workloads much, Variations seen on Random workloads & Drive states
- Across all HMB Sizes allocated, 4Q configurations has been optimal.

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

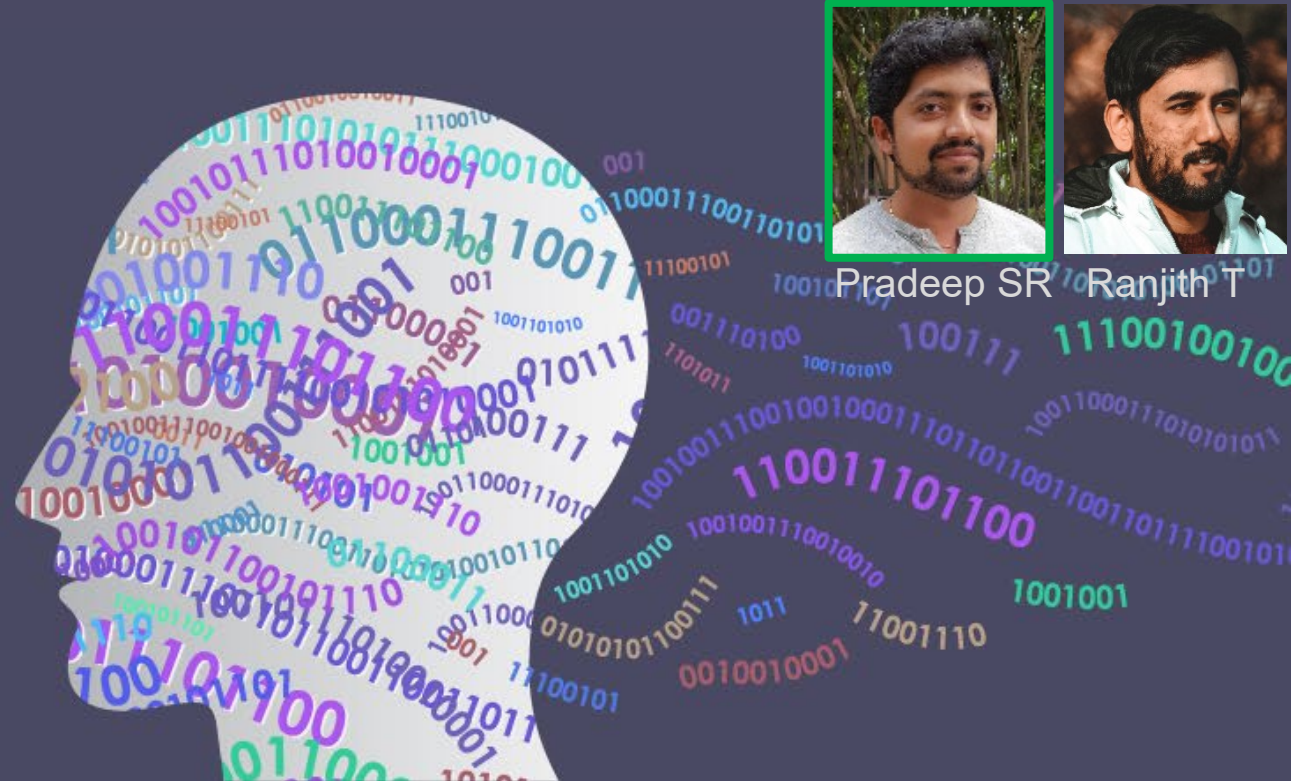

Pradeep SR Ranjith T

Please take a moment to rate this session.

Your feedback is important to us.

