

Shaping NVMe SSD IO Performance in Multi- Virtual Environments

SSDT-203-1: Performance Optimization & Modeling Techniques for SSDs

Gary Adams

AVP of Enterprise Marketing

Silicon Motion Technology Corp.

Legal Notice and Disclaimer

- The content of this document including, but not limited to, concepts, ideas, figures and architectures is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Silicon Motion Inc. and its affiliates. Silicon Motion Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in the informational content contained in this document.
- Nothing in these materials is an offer to sell any of the components or devices referenced herein.
- Silicon Motion Inc. may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Silicon Motion, Inc., the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.
- © 2024 Silicon Motion Inc. or its affiliates. All Rights Reserved.
- Silicon Motion, the Silicon Motion logo, MonTitan™, the MonTitan™ logo are trademarks or registered trademarks of Silicon Motion Inc.

8TB SRIOV SSD with Performashape™ Experiment

- Test configuration:
 - SSD hardware : Silicon Motion 8TB PCIe Gen5 Enterprise SSD with Micron B58R QLC NAND in U.2
 - SSD firmware : FIO SR-IOV FW supporting 4x VFs
 - Test platform : ASUS PRIME Z690-P (I5-12500 16GB)
 - Test program : FIO
 - Test workload : Supporting 4x VFs; using PerformaShape technologies for QOS improvement for multiple




* Picture from ASUS website

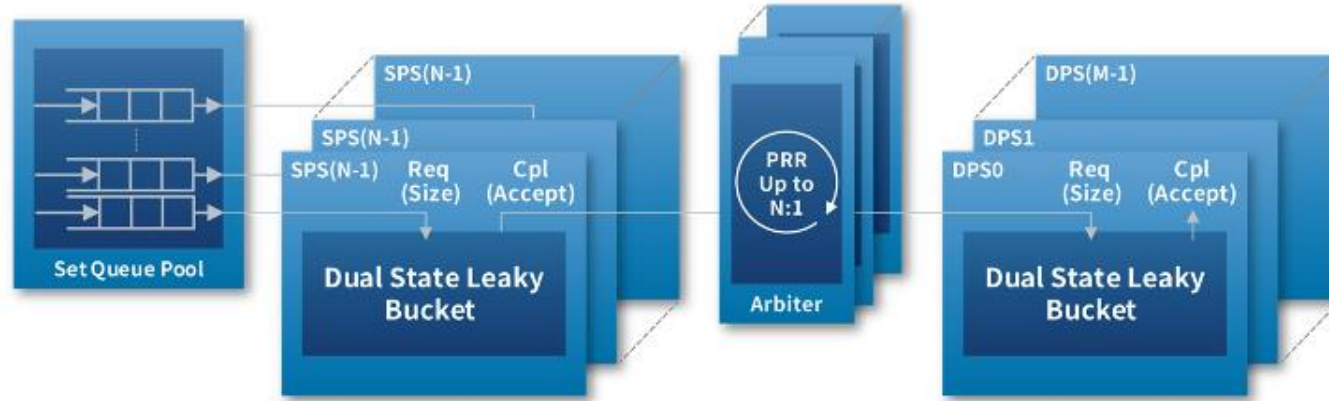
Over Subscribed SRIOV

VF/NS	Measurement	Performance Shapping Engine		Host Setting
		SPS Setting	DPS - ID4	
VF1/NS1	5.2 GB/s			6GB/S (5723MiB)
VF2/NS2	5.2 GB/s			6GB/S (5723MiB) - noisy
VF3/NS3	1.95 GB/s			2GB/S (1908MiB)
VF4/NS4	1.95 GB/s			2GB/S (1908MiB) - noisy

- ✓ 16GB/S Read Requests from Host in 13GB/S system
- ✓ Noisy and Limited performance on VF1 and VF2



Multi-Tenant Quality-of-Service (QoS) Improvement Technology



- To further enhance SSD QoS in multi-tenant environment, QoS management techniques can be deployed together with FDP
- PerformaShape™, as an example for QoS management, is developed based on “Dual State Leaky Bucket” algorithm where each QoS set is assigned with two token buckets and IO traffic flows when token is available
- PerformaShape™ benefits
 - Smooth out fluctuations
 - Isolate noisy neighbors
 - Fully utilize the SSD bandwidth

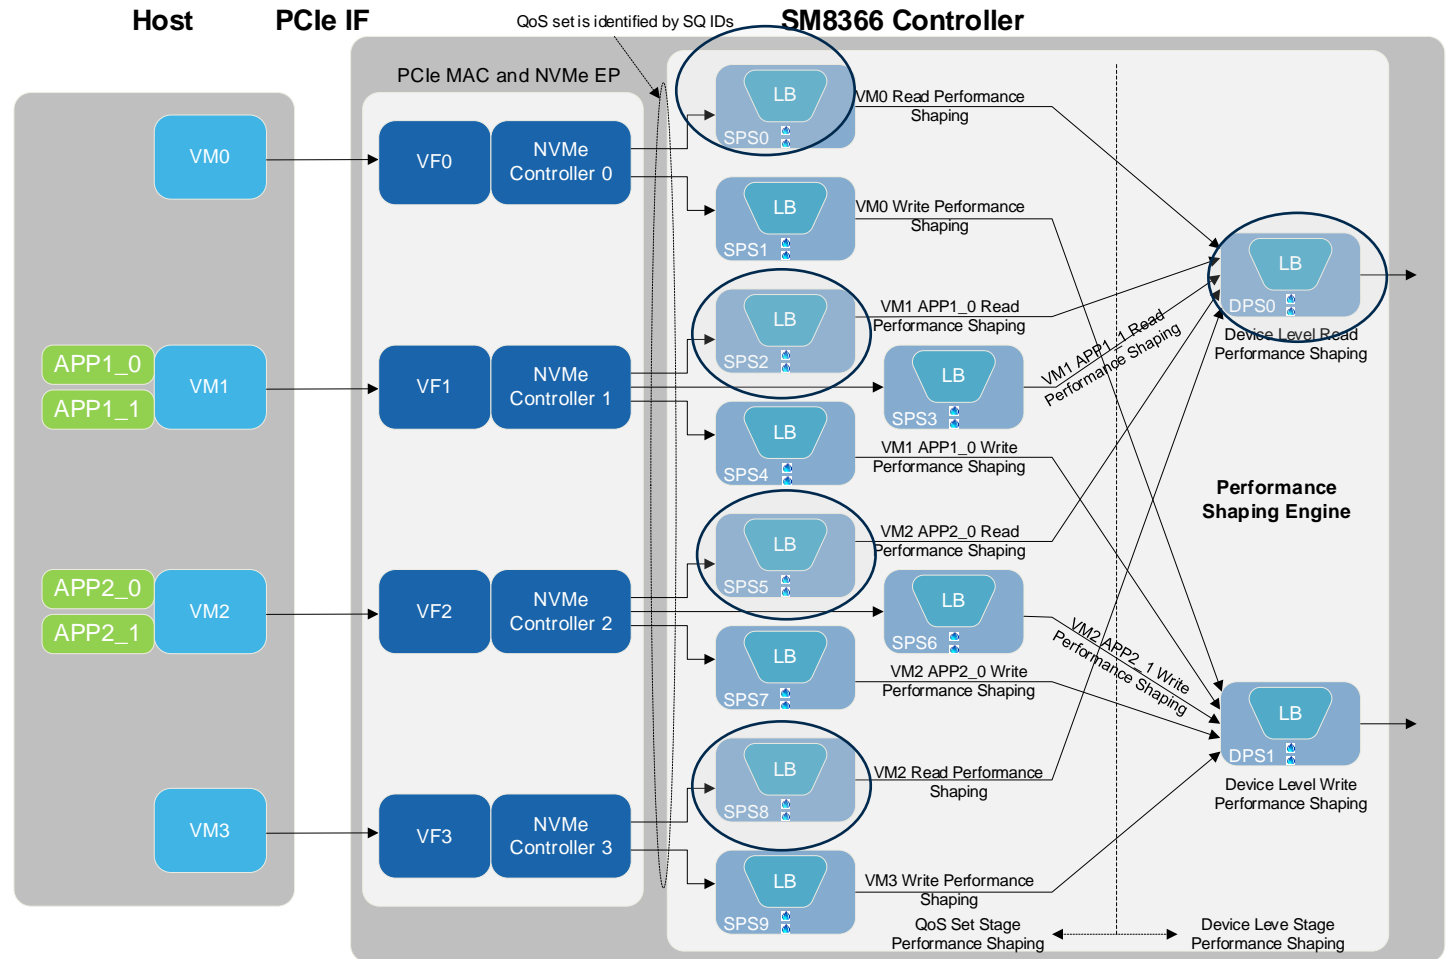
PerfomaShape™ on Virtual Functions

Set Selection

- Per NVMe Controller (per Virtual Machine)
- Per NVMe Namespace (per Container)
- Per Application
 - Per active zone with ZNS
 - Per Reclaim Unit with FDP
 - Per submission queue

Set Configuration

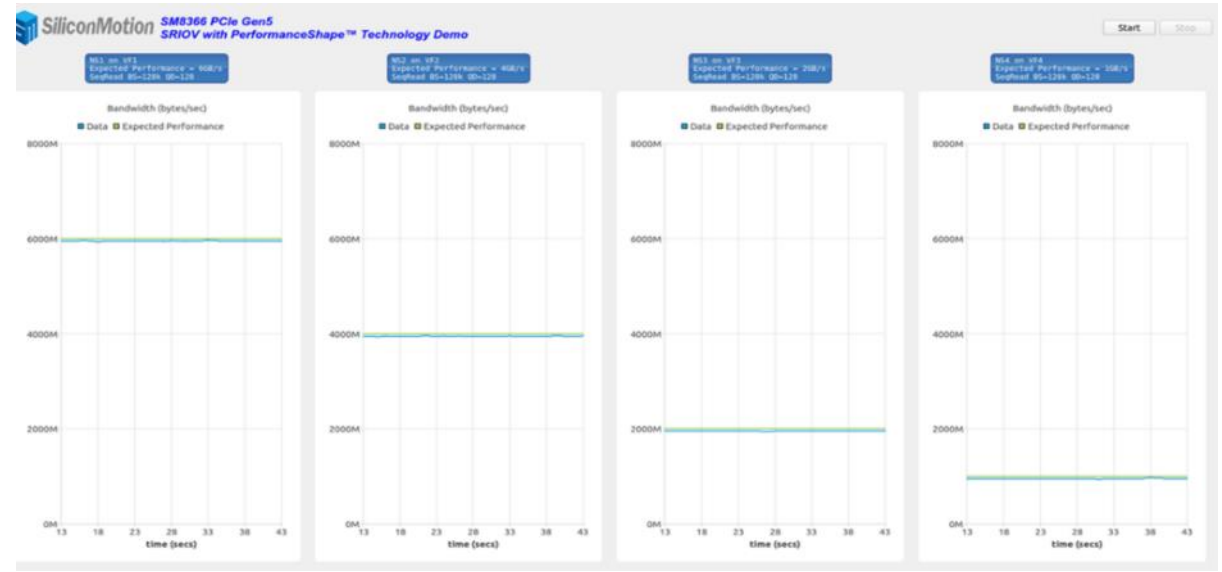
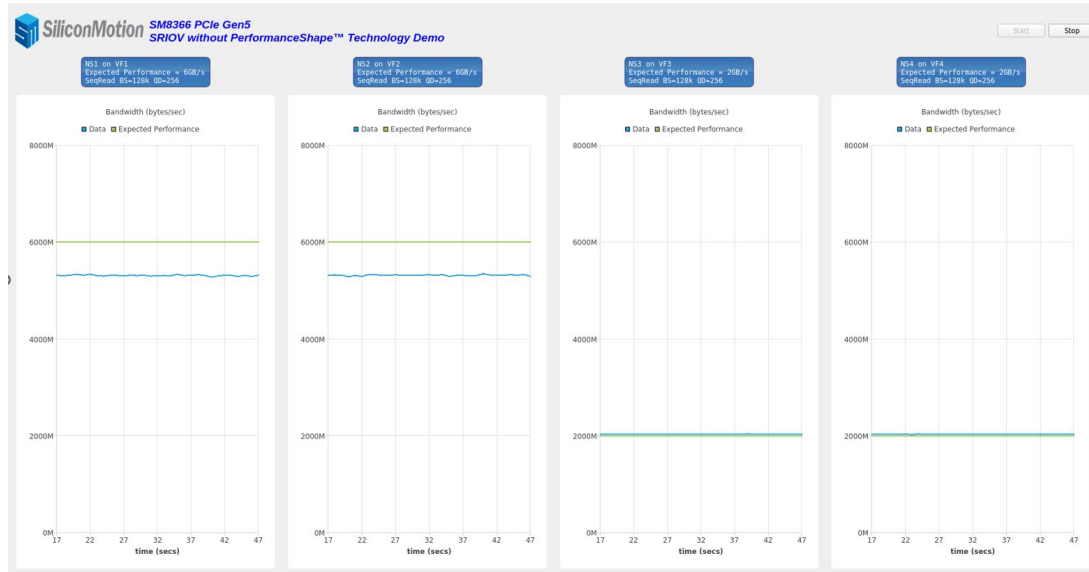
- Access Pattern: Random (IOPS); Sequential (KB/S)
- Access Type: Read



Managing SRIOV with PerformaShaping™

VF/NS	Measurement	Performance Shapping Engine		Host Setting
		SPS Setting	DPS - ID4	
VF1/NS1	5.97GB/s	6GB/S (8083)	12.9 - 13GB/s	6GB/S (5723MiB)
VF2/NS2	3.98GB/s	4GB/S (12125)		6GB/S (5723MiB) - noisy
VF3/NS3	1.99GB/s	2GB/S (24250)		2GB/S (1908MiB)
VF4/NS4	0.96GB/s	1GB/S (48500)		2GB/S (1908MiB) - noisy

- ✓ Mapped out ~ 13G Read requests from the PerformaShape™
- ✓ Isolates and Guarantees Performance per Tenant
- ✓ Removes Noisy Neighbors



Summary

- PerformaShape™ technology further improve QoS in multi-tenant applications
- We have a good starting point on looking Read Bandwidth control and look at mixed work loads next



Meet us at booth #315

Driving AI Innovation in Flash Storage

Scan to learn more!

