



# **NVMe™ Performance Local vs. Remote**

Oscar Pinto, Ming Lin, Gunna Marripudi  
Samsung Semiconductor Inc.



## Legal Disclaimer

This presentation is intended to provide information concerning SSD and memory industry. We do our best to make sure that information presented is accurate and fully up-to-date. However, the presentation may be subject to technical inaccuracies, information that is not up-to-date or typographical errors. As a consequence, Samsung does not in any way guarantee the accuracy or completeness of information provided on this presentation.

The information in this presentation or accompanying oral statements may include forward-looking statements. These forward-looking statements include all matters that are not historical facts, statements regarding the Samsung Electronics' intentions, beliefs or current expectations concerning, among other things, market prospects, growth, strategies, and the industry in which Samsung operates. 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 contained in this presentation or in the accompanying oral statements. In addition, even if the information contained herein or the oral statements are shown to be accurate, those developments may not be indicative developments in future periods.

The logo for Flash Memory Summit 2016, featuring a yellow sunburst icon above the text "Flash Memory" in red and black, with "SUMMIT" in white on a blue rectangular background below it.

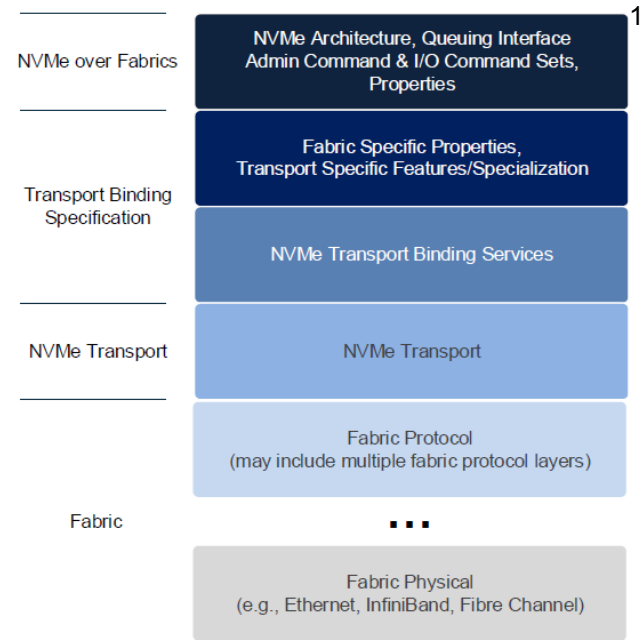
# Flash Memory Summit **Agenda**

- NVM Express<sup>®</sup> over Fabrics Overview
- Test Configuration
- Performance Comparison



# NVM Express<sup>®</sup> over Fabrics Overview

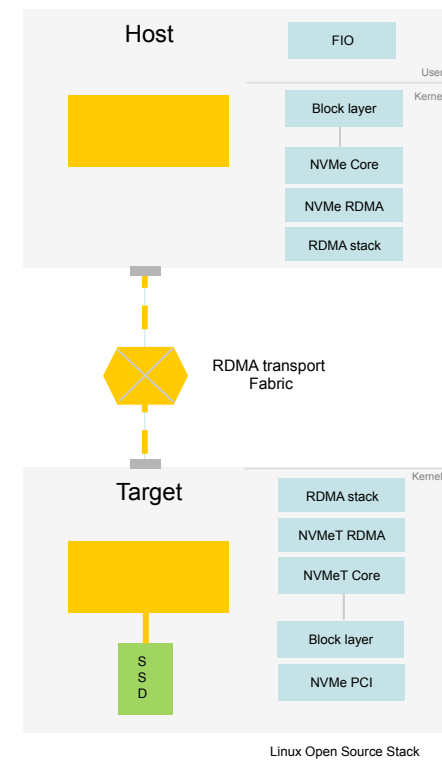
- Maintains consistency with base NVMe<sup>™</sup> definition but for fabrics support\*
- Support for multiple transport types
- Exposes NVMe<sup>™</sup> parallelism to host
- Performance closer to local NVMe drives





# NVM Express<sup>®</sup> over Fabrics Linux Stack

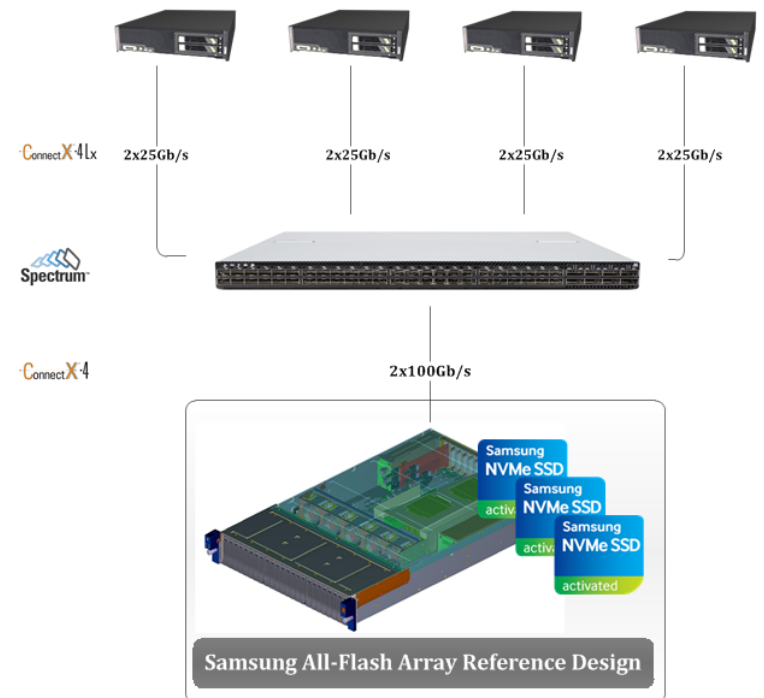
- Kernel stack
- Leverages OFED RDMA stack
- NVMe<sup>™</sup> command processing in
  - Initiator: *nvme\_core*
  - Target: *nvmet\_core*
- NVMe<sup>™</sup> transport binding in
  - Initiator: *nvme\_rdma*
  - Target: *nvmet\_rdma*






# Test Configuration

- Configuration
  - 1x NVMe-oF™ target
    - 24x Samsung NVMe™ Datacenter 2.5” SSDs
    - 2x 100Gb/s ConnectX®-4 EN
  - 4x initiator hosts
    - 2x25Gb/s each
  - Ubuntu 14.04.4 LTS Linux 4.7.0-rc2 kernel
  - Open Source NVMe-oF™ kernel drivers





# Latency Comparison

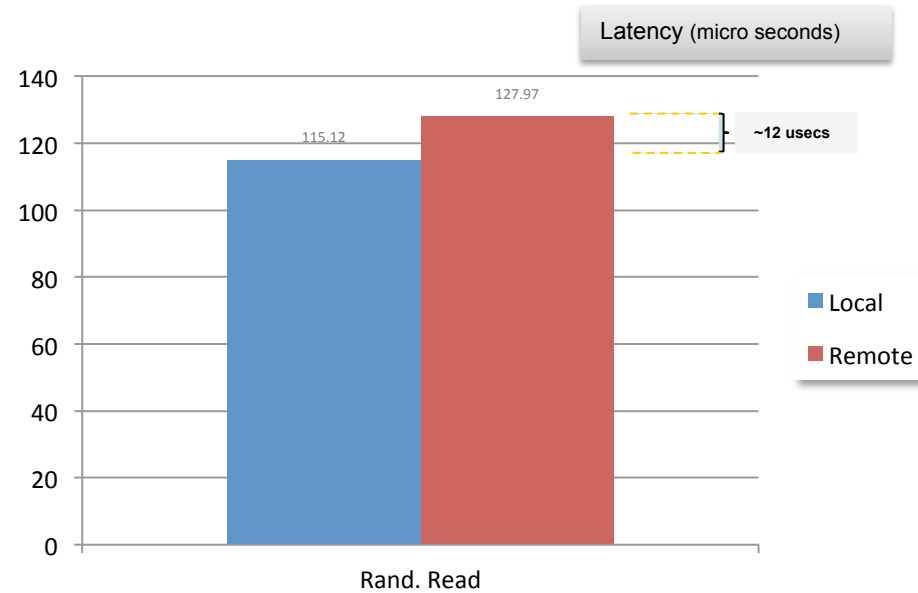


## Random IO at QD1, 1 job

- Round-trip delta: Reads ~17usecs; Writes ~9usecs



## Latency Comparison - Loaded




### Random IO at QD16, 2 jobs

- Performance delta: ~12 usecs





# Latency Comparison - Scale




## Random IO QD16

- NVMe-oF™ latency scales linearly from 1 to 24 drives




# Performance (24 SSDs)



- High aggregate NVMe-oF™ performance: 4.3M IOPS & 21.5GB/s throughput
- Further optimizations needed for performance to scale



# IOPS Scaling




## Random IO at QD16, 2 jobs per drive

- NVMe-oF™ stack scales linearly for Random IO



# Throughput Scaling



## Sequential IO at QD32, 1 job per drive

- NVMe-oF™ stack scales linearly for Sequential IO



## Summary: NVMe™ Local vs. Remote

Performance Delta		1-drive	24-drive
Latency (QD1)	Read	~17usecs	-
	Write	~9usecs	-
IOPS	Read	10%	12%
	Write	On par	2%
Throughput	Read	On par	18%
	Write	On par	On par



## Call for Action

- Further performance analysis & tuning of Linux Open Source stack
- Transport binding optimizations
- Feature enhancements
  - Faster failure recover scenarios
  - Reservations support



Thanks



# Test Methodology

- Host – Target Setup
  - Ubuntu 14.04.4 LTS Linux 4.7.0-rc2 kernel
  - Each host mapped with 2 subsystems each with 3 SSDs
  - Each subsystem mapped to 1x 25Gb/s NIC
- Benchmark tool
  - fio 2.6-20-g2caf
  - ioengine=libaio
  - Random IOPS: 4k, iodepth=16, numjobs=2
  - Sequential IO: 128k, iodepth=32, numjobs=1
  - Latency: random IO 4k, iodepth=1, numjobs=1