Host Memory Buffer (HMB) based SSD System

Forum J-31: PCIe/NVMe Storage

Jeroen Dorgelo

Mike Chaowei Chen

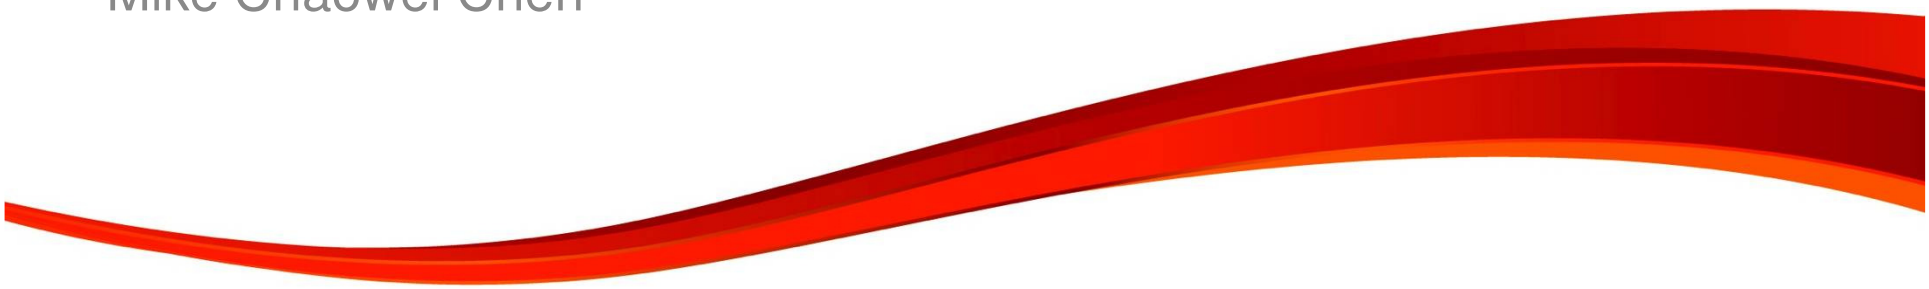

Agenda

- Market Trends in Client SSD
- Introducing DRAMless architecture for SATA & PCIe
- Host Memory Buffer(HMB) scales up DRAMless PCIe SSD
- HMB DRAMless vs. DRAM-based SSD
- Ecosystem development

Market Trends

- Main drivers for SSD adoption have been:
  - Performance
    - PC Mark8 score from HDD to SSD: 1,200 to 4,800 points (SATA)
    - Latest PCIe Gen3x4 NVMe drives score over 5,100 points
  - Form Factor:
    - Ultrabooks require space to fit in battery
    - SSD can shrink to module size (M.2)
- BUT all came with a Premium to pay

Market Trends

- Now: cost parity of 128GB SSD and 500GB Mobile HDD

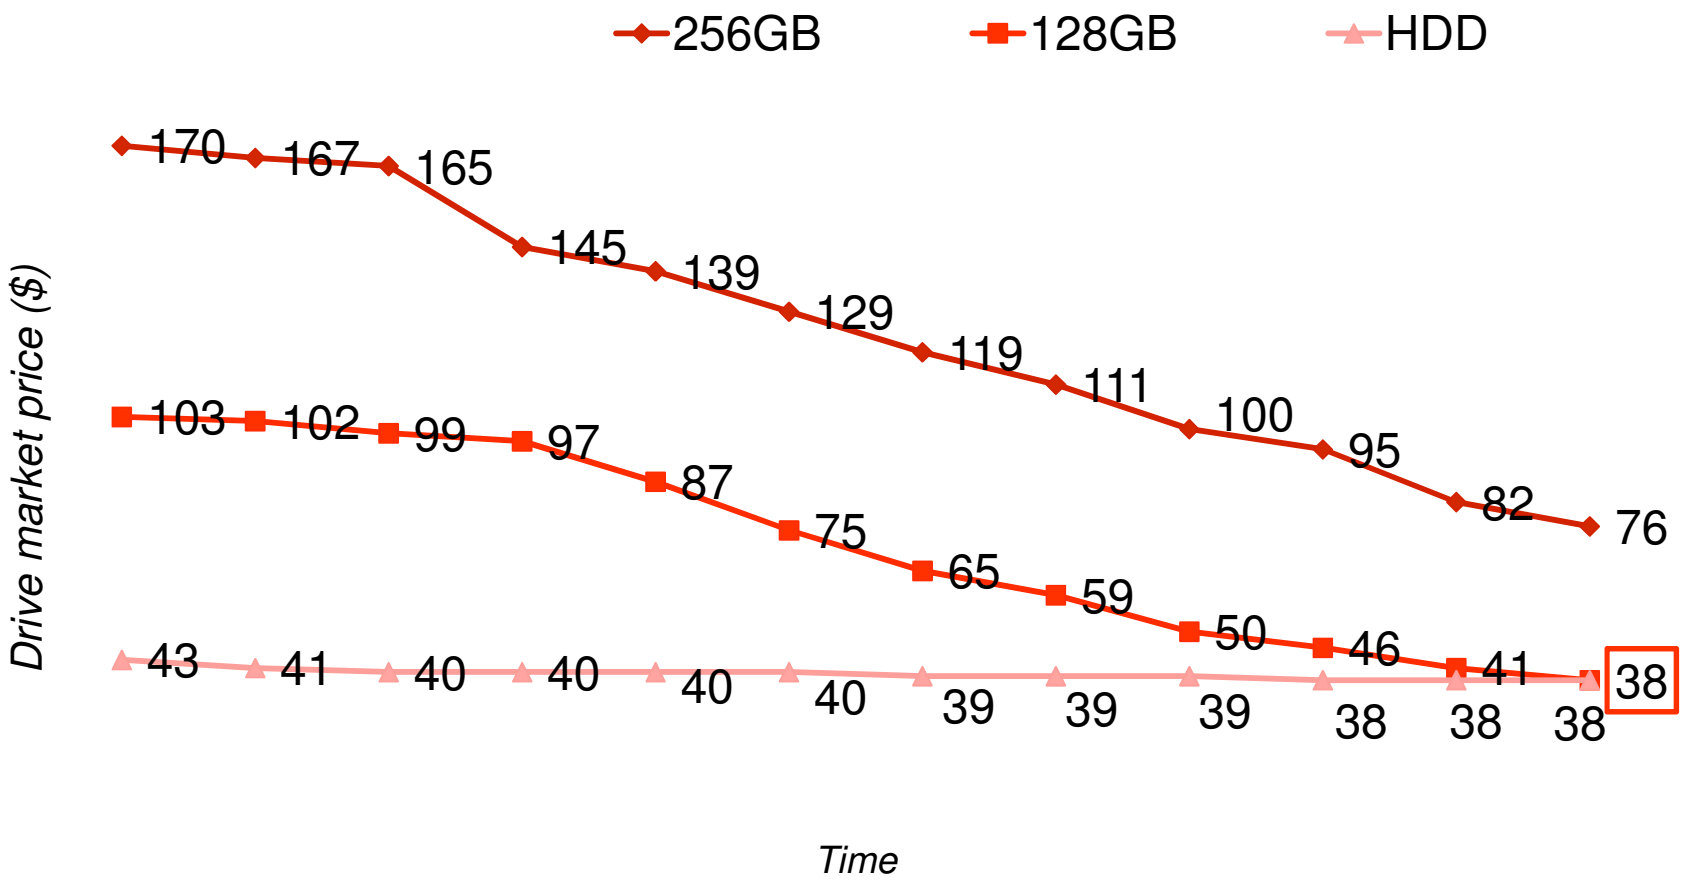

Developing DRAMless architecture

- 2015 SSD Architecture: Mainstream vs Entry-level

128GB SSD	Mainstream	Entry Level
NAND CH	4	**2**
NAND speed	333MT/s	**400MT/s**
Capacity	64Gb MLC 16dies	**128Gb TLC** **8dies**
Memory	16bit DRAM I/F	**DRAMless**

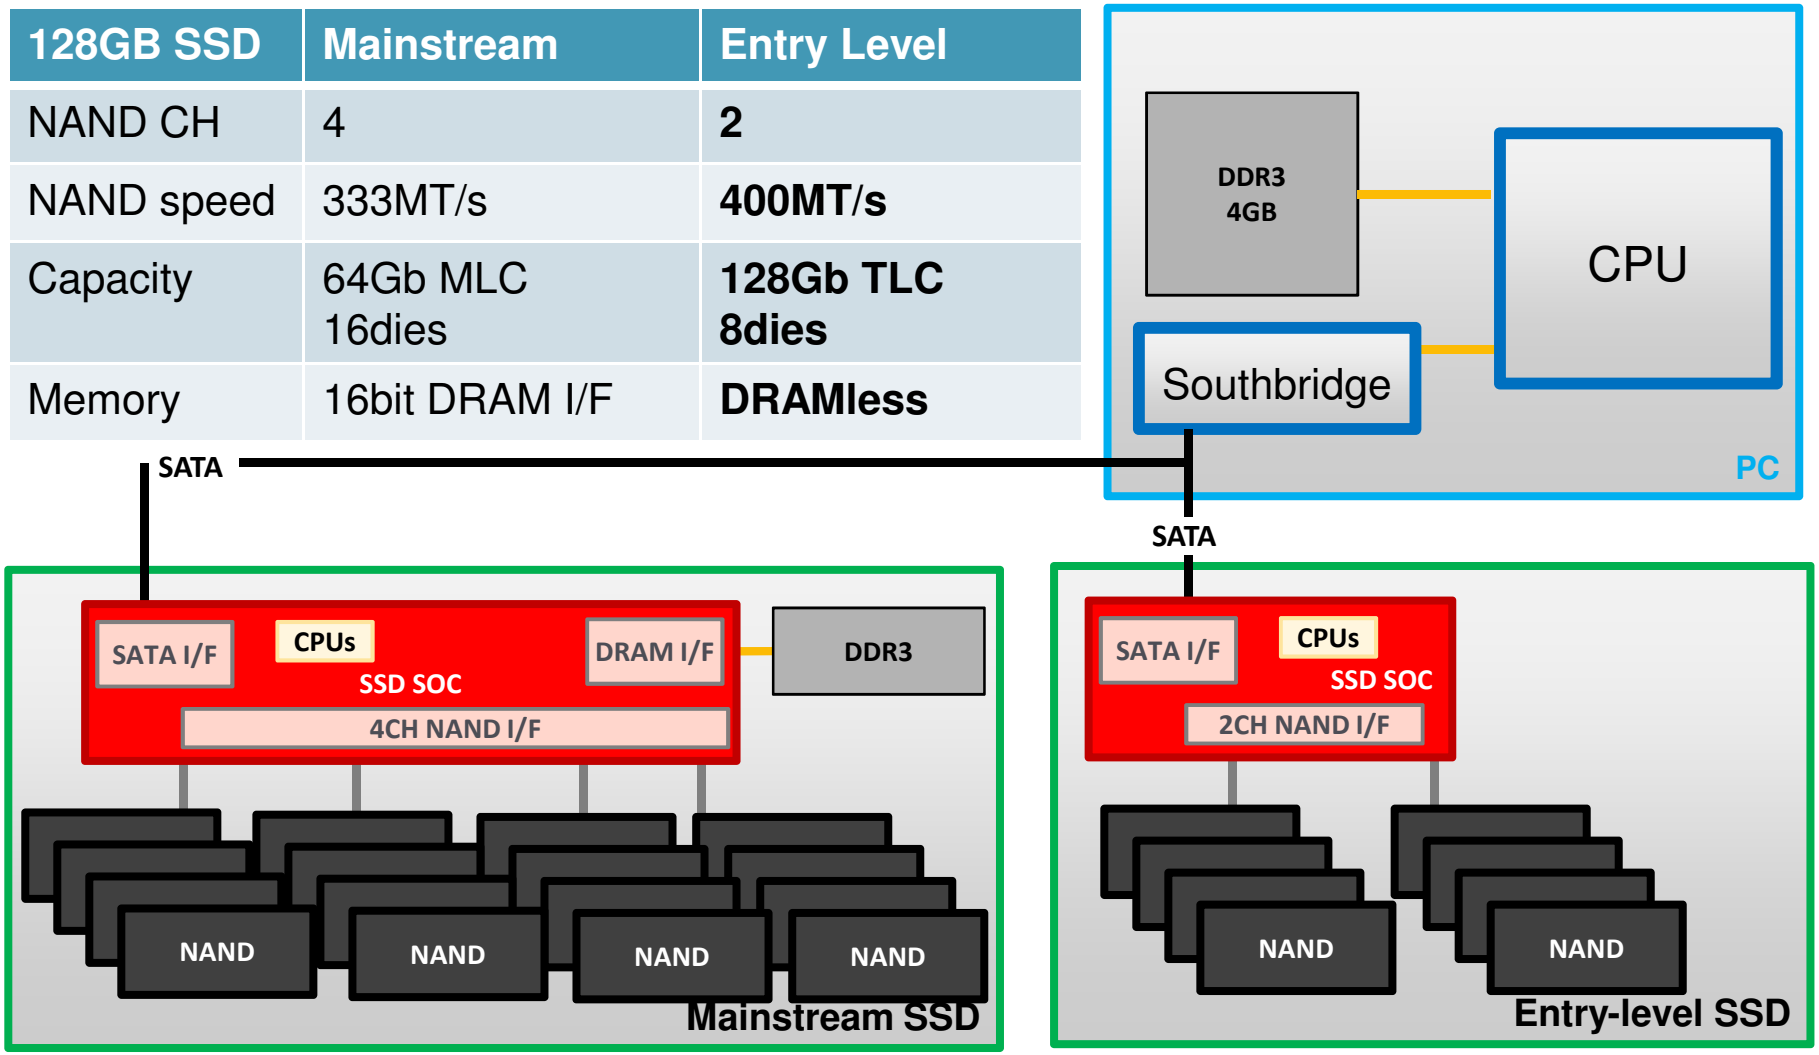

What I can get from a $30ish SSD?

- 20X HDD performance in full testing range
- Or 100X HDD IOPS in small 4GB or 8GB range
- And 1/20 space with M.2 2230 or even smaller with BGA
  - 1/10 in XY
  - 1/2 in Z-height
- And <2mW in Devslp mode

Enabling Host Memory Buffer

- Same BOM cost as Entry Level SATA drive!
- PCIe Gen3x1 NVMe SSD requests Host to allocate Memory space
  - In this case study: 128MB for the Look-up-Table
  - But it can claim any custom size of Host DRAM (partial LUT)

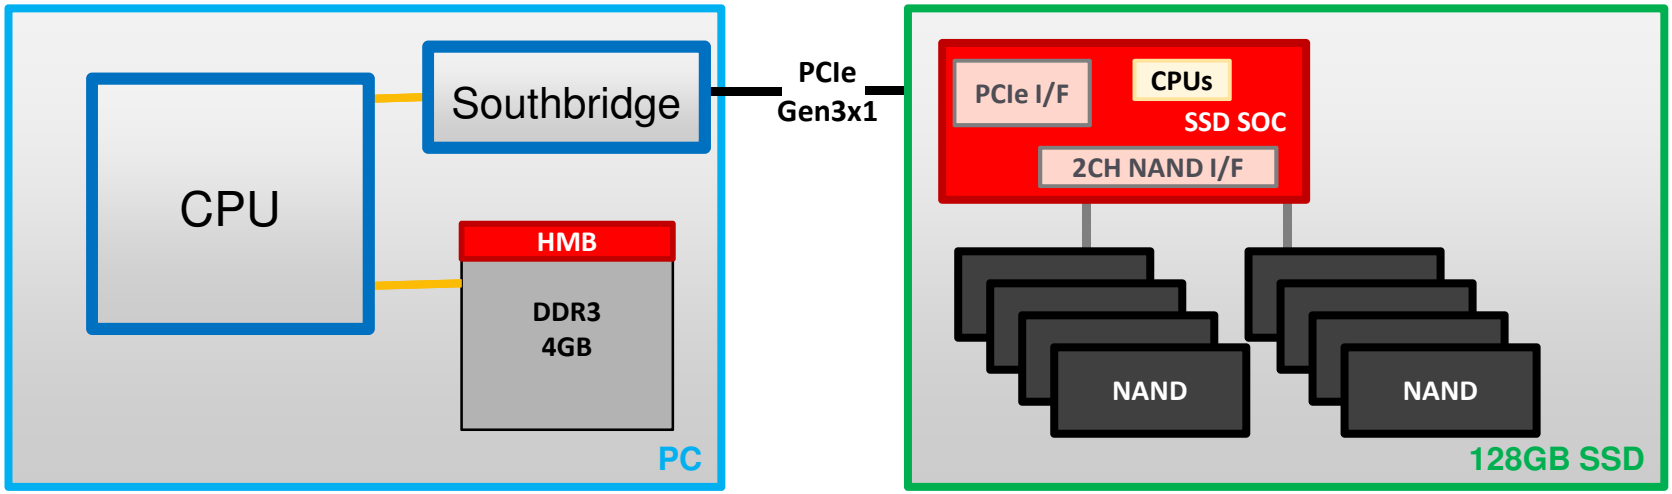

drawing not to scale

Performance benchmark 128GB TLC SSD

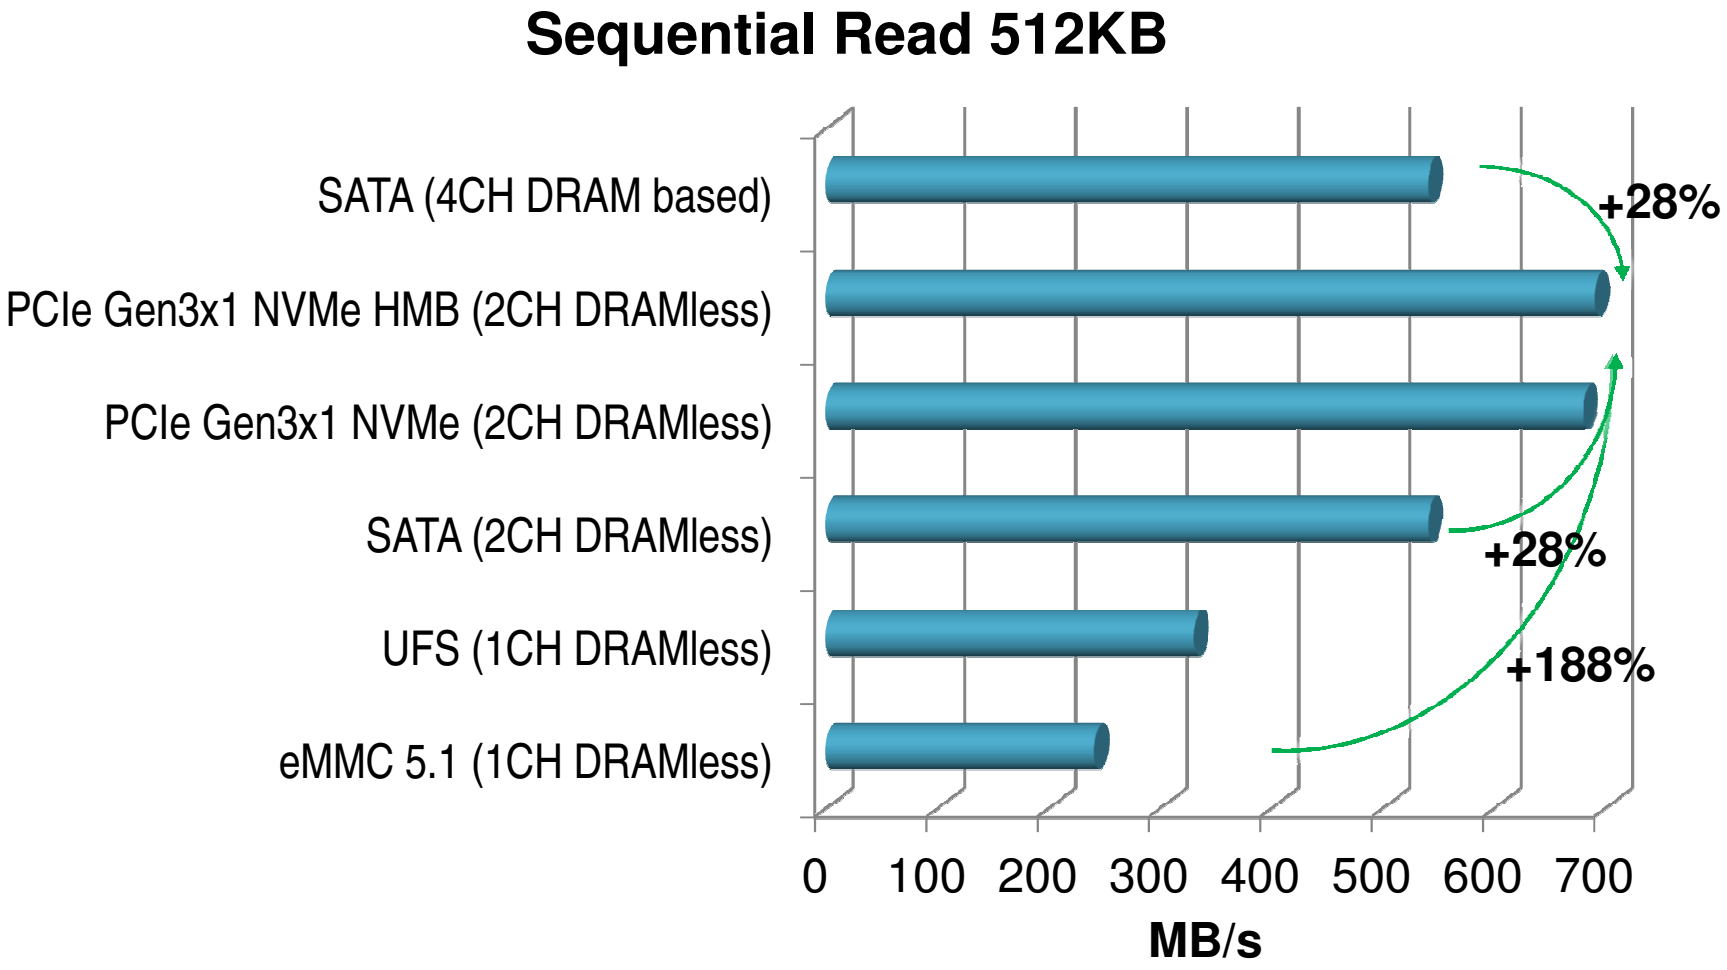

Test configuration: Chipset(Z97), Windows 8.1, Intel NVMe driver, CrystalDiskMark

Host Memory Buffer reduces latencies

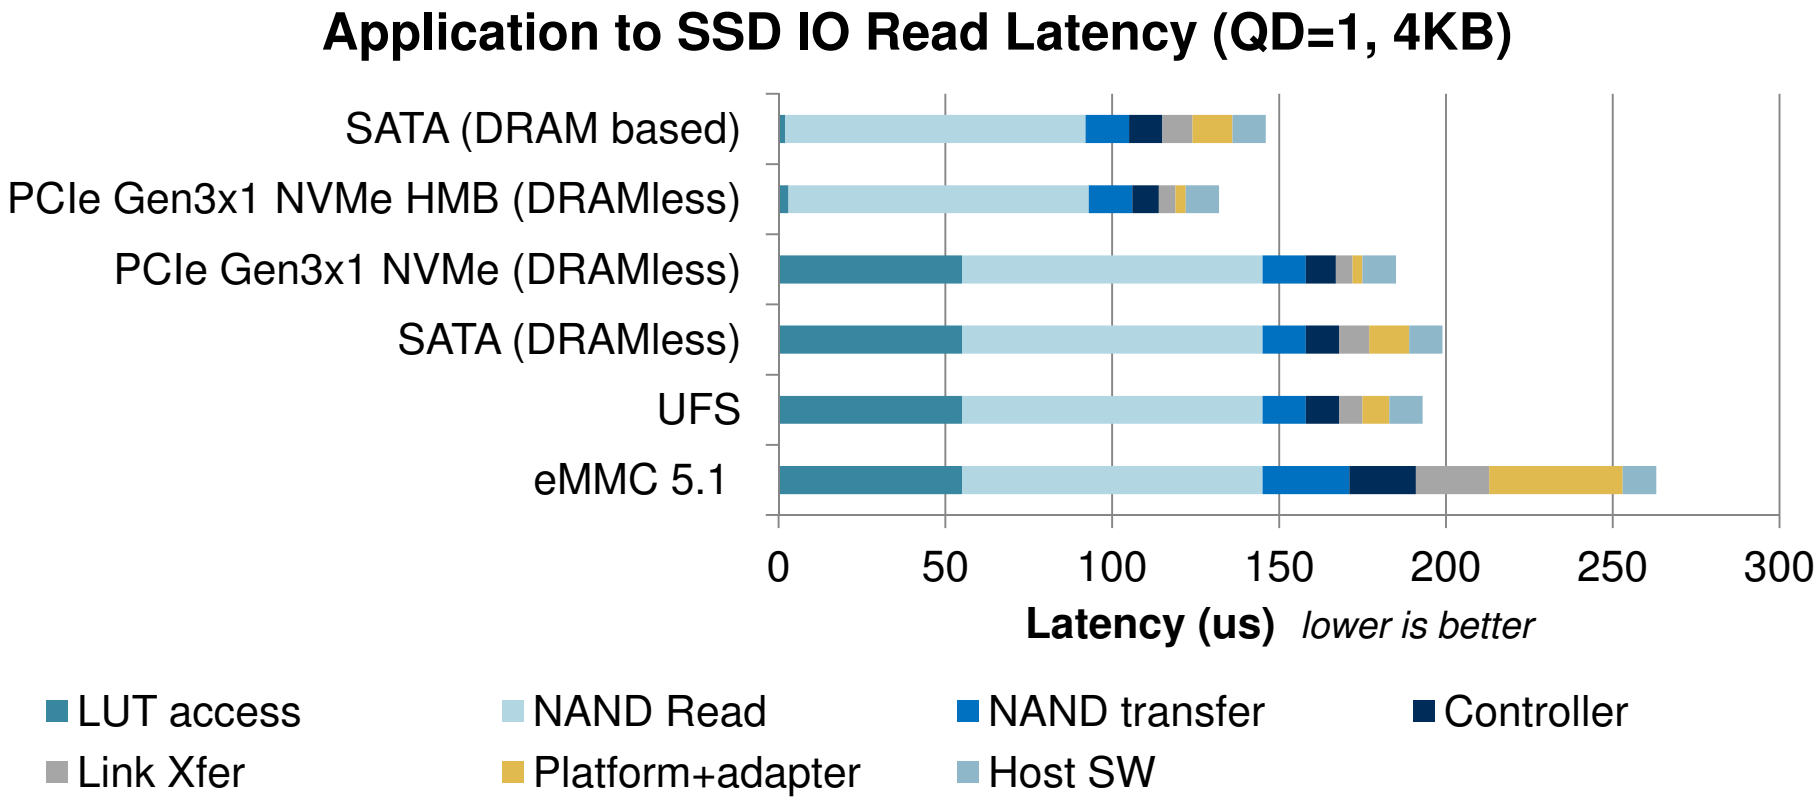

- Host Memory Buffer significantly reduces latencies:
  - Boosts IOPS with 50% over SATA and 100% over eMMC

Performance benchmark 128GB TLC SSD

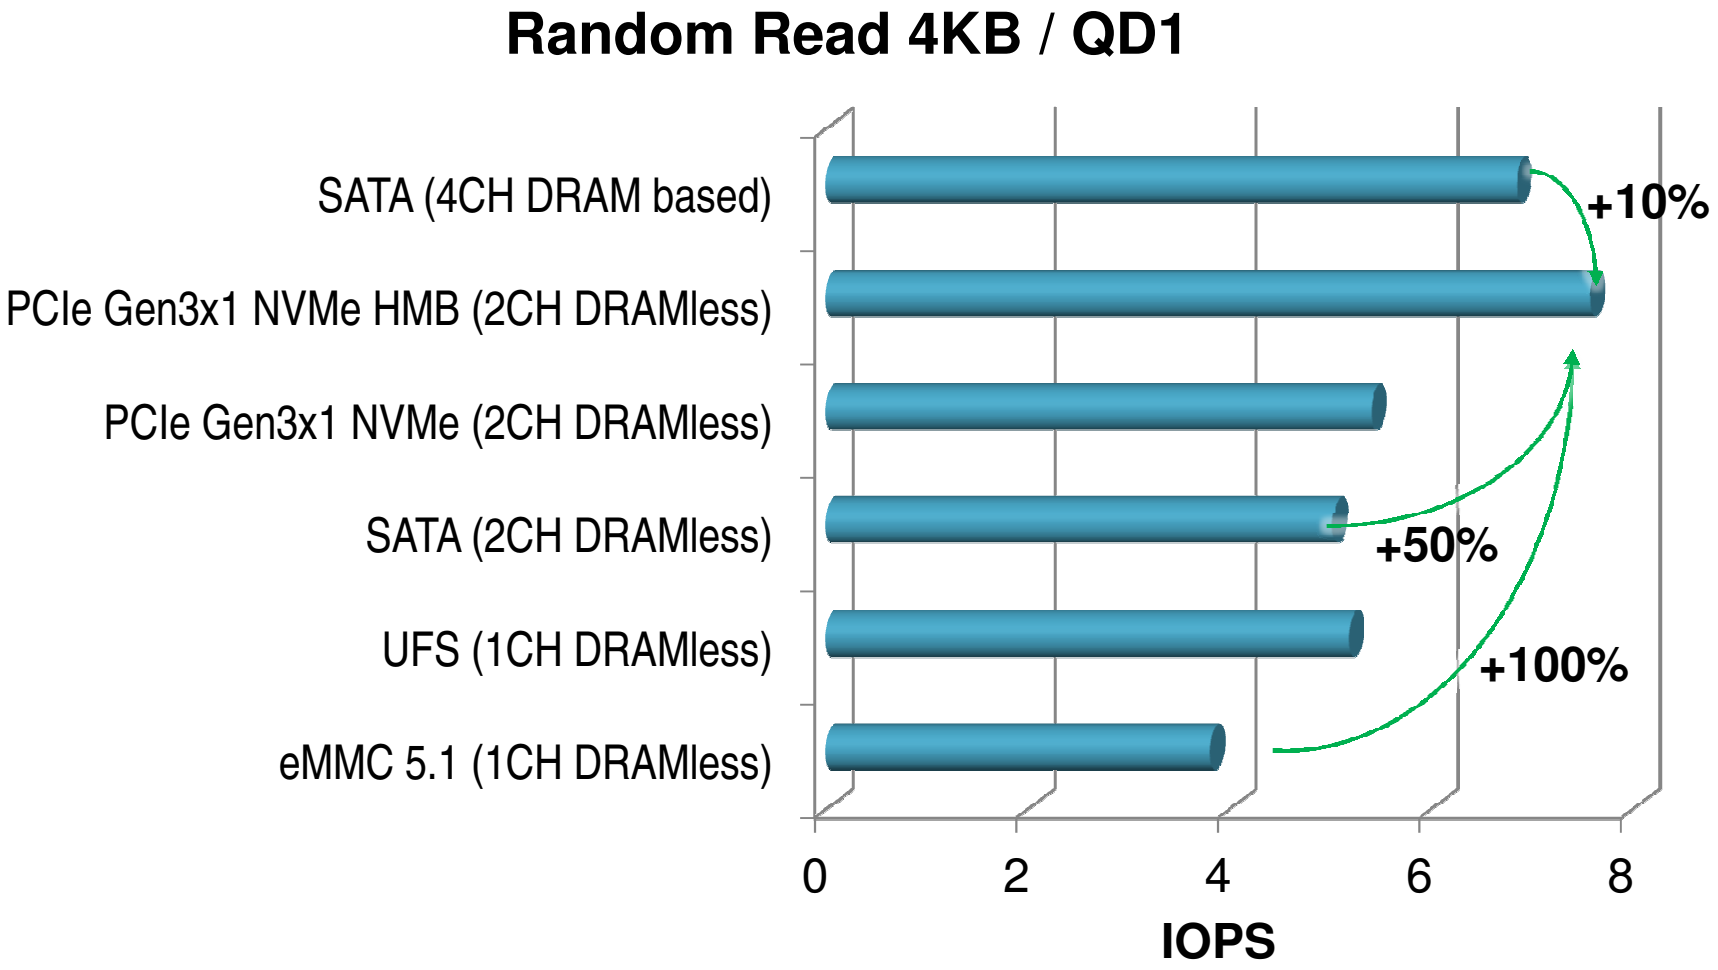

Test configuration: Chipset(Z97), Windows 8.1, Intel NVMe driver, CrystalDiskMark

Performance benchmark 128GB TLC SSD

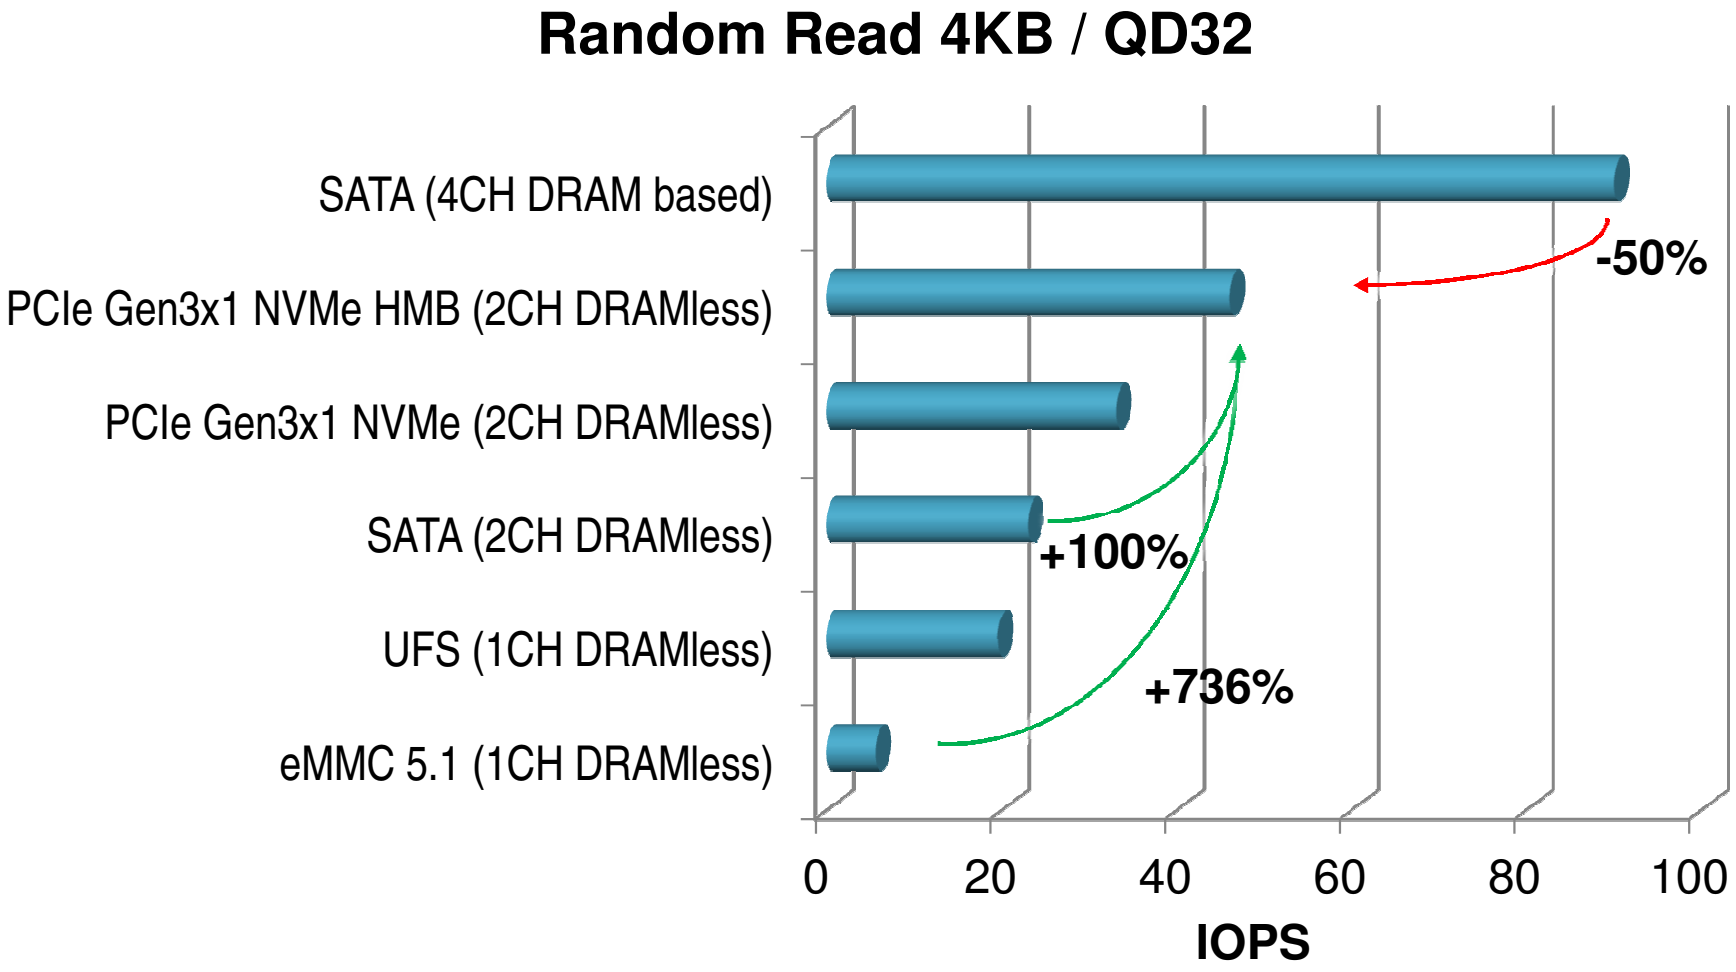

Test configuration: Chipset(Z97), Windows 8.1, Intel NVMe driver, CrystalDiskMark

Client Workloads are light weight

- Review sites have developed trace based benchmarks
  - Capturing user behavior on Windows machine
  - Feeding commands to the SSD
- Example of queue depth distribution:
  - QD1-4 cover >90% of all cases

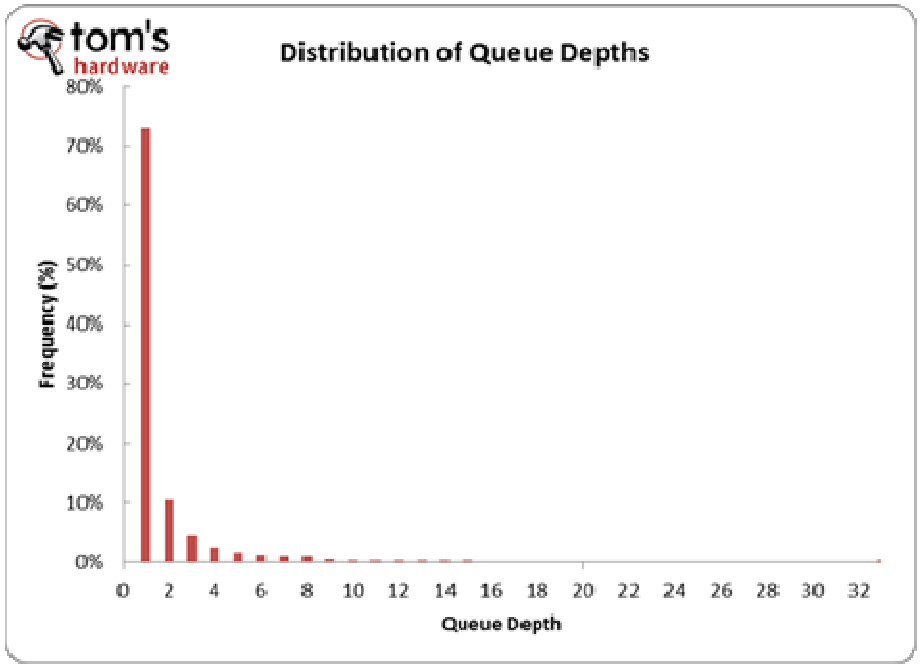

HMB Ecosystem development

- Operating system inbox driver & Bios support
  - Windows, Linux, Android, etc
- Chipset to support PCIe as storage interface
  - Coming in this year!
- Leading-edge SSD controller to support HMB
  - Marvell 88NV1140
- Deployment to PC OEMs
  - Configuration, integration and qualification
- Marvell is working in all areas to drive HMB

Summary

- A DRAMless SSD delivers good-enough performance and capacity for an entry-level client PC system
- Moving from SATA to PCIe Gen3 NVMe increases sequential read speeds and reduces latencies
- Enabling Host Memory Buffer boosts IOPS performance significantly
  - 50-100% compared to other DRAMless solutions
- Marvell's 88NV1140 enables this without adding BOM cost
  - We are working with PC OEM to launch HMB soon

The Opportunity...

- ...is bigger than you'd expect:
  - Tablets/Convertables become productivity devices
  - eMMC performance doesn't scale
  - PCIe Gen3x1 is a strong alternative
  - Advanced Marvell NVMe SSD controllers under development!

- The Future of NVMe is NOW!
  - Marvell 88NV1140
  - PCIe Gen3x1 NVMe 1.2 with Host Memory Buffer support

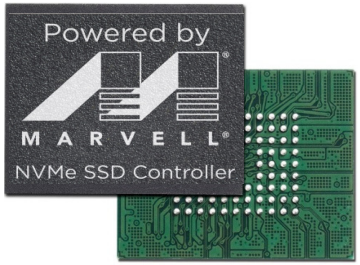

BGA SSD

$0.25 size

M.2230 SSD


MARVELL