Stockage non volatile update sur les technologies et interfaces

Stanislas Odinot, avant vente, division entreprise
Intel Corporation

Legal Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Le stockage un élément clé de l'évolution des usages

Usages et Architecture

(high-end) Scale UP

Vs. (high-end) Scale-OUT

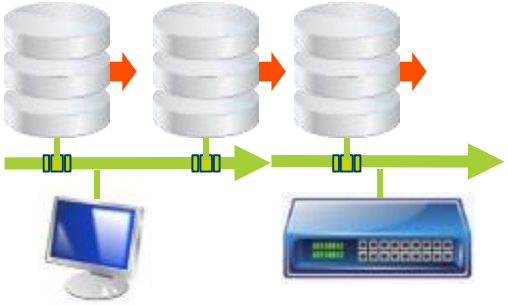

Interface

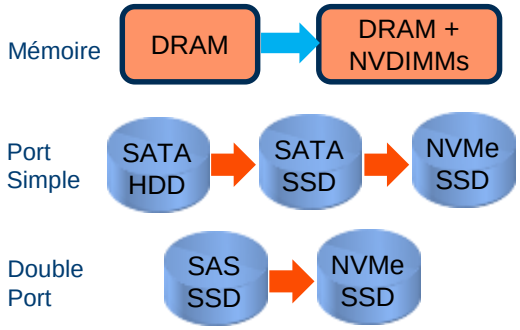

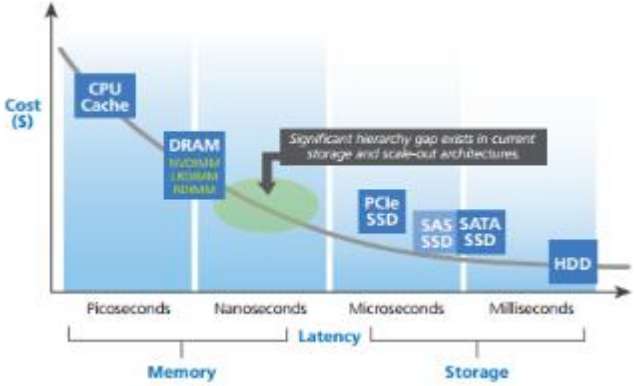

Gap performance/prix

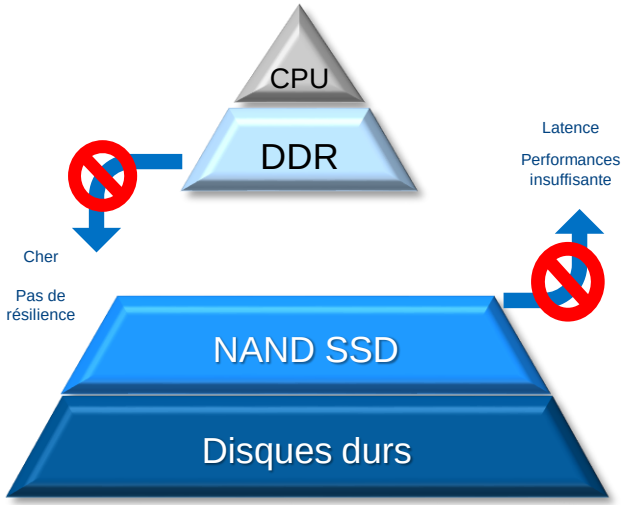

Chronologie du NVMe™
(sous l'angle Datacenter)

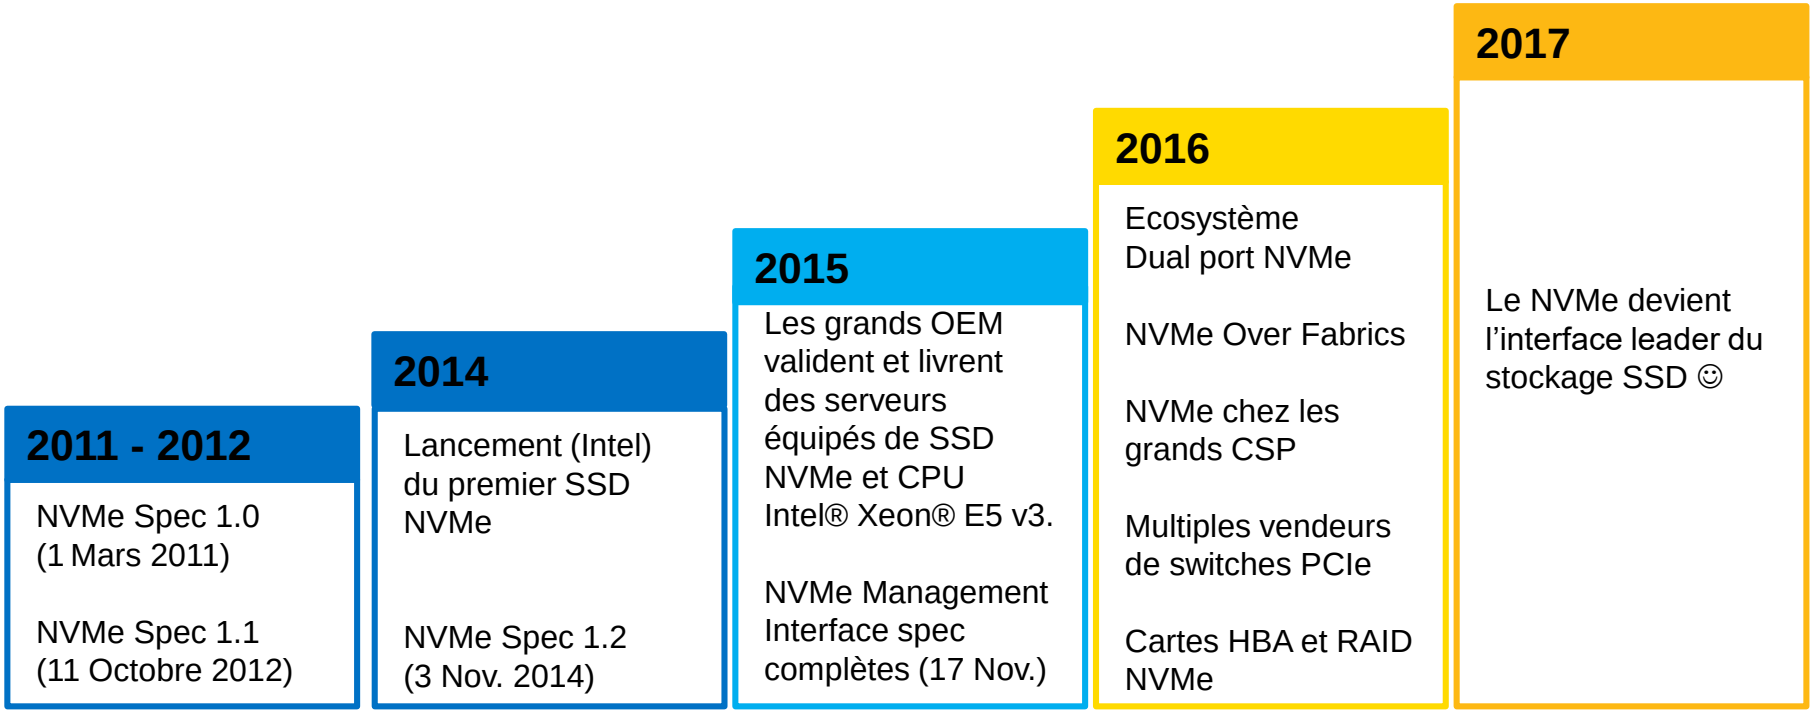

Côté matériel et pilotes

Formats

AIC

2.5”

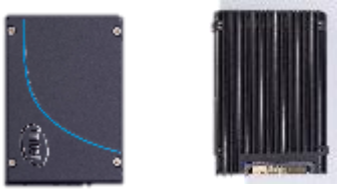

M.2

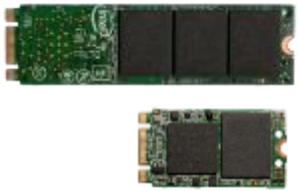

42, 80, and 110mm

Connecteur U.2 (SFF8639) & Intégration serveur

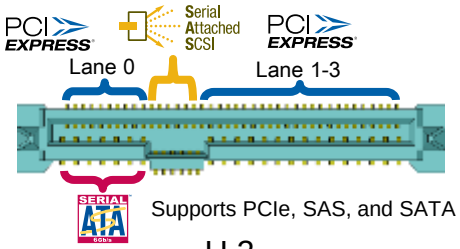

Supports PCIe, SAS, and SATA

U.2
connector and backplane

U.2 drive / U.2 SSD

U.2
cable

U.2
host connector

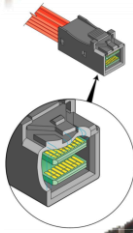

miniSAS HD

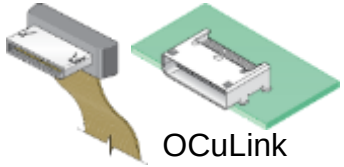

OCuLink

Pilotes

Natif / “in-box”

RedHat 6.5+, 7.0+

SUSE SLES 11 SP3 SLES 12

Ubuntu 13, 14

FreeBSD

Solaris 11.2

VMware ESXi 6.0

EFI

Windows Server 2012 R2 Certified · Windows 8 Compatible · Windows 10

Windows 8.1

Drivers NVMe à installer

Windows Server 2012 Certified · Windows Server 2008 R2 · Windows 7

VMware ESXi 5.5

U.2 PCI Express*: Un écosystème fort et en croissance

2U 24 bay NVMe Storage

24x U.2

SYS-2028U-TN24R4T+ SSG-2028R-NR48N

24x U.2 48x U.2

QuantaGrid D51B-2U

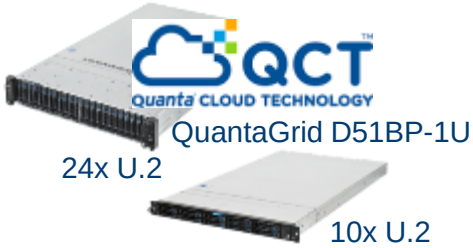

24x U.2

QuantaGrid D51BP-1U

10x U.2

2U All-Flash NVMe* Storage

24x U.2

ASR-3100

16x U.2 in 1U!

“Lightning”

2U 60x U.2 Storage!

TN70A – B7086

24x U.2 + 4 AIC in 2U

SB122A-PH

10x U.2 in 1U

2U 12 bay NVMe Storage

12x U.2

2U 24 bay NVMe Storage

24x U.2 in 2U

NVMe Dual Port, c'est dispo.

Intel D3700 et D3600 annoncés le 31 mars 2016

Accès redondants aux SSD NVMe

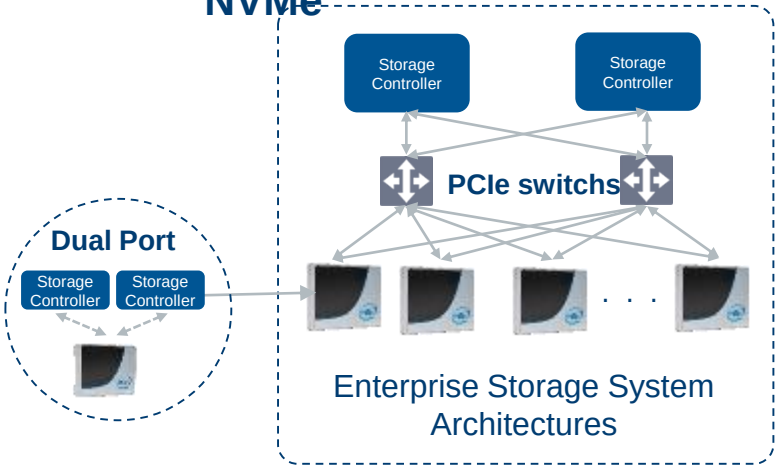

Reservations

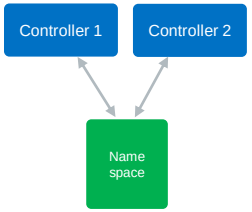

Manages who can write to the device and when

Multiple Namespace

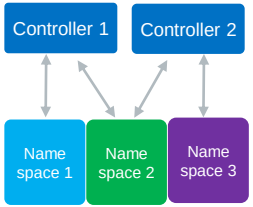

Smaller granularity in managing data across drives

Scatter Gather List

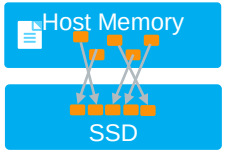

Connects scattered data in host memory to reduce IO operations

In Controller Memory Buffer

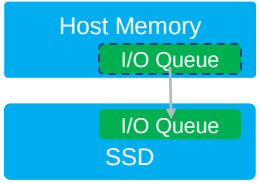

Moves IO queue to the SSD to reduce latency

NVMe over Fabric, ça progresse

- **Les spécifications sont complètes** et disponibles sur nvmexpress.org le 16 Mai 2016 (c'est l'objectif)
- **Une fois publié, les Host et Target seront disponibles** puis upstreamées dans le kernel d'ici ~ la 4.8

- Pour garantir que le « transport » soit **fabric agnostique**, il a été testé sur des cartes InfiniBand™*, iWARP, RoCE, RDMA
  - La partie Host a aussi été testée avec plusieurs target.

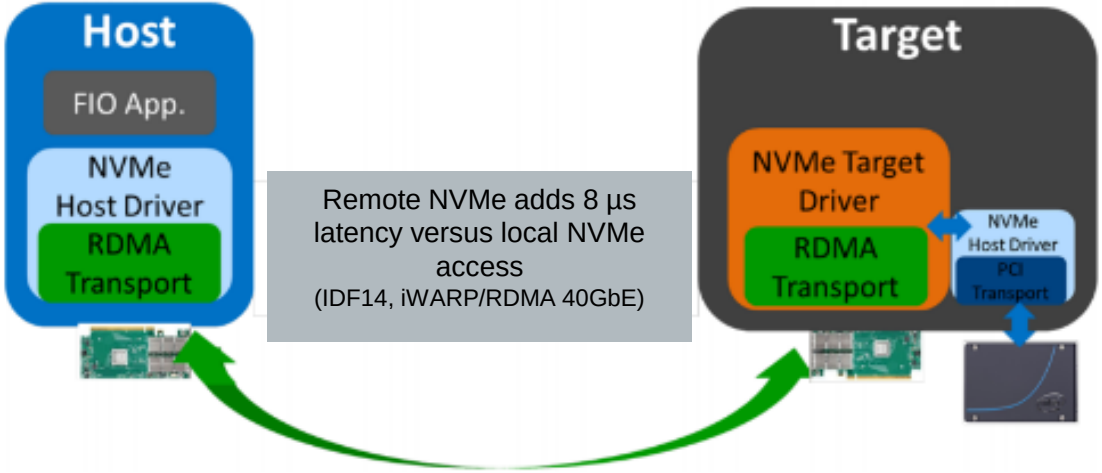

NVM Express™*

Plus de bande passante et d'IOPS

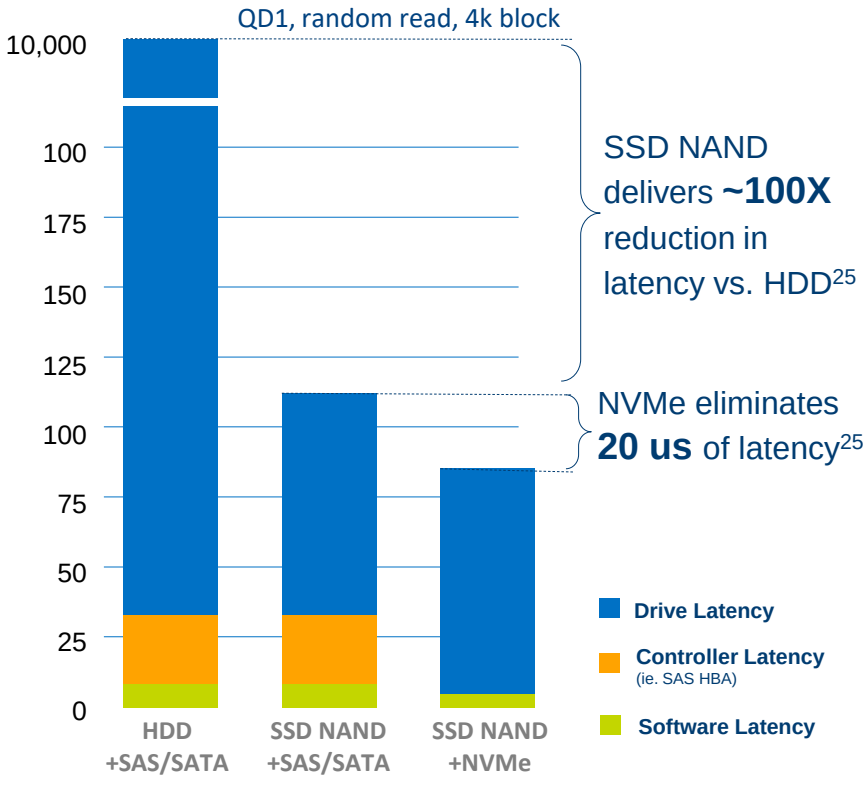

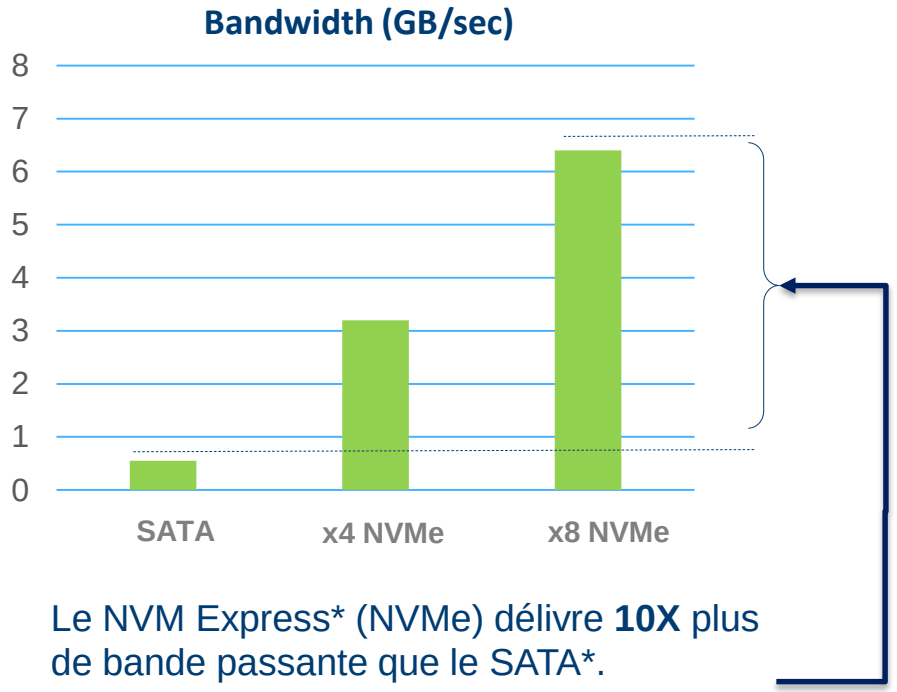

Le NVM Express* (NVMe) délivre **10X** plus de bande passante que le SATA*.

Et encore plus avec le PCIe Gen 4

See "Footnotes and Disclaimers" page for note

Comparaison des SSD Intel PCIe x8, PCIe x4, SATA 3.0

IOMeter + larges QD sur les tests de bande passante

	S3710	P3700	P3608
			
	SATA 3.0 (6Gb/s)	PCIe 4x	PCIe 8x
Workload	**Capacity: 1.2 TB**	**Capacity: 1.6 TB**	**Capacity: 1.6TB**
Random 4k Read	Up to **85k** IOPS	Up to **450k** IOPS	Up to **850k** IOPS
Random 4k Write	Up to 45k IOPS	Up to 150k IOPS	Up to 150k IOPS
Random 4k 70/30 R/W	Up to 63k IOPS	Up to 240k IOPS	Up to 300k IOPS
Sequential Read	Up to **550** MB/s	Up to **2800** MB/s	Up to **5000** MB/s
Sequential Write	Up to 520 MB/s	Up to 1900 MB/s	Up to 2000 MB/s

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Configurations: Performance claims obtained from data sheet, P3608 all random test QD256/8 workers, sequential QD128/128k. S3710, all random test QD32/4K on 70/30, sequential QD32/128K. P3700 all random test QD32/4 workers, sequential QD128/128k

Les PCI Express* SSD plus d'IOPs dans les applications

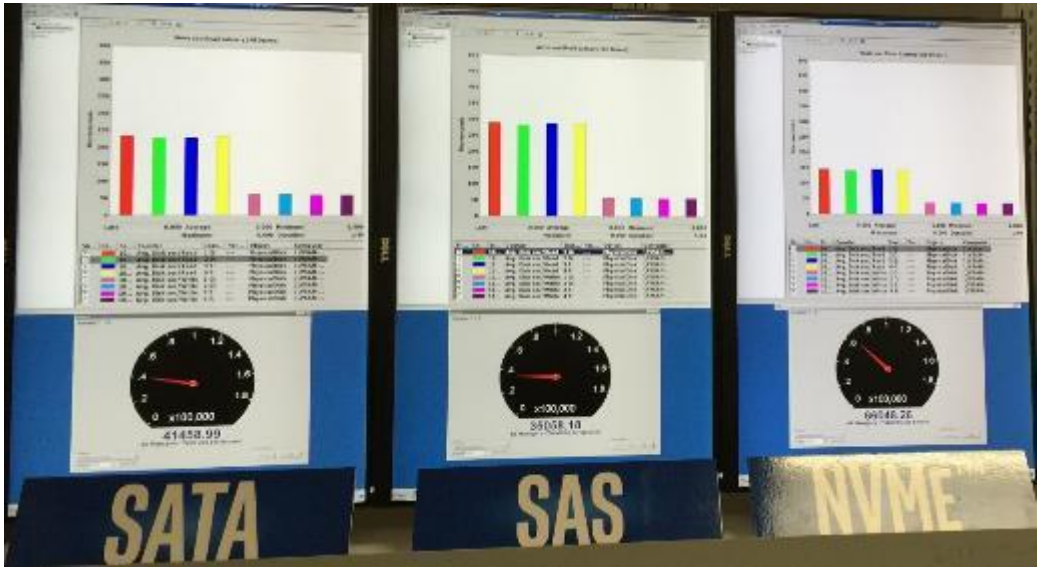

SSD Metric	6 Gbps SATA	12 Gbps SAS	PCIe* NVMe*
Total IOPs	**41.5k**	**35.1k**	**66.0k**
Avg Read Latency	~240 µS	~300 µS	~140 µS
Avg Write Latency	~50 µS	~50 µS	~40µS

Performances des applications avec ESXi* associé au NVMe :

1.9x plus rapide que le SAS
1.6x plus rapide que le SATA

4kB 70/30 R/W **QD8** using IOMeter, 4 workers

Vmware* ESXi 6.0 Update 1. Three identical Windows* Server VMs running 4kB 70/30 R/W QD8 using IOMeter. 4 workers per SSD device. SSDs used: NVMe- Intel® SSD Data Center P3700 Series (800 GB). SAS- 12 Gb/s SAS SSD (800 GB) on dedicated 12 Gb/s SAS HBA. SATA- Intel® SSD Data Center S3710 Series (800 GB) on dedicated 12 Gb/s SAS HBA. Supermicro* 2U SuperServer® 2028U-TNR4T+. Dual Intel® Xeon® Processor E5-2699 v3 (45M Cache, 2.30 GHz). 192 GB DDR4 DRAM. Boot drive: Intel® SSD Data Center S3710 Series (200 GB). Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

L’arrivée de la 3D NAND…

On part de ceci…

2D NAND

…arrive à cela…

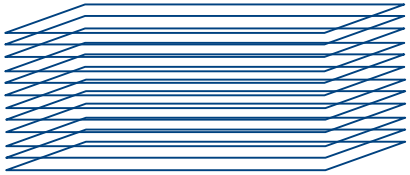

3D NAND

…délivre:

- ✓Des lectures et écritures **plus rapides**
- ✓Meilleure **efficacité énergétique**
- ✓**Améliore l’endurance** et offre à la TLC différents niveaux d’endurance

See “Footnotes and Disclaimers” page for note

NAND Versus la technologie 3D XPoint™
SSD IOPS vs. Queue Depth

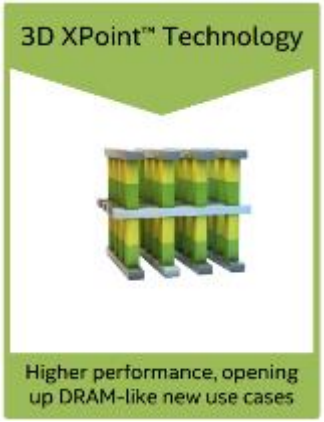

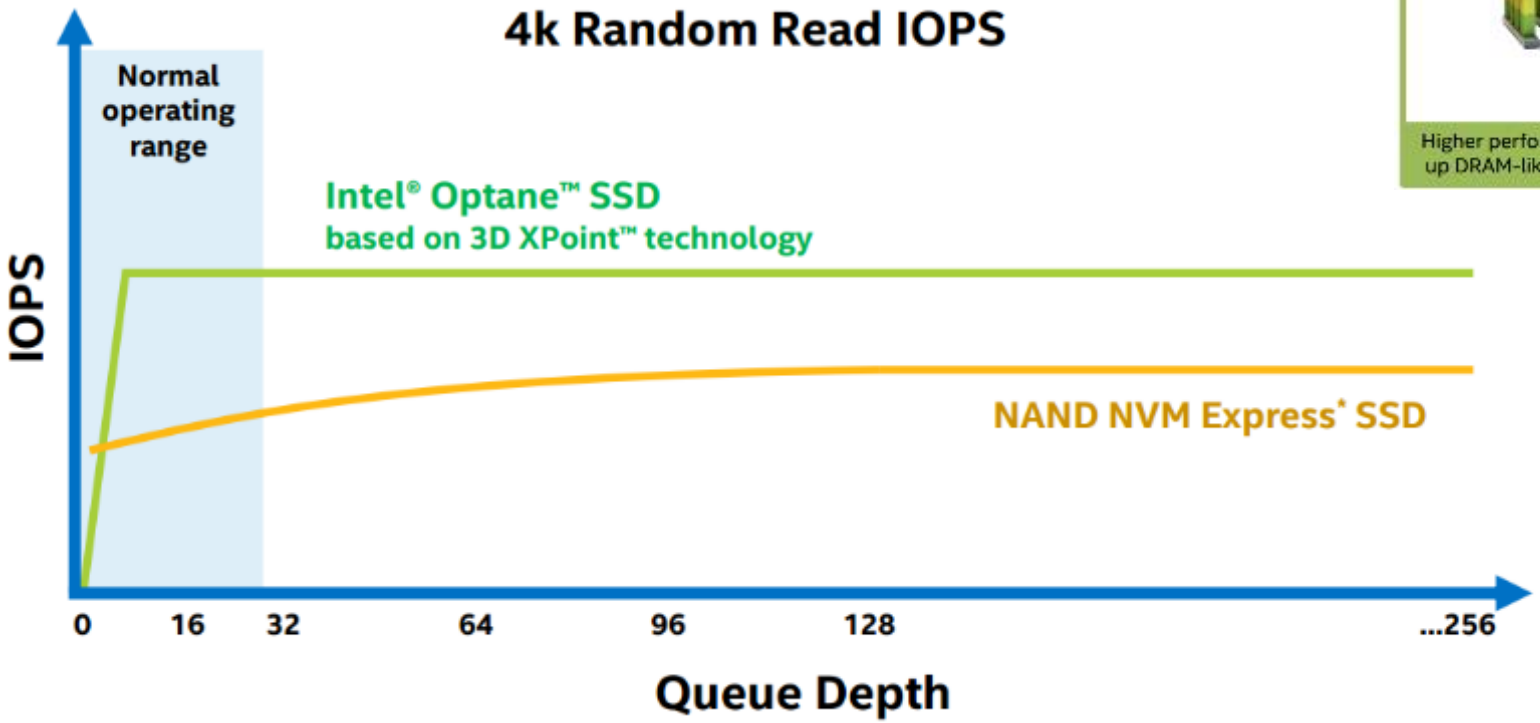

NVMe avec la technologie 3D Xpoint™

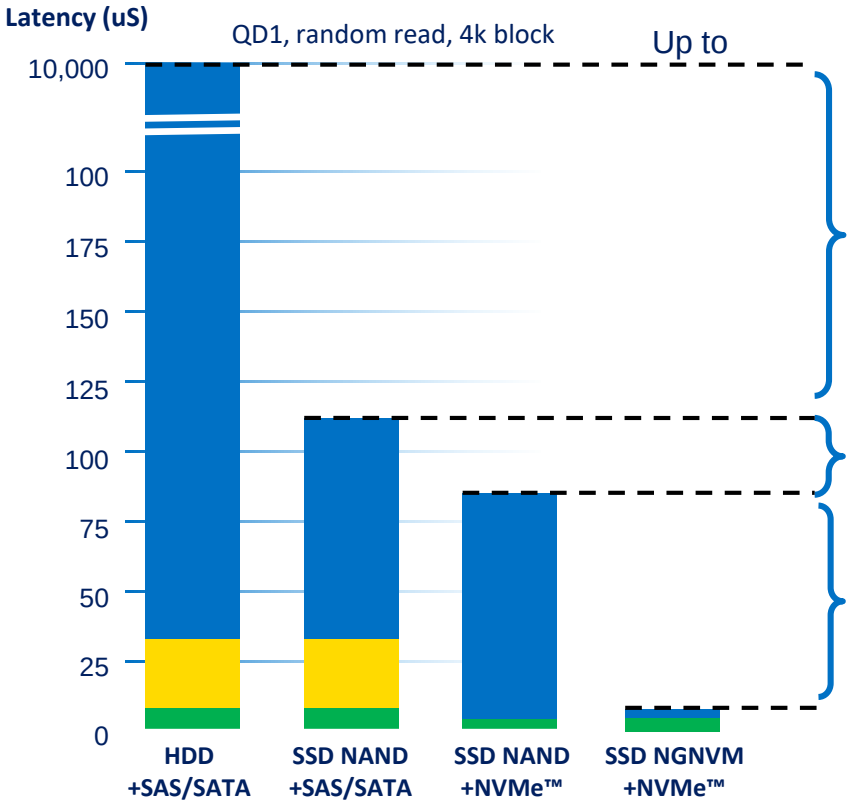

Les SSD base technologie NAND offrent une latence **~100X** inférieure à celle des HDD

NVM Express*(NVMe) éliminent aujourd'hui **20 µs** of latence

La technologie 3D XPoint™ offre au NVMe une latence **~10x** plus faible comparée aux SSD base NAND

Drive Latency

Controller Latency (ie. SAS HBA)

Software Latency

Source: Storage Technologies Group, Intel

Intel® Optane™ (prototype) vs Intel® SSD série DC P3700 à QD=1

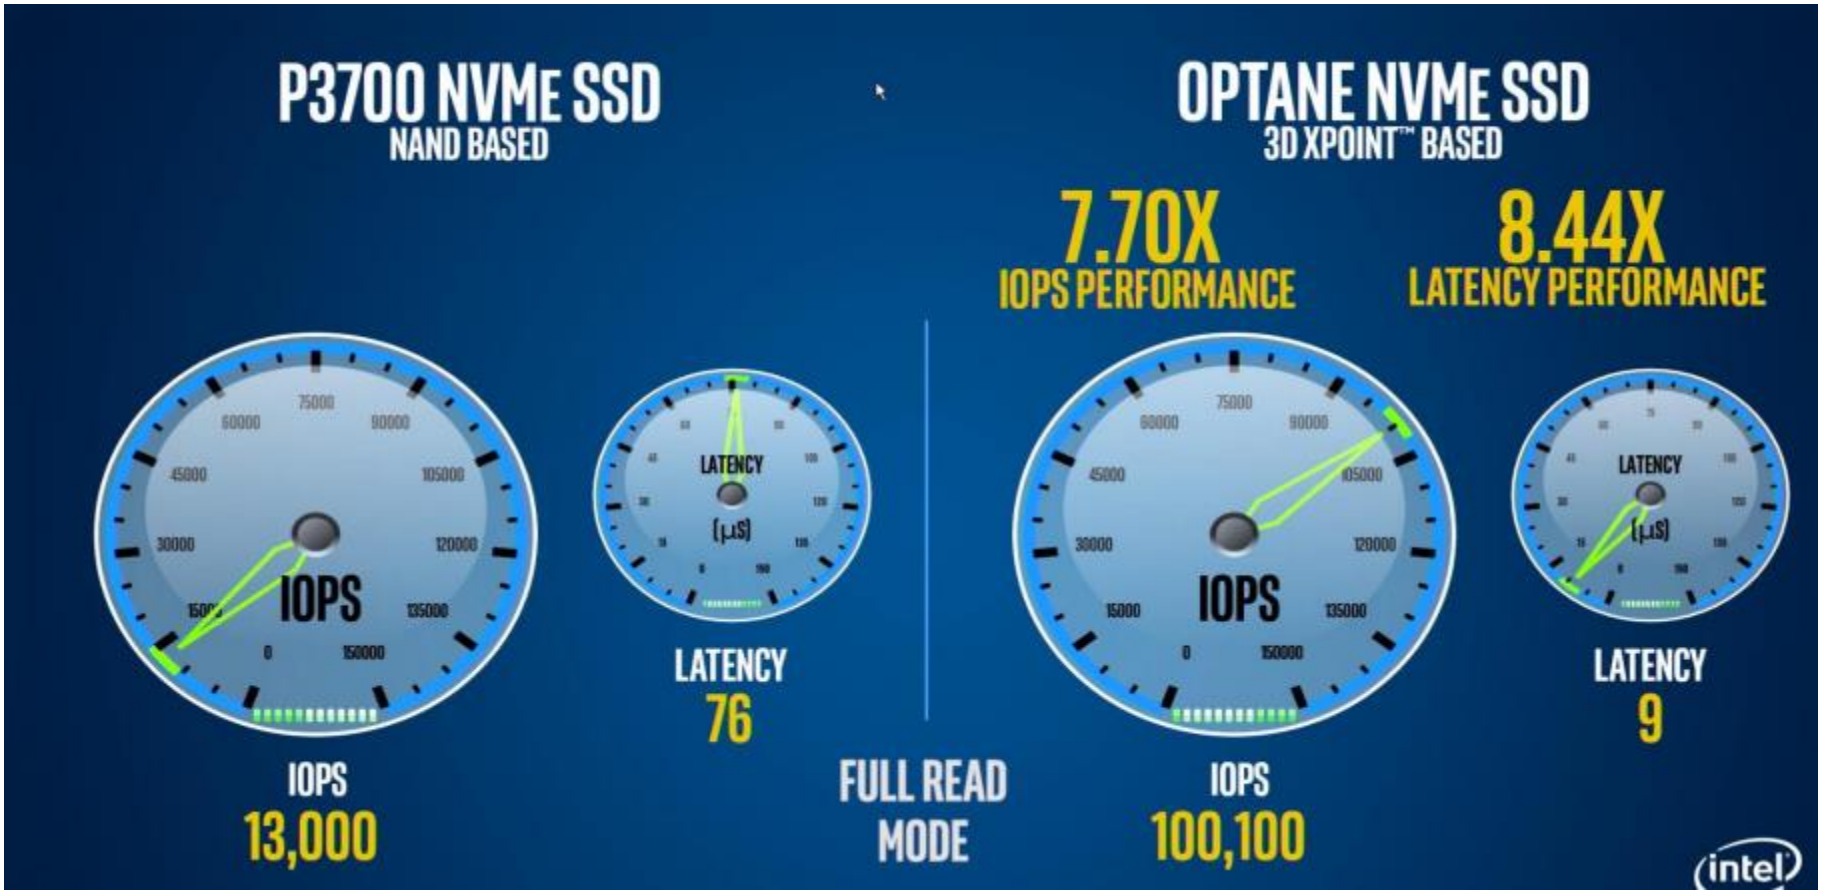

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance. Server Configuration: 2x Intel® Xeon® E5 2690 v3 NVM Express* (NVMe) NAND based SSD: Intel P3700 800 GB, 3D Xpoint based SSD: Optane NVMe OS: Red Hat*7.1

Intel DIMM basée sur la technologie 3D Xpoint™

- DDR4 physiquement et électriquement compatible
- Supporté sur la prochaine génération de plateforme Intel® Xeon™
- Jusqu'à 4X plus de capacité mémoire par système à un prix significativement plus faible que la DRAM
- Peut délivrer les bénéfices d'une grande capacité mémoire sans modification de l'OS ou des applications

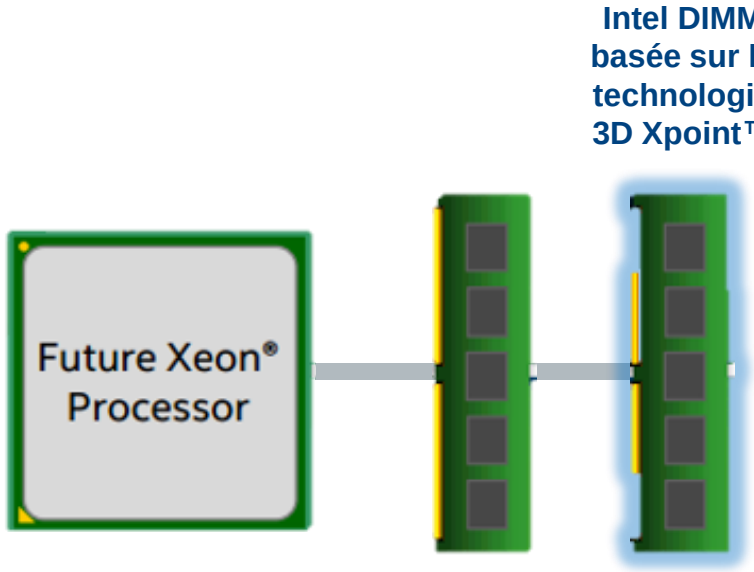

J’espère vous avoir convaincu…

- Qu’en 2016 le protocole et les équipements à base de PCIe/NVMe arrivent à maturité.
- Que les SSD PCIe/NVMe surpassent les meilleurs SSD SATA et SAS sur de nombreux critères dont la performance.
- Que les meilleurs SSD PCIe/NVMe équipés de NAND sont encore loin des performances de la DDR4, mais que la 3DXPoint devrait réduire le gap.

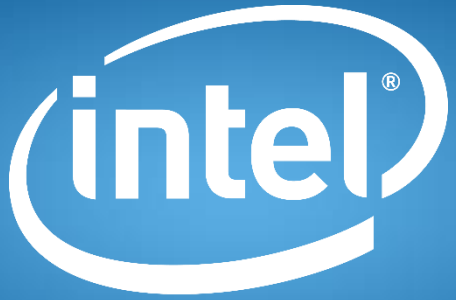

experience
what's inside™

Le PCIe* prépare l'arrivée des SSD Intel® Optane™

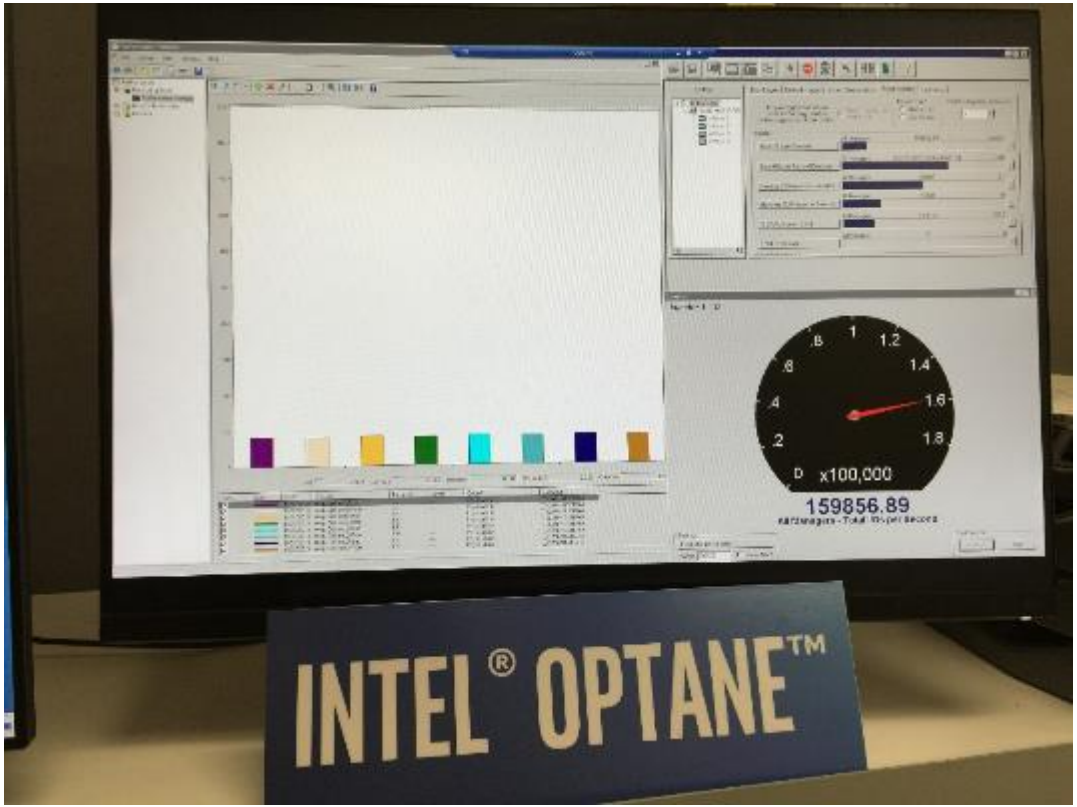

Performances des applications avec ESXi* associé à l'Intel® Optane™
2.4x plus rapide que de la NAND sur PCIe!

IOPS: 4kB 70/30 R/W **QD8** using IOMeter. 4 workers

Intel® Optane™
+159k IOPS

Intel DC série P3700
66k IOPS

Vmware* ESXi 6.0 Update 1. Windows* Server VMs running 4kB 70/30 R/W QD8 using IOMeter. 4 workers per SSD device. SSDs used: NVMe- Intel® SSD Data Center P3700 Series (800 GB) achieving 66k IOPs (shown on slide 10), and Intel prototype SSD using Intel® Optane™ Technology (shown here). Supermicro* 2U SuperServer® 2028U-TNR4T+. Dual Intel® Xeon® Processor E5-2699 v3 (45M Cache, 2.30 GHz). 192 GB DDR4 DRAM. Boot drive: Intel® SSD Data Center S3710 Series (200 GB). Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

*PCIe = PCI Express. Other names and brands may be claimed as the property of others