

# SSDs are Here – The Next Wave in Non-Volatile Memory Driven Storage Modernizations

Greg Matson

Director of Strategic Planning, Non-Volatile Memory Solutions Group, Intel Corporation

Guest Speaker: Ruiping Sun, Principal Architect, Yahoo!

**SSDS002**

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A



**STORAGE DEMANDS INCREASING  
CLOUD TRAFFIC QUADRUPLES BY 2018,  
AND DIGITAL UNIVERSE IS...**


**8ZB IN 2015 GROWING TO 44ZB BY 2020**



# Agenda


- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# The Future is Here – NVMe Express™





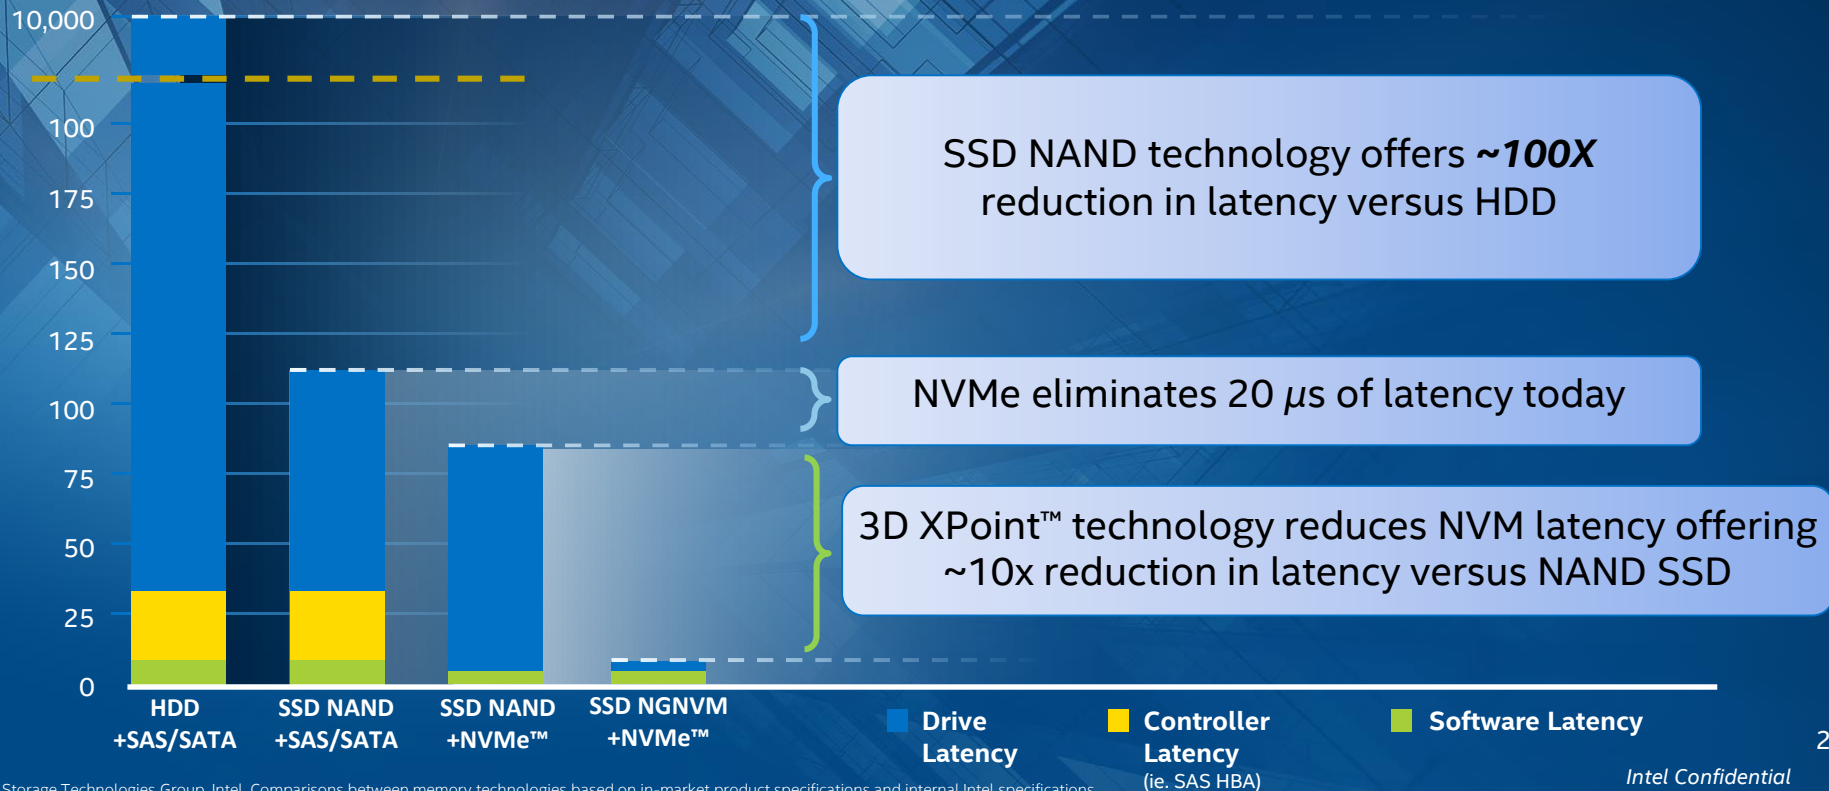
# Eliminates Storage Interface Bottleneck with NVMe™ and PCI Express®



**NVMe provides more than 10X the Bandwidth of SATA. Even More with Gen 4**

# NVMe™ is the next Quantum Leap!

Latency (uS)





# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# How are Data Center SSDs Being Used Today?



TEMP



TIER



CACHE




BOOT

# SSDs are enabling **CLOUD COMPUTING**



Instant access to  
products and services

## **End User**



Enterprise grade  
security and reliability

## **Private Cloud**



Fast development, time to  
market, and SLAs

## **IaaS/PaaS**



Providing responsive apps  
**SaaS**



Low cost, open source,  
built for software defined  
**Cloud Architecture**

# SSD Cloud Usage for AWS

Amazon EBS Volume Types


Volume Type	EBS General Purpose (SSD)	EBS Provisioned IOPS (SSD)	EBS Magnetic
Use Cases	Boot volumes	I/O intensive	Infrequent Data Access
	Small to Med DBs	Relational DBs	
	Dev and Test	NoSQL DBs	
Storage Media	SSD-backed	SSD-backed	Magnetic disk-backed
Max Volume Size	16 TB	16 TB	1 TB
Max IOPS/volume	10,000	20,000	40 - 200
Max IOPS Burst Performance	3,000 for volumes <= 1 TB	-	Hundreds
Max throughput/volume	160 MBps	320 MBps	40 - 90 MBps
Max IOPS/instance	48,000	48,000	48,000
Max throughput/instance	800 MBps	800 MBps	800 MBps
API Name	gp2	io1	standard
Price*	\$.10/GB - Month	\$ .125/GB - Month	\$.05/GB - Month
		\$.065/provisioned IOPS	
			\$.05/million I/O

## SSD Tier in Amazon Elastic Block Storage

SSD optimized instances (I2) are popular amongst respondents as a whole

## TYPES OF EC2 INSTANCES

■ Less than 1000 (N=95) ■ 1000-9999 (N=123) ■ 10000+ (N=48)



N=266

Source: Amazon AWS Infrastructure as a Service (IaaS) Survey, IDC, November, 2014

# SSD Cloud Usage for AWS

## Amazon EBS Volume Types

Volume Type EBS General Purpose (SSD) EBS Provisioned IOPS (SSD) EBS Magnetic

### EBS General Purpose SSD

Boot Volumes  
Small to Med DBs  
Dev and Test

### EBS Provisioned IOPS SSD

I/O Intensive  
Relational DBs  
noSQL DBs


Max throughput/instance	800 MBps	800 MBps	800 MBps
API Name	gp2	io1	standard
Price*	\$.10/GB - Month	\$.125/GB - Month \$.065/provisioned IOPS	\$.05/GB - Month \$.05/million I/O

## SSD Tier in Amazon Elastic Block Storage

SSD optimized instances (I2) are popular amongst respondents as a whole

## TYPES OF EC2 INSTANCES

■ Less than 1000 (N=95) ■ 1000-9999 (N=123) ■ 10000+ (N=48)



N=266

Source: Amazon AWS Infrastructure as a Service (IaaS) Survey, IDC, November, 2014



# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# REAL WORLD DEPLOYMENT

## YAHOO! ACCELERATING CEPH WITH SSD CACHING

RUIPING SUN  
PRINCIPAL ARCHITECT, YAHOO!

# CEPH Challenges:

High latency and low throughput as hot storage solution

## Erasure Code

Algorithm 8+3

- 1M photo: become 11 x 128K files → more small file IO
- Higher utilization: 63% (VS. 3 replications: 25%)
  - Example: 8T disk, EC usable space: 5T (3 Replications: 2T)

## CEPH write: twice

## Huge number of small files

- Number of files: 64 – 128 millions files on 8T file system
- One file access: become 3-4 disks accesses

## IO Performance:


Latency formula:

$$L = 1000 * 15 \text{ MS} / (1000 - n * 15 \text{ MS})$$

- n : number of IO request in 1 sec
- 15 MS: average disk latency

Tail latency:

- All 8 EC chunks(files) must be received: weakest link



# Yahoo CEPH Deployment Objectives:


Provide Reliable, Operable and Inexpensive Object Storage

## 2014 December

Strategy: Pickup low hanging fruit

Achieved: Yahoo Flickr 10PB

- Old photo
- Cold Storage: Access rate is low
- Infrastructure: 6T Disk, Used servers




## 2015 August

Strategy: Pickup low hanging fruit and Eye on the higher hanging fruit

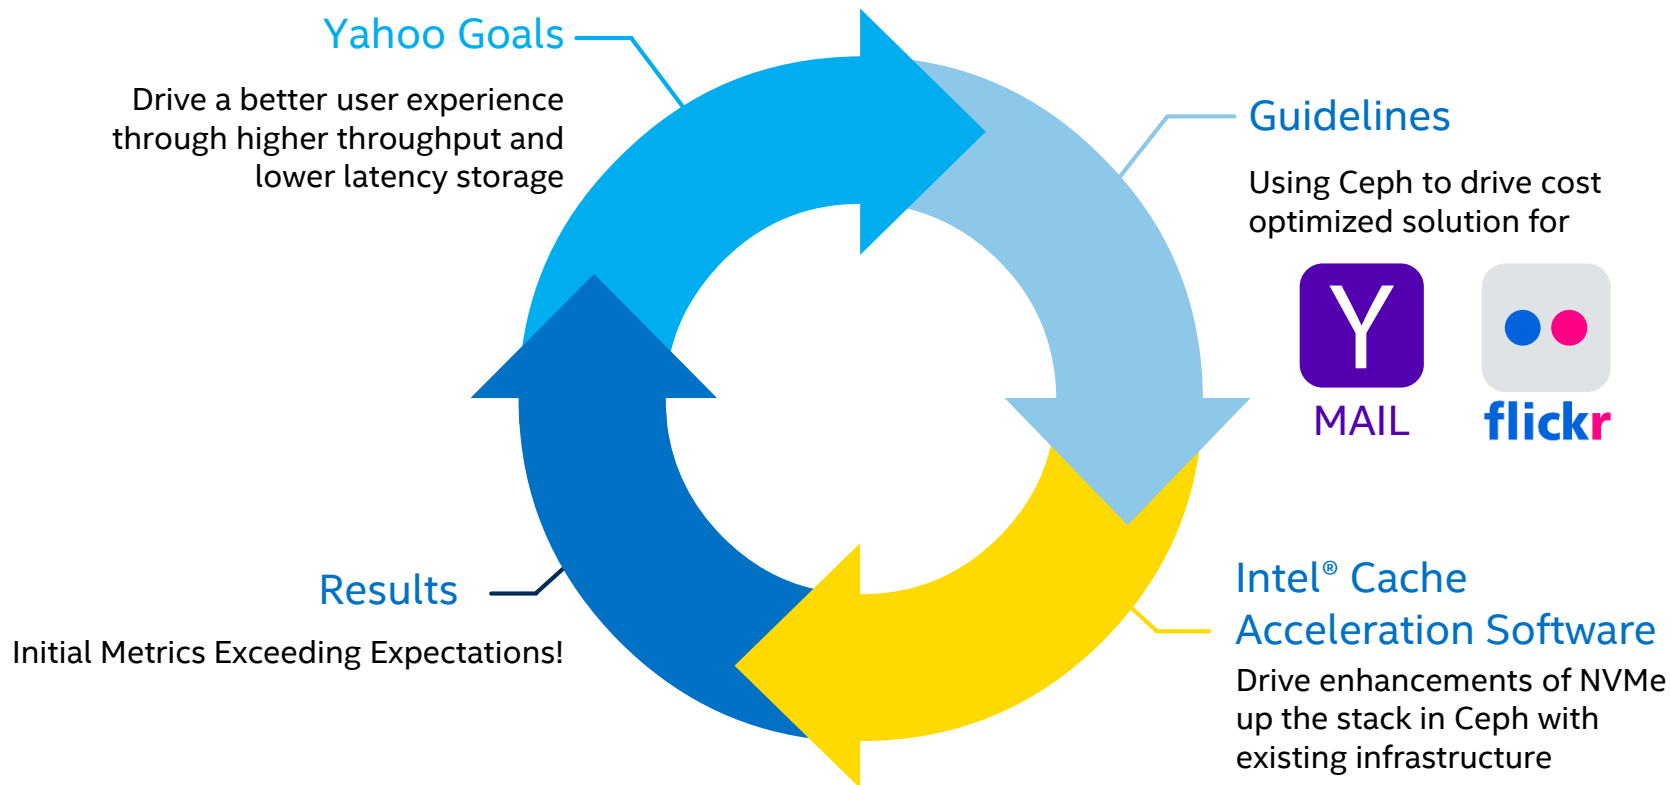
Archived: Yahoo Flickr 10PB

- New photo: 400 new photos/sec(upload)
- Warm Storage: Access rate is moderate
- Infrastructure: 6T Disk, Used servers




## Ceph Performance Improvements with NVMe™ Caching

## Yahoo! &amp; Intel – Better Together



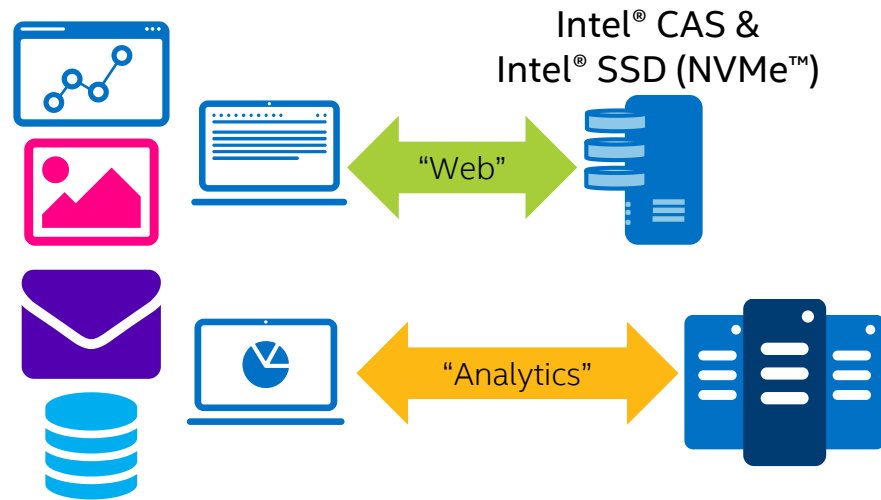


# BEFORE



# AFTER

YAHOO!



*Intel® CAS 3.0 featuring differentiated storage services hinting technology*

# Yahoo Evaluating Caching based Solution

Intel® CAS with hint feature, is a good potential solution

## Only Caching Metadata

- File access is pure random
- Attempting to cache file is too expensive

## Objectives (Improvement Goal)


- Ceph GET latency **30-70%**
- Ceph PUT latency **20-50%**
- Ceph GET throughput **200%**
- Ceph PUT throughput **100%**

## 2016

Strategy : Pickup both low and higher hanging fruit

Infrastructure:

- Cold and Warm storage: 8T Disk, High density server.
- Hot storage: 8T Disk, High density server, SSD caching



# Where Do We Go From Here?



Visit demo & animation in  
booth #571 for more  
details

- Integration of Intel CAS 3.0: Production Software
- Plan to expand into other Yahoo portfolio products in 2016
- Deliver contributions to open source Ceph
- Continued team effort




**Insights Driving Business Optimizations  
Resulting in a Great Customer Experience**

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# Intel Continues to Drive Technology

## Accelerating Solid State Storage in Computing Platforms



32 Tiers

### CAPACITY

Enables high-density flash devices

### COST

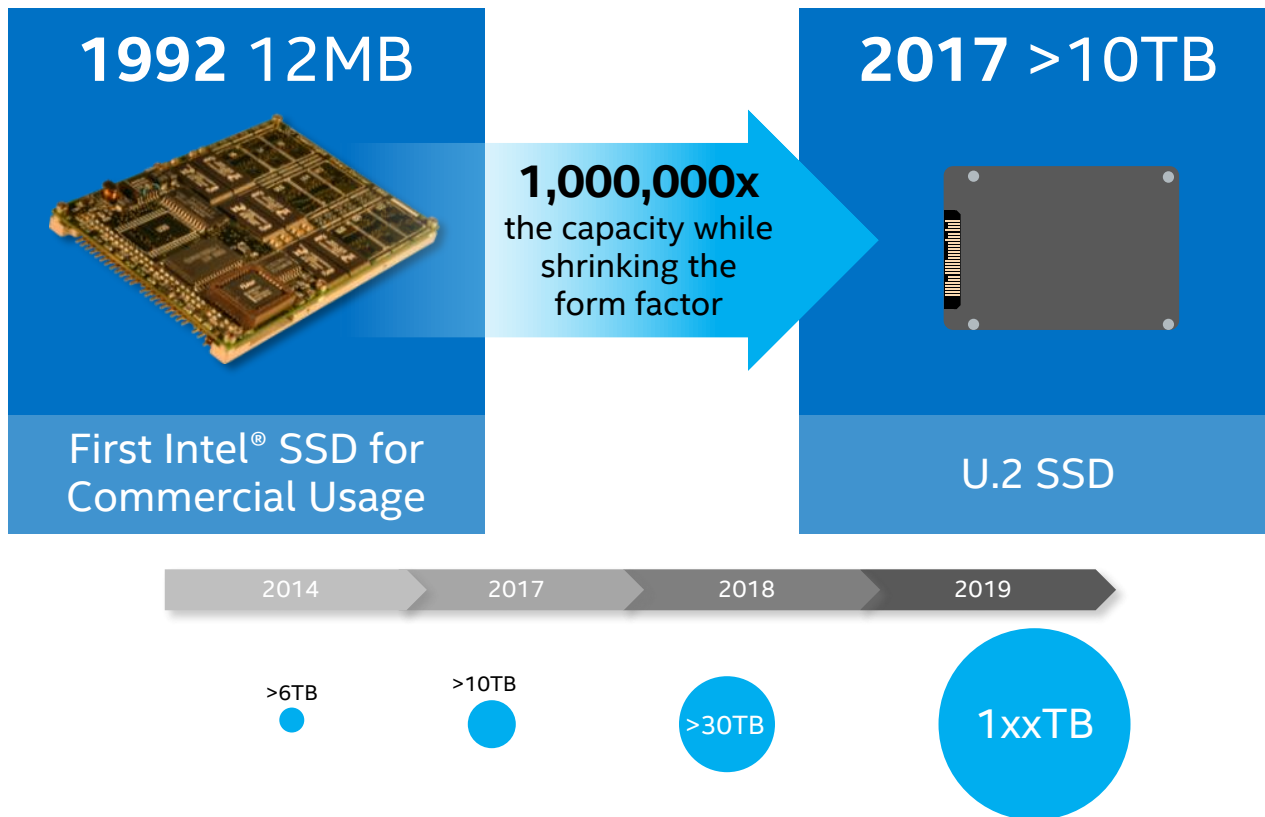
Achieves lower cost per gigabyte than 2D NAND at maturity

### CONFIDENCE

3D architecture increases performance and endurance



# Moore's Law Continues to Disrupt the Computing Industry



# THE BREAKTHROUGH

## A NEW CLASS OF NON-VOLATILE MEMORY



1000X

FASTER

THAN NAND

1000X

ENDURANCE

OF NAND

10X

DENSER


THAN DRAM

## NAND-LIKE DENSITIES AND DRAM-LIKE SPEEDS

# NAND Flash vs 3D XPoint™ Technology


## Economics of Storage Media

### 3D MLC and TLC NAND



Enable higher capacity SSDs at lower price

### 3D XPoint™ Technology

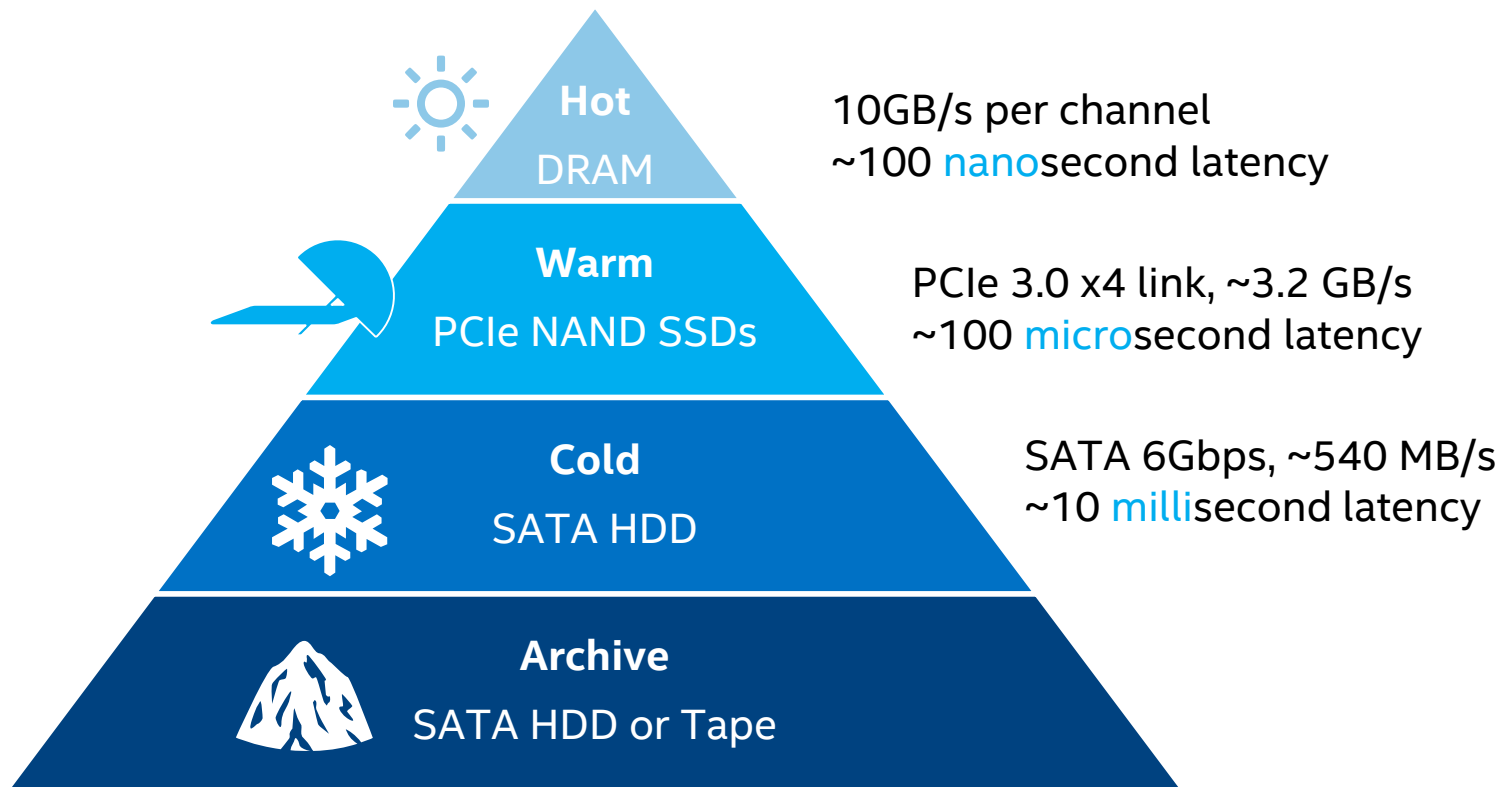


Higher performance, opening up new use cases for SSD

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A


# Storage and Memory Hierarchy Today



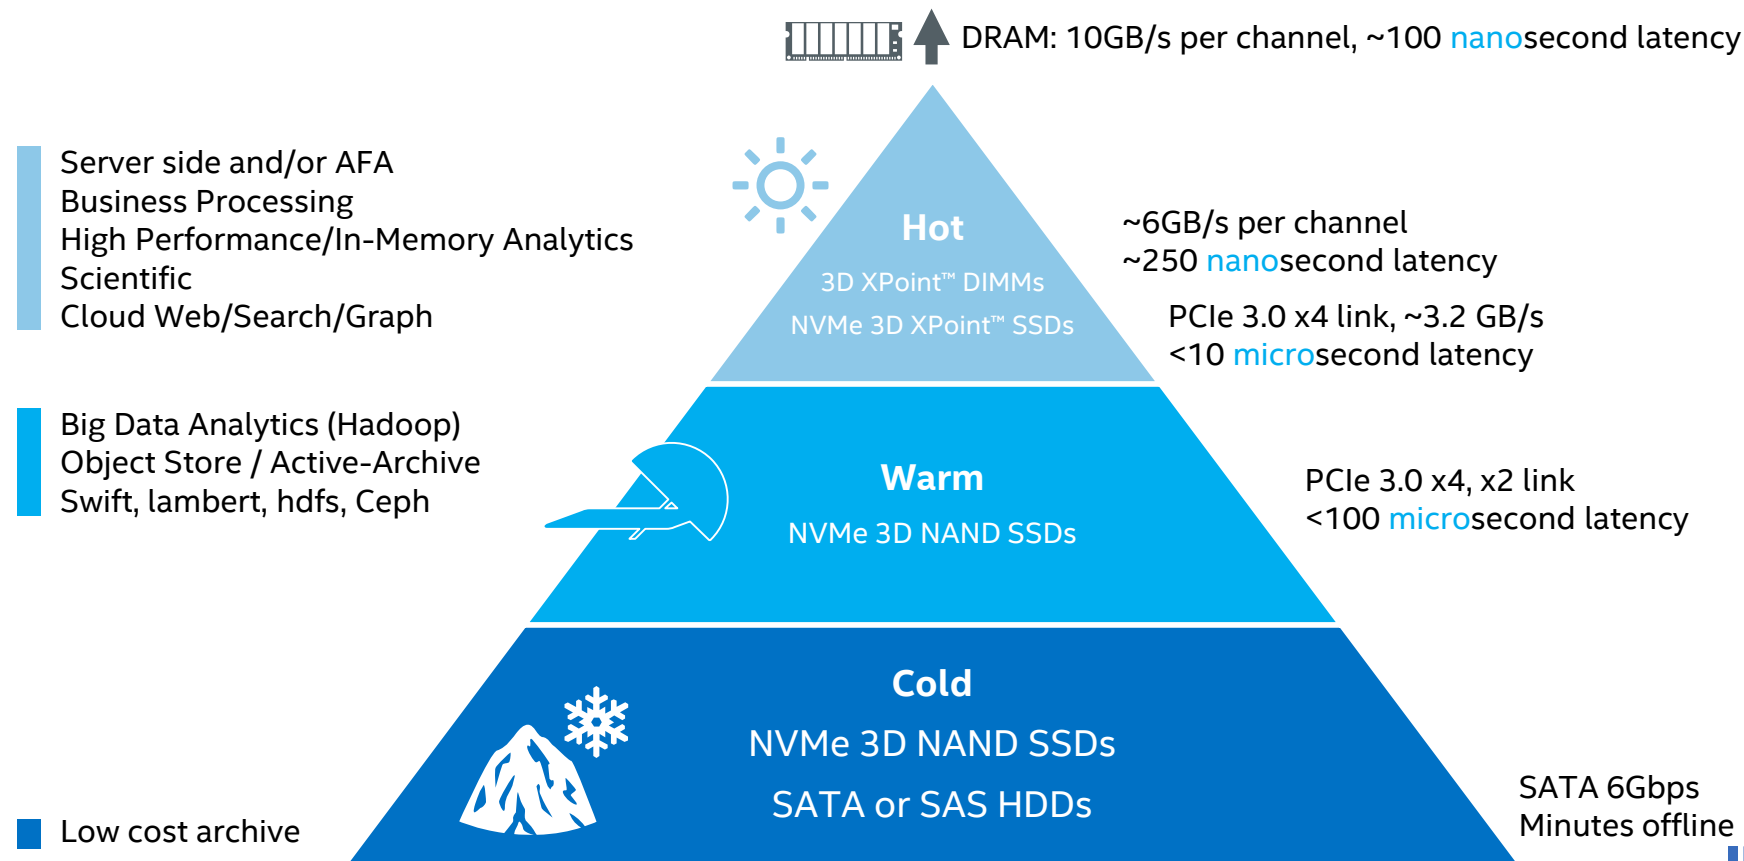


# IOPS/TB: HDD v. SSD

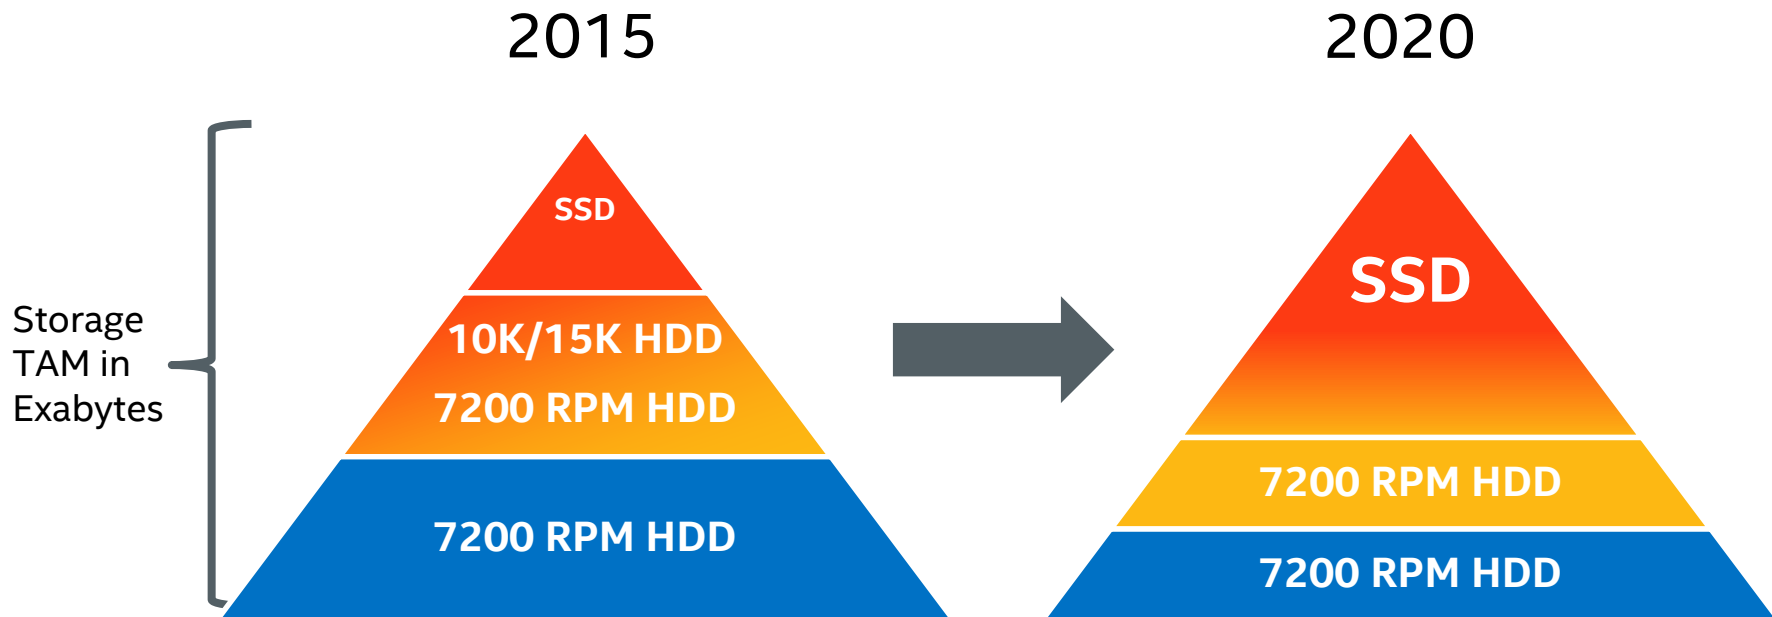
IOPS / TB: Queue Depth 1 4k Random Read



# Storage Hierarchy Tomorrow



# SSDs are Moving Beyond Hot Storage



TCO includes:  
performance, failure rate, power, data  
center efficiency, rack space, capx, opx

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVM Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# New Faster SSDs Emerge, Where Will They be Used?




## Extend DRAM

In memory database

Key value store, memcache

goal is to lower TCO and execute larger datasets



## Faster SSD

Real time analytics – noSQL databases,

Fraud detection, ad bidding, real time decision making, trading

Cloud hosting – amazing QoS for better application SLAs

Ultra High Definition all professional video production

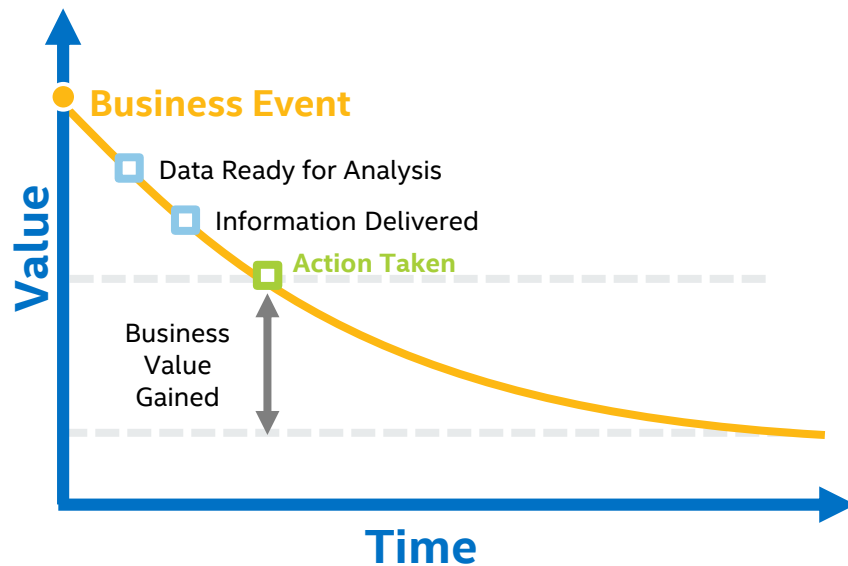
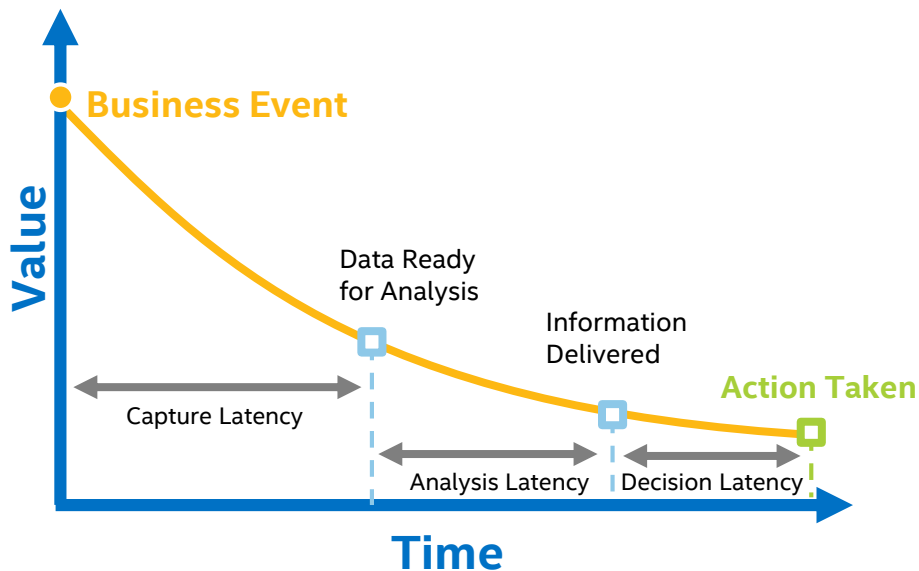


## Caching


Cloud hosting, all flash array, SAN

Write buffer for RAID array

# Lower Latency Opens Up Business Value

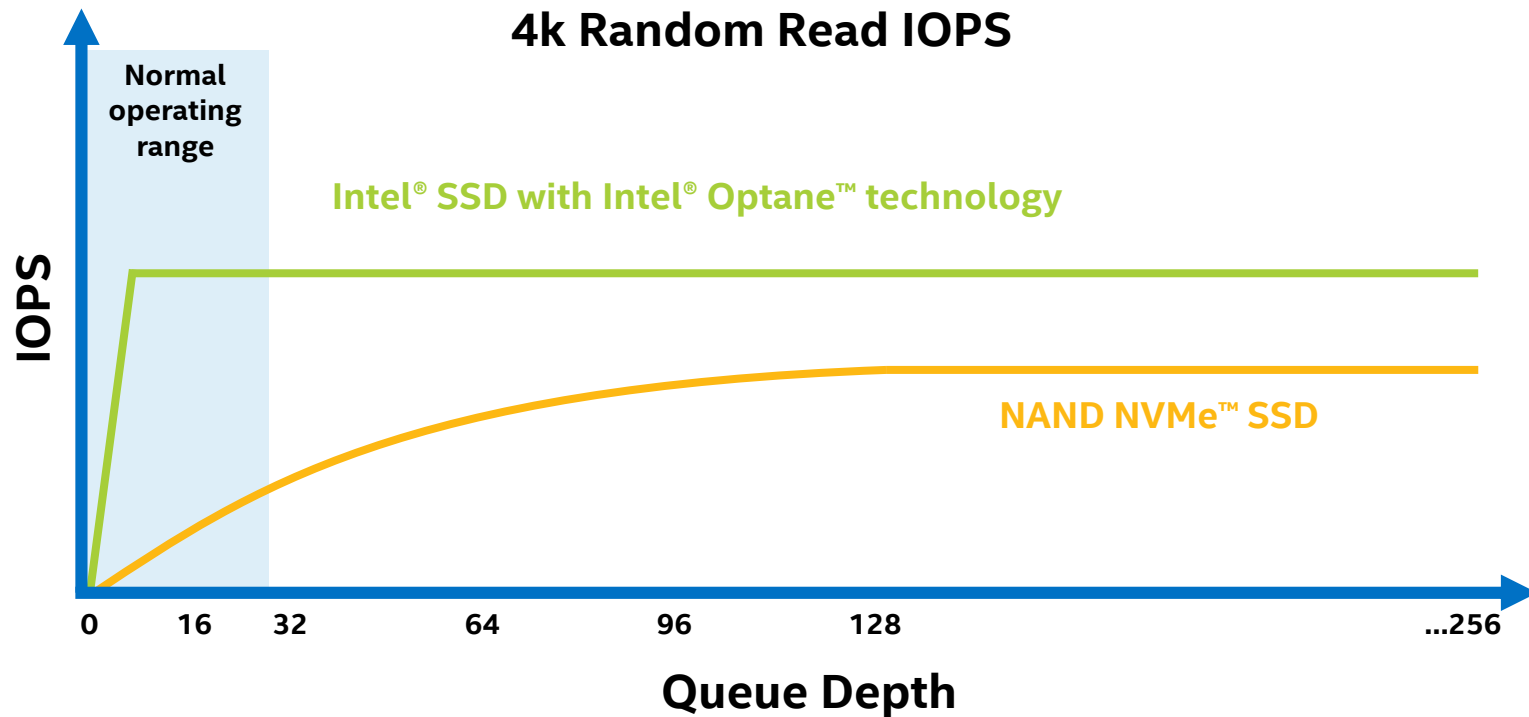


# Opening Up New Application Space






# NAND vs 3D XPoint™ Technology: SSD IOPS vs Queue Depth



# SSD Write IOPs are Approaching Free?!



Go ahead, write to disk!

# Agenda

- Data growth driving fundamental changes to storage technologies
- The future of SSDs is here today – NVMe Express™
- How are SSDs being used today?
- Special Guest Speaker! SSDs in Cloud Accelerating Ceph Object Store
- 3D NAND and 3D XPoint™ technology to build new types of SSDs
- Transformation of the storage hierarchy
- New use case for next generation SSDs
- Summary and Q&A

# Summary and Next Steps

- Understand the storage market and memory tiers to adopt the right SSD(s) for your workload
- NVMe™ was build for high-performance SSDs with the future in mind - and ready today!
- 3D NAND is the building block for high capacity, low cost SSDs while 3D XPoint™ technology delivers high performance and low latency
- Faster SSDs open up completely new use cases which require software applications that can take full advantage – start optimizing today!

# Additional Sources of Information

- A PDF of this presentation is available from our Technical Session Catalog: [www.intel.com/idfsessionsSF](http://www.intel.com/idfsessionsSF). This URL is also printed on the top of Session Agenda Pages in the Pocket Guide.
- Demos in the showcase – [Intel® SSD Pavilion](#)
- Additional info in the ABC community – [NVM Express™ Community](#)
- More web based info:
- [www.intel.com/ssd](http://www.intel.com/ssd)
- [www.intel.com/nvm](http://www.intel.com/nvm)

# Other Technical Sessions

Session ID	Title	Day	Time	Room
<b>SSDS003</b>	What You Need to Know to Win the Storage Transition – Preparing for NVM Express™ in the Data Center	Tues	4:00	2008
<b>SPCS006</b>	Technology Insight: Intel Non-Volatile Memory Inside. The Speed of Possibility Outside	Tues	5:15	3016
<b>SSDL001</b>	Hands-on Lab: How to Unleash Your Storage Performance by Using NVM Express* based PCI Express* Solid-State Drives	Wed	1:15; 4:00	2010
<b>SSDC001</b>	Tech Chat: Benchmarking Data Center Solid-State Drives – Insights Into Industry-Leading NVM Express* SSD Performance Metrics	Wed and Thurs	10:30 Wed 9:30 Thurs	Tech Chat Station 1
<b>SSDC002</b>	Tech Chat: Insights into Intel® Solid-State Drives Data Retention and Endurance	Wed and Thurs	10:30 Wed 9:30 Thurs	Tech Chat Station 2
<b>SSDC003</b>	Tech Chat: NVM Express* Features for High Availability and Storage Eco-System	Wed and Thurs	10:30 Wed 9:30 Thurs	Tech Chat Station 3
<b>SSDS004</b>	The Future of Storage Security	Thurs	1:00	2006
<b>SSDS005</b>	New Software Capabilities and Experiences Through Innovation in Storage Architecture	Thurs	2:15	2006

# 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.

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>.

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.

© 2015 Intel Corporation.

# Risk Factors

The above statements and any others in this document that refer to plans and expectations for the second quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "plans," "believes," "seeks," "estimates," "may," "will," "should" and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel's actual results, and variances from Intel's current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be important factors that could cause actual results to differ materially from the company's expectations. Demand for Intel's products is highly variable and could differ from expectations due to factors including changes in business and economic conditions; consumer confidence or income levels; the introduction, availability and market acceptance of Intel's products, products used together with Intel products and competitors' products; competitive and pricing pressures, including actions taken by competitors; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Intel's gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; segment product mix; the timing and execution of the manufacturing ramp and associated costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; and product manufacturing quality/yields. Variations in gross margin may also be caused by the timing of Intel product introductions and related expenses, including marketing expenses, and Intel's ability to respond quickly to technological developments and to introduce new products or incorporate new features into existing products, which may result in restructuring and asset impairment charges. Intel's results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Results may also be affected by the formal or informal imposition by countries of new or revised export and/or import and doing-business regulations, which could be changed without prior notice. Intel operates in highly competitive industries and its operations have high costs that are either fixed or difficult to reduce in the short term. The amount, timing and execution of Intel's stock repurchase program could be affected by changes in Intel's priorities for the use of cash, such as operational spending, capital spending, acquisitions, and as a result of changes to Intel's cash flows or changes in tax laws. Product defects or errata (deviations from published specifications) may adversely impact our expenses, revenues and reputation. Intel's results could be affected by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues. An unfavorable ruling could include monetary damages or an injunction prohibiting Intel from manufacturing or selling one or more products, precluding particular business practices, impacting Intel's ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. Intel's results may be affected by the timing of closing of acquisitions, divestitures and other significant transactions. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Form 10-Q, Form 10-K and earnings release.