IDF2012
INTEL DEVELOPER FORUM

nvm EXPRESS

NVM Express and the PCI Express* SSD Revolution

Danny Cobb, CTO Flash Memory Business Unit, EMC
Amber Huffman, Sr. Principal Engineer, Intel

SSDS003

Sponsors of Tomorrow. intel

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

The PDF for this Session presentation is available from our Technical Session Catalog at the end of the day at:

intel.com/go/idfsessions

URL is on top of Session Agenda Pages in Pocket Guide

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

NVM Express (NVMe) Overview

- NVM Express is a scalable host controller interface designed for Enterprise and client systems that use PCI Express* SSDs
- NVMe was developed by industry consortium of 80+ members and is directed by a 13-company Promoter Group

- NVMe 1.0 was published March 1, 2011
- Product introductions later this year, first in Enterprise

Technical Basics

- The focus of the effort is efficiency, scalability and performance
  - All parameters for 4KB command in single 64B DMA fetch
  - Supports deep queues (64K commands per Q, up to 64K queues)
  - Supports MSI-X and interrupt steering
  - Streamlined command set optimized for NVM (6 I/O commands)
  - Enterprise: Support for end-to-end data protection (i.e., DIF/DIX)
  - NVM technology agnostic

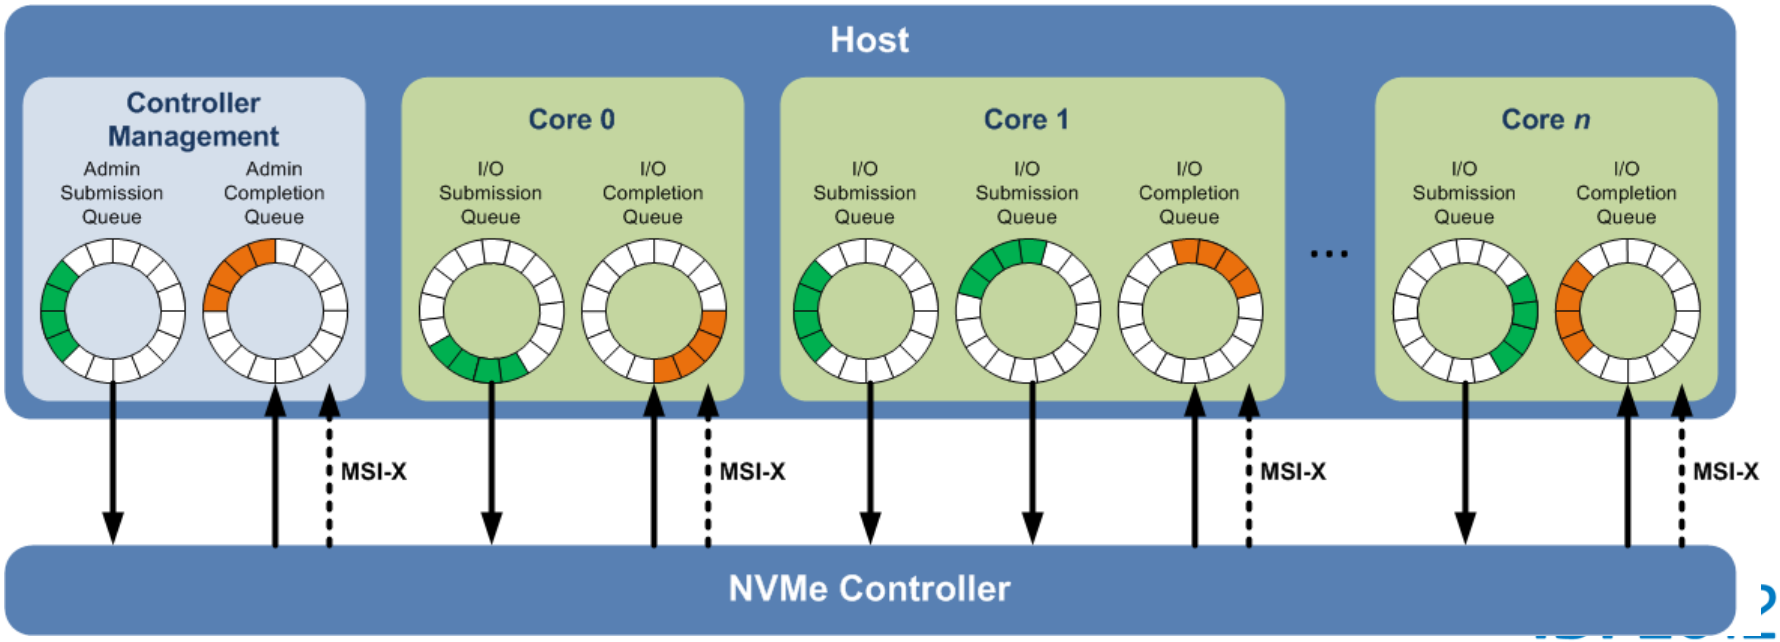

NVMe = NVM Express

NVMe Command Execution

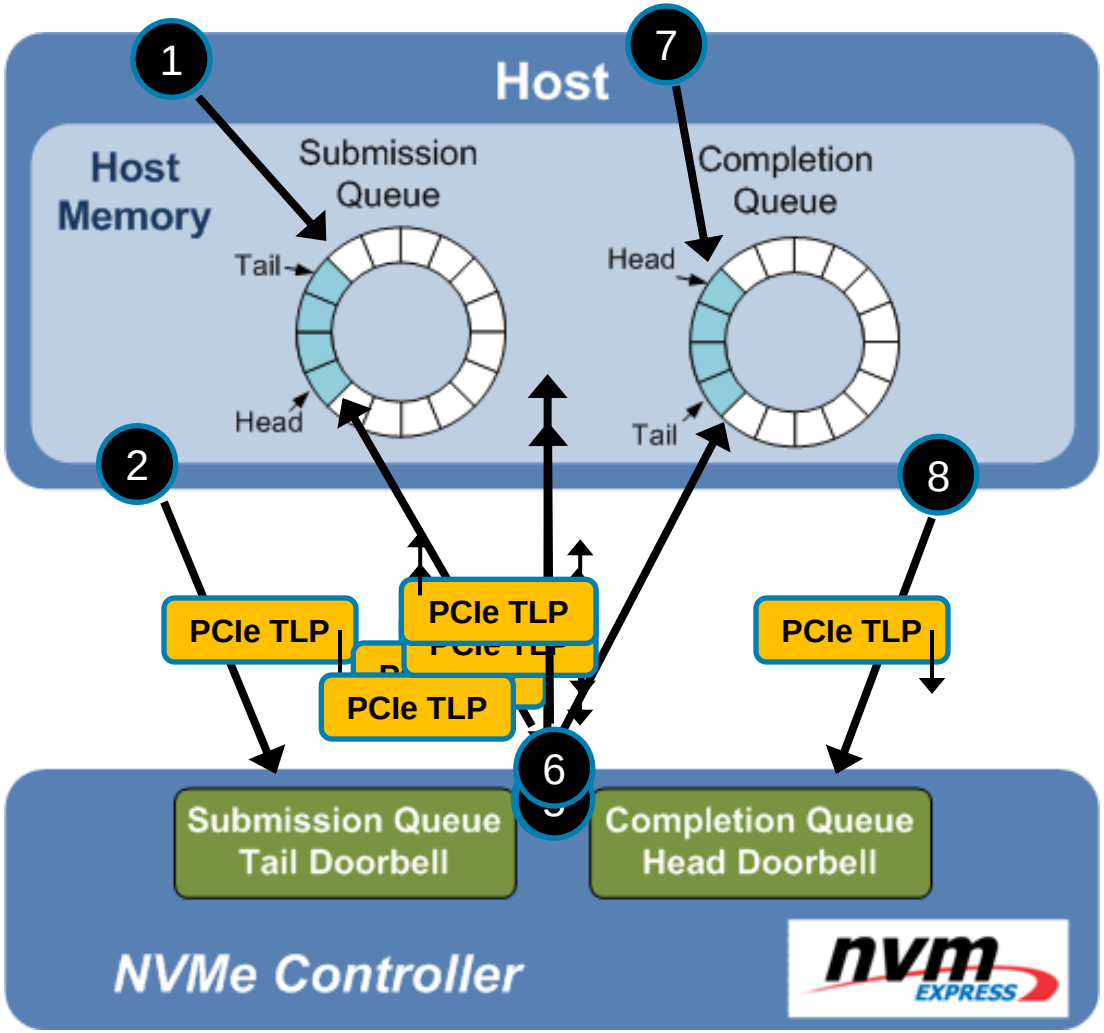

1) Queue Command(s)
2) Ring Doorbell *(New Tail)*
3) Fetch Command(s)
4) Process Command
5) Queue Completion(s)
6) Generate Interrupt
7) Process Completion
8) Ring Doorbell *(New Head)*

NVMe = NVM Express
PCIe = PCI Express*

1.0 Command Set Overview

Management Commands for Queues & Transport

Admin Command	Description
Create I/O Submission Queue	Queue Management
Create I/O Completion Queue	
Delete I/O Submission Queue	
Delete I/O Completion Queue	
Abort	
Asynchronous Event Request	Status & Event Reporting
Get Log Page	
Identify	Configuration
Set Features	
Get Features	
(Optional) Firmware Activate	Firmware Management
(Optional) Firmware Image Download	
(Optional) Security Send	Security
(Optional) Security Receive	
(Optional) Format NVM	Namespace Management

I/O Commands for SSD Functionality

NVM Command	Description
Flush	Data Ordering
Read	Data Transfer, Including end-to-end data protection & security
Write	
(Optional) Write Uncorrectable	
(Optional) Compare	
(Optional) Dataset Management	Data Usage Hints

13 Required Commands Total (10 Admin, 3 I/O)

NVM Express: Architecting for Performance & Power Efficiency

	AHCI	nvm EXPRESS
Uncacheable Register Reads Each consumes 2000 CPU cycles	4 per command 8000 cycles, ~ 2.5 µs	**0** per command
MSI-X and Interrupt Steering Ensures one core not IOPs bottleneck	No	Yes
Parallelism & Multiple Threads Ensures one core not IOPs bottleneck	Requires synchronization lock to issue command	No locking, doorbell register per Queue
Maximum Queue Depth Ensures one core not IOPs bottleneck	1 Queue 32 Commands per Q	64K Queues 64K Commands per Q
Efficiency for 4KB Commands 4KB critical in Client and Enterprise	Command parameters require two serialized host DRAM fetches	Command parameters in one 64B fetch

NVM Express is optimized for SSDs, replacing the decade old AHCI standard designed for the hard drive era

Proof Point: NVMe Latency

- NVMe reduces latency overhead by **more than 50%**
  - SCSI/SAS: 6.0 µs 19,500 cycles
  - **NVMe: 2.8 µs 9,100 cycles**
- NVMe is designed to scale over the next decade
  - NVMe supports future NVM technology developments that will drive latency overhead below one microsecond
- Example of latency impact: Amazon* loses 1% of sales for every 100 ms it takes for the site to load

Chatham NVMe Prototype

Prototype Measured IOPS

Cores Used for 1M IOPs

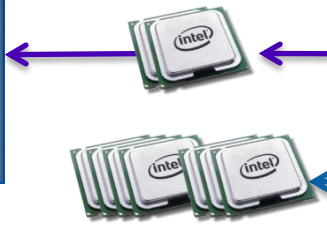

Linux* Storage Stack

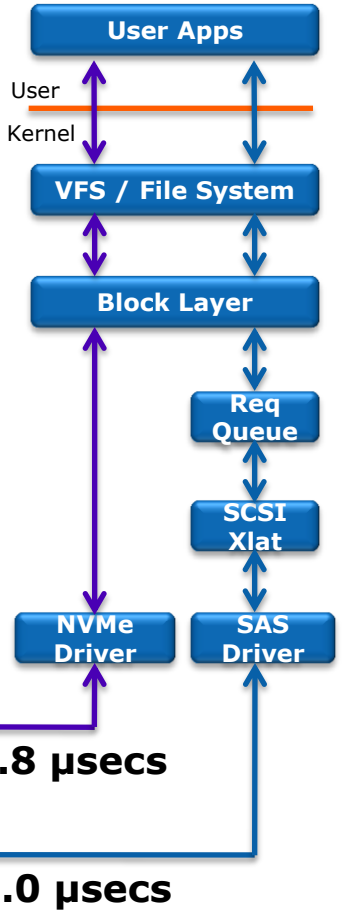

IDF2012 INTEL DEVELOPER FORUM

NVMe = NVM Express

Measurement taken on Intel® Core™ i5-2500K 3.3GHz 6MB L3 Cache Quad-Core Desktop Processor using Linux RedHat* EL6.0 2.6.32-71 Kernel.

Proof Point: NVMe Efficiency & Power

- NVMe prototype delivers lower clocks per I/O while at the same time delivering higher performance on the workloads
- Lower clocks per I/O is a proxy for efficiency and lower power – the CPU & system can go to a sleep state more quickly

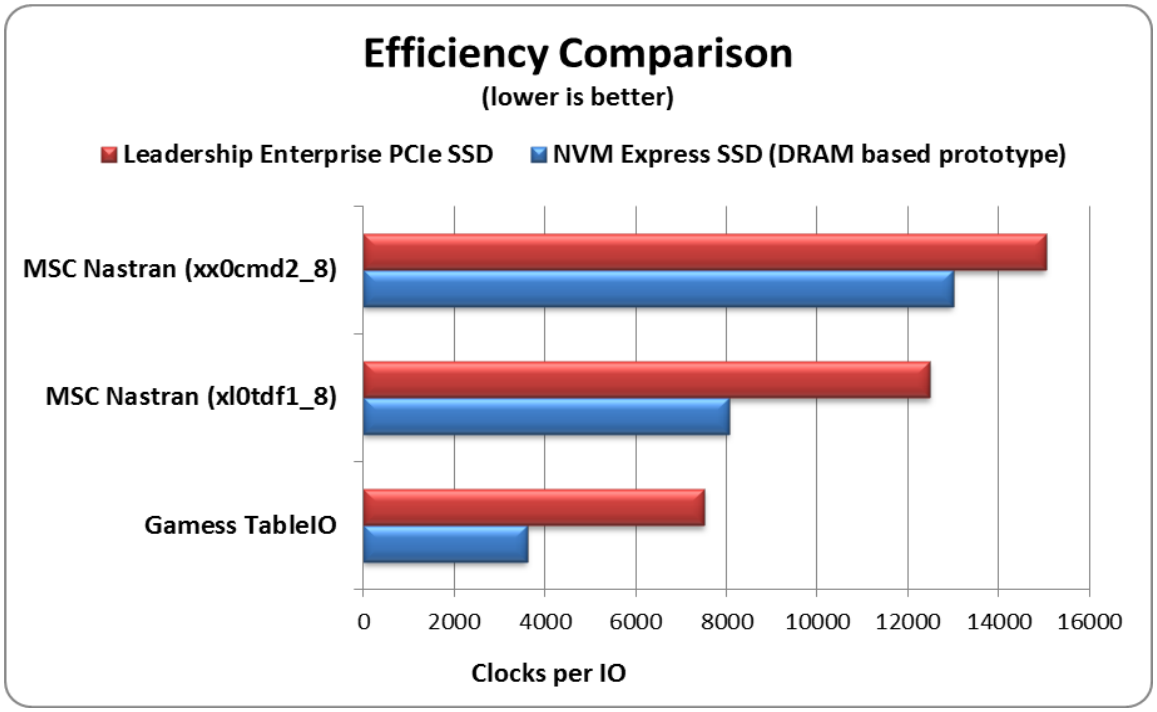

NVMe = NVM Express

Charts compare NVM Express (NVMe) and Leadership Enterprise PCIe *SSD. NVMe utilized DRAM to push protocol to limits. Leadership Enterprise PCIe SSD utilizes NAND making runtime comparisons inappropriate. Gamess TableIO workload is a computational test.

Proof Point: NVMe in a SAN

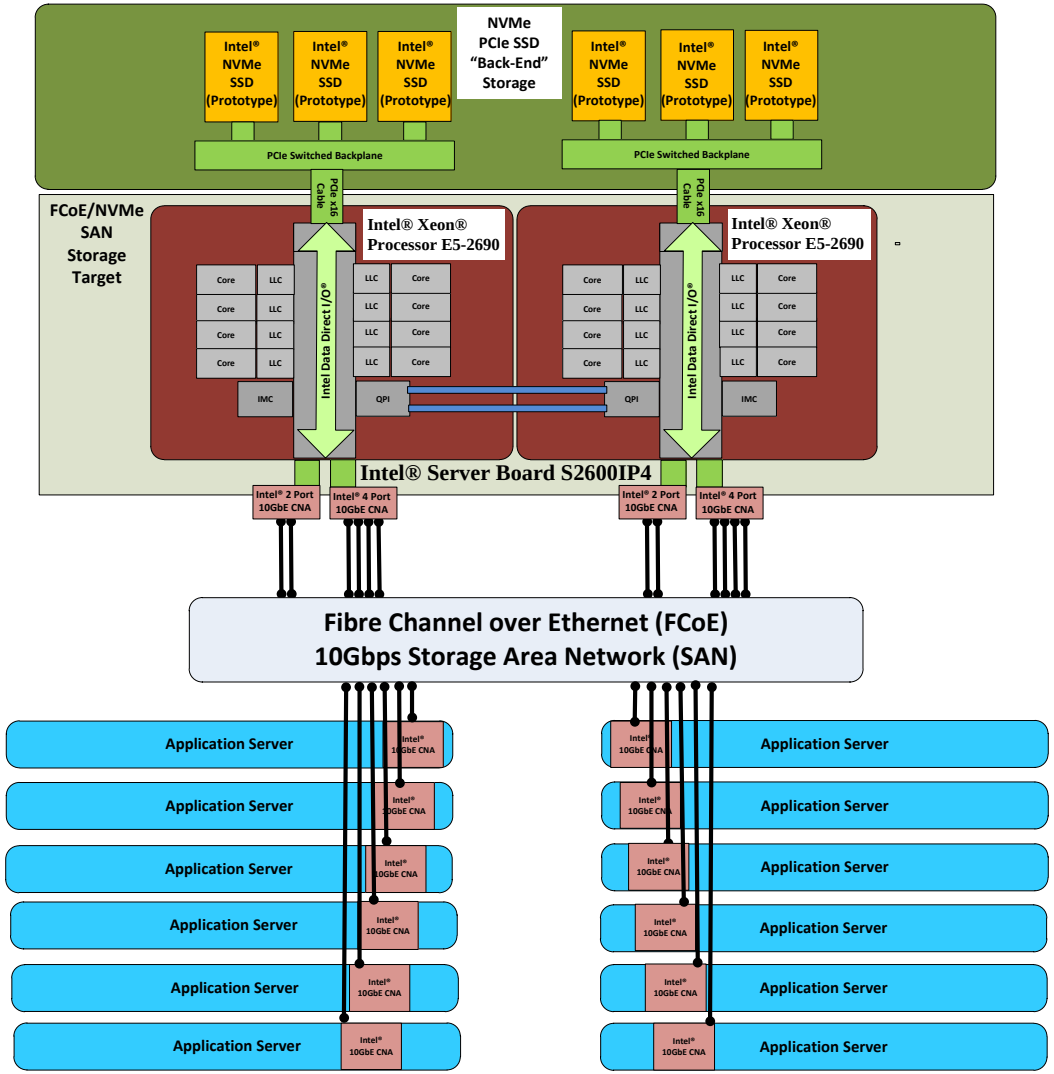

- Demo combines NVMe with existing ingredients to deliver > 3.1M 4K IOPs
- The performance of direct attached (DAS) NVMe SSDs married to an FCoE SAN
- Next generation SAN is possible today by use of highly efficient interfaces
- Check out the demo in the NVMe Community, available until **2pm**

SAN with NVMe: 3.1 Million 4K IOPs on 120Gbps FCoE

- Storage target configuration: Intel® S2600IP4 Server Board, Intel® Xeon® Processor E5-2690 2.9GHz, 8-16GB DDR3 1033 DIMMs, RH EL-6.2 – 3.3.0-RC1 kernel, TCM storage target, , 4 Intel® Ethernet Server Adapter X520 (10 Gbps CNA).
- Initiator configuration: 12 initiators: Intel® Xeon® Processor 5650 2.67GHz, RH EL-6.2 – 3.3.0-RC1 kernel.
- Test configuration: (per initiator) Linux fio V21.0.7, 4K Random Read, QD=8, Workers=16, 8 FCoE LUNs.

NVMe = NVM Express

NVM Express (NVMe) Community

18 Industry Booths

OEMs

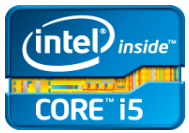

IT Value:
Standards-Based
Performance, Efficiency, Scalability, Flexibility

Controllers, Tools, Silicon IP, Software

SSD vendors

NVMe transitions from "Specs" to "Deployment"

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

Multi-Path I/O and Namespace Sharing

- An NVMe namespace may be accessed via multiple "paths"
  - SSD with multiple PCI Express* ports
  - SSD behind a PCIe switch to many hosts
- Two hosts accessing the same namespace must coordinate
- The NVMe Workgroup added capabilities in NVMe 1.1 to enable Enterprise multi-host usage models

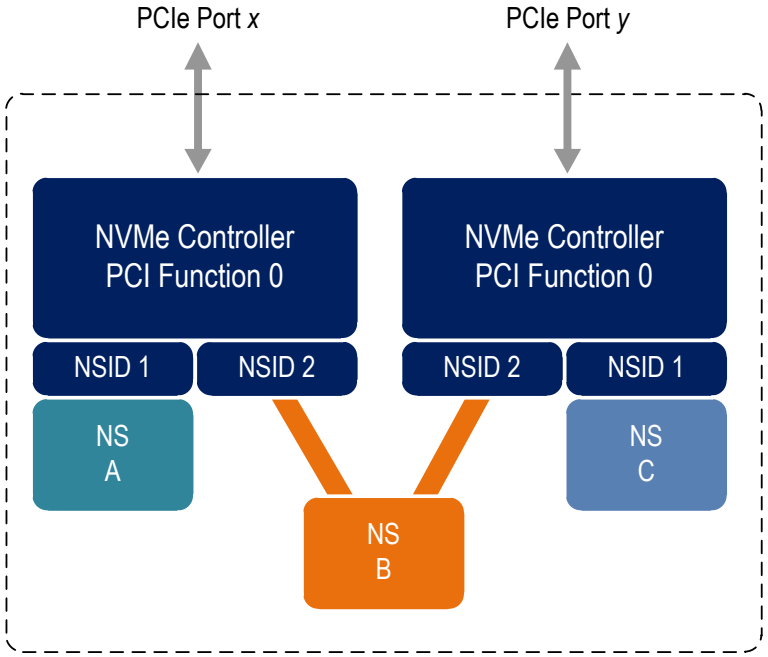

NVMe = NVM Express

Uniquely Identifying a Namespace

- How do Host A and Host B know that NS B is the same namespace?
- NVMe 1.1 adds unique identifiers for:
  - The NVMe Controller
  - Each Namespace within an NVM Subsystem
- These identifiers are guaranteed to be globally unique

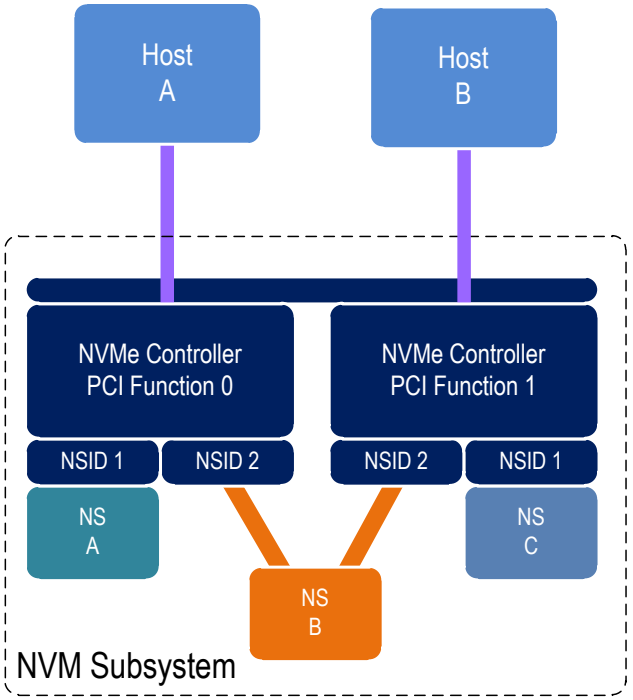

Unique NVMe Controller Identifier =
2B PCI Vendor ID + 20B Serial Number + 40B Model Number + 2B Controller ID

Unique Namespace Identifier =
64B Unique NVM Subsystem Identifier + 8B IEEE Extended Unique Identifier

NVMe = NVM Express

NVM Subsystem Reset

- Resets in NVMe 1.0 are controller based
- NVMe 1.1 adds a capability to reset the entire NVM Subsystem
  - E.g., new firmware needs to be applied to both controllers
- To perform an NVM Subsystem Reset, write the value “NVMe” to the register

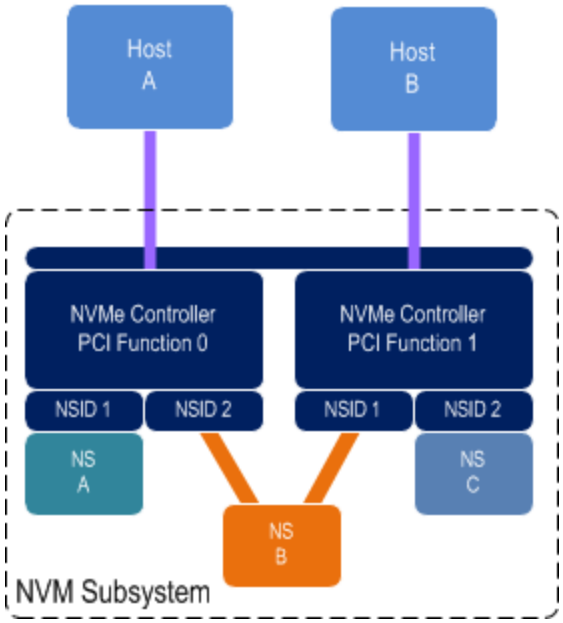

NVM Subsystem Register

Bit	Type	Reset	Description
31:00	RW	0h	**NVM Subsystem Reset Control (NSSRC):** A write of the value 4E564D65h ("NVMe") to this field initiates an NVM Subsystem Reset. A write of any other value has no functional effect on the operation of the NVM subsystem. This field shall return the value 0h when read.

NVMe = NVM Express

Reservations

- In some multi-host environments, like Windows* clusters, reservations are used
- NVMe 1.1 includes a simplified reservations mechanism that is compatible with implementations that use SCSI reservations
- What is a reservation? Enables two or more hosts to coordinate access to a shared namespace.
  - A reservation may allow Host A and Host B access, but disallow Host C

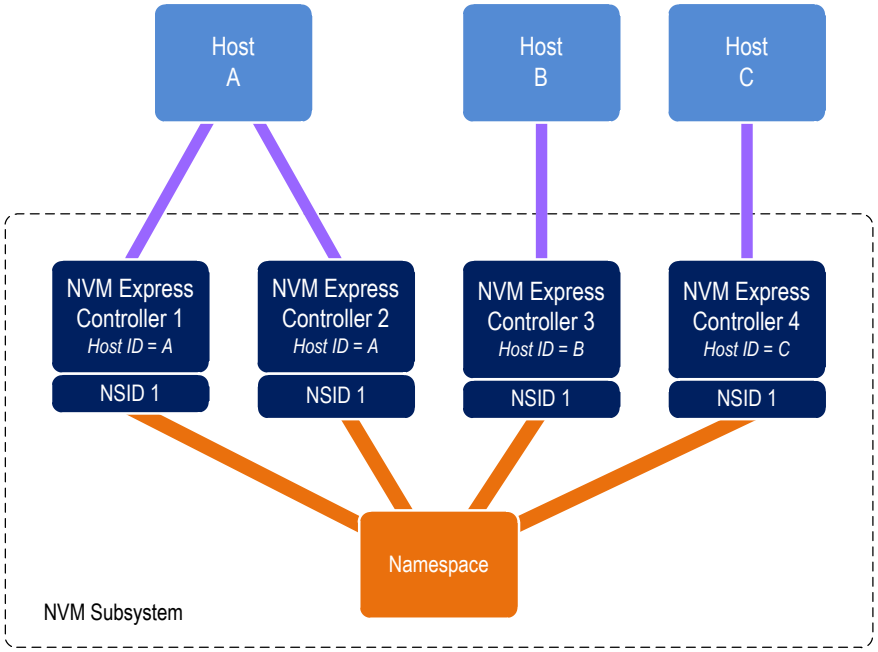

NVMe = NVM Express

NVMe Commands for Reservations

- NVMe has mapped the 12 SCSI reservation service actions to four optional commands:
  - Report Reservation, Register, Acquire Reservation, Release Reservation

Persistent Reserve In

Service Action	Service Action Description	NVMe 1.1 Command
Read Keys	Read current registered reservation keys	Reservation Report
Read Reservation	Read current persistent reservation	Reservation Report
Report Capabilities	Report persistent reservation capabilities	Identify Namespace Data Structure
Read Full Status	Read detailed status	Reservation Report

Persistent Reserve Out

Service Action	Service Action Description	NVMe 1.1 Command
Register	Register or unregister a reservation key	Reservation Register
Reserve	Create a persistent reservation	Reservation Acquire
Release	Release a persistent reservation	Reservation Release
Clear	Clears all reservation keys and all persistent reservations	Reservation Release
Preempt	Preempt persistent reservations and/or remove registrants	Reservation Acquire
Preempt and Abort	Same as preempt but also abort all commands	Reservation Acquire
Register and Ignore Key	Register a reservation key and unregister a reservation key	Reservation Register
Register and Move	Register a reservation key/reservation for another I_T nexus	n/a

NVMe = NVM Express

Reservations in Action

- Example: Host A and B have read/write access and host C has read-only access to the shared namespace

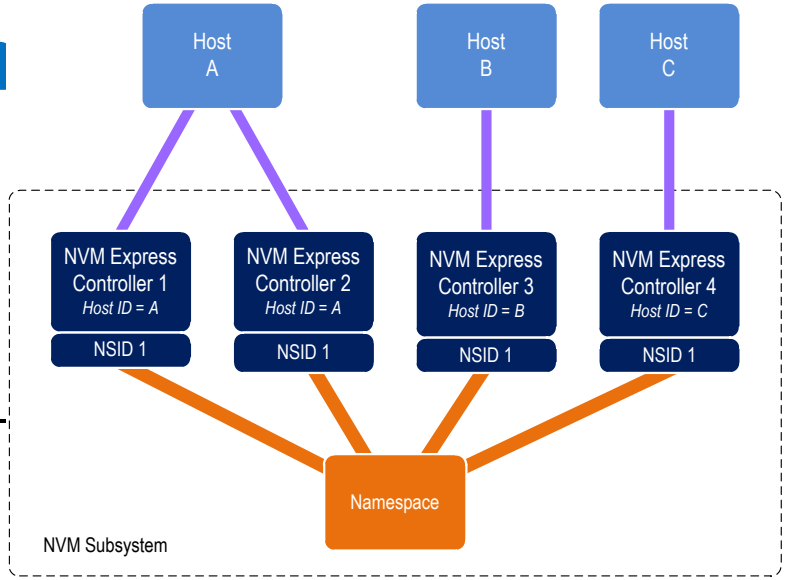


```
HostA-SetFeatures (HostID_A) -> OK
HostB-SetFeatures (HostID_B) -> OK
HostC-SetFeatures (HostID_C) -> OK
…
HostA-Register(NSID,Key_A) -> OK
HostB-Register(NSID,Key_B) -> OK
HostA-AcquireReservation(NSID, Reservation, WriteExclusiveRegistrantsOnly,Key_A) -> OK
HostC-AcquireReservation(NSID, Reservation, WriteExclusiveRegistrantsOnly,Key_C) ->
          Error – Reservation Conflict
…
HostA-Write(NSID)  -> OK
…
HostB-Read(NSID)  -> OK
…
HostB-Write(NSID)  -> OK
…
HostC->Read(NSID) -> OK
HostC->Write(NSID) -> Error – Reservation Conflict
…
HostA-ReleaseReservation(NSID,Key1) -> OK
HostC-Write(NSID) -> OK
…
```

Windows* 8 Drives Client Power Lower

- For Windows* 8 Connected Standby, Microsoft* has specified performance & power requirements
- For idle power, storage must be <= 5 mW in Connected Standby
- NVMe is meeting this challenge using autonomous power state transitions, added in NVMe 1.1

System.Fundamentals. StorageAndBoot.BootPerformance

msdn.microsoft.com/library/windows/hardware/hh748188

Feature	Specification
Power	
Max Idle Power	<= 5 mW
Random Performance	
4 KB Write IOPs (measured overa 1GB area)	>= 200
4 KB Write IOPs (measured over a 10GB area)	>= 50
64 KB Write IOPs (measured overa 1GB area)	>= 25
4 KB Read IOPs (measured overa 10GB area)	>= 2000
4 KB 2:1 read/write mix IOPs (measured overa 1GB area)	>= 500
4 KB 2:1 read/write mix IOPs (measured overa 10GB area)	>= 140
Sequential Performance	
Write speed (64 KB I/Os) (measured overa 10GB area)	>= 40 MB/s
Write speed (1MB I/Os) (measured overa 10GB area)	>= 40 MB/s
Read speed (64 KB I/Os) (measured overa 10GB area)	>= 60 MB/s
Device I/O Latency	
Max Latency	< 500 milliseconds

Additional I/O Latency requirement:

- Maximum of **20 seconds** sum-total of user-perceivable I/O latencies over any 1 hour period of a user-representative workload, where a user-perceivable I/O is defined as having a latency of at least 100 milliseconds.

NVMe = NVM Express

Achieving Low Idle Power

- NVMe 1.1 added the Autonomous Power State Transition feature
- Without software intervention, the NVMe controller transitions to a lower power state after a certain idle period
  - Idle period prior to transition programmed by software

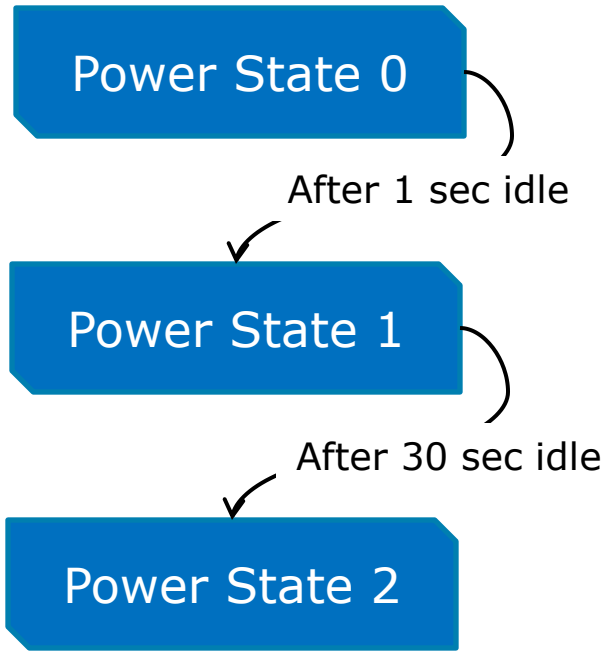

Example Power States

Power State	Opera-tional?	Max Power	Entrance Latency	Exit Latency
0	Yes	4 W	10 µs	10 µs
1	No	25 mW	10 ms	5 ms
2	No	3 mW	15 ms	30 ms

NVMe is delivering the features needed for leadership Enterprise and Client solutions

NVMe = NVM Express

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

Reference Drivers for Key OSs

- Linux*
  - Already accepted into the mainline kernel on kernel.org
  - Open source with GPL license
  - Refer to http://git.infradead.org/users/willy/linux-nvme.git
- Windows*
  - Baseline developed in collaboration by IDT*, Intel, and LSI*
  - Open source with BSD license
  - Maintenance is collaboration by NVMe WG and Open Fabrics Alliance
  - Refer to https://www.openfabrics.org/resources/developer-tools/nvme-windows-development.html
- VMware*
  - Initial driver developed by Intel
  - Based on VMware advice, "vmk linux" driver based on Linux version
  - NVMe WG will collaborate with VMware on delivery/maintenance

NVMe = NVM Express

Reference Drivers for Key OSs (cont.)

- Solaris*
  - There is a working driver prototype
  - Planned features include:
    - Fully implement and conform to 1.0c spec
    - Efficient block interfaces bypassing complex SCSI code path
    - NUMA optimized queue/interrupt allocation
    - Reliable with error detect and recovery fitting into Solaris* FMA
    - Build ZFS with multiple sector sizes (512B, 1KB, 2KB, 4KB) on namespaces
    - Fit into all Solaris disk utilities and fwflash(1M) for firmware
    - Boot & install on SPARC and X86
    - Surprise removal support
  - Plan to validate against Oracle* SSD partners
  - Plan to integration into S12 and a future S11 Update Release
- UEFI
  - The driver is under development
  - Plan to open source the driver in Q1 '13, including bug/patch process
  - Beta quality in Q1'13, production quality Q2'13

OFA NVMe Driver Release Plans

- Windows* reference drivers are targeting two releases per year
- Release 1 is available, and 1.1 work underway

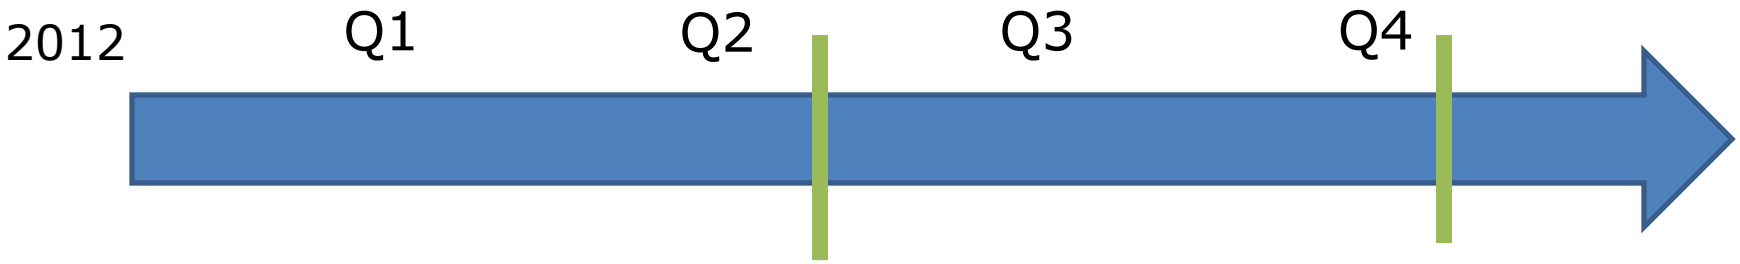

Windows Release 1

- Initial release
- NVMe 1.0b support
- Win7,2008R2

Already Released!

Windows Release 1.1

- NVMe spec updates
- Public IOCTLs
- Win8 Storport specific
- Bug Fixes
- Win7,2008R2,Win8

NVMe = NVM Express

Robust Driver Update Criteria

- The NVMe community is committed to robust reference drivers
- For the Windows* NVMe driver maintained with the OpenFabrics Alliance, there is a detailed update process:

Review Criteria:

- Patches submitted by anyone, email to distribution list
- Patch submission should include time sensitivity/expectations and justification for patch (what value it will add, any tradeoffs to consider)
- Patch must be reviewed by at least three NVMe company representatives
- Reviews include compliance with coding guidelines as well as logic

Testing Criteria:

- All patches and release candidates require, at a minimum, the following;
  - 1 hour of data integrity testing using sdstress (Microsoft Tool)
  - 1 hour of heavy stress testing using IoMeter covering, at least, 512B, 4KB and 128KB ranging from 1 OIO to 64 OIO both sequential and random
  - Quick and slow format of both MBR and GPT partitioning
  - Microsoft SCSI Compliance with no failures
- Testing done for all supported OSs for the release

NVMe = NVM Express

Driver Ecosystem Goals

- The long-term goal is for each major OS to ship with a standard NVM Express driver
- The short term goal is to allow NVMe device manufacturers to provide the drivers they need with their products leveraging the reference drivers
- The reference drivers provide high performance, validated and fully compliant drivers to the ecosystem with reasonable licenses (e.g., GPL, BSD)
- "Fork and Merge" to achieve short-term with reference drivers
  - Each NVMe device manufacturer "forks" the reference driver
  - Each NVMe device manufacturer adds in any product specific features
  - Each NVMe device manufacturer "merges" industry-wide applicable changes back to the reference driver

NVMe = NVM Express

Linux* "Fork and Merge"

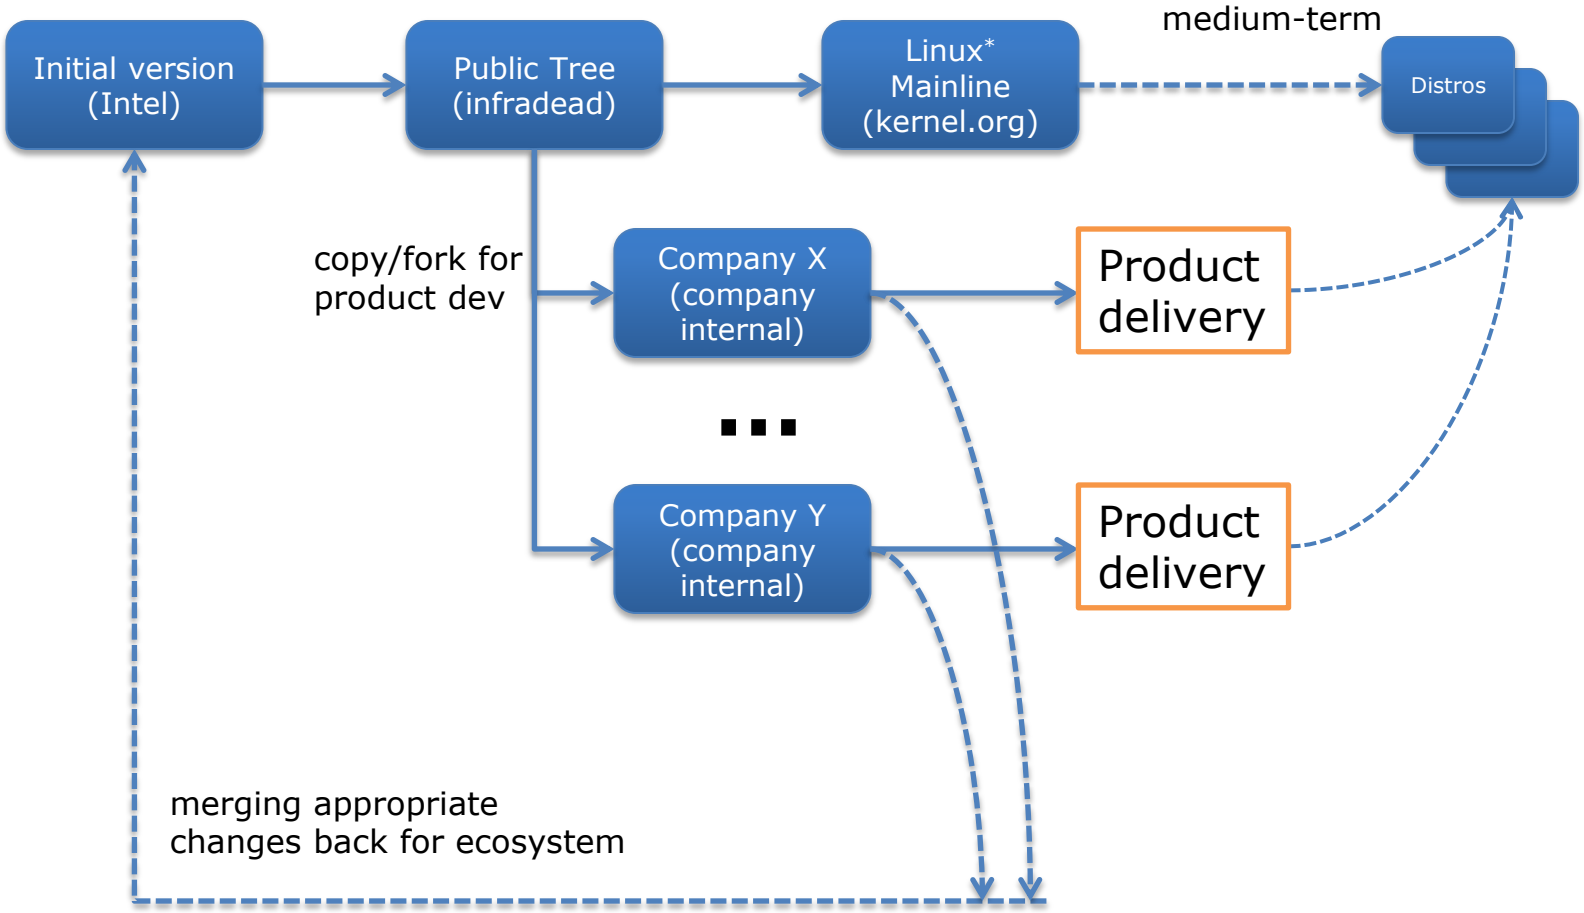

Windows* "Fork and Merge"

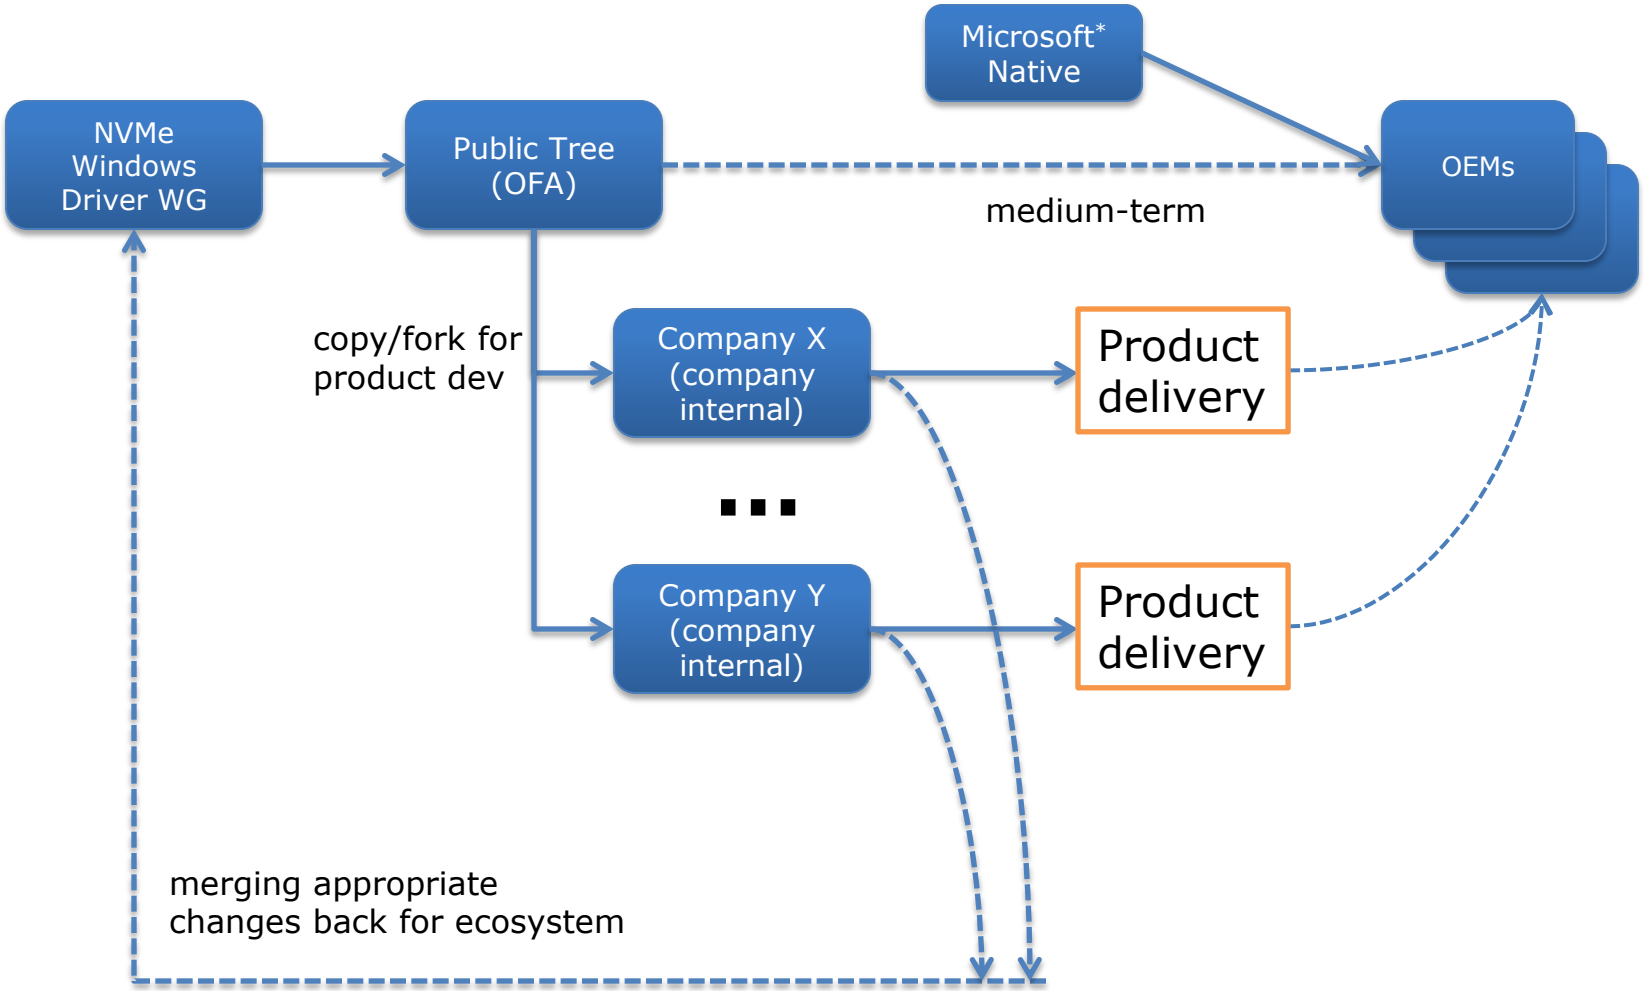

NVMe = NVM Express

VMware* (non-native) "Fork and Merge"

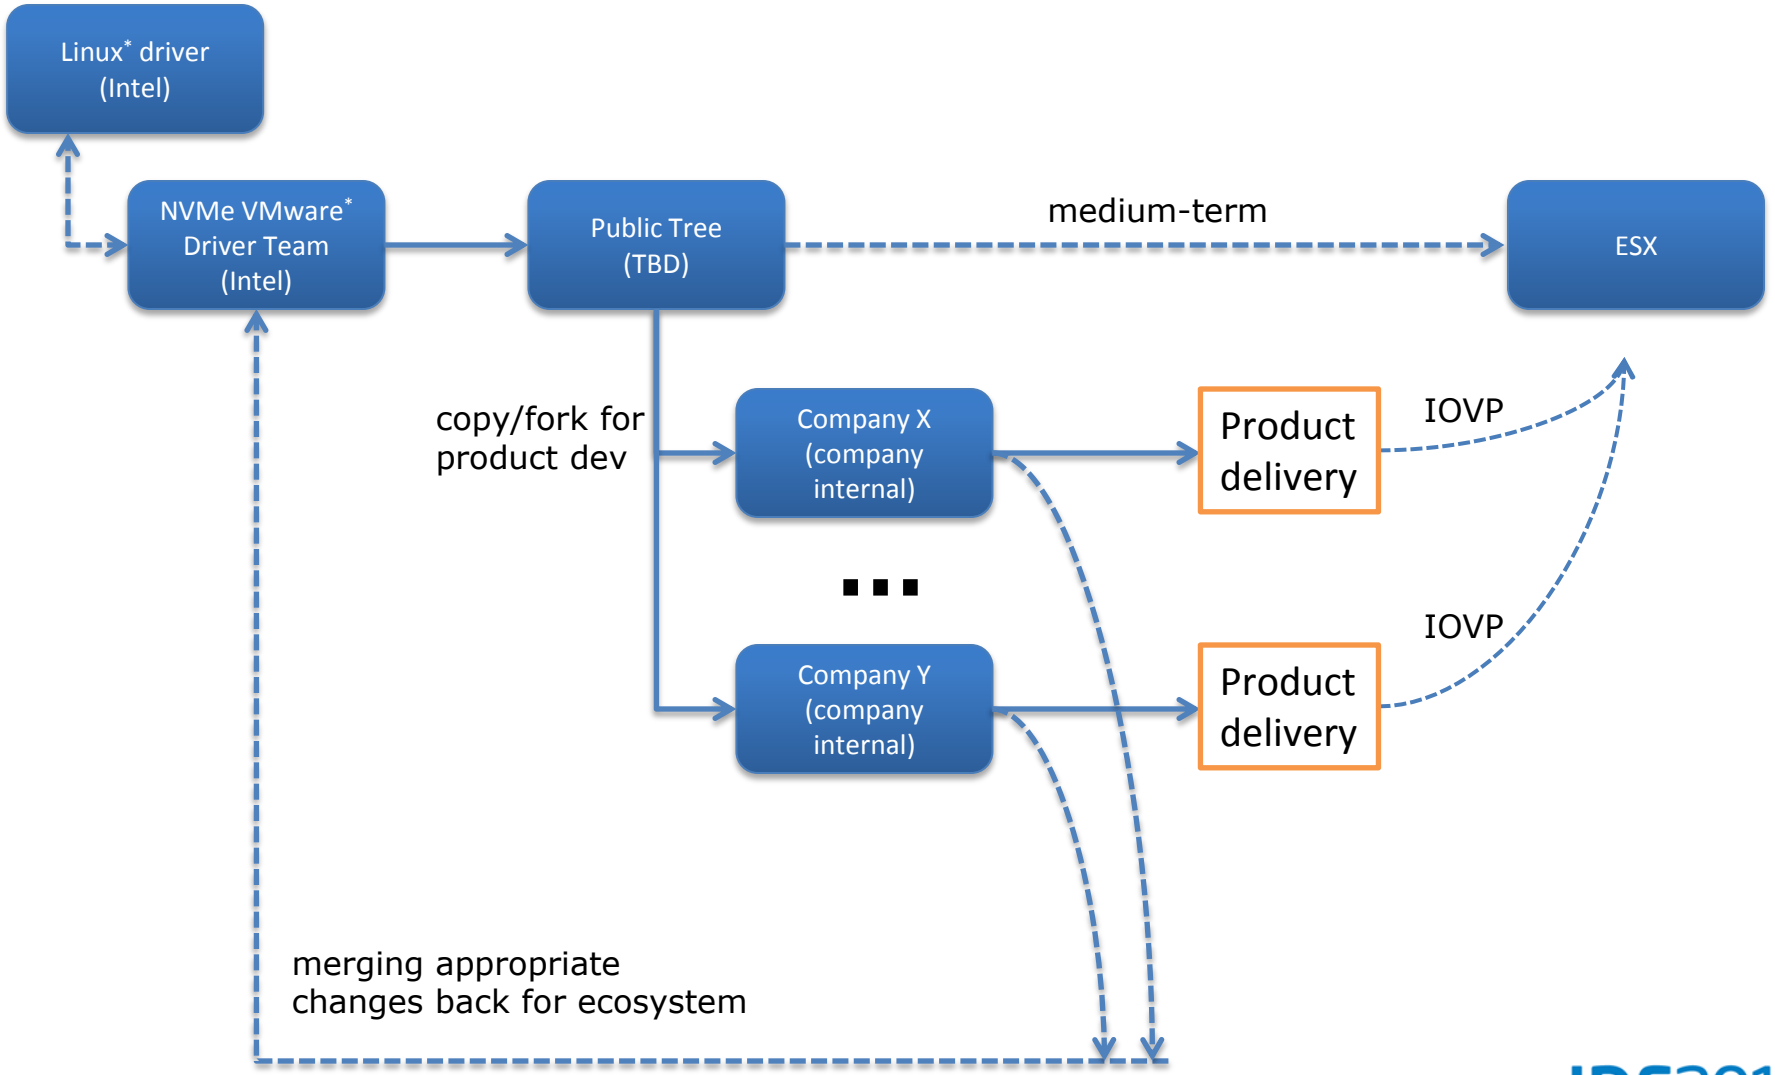

NVMe = NVM Express

Fork and Merge Strategy Summary

- The benefits of the "Fork and Merge" strategy include:
  - Maximize re-use of reference code
  - Enable continuous improvement of the reference drivers
  - Enable product team to focus on delivery goals

- NVM Express (NVMe) device manufacturer responsibilities:
  - Ensure drivers delivered with products have unique binary names
  - Ensure drivers delivered with products bind only to those products (e.g., Micron* version of driver only loads against Micron NVMe SSDs)
  - Merge changes back to the reference driver when appropriate

- The NVMe community is available to support manufacturers:
  - Reference driver maintainers are available to review driver changes
  - Reference driver maintainers will provide guidance that maximizes the ability for a change to be general and flexible for eventual inclusion in the ecosystem driver

Take advantage of reference drivers, and then "give back" to further improve the ecosystem.

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

Interoperability Program Underway

- The NVM Express Workgroup is collaborating with an industry leader, UNH-IOL, to develop the NVMe Interoperability program

March 12, 2012 08:00 AM Eastern Daylight Time

Industry Leaders Develop New, High-Performance PCIe SSD Solutions at UNH-IOL

Lab Now Accepting Founding Member Companies for NVMe Consortium

DURHAM, N.H.--(BUSINESS WIRE)--**The** University of New Hampshire InterOperability Laboratory **(UNH-IOL), an independent provider of broad-based testing and standards conformance services for the networking and storage industries, is accepting founding members for the laboratory's new** Non-Volatile Memory Express (NVMe) Consortium**.** The NVMe Consortium will focus on developing an interoperability test suite for NVMe compliant software and devices. **Founding members will join industry leaders** Dell, EMC, IDT, Intel, LSI Corporation, NetApp, Oracle and SanDisk to create new, innovative, high-performance storage solutions based on the NVMe standard for PCIe SSDs.

University of New Hampshire InterOperability Laboratory

- UNH-IOL has extensive experience in conformance and interop test services for leading industry standards in storage & networking (SATA, SAS, Fibre Channel, etc.)

- Since late 2011 UNH-IOL has been working with the NVMe Promoter Group to develop NVMe test documentation and tools

NVMe is working with UNH-IOL to ensure an interoperable ecosystem that OEMs can count on

NVMe = NVM Express

Specification Conformance

- Conformance: Proving that a product meets all the requirements defined in the specification
- Conformance provides two major benefits:
  - Establishes a foundation that allows future generations of products to be backwards compatible
  - Builds confidence in a new technology
- A conformant product does not imply interoperability
  - Conformance is focused on a single device, interop is system level
- The methods to establish conformance are:
  - Open documentation of conformance requirements
  - Common tools

Resources for NVMe Conformance

- UNH-IOL has created an NVMe conformance test suite document that extracts the requirements of the specification and defines an algorithm for how to test them
  - Feedback is welcome, available for download at: https://www.iol.unh.edu/services/testing/NVMe/testsuites/

- UNH-IOL is delivering T.N.T Software for conformance testing
  - Based on tNVMe tool that Intel contributed to github
  - Available to UNH-IOL members

- LeCroy* and UNH-IOL are collaborating on scripts that can be used with LeCroy's Summit product to test conformance

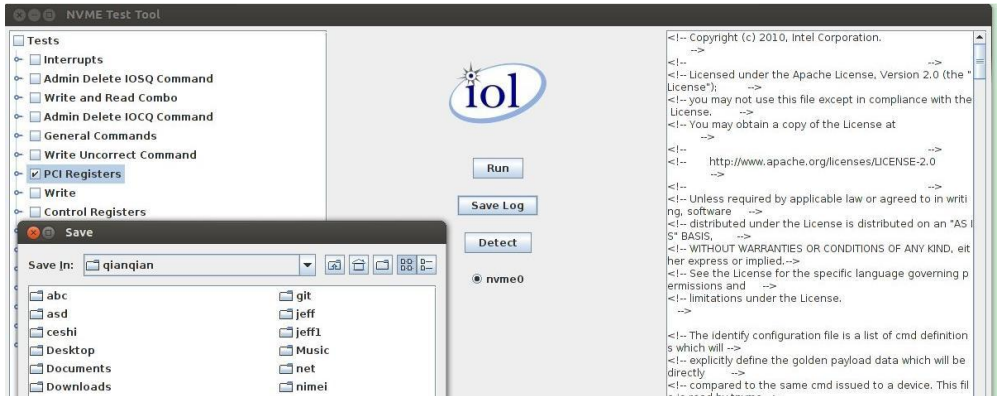

NVMe = NVM Express

Path to an NVMe Integrator's List

- Interoperability testing is at the system level encompassing the NVMe device, software, host chipset, cabling, etc.
  - Conformance shows if a product implemented a feature correctly
- UNH-IOL has created an Interoperability Test Spec to define an interop metric to set the bar for features working together
  - Feedback is welcome, available for download at: https://www.iol.unh.edu/services/testing/NVMe/testsuites/
- UNH-IOL and the NVMe Workgroup are targeting 1H' 2013 for an NVMe plugfest
  - The NVMe Interop Test Suite document is the basis for the plugfest test plan

The NVMe Plugfest in 1H'2013 will be used as the basis for the initial NVMe Integrator's List

NVMe = NVM Express

Agenda

- NVM Express (NVMe) Overview
- New NVMe Features in Enterprise & Client
- Driver Ecosystem for NVMe
- NVMe Interoperability and Plugfest Plans
- EMC's Perspective: NVMe Use Cases and Proof Points

Benefit of Industry Standards

- For developers: a framework for innovation
- For architects: a blueprint and roadmap
  - Relevant today
  - Ready for tomorrow
- For products: an accelerant to broad adoption

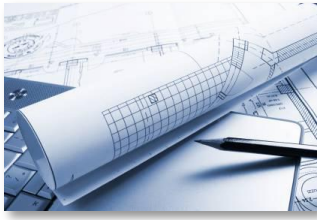

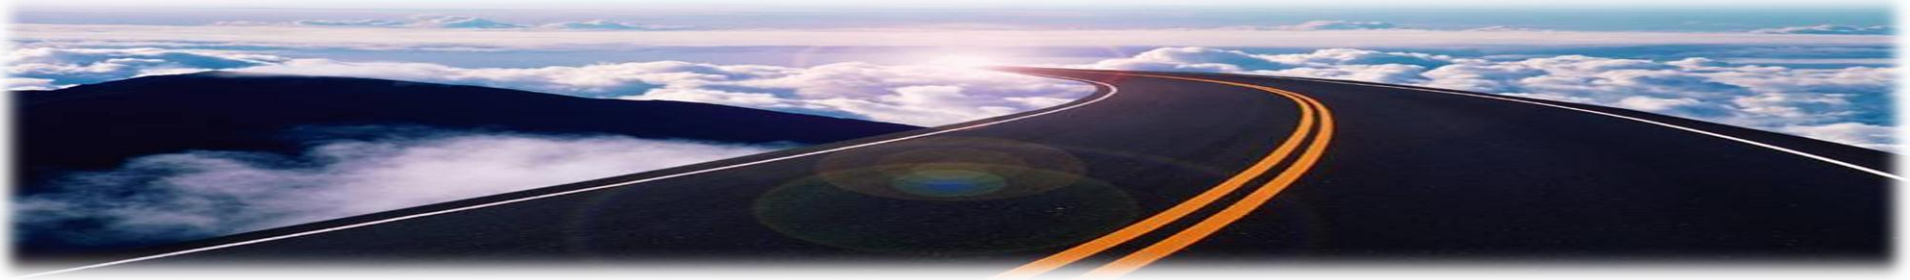

Performance

- Parallelism:
  - Multi-core drives multi-thread drives more operations in flight
- Throughput:
  - More work done, per unit time
- Latency:
  - Governed by Hz, not RPMs
  - From milliseconds to microseconds to...

Designing with the Right Pieces

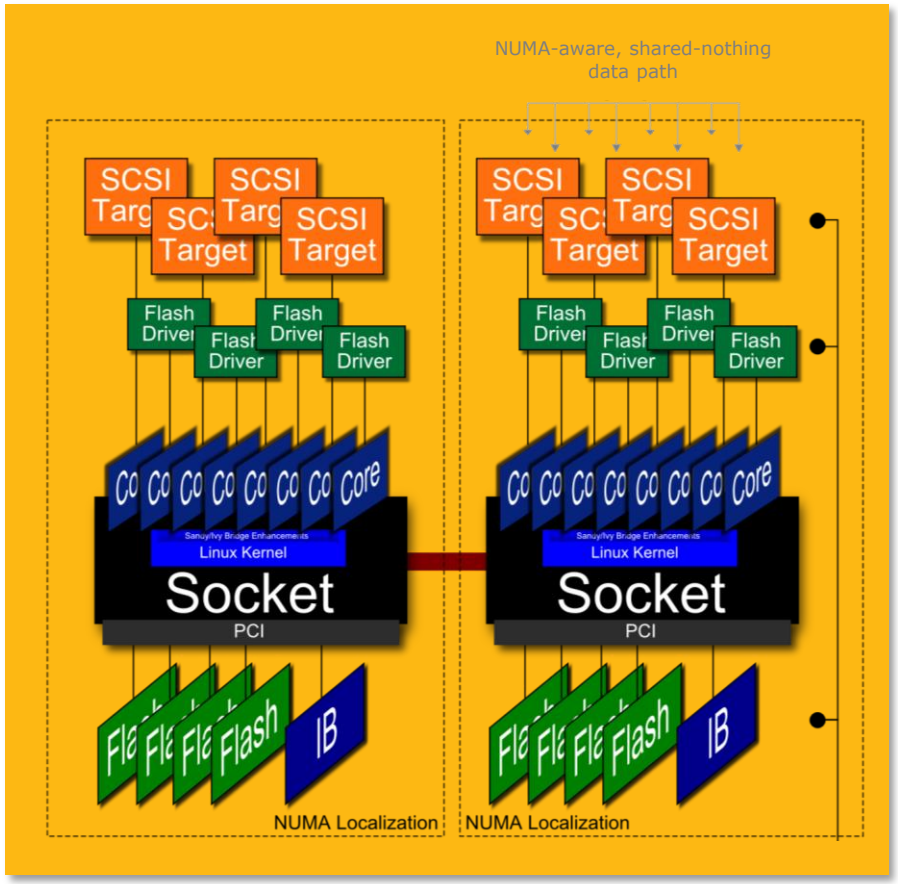

multi-core parallelism
+ PCI Express* Gen3 bandwidth
+ NUMA-aware software
+ NVMe flash
= millions of IOPs...

...at microsecond latencies

new math for storage platforms

NVMe = NVM Express

Proof Points in EMC* Products

Project Thunder: A networked non-volatile memory appliance

- Building block design center:
  - 10-20TB of flash capacity
  - Consistent 2.5M IOP throughput @ 150us
- NVMe ready:
  - Today: improved latency, reduced processor overhead
  - Tomorrow: ready for nano-second class NVM
  - On the showcase floor until **2pm**

NVMe = NVM Express

Use Cases Important to EMC*

- Ongoing: storage platforms
  - Better, faster, bigger
- Growing: Big-data
  - Speed, size, variety, social, mobile
- Emerging: High-frequency storage
  - **Time is money:** high-frequency applications
  - **Decisions as a service:** risk-management, confidence-management

Summary

- NVM Express is a scalable host controller interface designed for Enterprise and client systems that use PCI Express* SSDs
- NVMe 1.1 adds new features for Enterprise and Client and has started formal 30-day ratification period
  - Multi-path I/O and namespace sharing for Enterprise
  - Lower power through autonomous transitions during idle for Client
- The NVMe driver ecosystem has solutions for Linux*, Windows*, VMware*, Solaris*, and UEFI coming online
- NVMe is working with UNH-IOL to enable an interoperable ecosystem that OEMs can count on; check out the plugfest in Q1
- EMC sees tremendous opportunity with NVMe for emerging high-frequency storage, growing big-data, and making existing storage platforms better, faster, and cheaper

Take part in the NVMe Revolution – nvmexpress.org

NVMe = NVM Express

Want More Info on SSDs?

- Attend or download these SSD-related sessions

  Wednesday, Sept 12th

  - SSDS002 - Data Center Solid-State Drive Requirements
  - Hands-on-Lab – PCI Express and SATA Intel SSDS

  Thursday, Sept 13th

  - SSDS003 - NVM Express and the PCI Express SSD Revolution
  - SSDS004 - Solid-State Drives (SSDs) Enabling Business Ultrabooks
  - SSDS005 - Solid-State Drives (SSDs) and Large-Scale Corporate PC Deployments: Learn from the Experts
  - SSDS001 – PCI Express Solid-State Drives (SSDs): Trends and Opportunities

- Visit Intel Booth #323 on Level 1 of the Tech Showcase
  - SSD vs. HDD comparisons, working Intel SSD 910 Series PCIe demo
- Visit Intel online at www.intel.com/go/ssd
  - Product briefs, datasheets, whitepapers, videos, technical support

Please Fill Out The Online Session Evaluation Form

Enter to win fabulous prizes including Ultrabooks™, SSDs and more!

You will receive an email with a link to the online session evaluation prior to the end of this session. Please submit the evaluation by 10am tomorrow to be entered to win.

Winners will be announced by email

Sweepstakes rules are available at the Help Desk on Level 2
All sessions evaluations must be submitted by Friday Sept 14 at 5pm

IDF2012
INTEL DEVELOPER FORUM

Q&A

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

- A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.
- Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.
- The products described in this document 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.
- Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor_number.
- Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.
- Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: http://www.intel.com/design/literature.htm
- Intel, Ultrabook, Sponsors of Tomorrow 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.
-

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 the important factors that could cause actual results to differ materially from the company's expectations. Demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including supply constraints and other disruptions affecting customers; customer acceptance of Intel's and competitors' products; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Uncertainty in global economic and financial conditions poses a risk that consumers and businesses may defer purchases in response to negative financial events, which could negatively affect product demand and other related matters. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; and Intel's ability to respond quickly to technological developments and to incorporate new features into its products. Intel is in the process of transitioning to its next generation of products on 22nm process technology, and there could be execution and timing issues associated with these changes, including products defects and errata and lower than anticipated manufacturing yields. The 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; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. The majority of Intel's non-marketable equity investment portfolio balance is concentrated in companies in the flash memory market segment, and declines in this market segment or changes in management's plans with respect to Intel's investments in this market segment could result in significant impairment charges, impacting restructuring charges as well as gains/losses on equity investments and interest and other. 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. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. Intel's results could be affected by the timing of closing of acquisitions and divestitures. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. 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. 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 Form 10-Q, Form 10-K and earnings release.

Rev. 5/4/12