Solid-State Drives in Computing

Delivering on the SSD Potential

Santa Clara, CA USA
August 2008

Agenda

Solid-State Drives in Computing

- Delivering on the performance potential
- Delivering on device and system power potential
- Delivering on endurance requirements

Storage Performance Vectors

- Sequential performance
  - “Transfer rate”
- Random I/Os per second (IOPS)
  - “Access time”
- Benchmark results
  - “What it brings the system”
- Power efficiency
  - “Power per work done” (or power per fixed work amount)

All vectors relevant, not just selected ones that look good

Relative Sequential Transfer Rates

This is most commonly cited metric.
Sequential transfer rate is not a simple scalar.

For write measurements, drive first filled by writing to all blocks on the device. For random write measurement, 10% of the device span exercised to represent client usage. Measurements made using IOMeter.

Performance measurements are made using specific computer systems and/or components and reflect the approximate performance of the technology as measured by those tests. Any difference in system hardware or software design or configuration may affect actual results.

Why IOPS Matter
(usually more than transfer rate)

>95% of total service time is mechanical latency

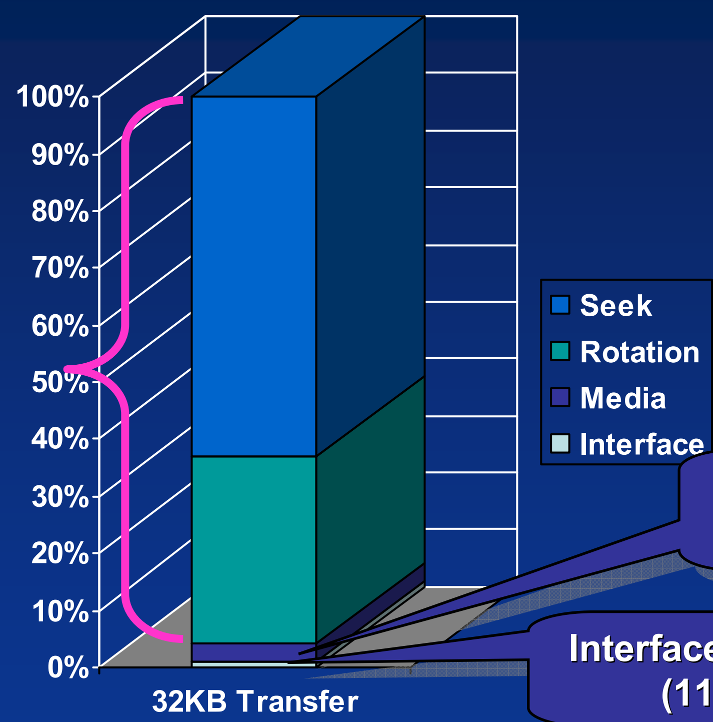

Approximate service time breakdown for 7200RPM HDD w/ 8ms average seek time and 75MB/s transfer rate performing 32KB random read operation.

For many accesses, the non-transfer time component is dominant. IOPS is one metric reflecting this component.

Relative Random IOPS Performance

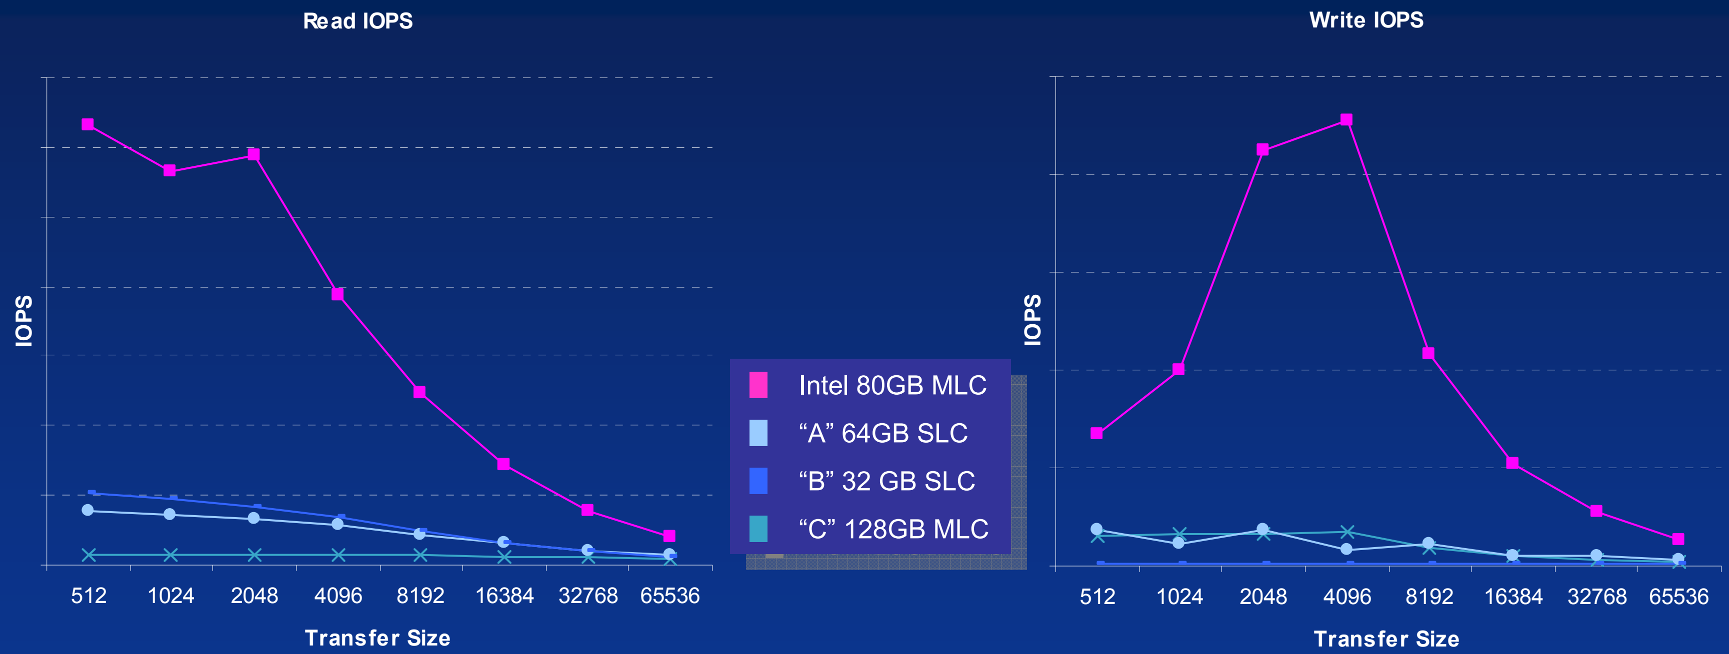

For smaller transfers, IOPS performance is better metric

For write measurements, drive first filled by writing to all blocks on the device. For client random write measurement, 10% of the device span exercised to represent client usage. Measurements made using IOMeter.

Performance measurements are made using specific computer systems and/or components and reflect the approximate performance of the technology as measured by those tests. Any difference in system hardware or software design or configuration may affect actual results.

Benchmark/System Performance

- Simplified view – Net result is workload-specific blend of the transfer rate and IOPS performance attributes
  - Both IOPS and transfer rate important
  - Very difficult to make up for one with the other
- Server workloads are highly weighted toward IOPS while client workloads have more weight toward transfer rate
  - This distinguishes tuning and attributes for server vs client drives

Benchmark Impact

Not all SSDs are created equal

Relative HDD Benchmark Performance

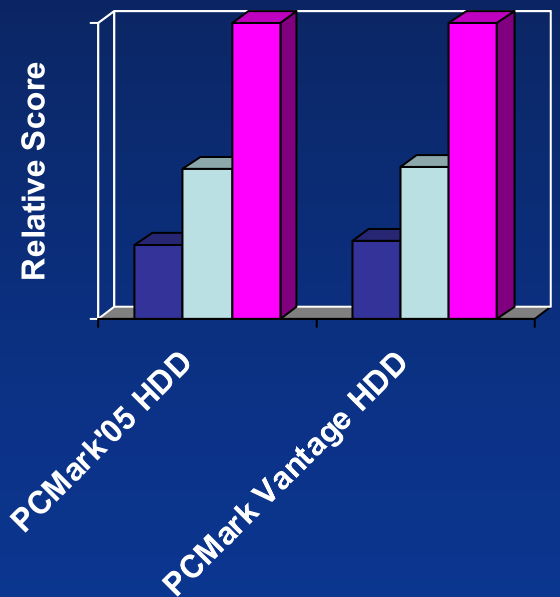

- Intel 80GB MLC
- "A" 64GB SLC
- "B" 32 GB SLC

Relative System Benchmark Gain

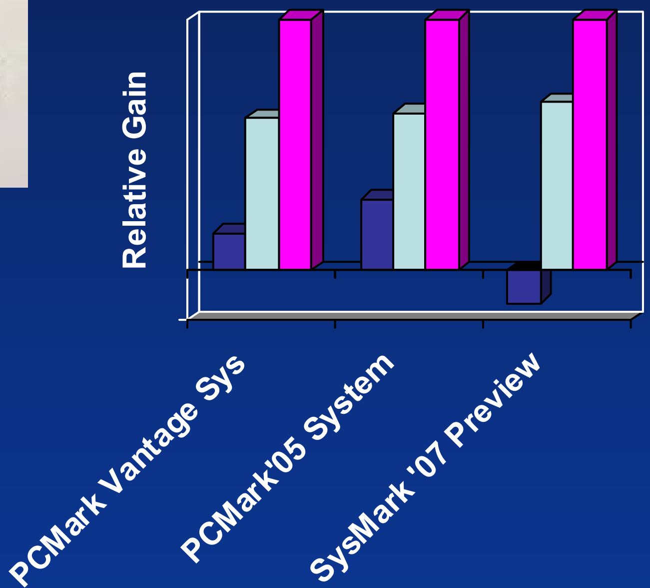

Good SSD performance translates to system gains

* Performance tests and ratings are measured using Dell D830 or HP 6910p SantaRosa notebook 2.0GHz with Merom processor and 2GB DRAM running Vista Enterprise Edition and reflects the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Data collected with 50nm NAND.
** Other brands and names are the property of their respective owners

Benefits of Faster Intel SSD Most Realized with Intense Workloads

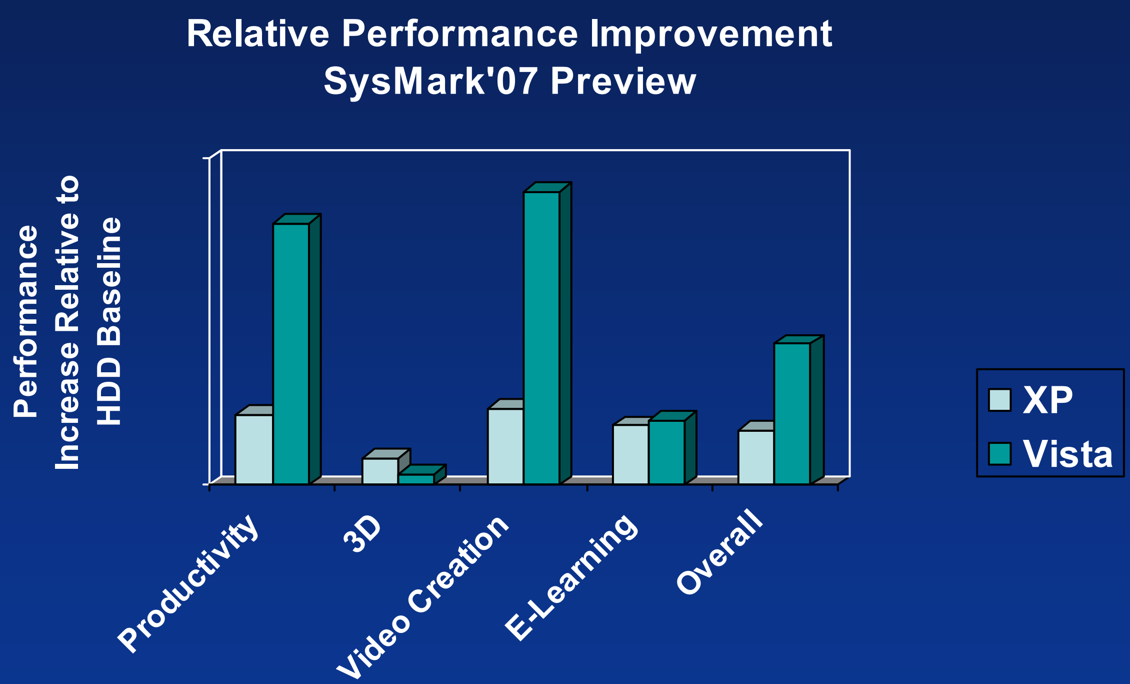

The benefits of a faster storage subsystem is most realized with disk-intensive environments that utilize it

* Performance tests and ratings are measured using Dell D830 or HP 6910p SantaRosa notebook 2.0GHz with Merom processor and 2GB DRAM running Vista Enterprise Edition and reflects the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Data collected with 50nm NAND.

SSD Power Efficiency

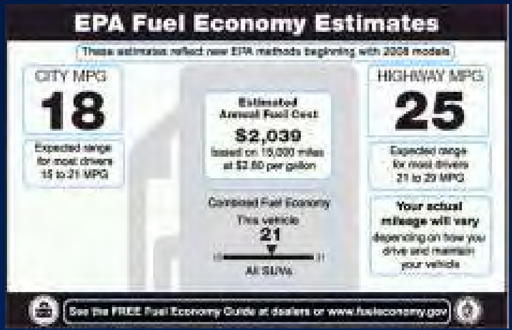

- Most HDDs and SSDs specify "Active" power
  - No standard definition of what "Active" means
  - "Active" doing what?
    - Analogous to "Gallons per hour at maximum speed"
    - Only useful comparison if everyone has same maximum speed
- What matters is IOPS/Watt (Enterprise) or the area under the power curve (Client).
  - For enterprise applications measure work done per watt consumed using a specified workload
    - Analogous to "Miles per Gallon (city)"
  - For client applications measure power for representative usage with benchmarks like MobileMark*
    - Analogous to "Gallons for standardized 200-mile road trip"

HDD Power During 2hr Mobile Workload

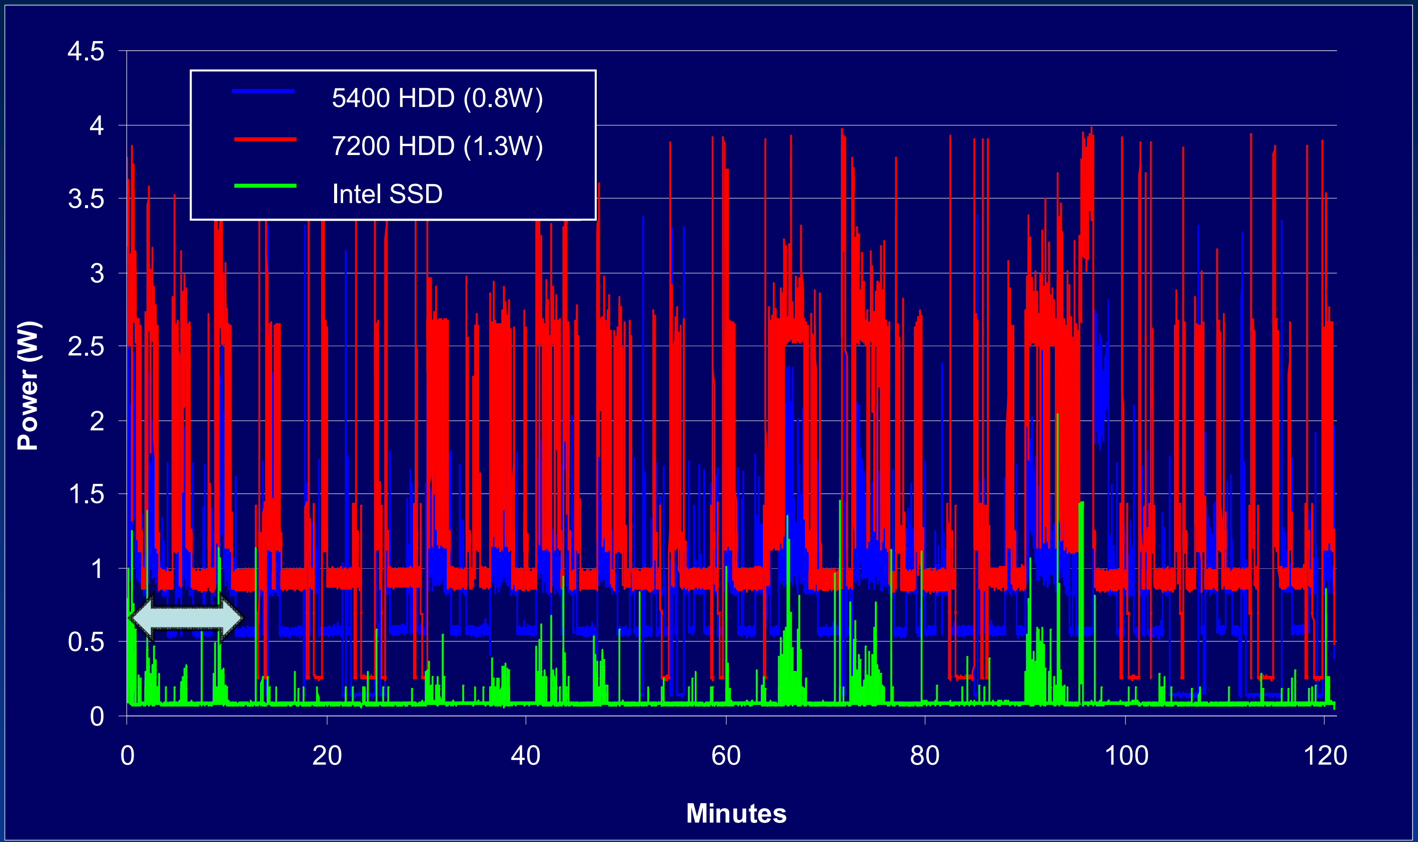

Performance measurements are made using specific computer systems and/or components and reflect the approximate performance of the technology as measured by those tests. Any difference in system hardware or software design or configuration may affect actual results.

HDD Power During 2hr Mobile Workload

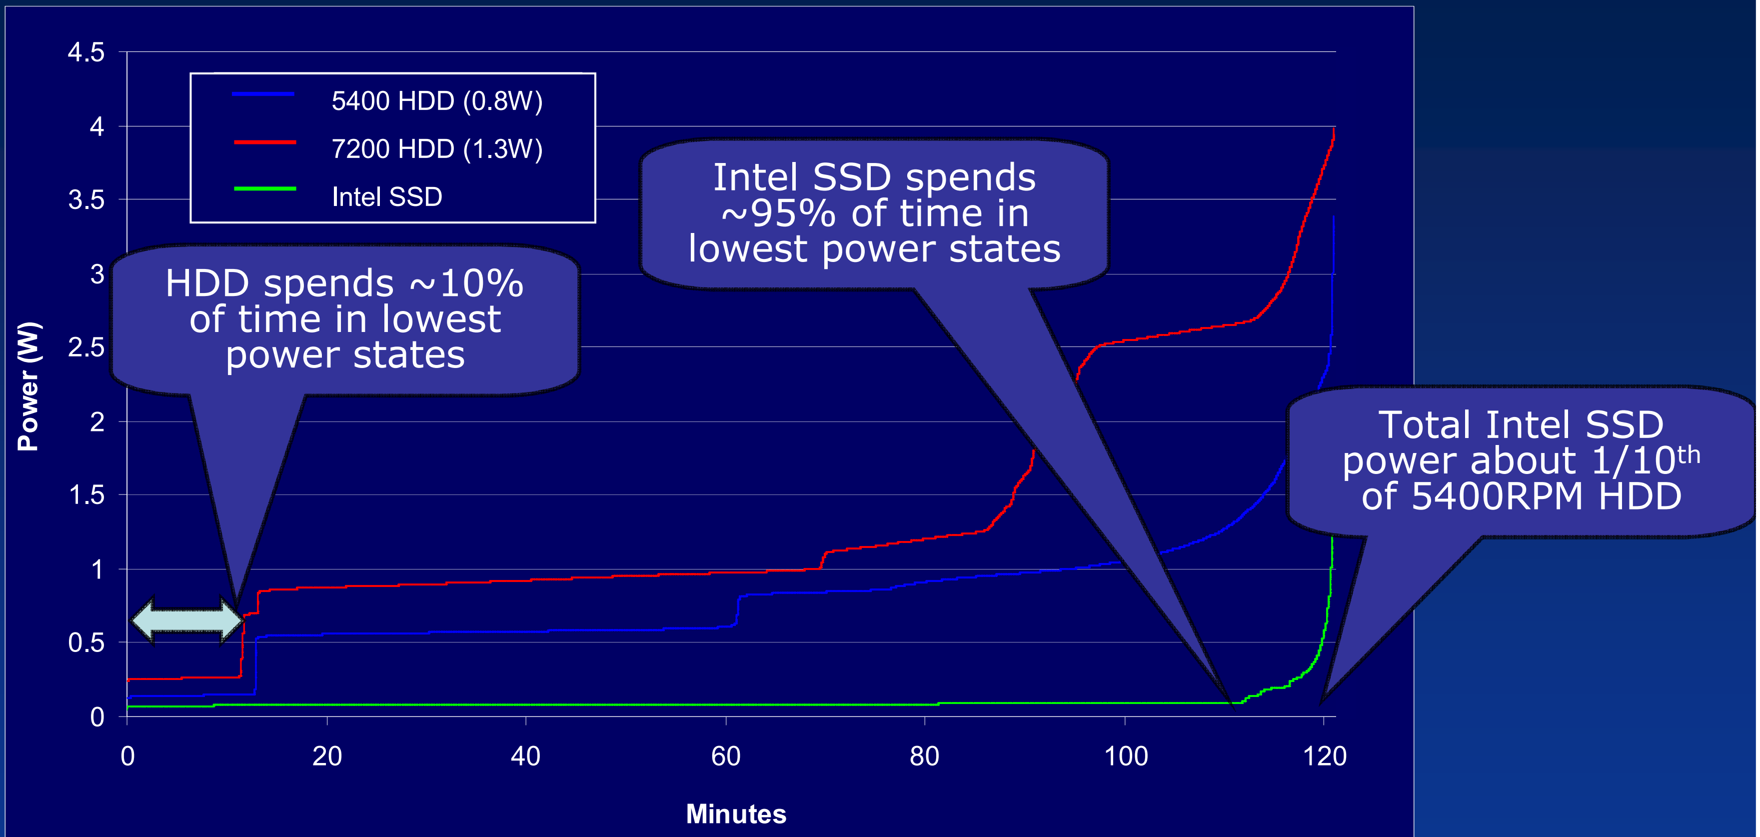

Combination of low power and high performance yields best energy efficiency

Performance measurements are made using specific computer systems and/or components and reflect the approximate performance of the technology as measured by those tests. Any difference in system hardware or software design or configuration may affect actual results.

System Power Savings are Real

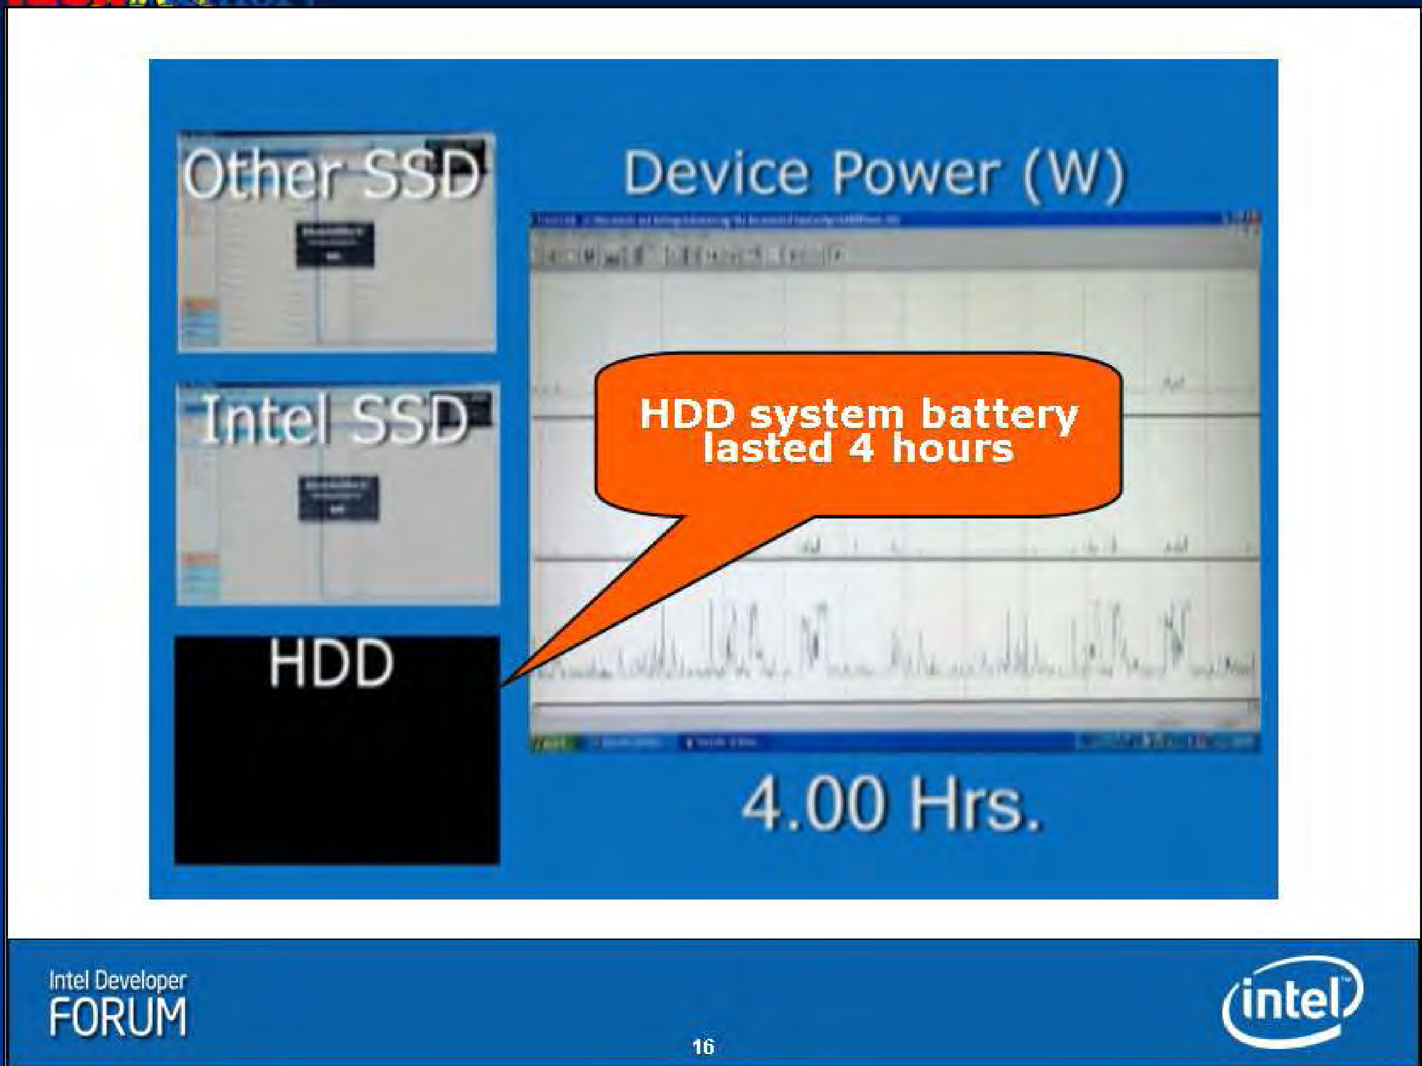

System Power Savings are Real

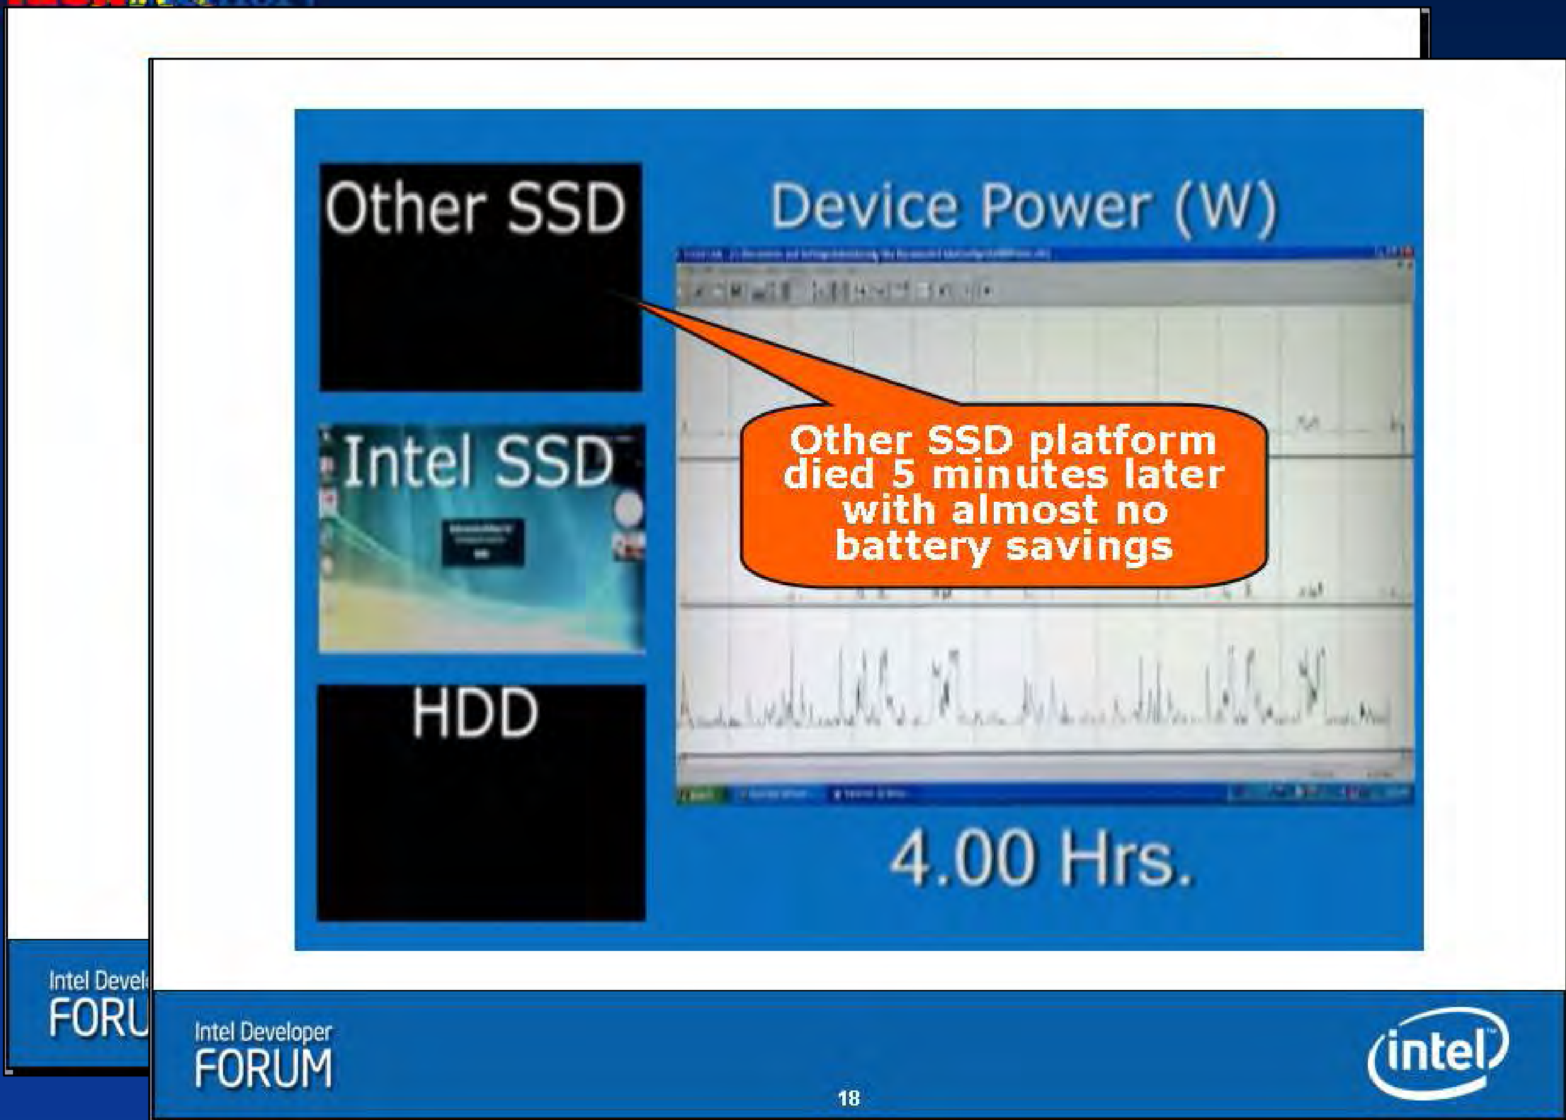

System Power Savings are Real

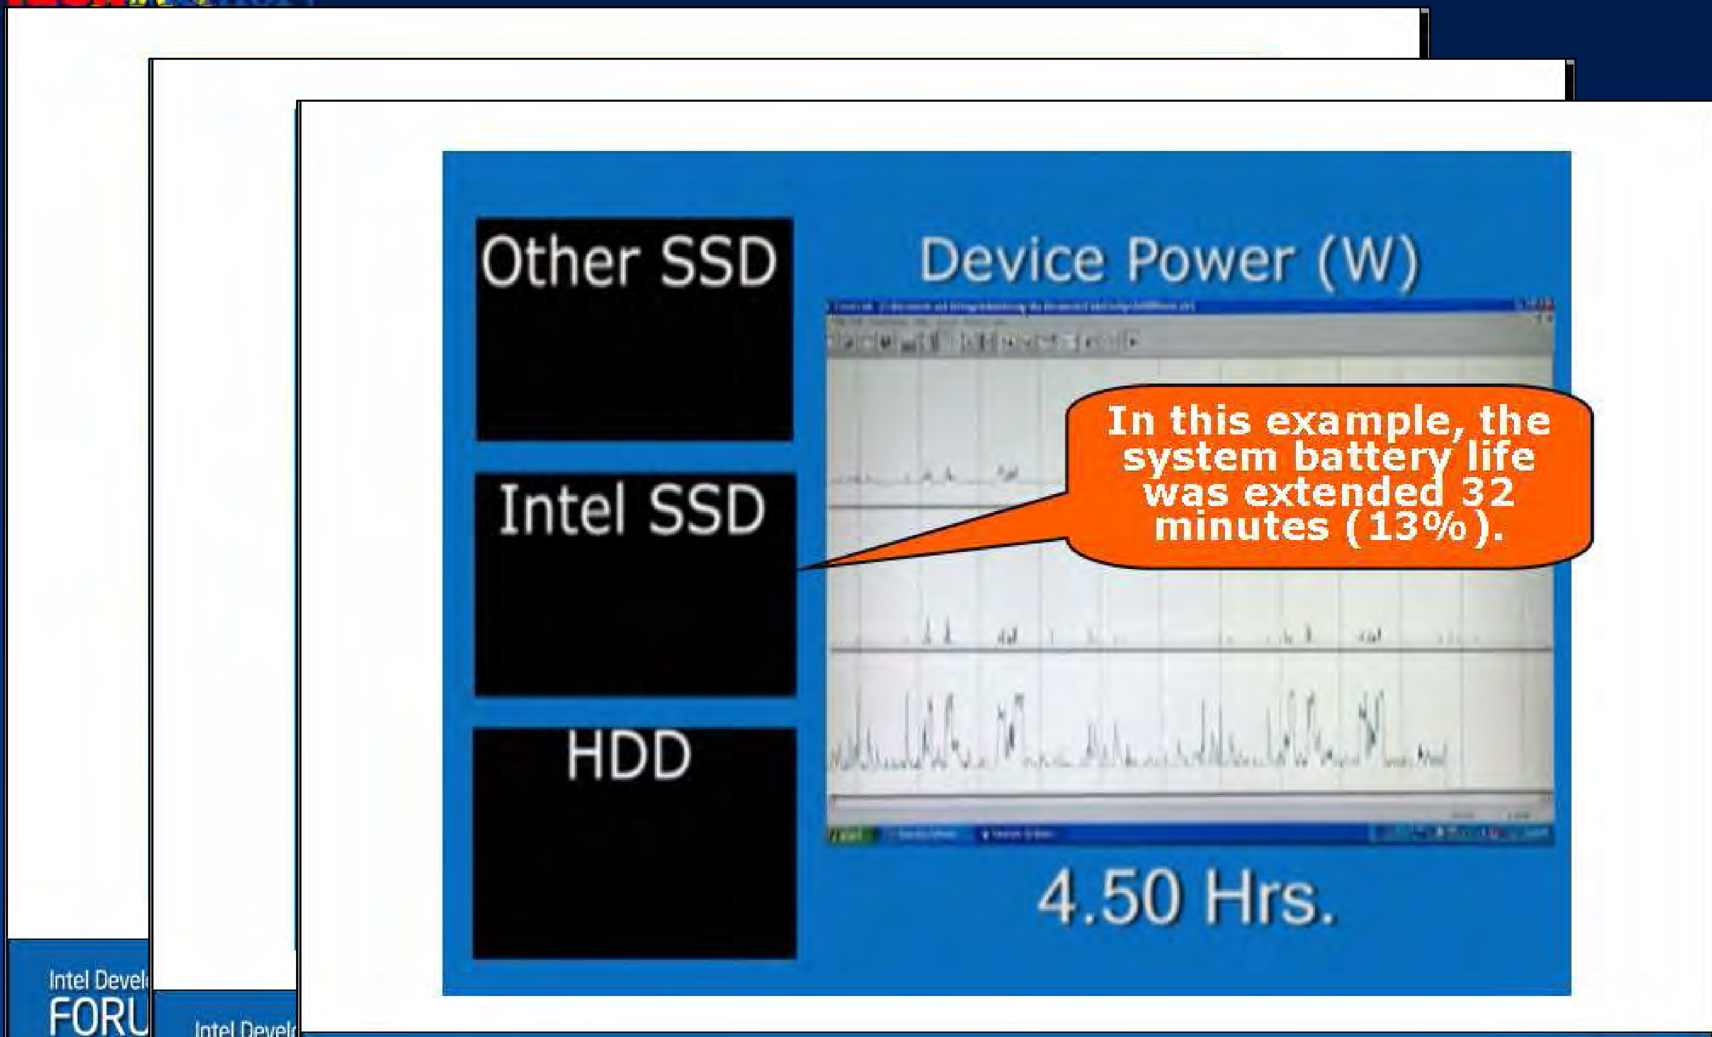

30 minutes (13%) battery life extension when HDD power consumption only 9%. Secondary effects compound savings.

Endurance – a Frank Discussion

Endurance & Cycling

- Number of NAND cycles is not equal to data written divided by SSD capacity
- Oversimplified view is typically off by *several orders of magnitude*

Intel Developer FORUM

Corollary: Writes NAND Cycles * Capacity

Slides from Intel Developer Forum spring 2008 in Shanghai, class MASS001: "SATA Solid-State Drives: Not All Drives are Created Equal"

Endurance – a Frank Discussion

End

Accurate Cycling Computation

- Realizable endurance and NAND cycling for given capacity stems from three factors
  - Amount of data written by host
  - Write amplification
  - Wear leveling efficiency
- The efficiencies of the algorithms largely determine the endurance of the device

$$Cycles = \frac{(Host\ Writes) * (Write\ Amplification\ Factor) * (Wear\ Leveling\ Factor)}{Device\ Capacity}$$

- Nu div
- Ove ord

Intel Developer FORUM

Intel Developer FORUM

10

Slides from Intel Developer Forum spring 2008 in Shanghai, class MASS001: "SATA Solid-State Drives: Not All Drives are Created Equal"

Endurance – a Frank Discussion

NAND Cycling Wrap-up

- Reliable NAND cycling is a combination of several factors
  - Usage model and user data written
  - Write amplification
  - Wear leveling performance

$$\text{Cycles} = \frac{(\text{Host Writes}) * (\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})}{\text{Device Capacity}}$$

- Example 32GB config using Intel vs traditional technology

$$\text{Intel Cycles} = \frac{(40\text{TB}) * (1.1) * (1.1)}{32\text{GB}} \cong 1{,}500\ \text{cycles}$$

$$\text{Other's Cycles} = \frac{(40\text{TB}) * (20) * (3)}{32\text{GB}} \cong 75{,}000\ \text{cycles}$$

Intel write amplification and wear leveling efficiency supports higher host write rates with less NAND cycling

Intel Developer FORUM

intel

16

Slides from Intel Developer Forum spring 2008 in Shanghai, class MASS001: "SATA Solid-State Drives: Not All Drives are Created Equal"

What This Means

- It means what is says: "*Intel write amplification and wear leveling efficiency supports higher host writes rates with less NAND cycling*"
- With Intel's technology, greater cycling headroom can be delivered with MLC than traditional technology may deliver with SLC
  - Traditional technology may be challenged to deliver reliable MLC solutions due to cycling efficiency
- Using Intel's technology, greater write performance can be safely delivered (especially IOPS)
  - Improved write efficiency naturally yields higher write performance, and improved cycling efficiency safely supports more write intensive applications that exercise this

Accurately Simplifying the Metric

- Total supported host writes is fine metric if accurately computed to account for all the significant factors. Basic algebra.

$$(\text{Cycles}) = \frac{(\text{Host Writes}) * (\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})}{(\text{Device Capacity})}$$

$$\frac{(\text{Cycles})}{(\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})} = \frac{(\text{Host Writes})}{(\text{Device Capacity})}$$

Accurately Simplifying the Metric

- Total supported host writes is fine metric if accurately computed to account for all the significant factors. Basic algebra.

$$\frac{(\text{Cycles})}{(\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})} = \frac{(\text{Host Writes})}{(\text{Device Capacity})}$$

$$\frac{(\text{Cycles}) * (\text{Device Capacity})}{(\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})} = (\text{Host Writes})$$

Accurately Simplifying the Metric

- Total supported host writes is fine metric if accurately computed to account for all the significant factors. Basic algebra.

$$(\text{Host Writes}) = \frac{(\text{Cycles}) * (\text{Device Capacity})}{(\text{Write Amplification Factor}) * (\text{Wear Leveling Factor})}$$

- Results based on same 32GB example as previous:

$$\text{Intel MLC Writes} \approx \frac{(5{,}000) * (32\text{GB})}{(1.1) * (1.1)} \approx 132TB$$

$$\text{Traditional SLC Writes} \approx \frac{(100{,}000) * (32\text{GB})}{(20) * (3)} \approx 53TB$$

- To be accurate, factors must be determined based on representative stimulus & steady state (filled drive). Since workload specific, client values will not apply to server workloads & environments.

Accurate value must comprehend algorithmic efficiency. Resulting value not applicable to other workloads.

Summary

- Well-rounded performance has substantial system performance impact
  - IOPS, Transfer Rate, Power all matter
  - Simple scalars can obscure true performance
- The performance and power potential of SSDs can be delivered, even with MLC technology
  - But not all SSD's are created equal
- Write amplification and wear leveling efficiency are fundamental to SSD cycling and endurance
  - Endurance simplifications must account for algorithmic efficiency to be accurate metrics, and results are specific to measured workload

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. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards 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.
- Any code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.
- Intel, Intel Inside, 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

This presentation contains forward-looking statements that involve a number of risks and uncertainties. These statements do not reflect the potential impact of any mergers, acquisitions, divestitures, investments or other similar transactions that may be completed in the future. The information presented is accurate only as of today's date and will not be updated. In addition to any factors discussed in the presentation, the important factors that could cause actual results to differ materially include the following: Demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including conditions in the credit market that could affect consumer confidence; 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. Intel's results could be affected by the timing of closing of acquisitions and divestitures. 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 new 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; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient supply of components from suppliers to meet demand. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; product mix and pricing; capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; excess or obsolete inventory; manufacturing yields; changes in unit costs; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, vary depending on the level of demand for Intel's products, the level of revenue and profits, and impairments of long-lived assets. Intel is in the midst of a structure and efficiency program that is resulting in several actions that could have an impact on expected expense levels and gross margin. Intel's results could be impacted by adverse economic, social, political and physical/infrastructure conditions in the countries in which 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. 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 and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the report on Form 10-Q for the quarter ended June 28, 2008.