



# Designing Solid-State Drives (SSDs) into Data Center Solutions

Tony Roug  
Principal Engineer  
Datacenter Group

Bhaskar Gowda  
Systems Engineer  
Datacenter Group

MEMS001

Sponsors of Tomorrow: 

# Agenda

- **Data center SSD opportunity**
- **ROI in data center applications**
- **Building a high performance datacenter configuration**
- **Q/A**



# Agenda

- **Data center SSD opportunity**
- ROI in data center applications
- Building a high performance datacenter configuration
- Q/A

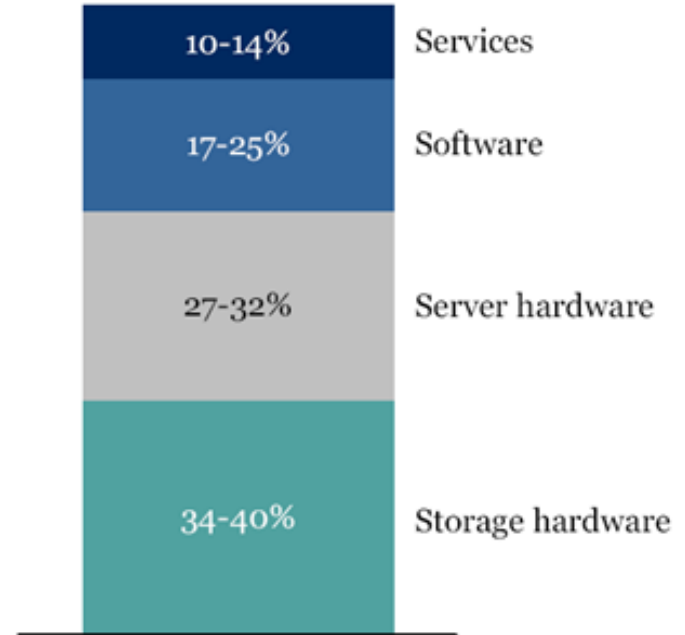


# Storage in Datacenters

*Intel estimates  
over 50% of 2009  
Server Workloads  
are I/O bound*

## Typical VMWare server consolidation project

100% = \$44,000-62,000

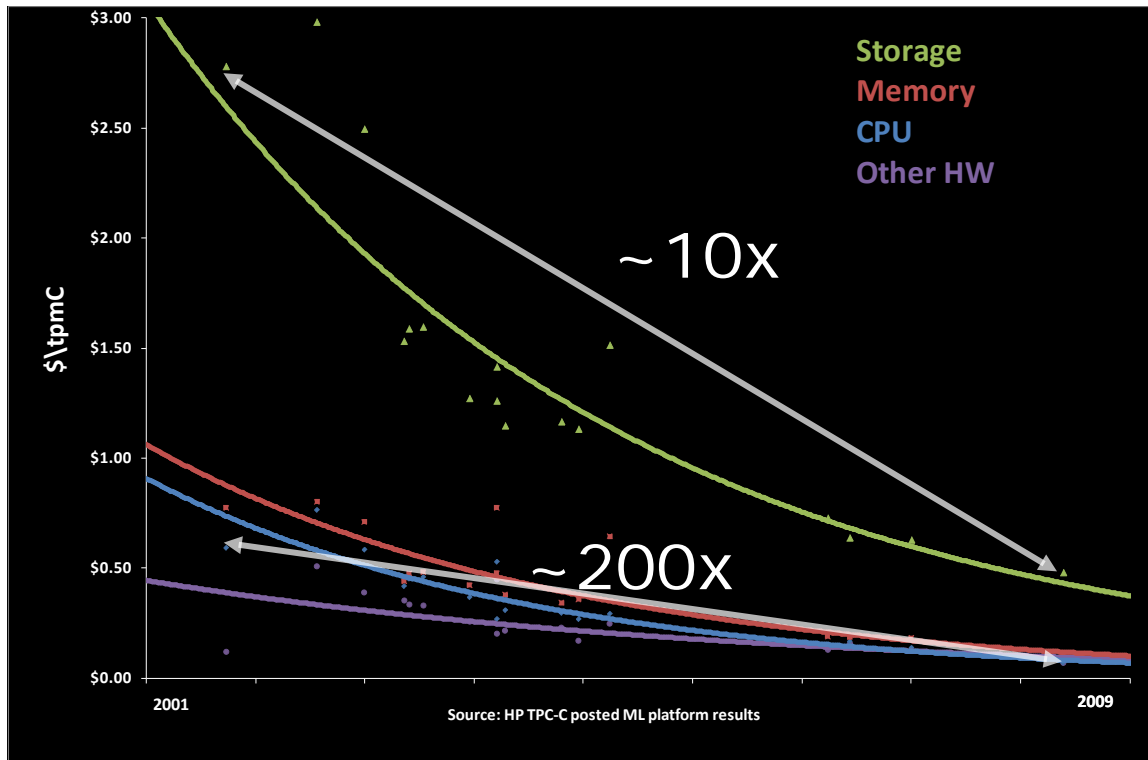


Source: TechAlpha analysis based on VMWorld interviews


***Storage hardware accounts for the single biggest share of datacenter project CAPEX***

# Storage in Datacenters

## TPC-C Case Study: Solution Costs



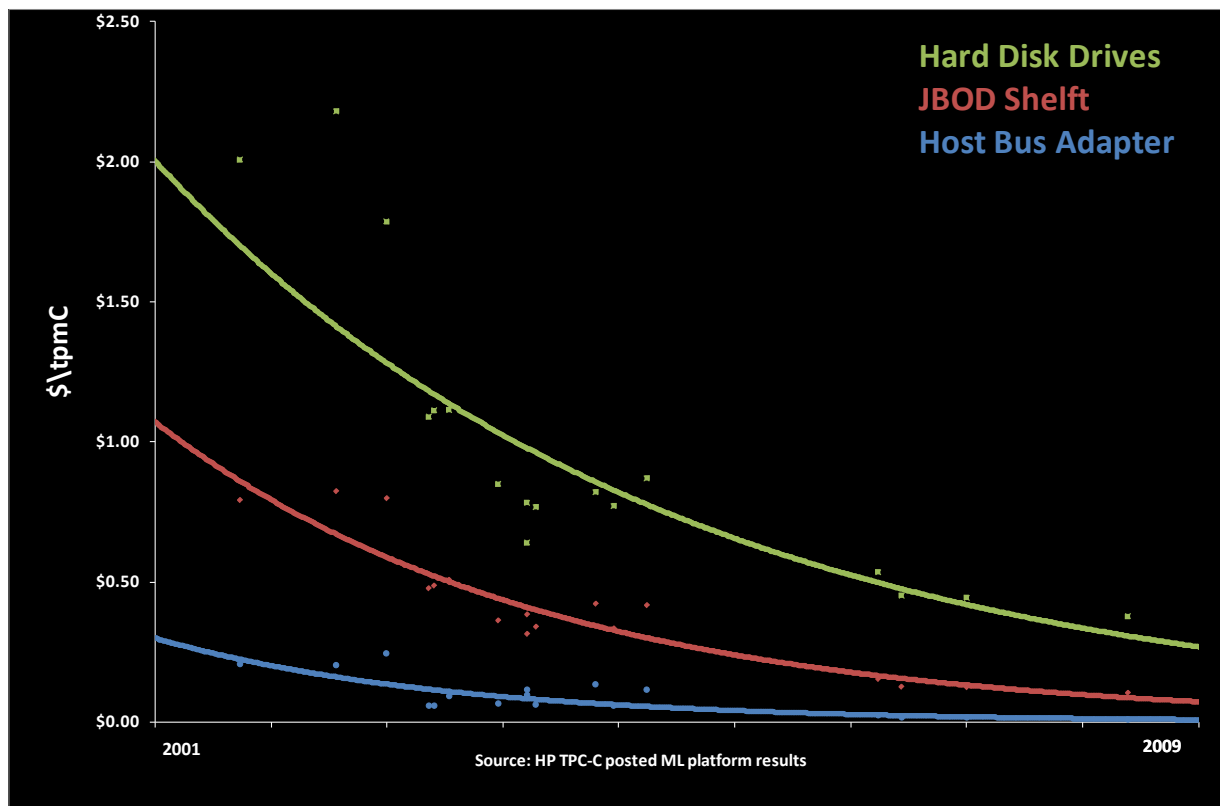
May 29, 2009  
Config




*CPU cost has reduce 20x more than storage*

# Storage in Datacenters

## TPC-C Case Study: Storage Cost Breakdown



May 29, 2009  
Config



*Hard-Disk Drives dominate: 77% of cost*  
*Hard-Disk Drives with JBOD: 98% of cost*

# Storage in the Datacenter



## Storage Usage

### Boot



### Cache



### Performance



### Capacity



## Central Storage (CAS, SAN, NAS)

### Boot Data:

(Operating System, Hypervisor, SWAP, Application Image)

### DRAM cache:

(e.g. OS block, hot application data)

### Hot Application Data

(Database, Email, etc)

### Cold/Luke-warm Application Data

(Data warehouse, Documents, Backups, etc)

## Network

### Network Boot:

(e.g. PXE)

### Network protocol:

(e.g. Fiber Channel (FC), FC over Ethernet (FCoE), iSCSI, NFS, CIFS)

## Server Attached (DAS)

### Boot Data:

(Operating System, Hypervisor, SWAP, Application Image)

### DRAM cache:

(OS, Database, etc )








### Hot Application Data

(Web, Database, Email, Search, Videos, etc)

### Luke-warm Application Data

(Web, Email, Videos, etc)

# Estimated 2009 HDD Opportunity

	 <b>Central Storage (CAS, SAN, NAS)</b>	 <b>Network</b>	 <b>Server Attached (DAS)</b>
<b>Storage Usage</b>			
<b>Boot</b> 	SATA 5.4K, SAS 7.2K <b>~ 1M HDDs*</b>		SATA 5.4, SAS 7.2K <b>~ 10M HDDs*</b>
<b>Cache</b> 			4-128GB
<b>Performance</b> 	SAS 10K, 15K <b>~ 20M HDDs*</b>		SAS 10K, 15K <b>~ 20M HDDs*</b>
<b>Capacity</b> 	SATA 5.4K, SAS 7.2K <b>~ 30M HDDs*</b>		SATA 5.4K, SAS 7.2K <b>~ 10M HDDs*</b>

*\*Intel projections linked to IDC Storage Tracker*

# Intel® High-Performance SATA Solid-State Drives

	Intel® X25-E SATA SSDs	Intel® X25-M and X18-M SATA SDDs	
<b>Architecture</b>	SLC, 10 channel	MLC, 10 channel	
<b>Capacity</b>	32G, 64G (50nm)	80G, 160G (34nm)	
<b>Sustained R/W</b>	R: up to 250MB/s W: up to 170MB/s	R: up to 250MB/s W: up to 70MB/s	
<b>Random R/W (4KB)</b>	R: up to 35K IOPs W: up to 3.3K IOPs	R: up to 35K IOPs W: up to 6.6K IOPs 80GB up to 8.6K IOPs 160GB	R: up to 35K IOPs W: up to 350 IOPs
<b>Lifetime</b>	3 years: random writes 32GB: 1 PB 64GB: 2 PB	5 years: random writes 35TB total up to 20GB/day	3 years: random writes 80GB: 7.5TB total, 6.5GB per day 160GB: 15TB total 13GB per day
<b>Workload</b>	<b>100% Disk Span</b>	<b>10% Disk Span</b>	<b>100% Disk Span</b>
<b>Usage Model</b>	<b>Datacenter</b>	<b>Client</b>	<b>Datacenter</b>

Intel® X25-M Mainstream SATA Solid-State Drives Intel® 18-M Mainstream SATA Solid-State Drives  
Intel® X25-E Extreme SATA Solid-State Drive

# Intel® Solid-State Drive Advantage

*Rules of thumb today*

Primary opportunity

## Storage Usage

### Boot Drive



### DRAM



### Performance



### Capacity



## Intel X25-E Advantage



faster boot  
lower power

< \$/G  
< Watt/G

>> IO latency  
< write endurance

>> read \$/IO  
>> writes \$/IO

>> \$/G

< Watts/G...

>> \$/G...

## Intel X25-M Advantage



faster boot  
lower power

<< \$/G  
< Watt/G

>> IO latency  
<< write endurance

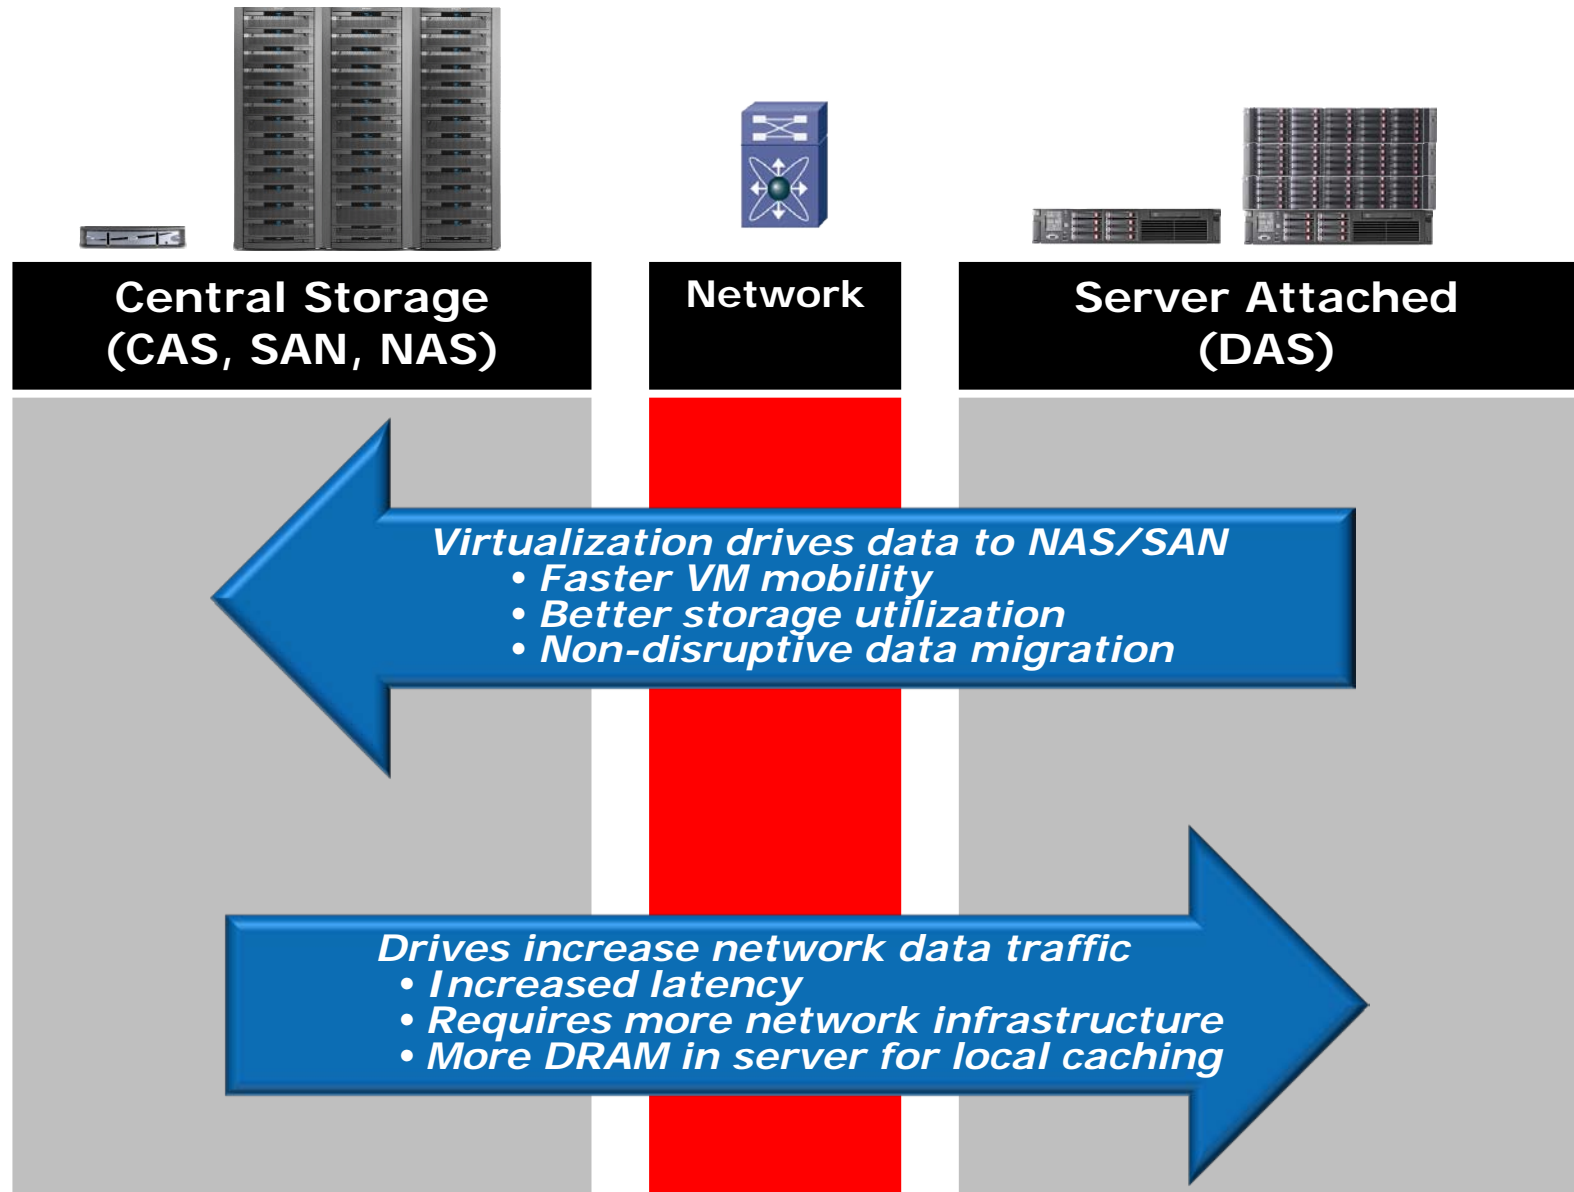
>> read \$/IO  
> write \$/IO

> \$/G  
< write endurance

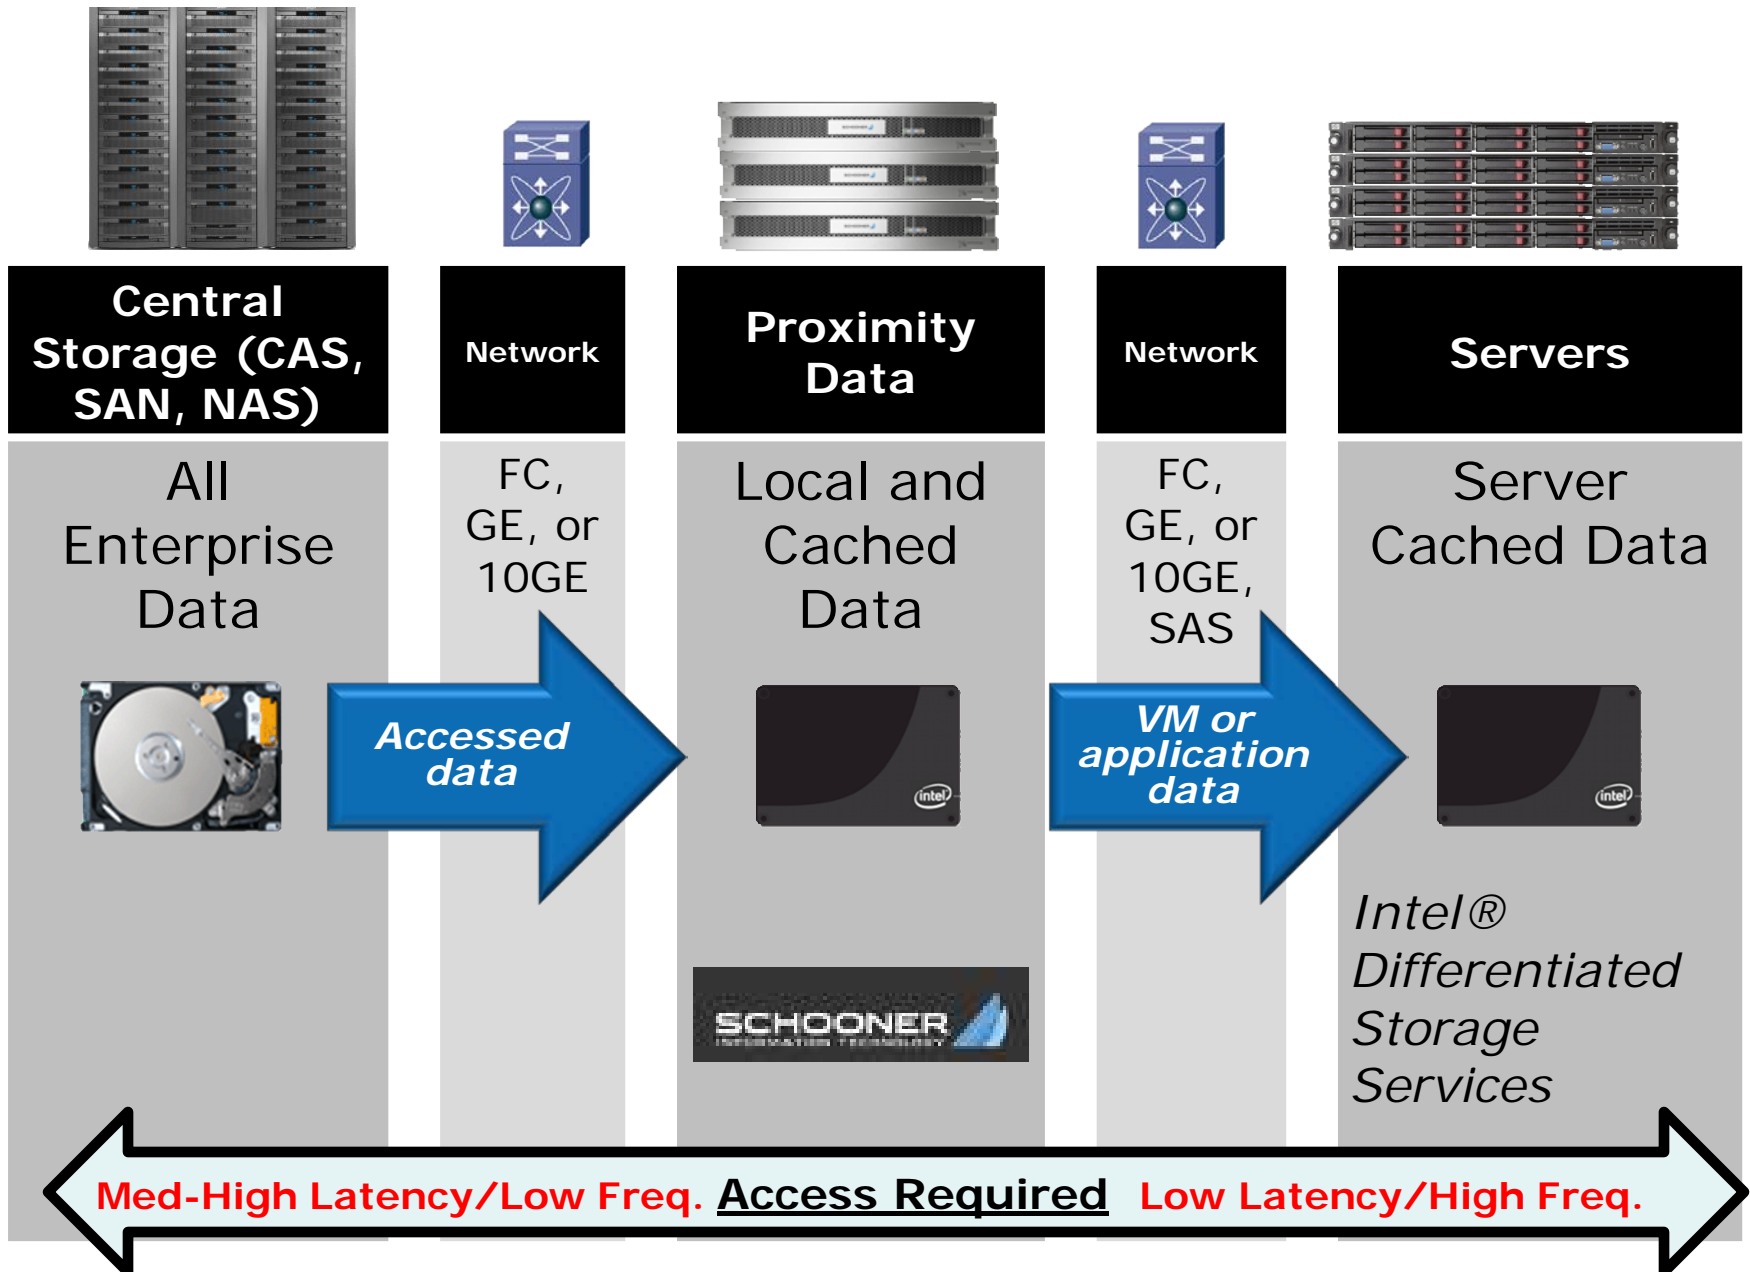
< Watts/G...

>> \$/G...

# Virtualization of Storage...



# Virtualization of Storage...



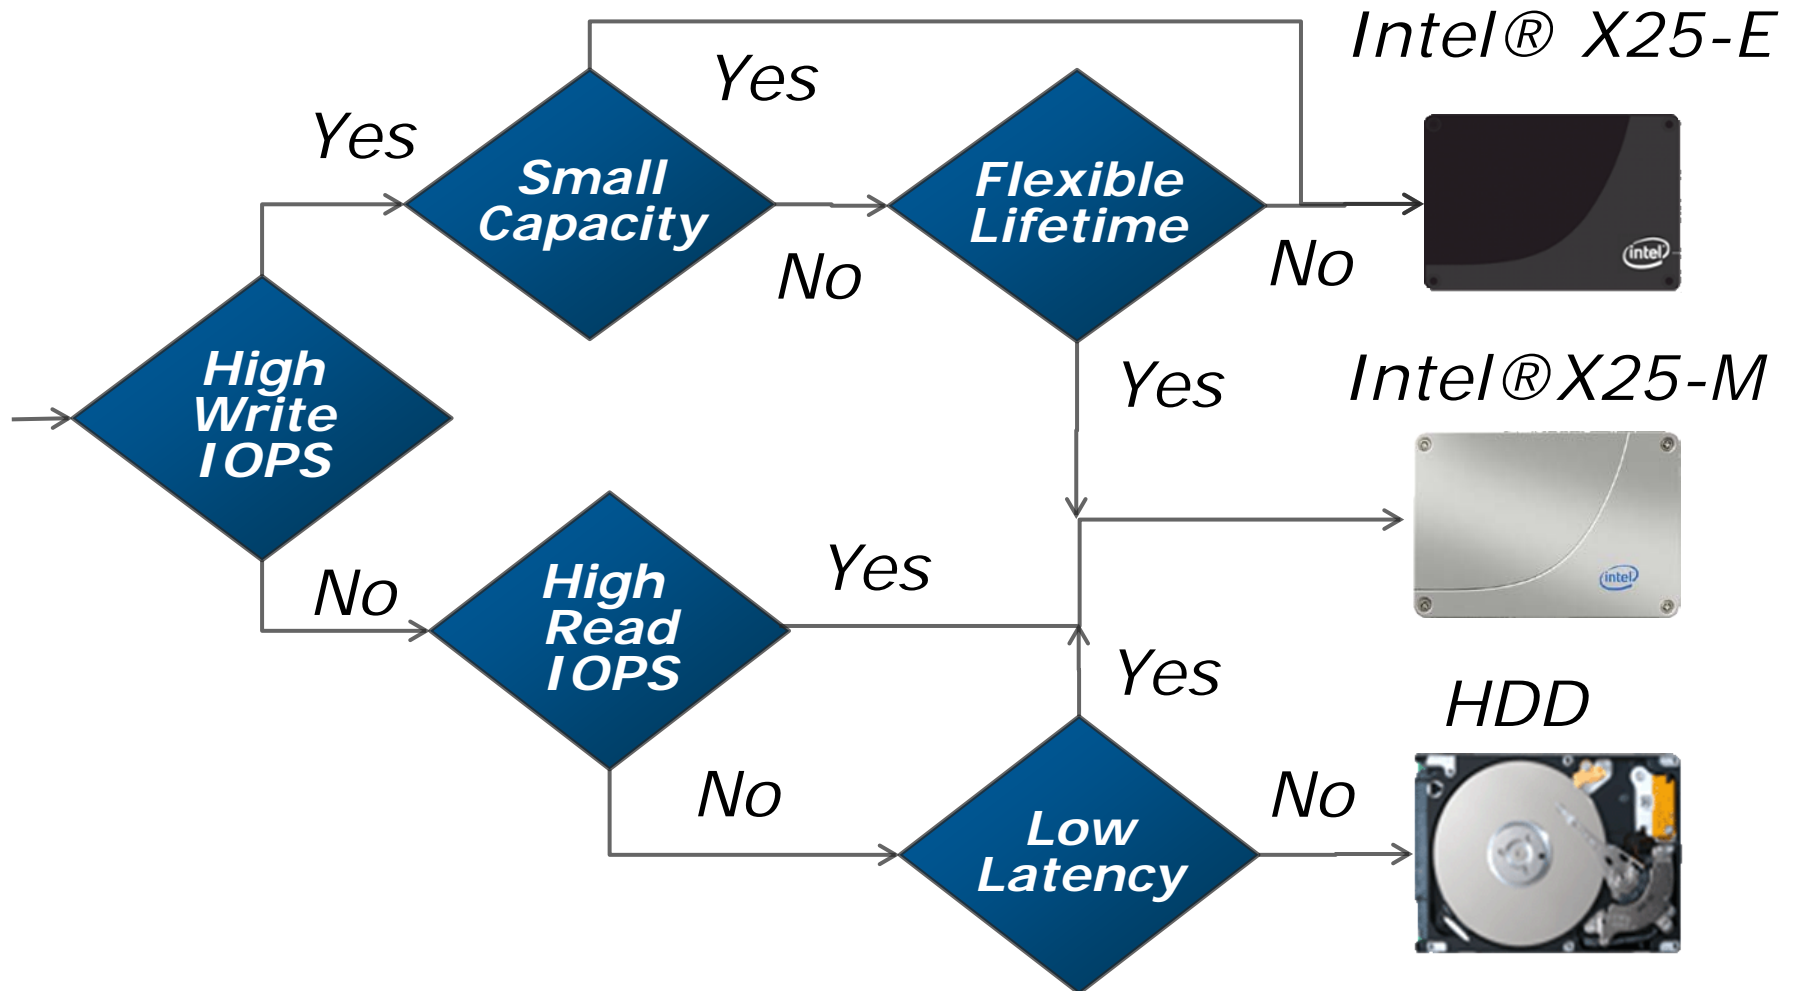
# Agenda

- Data center SSD opportunity
- **ROI in data center applications**
- Building a high performance datacenter configuration
- Q/A



# The ROI Basics - Decision Graph

## Performance Hard Disk Replacement





# Intel SSD Case Study: Level 3 Communications

Michael Peterson, Senior Network Engineer  
September 23, 2009





## Introduction to Level 3

Level 3 customers continually push the boundaries of content delivery:

- Long-form video distribution
- High-Definition streaming (on-demand and live)
- Online games
- Software downloads and patches
- User-generated content
- Large web portals




*Level 3 provides a single-source solution for content delivery that enables us to build out our online presence...to reach fans around the world.*

- Grant Nodine, VP of Technology, NHL

## The Level 3 Content Delivery Network



The Level 3 CDN required an alternative technology solution to provide a higher quality, better performing, lower cost environment for small-file content provider's libraries. The Intel SSD platform solved that requirement.



**Previous Config:**  
*Dual-Core Intel® Xeon®  
Four SATA HDDs  
24 - servers per rack*

**New Config:**  
*Dual-Core Intel® Xeon®  
Three SATA HDDs  
Intel® X25M SATA SSD  
24 - servers per rack*

### The Bottom Line Results



# Agenda

- Data center SSD opportunity
- ROI in data center applications
- **Building a high performance datacenter configuration**
- Q/A




# Building high performance DAS

*Can you get to 1M IOPS on a server?*



**Intel® Solid-State Drives**

**breakthrough performance compared to HDD**



**Storage Controllers (HBA)**

**optimized for HDD performance & latency**



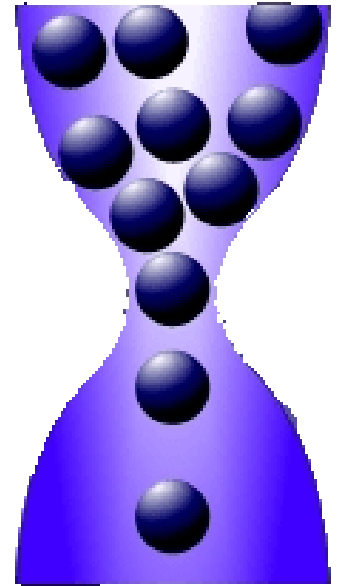
**Disk Enclosure (JBOD)**

**Optimized for HDD performance; expansion & power**



**Intel® Server Processors and Platforms**

**up to 9x performance gain over single-core servers**



# Breakthrough Storage Performance

- **Over 1 million storage I/Os per second<sup>1</sup>**

<sup>1</sup> Depending on application and server configuration

- Demo technology:
  - Intel® Xeon® processor 5500 series
  - Intel® Server Board S5520UR
  - LSI\* 9210-8i host bus adapter
  - Intel® X25-M Mainstream SATA Solid-State Drive - 160G 34nm
  - Supermicro\* 2U SuperChassis 216 SAS/SATA HD BAY


# Dynamo to SSD configuration



19 dynamo workers




Windows 2008 Server



x8  
PCIe G2

5 LSI  
HBAs



SATA II  
(3Gb/s)

28 SSDs



IOMeter 2006-07-27

Windows 2008 Server  
Enterprise Edition

Intel® Server Board  
S5520UR  
Intel® Xeon® 3.2 GHz  
W5580  
12GB DRAM

LSI\* SAS9210-8i  
firmware ver. 2.250  
driver ver. 2.8

Intel® X25-M 160G, 34nm  
config'd as 146G

# The test...

## ***IOMeter Configuration:***

*Transfer Request Size: 4K*

*100% span*

*100% read*

*100% random*

*19 dynamo workers*

*Raw drives*

*32 IOs outstanding per drive*

## ***IOMeter Results:***

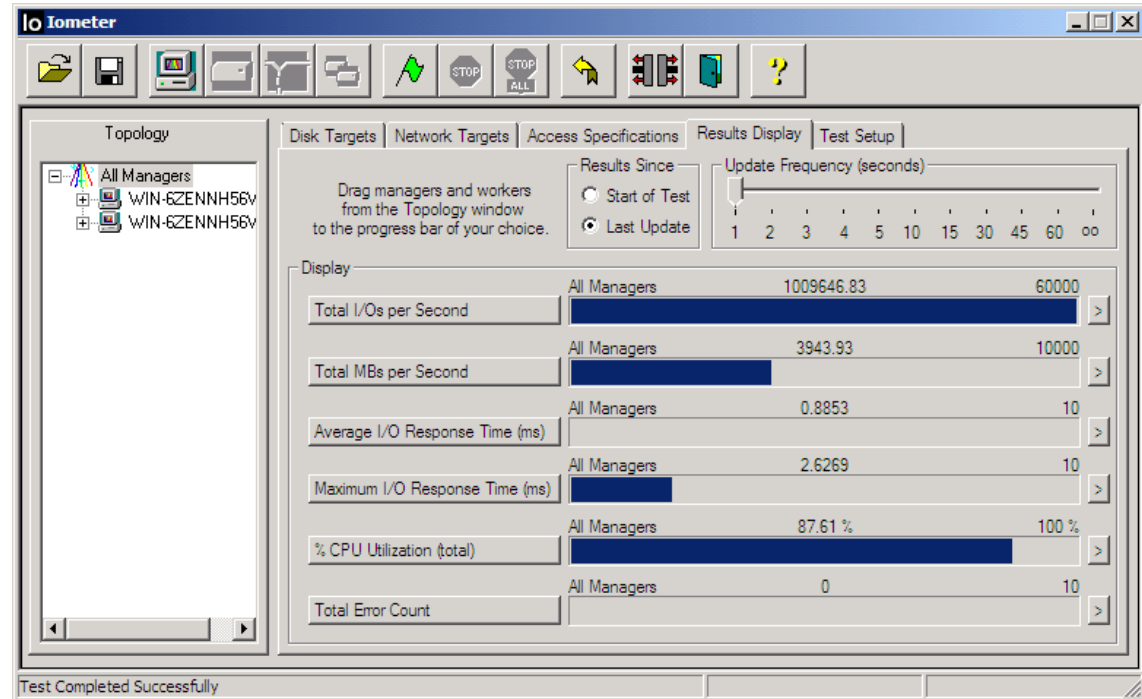
*Test Duration: 1 Hr*

*Average IOPS: 1.08M*

*Average Throughput: 3.9GB/s*

*Average IO Response: .88ms*

*Average CPU Utilization: 85%*



# Building high performance DAS



<b>Intel® X25-M Mainstream SATA Solid State Drive</b>	<b>Over 35K Read IOPS per SSD</b>
<b>LSI * 9210-8i host bus adapter</b>	<b>Over 200K Read IOPS per HBA</b>
<b>Supermicro* 2U SuperChassis 216 SAS/SATA HD BAY</b>	<b>Over 500K Read IOPS per 2U, 24 SFF drive JBOD</b>
<b>Intel® Xeon® processor 5500 Series</b>	<b>Over 1M Read IOPS</b>


*1M IOPS possible now!*

# Summary and Next Steps

- Opportunity for SSDs in datacenter is large!
  - Engage with Intel SATA SSDs today
- Both Intel® X25-E and Intel® X25-M drives apply
  - Understand the Intel® X25-E versus Intel® X25-M performance and reliability tradeoff – MEMS002 and MEMS003
- Breakthrough performance is possible today

# Agenda

- Data center SSD opportunity
- ROI in data center applications
- Building a high performance datacenter configuration
- **Q/A**



# Want More Info on SSDs?

- Attend or download these SSD-related sessions

## Tuesday, Sept 22<sup>nd</sup>

- EBLS001 - Extending Battery Life of Mobile PCs: An Overview

## Wednesday, Sept 23<sup>rd</sup>

- MEMS001 - Designing Solid-State Drives into Data Center Solutions
- MEMS002 - Understanding the Performance of Solid-State Drives in the Enterprise
- MEMS003 - Enterprise Data Integrity and Increasing the Endurance of Your Solid-State Drive
- MEMS004 - Future Solid-State Drive Innovations
- MEMQ002 - Open Q&A for SSD sessions

## Thursday, Sept 24<sup>th</sup>

- MPTS006 - Extreme Notebook Design: Architecting the Most Powerful Mobile Platforms for Gaming & Workstation Applications
- RESS006 - Differentiated Storage Services: Making the Most of Solid-State Drives
- STOS004 - Intel® Modular Server with Intel® Solid-State Drives

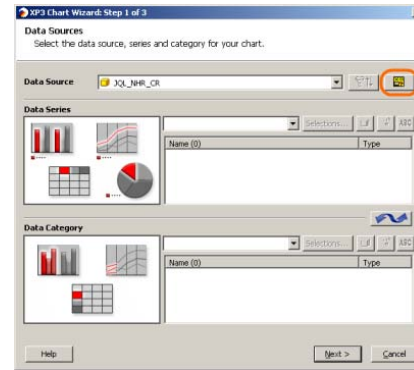
- Visit our Booth #532 on Level 1 of the Tech Showcase
  - SSD vs HDD comparisons, gaming demo and more!
- Visit us online at [www.intel.com/go/ssd](http://www.intel.com/go/ssd)
  - Product briefs, datasheets, whitepapers, videos, technical support

# Business Intelligence:




## Business data from industry

- Digital Media
- Energy and Utilities
- Financial Services
- Health and Life Sciences
- Retail and CPG
- Telecommunications




*Speed of result is differentiation*



*Kognitio WX2 scales using COTS technology*

# Business Intelligence:

*Financial services example*



HP Proliant\* BL460c G6  
Quad-Core  
Intel® Xeon® processor  
55xx  
48G DRAM

**HDD Config:**  
2 - Seagate\* 146G  
15K SAS HDDs

**SDD Config:**  
2 - Intel® X25-M  
160G  
SATA SSDs  
(config as 146G)

## The Bottom Line Benefits

Time to complete  
~6x  
Decrease

X25-E  
Solution Cost  
~1.5x  
more

X25-M  
Solution Cost  
~1.2x  
more

For notes and disclaimers, see legal information slide at end of this presentation.

# 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.
- Nehalem, Westmere, Boxboro, Postville, and other 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, Xeon, 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.
- Copyright © 2009 Intel Corporation.

# Risk Factors

The above statements and any others in this document that refer to plans and expectations for the third quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. 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 corporation's expectations. Ongoing uncertainty in global economic conditions pose a risk to the overall economy as consumers and businesses may defer purchases in response to tighter credit and negative financial news, which could negatively affect product demand and other related matters. Consequently, 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 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. Additionally, Intel is in the process of transitioning to its next generation of products on 32nm process technology, and there could be execution issues associated with these changes, including product defects and errata along with lower than anticipated manufacturing yields. 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; and Intel's ability to respond quickly to technological developments and to incorporate new features into its products. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; capacity utilization; start-up costs, including costs associated with the new 32nm process technology; variations in inventory valuation, including variations related to the timing of qualifying products for sale; excess or obsolete inventory; product mix and pricing; 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. 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. The current financial stress affecting the banking system and financial markets and the going concern threats to investment banks and other financial institutions have resulted in a tightening in the credit markets, a reduced level of liquidity in many financial markets, and heightened volatility in fixed income, credit and equity markets. There could be a number of follow-on effects from the credit crisis on Intel's business, including insolvency of key suppliers resulting in product delays; inability of customers to obtain credit to finance purchases of our products and/or customer insolvencies; counterparty failures negatively impacting our treasury operations; increased expense or inability to obtain short-term financing of Intel's operations from the issuance of commercial paper; and increased impairments from the inability of investee companies to obtain financing. The majority of our 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 our 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 impacted 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. 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 risk 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 27, 2009.