Server Based Storage

Faster…
Economical…
Reliable…

Steve Mattos
Strategic Program Manager
steve.mattos@intel.com

Flash Memory Summit 2012
Santa Clara Convention Center
August 21, 2012

Legal Notices and 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.

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.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. Test Configuration: 1. Desktop system with Intel® DX58SO motherboard, Intel® Core™ i7-975 @ 3.33GHz CPU, 3X2 GB DDR3-1333 memory, GeForce* GTS 250 - 512MB graphics card, Intel(R) Matrix Storage Manager 8.9 storage driver and Microsoft* Windows* 7 Ultimate 64 bit O/S. 2. Intel® P67 desktop system+ Intel® Core™ i7 CPU @ 3.40GHz + 4GB RAM+ Microsoft Windows* 7 Enterprise. For more information go to http://www.intel.com/performance

Intel, Intel Inside, the Intel logo, Centrino, Intel Core, Intel Atom, Pentium and UltraBook are trademarks of Intel Corporation in the United States and other countries.

Material in this presentation is intended as product positioning and not approved end user messaging.

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

Agenda

- Datacenter Growth and Management Decisions
- Server Based Storage
- Reference Materials
- Intel at the Flash Memory Summit

MEETING THE DEMAND FOR INTELLIGENT STORAGE

Cloud Computing Driven by:

More users, more devices, more data, more storage, more traffic...

55%

Growth in Unstructured data (e-mail, video, images, social media etc.)

75%

Growth in Public Cloud

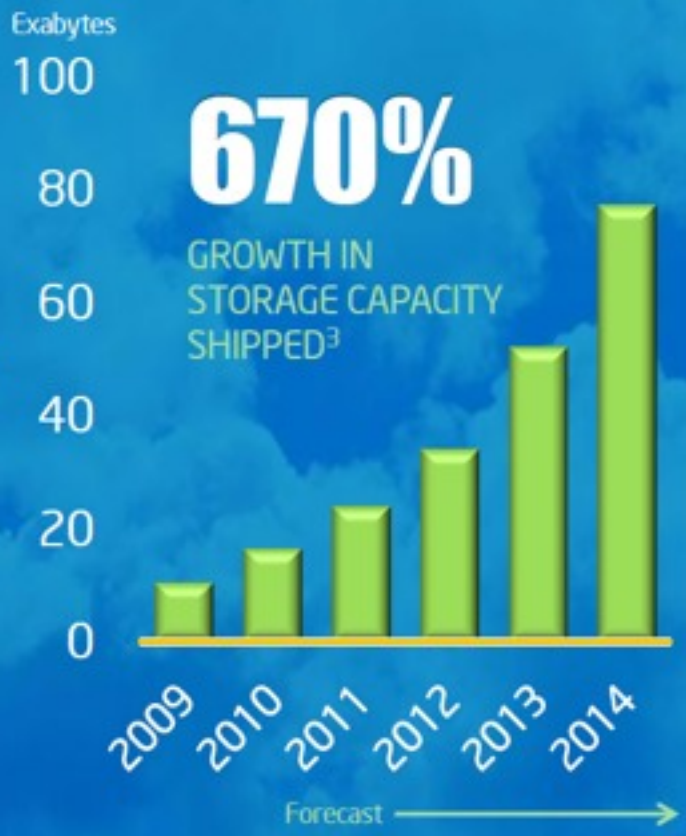

670%

GROWTH IN STORAGE CAPACITY SHIPPED[3]

>1000

EXABYTES OF TRAFFIC[4]

1. IDC "The Internet Reaches Late Adolescence" Dec 2009, extrapolation by Intel for 2015
2. ECG "Worldwide Device Estimates Year 2020 - Intel One Smart Network Work" forecast
3. IDC
4. Source: http://www.cisco.com/assets/cdc_content_elements/networking_solutions/service_provider/visual_networking_ip_traffic_chart.html, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html extrapolated to 2015

Datacenter Growth and Management Decisions

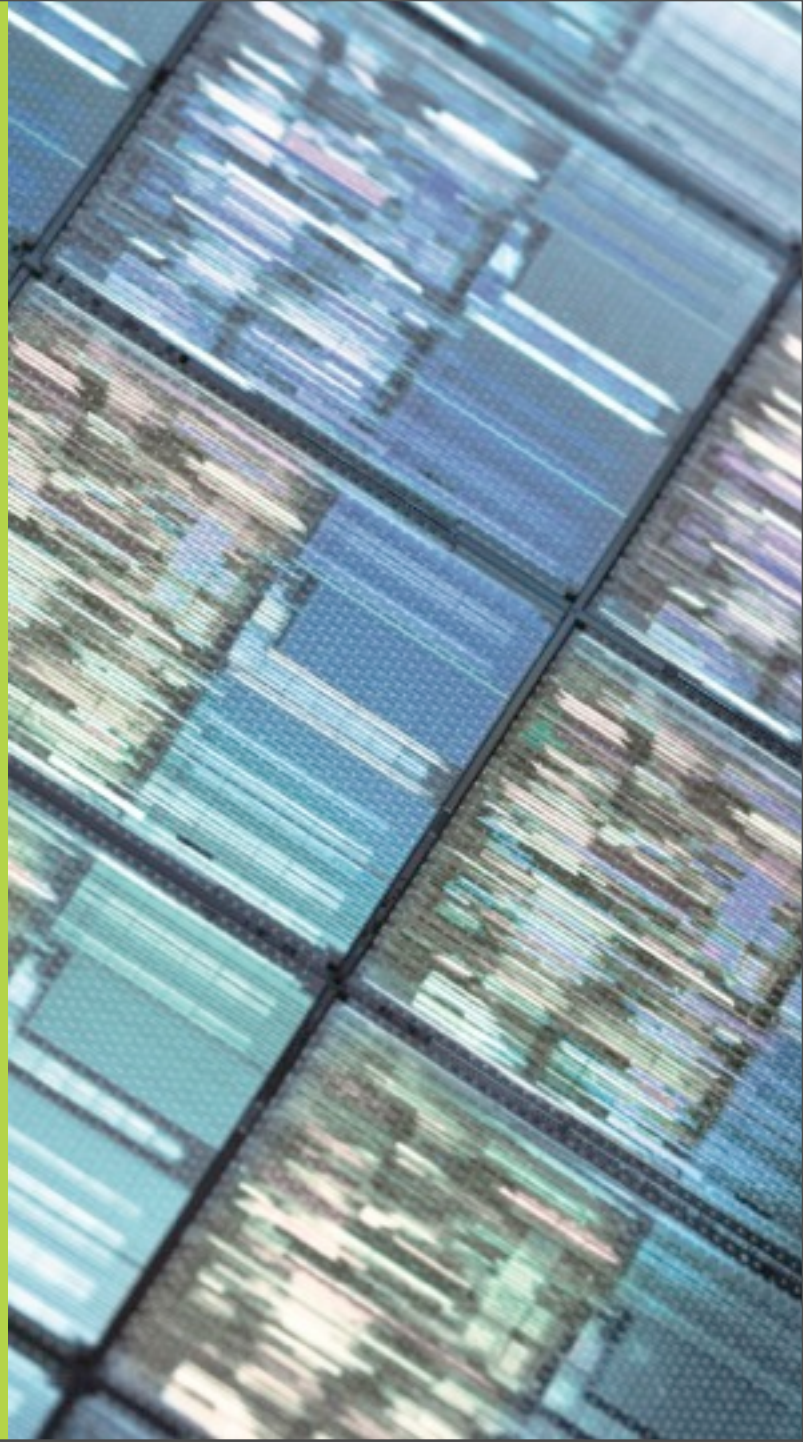

Agenda

- **Datacenter Growth and Management Decisions**
- Server Based Storage
- Reference Materials
- Intel at the Flash Memory Summit

Datacenter Storage Escalation

By 2015…

More Users

1B more netizens[1]

More Devices

15B connected devices[2]

More Data

1,000 Exabytes traffic[3]

Evolving Storage Paradigm

1. IDC "Server Workloads Forecast" 2009. 2.IDC "The Internet Reaches Late Adolescence" Dec 2009, extrapolation by Intel for 2015
2. ECG "Worldwide Device Estimates Year 2020 - Intel One Smart Network Work" forecast
3. Source: http://www.cisco.com/assets/cdc_content_elements/networking_solutions/service_provider/visual_networking_ip_traffic_chart.html extrapolated to 2015

Current Solution Methodology

- Scale out storage with HDD's and systems
  - Expands storage with additional HDD's in existing systems…
  - Alternative: Add additional storage subsystems
- Engineer ultimate component reliability
  - Use SAS drives for dual port redundancy and ultimate uptime reliability
- HDD over-provisioning
  - Increase HDD performance by trading capacity for rotational performance
- Add improved networking cards
  - Increasing bandwidth and network level performance

Summary of Current Solutions

SOLUTION	PROBLEMS
Scale out HHD's and systems	Addresses capacity, but adds power, heat
Ultimate component reliability	Burdens reliability cost on each server
HDD over provisioning	Underutilized drives and gobbles up space
Improved networking cards	Individual cards on systems do not improve
New and faster servers	Commonly used as first approach to improve

Smart choices must be made to remain competitive

Other Issues to Consider

- Are server CPU's fully utilized?
- Where are the current I/O performance bottlenecks?
  - Storage, Networking, Memory, System CPU's?
- What is the appropriate socket and core configuration?
  - And the associated impact of software licensing?
- How much system memory is optimal?
- What is the workload and usage assumption?
- MTBF degradation due to increased infrastructure?
  - Not linear

Easy to overlook key issues that impact cost and performance

Server Based Storage

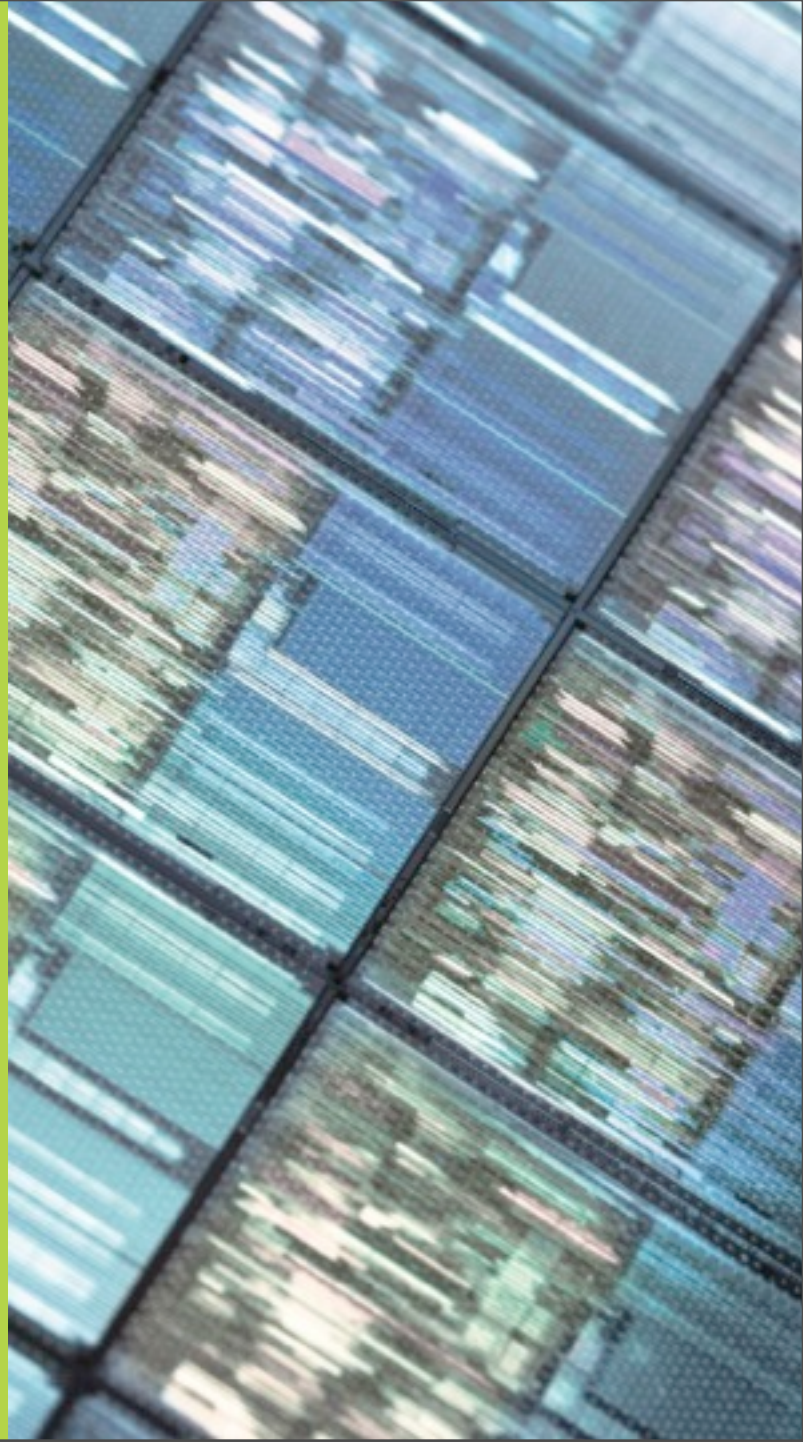

Agenda

- Datacenter Growth and Management Decisions
- **Server Based Storage**
- Reference Materials
- Intel at the Flash Memory Summit

Server Based Storage Advantages

- Industry moving to place primary storage inside servers
  - Reduces cost
  - Saves space
  - Improves performance and rebalances system I/O
  - Reduces thermal load & power demands
  - Reduces routing & switching overhead
  - Eliminates external storage subsystems
  - Minimizes storage management software diffusion
  - Reduces points of failure while increasing MTBF
- Takes advantage of existing hardware
  - SSD's and faster NIC's are enabling technologies
- Clustered servers used to provide fault tolerance
  - Different than Scale Out Storage

Another Way to Architect A Solution? Local Storage with SSDs

REQUIREMENTS	SOLUTION
Need more storage capacity	Add SSD's and HDD's to server enclosures to
Increased performance	Add SSD's to existing HDD configuration to
Limited space, power and air	Add storage to existing servers to increase
Need fastest possible storage	Add SSD's to servers and use as primary storage
Simplify storage management software	Rely upon emerging converged enterprise software and Cloud

Server Based Storage

Achieves Greater Business Value	
Performance & efficiency	Capacity
Availability	Simplified management

Storage Components
- Enclosure
- Backplane
- SATA/SAS Cables
- IOC/HBA
- Solid State Storage
- Hard Disks

Server Components
- Server board
- Server chipset
- Server processor
- Server memory
- Network card
- Power supply

Metadata Server

Application Server

Storage Server

Converged storage servers deliver a cost-effective storage platform

Solid State Advantage

- Server-based storage relies upon Solid State storage for hierarchical caching and performance acceleration
  - SSDs may also be configured for primary storage
  - Conventional HDDs provide lower-performance long term storage
- Capacity ratio of 1:10 is common
- Conventional 2.5” form factors offer the best removability
  - Hot plugability subject to hardware and software enabling and certification

SSD’s in Servers Offer Untapped Potential for Caching and Mainline Storage

SSD Datacenter Placement

Key Take-Aways

- SSDs displacing **15k HDDs**
- 2.5" HDDs displaced 3.5"

Why?

- Reliability
- Performance
- Power savings
- Density

SSDs

- Reliability beyond mechanical hard drives
- Performance AND endurance for rigorous datacenter workloads
- Power savings multiplied beyond SSD to

Better, Faster, Cheaper
and more Reliable

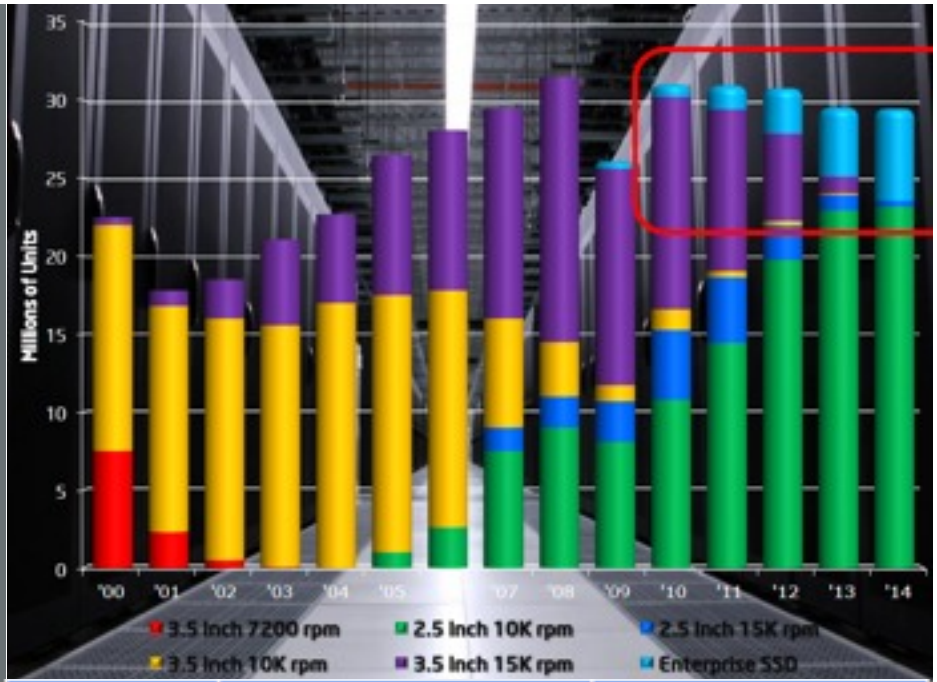

Segment	Intel Vision	Intel Value
Datacenter	*Every high-performance HDD will be replaced by an SSD*	•Quality & Reliability •Quality of Service •Validation & Support
Business Client	*Every Corporate IT Client will have an SSD*	•Quality & Reliability •TCO •Tools & Support
Enthusiast Upgrade Embedded	*The best system Upgrade is an SSD*	•Channel Reach •Quality/Brand •Performance

Cluster Configuration

- At least 3 to 4 servers required to support fault tolerance
  - Sometimes referred to as Network RAID
- Data is sent to JBOD drives based upon provisioning and data workload profiles
  - Ex: High bandwidth, low latency, directory data, archival, log files
- Drives within each server are not required to be the same capacity
- Data duplicated across server based upon real-time heuristic analysis and modeling

Data Mapping

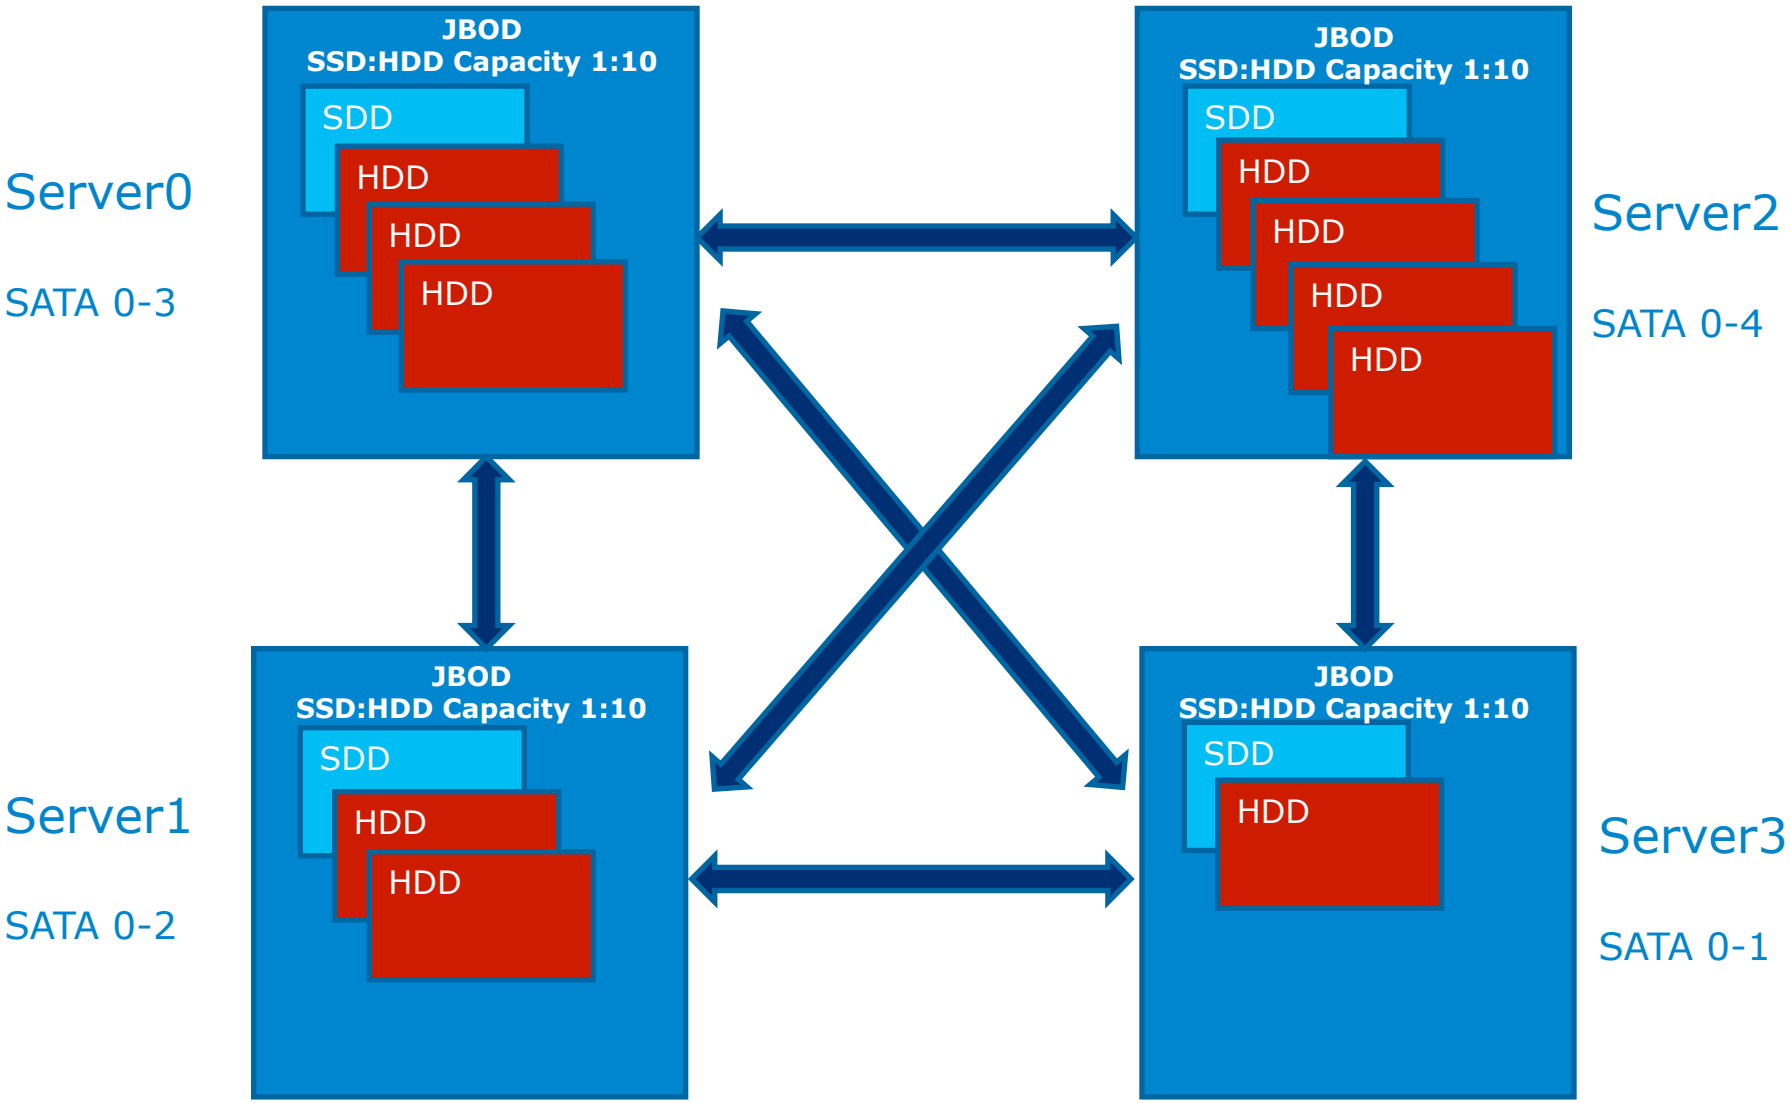

PCIe Add-In Option

- Add-in PCIe solutions may be used as conventional SSD form factor alternative
- Enables faster access to data
- Compact vs. conventional HDD + HBA
- Boosts storage performance vs. HDD's

Form factor must be consistent with industry standards

Network Requirements

- At least one 10GbE port required to support bandwidth
- Dual 10GbE ports recommended for enhanced reliability
  - 1GbE does not offer enough bandwidth
- Used to transfer data between server cluster nodes
  - Also serves conventional LAN and WAN traffic

Simplify with Ethernet 10GbE

GbE Server Connections

10GbE Server Connections

Up to 45% Reduction in power per rack

Up to 80% Reduction in cables and switch ports

Up to 15% Reduction in infra-structure costs

Up to 2x Improved bandwidth per server

Source: Intel 10GbE ROI Calculator. This ROI calculator is a cost comparison for a highly virtualized solution, using multiple 1GbE connections versus a dual port 10GbE implementation.
http://www.event-management-online.de/LAD/calculator.aspx

Solution Strategy

- Server based storage is emerging as an alternative to external storage subsystems
- Appropriate for many datacenter production environments
  - Cloud, virtualization, virtual desktop, decision support, etc.
- Not specifically designed to support big data usage models
- Consolidates hardware and software
- Reduces all major cost factors

Architectural Guidance

- Configurations of HDDs & SDDs balance cost and performance
  - SDD to HDD ratio contingent upon specific solution and workload
- Emerging PCIe add-in form factors can be used as alternative to conventional SSD form factors
- SAS ports no longer required to insure reliability
  - Fault tolerance via server clusters
  - SATA is acceptable and potentially optimal
- Embedded SATA ports offer lowest cost
- Add-in SATA RAID configured for JBOD
  - Software management console useful to monitor drive provisioning, health, status and performance
- Consumer-class drives may be acceptable for this architecture
  - But…requires understanding of workload…read to write ratios and performance requirements are key dependencies

Solid State Is Mainstream

Worldwide SSD Unit Sales[1]

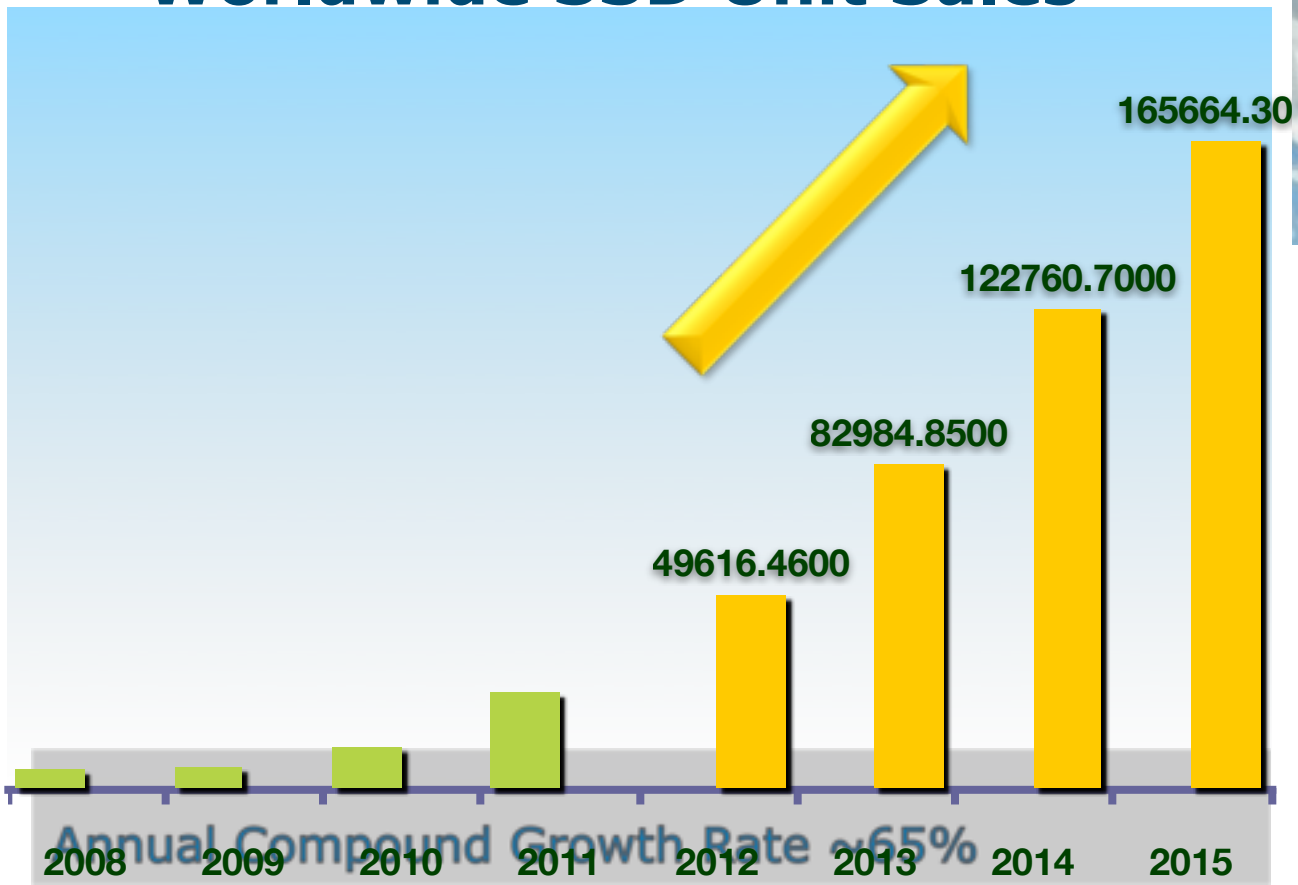

SSDs are Positioned for Strong Growth

Server Based Storage Summary

- Existing solutions for Data Center growth may create more problems than they solve
- Multiple approaches can be used to address growing storage requirements while maximizing performance and reducing costs etc.
- Moving SSDs inside servers maximizes performance and increases reliability while reducing power and thermal load

SSD momentum in the Enterprise will accelerate as strategies are updated

Agenda

- Datacenter Growth and Management Decisions
- Server Based Storage
- **Reference Materials**
- Intel at the Flash Memory Summit

Reference Materials

- Intel Solid State Technology
  - www.intel.com/go/ssd
- Storage Networking Industry Association
  - www.snia.org
  - http://snia.org/forums/sssi/programs/twg
- Differentiated Storage Services
  - www.intel.com
  - Search for "Differentiated Storage Services"

Agenda

- Datacenter Growth and Management Decisions
- Server Based Storage
- Reference Materials
- **Intel at the Flash Memory Summit**

Tuesday, Aug. 21

- **How SSDs Fit in Different Data Center Applications, Tutorial A-11: Flash in Data Centers (Enterprise Storage Track) 8:30 a.m. to 11:25 a.m.**
- Tahmid Rahman, Technical Marketing Engineer

- **Server-Based Storage: Faster, Economical, Reliable, Tutorial C-11: Enterprise SSDs (SSDs Track) 8:30 a.m. to 9:50 a.m.**
- Steve Mattos, Strategic Program Manager Storage Solutions

- **Merits and Methods of IO Traced Based Performance Benchmarking of SSDs, Open Session 101-B: Mobile Applications (Applications Track) 8:30 a.m. to 9:50 a.m.**
- Harry Pon, NAND Product Development

- **The Transition to PCIe for Client SSDs, Open Session 102-C: Standards (SSDs Track) 10:10 a.m. to 11:25 a.m.**
- Amber Huffman, Senior Principal Engineer

- **Write Atomicity and NVM Drive Design, Tutorial B-11: Flash Memory-Based Architectures (Architectures Track) 10:10 a.m. to 11:25 a.m.**
- Andy Rudoff, Enterprise Storage Architect
-
- **Verification and Management of Endurance in NAND SSDs, Tutorial C-12: SSD Technology (SSDs Track) 2:10 p.m. to 4:45 p.m.**
- Venkatesh Vasudevan, Director Quality and Reliability Engineering

- **Data Integrity on 20nm NAND SSDs, Tutorial C-12: SSD Technology (SSDs Track) 2:10 p.m. to 4:45 p.m.**
- Robert Frickey, Product Development Engineer

Wed., Aug. 22

- **Which Way are We Headed?, Open Session 201-A: Future Storage Interfaces (Interfaces Track) 8:30 a.m. to 9:40 a.m.**
- Jim Pappas, Director Initiative Marketing

- **Data Recovery Survival Tips and Realities, Open Session 201-B: Data Recovery of SSDs (SSDs Track) 8:30 a.m. to 9:40 a.m.**
- David Blunden, Applications Engineer

- **Thunderbolt, Open Session 204-D (Interfaces Track) 4:30 p.m. to 5:30 p.m.**
- Brett Branch, Software and Ecosystem Enabling

Thursday, Aug. 23

- **Industry Standards for PCIe SSD Storage, Session 301-A: PCIe Storage-1 (PCIe Storage Track) 8:30 a.m. to 9:40 a.m.**
- Jim Pappas, Director Initiative Marketing

- **Exploitation of Rber Diversity over Dies to Improve ECC Performance in NAND Flash Drive,**
  Session #301-B: SSD Technology (SSDs Track) 8:30 a.m. to 9:40 a.m.
- Ravi Motwani, ECC/DSP Architect
-

- **Intel Ultrabook Responsiveness and NVM Caching, Tutorial A-31 (Enterprise Storage Track) 8:30 a.m. to 10:50 a.m.**
- Dale Juenemann, Storage Architect

- **Solid State Drives - From Disruptive to the New Normal, Open KEYNOTE 10: 2:00 p.m. to 2:30 p.m.**
- Robert Crooke, Vice President/General Manager – Non-Volatile Memory Solutions Group

- **PCIe SSD Roundtable, Session 303-B: (PCIe Storage Track) 3:10 p.m. to 4:25 p.m.**
- Mark Meyers, Server Platform Architect

- **Top Ten Things You Need to Know about Flash Memory Today, Open Session: 304-A:**
  Closing Panel 4:40 p.m. to 6:00 p.m.
- Knut Grimsrud, Director Storage Architecture/Intel Fellow

Questions?

Thank You


intel

Abstract

- Customers are demanding an alternative to expensive and slow external storage subsystems for workloads that maximize the performance of server hardware.

- In response, ecosystem solution providers are developing solutions that reduce the need for external storage subsystems while improving performance and throughput as well as reducing overall power consumption. These new solutions integrate hybrid drive configurations inside the server enclosure to offer the best characteristics of solid state drive (SSD) performance balanced against the low cost and capacity of conventional hard drives.

- This session outlines the detail behind this trend toward internal SSD storage, provides tips for system integration and offers guidance for building optimized server solutions. Tradeoffs of Enterprise versus Consumer class SSD's will be discussed along with the latest trend in PCIe SSD's.

Flash Memory Summit
Event Logistics

- Date: Tuesday, August 21.
- When: 8:30 to 9:50 or 10:10 to 11:25 (30 minute tutorial session)
- Where: Santa Clara Convention Center

Tutorial C-11: Enterprise SSDs (SSDs Track)
Organizers: Tom Coughlin, President Coughlin Associates and Lakshmi Mandyam, Director Enterprise Segment Marketing, ARM, Scott Shadley, Senior Manager Product Marketing, Micron Technology
Chairperson: Pallab Chatterjee, CTO, SiliconMap

Instructors:

It Takes Guts to be Great
Sean Stead, SSD Technical Marketing, STEC

The Demise of SLC in the Enterprise
Bernie Rub, VP / CTO, Smart Storage Systems

MLC Media Discussion
Scott Shadley, Senior Manager Product Marketing, Micron Technology

The SSD Endurance Race: Who's Got the Write Stuff?
Ulrich Hansen, Director Market Development, HGST

TBD
Frank Berry, President, IT Brand Pulse

Server-Based STorage: Faster, Economical, Reliable
Steve Mattos, Strategic Program Manager - Storage Solutions, Intel

Tutorial Description:
SSD adoption keeps rising throughout the enterprise market. Hence storage designers and engineers need to understand how NAND flash and controllers are chosen, used, and optimized to select the right products for their applications. In particular, they must know about the properties of NAND flash in general and the differences between SLC and the emerging MLC technologies.

- Feature sets needed in enterprise SSDs to enable robust designs
- Usability of MLC flash despite its lower endurance compared to SLC
- Use of SDs to extend storage
- How to measure, determine, and work with limited SSD endurance

Program At-A-Glance

Tuesday, August 21st

	Enterprise Storage	Architectures	SSDs	Hardware	Applications	SSDs
8:00-8:30am	Registration & Continental Breakfast					
8:30-9:50am	Tutorial A-11 Flash in Data Centers	Tutorial B-11 Flash-Memory Based Architectures: A technical Discussion Part 1	Tutorial C-11 Enterprise SSDs	Session 101-A 3D Flash: The Next Dimension	Session 101-B Mobile Applications OPEN	
9:50-10:10am	Break					
10:10-11:25am	Tutorial A-11 Flash in Data Centers (cont.)	Tutorial B-11 Flash Memory-Based Architectures: A Technical Discussion Part 1 (cont.)	Tutorail C-11 Enterprise SSDs (cont.)	Session 102-A Flash Technology Trends	Session 102-B Consumer Applications OPEN	Session 102-C Standards OPEN
11:30am-Noon	OPEN - Keynote 1 SSDs: Enabling the Next Wave of Growth in the PC Industry Kevin Conley, Vice President and General Manager - Client Storage Solutions BUSanDisk					
Noon-1:00pm	Lunch					