Intel® I/O Acceleration Technology (Intel® IOAT) Overview

Lily Deng
Enterprise Marketing Operation - China
Digital Enterprise Group

Agenda

- What is Intel IOAT?
- Next Generation – IOAT2
- Real world Application Benefits with I/OAT

Network I/O Growing Rapidly

- More network data flowing in/ out of servers
  - Demand stretching resources

More LAN Traffic

IDC: iSCSI storage systems and connected servers … can drive a 600% packet load increase on Ethernet.**

Virtualization: More Apps / Server

IDC Survey of Senior IT:
"22% virtualized today with 45% of all planned server deployments virtualization candidates"*

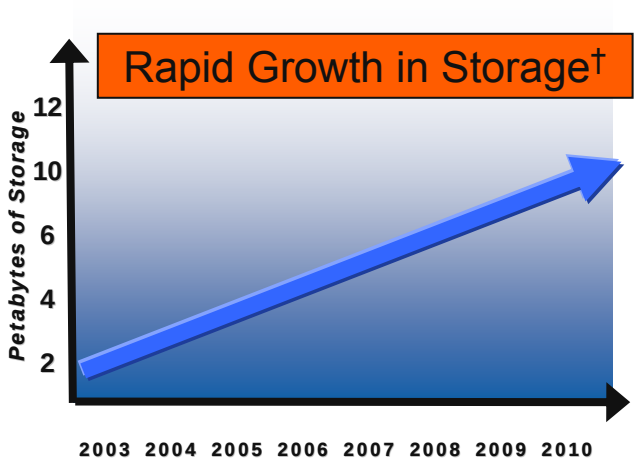

*Source: IDC WW Server Platform Scorecard November 2005

**Source: 2005 Storage I/O Traffic Set to Dominate Ethernet LAN Packets

†Source: IDC Estimates

Intel® I/O Acceleration Technology

Network Acceleration on the Intel® 5100 Series Server Platforms

Network I/O inefficiencies
✓ CPU makes multiple "forced" trips to memory
✓ System Overheads from Device ↔ App signaling
✓System Resource Utilization

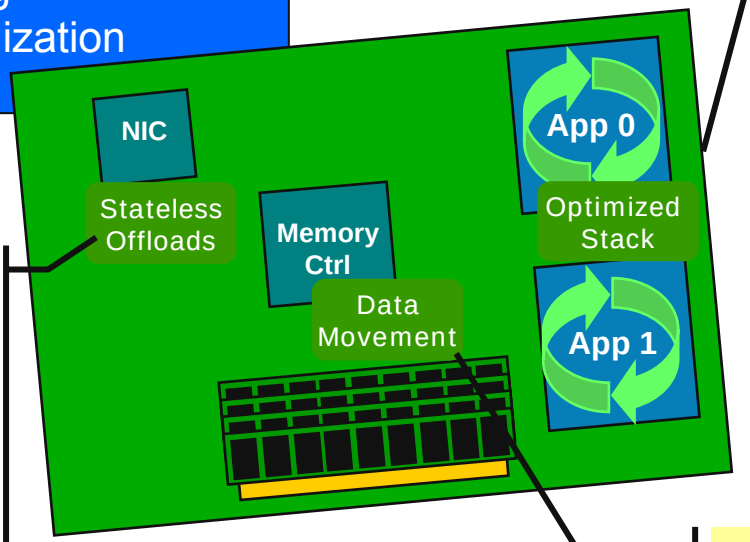

CPU
- SW interface enhancements
- SW pre-fetches
- Affinitize flows
- Native OS support

Network
- Segmentation & checksum offload
- Optimized Interrupt Moderation
- Multiple Queues & Flow Affinitization
- Split Header/Payload

Chipset
- DMA hardware in Memory Controller Hub offloads data movement

Intel I/OAT Improves Server Network Efficiency

Agenda

- What is Intel IOAT?
- Next Generation – IOAT2
- Real world Application Benefits with I/OAT

Intel® I/OAT Technical Overview

I/OAT1: Bensley (now)

I/OAT2: '07 Caneland, Stoakley, Cranberry Lake

Network I/O inefficiencies addressed:

- ✓ CPU makes multiple "forced" trips to memory
- ✓ System overheads in Device ↔ App signaling
- ✓ System resource utilization

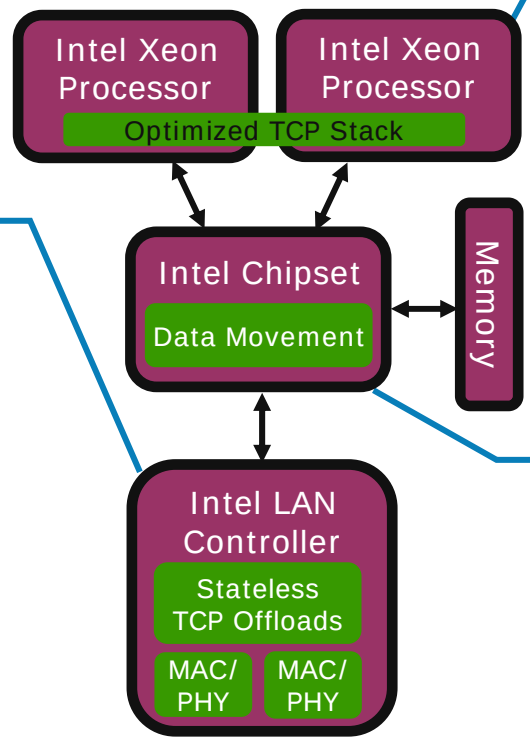

I/OAT1:
CPU (host)
- •SW interface enhancements
- •SW pre-fetches to CPU
- •Affinitize data flows
- •Native OS support

I/OAT2: Direct Cache Access (DCA)

I/OAT1:
LAN Controller (NIC)
- •Stateless Offloads
- •TCP Segmentation & checksum offload
- •Multiple Queues & Flow Affinitization
- •Split Header/Payload

I/OAT2: Direct Cache Access (DCA), MSI-X, Low Latency Interrupts, Header Splitting/ Replication

I/OAT1:
Chipset
- •Intel® QuickData technology in Memory Controller Hub offloads data movement & memory copies

I/OAT2: Direct Cache Access (DCA), MSI-X

IOAT Roadmap

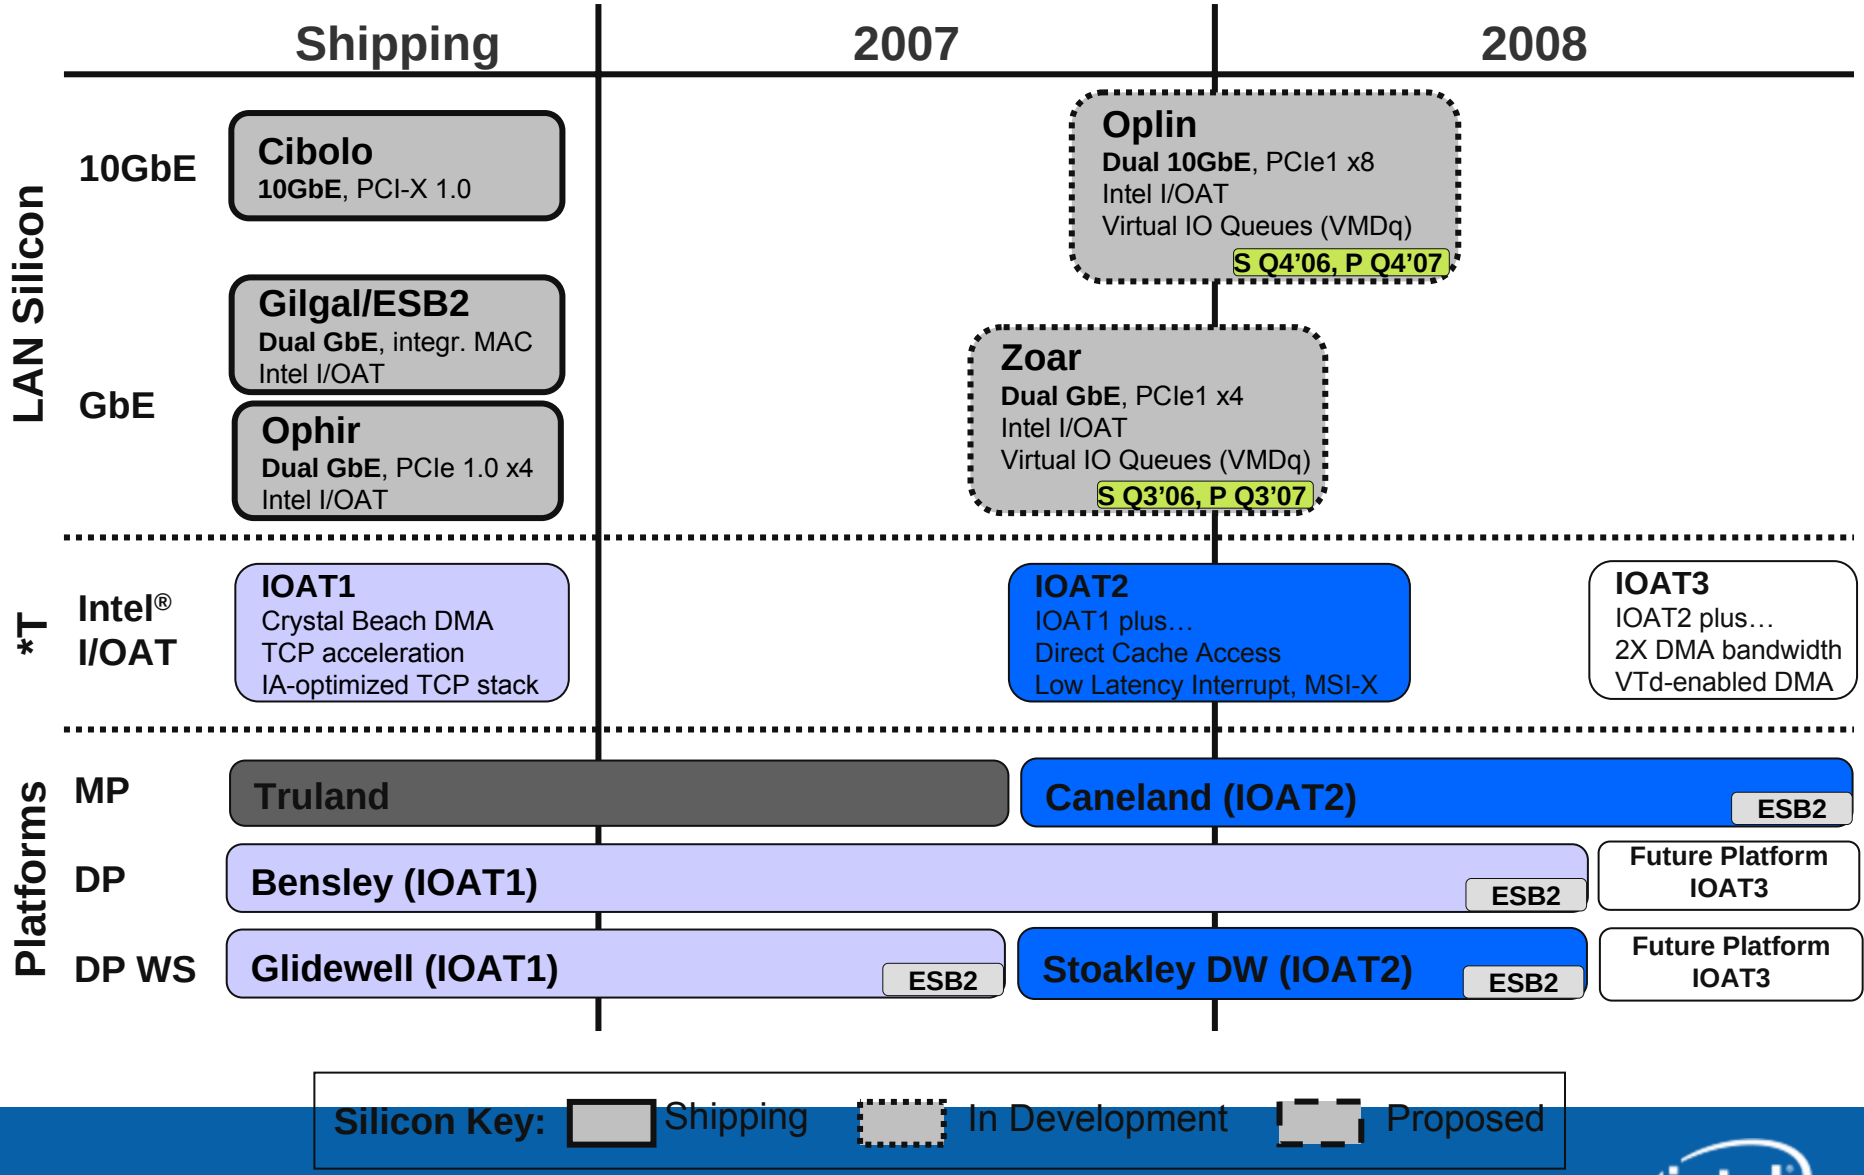

Intel® I/OAT Generations

Feature	IOAT 1 (Bensley)	IOAT2 (Stoakley Caneland)	IOAT3 (Next Gen Platform)
Intel QuickData Tech (Data Movement Engine) BW	2GB/s	2GB/s	4GB/s
Number of DMA Channels	4	4	8
LAN stateless offloads (Header/data split, Receive Side Scaling, TX/RX checksums, TCP segmentation)	✓	✓	✓
Message Signaled Interrupts	MSI	MSI-X	MSI-X
Direct Cache Access		✓	✓
Low Latency Interrupt		✓	✓
Optimized Header-Splitting / Replication		✓	✓
Multi-VM Direct Assignment of Data Movement Engine			✓
Required LAN Si	**IOAT1**	**IOAT2**	**IOAT3**
Gilgal Dual GbE PHY	✓		
Zoar Dual GbE MAC/PHY	✓	✓	✓
Kawela(Adoram) Dual GbE MAC/PHY	✓	✓	✓
Oplin Dual 10 GbE MAC	✓	✓	✓
Niantic(Hadar) Dual 10 GbE MAC	✓	✓	✓ w/RSC

DCA Usage Models on Stoakley Platform

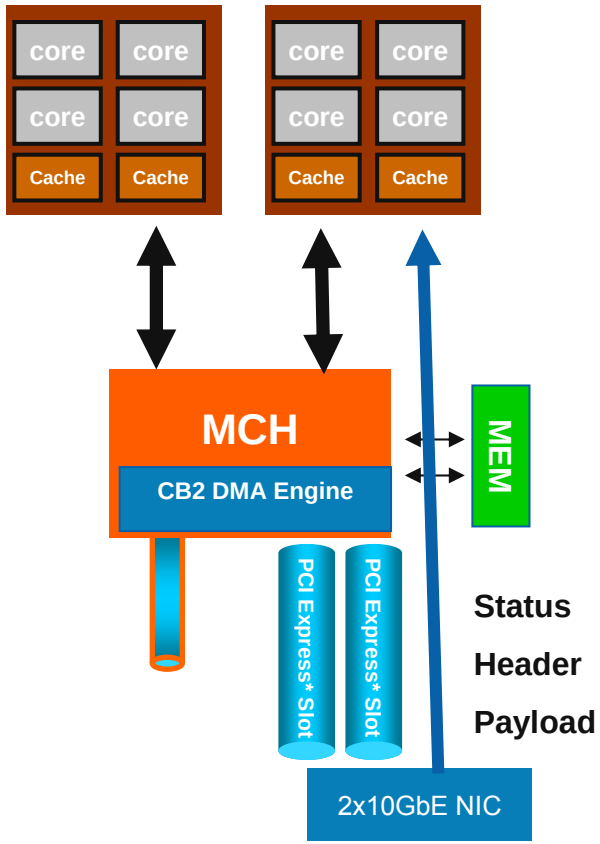

Small I/O:

DCA of all data from NIC to CPU

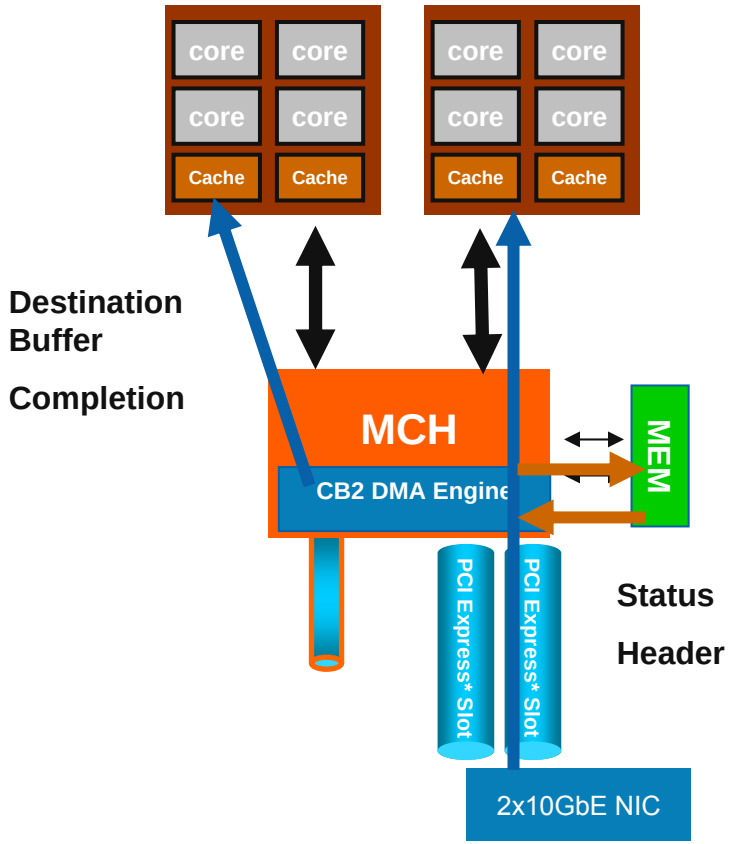

Large I/O:

DCA combined with Data Movement Engine (Crystal Beach 2 only)

Improving Ethernet Latency

- Ethernet Latency can be high
  - > Primary cause is interrupt moderation time
- Latency without interrupt moderation approaches Infiniband latency
  - > See graph
  - > No interrupt moderation = high CPU utilization (~100%)
- Low Latency Interrupt moderation addresses Ethernet latency

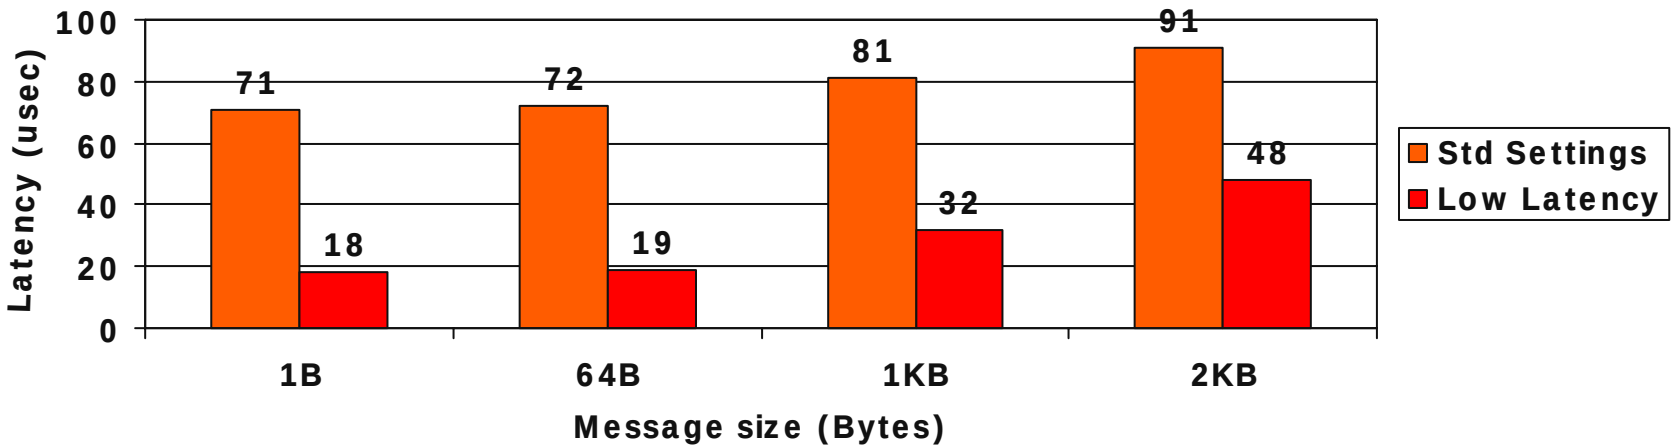

Bensley Platform, Red Hat RHEL4 - Kernel Version 2.6.12.4, Netperf Latency Test - Driver ver. 7.0.16, Driver Parameters Tuned for Low Latency - Interrupt Moderation disabled

Virtualized Server Networking

Problem: Server virtualization has a significant I/O performance penalty due to VMM software overhead of sharing NIC ports across multiple VMs

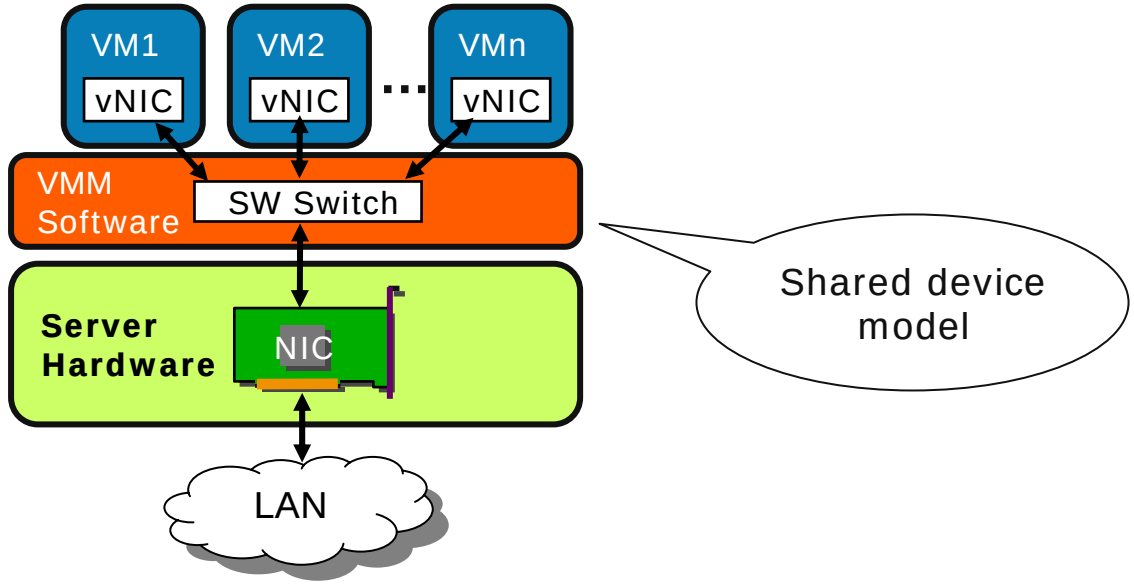

Solution: Platform and NIC hardware improvements for faster, more efficient networking in virtual servers

- I/OAT - Moves network data more efficiently through a virtualized system to provide Fast, Scaleable, and Reliable networking
- VMDq – NIC acceleration of VMM SW switch functions with multiple HW queues

Intel's goal: Narrow the networking performance gap between virtualized & single-OS servers with HW assistance

Virtual Machine Device Queues (VMDq)

More effective NIC sharing by sorting and grouping packets

- Receive Path:
  - Packets sorted into queues for destination VMs
  - Packets sent in groups to the VMM switch
  - Reduces number of times VMM switch code executes

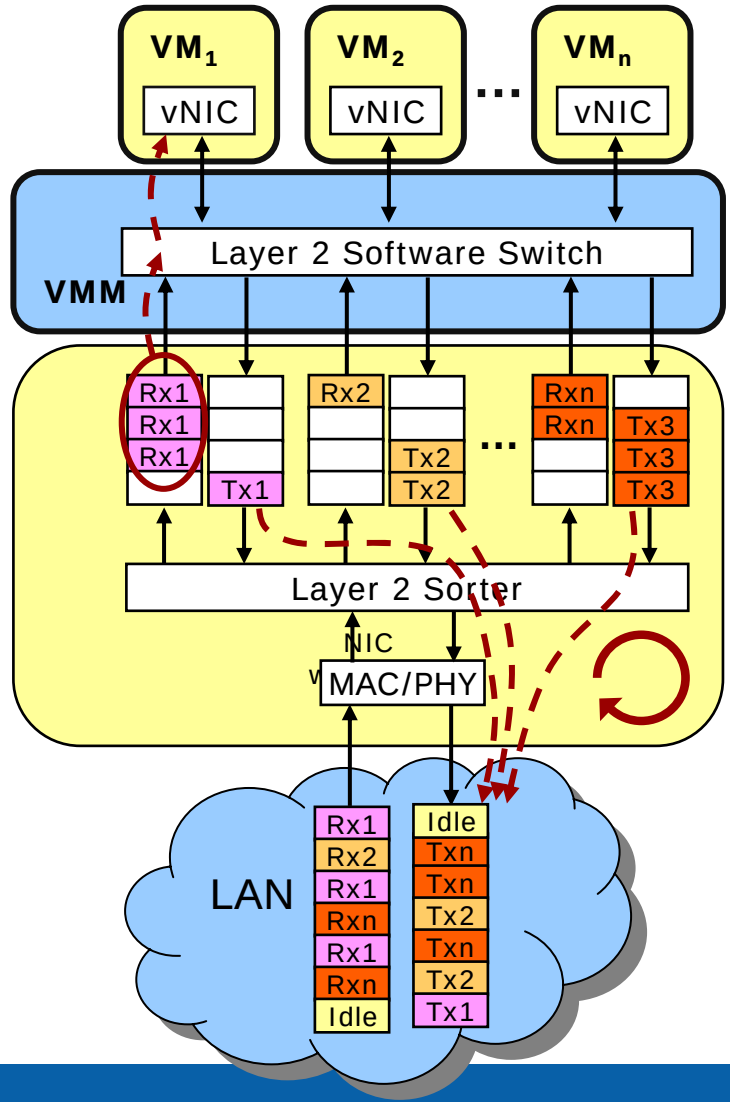

Transmit Path:

- Round-robin servicing of the transmit queue
- Ensures transmit fairness
- Prevents head-of-line blocking

Agenda

- What is Intel IOAT?
- Next Generation – IOAT2
- Real world Application Benefits with I/OAT

Software Ecosystem Support for I/OAT

Vendor	Product Version	Available
Microsoft	Microsoft Server 2003 Scalable Network Pack	Now
Linux	Linux Kernel 2.6.18	Now
suse	SuSE Enterprise Linux Server 10 (SLES10)	Now
redhat	RedHat Enterprise Linux 5.0 (RHEL5)	Now
vmware	VMWare ESX Server 3.5 (target)	2H'07

Intel I/OAT is tightly integrated into popular OS & VMM products
–Safe and flexible I/O acceleration choice for IT customers
–Avoids support risks of "new" 3rd-party network stacks
–Preserves existing network requirements – Teaming, failover, VLAN

Improving Media Server Performance

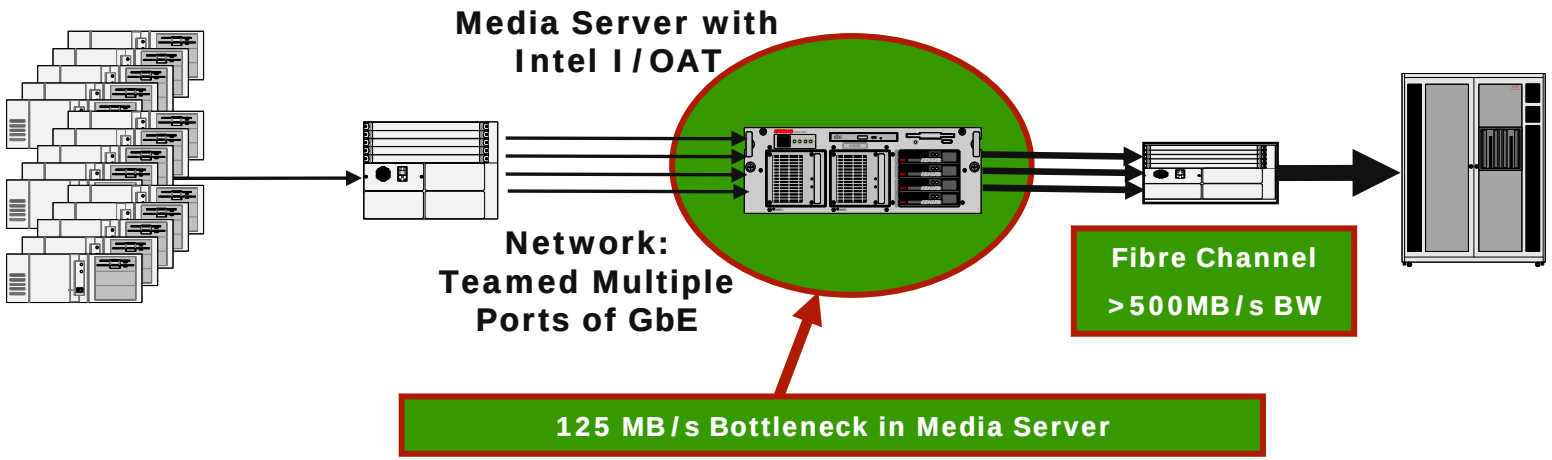

Results with Intel I/OAT + Teamed GbE NICs

- Faster Backups
- Media Servers Handle More Clients

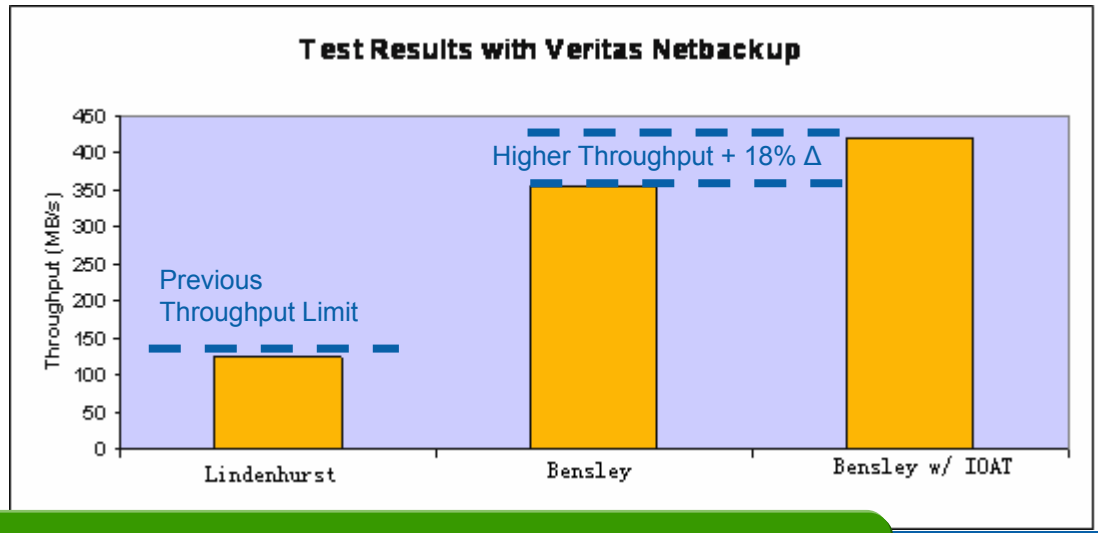

Helping IT 'Do More with Less'

Intel® I/OAT Real World Application Benefits

Front-End Server Throughput

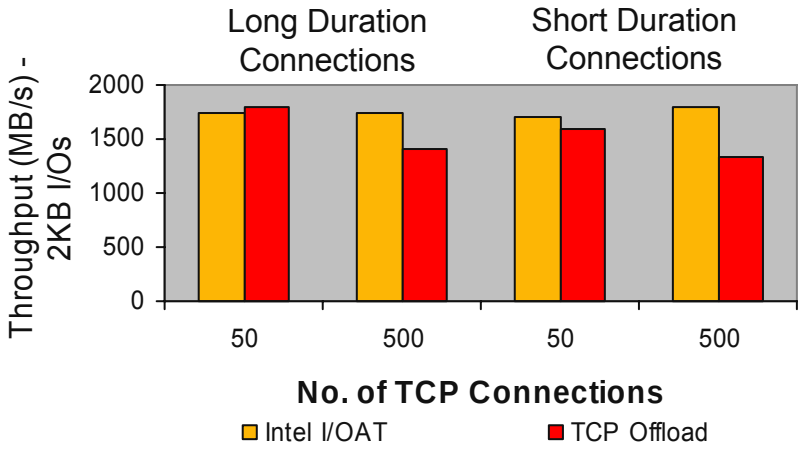

Data Source: Intel Labs, Verified by Ohio Supercomputer Center

Veritas* Network Backup Performance

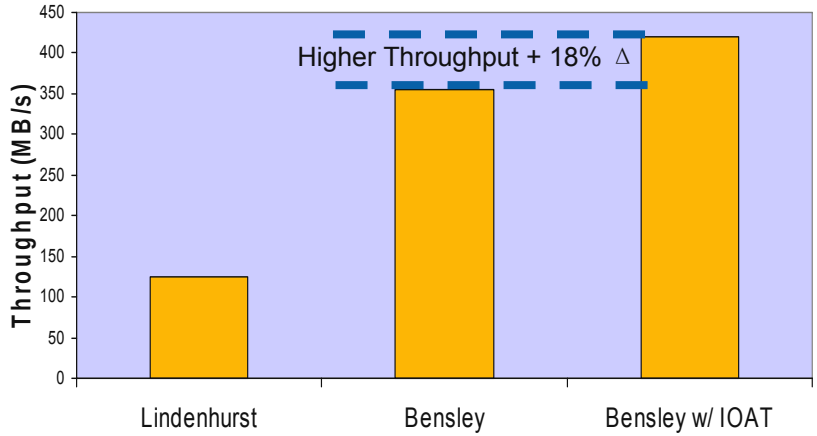

Data Source: Intel IT and Symantec

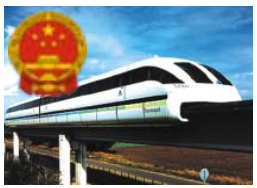

E-Document Delivery System

- Document distribution system
- Serves 1600 customers
- 130 remote offices & stations

- I/OAT reduced response time by **20%**
- Database back-up time reduced **28%**

Data Source: China Ministry of Railroad, Shanghai Administration Bureau

iSCSI Network Storage

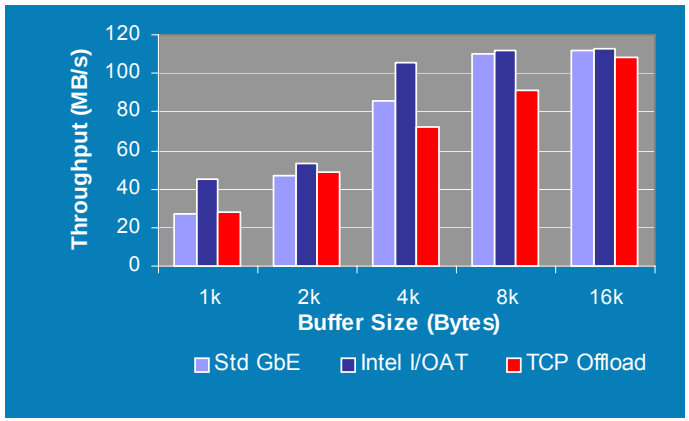

Data Source: Intel Labs, Preliminary Verification by Quantum

Summary

- Network I/O Demand is Increasing Rapidly
- IOAT helps IT to “Do More with Less”
- Next Generation IOAT2
  - Scales with 10GBE, Faster CPU and More cores
  - Low Latency, DCA, Header Replication/ Splitting
  - Better Virtualization Support
  - Tightly Integrates with Major OSes


intel
Leap ahead™