


The Economics of InfiniBand Virtual Device I/O

IBTA's Technical Forum '08: InfiniBand and the Enterprise Data Center

Jacob Hall
Wachovia Corporate & Investment Banking
VP, Chief Architect
Technology Products Group

We Must Design to Save



**Design your technology in a way
which saves cash, time, energy, and
the environment automatically**

Only the Economics Matter



...we must accelerate the adoption of better performing per watt technologies sooner. By increasing the rate of change, we will transfer \$\$ from energy and facilities to improved capability and competitive advantage.

Economics is Eco Friendly



Eco-nomics – *noun*

1. (used with a singular verb) the science that deals with the production, distribution, and consumption of goods and services, or the material welfare of humankind.
2. (used with a plural verb) financial considerations; economically significant aspects: What are the economics of such a project?

Futures

Processor Arrays



I/O is Decoupled
from the Chassis

Smarter Way To...

Take Inventory

Enables Recycling

Easier Reuse

Low Cost Chassis

Less Eco Impact

Better Cooling

Power Efficiency

Best Packaging


Desktop Computing Appliances

Top 10 US Banks

Processors = 1,507,682

Top 10 Super Computers

Processors = 467,861



Virtualized Video / USB ...it
enables >10 Gbps sooner

The I/O Economy

Economic Report – A State of Disutility



48 Machine Configuration 2 Socket Quad Core / 2U

VM Environment

- 2U Compute Device
- 10 x Cables
- 6 x Ethernet Ports
- 2 x 4Gbps FC HBA
- 1 x 3Gbps SAS
- 2 x 70GB HD's
- 3+ Device Drivers
- 1 Firmware

x 48

Totals

Capital Cost: \$549,578.00

Weight: 2916 lbs

I/O Weight: 516 lbs

Fuel for 3000 Miles: \$1021.00

Space: 106U

Watts: 16,128

Cables: 592

Bandwidth: 816 Gbps

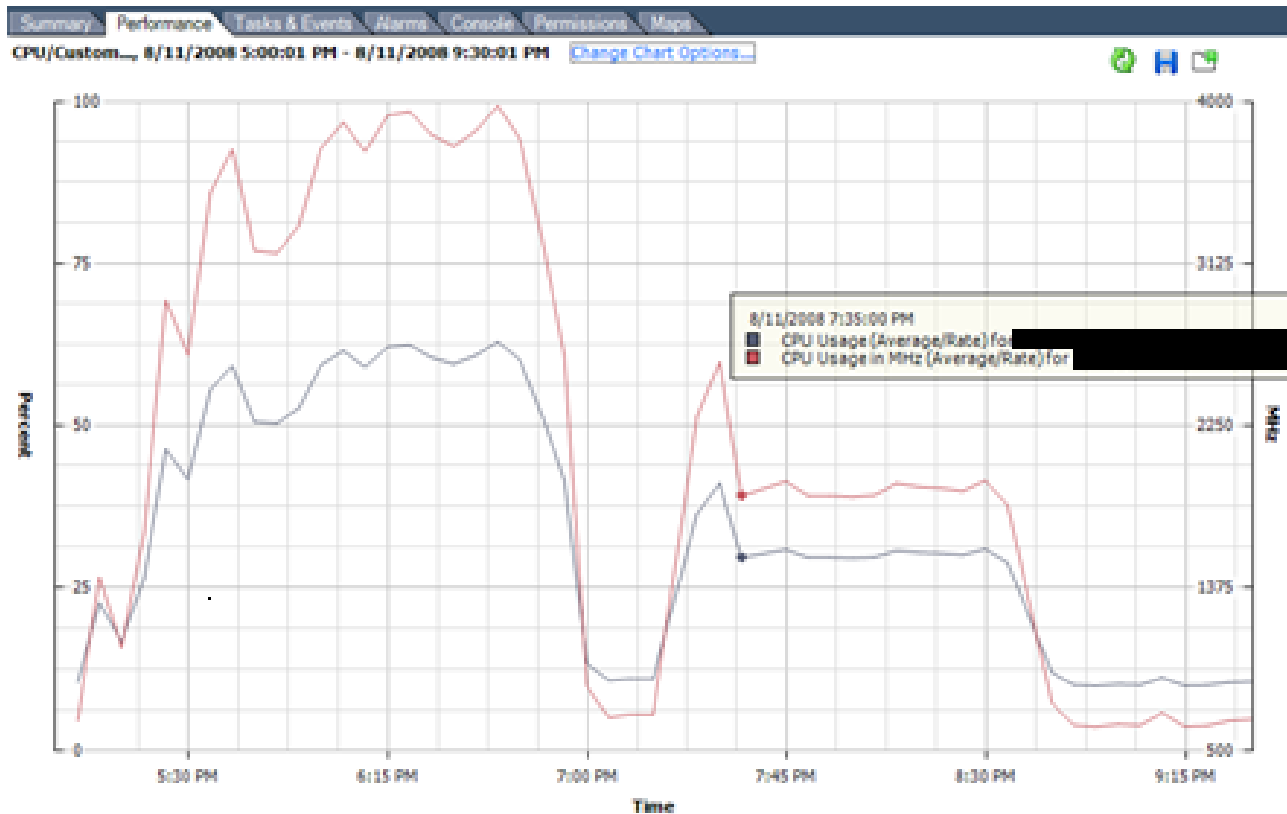
People: 

* Please review these facts. Quantified using readily available retail prices and data from the web.

Economic Trend - Latency Affects Everything



* Removing 240ms of 1Gbps Ethernet I/O Latency can Double Processor Efficiency



Performance Chart Legend

Key	Object	Measurement	Units	Latest	Maximum	Minimum	Average
■		CPU Usage (Average/Rate) for [redacted]	Percent	10.5	62.99	9.85	33.25
■		CPU Usage in MHz (Average/Rate) for [redacted]	MHz	663	3980	622	2100.40

J2EE to Database Server using VM's.


1/2 CPU Energy Consumption

Greater Efficiency, Less Servers, Faster Response

Upgrading I/O Takes To Long

Historically, It Has Taken too Long (7-8 Years) to Completely Transition to the next Fastest I/O.

To Improve Efficiency, I/O Designs Need to Help Accelerate the Adoption of Better Performing per Watt Technologies.



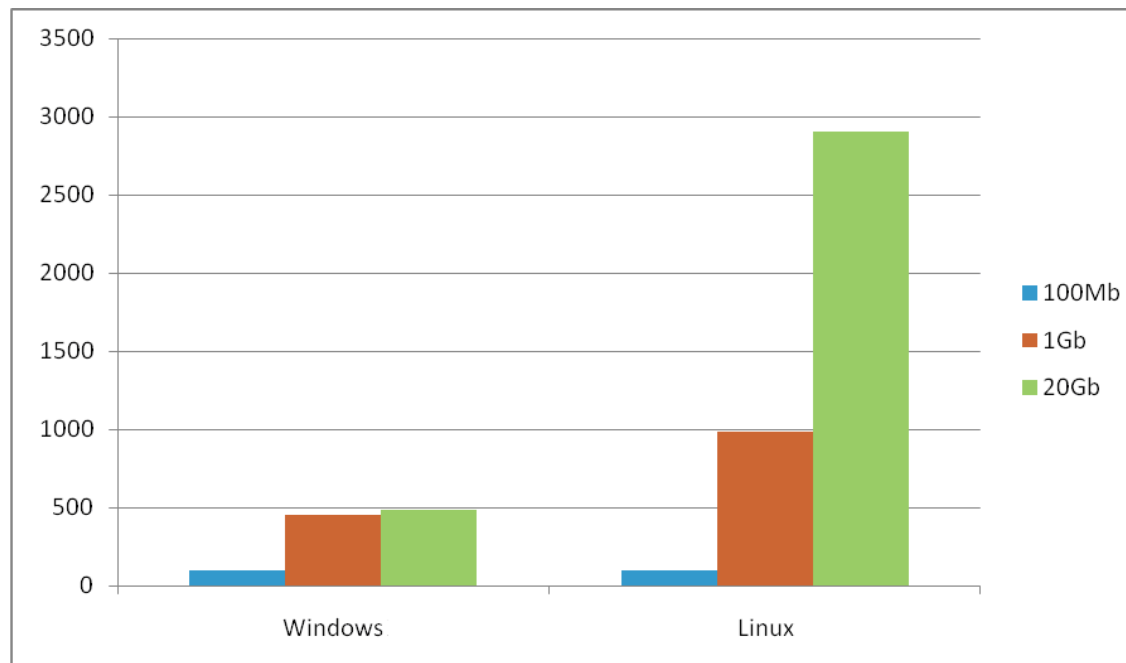
Economic Trend – Less Wires = Faster I/O



Two VM's on Separate Machines

PCI-Express 1.0

Utilized Netperf Network Benchmark Tool

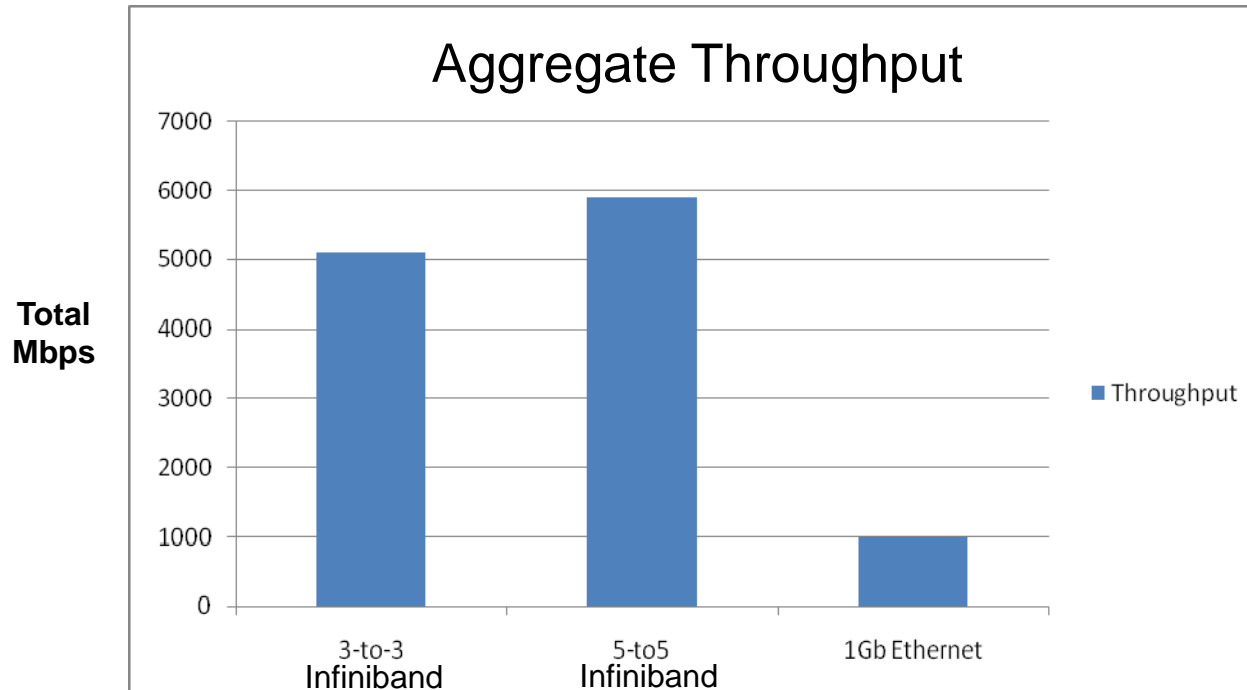


* Using the same VM configuration, Windows did not show a significant performance gain. We will be retesting with updated drivers soon and we expect to see performance gains for a single VM machine to machine test.

Economic Trend – VM Provides I/O Efficiency



By using Unified I/O under Hypervisors; we can justify the expense of faster I/O sooner and save =>80 Watts of power per machine, which is enough to add more processors, ram, or other. The savings over time are significant.




* Benchmark performed utilizing NetPerf and multiple VM's on separate physical machines

- One 20Gbps Infiniband port replaces six 1Gbps Ethernet connections (w/ bandwidth to spare for storage I/O)

Economic Outlook = Virtual Device I/O

48 Machine Configuration 2 Socket Quad Core



VM Environment

Totals

Capital Cost: \$324,900.00

Weight: 1568 lbs

I/O Weight: 128 lbs

Fuel for 3000 Miles: \$549.00

Space: 52U

Watts: 11,248

Cables: 200

Bandwidth: 1920 Gbps

People: 

* Please review these facts. Quantified using readily available retail prices and data from the web.

Economic Outlook - Unified I/O Savings



	Non-Unified I/O	Unified I/O	Unified I/O Savings
Capital Cost	\$549,578.00	\$324,900.00	\$220,000.00
Weight	2916 lbs	1568 lbs	1300 lbs
I/O Weight	516 lbs	128 lbs	380 lbs
Fuel to Move 3000 Miles	\$1020.00	\$548.80	\$470.00
Space	106U	54U	50U
Watts	16,128	11,248	4,880
Cables	592	200	390
Bandwidth	816 Gbps	1920 Gbps	1000 Gbps Faster

* These estimates were gathered using vendor websites. Results may vary depending on the vendor.

** 48 Port Ethernet and 24 Port Infiniband switches were utilized

*** Infiniband configurations assume no local disks.

Let This Influence Your I/O Design

Value of Use = Water

Very useful, serves many purposes
Could not survive without it
Water is free, but it is bottled
Low cost of change

Value of Exchange = Diamonds

Nice to look at
Not as useful, specific design
Maintains a high cost of change

* We need a low cost of change

You May Ask: How Do I Differentiate?

Add Significant Value / Innovate

- Design for Zero Transition Cost
- Enable Recycling
- Integrate Today's Add-On Features
- Integrate New Hardware Concepts

Simplify

- Reduce Touch Points
- Reduce the Weight

Efficiency

- Improve Performance Per Watt
- Improve Utilization per Watt
- Reduce the Wait



Infiniband in Our Enterprise Progress Update

What Affects Enterprise Adoption?



“The real price of every thing, what every thing really costs to the man who wants to acquire it, is the toil and trouble of acquiring it. What every thing is really worth to the man who has acquired it, and who wants to dispose of it or exchange it for something else, is the toil and trouble which it can save to himself, and which it can impose upon other people.” (Wealth of Nations Book 1, chapter V) 18

- Adam Smith



Infiniband as NAS Appliance Interconnect
Evaluating Infiniband for Mainframe SysPlex
Interconnect

Production Development Environment, hypervisor on
Infiniband with Fabric Management

Enterprise Opportunity for Infiniband

Hypervisor + Infiniband Virtual Device I/O

Consolidated Virtual Device I/O Everywhere

VM Memory, Infiniband HCA Disk Controller (DRAM & Flash SSD) – Swap Files, Non-Critical Data, High Speed Caching

Infiniband Attached Storage Array for DASD Consolidation, OS Booting



Why Enterprises Will Embrace Infiniband



“Perfect competition describes a market in which no buyer or seller has market power. Such markets are usually allocatively and productively efficient.”, Wikipedia

Ethernet is a monopoly, but is easily encapsulated as a virtual device
Infiniband supports 40Gpbs QDR; greater than 10 Meters
Infiniband 12x QDR Uplink @ 120 GB/s; 300 meters
Infiniband Supports large file sizes for Storage I/O
DCE / CEE will be a forklift upgrade anyway
Enterprises and Cloud / Utility Computing seeks “perfect competition”
Infiniband is a change agent and good for competition
It's Designed to Improve Operational Efficiency

Questions for the Enterprise



u·til·i·ty –adjective, having or made for a number of useful or practical purposes rather than a single, specialized one: a utility knife.

Which I/O Silicon provides all features and capabilities allowing us to maximize our consolidation efforts?

Does the labor (or disutility) required to retrain resources for Infiniband Unified I/O, exceed the cost of labor to maintain three plus independent I/O technologies?

Was Ethernet designed to be a Unified I/O Utility?

If Ethernet was not designed for storage, should we move storage I/O to Ethernet?

Does it make sense for I/O to evolve like Virtual Machines and Multi-Core Computing? Is there value in not evolving?

Which technology supports RDMA and is the most energy efficient?
Which provides the best processor efficiency?

Is there value in arguing over protocols ported to a protocol (Ethernet) which was not designed for storage?

The industry argument shouldn't be focused on iSCSI versus FCoE over legacy Ethernet for I/O consolidation. The discussion should be focused on which Unified I/O Silicon and Protocol Stack provide the best and most complete capabilities for all of our I/O demands and a reduced cost of change. More importantly, it should also focus on the Economics.

Thank You
