

Server and I/O Perspectives

Issues and Industry Trends in a Virtual World

Greg Schulz – Founder and Sr. Analyst
The StorageIO Group

Server and I/O Perspectives Issues and Industry Trends in a Virtual World

There is an old saying that the best I/O is one that you never has to occur which may be reality in fantasy land or even in a virtual world. Also consider that not all I/Os are storage related I/Os.

However the reality is that virtual environments provide a false sense of serenity to the user and application developer when it comes to actual IT resource utilization and impact.

This presentation takes a look at current server I/O capabilities and issues arising from various workloads along with industry trends to improve I/O performance.

Virtualization – Abstracting Storage and Data Resources

- Storage virtualization means many different things to different people
- In general, the logical abstraction of physical devices and their location
- Can be implemented on host server, in the network, in storage system
- Resources are required to perform virtualization and resources result in overhead

| Function | Activities | Benefits |
|-------------|--|---|
| Abstraction | Masking, mapping, redirection, device emulation (Virtual tape, disk volumes), protocol conversion. | Layer of transparency between applications and technologies and their physical locations. Support interoperability. |
| Aggregation | Volume management (pooling). NAS global name space, dynamic storage growth. Heterogeneous storage access. | Unified storage view, simplify storage management and data protection; enhance storage resource utilization and usefulness. |
| Migration | Movement of data from source to a new destination while data is being accessed. | Avoid downtime for relocation of data for technology replacement and ILM tiered storage |
| Protection | Replication, snapshots and point-in-time (pit) copies, continuous data protection (CDP), backup, archiving and security. | Support DR/BC, compliance and regulatory needs, simplified recovery, address diverse RTO and RPO needs. |
| Management | Performance, capacity planning, resource usage (aka SRM), policy management | Provide insight into resource usage activities and policy based management. |

Non Storage I/Os – Why You Should Care About Them

Non Storage I/Os include:

- Server to Server for general purpose networking
- Server to Server for data movement, mirroring and remote copy
- Server to Server for locking and data synchronization

Other I/Os involving storage include:

- Scanning for virus, spyware, compliance and other discovery
- Non storage related I/Os and why you should care about those I/Os
- Can you really eliminate I/Os? Is there such a thing as a virtual I/O?

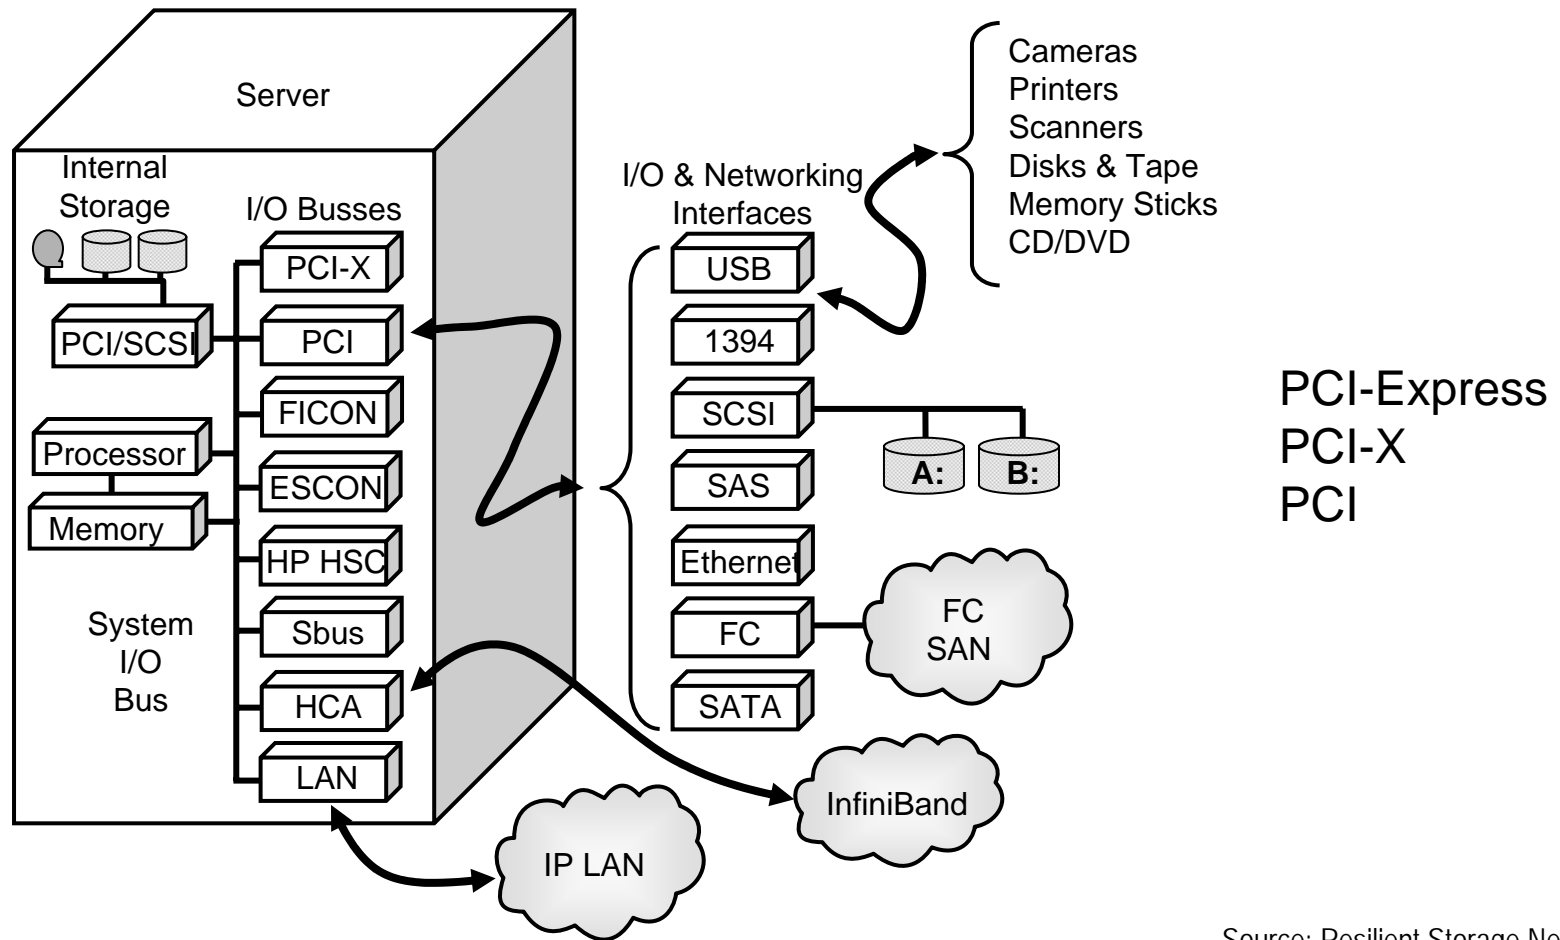
Why You Should Be Concerned

- These I/Os require memory, cpu/processor, backplane and bus adapter bandwidth

Can You Really Eliminate I/Os

- Virtual vs. Real I/O, what's the difference
- What about moving or “hiding” the I/Os?
- Who pays the price and what is the price?
- So can you really eliminate I/Os?
- Let’s talk about the “Scan Tax” or the “Scan Scam”
 - Overhead of compliance, email archiving, ILM, etc...
- Is there such a thing as a “virtual” non-existent I/O?

Where Are The Bottlenecks



Source: Resilient Storage Networks

Storage Devices and Media Options

| Technology | Feature/Function | Benefit or enables |
|-----------------------|---|--|
| 2.5" Form Factor Disk | Incremental drive improvements FC, SATA, SAS, SCSI | Economics, reliability, capacity backwards compatibility |
| SATA | MTBF improvements, more capacity, enterprise class | Lower cost storage for near-line & off-line tiered storage and retention |
| Capacity | 400GB now, 500GB on the way more RPM on some slower disks | More storage to manage more data and use for D2D and D2D2D |
| Cache Centric Array | More speed, capacity, features | Scalability and management ease |
| Distributed Array | Turn the crank, more.... | Scalability and co-existence |
| Clustered/Grid | Scale performance, capacity, RAS. Blurred differences between grids and clusters. | Scalability, leverage lower cost and commodity components, simplify and enable management. |

The trend: – Smaller, faster, larger capacity devices, scale out and up
The buzz: – Grid confusion vs. clusters, virtualization, partitions
Watch out for: – Grid hype to rival that of virtualization hype
Important questions: – Enterprise vs. desk top SATA?, is it a grid, or is it grid like?


Storage Interfaces and Protocol Options

| Technology | Feature/Function | Benefit or enables |
|-------------------------|-------------------------------------|-----------------------------------|
| Fibre Channel | 1Gb, 2Gb, 4Gb, 8Gb, 10Gb | Low latency, Scalable, Flexible |
| Ethernet | 10/100/100Mb & 10Gb | Ubiquitous, Cost, Knowledgebase |
| FICON | Point-point, switched, cascade | Devices, performance, distance |
| InfiniBand | Cluster interconnect, Low latency | Low response time, unified access |
| IP | TCP/IP – iSCSI, NAS, iFCP, FCIP | Distance, low cost, flexible |
| iSCSI | Low cost block storage access | Distributed and low cost |
| PCI, PCI-X, PCI-Express | Network and storage adapters | Access to I/O and storage |
| SAS | Disk drive interface | Performance, flexibility, cost |
| SATA | Disk drive interface, high capacity | Low cost storage |

The trend: – Faster interfaces, lower cost, more flexibility, more IP
The buzz: – InfiniBand, RDMA/iSER, SAS, Fibre Channel is dead
Watch out for: – Leverage the appropriate interface for the task at hand
Important questions: – What does an interface enable for your business?
 What are next speeds (e.g. 8Gb FC, 10Gb, etc.)?

Performance Enhancing Substances “Technologies”

- 10Gb Ethernet is here, however, 10Gb “Copper” Ethernet will be affordable
- InfiniBand and associated ULPs (e.g. iSER/RDMA, SRP, TCP, xDAPL) are here
- Start planning for PCI-Express (PCIe), more memory and faster CPUs
- Not all FC disks are fast, not all SATA disks are slow, watch for SAS
- Partitions and virtualization (e.g. consolidation) may not mean performance
- Clusters can be for scaling of resources and availability (not just servers)
- Storage QoS, policy based movement, migration will help with backup
- 4Gb Fibre Channel can benefit non performance (bandwidth) applications
- Don’t forget about SSD as it remains the most effective for \$/IO...




Courtesy: www.darwinawards.com

Closing Comments

- Learn and understand where the bottlenecks are and why they occur
- Some of the performance issues are the cost of doing “virtual” business
- Virtual computing can reduce some costs, however it can shift others
- The politics of storage can add to the cost of storage and bottlenecks
- Learn to speak the language of the business and of developers
- Interoperability is in the eye's of the beholder
 - Your view of interoperability will greatly determine if it exists and what it is
- Of course the obligatory plugs:
 - Resilient Storage Networks – Elsevier
 - www.storageio.com (see portfolio for articles etc)

*SNIA
Recommended
and endorsed
Reading*



Questions

Thank You

Greg Schulz
The StorageIO Group

A Copy of This Presentation Is Available At:

www.storageio.com/downloads.htm

Username: CMG Password: CMG