

High Performance Storage for Bioinformatics Research

Zacharias Stankiewicz

Kulathinal Lab

July 19, 2010

Definition of Scope

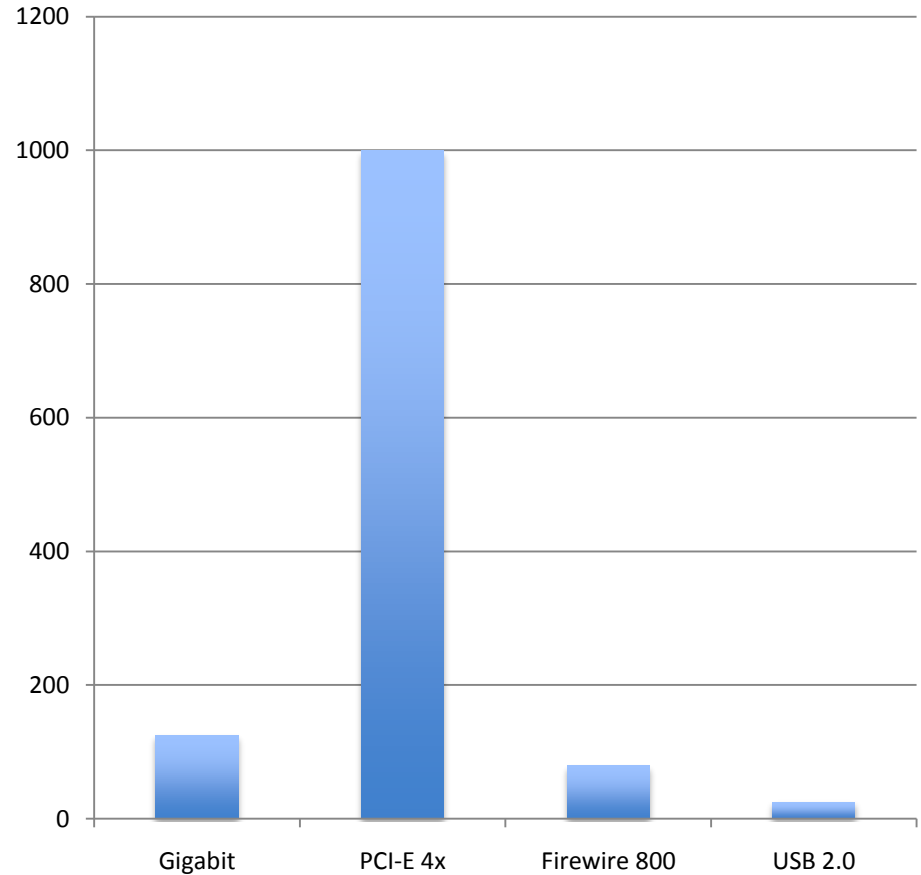
- Determine optimal combination of price and value for a locally-attached storage system
- Target: new Bioinformatics server.
- Ideally, the solution will be:
 - A good value
 - Scalable
 - High-Performing

FIRST QUESTION

How should the storage interface with the server?

Benefits of locally-attached storage in comparison to networked storage

- Speed
- Speed
- Speed
- Greasy fast lightning speed.
- Still share over a LAN (the server is up to the task)



Connectivity Options

Local Bus

Drives tethered to system bus and mounted inside server.

- Fastest access time
- Lowest CPU overhead
- Limited to finite amount of internal drives
- RAID options limited
- Space and heat issues may arise

External Bus

Drives in an enclosure connected to the server

- Expandable
- Cooling and power options
- Most methods of connection too slow (i.e. Firewire or USB)
- More computing overhead
- External enclosures can be very expensive

Connectivity Options - Solution

Compromise Solution:

External enclosure connected to system bus via PCI-E interface drive controller card

- Almost as fast as internal interface (card offsets a lot of cpu overhead)
- Can (must) use with an external enclosure, more options
- RAID options intact



*Pictured:
3ware 9650SE-4LPML*

SECOND QUESTION

Which enclosure should we use?

Enclosure - Solution

Enclosure:

In order to maximize the lifespan of drives a well cooled and powered enclosure is necessary

- Well-cooled
 - Room for even more fans
- SATA backplane interface – works well with SAS card
- Easy access – no tools needed



*Pictured:
Habey DS-1220N 12-bay 3U
Rackmount Storage Enclosure*

THIRD QUESTION

What options do we have for the storage system?

Drive Option – Consumer-Grade

3.5” Consumer-Grade Hard Drives

Typical hard drives as seen in most desktop computers.

- By far, the most cost-effective solution (14 GB for \$1)
- Standard drive controller and file system friendly
- Slowest read speeds
- Longer-term problems with vibration and heat tolerance



*Pictured:
WD Black 1.5TB 7200
RPM*

Drive Option – Solid-State

2.5” Flash-Based Storage

RAM-like storage has no moving parts, data storage in transistors

- Fastest read and write speeds of any option
- Nigh-indestructible (physically)
- By far, the most expensive drives (0.53 GB for \$1)



*Pictured:
Kingston 128GB SSD*

Drive Option – 2.5” Performance

2.5” Performance Drives

A smaller hard disk with high performance (spins faster) inside a standard-sized heat sink

- Faster sequential read speeds
- More resistant to heat, vibration than 3.5” drives
- More expensive than 3.5” drives (only 2.14 GB for \$1)



*Pictured:
WD Velicoraptor 600GB 10K
RPM*

FOURTH QUESTION

How do we determine the best drives for our use case?

Objective Performance Measurement

P/V Metric = GB/\$ * Sequential Read

*GB/\$ - Using Amazon.com pricing for “typical drive”
(pictured)*


*Sequential read speed is benchmarked for a single drive
and in megabits (8 megabits = 1 megabyte) per second*

- This metric equally weights the value of performance and value
- Changing the relative weights of these two values is irrelevant until ratios approaching 6:1 in favor of read speed performance.

Objective Performance Measurement - Results

drive	GB	cost	gb/\$1	sequential read	PV Score
wd black 1.5 tb 7200 rpm	1536	109.99	13.96	148	2066
wd 1.5 tb "enterprise drive"	1536	244.99	6.26	144	901
wd velicoraptor 2.5 600gb 10000rpm	600	279.99	2.14	145	310
kingston 128gb ssd	128	239.63	0.53	222	68

mb/sec



Objective Performance Measurement – Additional Considerations

- Relative performance on the basis of sequential read speed does not justify purchase of higher-performance drives.
- Higher performance drives are more competitive on the basis of random access time and write speed.
- In a RAID array, additional overhead is added by each layer of abstraction from the system bus.
- A catastrophic failure rate is needed to justify purchasing drives based on reliability data provided by vendors.
- Issues such as data density preclude using 2TB drives, and may even recommend using 1TB drives in lieu of 1.5TB drives

WRAPPING IT ALL UP

What is the recommended complete solution?

Complete Solution

- 4 – 8 1.5TB consumer-grade hard disks in RAID 5 array
 - Yielding either ~4-8TB of high-performance storage
- Drives in an external enclosure
- Drives interfaced with server via a PCI-E RAID controller

