Whither Hard Disk Archives?

Dave Anderson

Seagate Technology

6/2016

Topics as They Relate to Large Storage Archives

Where Topology might go

Basic HDD Topologies – advantages & disadvantages

- Hyper converged
- Networked Storage

Networking Considerations

Where Capacity might go

- Platter capacity = areal density
- Platter size
- Platter count = New form factors

Where Intelligence might go

One more thing

(Hyper)converged Architectures

Combine in a single system unit:

- Processors & memory
- Storage
- Networking

Advantages

- Interface & architecture simplicity
- Local storage management

Disadvantages

- Cost – CPU complex for each set of HDDs
- Inflexibility – limited variability in CPUs/HDDs relationship
- A more complex cooling problem, perhaps

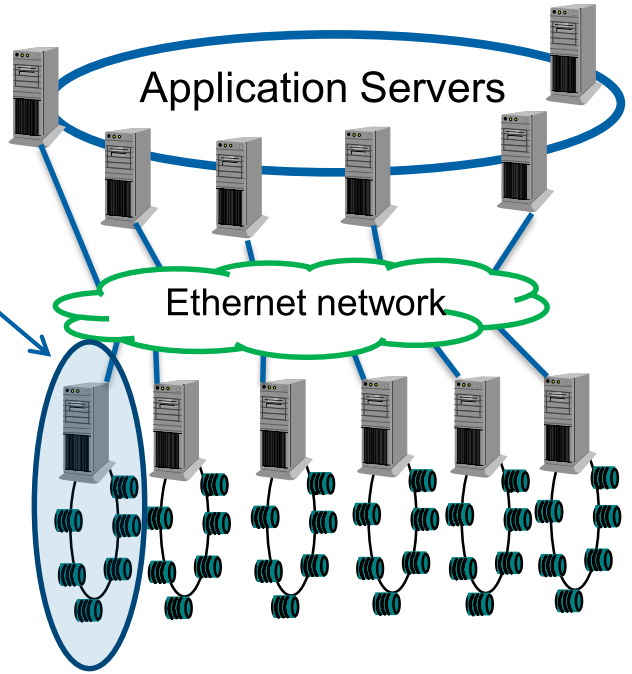

Networked Storage

Common Storage Pool:

- Attached via network to processors

Advantages

- Lower cost, no storage servers
- More redundancy freedom
- More freedom in CPUs/HDDs investment
- Simplifies software stack (no storage servers)
- Perhaps lower latency

Disadvantages

- Management practice not as developed
- Not as well developed a software stack
- Network picture not fully developed
- Relatively high latency interface
- Need low latency network for shared SSDs

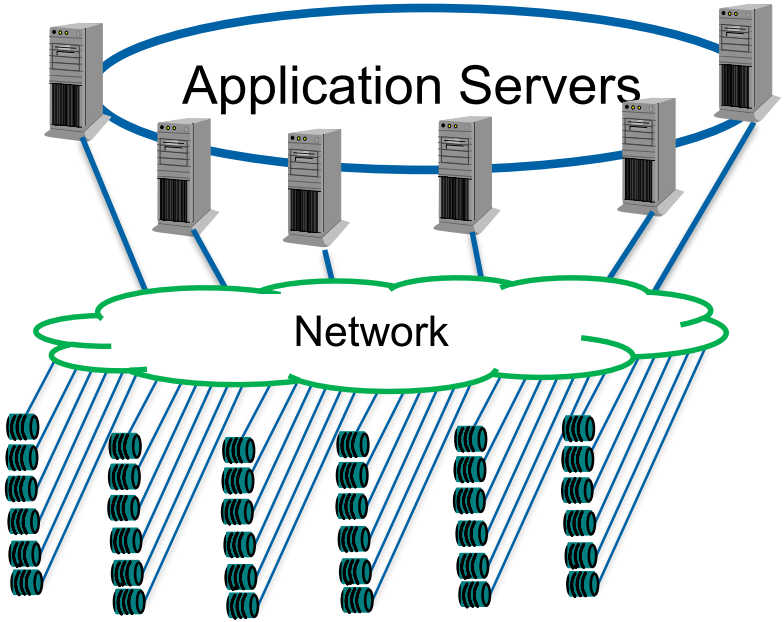

What about that Networking Part of Networked Storage?

Today's choices: Ethernet, Infiniband, Fibre Channel

- Ethernet has software-based protocol processing = more overhead
- Nondeterministic overhead – occasional dropped frames
- Infiniband not nearly as widely deployed, not an HDD interface

Need: low latency network (no good choice today)

- Enables networked, shared solid state storage option
- PCIe does not scale well
  - Cannot connect large numbers of drive economically
  - No good dual port (yet)
- Ideal low latency network would support:
  - Link types: optical & electronic
  - Protocol types: blocks & objects

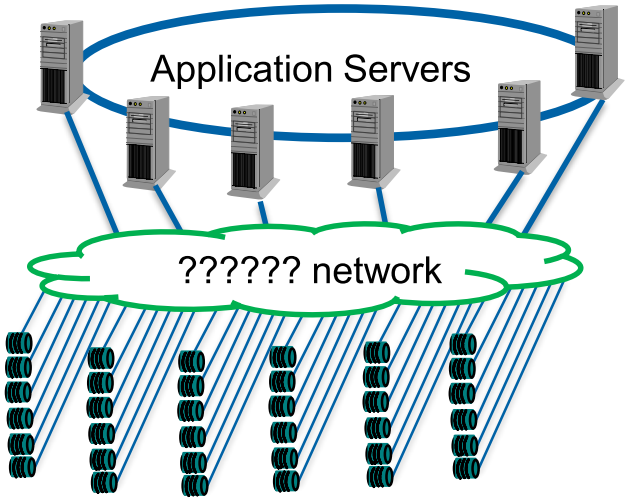

Where could this go? Check out UC Berkeley's FireBox concept

Where Areal Density Might Go:

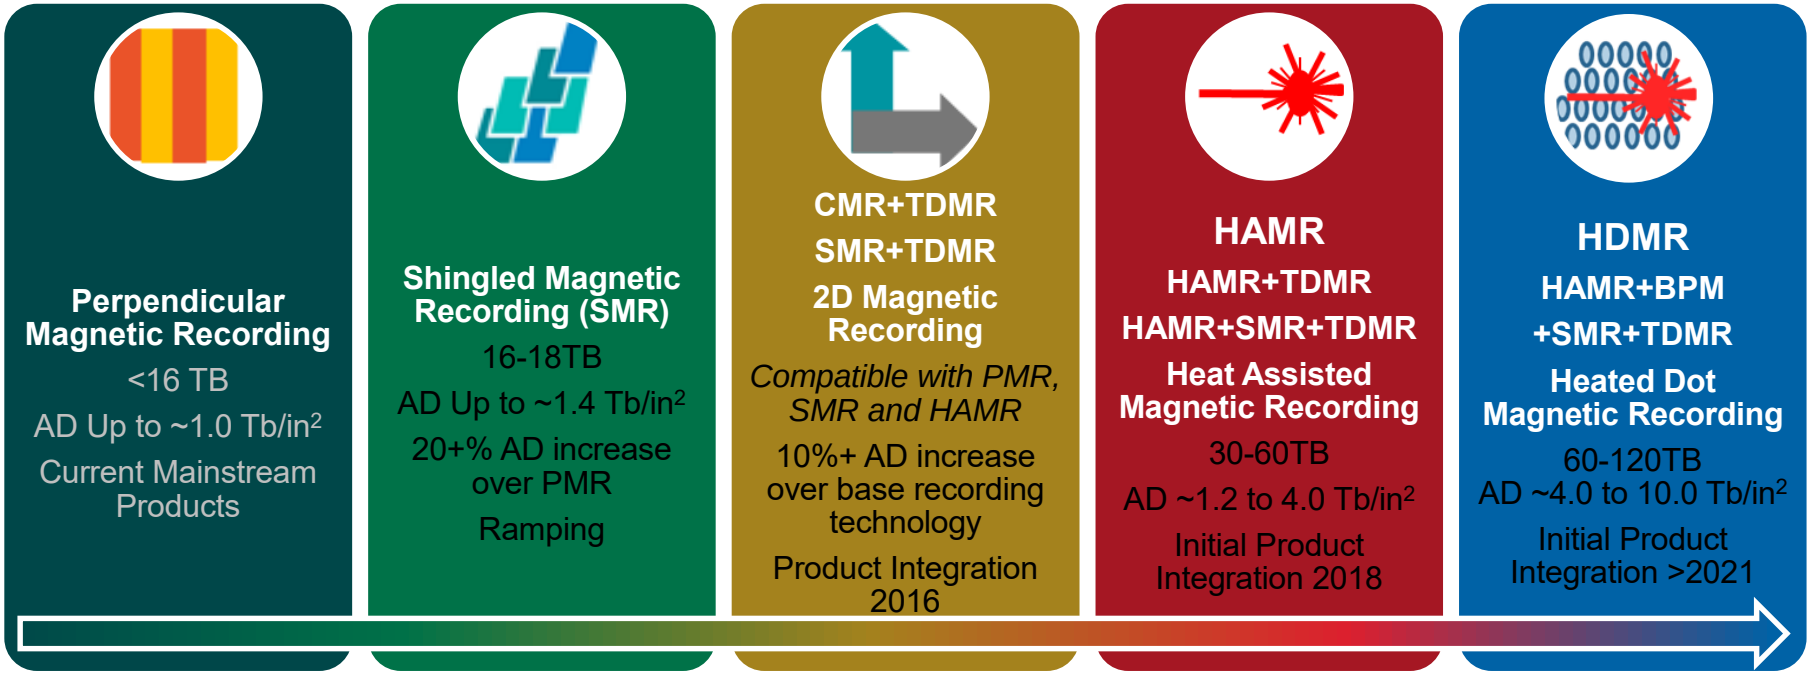

Seagate is exploiting many different technologies for capacity growth

More Capacity per HDD: The Form Factor Factor

History is littered with old HDD form factors:

- >5.25” - 5.25”– 3.5” – 2.5” – 1.8” – 1.x”

- Just because you built it in the past ,
- doesn’t mean you can build it again

Helium enables more platters in current form factor

A New Form factor is VERY expensive

- Changes in cabinets & chassis
- Changes in Component suppliers’ products
- Changes in drive manufacturing

Most feasible is not changing media size

- 3.5” x 1.6”?

3.5” x 1.0”

3.5” x 1.6”

One More Thing: Placing a little Computing Power with the Data

Enable application processing at the storage device (HDD & SSD)

First - sort of - product by ICL in 1979

Published in 3 academic research papers in 1998-2000

Why now:

- Movement to unstructured data
- Massive data sets
- Movement to storage objects

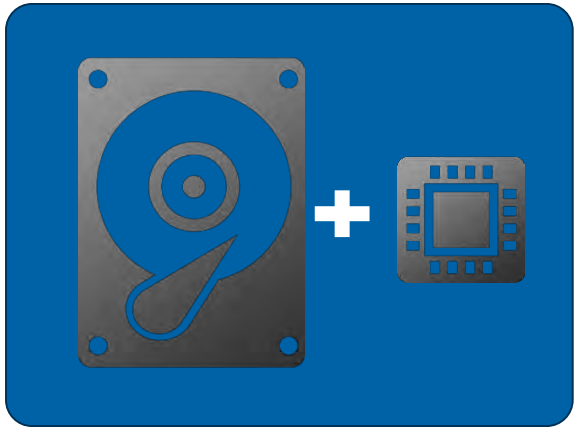

Active Disks: to Scale Search with the Data Size

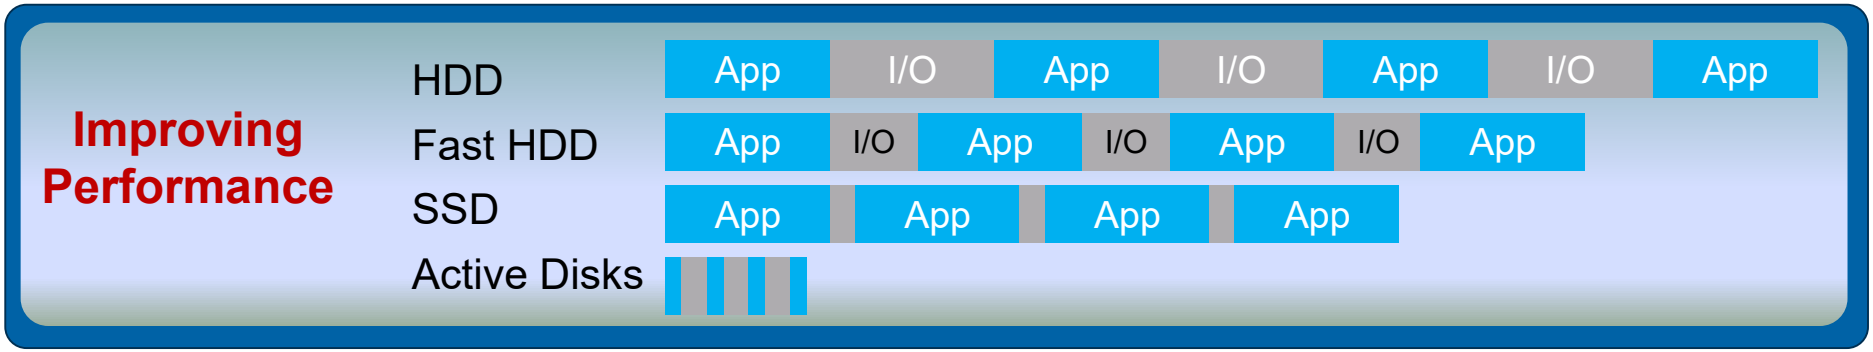

Motivation of this architecture:

- Parallelize analysis of data
- Reduce host data transfers
- Reduce application run time

Scale data processing with data size!

- Note the effect of spreading data across more drives!
- May impel wide declustering of data

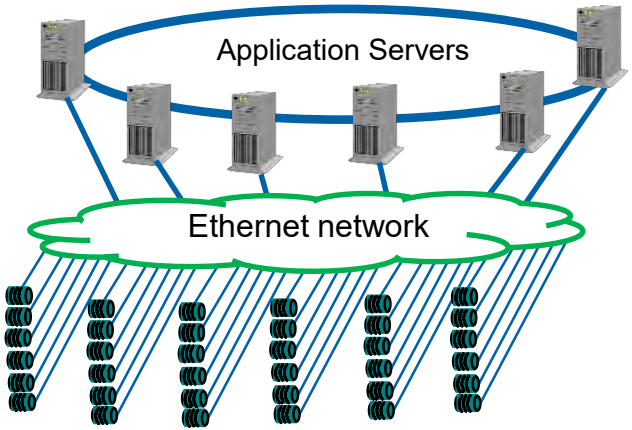

Research Papers

From Archaya: http://www.vldb.org/conf/1998/p062.pdf
Other papers: http://www.cs.umd.edu/~hollings/cs818z/s99/papers/activeDisks.pdf
http://redbook.cs.berkeley.edu/redbook3/idisk.pdf

Quantifying the Active Disk Benefit

Execution time Reduction:

4 active disks: up to 60%

32 active disks: up to 95%!

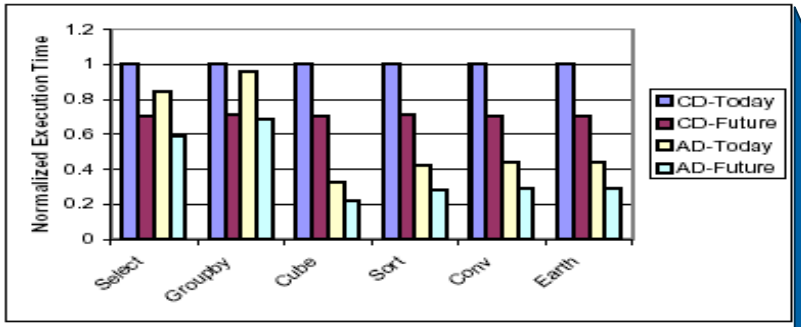

(a) 4-disks

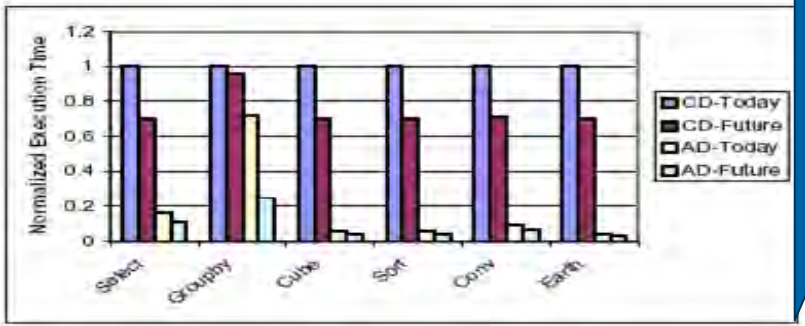

(b) 32-disks


Active Storage For Large-Scale Data Mining and Multimedia

Erik Riedel
Garth Gibson, Christos Faloutsos

Parallel Data Laboratory,
Center for Automated Learning and Discovery
Carnegie Mellon University
www.pdl.cs.cmu.edu/Active

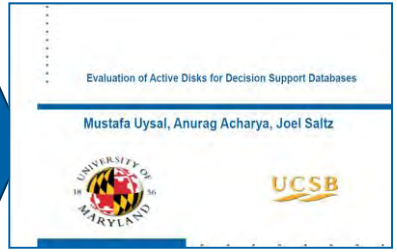
Evaluation of Active Disks for Decision Support Databases

Mustafa Uysal, Anurag Acharya, Joel Saltz

UNIVERSITY OF MARYLAND

UCSB

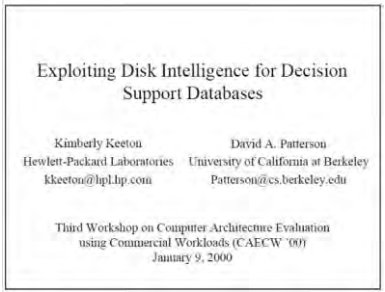
Exploiting Disk Intelligence for Decision Support Databases

Kimberly Keeton
Hewlett-Packard Laboratories
kkeeton@hpl.hp.com

David A. Patterson
University of California at Berkeley
Patterson@cs.berkeley.edu

Third Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW '00)
January 9, 2000

Summary

(Hyper)converged - today's dominant topology

Strong interest in Networked Storage

- Several issues need addressing:
- Holds a promise of enabling new architectures

Areal density (capacity per platter) will be increasing

New form factors are expensive, choosing one cannot be done lightly

Large Archive focused innovation looms over the horizon