My 10 Year Retrospective

1992: published dissertation on RAID,
founded PDL research lab at CMU

86-95:disk arrays

- parity logging, declustering, RAIDframe rapid prototyping

92-01: parallel prefetching filesystems

- disclosure, cost-benefit allocation, speculative execution

95-04?: network attached secure disks

- direct disk access, capabilities, objects with attributes, asynchronous oversight (out-of-band),active disks
- NSIC NASD WG, SNIA OSD, T10/SCSI OSD proposal

10 Year Retro/Prospective

3.5” disk	92	02	12?
Capacity (B)	320 M	100 G	30 T
$/Byte	2 / M	2 / G	2 / T
Track (KB)	32	330	3400
Access msec	20	10	5
LAN bits/s	10 M	1 G	100 G
IF bits/s	80 M	1.6 G	10 G
Read MB/s	1.7	40	800
Read once	3 min	40 min	11 hrs

Array box	92	02	12?
Capacity (B)	22 G	28 T	36 P
Actuators	12	280	1000
IO/sec	600	3,000	20,000
MB/sec	15	200	400,000
Read once	24 min	41 hrs	24 hrs

Disks get colder

But arrays with out-of band access need not.

10 Year Retro/Prospective

	1992-2002	2002-2012
Role in Hierarchy	Secondary	Secondary and tertiary
Scaling	Bigger boxes	Clusters of boxes in single mgmt space
Fault Tolerance	RAID 1 or 5	MTTDL = X years
Cost	How many IT staff?	What is not automatic?
Virtualization	Create/Resize volume	MTTDL = X years
Policy	RAID 1 or 5	MTTDL = X years
Performance	Bandwidth finally matters	Bandwidth, balancing, adaptive resources