SCSI vs. ATA – More than an interface

Dave Anderson, Jim Dykes, Erik Riedel
Seagate Research
April 2003

Seagate
We turn on ideas

Outline

Marketplace
- personal – desktop, low-end servers
- enterprise – servers, workstations, arrays
- consumer – appliances
- mobile – laptops

Mechanics & Electronics
- choices, comparison

Performance
- the direct impact

Reliability
- factors

Summary

Disc characteristics

User-visible characteristics

- Data rate ~ *(2πr) x (density) x (rpm)*
- Capacity ~ *(πr^2) x (density)*
- Seek time ~ *r*

Internal characteristics

- Areal density – *density*
- Platter size – *r*
- Spindle speed – *rpm*

Areal density

Data rate ~ *(2πr) x (density) x (rpm)*

Capacity ~ *(πr^2) x (density)*

Seek time ~ *r*

higher density, higher data rate

higher density, higher capacity

Areal density – *density*

- how many bits you can squeeze in

The bad news

- requires more signal processing
- tolerates less “noise”
- harder to do track-following (*servo*)

Platter size

Data rate ~ *(2πr) x (density) x (rpm)*

Capacity ~ *(πr²) x (density)*

Seek time ~ *r*

larger platter, higher data rate

larger platter, higher capacity

bad news – larger platter, higher seek time

Platter size – *r*

- large, smaller, smallest

More bad news

- larger platter, more power

Spindle speed

faster platter, higher data rate

Data rate ~ *(2πr) x (density) x (rpm)*

Capacity ~ *(πr^2) x (density)*

Seek time ~ *r*

Spindle speed – *rpm*

- slow, fast, very fast

The bad news

- creates more “noise”
- more (bad) vibration
- more speed, more power

Comparing mechanics

Cheetah 10K.6
enterprise

Barracuda 7200.7
personal

more
rigid
structure

larger actuator assembly

Comparing mechanics (2)

Cheetah 10K.6
enterprise

Cheetah 15K.3
enterprise

smaller platters – less mass, shorter seeks => less capacity

Mechanics summary

Basic design choices

- how high data rate
- how much capacity
- how small seek time

Each one affects which parts you pull off the shelf

- how you put them together depends on how they will be used
- some decisions driven by the marketplace
- some driven by cost

Outline

Marketplace
- personal
- enterprise

Mechanics
- choices – pick your parts
- the choices made to date – comparison

Electronics
- comparison

Performance
- the direct impact

Summary

Spindle speed

- highest spindle speed will be perfected in enterprise drives
- when it becomes cheap enough, it becomes the norm in personal drives as well

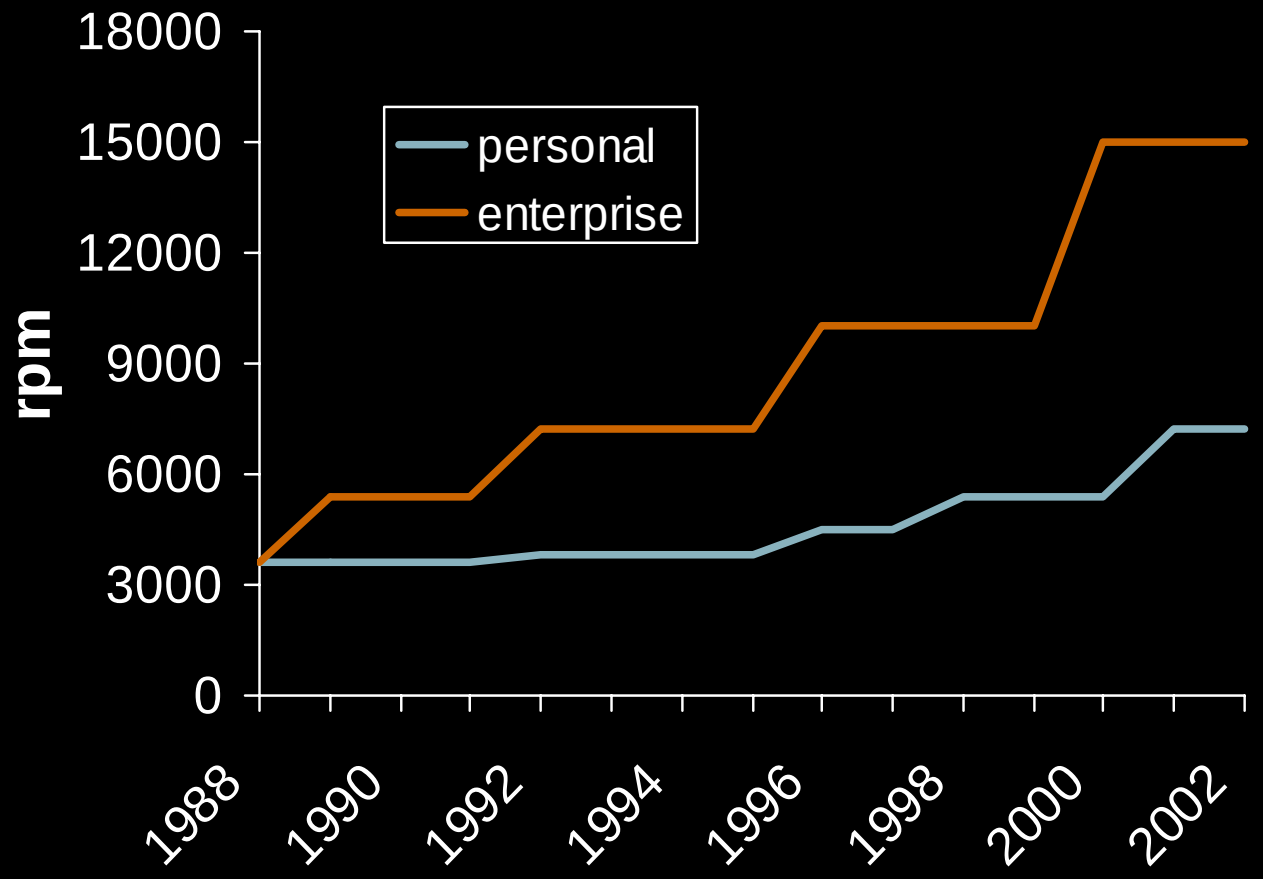

product information for Seagate and Control Data disc drives since 1988

Capacity

- wider spacing for enterprise drives
  - wait until you get 2x
- this can confuse comparisons based on equivalent capacity
- highest capacity was always in enterprise drives, until 2003

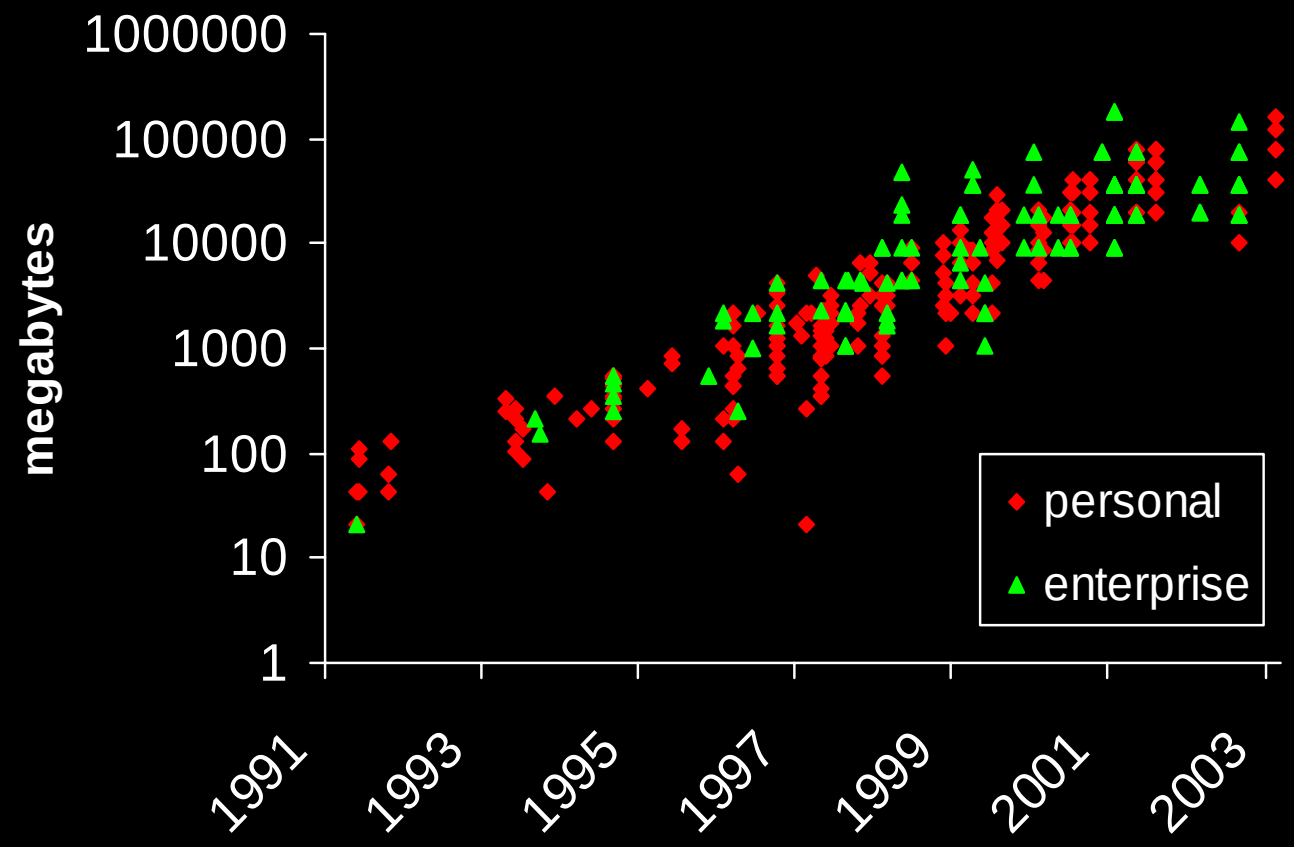

product information for Seagate disc drives introduced since 1991

Seek time

- enterprise always more aggressive on seek performance
- trend is toward further separation
- sensitive to both the mechanics and the signal processing
  - moving the arm fast enough (starting)
  - staying on track (stopping)

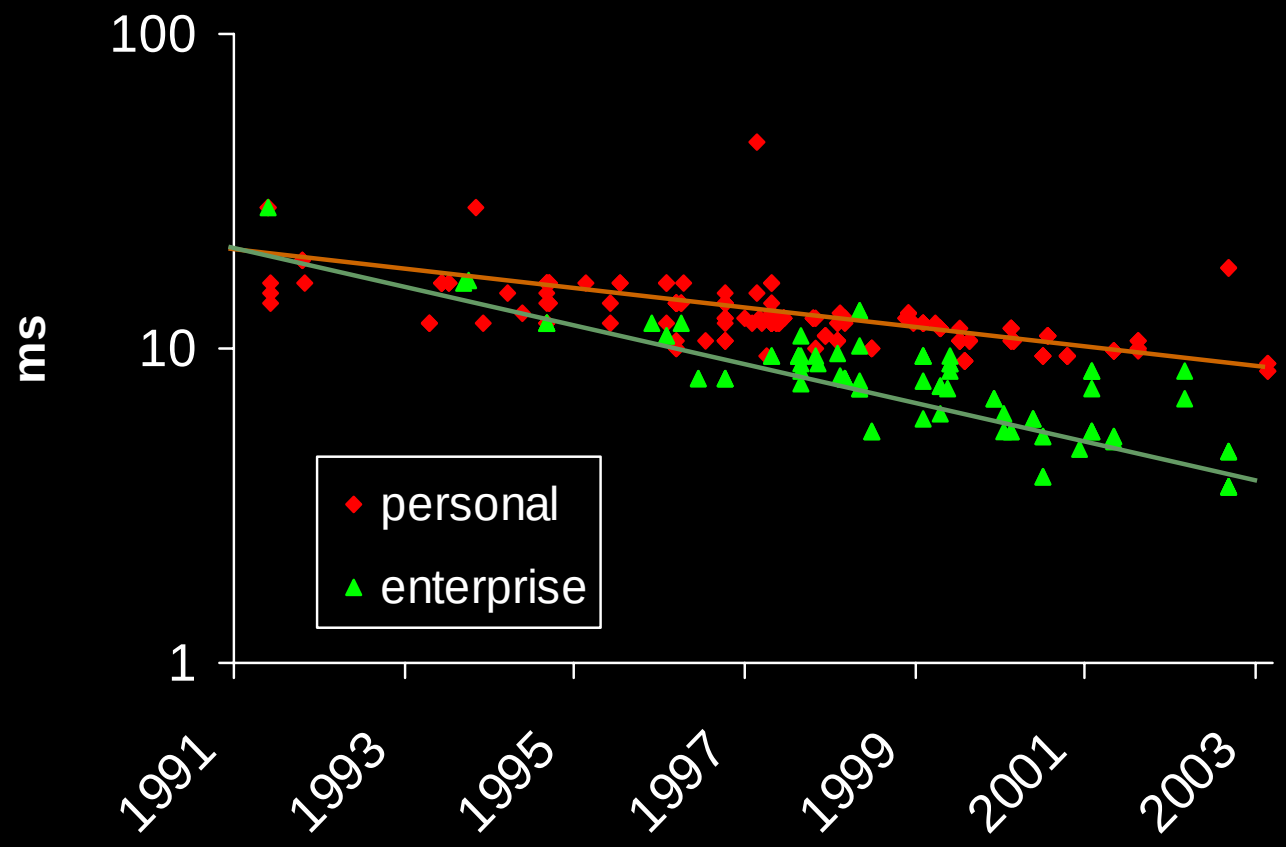

product information for Seagate disc drives introduced since 1991

Rotational vibration

- rotation of one drive affects neighboring drives
- personal drives not designed to reject the external energy
- manifests itself as a *performance* problem

- state of the art cabinets have 15 adjacent drives
- measurement of cabinets shows vibration up to 45 rad/s^2 (best cabinets at 5 rad/s^2)

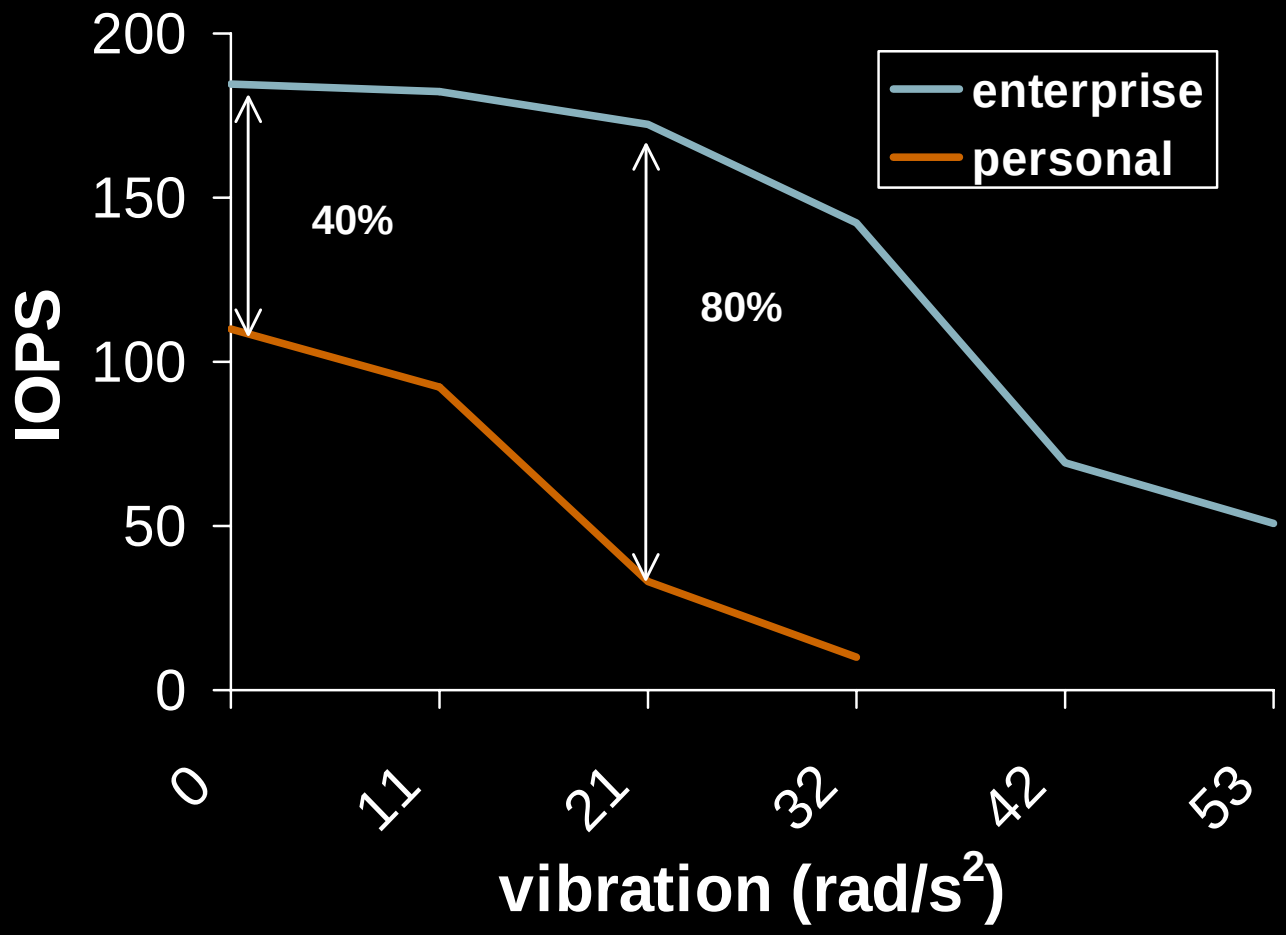

comparison of Cheetah 18LP and Barracuda III

Outline

Marketplace
- personal
- enterprise

Mechanics
- choices & comparison

Electronics
- this is the *only* place where interface matters directly

Performance
- the direct impact

Summary

Electronics comparison

Servo processor – track following

- more tracks require more processing
- enterprise drives will use two processors

Data mover ASIC – for all common-case data transfer

- more complex interface requires more gates

Program & data memory – every KB counts here!

- more complex interface requires more RAM & flash
- command queuing, multi-host support, parallel tasks
  - requires more code, more data space
- we only put in the functions users *need* and *use*

Command queuing

- enterprise has a much larger improvement due to more sophisticated queuing
  - requires more memory
  - and more code
- note – queuing benefit on writes would be even larger without (unsafe) caching

comparison of Cheetah 73LP and Barracuda IV

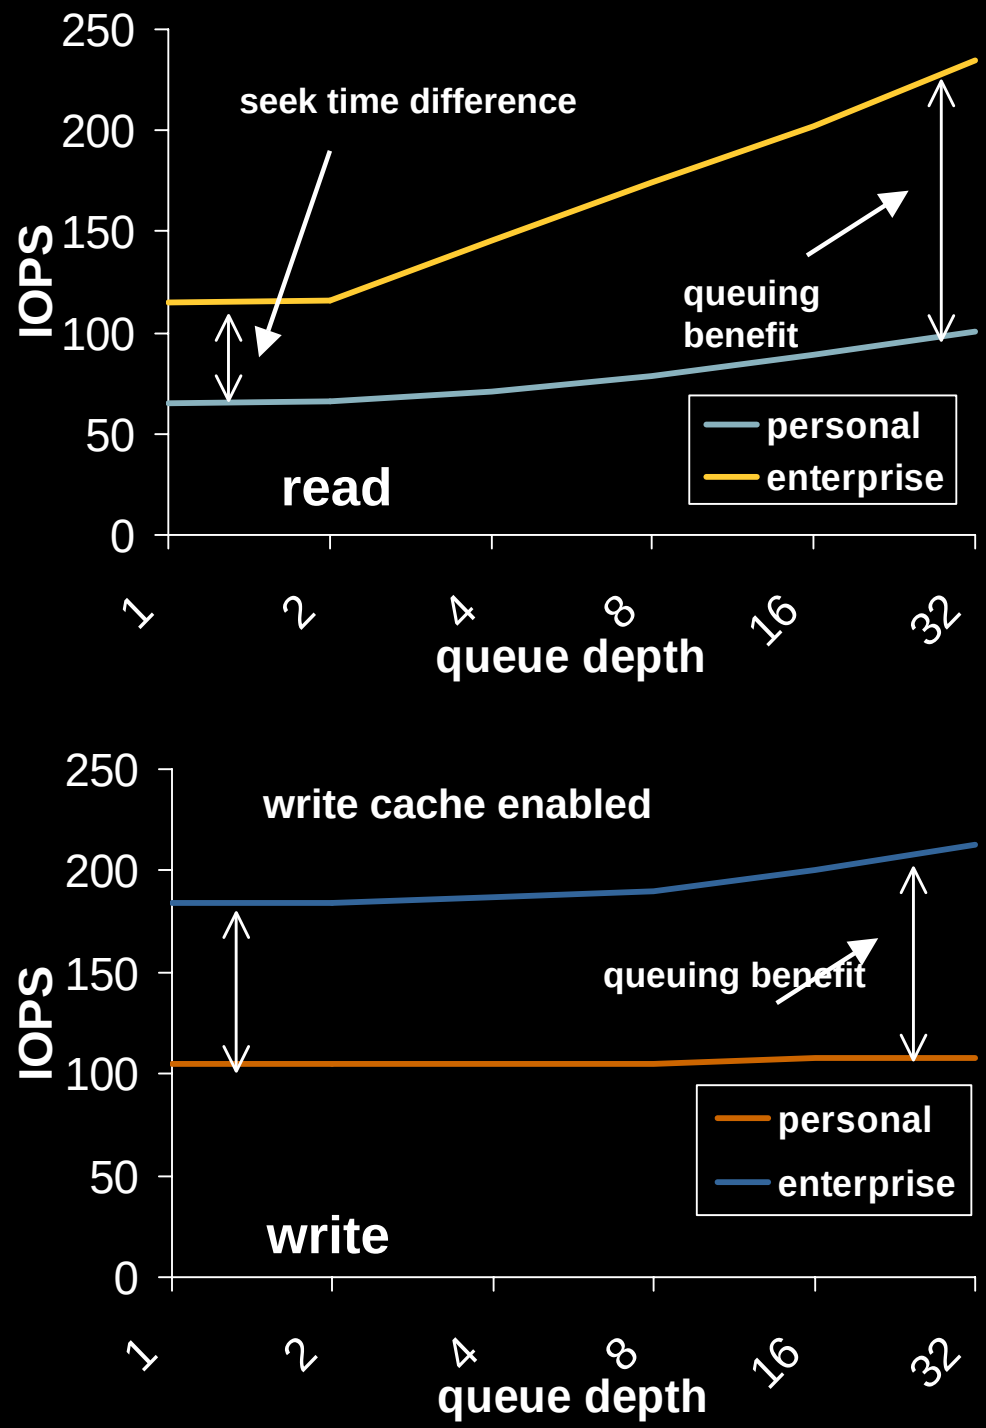

Outline

Marketplace
- personal
- enterprise

Mechanics
- choices & comparison

Electronics
- this is the *only* place where interface matters directly

Performance
- the direct impact

Summary

Performance comparison

		cap	rpm	seek	density	dia	p	ext bw
UltraStar 36 LZX	SCSI	36 GB	10000	4.9 ms	7.0 Gb/in^2	84 mm	6	36 MB/s
DeskStar 75	ATA	30 GB	7200	8.5 ms	11.0 Gb/in^2	95 mm	2	37 MB/s

slower spindle

higher areal density

equivalent areal density

larger platters

personal bandwidth slightly higher

enterprise bandwidth much higher

		cap	rpm	seek	density	dia	p	ext bw
UltraStar 36Z15	SCSI	36 GB	15000	4.1 ms	10.7 Gb/in^2	65 mm	6	53 MB/s

faster spindle

faster seek

smaller platters

more platters

Summary and conclusions

Drives are designed from the ground up to meet a specific set of usage characteristics

- more sophisticated than just $ / GB

If you want to understand the state-of-the-art

- make sure you look at enterprise drives
- and make sure you are comparing apples to apples

There is room for a wider variety of models

- tell us what points are worthwhile

More Details?

Just ask.

dave.b.anderson@seagate.com
james.e.dykes@seagate.com
erik.riedel@seagate.com

Detail Slides

Disc Drive Internals

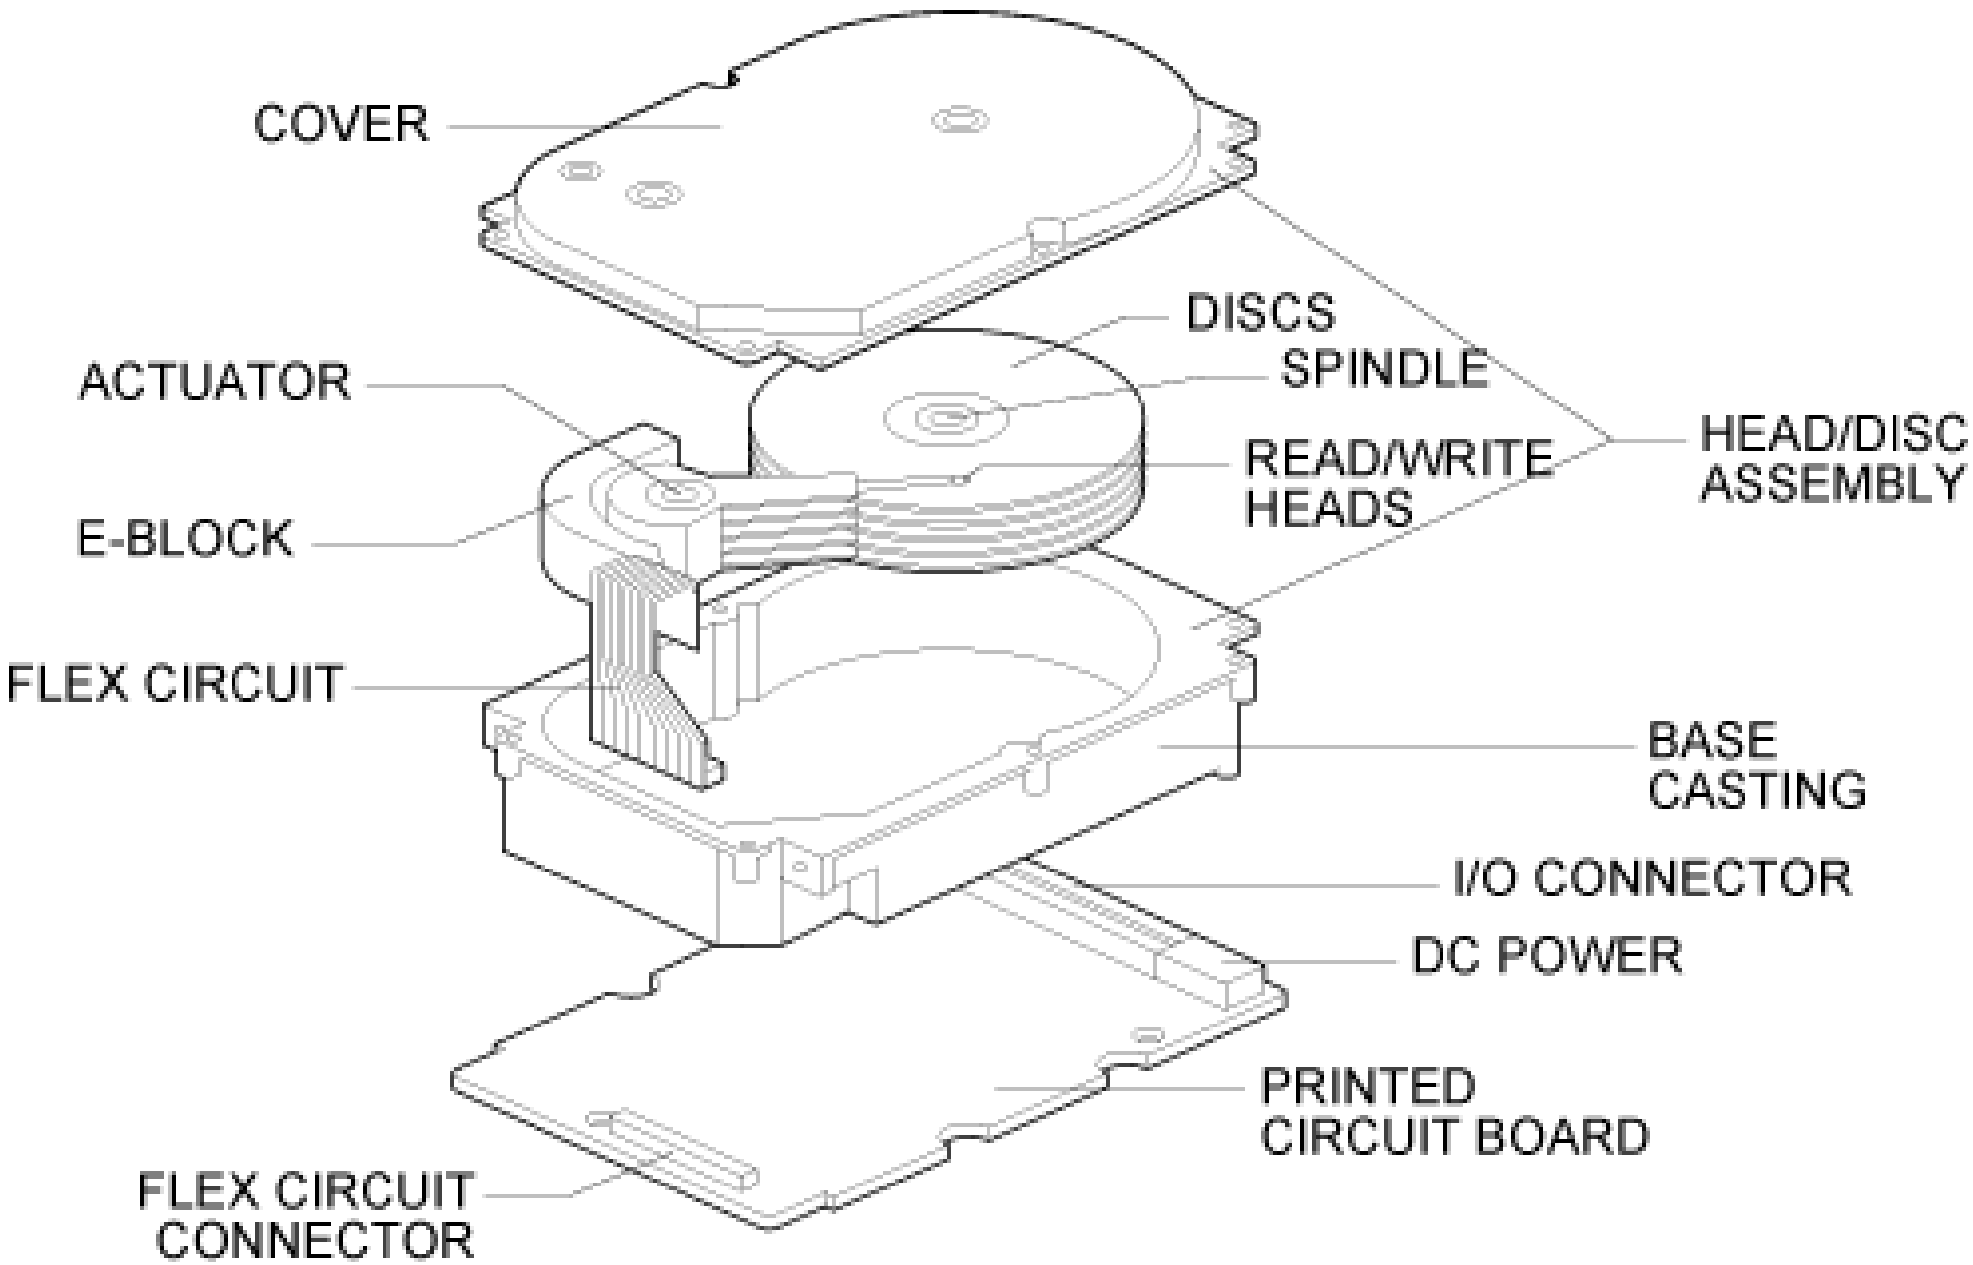

Number of platters

- trend toward depopulated drives, as users trade capacity for performance and reliability
- even to odd-numbered head and single head drives, to save head costs

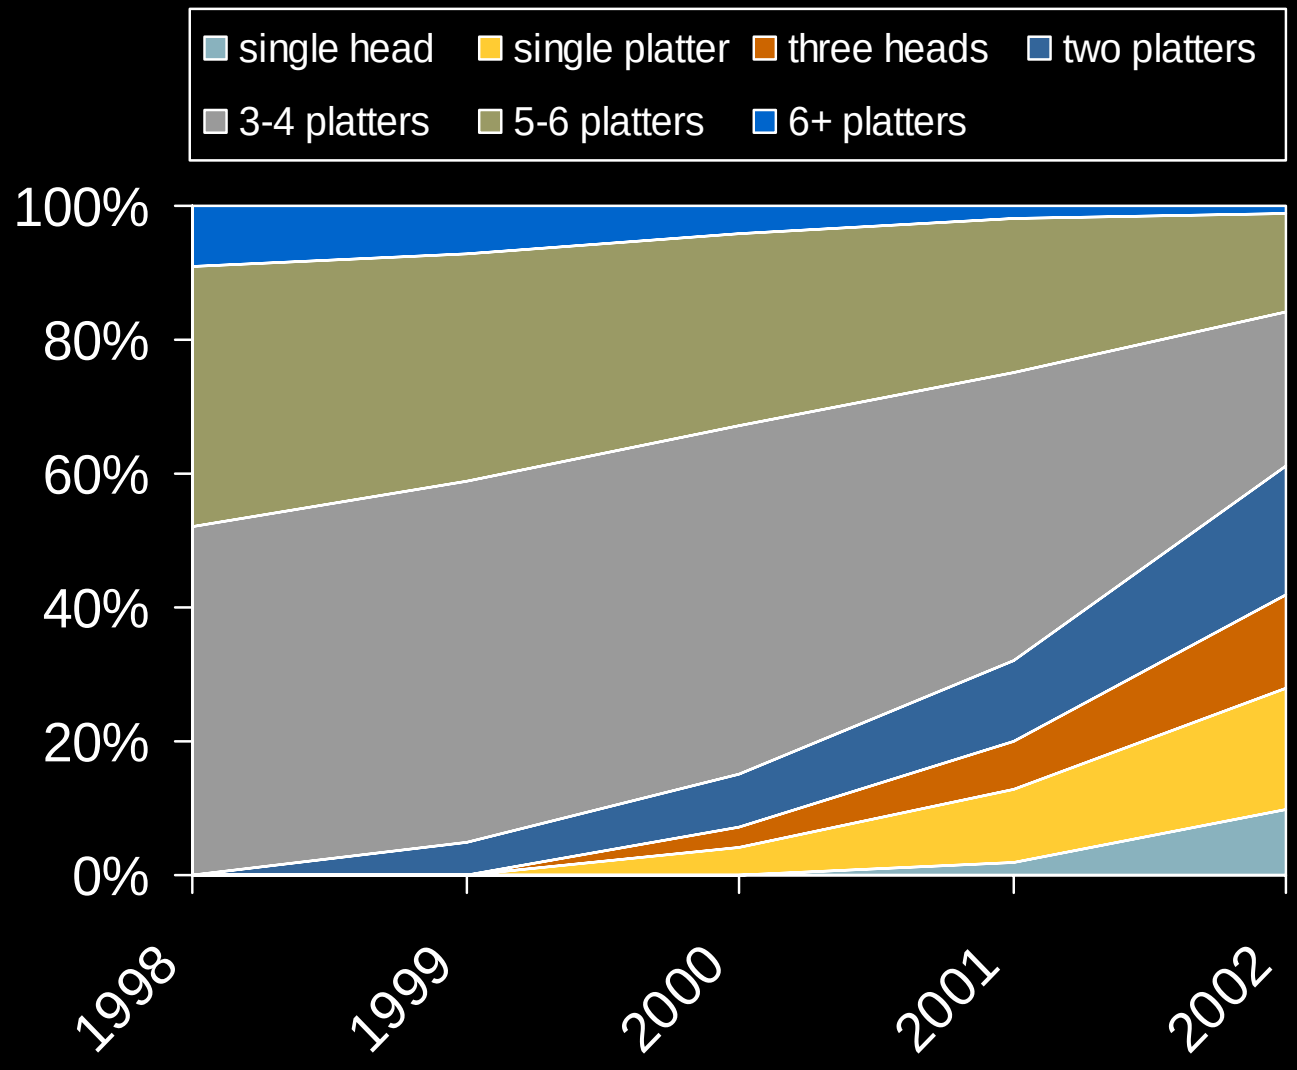

product information for entire industry since 1998

Areal density vs. linear density

- most of the recent improvement due to higher tracks per inch (tpi)
- linear density, bits per inch (bpi) is much higher, but growing more slowly

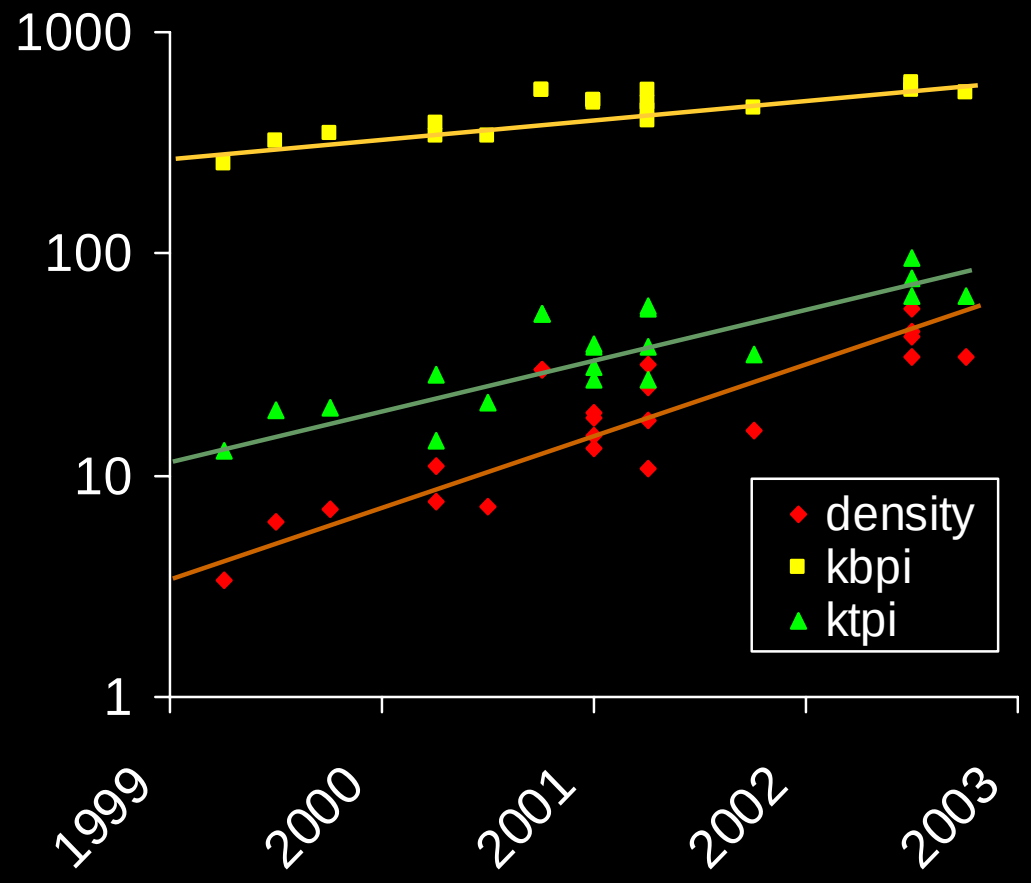

product information from the Appendix of "SCSI vs. ATA – more than an interface" in 2nd FAST Conference, April 2003

Linear density

- enterprise lags personal by a little
  - takes longer to get the signal processing into the right tolerances
- data rate and capacity are the drivers in personal

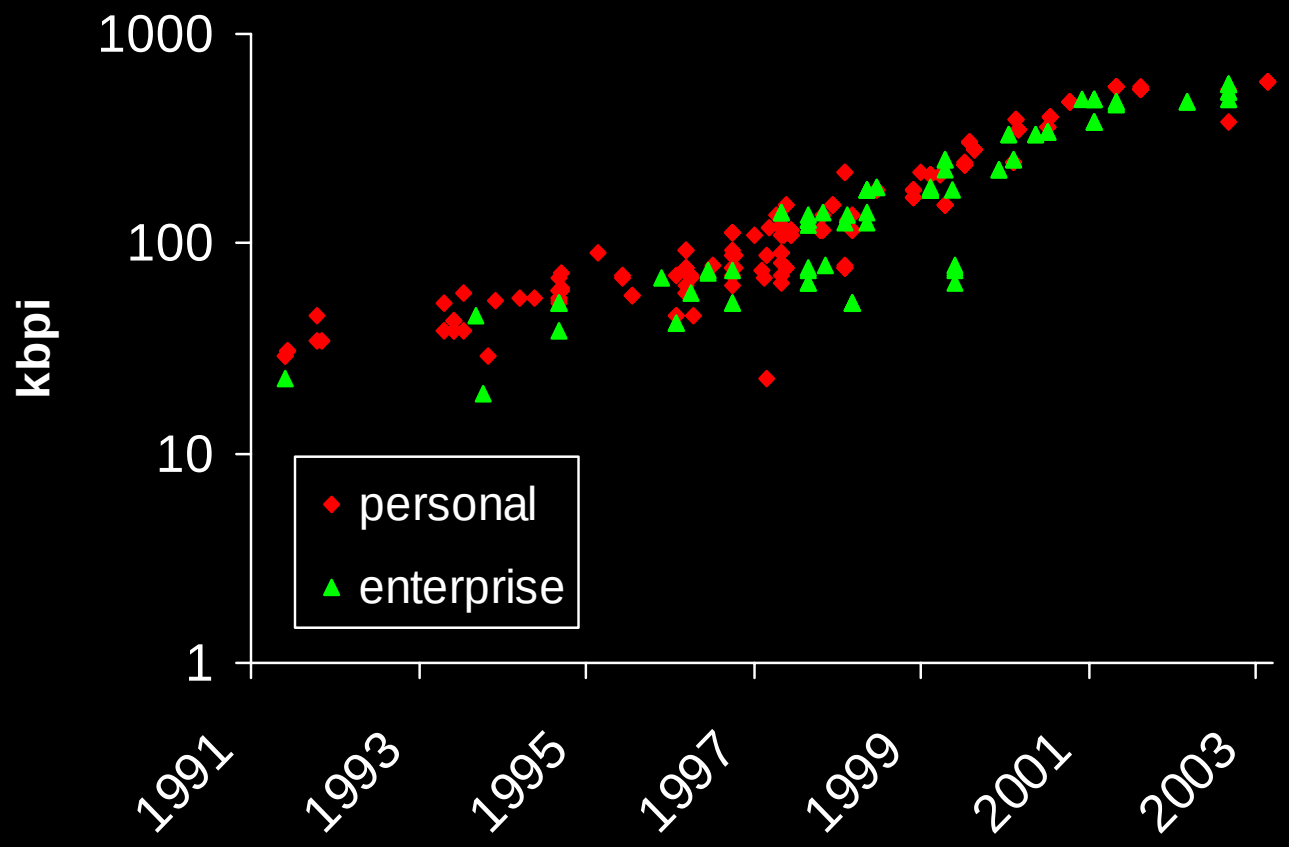

product information for Seagate disc drives introduced since 1991