Digital Storage

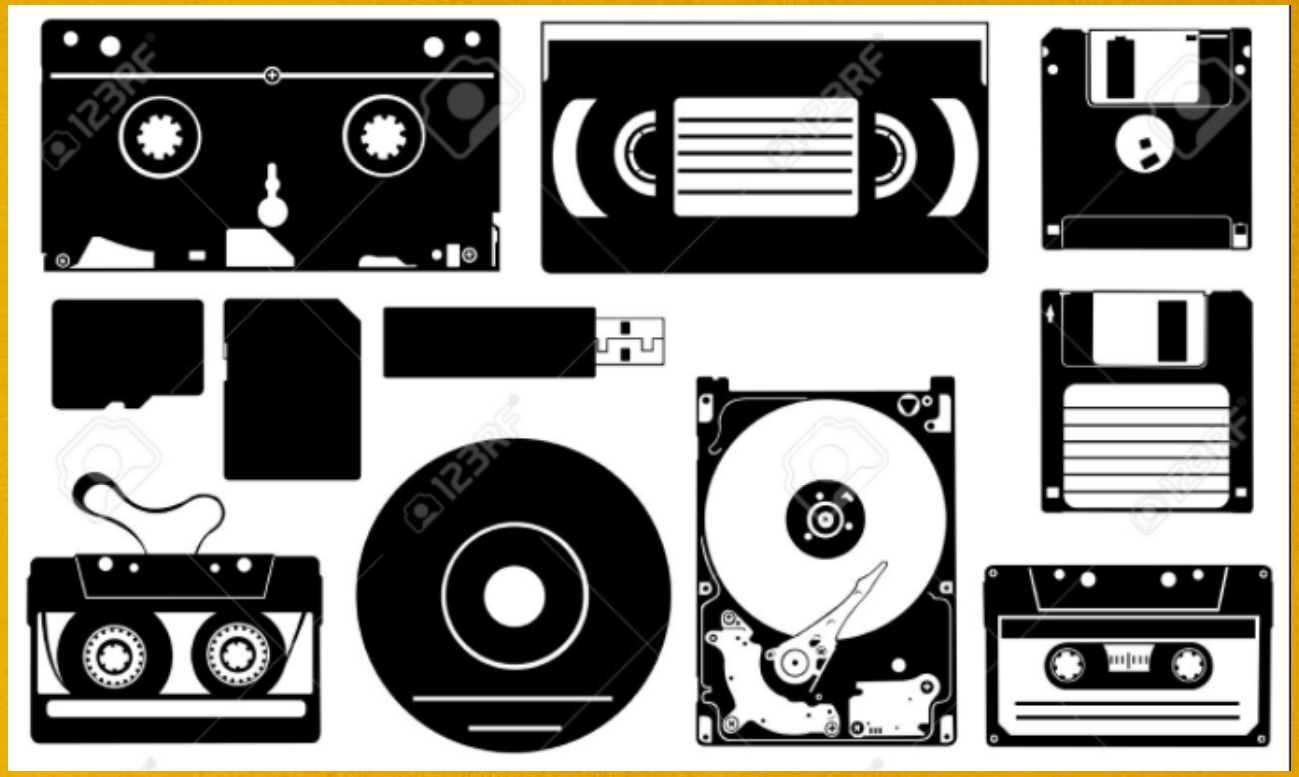

Lorrin R. Garson
OPCUG & PATACS
July 16, 2016

Topics Covered…

- An overview of digital storage devices
  - Discussion of consumer products, not commercial/enterprise hardware
  - History and progress
  - Some stories from 40 years experience
  - Pros and cons of various devices
  - Prices for common devices
  - What the future holds
  - Have a bit of fun!

First a little history…

Earliest Know Digital Storage

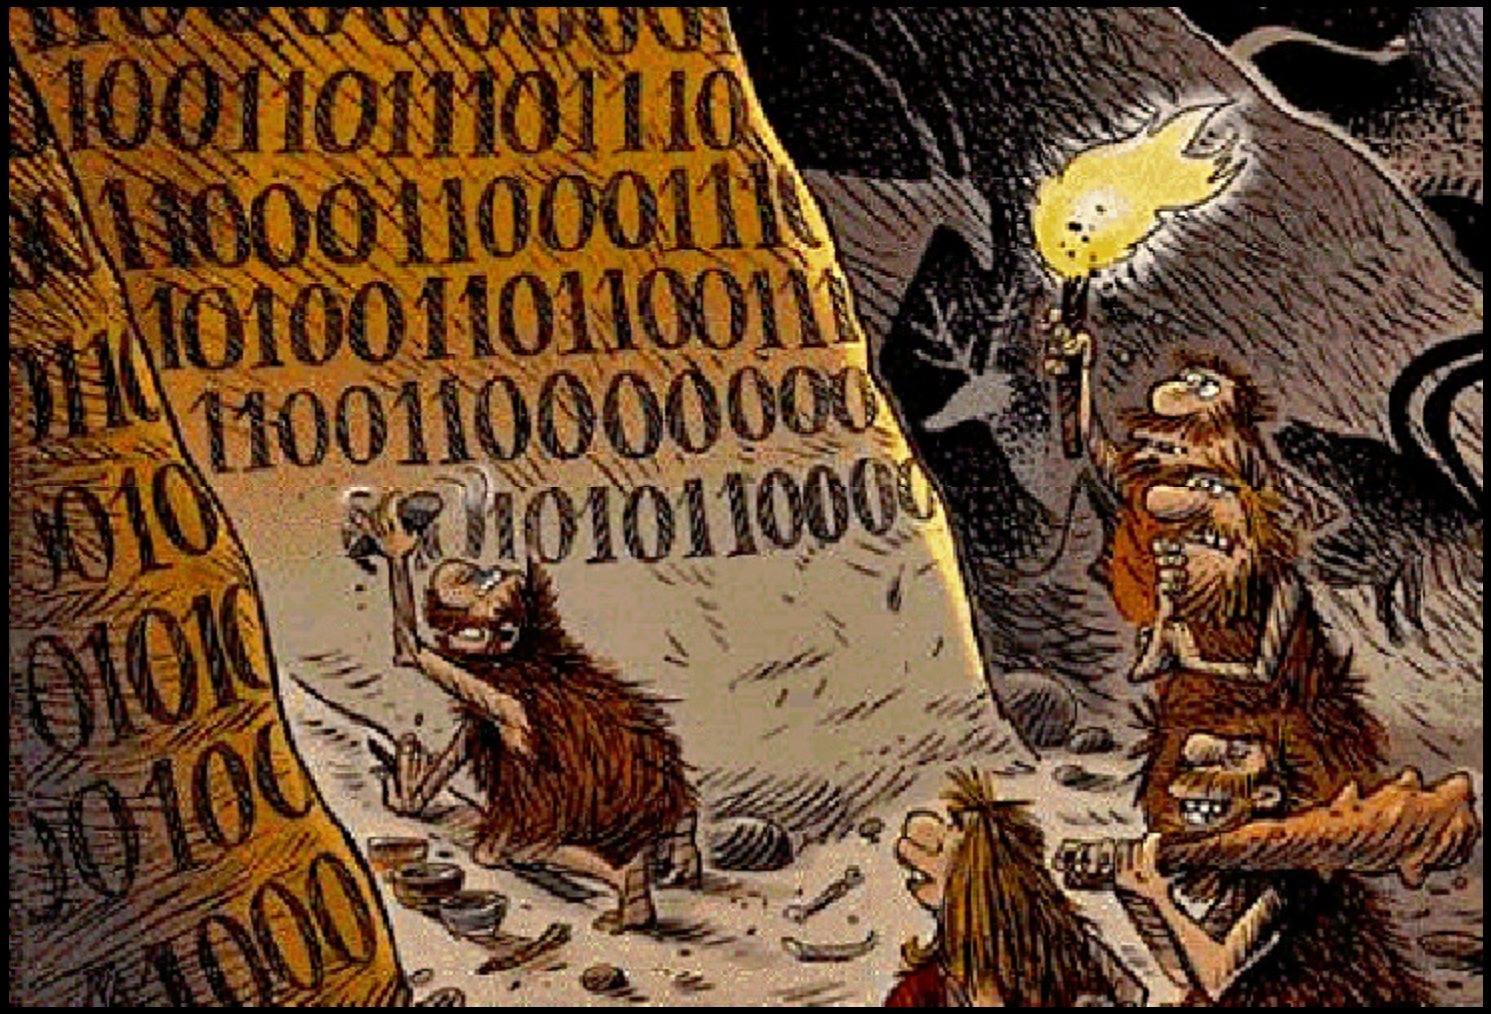

So…

How many boulders do you need?

- How much digital storage do you need?
- What kind?

My Computer

- Apple iMac 21.5-inch (bought Oct 2014)
- 3.1 GHz Intel Core i7 CPU
- 16 GB RAM
- NVIDIA GeForce GT 750M 1024 MB graphics card
- **500 GB internal solid state drive**
- **2-TB External USB-3 hard disk drive**
- **Two 2-TB USB-3 hard disk drives for backup**

How My Storage Is Used

- 500 GB Solid state drive
  - 154 GB used (31%)
  - 346 GB available (69%)
- Use of 154 GB
  - **100 GB user "Lorrin"**
  - 9 GB Library (for user "Lorrin")
  - 37 GB Applications
  - 8 GB System (OS X El Capitan)
- 2-TB External hard disk drive
  - 0.57 TB used (29%)
  - 1.4 TB available (71%)

My 100 GB of Stuff

- 55 GB "Documents"
  - 40 GB presentations
  - 5 GB computer information and documentation
  - 10 GB various documents (in 100's folders)
- 31 GB Music
- 3 GB Encrypted "Vault"
- 11 GB Miscellaneous

- 0.57 TB photos (~150,000 pictures)*

*Stored on an external hard disk drive

Kinds of Mass Storage

- Magnetic
  - Hard disk drives (HDD)
  - Tape
  - Floppy drives (8", 5.25", 3.5")
  - Zip drives ("super floppy")
- Optical (CD & DVD)
- Solid state devices
  - Solid state drives (SSDs)
  - Thumb drives
  - Memory cards

Types of Mass Storage

- Drive types:
  - Internal
  - External
  - Network Attached Storage (NAS)
  - Servers

Cloud Storage

- $204 billion/year business
- 17% growth expected in 2016
- Use dominated by businesses

URL ← Active link

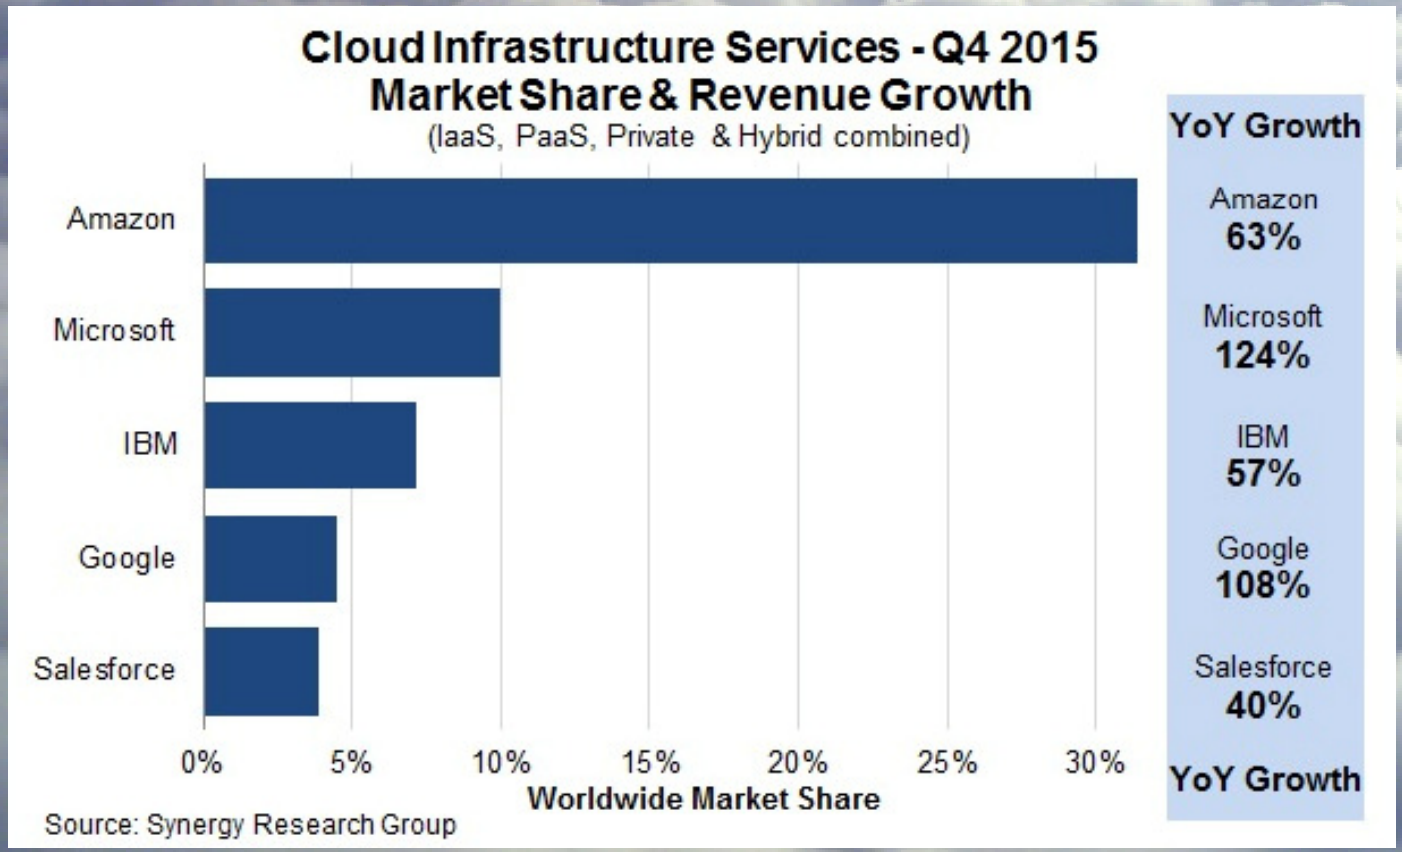

From URL


ARE YOU SURE THIS IS HOW WE GET DATA INTO THE CLOUD?
IT

Hard Disk Drives (HDDs)

- Introduced by IBM in 1956 (3.75 MB)
- Common form factors: 2.5-in & 3.5-in
- Common interface for internal drives:
  - SATA (serial ATA)
  - PATA (IDE & EIDE)

- Manufacturers:
  - Western Digital
  - Seagate
  - Toshiba
  - LaCie (Thunderbolt)

HDDs: Factors Impacting Price and Performance

- Rotation speed (RPM):
  - 5,400
  - **7,200**
- Seek time (5 to 15 ms)
- Cache (8 to 128 MB)
- Connection type

Hard Disk Drives: Prices*

*Prices as of February 24, 2016

Internal Hard Drives

Size (TB)	Low Price	Median Price	High Price
1	$35	$65	$110
2	60	95	150
3	80	100	195
4	110	170	250

Sales of HDD dropped by ~10% from 2014 to 2015

Hard Disk Drives: Prices

External Hard Drives

Size (TB)	Low Price	Median Price	High Price
1	$55	$59	$150
2	78	90	180
3	90	115	145
4	106	130	210

Not quite ancient history…

75-MB Disk Drive Alpha Microsystem

- **1979**
- $12,500 ($40,800 in 2016)
- Storage costs*
  - ✓ $544 per MB
  - ✓ $544,000 per GB
  - ✓ $544,000,000 per TB
- 2016
  - ✓ 3 TB Seagate SATA Internal Drive — $85.59
  - ✓ $28.59 per TB

*2016 dollars

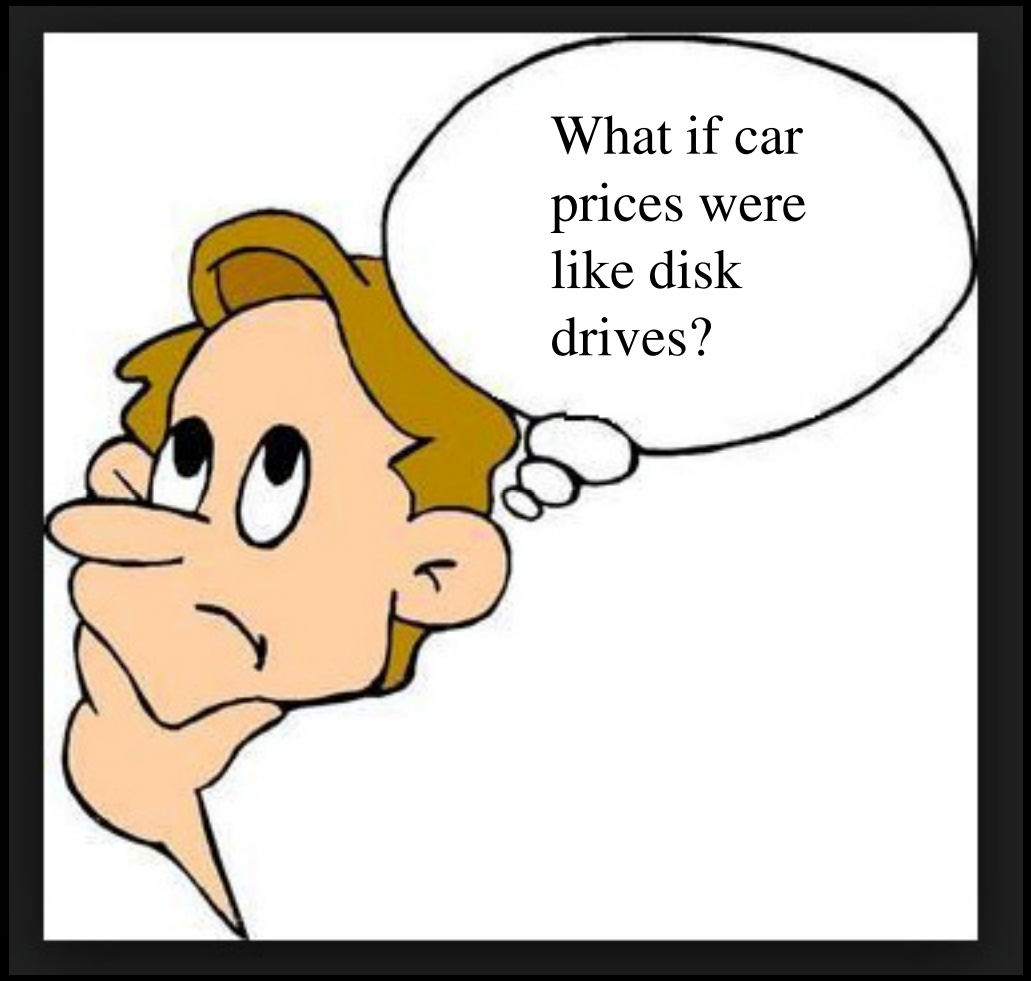
What if car prices were like disk drives?

1979 Trabant 601 Limousine

1979 Average car price $18,829*

2016 Volkswagen CC

2016 Average car price $33,560

*$5,770 1979 dollars

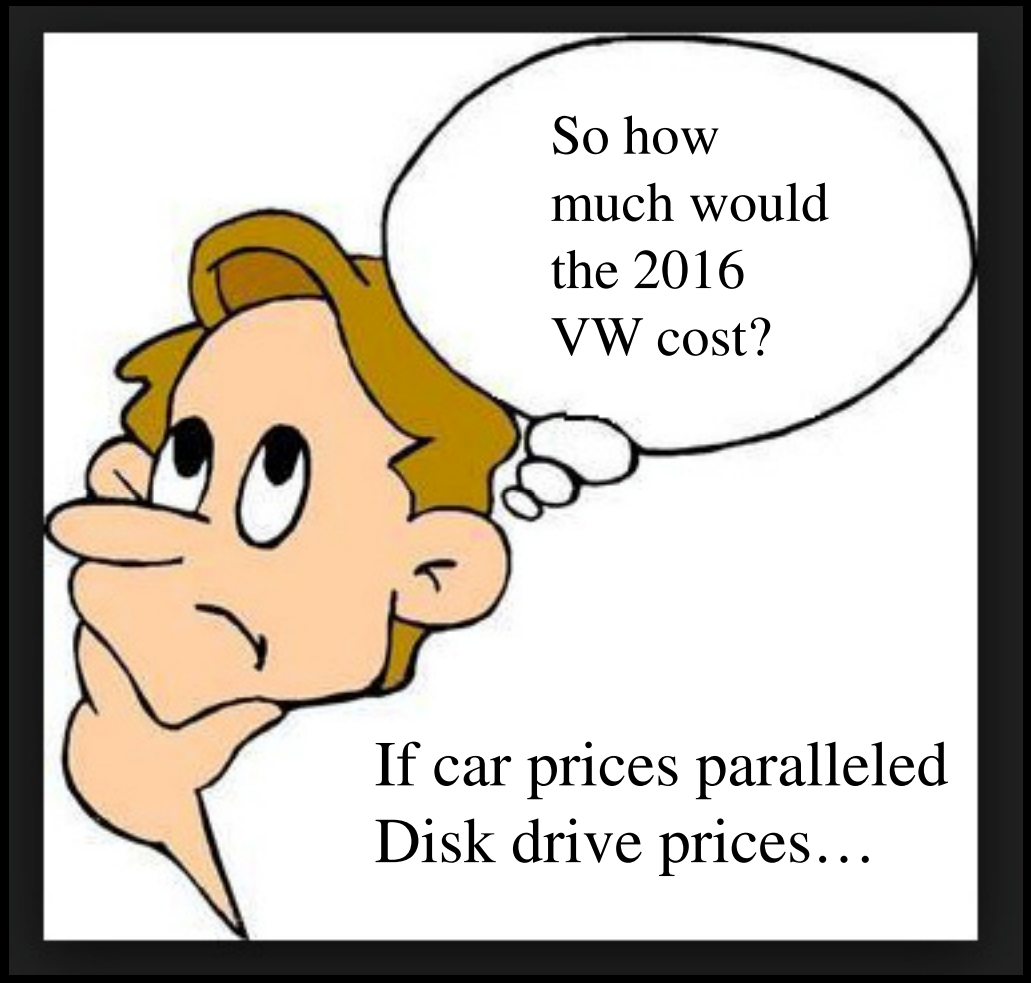
So how much would the 2016 VW cost?
If car prices paralleled Disk drive prices…

$0.00176

or…

2¢ per dozen

Life Expectancy
Hard Disk Drives (HDD)

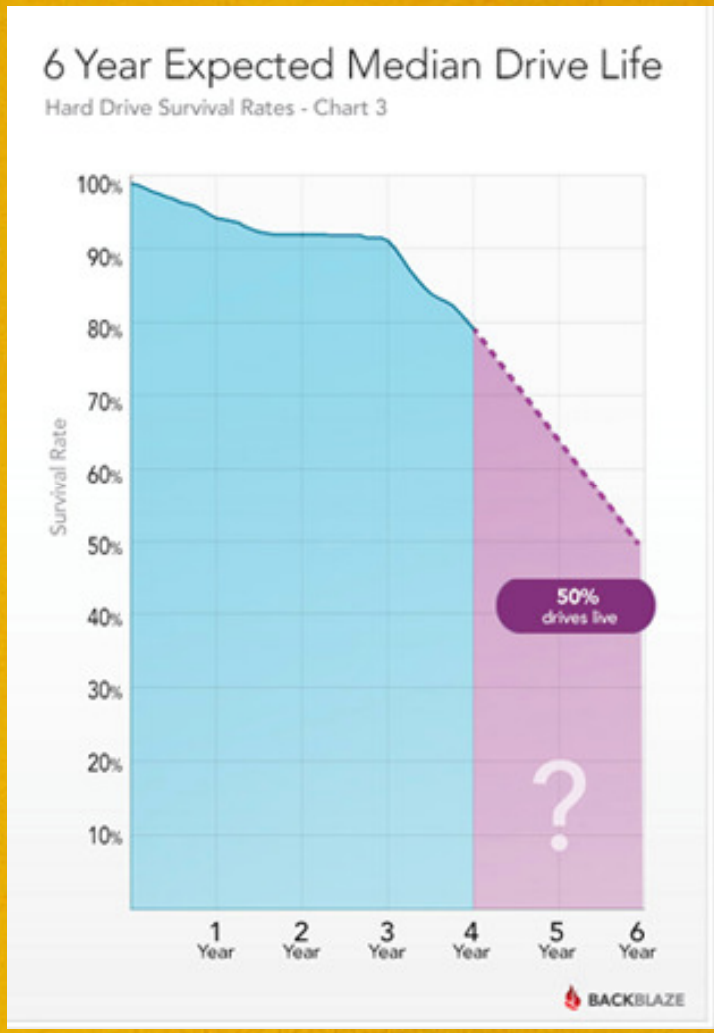

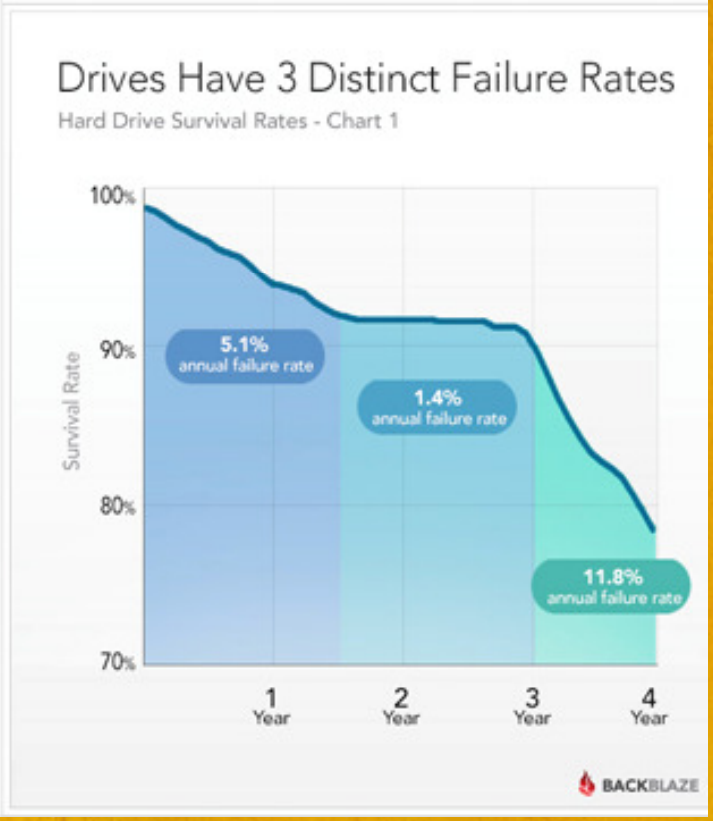

See URL URL

Failure Rate by Manufacturer Hard Disk Drives (2013-2015)

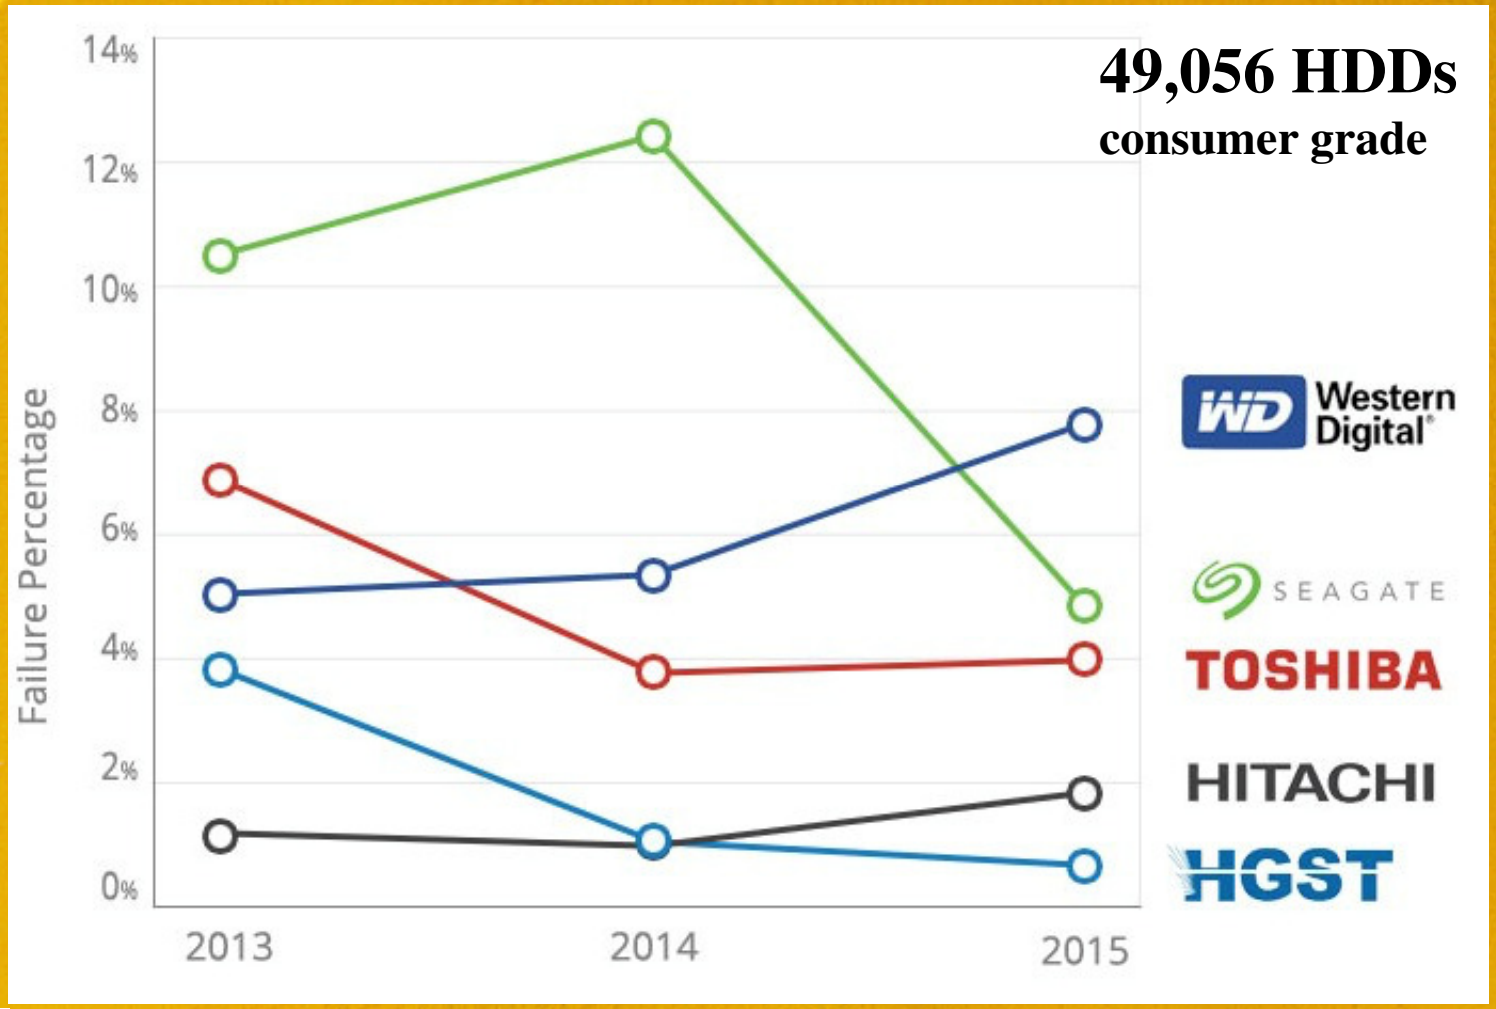

See URL

In 2015 Hitachi HDD became part of HGST, which is owned by WD

Hardware Interfaces
Internal Hard Disk Drives (HDD)

- Internal Storage:
  - **PATA** (**P**arallel **A**dvanced **T**echnology **A**ttachment)*
  - **SATA** (**S**erial **A**dvanced **T**echnology **A**ttachment)

Interface	Speed (Gb/Sec)		
PATA	0.66	1	1.3
SATA	1.5	3	6


SATA

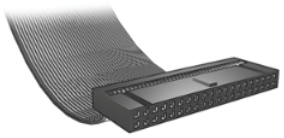
PATA

Hardware Interfaces
External Hard Disk Drives* (HDD)

- External Storage:
  - USB (1.0, 1.1, 2.0, 3.0, and C)
  - eSATA
  - Gigabit Ethernet
  - FireWire (abandoned by Apple)
  - Thunderbolt (to be abandoned by Apple?)

*Solid state drives also use these interfaces, but there is a need to enhance the standards for SSDs

USB

Interface	Description	Transfer Rate	Best For	Most Popular Environment
USB 1.0/1.1	Hot-swappable*	12 Mb/s maximum burst transfer rate	Connecting to different systems quickly and easily	PCs
USB 2.0	Hot-swappable*	480 Mb/s maximum burst transfer rate	Connecting to different systems quickly and easily	PCs
USB 3.0	Hot-swappable* full-duplex asynchronous protocol	4.8 Gb/s maximum burst transfer rate	High-def. video cameras, mobile media players and phones, and terabyte storage devices	PCs

USB-C type 10 Gb/s

USB Connectors

eSATA

Interface	Description	Transfer Rate	Best For	Most Popular Environment
eSATA	Fast and physically secure connection for external hard drives	Up to 6 Gb/s maximum sustained transfer rate—nearly equivalent to internal drive performance	Transferring large amounts of data frequently or for audio/video programs	PCs

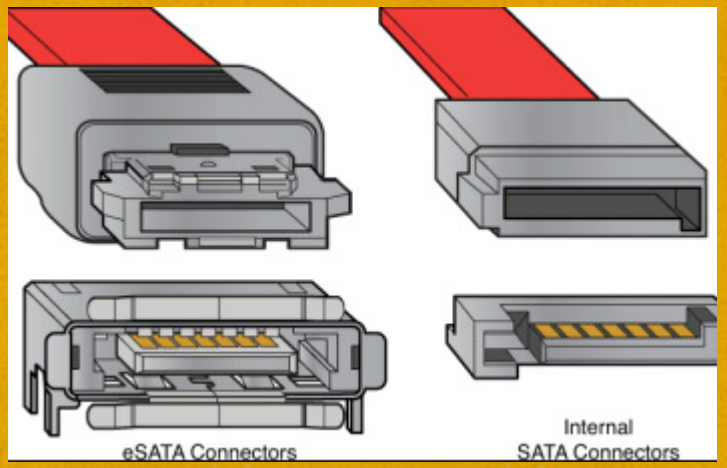

Gigabit Ethernet

Interface	Description	Transfer Rate	Best For	Most Popular Environment
Gigabit Ethernet	Can attach a device to a system's network adapter, router, switch, or hub port	1,000 Mb/s maximum sustained transfer rate—ten time faster than Fast Ethernet	Remote file access and sharing large amounts of data and/or large audio/video files across a network	PCs

Firewire

Interface	Description	Transfer Rate	Best For	Most Popular Environment
FireWire 400	Hot-swappable*	400 Mb/s maximum sustained transfer rate	Fast transfer of large amounts of data frequently or for audio/video applications like home video editing	Macintosh® computers
FireWire 800	Hot-swappable*	800 Mb/s maximum sustained transfer rate	Multiple-stream digital video and high-resolution digital audio and video applications	Macintosh® computers

FireWire 400 = IEEE 1394A
FireWire 800 = IEEE 1394B

Thunderbolt

Interface	Description	Transfer Rate	Best For	Most Popular Environment
Thunderbolt	Supports daisy chaining of multiple devices through a single connector so you can multitask at Thunderbolt technology speed without having to use a switch or hub	Two bidirectional channels (per port), with theoretically up to 10 Gb/s of bandwidth	Fast transfer of large amounts of data. Multiplestream digital video and highresolution digital audio and video applications	Macintosh® computers

USB-C type @ 10 Gb/s likely to dominate

IEEE 1394; aka iLink (Sony) &
Lynx (Texas Instruments)

Solid State Drives (SSDs)

- Very fast
- No moving parts
  - Quiet
  - Low power consumption—less heat
- Can only sustain a finite number of writes—then they wear out
- Wear-leveling prolongs life
- Relatively expensive 7X HDD
- Data deteriorate left unplugged
- For more information see

Comparing HDDs to SSDs

Factor	HDD	SSD	Comments
Price	$0.049/GB	$0.343/GB	SSD cost 7X HDD
Performance	1	10-20	URL
Life Span	4-7 years	10-50 years	URL URL URL URL
Fragility	1	17 URL	SSD 1/17th as fragile
Common Sizes	1000-3000 GB	120-1000 GB	

For more information see URL URL

Solid State Drives (SSDs): Prices

Internal Solid State Drives

Size (GB)	Low Price	Median Price	High Price
120-128	$43	$56	$93
240-256	65	80	135
480-512	81	147	224
960-1000	239	295	447
2000	600	744	929

A cheap theatrical!

Hybrid Drives (SSHD)

- Combination of HDD and SSD (Fusion Drive)
- Most frequently accessed files are stored on the SSD component
- Small amount of SSD provides significant performance gains
- Performance: HDD < SSHD < SSD

  1 < 5 < 15

- Hybrid drives somewhat more expensive than HDD but much less than SSD
- For more information see  

Network Attached Storage (NAS)

- A data storage server connected to a network that provides file-sharing among PCs on the network
- Hard wired (Ethernet) & wireless (WiFi)
- Considerations:
  - Capacity
  - Performance
  - Price
  - **Backup of NAS itself**
  - Ease of installation and maintenance
  - Noise and heat
  - **Number of bays**

Network Attached Storage (NAS)

Model/Mfg	Capacity	Comments	Ref
Apple Time Capsule	2 or 3 TB	With WiFi router	URL
Buffalo Linkstation 420	2-8 TB	4 models	URL
Netgear ReadyNAS 102	2-12 TB	Many related models	URL
QNAP TS-251	2-8+ TB	Several related models	URL
Synology Disk Station DS214se	4-6 TB	Several models	URL
WD My Cloud	2-16 TB	3 models	URL

All models can be used with PCs & Macs and can be remotely accessed (Internet)

RAID

- RAID = **R**edundant **A**rray of **I**ndependent **D**isks

Common RAID Modes

Mode	Advantage	How It Works	Comments
0	Better performance	Data distributed between multiple disks	Greater risk of data loss
1	Data redundancy	Same data distributed between two disks (mirroring)	Double cost of storage
5	Better performance and data redundancy	Stores parity data to reconstruct actual data	Requires minimum of 3 disks
JBOD	No increase in speed or security	Disk 1 is filled, then disk 2, then disk 3…	Multiple disks act as one Not truly RAID

For more information see 

Wireless Hard Disk Drives

- Type of NAS drive
- 16 GB to 4 TB
- Priced somewhat higher than HDD
- Some portable, 4-12 hr. battery
- Almost all 802.11bgn WiFi*
- Wide variety of manufacturers
- For more information see

***Apple Time Capsule uses 802.11ac**

Home Servers

- Never caught on in the consumer market
- NAS servers have taken the place of home servers
- In December 2014 Microsoft stopped selling Windows Home Server 2011

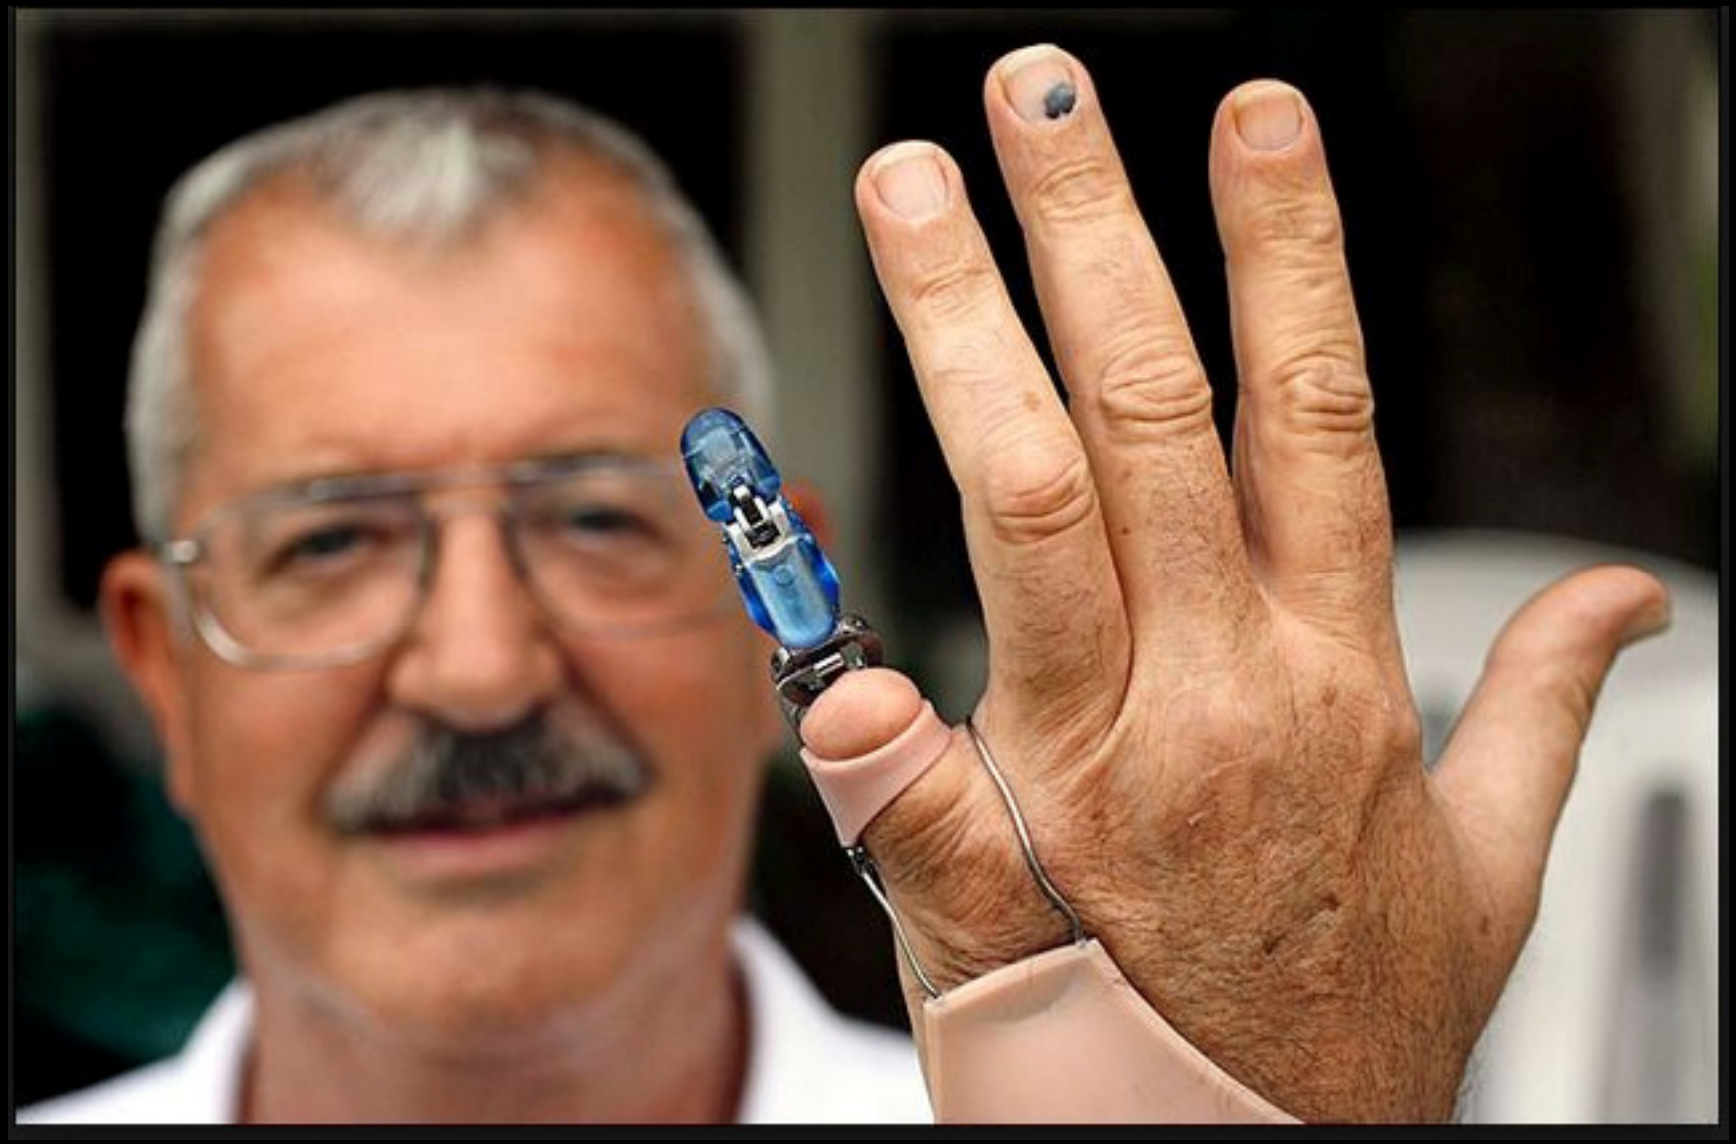

Thumb Drives*

- Many sizes and shapes
- Connectors/interfaces
  - USB-2
  - USB-3
  - USB-C
- Essentially commodity devices
- Drives with encryption much more expensive

*aka flash drive, USB drive, jump drive, pen drive

Prices USB-2 Thumb Drives

Size (GB)	Low Price	Median Price	High Price
8	$2	$5	$8
16	5	7	23
32	8	10	22
64	15	17	40
128	14	28	38
256	13	26	40

Prices USB-3 Thumb Drives

Size (GB)	Low Price	Median Price	High Price
8	$4	$6	$12
16	4	7	12
32	9	11	20
64	15	20	40
128	28	34	100
256	60	70	140

Prices USB-C Thumb Drives

Size (GB)	Low Price	Median Price	High Price
16	$11	$12	$20
32	12	18	30
64	25	28	50
128	50	50	90

More expensive than USB-3 drives

Memory Cards

Memory Cards

- Associated with digital cameras
- Types:
  - SD (**S**ecure **D**igital)
  - SDHC (**S**ecure **D**igital **H**igh **C**apacity)
  - SDXC (**S**ecure **D**igital '**X**tra **C**apacity)
  - CD (**C**ompact **F**lash)
  - Micro SD
  - xD Picture
  - Memory Stick
  - MMC (**M**ulti **M**edia **C**ard)
  - UHS-II

- For more information see URL URL

From  Wikipedia

Name	Abbreviation	Form factor	DRM
PC Card	PCMCIA	85.6 × 54 × 3.3 mm	No
CompactFlash I	CF-I	43 × 36 × 3.3 mm	No
CompactFlash II	CF-II	43 × 36 × 5.5 mm	No
SmartMedia	SM / SMC	45 × 37 × 0.76 mm	No
Memory Stick	MS	50.0 × 21.5 × 2.8 mm	MagicGate
Memory Stick Duo	MSD	31.0 × 20.0 × 1.6 mm	MagicGate
Memory Stick PRO Duo	MSPD	31.0 × 20.0 × 1.6 mm	MagicGate
Memory Stick PRO-HG Duo	MSPDX	31.0 × 20.0 × 1.6 mm	MagicGate
Memory Stick Micro M2	M2	15.0 × 12.5 × 1.2 mm	MagicGate
Miniature Card		37 × 45 × 3.5 mm	No
Multimedia Card	MMC	32 × 24 × 1.5 mm	No
Reduced Size Multimedia Card	RS-MMC	16 × 24 × 1.5 mm	No
MMCmicro Card	MMCmicro	12 × 14 × 1.1 mm	No
P2 card	P2		No
Secure Digital card	SD	32 × 24 × 2.1 mm	CPRM
SxS	SxS		No
Universal Flash Storage	UFS		Unknown
miniSD card	miniSD	21.5 × 20 × 1.4 mm	CPRM
microSD card	microSD	15 × 11 × 0.7 mm	CPRM
xD-Picture Card	xD	20 × 25 × 1.7 mm	No
Intelligent Stick	iStick	24 × 18 × 2.8 mm	No
Serial Flash Module	SFM	45 × 15 mm	No
μ card	*μ*card	32 × 24 × 1 mm	Unknown
NT Card	NT NT+	44 × 24 × 2.5 mm	No
XQD card	XQD	38.5 × 29.8 × 3.8 mm	Unknown

Average Price Common SD Type Cards

Size (GB)	SDHC	SDXC
8	$5	
16	$10	
32	$17	$25
64		$32
128		$50
256		$126

Digital Tape

- Introduced by IBM in 1952, Model 726 (2 MB/tape)
- Widely used by businesses with large amounts of data to backup
- 2.5 TB/tape cartridge (6.25 TB with compression)*
- Inexpensive
- Data recovery is slow

*FujiFilm has a prototype that stores 220 TB/cartridge

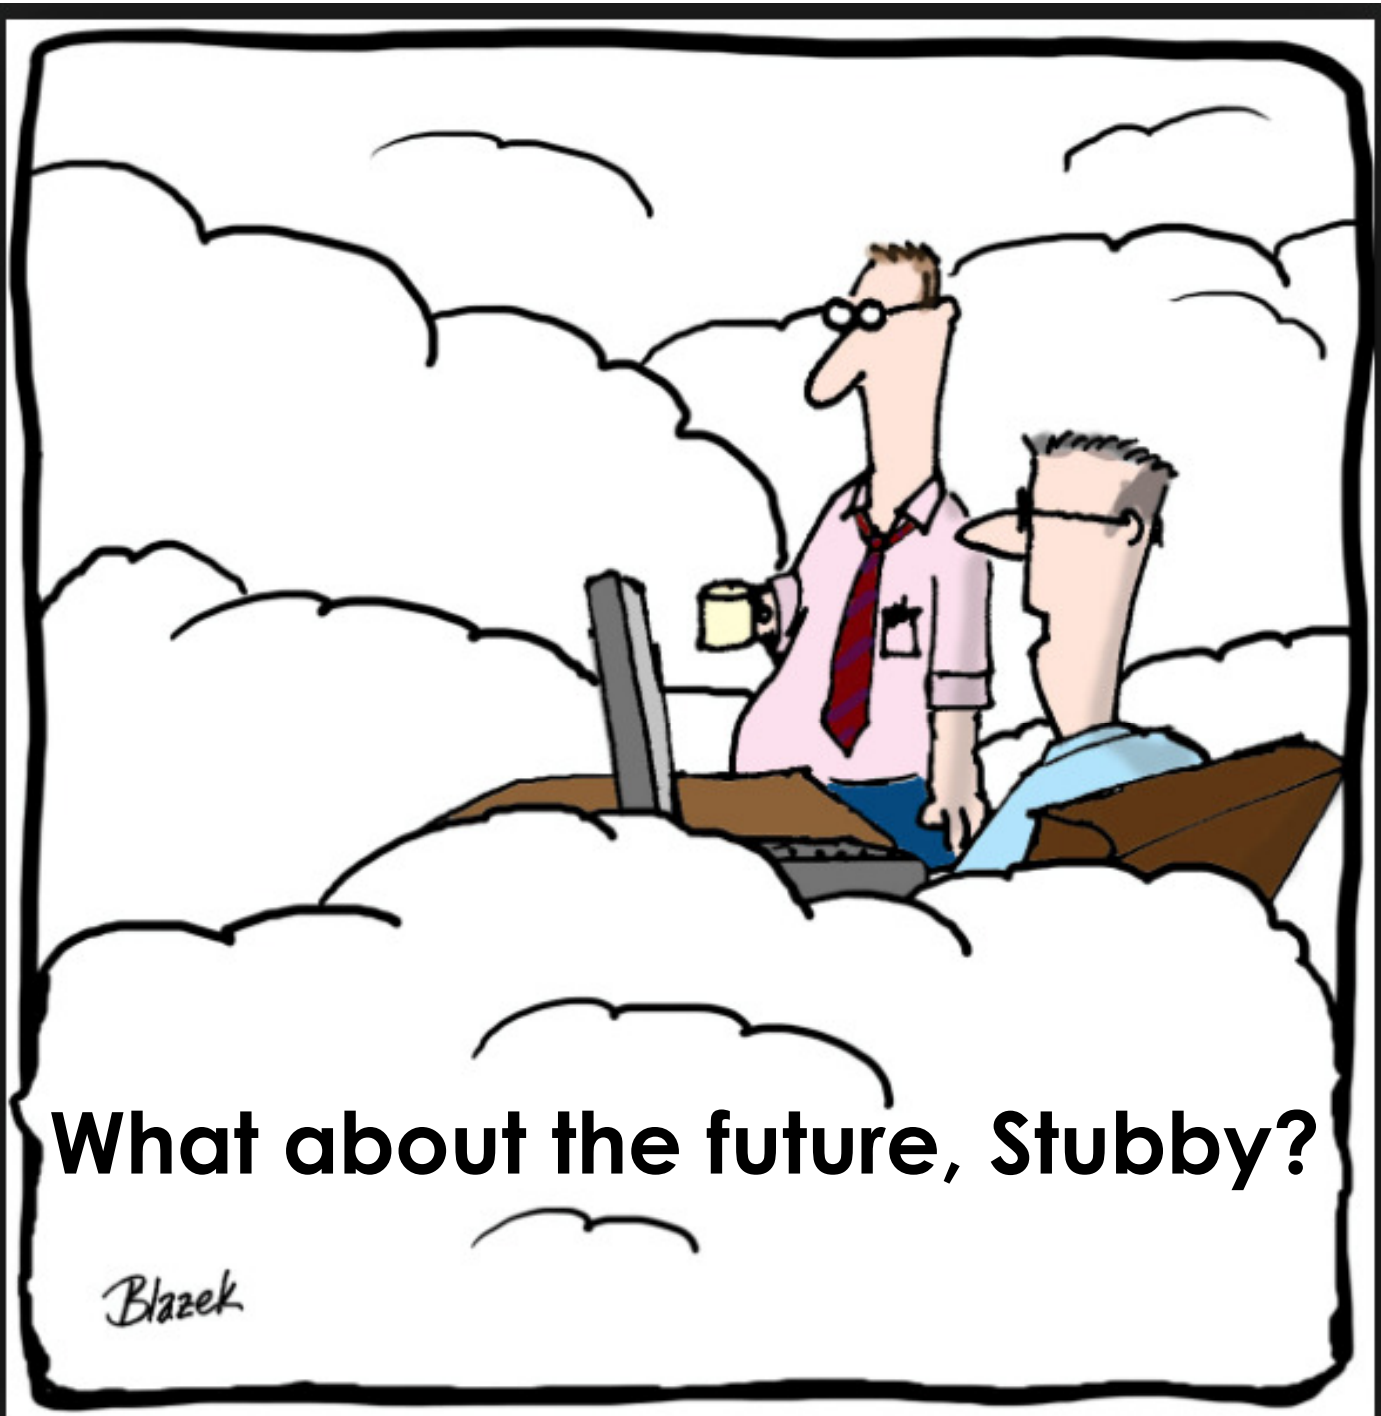
What about the future, Stubby?
Blazek

New Solid State Drives

- From Intel—soon 
  - Super fast—10 TB
  - Using 3D NAND chips
  - Greater durability
- Non-volatile NAND chips from Micron 
  - 1000X faster
  - "Exponentially greater endurance"
- Toshiba and Samsung likely to follow soon
- Interface standards will need upgrading
- For more information see 

What about using DNA for digital storage?

Deoxyribonucleic Acid (DNA)

- The "blueprint of life" which defines who we are 
- Extremely stable
  - 419 million-year-old DNA of bacteria found 
  - Can be stored for a million years 
- Specific molecular bonding can represent binary 0 and 1
- Well-established DNA sequencing processes can retrieve 0s and 1s
- Potential to store offline, archival data

DNA Double Helix

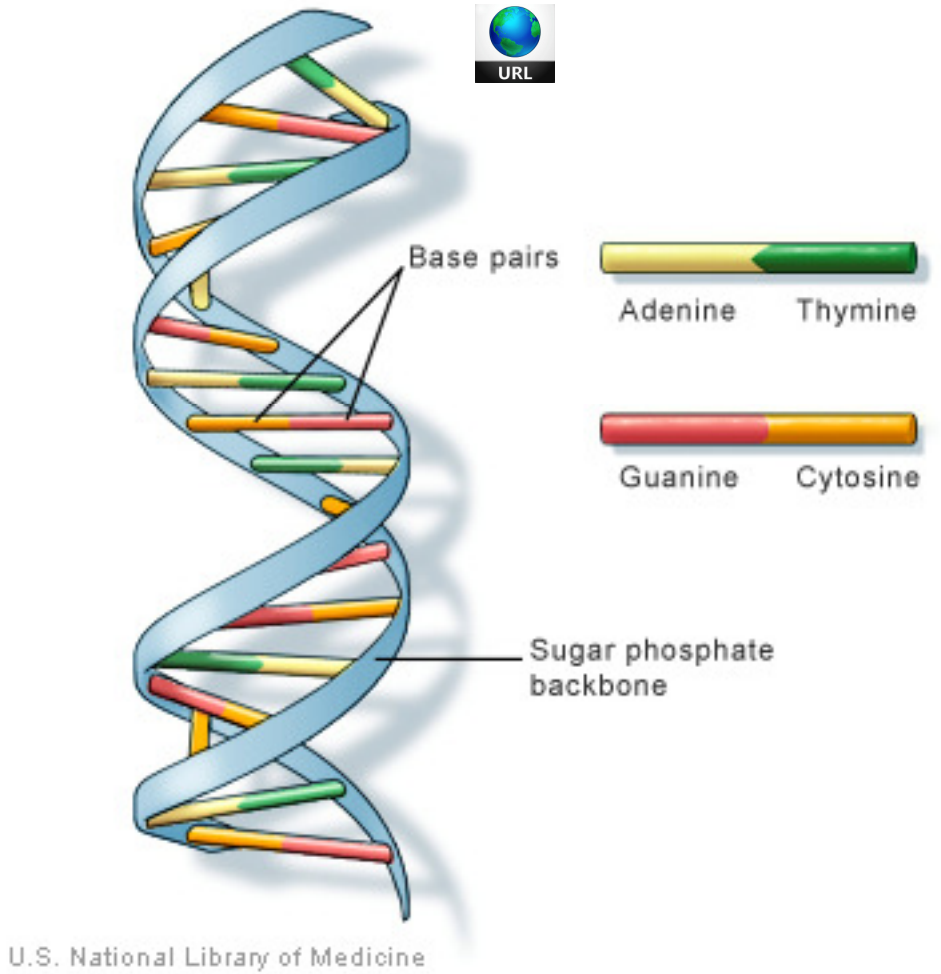

DNA Nucleotide Pairing

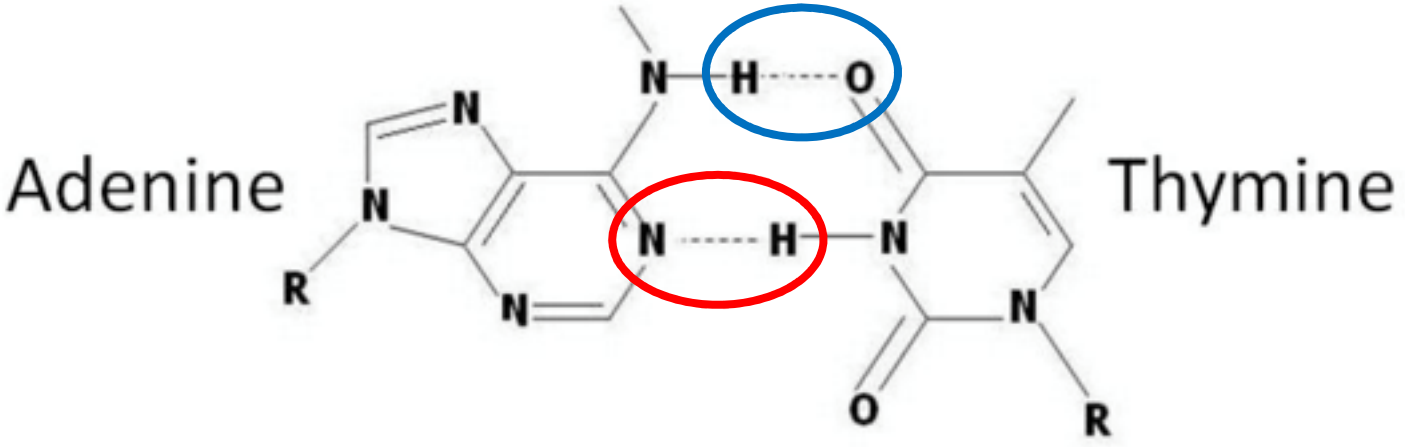

A + T = 1

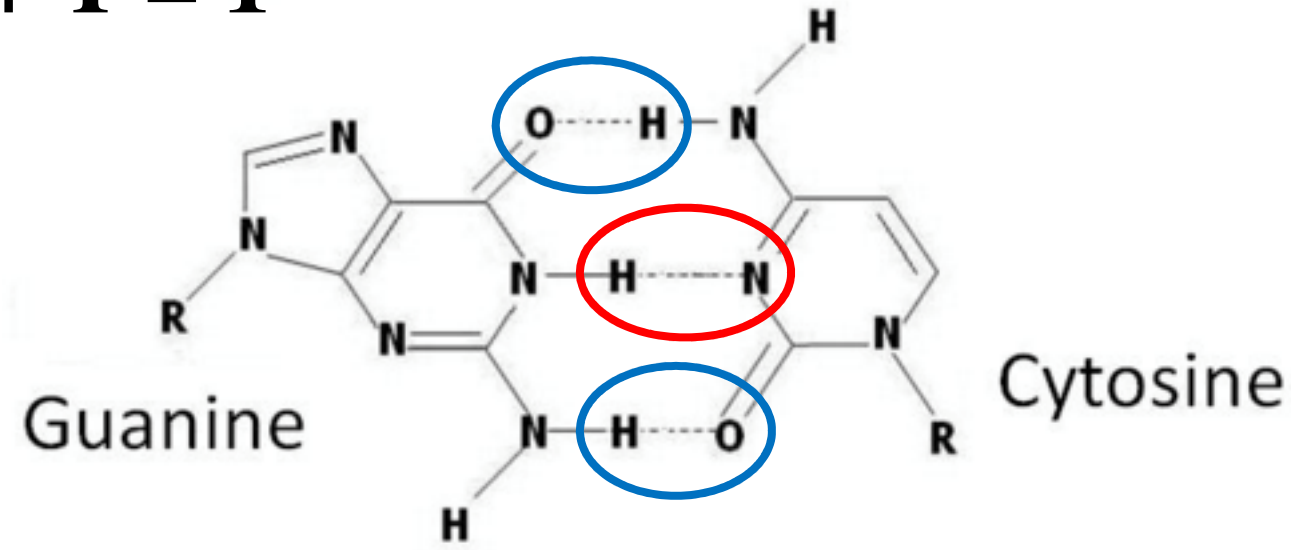

G + C = 0

Deoxyribonucleic Acid (DNA)

- Church and Kosuri (Harvard) have stored 700 TB in 1g DNA 
  - 700 TB = 233 3-TB drives weighing* 151,000 g
  - Data were stored in DNA, retrieved and reproduced on magnetic medium
- Magnetic and optical storage are 2D
- DNA storage is 3D—at the molecular level
- The cost for writing the data to DNA is still too high to be practical ($12,000/MB) 
- **Stay tuned!**

*151,000 g = 335 pounds

Holographic Storage on Quartz

- Digital data stored in three layers as nanostructured dots
- Two additional dimensions:
  - Size of data dot
  - Alignment of dot
- Capacity: 360 TB on CD-size platter (online)
- Life expectancy 13.8 billion years at 190^0 C*
- For more information see URL URL

***In ~1 billion years the sun will boil off all water on earth**

Thanks for your attention…

The End!!!