

WHY ARE WE STILL USING TAPE?

--AND--

CAN WE DO IT IN THE CLOUD?



ALAN D. HALL, OPERATIONS MANAGER/SYSTEM OWNER
NOAA'S OFFICE OF SATELLITE AND PRODUCTS OPERATIONS
COMPREHENSIVE LARGE ARRAY STORAGE SYSTEM (ARCHIVE)
ALAN.HALL@NOAA.GOV



WHAT DO I WANT TO TALK ABOUT?

Who am I?

Zetabyte Apocalypse

NOAA's Data Tsunami


Tape Migration

The Case for Tape

Archive in the Cloud

My Estimate for Cloud

Put it all together



40+ Year IT Veteran

Wrote first program in 1976 (high school)

Distributed mini-computer systems (Data Generals)

Enterprise System Admin tools

Enterprise Ingest System

Migration of optical disks and 8mm tapes to new technology

World Meteorological Organization - Ocean Observations

Operations Manager for NOAA's Archive (CLASS)

MY STORAGE EXPERIENCE




N
S
X



* By Kevin586 - en.wikipedia.org: 17:32, 22. Mai 2005 . . Kevin586 (Talk) . . 976x592 (62519 Byte) (A laserdisc (left) compared with a DVD (right).), CC BY-SA 3.0. <https://commons.wikimedia.org/w/index.php?curid=9502131>


ENGINEERING AXIOM

NOTHING EVER GETS BUILT ON SCHEDULE
OR WITHIN BUDGET.



ZETABYTE APOCALYPSE

CIRCA TODAY! ~170ZB ANNUALLY BY 2025!



- 500 million tweets are sent
- 294 billion emails are sent
- 4 petabytes of data are created on Facebook
- 4 terabytes of data are created from each connected car
- 65 billion messages are sent on WhatsApp
- 5 billion searches are made

By 2025, it's estimated that 463 exabytes of data will be created each day globally

Source: *How Much Data is Generated Each Day?*
By Jeff Desjardins, March 13, 2019
<https://www.visualcapitalist.com/how-much-data-is-generated-each-day/>

ZETABYTE APOCALYPSE

EXAMPLE - HAWK-EYE VISION PROCESSING

Hawk-Eye Electronic Line Calling - Tennis

Collects:

- Location of the ball and player
- Spin of the ball
- Speed and flight of the ball

At 300 Frames/second, Video is blended into arc-trajectories and impact ellipses

Average Rally: ~10 Seconds

Average ATP Match Duration: 2 ½ - 3 Hours

Quick Math:

4K Video
25MB/frame

300fps

7.5GB/sec

10 Cameras

750GB/sec

60 Seconds

45TB/minute

2.7PB/Hour

8.1PB/Match

NOAA's CLASS

COMPREHENSIVE LARGE ARRAY STORAGE SYSTEM

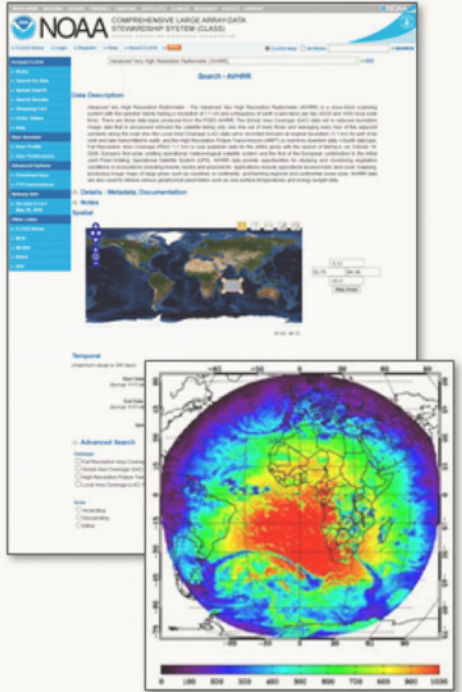
An **ARCHIVE** of environmental data and products from many sources

Users can sign-up for an account and **ORDER** by date, type, data

Orders are **DISSEMINATED** asynchronously via HTTP and FTP

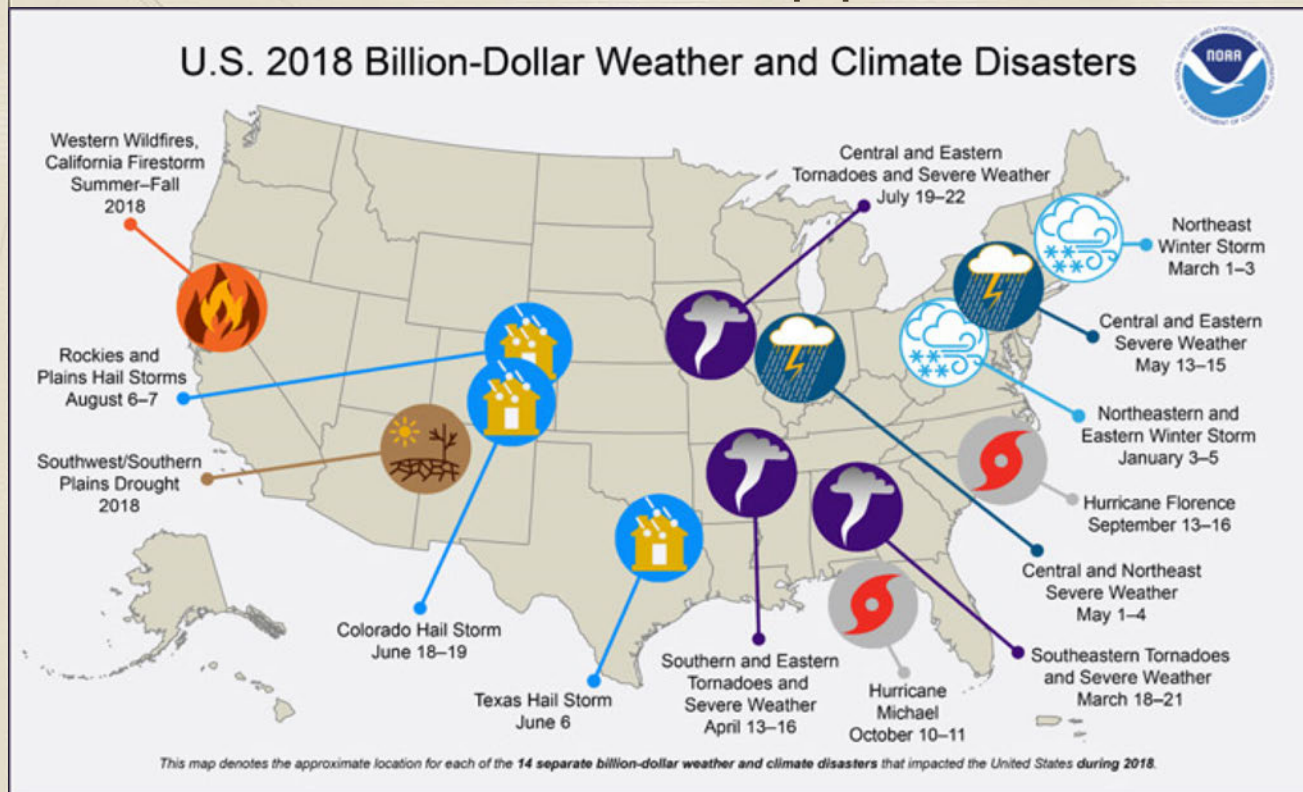
SUBSCRIPTIONS are available for automatic dissemination of new data

BULK orders are available for large packages of data



The screenshot displays the NOAA Comprehensive Large Array Storage System (CLASS) website. The page title is "NOAA COMPREHENSIVE LARGE ARRAY STORAGE SYSTEM CLASS". A search bar contains the text "APR08". Below the search bar, there is a "Data Description" section with a detailed paragraph about the system. A world map is visible, and a large circular temperature map is overlaid on the bottom right. The temperature map uses a color scale from blue (cold) to red (hot), with a legend at the bottom ranging from 100 to 1500. The URL "www.class.noaa.gov" is displayed at the bottom of the screenshot.

NOAA's BILLION \$\$ PROBLEM



NOAA Needs to:

Research-to-Applications
Faster

Delivery of Earth
Observations Faster

-TO-


Predict Faster

Alert Faster

Protect Life and Property

The United States experienced 14 billion-dollar weather disasters in 2018, which resulted in 247 deaths and ~\$91 billion in damage. (NOAA NCEI)

ENGINEERING PICK-UP LINE



WANNA SEE THE PROGRAMS ON
MY HP-48GX?
(BTW, IT NEVER EVER WORKED!)


NOAA's LEGACY DATA

Satellites!

- ~45 satellites have been launched since the 1970's
- First Polar Orbiting Satellite launched in 1970 (POES)
- First Geostationary Satellite launched in 1975 (GOES)
- Some are still operation today

Data Generated:

MBs to a few hundred of GBs per day



NOAA's NEW GENERATION SATELLITES



Joint Polar Satellite System (JPSS)

- Suomi NPP launched 2011
- JPSS-1 launched in 2017

Data Generated: 3.2TB per day

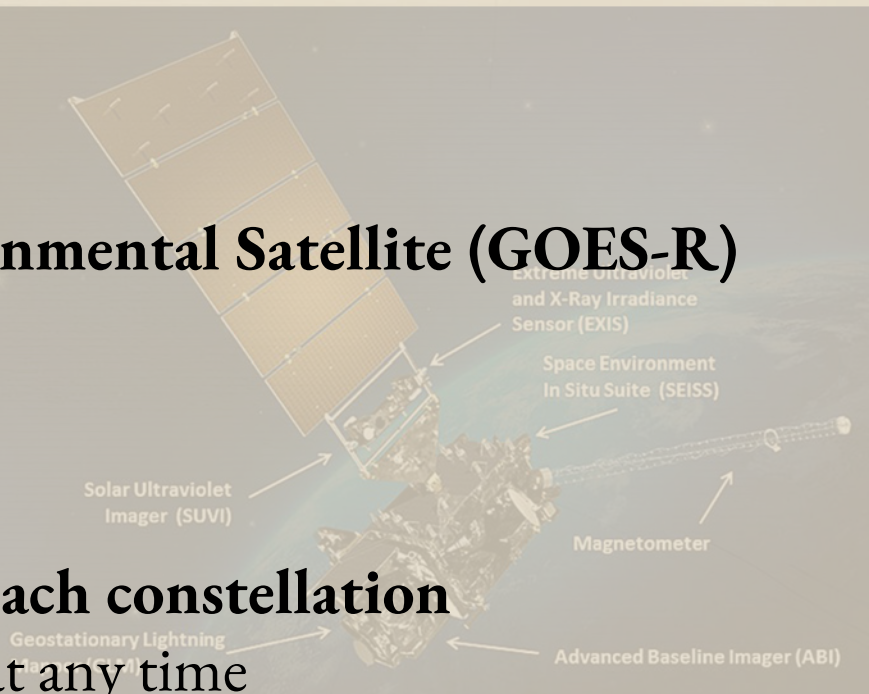
Geostationary Operational Environmental Satellite (GOES-R)

- GOES-16 Launched in 2016
- GOES-17 Launched in 2018

Data Generated: 1.4TB per day

Future Launches are planned for each constellation

Only two operational from each at any time




NOAA's DATA TSUNAMI




NOAA's DATA-JPSS JOINT POLAR SATELLITE SYSTEM

VON KARMAN
VORTICES-CANARY
ISLANDS



NOAA's DATA-JPSS JOINT POLAR SATELLITE SYSTEM

AURORA BOREALIS



NOAA's DATA- GOES

GEOSTATIONARY
OPERATIONAL
ENVIRONMENTAL
SATELLITE

HURRICANE DORIA



NOAA's DATA- GOES

GEOSTATIONARY
OPERATIONAL
ENVIRONMENTAL
SATELLITE

WILDFIRE DETECTION



TOTAL DATA INGEST

Historical Data Ingest (up to ~2009):

Legacy Satellites: ~600GB/day or less

Current Data Ingest:


Next Generation Satellites: 4.6TB/day

Legacy/Other Satellites and Non-Satellite Data: 1.7TB/day

Total: 6.3TB/day

Data Tsunami:

Increase of just over 10 fold from just a few years ago!



BUT WAIT! THAT'S NOT ALL!



New Launches:

Both Polar and Geostationary Launches

Other People's Data:

New Plan! Stop launching expensive Satellites!

Make use of Other Countries' or Commercial data

Current Possibilities:

EUMETSAT METOP-SG: 4-8TB/day

ISRO SCATSAT: 1TB/day

Total: ??TB/day

Increase of ??????!



REAL ENGINEERS

WEAR BADGES SO THEY DON'T FORGET
WHO THEY ARE.



Wile E Coyote
ACME Engineering Services

If lost, return to me!

NOAA's ARCHIVE



THE NUMBERS

~13PB of data stored on Spectra Logic T-Finity Libraries

Ingest: ~6.3 TB/day


Dissemination: ~18TB/day

Current Annual growth: ~3PB

Two identical libraries: Separate Geographic Locations
Asheville, NC & Boulder, CO

Total of ~26PB Stored

NOTE: Data is almost NEVER deleted!



NOAA's ARCHIVE

THE DEEP STORAGE



CLASS-AVL: 84PB Capacity

**10 Frame Spectra Logic T-
finity Library**

36 LTO-8 Tape Drives

**6,940 LTO-8 Tapes
(12TB Native)**

CLASS-BOU: 18PB

Capacity




**8 Frame Spectra Logic T-
finity Library**

36 LTO-6 Tape Drives

**6,940 LTO-6 Tapes
(2.5TB Native)**


NOAA's ARCHIVE

TOTAL ARCHIVE OF ~13PB: BYTES



NOAA's ARCHIVE

DATA STATS - DISSEMINATION AGE

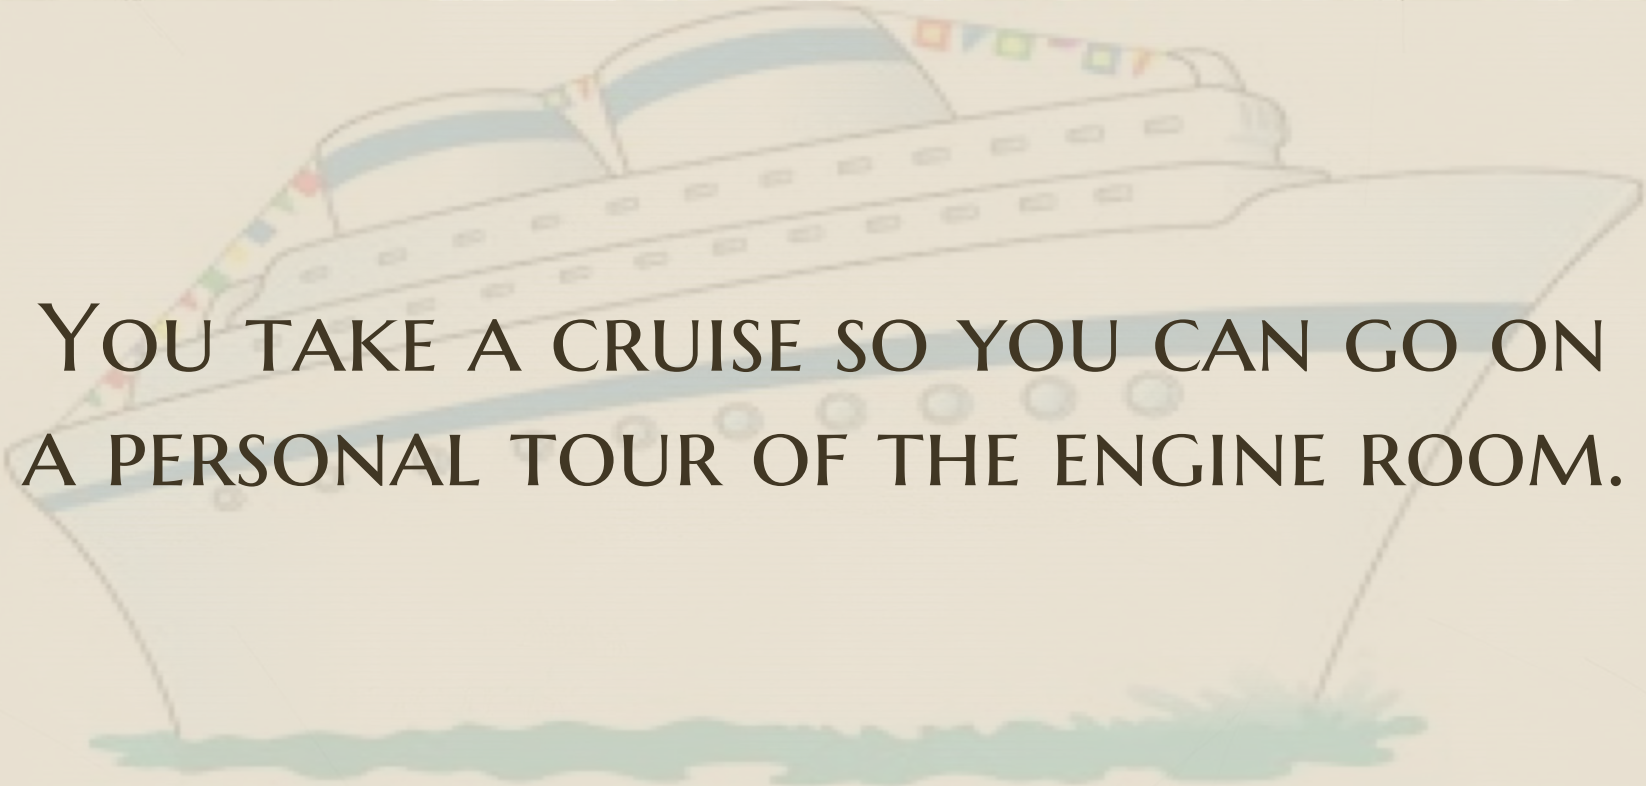


83.7% of data is delivered at Ingest via subscriptions

More than 50% of data is delivered within the first year after Ingest

Significant Wx events and Research drives data retrievals that are over 3 years old


YOU MIGHT BE AN ENGINEER IF...



YOU TAKE A CRUISE SO YOU CAN GO ON
A PERSONAL TOUR OF THE ENGINE ROOM.

NOAA's ARCHIVE

TAPE MIGRATION



CLASS-AVL: From 18PB (LTO-6) to 84PB (LTO-8)

Completed LTO-6 to LTO-8 Migration

Cost: ~\$900K

Operations were shifted to CLASS-BOU

CLASS-AVL dedicated to Migration

Only new data was written to tape in CLASS-AVL

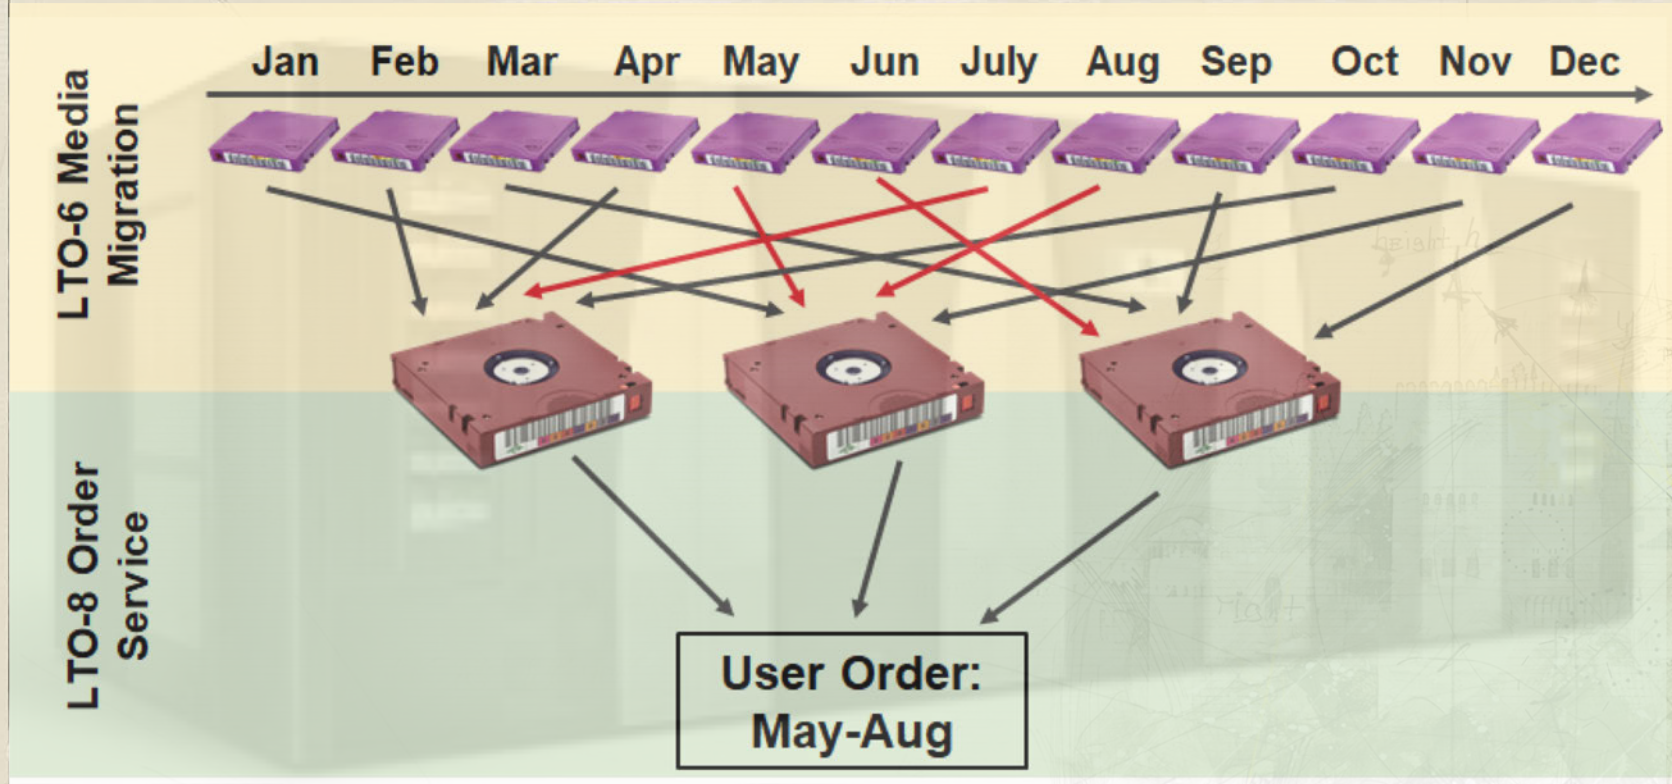
Challenge: Optimized Retrieval

Retrieval from archive is by date: Example order May-Aug

Migration planned for optimal retrieval by date


NOAA's ARCHIVE

TAPE MIGRATION: AUTOMATED REPACK - BRUTE FORCE




NOAA's ARCHIVE

TAPE MIGRATION: REPACK OPTIMIZED FOR DATE RETRIEVAL



LTO-6 Media
Migration

Jan Feb Mar Apr May Jun July Aug Sep Oct Nov Dec



LTO-8 Order
Service

User Order:
May-Aug


NOAA's ARCHIVE



TAPE MIGRATION CHALLENGE: MAXIMIZE SPACE

LTO-6: 2.5TB

LTO-8: 12TB




LTO-6 (2.5 TBs)	LTO-6 (2.5 TBs)	LTO-6 (2.5 TBs)	LTO-6 (2.5 TBs)	Unused (2TB)
-----------------	-----------------	-----------------	-----------------	--------------

The Goal:

Minimizing unused space on the destination tapes

NOAA's ARCHIVE

TAPE MIGRATION: MAXIMIZE SPACE: GROUPING




Grouping both **Source** (LTO-6) and **Destination** (LTO-8) tapes maximizes tape utilization

Unused Portion of Tape: **1.4% Avg**



NOAA's ARCHIVE



TAPE MIGRATION: DETAILS

Operations Shifted to Writing LTO-8 **Dec 12, 2018**

Migration Started: **Jan 22, 2019** Completed: **July 5, 2019**

165 Days; Avg: **67TBs/day**; Total: **11.1PB**

384.3M Files; ~**2.3M** files/day (Lots of small files)

19.7 LTO-8/day; **3,230** LTO-6 Migrated to **782** LTO-8 (4:1 Ratio)

Verification (checksums) Completed **Aug 20, 2019**

NO TAPE ERRORS!

Earliest manufacture date of the LTO-6 media: May 2013


YOU MIGHT BE AN ENGINEER IF...

YOU SEE A GOOD DESIGN AND STILL HAVE
TO CHANGE IT.

Before
Engineering
Prototype




After
Engineering




THE CASE FOR TAPE

STORAGE PYRAMID - ORIGINAL BEFORE GOOGLE

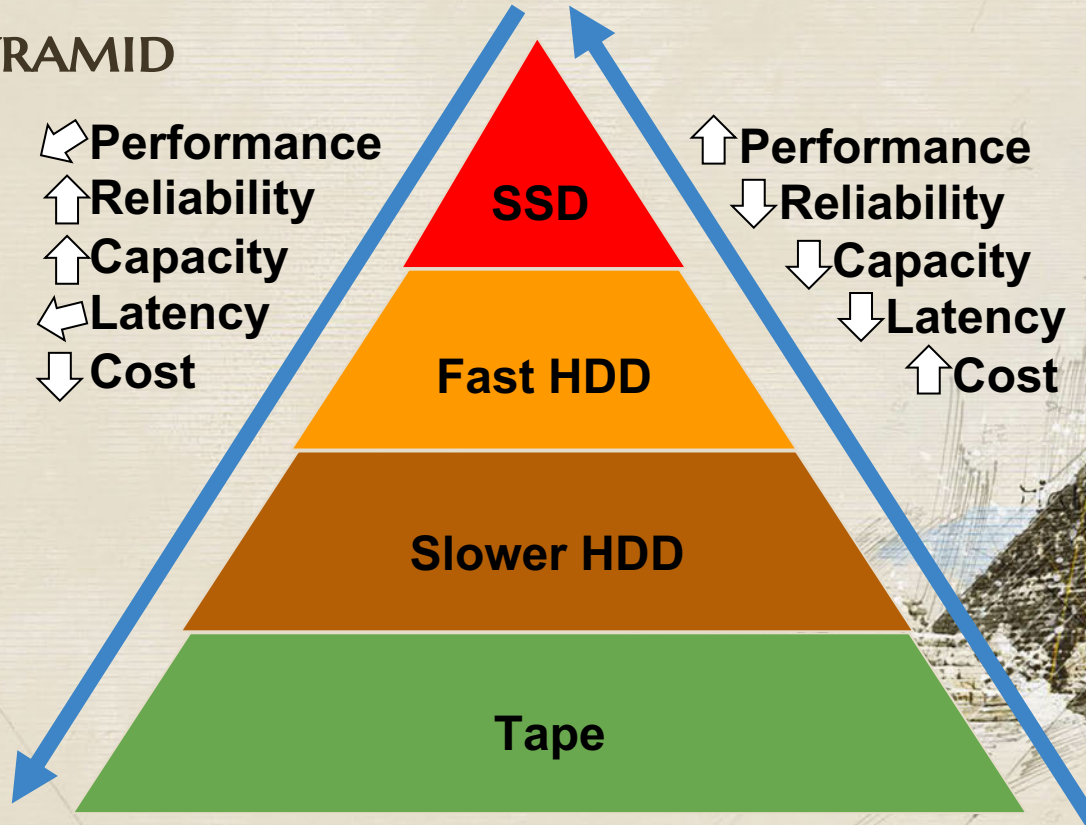


Who knew this was
tattooed on Fred Moore's
back?



THE CASE FOR TAPE

STORAGE PYRAMID



THE CASE FOR TAPE

WHAT IF IT'S ALL ON SPINNING DISK?

How much Disk do we need?

USING TAPE?"

13PB + (3PB/yr*3Yrs)=22PB for Archive

Ingest/Dissemination/Overhead: 2PB

Minimum of 24PB * 2copies = **48PB Usable**

Two Geographic Locations

Disk Cost is 155/TB/yr*48,000=**\$7.4M/yr**

5yr Lifecycle Refresh

Plus Maintenance/Facilities/Communications Costs: **~\$26M/3yrs**

Over 3 yrs

Interpolated Estimate based on actual annualized purchases

33 implementation



“WHY ARE WE STILL

TOTAL

Current


THE CASE FOR TAPE

THE FACTS ABOUT TAPE

The Way It Is (Modern Tape Era)

- Surpassed HDDs Reliability (BER)
- Data Rates >2x Faster Than HDDs
- Media Life 30+ Years
- Exabyte+ (1×10^{18}) Libraries
- RAIT, RAO and TAOS Arrive (Perf.)
- Lowest Energy Consumption and TCO
- Intelligent Robotics, Smart cartridges
- No Known Recording Limits

Fred Moore, STS2019 “Moving from Storage Hype to Storage Reality”



THE CASE FOR TAPE

CASE CLOSED!



Given that:

- ★ Most of our data is accessed infrequently after the first year
- ★ Spinning disk is cost prohibitive and not as reliable
- ★ Tape is the most reliable; lowest TCO; and fast retrievals

TAPE WINS!

YOU MIGHT BE AN ENGINEER IF...

YOUR IDEAL EVENING CONSISTS OF FAST-FORWARDING THROUGH THE LATEST SCI-FI MOVIE LOOKING FOR TECHNICAL INACCURACIES.

THE CLOUD

AN INDUSTRY ESTIMATE

Matt Star

Spectra Logic

STS2019: “Rub Some Math On It”



20 PB local copy and backup copy in the cloud

- Ten year running total, \$11,313,320
- Cloud Storage cost for 10 years \$6.3 million
- No Retrieval Pipe cost (\$2.86 million)
- No Cloud Out charges (\$2.33 million)

- On prem copy for 10 years \$3.94 million
- Cloud, Second copy \$6.292 million
- Total for two copy 10.2 million

THE CLOUD

MY EXAMPLE

13PB to upload over a year

Growth at 3PB annually

Dark copy almost **NEVER** accessed

- Bad Tape (damaged by over zealous robot)
- Catastrophic Recall (Asheville is a smoking hole)

On-Prem still primary for servicing



THE CLOUD

RUB SOME MATH ON IT (YMMV)




Budget Summary	Budget Summary	Cloud	On-Prem	3Yr Savings
Direct Costs*	Direct Costs*	\$7,120,011	\$9,405,609	\$2,285,597
Cloud Storage (AWS)	Cloud Storage (AWS)	\$3,147,200	\$0	(\$3,147,200)
Tech Refresh	Tech Refresh	\$4,616,355	\$7,122,000	\$2,505,645
Totals	Totals	\$14,883,566	\$16,527,609	\$1,644,042

Direct Costs	\$5,645,666	\$2,666,355	\$2,666,355	\$7,120,011
Cloud Storage (AWS)	\$575,754	\$1,187,640	\$1,383,806	\$3,147,200
Tech Refresh	\$1,171,020	\$2,164,200	\$1,281,135	\$4,616,355
Total	\$4,789,774	\$5,360,220	\$4,733,572	\$14,883,566


*Maintenance Contracts; Facilities; Communications

PUT IT ALL TOGETHER

MY ESTIMATES!!



WHAT IS THE DIFFERENCE BETWEEN ENGINEER BOOTS AND COWBOY BOOTS?



COWBOY BOOTS HAVE THE BULL-CRAP
ON THE OUTSIDE!


-PETE DONEIS



THANK YOU!




I deny everything in this presentation!





BACKUP SLIDES



ZETABYTE APOCALYPSE CIRCA TODAY!



<https://wearesocial.com/blog/2019/01/digital-2019-global-internet-use-accelerates>

<https://www.visualcapitalist.com/what-happens-in-an-internet-minute-in-2019/>


2019 This Is What Happens In An Internet Minute



Created By:
@LoriLewis
@OfficiallyChadd

ZETABYTE APOCALYPSE

EXAMPLE - AUTONOMOUS CARS



Quick Math: By 2020

10 Million Autonomous Cars
4TB/Day * 10M

Level 2 Autonomous Driving

40PB/Day

Level 3 and above would
generate much more data!

<https://www.unicomengineering.com/blog/the-key-to-autonomous-driving-the-5g-network/>

ZETABYTE APOCALYPSE

EXAMPLE - AVIATION



By 2026:

→ 98EB Generated from global aviation fleet annually

→ 5-8TB per flight

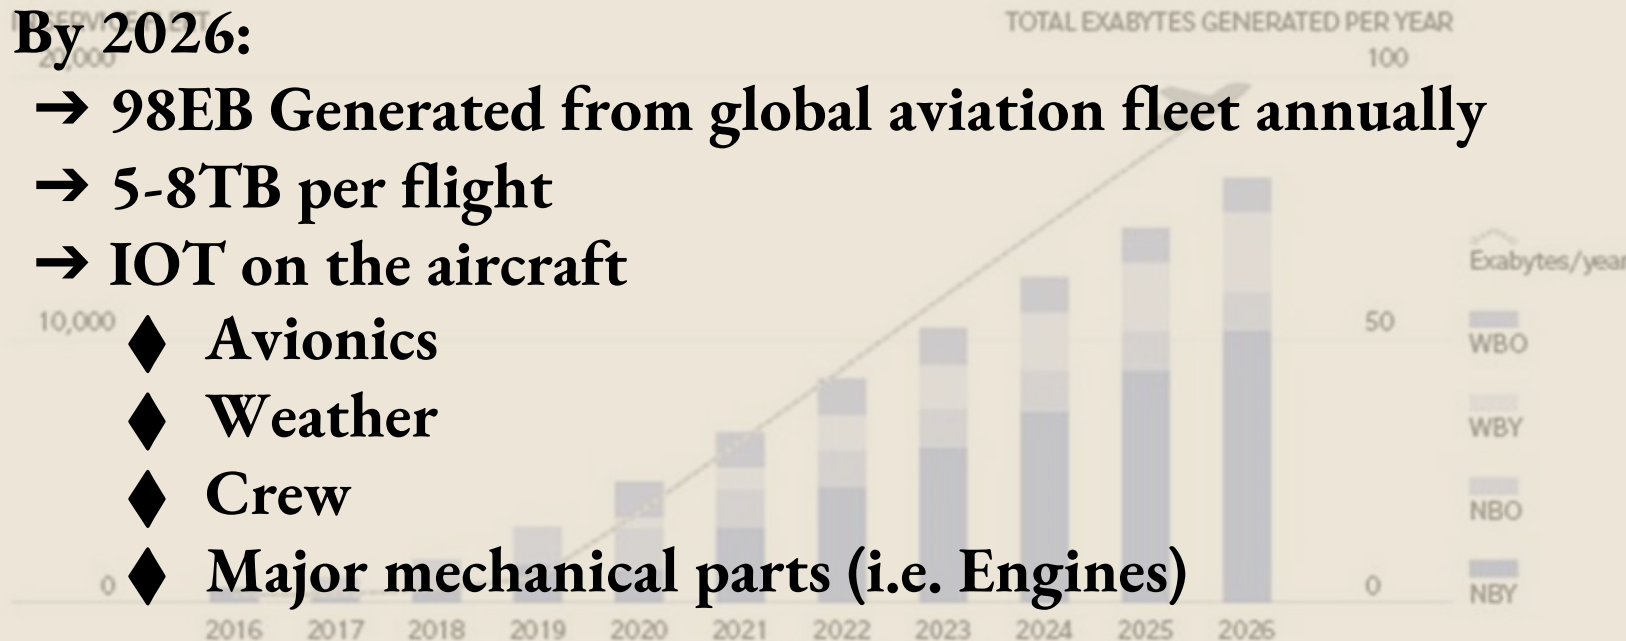
→ IOT on the aircraft

◆ Avionics

◆ Weather

◆ Crew


◆ Major mechanical parts (i.e. Engines)




Source: Oliver Wyman Fleet & MRO Forecast, www.planestats.com/betterinsight

ZETABYTE APOCALYPSE

CIRCA TODAY!



Annual Size of the Global Datasphere

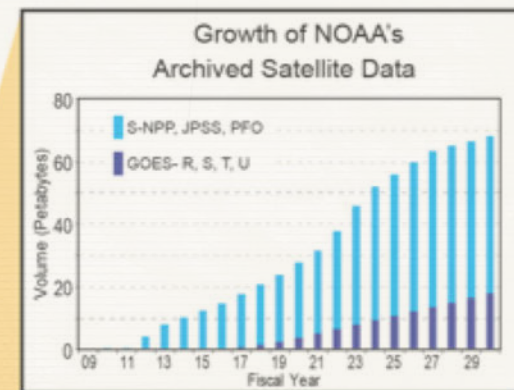
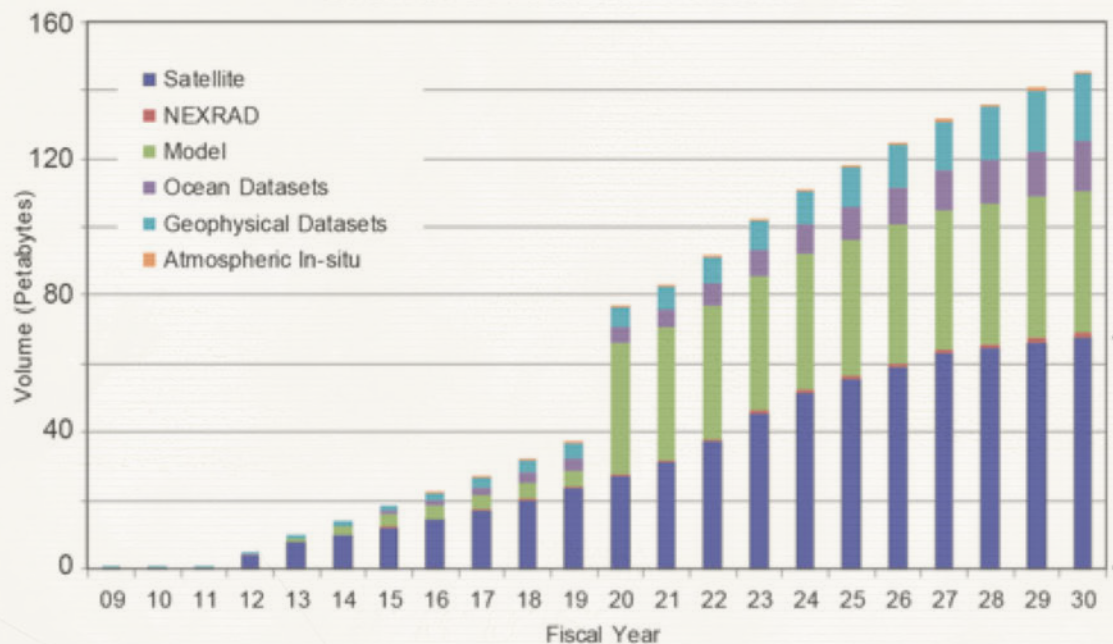


<https://www.storagenewsletter.com/2018/11/28/global-datasphere-from-33zb-in-2018-to-175zb-by-2025/>

DATA GROWTH (CIRCA 2016)






Growth of NOAA's Archive



NOAA's ARCHIVE

DISSEMINATION OVER LAST 12 MONTHS - BYTES





SlidesCarnival icons are editable shapes.

This means that you can:

- ◆ Resize them without losing quality.
- ◆ Change fill color and opacity.
- ◆ Change line color, width and style.

Isn't that nice? :)

Examples:





Now you can use any emoji as an icon!
And of course it resizes without losing quality.

How? Follow Google instructions
<https://twitter.com/googledocs/status/730087240156643328>

