

**OCIO**

Office of the  
Chief Information Officer  
LIBRARY OF CONGRESS

# Library of Congress Storage Architecture

OCIO/ITSO  
Production Engineering  
and  
Operations & Maintenance





# Overview

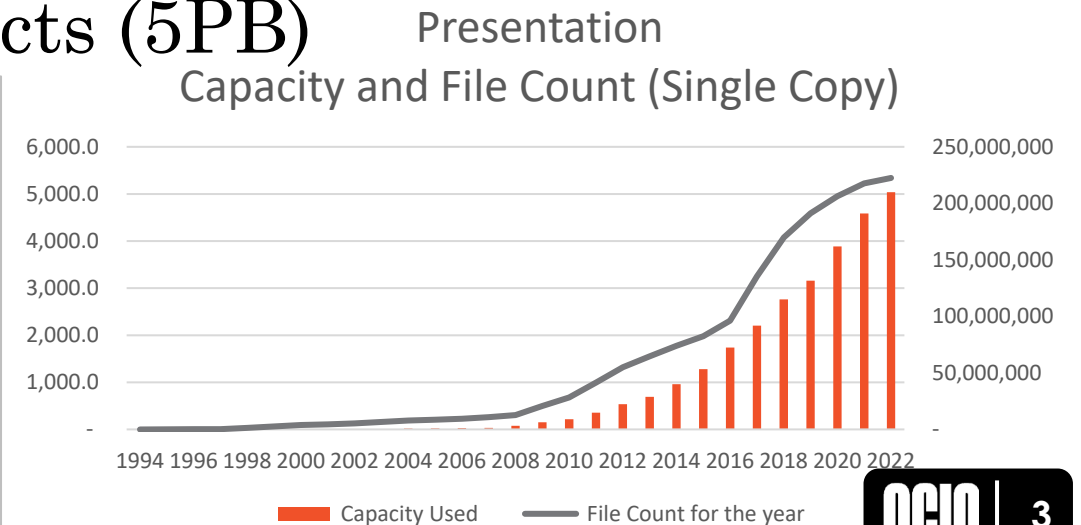

- Environment Stats
- Current Content Storage Architecture
- Upcoming Changes to the Architecture
- Perpetual Migration is Here
- Moving to the Cloud
- Large Scale Propagation – Lesson Learned and Pitfalls
- Changing to Cloud Based and Pull to On-site

# Content Storage Environment Statistics

- Preservation Single Copy 932M Objects (25.7PB)



- Presentation Single Copy 222M Objects (5PB)

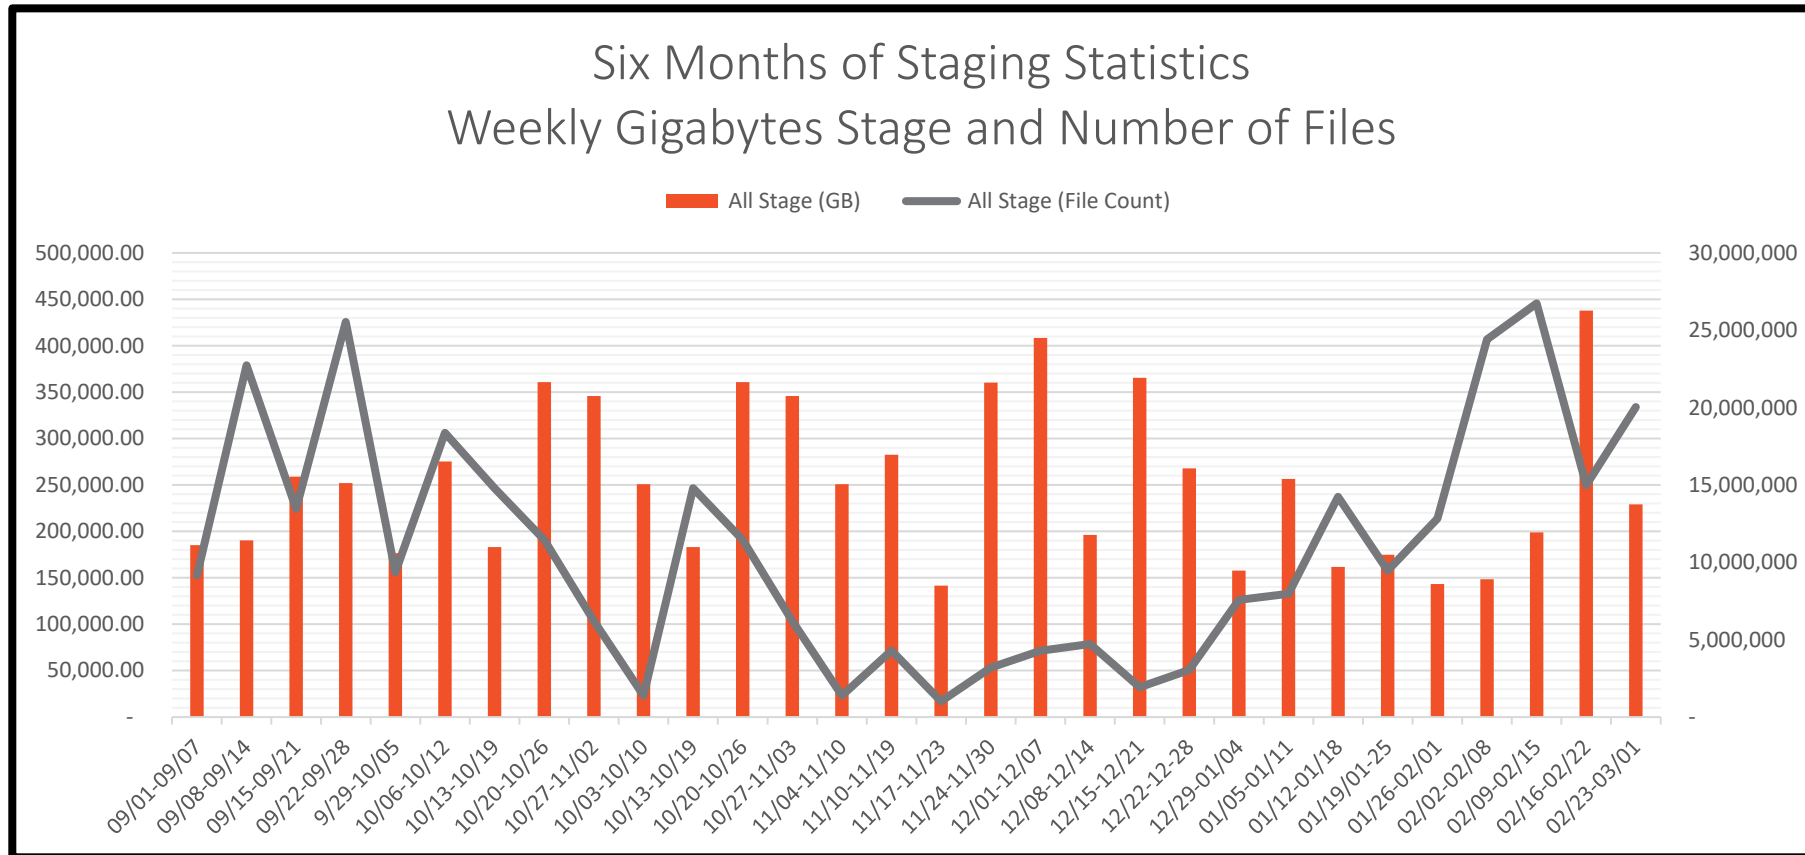


# Content Storage Environment Strategies

- Preservation Deploying a 3:2:1 strategy
  - Three (3) copies of the data (2 tape copies and one cloud copy)
  - Two (2) technologies (Versity and AWS)
  - One (1) off premises copy (AWS Glacier)
- Presentation Deploying a 2:2:1 strategy
  - Two (2) copies of the data
  - Two (2) technologies (Spectrum Scale [gpfs] and AWS S3)
  - One (1) off premises copy (AWS S3 Intelligent Tiering)

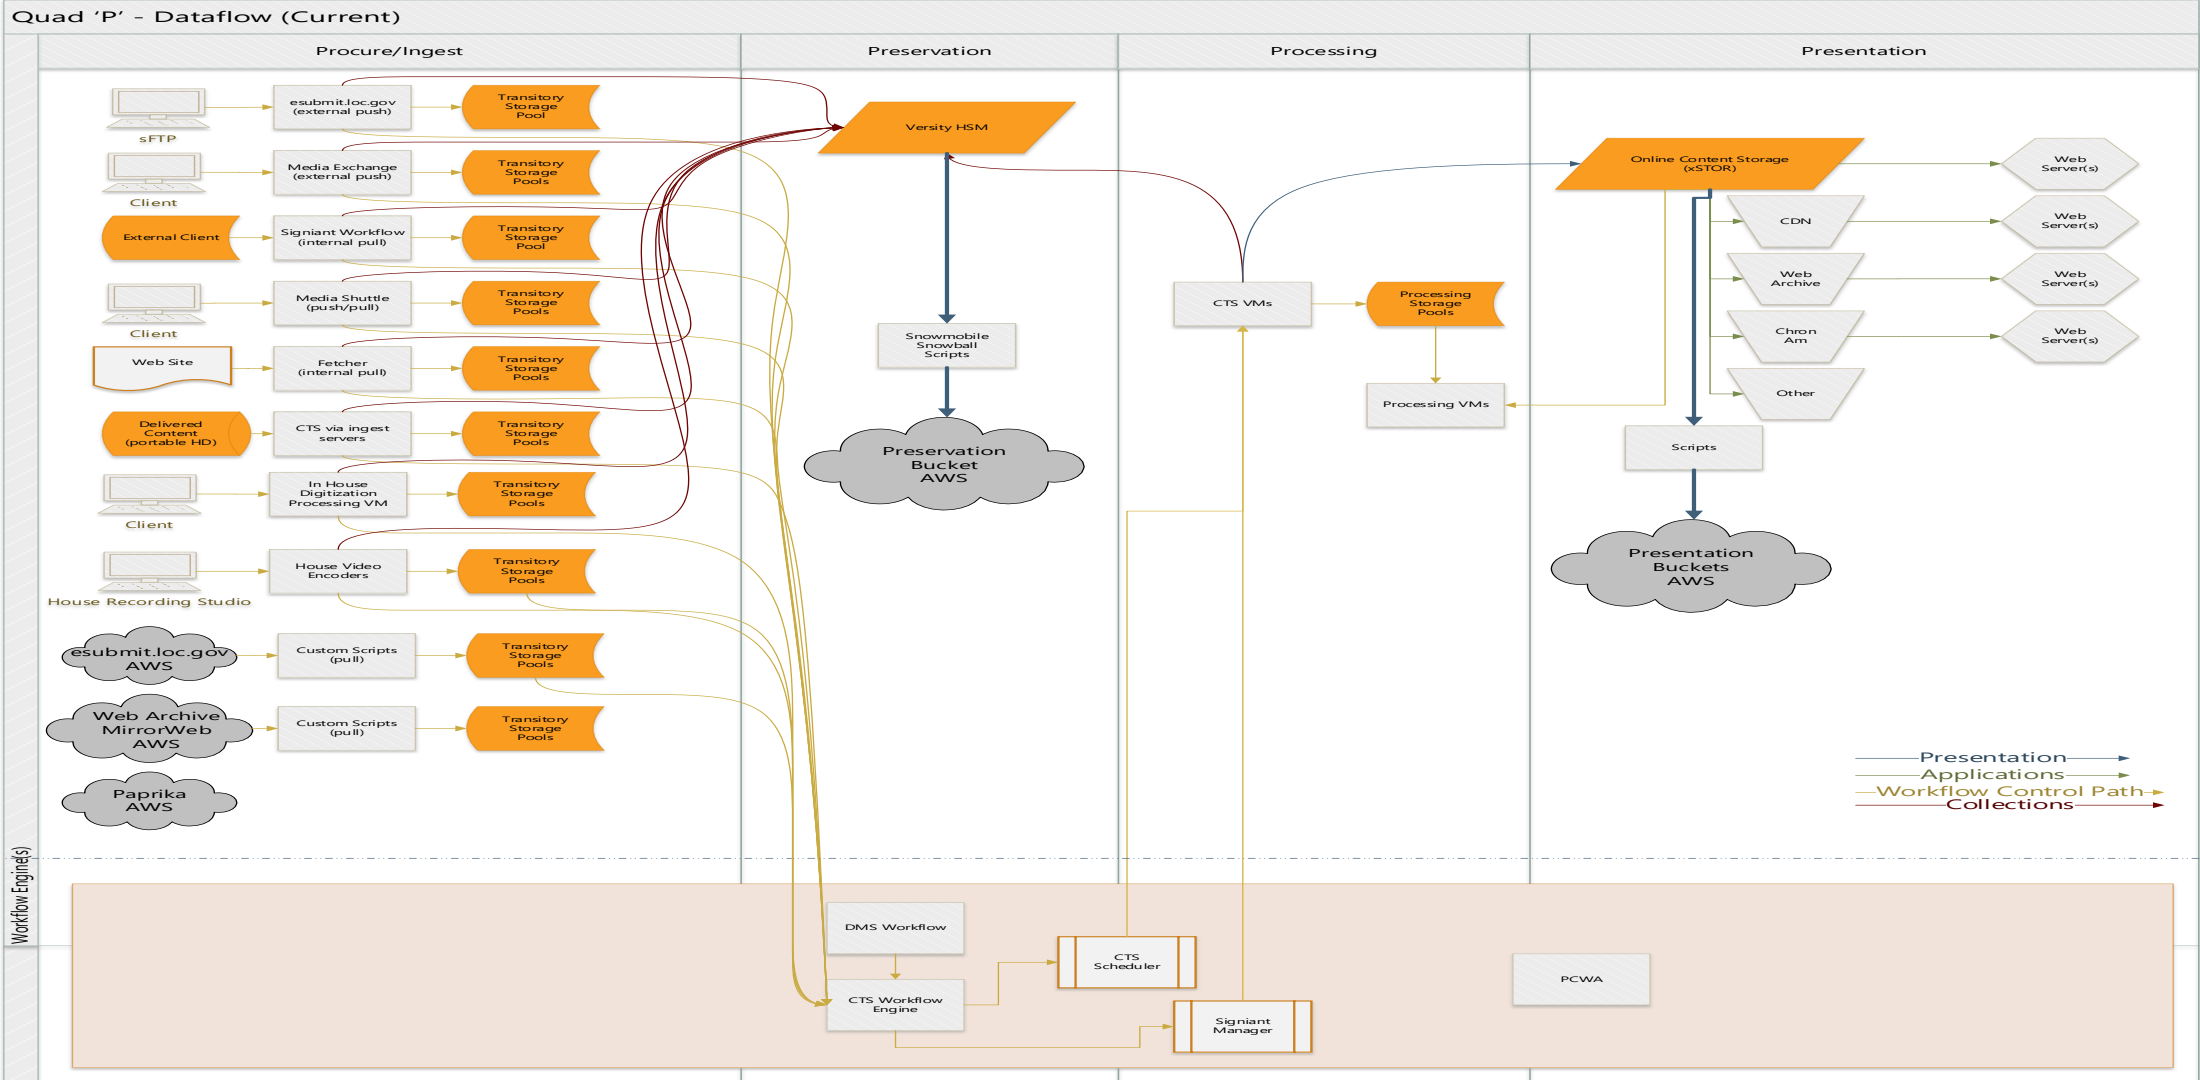
# How Active is Preservation?

- Over a six (6) months period we have staged 7.5PB and over 326M files.
- That is roughly 29% of 25.7PB (or 35% of the 932M files) for preservation.
- 78% of this was in data migrations, 4% in tape rebuilds and 18% in normal operations



# Current Environment

## Data Flow



# Still Looking to add Content Abstraction Layer

- **Content Abstraction Layer (CAL) would provide:**
  - Manage the movement of data to and from multiple sources
  - Manage the preservation of content:
    - File fixity checking
    - File validation / usability testing
  - Manage the movement / orchestration of data across multiple
    - Systems
    - Data centers
    - Cloud providers
    - External entities
  - Provide a persistent namespace and access method to data

# Content Abstraction Layer

## Challenges and Lessons Learned thus far

### Started Parallel Builds @ Summer 2019

#### 1. Loss of Storage Cohesion

- Storage backend became altered
- Reference to previous data layout lost when moving to inline/in-band data accesses

#### 2. Could Not Index Files on AWS S3

- currently 3+ years behind schedule

#### 3. Upgrading Versions Caused Unexpected Object/file Versioning

- Vendor made a choice to change how data was ingested causing source and destination inconsistency

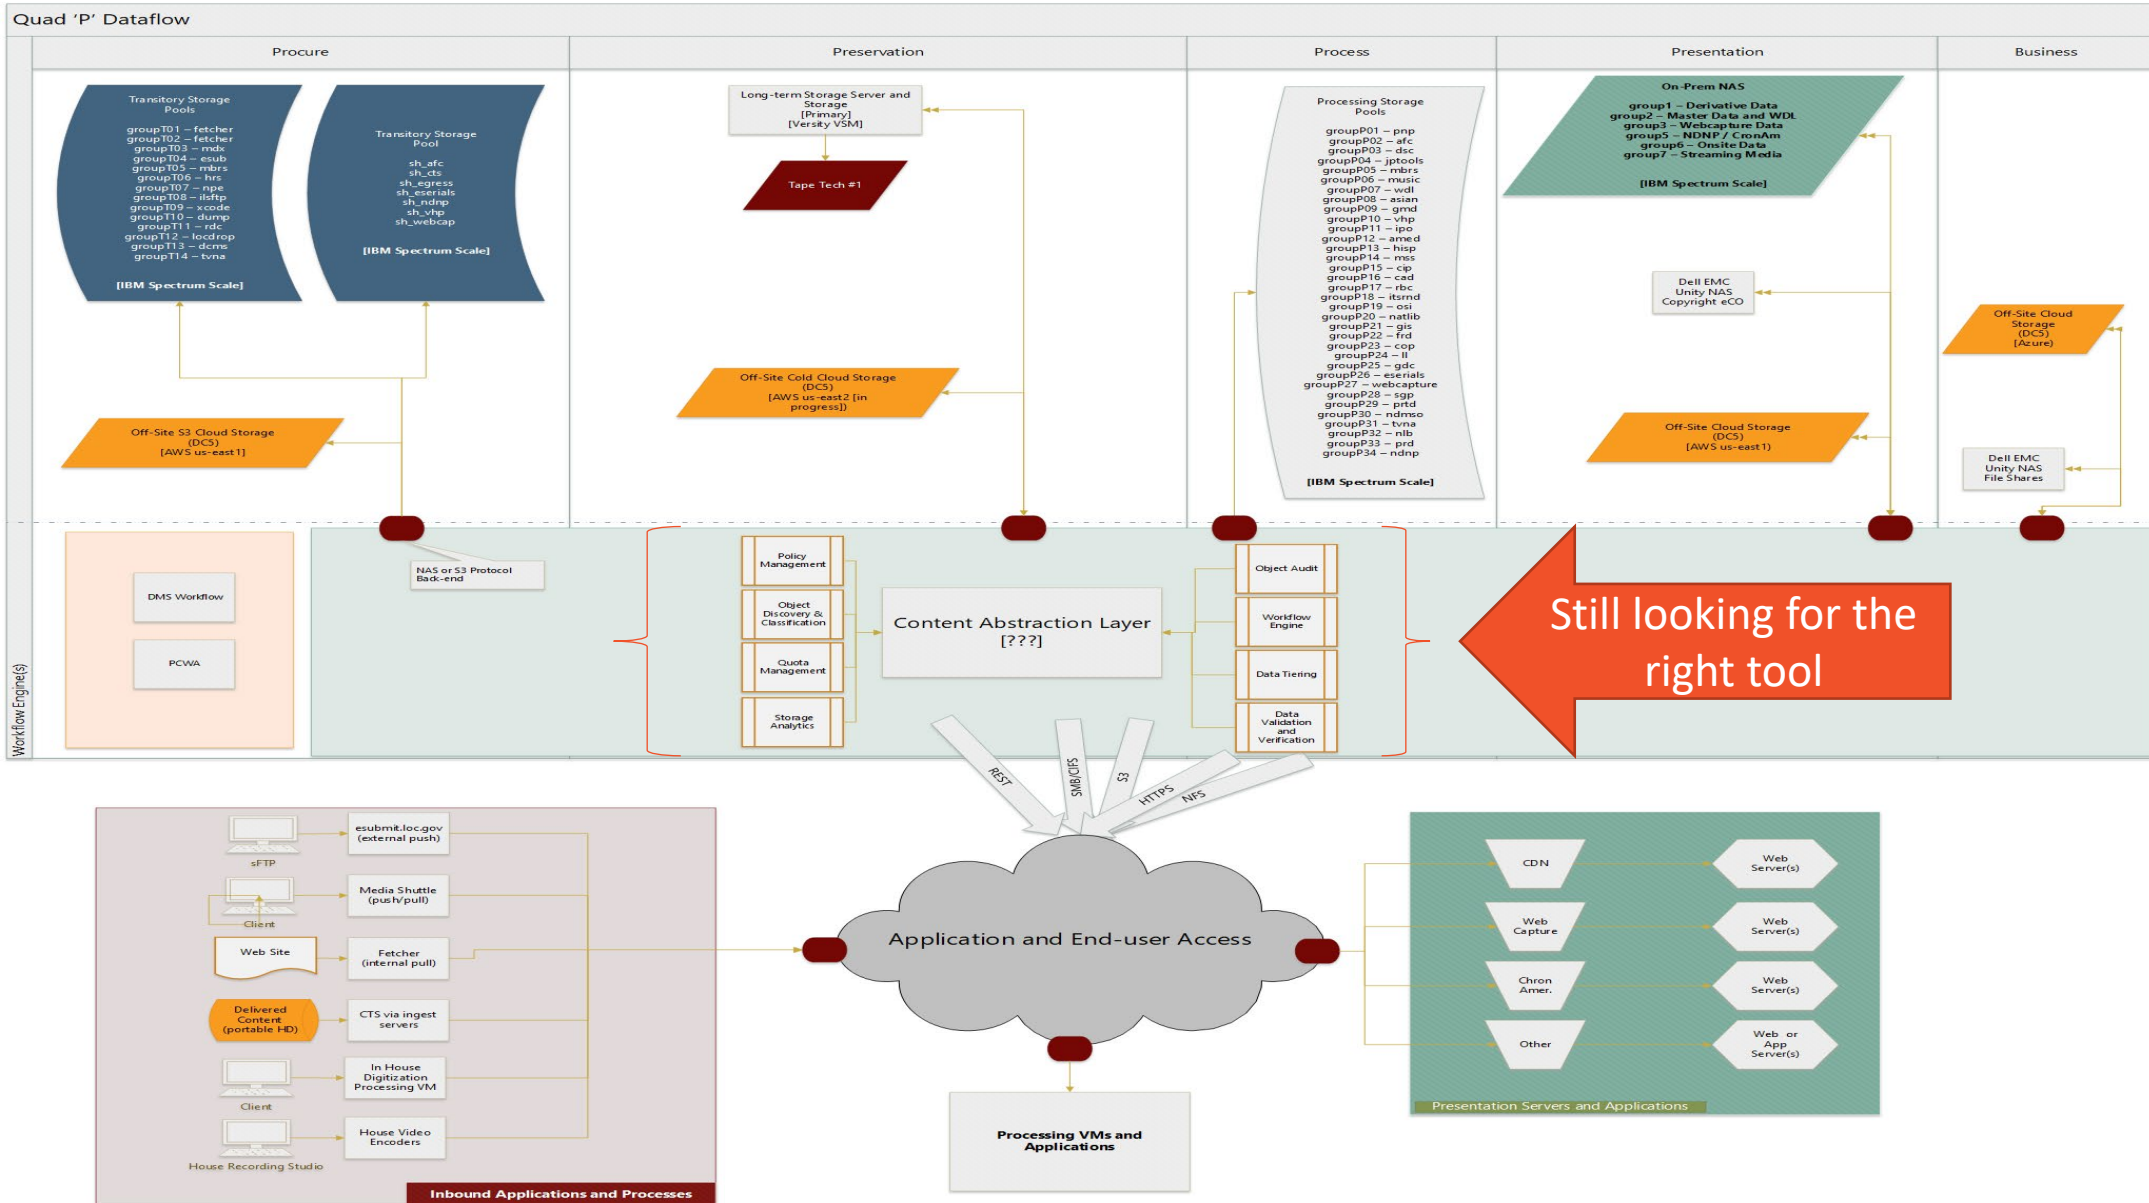
#### 4. Indexing Workflow Unforgiving

- Failure to follow strict processing caused 42% of the data not to be hashed properly

#### 5. Multiple Operational Issues Occurred After Latest Version Upgrade



# Proposed Storage Architecture (adding the Content Abstraction Layer)



# Perpetual Migration is Here

- **Software Obsolescence**
  - Oracle ended Oracle Hierarchical Storage Manager
- **Technical Refresh**
  - Old server
- **End-of-Service Life**
  - Storage hardware
- **Changing Resources**
  - Move to the cloud



Continuous Migration From the 2017 Conference  
From the 2012 Conference



"Working on the cure for those suffering with digital Disposophobia"

6

# Moving to the Cloud

Since 2019, we have propagated/replicated the following:

- **Presentation Storage**
  - Effort established July 2019
  - 55 AWS Snowballs utilized over 18 months
  - Currently in maintenance via daily and/or weekly sync of differentials
- **Preservation Storage**
  - Currently 4 years into a 2 year project to sync to AWS us-east2
  - Effort started July 2023
  - AWS Snowmobile (Received 10PB minus 3PB of overhead)
    - 7.2PB sent over a 12 month period
  - AWS Snowballs (100TB unit is actual 80TB usable per unit)
    - 96 filled over the past 24 months
    - Planning for 15 more over the next 9 to 12 months
    - Small files and Snowballs don't work well
  - Over the wire (Internet2 connection)
    - 1Gbps (981Mbps after overhead)
  - Estimated completion data at current rate early 2027



# Large Scale Propagation – Lessons Learned (thus far)

- **AWS Snowmobile:**
  - Vendor could not deliver a unit that would hold the LC current capacity
  - In 2019, the unit would not work with standard AWS 'cli'
  - Required special programming to put data to trailer
- **AWS Snowball:**
  - Small objects are not Snowball friendly
  - 1500+ operations per sec can cause system errors
- **AWS Elastic File Service (EFS)**
  - In 2019, 256 concurrent connections per instance
  - This has been to changed to 32K
- **Loss of file provenance data when moved to object storage**
  - File stat data is lost during the upload process
  - AWS throttles metadata change operations
  - 12 sec per object metadata change
  - 15 years would be needed to add metadata tags for every file
- **AWS ETAG – two (2) types**
  - An Object – MD5 Hash (example: 4ed3a1d36e630ccf1ea755778aa7ce5d)
  - Additive by object blocks – Cumulative MD5 (example: 0d9b509efd05aaae6632e9bc2d85e34f-5369)
    - Object size 41.9GB in size \* split in 5369 chunks @ a chunk size of 8MB



# Changing to Cloud based Ingest and Pull to Onsite

- **External Vendors now have the ability to send content to the Library via AWS**
  - Web Archiving via bucket replication
  - Copyright Mandatory Submissions via AWS sFTP Transfer Service
  - Cataloging in Production via AWS sFTP Transfer Service
  - External content scanning efforts submitted via AWS sFTP Transfer Service
- **Library of Congress moving to cloud centric workflows**
  - Several efforts being worked on that acquire, index and present content within AWS

**OCIO**

Office of the  
Chief Information Officer  
LIBRARY OF CONGRESS

Questions?

**OCIO**

Office of the  
Chief Information Officer  
LIBRARY OF CONGRESS

**Thank you!**

## **For More Information**

Web: <https://digitalpreservation.gov/meetings/>

## **Points of Contact**

Carl Watts, [cwat@loc.gov](mailto:cwat@loc.gov)