



Archive Consistency US ATLAS Tier 1

Scientific Data and Computing Center (SDCC)
Brookhaven National Laboratory

April 27, 2022



US ATLAS Tier 1 Archival Storage

- Comprised of three loosely coupled systems
 - ATLAS Rucio based at CERN
 - US ATLAS Tier 1 dCache disk cache at the SDCC at Brookhaven (BNL)
 - SDCC HPSS mass storage system [1]
- Each system maintains its own database to track stored files
- Database equivalence needed to ensure the following
 - No ATLAS files are lost
 - Orphaned files are not create

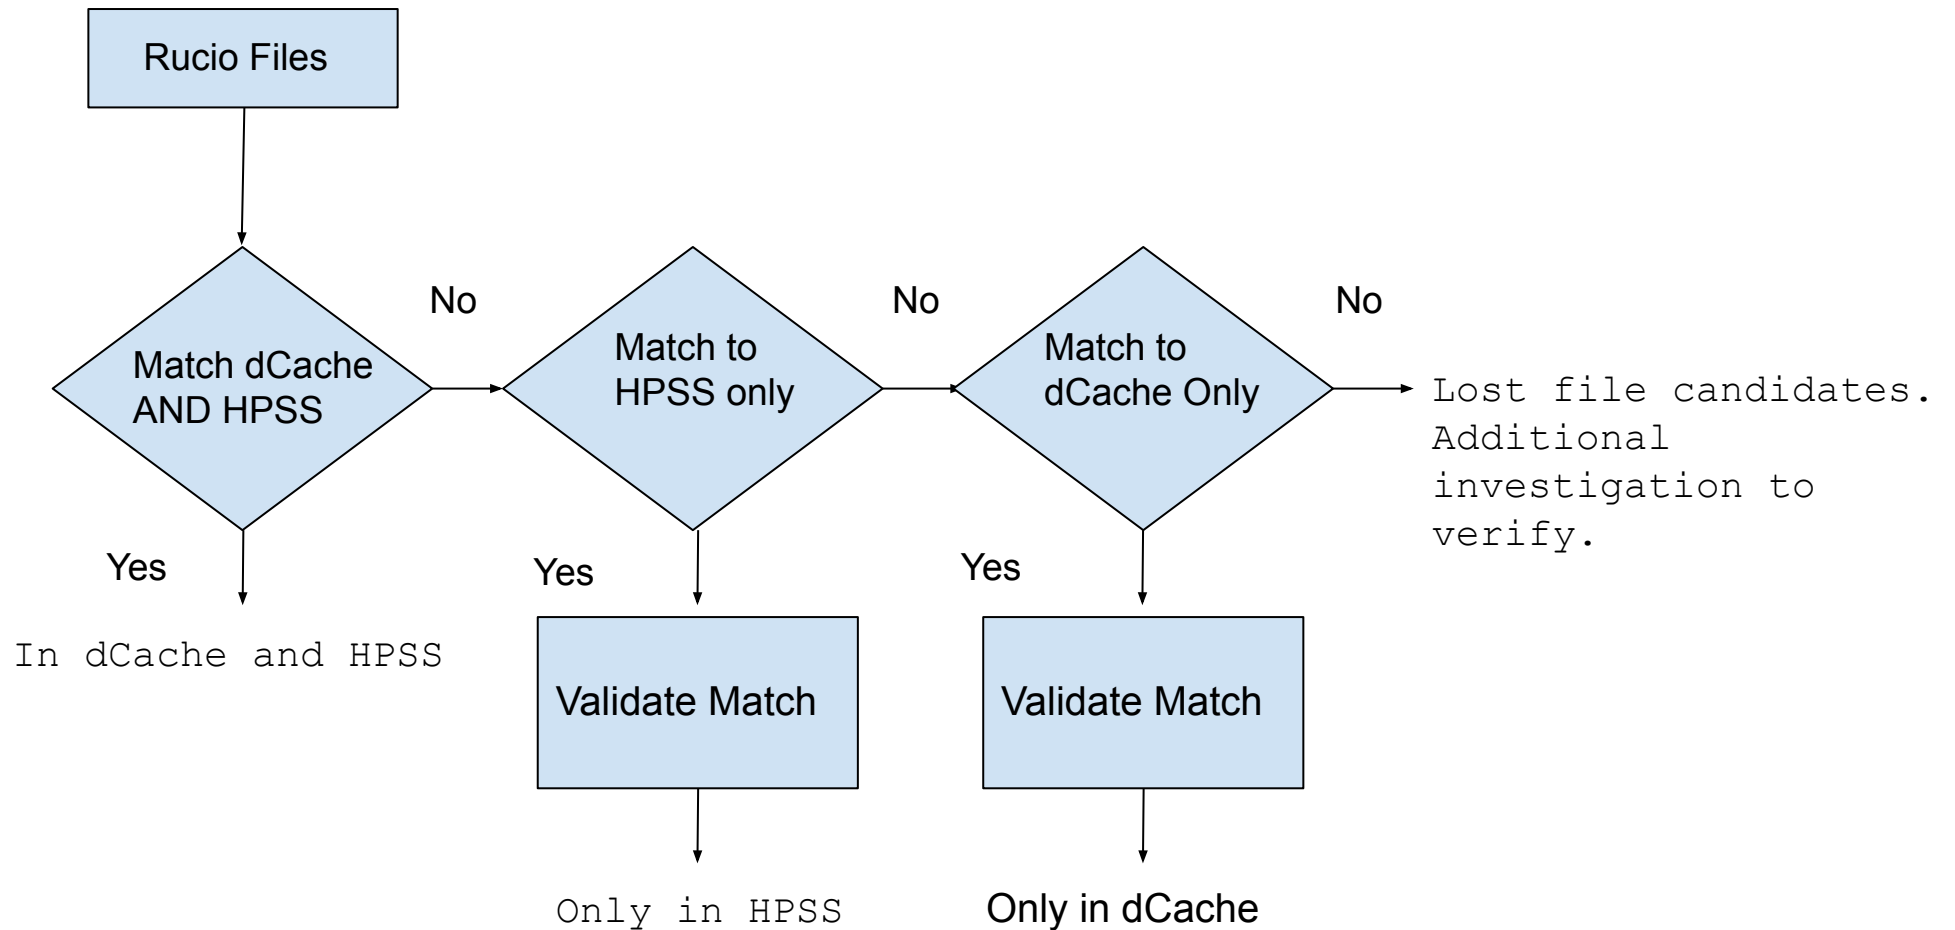
[1] HPSS - High Performance Storage System (<https://www.hpss-collaboration.org>)

US ATLAS Tier 1 Archive Check

- CERN Rucio is the reference database
 - dCache and HPSS namespaces were compared to Rucio database of files stored at BNL
 - Relevant Rucio RSE were DATATAPE, MCTAPE, and GROUPTAPE
 - These are TAPE token storage areas
 - Criteria for a match between Rucio and dCache
 - Matching file location, name, size, and Adler32 checksum
 - Criteria for a match between dCache (or Rucio) and HPSS
 - Matching file location, name and size.[1]

[1] File checksums currently not recorded in HPSS. Checksums in HPSS are being investigated

File Matching Flowchart



Namespace Analysis Results

| | Comments | DATATAPE | MCTAPE | GROUPTAPE |
|--|----------|------------|------------|-----------|
| Files in Rucio | | 16,529,567 | 22,616,519 | 3,095,620 |
| Rucio Files in dCache and HPSS | [A] | 16,518,808 | 22,606,229 | 3,086,960 |
| Rucio Files not in dCache and HPSS | | 10,759 | 10,290 | 8,660 |
| Rucio Files Found Only in HPSS | [B] | 5,923 | 0 | 0 |
| Rucio Files Not in HPSS | | 4,836 | 10,290 | 8,660 |
| Rucio Files Found Only in dCache | [A] | 3 | 2 | 0 |
| Rucio Files Not in HPSS or dCache | | 4,833 | 10,288 | 8,660 |
| Rucio Files Not Matching dCache CkSum | [B][1] | 0 | 0 | 0 |
| Rucio Files Not in HPSS or dCache | [2] | 4,833 | 10,288 | 8,660 |
| Rucio Files "In Flight" | [3] | 4,226 | 7,286 | 0 |
| "Old" Unmatched Rucio Files | | 607 | 3,002 | 8,660 |

- A. Directory, file name, size and checksum match
- B. Directory, file name and size match
 1. Files in Rucio not in 1, 2, or 3, but found in dCache with non-matching checksum
 2. Files in Rucio not found in HPSS or dCache, even when ignoring checksum differences
 3. Rucio files that have yet to be received at the SDCC

Requires further investigation. See next two slides.

Files in Rucio missing at the US ATLAS Tier 1. Count is altered by results of investigation into Rucio files found only in dCache or HPSS

Checksum/File Size Discrepancies

| | DATATAPE | MCTAPE | GROUPTAPE |
|---|----------|--------|-----------|
| Rucio Files Only in HPSS [1] | 5,923 | 0 | 0 |
| Bad Rucio Checksum (Checksum="") [2] | 1,244 | 0 | 0 |
| Exist in dCache but mismatched Checksum [3] | 18 | 0 | 0 |
| Rucio Files definitively only in HPSS (derived from above) [4] | 4,661 | 0 | 0 |
| Rucio Files Only in dCache [5] | 3 | 2 | 0 |
| Found in HPSS but with file size mismatch | 3 | 2 | 0 |

1. Files not matching directory, filename, size and checksum in dCache
2. Checksum value in Rucio is"" (double quote).
3. File exist in dCache but checksums in Rucio and dCache don't match
4. Rucio files in HPSS that definitely don't match files in dCache
5. Files found in Rucio that match dCache files but not HPSS files

Files effectively lost

Files Requiring Additional Investigation

Rucio files found only in HPSS: Checksum Verification

| | DATATAPE | MCTAPE | GROUPTAPE |
|------------------------------|----------|--------|-----------|
| Rucio Files Only in HPSS [1] | 4,661 | 0 | 0 |
| Verified as Correct [2] | 4,641 | 0 | 0 |
| Files with Problems | 20 | 0 | 0 |

1. Files with valid checksums in Rucio that are found only in HPSS
2. Recalculated checksum of file in HPSS and compare with Rucio

Files effectively lost

Rucio State

| | DATATAPE | MCTAPE | GROUPTAPE |
|--------------------------------|----------|--------|-----------|
| Rucio Files only found in HPSS | 5,923 | 0 | 0 |
| Rucio State : "Available" | 5,923 | 0 | 0 |

- A. All Rucio files found only in HPSS are marked as "Available" in Rucio. (includes files with invalid checksums in Rucio)

Files Requiring Additional Investigation

Rucio State of recently created files found in Rucio but not dCache or HPSS

| | DATATAPE | MCTAPE | GROUPTAPE |
|-----------------------------|----------|--------|-----------|
| # In Flight Files | 4,226 | 7,286 | 0 |
| Rucio State : "Available" | 0 | 0 | 0 |
| Rucio State : "Copying" | 4,201 | 7,286 | 0 |
| Rucio State : "Unavailable" | 25 | 0 | 0 |
| Rucio State : "Bad" | 0 | 0 | 0 |

“Old” Unmatched Rucio Files

Old (not recently created) files found in Rucio but not dCache or HPSS

| | DATATAPE | MCTAPE | GROUPTAPE |
|-----------------------------|----------|--------|-----------|
| “Old” Unmatched Rucio Files | 607 | 3,002 | 8,660 |
| Rucio State : "Available" | 350 | 325 | 226 |
| Rucio State : "Copying" | 7 | 50 | 29 |
| Rucio State : "Unavailable" | 250 | 2,623 | 8,405 |
| Rucio State : "Bad" | 0 | 4 | 0 |

Unmatched Rucio Files

Additional information about unmatched Rucio Files

| | DATATAPE | MCTAPE | GROUPTAPE |
|-----------------------------|----------|--------|-----------|
| "Old" Unmatched Rucio Files | 607 | 3,002 | 8,660 |
| # Datasets | 407 | 2653 | 8324 |
| Single file Dataset | 292 | 2301 | 8187 |
| Data Volume | 342 GB | 4 TB | 16 TB |

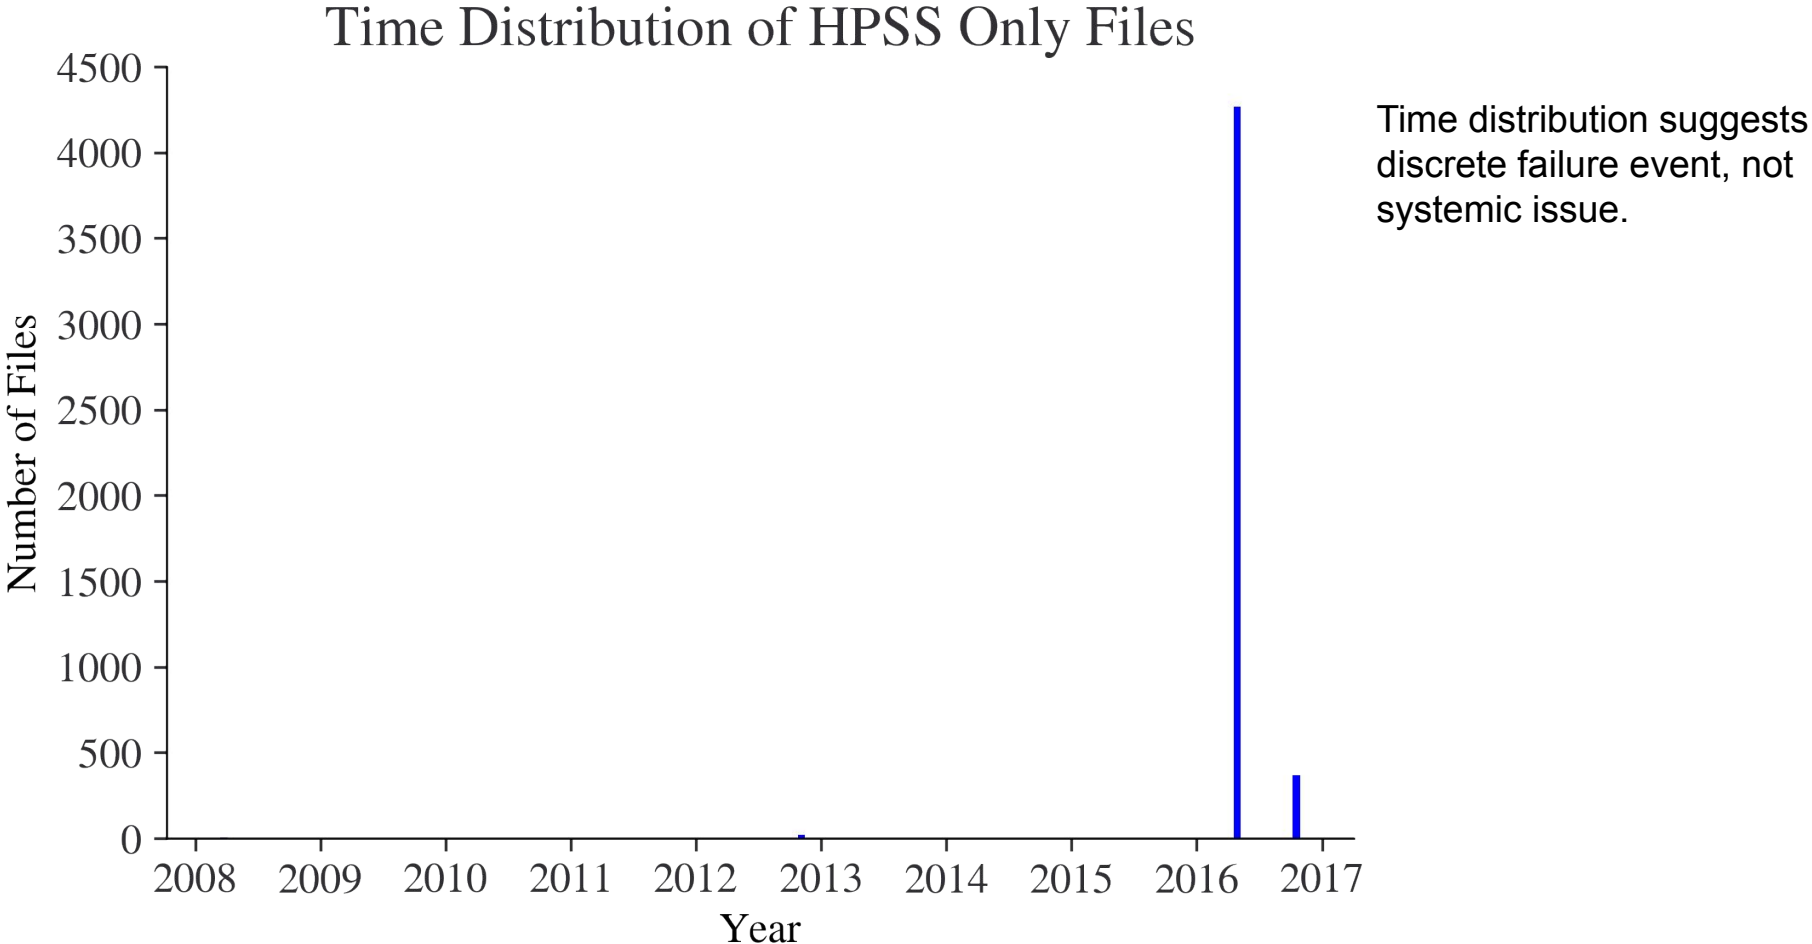
Current Status

Files requiring re-replication (if available at other sites)

| | DATATAPE [A] | MCTAPE [B] | GROUPTAPE |
|---|--------------|------------|-----------|
| Unmatched Rucio Files | 648 | 3,004 | 8,660 |
| % Unmatched Rucio Files | 0.0039% | 0.0133% | 0.2805% |
| % Unmatched Rucio Files (Aggregated) | 0.0291% | | |

- A. DATATAPE - Sum of unmatched Rucio files (607) and checksum mismatch with HPSS/dCache files (20 + 18 + 3)
- B. MCTAPE - Sum of unmatched Rucio files (3002) and Rucio files only in dCache (2)

Creation Time of Rucio Files Found Only in HPSS



Orphaned HPSS Files

Files found in HPSS but not registered in Rucio

| | DATATAPE | MCTAPE | GROUPTAPE |
|--|----------|---------|-----------|
| HPSS Files not in Rucio | 466,082 | 333,594 | 1,320 |
| # Files not in Rucio Datasets | 424,834 | 297,324 | 556 |
| # Files in Rucio Datasets | 41,248 | 36,270 | 764 |
| | | | |
| Total # Directories | 9,351 | 3,576 | 615 |
| # Directories that don't map to Rucio Datasets | 8,210 | 3,026 | 169 |
| # Directories that map to Rucio Datasets | 1,141 | 550 | 446 |


Orphaned dCache Files

Files found in dCache but not registered in Rucio

| | DATATAPE | MCTAPE | GROUPTAPE |
|--|----------|---------|-----------|
| dCache Files not in Rucio | 502,143 | 309,502 | 1,301 |
| # Files not in Rucio Datasets | 463,792 | 274,008 | 556 |
| # Files in Rucio Datasets | 38,351 | 35,494 | 745 |
| | | | |
| Total # Directories | 8,854 | 10,937 | 597 |
| # Directories that don't map to Rucio Datasets | 7,543 | 10,480 | 169 |
| # Directories that map to Rucio Datasets | 1,311 | 457 | 428 |

Selected File Creation Times: DATATAPE

Histogram of File Creation Time




Unmatched Rucio Files

Orphaned dCache Files

Orphaned HPSS Files

Selected File Creation Times: MCTAPE

Histogram of File Creation Time




Unmatched Rucio Files

Orphaned dCache Files

Orphaned HPSS Files

Selected File Creation Times: ATLASGROUPTAPE

Histogram of File Creation Time



Unmatched Rucio Files

Orphaned dCache Files

Orphaned HPSS Files

Considerations For The Future

- Automate namespace verification process, run periodically
 - Need better mechanism to extract relevant metadata from HPSS (Rumbler API in HPSS 8.2.0 ?)
 - Develop platform and processes to implement verification
- Investigate different checksum usages in HPSS
 - HPSS end to end data integrity (E2EDI)
 - HPSS user defined metadata for checksum storage if not full E2EDI
 - Ramification of backporting to existing data
 - There are pros and cons

Summary

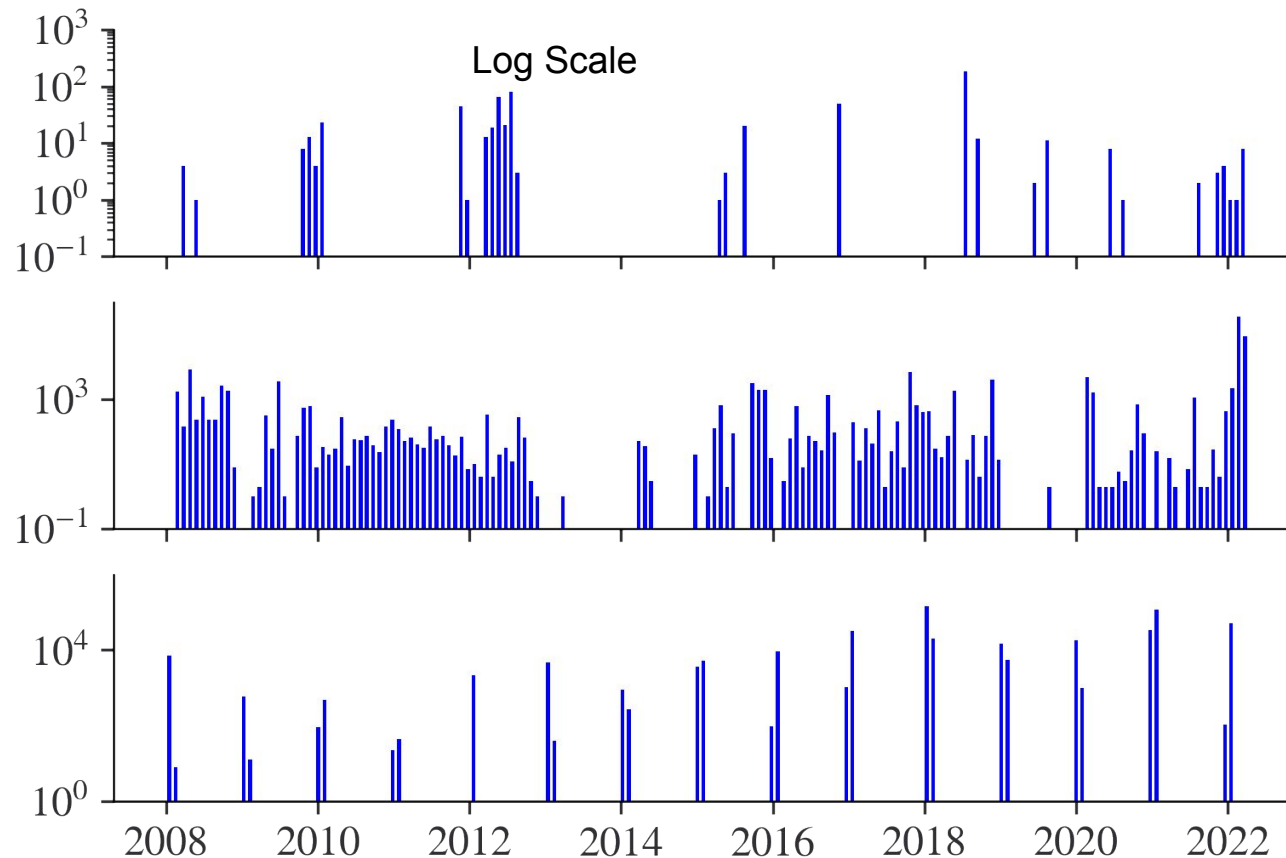
- 99.97% of files sent to the US ATLAS Tier 1 facility for archiving have been verified to be correctly archived
 - 99.996% of DATATAPE correctly archived
 - 99.987% of MCTAPE correctly archived
 - 99.72% of ATLASGROUPTAPE correctly archived
- Other issues identified by the namespace analysis
 - Adler32 and MD5 checksums computed on 1.19% of files

Backup Slides

1. Creation time histograms for DATATAPE, MCTAPE, and ATLASGROUPTAPE shown in log scale to highlight low level activity over time.
2. Creation time distribution for files stored in dCache with both MD5 and Adler32 checksums

Selected File Creation Times: DATATAPE

Histogram of File Creation Time




Unmatched Rucio Files

Orphaned dCache Files

Orphaned HPSS Files

Selected File Creation Times: MCTAPE

Histogram of File Creation Time



Unmatched Rucio Files

Orphaned dCache Files

Orphaned HPSS Files

Time Distribution of MD5 Generation

