



ATLAS Computing Status

Richard P Mount
Eric Lançon

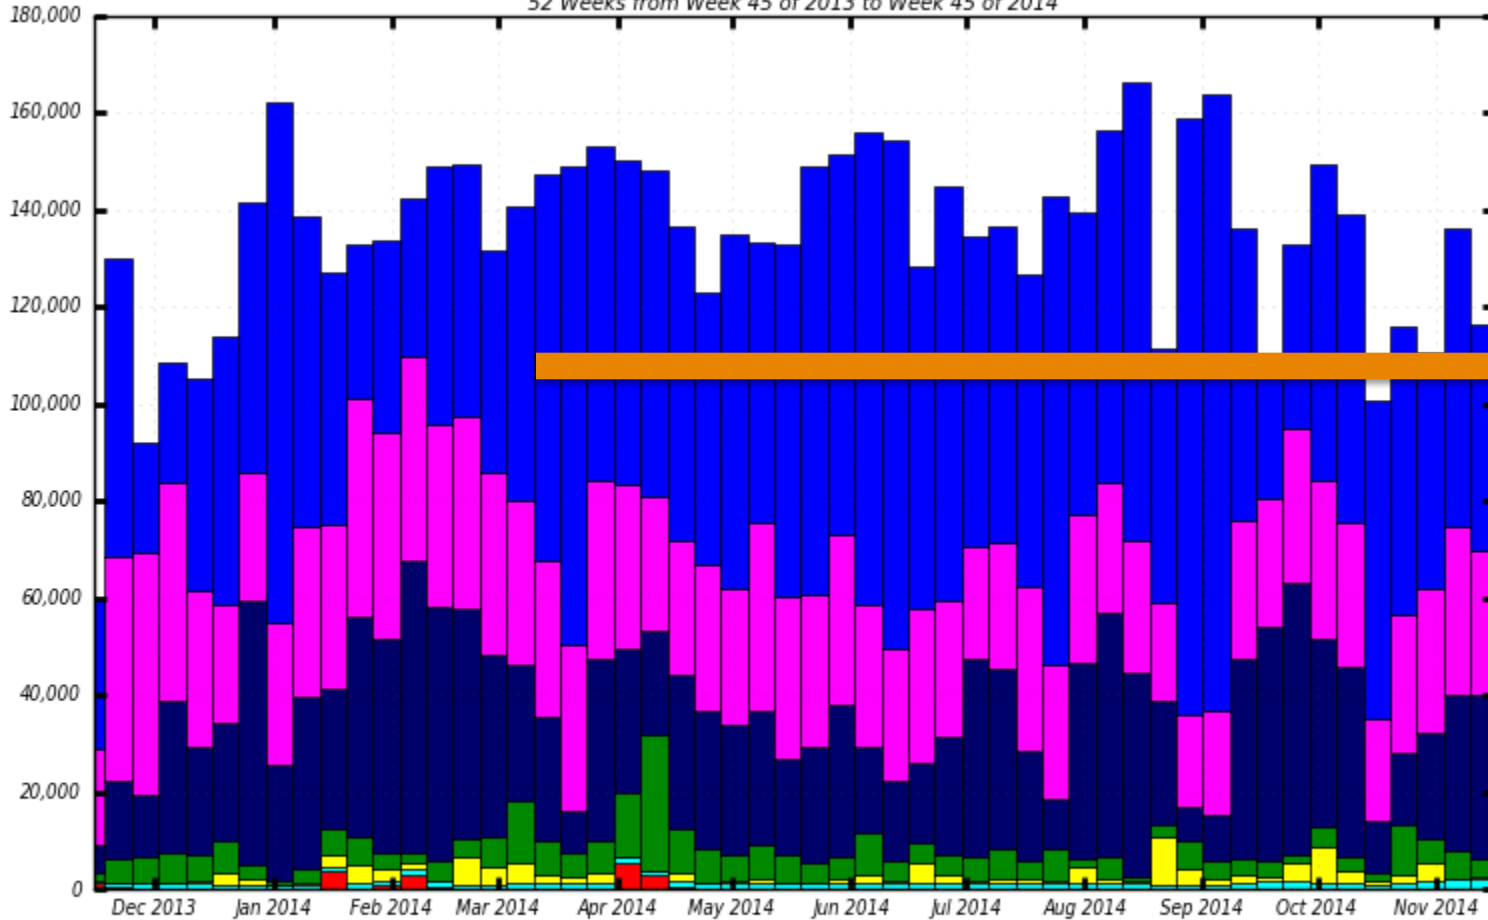


Tier 1 + Tier 2 CPU Usage in the last year



Slots of Running Jobs

52 Weeks from Week 45 of 2013 to Week 45 of 2014



Pledge

- MC Simulation
- Analysis
- MC Reconstruction
- Group Production
- Data Processing
- Others
- Extra Production
- unknown

Tier 1 + Tier 2 Disk Usage in the last year (excluding production buffers)




SLAC

← Total Pledge
T1+T2 = 90 PB

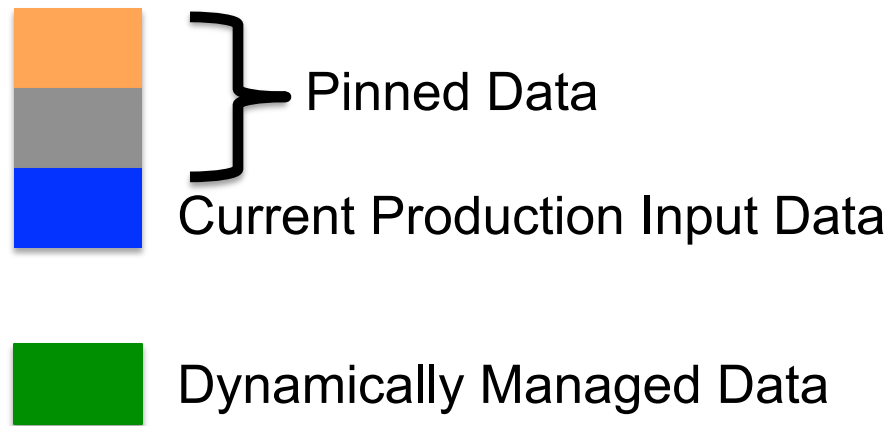
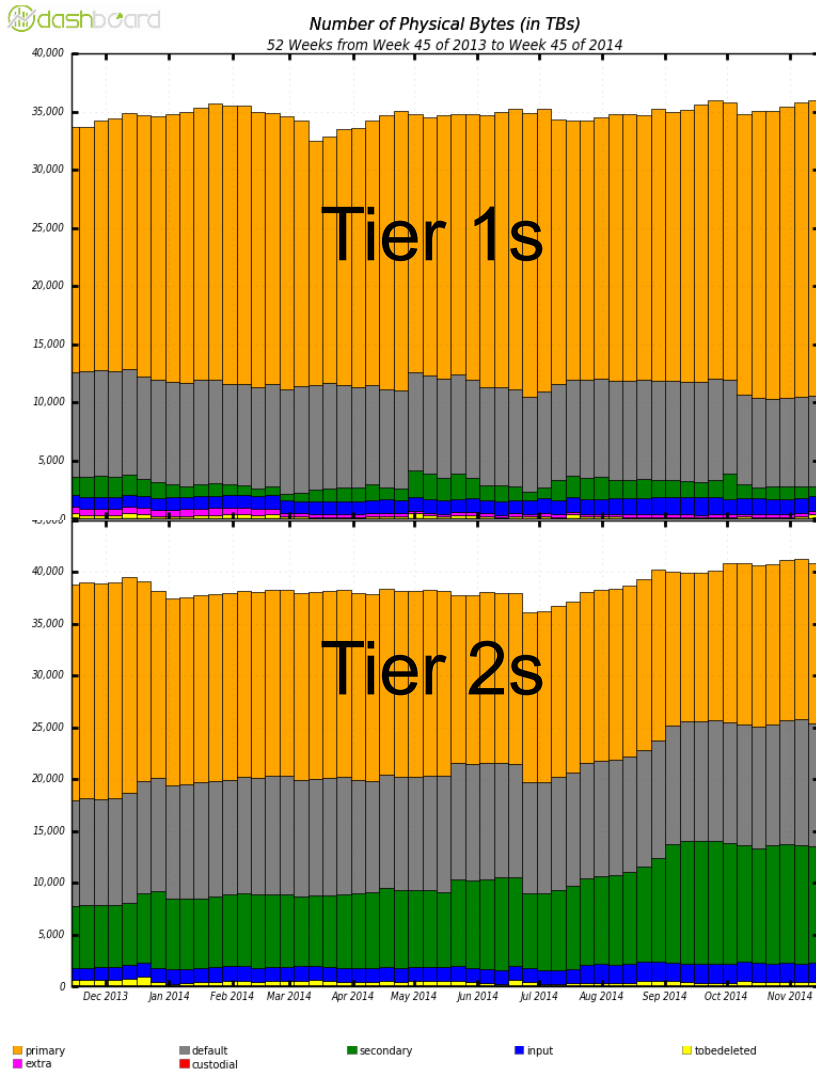


Number of Physical Bytes (in TBs)
52 Weeks from Week 45 of 2013 to Week 45 of 2014



Maximum: 77,045, Minimum: 0.00, Average: 72,047, Current: 76,884

Pinned and Dynamically Managed Data



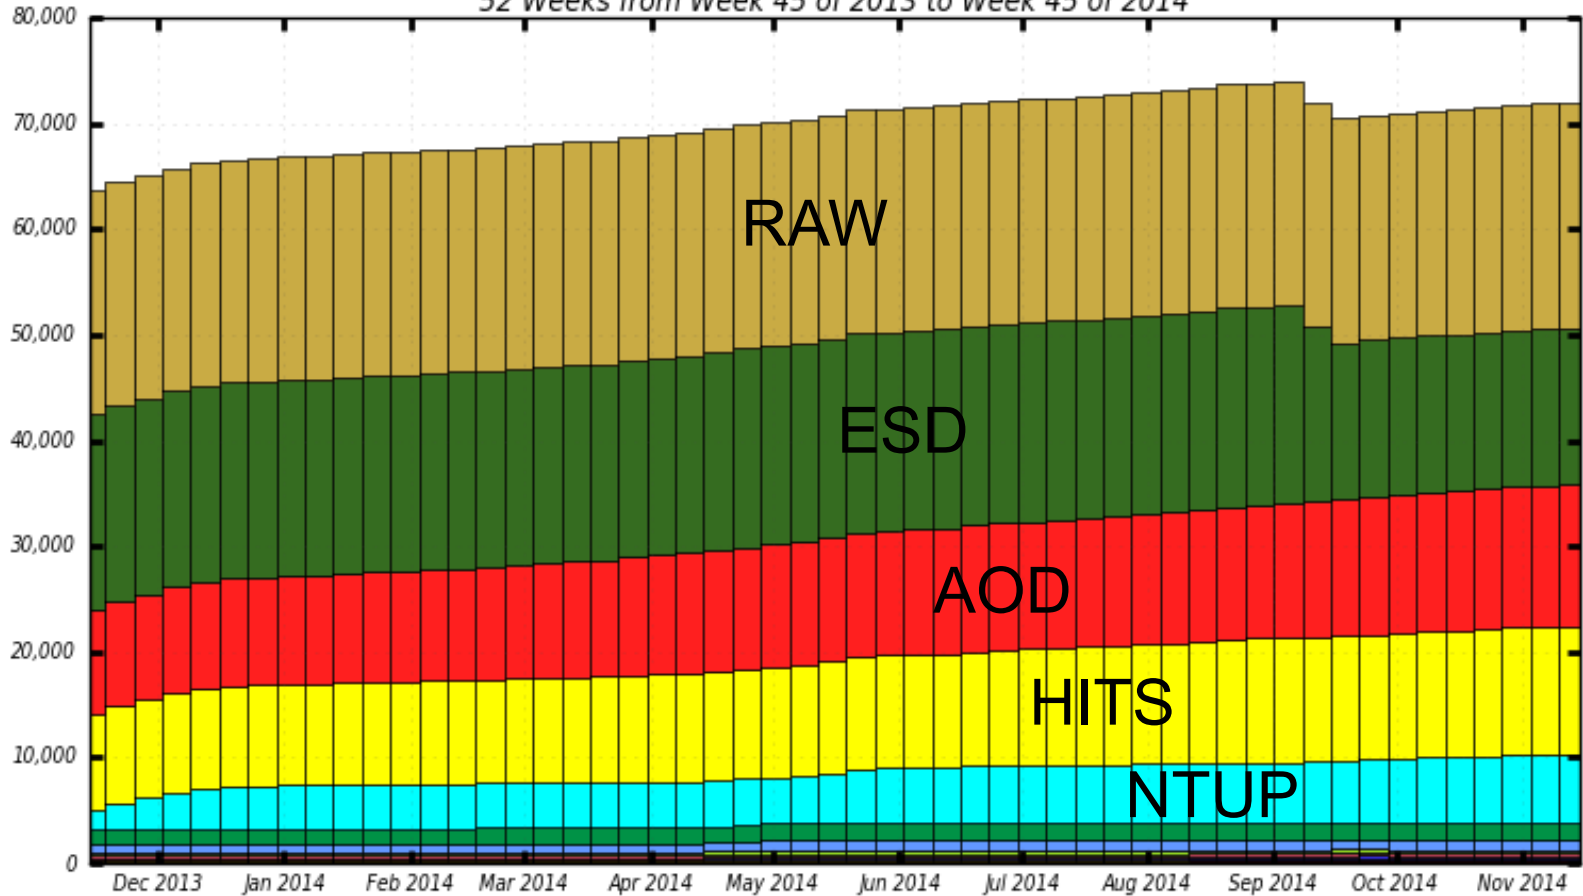
Disk-resident data with expired Lifetimes are marked as “Dynamically Managed” and will be deleted when the space is needed.

Additional replicas of popular datasets are also marked as “Dynamically Managed”

Tape Usage in the last year



Number of Physical Bytes (in TBs)
52 Weeks from Week 45 of 2013 to Week 45 of 2014





Addressing the Storage Issues (Disk and Tape)

New ATLAS Distributed Data Management Strategy

Approved by ATLAS in July

Cautious implementation in progress

1. All data to have a Lifetime. Delete from disk and tape at end of Lifetime. (ATLAS Physics has now proposed Lifetimes)
2. Disk versus tape residency to be managed algorithmically by ATLAS Distributed Computing (ADC)

Cautious Implementation – First model the effects – MC

MC to delete

(after applying the lifetime policy)

(older than lifetime & not touched for at least lifetime+6months)

| policy | lifetime | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3_ DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|--------|----------|--------------|--------------|--------------|------------------|---------|---------|-------------|
| RAW | 1200 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ESD | 2 | 3166 | 1544 | 1042 | 305 | 0 | 150 | 287 |
| RDO | 2 | 1747 | 0 | 140 | 1329 | 1 | 322 | 87 |
| HITS | 12 | 7819 | 54 | 733 | 0 | 242 | 7222 | 6 |
| AOD | 12 | 2247 | 55 | 261 | 378 | 63 | 1933 | 2 |
| DAOD | 12 | 5 | 0 | 1 | 0 | 0 | 0 | 7 |
| rest | 12 | 2226 | 42 | 942 | 23 | 0 | 654 | 1109 |
| | | | | | | | | |
| | | 17210 | 1695 | 3119 | 2035 | 306 | 10281 | 1498 |

Cautious Implementation – First model the effects – MC

MC to delete

(after applying the lifetime policy)

(older than lifetime & not touched for at least lifetime+6months)

| policy | lifetime | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3_ DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|--------|----------|--------------|--------------|--------------|------------------|---------|---------|-------------|
| RAW | 1200 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ESD | 2 | 3166 | 1544 | 1042 | 305 | 0 | 150 | 287 |
| RDO | 2 | 1747 | 0 | 140 | 1329 | 1 | 322 | 87 |
| HITS | 12 | 7819 | 54 | 733 | 0 | 242 | 7222 | 6 |
| AOD | 12 | 2247 | 55 | 261 | 378 | 63 | 1933 | 2 |
| DAOD | 12 | 5 | 0 | 1 | 0 | 0 | 0 | 7 |
| rest | 12 | 2226 | 42 | 942 | 23 | 0 | 654 | 1109 |
| | | 17210 | 1695 | 3119 | 2035 | 306 | 10281 | 1498 |



Cautious Implementation – First model the effects – MC

SLAC

MC to delete

(after applying the lifetime policy)

(older than lifetime & not touched for at least lifetime+6months)

| policy | lifetime | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3_ DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|--------|----------|--------------|--------------|--------------|------------------|---------|---------|-------------|
| RAW | 1200 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ESD | 2 | 3166 | 1544 | 1042 | 305 | 0 | 150 | 287 |
| RDO | 2 | 1747 | 0 | 140 | 1329 | 1 | 322 | 87 |
| HITS | 12 | 7819 | 54 | 733 | 0 | 242 | 7222 | 6 |
| AOD | 12 | 2247 | 55 | 261 | 378 | 63 | 1933 | 2 |
| DAOD | 12 | 5 | 0 | 1 | 0 | 0 | 0 | 7 |
| rest | 12 | 2226 | 42 | 942 | 23 | 0 | 654 | 1109 |
| | | 17210 | 1695 | 3119 | 2035 | 306 | 10281 | 1498 |

Group Disks: temporary technical stay of execution

Cautious Implementation – First model the effects – MC



SLAC

MC to delete

(after applying the lifetime policy)

(older than lifetime & not touched for at least lifetime+6months)

| policy | lifetime | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3_ DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|--------|----------|--------------|--------------|--------------|------------------|---------|---------|-------------|
| RAW | 1200 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ESD | 2 | 3166 | 1544 | 1042 | 305 | 0 | 150 | 287 |
| RDO | 2 | 1747 | 0 | 140 | 1329 | 1 | 322 | 87 |
| HITS | 12 | 7819 | 54 | 733 | 0 | 242 | 7222 | 6 |
| AOD | 12 | 2247 | 55 | 261 | 378 | 63 | 1933 | 2 |
| DAOD | 12 | 5 | 0 | 1 | 0 | 0 | 0 | 7 |
| rest | 12 | 2226 | 42 | 942 | 23 | 0 | 654 | 1109 |
| | | 17210 | 1695 | 3119 | 2035 | 306 | 10281 | 1498 |

T0,1,2 Disk: delete 7 of 50 PB

Cautious Implementation – First model the effects – MC



SLAC

MC to delete

(after applying the lifetime policy)

(older than lifetime & not touched for at least lifetime+6months)

| policy | lifetime | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3_ DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|--------|----------|--------------|--------------|--------------|------------------|---------|---------|-------------|
| RAW | 1200 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ESD | 2 | 3166 | 1544 | 1042 | 305 | 0 | 150 | 287 |
| RDO | 2 | 1747 | 0 | 140 | 1329 | 1 | 322 | 87 |
| HITS | 12 | 7819 | 54 | 733 | 0 | 242 | 7222 | 6 |
| AOD | 12 | 2247 | 55 | 261 | 378 | 63 | 1933 | 2 |
| DAOD | 12 | 5 | 0 | 1 | 0 | 0 | 0 | 7 |
| rest | 12 | 2226 | 42 | 942 | 23 | 0 | 654 | 1109 |
| | | 17210 | 1695 | 3119 | 2035 | 306 | 10281 | 1498 |

T0,1 Tape: delete 11 of 27 PB

Cautious Implementation – First model the effects – Data

Data

| policy | life time | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3 DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|----------------|-----------|--------------|--------------|--------------|-----------------|---------|---------|-------------|
| Express_stream | 4 | 309 | 0 | 0 | 0 | 300 | 1 | 1 |
| ESD_ph_str | 3 | 33 | 0 | 21 | 0 | 0 | 6 | 0 |
| AOD_ph_str | 24 | 25 | 0 | 3 | 0 | 20 | 0 | 0 |
| ESD_AOD_hi | 24 | 697 | 36 | 13 | 0 | 200 | 482 | 0 |
| ESDAOD_sm_str1 | 24 | 414 | 2 | 2 | 2 | 367 | 27 | 16 |
| ESDAOD_sm_str2 | 36 | 1 | 0 | 0 | 0 | 1 | 0 | 0 |
| TAG | 60 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| log_HIST | 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DESD_DAOD | 24 | 850 | 5 | 227 | 5 | 530 | 113 | 19 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| NTUP1 | 12 | 312 | 9 | 27 | 0 | 133 | 3 | 200 |
| NTUP2 | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | | 2641 | 52 | 293 | 7 | 1551 | 632 | 236 |

Months

Cautious Implementation – First model the effects – Data

Data

| policy | life time | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3 DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|----------------|-----------|--------------|--------------|--------------|-----------------|---------|---------|-------------|
| Express_stream | 4 | 309 | 0 | 0 | 0 | 300 | 1 | 1 |
| ESD_ph_str | 3 | 33 | 0 | 21 | 0 | 0 | 6 | 0 |
| AOD_ph_str | 24 | 25 | 0 | 3 | 0 | 20 | 0 | 0 |
| ESD_AOD_hi | 24 | 697 | 36 | 13 | 0 | 200 | 482 | 0 |
| ESDAOD_sm_str1 | 24 | 414 | 2 | 2 | 2 | 367 | 27 | 16 |
| ESDAOD_sm_str2 | 36 | 1 | 0 | 0 | 0 | 1 | 0 | 0 |
| TAG | 60 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| log_HIST | 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DESD_DAOD | 24 | 850 | 5 | 227 | 5 | 530 | 113 | 19 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| NTUP1 | 12 | 312 | 9 | 27 | 0 | 133 | 3 | 200 |
| NTUP2 | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | | 2641 | 52 | 293 | 7 | 1551 | 632 | 236 |

Months

T0,1,2 Disk: delete 0.3 of 24 PB

Cautious Implementation – First model the effects – Data

Data

| policy | life time | logical size | T0 DATA DISK | T1 DATA DISK | T2_T3 DATA DISK | T0 TAPE | T1 TAPE | GROUP DISKS |
|----------------|-----------|--------------|--------------|--------------|-----------------|---------|---------|-------------|
| Express_stream | 4 | 309 | 0 | 0 | 0 | 300 | 1 | 1 |
| ESD_ph_str | 3 | 33 | 0 | 21 | 0 | 0 | 6 | 0 |
| AOD_ph_str | 24 | 25 | 0 | 3 | 0 | 20 | 0 | 0 |
| ESD_AOD_hi | 24 | 697 | 36 | 13 | 0 | 200 | 482 | 0 |
| ESDAOD_sm_str1 | 24 | 414 | 2 | 2 | 2 | 367 | 27 | 16 |
| ESDAOD_sm_str2 | 36 | 1 | 0 | 0 | 0 | 1 | 0 | 0 |
| TAG | 60 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| log_HIST | 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| DESD_DAOD | 24 | 850 | 5 | 227 | 5 | 530 | 113 | 19 |
| DRAW | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| NTUP1 | 12 | 312 | 9 | 27 | 0 | 133 | 3 | 200 |
| NTUP2 | 36 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | | 2641 | 52 | 293 | 7 | 1551 | 632 | 236 |

Months

T0,1 Tape: delete 2 of 22 non-RAW PB



Summary of expected Lifetime-driven deletions

Disk: ~7 PB

Tape: ~13 PB

| Space available at the start of Run 2: | Disk PB | Tape PB |
|--|---------|---------|
| Freed space as above | 7 | 13 |
| New 2015 pledged space | 8 | 18 |



SLAC

Run 2 (and beyond) Model



Lifetime-based Storage Model

1. Use updated version of current model to calculate CPU usage and dataset volumes
2. Keep track of creation date (quarter) of each dataset
2. Remove each dataset from disk and tape when it has exceeded its (type-specific) Lifetime.
3. Adjust total request for disk and tape to be sufficient to accommodate the requested lifetimes.
4. Assign an experience-driven fraction of the disk+tape storage to disk.

Model is now being refined

Will be used to inform the ATLAS resource request for 2016 in February 2015



Brief Software Status (1)

Reconstruction:

- Factor 3 speedup goal achieved. (old news)

New Analysis Model (more efficient and more user-friendly):

- “Dual-use” xAOD replaces separate ATHENA-readable and Root-readable formats. (in use)
- Derivation Framework (trains with carriages provided by groups) replaces incoherent “group production” (being validated)
- Analysis Framework, supporting standardized use of performance group recommendations (jet energy scale etc.) (Framework in use, performance group tools on track)

Brief Software Status (2)

Simulation:

- Integrated Simulation Framework : (in use)

Full Geant4 Simulation:

- Remains the (slow) “gold standard”
- Ongoing work to speed up the ATLAS implementation
- Geant4 version 10 (multithreaded) implementation (not for Run-2 start; timetable not clear yet)

Fast Simulation:

- New version of FastCaloSim (“good enough for taus”) (development ongoing, completion of validation second half of 2015)
- Fast digitization, truth seeded reconstruction (development ongoing)



Distributed Computing Software Status

New Production System ProdSys2:

- Will replace ProdSys1 on December 1, 2014

New Distributed Data Management System Rucio:

- Will replace DQ2 on December 1, 2014



Summary

Usage

CPU – Well above pledge level

Disk – Tier 1s still have too much pinned data

Disk – Tier 2s are “healthy”

Tape – “Full”, but cleanup in progress

New Distributed Data Management Strategy

- Cautious implementation is underway

Run 2 Computing Model

- Under development

Software developments

- Great progress: (all except FastSim on track for Run 2)