



EOSPUBLIC → EOSAMS02

migration strategy discussion

Jaroslav Guenther
Thursday 8.4.2021
(CERN IT-ST-PDS)

A decorative graphic in the bottom left corner consisting of a complex network of thin, light blue lines radiating from a central point, resembling a particle detector or a data network.

> EOSPUBLIC

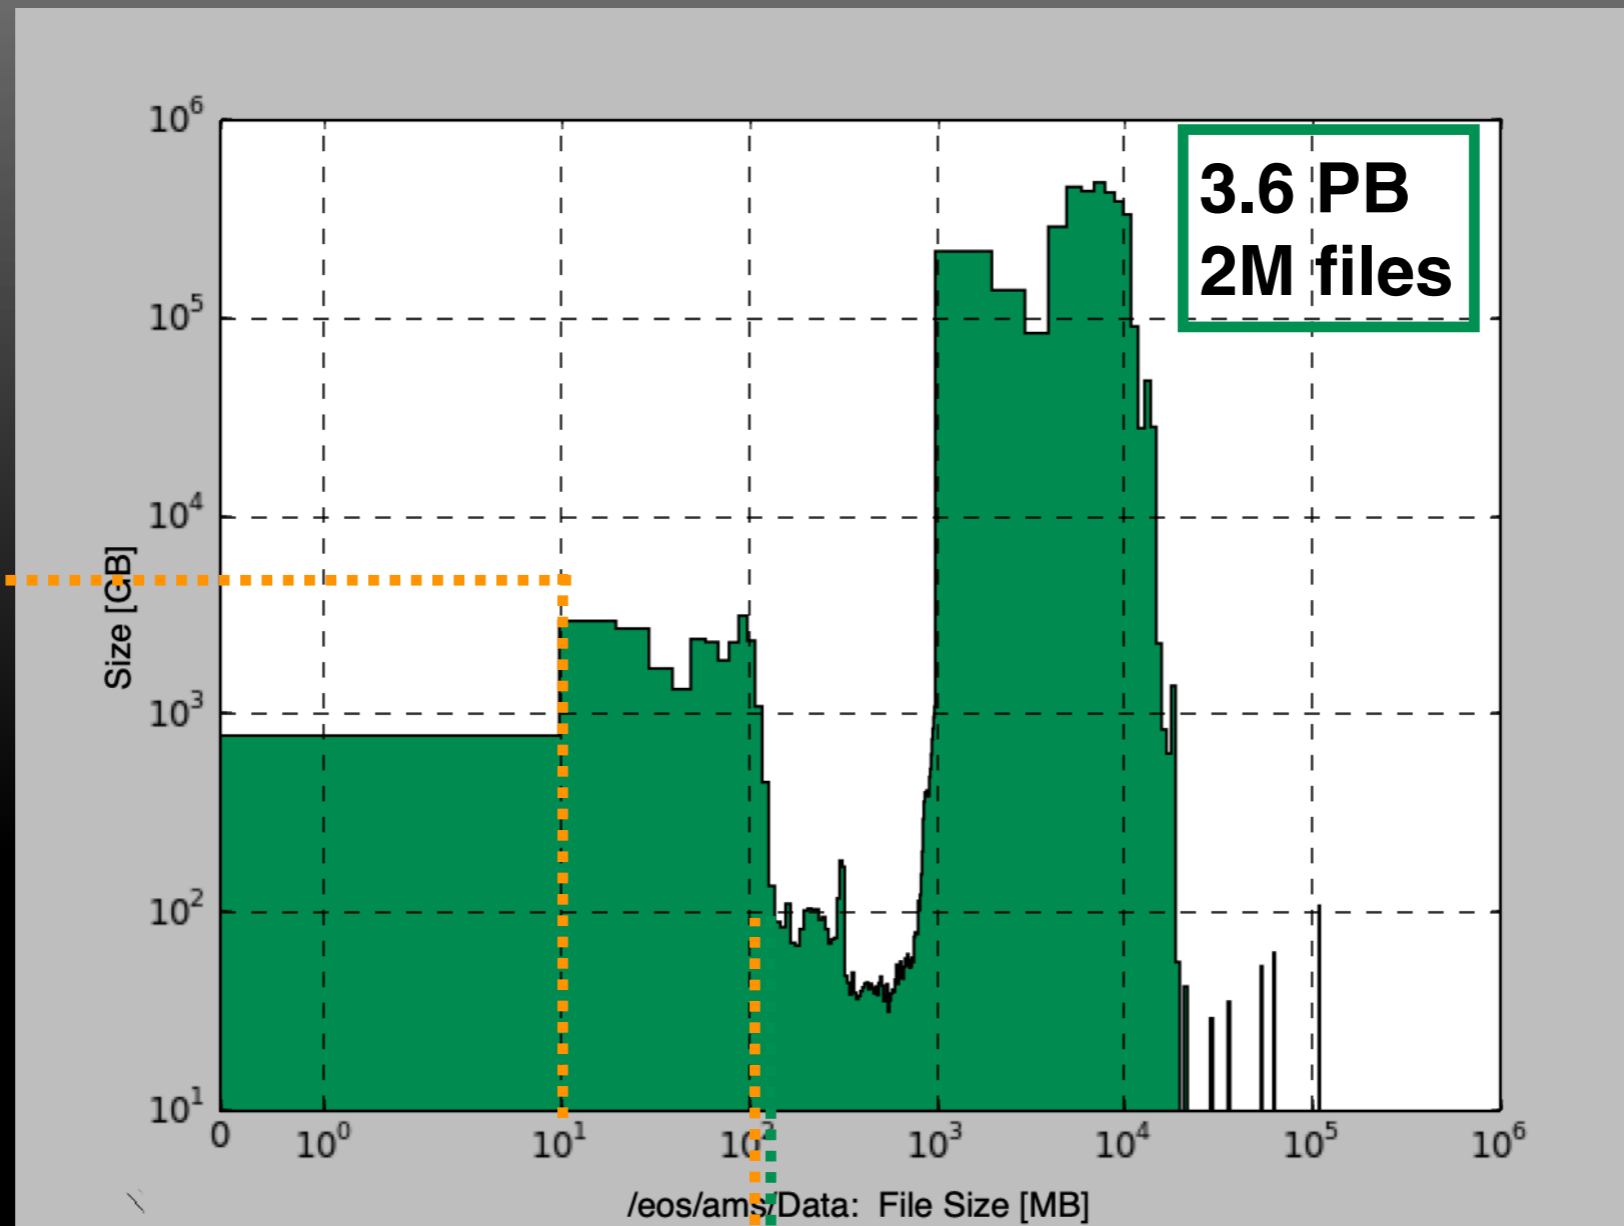
quota node	size	files	quota
/eos/ams/	8.8 PB	22 M	12 PB
~/user	574 TB	18.4 M	875 TB
~/group/mitep/	371 TB	2.8 M	467 TB
~/group/mit/	254 TB	1.1 M	300 TB
~/group/dbar/	114 TB	4 M	120 TB
~/group/alpha/	4.8 TB	66 K	5 TB
~/afs/	2.4 TB	2.7 M	30 TB
Total	10.1 PB	51 M	13.8 PB

	EOSPUBLIC		EOSAMS02	
quota node	size	files	size	Files
/eos/ams/	8.8 PB	22 M	632 TB	5 M
~/user	574 TB	18.4 M	510 TB	16 M
~/group/mitep/	371 TB	2.8 M	371 TB	2.8 M
~/group/mit/	254 TB	1.1 M	238 TB	1.1 M
~/group/dbar/	114 TB	4 M	114 TB	4 M
~/group/alpha/	4.8 TB	66 K	4.8 TB	66 K
~/afs/	2.4 TB	2.7 M	2.4 TB	2.7 M
Total	10.1 PB	51 M	1.9 PB	31.6 M

> EOSAMS

- not enough space (6.9 PB raw space)
- considering erasure encoding for big files under /eos/ams

File Size Distribution /eos/ams/Data



< 10 MB
~1 MB / file
1 TB / 700 k

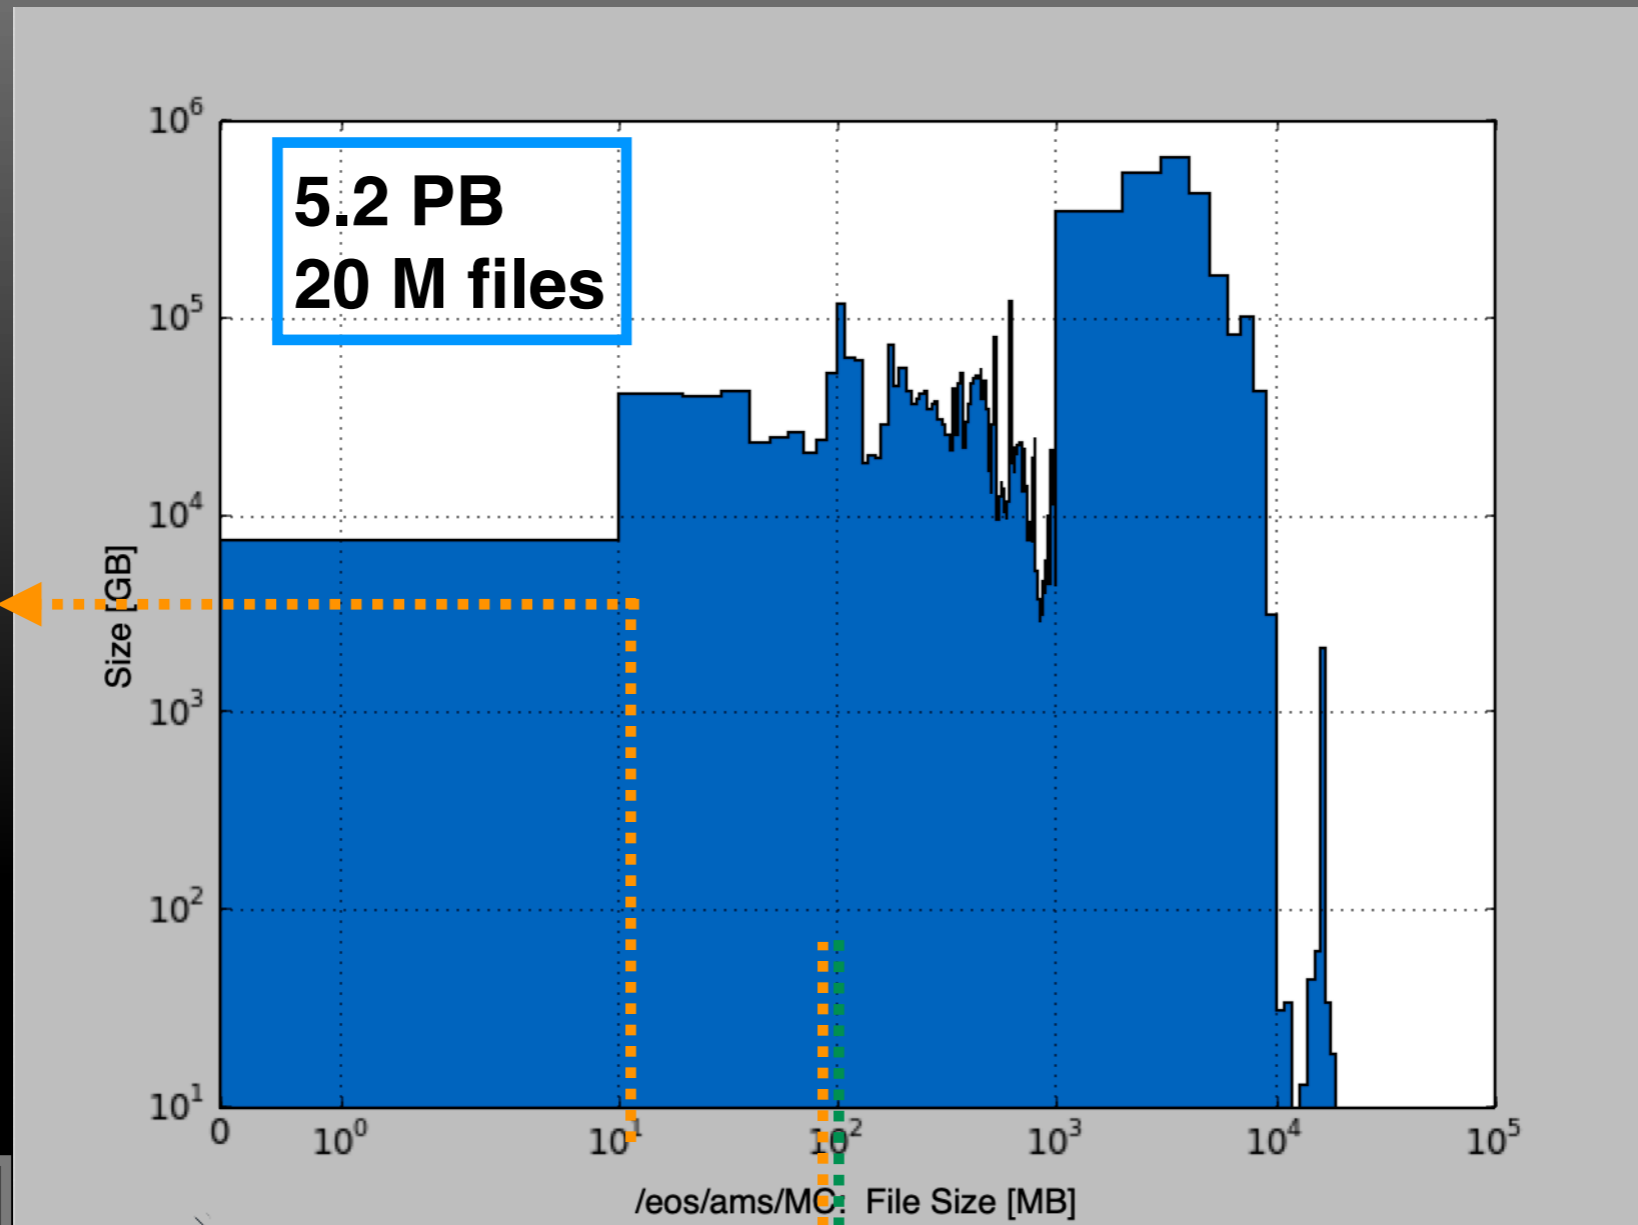
< 100 MB
~18 MB / file
22 TB / 1.2 M

> separate small file areas:
0-size, calibration, root files,
logs ... but also some RawData

> most volume from big files
→ suited for erasure encoding

> 100 MB
4.5 GB / file
3.6 PB / 800 k

File Size Distribution /eos/ams/MC



< 10 MB
~3 MB / file
8 TB / 2.3 M

< 100 MB
~28 MB / file
310 TB / 11 M

5.2 PB
20 M files

> most volume from big files
→ suited for erasure encoding

> 100 MB
636 MB / file
4.9 PB / 7.7 M

- > EOSAMS does not have enough capacity
 - > new HW delivery late June/August
 - > erasure encoding for big files
 - helps with the lack of disk space
 - read performance increase
 - > moving new pledge from EOSPUBLIC to EOSAMS
 - > freeze EOSPUBLIC → use EOSAMS02 and “add what is missing”
 - > use EOSPUBLIC → gradually migrate to EOSAMS02 behind EOSAMS alias in parallel (might not be possible for all areas) ...
(freezing AMS sub-trees of EOSPUBLIC)




> migration strategy is use-case dependent

link to some statistics from last year

- areas where we could switch to EOSAMS02 already ?
- areas of “write once read many” ?
- up-coming campaigns of data rewrites ?

> final sync - EOS downtime vs AMS buffer ?




> Who can we contact to resolve:

- entries owned by non existing user
- entry owners who lost write / browse permissions etc
- too many sub directories created

../2011/2011/2011/2011/2011/2011/2011;

- AFS - afs space migration was a test
 - make fresh copy from source again ?

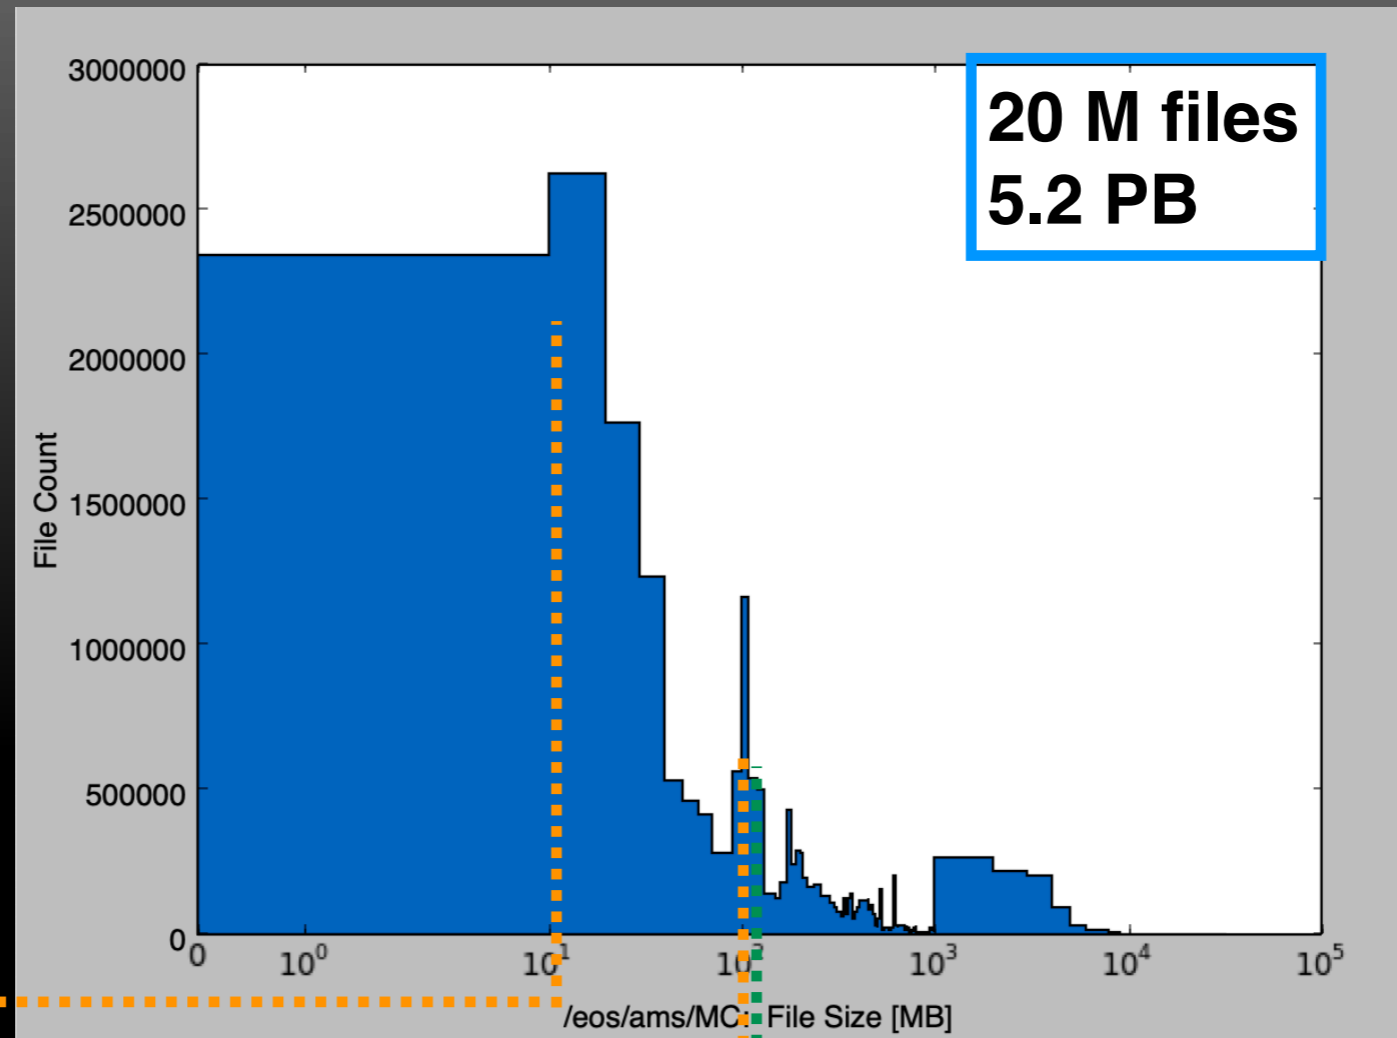


Thank you for your time!

A decorative graphic in the bottom-left corner consisting of a complex, overlapping grid of thin white lines forming a semi-circular shape.

Many Thanks for help @Cristian, @Elvin, @Luca, @Andreas !

File Size Distribution /eos/ams/MC



~3 MB / file
8 TB / 2.3 M

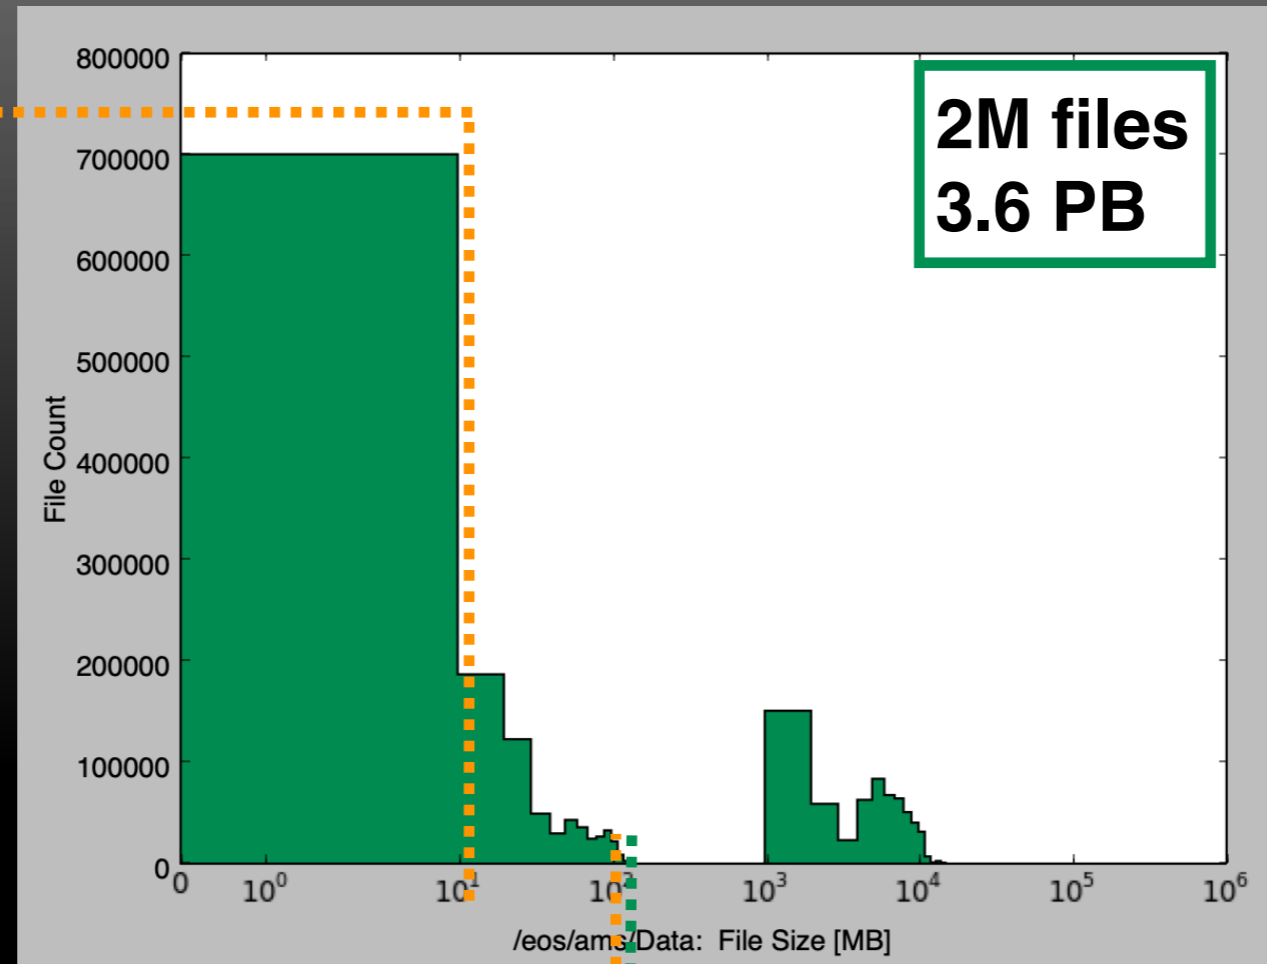
~28 MB / file
310 TB / 11 M

636 MB / file
4.9 PB / 7.7 M

File Size Distribution /eos/ams/Data



< 10 MB
~1 MB / file
1 TB / 700 k



< 100 MB
~18 MB / file
22 TB / 1.2 M

> 100 MB
4.5 GB / file
3.6 PB / 800 k

> small files: 0-size, calibration,
root files, logs ... but also some
RawData