FTS Community talk: LHCb

XrootD/FTS workshop 2023

Ben Couturier, Christophe Haen

FTS at the core of LHCb DMS

- All Third Party Copy transfers of LHCb go through FTS
  - Transfers dominated by Real Data flows
  - Jobs upload data to their final destination (i.e. no rebalancing campaign)
- Everything orchestrated by DIRAC
- As usual, follow the KISS principle
  - Little to no use of “fancy features”
- Reliability, performance & stability are our main requirements

Real Data distribution (most common workflow)

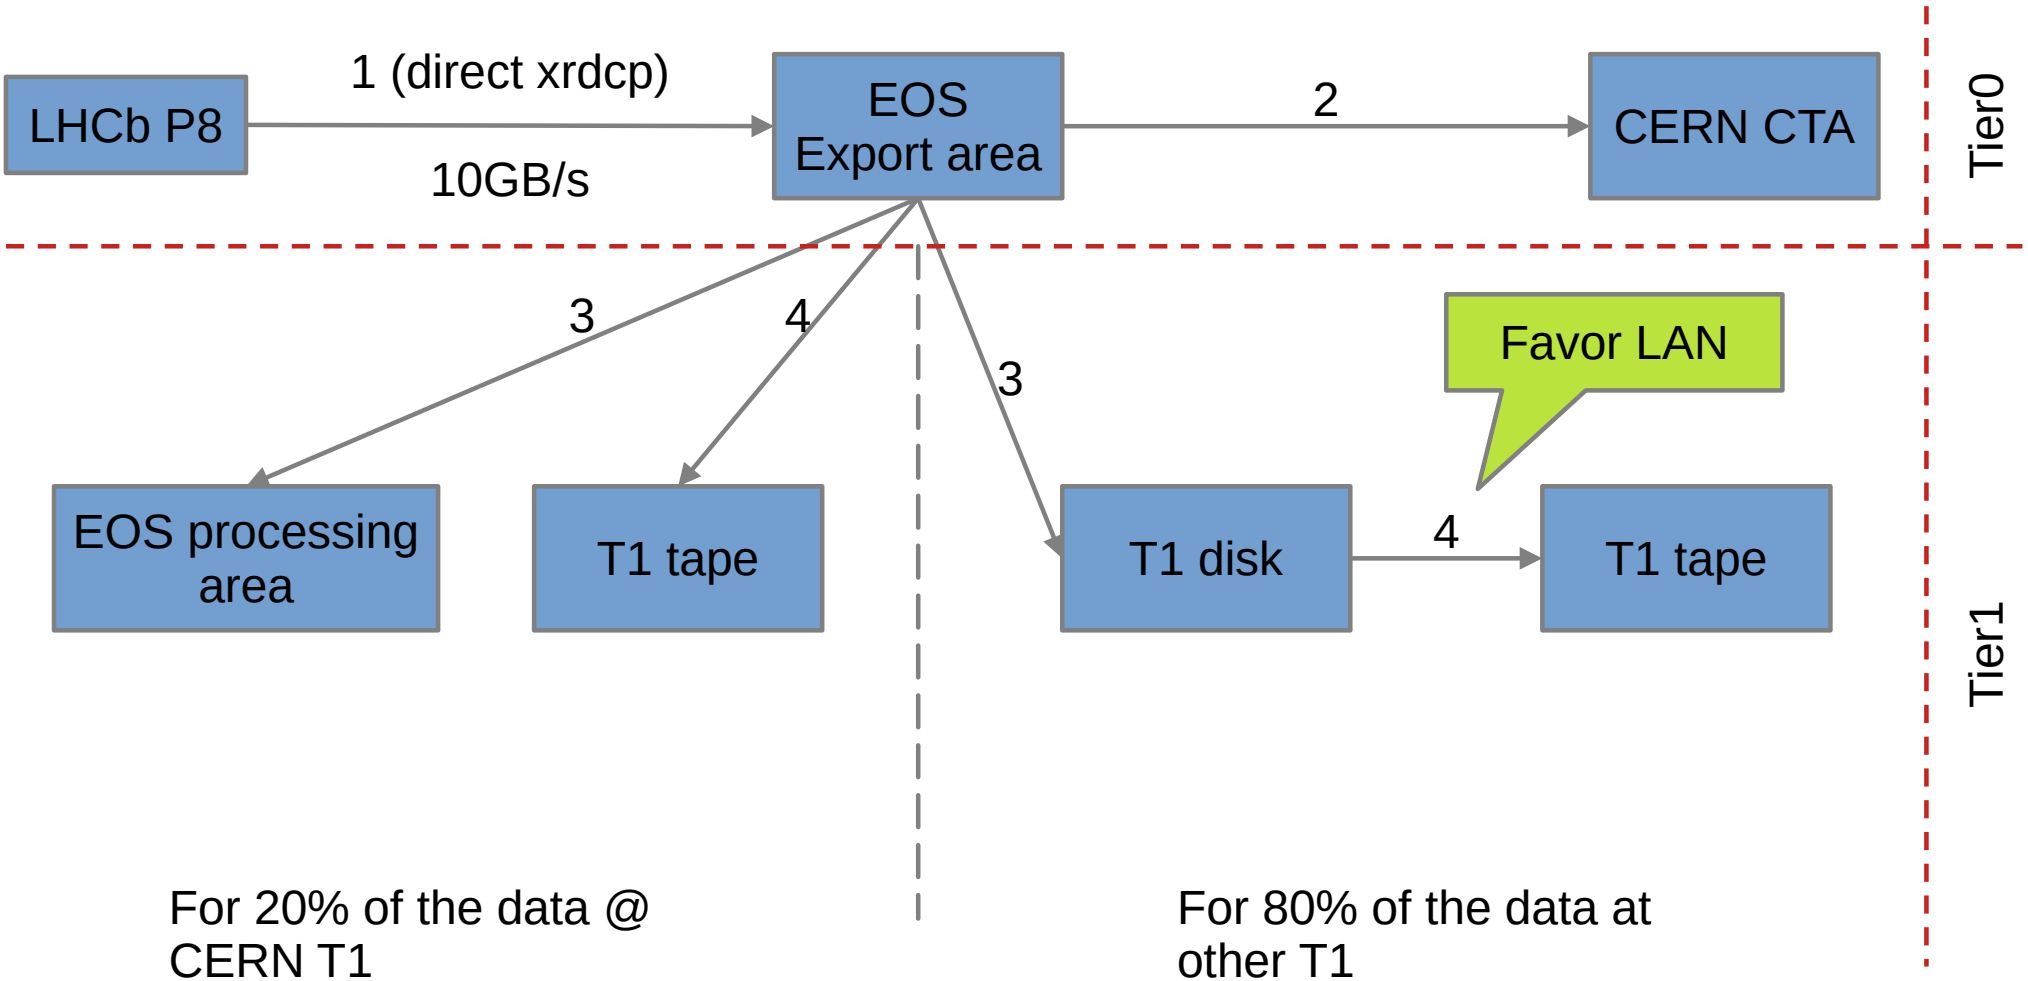

End of year reprocessing

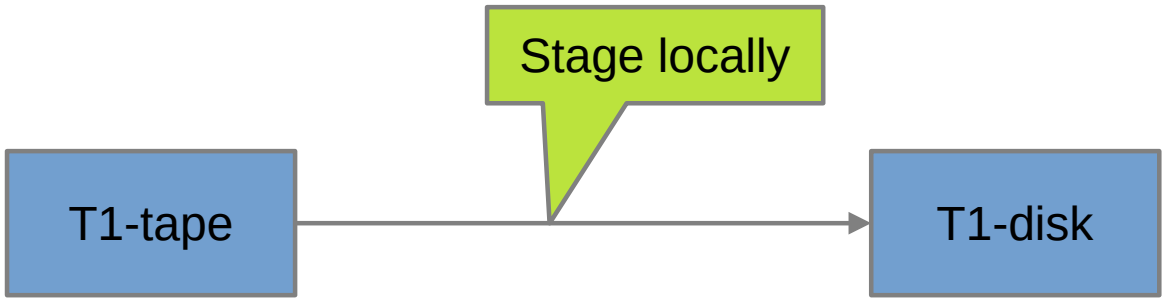

DC 2022 figures

- Again, favor LAN transfers
- Valid for T0 too

Write tests: CERN disk → T1 disk → T1 tape	
Site	expected Speed (GB/s)
CERN	11
CNAF	1.72
GRIDKA	2.23
IN2P3	1.25
NCBJ	1.32
PIC	0.2
RAL	2.96
RRCKI	0.25
SARA	1.07

Read tests T1 tape → T1 disk	
Site	expected Speed (GB/s)
CERN	1.90
CNAF	1.35
GRIDKA	1.36
IN2P3	0.98
NCBJ	0.91
PIC	0.17
RAL	1.93
RRCKI	0.21
SARA	0.74

Final data/MC distribution

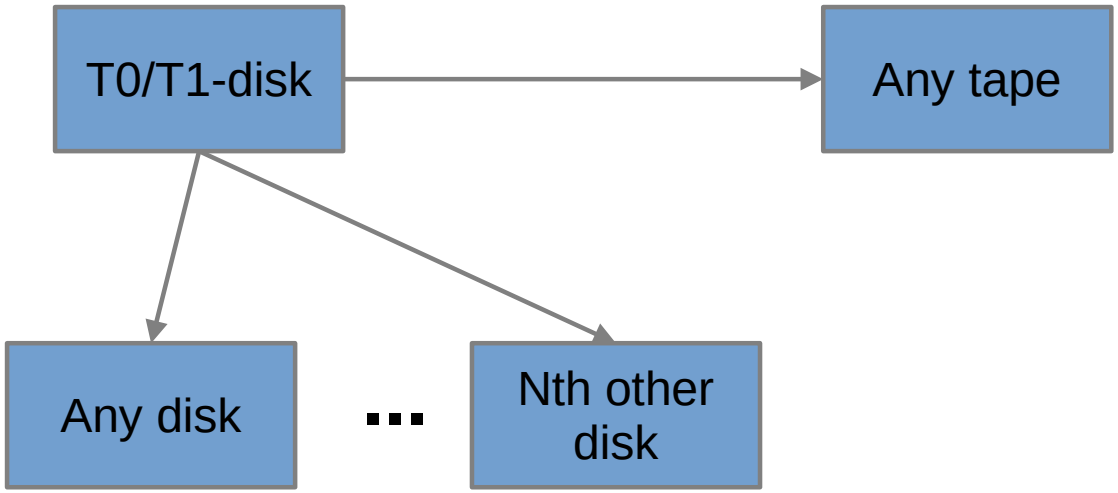

- Jobs upload their output on 1 disk
- Replicate from that disk with FTS
  - 1 archive to any tape
  - N replicas to any disk (T0,T1,T2)
  - Initially, N=2 (can be reduced depending on popularity)

Protocols

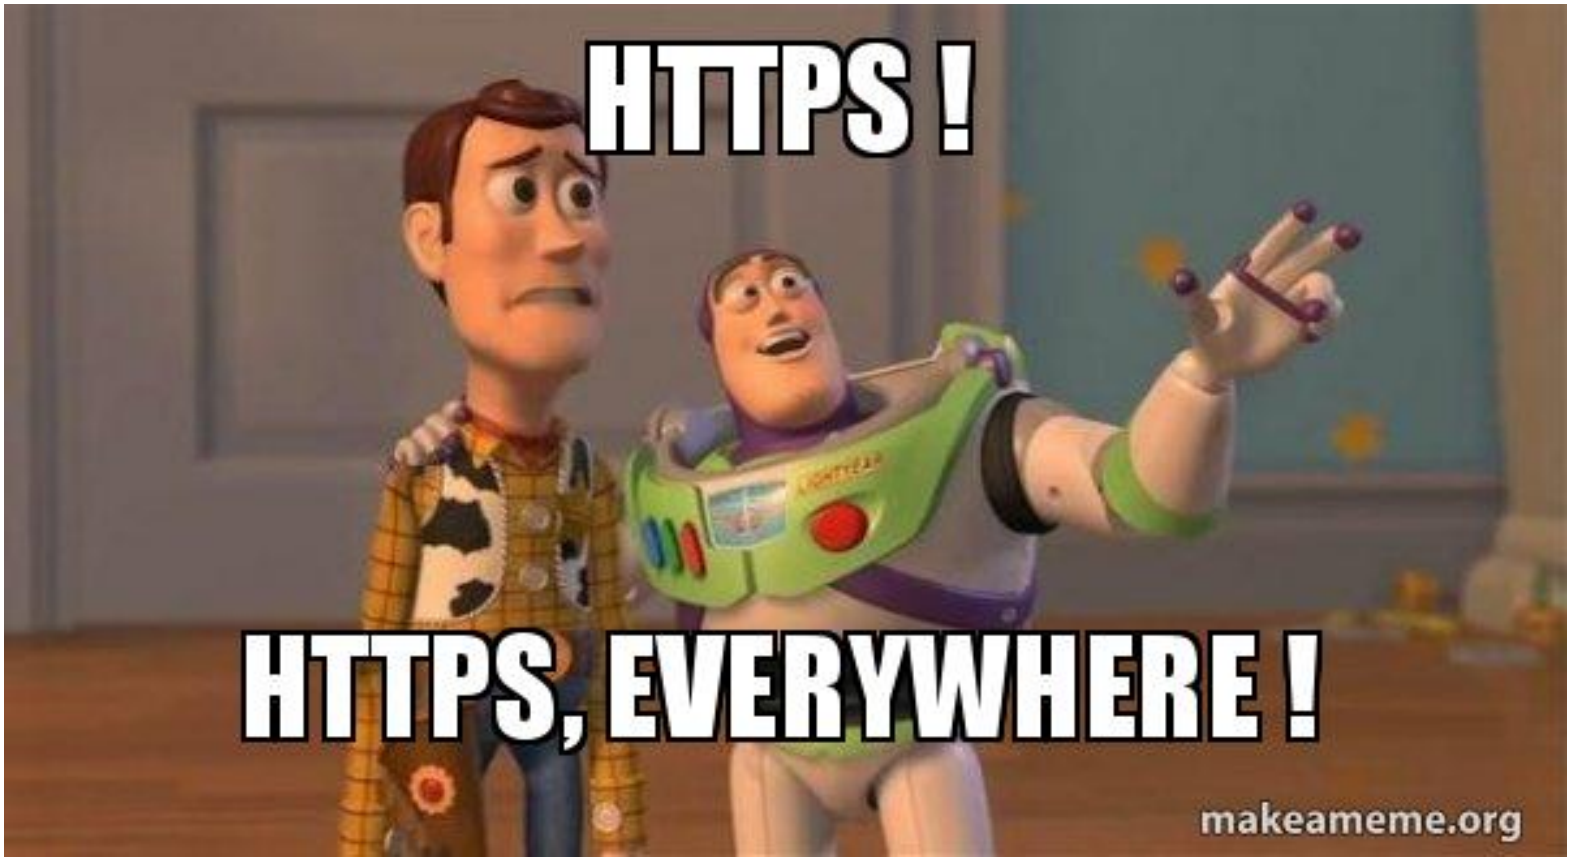

Almost….

Protocols

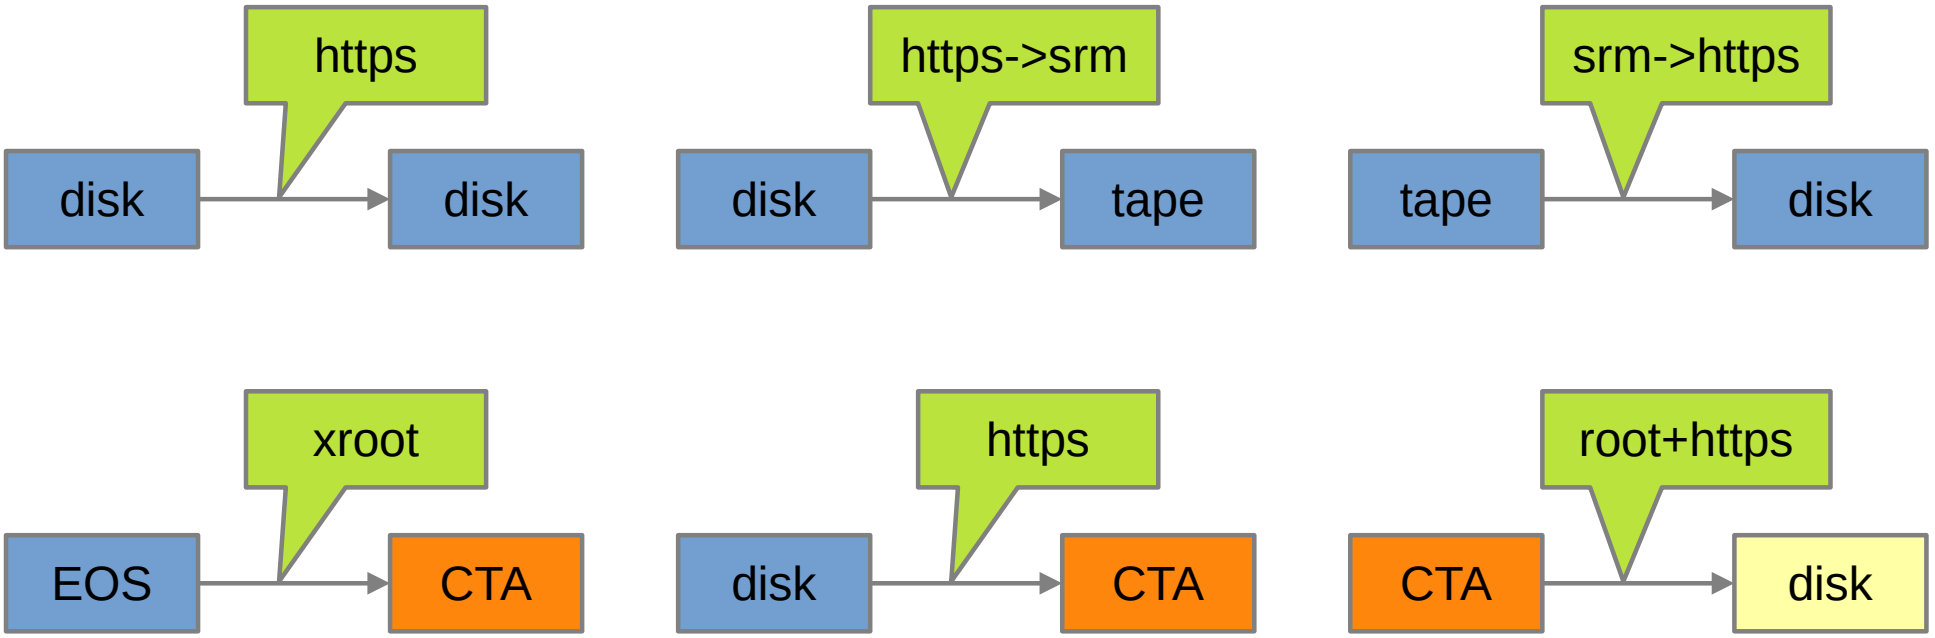

- CTA → disk: special multihop transfer
  - root://cta/myFile.raw → root://cta/myFile.raw
  - https://cta/myFile.raw → https://disk/myFile.raw
- Replaced for 1 week with HTTPs Tape API

Configuration

- Default except for T0 T1 links (empirical outcome of Data Challenge)

Link configuration

Parameters per link. If only source or only destination is specified, it applies to any transfer from/to that storage.

	Symbolic name	Source	Destination	Streams	Min Actives	Max Actives	Optimizer Mode	TCP buffer size	Disable delegation
	*	*	*	0	20	200	2	0	No
	gsiftp://EOS	gsiftp://eosll	*	1	20	200	2	0	No
	https://ANTAF	https://antar	*	1	20	200	0	0	No
	https://ANTAF	https://antar	https://webdav.	1	200	200	0	0	No
	HTTPCTALH(	https://eosc	*	0	50	200	3	0	No
	https://eos-nc	https://eoslh	https://se.cis.gc	1	150	300	0	0	No

~~Misslanous,misclanous,misclenaous~~, OTHER !!

- Single FTS instance at CERN
  - Proved to be very reliable and sufficient
  - No configuration discrepancy
  - If lasting downtime: just recreate another cluster
- Plan to use FTS for pinning files on tape for jobs (i.e. tape → tape cache)
- Activity: currently used only for DC, but interested in adding more
- Priority: not now, but if needed, we can
- Tape metadata: no immediate use case, will think of it

Feedback

- Very happy
  - At least, as happy as one can be doing DMS….
- Very stable and reliable
- Looking forward to use the new “explain” feature during the next DC
- No major request (for now :-))
- FTS is absolutely paramount to LHCb computing
  - This is a subliminal message to the IT dpt management…