

Storage Class Use

CERN

Tier-1 centres

Breakdown of access patterns/rates

- CERN
- a “typical” Tier-1

CERN

Tape1Disk0

- RAW data
- rDST data processed at CERN
- DST & MC data for archival

~50% RAW data will be accessed again for re-processing in a year

All rDST data will be accessed additional twice a year for stripping

CERN

Tape1Disk1

- DST data for analysis
- MC data

DST data will migrate from T1D1 class to T1D0 as later copies of the stripping are performed

CERN

Tape0Disk1

- Shared user analysis output files

Envisage a user/group quota

Large turnover of data

Tier1

Tape1Disk0

- RAW data
- rDST data processed at centre
- DST & MC data for archival

~50% RAW data will be accessed again for re-processing in a year

All rDST data will be accessed additional twice a year for stripping

Tier-1

Tape1Disk1

- DST data for analysis
- MC data

DST data will migrate from T1D1 class to T1D0 as later copies of the stripping are performed

Tier-1

Tape0Disk1

- Shared user analysis output files

Envisage a user/group quota

Large turnover of data

- (partial) DST data for analysis

will be removed from disk as later version become available

Tier-2

In general small requirement on disk at Tier-2's - buffer space for MC production

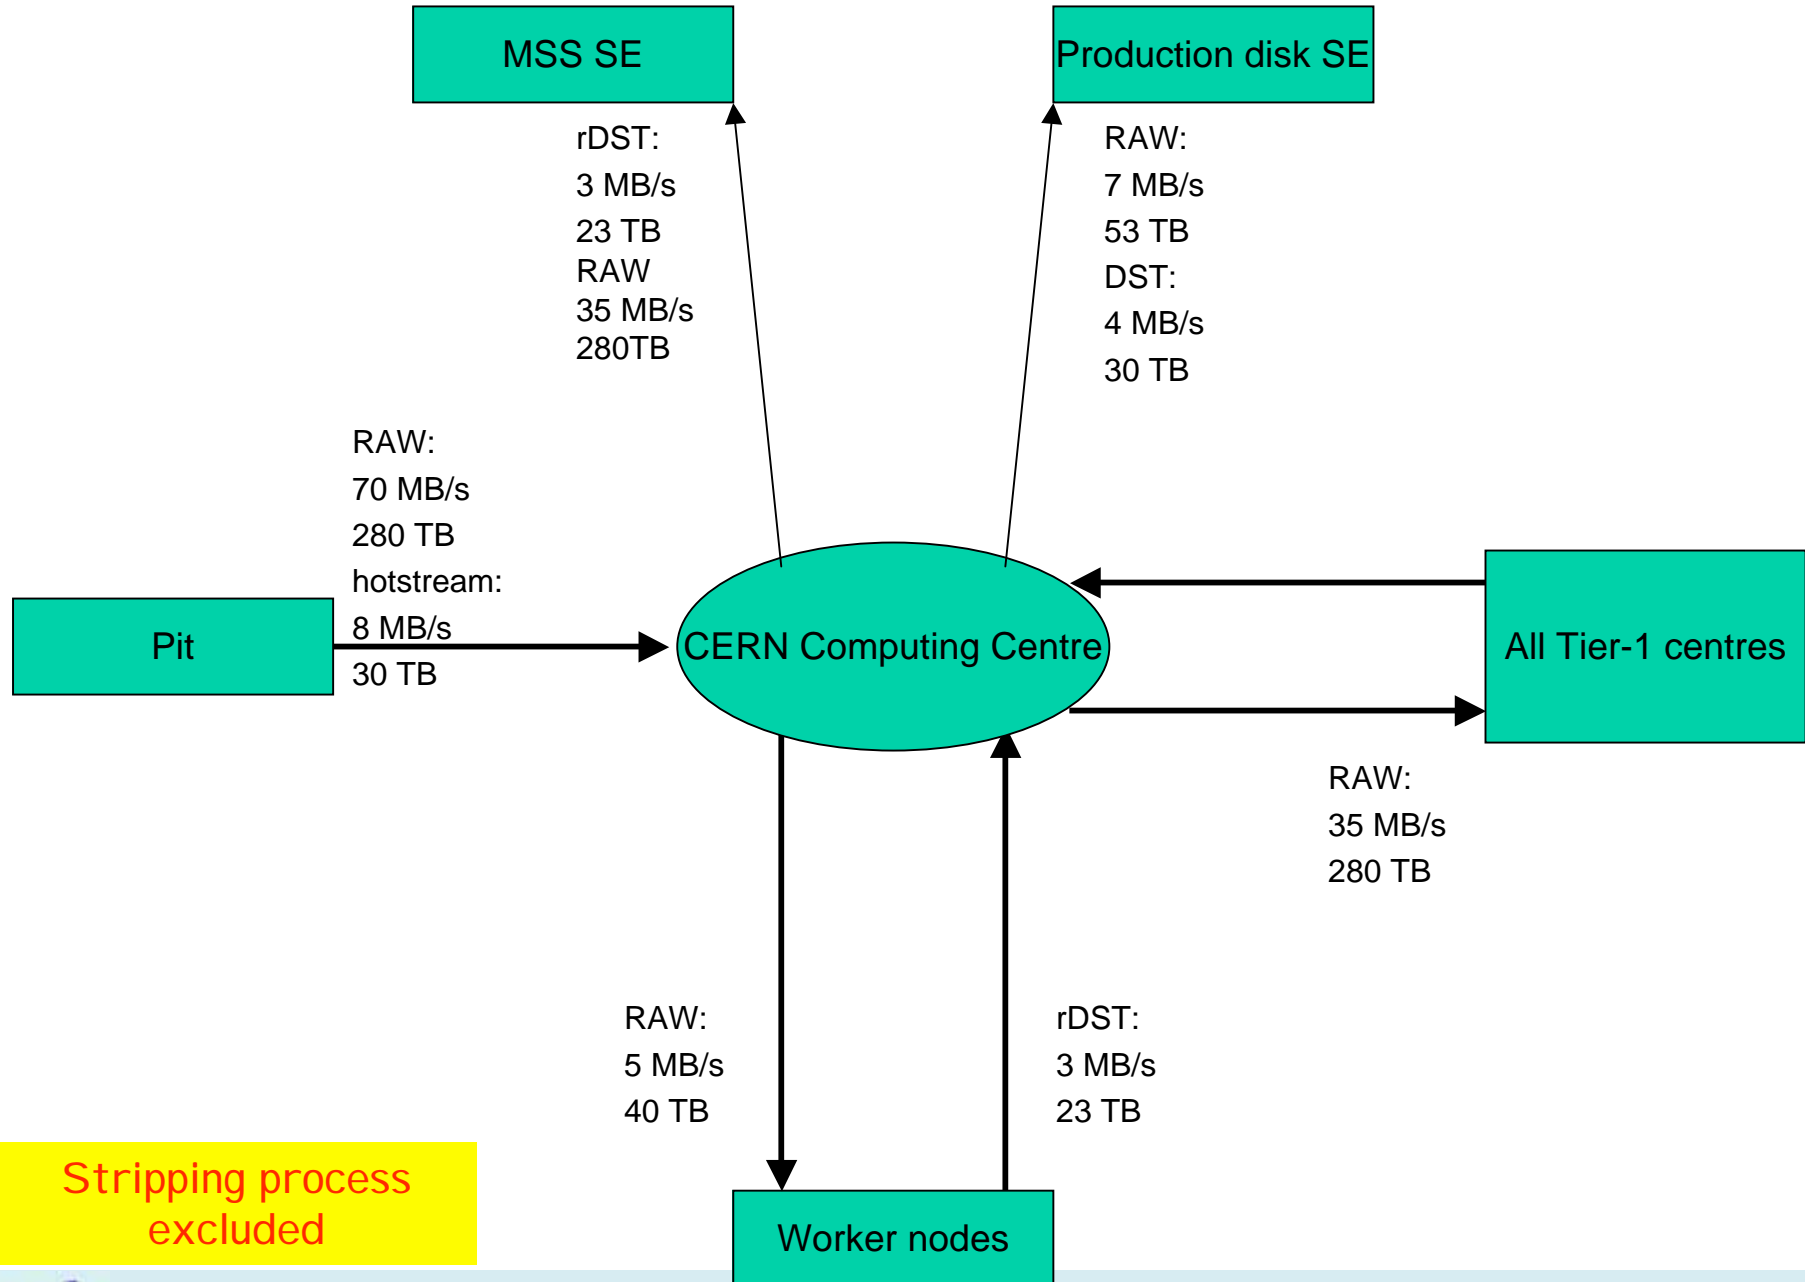
2 LHCb Tier02's request (have resources) to support analysis

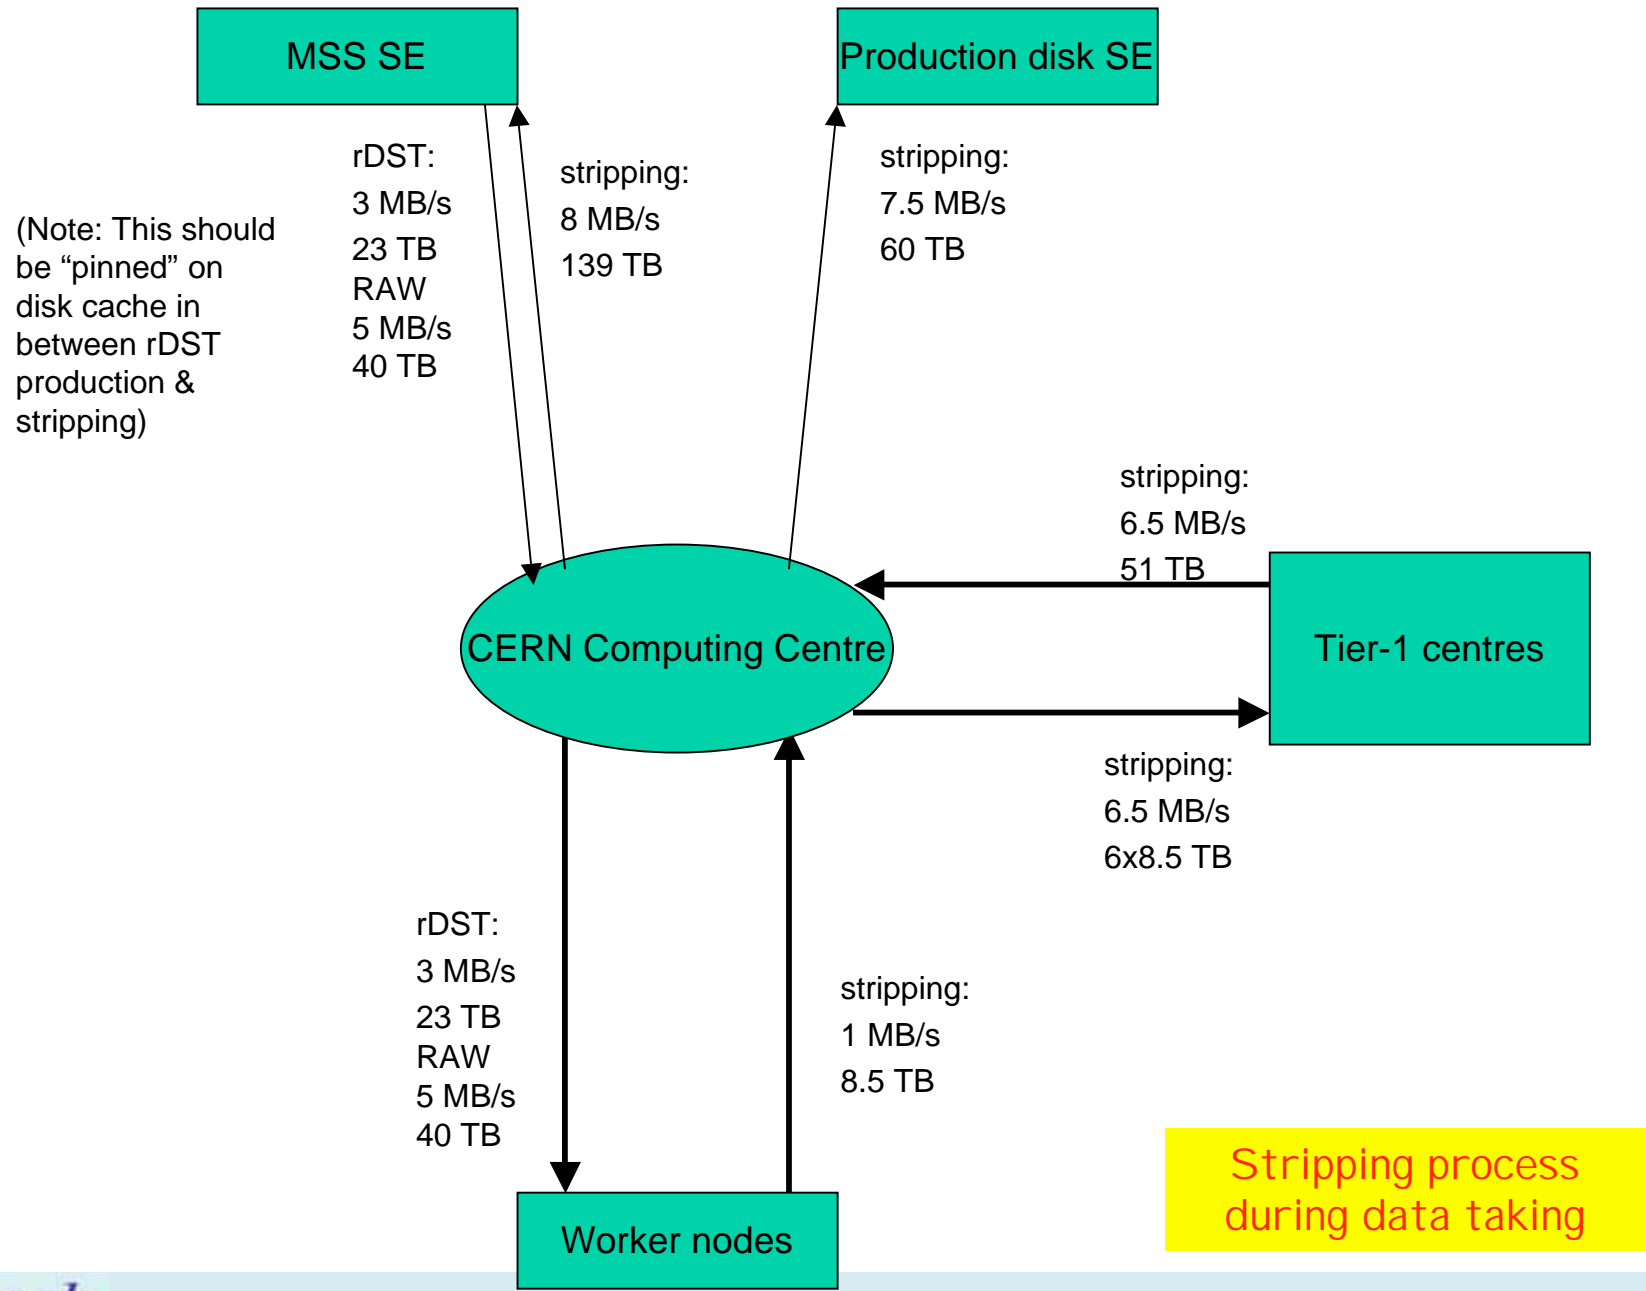
- Minimum ~90 TB of disk in 2008
- Depending on storage available datasets replaced up to 4 times per year

LHCb Computing Model - A Pictorial Guide 2008

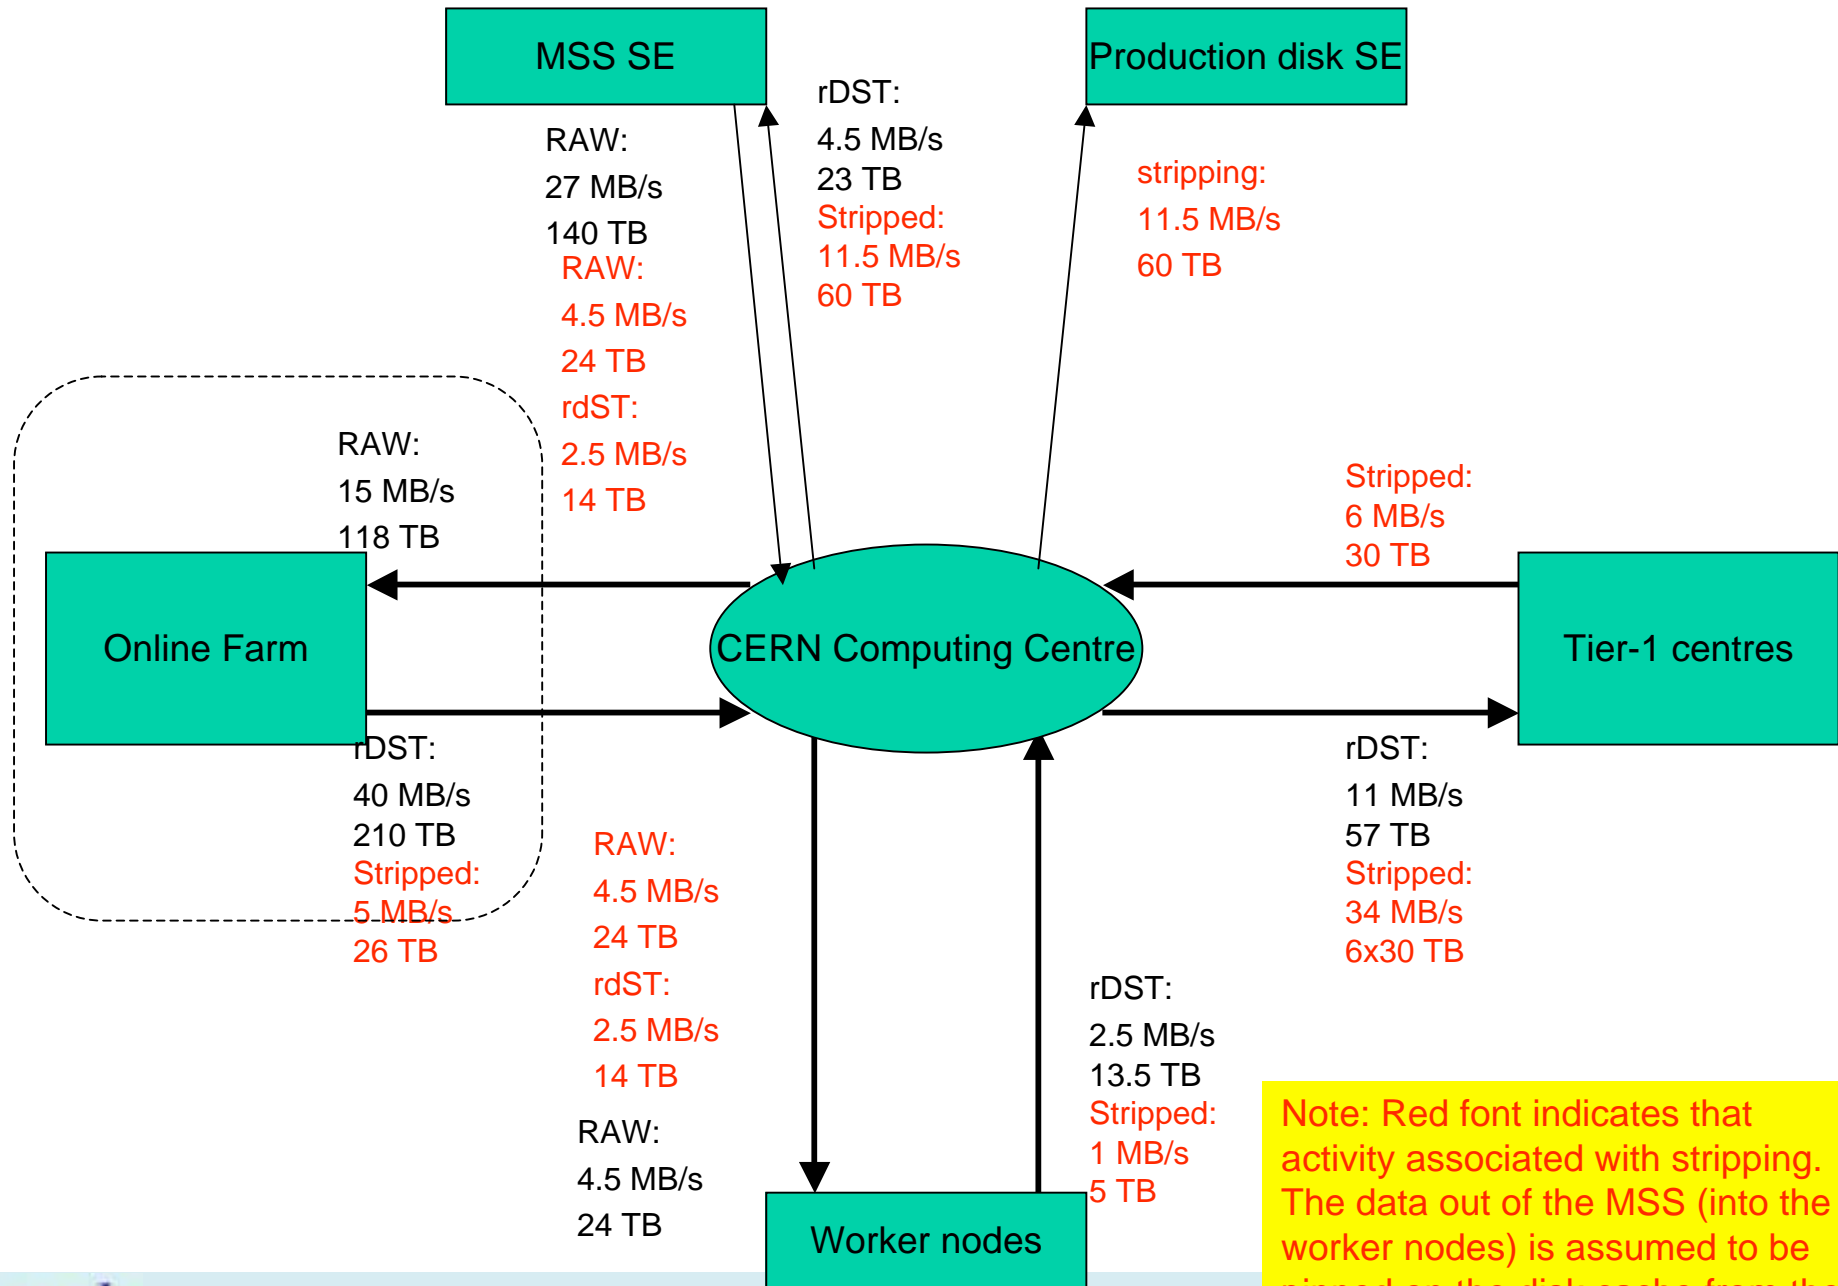


CERN Data taking





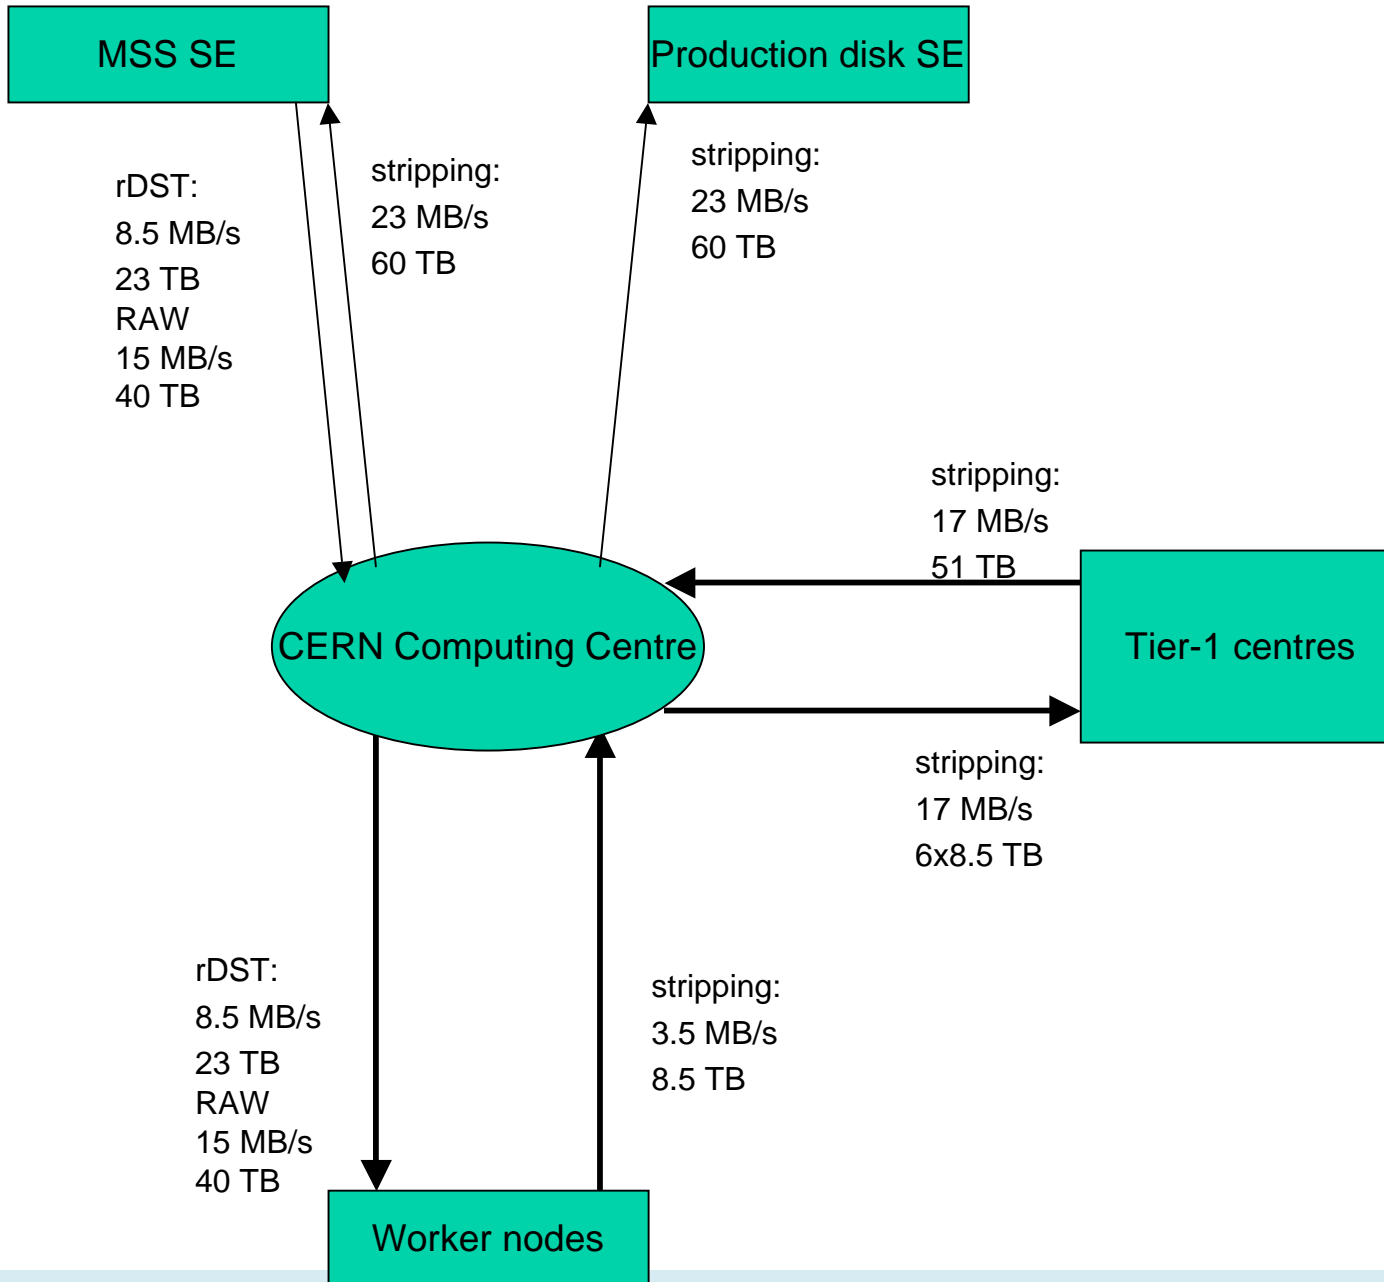
CERN Reprocessing (2 months)



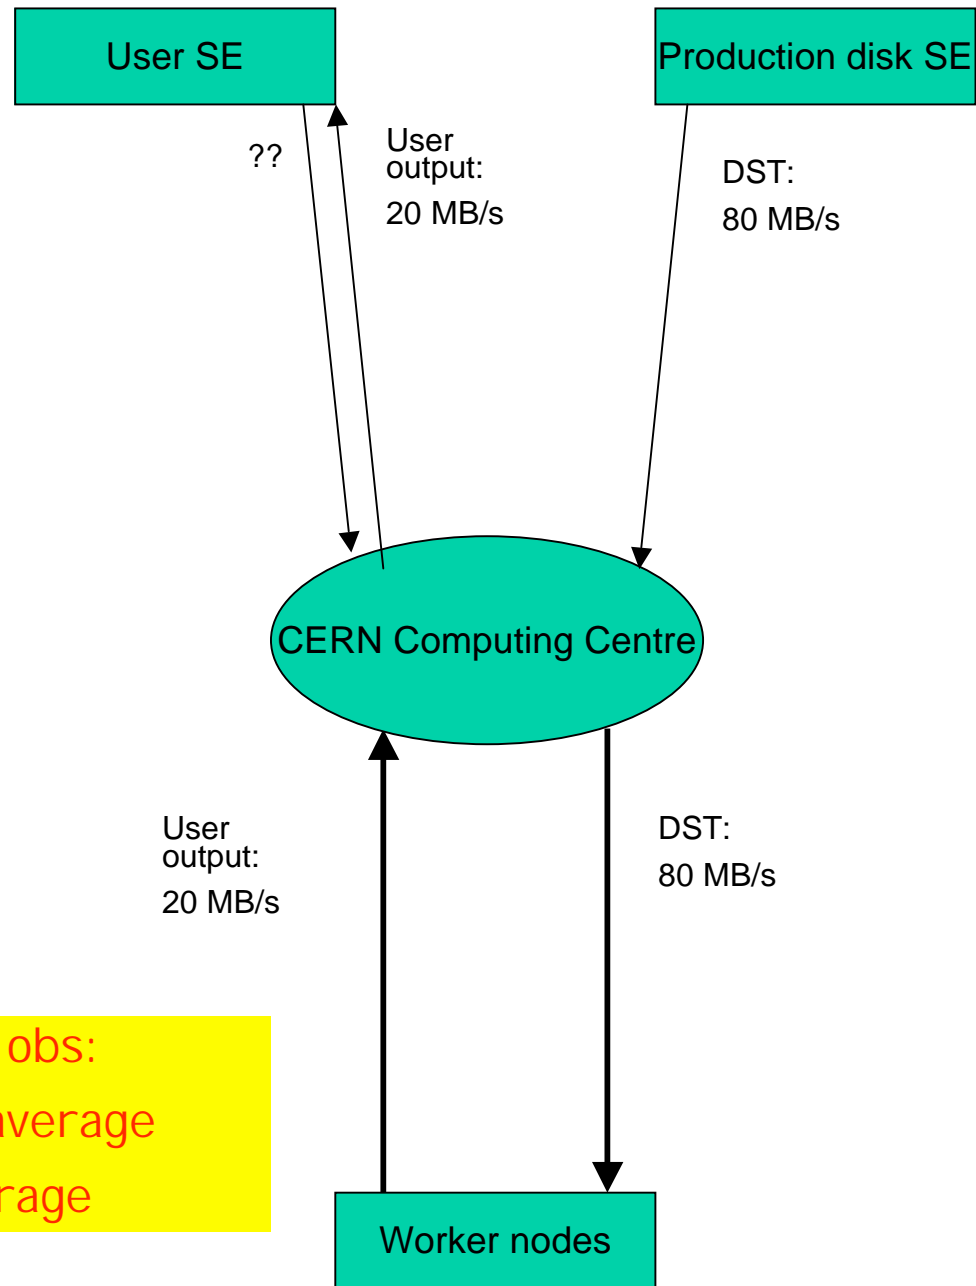
Note: Red font indicates that activity associated with stripping. The data out of the MSS (into the worker nodes) is assumed to be pinned on the disk cache from the reconstruction until stripped

CERN Stripping

(1 month - twice a year)

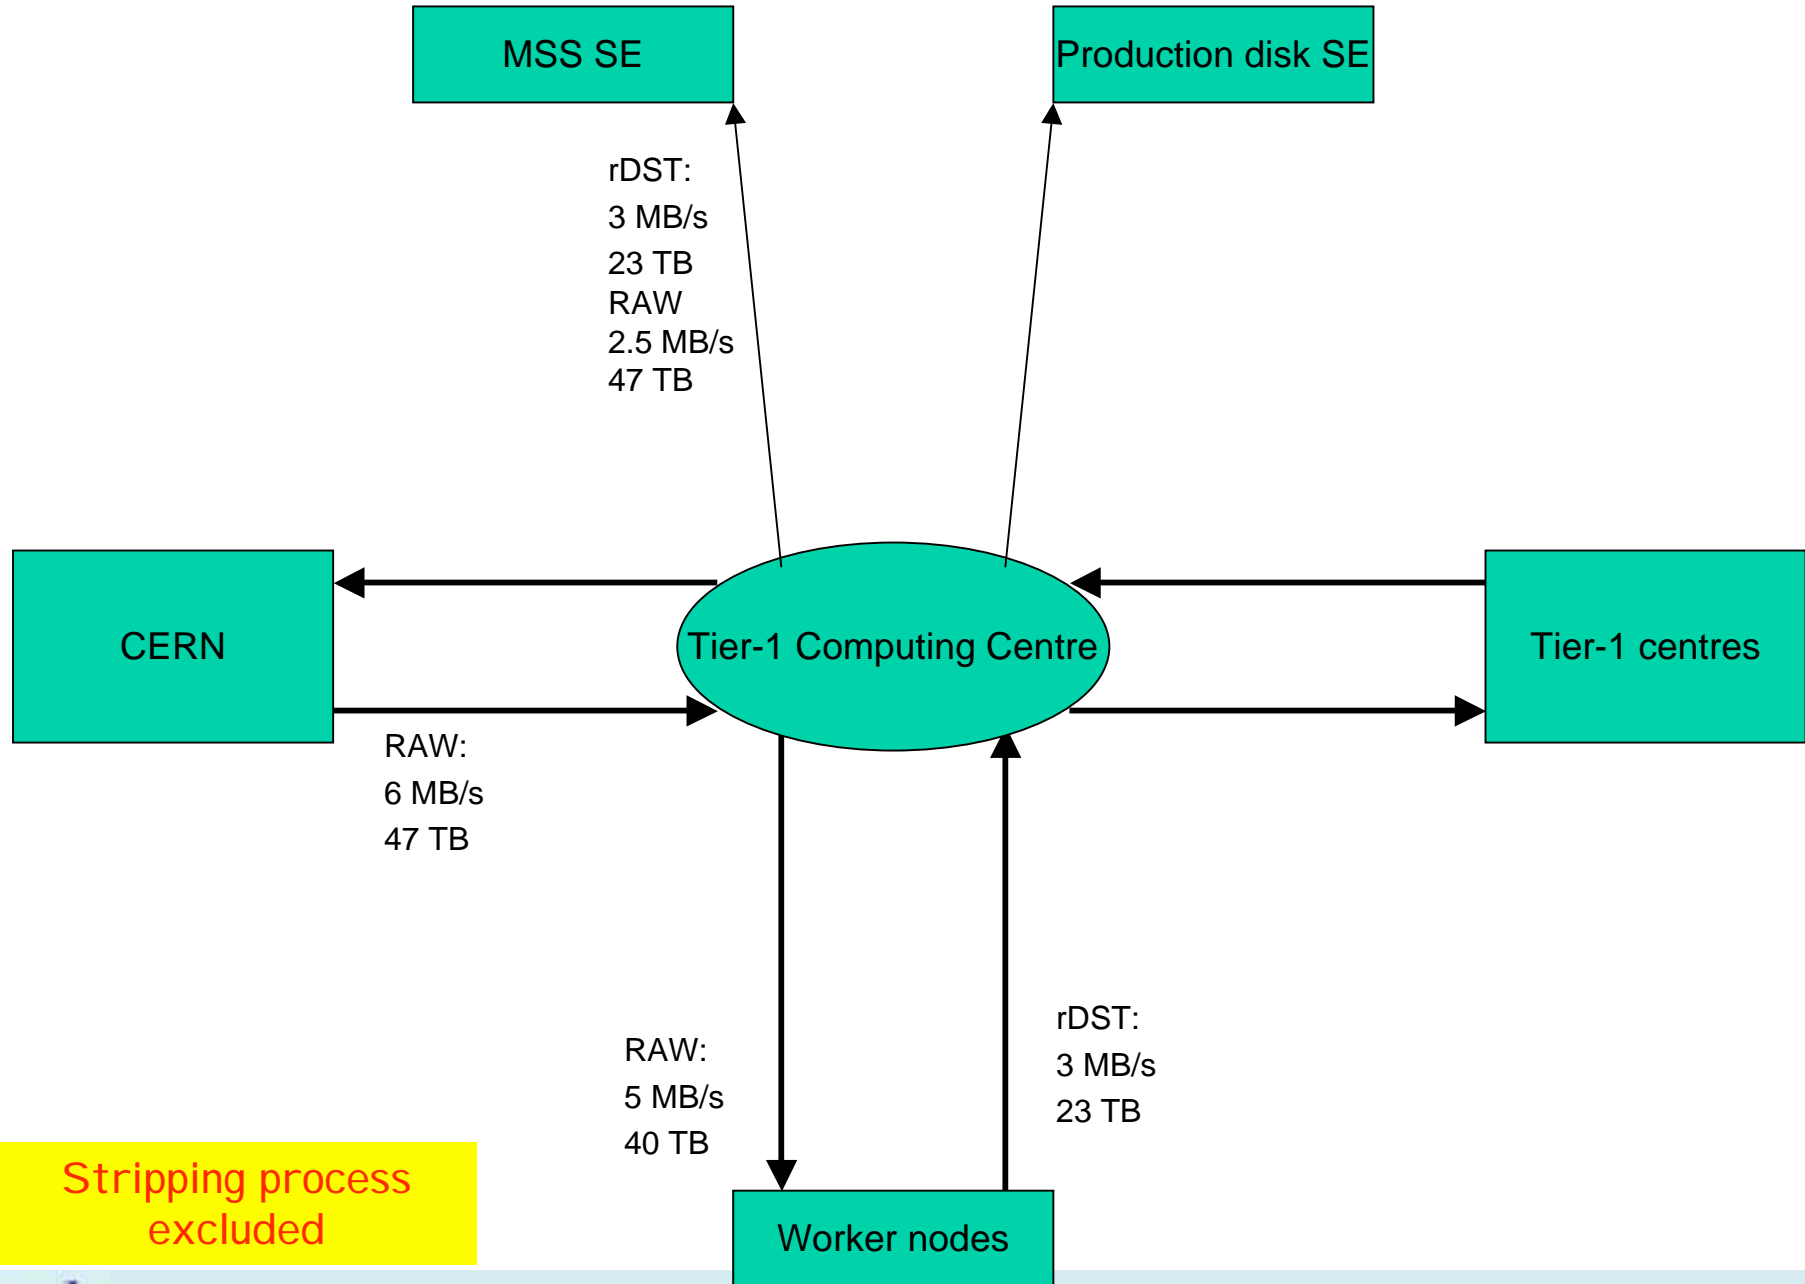


CERN Analysis

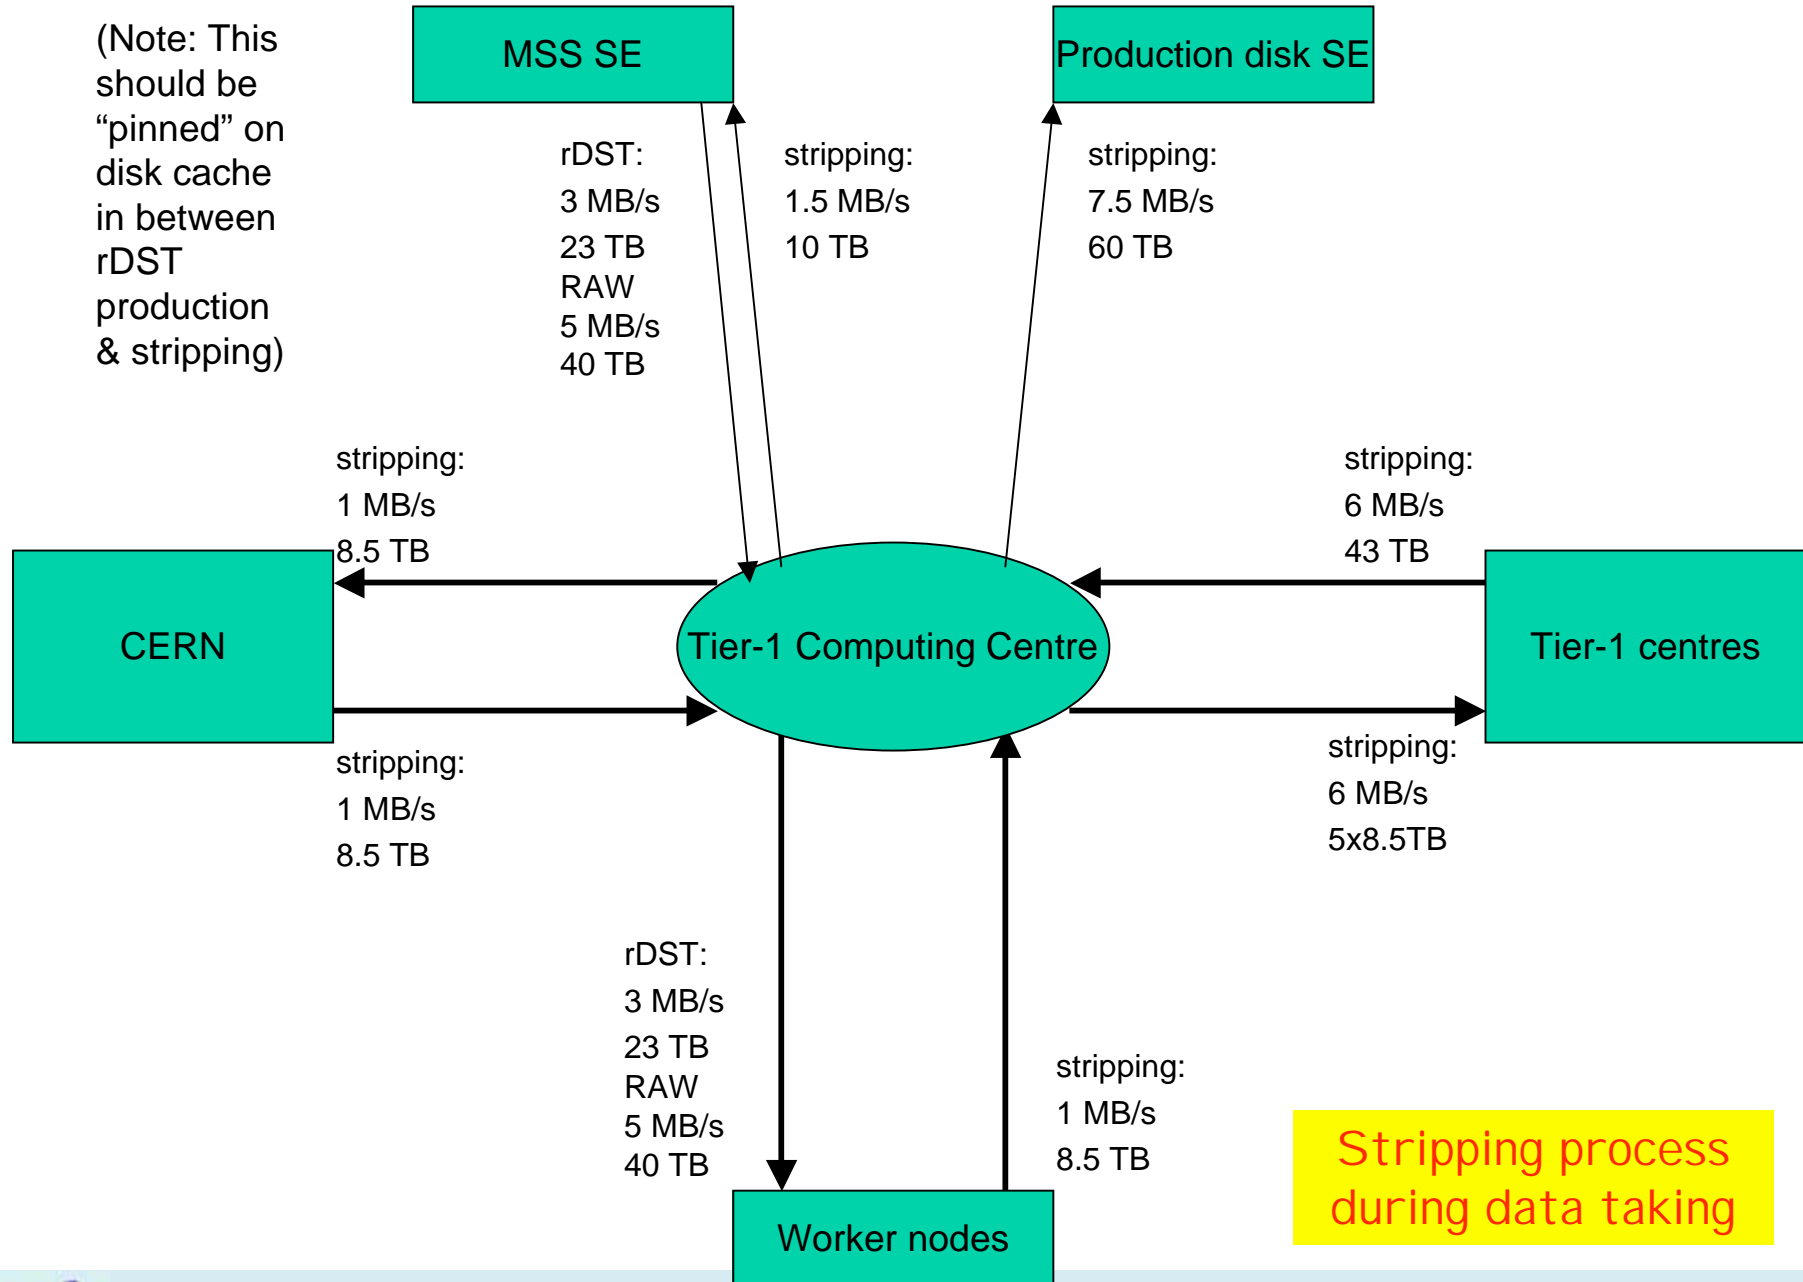


~35 concurrent jobs:
2.3 MB/s input on average
0.6 MB/s on average

“Typical” Tier-1 Data taking

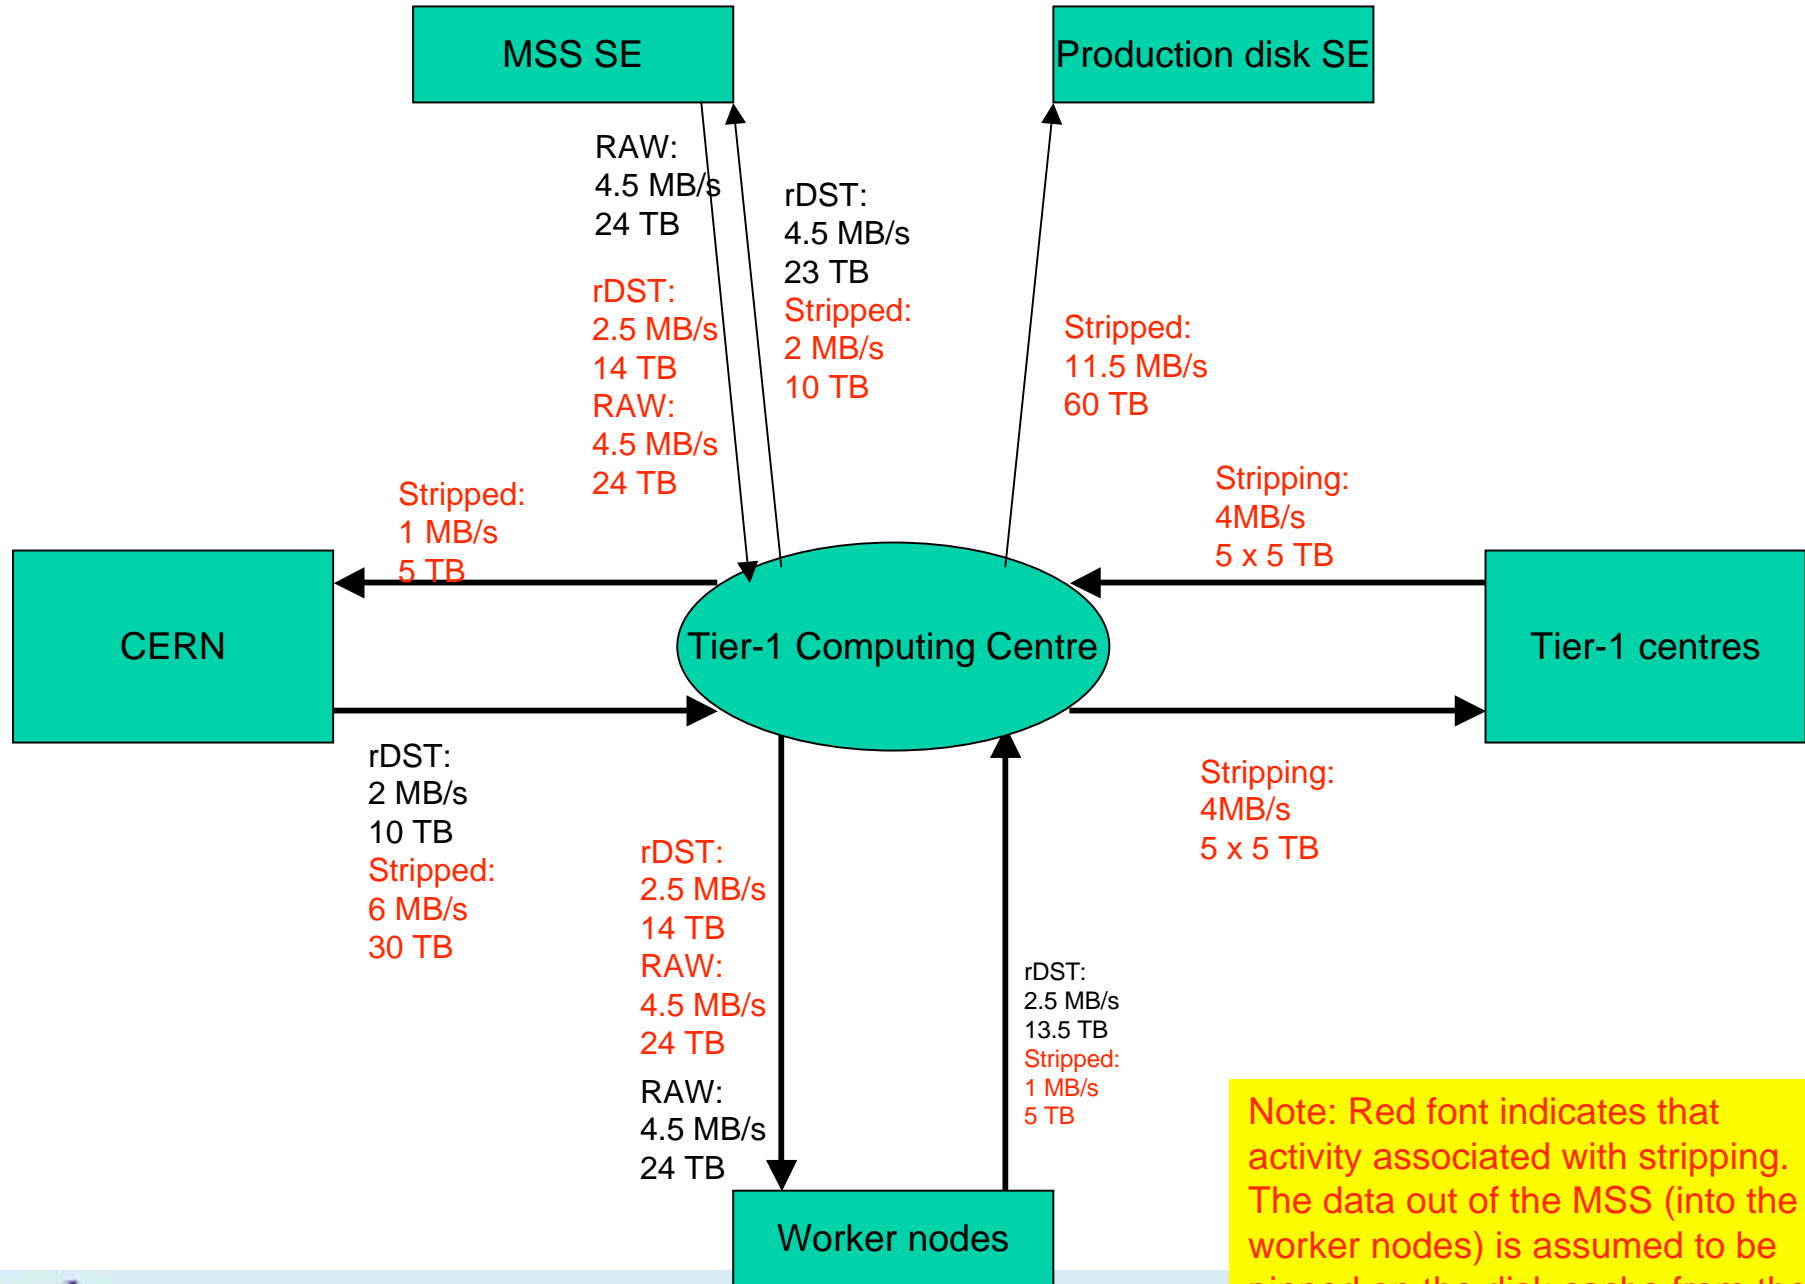


(Note: This should be "pinned" on disk cache in between rDST production & stripping)



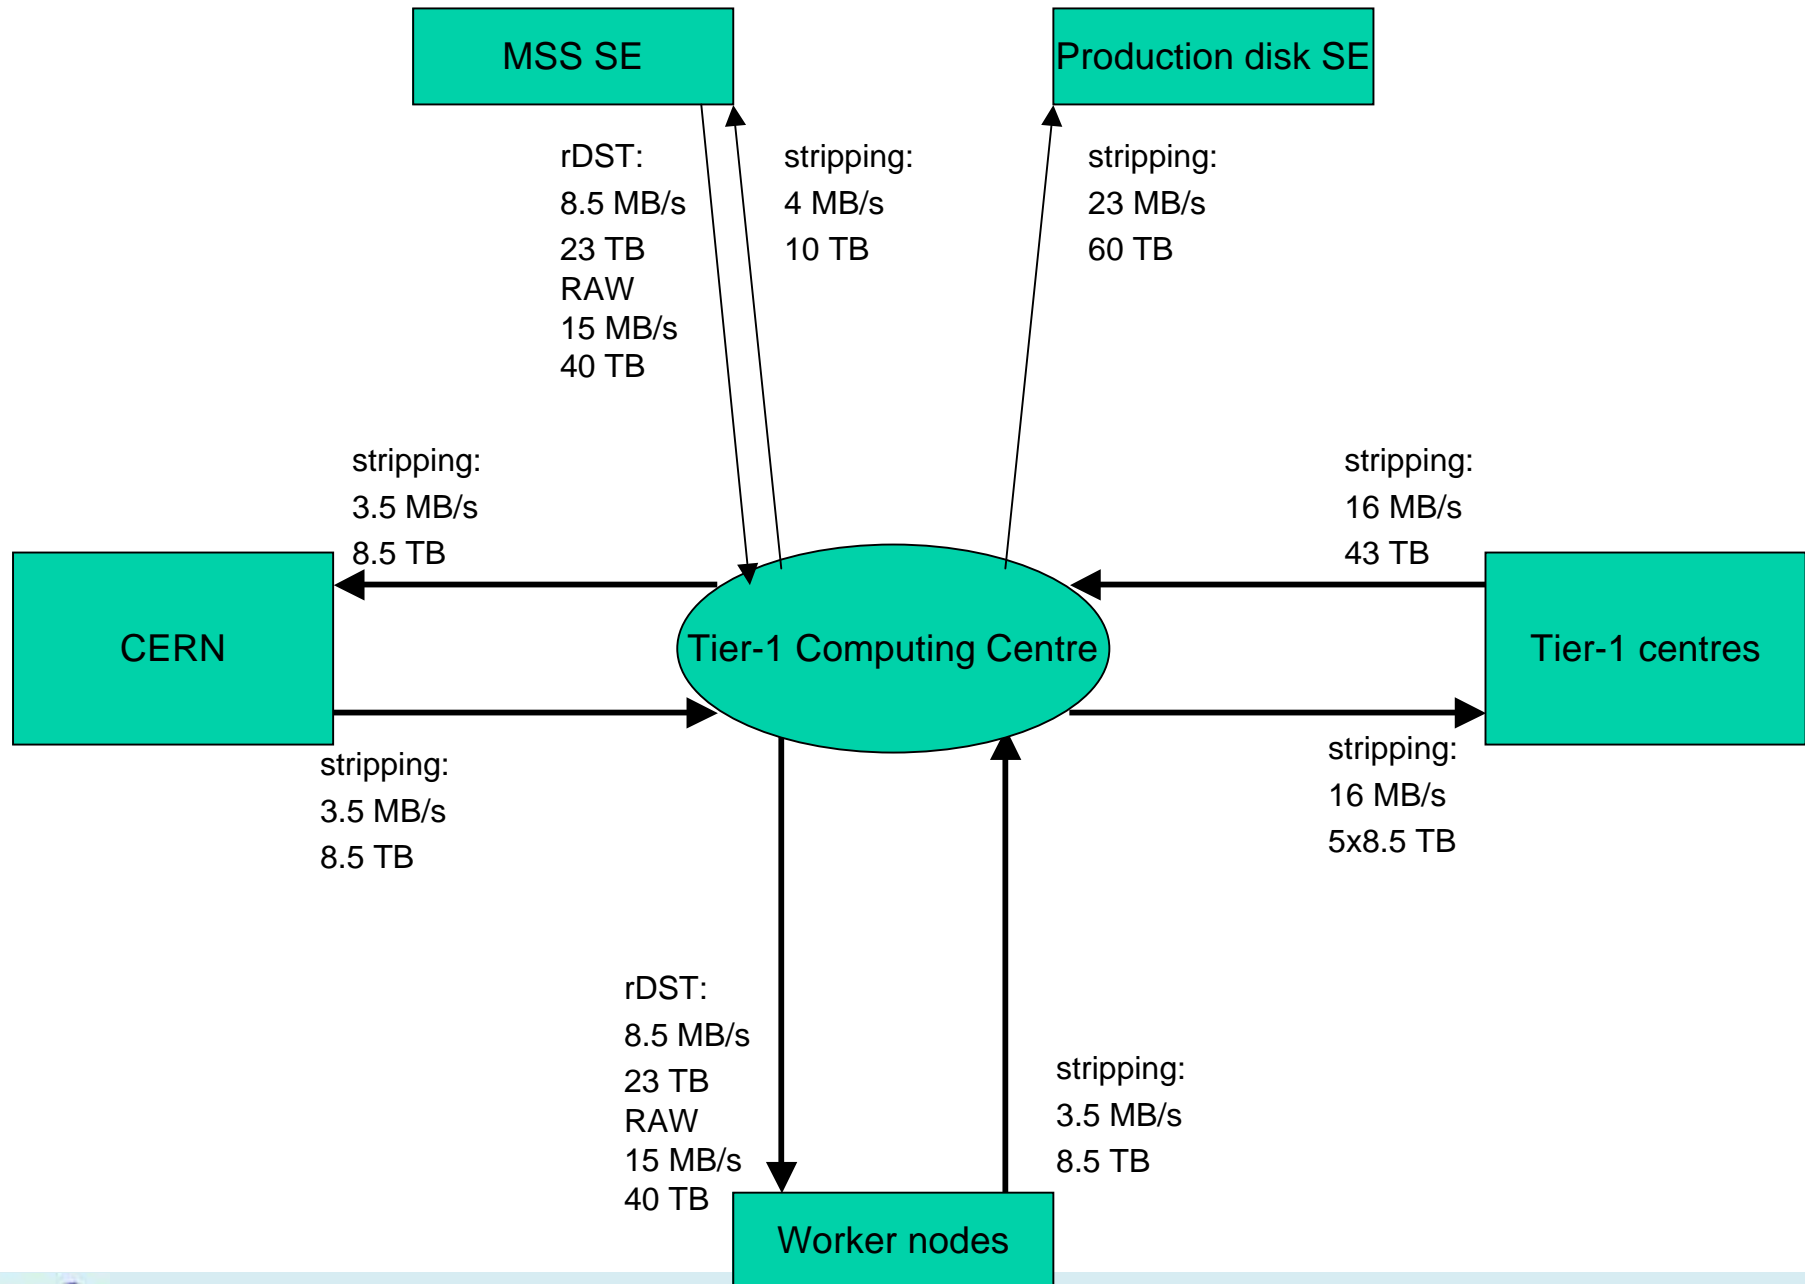
Stripping process during data taking

“Typical” Tier-1 Reprocessing (2-months)

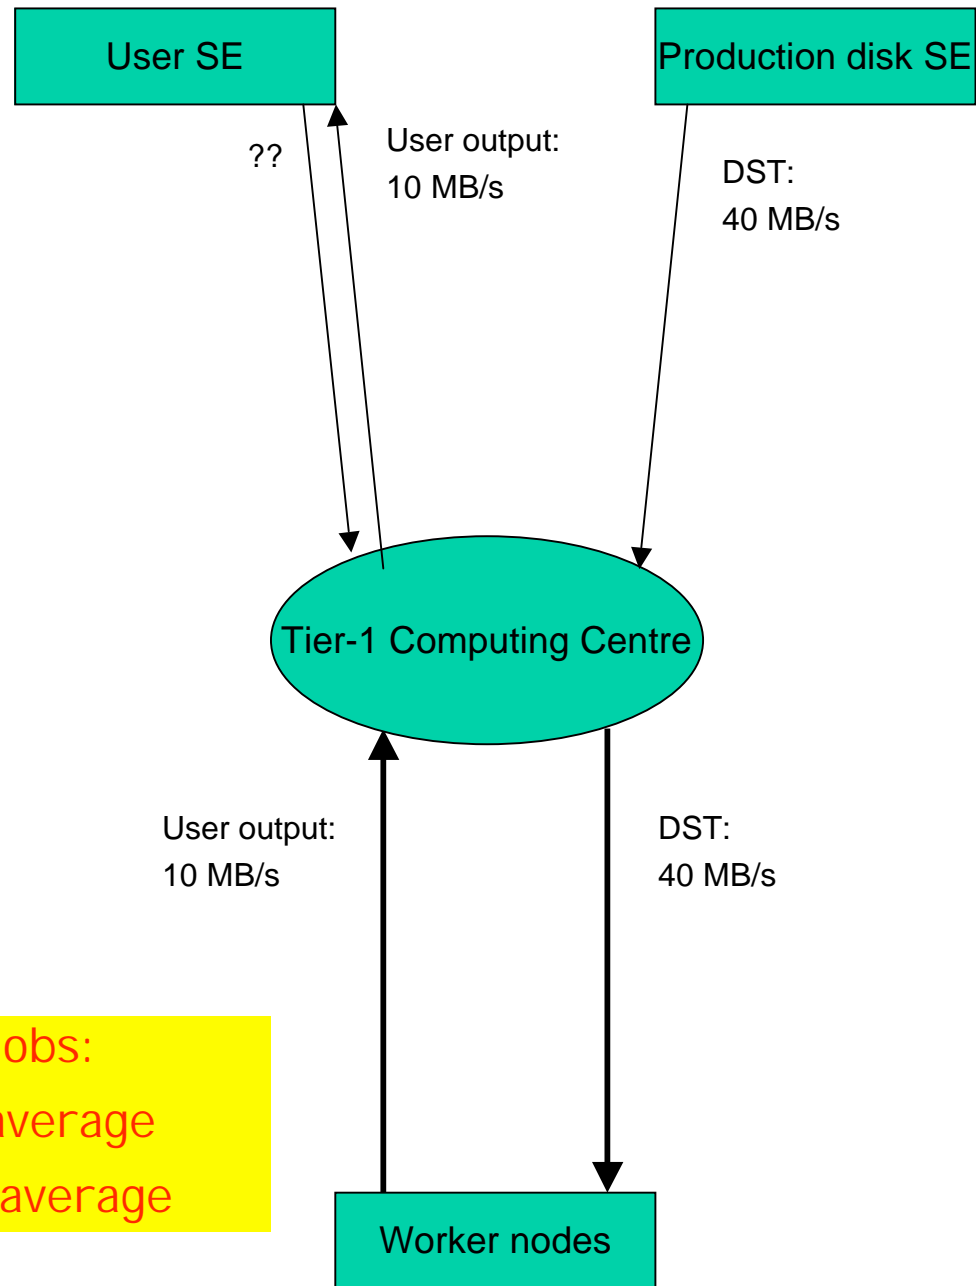


Note: Red font indicates that activity associated with stripping. The data out of the MSS (into the worker nodes) is assumed to be pinned on the disk cache from the reconstruction until stripped

“Typical” Tier-1 Stripping (1 month - twice a year)



“Typical” Tier-1 Analysis



~18 concurrent jobs:
2.3 MB/s input on average
0.6 MB/s output on average