LHC Run 3 tape infrastructure plans

Vladimír Bahyl
IT Department
March 2021

Agenda

Infrastructure overview

LHC Run 3 requirements

Tape Infrastructure plans

Conclusion

Past 5 years ...

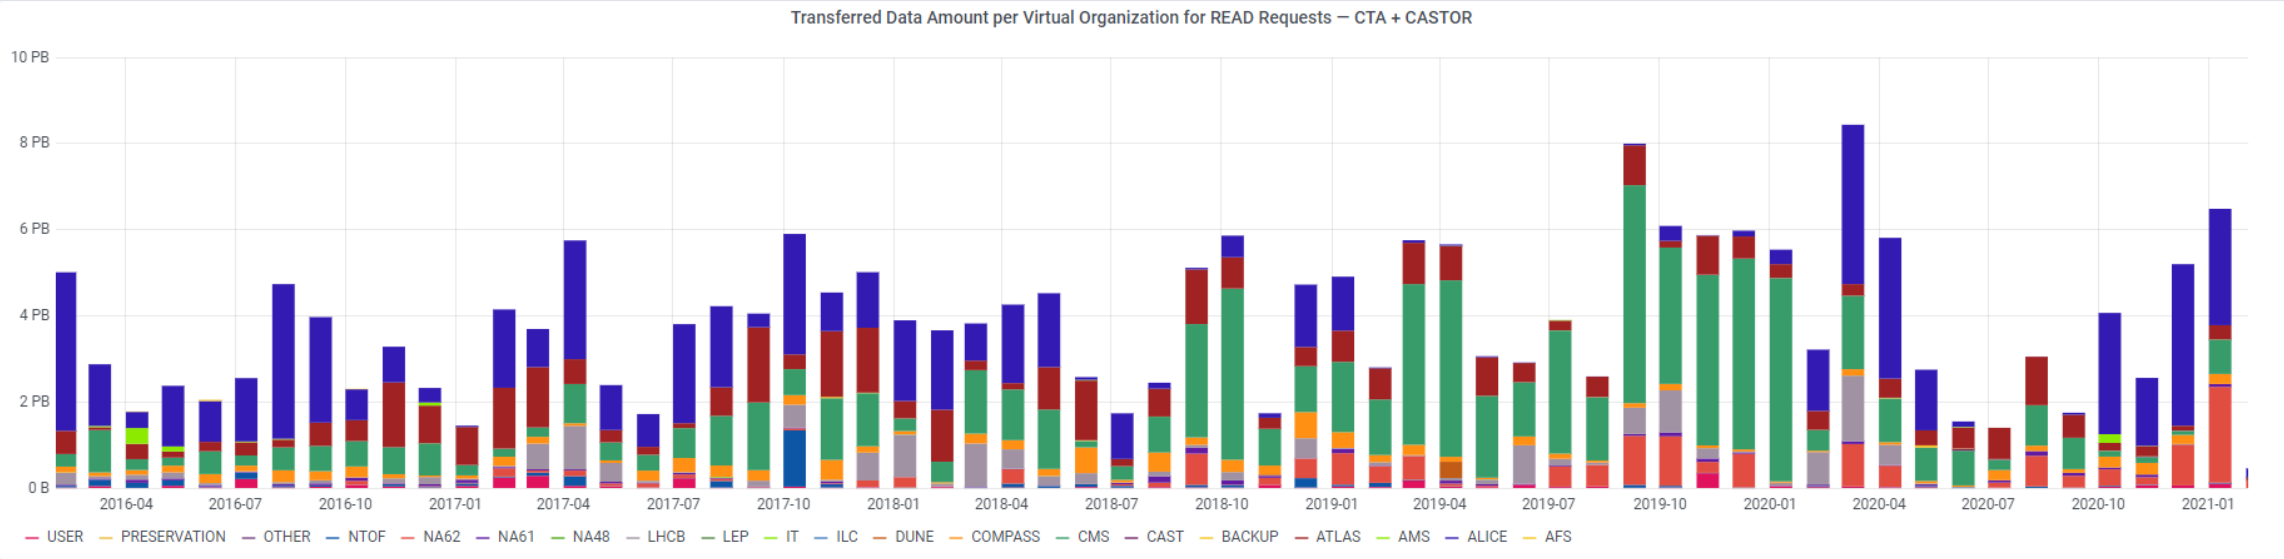
Transferred Data Amount per Virtual Organization for READ Requests — CTA + CASTOR
10 PB
8 PB
6 PB
4 PB
2 PB
0 B
2016-04
2016-07
2016-10
2017-01
2017-04
2017-07
2017-10
2018-01
2018-04
2018-07
2018-10
2019-01
2019-04
2019-07
2019-10
2020-01
2020-04
2020-07
2020-10
2021-01
USER
PRESERVATION
OTHER
NTOF
NA62
NA61
NA48
LHCB
LEP
IT
ILC
DUNE
COMPASS
CMS
CAST
BACKUP
ATLAS
AMS
ALICE
AFS

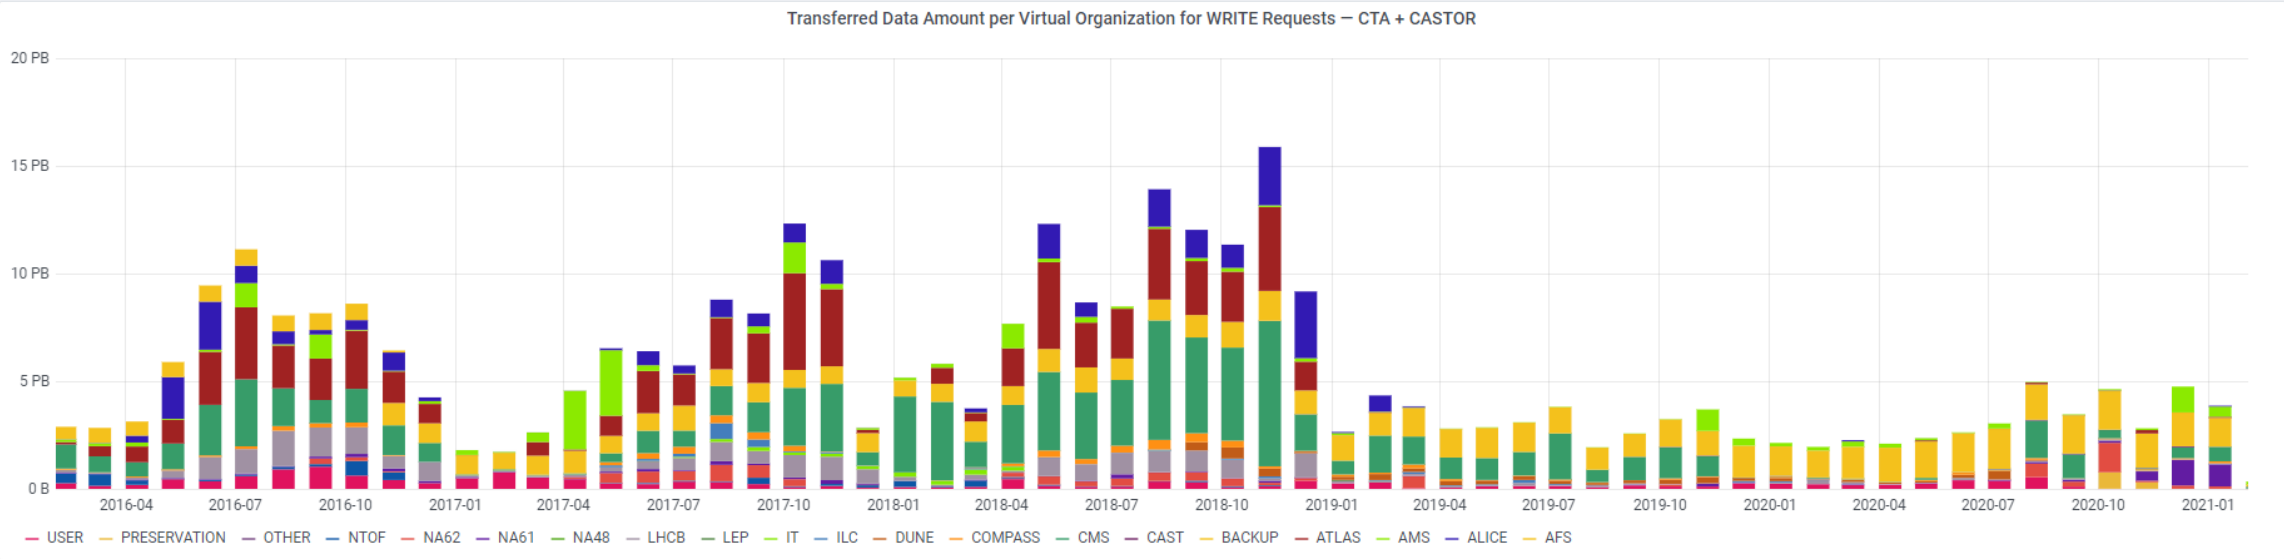
Transferred Data Amount per Virtual Organization for WRITE Requests — CTA + CASTOR
20 PB
15 PB
10 PB
5 PB
0 B
2016-04
2016-07
2016-10
2017-01
2017-04
2017-07
2017-10
2018-01
2018-04
2018-07
2018-10
2019-01
2019-04
2019-07
2019-10
2020-01
2020-04
2020-07
2020-10
2021-01
USER
PRESERVATION
OTHER
NTOF
NA62
NA61
NA48
LHCB
LEP
IT
ILC
DUNE
COMPASS
CMS
CAST
BACKUP
ATLAS
AMS
ALICE
AFS

Data volume vs. Tape Technology

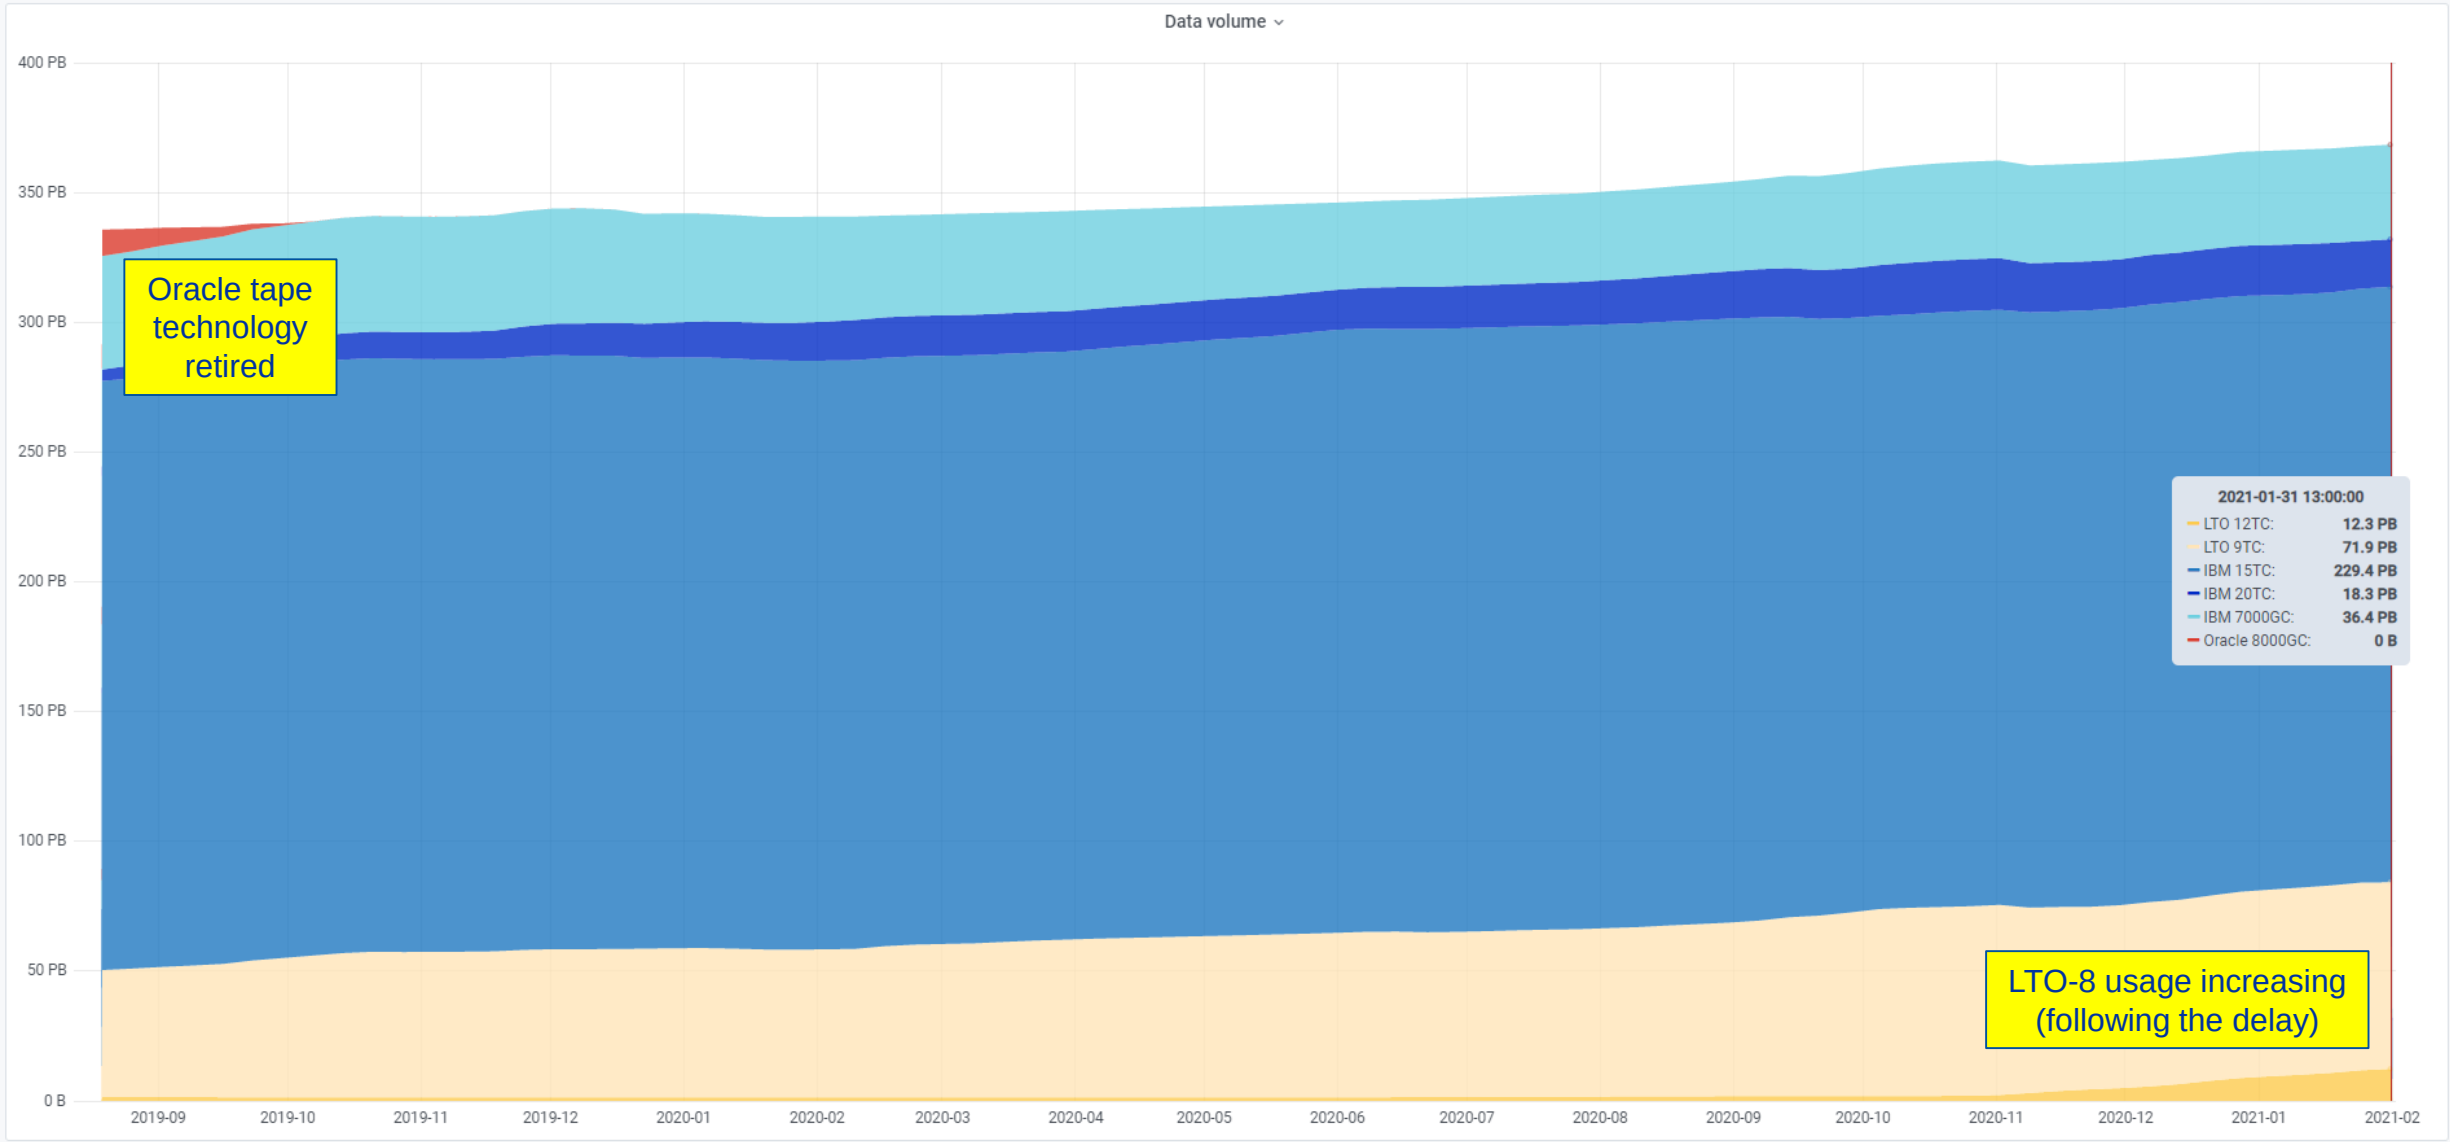

Tape Infrastructure (February 2021)

- CASTOR/CTA archive:
  - IBM : 3 x TS4500
    - 10 x TS1160, 46 x TS1155, 12 x TS1150
      20 x LTO-8
    - 1000 x 3592JE media (20 TB)
      15028 x 3592JD media (15 TB)
      6009 x 3592JC media (7 TB)
      960 x LTO-8 media (12 TB)
      7110 x LTO-7M media (9 TB)
  - Spectra Logic : 1 x TFinity
    - 10 x LTO-8
    - 1450 x LTO-8 media (12 TB)
      689 x LTO-7M media (9 TB)
  - ~370 PB; ~646 M files
  - ~20 PB of free space
- IBM Spectrum Protect backup:
  - IBM : 2 x TS3500
    - 55 x TS1140
    - 199 x 3592JC media (7 TB)
      12926 x 3592JB media (1.6 TB)
  - 12.5 PB; ~2700 M files
  - 19 x servers running version 8.1.10

Spectra Logic TFinity for LTO

Tape Infrastructure topology

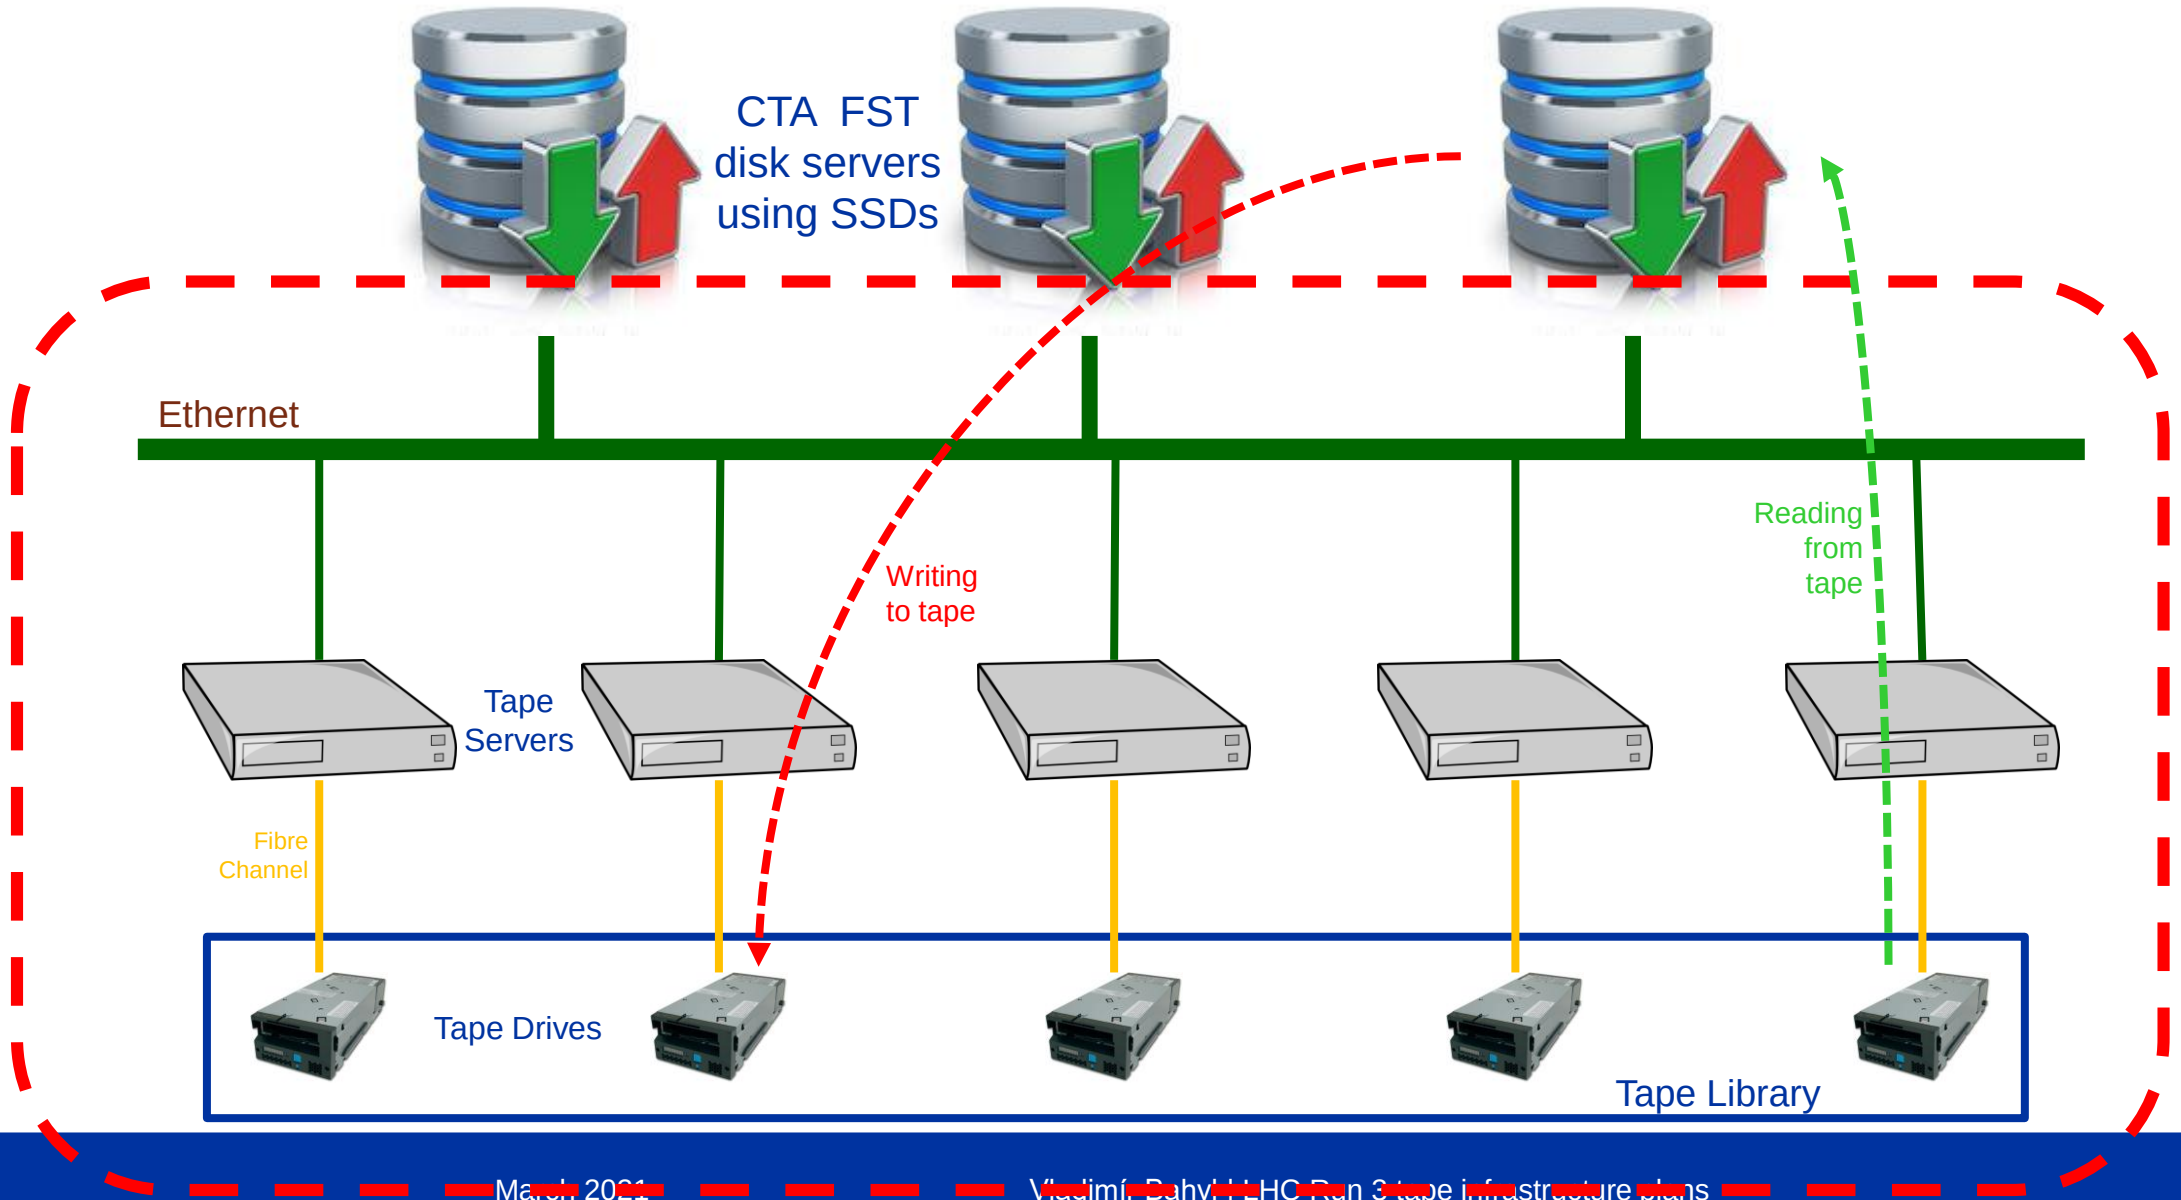

LHC Run 3 – boundary conditions

Physics requirements:

- Data amount per year: 150 PB LHC + 30 PB non-LHC = 180 PB
- Total data amount for 3 years: 540 PB
- Transfer rate requested: 10 GB/s for *each* of the 4 LHC experiments

Infrastructure limits:

- 2 LTO libraries, 2 IBM 3592 enterprise libraries
- 24900 free LTO slots; 9300 free IBM 3592 enterprise slots
- 48 tape drive slots per each library

LHC Run 3 – tape infrastructure plans 1/2

Space requirements – Tape Media:

- Fill free 9300 IBM enterprise slots with 3592JE (20TB) media = 186 PB
- Fill free 24900 LTO slots with LTO-9 (18TB) media = 448 PB
- = 634 PB of total available tape capacity will be sufficient for next 3 years

Throughput requirements – Tape Drives:

- IBMLIB3:
  - Total 38 x IBM TS1160 and 10 x IBM TS1155
- IBMLIB4:
  - Total 38 x IBM TS1160 and 10 x IBM TS1155
- Total 162 (new) drives: 76 x IBM TS1160 and 86 x LTO-9

- Expected throughput per tape drive ~300 MB/s
- Total >45 GB/s which should be sufficient for LHC Run 3

- IBMLIB1:
  - Total 38 x LTO-9 and 10 x LTO-8
- SPECTRALIB1:
  - Total 48 x LTO-9

Tape Infrastructure topology – repack use case

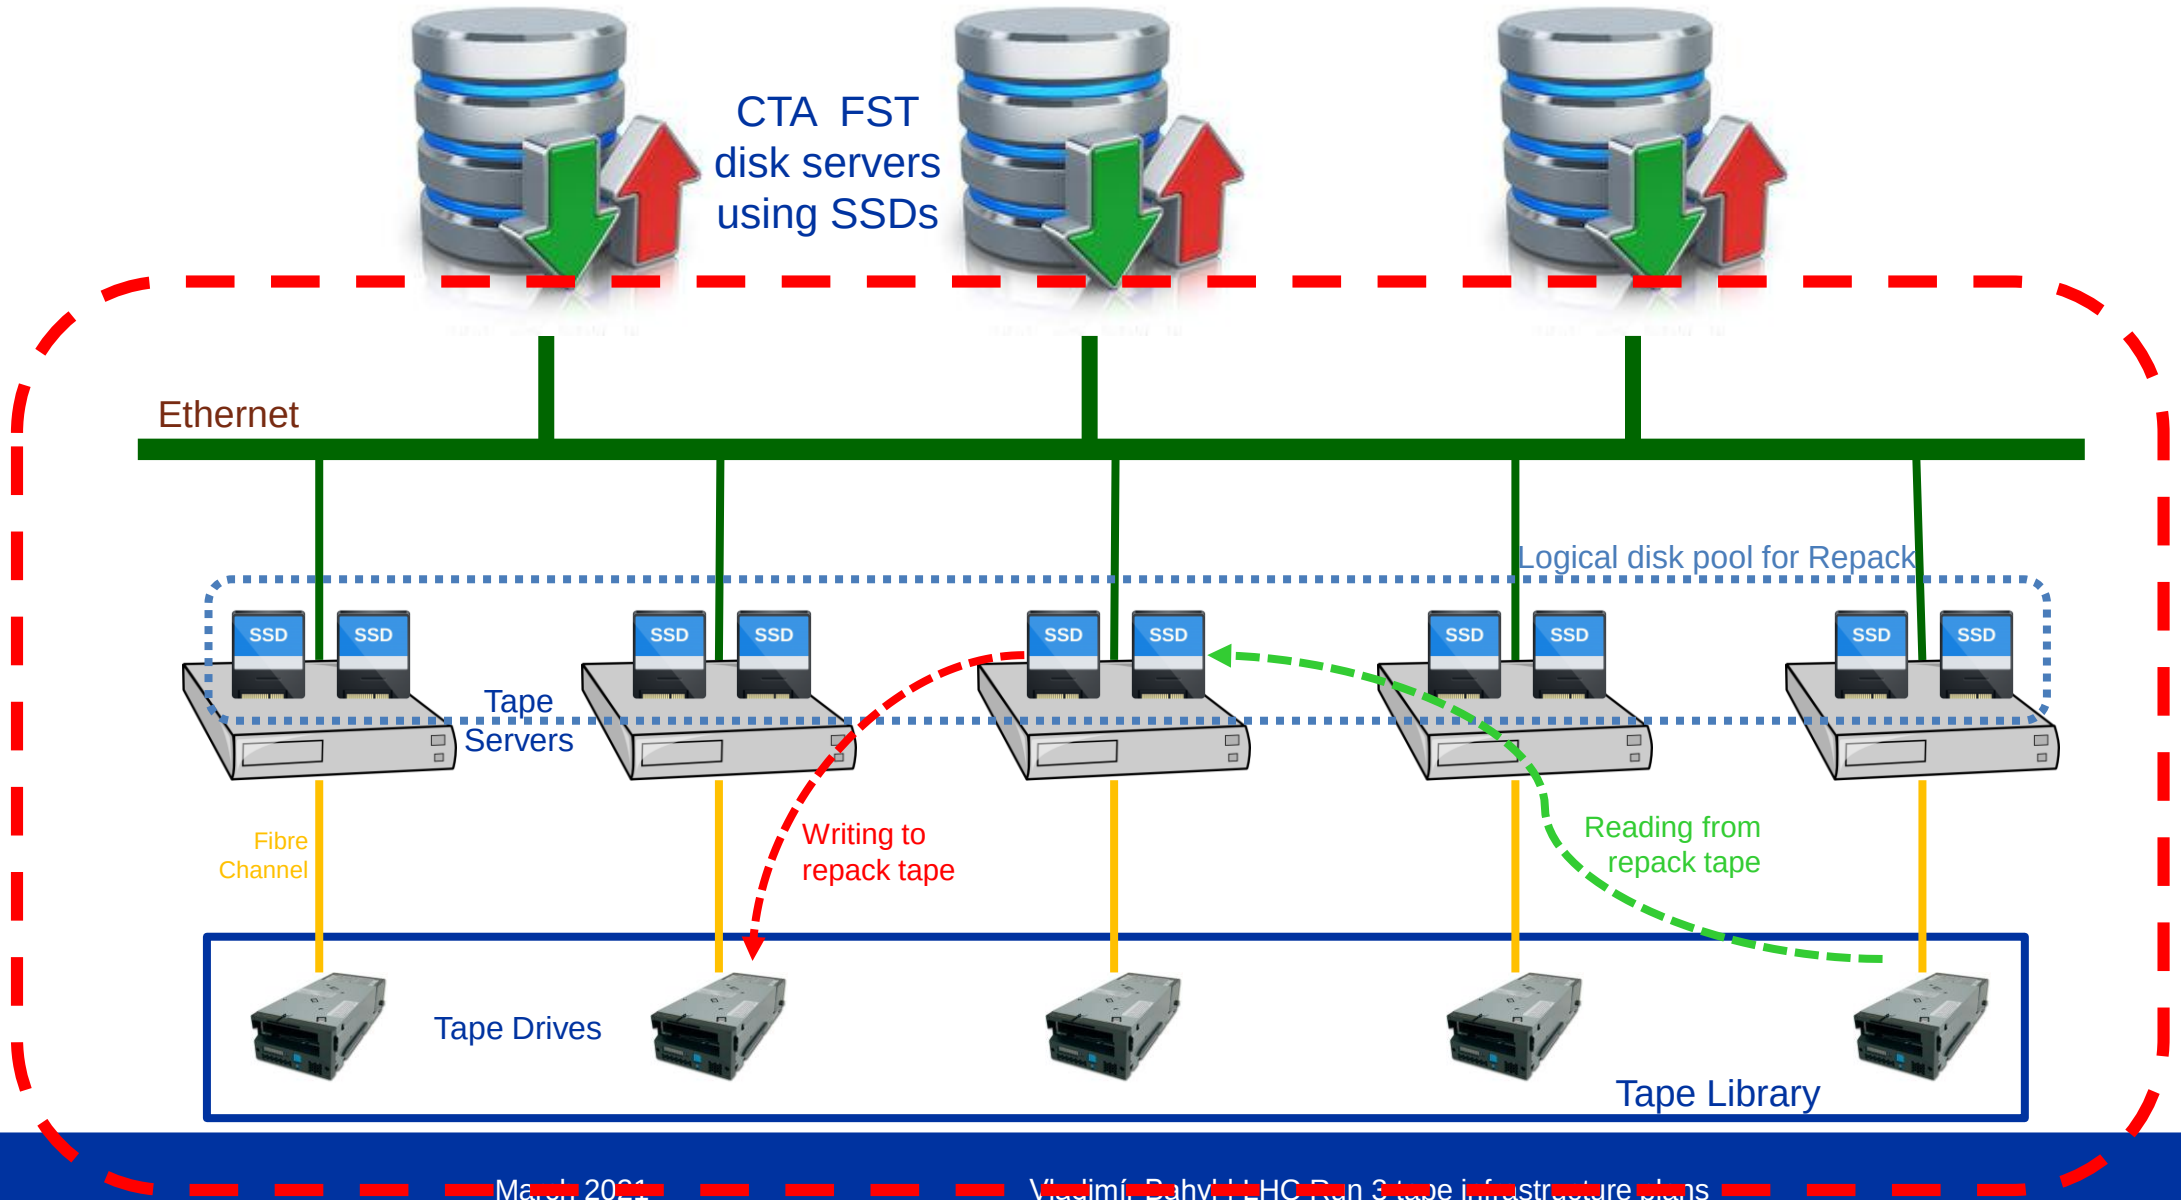

LHC Run 3 – tape infrastructure plans 2/2

Tape Servers:

- 100 new tape servers on order:
  - CPU: AMD 7302 providing 64 SMT cores in total
  - Memory: 256 GB
  - Disk: 1 x 800 GB NVMe SSD for operating system, 2 x 1.9 TB NVMe SSDs for storage
  - Network: 2 x 25 Gbps
  - Fibre Channel HBA: 2 x 16 Gbps
- Around 100 other older tape servers:
  - Memory: 128 GB or 64 GB
  - Disk: 4 x 1 TB for operating system and storage
  - Network: 2 x 10 Gbps
  - Fibre Channel HBA: 2 x 16 Gbps or 2 x 8 Gbps

Network:

- Switches: Juniper QFX5120-48Y, with 25G-SR optics
  - 48 x 25Gbps server ports and 8 x 100Gbps uplink ports
  - Usually deploy 16 servers per switch (due to rack space constraints) but this can be increased

Conclusion

We need to be conservative and base our plans on existing or already announced products.

CERN tape infrastructure is prepared to accommodate all LHC Run 3 data requirements.

- We have enough cartridges slots for IBM 3592JE and LTO-9 media.
- We have enough tape drive slots to sustain the new data throughput but also be backwards compatible.

We have contracts and finances in place to purchase the required tape equipment once LTO-9 is available.

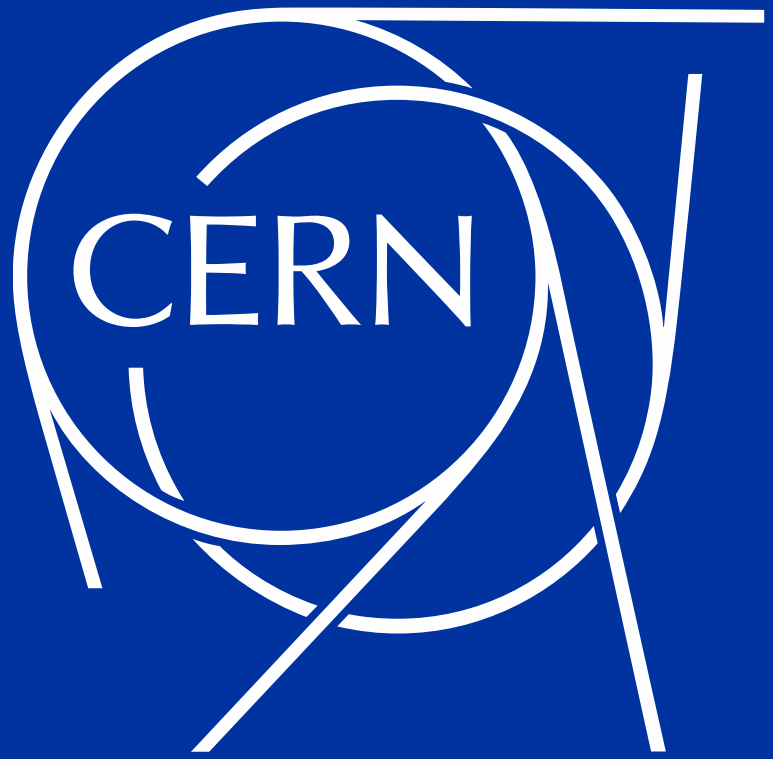
CERN