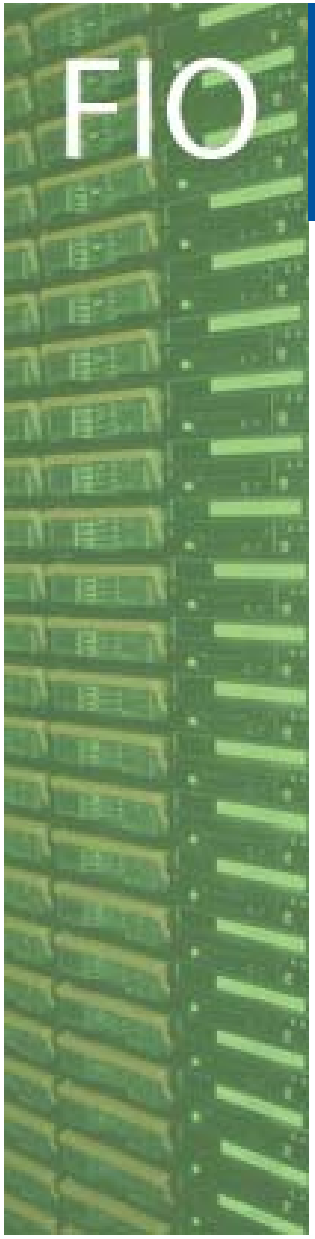

Fabric Infrastructure and Operations

CERN IT Department

CCRC Review of Tape at CERN

Tim Bell

February 2008

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Agenda

- Tape usage during CCRC
  - Volumes
  - Efficiency
- Review of CCRC from a tape perspective

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CCRC Data Written

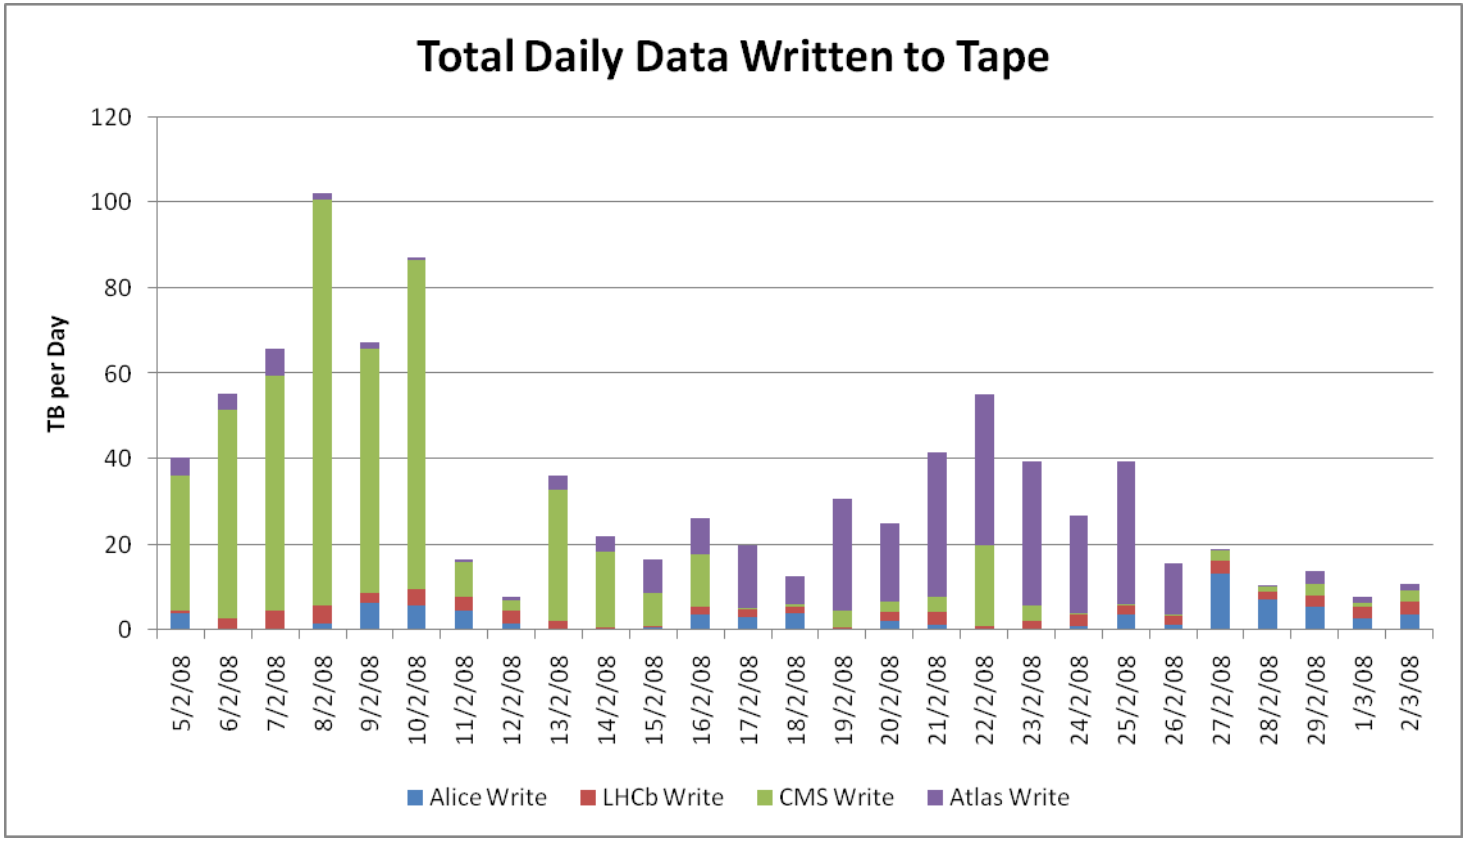

- CMS high volume during 1st two weeks and then Atlas

CCRC Data Read

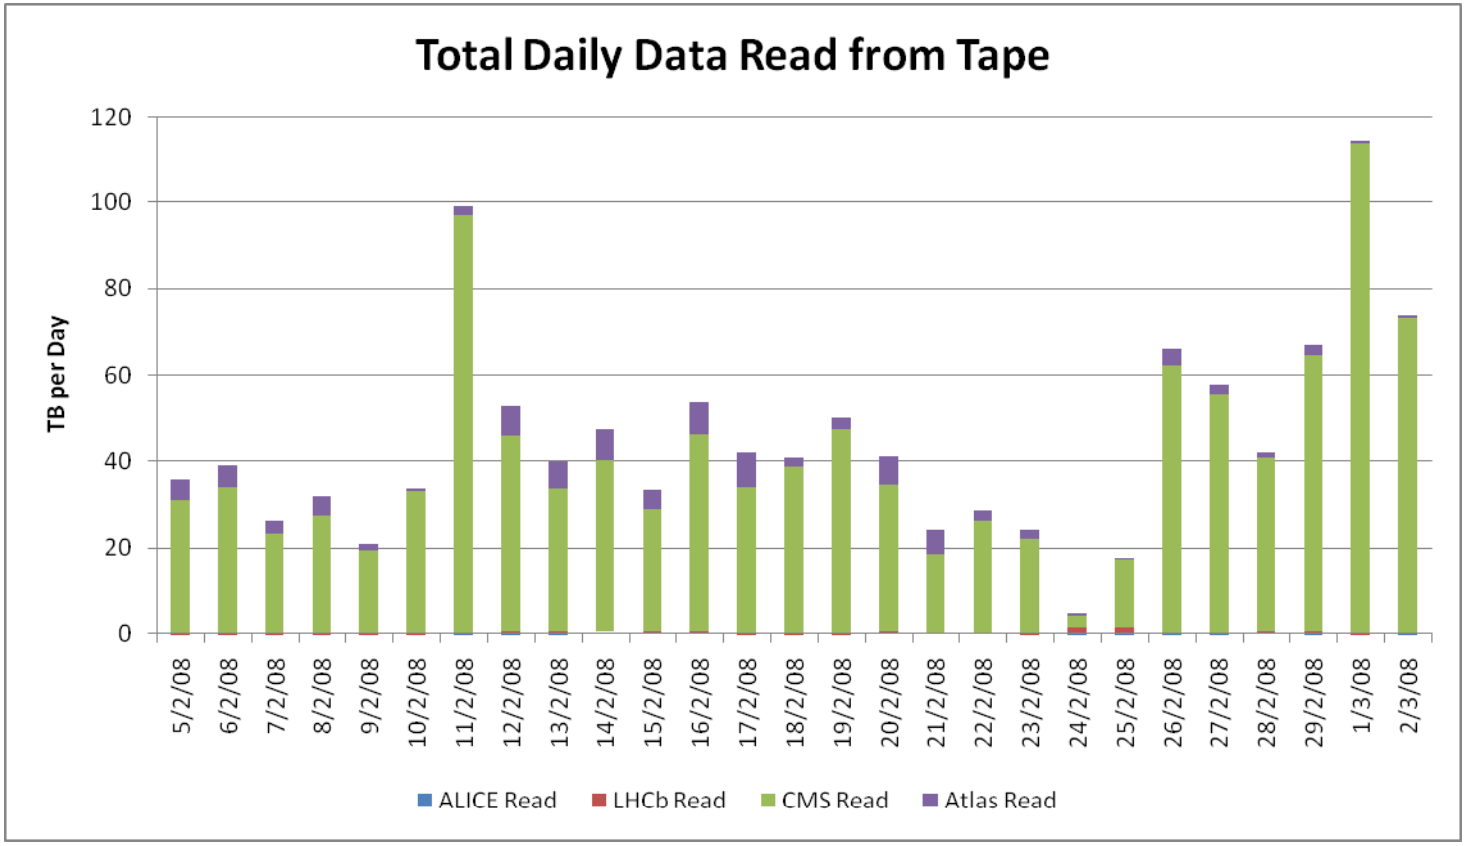

- CMS high volume throughout

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

File size and performance

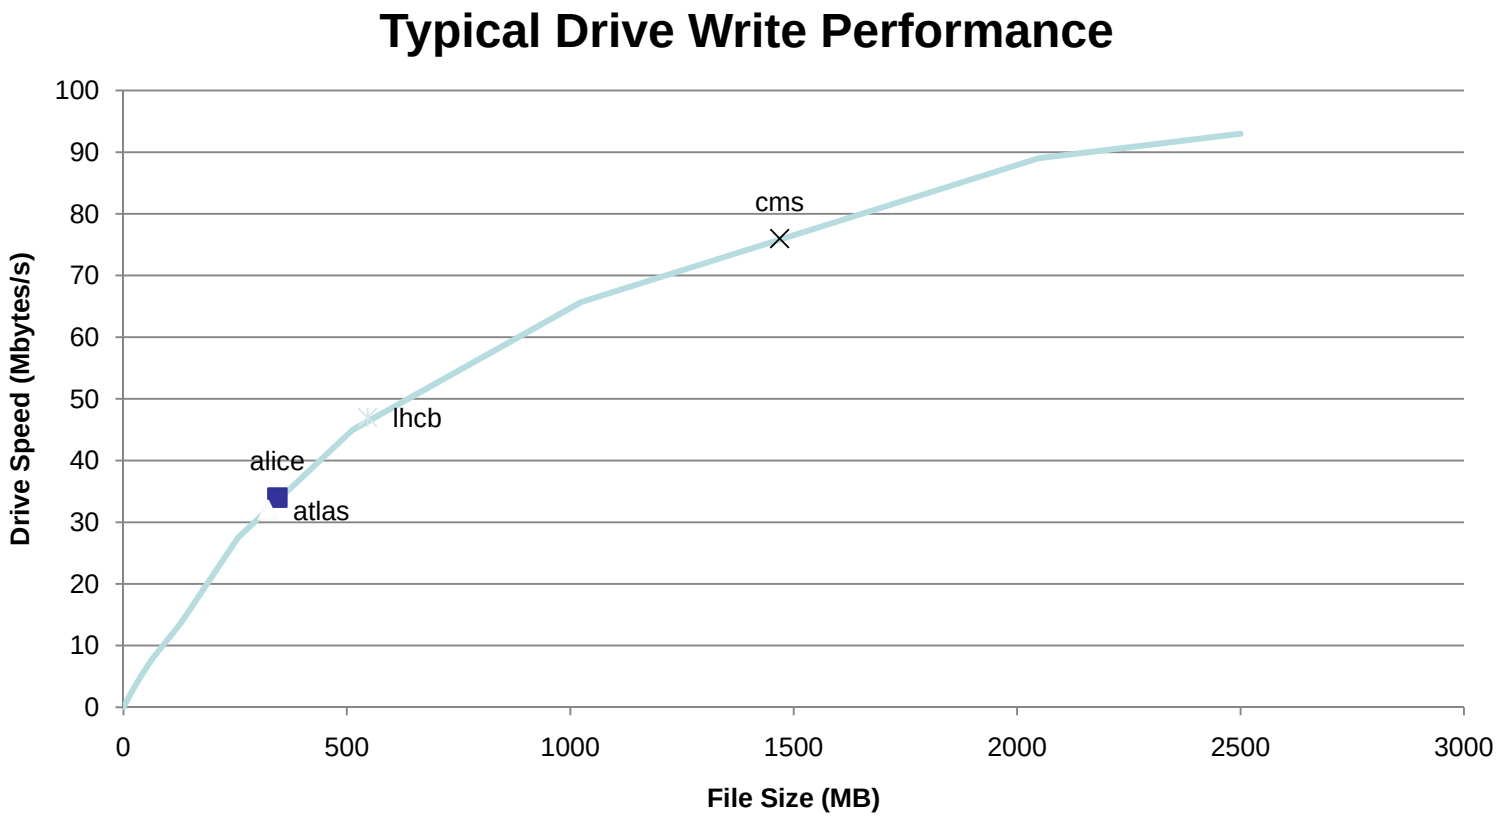

Data	Alice	Atlas	CMS	LHCb
CCRC Feb '08	340 MB	320 MB	1470 MB	550 MB
Jan '08	200 MB	250 MB	2000 MB	200 MB

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Total performance to tape

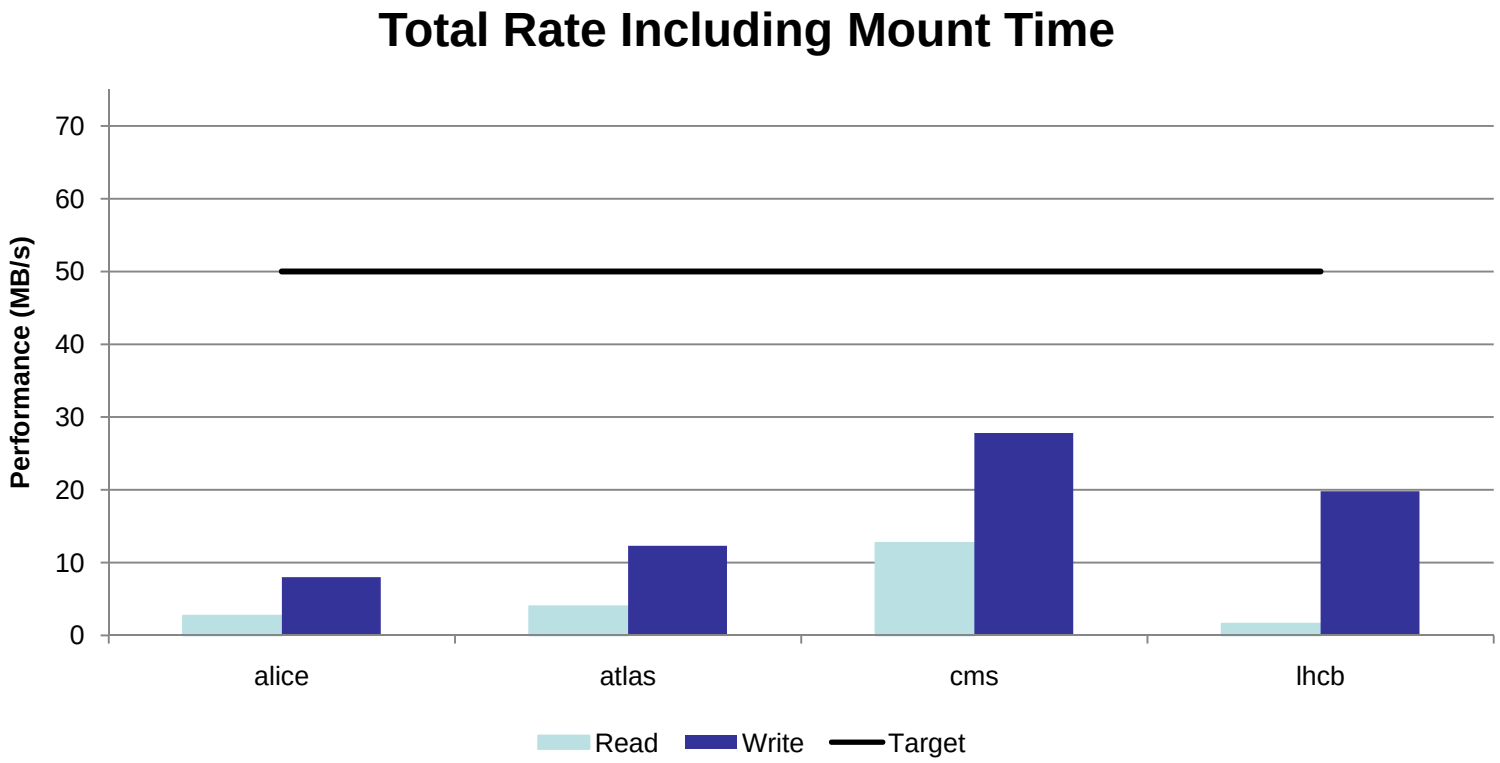

- Alice and LHCb running with Castor 2.1.4 so expected write rates when migrated are around double
- Atlas write rates were up to 30 MB/s during Week 8 as file size increased
- CMS write rates have doubled since January
- Read remains inefficient for all VOs

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Data per Mount

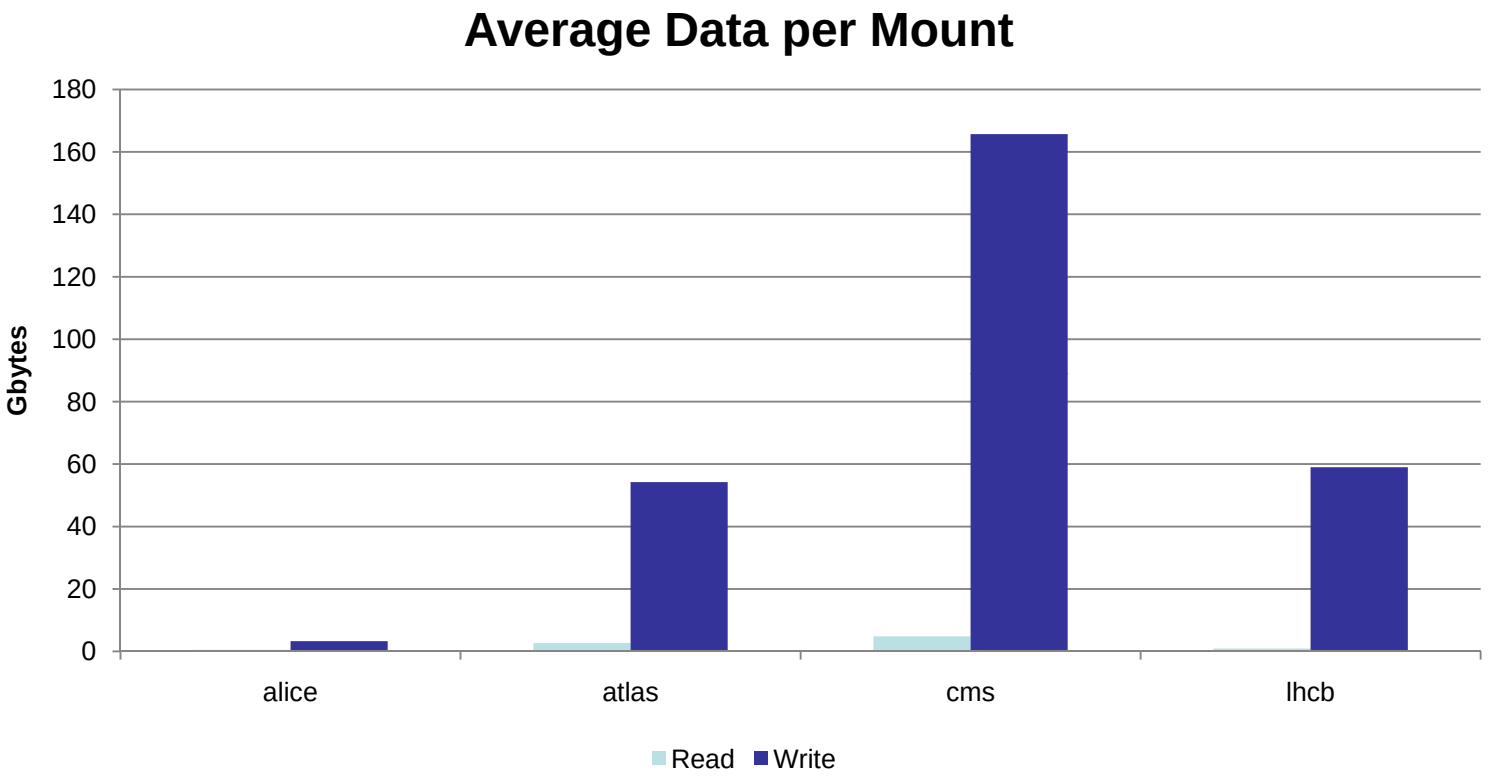

- Write data volumes improving with new policies
- Read data volumes remains a concern
  - Disk cache size versus Garbage collection policy

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tape usage read dominated

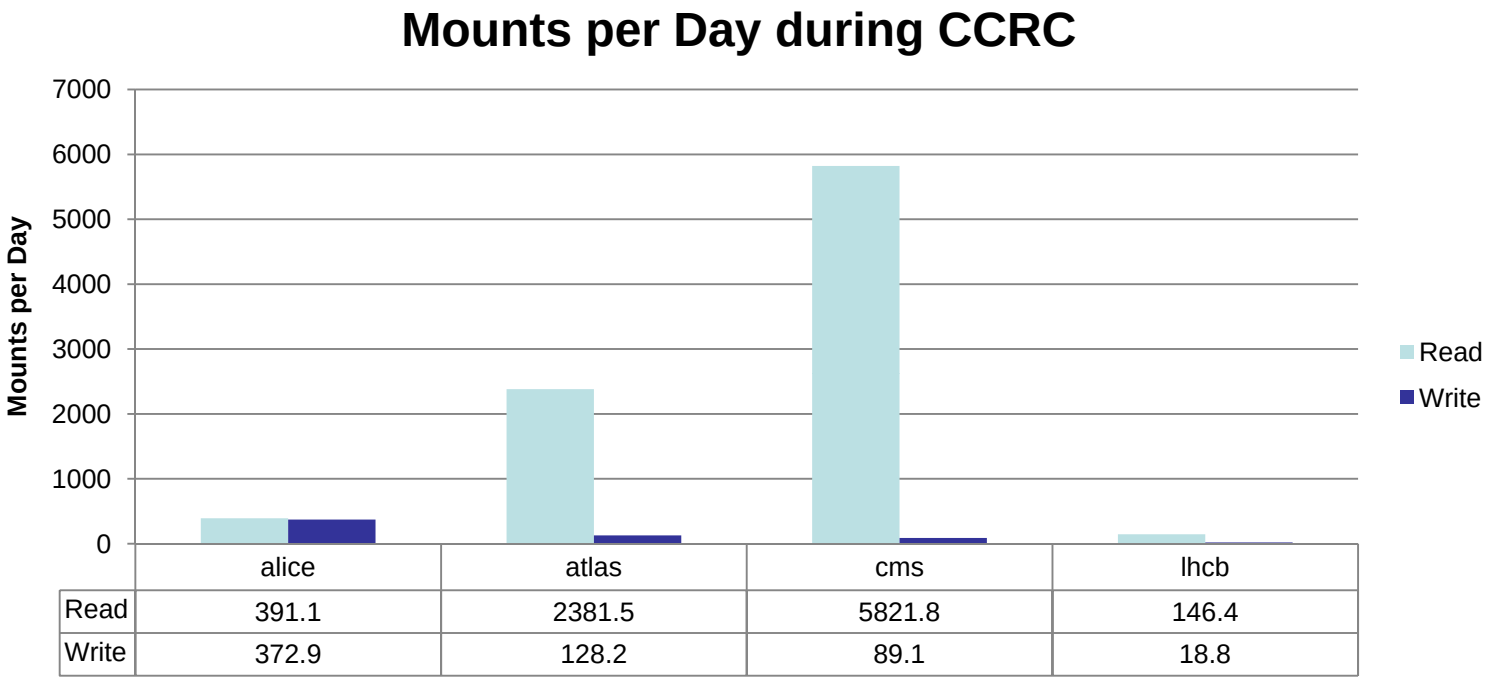

	alice	atlas	cms	lhcb
Read	391.1	2381.5	5821.8	146.4
Write	372.9	128.2	89.1	18.8

- Random read dominates drive time (90% reading)
- Writing under control of Castor policies
- Reading much more difficult to improve from the Castor side

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Production vs Users - CMS

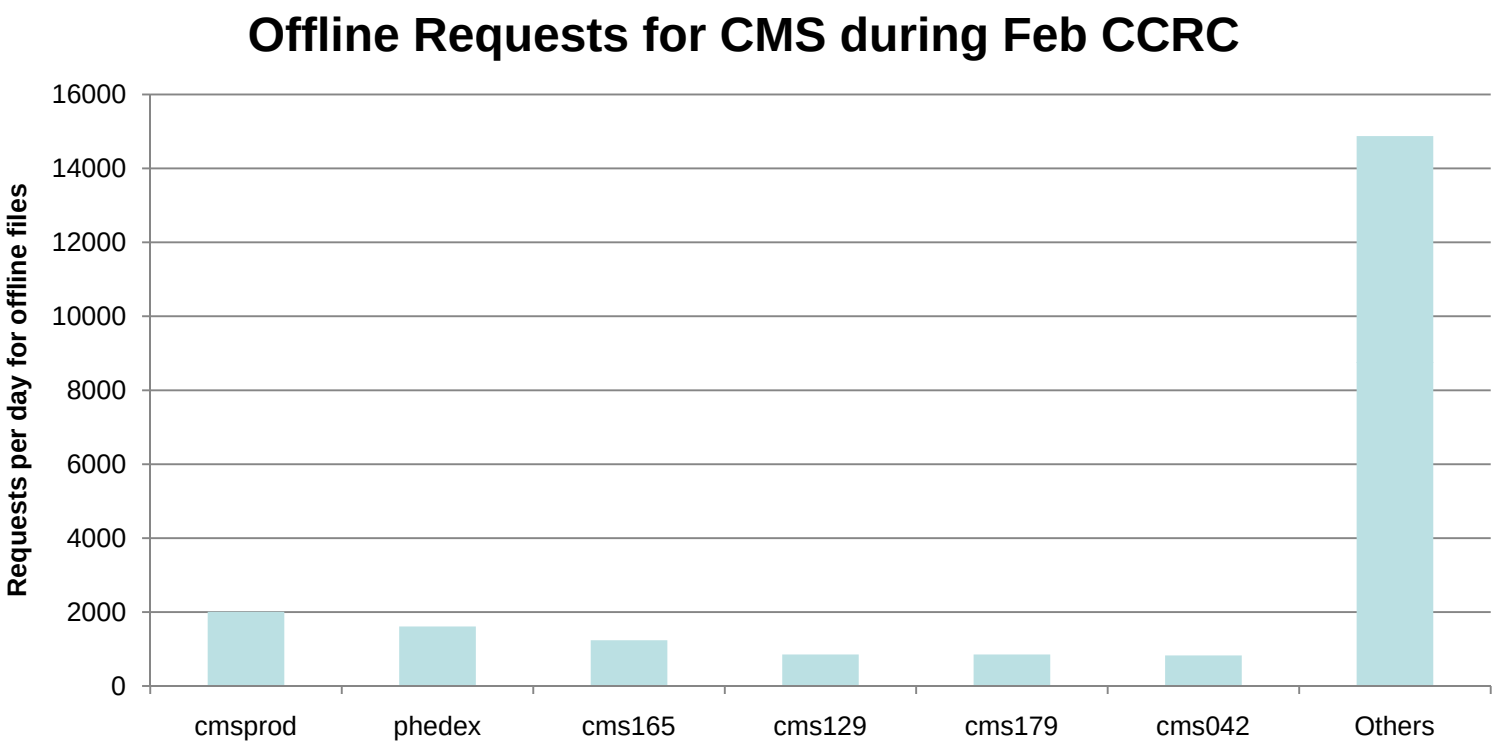

- Counts of requests for production files which were not on disk during CCRC period for CMS
- CMS production is under cmsprod and phedex (16% total)
- Requests for tape recalls dominated by non-production
- Full user list available on request

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Production vs Users - Atlas

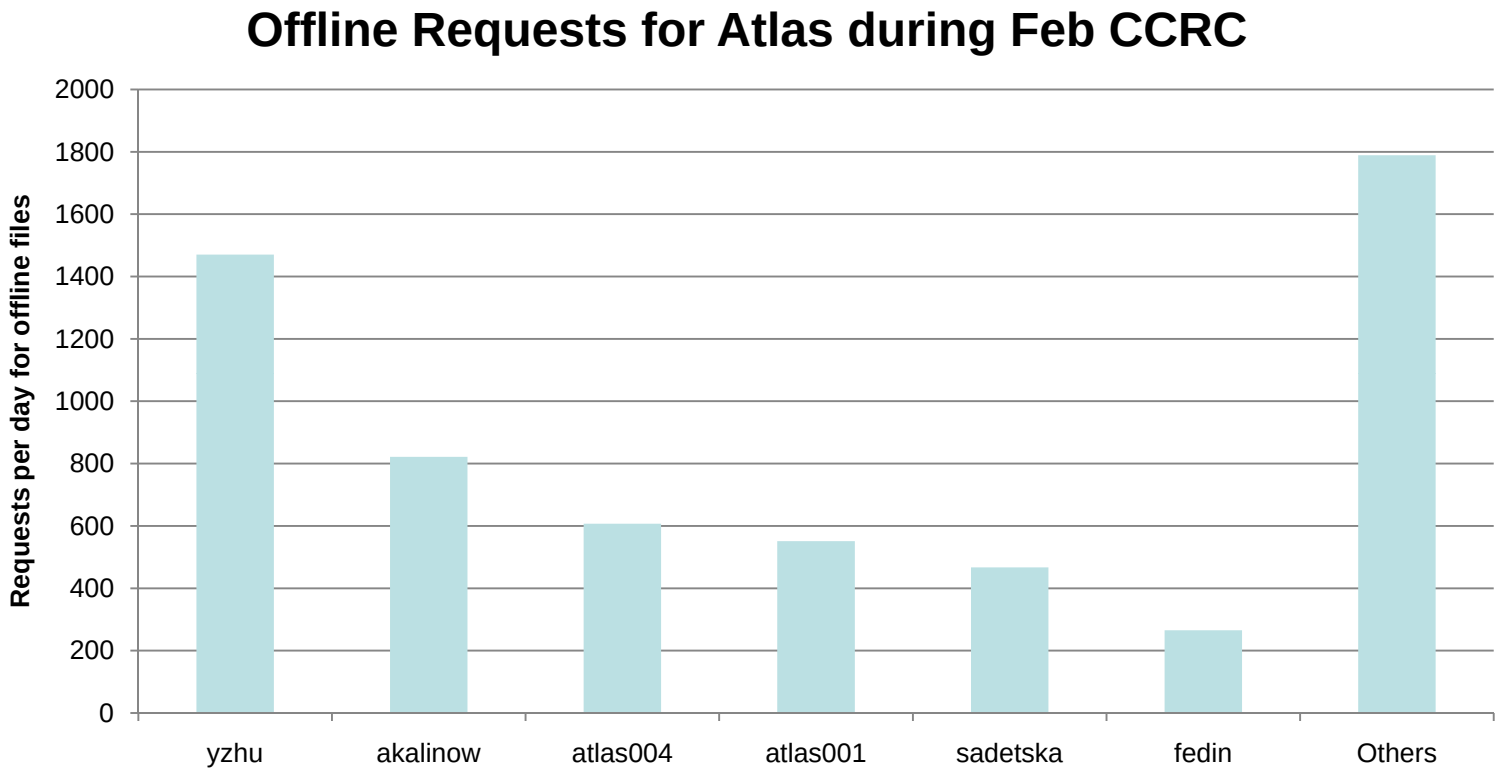

- Count of requests for production data files which were not on disk by user
- Requests for tape recalls dominated by non-production

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tape Review of CCRC

- Service ran well during CCRC
  - One robot failure was transparent to end users
  - Tape server and drive maintenance was transparent
  - Peaks of 4GB/s writing, 5 GB/s reading
- Tuning approach successful for write
  - New write policies doubled write performance for CMS and Atlas
  - Atlas performance improved when large simulation files were used
  - LHCb and Alice will improve with Castor 2.1.6
- Read mount share remain high
  - 90% of the mounts but only 45% data transferred compared to write
  - Production users competing for resources with less efficient end users
  - **End users using Tier-0 resources ?**
  - **Is the contents/size of the disk caches correct ?**
- Dedicating tape resources may be required
  - Allocate drives / robots to each VO to ensure fair share
  - Reduce resilience as drive or robot failure has larger impact
  - Monitoring to continue with the implementation of read policies during March through May CCRC

Backup Slides

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Metrics for Tape Efficiency

- File size
  - Average size of files to/from tape
- Repeat mount rate
  - Average number of times a tape is mounted for each tape touched that day
- Data transfer per mount
  - Average volume of data transferred for each mount
- Total Rate
  - Data read/written per-VO divided by total time on drives including mount, unmount and data transfer.

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Repeat Mounting

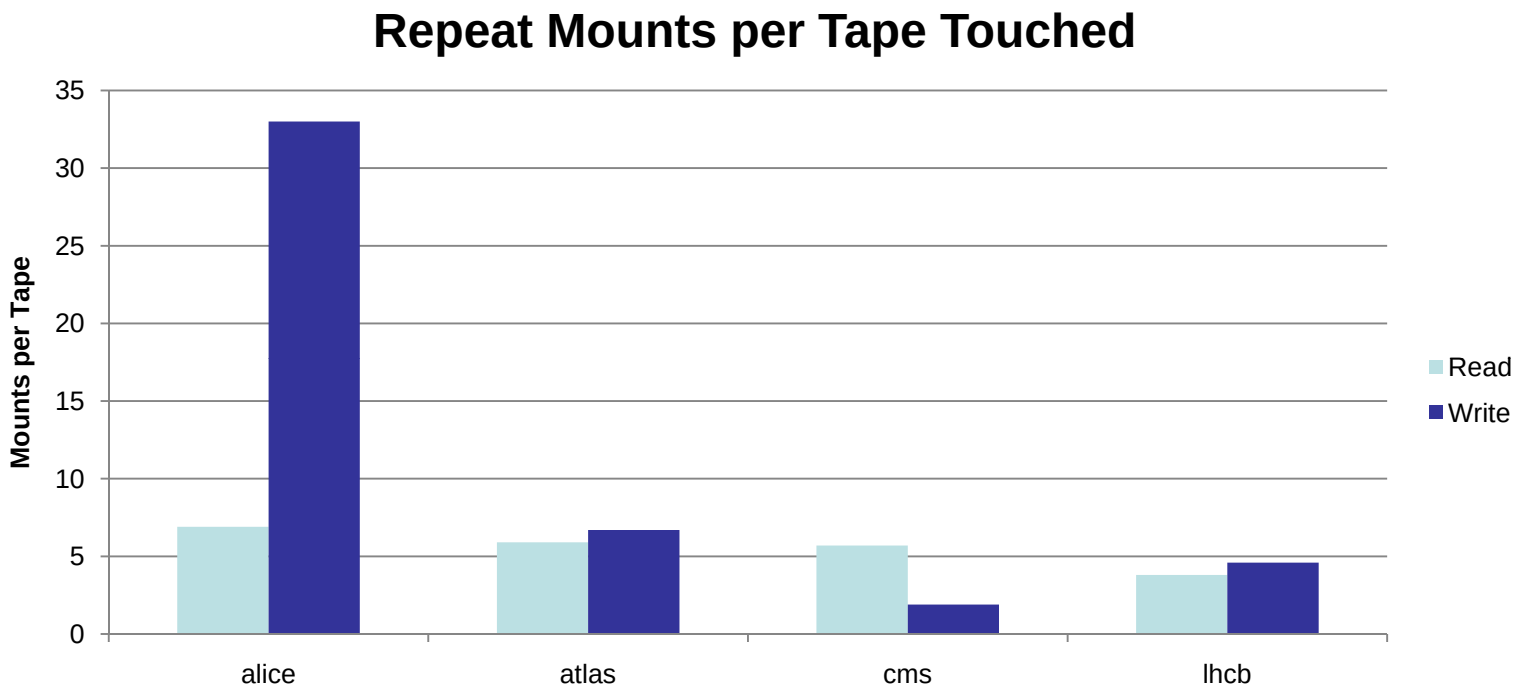

- Alice write repeat mounts will drop to at least 5 with Castor 2.1.6
- Atlas write performance limited by smaller files

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CCRC Atlas

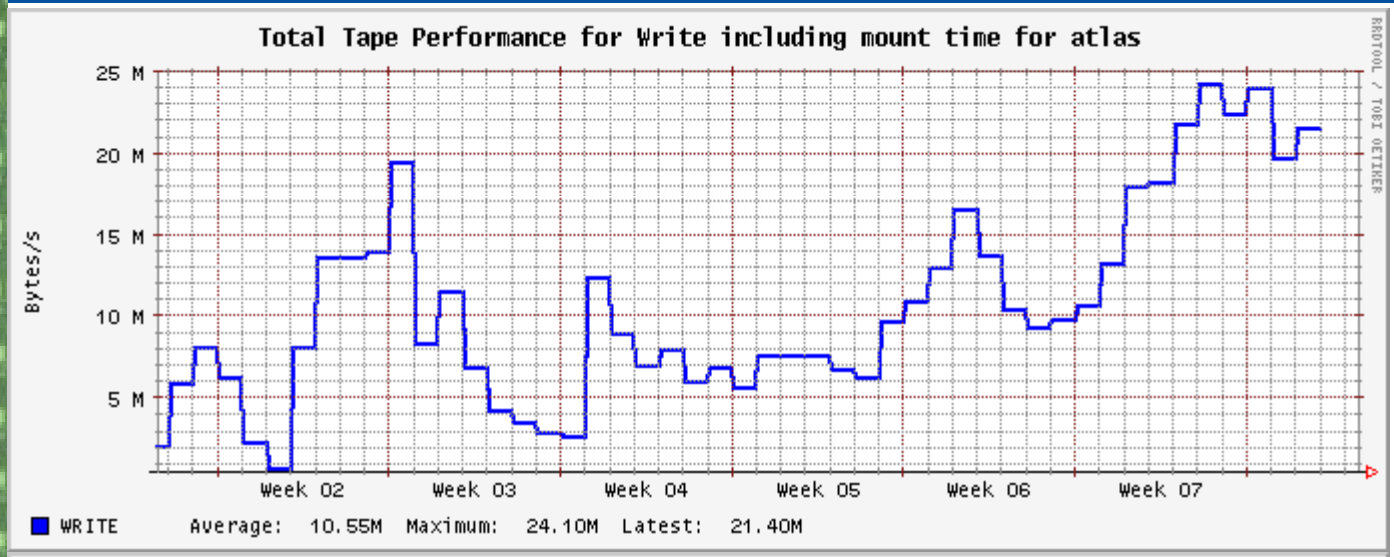

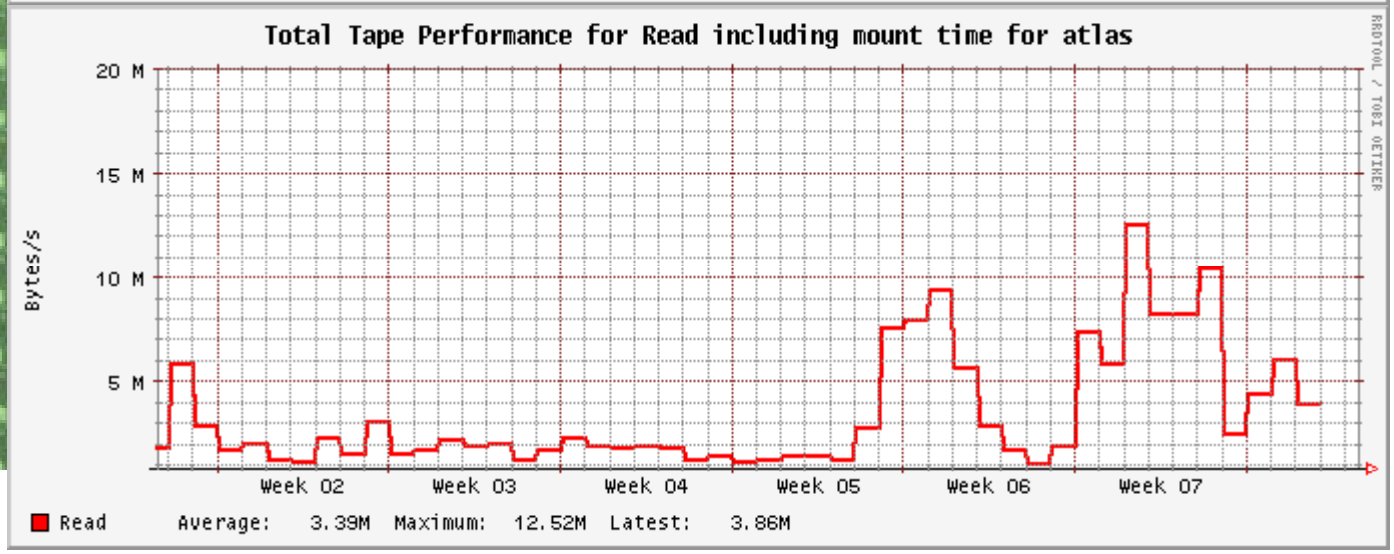

- Simulated data during week 7

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CCRC CMS

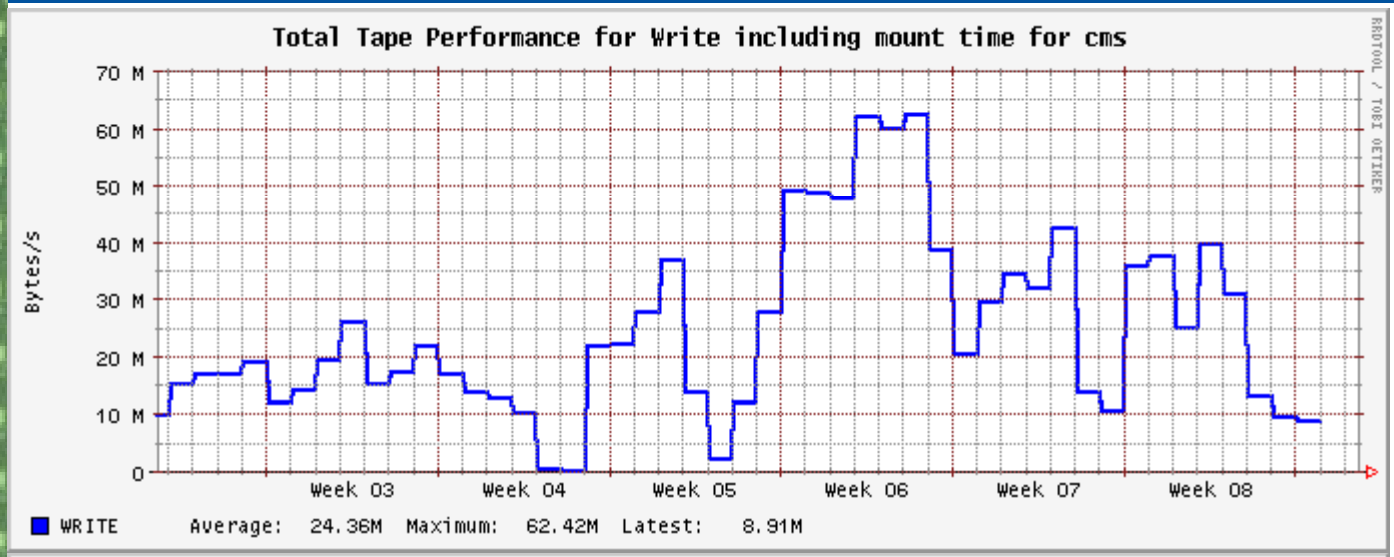

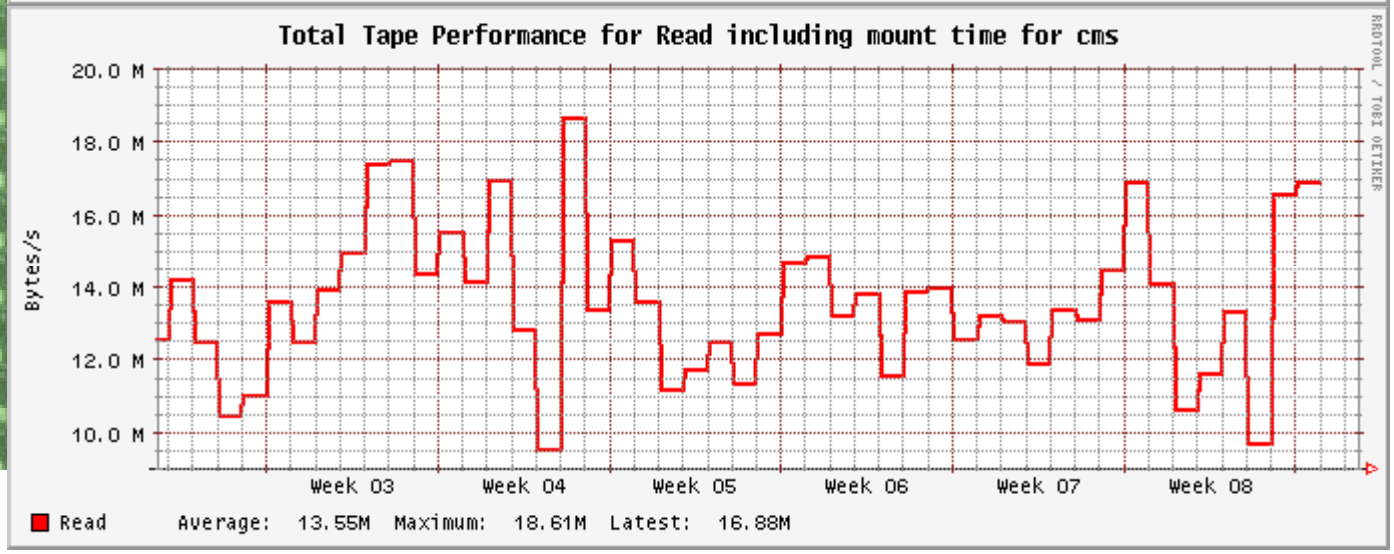

- Note scale difference for read vs write

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Additional Information

- Metrics Definition
  - https://twiki.cern.ch/twiki/bin/view/LCG/MssEfficiencyCERN
- Tape Efficiency Summary
  - https://twiki.cern.ch/twiki/bin/view/LCG/MssEfficiencyCERN