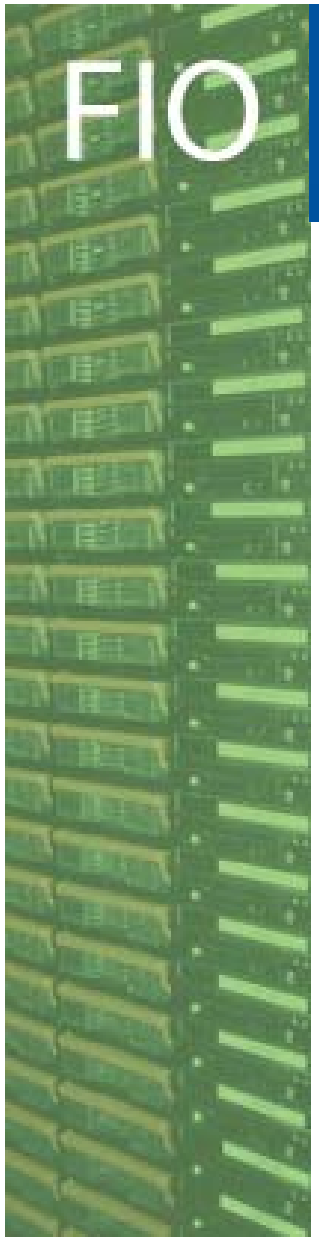

Fabric Infrastructure and Operations

CERN IT Department

Tape Efficiency CERN

Tim Bell

GDB February 2008

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

FIO

Metrics for Tape Efficiency

CERN IT Department

- 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
  - Volume of data read/written divided by total time on drives including mount, unmount and data transfer.

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

Proposal from F2F

- Experiments should
  - Move to 2GB files for tape transfers
  - Ensure that pre-staging is standard for all applications
- Castor Operations will change policies for CCRC
  - Write policy of at least one tape of data with 8 hours maximum delay
  - Limit mounting for reads unless at least 10GB or 10 files requested for each read mount or if a request is 8 hours old
- Monitor February CCRC performance and cover shortfall with
  - Major drive purchases and dedication for experiments
    - Fixed budget! Implies reduction in CPU/disk capacity

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

Changes Made to Castor

- Write policy changes
  - Checking for files to migrate moved from 30 minutes to 2 hours
  - Migration policy changed for CMS and Atlas with Castor 2.1.6 to migrate when 500 files are available to be written
  - Gradual approach to understand impacts
- Read policy will be implemented after write changes have been validated
- Awaiting CCRC data load to confirm results
  - Automatic collection of data for metrics in place
  - Weekly results available on Twiki

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

CCRC and Write Migration

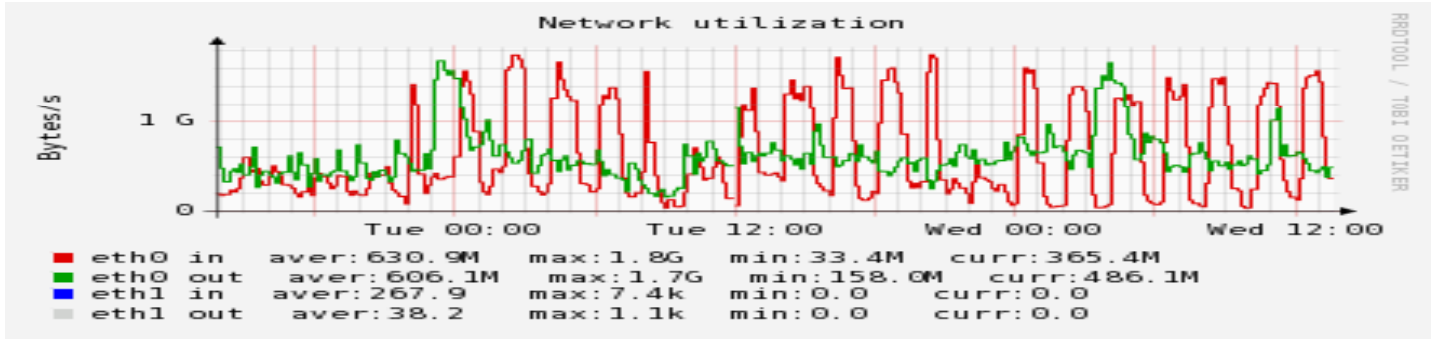

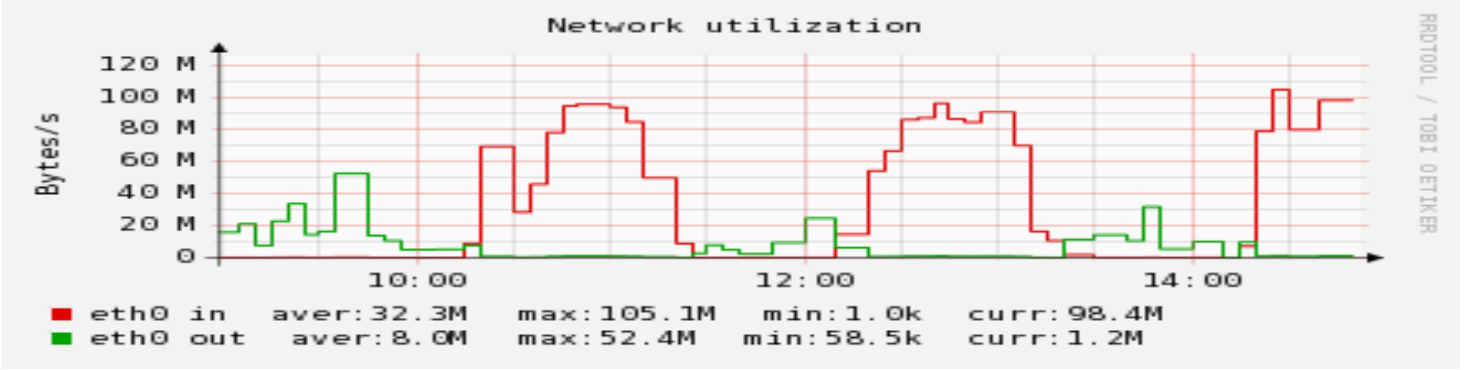

- Regular pulses as the CMS write migration starts
- Write policy is sufficient as it drops to zero before next migration
- Extra capacity available for other activities such as reading

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

FIO

Other Improvements

CERN IT Department

- Atlas file sizes have increased from 150MB to 250MB average
- Alice have identified the 22 byte file problem
- Tape pool definitions for CCRC and production are being made with consideration for recall strategies

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

Discussion summary

- Old Repeat Mount metric was confusing
  - Simplified the definition as average mounts for each tape touched
- Data is not representative of future load
  - CCRC should be representative so full analysis will be made after on completion of CCRC
- Users write small files
  - Only production pools are considered in CERN reporting
- Limited control on distribution on tape
  - Improve writing policies to get better co-location
  - Pre-staging needs to be larger than per-job to gain efficiency
- Historical files are small so current file sizes are small
  - Should only effect read file size and should gradually improve as write file sizes improve

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

Current Statistics

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

File size and performance

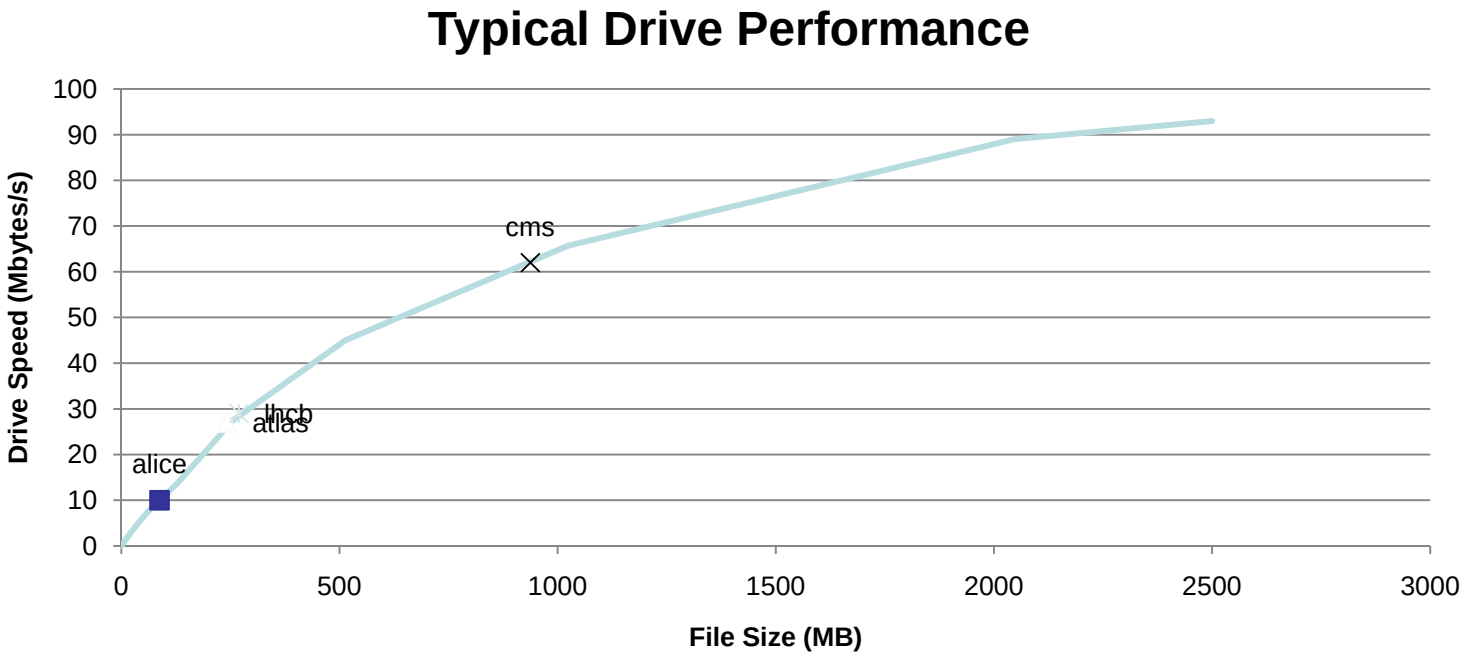

Alice	Atlas	CMS	LHCb
90 MB	250 MB	950 MB	270 MB

- Atlas file sizes improved from 150MB to 250MB
- Awaiting CCRC for representative work load

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

FIO

File size

CERN IT Department

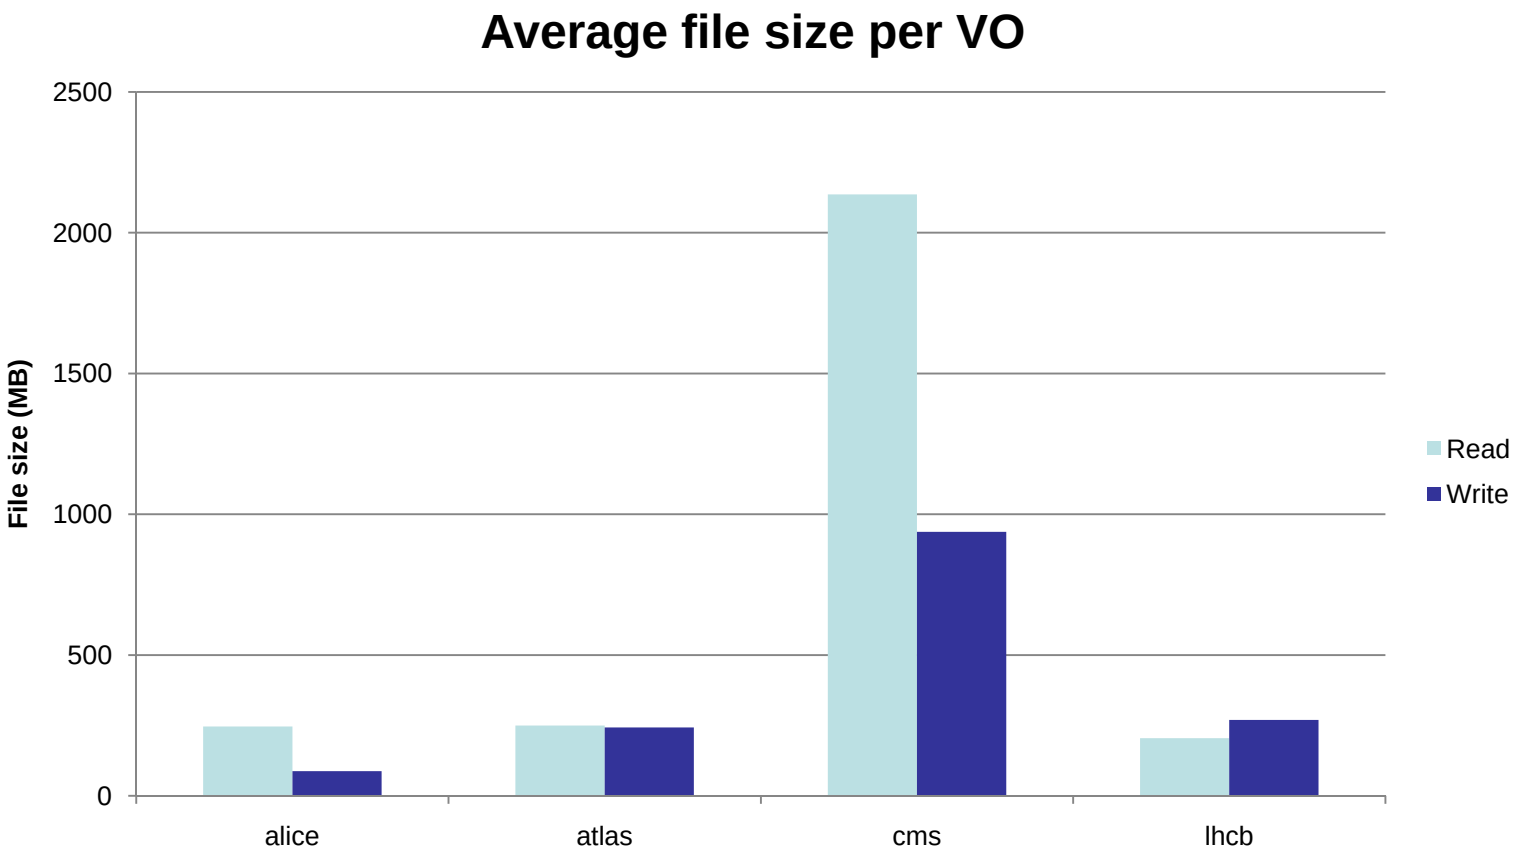

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Repeat Mounting

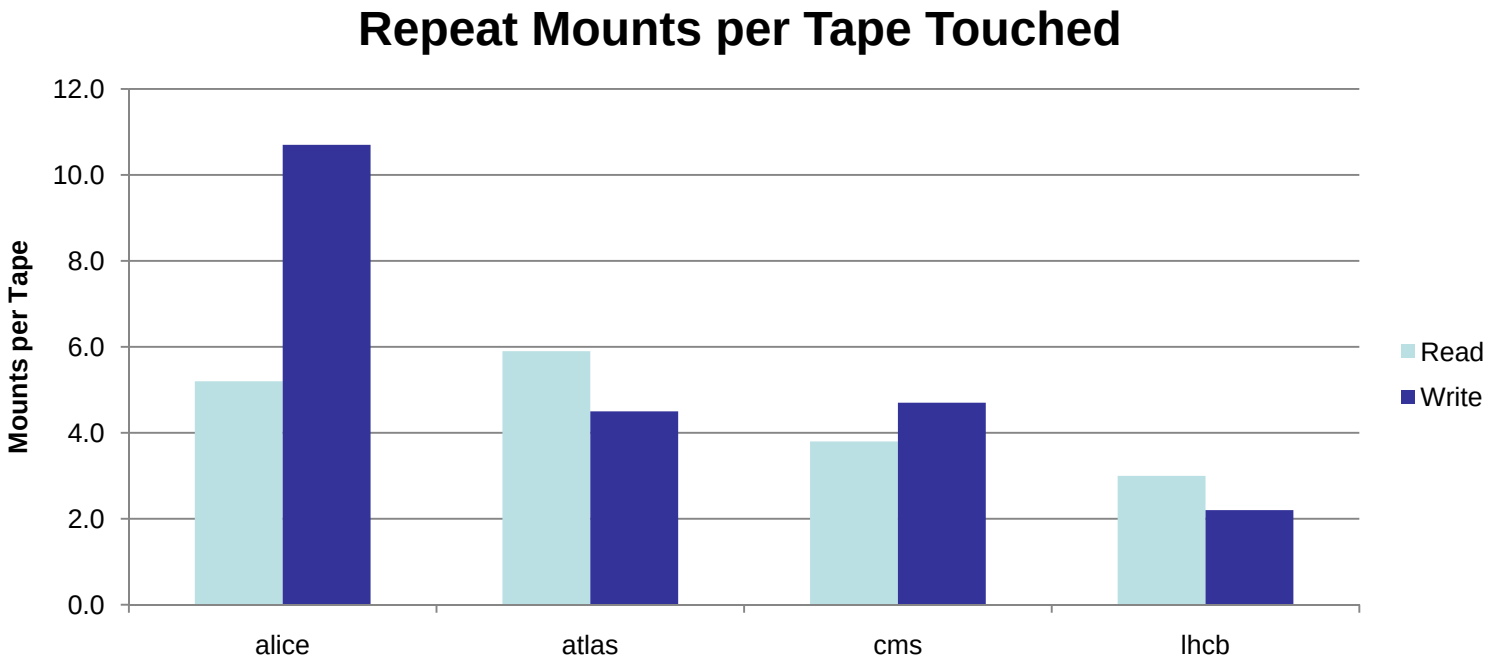

- Improved metric to ease understanding
- Requires pre-staging for reading and policy changes for writing

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

Total performance to tape

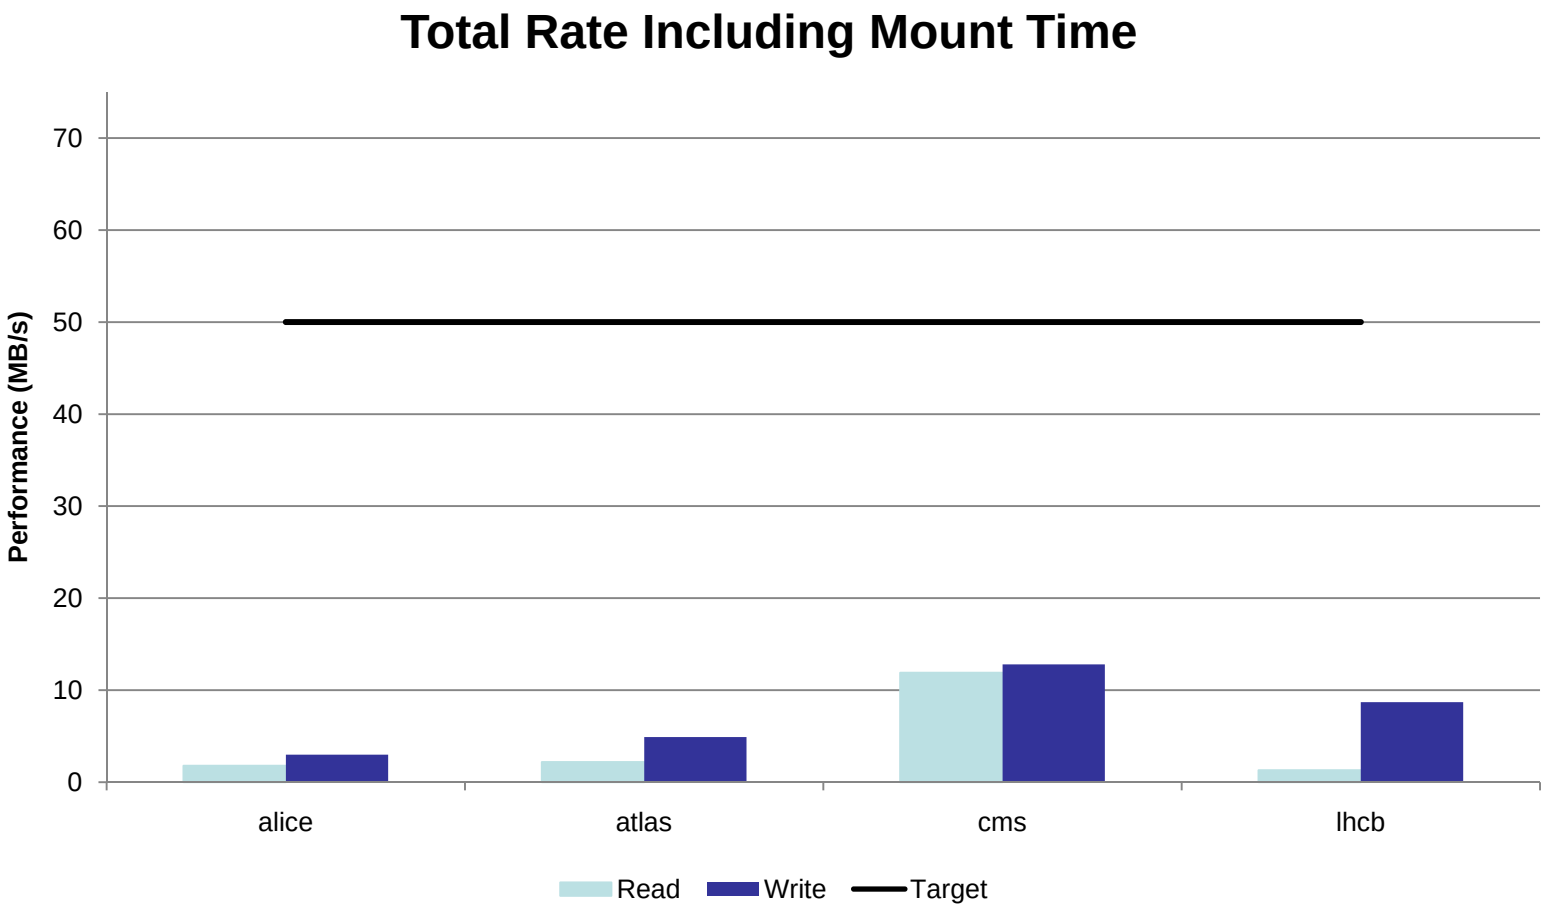

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Data transfer per mount

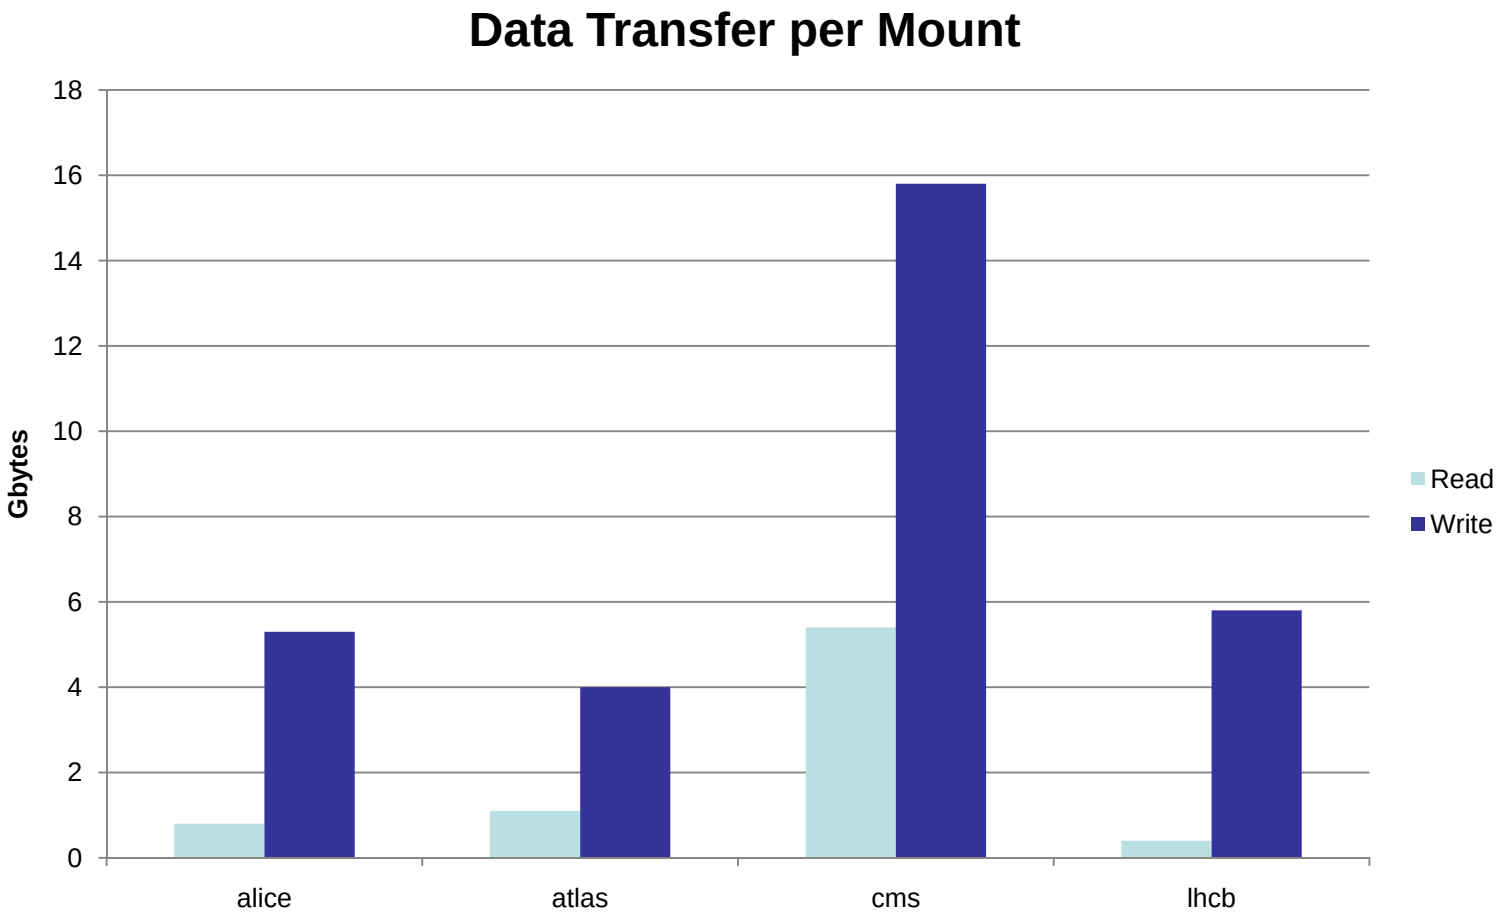

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

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

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