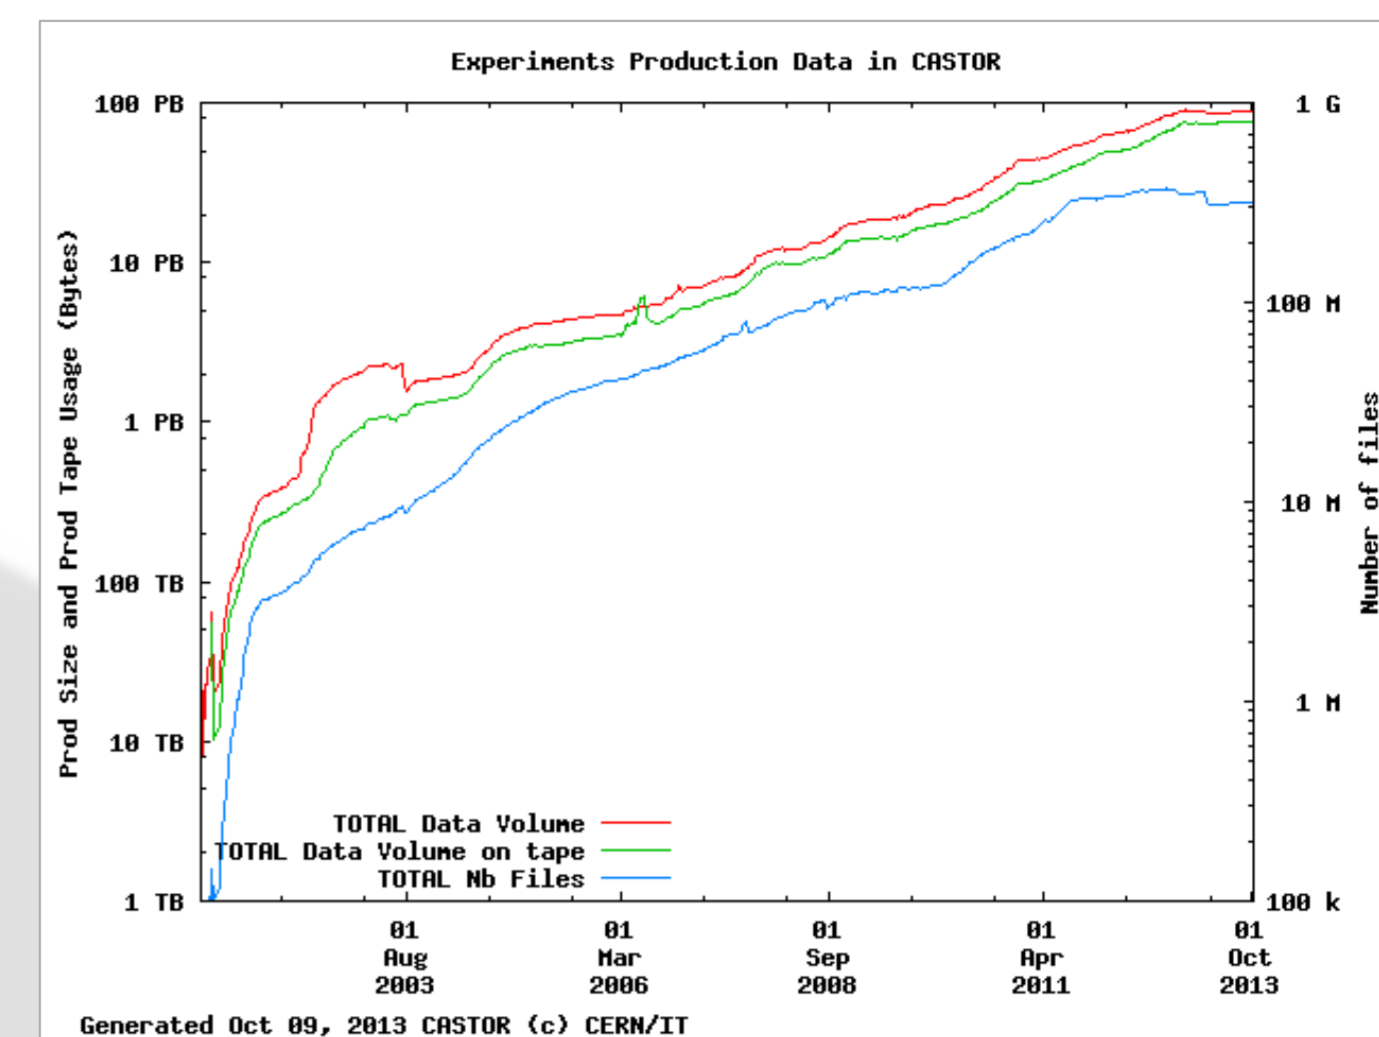


# The data torrent

- **92 PB** total archive, 75 PB since LHC start
- Doubling size every **19 months** since **2002**
- **1 PB/week** sustained during LHC operations



semilog scale!

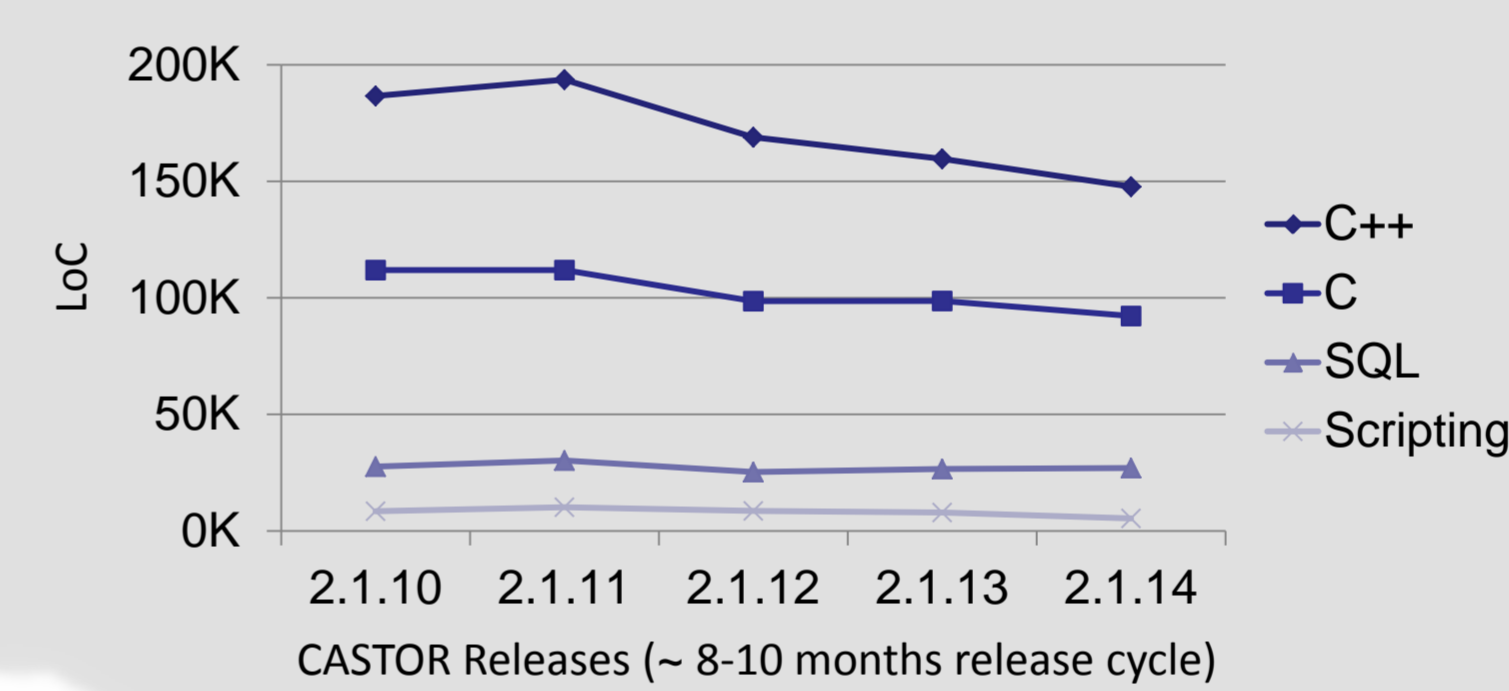


## Streamlining Operations

- Draining, Rebalancing, Read-Only mode
- Bulk metadata processing
  - Logic (PL/SQL) closer to the metadata
- Improved scheduling throughput
  - **250 Hz** write operations
  - **1 kHz** namespace operations
- Ready for the ‘Repack Challenge’:
  - ~**100 PB** of data to be moved in **12 months**
  - (more on D. Kruse, The Repack Challenge, CHEP2013 Poster)

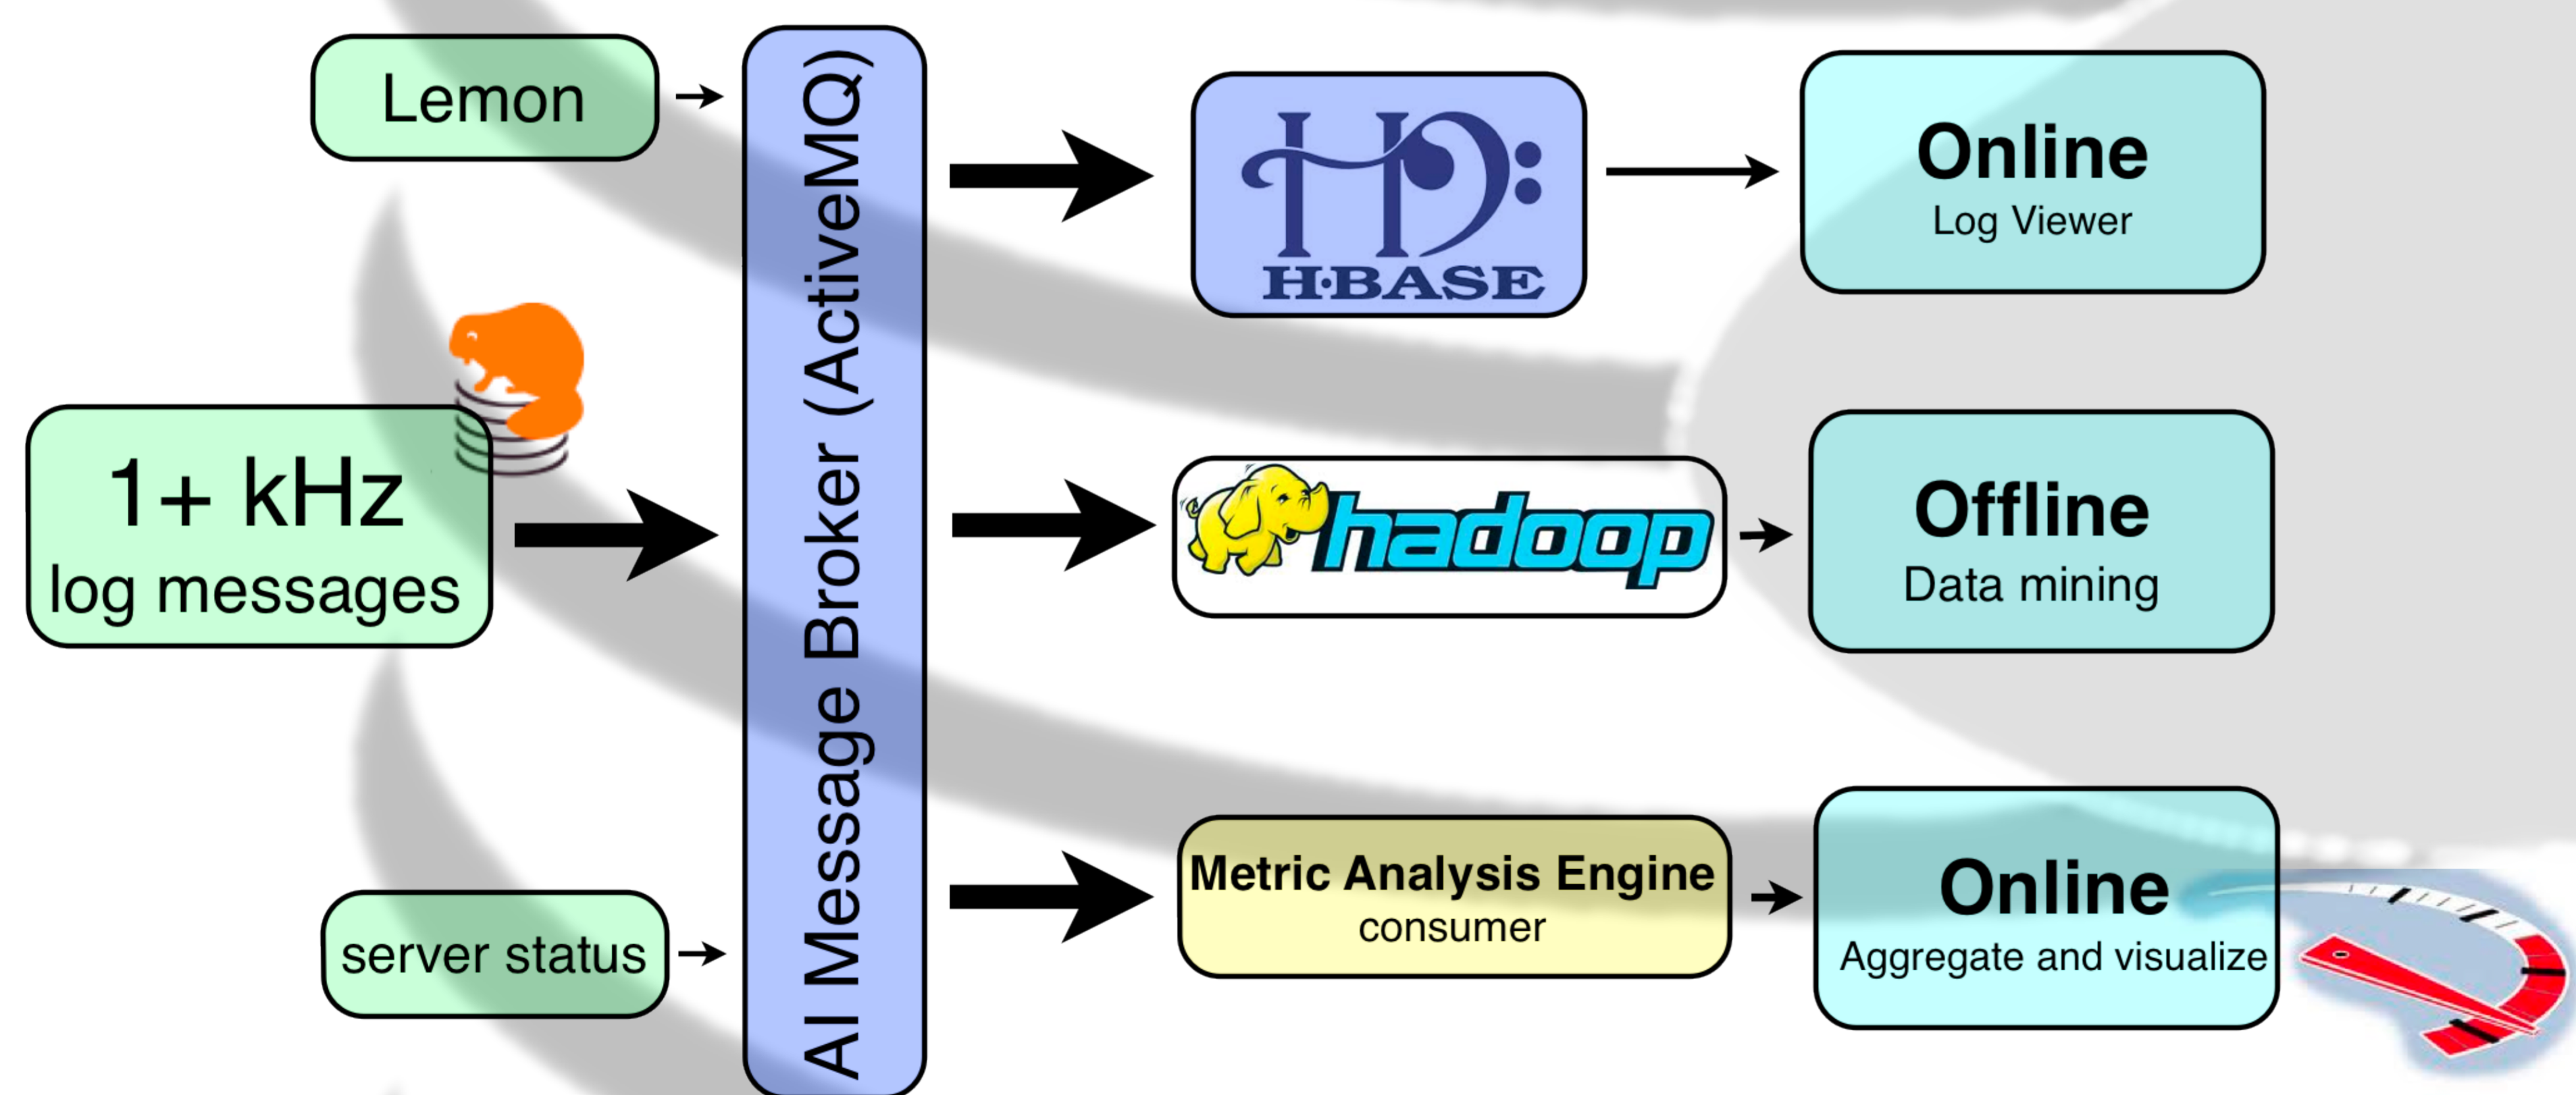
## Specialization

- Towards a **Tape Archive** system
  - Data placement managed by the experiments
- Secured namespace
- Prototyping an **AFS Backup** solution
- Reducing code base, whilst ensuring continuous improvements

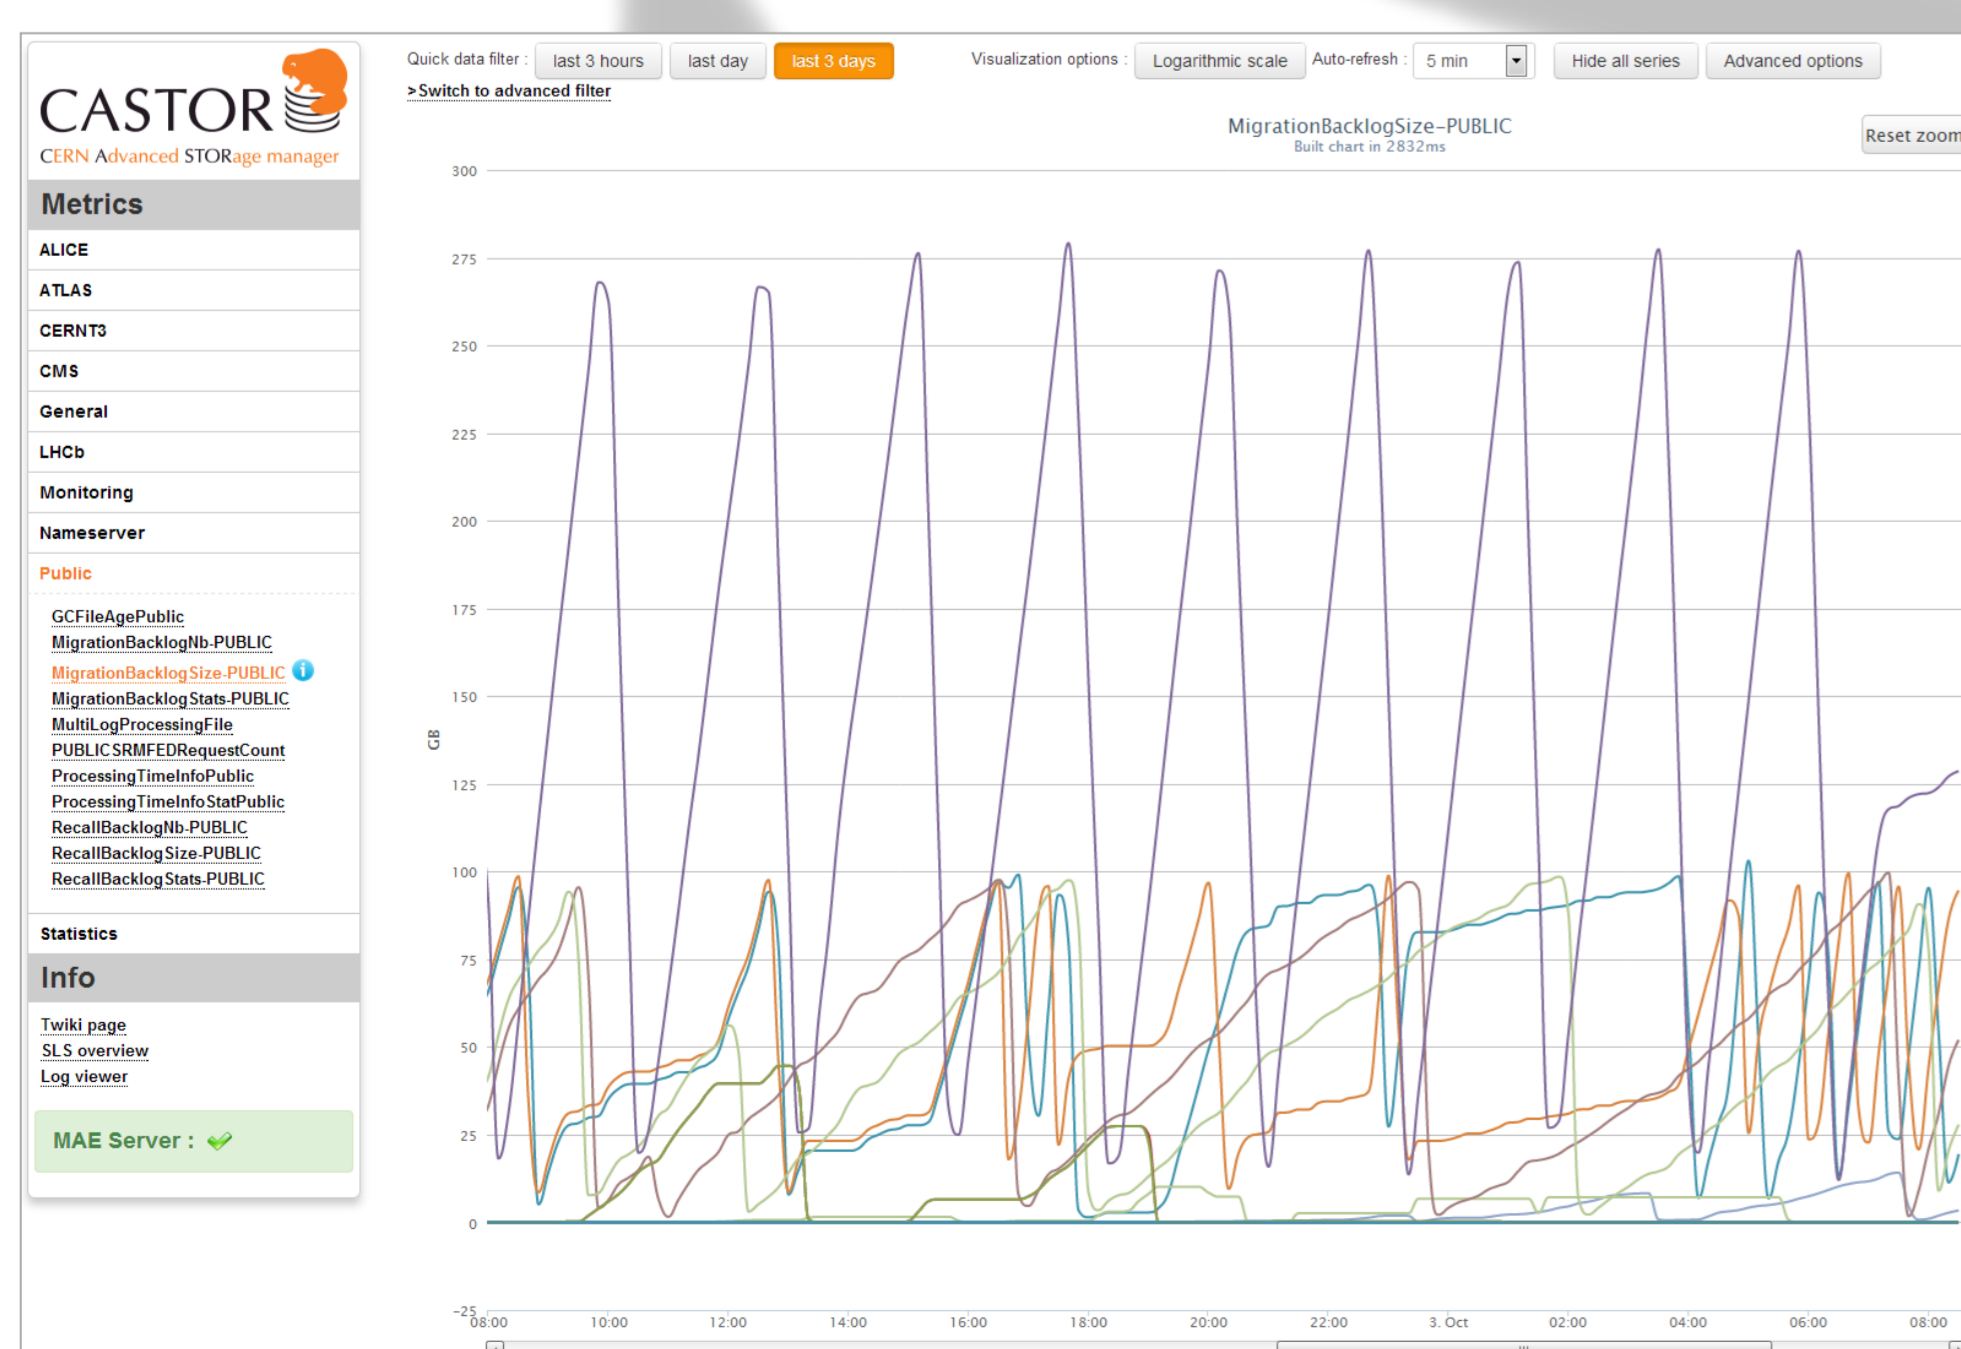


e.g. refactoring of the disk-to-disk copy logic in version 2.1.14, enabling operational improvements for draining and rebalancing.

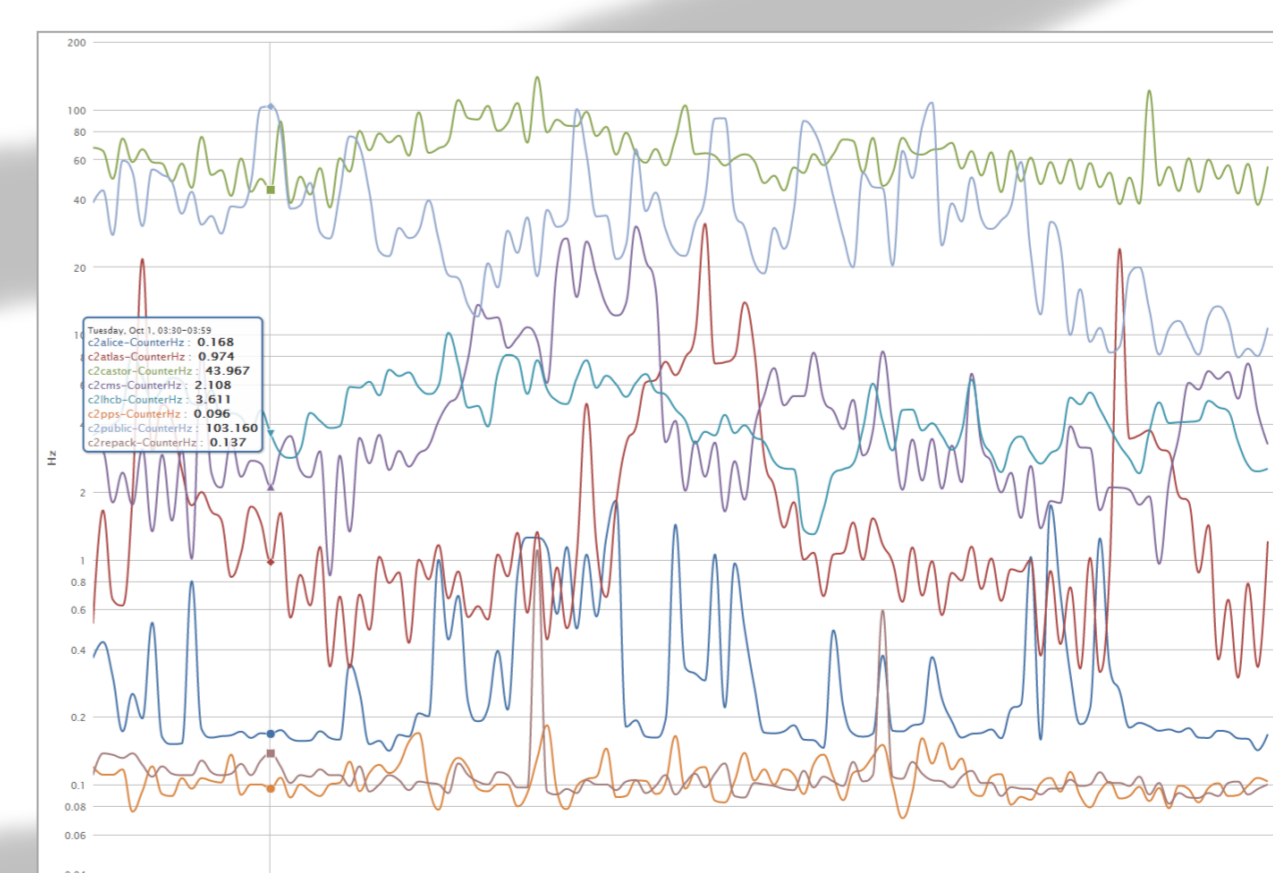
## A new Monitoring Infrastructure



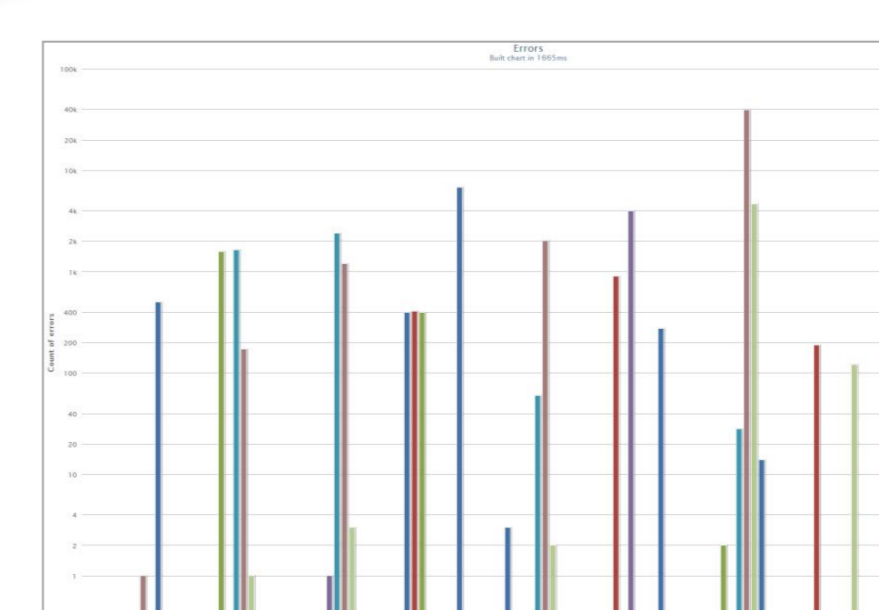
- Coherent architecture based on IT Agile Infrastructure common tools
- Exploiting emerging paradigms
  - **Apache HBase** for fast indexing and retrieving of log messages
  - **Hadoop MapReduce** for offline trend analysis



The cockpit interface, featuring a number of real-time metrics: here the migration-to-tape queue during 24 hours. Migration is typically only triggered when more than **200 GB** of data is waiting, or when **4 hours** have passed since the data were written.



The throughput in terms of requests per second for all CASTOR instances during one day.



The count of log messages classified as errors for the different components in each CASTOR instance.