

Experiences and challenges running CERN's high capacity tape archive

Germán Cancio, Vladimír Bahyl, Daniele Francesco Kruse, Julien Leduc, Eric Cano and Steven Murray

CERN – European Organization for Nuclear Research, CH-1211 Geneva 23, Switzerland

E-mail: {German.Cancio,Vladimir.Bahyl}@cern.ch

Abstract. CERN's tape-based archive system has collected over 70 Petabytes of data during the first run of the LHC. The Long Shutdown is being used for migrating the complete 100 Petabytes data archive to higher-density tape media. During LHC Run 2, the archive will have to cope with yearly growth rates of up to 40-50 Petabytes. In this contribution, we describe the scalable setup for coping with the storage and long-term archival of such massive data amounts. We also review the challenges resulting and mechanisms devised for measuring and enhancing availability and reliability, as well as ensuring the long-term integrity and bit-level preservation of the complete data repository. The procedures and tools for the proactive and efficient operation of the tape infrastructure are described, including the features developed for automated problem detection, identification and notification. Finally, we present an outlook in terms of future capacity requirements growth and how it matches the expected tape technology evolution.

1. Introduction

The CERN Advanced STORage manager (CASTOR) [1] is a hierarchical storage management system developed and used at CERN for storing physics data. Its main components comprise a central name server containing metadata about all files; a stager managing a disk cache layer; and a tape-based backend for permanent data archiving.

During the period of the first LHC run (2010 to 2012), over 70 Petabytes of new data were collected from the various experiments. Together with previously archived data, the capacity of the CERN data archive recently exceeded the 100 Petabytes mark.

In terms of infrastructure, 4 Oracle StorageTek SL8500 tape libraries located in two different buildings provide in total 40 000 cartridge slots. In each building there are 2 of those libraries interconnected in a complex, each complex equipped with 20 Oracle T10000D tape drives. In addition, 3 IBM TS3500 tape libraries provide another 26 000 cartridge slots and are equipped with 40 IBM TS1150 drives.

2. Large-scale media migration

By the end of 2013, the CERN physics archive comprised over 50 000 tape cartridges of 4 different types from 2 generations of tape drives: 5TB and 1TB cartridges from Oracle, and 4TB and 1TB cartridges from IBM. As new drive generations were arriving on the market and 1TB media was becoming obsolete, the LHC shutdown period offered an ideal opportunity to migrate (or “repack”)



the archive data from legacy cartridges and formats to higher-density ones. The preparations for this large-scale repack exercise were already discussed in a CHEP'2012 paper [2]; here we focus on the experiences observed so far. This media migration was subject to multiple constraints, notably:

1. The repack activity should be non-intrusive and transparent to normal user activity. This was achieved by granting user jobs higher priority, and implementing an automatic drive dedication system that ensures a percentage of drives available for user activity whenever there are user jobs in the queue.
2. Temporal file collocation needs to be preserved, as closely related data sets stored on the same tape are likely to be requested together in the future. For this, we simply make sure that the input cartridges were recalled in the same exact time sequence as they became full.
3. The migration process should finish before the LHC Run 2 starts so the infrastructure is fully available by the time of data taking. Our calculations indicated that this would be achievable without purchasing additional equipment, provided that the disk and tape layer can be efficiently exploited close to nominal transfer speeds.

Enterprise-class tape technology allows for reformatting existing 4TB/5TB media to higher capacity. The initial focus was therefore to liberate space on existing Oracle 5TB cartridges by reformatting and self-repacking them to 8 TB, and then move 1TB media contents onto the freshly liberated space. The same process was later applied with the IBM 4TB media that were reformatted to 7TB.

Such a large data migration operation involving multiple drive generations can be divided in two main stages.

In a first stage, legacy tape drives can be used for reading, while the output streams can benefit from faster new-generation tape drives being all available for writing. In order to avoid starvation of the faster new-generation destination drives, it is possible to use the newer tape drives for reading (if they are backwards read compatible); this however reduces the number of drives available for writing new media. In addition to that, we also verify every freshly filled tape by re-reading it completely (more details in section 4), which brings further down the number of available destination drives. In order to avoid overflows and/or saturation of output streams, a recall trimmer was implemented. This trimmer periodically monitors the occupancy of the repacking disk cache, and if necessary limits the number of concurrent tape recalls.

In a second stage, with more and more data residing on new media, user jobs also follow the destination where the data resides so over time an increasing fraction of those requests will be directed to the new-generation drives. With user activity significantly increasing towards the end of the LHC shutdown period, this was risking to further reduce the number of available repacking destination drives. Fine-tuning of priorities and drive allocations between user, repack and verification activities allowed maintaining an appropriate balance and avoiding blockages. In particular, verification activity was delayed to a large extent; this was possible given that in case of media write problems, the source repack tapes were still available for re-reading.

All of the Oracle tapes were repacked by end of September 2014. In December 2014, we installed 35 newly released IBM TS1150 tape drives and started IBM media migration, which will have completed by the time LHC data taking restarts in June 2015.

Figure 1 shows the activity of the CERN's tape infrastructure in the period between 2008 until March 2015. For both reading and writing, repack activity exceeded that of the LHC Run 1 by a factor of two. The sustained migration rate was above 3GB/s, often reaching peaks of 10GB/s, which gives confidence that the archive will be able to sustain the performance required for LHC Run 2.




Figure 1. Tape writing and reading during LHC Run-1 and during the repacking exercise.

By the end of the repack exercise, over 30 000 obsolete 1TB tape cartridges from the tape libraries will have been decommissioned from the CERN tape libraries, with less than 20 000 tapes remaining in the robots. Those liberated slots will provide significant room for growth during subsequent LHC runs.

3. Media Contamination Incident

During the repack exercise, 13 tape cartridges were identified on which a significant portion of the data could not be read. Inspection of the cartridges at the vendor's data recovery lab revealed many scratches and holes on long stretches of media. Foam or concrete particles (potentially caused by drilling) were suggested to be the possible source of contamination. We isolated the incident in time and physical location by verifying all tapes in all libraries in the computer room; investigations involving detailed log analysis allowed to trace down the problem source back to two tape drives, which were extracted from the library and sent to the vendor for analysis. Following low-level recovery at CERN and at the vendor lab, we were able to recover 94% of data on the affected tapes, leaving 113 files as permanently lost.

In order to prevent a similar incident to happen again, we devised a number of measures, including: a detailed measurement of dust levels by an external company; performing an in-depth internal cleaning of cartridges and library slots; and developing a tape library specific environmental monitoring device. The environmental monitoring device will be hosted in an empty tape drive tray and placed inside an Oracle SL8500 tape library. It primarily consists of the following inexpensive off-the-shelf components:

- Raw HVAC sensor component detecting particles $> 1 \mu\text{m}$
- Precise temperature and humidity sensor
- Arduino processor reading the sensors
- Raspberry Pi processor communicating with external infrastructure

The aim of this environmental sensor prototype is not to provide calibrated measurements, but rather to alert the operations team when an environmental anomaly is observed. It will continuously sample the surrounding air and issue an SNMP trap if airborne particle density, humidity or temperature crosses configurable thresholds. We have demonstrated that the above prototype behaves comparably as commercially available and calibrated air quality monitoring systems, but at a small fraction of the cost. We plan to deploy those devices in all our tape libraries. Further details on the prototype can be found in [3].

4. Archive reliability

One of the characteristics of the CERN data archive is the longevity of its custodial data. Data is stored indefinitely and not deleted from the archive unless done so by the user. Some of the archive files have been migrated across different software and media generations for over 30 years. In order to minimise data loss and improve the archive reliability, a number of strategies have been devised, implemented and refined over the last years. These include notably:

- **Systematic verification:** Every time a tape is filled, the complete tape is verified by re-reading all its content and validating its metadata (such as file sizes and checksums). The same procedure is applied to “cold” tapes i.e. cartridges that have not been mounted recently, therefore ensuring the correctness of older repository data. Around 10% of the drives are continuously busy running verification jobs. These jobs run at lower priority than user-initiated read or write requests, for not interfering with user activity.
- **Reducing tape mounts:** In order to reduce tape mounting and increase efficiency, a policy-driven engine examines each tape read request and decides on whether to grant a tape mount or postpone it. This takes into account criteria such as user/group priority, number of files and amount of volume to be read, waiting time, and concurrent drive usage by the user/group. Since deployment in 2010, and despite continuous file and volume recall increases, the average number of daily tape read mounts has been reduced from over 4000/day to 1500/day.
- **Eliminating tape back-hitch:** Back-hitching (or “shoe-shining”) refers here to the stop and backward motion a tape drive makes whenever it writes an end-of-file tape mark. By default, these file marks are synchronizing operations that require a drive to empty its internal buffers to tape, stop motion and then rewind to the position following the end of the previous file. Repeating this operation creates a high physical strain on the media as such an operation is executed 3 times per file with CASTOR’s AUL-based tape format. As described in detail in [4], “soft” or “immediate” tape marks were added on CERN’s request to the Linux “st” tape driver. Synchronising file marks are nowadays only issued after writing larger amounts of data on tape, which reduces the number of back-hitches by two orders of magnitude (from over 3000 to around 30 per terabyte written).
- **Media lifetime:** Media at CERN is typically kept in production for not longer than two drive generations (typically 6-8 years). This is well below the physical media lifetime, which is around 20-30 years; data is moved to new media technology, which avoids obsolescence and increases reliability over older generations.
- **Decommissioning often-mounted media:** We monitor the number of times a cartridge gets mounted over its lifetime, and decommission it after 5000 mounts. In practice, we only get to decommission 3-4 cartridges / year as most of the tapes are well below the threshold for the reasons outlined above.
- **For smaller virtual organisations (VO’s) secondary file copies can be created.** These second data copies are stored in a separate library residing in a different physical building. This operation has been done in particular for all legacy LEP experiments.

With these measures in place, over the last 5 years we have been able to reduce the number of annual file losses by at least two orders of magnitude (figure 2). The annual bit error loss rate in the period 2012-2015 is in the order of $O(10^{-16})$, which is still three to four orders of magnitude above the bit error rate for enterprise tape drives [5].




Figure 2. Evolution of archive file losses at CERN

In order to further increase reliability, we have investigated ANSI T10 SSC-4 Logical Block Protection (LBP) [6], which is supported by recent LTO and enterprise tape drives. Logical Block Protection protects the path between the data source and the tape media (e.g. FC interface and physical link, internal drive data channels, etc.) against errors such as link-level bit flips. It works by pre-calculating and appending a CRC code for each data block sent to the tape drive. This CRC is stored with the data on tape and re-validated by the drive via a read-after-write process. The CRC is also recalculated and checked every time the block is read from tape. We have measured the drive-level CRC calculation overhead to be minimal (below 1%) and therefore not having any performance impact. Another advantage of LBP is that the drive can do basic verifications autonomously (using SCSI VERIFY) and at maximum drive speed, given that it stores data blocks with their checksums, and no data needs to be transferred out of the drive. This can be used in complement to more extensive consistency checks performed by our verification engine, checking application-specific metadata (such as file-level checksums, the number and location of existing replicas, etc.).

5. Software and management tools

In 2014, a new CASTOR tape software layer was developed and deployed into production, replacing the previous daemons originally developed almost 20 years ago. The architecture has been completely redesigned using object-oriented design and C++ as implementation language. In terms of functionality, stumbling blocks for performance have been removed and the error detection and handling has been greatly enhanced by including full support for SCSI tape alerts (see below); Logical Block Protection (see previous point) is to be added soon. The disk access interface now supports the XROOT protocol in addition to RFIO. Further details on the new tape software layer can be found in [7].

In parallel to the development of the new tape software, the CASTOR tape alarm and logging system was reworked and significantly extended. This system is responsible on one hand for the monitoring of media, drives and libraries, and on the other, for automated fault detection and root cause isolation. Using a configurable rule-based engine, problematic elements are automatically identified and if necessary disabled. A ticket for the appropriate handler (e.g. tape operator, vendor engineer or service manager) gets created which contains a reference to the failure and the specific recovery procedure to follow. Typical fault examples include errors when reading or writing media; mechanical library failures when loading/unloading cartridges; disk server failures; or network issues.

For the cases where the failing element cannot be determined unambiguously, a rolling window of repeatedly failing requests is kept; if these accumulate for a given cartridge, drive and/or library above configurable threshold levels, the element is disabled and the appropriate support team notified which then has to manually investigate and confirm the potential cause. Until recently, most of the problems were identified by acting on repeatedly failing requests. However, thanks to SCSI tape alerts [8] that are reported via the new CASTOR tape software, the alarm system now identifies most of the issues at “first catch”. There are around 60 different SCSI tape alerts that can be reported by a tape drive, providing accurate information about warning or failure events and their root cause. Examples include warnings on performance drops when writing or reading media, physical media or drive failure or even damage, media approaching or reaching its end of life, media capacity loss, predictive drive failure alerts, cartridge memory chip failures, drive head requiring cleaning, power or cooling anomalies, drive microcode failure, etc. Figure 3 shows how the number of incident tickets created by the alarm system has evolved, with a clear increase of “first catch” type incidents since the deployment of the new tape software beginning of February 2015.

In the case of media unavailability or damage, a comprehensive repair workflow has been defined and refined over the last years. Several repair levels are defined, from local recovery up to sending the media to vendor. Users are regularly informed about status and recovery progress of their files.




Figure 3. Tape incident tickets created at CERN (May 2014 to March 2015)

6. LHC Run-2 and beyond

For the last 10 years, the data volume stored on tape at CERN has been growing at almost exponential rate. This trend is expected to continue over the next years.

For the period of LHC Run-2 (2015-2018), the forecast is up to 50PB of new data per year. The dominant fraction of this volume will consist of LHC raw and derived data (ESD, AOD) [9], complemented by data from other CERN-based or related experiments (such as AMS, NA62, or nTOF). For LHC Run-3 (2019-2022), the amount of raw data will grow even further, reaching around 100PB/year, to which approximately 50PB/year of derived and non-LHC data needs to be added. Looking further into the future, Run-4 (2023 onwards) will yield some 400PB/year of raw data, which can roughly be extrapolated to 600PB/year of total data. DAQ rates to tape peaking at 80GB/s are predicted to already be reached during Run-3. Figure 4 shows the integrated amount of data that would be stored at CERN, which would continue growing following an exponential shape and enter Exabyte-scale around 2023.

In order to assess the feasibility of these predicted data growth rates, it becomes necessary to extrapolate storage technology evolution over the next 10-15 years. For this we can refer to historical trends and the current status and results of storage research and development. According to the roadmap of the INSIC consortium [10], tape and disk storage areal density will continue growing at around 30% and 20% per year, respectively. Prototypes are already now proving the viability of technology for products in approximately one decade: Early April 2015, IBM and Fujifilm demonstrated the feasibility of reaching 123 Gigabits per square inch, which would correspond to a 220TB LTO-format tape cartridge [11]. Today's highest-capacity cartridge on the market offers 10TB; at a 30% annual growth rate, it would take approximately 11 years to reach 220TB. It should be noted that the areal bit density has the potential to grow much further: 123Gb/in² is below the density of today's highest-capacity hard disks; the ASTC consortium disk technology roadmap [12] yields for 10Tb/in² in around 10 years.

The HEPiX Bit-Preservation Working Group [13] has defined a model for long-term dimensioning and cost estimations of large-scale data archives. This model is based on a number of assumptions directly applicable to the CERN data archive, such as: tape-based archive with a small disk cache front-end; usage of enterprise-class tape equipment with the possibility of reusing media at higher density; lifetimes of 6 and 3 years for tape media and tape drives, respectively; or media migrations every 3 years to higher density or new media generations.




Figure 4. Predicted data growth at CERN, 2015-2029




Figure 5. Predicted per-tape capacity and number of tapes required at CERN, 2015-2029

The predicted CERN archive data evolution and expected capacity growth rates discussed above can be used as input parameters to the model. Figure 5 shows the predicted cartridge capacity growth, as well as the evolution of required cartridges at CERN for the 3-year periods between 2015-2029. Overall, the number of required tapes per 3-year period slightly increases up to the year 2026, but then shows a clear decrease, as tape capacity will start outgrowing CERN's data archival increases.

Following this model prediction, CERN would even be able to reduce its current tape library cartridge slots by almost half (or accommodate even higher archival requirements).

In addition to capacity, other important aspects need to be taken into account for forecasting. Tape transfer speed is predicted to increase at a lower annual rate than capacity growth (25%, versus 30%). This would imply an increase of the current number of tape drives to approximately 110 in order to sustain LHC Run-3 peak rates (80GB/s) expected around 2020. Costing is obviously another important element. The HEPiX model takes into account cost elements such as library, drive and media purchases, hardware maintenance fees, or power consumption costs. For confidentiality reasons these results are not part of this paper but in general terms, the expected cost is proportional to the number of cartridge purchases: on one hand, media cost dominates other elements such as drive purchases or maintenance fees, and on the other, past evolution shows that tape capacity per USD has evolved at a similar rate than the areal density.

It remains to be seen whether cloud-based storage solutions could result in contractions of the already shrinking tape market up to the point of affecting large-scale enterprise vendors, who could reduce or delay investments in R&D, or even retire from the tape business. It is therefore important to be attentive on the market evolution and if appropriate, devise alternatives such as cost-effective, high-reliability disk archival or commercial cloud based solutions.

7. Conclusion

The CERN Tape Archive, after successfully dealing with LHC Run-1, is about to complete a large media migration exercise during the LHC shutdown period. The longevity of the archived data needs to be preserved, and this is becoming a key and long-term activity. This notably includes improving reliability, performing bit-level data preservation and ensuring adequate environmental conditions. The LHC shutdown has been an opportunity to refresh the tape layer software and the associated operational tools, improving scalability and performance. With future LHC runs, data storage requirements at CERN will continue unprecedented growth rates and soon reach Exabyte-scale. Assuming a continued presence of tape on the market, its predicted technological evolution makes it a perfect match for future and long-lasting LHC archival requirements.

References

- [1] CASTOR homepage <http://cern.ch/castor>
- [2] Kruse D F 2013 The Repack Challenge *Jour. Phys.: Conf. Ser.* **513** 042028
- [3] Leduc J 2015 Dust sensors for long term data preservation *HEPiX Spring 2015 Workshop*, <https://indico.cern.ch/event/346931/session/7/contribution/44>
- [4] Murray S et al. 2012 Tape write-efficiency improvements in CASTOR *Jour. Of Phys.: Conf. Ser.* **396** 042042
- [5] StorageTek T10000D Tape Drive Data Sheet <http://www.oracle.com/us/products/servers-storage/storage/tape-storage/t10000d-ds-1991052.pdf>
- [6] Butt K 2007 Tape end-to-end data protection proposal <http://www.t10.org/ftp/t10/document.07/07-374r0.pdf>
- [7] Cano E et al. 2015 The new CERN tape software – getting ready for total performance *Unpublished paper presented at CHEP 2015, Okinawa*
- [8] Tape Alert v. 3.0 specification <ftp://www.t10.org/t10/document.02/02-142r0.pdf>
- [9] Bird I et al. 2014 Update of the Computing Models of the WLCG and the LHC Experiments *CERN-LHCC-2014-014*
- [10] INSIC 2012-2022 International Magnetic Tape Storage Roadmap <http://www.insic.org>
- [11] IBM Research sets new Record for Tape Storage <https://www-03.ibm.com/press/us/en/pressrelease/46554.wss>
- [12] ASTC Technology Roadmap http://www.idema.org/?page_id=5868
- [13] HEPiX Bit Preservation Working Group <http://w3.hepixon.org/bit-preservation>