Bit Preservation at CERN and ISO-16363

Introduction

CERN manages the largest scientific data archive in the High Energy Physics (HEP) domain. The archive currently holds over 185 Petabytes of custodial data from past and present HEP experiments, with some of its data being 40 years old, and most of it to be preserved "ad aeternum". The amount of data has and will continue to grow exponentially. Influx rates from CERN's Large Hadron Collider experiment are expected to augment from currently 40 Petabytes / year to around 600 Petabytes / year in a few years' time, therefore reaching archive volumes in the Exabyte-scale. CASTOR (for CERN Advanced STORage manager) (http://cern.ch/castor) is the hierarchical storage management system developed and used at CERN and other physics research sites for long-term data storage. Its main components comprise a central name server containing file metadata; a stager managing a disk cache layer; and a tape-based backend for permanent data archiving. At CERN, around 90 enterprise-class tape drives provide access to 9 tape libraries (4 Oracle StorageTek SL8500 and 5 IBM TS3500) distributed over two buildings and offering over 66'000 slots.
In order to minimise data loss and improve archive reliability, CERN is continuously devising, implementing and refining storage verification and preservation strategies and policies.
All operational procedures are defined under https://twiki.cern.ch/TapeOperations

Bit-preservation related metrics of ISO-16363 and how they are addressed by the CERN physics tape-based data archive (CASTOR)

3.3.5 The repository shall define, collect, track, and appropriately provide its information integrity measurements.
Supporting Text
This is necessary in order to provide documentation that it has developed or adapted appropriate measures for ensuring the integrity of its holding.
Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement
Written definition or specification of the repository's integrity measures (for example, computed checksum or hash value); documentation of the procedures and mechanisms for monitoring integrity measurements and for responding to results of integrity measurements that indicate digital content is at risk; an audit process for collecting, tracking, and presenting integrity measurements; Preservation Policy and workflow documentation.
Discussion
The mechanisms to measure integrity will evolve as technology evolves. The repository may provide documentation that it has developed or adapted appropriate measures for ensuring the integrity of its holdings. If protocols, rules and mechanisms are embedded in the repository software, there should be some way to demonstrate the implementation of integrity measures.

All data written to the CERN CASTOR archive is protected with an adler32 [X] file-level checksum that is generated at file creation. File-level checksums should also be pre-created by the experiments prior to file uploads and are then compared to the computed checksum in order to detect ingestion-time corruptions. File-level checksums are then also used for protecting against corruptions in data transfers between the user and the different CASTOR software layers. The checksum is stored separately from the data (cf also 4.4.1.2) and is checked every time a file is recalled or moved within CASTOR. In case of a mismatch between checksum and received data, the file transfer fail and a corresponding error message is sent back to the user.

We continuously investigate new developments for enhancing integrity.

A recent (2015) addition to CASTOR is support for ANSI T10 SCSI Logical Block Protection (LBP) [X], covering the elements in the full data path down from the application server via the tape drive to the tape media (e.g. Linux kernel to Fibre Channel controller, physical link, internal drive data channels, media recording heads etc.). Logical Block Protection works by CASTOR pre-calculating a 4-byte CRC code that is appended to each data block sent to tape. This CRC is recalculated and validated by the tape drive when receiving the data and when writing to media (via read-after-write heads). The benefit is that transmission errors are immediately discovered and notified to CASTOR, which can then act on the error and re-send the data. The CRC is recalculated and checked by the drive every time the data block is read back from tape.

In 2014, CASTOR was extended to obtain event descriptors via the ANSI T-10 SCSI-3 tape alerts specification [X] available on new-generation tape drives. 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 media read/write errors, warnings on performance drops, 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.

Any resulting tape reading or writing error or tape alert will cause events that are collected by the tape monitoring system and a ticket is immediately created in the CERN problem management system (SNOW) and assigned to a service operator.
Procedures and workflows for handling tape errors and alerts are defined under [X]. A short description of the tape monitoring system can be found in [X]. Currently, a major rework of the tape monitoring system is underway, with the aim to include low-level and vendor-specific SCSI reporting information into the monitoring and thus provide a deeper insight on the state of tape media and equipment ([X]).

[X] Adler32: https://en.wikipedia.org/wiki/Adler-32
[X] ANSI T10 SCSI Logical Block Protection:
[X] CERN procedures for handling tape errors and alerts: https://twiki.cern.ch/twiki/bin/view/TapeOperations/TapeAlerts
[X] CERN Tape monitoring system production instance: http://cern.ch/tape-stats
[X] CHEP 2015: “Experiences and challenges running CERN’s high capacity tape archive”: http://iopscience.iop.org/article/10.1088/1742-6596/664/4/042006
[X] CHEP 2016: “Tape SCSI monitoring and encryption at CERN” (being published)

5.1.1.3 The repository shall have effective mechanisms to detect bit corruption or loss.
Supporting Text
This is necessary in order to ensure that AIPs and metadata are uncorrupted or any data

losses are detected and fall within the tolerances established by repository policy (see 3.3.5).
Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement
Documents that specify bit error detection and correction mechanisms used; risk analysis; error reports; threat analysis; periodic analysis of the integrity of repository holdings.
Discussion
The objective is a comprehensive treatment of the sources of data loss and their real-world complexity. Any data or metadata that is (temporarily) lost should be recoverable from backups. Routine systematic failures must not be allowed to accumulate and cause data loss beyond the tolerances established by the repository policies. Mechanisms such as checksums (MD5 signatures) or digital signatures should be recognized for their effectiveness in detecting bit loss and incorporated into the overall approach of the repository for validating integrity.

The computing model adopted by the LHC experiments within the context of WLCG is such that all data collected by the experiments is persistently stored at CERN's Tier-0, and a copy is replicated and persistently stored across a number of collaborating Tier-1 sites. This allows for data redundancy and therefore higher reliability as experiments can recover data from a different site in case of data loss. Each site is responsible to fulfil the resources and service levels that are agreed in the form of a MoU [X].

Protection of Metadata: All of CASTOR's metadata is stored in Oracle databases operated by a dedicated team of database experts ([X]). All these databases have been setup with redundancy and recoverability in mind. Databases are backed up using a comprehensive, multi-level backup infrastructure following a well-established policy [X].

Detected data integrity issues are immediately handled as described in 3.3.5. In order to do a comprehensive check of the CERN archive contents, given that only around 20% of the data written to tape is ever read back by the users, a media verification process has been defined that regularly and proactively checks the integrity of all data in the archive, independently of user access. This process ensures that media cartridges can be mounted, and that the contained data can be read back and validated against metadata (such as checksums and file sizes). The process uses all the integrity checks described in 3.3.5. Tapes will typically be verified on a different drive than the one used for writing. Operators will be notified following standardized tape alert procedures (cf. 3.3.5) if drive or media warnings, inconsistencies or hard errors are detected.
Two verification modes exist, a "full" mode checking the complete tape data, and a fast "light" mode that concentrates on a data subset.
A "full" tape verification is triggered in the following cases:

- once a tape is completely filled with data, and
- whenever a tape hasn't been mounted for a long period of time, therefore ensuring the correctness of older (or "colder") repository data.

The "light" verification typically completes in 2-3 minutes as it concentrates on the most critical tape areas, by reading out the first and last ten files and another ten random files picked across the complete tape. This verification mode kicks in immediately after a tape has been written to, or if a tape hasn't been scanned for 60 days or more.
The verification system is documented under [X]. Operator procedures are found under [X].

[X] WLCG Memorandum of Understanding: http://wlcg.web.cern.ch/collaboration/mou

[X] Backup Service at CERN: http://cern.ch/service-backup
[X] Oracle Database Backup and Recovery Policy:
https://edms.cern.ch/ui/file/1450735/1/IT-DB_backup_rules.pdf
[X] tape verification system:
https://twiki.cern.ch/twiki/bin/view/TapeOperations/TapeVerification
[X] tape verification operator procedures:
https://twiki.cern.ch/twiki/bin/view/TapeOperations/TapeVerifyProblemMedia

5.1.1.3.1 The repository shall record and report to its administration all incidents of data corruption or loss, and steps shall be taken to repair/replace corrupt or lost data.
Supporting Text
This is necessary in order to ensure the repository administration is being kept informed of incidents and recovery actions, and to enable identification of sources of data corruption or loss.
Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement
Procedures related to reporting incidents to administrators; preservation metadata (e.g., PDI) records; comparison of error logs to reports to administration; escalation procedures related to data loss; tracking of sources of incidents; remediation actions taken to remove sources of incidents.
Discussion
Having effective mechanisms to detect bit corruption and loss within a repository system is critical but it is only the initial step of a larger process. In addition to recording, reporting, and repairing as soon as possible all violations of data integrity, these incidents and the recovery actions and their results must be reported to administrators and made available to all relevant staff. Given identification of the sources of data loss, an assessment of revisions to software and hardware systems, or operational procedures, or management policies is needed to minimize future risk of data loss.

All tape errors and alerts are registered in the CERN tape monitoring system (see 3.3.5); statistical and detailed information on cartridge, drive and media type reliability is generated and made available to service operators and management.
This information is also shared as appropriate with the respective technology vendors and used as input for technology replacement, upgrade and purchase decisions. Based on this information, CERN has repeatedly agreed with vendors on large-scale replacements of tape drives or media cartridges identified as problematic.
Over the years, CERN has defined a comprehensive workflow for dealing with media problems [X]. This workflow comprises several repair levels (from local operator to service manager to vendor handling). The workflow comprises steps for keeping users informed about the status of their data and of the recovery process (which may take hours to weeks). Instead of waiting for recovery, experiments having multiple data copies across sites can decide to erase the faulty files and replicate them from external Tier-1 sites.
Incidents affecting the tape service are in addition communicated to and discussed by management at CERN IT’s weekly service coordination meeting (C5, [X]); information relevant to the LHC user community is brought forward to the weekly WLCG experiment operation meetings [X].
The physical environment in which the CERN tape archive is located is also closely

monitored. In fact, today's tape bits are now smaller than most airborne dust particles or even bacteria. Therefore, tape media is now more sensitive to contamination from airborne dust particles. These can scratch the exposed magnetic substrate when mounted in the tape drive resulting in the loss of significant amounts of data. To mitigate this threat, CERN has prototyped and built custom environmental sensors [X] hosted in the production tape libraries that sample the same airflow as the surrounding drives. These sensors can measure bursts of tiny particles or the presence of larger particles, along with temperature, humidity and air flow. Whenever a problem is detected, a ticket is immediately created in the CERN problem management system (SNOW) and assigned to a service operator that will follow a predefined procedure.

All deployments of new CASTOR software releases are subject to comprehensive testing and staged deployment using CERN-IT's standard Agile Infrastructure toolset (based on Puppet). Comprehensive software and/or service reviews involving service representatives, users, and external experts can be triggered as a consequence of significant data loss incidents (example: ALICE data loss incident in 2010 [X]).

[X] Workflow for dealing with media problems: https://twiki.cern.ch/twiki/bin/view/TapeOperations/ProblematicTapeWorkflow
[X] http://cern.ch/c5
[X] https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings
[X] CERN DCES (Data Centre Environmental Sensor): http://www.ohwr.org/projects/dces-dtrhf-ser1ch-v1/wiki
[X] CASTOR external review, April 2010: http://indico.cern.ch/event/104191/overview

5.1.1.5 The repository shall have defined process for storage media and/or hardware change (e.g., refreshing, migration).

Supporting Text

This is necessary in order to ensure that data is not lost when either the media fail or the supporting hardware can no longer be used to access the data.

Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement

Documentation of migration processes; policies related to hardware support, maintenance, and replacement; documentation of hardware manufacturer's expected support life cycles; policies related to migration of records to alternate hardware systems.

Discussion

The repository should have estimates of the access speed and the quantity of information for each type of storage media. Then with estimates of the reliable lifetime of the storage media and information of system loading, etc., the repository can estimate the time required for storage media migration, or refreshing or copying between media without reformatting the bit stream. The repository can then set triggers for initiating the action at an appropriate time so the actions will be completed before data is lost. Copying large quantities of data can take a long time and can affect other system performance metrics. Repositories should also consider the obsolescence of any and all hardware components within the repository system as potential trigger events for migration. Increasingly, long-term, appropriate support for system hardware components is difficult to obtain, exposing repositories to risks and liabilities should they choose to continue to operate the hardware beyond the manufacturer or third-party support warranties. Repositories will likely need to perform media migration off of some types of media onto better supported media based on the estimated lifetime of hardware support rather than on the longer life expected from the media. It is important that *the process includes a check that the copying has happened correctly.*

One of the most important activities in CERN's data archive is the periodic migration (also called "repacking") of its data to newer storage technology generations, every 3-5 years. The capacity of tape cartridges has been growing exponentially over the last 20 years, and has the potential to continue at this rate for at least the next 10 years. By migrating existing data to newer-generation media, less cartridges will be required and expenses in additional tape libraries and floor space can be avoided.

Another reason is media reuse: Newer-generation enterprise-class tape drives often allow for the reformatting of existing media at higher density, allowing capacity gains of over 50%. Re-formatting existing media and recopying their contents therefore liberates additional space.

Technology evolution is a core element in CERN's archive resource planning. In order to avoiding obsolescence, CERN replaces the tape hardware infrastructure every 4-5 years, as after it may become difficult to find replacements for tape drives, firmware patches or software drivers for new operating system versions. In addition, new tape drive and media generations usually provide higher capacity, speed and increased reliability over older generations. CERN closely works with storage vendors in order to discuss technology roadmap previews (via NDA agreements) and get early access to new-generation storage hardware and media. We regularly participate to beta-testing of media and hardware in order to assess performance and functionality against CERN requirements and provide feedback to vendors. One example of such testing can be found under [X]. As another example, a vendor note regarding testing of new tape libraries by CERN can be found here: [X].

CERN also keeps regular exchanges with other large-scale archive sites and the world-wide tape user community. CERN has regularly been represented in the Large Tape User's Group (LTUG, [X]) board. CERN has also co-chaired, as well as contributed experiences and planning resources to the HEPiX bit-preservation Working Group [X].

All relevant information obtained from vendors, industry consortia (such as the INSIC roadmap [X]), peer institutes and tape user community, CERN's own analysis (for example [X]) etc. is used as input for CERN's archive resource planning and for preparation of tender and purchase orders.

CERN also investigates alternatives to tape for long-term storage, such as using commodity disk systems with software-provided redundancy such as erasure coding within the CERN EOS disk management system [X].

In order to minimise the risk of data loss due to media wear-out, we monitor the number of times a cartridge gets mounted over its lifetime, and decommission it after 5000 mounts (well below the typical lifespan of 20'000 mounts). In practice, we only get to decommission 3-4 cartridges / year as most of the tapes are well below that threshold.

All migrated data is subject to systematic verification as described in 5.1.1.3 as integral part of the media migration exercise.

CERN fully takes into account verification and migration costs (in terms of required time and resources) for its resource and capacity planning. In its large archive, the active data (influx of new data and files being read from the archive) represents only a small fraction of the total archive size. However, for a large-scale migration exercise, on top of the active data rates, the totality of the archived data will have to be first read, then written, and again read for verification. Given that tape drives can become a bottleneck as the required data rates

for the migration exercise may easily increase by orders of magnitude, CERN overcomes this problem by either procuring additional tape drives for the duration of the migration, and/or stretching the migration in time.
Details on the preparation and execution of recent media migration exercises at CERN can be found in [X] and [X].

[X] Oracle StorageTek T10000D beta-testing report: http://indico.cern.ch/event/282323/contributions/639674/attachments/519219/716344/T10000D-CERN-beta-test-summary.pdf
[X] https://www.spectralogic.com/2011/11/21/cern-evaluates-spectra-logic-tape-technology-to-store-large-hadron-collider-research-data/
[X] Large Tape User's Group (LTUG): http://www.ioug.org/p/cm/ld/fid=148&gid=632
[X] HEPiX bit-preservation Working Group: http://w3.hepix.org/bit-preservation/doku.php?id=bit-preservation:documentation
[X] Information Storage Industry Consortium (INSIC) roadmap: http://www.insic.org/news/2015%20roadmap/15_25roadmap.html
[X] CERN IT CTO: "Computing Evolution: Technology and Markets": https://indico.cern.ch/event/570249/contributions/2404412/attachments/1400426/2137004/2017-01-23-HSFWorkshop-TechnologyEvolution.pdf
[X] CHEP 2015: "Integrating New Storage Technologies into EOS": http://iopscience.iop.org/article/10.1088/1742-6596/664/4/042043/pdf
[X] CHEP 2013: "The Repack Challenge": http://iopscience.iop.org/article/10.1088/1742-6596/513/4/042028/pdf
[X] CHEP 2015: "Experiences and challenges running CERN's high capacity tape archive". http://iopscience.iop.org/article/10.1088/1742-6596/664/4/042006

5.1.2 The repository shall manage the number and location of copies of all digital objects.
Supporting Text
This is necessary in order to assert that the repository is providing an authentic copy of a particular digital object.
Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement
Random retrieval tests; validation of object existence for each registered location; validation of a registered location for each object on storage systems; provenance and fixity checking information; location register/log of digital objects compared to the expected number and location of copies of particular objects.
Discussion
A repository can have different preservation policies for different classes of objects, depending on factors such as the producer, the information type, or its value. Repositories may require a different number of copies for each class, or manage versions needed to meet access requirements. There may be additional identification requirements if the data integrity mechanisms use alternative copies to replace failed copies. The location of each digital object must be described such that the object can be located precisely, without ambiguity. The location can be an absolute physical location or a logical location within a storage media or a storage subsystem. Provenance information about copying and moving the data must be maintained/updated, including the identification of those responsible. This is

necessary in order to track chain of custody and assert that the repository is providing an authentic copy of a particular digital object. The repository must be able to distinguish between versions of objects or copies and identical copies. This is necessary in order that a repository can assert that it is providing an authentic copy of the correct version of an object.

CASTOR provides a metadata system which includes not only all file-level metadata, but also the exact location of each file on tape, as well as the status and exact slot location of each tape within each of the tape libraries. All tapes are always residing inside a tape library; no shelving is done at CERN. By default, a single copy of a file written to the CERN archive is stored on tape media. Secondary copies can be generated with a directory-level configuration option. Data copies of the same file will be stored in separate libraries by different vendors residing in different physical buildings. In case of a corrupt copy, CASTOR will discard it and recreate a copy of the correct file. Note that the LHC experiments do by themselves export and maintain physics data copies outside CERN; however, for critical data sets that require particularly high reliability they also store second copies in the CERN archive. Second data copies are also interesting for smaller (non-LHC) experiments that do not want to keep, or cannot afford to keep, off-site replicas of their data.

5.1.2.1 The repository shall have mechanisms in place to ensure any/multiple copies of digital objects are synchronized.

Supporting Text

This is necessary in order to ensure that multiple copies of a digital object remain identical, within a time established as acceptable by the repository, and that a copy can be used to replace a corrupted copy of the object.

Examples of Ways the Repository Can Demonstrate It Is Meeting This Requirement

Synchronization workflows; system analysis of how long it takes for copies to synchronize; procedures/documentation of synchronization processes.

Discussion

The disaster recovery plan should address what to do should a disaster and an update coincide. For example, if one copy of an object is altered and a disaster occurs while the second is being updated, there needs to be a mechanism to assure that the copy will be updated at the first available opportunity. The mechanisms to synchronize copies of digital objects should be able to detect bit corruption and validate fixity checks before synchronization is attempted.

Files stored in CASTOR are immutable, meaning that updating files is not supported. New versions of files can only be generated by prior deletion of the previous copy. The CASTOR software internally handles consistency of files between the disk cache and the tape archive. While no user-level “undelete” operation exists in CASTOR, service managers can attempt recovery after accidental/undesired deletion incidents, as by the “append-only” nature of tape technology, deleted data remains physically on tape (until media migration or decommissioning); and no overwriting does apply. The CASTOR metadata server logs all file operations therefore allowing for full reconstruction of metadata and therefore restoring the data itself.