



ARCHIVER Project

Technical Summary

PIC LARGE FILE REMOTE STORAGE

Problem Definition:

We want to substitute the current in-house tape library storage by an off-premise commercial service. Each instance of the service to be purchased consists in the 5-year safe-keeping of a yearly dataset from a single source according to the specifications below.

Lifecycle - Workflow Characteristics:

- A yearly dataset of 300 TB of new data is produced from a single source as a set of daily datasets over a one year period that must be stored for 5 years
- A daily dataset is 1 TB/day on average, with x2 fluctuations (from none to 2 TB/day)
- The source makes a daily dataset available all at once. Remote storage of a daily dataset must be completed within 8 hours of notice of its availability (250 GB/hour = 69 MB/s ≈ 0.5 Gbps)
- Source of data would initially be [PIC](#), possibly moving later to [ORM](#), Canary Islands
- Source is connected to the Spanish academic network (RedIRIS) and through it to Géant at a nominal speed of 10 Gbps
- A remotely stored yearly dataset is recalled as a whole (except for compliance verification, see below). Recall rate is low and usually takes place less than once per year.
- A prior notice could be given to the supplier prior to the start of a recall (to be negotiated)
- Recall must bring an entire yearly dataset back to disks at PIC in 30 days (10 TB/day=417 GB/hour=116 MB/s≈1 Gbps)

Authentication and Management Functions:

The producers and consumers of the data will be either human experts from PIC or automated processes running on PIC's servers (see Interface Characteristics below). A very limited number of identities is needed but they should be quite secure. Any reasonable industry-standard mechanism for authentication which is compatible with the Interface Characteristics can be considered.

Management functions: Web and Command Line interfaces should allow a manager to initialize and otherwise manage each instance and get basic information about the remote storage such as number of files stored, total volume stored, list of files stored, access to compliance evidence. Alarm and monitoring information can be sent by email. A desirable but not mandatory feature is to provide endpoints to PIC's Icinga alarm and monitoring system.

Data and Metadata Characteristics:

- Yearly dataset volume per instance: 300 TB
- Data are immutable once produced (Read Only)
- Data are not compressible (they are already compressed on origin)
- Data are organized in 2 GB binary files (one year≈150K files, one day≈750 files)
- Their content is in proprietary format (supplier doesn't need to look inside)
- Metadata: Filename: 256 character string, Checksum: 8 character Adler32 checksum
- Metadata are immutable once produced (Read Only)
- Neither data nor metadata contain personal or sensitive information
- Data are private (should only be accessible by authenticated authorized users)
- Data and metadata must be kept private, only accessible to authorized PIC servers

Interface Characteristics:

- put/get interface with full error recovery/correction using filename as id (see above)
- Programmatic interface to initiate a put/get operation and to check the status. Python bindings and REST-like API are desirable but not mandatory.
- Fully documented Command Line Interface to allow manual or bash script put/get
- Full documentation in order to allow us as customers to write an interface to emulate a Tape Library connected to PIC's Enstore or dCache software.
- Emulation as a Virtual Tape Library is desirable but not mandatory.

Reliability Requirements:

- A single-bit error anywhere within a file renders it completely useless
- Maximum tolerable loss over the 5-year storage period is one file out of the production of each year (≈150k files)

Compliance and Verification:

- Scrubbing: Supplier must provide evidence that every stored file has been read and its checksum freshly calculated at least once per year, without any intervention by the buyer. An alarm should be raised if the freshly calculated checksum does not match the original.
- Random 1% buyer's monthly check: Each month, the buyer will give a one week notification simulating a recall of a yearly dataset. The buyer will then ask to recall a randomly chosen 1% subset of the dataset, which must be accomplished within 24 hours.

Cost Requirements:

The total cost of an instance of the service must be less than 30k € total for the nominal data volume of 300 TB stored for 5 years. A desired cost structure would consist of a fixed fee plus a variable fee proportional to the volume of data actually stored.

Initial Data Management Plan:

DMP Topic	What needs to be addressed
Data description and collection or re-use of existing data	Raw, calibration and derived datasets will be handled in the ARCHIVER project. These data will be from scientific projects in which IFAE participates and which have explicitly given their consent. Initially, data will be from the MAGIC Telescopes https://magic.mpp.mpg.de/ . Data will be re-used in the context of multi-instrument gamma-ray astrophysics.
Documentation and data quality	Data documentation will be done by adapting the internal Data Management and Preservation Plans (DMPP) of the scientific projects who own the data. MAGIC has already given permission to use their DMPP plan. Data quality is measured continuously by the scientific instruments and data quality tags will be available as metadata.
Storage and backup during the research process	Today, PIC is using its dCache/Enstore based service to archive 'cold data' in a HSM like (disk + automated tape) service. That layer is going to be replaced/extended with ARCHIVER services using an additional copy of the data. (cleanup at the end of the project requires no further data movement).
Legal and ethical requirements, codes of conduct	Citation should be given when data is used by permission.
Data sharing and long-term preservation	Easy to use data sharing, both amongst access-controlled groups as well as Open Data, is a major objective of placing the data in the ARCHIVER-developed platforms. Long-term preservation will be tested by

	simulating in an accelerated manner the Data Management and Preservation cycle.
Data management responsibilities and resources	High-level data management will remain the responsibility of IFAE, while low-level data management and resources will be provided by the vendors.