ARCHIVER Project

Technical Summary

EMBL ON FIRE

Problem Definition:
FIRE is the ‘FIle REplication’ service, our current archiving system. It consists of two complete copies of the data, one is a distributed object store (across three of our own dedicated data centres), the other is a tape archive hosted at Hinxton.

The tape system is WORN (Write Once, Read Never), and is essentially a disaster-recovery mechanism only, not an actively used data store service. This tape system will be difficult to scale in the near future, so the problem is to find a way to host this second copy in the cloud, possibly distributed across multiple providers.

The cloud replica should be available for serving to users, rather than just a one-way archive that is never read. That said, given that it will be one of two complete replicas, there is some flexibility in having one replica colder - and therefore cheaper - than the other. So while our current tape archive is for disaster-recovery only, we would like its replacement to be actively usable, with some clearly defined SLA for access depending on the use of ‘cold’ storage tiers.

Lifecycle - Workflow Characteristics:
The initial upload will be of about 30 PB of data, with about 1 PB per month after that, growing over time. We note that the initial setup of an archive with this much data will be a lengthy process, not a brief event.

Authentication and Management Functions:
The archive will be managed by a service account, files in the archive will not be owned by regular users.

Users will authenticate through standard AAI techniques. Some of the data will be public, not requiring identification or authorisation for downloading - i.e. anonymous access must be allowed for that data. Other data will be strictly controlled, requiring that users belong to given groups of people who have been granted the right to access it. Users will not have the right to modify data, though some data may be versioned in the sense that a new file is uploaded and linked to the previous version.

Data and Metadata Characteristics:

Data consists of files from several kB up to 500 GB, with the vast majority in the range of 1 to a few GB. The current archive is 20 PB in size, and is currently doubling every two years. We do not control the sources of our data - any research team producing biological data anywhere in Europe, or beyond, can upload data to us - so we expect this growth doubling to continue for at least the next several years.

Some of the data will contain important PII, and we are under a legal obligation to protect it. We currently encrypt such data at rest, and the keys are strictly controlled. In particular, the people who manage the store should not need to have the decryption keys in order to perform their function.

Data may be uploaded from practically anywhere in the world. The majority comes from several sites around Europe and the USA. Anyone who produces genomic, proteomic, or other biological data can, in principle, upload it to EBI. The number of people doing this is going to increase over time, as noted above, we do not control the sources of our data.

Interface Characteristics:

The data store itself will be compatible with industry standards for large data stores - in practice this means S3 compatible. API access to the store itself is sufficient, we provide domain-specific search portals that return URIs to the user from which they can retrieve the data.

Web-based protocols will then allow direct download through the portal, or the user can take a list of URIs and hand them to a workflow manager for processing.

Users have no a-priori reason to care where the data is stored. Only if they need to process large volumes of data will they have cause to consider its location, so they can compute close to the data.

Reliability Requirements:

We have an obligation of custodial care for our data, i.e. we need to keep it ‘forever’. In practice, we expect a solution to target a lifetime of at least 10 years, with a suitable exit-strategy that allows us to extract our data within a defined time and cost envelope, should that be needed.

Compliance and Verification:

We do have legal requirements to protect our data. Currently we use encryption at-rest, and users who are granted access to restricted data are given separately-encrypted copies, so even if they get hold of other files they cannot decrypt them.

In practice, we expect to manage encryption keys outside of the physical archive, as we do today, so the archive will only contain opaque data. There will be relationship links ('file A and file B are equivalent, but not identical'), but that's all.

Cost Requirements:

The baseline to compare against would be a tape store of equivalent capacity, including the need to migrate to new media every few years. The solution would also allow us to return to operation after a disaster much faster than recovering data from tape. E.g instead of doing a bulk restore from tape, we could restore on-demand from the archive, which would be a vast improvement.

Initial Data Management Plan:

DMP Topic	What needs to be addressed
Data description and collection or re-use of existing data	The data to be used is primarily DNA sequence data from the European Nucleotide Archive (ENA), located in the EMBL-EBI data centres. Data is submitted to ENA by research teams around Europe.
Documentation and data quality	Data quality is measured as part of the DNA sequencing process, the quality measurement is included with the raw data. The origin of the data (organism, sample etc) forms part of the accompanying metadata collected during the submission process.
Storage and backup during the research process	The FIRE archive maintains primary responsibility for the safe storage of our data, using a distributed object store and a tape archive for redundant copies. The ARCHIVER project will explore the possibility of replacing the tape store with a cloud-based archive, but any decision to do that will not be taken until after the conclusion of the ARCHIVER project itself.
Legal and ethical requirements, codes of conduct	Data is encrypted before being uploaded into FIRE, if there is any requirement that it be protected. FIRE itself does not manage access keys, the data is opaque bits as far as it's concerned.
Data sharing and long-term preservation	Long term preservation remains the primary responsibility of FIRE, using the redundant

	copies in different technologies for safety. Data is made available via FTP, HTTP and other protocols.
Data management responsibilities and resources	Management of the original copy of the data will remain the responsibility of EMBL-EBI. We make redundant copies of the data on different storage technologies (tape and object store) to minimise risk of loss.