

Digital library storage in an iRODS data grid

Mark Hedges

King's College London, Centre for e-Research

mark.hedges@kcl.ac.uk

Digital library software provides a powerful and flexible infrastructure for managing and delivering complex digital resources and metadata. However, issues can arise in managing the very large, distributed data files that may constitute these resources. The Centre for e-Research, King's College London, is implementing an approach that combines the Fedora digital library software with a storage layer implemented as a data grid, using the iRODS middleware developed by the San Diego Supercomputer Center as the successor to SRB.

This approach allows us to use Fedora's flexible architecture to manage the structure of resources and to provide application-layer services to users. The grid-based storage layer provides efficient support for managing and processing the underlying distributed data objects, which may be very large (e.g. video). The Rule Engine built into iRODS is used to integrate complex workflows at the data level that need not be visible to users, e.g. digital preservation functionality.