

Digital Library Storage using iRODS Data Grids

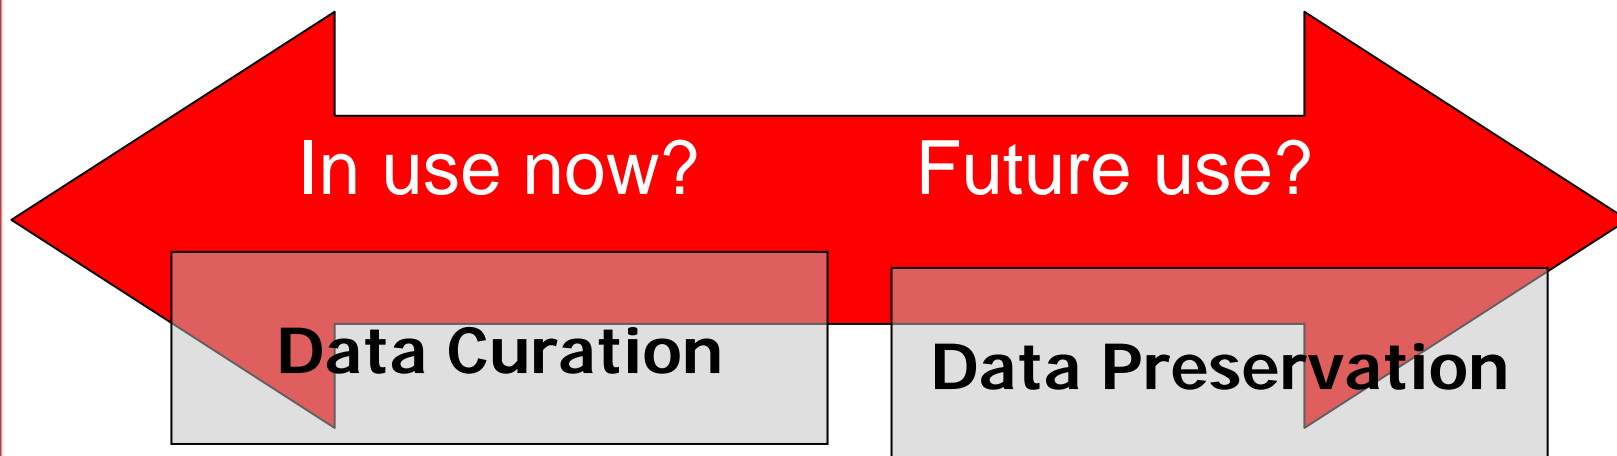
Mark Hedges, Tobias Blanke
Centre for e-Research, King's College London
Arts and Humanities Data Service
Arts and Humanities e-Science Support Centre

Adil Hasan
Rutherford Appleton Laboratory
Science and Technology Facilities Council

- Background: AHDS and Centre for e-Research
- Background: data deluge and broader data challenge
- Digital libraries and e-research infrastructures
- Digital libraries and data grids (SRB/iRODS)

- Arts and Humanities Data Service
- Established 1996, funded until 2008
- Distributed structure
- Mission: to collect, preserve and distribute digital resources produced by and for arts and humanities research (mainly in the UK)

- Centre for e-Research at King's College London
- Established 2007
- Incorporates staff and expertise of AHDS and other groups such as AHeSSC
- Continuity, but change of focus

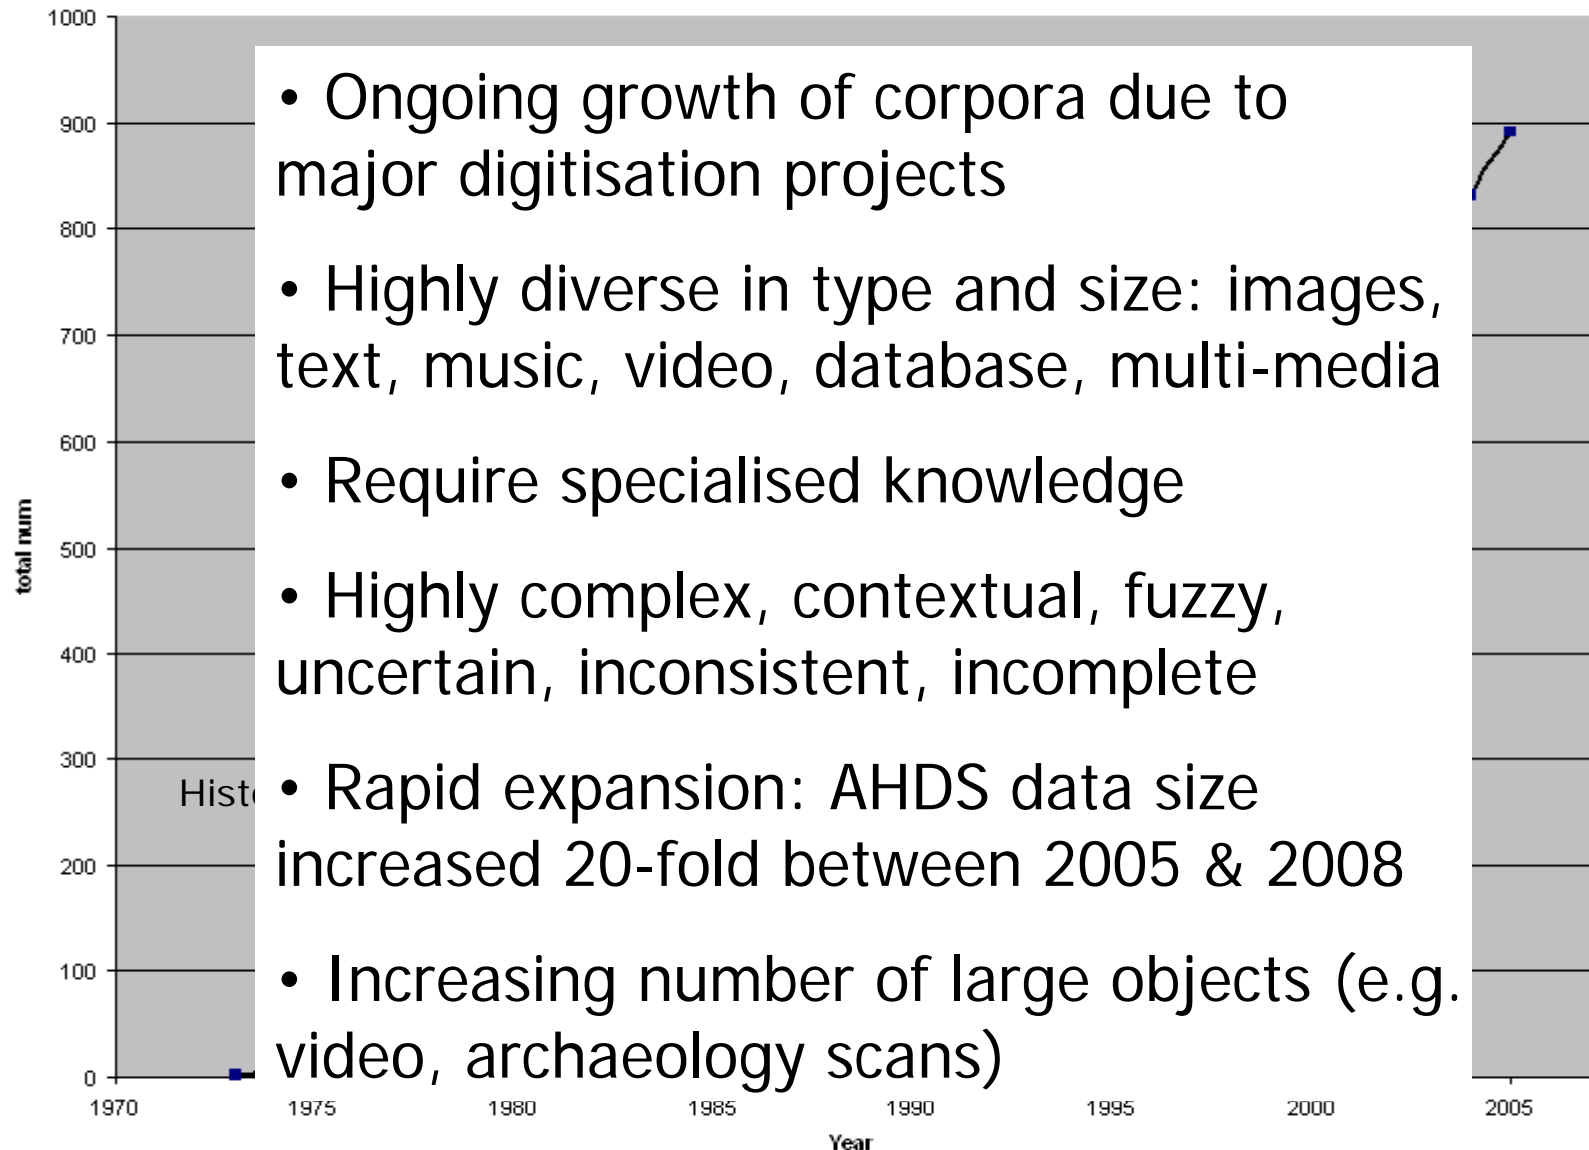


- **Curation:**

The activity of managing and promoting the use of data from its creation to ensure it is fit-for-purpose and remains available for discovery and re-use.

- **Preservation:**

An archiving activity in which data are maintained over time so they can still be accessed and understood through changes in technology



- Fedora Commons (at AHDS/CeRch)
- Supports digital resources that are diverse and structurally complex
- Flexible metadata management
- Disseminator framework supporting more complex and application specific processing of digital resources
- Not a stand-alone DL, but a component of an integrated research infrastructure

- Focuses on support for structure/ complexity rather than storage issues
- Doesn't natively support distribution of data
- Performance limitations when processing large objects

- Storage Resource Broker (SRB), a widely-used data grid technology developed by the San Diego Super Computer Center
- Addresses storage issues for digital repository and preservation environments
- Provides uniform, searchable access to virtualised, distributed resources, so DL is insulated from:
 - physical location of data
 - types of storage
 - migrating to new hardware
- Scalable – as library grows, new resources can be added dynamically
- Auditing facilities

- Not open source
- Not easy to exclude unwanted services
- Very effective for storage management, but not integrated with wider infrastructure.
- Not easy to integrate application-specific requirements (either change the core code, or implement in client, or use proxy commands)
- No built-in implementation of workflow (have to script this outside SRB, whether server or client side), or of asynchronous processing.
- Requires choreography between SRB admin and person running workflow.
- Relatively restricted support for metadata extension (Fedora supports but how to integrate)



- The open source successor to SRB
- Provides similar data virtualisation
- Rule Engine allows data management policies to be defined and realised as rules
- Policy virtualisation – insulation from how policies are implemented
- Execution of rules driven by events
- System level rules have great potential to ‘hide’ required data management operations from user/application level
- Event-condition-action model

- Rules (or policies) are sets of operations that you want to impose on an object (file, user, resource, etc).
 - The operations are called “micro-services”
 - Each micro-service is a C-app that executes and does something (e.g. checksum data, convert a file from one format to another).
 - Micro-services are transactional (recovery operations created for each micro-service).
- In most cases you can define server-side workflow as a rule controlling a set of {micro-services, rules}.

- Rule cast as {event: condition: action set: recovery set:}.
 - Can build rules of rules.
 - Allows you to model complex workflows.
- Supports execution of rules on most convenient resource (usually run on server connect to).
- Supports delayed execution of rules (i.e. “run this rule this evening”).
- Supports periodic execution of rules (i.e. “run this rule every evening”).

The components of a rule definition are as follows:

actionDef | *condition* | *workflowChain* | *recoveryChain*

Where:

- *actionDef* identifies the action to be carried out
- *condition* is necessary condition for execution
- *workflowChain* is sequence of actions to be executed
- *recoveryChain* is corresponding sequence of recovery actions (to ensure consistent state).

Rule can be built up cumulatively from other rules.

Data passed into/within rules (via parameters/context).

Note: syntax may change in near future.

Executed when an object has been ingested

acPostProcForPut | |

acCheckObjectIntegrity##

acAnalyseObject##

acNormaliseObject##

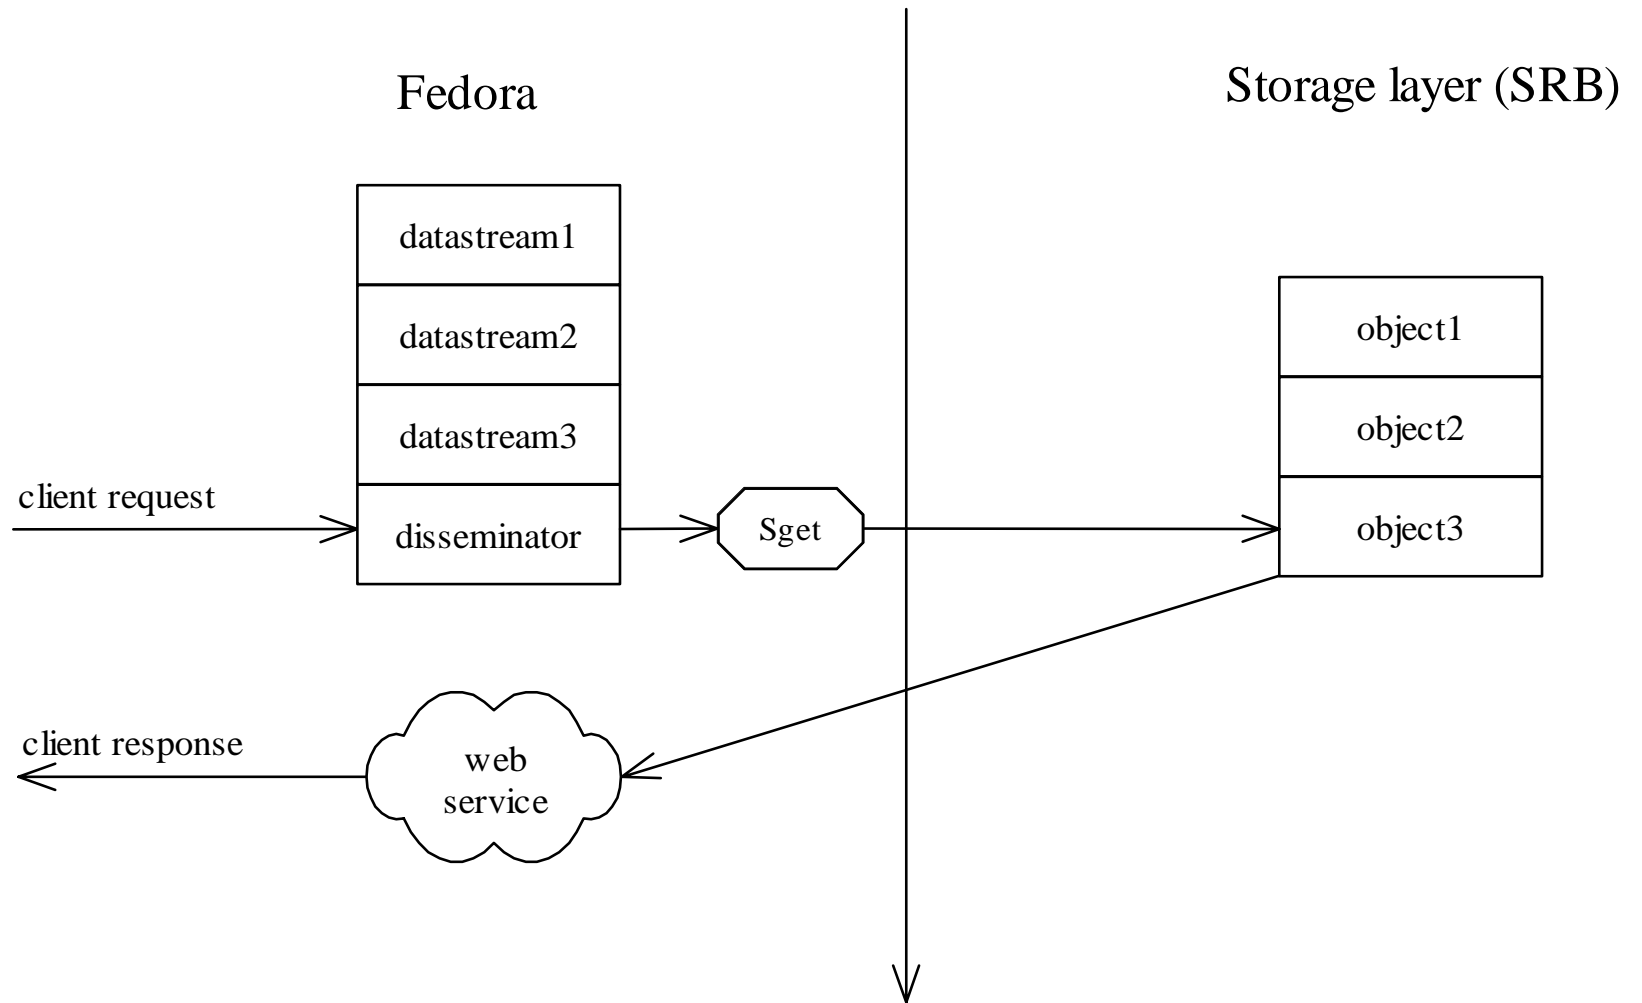
msiSysReplDataObj(PresRescGrp,all) |

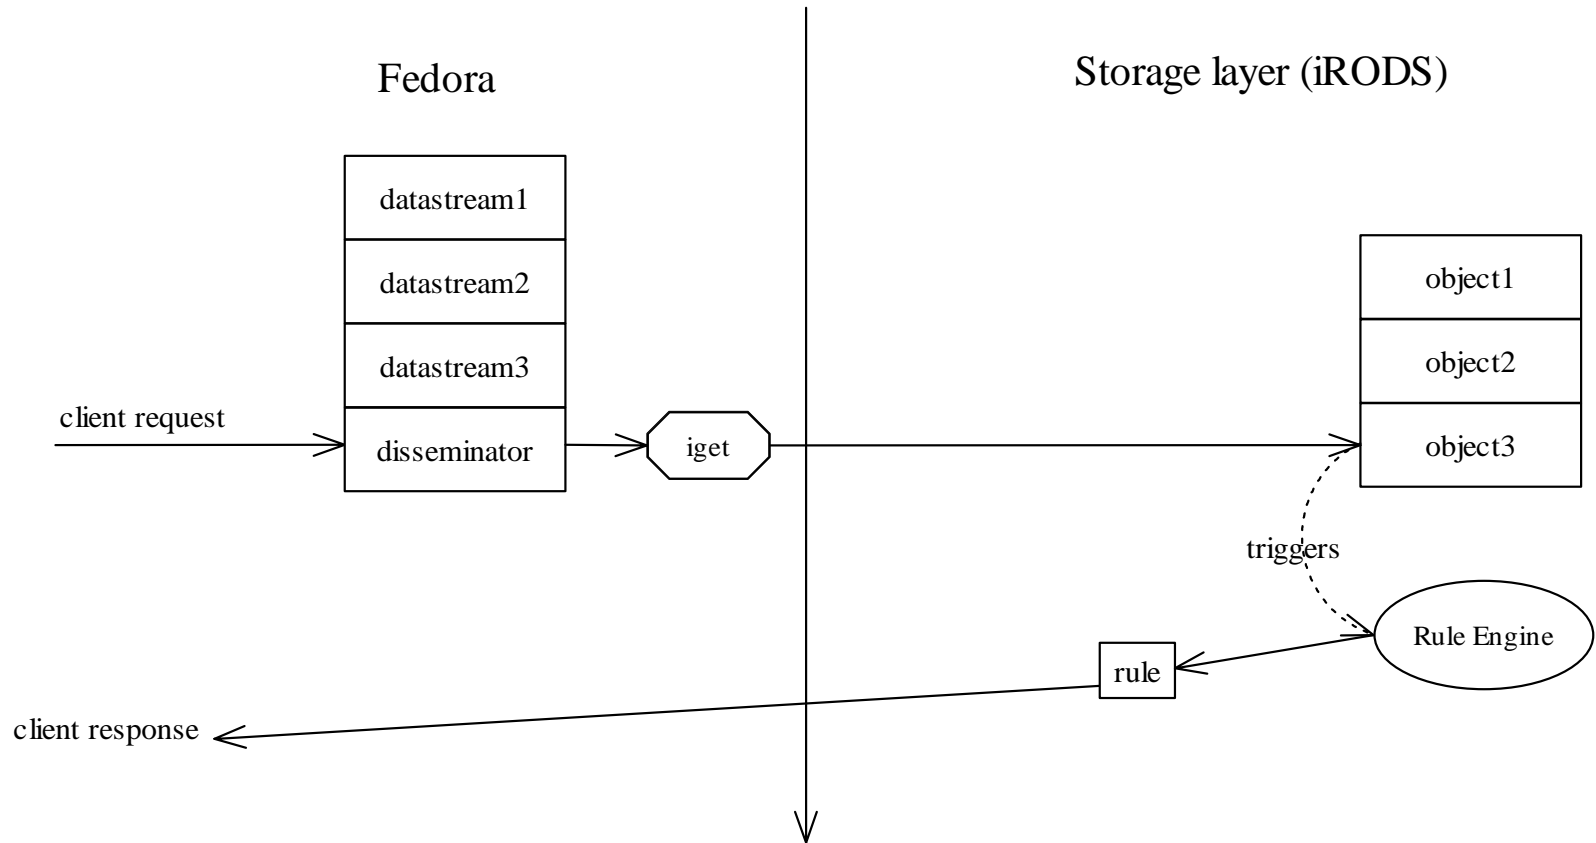
nop###nop###nop###msiCleanUpReplicas

Executed when an object has been ingested

```
acPostProcForPut |  
$format == "image/tiff" &&  
  $objectcategory="highResMS" |  
msiCheckForJPEG Tiling##  
  msiTiffToJPEG Tiling##  
  msiValidateTiffToJPEG Tiling |  
nop##msiCleanUpJPEG Tiling##  
  msiCleanUpJPEG Tiling
```


- Retrieving large objects for processing
- Retrieving entire object not always necessary, and can be inefficient
- Move the processing to the data
- Disseminators -> rules





- Prototypes -> production
- Developing more comprehensive set of rules for managing digital objects
- Jobs requiring data from multiple locations
- Dynamic deployment of jobs
- Virtual workspaces

mark.hedges at kcl.ac.uk
tobias.blanke at kcl.ac.uk
a.hasan at rl.ac.uk