

Amherst: async digital preservation to external storage

Title (Goal)	Amherst - async digital preservation to external storage
Primary Actor	Developer
Scope	Component
Level	
Author	Unknown User (acoburn)
Story (A paragraph or two describing what happens)	We would like to be able to store compound digital objects (e.g. a manuscript containing all original binary content, descriptive and structural metadata) in a container (e.g. bagit) and persist that container in some external storage location. That external storage could be a cloud-based system, such as AWS glacier.

Web Resource interaction

This service would interact with Fedora in two ways. First, it would react to Fedora's event stream (JMS or other). It would also expose its own HTTP endpoint to make it possible to regenerate compound digital objects for preservation.

Deployment or Implementation notes

This service would be deployed separately from fedora, probably on a separate machine. I envision that this would be implemented as a camel route that can be deployed in any OSGi container, written in Java and Blueprint XML. The implementation requires access to Fedora's HTTP API and event stream.

API-X Value Proposition

The primary use of this service would be for supporting asynchronous workflows/background workers.

In addition, API-X would allow for service discovery.