Amherst: binary derivative generation

| Title (Goal) | Amherst - binary derivative generation |
|---|---|
| Primary Actor | Developer |
| Scope | Component |
| Level | |
| Author | Unknown User (acoburn) |
| Story (A paragraph or two describing what happens) | For binary objects that are added to Fedora, we would like to have a service that will generate derivatives of various formats. An example could include extracting thumbnails from PDFs, MPEGs, or larger (TIFF, JP2) image formats. For many binary formats, multiple derivatives would be needed, e.g. a medium-sized image. These derivatives would be generated on-demand by invoking some appropriate external service. This should satisfy use cases that store the derivatives in Fedora as well as use cases that store the derivatives in external storage locations. That is, the service only concerns itself with converting Binary resources from A -> B, what is done with B is up to the implementation. Deployments of this service should be able to easily wire this service into an asynchronous message flow (e.g. via JMS events) or make the derivative generation part of an on-demand service offered to clients. This service should also be something that can easily be distributed across N servers. |

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 derivatives. For example, for resource /rest/path/to/binary, the service could be available at /rest /path/to/binary/svc:thumbnail or /rest/path/to/binary/svc:resize?w=200&h=300. The service needs to be aware of the MIMEtype of the resource and make appropriate endpoints and/or options available.

Deployment or Implementation notes

This service would be deployed separately from Fedora, probably on one or more separate machines. The current draft implementation runs as a combination OSGi service and camel route that can be deployed in any OSGi container. The implementation is currently 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 transformation of Binary resources.

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