Authenticated Access to Distributed Image Repositories

Robert Sanderson, Stanford University, azaroth@stanford.edu, http://orcid.org/0000-0003-4441-6852
Jon Stroop, Princeton University, jstroop@princeton.edu, http://orcid.org/0000-0002-0367-1243
Simeon Warner, Cornell University, simeon.warner@cornell.edu, http://orcid.org/0000-0002-7970-7855

Introduction

An increasing percentage of the world's cultural heritage is online and available in the form of digital images, served from open repositories hosted by memory, research and commercial organizations. These images are used for remote scholarship on objects that would otherwise be prohibitively expensive to travel to see, or sometimes too fragile to handle. However, access to the digital surrogates may be complicated by a number of factors: there may be paywalls that serve to sustain the host institution, copyright concerns, curatorial arrangements with donors, or other constraints that necessitate restrictions on access to high quality images. For many use cases open but degraded access is sufficient, and certainly better than nothing.

Images are also often the carrier for scientific and research information, particularly in the medical and biological domains. In many of these cases the images cannot be openly available because of personal privacy, and access must be restricted to health care providers or similar professionals. Archives also have this concern for individuals who may be still living but are mentioned in letters, photographs or other documents. In these cases, degraded access is not appropriate, and rights and authorizations need to be managed extremely carefully.

The International Image Interoperability Framework (IIIF)[1] has made great strides in bringing the world's image repositories together around a common technical framework. Now with its membership boasting nine national libraries, many top tier research institutions, national and international cultural heritage aggregators, plus commercial companies and other projects, use cases such as those above have raised authentication and authorization to the top of the “must-have” list of features to ensure continued rapid adoption. The release of new major revisions of both the IIIF Image[2] and Presentation[3] APIs in September of 2014[4] has set a solid framework, based on three years of experience and improvements, onto which additional services such as authentication can be built.

This presentation will focus on description of the IIIF authentication use cases and challenges, and then outline and demonstrate the proposed solution.

[1] http://iiif.io/
[2] http://iiif.io/api/image/2.0/
[3] http://iiif.io/api/presentation/2.0/
[4] http://iiif.io/news/2014/09/11/version-2-published/

Image Authentication Use Cases and Challenges

Authentication is always the topic that gets put off in standards or shared framework discussions, typically until late revisions when it suddenly becomes a barrier to adoption. Authentication for images intended to be re-used by other applications is particularly challenging compared to the general case of web authentication. The IIIF use cases are further complicated by the need to re-use images from multiple institutions in a single remote viewing environment; for example in virtually re-assembling a manuscript where the leaves are widely distributed. Additionally, there should be no need for prior arrangements beyond the implementation of the specifications.

Images are generally secondary resources in a web page or application. In the case of web pages, images are embedded in the HTML <img> tag, and are retrieved via additional HTTP requests. When a user cannot load a web page, it is possible—and a generally accepted behavior—to redirect the user to another page and offer the opportunity to authenticate. This is not an option for secondary resources such as images, and the user is instead simply presented with the much-hated broken image icon.

Authentication systems that span multiple domains are also complex, particularly with a Javascript client served from yet another domain, rather than from where the authentication challenge must be performed. This is the case for the majority of IIIF implementations because images of cultural heritage objects are usually hosted by the institution that owns the physical subject of the images. Details such as passive mixed content (the mixture of HTTP and HTTPS), cross origin resource sharing (enabling the ability to request data from different domains), and the desire not degrade the user experience with unnecessary authentication popups provide further challenges in this space.

Most users will access images via a browser-based client written in Javascript, and this makes some aspects easier while also enforcing particular constraints on the design of any solution. Notably, Javascript silently follows HTTP redirections so no information can be used from a response that redirects the user agent, and the HTTP headers from image requests cannot be retrieved (headers are available to Javascript code only for an AJAX XMLHttpRequest). So while simply putting an HTTP Link header on the image pointing to the authentication service might work in some environments, it would not be visible to browser-based applications.

Finally, the requirement to provide a degraded option[5], instead of simply yes-or-no access control, and serving the appropriate technical metadata (i.e. that which reflects the properties of the degraded version), complicates an already difficult situation. Both browsers and the web infrastructure have caches that will continue to serve the degraded content even after the correct credentials have been supplied unless the architecture of the authentication solution is designed with this in mind.

[5] Perhaps a grayscale version instead of color; a version of the image with a watermark; a version with more compression; or a smaller size.

Authentication for IIIF

Given the requirements and challenges described above, a system was designed to enable IIIF based systems to provide a coherent authentication protocol.

Due to the requirement for degraded images to be served, but not used from caches when the user is authorized for a better quality, the first design choice was to have separate URIs for the full and degraded images. In the terms of IIIF, the two have different image identifiers, but may have the same region, size, rotation, quality and format parameters. The second choice was to require a client to retrieve the image technical metadata, a short JSON description associated with the image (in IIIF terms this is the image information request or info.json), in order to determine whether authentication was required and if so where it could take place.

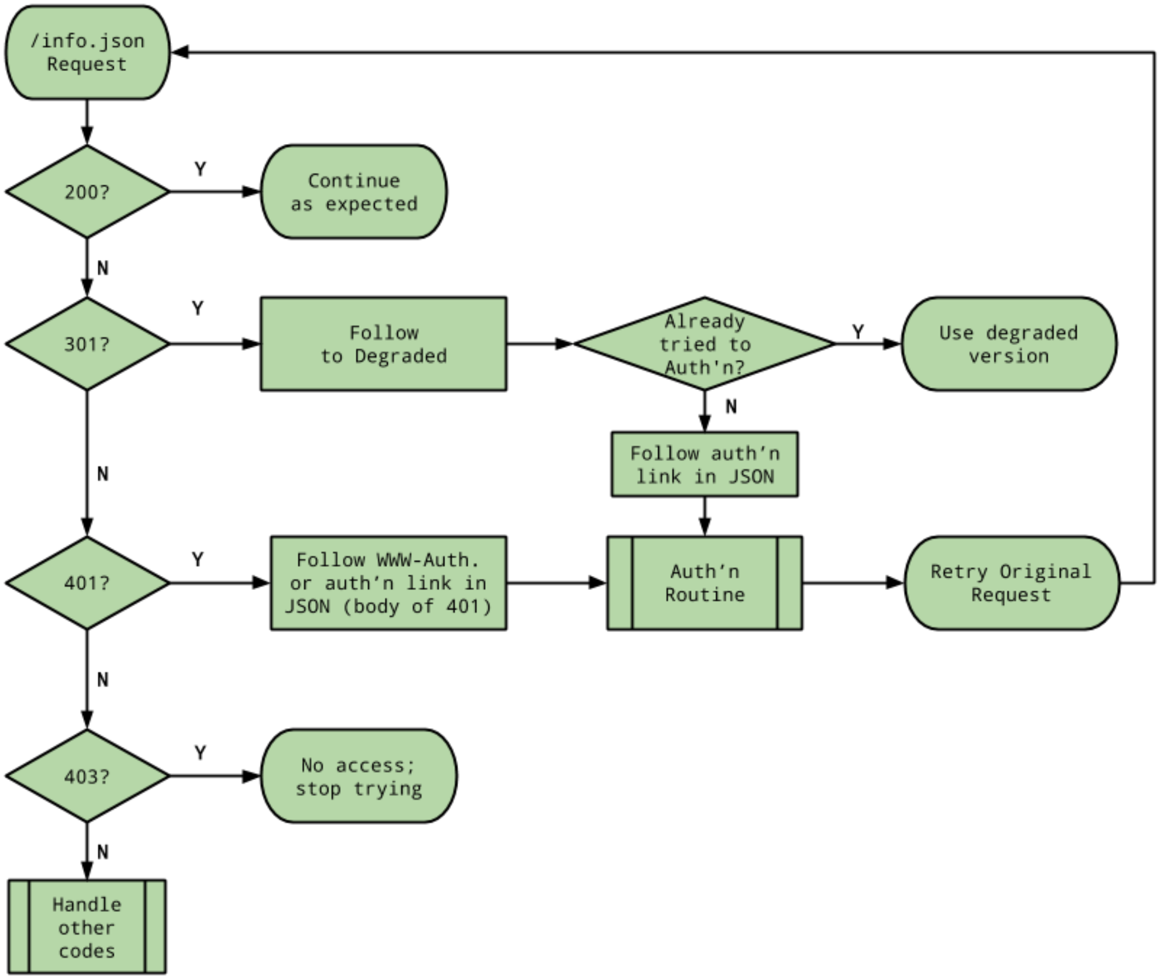

Figure 1. IIIF Authentication Flowchart (Client Perspective)

If the user is authorized to retrieve an image, then the server returns the info.json for the full image in response to the image information request, and the client can go ahead and use the image API. If authorization is required and there isn't a degraded image, then the response from the server is the typical 401 Unauthorized error with the info.json data as the body. This response is never cached, and so does not require a redirection and another round-trip. The response contains the link to the authentication service for the client to present to the user. Once authenticated, the client has all of the information needed to work with the image API so no further requests are needed.

Turning the silent redirections to our advantage, when an unauthorized request is received for the technical metadata, and a degraded image is available at the user's current level of authorization, the server redirects to an info.json document that describes what is currently available. That response contains a link to the authentication service to allow the user to obtain the needed authorization if she can. Once authenticated, the client re-requests the original image information to see if the user is now allowed to see the image. If so, then the JSON will be returned, and if not, then the server will once again redirect the client to a response with an authentication service link. The client should then just use the degraded image, if it hasn't already.

The actual mechanisms for authentication are not specified by IIIF, allowing any authentication system to be used. This allows for images to be protected by OAuth, local campus single sign-on systems, or simply basic web authentication. The only requirement is that authentication happen in another window or iframe and that this window close itself once the authentication has been completed, successfully or not. This is the signal to the client that it should re-request the image information (info.json) to determine any new options.

As part of the presentation at Open Repositories, we will demonstrate a version of the Mirador[6] viewer that can process this workflow and retrieve authenticated images from three different sites with different authentication systems. This will involve an open source implementation using the popular OpenSeadragon[7] client library, allowing other adopters to also make use of the work in other contexts.

Conclusion

While IIIF-compliant repositories already facilitate an unprecedented level of access to distributed image resources, the need to incorporate access controls reflects common institutional constraints. The authentication pattern described and implemented fulfills the needs of the IIIF community within the context of the web architecture. An IIIF client can “follow its nose” and find the best quality representations available to the user across different authentication domains. This distributed infrastructure enables effective re-use of repository resources within a growing suite of image applications.

[6] https://github.com/iiif/m2

[7] https://openseadragon.github.io/