Considerations and potential solutions for digital media stream:

Current needs and available resources for streaming
Audio and video content

Overview

Rutgers University Libraries is charged with delivering video and audio content for at least two specific missions and purposes. The first mission is for digitized and born-digital content which the Libraries intends to digitally preserve on the RUcore platform. The second mission is a curriculum support role, in which media in our collections must be made available to students and faculty as part of their research and coursework, but may not necessarily fall under a mandate for long-term digital preservation.

While in some cases, these two missions may converge (e.g. a video which is being preserved but available to students, and vice versa), there are a number of distinct requirements that each mission has that RUcore must meet. We currently have implemented a streaming solution for RUcore's preservation platform, but there are specific needs in the non-preservation content-delivery arena that may dictate a more simple approach. To that end, the technical group within the Media Task Force is charged with analyzing and selecting a solution that can effectively deliver content in a way that both missions are well served.

The document will analyze the needs as we know them for both missions, and list possible solutions that can be adopted to serve the unique needs of each use case.

Use case 1: RUcore streaming content delivery for academic and cultural heritage purposes

The Rutgers University Community Repository (RUcore) has already been actively engaged in digitally preserving and delivering streaming media. Our involvement and technological/software contributions with the New Jersey Video Portal (NJVid) have necessitated the development and adoption of standards and basic infrastructure for content delivery. In short, RUcore requires that a stable, high-resolution, and preferably uncompressed (or digitally native) copy of a video object be stored on our mass storage platform, accompanied with a suitably compressed and formatted presentation copy intended for online streaming by site visitors. Examples of such content include research videos such as the Robert B. Davis/Video Mosaic Collaborative collection, and various academic lectures and presentations conducted by faculty members and recorded on campus. Cultural heritage video and audio, such as interviews, recordings of historic events and oral histories, also fall under the purview of RUcore.

The current standards and delivery requirements for RUcore can be found in the video/moving images section of the RUcore standards documentation.

In short, some of the requirements for RUcore and preservation streaming platform include:

- The use of open technologies when possible to complement the open-source nature of the RUcore codebase
- Large amounts of storage to accommodate the long-term archiving of preservation-level, native born-digital or uncompressed digitized video sources, as well as the compressed presentation copies.

- Scalability and redundancy: a fault tolerant solution that can continue to serve content even during mild to moderate faults. Perhaps a load-balanced or redundant pair of similarly-equipped streaming servers.
- Access and control: Most cultural heritage videos are unrestricted, open content for which we have secured the rights to distribute publicly. However, some academic and research videos may need to restricted. This need may be handled by RUcore's Authentication and Authorization system using Shibboleth, and does not necessarily need to be a function managed directly by the streaming server.

RUcore's current streaming implementation consists of a single Apple Xserve running Quicktime Streaming Server, with 1TB of RAID1 mirrored storage. From a standards and content-delivery standpoint, this architecture has served us reasonably well. The system is also adaptable to future planned enhancements such as migration to the MP4 container format for video and mobile device streaming support. However, there are a number of issues which need to be addressed. In particular, only one such server is in operation, which does not afford us prompt fault recovery or redundancy should we encounter a serious hardware issue. Backup images of the operating system and streaming files do exist however, and restoration of services is possible within a few hours, depending on availability of repaired or replacement hardware.

A second major issue relates to the VMC and the RUanalytic Tool. Currently Flash-based, the tool requires use of the Real Time Messaging Protocol (RTMP), in order to successfully use streaming media. Our current versions of the streaming platform do not support this capability.

Use case 2: Streaming for general reserve holdings and curriculum support

Our second use case for content delivery involves videos that we do not own, but have obtained the rights to stream in the support of coursework and university curriculum. Typically, these are commercially produced videos for which Rutgers University Libraries have purchased a license to deliver to students. The restrictions on these videos can vary. For some, we may have perpetual rights to stream the content; in other cases, our permission to deliver these videos will lapse after a set period of time. The permitted audience may also vary: some videos may be viewable by all members of the university community, while others may only be viewed by students enrolled in only a very few courses at a time.

Aside from the rights restrictions, another attribute making this mission unique from RUcore videos, is the lack of a preservation stream. Because these videos are not owned or produced by the university, we do not typically have any preservation-grade digital copies associated with these videos, nor do we expect to store them in the same capacity as we would our RUcore repository.

For audio objects (typically reserves for the Music Library), a Helix streaming server is used with limited password protection capabilities. We do not currently have a formal solution in place to serve commercial videos locally, however the development/test server for RUcore has been pressed into service to provide streaming for the "Journey to Planet Earth" series, of which the School of Arts and Sciences (SAS) has purchased a three year streaming license.

Available streaming architectures

There are multiple solutions, of varying complexity and architecture, available to provide streaming services. The options range from proprietary standalone devices, to commercial or open source software packages that can be installed and run on server hardware. The following table lists some of the available options and the features and capabilities they provide.[1] [2] [3] [4]

[1] http://en.wikipedia.org/wiki/Comparison_of_streaming_media_systems
[2] http://www.haivision.com/products/furnace

Solution Name	Vendor/Developer	License	Cost (Software)	Media Types	Supported Formats	Operating Systems (Server)	Embed DRM Support
Adobe Flash Server	Adobe Systems	Proprietary	$4,500	Video	MP4, FLV	Windows, Linux	No
Darwin Streaming Server	Apple	APSL (Open Source)	Free	Audio Video	MOV, MP4, M4V, MP3	Windows, Linux, OS X, Solaris, BSD, other Unix	No
Helix Server	RealNetworks	RCSL/RPSL (Open Source)	Free	Audio Video	MOV, MP4, FLV, MP3, OGG	Windows, Linux, Solaris, BSD	No
IceCast	Xiph.org Foundation	GPL	Free	Audio	MP3, OGG	Windows, Linux, Solaris, BSD, other Unix	No
Video Furnace	Haivision	Proprietary	Rented, Variable cost, Subject to bandwidth usage	Video	Proprietary IPTV	Closed System (appliance)	Yes
Wowza Media Server	Wowza Media Systems	Proprietary	$999	Audio Video	MOV, MP4, M4V, MP3, FLV	Windows, Linux, OS X, Solaris, BSD, other Unix	Yes (Optional)

Additional Considerations

- MP4 (H.264) is a recommended RUcore standard for streaming video going forward. This is because it is supported on all major desktop operating systems (Linux, Windows, Mac), as well as on all major mobile device platforms and hardware (Android, Blackberry OS 5.x and later, iOS, Maemo, newer versions of S60 devices). Full implementation of MP4 support is slated in RUcore for version 5.2.2 in mid to late-2011.

- Legacy video objects on RUcore use a combination of Quicktime H.264 (MOV Container) and Flash video (FLV container).

- Providing mobile device support is a suggested requirement by members of the Cyber Infrastructure Working Group (CISC). The use of embedded Digital Rights Management (DRM) would severely curtail our abilities to provide cross platform mobile support, and has been known to cause significant end-user frustration,[5] while the effectiveness of the technology against unauthorized viewing is questionable.[6] A less intrusive LDAP or shibboleth-based front end authentication approach would be advisable.

[3] http://dss.macosforge.org/
[4] http://www.wowzamedia.com/
[5] Dr. Grace Agnew - Restrictions of Digital Rights Management - http://www.youtube.com/watch?v=aRfX2gPwXMo
[6] http://www.eff.org/issues/drm

Recommendations

While there are two disparate content delivery requirements that RUL must serve, it is still possible and logical to make use of similar streaming architecture solutions that are flexible enough to serve both needs. A reasonable approach would be to continue to build off the already-existing RUcore streaming architecture, as this system has already been effectively used to deliver video content on both a preservation platform (RUcore itself) and a primarily presentation and controlled-access focus (NJVid). This does not mean that a full FEDORA/WMS/RUcore implementation needs to be built-out. However, it stands to reason that a commercial/instructional video delivery platform can share elements of the storage platform (on a separate mount point and file system), and share the existing codebase that have been developed for video delivery modes

For all intents and purposes, this would mean that Use Case 1 is well underway to being effectively served. Use Case 2 can take advantage of the existing framework, perhaps even using methods already implemented by the existing NJVid architecture, for its needs.

Regardless of which streaming platform we select, some level of software development will be required to implement web-based delivery for commercial videos.

A possible implementation scenario

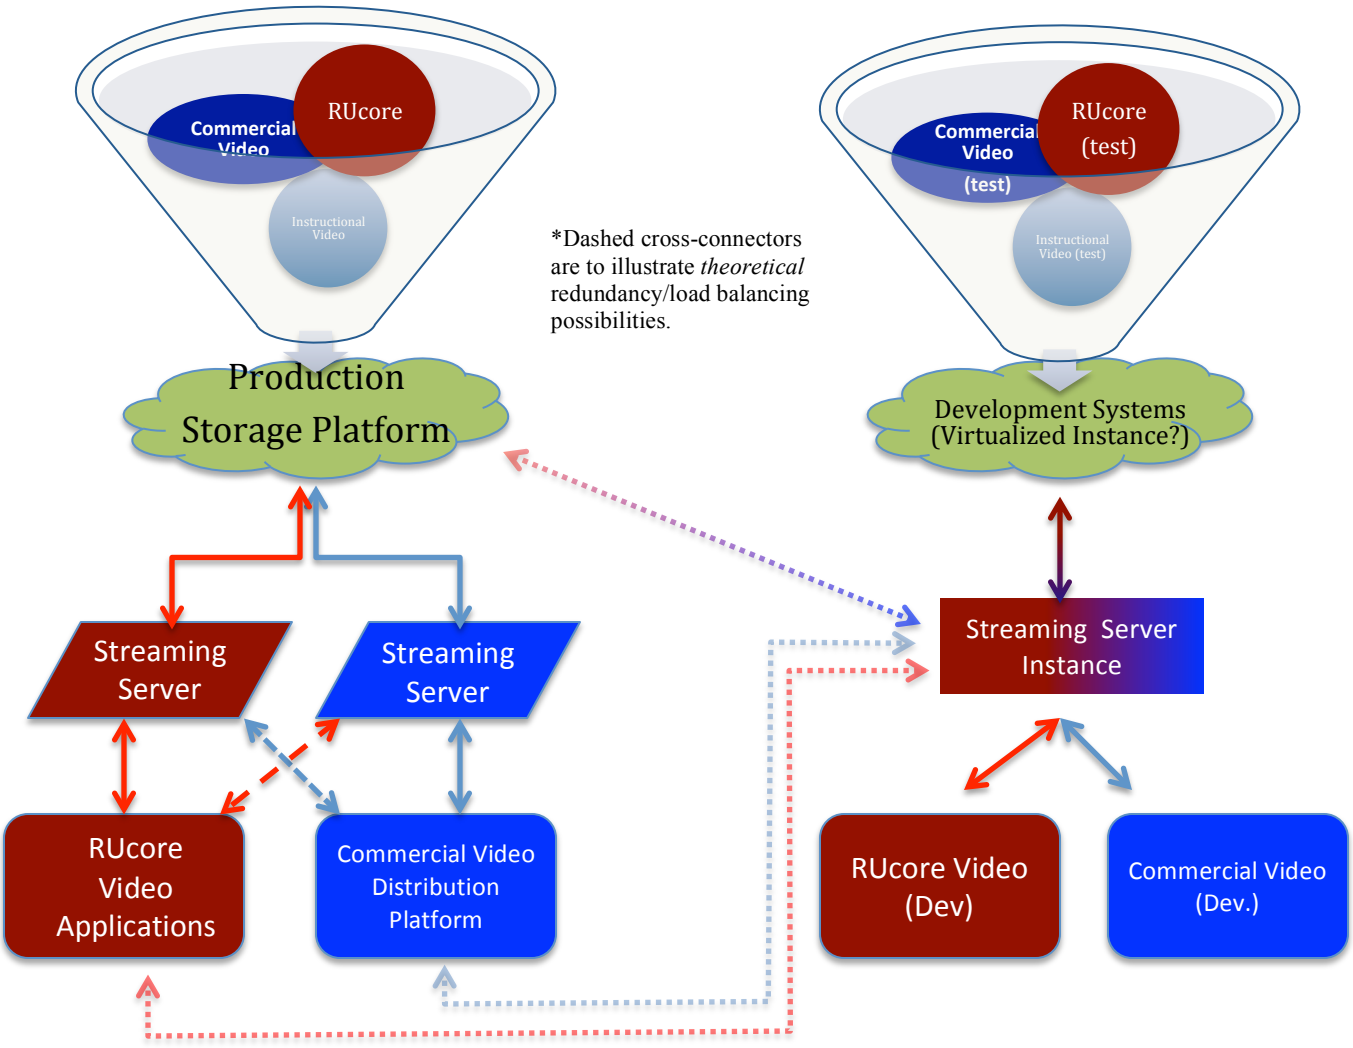

The scenario depicted above illustrates an ideal solution, where a streamlined hardware architecture is capable of serving both RUcore preservation and Commercial/Instructional video requirements. This would entail the use of a flexible, easily programmable streaming server architecture. In this configuration, a single streaming server with sufficient hardware, memory and network capacity could serve the functions of both applications, by accessing and making available separate mountpoints for each. A more robust and desirable solution, however, would be to have two or more servers operating in parallel for redundancy and load-balancing purposes. Each could be set as the primary streaming server for one application, with the ability for the software to change over to the other as a backup, as needed.

Additionally, a separate set of virtualized instances of these platforms could be created for development purposes. As testing and software development would not normally involve particularly high loads, a single development instance of the streaming platform could serve as a testing ground for both applications. The instance may also prove serviceable as a "last resort" backup, if given provisional access to the production storage platform in the unlikely event that the production servers are not operational, or are overloaded with traffic.

Hardware/Software Recommendations

- **Servers (Production)**
  - If not virtualized: Rack mount systems preferred.
  - Multiple CPUs required for multithreading (e.g. dual quad-core Xeons or better)
  - Significant RAM for buffering and file access (8GB or better)
  - Dual gigabit ethernet
    - 1 port dedicated for storage platform access (iSCSI)
    - 1 port for public streaming access and general internet connectivity
  - File System
    - Standard SATA hard drive storage configuration *or* standard virtual storage, sufficient for OS and streaming server software installation
    - *Either* access to storage space on the production storage platform, OR
    - 2-4TB of RAID 1 or RAID 5 for storage of streaming videos, expandable
- **Server (development)**
  - *Either* matches the same specs as above, OR
  - A virtual instance of a server matching the operating system and streaming software installed on the production system(s)
  - 1-4TB of storage for streaming, OS install, Streaming Server Software

- **Software**
  - **OS:** Linux recommended, though choices listed below will run on multiple operating system platforms.
  - **Streaming recommendation:** Wowza Streaming Server, production licenses on production machines ($1,000 each instance, perpetual), and development licenses on development platform (free, limited to 10 concurrent connections).
    - Supports streaming of the widest range of formats: MP4, MOV, FLV, MP3 audio.
    - Supports RTMP streaming (for use with integrated flash players), and encrypted authentication tokens to protect commercial content, with JWPLayer 5. Example: http://www.wowzamedia.com/forums/content.php?51
    - Mobile device support

- Should easily drop-in to the existing RUcore code base with minimal changes to code

Architectural (software) considerations

There is a significant desire not to have to re-invent the wheel for commercial videos. A large amount of effort on both the development and system administration side has gone into ensuring that key features that are useful for content delivery, authentication and access restrictions go into the RUcore software platform:

- A workable content model for streaming exists and is in the process of being further modernized
- Significant effort is being expended to make Shibboleth functional for authentication purposes
- Embargoes and access restriction capabilities exist and are already being implemented
- Cataloging capabilities. Although these videos are not part of our instutional objects, some level of cataloging and metadata entry will still need to be performed.

However, there is also a concern about what exactly goes into RUcore *per se*. That is to say that we do not want non-scholarly works to appear on RUcore scholarly search results and general portals.

That said, it seems like the best method for going forward would be to be make the most use of existing aspects of the repository software – the content delivery, the access control, the Workflow Management System, and the authentication – without actually making these commercial videos part of the institutional collections. It was decided at the September 28 Cyber Infrastructure Working Group meeting that this be achieved through a partition of the content through the use of separate collections (similar to the New Jersey Digital Highway being separate from, say, Faculty Deposits or Electronic Theses and Dissertations).

In any case, going with a separate solution will *still* require that we engage in significant development work to make web applications that tie in the needed requirements for access-controlled video content delivery.