summer conference **2008**

sparking innovative learning & creativity

Conference Proceedings

Rachel S. Smith, Editor

DAM If You Do!

BlueStream Digital Asset Management Infrastructure

Louis E. King | The University of Michigan

DAM If You Do!

BlueStream Digital Asset Management Infrastructure

Louis E. King | The University of Michigan

Digital media is emerging as an essential component of modern discourse (Lessig, 2008). Rapid expansion of the digital universe (Gantz, 2008; Weber, 2007) is driving a shift in media creation from select producers to the general public. Evidence of this abounds in large online forums such as YouTube, Flickr, FaceBook, WordPress, and Blogger, as well as secured collaborative environments – such as Sakai, Blackboard, Drupal, Joomla, and MediaWiki.

High social and technical barriers frequently deter a broader use of media across multiple venues and hinder the ability of individuals to engage in meaningful ways. These barriers are no more evident than they are in higher education where the complex social construct consistently raises concerns regarding privacy, copyright, intellectual property, students' rights, and patients' rights. Technical issues further aggravate the landscape by isolating pools of digital information in segregated technical "silos" with no manner of management, interoperability, or reuse possible.

The University of Michigan's BlueStream Project is developing enterprise infrastructure to lower these barriers, to foster innovation, and to increase productivity in media intensive learning, teaching, research and service. Robust media services streamline the construction and collection of rich digital assets and extend the digital repository into the numerous and varied working environments of the community. Researchers, instructors, learners, and staff spend less time building technical platforms and focus their attention on advancing academic work in innovative ways.

Rich Digital Assets

It is helpful to explore the BlueStream concept by beginning with the basic building block – the digital asset. Many systems consider the digital asset as a single digital file. BlueStream redefines the digital asset in a more comprehensive way, similar to that described in *A Framework of Guidance for Building Good Digital Collections* (NISO, 2005). It consists of the original digital file, subsequent versions, alternate file formats, data generated through media analysis, file level metadata, time-coded metadata, and fine-grained access control (See Figure 1: Anatomy of a Digital Asset). A digital asset is not solely the original file. It is a construct that correlates the original to its associated versions, derivatives, analysis data, and metadata.

Figure 1: Rich Digital Assets

This data model provides a mechanism for accumulating and managing every version and format of a file in a single place through a common management interface. It truly streamlines the workflow and reduces confusion that can arise from storing multiple versions of multiple files in multiple places. In addition, once centralized, it is much easier to build specialized services for analyzing and transcoding digital media. Automating these processes eliminates a major barrier by shifting "heavy lifting" from end-user desktop machines to enterprise servers. This results in large productivity gains in media intensive work.

Extensible Digital Media Repository

The power of this data model is not in productivity alone. It supports the collection, organization, search, and retrieval of assets in BlueStream's working repository. The term "working" is used to differentiate BlueStream from the Digital Library and the Institutional Repository. Unlike these other repositories, the contents in BlueStream are the every-day working materials of the academy that have not yet been vetted and chosen for long-term preservation. In fact, many of these assets may have extremely ephemeral life spans. If they should be chosen for inclusion in the permanent collections of the University there is a BlueStream workflow to handle that. Interestingly, the Institutional Repository uses BlueStream to create and host the "circulation" copies of digital video while maintaining Library stewardship over the preservation copies.

It is much easier to derive value out of a digital repository if it is embedded in the working environment. No one enjoys the cognitive disruption of changing tasks to hunt down a needed resource (Buchanan, Blandford, Thimbleby, Jones, 2004). BlueStream extends the repository into any of the working environments of the institution through web services architecture. This provides a quick and easy way to interoperate. For instance, a simple query or selection can call the materials needed for a particular web environment, pare them down to the optimal components desired, and remix them into an existing or customized environment (See Figure 2: Extensible Repository). Web services architecture provides a simple method for easily building effective interfaces to shape the academic experience (Suleman, Feng, Mhlongo, Omar, 2005). The combination of an extensible media repository and easy to build interfaces represents a tremendous opportunity for innovation.

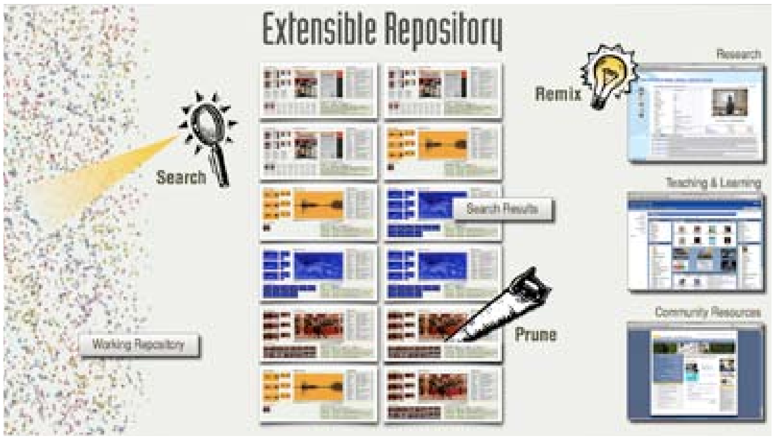

Figure 2: Extensible Repository

Bottom-Line Productivity

The academic environment is a sophisticated blend of face-to-face and technology mediated activities. The technology is ever changing and wide-ranging across physical facilities, communications networks, and the Internet. Customized environments are used to harness the power of human and information networks to support best practices in specialized endeavors.

The current practice in custom software development is for each application to invest in building a complete technology stack from the ground up. Every media-savvy application must develop its own software stack to store control, process, organize, and describe media objects. This translates to academic teams spending a great deal of time building technology at the expense of innovating (See Figure 3: Current Development Practice).

BlueStream provides a common digital media infrastructure for any application to build on. Developers can quickly and easily access these capabilities through web services and bring robust media handling to their own constituents. The whole community can coalesce around a common infrastructure to fulfill a common need and do so by applying sufficient resources to this solution rather than spreading resources thinly across numerous, weak, and partial attempts.

Scaffolding academic activities through a common media infrastructure drives technical development and academic production efficiencies that deliver sorely needed bottom line results in productivity to higher education.

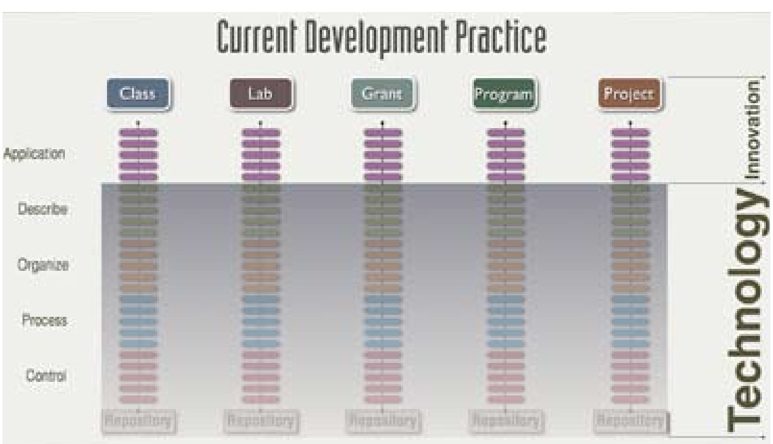

Figure 3: Current Development Practice

Investment Assurance

The shift to "born digital" academic materials coupled with increasing productivity is swelling the collection of digital working materials to unprecedented levels with no end in sight. These working documents, rough as some may be, are evidence of a prodigious individual and institutional investment. Created, vetted, and used daily, by everyone in the academy, this is where "the rubber hits the road." As these digital materials amass within the institution, and they are, it sets an imperative to consider what needs to be done to protect this investment.

A number of factors make this difficult. 1) The academic work itself, distributed across thousands of classes, research teams, projects, institutes, museums, and libraries is enough to confound and confuse any effort to organize the digital workspace. 2) This environment has spawned hundreds and hundreds of computer systems each having independent software, hardware, and networking lifecycles to manage. 3) Academic technology is managed by a corps of experts loosely knit into a structure akin to the academic units they serve. 4) Academics and technical staff have sufficient transience to disrupt institutional memory, leaving orphaned works and inoperable systems behind.

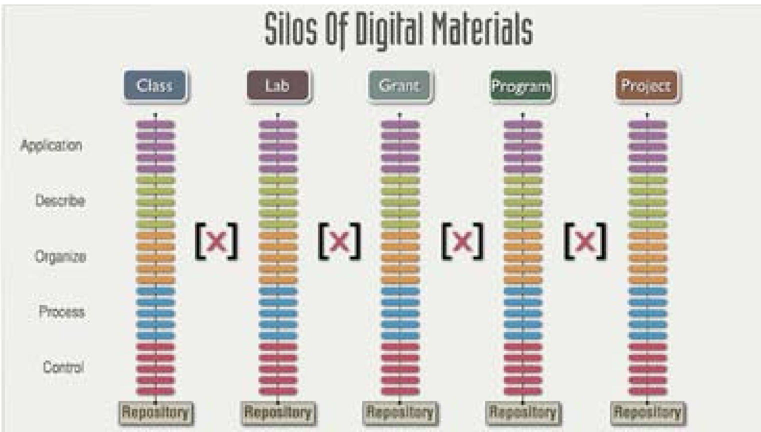

Figure 4: Silos Of Digital Materials

It is not surprising that these conditions have led to digital materials being scattered willy-nilly across the academy. The distributed nature of the institution and specialized nature of the work has naturally spawned a distributed systems development approach resulting in thousands of independent systems. Each system represents an isolated silo of potentially valuable digital materials nearly absent of institutional management, interoperability, or reuse (See Figure 4: Silos Of Digital Materials).

Digital materials used in the daily practice of the academy are at risk!

Ecology of Innovation

The BlueStream infrastructure supports a paradigm shift in which academic teams can build discipline-specific applications atop a sophisticated and robust digital media infrastructure. They do not need to spend time and money building and rebuilding technology stacks to support digital materials. The infrastructure, akin to plumbing for digital media, will handle that (See Figure 5: Digital Media Infrastructure). As with plumbing, it is not necessary to understand the minutiae of the thousands of components that make this happen. After all, who really cares? As long as the media flows when the tap is opened, not a soul.

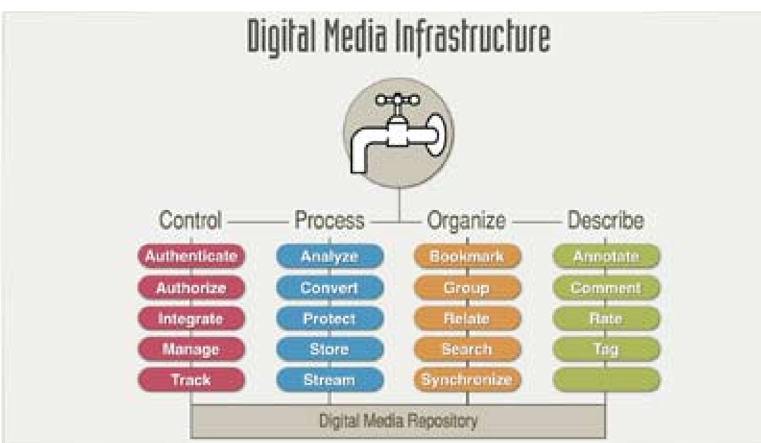

Figure 5: Digital Media Infrastructure

Freed from the burden of building technology from the ground up, freed from materials trapped in silos, freed from the cumbersome characteristics of media files, and armed with the power of new infrastructure, academic teams can turn their attention to shaping experiences that achieve

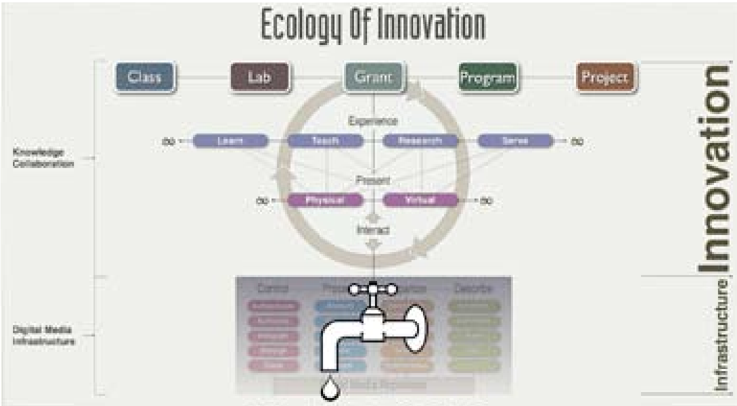

Figure 6: Ecology Of Innovation

their academic goals. In turn, these goals continue to drive the ongoing development of infrastructure. Nardi and O'Day (1999) describe this iterative process as a fundamental component in what they call the Information Ecology. The BlueStream project has embraced this concept and has formalized processes to foster these iterations and in so doing, support ecology of innovation (See Figure 6: Ecology Of Innovation).

Works In Progress

The BlueStream project accepts 25 to 50 applications per year for project space and support. Members of the university community involved in driving or supporting academic outcomes in research, teaching, learning, and service may apply. A portion of the funding line from the Provost's office is set aside to support smaller projects. This streamlines the process of rapidly producing results for known and contained media-intensive activities at the grassroots level. Larger projects are produced in partnership with academic units and are financed through a combination of unit level and BlueStream funds.

In addition to funding, participating units are required to assign academic technology support within the unit to the project. This "BlueStream Affiliate" is a key player in the train-the-trainer model. Affiliates provide a technologically conversant liaison to local culture, a conduit for units to influence the specification of BlueStream, and support for local knowledge exchange that spreads technical and media literacy via a sustainable model.

The following projects provide a glimpse into the type of work BlueStream enables. Take note of the depth and breadth of just these projects and keep in mind that there are currently over 35 active projects like these. In the long run the architecture currently in place can scale to a comprehensive enterprise infrastructure for the entire University of Michigan community.

Projects

Latin Tinge

Latin American Culture – Literature, Science, and the Arts

Professor Jesse Hoffnung-Garskof has always used media in teaching his courses in Latin-American culture, albeit with a considerable amount of fumbling over disks, tapes, and equipment.

He saw the BlueStream digital media infrastructure as an opportunity to do more than improve his delivery of lectures. It was an opportunity to transform the learning experience, to attract a greater number of students to Latin-American studies, and to feed the master's level program in his department.

Figure 7: Latin Tinge Course Website

The new Latin Tinge course has a corpus of over five hundred video and audio artifacts, carefully chosen to support student exploration of culture through sequences of musical forms – rumba, tango, samba, son, mambo, and salsa. The course web site provides a structured, week-by-week, presentation of the materials within cultural contexts presented by the professor. Students exchange their thoughts on required listening through blog postings. Students can navigate freely through the collection to synthesize understanding and find artifacts that support their own viewpoints. Students are required to find interesting and relevant cultural artifacts, add them to the collection, and tag them appropriately. Collectively, they add about 2,000 items per assignment and use them in conjunction in response to writing assignments.

Figure 8: Latin Tinge Student Submissions

High quality lecture materials are downloaded from the repository in advance of the class (no fumbling required). More importantly, much of the listening is done outside of the classroom. In its stead, drum and dance experts are invited in and the whole class learns the steps and beats of Latin American culture in much the same way that the heritage has been passed down from one generation to the next.

Learning objectives are being met and the class is very popular. The learning environment and teaching methods have scaled nicely to support a class size of 150 students. Building upon the BlueStream infrastructure, Professor Hoffnung-Garskof has transformed his undergraduate class into a signature course in Latin Culture that attracts large enrollments, prepares students for critical thinking in advanced study, and serves as a gateway to the graduate program.

Standardized Patient Instruction – Reflective Learning
School of Dentistry

The School of Dentistry has an eight-year record of proven efficacy in improving health care provider communication skills through a Standardized Patient Instructor (SPI) program. SPIs are utilized to simulate patient interactions for students. Typically, students are given an exercise to prepare for the simulation, go through the experience with the SPI, complete a self-assessment, and discuss the assessment with the SPI immediately after the simulation.

A pilot project was established to record video of these exercises and see if the recorded materials could advance the learning experience further. Dozens of exercises were taped. Unfortunately, the usual barriers to using media – recording, storing, securing, processing, and retrieving – thwarted the use of all but a handful of "best-case" examples.

This semester Professor Mark Fitzgerald will takes advantage of the BlueStream infrastructure to support the video component. Students will prepare for their simulation online by watching staged examples, will identify different segments that demonstrate both good and poor practice, and will write justifications for their conclusions. Hour-long simulations for 50 third-year dental students and 20 SPIs will be recorded, automatically uploaded to BlueStream, tagged with metadata, analyzed for voice-to-text indexing, secured with access control, and made available through the class web site. Students will review their own simulations, select clips that demonstrate when they were effective and when less effective, and write explanations of their conclusions. In turn, SPIs provide feedback on these assessments.

Figure 9: SPI Simulations

Figure 10: SPI Online Clip-making

Professor Fitzgerald summarizes the advantages to this approach as follows:

- Improved self-assessment skills of students through review of past performance and planning changes for future performance
- Improved performance skills by facilitating "reflective practice" prior to SPI interaction sessions
- Documented "competency" in various communication skills as a required step toward graduation
- Improved effectiveness of SPI feedback by allowing indexing and annotation of specific examples of effective or ineffective performance by students in simulated and real interactions
- Enhanced student e-portfolios through inclusion of videos demonstrating student personal communication skills and their associated reflections.

BlueStream makes it possible to produce this project of substantial scale and scope with very little lead-time. The only considerations beyond the infrastructure are metadata schemas, workflow, customized interfaces to support the learning method, and training and support for the community. In the end, the course assessment will determine if the learning objectives have been met. In the meantime, BlueStream affords another opportunity to explore compelling and innovative methods.

Grant Opportunities [Collaborative Spaces]
Digital Media Commons – Office of the Provost

GROCS is a program to fund interdisciplinary student research in the use of digital media in collaborative learning. Academic outcomes and collaboration are a significant and integral part of each exploration. Four to five teams are funded per year. Students receive a cash award, equipment, and a flexible learning environment to assist them in meeting their rigorous objectives within a single semester.

One of the most frequently noted benefits of sharing a learning lab is the serendipitous crossing of paths that bring people together who ordinarily would not meet. These connections enhance the collaborative nature of the physical environment and build cross-team communications. Careful observation has revealed that many of these connections are made outside of planned programmatic activities. In fact, it is the evidence of work left behind in the form of artifacts that attract interest and spawn connections. In this environment the artifacts might be just about anything –white boards, screen presentations, models, artworks, photos, costumes, circuit boards, toys, and specialized equipment.

Student research continues all week long although time spent in the lab could be fleeting. Therefore, it is essential that collaboration and discovery fluidly traverse the barriers of the physical learning space and the online learning environment without interruption.

BlueStream has been instrumental in easing the traditional difficulties encountered in such circumstances. Students are able to rapidly drag and drop video, images, audio, and texts into the digital repository. A minimum of metadata is automatically added, access control is automatically applied, and multiple formats are automatically made. Since the working repository is metadata driven, students use search terms to locate their materials. A search for a term, as opposed to a team name, is likely to produce a results list with artifacts from other teams that have been tagged in similar ways. The fuzziness of the search, in some way mimics the serendipity of paths crossing in the physical environment.

Figure 11: GROCS Blog

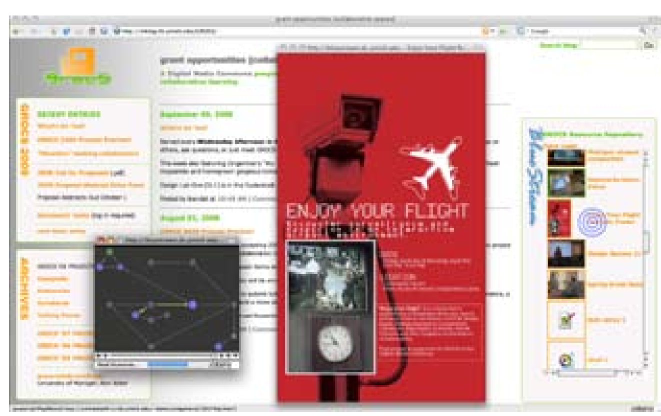

Figure 12: GROCS Blog

BlueStream also strengthens team collaboration and communication by extending the digital repository into the team and program blogs. The University Library hosts these blogs but does not support the large media files typically used by this community. Teams need only change the access control flag to public view and the media will automatically appear in a media tray on the blog. Media hyperlinks are not used in the blog. Instead, search terms are shown in bold and when executed return the specific file as well as files with similar tags from other groups. In this way, the blog perpetuates the collaborative nature of these investigations through their public-facing communications.

Secondary Master of Arts with Certification (MAC)
School of Education

It is no longer sufficient to prepare students of education to be excellent teachers! Simply put, it isn't possible to train and place enough high-quality teachers to make sufficient improvement. The University of Michigan School of Education addresses this concern by developing students to be skilled, critically reflective, and inquiring educators capable of becoming change agents wherever they are placed.

One proven method for developing all of these skills revolves around the synthesis of understanding through "records of practice." Records of practice include a wide range of materials that are central to classroom teaching. These include traditional records such as attendance, grades, and assignments, and non-traditional documentation such as video of lessons, recordings of student-teacher interviews, photos of board work, bulletin boards, posters, and project work, scans of student homework, tests, written work, teachers' lesson plans, and teachers' notes.

The Secondary MAC program requires all 50 students to carefully collect, share, and study records of practice from observations of a mentor teacher as well as from their own teaching. BlueStream supports the great variety of media types involved, streamlines the collection process, and drives collaborative use of the class-wide collection of records. The digital media collection focuses classroom discussion, supports small group work, and provides evidence in the form of rich media to support arguments made in student writing assignments.

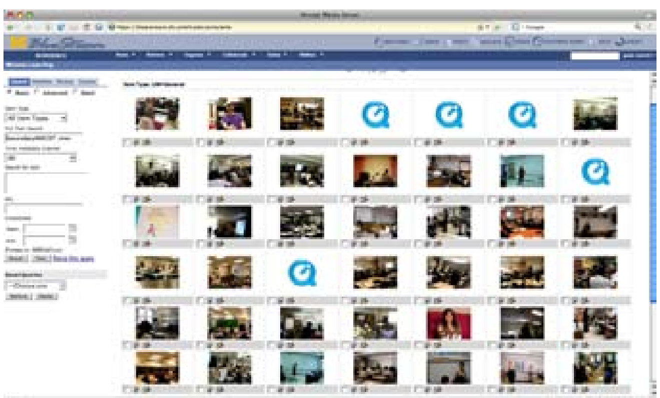

Figure 13: Secondary MAC Collections

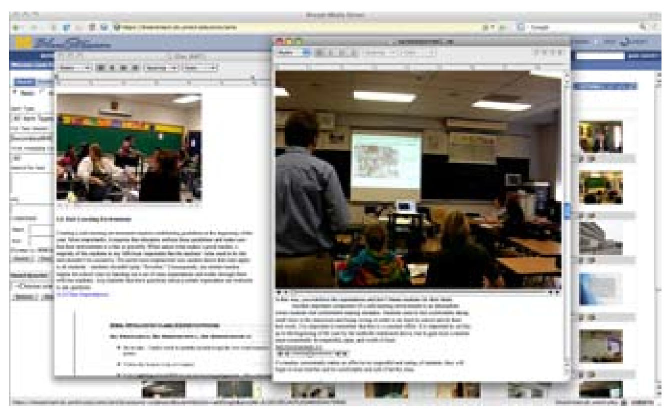

Figure 14: Secondary MAC Multimedia Writing Assignment

Many of the programs in the School of Education are adopting the "Records of Practice" method. The activities are engaging, the time spent on technology is minimal, and the learning outcomes are fantastic. Ordinarily, success of this kind is difficult to scale to ever-growing audiences. In this case, building atop the BlueStream infrastructure, scaling exponentially will not be a problem.

Conclusion

These projects demonstrate that making the digital repository present in both physical and virtual environments opens the door to improved academic discourse and collaboration. Individuals are able to seamlessly move from one social construct to the next, whether in-person, online, or by cell phone, and always have essential materials available.

It is also the case that the repository itself becomes an extensible dataset that can be combined with data from other sources to build more sophisticated collections of information. These "data mash-ups" give knowledge workers the opportunity to synthesize data in new ways and from different perspectives that can lead to greater understanding and new knowledge.

Results from the BlueStream project suggest that enterprise digital media infrastructure provides essential technical scaffolding to media-intensive academic work. It increases productivity of faculty, students, and staff. It improves academic discourse. It advances data collection and synthesis. It secures the community investment in digital materials. Overall, BlueStream provides a tremendous opportunity for innovation in media-intensive academic work in the higher-education information ecology.

Appendix 1: Technical Overview

BlueStream provides the organizational structure and facilities needed to support media intensive work in the higher education institution. An in-depth look into the full suite of technologies would lead to a technical tome. Instead, this section offers a glimpse into key components of the software, hardware, and architecture.

The BlueStream infrastructure is a modular system of components. One of the core architectural strategies is to recognize that each component has a lifecycle. Over time, each will need to be replaced by an alternate solution. This approach increases the complexity of dependencies between systems making it somewhat more difficult to upgrade than single vendor solutions. However, in the long run, the open architecture provides a mechanism for swapping components and integrating new capabilities to maintain "best of breed" quality across the entire system. Modularity across multiple vendors mitigates the risk of volatile purchases and acquisitions in the technical sector that can have a severe negative impact. Since the practice of digital media-intensive work is still in its early stages, it is also important to keep doors open to the potential of emerging technologies and the need to integrate additional components.

BlueStream Features: Highlights

Rapid Acquisition

Media files are large and unruly. One hour of video in standard DV format takes 13 GB of storage. Difficulty in uploading files to an asset management system is a showstopper for most end-users. BlueStream provides streamlined upload (See Figure 15) with drag-and-drop-from-the-desktop ease as well as web upload when that is more convenient. End users need not think about the best place to store the file and all of its derivatives. Based on the user identity, file type, and pre-established policy, all the files and related data are automatically stored in the optimal storage systems within the managed storage environment (See Figure 16).

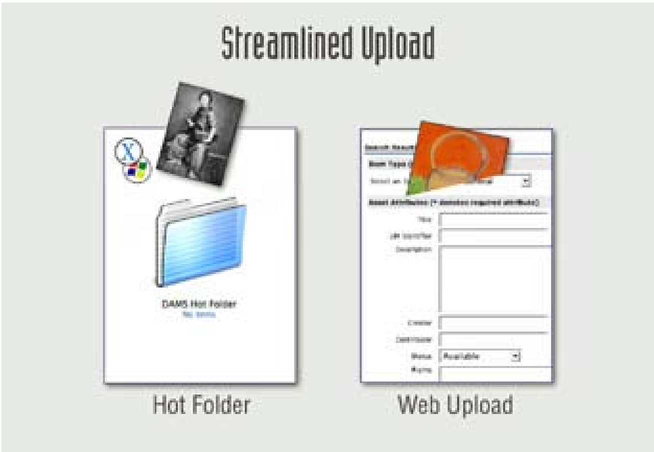

Figure 15: Streamlined Upload *http://media.nmc.org/2008/proceedings/king/StreamlinedUpload.mov*

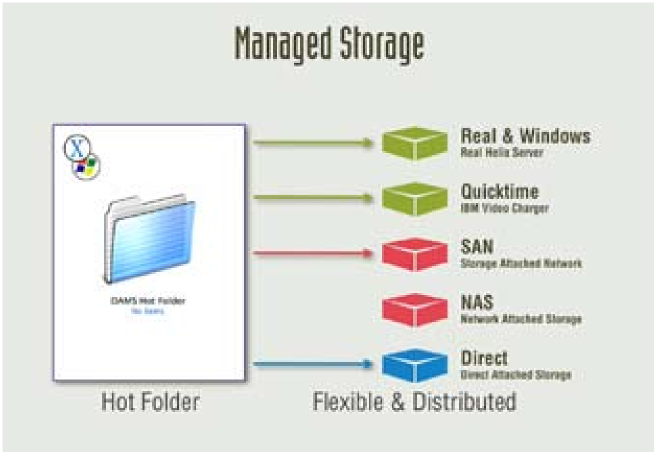

Figure 16: Managed Storage *http://media.nmc.org/2008/proceedings/king/ManagedStorage*

Flexible Metadata

A great deal of the power extracted from a digital repository is derived through metadata. Metadata provides technical, descriptive, and structural information that drives end-user search, retrieval, and use.

BlueStream eliminates early stumbling that can occur over this new concept. Metadata such as EXIF and XMP, encoded within the header of certain file types, is automatically read. Metadata from an external database, a spreadsheet, or an xml file can be automatically matched and associated to assets (See Figure 17). This provides a variety of methods for people and other systems to automate metadata acquisition. When automation is not possible, users can use the web interface to tag metadata.

The Dublin Core metadata schema is provided by default. Dublin Core is sufficient to support a majority of the metadata needs. When Dublin Core is not sufficient, project teams work with a metadata specialist from the University Library to identify standardized metadata schemas used for specific disciplines or when necessary, develop custom schemas (See Figure 18). Standards are adopted whenever possible to ensure that assets stored in BlueStream will "play well" with other systems and easily "migrate forward" to new components within the BlueStream infrastructure.

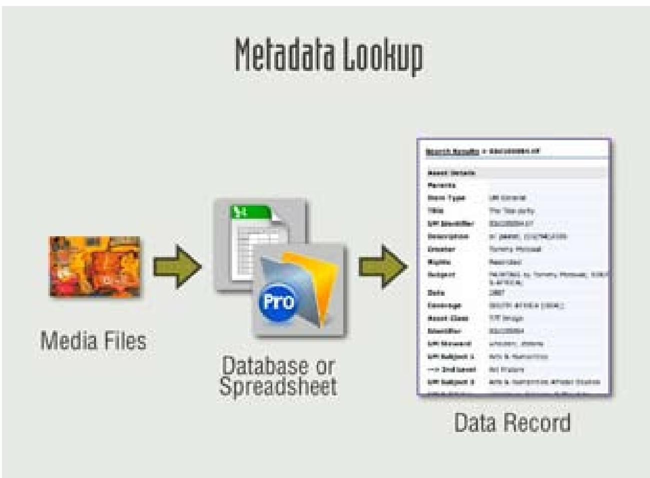

Figure 17: Metadata Lookup *http://media.nmc.org/2008/proceedings/king/MetadataLookup.mov*

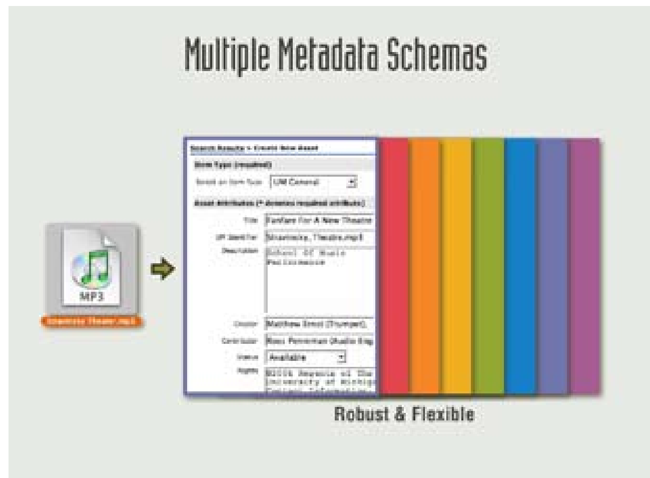

Figure 18: Multiple Metadata Schemas *http://media.nmc.org/2008/proceedings/king/MultipleMetadataSchemas.mov*

Organization & Control

BlueStream search capabilities takes advantage of full text indexing of file-level metadata, time coded metadata, and text-based documents. Users are able to use simple terms and Boolean strings to search across available assets and return query results in order of relevance. Collections allow users to loosely associate related files based on their own organizational constructs. Since collections are essentially a collection of "file pointers," a single asset can be represented in any collection without duplicating the file itself. Highly structured relationships such as the elements of a learning object – metadata, time-coded metadata, structural xml, documents, images, audio, and video – are supported through parent-child associations (See Figure 19).

Figure 19: Search, Organize, Relate *http://media.nmc.org/2008/proceedings/king/SearchOrganizeRelate.mov*

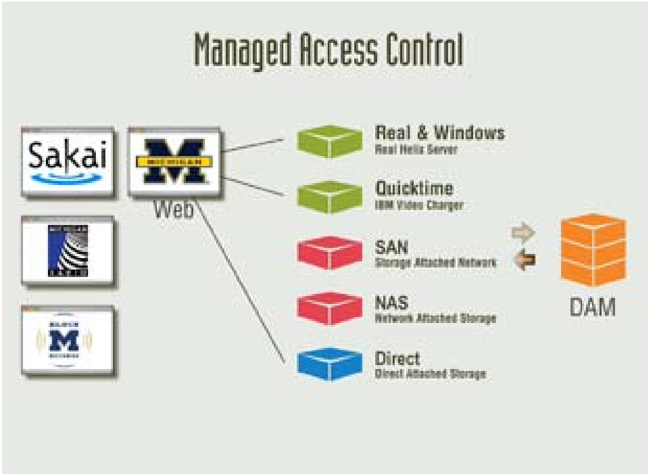

Figure 20: Managed Access Control *http://media.nmc.org/2008/proceedings/king/ManagedAccessControl.mov*

BlueStream secures every part of the asset, metadata, versions, and derivatives, through Access Control Lists. ACLs define who has access to an asset and what privileges the user has in regard to seeing, modifying, and deleting it. Fine-grained access control, set at the asset level, permits sophisticated approaches to supporting collaboration across different publishing and collaborative venues (See Figure 20).

Media Processing

Media processing is flat out too processor intensive and too time consuming for end-users to do. BlueStream takes this heavy lifting off of desktop machines and shifts it to powerful enterprise servers. Automated image conversion (See Figure 21), and video and audio transcoding (See Figure 22), assure that new assets have a baseline set of formats available. Additional formats needed to serve particular audiences can easily be ordered online. BlueStream will do all the work and automatically add the additional formats to the asset once the processing is complete.

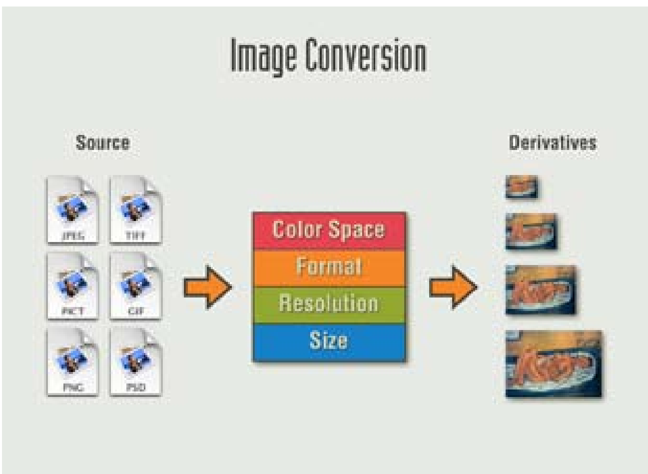

Figure 21: Image Conversion *http://media.nmc.org/2008/proceedings/king/ImageConversion.mov*

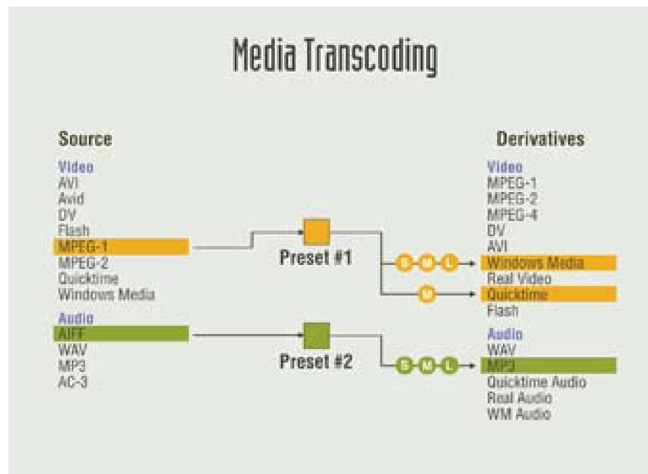

Figure 22: Media Transcoding *http://media.nmc.org/2008/proceedings/king/MediaTranscoding.mov*

Media Analysis

Many people consider metadata tagging to be a dreaded task. This attitude can be catastrophic to a digital repository since metadata is fundamental to the ability to find assets. Media analysis alleviates this dichotomy by using computer analysis to automatically generate metadata that enhances search and find capabilities.

Voice-To-Text analysis (See Figure 23) is particularly good for generating time-coded metadata of the spoken word. Accuracy levels of between 60% and 80% of key words can be achieved from quality recordings of non-accented speakers without background noise. This is not sufficient quality to support transcription. However, time-coded keywords add a valuable resource to the search index, allowing users to search within video and audio content without hand tagging metadata. Once the desired content is found the media will start playing from that exact point.

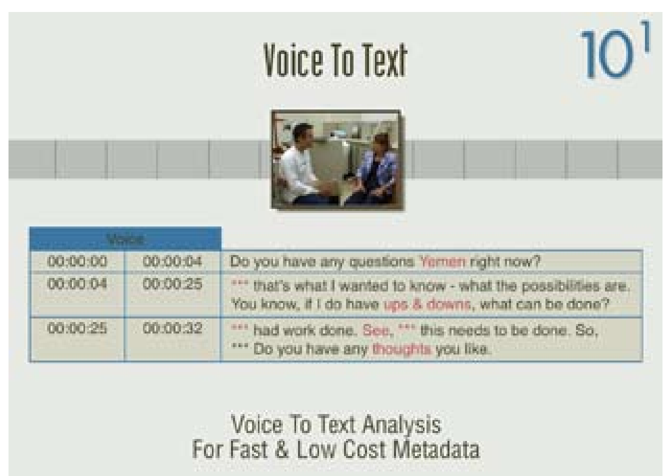

Figure 23: Voice-To-Text Analysis *http://media.nmc.org/2008/proceedings/king/VoiceToText.mov*

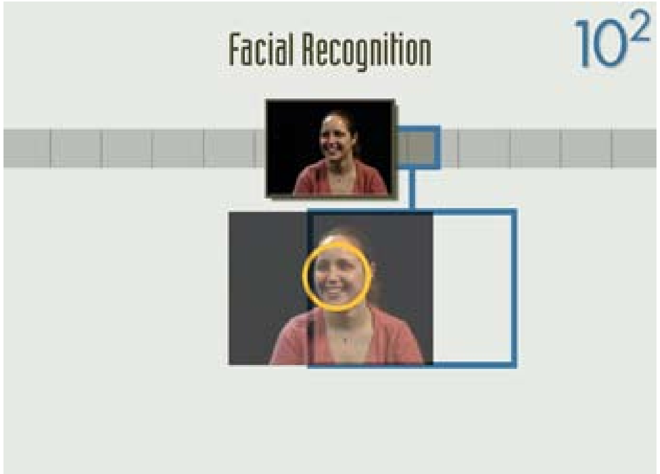

Figure 24: Facial Recognition *http://media.nmc.org/2008/proceedings/king/FacialRecognition.mov*

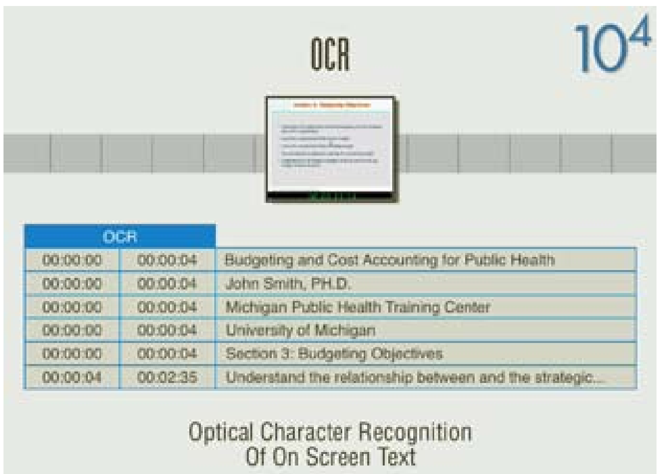

Figure 25: Optical Character Recognition *http://media.nmc.org/2008/proceedings/king/OpticalCharacterRecognition.mov*

Facial recognition (See Figure 24) compares video images to a library of faces. If a match is found, the name of the person is time coded and added to the search index. This is a particularly useful tool for indexing videos of public figures and panels of speakers. End users can easily search through video content to find precisely where a particular speaker appeared. Privacy and surveillance issues mandate that this feature be used with extreme care, in a manner that is transparent to the community, and with the consent of those who appear.

Optical character recognition (See Figure 25) analyzes English letters appearing in video, converts them to text, time codes the text, and adds it to the search engine. This is particularly useful for slide presentations, titles, and subtitles.

Scene change recognition generates a thumbnail of the video whenever the composition of the image changes dramatically (See Figure 26). The thumbnails and time code are saved as another part of the rich data model. They are used to present visual navigation through the content. The end-user can rapidly view the thumbnails, click on the one they are interested in, and play the video starting from exactly that point.

Media analysis is a machine process that has all the expected limitations that machines have in deriving meaning from human action. Nonetheless, the analysis is able to layer multiple channels of data resulting in a very useful ability to search within the content of media files (See Figure 27). This is very effective for lectures and other speaking content. In economic terms, automating metadata tagging, even with its inherent failings, is a huge advantage.

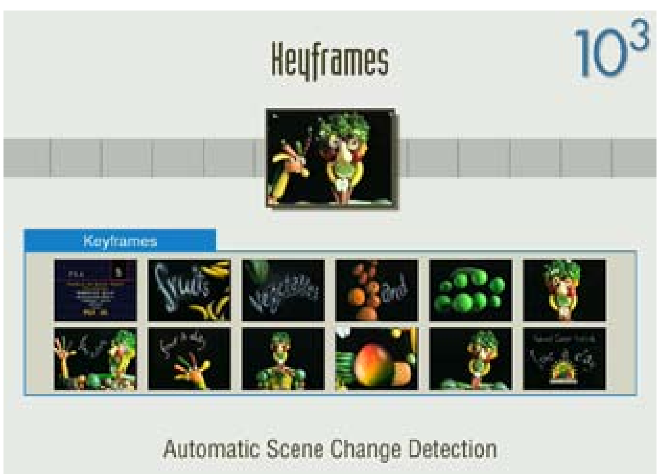

Figure 26: Scene Change Recognition *http://media.nmc.org/2008/proceedings/king/ScreenChangeRecognition.mov*

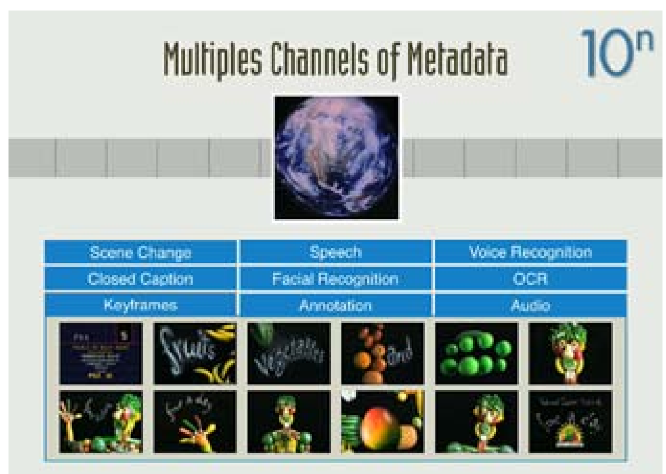

Figure 27: Multiple Channels *http://media.nmc.org/2008/proceedings/king/MultipleChannelsOfMetadata.mov*

Transcripts & Annotation

Automated, time-coded metadata generation is not always sufficient to meet academic requirements. Certain subject matter, outcomes, or activities might require higher standards. This is usually the case with video and audio data used for research. Transcripts, annotation, and event markup must be perfect to ensure that statistical conclusions are accurate and correct. BlueStream supports distributed research teams by easing the barriers to media markup. Required media formats are automatically generated and made available online. Researchers download these files via any Internet connection, use a markup tool to generate time-coded metadata (See Figure 28), and simply drag-and-drop the metadata file back into the system. BlueStream automatically parses the file and synchronizes it to the media (See Figure 29).

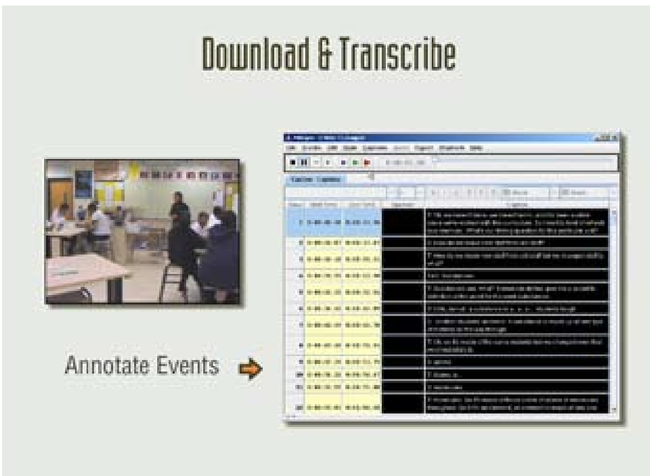

Figure 28: Download & Transcribe *http://media.nmc.org/2008/proceedings/king/DownloadTranscribe.mov*

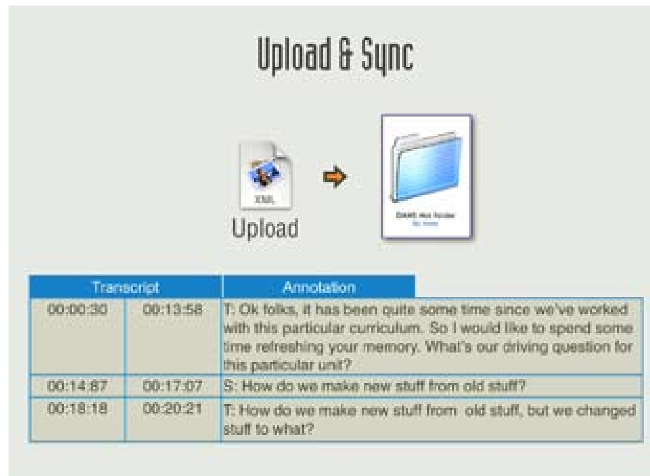

Figure 29: Upload & Sync *http://media.nmc.org/2008/proceedings/king/UploadSynch.mov*

Appendix 2: BlueStream Architecture

In April of 2003, IBM as primary vendor in conjunction with Ancept, Telestream, and Virage, was awarded a contract to partner with the University of Michigan in building a Digital Asset Management Systems (DAMS) Living Lab. The lab was to be a sandbox environment to support media intensive academic work in order to shape a larger solution for enterprise infrastructure for media-intensive work.

By 2006 the environment was sufficiently robust to rename the project for a broader academic use and BlueStream was born. Today, BlueStream supports 25-50 of the most media-intensive projects in the University and interoperates with the other enterprise level academic tools and environments. A brief outline of the technical architecture follows.

Software

1. IBM Middleware
   a. Content Manager Library
   b. WebSphere
   c. DB2 Database
   d. Tivoli Storage Manager
   e. Linux
2. Ancept Media Server
   a. Asset Management End-user Application
   b. Web Services Portal
   c. Media Services Integration
   d. Linux

3. UM Campus Infrastructure
   a. Authentication – Cosign
   b. Streaming Media – Real Networks Helix Server, Darwin Streaming Server
4. Autonomy Virage Video Logger – Media Analysis, Windows Server
5. Handmade Software, Image Alchemy – Image Conversion, Linux
6. Telestream Flip Factory – Media Transcoding, Windows Server
7. Quantum StorNext Advanced Integrity Module – File System

Hardware

1. Chassis – IBM eServer BladeCenter
   a. IBM eServer BladeCenter Chassis
   b. Nortel Networks Layer 2/3 Copper GbE Switch Module for BladeCenter
   c. IBM BladeCenter Optical Pass-thru Module
2. Blades – IBM HS21
   a. HS21, Xeon Quad Core E5430 2.66GHz/1333MHz/12MB L2
   b. Quad-Core Intel Xeon Processor E5430 80w 2.66GHz/1333MHz
   c. 2GB (2x1GB) PC2-5300 CL5 ECC DDR2 Chipkill FBDIMM Memory Kit
   d. QLogic 4Gb Fibre Channel Expansion Card (CFFv) for IBM BladeCenter
   e. IBM 73.4GB 10K SFF SAS HDD
3. Switch, 32 Ports – IBM System Storage San32B-3
4. Storage – IBM System Storage DS4200 Express Model 7V
   a. Expansion Units – IBM DS4200 EXP420
   b. Disks – IBM 750GB DS4200 SATA Enhanced Value Disk Dr

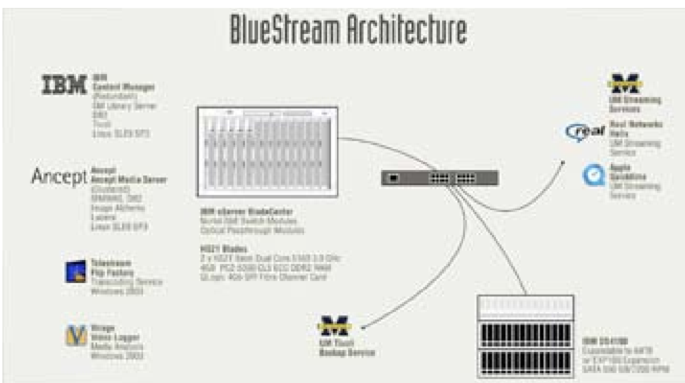

Figure 30: BlueStream Architecture

References

Buchanan, G. R., Blandford, A., Thimbleby, H., & Jones, M. (2004). Supporting information structuring in a digital library. RESEARCH AND ADVANCED TECHNOLOGY FOR DIGITAL LIBRARIES, 3232, 464-475.

Gantz, J. F., Chute, C., Manfrediz, A., Minton, S., Reinsel, D., Schlichting, W., et al. (2008). The Diverse and Exploding Digital Universe. EMC2, from http://www.emc.com/collateral/analyst-reports/diverse-exploding-digital-universe.pdf

Lessig, L. (2008). Keynote: Symposium for Teaching and Learning with Technology. Retrieved 04/15/2008, from http://symposium.tlt.psu.edu/content/lawrence-lessigs-keynote-presentation-version-1

Nardi, B. A., & O'Day, V. (1999). Information Ecologies: Using Technology with Heart: MIT Press.

NISO. (2004). A Framework of Guidance for Building Good Digital Libraries. Retrieved April 25, 2008, from http://www.niso.org/framework/Framework2.html

Suleman, H., Feng, K., Mhlongo, S., & Omar, M. (2005). Flexing digital library systems. DIGITAL LIBRARIES: IMPLEMENTING STRATEGIES AND SHARING EXPERIENCES, PROCEEDINGS, 3815, 33-37.

Weber, L. (2007). Computer Use Expected To Top 2 Billion. Inc.com. Retrieved April 25, 2008, from http://www.inc.com/news/articles/200707/computers.html

About the Author

Louis E. King serves as the Managing Producer of Digital Asset Management Systems in the Digital Media Commons at the University of Michigan. His current project, BlueStream, is developing campus infrastructure to support media-intensive teaching, learning, and research. This work is a continuation of fourteen years of progressive leadership in shaping academic tools and cyber-infrastructure to support emerging practice in the academy. As Director of the Office of Instructional Technology he spearheaded the development of a seminal course management system. This established a foundation for the eventual emergence of Sakai. Louis also served as Manager of New Media Initiatives where he established one of the original thirteen New Media Centers. Prior to joining the University of Michigan Louis spent 12 years in advertising as partner of a boutique creative house where he specialized in business-to-business communications.

Louis can be reached at leking@umich.edu.

Further information on BlueStream is available at http://sitemaker.umich.edu/bluestream