

University of Southern California

From the Selected Works of Win Shih

August 14, 2012

Disaster Planning and Business Continuity Goes Virtual

Win Shih, *University of Southern California*



Available at: https://works.bepress.com/win_shih/26/

Disaster Planning and Business Continuity Goes Virtual

Win Shih, Director of Integrated Library Systems
University of Southern California

Abstract

Located in earthquake-predisposed geography, University of Southern California (USC) Libraries have long made every conscious and strategic decision to fortify disaster preparedness as well as to diversify their IT facilities and operations in order to minimize risk and potential impact of a large-scale temblor. In this paper, we will discuss USC Library's cloud-centric IT infrastructure and holistic business continuity model; how it fits into our IT operations; and how we mediate vulnerabilities and minimize risk.

Introduction

As our world continues spearheading its digital universe – where information is created, stored and shared in massive arrays of ones and zeroes, bits and bytes and “Brontobytes Beyond” – we are collectively faced with myriad fundamental, lifestyle challenges and irrevocable changes. Our “digital vault” has grown at astonishing speeds and capacities. According to the annual IDC Digital Universe Study¹, in 2011 the amount of digital information created and replicated exceeded 1.8 *zettabytes* (1.8 trillion gigabytes) – having burgeoned by a factor of *nine* in five years alone.

The increase of digital assets is no less evidenced than within academia. Libraries are experiencing rapid expansion and migration to digital-based collections, services and infrastructure. We are diverting more and more of our resources toward acquiring and managing electronic collections at the demand of our patronage². Increasingly, scholarly works and publications are now born digital; researchers are required to make their publications and data openly available by mandate of their institutions and funding agencies³. Academic

¹ The 2011 Digital Universe Study: Extracting Value from Chaos, <http://www.emc.com/collateral/demos/microsites/emc-digital-universe-2011/index.htm> (June 2011)

² Ithaka S+R Library Survey 2010: Insights from U.S. Academic Library Directors, <http://www.sr.ithaka.org/sites/all/modules/contrib/pubdlcnt/pubdlcnt.php?file=http://www.sr.ithaka.org/sites/default/files/reports/insights-from-us-academic-library-directors.pdf&nid=360> (April 2011)

³ ROARMAP: Registry of Open Access Repositories Mandatory Archiving Policies, <http://roarmap.eprints.org/>

institutions now configure digital repositories to collect, provide access to and preserve the intellectual output and conserve of their members. For curation, preservation, and accessibility purposes, libraries are investing heavily in digitizing unique, often obscure and underutilized, special collections and archives⁴.

Challenges confronting libraries are not solely confined to the exponential growth of binary data, but also the diversity of data formats and associated metadata. As we rapidly continue digitizing audio and video information, we are not only experiencing the immediate growth of information file sizes and databases, but also a lack of standards and best practice for preserving and sharing an ever-exploding datastore of audiovisual, information content.

Throughout this paper, we shall explore: University of Southern California (USC) Library's cloud-centric IT infrastructure and holistic business continuity model; how each fits within our IT operations; and how we mediate vulnerabilities towards minimizing risk. We shall further discuss our process for developing comprehensive disaster response, recovery and business continuity planning based upon three key, interdependent, interlocking elements: Technology, People, and Communications.

University of Southern California Libraries

Located in Los Angeles, the University of Southern California (USC, <http://www.usc.edu/>), established in 1880, is now one of the world's leading, top-ranked private research universities. With a total student population of 37,000, USC enrolls more international students (8,000+) than any other American university.

USC Libraries is a member of the Association of Research Libraries (<http://arl.org/>), consisting of 23 constituent libraries and information centers (<http://www.usc.edu/libraries/locations/>), along with the USC Digital Library (<http://digitallibrary.usc.edu/>).

Fueled by an annual budget of approximately \$40.2 million and supported by a staff of 220 professionals, USC Libraries boasts rich, format-diverse collections housing content dispersed across physical and virtual space. Our print-based collection consists of 4.1 million typeset volumes, nearly 50,000 linear feet of manuscripts and archives, and more than 160,000 items including maps and cartographic research aids. Subscribed electronic resources, by comparison, contain more than half-a-million electronic books and journals. Moreover, USC

⁴ ACRL Research Planning and Review Committee. 2010 top ten trends in academic libraries: A review of the current literature. College & Research Libraries News, <http://crln.acrl.org/content/71/6/286.full> (June 2010)

owns a wealth of film and video recordings, audio recordings, microform units and other visual materials.

The USC Digital Library (<http://digitallibrary.usc.edu>) comprises a trove of historical photographs, oral histories and sound recordings, moving images, drawings, maps, documents, physical objects and other materials owned by USC and collaborating institutions. 33 unique collections consist of a total of 300,000 digitally archived items. The size of these collections varies from 14 object items in The Hancock Collection, to over 62,500 items in The International Mission Photography Archive. Total size of USC's Digital Library currently exceeds 12 terabytes (2 million megabytes).

Disaster Preparedness, Recovery, and Business Continuity

Situated dangerously close to earthquake-predisposed geologic topography, USC and its surrounding campuses have long made every conscious and strategic decision to fortify disaster preparedness and business continuity in the event of catastrophe. The University has devoted considerable resources to planning and preparing for disasters, both nature-made and human-instantiated. In tandem with Southern California's Earthquake Center, headquartered at USC, the University's Emergency Planning Office (EPO) maintains a campus-wide, emergency operations plan. Concomitantly, each school and department, under the coordination of EPO, maintains an internal emergency contingent guiding its own operations.

Within the IT sphere, USC Libraries has cooperated and corroborated with campus Information Technology Services (ITS), along with the Center for Scholarly Technology, in diversifying their combined IT facilities, resources, and operations in order to reduce risk and the potential impact of a large-scale temblor. Campus ITS has already set up a disaster recovery "hot site" located at a SunGard facility in Scottsdale, AZ. Redundant systems and scheduled synchronization between production and backup systems are prepared at this site for University-wide, mission-critical services. These services include enterprise resource planning systems, the identity management system, USC's E-mail system, and the Blackboard learning system. A Wide Area Network (WAN) connection between the two sites supports operational bandwidth during a serious emergency of up to 10 Gigabits per second.

For systems managed by USC Libraries, we have similarly decided to investigate ancillary hosting options supporting redundancy (disaster "alacrity") and agility, with business continuity. Since 2010, we have implemented the following, seven crucial modifications in the exigent interest of "risk readiness:"

Transformed

Traditionally, USC Libraries placed all of its servers in Campus ITS's data center and took advantage of their state-of-art, highly-secured facility, enjoying: expandable storage; redundant power supply and networking facilities; high security measures; copious bandwidth; temperature control; taped backups; and 24 x 7 onsite support. However, such centralized collocation precluded any "safety net" of backup facilities external to our precarious geographic base. Ergo, we conducted cost analyses for our Symphony Integrated Library Systems: "Three scenarios implemented for five years" were analyzed, compared and contrasted:

1. On-premise system with no mirror site;
2. On-premise system with a mirror site at SunGard facility, AZ;
3. Migration to SirsiDynix *Software as a Service* (SaaS) platform with a mirror site.

Our result was that Option 3 proved *most* cost-effective. The vendor's facility offers the same perquisites as our Campus ITS data center without the potential risks. Additionally, SirsiDynix takes care of all system upgrades, previously performed by our in-house staff. Further, SaaS follows a leasing model, where USC Libraries' annual cash outflow is much more manageable than traditional upfront, capital outlay for brand new hardware.

Thus, we migrated our ILS to the cloud-based Software as a Service environment in the summer of 2011. After operating in this new environment, we've benefited not only by having our mirror site in SirsiDynix's backup facility in Utah, but also enjoy: a robust, scalable system; responsive professional support from vendor's SaaS Team; manageability of and flexibility for future system expansion and growth; more frequent and smoother system upgrades. In one year, we have already performed two major ILS software upgrades. Running the latest version of SirsiDynix empowers us to take fullest advantage of the latest features and enhancements of our strategically-chosen SaaS solution.

Via a comparable studied method, we then moved two more library systems into the cloud:

1. EZproxy Authentication: This service validates legitimate patrons from off-campus to access our subscribed electronic resources and is one of our most heavily used services in addition to our Libraries Discovery service (Serials Solutions' Summon)

and online catalog (SirsiDynix Symphony). We transferred EZproxy operations from three virtual servers in the Campus ITS data center to an OCLC facility in Ohio. During a full year of applying this new SaaS, we experienced zero down time and zero service degradation.

2. USC Digital Library: Our digital library was originally built upon EMC's *Documentum* platform, designed as an enterprise content management system. Our software developers added additional customization to make it easier for external users. After assessing the pros and cons, we decided to adopt a system that is tailored for managing library digital collections and needs. Hence, we are now in the process of migrating our metadata and digital objects to OCLC's *CONTENTdm* hosting service.

Other important library services and tools, including our discovery service (Serials Solutions' *Summon*), interlibrary loans (OCLC's *ILLiad*), electronic reserves (Atlas System's *ARes*), and subject guides (Springshare's *LibGuides*) are all *SaaS*-based and vendor hosted. Consequently, we are almost exclusively "operating in the clouds" today, two years after our initial risk study, with no library production servers on USC's geographically vulnerable premises.

Built

As our digital collection continues to swell, we are looking for a long-term data storage solution to house preservation copies of digital files. With provost approval, USC Libraries partnered with USC ITS and USC Shoah Foundation Institute and created the USC Digital Repository (USCDR, <http://repository.usc.edu>) in 2011. With over 70 petabytes (70 million gigabytes) of capacity, USCDR allows campus units to consolidate all digital assets and storage management into one single operations facility with tighter security, greater reliability of data storage integrity, and supra-stringent backup operations.

The USC Shoah Foundation Institute (<http://dornsife.usc.edu/>) was originally founded and funded by Steven Spielberg to record testimonies in video format of survivors of and witnesses to The Holocaust. Its historically significant collection encapsulates more than 52,000 videos (105,000 hours) of Holocaust survivor and witness aural-visual testimonies, amounting to over 8 petabytes (8 million gigabytes) of digital data.

The USCDR is not only a warehouse of digital assets; it provides a full range and sterling example of cloud-based services effectively managed and maintained, per the following

organizational steps for the storage (conversion), preservation (integrity) and share-ability (online access) of invaluable digital content:

1. Digitization—converting physical or electronic collections to standard digital formats for preservation and online access;
2. Cataloging—creating metadata that allows users to locate and search digital collections in a consistent, meaningful way;
3. Preservation—ensuring the long-term accessibility and integrity of digital collections;
4. Web Access—providing optimal online access to digital collections through high-speed networks and multimedia productions;
5. File-server access—offering high-quality, high-bandwidth file management capabilities to access, manage, and manipulate large digital collections.

USDR’s mirrored “cloud center” is in remote Nevada, with nearly infinite digital storage capacity.

Remedied

In the cloud environment, our servers are typically located thousands of miles away from the campus hub. Our tests indicate that there is discernible network latency for accessing our ILS servers in Atlanta, GA. During busy times of day when the Internet is congested, the network can become “the weakest link” or “single point of failure.” To minimize this bottleneck, we work with Campus Network Services and SirsiDynix to set up site-to-site VPN tunneling between our campus and their data facilities to ensure adequate bandwidth is maximized for such traffic. The data transmitted across our private network is encrypted and secured.

Automated

There is a transitional switch from USC Library’s main website to an emergency website when needed (<http://libraries-emergency.usc.edu>). Along with the University’s emergency website (<http://emergency.usc.edu>), we use Blogger (<http://www.blogger.com>), a free service, to host the emergency site. When the main site goes down, the toggle switch automatically redirects users to the emergency site. This site serves two principal functions:

1. An Information and communication portal: The main section of the site offers news and updates in the event of emergency. It also provides contact information for emergency service throughout the city, University, and library,

along with twitter feeds from the University, libraries, and the University's Safety Alert Service.

2. A resource portal: It offers links to the following key library resources -
 - Search tools: discovery service and library catalog;
 - Electronic resources: database, e-journals/e-books, electronic reserves, and subject guides;
 - Services: interlibrary loan and "Ask-a-Librarian" services;
 - Library emergency procedures documents (for staff only, login required).

Select members of the library are granted administrative rights to update the content of emergency website. These individuals may also perform updating via mobile device.

Inculcated / Practiced

At USC there is redundancy not only regarding IT infrastructure but also pertaining to IT personnel knowledge and responsibilities, vis-à-vis cross-training and annual disaster drills.

Based on a template from the University's Emergency Planning Office, USC Libraries have developed a written Disaster Recovery and Business Continuity Plan. Led by the Senior Associate Dean of Libraries, this plan embraces every functional unit in the library and has been made available for all staff members on our library's intranet. Working with the University's Emergency Operations Center, USC Libraries' Recovery and Continuity Team is coordinated by a team leader and composed of the following eight sub-teams:

- Collection Recovery
- Communications
- Facilities and Space Management
- Finance / Business Office
- Human Resources
- Information Technology
- Logistics / Resource
- Public Services

Every spring, the entire team meets to review and update the plan. Then we go over a tabletop exercise with a fictitious emergency scenario. Each sub-team then gathers to develop a disaster recovery plan and report back to the whole group.

In the ILS unit, we also work with our vendors to perform annual disaster recovery tests each summer. We test major functions on our backup server in Utah, to ensure everything functions “smoothly” in the event of any hypothetical calamity.

Grown

Moving ILS to a cloud environment, thereby forfeiting some system administrative control to our vendors, and relocating data away from our domain – these were drastic changes from our traditional operations. The transition entailed new procedures and protocols working with vendor technical support in troubleshooting, emergency response and delineating coverage and duties. Fortunately, our informed optimism has never waned: Developing rapport with vendor’s IT support personnel and learning more about their products. We also feel we have less to worry about now, receiving fewer phone calls after hours when there is a system problem. Staff may now devote their time and efforts from managing the system, to other areas which – prior to the migration – they had been unable to focus upon, let alone address meaningfully.

Conclusion

As Ben Franklin pithily observed, “An ounce of prevention is worth a pound of cure⁵.” Since the traumatic aftermath of Hurricane Katrina in 2005, Disaster Recovery and Business Continuity has been rated among the top-ten vital issues facing higher education chief information officers in EDUCAUSE’s annual “Current IT Issues Survey⁶.” However, according to Acronis’ second annual Disaster Recovery Index survey, less than 50 percent of small and medium businesses in U.S. feel confident in their backup and disaster recovery plans and 76 percent of all respondents indicated that the best way to improve their plan would be to have an integrated system encompassing physical, virtual, and cloud data assets⁷. With rising expectations of

⁵ The Electric Ben Franklin, <http://www.ushistory.org/franklin/quatable/quote67.htm> (December 2011).

⁶ Ranking, in parentheses, of Disaster Recovery and Business Continuity, for each of past 6 year’s EDUCAUSE Survey: 2011 (9); 2010 (6); 2009 (8); 2008 (6); 2007 (5); 2006 (4), <http://www.educause.edu/issues>

⁷ The Acronis Global Disaster Recovery Index: 2012, http://eu.acronisinfo.com/?attachment_id=499

“everywhere, any time, all the time” availability and accessibility of library resources and services by our global patrons, there is no present like the time for libraries to begin developing an integrated disaster recovery plan with business continuity, gazing forward to “the day we hope never comes.” Libraries’ heightened investment in and commitment to online services; increasing complexities of campus IT infrastructures; advanced technologies employed to maintain digital assets; potential damage to an academic institution’s image and brand if (or when) catastrophe strikes – what more pressing reasons exist; what more urgent window than now for librarians and IT professionals to cross aisles, join as one unified force, communicating and codifying an integrated risk-readiness package to safeguard our physical and digital assets. No solution is omnibus, perfect or foolproof, but one which protects our collected legacy of human understanding across millennia and into the future is worth the time, attention, and ongoing archival efforts.