

1-1-2011

Report on the Heads of Library Technology  
Interest Group Panel Presentation: Emerging  
Technologies: Virtualization in Libraries, American  
Library Association Annual Conference,  
Washington, DC, June 2010

Edward Sanchez

Marquette University, [edward.sanchez@marquette.edu](mailto:edward.sanchez@marquette.edu)

# Report on the Heads of Library Technology Interest Group Panel Presentation: Emerging Technologies: Virtualization in Libraries, American Library Association Annual Conference, Washington, DC, June 2010

Edward Sanchez

*Raynor and Memorial Libraries, Marquette University  
Milwaukee, WI*

## **Opening Comments by Incoming Chair Edward Sanchez (Marquette University)**

The purpose of Heads of Library Technology (HoLT) is to provide a forum and support network for people with administrative responsibility for computing and technology in a library setting. Last year during ALA Annual, participants in the HoLT Interest Group meeting selected the topic on "Virtualization in Libraries" from a long list of challenges facing libraries. Many of us have heard compelling reasons for putting library resources into the cloud during other sessions at ALA. This panel presentation, however, is on the in-house use of virtualization technologies in large research, midsize, and public libraries. Our speakers represent a variety of backgrounds and types

of libraries, and will describe how leveraging virtualization in their particular setting has led to increased capacity, while reducing costs, downtime, and management headaches. It is our hope that by sharing our virtualization successes we will broaden the discussion on when to build locally and when to move to the cloud.

## **Virtualization in Large Research Libraries, Dave Pcolar (University of North Carolina at Chapel Hill)**

### *Library Information Technology Environments*

Library information technology (IT) environments are faced with a variety of challenges including increasing demands and decreasing resources, increased complexity, aging facilities, increased management expectations, and a market demographic moving rapidly from digital immigrants to digital natives.

Increasing demands are taking the form of web services outside the catalog, remote access, and digitization services like Special Collections that include manuscripts/maps, photographs, audio recordings, and archives. The growth of born digital collections, such as Electronic Theses & Dissertations (ETDs), and scholarly communications has given rise to digital repository services and the challenges of long-term digital preservation.

In libraries across the country decreasing resources are reflected in the reduction in continuation budgets which has given rise to one-off purchasing and difficulty in sustaining maintenance contracts. With less to go around internally libraries are looking externally for grant and specially funded project resources.

Complexity in library IT departments has increased as support of open source development continues alongside new development in JAVA, Tomcat, Drupal, and Joomla. On the hardware side server and storage continues to evolve with heterogeneous new products appearing with increasingly complex requirements.

Libraries with aging machine rooms and physical infrastructure are further challenged with increased costs for cooling power. According to the International Data Corporation Worldwide Server Power and Cooling Expense 2006–2010 Forecast, energy costs may

increase from 10% of the typical IT budget today to more than 50% in the next few years. Many departmental server rooms and data centers struggle to justify themselves in the face of campus green initiatives and space allocations.

Library management teams driven by innovation are pushing to extend lifecycle replacement on equipment to make funds available for promising technologies while simultaneously placing additional stress on support and service delivery units.

All of this is market driven by a demographic of 18–24 year olds whose expectation of online access is conditioned by its presence all their life as well as an increase in collaborative projects/assignments, and a significant change in expectations driven by “internet time” pressures.

### *Why Virtualization?*

Gartner states that during a 24-hour period, less than 10% of the available computing power is utilized. Current IT infrastructures are very inefficient due to underutilization, especially with x86/x64 servers. Hot/Cold server pools for critical applications double hardware costs and isolate resources to specific applications. Servers dedicated to proprietary software, or separate operating systems(OS) requirements, or that “do not play well with others” along with low demand but mission critical services with low transaction rates, all create inefficiencies that add up. On the storage end, storage allocations that are locked to specific machines or fully provisioned at service inception can waste valuable disk space and availability.

### *Infrastructure Choices*

We chose VMware as our enterprise platform because it has a proven track record in Industry, a comprehensive product suite, and supports a wide range of hardware. Our design team decided on a mixed storage environment with Tier 1 for critical applications (currently NetApp), Tier 2 for large data applications (Sunfiber arrays, Nexsan SATAbeast), and Tier 3 for disaster recovery/replication [central campus IT services with tape (SAM-FS based) and an Iron Mountain service agreement.

## *Implementation Details*

In 2007 we had 60+ stand alone hosts, 35 server instances, manual failover, 20TB of mounted data, and 2 machine rooms. In 2010 we have 18 standalone hosts, 7 physical VMware hosts, 68 server instances, fully automated failover, 150TB of mounted data, a single machine room, central IT outsourcing, and a cost avoidance of \$175,000.

In general, we see the overall cost reduction, reduced downtime, patch management, disaster recovery, and rapid deployment as the upside of the project. On the downside are licensing/maintenance costs, increased complexity, and increased coordination requirements.

## **Desktop Virtualization in Mid-sized Research Libraries, Stu Baker (Northwestern University)**

The challenge of desktop virtualization in our setting is that there is no one-size-fits-all solution. It may work for some users and not others, and IT staff needs to match individual users with the right technical solution. This inevitably involves considering four key items: performance, peripherals, scalability, and cost.

Desktop virtualization technologies may be provisioned in a variety of formats: hosted shared desktops, hosted VM desktops (VDI), hosted blade PC desktops, physical desktops with on demand apps, local streamed desktops, and local VM-based desktops (offline).

(1.) Hosted Shared Desktops: This solution utilizes a local machine running a server based OS that is not managed, using some sort of client through a web browser to connect to a server that publishes the operating system and on-demand applications. It allows for user personalization of the applications.

(2.) Hosted VM Desktops (VDI): there are multiple ways to deploy VDI

- Thick Client—where an existing PC or other boot device serves as host,

- Thin Client—depends on some other computer, usually a server, to run the OS and applications,
- Type Zero Client—bare bones, no installed software or local caching of information. Used in high-security operations.

(3.) Hosted Blade PC Desktops: Used for higher performance and complex applications (e.g., 3D, CAD). There is a one-to-one ratio of Apps for each VM providing dedicated memory and processing.

(4.) Local Streamed Desktops: This is typically a Netboot situation where the local host OS is launched from a networked image. Allows for more personalization. Best to store user data on a separate network drive.

(5.) Physical On-Demand Apps: Utilizes a local machine running a local OS that is not managed. Connection done using some sort of client through a Web browser to connect to a server that publishes on-demand applications that are locally installed. Data synced back to server. Good for mobile situation.

When determining the right fit, each of these technologies is matched to a particular user type on a continuum from task worker to mobile user with the former more dependent on server-side resources and the latter on clientside resources. Managing applications and an OS on every device, including changes, additions, and patches is time consuming. Dealing with security and data recovery is more complex, and hardware replacement cycles are expensive.

The advantages of virtual desktops are:

- Flexibility for users and IT staff,
- Use any device anywhere and anytime,
- Run multiple desktops from a single device,
- Faster data recovery and ideal for business continuity strategies, and
- Better security, space utilization and energy savings.

The recommended takeaways from this session are: (a) start by defining your service models and gather user/business requirements; (b) assume you will have a mixed environment; (c) align use cases/requirements to the appropriate technology; and (d) be aware

that the user/patron experience is the most important driver of success.

## **The Nuts and Bolts of Virtualization in a Mid-sized Public Library, Matthew Hamilton (Boulder Public Library)**

When Hamilton started at the Boulder Public Library (BLP) in January of 2009, he walked into an infrastructure that had been patched together over the years, and there was inconsistent planning for the growth of the infrastructure. There was no established replacement cycle or formal disaster recovery plan in place.

When Hamilton arrived on the job, BPL had 12 physical servers in place. At the time, it was a completely Windows Server 2003 shop with a variety of needs not best handled by this infrastructure. The physical hardware was aging and inconsistent, for example, a DNS server running on an old desktop machine. Most were Dell PowerEdge 1850s or PowerEdge 2850s that were a year past their end of life and were limping along on extended warranty.

Hamilton surveyed what BPL had and found a lot of processing capacity going unused. Windows servers were running simple, single tasks as their sole function that in some cases were never consuming more than 3%–4% of their processing power. Others were more demanding, but still didn't use more than 28% at their peaks. Hamilton knew that Linux could handle many of these functions with a much smaller footprint.

What Hamilton didn't have was a large pile of cash. On the contrary, within the first two weeks of coming on board, he was asked to take a \$15,000 reduction in budget. Additionally, BPL lost a contract employee who'd been with the library for 11 years and held a large portion of the web development and server administration knowledge.

What he needed was a strategic deployment of their resources to get the job done. He needed an infrastructure that was easy to manage with rock-solid reliability in terms of business continuity and disaster recovery.

Hamilton was also charged with revamping and ramping up BPL's web development efforts. Similar to what he found with the data

center, the web infrastructure was fractured, inconsistent, and out of date. They had websites developed by outside contractors in ColdFusion and ASP without the expertise in-house to support them as well as static websites running on IIS with dynamic content expectations. BPL's web developers were spending their time updating content instead of on development. High hopes and expectations abounded as Hamilton was charged with updating the look and capabilities of all five external and the internal websites.

So while spending less money, he needed to essentially re-establish the physical infrastructure of the data center, while improving on its capabilities. His three main objectives were to (1) simplify disaster recovery, (2) simplify server management, and (3) lower the barriers to innovation.

Knowing that Drupal could solve all of the web development problems they had, Hamilton made the conscious choice to take efforts off of continually patching outdated web services and move toward replacing them with a more modern platform. This involved asking a staff member who had exclusively used Microsoft products for years to take on the task of learning and managing Linux servers. Gradually they rolled out new versions of each of the websites on a Linux, Apache, MySQL, PHP (LAMP) platform and bring down the Windows IIS servers. Hamilton needed to make this transition as easy as possible for staff members while moving quickly enough to meet the demands of a reference and administrative staff hungry for long overdue enhancements to web services.

Today BPL has five physical servers: three virtual host machines (two production and one test server), an ILS server (they were not ready to take the plunge with that one yet), and a storage server for their local history digitization projects. Though the job is far from finished, with the server infrastructure part resolved, BPL had the opportunity to move on to web development.

Hamilton first met with a colleague, Eric Sisler (Westminster Public Library), who had been doing virtualization for years, and grilled him on how it worked for him in his even smaller library. He used the free Linux-based Vserver platform and it had served him well, developing organically in his environment, but BPL had more Windows



servers to deal with and had the opportunity to start essentially from scratch.

The city IT department contracts extensively with a local firm, Applied Trust, and Hamilton turned to them as consultants on this project. They gave them a catalog of the services running on each of their machines and for a couple of weeks monitored the peaks and valleys of processor demand. Based on BPL's needs in terms of ease of management, where they were going, and their budget, they decided on VMware's free bare metal hypervisor, ESXI. It installed quickly and easily with a very small draw on resources, but could scale as necessary into the future.

After determining the amount of resources needed and planning the conversion, they moved on to the actual migration. Two PowerEdge 2950 servers under warranty were sitting unused. Hamilton's predecessor had been retired for several months before he came on board so his staff was waiting to see what direction he would take when he came on board and nothing had yet been done with the new machines other than some preliminary investigation into Microsoft's HyperV product and some initial testing of Windows Server 2008. They beefed up the RAM in each of the servers to 32GB and maxed out the storage capacity in each. It wasn't necessary to go with 15k rpm drives, but instead they used Seagate Savvio 10k 2.5 Serial Attached SCSI drives.

One of Hamilton's requirements from the beginning was the ability to failover or recover quickly in the event of a disaster or other service interruption, so he needed two hosts at minimum that could be separated into two physical locations. They built the capacity into both of these 2950s to handle all of their servers if necessary. However, in deployment they split the VMs between the two. They have one that provides primarily out-facing services, the web server, SMTP, etc., and one that is primarily dedicated to internal services such as the Intranet and reservation/print server for the public access computers. "P to V" (physical to virtual) conversion took place over less than a week. They identified servers that were least likely to impact public service and started from there. Using the VMware Standalone converter, each machine took between 2-4 hours to be converted to a VM.

The amount of labor saved was immeasurable, and they saw about a 26% reduction in energy use for the data center. Total hardware costs were about a third of what it would have been if even six of the servers that were aging out were replaced. Replacing all of the physical servers one-for-one would have cost upwards of \$22,000 even with consolidating services as much as possible. Even with \$600 spent for two years of VSphere Essentials, the equivalent of \$14,000 savings for the library was realized—or more than Hamilton's entire hardware budget for the year.

The end product looks something like this: a single interface to manage multiple VMs on each host that could monitor resource usage and adjust allocation as necessary. They also enabled the command line interface on the ESXI hosts, which enabled the use of scripts that were shared in the VMware community forums to create automated backups of the VMs on a nightly basis. In order to do this and to allow for easier transition between hosts they purchased a low-cost NAS unit and use that as a central data store. The VMs run directly off the hosts, but they keep nightly clones on the NAS that we can be spun up on another host at any time. Purchasing licensed products from VMware would add to the library's capabilities, including automatic failover or migration of live servers, but the cost isn't worth it for BPL's modest needs.

Currently BPL enjoys more centralized management of the servers. They have a much quicker disaster recovery process than before, and Hamilton has started turning his attention to the real goals of his department—enabling development of digital services for staff and customers. The wish list was long and, as anyone who works in IT knows, the demands for new applications and services can very quickly exceed the capacity and skill-set on hand to support them.

Something they hadn't planned on, but quickly became aware of, was the ability to leverage this technology to allow for cheaper and easier testing of new products. Virtual machines made rollouts and cloning painless. This "side benefit" to virtualization quickly became one of the most exciting features. Suddenly, BPL could provide test environments quickly and cheaply with a minimum of risk because these were isolated from the rest of the network and could be turned off or rolled back to an earlier snapshot at a moment's notice without affecting core services.

Markus Stobbs, who is a Systems Administrator for NCAR's (National Center for Atmospheric Research) data center in Boulder, introduced Hamilton to this idea by showing him how NCAR developed and rolled out individual virtual machines for their Drupal developers. Taking a look at the VMware virtual appliance marketplace gives an idea of the types of prepackaged appliances that can be downloaded: web servers, e-mail servers, content management systems, firewalls, domain controllers, and more—most of which are free. Hamilton soon found many other communities on the net providing an even wider range of appliances, which was instrumental in helping his staff become comfortable with Linux machines. He could provide them with a VM they could explore without being afraid of causing service interruptions or outages. They could take on a new project and learn at their own pace.

While virtualization has afforded BPL these and other significant benefits, Hamilton believes its greatest benefits have not yet been realized by the majority of library users. Not yet tapped are its utility for packaging and distributing freely available, and in some cases "best of breed" library applications for demo or production purposes; its potential for the distributed management and preservation of burgeoning digital collections; and finally its utility in moving from local to cloud-based systems. Hopefully this panel discussion gave a greater appreciation of virtualization in libraries and possibly some ideas to share with decision makers at home organizations.