Levitating Libraries to the Clouds: A Strategy for Academic Libraries

Mark Dehmlow
University of Notre Dame, mdehmlow@nd.edu

Follow this and additional works at: <https://docs.lib.purdue.edu/charleston>

Part of the [Business Administration, Management, and Operations Commons](#), [Computer and Systems Architecture Commons](#), and the [Human Resources Management Commons](#)

An indexed, print copy of the Proceedings is also available for purchase at:

<http://www.thepress.purdue.edu/series/charleston>.

You may also be interested in the new series, Charleston Insights in Library, Archival, and Information Sciences. Find out more at: <http://www.thepress.purdue.edu/series/charleston-insights-library-archival-and-information-sciences>.

Mark Dehmlow, "Levitating Libraries to the Clouds: A Strategy for Academic Libraries" (2015).
Proceedings of the Charleston Library Conference.
<http://dx.doi.org/10.5703/1288284316305>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Levitating Libraries to the Clouds: A Strategy for Academic Libraries

Mark Dehmlow, Program Director, Information Technology, University of Notre Dame

Abstract

The University of Notre Dame's Office of Information Technology is in the process of implementing a "Cloud First" strategy through which it intends to move 80% of its core technical infrastructure into the cloud by the end of 2017. The strategy advocates a tiered prioritization structure that recommends the hosting (SaaS) model for most services, the AWS (IaaS) model for fewer services, and finally on premises for a handful of the remaining services. As a campus technology partner, the Hesburgh Libraries has begun planning for moving many of our services and infrastructure into the cloud. This initiative represents a radical shift in mindset for technology planning, moving from thinking about technology as assets that need replacement every 5 to 7 years to thinking about IT as a utility, paid monthly based on actual usage. This presentation will cover how the Hesburgh Libraries is devising our plan to move a significant amount of our infrastructure to the cloud, and the phases we have outlined to meet the OIT's 3-year plan. We will cover:

- Our experimentation in Amazon Web Services and how AWS differs from our current infrastructure
- The assessment of library service catalogs, both in terms of function and usage
- Our determination of which hosting models meets the needs of our services (SaaS, PaaS, IaaS, or on premises)
- Planning for the migration and roll out
- Any other considerations libraries may need to evaluate in this process

Introduction

The University of Notre Dame's Office of Information Technology is in the process of implementing a "Cloud First" strategy through which it intends to move 80% of its core technical infrastructure into the cloud by the end of 2017. The strategy advocates a tiered prioritization structure that recommends the hosting (SaaS) model for most services, the AWS (IaaS) model for fewer services, and finally on premises for a handful of the remaining services. As a campus technology partner, the Hesburgh Libraries has begun planning for moving many of our services and infrastructure into the cloud. This initiative represents a radical shift in mindset for technology planning, moving from thinking about technology as assets that need replacement every 5 to 7 years to thinking about IT as a utility, paid monthly based on actual usage.

Cloud First

The central strategy underlying the "Cloud First" program (see Figure 1) is to leverage the efficiencies of hosting applications in the cloud and following a tiered decisioning model that starts with moving most applications to a SaaS hosted model where vendors are not only responsible for infrastructure, but also the entire application stack. This allows organizations to focus on basic configuration, data, and implementation. The "Cloud First" strategy then focuses the next tier of applications being hosted on an infrastructure platform such as Amazon's AWS, so that the underlying infrastructure is provided by the platform and can be implemented quickly, at scale, and without significant redundancy. Whatever doesn't fit the prior models will continue to be hosted on-site, at a data center close to campus (see Figure 2).

Figure 1. University of Notre Dame Cloud First Marketing Materials (Office of Information Technologies (Cloud First). <http://oit.nd.edu/cloud-first/>).

Amazon was selected by the University of Notre Dame as its preferred infrastructure provider because they are currently the cloud infrastructure leader in the marketplace, providing some of the most advanced data centers combined with an expansive and extendable set of tools for managing the utilization of their infrastructure.

IT Evolution and Implications of the Cloud

IT infrastructure provisioning has followed a similar arc of abstraction over the years, not unlike that of programming. Early programming focused on writing code to interface with hardware (machine language). Over the years, programmers have abstracted languages—assembly, procedural, and now object oriented, to take advantage of common lower level programming code that often does not need to be rewritten. Increasingly, even within Object Oriented languages, scripting languages, and CSS bundles, coders are writing frameworks and

SaaS >> IaaS >> On Premise

Figure 2. Tiered decisioning model.

common modules that abstract functionality so that code can be reused and written more quickly. IT Infrastructure is following a similar path, having once relied on mainframe systems which required management by an operator, to racks filled with bare metal servers (one physical server per system), to virtualized systems (one physical server per multiple systems), to now cloud-based infrastructure, an

abstracted management environment that allows for all systems to be managed from a web-based dashboard.

Moving to this model has several potential benefits. The first is cost. When looking at calculating the cost of IT infrastructure, it is important to consider the total cost of ownership for the whole organization, not just the unit provisioning a service. For example, our library does not pay for electricity, the capital costs to provision a data center, networking, data center operations, or many other costs associated with providing IT infrastructure on campus. These costs are absorbed centrally by the University or the Office of Information Technology. As we begin looking at the cost for us to provision our systems in the cloud, we are looking at those costs through a campus-wide stewardship lens. While it is likely to cost the Libraries more to provision systems in the cloud because our IT on campus has always been subsidized centrally, we are still considering the overall impact for the University. Our Office of Information Technology is committed to working with us as we move out into, what is for us, a more expensive IT model.

A more impactful cost consideration is in the repurposing of staff. At our organization, staff are

generally more difficult to get funding for because of the long term commitment and added cost of benefits, etc. As new needs continue to emerge, being able to repurpose some of our most skilled, talented, and trainable staff is a huge benefit. Beyond the baseline cost factor, true cloud infrastructure will offer features that are considerably expensive and time-consuming to provide for a single institution—almost infinite scalability, the ability to provision services on-demand, and high availability that is realized by a geographically diverse infrastructure, in some cases implemented on a global scale.

Moving to the cloud, though, requires some shift in thinking about how to acquire and provision infrastructure. Instead of working with fixed budgets that are spent in 5-year cycles for purchasing equipment, the cloud model is a utility model based entirely on usage. It is more efficient in the sense that you use exactly what you pay for and you pay for exactly what you use—with locally implemented servers, there is frequently a lot of idle time where hardware is working at a fraction of capacity. Calculating utility-based estimates and budgeting for unfixed costs is a lot more nuanced and challenging than our historically fixed budgeting model.

Assessment: Preparing for the Cloud

Preparing for a broad infrastructure change required us to undertake a comprehensive service inventory. Like many institutions, this was an activity that we hadn't done for a long time, and it

had the benefit of us looking at our environment holistically. We started thinking about what should remain solo, what could be consolidated, and in some cases, what could be decommissioned. In addition to cleaning house and taking inventory, we assessed the service criticality, risk, dependencies, and impact of each service so we could, for the first time, begin to formally assign service-level definitions for each service we run. These definitions include things like who is impacted by an outage and how quickly should we bring services back online. This activity alone is helping to make us a more efficient IT organization.

We then began developing a rubric for how to determine what should go into AWS, what was a candidate for hosting with a vendor, and what we should keep on premise. For AWS, we asked questions like: Is it a web service? (Amazon is optimized for web-based applications.) Does it need extra reliability, scalability, and redundancy? How is it architected? (Applications developed with smaller, separable components are less expensive in the Amazon model.) What is the projected cost to run it in Amazon? For systems we thought could be hosted by a vendor, we asked: What software do they develop and maintain? Can they agree to and meet expected performance such as uptime, data throughput, etc.? How well has their support organization performed?

We looked at several examples and our current thinking about where we felt it was best to host the software.

Figure 3. Five–seven year cycle and monthly/yearly.

In addition to evaluating the qualities of our systems that make them a good or bad candidates for the cloud, we have also been monitoring system performance and resource utilization (CPU, memory, network

Table 1. System hosting assessment.

Aleph (Integrated Library System)	
System Characteristics: <ul style="list-style-type: none"> • client-based • large database • legacy software • vended software 	Candidate for: <i>On Premise Hosting</i>
Primo (Discovery System)	
System Characteristics: <ul style="list-style-type: none"> • web-based • large indices and database • mature but relatively new software • vended software • vendor has hosting model 	Candidate for: <i>Vendor Hosting</i>
Inventory Management System	
System Characteristics: <ul style="list-style-type: none"> • web-based • small database • low amount of input/output • newer software and architecture • open source, locally developed 	Candidate for: <i>Amazon Web Services</i>
Hydra/Fedora Based Institutional Repository	
System Characteristics: <ul style="list-style-type: none"> • web-based • medium database • diverse and complex storage needs • newer software 	Candidate for: <i>On Premise, but later Amazon Web Services</i>

bandwidth, etc.) using software called Ganglia.¹ This is helping us develop predicted costs in Amazon.

The migration plan is then fairly simple. Anything new that we develop, we should look first to AWS

as a host. For our existing systems, we will follow our system assessment recommendations.

Amazon Components and Billing Model

Amazon provides many different infrastructure components that can be implemented together to provide comprehensive hosting solutions. The

¹ Ganglia Monitoring System.
<http://ganglia.sourceforge.net/>

following list gives some examples of those components as well as some of the factors that go into calculating the cost of the service.

Amazon Elastic Compute Cloud (EC2)—Server instances

Cost based on:

- Server configuration (OS, Memory, CPU)
- Time in use
- Number of instances

Amazon Simple Storage Service (S3)—Object-based storage

Cost based on:

- Amount of storage
- Data transferred out
- Number of requests against data
- Backups and redundancy

Amazon Relational Database Service (RDS)—Database instances

Cost based on:

- Database size
- Amount of storage
- Number of instances
- Number of requests

Amazon Elastic Load Balancing (ELB)—Load balancing

Cost based on:

- # of load balancers per hour
- Amount of data processed by load balancer

In addition to the technological components for server administration, Amazon also provides a several tools for managing infrastructure including APIs and a tool called Trusted Advisor, which recommends how to optimize service utilization and minimize cost. There are also many third party companies who are developing services like system migration and detailed billing and monitoring services. For a full list of Amazon

Table 2. New IT roles for the cloud.

Title	Role
Systems Architects	Staff who design environments to maximize use of the cloud while ensuring applications work in that environment
DevOps	Hybrid staff who understand both applications development and systems operations and can work as consultants for cloud deployment
Financial Engineers	Staff who can take requirements, estimate usage, and calculate potential costs in the cloud
Business Analyzers	Those who can take requirements and consider the hosting models as well as existing solutions to help answer the questions: Should we build or license, and should we host in the cloud or on premise?

services and a pricing overview, as well as a calculator to help estimate cost, see the following:

- <https://aws.amazon.com/products/>
- https://media.amazonwebservices.com/AWS_Pricing_Overview.pdf
- <http://calculator.s3.amazonaws.com/index.html>

Challenges and Organizational Changes

With every solution comes not only benefits but also challenges. For Amazon, these challenges include:

- Vendor lock-in—applications need to be architected appropriately to maximize Amazon's infrastructure and minimize cost
- Knowledge drain—the longer systems administrators are not exercising their traditional skill set, the more they lose knowledge for implementing systems locally
- The complexities of maximizing efficiency of resource utilization in Amazon—what is the right sizing for services, what can be turned off during non-business hours, and can applications be broken into smaller components to reduce large system utilization?

Moving into the cloud will undoubtedly save staff time, but there is a need for new roles to maximize the potential of the cloud. We have begun to identify several new emerging roles.

All of these new roles require a fundamental understanding of technology as well as critical soft skills for working across teams. Increasingly, the necessary roles are interdisciplinary and require flexibility in thinking as well as flexibility in placement within any organization.

Perhaps the most substantial barrier to cloud adoption is organizational culture. The nature of system administrators is to provide for maximal stability, and abstracting infrastructure can obfuscate one's core understanding of that technology and leave staff who work in fairly

traditional IT roles feeling uneasy about the future. It is important for organizations to have a strategy for working with staff on such a large logistical and mental transition. First, it requires building a foundation around why. Moving from on-premise to the cloud means sacrificing some flexibility to gain efficiency. It is important to develop a change conversation in the organization and introduce new perspectives. It is also important to challenge assumptions that may be based on old models of thinking. It is good to combat the tendency to ignore sunk cost; people often will talk about how much investment you have made as a means for rejecting change. Helping them understand that cost is cyclical and every significant change has upfront costs that are hopefully offset by the potential benefits is critical. The most important thing is investing in your staff both in terms of assurance and training. For staff, change is scary, especially radical change that appears to be replacing the value they provide for your institution. Building a plan to bring them along with the change and assuring them you are committed to their development and growth so that they can become a core part of the future will help assuage fear, especially as they see you making good on those promises to develop them. Moving into any new way of doing things requires real training and time to experiment with new solutions. Making a goal out of developing competencies in new technologies and investing in their learning will help mitigate demoralization and boost a sense of co-ownership for moving toward new solutions.

Recommendations

After a year of preparatory work and some experience working with Amazon Web Services, we have learned the importance of organizing the transition to the cloud. Ensuring that there is time to conduct the migration in a sane manner is critical for giving staff a chance to get acclimated to a new way of doing things. Giving staff a chance to experiment is also important; complex migrations may not be plug-and-play and how things work in the cloud may feel counterintuitive at times. Budget for experimentation and for making mistakes. With any new technology, it can take some time to fully understand it and the

implications of its use. Experimenting at a small scale can mitigate exorbitant costs while providing staff a sandbox to learn in. It is critical to not underestimate the value of change management in the process. Anxieties and demoralization can be tempered considerably through reassuring staff and providing professional growth so that they can clearly see themselves in the future of your organization. It is also important to get application developers involved and provide them with a strong understanding of how a particular cloud environment works. Optimizing applications

in the cloud will likely require intentional architecting of applications to work most effectively and efficiently in the cloud. And finally, develop an exit strategy. Who knows exactly what the next major change in infrastructure management will be or when it will come. What is certain is that there will be a new way to provide IT infrastructure more efficiently and effectively in the future at some point, and implementing solutions in a flexible and standardized way can help mitigate problems with future transitions.