

Bold Structures: Early Thoughts on Cloud Architecture for Higher Education

Original Editor: Oren Sreebny, University of Chicago

Original Authors: Satwinder Singh, Columbia University; Mark Personett, University of Michigan; Christopher Hubing, The Pennsylvania State University; Aron Bucher, University of Minnesota; Bob Winding, University of Notre Dame; Erik Lundberg, University of Washington; Brad Greer, University of Washington; William Allison, University of California, Berkeley; Steven Zoppi, Internet2; Sara Jeanes, Internet2; Oren Sreebny, University of Chicago

Beware of overconfidence; especially in matters of structure. – Cass Gilbert

Architecture is a dangerous mixture of power and impotence. - Rem Koolhaas

Never talk to a client about architecture. Talk to him about his children. That is simply good politics. He will not understand what you have to say about architecture most of the time. An architect of ability should be able to tell a client what he wants. Most of the time a client never knows what he wants. – Ludwig Mies van der Rohe

Introduction

Higher education institutions were among the first large enterprises to adopt software services in the cloud, especially email, calendaring, and other collaboration tools. Widespread adoption of cloud infrastructure by central IT units, however, is only starting to become a factor at US universities.

In the summer of 2014 IT practitioners from a number of leading US research universities gathered together for two days in Chicago to discuss creating high level strategy for use of cloud technologies in our institutions. The output of that work was published in the ECAR publication “Cloud Strategy for Higher Education: Building a Common Solution” (<https://net.educause.edu/ir/library/pdf/erb1413.pdf>). In the year since those discussions several institutions have made substantial progress towards implementation of significant deployments on cloud infrastructure, and others have begun planning for such implementations.

In that work it has become evident that careful thinking about the architecture of institutional presences in cloud infrastructure is necessary to avoid future large-scale chaos and creation of unmanageable situations which are neither secure nor sustainable. A number of institutions expressed interest in coming together again to discuss cloud architecture in some detail. In July, 2015, 23 people from thirteen institutions, Internet2, and Amazon Web Services gathered in Chicago for a two day discussion of architectural topics. This document is an attempt to distill and elaborate on those discussions.

We recognize that we are early on in implementation of large-scale use of cloud infrastructure in higher education. This document represents a snapshot in time of where some institutions are at present, and perhaps sheds some light on future directions. The document itself will undoubtedly prove to be ephemeral, but we hope that documenting the thinking of these institutions on these topics will be useful to others, and that the collaborative nature of the conversation and partnership between our institutions and the cloud vendors will continue indefinitely.

A Note About Cloud Vendor Participation

As noted above, technical staff from Amazon Web Services participated in this workshop. Many technical architecture topics are not necessarily vendor or platform specific, and the intent is to generalize the work of this group as much as possible. Most of the institutions represented have done initial cloud deployments in AWS, though several have also worked in additional environments.

Amazon's technical participation was much appreciated by the group, who specifically noted the willingness of the AWS staff to engage as technical peers without the presence of sales tactics during the workshop. Not only did they bring ideas for how AWS services might be used to address certain problems, but they also heard first-hand about the challenges we're facing and gaps in their service set. The spirit of collaboration is an important value in higher education IT, and the ability of vendors to take part in these discussions in that spirit is very much to their credit.

The group looks forward to future gatherings with providers of other cloud platforms.

Campus Introductory Remarks

A few of the institutions are notably ahead in the construction of cloud infrastructure architectures for building and hosting campus services and applications. Notre Dame, Columbia, and Harvard shared architectural documents that show not only the decisions they have made, but illustrate the technical concerns that are being addressed as the architecture is built.

Some common observations from the introductory remarks:

- It is critically important to properly define the setup of virtual private clouds (VPCs) at the beginning. It is difficult to change these later on.
- Many campuses have overly complex pairwise sets of relationships between services within the campus data centers. Designing cloud architecture offers an opportunity to step back and think about how sets of services should be located and deployed more simply.
- It is important to think about bridging on-premise resources with resources in the cloud. One example is applications on one side accessing a database server on the other.
- Planning for new kinds of security roles is critical! In an era where infrastructure is defined by code, an entire data center can be deleted with a single command.
- There may be opportunities for the creation of shared toolsets that could be used by multiple institutions.

The rest of this document will dig into details, use cases, and examples. While many of these grew out of the discussions, the material is not limited to what happened in those two days. It is intended that this be a living document and set of accompanying materials. If you are at an institution that would care to contribute, we welcome your participation.