ISSN: 2473-3857
v2 n4224

Teaching Case

Cloudy Decisions: How the Cloud Affects Infrastructure Decision Making

David M. Woods
woodsdm2@miamioh.edu
Computer & Information Technology Department
Miami University Regionals
Hamilton, OH 45011, USA

Abstract

The availability of "as a service" cloud offering for infrastructure, software and other technology products is having a significant impact on how IS/IT systems are developed. Decision making processes for acquiring infrastructure have been drastically changed by the availability of cloud offerings. This case presents a seemingly simple need for server and storage infrastructure to support a software tool being used in a data mining course at the college level and explores how cloud offerings may alter IS/IT decision making processes. The case initially explores physical infrastructure and similar options available through Amazon Web Services (AWS). The case then explores several additional use cases that are complex and possibly cost prohibitive with physical infrastructure, but simple to address using AWS. Finally, the case looks at typical funding sources available at a college and considers how well each works with physical and cloud hosted infrastructure.

Keywords: IT Strategy, cloud computing, infrastructure-as-a-service

1. INTRODUCTION

Cloud computing is a rapidly evolving topic that offers interesting opportunities and challenges for expanding the IT/IS curriculum (Lawler, 2011). Cloud computing offers abundant opportunities to engage students in accessing vendor information, reviewing discussions in the industry press, and exploring other resources to gather and evaluate information, a skill that students can expect to use throughout their professional careers.

In addition to the need to learn about cloud computing technologies and consider potential applications, cloud technologies require fundamental changes to strategy and decision making aspects of IS/IT management.

This case study is structured around an everyday problem of replacing the server supporting an application. Adding the potential of using cloud infrastructure, specifically Amazon Web Services (AWS) (Amazon Web Services, n.d.), rather than a locally hosted physical server, adds new dimension to the problem.

This case scenario starts with a simple question of acquiring infrastructure with an investigation of cloud infrastructure as an alternative for locally hosted physical infrastructure. This portion of the case is suitable for use with introductory level courses. The remainder of the case introduces more complex considerations that are more appropriate for upper level course.

While this case focuses on how the potential use of cloud technologies can shift decision making and other aspects of managing IT infrastructure, there are also several places in the case where technical activities could be included to give students hands on experience with AWS. These are discussed further in the teaching notes.

ISSN: 2473-3857
v2 n4224

2. BACKGROUND

Professor Smith has just completed teaching a new course on the topic of data analytics at a mid-sized university. Many of the course activities and assignments involved using a large software package to explore different techniques used in data analytics. Throughout the course, students worked with a variety of data sets using the software.

While preparing to teach the course, Professor Smith contacted the University Technology Services (UTS) division about a server to host the software needed for the course. UTS identified a Linux server that had been recently retired during the university's move to a cloud hosted e-mail system and made this available to Prof. Smith. The server is hosted in the university's data center, and with some assistance from Prof. Smith and online help forums, one of the UTS Linux system administrators was able to install the needed software on the server.

After teaching the course for the first time, Professor Smith met with the staff from UTS to discuss plans for teaching the course in the upcoming fall semester. The main concern is that the server being used had several hardware issues during the semester that meant that it was not available for most of a week during the semester.

3. HARDWARE REPLACEMENT

At the end of semester meeting with UTS, the discussion quickly focuses on the hardware reliability problems. UTS offers two options – Professor Smith can continue to use the current server with UTS performing maintenance as needed. UTS has several retired servers identical to the one being used by Prof. Smith that can be used as a source for spare parts.

The other option offered by UTS is to purchase a new server. Prof. Smith thinks this would be a better long term solution and would also minimize the risk of interruption to the course. UTS suggests that once Prof. Smith can provide a funding code, they can get a price quote from their preferred vendor.

Since UTS is not charging for the current server, the need to identify a funding source surprises Prof. Smith. The technology needed to support a course has always been something that was just provided when requested. Prof. Smith notes this as an additional task and makes a mental note to talk to the department chair about funding sources.

Asking for an estimate of the server cost results in a long complicated explanation from the UTS staff involving RAM, redundant power supplies, cores, processor families, STATA or SAS drives, RAID configurations, support levels, etc. and that Prof. Smith will need to specify all of these technical details before submitting a request for a quote. A ballpark cost of $5,000 - $15,000 is finally offered.

At this point, Robin, a recent college graduate who has just joined the staff of UTS speaks up to ask, "What about using the cloud?" Prompted by Prof. Smith, Robin explains that the cloud could allow Prof. Smith to use Infrastructure-as-a-Service (IaaS) offered by vendors such as Amazon Web Services (AWS) rather than purchasing physical hardware that would be configured and supported by the UTS staff.

Questions: Could AWS be a viable option for Prof. Smith's course? Explore the AWS case studies to learn about the benefits of using AWS (Case Studies, n.d.). Think about which of these benefits might apply to this situation. Explore the AWS web site to learn more about the Elastic Compute Cloud (EC2) (Elastic Compute Cloud, n.d.) and Elastic Block Storage (EBS) (Amazon Elastic Block Store, n.d.) for EC2.

Activity: Using the information that the current server has eight CPU cores, 32 GB of memory, and 2 TB of storage, develop a cost estimate for using AWS EC2 and EBS resources for the course. Consider the different EC2 pricing options (on-demand and reserved) (Purchasing Options, n.d.).

Activity: Use the AWS TCO cost estimator (TCO Calculator, n.d.) as another tool for exploring the costs of the cloud and physical server options.

Develop Scenarios: The future is uncertain, so developing several scenarios and exploring how well different solutions will work in these scenarios is a productive way to understand the risks and benefits of different proposed solutions.

In the recently completed semester, the class Prof. Smith taught had an enrollment of 24 students. More students were interested, but the computer classroom used only had seating for 24 students. Currently one section of the course is

being offered in the fall, with an enrollment limit of 24. The department is considering adding a second section in the fall and possibly teaching the course in winter and summer terms.

Develop scenarios for how often the course might be taught over the next year and the next three years. For each scenario, develop cost and resource estimates for both alternatives – local, physical infrastructure and use of AWS cloud infrastructure. In addition to actual dollar costs, consider other resources that might be needed. Can you think of additional information it would be useful to have to develop the scenarios? Can you build a model that would let you quickly estimate the cost for any scenario?

In addition to scenarios that assume the class will be taught regularly during the next three years, try to identify other scenarios to consider. For example, a local company partners with Prof. Smith to teach the course in a project based format where students work with the company's data and infrastructure.

Evaluate Scenarios: Select what you consider to be the most reasonable scenarios. Which approach (cloud or physical servers) works best for each scenario? Considering all of the scenarios, can you identify specific strengths and weaknesses of the physical infrastructure and cloud infrastructure approaches?

4. PROVIDNG FLEXIBILITY

To learn more about Prof. Smith's requirements, Robin schedules a follow up meeting. In this meeting, Prof. Smith discussed how the class introduced students to the data mining software. It is a complex package with a large number of add-on modules available. Prof. Smith wanted to ensure that the complexity of the package didn't overwhelm students as they were getting started in the course

To accomplish this goal, Prof. Smith started the semester having students work with a very basic version of the package as they learned to navigate the tool and make use of basic data mining concepts. Prof. Smith's original plan was then to update the software installation to introduce additional modules one at a time.

However, this plan had to be revised. When Prof. Smith contacted UTS about installing additional modules, this was seen as a low priority task and the installation and configuration of an additional module took over a week to complete. In addition, since the UTS Linux system administrator was not familiar with the data mining software, Prof. Smith had to walk the system administrator through the installation and configuration process.

Prof. Smith suggested that since he was familiar with the software and the process for installing and configuring additional modules, he could handle these tasks. Since these tasks would require system administrator (known as root) access on the Linux server, UTS was not willing to permit this. Root access could allow Prof. Smith to make changes to other aspects of the server. If these changes rendered the server inoperable, this would impact the course and might require an immediate response from the UTS Linux system administrator. As a result, Prof. Smith decided to have all of the modules he might want students to work with installed during the third week of the semester.

Questions: Review the AWS documentation and learn more about how users can create and customize Amazon Machine Instances (AMIs) and then select which AMI to use when starting and EC2 instance (Amazon Machine Images, n.d.). Could the use of AWS instead of a physical server offer different solutions to the software installation and configuration problem? Also, how would the use of AWS change the risks associated with allowing Prof. Smith to install and configure the data mining software?

During the follow up discussion with Prof. Smith, Robin also asked if there were other ways in which the current server configuration was limiting what could be done in the course. Prof. Smith immediately mentioned that it would be useful if students could get experience with installing and configuring the data mining software.

This would have several benefits. First, the student would be better prepared for a job opportunity where an organization was just starting to explore the potential of data mining and did not have established infrastructure to support data mining. Students with actual experience installing and configuring a data mining software package could be better prepared for this type of opportunity.

The ability to install and configure additional modules would provide more flexibility as

ISSN: 2473-3857
v2 n4224

students worked on their final projects in the course. Since the data mining software supports a large and ever expanding set of modules, including user contributed modules, student projects wouldn't be limited to working with the modules setup at the beginning of the semester. Allowing students to modify configurations could give students experience tuning the application for specific uses.

Questions: The initial discussion assumed one server would be used to support all students in a course section. How reasonable would it be to provide an individual server for each student? Consider both the option of using physical infrastructure hosted by UTS and cloud infrastructure from AWS. In addition to cost, consider other resource needs such as time spent by the UTS staff.

Activity: Explore the AWS Educate program and see whether this program could provide students direct access to AWS resources (AWS Educate, n.d.).

Questions: Assuming each student could be provided their own infrastructure, how could the desire to allow students to gain experience with installing and configuring the software be supported? Again, consider both physical and virtual infrastructure options. What risks would be associated with allowing students to install and configure the software? Specifically, consider what steps would need to be taken to recover from a situation where a student error renders a physical or virtual server unusable.

5. EXPLORING MORE TOOLS

The discussion between Robin and Prof. Smith also included a discussion of possible future plans for the course. The first concern that Prof. Smith mentioned is that while the class is currently being taught using a specific software package, the main focus of the course is current topics in data mining, not a specific software package, so the ability to teach the course with a variety of software packages is important to ensuring that course content could evolve as data mining techniques continued to develop.

Prof. Smith mentioned a number of ideas for developing the current course along with the idea of developing an advanced course in data mining to explore advanced topics and additional tools. In addition to exploring additional topics, the advanced course could also build upon knowledge gained in the initial course to allow students to work with larger, more complex data sets and explore student developed questions as part of a project type assignment.

Prof. Smith also mentioned that a few students from the recently completed class had already asked about the possibility of doing more extensive work with data mining as part of their yearlong senior capstone projects.

Questions: Revisit the previous discussions about locally hosted, physical infrastructure versus the use of AWS cloud hosted infrastructure. With each approach, what would need to be done to accommodate the idea of using a variety of tools and exploring larger data sets in an advanced data mining course? How could capstone projects involving data mining be supported by physical or cloud infrastructure?

Update Scenarios: Develop additional scenarios based on the potential for developing an advanced data mining course and enabling student exploration of data mining in capstone projects. Summarize the strengths and weaknesses of the two potential approaches. Could a mix of physical and cloud infrastructure be a useful option for supporting the expanded set of requirements?

Additional Cloud Options: Explore the AWS Marketplace and search for available AMIs that could be useful for a data mining course. Note that there are several different pricing plans for AMIs in the AWS Marketplace. These include free, hourly costs, monthly costs, and BYOL (bring your own license). AWS offers an entire category of products for "Analytics" (AWS Marketplace, n.d.). Explore these offerings and identify any that could be used in a data mining course or capstone project. Consider what level of physical infrastructure and support services would be needed to offer the same range of options using locally hosted resources.

6. COST ASPECTS OF DECISION MAKING

After working with Robin to explore the different options available for using locally hosted physical infrastructure and AWS cloud infrastructure, Prof. Smith now has a much better idea of the options available and potential costs.

As discussions about the different options for providing infrastructure to support the data mining course have been ongoing, Prof. Smith

ISSN: 2473-3857
v2 n4224

has been looking for potential funding sources. Matching potential funding to the different infrastructure options presents challenges.

Several funding sources such as grants for course development provide a one-time source of money. These are a good match to the physical infrastructure option with large up-front costs to purchase and configure a server. There would be smaller ongoing costs for maintenance and support, but there are options to include these in the initial purchase cost.

Other funding sources such as the university's technology fee or a per course lab fee would work better for using AWS cloud infrastructure where costs are charged on a periodic, subscription like plan. Additionally, the costs for teaching a specific course section will have a direct relation to the number of students enrolled in the course.

Explore your options: Suppose this scenario was occurring at your school. Investigate the options available for funding technology for a course like the data mining course Prof. Smith is teaching. Remember to consider the AWS Educate program. Could having your school join the AWS Educate program be helpful?

In addition to the options discussed above, what funding options can you identify? Which infrastructure model (physical or cloud) would work best for the funding sources you identify?

Put it all together: Review the options for the data mining course and the potential funding sources you have identified and develop a proposal for Prof. Smith to submit. Include the options that will give Prof. Smith the most flexibility for continuing to develop the course and providing the biggest benefit to students within any constraints imposed by the available funding options.

7. CONCLUSIONS

This case explores how the availability of cloud based infrastructure and software changes the decision making processes for acquiring technology infrastructure.

Cloud services like AWS provide flexibility and elasticity that allow infrastructure capacity to quickly respond to changing needs. Cloud infrastructure can easily expand and contract and needs change.

AWS infrastructure also decouples the software on a server from the actual server infrastructure through the AMI that can be made available in a known state. Additionally, AWS offers the ability to access a wide range of existing AMIs, allowing users to quickly get up and running with a software package.

Cloud services also alter financial decisions about acquiring infrastructure. Physical infrastructure typically has large up-front acquisition costs. Usage costs must be estimated using long range estimates – for example projected enrollments in a specific course for the next several years. Cloud infrastructure costs are incurred as the resource is used and can often be structured to scale with usage. For example, if course enrollments decrease, infrastructure costs also decrease.

Considering the wide range of options enabled by the flexibility and elasticity of cloud infrastructure and matching these to the specific funding options available in an organization is a complex process that many organizations will face as the use of cloud infrastructure continues to grow (Kim, 2015).

8. REFERENCES

Amazon Elastic Block Store (EBS) – Block Storage. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/ebs/.

Amazon Machine Images (AMI). (n.d.). Retrieved May 23, 2016, from https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html.

Amazon Web Services (AWS) - Cloud Computing Services. (n.d.). Retrieved May 23, 2016, from http://aws.amazon.com/.

AWS Educate. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/education/awseducate/.

AWS Marketplace: Find and Buy Server Software and Services that Run on the AWS Cloud. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/marketplace/.

Case Studies & Customer Success. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/solutions/case-studies/.

ISSN: 2473-3857
v2 n4224

Elastic Compute Cloud (EC2) Cloud Server & Hosting – AWS. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/ec2/.

Kim, E. (2015). This one chart gives you an idea of how crazy Amazon's cloud growth really is. Retrieved May 23, 2016, from http://www.businessinsider.com/amazon-aws-growth-rate-is-far-outpacing-other-enterprise-vendors-2015-12.

Lawler, J. (2011). Cloud Computing in the Curricula of Schools of Computer Science and Information Systems. *Information Systems Education Journal*, 9(2), 34-55.

Purchasing Options. (n.d.). Retrieved May 23, 2016, from https://aws.amazon.com/ec2/purchasing-options/.

TCO Calculator. (n.d.). Retrieved May 23, 2016, from https://awstcocalculator.com/.