

Seeding the Cloud with Students from Every Computer Science Undergraduate Program

Joel Hollingsworth and David J. Powell

Department of Computing Sciences, Elon University, Elon, North Carolina, U.S.A.

Keywords: Cloud Computing, Platform as a Service, Enterprise Programming, Server Side Programming, IBM Smartcloud Enterprise, Google App Engine.

Abstract: The introduction of enterprise-level cloud computing to the computer science curriculum is necessary to remain consistent with the industry. Cloud computing is rapidly being used throughout industry and there is a current and projected shortage of IT developers with cloud computing skills. University computer science curricula need to provide cloud computing courses to prepare graduates to enter the work place. This paper describes a required cloud based enterprise-level programming course at Elon University that used free cloud based software and hardware and free training materials. This cloud based development course provides an example that could be immediately used at any university to give their students the cloud based skills needed to be successful in the IT market.

1 INTRODUCTION

The two major survey results released by IDC (International Data Corporation) and IBM in December 2012 indicated a severe talent shortage in cloud computing skills in the IT workplace (Helms, 2012; IBM, 2012). The Microsoft commissioned IDC surveyed 600 hiring managers from around the world. The study predicted the IT market to grow roughly 1.1 to 2.7 percent through 2020 and cloud related skills to be the primary growth opportunities in IT employment. The surveyed managers indicate that finding people with the right skill mix of cloud skills will be the #1 challenge for companies as the work force around the world lack cloud computing skills. The IBM survey had two parallel portions with one directed to 1200 IT and business decision makers and the other to academia consisting of 450 students and 250 faculty members from around the world. The IT decision maker results indicated a 25% major skill gap in hiring cloud computing skills and a 60% moderate to major skill gap. The Academia results are even higher with 46% reporting a major skill gap and 78% a moderate to major skill gap. Only 6% of graduates have the skills needed in cloud computing to be successful. The IBM survey indicated that Academia is moving slowly to shrink the existing gap with 31 percent of institutions having no plans to teach cloud

computing and 20 percent planning to teach cloud computing in the next two years. The biggest obstacles are cost and lack of access to training.

The hypothesis of this paper is that all undergraduate CS programs should be immediately offering a core course in cloud computing. Elon University has taught a core course in cloud computing since 2010. There are no obstacles involving cost as both hardware and software are free. There are no obstacles in training as training resources are also freely available. The only requirement is a faculty member with the desire to learn and teach cloud computing. This paper discusses a core course in cloud computing at Elon University that other institutions may emulate to get started. The acquisition, the use and the advantages of free state of the art hardware and software are enumerated. The course curriculum is then presented and discussed in terms of format, free instructor training, student materials and assessment. Finally, the paper concludes with a brief discussion of a follow on elective cloud computing course that builds off of faculty expertise teaching the core course and future directions.

2 CORE COURSE ON CLOUD COMPUTING

Google is a fast moving leader in web innovation and open standards. Starting in 2008, Google has held an annual Google IO developer conference at the Moscone Center in San Francisco. Each annual conference has started with a keynote address by the Google Senior Vice President of Engineering, Vic Gundotra, in which he presents a high level overview of Google's vision of the web and the key trends and tools that they are encouraging developers to use to move the web forward. Two Elon University faculty attended the 2009 annual conference and were immediately impacted by the keynote address that the "web has become the dominant platform of our era" and that since 2004, there have been no popular desktop applications developed other than games and browsers used by millions of users. Application development had clearly moved to web applications such as Gmail, Facebook and Wikipedia. Three key trends and tools described at the 2009 conference were increased user access to cloud computing, increased client functionality with HTML5 and increasing numbers of mobile devices accessing the web. After attending the 2009 conference, Elon University changed the content and description of a junior level, required Computer Science III course to be based on enterprise programming using a PAAS Cloud. The course has been successfully taught and continually updated based on industry trends for the past three years and has proven to be spot on in preparing the students with the skills needed by employers as shown in the IDC and IBM 2012 surveys and recent data provided at the Google IO 2012 conference that the App Engine cloud had 1,000,000 active applications with over 7.5 billion page hits daily (Magnusson, 2012).

The Elon Computer Science III course is a junior level course and students entering the course have already had two programming courses using Java SE and are familiar with the Eclipse Interactive Development Environment. The enterprise programming course was designed to focus on the most recent industry client standards of HTML, CSS and JavaScript, the most recent Java EE 6 JavaServer Faces standard and the SQL standard for creating, retrieving, updating and deleting data from a persistent data store. Students learn how to develop, test and debug on the cloud using three separate processes representing the client process, the server process and the datastore process. The use of multiple languages and multiple processes is a

significant increase in complexity for the student. PAAS allows the course to focus on development and not on system administration and configuration.

3 PAAS HARDWARE AND SOFTWARE

This section describes the goals, acquisition, software, configuration and advantages of PAAS for Computer Science III.

3.1 Goals

Elon University wanted to offer the previously described course using state of the art software and hardware. Elon University identified the following goals for the platform:

1. The client process, server process and data store process must run on different physical or virtual machines. The client should have the latest Firefox and Chrome browsers. The server should support a fully compliant Java EE 6 Application Server and the database should support the SQL standard and be network accessible.
2. The software and hardware must be available 24x7 and accessible from both on an off campus with no firewall restrictions.
3. The software must be pre-installed and not require any installation, system administration or configuration by the student or by the instructor.
4. The software must be widely used in industry and recognized as an industry leader.
5. The ability to create user accounts for students should be under the control of the faculty member, easy to do with a short turnaround time less than 24 hours from request to receipt of user account.
6. Each student account should be in a sandbox and not visible by other members of the class.
7. The IDE must be based on Eclipse and support development, testing and deployment from within the Eclipse environment.
8. The software and hardware must be free or at a very low price.

For 2010 and 2011, the Google App Engine came the closest to meeting all of these goals. It is a superb PAAS offering and was very successfully used in the Computer Science III course (Hollingsworth, 2010). The only minor limitations were:

1. It did not fully support Java EE 6. JavaServer

Faces 2 (JSF) was not provided with App Engine. However, there were workarounds that allowed third party libraries to support most of JSF 2. Threads were not supported.

2. It did not have free support for SQL but rather supported an underlying NOSQL Google Big Table database with a partially compliant JPA 1.0 interface.
3. The local testing environment was a subset of the deployment environment. Students could have code run in a local Jetty environment but fail when deployed to App Engine.

Elon University would still be using App Engine today if IBM did not provide us with a better PAAS through the IBM Academic Initiative and the IBM SmartCloud Enterprise in 2011. This environment was successfully used in 2012 and fully meets all of the course goals. The IBM environment will be the focus for the remainder of this paper.

3.2 Acquisition

For many years, IBM has offered universities an IBM Academic Initiative program that provides access to download and freely use course materials and software. Faculty can quickly register and gain access through a simple online application. A drawback of this program was that the faculty member had to download, configure and administer the software on university hardware. However, starting in 2010 with 20 institutions and expanding in 2012 to 80 institutions in 17 countries, IBM added the ability through the Academic Initiative program for faculty and students to gain access to the IBM Academic Skills Cloud (IBMa, 2012). The Academic Skills Cloud provides access to the IBM SmartCloud Enterprise and allows faculty and students to create virtual machine instances from over 20 pre-configured IBM images running the latest versions of IBM software such as Rational Application Developer and WebSphere

In 2011, Elon University became part of the pilot program on the IBM Academic Skills Cloud. After using the Google App Engine PAAS for two years, Elon faculty knew exactly what they wanted and needed in a PAAS from IBM. The IBM Academic Initiative Relationship Manager worked with Elon faculty and had an IBM image created and configured with the software needed to specifically meet the goals of the Computer Science III course. Elon faculty had the ability to use a simple IBM SmartCloud Enterprise web form to request the creation of a virtual machine instance of this image running on the IBM cloud. Within minutes after

submission, the virtual machine is created. A key aspect of the cloud instance is that the development environment is on the cloud and the deployment environment is on the cloud.

3.3 Virtual Machine Software and Configuration

For the Computer Science III course, faculty members created a separate, individual virtual machine for each student on the cloud running preconfigured software to provide a PAAS for developing enterprise Java EE 6 applications. Each student had their own machine IP address with associated login and password for developing and deploying enterprise web applications. All of the needed course software was pre-installed and configured to eliminate any need for software installation or configuration by each student. Each student's virtual machine was identical so faculty knew each machine was properly configured and could focus on the application and not installation related issues. The student could access the virtual machine through Remote Desktop Connection, bring up Rational Application Developer and immediately start development. Remote Desktop Connection is a software program that allows the students to access and develop code on the cloud virtual machines with the impression that they are developing on their own computer. Remote Desktop Connection is free and comes preinstalled on every Windows machine. For Apple machines, Remote Desktop Connection is automatically installed with Microsoft Office 2011 for Mac.

The IBM cloud offers a choice of a Linux or Windows Virtual Machine. Elon chose a Windows Virtual Machine since the Elon students were most familiar with Windows and the ease of access to the virtual machine using Remote Desktop Connection. Each virtual machine has 60 gigabytes of disk space. The software pre-installed on each machine included DB2 Database, Derby Database, WebSphere Application Server (WAS), Rational Application Developer (RAD), and Firefox. The DB2 database software and Derby database software fully support SQL. The lightweight Derby database was the focus in Computer Science III as it is simpler to use in learning database skills. The DB2 database is provided with the anticipation of using it in follow up elective courses and student projects. The DB2 database automatically starts upon machine startup. The Derby Database has a desktop shortcut for the user to manually start. WebSphere is one of the leading fully compliant Java EE 6 Application

Servers. To get a sense for its industry acceptance, there are more job opening requiring WebSphere skills on Indeed.com as shown in Figure 1 for five consecutive years than any other Java EE 6 application server (e.g. JBoss, WebLogic or App Engine). WebSphere is installed to automatically start at Windows start up. Rational Application Developer is an Eclipse Based Interactive Development Environment. For the second consecutive year, it has been rated the Best Software Development Platform in the annual Evans Data Corporation survey (Taft, 2012). RAD is installed and has a desktop shortcut for easy access. RAD has been preconfigured to start up with the Java EE perspective with two available WebSphere servers to which RAD can deploy an enterprise application project with a simple mouse click. One WebSphere server is running on the current cloud virtual machine. The other WebSphere server is running on a course deployment machine in the cloud. Two WebSphere servers were provided to give the student the chance to develop and test locally and to then deploy to a course deployment server. RAD has had the data perspective configured with database connections to a DB2 database and Derby database on the same virtual machine and connections to a DB2 database and Derby database on the course deployment server. The preconfigured connections facilitate the development and testing of SQL code to create, read, update and delete tables.

3.4 Advantages

The IBM SmartCloud Enterprise PAAS provides the course instructor and student with many advantages. It provides free hardware and software and removes the largest obstacle of cost reported in the IBM survey. It does not require any software installation or configuration which allows the student to focus on the challenging task of learning enterprise programming and not becoming frustrated with many lost hours due to an IT configuration issue. It is available 24x7 from on and off campus. Students have their own virtual machine so any mistakes made by others will not impact them and any mistakes they make will not impact others. It is a Windows virtual machine and when accessed from Remote Desktop Connection will have the same look, feel and familiarity of their own laptop. The WAS Application Server and RAD IDE are industry leaders. The RAD environment provides many wizards, editing and development tools that aid a student to develop code in SQL, HTML, CSS, JavaScript and Java within a single IDE.

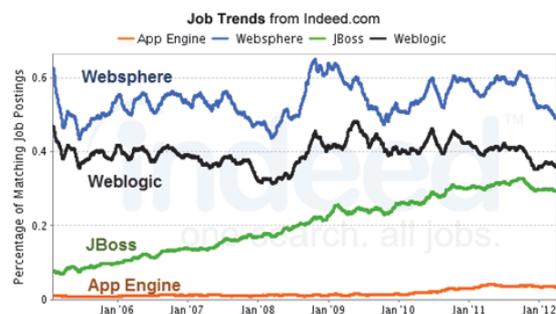


Figure 1: Current job postings for Java EE server skills.

4 CURRICULUM FORMAT

The course was a very fast paced course and covered a lot of concepts and languages. The fall 2012 semester course was held three times a week for approximately 14 weeks. The course was structured into a client segment followed by a server segment followed by a data persistency segment. The client segment was the first one-sixth of the course (8 classes) and covered HTML 5, CSS and JavaScript. The intent was to quickly cover these languages as they would be re-used and reinforced during the server segment of the course. The server segment was the largest segment of the course and took the next four-sixths of the course (20 classes). The server segment covered JavaServer Faces and Primefaces. The final segment for the last one-sixth of the course (8 classes) was the data persistency segment and covered database design and programming using SQL and the JDBC API.

5 INSTRUCTOR TRAINING

According to the IBM Survey, the second leading obstacle to introducing cloud computing is lack of access to training. There are many sources of free training available to faculty that can be subsequently used to provide many of the student materials for a course. This section discusses the free training resources used by Elon faculty to learn the materials to offer the class.

The client segment of the course consisted of HTML, CSS and JavaScript. There are many free web accessible resources available for learning these technologies. Elon University has free access to tutorials at Lynda.com. The faculty used three videos titled "HTML Essential Training", "CSS for Developers" and "JavaScript Essential Training" supplemented by tutorials at W3Schools on HTML,

CSS and JavaScript to learn these languages. Each of the video tutorials is approximately 5 hours in length. If your university does not have access to Lynda.com then you can take these three tutorials using the free one week Lynda trial or pay 25 dollars for a one month membership.

The server segment of the course, which is the largest segment, highly leveraged the outstanding tutorials provided by Marty Hall. Marty is the president of a training and consulting company focusing on Java EE, Rich Internet Applications and Android Development (Hall, 2012). In this role, Marty teaches on site courses around the world to major corporate and government customers in many disciplines (e.g. Energy, Healthcare, Finance, and Automotive). Each year, he typically teaches at least two courses on Java Server Faces and two courses on JavaScript and Ajax Basics. Each course contains a complete set of PowerPoint slides, Eclipse project code of examples used in the slides, course exercises, Eclipse course exercise solutions and links to key JSF references. Marty makes all of this material freely available to course instructors. After each course presentation, Marty updates the materials and format of the course to make it better. Elon University faculty have used the continually updated JSF and JavaScript course materials for three years to learn JSF 2, Primefaces and JavaScript. Marty's code is extremely well written and was ported to run in RAD and WebSphere without any modifications.

The database segment of the course leveraged the 14 article tutorial series on the Derby database by Robert Brunner (Brunner, 2006). Each article in the series contains downloadable code and builds off previous articles. SQL is sequentially covered starting with using it in Derby as a data definition language to create tables to using it as a data manipulation language for writing queries. After covering SQL, Brunner discusses JDBC and the design concepts of using a Data Access Object. The JDBC section is nicely complemented by the materials on "Accessing Databases with JDBC" provided by Marty Hall as a module in coreservlets.com's Java 7 course. Two additional web sites that proved valuable in learning SQL were sqlzoo.net which has interactive tutorial exercises and w3schools.com.

6 STUDENT MATERIALS

For the most part, the same materials used for instructor training were used for student learning.

The first client segment of the course used the same Lynda.com videos and on-line resources and was supplemented by instructor developed PowerPoint slides.

There is only one required text book for the course. The text book is Core JavaServer Faces Third Edition and is used for the second segment of the course (Geary, 2010). The text book is superbly complemented by the Marty Hall JSF slides. In fact, in the Marty Hall on line references, he lists the Core JavaServer Faces book as his favorite JSF reference. The text book stresses good programming practices and continually builds on previous practices in each example. Instead of following the text book chapter by chapter, Elon followed the topic order used by Marty Hall in teaching his JSF 2 course and mapped the reading assignments in the book to each topic. The slides provided by Marty Hall were customized and slightly enhanced for RAD and WAS.

The final database segment of the course used the Derby tutorial series by Brunner which contained 14 articles and was supplemented by instructor developed slides.

7 ASSESSEMENT

Each of the three course segments was assessed by a comprehensive exam at the end of each segment and with multiple homework assignments. The exam at the end of each segment was developed in a multiple-choice format similar to that used on certification exams and comprehensively tested the segment topics. During each segment, multiple homework assignments were assigned. For each segment, students used pairwise programming. Pairs were randomly assigned for the segment and then changed for the next segment. Pairwise programming was not used in 2010 and 2011 but was very successfully used in 2012 and proved to increase productivity, increase quality, reduce frustration, increase knowledge sharing, develop interpersonal skills and increase confidence. Elon has found pairwise programming to be extremely beneficial in enterprise programming. The complexity of developing in multiple languages and multiple processes is significantly increased over simple Java application development in Computer Science I and Computer Science II. Pairwise programming enabled students to more effectively deal with the complexity. The average homework assignment grade using pairwise programming increased from 85 in 2011 to 91 in 2012.

8 CONCLUSIONS AND FUTURE DIRECTIONS

Cloud based computing is rapidly being used throughout industry and there is a current and projected shortage of IT developers with cloud computing skills. University computer science curriculums need to provide cloud computing courses to prepare graduates to enter the work place. This paper described a required cloud based enterprise programming course at Elon University that used free cloud based software and hardware and free training materials. This cloud based development course provides an example that could be immediately used at any university to give their students the cloud based skills needed to be successful in the IT market.

The computing field is extremely dynamic and constantly undergoing rapid change. Elon University will keep up with the rapid change by continually monitoring the Google IO annual conference and Industry trends to update the core Computer Science III course and to use the core course as a building block for new elective offerings. Near term enhancements to the core course will be to use the upcoming, 2013 Java EE 7 release which has an upgrade to JSF to better support HTML 5. In addition, an elective cloud computing course will be offered in spring 2014 focused on HTML 5, Servlets 3.1, Restful Web Services and NOSQL data persistence.

REFERENCES

- Brunner, R. (2006) 'Developing with Apache Derby' [Online] Available from: <http://www.ibm.com/developerworks/opensource/library/os-ad-trifecta1/index.html> [Accessed 12/03/2012].
- Geary, D. and Horstmann, G. (2010) Core JavaServer Faces. Third Edition. Prentice Hall.
- Hall, M. (2013) [Online] available from: <http://www.courses.coreservlets.com/about-instructor.html>. [Accessed 12.20.2012].
- Helms, R. (2012) IT Workforce Can't Meet Demand For Cloud Skills. CIO. [Online] available from: http://www.cio.com/article/724335/IT_Workforce_Can_t_Meet_Demand_For_Cloud_Skills [Accessed 12.10.2012].
- Hollingsworth, J and Powell, D. (2010) 'Teaching Web Programming Using the Google Cloud'. Proceedings of the 48th ACM Southeast Conference.
- IBM. (2012) Academic Skills Cloud. [Online] available from: http://www.ibm.com/developerworks/university/cloud/asc_index.html [Accessed 12.10.2012].

IBM. (2012) 'Fast Track to the future, The 2012 IBM Tech Trends Report' [Online] available from: <http://ibm.com/developerworks/techtrendsreport> [Accessed 1.12.2013].

Magnusson, P. and D'Alesandre, G. (2012) 'Google I/O 2012 - App Engine Overview' [Online] available from: <http://www.youtube.com/watch?v=uy0nALQEAM4> [Accessed 1.15.2013].

Taft, D. (2012) 'IBM Rational Repeats as Top Dev Tools Platform: Survey' [Online] available from: <http://www.eweek.com/c/a/Application-Development/IBM-Rational-Repeats-as-Top-Dev-Tools-Platform-Survey-135226/> [Accessed 1.12.2013].