Google and IBM Clouds Make Enterprise Computing Available to all Undergraduate CS Majors

David J. Powell and Joel Hollingsworth
Department of Computing Sciences, Elon University, Elon, NC, USA

Abstract - *Cloud platform as service offerings (PAAS) from Google and IBM make it possible for all universities (large and small) to offer state of the art curriculum for enterprise level, server side software development. Google and IBM provide free access to the same hardware and software used by their corporate customers for software development and hosting of server side applications. Universities can easily and safely leverage the cloud to offer courses on enterprise computing to make their graduates more competitive in the market place. The Google cloud (App Engine) has been successfully used at Elon University for two years and the IBM cloud (SmartCloud Enterprise) for one year to teach Java Enterprise Level Programming of Servlets, JSP and JSF. This paper describes the advantages of using the IBM Smart Enterprise cloud as a platform for a course on Servlets 3.0 and JSP 2.2.*

Keywords: Cloud Computing, Enterprise Programming, Server-side Programming

1 Introduction

The computing field is constantly changing. Computer Science departments are challenged to continually analyze changes in the field to provide students with a curriculum that provides the key foundational concepts to support a career of lifelong learning while using state of the art software tools, frameworks and packages for the students to be competitive in the job market for internships and jobs after graduation. The rapid emergence and widespread use of mobile and cloud based computing where information is available to anyone, anywhere and at any time is changing the way we live, work, play and learn. Graduating students need to have a solid foundation in mobile and cloud computing to be competitive in today's marketplace [1]. Elon University made cloud based, server side programming, a part of the core curriculum in 2010. For the first two years (2010 and 2011), the department offered a core and an elective course in enterprise level computing using the Google cloud (Google App Engine). In 2012, the Google cloud was replaced with the IBM cloud (SmartCloud Enterprise). The enterprise level programming courses have enhanced student resumes and made the graduates competitive in the marketplace. Every graduate for the past three years has received at least one job offer and all rising seniors have accepted job internships for the summer.

The Google and IBM clouds are easy to use for teaching enterprise computing, offer students a competitive advantage in the job market and can be immediately integrated into an undergraduate curriculum of any size. This paper supports this hypothesis in four sections. The first section identifies the skills needed by students in enterprise computing and the challenges faced by a small university in offering a course to provide these skills. The second section describes the advantages and disadvantages of using the Google cloud as the platform for the course. The third section describes the configuration and use of the IBM cloud to provide an improved platform over an already excellent Google cloud. The fourth section provides details of the course delivery at Elon to provide readers a benchmark to use for developing a similar course. Finally, the paper concludes and identifies some future directions.

2 Curriculum

Ten years ago, most software development was for stand-alone applications running on a desktop machine with a single processor. There were many jobs for a software developer that was well skilled in a single language such as Java, C#, C++ or Visual Basic. Core computer curriculums were structured to support this single machine model. Students learned how to design, develop, debug, test, and maintain software for a single machine. However, during the past decade, the web platform has become the dominant development platform [2]. Developing for a web platform requires developers to be skilled in software development for multiple computers, multiple processes and multiple languages. The student needs to understand how to design, develop, debug, deploy and test an internationalized application that involves the client (either a mobile app or browser application), the web application server and the database server.

The number and types of languages and frameworks for enterprise programming is rapidly changing and will continue to change. At Elon University, we have selected mainstream, industry standards and software packages that implement the standards to provide the student the background to quickly learn other languages and frameworks while making them

attractive to potential employers in the job market. We have selected Java as the foundational language for computer science majors and Eclipse as the primary IDE. As a minimum, we expect every graduate to have core skills for the client browser of HTML 5, CSS and JavaScript. For the web application server, we have selected Java EE 6 which is used to develop core skills in Servlets, JavaServer Faces, JavaServer Pages, JPA and JDBC. For the database tier, we expect all graduates to have core skills in SQL.

The field of Computing Sciences continually changes at a rapid pace. Computer Science Departments cannot simply add an additional course on cloud computing and purchase the hardware, software and technical support needed to support it. Departments have a fixed set of resources and constraints that must be optimized to provide maximum benefit to the student both in the short and long term. Some of the fixed set of resources and constraints are curriculum guidelines, budget and availability. Each of these constraints will be briefly discussed along with how Elon has made tradeoffs to add cloud computing into the curriculum.

2.1 Curriculum guidelines

Currently, there are over 265 ABET accredited undergraduate computer science programs. Each of these programs has met the accreditation constraints of having at least 40 semester hours in computer science and 30 semester hours of science and mathematics [3]. Within the 40 semester hours of computer science courses, programs are required to have coverage of algorithms, data structures, design, programming languages and computer organization and architecture. The total of 70 semester hours is a lot of hours and cannot be simply increased without impacting the students' ability to satisfy general education requirements at the university as well as pursue additional minors or majors. If a department is going to add cloud computing then it will come at the expense of removing other important topics from the 40 semester hour of computer science courses. Every item in the curriculum is very important and making relative tradeoffs based on perceived importance is difficult and is made based on the educational objectives for the graduates from each university.

Elon University is on a 4 credit hour system and each Computer Science major must complete 44 hours (11 courses) of Computer Science courses to graduate. Eight of these eleven courses are designated as core courses that each student must take to graduate. Elon considers enterprise level, server side programming to be a fundamental core course required of all students and has made the course called Computer Science III part of the core eight courses. The Computer Science III course has a prerequisite of Computer Science I and II where the student learns problem solving using Object Oriented Design and Programming with Java SE in an Eclipse environment. The Computer Science III course is positioned to be taken during the fall semester of the junior year to provide students the knowledge to be competitive and productive in an internship during the summer between their junior and senior years. The Computer Science III course is a standards based course and is based on Java EE 6. The course content covers client, server and database machine basics. After Computer Science III, students can design, develop, debug, deploy and maintain a multiple tier web based application with:

- a client browser using HTML 5, CSS, JavaScript
- a web application server using servlets 3.0, JSP 2.2, JSF 2.0, JPA 2.0 and JDBC
- a database server using an SQL database.

2.2 Budget

The economic crisis has frozen or reduced the department budgets at many universities. There are many competing requests for the fixed budget. The chances of submitting a successful request for server hardware, server software, database software and system administrative personnel to support cloud computing would be difficult. Moreover, even if the request is funded then most likely it would not be at level to match the free PAAS offerings on the Google and IBM clouds. Both of these clouds provide universities with free hardware, free software with no maintenance fees and free system administration. The free hardware is accessible 24 x 7 from anywhere on or off campus, has no firewall restrictions, has ample storage and excellent performance. The software is upgraded on a routine basis and the faculty member can easily apply/use the upgrades at a convenient time either during or between semesters. Both clouds allow users to be added with no need for a third party to create user ids and passwords. Google allows the user to directly request a user account and receive it within seconds via a confirmation password on their mobile phone. IBM gives the faculty member a GUI application to easily request a virtual machine instance with a username and password. The instance is then queued for creation and available within minutes.

2.3 Availability

Students have different class schedules and a variety of different constraints from on campus and off campus activities such as internships, sororities, student organizations and sports. An ideal computing environment would be available to the students at any time from any place. Students should not need to be within the university firewall or in a certain lab to develop their cloud applications. If a server goes down during the week or weekend then personnel should be available to immediately address the problem. At a small campus where system administrators have multiple job functions and support multiple departments, 24 x 7 access, service support and uptime is not pragmatic. However, both Google and IBM offer

24 x 7 access and have no firewall restrictions. In three years of using the Google cloud and one semester of using the IBM cloud there has been no measurable downtime.

3 Google cloud (Google App Engine)

Elon has used the Google cloud to teach the Computer Science III course on Enterprise Computing in the Fall 2010 and Fall 2011 semesters. In addition, the Google cloud was used to teach an Elon elective course on Servlets 2.5 and JSP in Spring 2010 and a SIGCSE mobile computing - cloud workshop in Spring 2011. The Google cloud has proven to be a superb cloud computing platform for the Elon course offering with many advantages. It is free, easily accessible through an Eclipse plugin for both Apple and Window machines, has generous quotas, requires no server system administration, has automatic backups, provides database connections with JPA or JDO, supports most of the servlet, JSP and JSF standards and has both an active development and user group. Until the recent release of pre-configured images of the IBM cloud, it was the platform of choice for all Elon cloud computing. For details of the course and the use of the Google cloud in the courses please see [4, 5].

As with any platform, there are some disadvantages or issues that you would like to see improved on. The primary disadvantage of Google App Engine is that while standards based it does not implement the full Java EE 6 set of standards released in fall 2009 and limits some standard Java SE functionality. By not fully implementing all Java EE standards for Servlets, JavaServer Pages, JavaServer Faces, JPA and JDBC, students get frustrated and spend many hours trying to get a feature to work that either is not available or is partially implemented and does not work as defined. Server side programming is hard to debug for beginners and these additional unexpected limitations adds to the complexity of teaching cloud computing.

Java EE 6 is made up of 28 specifications. Some of the key specifications which the Google cloud does not fully support are Servlets 3.0, JSF 2.0, JPA 2.0 and Java SE 6. The Google cloud does not support Servlets 3.0. It is still on Servlets 2.5 which does not support annotations. The Google cloud does not directly support JSF 2. However, with the addition of two additional jar files and some web.xml configuration settings, Google will support most of JSF 2.0. A major JSF limitation is the use of an earlier version of the expression library (EL 1.1) which does not allow parameters to be passed to actions. The Google cloud does not recognize the El 2.2 library. The Google cloud does not fully support JPA 2.0. Their partial implementation introduces major restrictions on the storage and retrieval of objects as the underlying BigTable Datastore is not based on SQL. Google does not support the full Java SE 6. The Google cloud has a well-defined whitelist that enumerates the Java SE 6 features not supported for performance limitations in making a highly scalable cloud. The key whitelist limitations are: no jdbc, no writing of files and no threads. The Elon faculty feel that the failure to support JDBC and to provide a freely available SQL database is a major shortcoming to students.

4 IBM cloud (SmartCloud Enterprise)

IBM has an Academic Initiative program "that facilitates the collaboration between IBM and educators to teach students the information technology skills they need to be competitive and keep pace with changes in the workplace" [6]. During the summer of 2011, Elon University faculty worked closely with their IBM academic alliance relationship manager with the goal of having preconfigured cloud images developed by IBM that would have all of the advantages of the Google cloud and none of the disadvantages. IBM developed two images that supported the full Java EE 6 specifications with no whitelist, required no software installation by the student and allowed each student to have both an individual virtual machine for development and a separate virtual machine for deployment. In addition to having all of the advantages and none of the disadvantages of the Google cloud, WebSphere expertise is more desirable in the marketplace for internships and full time positions after graduation.

4.1 Development and deployment images

IBM created a development image and a deployment image. Both images use the windows operating system which Elon students are familiar with and required no learning curve. The development cloud image had Derby, WebSphere V8 (WAS) and Rational Application Developer V8 (RAD) installed and pre-configured. The pre-configuration was modeled after the ease and use of the Google cloud. The goal was to allow a beginner student in server side application development to develop their application in an Eclipse environment and to deploy it to an application server with a single mouse click. The pre-configuration involved the following key features:

1. The Derby database was installed with two shortcuts on the desktop. The first shortcut started the derby server as a network database server accepting all connections on port 1527. The student just had to double click on the shortcut to start the database. The second shortcut was to shut down the database.

2. A WebSphere profile was defined with a DataSource called ElonDB. The named DataSource made it easy for Servlets, JavaServer Pages and JavaServer Faces to connect to it and provided a nice abstraction to hide the underlying details of connecting to the external Derby database. Two shortcuts were created on the desktop to start up and shutdown WebSphere with the defined profile.

3. RAD had a workspace preconfigured with a Data Perspective and two servers. This section will first discuss the Data Perspective in detail before discussing the two servers.

When the student started RAD, RAD automatically came up on the predefined workspace. The Data Perspective had two database connections profiles defined. One connection called "ElonDB – localtest", shown in Figure 1, connected to a predefined Derby database on the Development image. The other predefined connection called "ElonDB – remote", connected to the predefined Derby database on the deployment machine.

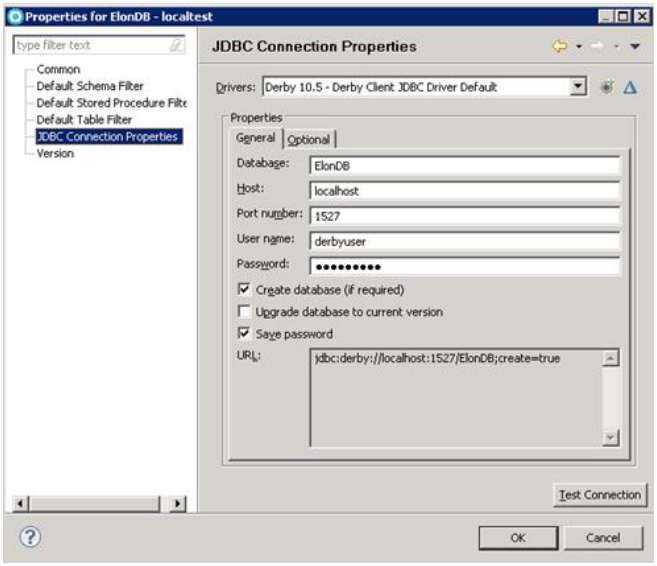

Figure 1: Eclipse configured database connection

The Database Connections in the Data Source Explorer View in the Data Perspective allowed the student to use the Data Source Explorer View wizards and GUI functionality to easily create, review, update and delete schemas, tables and data. The wizards and GUI functionality proved to be of great value when learning SQL and verifying that executing SQL statements through JDBC worked properly and to interactively develop the SQL to eventually call through JDBC.

Two WebSphere server instances are preconfigured in the Server Tab of the Java EE Perspective. The WebSphere instance running on the Development machine is called "WebSphere Application Server v8.0 at localhost with Derby". The WebSphere instance running on the Deployment machine is called "WebSphere Application Server V8.0 with Derby". The student right clicks the server instance that the student would like to publish an Enterprise Application on and selects the menu option to "Add" or "Remove". An Add or Remove Wizard is displayed for the user to select the Enterprise Application. Once added then the student can invoke the desired annotated servlet or JSP page by supplying the full URL in a browser.

4. RAD is based on Eclipse and required no learning curve for the Elon students. RAD provided additional wizards and tools that fully supported creation of Java EE 6 server side applications.

4.2 No software installation

The image instances required no software to be installed by the student. Both instances were directly accessible using the Microsoft Remote Desktop Connection from any Windows machine on or off campus or any Mac machine running Microsoft Office. The student only needed to supply the IP address of the image instance and the remote machine console was displayed as if the student were seated at the remote machine. The student could easily access and transfer files between machines using cut and paste or drag and drop.

4.3 Two virtual machines per user

The IBM SmartCloud Enterprise GUI allows the instructor to click on an image to request the naming and creation of an instance. The instructor created two individual virtual machines for each student using a login of the student's first and lastname and a password based on the student's unique and private student identification number. By using a standard naming convention, each student was guaranteed privacy from other students while providing the instructor the ability to remotely connect to a student machine to grade or answer questions. One machine had an instance of the development image and the other machine had an instance of the deployment image. With both instances, the student had full access to the WebSphere Application Server GUI administrative console. This allowed the student to get a sense for the tasks involved by a system administrator in adding user roles for application security, defining database schemas and tables for an application and installing an application into an environment with other running applications.

4.4 WebSphere job opportunities

The Google and IBM clouds are easy to use, extremely reliable and provide a standards based PAAS for use at all universities to teach Enterprise Computing. The IBM cloud with WebSphere has the added advantage of providing increased internships and job opportunities for graduates. Indeed.com and Dice.com are recognized computer science job search sites. Both are easily searchable to determine the number of currently open jobs looking for experience with the Java EE 6 servers for WebSphere from IBM, Weblogic from Oracle, JBoss from Redhat and App Engine from Google. Figure 2 shows the number for job openings posted by Indeed.com on May 22, 2012. The clear leader is WebSphere followed by Weblogic, Redhat and App Engine. There are significantly more jobs for WebSphere than App Engine. A

similar query was run on Dice.com with results showing a similar relative ranking of application servers. There were 3209 WebSphere postings, 2156 Weblogic postings, 1856 JBoss postings and 155 App Engine postings. Eleven Elon students took the Spring 2012 enterprise computing course on the IBM Cloud. Four received internships or full time positions using WebSphere and one received a full time position using JBoss.

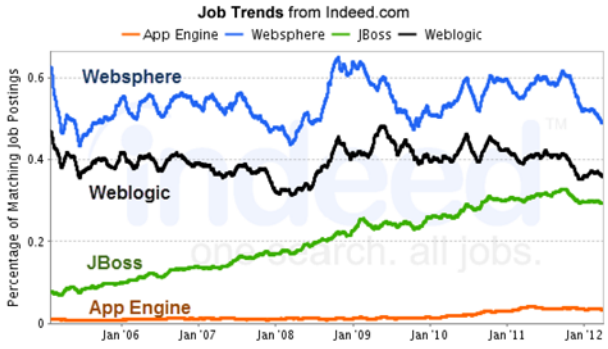

Figure 2: Current job postings for Java EE server skills

5 Course content

The Spring 2012 Elon course on enterprise computing focused on teaching server side programming using Servlets 3.0 and JSP 2.2 with an external SQL compliant database. When HTML was generated by servlets or JavaServer Pages, it was HTML 5. The IBM cloud discussed in section 4 was used for all development and deployment. This section describes the course materials, course format and course assessment.

5.1 Course materials

The course had two required text books and was supplemented by online materials. The required text books were Core Servlets and JavaServer Pages Volume 1 and 2 [7][8]. The author of these textbooks, Marty Hall, is president of a training and consulting company, coreservlets.com, that focuses on server-side technology and is an adjunct professor at John Hopkins University. He has extensive experience teaching Servlets, JavaServer Pages and JavaServer Faces at major corporations around the world. He freely supplies all of his course materials to faculty [9]. The materials include PowerPoint slides, example code, class exercises and class exercise solutions. His code is extremely well written and both uses and reinforces good programming practices. Every example that he provided worked without modification on the IBM cloud. This is not the case when using the Google cloud due to its whitelist limitations. Marty Hall updates and improves course materials after each of his courses based on feedback from corporate customers and students. The course materials complement his textbooks and are updated to the latest release of Servlets 3.0 and JavaServer Pages. This is critical. With the exception of one text book, all of the leading Servlet and JavaServer Page texts are based on Servlets 2.5.

Online materials were used to supplement the books and provided additional information on servlets 3.0, databases and HTML 5. For additional information on servlets 3.0, the students used the Servlet 3.0 Cookbook [10]. The site provided a nice overview of the new annotations used in Servlets 3.0 and provided illustrative examples. For learning SQL and the Derby database, the students were directed to read a 14 article series on IBM developerworks by Robert Brunner [11]. The hands on series starts by introducing the student to the Derby database by building a schema and table using SQL interactively with the Derby IJ Tool. The next seven articles cover the SQL data definition and data manipulation language. The last six articles cover JDBC. HTML 5 was used primarily for the new form capabilities. W3Schools provided a good online resource to introduce the new form attributes and elements [12].

5.2 Course format

The Spring 4 semester hour course met twenty eight times over 14 weeks. Each class session was 100 minutes in length and held in a lecture/lab setting. Each class typically consisted of a lecture followed by an in-class computer based exercise on the covered material using the IBM cloud. With the exception of database coverage, the course primarily followed the chapter sequence in the course texts and had four major parts. The first part of the course covered basic servlets, servlet annotations, request and response headers, cookies, sessions and file uploads. The second part of the course had a database focus. The Derby database was introduced with coverage of SQL, table design, normalization, JDBC and Data Access Objects. Exercises were assigned to have Servlets perform create, read, update and delete operations (CRUD) from Derby databases created by the students. The third part of the course covered JavaServer Pages, scriptlets, Java Beans, the JSP 2.0 expression language and using MVC for application design. Most exercises in this part required database CRUD operations to reinforce the heavy use of a database in server side applications. The fourth part of the course focused on advanced technologies covered in CoreServlets and JavaServer Pages Volume 2 for deployment descriptors, declarative and programmatic security, listeners, filters and JSTL. Exercises in this part required students to use their development machine for developing and testing the web application and database and to then deploy both the database tables and the web application to their deployment machine. Students had a chance to use the Administrative Console of WebSphere to add user names, passwords and groups and to associate these names against security roles defined in the web application to control resource access.

5.3 Student assessment

Students received frequent and continual assessment throughout the course with quizzes, homeworks and a final exam. The students had three major quizzes during the course. There was one quiz after part 1, one quiz after part 2 and one quiz after part 4. Each quiz had questions focused only on the material for the part of the course that it was assessing. The questions were taken from a test bank of questions for the concept being tested with a difficulty similar to that found on the Oracle Java EE Web Component Developer Certified Expert Exam. Students were provided links to various mock certification tests to prepare for each quiz. The goal of these quizzes was to have the student supplement the book reading with an in depth study of appropriate APIs in preparing for each quiz.

Individual homework programming assignments on the IBM cloud were assigned weekly. Each assignment typically took 2 – 5 hours to complete. During part one of the course, each assignment focused on the topic covered for the week. During parts three and four of the course, each homework assignment built on the previous assignment. The building block approach allowed each student to build a large, working web application that mimicked the play.google.com site for maintaining a list of Android Applications. The increasing complexity of each homework enabled students to understand the importance of the Model View Controller architecture, security restrictions and data access objects. The overall high average of 91 on assignments indicated that students had developed the ability to design, develop, debug, test and deploy an enterprise level application.

The final exam was a complete Java EE 6 Web Component Developer Certified Expert Exam from enthuware.com. The students were given the same number of questions (57) and the same time limit (120 minutes) as the actual exam. The use of the three earlier quizzes helped the students prepare for this exam as they understood the level of difficulty of the questions and the format of the exam.

6 Conclusions

The web platform is the dominant platform of our time. Elon University believes that enterprise computing should be a required course for all CS majors and has it made it a core course in its curriculum. The Google cloud and the IBM cloud have been successfully used at Elon and have provided students free, 24 x 7 access, with high reliability. The recently adopted IBM cloud with pre-configured images has proven to be as easy to use as the Google cloud while offering the advantages of being a fully Java EE compliant environment. These clouds are available to all universities and provide an ideal PAAS for teaching enterprise skills and make the student more competitive in the job market for internships and after graduation.

7 References

[1] Rosenberg, J. 2012. Skype Chief Technologist, April 16, 2012, Presentation to Computer Science Forum at Elon University.

[2] Gundotra, V., Picai, D., Google IO 2010 - Keynote Day1. http://www.youtube.com/watch?tv=mbVVDDu8f9k&feature=related

[3] Criteria for Accrediting Computing Programs 2012-2013. http://www.abet.org/computing-criteria-2012-2013/

[4] Hollingsworth, J. and Powell, D. 2011. Requiring Web-based Cloud and Mobile Computing in a Computer Science Undergraduate Curriculum. Proceedings of the 49th ACM Southeast Conference.

[5] Hollingsworth, J. and Powell, D. 2010. Teaching Web Programming Using the Google Cloud. Proceedings of the 48th ACM Southeast Conference.

[6] IBM Academic Initiative. https://www.ibm.com/developerworks/university/academicinitiative/

[7] Hall, M. and Brown, L., Core Servlets and JavaServer Pages: Core Technologies, Vol. 1 (2nd Edition), Prentice Hall, 2003.

[8] Hall, M., Core Servlets and JavaServer Pages: Advanced Technologies, Vol. 2 (2nd Edition), Prentice Hall, 2007.

[9] Hall, M., Intermediate Servlets & JSP Tutorials, http://www.corservlets.com/Course-Materials/csajsp2.html

[10] Servlet 3.0 Cook Book. http://www.servletworld.com/servlet-tutorials/servlet3/cook-book.html

[11] Brunner, R. 2006. Developing with Apache Derby. http://www.ibm.com/developerworks/library/os-ad-trifecta1/index.html

[12] HTML 5 Tutorial. http://www.w3schools.com/html5/

IBM cloud computing is a set of cloud computing services for business offered by the information technology company IBM. IBM Cloud includes infrastructure as a service (IaaS), software as a service (SaaS) and platform as a service (PaaS) offered through public, private and hybrid cloud delivery models, in addition to the components that make up those clouds. Learn Introduction to Enterprise Computing from IBM. Large Scale Enterprise Computing powers all major transactions, and the Mainframe is responsible for 87% of all credit card transactions and enables 71% of all Fortune 500 companies. Mainframes ...Â Mainframes are fundamental on how we do business, and IBM Z is the only production mainframe sold today. Virtually everyone depends on it. This is your introduction to the hardware, operating systems, security, and features that make this possible. On successful completion of this course, learners are eligible to earn their Introduction to Enterprise Computing badge. More information can be found here: https://www.youracclaim.com/org/ibm/badge/introduction-to-enterprise-computing. In January, IBM made waves when it announced its IBM Q System One, the world's first gate model quantum computer available to businesses â€" a system housed in a sleek, 9-cubic-foot glass case. It's a major milestone for quantum computers, which had to date mostly been found in research labs.Â Intel, IBM, Google, and the quantum computing startup Rigetti are building systems designed around superconducting circuits, taking their cues from today's conventional supercomputers. Microsoft is taking a completely different, and perhaps riskier approach, in trying to build a better qubit. The topological qubit that Microsoft is trying to build would fragment electrons to store information in multiple places at the same time, making it more stable and less prone to collapse.