CS XXXX: Cybersecurity in the Cloud

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Ranjodh Singh Sandhu

Spring, 2021.

Technical Project Team Members

Karanvir Singh Jassal

Justin Hoon Kim

John Daniel Light

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

David Wu, Department of Computer Science

Charles Reiss, Department of Computer Science

CS XXXX: Cybersecurity in the Cloud

A Technical Report on the Capstone Project

John Light
University of Virginia
jdl4cx

Hoon Kim
University of Virginia
jhk4tt

Ranjodh Sandhu
University of Virginia
rss6py

Karanvir Jassal
University of Virginia
ksj8abg

Abstract

In an age where speed and efficiency have been the primary measures for progress, there has been a mass migration of information and processes to online platforms, specifically the cloud. As such, there are those who seek to exploit this shift in medium, attacking cyber structures and frameworks within the cloud, for their own private gain. With each iteration of vulnerabilities and patches, attacks are only becoming more ingenious and harder to detect. Thus, it is no secret that traditional cybersecurity practices have not been able to keep up with the pace of growth in the cloud.

While schools have attempted to teach students these traditional practices as foundational in their own local projects, little effort has been made to apply them to the cloud computing environment. Introductory cybersecurity courses tend to focus more on understanding the basics of advanced topics in cybersecurity including encryption, digital forensics, binary exploits, and networks. These naturally tend to be more focused on programming against certain attacks. While that is appropriate in its own regard, with cloud environments, many of the security components are based on configuration. A user does not need to code the solution but use the available tools and resources to best secure the system in its particular use case. This is crucial as oftentimes, entire processes run on the cloud so there is a lot at stake when working on cloud platforms.

In order to bridge this gap, we propose a class that incorporates aspects from two different computer science courses at the University of Virginia -- CS 3710: Introduction to Cybersecurity and CS 4740: Cloud Computing. To develop this new class, we examined the documentation for these courses and expanded on their overlap. As former students of these courses, we reflected on their strengths and weaknesses in order to improve the structure of the new class and provide an experience that is tailored to cloud computing. As a result, students will better understand threats to the cloud structure and learn the safe practices that will help mitigate the associated risks. They will also be better equipped to handle cybersecurity issues when working in the cloud for their future careers, creating a safer and more secure environment for their clients and stakeholders.

While the class will have a focus on cybersecurity in the cloud environment, it will also aim to highlight some of the growing topics within cloud computing. Two growing fields that have had much controversy around them are Artificial Intelligence and the Internet of Things. While these two technologies have a tremendous amount of potential, they also pose some security concerns for users that have held them back. This course

would allow students to explore some applications of these technologies while also educating students on how to develop solutions in an ethical and secure manner. Both the Internet of Things and Artificial Intelligence are data-driven and many people do not have confidence that the data they are using is secure and collected ethically. So this course will aim to highlight practices that can keep the data in the hands of those that are meant to see it and how to gather data in ways that gain trust of the users. A student who has finished this course should be able to understand the concerns of users while developing solutions in the cloud environment. The generalized overview that CS 3710 and 4740 both bring is valuable so this course is not designed to replace them. Instead, this course is designed to be taken after completing CS 4740 because it will be more specialized.

During the first two weeks there will be a big review period of basic cloud terminology and essentials that were taught in 4740 and also a review on the class 3710. This will also go over new topics such as user and group permissions, authentication issues, and also cloud security basics. Then the next few weeks will be about network security revolving around the cloud including network addressing, network protocol layers, and network devices. This will give the student a better understanding of how different hosts interact within a network. Then the next several weeks will focus on Public-Key Crypto, Hashing for Authentication, Network Authentication, Authenticating Network Servers, Digital Signatures. Then the course will end on introducing more miscellaneous topics including how artificial intelligence and the internet of things are kept secure and good practices when it comes to security when dealing with those spectrums.

Introduction

Cloud computing is one of the fastest growing areas of computer science, and there needs to be more than one class available on the subject. Cyber security is important everywhere there is computer science, and this class helps students see where it's applied in the cloud. Currently, in order to get an understanding of cyber security, students need to take CS 3710 for an introduction, and also CS 4630 Defense Against the Dark Arts (DADA). Our class aims to dig deeper into cyber security, similar to 4630, but with a specific focus on the cloud. This way, instead of taking a generalized cybersecurity course, if students and particularly interested in cloud computing, they now have an option. The four of us have all taken CS 3710 and 4740, and three of us have taken 4630. When deciding what kind of class we wanted to create, the overlap between the courses had a lot of potential information we thought could be interesting to learn.

With many companies transitioning either completely or partially into the cloud, many companies have an increased reliance on the cloud. Cloud services are supporting entire operations and even contain sensitive information. With the shared responsibility model in the cloud, it is important that all users follow secure practices when operating in the cloud. This can help companies save resources and keep their information safe. As the liability can fall on the consumer rather than the provider, we believe that focusing on fundamentals will enable students to be successful in their careers. The experience and knowledge that students gain from this course will impact how they interact with the cloud.

Background

To understand the need for this course, one needs to know the course offerings for computer science at UVA. There is currently no class that covers this material, though there exist the ones mentioned above that are related. However, they don't go into the specifics of cybersecurity in cloud computing.

CS 4740 provides an introduction to cloud computing as it consists of a general overview, but since cloud computing encompasses so many different services, it is not possible to go in depth. Currently, if a student takes that course and

decides that they are interested in pursuing a career in cloud computing, they don't have many course options. A focus on security in the cloud will enable students to begin going more in-depth with cloud computing. We believe that this course is the first cloud computing course that should be taken after CS 4740 because of the complexity of vulnerabilities that can arise in the cloud environment. Before advancing into the different realms that exist within cloud computing, it is important that students have a solid foundation of this topic that is relevant everywhere in the cloud. Whether a student decides to focus more on artificial intelligence or cloud storage, they must first know how to secure the data and resources that they are working with.

While CS 3710 is an introduction to cybersecurity, the aim of this class is to show students how specifically the cloud environment is vulnerable to attacks. CS 3710 gives students a very broad idea of security but does not cover any attacks such as SQL injection attacks, cryptography, and network security as it relates to specifically cloud computing. Furthermore, this is becoming a more and more relevant field as more and more companies are moving towards cloud computing and adopting it. This will create a need for cloud security experts that can prevent hackers from exploiting companies' data.

Related Work

When looking up information for this course, we saw that other universities had courses that covered what we wanted to cover, reassuring us that this need does exist and that adding this class would be beneficial for UVA.

System Design

Our project was designed with the goal of taking information from two courses, CS 3710 and 4740, and digging deeper into their overlap. This class is for people who have an interest in cybersecurity and how it applies to cloud computing. The objectives of this class are for students to learn about threats to the cloud structure and learn the safe practices that will help mitigate the associated risks, understand how to handle cybersecurity issues when working in the cloud for their future careers, creating a safer and more secure environment for their clients and stakeholders, and explore the applications of IoT and AI while developing solutions in an ethical and secure manner. Our course design was a 16 week class with weekly quizzes, homeworks, and one final exam: The assignments are worth 60% of the grade, tests 20%, and quizzes 20%. We believe that the best learning takes place when students are applying the concepts. We wanted our grading structure to reflect a structure that promotes learning over the pressure of doing well on an examination.

Week 1 reviews the syllabus and touches on some concepts from 3710/4740. Week 2 and 3 go over some cloud data security information. Week 4 goes over session management and application security. In our final submission, these are the four weeks we created powerpoint lectures, quizzes, and assignments for. We split up our powerpoints so there would be two lectures per week. In the last 12 weeks, week 5-7 go over cryptography, network controls, and providers/scripts respectively. Weeks 8 and 9 go into detail about attacks like SQL injection and cloud based attacks. Weeks 10 and 11 talk about how to manage vulnerabilities and protect your data. Weeks 12-16 aren't specific about cybersecurity and the cloud, but discuss two other important topics in artificial intelligence (AI) and Internet of Things (IoT). These last two topics give students the opportunity to apply the knowledge they have gained throughout the semester in a fun way. AI and IoT are two technologies that have been growing rapidly and are major use cases for the Cloud. Working with these technologies is a way of rewarding the students at the end of the semester as they get to develop in the cloud while also making sure they are being safe and secure. Similar to current CS courses, lectures are recorded and available to listen to outside of class. Most homework assignments will require students to use AWS, some scripts, and Python to complete. To align with UVA honor code policy, this course's

honor code is also single sanction; if you are caught cheating, you will fail.

We decided that, for each week, it was important to include lecture material through the form of a powerpoint, a homework assignment, and a weekly quiz. We thought that all of these materials combined would qualify for enough classwork and would amount to the proper number of credit hours for a three credit course. Additionally, other classes we have taken generally assign a similar amount and type of work per week. The students would be able to learn and receive information from the powerpoints. The instructor would ideally lecture based on the slides. To test the students' learning and mastery over the week's material, a weekly quiz is designed to ensure each student took away the most important elements of the week's lesson. The quizzes would be open notes, since the purpose of the class is not to rank students based on grades and performance but to confirm that they are learning the material. If it seems that students are performing poorly on the quizzes, supplementary activities should be provided or additional office hours should be held. Finally, the assignments are to allow students to practically apply their educational knowledge. In our college experience, we have found that homeworks and assignments are the most vital piece of education for learning and information retention. The practice not only gives students the chance to experience potential real world applications, but the act of applying knowledge to an assignment also forces students to actually understand the material. Thus, we designed each week to comprise these activities.

For Week 1, we wanted to ease into the transition of classes and do what most classes do during week one. Essentially, we thought it would be beneficial to the students to go over the syllabus, what topics we will be covering during the class, honor policy, and expectations. Going through this information first is very important as students can get a feel of what the class will behave like and what the workload will consist of. Furthermore, a requirement to continue the class will be for students to sign the syllabus and return it which will also contain the agreement to uphold the honor system. Also in the first week, there will be an emphasis on logistics and how the class will be run. This will include introducing the TA's and their office hours. After all the class information is taken care of, the class will start diving into a basic review of Cloud Computing and Security. Firstly, the review will start with This will include covering topics such as PaaS, IaaS, and SaaS and discussing the main differences between the three services. After that, the content for the week will shift focus on the security basics. This will include a rundown of some common techniques used to exploit vulnerabilities and data such as phishing, distributed denial-of-service (DDoS) attack, and a brute force attack. The assignment will then focus on the students writing a python script and using a password cracking (brute force) technique. Along with that coding assignment there will be an essay due that will have the students research about past incidents of where these attacks have actually occurred and used to steal data from tech companies.

For Week 2, we thought that it would be best to focus on cloud data security, which consisted of privacy and protecting data. Since this would only be the second week of classes and students may still be adjusting their schedules, beginning with these concepts would be a good way to ease students into the more complex topics. The first topic introduced in this week is protecting data. We decided to focus on the different types of data that exist and how they can be protected. Many clients and companies use the cloud for data storage so it is important to be able to distinguish the information that exists in each type and how it can be secured. In addition, this course highlights the importance of keeping company and data secure by going through some examples of fraud throughout history. After that, the content for the week focuses on privacy laws and external regulations for specific sectors. After learning the different concepts relevant to the week's topics, the assignment aims to help students apply those same concepts in the cloud environment. Since we believe that the best way to learn is by doing, we get students to sign up for

AWS during the second week of the course so that they can dive in early on. After signing up, the students are tasked with securing an S3 bucket to limit access to specific users that they select. The goal of this assignment is to have students create a mindset of limiting access to only those who need it across all services. This week emphasizes the importance of understanding the domain of the work being done. Computer science, as a whole, has the potential to impact many different fields from medicine to business. This, however, means that as someone working with these different areas, an engineer must be aware of the expectations of those domains as well. Talking about external regulations and privacy laws helps establish that while standards exist across different fields, the depth and extent of the privacy and security varies. It is not a "one size fits all" approach when it comes to privacy.

For Week 3, we wanted to transition from understanding cloud data security and its broad practices to practical and specific implementations of cybersecurity in the cloud. The first and most common practice is Identity Access Management, or IAM. Naturally, we decided to focus on this specific subject matter first. Given the vast topic of IAM, we thought it would be sufficient enough to cover a week-worth of material. The overall objective of this week was to help students become familiar with IAM, understand identities and permissions for AWS resources, and be capable of maintaining a secure AWS environment. Thus, by the end of this week, students should be able to set up and configure users, groups, and roles; implement multi factor authentication; create IAM policies; implement password policies for security controls; and understand the AWS Key Management Service. Looking at the information concerning IAM, we decided we could break it down into six sections to make it more digestible for students. The first section is the overall background of IAM -- defining IAM itself, its purpose, and any terms that are commonly seen when dealing with IAM. Next, students would learn arguably the most important part of IAM, creating the actual objects that would be given or restricted access. During this portion, students would learn how to create and manage users, groups, and roles and understand the nuances between each object. Once students have a concrete understanding of creating objects, the curriculum focuses on the creation and customization of IAM policies themselves. Then, the students would have time to learn about additional AWS security features associated with IAM such as multi factor authentication, identity federation, and other miscellaneous security techniques. This material is supplemented by a quiz that has a mix of recall and application. We thought the couple recall questions included on this quiz were based on facts that any computer science student dealing with cloud and security ought to remember at all times. Overall, the intent of the quiz was to make sure that each student thoroughly understands each of the objectives for the week. Finally, there is a homework assignment attached with this week's curriculum. The assignment is designed to guide students to be able to create users, groups, and policies and attach permissions to these objects, essentially using the lectures to be able to implement IAM. Thus, through the powerpoint, quiz, and assignment, we decided that these would best teach students on how to properly implement and use AWS IAM to limit access and why this is important.

In Week 4, we thought that after exploring the cloud the previous two weeks, that it would be a good time to introduce some basic application security and session management to begin incorporating cybersecurity information. We start by discussing an important clarification between authentication and authorization as a background for session management, which is discussed in the second lecture this week. After clarifying the difference, we talk about ten common security risks for applications, and go into further detail about two of them: Cross site scripting and SQL injection attacks. In the classes we've taken, we felt like doing these attacks as homework really helped our understanding; we hope that the homework assignment for this week, which is an XSS and SQL injection attack on a practice website, would be beneficial. In the second lecture, we go over session management, the importance of session

IDs, and the roles that cookies play in session management. With week 2 and 3 providing the groundwork for cloud computing, this week aims to lay the foundation for cybersecurity. To end the week, we have a quiz that tests the students' understanding of authentication vs. authorization, common attacks and how to defend against them, key aspects of session management, and SQL/XSS attacks.

Results

After showing this course to several students that have previously taken at least one of the two courses (CS 3710 or CS 4740), we have gotten some promising results. A survey was conducted to see if this course was effective in teaching security in cloud computing given the four weeks of material to 5 students, and these are the results. There were four main questions that were asked after the students reviewed the material from the four weeks and looked at the quizzes and assignments.

The questions consisted of:

How much of the information did you already know? (1-5)

Do you think this course would help you gain knowledge in security relating to cloud computing? (1-5)

How interesting are the topics that are covered in the course? (1-5)

How likely are you to recommend a fellow student to take this course? (1-5)

The graphs for the following questions are presented below:

Question: How much of the information did you already know? (1-5)

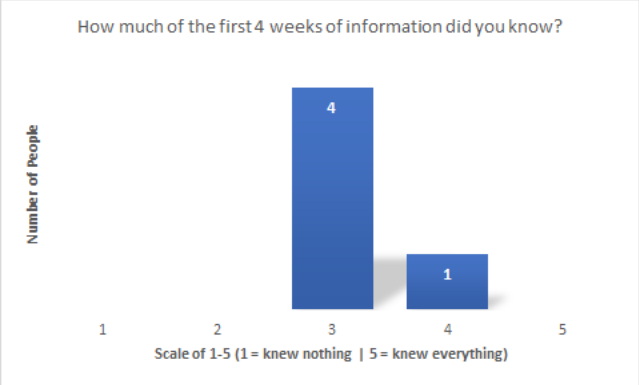

Most of the students said they knew some of the information, but not all of it. This shows us that while there is some overlap with the courses this class is inspired by, there is still new information that students haven't seen.

Question: Do you think this course would help you gain knowledge in security relating to cloud computing? (1-5)

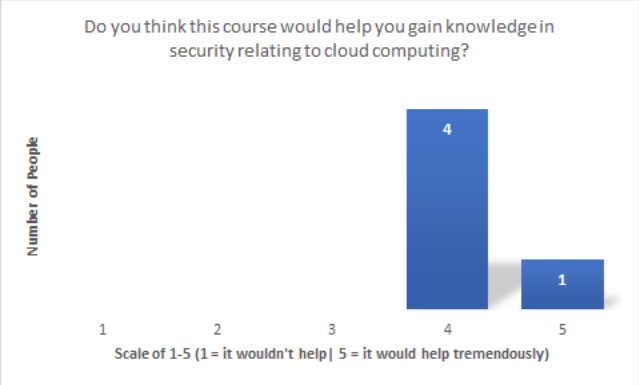

Most of the students rated that this course stated that this course would indeed help them learn about security measures relating to cloud computing, and said the structure of material was designed well.

Question: How interesting are the topics that are covered in the course? (1-5)

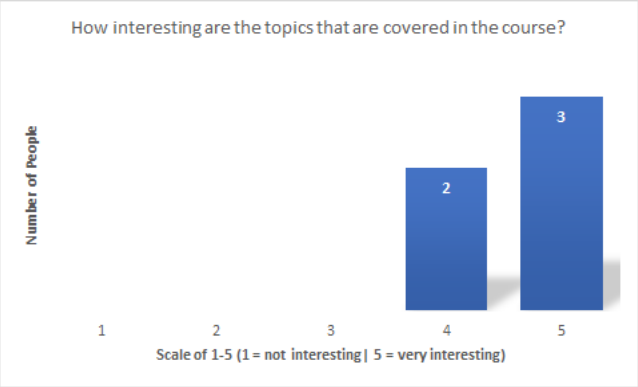

The response here was overwhelmingly positive, with comments about how cybersecurity is an interesting field and any more teachings on it are welcome.

Question: How likely are you to recommend a fellow student to take this course? (1-5)

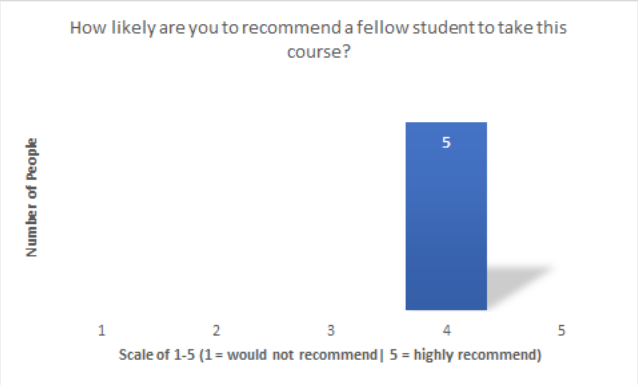

The responses for this question were all the same. Our interpretation of why this was is because this class would be an upper level course, students here could potentially already be well versed in some of these topics, and it wouldn't make sense for them to take it. However, for people without that knowledge and that have an interest in these topics, it would make a lot of sense to take this course.

Conclusions

We designed a class to address the need for expanded information where CS 3710 and CS 4740 overlap. Our course begins with a review of concepts from those two courses, then goes into detail about cyber attacks and defenses in the cloud, and ending with a discussion of AI and IoT. Students are tested weekly through quizzes and homework assignments, ending with a final exam. We believe that this course is necessary, as similar courses exist in other universities, and that it will provide a more comprehensive understanding of cloud computing and cybersecurity. This course will provide students with valuable information that they will be using when they enter the workforce, given that the role uses cloud computing services. Many people and companies have critical data stored on the cloud and any misconfiguration can be exploited. With the shared responsibility model that exists with cloud computing services, the liability of keeping the data secure on a digital level is in the hands of the user.

Before focusing on any specific area of cloud computing, it is important that students have a solid foundation on a topic that is prevalent everywhere in the cloud. By providing students with this fundamental skill and mindset, this course will enable students to use cloud services safely in future courses or work. The course also touches on ethics as cloud computing can be data-driven. This is done with the hopes that students will understand how to work ethically and within the regulations that exist within different fields of work. Working with cloud services may entail handling data from other sectors such as medical or financial. By having a better understanding of regulations and laws through this course, students will understand the security expectations when handling different types of data and have a better idea of how to proceed with the information.

Future Work

Currently, this project consists of a general overview of the course as well as material for the first 4 weeks of classes. These materials include a slide deck, a take-home assignment, and a quiz. If this course were to be implemented at the University of Virginia, the materials for the rest of the weeks would need to be developed as well as an examination. In addition, the first couple of semesters of this course's offering would bring more information and feedback on how this course could be adjusted.

The main work that needs to be done in the future is flushing out the assignments. Currently, the four weeks assignments are sufficient but creating the other 3-6 assignments for the other weeks will take some more creativity and time. Furthemore, we have not written any tests for this class as of right now and in the future that will need to be taken into consideration. Lastly, we are considering adding a big group project for this class that would help students realize the importance of teamwork in both cloud computing and security.

REFERENCES

[1] Amazon. (2003). *IAM User Guide*. Amazon. https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html.

[2] Amazon. (2003). *IAM*. Amazon. https://aws.amazon.com/iam/.

[3] Amazon. (2011). *Amazon Rekognition Developer Guide*. Amazon. https://docs.aws.amazon.com/rekognition/latest/dg/setting-up.html.

[4] *HIPAA Security Series*. U.S. Department of Health & Human Services. (n.d.). https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/smallprovider.pdf?language=es.

[5] *OWASP Top Ten*. OWASP. (n.d.). https://owasp.org/www-project-top-ten/.

[6] *PCI DSS Quick Reference Guide*. PCI Security Standards. (n.d.). https://www.pcisecuritystandards.org/documents/PCI_DSS-QRG-v3_2_1.pdf.

[7] *Session Management Cheat Sheet*. Session Management - OWASP Cheat Sheet Series. (n.d.). https://cheatsheetseries.owasp.org/cheatsheets/Session_Management_Cheat_Sheet.html.

[8] Tierney, M. (2020, July 2). *Data Security in Cloud Computing: Key Components*. Netwrix. https://blog.netwrix.com/2020/07/02/cloud-data-security/.

[9] *What Is Cloud Data Protection?* Palo Alto Networks. (n.d.). https://www.paloaltonetworks.com/cyberpedia/what-is-cloud-data-protection#:~:text=Cloud%20data%20protection%20is%20the,externally%20by%20a%20third%20party.

[10] YouTube. (2017, December 4). *AWS IAM | AWS IAM MFA | How to Create User, Group, Role,Policy and MFA*. YouTube. https://www.youtube.com/watch?v=Ihpkf3xwuJo.