UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
(Alexandria Division)

ANDREW BRODERICK, JACQUELINE BURKE, SUSAN CORLEY, LYNN FIELDS, KIMBERLY HERNANDEZ, KRISTINA MENTONE, MARK MILLER, MORDECHAI NEMES, RYAN OLSEN, DEBRA POTZGO, SHAWN SPEARS, JANETT STOUT, COLE STUDEBAKER, and JONATHAN WONG, each individually and on behalf of all others similarly situated,

Plaintiffs.

v.

CAPITAL ONE FINANCIAL CORPORATION, CAPITAL ONE BANK (USA) N.A., AMAZON.COM, INC., and AMAZON WEB SERVICES, INC.

Defendants.

Civil Action No. ______________

CLASS ACTION COMPLAINT
<u>AND DEMAND FOR JURY TRIAL</u>

Brian J. Dunne (CA 275689)
bdunne@piercebainbridge.com
PIERCE BAINBRIDGE BECK PRICE & HECHT LLP
355 S. Grand Avenue, 44th Floor
Los Angeles, CA 90071
Tel: (213) 262-9333

Andrew M. Williamson (VA 83366)
awilliamson@piercebainbridgecom
Andrew J. Pecoraro (VA 92455)
apecoraro@piercebainbridge.com
PIERCE BAINBRIDGE BECK PRICE & HECHT LLP
601 Pennsylvania Avenue, NW
South Tower, Suite 700
Washington, D.C. 20004
Tel: (202) 318-9001

Yavar Bathaee (NY 4703443)
yavar@piercebainbridge.com
Michael M. Pomerantz (NY 2920932)
mpomerantz@piercebainbridge.com
David L. Hecht (NY 4695961)
dhecht@piercebainbridge.com
Max P. Price (NY 4684858)
mprice@piercebainbridge.com
Michael K. Eggenberger (NY 5288592)
meggenberger@piercebainbridge.com
PIERCE BAINBRIDGE BECK PRICE & HECHT LLP
277 Park Avenue, 45th Floor
New York, New York 10172
Tel: (212) 484-9866

Attorneys for Plaintiffs

TABLE OF CONTENTS

PAGE

TABLE OF CONTENTS

PAGE

Plaintiffs, based on personal knowledge, and upon information and belief as to all other matters, allege as follows:

INTRODUCTION[1]

1. In March 2019, Capital One was the subject of one of the largest data thefts in history. The attacker, a former employee of Amazon Web Services, was caught and indicted. As information came to light about the nature of the attack, a striking set of facts began to emerge—not about the attacker, but about Capital One and Amazon. They had together, over several years, orchestrated a massive migration of highly sensitive data to a public cloud under the cover of false statements and Potemkin security software that Capital One and Amazon jointly created and jointly marketed to customers, regulators, and to the public as a means of keeping the data safe. But it was all a lie—and unbelievably, ***the precise conditions created by Defendants that gave rise to the March data theft persist to this day***.

2. This case is about a fraud by Capital One and Amazon—not the data theft that revealed it. And at base, it is about millions of Capital One customers who entrusted their most sensitive data—data that can be used by a thief to assume those customers' economic identity—to a bank and a cloud computing company based on a lie. Capital One and Amazon thoroughly monetized (and continue to monetize) sensitive Capital One customer data, mining it for every edge and insight about the behavior of Capital One's customers. But in order to obtain that data and the lucrative interest and fees those customers generated, Capital One promised customers that their data was safe and protected. Both Capital One and Amazon assured people around the country that this was the case. Those assurances have now been shown to be indisputably, willfully false and misleading—and they continue to be false, as were the statements Defendants made together

[1] Terms not defined in this Introduction are defined in the body of the Complaint.

over the years about the safety of Amazon's AWS public cloud for storage and processing of sensitive financial data.

3. As a result of these lies, Plaintiffs have paid billions of dollars in interest and fees to Capital One that they never would have paid had they known the truth: Their sensitive personal data was being pooled in a giant "data lake" on the world's most notoriously insecure public cloud, trawled by machine learning tools while at risk of theft via a well-known, unfixed Server Side Request Forgery ("SSRF") attack vector.

4. Defendants continue to aggregate and mine that data under the same perilous conditions that existed eight months ago. Customer data—years of it—is even today being aggregated and shared across hundreds of data mining systems, a simple SSRF attack away from another massive theft. That unsafe aggregation of data is not a bug; it is a feature. It is how Capital One makes money, and it is how Amazon sells its cloud computing services. Without years' worth of aggregated customer data, both companies would lose a competitive advantage.

5. Defendants know that there is no fix. They know that there is no setting they can change, or automated software they can write, to eliminate the risks that they intentionally force on their customers.

6. This fraud must stop. Plaintiffs seek damages and an injunction ordering the removal of sensitive Capital One customer data from Amazon's public cloud servers.

* * *

7. By the end of 2014, Capital One had collected an unprecedented amount of data about its customers. That data could tell Capital One how risky its credit card users were to lend to, how often they spent, what they spent on, and even where they went and what they cared about. The problem, however, is that significant amounts of hardware and software infrastructure were

needed to mine that data. Capital One needed data centers, storage, and computation power—all with the airtight security befitting a major financial institution.

8. This same opportunity was not lost on Capital One's competitors. They mined information from their customers by creating their own massive data centers, which they would upgrade, maintain, and secure at their own significant expense. Capital One had done the same for years, and in fact, had established its own data centers in Virginia by 2014. The cost, however, was too high for Capital One. Scaling would require more investment, and if the scaling was wrong, there was no inexpensive way to scale down.

9. Amazon's AWS presented a potential solution. AWS would allow Capital One to buy only as much computing power and storage as it needed. More importantly, it allowed Capital One to leverage Amazon's data scientists and machine learning tools, as well as arrays of the graphics processing units capable of the massive simultaneous calculations needed for machine learning.

10. There were significant problems, however, with using AWS to mine customer data. Machine learning models required massive amounts of historical data to train. If the data was insufficient, the models would not be accurate. In other words, Capital One would need to place years (and potentially over a decade) of sensitive customer information on the AWS cloud. But the potential damage from a security breach compromising a large trove of historical data would be incalculable.

11. Other large financial institutions knew this risk was too great. Both JP Morgan and Bank of America expressed and exercised extreme caution around customer data and refused to place their customers' data in the hands of a cloud provider. Banking regulators also had not yet weighed in on best practices and standards for aggregating data on a public cloud.

12. Capital One needed cover for its migration. At about the same time, AWS was searching for a large financial institution to adopt its ecosystem. AWS's business was being adopted by technology companies, startups, and other unregulated or less-regulated enterprises. The prize, however, was a large financial institution—one whose adoption of AWS would signal to other apprehensive financial institutions that it was okay to make the transition to the public cloud.

13. In 2015, when no other bank would, Capital One took the plunge and announced that it would migrate its user data and applications to the AWS cloud. It would move entire swaths of customer data to AWS's S3 servers to form a "data lake," a single source of data that Capital One's applications and machine learning models could all draw from. That data lake included ***over fifteen years*** of customer application data in order to better allow AI and machine learning algorithms to monetize that data for Capital One and Amazon.

14. This unprecedent aggregation of sensitive consumer data would, however, have to be sold as safe to Capital One's current and prospective customers. If those customers did not believe their information was safe, they would never agree to apply for, or use, a Capital One credit card. Capital One, with AWS's assistance, set out to assuage those fears by making false and misleading representations and omissions to current and potential customers, even developing its own software to manage the permissions of its internal computers and customer-facing applications to access the shared data lake. In other words, Capital One and AWS represented that they were able to guard against the inherent risk of pooling massive amounts of sensitive customer data for mining on the public cloud.

15. For years, however, AWS suffered from a widely known flaw. AWS servers, unlike those run by its competitors (*e.g.*, Google), were not secured against an SSRF attack, which would

allow an attacker to get inside a firewall and make requests to the data lake, including requests to pipe the data outside of the firewall to a third-party server. Year after year this flaw was the subject of some of the largest cybersecurity conferences in the United States. Each year, presentations were made expressly calling out AWS's particular SSRF vulnerability. Capital One ignored all of it.

16. To provide additional cover for its migration to the public cloud, Capital One created software, called Cloud Custodian, which it jointly showcased and marketed with Amazon. It was described as a "rules engine" that allowed Capital One to set specific policies within AWS that would apply in real time to the various servers that accessed its data lake. The software would, among other things, purportedly automatically scan Capital One's internal systems to ensure that all of the servers and permissions were set according to defined policies. Thus, when a computer wanted to access data from the data lake, it would assume a defined "role" that would then give it access to some portion or all of the data in the data lake.

17. These Identity and Access Management ("IAM") roles are used on AWS to allow various computers to access particular resources on a dynamic basis. A computer on Capital One's system with an IAM role configured to allow broad access, as required to train and deploy machine learning algorithms, could potentially allow that computer to access the entire data lake. Cloud Custodian would purportedly ensure that IAM roles were given the proper permissions to minimize the risk of a data breach; in other words, Could Custodian would grant the minimum amount of access necessary to complete a given task. For example, a customer-facing application such as a credit card application program would need to access systems to input the customer's data into the appropriate tables and then receive information about whether that applicant was approved and the

terms of the approval, but it would not need to access information about Capital One applicants from 2006.

18. The reality was that Cloud Custodian was not a solution to the serious problems posed by the mass aggregation of sensitive data and the open and dynamic access of countless servers to that data. Cloud Custodian's supposed benefit—ensuring the minimum amount of access necessary to complete a task—is at cross purposes with the goal of aggregating and mining broad swaths of customer data for profit. This is because in order to train and apply machine learning and AI systems, those systems need broad and dynamic access to user data, and that data must span years to ensure the accuracy and power of the AI and machine learning models.

19. A version of Cloud Custodian designed to minimize risk, then, would not serve Capital One's purpose for migrating to AWS's servers in the first place, which was the monetization of its customers' data. Accordingly, Cloud Custodian could not, and did not, solve the risk presented by the massive aggregation of data for exploitation on a public cloud server.

20. All that stood between an attacker and Capital One's data lake was a firewall, a system designed to block unauthorized access while permitting outward communication. The firewalls on Amazon's AWS cloud that guarded web applications, however, were known to be, and continue to be, vulnerable to a an SSRF attack. Other cloud providers have implemented additional precautions to ensure that requests from outside the firewall cannot be used to command resources on the inside, but AWS did not implement such precautions and has not done so to this day.

21. The net effect is that once an attacker obtains access to a server or system inside an AWS firewall, such as a firewall that protects a customer-facing web application, the attacker has access to all the data available to that server or system. If the attacker obtains access to a single

system that can assume a broad IAM role that permits it to access to the data lake, such as those that conduct machine learning tasks, all of that data can be transferred outside of the firewall at will.

22. Of course, Cloud Custodian could do nothing to prevent any of this, notwithstanding Defenadnts' statements otherwise. It did not matter to Defendants. AWS and Capital One jointly promoted Cloud Custodian as the solution to risk. This was a peculiar move for Amazon in particular because promotion of Cloud Custodian made no economic sense for Amazon.

23. ***First***, AWS already had a suite of tools that would purportedly ensure the proper configuration of IAM roles and monitor data access. In fact, AWS made money selling these tools to the users of its cloud. Nonetheless, AWS agreed to help Capital One promote Cloud Custodian, which competed with AWS's own tools.

24. ***Second***, Cloud Custodian was both open source and cross-platform, meaning that it could be migrated to competing cloud services, such as Microsoft's Azure or Google's GCP. Accordingly, the relationship between Capital One and Amazon was far from an ordinary business relationship between a cloud provider and one of its customers. A customer that adopted Cloud Custodian could more easily move its operations to a competing provider than one that relied on Amazon's own cloud management and security ecosystem. The only reason that AWS was willing to make that concession was to coax Capital One, a major financial institution, onto its platform, thus luring other financial institutions to join it.

25. Amazon also promoted Capital One's migration to AWS and the Cloud Custodian program. In late 2018, AWS hosted several web pages and videos touting its partnership with Capital One, the migration of Capital One's data to its cloud, Capital One's use of AWS to perform

machine learning on its user data at scale, and Cloud Custodian as a tool to keep the data safe. None of that promotion mentioned that Capital One and AWS had not dealt with the longstanding SSRF vulnerability peculiar to AWS.

26. Put simply, the only reason for AWS's decision to misleadingly promote a competing product was the immense value of attracting a large bank to its platform when other financial services companies refused to migrate their sensitive customer data to the public cloud. Capital One's use of AWS would demonstrate the safety of the cloud to financial services companies that sought to mine sensitive customer data. In exchange for this, Capital One would receive cover for its risky migration to the cloud, the pooling of customer data into the data lake, and the vast data mining operations it could conduct on its customers' personal information. Together, by developing and promoting Cloud Custodian, Capital One and AWS lulled regulators and customers into a false sense of security and created precedent for other large companies to adopt the AWS public cloud, thereby enhancing AWS's cloud ecosystem.

27. Capital One and Amazon knew about the inherent flaw in the architecture Capital One would have to deploy in order to exploit AWS's machine learning and AI tools and hardware, including the SSRF vulnerability. Both companies nevertheless falsely touted Cloud Custodian as the solution. In 2016, Amazon and Capital One posted the open source software on Amazon's AWS website, along with detailed documentation and marketing. But as both companies marketed Cloud Custodian as the solution to the risks of the data lake approach, they knew that Cloud Custodian was no solution at all.

28. For example, in December 2018, Kapil Thangavelu, Capital One's developer in charge of Cloud Custodian, gave a presentation at Amazon's AWS re:Invent conference. His presentation, entitled "Cloud Custodian—Open Source Security & Governance," touted Cloud

Custodian as a solution for the intractable task of maintaining appropriate permissions across several applications sharing aggregations of data. In an alarmingly prescient part of his speech, he discussed IAM roles and the precise vulnerability with poorly secured S3 servers that would later result in a breach of Capital One's own systems. He then falsely touted Cloud Custodian as a cure for that vulnerability.

29. Capital One and Amazon's statements proved false in March 2019, when a former Amazon employee scanned servers belonging to dozens of companies that had hosted their web applications on AWS and found a vulnerable entrypoint in Capital One's credit card application processing system. Using an SSRF attack, the attacker tricked one of Capital One's servers into sending information from Capital One's data lake to TOR nodes outside of Capital One's firewall and then to a server she controlled (the "Data Theft").

30. The scope of the breach was staggering, with compromised data going back to 2005. It was clear that Capital One had aggregated customer data on an unprecedented scale, and the compromise of one of the systems inside its firewall meant the complete compromise of over a decade of sensitive customer data.

31. Not only did Cloud Custodian fail to stop the Data Theft, it failed to even detect that it had happened at all; it wasn't until a July 2019 email from a third party that Capital One realized that it had suffered from the devastating attack. It was clear that Cloud Custodian was either a sham, designed to lull customers and regulators into a false sense of security, or it was never configured to limit access to years of historical data and found no anomalies to detect. Either way, all of Capital One and AWS's statements about Cloud Custodian were revealed to have been false and misleading.

32. Because the attack threatened to expose a more existential problem with Capital One's cloud operations, Defendants continued to lie about the root cause. Both Capital One and Amazon blamed a misconfigured firewall for the Data Theft, but that assertion is untrue. The problem is inherent in the architecture that Capital One chose and AWS enabled. Neither company addressed the fact that the architecture employed by Capital One on AWS was and is inherently at risk of a widespread data breach, including from an SSRF attack. Nor did either company address that, by design, Cloud Custodian, their touted solution to data vulnerability, was unable to detect or stop the attack.

33. Instead, Capital One and Amazon appear content to do nothing. AWS has not fixed its systemic vulnerability to the particular form of attack used in the Data Theft. Capital One has not fixed its aggregation-based, data-lake architecture that allows a simple hack to have devastating consequences. Both companies continue to profit on risking customers' valuable personal information.

34. Capital One, with AWS's knowing assistance, lied by stating that it would use industry-standard practices to protect its customers' personal information. They lied about the capability of Cloud Custodian. They lied about the Data Theft. And they are continuing to lie about the security of the personal information in the data lake.

35. If Plaintiffs knew the truth, they would not have paid interest and fees to Capital One, and they would not have applied for a Capital One credit card. More importantly, Defendants must be stopped from continuing their fraudulent scheme.

PARTIES

I. DEFENDANTS

36. Defendant Capital One Financial Corporation ("Capital One"), is a Delaware corporation with its principal executive offices located at 1680 Capital One Drive, McLean,

Virginia. It is a financial services holding company that offers an array of financial products and services to consumers, small businesses, and commercial clients, including the credit card products at issue in this lawsuit. Capital One reported $28 billion in revenue in 2018 and profits of over $6 billion after accounting for reserves, expenses, and taxes. Alone, Capital One's domestic credit card business generated $16 billion in revenue and $2.9 billion in profit.

37. Defendant Capital One Bank (USA), National Association ("COBNA") is a national bank headquartered at 4851 Cox Road, Glen Allen, Virginia. It offers credit and debit card products, including the credit card products at issue in this lawsuit, as well as other lending and deposit products. COBNA is one of Defendant Capital One's principal wholly owned subsidiaries. As such, references to "Capital One" herein are, unless otherwise noted, intended to encompass COBNA.

38. Defendant Amazon.com, Inc. ("Amazon.com") is a corporation existing under the laws of the State of Delaware with its headquarters and principal place of business located at 410 Terry Ave. North, Seattle, Washington.

39. Defendant Amazon Web Services, Inc. ("AWS") is a corporation existing under the laws of the State of Delaware with its headquarters and principal place of business located at 410 Terry Ave. North, Seattle, Washington. AWS is a subsidiary of Amazon.com.

40. Virginia is the largest market for data center space in the United States, and AWS maintains large data centers throughout the state. AWS operates its Virginia data centers directly or through a subsidiary called Vadata, Inc., which has operations in Ashburn, Haymarket, Manassas, Warrenton, Lorton, Culpeper, and Chantilly, VA. Either directly or through Vadata, Amazon leases 3.5 million square feet of space in Northern Virginia for its data centers.

41. Defendants AWS and Amazon.com are referred to collectively in this Complaint as "Amazon."

II. PLAINTIFFS

42. Plaintiffs, in the course of applying for Capital One credit cards, entrusted their personal information to Defendants with the understanding, based on Defendants' statements and representations, that Defendants would keep their information secure and employ reasonable and adequate security measures to ensure that it would not be compromised. Plaintiffs' expectation that their data would be secured was both reasonable and based on explicit promises made to them by Capital One and Amazon.

43. If Plaintiffs knew that Capital One and Amazon would not safeguard their information, they would not have applied for Capital One cards, and they certainly would not have paid the rate of interest and/or accepted the level of rewards associated with their cards. For most if not all Plaintiffs, the protection of their data was an indelible premise of applying for, and using, a particular credit card.

44. Plaintiffs' highly sensitive personal data remains in jeopardy to this day because Capital One continues to aggregate years of historical data on AWS's inherently flawed systems using an inherently flawed cloud architecture. As currently stored and maintained, Capital One and Amazon continue to breach their promises to Plaintiffs, and their statements about the safety of Plaintiffs' and other customers' data remain false and misleading. Plaintiffs require injunctive relief to abate their continuing injuries.

45. Plaintiff Andrew Broderick ("Broderick"), a resident of Texas, applied for three credit cards from Capital One, supplying personal information required by Capital One. Because Broderick lives in apprehension that his identity may be stolen as a result of Capital One and

Amazon's aggregation and maintenance of his data, as well as the Data Theft, he has expended effort to secure his identity, including obtaining credit monitoring.

46. Plaintiff Jacqueline Burke ("Burke"), a resident of South Carolina, applied for three credit cards from Capital One, supplying personal information required by Capital One. Capital One approved her applications, and Burke has maintained her accounts with Capital One to the present. Burke pays an annual fee on at least one of her Capital One cards, and she has paid interest on her balances to Capital One. Burke was informed by the IRS that she was the victim of a security breach and identity theft. Despite paying interest and other charges to Capital One for what she believed to be secure products, Burke lives in apprehension of identity theft as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft. Burke has expended effort to protect herself, including purchasing ten years of credit monitoring from Experian.

47. Plaintiff Susan Corley ("Corley"), a resident of Florida, applied for two credit cards from Capital One, supplying personal information required by Capital One. After she applied for credit from Capital One, an unknown party opened a credit card account in her name and gained access to her bank account. Corley remains in apprehension that her identity and personal information may again be stolen or compromised as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft. Corley has expended effort to protect herself, including using Credit Karma to monitor for further data thefts.

48. Plaintiff Lynn Fields ("Fields"), a resident of Wisconsin, applied for a credit card from Capital One, supplying personal information required by Capital One. Capital One approved the application, and Fields has maintained her account with Capital One to the present. Fields has been subject to at least one attack by unknown hackers that compromised her bank card. Fields

lives in apprehension that her identity may be stolen as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft.

49. Plaintiff Kimberly Hernandez ("Hernandez"), a resident of New Jersey, applied for a credit card from Capital One, supplying personal information required by Capital One. Capital One approved the application, and Hernandez has maintained the account with Capital One to the present. Hernandez has paid interest on her balances to Capital One. Despite paying interest to Capital One for what she believed to be a secure product, Hernandez lives in apprehension that her identity may be stolen as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft.

50. Plaintiffs Kristina Mentone ("Mentone") and Cole Studebaker ("Studebaker," and collectively with Mentone "Mentone/Studebaker"), residents of Connecticut, applied for joint credit cards from Capital One, supplying personal information required by Capital One. Capital One approved their application, and Mentone/Studebaker have maintained their account with Capital One to the present. When inquiring as to why his card was declined recently, ***Capital One informed Studebaker that there had been "a data breach," and that his account in particular "was affected by the breach."*** Capital One blocked and reissued Studebaker's card as a result of the Data Theft. Mentone/Studebaker have not only paid interest on their credit card balances to Capital One, but have also opened and maintained several other financial accounts with Capital One, in part for what they believed to be secure products. Because of the Data Theft, Mentone/Studebaker's entire portfolio of accounts and data has been put at risk. Mentone/Studebaker entrusted Capital One to keep all of their accounts and data safe. Because Capital One has already informed Mentone/Studebaker that they were in fact affected by the Data Theft, Mentone/Studebaker continue to live in apprehension that their identity, as well as sensitive

personal information across all of their financial accounts, may be stolen as a result of Capital One and Amazon's aggregation and maintenance of their data, as well as the Data Theft.

51. Plaintiff Mark Miller ("Miller"), a resident of Ohio, applied for two credit cards from Capital One, supplying personal information required by Capital One. Capital One approved the applications, and Miller has maintained the accounts with Capital One to the present. Since applying for Capital One cards, Miller has received several unsolicited credit cards—taken out in his name—for which he had never applied. Miller did not solicit these additional credit cards, and only someone who had access to the sensitive personal information that Miller supplied on his Capital One credit card application could have been able to apply for these cards. Miller expended significant effort to cancel each of these unsolicited credit cards. Because of the apparent theft of his sensitive personal information that was used to apply for these cards, Miller lives in apprehension that his personal data may be stolen as a result of Capital One and Amazon's aggregation and maintenance of his data, as well as the Data Theft.

52. Plaintiff Mordechai Nemes ("Nemes"), a resident of New York, applied for a credit card from Capital One, supplying personal information required by Capital One, including his Social Security Number and his business Employer Identification Number. Capital One approved the application, and Nemes has maintained the account with Capital One to the present. Nemes lives in apprehension that his personal identity or company data may be stolen as a result of Capital One and Amazon's aggregation and maintenance of his data, as well as the Data Theft.

53. Plaintiff Ryan Olsen ("Olsen"), a resident of Ohio, applied for a credit card from Capital One, supplying personal information required by Capital One. Olsen lives in apprehension that his identity may be stolen as a result of Capital One and Amazon's aggregation and

maintenance of his data, as well as the Data Theft. Olsen has expended effort to secure his identity, including obtaining credit monitoring.

54. Plaintiff Debra Potzgo ("Potzgo"), a resident of Pennsylvania, applied for three credit cards from Capital One, supplying personal information required by Capital One. Capital One approved her applications, and Potzgo has maintained her accounts with Capital One to the present. Within the past year Potzgo has suffered hacking attacks on her credit card accounts, resulting in numerous fraudulent transactions and charges posted to her account by hackers. Potzgo expended effort to have these fraudulent charges removed by Capital One. Despite paying interest and other charges to Capital One for what she believed to be secure products, Potzgo lives in apprehension of these and future hacking attacks against her as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft.

55. Plaintiff Shawn Spears ("Spears"), a resident of South Carolina, applied for two credit cards from Capital One, supplying personal information required by Capital One. Since applying for credit from Capital One, Spears has found numerous transactions on her credit report that she does not recognize as hers. An unknown party also secured cable and internet service in her name without her consent. Spears remains in apprehension that her identity and personal information may be stolen as a result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft.

56. Plaintiff Janett Stout ("Stout"), a resident of Indiana, applied for two credit cards from Capital One, supplying personal information required by Capital One. Capital One approved the applications, and Stout has maintained the accounts with Capital One to the present. Stout pays annual fees to Capital One for both cards. Despite paying these fees to Capital One for what she believed to be secure products, Stout lives in apprehension that her identity may be stolen as a

result of Capital One and Amazon's aggregation and maintenance of her data, as well as the Data Theft.

57. Plaintiff Jonathan Wong ("Wong"), a resident of Massachusetts, applied for two credit cards from Capital One, supplying personal information required by Capital One. Capital One approved the applications, and Wong has maintained one of the accounts with Capital One to the present. Wong has been subject to at least one email-based attack that used the password he had previously used on such applications. As a result, Wong has expended effort to secure his identity, including enabling two-factor authentication and geolocation on his accounts. Wong remains in apprehension that his identity and personal information may be stolen as a result of Capital One and Amazon's aggregation and maintenance of his data, as well as the Data Theft.

JURISDICTION AND VENUE

58. This Court has subject matter jurisdiction over this action pursuant to the Class Action Fairness Act ("CAFA"), 28 U.S.C. § 1332(d), because this is a civil action filed under Fed. R. Civ. P. 23, on behalf of a class exceeding 100 members, in which the amount in controversy exclusive of interests and costs exceeds $5 million, and where at least one member of the class (including, in fact, Plaintiffs themselves) are citizens of a different state than at least one of the Defendants.

59. This Court has personal jurisdiction over Defendants because their principal place of business is located within the state of Virginia and this district. They are at home in this forum and conduct significant business within it. They also have sufficient minimum contacts with this state and district. Indeed, Defendants' contacts with the state of Virginia are so pervasive that there is general personal jurisdiction over them. Moreover, much of the relevant conduct alleged herein occurred in Virginia and this district.

60. Venue in this District is proper pursuant to 28 U.S.C. § 1391(b)(3) in that Defendants are subject to personal jurisdiction in this district. Defendants Capital One and COBNA maintain their principal headquarters in, respectively, McLean and Glen Allen, Virginia.

FACTUAL ALLEGATIONS

I. CREDIT CARDS AND SENSITIVE PERSONAL INFORMATION—THE QUID PRO QUO

61. Modern credit card issuers are entrusted with copious amounts of sensitive personal information. In fact, the information provided by credit card applicants to issuers like Capital One includes some of the most sensitive, personal data one can imagine.

62. For example, Capital One asks applicants for their name, date of birth, social security number, and citizenship status:

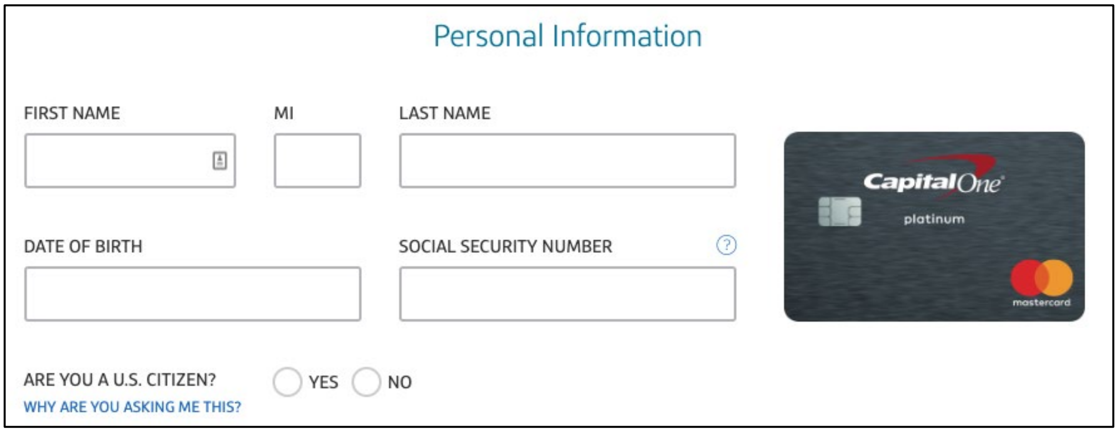

Personal Information

FIRST NAME

MI

LAST NAME

DATE OF BIRTH

SOCIAL SECURITY NUMBER

ARE YOU A U.S. CITIZEN? ○ YES ○ NO

WHY ARE YOU ASKING ME THIS?

63. Capital One also asks for an applicant's residential address, email address, and phone number:

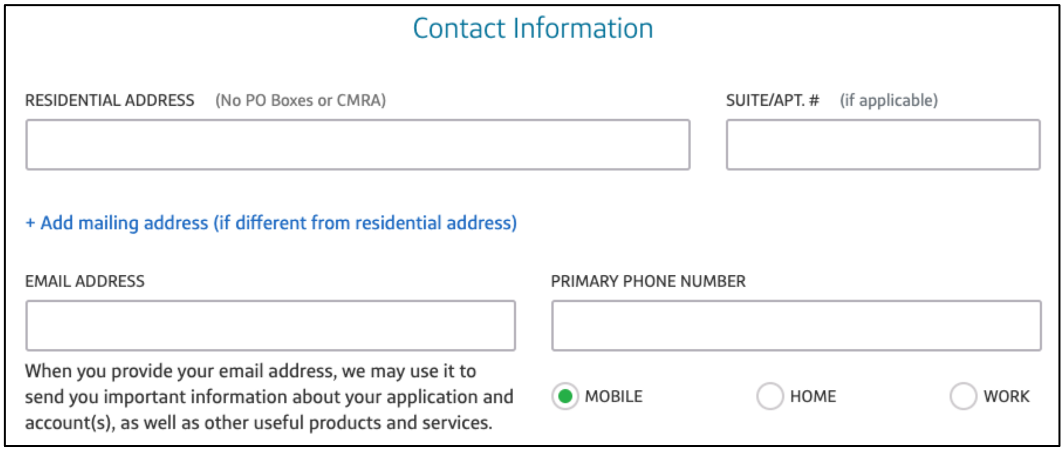

Contact Information

RESIDENTIAL ADDRESS (No PO Boxes or CMRA)

SUITE/APT. # (if applicable)

+ Add mailing address (if different from residential address)

EMAIL ADDRESS

When you provide your email address, we may use it to send you important information about your application and account(s), as well as other useful products and services.

PRIMARY PHONE NUMBER

● MOBILE ○ HOME ○ WORK

64. And Capital One asks credit card applicants for granular financial and employment information, including bank account information, employment status, annual income, and rent/mortgage information:

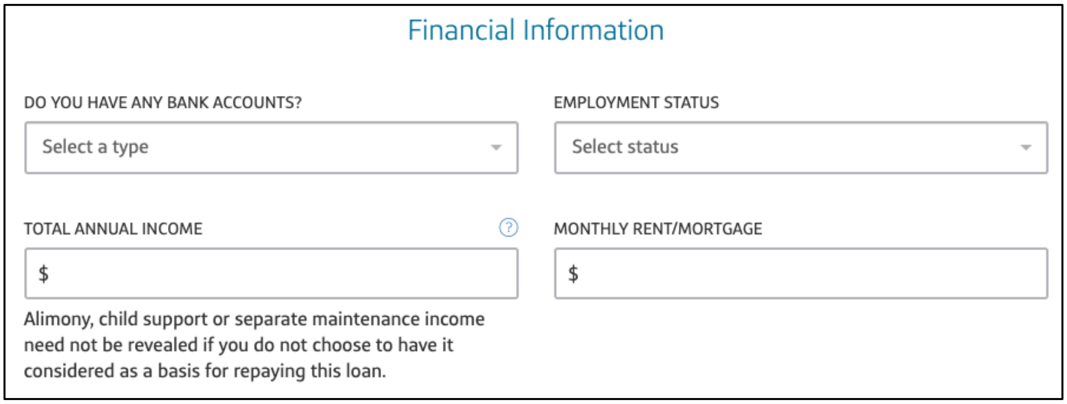

Financial Information

DO YOU HAVE ANY BANK ACCOUNTS?

Select a type

EMPLOYMENT STATUS

Select status

TOTAL ANNUAL INCOME

$

Alimony, child support or separate maintenance income need not be revealed if you do not choose to have it considered as a basis for repaying this loan.

MONTHLY RENT/MORTGAGE

$

65. Moreover, the sensitive personal information demanded by credit card issuers like Capital One isn't just received and stored: it's used to gather ***even more*** sensitive personal information on credit card applicants, including data from credit bureaus.

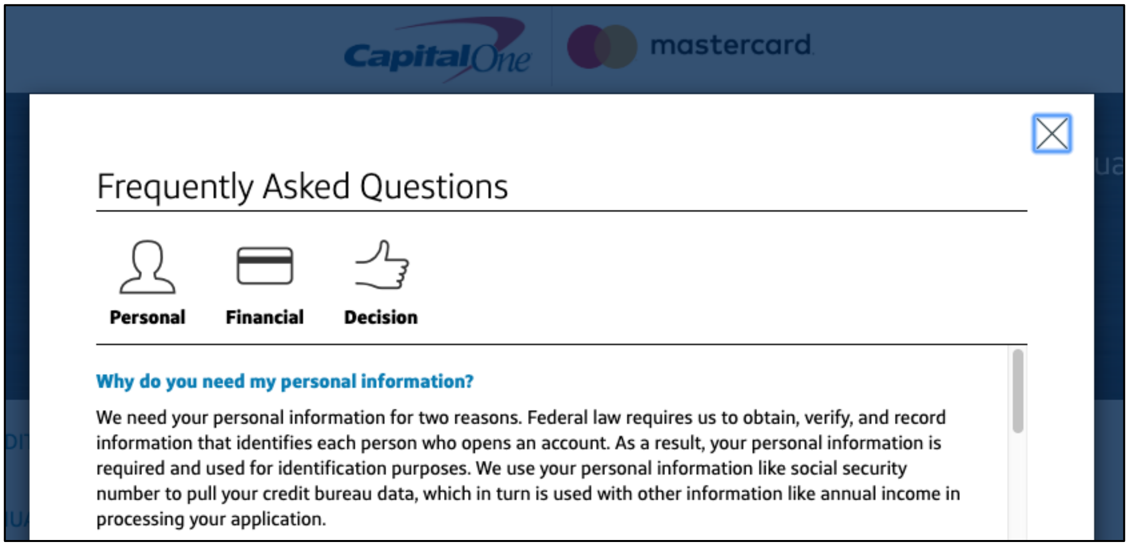

Frequently Asked Questions

Why do you need my personal information?

We need your personal information for two reasons. Federal law requires us to obtain, verify, and record information that identifies each person who opens an account. As a result, your personal information is required and used for identification purposes. We use your personal information like social security number to pull your credit bureau data, which in turn is used with other information like annual income in processing your application.

66. For example, an issuer like Capital One uses information supplied by credit card applicants to run what is called a "credit check" through one or more credit bureaus like Equifax, TransUnion, and Experian. This check sends an applicant's personal information to a credit bureau and returns a credit report that is then used by the issuer (*e.g.*, Capital One).

67. The credit report is based on information provided to the credit bureaus by other lenders and creditors. Altogether, the sensitive personal information gathered by an issuer like Capital One forms a credit history, which is used by the issuer to determine how much to lend (*i.e.*, credit limit) to an applicant, at what interest rate, and what fees to charge for use of the issuer's credit card.

68. In short, credit card issuers like Capital One use applicants' sensitive personal information to make money. The more personal information a credit card issuer has about its applicants, the more precisely it can target credit risk (and shore up its bottom line) through higher interest rates, low credit limits, and miscellaneous fees.

69. And Capital One is the best in the business at making money from granularly targeted fees and interest. Indeed, it is this ability to target people by risk level that has allowed Capital One to profit from the riskiest borrowers. In the past decade, Capital One's credit card

business has repeatedly been fined by federal and state regulators for unlawfully aggressive sales and monetization tactics. Between July 2011 and March 1, 2017, the Consumer Finance Protection Bureau ("CFPB") received more than 12,000 complaints directed toward Capital One's credit cards:

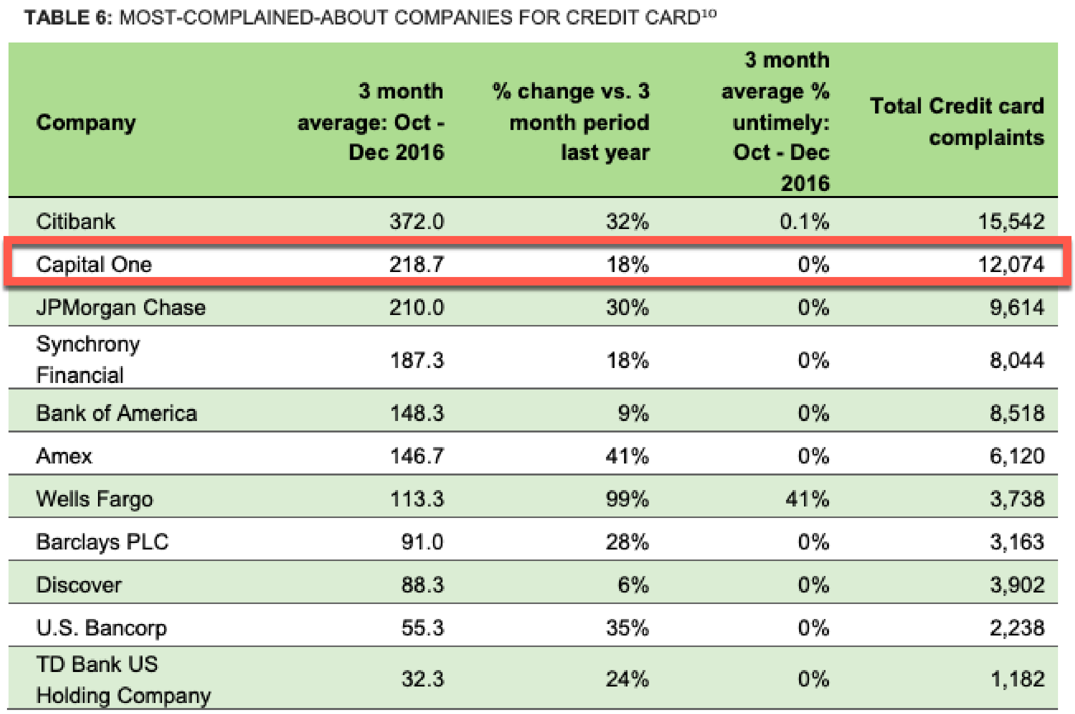

TABLE 6: MOST-COMPLAINED-ABOUT COMPANIES FOR CREDIT CARD[10]

Company	3 month average: Oct - Dec 2016	% change vs. 3 month period last year	3 month average % untimely: Oct - Dec 2016	Total Credit card complaints
Citibank	372.0	32%	0.1%	15,542
Capital One	218.7	18%	0%	12,074
JPMorgan Chase	210.0	30%	0%	9,614
Synchrony Financial	187.3	18%	0%	8,044
Bank of America	148.3	9%	0%	8,518
Amex	146.7	41%	0%	6,120
Wells Fargo	113.3	99%	41%	3,738
Barclays PLC	91.0	28%	0%	3,163
Discover	88.3	6%	0%	3,902
U.S. Bancorp	55.3	35%	0%	2,238
TD Bank US Holding Company	32.3	24%	0%	1,182

70. At one point in 2012, an astounding 22% of ***all*** credit card complaints received by the CFPB were against Capital One.

71. But user-targeted fees and interest are only part of the story. In recent years, credit card issuers like Capital One have developed an even more broadly sweeping way to make money from users' personal information: rewards programs. Specifically, card issuers like Capital One use rewards programs to maximize revenue from interchange fees (described below), and they use the personal information of applicants and cardholders to optimally target and shape these rewards programs.

72. Credit card companies make money not only from fees and interest paid by their cardholders, but also from processing fees paid by merchants. These fees are typically a flat rate plus a percentage of the total sale. This money is referred to as interchange income, and it is directly tied to the number and size of transactions a cardholder makes on their credit card. Interchange income represents 70% to 90% of the total fees paid to issuers by merchants.

73. In order to maximize credit card transaction volume (and thus interchange income), credit card companies like Capital One offer reward programs. These reward programs may create direct financial incentives (for example, "cash back"), restaurant gift cards, or airline miles to incentivize cardholders to make purchases using the issuer's credit card, thereby increasing interchange income.

74. At the same time, however, rewards programs create significant risks for issuers, from the out-of-pocket costs to cover the rewards to the risks associated with increased borrowing by cardholders. As a result, credit card issuers like Capital One aggressively compete to identify and attract high-purchase-volume, low-default-risk applicants. The secret sauce in this battle for rewards-program profits is granular, detailed personal information about applicants, which enables precise risk and reward targeting by card issuers. For example, knowledge of a cardholder's proclivity for fine dining can be used to target a rewards program that incentivizes and rewards dining out.

75. In 2018, Capital One's net income from interchange fees was approximately $2.8 billion. Capital One's 2018 annual filing with the SEC reported that the interchange fees it collected had increased for the year because of "higher purchase volume." Capital One's rewards program—the subject of its well-known, and extremely expensive, "What's In Your Wallet?"

national advertising campaign—exists to increase that volume. Indeed, Capital One nets its interchange fees against the cost of its rewards program, which in 2018 was $4.4 billion.

76. In short, the personal information collected from card applicants is vital to every aspect of a credit card issuer's lending business. Personal information is used to: (1) gauge risk; (2) set limits, fees, and interest; and (3) determine the type and overall level of rewards to both attract cardholders and incentivize maximum card use. And the role of personal information is non-binary: because personal information is integral to both revenue maximization ***and*** risk minimization, there is a direct, positive correlation between the amount and granularity of personal information a credit company collects and its expected profits from cardholders.

77. In sum, the more granular and accurate the information a credit card issuer is able to obtain about a borrower, the more predictable and stable its profits become. That is why credit card issuers demand highly sensitive information from applicants—it is integral to their bottom line.

78. Borrowers, however, do not provide sensitive information about themselves—especially the detailed personal information demanded by credit card companies—to every company that asks for it. For years, the government, the media, and other entities have warned Americans about identity theft, data breaches, and other risks and pitfalls of the modern information economy.

79. Data security is important to consumers—so important that credit card companies like Capital One make the promise of electronic safety and security a prominent part of their card offerings from the very first page of the credit application:

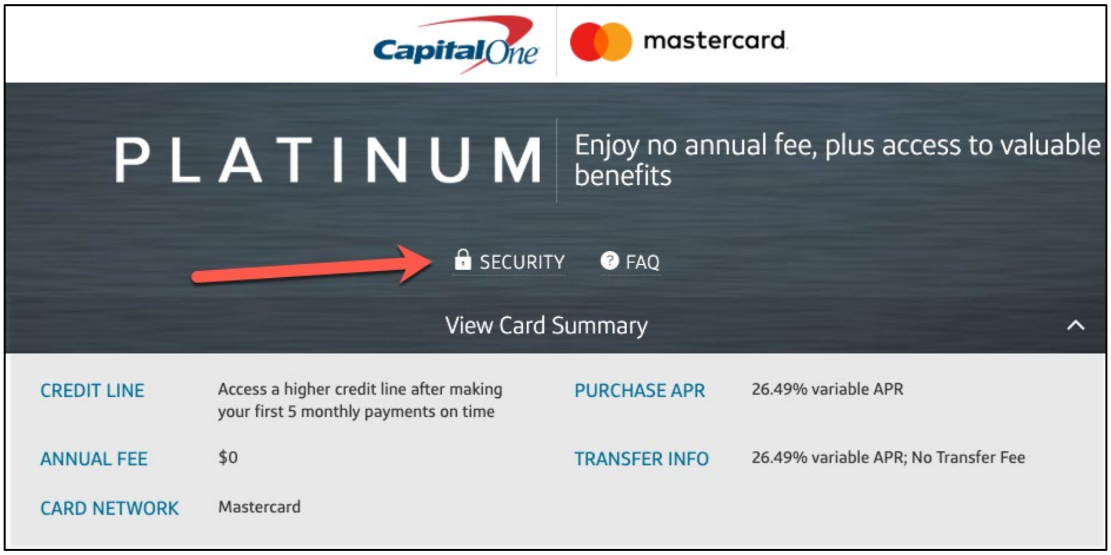

Capital One | mastercard

PLATINUM

Enjoy no annual fee, plus access to valuable benefits

SECURITY FAQ

View Card Summary

CREDIT LINE	Access a higher credit line after making your first 5 monthly payments on time	PURCHASE APR	26.49% variable APR
ANNUAL FEE	$0	TRANSFER INFO	26.49% variable APR; No Transfer Fee
CARD NETWORK	Mastercard		

80. In a page directly linked to the first page of its online credit card application, Capital One assures applicants that "[w]e're committed to protecting your personal and financial information" and that "[y]our security is a top priority."

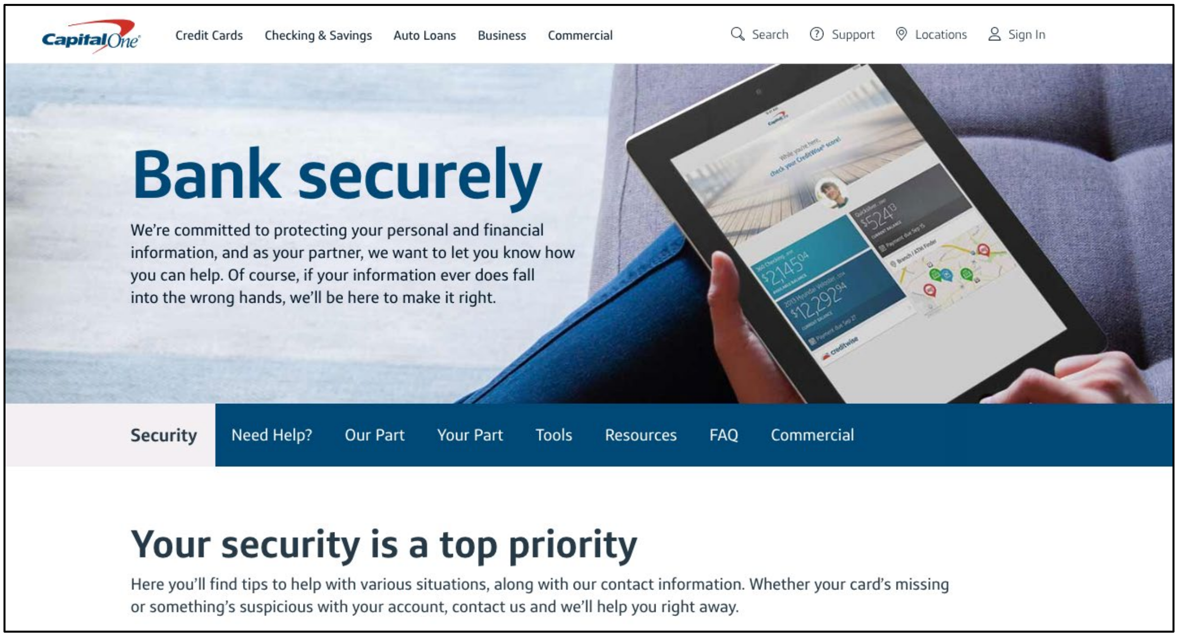

Capital One Credit Cards Checking & Savings Auto Loans Business Commercial Search Support Locations Sign In

Bank securely

We're committed to protecting your personal and financial information, and as your partner, we want to let you know how you can help. Of course, if your information ever does fall into the wrong hands, we'll be here to make it right.

Security Need Help? Our Part Your Part Tools Resources FAQ Commercial

Your security is a top priority

Here you'll find tips to help with various situations, along with our contact information. Whether your card's missing or something's suspicious with your account, contact us and we'll help you right away.

81. Even after the credit card application process, the stream of cardholder data continues to pour in. Credit card charges allow credit card companies to predict the expected amount of rewards that they will have to pay out, the amount of interchange income they can

expect, the risk of cardholder default, and even complementary products and services that can be marketed to cardholders.

82. Put simply, there is an important bargain at the heart of the credit card lender-borrower relationship: the card holder agrees to provide information that the card issuer needs to ensure that its business is profitable and predictable, and in return, the card issuer agrees to safeguard that sensitive customer information.

83. Capital One is no exception; it needs granular borrower data. In fact, one of the risk factors Capital One routinely discloses to its investors is a failure to accurately estimate its losses:

> *Estimates of Inherent Losses:* The credit quality of our portfolio can have a significant impact on our earnings. We allow for and reserve against credit risks based on our assessment of credit losses inherent in our loan portfolios. ***This process, which is critical to our financial condition and results of operations, requires complex judgments, including forecasts of economic conditions.*** We may underestimate our inherent losses and fail to hold an allowance for loan and lease losses sufficient to account for these losses. Incorrect assumptions could lead to material underestimations of inherent losses and inadequate allowances for loan and lease losses. . . .

(emphasis added).

84. As Capital One's investor disclosures explain, its business depends on the ability to make judgments and forecasts about likely losses. For that, Capital One relies heavily on accurate and timely data about its customers.

85. Without the assurance that Capital One would safeguard their sensitive personal information, creditworthy applicants simply wouldn't provide this information to Capital One. Potential customers would not apply for, let alone use and pay for (through interest, fees, and foregone rewards from other issuers), a card from an issuer that did not protect the sensitive information provided by the customer. This in turn would significantly harm—even decimate—

Capital One's credit card profits. Indeed, Capital One warned of precisely this risk in its 2019 annual report to shareholders:

> Negative public opinion or damage to our brand could also result from actual or alleged conduct in any number of activities or circumstances, including lending practices, regulatory compliance, ***security breaches (including the use and protection of customer information)***, corporate governance and sales and marketing, and from actions taken by regulators or other persons in response to such conduct. ***Such conduct could fall short of our customers' and the public's heightened expectations of companies of our size with rigorous data, privacy and compliance practices, and could further harm our reputation.*** In addition, our cobrand and private label partners or other third parties with whom we have important relationships may take actions over which we have limited control that could negatively impact perceptions about us or the financial services industry. The proliferation of social media may increase the likelihood that negative public opinion from any of the events discussed above will impact our reputation and business.

(emphasis added).

86. In a saturated market for credit cards, credit card companies fiercely compete for borrowers with good credit history. A *sine qua non* of this competitive process is the promise to electronically protect an applicant's most sensitive personal information using (at a minimum) industry-standard data security practices. As detailed in this Complaint, this is a promise that Capital One made repeatedly—and continues to make—to credit card applicants and cardholders, in numerous places and contexts, to obtain the valuable personal data that drives its bottom line. It is a promise bolstered by Capital One's co-conspirator Amazon. And it is a promise that was and is knowingly false.

II. CAPITAL ONE'S EXPRESS PROMISE TO SAFEGUARD SENSITIVE CUSTOMER DATA

87. Because Capital One needs sensitive borrower information for its business, it induces customers to provide that information by making representations about how it maintains and secures it. Among many other places, these representations are expressly set forth in Capital

One's various agreements with its customers and in the incorporated representations it makes in various disclosures and webpages.

88. For example, Capital One's Customer Agreement (the "Agreement") states that seven categories of documents, including the Agreement itself and certain "privacy notices," govern the relationship between the cardholder and COBNA. The Customer Agreement also states that "Capital One supports information privacy protection," and refers the reader to Capital One's website at www.capitalone.com.

89. Capital One's website contains several pages of representations about how data is collected and maintained by the Defendants. For example, Capital One purports to restrict access to Social Security numbers except when required for business purposes. Indeed, Capital One's "Identity Protection Commitment" on its website expressly represents, among other things, that "[w]e prohibit the unlawful disclosure of your Social Security number" and "[w]e restrict access to your Social Security number except when required for an authorized business purpose."

90. Capital One's representations to credit card applicants and cardholders include an express Privacy Notice. According to Capital One, the purpose of the Privacy Notice is to "let our customers know how we collect and use their information" and "how we keep information confidential and secure," so that "customers and potential customers [can] make informed decisions" about providing sensitive personal information to Capital One in exchange for its credit card services.

91. Capital One further represents to credit card applicants and cardholders that it maintains electronic safeguards, "such as passwords and encryption," to protect customer information.

92. Capital One also promulgates an express Privacy Policy through its website. The Privacy Policy makes further representations about Capital One's data and privacy security measures: "To protect your personal information from unauthorized access and use, we use security measures that comply with federal law. These measures include computer safeguards and secured files and buildings."

93. The Privacy Policy makes clear that Capital One collects "personal information" when a customer "[o]pen[s] an account or deposit[s] money," and expressly promises to protect that information. For example, Capital One's website states that it will "protect [cardholders'] personal information from unauthorized access and use," through measures that include "computer safeguards," "secured files," and "[secured] buildings."

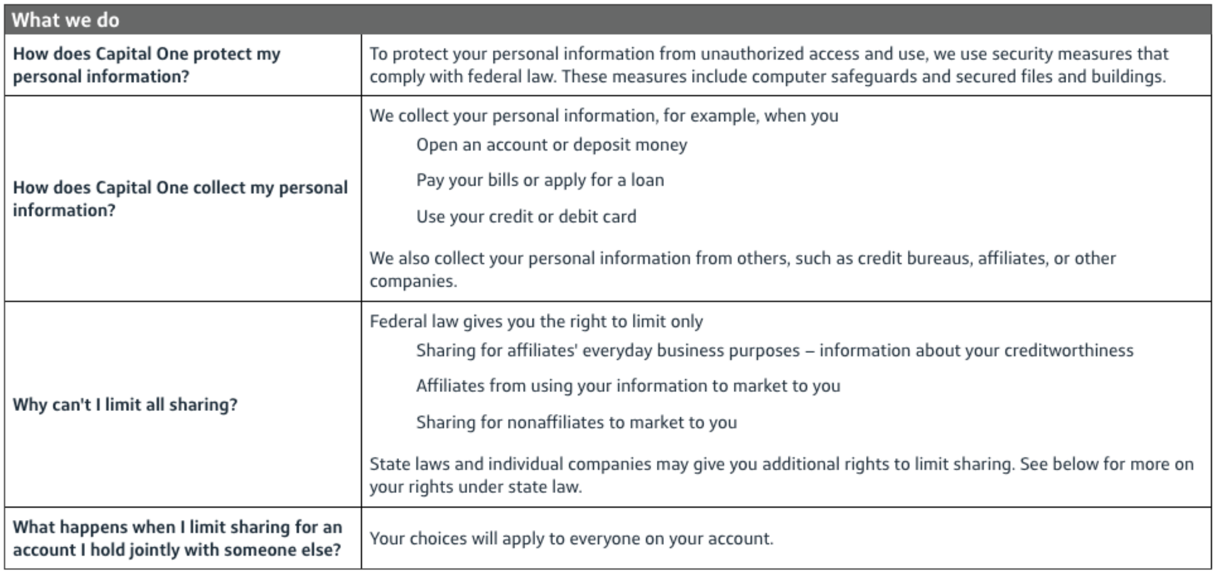

What we do	
How does Capital One protect my personal information?	To protect your personal information from unauthorized access and use, we use security measures that comply with federal law. These measures include computer safeguards and secured files and buildings.
How does Capital One collect my personal information?	We collect your personal information, for example, when you Open an account or deposit money Pay your bills or apply for a loan Use your credit or debit card We also collect your personal information from others, such as credit bureaus, affiliates, or other companies.
Why can't I limit all sharing?	Federal law gives you the right to limit only Sharing for affiliates' everyday business purposes – information about your creditworthiness Affiliates from using your information to market to you Sharing for nonaffiliates to market to you State laws and individual companies may give you additional rights to limit sharing. See below for more on your rights under state law.
What happens when I limit sharing for an account I hold jointly with someone else?	Your choices will apply to everyone on your account.

94. Capital One's website further represents that Capital One is "committed to maintaining the privacy and security of your information," that it "build[s] security into all of our systems and networks," that its "experts perform internal and external tests on all our applications and systems to safeguard your information," and that it "leverage[s] multiple preventative and

detective methods to mitigate risks and protect access to your information through a layered security program."

95. Capital One echoed this statement in its 2018 Annual Report, dated February 20, 2019, stating that it "safeguard[s] [its] customers' and [its] own information and technology, implement[s] backup and recovery systems, and generally require[s] the same of [its] third-party service providers," and that it "take[s] measures that mitigate against known attacks and use[s] internal and external resources to scan for vulnerabilities in platforms, systems, and applications necessary for delivering Capital One products and services."

96. Most notably, Capital One made (and continues to make) an unequivocal promise about its data security technology to its customers:

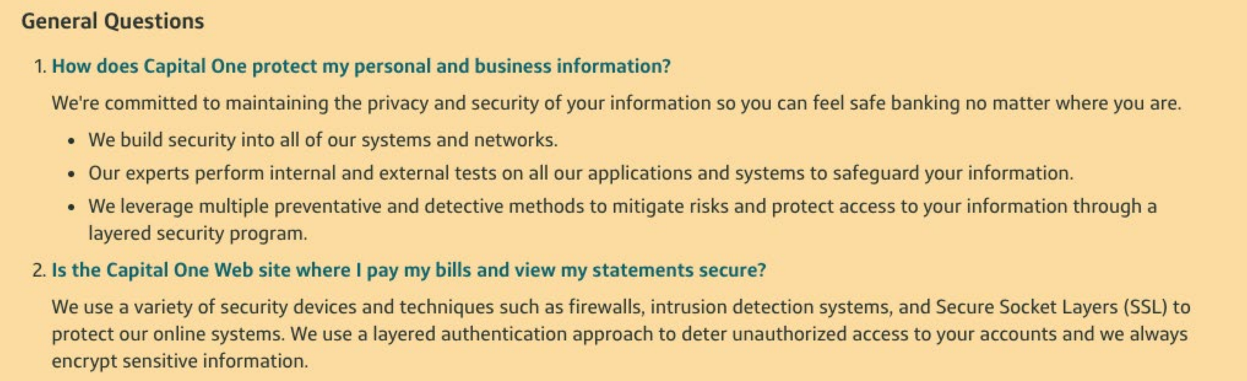

General Questions

1. **How does Capital One protect my personal and business information?**

   We're committed to maintaining the privacy and security of your information so you can feel safe banking no matter where you are.

   - We build security into all of our systems and networks.
   - Our experts perform internal and external tests on all our applications and systems to safeguard your information.
   - We leverage multiple preventative and detective methods to mitigate risks and protect access to your information through a layered security program.

2. **Is the Capital One Web site where I pay my bills and view my statements secure?**

   We use a variety of security devices and techniques such as firewalls, intrusion detection systems, and Secure Socket Layers (SSL) to protect our online systems. We use a layered authentication approach to deter unauthorized access to your accounts and we always encrypt sensitive information.

97. Capital One's technology is touted as so secure that Capital One frames it as the "Technology Guarantee": "We build information security into our systems and networks using ***internationally recognized security standards***, regulations, and ***industry-based best practices***."

98. ***Everyone*** who applies for a Capital One credit card encounters ***multiple*** assertions by the company about its data security practices, including Capital One's "Technology Guarantee."

99. As explained below, the representations by Capital One recited in this section were, and are, false. Capital One does not now, and has not for several years, followed "internationally

recognized standards" or "industry-based best practices" to build information security into its systems and networks, as it expressly promises applicants and cardholders who give Capital One their sensitive personal information. To the contrary, since 2014, Capital One took risks with its customers' data that were not in accordance with "industry-based best practices." Capital One was aware of the standards of the industry, but rejected those standards in favor of its own approach that elevated its profits above the safety of its customers' sensitive data. Instead, Capital One has developed a unique and troubling set of practices that fall far short of industry standards for the security of customer information.

III. CLOUD COMPUTING

100. Developing and maintaining data centers to store and process vast troves of sensitive client information is prohibitively expensive. For example, the banking industry spends $2,300 per employee annually on cybersecurity defense as part of this cost. For Capital One, the cost of cybersecurity is over $500 million a year.

101. To store, process, and mine sensitive customer data, banks like Capital One traditionally use a dedicated-server or private-cloud solution for their storage and processing needs. Dedicated servers assign specific hardware and software to perform specific tasks, while private clouds allow hardware and software to be assigned dynamically. In both scenarios, the equipment is dedicated to a single company that exercises control over the infrastructure. A private cloud is cloud-based infrastructure—such as servers, applications, and other equipment—that is dedicated to a particular business that exercises control over the infrastructure, which it often owns and operates. Private clouds offer greater degrees of security and control than hosting data on public clouds, but at an increased cost. Private clouds are not dynamically scalable, meaning the company pays for storage and processing capacity when that capacity is not in use or even no

longer needed. Companies typically develop and maintain their own private cloud, dedicating resources to its development and maintenance.

102. In contrast, public clouds are hosted and run by third parties such as Amazon AWS, Microsoft Azure, IBM Cloud, and Google Cloud. Those third parties own and maintain the infrastructure, which is then leased on a scalable, dynamic basis to multiple businesses. Because resources can be scaled to meet demand, the business only pays for the services that it uses, potentially saving money. Public clouds also allow companies to focus on its applications and services rather than developing its infrastructure, cutting the time to market and deploy those services to its customers. The primary downside of public clouds is the increased risk inherent in their use, and the related difficulty of meeting regulatory hurdles regarding the security of sensitive information.

103. Due in part to the risks associated with public clouds such as AWS, including regulatory hurdles, financial institutions and banks in particular have been reluctant to migrate their storage and processing of sensitive customer information to public clouds. For example, a 2016 report by Deutsche Bank revealed that public cloud adoption was “very small” among big banks such as Capital One despite the eagerness with which those same banks jumped on other technological innovations, including open source programming and processing Big Data. A 2018 Accenture report noted that only 34% of banks surveyed had complete plans for addressing issues related to security and compliance related to public cloud services—let alone had completed the migration to the public cloud. This was a red flag for more widespread adoption within the industry. As of 2015, no bank anywhere near the size of Capital One had migrated its customers’ personal information to a public cloud provider.

104. In recent years, more banks have opted to transfer only some of their services to public clouds, adopting a hybrid cloud approach that keeps their most sensitive information in the bank's private cloud and under the bank's direct control.

105. In addition to offering cheaper but less secure storage and processing capability, some public cloud services offer AI and computer learning services at a scale that is impossible for even the largest banks to replicate in a cost-effective manner. These services allow an issuer to more effectively monetize their customers' sensitive personal information by identifying patterns that are invisible to all but the most sophisticated algorithms.

A. Amazon and AWS

1. Amazon Develops AWS

106. During a 2003 executive retreat at Amazon CEO Jeff Bezos's house, the Amazon leadership team was asked to identify the company's core strengths. The response, which came after a fair amount of discussion, was obvious—it was Amazon's infrastructure. In particular, Amazon had built scalable systems to develop and deliver its customer-facing web applications, which were websites that dynamically generated content for individual users. Information processed on Amazon's servers could seamlessly be presented to both customers and developers and could handle fluctuations in traffic dynamically. Amazon's infrastructure also allowed for centralized storage that could be accessed dynamically across different web applications.

107. Amazon decided to sell its infrastructure to other companies that needed to develop and deploy web applications. On November 9, 2004, Amazon announced the forthcoming Amazon Web Services project ("AWS"). On March 19, 2006, AWS began offering a suite of services, including: (1) Simple Storage Service ("S3"), a cloud-based storage service that could scalably store large amounts of data accessible by multiple servers at the same time; and (2) the Elastic

Computer Cloud ("EC2"), a cloud-based server that could be configured and deployed dynamically based on particular needs (*e.g.*, memory and/or computation power).

108. Amazon touted its newly announced products on its AWS blog in a post, dated August 24, 2006:

> Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. Just as Amazon Simple Storage Service (Amazon S3) enables storage in the cloud, Amazon EC2 enables "compute" in the cloud. Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use.

109. Amazon's AWS product solved several important problems inherent with the dedicated server or private cloud approach. Internally hosting data for a company operating at scale requires massive amounts of infrastructure and expense. Not only would such a company need highly-trained personnel to estimate how much computing power and storage was needed for a particular application, but an entire department would often be needed to maintain and secure the infrastructure. And developing the infrastructure required to host and process massive amounts of sensitive customer information takes time, resulting in a lag between need and capability.

110. In addition, the age of Big Data had arrived. Although many banks had accumulated massive amounts of data, they lacked the computing power, often requiring highly specialized equipment, to process and harvest that information. But processing large troves of user data was the ultimate competitive edge, as Amazon itself had demonstrated. A real-time understanding of a customer's behavior, needs, and desires meant a keen ability to sell them precisely what they would

likely buy, at an optimized price point. Mining user data also meant that behavior patterns difficult to perceive by humans could be seized upon by a computer. All of this, however, required computation on a scale that would normally be cost-prohibitive for most businesses, even giant banks.

2. AWS and the Machine Learning Edge

111. AWS provided an answer to the Big Data problem—namely, how does a company leverage the data it collects from its customers without a massive infrastructure investment? AWS allows a company to launch as many "instances" of a server as it requires, meaning precisely configured servers can be spun up on demand, depending on what is needed. In other words, if a company needs additional computing power, it can instantly purchase a dynamically created server—by the hour, if it wants to. Because these services are paid for as they are used, there is no need for a massive upfront investment in the specialized equipment and people necessary to fully harvest that information. Companies running traditional dedicated server or even private cloud infrastructure would have to make an additional investments—into expensive graphics processing units ("GPUs"), related hardware, and specialists such as data scientists—in order to leverage machine learning to the massive amount of data being collected by a company through its operations. AWS had already made that investment into the specialized equipment and people necessary for large scale machine learning, and it leased that capability out to users of its public cloud services.

112. In an April 9, 2015 blog post, AWS's Chief Evangelist, Jeff Barr, explained precisely this:

> Today, it is relatively straightforward and inexpensive to observe and collect vast amounts of operational data about a system, product, or process. Not surprisingly, there can be tremendous amounts of information buried within gigabytes of customer purchase data, web

> site navigation trails, or responses to email campaigns. The good news is that all of this data can, when properly analyzed, lead to statistically significant results that can be used to make high-quality decisions. The bad news is that you need to find data scientists with relevant expertise in machine learning, hope that your infrastructure is able to support their chosen tool set, and hope (again) that the tool set is sufficiently reliable and scalable for production use.

113. AWS provides machine learning computation and knowhow that would require a massive infrastructure investment for companies seeking to mine troves of customer data. In the same April 2015 blog post, Barr announced Amazon Machine Learning—a set of tools that would allow common machine learning operations such as regressions and classifications of data to be performed on AWS-hosted data.

114. As machine learning has become more sophisticated, ordinary servers are no longer sufficient to perform the complex mathematical computations needed. New technologies, such as deep networks of artificial neurons, require the processing of large clusters of numbers at once as "vectors."

115. The only computation mechanism to do so with adequate speed comes from the world of video games. Graphics Processing Units ("GPUs") are designed to perform complex mathematical calculations on vectors at rapid speeds in order to, for example, render 3D video games with dynamic lighting and realistic physics. Enterprise-grade GPUs, like those owned and operated by AWS, are able to process hundreds of thousands of small programs at once, precisely the sort of work required to process, and therefore monetize, Big Data.

116. AWS allows developers access not just to servers equipped with GPUs, but to cloud servers equipped with entire arrays of enterprise-grade GPUs. Giant arrays of GPUs in the AWS cloud allow machine learning to occur at scale. In other words, by 2015 AWS presented for the first time a clear path to data-mining valuable user information without a massive upfront infrastructure investment.

3. The AWS Business Model and the Adoption Feedback Loop

117. AWS charges for the infrastructure services it offers in a variety of ways, such as hourly or even yearly prices. Some of the resources it provides customers are spun up dynamically as they are needed, and Amazon bills them according to the use of those resources. Put simply, AWS makes money by selling computer time, storage space, and its own cloud-management and security tools.

118. This business model relies on attracting customers to the AWS platform and keeping them on it. The model faces two hurdles: (1) convincing customers to adopt the public cloud; and (2) convincing customers to tie themselves to Amazon's ecosystem rather than a competitor's. In other words, to grow AWS, Amazon has to convince customers of the efficacy and security of the public cloud, attract customers to its platform, and then provide tools that keep them on the platform.

119. Banks in particular are juicy targets for AWS's cloud storage and machine learning services. By 2016, banks were investing ***over $20 billion*** annually in Big Data services, such as data mining, and Amazon wanted a piece of that growing pie.

120. A 2016 *Wall Street Journal* article noted that Amazon was attempting to increase its footprint with banks, and that "[l]anding a big Wall Street bank would give Amazon extra credibility around security and privacy safeguards." However, banks proved to be reluctant adopters of public cloud services. JP Morgan's Chief Operating Officer noted that moving to the public cloud would require addressing "key controls" such as access, encryption, and legal and compliance issues. Any bank migrating to a public cloud service would have to feel comfortable about security, including assuaging the concerns of both its customers and its regulators. Convincing a single large bank, such as Capital One, to move to AWS and use and adopt AWS's

tools could trigger a wave of other large clients with troves of sensitive customer data to migrate to AWS as well.

121. Moreover, engineers, developers and IT professionals will only train for Amazon's cloud ecosystem if there is sufficient demand for such training in the job market. The more companies that adopt Amazon's cloud infrastructure, the more valuable training on the AWS platform will be for engineers, developers, and IT professionals. In turn, the more professionals that are trained for the AWS cloud system, the easier the transition will be for companies seeking to move to AWS.

122. By 2014, Amazon knew that if it could reach a critical mass of large companies using its cloud ecosystem, a virtuous circle would emerge—a feedback loop. Adoption by Big Data users would lead to the refinement of AWS's tools and to the training of more AWS professionals—both of which would clear the path for other companies, large and small, to migrate to AWS.

123. For Amazon, the next move was clear—they would need to expend significant resources to attract a critical mass of large companies to their platform. Sure enough, AWS's marketing in this period focused on touting the fact that several large companies had made successful migrations to the cloud. For example, in an October 16, 2014 blog post, AWS's Stephen Organ told his story:

> In 2011, I had the opportunity to develop private cloud infrastructure at Bloomberg for their web properties. While these efforts drove down costs and sped the delivery of infrastructure by several orders of magnitude, it was impossible to keep up with the pace of innovation that AWS has clearly demonstrated. Subsequently, one of the major prongs of the technology strategy I employed at Dow Jones (and by extension News Corp) was to become a cloud-first enterprise. Over the last several years we migrated substantial portions of infrastructure to AWS while enabling new product development to happen much more rapidly with far less investment.

> Through these experiences, I have seen that cloud computing is among the best, if not the best opportunity for enterprises to become more agile, drive down costs, and free up resources to focus on their business—the products and services that bring in revenue.

124. AWS's pitch was clear: the cloud was the future and a surefire way to reduce costs while gaining flexibility and agility. AWS's goal was to bootstrap the feedback loop it needed to build out its ecosystem.

4. The Bug Is a Feature: The Dynamic Access, Data Pooling, and Server-Side Request Forgery Problems

125. The promise of AWS, and in particular its machine learning capabilities, comes with an important caveat. To scalably apply machine learning, fast and on-demand access to large amounts of data is critical. Data must be harvested from central databases, cleaned and organized, then fed into machine learning models for training and testing. To train a sophisticated machine learning model, such as a deep neural network, large datasets have to be used, and to ensure that the models work correctly, the data used to train them must span broad swaths of time.

126. A machine learning model is not like an ordinary computer program, which is often nothing more than a series of instructions given to a computer for rote execution. Machine learning models learn directly from data. They spot patterns in data and make decisions based on those patterns. There is often no instruction that encapsulates a decision—the decisions flow directly from the data observed by the model during training.

127. For example, a machine learning model designed to decide whether to provide a person with a credit card based on their personal information would do so based on past credit card applicants and particular events or outcomes associated with them, such as bankruptcy or default on debt. If a machine learning model is used, there may often be no computer programmer who sets hard rules; what has worked in the past is used by the computer to make decisions about the present.

128. If a machine learning model is trained with too narrow of a dataset, it may make poor decisions in data contexts that differ from the training set. For example, if a model is trained on credit card applicants from only a three-month period during the 2008 financial crisis, it may poorly predict credit card-holder outcomes during prosperous times.

129. That is why machine learning models demand large amounts of data. The more data they train on, the more powerful and accurately predictive the models become. Thus, for a system of machine learning models to be properly trained, all of the models need access to historical data. For the models to continue to function, they have to continue to train on new data as it is collected. In short, data has to be accumulated into a central place, and the machine learning algorithms require access to all of that data.

130. AWS's S3 servers allow for such machine learning by allowing large amounts of data to be pooled into what is referred to as a "data lake." Different web applications all draw from the same data lake on a dynamic basis as required—regardless of whether those particular applications require access to the entire broad swath of data in the lake.

131. To restrict web applications to parts of the data lake that they need (and nothing more), AWS requires the configuration of access "policies" as part of predefined roles that can be assumed by applications. One way to do this is through Identity and Access Management ("IAM") roles.

132. AWS allows resources on its cloud to be configured to assume IAM roles, and then access the resources they need based on policies associated with those roles. AWS describes some of the potential uses of IAM roles in its documentation:

> You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you might want to grant users in your AWS account access to resources they don't usually have, or grant users in one

> AWS account access to resources in another account. Or you might want to allow a mobile app to use AWS resources, but not want to embed AWS keys within the app (where they can be difficult to rotate and where users can potentially extract them). Sometimes you want to give AWS access to users who already have identities defined outside of AWS, such as in your corporate directory. Or, you might want to grant access to your account to third parties so that they can perform an audit on your resources.
>
> For these scenarios, you can delegate access to AWS resources using an IAM role. This section introduces roles and the different ways you can use them, when and how to choose among approaches, and how to create, manage, switch to (or assume) and delete roles.

133. Managing IAM roles and the policies applicable to these roles, however, becomes a monumental task at a sufficiently large scale. Once the number of applications sharing the same data lake become numerous, an incredibly robust management system is required to: (a) ensure that IAM roles are narrowly scoped, allowing access only to resources necessary to applications assuming a given role; and (b) ensure that no IAM roles are misconfigured. The more complex a cloud-based system becomes, the harder it is to manage all of the resources accessing the pooled data.

134. The power of maintaining a data lake thus comes with an important cost—security. It is the curse of data centralization and dynamic access that if just one application can be accessed from the outside—from the Internet—then the entire data lake can be at risk.

135. To guard against this, the data lake, servers, and applications are placed behind a firewall. A firewall, among other purposes, ensures that sensitive resources on a computer network are not exposed directly to the Internet. For web applications that need to pass data to and from a user on the open Internet, such as a credit card application, a Web Application Firewall ("WAF") is used. A WAF filters, monitors, and blocks web traffic to and from a web application. By inspecting web traffic, a WAF can be used to prevent application security flaws, such as SQL injection, cross-site scripting (XSS), file inclusion, and security misconfigurations.

136. One particularly nasty vulnerability—a vulnerability that can expose an entire data lake to significant risk—is a Server-Side Request Forgery ("SSRF") attack. In an SSRF attack, the attacker abuses functionality on a server to read or update internal resources that the attacker should not legitimately be able to access—such as sensitive data sitting in a data lake.

137. To understand an SSRF attack, it is first necessary to understand how modern servers fulfill requests made by web applications. Servers often expose application programing interfaces ("APIs"), which allow outside users to obtain information from the server by sending it formatted requests. In modern web applications, API requests are performed by submitting HyperText Transfer Protocol ("HTTP") requests—just like what a web browser sends to a remote server when a user clicks a link on a webpage. However, when a server receives an API request via HTTP, it may return structured data (for example, a JSON object) instead of a webpage.

138. In an SSRF attack, the attacker tricks a server, including the WAF, into thinking that the attacker is permitted to make a change to data on the server or to request data from the server. This sort of attack is typically used to obtain access to resources that are behind a WAF. By tricking a server into thinking that it is receiving a legitimate request for resources from inside the firewall (rather than an illegitimate request from outside), the attacker obtains a foothold inside the protected network.

139. Thus, an SSRF attacker can imitate a user-facing web application that makes internal requests to a server from behind a firewall, thereby obtaining access to shared resources within the firewall—such as a data lake.

140. Such an attack is an example of one of the vulnerabilities of an architecture that allows different applications to dynamically access data on a network, and it also lays bare the risks of pooling data to be shared across several applications and network resources. Put simply,

although data pooling and dynamic access allow for machine learning at scale, they necessarily expose an enterprise's internal resources and data to greater risk.

141. AWS has no protections built into its systems against this sort of vulnerability. Because Amazon uses IAM roles to control access to sensitive resources, an attacker who gains access to a resource behind a firewall that can assume a privileged IAM role can gain access to whatever is permitted under the policy that applies to that role.

142. This is a well-known flaw in AWS-based systems. Detailed guides exist online that demonstrate how to exploit AWS resources with SSRF attacks.

143. How well-known is this security flaw in AWS? In 2014, ***three separate presenters at two of the world's preeminent computer security conferences presented the SSRF security flaw onstage to thousands of computer security professionals around the world***. These presenters walked through—in detail—the devastating, systemic, and nefarious SSRF risk endemic to the AWS infrastructure.

144. For example, on March 21, 2014, at the Insomni'hack information security conference in Geneva Switzerland, security research Nicolas Gregoire gave a talk called "Lurking in Clouds: Easy Hacks for Complex Apps." In Gregoire's talk, he laid out what he called a "huge hole" in an AWS-based web application's security: an SSRF attack that used data leaked from an AWS metadata server and a privileged IAM role to obtain access to private data stored on AWS.

This export feature still has a ***huge*** hole

Any idea?

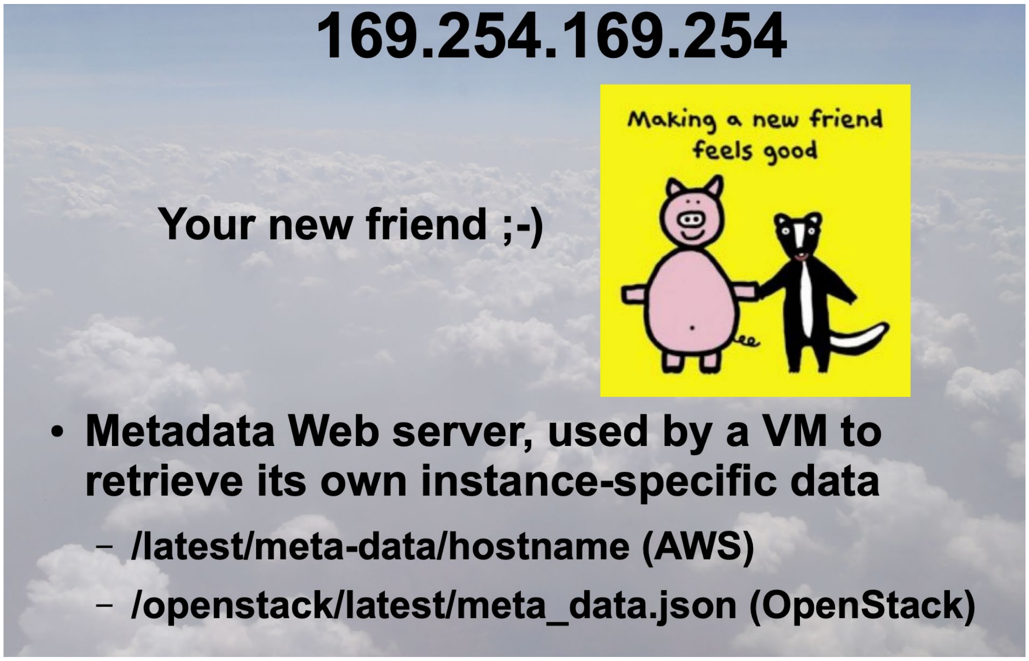

169.254.169.254

Your new friend ;-)

Making a new friend
feels good

- Metadata Web server, used by a VM to retrieve its own instance-specific data
  - /latest/meta-data/hostname (AWS)
  - /openstack/latest/meta_data.json (OpenStack)

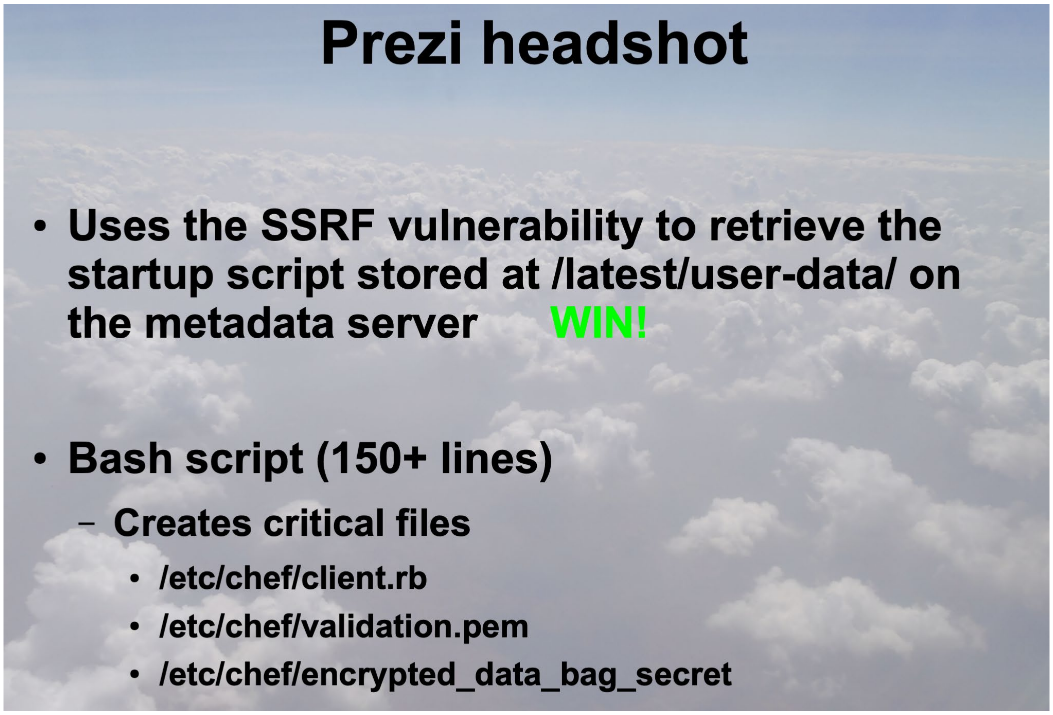

Prezi headshot

- **Uses the SSRF vulnerability to retrieve the startup script stored at /latest/user-data/ on the metadata server WIN!**
- **Bash script (150+ lines)**
  - **Creates critical files**
    - **/etc/chef/client.rb**
    - **/etc/chef/validation.pem**
    - **/etc/chef/encrypted_data_bag_secret**

145. In October 2014, at the BlackHat Europe conference in Amsterdam—the European edition of the world's preeminent computer security conference—security researcher Erik Peterson gave a lecture titled "Bringing a Machete to the Amazon." Peterson's lecture was seen live and/or via webcast by thousands of information security professionals around the world.

146. Peterson emphasized the unique security risks of AWS, including the risks inherent in large-scale migrations to the public cloud. He called this "Emergent Insecurity."

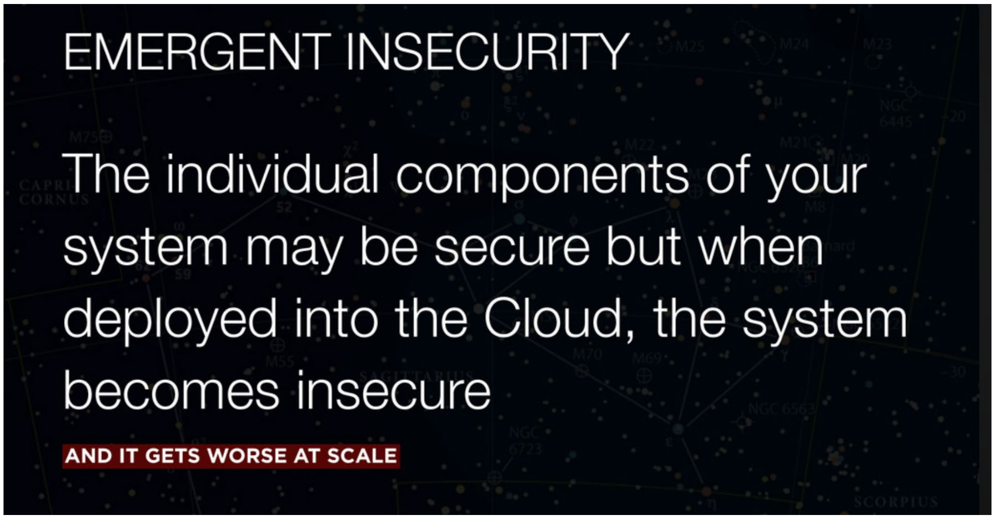

EMERGENT INSECURITY

The individual components of your system may be secure but when deployed into the Cloud, the system becomes insecure

AND IT GETS WORSE AT SCALE

147. Peterson warned that “AWS API’s operate outside of traditional security controls, [and] can make all existing controls irrelevant.” And he warned:

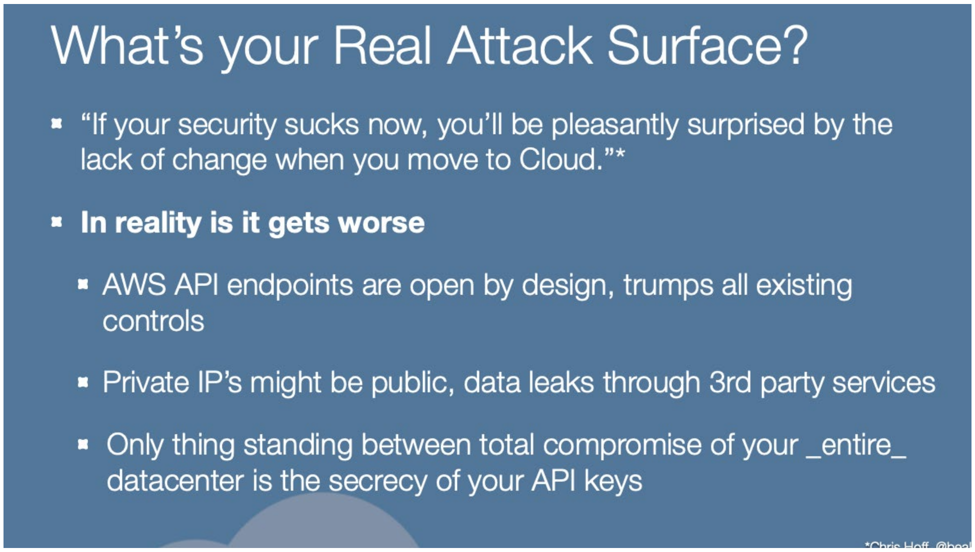

What’s your Real Attack Surface?

- “If your security sucks now, you’ll be pleasantly surprised by the lack of change when you move to Cloud.”*
- **In reality is it gets worse**
  - AWS API endpoints are open by design, trumps all existing controls
  - Private IP’s might be public, data leaks through 3rd party services
  - Only thing standing between total compromise of your _entire_ datacenter is the secrecy of your API keys

*Chris Hoff @beal

148. Then Peterson moved into particular vulnerabilities of the AWS system. He specifically, lengthily highlighted one particular vulnerability of AWS: the risk of leaked AWS metadata and improper IAM scoping allowing a devastating SSRF attack that “***can result in a total data center compromise in AWS***.”

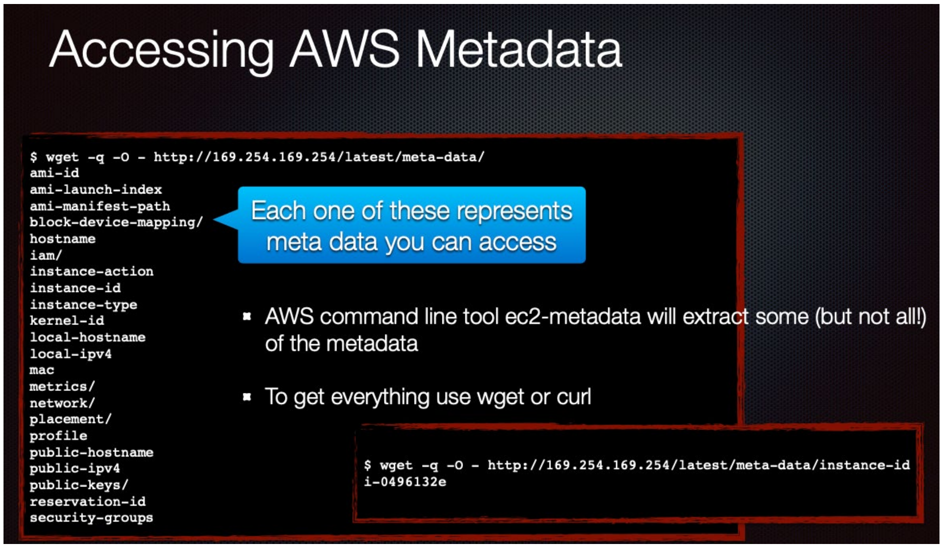

Accessing AWS Metadata

```
$ wget -q -O - http://169.254.169.254/latest/meta-data/
ami-id
ami-launch-index
ami-manifest-path
block-device-mapping/
hostname
iam/
instance-action
instance-id
instance-type
kernel-id
local-hostname
local-ipv4
mac
metrics/
network/
placement/
profile
public-hostname
public-ipv4
public-keys/
reservation-id
security-groups
```

Each one of these represents meta data you can access

- AWS command line tool ec2-metadata will extract some (but not all!) of the metadata
- To get everything use wget or curl

```
$ wget -q -O - http://169.254.169.254/latest/meta-data/instance-id
i-0496132e
```

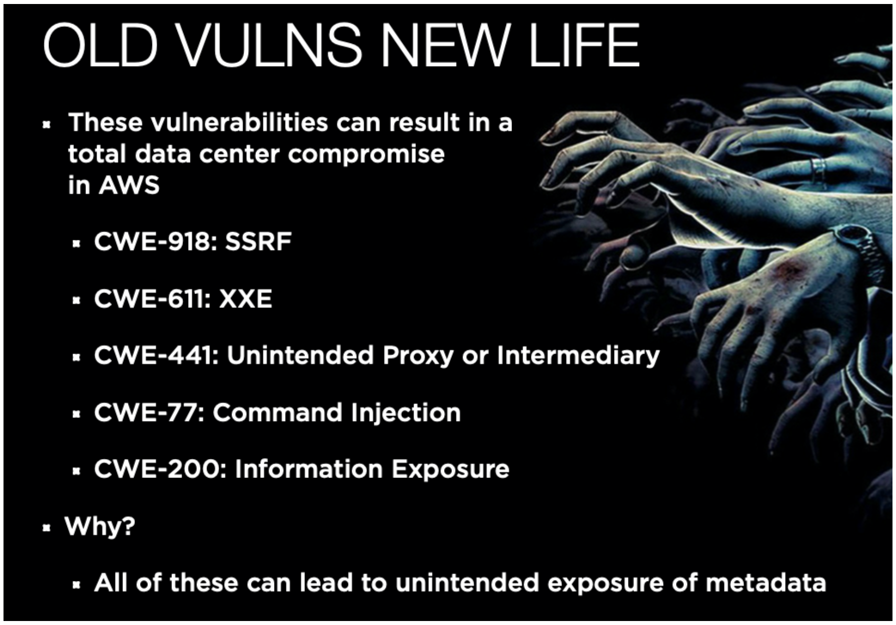

OLD VULNS NEW LIFE

- These vulnerabilities can result in a total data center compromise in AWS
  - CWE-918: SSRF
  - CWE-611: XXE
  - CWE-441: Unintended Proxy or Intermediary
  - CWE-77: Command Injection
  - CWE-200: Information Exposure
- Why?
  - All of these can lead to unintended exposure of metadata

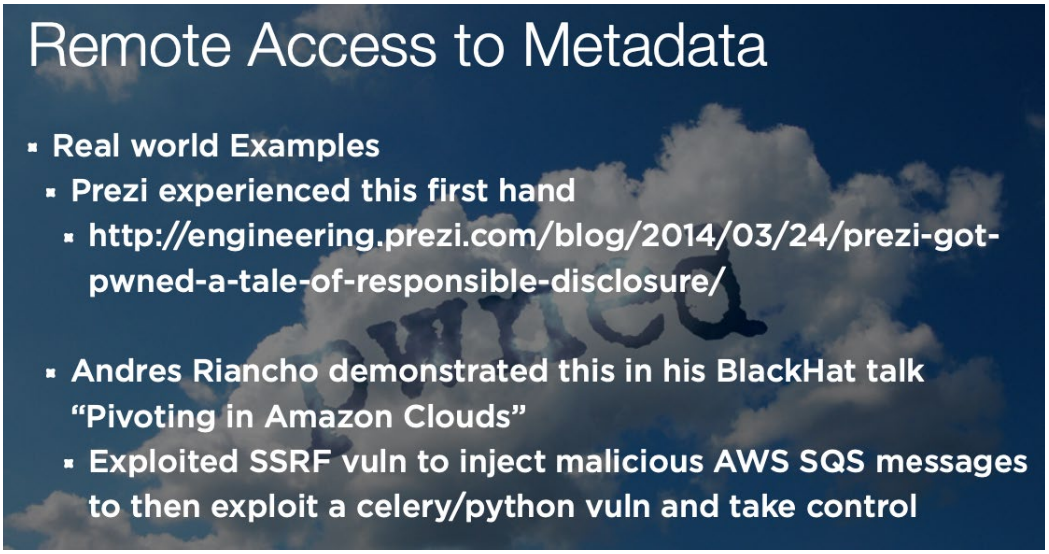

Remote Access to Metadata

- Real world Examples
  - Prezi experienced this first hand
    - http://engineering.prezi.com/blog/2014/03/24/prezi-got-pwned-a-tale-of-responsible-disclosure/
  - Andres Riancho demonstrated this in his BlackHat talk "Pivoting in Amazon Clouds"
    - Exploited SSRF vuln to inject malicious AWS SQS messages to then exploit a celery/python vuln and take control

149. Peterson's BlackHat Europe 2014 presentation on the SSRF vulnerability in AWS referenced ***another*** BlackHat presentation on ***the same*** security vulnerability from that year. Specifically, in August 2014 at the BlackHat USA 2014 conference in Las Vegas, security researcher Andrés Riancho actually ***demonstrated on stage*** an SSRF exploit through which leaked metadata and a privileged IAM role could be used to take control of all resources in an AWS account. Like Peterson's BlackHat Europe 2014 presentation, Riancho's presentation at BlackHat USA 2014 was seen by thousands of security researchers.

150. Riancho's BlackHat USA 2014 presentation on AWS security vulnerabilities, including the systemic SSRF risk baked into the AWS access management architecture, was titled "Pivoting in Amazon Clouds." Riancho explain his motivation for giving the lecture: "Mission critical applications are being deployed to the Amazon cloud and most information security experts have no clue about what needs to be inspected to make sure they are secure." According to Riancho, "classic security testing is not enough, knowledge about Amazon's EC2 instance life-cycle, user-data, IAM roles, and other Amazon cloud services are required when testing and exploiting Amazon cloud architectures." In his presentation, Riancho described the SSRF vulnerability in AWS as allowing an attacker "the keys to the kingdom."

151. In November 2014, Riancho gave ***another*** BlackHat presentation on AWS security—this one a BlackHat webcast. In this webcast, titled "Amazon AWS Security Basics: Escalating Privileges from EC2," Riancho described how "[m]isconfigured IAM profiles can be used to elevate the AWS user's privileges, perform DoS attacks and access private information."

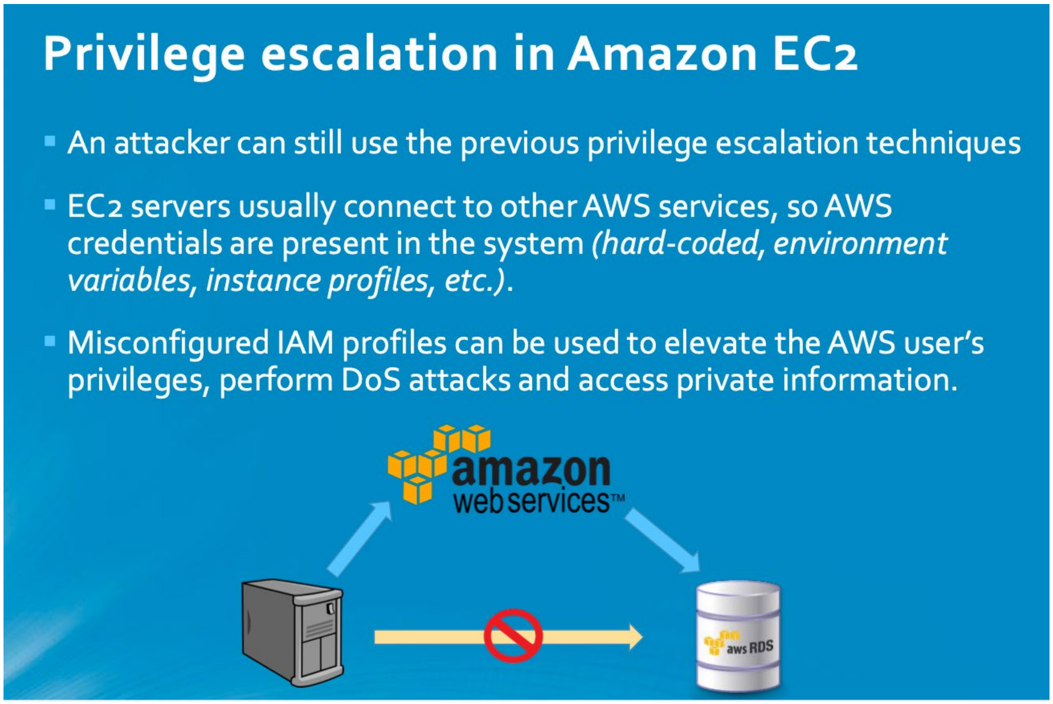

Privilege escalation in Amazon EC2

- An attacker can still use the previous privilege escalation techniques
- EC2 servers usually connect to other AWS services, so AWS credentials are present in the system *(hard-coded, environment variables, instance profiles, etc.)*.
- Misconfigured IAM profiles can be used to elevate the AWS user's privileges, perform DoS attacks and access private information.

152. In short, by the end of 2014, ***everyone*** in the computer security world knew about AWS's security shortcomings—and in particular, a gaping, endemic security hole baked into the very structure of AWS: ***a single improperly-scoped IAM role can give an attacker "the keys to the kingdom" for all resources in an AWS account, and an improperly-scoped IAM role is a virtual certainty in a complex, dynamically scaling AWS environment***.

153. AWS has been unwilling to fix this systemic vulnerability. For example, Google's cloud system allows the creation of special HTTP headers—additional information that is appended to a request to a server—that help ensure that only authorized requests are fulfilled. This, however, comes at a price. There must be strictly defined polices associated with such a header system, and most importantly, resource access cannot be dynamically modified with the same ease as an access control system in which general IAM roles are defined, and any resource can simply assume them. In other words, fixing the IAM (and thus SSRF) vulnerability would reduce the effectiveness of the machine learning capabilities of AWS. For AWS, which profits greatly from increased use of its machine learning tools and related services, the "bug" is a feature.

154. As a result of AWS's unwillingness to compromise profits for security, complex, dynamic AWS environments have continually been vulnerable to the well-known, devastating SSRF attack vector ***for more than five years after its widespread public disclosure***—and AWS is ***still just as vulnerable to this attack today***.

155. For example, in 2018, ***four years*** after this specific SSRF vulnerability was described at ***three different security conferences***, including BlackHat USA and BlackHat Europe, ***another BlackHat presentation described the same AWS vulnerability*** to SSRF.

156. In August 2018 at the BlackHat USA conference in Las Vegas, Netflix security researcher William Bengston gave a presentation titled "Detecting Credential Compromise in

AWS." Bengston noted that an attacker executing an SSRF attack on AWS could assume the IAM role of the targeted system, request information from inside the firewall, export that information, and avoid detection:

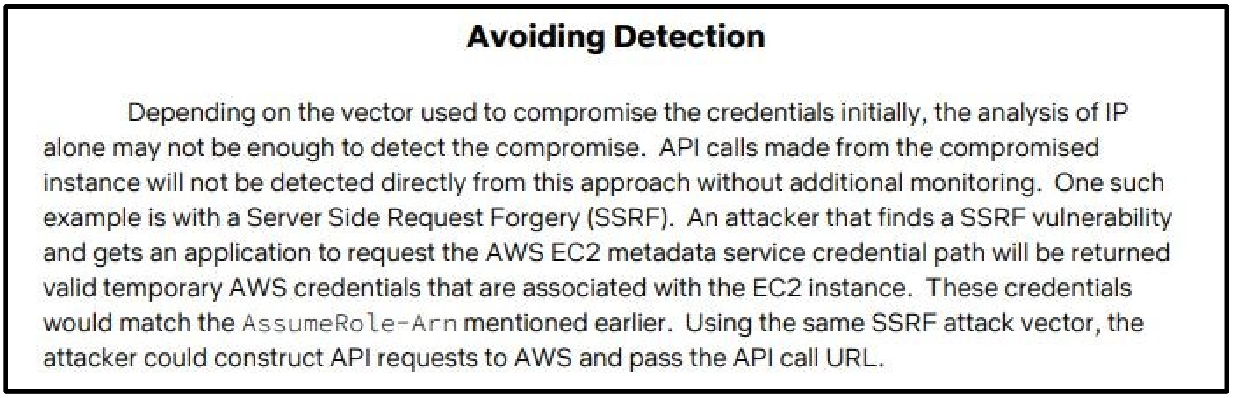

Avoiding Detection

Depending on the vector used to compromise the credentials initially, the analysis of IP alone may not be enough to detect the compromise. API calls made from the compromised instance will not be detected directly from this approach without additional monitoring. One such example is with a Server Side Request Forgery (SSRF). An attacker that finds a SSRF vulnerability and gets an application to request the AWS EC2 metadata service credential path will be returned valid temporary AWS credentials that are associated with the EC2 instance. These credentials would match the `AssumeRole-Arn` mentioned earlier. Using the same SSRF attack vector, the attacker could construct API requests to AWS and pass the API call URL.

157. Once inside, the attacker "can execute commands on the system directly," with all the IAM role permissions of the hacked system:

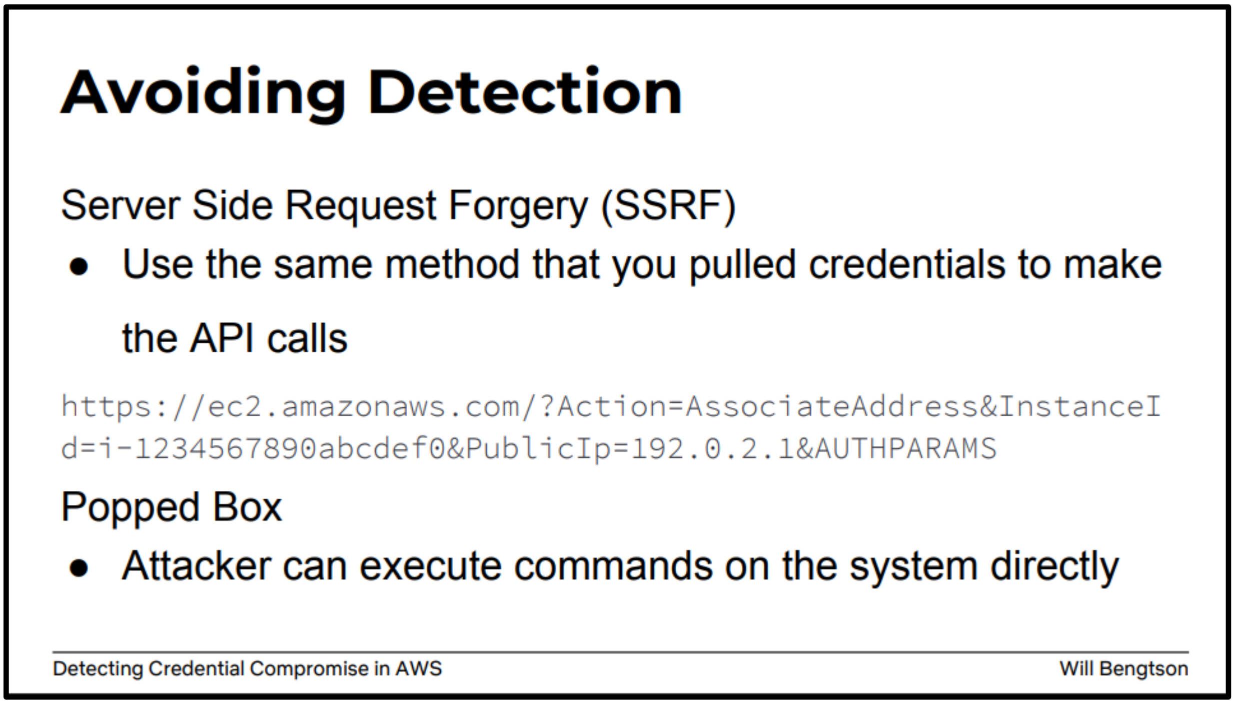

Avoiding Detection

Server Side Request Forgery (SSRF)

- Use the same method that you pulled credentials to make the API calls

```
https://ec2.amazonaws.com/?Action=AssociateAddress&InstanceI
d=i-1234567890abcdef0&PublicIp=192.0.2.1&AUTHPARAMS
```

Popped Box

- Attacker can execute commands on the system directly

Detecting Credential Compromise in AWS

Will Bengtson

158. Amazon has not fixed this type of AWS vulnerability, despite being on alert of the risk, because the openness and unfettered access to pooled data is a feature, not a bug. It allows AWS to sell dynamically created computing and storage resources and allows its machine learning

tools to be applied broadly to large pools of data. Granular and static access controls required to minimize the risk of data thefts would be inimical to the very design of the AWS cloud ecosystem.

159. As described in more detail later in this Complaint, Capital One and Amazon specifically—and falsely—stated this SSRF attack was fixed in their computer systems several times between 2015 and the present. And (as also described in more detail later in this Complaint) it was ***this exact SSRF attack vector*** that led to the devastating theft of Plaintiffs' and Class members' sensitive information from Capital One's AWS data lake.

B. Capital One Knew About the Risks of Pooling Sensitive Data in the AWS Cloud.

160. By the end of 2014, Capital One was determined to mine the massive amount of customer data it received as part of its operations.

161. Like other banks at the time, Capital One had traditionally relied on a dedicated server system for its storage and computing needs. For example, on March 12, 2014, Capital One opened a new data center in Chesterfield, VA, a 242,000 square foot facility that took 14 months to build. Such data centers are prohibitively expensive in both construction and upkeep.

162. In Capital One's peripheral vision, however, was a far lower cost option—the public cloud. Amazon was already operating massive data centers in Virginia, near where Capital One had built its data centers. AWS required a far lower upfront investment, had the option to purchase as much or as little computing and storage resources as needed, and allowed for data mining at a scale and price that Capital One could not hope to meet with a traditional dedicated server or private cloud model.

163. The risk of opting for AWS over Capital One's own data centers was clear to Capital One. It would mean: (1) moving sensitive customer data offsite to a third-party's servers; (2) ensuring the security of that data in a new and complicated system outside Capital One's direct

physical control; and (3) most importantly, convincing its customers, potential customers, and regulators that sensitive user data would be safe on the public cloud.

164. As an early adopter—in fact, the first large bank to migrate customer data to a public cloud for mining—Capital One was forced to rely heavily on Amazon's AWS trained engineers. As late as 2018, only 0.5% of IT professionals held an AWS certification, a number that was even lower when Capital One decided to migrate in 2015. This increased the risk to Capital One's valuable customer data.

165. That risk was further heightened by Capital One's plan to (i) migrate more than fifteen years (at least) of historical data on customers and potential customers to AWS to facilitate enhanced machine learning, (ii) pool all of the sensitive data it possessed into a data lake, and (iii) provide broad access to that data lake so that it could apply machine learning models to the data. Data pooling and mining on the scale offered by AWS would allow Capital One to glean insights about its current and prospective customers, providing it a competitive edge over other banks.

166. Capital One's new CIO, Rob Alexander, was also a proponent of Agile Software Development, a trendy and vague framework for developing software based on oft-repeated mantras, like, "responding to change over following a plan" and "[i]ndividuals and interactions over processes and tools." Alexander wanted to bring Agile to Capital One, and AWS allowed him to do so—namely, by enabling developers to write small software experiments without attention to institutional and procedural constraints.

167. To accomplish agile development, there would have to be open access to data. As Capital One's Linda Aplsey, VP of data engineering, told the publication CIO Dive in a January 31, 2019 interview, the transition enabled rapid and agile development, but Capital One developers were responsible for what she referred to as the four Ds:

> - Data movement: Data ingestion coming in and out of the company or moving within Capital One.
>
> - Data storage: Capital One has to store a large amount of data in secure, resilient platforms so it can perform as required.
>
> - Data discovery: The company has to ensure data is accessible.
>
> - Data cleansing: To make sure data is ready to use at scale, it required deduping, cleansing, and getting rid of training zeros.

168. Apsley, however, knew that the open access required for rapid and agile software development would come at the price of security. To assuage fears, she told CIO Dive that any security risk could be managed:

> Standardizing across languages, Capital One can use common tools with data across a much larger ecosystem. By extension, data management becomes easier. For example, a vulnerability scan could occur to make sure data does not have problems when it comes in, Apsley said.

169. Apsley's assurances were false and misleading. The larger the ecosystem, the more difficult it becomes to enforce policies and to manage massive amounts of data. It did not become "easier" as she had stated to manage more interacting resources.

170. As Erik Peterson put it in his BlackHat Europe 2014 presentation on AWS security risks:

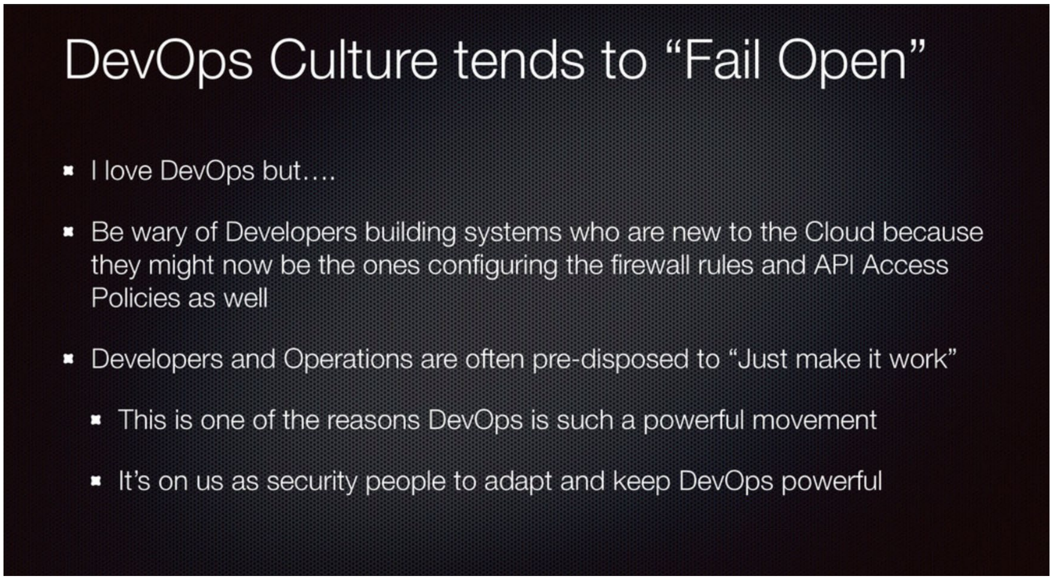

DevOps Culture tends to “Fail Open”

- I love DevOps but….
- Be wary of Developers building systems who are new to the Cloud because they might now be the ones configuring the firewall rules and API Access Policies as well
- Developers and Operations are often pre-disposed to “Just make it work”
  - This is one of the reasons DevOps is such a powerful movement
  - It’s on us as security people to adapt and keep DevOps powerful

171. Contrary to Apsley’s false statements, in fact Capital One’s move to the cloud and its adoption of agile software development implicated “emergent insecurity.”

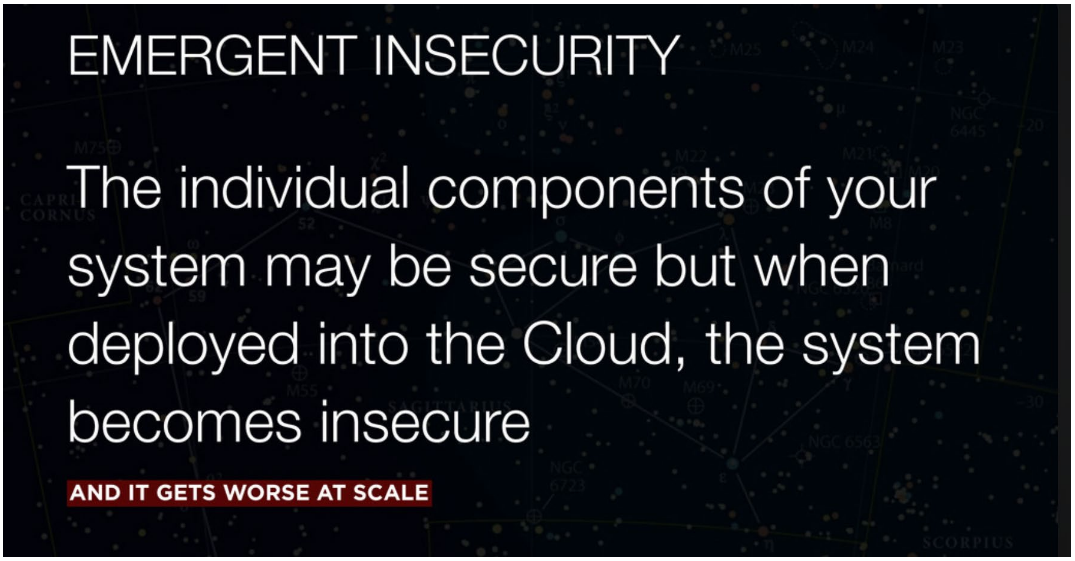

EMERGENT INSECURITY

The individual components of your system may be secure but when deployed into the Cloud, the system becomes insecure

AND IT GETS WORSE AT SCALE

172. In reality, migrating to the cloud and adopting common AWS tools doesn’t make things easier, more manageable, or more secure for the protection of sensitive data: “In reality . . . ***it gets worse***.”

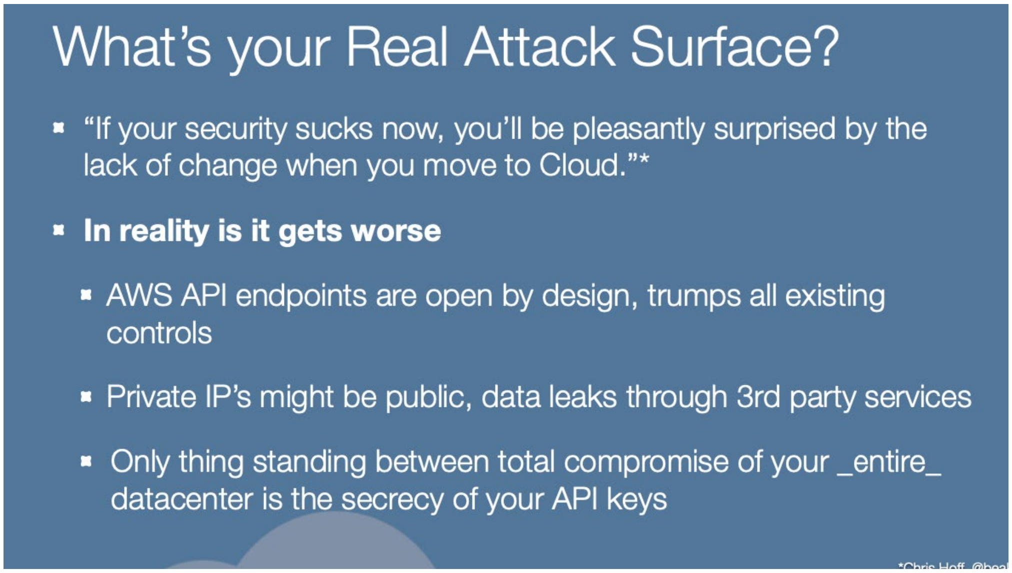

What's your Real Attack Surface?

- "If your security sucks now, you'll be pleasantly surprised by the lack of change when you move to Cloud."*
- **In reality is it gets worse**
  - AWS API endpoints are open by design, trumps all existing controls
  - Private IP's might be public, data leaks through 3rd party services
  - Only thing standing between total compromise of your _entire_ datacenter is the secrecy of your API keys

*Chris Hoff, @bea

173. The actual risks and difficulties of migration to the cloud and agile software development were well-known to Apsley at the time she made her statements to CIO Dive: she lied, in order to conceal the very real, very substantial risks and problems of Capital One's migration to the cloud and internal development policies and practices.

174. Apsley also falsely touted the ability to perform automated vulnerability scans to ensure that data would be secure. All of this was to provide the bank with cover so it could obtain a competitive edge over its competitors.

175. That competitive edge, however, only existed because Capital One's competitors recognized and refused the risks of migrating sensitive data to a cloud ecosystem like AWS. For example, Bank of America, the second largest bank in the United States, had been reticent to use the public cloud, and industry executives cautioned that putting sensitive user data on the cloud could put that data at risk of both a data breach and compliance issues.

176. Capital One clearly understood those risks: A cloud-based strategy meant placing massive amounts of that sensitive user data on the public cloud—facing hackers, terrorists, and

foreign governments. Indeed, that is precisely what Capital One told its investors in its 2016 Annual Report:

> Further, cyber and information security risks for large financial institutions like us have generally increased in recent years in part because of the proliferation of new technologies, the use of the Internet and telecommunications technologies to conduct financial transactions and the increased sophistication and activities of organized crime, perpetrators of fraud, hackers, terrorists, activists, formal and informal instrumentalities of foreign governments and other external parties.

177. As Capital One made clear to its investors, cybersecurity threats by a myriad of bad actors were increasing in speed, scale, and sophistication. As Capital One's CIO Rob Alexander explained, the migration to AWS would have to address the threat posed by cybercriminals and address cybersecurity head on:

> Of course, security is critical for us. The financial services industry attracts some of the worst cyber criminals. So we work closely with the Amazon team to develop a security model, which we believe enables us to operate more security in the public cloud than we can even in our own datacenters.

C. Capital One Moves to Amazon's AWS

1. Capital One and Amazon Partner to Move Capital One's Data to the Cloud.

178. Capital One was motivated to migrate all of its sensitive data to the AWS public cloud in order to leverage that data using the maching-learning tools and computation power that AWS provided. AWS was motivated to attract a giant bank like Capital One, lock it into the AWS ecosystem, and supercharge the positive feedback loop. Both Capital One and Amazon knew that pooling sensitive customer data was risky given the inherent vulnerabilities in data lakes and AWS's architecture. Moreover, they knew that Capital One's customers, potential customers, and regulators would recognize the risk Capital One was running by being the first major bank to migrate to the public cloud.

179. In view of the above, Capital One and Amazon knew they had to assuage customer fears of the public cloud any way they could. The way they settled on was more deception. Specifically, Capital One and Amazon worked together to create a smokescreen: a software tool that Capital One and Amazon would falsely announce had solved the security problems inherent in using the AWS cloud for machine learning at scale—specifically, the SSRF risks tied to dynamic access and data pooling on the AWS platform. And thus was Cloud Custodian born.

2. Cloud Custodian: Amazon and Capital One's Potemkin Village

180. On October 7, 2015, Capital One became one of the largest corporations, and the only large bank, to publicly adopt AWS as its public cloud platform. It unveiled its cloud strategy on the stage of Amazon's 2015 AWS re:Invent Conference—a conference dedicated to Amazon's AWS cloud computing platform. At the conference, Capital One's CIO Rob Alexander showcased an unorthodox plan to migrate its customers' data to Amazon's owned and operated servers. He announced:

> We've expanded our use of AWS from development and test more broadly and this year we've taken a more aggressive stance, recognizing that we can deploy some of our most critical production workloads on the AWS platform. This is a game changer for Capital One. We recognize that we want to be in the business of building great applications for our customers, not in investing to build costly and complex infrastructure. As a consequence, we are focusing on consolidating and rationalizing our datacenter footprint. . . .

181. The strategy was a hasty move into uncharted territory for a major bank. It would require a ***complete*** migration to Amazon's cloud. For the first time, user data would no longer be in the bank's custody. It would be in the hands of a third-party partner, Amazon.

182. And this came on the heels of widely-publicized, ***repeated*** security warnings about a particular, systemic, and by now well-known data theft risk endemic to AWS: the near-

impossibility of protecting sensitive data against devastating SSRF attacks in a dynamically accessible, data-pooled AWS environment.

183. Because of the significant—and by now well-known—systemic security risks associated with the migration of Capital One's sensitive customer data to the AWS cloud, Amazon and Capital One devised a plan to calm the masses: they would create a software tool that would purport to eliminate the security risks posed by dynamic access and data pooling. On an aws.amazon.com website, Capital One and Amazon together described Capital One's transition to the cloud, stating that the companies had together implemented a "risk framework" that "met the same high bar for security and compliance" met in on-premises environments:

> To implement the resulting cloud risk framework, Capital One relied on both people and technology. "One key early step we took was to establish a cloud governance function, consisting of risk managers and cloud engineers, to curate capabilities and controls that would keep us well managed as we moved applications into the cloud," says Brady, adding that this team has continued to update and refine the cloud-risk-control framework quarterly. "We developed and open-sourced a compliance enforcement engine called Cloud Custodian, to automate detection and correction of policy violations so we could keep our teams inside the guardrails without restricting their ability to work creatively and innovate for our customers. We also built a reporting portal where we can see and measure compliance in the entire fleet of services throughout our complex, multi-account environment."

184. Specifically, AWS and Capital One announced their joint development and marketing of the Cloud Custodian system, which was released in April 2016 as an open source project developed by both Amazon and Capital One and maintained by Capital One.

185. Cloud Custodian would purportedly serve as an overarching system that would manage applications and their access to underlying customer data. The central purpose of the system was to manage IAM roles and ensure that policies for access to customer data stored in

Amazon's S3 storage systems would be correct and that data would not be exposed or queried beyond what was necessary for the ordinary operations of Capital One's applications.

186. AWS and Capital One, however, knew that their statements about the new software were false and misleading. In particular, Cloud Custodian would not actually solve the problems endemic to ***Capital One's*** AWS ecosystem, which relied on massive scaling, rapidly-changing dynamic access through IAM roles, and vast pooling of customer resources on Amazon's S3 servers to facilitate machine learning. These features of Capital One's use of AWS left massive amounts of sensitive data vulnerable to attack and theft from the same well-known SSRF vector AWS and Capital One purported to resolve through Cloud Custodian. AWS and Capital One's statements that they were developing a "compliance enforcement engine" that would "automate detection and correction of policy violations so we could keep our teams inside the guardrails without restricting their ability to work creatively and innovate for our customers" was not feasible for Capital One's large scale and sensitive data.

187. In fact no one had developed—or could develop—such software for an AWS customer like Capital One—not even Amazon. This is because the ***AWS architecture itself*** allows for HTTP requests to internal resources from outside of a firewall—and once this access is obtained, IAM roles broad and flexible enough to facilitate machine learning and data sharing at scale cannot prevent an attacker from accessing a wide range of customer data.

188. The very premise of the Cloud Custodian was flatly misleading. The most devastating security risk for Capital One data on AWS was not (and is not) that IAM roles would be misconfigured with the wrong policies (although that was and remains a significant risk). The most devastating—and unfixable, under AWS architecture—security risk for Capital One data on AWS is that ***IAM roles have to be broadly defined to quickly and dynamically allow machine***

learning systems to access the data lake. If Capital One had configured IAM roles with the broad policies necessary for it to conduct the data mining it wanted to do on pooled customer data—its central reason for migrating to AWS in the first place—there would be nothing for Cloud Custodian to detect and "fix." If an attacker obtains access to resources within the Capital One AWS firewall, the bank's entire data lake is potentially vulnerable, and Cloud Custodian would and will do nothing to mitigate that risk. In other words, Cloud Custodian was in no way a "risk framework" that would permit Capital One's safe entry to the public cloud in the manner described by Capital One and Amazon.

189. The Cloud Custodian documentation describes the system as a rule and policy engine that was designed to automatedly scan all of Capital One's cloud resources to ensure that permissions for resources were properly set and that access controls, including firewall settings, were correctly configured:

> Cloud Custodian is a tool that unifies the dozens of tools and scripts most organizations use for managing their public cloud accounts into one open source tool. It uses a stateless rules engine for policy definition and enforcement, with metrics, structured outputs and detailed reporting for clouds infrastructure. It integrates tightly with serverless runtimes to provide real time remediation / response with low operational overhead

190. Importantly, Cloud Custodian allows the application of bulk actions to a set of resources. Thus, if there is a misconfiguration, Cloud Custodian purportedly allows Capital One to detect it and fix it in real time. But again, the very premise of this functionality was flawed. A deliberately but broadly configured IAM role—required for machine learning at scale on pooled Capital One customer data—will not appear anomalous to Cloud Custodian ***because it is a feature and not a bug***.

191. In a podcast interview on July 14, 2017, Capital One Senior Distinguished Engineer Kapil Thangavelu misleadingly explained the purported central purpose of Cloud Custodian:

> Q. For people out there who are not familiar. What is Cloud Custodian and why did you develop it and what is its role today?
>
> A. We are as an industry and across industries moving into cloud and we want developers to use the cloud, ***but developers sometimes do silly things like leave their database publicly accessible on the Internet or other things along those lines.*** We want to let them use the native cloud experience. We don't want to hide the cloud so to speak and get them down to the lowest common denominator. But we want to put guard rails on. And we want to do guard rails not through process and checklists but through automation that's real time and reactive that puts guard rails on that puts developers in a safe space to be productive, without having to do the manual side. Automated guard rails.

192. Thangavelu was clear that the purpose of Cloud Custodian was to ensure that Capital One's developers would not do "silly things" such as leave a resource containing sensitive data improperly accessible. Cloud Custodian would serve as the guard rails for the cloud infrastructure and for Capital One's developers as they designed Capital One's customer-facing applications.

193. Thangavelu's statements were false and misleading. The pooling of sensitive data on the public cloud meant that once an attacker outside of a firewall obtained access to internal resources, the data lake was vulnerable: all of its information was for the taking if the attacker could assume a broad IAM role. Cloud Custodian did not fix this problem. Moreover, Amazon's particular vulnerability to SSRF attacks meant that obtaining access to systems inside the firewall remained simple for most attackers—indeed, it was a vulnerability that had been exploited by hackers for years, beginning well before Capital One migrated to AWS. Thangavelu's statements were misleading because they ignored all of this in favor of a false narrative that any risk could be mitigated with Cloud Custodian.

194. At Amazon's yearly re:Invent conference in November 2018, Capital One's Thangavelu gave an important presentation showcasing Cloud Custodian. The presentation

highlighted all of the important features that Cloud Custodian had to offer. Several minutes into his presentation, Thangavelu began discussing Identity and Access Management ("IAM") roles—a system of roles associated with resources on AWS's cloud system.

195. These IAM roles allow a server, application, or other cloud resource to assume a user role with a specified set of permissions. For example, an AWS virtual server, called an EC2 server, can assume an IAM role providing it access to all of the data stored on a storage server, called an S3 server.

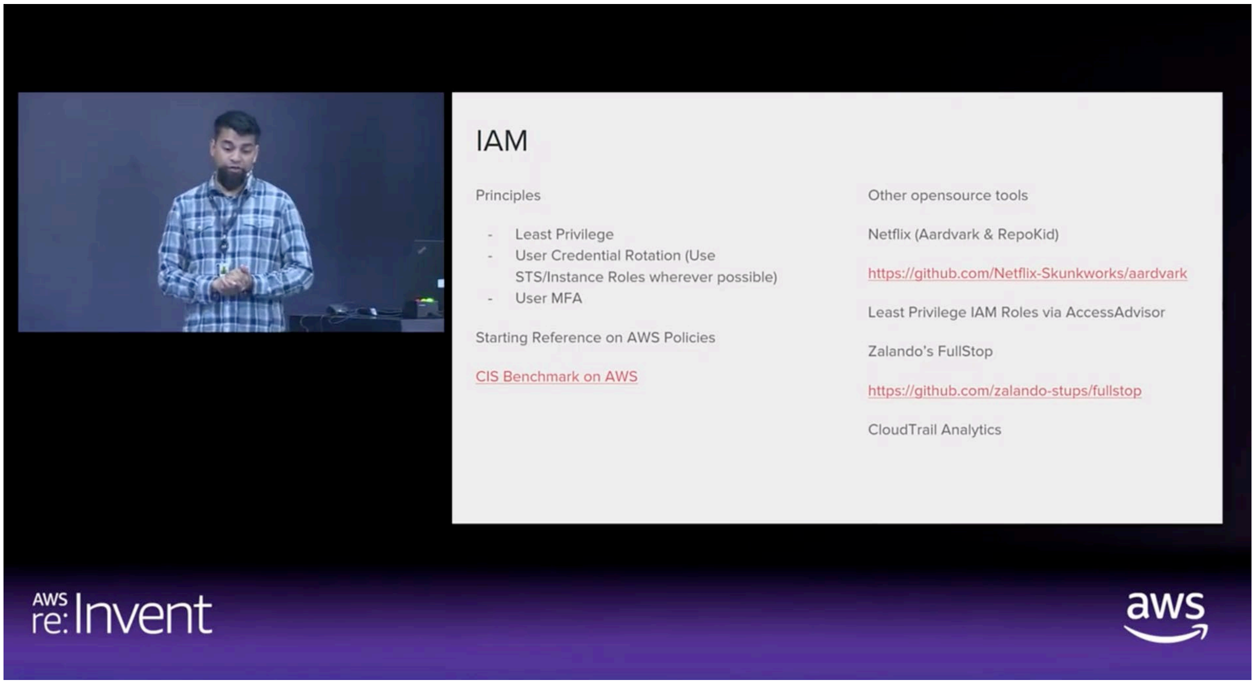

IAM

Principles

- Least Privilege
- User Credential Rotation (Use STS/Instance Roles wherever possible)
- User MFA

Starting Reference on AWS Policies

CIS Benchmark on AWS

Other opensource tools

Netflix (Aardvark & RepoKid)

https://github.com/Netflix-Skunkworks/aardvark

Least Privilege IAM Roles via AccessAdvisor

Zalando's FullStop

https://github.com/zalando-stups/fullstop

CloudTrail Analytics

196. In an astonishingly prescient part of his speech, Thangavelu described the precise vulnerability that would result in a massive data theft ***of his own customers' data*** the next year:

> In the cloud, all these resources are just available via URL and those are part of your network boundary. And ***those resources that have embedded IAM policies need special care and attention because they can be enabled to be accessible outside of your account***. I think ***everyone is familiar with some of the things around S3*** but that extends out to a lot of the other resources I called out a couple here.

197. The significance of this statement cannot be overstated—Capital One's senior engineer was clear that IAM policies needed "***special care***" precisely because outside accounts should not have access to certain resources. He then made clear that the most open and notorious issue was improper configurations that would provide access to data stored on one of Amazon's S3 storage servers.

198. Indeed, Mr. Thangavelu plainly stated in his 2018 presentation that "***everyone is familiar***" with the potential to misconfigure an IAM role to allow an outside resource full and unfettered access to sensitive data stored on Amazon S3 storage systems.

199. In January 2019, approximately one month after his presentation at re:Invent 2018, Thangavelu left Capital One and joined Amazon AWS as a Principal OpenSource Technologist.

200. All of this shows that Capital One was aware of the potential for precisely the problem that would later allow a data theft of Capital One's customer data in March 2019. It was this vulnerability—known to ***everyone***, requiring ***special care***, and that Amazon and Capital One had falsely stated Cloud Custodian would detect and remediate in real time—that would allow the theft.

201. There is, therefore, no question that Capital One was aware of this risk, knew that others were aware of the risk, and pretended to design special software to deal with the risk.

202. Capital One and AWS clearly appreciated an obvious risk that a broadly configured IAM role, if assumed from inside the firewall, would grant full access to user data stored in its data lake, and Capital One brandished its Cloud Custodian infrastructure as a purported means of preventing precisely the sort of attack that occurred.

203. Cloud Custodian did no such thing for Capital One's dynamically accessible, pooled user data on AWS—and Capital One and AWS knew it. Cloud Custodian was nothing

more than a means of lulling customers into a false sense of security. No automatic policy scanning and correction was possible when the very design of Capital One's AWS cloud architecture called for dynamic access across a large number of applications and the reckless pooling of massive amounts of sensitive data in one place. And nothing Cloud Custodian did mitigated the everpresent risk of an SSRF attack on an AWS web application.

204. That, of course, did not matter—because minimizing the risk was not the real purpose of Cloud Custodian. The real purpose was to assuage the concerns of customers and regulators so that Capital One and AWS could make immense profits from mining Capital One user data using AWS machine learning tools.

205. It made no sense for AWS to help develop Cloud Custodian as a stand-alone, open-source tool, because AWS had already developed similar cloud tools that it was selling to its customers. Amazon Macie, for example, was a machine learning system designed specifically to, among other things, protect personally identifiable data stored on Amazon S3 servers. As Amazon explains on its website:

> Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. ***Amazon Macie recognizes sensitive data such as personally identifiable information (PII)*** or intellectual property, and provides you with dashboards and ***alerts that give visibility into how this data is being access or moved***. The fully managed services continuously monitors data access activities for anomalies, and generates detailed alerts when it ***detects risk of unauthorized access or inadvertent data leaks***. Today, Amazon Macie ***is available to protect data stored in Amazon S3***, with support for additional AWS data stores coming later this year.

206. AWS's Access Advisor purportedly helps identify and reduce permissions and privileges to no more than necessary. AWS's GuardDuty purportedly alerts companies when someone is scanning for potentially vulnerable systems or moving unusually large amounts of data

to or from unexpected places. The AWS web access firewall purportedly detects common exploitation techniques, including server-side attacks like the one that led to the Data Theft.

207. Indeed, in a presentation on May 3, 2019, near the time of the theft at issue here, Amazon's Jonathan Allen, Enterprise Strategist at AWS gave a presentation at the AWS Summit. His presentation listed in detail all of the "tools and automated tasks" that "enhance security" on Amazon's public cloud.

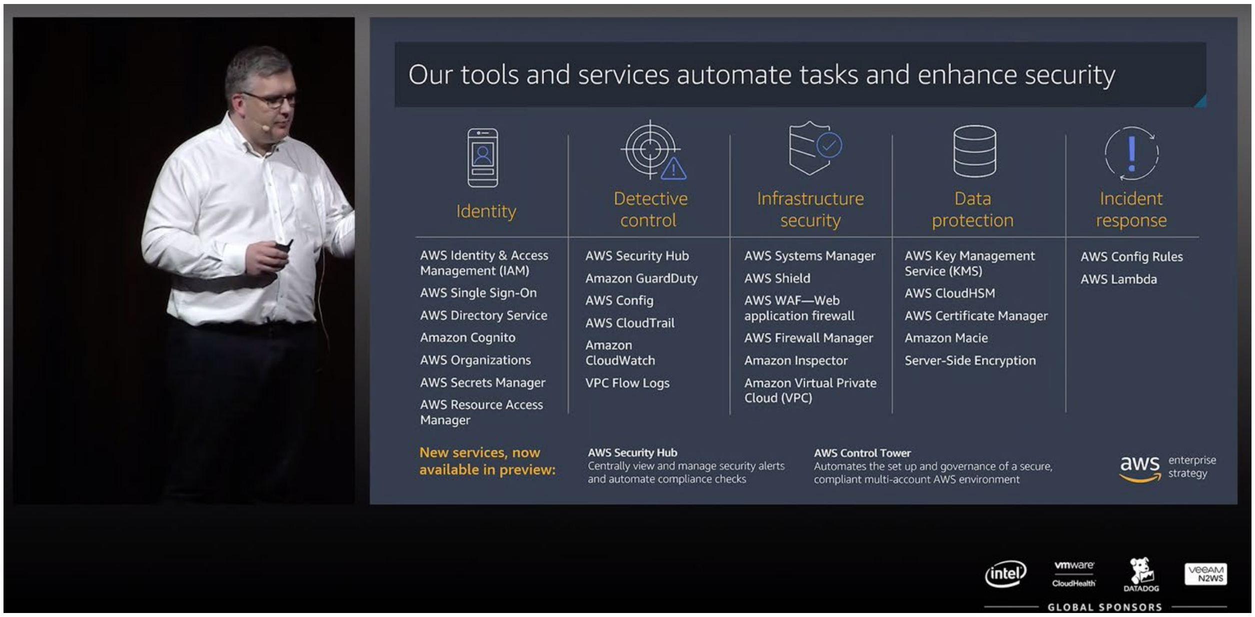

Our tools and services automate tasks and enhance security

Identity
- AWS Identity & Access Management (IAM)
- AWS Single Sign-On
- AWS Directory Service
- Amazon Cognito
- AWS Organizations
- AWS Secrets Manager
- AWS Resource Access Manager

Detective control
- AWS Security Hub
- Amazon GuardDuty
- AWS Config
- AWS CloudTrail
- Amazon CloudWatch
- VPC Flow Logs

Infrastructure security
- AWS Systems Manager
- AWS Shield
- AWS WAF—Web application firewall
- AWS Firewall Manager
- Amazon Inspector
- Amazon Virtual Private Cloud (VPC)

Data protection
- AWS Key Management Service (KMS)
- AWS CloudHSM
- AWS Certificate Manager
- Amazon Macie
- Server-Side Encryption

Incident response
- AWS Config Rules
- AWS Lambda

New services, now available in preview:

AWS Security Hub
Centrally view and manage security alerts and automate compliance checks

AWS Control Tower
Automates the set up and governance of a secure, compliant multi-account AWS environment

208. Allen's presentation addressed, for example, "AWS Identity & Access Management" and AWS's firewall settings and configuration rules—precisely what Cloud Custodian was purportedly built to accomplish. All of these tools were accessible and configurable through Amazon's existing architecture, and Amazon made money selling these tools.

209. Cloud Custodian also made no sense to AWS given its stated portability and open source nature. The system purportedly works with Google's GCP and Microsoft's Azure in addition to Amazon's AWS. Amazon had no rational economic interest in facilitating the use of competing public cloud ecosystems.

210. Amazon's bargain, then, was different than a normal service provider-to-customer relationship: In exchange for helping Capital One migrate to the AWS ecosystem under the cover of a software system that would keep their data safe, AWS received a large bank as a customer, which not only would increase the likelihood that others would also adopt the platform, but would also result in significant fees for AWS—fees Capital One could only pay AWS if its customers were lulled into a false sense of security.

3. Capital One Migrates to the AWS Cloud and Applies Machine-Learning to Customer Data under the Cover Provided by Cloud Custodian

211. Capital One migrated ***all*** of its most sensitive data to AWS's systems under the cover of its Cloud Custodian façade, creating an S3-based data lake. The migration meant moving ***everything*** to the AWS cloud and employing the inherently vulnerable dynamic access and data pooling architecture.

212. As Capital One's VP of data engineering, Linda Apsley, later explained in a 2017 interview, the strategy required a widespread migration to the cloud:

> It's been an amazing journey, I think, for the company to take this step of saying we're going to take our application and our entire eco-system and go cloud native—***put everything in the cloud***.

213. Applying machine learning to the data Capital One's customers are forced to provide as part of their application for, and use of, Capital One's products, allows Capital One to glean immense insight from its users' actions.

214. As Apsley explained in the same 2017 interview:

> The key thing on that question of where we're going is really machine learning and our anticipation of how that will help us understand the most important and innovative solutions we can give to our customers.

215. Specifically, understanding a customer's spending habits leads to a more profitable interchange fee business, allows precise calibration of rewards needed to induce a cardholder to spend using the card, and allows a better measurement of lending risk. Machine learning also allows an AI-driven system to interact with customers, reducing overhead costs.

216. Indeed, Amazon and Capital One's promotional video entitled, "AWS Technology Enables Capital One's Move to Machine Learning" (posted on Youtube on August 24, 2018) jointly articulated that vision. In that video, Capital One's Steph Hay, Head of the Company's "Conversation Design," explained that cloud-based access to data was critical for the Company's virtual assistant program, a system Capital One was developing to interact with customers:

> The elements that are critical to [Capital One's personal assistant] success are access to data—rich data that enables us to see across time and channels, the kind of customer behaviors that would allow us to design human experiences.

217. In that same video, the companies boasted about Capital One's transition to the cloud and its ability to leverage artificial intelligence to learn from the data Capital One collects about and from its customers. As George Brady, EVP and Chief Technology Officer of Capital One explained:

> AWS technologies have enabled our move to machine learning in a number of ways. One our – ***what underlies our data lake is S3, so just opening up an ability to store so much more historical information for us, which is just fuel for our models***, and models we've already built and will build in the future. The second area if I tie it to the GPU capabilities that Amazon has released and the services around that—the ability to train our models and continue to learn is really really important.

218. AWS had provided Capital One with a data lake that would allow it to pool the historical data it had collected and to leverage arrays of GPUs hosted on AWS to perform the mathematical operations required for machine learning.

219. The joint AWS and Capital One statements were misleading because they stated the benefits of their architecture by lying about, and omitting, many of the risks. That is, the companies failed to disclose that this architecture came at the price of the security of sensitive customer data. Data pooled into a data lake was broadly accessible by other resources that ***required*** access to the data for machine learning. If any system that was permitted to access the data lake was compromised, massive amounts of historical customer data would necessarily be compromised as well. Cloud Custodian, which theoretically could be configured to allow only minimal access to sensitive data, did nothing to mitigate this risk.

220. By the time of Capital One's 2017 Annual Report, the company's mission to capture the AI and machine learning advantage was clear. Capital One disclosed to investors that it was adopting a new approach to development—one that harnessed the power of artificial intelligence:

> We also have transformed how we deliver software. In 2017, we made significant strides in promoting our DevOps culture and building APIs and micro-services.We continued to deploy and contribute to open source across the company. We're harnessing the power of artificial intelligence, such as our natural language processing engine, which is the backbone for new customer experiences.

221. At the end of 2018, Capital One boasted in its year-end report to investors that it had obtained leadership in the use of cloud computing and rapid development in the banking sector:

> The vast majority of our operating and customer-facing applications operate in the cloud, which unleashes our associates to design real-time, intelligent experiences that work backwards from our customers needs. We are now considered one of the most cloud-forward companies in the world. In 2018, we made significant progress on our technology journey, and we will continue to invest to transform our infrastructure, data and technology tools.

222. Capital One, competing within an industry that was inherently and rightfully skeptical of the public cloud, had within a few years purportedly completed its transition to be "considered one of the most cloud-forward companies in the world." The massive amount of information it collected and maintained could now be mined and leveraged to increase profitability and reduce overhead. Its rapid entry into the world of machine learning and cloud computing, however, came at the price of the security of its customers' most sensitive information.

IV. THE 2019 DATA THEFT

A. Hacker, Paige Thompson, Exploits Capital One's Inherently Flawed Cloud-Based System.

223. On March 12, 2019 a hacker obtained a foothold behind Capital One's firewall, gaining access to an EC2 server connected to Amazon's data lake.

224. Several computer security professionals concluded after the attack that Thompson likely compromised Capital One's EC2 server by using an SSRF attack. As well-known computer security reporter, Brian Krebs, reported on his blog on August 19, 2019:

> The type of vulnerability exploited by the intruder in the Capital One hack is a well-known method called a "Server Side Request Forgery" (SSRF) attack, in which a server (in this case, CapOne's WAF) can be tricked into running commands that it should never have been permitted to run, including those that allow it to talk to the metadata service.

225. Thompson, a former AWS employee, had developed her own scanning software that would search AWS servers for systems vulnerable to an SSRF attack. She quickly honed in on Capital One's web applications, among those of several dozen other companies.

226. Once she obtained access to resources inside Capital One's firewall, she assumed an IAM role that allowed access to user data stored on Amazon S3 servers that comprised the data lake. Specifically, on or about March 22, 2019, she was able to list all of the resources available

to the IAM role she had obtained, then did so again on or about April 21, 2019. The data available to her included ***over fifteen years of credit card applications***.

227. The hacker then proceeded to route the data she obtained through The Onion Router ("TOR") nodes—a system originally developed by the United States Navy to obfuscate the source and destination of data traveling over the internet. The TOR network disguises a user's identity by moving traffic across different TOR servers and encrypting that traffic so it cannot be traced back to the user.

228. The standard industry practice is to block a cloud server's access to TOR's "exit" nodes, preventing data from leaving the server through an obfuscated and encrypted channel. These exit nodes are publicly identified and traffic from them is frequently blacklisted. Indeed, of the blacklists surveyed by academics, 88% of TOR relays appeared on the blacklists.

229. Capital One, however, had apparently failed to block TOR exit nodes, as the hacker was able to freely and anonymously route information through them. The hacker could do this because once she gained control over a machine with an IAM role that provided access to the data lake, she could route the data to servers outside of Capital One's firewall, including TOR nodes.

230. Put simply, because Capital One had pooled massive amounts of data onto Amazon's S3 servers (i.e. the data lake) and broadly permissioned servers to access that data in order to more efficiently data mine, obtaining access to a server inside the firewall meant obtaining access to ***all*** of the data in the data lake.

231. On June 18, 2019, the hacker identified herself on Twitter, bearing the handle, "erratic/Thompson." She sent a Twitter message stating, "Ive [*sic*] basically strapped myself with a bomb vest, fucking dropping capitol ones [*sic*] dox and admitting it." The word "dox" is an

informal term of art in the hacking community that refers to the publication of private information, typically with malicious intent.

232. erratic/Thompson then wrote "I wanna distribute those buckets [of Personal Information] I think first." Two minutes later, erratic/Thompson followed up, saying "[t]here [are] ssns"—referring to social security numbers—"with full name and dob [dates of birth]." The FBI understands these communications to mean that erratic/Thompson intended to disseminate data stolen from Capital One in the Data Theft.

233. On or about June 26, 2019, erratic/Thompson posted a message—later reviewed by the FBI—in the cloud-based collaboration service Slack, in which erratic/Thompson indicated that she was in possession of files she had extracted from Capital One.

B. Capital One Discovers the Data Theft

234. Capital One solicits and receives disclosures of vulnerabilities in its computer systems through a Company email address, "responsibledisclosure@capitalone.com."

235. On July 17, 2019, an individual contacted the Company at this email address, stating that "[t]here appears to be some leaked s3 data of yours in someone's github." "s3" refers to the "Simple Storage Service" offered by Amazon Web Services. Github refers to GitHub Inc., a subsidiary of Microsoft Corp., that provides webhosting for software development and allows users to manage and store files.

236. The email message provided the web address of a GitHub file containing data exfiltrated from Capital One or otherwise related to the Data Theft.

237. Prompted by the July 17, 2019 email, Capital One began an internal investigation, which led to the discovery of the Data Theft on July 19, 2019. It was at this point that Capital One finally fixed the particular misconfigured user rolls and web application firewall that had been exploited in the Data Theft.

238. Capital One referred the matter to the Federal Bureau of Investigation soon after discovering the Data Theft. Within days, the FBI had identified electronic communications linking erratic/Thompson to the Data Theft.

239. On July 29, 2019, Capital One announced the Data Theft. The United States arrested and charged erratic/Thompson, whose real name was Paige A. Thompson, that same day.

C. Capital One's Response

240. On August 4, 2019, Capital One issued a press release announcing the Data Theft and the scope of the data taken from its servers:

> **What happened**
>
> On July 19, 2019, we determined that an outside individual gained unauthorized access and obtained certain types of personal information about Capital One credit card customers and individuals who had applied for our credit card products.
>
> **What we've done**
>
> Capital One immediately fixed the issue and promptly began working with federal law enforcement. The person responsible was arrested. Based on our analysis to date, we believe it is unlikely that the information was used for fraud or disseminated by this individual. However, we will continue to investigate.

241. Despite Capital One's self-assured representation that it was unlikely that any of the stolen information was used or disseminated, the facts make clear that Capital One's statement was false and miskeading because it lacked any reasonable basis. For a period of months, the hacker had provided the public with a roadmap to pull massive amounts of user information off of Capital One's servers using an obvious and well-known misconfiguration exploit, and Capital One had no way to track who did so.

242. The August 4th release was also false and misleading when it said that "Capital One immediately fixed the issue." The real vulnerability arose from Capital One's and Amazon's

decision to pool sensitive user data in one place, allowing it to be openly shared across resources. Capital One may have sealed the *particular* hole in its firewall that the hacker used, but the vulnerability is an inherent part of its architecture. That is, the very same design that allowed resources to share pools of sensitive customer data using broad IAM roles was the reason the hacker was able to compromise so much data once inside the firewall. That problem had not been fixed and Capital One knew it.

243. Capital One further described the likely impact of the Data Theft, disclosing a theft of information of unprecedently massive scale:

> **What's the impact**
>
> Based on our analysis to date, this event affected approximately 100 million individuals in the United States and approximately 6 million in Canada.

244. The primary victims of the Data Theft were small business and consumer applicants for Capital One's credit cards for a massive time period—from 2005 to 2019:

> The largest category of information accessed was information on consumers and small businesses as of the time they applied dor one of our credit card products from 2005 through early 2019. This information included personal information Capital One routinely collects at the time it receives credit card applications, including names, addresses, zip codes / postal codes, phone numbers, email addresses, dates of birth, and self-reported income.

245. In addition, the stolen data included "credit scores, credit limits, balances, payment history, contact information," and even "transaction data from a total of 23 days during 2016, 2017, and 2018."

246. Mentioned in passing at the end of its statement—perhaps to bury the most damning admission—Capital One disclosed that social security numbers and bank accounts had been stolen:

> The individual also obtained the following data:

> - About 140,000 Social Security numbers of our credit card customers.
>
> - About 80,000 linked bank account numbers of our secured credit card customers.

247. The scope of the damage was even more staggering for Capital One's Canadian customers—"approximately, 1 million Social Insurance Numbers were compromised" in the incident.

248. The jaw dropping scope of data stolen was a consequence of the pooling problem. Capital One and Amazon had pooled large amounts of historical data into a data lake and deliberately failed to narrowly tailor which of its applications would have access to the entire set of data. If Capital One had not pooled all of its historical data to be accessed using broadly configured IAM roles, an attack of such scope would have been impossible. Capital One, however, did so because it wanted to mine customer data with ease across numerous hastily developed and deployed applications and businesses.

249. The Chairman and CEO of Capital One, Richard D. Fairbank issued an apology:

> While I am grateful that the perpetrators has been caught, I am deeply sorry for what has happened. . . . I sincerely apologize for the understandable worry this incident must be causing those affected and I am committed to making it right.

250. At no point in the statement did Capital One explain why its often-touted collaboration with Amazon, Cloud Custodian, had failed to mitigate any of the risks that manifested. He did not explain how Cloud Custodian would allow the hacker to repeatedly exploit the obvious configuration error without detection. He did not mention that the cause of the hack was an SSRF vulnerability with AWS that had existed and persisted for several years before the attack.

251. Notably, it was not until someone contacted Capital One months after the Data Theft that the Company even realized that it had been hacked. Cloud Custodian was represented "to automate detection and correction of policy violations." Cloud Custodian was unable to detect, let alone prevent, the Data Theft because the system was configured exactly the Capital One and Amazon intended that it be configured, with broadly defined IAM roles that emphasized data mining over data security.

252. The truth had been revealed: Either (i) there was nothing for Cloud Custodian to detect and fix because Capital One had configured the IAM roles accessible to resources within the fire wall broadly, or (ii) Cloud Custodian was simply not used to monitor IAM role policies notwithstanding Capital One and Amazon's statements to the contrary.

253. In both events, Cloud Custodian was exposed for what it was—a sham designed to convince credit card users that their information was safe on the AWS cloud. The inherent flaw in Capital One's aggregation of sensitive historical data and broad access policies was also readily apparent given the staggering scope of the attack.

D. Amazon's Response

254. Immediately after the Data Theft, Amazon took to the press to deflect blame for the Data Theft. An Amazon spokesman made a statement to Newsweek on July 30, 2019 that the incident was caused by a misconfigured web application:

> AWS was not compromised in any way and functioned as designed. The perpetrator gained access through a misconfiguration of the web applications and not the underlying cloud-based infrastructure. As Capital One explained clearly in its disclosure, this is type of vulnerability is not specific to the cloud.

255. The statement was false and misleading. The Data Theft, particularly its staggering breadth, was not caused by a mere misconfiguration of the web application as Amazon contended. Rather, it was the result of the cloud-based architecture Amazon and Capital One had implemented

in order to allow Capital One to rapidly mine user data. Massive amounts of data—more than a credit card application program could possibly need—had been centralized in the data lake. Then, the IAM roles were broadly defined to allow access to the entirety of that data lake to better facilitate data mining. If the hacker accessed resources inside the firewall due to a misconfiguration, that was part of the means and not the cause of the Data Theft, and certainly not the cause of the massive scope of the Data Theft.

256. Amazon did not disclose in its statement that the attacker was a former employee of Amazon that knew that data had been pooled for not only Capital One's web applications, but others as well. Thompson had developed a tool that scanned AWS looking for a way through web application firewalls. Thompson knew that once she made it through the firewalls, she could potentially assume IAM roles that would provide her access to large aggregations of data stored on shared S3 servers. Thus, any breach would reap vast amounts of sensitive data, not only increasing the damage of that particular breach but incentivizing others to attempt similar breaches.

257. Amazon was also silent about Cloud Custodian. There was no explanation as to why Cloud Custodian did not detect any potentially "misconfigured" web application. Indeed, AWS and Capital One together touted in their joint marketing that Cloud Custodian was developed to scan for misconfigurations and correct them. There was no explanation for why Cloud Custodian appeared to have done nothing at all. The truth was nonetheless now obvious—the broadly configured IAM role exploited by Thompson was a feature, not a bug.

V. THE FALLOUT

A. The Breadth of Data Compromised In the Theft Makes Clear That Capital One Was Pooling Sensitive Customer Data and Defining Broad IAM Roles That Allowed for Dynamic Access.

258. The scope of compromised data speaks volumes about the fatal flaw in Capital One's AWS-based systems. There is no legitimate reason for a credit card application system or

any public-facing system to have access to massive amounts of user data spanning more than a decade.

259. Capital One began its cloud migration in 2015, but the data compromised—as Capital One admitted in its press release—***spans from 2005 until 2019***. There is no legitimate reason to house so much past data on its cloud servers other than to facilitate data mining across applications. The large breadth of user data makes clear that Capital One did not restrict the amount of data that applications could access to the minimum amount necessary to complete a task. Instead, Capital One designed its systems, including public-facing web applications, to have broad access to a full data lake to train and use machine learning models.

260. Capital One's decision not to minimize the cross-section of data exposed to public-facing resources on its cloud servers was objectively unreasonable, reckless, far short of any industry or international standard, and a far cry from any notion of common sense.

261. In fact, it was a breach of the agreement Capital One made with its customers. For example, Capital One's Privacy Commitment stated, "We restrict access to your Social Security number except when required for authorized business purpose." The fact that 140,000 social security numbers and more than 1 million Canadian Social Insurance numbers sat on Capital One's AWS servers—ready for the taking by public-facing applications—is a clear violation of the agreement that access would be restricted to the extent required for business purposes.

262. There was simply no plausible business purpose for maintaining such a large store of information on a public-facing cloud resource. And the access to the resources were in no way tailored to the business use for which the data was maintained.

B. The Data Theft Made Clear that Cloud Custodian Was a Façade Designed to Falsely Signal Security to Customers.

263. The March 2019 Data Theft also made clear that Capital One's statements about Cloud Custodian were false. Specifically, Capital One—through statements broadly disseminated, including at conferences and on the internet—represented that it had developed a risk-mitigating infrastructure that was designed to automatedly scan for precisely the sort of vulnerability that allowed for the Data Theft.

264. For example, in its joint statement with Amazon, entitled "Capital One on AWS," Capital One boasted that the Cloud Custodian software "eliminat[ed] the need to manage hundreds or thousands of scripts and policies, and offer 'real-time' compliance and cost management at scale." That statement was false and misleading—either because Capital One failed to include the resources that were ultimately compromised within the purview of its Cloud Custodian system or because the systems themselves were designed for widespread and broad access to centralized stores of sensitive customer data.

265. Both Capital One and Amazon also published an architecture diagram in their joint statement that plainly represented that Cloud Custodian monitored EC2 and S3 policies.

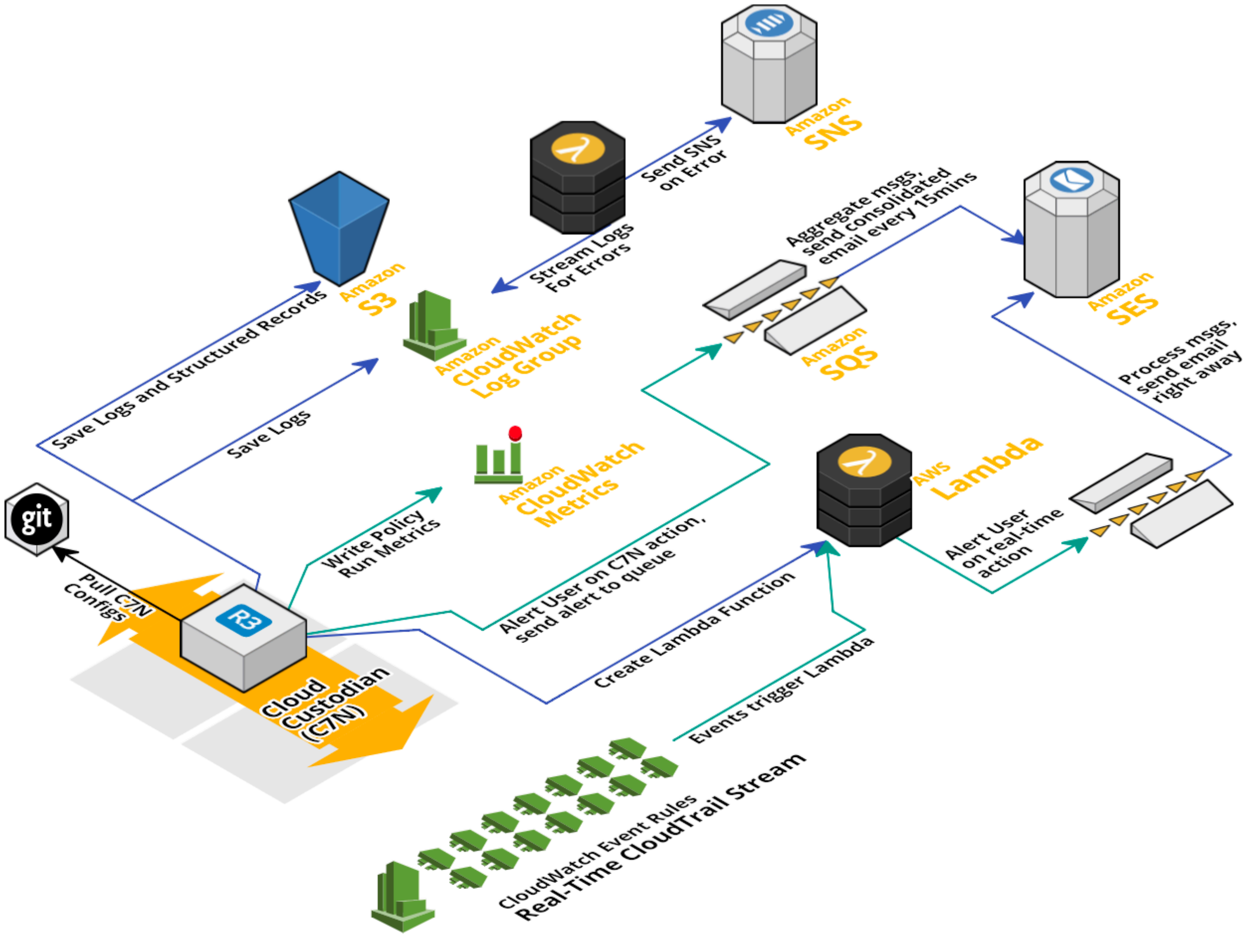
Amazon SNS
Send SNS on Error
Stream Logs For Errors
Aggregate msgs, send consolidated email every 15mins
Amazon SES
Amazon S3
Save Logs and Structured Records
Amazon CloudWatch Log Group
Save Logs
Amazon SQS
Process msgs, send email right away
Amazon CloudWatch Metrics
Write Policy Run Metrics
AWS Lambda
Alert User on real-time action
git
Pull C7N Configs
Cloud Custodian (C7N)
Alert User on C7N action, send alert to queue
Create Lambda Function
Events trigger Lambda
CloudWatch Event Rules
Real-Time CloudTrail Stream

266. The architecture diagram plainly implied that requests to S3 resources would be made through Cloud Custodian, not directly from resources such as EC2 servers or AWS Lambda-based applications. This architecture diagram was false, as the user data on Capital One's S3 servers, the data lake, was accessed directly from an EC2 server that assumed an improperly permissioned IAM role. To wit, when the Data Theft occurred, there was no sign of Cloud Custodian. It was certainly not the gatekeeper that Capital One represented it to be in its false and misleading architecture diagram.

267. Capital One's Thangavelu represented in a statement on July 14, 2017, that Cloud Custodian was designed "to put guard rails on" to ensure that policies and roles were correctly configured as to avoid allowing developers to "do silly things like leave their database publicly accessible on the internet." These "guard rails" would purportedly be automated precisely to avoid system vulnerability through "process and checklists."

268. Thangavelu's statements were false and misleading. To begin with, there is no evidence that Cloud Custodian had automatedly scanned the resources that were compromised in the summer of 2019. If it had and the resources were misconfigured, Cloud Custodian should have corrected the problem. And if Capital One had configured its IAM roles to allow access to its data lake, then there would be nothing for Cloud Custodian to correct—everything would be as it was designed to be.

269. Moreover, Thangavelu's statement about "guard rails" was also false and misleading. There would be no guard rails if Cloud Custodian was never brought to bear on a public-facing resource or if Capital One deliberately designed its system for open access to pooled resources. Capital One's developers were free to improperly configure a web application and the roles the application could assume, thereby allowing unconstrained access to user data on S3 servers, and if they deliberately did so, Cloud Custodian would have ***enforced*** their decision.

270. The truth was that Cloud Custodian was in fact a mere façade—a complex and massive piece of software with the stated purpose of providing security on the cloud, but with no substance in practice.

271. It nonetheless served its true purpose—to induce customers to provide Capital One with their most sensitive information so that Capital One could harvest it with its AI and machine

learning to obtain a competitive edge. Cloud Custodian was nothing more than a red herring designed to lull users and developers into a false sense of security.

272. In fact, both Capital One and Amazon knew that the real vulnerability was a combination of several factors, none of which was addressed by Cloud Custodian: (1) AWS was vulnerable to an SSRF attack which would allow the attacker to gain the IAM roles of the compromised system; (2) Capital One's data lake contained years of sensitive personal information; and (3) IAM roles were deliberately set to be broad to enable Capital One to data mine.

273. The risk posed by SSRF attacks, and Cloud Custodian's inability to address that vulnerability, was known to both Capital One and Amazon throughout the relevant period. That risk was disclosed at least as early as the 2014 Black Hat presentations. In 2018, Capital One's Chief Architect, speaking onstage at AWS's own conference, described the risk, and falsely claimed that Cloud Custodian would prevent such an attack. 2018 Black Hat presentations confirmed that the SSRF vulnerability still existed in AWS. And the 2019 Data Theft confirmed that the risk was more than theoretical. That risk is ongoing today.

C. Capital One's Representation (and Promise) that It Used Encryption Was False and Misleading.

274. Capital One stated in its privacy notice to customers that it would implement electronic safeguards, "such as passwords and encryption," to protect customer information.

275. That statement was false and misleading. Any encryption used by Capital One was pointless because encryption was based on user credentials. Obtaining usernames and passwords therefore ensured that data could be decrypted.

276. Capital One admitted as much. In its press release announcing the Data Theft, Capital One wrote:

> Was the data encrypted and/or tokenized?
>
> We encrypt our data as standard. Due to the particular circumstances of this incident, ***the unauthorized access also enabled the decrypting of data***.

277. Encryption in which a basic username and password is all that is required to decrypt is virtually useless because hacks by their very nature often involve compromise of those very same security credentials.

278. Capital One's failure to implement a process or procedure that would effectively prevent this simple decryption of data was objectively reckless and far short of any industry standard or notion of best practices, rendering its statements to customers false.

279. Moreover, because of the data pooling problem, poor encryption meant that access to a resource capable of assuming a broad IAM role meant access to massive amounts of sensitive customer information in an unencrypted form. The failure to properly encrypt the data meant that unprecedented amounts of information could be compromised.

D. The Flaws in Capital One's Architecture Still Exist and Capital One Should Be Required to Move Sensitive Customer Data Off of the AWS Cloud.

280. From Capital One's perspective, the data pooling and dynamic access problems that resulted in the Data Theft are features, not bugs. They are part of the design used by Capital One to allow it to rapidly mine user data that it collects. This means internal resources are given broad access to centralized data. Moreover, large aggregations of data—spanning several years—are needlessly and recklessly exposed to web applications that face the public. A foothold inside Capital One's firewall can mean unfettered access to an entire data lake of sensitive customer information. This problem continues.

281. Despite Capital One and Amazon's assurances to the public that the allegedly "misconfigured" firewall settings have been corrected, the inherent flaws in the architecture of

Capital One's cloud systems remain. Broadly defined IAM roles that allow web applications to access the data lake are part of the design of Capital One's applications. Nothing appears to have changed in that respect since the Data Theft.

282. Moreover, nothing prevents other attacks from obtaining the same broad swath of sensitive customer data if those attacks allow access to systems with a broadly defined IAM role. For example, a successful SSRF attack would leave internal resources just as vulnerable as they were in the Data Theft.

283. The SSRF vulnerability remains to this day. As Evan Johnson, the manager of a product security team at a large software company, explained in a post, "Preventing the Capital One Breach," SSRF is a well known and serious vulnerability—one that AWS has no mitigations for:

> Every indication is that the attacker exploited a type of vulnerability known as Server Side Request Forgery (SSRF) in order to perform the attack. SSRF has become the most serious vulnerability facing organizations that use public clouds. SSRF is not an unkown vulnerability, but it doesn't receive enough attention and was absent from [a top 10 list of vulnerabilities].
>
> SSRF is a bug hunters [sic] dream because it is an easy to perform attack and regularly yields critical findings, like this bug bounty report to Shopify. The problem is common and well-known, but hard to prevent ***and does not have any mitigations built into the AWS platform***.

(emphasis added).

284. SSRF remains a known vulnerability on AWS, and Capital One's use of customer data stored on AWS for data mining and machine learning means that a future attack will be able to access the same broad swath of data pulled in the Data Theft. Because Capital One does not have full control over the servers and allows public-facing web applications to access entire

historical databases of customer data, there is no way to fix the problems that led to the success and scope of the Data Theft short of a redesign of Capital One's entire cloud architecture.

285. Plaintiffs have been, and continue to be, harmed because Capital One maintains their personal data as part of large aggregations of data or data lakes. Defendants should be enjoined from maintaining centralized stores of highly sensitive customer data on cloud servers.

286. Plaintiffs have been, and continue to be, harmed because Capital One's cloud-based applications enjoy broad access to centralized and sensitive customer data. Defendants should be enjoined from maintaining web applications with permission to assume broadly defined IAM roles that provide data lake access.

287. Plaintiffs have been, and continue to be, harmed by Defendants' façade—Cloud Custodian. Defendants should be enjoined from falsely or misleadingly touting Cloud Custodian as a means of mitigating the risk of broad access to central data or overbroad policies for IAM roles.

288. An injunction preventing Defendants from maintaining Plaintiffs and class members' sensitive customer data on the AWS cloud would be just and equitable, and there is no other adequate remedy for the prospective risks posed by Defendants' flawed design of Capital One's cloud architecture.

CLASS ACTION ALLEGATIONS

289. The Classes' claims all derive directly from a course of conduct by Defendants. Defendants have engaged in uniform and standardized conduct toward the class. They did not differentiate, in degree of care or candor, in their actions or inactions, or in the content of their statements or omissions, among individual Class members. The objective facts on these subjects are all the same for all Class members. Within each Claim for Relief asserted by the class, the same legal standards govern. Additionally, many states, and for some claims all states, share the

same legal standards and elements of proof, facilitating the certification of multistate or nationwide class or classes for some or all claims. Accordingly, Plaintiffs bring this lawsuit as a class action on their own behalf and on behalf of all other persons similarly situated as members of the proposed class pursuant to Federal Rules of Civil Procedure 23(a) and (b)(3) and/or (b)(2) and/or (c)(4). This action satisfies the numerosity, commonality, typicality, adequacy, predominance, and superiority requirements of those provisions.

The Nationwide Credit Card Customer Class

290. Plaintiffs bring this action and seek to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of themselves and a Nationwide Credit Card Customer Class defined as follows:

> All individuals who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

291. Excluded from the Nationwide Credit Card Customer Class are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Connecticut Subclass

292. Plaintiffs Kristina Mentone and Cole Studebaker bring this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of themselves and a Connecticut Subclass defined as follows:

> All individuals who reside in Connecticut who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

293. Excluded from the Connecticut Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Florida Subclass

294. Plaintiff Susan Corley brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of herself and a Florida Subclass defined as follows:

> All individuals who reside in Florida who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

295. Excluded from the Florida Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Indiana Subclass

296. Plaintiff Janett Stout brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of herself and an Indiana Subclass defined as follows:

> All individuals who reside in Indiana who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

297. Excluded from the Indiana Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Massachusetts Subclass

298. Plaintiff Jonathan Wong brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of himself and a Massachusetts Subclass defined as follows:

> All individuals who reside in Massachusetts who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

299. Excluded from the Massachusetts Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The New Jersey Subclass

300. Plaintiff Kimberly Hernandez brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of herself and a New Jersey Subclass defined as follows:

> All individuals who reside in New Jersey who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

301. Excluded from the New Jersey Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The New York Subclass

302. Plaintiffs Mordechai Nemes brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of himself and a New York Subclass defined as follows:

> All individuals who reside in New York who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

303. Excluded from the New York Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Ohio Subclass

304. Plaintiffs Mark Miller and Ryan Olsen bring this action and seek to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of themselves and a Ohio Subclass defined as follows:

> All individuals who reside in Ohio who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

305. Excluded from the Ohio Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Pennsylvania Subclass

306. Plaintiff Debra Potzgo brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of herself and a Pennsylvania Subclass defined as follows:

> All individuals who reside in Pennsylvania who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

307. Excluded from the Pennsylvania Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries

or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The South Carolina Subclass

308. Plaintiffs Jacqueline Burke and Shawn Spears bring this action and seek to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of themselves and a South Carolina Subclass defined as follows:

> All individuals who reside in South Carolina who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

309. Excluded from the South Carolina Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Texas Subclass

310. Plaintiff Andrew Broderick brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of himself and a Texas Subclass defined as follows:

> All individuals who reside in Texas who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

311. Excluded from the Texas Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

The Wisconsin Subclass

312. Plaintiff Lynn Fields brings this action and seeks to certify and maintain it as a class action under Rules 23(a); (b)(2); and/or (b)(3); and/or (c)(4) of the Federal Rules of Civil Procedure on behalf of herself and a Wisconsin Subclass defined as follows:

> All individuals who reside in Wisconsin who applied for and/or received a credit card from Capital One from October 7, 2015, to the present.

313. Excluded from the Wisconsin Subclass are Defendants, their employees, officers, directors, legal representatives, heirs, successors, and wholly or partly owned subsidiaries or affiliates of Defendants; class counsel and their employees; and the judicial officers and their immediate family members and associated court staff assigned to this case.

Numerosity and Ascertainability

314. This action satisfies the requirements of FED. R. CIV. P. 23(a)(1). There are hundreds of thousands or more Capital One cardholders nationwide, and at least thousands in each of the States. Individual joinder of all Class members is impracticable.

315. The Class is ascertainable because its members can be readily identified using registration records, sales records, production records, and other information kept by Defendant or third parties in the usual course of business and within their control. Plaintiffs anticipate providing appropriate notice to the certified Class, in compliance with FED. R. CIV. P. 23(c)(1)(2)(A) and/or (B), to be approved by the Court after class certification, or pursuant to court order under FED. R. CIV. P. 23(d).

Predominance of Common Issues

316. This action satisfies the requirements of FED. R. CIV. P. 23(a)(2) and 23(b)(3) because questions of law and fact that have common answers that are the same for the Class

predominate over questions affecting only individual Class members. These include, without limitation, the following:

a. Whether Defendants engaged in an enterprise to defraud Plaintiffs and Class members into providing personal information to Capital One;

b. Whether Defendants falsely claimed that Cloud Custodian would detect and prevent misconfigured IAM roles and policy-based permissions;

c. Whether Defendants knew or should have known about AWS's SSRF vulnerability;

d. Whether Defendants knew or should have known that its web application firewall was vulnerable to attack, including by an SSRF.

e. Whether Defendants knowingly or recklessly made false or misleading statements and/or omissions about the security of Capital One's customer data on the AWS cloud;

f. Whether Defendants knowingly or recklessly made false or misleading statements about the use of customer data on the AWS cloud and the breadth of data that would be stored there;

g. Whether Defendants engaged in unfair, deceptive, unlawful, and/or fraudulent acts or practices in trade or commerce by failing to disclose that Capital One's cloud architecture was inherently flawed and that AWS was ill suited for highly sensitive customer data;

h. Whether Defendants' conduct, as alleged herein, was likely to mislead a reasonable consumer;

i. Whether Defendants' statements, concealments, and omissions regarding the security of customer data were material in that a reasonable consumer could consider them important in applying for, and using, a credit card;

j. Whether Defendants misrepresented that the Class's sensitive personal data was safe;

k. Whether Defendants violated each of the States' consumer protection statutes, and if so, what remedies are available under those statutes;

l. Whether Defendants have been unjustly enriched by their conduct;

m. Whether Defendants failed to comply with internal company policies and applicable laws, regulations, and industry standards relating to data security;

n. Whether Defendants continue to fail to comply with internal company policies and applicable laws, regulations, and industry standards relating to data security;

o. Whether Defendants knew or should have known that Capital One did not employ reasonable measures to keep Plaintiffs' and Class members' personal information secure and prevent the loss or misuse of that information;

p. Whether Defendants should have discovered the Data Theft prior to the external security researcher's report to the Company on July 17, 2019;

q. Whether Defendants made false or misleading statements and/or omissions in connection with the Data Theft.

r. What aggregate amounts of statutory penalties are enough to punish and deter Defendants and to vindicate statutory and public policy;

s. How penalties should be equitably distributed among Class members;

t. Whether Defendants conspired together to violate RICO;

u. Whether Defendants associated with any enterprise engaged in, or in the activities of which affect, interstate or foreign commerce, to conduct or participate, directly or indirectly, in the conduct of such enterpris's affairs through a pattern of racketeering;

v. Whether Plaintiffs and the Class members are entitled to actual damages or other forms of monetary relief; and

w. Whether Plaintiffs and the class are entitled to equitable relief, including, but not limited to, injunctive relief and restitution.

Typicality

317. This action satisfies the requirements of FED. R. CIV. P. 23(a)(3) because Plaintiffs' claims are typical of the claims of other Class members and arise from the same course of conduct by Defendants. The relief Plaintiffs seek is typical of the relief sought for the absent Class members.

Adequate Representation

318. Plaintiffs will fairly and adequately represent and protect the interests of the Class. Plaintiffs have retained counsel with substantial experience in prosecuting consumer class actions, including actions involving defective products and misleading / fraudulent services.

319. Plaintiffs and their counsel are committed to vigorously prosecuting this action on behalf of the Class. Neither Plaintiffs nor their counsel have interests adverse to those of the Class.

Superiority

320. This action satisfies the requirements of FED. R. CIV. P. 23(b)(2) because Defendants have acted and refused to act on grounds generally applicable to the Class, thereby making appropriate final injunctive and/or corresponding declaratory relief with respect to each Class as a whole.

321. This action satisfies the requirements of FED. R. CIV. P. 23(b)(3) because a class action is superior to other available methods for the fair and efficient adjudication of this

controversy. The common questions of law and fact regarding Defendants conduct and responsibility predominate over any questions affecting only individual Class members.

322. Because the damages suffered by each individual Class member may be relatively small, the expense and burden of individual litigation would make it very difficult or impossible for individual Class members to redress the wrongs done to each of them individually, such that most or all Class members would have no rational economic interest in individually controlling the prosecution of specific actions, and the burden imposed on the judicial system by individual litigation by even a small fraction of the Class would be enormous, making class adjudication the superior alternative under FED. R. CIV. P. 23(b)(3)(A).

323. The conduct of this action as a class action presents far fewer management difficulties, far better conserves judicial resources and the parties' resources, and far more effectively protects the rights of each Class member than would piecemeal litigation. Compared to the expense, burdens, inconsistencies, economic infeasibility, and inefficiencies of individualized litigation, the challenges of managing this action as a class action are substantially outweighed by the benefits to the legitimate interests of the parties, the court, and the public of class treatment in this Court, making class adjudication superior to other alternatives under FED. R. CIV. P. 23(b)(3)(D).

324. Plaintiffs are not aware of any obstacles likely to be encountered in the management of this action that would preclude its maintenance as a class action. Rule 23 provides the Court with authority and flexibility to maximize the efficiencies and benefits of the class mechanism and reduce management challenges. The Court may, on motion of Plaintiffs or on its own determination, certify nationwide, statewide, and/or multistate classes for claims sharing common legal questions; utilize the provisions of Rule 23(c)(4) to certify any particular claims, issues, or

common questions of fact or law for class-wide adjudication; certify and adjudicate bellwether class claims; and utilize Rule 23(c)(5) to divide any class into subclasses.

CLAIMS FOR RELIEF

I. NATIONWIDE CLASS CLAIMS

COUNT ONE:
Violation of 18 U.S.C. § 1962(c), the Racketeer Influenced and Corrupt Organization Act ("RICO") (All Plaintiffs on behalf of the Nationwide Credit Card Customer Class against All Defendants)

325. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

326. Plaintiffs bring this Count on behalf of the Nationwide Credit Card Customer Class.

327. Plaintiffs are natural persons, and as such are "persons within the meaning of 18 U.S.C. § 1961(3).

328. Defendants are "persons" within the meaning of 18 U.S.C. § 1961(3).

329. Defendants violated 18 U.S.C. 1962(b) by participating in or conducting the affairs of the Capital One-Amazon association-in-fact through a pattern of racketeering activity.

The Capital One-Amazon RICO Enterprise

330. The following persons, and others presently unknown, have been members of an constitute an "association-in-fact enterprise" within the meaning of RICO, and will be referred to collectively here as the "Capital One-Amazon RICO Enterprise":

331. Capital One, which participated in the design of, launched, purchased, marketed to consumers, and operated Capital One's "data lake" on AWS, knowing that it was fundamentally incapable of keeping data secure, and which actively concealed the scope and nature of this defect from and lied to the public.

332. Amazon participated in the design of, launched, marketed, and operated AWS data hosting services, including services involving the machine learning services and "data lakes" at issue here, and actively concealed from the public the scope and nature of this defect therein.

333. The Capital One-Amazon RICO Enterprise, which engages in and whose activities affect interstate and foreign commerce, is an association-in-fact of corporate entities within the meaning of 18 U.S.C. § 1961(4) and consists of "persons" associated together for a common purpose. The Capital One-Amazon RICO Enterprise has an ongoing organization with an ascertainable structure and functions as a continuing unit with separate roles and responsibilities.

334. While Capital One and Amazon participate in the conduct of the Capital One-Amazon RICO Enterprise, they have an existence separate and distinct from the Capital One-Amazon RICO Enterprise. Further, the Capital One-Amazon RICO Enterprise is separate and distinct from the pattern of racketeering in which Capital One and Amazon engage.

335. At all relevant times, Capital One has operated, controlled, or managed the Capital One-Amazon RICO Enterprise, through various actions. Capital One's participation in the Capital One-Amazon RICO Enterprise is necessary for the operation of its scheme to defraud because (among other reasons) Capital One assisted in the design of, launched, and operated, Capital One's "data lake" on AWS, knowing that it was fundamentally incapable of keeping data secure; concealed and lied about the risks of Capital One's "data lake" on AWS; and has profited from and is profiting from its concealment and lies.

336. At all relevant times, Amazon has operated, controlled, or managed the Capital One-Amazon RICO Enterprise, through various actions. Amazon's participation in the Capital One-Amazon RICO Enterprise is necessary for the operation of its scheme to defraud because (among other reasons) Amazon designed, manufactured, and sold Capital One's "data

lake" on AWS, knowing that it was fundamentally incapable of keeping data secure; concealed and lied about the risks of Capital One's "data lake" on AWS; and has profited and is profiting from its concealment and lies.

337. Capital One and Amazon, as members of the Capital One-Amazon RICO Enterprise, serve a common purpose to monetize the personal information of bank customers and to obtain other benefits described throughout this Complaint..

338. For example, Capital One monetizes the personal data of its own customers through data mining by means of AWS's public cloud capabilities. Amazon monetizes the personal data of Capital One's customers by charging Capital One for the use of their public cloud resources. Amazon additionally monetizes the personal data of bank customers by attracting other banks and charging those banks for the use of their public cloud services.

Pattern of Racketeering Activity

339. Capital One and Amazon conduct and participate in the conduct of the affairs of the Capital One-Amazon RICO Enterprise through a pattern of racketeering activity that has consisted of numerous and repeated violations of the federal mail and wire fraud statutes, which prohibit the use of any interstate or foreign mail or wire facility for the purpose of executing a scheme to defraud, in violation of 18 U.S.C. §§ 1341 and 1343.

340. For Capital One, the purpose of the scheme is detailed throughout this Complaint, including for example, at paragraphs 153-172, *supra*. For example, Capital One conducts and participates in the Capital One-Amazon RICO Enterprise for the purpose of, *inter alia*, concealing the scope and nature of the fundamental data security flaws in AWS, including the AWS servers on which its customers' and applicants' data was hosted, in order to induce more individuals and entities to apply for its credit cards, offer credit card terms at greater interest rates, and/or to reduce

costs, including those associated with credit card rewards necessary to induce credit card applications. By concealing the scope and nature of the fundamental data security flaws in AWS, including the AWS servers on which its customers' and applicants' data was hosted, Capital One also maintained and boosted consumer confidence in the Capital One brand and in the AWS brand, and avoided remediation costs and negative publicity, all of which furthered the scheme to defraud and helped Capital One to induce more individuals and entities to apply for its credit cards, offer credit card terms at greater interest rates, and/or to reduce costs, including those associated with credit card rewards necessary to induce credit card applications.

341. For Amazon, the purpose of the scheme is detailed throughout this Complaint, including for example, at paragraphs 110-17 and 171-203, *supra*. For example, Amazon conducts and participates in the Capital One-Amazon RICO Enterprise for the purpose of concealing the scope and fundamental data security flaws in AWS in order to sell more of its cloud computing services, sell those services at a higher price and/or for a higher profit, and to, *inter alia*, avoid incurring costs associated with designing and testing a method or methods to resolve or protect against the fundamental data security flaws in AWS. Additionally, Amazon helped develop and promote Cloud Custodian as a solution to the inherent risk posed by its architecture. By concealing the scope and nature of fundamental data security flaws in AWS, Amazon also maintains and boosts consumer confidence in the AWS brand and in the Capital One brand, and avoids remediation costs and negative publicity, all of which furthers the scheme to defraud and helps Amazon sell more of its cloud computing services than it otherwise would sell, and to sell those services at a much greater price and/or for greater profit.

342. As detailed in the general factual allegations, Capital One and Amazon were aware of the risk of data thefts posed by fundamental data security flaws in AWS, but they intentionally

subjected Plaintiffs and Class members to that risk and consciously disregarded that risk in order to maximize their profits.

343. To further the scheme to defraud, Capital One and Amazon repeatedly misrepresented and concealed the nature of the fundamental data security flaws in AWS.

344. To carry out or attempt to carry out the scheme to defraud, Capital One and Amazon have conducted or participated in the conduct of the affairs of the Capital One-Amazon RICO Enterprise through a pattern of racketeering activity that employees the use of the mail and wire facilities, in violation of 18 U.S.C. § 1431 (mail fraud) and § 1343 (wire fraud), including, for example:

a. Capital One and Amazon devised a scheme to defraud by use of the mail, telephone, television, and Internet, or caused to be transmitted by means of mail and wire communications traveling in interstate or foreign commerce, writing(s) and/or signal(s), including Capital One websites, Amazon websites, communications between Capital One and Amazon, statements to and agreements with credit card applicants and cardholders, as well as advertisements and other communications to Capital One customers, including Plaintiffs and Class members; and

b. Capital One and Amazon utilize and have utilized the interstate and international mail and wires for the purpose of obtaining money or property by means of the omissions, false pretenses, and misrepresentations described therein.

345. Capital One and Amazon's pattern of racketeering activity in violation of the mail and wire fraud statutes includes, but is not limited to, the conduct alleged throughout this

Complaint, including (for example), the conduct alleged in paragraphs 177-219 and 237-254, *supra.*

346. Capital One and Amazon's conduct in furtherance of their scheme was intentional. Plaintiffs and Class members were directly harmed as a result of Capital One and Amazon's intentional and wrongful conduct in that their personal and private data was put at risk of being compromised—and in many cases was compromised—as a result of Capital One's and Amazon's conduct as alleged herein.

347. As described throughout this Complaint, Capital One and Amazon engaged in a pattern of related and continuous predicate acts and are likely to continue to do so. The predicate acts constituted and constitute a variety of unlawful activities, each conducted with the common purpose of defrauding Plaintiffs and other Class members and obtaining significant monies and revenues from them. The predicate acts also had the same or similar results, participants, victims, and methods of commission. The predicate acts were related and not isolated events.

348. The predicate acts all had the purpose of generating significant revenue and profits for Capital One and Amazon at the expense of Plaintiffs and Class members. The predicate acts were committed or caused to be committed by Capital One and Amazon through their participation in the Capital One-Amazon RICO Enterprise and in furtherance of its fraudulent scheme, and were interrelated in that they involved the relationship and common purposes described throughout this Complaint, including (for example) in paragraphs 177-219 and 237-254, *supra.*

349. By reason of and as a result of the conduct of Capital One and Amazon, and the pattern of racketeering activity engaged in on behalf of the Capital One-Amazon RICO enterprise, Plaintiffs and Class members have been injured in their business and/or property—for example, through paying credit card interest and fees for Capital One credit cards they never would have

obtained (or even applied for) had they known the sensitive personal data they were required to submit as part of their credit card applications would placed in a mining-friendly "data lake" on an insecure public cloud server with widely-known, unremediated access vulnerabilities, which, unknown to Plaintiffs and the Class members at that time, put their data at risk of being easily compromised.

350. Capital One and Amazon's violations of 18 U.S.C. § 1962(c) have directly and proximately caused injuries and damages to Plaintiffs and Class members, and Plaintiffs and Class members are entitled to bring this action for three times their actual damages, as well as injunctive/equitable relief and costs and reasonable attorneys' fees pursuant to 18 U.S.C. §§ 1964(a) and 1964(c).

COUNT TWO: Violation of 18 U.S.C. § 1962(d), the Racketeer Influenced and Corrupt Organization Act ("RICO") (All Plaintiffs on behalf of the Nationwide Credit Card Customer Class against All Defendants)

351. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

352. Plaintiffs bring this Count on behalf of the Nationwide Credit Card Customer Class.

353. In addition to the General Factual Allegations re-alleged and re-incorporated here through the general Reallegation and Incorporation by Reference paragraph above, Plaintiffs reallege and incorporate by reference the allegations set forth in Count One.

354. At all relevant times, Capital One and Amazon have been and continue to be associated with the Capital One-Amazon RICO Enterprise, and have agreed and conspired to violate 18 U.S.C. § 1962(c), that is, agreed to conduct and participate, directly and indirectly, in the conduct of the affairs of the Capital One-Amazon RICO Enterprise, through a pattern of racketeering activity, in violation of 18 U.S.C. § 1962(d).

355. Capital One and Amazon knew that their predicate acts of wire fraud and mail fraud were part of a pattern of racketeering activity and agreed to the commission of those acts to further their scheme to defraud Plaintiffs and Class members.

356. As a direct and proximate result of Capital One and Amazon's conspiracy, and the multiple overt acts taken by Capital One and Amazon in furtherance of that conspiracy, Plaintiffs and Class members have been injured in their business and/or property.

COUNT THREE:
Fraud by Concealment
(All Plaintiffs on behalf of the Nationwide Credit Card Customer Class against All Defendants)

357. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

358. Plaintiffs bring this claim on behalf of the Nationwide Credit Card Customer Class under the common law of fraud by concealment, as there are no true conflicts (case dispositive differences) among various states' laws of fraud by concealment. In the alternative, Plaintiffs bring this claim under the laws of the states where Plaintiffs and Class members reside and/or applied for Capital One credit cards.

359. As alleged above, Defendants concealed and suppressed material facts regarding the scope and nature of the fundamental data security flaws in AWS, including the AWS servers on which Capital One's customers' and applicants' data was hosted. These fundamental flaws include, but are not limited to, the flaws inherent in pooling data into a "data lake" in order to facilitate machine learning.

360. Defendants took steps to hide the fundamental data security flaws in AWS, including the AWS servers on which Capital One's customers' and applicants' data was hosted.

361. Defendants had a duty to disclose the fundamental data security flaws in AWS, including the AWS servers on which Capital One's customers' and applicants' data was hosted because they:

a. Had exclusive and/or far superior knowledge and access to the facts than Plaintiffs and Class Members, and knew that the facts were not known to or reasonably discoverable by Plaintiffs and Class Members;

b. Intentionally concealed the foregoing from Plaintiffs; and

c. Made incomplete and misleading representations about the fundamental data security flaws in AWS, including about the AWS servers on which Capital One's customers' and applicants' data was hosted, while purposefully withholding material facts from Plaintiffs that contradicted these representations.

362. These omitted and concealed facts were material because they would be relied on by a reasonable person applying for or continuing to use a Capital One credit card. Whether a service provider has affected measures reasonably designed to protect consumers' sensitive personal information is a material concern to a consumer. Plaintiffs and Class Members trusted Capital One to provide this baseline level of data security.

363. Defendants concealed and suppressed these material facts to falsely assure consumers that consumer data entrusted to Defendants was safe from hackers and other malicious actors.

364. Defendants actively concealed and/or suppressed these material facts, in whole or in part, in order to protect and increase their profits. Defendants concealed these facts at the expense of the security of Plaintiffs and Class Members' sensitive personal information.

365. Plaintiffs and Class Members were not aware of these omitted material facts and would not have applied for or continued to use Capital One credit cards as they did if they had known of the concealed and/or suppressed facts.

366. Statements disseminated by Defendants in marketing materials and in other public statements referring to or otherwise related to Capital One's data storage infrastructure and cybersecurity practices would have been seen by potential consumers throughout the country—including by Plaintiffs and Class Members. Plaintiffs and Class Members relied on statements by Defendants regarding the security of the sensitive Personal Information that they entrusted into Defendants' care. Had they been aware of the fundamental data security flaws in AWS, including the risks inherent in the AWS servers on which Capital One's customers' and applicants' data was hosted, Plaintiffs and the Class would not have applied for or continued to use credit cards from Capital One.

367. Because of the concealment and/or suppression of the facts, Plaintiffs and the Class sustained damage because they paid for services that were, in fact, worth significantly less than what they paid.

368. Accordingly, Defendants are liable to the Class for their damages in an amount to be proven at trial.

369. Defendants' acts were done maliciously, oppressively, deliberately, with intent to defraud, and in reckless disregard of Plaintiffs' and the Class's rights and well-being, and with the aim of enriching Defendants. Defendants' conduct therefore warrants an assessment of punitive damages in an amount sufficient to deter such conduct in the future, which amount is to be determined according to proof.

COUNT FOUR:
Fraud by Misrepresentation
(All Plaintiffs on behalf of the Nationwide Credit Card Customer Class against All Defendants)

370. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

371. Plaintiffs bring this claim on behalf of the Nationwide Credit Card Customer Class under the common law of fraudulent misrepresentation as there are no true conflicts (case dispositive differences) among various states' laws of fraudulent misrepresentation. In the alternative, Plaintiffs bring this claim under the laws of the states where Plaintiffs and Class Members reside and/or applied for Capital One credit cards.

372. Defendants marketed and sold their services to consumers—including Plaintiffs and Class members—in a manner that was intentionally designed to deceive them into believing that Capital One's data storage infrastructure and cybersecurity measures on AWS were sufficient to protect consumers' sensitive personal information.

373. Statements disseminated by Defendants in marketing materials and in other public statements about Capital One's data storage infrastructure and cybersecurity measures on AWS would have been seen by potential consumers throughout the country—including by Plaintiffs and Class Members. Plaintiffs and Class Members relied on statements by Defendants regarding the Capital One's data storage infrastructure and cybersecurity measures on AWS when deciding to apply for and/or use Capital One credit cards.

374. Defendants' misrepresentations are material because reasonable consumers attach importance to statements regarding the security of their sensitive personal information and are induced to make purchasing decisions based on such representations.

375. At all relevant times when such misrepresentations were made, Defendants knew that the misrepresentations were false and misleading.

376. Plaintiffs and Class members reasonably and justifiably relied on Defendants' fraudulent misrepresentations when deciding to apply for Capital One credit cards, and had the correct facts been known, would not have so applied.

377. Therefore, as a direct and proximate result of Defendants' fraudulent misrepresentations, Plaintiffs and Class Members have suffered economic losses and other general and specific damages, including inflated prices paid for Capital One's services, in an amount to be proven at trial.

COUNT FIVE:
Unjust Enrichment
(All Plaintiffs on behalf of the Nationwide Credit Card Customer Class against Capital One and COBNA)

378. Plaintiffs bring this claim against Capital One on behalf of the Nationwide Credit Card Customer Class under the common law of unjust enrichment as there are no true conflicts (case dispositive differences) among various states' laws of unjust enrichment. In the alternative, Plaintiffs bring this claim under the laws of the states where Plaintiffs and Class Members reside and/or applied for Capital One credit cards.

379. Capital One has received and retained a benefit from Plaintiffs and the Class and inequity has resulted.

380. Capital One benefitted through its unjust conduct by selling its services, at a profit, for more than those services were worth to Plaintiffs who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

381. It is inequitable for Capital One to retain these benefits.

382. Plaintiffs and Class Members do not have an adequate remedy at law.

383. As a result of Capital One's conduct, the amount of its unjust enrichment should be disgorged, in an amount to be proven at trial.

II. STATE SUBCLASS CLAIMS

COUNT SIX:
Violation of Connecticut Unfair Trade Practices Act, C.G.S. § 42-110, (Connecticut Subclass against All Defendants)

384. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

385. Defendants are "persons" as defined by C.G.S § 42-110a(3).

386. Defendants are engaged in "trade" or "commerce" as those terms are defined by C.G.S. § 42-110a(4).

387. Upon commencement of this action, Plaintiff will send notice to Connecticut's Attorney General and to the state's Commissioner of Consumer Protection pursuant to C.G.S. § 42-110g(c).

388. In violation of C.G.S. § 42-110g(a), Defendants engaged unfair and deceptive acts and practices in the conduct of trade or commerce, as defined in C.G.S. § 42-110b, including:

a. representing that services have sponsorship, approval, characteristics, ingredients, uses, benefits, or qualities that they do not have;

b. representing that services are of a particular standard, quality, or grade, or that goods are of a particular style or model, if they are of another; and

c. engaging in any other unconscionable, false, misleading, or deceptive act or practice in the conduct of trade or commerce.

389. Defendants' unlawful, deceptive, and unconscionable acts include:

a. failing to implement and maintain reasonable security and privacy measures to protect Connecticut Subclass members' Personal Information, which was a direct and proximate cause of the Data Theft;

b. failing to identify foreseeable security and privacy risks, remediate identified security and privacy risks, and adequately improve security and privacy measures following previous cybersecurity incidents, which was a direct and proximate cause of the Data Theft;

c. failing to comply with legal duties pertaining to the security and privacy of Connecticut Subclass members' Personal Information, including legal duties imposed by the Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, and, in the case of Capital One, duties imposed by Section 501(b) of the Gramm-Leach-Bliley Act, 15 U.S.C. § 6801(b);

d. misrepresenting that they would protect the confidentiality of Connecticut Subclass members' Personal Information, including by implementing and maintaining reasonable security measures;

e. misrepresenting that they would comply with legal duties pertaining to the security and privacy of Connecticut Subclass members' Personal Information;

f. concealing the material fact that they did not reasonably or adequately secure Connecticut Subclass members' Personal Information; and

g. concealing the material fact that they did not comply with common law and statutory duties pertaining to the security and privacy of Connecticut Subclass members' Personal Information.

390. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, Connecticut Subclass members have been harmed in that they purchased services from Defendants for more than those services were worth to the Connecuticut Subclass members, who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital

One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

391. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, have suffered and will continue to suffer injury which they could not reasonably avoid, including ascertainable losses of money or property, and other damages as alleged above.

392. Defendants' violations of Connecticut law were done with reckless indifference to the rights of Connecticut Subclass members or were an intentional or wanton violation of those rights.

393. Connecticut Subclass members seek, pursuant to C.G.S § 42-110g, all monetary and non-monetary relief allowed by law, including actual damages, punitive damages, disgorgement, injunctive relief, and attorneys' fees and costs.

COUNT SEVEN:
Violation of New Jersey Consumer Fraud Act, N.J. Stat. §§ 56:8-1, *et seq.*, (New Jersey Subclass against All Defendants)

394. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

395. The New Jersey Subclass is made up of "persons" as defined by N.J. Stat. Ann. § 56:8-1(d).

396. Defendants sell "merchandise" within the meaning of N.J. STAT. ANN. § 56:8-1(c).

397. In violation of N.J. STAT. ANN. § 56:8-2, Defendants engaged in unconscionable commercial practices, deception, fraud, false pretense, false promise, misrepresentation, or knowing concealment suppression, or omission of material fact. These acts include:

a. failing to implement and maintain reasonable security and privacy measures to protect New Jersey Subclass members' Personal Information, which was a direct and proximate cause of the Data Theft;

b. failing to identify foreseeable security and privacy risks, remediate identified security and privacy risks, and adequately improve security and privacy measures following previous cybersecurity incidents, which was a direct and proximate cause of the Data Theft;

c. failing to comply with legal duties pertaining to the security and privacy of New Jersey Subclass Members' personal information, including legal duties imposed by the Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, and, in the case of Capital One, duties imposed by Section 501(b) of the Gramm-Leach-Bliley Act, 15 U.S.C. § 6801(b);

d. misrepresenting that they would protect the confidentiality of New Jersey Subclass Members' personal information, including by implementing and maintaining reasonable security measures;

e. misrepresenting that they would comply with legal duties pertaining to the security and privacy of New Jersey Subclass Members' personal information;

f. concealing the material fact that they did not reasonably or adequately secure New Jersey Subclass Members' personal information; and

g. concealing the material fact that they did not comply with common law and statutory duties pertaining to the security and privacy of New Jersey Subclass Members' personal information.

398. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, New Jersey Subclass members have been harmed in that they purchased services from Defendants for more than those services were worth to the New Jersey Subclass members, who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital

One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

399. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, New Jersey Subclass Members have suffered and will continue to suffer injury which they could not reasonably avoid, including ascertainable losses of money or property, and other damages as alleged above.

400. New Jersey Subclass Members seek all monetary and non-monetary relief allowed by law, including, pursuant to N.J. STAT. ANN. §§ 56:8-2.12 and 56:8-19, ascertainable losses of moneys or property, treble damages, restitution, injunctive relief, attorneys' fees, filing fees, and costs.

COUNT EIGHT:
Violation of New York General Business Law §§ 349, *et seq.*,
(New York Subclass against All Defendants)

401. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

402. In violation of N.Y. GEN. BUS. LAW § 349, Defendants engaged in deceptive acts or practices in the conduct of Capital One's business, trade, and commerce or furnishing of services. These acts include:

a. failing to implement and maintain reasonable security and privacy measures to protect New York Subclass Members' personal information, which was a direct and proximate cause of the Data Theft;

b. failing to identify foreseeable security and privacy risks, remediate identified security and privacy risks, and adequately improve security and privacy measures following previous cybersecurity incidents, which was a direct and proximate cause of the Data Theft;

c. failing to comply with legal duties pertaining to the security and privacy of New York Subclass Members' personal information, including legal duties imposed by Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, and, in the case of Capital One, duties imposed by Section 501(b) of the Gramm-Leach-Bliley Act, 15 U.S.C. § 6801(b) and duties imposed by New York General Business Law § 399-ddd (relating to the "Confidentiality of social security account number(s)");

d. misrepresenting that they would protect the confidentiality of New York Subclass Members' personal information, including by implementing and maintaining reasonable security measures;

e. misrepresenting that they would comply with legal duties pertaining to the security and privacy of New York Subclass Members' personal information;

f. concealing the material fact that they did not reasonably or adequately secure New York Subclass Members' personal information; and

g. concealing the material fact that they did not comply with common law and statutory duties pertaining to the security and privacy of New York Subclass Members' personal information.

403. Defendants acted willfully and knowingly in committing the deceptive acts and practices alleged herein.

404. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, New York Subclass members have been harmed in that they purchased services from Defendants for more than those services were worth to the New York Subclass members, who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital

One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

405. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, New York Subclass Members have suffered and will continue to suffer injury which they could not reasonably avoid, including ascertainable losses of money or property, and other damages as alleged above.

406. New York Subclass Members seek, pursuant to N.Y. Gen. Bus. Law § 349(h), all monetary and non-monetary relief allowed by law, including actual damages or statutory damages in the amount of $50 per violation (whichever is greater), treble damages for willful or knowing violations, injunctive relief, and attorneys' fees.

COUNT NINE:
Violation of Pennsylvania Unfair Trade Practices and Consumer Protection Law, 73 Pa. Cons. Stat. §§ 201-1 *et seq.*
(Pennsylvania Subclass against All Defendants)

407. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

408. The Pennsylvania Subclass Members purchased services in "trade" and "commerce," primarily for personal, family, and/or household purposes, within the meaning of 73 Pa. Cons. Stat. § 201-9.2(a).

409. Defendants are "person[s]," within the meaning of 73 PA. CONS. STAT. § 201-2(2).

410. In violation of 73 PA. CONS. STAT § 201-3, Defendants engaged in unfair or deceptive acts or practices in the conduct of its trade and commerce, including the following:

a. representing that its goods and services have characteristics, uses, benefits, and qualities that they do not have (73 PA. STAT. § 201-2(4)(v));

b. representing that its goods and services are of a particular standard or quality if they are another (73 PA. STAT. § 201-2(4)(vii)); and

c. advertising its goods and services with the intent not to sell them as advertised (73 Pa. Stat. § 201-2(4)(ix)).

411. Defendants' unfair or deceptive acts or practices include:

a. failing to implement and maintain reasonable security and privacy measures to protect Pennsylvania Subclass Members' personal information, which was a direct and proximate cause of the Data Theft;

b. failing to identify foreseeable security and privacy risks, remediate identified security and privacy risks, and adequately improve security and privacy measures following previous cybersecurity incidents, which was a direct and proximate cause of the Data Theft;

c. failing to comply with legal duties pertaining to the security and privacy of Pennsylvania Subclass Members' personal information, including legal duties imposed by the Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, and, in the case of Capital One, duties imposed by Section 501(b) of the Gramm-Leach-Bliley Act, 15 U.S.C. § 6801(b);

d. concealing the material fact that they did not reasonably or adequately secure Pennsylvania Subclass Members' personal information; and

e. concealing the material fact that they did not comply with common law and statutory duties pertaining to the security and privacy of Pennsylvania Subclass Members' personal information.

412. Defendants intended to mislead the Pennsylvania Subclass and induce them to rely on their misleading omissions. Defendants acted intentionally, knowingly, and maliciously to

violate Pennsylvania Unfair Trade Practices and Consumer Protection Law, and recklessly disregarded Pennsylvania Subclass members' rights.

413. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, Pennsylvania Subclass members have been harmed in that they purchased services from Defendants for more than those services were worth to the Pennsylvania Subclass members, who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

414. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, Pennsylvania Subclass Members have suffered and will continue to suffer injury which they could not reasonably avoid, including ascertainable losses of money or property, and other damages as alleged above.

415. Pennsylvania Subclass Members seek, pursuant to 73 PA. CONS. STAT. §§ 201-9.2, all monetary and non-monetary relief allowed by law, including actual damages or statutory damages of $100 (whichever is greater), treble damages, attorneys' fees, filing fees, costs, injunctive relief, and such additional relief as the court deems necessary or proper.

COUNT TEN:
Violation of Wisconsin Deceptive Trade Practices Act, Wis. Stat. § 100.18
(Wisconsin Subclass against All Defendants)

416. Plaintiffs repeat and reallege all preceding paragraphs, as if fully set forth herein.

417. The Wisconsin Subclass is made up of members of "the public" within the meaning of WISC. STAT. § 100.18(1).

418. Defendants are "person[s], firms[s], corporation[s] or association[s]," within the meaning of WISC. STAT. § 100.18(1).

419. With intent to sell, distribute, or increase consumption of their services or anything else they offered to members of the public for sale, use, or distribution, Defendants made, published, circulated, placed before the public in Wisconsin —or caused to be made, published, circulated, or placed before the public in Wisconsin—advertisements, announcements, statements, and representations to the public which contained assertions, representations, or statements of fact which are untrue, deceptive, and/or misleading, in violation of Wis. Stat. § 100.18(1).

420. Defendants also engaged in the above-mentioned conduct as part of a plan or scheme, the purpose or effect of which was to sell, purchase, or use services not as advertised, in violation of WIS. STAT. § 100.18(9).

421. Defendants' deceptive acts, practices, plans, and schemes include:

a. failing to implement and maintain reasonable security and privacy measures to protect Wisconsin Subclass Members' personal information;

b. failing to identify foreseeable security and privacy risks, remediate identified security and privacy risks, and adequately improve security and privacy measures following previous cybersecurity incidents;

c. failing to comply with legal duties pertaining to the security and privacy of Wisconsin Subclass Members' personal information, including legal duties imposed by the Section 5 of the Federal Trade Commission Act, 15 U.S.C. § 45, and, in the case of Capital One, duties imposed by Section 501(b) of the Gramm-Leach-Bliley Act, 15 U.S.C. § 6801(b);

d. misrepresenting that they would protect the confidentiality of Wisconsin Subclass Members' personal information, including by implementing and maintaining reasonable security measures;

e. misrepresenting that they would comply with legal duties pertaining to the security and privacy of Wisconsin Subclass Members' personal information;

f. concealing the material fact that they did not reasonably or adequately secure Wisconsin Subclass Members' personal information; and

g. concealing the material fact that they did not comply with common law and statutory duties pertaining to the security and privacy of Wisconsin Subclass Members' personal information.

422. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, Wisconsin Subclass members have been harmed in that they purchased services from Defendants for more than those services were worth to the Wisconsin Subclass members, who would not applied for or used Capital One credit cards at all, or at the terms offered by Capital One, had they been aware that their sensitive personal information would reside in Capital One's "data lake" on AWS, which was fundamentally incapable of keeping data secure.

423. As a direct and proximate result of Defendants' unfair and deceptive acts and practices, Wisconsin Subclass members have suffered and will continue to suffer injury which they could not reasonably avoid, including ascertainable losses of money or property, and other damages as alleged above.

424. Wisconsin Subclass Members seek, pursuant to WIS. STAT. § 100.18(11)(b), all monetary and non-monetary relief allowed by law, including pecuniary losses, attorneys' fees, filing fees, and costs.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs pray that this Court:

A. Enter an order certifying this case as a class action pursuant to Federal Rule of Civil Procedure 23;

B. Enter a judgment declaring that Defendants have committed the violations of law alleged in this case;

C. Award actual, compensatory, statutory, consequential damages;

D. Award punitive and treble damages;

E. Award equitable monetary relief, including restitution and disgorgement of all ill-gotten gains, and the imposition of a constructive trust upon, or otherwise restricting the proceeds of Defendants' ill-gotten gains, to ensure an effective remedy;

F. Award Plaintiffs the costs of this action, including reasonable attorneys' fees and expenses and expert fees and costs;

G. Enjoin Defendants from continuing to falsely market and advertise, conceal material information from the public, and commit unlawful and unfair business acts and practices; and order Defendants to engage in a corrective notice campaign;

H. Enjoin Defendants from maintaining web applications with permission to assume broadly defined IAM roles that provide access to the data lake, as alleged herein;

I. Enjoin Defendants from maintaining Class Members' sensitive personal information on the AWS cloud;

J. Award declaratory relief;

K. Award pre-judgment and post-judgment interest at the highest rate allowed by law; and

L. Grant such further relief as this Court may deem just and proper.

JURY DEMAND

Plaintiffs demand a trial by jury on all claims so triable as a matter of right.

Dated: November 15, 2019

Respectfully Submitted,

/s/ Andrew M. Williamson

Andrew M. Williamson (VSB No. 83366)
Andrew J. Pecoraro (VSB No. 92455)
PIERCE BAINBRIDGE BECK
PRICE & HECHT LLP
601 Pennsylvania Avenue, NW
South Tower, Suite 700
Washington, DC 20004
202-318-9001 – Telephone
202-463-2103 – Facsimile
awilliamson@piercebainbridge.com
apecoraro@piercebainbridge.com

Yavar Bathaee (*pro hac vice* forthcoming)
Michael M. Pomerantz
(*pro hac vice* forthcoming)
David Hecht (*pro hac vice* forthcoming)
Maxim Price (*pro hac vice* forthcoming)
Michael K. Eggenberger
(*pro hac vice* forthcoming)
PIERCE BAINBRIDGE BECK
PRICE & HECHT LLP
277 Park Avenue, 45th Floor
New York, New York 10172
(212) 484-9866
ybathaee@piercebainbridge.com
dhecht@piercebainbridge.com
mprice@piercebainbridge.com
mpomerantz@piercebainbridge.com
meggenberger@piercebainbridge.com

Brian J. Dunne (*pro hac vice* forthcoming)
PIERCE BAINBRIDGE BECK
PRICE & HECHT LLP
355 South Grand Avenue, 44th Floor
Los Angeles, California 90071
(213) 262-9333
bdunne@piercebainbridge.com

Counsel for Andrew Broderick, Jacqueline Burke, Susan Corley, Lynn Fields, Kimberly Hernandez, Kristina Mentone, Mark Miller, Mordechai Nemes, Ryan Olsen, Debra Potzgo, Shawn Spears, Janett Stout, Cole Studebaker, and Jonathan Wong, each individually and on behalf of all others similarly situated.