Nerd For Tech · Follow publication

Capital One Data Breach — 2019

8 min read · Dec 4, 2022

Tanner Jones Follow

Listen Share More

Introduction

Capital One is one of the largest credit card issuers, and their data breach exposed records belonging to over 100 million individuals. The threat actor infiltrated the network through a misconfigured open-source web application firewall (WAF). The multi-stage attack allowed the perpetrator to exfiltrate data stored in an AWS S3 storge bucket. What data was stolen? A statement by Capital One indicates the largest amount of data compromised was personal information including, names, addresses, zip codes/postal codes, phone numbers, email addresses, dates of birth, and self-reported income. (Wolf, 2022) Along with personal information, 140,000 social security numbers and 80,000 linked bank account numbers were compromised as well as approximately 1 million Canadian social insurance numbers. (One, 2022) Customer status data was exposed including, credit scores, credit limits, balances, payment history, contact information, and fragmented transaction data. (One, 2022) The paragraphs below will provide a brief technical outline of how the data breach occurred, Capital One's response to the security incident, their failed security controls, and how a cloud service provider (CSP) such as Amazon Web Services (AWS) could improve their services.

Technical Outline

Capital One is the fifth largest bank in the United States. They employ nearly 50 thousand people and have a revenue of 32 billion. (Trends, 2022) They were one of the first banks in the world to invest in migrating their on-premises datacenters to a cloud environment. (Khan, 2022) How does such a large-scale data breach occur to such a large institution that is heavily regulated? Capital One has a significant security budget, as they are a large financial institution. The attack wasn't performed by a nation state threat actor but rather by a single individual. The attacker's name is Paige Thompson, and she is a former AWS employee. She entered the network by creating a scanning tool to scan cloud infrastructure and look for misconfigured firewalls. Capital One wasn't the only company that she had compromised. She also leaked data from over 30 organizations including private and government entities. (Khan, 2022) She was later arrested after an anonymous email sent to Capital One, which aided the FBI in the discovery of the GitHub repository and social media posts. (Waldman, 2022) Below is an outline of the attack and response:

- March 22 and 23, 2019: The attacker gained access.
- July 17, 2019: Capital One discovered the incident.
- July 19, 2019: They determined that an external user gained unauthorized access.

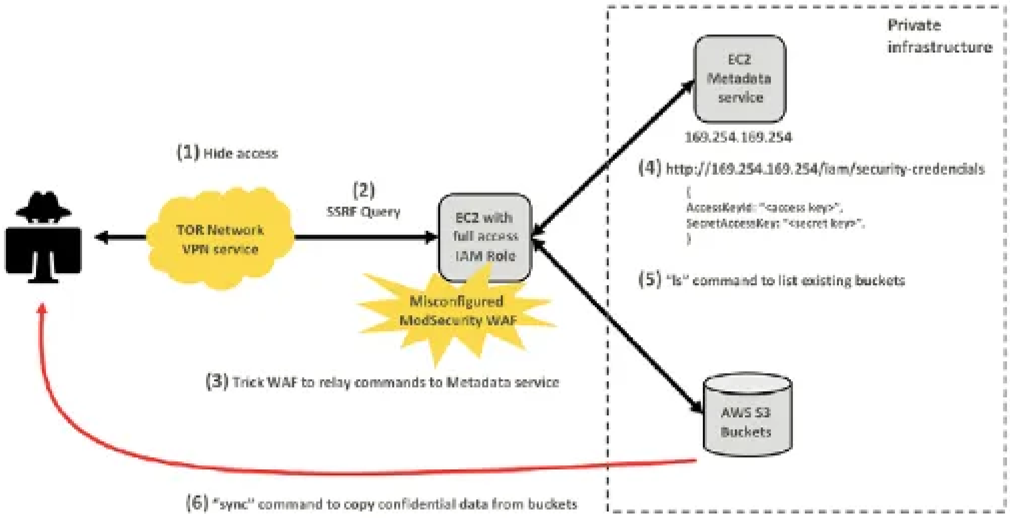
Private
infrastructure
EC2
Metadata
service
169.254.169.254
(1) Hide access
(2)
SSRF Query
TOR Network
VPN service
EC2 with
full access
IAM Role
(4) http://169.254.169.254/iam/security-credencials
{
AccessKeyId: "<access key>",
SecretAccessKey: "<secret key>",
}
(5) "ls" command to list existing buckets
Misconfigured
ModSecurity WAF
(3) Trick WAF to relay commands to Metadata service
AWS S3
Buckets
(6) "sync" command to copy confidential data from buckets

Figure 1: Attack on AWS Infrastructure

Figure 1 outlines the attack performed to gain access and exfiltrate the data. (1) The threat actor accessed the network via anonymizing services such as TOR and VPN. (2) They then performed a server-side request forgery or SSRF attack that tricked the server into executing commands as a remote user. (3) Next, a scanning tool identified misconfigured WAF and relayed commands to gain access to the AWS platform and receive temporary credentials to the environment. (4) The SSRF and WAF misconfiguration were combined, and the attacker obtained valid credentials, which allowed her to run commands through the AWS CLI. (5) She identified data stored in S3 buckets. (6) Lastly, she copied 30 GB of customer data across 700 different S3 buckets.

Responses

Once Capital One was aware of the data breach, they were quick to fix the issues and alert customers of the security incident. This partly is attributed to the finance industry and the regulations and requirements that they must follow. Ultimately, they were quick and transparent with the security incident. They provided customers with details by notifying the affected individuals and provided two years of free credit monitoring and identity protection at no charge. Due to the nature and size of the data breach, the criminal responsible was quickly found and arrested. This allowed authorities to identify 30 other security incidents for which the attacker was responsible. Capital One was transparent and has provided an in-depth analysis of the incident which is very helpful. That information allows other companies and individuals to learn from their mistakes and make necessary improvements. Their website is clear and concise on the data breach by providing ongoing updates and providing conclusive information in the 2019 Cyber Incident Settlement. (One, 2022)

Response Improvement

Overall, Capital One responded to the data breach rapidly and responsibly once they were aware of the incident, although there are a few areas that could be improved. The number of individuals that had their social security numbers comprised was more than what they had originally stated. After further investigation on January 27, 2021, they discovered that 4,700 social security numbers were compromised. (One, 2022) It took a substantial amount of time to discover and notify the impacted customers. They could have had the proper tools in place to detect the incident sooner and mitigate the damage. They could have been quicker to reevaluate their cloud architecture and implement the proper controls. Lastly, the company could have looked at security with a different lens and not confuse compliance with security. The financial sector is highly regulated and requires continuous audits and evaluations, but these flaws weren't discovered until after over 100 million individuals lost sensitive information.

Biggest Mistake

The biggest mistake by Capital One was the lack of proper security controls to monitor, detect, and alert intrusions. They needed proper mechanisms to provide more visibility into what was happening on the network. Unfortunately, the data breach could have been adverted if proper controls were implemented. Access to cloud resources and services is done using Access Management (IAM) controls. The role that was used to carry out the attack was overprovisioned and shouldn't have had permissions to have access to encryption keys, S3 buckets and WAF configurations. The threat actor was able to enter the network undetected, perform attacks, scan external security controls, identify misconfigurations, and exfiltrate data without being detected. To prevent attacks from gaining access and gaining access credentials, proper technical controls are required to restrict the use of user accounts and administrative privileges.

Along with the lack of controls, they had misconfigured security controls too that didn't follow a least privilege or a zero-trust architecture. They should have leveraged proper IAM roles, groups, and permissions to restrict access to sensitive data and services. IAM controls should be audited to ensure proper permissions are given to roles, users, and groups. Proper evaluation of IAM should be conducted regularly to ensure roles, users, or groups are needed or properly disabled or deleted if not needed. How often is the WAF reconfigured? Should that role be active all the time or should it be disabled if not used within 30 days? Proper use of IAM should disable roles if not currently in use or delete them if they are no longer needed.

Lastly, data loss prevention (DLP) was not properly implemented to restrict the ability to exfiltrate data from the network. Visibility into the network and proper tools could have detected data trying to leave from over 700 different S3 buckets. Technical controls could have aided Capital One in detecting, mitigating, and preventing the attack. They leveraged the cloud for their customer facing applications and their large network creates an equally large attack surface. Proper processes, procedures and testing would have helped them in catching the misconfigured WAF through audits and offensive/defensive security engagements such as red teaming, tabletop exercises, and bug bounty.

AWS Support

I believe that Capital One was responsible for the data breach. They had tools at their disposal from AWS, 3rdparties, and open-source projects. Ultimately, the attack could have been prevented, and according to the cloud responsibility model, the customer is responsible for misconfiguring security controls and AWS is not at fault. I do think that AWS could have made S3 buckets more secure by default. For example, maybe 700 buckets shouldn’t be so easily accessible by one role. I think tools should be provided by AWS to aid users with proper configuration of applications and services. The services shouldn't contain data related to security configurations. They could implement security best practices, testing, and alert users during the configuration process. They also could provide tools to scan the application or service after it is configured. In VMware's 2021 State of Cloud Security report, 1 in 6 companies surveyed experienced a cloud data breach due to a misconfiguration in the past year. (Cozens, 2022) AWS should support its users by providing mechanisms to check for misconfigurations and security issues.

Conclusion

In conclusion, the attack was successful in consequence to the failure of five distinct security controls. Capital One discovered that an external user accessed their network and took proper actions to secure the network and inform the impacted individuals. The 5 failed security controls are outlined below (Khan, 2022):

1. A misconfigured WAF called Modsecurity WAF

2. Weakness of the cloud infrastructure that enabled querying the metadata service, allowing the attacker to gain access through temporary credentials

3. The lack of least-privilege and allowing an IAM role access to the S3 storage buckets

4. Improper use of cryptography and keys to properly encrypt data

5. The ineffectiveness of the intrusion detection (IDS or IPS) and monitoring capabilities

Capital One has faced fines and penalties of 80 million dollars because of the data breach and has suffered reputational and other non-monetary consequences. (Press, 2020) This is a good example of how a well-funded and highly regulated organization is not immune to cyber-attacks. Proper due diligence needs to be done to ensure their customers and their organization is properly secured. In addition, cloud service providers can do better to make their services more secure by default and provide mechanisms to evaluate the security posture of services and applications through tools and configurations.

References

Cozens, B. (2022, June 9). *Cloud data breaches: 4 biggest threats to Cloud Storage Security*. Malwarebytes. Retrieved November 20, 2022, from https://www.malwarebytes.com/blog/business/2022/06/cloud-data-breaches-4-biggest-threats-to-cloud-storage-security

Khan, S., Kabanov, I., Hua, Y., & Madnick, S. (2022, February 1). *A systematic analysis of the Capital One Data Breach: Critical Lessons Learned*. ACM Transactions on Privacy and Security. Retrieved December 2, 2022, from https://dl.acm.org/doi/10.1145/3546068

One, C. (2022). *2019 capital one cyber incident: What happened*. Capital One. Retrieved December 2, 2022, from https://www.capitalone.com/digital/facts2019/

Press, T. A. (2020, August 7). *Capital one is fined $80 million for huge data breach*. Fortune. Retrieved December 2, 2022, from https://fortune.com/2020/08/07/capital-one-fined-80-million-data-breach/

Trends, M. (2022). *Capital One Financial Revenue 2010–2022: COF*. Macrotrends. Retrieved December 2, 2022, from https://www.macrotrends.net/stocks/charts/COF/capital-one-financial/revenue

Waldman, A. (2022, June 20). *Paige Thompson found guilty in 2019 capital one data breach: TechTarget*. Security. Retrieved December 2, 2022, from https://www.techtarget.com/searchsecurity/news/252521775/Paige-Thompson-found-guilty-in-2019-Capital-One-data-breach

Wolf, A. (2022, June 18). *The capital one breach: Its impact and what it teaches us about cybersecurity*. Arctic Wolf. Retrieved December 2, 2022, from https://arcticwolf.com/resources/blog/the-capital-one-breach-its-impact-and-what-it-teaches-us-about-cybersecurity/

AWS Security Business Learning Banking