Understanding the CapitalOne Data Breach

Prepared by -
Gaurav Suryawanshi
[Lead Cloud Security Consultant]
as a part of
CYBERSEC 590.03.Fa24
Mid-term report

Document Version Control

Version	Date	Changelog summary	Reviewed by	Approved by	Comments
v1.0	15th October, 2024	Uploaded the initial case study report for mid-term submission	Gaurav Suryawanshi (Self)	Gaurav Suryawanshi (Self)	Waiting for Prof. David's feedback and comments.

Table of Contents

Understanding CapitalOne’s historical context

CapitalOne is one of the leading North American banks, with approximately 50,000 employees and $37 billion in revenue in 2024. The company operates in a highly regulated industry and complies with various regulations, including those related to corporate governance, cybersecurity, and financial protection. CapitalOne values technology and has been an early adopter of cloud computing, with a majority of its applications operating in the cloud.

The events that led to the breach can be traced back to the implementation of CapitalOne's Cloud Strategy in 2014, coinciding with George Brady's appointment as the bank's Chief Technology Officer. During this period, CapitalOne made daring and unconventional decisions for a company operating in a heavily regulated industry such as finance. These decisions included a commitment to open-source technology, the adoption of agile development principles, and a shift to the public cloud (specifically AWS instead of a private cloud). This strategic transformation was accompanied by an aggressive recruitment of tech talent. Furthermore, the bank actively cultivated a public image as a forward-thinking technology company, rather than a traditional bank, focusing on rapidly developing new capabilities to enhance customer experience.

Figure 0.1 – Humor to denote CapitalOne’s aggressive approach and early adoption of AWS against a private cloud

The company has worked closely with AWS to develop a security model for operating securely in the cloud. However, this move to the cloud also played a key role in a data leak incident in 2019. [3] The cyberattack on CapitalOne in 2019 was one of the biggest data breaches in the financial sector, where a hacker illegally accessed the details of more than 100 million customers in the US and roughly 6 million in Canada.

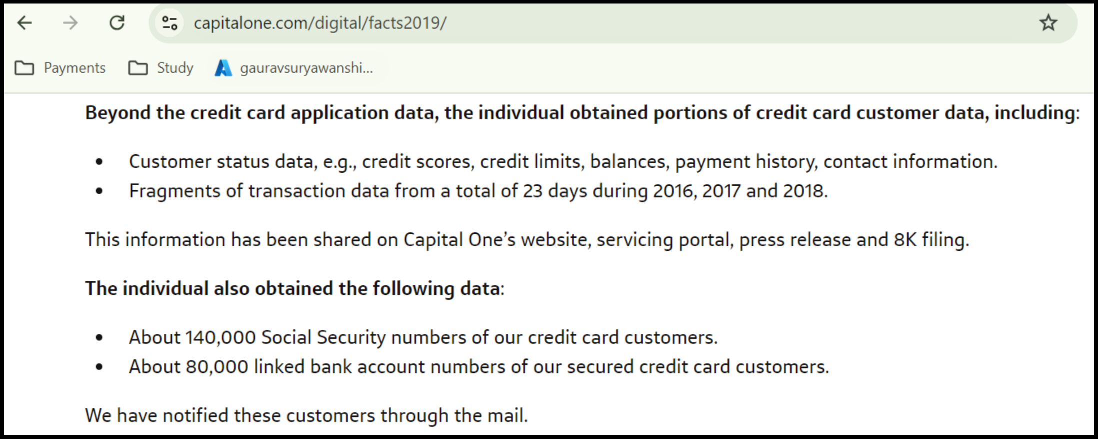

Beyond the credit card application data, the individual obtained portions of credit card customer data, including:

- Customer status data, e.g., credit scores, credit limits, balances, payment history, contact information.
- Fragments of transaction data from a total of 23 days during 2016, 2017 and 2018.

This information has been shared on Capital One’s website, servicing portal, press release and 8K filing.

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.

We have notified these customers through the mail.

Figure 0.2 - Snippet from CapitalOne’s website disclosing the specifics of the data which was breached

This breached data encompassed information such, as names, addressess phone numbers, email addresses, date of birth, self reported income credit scores, credit limits, balances and payment history. It is important to highlight that no credit card account numbers or login credentials were exposed during this incident. The breach occurred mainly due to a SSRF attack done against a firewall misconfiguration sitting on a EC2 instance - hosted on AWS that the hacker took advantage of, who was also a former employee of AWS.

Chapter 1 - Diving deep into the intricacies of the data breach

What happened and how did they discover the attack?

On a busy Wednesday in July 2019, the security team at CapitalOne had a meeting regarding enhancing the security of its systems. The team was in the process of creating a phishing email for all staff members when they were surprised by an email in their responsible disclosure mailbox. This mailbox was meant for people to report any security vulnerabilities they came across CapitalOne's assets or internal/external digital infrastructure.

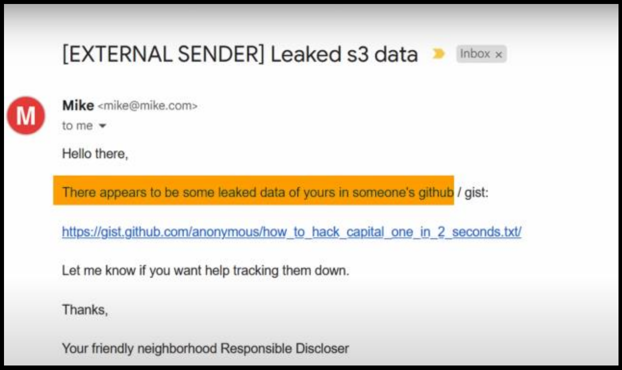

[EXTERNAL SENDER] Leaked s3 data Inbox ×

Mike <mike@mike.com>
to me

Hello there,

There appears to be some leaked data of yours in someone's github / gist:

https://gist.github.com/anonymous/how_to_hack_capital_one_in_2_seconds.txt/

Let me know if you want help tracking them down.

Thanks,

Your friendly neighborhood Responsible Discloser

Figure 1.1 – Snippet of the email received (original sender's name redacted) [5]

CapitalOne's security team also noted that they could be reached via Twitter DMs, if anyone wants to report any vulnerabilities and/or flaws in their infrastructure. Although initially concerning, such reports often pertained to minor bugs that could be resolved before they were exploited. However, the email indicated that there was a potential leak of S3 data on someone's GitHub. It was apparent that the damage had already been done.

Figure 1.2 - Humor to denote the falling CapitalOne's shares after the attack details went public

How does a SSRF attack work?

A SSRF, or Server-Side Request Forgery attack, tricks a server into fetching or manipulating data from an unintended location. In the CapitalOne breach, the attacker exploited a vulnerability in a web application hosted on an EC2 instance.

Here's how it worked,

1.) **EC2 Instance:** CapitalOne was using an EC2 instance to run part of their application. This instance had an IAM role attached to it, which determined what AWS resources the instance could access.

2.) **Vulnerability Exploitation:** The attacker sent specially crafted requests to the application, exploiting the SSRF vulnerability. This allowed the attacker to make requests to the metadata service of the EC2 instance.

3.) **Accessing the Metadata Service:** The metadata service provides information about the EC2 instance, including IAM role credentials. By accessing this service, the attacker was able to retrieve temporary security credentials associated with the IAM role.

4.) **Gaining Unauthorized Access:** With those credentials, the attacker could access resources within the permissions granted to that IAM role. For example, if the role had permission to read data from certain S3 buckets, the attacker could do the same. This is likely how the attacker accessed the vast amounts of data stored by CapitalOne.

In simpler terms, by exploiting a weakness in how the server handled incoming requests, the attacker tricked the server into giving away its secrets, which included the keys to access sensitive information.

Let's use a classroom analogy to explain a SSRF attack.

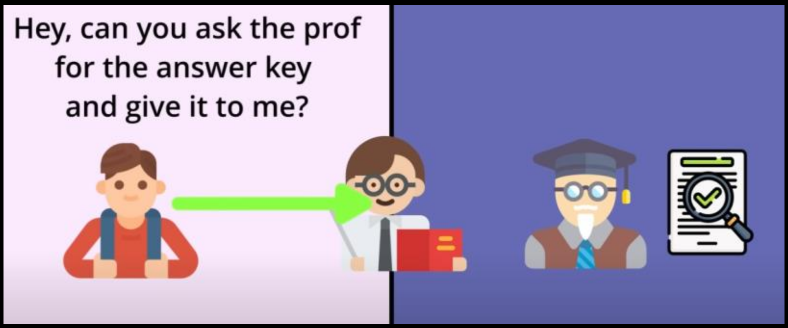
Hey, can you ask the prof for the answer key and give it to me?

Figure 1.3 – Classroom analogy [5]

Imagine a teacher (the server) has a teaching assistant (the EC2 instance) who helps manage the classroom's resources, like books and supplies (the S3 buckets). The assistant has special permissions (IAM role) to access these resources, given by the teacher.

Now, imagine a mischievous student (the attacker) who wants to get access to the answer keys for the exams (sensitive data). The student can't ask the teacher directly because they know they're not allowed to have it. So, they devise a plan.

The student tricks the teaching assistant into thinking they're asking for a different resource (a normal request). The assistant, not realizing the trick, goes to the teacher and asks for the answer key instead of the intended resource. The teacher unknowingly hands it over because it came through the trusted assistant.

In this analogy,

- The teacher is the server.
- The teaching assistant is the EC2 instance with special permissions.
- The student is the attacker.
- The answer key is the sensitive data.

The student, through clever manipulation, gets the assistant to ask the teacher for something they are not allowed to obtain. The teaching assistant (EC2 instance) thinks they’re doing a regular task, not realizing they’ve been duped into requesting something they shouldn’t have access to. The teacher (the server), trusting the assistant’s request, unwittingly hands over the sensitive data. The student (attacker) now has the answers (sensitive data) without the teacher ever knowing they were involved in something they shouldn’t have been. Similarly, in this data breach - the attacker (Ms. Paige Thompson) was the student, the misconfigured WAF webserver (ModSecurity EC2 instance) with over permissive IAM privileges was the TA, the (EC2 metadata service) was the professor who handed out the answer key, was the customer data.

In a real-life SSRF attack, just like the student tricking the teaching assistant, the attacker tricks the server into asking for something it shouldn't. The server, believing it's a legitimate request, ends up handing over sensitive information. And just like in the classroom scenario, it all happens without the main authority (the server/teacher) realizing the error.

Who was the threat actor?

During the CapitalOne data breach incident, a former AWS employee, by the name of Paige Thompson managed to take advantage of a firewall misconfiguration on a server belonging to CapitalOne. This security loophole granted her access to the information of millions of people. Under the alias *'erratic'* Thompson was known for her involvement in cyber-attacks. She had prior experience in software engineering and cloud computing through her work and insider knowledge at Amazon Web Services, it was likely that her technical background equipped her with the artifacts needed to detect and exploit the flaw.

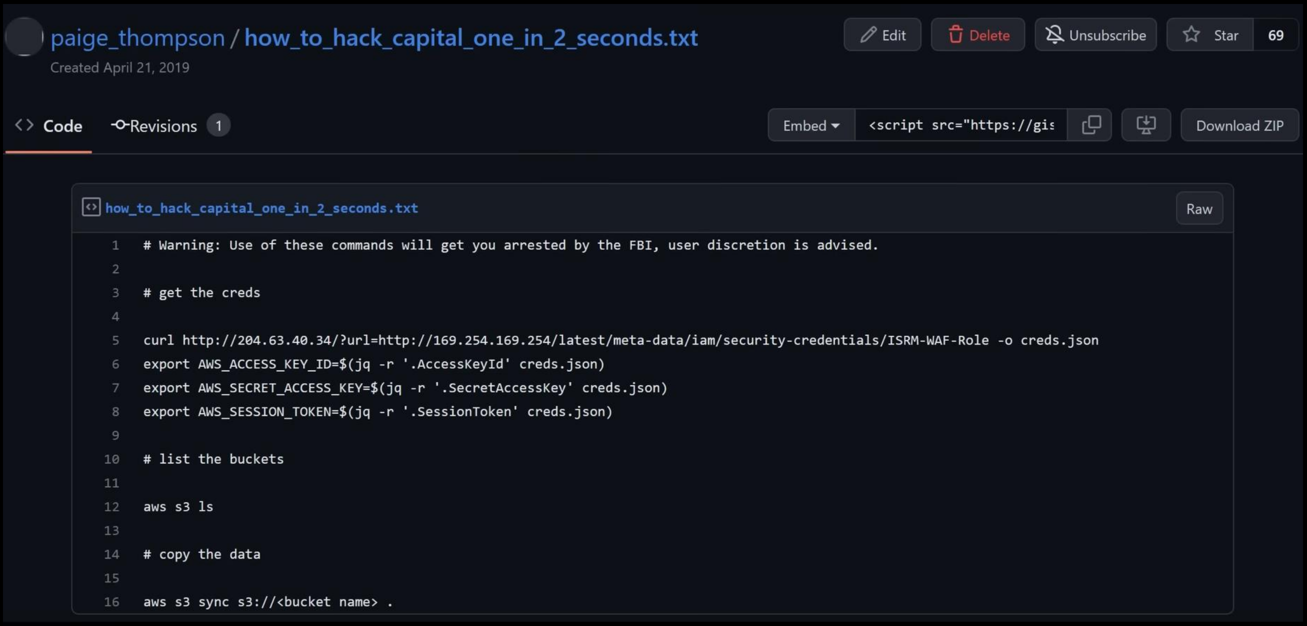

```
# Warning: Use of these commands will get you arrested by the FBI, user discretion is advised.

# get the creds

curl http://204.63.40.34/?url=http://169.254.169.254/latest/meta-data/iam/security-credentials/ISRM-WAF-Role -o creds.json
export AWS_ACCESS_KEY_ID=$(jq -r '.AccessKeyId' creds.json)
export AWS_SECRET_ACCESS_KEY=$(jq -r '.SecretAccessKey' creds.json)
export AWS_SESSION_TOKEN=$(jq -r '.SessionToken' creds.json)

# list the buckets

aws s3 ls

# copy the data

aws s3 sync s3://<bucket name> .
```

Figure 1.3 – Capture of CapitalOne attack PoC from attacker (Ms. Thompson's GitHub) [5]

Thompson searched the internet for secured firewalls and identified a weakness, in CapitalOne's systems which allowed her to breach their data stored on AWS servers. She then published the PoC, and steps needed to extract and access this information which she exfiltrated from CapitalOne's assets, on GitHub. GitHub is a platform, which is generally used within the developer community and enterprises for creating and sharing code. This left clues and traces that eventually led the investigating authorities such as FBI to her whereabouts. This incident highlighted the significance of setting up and safeguarding cloud systems when dealing with such sensitive data. Thompsons expertise, in technology and familiarity with AWS frameworks were factors, in the effectiveness of her breach.

Timeline of Key Events

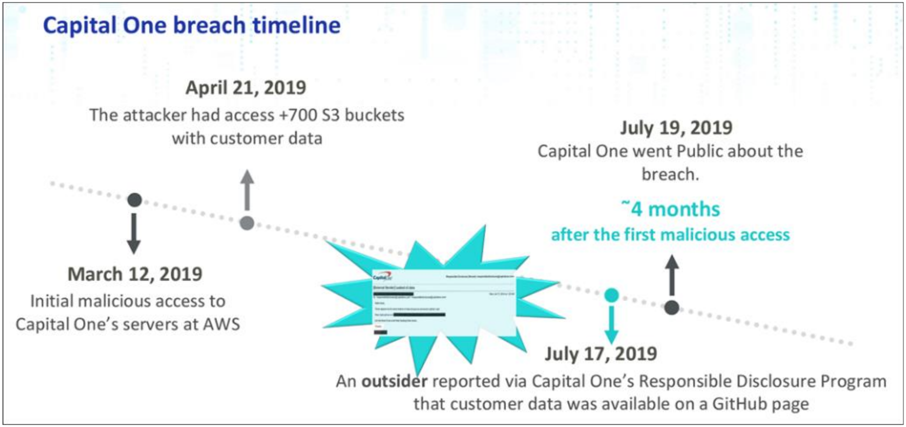
Capital One breach timeline
April 21, 2019
The attacker had access +700 S3 buckets with customer data
July 19, 2019
Capital One went Public about the breach.
~4 months
after the first malicious access
March 12, 2019
Initial malicious access to Capital One's servers at AWS
July 17, 2019
An outsider reported via Capital One's Responsible Disclosure Program that customer data was available on a GitHub page

Figure 1.4 - Timeline of the attack events during the on-going data breach [4]

March 12th, 2019

The attacker Ms. Paige Thompson, a former Amazon employee was able to obtain initial access to the CapitalOne's infrastructure which was hosted on the AWS servers.

April 21st, 2019

Attacker had access to more than 700+ AWS S3 buckets which contained sensitive data of the customers belonging to CapitalOne services and offerings.

July 17th, 2019

Three months later, CapitalOne failed to detect this ongoing attack through their regular cybersecurity operations team. However, CapitalOne discovered the data breach and ongoing data exfiltration through their VDP (Vulnerability Discovery Program), when one of the bug bounty hunter (who was an outsider) reported this issue through their VDP portal.

July 19th, 2019

CapitalOne decided to go public and disclose this incident to the media and external stakeholders, after they started the investigation on July 17th, 2019 – which was 4 months after the initial foothold by the attacker.

Chapter 2 – What went wrong?

In the CapitalOne breach, the key AWS services which were vulnerable were:

Amazon S3 (Simple Storage Service): This is a widely used storage service for data. In this breach, CapitalOne stored sensitive information in S3 buckets. The attacker exploited a misconfiguration in the firewall of a web application, enabling her to communicate with CapitalOne 's S3 buckets through their web application's backend.

AWS IAM (Identity and Access Management): IAM securely manages access to AWS services and resources. The attacker found a way to escalate privileges within IAM, granting her broader access than she should have had. This misconfiguration played a crucial role in her access to the S3 data.

AWS EC2 (Elastic Compute Cloud): EC2 offers scalable computing capacity. The attacker utilized an EC2 instance to execute commands after gaining access through the misconfigured web application firewall. This EC2 instance was pivotal in her ability to access and exfiltrate data.

Attack Vector	**Service/Component**	**Service/Component belongs to AWS? (Y/N)**	**Vulnerability**	**Impact**
Existence of a reverse proxy	ModSecurity Web Application Firewall (WAF)	N	Misconfiguration allowed Server-Side Request Forgery (SSRF) attacks.	Enabled unauthorized access to AWS metadata service, leading to credential theft and data exfiltration.
Architecture gap of the cloud infrastructure that enabled querying the metadata service	AWS Metadata Service	Y	Accessed via SSRF vulnerability through the WAF.	The queried temporary login details (AccessKeyId & SecretAccessKey) were utilized to access AWS services.
Existence of an over-provisioned IAM role which granted access to the S3 storage buckets	AWS Identity and Access Management (IAM) Role [ISRM-WAF-Role]	Y	Over-permissioned IAM role associated with the WAF.	Allowed the attacker to list and read data from Amazon S3 buckets beyond necessary permissions.
Ineffectiveness of the encryption method which was being used	AWS S3 (Simple Storage Service) Buckets	Y	The key:value pair for the KMS (Key Management System) was loosely set.	It was speculated that the role attached to the instance (ISRM-WAF-Role) also allowed decryption of data, since it most likely had kms:decrypt privilege as well.
Ineffectiveness of the intrusion detection and monitoring systems in place	n/a	*not found*	Absence of IDS and the failure in monitoring security events	The IDS was not operational which failed to inform the security / SOC teams to inform about the ongoing attack, where the damage could have been reduced.

Table 2.1 – Analysis of AWS and third-party services and/or modules impacted

Technical Analysis of the data breach

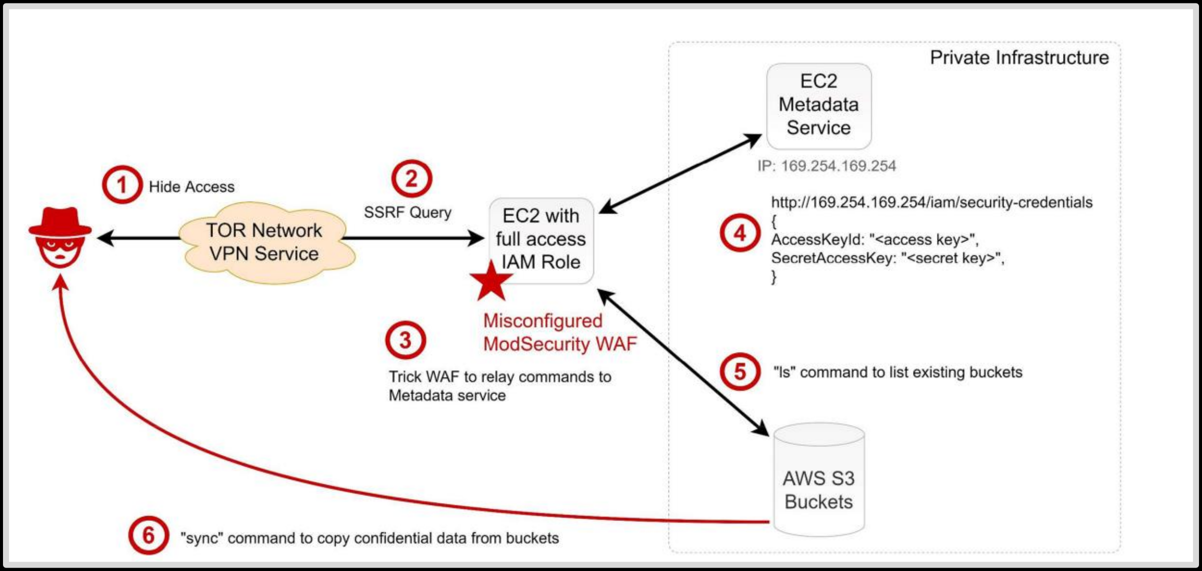
Private Infrastructure
EC2 Metadata Service
IP: 169.254.169.254
1 Hide Access
2 SSRF Query
TOR Network VPN Service
EC2 with full access IAM Role
Misconfigured ModSecurity WAF
3 Trick WAF to relay commands to Metadata service
4 http://169.254.169.254/iam/security-credentials { AccessKeyId: "<access key>",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

Figure 2.1 – Technical breakdown and specifics of the constructed attack [1]

1. The FBI's cyber-experts and CapitalOne's internal analysis, identified multiple unauthorized accesses through anonymizing services such as the TOR network and the VPN service provider IPredator. These services were used to conceal the source IP address of the malicious accesses.

2. The SSRF attack enabled the attacker to deceive the server into executing commands as a remote user, granting the attacker access to the private server.

3. A misconfiguration in the WAF allowed the intruder to manipulate the firewall into relaying commands to a default back-end resource on the AWS platform, specifically the metadata service with temporary credentials for the environment (accessed through the URL http://169.254.169.254).

4. Through a combination of the SSRF attack and the WAF misconfiguration, the attacker utilized the following URL endpoint - http://169.254.169.254/iam/security-credentials to obtain the AccessKeyId and SecretAccessKey credentials from a role referred to in the FBI indictment as "*****-WAF-Role" (with the name partially redacted). The resulting temporary credentials granted the attacker the ability to execute commands in the AWS environment via API, CLI, or SDK.

5. Using the obtained credentials, the attacker executed the "ls" command multiple times, which provided a comprehensive list of all AWS S3 Buckets associated with the compromised CapitalOne account ("$ aws s3 ls").

6.) Finally, the attacker utilized the AWS "sync" command to duplicate approximately 30 GB of CapitalOne 's credit application data from these buckets to the attacker's local machine ($ aws s3 sync s3://bucketone). As indicated in the FBI report, this command allowed the attacker to gain access to over 700 buckets.

In summary, the attack succeeded due to five specific control failures:

(1) A misconfigured reverse proxy (ModSecurity WAF)

(2) Vulnerabilities in the cloud infrastructure that allowed access to the metadata service and temporary credentials

(3) An over-provisioned IAM role that provided access to S3 storage buckets,

(4) Ineffective encryption methods

(5) Inadequate intrusion detection and monitoring systems.

Mapping of System hazards and Constraints violations to Cyber Kill Chain

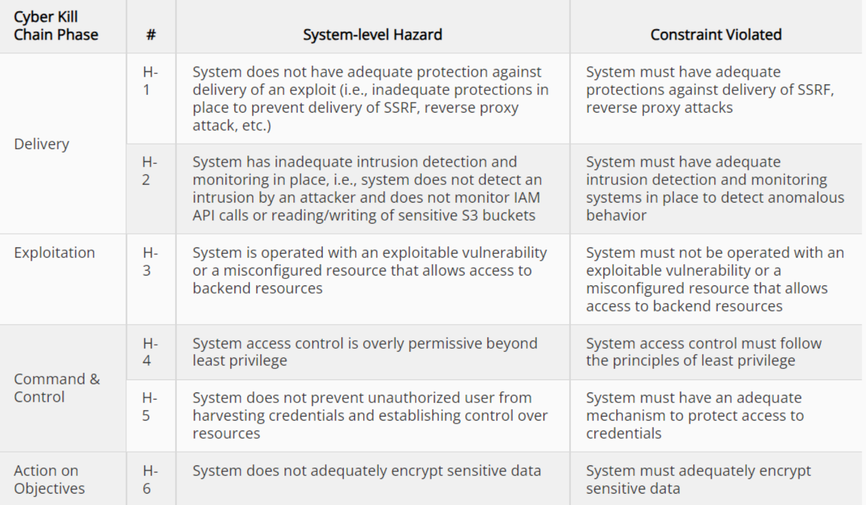

Cyber Kill Chain Phase	#	System-level Hazard	Constraint Violated
Delivery	H-1	System does not have adequate protection against delivery of an exploit (i.e., inadequate protections in place to prevent delivery of SSRF, reverse proxy attack, etc.)	System must have adequate protections against delivery of SSRF, reverse proxy attacks
	H-2	System has inadequate intrusion detection and monitoring in place, i.e., system does not detect an intrusion by an attacker and does not monitor IAM API calls or reading/writing of sensitive S3 buckets	System must have adequate intrusion detection and monitoring systems in place to detect anomalous behavior
Exploitation	H-3	System is operated with an exploitable vulnerability or a misconfigured resource that allows access to backend resources	System must not be operated with an exploitable vulnerability or a misconfigured resource that allows access to backend resources
Command & Control	H-4	System access control is overly permissive beyond least privilege	System access control must follow the principles of least privilege
	H-5	System does not prevent unauthorized user from harvesting credentials and establishing control over resources	System must have an adequate mechanism to protect access to credentials
Action on Objectives	H-6	System does not adequately encrypt sensitive data	System must adequately encrypt sensitive data

Figure 2.1 - Mapping of the takeaways, implementations, actions post the data-breach framework [1]

During the assessment of the CapitalOne breach, we have observed the seven significant risks associated with different stages of the Cyber Kill Chain.

Mapping of the CapitalOne's data breach to MITRE ATT&CK framework

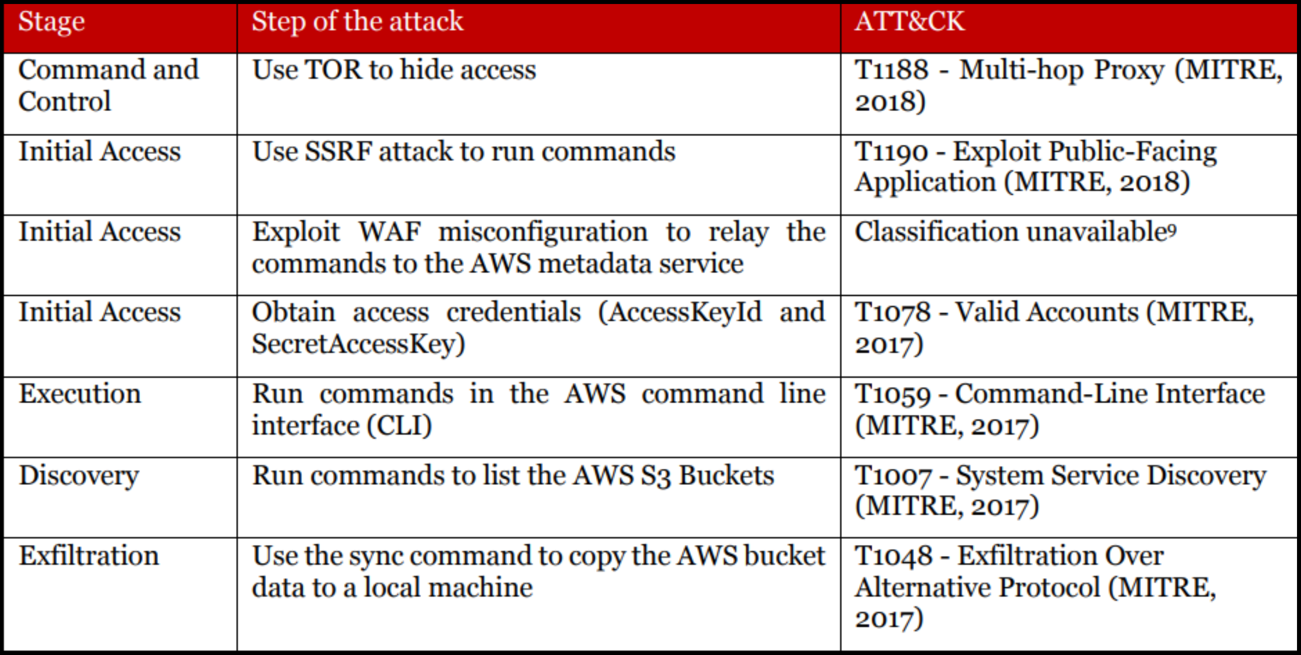

Stage	Step of the attack	ATT&CK
Command and Control	Use TOR to hide access	T1188 - Multi-hop Proxy (MITRE, 2018)
Initial Access	Use SSRF attack to run commands	T1190 - Exploit Public-Facing Application (MITRE, 2018)
Initial Access	Exploit WAF misconfiguration to relay the commands to the AWS metadata service	Classification unavailable[9]
Initial Access	Obtain access credentials (AccessKeyId and SecretAccessKey)	T1078 - Valid Accounts (MITRE, 2017)
Execution	Run commands in the AWS command line interface (CLI)	T1059 - Command-Line Interface (MITRE, 2017)
Discovery	Run commands to list the AWS S3 Buckets	T1007 - System Service Discovery (MITRE, 2017)
Exfiltration	Use the sync command to copy the AWS bucket data to a local machine	T1048 - Exfiltration Over Alternative Protocol (MITRE, 2017)

Figure 2.2 – Mapping of the attack steps and phases to MITRE ATT&CK enterprise framework [2, 3, 4]

Chapter 3 – How could have CapitalOne detected the incident?

CapitalOne's data breach could have been detected earlier through the implementation of several proactive measures, all of which encompass close monitoring, alerting systems, and regular reviews of their cloud infrastructure. This can be explained in the following manner -

Enhanced Monitoring and Logging

CapitalOne could have established more advanced monitoring and logging systems for their AWS services. For instance, logging every instance of access to sensitive data stored in S3 buckets would have provided valuable insights. Tools such as AWS CloudTrail enable organizations to maintain a comprehensive record of all activities on their AWS accounts, including data access, timestamps, and the source of access. Vigilant monitoring of these logs would have revealed unusual patterns, such as access requests from unfamiliar locations or repeated attempts to access substantial amounts of data. Real-time monitoring would have allowed the identification of abnormal activities, such as sudden spikes in data access from unauthorized sources, triggering timely alerts for further investigation by the security team.

Intrusion Detection Systems (IDS)

An Intrusion Detection System (IDS) functions as a vigilant network alarm system. It can identify and raise alerts or block suspicious activities, such as unauthorized attempts to access data. In the case of CapitalOne, an IDS could have detected the hacker's attempts to communicate with internal servers through the misconfigured firewall. The implementation of an IDS would act as an early warning system, preventing unauthorized access to sensitive data.

Regular Audits and Security Reviews

Conducting routine security audits to ensure proper configuration is a fundamental preventive measure. In the context of CapitalOne, regular audits of their firewall and permissions could have identified the vulnerability in their EC2 instance. By analogy, similar to conducting routine checks on a home security system to identify and rectify vulnerabilities, consistent cybersecurity audits would have potentially identified the misconfiguration that allowed the attacker to access their data.

Real-Time Threat Detection

Utilizing tools like Amazon GuardDuty, CapitalOne could have established continuous real-time monitoring for suspicious activities, such as unauthorized attempts to access specific services. This tool automatically identifies unusual behaviors and can promptly alert the security team. Configured to monitor abnormal requests, such as those pertaining to EC2 metadata, Amazon GuardDuty might have intercepted the attacker before they escalated their privileges and exfiltrated data.

Response Drills and Incident Planning

Despite robust security measures, breaches can occur. The ability to respond promptly is crucial. CapitalOne could have conducted regular incident response drills to practice coordinated responses to potential breaches. A well-practiced response plan would have facilitated swift actions to mitigate the breach, including isolating affected instances, revoking permissions, and initiating internal investigations before the situation escalated.

In summary, CapitalOne could have detected the breach earlier by:

- Vigilantly monitoring AWS logs.

- Implementing an IDS to flag suspicious activities.

- Conducting regular security configurations and permissions reviews.

- Utilizing real-time threat detection tools like Amazon GuardDuty.

- Practicing incident response drills for swift reactions.

Adhering to these measures could have significantly increased their ability to detect and respond to the breach before the hacker could exfiltrate millions of records.

Additionally, to help detect and potentially prevent incidents like the CapitalOne data breach, a RACI matrix can be used to define roles and responsibilities across a Security Operations Center (SOC), Network Operations Center (NOC), and Global Operations Center (GOC). The RACI model clarifies who is Responsible, Accountable, Consulted and Informed for various tasks. Here's how these roles could be aligned:

Task/Activity	SOC	NOC	GOC
Threat Detection and Monitoring	R (Responsible)	C (Consulted)	I (Informed)
Incident Response Coordination	A (Accountable)	C (Consulted)	I (Informed)
Network Performance Monitoring	C (Consulted)	R (Responsible)	I (Informed)
Configuration Management	C (Consulted)	A (Accountable)	I (Informed)
Security Policy Implementation	A (Accountable)	C (Consulted)	I (Informed)
Communication with Stakeholders	I (Informed)	I (Informed)	R/A (Responsible/Accountable)
Resource Allocation for Incident Handling	C (Consulted)	C (Consulted)	R/A (Responsible/Accountable)
Log Analysis and Anomaly Detection	R/A (Responsible/Accountable)	C (Consulted)	I (Informed)

Table 3.1 – RACI Matrix in alignment to the proposed SOC, NOC and GOC

Security Operations Center (SOC): The primary responsibility of the SOC is to detect threats, analyze security incidents, and coordinate responses. The team is accountable for implementing security policies and conducting log analysis to identify anomalies.

Network Operations Center (NOC): The NOC focuses on maintaining network performance and ensuring that any disruptions or anomalies are promptly addressed. They work closely with the SOC on configuration management to ensure that security measures are integrated into network operations.

Global Operations Center (GOC): The GOC serves as the central hub for communication and coordination across the organization. They are responsible for allocating resources during incidents and ensuring that all stakeholders are informed.

This structured approach ensures comprehensive coverage of all aspects of security monitoring, detection, and response, potentially reducing the impact of breaches such as the one experienced by CapitalOne.

Chapter 4 – How could have CapitalOne prevented the incident?

CapitalOne could have took several steps to prevent this incident from its occurrence, such as

Aggressive patch management and proactive auditing

The SecOps team should have regularly audited their infrastructure and configurations to ensure security measures were, up to date and effective. For instance examining the permissions assigned to IAM roles and inspecting the security groups linked to EC instances might have uncovered any settings.

Triggering incident response playbooks

Setting up automated incident responses to specific triggers could have helped contain a breach more efficiently. For instance, if an IAM role suddenly requested access to resources it doesn't typically use, an automated response could have promptly revoked its permissions.

Employing the Principle of Least privilege

Applying this principle involves making sure that every user and application has enough access rights they require – no more, than necessary. For instance if there's a program running on EC2 that only needs permission, for reading buckets it shouldn't have the capability of listing or deleting items.

Isolation of large network ranges and subnets

Network Segmentation involves the practice of dividing systems and data repositories, from vulnerable areas within the network structure to enhance security measures effectively safeguarding against potential breaches in case one segment is compromised by an unauthorized party; hence ensuring that the entire system remains protected from complete exposure to threat actors. For example, it includes organizing servers responsible, for managing public interactions separately from those handling confidential customer information.

Usage of an enterprise-graded WAF over ModSecurity's open-source WAF

Using a Web Application Firewall (WAF) can be beneficial as it filters and monitors HTTP requests, for a web applications security purposes. Through the implementation of a WAF system in place could have potentially prevented the requests that triggered the SSRF attack scenario. As an example suppose there is an attempt to reach the EC2 metadata service through a request – something that generally does not occur in operations, in instances the WAF would be capable of either flagging or stopping such requests promptly.

Restricting access to metadata services and environments variables

It is important for CapitalOne to consider implementing restrictions on access to the EC2 metadata service, such as through the use of IAM policies or by limiting access to specific IP addresses. For instance, applications running on EC2 that do not require access to metadata should explicitly be denied access.

Providing security awareness and trainings to stakeholders and employees

Additionally, providing regular security training for developers and system administrators to educate them on secure coding practices and AWS security features, including the risks of SSRF and how to mitigate them through code, could have been beneficial.

CapitalOne stated that the initial and key reason for the success of this attack was the misconfigured WAF which was deployed as a reverse proxy on their AWS EC2 instance. This would have been prevented and patched if there was an active vulnerability scanning and reporting of misconfigurations ahead of time.

However, it is believed that two of the major reasons for the attacker's success were as follows -

1.) Ms. Paige Thompson was able to trick the metadata service to request access credentials *AccessKeyId* and *SecretAccessKey* (similar to "root access"), which allowed them to run commands in the servers hosted in the CapitalOne's AWS environment. Following NIST controls would be able to prevent the attacker from getting the access to temporary credentials by monitoring and auditing the use of administrative accounts:

- PR.AC-1: Identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users and processes;
- PR.AC-4: Access permissions and authorizations are managed, incorporating the principles of least privilege and separation of duties;
- PR.AC-6: Identities are proofed and bound to credentials and asserted in interactions;
- PR.AC-7: Users, devices, and other assets are authenticated (e.g., single-factor, multi-factor) commensurate with the risk of the transaction (e.g., individuals' security and privacy risks and other organizational risks);
- PR.IP-1: A baseline configuration of information technology/industrial control systems is created and maintained incorporating security principles (e.g. concept of least functionality);
- PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy;
- PR.PT-3: The principle of least functionality is incorporated by configuring systems to provide only essential capabilities;
- DE.AE-3: Event data are collected and correlated from multiple sources and sensors;
- DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events;
- DE.CM-7: Monitoring for unauthorized personnel, connections, devices, and software is performed;
- DE.DP-2: Detection activities comply with all applicable requirements. [1, 2, 4]

2.) Ms. Paige Thompson used a synchronization (sync) command on CapitalOne's server located on the AWS cloud platform to exfiltrate 30GB+ of confidential data by transferring information from AWS S3 buckets to their personal computer. Following NIST controls would have helped in preventing data exfiltration by restricting remote access and by monitoring outbound traffic:

- ID.AM-3: Organizational communication and data flows are mapped;
- PR.AC-3: Remote access is managed;
- PR.DS-1: Data-at-rest is protected;
- PR.DS-5: Protections against data leaks are implemented;
- PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy;
- PR.PT-3: The principle of least functionality is incorporated by configuring systems to provide only essential capabilities;
- DE.AE-1: A baseline of network operations and expected data flows for users and systems is established and managed;
- DE.AE-3: Event data are collected and correlated from multiple sources and sensors;
- DE.CM-1: The network is monitored to detect potential cybersecurity events;
- DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events;

- DE.CM-7: Monitoring for unauthorized personnel, connections, devices, and software is performed;
- DE.DP-2: Detection activities comply with all applicable requirements [1, 2, 4]

Chapter 5 - Lessons and Takeaways from the CapitalOne's Data Breach

The CapitalOne case revealed several technical lapses in configuration and programming that contributed to the breach. It is recommended to prioritize a robust application review and vulnerability assessment process to identify and patch vulnerable applications. Additionally, understanding the shared responsibility model and implementing the "trust but verify" security principle is crucial, particularly in cloud security. Enforcing the least privilege security principle for all IAM policies and roles is essential to prevent excessive permissions that may lead to breaches.

Moving to operational controllers, it is imperative to implement an effective Intrusion Monitoring/Detection system to raise alarms for anomalous behavior, particularly in IAM and AWS Security Token Service (STS) API calls. Adhering to the "trust but verify" security principle and approaching security as a system issue are also critical for operational controllers, such as DevOps and Security Architects.

Furthermore, both cloud service providers and organizations should prioritize security and ensure coordination and communication between security and development teams. Cloud service providers should focus on secure architecture and simplicity of design, revisiting the pace of releasing new services to minimize technical debt and knowledge gaps.

Overall, it is necessary to address these recommendations to strengthen security measures and prevent similar breaches in the future.

Following are the changes which CapitalOne made, post the incident:

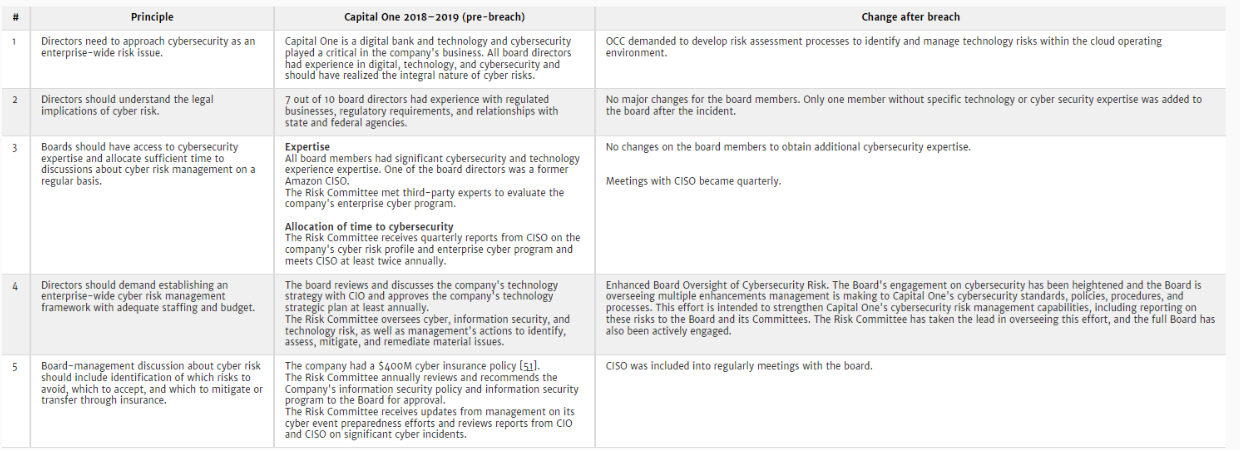

#	Principle	Capital One 2018–2019 (pre-breach)	Change after breach
1	Directors need to approach cybersecurity as an enterprise-wide risk issue.	Capital One is a digital bank and technology and cybersecurity played a critical in the company's business. All board directors had experience in digital, technology, and cybersecurity and should have realized the integral nature of cyber risks.	OCC demanded to develop risk assessment processes to identify and manage technology risks within the cloud operating environment.
2	Directors should understand the legal implications of cyber risk.	7 out of 10 board directors had experience with regulated businesses, regulatory requirements, and relationships with state and federal agencies.	No major changes for the board members. Only one member without specific technology or cyber security expertise was added to the board after the incident.
3	Boards should have access to cybersecurity expertise and allocate sufficient time to discussions about cyber risk management on a regular basis.	**Expertise** All board members had significant cybersecurity and technology experience expertise. One of the board directors was a former Amazon CISO. The Risk Committee met third-party experts to evaluate the company's enterprise cyber program. **Allocation of time to cybersecurity** The Risk Committee receives quarterly reports from CISO on the company's cyber risk profile and enterprise cyber program and meets CISO at least twice annually.	No changes on the board members to obtain additional cybersecurity expertise. Meetings with CISO became quarterly.
4	Directors should demand establishing an enterprise-wide cyber risk management framework with adequate staffing and budget.	The board reviews and discusses the company's technology strategy with CIO and approves the company's technology strategic plan at least annually. The Risk Committee oversees cyber, information security, and technology risk, as well as management's actions to identify, assess, mitigate, and remediate material issues.	Enhanced Board Oversight of Cybersecurity Risk. The Board's engagement on cybersecurity has been heightened and the Board is overseeing multiple enhancements management is making to Capital One's cybersecurity standards, policies, procedures, and processes. This effort is intended to strengthen Capital One's cybersecurity risk management capabilities, including reporting on these risks to the Board and its Committees. The Risk Committee has taken the lead in overseeing this effort, and the full Board has also been actively engaged.
5	Board-management discussion about cyber risk should include identification of which risks to avoid, which to accept, and which to mitigate or transfer through insurance.	The company had a $400M cyber insurance policy [51]. The Risk Committee annually reviews and recommends the Company's information security policy and information security program to the Board for approval. The Risk Committee receives updates from management on its cyber event preparedness efforts and reviews reports from CIO and CISO on significant cyber incidents.	CISO was included into regularly meetings with the board.

Figure 5.1 – Takeaways, changes and implementations after the breach by the CapitalOne's Cybersecurity Leadership

While there is a high demand for cybersecurity skills and companies are eager to hire top talent, it's important to note that weak leadership and a toxic work culture can quickly lead to employee retention issues. This is not just a technological risk, but also a management risk that can significantly impact the crucial operations of an organization. In the aftermath of the incident, CapitalOne not only faced negative consequences for its reputation and stock value, but also made changes to its chief information security officer. However, there were no identified consequences for other compliance, audit, or technology employees. Following are the key takeaways from our case study and analysis -

- **Cloud Configuration Risks:** The breach highlighted the potential exploitation of misconfigured cloud infrastructure. It's crucial to ensure that firewalls, permissions, and access controls are properly set up. For instance,

a similar incident occurred in 2023 when Tesla's misconfigured Kubernetes containers were exploited for crypto jacking.
- **Importance of Least Privilege:** The attacker took advantage of overly permissive IAM roles. Limiting permissions to the essentials can minimize the impact of a breach. In 2024, a healthcare provider faced a comparable issue when excessive permissions allowed an attacker to access confidential patient data.
- **Need for Threat Detection and Response:** Swift detection and response can help mitigate damage. The CapitalOne incident underscored the necessity for real-time monitoring. A financial firm faced a similar situation in 2023, but rapid incident response helped contain an insider threat.
- **Value of Security Audits and Penetration Testing:** Regular audits and testing can uncover vulnerabilities before attackers exploit them. CapitalOne 's case emphasizes this necessity. In 2023, a tech company averted a major breach when a penetration test revealed weaknesses in their API security.
- **Human Element and Insider Threats:** The attacker was a former AWS employee. Organizations should have robust monitoring of employees and ex-employees with access to sensitive systems. In 2024, a similar insider threat incident occurred in the retail sector, highlighting the need for vigilant internal security measures.
- **Comprehensive Incident Response Plans:** Having a well-rehearsed plan can significantly minimize breach damage. In 2023, a manufacturing company effectively contained a breach because of their detailed and regularly practiced incident response plan.

We should understand that insider threats from within organizations can stem from motivations like personal issues or wanting to profit financially or expose security weaknesses. Employees who are unhappy or feel unappreciated can become insider threats, which emphasizes the need for a supportive workplace and open dialogue. To address these risks businesses should think about implementing the following strategies -

- Encourage a nurturing environment by fostering communication and addressing employee issues to reduce negative emotions among the team members.
- Keeping an eye out for any mischievous activities is crucial in detecting possible insider threats early on after an employee leaves the organization.
- Ensure offboarding procedures by promptly revoking access when an employee departs to prevent unauthorized entry.

These lessons underscore the importance of a multi-layered approach to cybersecurity, integrating technical defenses with strong policies and a culture of security awareness. By learning from these incidents, organizations can better prepare for future threats.

Chapter 6 – Useful cloud security frameworks and standards

Following frameworks and cloud security standards would have helped to contain, if not prevent CapitalOne’s data breach, if those were mandated and enforced strictly -

1. Cloud Security Alliance (CSA) - Cloud Controls Matrix (CCM)

The **CSA CCM** provides a detailed set of controls for securing cloud environments. For CapitalOne, the most relevant controls could include:

- **IAM Controls**: The CapitalOne breach was a result of compromised IAM (Identity and Access Management) roles. CSA CCM emphasizes strong IAM controls, ensuring only authorized individuals have access to critical resources.

- **Configuration Management**: CSA CCM stresses the importance of properly configuring cloud environments, including firewalls and other security mechanisms. If CapitalOne had followed this, the misconfigured firewall that led to the breach could have been avoided.

2. Cloud Security Maturity Model (CSM)

The **CSM** framework focuses on maturity levels for cloud security practices. If CapitalOne had implemented CSM, it could have helped by:

- **Risk Assessment**: At higher maturity levels, CSM stresses frequent risk assessments, which would have identified the misconfigured firewall as a risk.
- **Continuous Monitoring**: As an organization grows in cloud security maturity, continuous monitoring of cloud assets becomes a priority. Had CapitalOne adopted this, it might have detected unusual access patterns early on.

3. ISO/IEC 27001 (Information Security Management Systems)

ISO 27001 is a widely recognized standard for managing information security. Implementing these controls could have helped CapitalOne by:

- **Access Control (A.9)**: ISO 27001 emphasizes strong access control, including limiting access to data and resources based on roles. This would have limited the attacker's ability to access sensitive data.
- **Risk Treatment (A.6)**: If CapitalOne had identified the risks posed by cloud misconfigurations, ISO 27001 would require the company to address these risks proactively through corrective measures.

4. NIST (National Institute of Standards and Technology) - Cybersecurity Framework (CSF)

The **NIST CSF** is a comprehensive framework that focuses on five core functions: Identify, Protect, Detect, Respond, and Recover. Controls from NIST that could have helped include:

- **Protect (PR.AC-5 - Network Integrity)**: NIST emphasizes protecting the network by ensuring configurations are correct. Properly configuring their firewall would have prevented the attacker from gaining unauthorized access.
- **Detect (DE.CM-1 - Continuous Monitoring)**: NIST promotes continuous monitoring of systems and assets. CapitalOne could have identified unusual behavior early, stopping the breach before it escalated.

5. GDPR (General Data Protection Regulation)

However, **GDPR** applies primarily to protecting personal data for EU residents, its principles could have helped CapitalOne in the following areas:

- **Data Access Controls (Article 25 - Data Protection by Design and Default)**: If CapitalOne had stronger data protection measures in place, like strict controls on who could access customer data, the impact of the breach would have been minimized.
- **Data Breach Notification (Article 33)**: GDPR requires rapid breach notifications, which might have prompted faster action in detecting and mitigating the breach.

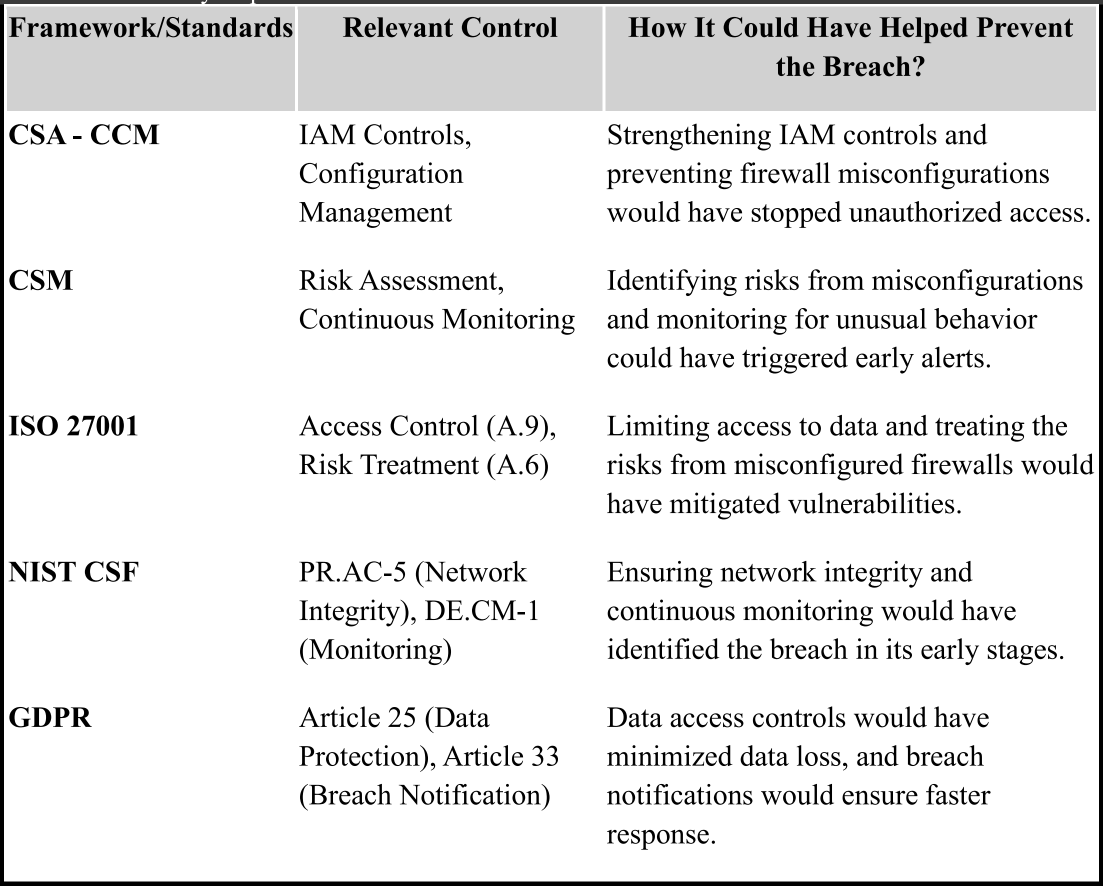

Framework/Standards	Relevant Control	How It Could Have Helped Prevent the Breach?
CSA - CCM	IAM Controls, Configuration Management	Strengthening IAM controls and preventing firewall misconfigurations would have stopped unauthorized access.
CSM	Risk Assessment, Continuous Monitoring	Identifying risks from misconfigurations and monitoring for unusual behavior could have triggered early alerts.
ISO 27001	Access Control (A.9), Risk Treatment (A.6)	Limiting access to data and treating the risks from misconfigured firewalls would have mitigated vulnerabilities.
NIST CSF	PR.AC-5 (Network Integrity), DE.CM-1 (Monitoring)	Ensuring network integrity and continuous monitoring would have identified the breach in its early stages.
GDPR	Article 25 (Data Protection), Article 33 (Breach Notification)	Data access controls would have minimized data loss, and breach notifications would ensure faster response.

Table 6.1 – Relevant Controls which could have helped to prevent the data breach

The study of the CapitalOne incident revealed that the company lacked proper security controls, and the NIST Framework could have helped mitigate the incident if compliance controls were in place. Many companies worldwide struggle to manage security in new cloud computing environments, even when compliance controls and vendor guidance exist. Regulatory agencies should ensure that proper compliance frameworks and regulations are in place to support local companies. In Latin America, the absence of legislation enforcing well-established standards like NIST or ISO frameworks means companies are not required to implement such controls, except when they take the initiative. Local banks in Brazil must comply with cyber security controls enforced through Central Bank Rule 4658 and existing laws like LGPD12, but these standards lack the completeness and comprehensiveness of NIST. It is recommended that companies adopt global governance frameworks with cyber security controls capable of addressing new technologies, considering the impact of local failures on the industry. This highlights the need for a global policy for data protection in an increasingly connected world. [1]

Glossary

Acronym	Definition
AWS	Amazon Web Services: A comprehensive cloud computing platform provided by Amazon.
EC2	Elastic Compute Cloud: A cloud computing platform providing virtual servers to run applications.
FBI	Federal Bureau of Investigation: The domestic intelligence and security service of the United States.
IAM	Identity and Access Management: A framework for managing electronic identities and access rights.
PoC	Proof of Concept: A demonstration to verify that a concept is viable in a real-world scenario.
S3	Simple Storage Service: A service offering scalable object storage through a web service interface.
SSRF	Server-Side Request Forgery: A vulnerability allowing an attacker to make requests on behalf of a server.
STS	Security Token Service: A web service that enables to request temporary, limited-privilege credentials for users.
TOR	The Onion Router: A privacy-focused network facilitating anonymous communication online.
WAF	Web Application Firewall: A firewall that filters and monitors HTTP traffic between a web application and the internet.

References

[1] Shaharyar, Khan, et al. "A Systematic Analysis of the Capital One Data Breach: Critical Lessons Learned." *Proceedings of the ACM on Computer Science*, 11 July 2022, https://dl.acm.org/doi/10.1145/3546068.

[2] Fitzgerald, Maggie. "Capital One Says Hacker Breached Accounts of Over 100 Million." *CNN Business*, 29 July 2019, www.cnn.com/2019/07/29/business/capital-one-data-breach/index.html.

[3] "2019 Cyber Incident Settlement Reached." *Capital One*, 22 Apr. 2022, www.capitalone.com/digital/facts2019/.

[4] "A Case Study of the Capital One Data Breach." *RSA Conference*, 18 May 2021, www.rsaconference.com/library/presentation/usa/2021/a-case-study-of-the-capital-one-data-breach.

[5] "Capital One Data Breach Explained." *YouTube*, uploaded by Kevin Fang, 29 July 2019, youtu.be/r7HV4s-4ksQ.

--- END OF THE DOCUMENT ---