VECTRA
SECURITY THAT THINKS.®


CASE STUDY

# When a SIEM Just Isn't Enough: Squashing a Sophisticated Cyberattack at a Global Healthcare Giant

As cloud footprints grow exponentially, attackers only need a single opening to infiltrate environments and create a means for persistence. This reality is driving an urgent need to look for signs of attacker behavior within one's environment, since existing security tools meant to identify internal threats are proving to be ineffective.

This scenario became a reality for one of the world's leading healthcare organizations as it witnessed an attack on their cloud infrastructure in late 2021. With over 100,000 employees, this organization leverages Amazon Web Services (AWS) to store sensitive medical data, while running a massive analytics workload that leverages the power of cloud compute. This has culminated into a multi-region deployment comprising of thousands of EC2 instances, hundreds of S3 buckets, millions of Lambdas and data warehousing and analytics services—a labyrinthine web subject to billions of actions every day.

Within the first two weeks of deployment, Detect for AWS was quickly put to the test as a would-be attacker made their way into the environment.

Their SIEM does a great job at log aggregation but attempting to configure custom detections within the SIEM for post-exploitation coverage hasn't yielded favorable results. The rules are often bypassed, and investigation of the few alerts that fire, take up considerable time, hindering their ability to respond quickly. This challenge prompted their selection of Vectra Detect for AWS to deliver seamless threat detection within their AWS footprint.

## Organization

Global Healthcare Giant

## Industry

Pharmaceutical

## Challenge

- Trouble configuring custom detections within the SIEM for post-exploitation coverage
- The rules are often bypassed, and investigation of the few alerts that fire, take up considerable time, hindering the team's ability to respond quickly

## Results

- While monitoring over 1 million identities, users and services, within the environment, Detect for AWS spotted behavior closely resembling an attacker probing the footprint for weaknesses, and their SIEM did not.
- Detect for AWS then observed the malicious principal attempting to disable security tools within the environment as a means to establish persistence

Within the first two weeks of deployment, Detect for AWS was quickly put to the test as a would-be attacker made their way into the environment. The actor began by stealing credentials, conducting heavy reconnaissance activities, and attempting to escalate privileges to access sensitive resources in the environment. What really stood out to the SOC, was the bulk of the actions taken by this principal were not bubbled up by the rules configured in the incumbent SIEM. Fortunately, Detect was ready.

## Attack Caught Red-Handed using Stolen Credentials

During its second monitoring of the organization's footprint, Detect spotted behavior closely resembling an attacker probing the footprint for weaknesses. At the time, Detect was monitoring over 1 million identities, users and services, within the environment.

## No Escaping Vectra

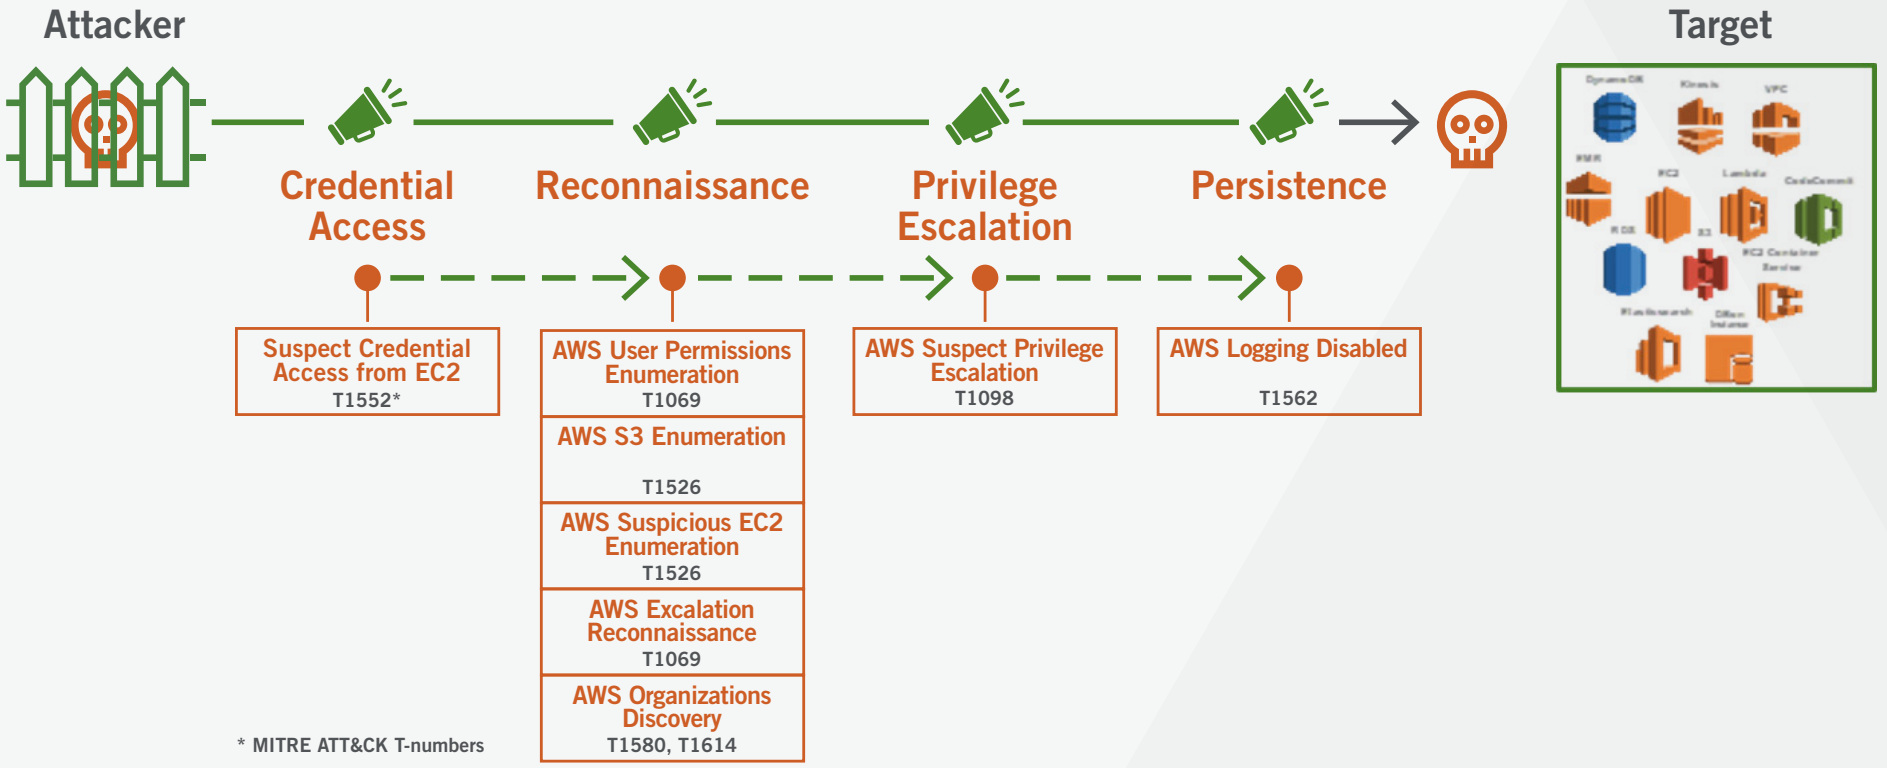

Detect first triggered an alert on an attempt by the attacker to gather EC2 instance metadata that would enable them to infiltrate the environment and further their attack. Following the usage of these compromised credentials to access the footprint, Detect fired on numerous reconnaissance activities being performed by the attacker. This included enumeration of permissions, reconnaissance of resources that the compromised account had access to, as well as discovery tactics to understand how the organization's account architecture was configured.

Detect then observed the malicious principal attempting to disable security tools within the environment as a means to establish persistence. There were also several attempts to use privilege escalation techniques within the environment—presumably to reach high-value services. The principal had nowhere to hide as they forayed through the cloud kill chain.

## How Did the SOC Respond?

As a result of the behaviors spotted against the malicious principal, Detect had elevated the principal's account ensuring that it immediately attracted attention by the SOC team. The account was one of only seven prioritized in the Detect console. To add some perspective, that's seven out of over six hundred accounts Detect was tracking activity against!

Using the Principal's account page on the Detect console, the SOC team had a view of all the behaviors that triggered detections. The team was also able to gather key insights surrounding the principal's activities outside the behaviors that triggered detections, through the Instant Investigation feature. These included:

- **Regions where the principal was active:** The SOC noticed the compromised principal attempting actions in seven different AWS regions, which is highly unusual for an account in their environment.
- **Services the principal interacted with, and actions taken:** Within each region, the SOC had visibility into the different AWS services this principal interacted with. They noticed attempts to spin up EC2 instances, probe IAM permissions, access S3 buckets and noted attempts to probe the KMS service.
- **Roles assumed by the principal:** A list of all roles assumed by the malicious principal across the environment was also provided to the SOC. Fortunately, the principal did not have access to assume any additional roles within the environment.

As soon as the principal was elevated, the SOC team was able relegate permissions, quarantine the account, and rotate credentials ensuring the attack was stopped in its tracks.

**For more information please contact us at info@vectra.ai.**

Email info@vectra.ai  vectra.ai