LogicHub

Intelligent Security Automation Use Case

AUTOMATING THREAT HUNTING IN AWS CLOUDTRAIL LOGS

Introduction:

LogicHub is the leading Intelligent Security Automation Platform that marries a powerful Decision Engine to a flexible Workflow Engine. Proven to deliver 10x the performance of traditional Security Orchestration Automation and Response (SOAR) solutions, it is the only platform of its kind to deliver analysis and decision-making automation to exponentially improve alert triage, incident response, and threat hunting.

As more companies are moving out of physical data centers and into cloud based solutions, analysts need to develop new ways to analyze their cloud based solutions for risks and threats. LogicHub has developed a playbook to hunt for risks in one such solution, AWS CloudTrail logs. This conducts seven investigations in parallel to identify risks within the CloudTrail logs.

CloudTrail Threat Hunting Challenges:

- Thousands of log entries per day to analyze
- Human analysts are limited in both quantity and availability
- It would take hours to search through a day's worth of data
- No easy way to retain learned threat intelligence and improve institutional knowledge

LogicHub Solution:

Automated threat hunting of AWS CloudTrail logs with LogicHub is a powerful and easy method to kick off your threat hunting campaigns by focusing on a smaller subset of important events. LogicHub is capable of reducing the noise in the data by identifying smaller subset of riskier entries.

Our Approach

Our approach breaks up the investigation into 7 parallel investigations that can output the results into two high level buckets (bad/malicious or needs further investigation) based on a

scoring model. An event scoring 10 is in the "bad/malicious bucket", and a score of 1-9 is in the "needs further investigation" bucket where the higher the score the more likely the event is malicious. Lower scoring events can be filtered out to further avoid alert fatigue.

Our investigation "branches" are as follows:

- Public S3 Buckets Branch
- S3 Bucket Downloads Branch
- Launch Instances Branch
- Open Security Group Access Branch
- Modification to IAM Branch
- Impossible Travel Branch
- Authentication Failures Branch (Brute Force Detection)

Public S3 Buckets Branch

In the last half of 2017 there were several newsworthy breaches involving companies leaving sensitive data in S3 Buckets misconfigured to allow unrestricted public access. Due to these misconfigurations, tens of millions of records were exposed.

To combat this, our approach analyzes S3 Buckets on three factors:

- Bucket sensitivity; a one-time score assigned to the bucket
- Whether the user who created or modified the bucket logged in within an approved location (typically country based)
- Whether the user who created or modified the bucket used MFA to login (as applicable)

Based on these factors, the outputted results can be prioritized as well as provide additional information needed to remediate the event if necessary.

S3 Bucket Downloads Branch

Continuing with the S3 Bucket woes, the second branch in our solution analyzes downloads from your S3 Bucket. For this, our approach analyzes S3 Buckets on two factors:

- File Size
- Bucket Sensitivity; a one-time score assigned to the bucket

Using these factors, events are prioritized based on a learning module, trained by your analyst to predict the final scores based off of the file size and bucket sensitivity allowing the outputted results to be prioritized. Added to every output is additional information needed to remediate the event if necessary.

If required, more analysis for this branch can be added. Whether it's a scored factor or a data point outputted such as location of user or keywords in the filename (if filename contains SSN, PAN, etc..), or any number of customizations to zero in on important events.

Launch Instances Branch

Running computationally powerful instances can be a costly endeavor as well as motivation for an attacker. In a financially motivated attacker were to make their way into your AWS environment, one possible lucrative post exploitation effort would be to set up an instance for bitcoin mining. For this reason, it important to ensure that the instances being launched in your AWS environment are authorized and scoped appropriately.

This branch scores events by:

- Machine power based on CPU and Memory allocated to the machine
- Whether the user who created or modified the bucket logged in within an approved location (typically country based)
- Whether the user who created or modified the bucket used MFA to login (as applicable)

Using these factors, analysts will be alerted to large instances being deployed or any instance launched from a likely compromised user account.

Open Security Group Access Branch

Similar to the concern surrounding unrestricted public access to S3 Buckets, companies need to ensure non-public AWS instances are not exposed to the public internet. In a larger company, as many different resources from many different teams have access to launch instances and configure security groups, it can be a bit of a cat and mouse game trying to ensure these instances stay private.

This branch monitors CloudTrail logs for updates to security groups that enable public access and alerts to the event. As with all branches, this branch can be calibrated to filter these alerts, for example, only alerting on a certain group of instances or only under more specific circumstances.

Modification to IAM Branch

During post-exploitation, one common goal is to either elevate the privileges of the compromised account or create a backdoor account with elevated privileges. This can also be a malicious insider attempting to gain more privileges outside the authorized processes. This branch alerts for modifications made to a user's permissions and ranks risk on the following factors:

- Modification type, risked ranked by the analysts
- Whether the instance was launched by a user located in a whitelisted (or not in a blacklisted) country
- Whether MFA was used

Using these criteria, the resulting prioritized list can lead analysts towards suspiciously modified accounts.

Impossible Travel Branch

Impossible travel is when two events from the same user occur in different locations but in a shorter timeframe than the distance between the events could be traveled. This branch identifies and analyzes the distance to time between events. Events are then ranked based on the speed needed in order to travel the distance between the two location in the amount of time given. For example, if the speed needed to cover the distances is below 45 mph, then the event is scored as 2. However, if the speed needed to cover the distances is greater than 550m mph (faster than the common speed of an airplane) the travel is deemed impossible and rank as a 10.

The resulting output filters out the non-suspicious events as well as shorter distances, in order to account for false positives.

Authentication Failures Branch (Brute Force Detection)

In our last branch, our flow monitors for excessive failed login attempts in order to identify potential attackers attempting to brute force their way in. These events are scored and ranked by number of failed attempts and whether there was a successful attempt made shortly after the excessive failed events. The latter indicating that the account is potentially compromised.

Additional Branches

As well as each branch being customizable to narrow in on suspicious events, more branches can be added as more use cases are identified. If a new attack or new concern with misconfigurations arise, new branches can be built to identify, validate and risk rank events to be added to your final output.

Conclusion

Automating threat hunting AWS CloudTrail logs with LogicHub is powerful, easy, and can help you detect attackers and threats otherwise easily missed in the mountain of data. SOC teams are able to improve their productivity and response times, while minimizing false positives and false negatives.

For more security automation use cases, visit www.logichub.com.