SEC301

Automating threat detection and response in AWS

Ross Warren
Security Specialist, AWS WWCS Geo Solution Architect
Amazon Web Services

Brandon Baxter
Solution Architect, ESS Security Specialist

AWS re:Invent

Workshop agenda

Use US West (Oregon)
us-west-2

Module 1: Environment setup (20 min.)

Module 2: Attack kickoff and presentation (40 min.)

Module 3: Detect, investigate, and respond (45 min.)

Module 4: Review, questions, and lessons learned (15 min.)

Don’t leave after Module 4 – Hands on with Amazon Detective

Please follow directions

AWS event engine

We will use the AWS event engine; please don't use one of your own accounts

https://dashboard.eventengine.run

You will enter the hash that is provided at the URL above

Event name: Automating threat detection and response in AWS

1. Click **AWS Console**

2. Click **Open Console**

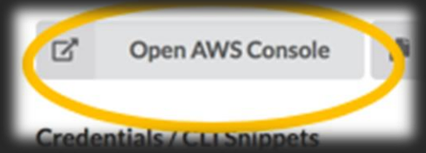

3. Verify that you are in the AWS Management Console in the **us-west-2** Region

Workshop scenario

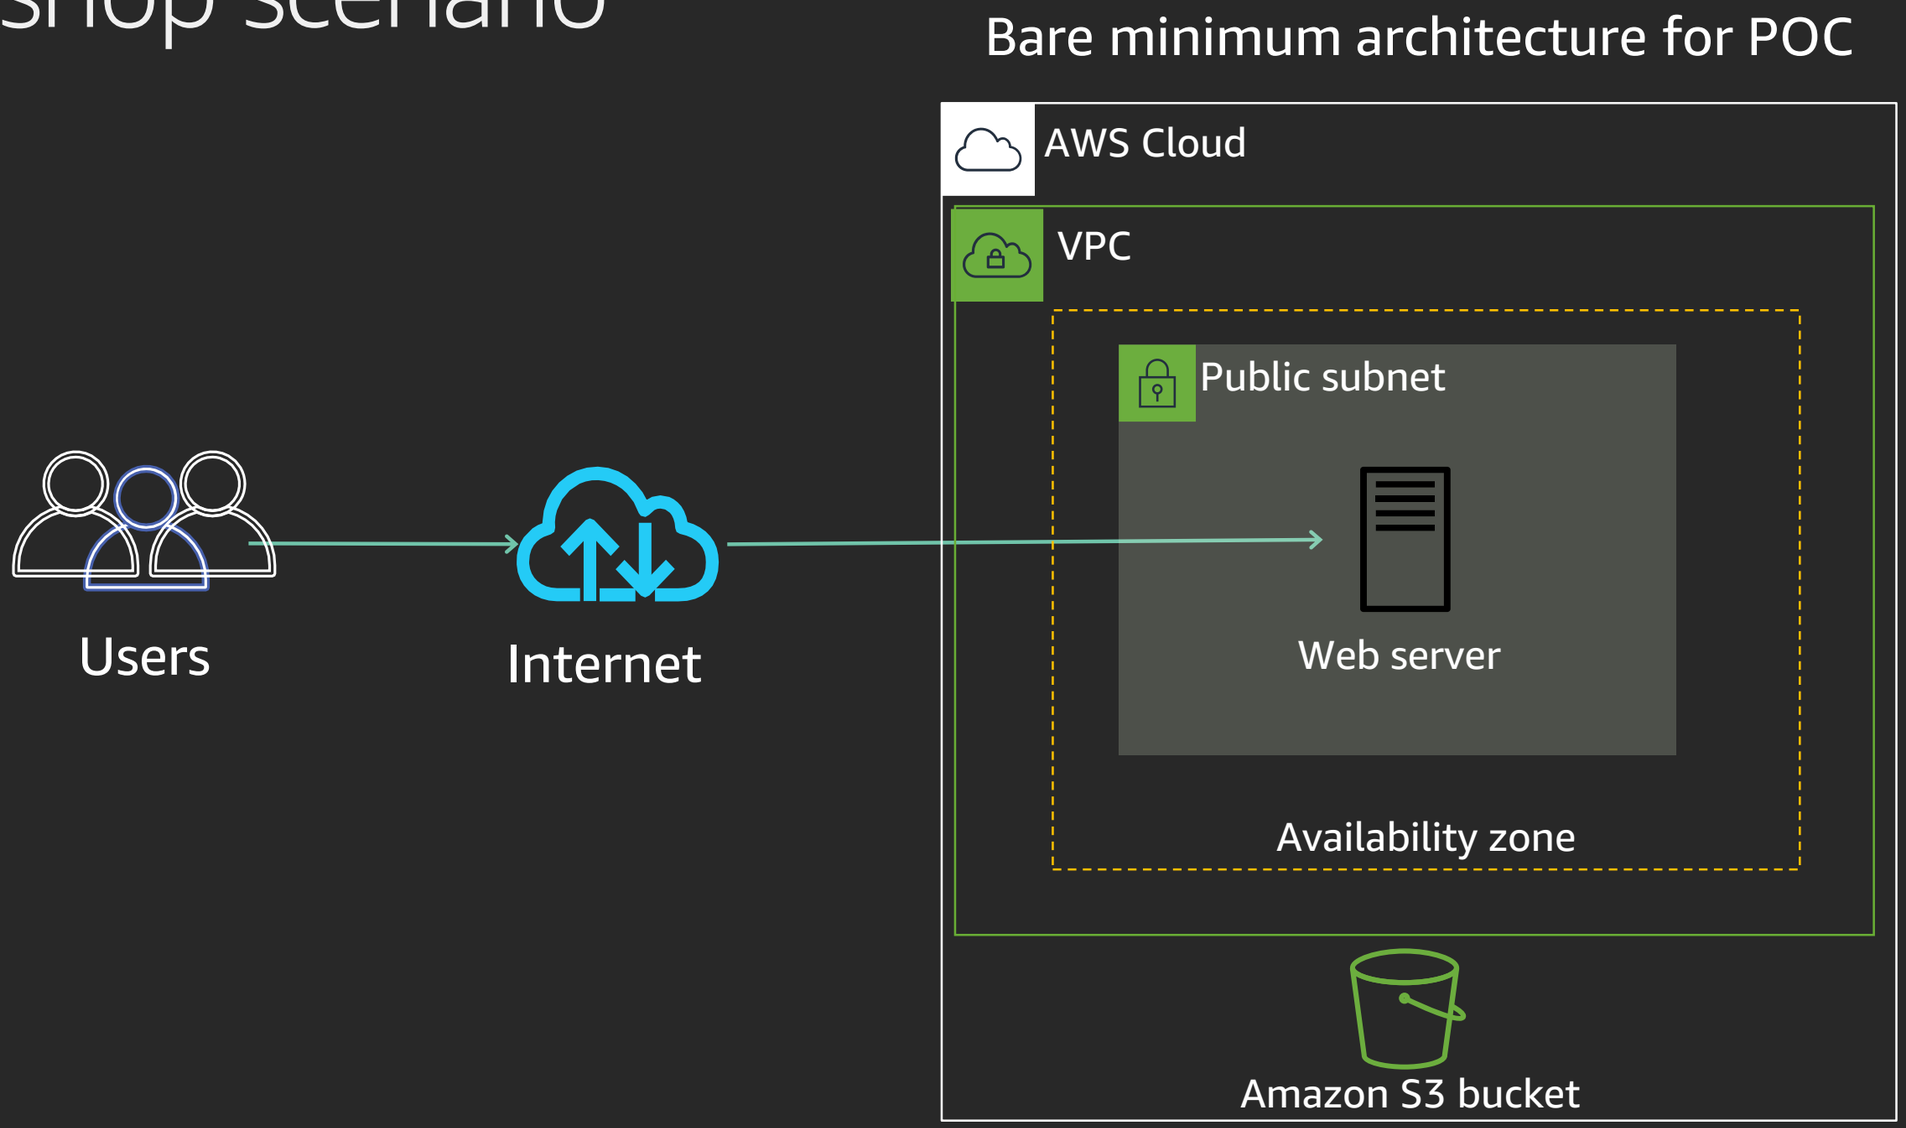

Module 1: Build detective controls

Run the AWS CloudFormation template (~5 min.)

- Before moving on, make sure the stack status = CREATE_COMPLETE
- You will get an email from Amazon SNS asking you to confirm the subscription; do this so that you can receive email alerts from AWS services during the workshop

Manual setup steps (~15 min.)

- Create a Amazon CloudWatch Event Rule
- Enable Amazon GuardDuty
- Enable AWS Security Hub

Module 1: Build detective controls

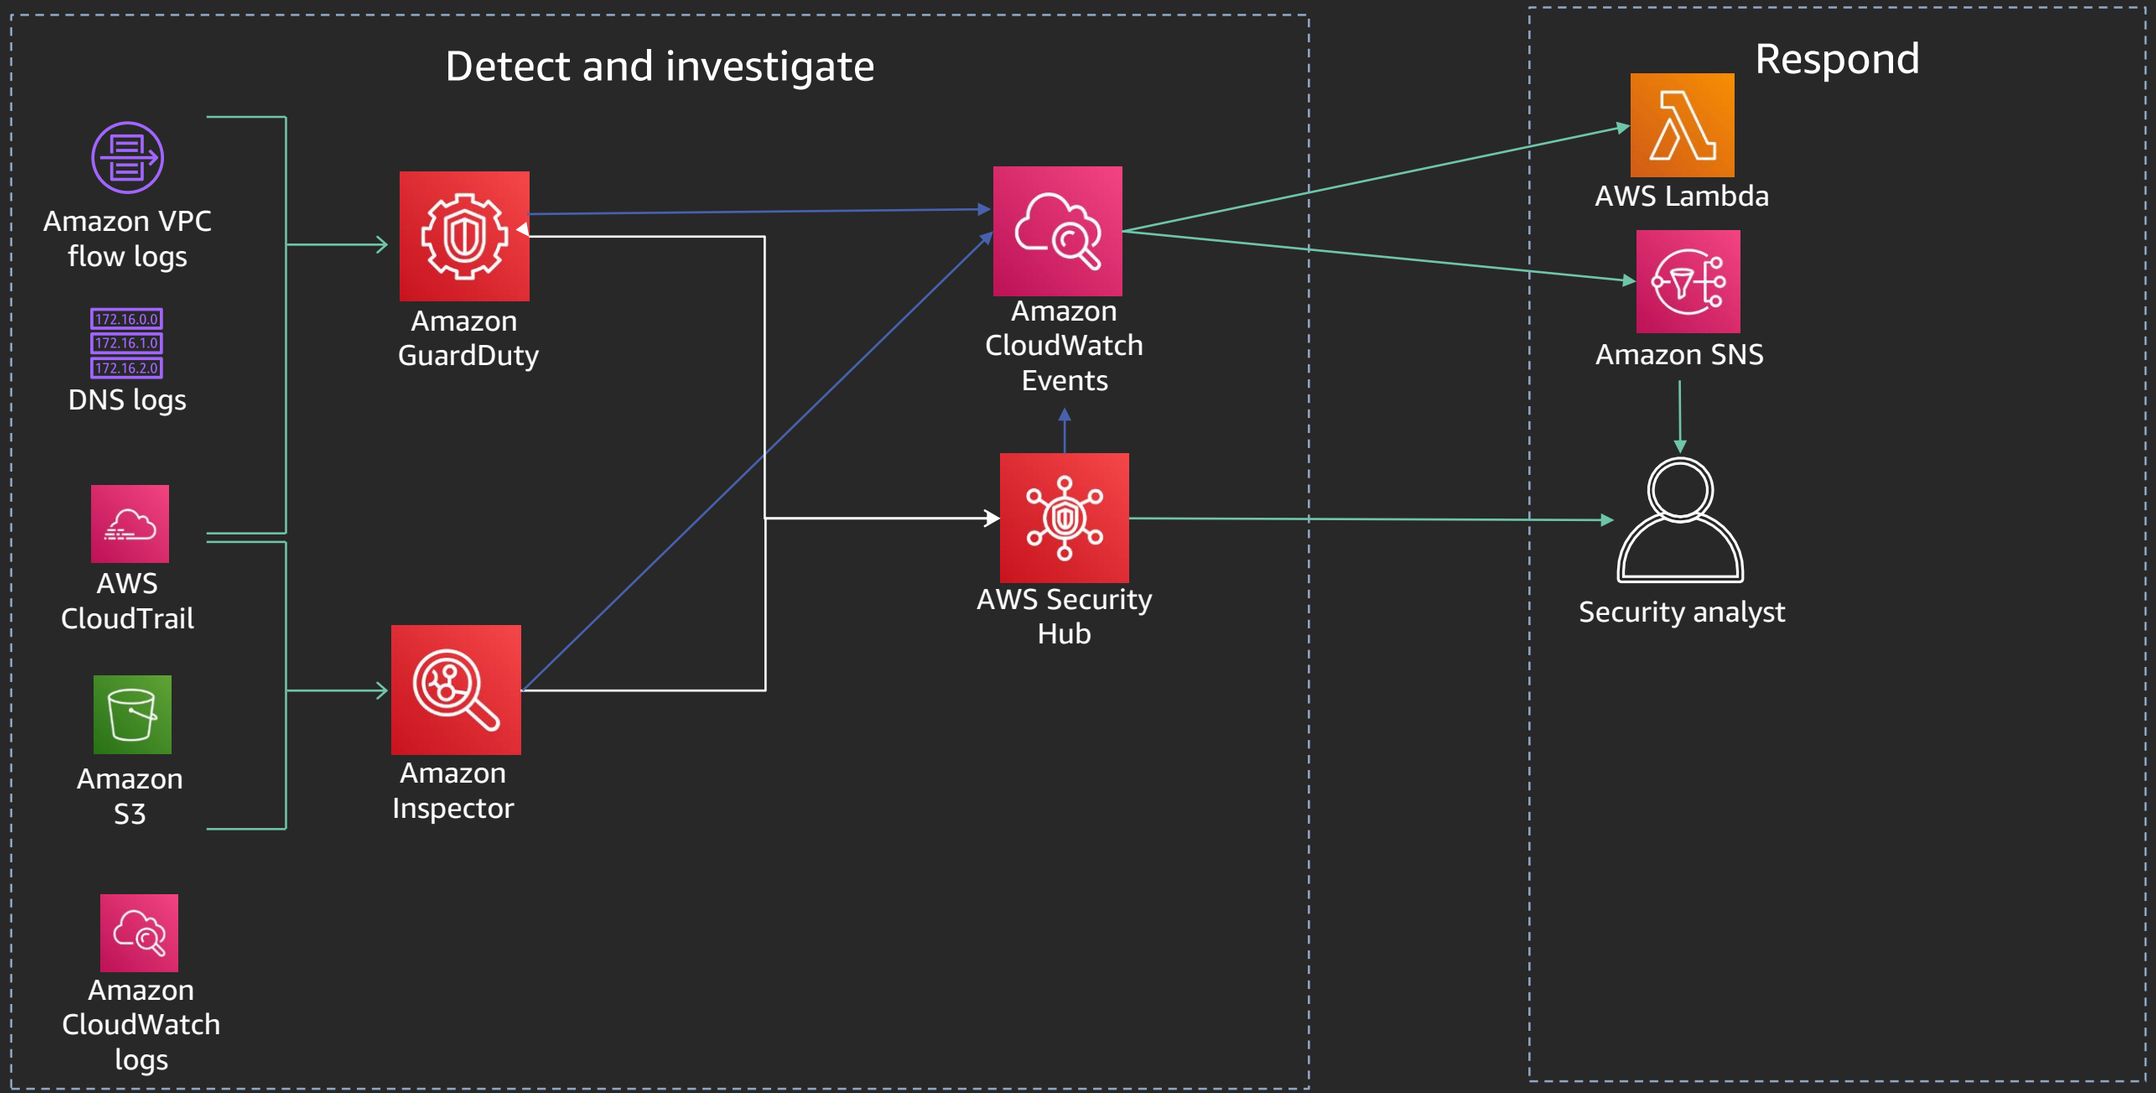

Module 1: Build detective controls

Use US West (Oregon)
us-west-2

https://dashboard.eventengine.run

https://automating-threat-detection.awssecworkshops.com/

Directions

1. Browse to the URL
2. Read through the workshop scenario
3. Choose **Module 1: Environment Build** in the outline on the left
4. Complete the module (~15 min.), and then stop

Please follow directions

Do not skip the manual steps

Module 2: Attack kickoff

Run the AWS CloudFormation template (~5 min.)

- Do not move on to Module 3

Threat detection and response presentation (~30 min.)

Workshop walk-through (~5 min.)

Module 2: Attack kickoff

Use US West (Oregon)
us-west-2

https://dashboard.eventengine.run

https://automating-threat-detection.awssecworkshops.com/

Directions

1. Browse to the URL
2. Choose **Module 2: Attack Simulation** in the outline on the left
3. Complete the module (~5 min.), and then stop
4. Do not move on to Module 3

Threat detection and response: Introduction

AWS re:Invent

Why is threat detection so hard?

Large datasets

Signal to noise

Skills shortage

Deep set of security tools

Identity

AWS Identity and Access Management (IAM)

AWS Single Sign-On

AWS Directory Service

Amazon Cognito

AWS Organizations

AWS Secrets Manager

AWS Resource Access Manager

Detect

AWS Security Hub

Amazon GuardDuty

AWS Config

AWS CloudTrail

Amazon CloudWatch

Amazon Virtual Private Cloud (Amazon VPC) flow logs

Amazon Detective

Infrastructure protection

AWS Systems Manager

AWS Shield

AWS WAF (web application firewall)

AWS Firewall Manager

Amazon Inspector

Amazon VPC

Data protection

AWS Key Management Service (KMS)

AWS CloudHSM

AWS Certificate Manager

Amazon Macie

Server-side encryption

Respond

AWS Config rules

AWS Lambda

AWS Systems Manager

AWS threat detection services

AWS re:Invent

Threat detection: Log data inputs

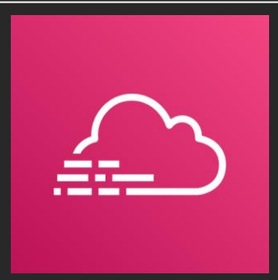

AWS CloudTrail

Track user activity and API usage

Flow logs

IP traffic to and from network interfaces in a VPC

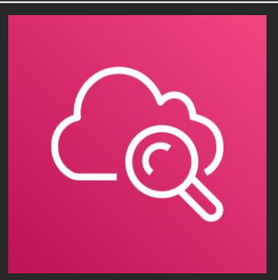

Amazon CloudWatch

Monitor applications using log data; store and access log files

DNS logs

Log of DNS queries in a VPC when using the VPC DNS resolver

Threat detection: Amazon GuardDuty

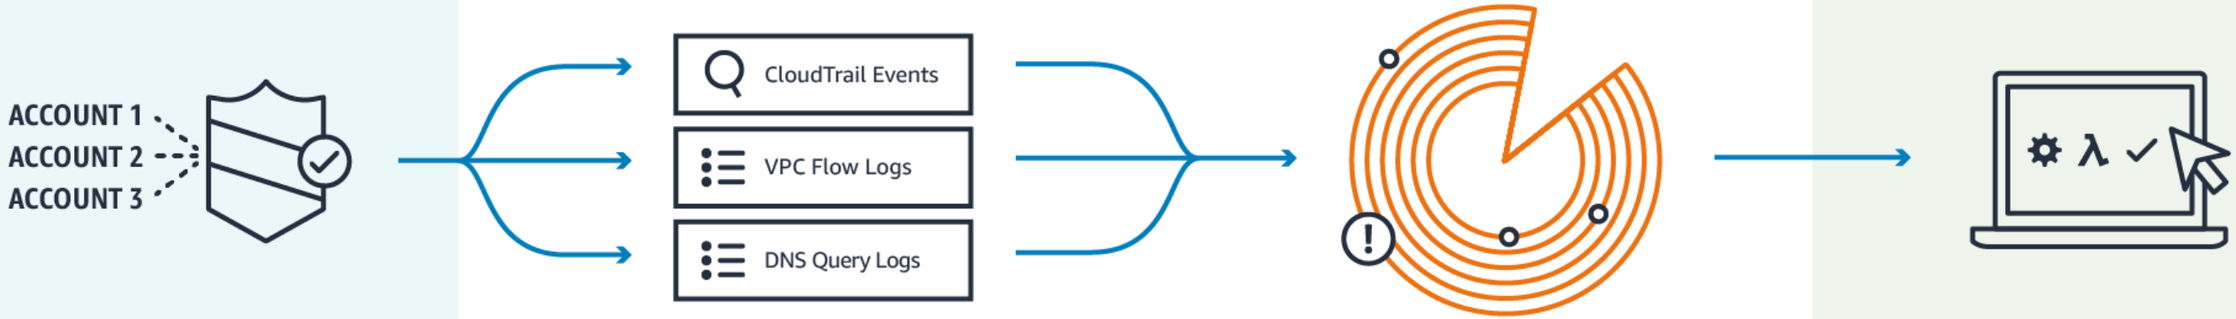

Enable GuardDuty

With a few clicks in the console, monitor your AWS accounts without additional security software or infrastructure to deploy or manage

Continuously analyze

Automatically analyze network and account activity at scale providing broad, continuous monitoring of your AWS accounts and workloads

Intelligently detect threats

Utilize managed rule-sets, integrated threat intelligence, anomaly detection, and machine learning to intelligently detect malicious or unauthorized behavior

Leverage actionable alerts

Review detailed findings in the console, integrate into event management or workflow systems, or trigger AWS Lambda for automated remediation or prevention

Threat detection: AWS Security Hub

Insights & standards

- Comprehensive view of your security and compliance state within AWS
- Aggregates security findings generated by other AWS security services and partners
- Analyze security trends and identify the highest-priority security issues

AWS Security Hub

Findings

Security findings providers

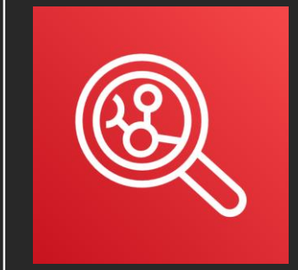

Amazon Inspector

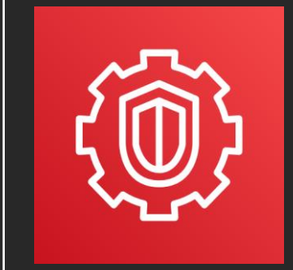

Amazon GuardDuty

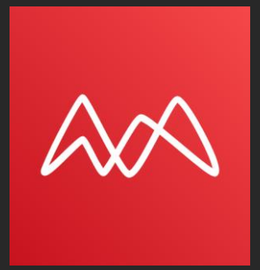

Amazon Macie

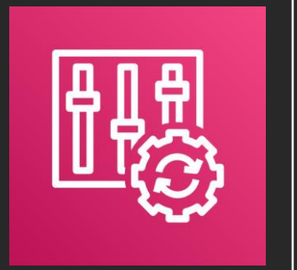

AWS Config

Partner solutions

Other

Threat investigation: Amazon Detective

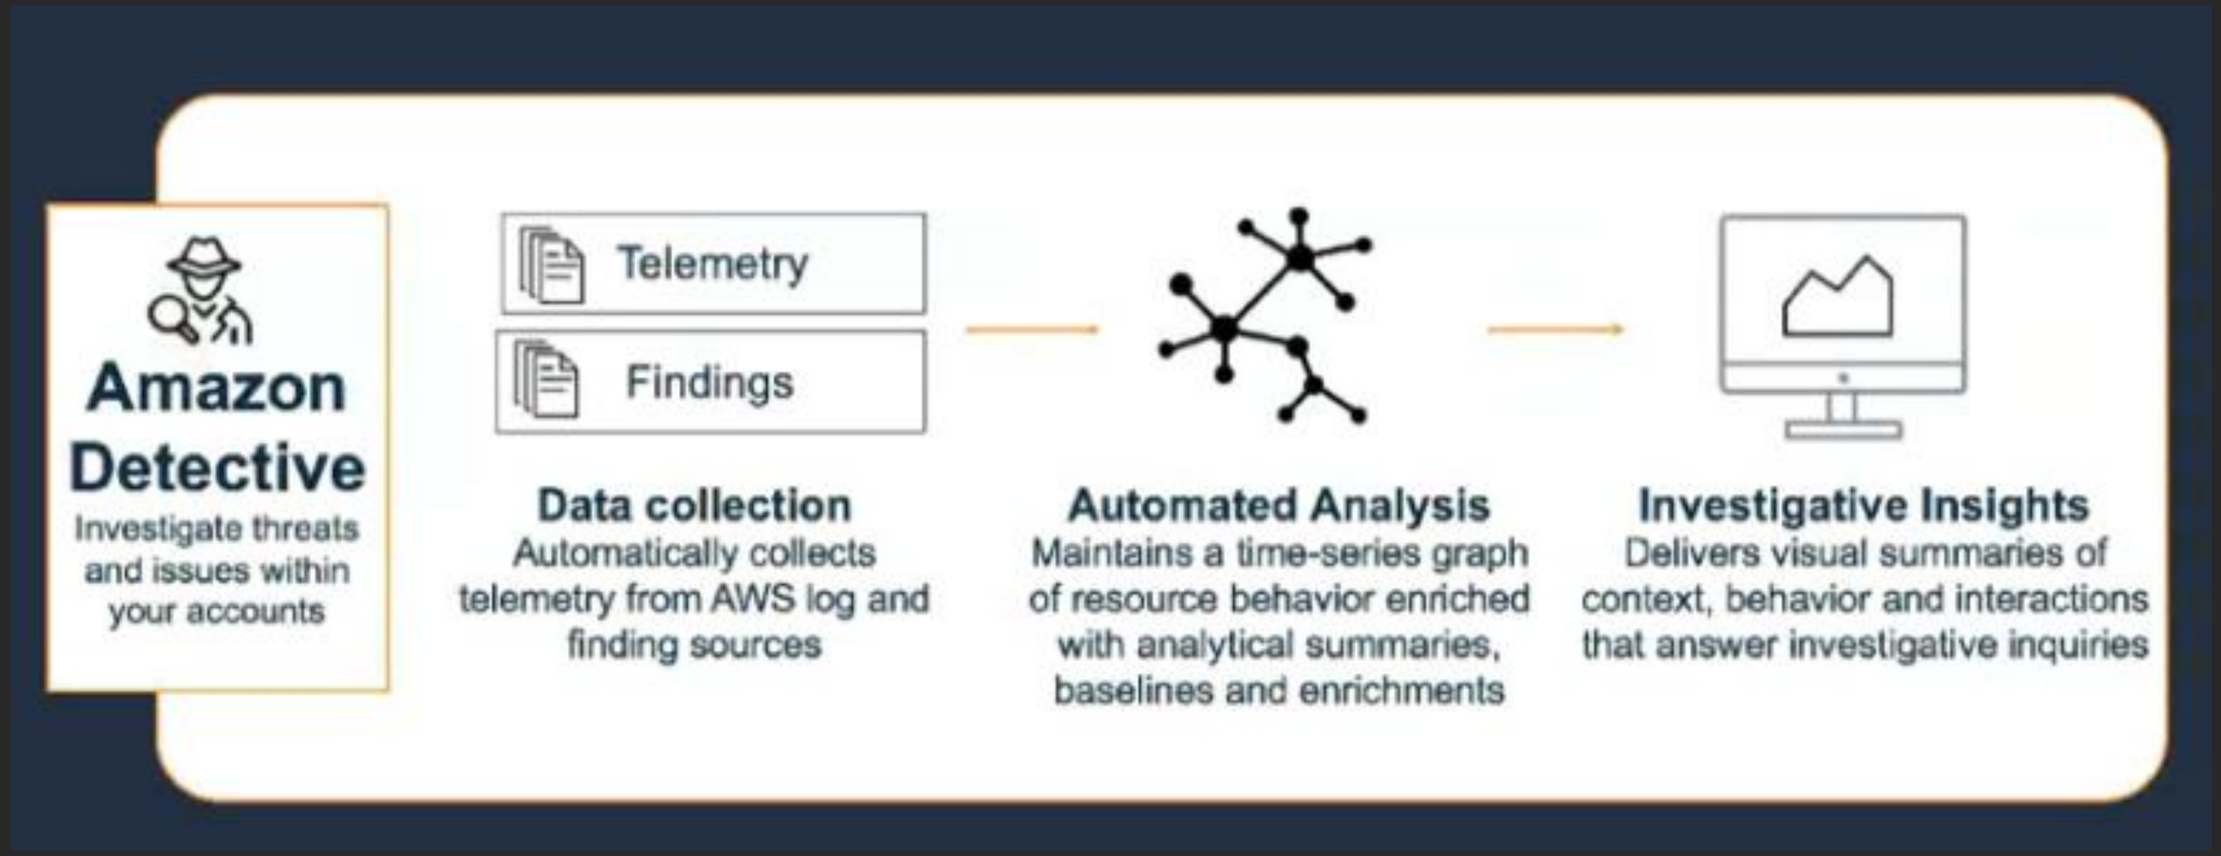

Attacker lifecycle: Stages

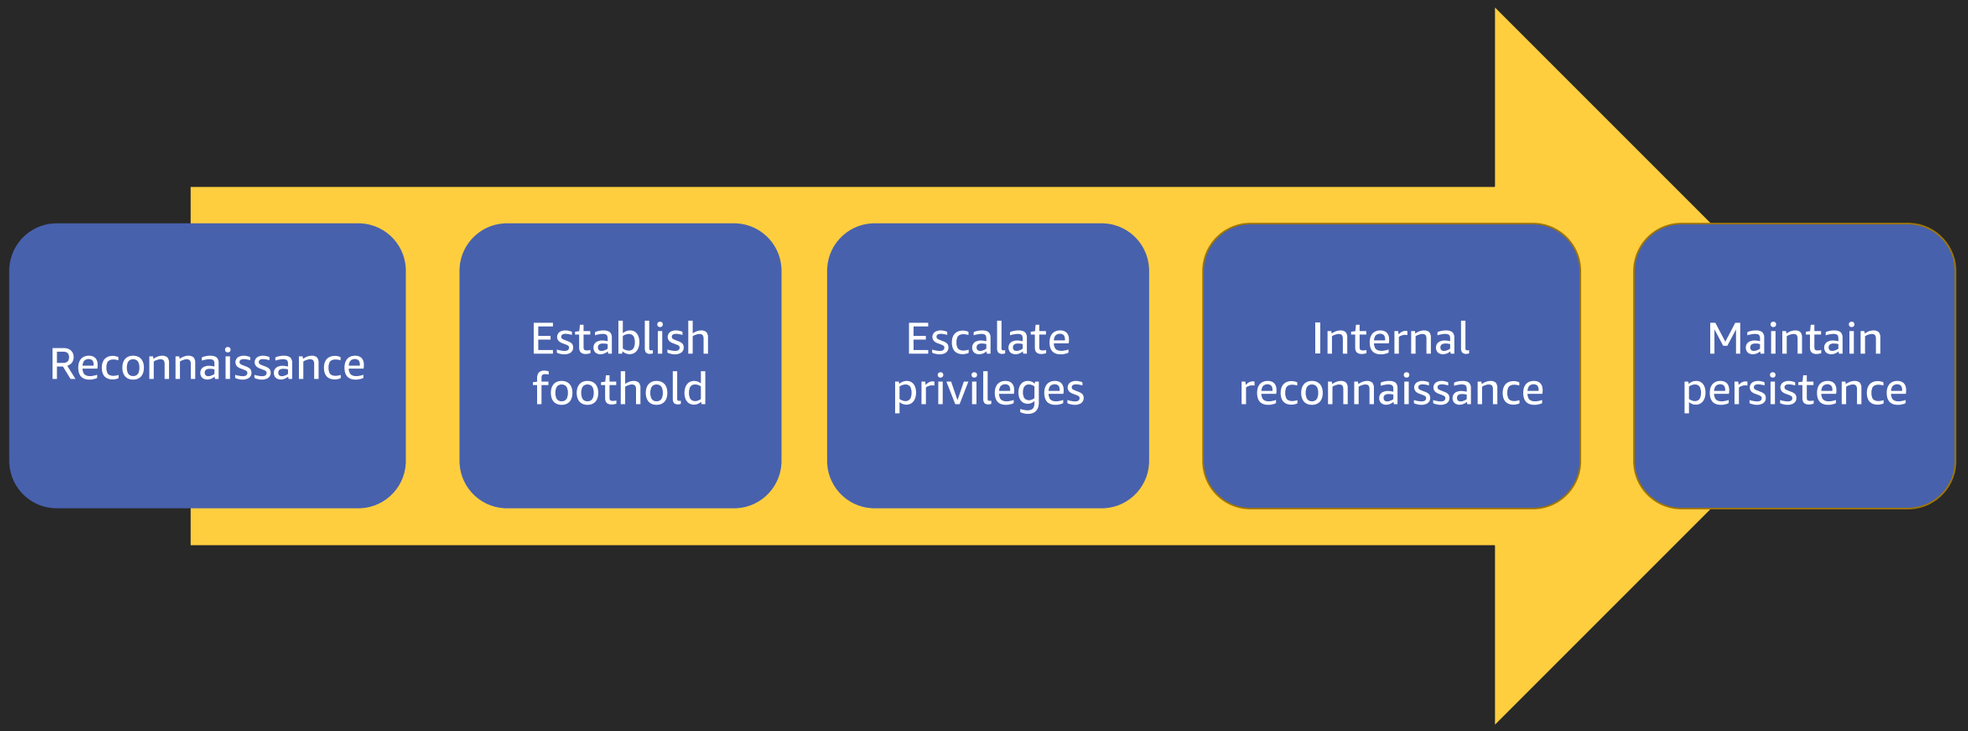

Attacker lifecycle: Attacker actions

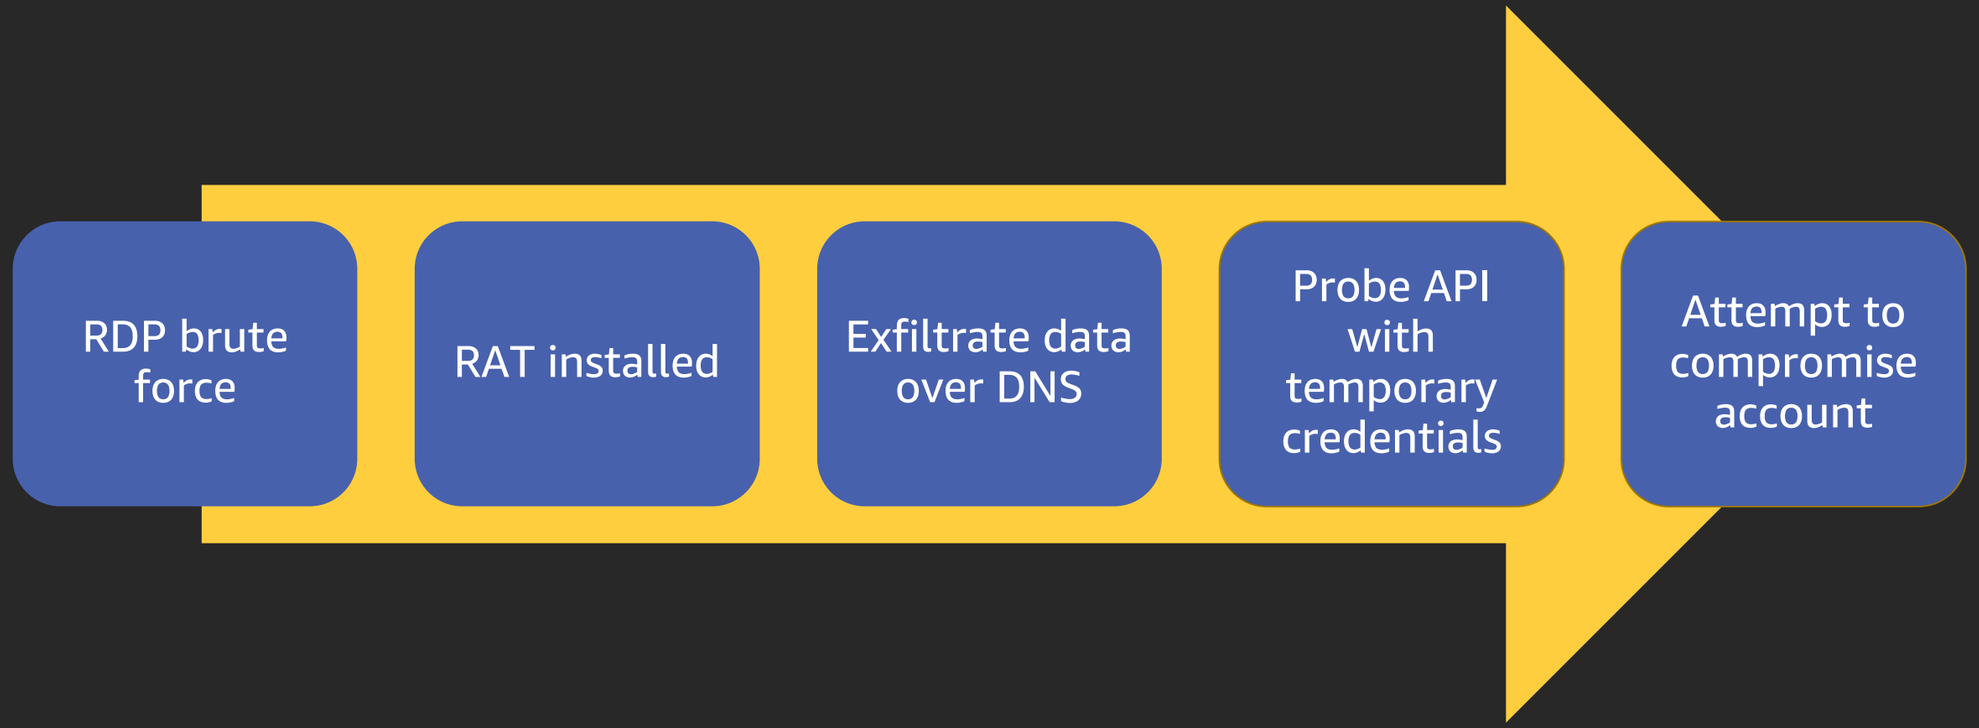

Attacker lifecycle: Amazon GuardDuty findings

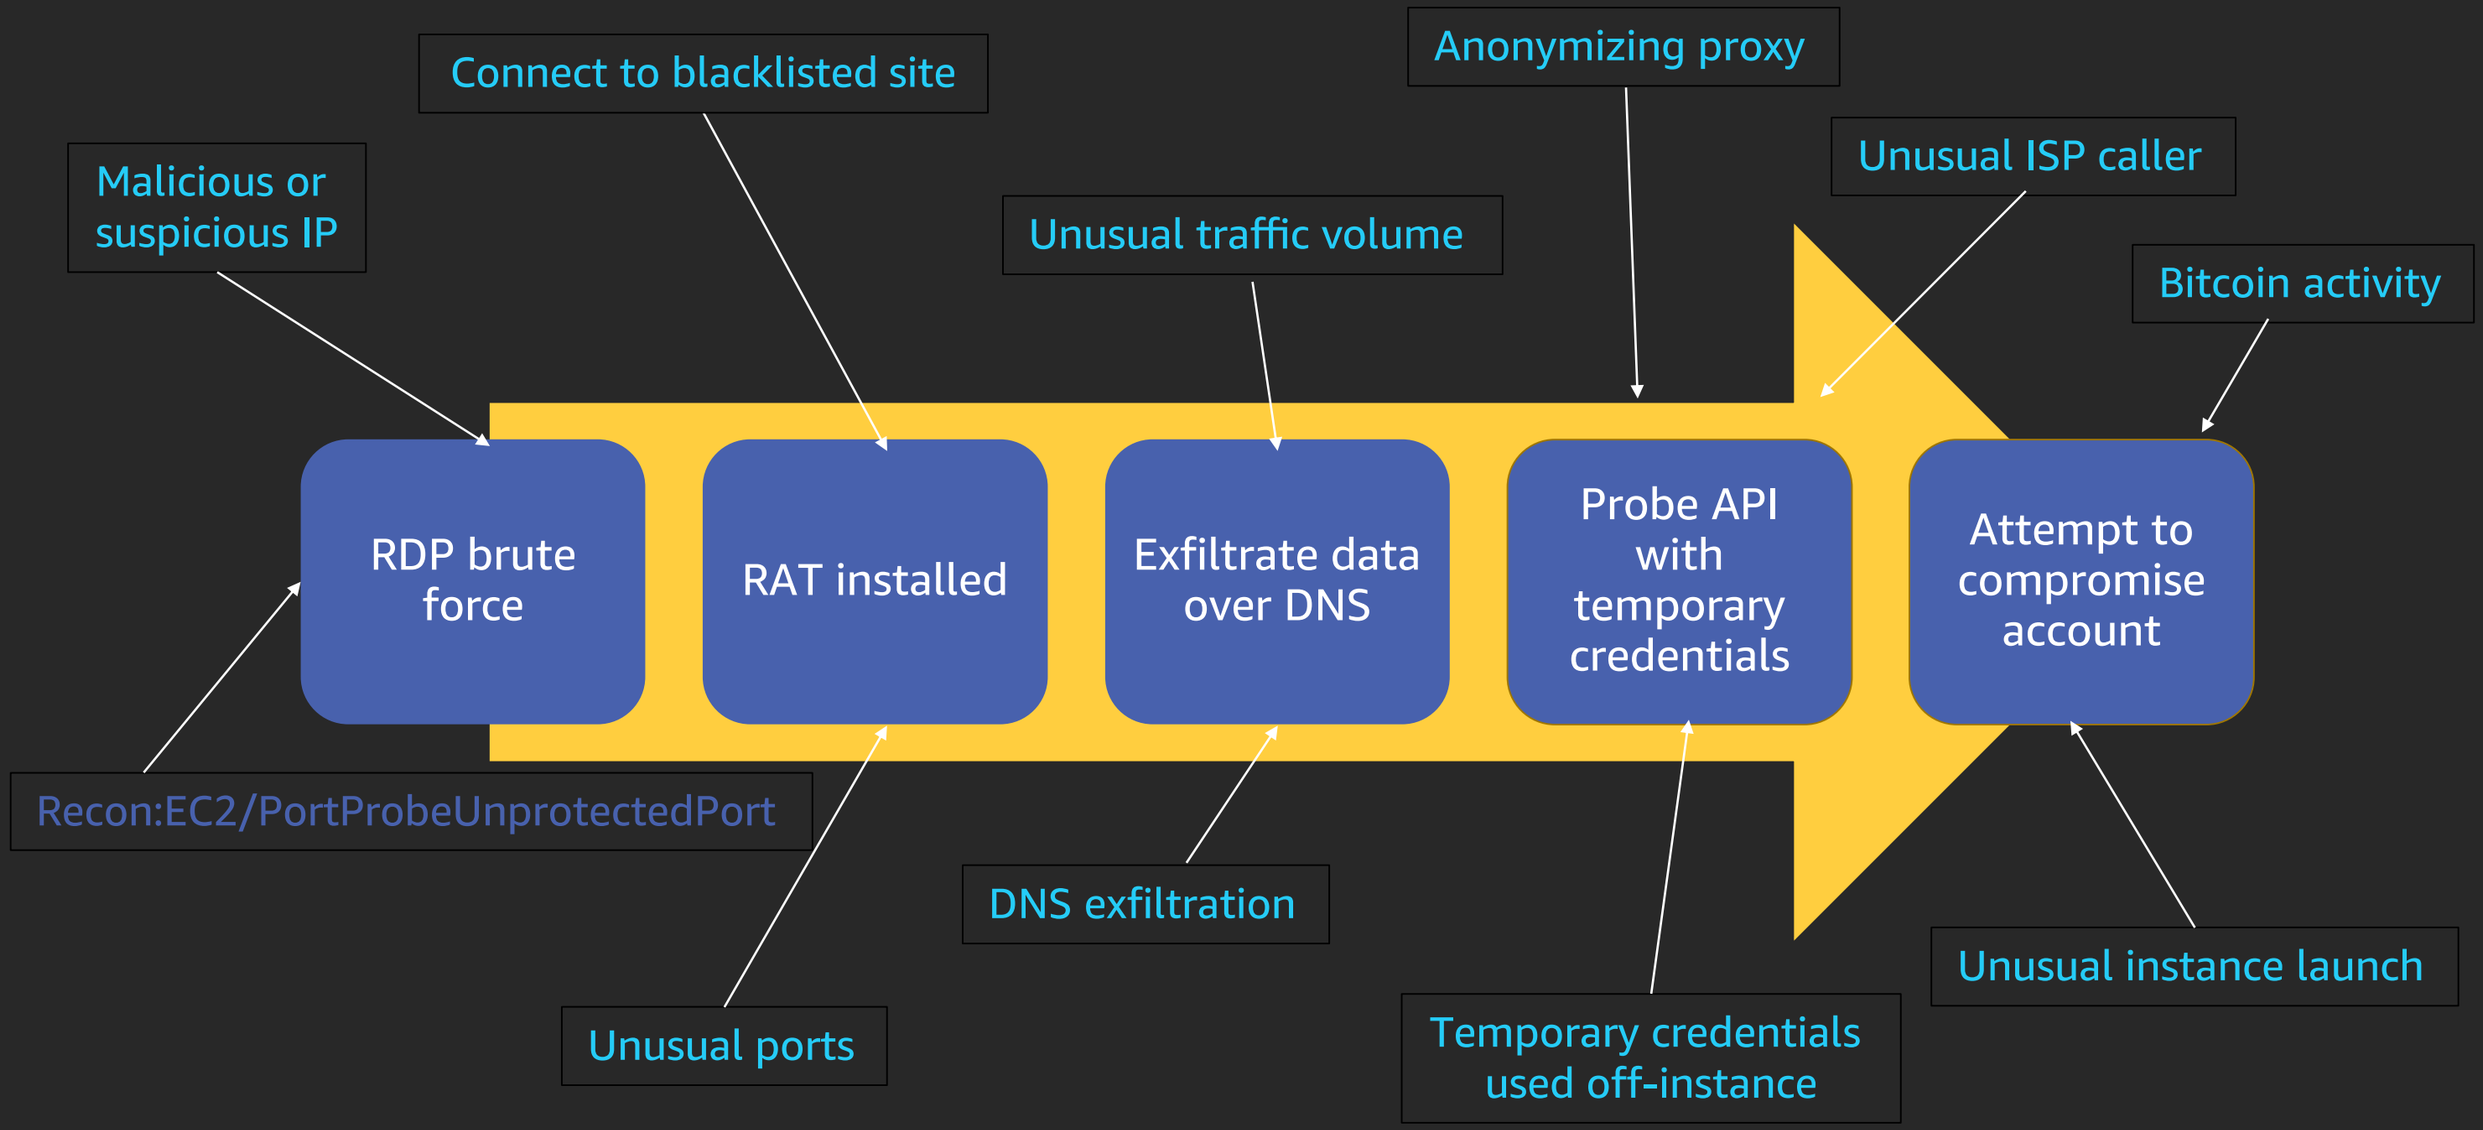

Respond

AWS re:Invent

Threat response: Amazon CloudWatch Events

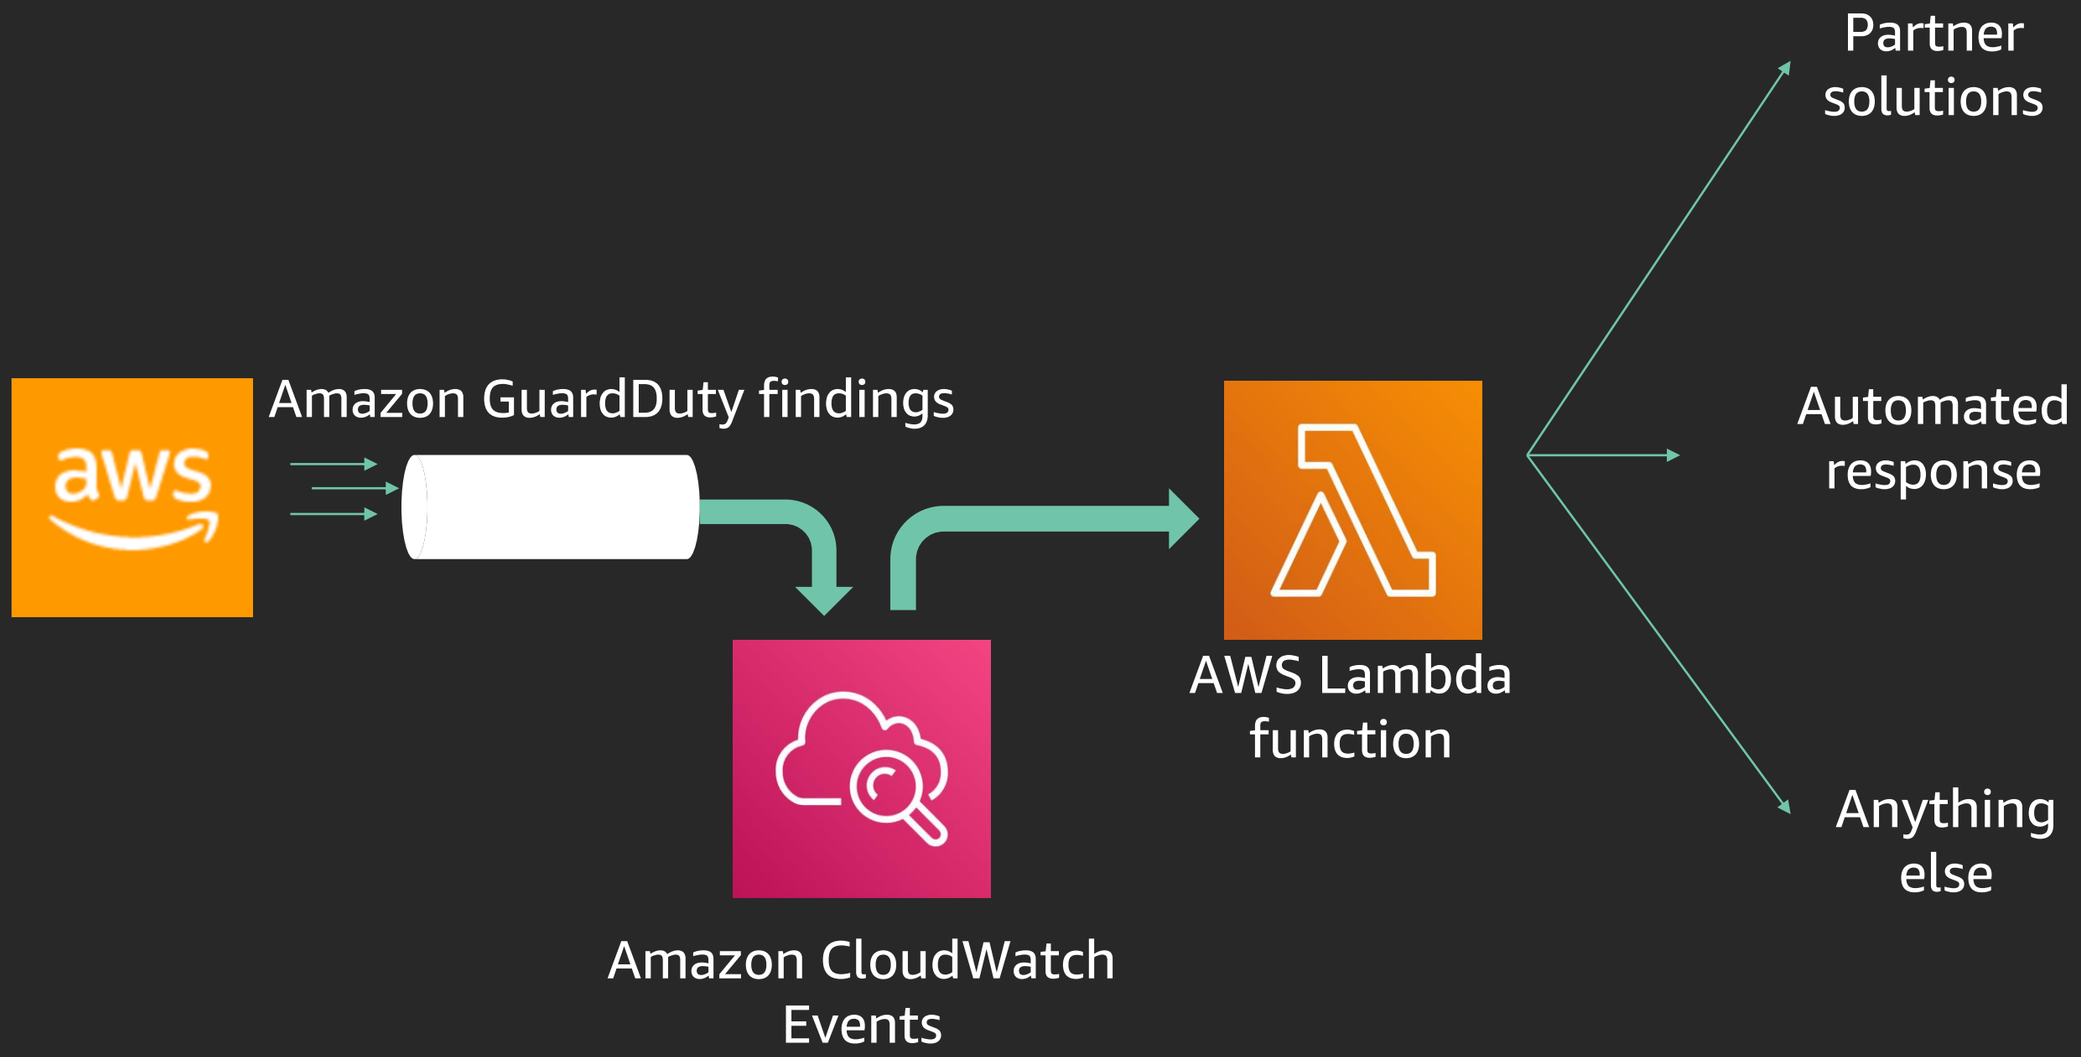

Threat response: Services

AWS Lambda

Run code for virtually any kind of application or backend service—zero administration

AWS Systems Manager

Gain operational insights and take action on AWS resources

Amazon Inspector

Automate security assessments of Amazon EC2 instances

Security response automation on AWS

Monitor: Your automated monitoring tools collect information about resources and applications running in your AWS environment.

Detect: When a monitoring tool detects a predefined condition—such as a breached threshold, anomalous activity, or configuration deviation—it raises a flag within the system.

Respond: When a condition is flagged, an automated response is triggered that performs an action you've predefined—something intended to remediate or mitigate the flagged condition.

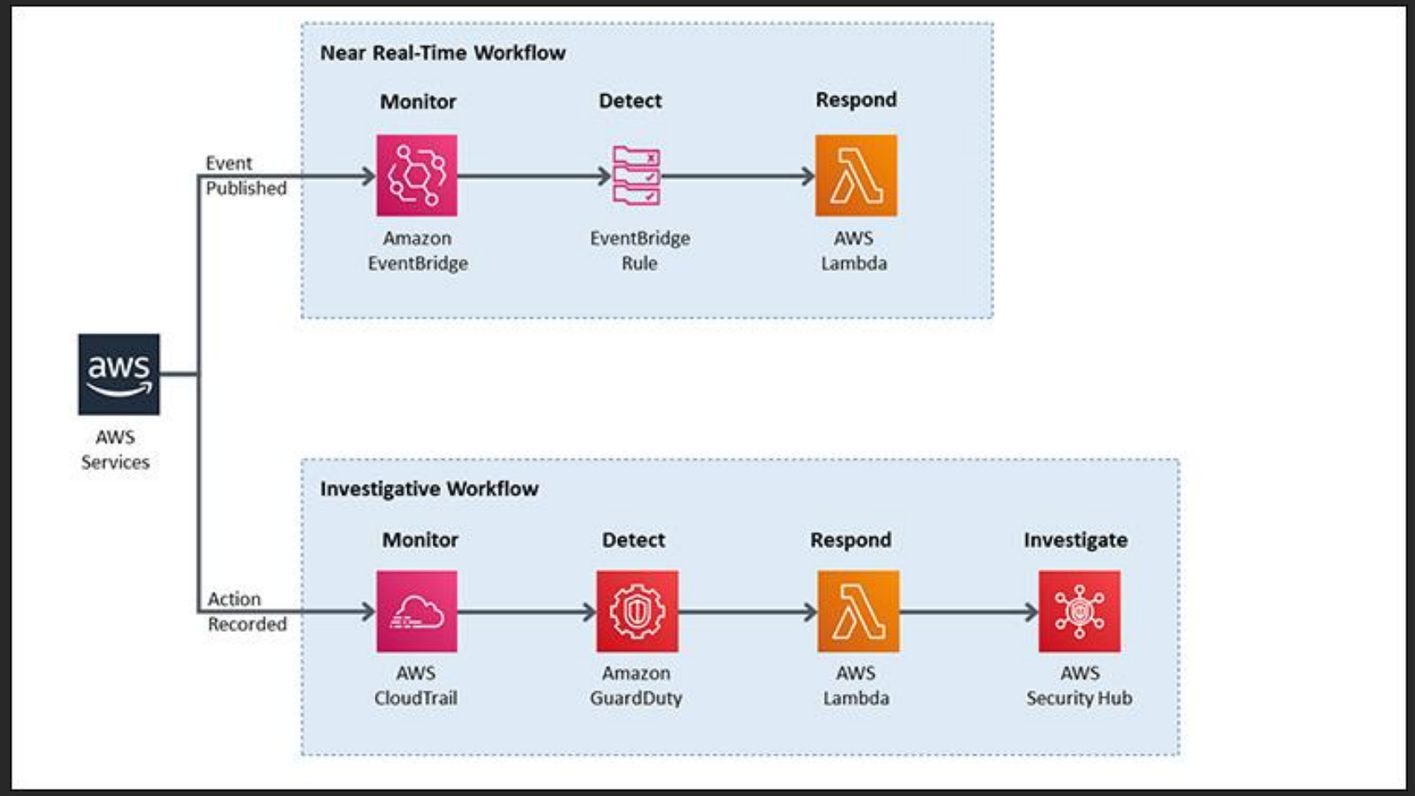

https://aws.amazon.com/blogs/security/how-get-started-security-response-automation-aws/

Security response automation on AWS

Monitor: Collection of AWS CloudTrail information about activities performed in your AWS account, usage metrics from your Amazon EC2 instances, or flow log information about the traffic going to and from network interfaces in your Amazon Virtual Private Cloud (VPC).

Detect: A triggering condition might be an anomalous activity detected by Amazon GuardDuty, a resource becoming out of compliance with an AWS Config Rule, or a high rate of blocked requests on an Amazon VPC security group or AWS WAF web access control list.

Respond: Examples of automated response actions might include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

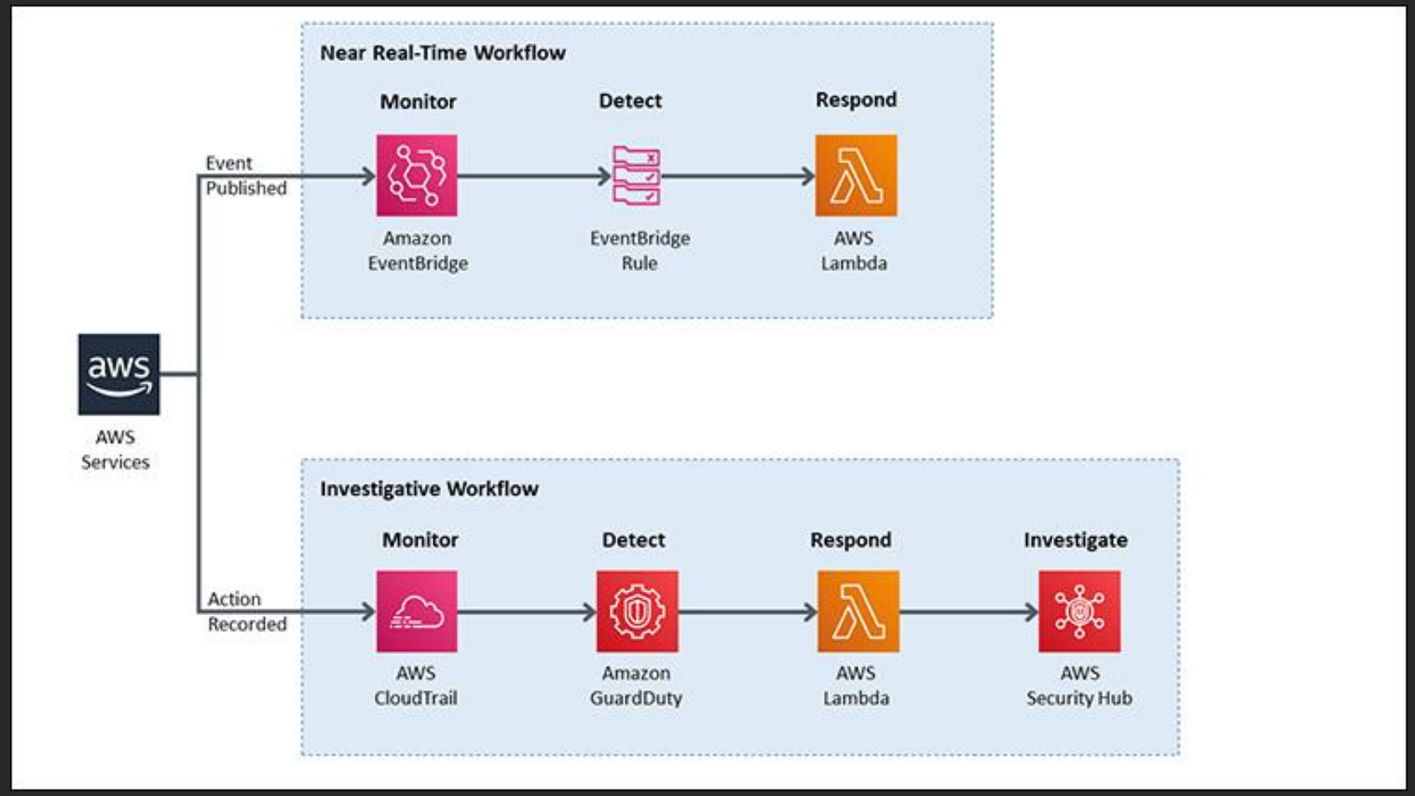

https://aws.amazon.com/blogs/security/how-get-started-security-response-automation-aws/

Threat response: Detailed playbook

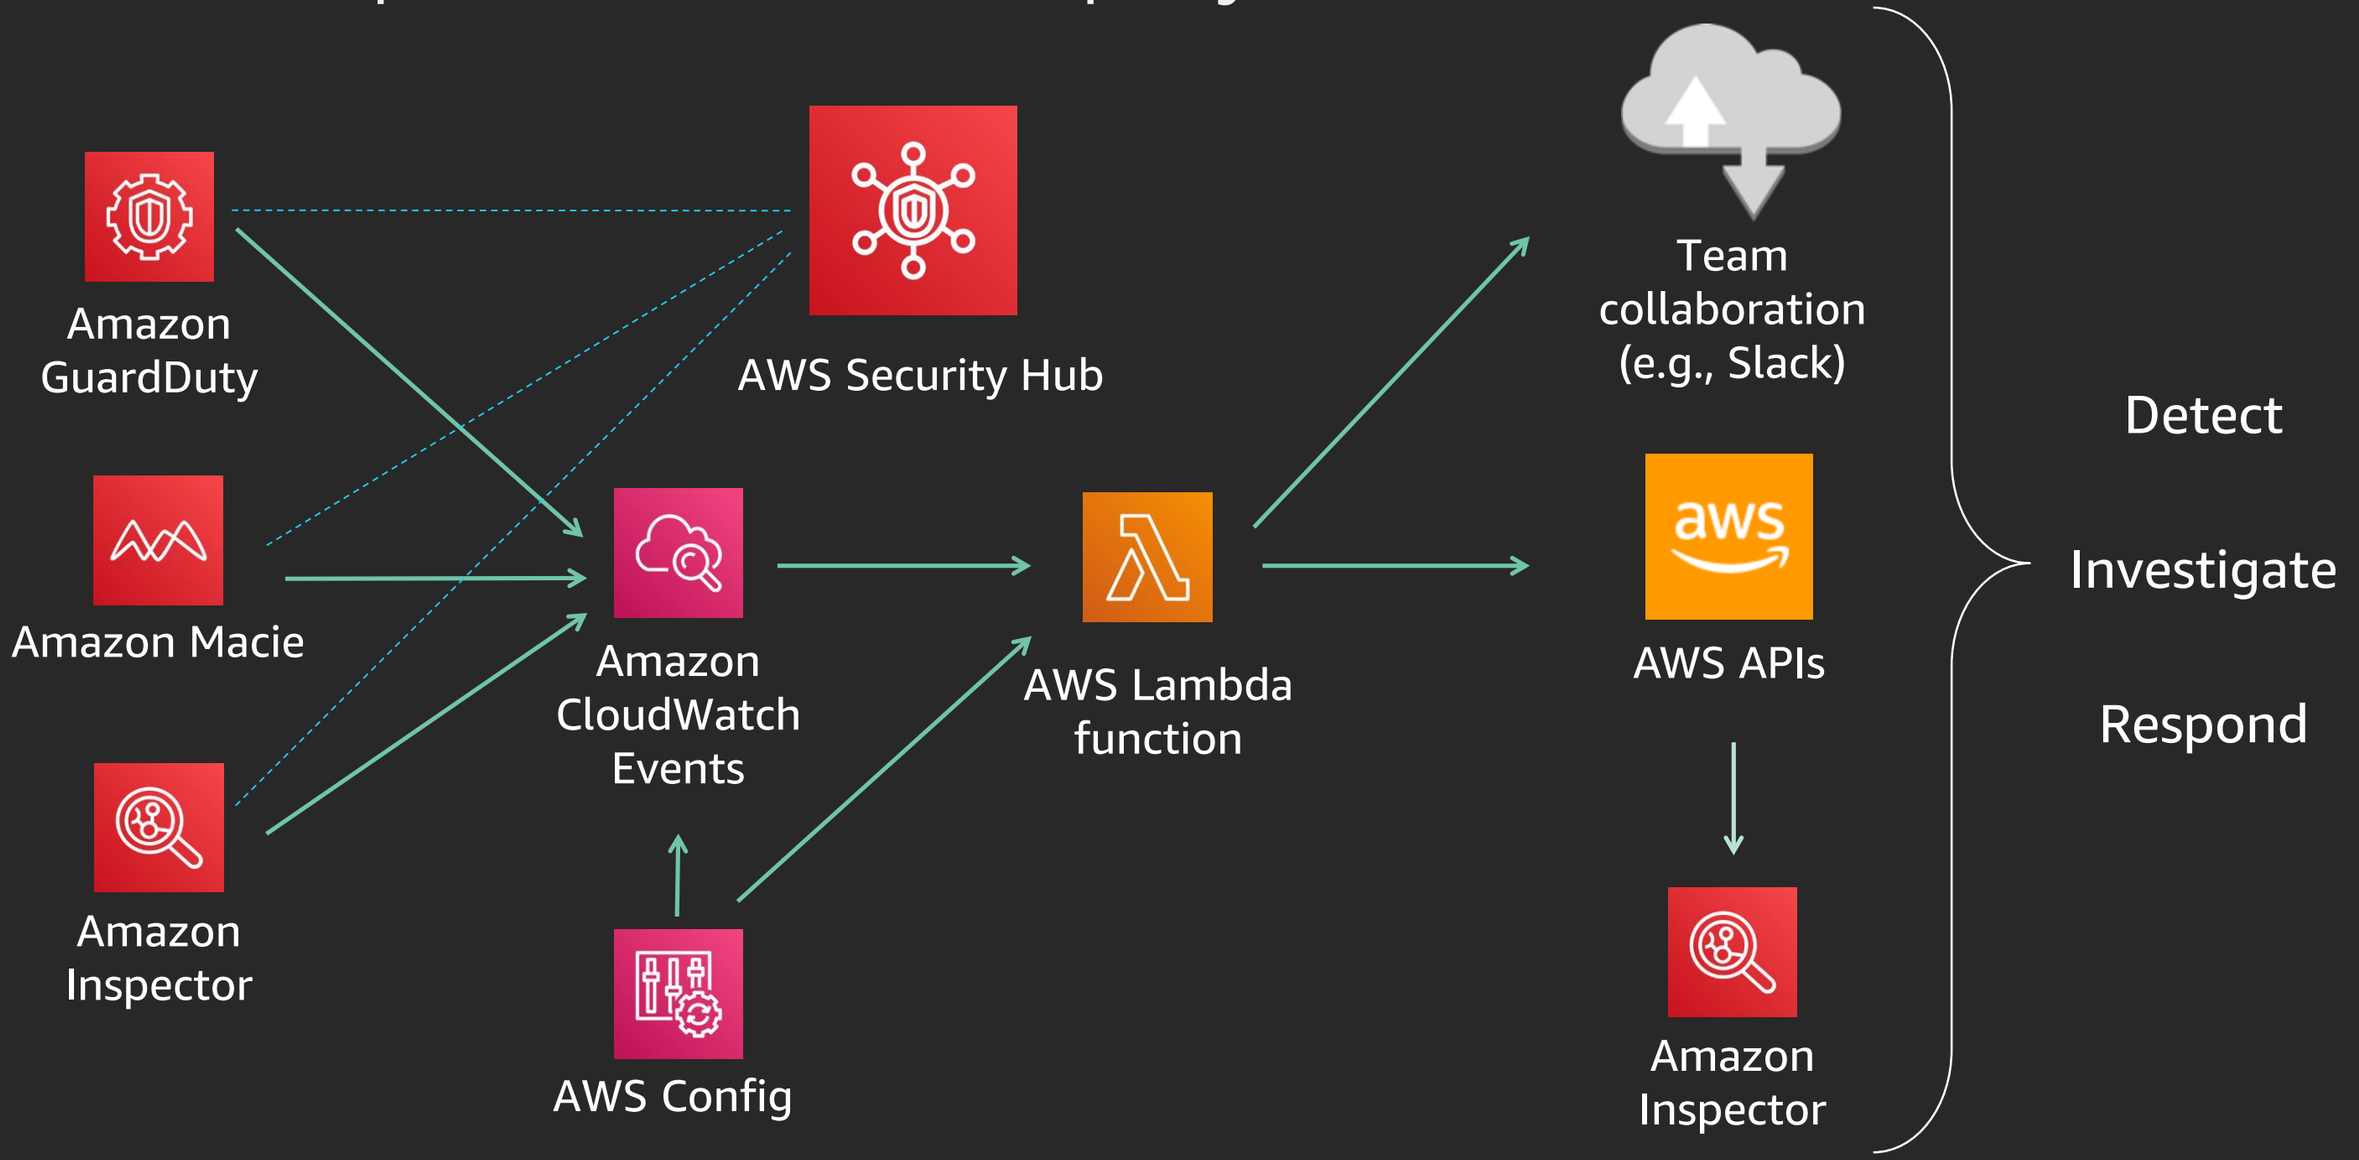

Workshop walk-through: What happened?

AWS re:Invent

Module 2: Attack target

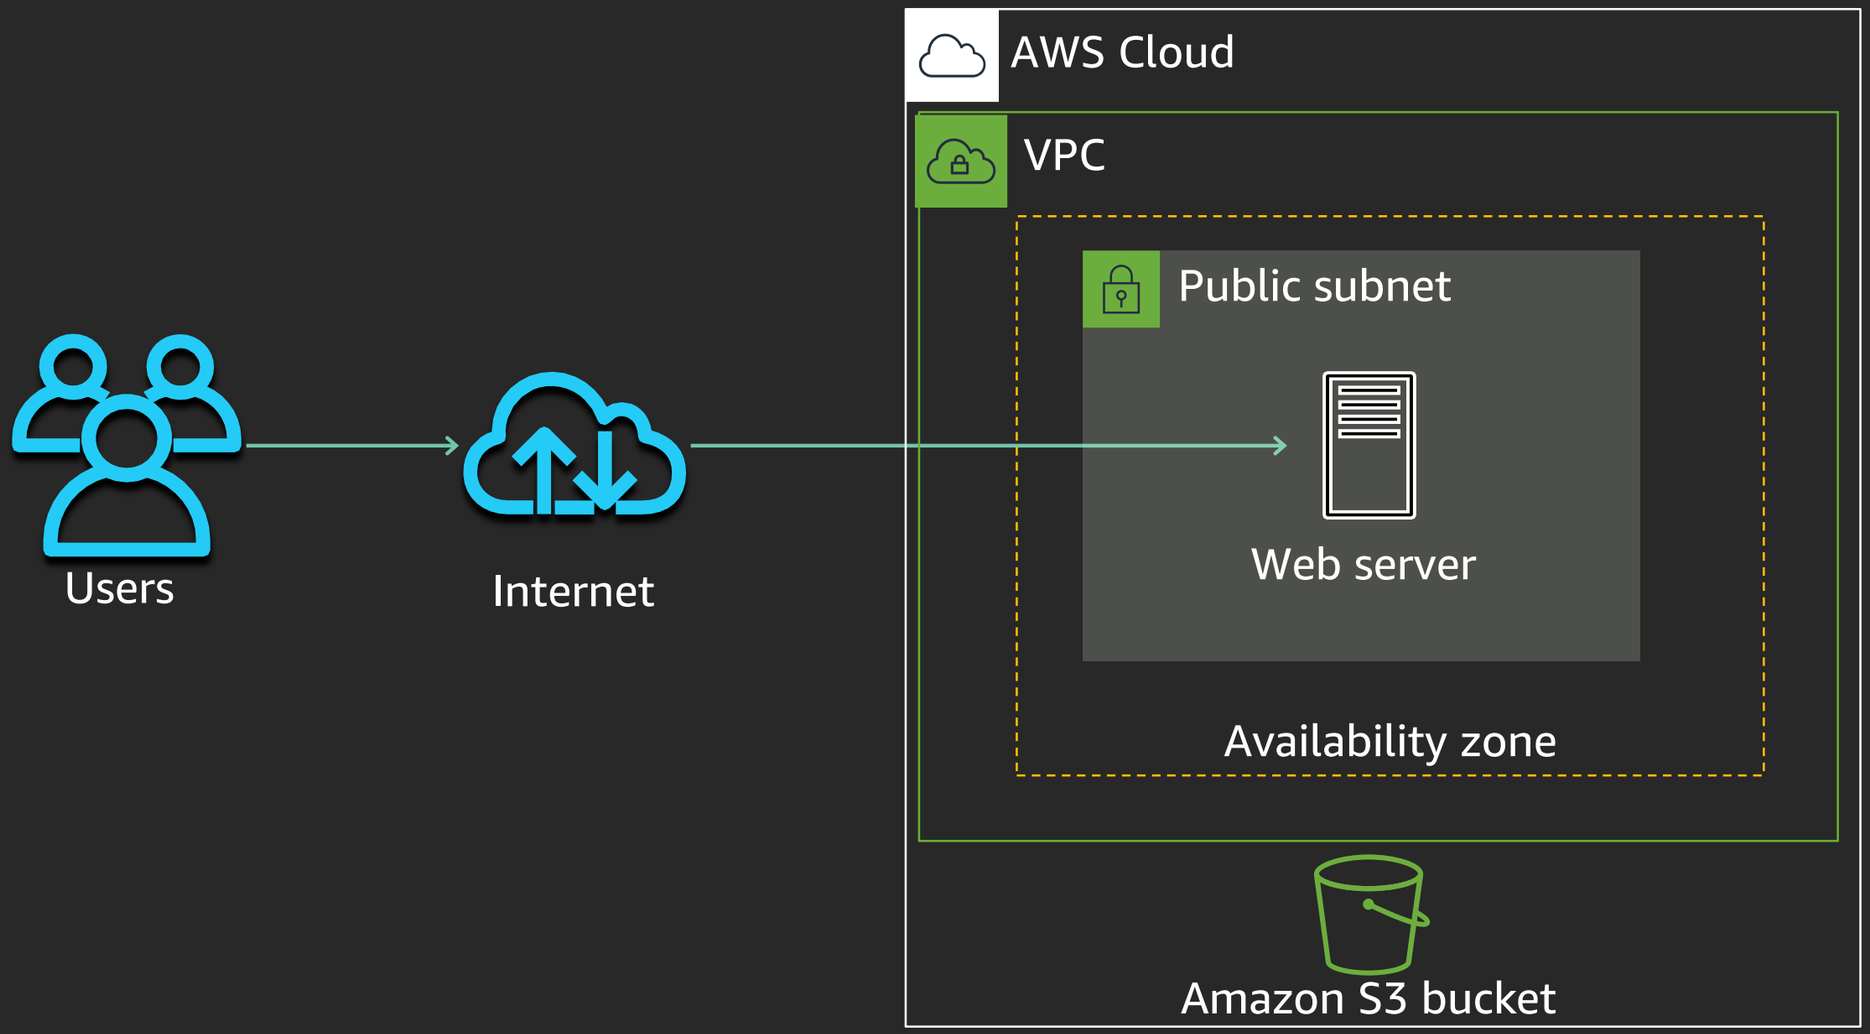

Module 2: Setup

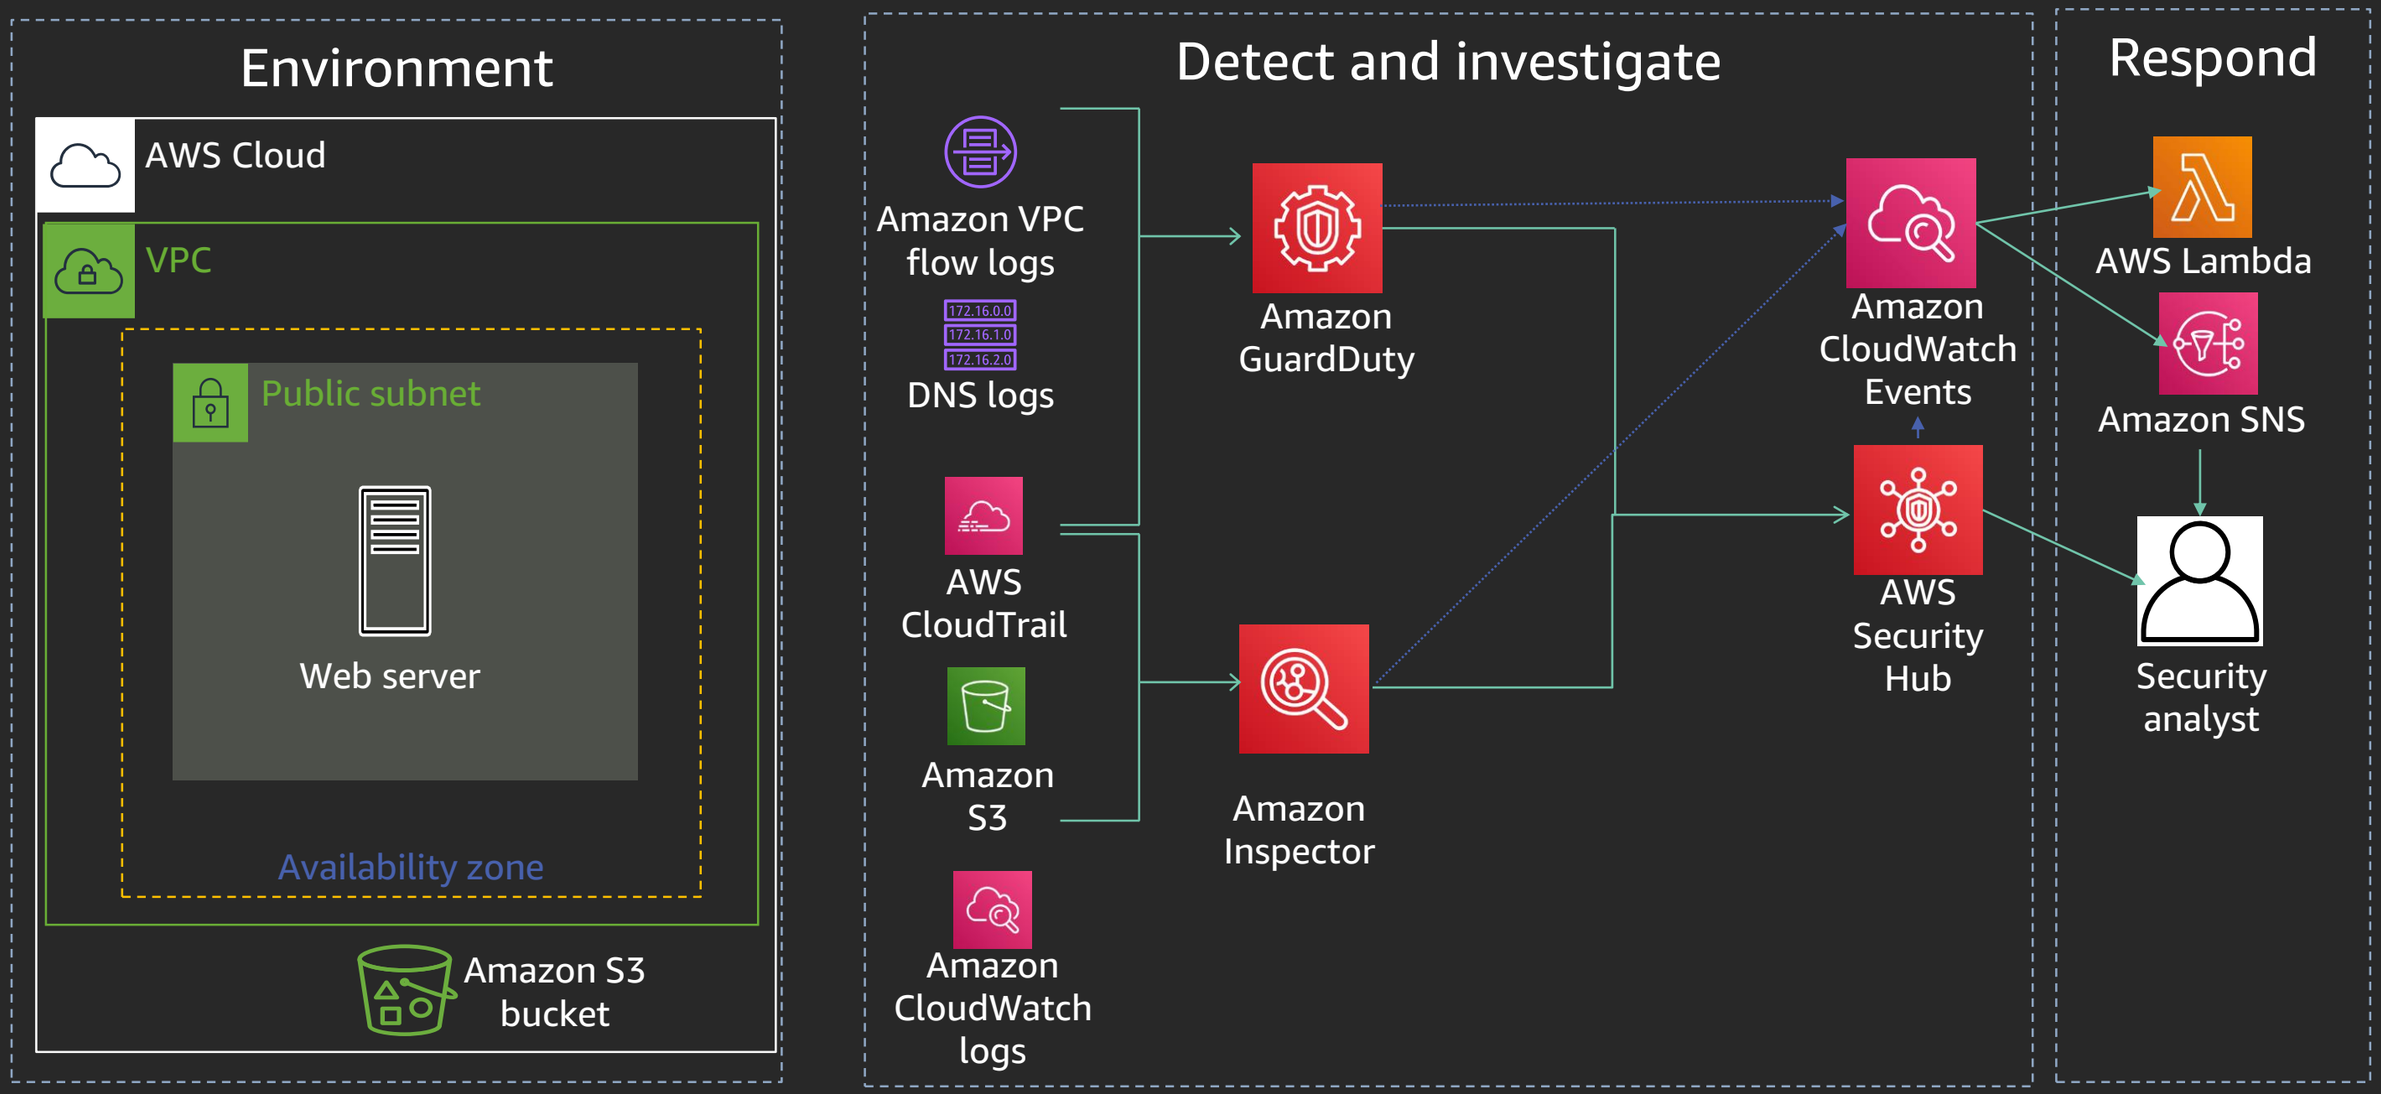

Module 2: The attack

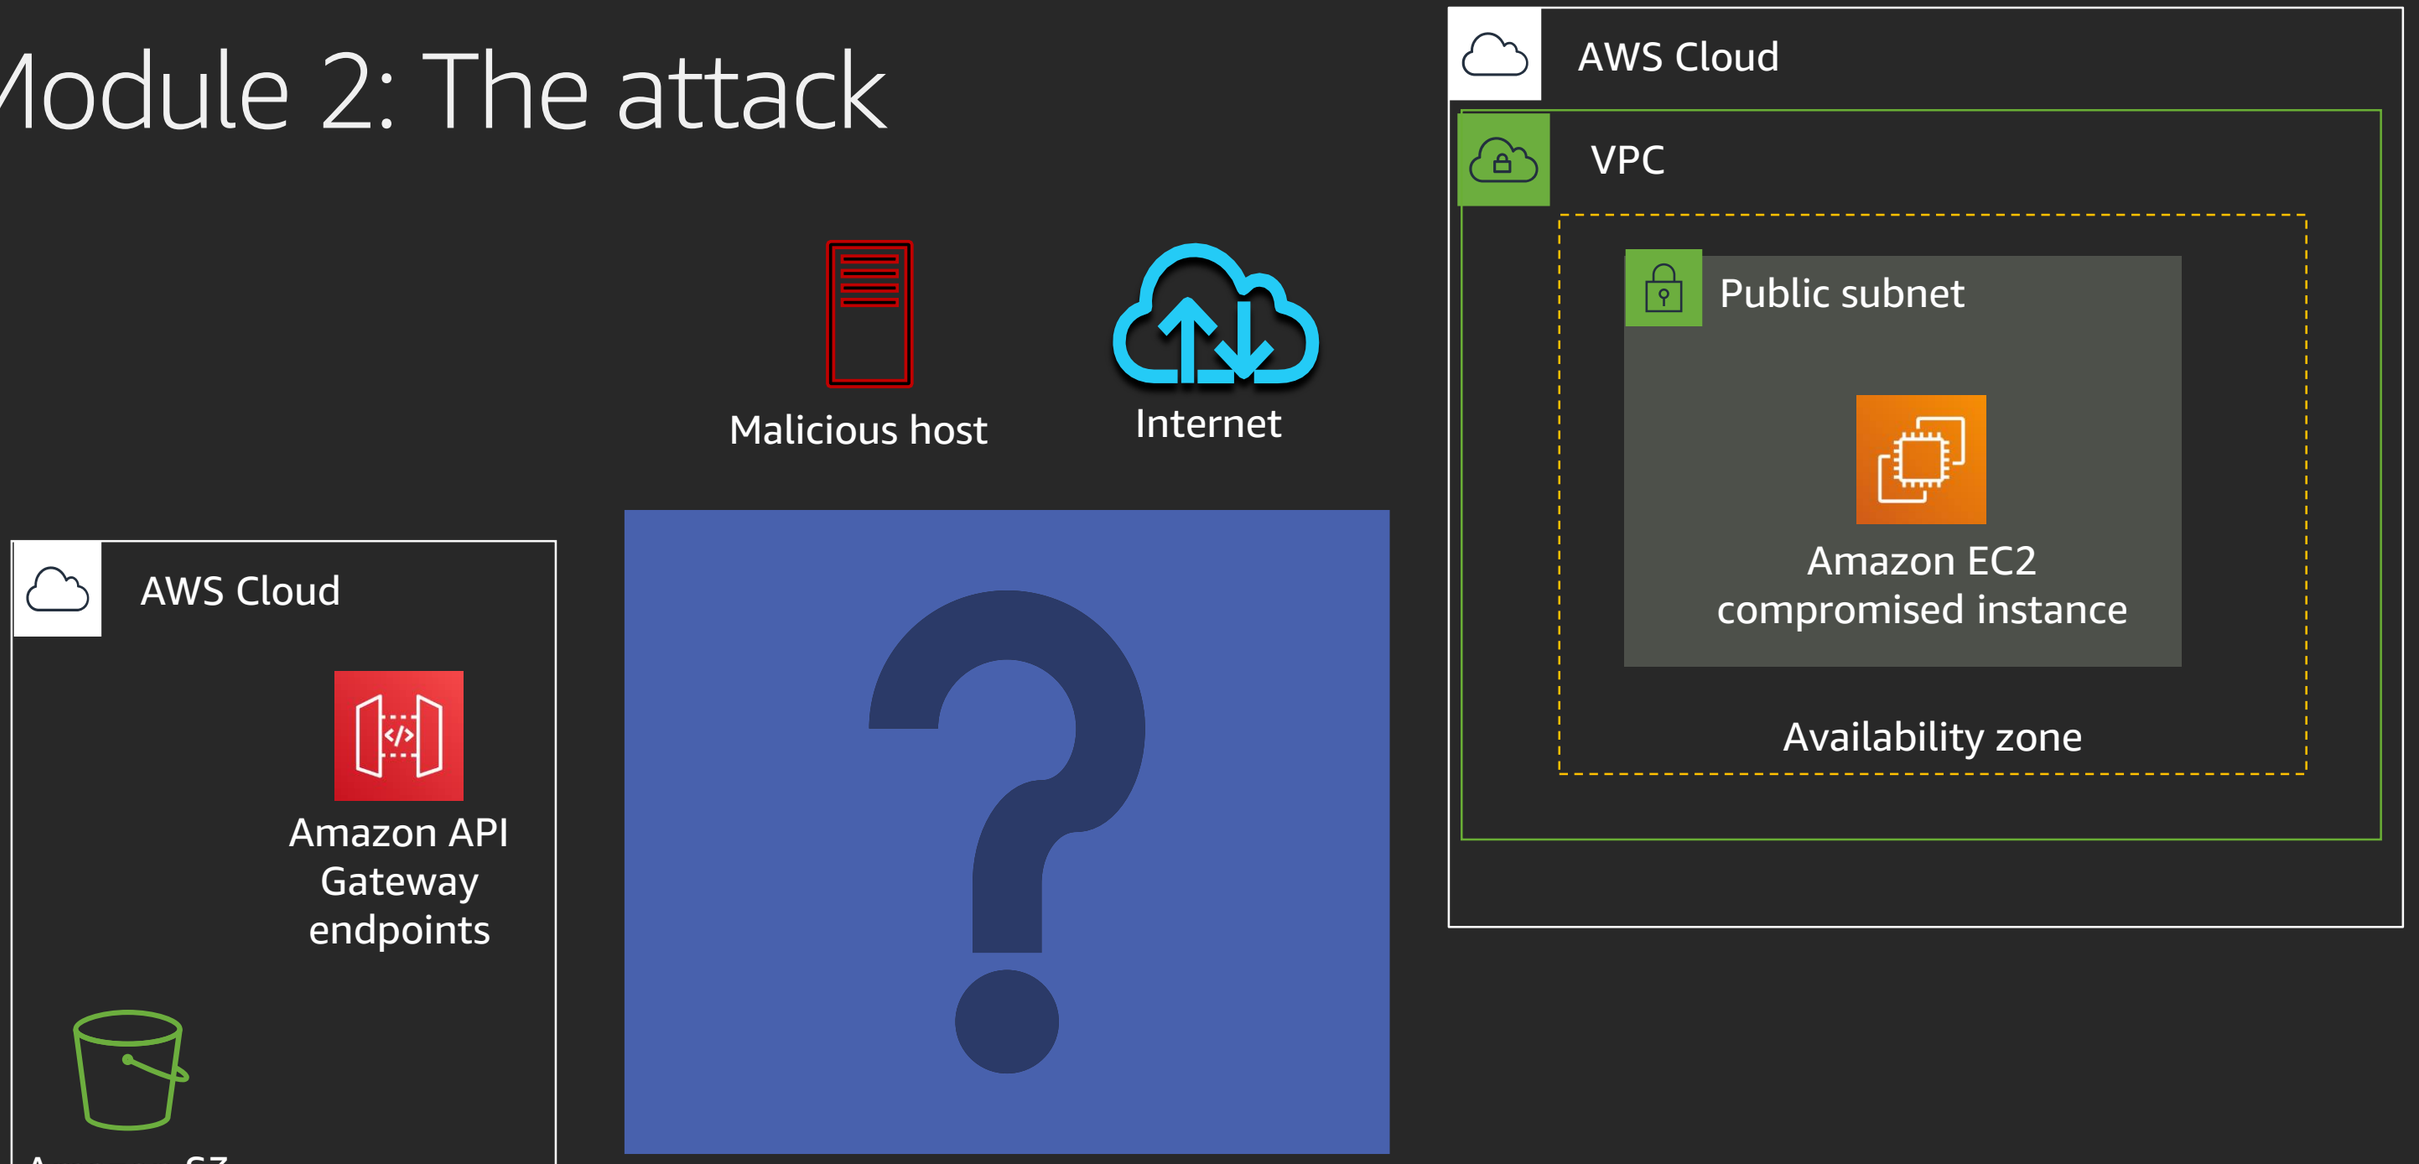

Module 3: Detect and response

Use US West (Oregon)
us-west-2

https://dashboard.eventengine.run

https://automating-threat-detection.awssecworkshops.com/

Directions

1. Browse to the URL
2. Choose **Module 3: Detect & Respond** in the outline on the left
3. Run through the module (~45 min.)

Review, questions, and lessons learned

AWS re:Invent

Module 4: What happened?

- Review (5 min.)
- Questions (10 min.)
- Lessons learned (5 min.)

Module 4: What happened?

Use US West (Oregon)
us-west-2

https://automating-threat-detection.awssecworkshops.com/

Directions

1. Browse to the URL
2. Choose **Module 4: Discussion** in the outline on the left
3. We will summarize the workshop, then answer questions
4. Do not need to do – Account lessons learned instructions

Don’t leave – Hands on with Amazon Detective is very soon.

Scenario discussion

Module 4: The attack

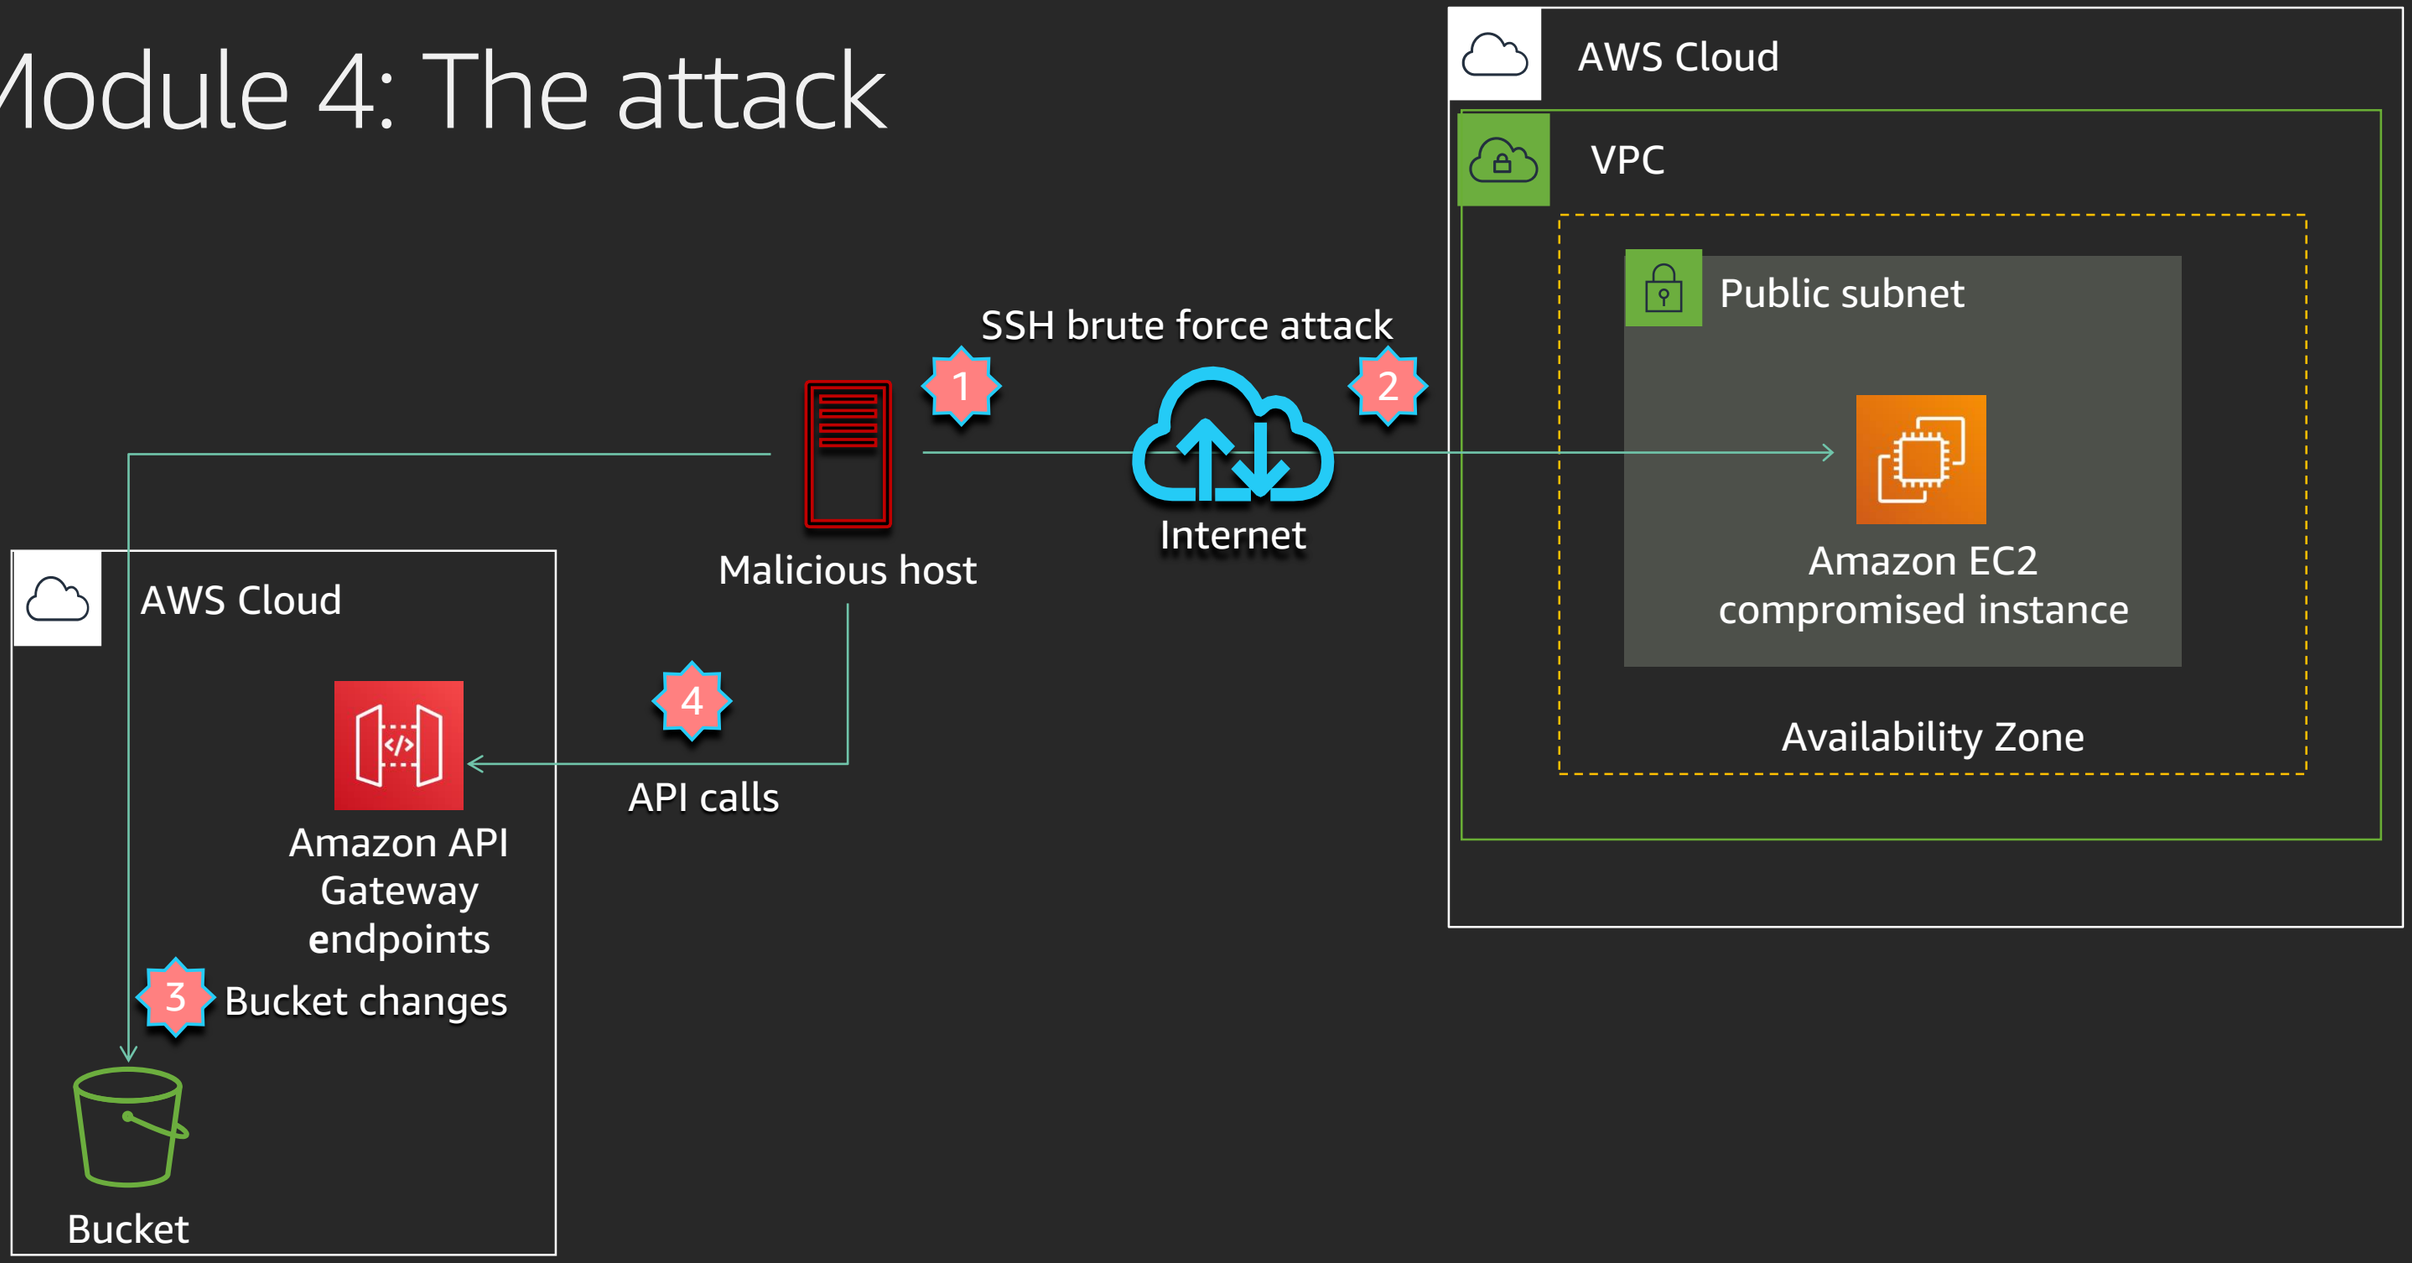

Module 4: What really happened?

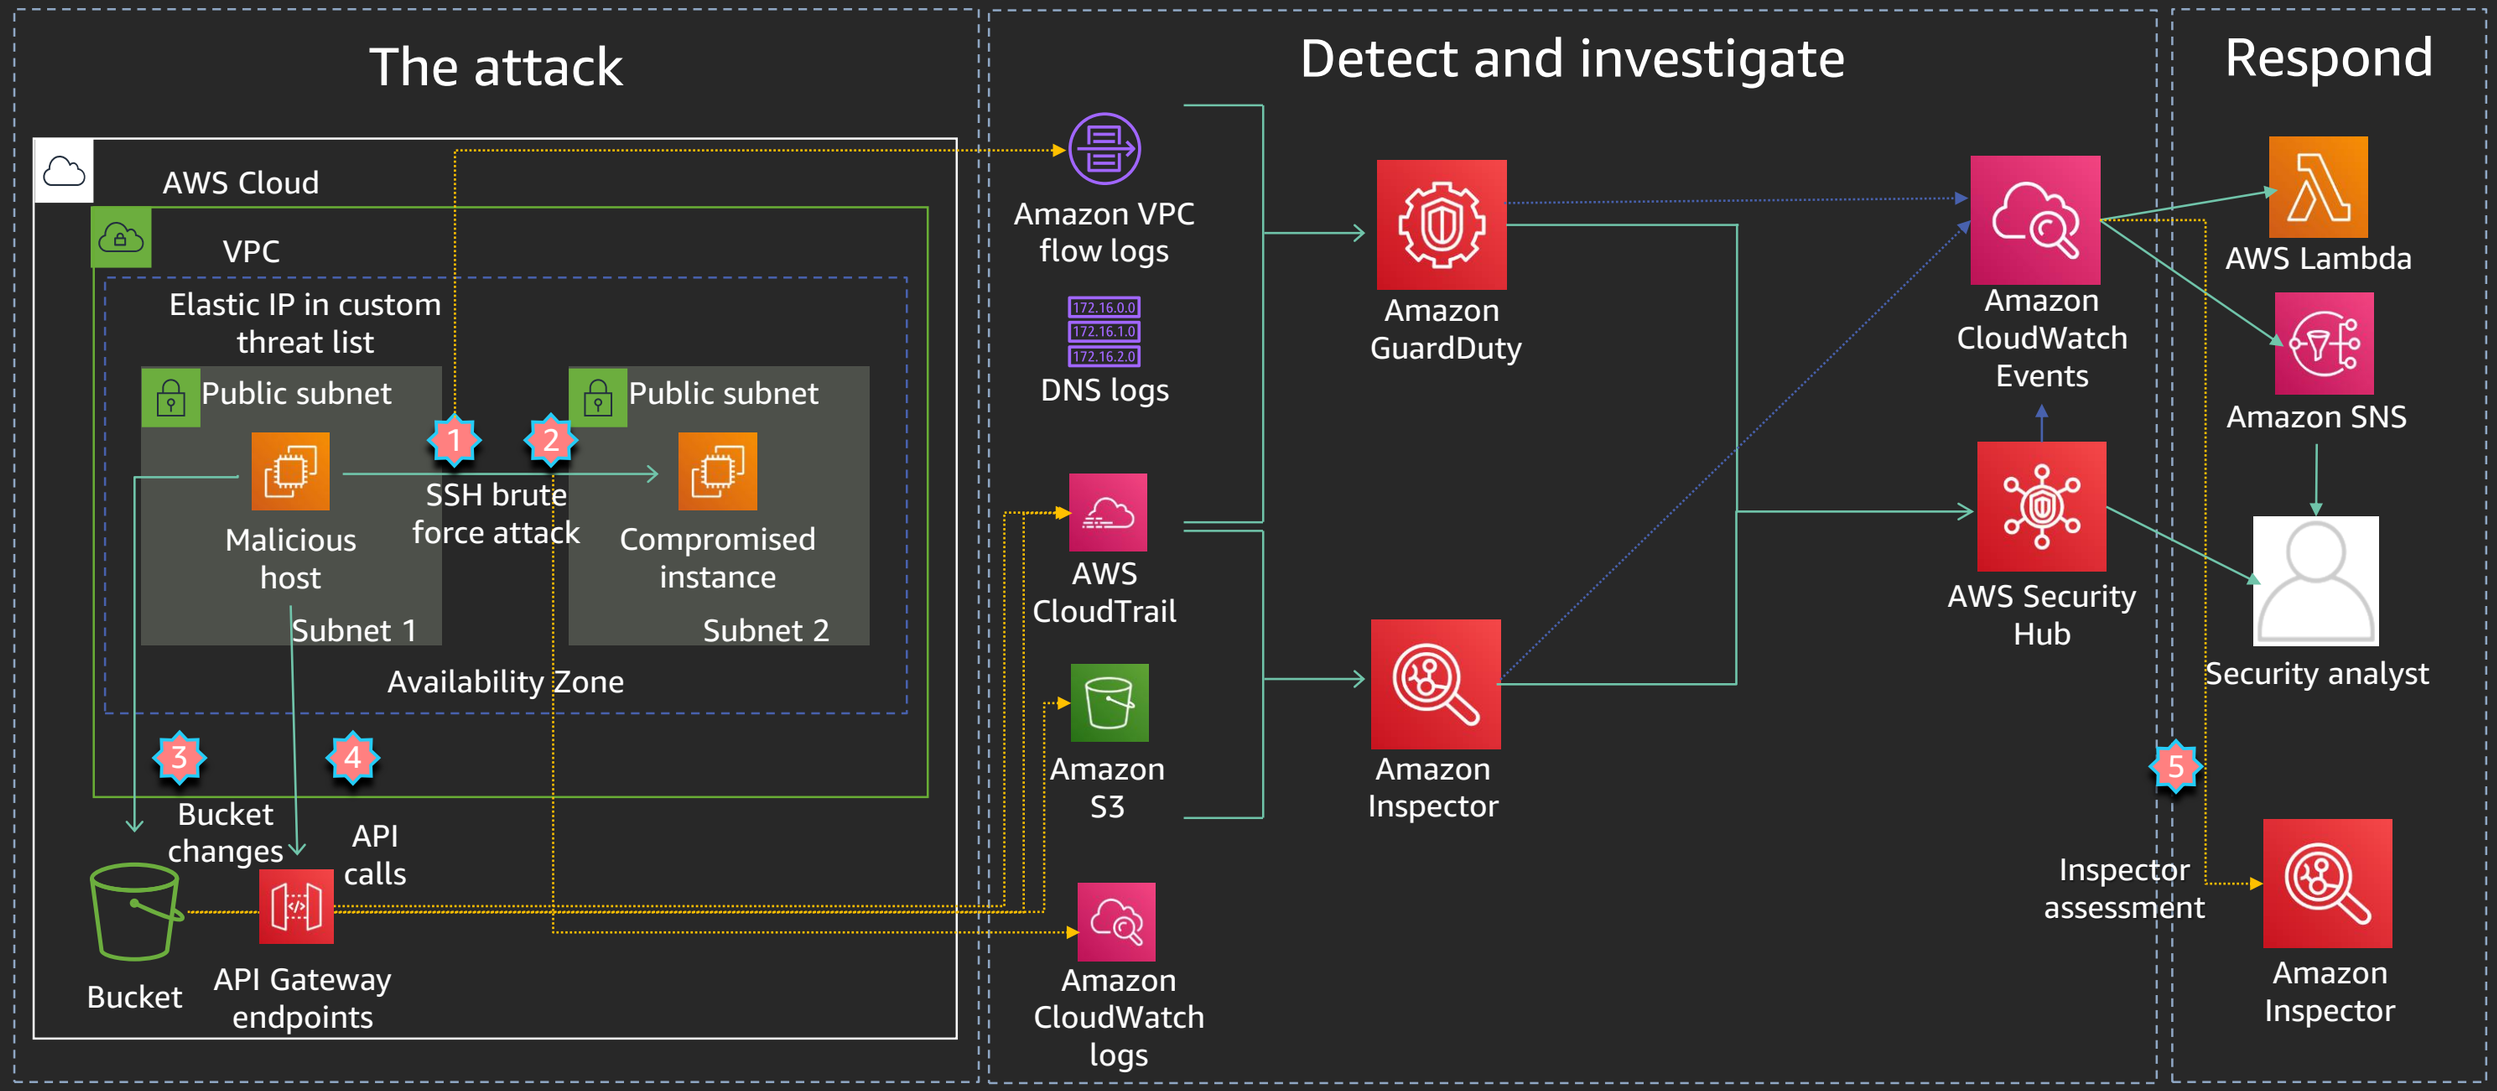

Questions

AWS re:Invent

Workshop questions 1

Which of the following AWS services have direct access to your Amazon EC2 instances?

What performance impact does Amazon GuardDuty have on your account if you have more than 100 VPCs?

How do you kick off notifications or actions based on events in Amazon GuardDuty?

Workshop questions 2

The lab mentions that you can ignore the high-severity SSH brute force attack finding; why?

Why did the API calls from the malicious host generate Amazon GuardDuty findings?

What is required for Amazon CloudWatch logs to capture evidence to help investigate an SSH brute force attack?

What key remediation step was missed regarding the SSH brute force attack?

Lessons learned from incident response

Use a strong tagging strategy.

Technical Tags

Name – Used to identify individual resources

Application ID – Used to identify disparate resources that are related to a specific application

Application Role – Used to describe the function of a particular resource (e.g. web server, message broker, database)

Cluster – Used to identify resource farms that share a common configuration and perform a specific function for an application

Environment – Used to distinguish between development, test, and production infrastructure

Version – Used to help distinguish between different versions of resources or applications

Tags for Automation

Date/Time – Used to identify the date or time a resource should be started, stopped, deleted, or rotated

Opt in/Opt out – Used to indicate whether a resource should be automatically included in an automated activity such as starting, stopping, or resizing instances

Security – Used to determine requirements such as encryption or enabling of VPC Flow Logs, and also to identify route tables or security groups that deserve extra scrutiny

Business Tags

Owner – Used to identify who is responsible for the resource

Cost Center/Business Unit – Used to identify the cost center or business unit associated with a resource; typically for cost allocation and tracking

Customer – Used to identify a specific client that a particular group of resources serves

Project – Used to identify the project(s) the resource supports

Security Tags

Confidentiality – An identifier for the specific data-confidentiality level a resource supports

Compliance – An identifier for workloads designed to adhere to specific compliance requirements

https://aws.amazon.com/answers/account-management/aws-tagging-strategies/

https://aws.amazon.com/blogs/aws/new-use-tag-policies-to-manage-tags-across-multiple-aws-accounts/

Lessons learned from incident response

- Questions to ask during the investigation
  - Is this finding a true positive?
  - is the event an unusual activity or more?
  - Where did the incident occur?
  - Who reported or discovered the incident?
  - How was it discovered?
  - Are there any other areas that have been compromised by the incident? If so, what are they and when were they discovered?
  - What is the scope of the impact?
  - What is the business impact?
  - Have the source(s) of the incident been located? If so, where, when, and what are they?

Amazon Detective

AWS re:Invent

Introducing Amazon Detective (Preview)

Analyze and visualize security data to rapidly get to the root cause of potential security issues

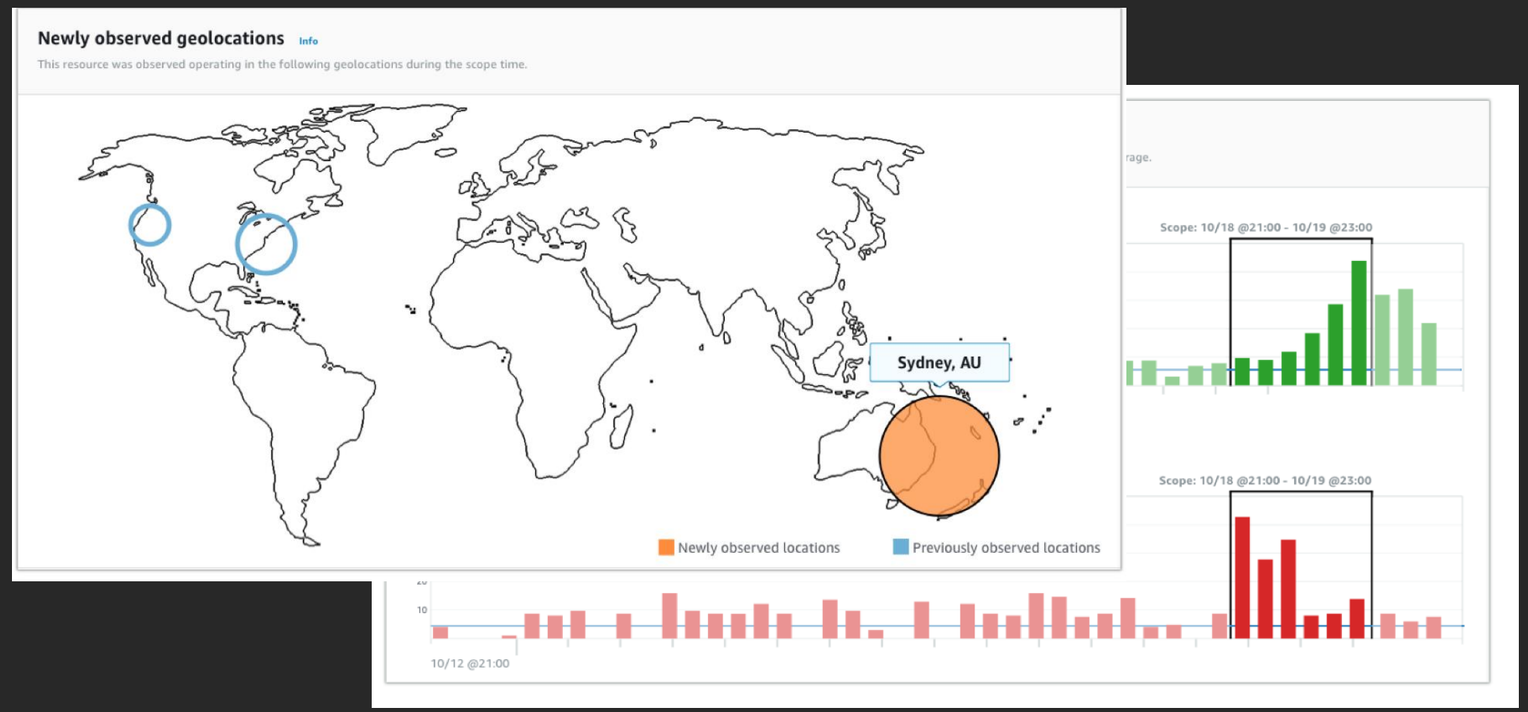

How Amazon Detective works?

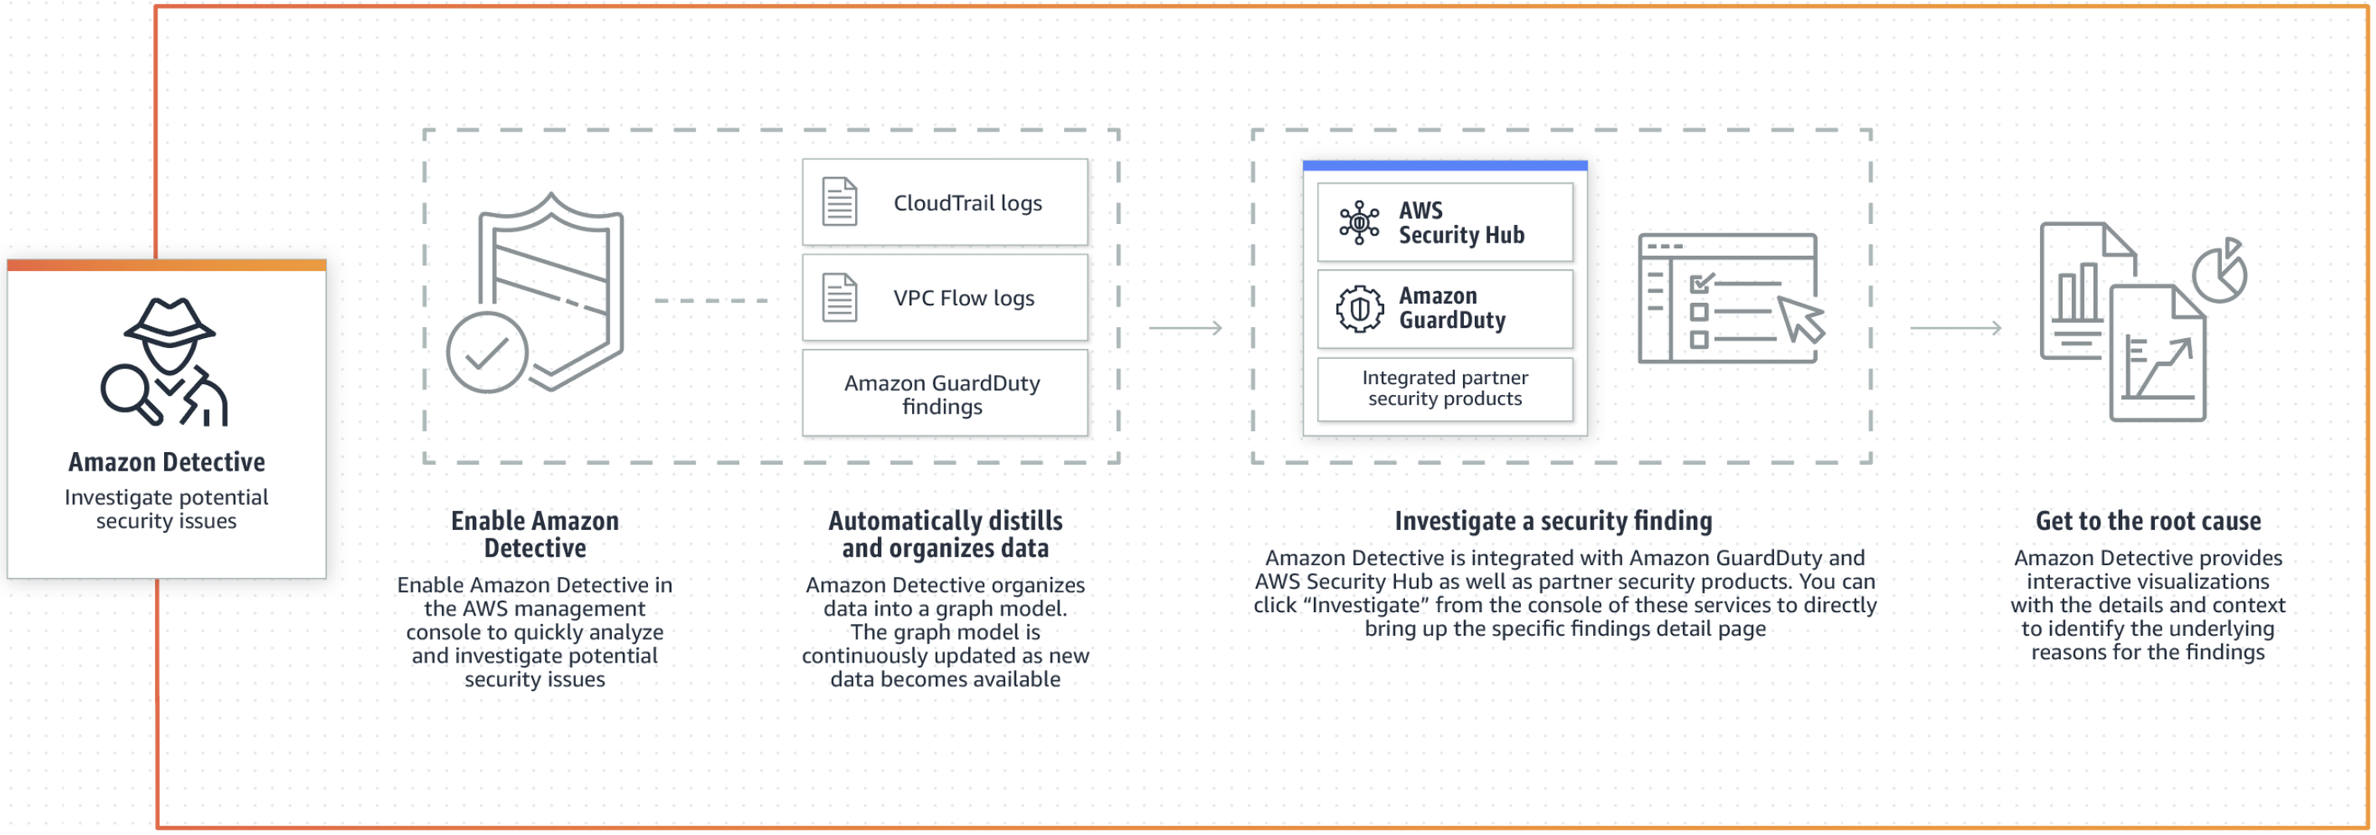

Pivot from GuardDuty, Security Hub and other consoles

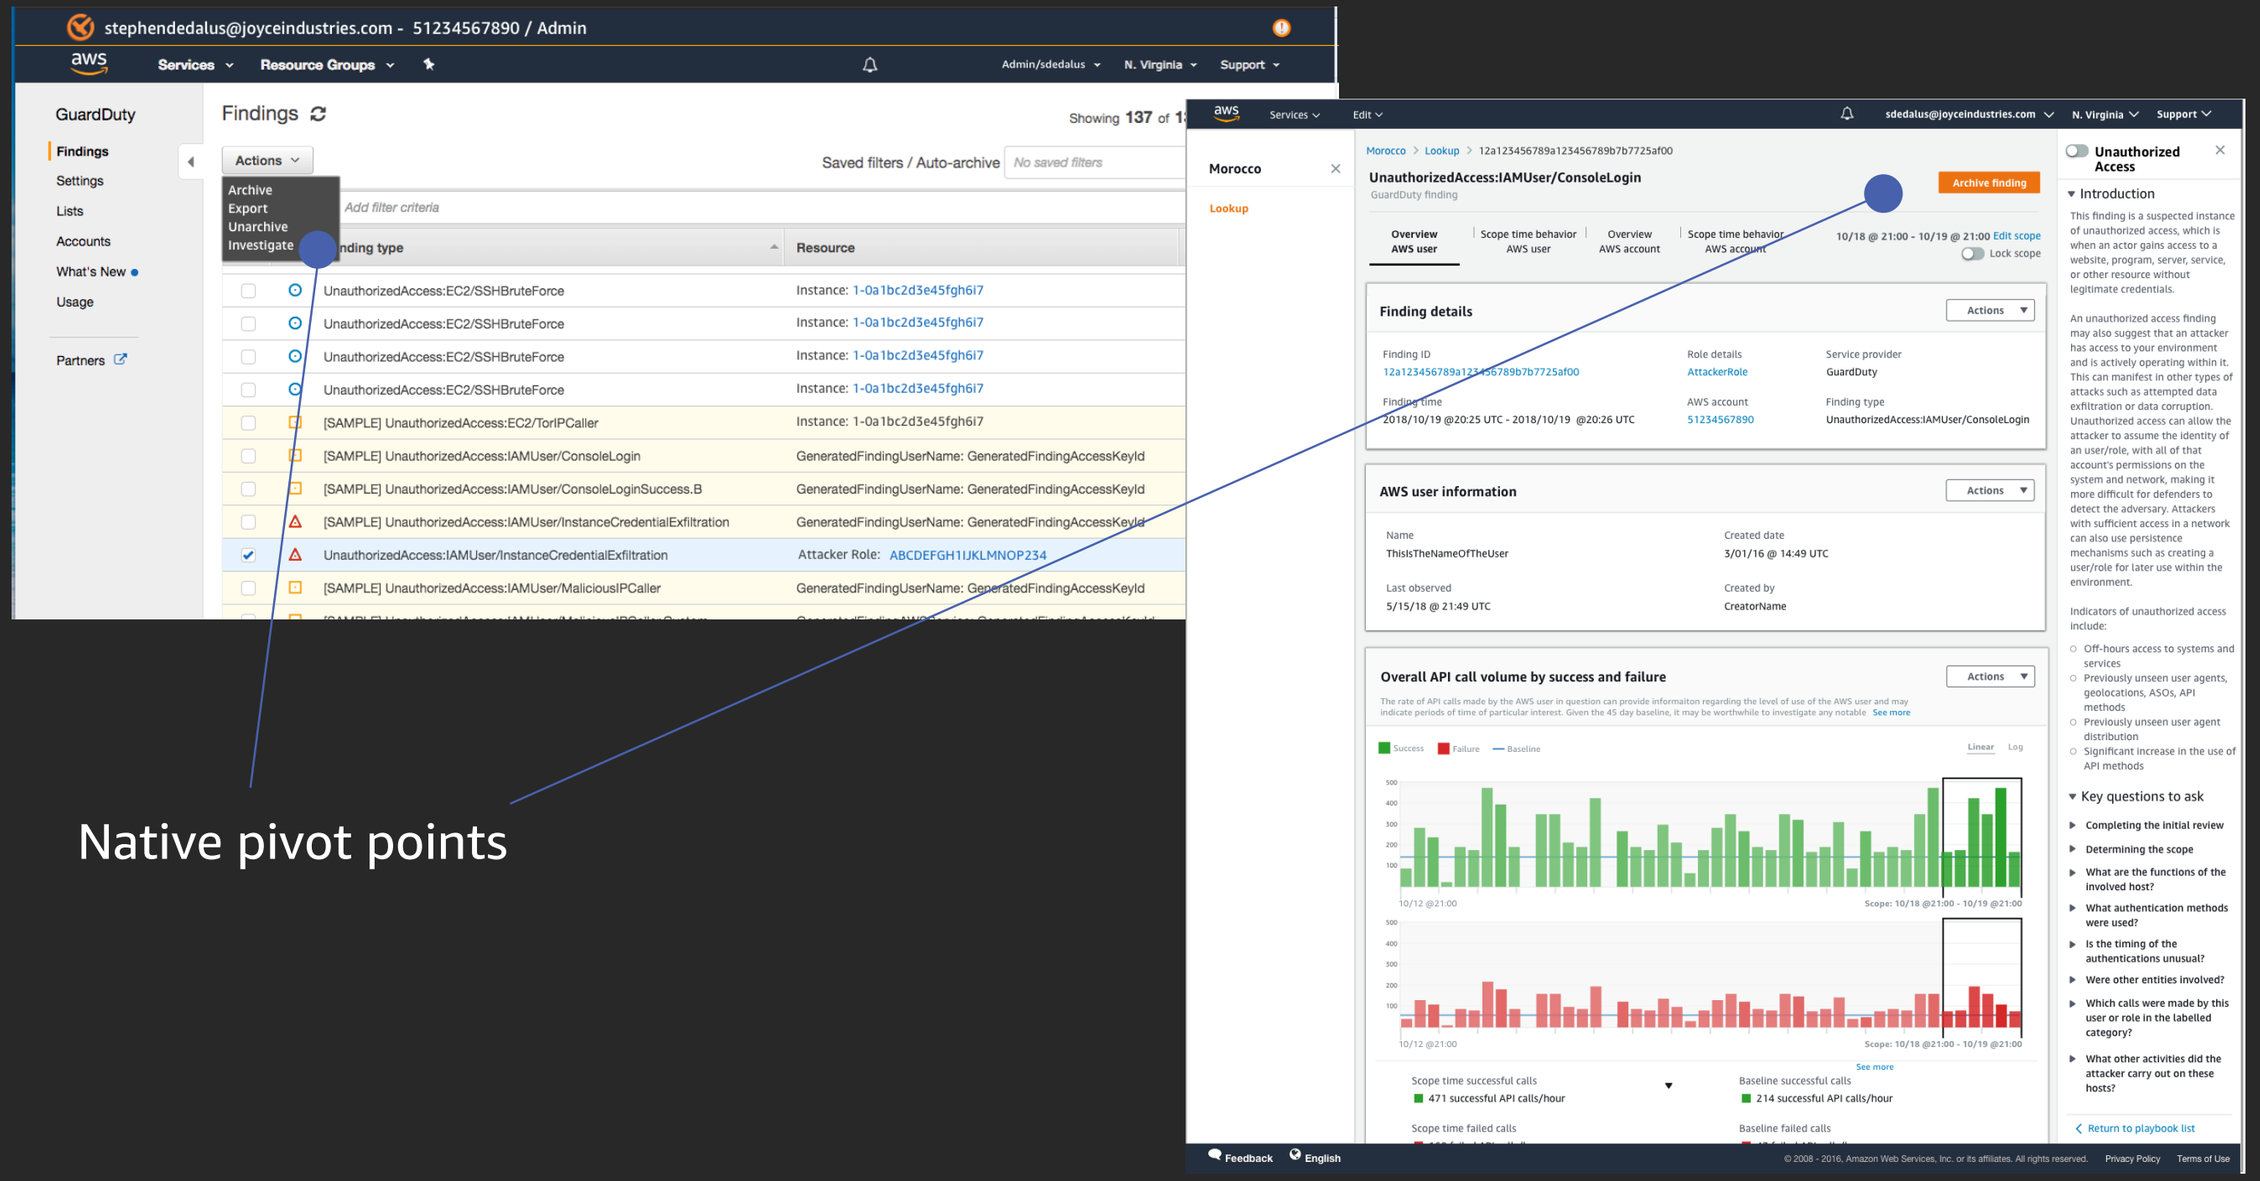

Native pivot points

Hands on with Amazon Detective

AWS re:Invent

Module 5: re:Invent 2019

Use US West (Oregon)
us-west-2

https://dashboard.eventengine.run

https://automating-threat-detection.awssecworkshops.com/

Directions

1. Browse to the URL
2. Choose **Module 5: re:Invent 2019** in the outline on the left
3. Run through the module (~15 min.) OR watch the presenter

Remember we are playing the part of a security analyst. Take good notes.

There will be prizes.

What happened?

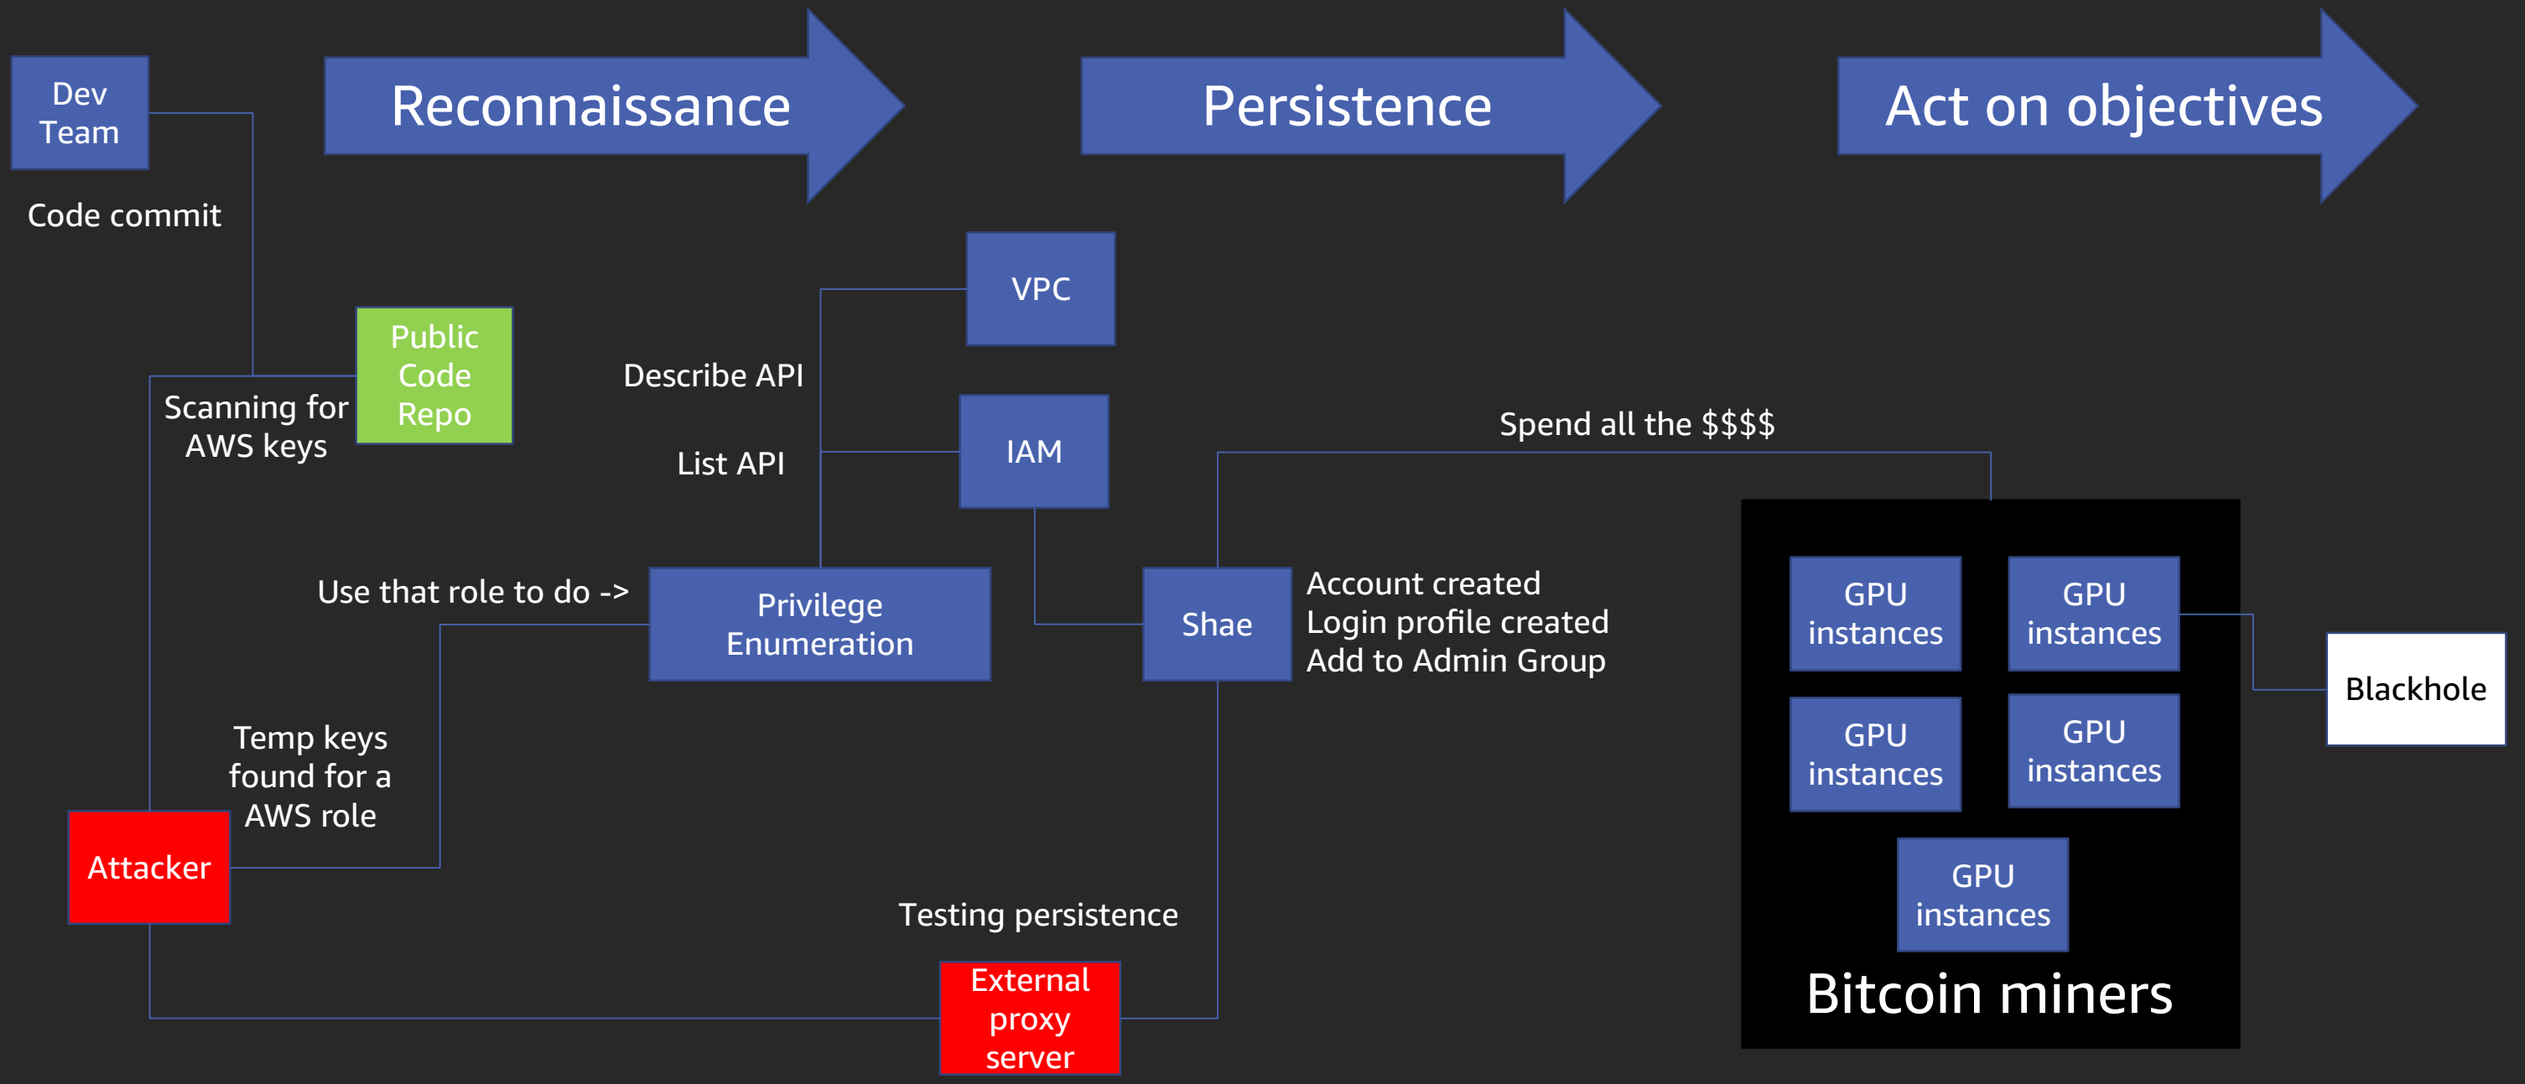

Recon:IAMUser/UserPermissions **Persistence:IAMUser/UserPermissions** **CryptoCurrency:EC2/BitcoinTool.B**

IoC List

VPC Flow

- 104.140.201.42 (bitcoin IP)
- 199.2.137.201
- 199.2.137.22
- 199.2.137.21
- 199.2.137.25
- 199.2.137.20 or 199.2.137.0/24

CloudTrail

- 78.109.29.23
- 185.247.211.91
- 45.9.237.16
- 45.9.236.13
- 165.231.161.12
- 165.231.161.4
- 94.176.148.35
- 86.105.9.94

Impacted EC2 instances

- i-08659bebb829b7ec8
- i-0567bc95c2e52acd9
- i-0d901bdde3d63c2d3
- i-0441b8673fab0f671
- i-0cd943f23b0da5cf0

User Agents

- Im the adversary. Catch me if you can.
- Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/78.0.3904.97 Safari/537.36 (Brave Browser Console Logins)
- aws-cli/1.16.148 Python/3.6.8 Linux/5.0.0-36-generic botocore/1.13.17 (cli usage)

IoC List

GuardDuty findings

- Recon:IAMUser/UserPermissions - c6b736ca2c529b1791634f5cad495c49
- Persistence:IAMUser/UserPermissions - 1cb736cdc0f4802c50c2d407e5c28100
- UnauthorizedAccess:IAMUser/ConsoleLogin - ccb736dc83e39538fcf2b2ad9d83f2dc
- Stealth:IAMUser/CloudTrailLoggingDisabled - f4b75079ce971335581bb30beac8e67f
- CryptoCurrency:EC2/BitcoinTool.B - 14b750779a4f35d3c9b05461257c2e79,58b7507784787cf256059fd2b9950a6b,3ab75077915c63a8f0c5db671abd8f61,d8b7507789c8645a29bda24776e650d2,40b750779b84348b2dc6349d8d7e6fe2
- Trojan:EC2/BlackholeTraffic - 14b750849deaedb2e08a8666eddf1ef7

Impacted user/roles

- AWSReservedSSO_AdministratorAccess_2647948df6af6adc - AROA55TXQP4FK3H45TMI5
- Shae - AIDA55TXQP4FDRFKD5N7E
- Jon - AIDA2ZZIW6DOFWLNRIIOB
- imopatis@corp.essos.com - AROA2ZZIW6DOPSLFDSASC

Useful links

https://aws.amazon.com/security/

https://enterprise.verizon.com/resources/reports/DBIR_2018_Report.pdf

https://www.nist.gov/cyberframework

https://d0.awsstatic.com/whitepapers/AWS_CAF_Security_Perspective.pdf

https://aws.amazon.com/security/penetration-testing/

Thank you!

AWS re:Invent

Please complete the session survey in the mobile app.