UNIVERSITY OF AMSTERDAM

Profiling abuse of exposed secrets in public repositories

Maurice Mouw
University of Amsterdam
mmouw@os3.nl

1 February - 2021

Supervisors:

Mick Cox
Deloitte

Fons Mijnen
Deloitte

Introduction

- In 2016 an **Uber breach** in which 57 million customer records where exposed was due to an **access key** being **accessible via github**.[1]

- A **Starbucks API key** for JumpCloud was **discovered** in October of 2019 by a researcher in a **public Github repository**.[2]

Research Question:

How are leaked credentials/secrets on code collaboration platforms like github, gitlab, sourceforge and bitbucket found and abused by malicious users in a cloud platforms like AWS?

[1] https://www.bloomberg.com/news/articles/2017-11-21/uber-concealed-cyberattack-that-exposed-57-million-people-s-data
[2] https://latesthackingnews.com/2020/01/04/starbucks-exposed-an-api-key-in-github-public-repository/

Related Work

- Meli et al. did a **large scale** systematic **study** on **Github**. The researchers estimate that **93.74%** of the **discovered API secrets** are **sensitive** and 76.24% of the found asymmetric keys are sensitive.[1]

- **Atlassian** created a open source project called **Spacecrab** with which they created and **published AWS tokens**. Of the published tokens on Github **82.38%** where **found** after **approximately 30 minutes** of publication.[2]

[1] https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_04B-3_Meli_paper.pdf

[2] https://i.blackhat.com/briefings/asia/2018/asia-18-bourke-grzelak-breach-detection-at-scale-with-aws-honey-tokens-wp.pdf

Background

- **Honey tokens** are digital entities that help identify malicious events. They can come in many forms, it could be an URL, database record, user account, fake executable or e-mail address.

- **Mitre ATT&CK** is a framework that attempts to describe the techniques attackers use against (cloud) environments and services and possible mitigation techniques.[1]

[1] https://attack.mitre.org/versions/v8/matrices/enterprise/cloud/

Methods: Lab environment (CloMo)

- Automated setup of Logstash, Elasticsearch, Grafana and jenkins to analyze Cloudtrail and filebeat log data in an AWS environment.

- Use terraform, saltstack and jenkins to automate the deployment and configuration of the AWS experimental environment.

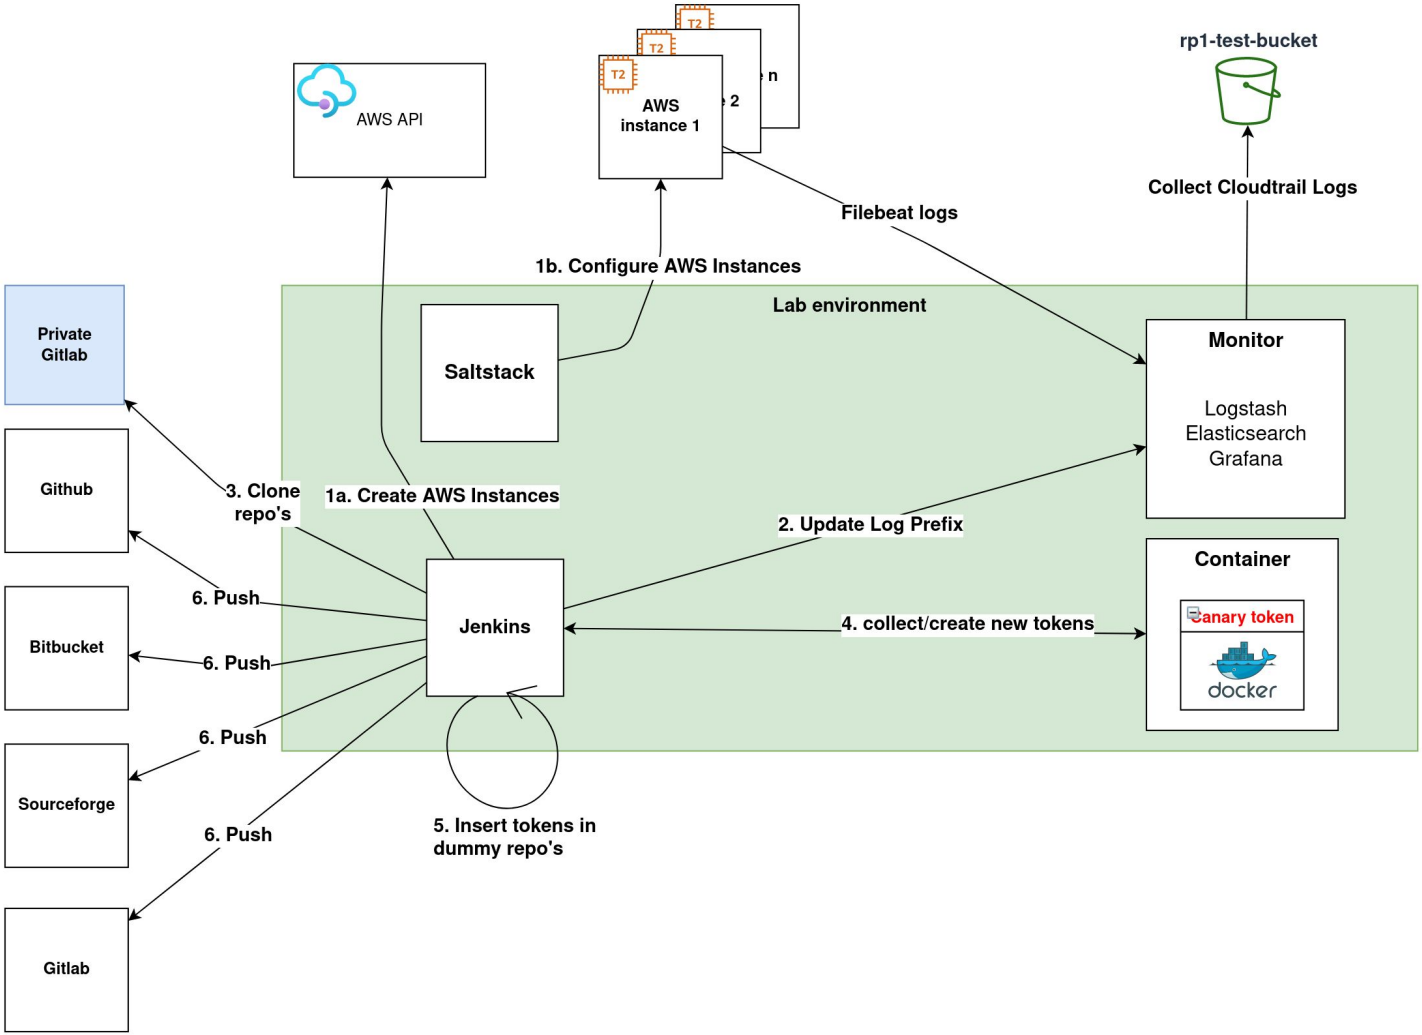
AWS API
T2
T2
T2
AWS instance 1
rp1-test-bucket
Collect Cloudtrail Logs
Filebeat logs
1b. Configure AWS Instances
Lab environment
Private Gitlab
Saltstack
Monitor
Logstash
Elasticsearch
Grafana
Github
3. Clone repo's
1a. Create AWS Instances
2. Update Log Prefix
Container
Canary token
docker
Bitbucket
6. Push
Jenkins
4. collect/create new tokens
6. Push
Sourceforge
6. Push
5. Insert tokens in dummy repo's
6. Push
Gitlab

Methods: Experiment-1

- Create 1 set of Github, Gitlab, Bitbucket and Sourceforge accounts with repositories.

- Create unique IAM account for each repository that has read access to multiple AWS services in eu-west-1 region.

- Publish code containing unique md5 hashed (weak) passwords per platform, AWS tokens and Canary (AWS) tokens. Tokens where published using the terraform syntax.

```
provider "aws" {
  region     = "eu-west-1"
  access_key = "AKIAXYZDQCEN5OYVW4OQ"
  secret_key = "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"
}
```

- Duration:
  - 5 iterations of 2 hours each before refreshing tokens.
  - 5 iterations of 24 hours before refreshing tokens.

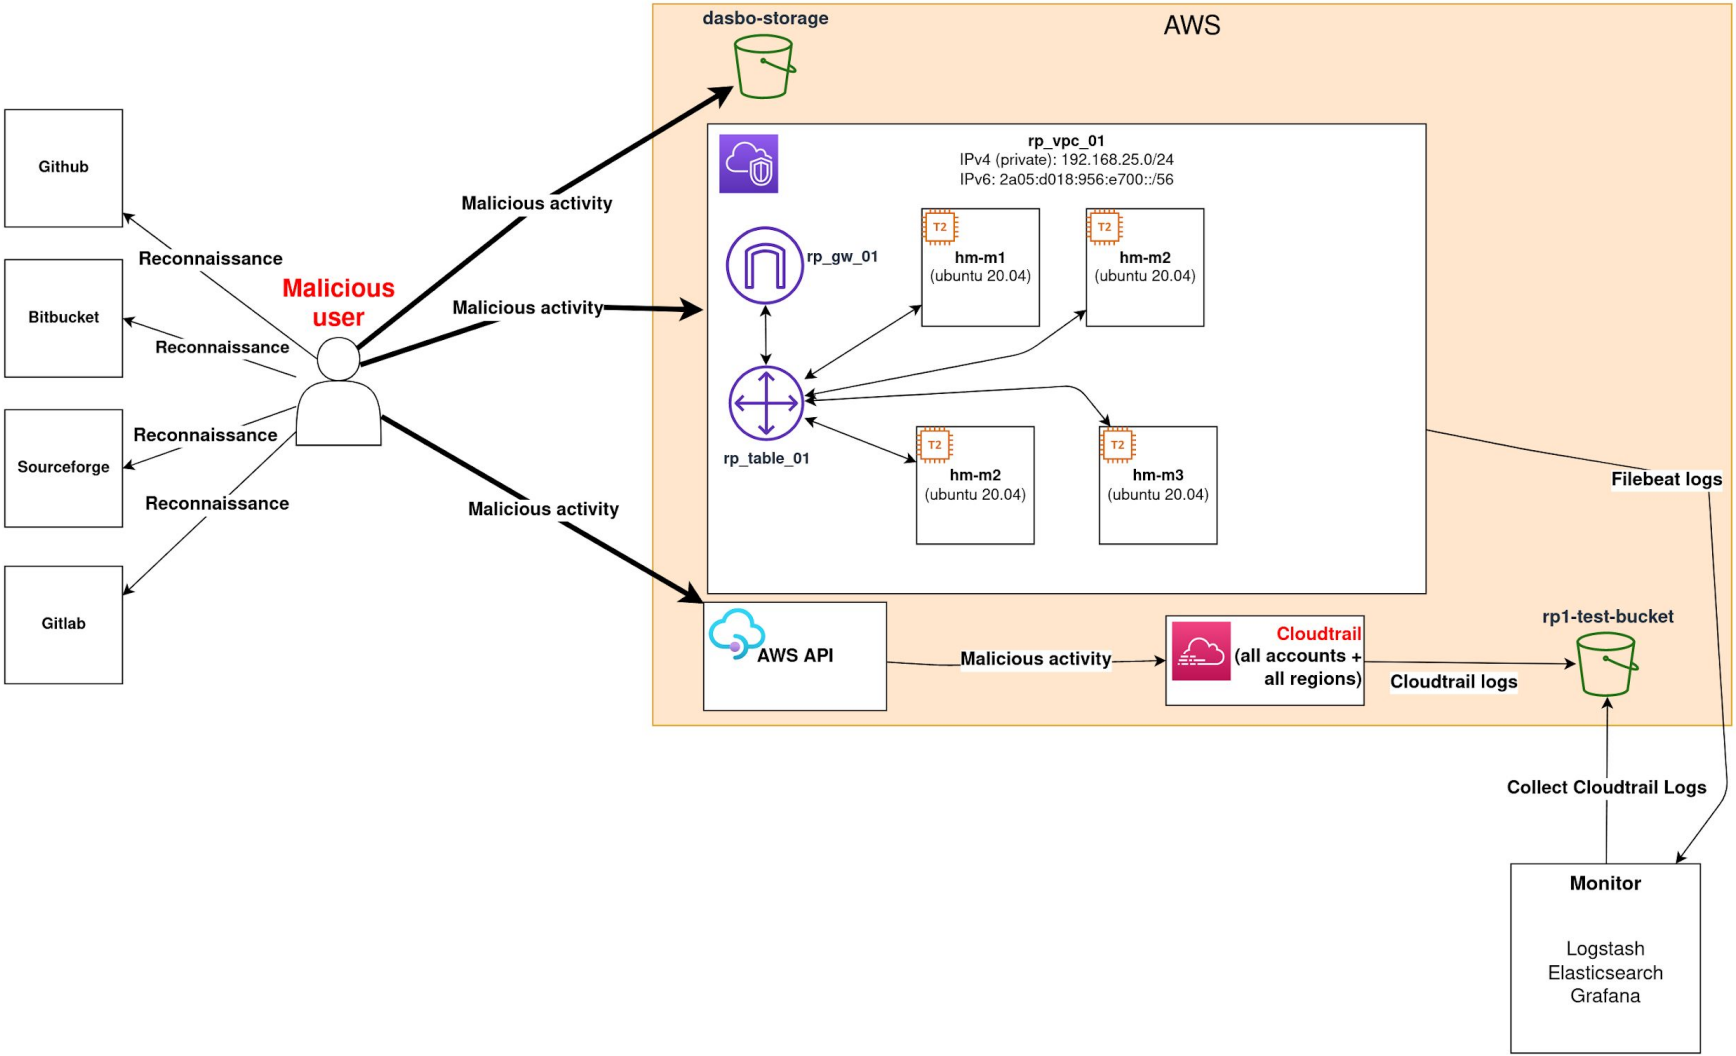
dasbo-storage
AWS
Github
Bitbucket
Sourceforge
Gitlab
Reconnaissance
Reconnaissance
Reconnaissance
Reconnaissance
Malicious user
Malicious activity
Malicious activity
Malicious activity
rp_vpc_01
IPv4 (private): 192.168.25.0/24
IPv6: 2a05:d018:956:e700::/56
rp_gw_01
rp_table_01
hm-m1
(ubuntu 20.04)
hm-m2
(ubuntu 20.04)
hm-m2
(ubuntu 20.04)
hm-m3
(ubuntu 20.04)
Filebeat logs
AWS API
Malicious activity
Cloudtrail
(all accounts +
all regions)
Cloudtrail logs
rp1-test-bucket
Collect Cloudtrail Logs
Monitor
Logstash
Elasticsearch
Grafana

Methods: Experiment-2

- Create 5 new Github, Gitlab, Bitbucket and Sourceforge accounts and 20 new AWS (IAM) accounts with read access rights to multiple services in all regions.

- Create new repositories on all platforms and publish code containing clear-text passwords and AWS tokens, use separate commits for AWS tokens/passwords. The AWS tokens where stored as environment variables.

```
aws_access_key_id="AKIASAI5KDNJDZNLXOQW"
aws_secret_access_key="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"
```

- 10x 72-hours publish times before expiring.

Results - Experiment 1

- 20% of the tokens on Github where found the fastest being 45 minutes and 1 token on sourceforge which happened 56 hours after publication on Sourceforge. None of the other tokens where found.

- None of the hashed MD5 passwords where ever used.

	Github	**Gitlab**	**Bitbucket**	**Sourceforge**
API tokens found	2/10	0/10	0/10	0/10
Canary tokens found	2/10	0/10	0/10	1/10
Percentage found	20%	0%	0%	5%

Results - Experiment 1 (Cont'd)

- The fastest a token was discovered was 2 minutes and 56 seconds.

- All registered attempts made started with the action 'DescribeInstances', most attempts were made using the official aws-sdk-nodejs (version 2) tooling in combination with Tor.

Action	**Used * times**	**Mitre ATT&CK technique**
DescribeInstances	7	Cloud Infrastructure Discovery
GetAccountAuthorizationDetails	2	Account Discovery
N.A.	4	N.A.

Results - Experiment 2

- 100% of tokens published on Github where found, none of the other tokens where found.

- None of the cleartext passwords where ever used.

	Github	**Gitlab**	**Bitbucket**	**Sourceforge**
API tokens found	10/10	0/10	0/10	0/10
Percentage found	100%	0%	0%	0%

Results - Experiment 2 (Cont'd)

- Most initial attempts started with the action 'DescribeInstances', if this worked automation kicked in, attempting to execute 'Runinstance' over varying durations. Few attempts to consolidate access.

Action	**Used * times**	**Mitre ATT&CK technique**
DescribeInstances	43	Cloud Infrastructure Discovery
DescribeSubnets	25	Cloud Infrastructure Discovery
DescribeVpcs	25	Cloud Infrastructure Discovery
RunInstances	1076	Modify Cloud Compute Infrastructure
S3 Bucket (multiple Actions)	54	Cloud Service Discovery

Discussion

- The exposure of a given repository, the size of a commit and the ‘keywords’ used to identify tokens seem to increase/limit the detection of AWS tokens.

- Github seems to be the platform attracting most attackers, which could be due to the search API.

Conclusions

How are leaked credentials/secrets on code collaboration platforms like github, gitlab, sourceforge and bitbucket found and abused by malicious users in a cloud platforms like AWS?

- Credentials are rarely found except on Github, the most likely reason for this is the REST API available for Github.

- Most attempts start with a Discovery technique using ‘DescribeInstances’.

- The second request varied depending on the response this could be additional Discovery actions like ‘GetAccountAuthorizationDetails’ or ‘RunInstances’ depending on the response of the first action.

Conclusions (Cont'd)

- The main motivation seems to be cryptomining the aws-sdk-js[1] tool seems to be favored in combination with Tor as this was used with most of the tokens that were used.

- Mitigation and limiting the actions malicious users can take could be achieved via:
  - AWS organization with Service Control Policies in combination with managed AWS accounts.
  - IAM users with specific policies only allows actions they need to be able to do.
  - Use tools like Vault to store secrets.
  - Use (pre-)commit hooks in combination with tools like shhgit or aws secret-scanning.

[1] https://github.com/aws/aws-sdk-js

Future Work

- More ‘established’ repositories and a longer period of time could be used to expose AWS credentials. This could result into getting better data on the actions a malicious user takes once they have access to a AWS account.

- Repeat the experiments after a period of time (e.g. 1 year).

- Less strict security policies could be given to a leaked honey tokens to get more/better data on the actions a malicious user takes.

Bonus Slide

Demo Grafana

Demo Jenkins

Link to CloMo:

https://github.com/Mandorath/CloMo

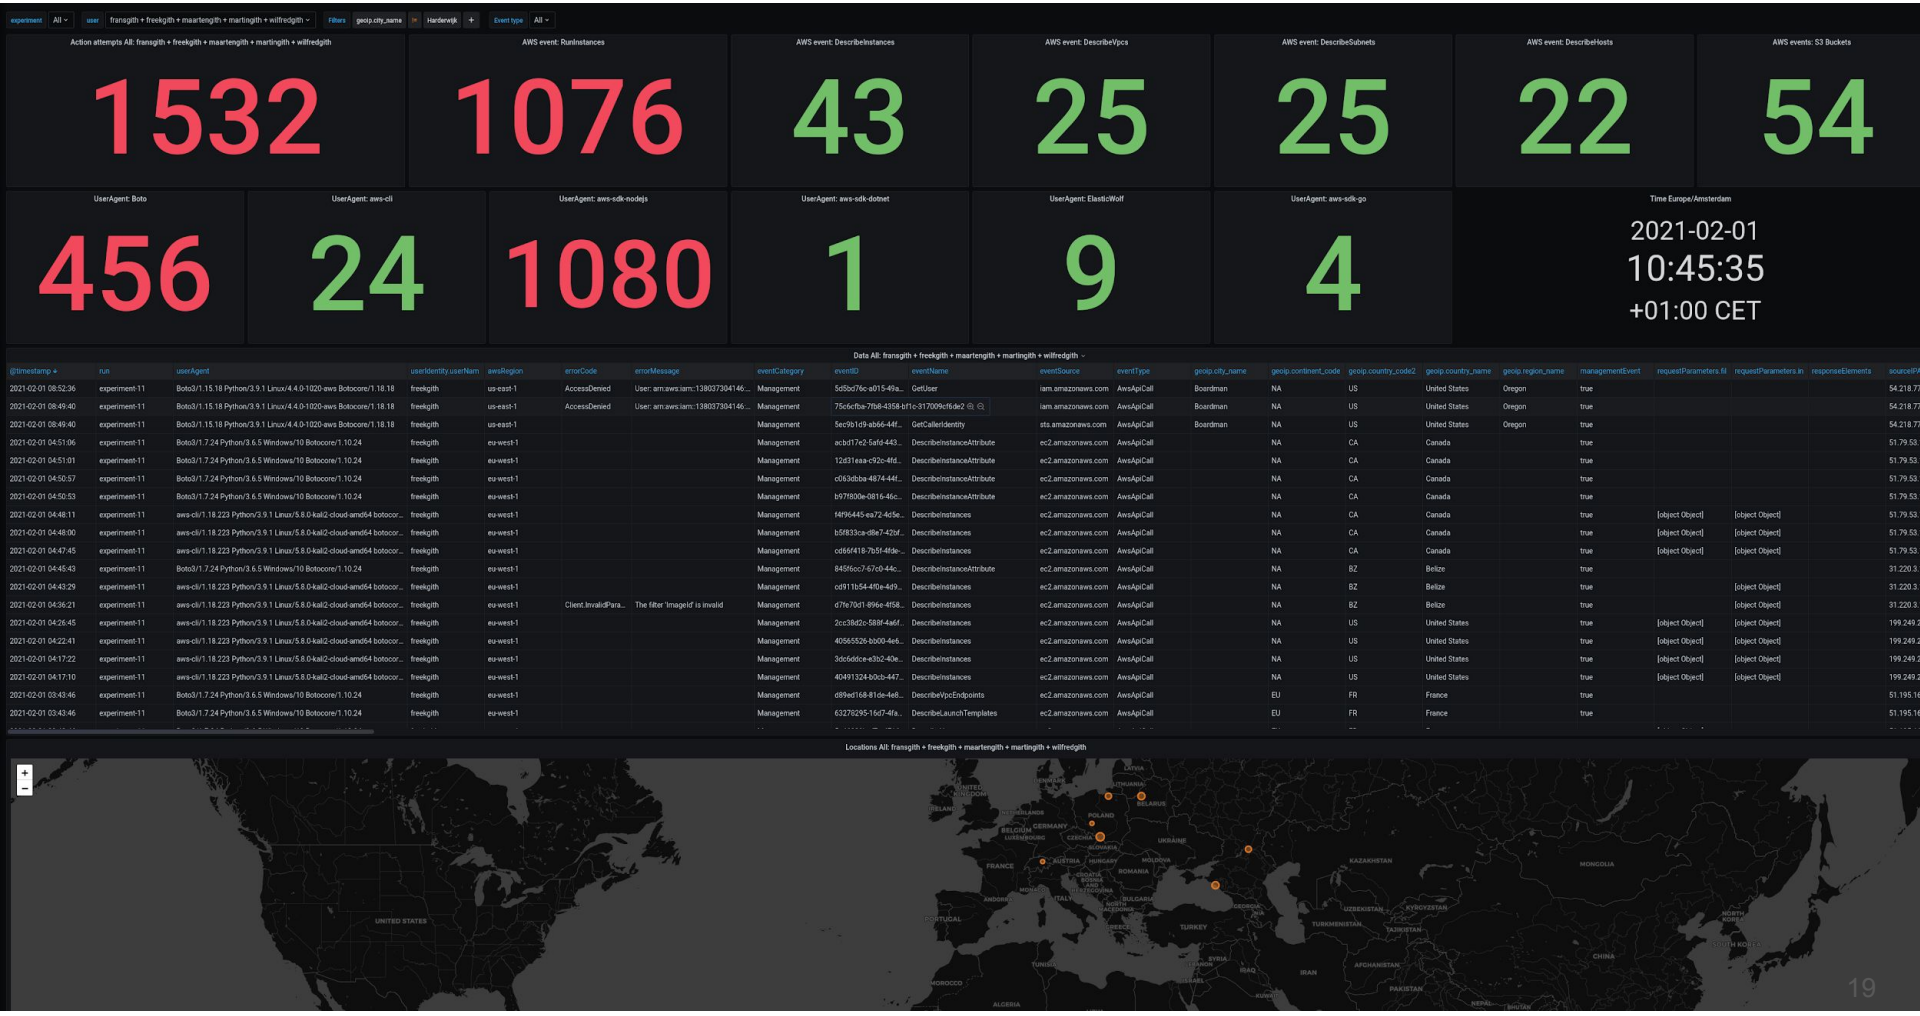
experiment All
user fransgith + freekgith + maartengith + martingith + wilfredgith
Filters geoip.city_name Harderwijk
Event type All
Action attempts All: fransgith + freekgith + maartengith + martingith + wilfredgith
1532
AWS event: RunInstances
1076
AWS event: DescribeInstances
43
AWS event: DescribeVpcs
25
AWS event: DescribeSubnets
25
AWS event: DescribeHosts
22
AWS events: S3 Buckets
54
UserAgent: Boto
456
UserAgent: aws-cli
24
UserAgent: aws-sdk-nodejs
1080
UserAgent: aws-sdk-dotnet
1
UserAgent: ElasticWolf
9
UserAgent: aws-sdk-go
4
Time Europe/Amsterdam
2021-02-01
10:45:35
+01:00 CET
Data All: fransgith + freekgith + maartengith + martingith + wilfredgith
@timestamp
run
userAgent
userIdentity.userNam
awsRegion
errorCode
errorMessage
eventCategory
eventID
eventName
eventSource
eventType
geoip.city_name
geoip.continent_code
geoip.country_code2
geoip.country_name
geoip.region_name
managementEvent
requestParameters.fil
requestParameters.in
responseElements
sourceIPA
2021-02-01 08:52:36 experiment-11 Boto3/1.15.18 Python/3.9.1 Linux/4.4.0-1020-aws Botocore/1.18.18 freekgith us-east-1 AccessDenied User: arn:aws:iam::138037304146:... Management 5d5bd76c-e015-49a... GetUser iam.amazonaws.com AwsApiCall Boardman NA US United States Oregon true 54.218.77
2021-02-01 08:49:40 experiment-11 Boto3/1.15.18 Python/3.9.1 Linux/4.4.0-1020-aws Botocore/1.18.18 freekgith us-east-1 AccessDenied User: arn:aws:iam::138037304146:... Management 75c6cfba-7fb8-4358-bf1c-317009cf6de2 iam.amazonaws.com AwsApiCall Boardman NA US United States Oregon true 54.218.77
2021-02-01 08:49:40 experiment-11 Boto3/1.15.18 Python/3.9.1 Linux/4.4.0-1020-aws Botocore/1.18.18 freekgith us-east-1 Management 5ec9b1d9-ab66-44f... GetCallerIdentity sts.amazonaws.com AwsApiCall Boardman NA US United States Oregon true 54.218.77
2021-02-01 04:51:06 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management acbd17e2-5afd-443... DescribeInstanceAttribute ec2.amazonaws.com AwsApiCall NA CA Canada true 51.79.53
2021-02-01 04:51:01 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management 12d31eaa-c92c-4fd... DescribeInstanceAttribute ec2.amazonaws.com AwsApiCall NA CA Canada true 51.79.53
2021-02-01 04:50:57 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management c063dbba-4874-44f... DescribeInstanceAttribute ec2.amazonaws.com AwsApiCall NA CA Canada true 51.79.53
2021-02-01 04:50:53 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management b97f800e-0816-46c... DescribeInstanceAttribute ec2.amazonaws.com AwsApiCall NA CA Canada true 51.79.53
2021-02-01 04:48:11 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management f4f96445-ea72-4d5e... DescribeInstances ec2.amazonaws.com AwsApiCall NA CA Canada true [object Object] [object Object] 51.79.53
2021-02-01 04:48:00 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management b5f833ca-d8e7-42bf... DescribeInstances ec2.amazonaws.com AwsApiCall NA CA Canada true [object Object] [object Object] 51.79.53
2021-02-01 04:47:45 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management cd66f418-7b5f-4fde-... DescribeInstances ec2.amazonaws.com AwsApiCall NA CA Canada true [object Object] [object Object] 51.79.53
2021-02-01 04:45:43 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management 845f6cc7-67c0-44c... DescribeInstanceAttribute ec2.amazonaws.com AwsApiCall NA BZ Belize true 31.220.3
2021-02-01 04:43:29 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management cd911b54-4f0e-4d9... DescribeInstances ec2.amazonaws.com AwsApiCall NA BZ Belize true [object Object] 31.220.3
2021-02-01 04:36:21 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Client.InvalidPara... The filter 'ImageId' is invalid Management d7fe70d1-896e-4f58... DescribeInstances ec2.amazonaws.com AwsApiCall NA BZ Belize true [object Object] 31.220.3
2021-02-01 04:26:45 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management 2cc38d2c-589f-4a6f... DescribeInstances ec2.amazonaws.com AwsApiCall NA US United States true [object Object] [object Object] 199.249.2
2021-02-01 04:22:41 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management 405655526-b000-4e6... DescribeInstances ec2.amazonaws.com AwsApiCall NA US United States true [object Object] [object Object] 199.249.2
2021-02-01 04:17:22 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management 3dc6ddce-e3b2-40e... DescribeInstances ec2.amazonaws.com AwsApiCall NA US United States true [object Object] [object Object] 199.249.2
2021-02-01 04:17:10 experiment-11 aws-cli/1.18.223 Python/3.9.1 Linux/5.8.0-kali2-cloud-amd64 botocor... freekgith eu-west-1 Management 40491324-b0cb-447... DescribeInstances ec2.amazonaws.com AwsApiCall NA US United States true [object Object] [object Object] 199.249.2
2021-02-01 03:43:46 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management d89ed168-81de-4e8... DescribeVpcEndpoints ec2.amazonaws.com AwsApiCall EU FR France true 51.195.14
2021-02-01 03:43:46 experiment-11 Boto3/1.7.24 Python/3.6.5 Windows/10 Botocore/1.10.24 freekgith eu-west-1 Management 63278295-15d7-4fa... DescribeLaunchTemplates ec2.amazonaws.com AwsApiCall EU FR France true 51.195.14
Locations All: fransgith + freekgith + maartengith + martingith + wilfredgith
UNITED STATES
CANADA
MEXICO
RUSSIA
KAZAKHSTAN
MONGOLIA
CHINA
IRAN
TURKEY
FRANCE
GERMANY
POLAND
UKRAINE
BELARUS
LITHUANIA
LATVIA
SWEDEN
FINLAND
NORWAY
UNITED KINGDOM
SPAIN
PORTUGAL
ITALY
GREECE
ROMANIA
BULGARIA
HUNGARY
AUSTRIA
CZECHIA
BELGIUM
NETHERLANDS
IRELAND
DENMARK
MOROCCO
ALGERIA
TUNISIA
LIBYA
EGYPT
SAUDI ARABIA
IRAQ
SYRIA
ISRAEL
LEBANON
AFGHANISTAN
PAKISTAN
INDIA
NEPAL
BHUTAN
UZBEKISTAN
TURKMENISTAN
KYRGYZSTAN
TAJIKISTAN
NORTH KOREA
SOUTH KOREA