Real world detection engineering in a Multi-Cloud environment

Aaron Jewitt

elastic

Detection Engineer
Elastic Infosec

Aaron Jewitt
@acjewitt@infosec.exchange
linkedin.com/in/aaron-jewitt/

Agenda

00 What is ‘Real World’ Detection Engineering?

01 Protecting a distributed workforce

02 Dangers in the cloud

03 Detecting real world cloud attacks

What do you mean by ‘Real World’ Detection Engineering?

3 types of Detection Engineers

Detection Engineer working for a Security Product company

- Creating the default detection rules for security products
- Endpoint rules are very tightly tuned to prevent false positives

Detection Engineer working for a Threat Intel company

- Creating detection rules and IOCs to detect specific threat actors
- They will detect threat actor behavior, but are often noisy

Detection Engineer working in a SOC - ‘Real World’

- Creating and tuning SIEM detection rules for a specific environment
- Responsible for balancing false positive noise with ability to detect real world attacks
- You get blamed if the rule is too noisy, and if it is too quiet

The ‘Real World’ can be messy

- You have limited resources
- Security is often secondary to business needs
- API keys and tokens are hard to manage
- You have Shadow IT, emergency fixes, large customers with one-off requirements, and more
- Software baselines change constantly
- Legitimate software will use the same TTPs as malware
- Each Cloud and SaaS provider has very different logs

My Core Principles of Real World Detection Engineering

They help manage the mess of the real world

1. An alert must be actionable
2. Your rules should be tightly focused
3. Provide as much alert context as possible
4. You will never achieve a zero FP rate, but that doesn’t mean you shouldn’t try

1: An alert must be actionable!

Every time an alert is placed in front of a human, action should be taken!

If you don’t have a policy preventing the action, document that risk and disable the notification

Low severity rules can still be used by other alerts and investigations

https://www.elastic.co/blog/detection-engineering-maximizing-analyst-efficiency-using-cardinality-threshold-rules-on-your-alerts-

2: Your rules should be tightly focused

Don’t make one rule where you could make two. Don’t make two where you could make four.

Rules that try to detect everything are a pain to investigate and prone to different types of false positives

A tightly focused rule can be more quickly investigated and is easier to automate

3: Provide as much context as possible

Include all important fields from the alert and enrich from internal and external sources

Get all of the information needed to triage an alert in front of the analyst, don’t make them hunt

- Check threat intel for process and file hashes
- Check for other alerts in the last 7d from the same host and user
- Enrich with asset info to include the system owner

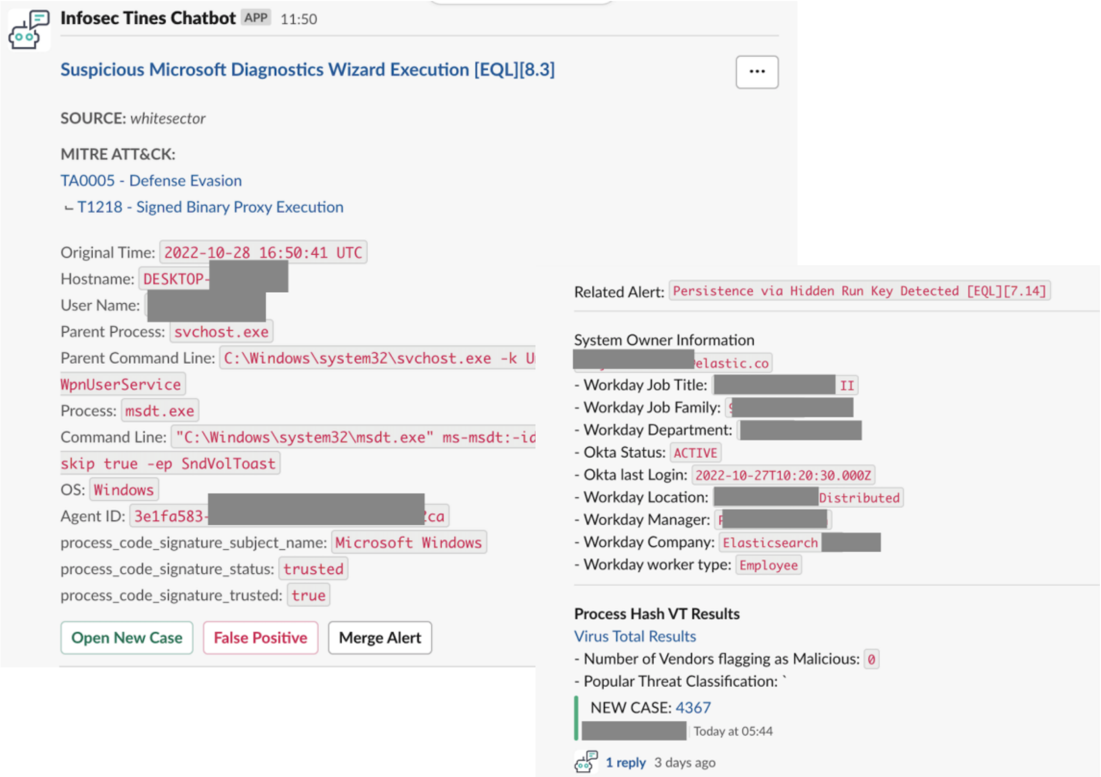

Infosec Tines Chatbot APP 11:50

Suspicious Microsoft Diagnostics Wizard Execution [EQL][8.3]

SOURCE: *whitesector*

MITRE ATT&CK:
TA0005 - Defense Evasion
⌙ T1218 - Signed Binary Proxy Execution

Original Time: `2022-10-28 16:50:41 UTC`
Hostname: `DESKTOP-`[redacted]
User Name: [redacted]
Parent Process: `svchost.exe`
Parent Command Line: `C:\Windows\system32\svchost.exe -k U`
`WpnUserService`
Process: `msdt.exe`
Command Line: `"C:\Windows\system32\msdt.exe" ms-msdt:-id`
`skip true -ep SndVolToast`
OS: `Windows`
Agent ID: `3e1fa583-`[redacted]`ca`
process_code_signature_subject_name: `Microsoft Windows`
process_code_signature_status: `trusted`
process_code_signature_trusted: `true`

Related Alert: `Persistence via Hidden Run Key Detected [EQL][7.14]`

System Owner Information
[redacted]`elastic.co`
- Workday Job Title: [redacted]`II`
- Workday Job Family: [redacted]
- Workday Department: [redacted]
- Okta Status: `ACTIVE`
- Okta last Login: `2022-10-27T10:20:30.000Z`
- Workday Location: [redacted]`Distributed`
- Workday Manager: [redacted]
- Workday Company: `Elasticsearch`[redacted]
- Workday worker type: `Employee`

Process Hash VT Results
Virus Total Results
- Number of Vendors flagging as Malicious: `0`
- Popular Threat Classification: `

NEW CASE: 4367
[redacted] Today at 05:44

1 reply 3 days ago

4: You will never have zero False Positives

But that doesn't mean you shouldn't try

Every week you should dedicate time tuning your most noisy detection rules

If an alert keeps going off and no action is taken, it may be time to turn off that alert

Don't let perfect get in the way of good

Protecting a distributed workforce

Elastic has a Globally Distributed workforce

Globally Distributed Workforce

>2,000

ELASTICIANS

>42

COUNTRIES

Cloud Native Implementation

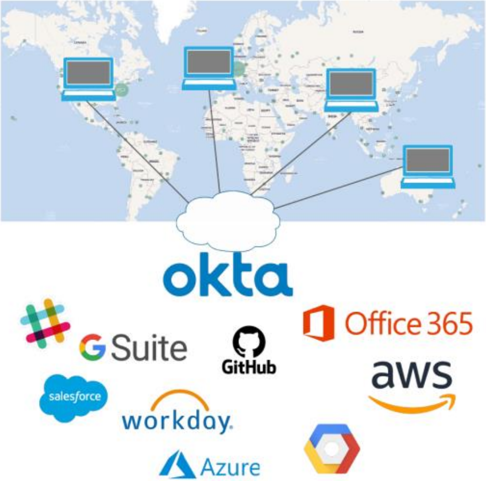

InfoSec by the numbers (Daily)

150TB

Security Data

Enables us to monitor for abnormal and security relevant activity

500TB

Observability Data

Proxy logs, performance logs, and other log sources that come in handy during an investigation

>500K

Endpoints

Globally dispersed cloud instances, virtual desktop environments, and user workstations

How do you monitor and protect a globally distributed, cloud SaaS company?

- Require FIDO2 phishing resistant SSO
- Protect your workstations with a high quality EDR and OSQuery
- Collect process, network, login, auditd, and file events from all workstations, server and container systems with Elastic Agent, sysmon, or Auditbeat
- Collect ALL of your third party and service logs
- Normalize the data and make it all searchable from a single location.
- Have a comprehensive Asset Management program
- AUTOMATE ALL THE THINGS!!

Identity is the perimeter

Mandatory FIDO2 MFA in your SSO

- Biometrics on your MacOS, IOS or Android
- Windows Hello
- FIDO2 USB device like Yubikey

All other MFA options can be phished

- Attackers can easily proxy your MFA push notifications or one time code authentications to steal your session tokens

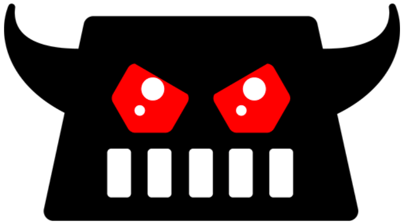

Protect your session tokens

Even FIDO2 MFA isn’t 100%

- Your SSO session tokens are just cookies in your browser
- Pay special attention to workstation alerts about accessing browser cookies

Set the max lifetime of a session to limit damage and risk from stolen cookies

Require re-authentication with FIDO2 on your most critical applications

Make all of your data visible in your SIEM

Don’t limit yourself to just ‘security’ alerts. Get the 3rd party logs, Asset Management, Vulnerability Management, and other data in one place

But don’t pay for shipping! Keep the data where it is created and use Cross Cluster Search (CCS) to let the SIEM search the events

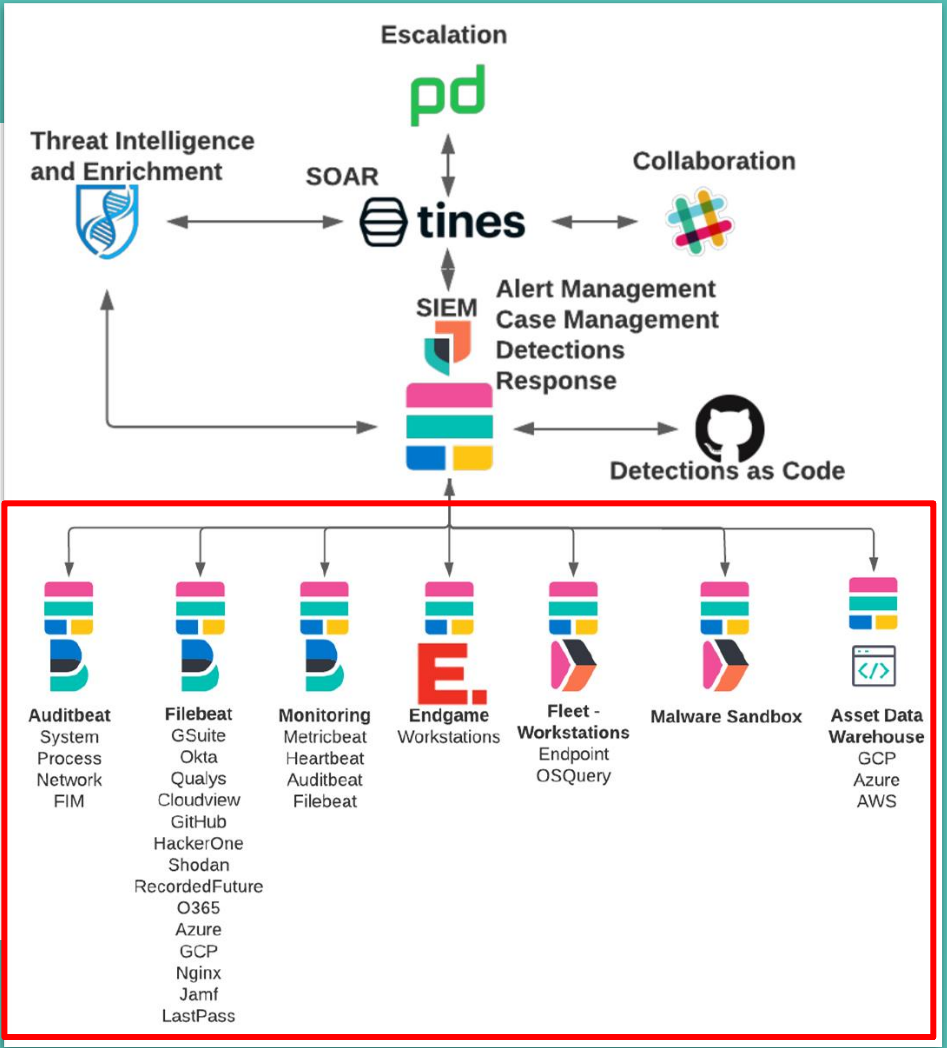

Standardize your data field names

Data like this is a pain to search through during a stressful response

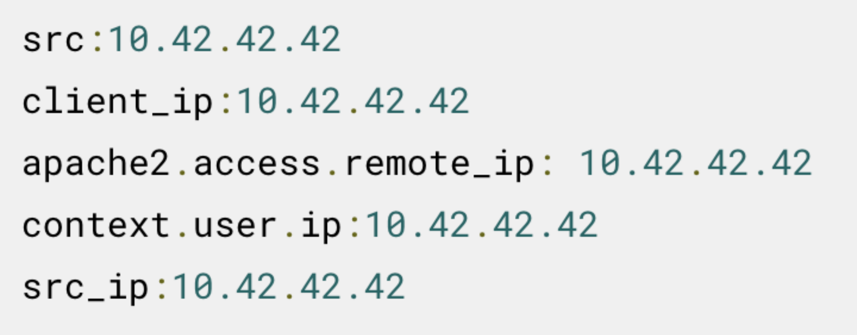

```
src:10.42.42.42
client_ip:10.42.42.42
apache2.access.remote_ip: 10.42.42.42
context.user.ip:10.42.42.42
src_ip:10.42.42.42
```

This is so much easier to build detections and dashboards

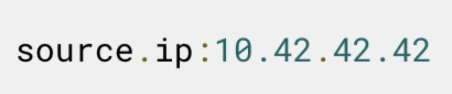

```
source.ip:10.42.42.42
```

Elastic Common Schema (ECS) is the default for all Elastic integrations. It is open source and it ‘just works’.

Automation

Because ain’t nobody got time for that!

Get the alerts out of your SIEM and to the people that can take action on the data

Enriching alerts speeds up triage time. Sometimes you can completely automate the investigation.

You can build almost anything with webhooks and APIs…

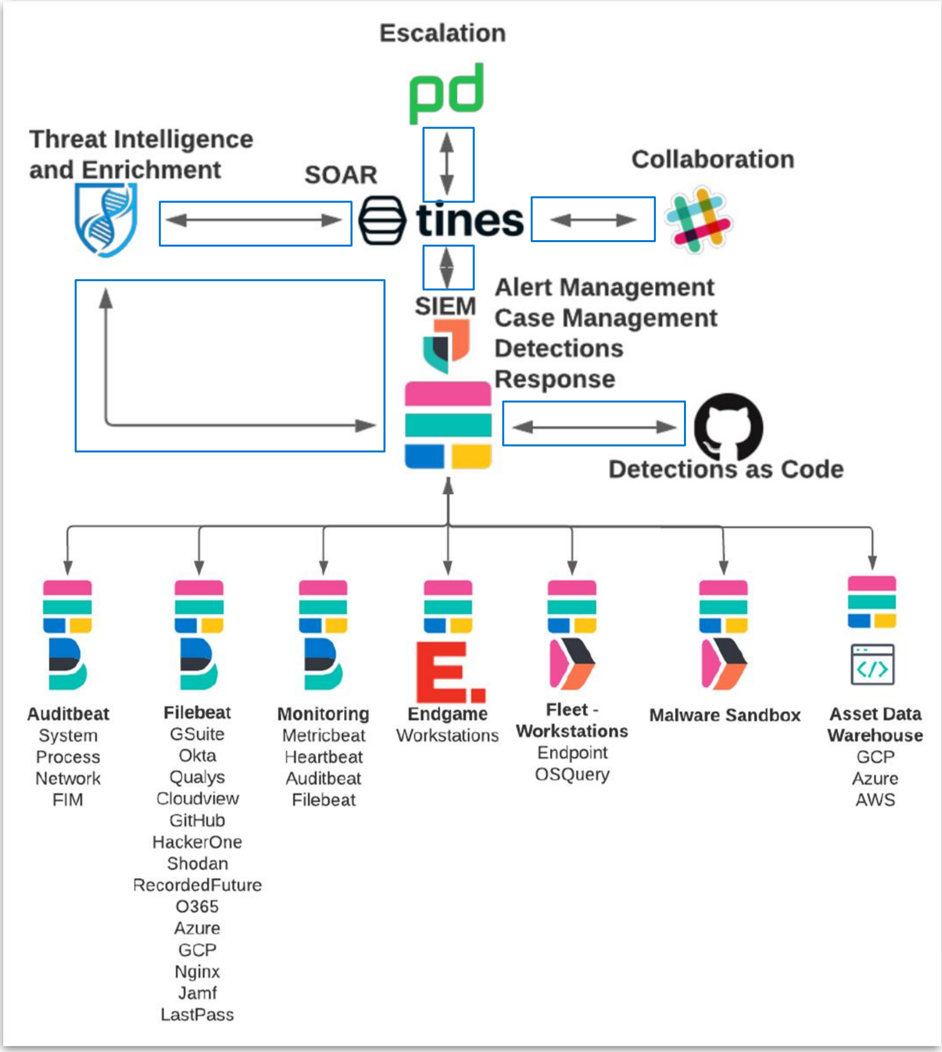

Automate Everything!

- Enrichment of high severity alerts in Slack
- Distribution of alerts to system owners via interactive Slack notifications
- Automated triage and closing or escalation of alerts
- Run OSQuery commands following a malware alert
- Quarantine cloud VMs by updating the VPC settings

Alert Context Automation (TriageBot)

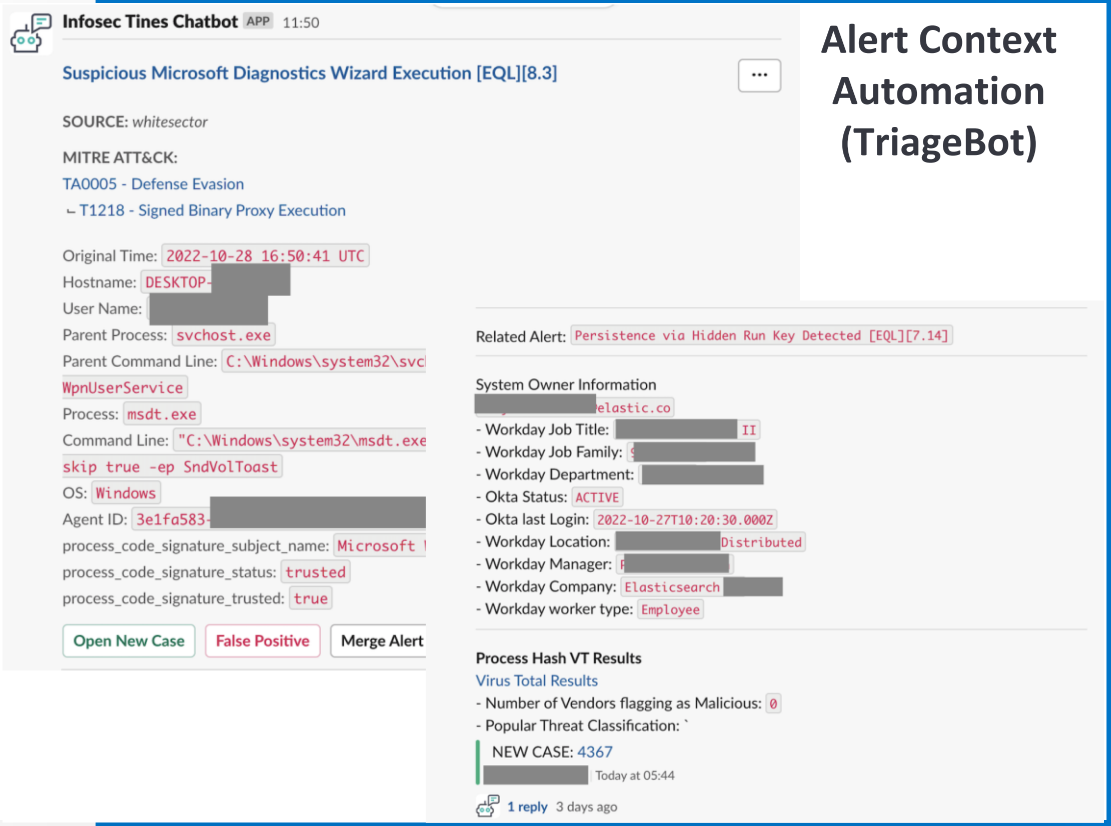

Infosec Tines Chatbot 11:50

Suspicious Microsoft Diagnostics Wizard Execution [EQL][8.3]

SOURCE: *whitesector*

MITRE ATT&CK:
TA0005 - Defense Evasion
⌞ T1218 - Signed Binary Proxy Execution

Original Time: 2022-10-28 16:50:41 UTC
Hostname: DESKTOP-[redacted]
User Name: [redacted]
Parent Process: svchost.exe
Parent Command Line: C:\Windows\system32\svc
WpnUserService
Process: msdt.exe
Command Line: "C:\Windows\system32\msdt.exe
skip true -ep SndVolToast
OS: Windows
Agent ID: 3e1fa583-[redacted]
process_code_signature_subject_name: Microsoft
process_code_signature_status: trusted
process_code_signature_trusted: true

Related Alert: Persistence via Hidden Run Key Detected [EQL][7.14]

System Owner Information
[redacted]elastic.co
- Workday Job Title: [redacted] II
- Workday Job Family: [redacted]
- Workday Department: [redacted]
- Okta Status: ACTIVE
- Okta last Login: 2022-10-27T10:20:30.000Z
- Workday Location: [redacted]Distributed
- Workday Manager: [redacted]
- Workday Company: Elasticsearch [redacted]
- Workday worker type: Employee

Process Hash VT Results
Virus Total Results
- Number of Vendors flagging as Malicious: 0
- Popular Threat Classification: `

NEW CASE: 4367
[redacted] Today at 05:44

1 reply 3 days ago

You don’t have time to triage every alert

- Automate the investigation and response completely if you can
- Distribute it to the system owners if possible
- If you can’t automate it, maybe you can speed it up with enrichment
- Set some alerts to low severity and then build alerts with your alerts
- Use Threshold rules to alert on multiple different alerts on a single entity: www.elastic.co/blog/detection-engineering-maximizing-analyst-efficiency-using-cardinality-threshold-rules-on-your-alerts-
- Use Indicator Match rules to escalate alerts from low to high based on the system owner or role of the user
- New Terms alerts to see new alerts that haven’t triggered in 30d

Asset Inventory

You can’t secure something if you don't know it exists!

Tell me the truth...I'm..I'm ready to hear it.
You need an accurate inventory of your assets before you can work on protecting it.

The Big Picture

We are building an inventory of all assets across Elastic, using native features of the Elastic stack (devices, apps, people, accounts, cloud resources, …)

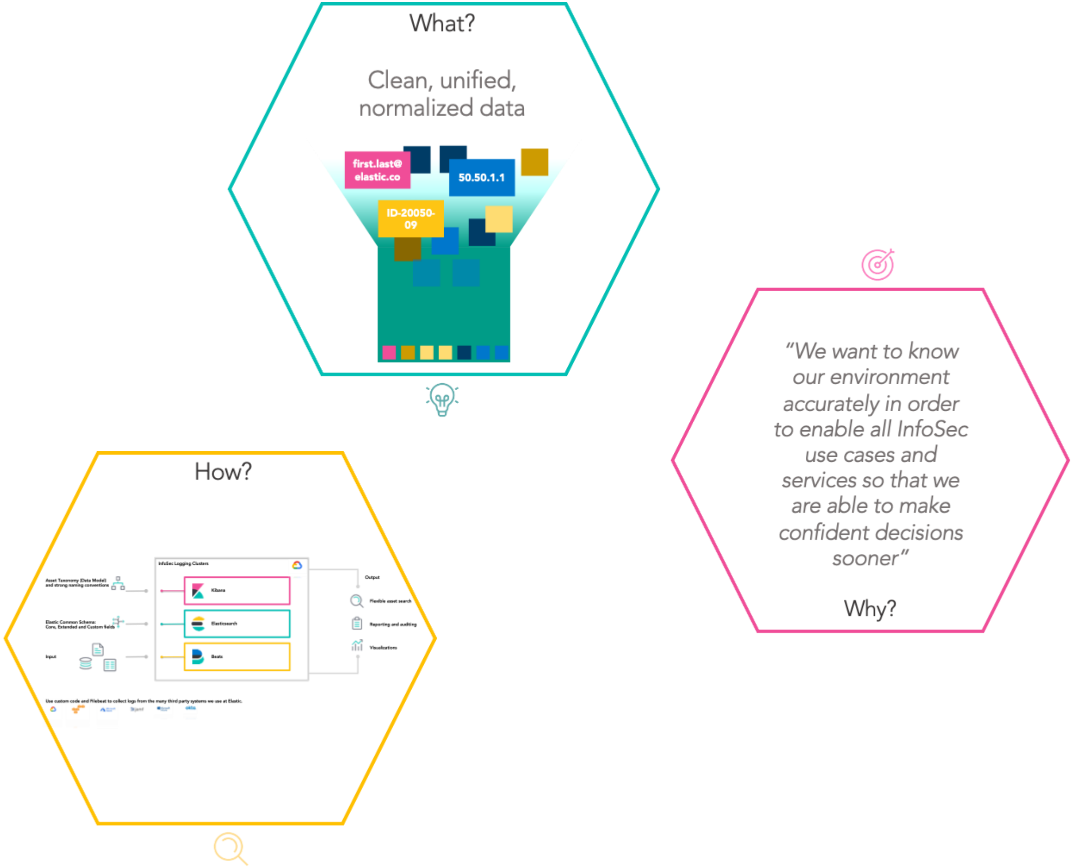

What?

Clean, unified, normalized data

How?

Why?

"We want to know our environment accurately in order to enable all InfoSec use cases and services so that we are able to make confident decisions sooner"

first.last@ elastic.co
50.50.1.1
ID-20050-09
InfoSec Logging Clusters
Asset Taxonomy (Data Model) and strong naming conventions
Elastic Common Schema: Core, Extended and Custom fields
Input
Kibana
Elasticsearch
Beats
Output
Flexible asset search
Reporting and auditing
Visualizations
Use custom code and Filebeat to collect logs from the many third party systems we use at Elastic.
okta

https://www.elastic.co/blog/how-to-build-a-cybersecurity-asset-management-solution-on-the-elastic-stack

High Level Architecture

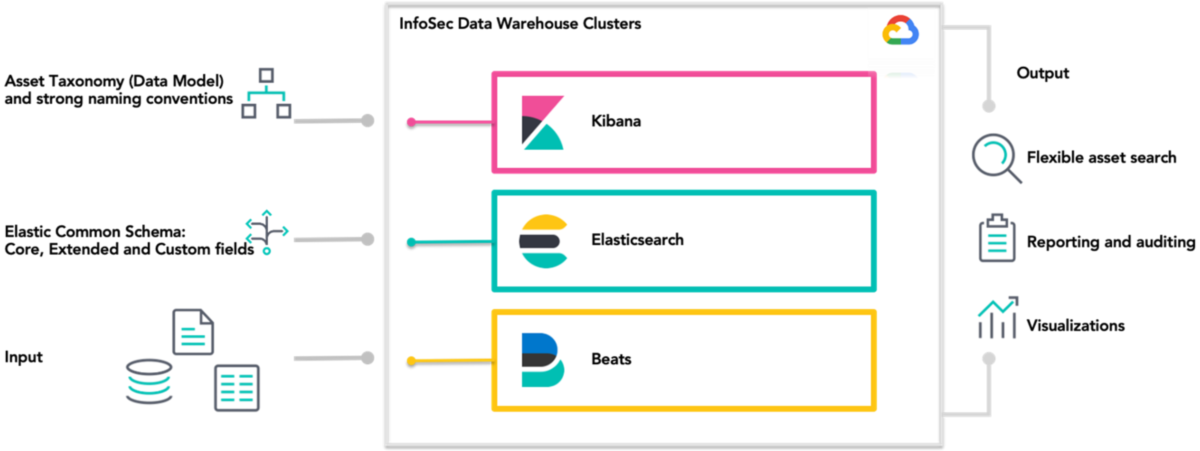

InfoSec Data Warehouse Clusters

Asset Taxonomy (Data Model) and strong naming conventions

Elastic Common Schema: Core, Extended and Custom fields

Input

Kibana

Elasticsearch

Beats

Output

Flexible asset search

Reporting and auditing

Visualizations

Use custom code and Filebeat to collect logs from the many third party systems we use at Elastic.

GitHub

Suite

Microsoft Azure

jamf

okta

Microsoft Intune

Asset Inventory with Detection rules

- The Asset Inventory can be used as an indicator index to alert on activity on critical assets or users
- Create Indicator match Detection rules that looks at all alerts and compares them to the Asset inventory
  - Alerts on an Admin account
  - Alerts on a Service Account
  - Alerts on a critical asset
  - Users outside of a specific team accessing that teams resources

Asset Inventory Enrichment

- Automate Enriching alerts with the Asset Inventory
- Use the host.id field in an alert to match workstation info
- Use the user.email or the user.name fields to gather data about work role, location, github username and more
- Enrich the Slack alert with the system owner information.

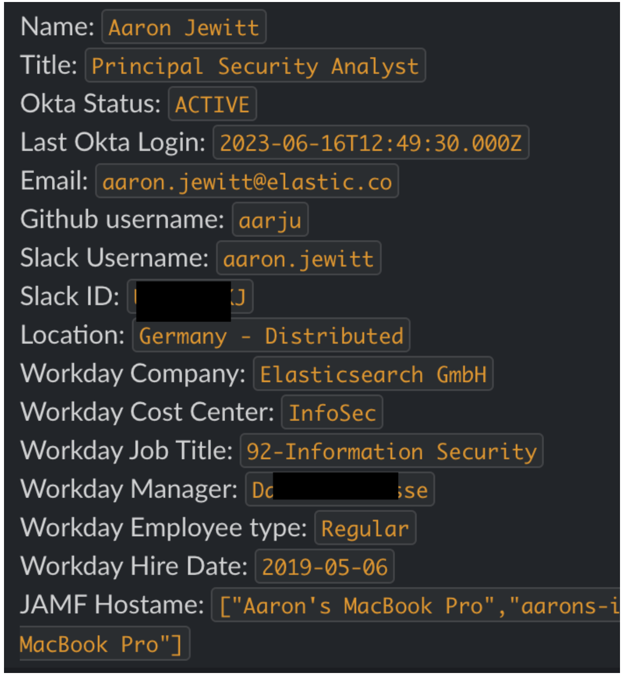

Name: Aaron Jewitt
Title: Principal Security Analyst
Okta Status: ACTIVE
Last Okta Login: 2023-06-16T12:49:30.000Z
Email: aaron.jewitt@elastic.co
Github username: aarju
Slack Username: aaron.jewitt
Slack ID: [redacted]
Location: Germany - Distributed
Workday Company: Elasticsearch GmbH
Workday Cost Center: InfoSec
Workday Job Title: 92-Information Security
Workday Manager: [redacted]
Workday Employee type: Regular
Workday Hire Date: 2019-05-06
JAMF Hostame: ["Aaron's MacBook Pro","aarons-i
MacBook Pro"]

Dangers in the cloud

What is the ‘Cloud’?

- Linux servers mostly
- Infrastructure as code, CI/CD systems
- Shadow IT and Tech Debt
- Interconnections that nobody really understands
- Lots of constantly changing technologies
- It’s all controlled by APIs with Single factor authentication
- The fastest way to give your $$$ to Amazon, Microsoft, and Google

What are a Detection Engineers three biggest Challenges in the cloud?

- Scale of the attack surface
- Speed of potential attacks
- Complexity of the environment

How many endpoints does a SaaS cloud provider have?

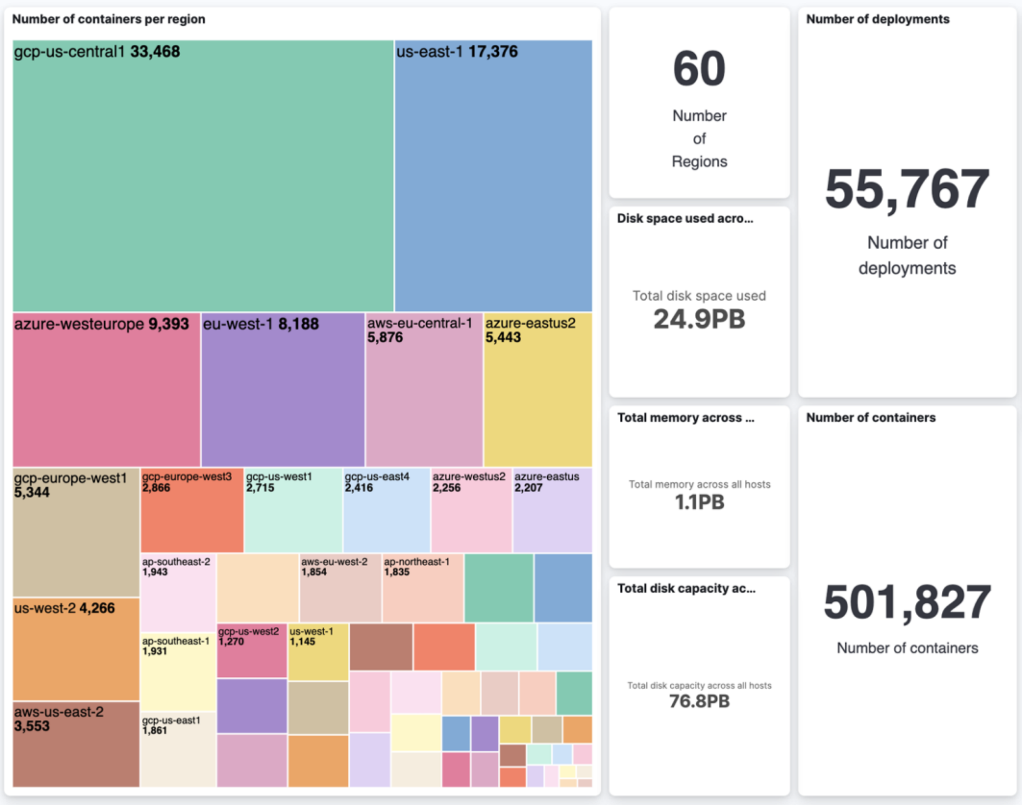

60

Number of Regions

Disk space used acro...

Total disk space used

24.9PB

Number of deployments

55,767

Number of deployments

Total memory across ...

Total memory across all hosts

1.1PB

Number of containers

501,827

Number of containers

Total disk capacity ac...

Total disk capacity across all hosts

76.8PB

Log collection in the Cloud

Challenge - scale

- Centralized logging only for critical security logs such as Cloud Trail, GCP Audit, Azure, and authentication logging
- Every cloud region contains its own monitoring cluster for collecting Observability and Security logs from that region
    - Netflow, Proxy, Metricbeat, Devops, Auditbeat, and more...
- Everything stays in the region it was created and we use CCS to query it

Cross Cluster Search (CCS) in Elastic Cloud

Because Data Transfer is expensive $$$

- 500TB of observability and security data collected per day
- Transferring 500TB per day to a single region would cost us over $40k per day
- CCS gives us distributed storage with centralized search
- The search query is sent to each remote cluster and only results are returned

Challenge - Speed of attacks

- You have to detect and respond with speed if you want to stop a cloud attack
- If you can write a script that creates or deletes thousands of containers in your cloud account, then so can the attacker
- You don’t have time to sed/grep/awk manually through logs during an attack. There is no time for traditional forensics
- Identify, contain, and eradicate at cloud speeds

Challenge - Complexity of the Cloud

Infrastructure as Code is magic, until something breaks

A single line of code can make big changes to your cloud

Do you know how it all really works? Does anyone?

Read the docs! Talk to your experts!

So many technologies!

Challenge - Complexity of the Cloud

This is the very simplified version…

How do the services interact with each other?

How does authentication work? Where are the admin accounts used?

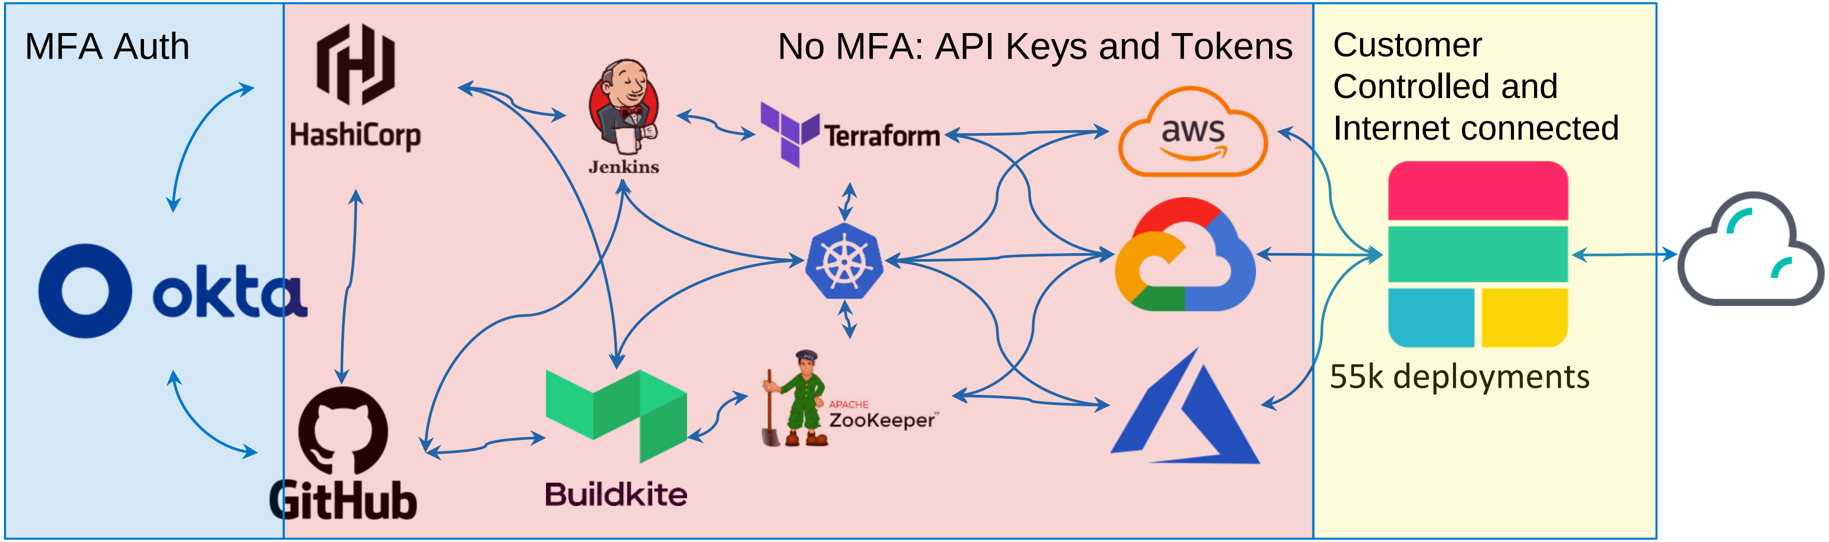

Detecting real world cloud attacks

What are a Detection Engineers biggest fears in the cloud?

- Exploitable services
- Misconfigurations
- Compromised tokens

Anything can be exploited

- Any process that handles user controlled data is vulnerable
- Log4Shell proved that’s not just internet connected services
- When a new 0-day is published the cloud systems are some of the first to be targeted
- People were probably using the 0-day before it went public

Detecting 0-days and Exploited Processes

Exploiting a process is only the first step, look for the follow on actions

Build custom detections for strange activity from all processes that handle untrusted data

- Strange child process
- Strange network connection for a process
- Strange process by a service account
- Rare error codes in cloud logs

Strange Child Process Example: NGINX

Strange Child process of web service

- Nginx is a popular and versatile web server that is often configured to be internet facing
- Start by searching for the normal child processes of nginx

```
process.parent.name:"nginx"
```

- Adjust the search to remove ‘known good’ events and create a Detection Rule for everything else

```
process.parent.name:"nginx"

AND NOT process.parent.title: "/bin/sh /etc/init.d/nginx upgrade"

AND NOT process.parent.title: "/bin/sh /etc/init.d/nginx rotate"

AND NOT process.name: "nginx"
```

Strange curl or wget Process Example

Detect the attack later in the kill chain

- After exploiting a vulnerable process attackers will often use commands such as wget or curl to download additional capabilities
- Lots of legitimate uses for these processes
- Adjust the search to remove ‘known good’ events and create a Detection Rule for everything else
- Create multiple rules for different scenarios
  - Strange parent process of curl
  - wget or curl by a strange user name

Misconfigurations

- S3 buckets filled with sensitive info left open to the public
- Overprivileged accounts
- Insecure automations
- Vulnerable containers left open to the internet
- Containers or pods that aren’t isolated properly

Misconfigurations

- Luckily there are a lot of pre-built detections for these misconfigurations such as creation of open S3 buckets
- Elastic’s Cloud Security Posture Management application can ingest your cloud configuration and show you the risks

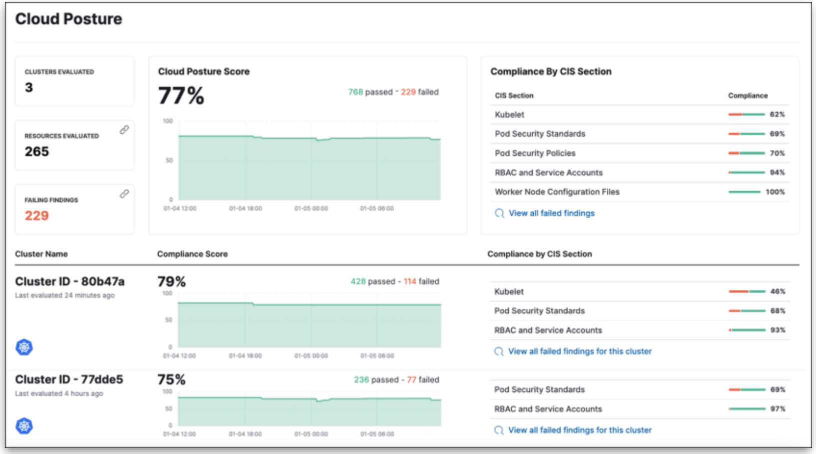

Cloud Posture

CLUSTERS EVALUATED
3

RESOURCES EVALUATED
265

FAILING FINDINGS
229

Cloud Posture Score

77%

768 passed - 229 failed

Compliance By CIS Section

CIS Section	Compliance
Kubelet	62%
Pod Security Standards	69%
Pod Security Policies	70%
RBAC and Service Accounts	94%
Worker Node Configuration Files	100%

View all failed findings

Cluster Name	Compliance Score	Compliance by CIS Section
Cluster ID - 80b47a Last evaluated 24 minutes ago	**79%** 428 passed - 114 failed	Kubelet 46% Pod Security Standards 68% RBAC and Service Accounts 93% View all failed findings for this cluster
Cluster ID - 77dde5 Last evaluated 4 hours ago	**75%** 236 passed - 77 failed	Pod Security Standards 69% RBAC and Service Accounts 97% View all failed findings for this cluster

My Biggest Fear - Compromised Tokens

- How are your tokens being used?
- Where are they stored?
    - Are you sure about that?
- You can usually find them in bash history
- ENV on a running container usually has a few
- They are stored in cleartext files like ~/.ssh or ~/.bashrc
- Sometimes they are written right into a config file or shared via slack or email
- Search your environment for tokens!
- Monitor Pastebin for your tokens!

How do you detect a compromised token?

- This is one of the most challenging things to detect
- Most cloud API tokens can be used from **anywhere in the world**
- Unless they are being destructive, attacker activity will look normal

How do you detect a compromised token?

What does an attacker do after they find a token?

They use the token to authenticate from their own systems. They don’t want to risk triggering an endpoint alert on our systems.

We used our SIEM Detection engine, our Asset DB, and our SOAR to create and automatically triage an alert on API activity from an IP address that doesn’t belong to us.

The goal is to detect any API key activity from an IP that we don’t manage

Detecting activity from unmanaged IPs

Step 1 - Alerting new IP activity

The first step is to create a New Terms detection rule for activity from an IP address not seen in the last 30 days

Use the alert action to send the contents of the alert to your SOAR webhook as a json

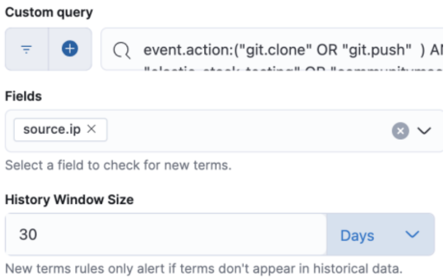

Custom query

event.action:("git.clone" OR "git.push") A

Fields

source.ip

Select a field to check for new terms.

History Window Size

30 Days

New terms rules only alert if terms don't appear in historical data.

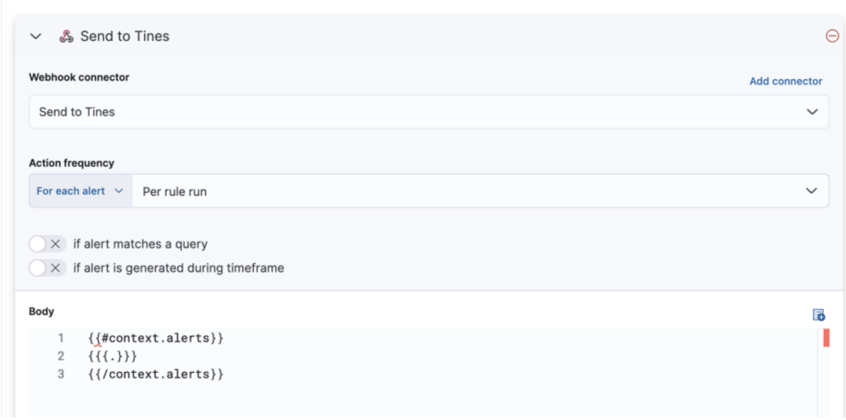

Send to Tines

Webhook connector

Send to Tines

Action frequency

For each alert Per rule run

○ ✕ if alert matches a query

○ ✕ if alert is generated during timeframe

Body

```
{{#context.alerts}}
{{{.}}}
{{/context.alerts}}
```

Detecting activity from unmanaged IPs

Step 2 - SOARing

Within our SOAR we take the source.ip from the alert and check for its existence in several other index patterns

If we find the IP we close the alert because we know we have an agent on that host or it belongs to us in some way

If not found then we escalate the alert

This is also great at finding ‘shadow IT’ and other systems not in our asset inventory

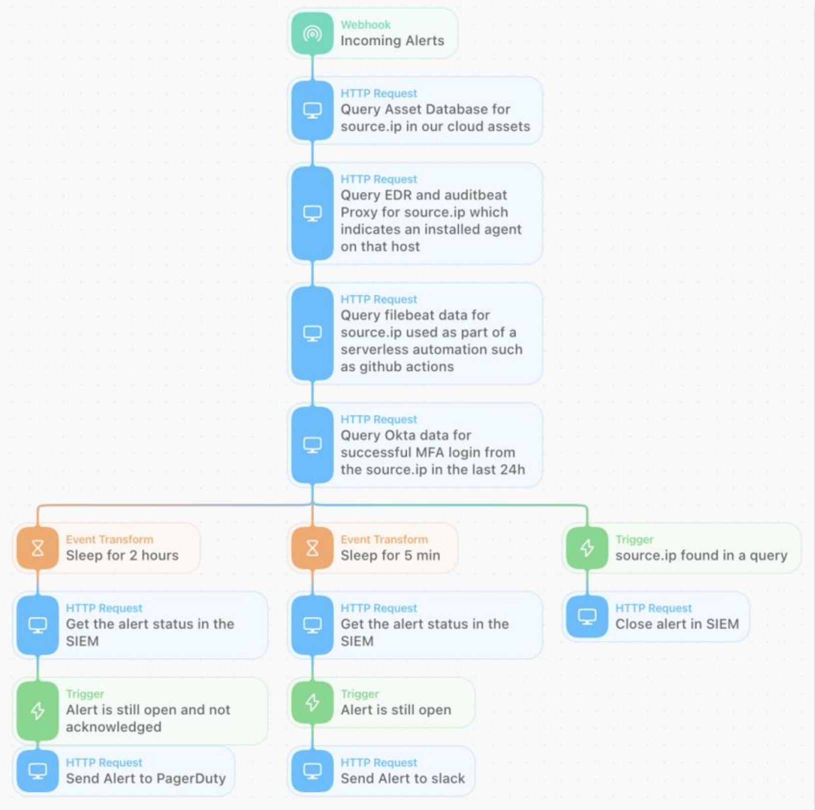

Webhook
Incoming Alerts

HTTP Request
Query Asset Database for source.ip in our cloud assets

HTTP Request
Query EDR and auditbeat Proxy for source.ip which indicates an installed agent on that host

HTTP Request
Query filebeat data for source.ip used as part of a serverless automation such as github actions

HTTP Request
Query Okta data for successful MFA login from the source.ip in the last 24h

Event Transform
Sleep for 2 hours

HTTP Request
Get the alert status in the SIEM

Trigger
Alert is still open and not acknowledged

HTTP Request
Send Alert to PagerDuty

Event Transform
Sleep for 5 min

HTTP Request
Get the alert status in the SIEM

Trigger
Alert is still open

HTTP Request
Send Alert to slack

Trigger
source.ip found in a query

HTTP Request
Close alert in SIEM

elastic

Additional Resources

My Blog posts at Elastic: www.elastic.co/blog/author/aaron-jewitt

Some other blog posts from our team members:

www.elastic.co/blog/how-the-elastic-infosec-team-uses-elastic-security

www.elastic.co/blog/how-to-build-a-cybersecurity-asset-management-solution-on-the-elastic-stack

www.elastic.co/de/blog/elastic-on-elastic-how-infosec-deploys-infrastructure-and-stays-up-to-date-with-eck

www.elastic.co/blog/beginners-guide-to-cloud-security

Questions?

Don't let the sadness
of the swamps get to you.
You have to try. You have to care.