amadeus

SIFT

SMART
INTELLIGENT
FINDING
TRIAGING

Benjamin Hilaire
Jean-Philippe CARLENS

08SEP26

Telecom valley | Animateur Azuréen Numérique

AZUR TECH Winter

2 DÉCEMBRE 2025 | SOPHIA ANTIPOLIS

LES TECHNOLOGIES DU NUMÉRIQUE

4 millions
passengers
boarded per day

100 000 tps
peak inbound
traffic

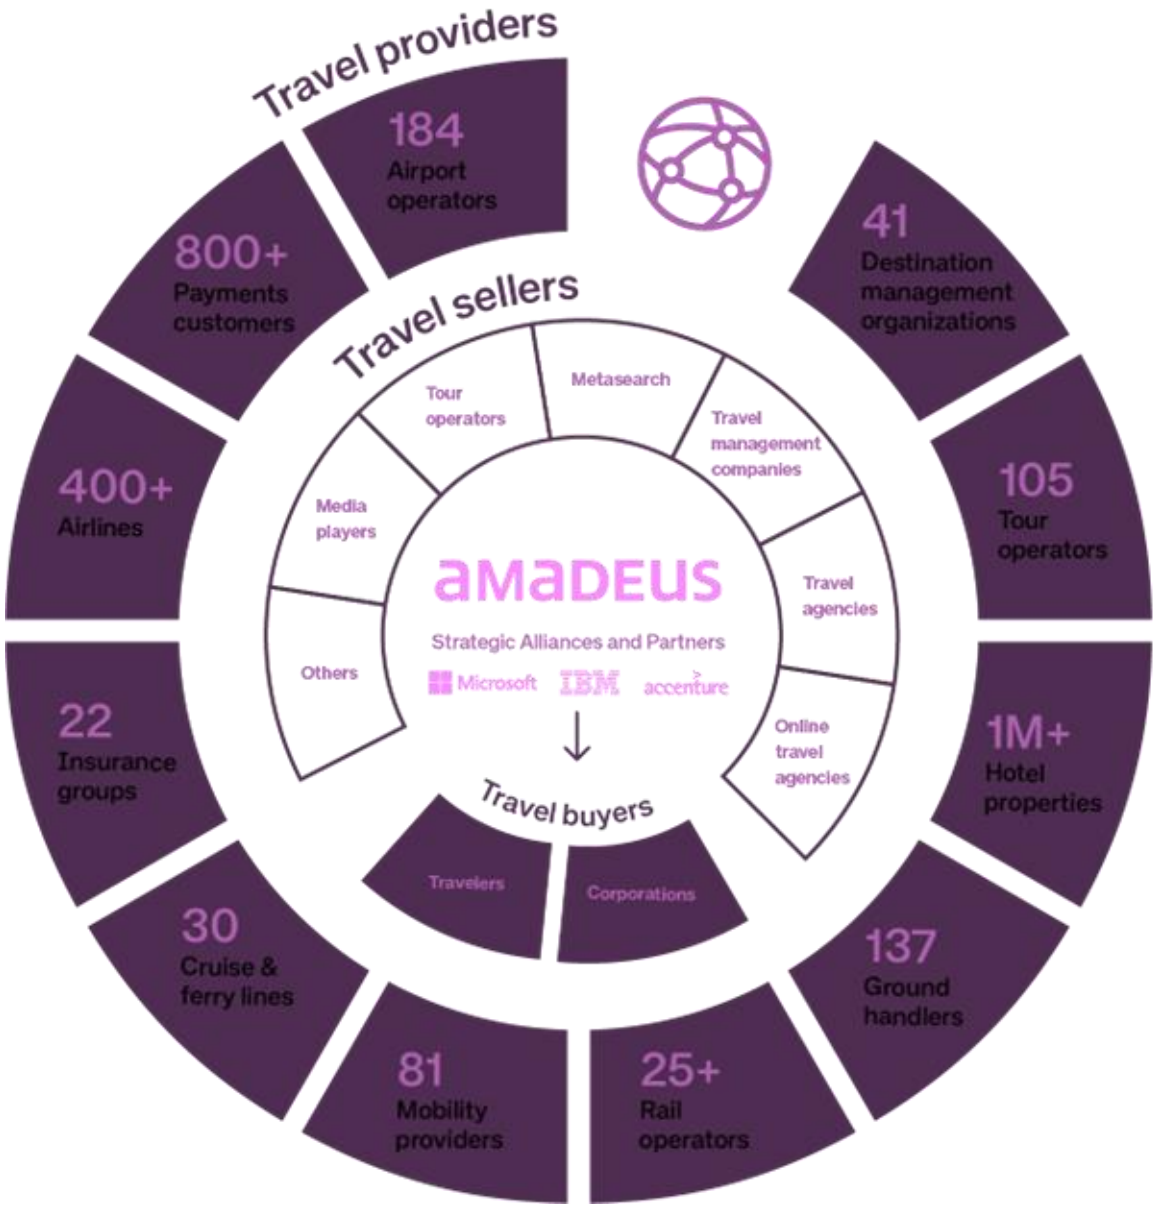

More than
18,000
people worldwide

190+
countries

150+
nationalities

60+
languages

Amadeus. It’s how travel works.

Serving customers in

190+ countries

6th consecutive year included in the Financial Times list of

Diversity leaders

Global team of

20,000+ professionals

One of the largest R&D investors in the software industry in Europe. Gross investment in 2024

€1,365 million

Bookings in 2024

470+ million

Passengers boarded in 2024

2.2+ billion

Payments processed

$120+ billion

Revenue 2024

€6,141.7 million

Source: Global Report 2024

We're a Team!

Dream team from

- Application Security, CISO
- Security scanner team, Engineering Toolchain

Mickaël Bridard

Jean-Philippe Carlens

Nicolas De Toffoli

Benjamin Hilaire

Federico Rizzo

AGENDA

Swimming

AppSec is critical

Biking

Scanning our code for secrets

Running

SIFT for smart triaging of findings

Celebration!

Disclaimer

Original work was done in

July 2025

(last century in AI Era)

Most of the slides are from

December 2026

SWIMMING

APPSEC IS CRITICAL

Oooooops...

CPO MAGAZINE HOME NEWS INSIGHTS RESOURCES

CYBER SECURITY NEWS · 4 MIN READ

US Treasury Breached by Chinese State-Sponsored Hackers Via Stolen API Key

SCOTT IKEDA · JANUARY 6, 2025

New information from US Treasury Department officials indicates that the December breach of its workstations was the work of China's state-sponsored hackers, and that they got in via a stolen API key.

The hackers stole unclassified documents during the raid, but little is known about what they accessed. The attackers appear to have obtained access via a third-party vendor called BeyondTrust, a security and technology provider for Treasury workstations. Both BeyondTrust and the Treasury have begun to work with CISA, the FBI, US intelligence agencies and third-party investigators, but the state-sponsored hackers appear to have had free reign of the system at this point.

Chinese state-sponsored hackers once again

A letter to lawmakers from Treasury officials, part of a [illegible] incident, indicates that the attribution to China [illegible] ters." More details are re[illegible]

Shai-Hulud 2.0 Supply Chain Attack: 25K+ Repos Exposing Secrets

Detect and mitigate malicious npm packages linked to the recent Shai-Hulud-style campaign. Over 25,000 affected repositories across ~350 unique users.

Wiz Customers: Pre-built detection query

Hila Ramati, Merav Bar, Gal Benmocha, Gili Tikochinski
November 24, 2025 · 7 minute read

Table of contents
- Key takeaways
- What is this campaign?
- Scope and Prevalence
- Preliminary Impact
- Payload analysis
- Which actions should security teams take?
- How Wiz Can Help
- Appendix
- References

Key takeaways

- **A new Shai-Hulud–linked npm supply-chain campaign compromised major packages**
  Popular projects from Zapier, ENS Domains, PostHog, and Postman were temporarily trojanized, leading to GitHub repos populated with stolen victim data. Some of these packages are highly prevalent, occurring in roughly 27 % of cloud and code environments scanned by Wiz.

The number of compromised packages on npm is steadily growing, currently at ~700 in total

Wiz Research and other vendors are monitoring newly added versions, but fortunately many packages have already been reclaimed by their owners, including those from Zapier, Posthog and Postman, and the malicious versions have been removed from npm.

- **The blast radius is already massive – 25,000+ malicious repos across ~500 GitHub users**
  Wiz Research has identified widespread automated replication tied to this campaign.
- **The attack is accelerating at ~1,000 new repos every 30 minutes**
  Newly compromised packages continue to surface, many containing files tied directly to this activity.

xAI Dev Leaks API Key for Private SpaceX, Tesla LLMs

May 1, 2025 · 36 Comments

An employee at Elon Musk's artificial intelligence company **xAI** leaked a private key on **GitHub** that for [illegible] have allowed anyone to query private xAI large language models (LLMs) [illegible] custom made for working with internal data from Musk's companies, [illegible] **Twitter/X,** KrebsOnSecurity has learned.

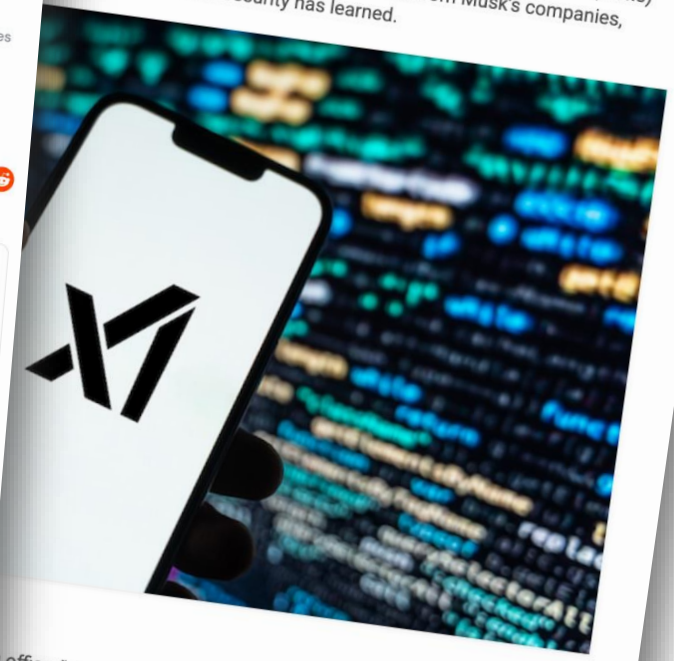

[illegible] g officer" at the security consultancy **Seralys**, was the first to publicize [illegible] application programming interface (API) exposed in the GitHub code [illegible] ember at xAI.

[illegible] ht the attention of researchers at **GitGuardian**, a company that [illegible] diating exposed secrets in public and proprietary environments. [illegible] scan GitHub and other code repositories for exposed API keys, and fire [illegible] sers.

[illegible] ebsOnSecurity the exposed API key had access to several unreleased [illegible] veloped by xAI. In total, GitGuardian found the key had access to at [illegible] Ms.

[illegible] cess the X.ai API with the identity of the user," GitGuardian wrote in an [illegible] AI. "The associated account not only has access to public Grok models [illegible] appears to be unreleased (grok-2.5V), development [illegible] ls (tweet-rejector, grok-spacex-

Sources
https://www.cpomagazine.com/cyber-security/us-treasury-breached-by-chinese-state-sponsored-hackers-via-stolen-api-key/
https://www.wiz.io/blog/shai-hulud-2-0-ongoing-supply-chain-attack
https://krebsonsecurity.com/2025/05/xai-dev-leaks-api-key-for-private-spacex-tesla-llms/

Open door for hackers

- What is a secret?
  - User password
  - API token
  - SSH key

- In 2024, 23M+ new hardcoded secrets added to public GitHub:
  - +25% vs. 2023
  - Source: GitGuardian

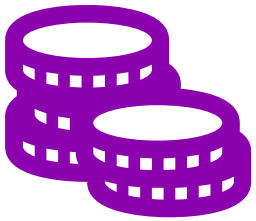

- Cost can be huge:
  - Business
  - Reputation
  - Fines
  - …up to billions of $

Shifting security left

A… s,
we … bout
SECRETS!

Secret scan is part of Amadeus Secure Development Lifecycle

Any company is concerned…

BIKING

SCANNING OUR CODE FOR SECRETS

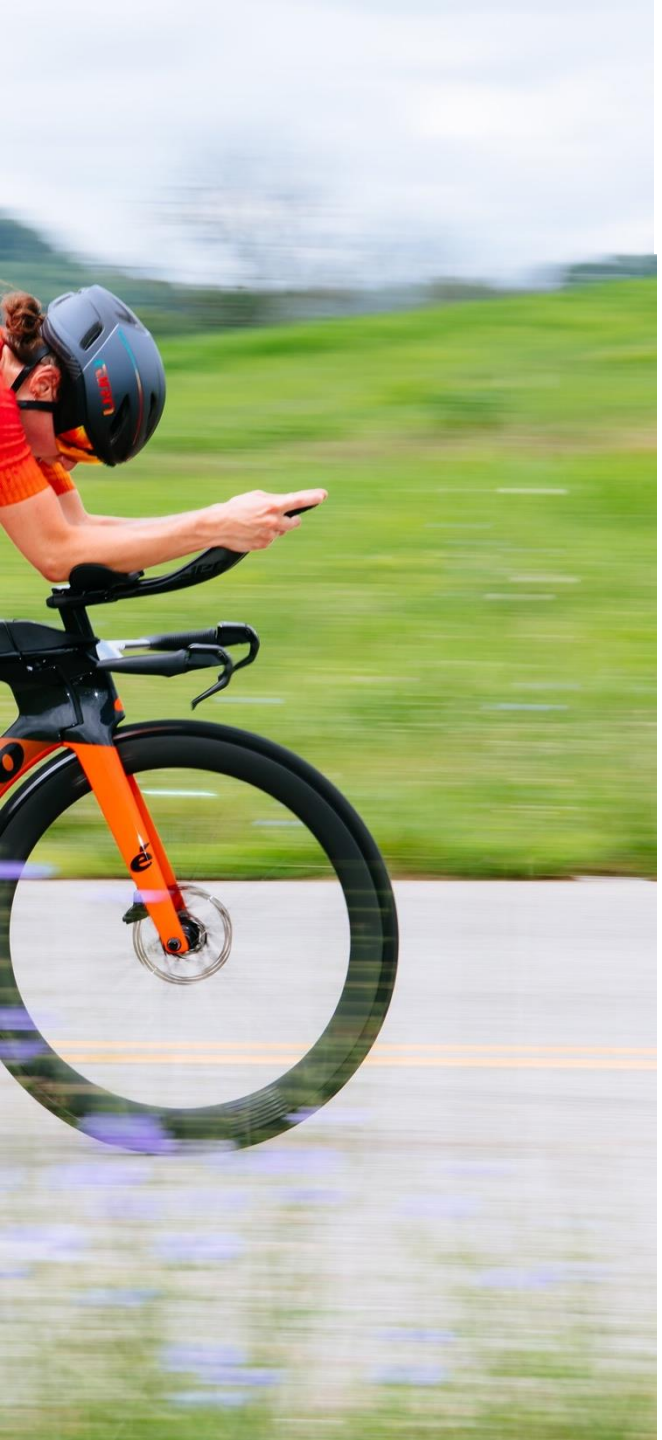

The bike leg

The longest and most strategic part...

Where keeping a steady pace is everything!

Which means scanning everything, regularly, everywhere!

We need a ~~bike~~ scanner!

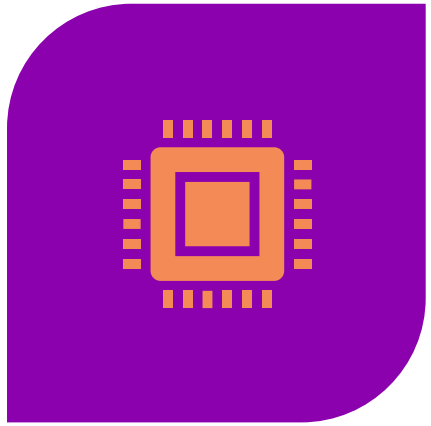

STATIC
APPLICATION SECURITY TESTING

DYNAMIC
APPLICATION SECURITY TESTING

SOFTWARE COMPOSITION
ANALYSIS

Multiple categories... like bikes

Software Composition Analysis

Detect third-party dependencies

Report their known vulnerabilities

Maybe the MTB bike?

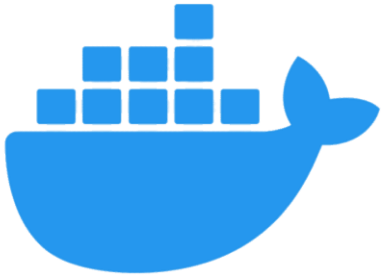

docker

Maven

pypi

nuget

Dynamic Application Security Testing

Tests while it's running

Runtime and misconfigurations issues

No source-code access

Static Application Security Testing

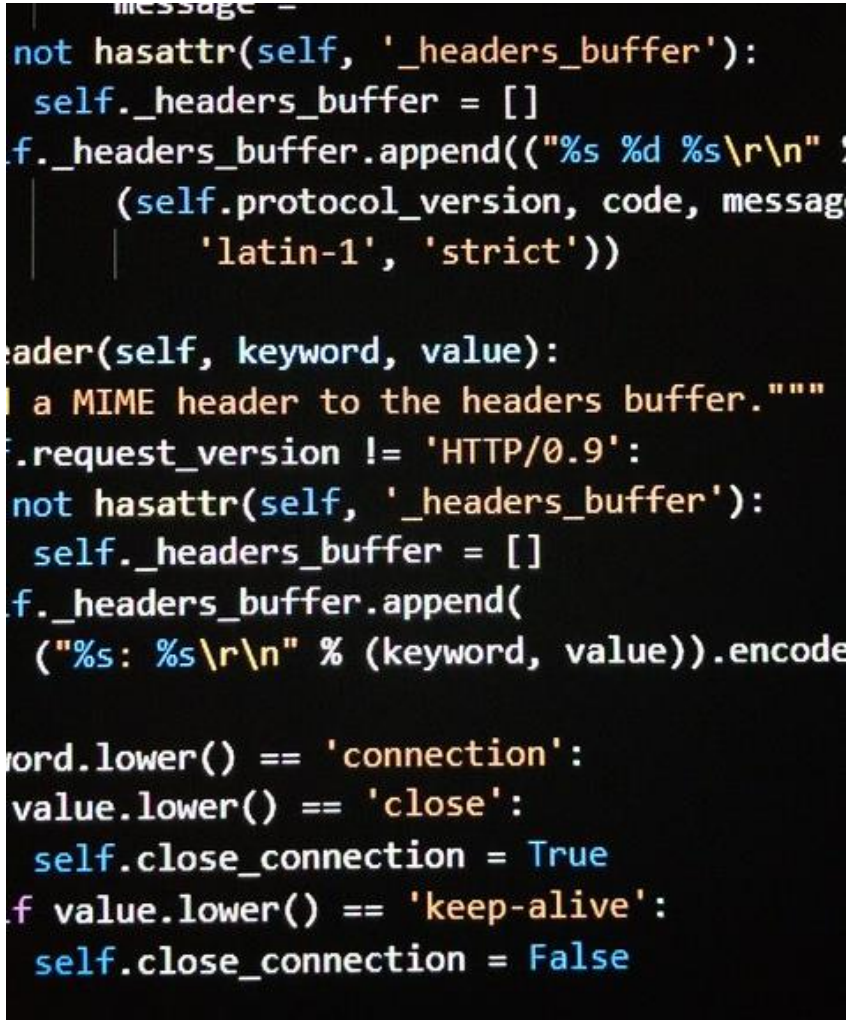

- Analysis of source-code
- Detect insecure code
- Shift-left practice

String values

main.go 2, M

main.go > ...

```go
package main

import "fmt"

var apikey string = "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"

func main() {
    fmt.Println("Hell
}

```

```
ggshield: apikey

Secret detected: GitHub Access Token
Validity: Invalid
Known by GitGuardian dashboard: NO
Total occurrences: 1
```

Configuration files

```
1     ####      Environment Settings      ####
2
3  masterkeyring:
4   driver : keyring
5   service: system
6
7  key_pass: zaCELgL0imfnc8mVLWwsAawjYr4Rx-Af50DDqtlx
8
9   master_sign_pubkey: True
10  signing_key_pass:
```

And more!

Cryptographic key files

Log files

Documentations

15th September 2025

23rd November 2025

The LiteLLM Supply-Chain Attack — TeamPCP "SANDCLOCK" CI/CD Credential-Harvesting Campaign Via A Backdoored Trivy GitHub Action

CYBER THREAT INTELLIGENCE

Mini Shai-Hulud Targets SAP npm Packages With a Bun-Based Secret Stealer

New Actors Deploy Shai-Hulud Clones: TeamPCP Copycats Are Here

Moshe Siman Tov Bustan

May 17, 2026 4 min read

BLOG

Shai-Hulud "Hades" Wave Hits Six PyPI Bioinformatics Packages via Stolen Tokens

Which model for our bike?

Directly a premium one?

Proprietary solutions:

- Managed service maintained by an external company
- Integrated
- Advanced detection models (built-in ML/AI, vendor-maintained rules)
- Faster onboarding at large scale

Which model for our bike?

Free, open-source solutions:

- Flexible, can run everywhere: laptop, CI/CD, automations, Gitlab, GitHub...
- Cost-effective: no license, simple to scale org-wide
- Vendor-neutral: no lock-in, no external data exfiltration concerns

We can win the race with a cheap one?

Our choice: gitleaks/**gitleaks**

Find secrets with Gitleaks 🔑

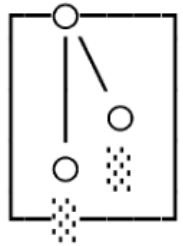

- Widely used in the industry
- Portable
- Configurable
- Simple

How does it work?

A set of named regular expressions, rules, describing how a secret looks like

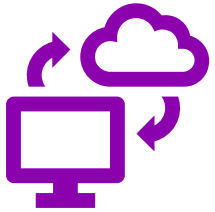

AWS access key

AKIA[0-9A-Z]{16}

GitHub token

ghp_[0-9A-Za-z]{36}

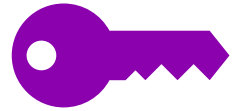

Private key

-----BEGIN (RSA|EC) PRIVATE KEY-----

Two types of wheels for the bike

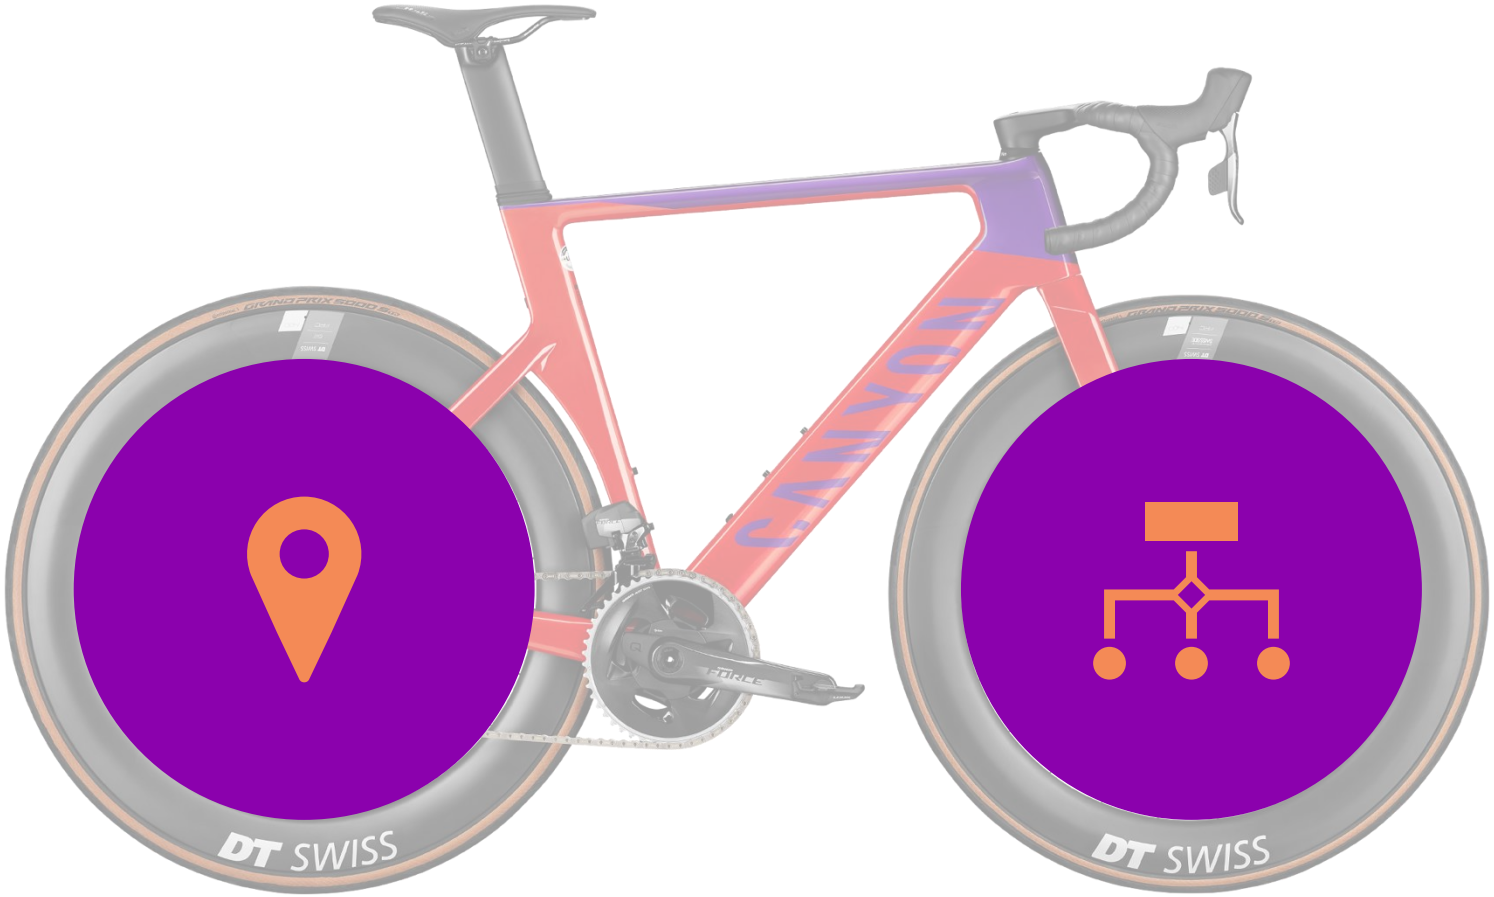

ENTIRE GIT HISTORY

PULL-REQUESTS

Report findings with context

← SQL connection string

Close alert

#123 Open in f457c475 • detected 15m ago

Locations

/src/app/connector.cs @ f457c754

```
38        connection.ConnectionString);
39      }
40    {
41      return "Data Source=np:$(SQLTarget);Initial Catalog=FABRIKAM;UID=sa;PWD=password123;Encrypt=$(Encrypt);Trust
42    }
43 }
44
```

Recommendation

Review evidence of possible plaintext (or base 64-encoded plaintext) secrets in versioned engineering content.

Follow the steps below before you close this alert:

1. Rotate the secret and store securely if it's in use to prevent breaking workflows.
2. Revoke this SQL connection string to prevent unauthorized access.
3. Check security logs for potential breaches.
4. Close the alert as revoked.

Severity

Critical

Confidence: other

Introduced

Jul 9, 2021

Finding details

Type
SQL connection string

ID
SEC101/200

We deployed at large scale

Prevent new leaks

- Scan triggered on pull-requests
- Changes only

Remediate existing leaks

- Recurrent scan of all repositories
- Full history

Great news!

Many findings

. . . including a lot of false positives

Regular expressions easily generate false positives

Generic rules are noisy: they match too broadly

Not generic: `ghp_abc123def456`

Generic: `String test_key = "value"`

Generic regular expression = false positive explosion

Annoying...

☐	curl-credentials	curl -X POST -u "KEY:SECRET"	SECRET	High

</> SECRET | COMMENTS & HISTORY | ATTACHMENTS

Details

Secret	curl -X POST -u "KEY:SECRET"
Entropy	3.9946804
Repository	ad - gitops-test-automation

Lines

```
curl -X POST -u "KEY:SECRET"
```

- Source : README.md *(line 42, commit 6887245b2aa0fa05331c51823775533052087347, at 2024-10-22T12 14 18+00 00)*

RUNNING

SIFT FOR SMART TRIAGING OF FINDINGS

Triaging findings is time consuming

Can you tell which of these are real secrets?

Filename: conf/system/prod/server.conf.old

Line 5: `Password="v9$Tq7#pLm2@Xc8!"`

Rule: generic-key

✓

Filename: etc/apps/graphql/schema/plugins.py

Line 12: `key="cvssV3TemporalVector"`

Rule: generic-key

✗

Filename: api/.npmrc

Line 2: `auth = T2hObzpUaGlzSXNBRmFrZQ==`

Rule: generic-api-key

✓

How did we reason as a human?

EXTRACT

```
Filename: api/.npmrc
Line 2:   auth = T2hObzpUaGlzSXNBRmFrZQ==
Rule: generic-api-key
```

REASON

- api/.npmrc → a valid NPM configuration file
- The key name is auth → self-explanatory ☺
- Base64 encoded string → looks like a valid credential
- ⚠ But once decoded: OhNo:ThisIsAFake → obviously not used in Amadeus

CONCLUDE

False Positive

Let's ask a LLM to perform the analysis!


```
Filename: api/.npmrc
Line 2:   auth = T2hObzpUaGlzSXNBRmFrZQ==
Rule: generic-api-key
```

>>> Tell me if it is a false positive or a potential secret and why

Model 1

The presence of an auth directive inside an .npmrc file strongly indicates a credential configured for npm registry access. The value looks like an encoded token, which matches typical secret patterns. Given the file path and the semantic meaning of auth in .npmrc, I would classify this as a potential secret.

Model 2

An .npmrc file carrying an auth = <value> line is typically used to store registry access tokens. The final decision should verify the actual auth value (e.g., by decoding or validating whether it is a placeholder).
I would classify this as potential secret.

Model 3

While .npmrc files often contain registry authentication, this specific entry looks like a non-sensitive placeholder. The value is Base64-encoded and decodes to OhNo:ThisIsAFake, which is clearly not a production credential.
I would classify this as a false positive.

How would different LLMs perform on a real project?

Human triage

- 1 project
- ✓ 10 true positives
- ✕ 44 false positives
- ⏳ ~45 min triage

A FEW HOURS LATER...

Results of an early assessment with several LLMs

Human triage

- 1 project
- ✓ 10 true positives
- ✕ 44 false positives
- ⌛ ~45 min triage

Models	True Positives Identified	False Positives Identified
codellama:13b	60%	20%
mistral-nemo	100%	27%
phi4	100%	27%
qwen2.5-coder:14b	70%	75%
mistral-small	100%	88%

Analysis performed in March 2025

Promising results… in a few minutes!

Turning a generic LLM into a finding triage assistant

ROLE

```
You are an expert in software and infrastructure security
You like simplicity and pragmatism
```

OBJECTIVE

```
You analyze security alerts and decide if you consider them as
potential secrets or false positive
You rely only on the alert inside the <SecurityScannerData> tags
Your answer must be binary: TRUE POSITIVE or FALSE POSITIVE
```

FORMAT

```
Your answer must be a machine-readable JSON object containing:
-  "result" : true or false
-  "reasons": a short explanation supporting your conclusion
```

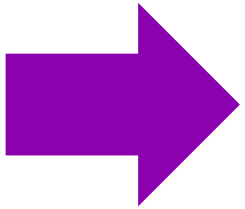

LLM's brain configuration

Turning our reasoning into automation: SIFT!

gitLeaks **EXTRACT**

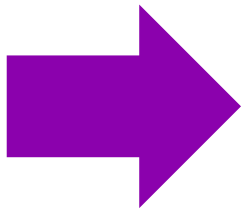

Smart Intelligent Finding Triaging

Prepare prompt

Call LLM

Produce result

```
"locations": [
 {
   "physicalLocation": {
     "artifactLocation": {
       "uri": “api/.npmrc"
      },
      "region": {
        "startLine": 2,
        "snippet": {
          "text": “auth = T2hObzpUaGlzSXNBRmFrZQ=="
        }
       }
     }
 }
],
```

GitLeaks result (raw finding)

```
"analysis": {
    "result": "false",
    "reasons": "The detected string is encoded in base64. The string
decodes to OhNo:ThisIsAFake which suggests it might be a placeholder
or example text. Despite it is present in an api/.npmrc file, this
is likely a false positive."
  }
```

SIFT result (human-like reasoning, at scale)

What's behind the scene "Call LLM"?

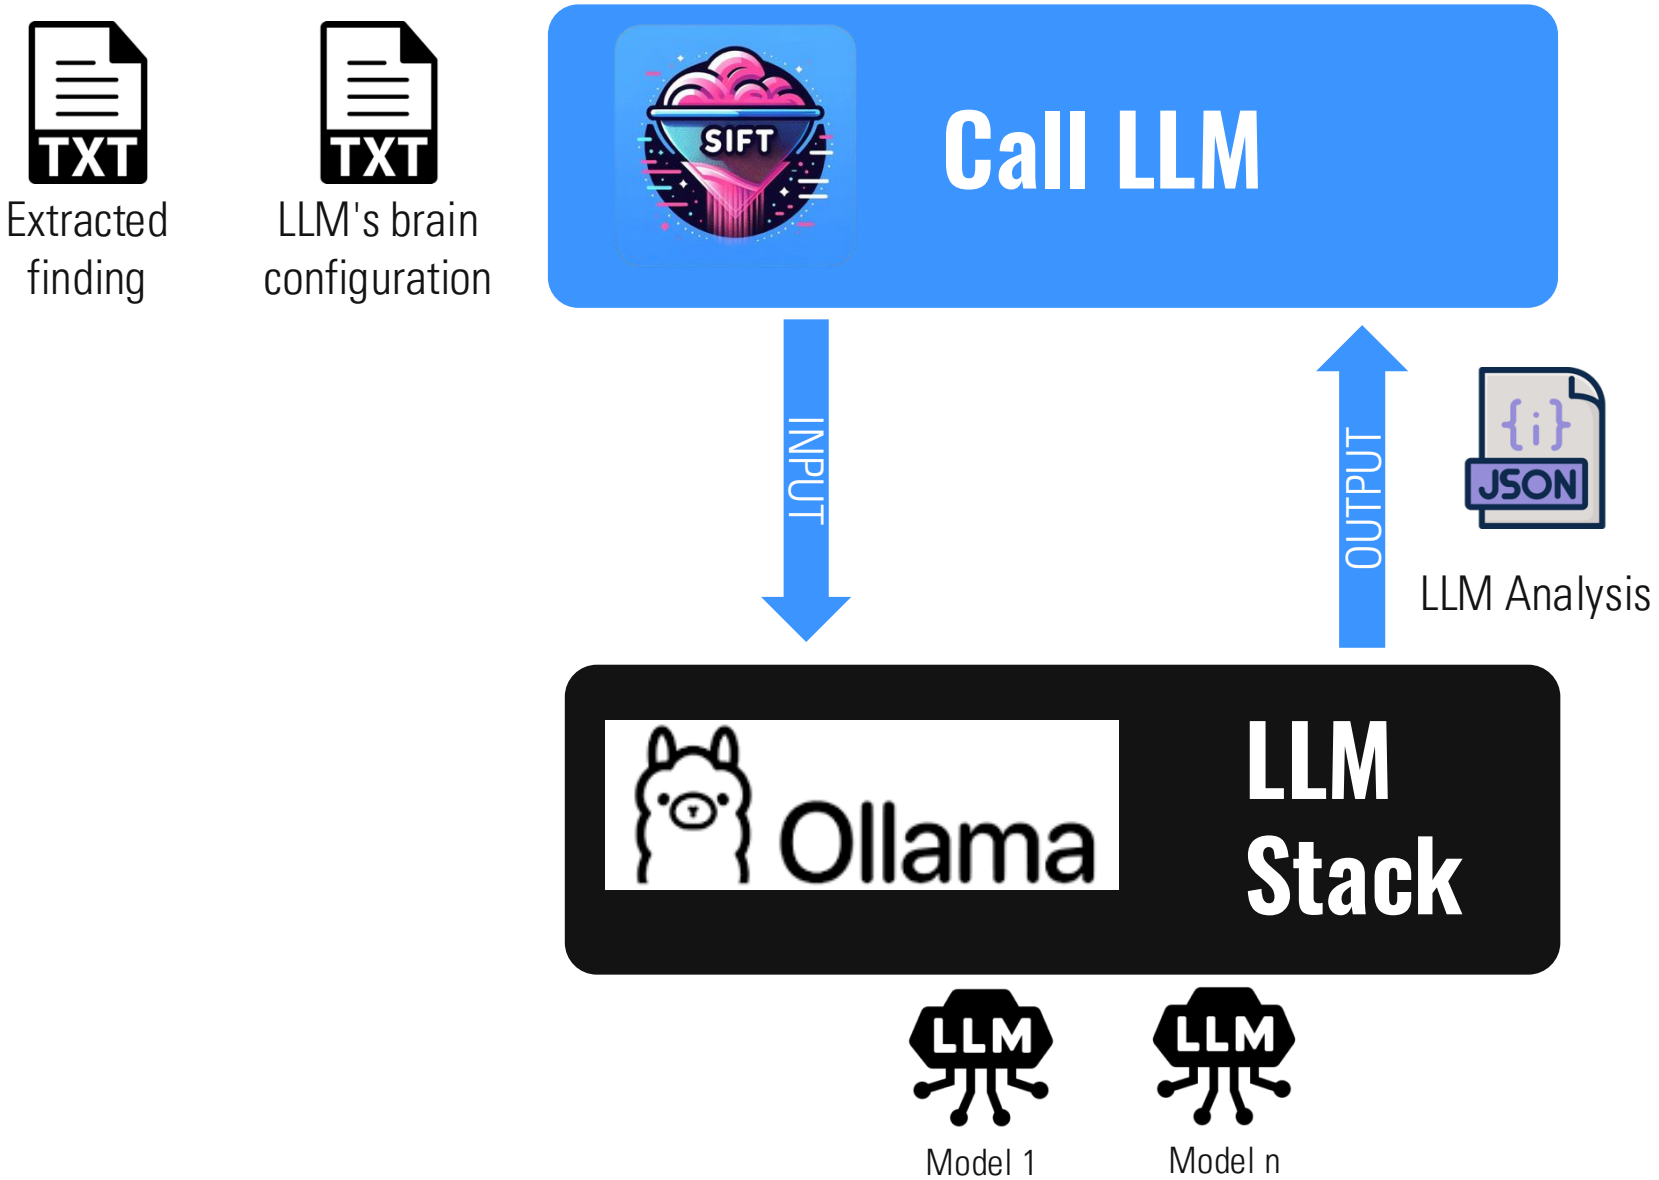

Let’s run SIFT on all already triaged findings!

Human triage

- 15,000 findings
- ✓ 1,500 true positives
- ✕ 13,500 false positives
- ⏳ ~59 estimated days

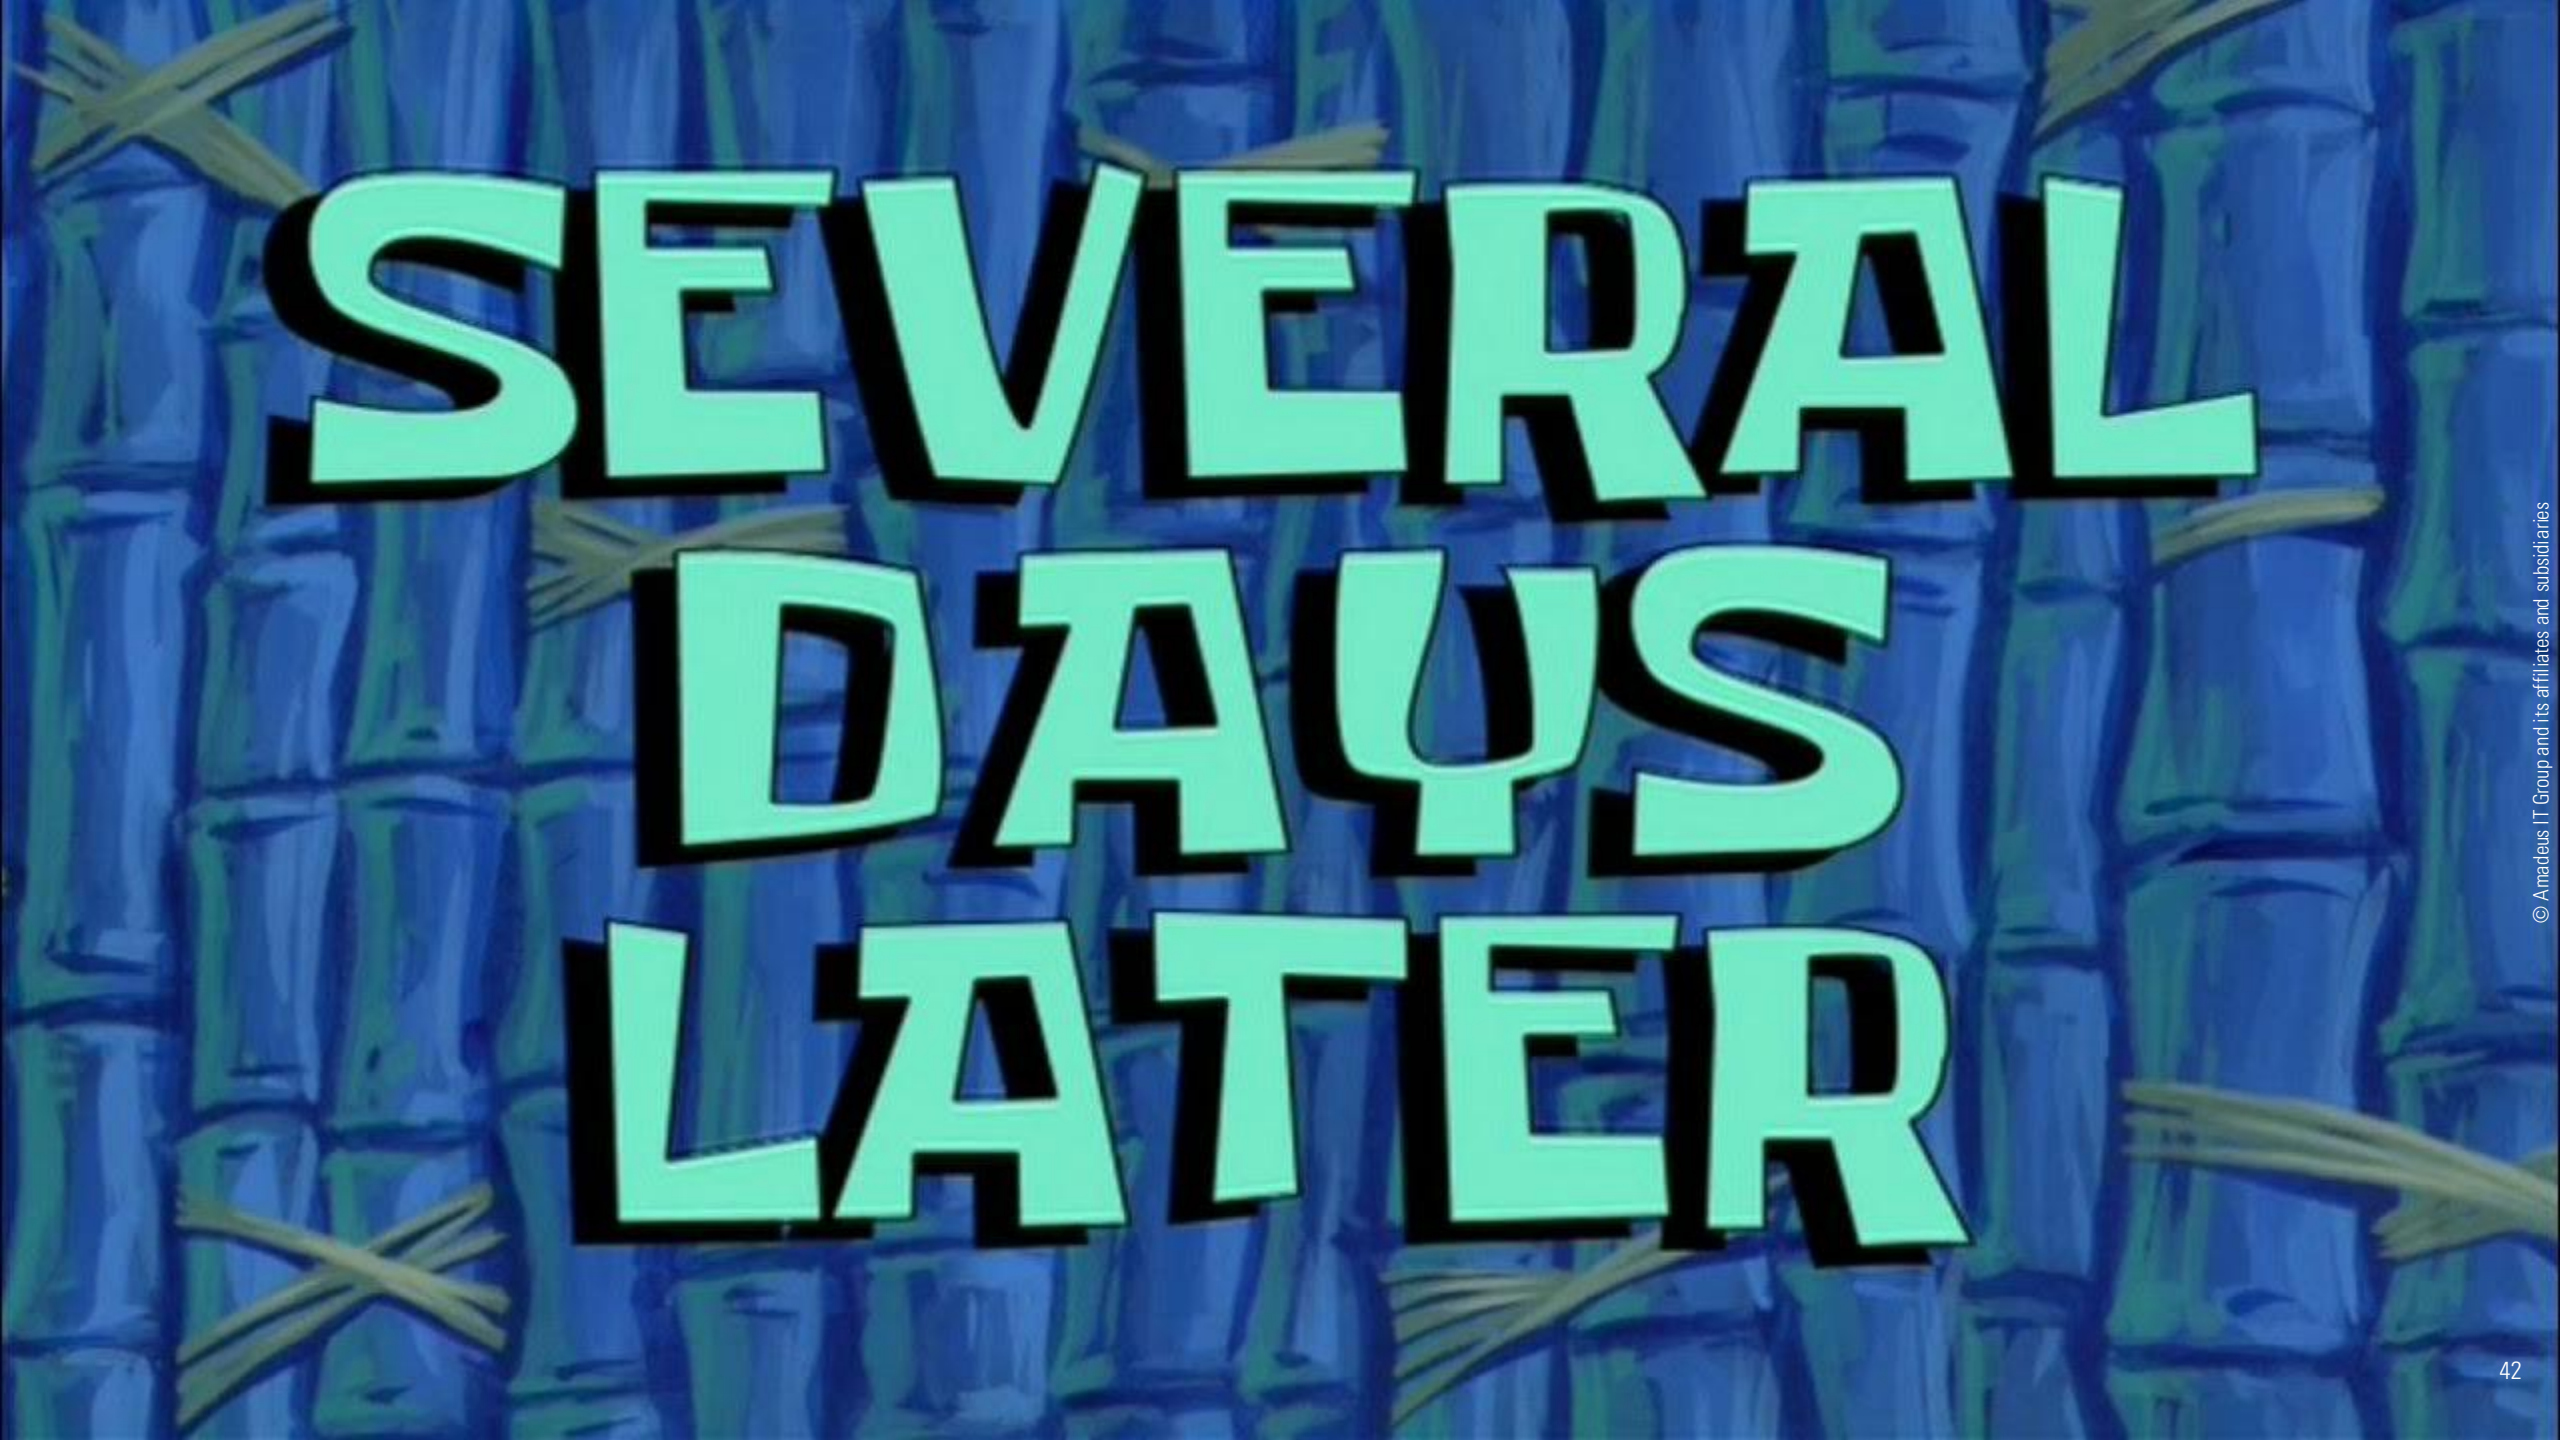

How did SIFT compete with humans on 15.000 findings?

1.500
10% of True Positives

13.500
90% of False Positives

475
Hours of triage
(~59 working days)

Figures shown are representative and do not reflect actual company

SIFT

97%
True Positive Agreement

82%
False Positive Agreement

83
Hours of triage
(~3 days - not sleeping)

We have a new product that will reduce needed effort to win the race against secrets findings.

SIFT: *Smart Intelligent Finding Triaging*

And it is 100% legal, secured and open-sourced!

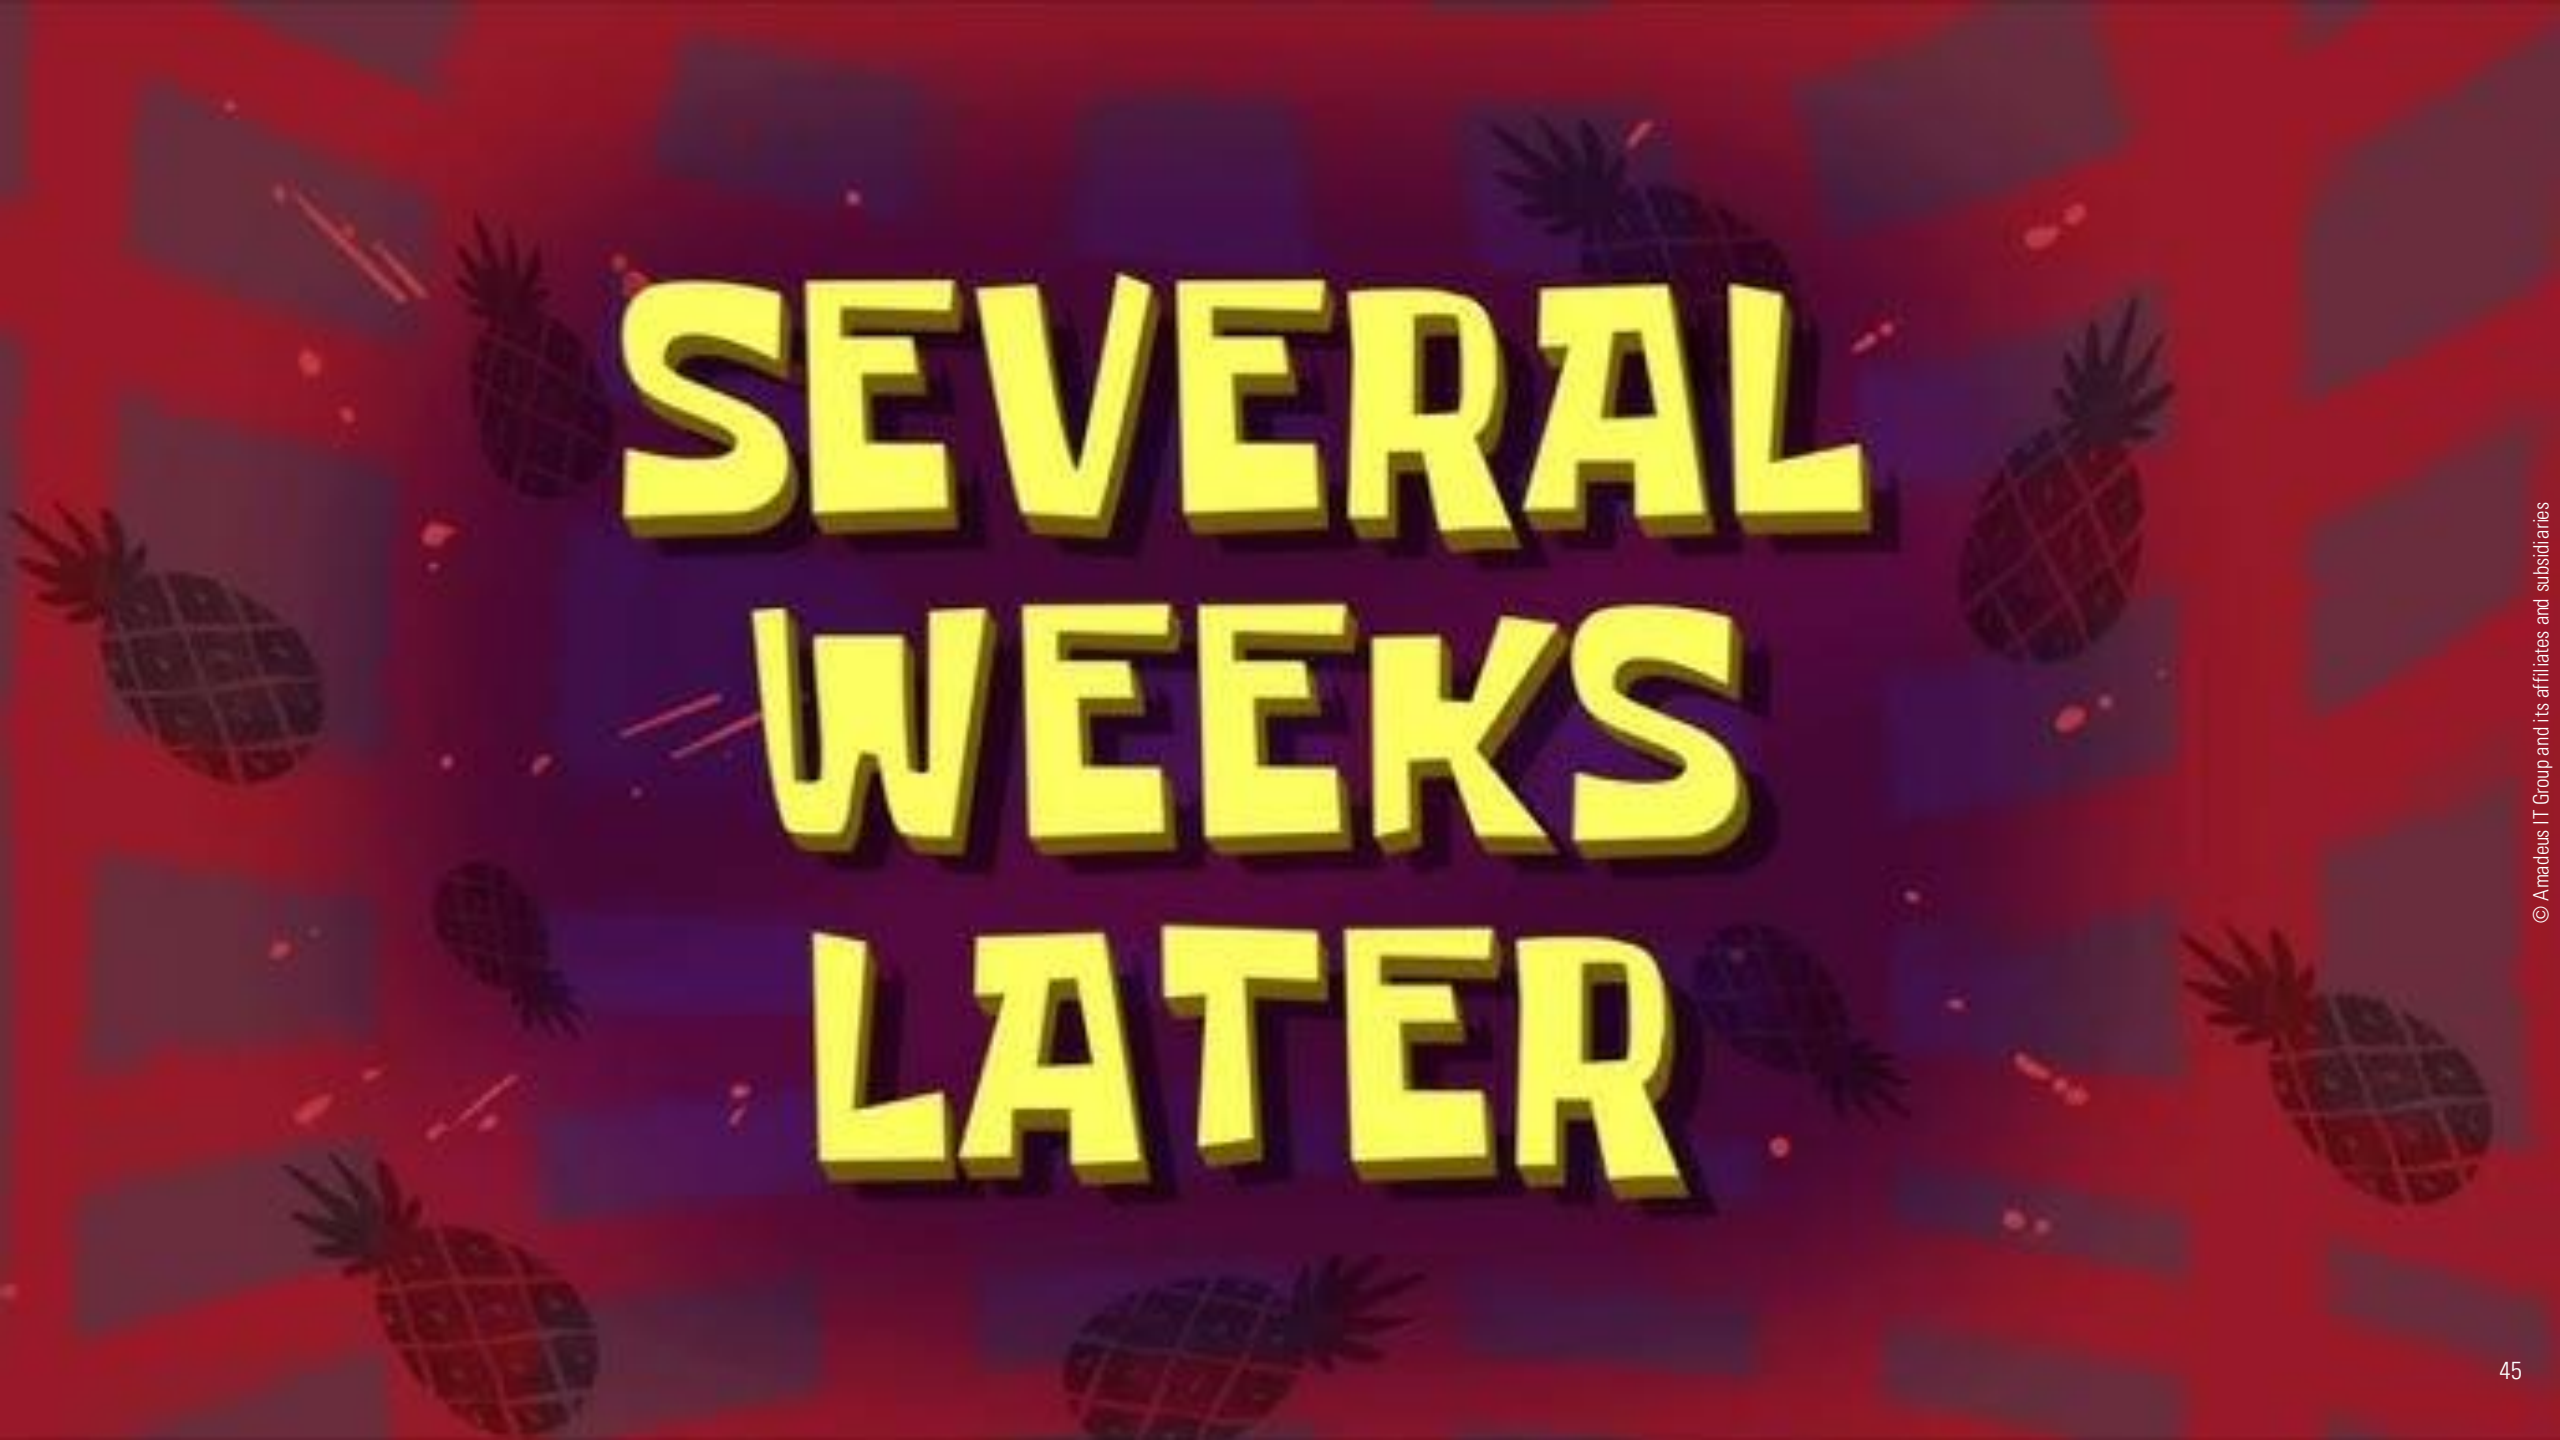

Bye false positives!

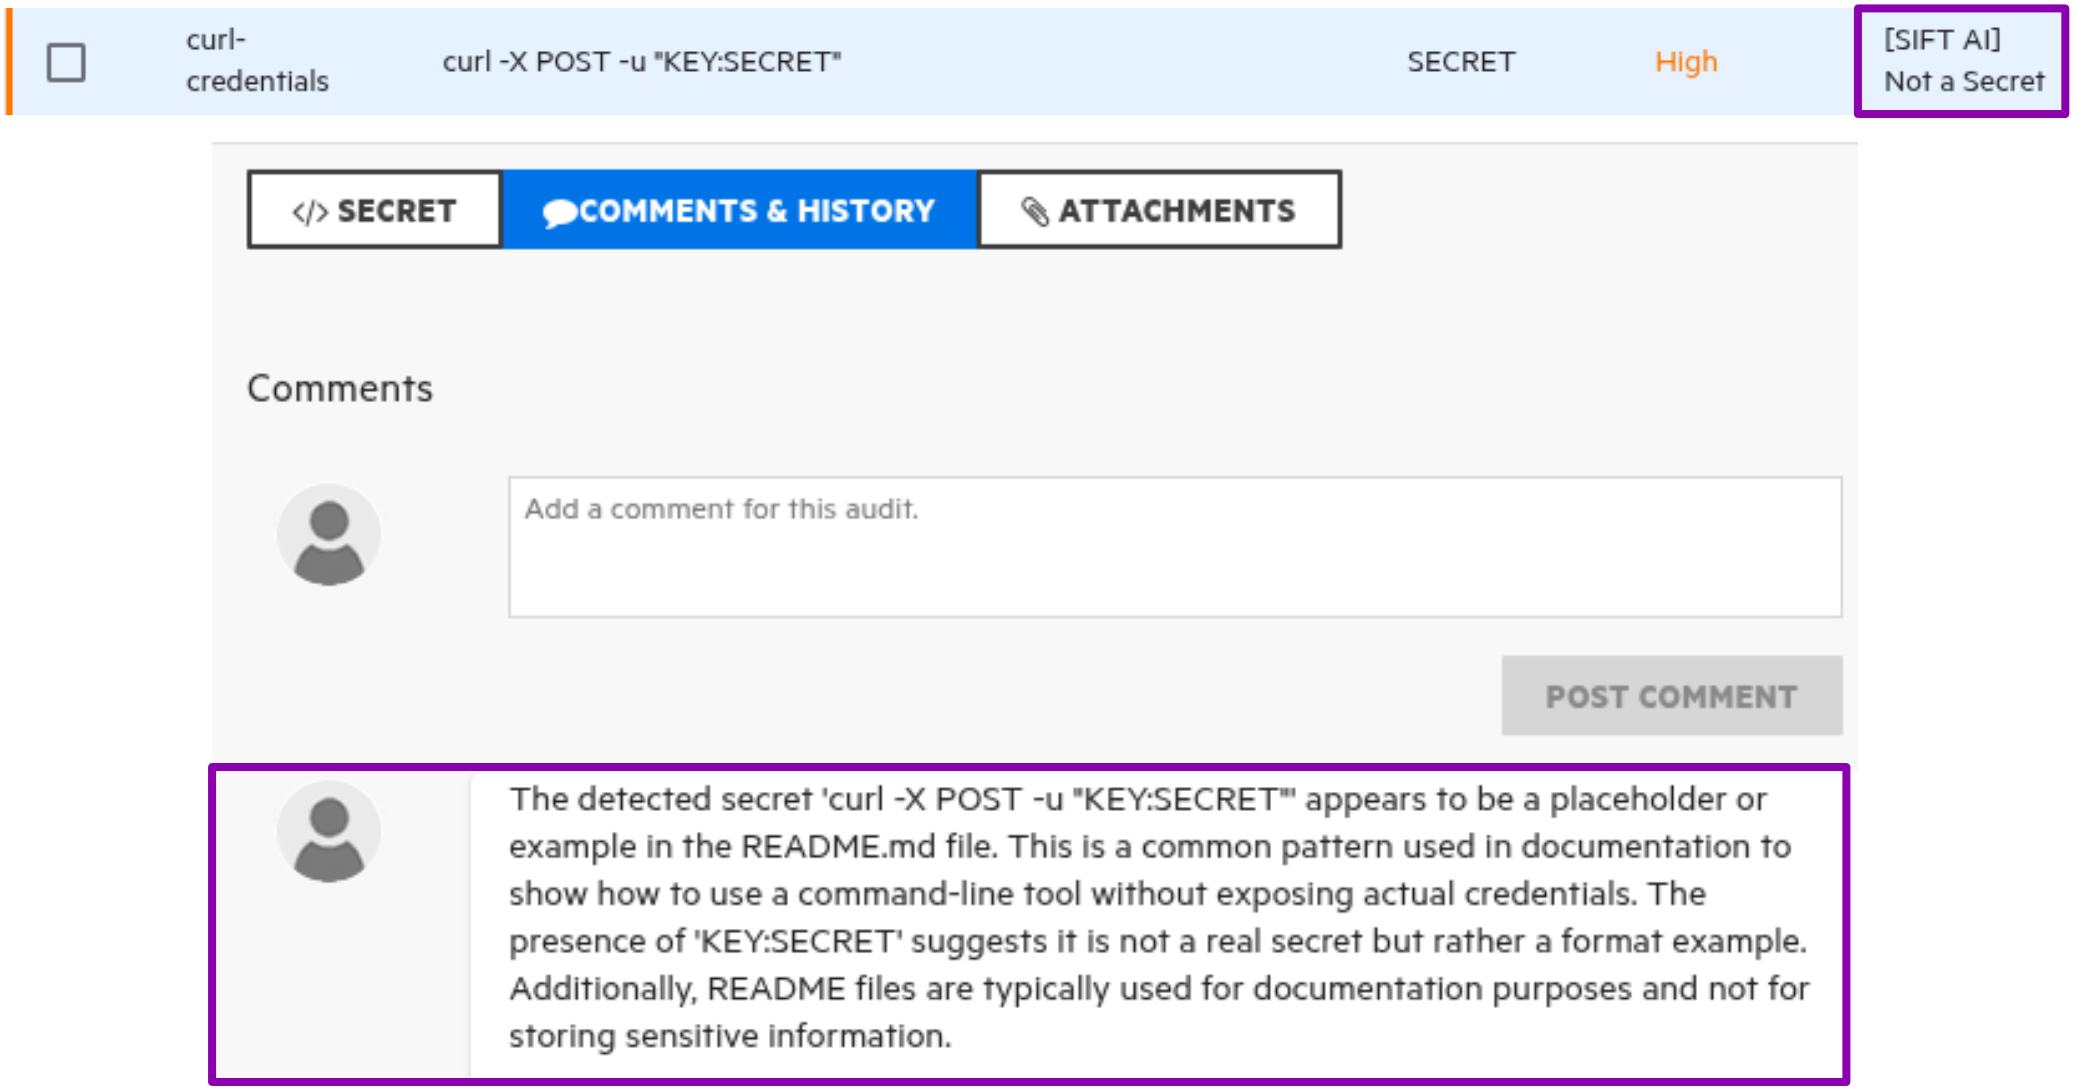

CELEBRATION!

Wrap-up

- At Amadeus, we do have some ideas…
  and now we have SIFT! 😄

- Significant improvements for:
  - Triage accuracy
  - User eXperience

- Open Source initiative: https://github.com/AmadeusITGroup/sift

September 2026 Update

What happened since ?


SEVERAL
MONTHS
LATER

Can we improve the secret scanner ?

Using Agentic Scanner

01

Plan

02

Scan

03

Validate and deduplicate

Identify non-trivial secret (based on code, parameter position, parameter name,…)

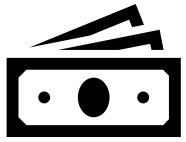

At a COST (requires ”medium” level model)

Can we extend SIFT to all scanner's cases ?

Yes…but

01

It worked with a **skill**

02

It’s much better with a **harness** (programmatic calling agentic)

03

It has bias, assumption and must be **reviewed by human**

Agentic flow

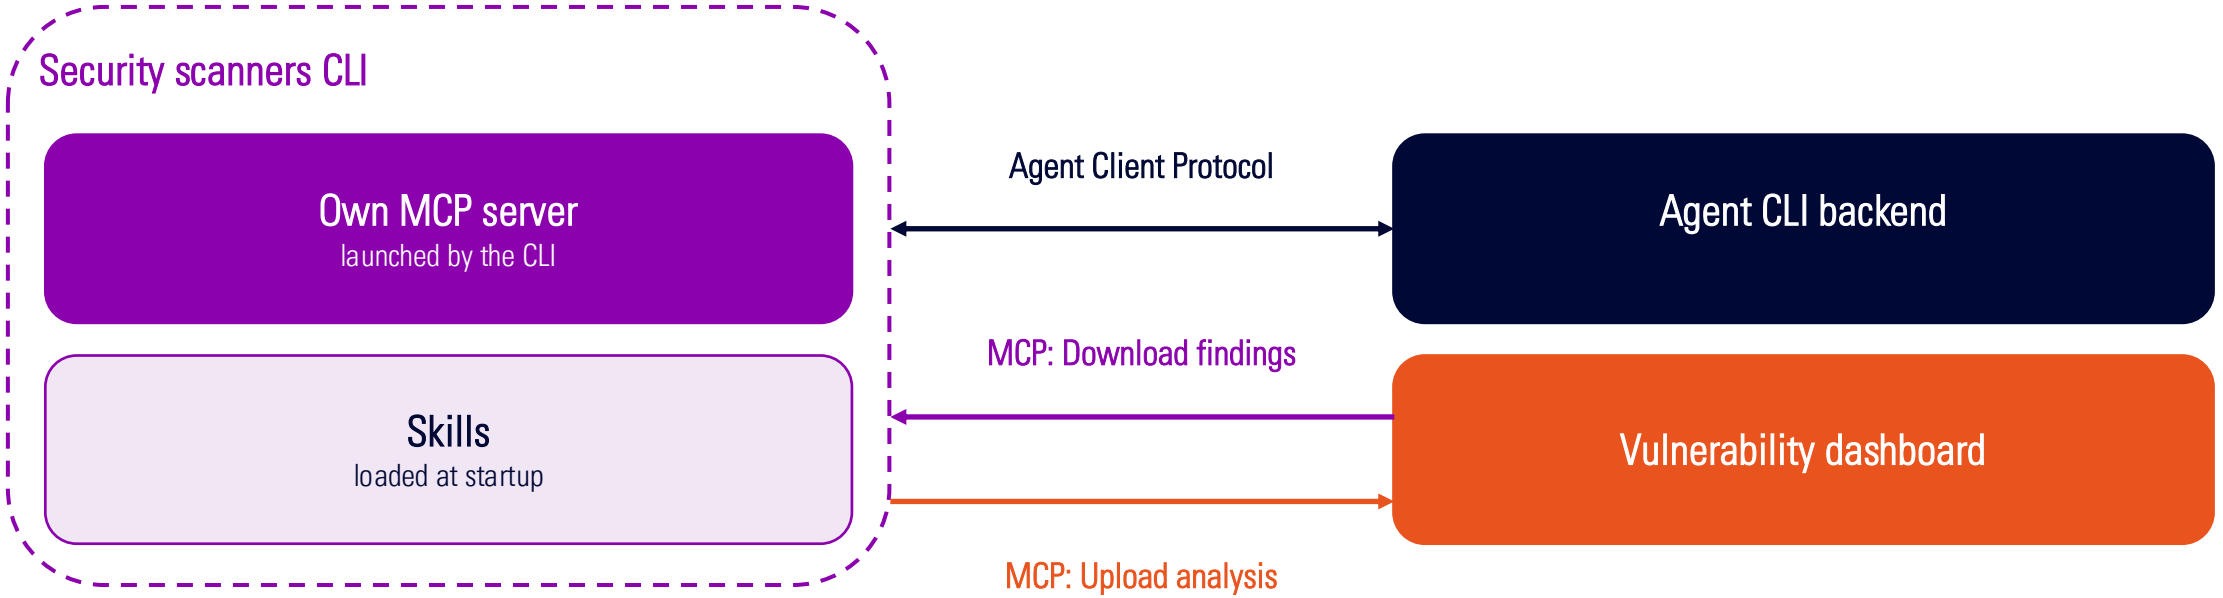

No central subscription, no shared secret

Runs on each developer's own coding agent

Time for questions